Positive and Negative Relevance

Jonathan Bennett

[From *American Philosophical Quarterly* 20 (1983), pp. 185–94.]

1. A sketch of the distinction

IF at some time a person's conduct can make the whole difference between probability = 0 and probability = 1 for the obtaining of some state of affairs S, then his conduct is relevant to (the obtaining of) S; or, as I shall sometimes say, for short, *he* is relevant to S. One can also be relevant to S through affecting its probability less drastically than that. Everything I say here can take in that territory as well, but it is helpful to think mainly in terms of all-or-nothing relevance.

A state of affairs may obtain (1) because somebody did such-and-such, or (2) because he did not do so-and-so. For example, the gate at the railway crossing was open (1) because he opened it, or (2) because he didn't close it. (2) doesn't apply unless he could have closed it: the Emperor of Japan didn't open the gate and couldn't have closed it, and so his conduct was not relevant at all to its being open. But I don't confine (2) to people who ought to have prevented S, or who had some duty or obligation to prevent it, or anything like that. I take someone's conduct to be *negatively relevant* to the gate's being open just so long as he could have closed it and didn't.

That makes each of us negatively relevant to millions of states of affairs. But why not? I have been told: 'It isn't useful to make the notion of negative relevance as weak as that. We should at least build into it that the person knew he could prevent S from obtaining.' But if we do, we'll make some important truths harder to say, and thus harder to see. If I didn't know that I could have closed the gate, that may excuse me; and if I didn't know I was opening it, that may excuse me. It won't be easy to compare those two excuses if we work with a concept of negative relevance according to which a person is not negatively relevant to S unless he does know what his conduct is leading to. It may be objected that the ignorance excuse falls very unevenly across the positive/negative divide, being available a million times on the negative side for every one on the positive. But that doesn't mean that comparisons are not worth making; and in any case, you can't even state that objection without using 'negative relevance' in my way. If we refuse to call someone's conduct negatively relevant to S unless he knew that he could make the difference to whether S obtained or not, we shall be definitionally abolishing the excuse 'He was negatively relevant, but he didn't know', stamping into the ground an interesting and perhaps important difference between the two sorts of relevance.

Similarly for the suggestion that we shouldn't call someone negatively relevant to S's obtaining unless it was his job to prevent it. We ought to be able to compare how the plea 'It was not his special responsibility' works on the negative side of the line with how it works on the positive side; but for that we need a concept of negative relevance which allows the plea to arise on that side of the line.

In short, we need a thin, clean positive/negative line; then we can then consider—separately, visibly, in a controlled manner—how it relates to other matters which we think morally important. So we ought not to work with the line between doing or causing and letting or allowing. The verbs 'to let' and 'to allow' are crammed with half-understood complexities, so that the attempt to do basic moral philosophy with them is comparable to trying to use 'earth' as a basic term in chemical theory.

The distinction between positive and negative is sometimes hard to draw. My *staying home* last night is my *not going out*; if something happened in consequence of that, was my relevance positive or negative? There are plenty of cases which do not raise this problem, but what are we to make of the ones which do?

A sceptic might say that the problematic cases warn us that there is no grounded distinction here at all, and that what makes some cases unproblematic—e.g. 'because he opened it' and 'because he didn't close it'—is not fit to be relied on. It is just an accident, he might say, that you have to use a word like 'not' if you are to say in English that someone didn't close a gate: we could have had a word 'ose' whose meaning was such that to ose something is not to close it; and then we could say that the accident occurred because the crossing keeper (positively) osed the gate rather than because he (negatively) did not close it. You might object that 'ose' is negative under the skin, since the sceptic explained it with the aid of 'close' and 'not.' But he could reply that things might have been different in such a way that we had 'ose' first, and explained 'close' in terms of 'ose' and 'not'—for someone to close a gate, we would then say, is just for him not to ose it. And so on. That fragment of the debate shows that there is a real question about how deep this positive/negative distinction lies. Do we pick it off the surface of the language we happen to have, or is there something deeper to it?

I am sure there is something deeper to it, and I have a theory about that. In an article some years ago,[1] I drew a line between the two sorts of relevance (wrongly calling it the line between doing and letting) without appealing to facts about the use of negating words like 'not.' That analytic endeavor drew some criticisms, and provoked some rivals.[2] The rivals need not be discussed here: they were addressed to the verb 'to let'; and although they were not very accurate about it, their whole strategy made them negligible—they were trying to describe one square inch of the surface of English rather than to construct something with which to do moral philosophy. (For example, they took as an unanalyzed primitive the wonderfully cloudy notion of 'alteration in the conditions'.) The chief criticism was that according to my analysis someone could, by not moving, be positively relevant to the obtaining of S. Well, in the course of a recent reworking of the analysis—in which it is deepened and sharpened and

[1] Jonathan Bennett, 'Whatever the Consequences', *Analysis* 26 (1966), pp. 83–102, at pp. 94–95, reprinted in B. Steinbock (ed.), *Killing and Letting Die* (Englewood Cliffs, N. J. 1980).

[2] For example, Daniel Dinello, 'On Killing and Letting Die', *Analysis* 31 (1971), pp. 83–86; Bruce Russell, 'On the Relative Strictness of Negative and Positive Duties', *American Philosophical Quarterly* 14 (1977), pp. 87–97; both reprinted in Steinbock, op. cit.

in other ways also immeasurably improved —I show why that is not a defect in it, and expose the error which makes people think that it is.[1]

(In that reworking, I speak of positive and negative 'instrumentality'—a poorly chosen word which I now replace by 'relevance.' This is just a change in labels: at no stage have I let 'instrumental' have the whole of its ordinary meaning; my use of it has been entirely controlled by the equation of 'X was instrumental in the obtaining of S' with' S obtains, and X could have so behaved that S didn't obtain' or with a probabilistic broadening of that. I am still addressing myself to the right hand side of the equation, but I now put it in such forms as '[the conduct of] X is relevant to [the obtaining of] S'. since that formulation requires less depletion of the ordinary meaning of 'relevant' than the other did of the ordinary meaning of 'instrumental'.)

The analysis I have mentioned does not let us divide all facts into positive and negative, but it does segregate most relevances of conduct into positive and negative, doing this without reference to occurrences of such negating words as 'not.' Because I have that analysis, I am not deterred by the problematic cases or by sceptical doubts about the unproblematic ones.

2. Does it make a moral difference?

There is a division of opinion on whether the difference between the two sorts of relevance—now taking the line between them to be vague common property, perhaps caught by my analysis but not defined by it—has any moral significance. Given that bad S obtains because of how I behaved, am I more to blame, other things being equal, if it is because I did do A than if it is because I did not do B? Am I in worse moral shape, other things being equal, if I was positively than if I was negatively relevant to S's obtaining?

Drop 'other things being equal' and the answer, uncontroversially, is 'probably yes.' Several morally powerful differences tend to accompany that between kinds of relevance, especially a difference having to do with difficulty. It is in general easier to avoid doing a specified thing than it is to do a specified thing; easier to obey 'Don't do A!' than to obey 'Do B!'; easier to avoid opening a gate than to close one. At times when I could have closed the gate but did not, I may have been lecturing or making love or closing another gate; or the gate might be so heavy that in closing it I would hurt myself; or I might have had to sprint to reach it in time. But if I caused the bad state of affairs by closing the gate, it is not likely that my not closing it would have involved any such inconvenience or pain or difficulty as that.

But that difference only tends to go with the line between the kinds of relevance. sometimes a person is negatively relevant to something bad, and could easily, safely, comfortably have avoided so behaving that the bad thing came about—e.g. when to close the gate he needed only to reach out and flip a switch.

The difference between the two kinds of relevance tends also, as I hinted earlier, to bring with it a difference in what is or could be known. To take an extreme example: a woman would have been saved from death by fire if someone had telephoned her at midnight and woken her. She died, then, because nobody woke her: I did not and you did not and nor did the Emperor of Japan. But this does not put us in moral trouble, because we had no reason to suspect that by calling

[1] 'Morality and Consequences', in S. M. McMurrin (ed.), *The Tanner Lectures on Human Value* II (Salt Lake City, 1981), pp. 45–116; especially pp. S4, 66–69. Better: 'Negation and Abstention: Two Theories of Allowing', *Ethics* 104 (1993), pp. 75–96. Best: *The Act Itself* (Oxford University Press, 1995).

that number at that moment we could rescue someone from death by smoke inhalation. Most of our negative relevances are like this: we could have prevented a bad thing from happening, but we had no evidence that this was so; whereas our positive relevances more often involve knowledge about the consequence in question.

It is not satisfactory just to announce that the difference in kinds of relevance is accompanied by differences in ease of avoidance or in likelihood of knowledge: we need an analysis of the positive/negative distinction which enables us to explain why it has these accompaniments. My analysis does that job, among others.

As well as erasing significant differences which systematically tend to go with the line between positive and negative, the 'other things being equal' rider is also meant, of course, to even out differences which happen to coincide with positive/negative in particular cases—e.g. between fairly bad consequences and very bad ones.

Even with the 'other things being equal' rider built in, most people think that it is worse to be positively than to be negatively relevant to the obtaining of a bad state of affairs. Some think otherwise. I have been inclined to think otherwise, on the strength of my analysis of the positive/negative distinction: I have never found anyone, and could not imagine there being anyone, who would attribute moral significance to the distinction defined by my analysis while being clear about its real nature. That is still where I stand, but a reason for doubt has turned up.

3. Trouble?

It came about in this way. Wanting to illustrate and confirm the conclusion I had reached by arguing from my analysis, I constructed a sliding scale of examples—a pair of stories which did not significantly differ except as positive to negative, containing parameters which could be varied so as to slide the conduct in either story anywhere from very wrong to splendidly right.[1] Irrelevant variables were filtered out with the help of an extended investigation of morally significant features which tend to accompany positive/negative.

(1 mention those details so as to distinguish my examples from certain brief, shallow, impressionistic ones which have been claimed to show that positive/negative lacks moral significance.) That spectrum of examples tended to confirm the moral neutrality of the positive/negative distinction: as the story pair was varied along the parameters, people's moral judgments varied accordingly, in the same way on each side of the positive/negative line. But I found that by slightly modifying some of this material I got a new sliding scale of examples which again exhibited the moral neutrality of the positive/negative difference, but in such a trivially obvious way as to suggest that the point had been missed. The new examples still involved the distinction caught by my analysis, but those cheap, easy victories make me wonder whether that is the distinction people have in mind when they imply that positive/negative makes a moral difference.

We cannot do anything with that doubt until a rival account of the distinction is produced. While we are waiting for that, however, another possibility should be examined. It is that the relevant distinction—which is perhaps the one caught by my analysis—sometimes makes a moral difference and sometimes doesn't. If that were so, then my argument from my analysis would be vulnerable. That argument displays a certain distinction and invites others to agree that it, now that we have it in focus, obviously lacks all moral significance. Perhaps that is really a trick:

[1] 4. Ibid., pp. 89–9S.

rather than commanding a clear view of the distinction's moral inertness, perhaps we are viewing it in a manner which hides its intermittent moral power. In looking at the distinction abstractly and in isolation, we may be floating up to an altitude where the details cannot breathe—I mean the detailed circumstances in which, when they do obtain, positive/negative does make a moral difference.

My principal topic in this paper is the thesis that some such distinction as that between positive and negative relevance—not necessarily the one caught by my analysis—sometimes makes a moral difference and sometimes does not.

4. Arguments from examples

If that middle position is right, it casts an ironic light on the labors of those who have thought it important to adduce cases where the distinction does make a moral difference, and those others who have adduced ones where it does not. But that is not why the middle position is important; for arguments from examples are bound to be ineffective anyway, even if the truth is an 'always or never' affair. I shall turn aside to explain why.

Here is a typical example of one side of the debate.[1] One man out of hatred and greed kills his wife by poisoning her; another finds that his wife has inadvertently taken poison and, out of hatred and greed, does not give her the available antidote. We are invited to agree that the two bits of conduct are morally on a par, although one relevance is positive and the other negative; which is supposed to show that the difference between the two sorts of relevance is morally inert.

That argument is far from conclusive; For one thing, it compares two bits of extremely bad behavior; and one might object that the moral difference between them is too small to be detected by our moral sensors in the presence of so much wickedness. Compare two vile men each spending several hours torturing a victim, while one of them did and the other did not make to his assistant a promise which he did not intend to keep. The latter torturer would not strike us as noticeably superior to the other; yet we think it bad to make promises you don't intend to keep.

How could the examples be changed so as to obviate this objection? It is no use removing all the moral input. What is in question is the thesis that moral force is augmented or diminished according as the conduct's relevance is positive or negative. One cannot test this by means of cases which are morally flat to begin with. We could move to a pair involving some good state of affairs: one man gives his wife the antidote, while the other does not prevent her from taking it. But that seems to count against the thesis that positive/negative is morally neutral! (On the credit side of the ledger, that is; things might be different on the debit side.) Most people would think that giving the antidote is at least slightly more commendable than merely not preventing its being taken. It may be replied that there are extra distorting factors at work, and that the moral difference is due to them and not to positive/negative. But that is just my point: many factors must be identified and explained and brought under control before an argument from examples can show much.

Furthermore, arguments from examples, just like my argument from my analysis, are vulnerable to the charge of using a positive/negative distinction other than the one people have in mind when they attribute moral significance to something in the positive/negative area. You cannot avoid the risk of picking on the wrong distinction merely by not

[1] See, for example, Michael Tooley, 'Abortion and Infanticide', *Philosophy and Public Affairs* 2 (1972), pp. 37–6S, at p. 59; James RacheIs, 'Active and Passive Euthanasia', *New England Journal of Medicine* 292 (197S), pp. 78–80, reprinted in Steinbock, op. cit.

saying what distinction you are picking on. As for the usual, casual employment of the verb 'to let' (as in 'I shall argue that there is no moral significance in the difference between killing and letting die')—that is at best a gesture in the right direction, not a precise capture of something fit to be used in basic moral thinking.[1]

The opposing position has also been supported with examples,[2] such as the famous transplant one. Five people are bound to die unless they get organ transplants today; their only chance of survival—and it makes their survival certain—is to take heart and lungs and both kidneys and liver from a single healthy person who has no obligations to any of them, is not offering this sacrifice, and so on. Most people would condemn a surgeon who killed this man and did the transplants, rather than letting the five die. They would prefer his being negatively relevant to five deaths to his being positively relevant to one; and, since five deaths are worse than one, that is evidence that positive relevance to something bad is much, much worse than negative.

Judith Thomson has cast doubt on this argument by comparing it with the trolley example: a runaway trolley is careening down the hillside on a line where there are five people who must die if it reaches them; the only difference we can make is to throw a switch to divert the trolley onto a line where it will kill one person.[3] Once again: negative relevance to five deaths, versus positive relevance to one. But everyone judges the trolley case with a very different intensity, and most judge it in the opposite direction, from the transplant one. Yet *modulo* the positive/negative difference they are the same case. So our intense conviction that the quintuple transplant would be wrong must have some other source than a moral discrimination between positive and negative relevance.

Of course that criticism of the transplant example depends on our having got hold of the right distinction. Perhaps what is at work in the transplant case is not positive/negative (morally neutral) plus something else (morally significant). but rather a subtle positive/negative distinction which is not present in the trolley example. But no one is entitled to that defense of the transplant against the trolley unless he tells us what distinction he is invoking—i.e. what the positive/negative difference is which one story exemplifies and the other does not.

So on this side of the debate, too; arguments from examples need help from a proper account of the nature of the distinction that is at issue; and when that account is available, the odds are that the examples can be demoted from argumentative to illustrative rank. That is why the middle position's importance does not come from its threat to arguments based on mere examples.

5. Reasons and universality

Here is an argument for saying that either the positive/negative distinction always makes a moral difference or it never does. Something makes a moral difference by providing a reason for some change in one's moral judgments; and what is ever a reason is always a reason, since reasons are essentially universal. For example, beliefs make a moral difference: if we are inclined to be down on someone for something bad (S) to which his conduct was relevant, we are given a reason for easing up on him if we learn that he did not think that his behavior would have S as a

[1] See Bennett, op. cit., pp. 48–53.

[2] See, for example, Dinello, op. cit.

[3] Judith Jarvis Thomson, 'Killing, Letting Die, and the Trolley Problem', *The Monist* S9 (1976), pp. 204–17, at pp. 206–07.

consequence—perhaps because he did not realize that it was in his power to affect whether or not S obtained. This may not be a strong reason, or a reason for a large change of judgment; and it could always be outweighed by something counting more strongly the other way. Still, the fact of ignorance is a reason, always a reason, for greater leniency of judgment. And so with reasons generally.

That argument cannot be right as it stands.

Even if it is somehow true that reasons are essentially universal, it is also somehow true that something may be a reason for acting in a certain way on one occasion but not on another. Shall I have a third drink? If I do, I will be tired tomorrow morning; and that is a reason against, because I am scheduled to give a lecture tomorrow and I will do it badly if I am tired. But the fact that a third drink last night would have left me tired this morning was not a reason against that drink, since I have no commitments requiring alertness today.

Most of our practical reasons are like that, applying in some circumstances but not in others. I infer that where they do apply they are not the whole reason, but combine with other facts to make up a whole reason. The thesis 'Once a reason, always a reason' is true not of partial reasons ('It would make me tired in the morning") but of complete ones ('He did not know', 'It would make him angry and resentful', 'It would hurt').

Indeed, it is true almost to the point of triviality. Suppose that the fact *Fa* is a reason for leniency of judgment on action *a*, whereas the fact *Fb* is not a reason for leniency towards action *b*. This difference must be explained by some other difference between *a* and *b*: they could not differ *only* in that *Fa* has a moral bearing on *a* which *Fb* does not have upon *b*, because moral properties are supervenient on non-moral ones. Let the relevant difference be the fact (*Ga* and not *Gb*). Then *Fa* is a reason for leniency because *Ga*; but that is what I mean by saying that *Fa* is a partial reason for leniency, the complete reason being (*Fa* & *Ga*). Perhaps '. . . is a partial reason. . . ' can be made to carry more meaning than that, so that it is less trivially true that what is not always a reason is at most a partial reason; but I am prepared to bet that any defensible strengthening of 'partial reason' would still leave the thesis true, if not trivially so.

6. Stepping-stones and signposts

There are two kinds of partial reason which will not help in defending the controversial thesis that positive/negative has moral significance sometimes but not always.

Firstly, there is the **stepping-stone**—a reason consisting in circumstances Q which are relevant only because they will lead to R which bears directly on the moral question. If (Q) I am tired in the morning, that will lead to (R) my wanting and being unable to lecture well; that is why Q is a reason for my not having a third drink. If the positive/negative difference ever yields partial reasons for a moral judgment, they are not of this stepping-stone kind, and here is why. In giving a stepping-stone reason for a practical judgment, we are calling attention to a fact about what will ensue from a given bit of conduct; but when kind of relevance is in question, the relevant upshot is built right into the question. We don't have monadic predicates '. . . is positive' and '. . . is negative' which we use to divide up bits of conduct independently of their consequences. Rather, we have the dyadic '. . . is positively/negatively relevant to. . . ' which relates conduct to upshots. Thus, positive/negative comes into play only in relation to given upshots, which means that 'The relevance was positive/negative' could never be a stepping-stone kind of reason for a moral judgment.

Secondly, there is the **signpost**—a reason consisting in circumstances Q which are evidence for R, which is more directly a reason, but not by bringing R about. There may be a causal flow the other way. Is it all right for me to open the windows when the phonograph is loudly playing Janacek? No, because (Q) the neighbors' car is in their driveway: that is a sign that (R) they are home, which in turn is a stepping-stone reason for not opening the window. Or Q may point to R without causal flow in either direction, as when facts about kind of relevance give signpost reasons. The fact (Q) that his relevance to S was negative raises the likelihood both that (R_1) he could not easily have avoided S and that (R_2) he could not have been expected to know about his relation to S; and R_1 and R_2 do bear on the moral evaluation of him or of his conduct. But that is the sort of thing we set aside with the 'other things being equal' clause: nobody questions for a moment that there is *that* sort of moral significance to the positive/negative difference.

In short, the positive/negative distinction could not possibly yield stepping stones and is uncontroversially agreed to yield signposts. Neither of those kinds of reason, therefore, has any place in the ongoing debate about the moral force of the distinction.

These two kinds of 'reason' are indeed barely entitled to be so called. Once stepping-stone Q has brought R into the world, R does all the work. And a signpost Q doesn't even *make* R the case, but merely lets us know about it; here it is even clearer still that R carries the whole load. If positive/negative is to be an interesting source of partial reasons for moral judgments, they must be 'reasons' in a fuller sense than that. The fact (Q) about kind of relevance must itself bear part of the moral load, rather than merely generating or pointing to something else which bears it. We need cases where (Q & R) gives a complete moral reason which collapses if either conjunct fails.

7. Rights

A premise about kind of relevance can in that way combine with another premise to yield a conclusion about a *right*. Or so it is maintained by those who give the concept of rights a large place in their moral philosophizing, since they seem to hold that I have a right not to be killed but not a right to be kept alive, a right to marry but not a right to have a wife found for me, and so on. They do not build an absolute positive/negative asymmetry into the very concept of a right: they allow us some rights to be helped (i.e. rights which could be violated through negatively relevant conduct), but hold that these are less extensive than our rights which could be violated only through positively relevant conduct. Philippa Foot, using the point specifically to argue that something like positive/negative makes a moral difference sometimes but not always, endorses 'the principle that the right to non-interference extends further than the right to have one's purposes furthered by others'.[1]

If that is correct, facts about kind of relevance can be partial reasons. We can have a complex fact (Q) about you and me and relations between us, including the fact that my conduct was relevant to your suffering a certain kind of reverse; to this we add the premise (R) that my relevance was positive; and out rolls the conclusion that I violated one of your rights, which does not follow from either Q or R alone. Here Q and R are partners: neither is a stepping stone or pointer to the other.

That, however, does not automatically give positive/negative the most interesting sort of moral

[1] Philippa Foot, 'Killing, Letting Die, and Euthanasia', *Analysis* 41 (1981), pp. 159–60.

significance—the sort which concerns how good or bad conduct is. Suppose we are told (Q) a lot about someone's relevance to a woman's death—his state of knowledge, his range of options, his relations with her, her status—everything that might be morally significant except whether the relevance was positive or negative. Now that final fact is added. If it is (R_p) that the relevance was positive, then I am supposing that that, conjoined with Q, implies the judgment (J_p) that he violated her right to life. But suppose instead that what we must combine with Q is the information (R_n) that his relevance was negative: then we get the judgment (J_n) which is the dual of J_p—that is, it's what you get if you take J_p in a form where it is analyzed out far enough to expose the concept of positive relevance, replace that by negative, and put the resulting package together again. So the fact about kind of relevance makes the difference between J_p and J_n, and I am assuming that J_p implies that a right was violated while J_n does not. But that is not to grant that J_n is a *less severe* judgment on the man than J_p is—merely that it is a *different judgment*. If J_p is true, he violated her right to life, and that is bad; if J_n is true he was the agent of a dual of violating her right to life, and perhaps that is no better.

Corresponding to *every* positive rights violation there is a possible bit of conduct which is its dual. Now, if the thesis of the moral neutrality of positive/negative is true then *no positive violation of a right is any worse than its dual*. That thesis does not attack the concept of a right, even if the concept has built into it some degree of positive/negative asymmetry. Even if positive rights-violations are no worse than their duals, there may still be good reasons for caring about whether a given malfeasance violates a right or not. I do not find the concept of a right helpful in my own thinking about basic moral matters; but someone might have a moral system in which it was useful to codify certain moral judgments under different headings depending on which sort of relevance of conduct was involved. For him a fact about kind of relevance, though it would not affect the gravity of a moral charge, might importantly affect the heading under which the charge was to be cataloged.

Of the moral philosophers who take rights seriously, few would accept this account. Most would say that when judgments about rights embody a positive/negative asymmetry, the latter acquires a significance that bears not just on moral category but on moral level. But why? One who takes that view about moral rights ought to defend the thesis that facts about kind of relevance sometimes affect how bad or wrong a bit of conduct is. It is no defense to point out that such facts can affect whether the conduct violates a right. That is being conceded. The question is whether rights violations are morally worse than their duals.

So far as I know, no rights theorist has tried to fill this gap in the argument. Some, I think, would respond as follows:

> 'There is a degree of positive/negative asymmetry in the concept of a moral right, or in the facts about what rights there are; and you have conceded this. To that I add my intuition—my insight, my knowledge—that rights violations are a specially grave kind of misconduct, and worse than their duals. Argument has to stop somewhere, and mine stops there.'

I can only reply that it seems to me deeply unsatisfactory for someone to insist that positive/negative, without always being morally significant, becomes so in the context of rights, while refusing even to speculate about how the trick is worked, what the chemistry is of this moral molecule, what rouses the distinction from its moral lethargy in this context in particular. Argument must indeed stop somewhere, but isn't this an extravagantly strange place to stop it'?

8. The difficulty excuse

Not long ago I encountered, in an unpublished typescript, another thesis about how facts about kinds of relevance yield partial reasons for moral judgments: 'Blame for a positive act does not depend on how difficult not acting would have been. Blame for an omission must take account of the difficulty, inconvenience, etc. of the omitted act.' Now, omissions are one species of negatively relevant conduct; or, more carefully put, if you omit to prevent S your conduct is negatively relevant to S (though the converse does not always hold). This writer has therefore implied that if someone's conduct was relevant to the obtaining of S which is bad, the excuse 'It would have been hard for him so to behave that S did not obtain' may reduce blame if the relevance was negative, but not if it was positive.

That fits the pattern I am investigating. The upshot's being difficult to avoid (Q) combines with the relevance's being negative (R) to generate a reason for a blame-reducing judgment on the person. The judgment comes from (Q & R), and each conjunct is needed. If you take away (R) the fact that the relevance was negative, then the other conjunct does not have the power to reduce blame, since 'blame for a positive act does not depend on how difficult not acting would have been'.

It is helpful that the writer does not derive this from something about rights, saying on the contrary: 'If someone is wronged, he has a right which is violated. But the wrongfulness of murder seems to be the basis of the right, rather than vice versa.' We have here a genuinely new proposal, not a disguised version of the rights thesis.

That accentuates the need for some explanation of *why* the positive/negative difference should have this sort of moral force. The thesis is that if my conduct is relevant to something bad, and if I could not easily have behaved so that the bad thing did not happen, the latter fact gets no leverage on my blameworthiness if my relevance was positive; as though the positiveness smoothes the moral surface so that facts about difficulty of avoidance can get no purchase on it. It is a striking and extremely interesting thesis, but we can reasonably ask why we should believe it.

There is really no question of our accepting it as it stands. The thesis implies that no facts about difficulty, inconvenience, etc. can reduce blame for positively relevant conduct; but that is too much to swallow, and it conflicts with something said earlier in the same paper: 'There is sometimes good excuse for. . . doing what will lead to deaths; e.g. having a system of rail or air transport.' Presumably the excuse for having such lethal systems is the inconvenience of not having them, so it seems that some restriction must be placed on the thesis that 'Blame for a positive act does not depend on how difficult not acting would have been. . . '.

1 said as much to the author in private correspondence, and when this material appeared in print, in a joint work whose individual authors are not identified,[1] the thesis in question had been weakened: 'Blame for a positive act seldom depends on how difficult not acting would have been. Blame for an omission must take account of the difficulty, inconvenience etc. of the omitted act.' The only change is the addition of the word 'seldom': from a philosophical point of view this is disappointing, since it weakens not by clear restriction but by softening; it fends off counterexamples but does not aid understanding. Still, we are at least being pointed to the possibility of some thesis of the form: 'Blame for an F positive act never depends on how difficult etc.

[1] 'Euthanasia and Clinical Practice: Trends, Principles, and Alternatives', The Report of a Working Party (The Linacre Centre, London, 1982), ch. 3.

whereas blame for any omission, even an F one, must take account of the difficulty etc.' The claim that some such thesis is true is substantive and interesting.

But it can hardly be evaluated except through evaluating a specific thesis of that kind, with a specific value of F. If I could think of any plausible values, I would produce and discuss them; but I cannot. Indeed, I cannot see the thesis as plausible even in the soft form in which it says that blame for a positive act *seldom* depends on how difficult not acting would have been. To offer this without argument is no worse than the comparable dogmatism of the rights theorist, but it is more puzzling. The notion of a moral right is at present enormously popular among moral philosophers, which makes it hard for them to see that its credentials might be asked for. But the thesis about reducing blame through a plea of difficulty of avoidance is not like that: it is a splendidly naked attribution of intermittent moral force to the difference between positive and negative relevance. I am puzzled by anyone's thinking it satisfactory to make such an attribution without giving reasons for it.

Perhaps, disappointingly, it is all a result of mere confusion. Perhaps an interesting, controversial thesis to this effect:

> A truth of the form 'It would have been difficult not to' is more likely to excuse an omission to prevent something bad than to excuse a positive causing of something bad

has been conflated or muddled with the tame thesis that

> A proposition of the form 'It would have been difficult not to' is more likely to be true of an omission to prevent something bad than to be true of a positive causing of something bad.

The latter thesis—that it is in general harder to obey 'Do A' than to obey 'Don't do B'—is uncontroversial: it is asserted even by writers who cannot explain why it is true, as I can through my analysis of the difference between positive and negative relevance. It is quite different from the former thesis. The former distinguishes the positive and negative territories in terms of what moral force they allow to facts about difficulty of avoidance; the other distinguishes them in terms of how many such facts they contain. The writing I am discussing certainly commits itself to the former thesis, not to the latter; but I suspect that the author has conflated the two. That would explain the confident, unargued assertion of something which, once it is seen clearly, does not look at all plausible.

9. Conclusion

The only candidates I have been able to find are those two: the one about rights, and the one about the force of the 'difficulty of avoidance' excuse. In neither case can I see any reason to agree. So I am still unable to find or construct an argued, articulated case for saying that a fact about kind of relevance of conduct can be a partial reason for a moral judgment, except by being a signpost to a reason. So far, my argument from my analysis still stands.