

Backward causation: harder than it looks

Abstract

According to Albert, there are certain situations where we can cause past events. In response to a well-known objection that we never observe backward causation, he argues that there are good reasons why we can't tell when it obtains. However, I identify a new difficulty with Albert's view: at face value, it has the unattractive consequence that backward causation is not just possible but downright rife. In this paper I show how this implication can be blocked. I then use my analysis to defend Albert's account from a second well-known objection, namely that it allows us to control the past.

1 Introduction

Albert (2000, 2014) has offered a broadly Lewisian account of causation on which the causal arrow ultimately stems from the record asymmetry - the fact that we have records of the past and not the future. A surprising corollary is that there are certain scenarios where we can cause events in the past. At first sight, this might seem inconsistent with the fact that we never observe backward causation. But by showing that this phenomenon is by its nature unobservable, Albert takes this criticism in stride.

Nevertheless, I identify a new difficulty. We ordinarily think backward causation is rare or non-existent. Although Albert suggests his account respects this widely held view, at face value it does not, for the following reason. On his account, any unrecorded past event is vulnerable to backward causation. But one would think that there are countless unrecorded minor events in our past, from coconuts falling in long-lost jungles to pterodactyls squawking in Jurassic skies. If this is right, then countless events are vulnerable to backward causation. This implies backward causation is not just possible, but rampant throughout the world.

To remedy this, I propose a mechanism whereby records persist whilst *seeming* to vanish. The implication is that despite appearances, there are *not* countless unrecorded minor events in our past, so backward causation is genuinely rare/non-existent after all. Not only does this square Albert's causal account with common sense, it also disarms a second well-known objection relating to control.

This paper is structured as follows. In section 2 I outline Albert's truth conditions for causal statements, and explain how we have probabilistic epistemic access to these via the 'Mentaculus'. In section 3 I explain how this picture is meant to yield a causal arrow that usually points forwards, but occasionally backwards. I also show how Albert responds to a well-known objection that we never observe backward causation: he claims it is unverifi-

able. In section 4 I identify a new problem: when Albert’s account is paired with commonsense ideas about records, it has the unattractive consequence that backward causation is extremely common. In section 5 I revise these commonsense ideas to overcome this issue. In section 6 I use my response to disarm a second well-known objection to Albert which says his account allows us to control the past. I conclude in section 7.

2 Causation in the Mentaculus

Albert wants to show why we are “likely to be in a position to influence much about the future and next to nothing about the past” (2014, 165). In context, it is clear that Albert uses ‘influence’ to mean ‘cause’. To understand his explanation, we must first get to grips with his framework for causal analysis. This runs as follows: causal statements are cashed out in terms of a factual statement plus a counterfactual statement, and an entity called the ‘Mentaculus’ gives probabilistic verdicts on these.¹ In what follows I lay this all out, starting with causal statements. This isn’t exactly how Albert presents his account, but I believe it is a fair reconstruction.

Suppose in our world, some particular event x occurs (I’ll always assume this), and we wish to know whether or not it causes some other particular event y in our future/past. On Albert’s account, as on Lewis’ (1973a, 1973b), x causes y iff both these Causal Criteria are satisfied:

Criterion 1: y occurs in our world

Criterion 2: Had $\neg x$ obtained, $\neg y$ would have obtained

Criterion 1 is a kind of ground zero requirement for causation, for it establishes that the relevant events actually occur in the first place. Meanwhile,

1. Loewer (2007, 2011, 2012) paints a similar picture, but analyses causation through the lens of possible human interventions; see Frisch (2007, 353-358) for a comparison. For consistency, I shall only focus on Albert’s proposal. See also Kutach (2002) for a similar account.

Criterion 2 is the definitive feature of a counterfactual analysis of causation. As Lewis himself recognised, this analysis faces various difficulties, and requires some refinement in order to get around them. Two particularly salient problems are pre-emption scenarios and backtracking counterfactuals.² But since these issues aren't directly relevant for us and would only complicate things, I shall set them aside.

In summary, a causal statement is determined by a factual statement (Criterion 1) and a counterfactual statement (Criterion 2). So, what dictates their truth values? Within the context of classical statistical mechanics, these are determined by the trajectories picked out by the relevant microstates in the universe's phase space Γ . This is straightforward in the case of Criterion 1: it is satisfied when the *actual* current microstate m_0 time-evolves forwards/backwards under the dynamics to satisfy y . But what is the 'relevant microstate' in the case of Criterion 2? Like Lewis (1973b), Albert (2014, 163) consults our closest possible world, which he interprets as the nearest nomologically possible microstate m_\star to m_0 that satisfies $\neg x$ (as measured by distance in Γ). Criterion 2 is then satisfied when m_\star time-evolves forwards/backwards to satisfy $\neg y$.³ This sums up the truth conditions for causal statements.

Although these truth conditions are conceptually clear, they are epistemically opaque. We can't literally envisage microstates, much less microstates of the *entire universe*, and we certainly can't track their trajectories in Γ . Instead, we must make do with probabilistic verdicts on the Causal Criteria given by the Mentaculus. This entity is a combination of three ingredients:

2. See Paul and Hall (2013, Ch. 3) and Reiss (2015, 110-114) respectively for discussions of these issues. See also Kim (1973), who argues that the Causal Criteria are insufficient to characterise causal connections.

3. Whereas Lewis (1973a, 1979) envisaged possible worlds satisfying antecedents via 'miracles' (violations of law) and identified the closest with recourse to his infamous metric, Albert does away with all this.

1. The universe's fundamental dynamical laws. For simplicity, these are taken to be Newton's deterministic laws.
2. The 'Past Hypothesis', the posit that the universe began in an extremely low-entropy macrostate M_{PH} . Since Albert (2000, Ch. 6) regards this as a law, all nomologically possible microstates are consistent with this posit.
3. The 'Statistical Postulate', a probability distribution that's uniform (according to the Lebesgue measure) over M_{PH} 's microstates.⁴

The result is a uniform distribution over trajectories in Γ , which Albert regards as the universe's nomologically possible trajectories. By conditionalising the Mentaculus on the actual macrostate M_0 or on the counterfactual macrostate M_\star (where $m_0 \in M_0$ and $m_\star \in M_\star$), we obtain y 's respective probabilities in the actual or counterfactual world *given* all macroscopic evidence in that world.⁵ The procedure for finding these probabilities is the same in both cases:

- i Apply a uniform distribution over M_0 or M_\star (this satisfies the Statistical Postulate).
- ii Discard any microstates that didn't evolve from M_{PH} (this satisfies the Past Hypothesis). This leaves us with the nomologically possible sets of present microstates $N_0 \subset M_0$ or $N_\star \subset M_\star$. These are just macrostates, but conditionalised on M_{PH} .
- iii Time-evolve N_0 or N_\star forwards/backwards under the dynamics to the time of the consequent.

4. Albert (2000, Ch. 7) argues that if the GRW theory is correct, we can dispense with the Statistical Postulate, as the universe's chanciness will be built directly into the dynamical laws. But since we're assuming a classical framework, we can put this idea aside.

5. Of course, we never observe *all* this evidence. But presumably, observing a large chunk of it brings our credences into rough agreement with the Mentaculus' probabilities.

- iv Determine the proportion of N_0 or N_\star that find themselves in a macrostate in which y occurs. There may be multiple such macrostates, but I shall gloss these as N_y .
- v This gives $\Pr(y | N_0)$ or $\Pr(y | N_\star)$, i.e. the chances of y occurring in the future/past of N_0 or N_\star .

This provides us with probabilistic handles on the Causal Criteria as follows. $\Pr(y | N_0)$ is the chance of Criterion 1 being *satisfied*, whilst $\Pr(y | N_\star)$ is the chance of Criterion 2 being *violated*. When $\Pr(y | N_0) \approx 1$ and $\Pr(y | N_\star) \approx 0$, both Causal Criteria are probably satisfied, and x probably causes y .

3 The Possibility of Backward Causation

We've now identified the conditions under which causation occurs, and our means of epistemic access to these conditions via macroscopic evidence. With this groundwork in place, we're ready to see how it yields a causal arrow, and why this may occasionally point backwards.

In order to derive the causal arrow, we must first acknowledge that the world exhibits a record asymmetry. At this point, two questions naturally arise. First, what *is* a record? This is notoriously hard to answer in a precise way. But for our purposes, we can stick with what seems to be Albert's own view: records are localised macrostates that are highly informative about other times,⁶ and it just so happens that these 'other times' always lie in our past.

6. See Albert (2000, Ch. 6) for details. Why are records macrostates and not microstates? I think there is a principled reason for this: in order for something to function as a useful record, it must i) be reasonably stable over time, and ii) enter into reliable correlations. Microstates, being full specifications of the positions and momenta of all particles, are constantly and erratically changing, and therefore lack these key features. Interestingly, Albert's conception of records hinges on what qualifies a macrostate in the first place - a matter than depends on our own physical constitution. See Hemmo and Shenker (2016) for details, and Stradis (2021) for a discussion of how this might bear on the record asymmetry.

Second, what explains the record asymmetry? According to Albert (2000, Ch. 6), it is the fact that the Mentaculus contains a Past Hypothesis but no analogous ‘Future Hypothesis’. For our purposes the details and difficulties of this explanation aren’t relevant,⁷ and I am simply going to take it as read. The implication is that not only the universe’s actual macrostate N_0 but also the counterfactual macrostate N_\star contains records of the past and not the future. Through piecemeal observations of *our* world’s records - which, recall, are localised macrostates - we are able to partially reconstruct N_0 and N_\star , allowing our credences to roughly track the Mentaculus’ probabilities.

If we grant the above, the causal arrow drops out as follows. As we saw, whether or not x causes y depends on the two Causal Criteria. If we take Criterion 1 for granted (i.e. both x and y occur in our world), then the question hinges on Criterion 2. This is a counterfactual statement that we gauge via the Mentaculus. Because of the record asymmetry, the Mentaculus gives different answers whether x precedes y or vice versa, as shown below.

When x precedes y , Criterion 2 is a forward counterfactual. To evaluate this, we time-evolve N_\star forwards and gauge the proportion that land up in N_y , yielding $\Pr(y \mid N_\star)$. Since N_0 lacks records of the future, it lacks records of y . But since N_\star mimics N_0 except in satisfying $\neg x$, N_\star generally also lacks records of y . It can therefore easily happen that $\Pr(y \mid N_\star) \approx 0$. In such circumstances, Criterion 2 is probably satisfied, and so x probably causes y .

When y precedes x , Criterion 2 is a backward counterfactual. To evaluate this, we time-evolve N_\star backwards and gauge the proportion that land up in N_y , yielding $\Pr(y \mid N_\star)$. Since N_0 contains records of the past, it generally *does* contain records of y . But since N_\star mimics N_0 except in satisfying $\neg x$, it generally *also* contains records of y . This implies $\Pr(y \mid N_\star) \approx 1$. In such circumstances, Criterion 2 is probably violated, and so x probably *doesn’t* cause y .

7. See Frisch (2007) for a useful overview and critique.

In summary, here's how the causal arrow emerges in Albert's account. When Criterion 1 is satisfied, x causes y iff Criterion 2 is satisfied. When x precedes y , Criterion 2 is a forward counterfactual; this may well be probable, so x can easily cause y . But when y precedes x , Criterion 2 is a backward counterfactual; this is generally improbable due to records, so x probably doesn't cause y .

As Albert (2014, 165-166) observes, the causal arrow that drops out of this isn't absolute; there are certain scenarios where it probably points backwards. To illustrate this, he offers the following example (credited to Adam Elga).

Suppose Atlantis existed in our past (y), and that I just kept my fingers still (x). Does x cause y ? Since Criterion 1 is satisfied, we must turn to Criterion 2. Because y obviously precedes x , Criterion 2 is a backward counterfactual, so we determine it as follows: time-evolve N_\star backwards to antiquity, and gauge the proportion that land up in an Atlantis-containing macrostate N_y . By hypothesis, N_0 happens to lack records of Atlantis, and since N_\star mimics N_0 except that I just clicked my fingers ($\neg x$), N_\star also lacks records of Atlantis. This implies $\Pr(y \mid N_\star) \approx 0$, which means Criterion 2 is probably satisfied. Therefore, my finger-resting probably causes Atlantis' existence. If we generalise from this example, there are certain scenarios where it's probable that we cause events in the past.

Incidentally, this example could be generalised a step further. Although the antecedent here happens to involve human action, this doesn't actually do any important work in the argument. My finger-resting is just a generic present event that could be interchanged for virtually anything, from a leaf falling off a tree to a soap bubble bursting. Therefore, the real upshot is that there are certain scenarios where it's probable that present events *in general* cause past events; *our* doing so is merely a special case of this. But since agency will become important in section 6, I shall stick to examples whose antecedents are human actions.

Elga's example poses a clear challenge to Albert's account: there are certain situations where we can probably cause past events, and yet we never observe backward causation. How can these ideas be reconciled?

To this end, Albert (2014, 166) draws a key distinction between a) the chance that backward causation is in play, and b) the chance *we assign* to it being in play, and rightly points out that these may diverge. This is apparent in the Atlantis case as follows. On the one hand, we saw that the value of a) is 1. This is because it was *stipulated* that Atlantis existed in our past, i.e. that Criterion 1 is satisfied. However, we don't have direct epistemic access to this fact. Instead, our credence comes from conditionalising the Mentaculus on the universe's current macrostate, and since this lacks putative records of Atlantis, it will *seem* like $\Pr(y \mid N_0) \approx 0$. This makes it seem like Criterion 1 *isn't* satisfied, making the value of b) negligible. So, even when backward causation is highly probable, we're none the wiser - thus eliminating any verified counterexamples to the orthodox view that backward causation is rare/non-existent.⁸

4 The Rampancy of Backward Causation

Albert's response is good as far as it goes. However, it runs into a different problem. Regardless of whether or not we observe backward causation, we take this phenomenon to be extremely rare or non-existent. Albert doesn't explicitly flag this up as an important feature of causation that he needs to respect. Nevertheless, he seems to acknowledge this implicitly, stressing at various points that we are "likely to be in a position to influence much about the future and next to nothing about the past" (2014, 165). In order to be palatable, Albert's causal account must satisfy the idea that backward causation is rare/non-existent. At face value however, it violates this idea, as I explain below.

8. This proposed barrier to backward causation is similar to Dummett's (1964). As I understand it, his 'dancing chief' example is meant to show that acting for the sake of the past is only worthwhile if we can't know what happened - which we generally *can*, at least in principle.

In arguing that we can unverifiably cause Atlantis' existence, Albert makes a key hidden assumption: most nomologically possible microstates lacking records of Atlantis didn't contain that city in their past. Without this assumption, we'd have no reason to expect only a tiny proportion of N_{\star} to time-evolve backwards into N_y , and hence no reason to expect $\Pr(y \mid N_0)$ to be minuscule, which was the rationale for saying Criterion 2 is probably satisfied in that scenario. Moreover, we'd have no reason to expect only a tiny proportion of N_0 to time-evolve backwards into N_y , and hence no reason to expect our credence in Atlantis' existence to be low, which was the rationale for saying backward causation is unverifiable.

Presumably, this is just a special case of a much more general assumption, which we can express as follows:

Fidelity Assumption: If a nomologically possible region N_i lacks records of y , then $\Pr(y \mid N_i) \approx 0$

This tells us that if some event occurred but has no records in N_0 or N_{\star} , then the Mentaculus ascribes that event a low probability in both cases. It would then follow a fortiori that if Atlantis existed but has no records in N_0 or N_{\star} , then the Mentaculus ascribes that city's existence a low probability in both cases.

If the Fidelity Assumption is indeed integral to Albert's reasoning, then Atlantis' vulnerability to backward causation is just one instance of a wider fact: *all* unrecorded past events are vulnerable to backward causation. As I explain below, this implies something radical about the frequency of this phenomenon.

If the chance of any given event leaving behind no records is tiny, then there would rarely/never be unrecorded events in our past, so backward causation would be rare/non-existent, in accord with common sense. There is certainly a large class of *major* events for which this is clearly true, such as the

existence of a city, the eruption of a volcano, or the impact of a large asteroid, all of which we'd expect to leave behind many records. However, it's not obviously true of *minor* events like a stone rolling off a cliff, a Neanderthal's cry, or a carp splashing on a lake. On the contrary, one would ordinarily think such events have a *high* chance of leaving behind no records - at least after an appreciable period of time. Any such event would be subject to backward causation by later events. But since the past is choc-a-bloc with minor events, this suggests that backward causation is not just possible, but rife.

For clarity, let's see how my objection takes shape with an example. Suppose sometime during the Tang Dynasty, a particular carp made a particular splash at a particular time on the Yangtze River (y), and that a moment ago I kept my fingers still (x). Does x cause y ? Since Criterion 1 is satisfied, we must examine Criterion 2 using the Mentaculus: time-evolve N_\star backwards by hundreds of years, and find the proportion that land up in the splashing-carp-macrostate N_y . Since N_0 lacks records of the carp's splash, and since N_\star mimics N_0 except that I just clicked my fingers ($\neg x$), N_\star also lacks records of the splash. This implies $\Pr(y | N_\star) \approx 0$, which means Criterion 2 is probably satisfied. Therefore, my finger-resting probably causes the carp to splash.

Just as in Atlantis case, the Fidelity Assumption entails that our credence in the splash is low, so it won't *seem* like my finger-resting causes it. But with a multitude of minor events in our past, it seems we can be confident that there is a vast array of anonymous past events that are caused by everything we do in the present.⁹

Rampant backward causation follows as a corollary of Albert's causal account plus a commonsense idea about how records vanish, and it contradicts the standard view that backward causation is rare/non-existent. To deal with

9. Like Elga's example, this is in fact a special case of a more general issue: countless minor past events are caused by any present event we choose to consider, whether it involves human action or not.

this, we face a trilemma: either give up the view that backward causation is rare/non-existent, revise our ideas about vanishing records, or abandon Albert’s causal account altogether. In the next section I motivate option two.

5 The Dispersion Mechanism

In this section I propose a mechanism whereby an event’s records persist even though they appear to vanish. I shall start by clarifying the statistical relationship between events and their immediate records. Then, by applying these ideas iteratively for later times, we will develop an idea of how events relate to their more temporally remote records.

Any event C raises the chance of a potential record of it, such as A :

$$\Pr(A \mid C) > \Pr(A \mid \neg C) \tag{1}$$

For example, a splashing carp raises the chance of a circular wave forming. However, we know from experience that C doesn’t guarantee A :

$$\Pr(A \mid C) < 1 \tag{2}$$

In our example, the wave could be suppressed by a sheet of glass, or a well-orchestrated gust of wind. What then *does* guarantee A ’s occurrence, prior to that event?

In a deterministic system, every event has a “determinant” at every other moment of time, a “minimal set of conditions jointly sufficient, given the laws of nature, for the fact in question” (Lewis, 1979, 474). A determinant is a time-slice of the event’s past or future light cone.¹⁰ A ’s prior determinant is

10. See Arntzenius (1990, 83-84) and Frisch (2005, Ch. 8). Lewis (1979) envisaged records playing the role of determinants, but since records are localised, this is inconsistent with the light cone time-slice picture.

therefore C plus some very complex, far-reaching, microscopically specified set of ‘background conditions’ U_A :

$$\Pr(A \mid C \wedge U_A) = 1 \tag{3}$$

This is illustrated in Figure 1.




Figure 1: Record A drawn with its prior determinant, ‘ C plus U_A ’ (the time axis runs from left to right). Solid lines represent causal correlations.

From experience, it seems plausible that C generally leaves behind multiple records, even though no particular item is guaranteed. This implies U_i are not mutually exclusive, and may be numerous. We can incorporate this into our example by imagining that in addition to U_A leading to the circular wave A , there also exists U_B which leads to a curious heron B . This is illustrated in Figure 2.




Figure 2: Records A and B drawn with their prior determinants, ‘ C plus U_A ’ and ‘ C plus U_B ’ respectively. The dotted line represents a non-causal correlation.

So, given C , multiple records tend to form depending on which U_i happen to obtain. But if we make the assumption that U_i are uncorrelated, the records that end up forming will *also* be uncorrelated - that is, uncorrelated given C .¹¹ The upshot is that in the wake of C , the records which *actually* materialise will seem to us randomly selected from the set which *potentially could* have materialised, i.e. the set normally associated with C .

To understand why this is significant, let’s reiterate this logic to see how events like C relate to their more temporally remote records. By construing one of these records (say, A) as a recordable event in its own right, we can give it the same treatment as C itself. This again produces various new records depending on which U_i obtain. For example, D (a startled dragonfly) might result if U_D obtains, whilst E (a wobbling reed) might result if U_E obtains. This is shown in Figure 3. Whereas A is a ‘first-generation’ record of C , we can think of D and E as ‘second-generation’ records of C .

11. What we are looking at here is the famous ‘fork asymmetry’. My assumption that U_i are uncorrelated seems plausible: many have derived the fork asymmetry (or similar structures) from independence conditions in the universe’s initial state, and there are various reasons to expect such conditions to hold. For more on these two points see Frisch (2014) and Lloyd (1994) respectively.




Figure 3: Second-generation records D and E , together with all background conditions that they depend on.

It's easy to see the pattern emerging: D for instance might cause F (well-fed frog) if U_F obtains, and G (awestruck fisherman) if U_G obtains. This is shown in Figure 4. Whereas F and G are immediate or 'first generation' records of D , they are 'third-generation' records of C .




Figure 4: Third-generation records F and G of C , along with all background conditions that they depend on. For simplicity, I've omitted the second-generation record E (which branches off from A) and the third-generation records it yields.

This ‘Dispersion Mechanism’ suggests that later-generation records become harder to utilise than early-generation records for two reasons. First, they are more ‘motley’ and hence harder to recognise as records of the original event. Second, even if these records could be recognised as such, they’re less reliable in their own right. Let me unpack these claims in turn.

What I call the ‘motley’ of later-generation records is a by-product of two things that happen over time. On the one hand, the range of possible later-generation records of C becomes larger. This can be appreciated by filling in my omission in Figure 4 (see caption). On the other hand, the chance of any given later-generation record becomes smaller, as its formation hinges on a longer list of U_i obtaining. For example, whereas A hinges on U_A alone, F hinges on U_A , U_D , and U_F . Taken together, these two factors (‘more diverse’ and ‘less certain to appear’) entail that later-generation records have a miscellaneous, motley character, making them less obviously symptomatic of the original event. This is clear in our example: we all recognise a circular wave as indicating a splashing fish, but who can say the same about a well-fed frog?

Meanwhile, the unreliability of later-generation records can be understood as follows. In order to be reliable, a putative record of C must have a veridical genealogy stretching back to that event. However, later-generation records have longer genealogies. This is clear from Figure 4: A only depends on the genuine occurrence of C , whereas F depends on the genuine occurrences of D , A , and C . Therefore, when we trace back their origins, later-generation records have more opportunities to be spurious than early-generation records. Our example helps illustrate the contrast: whereas a circular wave misleads us about C if it’s caused by a falling pinecone (and not the splashing carp),

a well-fed frog misleads us about C if the frog ate a mosquito (and not a dragonfly), *or* if the dragonfly was startled by a swan (and not the wave), *or* if the wave was caused by a pinecone (and not the carp).

So, here is a short summary of the above. Events generally produce multiple records, none of which is guaranteed, and since U_i are uncorrelated, those records that actually materialise seem randomly realised from the set of possibilities. By reiterating this logic for each of these records, we obtain a picture in which C 's late-generation records are motley and unreliable. This means that even long after an event, records still remain even though they seem to have vanished.

For clarity, let's see how this blocks backward causation. Suppose there's an event y in our past. What's the probability of a present event x causing it? This is high when the Causal Criteria are satisfied, and since Criterion 1 is satisfied, we must turn to Criterion 2. However, y occupies the role of C in the above picture, and therefore leaves records in the present via the Dispersion Mechanism. Since N_\star also contains these records, $\Pr(y \mid N_\star) \approx 1$, and so Criterion 2 is probably false. This means x probably *doesn't* cause y . Since this logic applies irrespective of whether y is a major or minor event, there *isn't* a vast backlog of past events that are subject to backward causation. Hence, this phenomenon is genuinely rare/non-existent after all.

One might raise the following objection. According to the above picture, a given very-late-generation record of y will be extremely unreliable. Therefore, when we conditionalise the Mentaculus on this record, this will raise the probability of y only marginally. But in order for it to be unlikely that a present event x can cause y , we require $\Pr(y \mid N_\star) \approx 1$, i.e. we require y 's counterfactual probability to be high. How can this be achieved?

The answer is that although records become less reliable with time, this is counterbalanced by the fact that they become more numerous. This follows from the fact that events leave behind multiple records; since this holds

true for records themselves, we would expect the sum-total of C 's records to mushroom with each iterative generation. Again, this is apparent in Figure 4. So, granted that conditionalising on any single very-late-generation record raises y 's probability only marginally, conditionalising on all of them raises this value significantly. Since N_{\star} embeds the whole lot, $\Pr(y \mid N_{\star}) \approx 1$ still stands. So, even long after y has occurred, the chance of a present event x causing it will remain low.

This section comes with a caveat. The Dispersion Mechanism rests on the notion that events do in fact leave behind multiple first-generation records. However, I did not prove this; I merely stipulated it as plausible. A sceptic might therefore push back by arguing that this isn't true of certain microscopic events - say, a particular Martian dust particle falling in the distant past - which are therefore vulnerable to backward causation. If this doubt is well-founded, what would it bode for my account?

My response is that microscopic events like these are not the sorts of ordinary events that usually appear in counterfactuals. Indeed, it looks like the degree to which an event might feasibly fail to leave behind records - and hence fail to be protected from backward causation - is the degree to which it doesn't feature in our standard causal discourse. So, even in this pessimistic scenario, my proposal would still block backward causation of ordinary minor events like splashing carps - and hence still benefit Albert's account, which at face value only avoids backward causation of major events like sunken cities.

6 Implications for Control

What does it mean for us to *control* something? It seems clear that in order to control an event, we need to cause it in the first place. At the same time, there are many things that we cause but do not control: happy accidents like roulette wins, unhappy accidents like wine spillages, and the innumerable events that we bring about through our daily activities but

aren't even aware of. This implies that whatever control amounts to, it must be some sort of a narrow species of causation. So, if it turned out that we can control the past, this would be even more surprising than if we could merely cause it.

Back in section 3, we saw that Albert defended his account by claiming backward causation is unverifiable. However, some have countered this by describing scenarios where *our awareness* of this phenomenon is a built-in feature. They then go on to argue that in such circumstances, we can not only cause past events, but *control* them. In this section I shall use the Dispersion Mechanism sketched above to argue that in these situations, backward causation probably isn't happening in the first place, so these alleged instances of backward control fall at the first hurdle.

Our first task is to pin down what it might take to control an event. There are few who doubt that causation is a necessary condition for control. The sufficient conditions, however, are less clear cut. According to Frisch (2010), these amount to 'causation plus epistemic access to the consequences'. According to Albert (2014) and Fernandes (2019), these amount to something stronger: 'causation plus profitability of the consequences'. The two examples below purport to satisfy these two sets of sufficient conditions (respectively), and thereby show how Albert's account allows situations where we can control the past. Let's look at these in turn.

The first case, from Frisch (2010), runs as follows. Suppose I'm playing the piano, and a certain melody appears in the score twice, where it's followed by two different endings. So, the score runs as follows: melody, first ending, melody, second ending. Now, during my performance, I can't remember whether I played the melody once or twice. However, I know from experience that I always play this song correctly. This means that when I play the first ending, this is a reliable record of my having played the melody once, and when I play the second ending, this is a reliable record of my having played the melody twice. Let's also assume that which ending I decide to play is

the *only* record of how many times I played the melody.

For argument's sake, imagine I play the first ending (x). This makes it highly likely that I played the melody just once (y), and so $\Pr(y \mid N_0) \approx 1$. x causes y iff the Causal Criteria are satisfied, and since Criterion 1 is probable, the key question is whether or not Criterion 2 is probable. But this *is* probable, for had I played the second ending ($\neg x$), it would be highly likely that I played the melody twice ($\neg y$), and so $\Pr(y \mid N_\star) \approx 0$. Since the Causal Criteria are probably satisfied, x probably causes y . Because there are no records relevant to y besides the antecedent itself, Albert's barrier to backward causation is overcome. Moreover, unlike in the Atlantis or carp scenarios, it's part and parcel of the situation that I'm aware of this probable causal influence. This is because the records responsible for causal influence are my own actions, and I'm aware of when I perform these. If the sufficient conditions for control are 'causation plus epistemic access to the consequences', then my decision to play the first or second ending in the present allegedly controls how many times I played the melody in the past.

Albert's (2014) response to this example is that even if we have epistemic access to the consequences of our actions, this doesn't count as control, since we can't profit from this knowledge. This brings us to Fernandes' (2019) example, which attempts to satisfy this stronger criterion. I lay this out below.

Imagine I am sat in a room that has a persistent fly buzzing round it, and there is a video camera which records its location at t_1 . Additionally, this video camera is connected to a screen, which reveals the fly's location a moment later at t_2 . This apparatus was rigged up by a sadistic scientist who upon observing the screen at t_2 gives me a reward iff the fly buzzes past my face at t_1 . But as it happens, I am an expert fly-swatter, which means my swatting at t_2 is a reliable record of a fly buzzing past my face at t_1 , and my remaining still at t_2 is a reliable record of a fly *not* buzzing past my face at t_1 . In this scenario, the only records of the fly's location at t_1 are my behaviour

and the image on the screen.

Now, let's suppose that a fly buzzes past my face (y) at t_1 , and I swat at it (x) at t_2 . Does x cause y ? Criterion 1 is satisfied by hypothesis. To probe Criterion 2, we consider the implications of me not swatting ($\neg x$) at t_2 : in all likelihood, this would mean the fly *didn't* buzz past my face ($\neg y$) at t_1 , so $\Pr(y \mid N_\star) \approx 0$. As in Frisch's example, both Causal Criteria are probably satisfied, so my swatting probably causes the fly to buzz past my face. But this time, whether or not I swat has an extra implication: it determines the image on the screen, and hence determines whether or not I get a reward. As in the piano case, my awareness of backward causation is built into the scenario. But this time, my action ($\neg x$) causes *other* present records of whether or not y occurred, from which I may profit. So, if the sufficient conditions for control are 'causation plus profitability of the consequences', then my decision to swat or not swat in the present allegedly controls the fly's location in the past.

Most discussions about whether or not these scenarios involve backward control relate to the sufficient conditions: what these are, and whether or not they're satisfied.¹² It's generally taken for granted that backward causation - a necessary condition - is in play. However, I shall deny this basic premise by appealing to the Dispersion Mechanism. This means we can discount these two alleged instances of backward causation whilst sidestepping this more in-depth debate about the sufficient conditions for control.

The above two examples share a key similarity: backward causation is allegedly in play because N_0 contains no records of the consequent that also feature in N_\star (in the piano case, this is because the only record is my action; in the fly case, this is because the only records are my action plus other records it causes). If such records existed, they would cement the matter of whether or not y occurred, preventing its probability from varying much between N_0 and N_\star and thus obstructing backward causation/control.

12. See for instance Loewer (2012), Albert (2014), and Loew (2017).

However, precisely such records generally *do* exist in the present. Most significantly, if my present actions are genuine records of what I did a few moments ago, then these cannot simply pop out of nowhere. They must have physical precursors that have persisted since that time - presumably, records that have been stored in my brain. But the Dispersion Mechanism suggests that there will probably be many other records at this time: light leaving the window, air currents, DNA traces, and so on. These are often relatively motley and unreliable, hence why we tend to ignore them when envisaging scenarios like the ones above. But they probably exist all the same. Therefore, whether or not I play the second ending or swat will probably *not* cause past events, since this will be fixed (with high likelihood) by these many records in the present. So, backward causation/control are unlikely in both scenarios.

It is tempting to respond as follows. When evaluating counterfactuals, the explicitly stated antecedent $\neg x$ isn't the *only* respect in which N_\star differs from N_0 . We almost always have to flesh out N_\star with additional facts. At the bare minimum, some gravitational fields must be redrawn to account for the fact that the antecedent entails a different matter distribution from our world - however slight.¹³ Therefore, perhaps N_\star is not just a world in which I *act* differently in the present, but also a world in which all records - including those in my brain - indicate a different past to the one indicated by N_0 .

The problem, however, is that this is asking for an awfully large gulf between N_0 and N_\star . In asking for all records of my melody-playing or the fly's location to be different, we're not just asking for a world in which I act differently in the present. We're also asking for a world in which my brain's contents are different, the light that left my window is different, the air currents are different, and so on. So, even if we grant that antecedents bring with them some extra changes (e.g. redrawn gravitational fields), a world *this* different from our own seems like a poor candidate for qualifying

13. It's on these sorts of grounds that Fernandes argues that had I not swatted, the screen would display a different image.

as N_{\star} .¹⁴

Of course, one is free to simply stipulate a scenario in which the Dispersion Mechanism fails, and there are genuinely no other records besides the ones Frisch and Fernandes describe. In these highly unusual circumstances, backward causation would indeed be likely, and my proposed obstacle to backward control would be lifted. However, we generally wouldn't know whether or not this situation obtains, since the sorts of records I've been describing are motley and unreliable. So once again, we run into a familiar distinction: even when the chance of backward causation is high (as in this instance), we'd be none the wiser, so Albert's original unverifiability defence rearises.

7 Conclusion

Albert's account allows backward causation whenever an unrecorded event lies in our past. One well-known objection is that we never observe backward causation. However, Albert's response is that this phenomenon is by its nature unverifiable, hence why we never seem to encounter it.

When we combine this account with the commonsense idea that minor events lack records in their distant futures, the result is that backward causation is ubiquitous. This is an unattractive feature of Albert's causal account when we read it at face value. To remedy this, I have proposed the Dispersion Mechanism - a process in which records seem to vanish whilst in fact proliferating.

This analysis also defends Albert's account from a second well-known objection, which is that it allows backward control. Although the sufficient conditions for this are much debated, it is widely accepted that backward causation is a necessary condition. Since my proposal obstructs this precondition, purported cases of backward control are nipped in the bud.

14. This is analogous to how Lewis (1979) defuses Fine's (1975) objection about Nixon and the button.

References

- Albert, David. 2000. *Time and Chance*. Cambridge, MA: Harvard University Press.
- . 2014. “The Sharpness of the Distinction between the Past and the Future”. In *Chance and Temporal Asymmetry*, ed. Alastair Wilson, 159–174. Oxford: Oxford University Press.
- Arntzenius, Frank. 1990. “Physics and Common Causes”. *Synthese* 82:77–96.
- Dummett, Michael. 1978. “Bringing About the Past”. In *Truth and Other Enigmas*, 333–350. Cambridge, MA: Harvard University Press.
- Fernandes, Alison. 2019. “Time, Flies, and Why We Can’t Control the Past”. In *Time’s Arrows and the Probability Structure of the World*, ed. Barry Loewer, Eric Winsberg, and Brad Weslake, forthcoming. Cambridge, MA: Harvard University Press.
- Fine, Kit. 1975. “Review of Lewis’ Counterfactuals”. *Mind* 84:451–458.
- Frisch, Mathias. 2005. *Inconsistency, Asymmetry, and Non-Locality: A Philosophical Investigation of Classical Electrodynamics*. Oxford: Oxford University Press.
- . 2007. “Causation, Counterfactuals, and Entropy”. In *Causation, Physics, and the Constitution of Reality: Russell’s Republic Revisited*, ed. Huw Price and Richard Corry, 351–395. Oxford: Oxford University Press.
- . 2010. “Does a Low-Entropy Constraint Prevent Us from Influencing the Past?”. In *Time, Chance, and Reduction: Philosophical Aspects of Statistical Mechanics*, ed. Andreas Hüttemann and Gerhard Ernst, 13–33. Cambridge: Cambridge University Press.
- . 2014. *Causal Reasoning In Physics*. Cambridge: Cambridge University Press.

- Hemmo, Meir, and Orly R Shenker. 2016. *Maxwell's Demon*. Oxford Handbooks Online. DOI: 10.1093/oxfordhb/9780199935314.013.63: Oxford University Press.
- Kim, Jaegwon. 1973. "Causes and Counterfactuals". *The Journal of Philosophy* 70 (17): 570–572.
- Kutach, Douglas N. 2002. "The Entropy Theory of Counterfactuals". *British Journal for the Philosophy of Science* 69 (1): 82–104.
- Lewis, David. 1973a. "Causation". *Journal of Philosophy* 70:556–567.
- . 1973b. *Counterfactuals*. Cambridge, MA: Harvard University Press.
- . 1979. "Counterfactual Dependence and Time's Arrow". *Noûs* 13:455–467.
- Lloyd, Seth. 1994. "Causal Asymmetries from Statistics". In *Physical Origins of Time Asymmetry*, ed. Jonathan J. Halliwell, Juan Prez-Mercader, and Wojciech Zurek, 108–116. Cambridge: Cambridge University Press.
- Loew, Christian. 2017. "The Asymmetry of Counterfactual Dependence". *Philosophy of Science* 84 (3): 436–455.
- Loewer, Barry. 2007. "Counterfactuals and the Second Law". In *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, ed. Huw Price and Richard Corry, 293–326. New York: Oxford University Press.
- . 2011. "The emergence of time's arrows and special science laws from physics". *Interface Focus* 2 (1): 13–19.
- . 2012. "Two accounts of laws and time". *Philosophical Studies* 160 (1): 115–137.
- Paul, Laurie Ann, and Ned Hall. 2013. *Causation: A User's Guide*. Oxford: Oxford University Press.

Reiss, Julian. 2015. *Causation, Evidence, and Inference*. New York and London: Routledge.

Stradis, Athamos. 2021. “Memory, the Fork Asymmetry, and the Initial State”. *Synthese* 199:9523–9547.