A Challenge to Common Sense Morality*

Jeff McMahan

One of the depressing injustices of contemporary intellectual life is that, if a book addresses moral issues of the profoundest importance in a serious, scrupulous, and rigorous way, then, apart from the recognition it may receive among a tiny coterie of academic philosophers, it is virtually assured a secure niche in oblivion. No trade publisher will touch it. Review sections in newspapers and magazines, swollen with extravagant accolades for the shallow, vaporous, and ephemeral offerings of the popular culture, will ignore it. It makes no difference if the book is accessibly, lucidly, and even powerfully written. If the logic is complex, requiring sustained concentration and careful thought, and if the reader is neither massaged with rhetoric nor drugged with sentimentality, the book is destined for obscurity.

This bitter reflection is of course entirely familiar; it cannot have escaped any observant reader of the *New York Times Book Review*. But one's chronic irritation over the neglect that good work in ethics suffers in popular forums is likely to be exacerbated by a reading of Jonathan Bennett's *The Act Itself*. Here is a book that addresses some of the deepest and most important ethical issues: whether there is a morally significant distinction between making things happen and allowing them to happen (e.g., between killing and letting die), whether it is worse to bring about an effect as an intended means than to bring it about as a foreseen side effect, whether morality demands extreme forms of self-sacrifice, whether it permits or requires partiality toward those to whom one is specially related, whether it may permit or even require one to commit an atrocity if that is necessary to avert an even greater calamity, and so on. Unless one has thought carefully about these issues, one cannot have well-informed moral views about such matters as distributive justice, either within or between societies, euthanasia and assisted suicide, war and terrorism, nationalism, and so on. Bennett is, moreover, a formidably intelligent philosopher; the book is concise and compressed; and the prose is never ponderous or dull but proceeds at a vigorous pace

with colorful and pungent phrasing (for example, a person who has com-
mitted an atrocity for the sake of the greater good is said to have a store
of memories with "an ugly intruder squatting immovably in the middle"
[p. 187]). This is just the sort of book that those who hope to understand
or contribute to the formulation of policy in such areas as medicine and
international affairs desperately need to study. Yet no one except a few
professional philosophers is ever likely to hear of it.

The book is divided into eleven crisp chapters. The first three are
devoted mainly to stage setting, with a good bit of material in the theory
of action. In the next three, Bennett develops an analysis of the distinc-
tion between making things happen and allowing them to happen. This
is followed by two further chapters in which rival analyses are either criti-
cized and rejected or absorbed into his own account. The burden of
these five middle chapters is to show that the distinction between making
and allowing is devoid of moral significance. This thesis presents a di-
lemma whose two horns are explored in the two subsequent chapters. If
we retain the common sense view that making a bad outcome occur is
very bad, and if allowing the same outcome to occur is just as bad, other
things being equal, then it follows that morality is vastly more demanding
than we have hitherto supposed. If, by contrast, we retain the view that
allowing a bad outcome to occur may not be terribly bad, and if making
the same outcome occur is no worse than this, other things being equal,
then it seems that morality permits various acts that we have hitherto
regarded as atrocities. In the final chapter, Bennett explores but rejects
one alternative way of forbidding certain atrocities—namely, by claiming
that it is specially objectionable to *intend* to bring about a bad outcome,
whether by making it occur or by allowing it to occur.

Although various details have been altered, the core of Bennett's
argument will be familiar to readers of his celebrated Tanner Lectures.[1]
Bennett believes that our intuitive discriminations between instances in
which one makes something happen and instances in which one allows
something to happen (e.g., between harming a person and allowing a
person to be harmed, killing a person and letting a person die, and so
on) are systematically tracking a single, sharp distinction. This distinc-
tion is not well captured by any of the standard contrastive locutions in
ordinary language (doing and allowing, acting and omitting, causing
and not preventing, and so on). Bennett therefore coins two technical
phrases to express the distinction. Conduct, he says, may be either *posi-
tively relevant* or *negatively relevant* to the occurrence of an outcome. The
details of his account of this distinction are intricate and sometimes quite
technical; I cannot reproduce them here. But the gist of the account is

1. Jonathan Bennett, "Morality and Consequences," in *The Tanner Lectures on Human
Values,* vol. 2, ed. Sterling McMurrin (Salt Lake City: University of Utah Press, 1981), pp. 47–
116.

really fairly simple. Suppose that if Agent moves his body in certain ways, outcome O will occur, whereas if he moves in certain other ways, it will not. If O occurs, and if the set of possible movements by Agent that would have resulted in the occurrence of O is small compared to the set of movements that would not have resulted in O, then Agent's behavior was positively relevant to O. If, by contrast, the set of possible movements that would have resulted in O is considerably larger than the set of movements that would have resulted in its not occurring, then Agent's behavior was negatively relevant to O.

Bennett illustrates the distinction with a set of examples that he uses throughout the book for a variety of purposes. In these examples, "a vehicle is on a ground that slopes down to a cliff top."

> *Push:* The vehicle stands, unbraked, on the slope; Agent pushes it; and it rolls over the cliff edge to its destruction.
> *Stayback:* The vehicle is already rolling; Agent could but does not interpose a rock which could stop it; and the vehicle rolls to its destruction.
> *Kick:* The vehicle is rolling to a point where there is a rock that can bring it to a halt. Agent kicks away the rock, and the vehicle rolls to its destruction. (p. 67)

In Push, only a tiny proportion of the many ways in which Agent could have moved would have resulted in the vehicle's destruction; his action was therefore positively relevant to its destruction. In Stayback, virtually all of the ways in which he could have moved would have resulted in the vehicle's being destroyed; his action was therefore negatively relevant to its destruction. Finally, in Kick, only a few of the ways in which Agent could have moved would have resulted in the destruction of the vehicle; hence Agent's action was positively relevant. In each of these cases, Bennett's distinction delivers the intuitively correct classification. Even though "it is natural to say that in Kick Agent lets or allows the vehicle to roll on down," it seems clear that what Agent does belongs with his conduct in Push on the positive side of the generic positive-negative divide (p. 67). This shows, Bennett notes, that the distinction we are after does not coincide with any contrast that simply stations a verb such as 'let' or 'allow' on the negative side.[2]

Bennett contends that it is this distinction that guides our intuitive discriminations between, for example, an agent's doing harm and his allowing harm to occur, between an agent's killing a person and his letting a person die, and so on. But the distinction is simply between acts that are among the great majority of possible movements that would re-

---

2. Because there are instances of action that are intuitively classified as "makings" but are also naturally described as allowing an outcome to occur, I will often refer to the general contrast that Bennett is exploring as the "generic" positive-negative distinction rather than as the making-allowing distinction.

sult in O and acts that are among the small minority of possible move-
ments that would result in O. This is clearly without any moral signifi-
cance. The moral status of an act that results in O could not possibly be
affected by whether the majority of the other possible movements that
the agent could have made instead would also have resulted in O. Ben-
nett thus "reluctantly" accepts that the distinction between making and
allowing is morally neutral (p. 139).

This, then, is the main argument of the book: our common sense
discriminations between making and allowing are based on a deeper dis-
tinction that is utterly lacking in moral significance; therefore our ten-
dency to find a moral difference between harming and allowing harm to
occur, or between killing and letting die, is misguided. Significant revi-
sions in our moral thinking appear to be required.

This argument invites obvious questions. For example, if the argu-
ment is correct, how can we account for the fact that we have been sys-
tematically misled into believing that there is a moral asymmetry be-
tween making and allowing? The idea that there is such an asymmetry is
ubiquitous. To the best of my knowledge, there are no societies, no cul-
tures, past or present, in which failing to prevent a harm is regarded as
morally equivalent, other things being equal, to causing the harm or in
which the failure to save a life is treated as equivalent, other things being
equal, to killing a person. If people in these societies have all along been
tracking the morally neutral positive-negative distinction, albeit uncon-
sciously, how have they managed to be so obtuse as to attribute deep
moral significance to its surface manifestations, such as the distinctions
between doing harm and allowing harm and killing and letting die?

There is an obvious reply to this challenge. Most of us accept that
mass delusion does occur. Throughout all of human history, virtually ev-
ery society has believed in the existence of a god or set of gods. Some of
us, myself included, think that all of these beliefs have been delusions.
And even those who believe in some particular religion are normally
committed to believing that the adherents of the many other known re-
ligions, past and present, have been systematically deluded. Credulity is
a besetting vice in human beings.[3] Nor is religion the exclusive province
of mass delusion; it occurs in morality as well. During most of the nine-
teenth century, for example, virtually all whites in the American South
believed in the permissibility of human slavery.

Producing other instances of mass delusion does not, however, fully
answer the challenge to Bennett's argument. For, even if religion is
a delusion, its widespread acceptance is readily explained. People are
strongly disposed to believe that things are the way they wish them to be,
and religion offers a vision of the world that is more to people's liking

---

3. For discussion, see Carl Sagan, *The Demon-Haunted World* (New York: Random
House, 1995).

than the one our senses and sciences reveal. A belief in the acceptability
of slavery is also unsurprising in view of the importance of that institution
to the interests of those who maintained it. But it is not so obvious why
anyone would believe that there is a moral asymmetry between making
and allowing if all that our intuitions are ultimately responding to is Ben-
nett's positive-negative distinction.

With a bit of work, however, one can begin to formulate reasonably
plausible explanations of why we believe that there is an intrinsic moral
difference between making and allowing (other than, of course, that we
are simply apprehending the truth that there is such a difference). It may
be, for example, that the asymmetry between making and allowing is one
of the mechanisms our morality has evolved for limiting the scope of
individual responsibility. The idea that we can, in general, avoid serious
wrongdoing by avoiding *doing* certain things is one way of carving out a
limited domain of personal responsibility, leaving substantial space in
which we are largely free from the imperious demands of morality. An-
other suggestion is that we have tended to converge on a norm that con-
demns doing harm more stringently than allowing harm to occur be-
cause it is in everyone's interests to do so. Gilbert Harman, who believes
that a morality just *is* a set of norms or conventions that self-interested
people have tacitly agreed to adhere to, has offered this account:

> Whereas everyone would benefit equally from a conventional prac-
> tice of trying not to harm each other, some people would benefit
> considerably more than others from a convention to help those
> who needed help. The rich and powerful do not need much help
> and are often in the best position to give it; so, if a strong principle
> of mutual aid were adopted, they would gain little and lose a great
> deal, because they would end up doing most of the helping and
> would receive little in return. On the other hand, the poor and
> weak might refuse to agree to a principle of noninterference or
> noninjury unless they also reached some agreement on mutual aid.
> We would therefore expect a compromise . . . [that] would involve
> a strong principle of noninjury and a much weaker principle of
> mutual aid—which is just what we now have.[4]

The trouble with these explanations, as well as others of the same
sort, is that they are doubtfully compatible with the view that the distinc-
tion between making and allowing is morally neutral. For, at least accord-
ing to various nonrealist understandings of morality (and Bennett is
avowedly a nonrealist), the sort of indirect significance that these expla-
nations attribute to the making-allowing distinction is really the only
kind of moral significance that a distinction can have.

Bennett himself offers a rather different account of how, in thinking

4. Gilbert Harman, *The Nature of Morality* (New York: Oxford University Press, 1977),
p. 111.

that the making-allowing distinction is morally significant, "we could easily be making a moral mistake." It might be that "whenever someone discriminates morally on the basis (she says) of the difference between doing and letting, causing and allowing, or the like, she has something in mind that she thinks is present in the case and believes makes a moral difference there; but she is wrong if she thinks that this difference comes from a single factor that is present and operative in all these cases" (p. 79). This diagnosis, as I understand it, is that when people imagine that their intuition is responding to the difference between making and allowing, it is in fact responding to some other factor or factors that they have misidentified and that this factor or set of factors may be different in different cases (that is, people may mistake a variety of other factors for the making-allowing distinction). Elsewhere Bennett identifies various factors that are contingently associated with either making or allowing and that might be doing the moral work that people mistakenly attribute to the difference between making and allowing (e.g., it is normally less costly to Agent to avoid making outcome O happen than it is to prevent it from happening; there is normally a higher probability of O's occurring if Agent undertakes to make it happen than if he decides not to prevent it from happening; and so on) (pp. 75–76).

Bennett is undoubtedly right that, in the broad range of cases in which we detect a moral asymmetry between making and allowing, we cannot be confident that we are responding to "a single factor that is present and operative" in them all. In part this is because, in some of these cases, our intuitions are being aroused and prodded by something quite different, as Bennett recognizes. But another part of the explanation may be that the difference between making and allowing is not reducible to a single factor—as it must be if Bennett's positive-negative contrast is the correct analysis of it. It might instead be that, while we are indeed responding to somewhat different factors in different cases, we are also right in detecting an asymmetry between making and allowing; for the distinction between making and allowing might be internally complex, compounded from the various factors that engage our intuitions (other than such obviously distinct, though contingently correlated, factors as comparative cost and probability). I will return to this possibility below.

To be persuasive, Bennett's suggested explanation requires considerable elucidation. First, the other distinctions that we supposedly confuse with the distinction between making and allowing need to be identified. Second, either it must be shown that these other differences are morally significant or an explanation must be offered of why they stimulate our intuitions. Otherwise, our attribution of moral significance to the making-allowing distinction will remain mysterious. Third, and finally, it needs to be shown that these other differences in fact coincide closely enough with the difference between making and allowing to ac-

count for our systematic conflation between them. In the absence of such an account of our mistake, Bennett's claim that we have been systematically deluded in attributing significance to what is in reality nothing more than the positive-negative distinction casts doubt on the plausibility of his analysis.

Is it plausible to suppose that the positive-negative distinction is the basis for our intuitive discriminations between making and allowing? Bennett concedes that critics may find counterexamples to his analysis. But he is prepared for this; for he holds that our discriminations between making and allowing follow the positive-negative distinction "only imperfectly, because [we] sometimes drift away from it and use the terminology of making/allowing in ways that have no solid conceptual support." And, anyway, "an intuition that goes against my analysis is not the same as a rival to it" (p. 100). But, even though he is reconciled to the appearance of the odd counterexample, he repeatedly insists that, in the almost thirty years since he published the earliest ancestor of the present analysis, only one counterexample has actually been advanced (pp. 85, 98, and 100).[5] This is an example in which his analysis implies that a person's remaining completely immobile is positively relevant to the occurrence of an outcome. Bennett, refusing to concede that this is in fact a counterexample, argues vigorously that his classification is plausible. I will not dispute this; the example seems to me to lie in a murky area where our intuitions are likely to be confused. There are, however, other counterexamples that I believe are more damaging. I will explore some of these, beginning with three cases that first appeared in Bennett's Tanner Lectures and reappear in the book.

> *Suit:* An African village is in need. Agent launches a lawsuit that deprives them of a thousand dollars they would otherwise have had.
> *Cancel:* Same village, same everything, but this time Agent learns that his accountant thinks he is supposed to sign away a thousand of Agent's dollars to the village, and Agent tells him not to.
> *No-help:* Same village etc. This time Agent could but does not give the villagers a thousand dollars. (P. 103)[6]

In all of these cases, the outcome is that the villagers suffer some deprivation. According to Bennett's analysis, Agent's behavior is positively relevant to that outcome in Suit and Cancel and negatively relevant only in No-help. He notes that "one might hope for a line having Suit

---

5. The first version of his analysis appeared in "Whatever the Consequences," *Analysis* 26 (1966): 83–102; reprinted in *Killing and Letting Die,* 2d ed., ed. Bonnie Steinbock and Alastair Norcross (New York: Fordham University Press, 1994).

6. The wording here is Bennett's, except that I have substituted "Agent" where he uses the first person.

on the left [of the positive-negative line] and the other two on the right: Cancel and No-help, one might think, are just two ways of *not giving money to the village,* and it does not matter that one does and the other doesn't require a fairly specific kind of movement" (p. 103). Cancel, then, looks like a counterexample, which Bennett takes as a challenge to his metric for dividing an agent's "behavior space." His distinction depends on dividing the logical space of possible behavior for an agent in terms of the specific ways in which the agent could move his body at a certain time. Cancel may be thought to require a different metric: "a behavior space that represents . . . things [Agent] could do with [his] money" (p. 103). Bennett considers this option, concluding that "it sounds natural, but it is tailored to fit that trio of cases; it is no use until we can generalize it, and I cannot see how to do that" (p. 104).[7]

The problem with this response is that it presupposes that our analysis of the making-allowing distinction has to employ the device of defining a behavior space for an agent at a time and then finding the right metric for carving it up. But there is more to our intuitive sense of what constitutes a making or an allowing than is captured by the dissection of a behavior space at a given time, no matter what the metric. Background factors matter. In Cancel, for example, it matters that what is in progress at the time Agent intervenes is a transfer of *his* resources to the villagers. If it were someone else's aid that was on the way, Agent's interception of it would clearly be positively relevant to the villagers' deprivation. This is not primarily a point about ownership or property rights but about the source of the aid or assistance that Agent blocks before it can become operative. In Cancel, a process has been initiated that, if it continues, will result in Agent's aiding the villagers, albeit involuntarily. Before the process culminates in the provision of aid, Agent terminates it. In relevant respects, this is analogous to an aborted rescue: Agent stops the process whereby he might have rescued the villagers from their plight; what he does is *not to prevent* their suffering a deprivation. His behavior is, in the generic sense, negatively relevant to the outcome in which they suffer a deprivation.

7. In an earlier discussion of these cases as they appeared in Bennett's Tanner Lectures, I interpreted Bennett as deploying them in an effort to establish the moral neutrality of the making-allowing distinction. But, while he did indeed expressly say that he regarded them as "establishing the neutrality thesis," he also had a further reason for producing them, which was to note that the absence of any moral difference between Cancel and No-help is so obvious that it must surely make one wonder whether the positive-negative distinction is what underlies our intuitive discriminations between making and allowing. I regret having failed to acknowledge this importantly self-critical aspect of his argument. See his "Morality and Consequences," pp. 89–95, and my "Killing, Letting Die, and Withdrawing Aid," *Ethics* 103 (1993): 258–61 (reprinted in Steinbock and Norcross). In the book, the doubt that Bennett raises about his own analysis no longer focuses on the absence of a moral difference between Cancel and No-help but relies instead on our intuitive sense that the agency in Cancel is generically negative.

It seems clear that our sense that the cases should be classified in these ways is not being driven by moral intuition—though Bennett suggests that it is (p. 104). We do not judge Agent's behavior in Cancel to be negatively relevant to the villagers' deprivation because we can detect no moral difference between it and Agent's negatively relevant behavior in No-help. We judge it to be negative because we recognize that to abort a process whereby one would have prevented the occurrence of an outcome simply cannot be interpreted as making the outcome occur.

What Cancel suggests is that Bennett's analysis ignores relevant background considerations. Indeed, attention to background considerations could, if we had further details, lead us to locate Agent's action on the negative side of the divide in Suit as well. For, if the lawsuit would succeed, that suggests that the money that the villagers have in their possession belongs by right to Agent. If he decides not to let them retain it, he simply denies them his aid or declines to prevent their deprivation. There may, of course, be further relevant complexities here. But that is the point: how Agent's behavior is relevant to the outcome cannot be determined by an apparatus that restricts our attention to a single behavior space at a single time.

This claim can be reinforced by consideration of a range of further examples. Again, we may build on Bennett's own offerings. Recall that, in Kick, Agent dislodges a rock that would otherwise have prevented the vehicle from rolling over the cliff. Here the background is that the rock was in place independently of Agent; thus the vehicle would not have been destroyed had it not been for his intervention. This is at least a major part of the reason we judge his behavior to be positively relevant to the vehicle's destruction, despite the fact that there is a clear sense in which he does not so much destroy the vehicle as allow it to be destroyed. But now let us tamper with the background to Agent's act of kicking away the rock.

> *Interpose & Kick:* The vehicle is already rolling; Agent interposes a rock that would stop it but immediately reconsiders, deciding that he would rather watch the vehicle go over the cliff. Agent kicks away the rock and the vehicle rolls to its destruction.

According to Bennett's analysis, when Agent kicks away the rock in Interpose & Kick, his action is positively relevant to the vehicle's destruction. For, at any point after he has positioned the rock in the path of the vehicle, only a small proportion of the ways in which he could move would result in the vehicle's destruction. This analysis treats Agent's kicking the rock in this case exactly the way it does in Kick. Perhaps this may be considered a virtue; for, at the time at which the kick occurs, everything, including Agent's physical movements, is exactly the same in both cases. But our intuitive sense is that it makes a difference that, in Interpose & Kick, Agent has himself interposed the rock that he then

removes. This background detail inclines us to locate this kick on the negative side of the generic positive-negative divide. Agent considers preventing the vehicle's destruction, even prepares the conditions for its survival, but then decides not to prevent the destruction after all. That he prepared but then withdrew a barrier to the vehicle's destruction does not mean that he destroyed the vehicle or made it go to its destruction. His little series of activities in the path of the vehicle leave everything as it would have been had he not intervened in any way at all.

This may be clearer if we consider a further variant.

> *Drop & Kick:* The vehicle is rolling. Agent is carrying a huge rock to a stone hedge he is building. Just as he is crossing the vehicle's path, he accidentally drops the rock. Believing that it would be easier to roll it the rest of the way to the hedge, Agent then kicks the rock in exactly the way he does in Kick and the vehicle rolls to its destruction.

Perhaps even more obviously than in Interpose & Kick, Agent does not destroy the vehicle or make it go over the cliff. But again Bennett's analysis ignores background factors. It considers the question of how the rock got there to be irrelevant and thus implausibly declares Agent's kicking the rock to be positively relevant to the destruction of the vehicle.

Again our intuitive sense that Agent's conduct belongs instead on the negative side of the generic divide is certainly not the product of moral intuition. I assume that it was precisely because they are devoid of moral content that Bennett chose to use cases in which what is at issue is only whether a vehicle will be destroyed.

There is a further curious feature of the way that Bennett's analysis treats Drop & Kick that deserves mention. Compare Drop & Kick with the following variant.

> *Carry:* The vehicle is rolling. Agent successfully carries a large rock to the hedge that he is building, crossing the vehicle's path as he does so. The vehicle rolls to its destruction.

It seems that in this case Agent's behavior is negatively relevant to the vehicle's destruction, according to Bennett's analysis. In Drop & Kick, however, his behavior seems negatively relevant up to the point at which he drops the rock, positively relevant at the time of the kick, then negatively relevant again thereafter. But surely the accidental and fortuitous dropping of the rock, unaccompanied by any change of Agent's purpose, cannot make the difference between his allowing the vehicle to go over the cliff (in Carry) and his making it go over (in Drop & Kick).

Alternatively, it is possible that, according to Bennett's analysis, it is irrelevant whether Agent drops the rock. For it is possible that, while he is carrying the rock across the path of the vehicle, most of the ways in which he could move his body—to pick a flower, dance a jig, and so

on—would involve dropping the rock. If that is so, then it seems that most of the ways he could move once he is in the path of the vehicle would result in the vehicle's being blocked. Hence, if the vehicle is destroyed, Agent's behavior while he was in the path must have been positively relevant to its destruction, *even in Carry*. This understanding of the implications of the analysis would eliminate the seemingly arbitrary difference in the way the analysis treats Drop & Kick and Carry. But the cost is that the analysis now misclassifies both cases.

Let us pass on to some further variants. Note that Kick is a case in which Agent removes a preexisting barrier to the destruction of the vehicle. Consider next another case in which this is true.

> *Lunge:* The vehicle is rolling. Agent, who is a bit hard of hearing, is enjoying a picnic in its path. If he stays right where he is, the vehicle will be stopped by running into him. Becoming aware of the vehicle's approach only a second before it would have hit him, Agent lunges to his left out of the vehicle's path and the vehicle goes over the cliff.

Assume that Agent was seated on the left edge of the vehicle's path, facing away from it. In the brief interval between his hearing the vehicle and the time when it would have hit him, the only movement he could have made that would have got him out of its path is the lunge to his left. According to Bennett's analysis, Agent's lunge is positively relevant to the vehicle's destruction. Perhaps, again, this might be regarded as a virtue, for Agent's behavior in Kick clearly belongs on the positive side; and in Lunge, just as in Kick, Agent simply removes a preexisting barrier to the vehicle's destruction—in this case his own body. Perhaps consistency requires that the cases be treated alike. Yet intuitively it is clear that Lunge belongs on the negative side of the generic divide. Some feature of the background to Agent's removal of the barrier in Lunge distinguishes his action in this case from that in Kick. And, again, it has to do with the source of the barrier. If Agent would himself be providing the barrier, as in Lunge, then his removal of it before it can be effective counts simply as not providing a barrier—i.e., not preventing the vehicle's destruction. But if the barrier is there independently of Agent, then in leaving it there he would not be preventing the vehicle's destruction; the rock would.[8] Therefore his removal of the rock cannot be seen as his simply not providing a barrier; when he removes the rock, he is doing more than not preventing the vehicle's destruction.

The defender of Bennett's analysis might remain adamant when confronted with Lunge, claiming that the removal of a barrier to the vehicle's destruction is an instance of positive relevance irrespective of

---

8. Bennett himself observes that "there is no opportunity in Kick for Agent to prevent the vehicle from rolling. What is true is that if he had left the rock in place, *it* would have prevented the vehicle from rolling" (p. 67).

the source of the barrier. If Agent's removal of the rock is positively rele-
vant to the destruction of the vehicle, then his removal of his own body
must be positive as well. To see that this response is inadequate, we need
a further example. This is not a gratuitous piling on of counterexamples
but will help us to develop a more perspicuous diagnosis of where the
weaknesses in Bennett's analysis lie.

    In Lunge, the interval during which Agent may act to evade the ve-
hicle is highly compressed. If we expand the interval, we produce a sur-
prising result.

> *Leisurely Evasion:* The vehicle is rolling. Agent, who is enjoying
> a picnic in its path, sees from a distance that it is rolling toward him.
> He insouciantly gathers his supplies and moves out of the path of
> the vehicle, which then rolls over the cliff.

In this case, when Agent moves out of the vehicle's path, it seems that
most of the movements he could make would result in the vehicle's de-
struction; for most of the movements he could make before the vehicle
reaches the place where he has been sitting would take him out of its
path. For example, while there are numerous places within the vehicle's
path that he could move to in order to pick a flower, dance a jig, or make
some other set of movements, there are, given the extended interval be-
fore the vehicle reaches the area between the picnic spot and the cliff,
many more places outside the path of the vehicle to which he could go
to make the same set of movements. So, if we examine Agent's behavior
space for the interval between the vehicle's starting to roll and the last
moment when he could prevent it from plunging over the cliff, it seems
that most of the movements he could have made would have resulted in
its going over the cliff. His behavior in Leisurely Evasion is therefore
negatively relevant to its destruction.

    If this is right, then Leisurely Evasion shows that Bennett's analysis
does not imply that, in all instances in which Agent removes a barrier to
the vehicle's destruction, his behavior is positively relevant to that out-
come. Hence it would be premature for the defender of Bennett's anal-
ysis to claim the virtue of consistency on the basis of examining Kick and
Lunge alone. Indeed, if the analysis implies that Agent's behavior is posi-
tively relevant in Lunge but negatively relevant in Leisurely Evasion, this
is a serious defect. For, in both cases, Agent simply moves out of the path
of the vehicle. He has more time to do it in the latter than in the former,
but surely that difference is irrelevant to whether he makes the vehicle
go over the cliff or merely allows it to do so. (We can imagine that when
Agent sees the approaching vehicle in the distance in Leisurely Evasion,
he immediately lunges to the left and lies there panting, just as he does
in Lunge. In that case, the same set of physical movements that are posi-
tively relevant in Lunge are negatively relevant in Leisurely Evasion, ac-
cording to Bennett's analysis.)

The defender of Bennett's analysis might respond in the following way. In Leisurely Evasion, some of the possible movements Agent might make would involve his staying within the path of the vehicle. Many more than this would involve his straying outside it. But still others would involve his moving outside the path and then coming back within it. Taking this last possibility into account, it seems that, on the whole, more of his possible movements would have him in the path of the vehicle at the time of its arrival in his neighborhood than would have him outside it. Therefore his actual evasion of the vehicle is positively relevant to its destruction.[9]

There are two objections to this response. First, even considering those possible movements that would take him out of the path and then back into it, it is still implausible to suppose that a large majority of possible movements open to him during the relevant interval would result in his being hit by the vehicle. (After all, some of his possible movements might get him out of the path, back into it, and then back out again.) At best, one might reckon that about half the movements he could make would get him out of the vehicle's path, while the other half would find him in it at the critical moment. But in that case—with his behavior space about equally divided between movements that would result in the vehicle's destruction and those that would not—his actual behavior would presumably be neither positively nor negatively relevant to the actual outcome (i.e., destruction). So the inconsistency between the classification of Lunge and the classification of Leisurely Evasion remains.

Second, even if one could twist or torture the analysis into implying that the Agent's behavior is positively relevant to the vehicle's destruction in Leisurely Evasion, one would be extracting a uniform classification of Kick, Lunge, and Leisurely Evasion at the cost of having the analysis yield utterly implausible classifications of the latter two. It is clear that, in both Lunge and Leisurely Evasion, Agent's behavior belongs on the negative side of the generic divide. Both are paradigm cases of allowing the vehicle to be destroyed. While Bennett's analysis probably does locate Agent's behavior on the negative side in Leisurely Evasion, it is hard to see how it could be induced to classify his behavior as negative in Lunge.

So far I have been assuming that, in Leisurely Evasion, Bennett's analysis would have us divide Agent's possible movements within a behavior space that covers the whole of the interval between the vehicle's beginning to roll and the last moment at which he might have prevented it from going over the cliff. But there seems to be an open question here about the specification of the appropriate temporal interval. Perhaps we should specify a series of behavior spaces, each indexed to a shorter tem-

9. I owe this suggestion to Bruce Bethell.

poral interval. But it is hard tó make sense of this suggestion in this case. Suppose we divide the interval between the vehicle's beginning to roll and the last moment it could be blocked from going over the cliff into ten shorter intervals of equal length. In that case, whether Agent's behavior is positively or negatively relevant during the tenth interval depends on where he is in relation to the path of the vehicle, which in turn depends on where he has moved during the preceding intervals. If, for example, he finds himself in the immediate path of the vehicle, the situation in the tenth interval may be identical with that in Lunge, in which case his behavior in that interval will be positively relevant if the vehicle is destroyed—even if it was negatively relevant during earlier intervals. This seems hopeless.

Alternatively, we might expand the interval in Lunge. Suppose we take Agent's behavior space to cover the interval between the vehicle's beginning to roll and the last moment that it could be stopped from going over the cliff, as we do in Leisurely Evasion, rather than the interval between Agent's becoming aware of the vehicle's approach and the last moment it could be stopped. In that case, the situation in Lunge is relevantly like that in Leisurely Evasion, and Agent's behavior in Lunge becomes negatively relevant as well. This is, of course, the desired result. Perhaps I have been mistaken all along to assume that Bennett's analysis applies to Agent's behavior in Lunge only at the time that Agent performs the act that seems to determine the fate of the vehicle. Perhaps instead one should see the entire range of Agent's possible movements over the longer interval as what determines the fate of the vehicle.

This seems plausible. But it would be a mistake to suppose that we have resolved the problem of the time-frame in a way that enables Bennett's analysis to avoid embarrassing counterexamples. For the problems can be reintroduced through minor modifications of the examples. In Lunge, for example, if the vehicle is parked immediately behind Agent's picnic spot, then the interval between the time that it starts to roll and the time it would hit him is again compressed, so that the only movement that will get him out of its path is a lunge to the left. Here Bennett's analysis implausibly implies that Agent's lunge from the vehicle's path is positively relevant to the vehicle's destruction and there seems to be no principled way of fiddling with the time-frame in order to avoid this implication. And, when we compare this variant of Lunge with Leisurely Evasion, we again find that Bennett's analysis fails to classify Agent's behavior the same way in each. Yet the only difference between the cases is that in Leisurely Evasion the vehicle is farther away when it starts to roll.

Because circumstances themselves may dictate an awkward or inconvenient time-frame, Bennett's analysis remains vulnerable to counterexamples that manipulate the temporal variable. Simply varying the length of the interval in which an agent may act may affect the way in

which his or her behavior is classified by Bennett's analysis, even if all other variables are held constant. This, together with the fact that it ignores relevant background considerations, seems fatal to the analysis.

My claim is that Bennett's positive-negative distinction fails to map onto our intuitive discriminations between making and allowing. As I mentioned, Bennett anticipates this objection in part by conceding that our intuitive discriminations sometimes stray from their mooring in the positive-negative contrast. It seems, however, that the divergence between the positive-negative distinction and our intuitive discriminations is too wide, and too well explained by flaws in the analysis, to be attributed to a few vagaries of our psychology. Nor is it plausible to suppose, in the cases we have been considering, that our taxonomic intuitions are being pulled off course by the gravitational tug of morality. The problems with Bennett's analysis lie in its leaving background factors out of account and its sensitivity to considerations, such as the amount of time in which an agent may act, that are in fact irrelevant. I conclude, therefore, that the positive-negative contrast is not what our intuitions about making and allowing are tracking. If that is right, Bennett's main argument collapses; for the fact that the positive-negative distinction is morally neutral does not show that our discriminations between making and allowing are. Our intuitive discriminations are being guided by something else, which may well have moral significance.

Bennett distinguishes three possibilities:

> (1) The ordinary uses of 'make' and 'allow', and of the other . . . [cognate] locutions, reflect an underlying jumble with no systematic core. (2) People use those locutions to express a single, clean, systematic distinction, and I have failed to describe it. (3) In their uses of the locutions in question, people are guided by a clean, deep concept, but only imperfectly. (Pp. 99–100)

I agree with Bennett that the first possibility is unlikely: our intuitive discriminations may be subtle and their source elusive, but they are far from arbitrary. In the third possibility, the clean, deep concept is of course the positive-negative distinction. If I am right that that is not the basis of our intuitive discriminations, this leaves the second possibility.

As we have seen, Bennett throws down the gauntlet to those who would dispute his claim to have uncovered the basis of our discriminations: "An intuition that goes against my analysis is not the same as a rival to it." If one wishes to defend the moral significance of the distinction between making and allowing, it is of course insufficient to argue simply that Bennett has failed to show that the distinction is morally neutral. It may be tempting, given that our intuitions strongly favor the idea that the distinction has significance, to locate the burden of justification with the person who claims that it lacks significance. And there is a little something to this: if our intuitions have any normative authority at all, the

burden of justification must be somewhat heavier for the person who denies what they resoundingly affirm. But obviously one cannot rest content in the belief that there is a moral asymmetry between doing and allowing by simply challenging the other fellow to prove that one is wrong.

So those of us who think it likely that the distinction does have significance have work to do—work that Bennett has facilitated in his chapters exposing the defects of various other proposed analyses. I think it is a mistake, however, to assume that a rival account of the basis of our intuitive discriminations must offer "a single, clean, systematic distinction." It seems likelier that the basis of our discriminations is complex and messy, incorporating an array of diverse factors that may not be systematically related. If morality in fact has multiple sources—conventions that arise as a result of tacit social bargaining, common dispositions of evaluation that have evolved through natural selection, the operation of sympathy, and so on—then perhaps one should expect that its elements will often be untidy.

Much as I would like to, I cannot offer a fully developed analysis; perhaps, however, I can nudge matters forward in a modest way. It seems to me that the most promising line of thought to which we may direct our efforts is a proposal that Bennett attributes to Frances Kamm and that has roots further back in the seminal work of Philippa Foot. Since I am not concerned with defending the details of Kamm's version of this proposal, I will refer to the general kind of account in which I am interested, and of which her proposal is a prominent variant, as the *Dependence Account*.

The proposal, in Bennett's presentation, has two elements. The first is that, for Agent to allow Victim to suffer an outcome O, Victim must antecedently be under a threat of O. (As Bennett points out, we want an analysis that covers good outcomes as well as bad—and indeed outcomes that are neither good nor bad, that combine good and bad features, and so on. But the exposition is simpler if we focus only on bad outcomes. If a promising analysis emerges, it can presumably be generalized.) The second element is that the alternative to Victim's suffering O must be that Victim avoids O "*via* Agent, or *through Agent's efforts*," whereas sometimes, in contrast, the alternative to Agent's *making* Victim suffer O is *not* Victim's avoiding O via Agent or through his efforts (p. 125).[10]

Bennett begins his critique of this proposal by asking what it is for Victim to be under threat of O. His initial suggestion is that this might mean that it is highly likely that Victim will suffer O. But he invites us to consider a case in which Agent is determined to prevent O but improbably changes his mind a tenth of a second before O is due to occur. In that case, he says, it was not highly likely that O would occur; hence Vic-

---

10. Again, I convert Bennett's first person to third person.

tim was not under threat; hence Agent did not allow O to occur. One could quarrel with this objection, but it is not worth the bother. For the moral that Bennett draws from the objection is (with a modification to be noted shortly) surely right—namely, that the relevant probability of O is "its probability relative to the state of the world apart from the agency of the protagonist" (i.e., Agent) (p. 125). He then states what he believes to be the best way to understand "under threat" for the purposes of this proposal: *"There is a threat that [O] will occur means [O's] occurrence is probable relative to (i) the facts of the non-human situation conjoined with (ii) the proposition that the agent will not exercise his agency"* (p. 125).

One can do better than this. The stipulation that the probability has to arise from "the facts of the non-human situation" rules out the possibility that Victim could be under threat from a person other than Agent. Perhaps Bennett inserts this clause in order to accommodate the following possibility.[11] Suppose a causal sequence is in progress that, if unaltered, will result in O. But suppose that there is also a very high probability that another agent will intervene to prevent O from occurring if Agent does not. This might be taken to mean that, even if Agent will not prevent O, it is not the case that it is probable that O will occur; hence Victim cannot be under threat of O; hence even if O does occur Agent cannot be said to have allowed it to occur. Perhaps the stipulation that the probability of O must be relative to the nonhuman situation is intended to preserve the obvious truth that Victim is under threat in these circumstances. But there is another way to ensure this without ruling out the possibility of Victim's being under threat from an agent. This is to say that, in order for Agent to allow victim to suffer O, Victim must be antecedently under threat of O in the sense that there is some probability that O will occur unless events attributable to Agent will prevent it. It is not necessary that the probability of O should be high. If it is very low (perhaps because some other agent is poised to prevent O), Agent may have a good excuse for having failed to prevent O if it does occur.

Clause (ii) of Bennett's proposal is meant to allow for the possibility that Victim can be under threat of O (so that it is possible for Agent to allow O to occur) even if there is effectively a zero probability that O will occur because Agent will in fact prevent it. This clause is nevertheless misleading because of its focus on Agent's *agency*. Recall the circumstances of Leisurely Evasion. There is a sense in which there is no threat to the vehicle if Agent will not move. If he does not move, there is (we may assume) no probability that the vehicle will go over the cliff. But

---

11. There is an alternative explanation. Bennett may want to reduce this element of the Dependence Account to a proposal of Alan Dónagan's to which he is more kindly disposed. The proposal is in Donagan's *The Theory of Morality* (Chicago: University of Chicago Press, 1977).

clearly we want to say that the vehicle is under a threat of destruction and that Agent will allow the destruction to occur if he moves. It is not that the probability of destruction exists only if he will not exercise his agency to prevent the destruction. On the contrary, he could prevent it by *not* exercising his agency—e.g., by falling asleep.

As I will presently suggest, it is an important point about the Dependence Account that it acknowledges that there are many ways in which people can prevent outcomes without exercising their agency or by *not* exercising their agency (in the commonly accepted sense of that phrase).

Bennett next turns to the second element of the Dependence Account: that the alternative to Agent's letting Victim suffer O is for Victim to avoid O via Agent or through Agent's efforts. If, for example, O is destruction, then the alternative to Agent's letting Victim be destroyed is for Victim to owe his continued existence to Agent's efforts. Bennett believes that this element is the proposal's undoing. He writes:

> Add this detail to Stayback: Agent is under great psychological pressure . . . to interpose the rock; not interposing it is the hardest thing he has ever done. In that case, it will not be true that *If the vehicle had survived this occasion, it would have owed its further life-span to Agent's efforts,* if this is taken in its plain meaning. Yet it will still be true that Agent's conduct relates to the vehicle's destruction in a negative or passive or allowing fashion. To get this analysis of Kamm's to fit the cases properly is to construe "through x's efforts" or "via x" as meaning "as a result of what x actively or positively did or made happen." So this analysans succeeds only when helped by the distinction that is supposed to be analyzed. (P. 126)

The Dependence Account need not be circular in this way. "Efforts" is definitely the wrong term for a generic description of Agent's contribution to Victim's avoidance of O—in part because, as the variant of Stayback shows, it is suggestive of straining, even if only psychological, but mainly because it is strongly suggestive of action, or doing, or making, which is precisely what leads to the charge of circularity. But the prevention of an outcome may be attributable to an agent (in the way required by the Dependence Account) even though the prevention is not accomplished through acting, doing, making, straining, or other cognate exercises of agency.

This point can best be illuminated with a pair of examples, both involving Agent and another individual called "Other."

> *Thwarted Theft:* Other has a medical condition that will cause various disfiguring changes in his appearance unless he has an operation within the next week. Having insufficient resources to pay for the operation and being unable to obtain the necessary funds any other way, Other seizes wealthy Agent's wallet, which contains

enough cash to pay for the operation. Agent knows of Other's plight but nevertheless grapples with him and retrieves the wallet. Other shortly undergoes the various changes in appearance.

Successful Theft: The facts are the same as in Thwarted Theft except that Agent does nothing to prevent Other from escaping with his wallet.

I believe it is clear that in Thwarted Theft the relevance of Agent's behavior to the outcome in which Other suffers the changes in appearance is on the negative side of the generic divide. Even though it requires strenuous action on his part to do so, he merely allows Other to suffer the changes. This is not, moreover, an esoteric variety of allowing, as in Kick, in which Agent allows the vehicle to be destroyed not by not preventing its destruction but by removing an obstacle to its destruction. In Thwarted Theft, Agent merely does not prevent Other from suffering the changes. He fails to prevent the outcome by resisting Other's attempt to extract aid or assistance from him by force.[12]

What, then, is the alternative to Agent's allowing Other to suffer the changes? This is the scenario in Successful Theft. Other avoids the changes, or retains his good looks, via Agent's aid, assistance, or support—though, of course, the aid or assistance is not provided fully voluntarily. (I say "fully" because there may be an element of voluntariness in Agent's conscious failure to attempt to retrieve the wallet.) But while Other avoids the bad outcome via Agent, he does *not* do so via Agent's *efforts, action,* or anything that Agent *does* or *actively makes happen.* Instead Agent *allows* his support for Other by not resisting its forcible extraction from him.

In summary, among the defining features of cases involving behavior that is negatively relevant (in the generic sense) to a bad outcome O are that (1) there is some antecedent probability, independent of what Agent might do, that Other will suffer O, (2) O occurs, and (3) the alternative to Agent's allowing O to occur was that Other would have avoided it via Agent's aid, assistance, support, or protection—that is, via forces or events relevantly attributable to Agent. This gives a clear and simple sense of the idea of being under threat and is not obviously circular. Nevertheless there are enormous problems here.

The main problem is of course that everything revolves around a capacious, unanalyzed notion of Other's being shielded from a threat by Agent's aid, assistance, support, help, protection, and so on. How are we to understand this amorphous notion? How do we determine when Other's avoidance of O would be attributable to Agent in the relevant

---

12. On one interpretation, Bennett's Suit is relevantly analogous to Thwarted Theft. There, too, Agent acts to prevent others from using his support to enable them to avoid a bad outcome. To have prevented the outcome, he would not have had to exercise his agency but merely to allow them to keep what was his by right.

way? I suspect that we will have to fill out this notion in piecemeal fash-
ion, by surveying cases and consulting our taxonomic intuitions. And I
predict that the result of such an inquiry will not be "a single, clean,
systematic" account of the nature of allowing but instead an untidy wel-
ter of considerations. If, however, these considerations are what in fact
guide our determination of whether Other's avoidance of O would be
attributable to Agent, then they are what we have to deal with. At least
we will be discovering what really underlies our discriminations between
making and allowing.

It seems clear from the cases we have reviewed thus far that Agent
can provide aid or assistance without actively doing anything—e.g., he
can provide aid for Other by not trying to retrieve his stolen wallet.
He can do so unintentionally or involuntarily—e.g., if it is impossible for
him to recover the wallet that Other has stolen. And he can even do so
unknowingly—e.g., when he is unaware that Other has stolen his wallet.
There are, however, large areas of vagueness here, cases in which it is
unclear whether Other's avoidance of a bad outcome O would be ascrib-
able to Agent's aid, assistance, support, or whatever. If, in these cases, O
occurs, there may be corresponding uncertainty about whether Agent
has allowed O to occur and perhaps some question about whether he
has made it occur.

Here is a case that may initially seem straightforward.

> *Previous Gift:* Other has a condition that will cause him to suf-
> fer disfigurement unless he has an expensive operation. Long be-
> fore his condition was diagnosed, Agent gave him a sum of money
> as a gift. Other saved the money and, as it turns out, it is just the
> amount he needs for the operation.

If Other spends the money on the operation, we would not say that his
avoidance of disfigurement is attributable to Agent or that he avoids it
via Agent or as a result of Agent's aid, assistance, or support. Thus, if
Agent were forcibly to take back the money just prior to Other's spend-
ing it on the operation, we would not classify this as Agent's letting Other
become disfigured. For Agent would not be thereby failing to save Other
from disfigurement; rather, he would be preventing Other from using
his own resources to avoid disfigurement—that is, preventing Other
from saving himself—and this falls on the positive side of the generic
positive-negative line.

But there are complications lurking here, which surface if we con-
sider a further variant.

> *Present Gift:* Other needs a certain sum of money now to pay
> for his operation, and Agent gives him the money now as a gift.

If Other has the operation, this would seem to be a case in which he
avoids disfigurement via Agent or through Agent's aid, assistance, or sup-

port. But if we say that Agent enables Other to avoid disfigurement in this case, why do we not say the same in Previous Gift? The two bequests have the same effect. Suppose, moreover, that in Present Gift Agent were forcibly to take the money back before Other was able to use it for the operation. One is inclined to make the same judgment here as in Previous Gift: that Agent would not be merely allowing Other to suffer disfigurement, that he would instead be preventing Other from saving himself from disfigurement, and that this falls on the positive side of the generic divide. This is because the resource that would otherwise have been used to prevent the disfigurement—the money—would, when Agent takes it, properly belong to Other rather than Agent. But now revert to the original story in which Agent gives Other the money now and Other uses it for the operation. If the resource he uses to save himself— the money—is properly his at the time of the operation, what ground is there for the claim that he has avoided disfigurement via Agent or with Agent's aid or assistance? If the reply is that Agent is the original source of the money, why does this reply not also apply in Previous Gift, making that too a case in which Other's avoidance of disfigurement is attributable to Agent?

In the variant of Present Gift in which Agent forcibly takes back the money, perhaps it is uncertain to whom the money really belongs at the time Agent repossesses it. The rules of ownership may be sufficiently fluid to allow for the possibility that benefactors do not fully relinquish their rights to property they give as a gift until a certain undetermined period of time has lapsed. If this were true, then Agent's taking the money back would count as letting Other suffer the disfigurement rather than as preventing him from saving himself. This possibility illustrates the way in which our determinations about whether behavior counts as allowing an outcome to occur may depend on independent beliefs about what rightfully belongs to whom. And this in turn means that in some cases whether an act counts as a making or an allowing will depend on prior moral beliefs. This alone, however, involves no circularity. Nor does it mean that the distinction between making and allowing is reduced to triviality because the way we draw the distinction in particular cases *expresses* our moral judgments rather than providing a basis for them. It is true that antecedent moral judgments sometimes guide our classification of behavior as making or allowing.[13] In particular, we are disposed to classify as an instance of killing any act of which we strongly disapprove that has death as its outcome. So, for example, when asked whether a parent who fails to feed her baby kills it or merely allows it to die, most

13. For an argument that some writers, myself included, develop their analyses of the distinction between making and allowing so that it conforms to common moral intuitions about, for example, killing and letting die, see Tracy L. Isaacs, "Moral Theory and Action Theory, Killing and Letting Die," *American Philosophical Quarterly* 32 (1995): 355–68.

people will say that she kills it, that failing to feed one's own child is murder. But this tendency (which I agree is mistaken) is not consistently encouraged by the occasional dependence of our classificatory judgments on prior moral beliefs about ownership or entitlement. So, for example, the fact that the food belongs to the parent may challenge rather than support the (in my view mistaken) claim that the parent kills the child by failing to feed it.

I will not attempt to sort out the puzzles raised by Previous Gift and Present Gift but will instead note one further pair of cases that illustrate the difficulties that beset the sort of account I find most promising. The background to the cases is that Agent and Other are both suffering from a disease that will shortly kill them unless they are treated with a certain medicine. They discover that there is one dose, sufficient to cure one of them, at the local pharmacy. There is no other source from which another dose can be obtained in time to save the other.[14]

> *Self-interest:* Agent has a car and Other does not. Agent therefore gets to the pharmacy before Other, buys the medicine, and takes it.

> *Sacrifice:* Agent refrains from going to the pharmacy so that Other can have the medicine.

Is Agent's action in Self-interest positively or negatively relevant (in the generic sense) to Other's death? This may depend on what we say about Sacrifice, which is the alternative in which Other survives. If we say that in Sacrifice Other survives independently of Agent or his aid, assistance, or support, then presumably we should say that in Self-interest the relevance of Agent's behavior to Other's death is positive. This is the case because Agent's intervention makes the difference between life and death for Other; had it not been for Agent's intervention, Other would have survived without Agent's assistance. If, by contrast, we say that in Sacrifice Other survives via Agent or through his aid or assistance, then presumably we should say that in Self-interest Agent simply fails to assist Other, fails to sacrifice himself for Other, so that Agent's behavior is (again in the generic sense) negatively relevant to Other's death.

But it is quite unclear which of these is the best description of what happens in Sacrifice. Given that Agent deliberately lets Other have the medicine when he could have it for himself, it is tempting to say that Other avoids death via Agent or through his aid. We would certainly say this if, instead of merely letting Other buy the medicine, Agent had got to the pharmacy first, bought the medicine, and then either given or sold it to Other. Nevertheless, there seems to be genuine indeterminacy

---

14. A generic version of this example is sketched in Shelly Kagan, *The Limits of Morality* (Oxford: Oxford University Press, 1989), p. 109.

about how the action in both Sacrifice and Self-interest should be classified. And perhaps we should not expect an account of the distinction between making and allowing to issue decisive classifications in cases of this sort. It is perhaps a virtue in an account that, when it finds a case in the deep shadows of a gray area, it leaves it there. Bennett's analysis, by contrast, imparts an adventitious brightness to these cases by unambiguously declaring that Agent's action is positively relevant to Other's death in Self-interest and negatively relevant to Other's survival in Sacrifice.

One could raise further challenges to the Dependence Account by considering further variants of these cases in which, for example, Agent learns of the medicine but conceals the information from Other, in which Agent actively impedes Other's ability to buy the medicine, and so on. Consideration of the latter might then lead to an exploration of other problematic cases in which Agent impedes Other's escape from a threat or, alternatively, facilitates his escape.[15] A careful review of cases of these and other sorts is, I think, necessary before one can be confident that the Dependence Account is on the right track. And even if the Dependence Account does accurately capture our ordinary discriminations between making and allowing, the question remains open as to whether it is plausible to attribute moral significance to the considerations that, according to the account, distinguish making from allowing. All that I can confidently claim here is that the Dependence Account is more faithful to the discriminations we in fact make than Bennett's analysis is.

As I noted earlier, Bennett devotes his final three chapters to some problems raised by the rejection of the moral significance of the making-allowing distinction. The central problem is that a morality that assumes the making-allowing distinction to be morally neutral appears to be intolerably demanding. Bennett rejects what he calls the "morality of self-interest," according to which agents are permitted to give significant priority to their own interests over the interests of others. Thus he is rightly unwilling to accept that we are permitted or required to harm or kill other innocent people in order to protect or promote our own lesser interests. But, if the making-allowing distinction is neutral, we should also not be permitted to *allow* other people to be harmed or to die in order to protect our own lesser interests; therefore we must sacrifice our interests when this is necessary to prevent greater harms to others.

There is an engaging candor in Bennett's confrontation with this problem. In effect, he accepts defeat. "There seems to be no way out of the difficulty," he writes, "except either to accord fundamental moral weight to the making/allowing distinction or to accept the tremendously

---

15. Cases of this sort are discussed in Noam Zohar, "Collective War and Individualistic Ethics: Against the Conscription of 'Self-Defense'," *Political Theory*, vol. 21 (1993), sec. III.

exigent morality's condemnàtion of our conduct." But he holds firm in denying moral weight to the distinction and also rejects the extreme demands of a highly exigent morality: "I am unwilling to hold myself to such a standard" (pp. 162–63). Given his nonrealist metaethics, this may not seem much of a collapse. Earlier in the book he announces that his own pursuit of consistency in ethics is little more than an aesthetic preference: "As a personal matter I want to be guided by rather general moral principles. This desire . . . seems to come from my wish to be whole and interconnected in my person, so that I can understand some of my attitudes as parts or upshots of other more general ones" (p. 21). So his inability to resolve the dilemma may leave him a little less whole than he would like to be, but he finds the personal fragmentation easier to cope with than the loss of those comforts that a highly demanding morality would require him to relinquish.

This strikes me as profoundly unsatisfactory. Perhaps he is right that there really is no credible way of reconciling our rejection of the permissibility of harming and killing to promote our own lesser interests with the conviction that morality leaves us ample space for the pursuit of our own interests and values. But we must not acquiesce in this conclusion without first engaging in an arduous struggle against it. Bennett's own reproach to one of his opponents is apposite here: "Argument must indeed stop somewhere, but this seems an unduly early place to stop it" (p. 82). The threat of the dilemma before which Bennett crumbles provides a powerful incentive to make every effort to determine whether the idea that there is a moral asymmetry between making and allowing can be plausibly defended. This is not to say—as Bennett claims that others have said—that, if accepting that the making-allowing distinction is morally neutral leads to an implausibly demanding morality, this demonstrates that the distinction cannot be neutral. It is only to say that we should not abandon the distinction without doing our very best to rescue it.

Thus far, I have said nothing about Bennett's final chapter on intention and the doctrine of double effect. The connection between this chapter and the rest of the book is that the appeal to the moral significance of intention is another traditional way of attempting to prohibit the commission of atrocities in the service of the greater good. The traditional view is that there is a special moral objection, independent of considerations of consequences, to either making a bad effect occur or allowing it to occur as an intended means to a good end. It comes as no surprise that Bennett vigorously attacks this view as well. Again this lands him in a rather uncomfortable position, and the book concludes with a feeble and factitious effort to find some way of distinguishing morally between acts of terrorism that seek the deaths of innocent people as a means of achieving political goals and acts of war that cause civilian ca-

sualties as foreseen but unintended side effects. But, despite the fact that his conclusions are again intuitively unsettling, the argument of this last chapter is brilliant and powerful, at least as good as anything that has ever been written on the issues it addresses. Anyone who hopes to defend those aspects of common sense morality that attribute significance to an agent's intentions will have to begin by confronting the formidable and ingenious challenge that Bennett has posed.