Abstract: According to Schaffer (2000a), "trumping preemption" is a category of redundant causation distinct from early and late preemption and from overdetermination. I show that the putative causal difference between causal processes in cases thought to be trumping preemption generates early preemption or overdetermination rather than trumping. I draw a novel lesson from cases thought to be trumping: that the boundary between preemption and overdetermination should be redrawn.

A Closer Look at Trumping[1]

Sara J. Bernstein

(penultimate version; final version forthcoming in *Acta Analytica*)

Schaffer (2000) introduces what he takes to be a new variety of redundant causation—*trumping preemption*—as a counterexample to counterfactual theories of causation. In this paper I show that trumping is not a new species of redundant causation. However, the possibility of trumping preemption does suggest that the line between causal preemption and causal overdetermination should be redrawn. Understanding differences between types of redundant causation gives us greater insight into the relationships between causation, counterfactual dependence, and completion of causal processes more generally.[2]

Schaffer's major illustrative case of trumping preemption, which I'll call Magic, is as follows:

> "Imagine that it is a law of magic that the first spell cast on a given day match the enchantment at midnight. Suppose that at noon Merlin casts a spell (the first that day) to turn the prince into a frog, that at 6:00 PM Morgana casts a spell (the only other that day) to turn the prince into a frog, and that at midnight the prince becomes a frog. Clearly, Merlin's spell (the first that day) is a cause of the prince's becoming a frog and Morgana's is not, because the laws say that the first spells are the consequential ones." [Schaffer, 2000]

[2] Here I will not be concerned with whether trumping preemption succeeds as a counterexample to counterfactual theories of causation, but rather with what trumping suggests about the relationship between completion of causal processes and the causal relation.

According to Schaffer, Magic isn't a case of late preemption because both causal processes "run to completion," and Magic isn't overdetermination because only one process, Merlin's spell, is intuitively the cause of the enchantment.

For trumping to be a causal category distinct from overdetermination, there must be a difference in causal status between Merlin's and Morgana's spells. There is a difference in causal status between redundant causes when there are metaphysical grounds for calling one event *the cause* and another event *not the cause*. Apart from stipulating such a distinction into the laws, establishing a difference between the causal statuses of both processes has proven difficult, leading some like Hitchcock (2011) and Paul and Hall (2013) to argue that trumping just is overdetermination. Many others have suggested that trumping cases are underdescribed in ways that mask their true causal structure.

In what follows I give a novel argument for these related conclusions and a careful articulation of their reasons: I show that one avenue of establishing a causal difference between the two causes, a difference between types of causal completion, is unpromising. I give two careful readings of trumping cases and show that neither distinguishes trumping from other sorts of redundant causation. I then suggest a new lesson from trumping cases: that the line between preemption and overdetermination should be relocated.

Roadmap: in section 1, I give criteria for distinguishing between different types of redundant causation. I show that the notion of causal completion upon which these distinctions normally rely is ambiguous. I define several different types of causal completion that have previously been run together, and I distinguish between the types. In section 2, I give two careful readings of trumping cases, and argue that each reading is something other than trumping. In section 3, I articulate a new notion of overdetermination that occupies a middle ground between preemption and traditional overdetermination, *asymmetric overdetermination*, and show what it teaches about causation more generally. In section 4, I reexamine Schaffer's cases and conclude that asymmetric overdetermination is a better category for such cases.

2

1. Redundant Causation: Axes of Difference

Redundant causation is characterized by the presence of multiple events c1 and c2 such that (i) at least one of them causes *e*, and (ii) either of them would have caused *e* if the other had been absent. For example: Billy and Suzy each throw a rock at a window. Billy's rock shatters the window, whereas Suzy's rock flies through the space where the window used to be. That is a case of *late preemption*, in which a preempting causal process brings about an effect before the preempted process can. It is to be contrasted with *early preemption*, in which one causal process inhibits another before it reaches completion. For example, suppose that Suzy sees Billy throw his rock at the window and thus does not throw hers, though she would have thrown the rock had he not thrown his. *Overdetermination* obtains when there are multiple causal processes sufficient to bring about an effect in the way that it occurs, as in a case where Billy and Suzy's rocks hit the window at precisely the same time and either would have been sufficient to shatter the window.

Species of redundant causation differ along two closely related dimensions: difference in the causal status of redundant causes, and completion of causal processes.

Differential causal status is a largely intuitive notion, but the general idea is that one cause *really* brings about the effect and the backup one doesn't. A difference in causal status is a hallmark of preemption, since only the preempting causal process actually contributes to the effect: in the case where Billy's rock shatters the window and Suzy's rock flies through the space where the intact window was, only Billy's rock causally contributes to the shattering. Overdetermination, in contrast, is typically characterized by no difference in causal status between the causes: the actual causal contribution of each cause is individually sufficient to bring about the effect in the way that it occurs.

Differences in causal status between redundant causes are normally taken to depend on differences in *completion* between causal processes. The relevant notion of "completion" is rarely articulated, but the intuitive idea is that a complete causal process is not interrupted in any way. In a case of early preemption, for example, the preempted process is "stopped" before it reaches the effect, whereas the preempting causal process reaches all the way to the effect. Thus the causal asymmetry between the preempting and

preempted processes lies in the fact that one is uninterrupted but the other isn't.

But this method is unreliable in evaluating trumping cases for two reasons. First: trumped processes *are* complete. Second: confusion over the many interpretations of "complete process" plagues the literature. Here are several different ways the term can be understood:

a. *Continuity completion*. A causal process is continuity complete if there is an uninterrupted chain of actual causal intermediary events between c and e, such that each event in the process is caused by its immediate predecessor. In an early preempted causal process, for example, there are no causally intermediary events between the point of inhibition and the effect, whereas the preempting process has an uninterrupted string of causally intermediary events running from the cause to the effect.

The above notion captures the simplest and most commonly discussed type of causal completion. But not every type of completion is so simple, or involves an actual causal relationship between the process and the effect. Consider the following types of completion:

b. *Spatiotemporal completion*. A causal process is spatiotemporally complete if the intermediary events between c and e are spatiotemporally contiguous with each other. For example, Billy throws his rock at the window, and every intermediary event between the throw and the window shattering is spatiotemporally contiguous with the last.

c. *Relational completion*.[3] Suppose that $c1$ is the cause of e, and that $c2$ is a backup process that exists at every point along $c1$'s trajectory. Call $c2$ "relationally complete", owing to its ability to bring about the effect at every point along the main process' route, should the main process fail. Relational completion is suggested by locutions like "at no point does there fail to be a backup process," used to describe both late preemption and overdetermination. According to this usage, causal completion is a relation borne from a "backup" causal process to a "main" causal process in virtue of the former's ability to bring about the effect during the main process' trajectory.

---

[3] See Lewis (1973, 1999) and Paul and Hall (2003) for uses of this understanding of completion.

Relational completion can also obtain relative to particular times that the main process occurs. Consider a case of late preemption in which Billy throws his rock at time $t1$ and Suzy throws her rock at time $t3$. From $t3$ to the effect, Suzy's process is relationally complete with respect to Billy's process: it acts as a backup process for the occurrence of the shattering. But from t0 to t3, there is no relational completion, since Suzy's process hasn't yet been initiated.

A sub-type of relational completion is *dependent completion*. A causal process $c2$ is *dependently* complete if, at every time at which the primary process $c1$ occurs, $c2$ "pays attention" to the primary process, and is ready to bring about the effect should the main process fail. For example, imagine a remote-controlled smart-but-lazy rock that is programmed to spring into action from the ground, should the other rock fail to do its job. Suppose that the smart-but-lazy rock is so quick that it can still bring about the shattering that Billy's rock would have brought about, were it to be called into action. Now imagine that Billy throws his normal rock at the window; Suzy activates the alertness of her smart-but-lazy rock resting on the ground; Billy's rock shatters the window; and Suzy's smart rock is never called into action from the ground. In this example, Suzy's rock is dependently complete with respect to Billy's causal process: at every point along Billy's rock's trajectory, Suzy's smart-but-lazy rock is ready to leap to action and shatter the window. At no point along Billy's rock's trajectory is there *not* a backup cause of the window's shattering. Suzy's process is not continuity complete, since there are no <u>actual</u> causal intermediaries between the smart-but-lazy-rock and the shattering of the window, but it is *dependently* complete, since it could have leapt into action at any time along Billy's rock's trajectory.

Note that dependent completion and continuity completion can co-occur. Consider a modified smart rock that will only continue on its trajectory if Billy's rock continues on its trajectory. At every moment, the modified smart rock "checks" to see what Billy's rock is doing, and if Billy's rock is still moving, the smart rock moves as well. Billy throws his normal rock; Suzy throws her modified smart rock; Billy and Suzy's rocks hit the window at exactly the same time. In this case, the process stemming from Suzy's smart rock is both continuity complete and dependently complete, unlike the original smart rock case in which there are no actual causal intermediaries between the smart rock and the shattering of the window.
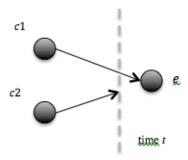
d. *Timely and late completion*. A causal process terminates either at the time the effect occurs, or runs to terminus after the preempting cause brings about the effect (as in late preemption). Call the former *timely completion* and the latter *late completion*.

A timely complete process leading from c_1 to e terminates exactly at the time that e occurs. Generally, a timely complete process terminates *by* causing the effect. Overdetermining processes, for example, are often understood as timely complete. Whether or not a timely complete process can count as *non*causal with respect to the effect is at the heart of the controversy about trumping. Put more precisely, that question is: can there be a causal process which is (i) sufficient to bring about the effect, and (ii) terminates at the exact time the effect occurs, but does not cause it? I address this controversy in later sections.

In contrast, a *late complete* process terminates after the effect occurs without causing it. Consider the following case of late preemption:

Suppose that c_1 is Billy's rock; c_2 is Suzy's rock; and e is the shattering of the window. The vertical line indicates the spatial location of Suzy's rock at the time the shattering occurs. Note that the causal process leading from Suzy's rock to the effect is uninterrupted in the sense that there are not any missing events between Suzy's rock throw and the terminus of the causal process originating from the rock throw—Suzy's rock merely flies through the space where the window used to be and lands. But Suzy's rock isn't *timely complete* in that it's not poised to bring about e at the time that it occurs. Being a few inches behind Billy's rock, Suzy's rock is positioned to bring about the shattering at a time shortly after Billy's rock brings it about, say, t+1. Thus characterizations of late preemption differ with respect to whether the preempted causal

6

process is complete[4] because "causal completion" is ambiguous between timely and late completion. For clarity, we can call the late preempted process late complete but not timely complete.

e. *Non-causal timely completion*. Trumping requires a special type of completion: *non-causal timely completion*. A process is non-causally timely complete if it is (i) timely complete and, (ii) spatiotemporally contiguous with the effect, but doesn't cause the effect. The general idea of non-causal timely completion is that the causal process is sufficient to bring about the effect in the way that it occurs and terminates at the exact time that the effect occurs, but does not have a causal relationship with the effect. Clearly, this is a very strange sort of causal process. Below, I examine this type of process in closer detail, and argue against its existence.

With a better understanding of the different types of causal completion, we can provide a more careful taxonomy of the paradigmatic types of redundant causation. Timely completion and completion *simpliciter* distinguish the three traditional types of redundant causation: early preemption involves one complete causal process and one incomplete one; late preemption involves one timely complete process and one late complete process; and overdetermination involves two timely complete causal processes.[5] We are also in a better position to examine the nature of putative trumping cases, which must exhibit a special kind of causal completion in order to be distinct from other types of redundant causation. I will now turn my attention to this topic.

2. The Mechanism of Trumping: Two Readings

Trumping cases are unified by three common features: (i) two complete causal processes, each sufficient to bring about the effect in the way that it occurs, (ii) a set of laws or rules governing the cause to which the effect responds, and (iii) an alleged

---

[4] Menzies (2008) describes the causal process in cases of late preemption as complete, whereas Hall and Paul (2003) and Hitchcock (2011) describe late preempted processes as incomplete.
[5] Here I assume that overdeterminers are, in fact, causes.

difference between the causes owing to the laws or rules.[6] Consider the following such case:

> (Time-Lock Safe) It is a rule of a remote-controlled time-lock safe that at midnight it responds to the first lock command of the day. At 6am, remote control A commands the safe to lock. At noon, the remote control B commands the safe to lock. At midnight, the safe locks.

In this case, both commands are individually sufficient to bring about the locking of the safe in precisely the same way; the rules govern the command to which the locking is responsive; and there is an alleged causal difference between the two commands, owing to the rules dictating the command to which the safe responds.

Two questions: First: what type of completion does the trumped causal process exhibit? Second, how does the trumped causal process differ from the trumping causal process?

A good starting point for answering these questions is careful consideration of the precise *mechanism of trumping*. In Magic, the mechanism of trumping is the magical law; in Time-Lock Safe, the mechanism of trumping is the device that controls the command to which the safe responds. As we'll see, the devil is in the details of these mechanisms.
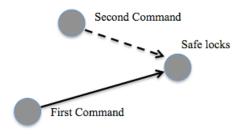
A trumping mechanism must work in one of two ways: either it causally privileges the first process at the moment that the process is initiated, thus fixing the causal privilege of the first process early on in the causal structure; or the mechanism "waits" until both processes run to completion to check which one was initiated first. In the Time Lock Safe example, either the mechanism privileges one command over the other at 6am, or it waits until midnight to check which command was initiated first, and "listens" to that command. After examining each of these procedures in closer detail, I will argue that the former way suggests an early preemptive structure and the latter way suggests an overdetermined structure.

---

[6] Schaffer's original case, Magic, is odd along two dimensions: first, it uses magical spells, about which it is easy to lack causal intuitions or to consider the classificatory verdict about such processes "spoils to the victor," in Lewis' well-known terminology. The idea is that intuitions about the case are too shaky a basis upon which to draw causal conclusions. Second, there are no intermediary events between magical spells and the enchantment, confounding attempts to analyze the structure of causal processes in trumping cases. To get a better grip on what makes trumping distinctive, I introduce a new case before returning to Magic.

2.1.1 Early Decision

First way: the device "locks in" the first process as the cause at 6am, the moment that it begins. Though the second causal process is initiated and its electronic command runs to the effect, the second process is never a "causal contender" for bringing about the effect owing to its being initiated after the first. Call this reading of the mechanism of trumping *early decision*, since the mechanism decides early on which cause will bring about the effect.

On the early decision reading, it makes sense to think of the first command as causally "active" and the second command as causally "inactive," since the mechanism of trumping has already chosen the first command of the day as the causally efficacious one. Being active signifies <u>actual</u> sufficiency to bring about the effect. That scenario is represented in the following diagram:
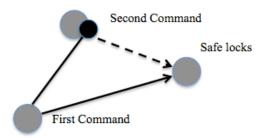

The solid line represents the "active" causal process that brings about the effect. The dotted line represents the trumped, "inactive" causal process. Note that while the "inactive" process is not continuity complete, it is still *dependently* complete, since we can assume that it would spring into action were the first process to halt, as in the case of the remote-controlled smart-but-lazy rock which can bring about the shattering of the window but which stays on the ground. In other words: rather than cutting off the second process entirely, the first command shifts the completion of the second process from what would have been continuity completion to merely dependent completion—it "watches" to make sure that the first process does its job, ready to spring into action should the second process fail.

If this is the correct reading of the mechanism of trumping, the friend of trumping could hold that the difference between the trumped and trumping causal processes consists in differences in types of causal completion: the trumping process is continuity

9

complete, whereas the trumped process is only dependently complete.

      The problem is that the difference between the two processes doesn't make the case trumping. If one process impairs the continuity completion of the backup process before it reaches the effect, then the case is best considered an outré case of *early* preemption, not trumping preemption. Recall that, in early preemption, the main causal process inhibits the stream of actual causal intermediaries between the backup process and the effect. If one command removes actual causal intermediaries in the process before it reaches the effect, then it is natural to understand the first command as an inhibitor, as follows:
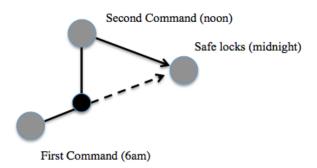

In this diagram, the second command is "switched off" by the first command at the time it begins. (This isn't merely a theoretical exercise: we can imagine that the electronic circuitry is programmed in such a way that the second process is switched off in the presence of the first in order to conserve energy.) So understood, the trumping process inhibits the trumped process, since the former removes the latter's actual causal sufficiency to lock the safe.

      Why think of one cause as an inhibitor in this case? Because the type of causal completion exhibited by the trumped process after it is inhibited is merely dependent completion: the trumped process is no longer actually sufficient to lock the safe, because the inhibited process has no actual causal intermediaries between the point of inhibition and the locking of the safe. To understand this idea, consider a trumping case in which the rule is reversed, as follows:

> (Reverse Time-Lock Safe) It is a rule of a remote-controlled time-lock safe that at midnight it responds to the latest lock command of the day. There are at maximum two commands. At 6am, remote control A commands the safe to lock. At noon, remote control B commands the safe to lock. At midnight, the safe locks.

The early decision reading takes the causal decision to be made exactly when the second

causal command is issued. Thus the first command is put on hold at the time the second command initiates, since the safe can only listen to the latest command of the day, and there are at most two commands. In other words, the first command is causally inhibited *because of* the one following it, as illustrated by the following neuron diagram:

Note that the second command doesn't cut off the completion of the first process simpliciter: we can assume that, were the first process to somehow stop, the second would spring back into action. That is: process A is continuity complete from 6am to noon, then dependently complete from noon onwards. At noon, the second command changes the type of causal completion of the first command from continuity completion to dependent completion. But because one redundant causal process hinders the string of actual causal intermediaries of another, this reading suggests an early preemptive structure: between noon and midnight, process A isn't *actually* sufficient to bring about the locking of the safe. Rather, process A is disarmed by process B far before it reaches the safe insofar as there are no <u>actual</u> causal intermediaries between process A and the effect. Given that the inhibitory structure is essential to the way we understand early preemption, the early decision mechanism of trumping generates early preemption rather than trumping.
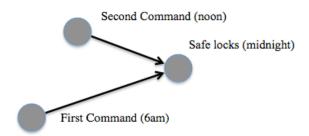
2.1.2 Late Decision

The second reading of the mechanism of trumping interprets both causal processes as timely complete and spatiotemporally complete.[7] On this reading, the

---

[7] That the Late Decision is the correct reading is a widespread assumption in the literature. Schaffer himself (2002) seems to assume this reading, as do Paul and Hall (2003), Hitchcock (2011), and Funkhouser (2009). This reading is natural given the assumption of only one type of causal completion.

mechanism (in this case, the electronic circuitry of the safe) waits until midnight to "check" the timestamp of the order to which it should respond. At midnight, the electronic device recognizes that the first command was issued at 6am, and executes that command. Such a scenario is represented in the following diagram:

Here, both causal processes are causally active up until midnight, when the device picks the command to which it will be responsive. Call this the *late decision* reading, owing to the fact that the mechanism waits until midnight to decide which command to implement.

There are several problematic features of the late decision reading. First, rules "checking" to see which causal instruction to heed *ex post facto* are absurd. Assuming that a rule of this sort is a stand-in for natural laws, this sort of law is implausible: even granting that the causal relation is inextricably bound up with laws, no law "waits" until causal processes are complete to listen to one over the other. With background conditions, laws in our world determine causation, but don't primitively prioritize certain causes over others given antecedent causal parity of the processes.

Second, this reading of trumping relies on the existence of a *non-causal timely complete* process; that is, a process that is sufficient to bring about the effect in precisely the way that it occurs, and which terminates at the effect, but doesn't cause it. What should we make of this type of process?

A few details will help. First, we must assume that the mechanism of trumping does not block the preempted cause from spatiotemporally connecting to the effect, or else it would undermine the completion of the trumped process. So the trumped process must be spatiotemporally complete: that is, it must make physical contact with the effect. But the trumped process must be spatiotemporally contiguous with the effect *without making a difference to the way the effect occurs*: if the trumped causal process changes the way the effect occurs, then the relationship between the trumped process and the

12

effect is straightforwardly causal in some way.[8] To get a grip on this idea, imagine that the trumped causal process *does* make a difference to the way the safe locks at midnight—for example, the circuitry issues the command to lock with slightly more electrical charge. Taking events to be modally fragile, the fact that the trumped causal process makes a difference to the way the effect occurs recategorizes the case as *joint causation underdescribed*: at a maximally specific level of description, both the trumped and the trumping processes are necessary to bring about the effect in precisely the way that it happens. This scenario would also undermine the distinctiveness of trumping. Even without the assumption of modal fragility, a trumped process that makes a difference to the effect causes *some* property of the effect, making the relationship causally suspect. A clear case of trumping requires that the trumped process isn't difference-making with respect to the effect.

For the trumped causal process to be spatiotemporally contiguous with the effect without making a difference to it, the effect is physically insensitive to the force of the extra process.[9] This means that the effect has a property that renders it impervious to the force of an extra cause. For example, the safe is programmed such that it doesn't lock any differently whether it is preceded by one or two locking commands. (In contrast, consider a physically sensitive safe whose locking is accompanied by a beep when it is impacted by a second process. Here, both causes are necessary to bring about the effect in exactly the way that it occurs.) In a case of putative trumping, the effect is physically impervious to the extra process.

Thus a non-causal timely complete process is characterized by (i) spatiotemporal contiguity with the effect, (ii) sufficiency to bring about the effect in exactly the way that it occurs, and (iii) connection to an effect that is insensitive with respect to its force. Simply stated: the trumped causal process is physically connected to the effect and is <u>able</u> to bring about the effect, but does not make it occur any differently.

Now, given that what accounts for the fact that the trumped process doesn't make a difference is *the mechanism or effect itself*, there is no reason to think of the *process*

---

[8] Moreover, there might be reason to consider trumping joint causation, in which multiple causes are necessary to bring about an effect in exactly the way it occurs.
[9] For more on physically insensitive effects, see [BLINDED].

leading up to the effect as special—that is, to believe in the existence of a special causal process that generates trumping. Rather, on a late decision reading, it seems that the mechanism "chooses," as it were, which process lays claim to being the cause <u>after</u> both processes have completed their trajectories. But nothing is distinctive about the trumped process itself; i.e., there is nothing about the causal process that distinguishes it from the trumping process. In support of this point, consider the following trumping-structured case that shows that the process *itself* makes no difference:

> (Course Enrollment) There is one open spot left in a course. Two students request enrollment from the registrar's computer system. Suzy submits her request at 6am; Billy submits his request at noon. At midnight, the system checks the timestamp on each request and permits Suzy to join the course.[10]

This example fits a late-decision structure of trumping because (i) Billy and Suzy's requests are individually sufficient to fill the empty spot in the course, (ii) the mechanism implements one request after both causal processes are complete, and (iii) there is an intuitive causal difference between the two requests (namely, Suzy's request is the cause of the spot being filled and Billy's request is the backup).

In this example, the special mechanism of the registration system privileges the causal contribution of one student over the other by making the system impervious to the later request. But up until the mechanism chooses which student can register, the causal processes themselves do not differ from each other: they are each timely and spatiotemporally complete. There's no reason to consider the trumped causal process special in any way. Rather, as in the Late Decision reading of Time Lock Safe, it's the mechanism of the *effect* (in this case, the computer system) that creates the alleged causal difference between the trumped and trumping causal processes, since the mechanism "decides," in one way or another, which potential cause is the actual causal contributor after both causal processes are completed.

Now, given that there is nothing special about the trumped causal process, the sufficiency of both causes provide strong theoretical pressure to view the case as overdetermination. If the relationship between an effect and a spatiotemporally complete

---

[10] Thanks to [BLINDED] for this example.

process whose actual causal contribution is sufficient to bring about the effect isn't *causation*, then what is it? As Hitchcock points out, Schaffer admits elsewhere[11] that complete processes and their effects are linked by a "causally suggestive relation." But in the absence of a difference in the physical impact of the preempting cause and the preempted cause, the best explanation is that the relationship *is* straightforwardly causal.

In support of this point, consider that the late decision reading of trumping already has the causal structure of overdetermination: there are two individually sufficient causes, each connected to respective timely and spatiotemporally complete processes, and each is available to bring about the effect in precisely the way that it occurs. If the device malfunctions and fails to invoke the timestamp-checking procedure at midnight, then the case is *incontrovertibly* overdetermination, given the presence of multiple sufficient causal contributions for filling the open spot. Only the rule governing the request to which the system responds privileges one cause over the other. Similarly with Time Lock Safe: if the rule in the system were to fail, then the commands would function as straightforward overdeterminers.

To establish that trumping holds, one must insist that the rule *just is* what distinguishes trumping from overdetermination. But primitive rules and laws play no role in distinguishing other forms of redundant causation from each other. That the bare "causal stuff" of the two scenarios—the events, and the processes attached to them— is precisely the same speaks against a rigid distinction between the two categories. Setting aside the laws, the causal processes are symmetric in virtue of having the same type of completion.

Since many have balked at the use of laws and rules in trumping cases[12], I will not further explore that issue here. Even if we accept that laws play a role in distinguishing trumping from other kinds of redundant causation, we can productively mine trumping cases for lessons on another important issue for redundant causation: where we draw the line between preemption and overdetermination. I turn now to this matter.

3. A Third Way: A New Notion of Overdetermination

---

[11] "Overdetermining Causes" (2003).

[12] See Funkhouser (2009) and Hitchcock (2011) for suspicions about the use of laws in trumping.

The friend of trumping holds that there is a causal difference between the trumped and trumping causal processes; the trumping skeptic holds that the role of laws is suspect and that the case should be categorized as overdetermination. The disagreement seems to revolve around whether the trumped cause is, in fact, a *cause* of the effect: intuition says no, yet the rote structure of trumping minus the laws says yes. Causal completion of processes can't play the only role in answering the question, since the late decision reading of trumping interprets both causal processes as complete in precisely the same ways.

Taxonomies of redundant causation adhere to two common dogmas. First dogma: that a difference in causal status between redundant causes is a hallmark of preemption. For example, if one redundant cause is actually sufficient to bring about the effect and the other is merely counterfactually sufficient, then the case is one of preemption. Second dogma: that there is straightforward overdetermination only when there are multiple *actual* causal contributions that are sufficient to bring about an effect. For example, if one causal contribution is actually sufficient and the other causal contribution is counterfactually sufficient, then the case cannot be overdetermination.

Abandoning these dogmas opens a promising middle ground. Rather than assimilate trumping to the category of traditional overdetermination, we should relocate the boundary between preemption and overdetermination.[13] That way, we can admit that there *is* a causal difference between the trumping and trumped causal processes, but that the difference between them isn't as stark as *being the cause* and *not being the cause*.

Ordinarily, overdetermination is taken to involve two causes whose *actual* causal contributions are each sufficient to bring about an effect. For example: Billy and Suzy each throw a rock through a window, and the actual impact of each rock is individually sufficient to shatter the window. In this paradigmatic case of overdetermination, both Billy and Suzy can equally lay claim to being the cause of the window's shattering.

But let us suppose that there is a special category of overdetermination that involves *counterfactual*, rather than actual, causal contribution of at least one cause. An event c_1 in a class of causes Y counterfactually overdetermines an effect e if: had other

---

[13]For more on this new notion of overdetermination, see [AUTHOR].

causes c_2, c_3, c_4, etc. in class Y not occurred, c_1 would have brought about e in precisely the way that it happened. Here, events can both be overdeterminers even with a difference in causal status: one cause is an actual causal contributor and the other is a counterfactual causal contributor, but both are overdeterminers. Call these sorts of cases *asymmetric overdetermination*.

We should not resist calling the trumped causal process an overdeterminer even if its causal contribution is merely counterfactual: our concept of overdetermination makes room for this new notion. Reconsider the course enrollment example, in which the system chooses which student can enroll in the course. The initial temptation is to call Suzy's request the actual cause of the filling of the course and Billy's request the backup, thus categorizing the case as preemption. But Billy's request isn't causally *irrelevant*, as with an early or late preempted cause: after all, it is physically connected to the effect and (were it not for Suzy's request) sufficient to bring about the filling of the course at exactly the time that it occurs. Yet it also seems wrong to say that Billy's request has the <u>same</u> causal status as Suzy's request. Intuitively and metaphysically, the causal status of Billy's request with respect to filling the spot is something in between being *the* cause of the course being filled, and being causally irrelevant. And the reason Billy's request has this intuitive status is that it is connected to a timely and spatiotemporally complete causal process. Were it not for the decision rule of the registrar's system, Billy's causal process has the "physical stuff" to fill the extra spot in the course.

This case suggests that our concept of overdetermination should be stretched to prize completion of a causal process over its actual ultimate causal contribution to the effect. In other words: the actual causal contribution of Billy's request to the filling of the course is less relevant to its categorization than the fact that (i) its causal process is spatiotemporally complete and timely complete, (ii) it could have brought about the filling of the course at precisely the same time that Suzy's request did, and (iii) it is a mere artifact of the effect's late decision procedure that Billy's request *wasn't* the cause of the course being full. The fact that Billy's causal process is complete and can bring about the effect suffices for categorizing it as a funny sort of cause, even if the mechanism prevents its actual contribution to the filling of the course.

Rather than think of Billy's rock as preempted, we should think of it as an *asymmetric overdeterminer*. An event is an asymmetric overdeterminer if (i) it is attached

to a spatiotemporally complete and timely complete causal process, and (ii) it could have brought about the effect in precisely the way that it actually occurred, had its causal contribution not been stopped *ex post facto*. Then the anatomy of a late-decision trumping case is as follows: there are two timely and spatiotemporally complete causal processes sufficient to bring about the effect. The effect is insensitive with respect to the extra force of the trumped cause. The causal contribution of the trumping cause is actual; the causal contribution of the trumped cause is counterfactual. The difference between the trumping and trumped processes lies in the respective actual causal contributions of the events, but not in completion differences between causal processes attached to them. Understanding the trumped cause as an asymmetric overdeterminer allows us to have our cake and eat it too: we can explain why the trumped cause is less than fully causal without judging it to be causally irrelevant or taking it to be as causal as the "chosen" cause. We locate the difference between the trumped cause and the trumping cause in the correct place: in the mechanism of the effect, rather than in the processes themselves.

Note that the claim is not the overly radical idea that *all* merely counterfactual contributors are causes. Rather, the idea is that cases in which a mechanism "chooses" the command to which it should adhere *ex post facto* are cases in which the non-chosen backup still counts as a cause owing to the completion of its causal process. Since all of the causal ingredients are there up until the decision, it is not a stretch to hold that the backup contributor *is* causal in virtue of its spatiotemporal contiguity with the effect. The idea is that the counterfactual contributor is classified as causal in virtue of its attachment to a spatiotemporally and timely complete causal process, and that this classification is independent of the rule that decides which one ultimately counts as the actual cause.

One might be worried that this debate is terminological. The worry is that *I* call the case "asymmetric overdetermination" and the friend of trumping calls it "trumping," but we both seem to agree on the causal structure of the relevant cases. Not so. For the friend of trumping holds that the trumped process is not a cause *at all*, whereas I hold that it *is* a funny sort of cause: a cause whose actual contribution is blocked by a late-decision mechanism but deserves the name "cause" nonetheless. The friend of trumping and I disagree on the causal status of the trumped cause, and hence on the lessons of the cases more generally.
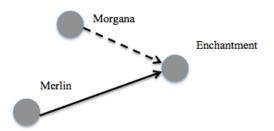
A summary will now be helpful. On the early decision reading of trumping, the

18

actual sufficiency of the trumped causal process is inhibited before it reaches the effect, thus rendering the case early preemption. On the late decision reading of trumping, the trumped causal process connects with the effect at the moment that it occurs and is sufficient to bring it about, thus structuring the case like overdetermination. No *detailed* reading of the trumping case distinguishes it from other forms of redundant causation. Even granting a causal difference between the trumped and trumping causes, there are compelling reasons to view trumping as a special kind of overdetermination rather than as a distinctive causal category. That trumping can be assimilated to overdetermination suggests an interesting general lesson: completion of causal processes is more important than actual causal contributions in classifying cases of redundant causation. Timely completion and spatiotemporal completion of causal processes suggests overdetermination even if a mechanism blocks the actual causal contribution of one of the causes.
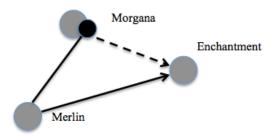
4. Schaffer's Cases

With these lessons in mind, let us return to Schaffer's example. Despite the strangeness of the case owing to the lack of causal intermediaries between the trumped cause and the effect, the case can be analyzed by modifying the treatment of both spells as causal processes rather than magical action-at-a-distance. We can thus apply the same two readings to the case and draw the same lessons as Time Lock Safe.

Early decision reading: the cause of the enchantment is "locked in" at the time the first spell is cast, rendering the second spell causally inert at every point after the first spell. The first spell is the only causally active one up until midnight. At midnight, the enchantment comes to fruition. On this reading, Merlin's spell is continuity complete and Morgana's spell is dependently complete, since the latter is in a sense "switched off" by the first spell of the day having been cast. Here is a diagram of the case. The solid line represents the causal process active until the moment of the enchantment; the dotted line represents the causal process deactivated after its initiation:
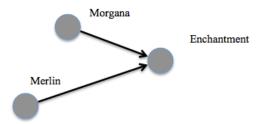
If the preempting process renders the preempted process causally inert before it reaches the effect, then as before, it is natural to classify the case as early preemption, not trumping. We can understand Merlin's spell as an inhibitor, as follows:

Here, Merlin's spell (in virtue of being the first spell of the day) causally disconnects Morgana's spell from the enchantment at the moment it begins. Loosely speaking, Morgana's spell does not have a chance, given the existence of Merlin's spell. Its causal deactivation from the moment it begins removes its causal relevance to the enchantment.

On the late decision reading, both causal processes are causally active until the enchantment. At midnight, the enchantment "checks" to see which spell of the day was first, and matches that spell. On this reading, both causal processes are timely complete and spatiotemporally complete, since both processes are poised, at the moment the enchantment occurs, to bring about the enchantment:

But as before, the difference between Merlin's and Morgana's spells lies not in the *processes,* but in the effect that "chooses" which spell is the cause after the processes are complete. Given that both processes are complete in precisely the same ways, we

should understand Magic as a case of asymmetric overdetermination: each causal process is timely complete and spatiotemporally complete, but the effect "decides" that Merlin's spell will make the <u>actual</u> causal contribution. Morgana's spell counterfactually overdetermines the enchantment.

Finally, we can briefly examine another one of Schaffer's trumping cases that I'll call Charge:

> "Imagine that in a world that could well be our own, the major and the sergeant stand before the corporal, both shout "Charge!" at the same time, and the corporal decides to charge. Orders from higher-ranking soldiers trump those of lower rank." [Schaffer, 2000][14]

Unlike the other underdescribed cases, Charge immediately lends itself to an early decision reading. Since it is a key stipulation of the case that the major and the sergeant stand before the soldiers, then we can assume that the soldiers notice the presence of both the major and the sergeant before the shouted orders, and antecedently know that if the major shouts, *his* is the order they should listen to.[15] Thus the sergeant's order is causally inactive from the moment it begins. Given this scenario, it is natural to see the major's order as inhibiting the sergeant's, since the latter process is immediately made causally irrelevant by the soldiers' antecedent recognition that the major's order trumps the sergeant's. Described carefully, this case has a clear early preemptive structure.

However, we can easily modify the case so that a late decision reading is plausible:

> (Modified Charge) Officers are positioned out of sight of the soldiers. Orders from higher-ranking soldiers trump those of lower rank. The major and the sergeant both shout "Charge!" at the same time, and the soldiers charge.

Here the late decision reading is as follows: supposing that the soldiers hear the two orders at the same time, they consciously decide which order to heed after hearing both. In that case, we can reapply the argument from above in order to classify the case. Given

---

[14] Let us suppose that the soldiers can recognize and distinguish the respective voices of the higher-ranking officers. Thanks to an anonymous referee for this clarification.

[15] Lewis (2000) picks up on the special features of the case, noting that this might be a case of "cutting" (early) preemption.

that both orders are timely complete and spatiotemporally complete, we can understand the case as one of asymmetric overdetermination: the major's order is the actual cause and the sergeant's order is the asymmetric overdeterminer. Once again, we do justice to the "middle ground" causal status of the sergeant's order by understanding it as a special kind of overdeterminer: had the major not shouted his order, the sergeant's causal process would have brought it about in precisely the same way in virtue of its timely and spatiotemporal completion.

3. Conclusion

Careful examination of the mechanism of trumping reveals that such cases are best classified as early preemption or asymmetric overdetermination. Even granting a unique role to laws in marking off trumping from other sorts of redundant causation, attention to the details of the cases along with theoretical considerations about the anatomy of redundant causation provide compelling reasons to deny the distinctiveness of trumping. Trumping does suggest, however, that the line between preemption and overdetermination is misdrawn. Reading late decision-structured trumping as asymmetric overdetermination balances intuitions about causation while respecting the metaphysical facts of such cases.

References

Funkhouser, E. 2009. Frankfurt Cases and Overdetermination. Canadian Journal of Philosophy. Volume 39, Number 3, pp. 341-369.

Hitchcock, C. 2011. Trumping and Contrastive Causation. Synthese 181:227–240.

Hall, N. and Paul, L. *Causation: A User's Guide*. 2013. Oxford: Oxford University Press.

---------- Hall, N. and L. A. Paul. (2003). Causation and Preemption. In P. Clark and K. Hawley (eds.), *Philosophy of Science Today*. Oxford: Oxford University Press.

 Lewis, D. Causation as Influence. 2000. Journal of Philosophy 97 (4):182-197.

-------- 1973. Causation. Journal of Philosophy, 70: 556–567.

Menzies, P. 2008. Counterfactual Theories of Causation. In The Stanford Encyclopedia of Philosophy (Summer 2008 edn), ed. E. Zalta. http://plato.stanford.edu/archives/sum2008/entries/causation–counterfactual.

Schaffer, J. 2000. Trumping Preemption. The Journal of Philosophy, Vol. 97, No. 4, Special Issue: Causation. pp. 165-181.

_____ Overdetermining Causes. Philosophical Studies 114: 23-45.