

Deviant Causation and the Law
Sara Bernstein, University of Notre Dame
Draft of November 2018

A gunman intends to shoot and kill Victim. He shoots and misses his target, but the gunshot startles a group of water buffalo, causing them to trample the victim to death. The gunman brings about the intended effect, Victim's death, but in a "deviant" way rather than the one planned. This paper argues that such causal structures, *deviant causal chains*, pose serious problems for several key legal concepts. A deviant causal chain occurs when an outcome is brought about through different or "deviant" means rather than the means intended. What the law makes of these cases, and how the law should handle them, is the focus of this paper.

Here is the plan. After a few preliminaries (§1), I show that deviant causal chains pose problems for the legal distinction between attempts and completed crimes, and also for the distinction between attempts and mere preparations (§2). I then suggest that cases of deviant causation undermine notions of legal liability for completed crimes in terms of *actus reus* and *mens rea*, and argue that natural accounts of negligence misclassify some cases involving deviant causal chains (§3). I diagnose several conceptual vulnerabilities in the law revealed by deviant causal chains (§4). Finally, I show that natural strategies for shoring up legal theories in the face of these examples are unpromising (§5).

1. Preliminaries

Deviant causal chains share a common structure: an agent intends to bring about x , and x occurs in a different way than that intended or expected by the agent.¹ Following Davidson (1963), I will call *basic deviant causal chains* those deviant causal processes initiated by mere intentions. An example of basic deviance, drawn from Davidson, is the following:

¹ Here I focus on cases of deviant causation initiated by agents, but the cases need not involve initiation by an agent. For example, a falling boulder might kill a victim by startling her to death rather than crushing her. Even in legal contexts, deviant causal chains need not be initiated by agents. For example, we may puzzle over whether to include climate change damage that occurs in peculiar ways, as when a homeowner is injured while packing his bags in anticipation of an upcoming flood. Thanks to Douglas Husak for the latter example.

(Mountaineer) A mountaineer intends to eventually rid himself of the man below him on the rope by dropping him. Merely forming the intention so unnerves him that he loosens his grip on the rope, dropping the fellow climber.

Here, the mere intention to murder the fellow mountaineer initiates a causal chain leading to the murder. This and similar cases are to be contrasted with *nonbasic deviant causal chains*, which stem from an intention and a particular action stemming from the intention. An example of nonbasic deviance is the following:

(Bad Shot) Assassin A is dispatched to kill Victim. Assassin A shoots at Victim, but misses. The shot wakes up Assassin B, who was independently dispatched by a different assassination agency. Assassin B shoots Victim and kills him. (Assassin B would not have shot had Assassin A not awakened him.) Had Assassin A not shot, Victim would not have died.

Here, Assassin A not only intends to kill Victim, but also takes a shot in service to doing so. The shot doesn't directly kill Victim, but does make it the case that Victim dies in virtue of causing Assassin B to shoot. Cases of nonbasic deviance are so-called because there is an action as well as an intention after which the causal chain goes awry relative to the agent's plan. I will largely focus my discussion on cases of nonbasic deviance, since poor intentions alone are neither criminal nor tortious.

Deviant causal chains can be further distinguished along several dimensions. There are those that contain causal processes very different in type from the one planned by the agent. One example is the one with which we began, in which a gunshot wakes up a herd of angry water buffalo. The causal process that includes stampeding water buffalo is a very different sort of causal process than the one that includes a direct gunshot. Contrast the water buffalo case with Bad Shot, in which Assassin A's shot causes Assassin B's shot, which results in Victim's death: one causal process involving a gunshot is swapped for a very similar sort of process, namely, a different gunshot. Call those deviant causal chains containing a very different sort of causal processes than the one planned *process-divergent*, and those containing a similar sort of causal processes to the one planned *process-similar*.

A parallel distinction can be drawn with respect to the nature of the outcome planned by the agent. Some deviant causal chains eventuate in an outcome unlike the one

planned by the agent, while still involving some shared essential planned part with the planned one, e.g., a particular victim's death. If Jane plots to kill Jacintha by gradually adding drops of slow-acting poison to her toothpaste, but Jacintha dies by choking on the same toothpaste that would have eventually killed her in any case, Jacintha dies a different sort of death than the one Jane had planned. Call this sort of event *outcome-divergent*. Contrast this sort with Bad Shot, in which Victim dies a very similar death to the one Assassin A had initially planned, i.e., via gunshot. Call this latter sort of case *outcome-similar*. As I shall suggest, these distinctions become important when attempting to draw lessons for the law from deviant causal chains.

Deviant causal processes pose problems for legal concepts because the “ingredients” for criminal and tortious liability are present, but the cases do not easily or correctly fit into existing legal taxonomies of crimes. Such cases also reveal important ambiguities in the law involving differences between an agent's intended outcome and the actual outcome.

Problems with deviant causal chains arise in part because the standard for establishing the presence of causation in tort law and in criminal law is the simple counterfactual “but-for” test: is it the case that but for the occurrence of *x*, *y* wouldn't have occurred? Since deviant causal chains are initiated by elements that pass the but-for test (i.e., if the legally relevant intention and action hadn't occurred, the deviantly caused outcome wouldn't have occurred), deviant causation becomes the vehicle for legal responsibility. On to the cases.

2. Deviant Causal Chains and Attempts

Here I will suggest that deviant causal chains cause several problems for legal treatments of attempts. Intuitively, an attempt is different than a completed crime insofar as a mere attempt does not result in the intended crime, whereas a completed crime eventuates in the intended outcome. For example, if Jane attempts to injure Jacintha by taking a hammer to her toe, but misses and succeeds only in scaring her, Jane would be legally culpable for an attempted crime rather than a completed one.

Roughly, a criminal attempt occurs when (i) there is specific intent to commit a crime, and (ii) the offender takes direct action toward completion of the crime. More

specifically, many US states share a definition of a criminal attempt according to which “A person commits the offense of attempt when, with intent to commit a specific offense, he or she does any act that constitutes a substantial step toward the commission of that offense.” ((720 ILCS 5/8-4) (from Ch. 38, par. 8-4). According to *Model Penal Code* (USA) (Part 1, General Provisions, Article 5, Inchoate Crimes), a person is guilty of an attempt to commit a crime if, acting with the kind of culpability otherwise required for commission of the crime, s/he:

- (a) purposely engages in conduct which would constitute the crime if the attendant circumstances were as he believes them to be; or
- (b) when causing a particular result is an element of the crime, does or omits to do anything with the purpose of causing or with the belief that it will cause such result without further conduct on his part; or
- (c) purposely does or omits to do anything which, under the circumstances as s/he believes them to be, is an act or omission constituting a substantial step in a course of conduct planned to culminate in his commission of the crime.

Satisfying the definition of a criminal attempt often depends on the establishment of *mens rea* and *actus reus*. Attempts differ from completed crimes insofar as the intended outcome does not occur in an attempt but does occur in a completed crime. Relying on this distinction, the law remains open with respect to what actually occurs as a result of *mens rea* and *actus reus*.

Or so one might think. Nonbasic deviant causal chains exploit this ambiguity, posing at least two specific problems for legal treatments of attempts. First, what would ordinarily be thought of as an attempted crime eventuates in the intended outcome, such that a deviant causal chain satisfies the definition of a completed crime rather than a mere attempt. In nonbasic deviant cases, an agent both forms an intention and performs an action that counts as a substantial step towards the crime. The intended outcome is even caused by the agent’s act. But the scenario does not (and should not) count as a completed crime by the agent, given that the means through which the outcome is brought about are different than those intended. In *Bad Shot*, Assassin A intends to cause Victim’s death, and takes substantial action in service to doing so. Moreover, Victim’s death occurs, and wouldn’t have occurred had it not been for Assassin A’s shot. Assassin A’s shot should count only as an attempt rather than a completed crime. Because the

intended outcome actually occurs, however, Assassin A counts as having initiated a completed crime rather than an attempt.² Exacerbating this problem is that many such cases also respect the concurrence requirement for *actus reus* and *mens rea*³: in Bad Shot, the intention and the act are concurrently timed. So the synchronicity requirement on *actus reus* and *mens rea* cannot be used to recategorize deviant causal chains.

Deviant causal chains also make trouble for the distinction between attempts and mere preparations. What one would ordinarily think of as merely preparing to attempt the crime instead counts as a full-blooded attempt or a completed crime, owing to fortuitous occurrence of the outcome. Consider the following case:

(Murderous Lover) Jane plots to murder Jacintha by poisoning her drink. While preparing the poison, Jane is so nervous that she breaks the bottle of poison. Very shortly thereafter, Jacintha slips on a splinter of glass from the broken poison bottle, and dies from the fall.

Intuitively, Jane was only preparing to attempt to murder Jacintha. If Jacintha's death had not occurred, Jane's actions would not have counted as anything but a mere preparation. But because Jane's mere preparation results in the ultimately intended outcome, preparing the attempt counts (minimally) as the attempt itself. Another example:

(Bad Driver) Jane plots to murder Jacintha with a hatchet. As she is driving to Sears to buy a hatchet, she fails to see a pedestrian in the pedestrian walkway, and hits and kills the pedestrian. Upon rushing out of the car, she realizes the pedestrian was Jacintha.

It seems clear that Jane has merely begun to prepare to murder Jacintha rather than carry out the attempt. The law, however, is forced to view such preparations as attempts *simpliciter*, due to the collective presence of *mens rea*, *actus rea*, and the intended

² Cases of basic deviance sometimes generate similar puzzles. In Davidson's basic deviance case, the mountaineer's murderous intention so unnerves him that it causally results in the desired outcome. But does Mountaineer's dropping his partner count as a full-blooded criminal attempt? If we say "yes", it should be uneasy: the mountaineer did not consciously take action in service to murdering his climbing partner. A deviant causal chain resulting from a mere intention should be punished differently than a full-blooded attempt, at the very least. Even calling it a half-hearted attempt seems too strong. If we say "no", it should also be uneasy: the mountaineer was poorly intentioned and should be held legally accountable for something. Involuntary manslaughter isn't a perfect fit, either, since an intention to harm is present. Voluntary manslaughter, perhaps, is the best fit, though this might be seen as too strong on many readings of the case.

³ Husak (2017) explores the concurrence condition in detail.

outcome.

Nor do existing philosophical accounts of attempts handle this problem well. According to Yaffe's (2010) proposed definition, an attempt includes a causal chain that begins with an intention that plays the proper causal role and ends with the world matching its content. But in failing to give a theory of what it is for the intention to play a proper causal role, the account classifies Murderous Lover and Bad Driver as attempts rather than mere preparations. Yaffe discloses this vulnerability in his account of attempts, admitting:

“What constraints must a causal chain that begins with the intention and ends with the world matching its content meet if it is to be an instance in which the intention plays the proper causal role? This is a very difficult question for which I do not have a complete answer, and I do not have an answer that “solves” cases of deviant causation.” (2010, p. 86)

The challenge is in specifying the appropriate connection between the intention and the way the outcome occurs. The most natural answer is that outcome must be caused by the intention. In deviant causal chains, however, outcomes are caused by intentions, and yet the outcome does not occur in the way intended: a fact relevant to, but not accounted for, by legal treatments of these cases.

3. Deviant Causal Chains and Negligence

It is tempting to think that deviant causal chains are so outré as to not warrant major concern. One might admit that these very strange cases present problems for existing legal taxonomies, yet think there is no major reason to worry about them given their rarity. But as I shall now argue, deviant causal processes play a major role in extremely common legal scenarios, negligence torts, and reveal a host of issues lurking in the background of more normal negligence cases.

Negligence occurs when there is “failure to act with a prudence that a reasonable person would have under the same circumstances,” and the failure to act results in harm. According to *Model Penal Code* (§2.02(2) (d)), “[a] person acts negligently...when he should be aware of a substantial and unjustifiable risk that the material element exists or will result from his conduct.” A negligent act that causes foreseeable harm or injury

counts as liable negligence in the law. Five specific elements are required for negligence: duty (whether a defendant owed a duty to a plaintiff), breach of duty (violation of what was owed), cause-in-fact (causation as given by the but-for test), proximate causation (foreseeable causal result of negligence), and damages (harm or injury to plaintiff).

Prima facie, legal responsibility for negligent deviant causation is an unstable concept since foreseeability is required for negligence torts. Given the unexpected nature of deviant causal chains, it might seem like such cases don't involve foreseeable outcomes, and so there is no such thing as deviant negligence. In most cases of negligence, moreover, cause-in-fact and proximate causation often go together. Consider the following example:

(Rollercoaster) Joe owns a fairground and the rollercoaster tracks have rusted away. Children are mounting the ride. One child, Jane, notices the flaw and runs off to tell someone, but trips and dies on the way. Joe's negligence caused Jane's death. But Joe is not legally responsible for Jane's death.

Here, Joe's negligence in letting the tracks rust away causes the child to trip and die. According to the law, Joe is not legally culpable: because Jane's death in that particular way was not foreseeable, Joe is not responsible for the death despite causing it.

But cause-in-fact and proximate causation do not necessarily come apart in cases involving deviant causal chains. There can be foreseeable outcomes caused by negligent actions that should not necessarily be legally classified as negligence. Consider the following example:

(Cardiac Arrest) Jacquise knows that Kaia is heart attack-prone. Jacquise jumps out from behind a pillar and scares Kaia. Kaia laughs at the dangerous prank, but in doing so, becomes distracted and trips over a planter. The surprise of tripping over the planter causes a heart attack.

This case has every element required for negligence. Jacquise has a duty to Kaia not to perform actions that might cause a heart attack, and breaches that duty by surprising her. Jacquise can foresee that Kaia might die of a heart attack if he surprises her. Had he not jumped out from behind the pillar, Kaia would not have died. And Jacquise causes injury by causing the heart attack. Intuitively, however, such a case should not count as negligence, since the heart attack was deviantly caused.

It is worth noting that case law yields mixed results on the lesson of this latter sort of case. Generally, if an outcome is caused in an unforeseeable way, the defendant is not liable for damages. In *Bunting v. Hogsett*, however, the law ruled in the opposite direction. Hogsett, a train operator, conducted a railroad engine far above the speed limit. In an attempt to avoid a crash with plaintiff Bunting's train, Hogsett reversed the throttle, shut off the engine, and jumped from the train, in the hopes that such actions would prevent a collision. The collision occurred nevertheless, and did not immediately injure Bunting. However the crash did reactivate the engine on the original train, which then struck Bunting, injuring him. The Pennsylvania Supreme Court ruled in the plaintiff's favor:

“[...] It was the engineer's negligence that caused the first collision [...] [The first collision] opened the throttle, and turned loose the destructive agency which inflicted the injuries complained of.” (21 A at 32)

Even though the particular route of harm was unforeseeable, the plaintiff recovered damages.

Similarly in *Sundquist v. Madison Railways Co.* That case concerns a victim of a negligently caused streetcar accident. Though she was physically unharmed by the accident, she was psychologically traumatized by it. Months after the traumatic accident, the victim heard a streetcar bell ring and became so frightened at the sound that she tripped, fell, and was paralyzed by the fall. Though neither the particular form of harm nor the route by which it occurred were foreseeable, the court ruled in her favor. Still, this sort of ruling is unusual. The law's internal conflict about the significance of the type of causal process leading to harm only reinforces the idea that there should be more legal thought on process-divergence and outcome-divergence. Ideally, the law should provide as detailed principles as possible for applying rules of culpability and responsibility, rather than leaving such decisions up to individual rulings. In the next section, I diagnose where and why the law goes wrong.

4.0 Diagnoses

It's not that the criminals or tortfeasors in these fictional and real cases are completely unaccountable. It's that their situations are either misclassified or

unclassifiable under current legal treatments of the relevant concepts. Where does the law go wrong? I hereby identify several problems exposed by deviant causal chains.

First, the law depends on a particular outcome's nonoccurrence or occurrence to undergird a distinction between an attempt and a completed crime. In cases of both basic and nonbasic deviance, however, the intended outcome actually occurs. Thus conditions for completed crimes and completed attempts are satisfied even when the cases are not good fits for existing legal categories.

Second, the law leaves open whether process-similarity or process-divergence effects legal culpability, which is to say, how closely an actual causal chain resulting from an instigator's action must adhere to the instigator's planned way of bringing about the outcome. The law leaves open, for example, whether a death brought about by instigating a water buffalo-trampling with a gunshot should result in the same culpability as an intended death by direct gunshot, where the *mens rea* and *actus rea* are the same in both cases. A causal chain diverging wildly from the planned one, but with the desired result, satisfies the same conditions for liability as non-deviant causal scenarios.

Third, the law leaves open how closely the nature of an outcome must hew to an instigator's plan for its occurrence—in my language, whether it is outcome-similar or outcome-divergent. The obvious conceptual vulnerability lies in what it is for a crime to “match” an agent's intention. In *Murderous Lover*, Jane's death is brought about in a very different way than the one planned. But since the law doesn't weigh in on how and whether such outcome-divergence affects criminal culpability, Jacintha can be legally responsible for Jane's death even when the outcome occurs differently than what she had intended.

Metaphysically speaking, this problem is about the modal fragility of events. An event is *modally fragile* if it couldn't have happened in a different manner without being a different event entirely. Clearly, outcomes for which people are criminally or civilly responsible are not maximally modally fragile: not even the best planner knows the precise second or spot at which the death will occur, or the exact manner in which it will. But the failure of legal thinking on this matter means that there are no implicit or explicit standards for the modal fragility of intended outcomes, leaving existing legal definitions open to intuitive counterexamples.

Fourth, the conditions for legal responsibility for completed outcomes are heavily

rooted in the key elements of but-for causation and concurrent *mens rea* and *actus reus*. Since all elements are present in nonbasic deviant causal chains, such cases end up satisfying definitions for more serious crimes than those committed. The most promising solutions might seem to be additions to the law that help eliminate such cases from counting as full-fledged crimes. But, as I shall suggest in §5, supplementing the requirements for legal accountability does not necessarily help.

Thus far, I have focused on how the law as it is currently structured is inadequate for the task of accounting for such cases correctly. But there is a more serious normative question lurking: what *should* the law say about these cases? How might the law be amended so that deviant causal chains do not pose obvious counterexamples? I do not have space to delve into this complex issue in depth, but I will conclude this section with a pessimistic observation.

Correctly classifying deviant causal chains in the law would require specific guidelines about causal process similarity and outcome similarity, i.e., how close a causal process and outcome must adhere to a criminal or tortfeasor's original intentions. The idea would be for the law to require that an outcome for which an initiator is held responsible be "the right kind" of outcome, i.e., that the outcome must adhere fairly closely to the relevant *mens rea*. Similarly, the law would hold that a causal process be "the right kind" of causal process, i.e., that it must adhere fairly closely to the sort of causal process the initiator had planned. With these sorts of supplementary requirements, outcome-divergent and even process-divergent deviant causal scenarios would not satisfy existing legal definitions in the problematic ways detailed above. In the water buffalo case, Bad Driver, and Murderous Lover, the causal processes and particular natures of the outcomes diverge significantly from the intentions of the initiators: the outcomes occur in different ways and via different causal routes than those planned.

However, such guidelines would face major theoretical and practical barriers. First, it would be very difficult to give a precise, principled metric about process similarity and outcome similarity. Consider the following case:

(Starry Night) Suppose that Maya plans to poison Tegan at 10pm when the storm clouds are rolling in, for extra ghoulish dramatic effect. Maya runs late attempting to mix the poison, however, and so arrives at Tegan's house at 10:45pm. Maya pushes Tegan off the balcony when the sky is starry and clear.

Exactly how much process-divergence and outcome-divergence is there in this example? A precise, detailed answer seems very hard to give. It is a mostly intuitive matter. Additionally, a law building such a metric into the standard for legal culpability would require direct epistemic access to an agent's intentions, as opposed to inferring them from what the agent said and did. Establishing *mens rea* is already a challenge, but ascertaining the contours of an agent's very specific intentions for a crime or tort seems nigh impossible in this setting.

Second, the law would need principled rules for when to emphasize process-similarity and outcome-similarity, and how much weight these should be given in legal culpability. In *Bad Shot*, Victim dies in very much the same way that Assassin A wanted, but the best classification of Assassin A's action is as an attempt rather than a completed crime. In the other direction, being the "right kind of outcome" should not be the major criterion for culpability or liability. Consider the following case:

(Gang Shootout) Two rival gangs, Blue and Grey, are engaged in a standoff. Shooter A from Blue breaking the standoff by shooting at Jelena, the leader of Grey, but misses. Shooter B from Grey shoots at Shooter A from Blue, but misses and ends up killing Jelena via friendly fire.

Here, Shooter A initiates a causal chain via gunshot, intending to directly kill his rival. That the causal chain eventuates in his rival being killed by a gunshot specifically isn't enough to make Shooter A culpable for Jelena's death, however. The fact that Jelena's death is "the right kind of outcome" seems less relevant than the complicated process by which it is brought about.

5. Possible Repairs

In this section, I consider several possible repairs to legal concepts in light of the problems posed by deviant causal chains. Though some are promising partial fixes, I suggest that none fully succeed.

5.1 Foreseeability

The most obvious fix is already implemented in tort law: inclusion of a

foreseeability constraint on legal responsibility. Though its contours are controversial, the basic idea of foreseeability is “whether one can see a systematic relationship between the type of accident that the plaintiff suffered and . . . the defendant’s [wrongdoing].”⁴ Since some deviant causal chains involve unforeseeable process-divergent and/ or outcome-divergent effects, a foreseeability constraint does solve some problems with deviant causal chains. This is particularly the case for scenarios like Rollercoaster, in which the child’s tripping and dying is not a foreseeable consequence of the rusting tracks.

The foreseeability strategy faces several problems, however. The first problem is that foreseeability is presently only a constraint in tort law rather than criminal law, and generally only with respect to negligence. Thus the law as it is currently written leaves out deviant causal chains initiated by “positive” events such as gunshots, and most cases of criminal liability. A natural response would be to add more foreseeability conditions to the law, including criminal law and cases other than negligence. However this strategy faces several problems as well.

First, adding foreseeability to the other conditions for legal responsibility is not always sufficient to generate the correct result. In *Bad Shot*, Victim’s death is foreseeable, and yet Assassin A is not legally responsible for Victim’s death. And in *Cardiac Arrest*, Kaia’s potential heart attack is foreseeable and negligently caused, but still shouldn’t satisfy the legal definition of negligence. In both cases, the outcomes are foreseeable, but the processes by which they occurred are not. Adding foreseeability will overgenerate cases of legal responsibility.

Second and relatedly, it is very hard to construct a foreseeability constraint that is not too stringent or too lax. The law has already struggled to do so, as evidenced by several landmark negligence cases. The landmark *Palsgraf vs. Long Island Railroad Co.* exemplifies the tangle of challenges surrounding foreseeability and causation. In that case, the plaintiff Helen Palsgraf was waiting to board a train to Long Island. As the train was pulling away, two men with a package raced to jump aboard. One man made it onto the train; the other, carrying a package, stumbled to catch the train. Several train employees rushed to help the package-carrier onto the train. The package, which the would-be passenger then dropped, turned out to contain fireworks. It exploded, causing

⁴ “The dominant test of proximate cause in torts makes a defendant liable when but only when the harm he in fact caused was, at the time he acted, foreseeable to him...” (Moore, 1997. P. 363)

injury to Palsgraf. Palsgraf sued the railroad company for negligence. The first ruling was in Palsgraf's favor: she was awarded \$6000 in damages. Upon appeal, however, the decision was overturned. A judge ruled that "Here, by concession, there was nothing in the situation to suggest to the most cautious mind that the parcel wrapped in newspaper would spread wreckage throughout the station." A dissenting opinion was also issued, suggesting that "given such an explosion as here, it needed no great foresight to predict that the natural result would be to injure one on the platform at no greater distance from its scene than was the plaintiff." The lasting upshot of the case is that a defendant who is found negligent is only liable for foreseeable harm or the injury rather than for every injury resulting from the negligent act.

In the British case *Re Polemis & Furness, Withy & Co Ltd*, a defendant's employees loaded cargo onto a ship. Due to an employee's negligence, a plank fell into the hold of the ship. The plank then caused a spark that ignited some gas vapor in the hold. The gas vapor caused an explosion that destroyed the ship. The court ruled that "a defendant can be held liable for all consequences flowing from the wrongful conduct regardless of how unforeseeable." The concept of strict liability for unforeseeable outcomes has now been stricken from British law, but its implications remain controversial even in contemporary legal thought.

Finally, in *Mauney v. Gulf Refining Co.*, an oil refinement company negligently set a fire. As people were frantically running from the refinery fire, a woman who operated a bakery across the street grabbed her child and started running. She tripped over a chair and had a miscarriage, later suing Gulf Refining Company for negligence. The court ruled in favor of Gulf Refining Company holding: "If the appellant didn't see a chair in her way in her own place of business it would impose an inadmissible burden on appellees to say that they should have seen from across the street and through the walls of a building on another corner what appellant did not see right at her feet."

Though this seems like the correct result, legal scholars differ on what the lesson should be from this and similar sorts of cases. Of *Mauney v. Gulf Refining Co.*, Hart and Honoré (1985) lament: "Here the court not merely insists on applying the practical test of foreseeability in order to determine the limit of responsibility but also requires such detailed provision that if similar criteria were generally followed, recovery for ulterior harm would seldom be allowed." (p. 267)

A related problem is that the law is ambiguous between in-principle foreseeability and in-practice foreseeability. Moore, for example, finds the foreseeability constraint too demanding, arguing, “[The rule of foreseeability] looks as if it meant that a motorist while driving should think not merely of the children whom he sees crossing the street, but of the adults who might dash out to rescue them, and of the surgeons who might operate on them negligently should they be run over and taken to a hospital.” (2010, p. 265) These obstacles are not insurmountable, but there is much more work to be done before a foreseeability condition can reliably solve problems about deviant causal chains.

5.2 Adding A Non-Deviant Causation Requirement to Legal Responsibility

Another tempting solution is to include a stipulation in the law that the criminal or tortious causal relationship between the act and the outcome be non-deviant. Yaffe (2018), for example, includes a non-deviance stipulation in his definition of the correct role of the intention in guiding an action, and also in an independent account of criminal culpability. Husak (correspondence) proposes that a defendant should not be liable when the outcome “is too remote or accidental in its occurrence to have a just bearing on the actor’s liability or on the gravity of his offense.”

There are several problems with this strategy, however. First: in some cases, deviant causes *should* be the subjects of legal responsibility. Consider the following case:

(Ricochet) Jane intends to murder Jacintha. Jane shoots at Jacintha, but the bullet ricochets off of a giant slab of marble, and then strikes and kills Jacintha.

Arguably, this is a deviant causal chain. Arguably, Jane is legally culpable for Jacintha’s death even though it did not occur in the way she had planned. Any blanket requirement for legal responsibility involving non-deviant causation is bound to run into similar counterexamples.

Second, philosophers struggle to draw a principled distinction between deviant and non-deviant causation⁵, and such a distinction would be essential to any legal theory based on it. Since the default/ deviant distinction is irreducibly tied to norms and intentions regarding a particular causal scenario, it is extremely implausible that a

⁵ See e.g. Peacocke 1979, Davidson 1980, Mitchell 1982, Schlosser 2007, McDonnell 2015, Blanchard and Schaffer 2017.

principled metaphysical distinction can be drawn in such cases. There is no mind-independent, structural metaphysical difference between Assassin A's and Assassin B's shots. The difference between the shots involves Assassin A's expectation about the causal process resulting from his shot. *Mutatis mutandis* for other deviant causal chains.

Third, since the difference between a deviant and non-deviant causal chain depends on a particular agent's intentions regarding the outcome, such intentions must be evidentially accessible rather than self-reported. One can easily imagine criminal actors and tortfeasors claiming "But I didn't intend for the death to happen in *that particular way!*" How these claims are to be legally evaluated poses an insurmountable obstacle to using such a distinction, and it's arguable, in some cases, that the particular nature of the outcome should matter so much.

5.3 Denying the Transitivity of Causation

Causation is transitive when, if *a* causes *b* and *b* causes *c*, *a* thereby causes *c*. Transitivity of causation undergirds deviant causation: the gunshot causes the water buffalo stampede that causes Victim's death, thus the gunshot causes Victim's death. Rejecting the transitivity of causation in cases of deviant causal chains, or rejecting automatic transitivity of causation more generally, is a tempting solution. In *Bad Shot*, for example, one might hold that Assassin A causes Assassin B to shoot, and Assassin B causes Victim's death, but Assassin A does not thereby cause Victim's death.

This strategy is problematic for several reasons, however. Denying the transitivity of causation only in deviant cases requires a principled deviant/ non-deviant causation distinction. As I suggested in §5.2, the prospects for such a distinction are poor: the distinction is irreducibly mind-dependent and context sensitive, and it is poorly suited to play such a central legal role.

Denying the transitivity of causation more generally is theoretically costly. Transitivity is a central intuitive datum about the nature of causation. Setting aside well-known independent problems with transitivity,⁶ many take the transitivity of causal chains to be a theoretical desideratum of a theory of causation owing to its significant explanatory power. Such explanatory power is particularly central to legal explanations,

⁶ See, for example, Sartorio (2005), Paul (2000), and more recently, McDonnell (2018).

in which causal relationships between proximate causes and their effects often have causal intermediaries. For example, an accomplice who hands his partner the gun thereby causally contributes to the crime.

Denying the transitivity of causation requires principles of demarcation for causal chains. For example, why does my driving my car to the coffee shop cause me to drink coffee, but not my purchase of my car a decade prior which ultimately led to my driving the car? If we are to deny that purchasing my car ten years ago caused my coffee-drinking, we must have a metaphysical or explanatory principle for ruling out the car purchase as a cause. But it is very difficult to give a principled account of when a causal chain stops and another begins. Moore (2009) holds that causation peters out as the number of events in a causal chain increases. But there are several reasons to reject this idea. Clearly there are causal relationships between spatiotemporally distant events, like the discovery of a slate mine 300 years ago and current slate prices, or exposure to asbestos and mesothelioma 25 years after exposure. Further, it seems very unlikely that we will be able to give a principled measure of the degree to which causation “weakens” over time.⁷ Such a measure would have to take a stand on the individuation conditions for events, answering questions like: was Jane’s preparing the poison one event or three events? Was the commencement ceremony one event or hundreds of little ones? The law would have to take a stand on the existence and persistence conditions for causal chains, weighing in on issues like why a causal chain counterfactually dependent on, say, Jane’s birth, begins at *that* particular moment far after her birth and ends at *that* other moment, versus a moment before and a moment later. Such a principle must also include implausibly broad generalizations about the boundaries of causal chains: causal chains peter out after n number of events or n years or n meters, etc. Denying transitivity of causation guts the explanatory power of many causal claims while incurring a massive, unmanageable explanatory burden of its own.⁸

⁷ See Beebe (2013), Bernstein (2017) and Kaiserman (2017), (forthcoming) for more on the prospects of making sense of “degrees of causation.”

⁸ McDonnell (2018) develops a strategy of using proportionality to pick out the causal relations that reliably change, without appealing to transitivity. According to his view, if an agent’s intention does not match an intermediate event in the right way, then a chained causal relation cannot be safely assumed. I do not discuss such a strategy here, but I do find it promising.

5.4 Modifying the type of Causation Required for Legal Responsibility

Since problems with deviant causal chains get off the ground due to the but-for test for causation, one seemingly attractive (if somewhat radical) solution is to modify the central concept of causation in the law. A natural alternative to viewing causation as counterfactual dependence is instead to view causation as a “productive” relation. According to productive accounts of causation, *c* causes *e* if *c* transfers some sort of physical energy, mark, or “oomph” to *e*. A paradigmatic case of productive causation is one domino hitting another. The legal test for causation would then be: is it the case that a physical process relates the putative cause and effect?⁹ This modification is tempting because in many deviant causal chain cases, the initiator does not transfer energy to the outcome. In *Bad Shot*, for example, Assassin A’s shot does not seem to physically interact with Victim. Thus it would not pass a modified legal test for causation.

On balance, however, such a modification wouldn’t help. Productive causation isn’t generally well-suited to be a legal test for causation, since there are countless physical interactions with each outcome that are not legally relevant. Suppose that, in *Bad Shot*, a hapless intern assassin had been hoping to eliminate the target by playing a very loud boombox. (And suppose, further, that boomboxes were still in production.) Then the auditory vibrations of the boombox constitute a plethora of minute physical interactions between the hapless intern’s activities and Victim’s death. Even in the original *Bad Shot* case, the vibrations of Assassin A’s shot might physically interact with Victim in minute, microscopic ways. But such interactions do not seem morally or legally relevant: the liberal, productive sense of “cause” is not alone suitable for legal use.¹⁰ Moreover, some deviant causal chains do contain unbroken chains of productive causation. Consider a modified version of *Bad Shot*:

(Boulder) An unskilled gunman intends to shoot and kill someone. He shoots and misses his target, but the gunshot strikes a boulder, which then rolls down onto

⁹ Hart and Honoré seem to agree that such a modification is warranted in some legal cases, writing “Only by identifying causation literally with the setting in motion of a physical process which has not come to an end can the cause be said still to be contributing to the event.” (p. 244)

¹⁰ See Bernstein (2017) for an argument that a productive theory of causation is not well-suited to moral and legal aims.

Victim, killing him.

In this case, the gunshot transfers energy to the boulder, which then transfers energy to the Victim, killing him. Deviant causal chains cannot be ruled out by modifying the notion of causation into a productive one.

5.5 Dropping Causation in Favor of Causal Verbs

Finally, one might be tempted to drop talk of causation altogether in favor of causal verbs. For example: for A to murder B, A doesn't just make it such that B's death occurs; A kills B. With respect to deviant causal chains, one can then say that the initiator makes it the case that the outcome occurred, but didn't cause it. In *Bad Shot*, for example, one might say that Assassin A made it the case that Victim died, but A didn't kill Victim. And in *Murderous Lover*, one might say that Jane made it the case that Jacintha died, but Jane didn't kill Jacintha.

As one might imagine, there are problems with this strategy as well. Presumably, what unites causal verbs is just that they are specific instances of causation: killing, poisoning, and shooting are just particular instantiations of the causal relation. They are determinates of a common determinable. Causal verbs cannot be explicated in non-circular, non-causal terms.

Further, unpacking the causal verbs will just bring back the problem of whether or not a causal chain is deviant. To claim that Assassin A brought about Victim's death but didn't kill Victim is just to stipulate that Assassin A initiates a deviant causal chain. As I discuss in §5.2, such a distinction is irreducibly anthropocentric and difficult to produce and apply in a principled way.

6. Conclusions

Deviant causal chains pose counterexamples for several central legal concepts, including the legal distinction between attempts and completed crimes, attempts and mere preparations, and the joint sufficiency of *mens rea*, *actus reus*, and actual causation for legal responsibility. In addition to providing counterexamples, deviant causal chains lay bare heretofore underexplored wide vulnerabilities in the law. Here I have focused on

attempts and negligence, but similar problems can be generated for many other legal concepts and distinctions, including (e.g.) the distinction between serious and non-serious criminal attempts, proximate versus non-proximate causation, completed versus abandoned criminal attempts, and substantial vs. non-substantial steps to complete crimes. Despite the existence of a long history of legal thought on causation in the law, the relevance of the particular nature of causal processes to legal responsibility remains underexplored. If I am right, much more investigation is warranted.¹¹

¹¹ Thanks to audiences at University of Barcelona, Columbia University, Bowling Green University, and Rutgers University for very helpful feedback on this paper. Thanks also to Neil McDonnell, Douglas Husak, Daniel Nolan, and Maytal Giboa for helpful feedback and conversation.

References

- Bernstein, Sara (2017). Causal Proportions and Moral Responsibility. In David Shoemaker (ed.), *Oxford Studies in Agency and Responsibility, Volume 4*. Oxford: Oxford University Press. pp. 165-182.
- Beebe, Helen. (2013) Legal Responsibility and Scalar Causation. *Jurisprudence*, 4: 102-8.
- Davidson, Donald (1963). Actions, reasons, and causes. *Journal of Philosophy* 60 (23):685-700.
- Hart, H. L. A. & Honoré, Tony (1959). *Causation in the Law*. Oxford University Press UK.
- Husak, D. (2017) “Courses of Conduct” in *The Ethics and Law of Omissions*, Oxford: Oxford University Press, eds. D. Nelkin and Sam Rickless.
- Kaiserman, A. (2017) Partial Liability. *Legal Theory* 23.
- Kaiserman, A. (forthcoming) “More of a Cause”: Recent Work on Degrees of Causation and Responsibility, *Philosophy Compass*.
- Moore, Michael S. (2009). *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics*. Oxford University Press.
- Moore, Michael S. (1997). *Placing Blame: A Theory of the Criminal Law*. Oxford University Press.
- Paul, L. A. (2000). Aspect causation. *Journal of Philosophy*, 97(4):235–256.
- Sartorio, C. (2005). Causes as difference-makers. *Philosophical Studies*, 123(1-2):71–96.
- McDonnell, N. (2018) Transitivity and Proportionality in Causation. *Synthese* 195 (3): 1211-1229.
- Yaffe, Gideon (2010). *Attempts: In the Philosophy of Action and the Criminal Law*. Oxford University Press.
- Yaffe, Gideon (2018). *The Age of Culpability: Children and the Nature of Criminal Responsibility*. Oxford University Press.