

MAKING Metaphysics

Thomas Byrne*

[Forthcoming in *Philosophers' Imprint*.]

We can cause windows to break and we can *break* windows; we can cause villages to flood and we can *flood* villages; and we can cause chocolate to melt and we can *melt* chocolate. Each time these can come apart: if, for example, A merely instructs B to break the window, then A causes the window to break without breaking it herself. Each instance of A breaking/flooding/melting/burning/killing/etc. something, is an instance of what I call *MAKING*. I argue that *MAKING* is an independent, theoretically important notion—akin but irreducible to causing—and metaphysicians should pay attention to it.

1. When the window breaks, it might be the case that

Annie caused the window to break,

and it might be the case that

Annie broke the window.

Often the two will go hand-in-hand—for instance, when Annie throws a rock at the window—but they needn't. If Annie only instructed Ben to throw a rock at the window, then while it wouldn't be the case that Annie broke the window, it would be the case that she caused the window to break (by causing Ben to break it).¹

Similarly, when the house floods, it might be that Annie caused the house to flood and it might be that Annie flooded the house. These can also come apart: floodwaters are heading towards the open doorway of Annie's house; Annie closes her door and so the waters instead enter Ben's house, one door down. Annie doesn't flood Ben's house, even though she causes it to flood.

And this is nothing special about Annie, since the same goes more generally. When the chocolate melts, it might be that the fire caused the chocolate to melt and it might be that the fire melted the chocolate. These can come apart, too: if the fire cuts power to the air-conditioning, then the fire

* Department of Linguistics & Philosophy, Massachusetts Institute of Technology.

¹ Not only does *Annie* cause the window to break₁ (an instance of so-called *thing* causation) but so too does *Annie's instructing Ben* cause the window to break₁ (*event* causation). I focus on the former, but, in doing so, I make no claim about which kind is fundamental. I also focus on whether A *causes* X to V₁ as opposed to whether A is a *cause* of X's V₁-ing (although, this distinction crops up briefly in §5); again, this is not a substantive commitment on my part.

doesn't melt the chocolate, even though it causes the chocolate to melt (by causing the summer air to melt it).

There's a pattern here. *Break*, *flood* and *melt* are all verbs that can be used both *intransitively* and *transitively*. A verb *V* is used intransitively in a given sentence when it doesn't allow a grammatical object (e.g. "the window broke"), while it's transitive when it does (whether explicitly, e.g. "Annie breaks the window," or implicitly, e.g. "stop stealing [cars]!"). If we mark these intransitive and transitive uses with the subscripts "I" and "T," then we can capture the pattern as follows.² For thing *X* and verb *V*, when

X V_I's (e.g. the window breaks),

it might be the case, for actor *A* (where *actor* ranges over both agents and things), that

(i) *A* caused *X* to *V_I* (e.g. Annie caused the window to break)

and it might be the case that

(ii) *A V_T*'d *X* (e.g. Annie broke the window).

Philosophers have spent much time thinking about instances of (i) and the language is familiar: causes, causing, etc.

Not so with instances of (ii). For lack of a better term, I will say that, when *X V_I*'s:

A MAKES X V_I just in case *A V_T*'s *X*.

Annie MAKES the window break_I just in case Annie breaks_T the window; the fire MADE the chocolate melt_I just in case the fire melted_T the chocolate; and so on. I will also talk of the things that MAKE—*viz.* MAKERS—and of the activity that MAKERS partake in, MAKING.

To what end? Well, we might, across various domains of inquiry, think that it's important to know what causes what to happen; or we might think it's important to know what MAKES what happen. Here are two examples from the normative domain.

When the ethicist is seeing to settle whether Annie acted wrongly, she might start by seeing what the various *consequences* of her action were (where *C* is a consequence of *E* just in case *C* counterfactually depends upon *E*). But then she might think it's important to additionally know, of

² I say, for instance, that the verb *melt* can be used both transitively and intransitively: *melt_T* and *melt_I*. I might instead have said that there are two verbs, *melt_T* and *melt_I*, that share a lexeme (as linguists call it). Nothing rests on this choice, except easier expression. Also note that while *melt_T* and *melt_I* are homographic and homophonic—looking and sounding the same—this won't always be the case. For example, when Annie kills_T Ben, it's not that *Ben kills_I*, but that *Ben dies_I*. Irregular verbs like these are common and, for my purposes, uninteresting.

Linguists call verbs that behave in this way—that is, those verbs where *V_T*'s object (e.g. the window) is *V_I*'s subject—*ergative*.

those consequences, which ones Annie caused. Alternatively, she might think it's important to additionally know, of those consequences, which ones Annie MADE. Since causing and MAKING come apart, the upshot will be moral theories with conflicting verdicts.

Similarly, to convict A of B's homicide, the lawyer must prove both that A possesses the requisite mental states for that crime and that B's death is a consequence of A's actions. But that is not sufficient for a just conviction (imagine: A poisons B's gin, B notices, heads to the shop to buy a replacement bottle and is hit by a car while crossing the road) and so there must be a further component of the crime. That further component might be that A caused B's death or it might be that A MADE B die—that A killed B. The upshot may be different people going to jail.³

For now, it suffices to say that it's an open, empirical question whether ethicists and lawyers in fact think that what matters is who caused what or whether they think that what matters is who MADE what. (Although, for what it's worth, I don't think it's a coincidence that ethicists have spent much time with principles like "it's impermissible to *kill* one to save five," yet no time with causing.⁴ Nor does it count for nothing that the gin example, above, is itself a counterexample to the idea that the lawyer's further component is causing, since A does cause B's death, yet nonetheless should not be convicted of B's homicide. More on these points later.)

I take that same open, empirical question to apply to other domains too. Do oncologists want to discover just what things cause cells to mutate or do they want to discover just what things mutate cells? Does the crash investigator want to know what caused the boat to sink or does she want to know what sunk the boat? This time, the upshots will be different programmes of inquiry.

Since these are open questions—certainly open for the metaphysician—we might expect metaphysicians to have thought about both causing and MAKING. But they haven't. They have spent endless time thinking about causing—"what's its nature?", "is it transitive?", "can omissions cause?", and so on—but almost no time thinking about MAKING.⁵ Why is that? As far as I know, no one has ever defended the disparity (nor, even, thought that it needed defending), but, asking around, the rationale seems to look like this:

³ Two points to avoid confusion. Firstly, it is the mental state required that distinguishes the various degrees of homicide. Roughly: intention and premeditation for first degree; intention for second degree; recklessness for involuntary manslaughter; and so on. Secondly, the law has brute provisos that deal with cases of preemption and overdetermination *vis-a-vis* the consequence requirement.

⁴ Carolina Sartorio (2010) is one exception. She considers whether causation might play an interesting role in ethics. Her conclusion is largely negative.

⁵ As far as I know, the only time that MAKING has been discussed as a serious peer of causing was in a talk by Brian Weatherson (2007). Unfortunately, the talk didn't progress to a paper. All other talk of MAKING has been oblique: concerning its grammar (see §2), concerning the specific instance of MAKING X die_i (see §10), concerning MAKING as a specific type of causing (see G.E.M. Anscombe (1971)), and so on.

“MAKING reduces to causing so even if MAKING features in various theories in various domains it only does so in virtue of the fact that it reduces to causing.”

And, indeed, when metaphysicians have paid attention to MAKING it has been in search of that reduction. In §5, I show that every proposed reduction of MAKING to causing fails. The rationale continues:

“Even though we haven’t found the correct reduction, we nonetheless have good reason to think that there must be such a reduction.”

In §6, I show that we do not have good reason for thinking that there must be such a reduction. And it continues still:

“There are in-principle objections to MAKING being an independent, interesting notion.”

In §7, I consider those objections and argue that none of them is compelling.

I conclude that if theorists, from various domains, are interested in what MAKES what, then metaphysicians should pay attention to MAKING. In §§9-11, I focus on the ethical and legal domains and argue that both ethicists and lawyers are interested in what MAKES what; and so metaphysicians should pay attention to MAKING. I end, in §12, by discussing just how that attention should be paid.

2. First, some more about MAKING, itself.

Above, I said that, when $X V_I$'s

A MAKES $X V_I$ (e.g. Annie MAKES the window break)

just in case

A V_T 's X (e.g. Annie breaks the window).

While this schema neatly captures the instances of MAKING we saw above, it might appear to not capture all possible instances. For starters, it might seem to not capture those instances of A MAKING $X V_I$ for verbs that have no transitive use in English. *Smile*, *laugh*, *sneeze*, and *jump* are all such verbs and for each of them, “A V_T 'd X ” is ungrammatical (e.g. #A smiled X).

Thankfully, that those instances of “A V_T 'd X ” are ungrammatical is no barrier to our determining whether they are true. Judith Thomson considers *smile*:

...suppose you sign on as apprentice to the chief photographer. As he is about to photograph a batch of children for their school yearbooks, he notices that a child in the back row looks gloomy, this being undesirable in such photographs. “Smile that child up there in the back row,” he says to you; “I want to see if you can.” You

know perfectly well that he does not want to see merely whether you can cause the child to smile—as, e.g., by paying a more experienced apprentice to cause the child to smile. You know he wants *you* to smile the child (1977: 129).

Accordingly, I see no problem with coining the transitive use of *smile* and stating that, when *X* smiles_I

A MAKES *X* smile_I just in case *A* smiles_T *X*.

And, in turn, doing the same for *laugh* (e.g. “Annie *laughed* the baby”), for *sneeze* (e.g. “the nurse *sneezes* the patient”) and for those other verbs that lack transitive uses.⁶

Just as we can coin transitive verbs from intransitive ones, so too can we coin intransitive verbs from certain transitive ones. For example, English recently started using *microwave* as a transitive verb—“Annie *microwaved*_T the potato”—and we might, in turn, coin the intransitive use without any loss of sense: “the potato *microwaved*_I in three minutes.” That means that we can say that, when the potato *microwaves*_I

A MAKES the potato *microwave*_I just in case *A* *microwaves*_T the potato.

The same goes for other verbs too. In particular, it's worth mentioning—since it's a central concept in ethics—that the same goes for *harm*. Like *microwave*, *harm* has only a transitive use in English (e.g. Annie *harms*_T Ben), yet we can happily coin its intransitive use and say that, when *X* *harms*_I

A MAKES *X* *harm*_I just in case *A* *harms*_T *X*.

(And, indeed, *harm* was historically used intransitively: e.g. “The men is fresh, too, and won't *harm*_I for a bit of exercise.”⁷)

But what goes for *microwave* and *harm* doesn't go for all those verbs that have no intransitive use in English. That is, for some transitive verbs, not only is “*X* *V*_I's” ungrammatical, but it's also nonsensical. For example, it might be the case that Annie *steals*_T the money or that Annie *ponders*_T the universe, yet “the money *steals*_I” and “the universe *ponders*_I” are nonsense. Since “Annie MAKES the money *steal*_I” and “Annie MAKES the universe *ponder*_I” are similarly nonsensical (what *could* they mean?), verbs like this are not counterexamples, but cautionary tales: we shouldn't coin with abandon.

Lastly, just as we can happily coin transitive verbs from certain intransitive ones—and *vice versa*—in order to capture certain instances of MAKING, so too can we coin verbs from certain adjectives in order to capture others. Repurposing a different example of Thomson's (*ibid.*), consider *blue*:

⁶ English also provides its own workaround for these instances of MAKING: *A makes*_T *X* *V*_I. Reconsider Thomson's story, replacing “smile the child” with “make the child smile” and notice how her conclusion still holds: you might make a child smile by telling them a joke, but not by having someone else tell them that joke. This workaround inspired my terminology.

⁷ “harm, v.” *OED Online*. Oxford University Press, March 2020.

If your customer says, “I like the work you do better than the work your assistants do; so would *you* blue the inside of this bluebottle they just tattooed on my arm?” you would know perfectly well that he does not merely mean that you are to cause the inside of the bluebottle to be blue—as, e.g., by getting one of your assistants to blue it.

And so we might say that, when X blues_I

A MAKES X blue_I just in case A blues_T X

(just as, when X reddens_I, A MAKES X redden_I just in case A reddens_T X).⁸ The same goes for other adjectives, too: the script doctor might *funny*_T the play—and, thus, MAKE the play funny_I—by adjusting the dialogue, but not by having someone else adjust it; and so on.

This is all to say that while there are certain linguistic barriers to our accepting the schema, above, those barriers do not run deep. We should ignore them.

Are there other barriers to our accepting the schema? Perhaps there are brute counterexamples to it, ones that aren’t confined by the caprice of our language. For example, I said earlier that the fire doesn’t melt the chocolate when it merely cuts power to the air-conditioning. But can’t we simply stipulate that the fire nonetheless *schmelts*_T the chocolate (where, roughly, A schmelts_T X just in case A cuts power to whatever was preventing X from melting); and, having done so, doesn’t that make it the case that the fire MAKES the chocolate schmelts_I? We sure can and it sure might, but this is no counterexample. After all, the relevant instance of the schema is that, when X schmelts_I

A MAKES X schmelts_I just in case A schmelts_T X,

and the case accords with it.

Ultimately, we should accept that, when X V_I’s

A MAKES X V_I just in case A V_T’s X.

3. By now, certain patterns, as to when this might MAKE that, have started emerging. For instance, if A causes X to V_I merely by causing B to (willingly) V_T X, then that is sufficient for it to be the case that A does not MAKE X V_I. But these patterns do not amount to a theory of MAKING. Moreover, I provide no theory in this paper.

Without such a theory you might think you are ill-equipped to determine exactly when this MAKES that. Not so. Since A V_T’s X just in case A MAKES X V_I, we can determine whether A MAKES X V_I by

⁸ As evidence for how happily we can coin such verbs, consider: “[it] has a rose to brownish red or reddish brown peridium, yellow context that blues_I when exposed”; and “... a redhead who had blue_T part of her hair.” Both from: “blue, v.1.” *OED Online*. Oxford University Press, March 2020.

determining whether A V_T's X. And, it turns out, we are all experts at determining whether A V_T's X.

We can flex that expertise by running through some subtle contrast cases:

LIGHTNING: Annie beats Ben and leaves him immobile. A freak thunderstorm rolls in and Ben is hit by a lightning bolt, which stops his heart.

As before, but this time there is no thunderstorm. Instead, the sun sets and Ben dies of exposure overnight.

Does Annie kill Ben?

FLOOD: heavy rain has raised the river to dangerous levels. To protect her house, Annie blocks the doorway with sandbags. The river soon floods and the floodwaters, unable to enter Annie's house, instead enter Ben's, one door down.

As before, but this time Annie is out of town, is unable to block her doorway and so her house floods. Later, she pumps the waters out of her flooded basement. That water flows down the street and into Ben's house.

Does Annie flood Ben's house?

JUMP: a runaway trolley with enough momentum to kill a single individual is heading towards Annie. Annie jumps off the track and the trolley instead hits and kills Ben who is tied to the track behind her.

As before, but this time it's not Annie on the tracks, but a barrier. Annie removes the barrier and the trolley instead hits and kills Ben who is tied to the tracks behind it.

Does Annie kill Ben?

DYE: Annie mixes some peroxide into Ben's pomade. Wily Ben realises and so he uses a different bottle instead. Unfortunately, the contents of the second bottle have (independently) fermented into a peroxide-like substance which turns Ben's hair white.

As before, but this time Ben doesn't realise at all. He unwittingly uses the pomade, his hair turns white.

Does Annie turn Ben's hair white?

WIND: the howling wind agitates the bear who jumps up and down on the tree branch until it snaps.

As before, but this time there is no bear. The wind buffets the branch until it breaks and falls to the ground.

Does the wind break the branch?

I trust that you answered as follows: Annie doesn't kill Ben when he's struck by the lightning, but does when he dies of exposure; Annie doesn't flood Ben's house when she blocks her doorway, but does when she pumps the water; Annie doesn't kill Ben when she jumps off the track, but does when she moves the barrier; Annie doesn't turn Ben's hair white when he instead uses the second, fermented bottle, but does when he unwittingly uses the bottle she tampered with; and the wind doesn't break the branch when it agitates the bear, but it does when it buffets the branch. Congratulations! You're an expert at determining whether A V_T 's X and, in turn, whether A MAKES X V_I .

4. We might be experts at determining, for various verbs V, whether A V_T 's X and, in turn, whether A MAKES X V_I in a given instance, but this alone doesn't license my claim that each of these instances are tokens of the metaphysical type MAKING. There might be no such (natural) type: Annie might flood the village and Ben might break the window, yet it's possible that those actions share little more than the grammatical form of their description. What should we make of such MAKING-scepticism? Not too much.

As already mentioned, patterns have started emerging across the cases. For example, whatever the verb V, A's merely causing B to (willingly) V_T X is sufficient for it to be the case that A does not MAKE X V_I . Why should that be? An explanation is required.

Just as there are patterns emerging across the cases we've looked at, there are also patterns to be found *within* cases. Recall FLOOD: waters head towards Annie's house, Annie blocks her doorway and the waters instead enter Ben's house; yet Annie does not flood Ben's house. Yet this is nothing special about waters nor flooding_T. Suppose instead that the paint factory explodes and a wave of blue paint heads towards Annie's house, Annie blocks her doorway and the paint instead enters Ben's house; Annie doesn't *blue* Ben's house. Or suppose that a plague of locusts head towards Annie's greenhouse, Annie closes the door and the locusts instead enter Ben's greenhouse; Annie doesn't *destroy* Ben's crop.

Similar patterns are found within the other cases, too. These patterns all require explanation. I submit that the most natural explanation is that MAKING is a unified metaphysical type. It falls to the MAKING-sceptic to provide their own, better explanation; I've not seen one.

5. When metaphysicians have said something about MAKING, it has tended to be that it reduces to causing in one way or other. If they're correct, then their focus on causing is justified. Here I show that each proposed reduction of MAKING to causing fails. The final two proposals have more going for them, so I spend more time with them.⁹

Attempted reduction 1: MAKING is *direct* causing. A natural idea is that while many things might cause X to V₁, only certain things will do so *directly*, where A₁ directly causes X to V₁ just in case it causes X to V₁, but doesn't do so by causing some A₂ to cause X to V₁. It then says that A MAKES X V₁ just in case A directly causes X to V₁. For instance, suppose that Annie instructs Ben to throw the rock at the window, Ben throws the rock, the window breaks. Friends of this reduction would say that since Annie causes the window to break only by causing Ben to cause the window to break, Ben MAKES the window break, but Annie doesn't. But this reduction immediately fails since Ben too causes the window to break only by causing something else—namely, the rock—to cause the window to break. So it, in fact, falsely returns that Ben doesn't break the window.

Attempted reduction 2: MAKING is *intended* (*iforeseen*) causation. This would say that A MAKES X V₁ just in case A intentionally (*iforeseeably*) causes X to V₁. Since inanimate objects can MAKE X V₁—boulders break windows, lightning burns trees, etc.—and I can (and have) broken things both unintentionally and unforeseeably, this is a non-starter.

Attempted reduction 3: MAKING is *morally loaded* causation. Shelly Kagan (1989) was struck by how we would never accuse a surgeon of killing her patient, even when the surgeon cuts into the patient's heart which causes the heart to stop, which causes the patient to die. As a result, Kagan suggested that whether A kills B is fixed, in part, by moral features of the situation. In particular, he suggested that the surgeon doesn't kill the patient because she acted permissibly. Mightn't we generalise this reduction (or some version of this idea) to MAKING in general? I don't think so. Firstly, whether I MAKE X V₁ cuts across every moral distinction: I can break the window either permissibly or impermissibly, either justifiably or unjustifiably, either with blame or without it, and so on. Secondly, we already saw in WIND, that wind might break a branch or it might merely cause that same branch to break, even though wind never does anything permissibly nor impermissibly (nor justifiably, etc.). Ultimately, this reduction is no good. (And what of Kagan's surgeon case? I suggest that what we *accuse* people of often differs from what they in fact do. After all, it would be natural for the surgeon herself to say "I feel awful, I killed a patient today.")

⁹ Much of this section's content has been said before by Thomson (1987) and, less so, by David Lewis (1986).

It's important not to confuse attempt 3 with a related idea. Namely, that it isn't that A's MAKING X V_I is *fixed* by certain moral facts but instead that it *fixes* certain moral facts. For instance, it could be that in order for to A infringe X's right against being V_I'd, A must MAKE X V_I. As it happens, I think this is correct (see §9), but for now it suffices to say that this is clearly not a reduction: entailment is not reduction.

Attempted reduction 4: MAKING is *causing*, as distinct from being *a* cause.¹⁰ Some authors—e.g. Peter Unger (1977) and Alex Kaiserman (2016)—take seriously the distinction between causing X to V_I and merely being a cause of X V_I-ing, where the mere causes of X V_I-ing are each, in Unger's terms, “a part of that which causes” X to V_I (184). For example, this view would have it that the speeding driver and the torrential rain are each causes of the shopfront breaking (when the car crashes through it) in virtue of their both being parts of what caused the shopfront to break—namely, the driver speeding in the torrential rain. If this is right and what causes X to V_I comes apart from the mere causes of X V_I-ing, then perhaps A MAKES X V_I just in case A causes X to V_I, after all.

Now, I have no idea whether this is the right way of thinking about causation, but I'm sure it isn't the right way of thinking about MAKING. After all, since it says that the speeding driver doesn't cause the shopfront to break, it would also say that the speeding driver doesn't break the shopfront, but clearly the speeding driver breaks the shopfront—he crashes straight through it!

Attempted reduction 5: MAKING is causing that is *insensitive* to other circumstances. This was David Lewis's idea. He said:

It may be that the causation depends on an exceptionally large and miscellaneous bundle of circumstances all being just right. If any little thing had been different, that cause would not have caused that effect. But sometimes causation is comparatively insensitive to small differences in the circumstances. When my strong recommendation causes ... deaths, that is comparatively sensitive causation—there are many differences that would have deflected the chain of events [and so I do not kill]. But if you shoot at your victim point-blank, only some very remarkable difference in circumstances would prevent his death [and so you do kill] (1986: 186).

And we might put the resulting reduction as follows: A MADE X V_I just in case A (sufficiently) insensitively caused X to V_I. Setting aside general worries—in particular, just what counts as sufficiently insensitively?—this reduction still returns the wrong verdict in FLOOD. Recall:

FLOOD: heavy rain has raised the level of the river to dangerous levels. To protect her home, Annie blocks the doorway with sandbags. The river soon floods and the floodwaters, unable to enter Annie's house instead enter Ben's, one door down.

¹⁰ Thanks to the reviewer for this journal for suggesting this as a possible reduction.

Annie insensitively causes Ben's house to flood—after all, once she blocks her doorway, it's determined that the waters will enter Ben's house—nonetheless, Annie doesn't flood Ben's house.

So insensitive causing isn't sufficient for MAKING nor, *a fortiori*, necessary and sufficient. We might wonder whether it's merely necessary. It isn't. Suppose that Robin fires an arrow through the crowds at Waterloo Station; it *just* misses every jostling passenger and breaks a window on the far side. Robin causes the window to break, but that he does so is sensitive to the precise speeds and trajectories of all those passengers that the arrow just misses. Nonetheless, Robin clearly breaks the window.

Attempted reduction 6: MAKING is causing *without intervening human acts*. Jennifer Hornsby (1980) says that A MAKES X V_I just in case A causes X to V_I and there exists no human that more proximately causes X to V_I. For instance, when Annie instructs Ben to throw the rock at the window, Ben more proximately causes the window to break and therefore Hornsby says that Annie doesn't MAKE the window break. This reduction returns the wrong verdict in both FLOOD and LIGHTNING. Recall:

LIGHTNING: Annie beats Ben and leaves him slowly dying of his injuries. A thunderstorm rolls in and Ben is hit by a lightning bolt.

Even though there exists no human that more proximately causes Ben to die, Annie does not kill Ben. Similarly, even though there exists no human that more proximately causes Ben's house to flood, Annie doesn't flood Ben's house.

So that there exists no human that more proximately causes X to V_I is not sufficient for MAKING, nor, *a fortiori*, necessary and sufficient. We might again wonder if it's merely necessary. It isn't: as we saw in DYE, Annie turns Ben's hair white even when it's Ben himself who applies the pomade (and, in doing so, more proximately causes his hair to turn white).¹¹

6. Each proposed reduction of MAKING to causing fails. Induction alone thus gives us pause to examine the assumption that things do reduce in that manner. So what might motivate that assumption? As far as I know, no one has ever motivated it, but here's how I imagine the argument would go: "that A MAKES X V_I entails that A causes X to V_I, but the converse does not hold. That is, MAKING entails causing, but causing doesn't entail MAKING. So MAKING reduces to causing."¹²

¹¹ Jonathan Bennett (1988) combines the previous two attempts and suggests that A makes X V_I just in case both A insensitively causes X to V_I and there is no more proximate human cause of X's V-ing. However, this conjunctive reduction still returns the wrong verdict in FLOOD.

¹² Between them, Jerry Fodor (1970) and John Morreall (1976) given four reasons why causing doesn't entail MAKING.

What should we make of this argument? Its sort has been criticised before by Timothy Williamson regarding knowledge's claimed reduction to true belief. I simply repeat his criticism here—namely, that the argument doesn't generalise and, therefore, should be rejected.

Something's being red entails that it's coloured, but not *vice versa*. The argument structure, above, would thus have it that we can reduce things as follows:

$$red = coloured + X.$$

But there is no reason to think that the above equation can be solved in a non-circular way (and its being solved in a circular way—where $X = red$ or *red's wavelengths*, for instance—is no evidence that *red* reduces to *coloured*). So that something's being red entails that it's coloured (but not *vice versa*) is no reason to think that being red reduces to being coloured and, in turn, that A MAKES X V_I entails that A causes X to V_I (but not *vice versa*) is no reason to think that MAKING reduces to causing (Williamson 2000: 3).

So we should be suspicious of that argument. But that's only one argument and perhaps there are others that could be given for the same conclusion. So best to change tack: here's an argument that MAKING doesn't reduce to causing.

For the sake of argument, suppose that MAKING does, in fact, reduce to causing. That is, suppose that our concept MAKING reduces to our concept causing, just as, for example, our concept *bachelor* reduces to our concept *unmarried*. Reductions like these come with an epistemic restriction: in the case of *bachelor*, it comes with the restriction that we cannot be more sure that X is a bachelor than we are that X is unmarried.¹³ The same goes for MAKING: if MAKING reduces to causing, then we can't be more sure that A MAKES X V_I than we are that A causes X to V_I.

And normally, we won't be. If Annie shoots Ben dead then it's clear both that Annie kills Ben and that Annie causes Ben to die. Similarly, if Annie leaves the hose running in the garden, it's clear both that Annie floods the lawn and that Annie causes the lawn to flood. However, that won't always be so. Consider:

POISON: Ben is set to go on a cruise on Monday. The day before, Annie poisons Ben, sending him to hospital instead. The cruise departs without him and sinks on Tuesday with all hands lost. The poison kills Ben the following week.¹⁴

¹³ There might be strange exceptions. For instance, I might tell my young nephew that George is a bachelor. My nephew might thereby come to know that George is a bachelor even though my nephew has no idea what it is to be a bachelor and, in turn, does not know that George is unmarried. I won't be looking at cases like this.

¹⁴ See Hart & Honoré (1959: 220 fn 1).

Annie kills Ben. Does she also cause Ben to die? On first glance, it's not at all obvious either way. (Note, the mere fact that she causes him to die *at a given time* is insufficient for it to be the case that she causes him to die. The paramedic's chest compressions might cause the patient to die at 3pm instead of 2pm, but the paramedic doesn't cause the patient to die.¹⁵)

What about on second glance? Isn't this just a case of, so-called, (*early*) *preemption*? Consider:

two assassins plan to kill Dictator. The first assassin will shoot and only if he misses will the second assassin shoot. The first shoots: it's a direct hit and Dictator dies.

The first assassin is said to preempt the second assassin since, by killing Dictator, he prevents the second assassin from doing so. Nonetheless, it's obvious that the first assassin causes Dictator to die. Since Annie's poisoning also prevents the cruise from causing Ben to die, why shouldn't we say of Annie what we say of the first assassin—namely, that she obviously causes Ben to die?

For starters, in poisoning Ben, Annie causes Ben to miss the cruise and, in turn, to live a week longer than he would otherwise have done. That fact alone sits ill with the claim that she causes him to die, and there is no corresponding fact *vis-a-vis* the first assassin. More importantly—and to repeat where I started—it just isn't obvious that Annie causes Ben to die, yet it is obvious that the first assassin causes Dictator to die.

Of course, we could convince ourselves that Annie does cause Ben to die; we could move from the clear verdict about the first assassin, via some theory of preemption or of causation, to that conclusion—and perhaps that is the right conclusion. But, regardless, that is a conclusion driven by theory and it doesn't undercut the fact that it's (much) clearer that Annie kills Ben than it is that she causes him to die. And that can't be if MAKING reduces to causing.

The same goes for other cases too. Consider:

BUSHFIRE: a bushfire is going to destroy both Annie's house and Ben's house. To protect her house, Annie opens the dam. Its waters quell the fire but also wash away Ben's house.¹⁶

Annie destroys Ben's house. Does she also cause Ben's house to be destroyed? As before, first glances are not definitive. (And, as before, the mere fact that Annie causes the house to be destroyed in a certain, watery way is not sufficient for it to be the case that she causes the house to be destroyed. I

¹⁵ See Carolina Sartorio (2006).

¹⁶ From Hart & Honoré (1959: 219).

can cause the firework to explode quietly—by throwing a blanket atop it—without causing the firework to explode.)

Of course, we could again move from our conclusion about the first assassin, via some theory, to the conclusion that Annie does cause Ben’s house to be destroyed. But, again, that reasoning doesn’t undercut the fact it’s (much) clearer that Annie destroys the house than it is that she causes it to be destroyed. And, again, that can’t be if MAKING reduces to causing.

You might think I’m cheating. I have, after all, deliberately picked cases with complex causal structures and then used that complexity to elicit these unclear causal judgements from you (“so of course our causal judgements will be unclear,” goes the complaint). But the point isn’t that those judgements are unclear, but that they are *less* clear than the corresponding MAKING judgements. From that perspective, this complaint makes no sense—after all, if MAKING *does* reduce to causing then that complexity should infect our MAKING judgements (at least) as much as our causing judgements. Yet, it doesn’t.

Ultimately, we should be very suspicious of the assumption that MAKING reduces to causing: no wonder the attempts of the previous section failed.

7. Before concluding that MAKING is an independent, interesting notion—just like causing—I consider three possible objections to that conclusion.

Objection 1: “if MAKING truly were an independent notion then it would, like causing, have a name. But it doesn’t: you just made one up.” Suppose, as this objection assumes, that the words we have somehow reflect reality (more realistically: the *concept* words we have somehow reflect reality). And it is true that “causing,” “cause,” and “causer” are all words in our language, while “MAKING,” “MAKE,” and “MAKER” are not.

But when we move from the generic words to specific ones, the opposite is true. There is no word for someone who causes people to die, but someone who MAKES people die is a *killer*; and there is no word for the action that people who cause people to die do, yet killers *kill*. Similarly, there is no word for the action of causing X to flood, but those that MAKE X flood, *flood*_T X, and we might call them *flooders*. Indeed, our language is littered with words for different instances of MAKING and for those who MAKE: *break*_T, *smelt*_T, *bake*_T, *cook*_T, *build*_T, *demolish*_T, *thief*, *builder*, *baker*, *smelter*, etc.

So while only causing has generic words, only MAKING has specific ones—thousands of them. At worst, that’s a wash for MAKING.

Objection 2: “MAKING is too messy and disorderly. What could possibly explain why Annie doesn’t flood Ben’s house in FLOOD but does in its variation? And what could possibly explain why Annie doesn’t kill Ben in LIGHTNING but does in its variation? And so on.” When we look back at all the cases, above, we might conclude that there is no order as to when A does and doesn’t MAKE X V_I. That smacks of defeatism at this early stage, but perhaps it isn’t. But even if it isn’t, disorder hardly leaves MAKING worse off than causing. Ted Sider describes things nicely:

Causation is a particularly unsavoury fundamental posit—at least if the posit is intended to closely match our ordinary concept of causation. It takes only a glance at the recent literature on causation to appreciate how arbitrary and baroque our ordinary concept of causation is (2011: 15-16).

And since this disorder has not dampened our interest in causing, nor should it dampen our interest in MAKING.

Objection 3: “whether A MAKES X V_I is affected by the mental states of agents. We saw this with DYE and its variations: Annie turns Ben’s hair white when he unknowingly uses the peroxide, but not when he does so knowingly. On the other hand, whether A causes X to V_I is a purely physical matter.”

This might all be true, but it isn’t obvious how it’s objectionable. Perhaps as follows: that MAKING is affected by the agents’ mental states renders MAKING unsuited for playing the roles in theories that might otherwise be played by causing. I suggest one sure-fire way of determining whether MAKING *is* suited for playing a certain role is by looking at whether it does, in fact, play that role (see §§9,11); and if it does, then this objection falls flat.

8. We know what MAKING is, we know it when we see it, we have no reason to think it reduces to causing, and there is no in-principle objection to its being an independent, interesting notion. But that might all be besides the point if theorists are not interested in what MAKES what.

What is it for them to be interested? Roughly, if theorists in a given domain draw a distinction that aligns with, for the relevant actors A, whether A MAKES X V_I, then those theorists are interested in what MAKES what. If, for instance, what distinguishes those cases of intentionally bringing about the death of another that *are* homicide, from those that are not homicide, is whether the defendant killed the victim, then lawyers are interested in what MAKES what.

And if lawyers or ethicists or oncologists or ornithologists or whoever *are* interested in what MAKES what, then it falls to metaphysicians to pay attention to MAKING: the more theorists, the more

that metaphysicians should pay attention (just how that attention should be paid, I turn to later). That is the conclusion of the preceding seven sections.

But are theorists interested in what MAKES what? As I said, this is an open, empirical matter that can only be settled one theory at a time. Obviously, there are too many theories to address them all here; I shall focus my attention on ethics and the law.

9. There are many ethical theories and their corresponding theorists are interested in different things. The *consequentialist*, for instance, thinks that what fixes whether or not a given action is permissible is whether its consequences were at least as good as any of the other actions available to the agent: if they were, then the action was permissible. Accordingly, consequentialists are interested in consequences.

Deontologists are also interested in consequences: that the child drowned, and wouldn't have done so otherwise, explains why it was impermissible for Annie to ignore the drowning child. But, in their quest to capture our common-sense morality, deontologists hold that there are additional constraints on how we are permitted to act, ones that go beyond what is a consequence of what.

Here is Philippa Foot's (2002) classic illustration of such a constraint:

RESCUE-I: the tide is rising and five people are trapped in one cove, while a single person, Ben, is trapped in another. Annie can only make it to one of the two coves in time to rescue the people trapped inside. Annie rescues the five; Ben drowns.

RESCUE-II: the tide is rising and five people are trapped in a cove. To reach the five in time, Annie would have to run down Ben who is trapped in the access road. Annie rescues the five; Ben is run down.

In both cases, Ben's death and the survival of the five are consequences of Annie's action, but only in RESCUE-II does Annie act impermissibly. And Foot's explanation of this is as follows: only in RESCUE-II does Annie kill Ben and it's impermissible to kill one in order to save the lives of five others.

Of course, that principle is just a specific instance of a more general one, and so Foot would also endorse:

[IMP] other things equal, it's impermissible for A to V_T B (/B's X) in order to prevent five others (/other Xs) from V_I -ing.

(E.g. other things equal, it's impermissible for Annie to flood Ben's house in order to prevent five other houses from flooding.)¹⁷

If Foot's right then MAKING cuts an important ethical distinction and, as I've put it, deontologists are interested in what MAKES what. Is she right? Her cases don't settle things since they don't differentiate between killing and, amongst other things, causing to die. But we can test [IMP] by looking at other cases and related principles—ones that better isolate MAKING.

For instance, we know that A's causing B to cause X to V_I is insufficient for A's MAKING X V_I. With that in mind:

RESCUE-III: the tide is rising and five people are trapped in a cove. Since it's dark, she will have to use her car's headlights in order to reach the five, the lights of which will enable Villain to pick an innocent target and shoot them dead. Annie rescues the five, Villain sees Ben by the lights of Annie's car and shoots him.

It's permissible for Annie to save the five and [IMP] agrees. On the other hand, the causal analogue of [IMP]—that it's impermissible to cause one to die in order to save five others—would wrongly predict that Annie acts impermissibly.

Just as certain common-sense permissibility verdicts are captured by the principle that it's impermissible to kill one to save five others, so too are other common-sense verdicts captured by the principle that it's impermissible to kill another in order to save oneself. Again, that principle is just a specific instance of a more general one:

[IMP2] other things equal, it's impermissible for A to V_T B (/B's X) in order to prevent herself (/her X) from V_I-ing.

As with [IMP], if [IMP2] is correct, then MAKING cuts an important ethical distinction. (Indeed, since the constraints on self-defence presumably parallel those on defence-of-others, [IMP] and [IMP2] come as a package.)

As it happens, we can test [IMP2] by the lights of cases we've already seen. Recall:

¹⁷ I have been assuming that A MAKES B die_I just in case A kills_T B. I think that's right, but the editors of this journal pointed out that this conflicts with a common usage of *kill*. Here is their example: if Godfather orders Henchman to kill Victim, then Godfather doesn't MAKE Victim die, but merely causes Victim to die (by causing Henchman to kill him), yet we might well say "Godfather killed Victim." What's going on here? I think what we mean is that Godfather *had* Victim killed. As evidence, suppose that Prosecutor asks for a list of all those killed by Godfather and suppose that Victim is included on that list; Prosecutor might well complain "no, I want a list of just the ones he literally killed"—and that complaint wouldn't make sense if Godfather literally killed Victim. That said, it's unclear to me just what the rules are for this (more figurative) use of *kill*. I suggest that it has something to do with crediting appropriate responsibility, since Godfather is, in some important sense, responsible for Victim's death. (Interestingly, the same goes for many verbs: we might say "Steve Jobs has built a million computers," since he is, in some important sense, responsible for their being built, even though only few Macs may rightly sit beneath a museum sign that reads "Built by Steve Jobs.") Anyway, "A MAKES B die_I just in case A kills_T B" should be understood as "A MAKES B die_I just in case A *literally* kills_T B."

JUMP: a runaway trolley with enough momentum to kill a single individual is heading towards Annie. Annie jumps off the tracks and the trolley instead hits and kills Ben who is tied to the track behind her.

FLOOD: heavy rain has raised the river to dangerous levels. To protect her house, Annie blocks the doorway with sandbags. The river soon floods and the floodwaters, unable to enter Annie's house, instead enter Ben's, one door down.

It is permissible for Annie to jump off the tracks and to block her doorway and, since she neither kills Ben nor floods his house, these verdicts are rightly predicted by [IMP2].

We can further test [IMP2] by looking at cases inspired by the variations of JUMP and FLOOD, from earlier:

a runaway trolley is heading towards Annie and Ben. Ben is currently protected by a barrier. Unfortunately, the barrier is only big enough to protect a single individual. Annie moves the barrier in front of herself. The trolley kills Ben, while Annie is unharmed.

the floodwaters flood Annie's house while she's away. When she returns, Annie pumps the water out of her basement. That water flows down the street and into Ben's house.

Annie acts impermissibly in both these cases. [IMP2] also gets these right since Annie kills Ben and floods Ben's house.

I think it's no coincidence that common-sense permissibility aligns so neatly with MAKING. I think [IMP] and [IMP2] are correct.

But I've been ignoring the obvious objection. Thomson (1976, 1985) gave the following case:

BYSTANDER: a runaway trolley is heading towards five workers. The only way Smith, a bystander, can save them is by switching the tracks and diverting the trolley down a spur. Smith does so. B is tied to the spur and is killed by the trolley.

And Thomson said that it's permissible for Smith to divert the trolley yet, in doing so, Smith kills B. If Thomson is correct, then BYSTANDER is a counterexample to [IMP]. Everyone took Thomson to be correct, so I must respond.¹⁸ As it happens, we've already seen the makings of that response.

What's so interesting—so vexing—about BYSTANDER's structure is that the trolley that kills B was already set to kill the five and Smith's role was merely to divert that trolley from the five and onto B. Note how we have already seen something very similar to this structure: in FLOOD, the floodwaters were already set to flood Annie's house and Annie's role was merely to divert those waters away from

¹⁸ As it happens, one of the few people who doesn't take Thomson to be correct is Thomson herself. Her current view is that it is *not* permissible for Smith to divert the trolley (2008). Her argument for that conclusion requires that Smith "makes B pay the cost of death" which, to my mind, requires that Smith kills B.

her house and onto Ben's house. Since Annie doesn't flood Ben's house (doesn't MAKE Ben's house flood) in that case, why would it be that Smith does kill B (does MAKE B die) in BYSTANDER?

The natural answer points to the structural difference that remains between the two cases—namely, that in BYSTANDER it's Smith, a third party, who diverts the trolley, while Annie herself diverts the floodwaters. But we can happily do away with this difference:

Heavy rain has raised the level of the river to dangerous levels. Jones sees the waters heading towards a gated hamlet of five houses. The gate is broken, but Jones happens to be carrying some sandbags. Jones blocks the gate with the sandbags and the waters instead enter Ben's house, just down the street.

Jones doesn't flood Ben's house. And if Jones doesn't flood Ben's house, then surely Smith doesn't kill B.¹⁹

There's another route, this time inspired by JUMP, to the same conclusion. Consider:

MALFUNCTIONING: a self-driving trolley has malfunctioned such that it now travels to and runs down the nearest person on connected track. It is currently heading towards Bloggs. Bloggs jumps off the track, so the trolley changes course down a spur towards B who is tied to the track (B now being the nearest person to the trolley). The trolley kills B.

Bloggs doesn't kill B. Is there an important difference between MALFUNCTIONING and BYSTANDER? I don't think so and I suggest that only someone intent on there being one could find it. And, if there isn't a difference, then we must conclude that Smith doesn't kill B in BYSTANDER, after all. (An interesting question is just why everyone has accepted that Smith does kill in BYSTANDER.²⁰ I have no answer, but I suggest it turns on the sort of pragmatic issue raised by Kagan's surgeon case.)

Once we understand killing as MAKING die—akin to MAKING flood, etc.—we realise that BYSTANDER is no counterexample to [IMP], after all. Ultimately, Foot was right about [IMP] and [IMP2] and MAKING cuts an important ethical distinction: ethicists *are* interested in MAKING.

10. A brief aside. If I'm right that MAKING cuts an important ethical distinction, you might wonder why the ethicists themselves don't have their own theories of MAKING. As it happens, theories have been proposed. Here are the three most popular and representative ones; I give a certain amount of

¹⁹ There are always further differences to point to, but each time the question is whether they *matter*. An example of one that surely doesn't matter is that Smith diverts the trolley, but Jones shields the hamlet. As Thomson herself says, "what difference could it be thought to make that a person deflects the trolley by putting a shield around the five instead of by throwing a switch? (2015: 119)" I suggest that the same will go for all further differences between the cases.

²⁰ Interestingly, one person who never accepted that Smith kills in BYSTANDER was Foot, herself. She thought that whether A kills B is fixed by whether A starts the fatal sequence of events that leads to B's death and, in BYSTANDER, "we have the diverting of a fatal sequence and not the starting of a new one" (2002: 85).

detail since I will return to them later. (Following their proponents, I present them as accounts of *killing* specifically. I leave it to the reader to generalise them, should they desire.)

The most straightforward theory of killing has its roots in Alan Donagan's work (1979). It says that A kills B just in case B would not have died had A *done nothing at all*. Borrowing two cases from Jonathan Bennett (1998) to illustrate, suppose that in PUSH, A pushes B off a cliff to his death. If A had done nothing, then B would not have died and so A kills B. On the other hand, suppose that in STAYBACK, B is tumbling towards the cliff edge and A could easily catch B, but refuses; B tumbles to his death. B would still have died had A done nothing, so the theory rightly says that A doesn't kill B. (Of course, the theory stands or falls with its definition of *doing nothing at all*—does ignoring B count? does napping? sneezing?—but we can set that aside.)

Warren Quinn (1989) grounds his theory of killing in the distinction between positive and negative rights. To do so, he distinguishes between cases where B's death is a consequence of something A does (of A's *positive agency*) and cases where it's a consequence of something that A does *not* do, but might have done (of A's *negative agency*). He then says that A kills B just in case B's death is a consequence of A's action and that action is an instance of positive agency. A's pushing B in PUSH is an instance of positive agency (it's something A *does*) and thus Quinn says that A kills B. On the other hand, A's refraining from saving B in STAYBACK is an instance of negative agency (it's something A does *not* do) and so Quinn says that A does not kill B here. (We might ask of positive agency what we asked of *doing nothing at all*—does ignoring B count? does napping?—but we can again set such questions aside.)

Lastly, Bennett (1998) says that when B's death is a consequence of A's action ϕ , A kills B just in case the majority of the ways A could have acted, at the time A ϕ 'd, would not have resulted in B's death. In PUSH, the majority of the ways A could have acted (dancing a jig, taking a nap, etc.) would not have resulted in B's death and thus Bennett rightly says that A kills B. On the other hand, when A fails to catch B in STAYBACK, the majority of the ways A could have acted (dancing another jig, taking another nap, etc.) would still have resulted in B's death and thus Bennett rightly says that A does not kill B. (We can also set aside the question of just how these majorities are to be determined.)

The problem for these theories is that they all suffer from multiple, clear counterexamples. For example, these theories all give the wrong verdict in

JUMP: a runaway trolley with enough momentum to kill a single individual is heading towards Annie. Annie jumps off the tracks and the trolley instead hits and kills Ben who is tied to the track behind her.

Quickly: had Annie not done anything then the trolley would have hit her instead of hitting Ben; similarly, Annie's jumping is an instance of positive agency; and lastly, so long as we suppose that Annie had to jump in a somewhat specific way to avoid the trolley, the majority of ways Annie could have acted would not have resulted in the trolley missing her and killing Ben instead. So Donagan, Quinn and Bennett each incorrectly say that Annie kills Ben.

For similar reasons, the generalised versions of their theories each also give the wrong verdicts in LIGHTNING, FLOOD and DYE. Ultimately, these theories are no good.

11. Legal causation fixes whether A is legally responsible for a given outcome (at least, in Britain and the USA). This legal responsibility is a central element of the law, both criminal and civil: for instance, to be guilty of murder, it's necessary that A is legally responsible for the victim's death; to be found liable in tort, it's necessary that A is legally responsible for the damage; and so on.

As a primer to the concept of legal causation, recall the case from earlier where A poisons B's gin, B notices, heads to the shop to buy a replacement bottle and is killed by a car while crossing the road. Although A both intended B's death and B would not have died had A not acted as she did, A is not guilty of murder. Why? Because A is not a legal cause—is not legally responsible—for B's death.

What is this legal causation? The received view is that it's a technical legal concept, but one that aligns closely with causing. For instance, H.L.A. Hart and Tony Honoré say:

“Over a great area of the law [the courts] have, in using causal language, sought to apply a group of causal notions embedded in common sense” (1959: 123).

And Michael Moore says:

“The central idea that organises this book is that causation as a prerequisite for legal [responsibility] is intimately related to causation” (2009: vii).

“...what criminal law and the law of torts mean by ‘cause’ is what we ordinarily mean by ‘cause’ as we explain the world” (*ibid*: 5).

The received view is wrong. It is wrong because legal causation aligns with causing only to the extent that causing aligns with MAKING since legal causation is, in fact, MAKING. That is, A legally causes X to V_I just in case A makes X V_I: A legally causes Ben to die (/the window to break/the house to be destroyed) just in case A kills Ben (/breaks the window/destroys the house).

Demonstrating this is simple enough: we need only examine a number of legal cases and see whether, each time, the court's judgement *vis-a-vis* whether the defendant legally caused the outcome—caused X to V_1 —aligns with whether the defendant MADE X V_1 .

And, as it happens, we have already examined many such cases: LIGHTNING, FLOOD, JUMP, DYE, BUSHFIRE and POISON (and, often, their variations) are all tidier version of legal cases or distinctions drawn therein.²¹ And, without exception, Annie MADE X V_1 just in case her defendant counterpart was found by the courts to have legally caused X to V_1 .

For instance, recall:

FLOOD: heavy rain has raised the level of the river to dangerous levels. To protect her house, Annie blocks the doorway with sandbags. The river soon floods and the floodwaters, unable to enter Annie's house instead enter Ben's, one door down.

as before, but this time Annie is out of town, is unable to block her doorway and her house floods. Later, she pumps the waters out of her flooded basement. That water flows down the street and into Ben's house.

Annie does not flood Ben's house in the first case and the court found that her counterpart does not legally cause the corresponding property to flood. Conversely, Annie does flood Ben's house in the second case, and the court found that her counterpart did legally cause the corresponding property to flood.²² The same goes for the other cases.²³

Ultimately, legal causation just is MAKING.²⁴ Lawyers are interested in what MAKES what.

12. Ethicists and lawyers are both interested in what MAKES what. It remains an open, empirical question whether theorists in other domains are similarly interested in MAKING. You might be

²¹ LIGHTNING: *Alphacell Ltd v. Woodward* [1972] AC 824; *Doss v. Town of Big Stone Gap* [1926], 134, S.E. 563 (Va.); *State v. Scates* [1858] 50 N.C. 409. FLOOD: see fn. 22 below. JUMP: (roughly) *Scott v. Shepherd* [1773], 96 Eng. Rep. 525. DYE: *Dalby* [1982] WLR 425, *Roberts* [1971] 65 Cr App R 95. BUSHFIRE & POISON: both discussed in Hart & Honoré (1959). For many more cases, see *Blackstone's Criminal Practice 2006* and Hart & Honoré (1959). Each time, the judgement aligns with MAKING.

²² The reasoning: "if an extraordinary flood is seen to be coming upon the land the owner of such land may fence off and protect his land from it, and so turn it away, without being responsible for the consequences, although his neighbour may be injured by it" (*Whalley v Lancashire and Yorkshire Railway Co* (1884) 13 Q.B.D. 131). See also: *Nield v London and North Western Railway Company* (1874-75) L.R. 10 Ex. 4; and for locusts instead of floodwaters, *Greyvensteyn v Hattingh* [1911] A.C. 355

²³ Perhaps you worry that I have just picked the cases that fit my claim. As it happens, things went the other way: I started with those legal cases that the legal textbooks use to illustrate the intricacies of legal causation—largely, intricacies regarding how it comes apart from causing. I then realised that legal causation comes apart from causing in just those cases that MAKING also comes apart from causation. That realisation was, in part, the inspiration for this paper.

²⁴ There are certain *policy-driven* exceptions. For example, if A negligently starts a fire in New York State, which spreads from one building to many others, A is only legally responsible for the damages done to the first building. As Hart & Honoré say, "such rules ... have nothing whatever to do with causation ... [Instead] they represent a particular *policy* which a particular legal system has adopted" (1959: 84). (See also *Environment Agency v. Empress Car Co. (Abertillery) Ltd* [1999].)

sceptical as to just how open that question truly is for other domains. For example, you might think that oncologists are surely not interested in MAKING since they are surely interested in what *causes* cancer. I think that such scepticism is rash: it takes only a glance at the commonly listed “causes” of cancer—smoking, exposure to radiation and asbestos, certain viruses, etc—to see that those “causes” are all things that *mutate* cells, that MAKE cells mutate. That alone is sufficient for the question *vis-a-vis* oncology to remain open. I think the same goes for the other domains, too.

That said, it doesn’t much matter for my conclusion in this paper. My conclusion is that metaphysicians should start paying attention to MAKING and the fact that ethicists and lawyers are interested in what MAKES what is sufficient reason for metaphysicians to do so.

Perhaps, for now at least, you prefer to leave things to the ethicists and the lawyers. But I don’t think that will work. Since MAKING is a general, metaphysical matter, attempting to theorise about it through a domain-specific lens is doomed to failure. The attempts we saw in §10 were examples of this, none clearer than Quinn’s theory: he proposed that whether A MAKES X V_I is, in part, fixed by rights, but something as general as MAKING could never be fixed by something so domain-specific. Instead, we should be theorising about it from a general, domain-neutral standpoint.

And so metaphysicians should pay attention to MAKING. What would that look like? In short, by theorising about MAKING in just the same way they’ve been theorising about causing: asking whether omissions can MAKE or whether MAKING is transitive, etc.; and, ultimately, seeking a theory of just when it is that A MAKES X V_I.^{25 26}

References

- Anscombe, G. E. M. 1971. *Causality and Determination: An Inaugural Lecture*. Cambridge: Cambridge University Press.
- Bennett, Jonathan. 1988. *Events and Their Names*. London: Hackett Publishing Company.
- — — . 1998. *The Act Itself*. Oxford University Press.
- Donagan, Alan. 1979. *The Theory of Morality*. University of Chicago Press.

²⁵ There might well be instances of MAKING that aren’t instances of A MAKING X V_I. For example when Annie throws the rock at the window, she MAKES the window break_I: doesn’t she also MAKE the rock break_T the window? There’s reason to think so. After all, when Annie instructs Ben to throw the rock she not only causes the window to break_I, but she does so *by* causing the rock to break_T the window. And symmetry would have the same go for MAKING—namely, that Annie MAKES the window break_I by MAKING the rock break_T the window. Perhaps that’s right. These are the sorts of issues that a theory of MAKING would seek to resolve.

²⁶ Many people helped me write this paper and I’m grateful to all of them: David Builes, Alex Byrne, Joseph Byrne, Kevin Dorst (who also suggested the title), Cosmo Grant, Daniel Muñoz, Agustín Rayo, Miriam Schoenfield, Kieran Setiya, Judy Thomson, Jake Wojtowicz, Steve Yablo, the audience at *MITing of the Minds 2020*, and the editors of this journal as well as their anonymous reviewer. I am doubly grateful to Caspar Hare and Brad Skow who helped me from beginning to end—thanks guys.

- Fodor, Jerry. 1970. 'Three Reasons for Not Deriving "Kill" from "Cause to Die"'. *Linguistic Inquiry* 1 (4): 429–38.
- Foot, Philippa. 2002. *Moral Dilemmas: And Other Topics in Moral Philosophy*. Oxford University Press.
- Hart, H. L. A., and Tony Honoré. 1959. *Causation in the Law*. Oxford University Press.
- Hornsby, Jennifer. 1980. *Actions*. Routledge and Kegan Paul.
- Kagan, Shelly. 1989. *The Limits of Morality*. Oxford University Press.
- Kaiserman, Alex. 2016. 'Causal Contribution'. *Proceedings of the Aristotelian Society* 116 (3): 387–94.
- Lewis, David. 1986. 'Postscripts to "Causation"'. In *Philosophical Papers Vol. II*, edited by David Lewis. Oxford University Press.
- Moore, Michael S. 2009. *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics*. Oxford: Oxford University Press.
- Morreall, John. 1976. 'The Nonsynonymy of Kill and Cause to Die'. *Linguistic Inquiry* 7 (3): 516–18.
- Murphy, Peter, and Eric Stockdale, eds. 2005. *Blackstone's Criminal Practice 2006*. London: Oxford University Press.
- Quinn, Warren S. 1989. 'Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing'. *The Philosophical Review* 98 (3): 287–312.
- Sartorio, Carolina. 2006. 'On Causing Something to Happen in a Certain Way Without Causing It to Happen'. *Philosophical Studies* 129 (1): 119–136.
- — —. 2010. 'Causation and Ethics'. In *Oxford Handbook of Causation*, edited by Helen Beebee, Christopher Hitchcock, and Peter Menzies. Oxford University Press.
- Sider, Theodore. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.
- Thomson, Judith Jarvis. 1976. 'KILLING, LETTING DIE, AND THE TROLLEY PROBLEM'. *The Monist* 59 (2): 204–17.
- — —. 1977. *Acts and Other Events*. London: Cornell University Press.
- — —. 1985. 'The Trolley Problem'. *Yale Law Journal* 94: 1395–1415.
- — —. 1987. 'Verbs of Action'. *Synthese* 72 (1): 103–22.
- — —. 2008. 'Turning the Trolley'. *Philosophy & Public Affairs* 36 (4): 359–74.
- — —. 2015. 'Kamm on the Trolley Problems'. In *The Trolley Problem Mysteries*, by Frances Kamm. The Berkeley Tanner Lectures. Oxford, New York: Oxford University Press.
- Unger, Peter. 1977. 'The Uniqueness in Causation'. *American Philosophical Quarterly* 14 (3): 177–88.
- Weatherson, Brian. 2007. 'Causation and Causatives'.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.