

# Open Research Online

---

The Open University's repository of research publications and other research outputs

## Two distinctions that do make a difference: the action-omission distinction and the principle of double effect

### Journal Item

How to cite:

Chappell, Timothy (2002). Two distinctions that do make a difference: the action-omission distinction and the principle of double effect. *Philosophy*, 77(2) pp. 211–234.

For guidance on citations see [FAQs](#).

© 2002 The Royal Institute of Philosophy

Version: Accepted Manuscript

Link(s) to article on publisher's website:  
<http://dx.doi.org/doi:10.1017/S0031819102000256>

---

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's data [policy](#) on reuse of materials please consult the policies page.

---

[oro.open.ac.uk](http://oro.open.ac.uk)

## **Two distinctions that do make a difference: the action/ omission distinction and the principle of double effect**

### *Abstract:*

Some performances that agents perform are more fully deserving of the name “actions” than others. Some performances are paradigm actions; some performances are not actions at all; many performances fall at intermediate points on a spectrum running from paradigm actions to outright non-actions. In brief, and with apologies for the neologism, there are degrees of actionhood.

Again, agents are more responsible for some of their performances than for others. Some performances are paradigms of agent-responsibility; for some performances, agents are not responsible at all; for many performances, the agent’s responsibility falls at some intermediate point on a spectrum running from paradigm responsibility to outright non-responsibility. In brief, there are degrees of responsibility.

Now degrees of actionhood explain degrees of responsibility; and AOD and PDE point us to cases where the degree of actionhood is lower. This is why it is reasonable to think that they also point us to cases where the degree of responsibility is lower.

**Philosophy 77** (2002), 211-234

### I

For my part, I am puzzled by both sides of the debate about the action/ omission distinction (AOD) and the principle of double effect (PDE). I am puzzled by the complexity and the lack of convergence which is evident in the arguments usually presented by proponents of AOD and proponents of PDE, and by the fact that the proponents of one of these theses are often not proponents of the other thesis as well. I am also puzzled by the claims, commonly made by opponents of these theses, that their philosophical motivation is utterly mysterious, and that their falsehood is utterly obvious.

My puzzlement arises from my perception that there is a single simple way of arguing for both theses, which their proponents do not much propound, and which their opponents do not much oppose. I can’t see why this relatively uncomplicated unitary strategy of argument for PDE and AOD has been so little deployed or discussed in the huge literature on these distinctions. In this paper I present and assess the simple strategy. A conclusion about how good the simple strategy’s chances are of vindicating either PDE or AOD will emerge by the end of the paper. The conclusion that, at any rate, the simple strategy has been unjustly neglected should become clear much sooner than that.

### II

As I understand them, AOD and PDE are the following theses:

**The action/ omission distinction** says that, *ceteris paribus*, we are less responsible (and

possibly not responsible at all) for our omissions than for our actions.

**The principle of double effect** says that, *ceteris paribus*, we are less responsible (and possibly not responsible at all) for our actions under those descriptions under which we do them knowingly but do not intend them, than we are for our actions under those descriptions under which we do them knowingly and do intend them.

If you want to defend these theses then, it seems to me, the obvious strategy is as follows:

STEP 1: Some performances that agents perform are more fully deserving of the name “actions” than others. Some performances are *paradigm* actions; some performances are not actions at all; many performances fall at intermediate points on a spectrum running from paradigm actions to outright non-actions. In brief, and with apologies for the neologism, *there are degrees of actionhood*.

STEP 2: Agents are more responsible<sup>1</sup> for some of their performances than for others. Some performances are *paradigms* of agent-responsibility; for some performances, agents are not responsible at all; for many performances, the agent’s responsibility falls at some intermediate point on a spectrum running from paradigm responsibility to outright non-responsibility. In brief, *there are degrees of responsibility*.

STEP 3: *Ceteris paribus*, degrees of actionhood determine and explain degrees of responsibility.

STEP 4: *Ceteris paribus*, any omission has a lower degree of actionhood than any corresponding action. And *ceteris paribus*, any performance under a known-but-not-intended description has a lower degree of actionhood than that same performance under a known-and-intended description.

STEP 5: *Ceteris paribus*, any omission bears a lower degree of responsibility than any corresponding action. And *ceteris paribus*, any performance under a known-but-unintended description bears a lower degree of responsibility than that same performance under a known-and-intended description.

Steps 1 and 2 of this strategy note phenomena which Step 3 offers to explain by connecting them. Step 5 results from *modus ponens* on Steps 3 and 4. Step 5 simply states the two theses, AOD and PDE, that we are discussing. So AOD and PDE are true if Steps 1-5 represent a sound argument. Do they?

### III

One reason why the strategy just outlined is such an obvious way of arguing for PDE and

---

1. Unless I say otherwise, “responsibility” in this paper always means “*moral* responsibility”.

AOD is that it is really very easy indeed to establish the first three Steps. That there are degrees of actionhood and of responsibility, and that degree of actionhood is in general explanatory of degree of responsibility: these are *platitudes*. What I mean by saying that I couldn't help pouring my coffee all over you, because the bang of the firework made me jump, is that my spilling of the coffee was *less of an action* than it perhaps seemed to you, and *therefore*, something for which I was *less responsible*. What you mean by saying that you acted out of character when you sexually harassed the waiter in the wine bar is that this was a slip of some sort, not an action that "the real you" caused in quite the normal way; hence it was not fully your *action*; hence it was not something for which you are fully *responsible*. What a defendant in court means by pleading "diminished responsibility" for a crime is (not, usually, that he bore *no* responsibility for that crime, but) that he was *less responsible than usual*. Such a plea is meaningless unless it is backed up by evidence that the defendant, at the time and in respect of the crime, was (not, usually, no agent at all, but) *less of an agent than usual*. In politics, law, and private life, our ordinary thinking about agency and responsibility is pervaded by the assumptions that agency generally explains responsibility, and that neither agency nor responsibility is all-or-nothing (although of course there are unequivocal cases of action and non-action, of responsibility and non-responsibility). Unless there is good philosophical reason to doubt it (and here there is none), it seems lost labour to produce a proof of what everyone already believes.

Even if we do not need a philosophical proof of the platitudes expressed by Steps 1-3, it would be nice to have a philosophical explanation of them. Presumably part of the explanation lies in the complexity of the criteria of action. There is more than one test whereby we decide whether or not some performance counts as an action. We ask, for instance, whether the agent *caused* her own performance, or some subsequent state of affairs. We ask whether the agent *intended* her performance, or some subsequent state of affairs (under some salient description). And we ask whether the agent had *the ability to do otherwise*<sup>2</sup>: that is, whether she had open to her a decent range of feasible alternatives to that performance and its salient consequences.

Even if causation, intention and ability to do otherwise are the only tests for action, and even if each of these three tests admits only of a Yes or No answer, this already means that there can be degrees of actionhood (though it does not yet mean that we have a defence of PDE or AOD). For there are surely possible cases where some performance passes one or two of our three tests, but not all of them. If I spill your coffee by falling accidentally out of a window on top of you, I cause your coffee to spill, but I do not intend it and could not do otherwise. If I try to spill your coffee by jogging your elbow, but my jab at your elbow misses, then I intend what I do (or something very like what I do), and moreover I could do otherwise than I do; but I do not cause your coffee to spill, because it does not spill. And so on. All of these cases, and indefinitely many others of utterly familiar sorts, will fail to be *fully* actions; but they will be action-like in *some* respects.

---

2. This remark is not meant to be controversial: I am not here raising questions about free will, alternative possibilities, or Frankfurt-style situations (see H.Frankfurt, "The Principle of Alternate Possibilities and Moral Responsibility", *Journal of Philosophy* 1969). Readers who have strong views about these issues should mutate the mutanda by saying that I *appear* to have such alternatives, or whatever.

Obviously there will be many more possible degrees of actionhood, if causation, intention and ability to do otherwise are *not* the only tests for action, or if some or all of these tests admit, not only of Yes or No answers, but of a *spectrum* of answers. I don't know whether there are other tests for action besides the three questions I listed two paragraphs back; so I shall leave that open. But clearly there is a spectrum of answers to each of the three questions. There are degrees of *causation*: in all sorts of imaginable ways, I can be the main cause of an action or a subsequent state of affairs, or just one contributory factor. There are degrees of *ability to do otherwise*: in all sorts of imaginable ways, I can have a generous range of highly feasible alternatives to my actual performance, or a niggardly range of nearly unfeasible alternatives.<sup>3</sup> *Intention* is trickier; yet perhaps there are even degrees of intention. Recall my example of action out of character. There can be questions about whether it is the *real me* who chooses to have a thirteenth pint of lager, or some laddish sub-persona in my head with whom "the real me" is properly less than one hundred per cent. identified. There can also be questions about how definitely and decisively I intend what I do; and there can be questions about the match or lack of it between the description of what I intend and the description of what actually happens.

So if we are looking for spectrums of causation or intention or ability to do otherwise, we not only find some; we find a bewildering variety of such spectrums. But if there are degrees of satisfaction of the three tests for actionhood, there must also be degrees of actionhood itself. (Even if we choose to say that there is a cut-off point below which performances do not count as actions, there will still be a spectrum in which this cut-off point makes a cut.) Likewise, by Step 3, there must also be degrees of responsibility. And so it turns out: the study of excuses, as J.L.Austin noticed long ago<sup>4</sup>, is a rich and varied study.

Heartened by this evidence that the simple strategy starts from mere platitudes, we may turn to the remaining two steps of the strategy, which are perhaps more controversial. Section IV will consider Steps 4-5 as they apply to AOD; Section V will consider Steps 4-5 as they apply to PDE. Section IV and Section V will take strongly contrasting lines: Section IV will attack an intuition that tells against AOD; Section V will defend another intuition that tells in favour of PDE.

#### IV

The intuition telling against AOD that the simple strategy should attack is the intuition that omissions are causes. The relevance of this manoeuvre to the simple strategy should be obvious. One of the three tests for actionhood presented in III is the *causation* test for actionhood. Provided omissions don't generally do better than actions on the other two tests for

---

3. Note that my contrast is between *generous* and *niggardly* ranges of alternatives, not between *large* and *small* ranges. There is a technical reason for this, namely my view that all ranges of alternatives that agents actually have are infinitely large: cp. my "Option Ranges", *Journal of Applied Philosophy* 2001.

4. J.L.Austin, "A Plea for Excuses", *Proceedings of the Aristotelian Society* 1956-7.

actionhood (which seems a safe assumption), we can show that we are in general less responsible for our omissions than our actions if we can prove the half of Step 4 that applies to AOD:

STEP 4A: *Ceteris paribus*, any omission has a lower degree of actionhood than any corresponding action.

We could perhaps prove Step 4A by a relatively cautious route: by showing merely that omissions are generally *less* causally influential than actions. But here I shall consider a bolder route of argument, which leads to the conclusion that omissions are not causes at all.

Let us begin with the word “corresponding”. In what sense, exactly, does any omission O “correspond” to some action A? The answer that people usually have in mind is shown by pairs of cases like the following:

*Bathtub*: When Harry’s father slips on the soap and goes under the bath-water, Harry holds his father’s face down; Harry’s father drowns. When Larry’s father slips on the soap and goes under the bath-water, Larry sits back and smiles; Larry’s father drowns.

*Clifftop*: When Jim wanders too close to the cliff-edge in the fog, Harry pushes him over; Jim falls and dies. When Tim wanders too close to the cliff-edge in the fog, Larry remains studiously silent; Tim falls and dies.

*Famine*: Harry sends poisoned food to the victims of an African famine; ten Africans die. Larry sends no money to the victims of an African famine; ten Africans die.

The general idea is clear enough. In Bathtub, Clifftop, and Famine, Harry and Larry achieve the same nefarious ends. The only difference between them is that Harry achieves his ends by three actions, whereas Larry achieves his by three omissions. And hence the usual suggestion, which is that actions and omissions correspond when they are apt to produce the same results.

One problem with this way of putting things is that it begs the question currently at issue, which is whether omissions have *any* results: after all, “results” is plainly causal language. But I’ll pass over that for present purposes. In this section I shall use the word “results”, contrary to standard usage, in a way which allows it to be neutral as to whether any omission’s “results” are caused by that omission, or are merely subsequent to it or likely to be subsequent to it.

Another problem with this contrast between actions and omissions appears when we compare Famine, as just described, with another case, Famine 2:

*Famine 2*: Harry sends poisoned food to the victims of an African famine; ten Africans die. Larry sends no money to the victims of an Ohioan famine; no Ohioans die.

Suppose that there is a famine in Africa but no famine in Ohio. Then it cannot be equally true that I *act so as to cause* a famine in Africa and a famine in Ohio; for there *is* no famine in Ohio. Nonetheless, it *can* be equally true that I *omit to prevent* a famine in Africa and a famine in

Ohio (e.g. by failing to take necessary<sup>5</sup> precautions in both cases); *despite* the fact that there is no famine in Ohio. The results of my actions have to *exist*; the results of my omissions, apparently, do not.

This point suggests the following general argument against the idea that omissions are ever causes.<sup>6</sup> The results of real causes have to be additive: if X causes E1, and X causes E2, then X causes (E1 and E2). But the results of omissions are not additive. If X omits to prevent E1, and X omits to prevent E2, then it doesn't follow that X causes (E1 and E2); it isn't even clear that it follows that X omits to prevent (E1 and E2). Therefore, omissions are not real causes as actions are.

Why doesn't additivity apply to omissions? Why doesn't "X omits to prevent E1, and X omits to prevent E2" entail "X causes (E1 and E2)"? Because, for one thing, the Ohio-famine case has just shown that, for any event at all, an omission to prevent it does not imply that that event even *happens*. For another thing, for many values of E1 and E2, E1 and E2 will be mutually exclusive in one way or another. First example: E1 = the situation that exactly £20 is stolen from your bank account and E2 = the situation that exactly £200 is stolen from your bank account. The lazy bank clerk does nothing whatever to protect your account's security from internet hackers. So he omits to prevent E1, and he omits to prevent E2. (And he omits to prevent lots of other Es too, corresponding to all the other sums you might have stolen from your bank account. This point reinforces my argument, but to keep things simple, I will ignore it.) If omitting to prevent is causing, it should follow that the lazy bank clerk causes the result that "You have been robbed of exactly £20 and of exactly £200". But it is not even clear that the lazy bank clerk has *omitted to prevent* the logically impossible combination (E1 and E2); still less that he has *caused* it.

---

5. "Necessary", not "*apparently* necessary": it can be genuinely true, and not just seem true, that I ought to do something to prevent an Ohio famine, even when we get lucky and the famine does not happen after all.

6. This argument develops a suggestion of Nicholas Denyer's, on p.44 of D.Oderberg and J.Laing, edd., *Human Lives* (London: Macmillan, 1997). My thanks to Nick Denyer.

Second example: a nefarious machine has trapped Alan and Bill. It will either kill Alan and release Bill, or kill Bill and release Alan, depending on whether Conor presses button A or button B. If Conor doesn't press either button, the machine will randomly kill one and release the other. If Conor presses button A and not B then Conor rescues Bill and omits to prevent Alan's death. If Conor presses button B and not A then Conor rescues Alan and omits to prevent Bill's death. What if Conor omits to press either button? Then Conor omits to prevent Bill's death, and Conor omits to prevent Alan's death. It does not follow that Conor omits to prevent (Bill's death and Alan's death); for Conor is not in a position to save them both. *A fortiori* it doesn't follow that Conor *causes* (Bill's death and Alan's death), because that result doesn't even happen.

You might respond to these examples that the bank clerk does not cause (E1 *and* E2) by omitting to prevent E1 and omitting to prevent E2. Rather, he causes (E1 *or* E2) by omitting to prevent E1 and omitting to prevent E2. Similarly, by omitting to press either button Conor does not cause (Bill's death *and* Alan's death). Rather Conor causes (Bill's death *or* Alan's death). But if you go this way, then "X omits to prevent E1, and X omits to prevent E2" does not entail "X causes (E1 and E2)". So this response does not allow omissions to be additive, as genuine causes must be. So it concedes that omissions are not genuine causes; which was the main point at issue.

Alternatively, you might say that what the bank clerk causes is whichever of E1 and E2 *does* happen, and similarly that what Conor causes is whichever death *does* happen. More generally, you might insist on restricting talk of the results of omissions to *existent* events. But after all, the omissions are real enough, and what the result of each omission is supposed to be is clear enough. Moreover, it's certain that our responsibility for omissions sometimes includes omissions which *don't* have existent results: you can be responsible for omitting to take the necessary precautions against famine in Ohio. So it isn't clear why anyone would be entitled to restrict talk of the results of omissions to existent events. (Nor is it clear why anyone would be entitled to another restriction, which restricts the scope of the argument to *intended* omissions: after all, unintended omissions are no less real than intended ones.)

Can you evade the present argument by looking for cases where there isn't a problem about realising all the results of all my omissions together? No, because there *are* no such cases: at any node of choice there are indefinitely many alternatives that I omit to take, indefinitely many of which are omissions of the same type as my omission to prevent famine in Ohio. Additivity fails for omissions *everywhere*.

Can you escape by denying the general truth of the principle that genuine causes have to be additive? Not if, as I suspect, this principle is true *a priori*. I certainly can't think of any counter-examples to it. This, for example, is no counter-example: I have a hand grenade just powerful enough to kill Franz, and just powerful enough to kill Hermann; I chuck it into Franz' and Hermann's dug-out; one of them dies. Here it may be unknown to me which result (which death) I cause; and I certainly do something which is sufficient to cause either death; but that doesn't mean that I do something that is sufficient to cause both; *ergo*, this isn't a case where "X causes A" and "X causes B" implies "X causes A *or* B" instead of "X causes A *and* B", because it isn't a case where "X causes A" and "X causes B" are both true. If there are indeed no counter-



examples to the additivity principle, then the only get-out left is to say that the principle is *in general* true, but that it mysteriously breaks down in the case of omissions; and that just seems like special pleading.

Here is another possible response: “Your examples involve a trick—call it over-specification. What the bank clerk fails to prevent, and so causes, is not the theft of two or more logically incompatible quantities of money from my account: it is the theft of *any* amount of money from my account. Likewise what Conor fails to prevent, and so causes, is not two physically incompatible deaths: it is *one or the other* of two physically incompatible deaths.” The trouble with this response is that “the theft of *any* amount of money from my account” is not a fully concrete effect, and neither is “one or the other of two deaths”; and the principle that genuine causes must have fully concrete effects seems no less plausible than the additivity principle. At any rate, I can think of no cases where a successful *action* fails to lead to any fully concrete effect, as successful omissions often do when they are analysed in the way just suggested. (Note: I am not denying that genuine causes have incompletely concrete effects; what I am denying is that a genuine cause can ever have *nothing but* incompletely concrete effects.) If there are indeed no counter-examples to the concreteness principle—the principle that genuine causes must have at least some fully concrete effects—then the only get-out left is to say that the principle is *in general* true, but that it mysteriously breaks down in the case of omissions; and that seems like another piece of special pleading.

The results of real causes have to be additive; the results of omissions are not additive; therefore, omissions are not real causes. That is the basic form of the bold argument for the claim that omissions have no causal powers. I believe the bold argument is now vindicated. Nonetheless, there is more to be said; for many of us feel a strong intuition that omissions *must* have causal powers. So next, I shall quickly dismantle two of the engines that drive this intuition.

The first reason why people often think that omissions must have causal powers is because they think that agents must be causally responsible for their omissions. They think that agents must be *causally* responsible for their omissions, because they think that agents must be *morally* responsible for their omissions. But this is a *non sequitur*. We can hold agents *morally* responsible for (some of their) omissions without having to say that they are *causally* responsible for them. According to AOD, although moral and causal responsibility *normally* go together, there are cases where they come apart. AOD’s claim is not that we are *always* morally responsible for our actions and *never* morally responsible for our omissions. It is that we are always morally responsible for our actions *unless there is a good reason why not*, and that we are never morally responsible for our omissions *unless there is a good reason why*. The point about the agent who is morally responsible for an omission is not what is usually true—that he is morally responsible because he is causally responsible. It is rather that he is morally responsible precisely because he is *not* causally responsible, when he should have been.

The second reason why people think that omissions must have causal powers is because they think that counterfactuals typically express causal relations, and they notice counterfactuals like these: “If you had given £10 to famine relief in Africa, ten deaths would have been prevented”. They therefore infer that your omission to give £10 *caused* ten deaths; and so, more

generally, that a counterfactual understanding of causation commits us to the view that omissions are causes.

The argument presented here for the claim that omissions have no causal powers does not depend on any particular understanding of causation. The argument will go through for *any* theory of causation that entails both the additivity principle and the concreteness principle. So far as I know there is nothing in counterfactual understandings of causation to block these entailments. So the argument should go through for counterfactual views too. And so it proves.

To see this, compare two pairs of counterfactuals:

- (A1) If you had not sent poisoned food to Africa Aid, ten deaths would not have occurred;
- (A2) If you had not sent poisoned food to Afghan Aid, ten deaths would not have occurred;
- (B1) If you had given ,10 to Africa Aid, ten deaths would not have occurred;
- (B2) If you had given ,10 to Afghan Aid, ten deaths would not have occurred.

If you send poisoned food to ten people in Africa and send poisoned food to ten people in Afghanistan, it follows by the principle of additivity that you cause twenty deaths. But if you omit to send £10 famine relief to Africa Aid and omit to send £10 famine relief to Afghan Aid, it does *not* follow that you omit to prevent twenty deaths. What your omission is, depends on what it is an omission *of*: that is, on what you could have done if you had *not* omitted. And maybe you only have £10: that is, maybe it is possible for you to save ten African lives *or* ten Afghan lives, but not possible for you to save ten African lives *and* ten Afghan lives. In this case, the most that you cause, or fail to prevent, is ten African deaths *or* ten Afghan deaths; in other words, additivity breaks down again.

“*At most* you cause ten African deaths *or* ten Afghan deaths”: actually, a little more thought about the case shows that you do not even cause that. The claim that “you cause ten African deaths *or* ten Afghan deaths” can be read two ways. On the first reading, it is a disjunction about what you cause: it says “either you cause ten African deaths, *or* you cause ten Afghan deaths”. Here we cannot affirm either disjunct; for if we do, there will be equally good reason to affirm the other disjunct instead. The only way out of this *impasse* is to deny the claim as read this way. But the other reading is no better. On that reading, the claim says that you cause a *disjunctive event*, the name of which is “either ten African deaths or ten Afghan deaths”. Even if there any disjunctive events, which seems doubtful, it seems certain that disjunctive events are not fully concrete. So on this reading the claim falls foul of the principle of concreteness, which says that there *is* no genuine causing of what is not fully concrete. (Except in virtue of the causing of something else that *is* fully concrete; but nothing else is on offer here.) So on either reading, the claim comes out false.

The moral is that the truth of counterfactuals like (B1) and (B2) does not show that the counterfactual theorist of causation must treat omissions as causes, any more than it shows that anyone else should treat omissions as causes. Rather, what it shows is that the counterfactual theorist of causation had better not assume, without further analysis, that a true causal statement corresponds to every true counterfactual. (Of course, the counterfactual theorist of causation had

better not assume that *anyway*: witness [”If seven was an even number, it would not be a prime”].) The counterfactual theorist of causation will face insuperable difficulties if he tries to treat omissions as causes. Witness, for example, the struggles of David Lewis in “Causation”, Appendix D (in Lewis, *Philosophical Papers* Vol.II (Oxford: OUP, 1986)), who, despite three attempts, can find no satisfactory way of squaring his usual account of causation with his unargued conviction that “it is not to be denied that there is causation by omission” (p.191). I believe Lewis is wrong about this. The right route for him to take in developing his counterfactual theory was not for him to struggle to accommodate omissions as causes. It was to say what I think everyone should say, no matter what their theory of causation: that omissions are not causal influencing at all, but *abstainings* from causal influencing.

After all this, do we still feel the urge to say that omissions are causes? Perhaps we do: perhaps because we see a sort of symmetry between action and omission. The idea is that wherever producing Y by acting is in my power, producing not-Y by omitting is equally in my power. Therefore, acting and omitting must be things over which I have an equal amount of causal influence. So actions and omissions must always display an *equal* degree of actionhood.

This is really a proposal to widen the scope of the word “action” so that it includes everything that is under the agent’s control. That proposal should be resisted: action and control are importantly distinct notions. The question that matters to the defence of AOD is not whether I *control* my omissions, but whether my omissions are no less *actions* than any other intentional performances that I produce. The defender of AOD does not in the least wish to deny that Larry is in control of his intentional performance when he lets Tim walk over the cliff. What he wishes to deny is that Larry’s letting Tim walk over the cliff is an *action*; or, at least, is *as much of* an action as Harry’s pushing Jim off the cliff.

“But why should ethics be about *action* rather than about *control*?” Ethics is, or should be, about both. It is only pure act consequentialism that focuses on control and nothing else. Whether that exclusiveness is right or wrong is a very big question. To answer it would be to give a verdict on the core thesis of pure act consequentialism: “Any agent is always morally obliged to bring about the best consequences within her control”. Elsewhere I have presented two separate lines of argument against this thesis: (1) that since there is an incommensurable plurality of goods, “the best consequences” is typically a non-referring expression; (2) that even if “the best consequences” is not a non-referring expression, its reference is unknowable to the less than omniscient, so that act consequentialism provides non-omniscients with neither a decision procedure nor a criterion of rightness<sup>7</sup>. Here my main concern is with a narrower question: whether AOD and PDE are intelligible and plausible. The contradictory of the act consequentialist’s core thesis is the non-consequentialist thesis that the different ways in which consequences (and other things) can be within or outside an agent’s control are themselves morally relevant (or can be). AOD and PDE are not ways of *proving* this non-consequentialist thesis; they are only ways of *articulating* it. So this paper is not a disproof of act

---

7. For (1) see my “The Implications of Incommensurability”, *Philosophy* 2001. For (2) see my “Option Ranges”, *Journal of Applied Philosophy* 2001.

consequentialism, or of consequentialism in general. It is only an attempt to show that the alternatives are a good deal more comprehensible than consequentialists typically allow.

At any rate, it is now clear that Section II's simple strategy for defending AOD and PDE has (at least) a good chance of vindicating AOD. In Section V, I shall ask whether the same simple strategy looks equally promising as a way of vindicating PDE.

## V

As with AOD, the crucial move in the simple strategy is Step 4. Applied to PDE, Step 4 is this thesis:

STEP 4B: *Ceteris paribus*, any performance under a known-but-not-intended description has a lower degree of actionhood than that same performance under a known-and-intended description.

I begin with two clarificatory comments on 4B. These comments also apply to the definition of PDE that I offered in 'II.

First comment: don't confuse "not intended" with "unintentional". No one is going to deny that what I do *unintentionally, by accident*, e.g. tripping over your feet as I arrive late at the seminar and try to find a chair, has a lower degree of actionhood than what I do *intentionally and not by accident*, e.g. kicking your feet to stop you falling asleep during the seminar. PDE and Step 4B are not theses about the distinction between the intentional and the unintentional. They are theses about the distinction between those descriptions under which we intend what we intentionally do, and those descriptions under which we do not intend what we intentionally do. (The idea of *intentionally* doing something under a *not-intended* description may sound self-contradictory. It isn't. Bear with me, and all will become clear.)

Second comment: the form of my statement of 4B corresponds to the form of my statement of PDE; but my statement of PDE is not the standard one. The standard statement of PDE makes it a thesis about foreseeing and intending causal consequences. But, as noted (implicitly) by Bennett, *The Act Itself* p.39, PDE is not just a principle about foresight of the causal consequences of an action; it is also a principle about awareness of the logical consequences of an action. Action-descriptions can extend the *diachronic* scope of what they describe (pulling the trigger vs. shooting the Archduke vs. starting the War): in this case causal consequences are relevant. Action-descriptions can also extend the *synchronic* scope of what they describe (making conversation vs. joking about silly wigs vs. inadvertently insulting my Head of Department): this need have nothing to do with causal consequences. (Corollary: there are *two* "accordion effects"<sup>8</sup>, not just one: there are two dimensions in which any action can be more narrowly or more widely redescribed.)

---

8. The term is Joel Feinberg's: see his "Action and Responsibility" in his collection *Doing and Deserving* (London: Princeton UP 1970).

This much to clarify what Step 4B is, and to clarify PDE, the principle for which 4B is part of the argument. How then might we argue for 4B? The obvious way to do it is to go back to 'III's three tests for actionhood, *causation*, *intention* and *ability to do otherwise*. This time, our proposal will not be that performances under known-but-not-intended descriptions score less well than performances under known-and-intended descriptions on the causation or ability to do otherwise tests. Instead it will be, unsurprisingly, that the crucial difference comes with the intention test. The crucial difference, the one which explains the difference in actionhood, will be that the known-and-intended description of my action *picks out the reason why I do it*. The other descriptions of my action pick out aspects of my action that I am perfectly aware of, and that I accept as applying to my intentional action when I choose to do it; but they do not pick out the intention that actually explains what I do.

So 4B will be a defensible thesis, if it is defensible to think this about at least some of your actions: While different descriptions of your action are all true, yet one description picks out the nature of your action more accurately than the others, because it describes your action *according to the plan which you had in mind in doing it*.<sup>9</sup> We may pick out your action as the execution of the intention that actually guided it; or we may pick it out as fulfilling some other purpose which you were aware the action achieved, and so *could* have had in doing it, but in fact did not.<sup>10</sup> 4B is simply the claim that the action picked out the first way scores higher on the intention test for actionhood than the action picked out the second way.

It seems to me *obviously* right to see a higher degree of actionhood in an action described as I plan it than in the same action described in any other way (even if I know that this other description is true of the action). Intention is about the fulfilment of *particular* plans and purposes in the world. I switch the light on *to find the catfood*; I do not switch the light on to find the catfood and add to peak-time use of the electricity supply and become the thirty-first person in Scotland to alter a light-setting in the last ten minutes and do the easiest action I have done all day... (etc.): not even if I know that these other descriptions are all true of what I do when I switch the light on.

So the distinction that PDE makes is reasonable as a distinction about *actionhood*. By Step 3, it ought to follow that it is also reasonable as a distinction about *responsibility*. *Ceteris paribus*, it ought to be true that the agent is more responsible for her intentional action *under that description under which she does it as part of her plan*, than she is responsible for her intentional action *under any other description, including any description under which she does it intentionally*.

Is this true as well? Well, at any rate, it seems to me, the opponents of PDE have given us

---

9. A recent *locus classicus* for the planning conception of intentions is the work of Michael Bratman: *Intentions, Plans, and Practical Reason* (Cambridge, Mass.: Harvard UP 1987).

10. This contrast is often called the contrast between the *direct* and the *indirect* intention, but this strikes me as a misnomer. What I do intentionally, but not as part of my intended plan, I do not intend *at all*, not even indirectly.

no good reason to think that it is not true. Sometimes they accuse PDE of implying that the agent is not responsible at all for what she does intentionally, but not in pursuit of the plan that guides her action. (For example, see the Singer quotation below.) But this is a misunderstanding of PDE. If my plan is to bomb a weapons factory which the enemy has deliberately located in a heavily populated part of a city, PDE agrees that I am at least partly responsible for the civilian deaths that I know will follow if my raid goes according to plan. (At least partly responsible”: the enemy is responsible too, because of where they put the weapons factory.)

Again, despite what some critics of PDE seem to think (again, see the Singer quotation below), it is not at issue that the agent can rightly be blamed for doing something intentionally that she does not intend. PDE agrees that if the fulfilment of the bombing plan will cause huge civilian losses and destroy only a handful of hand grenades, then the agent who carries out the bombing can rightly be held culpable even though the huge civilian losses were not what she *planned* to produce. (This is PDE’s “requirement of proportionality”.)

What *is* at issue in the debate about PDE is this question: “Is there any difference in responsibility between the bomber who targets the weapons factory, knowing that she will also kill civilians, and the bomber who targets the civilians *as well as* the weapons factory, or indeed just the civilians?” Intuitively, the correct answer to this question is the answer that PDE delivers: Yes.

As I said at the end of III, my discussion of PDE simply aims to show how to defend this intuition. Part of the defence of that intuition can be an appeal to the simple strategy (II) to show that PDE is not a free-standing philosophical curiosity, but a thesis with a deep justification in the theories of action and responsibility. Another way to defend that intuition is to point to the failure of the arguments given by the opponents of PDE. I will now complete my discussion of PDE by looking at four such opponents: James Rachels, Peter Singer, David Bostock, and Jonathan Bennett. In each case, as we shall see, the argument given against PDE rests on serious mistakes.

The mistakes are most obvious in the case of **James Rachels**, who defends the counter-intuitive conclusion about the tactical bomber (*The End of Life* (Oxford: OUP, 1986) p.92):

[Consider] a bombing raid that would destroy both an armaments factory and some surrounding homes. The commander may order such a raid with different intentions... [But] the same civilians are killed, and the same factory is destroyed, *regardless of intention*. If the act is wrong with one intention, how can it be right with another? It is hard to see how the transformation from wrong to right can be made simply by purifying the intention.

I assume Rachels must be arguing like this:

- (P) In respect of their outcomes, the two commanders’ actions are morally indiscernible; therefore
- (C) The two commanders’ actions are morally indiscernible.

As it stands, this argument is a blatant *non sequitur*. To be valid, it needs the following extra premiss:

(P2) The only respect that has any effect on moral discernibility/ indiscernibility is outcome.

But (P2) is exactly what defenders of PDE deny, since they think other things matter too; for instance, they think that intention matters as well as outcome. So Rachels' argument is blatantly invalid without (P2), and flagrantly question-begging with (P2).

A subsidiary point: Rachels' talk of purifying the intention is (as Rachels explains) an allusion to Pascal's satire on abuses of PDE, in the *Provincial Letters*. Pascal said that these abuses took it that all you had to do to avoid blame was learn to be careful where you pointed your internal acts of intending. PDE itself should not be equated with these abuses of PDE. Intending is not an internal act, a little speech that you make to yourself before acting; it is for instance possible to be *mistaken* in your beliefs about what you intend. Intending is rather a *set of the whole agent*. (Thinking about Kavka's Toxin Puzzle (Kavka, *Analysis* 1983) may help to clarify this point. What would the millionaire have to *know* about me to know that at midnight I intended to drink the toxin? Whatever the right answer to this might be, it definitely seems to be the wrong answer to suggest that all the millionaire needs, to find out what I intended, is a sufficiently sophisticated brain-scanner.)

**Peter Singer's** discussion of PDE makes the same mistakes as Rachels. But he also throws in, at no extra cost, some further mistakes of his own (*Practical Ethics* (Cambridge: CUP, 1979), pp.209-210):

[Proponents of PDE] argue that as long as the *directly intended* effect is the beneficial one that does not violate an absolute moral rule, the action is permissible.<sup>11</sup> Though we foresee that our action (or omission) will result in the death of the patient, this is merely an unwanted side-effect. But the distinction between directly intended effect and side-effect is a contrived one. We cannot avoid responsibility simply by directing our intention to one effect rather than another. If we foresee both effects, we must take responsibility for the foreseen effects of what we do... for example, a chemical company might want to [dump some toxic waste in a river]. Would we allow... the company to say that all they directly intended was to improve the efficiency of the factory? In rejecting [such an excuse, the defenders of PDE] would have to rely upon a judgement that the cost is disproportionate to the gains. Here a consequentialist judgement lurks behind [PDE].

---

11. This statement of PDE is an obvious *misstatement*. Proponents of PDE do not argue that all's well just so long as we do not directly intend a harmful consequence; they require proportion between the beneficial and harmful consequences. Singer first ignores the part of PDE about proportionality, and then pretends that it is consequentialist, and so (?) inconsistent with the rest of PDE.

For a start, PDE does not need the distinction between intended effect and side-effect that Singer finds “contrived”. (My statement of PDE in II managed without it.) PDE *does* need the distinction between what is in an agent’s intentional plan either as means or end, and what is not in the agent’s intentional plan at all, even though he is aware of it. While this distinction is not contrived but a natural part of our thinking about agency, it is sometimes a hard one to draw. However, I am not arguing that PDE is always easy to apply; I am arguing that it is generally intelligible. Moreover, this distinction does not seem to be what Singer is attacking, though no doubt it is what he ought to be attacking. (So what *is* he attacking? Heaven knows.)

Again, as I have already pointed out, Singer’s claim that “we must take responsibility for the foreseen effects of what we do” is simply beside the point, because PDE does not deny this. PDE only claims that we are normally *less* responsible for the action-descriptions under which we don’t intend what we do.<sup>12</sup> This seems to be the right thing to claim. Heavy as the chemical company’s responsibility is for poisoning the river by carelessness, their responsibility would be even heavier if poisoning the river was *part of their plan*. (Here Singer apparently muddies the issue by implying that the chemical company *do* do it on purpose, and are disingenuous in their appeal to PDE. This seems to be the force of his rhetorical question “Would we *allow* the company to say...?”) But this is a distraction from the main question about PDE and responsibility. The main question is “Should we say anything different if the chemical company *are* sincere in their claim not to have intended this foreseen outcome?” And the answer to this question is Yes: as I say, there is plainly a difference between responsibility for carelessness and responsibility for deliberate environmental vandalism.)

---

12. Antony Duff has pointed out to me that there is evidence of a lack of parallel here between AOD and PDE. Since my omissions can be entirely unintentional, they can be things for which my responsibility really is zero. The parallel point does not seem to hold for PDE. What I do intentionally, but under a not-intended description, I do *intentionally*; and I cannot have zero responsibility for anything that I do intentionally. So then, Duff suggests, PDE begins to look like a distinction about *culpability* rather than about *responsibility*. (Cp. Sophie Botros, “An error about the doctrine of double effect”, *Philosophy* 1999, who argues that the main point of PDE is not “prohibitive” but “justificatory”.)

I do not agree with Duff about all of this. I accept that PDE never leaves room for zero responsibility (so that the bracketed clause beginning “and possibly...” in my statement of PDE becomes otiose). But I still think that PDE is a distinction about responsibility, and that it is right to say that we are more responsible for our actions described as we plan them than for the same actions otherwise described.

Even if Duff was right, we would still have to ask *why* one is (sometimes) not culpable, or less culpable, for what one does intentionally but under a not-intended description. To answer this, I think, we need to be able to point to some distinction about the conditions of actionhood. So even if Duff was right, then the main difference it would make to my argument would be that strictly speaking my Step 3 should say that “Degrees of responsibility *and/ or* culpability are proportionate to degrees of actionhood”. There would then be some intricate work to do to cash out this “and/ or”; but I shall not attempt this here.



Finally, Singer makes a crude mistake in saying that “the judgement that the cost is disproportionate to the gains is a covert consequentialist judgement”.<sup>13</sup> No: the judgement is (at most) *consequential* not *consequentialist*. Consequentialism is the view that consequences are the *only* thing that count morally. So to judge *some* cases by their consequences is hardly to be a consequentialist, especially not when the reason why you do so is precisely because you think that these cases, unlike others, are not settled by other factors, e.g. by absolute rules. (In fact, it is hard to see how a consequentialist could have got into the position where he *needed* to make such a judgement as the one Singer calls consequentialist.)

**David Bostock’s** reasons for rejecting PDE are rather less crude (*Aristotle’s Ethics* (Oxford: OUP, 2000), p.104):

Suppose I intend to assassinate the Queen... by placing a bomb in the saddle of [her] horse. Of course I foresee that if this plan succeeds then the horse will die too. Do I count as *intending* to kill the horse? Some might say no, on the ground that my *goal* was to kill the Queen and not the horse, so that that and only that was what I intended. But, since I did foresee the horse’s death, they might concede that I did kill the horse intentionally. On this view its death was intentional but not intended. But this distinction is irrelevant...To see this clearly, just reverse the example, and suppose that my goal was to kill the horse, and it was merely incidental to this that the Queen would die too. Obviously I would be held responsible for the Queen’s death... The distinction between the goal itself, what is a means to it, and what is merely a side-effect, seems to me irrelevant to responsibility and blame; what matters is just whether these effects of my action were foreseen by me. If they were, then I shall count them as both intentional and intended.

A nice story; but it contains at least two *non sequiturs*. The first *non sequitur* is to move from a verdict about this *single* pair of stories to a generalisation about *every* story. Bostock does not even try to justify this wild leap. The second *non sequitur* is as follows. In Bostock’s reversed example, I would indeed be held responsible for the Queen’s death. That does not imply that there is no difference in responsibility between the two stories. Even if I am *exactly equally* responsible in the two stories, the reason *why* I am responsible can still be different. And this is just what PDE says. In the first story, PDE says that I am responsible for the Queen’s death *because it is what I plan to bring about*. In the second story, PDE says that I am responsible for the Queen’s death *because it is a side-effect of what I plan to bring about*. Here the point is that this side-effect is not justifiable by the requirement of proportionality, even if what I plan to bring about (the horse’s death) is itself morally justifiable, which seems unlikely to be true.

Finally, **Jonathan Bennett’s** discussion returns us to the case of the bombers. Bennett

---

13. Jonathan Glover, *Causing Death and Saving Lives* (London: Penguin, 1977), p.90, makes the same crude mistake about the idea of proportionality: “Having allowed this concession to utilitarian calculation, where is a line to be drawn, and why?”. For more about why this is a mistake, and why it is a crude one, see David Oderberg, *Moral Theory*, pp.97 ff.

agrees with PDE that there is an important distinction between doing something *intentionally* and doing it *as an intended part of the plan which you are acting on* (*The Act Itself* (Cambridge: CUP, 1995), p.201): “What someone intends to achieve by  $\phi$ ing depends on which of his beliefs about his  $\phi$ ing’s consequences explain his behaviour”. His acknowledgement of this point leads Bennett to suggest that the way to tell whether someone intends something, is to ask them questions like this one (asked of a commander who ordered a bombing raid that killed lots of civilians): “If you had thought that your raid would not cause civilian deaths, would you have called it off?”.<sup>14</sup>

---

14. For another critique of PDE’s deployment of this sort of counterfactual question which is in some ways reminiscent of Bennett’s, see Helga Kuhse, *The Sanctity of Life Doctrine in Medicine* (Oxford: OUP 1987), at p.134, with J.L.A.Garcia’s comment (Garcia, “Intentions in Medical Ethics”, in Oderberg and Laing, *Human Lives* (London: Macmillan 1997), p.168): “[Kuhse] has misunderstood the point of the test. The relevant issue is what is part of the agent’s actual plan. What the proponent of intention-sensitive medical ethics wants to show is that death is not part of the agent’s plan in permissible instances of letting die. The test helps to show this because changing her expectation of death tends to change the mercy-killer’s behaviour, because it forces a change in her plan to end her patient’s pain; while it tends not change the behaviour of the physicians [who are not prepared to intentionally kill the patient] because it forces no change in their plans, since their patients’ deaths were not planned in the first place.”

This question of Bennett's discriminates the tactical bomber from the terror bomber and the punitive bomber in just the way that PDE wants. For the tactical bomber was only after the munitions factory in the middle of the civilians, so his answer is No. Whereas the terror bomber intended to terrorise the enemy (his end) by killing lots of civilians (his means), so his answer is Yes. And the punitive bomber intended to kill lots of civilians *as a punishment*, so killing them was his end, and his answer is Yes too.

"What more could we want?" asks Bennett (p.216). Quite so. However, Bennett does not mean this as a rhetorical question:

Well, we could want more precision. What is the possible state of himself that [the bomber] is asked to entertain? There are three ways we might go on:

(a) [the fact that no civilians are killed is the only difference from how things actually are:] In that case, we are leaving the tactical man with his belief that his raid will destroy the factory, and the terror man with his belief that his raid will lower morale. [So only the punitive bomber will answer Yes.] This question, then, does not yield the grouping that the [defender of PDE] wants...

Bennett goes on to put up two other more precise alternatives, (b) and (c). He then argues that only (c) gets the tactical bomber differentiated from the other two bombers in the way that PDE requires, and that (c) has no moral significance. So, he concludes, PDE has no moral significance.

We can ignore (b) and (c); for the crucial move comes under (a). The crucial move is Bennett's claim that the terror bomber will answer No to the question as made precise by (a). This claim is wrong unless the terror bomber changes his intention. So long as the terror bomber intends *to lower morale by killing civilians*, he is bound to answer Yes to the question specified by (a). For he cannot *lower morale by killing civilians* without *killing civilians*; it is part of his plan to kill civilians. Of course, he might drop this means to his end; his end being only to *lower morale*. But if he does drop that means, then he abandons his plan *to lower morale by killing civilians*. And if he abandons this plan, then he is no longer the terror bomber of the hypothesis.

So the terror bomber will *not*, so long as he is still a terror bomber, answer No to the question "If you had thought that your raid would not cause civilian deaths, would you have called it off?", made more specific in the way that (a) indicates. He will answer Yes. So that question groups the terror bomber as PDE wants: with the punitive bomber.

Is it necessary to add an argument that the question as spelled out by (a) is not a morally insignificant question? There is no hint that Bennett thinks we need such an argument. His position seems rather to be that (a) would be fine for PDE if only it could be made to work. But I have shown that it can be made to work.

## VI

I conclude that PDE, developed and defended according to my simple strategy, survives all the objections put up against it by the four critics whom I have considered (Rachels, Singer, Bostock, and Bennett). Since the standard of these objections to PDE is, in my view, fairly representative of the anti-PDE literature in general, it follows that the simple strategy makes the prospects for PDE pretty favourable. The verdict of 'IV was that the simple strategy was also a promising one for the defence of AOD. So the overall conclusion is that there is good reason to be optimistic about the prospects for defending both PDE and AOD by way of the simple strategy.

And why, to close with a more general question, would anyone want to defend AOD and PDE? No doubt there are lots of possible motivations; but let me close by revealing my own.<sup>15</sup>

My motivation comes from the fact that I hold both of two views. First, I am a non-cardinal pluralist. This means that I hold (like Isaiah Berlin, Charles Taylor, Martha Nussbaum, Aristotle, and Aquinas) that there is a irreducible plurality of different goods at least some of which are of the same basic importance, but which stand in no cardinal-numerical commensurability relations to each other. Second, I also hold (like pretty well everyone) that we ought to promote the good.

These two views seem to be in tension. For what can promoting the good mean for a non-cardinal pluralist, who recognises a number of *different* goods (call them G1-G3), all of which he takes to be equally important, and for which there is no way of assigning numerical values to different possible trade-offs of those goods?

The injunction cannot mean that the non-cardinal pluralist should promote all of G1-G3 *simultaneously*. Since G1-G3 are different goods, it cannot be possible to promote them all at once in every possible (or even likely) situation. Nor can the injunction mean that the non-cardinal pluralist should promote *some one* of G1-G3, say G1, *at all costs*. For the costs in question are (or will often be) costs *to G2 and G3*. And by definition, the non-cardinal pluralist holds that G1 and G2 are just as much goods, and are just as important, as G1 is. So he has no more reason to select G1 rather than G2 or G3 as the single good to promote at all costs. Given that he thinks all three of G1-G3 *are goods*, it is arbitrary and irrational of him to select any of the three goods (or any two out of the three goods, come to that) for promotion while ignoring the others.

---

15. I first presented the argument that follows in my *Understanding Human Goods* (Edinburgh: Edinburgh UP, 1998), Ch.3. See also my "The Implications of Incommensurability", *Philosophy* 2001.

Therefore, what the injunction to promote the good must mean for the non-cardinal pluralist is this: that he is free to promote *any* of G1-G3, *but not at just any cost to the unpromoted goods*. And so the non-cardinal pluralist is rationally bound to recognise both of two sorts of limits: limits on the *demands* that any of the goods can place on him; and limits on the *costs* that he should accept in losses of any one of the goods he recognises. The limitation of the demands of each goods means that the non-cardinal pluralist has a kind of freedom to choose between the goods. His negative responsibility to ensure that each of the goods is promoted is limited: in going beyond the minimum that he is bound to do in respect of any good, he acts *supererogatorily*. The limitation of the costs to each good means that the non-cardinal pluralist has to recognise constraints. His freedom to treat any of the goods any way he likes while pursuing other goods is also limited: in falling short of the minimum that he is bound to do in respect of any good, he breaks a moral rule that that good imposes on him.

It then becomes an interesting question for the non-cardinal pluralist exactly where these limits are. The non-cardinal pluralist cannot explain this by contriving a quantitative system of weighing and measuring costs to goods against each other. For that would be tantamount to denying the central claim of non-cardinal pluralism, that there *is* no quantitative system for commensurating goods.

Does this mean that the non-cardinal pluralist can give *no* explanation of how to identify the limits of how we may treat any good? Of course not: the non-cardinal pluralist can identify those limits not by a quantitative but by a qualitative method (so that in fact the metaphor of “costs” falls away). He can say that the limits are set by the answer to the question whether our actions *honour* each of the goods involved, or show *respect* for each of the goods involved. The answer to *that* question is going to depend on how clear it is that our actions are directed *against* those goods, in the sense that, for instance, an act of murder is directed against the good of life. The way to decide whether an action is more or less, or not at all, “directed against” some good will be to examine how the conditions of actionhood apply to this action. And if the argument of the present paper is correct, part of what this will involve is looking at whether or not the action in question falls within the scope of PDE, or AOD, or both.<sup>16</sup>

*Timothy Chappell*  
*Department of Philosophy*  
*The Open University*  
[t.chappell@open.ac.uk](mailto:t.chappell@open.ac.uk)

---

16. My thanks to José Bermúdez, Nick Denyer, Antony Duff, Basil O’Neill, Joseph Shaw, and Chris Tollefsen for written comments, and to audiences in Stirling, Glasgow and Leeds for helpful discussion.