

Killing

There has been a good deal written about the moral significance of the contrast between killing and letting die. The canonical contrast here is between cases where someone's death is the result of some positive action, as opposed to cases where someone's death is the result of a failure to do something.¹ But there has been rather less written about the distinction between cases where someone's death is counterfactually dependent on some positive action by an agent and cases where

¹ Philippa Foot (1967) did much to bring this kind of contrast to the attention of philosophers. There is a vast literature on the topic of killing and letting die, as well as on the more general distinction between doing and allowing. Particularly influential works in this area include (in addition to Foot 1967), Bennett (1995), Foot (1984) Quinn (1989) and Thomson (1986). We note that Foot distinguished two kinds of uses of 'allowing'. The use she was "concerned" with applies where an agent could "intervene" to stop a sequence of events "somehow already in train", but does not intervene, as when one "could stop a leaking tap but allows the water to go on flowing." The other "is roughly equivalent to *enabling*," as when "someone may remove a plug and *allow* water to flow" (1967, 3). Neither use maps on to the key distinction that animates this paper.

an agent *kills* someone, and there has in particular been little extended discussion of the ethical significance of this latter distinction. Section 1 introduces the distinction. Section 2 explores its conceptual underpinnings. Section 3 explores the moral significance of the distinction. Some concluding remarks follow.

I

We are all going to die. But some of us will die sooner than others. And the actions of other agents make a difference to when we die. Let us say that *x* is a *factor* in *y*'s death at *t* just in case *x* does something such that *y* would have died later than *t* had they not done that. (We are using 'factor' as a term of art here – it has obviously been chosen because it bears *some* resemblance to ordinary usage, but it is not important to us that it matches ordinary usage). And to ensure that our current topic does not get too muddled with doing/allowing discussions, let us not be too expansive about what counts as doing something. In particular we shall not count omissions, failures to do something, as a kind of doing something.

Someone can kill someone else and not be a factor (in the defined sense) in that person's death. Suppose Jones pushes a button that releases a trapdoor that sends Smith plummeting into a canyon to his death. If Jones had not pushed Smith, Smith would have been killed by a sniper right after the time of the actual trapdoor release. It is a deep canyon: Smith dies later than he would have had Jones not

released the trapdoor. Jones doesn't know that – they just want Smith dead.

Obviously, Jones killed Smith, but Jones wasn't, in the sense defined, a factor in Smith's death.²

Our focus here will be on failures of entailment in the *other* direction. It is clear that one can be a factor in someone's death without killing them. Suppose Jones takes Smith to play golf and a lunatic sniper, whose presence Jones did not foresee, kills various people on the golf course, including Smith. Jones did not kill Smith (the sniper did), and this is true even if Jones coerced Smith into going to the golf course. But it is easy to fill in the details so that Jones was a factor in the defined sense – in particular, it is easy to fill in the details so that if Jones hadn't taken Smith to the golf course, Smith would have died later.

When x is a factor in y's death, y would have lived longer had x not acted the way they did. How much longer? This varies from case to case, but it bears emphasis that the time differential provides no clean guide to whether x killed y or not. If x shoots person y in the head, x kills y, but the story can be filled in so that y would have lived only a tiny bit longer had they not been shot (as when one shoots a person who is close to death anyway); and of course it can also be filled in so that

² That is one reason why the intended meaning of our term of art does not quite match ordinary uses of 'being a factor in someone's death'.

y would have lived a great deal longer had they not been shot (as when a person is fatally shot in their youth and would have otherwise lived a long life). The relevant time differential is not going to be the key to the contrast between cases where one is a mere factor in someone's death and cases where one kills someone. What does make for that contrast?

2

'To Kill' is an example of what linguists and grammarians call a 'causative verb', a kind of verb that implies causality. (Where V is a causative verb, 'x V'd y but x had no effect on y', makes no sense.) We wish to draw attention to four plausible ideas concerning how (at least a large class of³) causative verbs work, all of which straightforwardly apply to 'kill'. The first two draw upon prominent themes in Hart and Honoré's classic *Causation in the Law*, an important resource for any study of the interplay of ethics and causation. The focus of those authors, of course, is on applications of causal ideology in a legal setting, but they emphasize that the courts' uses of causal language have "have their roots in certain features of a variety of concepts which permeate the daily non-legal discourse of ordinary men"

³ Our reason for this qualification will become clear (see the discussion of 'hasten' and 'accelerate' below).

(1985,1). And while their focus is not on causative verbs *per se*, their discussion can easily be adapted into some promising heuristics concerning how such verbs are put to work.

The first idea is this. If *x*, by doing *F*, is a factor in *y*'s death, but *y*'s death also crucially depends on some later intervention by a voluntary agent that exploits *x*'s doing *y*, then *x* does not kill *y*. This applies a general idea about causation in the law that Hart and Honoré put as follows:

. . .free, deliberate and informed act or omission of a human being, intended to exploit the situation created by defendant, negatives causal connection
(1985, 136)

And here is one of their vivid examples:

A throws a lighted cigarette into the bracken which catches fire. Just as the flames are about to flicker out, *B*, who is not acting in concert with *A*, deliberately pours petrol on them. The fire spreads and burns down the forest. *A*'s action, whether or not he intended the forest fire, was not the cause of the fire: *B*'s was (1985, 74).

The ‘intends to exploit’ in their formulation is crucial to this idea. As Hart and Honoré note, “the act of a consumer who uses a product without knowing that it is defective will not bar an action against a negligent manufacturer” (1985, 149). Similarly, if x plants a car bomb and y, ignorant of the danger, voluntarily turns the key, resulting in the death of y’s family, that does not stop us saying that x killed y’s family.

We find *something* like this idea in some more recent work on causatives. For example, Richard Thomason (2014) distinguishes ‘efficient agents’ from ‘automatic agents’ (where paradigms of the latter include the wind or a fire) and offers the generalization that “An efficient agent is precluded [from being in the relevant causal relation to the outcome of a process] by another efficient agent between it and the body [of a process]” (2014,70. Note, though, that for reasons just given, Thomason’s generalization is too incautious.)

Hart and Honoré note that intermediate agents that are acting out of ‘self-preservation’ (1985,144), so that significant harm will come to them by acting otherwise, or who are acting under some legal or social obligation, are not intended to count as engaging in “free, deliberate and informed acts”. In this way historians often give some number in the millions when asked how many people Stalin

killed.⁴ As Thomason would put it, when we use ‘kill’ in this kind of way, various intermediate animate agents are ‘automatized’.

Here is a rough first pass at the second idea: if x is a factor in y’s death by doing F, but the death crucially depends upon an abnormal intervention between x’s Fing and y’s death, then x does not kill y. Hart and Honoré offer a version of this idea, generalized to causal responsibility in general:

The basic principle is that normal physical events, even subsequent to the wrongful act, do not relieve a wrongdoer of responsibility but that an abnormal conjunction of events (in this case the wrongful act and the third factor) negatives causal connection, provided that the conjunction is not designed by human agency (1985, 162-163).

They emphasize that it is crucial that the abnormality is posterior to the action. Here are few contrast cases, inspired by their discussion:

Pair 1.

⁴ See for example <https://news.stanford.edu/2010/09/23/naimark-stalin-genocide-092310/>

Scenario 1. Jones pushes Smith over and Smith falls on to a spike in the ground at the spot where he lands, with fatal effect.

Scenario 2. Jones pushes Smith over and a tree falls on top of Smith while Smith is laying on the ground. Smith is crushed and dies.

Intuitively, Jones kills Smith in scenario 1 but not scenario 2.

Pair 2.

Scenario 1: Jones drops a brick out of a second story window, landing on Smith's head (who is standing beneath the window), with fatal effect.

Scenario 2: Jones drops a brick out of his second story window, and as it is falling, a bolt of lightning hits it, redirecting it to the neighboring garden, where it lands on Smith's head, with fatal effect.

Scenario 1 is a normal continuation scenario. 2 is not. We are much more comfortable saying that Peter killed Steve in scenario 1 than in scenario 2.

We glossed our presentation of the second idea as a “rough first pass”. Hart and Honoré's caveat “provided that the conjunction is not designed by human agency” points to some need to qualify. Suppose Jones knows that a tree is about to fall because of some abnormal weather conditions and knows where it is going

to land. Jones pushes Smith so that Smith lands on the relevant spot just prior to the tree's breaking and falling on that spot. Here the conjunction of falling and then being hit by a tree is "designed by" Jones, so the abnormality of the breaking does not block causality.⁵

The idea that abnormal interventions block causality has been taken up in some recent literature, including the substantial literature on causal modelling.⁶ For example, as a prelude to a formal model of causal strength, Icard, Kominky and Knobe offer the principle of 'Abnormal Inflation' that says that if C and A are both factors in an outcome, "people will be more inclined to say" C is the cause when "they regard C as abnormal than when regard C as normal" (2017, 81) and the principle of 'Supersession' that says that people are less inclined to treat C as the cause when A is abnormal than when A is normal. (We note in passing, though, that temporal order does not figure in these principles: arguably principles that take account of temporal order will be even more predictive of folk practice).

⁵ The example of x pushing y to the ground just as a tree is about to fall on them is discussed at length in Hart and Honoré, Chapter 3. There they make the point that foreknowledge of the event makes a difference.

⁶ See, for example, discussions of normality and deviance in Gallow 2022, Hall 2007 and Halpern 2008.

Elsewhere, Knobe remarks that folk judgments of causation are driven by which counterfactuals they focus on rather than ignore. Moreover, he goes on to suggest (citing Kahneman & Tversky 1982; Roese 1997) that people are inclined to “. . . .consider counterfactuals in which events of *unusual* types are replaced by events of *usual* types“(2009, 241).

Let us juxtapose this abnormality driven idea with David Lewis’ suggestion about what is distinctive about killing, namely,

. . .insensitivity to circumstances... if you shoot at your victim point—blank, only some very remarkable difference in circumstances would prevent his death. The same is true if you set... a delayed action-bomb, working inexorably towards its lethal outcome. The case of a bomb with a randomizer is comparatively insensitive: the bomb might well have chanced to go off, but it isn’t the fine details of the circumstances that would make the difference (1986, 186).

The idea (adapted to the ideology of factors) is that when you are a factor but don’t kill, the death depends much more on the fine details of circumstances than when you kill. We don’t think this does as well as the normality-theoretic

idea, at least when it comes to understanding folk practice. Suppose Jones takes Smith to Spain. A powerful atomic bomb has already been directed to Spain, big enough that there is no chance of anyone in Spain surviving when it hits. Smith's subsequent annihilation does not seem to depend very much on the fine details. Compare this with a case where Jones shoots at Smith from a distance while Smith is running. Jones is not highly skilled. If Smith had slowed down just a tiny bit or speeded up just a bit, or if Jones had aimed just a tiny bit to the left or right, Smith would not have died. Comparatively speaking there is more sensitivity to fine details. (And we can ramp up the sensitivity even more in the second case by having the bullet enter a small hole in a slightly defective bullet proof vest.) But in the latter case, Jones kills Smith, whereas in the former Jones does not. Contrast similarly a case where Jones takes Smith into a car that has a car bomb with no randomizer with a case where Jones pushes Smith who lands on a spike. The death of Smith may depend on fine details of spike location and the trajectory of the fall, while death subsequent to the turning of the key may be comparatively less sensitive to details. At least if we are trying to understand what drives ordinary judgments, normality-theoretic ideas seem superior to Lewis'.

A third theme—one that is prominent in work (some co-authored) by Joshua Knobe is that *ceteris paribus*, people are more inclined to consider bad actors the

cause rather than good actors even if both are factors.⁷ Here is an example much in the spirit of this discussion, adapted to killing⁸:

A patient needs 50 mls of drug A each day. However, if that patient receives 100 mls over the space of an hour, they will die. One morning, within the space of an hour, the nurse gives the patient a 50 ml dose and someone who shouldn't be handling hospital drugs also gives the patient a 50 ml dose. (To rid ourselves of distractions due to temporal order, we can make the injections simultaneous – perhaps the nurse cannot see the other injector because they are separated by a screen or can see the other injector, who is dressed as a nurse and who tells the real nurse they are injecting something else).

⁷ This a central theme of Knobe (2009), where draws on the principle that “people are inclined to consider counterfactuals in which bad events are replaced by good events” (2009. 240), which then has knock on effects for judgments of causality. See also Hitchcock and Knobe (2009).

⁸ Knobe and Hitchcock (2009) present cases very similar to the one that follows.

Clearly, we are more inclined to say that the imposter killed rather than the nurse. Of course in this case there is obviously a clear normality contrast as well – what the nurse is doing is normal, but what the imposter is doing is abnormal. But the case can be tweaked so that the relative normality of what is being done is evened out or reversed. Suppose for example the nurse normally administers drugs in the evening and the imposter would have done their thing in the morning come what may. Here the fact that the nurse is injecting at an abnormal time is highly relevant to the patient’s death, and so both interventions are abnormal. Another example: suppose that the hospital is full of sickly enemy soldiers and the administration quite frequently allows government agents to interfere by adding fatal extra doses here and there to supplement drug A medication. Here the intervention by the non-nurse, while bad, may be quite normal. If we suppose in addition that the nurse normally administers drugs in the evening, then it may be that the nurse’s intervention was more abnormal than that of the non-nurse. The asymmetry in badness thus need not come with an asymmetry in abnormality, and still the effect of badness on causal and causative judgments persists.⁹ (Of course there is still a

⁹ For some related points, see Knobe and Fraser’s (2008) where they critically discuss Julia Driver’s (2008) suggestion that it is atypicality and not morality per se that explains the relevant data.

sense in which the bad actors are acting abnormally: Icard et al., 2017, 81, subsume cases of bad actors under abnormality principles by including violations of *prescriptive* norms as well as statistical abnormality within the category of the abnormal.)¹⁰

It is also worth noting that the asymmetry persists even if we set things up so that the protagonists have the same subjective risk of doing something that has the

¹⁰ It's a good question how the bad/good principle interacts with the other ideas discussed here. For example are we more willing to say x killed y when x kidnaps y and a tree freakishly falls on the back of the car while x is driving y than when x is taking y to the hospital and a tree freakishly fall on the back of the ambulance, killing y? Hart and Honoré cite plenty of case law in which an abnormal intermediary blocks causation by some agent even though some bad action by that agent is a factor. (See for example their discussion of *Toledo & Ohio Central R. Co. v. Kibler & Co* on 163-164). This suggests that making the initial action bad has little effect on the capacity of subsequent abnormal interventions to block killing. We shall not pursue this further here.

patient's death as a consequence.¹¹ Suppose for example, that the imposter has a million vials, knows one of them has drug A (the rest being saline), but doesn't know which. Hoping to kill the patient, the imposter (who has only time to use one vial before making their escape) by chance injects the vial containing drug A. Meanwhile the nurse knows there is an overdoser on the loose. But it's a big city and the temporal window for fatal double dosing is quite short, so their credence is one in a million that the overdoser will be a factor on this occasion. Here the subjective risks are the same but there is a clear asymmetry in our disposition to make judgments about who killed who.¹²

Let us turn to the fourth idea. The key thought here is that when it comes to judgments of killing, there is a big difference between cases where one is a factor

¹¹ We use 'consequence' in a way that is intended to be neutral between killing and not killing. That seems to fit with ordinary usage. If Jones drops a brick out the window which is then transported to Smith's head by a freak weather event, then even though we demur from Jones killed Smith', we are much happier with 'Smith's death was a consequence of Jones' dropping a brick out of the window'.

¹² Of course if the nurse knows that the imposter is administering A then if they administer A themselves they are also at fault. So we don't want to even up subjective risk in *this* kind of way, since that would remove the moral asymmetry.

because one accelerates a death-culminating process that has already been initiated, and cases where one *initiates* a process.¹³ Consider the following pairs:

A1 Smith has been bitten by a Death Adder. Jones (perhaps out of a desire that Smith does not suffer too much) squeezes Smith's extremities in such a way as to accelerate the action of the venom. As a result Jones dies one minute earlier than if Jack did nothing.

A2 Smith has been bitten by a Death Adder. Jones (perhaps out of a desire that Smith does not suffer too much) injects a faster acting poison into Jones's bloodstream (or alternatively, shoots Smith through the head): As a result, Smith dies one minute earlier than if Jones did nothing

B2 Smith is in a pot of water over a flame. Jones, a cannibal, wants dinner quickly. Jones adds large amount of salt to the pot (which has been placed

¹³ The ideology here is rather similar to Foot's concept of a sequence "thought of as somehow already in train" (1967, 3).

over the flame by other cannibals).¹⁴ Smith dies one minute earlier than they would have if Jones hadn't added the salt.

B2 Smith is in a pot of water over a flame. Jones, a cannibal, wants dinner quickly. Jones (who is very strong) lifts the pot and moves it to a red-hot griddle some yards away. Smith dies one minute earlier than they would have if Jones had left the pot alone

In A2 and B2, we are more inclined to say that Jones killed Smith than in A1 and B1. (Similarly for various other causatives. For example, we are more inclined to say that Jones *boiled* Smith in B2 than in B1, and, assuming that both the venom and the faster acting poison work by paralysis, we are more inclined to say that Jones *paralyzed* Smith in A2 than in A1).

¹⁴ We realize that it is something of a folk myth that salt does much to speed up boiling. Readers who care about this should substitute some chemical that really does substantially speed things up.

Clearly this kind of contrast is playing a role in discussions about when various sorts of actions during palliative care constitute killing, suicide and so on. Here is the Michigan Supreme Court:

There is a difference between choosing a natural death summoned by uninvited illness or calamity, and deliberately seeking to terminate one's life by resorting to death-inducing measures unrelated to the natural process of dying (cited by Cantor, 2006, 411).

Here is the idea. Accelerating the “natural death” by, say, switching off life support or voluntarily lessening one’s food and water intake, is a matter of the “natural process of dying” playing itself out. By contrast, measures that kill in a way that does not merely accelerate the natural dying process already underway, but which induces death by different means are a completely different matter.

Hart and Honoré have a variant of the idea presented here: It is that one merely accelerates death (as opposed to ‘causes death’) when (a) one slightly accelerates death and (b) one does something that “would not be sufficient to kill a person in a normal state of health” (1985, 344). This may in one way suggest an improvement to our own formulation, but in one other way it does worse. In our cases, the process accelerated makes only a modest amount of difference to the

time of death. But what if it made a difference of days rather than minutes? Here we do seem to be more willing to count the relevant action as a killing, but intuitions are still pretty murky. Suppose a snake venom was untreatable, but unaccelerated would take two agonizing weeks to kill someone. Someone massages their extremities, accelerating the passage of venom, and the victim dies a full week earlier than they would have done. How much more comfortable are we saying the massager killed the victim? We don't think the judgment is especially clear. Here it is worth noting that various palliative care activities that accelerate death by days rather than minutes are nevertheless *not* treated as killing (though it is unclear how much self-delusion is going on here)¹⁵. Thus, while the length of time by which a life is shortened may make some difference in our willingness say that x killed y, the judgments are not sharp enough to warrant further epicycles on our initial articulation of the idea. In any case, the ideas articulated here mark generic pressures for and against judgments of killing rather than clean exceptionless generalizations, and should be taken in that spirit.

¹⁵ Thus Cantor writes: "It is self deception if people think they are not killing anyone when they deliberately choose a regimen of treatment which they know will result in the patient's death when there is an alternative which will keep the patient alive" (2006, 411).

In another way our articulation is superior. Suppose Smith is in a hospice. Jones wants Smith dead and lets a venomous snake into the bed. This is the kind of snake – a Pygmy Rattlesnake perhaps -- whose bite will be survived when suffered in a good state of health, but is fatal to people who are as weakened as Smith is. Smith, as expected, dies from the bite. Even if it was true that Smith would have in any case died a very short while later from the illness that put them in the hospice, Jones killed Smith. But Hart and Honoré’s formulation suggests a different (and incorrect) judgment, since a bite from that snake “would not be sufficient to kill a person in a normal state of health”.

We note in passing that when someone’s action is a factor but not a killing we are – as Hart and Honoré’s discussion suggests -- at least reasonably comfortable as describing the case as one where one agent ‘accelerates’ or ‘hastens’ the death of another. (And we note that the expression ‘hastening death’ is quite pervasive in the palliative care literature.¹⁶) Could our main topic have indeed been presented as the contrast between killing and merely hastening death (or as the contrast between killing and merely accelerating death)? We are not at all averse to this gloss, but we chose our official formulation of the central topic in part because ‘factor’ is clearly defined and in part to sidestep one tricky issue.

¹⁶ See for example, Cavanaugh 1996, Cantor 2006 and Billings 2011

Verbs like ‘hasten’ and ‘accelerate’ are typically treated by grammarians as causatives. Yet while the ideas about killing presented above apply to all sorts of paradigmatic causatives, they do not seem to apply so smoothly to ‘hastening’ and ‘accelerating’. In the cases where someone action is a factor in another death and they don’t kill them, we seem to be more comfortable saying they accelerate death and that they hasten death.¹⁷ Does this mean that the classification of them as causatives is not quite right (so that hastening death does not require causing a death to come about)? Or does it mean that what Hart and Honoré say about ordinary concepts of causation is really only true of an important category of causatives? (One version of this idea has it that there is a special relation semantically associated with this category of causatives but it is not one generally expressed by the ordinary verb ‘cause’; and that it is *this* relation, rather than causation, that is the real target of Hart and Honoré’s meditations.)¹⁸ Or are these

¹⁷ Of course in certain end of life discussions one suspect that the words ‘hasten death’ are in play simply to avoid admissions of guilt rather than to mark an interesting distinction.

¹⁸ We note that ‘x boils y’ and ‘x kills y’ make sense when y is an object, but ‘x hastens y’ only seems to make good sense when y is an event. We suspect that this may be highly relevant to the issues raised in this paragraph. We also note that

uses of ‘hastening’ somewhat metaphorical, one that goes beyond a core meaning?¹⁹ We did not want to get into this. For convenience, though, we shall in the following discussion take it that when it is true that one is a factor in someone’s death, one accelerates their death.²⁰

There is a related question. Could our core topic have been posed as the difference between being a factor in someone death and causing someone to die? The discussions cited by Hart and Honoré and by Knobe would seem to suggest that the kinds of consideration that make us fall short of calling a factor a killing

there is substantial literature within linguistic on the relation between causatives and so called ‘periphrastic causation’ constructions, where an effect combined with a making relation are lexicalized (e.g. The sergeant made him march’, ‘The director caused him to quit’), where the consensus is that periphrastic constructions are also less demanding than paradigmatic causatives. For a helpful recent survey of this and related issues, see Martin and Schafer (2014).

¹⁹ Some informants felt that one would be taking semantic liberties if one said ‘x hastened y’s death’ in a case where x takes y to a golf course and a sniper shoots y.

²⁰ Note though that we do not want to collapse the distinction between accelerating someone’s death and accelerating their death by accelerating a process that has already been initiated. The latter are a subspecies of the former.

would also make us reluctant to say that the relevant agent caused someone to die by performing the relevant action. But the issues are delicate: Take the case where Jones's brick is redirected by lightning. Consider: (i) Jones killed Smith. (ii) Jones was the cause of Smith's death. (iii) Jones caused Smith's death. (iv) Jones was a cause of Smith's death. The fourth sounds rather quite a bit better than the first to our ear and the third a bit better (though we may have been corrupted by philosophical fashion). A famous example from Katz has generally convinced linguists that 'cause to die' and 'kill' can't line up perfectly: Suppose a gunsmith's defective work causes a sheriff's gun to jam, with the consequence that the sheriff is gunned down. 'The gunsmith caused the death of the sheriff' sounds a lot better than 'The gunsmith killed the sheriff'. A common reaction among linguists is to suppose that 'kill' requires "direct causation" but as Neeleman and Van de Koot (2012) point out, this can't be right: if we add to the story that the gunsmith plotted the death of the sheriff and deliberately tampered with the gun, the causal connection is no less indirect, but the "The gunsmith killed the sheriff" is much more acceptable. (Recall also that Stalin killed millions.) It is also not clear that a lack of intention systematically induces a refusal to apply 'kill' in more indirect cases: a recent discussion by a lawyer of product liability contains the turn of

phrase ‘when a company kills through its negligence’.²¹ And similar examples are easy to find.

How great a mismatch is there between ‘kill’ and ‘cause to die’? Is the kind of causal relation implicit in causatives like ‘kill’ and ‘boil’ *vastly* more demanding than the relation typically expressed by ‘cause’ (in which case, ‘cause’, though a causative is perhaps a bit more like ‘hasten’). This is something like Lewis’ view, according to which causing people to die is a commonplace achievement: “I am sure that I – and likewise you, and each of us – have caused ever so many people to die, most of them yet unborn. But I have never killed anyone – I hope... so killing must be a special kind of causing to die.” Or is the relation expressed by ‘cause’ in ordinary discourse much more demanding than that suggested by approaches like David Lewis (1973), so that the extension of ‘caused y to die’ and that of ‘killed y’ is at least quite close?²² Because we don’t wish to try to sort all this out here, we

²¹ <https://www.fieldfisher.com/en/injury-claims/insights/new-law-to-prosecute-negligent-companies>

²² Note that Lewis (1973) famously took counterfactual dependence between events x and y to be sufficient (though not necessary) for the truth of ‘x caused y’. If he is correct, then ‘Kyle’s dropping of the brick caused Steve’s death’ is true in both versions above. Once we have got that far it is hard to resist accepting ‘Kyle

shall continue to focus on the distinction between cases where one kills and cases where one is a factor in a death but doesn't kill, and will not pass judgment on the relation between 'killing' and 'causing death'.

It is also worth making salient one reason why our taxonomy can only give rather soft predictions about folk judgments using causatives. Suppose that Jones makes a defective television that then explodes in the presence of Smith. Smith dies. Candidate uses of 'kill' include 'The television killed Smith' and 'Jones

caused Steve's death'. (A natural bridge principle is that if x is the agent of an event that caused some event y, then 'x caused y' is true.) It bears emphasis that 'Jack caused Jill to die but Jack didn't kill Jill' doesn't exactly roll off the tongue, which one might expect it to given Lewis' vision. However Katz's Wild West story suggests, *prima facie*, that the relation between 'cause to die' and 'kill' is not quite entailment. Knobe's (2009) picture is that folk judgments of causality are driven by a selective filtering of counterfactuals – it's a matter of the right counterfactuals breaking a certain way. The approaches can be made compatible if we suppose that the truth of causal claims in English come apart quite dramatically from ordinary folk's willingness to accept them. It is beyond the scope of this paper to inquire as to which aspects of ordinary practice should be semanticized.

killed Smith'. But contexts have a 'pick one' flavour.²³ Here there is a choice between an 'automatic agent' (recalling Thomason's ideology), namely the television, and a bona fide agent, Jones. But it is hard to pick both: The speech 'The television and Jones both killed Smith is very awkward'. (Contexts in which one says 'The cause' also have a 'pick one' flavor: in this case, the 'pick one' aspect is particularly salient because of the appearance of the definite.) For the purposes of product liability practice in tort law – one kind of setting which animates Hart and Honoré – questions framed directly using causatives like 'kill' are often sidestepped in settling liability in favor of expressions along the lines of 'is responsible for', a phrase that figures prominently in Hart and Honoré's own discussion. But if one is in a context where one wants to deploy the causative 'kill', then one will have to pick one out of the television and Smith to fill in 'X killed Smith'. And the above discussion doesn't really provide that much guidance as to how ordinary folk resolve the competition between the television and Smith. For what its worth we think in this case it is much more likely they will go for 'The television killed Jones', but the factors resolving the competition between animate and inanimate agents as the preferred agent of the causative seem quite subtle and

²³ We also think it plausible that even 'X caused the death of Smith' also has a pick one flavour too.

we cannot hope to fully catalogue them here.²⁴ When the activity of the inanimate object is subsequent to the animate agent and is itself an abnormal intervention, there will be strong pressure to supply the inanimate agent. Thus, when a freak wind or lightning bolt carries the brick to Smith's demise, we say it is the wind or the lightning bolt that killed Smith, not the person who dropped the brick out the window. But where there is no abnormal intervention of that sort we may often *still* choose the inanimate factor over the animate one on account of its temporal proximity.

Of course one might think that 'Smith killed Jones' is true, even if not asserted, in a context when one says 'The television killed Jones'. But that makes it puzzling why 'The television and Smith killed Jones' is not felicitously assertable.

²⁴ Consider the following pair, which also point to some subtleties, this time involving animate agents: (a) Smith throws Jones through a pane of glass: Here 'Smith broke the glass' and 'Jones broke the glass' both seem passable. (b) (adapting a case from Hart and Honoré, though their focus is on 'caused the glass to break' rather than 'break the glass': Smith hits Jones, "who staggers" and then falls against the glass. Here, even though "the second agent's role is hardly an 'action' at all," (1985, 76) 'Smith broke the glass' does not sound acceptable.

Granted ‘The television but not Smith killed Jones’ is awkward: But ‘The television and not Smith was the cause of Smith’s death’ is awkward too – and in the latter case it seems fairly clear that there is a semantic requirement of uniqueness. Notice also that the ‘pick one’ phenomenon even applies to many cases where there is a pair of candidate *animate* agents. Suppose Stalin sends a note to Jones instructing Jones to kill Smith. There is a natural context in which we say ‘Stalin killed Smith’ and one in which we say ‘Jones killed Smith’. But ‘Stalin and Jones killed Smith’ is only natural in contexts where neither’s contribution was sufficient and it took teamwork to kill. We shall not pass final judgment on the extent of context-dependence here. It is enough to notice that when it comes to making causative judgments, folk often find themselves choosing between certain animate agents and temporally more proximate agents, and the resolution of this competition is quite delicate.

We now have a preliminary taxonomy of cases where x is a factor in y’s death but does not kill y: let us call the first kind of case above ‘mediating agent cases’, the second ‘abnormal event cases’, the third ‘bad action cases’ and the fourth ‘accelerated process cases’. Perhaps the taxonomy is not completely exhaustive. But it will provide us enough to work with when probing the moral significance of killing.

In what follows we shall be focusing our discussion on the second and fourth of these categories. In a mediating agent case, it is very natural to suppose that the badness, responsibility, level of guilt etc. of the original agent is somewhat diminished vis a vis cases of killing, since the mediating agent takes on at least some of the guilt and so on. Meanwhile, if badness of the agent contributes to whether an action counts as a killing then it is quite obvious that in *this* way, being a killing rather than a mere factor can have moral significance, but for uninteresting reasons, since in bad action cases it is the badness of the agent that explains why the event is a killing rather than the other way around. But how about the second and fourth subcases? In the next section we look at these two subcategories, with an eye to exploring the moral importance of the distinction between those cases and cases of killing.

3

Is there is a morally important distinction between the class of abnormal event cases and the class of cases that are killings? And is there a morally important distinction between accelerated process cases and the class of cases that are killings? In this section we attempt to probe these questions, tentatively arriving at negative conclusions. Of course we haven't providing anything like informative necessary and sufficient conditions for 'to kill' and there may also be some

measure of semantic context-dependence and vagueness. Yet we have provided what we hope is a helpful guide to the sorts of considerations that help contour the extension of ‘to kill’ in the mouths of English speakers, and our investigation will proceed on the assumption the truth conditions of ‘kill’ sentences are sensitive to the presence and absence of abnormal interventions and to the folk distinction between initiating and accelerating a process.

We think it might be helpful to look at cases where there are choices that are symmetric with respect to both the risk of death and whether the potential victim’s death is motivating, but where there is a risk of killing in one case and a risk of mere acceleration in other. We begin with cases where the protagonist does not want an early demise for the potential victim. Call these ‘mere risk cases’. We then look at pairs of cases where the victim is intentionally accelerated on one option and intentionally killed in another, though where, in neither case is the outcome foreseen. Call these ‘intentional but unforeseen cases’. We finally at a choice between actions that are known to accelerate death but where only one is a killing. Call such cases ‘knowledge cases’

3.1 Mere Risk Cases

If one takes a small risk of producing an outcome but are not in any way motivated by that outcome, one neither produces that outcome knowledgeably nor intentionally. Is there an important ethical difference between risking killing someone and risking being a factor in their death?

Consider the following case. (It is obviously extremely stylized, but extremely stylized cases are often useful for probing intuitions):

One is skiing out of control down a ski slope. One can't but collide with Smith but one can control whether one knocks Smith down the right or left slope of a hill. At the bottom of the left slope it is known that there are no spikes but just very occasionally, trees fall on to that area. At the bottom of the right slope there is no chance of having a tree fall on someone, but just very occasionally, people fall on spikes. Suppose, as one's credences go, if one knocks Smith to the left, there is zero chance Smith will die from spikes and a one in a million chance that Smith will die from a tree falling on them just after they reach the bottom, and if one knocks Smith to the right, there is zero chance they will die from falling trees but one in a million chance they will die from a spike.

If one attached special disvalue to killing, one would have a clear preference for risking death from tree fall to death from a spike. Indeed, if one attached special disvalue to killing, one would presumably go for the tree-risk option over the spike-risk option even in certain circumstances where the risk of death from tree fall is *higher than* the risk of death from spikes. (Of course, how much higher a risk one would tolerate would depend on how much extra disvalue one attaches to killing in particular). Suppose for example that we had the choice between giving Smith a one in a million risk of death from a spike or instead one in half a million risk of death from a tree falling on them. We submit that it seems quite strange to go for the tree risk over the spike risk on the grounds that the eventuation of the tree risk would not constitute a killing and the eventuation of the spike risk would. Its not that we can't imagine a person who attached special disvalue to killing per se in this way. Such a person would say 'Granted, its twice as likely that Smith will die if I avoid the spike risk in favor of the tree risk. But there is no risk of my killing them with the latter option, only of my accelerating their death by pushing them to where a tree is about to fall; so that is obviously preferable.' But we submit that this is a rather strange way to think in a scenario like this. This preference for tree risks seems quite fetishistic when it involves putting someone's life at greater risk.

The pattern persists even when the risk differential is reduced. Suppose a morally good skier was in the unfortunate position of having to either expose Smith to a one in a million chance of death from spikes or a one in eight hundred thousand chance of death from tree fall. It would be rather shocking for that skier to expose Smith to a greater risk. This indicates to us that when one looks head on at the distinction between killing and being a mere factor (thanks to an abnormal event), it is hard to give it moral weight.

Let us turn to a case that involves the choice between a risk of killing and a risk of accelerated process (without killing).

A heavy cannonball is rolling down a chute from the top of a hill to the bottom, placed there by a malevolent agent who has ill designs for Smith. One knows it is heading towards Smith and that if one does nothing, Smith, who is tied up at the bottom of the hill, will die at 2pm. One is motivated to save Smith. One has a choice. (a) One can pull a lever which simultaneously lifts the top of the chute at the top of the hill and moves it slightly rightward. One is very confident that this will save Smith by redirecting the chute away from him. But there is a 1/100 chance that it will simply accelerate Smith's death by making the chute steeper and thus accelerating the descent of the ball but failing to redirect it away from Smith. In this case Smith will die at

1.55 (b) One can direct a second cannonball down a second chute. One is 99/100 confident that things are set up so that this cannonball will, further down the hill, collide with the first chute, redirecting that chute and saving Smith. But there is a 1/100 chance that things are set up so it will instead head towards Smith at an even faster rate, beating the first cannonball to the bottom of the hill, so that Smith dies at 1.55.

Here the dialectic is rather similar to that of the previous case. Suppose one attached disvalue to killing vis a vis death by accelerating a process. Then one should presumably opt for the lever over the second cannonball and maintain that preference even if the risk of death is a bit higher.²⁵ Disvalue killing enough and

²⁵ Admittedly there is logical space for a view that attached special disvalue to killing but which opted for the least risk of death when there was a risk of death mismatch among options – on this view, any special disvalue of killing only comes into play when the risks of death *exactly match*. As a toy model, imagine a lexical priority view where only a perfect match in expectation of death would lead one to proceed down to the lexically secondary question of whether an event, if a death, is a killing. We invite readers to explore variants on this lexical priority view for themselves.

one should prefer the lever over the cannonball even if the risk of Smith dying is double on that option. And more generally, if one attaches special disvalue to killing, one would expect there to be some extra risk r such that one would prefer the lever over the second cannonball even though the former add a risk r of death vis a vis the second. Once again, it seems very strange to imagine opting for the riskier option on such grounds. Wouldn't we look askance at an agent who exposed Smith to the greater risk when a second cannonball, with lesser risk, was available?

3.2 Intentional but Unforeseen Cases

One can intentionally do things without knowing that one will do them. Suppose for example that Jones wants Smith dead. Jones has a crossbow. They knows they are a bad shot – at the kind of distance that Smith is, they tend to miss nine times out of ten. But Jones gives it their best try. They do their best to aim, pull the

trigger, and the arrow pierces Smith's heart, with lethal effect. Jones doesn't know that they will succeed in killing Smith but Jones does so intentionally.²⁶²⁷

One might take the view that while killing does not in general have disvalue over accelerating process and abnormal event cases, there is special disvalue to

²⁶ A more controversial thesis is that in these cases one intends the outcome. In the theory of intentional action, the 'Simple View' that intentionally F-ing entails an intention to F is hotly contested. (The label 'The Simple View' is due to Bratman 1984, who contested it. For proponents of the Simple View (including Adam, 1986 and McCann 1991), the inference from intentionally killing to having an intention to kill is utterly straightforward in any particular case, because it is underwritten by entailment. For the denier, the issue is less straightforward.

²⁷ Seen through the lens of a Hart and Honoré style framework, the reason this intuitively counts as a killing is that while it is abnormal for her aim to be accurate, the fact that she is pointing in the right direction just as she is finishing pulling the trigger is part of the scene, not an abnormal event subsequent to the pulling of the trigger. If she aimed in the wrong direction but a bolt of lightning or a freak tornado redirected the arrow, an attribution of killing would be much more dubious.

intentionally killing vis a vis intentionally accelerating death via abnormal events or by accelerating lethal processes.

Here we can look at some pairs that manifest this contrast.

3.21 Accelerated Process Version

A heavy cannonball is rolling down a chute from the top of the hill to the bottom, placed there by someone who has ill designs for Smith. Jones knows it is heading towards Smith and that if Jones does nothing, Smith, who is tied up at the bottom of the hill, will die at 2pm. Jones wants Smith's death to come sooner rather than later. Jones has two ways of trying to achieve this. (a) Jones can pull a lever which lifts the top of the chute. She is .2 confident that it will accelerate the cannonball by sharpening the steepness of the chute, so that Smith dies five minutes earlier than if Jones had done nothing, and is .8 confident that it will make no difference. (b) Jones can direct a second cannonball down a second chute. Jones is .2 confident that it will travel much faster to the bottom of the hill than the first cannonball, so that Smith dies five minutes earlier than if Jones had done nothing, and .8

confident that it will be slower and that the first cannonball will still kill Smith at 2pm.

3.22 Abnormal Event version.

Jones wants to kill Smith. Jones is deliberating between two options. Jones can push Jill down the right slope of a hill into Avalanche Valley where Jones is 1/10000 confident that an avalanche will occur shortly after Smith gets to the bottom of the slope from which she will not survive.²⁸ Or Jones can push Jill down the left slope of a hill into Spike Valley where she is 1 in 10000 confident that Jill will fall on a spike and die. (Jones is certain there will be no avalanche in Spike Valley and that there are no spikes in Avalanche Valley.)

²⁸ ‘Avalanche Valley’ is so called to help the reader remember the case. The idea isn’t the avalanches are completely normal there. Lets imagine that there is only an avalanche only once every three years, though when there is one it kills everyone at the bottom of the valley.

Again, holding Jones's desire for Smith's early demise fixed, we struggle to see a significant moral difference between these cases. Suppose, for example, Jones opts for the lever and Bill dies at 1.55. It would sound rather hollow, as an attempt to argue for a lesser gravity of sin, for Jones to make the speech 'Well at least I didn't kill Smith'. Now of course the moral gravity of the cases can be varied by varying Jones's motivational structure. A more sinister (and natural) version of the case is one where Jones has ill will towards Smith and wants Smith dead sooner rather than later for malicious reasons. But alternatively, we could fill in the details so that a fast-moving cannonball produces a more painless death and Jones wants to accelerate death because the death then involves less pain.²⁹

Varying the cases along a parameter like this makes a difference to the cases. But so long as the cases are symmetric along this parameter, the mere contrast between intentionally killing and intentionally accelerating death does not seem to have much of a moral grip.

²⁹ In this variant, the claim that the death is accelerated intentionally is perhaps a bit more tendentious. But for what its worth our judgment is that it is still straightforwardly correct. (Note that only in the presence of the "Simple View" -- see fn. 27 -- does this judgment conflict with Double Effect sensibilities.)

3.3 Knowledge cases.

One could in principle opt for the view that while there is no value contrast *per se* between the class of killings and the class of accelerations, there is an important moral contrast between knowingly bringing about a killing and knowingly bring about a death that is not a killing. Does a suggestion along those lines hold up?

When it comes to the knowledge cases, it is hard to find versions of the mere acceleration type that fit the ‘abnormal event type’. Recall that when the agent has advance knowledge of the relevant abnormality, the conjunction of the death and the abnormality will typically be ‘designed’ by human agency, and so the abnormality downgrade from killing to mere factor will be hard to achieve. But it is easy to contrive foreknowledge cases that are accelerated process cases and not killings. For example, one can easily tweak the cannonball case so that one has a choice between a pressing a lever that one knows will accelerate the initial cannonball and rolling a second cannonball that one knows will arrive ahead of the first cannonball. Once again, we can fill out the case either by imputing wholly malicious motives (“I want Smith dead as soon as possible because I can’t stand Smith”) or more noble ones (“If I do nothing this death will be a slow and agonizing one and a fast-moving cannonball will be far less painful”). But holding

fixed the motives, we find it extremely difficult to see a moral difference between the two cases. As we vary motives across cases, we get markedly different reactive attitudes: Our reaction to someone who is concerned about agony is very different to someone who values the shortening of a life as an end in itself. But if we hold motives fixed and the known duration of acceleration fixed, and merely vary whether there is known acceleration without killing or known acceleration with killing, we find it much harder to discern a palpable moral difference.

Now, of course, for many warm up illustrations of the contrast between killing and mere counterfactual dependence, there is a marked epistemic contrast. When Jones takes Smith to play golf and Smith dies from sniper fire, Smith's death is utterly unforeseeable by Jones. By contrast when Jones shoots Smith through the head, Smith's death is utterly foreseeable. Nevertheless, once we ensure epistemic parity by finding cases of killing and cases of mere acceleration when the outcomes are known in each case, a moral difference between killing and mere acceleration is difficult to discern.

Our tentative conclusion is that the distinction between being a mere factor in a death and killing someone is not a distinction with any straightforward moral significance. We hope that those who are not yet quite convinced will still find the preceding reflections useful: They do something to display how ordinary practice

drives a wedge between killing and being a mere factor, and present a series of helpful thought experiments that probe the moral import of that wedge. It is worth remarking that our judgments about the moral insignificance of the killing/factor distinction are plausibly reinforced by the realization that the distinction between killing and being a factor implicit in ordinary practice is extremely soft from a metaphysical point of view. If Jones pushes Smith onto a spike, Jones kills Smith. The spike may have been left by a malicious person who was titillated by the idea of someone falling it on it. Nevertheless, even when informed of this, we suspect people will judge that Jones killed Smith. Suppose instead that Jones kidnaps Smith, turns the key of the car, and a car bomb that Jones was unaware of explodes, placed by a person that was titillated by the idea of the people in the car getting killed. Here we suspect that ordinary people will be less inclined to say that Jones killed Smith. (We made it a kidnapping so that in each case Jones did something bad). But the contrasts driving these judgments seem insubstantial. Hart and Honoré's discussion indicates why the judgments might be breaking this way: the explosion is an abnormal event subsequent to Jones doing what they did. But the spike was "an existing state of affairs rather than an intervening event". The thought driving all this is presumably that the spike didn't have to do anything subsequent to the push – *it just sat there*, whereas the bomb had to do something subsequent to pushing Smith into the car, namely explode. But we find it very

difficult to think that this kind of folk distinction can stand up to much scrutiny. Suppose Jones had pushed Smith into a car that happened to be full of lethal gas (again with no lethal intent). We can easily imagine a significant portion of ordinary people assimilating the gas to the “existing state of affairs” category and judging Jones to have killed Smith. But there is a very natural sense in which the gas had to do something subsequent to the pushing in order for Smith to die – it had to enter Smith’s lungs and cause chemical burns etc. Another example: suppose Jones takes Smith to Lichtenstein and there is a large asteroid that, unbeknownst to Smith, is already on a collision course with Lichtenstein. Again, many will demur from saying that Jones killed Smith when the asteroid kills Smith (along with the rest of the population of Lichtenstein). But to think that the spike just sat there and didn’t have to do anything while the asteroid had to do something does not seem very respectable, even if folk judgments are driven by initial takes on the cases along those lines.

Hart and Honoré are themselves certainly live to the possibility that the key folk ideas about causation that drive doctrines about causation in the law will not look good once fully understood:

It may, of course, well be that when we thoroughly understand the common-sense notions of causation we should no longer wish our thought on any

matters, let alone legal judgments of responsibility, to be dominated by them: we may think that they are vague, crude, or anthropomorphic, or all of these: they may be ‘the metaphysics of the Stone Age’ which should be replaced by modern notions of probability or ‘risk’ (1985, 1-2).

In our view, such concerns are far from misplaced.

Let us briefly address two objections to the pessimistic lines of argument in this section.

First, one might worry that the preceding reflections suggest that there is something wrong with hoping not to kill people (as opposed to hoping to not be a factor in a death). Even if we are certain, or almost certain, that we will be a factor in some death, is it not still reasonable to desire/hope and so on that we will not be killers? If the preceding reflections make such hopes and desires misplaced, this would certainly be jarring. In response, note that our central normative claim was that there is no difference in badness between killing and being an abnormal event or accelerated process factor when things are otherwise *epistemically and motivationally symmetric*. But that does not mean that the news that someone kills tells us nothing about their epistemic and motivational situation. It is very easy to be a factor in someone or other’s death and almost inevitable that each of us will

be. Conditional on the news that we will be a factor, that doesn't do much to increase our expectation that we will be bad people in the future. But conditional on the news that we will kill, matters are different. A factor is much more likely to count as a killing if the relevant death is proximate to the action. Close the temporal gap between the action and the death and there is likely to be less reliance on abnormal intervening events. Make the gap great, and it is more likely that abnormal intervening events play a role. Of course, this isn't always the case (imagine we plant a bomb set to go off in a year). But killings *tend* to be much more proximate to the relevant death than that. Further, conditional on a death being proximate to our actions, it is more likely that it is intended or at least foreseen – and in those latter scenarios our moral status tends to be compromised. And let us not forget our first and third themes: for those reasons as well, the conditional probability of our being bad conditional on killing is higher than conditional on being a mere factor but not killing. Compatible with everything we have said, it is reasonable to hope that we don't kill.

Second, one might fall back on the idea that even if the killing/not killing distinction does not have moral significance per se, there is a distinction in the vicinity that does have moral purchase, namely the distinction between causing someone to die on the one hand, and accelerating their death without causing them to die on the other. Now in the context of discussions of killing and letting die

some have suggested that we have a basic aversion to causing death and that this underlies our moral sensibilities.³⁰ And one might be concerned that the above reflections in effect try to show too much. Even if ‘cause to die’ doesn’t quite line up with ‘kill’ (as the Katz vignette, discussed earlier, would suggest), certain of the cases above are plausibly asymmetric with regard to ‘cause’. After all, it is much more natural to say that one caused Smith’s death if he dies at 1.55 pm from a second cannonball than if he died from one’s accelerating a cannonball that had been initially released by someone else. Of course there are those like Lewis (1973) that will, on reflection, insist on a causal symmetry between those cases.³¹ But what if one opts for a view that comports with our ordinary disposition to ascribe causality in one case and not the other?

We certainly don’t want to be too dismissive to those who insist that causality is of fundamental moral significance. And that is part of why our conclusions are offered in a somewhat tentative spirit. But we do think it worth pressing whether we really do feel a difference between, say, the lever and “second cannonball” versions of accelerated death, where both involve a positive

³⁰ See for example, McMahan (1993), 277-279, where he argues that we care about the “form and degree of an agent’s causal responsibility”.

³¹ See fn. 23.

intervention by the agent, and so where there is no act/omission contrast to muddy the waters? Think back to the risk versions of the cases for example. Do we really expect good agents to take on a greater risk of Smith's dying on the grounds that this is compensated for by a lesser risk of the agent's counting as causing the death (as opposed to merely accelerating it)? It seems to us quite odd to countenance any increased risk of death on these grounds: If the lever is more likely to end in death than the second cannonball, that would seem to settle the matter, assuming the relevant symmetries in time of death. Cases like this, with a pair of positive interventions to choose between but some asymmetry in how folk would classify the outcomes using causal language, may indicate that causal asymmetries do not quite carry the moral weight that the objector anticipates. Meanwhile, if the causal truths are symmetric (as Lewis 1973 would suggest), the objection cannot get going in the first place.

It may also be worth mentioning that a theme from the reply to the first objection may also be relevant here. Suppose, in line with the preceding discussion, that causation has no moral import per se when it divides acceleration cases that are otherwise symmetric. It does not straightforwardly follow that one has no reason to hope that one doesn't cause someone's death. Supposing only a subspecies of cases of being a factor in a death are cases of causing, it may, depending on the details, be more likely, conditional on causing one's death, that

one knows that one is doing it, or that one is at least negligent or reckless in some way, than conditional on being a mere factor without causing. And for that reason, there be nothing wrong with hoping that one doesn't cause someone to die even if one is sure that one will be a factor in someone death.

Concluding Remarks

Supposing you agree with us, a range of further interesting questions arise, of which we mention two.

First, how does this bear on absolutist ideas in ethics? If no moral distinction is instituted between killing and being a mere factor (at least when it comes to accelerated process and abnormal event cases), the absolutist will need to extend their prohibitions to factor-theoretic analogues. Certain kinds of packages of this sort seem remarkably unappealing. Suppose one accelerated someone's death by a second by giving them a delicious but buttery meal that slightly accelerates the ongoing blockage of the arteries that is already close to a fatal conclusion. That does not seem like something we should ban outright, even in cases where one knows one is doing it. But there is a kind of absolute prohibition that does make eminent sense to us. Valuing killing someone as an end itself seems despicable.

And so does valuing accelerating someone's death as an end in itself. An absolute moral prohibition on certain kinds of preference structures seems eminently defensible with respect to the acceleration of death, even if the distinction between killing and accelerating death without killing (because of abnormal case and accelerated process phenomena) has no general moral import.

Second, what should we make of current legal practice in the light of the above? If the killing/mere factor distinction is ethically unimportant, can one nevertheless, as a practical matter, justify giving it a prominent role to play in tort and criminal law? This is a large and important question. But it is a question for another time.

Frederick Adams (1986). "Intention and Intentional Actual: The Simple View,"
Mind and Language 1,4 (1986): 281-301

Jonathan Bennett. *The Act Itself*. (Oxford: Clarendon Press, 1995).

Billings, J. Andrew, "Double effect: a useful rule that alone cannot justify hastening death," *Journal of Medical Ethics*, 37, 7 (2011): 437-40.

Bratman, Michael, "The Two Faces of Intention," *The Philosophical Review*, 93, 3 (1984): 375-405

Norman L. Cantor, "On Hastening Death without Violating Legal and Moral Prohibitions," *Loyola University Chicago Law Journal*, 37, 2 (2006): 407-431.

Thomas A. Cavanaugh, "The ethics of death-hastening or death-causing palliative analgesic administration to the terminally ill," *Journal of Pain Symptom Management*, 12, 4 (1996): 248-254.

Julia Driver, "Attributions of causation and moral responsibility," in W. Sinnott-Armstrong (Ed.), *Moral psychology, Vol. 2. The cognitive science of morality: Intuition and diversity*, (MIT Press, 2008: 423-439).

Philippa Foot, "The Problem of Abortion and the Doctrine of Double Effect," *Oxford Review*, 5 (1967): 1-5.

Philippa Foot, “Killing and Letting Die,” in Jay L. Garfield and Patricia Hennessey (eds.) *Abortion: Moral and Legal Perspectives* (Amherst: University of Amherst Press), 177-183.

J. Dmitri Gallow, “A Model Invariant Theory of Causation,” *The Philosophical Review* 130, 1 (2021): 45-96

Ned Hall, “Structural Equations and Causation”, *Philosophical Studies*, 132, (2007): 109–136

- Joseph Halpern, Joseph 2008 “Defaults and Normality in Causal Structures” in G. Brewka and J. Lang (eds.), *Principles of Knowledge Representation and Reasoning* (Palo Alto: AAAI Press, 2008): 198–208.

Herbert L.A. Hart and Anthony Honoré, *Causation in the Law, Second edition.* (Oxford University Press, 1985. First published 1959).

Christopher Hitchcock & Joshua Knobe, “Cause and Norm,” *Journal of Philosophy* 11, (2009): 587-612.

Thomas Icard, Jonathan Kominsky, and Joshua Knobe, J, “Normality and Actual Causal Strength,” *Cognition* 161: 80-93.

Daniel Kahneman and Amos Tversky, “The simulation heuristic,” in Kahneman, Slovic, and Tversky (eds.), *Judgment under uncertainty: Heuristics and biases* Cambridge, UK: Cambridge University Press, 1982): 201-210.

Jerrold Katz (1970) “Interpretive Semantics vs Generative Semantics,” *Foundations of Language* 6,2 (1970): 220-259.

Joshua Knobe and Benjamin Fraser, “Causal judgment and moral judgment: Two experiments,” in W. Sinnott-Armstrong, *Moral psychology, Vol. 2. The cognitive science of morality: Intuition and diversity* (MIT Press, 2008): 441–447.

Joshua Knobe, “Folk Judgments of Causation,” *Studies in the History and Philosophy of Science* 40,2 (2009): 238-242.

David Lewis, “Causation,” *Journal of Philosophy* 70,17 (1973): 556-567.

David Lewis, "Postscripts to Causation," in *Philosophical Papers Volume II*, Oxford University Press (1986).

Jeff McMahan, "Killing, Letting Die and Withdrawing Aid," *Ethics* 103,2 (1993): 250-279.

Neeleman and Van de Koot (2012) 'The Linguistic Expression of Causation,' in Everaert, Siloni and Marelj, eds., *The Theta System: Argument Structure at The Interface* (Oxford University Press, 2012): 20-51

Hugh McCann, "Settled Objectives and Objective Constraints," *American Philosophical Quarterly* 28,1 (1991) 25-36

Fabienne Martin and Florian Schafer, Florian, 2014 'Causation at the Syntax-Semantics Interface,' in Copley and Martin (eds.), *Causation in Grammatical Structures* (Oxford University Press, 2014): 209-244

Warren S. Quinn, 1989, "Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing," *Philosophical Review*, 98,3 (1989): 287–312.

Neal Roese, “Counterfactual thinking,” *Psychological Bulletin* 121 (1997): 133-148.

Richard Thomason, “Formal Semantics for Causal Constructions,” in Copley and Martin eds., *Causation in Grammatical Structures* (Oxford University Press 2014): 58-75

Judith Jarvis Thomson, “Killing, Letting Die and the Trolley Problem,” in W. Parent (ed.), *Rights, Restitution, and Risk: Essays in Moral Theory*, (Cambridge, MA: Harvard University Press, 1986): 78–93.