

Choosing your poison and the time of a killing

Auke J. K. Pols

Published online: 29 June 2012

© The Author(s) 2012. This article is published with open access at Springerlink.com

Abstract The problem of the time of a killing is often cited as providing grounds for rejecting the action identification thesis favoured by Anscombe and Davidson. In this paper I make three claims. First, I claim that this problem is a threat to the action identification thesis because of two assumptions the thesis makes: since the thesis takes actions to be a kind of doings, it has to assume that agents' doings last as long as their actions and vice versa. Second, I claim that not making both of these assumptions necessarily leads to another problem, the problem of the acting dead. This means that any theory of action has to choose its poison and face either one of these unresolved problems. Third, I claim that the solution to the problem of the time of a killing can be found by heeding linguistic arguments that 'kill' cannot mean 'cause to die,' as is commonly assumed, but instead has to have a more complex meaning. I discuss an alternative, more complex proposal and show how it allows us to keep the action identification thesis, fits colloquial usage of 'kill' and deals with the problem of the time of a killing.

Keywords Action theory · Time of a killing · Action individuation · Anscombe, G.E.M. · Davidson, D

1 Introduction

There is a notorious argument against the action identification thesis, or AIT, that is called the problem of the time of a killing (Thomson 1971). The thesis says, roughly, that when an agent does something by doing something else, the agent only

A. J. K. Pols (✉)
School of Innovation Sciences, Eindhoven University of Technology, IPO 1.09,
P.O. Box 513, 5600 MB Eindhoven, The Netherlands
e-mail: A.J.K.Pols@tue.nl

performs one action, but this action can be described in different ways.¹ The problem is that the thesis is confronted with the following scenario: Suppose that Smith shoots Jones on Saturday; Jones collapses, bleeding, and finally dies on Sunday. The question we can now ask, is: ‘When did Smith kill Jones?’ We can ask the same question for the place of a killing if Jones manages to leave the place of the shooting before collapsing: though I will focus on establishing the time of a killing, it will turn out that one can establish where the killing occurred in a similar way.

According to Thomson and various authors after her, most notably Mackie (1997), Mossel (2001), Weintraub (2003) and Sandis (2006), those who accept AIT are committed to saying that since Smith has killed Jones by shooting him, the shooting is the killing. If that were the case, the killing would have happened on Saturday, a day before Jones died, which would be absurd.

In this paper I make three claims. The first claim widens the scope of the problem: it is that establishing the time of a killing is a problem for AIT because the thesis sees actions as a kind of doings, as things that an agent does (Davidson 1980, essay 3). This means that it has to make two assumptions: that agents’ actions last as long as their doings, and vice versa. The second claim is a challenge to those theories that do not make both assumptions: it is that those theories have to deal with another problem that I will call the problem of the acting dead. This means that any theory of action has to choose its poison and confront either the one problem or the other. The third claim, finally, is that the problem of the time of a killing can be overcome by drawing on linguistic research into what ‘kill’ means, and why its usual interpretation as ‘cause to die’ cannot be correct. Though I will thus accept a proposed revision of what ‘kill’ means, my aim in this paper is to deal with the problem of the time of a killing without accepting a definition that goes against colloquial usage of the word ‘kill.’ This includes definitions that allow that there can be killings before dyings, like those endorsed by Davidson (1980, essay 8), Lombard (1978) and Heinaman (1983).

The paper is structured as follows. In section two, I explicate both AIT and the problem of the time of a killing, examine some earlier attempts at dealing with it, and argue that the problem hinges on the assumptions that agents’ doings last as long as their actions, and vice versa, that their actions last as long as their doings. In section three, I show that any theory that does not make these two assumptions necessarily falls prey to the problem of the acting dead. I investigate several proposals to deal with the problem by authors who do not make both assumptions, particularly Mackie (1997) and Mossel (2001), and argue that their proposals do not adequately solve the problem. In section four, I return to the problem of the time of a killing and show how drawing on linguistic research into what ‘kill’ means can help us deal with that problem. I follow Fodor (1970), Wierzbicka (1975) and Morreal (1976) in arguing that the meaning of ‘kill’ cannot be, as is commonly assumed, ‘cause to die,’ due to differences in linguistic characteristics and usage of both phrases. Instead, the meaning of ‘kill’ is rather more complex. It has already been observed by Thomson (1971), Mossel (2001) and Sandis (2006) that the problem of the time of a killing

¹ The thesis goes by various names, including the ‘Action Sequence Identity Thesis’ (Mackie 1997), the ‘Anscombe thesis’ (Bennett 1994; Schnieder 2009), and the Davidson/Anscombe thesis about act identification (Wilson 2009).

vanishes if we do not try to locate killings in time very precisely.² Explicating the meaning of ‘kill’ can help us see why this is looking for the solution in the right direction, rather than evading the problem: this paper suggests a definition that only allows us to describe an action as a killing at time t if t includes both the action and its consequence, the death. It will be shown that this solution applies to the whole class of causative verbs, rather than just to ‘to kill,’ though the discernible effects of this are limited.

2 The problem of the time of a killing

In this section I first present the problem of the time of a killing, and then examine each of its conditions in turn. The scenario is as follows:

- (a) Smith shoots Jones on Saturday. Jones collapses, bleeding.
- (b) Jones dies from blood loss on Sunday.

The problem is basically a *reductio ad absurdum*, though the result is a colloquial, rather than a logical, absurdity:

- (1) Smith shot Jones on Saturday.
- (2) According to AIT, if an agent does something by doing something else, the agent has only performed one action. Since Smith kills Jones by shooting him, the shooting is the killing.
- (3) Therefore, Smith killed Jones on Saturday.
- (4) Therefore, Jones was killed a day before he died, which is absurd.

With the exception of step (1), which is simply taken from the scenario, all of these steps have been contested. I will examine each of these steps in turn, starting with explicating AIT.

According to Davidson (1980, essay 3), events involving agents can be either doings, things that the agent does, or happenings, things that happen to the agent. Davidson does not offer a clear demarcation between the two, and indeed argues that there are borderline cases, e.g. dozing (see also Thalberg 1967). I will assume here that an agent’s doings are those events that can be described using an active verb with the agent as its subject. Though this definition includes borderline and otherwise disputable cases, it suffices for my present purposes, as I am primarily interested in those doings that are actions.

Davidson further argues that if there is some description under which what we do is intentional, we can be said to perform an action. In case this is something I do directly,³ like blinking or raising my hand, it is clear that only one action has been

² Indeed, any definition of ‘kill’ mentioned in this paper *allows* us to say that Smith killed Jones in the second millennium. The reasons for narrowing down the time frame seem to be pragmatic rather than conceptual.

³ There has been disagreement about whether these ‘basic’ or ‘primitive’ actions that are done directly are bodily movements (Davidson 1980, essay 3), or rather willings (Ginet 1990), tryings (Hornsby 1980) or could include artefact use (Annas 1978). Though I will use bodily movements as examples of basic actions throughout this paper, my point does not depend on any specific conception of basic actions.

performed. Sometimes, however, agents do things by doing other things, like the man from Anscombe's (1957/2000) well-known example who pumps water by moving his arm up and down with his fingers around the pump handle, or Smith who kills Jones by shooting him. Both Anscombe and Davidson argue that in these kinds of cases the agent also only performs one action, for the pumping man does not have to do anything else for his moving of his arm to be his pumping of the water, and Smith does not have to do anything else for his shooting of Jones to be his killing of Jones. Neither Davidson nor Anscombe gives a formal definition of the action identification thesis implied in this claim, but I will adopt one from Bennett 1994 here:

AIT. If someone φ s by ψ -ing, and F is the act which makes it the case that she φ s, and P is the act which makes it the case that she ψ s, then F is P (29).

So if Smith kills Jones by shooting him, Smith performs an act F which makes it the case that he kills, and Smith performs an act P which makes it the case that he shoots, and F is P. According to Davidson, this act is, or can properly be described in terms of, what Smith does directly, not by doing anything else. In this case, this would be Smith's moving of his finger on the trigger (but see footnote 3).

Mackie (1997) expresses concern that AIT would not work because the is-relation between actions is symmetric and reflexive (F is P implies P is F), while the by-relation is asymmetric and irreflexive (Smith kills by shooting; but he does not shoot by killing). The worry is that, since we can swap F and P at will, the same could happen to φ and ψ . However, close examination of AIT can put this worry to rest: φ -ing and ψ -ing are ways to describe an action, while the identity relation claims that there is only one action to which both descriptions apply. The important point here is that whether a certain description applies to an action depends on the particular circumstances in which the action is performed. Killings can be shootings-in-particular-circumstances, namely, if the shooting results in a death. Shootings, on the other hand, cannot be killings-in-particular-circumstances. Or as White (1980) puts it: a killing can be a kind of shooting (namely, a shooting that results in a death), but a shooting cannot be a kind of killing. This ensures that the asymmetry of the by-relation is not threatened, even though the action described is in both cases the same.

As AIT assumes that actions are a particular kind of doings, it has to make two assumptions. The first is that agents have to be doing something for the duration of their actions, or that agents' doings last as long as their actions. The second is that agents cannot be doing something after they have ceased to be acting, or that agents' actions last as long as their doings (after Hornsby (1980, p. 29)). It is this combination of assumptions that makes step (2) of the problem valid and opens up AIT to the problem of the time of a killing. For once Smith has moved his finger, he has done something. According to the assumptions, Smith's action is over now that he has done something, and indeed, Smith needs to do nothing else in order to kill Jones. Yet if Smith has performed the action that can be described as killing Jones, it seems reasonable to infer that he has killed Jones. And if Smith has killed Jones, Jones has been killed, and therefore, Jones is dead—but he isn't yet.

Bennett (1973) has challenged the step from (2) to (3) by arguing that, just like objects, events (including actions) can also acquire new properties. He argues that

Smith's shooting was indeed a killing, though it did not *become* a killing until Jones died on Sunday. Anscombe (1979, p. 227) similarly argues that some descriptions of acts can only become true later, and that careful formulation will solve our problems here. For example, since Smith's action caused the gun to go off, we could say: "That act which (as things turned out) was the killing of [Jones] by [Smith] caused the gun to go off."

Yet there are at least three problems with such an approach. First, Smith's shooting of Jones did not acquire a new property when Jones died. Rather, the shooting had the property of being a cause of death all along. On Saturday, the shooting is the event that causes Jones to die on Sunday, or that makes it the case that Jones dies on Sunday. We might not know this to be true until Jones does indeed die on Sunday, but those facts hold regardless of whether Jones's death has already occurred or not. Assuming that 'kill' is 'cause to die,' the action must have had the property of being a killing all along.

Second, Lombard (1989) has argued that indeed there cannot be a killing *unless* there is a death, but that this does not imply that there cannot be a killing *until* there is a death, which Bennett assumes. The core of Lombard's argument is that, logically, 'there cannot be a killing unless there is a death' means: 'If there is no death, then there cannot be a killing.' This, again, means: 'If there is a time T at which there is no death, then there is a time T^* at which there cannot be a killing.' It does *not* mean: 'If there is no death *now*, there cannot be a killing *now*,' which would be required in order to state that there cannot be a killing *until* there is a death. So from the proposition that there cannot be a killing *unless* there is a death, we cannot derive the proposition that there cannot be a killing *until* there is a death.

Third, Mossel (2001) has argued that Bennett's solution only postpones the problem, not solves it. He notes that as soon as Jones has died on Sunday, it becomes true that Smith's shooting is a killing, and thus that Smith has killed Jones on Saturday, a day before Jones died. In short, Bennett's argument does not seem to be able to prevent the move from (2) to (3).

I do not know of anyone who has challenged the move from (3) to (4), or from 'Smith has killed Jones' to 'Jones was/has been killed.' Indeed, since each step in the problem seems plausible, it has been suggested that we should bite the bullet and accept the intuitively strange consequence of having killings before dyings (e.g. Davidson 1980, essay 8). There are two pragmatic arguments against taking this position, however. First, this is not how we use the phrases 'kill,' 'has been killed,' etc. If we hear that Jones has been killed, we would assume that he is dead, and we would at least be surprised to find out otherwise. Second, White (1980) has observed that we do have a separate way to describe those woundings that have not yet resulted in death, but will eventually do so: as fatal woundings.

As my aim in this paper is to deal with the problem of the time of a killing without accepting a solution that goes against colloquial usage of the word 'kill,' I will not settle for accepting killings before dyings, but rather argue that the problem can be solved by establishing the meaning of 'kill.' Before I do so, however, I will deal with those who would address the problem of the time of a killing by arguing that it implies that we should abandon AIT, or more precisely, either of the two assumptions that I have claimed both support AIT and open it up to the problem.

Specifically, I argue that those who abandon at least one of those two assumptions open themselves up to the problem of the acting dead, and that no satisfying solution to this problem has been put forward thus far.

3 The problem of the acting dead

The two assumptions that lie behind AIT are that agents' actions last as long as their doings, and vice versa. Either one or both of those assumptions can be renounced. The first assumption is renounced by Thomson (1971), who holds that the duration of an action of killing extends from the shooting up to and until the death of the victim, regardless of what the murderer does after the shooting. The second assumption is renounced by Mackie (1997) and Mossel (2001), who hold that agents can be doing something even though they are not active. For example, an agent could legitimately say: "I'm baking bread, doing the laundry and executing a database query," while sitting in a comfortable chair, doing nothing but waiting for the dough to rise, and for the washing machine and the computer to finish.

No matter which assumption is renounced, those who renounce it run into trouble if they claim to be able to handle this scenario adapted from the time of a killing:

- (a) Smith shoots Jones on Saturday. Jones collapses, bleeding.
- (b) Smith dies of a heart attack, moments after firing the shot.
- (c) Jones dies from blood loss on Sunday.

For those who assume that actions can extend beyond agents' doings, the problem of the acting dead goes as follows:

- (1a) Smith has done something, namely, shooting Jones.
- (2a) Smith's action of killing Jones lasted from the shooting until Jones's death (Thomson's assumption).
- (3a) Therefore, Smith's action of killing Jones took place (for the most part) when Smith was already dead. But that would be absurd.

As Heinaman (1983) has already attacked Thomson's position as elaborated in (Thomson 1977) and shown the problems in extending actions beyond doings, I will concentrate here on showing the problem in extending doings beyond actions.

For those who assume that doings can extend beyond agents' actions, the problem is very similar:

- (1b) Smith has performed an action, namely, shooting Jones.
- (2b) Smith was doing something from the shooting until Jones's death (Mackie's/ Mossel's assumption)
- (3b) Therefore, Smith was doing something while (for the most part) Smith was already dead. But that would be absurd.

There seems to be no good reason to object to assumption (1a/b), regardless of one's position. Smith does something, namely, he is shooting; and he performs the action of shooting Jones. Once Jones has been shot, everyone will agree that both

this doing and this action are now over—and that this shooting is in some way or other related to Smith's killing of Jones.

(2a/b) would probably be the easiest step to deny, for the simple reason that not accepting one of AIT's two assumptions does not imply accepting that *all* doings extend beyond an agent's actions or vice versa, only that *some* do. It could then be argued that in this particular case the doing does not extend beyond the agent's action or vice versa. Of course, this would place the burden of proof on the arguer, who would have to come up with a criterion for when these doings or actions do extend beyond the other, and why they do not do so in this case.

Mackie (1997) suggests such a criterion. He presents us with the following case: Brown prints a paper by pressing a combination of keys on his computer, then waits for the printer to finish printing the document. Mackie argues that here we can say that Brown is printing the paper while he waits for the printer, and this is because Brown is at least doing *something*. Mackie suggests that if Brown had died immediately after pressing the combination of keys, we would not say now that he is doing anything at all, not even printing out his paper. So an agent can be doing something without being active, as long as that agent is still around.

It does not seem strange for Brown to say that he is printing the paper while he is waiting for the printer to finish, but this example is problematic for two reasons. First, we cannot obviously adapt Mackie's example to the original time of a killing-scenario. If Smith shoots Jones and then waits for him to die, we would not say that he is killing Jones during that time. This does not necessarily render Mackie's point invalid, but it does at least show that additional criteria are needed. Mackie suggests that the notion of control might be relevant here: Brown is in control of the printing process in a sense that Smith is not in control of Jones's dying. Mossel (2001) adds that Brown has to actively supervise the process of printing, which may itself not require any physical actions. However, it is not clear that these criteria would do the required work: Smith might have saved Jones's life by calling an ambulance, or by providing him with a miracle cure, but that would not sway our intuition that Smith was not actively killing Jones, even if he stood by, supervising Jones's demise. At best, Smith was making sure that Jones would die (as a result of his shooting); but that does not imply that Jones was killing Smith during that time. So positing control of and supervision over the process is not sufficient either as a criterion for extending actions beyond the (direct) doings of the agent.

Second, one might wonder whether Brown's printing of the paper while he is waiting is the same kind of doing as his printing of the paper while he is pressing the combination of keys. Consider the following change in Mackie's example: Suppose that Brown dies immediately after pressing the combination of keys. The printer dutifully does its job. Later, a colleague remarks, shocked: "The last thing he did was print out his paper." This seems a correct observation: the intended result, the paper having been printed, has been achieved. In both cases, Brown does something in order to print out a paper, resulting in him having printed out a paper. Since we can say that Brown is printing the paper in the case where he remains alive, but not in the case where he is dead, even though the result is in both cases the same, this suggests that 'is printing' refers to a plan or project that Brown has, rather than to the kind of doing that is directly involved in actions.

It is true that Brown has a plan or project to have a paper printed in Mackie's example, but not (anymore) in my example. What this suggests, however, is not so much that one of the assumptions behind AIT is false, but instead that the phrase 'is doing' can be used in different ways: to denote actions, or what the agent is doing *now* (Brown is typing; stretching; smiling at his boss), and to denote plans or projects, or what the agent is doing in an extended sense. (Brown is preparing a journal submission; protecting his health; setting himself up for a promotion) (see also Heinaman 1983; Falvey 2000). Or to be more specific: some things that we do intentionally we do either by doing something else, or directly. According to AIT, those doings are actions, and those are also the doings that are referred to in the problem of the acting dead. On the other hand, some things we do by doing something else *and* something else again *and...*, like making falafels or preparing a journal submission, and those doings are plans that are being executed. AIT, as defined, is only about those doings that are actions; otherwise it would simply be false. For example, Smith could have escaped from the police by masquerading as a construction worker, but there was no action of escaping, let alone one that was also an action of masquerading. This may appear to beg the question about what actions are; but the point is rather to state that there are two conceptually different entities, actions and plans, and 'is doing' may refer to the execution of either of them. Showing that people can be engaged in doings in the plan sense is neither a rejection of AIT, nor a proper rejection of step (2) in the problem of the acting dead, which is about doings in the action sense.⁴

Two factors further complicate this picture. The first is that a plan can consist of only one action (and waiting for its results, as with Brown's plan for printing a paper), so that it may not always be possible to distinguish a doing in the action sense from a doing in the plan sense. The second is the status of activities. Activities are doings that stretch out over time, yet resemble actions more than plans, as the agent has to keep doing something for the duration of the activity (Anscombe's man pumping water; Jones's wife killing Smith by strangling him). It might be that because 'kill' can denote both an action and an ongoing activity, 'Smith is killing Jones' sounds strange in the sense that 'Brown is printing a paper' does not. One could require the use of a verb that cannot refer to both an activity and an action as a criterion for whether a doing can extend beyond an action. However, this criterion would clearly not be of the same kind as Mackie's control and Mossel's supervision. Moreover, it would be too vague to be useful: even 'printing' can denote an activity, namely, the operation of a manually operated printing press.

The final strategy for dealing with the problem of the acting dead is to challenge step (3) and argue that this inference is either invalid or not absurd. The latter possibility seems most promising, as there is a sense in which we can do things after our deaths, for example: 'Jones surprised his wife by leaving her his secret falafel recipe.' This suggestion is not compatible with the criteria proposed by Mackie and Mossel, however: Jones was not in control of his wife's being surprised, nor did he

⁴ This observation also takes care of Mackie's point that agents do not always have to be doing something for the duration of their action. He gives the example of slowly killing your flatmate by repeatedly administering small doses of poison over a long period. As I have argued, this is rather a killing in a plan sense, consisting of multiple actions (the poisonings), none of which itself is the killing.

supervise the whole process that led to her being surprised. Furthermore, it seems that these kind of cases can in principle be reduced to cases of an agent doing something by doing something else during his lifetime: by signing his will on Friday, Jones ensured that his wife would inherit his secret falafel recipe after his death (by which, in turn, he would surprise her.) All the things that agents can be said to do after their deaths have to be arranged by them while they are still alive.

In this section I have argued that the problem of the acting dead arises for anyone who does not accept one or both of the assumptions made by AIT, that agents' actions last as long as their doings and vice versa. I have also argued that distinguishing between doings as actions and doings as plans is important both to understand AIT and to understand why Mackie's and Mossel's solutions to the problem are unsatisfactory as they stand. This does not imply that the problem of the acting dead cannot be overcome, or that it would be more difficult to overcome than the problem of the time of a killing. It does imply that those who are quick to reject AIT (or rather, at least one of the two assumptions underlying it) based on the problem of the time of a killing have their own problem to solve if they wish to offer a plausible alternative.

In the final section I revisit the problem of the time of a killing and suggest a possible solution based on what 'kill' means. Thomson (1971) has already expressed doubts about interpreting 'kill' as 'cause to die' *simpliciter*. Similarly there are problems with Davidson's (1980) interpretation of 'kill' as 'do something that causes someone to die,' that for example Lombard (1989) and Weintraub (2003) endorse. I will now examine problems with both interpretations and their consequences for the problem of the time of a killing.

4 Killing and causing to die

At first glance, it seems quite reasonable to say that 'kill' means 'cause to die.'⁵ However, various problems arise on closer scrutiny. For example, this formulation does not make clear what it would mean to say 'Smith caused Jones to die,' or whether event or agent causation would be involved. Furthermore, Fodor (1970) has shown that some transitive verbs may be analysed in this way (e.g. 'melt' = 'cause to melt'), but not 'kill.' He has also mentioned that a sentence of the form 'Smith caused Jones to die on Sunday by shooting him on Saturday' would be well-formed, while 'Smith killed Jones on Sunday by shooting him on Saturday' would not. Morreal (1976) adds that 'cause' is not an action verb and cannot be modified, whereas 'kill' is and can. For example, we can say: 'Smith killed Jones slowly and painfully,' but not: 'Smith caused slowly and painfully Jones to die.' We can say: 'Smith caused Jones to die slowly and painfully,' but here it is the dying that is modified, not the killing.

As an alternative, Davidson (1980, essay 3, pp. 48, 49) has suggested that 'kill' means 'do something that causes someone to die,' where the causation involved is

⁵ It seems that Sandis (2006) endorses this definition, as it would be a natural consequence of his claim that a killing is the causing of a death.

event causation. This takes care of Fodor's and Morreal's objections, yet if we accept Davidson's definition, the problem of the time of a killing arises again: Smith did something on Saturday that caused Jones to die (on Sunday), namely, he shot him. This implies that the killing took place on Saturday, a day before Jones died, which means that we have the absurdity again.

Weintraub (2003) accepts Davidson's definition of 'kill' as 'do something that causes someone to die,' as she argues that there is nothing absurd about a cause (the doing) being distant from its effect (the death), but that our intuitions are doubly mistaken. First, because we conflate two readings of the question: 'When did Smith kill Jones?' namely, 'When did Smith do something that caused Jones to die?' and 'When did the causing of Jones's death occur?' As both cause and effect are involved in the causing, the second reading suggests that both should have a place in the answer. Second, we make a mistake when we think that this causing is an event that can be temporally localised, because it isn't. This makes the reading 'When did the causing of Jones's death occur?' incoherent. What remains is the straightforward first reading, a reading that has a straightforward answer. Thus, the problem of the time of a killing is defused by showing that, while having killings before dyings may be intuitively absurd, this is because our intuitions are mistaken.

I agree with Weintraub that there is nothing paradoxical about causes being distant from their effects, and that the answer to the question 'When did Smith do something that caused Jones to die?' is: 'On Saturday.' Rather than assuming that our intuitions are mistaken and explaining why this is so, this paper aims instead to discover whether it is possible to provide a definition of 'kill' that does not go against its colloquial usage, as is the case with 'do something that causes someone to die.' There are at least two reasons for attempting to find such a definition. First, as White has noted, it seems that there is a difference between 'kill' and 'fatally wound' that would be difficult to explicate if we were to use this definition of 'kill,' at least for all killings that involve woundings. Second, the stakes in this debate seem to be higher than our intuitions about just one word. 'Kill' belongs to the class of causative verbs, verbs that express creation, change or annihilation. Our intuitions seem to balk in many cases where causative verbs are used in relation to effects distant from the actions that cause them. For example, it would seem strange to say that Smith's doctor cured Smith on Monday, when all she did was inject Smith with a slow-working medicine that did not improve Smith's health until Thursday; or to say that Jones steamed the Christmas pudding the moment he put it in the steamer. While Weintraub's analysis of why our intuitions are mistaken might well be extended to cover these cases, it would mean assuming that our intuitions are structurally, rather than incidentally misguided.

Wierzbicka (1975), for one, does not accept the equation of 'kill' with 'cause to die,' or 'do something that causes someone to die.' She reasons as follows: A sentence like 'Smith killed Jones at time *t*' implies that all the events required in order to describe the action as a killing took place simultaneously, that is, at time *t*. Wierzbicka then claims that two events in particular are required, Smith's doing something and Jones's dying. Without either of them there cannot be a killing. Wierzbicka supports her claim by pointing out that many other causative verbs seem to require two simultaneous events to merit application. For example, if Jones

sharpened his cooking knife, Jones did something and his knife was becoming sharper because of that. And if Smith cleaned his gun, then Smith did something and his gun was becoming cleaner because of that. This leads Wierzbicka to define the meaning of a sentence like ‘Smith killed Jones’ as follows: ‘At one time Smith and Jones were in some one place; Smith was doing something and something was becoming in contact with Jones’s body because of that; Jones was dying because of that; after that, Jones was dead because of that’ (cf. *ibid.*, pp. 498, 499). This is rather more complex than ‘cause to die,’ but given that ‘kill’ is used in a rather complex way, Wierzbicka states that this is only to be expected.

Before proceeding, I would like to suggest two amendments to this definition. First, it seems that the requirement of something coming into contact with Jones’s body (a bullet, a knife) is unnecessarily restrictive and should be removed: Smith could also have killed Jones by depriving him of water or oxygen, or making him March until he collapsed from exhaustion. Second, strictly speaking, the mention of ‘one time’ and ‘one place’ is unnecessary, as every cause–effect pair will occur at some (suitably broadly defined) one time and one place. I will leave those phrases be for now, however, in order to highlight this definition’s consequences for establishing the time and place of a killing.

Supposing that we accept Wierzbicka’s amended definition and her analysis of the structure of causative verbs, and also that we assume that it is sensible to inquire as to the time of a doing, however that doing is described, the question ‘When did Smith kill Jones?’ can be then read as: ‘When was the one time in which Smith and Jones were in some one place; Smith was doing something and Jones was dying because of that; after that, Jones was dead because of that?’ The answer is then: ‘during the weekend,’ as two events are required to happen over that time span in order to describe the action as a killing: the action and its consequence. In order to establish the time and place of a killing, these events have to be regarded as happening simultaneously, or rather, at that particular time and in that particular place. Thus, if the shooting took place on Saturday but the death did not, then there was no killing on Saturday because only one of the relevant events took place on that day. This explains why we can, in principle, establish the time of a killing only imprecisely: we have to find a time frame during which both events occurred, rather than establish the time at which a single event occurred. To put it more simply: the time given in the answer should at least include the time of the shooting and the time of the death.⁶

One objection to this solution could be that it will often yield answers that are vague or imprecise. However, if that strikes us as a problem, it might be that we were really looking for the answer to a different question, e.g. ‘When did Smith shoot Jones?’ It should also be noted here that while this solution has consequences for establishing the time of any doing or happening that is described using a causative verb (including ‘cure’ and ‘steam’), in many cases it will have no

⁶ Strictly speaking, according to Wierzbicka’s amended definition three events are involved in the killing, and those should be treated as a trinity: the action, the victim’s dying and the victim’s death. However, any time frame including an action and a resulting death to which the definition is applicable, will necessarily also include the intermediate event of the dying.

discernible effect. This is because for many causative verbs, such as ‘sharpen’ or ‘clean,’ there is no temporal interval between the action and the event that is its consequence.⁷

A second objection might be that, while inquiring as to the time of a single event occurring is fine, inquiring as to the time ‘in which an event occurs after which another event occurs because of that’ sets a dangerous precedent. One could easily come up with insensible examples of such questions, like: ‘When did Jones pet his dog, after which his dog ran around the garden three times in joy because of that?’ In order to avoid this, one could drop the assumption that it is sensible to inquire as to the time of a doing, however described, but this seems rather ad hoc; or one could propose a different underlying structure for causative verbs, thereby revising the definition of ‘kill.’ However, it seems that there are two reasons why we should not worry too much about precedent-setting. First, it is common practice to describe actions and other events in terms of their consequences and it would be surprising if this were not reflected in the linguistic structure of (especially) causative verbs, or if it did not have consequences for establishing the time of a doing described using such a verb. Second, if we assume that it is sensible to inquire as to the time of a doing, however described, we only assume that it is sensible to inquire as to the time of *some* events occurring after which other events occur, namely, when the event in question can be described using a causative verb. We do not have to assume that it is sensible to inquire as to the time of *any* event occurring after which other events occur.

A third objection could be that this seems to be an abandonment of AIT. For if two actions (or events in general) are the same, they should occupy the same stretch of time. Yet Wierzbicka tells us that the shooting occurred on Saturday and the killing occurred during the weekend.

The important point here, as already noted by Anscombe (1957/2000), is that one action is (or can be described as) another only *in* certain circumstances, though it does not *include* those circumstances. To borrow a comparison from White (1980): a married man is not a man *and* his partner, but a man to whom the description ‘married’ applies in the context (whether explicit or implicit) in the sentence. Wierzbicka has spelled out the circumstances in which we can describe an action as ‘to kill,’ and it turns out that the description ‘killing’ can only be applied to the action if the time frame (and place) referred to when using the description, whether directly or indirectly, contains both the action (the shooting) and its consequence

⁷ When the action is temporally very distant from its consequence, the definition may yield results that seem awkward. For example: suppose that Smith’s bullet gets lodged in Jones’s brain. Jones survives and lives on for another ten years, after which the bullet damages vital brain tissue, resulting in Jones’s death. The definition would then commit us to saying that Smith killed Jones during the last decade. I will simply accept this consequence for the following reasons: First, an alternative definition like Davidson’s would yield the no less awkward result that Jones lived happily on for ten years after Smith had killed him. Second, this kind of case seems to be the exception rather than the rule. As time goes on, the causal chain between action and consequence can be expected to become more and more tenuous as other causal factors start to influence the fate of Jones, which raises doubts about whether the case could still properly be described as a killing.

(the death). This is not incompatible with stating that there is only one action described, rather than a pair of events. It means that whenever we describe an action by making a statement of the kind 'X killed Y at time t ,' the time (interval) t should contain both the action and its consequence in order for the statement to be true. Straightforwardly describing the shooting on Saturday as a killing, and then claiming the killing occurred before the dying, is simply using 'kill' in a way that does not reflect its meaning. That there is only one action does not imply that the ways to describe it in a sentence can be substituted at will without affecting that sentence's truth value.

Returning to our original problem scenario, we can say that this elaboration of the meaning of 'kill' renders the move from step (2) (accepting AIT) to step (3) (saying that Smith killed Jones on Saturday) invalid. This solution differs from that of Bennett (1973) and Anscombe (1979) in that it does not depend on old events retroactively gaining new properties, or a description of the action becoming true later, after Jones's death, but on stating conditions for use of the word 'kill' that must be fulfilled whenever the word is used. This overcomes Mossel's (2001) worry that as soon as Jones has died, it becomes true that Smith's shooting of Jones on Saturday was Smith's killing of Jones on Saturday. After Jones's death we can certainly say that Smith has killed Jones by shooting him, but not that Smith killed Jones on Saturday, or before Jones was dead: the definition of 'kill' does not allow this.

One worry that remains is Lombard's (1989) claim that there cannot be a killing *unless* there is a death, but that this does not imply that there cannot be a killing *until* there is a death. Lombard's claim is reasonable if one assumes that 'kill' means 'cause to die': one cannot cause a death unless there is a death, but to say that one cannot cause a death until there is a death would be wrong, as causes precede effects. Assuming Wierzbicka's definition, however, there can be no killing *unless and until* there is a death. The important point here is that Wierzbicka does not suggest this definition because she makes the unless-until inference, which Lombard argues is fallacious; her reasons are observations, rather, on the common usage of 'kill,' or what 'kill' implies, and the problems that arise from equating 'kill' with 'cause to die.' We can say, though, following White (1980), that there cannot be a fatal wounding *unless* there is a death. Again, the fact that there is a difference between 'kill' and 'fatally wound' suggests that there might be more to 'kill' than just 'cause to die.' Accepting this can help us deal with the problem of the time of a killing while preserving the action identification thesis.

I will end by noting that, while this definition of 'kill' helps us to deal with the problem of the time of a killing, it does not help us to deal with the problem of the acting dead. In particular, remember that Mackie and Mossel wanted to be able to say that, in some circumstances, agents can be doing something while they are no longer active. Suppose that they wanted to say this about Smith killing Jones, they would have to say that Smith was doing something (killing) *over* the weekend. Wierzbicka's definition, however, only allows us to say that Smith was doing something (killing) *during* the weekend.

To this, Mackie and Mossel might reply that there are differences between Smith killing Jones and Brown printing a paper, such that they would want to say that

Brown was printing a paper *over*, say, the whole afternoon, but not that Smith was killing Jones over the weekend. However, Wierzbicka has argued that her analysis of ‘kill’ holds for all causative verbs, including ‘printing.’ So if Brown prints a paper by pressing the right combination of keys, and his printer takes the whole afternoon to print the document, according to Wierzbicka we can only say that Brown printed the paper *during* the afternoon, not *over* the afternoon, whether or not he died after pressing the right combination of keys.⁸

Wierzbicka’s claim has two important implications. This first is that by following Wierzbicka, Mackie and Mossel would not even have a starting example, let alone be able to solve the problem of the acting dead. The second is that, as already noted, Wierzbicka’s claim makes it possible to generalise the problem of the time of a killing (and its solution) and the problem of the acting dead from ‘kill’ to all causative verbs. Questions such as: ‘When did Brown print the paper?’ ‘When did Jones microwave the falafels?’ or ‘When did Jones insult Smith?’ might be less likely to require a precise answer than questions about killings, but the problem structure, and that of the solution, is in all cases the same.

5 Conclusion

The problem of the time of a killing is often presented as a reason to abandon the action identification thesis, as it would have absurd implications. In this paper I have argued for three claims. First, that the problem threatens the action identification thesis because of two assumptions the thesis makes: that agents’ actions last as long as their doings and vice versa. Second, that abandoning at least one of these assumptions necessarily means confronting another problem, the problem of the acting dead, to which no satisfactory solution has thus far been provided. Third, that by examining the meaning of ‘kill’ and acknowledging that it is not ‘cause to die,’ we can come to a definition of ‘kill’ that does not lead to any absurd consequences, that adequately deals with the problem of the time of a killing and that allows us to keep the action identification thesis.

Acknowledgments The author would like to thank Joel Anderson, Jonathan Dancy, Anthonie Meijers, Philip Nickel, Kevin Scott Dunn, Bart Streumer, Wybo Houkes and an anonymous referee for this journal for their helpful comments on earlier drafts of this paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

⁸ The exception to this is, of course, cases where the relevant verb denotes an activity rather than an action: Smith could have killed Jones *over* the afternoon by slowly strangling him, and Brown could have printed his paper *over* the afternoon by labouriously operating an antique printing press. These activities seem to be precisely temporally localisable. Here, however, the doing does not extend beyond the agent’s action/activity, which is the kind of case that Mackie and Mossel are interested in, but the agent remains active for the full afternoon until the intended consequence of the activity is achieved.

References

- Annas, J. (1978). How basic are basic actions? *Proceedings of the Aristotelian Society*, 78, 195–213.
- Anscombe, G. E. M. (1957/2000). *Intention*. Oxford: Basil Blackwell/Cambridge: Harvard University Press.
- Anscombe, G. E. M. (1979). Under a description. *Noûs*, 13(2), 219–233.
- Bennett, J. (1973). Shooting, killing and dying. *Canadian Journal of Philosophy*, 2(3), 315–323.
- Bennett, J. (1994). The “namely”-analysis of the “by”-locution. *Linguistics and Philosophy*, 17, 29–51.
- Davidson, D. (1980). *Essays on actions and events*. Oxford: Clarendon Press.
- Falvey, K. (2000). Knowledge in intention. *Philosophical Studies*, 99, 21–44.
- Fodor, J. A. (1970). Three reasons for not deriving “kill” from “cause to die”. *Linguistic Inquiry*, 1(4), 429–438.
- Ginet, C. (1990). *On action*. Cambridge: Cambridge University Press.
- Heinaman, R. (1983). House-cleaning and the time of a killing. *Philosophical Studies*, 44, 381–389.
- Hornsby, J. (1980). *Actions*. London: Routledge & Kegan Paul.
- Lombard, L. B. (1978). Actions, results, and the time of a killing. *Philosophia*, 8(2–3), 341–354.
- Lombard, L. B. (1989). ‘Unless’, ‘until’, and the time of a killing. *Pacific Philosophical Quarterly*, 70, 135–154.
- Mackie, D. (1997). The individuation of actions. *The Philosophical Quarterly*, 47(186), 38–54.
- Morreall, J. (1976). The nonsynonymy of kill and cause to die. *Linguistic Inquiry*, 7(3), 516–518.
- Mossel, B. (2001). The individuation of actions. *Australasian Journal of Philosophy*, 79(2), 258–278.
- Sandis, C. (2006). ‘When did the killing occur?: Donald Davidson on action identification. *Revista de Filosofía*, 37, 179–183.
- Schnieder, B. (2009). ‘By’. A refutation of the Anscombe thesis. *Linguistics and Philosophy*, 31, 649–669.
- Thalberg, I. (1967). Verbs, deeds and what happens to us. *Theoria*, 33(3), 259–277.
- Thomson, J. J. (1971). The time of a killing. *The Journal of Philosophy*, 68(5), 115–132.
- Thomson, J. J. (1977). *Acts and other events*. Ithaca: Cornell University Press.
- Weintraub, R. (2003). The time of a killing. *Analysis*, 63(3), 178–182.
- White, A. R. (1980). Shooting, killing and fatally wounding. *Proceedings of the Aristotelian Society*, 80, 1–15.
- Wierzbicka, A. (1975). Why “kill” does not mean “cause to die”: The semantics of action sentences. *Foundations of Language*, 13(4), 491–528.
- Wilson, G. (2009). Action. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2009 Edition). Retrieved June 12, 2012 from <http://plato.stanford.edu/archives/fall2009/entries/action/>.