

THE JOURNAL OF PHILOSOPHY

FOUNDED BY FREDERICK J. E. WOODBRIDGE AND WENDELL T. BUSH

Purpose: To publish philosophical articles of current interest and encourage the interchange of ideas, especially the exploration of the borderline between philosophy and other disciplines.

Editors: Bernard Berofsky, Akeel Bilgrami, John Collins, Arthur C. Danto, Kent Greenawalt, Patricia Kitcher, Philip Kitcher, Isaac Levi, Wolfgang Mann, Mary Mothersill, Christopher Peacocke, Carol Rovane, Achille C. Varzi. *Consulting Editors:* David Albert, James T. Higinbotham, Charles D. Parsons, Wilfried Sieg. *Managing Editor:* John Smylie.

THE JOURNAL OF PHILOSOPHY is owned and published by the Journal of Philosophy, Inc. *President,* Arthur C. Danto; *Vice President,* Akeel Bilgrami; *Secretary,* Daniel Shapiro; *Treasurer,* Barbara Gimbel; *Other Trustees:* Lee Bollinger, Leigh S. Cauman, Kent Greenawalt, Michael J. Mooney, Lynn Nesbit.

All communications to the Editors and Trustees and all manuscripts may be sent to John Smylie, Managing Editor, Mail Code 4972, 1150 Amsterdam Avenue, Columbia University, New York, New York 10027. FAX: (212) 932-3721.

You may also visit our website at: www.journalofphilosophy.org

THE JOURNAL OF PHILOSOPHY

2006

SUBSCRIPTIONS (12 issues)

| | |
|---|----------|
| Individuals | \$45.00 |
| Libraries and Institutions | \$100.00 |
| Students, retired/unemployed philosophers | \$20.00 |
| Postage outside the U.S. | \$15.00 |

Payments only in U.S. currency on a U.S. bank. All back volumes and separate issues available back to 1904. Please inquire for price lists, shipping charges, and discounts on back orders. Please inquire for advertising rates; ad space is limited, so ad reservations are required.

Published monthly as of January 1977; typeset by Dartmouth Journal Services, Waterbury, VT, and printed by The Sheridan Press, Hanover, PA.

All communication about subscriptions and advertisements may be sent to Pamela Ward, Business Manager, Mail Code 4972, 1150 Amsterdam Avenue, Columbia University, New York, NY 10027. (212) 866-1742

All materials published in *The Journal of Philosophy* are copyrighted by the *Journal* and are protected by U.S. copyright law. All rights reserved. For more information about our copyright policies, please consult our January issue or contact us.

The Journal of Philosophy (ISSN 0022-362X) is published monthly by The Journal of Philosophy, Inc., MC4972, Columbia University, 1150 Amsterdam Avenue, New York, NY 10027. Periodicals postage paid at New York, NY, and other mailing offices.

POSTMASTER: Please send address changes to *The Journal of Philosophy* at MC4972, Columbia University, 1150 Amsterdam Avenue, New York, NY 10027.

© Copyright 2006 by the Journal of Philosophy, Inc.

ISSN 0022-362X

THE JOURNAL OF PHILOSOPHY

VOLUME CIII, NO. 6, JUNE 2006

ABSOLUTIST MORAL THEORIES AND UNCERTAINTY*

You have promised to vote for Mary in the election for party president. You then learn that Fred's policies are much better. What to do? Most of us grant that if Fred's policies are a great deal better, you should break your promise. Absolutists about promise keeping demur. Bad consequences from keeping a promise never justify breaking it no matter how bad they may be. That is a very familiar issue. Our concern in this paper is with an issue that seems to have slipped under the radar. Suppose that you are pretty sure but far from certain that you promised to vote for Mary. It was some time ago; it was quite a party; you know that you indicated that you would vote for her, but did you actually *promise* to do so? What should absolutists about promise keeping say about how the element of doubt affects the question of what you ought to do in such a case?

We will be concerned with this question as it arises for absolutist theories in general. An absolutist theory absolutely prohibits actions of kind *K*, where *K* may be promise breaking, lying (as in Kant¹), intentionally killing the innocent (as in some versions of traditional Catholic ethical theory), or else see the list of "things forbidden whatever consequences threaten" in the Hebrew-Christian ethic according to Elizabeth Anscombe.² Here we need to understand *K* as a property of an action as opposed to a relation between an action and available alternatives to that action. Classical utilitarianism absolutely prohibits doing actions that fail to maximize utility. But an action's failing to

*We are indebted to many audiences and the paper is substantially changed from the various presented versions. Special thanks to Andy Egan, Laura Schroeter, Karen Jones, Jim Griffin, and the editors of this JOURNAL.

¹Kant, "On the Supposed Right to Lie from Philanthropic Motives," in *The Metaphysics of Morals* (*Die Metaphysik der Sitten*, 1797).

²Anscombe, "Modern Moral Philosophy," in *Ethics, Religion, and Politics: The Collected Papers of G.E.M. Anscombe*, Volume III (Minneapolis: Minnesota UP, 1981), pp. 26-42.

maximize utility is a relation the action has to available alternatives. The distinctive feature of the kind of absolutism that we find, for example, in Kant and the Catholic tradition is that the absolutely prohibited kind is independent of the nature of any available alternatives. Of course, an important question for absolutisms of this kind is what to say about cases where *all* the available actions are of the prohibited kind. There are, say, only two things you can do and both involve breaking a promise (a different promise in each case). Opinions differ on how much of a problem this is for absolutism and our discussion will steer clear of the question except for a short mention later in the context of the bearing of strong dilemma views on our argument. Our examples will all be cases where at most one option is of the prohibited kind, *K*.³

More particularly, our concern will be with the following question: agent *S* has to choose between doing *A*, where *A* has *p* chance of being *K* and $1-p$ chance of being not-*K*, and doing *B*, where *B* is certain to have such-and-such consequences. What should a theory that absolutely prohibits doing *K* say concerning what *S* ought to do, as *p* rises and falls? We know what an absolutist theory says when $p = 1$. It says that no matter what the such-and-such consequences are, the action is forbidden. What about values of *p* other than 1?⁴

Mostly, we will frame the argument to come around theories that absolutely prohibit intentionally killing the innocent. They are more attractive and more popular than absolutisms that prohibit lying or promise breaking. However, it will be clear, we trust, that our discussion applies, *mutatis mutandis*, to absolutist theories in general. Our discussion will not be especially friendly to absolutism; the message will be that absolutism has serious trouble giving a plausible answer to our question. But before we enter the argument proper, we need to remark on a key distinction among "oughts."

1. TWO SENSES OF "OUGHT"

How should the epistemic state of an agent affect what they ought morally to do? In one sense of "ought," the answer is not at all; in another, the agent's epistemic state is crucial. We will make the point with a simple example.

³ Thanks here to audiences and the editors.

⁴ We could complicate the discussion by considering cases where there is doubt about the such-and-such consequences but we will not as the points we wish to make can be made without the complication.

A doctor has a patient with a minor but unpleasant skin condition on his leg. She knows that there is only one drug that is capable of curing the skin condition completely. She knows that in 95% of cases it does so without any deleterious side effects but that in 5% of cases it leads to the loss of the affected limb. There is no way of knowing which group any given patient falls into but, as it happens, our patient is one of the majority. What ought the doctor to do?

It is a mistake to insist that there is only one answer. There are two correct answers, each correct for a different sense of "ought." In one sense she ought not prescribe the drug. A 5% risk of limb loss is not justified by the potential benefit involved. If our doctor prescribed the drug, she would rightly be struck off the medical register. This is the sense of "ought" on which the agent's epistemic state matters. It is the sense of "ought" often called the subjective one. However, there is also a central sense of "ought" on which she ought to prescribe the drug. Suppose that the drug is not prescribed but that, some time in the future, a fail-safe test for whether a patient falls into the 95% or the 5% group is developed. It reveals that the patient in our example is in the 95% group. We will rightly say, looking back, something like "The right decision was made at the time, but we can now see that it would have been best to have prescribed the drug." But if it would have been best, there must be some good sense in which it ought to have been done. This is the sense in which what ought to be done has nothing to do with the agent's epistemic state, and is the sense of "ought" sometimes called the objective sense.

The subjective-objective terminology is unfortunate. It suggests that there is really one concept, namely, of *that which ought to be done*, and the distinction is between what *in fact* falls under that concept, which is what objectively ought to be done, and that which *probably* falls under that concept, which is what subjectively ought to be done.⁵ However, what our doctor ought to do at the time of decision is not that which is most likely to be objectively right. It is 95% probable that prescribing the drug is the action which is objectively right, and that is precisely the course of action which she ought not do given what she knows at the time. There are two distinct concepts we need to inventory and when it is important to bear the distinction in mind, we will use "objective-ought" for the one according to which the drug in our example ought to have been prescribed, and "decision-ought" for the one on which the

⁵ There is a distinction between that which is square in fact and that which is likely square but there are not two concepts *square*.

drug ought not to have been prescribed; that is, she objective-ought to have prescribed the drug and decision-ought not to have prescribed the drug. The "decision-ought" terminology has the drawback of possibly suggesting that we are talking about the right decision process to follow, whereas we are in fact talking about what a subject ought to do given their epistemic situation, but it has the signal advantage of capturing the key idea that it is the "ought" that matters most directly for what a subject ought to do at the time of acting and which determines the question of an agent's responsibility for an action.

The example we used as a reminder of the objective-ought/decision-ought distinction is an example where what ought to be done is determined by consequentialist considerations. But it would be wrong to think of the distinction as an in-house issue for consequentialists that is of no interest to absolutists. To start with, almost every moral theory allows that sometimes what ought to be done is settled by the likely consequences, and whenever that is the case, it is transparent that the distinction between the two "oughts" matters. But more importantly, the question of what to say when there is some degree of doubt about the category a contemplated action falls under is one that has to be addressed by any moral theory. As we noted at the beginning, there is often some degree of doubt about the category an action falls under. This is especially the case when the relevant category is whether or not an action is the intentional killing of the innocent: the probability of innocence very often is a matter of degree. Of course, absolutists are free to insist that their theory is concerned solely with objective-oughts, and so that the discussion (and trouble-making) to come in later sections does not concern them. But if they do this, they turn their backs on an important issue that needs to be addressed by anyone of an absolutist stripe. There will need to be a kind of ethical theory that agrees with them in being absolutists about objective-oughts and, in addition, has something pertinent to say about decision-oughts. Call it "absolutism*" if you like, in which case this paper is a discussion of absolutism*, but we will stick with the shorter name. As a matter of record, the "turning the back" stance is not taken by one prominent group of absolutists: namely, those Catholics who oppose early stage abortion while granting that it is far from certain that an early stage foetus is a person. They argue that it is possible that early stage abortion is not the intentional killing of an innocent person because it is not the killing of a person, and, therefore, that it may not be ruled out by absolutist prohibitions against the intentional killing of innocent persons. But, all the same, early stage abortion ought not to be allowed because we cannot be sufficiently

confident that the early stage foetus is not a person. Their position is, precisely, that early stage abortion may or may not be something that is objectively wrong but it is certainly something that decision-ought not to be done.

In sum, this paper is concerned with the decision-ought question for absolutist theories, but except where required by clarity, we will drop the prefix "decision" in what follows. Indeed, in everyday life, we typically read the "oughts" we discuss with the prefix taken for granted: as in "If there is doubt about the safety of an inoculation, we ought not make it compulsory" or "It is true in general that we ought not take unjustified risks." These contentions only make sense as ones directed to what decision-ought not to be done. As heralded, our discussion will be largely framed in terms of absolutisms that have an absolute prohibition against intentionally killing the innocent, and we will often refer to the case we now describe.

II. A CONCRETE EXAMPLE

A skier is heading in a direction you know for sure will trigger an avalanche that will kill ten people. You know the only way to save the ten people is for you to shoot him. The probability that the skier intends to trigger the avalanche and kill the ten people is $1-p$. We can agree that our target absolutist theory says it is right for you to shoot if it is certain that the skier intends to kill the ten, that is, if $p = 0$, for in that case you would not be killing someone innocent—you would be protecting the ten in the only way possible from an unprovoked attack. We can agree that our target theory says that it would be wrong for you to shoot if you are certain he simply happens to be skiing in that direction, that is, if $p = 1$, for then you would be intentionally killing someone innocent, and that is never right no matter how many you will be allowing to die by your failure to shoot. The number of lives that would be saved in the example as described is ten, but of course the distinctive position of absolutism is that the number does not matter: it is never right intentionally to kill the innocent no matter how many lives would be saved by doing so. Our question is, What should the theory say for other values of p ?

III. THE INFINITE DISVALUE APPROACH

Perhaps the simplest absolutist answer to our question is to hold that whenever there is any chance that an action violates an absolute prohibition, the action ought not to be performed. This is the answer suggested by the absolutists' case against early stage abortion summarized above. In our example, the answer would prohibit shooting the skier whenever there is any chance that he is innocent, whenever,

that is, $p < 1$. One way of implementing this answer is to assign infinite disvalue to intentionally killing the innocent and some finite disvalue to allowing people to die. For then the expected disvalue of the shooting—that is, the product of the disvalue of intentionally killing the innocent times the chance that the shooting is an intentional killing of the innocent—will exceed the disvalue of allowing others to die, no matter how many others die and how certain it is that they will die, provided there is some chance that the shooting is indeed an intentional killing of the innocent. It will, on this approach, be impossible to make the action that has some chance of being the intentional killing of someone innocent the right thing to do by making the number allowed to die by refraining from shooting large enough—the numbers allowed to die will be irrelevant, just as absolutists typically say.⁶

The trouble with this response is that there is nearly always some greater than zero chance that someone is innocent. All the evidence may be against them, but induction from the past record of overturned verdicts in cases that looked watertight at the time tells us that there is nearly always some chance that someone who looks clearly to be guilty is in fact innocent. We will get the result that it is never, or hardly ever, right to shoot the skier. Indeed, it will be hard to find any cases where it is right intentionally to kill someone as there is always some chance that the someone is innocent, and a small chance times an infinite disvalue equals an infinite disvalue. We will have a quick (too quick) argument from absolutism against intentionally killing the innocent to an extreme kind of personal pacifism.

The obvious response to this difficulty is to assign a very high but finite disvalue to intentionally killing the innocent. It will then be possible for an intentional killing that is *very* likely the intentional killing of someone guilty to be the right action in the appropriate circumstances. There are two problems with the “very high but finite”

⁶ We get the same result if we work with a utility function that has 0 rather than minus infinity as its lowest value (greatest disvalue) and 1 as its highest value, and assign value 0 to intentionally killing the innocent and some value between 0 and 1 to allowing people to die. For it will of course still be the case that the expected disvalue of the shooting—that is, the product of the disvalue of intentionally killing the innocent times the chance that the shooting is an intentional killing of the innocent—will exceed the disvalue of allowing others to die, no matter how many others die and how certain it is that they will die, provided there is some chance that the shooting is indeed an intentional killing of the innocent. In the one case, the point is that minus infinity times a positive number is minus infinity; in the other that zero times a positive number is zero. A similar point applies to the argument in the next section. We are indebted to the editors.

disvalue approach, depending on the position taken on the mathematics of disvalue.

IV. LARGE BUT FINITE DISVALUE WITH NO BOUND ON DISVALUE IN GENERAL

Any reasonable version of absolutism allows that it is bad to act in a way that leads to death. Refraining from throwing a life belt to a drowning woman is not plausibly excused by saying that it is merely an allowing to die and not an intentional killing. The position is not that there is nothing bad about allowing to die. The absolutist contention is rather that, no matter how bad allowing to die is, it never justifies intentionally killing the innocent.

Now suppose that the way the mathematics of disvalue works is that the disvalue of allowing to die increases as the number that will die increases without any upper limit (that is, the *value* tends to minus infinity). In this case absolutism cannot handle the difficulty created by treating the disvalue of intentionally killing the innocent as infinite by making its disvalue very large but finite. The reason is that by considering a case where enough are allowed to die—we have ten being allowed to die in our skier example as it stands but we could make the number as big as we please—we will be able to construct a case where allowing to die is worse than intentionally killing the innocent. Suppose that the disvalue of intentionally killing the innocent is some large but finite number N . Suppose that in our skier example it is all but certain that the skier is innocent so we can set the expected value at almost minus- N . All we need to do is make the number certainly allowed to die, X , large enough to create a disvalue greater than N (that is, a value less than minus- N). This result is, of course, unacceptable to absolutists; it is precisely what they deny. It makes the numbers allowed to die critical for what ought to be done.

V. THE BIG BAD NUMBER

Absolutists sometimes in effect refuse to discuss the very hardest of choices: the choices that arise in the horrific cases consequentialists describe when seeking to make the case that there are no defensible absolute prohibitions. According to consequentialists, for every kind of action one might nominate as absolutely forbidden, there are horrific cases where following the prohibition has consequences sufficiently appalling to show that the prohibition should not be given absolute status. One way of thinking about the absolutists' (sometime) refusal to discuss these cases is as embodying a certain view about moral disvalue. On this view there is a big bad negative number B which marks a lower bound, a floor, on expected disvalue

in the area of action. Things can get worse and worse but can never drop below B . On any view like this, absolutists have, it seems, a way out of the objection of the previous section.⁷ It need not be the case that making the number of people allowed to die in the skier example big enough can mean that allowing that number to die is worse than that of taking such-and-such a chance that one's action is intentionally killing the innocent. The (expected) disvalue of that chance of one's action being the intentional killing of the innocent may be at the "floor," namely, B .

In order to discuss this proposal we need to spell out in a little more detail the way probability and value interact on the proposals we are considering as possible replies by absolutists to our question about how absolutism should handle the problem of uncertainty.

In the skiing example, you have the choice of shooting or not shooting. There are two possible outcomes of shooting: intentionally killing someone innocent and intentionally killing someone guilty (that is, someone who without justification is seeking to kill by triggering an avalanche). There is one possible outcome for not shooting: allowing ten people to die. The expected value of shooting is:

$$(A) [p \times V(\text{intentionally killing the innocent}) + (1-p) \times V(\text{intentionally killing the guilty})]$$

The expected value of not shooting is:

$$(B) [1 \times V(\text{allowing 10 to die})].$$

The problem noted in section III with giving intentionally killing the innocent an infinite disvalue is that whenever p is nonzero, that is, nearly always, (A) will come out worse than (B), which means that it is never right for you to shoot in the skier case—clearly the wrong answer by all lights except extreme personal pacifism, as we noted. The problem noted in section IV with giving intentionally killing the innocent a finite but very large disvalue is that it threatens to make it the case that increasing the number allowed to die in (B) can make (B) worse than (A) even when p , that is, the probability that shooting is indeed the killing of the innocent, is high (in fact, as high as you like). This would mean that we lose the distinctive feature of absolutism, as we noted.

If the big bad number approach is to avoid each of these (disastrous for absolutists) results, it must insist that there is a high prob-

ability value, p^* , that has two properties. First, if the probability that the shooting is the shooting of someone innocent is below p^* , then the product of that probability with the value of intentionally killing someone innocent is greater than the big bad number B (the disvalue is smaller). That way we allow the possibility that it is sometimes right to shoot—namely, when what would happen, were you not to shoot, has a value below this value while still being above B . For the big bad number theory must allow that when the probability of innocence falls below a certain value, it is right to shoot. Second, it must also insist that when the "activation" probability, p^* , is reached, the absolute prohibition kicks in. On the big bad number theory, that is to say that when $p = p^*$, $p \times V(\text{intentionally killing the innocent}) = B$. The reason is that otherwise we will have the consequentialist-type result that by increasing the number allowed to die from ten to whatever is called for, we can make (B) worse than (A), so making what ought to be done in part a function of how many are allowed to die.

The upshot is that the big bad number approach has to be understood as a *threshold* approach to the question we are posing absolutists. Although it would be wrong to insist that the absolute prohibition on intentionally killing the innocent only comes into play when a contemplated course of action is certain to be the intentional killing of someone innocent—that makes absolutism largely irrelevant in practice as there is almost always some element of doubt—there is, runs the idea, a threshold at which the absolute prohibition kicks in (p^* above); this is the probability level which might be called that at which it is a *moral certainty* that the contemplated action falls under the category in question, as we might put it.⁸ In the skiing example, shooting is absolutely prohibited no matter how many would be saved when it is a moral certainty that the shooting is the intentional killing of someone innocent.

We hazard that many absolutists will find this version the most attractive.⁹ We now turn to its troubles.

VI. THE THRESHOLD VERSION OF ABSOLUTISM

Any threshold version has a delicate balancing act to perform. Set the bar too high, and the version becomes irrelevant in practice; in real life there is always an element of doubt. Set the bar too low, and you make absolutism about not intentionally killing the innocent

⁷ Thanks here to Andy Egan but he should not be held responsible for our formulation.

⁸ Thanks to Kit Fine here.

⁹ It was suggested to Smith by Tony Coady in discussion. We do not know if it is his considered view, however.

too like personal pacifism. But suppose we agree that setting the threshold at 95% is plausible. An obvious question is how an absolutist might give a principled justification for this value. For the difference between 94% and 96% will correspond to a complete shift in ethical outlook on the threshold version of absolutism. At 94% consequentialist considerations can come into play and settle what ought to be done; at 96%, they are absolutely ruled out—no amount of avoided death and suffering justifies intentionally killing someone whose chance of being innocent is 96%. However, we will not pursue this question. But before we leave it, we should emphasize that it would be a mistake to think of the threshold as a function of the available options—or, more precisely, this is not something absolutists can do. A natural thought is that the bigger the numbers allowed to die, the higher we should set the threshold. But that is giving the numbers allowed to die the kind of central role that absolutism precisely rejects.

The problem we especially wish to press for the threshold view is that it cannot handle actions involving multiple people. We will make the point using a modification of the skiing example. Suppose that there are two skiers, *X* and *Y*, whose probability of innocence is individually just below the threshold, 0.95 as we are supposing, and whose paths threaten separate groups of people. On the threshold view, you ought to shoot *X* if the number saved is large enough because the probability of innocence is below the level at which the absolute prohibition kicks in; ditto for *Y*. But, on the version of absolutism under discussion, you ought not shoot *X* and *Y* no matter how many would be saved, as the probability that shooting both is to shoot at least one innocent person will exceed the threshold, or so we may suppose consistently with their individual probabilities of innocence being below the threshold. This means that the threshold version of absolutism violates agglomeration. It allows the possibility that you ought to shoot *X*, and that you ought to shoot *Y*, when it is false that you ought to shoot *X* and *Y*.¹⁰

Agglomeration is very appealing and is widely accepted and many will regard this as trouble enough for the threshold version of absolutism. However, we accept the minority position that agglomeration fails in certain special cases. For us the problem is that the threshold view makes the modified version of our skiing example into an implausible kind of moral dilemma. There is no course of action

¹⁰ The difficulty raised here for a threshold version of absolutism is similar to the one raised by the lottery paradox against analyses of a “yes-no” concept of belief, the one sometimes called *acceptance*, in terms of degree of belief reaching a specified high level

open to you (the potential shooter) which is such that by following it you do all that you ought to do. Here are the courses of action open to you in the example.

- (1) Shoot *X* and do not shoot *Y*.
- (2) Do not shoot *X* and do shoot *Y*
- (3) Do not shoot *X* and do not shoot *Y*
- (4) Do shoot *X* and do shoot *Y*.

Here are the verdicts the threshold view gives for each. If you adopt (1), you do the wrong thing by virtue of not shooting *Y* (although you do the right thing in shooting *X*). If you adopt (2), you do the wrong thing by virtue of not shooting *X* (although you do the right thing in shooting *Y*). If you adopt (3), you do the wrong thing by virtue of not shooting *X*, and also the wrong thing by virtue of not shooting *Y*. If you adopt (4), you do the wrong thing by virtue of shooting *X* and *Y*—as there is no way of shooting *X* and shooting *Y* without shooting *X* and *Y*.

This is serious trouble. It is the job of a moral theory to delineate what ought to be done in a way that makes at least one option correct (“at least one” to cover the possibility of ties). We are talking here about the decision-ought of course, and it is part of decision-ought’s job description to mark out one or more courses of action as correct (although it may at the same time be the case that every course of action available to the agent is bad). There is of course an interesting debate about whether there are situations in which every available course of action is both correct and incorrect in the decision sense: so-called strong moral dilemmas. (Smith is sympathetic to this possibility; Jackson is not.) However, the situation delivered by the threshold version of absolutism is *not* one where every available course of action is both correct and incorrect. There are two available actions that the threshold version says that you ought to carry out, namely, shooting *X*, and shooting *Y*, but it does not in addition say that you ought not shoot *X* or that you ought not shoot *Y*. Instead it says you that you ought not shoot *X* and *Y*. The problem here is not that the threshold view says that this is a strong moral dilemma; the problem is rather that it puts value and probability together in a way that leads to paradox.

(threshold). What happens if we replace mention of degrees of belief by a “yes-no” notion like acceptance in our discussion of absolutism? The answer is that we change the subject. There is a question for absolutism arising from the often realized possibility that we have one or another nonextreme degree of belief that some action falls under the relevant characterization, and you do not answer that question by discussing another (interesting) one.

There are three things that absolutists might say in reply.¹¹ First, they may reply that what matters is the status of each action; if each shooting is morally right, no problem if the conjunction is not right. Getting the correct answers action by action suffices. But then it would matter whether you fired two shots or threw a single hand grenade that would kill both *X* and *Y*. Or maybe the case is one where a single shot could kill both *X* and *Y* by passing in turn through their bodies. We would get the bizarre result that it matters greatly whether you shoot with a single bullet or use two. Second, they may say that there is an available course of action that is right, namely a disjunctive one: shooting *X*-or-*Y*-but-not-both. The trouble for this suggestion is that there exactly two ways of shooting *X*-or-*Y*-but-not-both—by shooting *X* and not shooting *Y*, and shooting *Y* and not shooting *X*, and both involve doing something wrong (not shooting *Y* in the first case, and not shooting *X* in the second). Third, they may say this is one of the cases where what ought to be done is a function of what else is done or not done. There are cases where you ought to do *A* and *B*, but if you do not do *A*, then you ought not to do *B*. A dentist ought to administer a local anaesthetic and then start drilling, but if they have not administered the anaesthetic, then they ought not to start drilling. But the trouble we have made for the threshold version of absolutism is that it has the consequence that in our two skiers' case, you ought to shoot *X* *independently* of whether or not you shoot *Y*, and conversely; at the same time, you ought not to shoot *X* and *Y*.

VII. ABSOLUTISM AS A THESIS ABOUT THE MIND IN ACTION

We have been treating absolutism as absolutely forbidding certain *actions*, but sometimes it is thought of as a thesis about the mind in action, a thesis about which *intentions* a moral agent can properly form and act on. In terms of our target version of absolutism, the contention is not that the *action* of intentionally killing the innocent is absolutely forbidden but that *forming the intention* of killing the innocent is.¹² When it is said that it is *never* right intentionally to kill the innocent, what is meant is that it is never right to form that intention. The suggestion is not, of course, that there is nothing to say about action *per se*;

¹¹ Thanks here and in the previous paragraph to the editors.

¹² As in some versions of the principle of double effect, see, for example, Anscombe, "War and Murder," in *Ethics, Religion and Politics*, chapter 6. The most worked-out version of a "mind in action" ethical theory we know of is Michael Slote's "agent-based" approach in *Morals from Motives* (New York: Oxford, 2001); however, he is not defending absolutism. Slote has detailed discussions of who can and who cannot be read as agent-based theorists. Thanks to Dave Sobel for discussion here.

indeed, the precise relation between the moral status of the intention and of action in accord with it is open to debate in such theories.¹³

There are troubles on two fronts for this suggestion. One arises from the possibility of uncertainty about one's intentions, the other from asking after the content of the intention.

Although it has been suggested to us in discussion that absolutisms of the kind in question avoid the troubles we have been raising on the ground that subjects can be certain of what they intend, this is a mistake. We do not have incorrigible knowledge of our mental states and the point is especially obvious and widely agreed in the case of intentions (and the same goes for motives). It is a commonplace that we can believe that our intention is to serve someone's interests by informing them of an unpalatable truth they need to know, when it is obvious to our friends that our real intention is to upset them. Likewise, subjects can sometimes be wrong about an action's being in self-defense, or about their "real" intention in making a devastating criticism in discussion. The problems we have been raising for absolutism about actions arising from uncertainty about their nature can be raised *mutatis mutandis* for absolutism about intentions.

The second problem arises from asking after the content of the intention. How might the absolutism in question—absolutism thought of as forbidding forming the intention to kill the innocent—specify the forbidden intention?

One possibility is that the forbidden intention is specified as forming the intention to kill someone who is innocent, without any reference to whether or not the agent has any opinion about whether the "someone" is or is not innocent; it is no part of the specification of the forbidden mental state that the agent has a view about the innocence or otherwise of the person or persons about whom the intention to kill may be formed.

This is a very unattractive option. Suppose that *A* is in fact innocent. It would be just as forbidden to form the intention of killing *A* when you have every reason to believe *A* guilty or every reason to believe *A* is attempting to kill you, as it would be to form the intention in a case where you had every reason to hold *A* innocent. Moreover, this specification would mean that absolutism as a thesis about intentions would conspicuously fail to address the decision-ought question we discussed earlier. An agent wondering whether or not it is permissible to form the intention to kill *A* is inevitably faced with that question

¹³ See, for example, the discussion of Kantian theories in C.D. Broad, *Five Types of Ethical Theory* (London: Routledge Kegan Paul, 1930).

in the context of one or another view about *A*'s innocence or guilt. Perhaps the agent's view is that *A* is certainly innocent, or probably guilty, or probably about to attack, or...?

The upshot is that absolutism's specification of the forbidden intention must include something about the agent's beliefs about the putatively intended victim. When you ask yourself what intention you ought to form according to absolutism in our skier example, your opinion about the skier's innocence has to enter the picture. What must be absolutely forbidden is forming intentions to kill people about whom the agent has such-and-such beliefs concerning their innocence. But now we will see is that the various possible spellings out of this schema raise problems broadly analogous to those raised in previous sections. Suppose, first, that the rule is:

R': it is absolutely forbidden to form the intention to kill someone you believe with degree of belief one to be innocent.

It is hardly ever the case that you will have such a degree of belief. Indeed, on certain views about degrees of belief, a degree of belief of one can only be possessed, or only possessed justifiably, when the belief is *a priori* true. Certainly, it will never be the case that you have such a degree of belief in the skier or like examples; there will always be some chance that he is guilty. But then *R'* fails to rule out any real-life intentions.

Suppose, second, that the rule is:

R'': it is absolutely forbidden to form the intention to kill someone you believe with degree of belief *p*^{*} or above to be innocent.

R'' is a "mind in action" version of the threshold view we discussed earlier and is exposed to the same problems. What principle sets the right value of *p*^{*}? And in cases of actions involving more than one person, we will get cases where the subjective probability that one or another person is innocent is above the threshold but the probability of each person being innocent is below the threshold.

Suppose, third, that the rule is:

R''': it is absolutely forbidden to form the intention to kill someone you believe with any degree of belief above zero to be innocent.

In this case, it will never be right to shoot the skier and, more generally, as there is always in real-life cases some chance that the person you intend to kill is innocent, *R'''* will imply that it is never right to form the intention to kill. Absolutism will be reduced in practice to a kind of personal pacifism.

A familiar theme in the philosophy of action is that an action can be intentional under one description but not under another. Killing

the man behind the arras may be the same action as killing Polonius, but it is only intentional under the description "killing the man behind the arras."¹⁴

Is there a halfway house for absolutism that sees it as a thesis not about intentions *per se*, and not about actions *per se*, but about actions under intentional descriptions, that avoids the troubles we have been raising?¹⁵ As far as we know, there is no extant version of absolutism framed in this style but it might run somewhat as follows, framed in the terms of our target example:

R^{iv}: It is absolutely forbidden to perform an action under the description "intentionally killing someone innocent."

But what, exactly, is *R^{iv}* forbidding? On the face of it, actions under certain descriptions, but what kind of animal is an action under a description?

The obvious answer is that an action under a description is simply an action that falls under the description. What is absolutely forbidden are certain actions that have the property of falling under the description "intentionally killing someone innocent." But does it matter whether or not the agent has any opinion about whether the "someone" is or is not innocent? To say it does not matter is to say that an action of killing someone you have every reason to believe guilty is as forbidden as killing someone you have every reason to believe innocent in the case where they are in fact innocent. Some will bite this bullet but the really important problem is that the halfway position so understood fails to address the decision-ought question we are concerned with. An agent wondering what to do is inevitably faced with that question in the context of one view or another about innocence or guilt. If absolutism is to address the decision-ought question, it must say something about actions forbidden under descriptions of the form "intentionally killing someone believed innocent to such-and-such a degree." (How does one act on a stock tip like "Put your money in insurance stocks"? By putting one's money in stocks *believed* to one degree or another to be insurance stocks.) But once degree of belief enters the picture, as it must, the same issues *mutatis mutandis* will arise as arose for *R'*, *R''*, and *R'''*.

How else might we understand *R^{iv}*? The view that one and the same action can be intentional under one description and fail to be intentional under another is the view that being intentional is not a prop-

¹⁴ Donald Davidson, "Agency," in *Essays on Actions and Events* (New York: Oxford, 1980), pp. 43–61.

¹⁵ Thanks to Bob Goodin for convincing us of the need to address this position. It is not one he holds.

erty of actions; instead being intentional under *D* is, for various *D*. Otherwise we would have a violation of Leibniz's law. Perhaps *R*^{iv} is in effect the view that actions do not have the property of being absolutely forbidden: not in the sense that strictly speaking intentions have that property—we have seen the serious problems with that position already—but in the sense that actions have the property of being absolutely-forbidden-under-the-description—"intentional killing of someone innocent." But now we have a theory that says nothing about which actions, and nothing about which intentions, are absolutely forbidden. Of course we might restore the needed property of delivering answers as to what is absolutely forbidden by allowing *detachment*. Detachment would say that from "It is absolutely forbidden to perform an action under the description 'intentionally killing someone innocent'," and "This action falls under the description 'is an intentional killing of someone innocent'," it follows that the action is absolutely forbidden. But now we are back with the problems rehearsed above. We need to ask questions like, "How likely is it that the action falls under the description?" and, "What degree of belief does the agent have that the person is innocent?" We are back where we started.

VIII. ARE WE PRESUMING CONSEQUENTIALISM'S CONCEPTUAL FRAMEWORK?

A natural thought is that we are setting our discussion in a framework that biases against absolutism. Are we not approaching examples like the skiing one in terms of a choice between the *option* of shooting and not shooting, and asking about the probability that the shooting would have as an *outcome* the killing of someone innocent? And is the conceptual apparatus of options and outcomes not the apparatus characteristic of consequentialism? But in fact the puzzle we have raised does not depend on thinking of the shooting as distinct from intentionally killing someone innocent in the case where the shooting is in fact of someone innocent. We do not have to think of the killing of someone innocent as in any sense the outcome of the shooting; "they" can be one and the same action. Take the way we put the problem in schematic form earlier:

Our concern will be with the following kind of question: agent *S* has to choose between doing *A*, where *A* has *p* chance of being *K* and $1-p$ chance of being not-*K*, and doing *B*, where *B* is certain to have such-and-such consequences. What should a theory that absolutely prohibits doing a *K* action say about what happens to what *S* ought to do as *p* rises and falls?

Our problem is generated not by a bifurcation between action and outcome—we can be completely neutral on that subject—it arises

from a point about the epistemology of action rather than its metaphysics. We can be and often are in doubt about whether some action does or does not fall under the description that the absolutist regards as crucial for its moral status. That is what generates our problem, not a potentially controversial metaphysics of action.

FRANK JACKSON

The Australian National University

MICHAEL SMITH

Princeton University