Causation Beyond Manipulation: Revisiting the Butterfly "Effect"

Brett Park

University of Pittsburgh, Department of History and Philosophy of Science

Abstract

Counterfactual dependence and probabilistic dependence are two criteria frequently used to analyze causation. "Mere correlations" — instances of probabilistic dependence and counterfactual independence — are a well-studied class of cases where these criteria diverge. In this essay, I provide an example of the opposite type of divergence: counterfactual dependence and probabilistic independence. The butterfly effect of chaos theory says that had a butterfly in the distant past not flapped its wings, but everything else was identical, it is possible (and indeed probable) that a present tornado would not have occurred. However, the math of chaos also tells us that whether or not the butterfly flaps its wings, the probability of the tornado is the same. I show how these two claims fit together, highlighting the distinct and unorthodox counterfactual origin of probabilistic independence in chaotic systems. Examining the case under different theories of causation, I find widespread disagreement about whether the butterfly's flap causes the tornado. I argue that this disagreement can be explained by an underlying semantic indeterminacy in our ordinary conception of causation. Rather than being exceptional, we should expect these types of relationships, and thus indeterminacies, to predominate in chaotic systems over long timescales.

1 Introduction

Difference-making accounts of causation tend to analyze causation in terms of two properties: probabilistic dependence and counterfactual dependence.[1] According to the probabilistic criteria, causes raise the probability of their effects. As a first pass, A causes B when $P(B|A) > P(B|\neg A)$. According to the counterfactual criteria, changing the cause changes whether the effect occurs. However, there is a well-known place where these criteria disagree. Cases of "mere associations" exhibit probabilistic dependence and counterfactual independence. The rooster regularly crows before the sunrise, but the sunrise is insensitive to changes in the crow. The counterfactual criteria correctly deems these cases as non-causal, and the basic probabilistic criteria must be amended, often by conditioning on relevant background conditions (Cartwright, 1979; Skyrms, 1980).

There are no well-studied examples of disagreement in the opposite direction: cases of counterfactual dependence and probabilistic independence. In this paper, I describe such a case: the butterfly effect of classical chaos theory. According to the popular example, a butterfly flaps its wings in Brazil and a tornado occurs in Texas a month later. In an otherwise identical world where the butterfly does not flap, the tornado does not occur. However, the math of chaos also tells us that whether or not the butterfly flap its wings, the probability of the tornado will be the same. The first part of the paper is an analysis of how this conjunction of properties happens in chaotic systems. Rather than being exceptional, these types of relations tend to be pervasive over long timescales in chaotic systems.

None of this tells us whether the butterfly's flap *causes* the tornado. I turn to this question in the second part of the paper. It turns out that different philosophical accounts of causation give dramatically different answers as to what causes what in chaotic systems. Rather than side with one of these accounts I attempt to locate the source of the disagreement. I argue that, along the manipulationist understanding of causation, chaotic systems exemplify many pathologies that prevent us from giving a clear causal description. Thus, whether the butterfly's flap "causes" the tornado is best thought of as being semantically indeterminate.

The paper is structured as follows. In §2, I describe the butterfly effect using the math of chaos theory, showing how probabilistic independence arises. In §3, I describe two senses in which the tornado counterfactually depends on the butterfly's flap. In §4, I argue that the butterfly effect is semantically indeterminate. I extend the discussion to causal chains in §5 and give concluding remarks in §6. A technical appendix is included in §7.

[1] For examples of the former, see Mellor (1995) and Suppes (1970). For the latter, see Lewis (1973). Another class of difference-making accounts not discussed here are regularity theories, such as Mackie's (1980) INUS account.

Note: In the rest of the paper, I will make frequent distinctions between *micro*-events and *macro*-events. A micro-event is the precise way that an event happened, and will be designated by lower case letters $a, b, c, ...$ and words $flap_1, flap_2, ...$ A macro-event is an approximate description of an event, comprised of many micro-events. Macro-events are designated by upper case letters $A, B, C...$ and words $Flap$, $Tornado$, etc.

2 The Butterfly Effect

2.1 Some Preliminaries

The enduring metaphor for chaos theory is Edward Lorenz's (1972) butterfly effect. Robert Bishop describes the butterfly effect as

> the flapping of a butterfly's wings in Argentina could cause a tornado in Texas three weeks later. By contrast, in an identical copy of the world sans the Argentinian butterfly, no such storm would have arisen in Texas (2024).

The description is a counterfactual one, resembling how philosophers talk about "actual causation" — i.e. the causal relations between actually occurring events. In future sections, we will make precise the counterfactual dependence that obtain in the case.

There is one interpretation of the counterfactual that is not consistent with the mathematics of chaos. Barry Loewer mistakes the example to mean that "whether or not a butterfly flaps its wings off the coast of Africa can make a difference to the probability of a storm occurring in the Caribbean" (2023, 24). In this section, I will give a gloss of chaos theory as it pertains to the butterfly effect. We will find out why this probabilistic reading is incorrect.[2] The discussion will be held at a largely informal, conceptual level. Additional mathematical details are provided in the appendix.

In understanding the case, I will assume that the weather is at least approximately chaotic in the ordinary mathematical sense described below, tabling some complex issues of how chaos is actually borne out in modern meteorology.[3] In

[2]This statement is also at odds with Lorenz's original talk on the butterfly effect, where he writes "If the flap of a butterfly's wings can be instrumental in generating a tornado, it can equally well be instrumental in preventing a tornado" (1972, 1)

[3]Lorenz's (1963) original weather "system" is chaotic in this sense. It is a highly simplified model of fluid convection featuring three coupled ordinary differential equations. However, the full equations for fluid dynamics, and thus the weather, are a set of partial differential equations called the Navier-Stokes equations. They add many subtleties to the case that we shall avoid. First, they are not continuously dependent on their initial conditions. This means that convergence of the initial conditions does not imply a convergence in future states. As a consequence, there is a finite time horizon beyond which future predictions cannot be extended, no matter the precision of initial data. This is a stronger form of chaos than what is discussed here and in Lorenz's 1963 paper, although it is discussed in his later

general, while chaos-like behavior is believed to be widespread in classical systems, rigorously proving an abstract dynamical system to be chaotic is a notoriously difficult exercise requiring many simplifying assumptions.[4] As usual, the inference from simple models to the real world is licensed by the fact that certain subsystems of the world behave in a way well-described by mathematical chaos. For example, we can use observational atmospheric data to estimate the exponential rate of divergence between nearby initial weather states, one of the hallmarks of chaos (Zeng et al., 1991). Other observable examples of chaos (e.g. double pendulums and various many-body systems) or quasi-chaos (e.g. dice rolls and coin flips) abound. Therefore, we can import many of the same conclusions drawn from this case there.

Finally, my analysis will be entirely restricted to the context of deterministic classical systems. Determinism is a standard assumption in both the physics literature on chaos (see Zuchowski 2017, 67) and the philosophical literature on counterfactuals (Lewis, 1979; Dorr, 2016; Loewer, 2023). It remains unclear to what extent the world's fundamental dynamics are deterministic.[5]

2.2 Dual Faces of Chaos

Chaos theory can be understood in two complementary ways. Most frequently, chaos is described in terms of sensitive dependence of initial conditions; slight differences in the initial state of a system lead to large differences in future states. Suppose we have a space of all possible states our system can take, call it Γ. This "phase space" is usually a very high dimensional space given by the possible values for all the system's degrees of freedom (e.g. positions and momenta of all the particles). The dynamics of the system will deterministically trace out a curve in Γ called a trajectory. Chaos tells us that if we have two initial states $x(0)$, $y(0)$ that start out close together in Γ, their trajectories will tend to grow farther apart at an exponential rate.[6] At some future time, they will occupy very different parts of Γ.

---

work (Lorenz, 1969). Secondly, in modern meteorological models, there is a separation of scales between the "small-scale" (spanning a few kilometers), mesoscale (tens to hundreds of kilometers), and the global scale (thousands of kilometers). Small scale error quickly multiplies, but the degree to which this error get amplified up to large scale error is highly variable (Palmer et al., 2014). Lastly, because the weather is not an isolated system, it involves forcing and damping and is only chaotic on a so called "strange attractor." More details of this are offered in §7.2 of the appendix.

[4] To get a sense for the difficulty of such an enterprise, see Yakov Sinai's (1970) seminal article.

[5] In quantum mechanics, this question largely turns on solutions to the quantum measurement problem. For a discussion of chaos in Everettian quantum mechanics, see (Wallace, 2012, 64-102). For a discussion of chaos in Bohmian mechanics, see Dürr et al. (1992). Of these two interpretations, the arguments of this paper are most easily ported to Bohmian chaos. However, it is also unsettled because we do not yet have a final fundamental theory of physics.

[6] 'Closeness' is measured in terms of a relevant phase space norm.

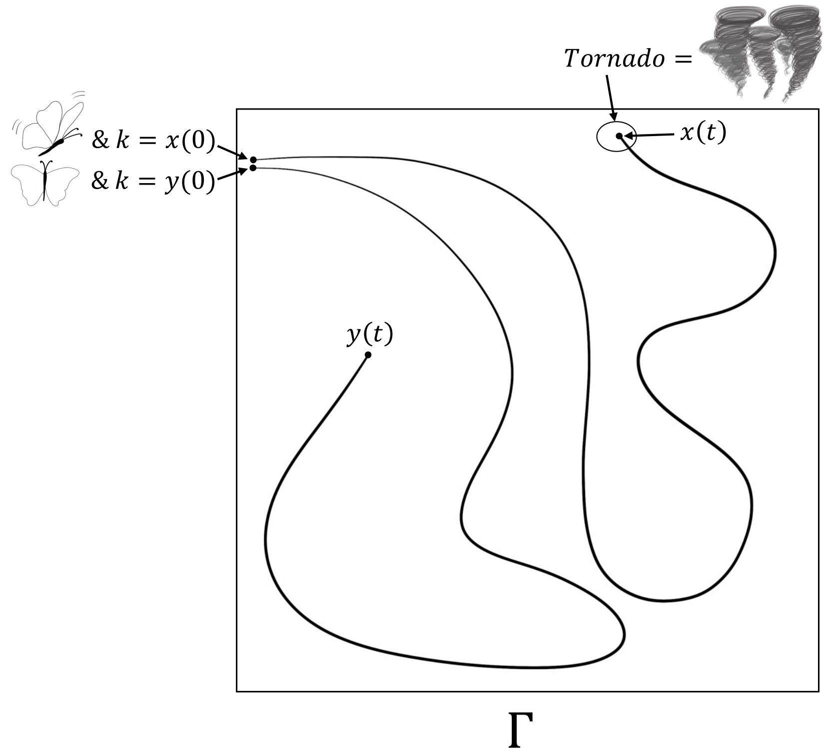

Figure 1: Sensitive Dependence on Initial Conditions in Phase Space Γ. Two initial states, $x(0)$ and $y(0)$, are identical except for a butterfly's flap. They quickly diverge under time evolution. At time t, $x(t)$ is in region associated with the macro-event of a tornado's occurrence (made up of many possible tornado configurations), and $y(t)$ ends up in a very different region of Γ. *Note: these trajectories represent the time evolution of the entire weather system, not the flight paths of the butterfly!*

For our purposes, let us take Γ to be the possible states of the global weather system. Call $x(0)$ the actual initial state. This is associated with a certain butterfly flapping in a precise way, $flap_a$, as well as the precise background conditions of the rest of the weather systems k. Call $y(0)$ some nearby initial state where the butterfly does not flap, $\neg flap_1$, that features the *same* background k (i.e. in an otherwise identical world). Finally, call $Tornado$ the macro-event of the tornado occurring at time t, comprised of many micro-realizers ($tornado_1, tornado_2...$). In Γ, $Tornado$ will be associated with a region of Γ consistent with the tornado's occurrence. Under the butterfly effect hypothesis, the dynamics will carry these initial states forward in time to $x(t)$ and $y(t)$, where we find that $x(t) \in Tornado$ and $y(t) \notin Tornado$. This is illustrated in Figure 1.

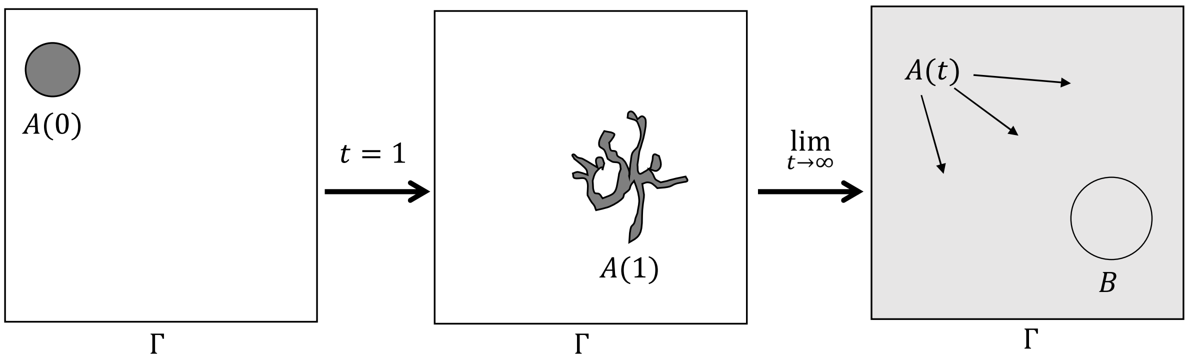

Figure 2: Mixing in Phase Space Γ. A region of initial states $A(0)$ begins to mix under the dynamics at $t = 1$, becoming fully mixed in the limit as $t \to \infty$. The lighter gray at $t \to \infty$ indicates that $A(t)$ does not change sizes, but is becoming stretched into arbitrarily thin filaments that are spread across all of phase space. It is worth taking a moment to consider how Figures 1 and 2 are complementary, where trajectory divergences in Figure 1 also describe the spreading of regions in Figure 2.

Carrying this analysis further would yield the unsurprising conclusion that counterfactual and probabilistic dependence can obtain between small antecedents and large consequents. It is also not relevant toward understanding the counterfactual at hand: "had the butterfly not flapped, the tornado would not have occurred." This statement does not uniquely specify the exact micro-state of the butterfly in the counterfactual scenario, but rather the macro-event of its non-occurrence.

When relating macro-events, chaos has a different face. While chaos is frequently characterized as the *divergence* of individual trajectories, it can also be understood as the *convergence* of probability distributions. Suppose that instead of considering an exact initial state of the system, we wanted to consider a "macro-event," $A(0)$, associated with a continuous region of exact initial states. We can we evolve $A(t)$ forward in time by evolving each of these exact initial states. If the system is chaotic, every initial $A(0)$ will have the same long-run behavior; $A(t)$ will be deformed into a highly complex filamented structure that uniformly spreads across all of Γ to greater and greater precision (Figure 2). Formally, this property is known as "mixing" because the process resembles a drop of ink being mixed into a glass of water (Gibbs, 1902, 144-145). No matter where drop originates, it will converge to the same spread-out distribution across the glass. Similarly, every initial probability distribution will spread out across Γ in the same coarse-grained fashion, approaching a single fixed distribution (with small exceptions outlined in the appendix). For this reason, chaos is sometimes identified with the condition that for all macro-events A, B:[7]

$$\lim_{t\to\infty} P(B^t|A^0) = P(B). \quad (1)$$

The prior probability $P(B)$ is obtained from the system's fixed distribution. In Figure 2, it can be roughly equated with the size of B relative to the whole Γ. The conditional probability is given by the proportion of initial conditions in $A(0)$ that end up in B at t (i.e. the relative overlap between $A(t)$ and B). (1) says that every event at time 0 grows probabilistically irrelevant for every event at t as $t \to \infty$. Here is another way to see (1); over the course of time, chaotic systems strip away all regularities between macro-events, leaving only the precise deterministic laws.[8] Even though the precise equality in (1) only holds in the infinite time limit, for practical purposes these distributions become indistinguishable in finite times.

Returning to the butterfly effect, instead of considering the exact initial state of the butterfly, we now consider a range of nearby initial states consistent with the macro-events of the butterfly flapping, $Flap$, and not flapping, $\neg Flap$.[9] Again, to probe the consequences of changing only the flap, we freeze background degrees of freedom at their precise values k.[10] Therefore, our initial macro-states are $[Flap \,\&\, k](0)$ and $[\neg Flap \,\&\, k](0)$. For large enough t, $[Flap \,\&\, k](t)$ and $[\neg Flap \,\&\, k](t)$ will both be uniformly spread out across Γ similar to Figure 2. For any macro-event at t, such as $Tornado$, the conditional probabilities given either of these initial regions will converge

[7] Werndl (2009, 215) shows that under certain common assumptions about probability in dynamical systems (see appendix), mixing is equivalent to (1). She also argues that mixing is both necessary and sufficient for chaos.

[8] This is assuming the macro-events occupy roughly equal measure in phase space. For example, in thermodynamics systems, chaotic microdynamics can lead to a reliable approach to equilibrium in the macrodynamics due the the equilibrium state occupying an overwhelming measure of phase space.

[9] $\neg Flap$ cannot be simply the set complement of $Flap$ in Γ because that will include very distant alterations of $Flap$, such as the butterfly being in Japan rather than Brazil.

[10] Barry Loewer has suggested that we use macro-background conditions to evaluate these types of counterfactuals (Loewer, 2023, 37). The problem with this approach is that it allows spurious counterfactual dependencies between dynamically unconnected events. Suppose the flap occurs on Earth and the tornado occurs on some very distant planet such that the events are spacelike separated in a relativistic spacetime (i.e. cannot influence one another). Evolving the initial macro-conditions of that planet's atmosphere forward, we will find a low probability of the tornado occurring. This would render the false judgment that had the butterfly not flapped, the distant tornado probably would not occur, even though the flap cannot influence the tornado at all. As such, the way I am analyzing counterfactuals is much closer in spirit to Lewis' (1979) account in that it heavily favors perfect match between worlds over approximate match.

to that event's prior probability:

$$\underbrace{P(Tornado^t | Flap^0 \,\&\, k^0)}_{\approx P(Tornado)} \approx \underbrace{P(Tornado^t | \neg Flap^0 \,\&\, k^0)}_{\approx P(Tornado)} . \tag{2}$$

The details of (2) are elaborated in the appendix. This tells us that, contra Loewer, *whether* the tornado occurs is probabilistically independent of *whether* the butterfly flaps.[11] Importantly, (2) will hold for virtually all macro-events at sufficiently distant times. The right hand side of (2) also reveals the distinctive counterfactual behavior for chaotic systems that will guide the rest of our discussion; *change virtually any macro-event in the sufficiently distant past, and you effectively re-roll the dice for all present macro-events.* The probability that any given event occurs will be in line with its overall objective probability; high probability events probably will happen, low probability events probably won't.

Here is a way to visualize the situation (Figure 3). At our initial time, separate Γ into the degrees of freedom associated with the butterfly $\Gamma_{butterfly}$ and the degrees of freedom associated with the rest of the system $\Gamma_{background}$ such that $\Gamma = \Gamma_{butterfly} \times \Gamma_{background}$, where $\times$ is the Cartesian product. $\Gamma_{background}$ is fixed at their actual values k, while $Flap$ and $\neg Flap$ are associated with regions of initial conditions in $\Gamma_{butterfly}$. Inspecting $\Gamma_{butterfly}$, each initial condition of the flap either leads to $Tornado$ occurring (light grey) or $Tornado$ not occurring (dark grey). You can imagine the initial regions as looking like a cookie with sprinkles, where the doughy regions are heading to $\neg Tornado$ at t and the sprinkled regions are heading to $Tornado$. The sprinkles are numerous, tiny, and randomly placed throughout both cookies (much more finely intermixed than illustrated). (2) says that the proportion of initial conditions in $Flap$ and $\neg Flap$ leading to $Tornado$ is approximately the same; they both approximate the relative size of $Tornado$ in Γ. This will be true even if we consider non-uniform distributions over $Flap$ or $\neg Flap$, so long as they are not allowed to vary too fast in a way that is pathologically tailored towards a certain outcome.[12]

[11] It might be objected that full probabilistic independence requires $=$'s and not $\approx$'s in (2). However, exact independence rarely, if ever, occurs in real world data. Therefore, (2) is consistent with the probabilistic independence we use in real-world applications. Furthermore, (2) grows arbitrarily close to equality as t increases.

[12] This is Henri Poincaré (1905, 224-226) celebrated method of arbitrary functions. For more detailed discussions, see Myrvold (2021) and Strevens (2011).

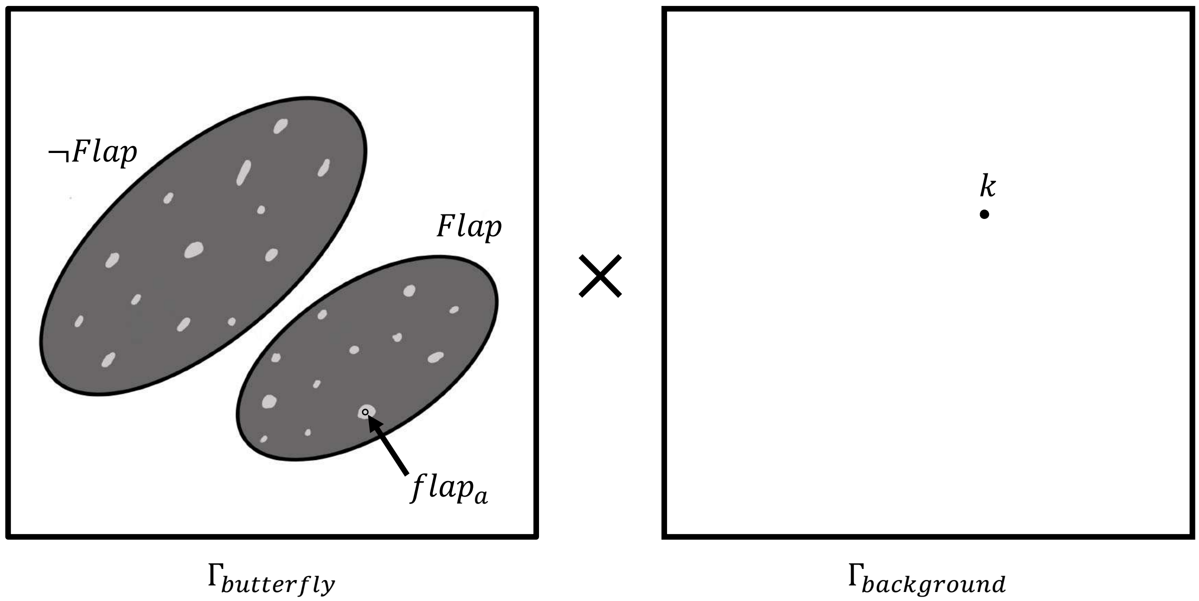

Figure 3: Phase Space Sketch of the Butterfly Effect at time 0. Background conditions are frozen at their actual values k. In the $Flap$ and $\neg Flap$ regions, the light grey initial conditions end up in $Tornado$ at t and the dark grey initial conditions end up in $\neg Tornado$. $flap_a$ represents the actual initial conditions of the flap. $Flap$ and $Tornado$ are probabilistically independent because $Flap$ and $\neg Flap$ have the same proportion of initial conditions heading to $Tornado$. See Diaconis et al. (2007, 216) for similar figure for coin flips.

3 Counterfactual Dependence

In the previous section, we described how probabilistic independence occurs in the butterfly effect due to chaos. This leads to a "whether-whether" probabilistic independence; whether $Flap$ occurs does not change the probability of whether $Tornado$ occurs, even given precise background conditions. In this section, I will describe how the tornado counterfactually depends on the flap.

It should be noted that counterfactual dependence has been defined many times over and not every definition will agree about the case. Indeed, some authors tie counterfactual dependence to chance raising, making the conjunction of probabilistic independence and counterfactual dependence an analytic impossibility.[13] However, one does not have to look far for counterfactual analyses that yield the conclusion that the tornado counterfactually depends on the flap. In this section, I will point to two distinct, non-trivial ways in which the tornado exhibits counterfactual dependence on

[13]See, for example, David Lewis (1986, 175-184) and Chrisopher Hitchcock (2003, 140)

the flap, despite the two events being uncorrelated. The first is a type of "how-whether" dependence, and the second is a "whether-whether" dependence.

3.1 Counterfactual Sensitivity

In his 2000 counterfactual account of causation, David Lewis presented the notion of "influence." A influences B if changing the specifics of A would have led to a large change in the specifics of B. This is a type of "how-how" counterfactual dependence, the specifics of how A occurs (or does not occur) change how B occurs (or does not occur). In our example, *Flap* will have a considerable influence over $Tornado$ insofar as the specifics of the tornado, and indeed whether it occurs at all, are sensitive to very precise specifications of the butterfly's flap. In this section, I am going to present a modified and stronger notion of "how-whether" counterfactual dependence, modeled on Lewis' notion of influence, that also applies to the case. The reason for this will become apparent in the next section; it gives us the resources to describe two very different counterfactual origins of probabilistic dependence.

Assume that in the actual world, *Flap* and $Tornado$ occur, with $flap_a$ being the precise initial conditions of the flap. Because the system is sensitively dependent on initial conditions, there will be many initial conditions nearby to $flap_a$ which will deterministically evolve into $\neg Tornado$. Thus, $Tornado$ is counterfactually sensitive to changes in $Flap$, where:

> **Counterfactual Sensitivity**: If A and B are two actually occurring events, and A is realized by a_a, then B is *counterfactually sensitive* to A iff there exists a wide range a_1, a_2,... of distinct but not-too-distant alterations of A such that if a_1 or a_2 or ... had occurred, then B would not have occurred.

Again, this can be described as "how-whether" counterfactual dependence: the precise details of *how* A occurs can change *whether* B occurs. For dynamical systems, counterfactual sensitivity tracks the following question. Are there a non-negligible amount of initial conditions in either A or $\neg A$ which lead to $\neg B$? If so, then B is counterfactually sensitive to A. What counts as negligible will be a contextually determined, but for our purposes just assume that it is some very small proportion of the initial conditions. As we can see from Figure 3, $Tornado$ is counterfactually sensitive to $Flap$. This is true even though there is not a systematic connection between the macro-events (e.g. $Flap$ being mostly one color and $\neg Flap$ being mostly the other).

Counterfactual sensitivity allows us to say something general about chaotic systems. As we have mentioned before, changing virtually any macro-event in the distant past would reroll the dice on the system's macro-future.[14] So long as the future

[14] This is assuming the macro-events are not extremely gerry-mandered.

event's objective chance is above some small value, then it will be counterfactually sensitive to nearly all macro-events at some sufficiently distant past time. For example, take *Sun* to be the event of sunshine in some patch of the Mojave desert at t, assumed to be a very probable event. Even this event will be counterfactually sensitive to changes in the distant past. For virtually all macro-events A at time 0, there will be a small portion of initial conditions in $A \,\&\, k$ and $\neg A \,\&\, k$ which lead to $\neg Sun$.

We have a situation that might be described as *counterfactual rampancy*: nearly all macro-events in the future are counterfactually sensitive to nearly all macro-events in the distant past.[15] Change virtually anything in the distant past, and virtually anything in the present might not occur. However, the changes made to the present would be effectively random; the probabilities for every present event would fall in line with their objective probability. Thus, this rampant counterfactual sensitivity cannot be used to tilt future distributions one way or another. We cannot fix the initial weather state precisely enough to raise or lower the probability of the tornado, and even if we did have god-like abilities of control we would also need god-like abilities of prediction to know which initial states lead to which outcomes.

3.2 Two Routes to Probabilistic Independence

We have just seen that *Tornado* is counterfactually sensitive to *Flap*, even though the two events are probabilistically independent. It is worth pausing for a moment to contrast this type of probabilistic independence with more "standard" cases. Here is a test. Imagine finding out that two events A and B are probabilistically independent, given precise background conditions k. My guess is that your default assumption is that the occurrence of the second event will be insensitive to changes in the first.

Here is an example. Suppose that just before the tornado touches down in Texas at t, John sneezes in London at $t-1$. The sneeze might exchange some minuscule conserved quantities with the impending tornado, but certainly not enough to change whether it occurs. Thus, whether John sneezes is probabilistically independent of whether the tornado occurs, given the background conditions. As we did with the butterfly flap, if we represented the weather system's phase space as $\Gamma = \Gamma_{John} \times \Gamma_{background}$, fixed the degrees of freedom in $\Gamma_{background}$ at their actual values k, and traced out the relevant *Sneeze* and $\neg$*Sneeze* regions, we would find *Sneeze* does not raise the probability of *Tornado*:

$$P(Tornado^t | Sneeze^{t-1} \,\&\, k^{t-1}) = P(Tornado^t | \neg Sneeze^{t-1} \,\&\, k^{t-1}). \tag{3}$$

[15] Counterfactual rampancy is something that distinguishes genuinely chaotic systems from quasi-chaotic ones such as coin tosses and dice rolls. Quasi-chaotic systems are sensitive to their initial conditions, but not arbitrarily so, and they are not mixing over their state spaces.

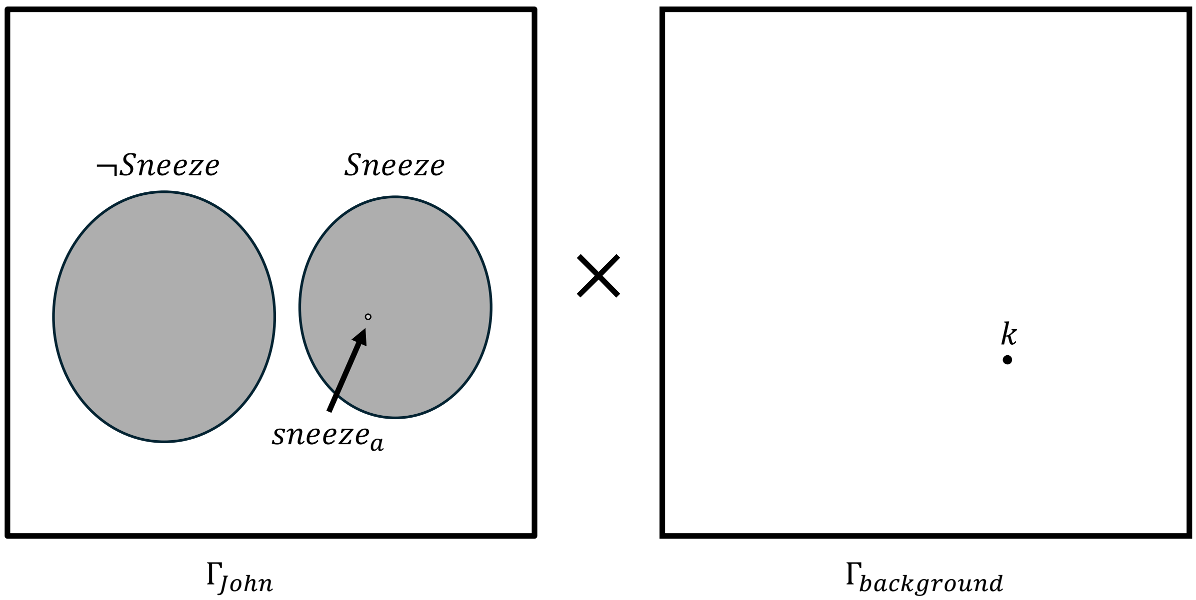

Figure 4: Phase Space Sketch of the sneeze at $t - 1$. Background conditions are frozen at their actual values k. In the *Sneeze* and ¬*Sneeze* regions, the light grey initial conditions end up in $Tornado$ at t and the dark grey initial conditions end up in $\neg Tornado$ (there aren't any). The actual initial conditions of the sneeze are represented by $sneeze_a$. *Sneeze* and $Tornado$ are probabilistically independent because *Sneeze* and ¬*Sneeze* have the same proportion of initial conditions heading to $Tornado$ (all of them).

But (3) has a much different origin than (2). (3) holds because *all* initial conditions in $Sneeze\,\&\,k$ and $\neg Sneeze\,\&\,k$ end up in $Tornado$ at t. That is, $Tornado$ is counterfactually **insensitive** to *Sneeze*. The contrast between this case and the butterfly effect is obvious when we compare Figure 4 with Figure 3.

Counterfactual sensitivity allows us to articulate two distinct counterfactual routes to probabilistic independence. Along the "typical" route, there are no relevant changes to the first event that would make the second event not occur. The second event is counterfactually **insensitive** to the first. However, sensitive dependence on initial conditions provides another path. This path can be characterized as counterfactual **hypersensitivity**: *there are many changes to the first event that would lead to second event not occurring, but these changes become too fine-grained to be picked up as a correlation between macro-events.* Thus, probabilistic independence contains a degeneracy of counterfactual possibilities. Merely knowing that two events are probabilistically independent, we do not know whether changing the first can change the second's occurrence.

3.3 Actual Counterfactual Dependence

So far, we have seen that the butterfly effect conjoins "how-whether" counterfactual dependence with "whether-whether" probabilistic independence. Now we will extend the analysis to include "whether-whether" counterfactual dependence.

Again, assume that in the actual world a certain Brazilian butterfly flaps and a certain Texas tornado occurs sometime much later. Given the previous discussion, we expect that changing some macro-event like $Flap$ will effectively re-roll the dice on the system's future; the probability distribution over future events will be close to their prior probabilities. The probability of a tornado occurring within a small geographic area at a specific time is going to be extremely low, so we should expect that $P(Tornado) \approx 0$. Therefore, had the butterfly not flapped it is *overwhelmingly likely* that the tornado would not have occurred, because chaos ensures that $P(\neg Tornado^t | \neg Flap^0 \,\&\, k^0) \approx P(\neg Tornado) \approx 1$. As we will see below, most ordinary counterfactuals with false antecedents must be justified by reference to overwhelming likelihood. Thus, the butterfly effect answers to the following notion of counterfactual dependence:

> **Actual Counterfactual Dependence**: If A and B are actual events, B counterfactually depends on A iff had A not occurred, B *very probably* would not have occurred.

In the rest of this section, I will show why this definition has become more or less the standard in the literature on counterfactual conditionals.

It is commonly assumed that B is counterfactually dependent on A iff

(I) $A \mathrel{\Box\!\!\rightarrow} B$

(II) $\neg A \mathrel{\Box\!\!\rightarrow} \neg B$.

where $A \mathrel{\Box\!\!\rightarrow} B$ is the counterfactual conditional "If A were the case, then B would be the case." Another standard assumption in the analysis of counterfactuals — see, for example, Stalnaker (1968, 104), Lewis (1973, 26), and Edgington (1995, 290) — is that when A and B occur in the actual world, then **(I)** is automatically true: $(A \,\&\, B) \rightarrow (A \mathrel{\Box\!\!\rightarrow} B)$.[16] Thus, when A and B are actual events, A counterfactually

[16] This principle is known as "Conjunction Conditionalization" **(CC)**. While widely adopted in the literature on counterfactuals (see Table 2 in Starr, 2021), **(CC)** has been the subject of debate. Numerous objections have been raised against **(CC)** (Bennett, 1974; McDermott, 2007; Woodward, 2023), citing unintuitive cases that follow from the principle. However, Walters and Williams (2013) have made a compelling case that **(CC)** is no idle wheel in the logic of conditionals. They show that various combinations of less contentious logical principles for conditionals entail **(CC)**, and so rejecting it is "no minor surgery, but a complicated operation involving the mutilation of far more entrenched aspects of the logic of conditionals" (2013, 584). Furthermore, Walters (2016) suggests that the unintuitive cases can be

depends on B iff $\neg A \,\Box\!\!\rightarrow \neg B$. Consequently, assuming that $Flap$ and $Tornado$ are actual events, whether $Tornado$ counterfactually depends on $Flap$ comes down to whether $\neg Flap \,\Box\!\!\rightarrow \neg Tornado$.

How can we determine the truth of $\neg Flap \,\Box\!\!\rightarrow \neg Tornado$? David Lewis analyzed counterfactuals with respect to a similarity ordering of possible worlds. According to Lewis, there will be many worlds that are equally similar to ours where $\neg Flap$ is true. $\neg Flap \,\Box\!\!\rightarrow \neg Tornado$ is true when all of these $\neg Flap$ worlds are also $\neg Tornado$ worlds.[17]

However, Lewis' definition turns out to be much too stringent. The world is objectively chancy. For one, quantum mechanics might imply that the dynamics of the world are chancy and not deterministic. But even if the world is fundamentally deterministic, determinism of the microphysics almost never implies determinism of the macrophysics. Accordingly, specifying some macrophysical change in the past is rarely sufficient to uniquely determine the occurrence of macrophysical events in the future. However, I would seem right to assert that

If I were to leave an ice cube on the counter, it would melt. (4)

despite the fact that in a small portion of the relevant initial conditions, the ice cube gets colder. Since the time of Boltzmann, we have known that even our most secure macroscopic regularities can only be probabilistically underwritten by the underlying microphysics.[18] As numerous authors have pointed out, if we follow Lewis in thinking that there are many closest possible worlds, and we take $A \,\Box\!\!\rightarrow B$ to mean that B is true in **all** closest A worlds, then almost all counterfactuals we utter in ordinary conversation will turn out to be false (Hoefer 2004, 107-110, Stalnaker 2019, 196, and Edgington 2008, 14). Thus, there must be a link between probabilities and counterfactuals that grounds our use of counterfactual conditionals in physics while respecting our ordinary judgments.

There are a number of different candidate proposals for how to square our use of

explained by the fact that if **(CC)** is true, then $A \,\&\, B$ is logically stronger than $A \,\Box\!\!\rightarrow B$. Under Grice's maxim of quantity, asserting $A \,\Box\!\!\rightarrow B$ typically carries a conversational implicature that A did not happen. Likewise, Lewis writes "It is conversationally inappropriate, of course, to use the counterfactual construction unless one supposes the antecendent false; but this defect is not a matter of truth conditions" (1986, 18). Thus, counterexamples to **(CC)** can be successfully explained pragmatically, while keeping the usual semantics of counterfactuals intact.

[17] Lewis' adds a slight wrinkle to this because because he believes it is possible for worlds to get indefinitely more similar to w. Accordingly, he says that $A \,\Box\!\!\rightarrow B$ is true at world w iff there is an $A \,\&\, B$ world more similar to w than any $A \,\&\, \neg B$ world. We will ignore this wrinkle for the purposes of exposition.

[18] The second law of thermodynamics has this character; there are always small, nonzero measure regions of initial conditions where entropy will decrease (Albert, 2000).

counterfactual conditionals with objective chances. Alan Hájek believes that most counterfactuals we utter are strictly false, although we can use high chance to "legitimately assert" them in the forgiving context of ordinary conversation (Hájek, 2014, 87). Robert Stalnaker famously asserted that the probability of a conditional being true equals the conditional probability of $P(B|A)$, an idea that has generated enormous debate (1970, 75).[19] Dorothy Edgington has argued that when the antecedent of a counterfactual is false, it has no truth value (2008). Instead, its acceptability goes by the conditional probability of B, given A. On all of these proposals, $\neg Flap \mathrel{\Box\!\!\rightarrow} \neg Tornado$ is on par with most counterfactual scenarios we routinely assert and accept in daily discourse; as expressions of the highly probable consequences of minimal changes to the actual world.

In summary, $Tornado$ counterfactually depends on $Flap$ in the same way we ordinarily evaluate counterfactuals dependence among actual events. Counterfactuals answer the question; how would the world be different if A had been true? If A is true in actuality, the answer is "no different." If A is false, the answer must be grounded in the probable, not inevitable, consequences that would follow were it true.[20] If we live in a world where a butterfly flapped its wings in the distant past and a tornado just occurred, then chaos tells us that had the butterfly not flapped its wings, the tornado very probably would not have occurred.

4 Causation

Here is a summary of the previous sections' claim that the butterfly effect exhibits both counterfactual dependence and probabilistic independence:

[19] This has come to be known as "Stalnaker's Thesis." Despite its initial appeal, David Lewis (1986, 136-139) famously put forward a series of results which show that if we allow for arbitrary nesting of $>$'s, there is no proposition $A > B$ that can satisfy Stalnaker's Thesis without trivializing the probability distribution in question. Lewis' results led to more triviality results (Carlstrom and Hill, 1978; Hájek, 2011). See chapter 4 and 5 of Bennett (2003) for a review of this history. Despite this, it has also been shown by van Fraassen (1976) and Bacon (2015) that a contextualized version of Stalnaker's Thesis does hold in a Stalnakerian semantic framework. Finally, subsequent literature has shown that the triviality results for Stalnaker's Thesis are actually a problem for connecting probability with epistemic modality, generally (Goldstein and Santorio, 2021).

[20] It may feel strange that we are allowed to use facts about the actual world to justify $Flap \mathrel{\Box\!\!\rightarrow} Tornado$, but need to rely on probabilities to justify $\neg Flap \mathrel{\Box\!\!\rightarrow} \neg Tornado$. $Flap \mathrel{\Box\!\!\rightarrow} Tornado$ would turn out false if we evaluated it probabilistically since $P(Tornado^t|Flap^0 \,\&\, k^0) \approx 0$. However, counterfactuals are intended to be *counter-to-the-facts*-uals. Only when their antecedents are false do we need to start looking at worlds beyond our own to ascertain truth conditions. Once we are looking beyond our world, we are confronted with the problem of objective chance. When dealing with a similar issue, Jonathan Bennett writes "It is expectable that the conceptual structures we use in subjunctive conditionals should make a special case of the actual world" (2003, 250).

1. **Probabilistic Independence:** Where $Flap$ and $Tornado$ are macro-events, given precise background conditions k for the rest of the weather system, conditioning on whether $Flap$ or $\neg Flap$ occurs does not raise or lower the probability that $Tornado$ occurs.

2. **Counterfactual Sensitivity:** If $Flap$ and $Tornado$ are actual events and $Flap$ is realized by $flap_a$, then there are many nearby initial conditions $flap_1, flap_2, ..., \neg flap_1, \neg flap_2, ...$ that lead to $\neg Tornado$.

3. **Actual Counterfactual Dependence:** If $Flap$ and $Tornado$ are actual events, had $Flap$ not occurred, $Tornado$ very probably would not have occurred.

Thus, we have "whether-whether" probabilistic independence, "how-whether" counterfactual dependence, and "whether-whether" counterfactual dependence all coexisting in a single case. None of these judgments have not been cut from whole cloth. The probabilities are supplied by dynamical systems theory and the analysis of counterfactuals is in line with mainstream approaches to the subject: counterfactual sensitivity is modeled after David Lewis' (2000) notion of influence, and actual counterfactual dependence fits within the frameworks offered by Stalnaker (1968; 1970) and Edgington (2008).

So far, I have not addressed the question of whether the butterfly's flap *causes* the tornado. But the butterfly "effect" is a patently causal claim. Since its introduction by Edward Lorenz (1972), it is routinely described in the philosophy and physics literature in explicitly causal terms (see Smith 1990, 247, Frisch 2014, 212, Bricmont 2022, 83, and Hilborn 2004).[21] When these authors describe the case, it is a simple binary relationship — the flap of the butterfly's wings causes the tornado. Can we make sense of this with any existing philosophical accounts of causation?

4.1 Probabilistic and Counterfactual Theories of Causation

I will focus on two broad types of causal theories — probabilistic and counterfactual. The primary insight of the probabilistic camp is that causes tend to raise the probability or chances of their effects. The fact that smoking causes lung cancer is evidenced by the fact that smoking raises one's chance of getting the disease. The insight of the counterfactual theory is that the second event changes when the first event changes.

Probabilistic or chance raising analyses of causation straightforwardly rule out the possibility that $Flap$ causes $Tornado$. The standard proposal is some variation of

$$A \text{ causes } B \text{ iff } P(B|A \cap K) > P(B|\neg A \cap K) \tag{5}$$

[21] This is not an appeal to the authority of these philosophers' and physicists' intuitions on *causation*. Rather, it is a way of excluding the intuitions of laypeople who have only ever encountered the butterfly effect as a causal claim and have no understanding of the underlying dynamics from which it originates.

for some background conditions K.[22] Chaos ensures that $Flap$ will not raise the probability of $Tornado$ for precise k or imprecise K background conditions. As $t \to \infty$, only conditioning on precise micro-specifications of the flap $\{flap_1, flap_2, ...\neg flap_1, \neg flap_2\}$ plus precise background conditions k, could you raise or lower the probability of $Tornado$.

Additionally, the case will elude accounts that use probability raising to analyze counterfactual dependence. For example, Christopher Hitchcock describes counterfactual dependence as

> the actual chance of e's occurrence, Ch(e), at the time of c's occurrence, is higher than Ch(e) would have been, at the same time, had c not occurred (Hitchcock, 2003, 140).

According to such a chance-raising account of counterfactual dependence, $Tornado$ at t will not counterfactually depend on $Flap$ time 0.

The most a chance-raising analysis can say is that $Tornado$ causally depends on $flap_a \,\&\, k$, the actual micro-conditions of the flap together with the precise background conditions.[23] On such a reading, we might say that causation reduces to micro-determinism in chaotic systems at $t \to \infty$. The macro-present causally depends

[22]See Suppes (1970) and Mellor (1995) for two probabilistic accounts of causation

[23]A probabilistic theory of causation would count the case as probabilistically and causally dependent if it compared the conditional probabilities given the actual microstate of the flap with those given the macrostate of the flap not occurring

$$\underbrace{P(Tornado^t|flap^0_a \,\&\, k^0)}_{1} > \underbrace{P(Tornado^t|\neg Flap^0 \,\&\, k^0)}_{\approx P(Tornado) \ll 1}.$$

Determinism implies that that the left hand side equals 1 or 0 (in this case 1), while chaos ensures the right hand side equals the prior probability of the tornado. Thus, a probabilistic theory of causation could go either way, depending on what initial events we are contrasting.

First, we should note that this is contrasting two levels of analysis, micro and macro, when actual counterfactual dependence occurs at the macro-physical level. For counterfactual dependence, we have a natural explanation for why the probabilistic symmetry between $Flap$ and $\neg Flap$ is broken, one of them occurs in the actual world while the other does not. We could try to recreate this outcome using a hybrid micro-macro account, but we are then comparing the criteria for causation at different event-levels for the two accounts.

More importantly, this strategy would not offer an independently viable conception of either probabilistic dependence or causal dependence. We are comparing probabilities conditioned on a precise microstate, either 0 or 1, with probabilities conditioned on a macrostate, which will generally be somewhere between 0 and 1. On this hybrid analysis, probabilistic independence will only show up in circumstances of deterministic macro-physics. For example, suppose I can choose between two coins to flip, a penny and a nickel. I pick the penny and flip heads. The outcome should be probabilistically independent of my choice in coin, with the probability of heads being 1/2 for each. However, if I am comparing the micro-event of the penny flip, *penny*, with the macro-event of nickel flip, $Nickel$, then this analyses will say that the outcome is "probabilistically dependent" on the choice in coin: $P(Heads^t|penny^0) = 1$ while on micro-events in the distant past but not macro-events.[24] While this would be provocative conclusion, it is too sure-footed. The idea that in chaotic systems there are no macro-causes of the present in the distant past is indeed surprising because we would expect there to be many. The deeper problem is that this approach treats counterfactual insensitivity and counterfactual hypersensitivity as causally on a par. At the very least, the notion that $Flap$ is not a cause of $Tornado$ (Figure 3) is not as obvious as the notion that $Sneeze$ is not a cause of $Tornado$ (Figure 4).

A strictly counterfactual theory of causation could explain why one *might* be inclined to say that the butterfly's flapping can cause the tornado's occurrence. Lewis' original 1973 account has this flavor. According to Lewis, B causally depends on A iff B counterfactually depends on A: $A \mathrel{\Box\!\!\rightarrow} B$ and $\neg A \mathrel{\Box\!\!\rightarrow} \neg B$ (Lewis, 1973, 563). If A and B are actual events, then $A \mathrel{\Box\!\!\rightarrow} B$ is automatically true and B causally depends on A iff $\neg A \mathrel{\Box\!\!\rightarrow} \neg B$. Furthermore, A causes B iff there is a chain of causal dependence between A and B (causal dependence is sufficient but not necessary for causation). We will return to the issue of causal chains in §5, but for now let us just consider whether $Tornado$ causally depends on $Flap$. As we have mentioned before, Lewis requires that $\neg B$ holds in *all* closest $\neg A$ worlds. However, given the prevalence of objective chance, it would be hard to find a single case of causal dependence using his semantics for counterfactuals. Here are two modifications to the counterfactual semantics worth entertaining:

> **(Prob-Lewis)**: a version of Lewis' 1973 theory where high probability can license $\neg A \mathrel{\Box\!\!\rightarrow} \neg B$. $Tornado$ causally depends on $Flap$ due to actual counterfactual dependence. So too would all low-probability macro weather events that co-occur with $Tornado$.
>
> **(Stal-Lewis)**: a pairing of Lewis' 1973 theory of causation and Stalnaker's (1968) counterfactual semantics. Stalnaker (1968, 104) holds that there will be a unique closest $\neg Flap$ world. If this world is a $\neg Tornado$ world, as it is very likely to be, then $Flap$ causes $Tornado$. Additionally, counterfactual sensitivity entails that this world might be a $\neg Tornado$ world (Stalnaker (1981) understands "might" in terms of epistemic possibility), and thus it

$P(Heads^t|Nickel^0) = 1/2$.

The direction of probabilistic dependence will also frequently misalign with expectations. For example, suppose I decide to ride my motorcycle to work today instead of the less dangerous option of taking the bus. If, given the exact initial conditions for my ride, I arrive safely, then riding my motorcycle (supposedly) raised the probability of me getting to work safely. Attaching a theory of causation to this analysis of probabilistic dependence will yield similarly strange judgments, such as my decision to smoke causing me to *not* develop lung cancer merely because I, in actuality, took up smoking and did not develop lung cancer (supposing this is a microdeterministic process).

[24] Or at least no macro-events on any natural coarse-graining of phase space.

> will be robustly the case that $Flap$ not occurring might have caused $Tornado$ to not occur.

Accordingly, simple modifications to Lewis' theory can support the causal interpretation of the butterfly effect. Problems with Lewis' 1973 theory tend to crop up on the necessity side; cases of late preemption and overdetermination show us that causation does not imply counterfactual dependence. However, counterfactual dependence is still thought to be sufficient, or nearly sufficient, for causation (Paul and Hall, 2013, 4).[25] Thus, we have our first solid support for the idea that the $Flap$ is a cause of $Tornado$.

The butterfly effect also shows up in Lewis' (2000) revised "influence" analysis, which relies on a "how-how" counterfactual dependence. Recall that according to this account, A influences B iff, roughly, nearby changes in the specifics of A (including changes where A still occurs and changes where A does not occur) leads to substantial changes in the specifics of B. Similar to before, B causally depends on A iff A influences B, and A causes B if there is a chain of stepwise influence from A to B. To illustrate, Lewis offers a helpful analogy:

> Think of influence this way. First, you come upon a complicated machine, and you want to find out which bits are connected to which others. So you wiggle first one bit and then another, and each time you see what else wiggles. Next, you come upon a complicated arrangement of events in space and time. You can't wiggle an event: it is where it is in space and time, there's nothing you can do about that. But if you had an oracle to tell you which counterfactuals were true, you could in a sense "wiggle" the events; it's just that you have different counterfactual situations rather than different successive actual locations. But again, seeing what else "wiggles" when you "wiggle" one or another event tells you which ones are causally connected to which (2004, 91).

In the weather system, wiggling any macro event at time 0 would significantly wiggle all weather events at t. Recall that there is widespread counterfactual sensitivity within the system, and counterfactual sensitivity is, in general, a stronger condition than influence. Thus, everything would be causally connected according to Lewis. But this wiggling would by highly unsystematic. For example, moving between initial conditions in $Flap$ and $\neg Flap$, we would see vast changes in what occurs at t, but these changes would not covary with $Flap$ or $\neg Flap$. Over the course of many wiggles, the relative frequency of each event at t would match its objective probability. Again, this is very different than $Sneeze$ at $t-1$, where wiggling would produce few notable changes at t.

[25] One potential counterexample on the sufficiency front is cases of double prevention, although their status remains controversial.

In summary, on certain modified versions of Lewis' 1973 counterfactual theory of causation, $Flap$ at time 0 causes $Tornado$ at t. On his 2000 influence account, practically all meteorological events at time 0 are causally connected with practically all meteorological events at t. According to probabilistic theories of causation, $Flap$ cannot cause $Tornado$ because it does not raise the probability of $Tornado$. However, the micro-conditions of the flap — $flap_a$ — could count as causing $Tornado$, if supplemented with precise background conditions. Evidently, there is extensive disagreement about this case. How should we proceed?

4.2 Causation and Manipulation

We might be tempted to think that, with the butterfly effect, science has furnished a counterexample to probability-raising accounts of causation. Or we might want to say that, given our best contemporary accounts of causation use probability raising, we were mistaken to ever believe that butterfly flaps can cause tornadoes. Yet another direction, the one I propose we take, is to step back and think not about what causation *is*, but what causation is *for*, and how the butterfly effect fits into that picture. This is the strategy adopted by "manipulationist" or "interventionist" theories of causation, exemplified by James Woodward (2003; 2021). I turn to this way of thinking not as yet another a way of adjudicating the case, but as a way of understanding the source of the disagreement between the aforementioned theories of causation.

On the manipulationists reading, our understanding of causation should account for why causation is an incredibly useful concept to beings like us. In everyday and scientific inquiry, we have pragmatic aims when identifying causal dependencies. It is not enough that we know what correlates with what, we want what Nancy Cartwright calls "effective strategies": ways to manipulate our environment to bring about desired ends (1979, 419). The rooster's crow regularly precedes the sunrise, but we cannot mute the rooster to prevent the sunrise. For Woodward, once we have observed a correlation, we are tasked with understanding whether and how this correlation can be "exploited for purposes of manipulation and control" (2007, 72). In short, the purpose of causation is to enable us to successfully manipulate the world in ways that tilt the future in our favor.

Since we are macroscopic beings, we are, by and large, only capable of observing macroscopic differences and implementing macroscopic interventions. Thus, we routinely rely on the assumption that the macroscopic, coarse-grained behavior of a target system is not dependent on the details of its microscopic realizers (Woodward 2007, 80; Weinberger et al. Forthcoming, 20). This assumption breaks down in chaotic systems in spectacular fashion. Even the smallest differences in initial states lead to macroscopically large future differences. Thus, even if $flap_a$ causes $Tornado$ on a probability-raising account, it is not the type of causal relationship that can be exploited by us. The types of variables we might hope to control, such as $\{Flap, \neg Flap\}$, cannot help us prevent future weather disasters because they yield the same future probability distributions.

Another property of helpful causal generalizations is their *stability* under various background conditions (Woodward, 2007, 77). Causal dependences are useful when they apply to a large range of relevant situations. It would be useless, for example, to find a causal relationship that only holds when the rest of the universe is in a very precise state, a state we are never going to see again. Something approximating this obtains in chaotic systems. Whether $flap_a$ "causes" $Tornado$ is dependent on the fine-grained details of the rest of the atmosphere. It also depends on even the smallest outside perturbations. Michael Berry provides a calculation for a simple chaotic system: two classical oxygen atoms colliding in a closed container under the gravitational effect of an electron at the edge of the observable universe (Berry, 1978, 95-96). By their 56th collision, the gravitational effect will have noticeably altered the location of their collisions. For turbulent fluids, these microphysical differences are hypothesized to rapidly cascade up to macrophysical differences (Bandak et al., 2024). Thus, we could generate a "butterfly effect" between a tornado and even the smallest and remotest influences that have reached us. Clearly, as timescales increase, probability-raising causal relationships in chaotic systems become increasingly unstable, and, therefore, useless.

So where does this leave us? Macro-events like $Flap$ cannot be used to tilt future distributions in our favor, and micro-events like $flap_a$ are physically inaccessible and require too much knowledge of background conditions to be of any use. Therefore, the relation between the flap and the tornado is too delicate and unpredictable for the purposes of manipulation and control. Despite this, certain counterfactual dependencies can be drawn from the case, and these have been enough for various authors to describe it in causal terms. How might we explain any lingering sense that there is something causal going on here?

4.3 Indeterminacy of 'Cause'

A recent paper by Woodward with Naftalie Weinberger and Porter Williams (Forthcoming) offers a clue. They identify several contingent features of the world that license and support causal reasoning, what they call "the worldly infrastructure of causation." One of those features we have already described; macroscopic, coarse grained behavior of target systems should be independent of their exact microscopic realizers. Again, this does not hold over long timescales in chaotic systems. When these infrastructure features break down, they write that:

> it is warranted to conclude that the behavior of such systems simply will not admit a straightforward causal interpretation, at least on anything like how

> we presently think about causation... some systems – at least when modeled at certain levels of analysis – may simply be "unfriendly" to causal analysis because important worldly infrastructure is not present (Forthcoming, 34).

This is a starting point for my proposal for the present case. The reason one might have conflicting intuitions about the butterfly effect, and the reason various philosophical accounts of causation diverge, is because, at the event-level, chaotic systems over long timescales become unfriendly to causal analysis.[26] We are taking the concept of causation beyond the domain it is tuned for; the domain of manipulation and control.

A stronger way of putting this is that 'causation' has become referentially indeterminate, or, exhibits "open texture" in the sense Friedrich Waismann (1945, 121).[27] According to Waismann, a concept is open texture insofar as there exists potential contexts under which there is no correct answer as to whether it applies. This is because there are two or more definitions of the term that are coextensive within its ordinary domain of use, but disagree in the novel context.

In particular, I contend that when pondering whether $Flap$ causes $Tornado$, our judgments are split between two competing conceptions of causation that typically align.[28] One conception is closely tracked by the conjunction of counterfactual and probabilistic dependence. Counterfactual dependence tells us that varying A will vary B, and probabilistic dependence ensures that they covary. Another conception is that causation is just something like counterfactual dependence. On this reading, probabilistic dependence is not part of the meaning of 'cause' but rather part of its ordinary context. This is because the counterfactual dependencies which are readily observed and controlled, the ones that make causation a useful concept to have, are the ones which are predictable. The butterfly effect provides a novel case where these two conceptions come apart. By ascribing referential indeterminacy to 'cause' in this case, we can explain conflicting intuitions, as well as the divergence in philosophical accounts.

To see how this could be the case, consider how we ordinarily discover causal relationships. First, we observe a correlation between some events or variables X and Y. Then, we investigate whether changing X in various ways can be used to reliably change Y, or vice versa. If we discover that varying X produces variations in Y, then we can label the relationship as causal. We will often use this knowledge about

[26] As David Danks and Maralee Harrell (2015) point out, chaotic systems also make us reconsider whether the right causal variables to consider are events, or higher level features of the system.

[27] On this point Woodward and I depart company. He insists on a monocriterial view of causation centered on manipulability (Woodward, 2003, 93) with the worldly infrastructure helping to account for the other criteria while relegating them to a subordinate semantic role. My view is more ecumenical among these competing criteria, best articulated by as the "amiable jumble" view of Brian Skyrms (1984, 254) described below.

[28] A similar story could be told about whether $flap_a$ causes $Tornado$, instead focusing on the presumed stability under different background conditions of causal connections.

type-level causal relationships to then make causal claims about actual events. The question the butterfly effect confronts us with is: should the initial correlation count as part of the conceptual content of ‘cause’ or merely a heuristic for identifying it? The path of discovery to the butterfly effect bypasses this identification process. We did not observe a correlation between butterfly flaps and tornadoes. Rather, the example is drawn out of the mathematics of chaos theory. Only after sensitive dependence was discovered in the dynamics of weather models did physicists begin speculating about counterfactuals like the butterfly effect and labeling them as causation.

This situation is not without precedent. Brian Skyrms (1984) describes something similar in the context of the violations of locality observed in entangled quantum systems. He writes that “our ordinary, everyday conception of causation is an amiably confused jumble” of many distinct causal theses, such as probabilistic dependence and the transfer of energy/momentum, which “in the noisy macroworld of everyday life ... go together” (ibid. 254). In the context of certain experiments on quantum systems, these ideas come decoupled, and we are unsure whether our concept applies. Thus, this would not be the first time that advances in physics have opened up a semantic indeterminacy for our ordinary understanding of causation.

There are some additional surprising features about the current case. First, the concept is coming apart along a new axis: the probabilistic and counterfactual. Second, it appears in classical physics, a primarily macro-domain that is less prone to clashing with the manifest image. Finally, the physical circumstances of the present indeterminacy are neither delicate nor unusual. In chaotic systems, they will predominate over long timescales.

But now that we have the indeterminacy in our sights, we can close it by fiat! I am inclined to agree with Skyrms that “It is better to stick with the amiable jumble. Anything else will seem inadequate” (ibid. 254). Except I would go even further. Whether or not we should try to modify our folk notion of cause will depend on whether we *even can* modify it in any meaningful way. My guess is that causation will turn out to be too deep in our conceptual stack for idle tinkering. Just as our folk notion of time cannot be brought into line with what relativistic physics refers to as ‘time’, our core folk notion of cause might refuse to fall in line with whatever scientists deem to be the most productive use of the word.

5 Long Timescales, Causal Chains, and Causal Explanations

Before ending, I will discuss the possibility of causal chains in chaotic systems. Since the time between the flap and the tornado was assumed to be large, it is natural to wonder what happens when we include the many mediating connections between the two events. Perhaps if we accounted for these connections, in the form of a causal chain, we would get a clear causal picture. I will handle this in two parts; I will first address the timescale question and then the question of causal chains.

First, there are many systems which are chaotic over much shorter timescales than the weather. The timescale on which a system displays chaos is known as the Lyapunov time. The weather has an estimated Lyapunov time of ∼14 days. A double pendulum's Lyapunov time is much shorter at ∼5 seconds. Additionally, there are many quasi-chaotic processes — e.g. dice rolls, coin flips, and roulette wheels — whose hypersensitivity to initial conditions creates probabilistic independence on short timescales, even though they are not mixing. Any of these systems could create a "butterfly effect" in short order. Thus, there is nothing in principal that keeps the type of behavior described confined to relatively long timescales, or compels us to consider mediating variables.

Now let us think about causal chains in the butterfly effect example. Similar to the case of standard causal dependence, the butterfly effect requires us to delineate two different interpretations of "causal chain" that are often blurred together.

It is true that you could specify a probability raising chain of weather "events" stretching from the butterfly's flap at time 0 to the tornado at time t. Considering a partitioning of the weather's phase space into macrostates, $M_1, M_2, ..., M_n$, such that each macrostate involves an approximate *macrophysical* specification of the atmosphere (local temperature, pressure, moisture, etc.) all the way down to a scale that can pick up the small differences left from the butterfly's flap. Each macrostate will consist of many microstates. Given the macrostate of the atmosphere at an instant $M_i^{t_n}$, there will be a macrostate $M_j^{t_m}$ that is highly probable a short time later; $P(M_j^{t_m}|M_i^{t_n}) \approx 1$. Given a different macrostate at t_n, the probability of $M_j^{t_m}$ is much lower. Thus, we could string a chain of macrostates from time of the butterfly's flap to the tornado's occurrence such that each macrostate raises the probability of the next. We might even be able to represent this as what advocates of the "Network Model of Causation" describe as "vast and mind-bogglingly complex 'neuron diagram"' (Beebee, 2004, 291) of localized weather events that are causally related to neighboring weather events. Thus, if by causal chain we just mean a chain of probability raising relations mediating $Flap$ and $Tornado$, then the case appears to meet this criteria.

However, often when we talk about causal chains we assume that the chain figures in a stable causal connection between macrophysical tokens or types. In other words, given the initial macrostate at time 0, there is a kind of regularity in the complex chain of events that leads to the macrostate at t. However, in this case, even if we observed the *same* weather macrostate in the future that we did at time 0, the course of events that would follow would (almost certainly) be extremely different. Unaccounted for microphysical differences would be constantly bubbling up and changing the course of macro-physical events in large ways. Thus, there is no *stable* causal chain that will connect the $Flap$ and $Tornado$.

Another way of looking at the situation is in terms of causal explanation. Immediately preceding the tornado, there are large macro-events that causally explain it, in the sense of making its occurrence unsurprising. For example, we could say that the tornado was caused by the large difference in vertical temperature immediately preceding it. This temperature difference was caused by dense cloud cover and warm air drifting off the Gulf of Mexico. These in turn have their own causes, which have their own causes, etc. What we cannot do is link these explanations together indefinitely to tell a convincing story about how macro-conditions in the distant past caused the tornado. The reasons are not purely combinatorial, where one effect has many causes (Lewis, 1986, 214-215). It is because at each step in the causal chain, we need to specify the past at a finer level of grain, so that our "macrostate" occupies a smaller region in phase space. Otherwise, conditioning on the past state will not raise the probability of the tornado's occurrence. Eventually, this level of grain will cease to be macro-physical. Accordingly, chaotic systems admit causal explanations between macro-events over short timescales, but not long timescales. Given a macrophysical description of the weather 10 years ago, you cannot, in principle, tell a clear causal story that leads to its current state.

6 Conclusion

What are we to make of Edward Lorenz's original question; can the flap of a butterfly's wings in Brazil cause a tornado in Texas?[29] The physical side of his question has been largely settled in the intervening years; yes the weather is highly chaotic. The semantic side of this question has received little attention. We have seen that chaos creates a situation where a tornado can counterfactually depend on the butterfly's flap, even though the two macro-events are probabilistically independent. Due to this odd combination of properties, I argued that there is no correct answer for whether the butterfly's flap "causes" the tornado. This is because there are two meanings of 'cause' — counterfactual dependence vs. predictable counterfactual dependence — which are coextensive inside the concept's normal domain, but come apart in this novel context.

We are left in an odd practical situation. While it may be true that changing something as small as a butterfly's flap in the distant past would have probably prevented an actual tornado from occurring, we cannot in practice manipulate butterfly flaps to achieve favorable meteorological outcomes. We have imperfect knowledge of the present, and changes we can make in the present are macro-sized. Such macro-sized changes cannot be used to tilt the probabilities of far-off future events in our favor; they reshuffle an already shuffled deck. However, while appreciating chaos in this way does not augment our abilities of control, it does deepen our understanding of where to

[29] Lorenz uses the phrase "set off" instead of "cause" in his original lecture (1972, 1).

situate 'causation' between us and nature.

7 Appendix

In this appendix, I will describe some standard mathematics of chaos theory as well as some technically involved details that were omitted in the main paper.

7.1 Measure-Preserving Dynamical Systems

A measure-preserving dynamical system can be defined by a quadruple $(\Gamma, \Sigma, \mu, \phi)$. Γ is the *phase space* of the system, the set of all possible states the system can take. The state of a dynamical system is represented as a point x in its phase space. Σ is a σ-algebra on Γ, defining the measurable sets. Measurable sets $A \in \Sigma$ are also called "events." $\mu : \Sigma \rightarrow [0, 1]$ is a measure where $\mu(\Gamma) = 1$. $\phi_t : \Gamma \rightarrow \Gamma$ is a surjective map given by the system's dynamics: $\phi_t(x)$ is phase point $x \in \Gamma$ evolved forward by time $t \in \mathbb{R}$ or $\mathbb{Z}$. For all $A \in \Sigma$, $\phi_t(A) = \{\phi_t(x) : x \in A\}$. A dynamical system is measure preserving iff for all measurable subsets $A \in \Sigma$, $\mu(\phi_t^{-1}(A)) = \mu(A)$, where $\phi_t^{-1}(A)$ are all the points that get mapped onto A.

A measure-preserving system is mixing iff for any two events $A, B \in \Sigma$,

$$\lim_{t \to \infty} \mu\left(\phi_t\left(A\right) \cap B\right) = \mu(A)\mu(B). \tag{6}$$

This says that the measure of A that ends up in B is the product of the measures of A and B. In other words, every region A of phase space eventually evolves towards the same spread-out distribution over every other region B. To state this probabilistically, it is typical to use the time-invariant, normalized measure as an objective probability measure for the system (Berkovitz et al., 2006, 673). The prior probability of event $A \in \Sigma$ is given by the measure of that event,[30]

$$\mu(A) = P(A) \text{ for all } A \in \Sigma. \tag{7}$$

The invariant measure can be thought of as providing the long-run objective chance for

[30]This can be justified by the fact that mixing systems are ergodic, meaning that for almost all $x \in \Gamma$,

$$\lim_{\tau \to \infty} \frac{1}{\tau} \int_0^\tau f(\phi_t(x))\, dt = \int_\Gamma f\, d\mu$$

for any measurable function $f : \Gamma \rightarrow \mathbb{R}$ such that $f \in L^1(\mu)$. This equation implies that, for almost all trajectories, the proportion of time spent in any region of phase space will be equal to the measure of that region. Note that ergodicity is strictly weaker than mixing, so there are systems (such as the simple harmonic oscillator) that are ergodic but do not satisfy (6) an event occurring. For any two events $A, B \in \Sigma$ occurring at times 0 and t,

$$P(B^t \,\&\, A^0) = \mu(\phi_t(A) \cap B). \tag{8}$$

The probability of B occurring after A is defined as the "amount" of A that ends up in B, or the measure of initial conditions in A that lead to B. From (7) and (8), the definition of mixing (6) implies:

$$\lim_{t\to\infty} P(B^t \,\&\, A^0) = P(A)P(B), \tag{9}$$

or using the standard ratio formula for conditional probability, $P(B|A) = P(A \cap B)/P(A)$,

$$\lim_{t\to\infty} P(B^t|A^0) = P(B) \tag{10}$$

which is equation (1) in the main text.

7.2 From Time-Invariant Measures to Snapshot Attractors

The time invariant measure μ is specific to the system in question. For example, in mixing thermodynamic systems, such as a box of hard sphere gas, the time invariant measure is the microcanonical measure which is uniform over the microstates with a given energy. This measure represents thermal equilbrium. Importantly, this would *not* be the relevant measure for the weather system or systems like it. The weather does not reach thermal equilibrium, but is constantly undergoing forcing and dissipation.

A better first approximation for understanding the weather system's dynamics is to consider the invariant measure of systems like the Lorenz system (see §7.4 below). These systems feature "strange attractors" that almost all nearby trajectories are asymptotically drawn towards, and the invariant measure is defined over the attractor region. These apply to non-isolated chaotic systems, like the weather, that undergo external forcing and dissipation. The invariant measure over these attractors are often non-uniform and can only be approximated numerically by letting ensembles of initial conditions relax towards the time-invariant distribution. Regarding these systems, Charlotte Werndl states that "for characterizing the unpredictability of motion dominated by strange attractors, it is widely acknowledged that *it suffices to consider the dynamics on attractors*, where natural invariant measures can be defined" (Werndl, 2009, 199).

An additional subtlety comes from the fact that the Earth's weather dynamics are always shifting as the forcing and dissipation parameters change over time (e.g. changes in the seasons, changes in the chemical composition of the atmosphere, etc.). Thus, the attractor region will also evolve over time as these parameters change, as will its natural measure. However, the same general lesson applies here; once the initial conditions have spread across this time-dependent attractor region, there will be no detectable correlation between $Flap\,\&\,k$ and $Tornado$. Both ensembles, $Flap\,\&\,k$ and $\neg Flap\,\&\,k$, will be similarly distributed across what is called a "snapshot" attractor (Romeiras et al., 1990), which is the time t attractor observed by taking an ensemble of trajectories in the distant past and evolving them to time t under the same equations of motion. Different initial ensembles converge to the same distribution over the snapshot attractor at t in much the same way as occurs in systems with time-invariant measures (see Drótos et al. 2015 and references therein). Thus, the prior probability $P(Tornado)$ for a weather model that accounts for these time-dependencies will be given by the natural probability measure over this snapshot attractor at the time the tornado occurs.

7.3 Conditioning on Measure-Zero Events

Many of the probabilistic claims made in the text implicitly rely on conditioning on events that are measure zero in μ. Therefore, the ratio formula used to derive (10) does not apply because it involves division by zero. Here, I propose a modified understanding of conditional probability for the cases at hand. This strategy is broadly in keeping with Hájek (2003) suggestion that the ratio formula (particularly when applied to the full phase space measure) as is not the only salient definition of conditional probability.

The easy case involves conditioning on a single micro-state x, corresponding to the exact initial conditions of a flap and exact background conditions. There is a natural definition of conditional probability suited for this task:

$$P\big(Tornado^t|\, x^0\big) = \begin{cases} 1 & \text{if } \ x(t) \in Tornado \\ 0 & \text{if } \ x(t) \notin Tornado. \end{cases}$$

Determinism will entail that counterfactual and probabilistic independence cannot run apart at the level of micro-states.

The more difficult case is the one involving the initial regions $Flap \cap k$ and $\neg Flap \cap k$, which involve a spread out region in $\Gamma_{butterfly}$ and a point in $\Gamma_{background}$. They will be measure zero according to the full phase space measure (think about how a line has zero area), and will remain so under time evolution. Thus, the relevant conditional probability cannot be defined by applying the ratio formula with this measure.[31] Instead, for our given k, we define a probability measure $P_k(\cdot)$ strictly over $\Gamma_{butterfly}$. What we want to know is the proportion of the initial conditions in $Flap$ and

[31] According to the full phase space measure,

$$P(Tornado^t|Flap^0\ \&\ k^0) = \frac{\mu(\phi_t(Flap\ \cap\ k) \cap Tornado)}{\mu(Flap\ \cap\ k)} = \frac{0}{0}.$$

$\neg Flap$ that lead to $Tornado$ at t, while holding k fixed. The relevant conditional probability is given by

$$P(Tornado^t | Flap^0 \,\&\, k^0) = \frac{P_k(\phi_t^{-1}(Tornado) \cap Flap \cap k)}{P_k(Flap)}. \tag{11}$$

In other words, we use the inverse transformation to find the set that maps to $Tornado$ after t and take the intersection with $Flap \cap k$. This will select the initial conditions in $Flap \cap k$ that end up in $Tornado$ at t. The rest is just an application of the ratio formula, only using the measure $P_k(\cdot)$. This is type of measure that can be tacitly assumed from Figure 3. Again, as stated in the main text, given how interspersed the regions of initial conditions that lead to different outcomes will be, just about any non-pathological $P_k(\cdot)$ will suffice to give probabilistic independence.

7.4 The Mixing Problem

Given this description of conditional probabilities, mixing does not directly imply that $Tornado$ and $Flap$ are probabilistically independent. Consider again the initial region $Flap \cap k$. Mixing alone does not guarantee that this region will spread out under time evolution. In chaotic systems, positive measure regions in phase space are exponentially stretched in some directions and exponentially contracted in others. The stretching directions correspond to unstable manifolds in phase space and the contracting directions stable manifolds (Hilborn, 2000, ch. 4). If the initial region $Flap \cap k$ is confined to a lower dimensional subspace purely along a stable manifold, then we will not get the chaotic behavior described in the case.[32]

However, if the hypothesis of the butterfly effect is true — i.e. changing only a butterfly's flap can change the tornado's occurrence — then this implies that the $Flap$ and $\neg Flap$ regions are not confined to the stable manifolds. As long as the equations of motion for the system are sufficiently coupled, virtually any initial macro-differences along one or more dimensions will be smeared out over all of phase space. To use a

[32]For example, consider the Baker's map

$$(x_{n+1}, y_{n+1}) = \begin{cases} (2x_n, y_n/2) & \text{if } 0 \le x < \frac{1}{2} \\ (2x_n - 1, (y_n + 1)/2) & \text{if } \frac{1}{2} \le x \le 1 \end{cases}$$

which is a chaotic 2D map from the unit square onto itself whose invariant measure is the 2D Lebesgue measure. Say our initial region is a vertical line consisting of an interval of values in the y direction that are fixed at a single value in the x direction. Even though the Baker's map is mixing for all regions with positive 2D Lebesgue measure, this lower-dimensional region will not spread out to fill the state space as $n \to \infty$. Rather, it will exponentially shrink under iteration.

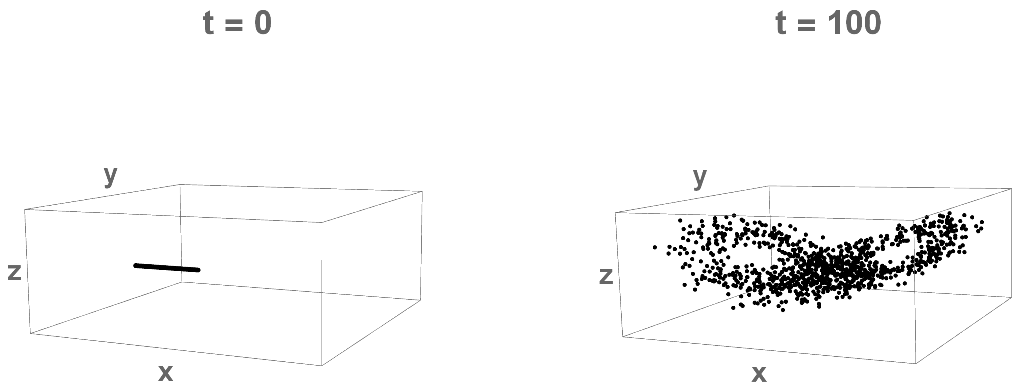

Figure 5: Simulation of the Lorenz system. A 1-dimensional ensemble of initial conditions spreads across the whole attracting region under time evolution.

relevant example, take Lorenz's (1963, 135) original model for atmospheric convection

$$\frac{\mathrm{d}x}{\mathrm{d}t} = 10(y - x), \tag{12}$$

$$\frac{\mathrm{d}y}{\mathrm{d}t} = x(28 - z) - y, \tag{13}$$

$$\frac{\mathrm{d}z}{\mathrm{d}t} = xy - \frac{8}{3}z, \tag{14}$$

which is known to be chaotic. A line of initial conditions with non-zero length in only the x, y, or z directions will spread out across the entire attractor region (Figure 5). This is because the equations are coupled in such a way that differences along any of these dimensions immediately creates differences along all of them.

References

Albert, D. Z. (2000). *Time and chance*. Harvard University Press.

Bacon, A. (2015). Stalnaker's thesis in context. *Review of Symbolic Logic*. *8*(1), 131–163.

Bandak, D., Mailybaev, A. A., Eyink, G. L. and Goldenfeld, N. (2024). Spontaneous stochasticity amplifies even thermal noise to the largest scales of turbulence in a few eddy turnover times. *Phys. Rev. Lett. 132*, 104002.

Beebee, H. (2004). Causing and nothingness In *Causation and Counterfactuals*, Collins, J., Hall, N., and Paul, L. A. (eds). MIT Press. Cambridge, MA. pp. 291–308.

Bennett, J. (1974). Counterfactuals and possible worlds. *Canadian Journal of Philosophy*. *4*(2), 381–402.

Bennett, J. (2003). *A philosophical guide to conditionals*. Clarendon Press.

Berkovitz, J., Frigg, R. and Kronz, F. (2006). The ergodic hierarchy, randomness and hamiltonian chaos. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*. *37*(4), 661–691.

Berry, M. V. (1978). Regular and irregular motion. *AIP Conference Proceedings*. *46*(1), 16–120.

Bishop, R. (2024). Chaos In *The Stanford Encyclopedia of Philosophy*, Zalta, E. N. and Nodelman, U. (eds). Metaphysics Research Lab, Stanford University. Winter 2024 edition.

Bricmont, J. (2022). *Making Sense of Statistical Mechanics*. Springer Nature.

Carlstrom, I. F. and Hill, C. S. (1978). Review of adams's logic of conditionals. *Philosophy of Science*. *45*(1), 155–158.

Cartwright, N. (1979). Causal laws and effective strategies. *Noûs*. *13*(4), 419–437.

Danks, D. and Harrell, M. (2015). Chaos, causation, and describing dynamics. *Minnesota Studies in the Philosophy of Science*. *21*.

Diaconis, P., Holmes, S. and Montgomery, R. (2007). Dynamical bias in the coin toss. *SIAM review*. *49*(2), 211–235.

Dorr, C. (2016). Against counterfactual miracles. *Philosophical Review*. *125*(2), 241–286.

Drótos, G., Bódai, T. and Tél, T. (2015). Probabilistic concepts in a changing climate: A snapshot attractor picture. *Journal of Climate. 28*(8), 3275–3288.

Dürr, D., Goldstein, S. and Zanghi, N. (1992). Quantum chaos, classical randomness, and bohmian mechanics. *Journal of Statistical Physics. 68*(1), 259–270.

Edgington, D. (1995). On conditionals. *Mind. 104*(414), 235–329.

Edgington, D. (2008). Counterfactuals. *Proceedings of the Aristotelian Society. 108*, 1–21.

Frisch, M. (2014). *Causal Reasoning in Physics*. Cambridge University Press.

Gibbs, J. W. (1902). *Elementary Principles in Statistical Mechanics: Developed with Especial Reference to the Rational Foundation of Thermodynamics*. C. Scribner's sons.

Goldstein, S. and Santorio, P. (2021). Probability for epistemic modalities. *Philosophers' Imprint. 21*(33).

Hájek, A. (2003). What conditional probability could not be. *Synthese. 137*(3), 273–323.

Hájek, A. (2011). Triviality pursuit. *Topoi. 30*(1), 3–15.

Hájek, A. (2014). Most counterfactuals are false. *Unpublished Manuscript*.

Hilborn, R. C. (2000). *Chaos and Nonlinear Dynamics: An Introduction for Scientists and Engineers*. Oxford University Press.

Hilborn, R. C. (2004). Sea gulls, butterflies, and grasshoppers: A brief history of the butterfly effect in nonlinear dynamics. *American Journal of Physics. 72*(4), 425–427.

Hitchcock, C. (2003). Routes, processes, and chance-lowering causes In *Cause and Chance: Causation in an Indeterministic World*, Dowe, P. and Noordhof, P. (eds). Routledge.

Hoefer, C. (2004). Causality and determinism: Tension, or outright conflict? *Revista de Filosofía (Madrid). 29*(2), 99–115.

Lewis, D. (1973). Causation. *The Journal of Philosophy. 70*(17), 556–567.

Lewis, D. (1973). *Counterfactuals*. Harvard University Press.

Lewis, D. (1979). Counterfactual dependence and time's arrow. *Noûs*. pp. 455–476.

Lewis, D. (1986). Causal explanation In *Philosophical Papers Vol. II*, Lewis, D. K. (eds). Oxford University Press. pp. 214–240.

Lewis, D. (1986). *Philosophical Papers Vol. II*. Oxford University Press.

Lewis, D. (2000). Causation as influence. *The Journal of Philosophy. 97*(4), 182–197.

Lewis, D. (2004). Causation as influence . In *Causation and Counterfactuals*. MIT Press. p. 75–106.

Loewer, B. (2023). The mentaculus: A probability map of the universe In *The Probability Map of the Universe: Essays on David Albert's* Time and Chance, Loewer, B., Weslake, B., and Winsberg, E. (eds). Harvard University Press. Cambridge MA.

Lorenz, E. (1972). Predictability: Does the flap of a butterfly's wing in brazil set off a tornado in texas? American Association for the Advancement of Science, 139th meeting.

Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of atmospheric sciences. 20*(2), 130–141.

Lorenz, E. N. (1969). The predictability of a flow which possesses many scales of motion. *Tellus. 21*(3), 289–307.

Mackie, J. L. (1980). *The Cement of the Universe: A Study of Causation*. Oxford University Press.

McDermott, M. (2007). True antecedents. *Acta Analytica. 22*, 333–335.

Mellor, D. H. (1995). *The Facts of Causation*. Routledge.

Myrvold, W. C. (2021). *Beyond Chance and Credence: A Theory of Hybrid Probabilities*. Oxford University Press. Oxford, United Kingdom.

Palmer, T., Döring, A. and Seregin, G. (2014). The real butterfly effect. *Nonlinearity. 27*(9), R123.

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

Poincaré, H. (1905). *Science and hypothesis*. The Walter Scott Publishing Co., LTD.

Romeiras, F. J., Grebogi, C. and Ott, E. (1990). Multifractal properties of snapshot attractors of random maps. *Phys. Rev. A. 41*, 784–799.

Sinai, Y. G. (1970). Dynamical systems with elastic reflections. *Russian Mathematical Surveys. 25*(2), 137.

Skyrms, B. (1980). *Causal Necessity*. Yale University Press. New Haven and London.

Skyrms, B. (1984). Epr: Lessons for metaphysics. *Midwest Studies in Philosophy. 9*(1), 245–255.

Smith, P. (1990). The butterfly effect. Proceedings of the Aristotelian society. pp. 247–267.

Stalnaker, R. (1968). A theory of conditionals In *Studies in Logical Theory*, Rescher, N. (eds). Blackwell. pp. 98–112.

Stalnaker, R. (1970). Probability and conditionals. *Philosophy of Science. 37*(1), 64–80.

Stalnaker, R. (2019). *Knowledge and Conditionals: Essays on the Structure of Inquiry*. Oxford University Press. Oxford, England.

Stalnaker, R. C. (1981). pp. 87–104 in *A Defense of Conditional Excluded Middle*. edited by Harper, W. L., Stalnaker, R. and Pearce, G. Springer Netherlands.

Starr, W. (2021). Counterfactuals In *The Stanford Encyclopedia of Philosophy*, Zalta, E. N. (eds). Metaphysics Research Lab, Stanford University. Summer 2021 edition.

Strevens, M. (2011). Probability out of determinism In *Probabilities in Physics*, Beisbart, C. and Hartmann, S. (eds). Oxford University Press. pp. 339–364.

Suppes, P. (1970). *A Probabilistic Theory of Causation*. North-Holland Pub. Co.

van Fraassen, B. (1976). Probabilities of conditionals In *Foundations of probability theory, statistical inference, and statistical theories of science*, Hooker, W. H. C. (eds).

Waismann, F. (1945). Verifiability. *Proceedings of the Aristotelian Society, Supplementary Volume 19*. pp. 119–150.

Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford University Press.

Walters, L. (2016). Possible world semantics and true-true counterfactuals. *Pacific Philosophical Quarterly. 97*(3), 322–346.

Walters, L. and Williams, J. R. G. (2013). An argument for conjunction conditionalization. *The Review of Symbolic Logic. 6*(4), 573–588.

Weinberger, N., Williams, P. and Woodward, J. (Forthcoming). The worldly infrastructure of causation. *The British Journal for the Philosophy of Science*.

Werndl, C. (2009). What are the new implications of chaos for unpredictability? *The British Journal for the Philosophy of Science. 60*(1), 195–220.

Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.

Woodward, J. (2007). Causation with a human face In *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, Price, H. and Corry, R. (eds). Oxford University Press.

Woodward, J. (2021). *Causation with a Human Face: Normative Theory and Descriptive Psychology*. Oxford University Press.

Woodward, J. (2023). Causation and Manipulability In *The Stanford Encyclopedia of Philosophy*, Zalta, E. N. and Nodelman, U. (eds). Metaphysics Research Lab, Stanford University. Summer 2023 edition.

Zeng, X., Eykholt, R. and Pielke, R. A. (1991). Estimating the lyapunov-exponent spectrum from short time series of low precision. *Phys. Rev. Lett. 66*, 3229–3232.

Zuchowski, L. C. (2017). *A Philosophical Analysis of Chaos Theory*. Palgrave MacMillan. Cham.