The Weirdness of the World

Chapter Nine
Almost Everything You Do Causes Almost Everything (Under Certain Not Wholly Implausible Assumptions); or Infinite Puppetry

with Jacob Barandes

Infinitude is bizarre. For example, on standard mathematical treatments, the set of all counting numbers (1, 2, 3, 4…) has the same cardinality, or number of elements, as the set of all squares of the counting numbers (1, 4, 9, 16…), despite the fact that the second is a proper subset of the first.[1] Or consider the "St. Petersburg game", a game where you flip a fair coin as often as necessary for it to come up heads exactly once, earning $\$2^n$ per flip. (Thus, you earn $2 if the first flip is heads, $4 if the results are tails then heads, $8 if it's tails-tails-heads, etc. Probably you'll make $32 or less.) This game not only has an infinite expected dollar return but also exactly the same expected return as an alternative version that pays $\$2^n$ per flip plus $1000 as a bonus no matter the outcome.[2]

In this chapter, we will explore the idea that if the universe is infinite and various other plausible assumptions about fundamental physics and cosmology hold, then almost everything you do causes almost everything.

1. It's Reasonable to Think That the Universe Is Infinite.

---

[1] This observation dates back to Galileo, who, however, instead of concluding that the bijection of the two sets was sufficient to establish their equal cardinality (now the standard view), concludes that "the attributes 'larger,' 'smaller,' and 'equal' have no place either in comparing infinite quantities with each other or in comparing infinite with finite quantities" (1638/1914, p. 33 [80]). For a philosophical introduction to infinitude, see Oppy, Hájek, Easwaran, and Mancosu 2021.

[2] Actually, things get even weirder if the returns are sometimes negative. See Nover and Hájek 2004; Fine 2008.

On recent estimates, the observable universe – the portion of the universe that we can detect through our telescopes – extends about 47 billion light years in every direction.[3] But the limit of what we can see is one thing, and the limit of what exists is quite another. It would be remarkable if the universe stopped *exactly* at the edge of what we can see. For one thing, that would place us, surprisingly and un-Copernicanly, precisely at the center.

But even granting that the universe is likely to be larger than 47 billion light years in radius, it doesn't follow that it's infinite. It might be finite. But if it's finite, then one of two things should be true: Either the universe should have a boundary or edge, or it should have a *closed topology*.

It's not absurd to think that the universe might have an edge. Theoretical cosmologists routinely consider hypothetical finite universes with boundaries at which space comes to a sudden end. However, such universes require making additional cosmological assumptions for which there is no direct support – assumptions about the conditions, if any, under which those boundaries might change, and assumptions about what would happen to objects or light rays that reach those boundaries.

It's also not absurd to think that the universe might have a closed topology. By this we mean that over distances too large for us to see, space essentially repeats, so that a particle or signal that traveled far enough would eventually come back around to the spatial region from which it began – like how when Pac-Man exists one side of the TV screen, he re-emerges from the other side. However, there is currently no evidence that the universe has a closed topology.[4]

[3] Gott et al. 2005.

[4] The connection between topology and local geometrical notions like spatial curvature is subtle. If the curvature of space turns out to be negative on large scales, then that would preclude the universe from having a closed topology. By contrast, if the curvature of space is

Leading cosmologists, including Alex Vilenkin, Max Tegmark, and Andre Linde, have argued that spatial infinitude is the natural consequence of the best current theories of cosmic inflation.[5] Given that, plus the absence of evidence for an edge or closed topology, infinitude seems a reasonable default view, and we will assume it for the remainder of this chapter. The mere 47 billion light years we can see is the tiniest speck of a smidgen of a drop in an endless expanse.

2. It Is Reasonable to Think That If the Universe Is Infinite, We Stand in Both Spacelike and Timelike Relation to Infinitely Many Sibling Galaxies.

Let's call any galaxy with stars, planets, and laws of nature like our own a *sibling galaxy*. Exactly how similar a galaxy must be to qualify as a sibling we will leave unspecified, but we don't intend high similarity. Andromeda is sibling enough, as are probably most of the other hundreds of billions of ordinary galaxies we can currently see.

The finiteness of the speed of light means that when we look at these faraway galaxies, we see them as they were during earlier periods in the universe's history. Taking this time delay into account, the laws of nature don't appear to differ in regions of the observable universe that are remote from us. Likewise, galaxies don't appear to be rarer or differently structured in one direction or another. Every direction we look, we see more or less the same stuff. These observations help motivate the Copernican Principle, which is the working hypothesis that our position in the universe is not special or unusual – not the exact center, for example, and not the

---

positive on large scales, then the magnitude of that curvature would naturally pick out a specific periodicity distance. The best estimates of the large-scale curvature of space are that it's approximately zero (Planck Collaboration 2014),

[5] Vilenkin 2006; Tegmark 2007; Linde 2015/2017.

one weird place that happens to have a galaxy operating by special laws that don't hold elsewhere.[6]

Still, our observable universe might be an atypical region of an infinite universe. Possibly, somewhere beyond what we can see, different forms of elementary matter might follow different laws of physics. Maybe the gravitational constant is a little different. Maybe there are different types of fundamental particles. Even more radically, other regions might not consist of three-dimensional space in the form we know it. Some versions of string theory and inflationary cosmology predict exactly such variability.[7] But even if our region is in some respects unusual, it might be common enough that there are infinitely many other regions similar to it – even if just one region in 10^{500}. Again, this is a fairly standard view among speculative cosmologists, which comports well with straightforward interpretations of leading cosmological theories. One can hardly be certain, of course. Maybe we're just in a uniquely interesting spot! But for purposes of this chapter, we are going to assume that's not the case. In the endless cosmos, infinitely many regions resemble ours, with three spatial dimensions, particles that obey approximately the "Standard Model" of particle physics, and cluster upon cluster of sibling galaxies.

Under the assumptions so far, the Copernican Principle suggests that there are infinitely many sibling galaxies in a *spacelike* relationship with us, meaning that they exist in places that are not reachable from our galaxy even by photons traveling at the speed of light. However, for the causal relationships we'd like to explore, we need to assume a bit more than this. We will also need it to be the case that there are infinitely many sibling galaxies in a *timelike* relationship

[6] One constraint on the Copernican Principle is the Anthropic Principle: Whatever region we occupy must be capable of supporting cosmological observers like us. Even if such regions are rare, we should we unsurprised to be in one. See Barrow and Tipler 1986; Peacock 1998.

[7] Linde 2015/2017.

to us – that is, existing in the future in locations that are, at least in principle, reachable by particles originating in our galaxy. And this requires thinking about heat death.

Stars have finite lifetimes. If standard physical theory is correct, then ultimately all the stars we can currently see will burn out. Some of those burned-out stars will contribute to future generations of stars, which will, in turn, burn out. Other stars will become black holes, but then those black holes also will eventually dissipate (through Hawking radiation).[8] Given enough time, assuming that the laws of physics as we understand them continue to hold, and assuming things don't recollapse in a "Big Crunch" in the distant future, the standard view is that everything we presently see will inevitably enter a thin, boring, high-entropy state near equilibrium – heat death.

But what happens *after* heat death? This is of course even more remote and less testable than the question of whether heat death is inevitable. But we can speculate based on currently standard assumptions – and why not speculate? After all, we're weird philosophers. If anyone has license to speculate, it's us! Let's think as reasonably as we can about this. So here's our best guess, based on standard theory, from Ludwig Boltzmann – recall the "Boltzmann brains" of Chapter 4 – through at least some time slices of Sean Carroll.[9]

For purposes of this argument, we will assume that the famously probabilistic behavior of quantum systems is intrinsic to the systems themselves, persisting post heat-death and not requiring external observers carrying out measurements. This is consistent with most current approaches to quantum theory (including most many-worlds approaches, objective-collapse

[8] Hawking 1974.
[9] Boltzmann 1895, 1897; Carroll 2010; Aguirre, Carroll, and Johnson 2012.

approaches, and Bohmian mechanics).[10] It is, however, inconsistent with theories according to which the probabilistic behavior is generated by external observers carrying out measurements (some versions of the "Copenhagen interpretation") and theories on which the post-heat-death universe would inescapably occupy a stationary ground state.[11] Under this assumption, post-heat-death and without observers, the universe will continue to support random fluctuations. That is, from time to time, particles will, by chance, enter unlikely configurations. This is predicted by both standard statistical mechanics and standard quantum mechanics. Seven particles will sometimes converge, by chance, upon the same small region. Or seven hundred. Or – very rarely! – seven trillion.

There appears to be no in-principle limit to how large such chance fluctuations can be or what they can contain if they pass through the right intermediate phases. Wait long enough and extremely large fluctuations should occur. Assuming the universe continues infinitely, rather than having a temporal edge or forming a closed loop, for which there is no evidence, then eventually some random fluctuation should produce a bare brain having cosmological thoughts – the Boltzmann brain idea, discussed in Chapter 4. Wait longer, and eventually some random fluctuation will produce, as Boltzmann suggested, a whole galaxy. If the galaxy is similar enough to our own, it will be a sibling galaxy. Wait still longer, and another sibling galaxy will arise, and another, and another….

For good measure, let's also assume that after some point post heat-death, the rate at which galaxy-size systems fluctuate into existence does not systematically decrease. There's

[10] In particular, our argument doesn't depend on whether this probabilistic behavior is objective, arising from wave functions stochastically collapsing, or subjective, as experienced by embodied observers in the "universal wave function" of the many-worlds interpretation.

[11] On the Copenhagen interpretation, see Faye 2008/2019. Regarding the possibility of a stationary ground state, see Boddy, Carroll, and Pollack 2016, 2017.

some minimal probability of galaxy-sized fluctuations, not an ever-decreasing probability with longer and longer average intervals between galaxies. This assumption will prove helpful later, and it appears to be the most natural interpretation of the post-heat-death situation. Fluctuations appear at long intervals, by random chance, then fade back into chaos after some brief or occasionally long period, and the region returns to the heat-death state, with the same small probability of large fluctuations as before. Huge stretches of not much will be punctuated by rare events of interesting, galaxy-sized complexity.

Of course, this might not be the way things go. We certainly can't *prove* that the universe is like this. But despite the bizarreness that understandably causes some people to hesitate,[12] the overall picture we've described appears to be the most straightforward consequence of standard physical theory, taken out of the box, without too much twisting things around.

Even if this specific speculation is wrong, there are many other ways in which the cosmos might deliver infinitely many sibling galaxies in the future. For example, some process might ensure we never enter heat death and new galaxies somehow continue to be born. Alternatively – and this will become relevant later – processes occurring pre-heat-death, such as the formation of black holes, might lead to new bangs or cosmic inflations, spatiotemporally unconnected or minimally connected to our universe, and new stars and galaxies might be born from these new bangs or inflations in much the same way as our familiar stars and galaxies were born from our familiar Big Bang.[13] Depending on what constitutes a "universe" and a "timelike" relation, those

[12] The apparent implication that there will be infinitely many Boltzmann brain observers strikes some theorists as so bizarre as to constitute a reductio ad absurdum. However, see the discussion in Chapter 4, Section 6.

[13] Frolov, Markov, and Mukhanov 1989; Garriga and Vilenkin 1998; Easson and Brandenberger 2001; Carroll and Chen 2004; Garriga, Vilenkin and Zhang 2016.

sibling galaxies might not exist in our universe or stand in a timelike relation to us, technically speaking, but if so, that detail won't matter to the core idea of this chapter. Similarly, if the observable universe reverses its expansion, it might collapse upon itself in a Big Crunch, followed by another Big Bang, and so on in an infinitely repeating cycle, containing infinitely many sister galaxies post-Crunch. This isn't currently the mainstream view, but it's a salient and influential alternative if the heat-death scenario outlined above is mistaken.[14]

We conclude that it is reasonable to think that the universe is infinite, and that there exist infinitely many galaxies broadly like ours, scattered throughout space and time, including in our casual future. It's a plausible reading of our cosmological situation. It's a decent guess and at least a possibility worth taking seriously.

3. If Infinitely Many Sibling Galaxies Exist, Counterparts of Almost Everyone Are Doing Almost Everything Somewhere.

For the remainder of this chapter, we will assume that given sufficiently many opportunities, any finitely probable event will almost certainly occur arbitrarily many times. (By "finitely probable" here, we mean any event whose probability is one over a finite number. This excludes zero probability events as well as infinitesimally probable events if there can be any, that is, events whose probability is non-zero but still less than one over any finite number.[15]) However unlikely an event might be – say a streak of 100 heads in a series of independent coin

[14] Steinhardt and Turok 2002; Penrose 2006.

[15] One plausible but controversial example of an infinitesimally probable event is drawing the number "4" in an infinite lottery of all the counting numbers. A more physically plausible case is two complex, unrelated systems being exactly identical in every spatial detail down to an infinite level of precision (if spatial properties are continuous and not quantized). See Benci, Horsten, and Wenmackers 2018; but for one problem, see Norton and Parker 2021.

flips – the odds are nearly 100% that it will eventually occur. It might take many, many lifetimes to achieve, but we can say with a high degree of confidence that approximately one in every 2^{100} coin flips will begin a series of a hundred heads. At a rate of one flip per second, that's about once every 4×10^{22} years, or about three trillion times the duration since the Big Bang. That's a long time to wait for a gamble to pay off, but of course it's peanuts compared to infinity. Given infinite time, almost all such unlikely events will occur over and over again, endlessly.[16]

The consequence for our sibling galaxies is that every type of object or event that has a finite chance of occurring will occur not just in one galaxy but in infinitely many. This consequence has frequently been noticed – for example, by Vilenkin and Tegmark in their popular treatments of cosmological infinitude.[17] Suppose you've lost your car keys. The infinite cosmos will contain infinitely many key-shaped chunks of metal that would fit your ignition and start your car. Most, of course, will be much farther away than your couch cushions. You'd like a diamond as big as the Moon? The infinite cosmos sparkles with them.[18] If we further assume that the evolution of intelligent life is a finitely probable event, then there are infinitely many space aliens, of all finitely probable forms. (So the supersquids and antheads of Chapter 2 are real after all, even if not nearly as close as Sirius or Antares.) Infinitely many aliens will be of broadly human form, assuming that broadly human form was a finitely probable consequence of galactic evolution.

Similarly, infinitely many of these space aliens will be similar to you, specifically. They will live out, with varying degrees of similarity, every finitely possible life that you or someone

[16] This relates to the law of large numbers in statistics and ergodicity in dynamical systems but we won't attempt to establish it formally here.

[17] Vilenkin 2006; Tegmark 2014.

[18] Actually, such a diamond might not even be very far: Kaplan et al. 2014.

very similar to you might have lived. There will of course also be infinitely many Shakespeare-counterparts writing infinitely many Shakespearean plays, some vastly better than Shakespeare's own plays. Every finitely possible work of philosophy, too, will be written somewhere, including the most maximally correct metaphysics and cosmology that can be typeset in three hundred pages. Cast aside this book and imagine instead that much better book!

This merits a moment's awe and reflection. If something like standard physical theory and cosmology is correct, and if our other assumptions also hold, then the universe teems with duplicates and near-duplicates of you and all your loved ones, on duplicates or near-duplicates of Earth, doing almost every imaginable thing. Some will even do extremely improbable but not physically impossible things, such as declaring they can fly and then, to everyone else's amazement, soaring calmly up into the air for minutes at a time on a chance series of gusts and landing gently down.

According to current physical theory, there's a tiny, tiny, but finite probability that a baseball thrown at a brick wall will simply pass or "quantum tunnel" through the wall, appearing on the other side with both baseball and wall unharmed. Our infinitely many galaxies will thus contain infinitely many baseballs freakishly passing through infinitely many brick walls. Virtually every type of event that is less than galaxy-sized, finitely specifiable, and has even the tiniest chance of occurring will occur somewhere, within some finite error tolerance – a trillionth of the radius of a proton, say, for every constituent part, if you want to be really fussy.

Well, that's true with one potentially hugely important exception: events that are conceptualized in such a way that makes them unique. Maybe when I use words like "I" and "here" and "now", I pick out a single individual, in a single time and place, in the whole of the infinite cosmos; and maybe I do so in a way that doesn't require my having, or this spatiotemporal region's having, any distinctive physical properties that aren't also instantiated in far-away duplicates. Maybe I can just tag myself as unique.[19] Alternatively, maybe this spatiotemporal region is unique in virtue of having a certain unrepeated location in absolute spacetime.[20] If so, then here's an entirely unique event that won't occur anywhere else: *This* baseball, the one in your hand right now, bouncing off this particular brick wall in front of you. Other far away balls might have identical physical configurations down to 10^{-100} meter. They might bounce off extremely similar brick walls to this one here. Some will freakishly pass though instead of bouncing. But on this way of thinking, only this ball in your hand is yours. And that ball won't be passing undisturbed through an ordinary brick wall even if you throw until your arm is sore – we bet you a million dollars. If people or locations can be uniquely tagged, nowhere in the infinite cosmos does that happen. Similarly, although there might be infinitely many canyons that look just like Arizona's Grand Canyon, only one is in fact Arizona's Grand Canyon. Infinitely many twins of Confucius might grow up in infinitely many humanlike cultures on infinitely many Earthlike planets, speaking languages that sound almost exactly like classical Chinese, saying just the types of things that Confucius said. But, on this view, when you or we use the word "Confucius", we don't mean any of those far away philosophers. Among all the things that Confucius might have done, he only did a small subset.

---

[19] See Perry 1979 on the essential indexical and Lewis 1979 on centered worlds. Relatedly, but somewhat differently, qualitatively different individuals might have different haecceities – that is, properties in virtue of which individuals are the particular individuals they are and no other, which wouldn't be shared even by another individual with exactly the same qualitative features (for a review, see Cowling 2015/2016).

[20] In the context of an infinite cosmos, this might require being an absolutist rather than a relationalist about spacetime, if every finitely specifiable relationship is duplicated somewhere. On absolutism versus relationalism, see Huggett, Hoefer, and Read 2006/2021.

From this, it might seem that, despite the infinitude of the universe, the scope of your interactions and the consequences of your actions will be very limited. You will likely never cause any baseball to pass through any brick wall, for example, and your life will only have the limited range of effects on Earth that you're normally inclined to think it has – and the same will be true for virtually every other counterpart of you elsewhere.

However, that's not quite right, as we will now endeavor to show.

4. Infinitely Rippling Particles.

Suppose you raise your right hand. As a result, photons, electrons, and other elementary particles that would otherwise have remained several centimeters distant from you instead enter the vicinity of your hand. Their behavior changes significantly, or their quantum-mechanical properties change significantly. A photon that would have been absorbed into your desktop instead reflects off your hand and through your window. A nitrogen molecule in the air floats differently, ending up on a complex trajectory that, fifteen minutes later, takes it and the elementary particles composing it under the gap beneath your door. Of the many, many elementary particles disrupted by your movement, a portion of them escape the Earth's atmosphere into interstellar space. Let's follow one of those particles.

The particle will eventually interact with something – a hydrogen atom, a chunk of interstellar dust, a star, the surface of a planet. Something. Let's call that something a *system.* The particle might be absorbed, reflected, refracted, or annihilated by an antiparticle, or it might decay into other particles. (If the particle passes through the system entirely unaltered, let's ignore that system and keep following the particle.) If it interacts with a system, it will change the system, maybe increasing the energy of the system if it's absorbed or annihilated, or altering the trajectory of another particle if it's reflected or refracted. The perturbed system will then emit, reflect, refract, or gravitationally bend another particle differently than it otherwise would have. Choose one such *successor particle.* This successor particle will now head off on Trajectory A instead of Trajectory B or instead of not being emitted at all. That successor particle will in turn perturb another system, generating another successor particle traveling along another trajectory that it would not otherwise have taken. In this way, we can imagine a series of successor particles, one after the next, perturbing one system after another. Let's call this series of particles and perturbations a *ripple.*

Might some ripples be infinite? We see a few ways in which they could fail to be.

First, the universe might have finite duration, or after a finite period of time it might settle into some unfluctuating post-heat-death state that fails to contain systems capable of perturbation by incoming particles (or maybe even fails to contain any particles at all); or if not quite that, it might enter a state in which perturbable systems grow ever more spatiotemporally sparse quickly enough over time that our traveling particle cannot be expected to encounter an infinite number of them even given infinite time. (This is essentially one infinitude unsuccessfully chasing a faster infinitude.) However, as discussed above, the most straightforward interpretation of current physical theory does not suggest that the universe is finite or headed for utter quiescence. Nor is there any reason to think that the probability density of particle-perturbable systems would continue to decrease over time after heat death in such a way as to dodge our ripple of particles.

Second, the ripple might end. For example, traveling particles might be absorbed by some systems without perturbing those systems in a way that has any effect on successor systems. Once again, this seems unlikely on standard physical theory. Even a particle that strikes a black hole will slightly increase the black hole's mass, which should slightly alter how the black hole bends the light around it or subtly alter its Hawking radiation. Alternatively, all particles, even photons and protons, might decay at extremely long intervals into something that cannot continue the ripple, contrary to standard physical theory. Still another ripple-ending event might be the perturbation of a system in exactly the same way, by freak chance, that the system would have been perturbed by a *different* particle had you not raised your hand. (This last case counts as a ripple-ender because after such an event, further perturbations down the line no longer cause the systems to behave differently than they would have if you hadn't raised your hand.)

So there are several ways in which ripples might hypothetically end. But such rare ripple-enders can presumably be avoided by always choosing a large enough number n of successor particles, leading to n^m successors after m interactions, minus the small proportion of stopped ripples. We will assume, we hope not implausibly, that (contingent upon our other assumptions) this n^m strategy is sufficient to ensure that virtually every hand raising generates at least one infinite ripple.[21]

Thus, the most straightforward interpretation of existing physical theory implies that when you raise your hand – try it now, if you like – you launch a succession of particles rippling infinitely through the universe, perturbing an infinite series of systems. Of course this is speculative and uncertain, but if the universe is infinite, the conclusion is more natural and physically plausible than its negation.

[21] Here's a toy model on which this assumption would be plausible. Suppose that we follow a sphere of ripples out from a center. At time 1, there's a 1 in 10^{100} chance that all the ripples stop. At time 2, there's a smaller chance because there are more ripples, so there is maybe a 1 in 10^{1000} chance that all the ripples stop. At time 3, there's a 1 in 10^{10000} chance. And so forth. This is a convergent series. As time goes to infinity, the cumulative chance that all the ripples stop is not much more than 1 in 10^{100}.

5. Almost Everything You Do Causes Almost Every Type of Non-Unique Future Event.

Thus, it is fairly plausible, and probably the most straightforward interpretation of current physical theory, to suppose the following: (1.) The universe is infinite. (2.) This infinitude continues temporally after heat death. (3.) Post heat-death, galaxies like our own will occasionally fluctuate into existence by freak chance, with some finite and not ever-decreasing probability. (4.) Ordinary actions of ours, like raising our hands, will cause an infinite series of traveling particles to ripple through this post-heat-death universe, interfering from time to time with the systems that fluctuate into existence, including those sibling galaxies.

If all of this is true, then any event that has a finite chance of occurring as a result of being perturbed by one of these successor particles will in fact eventually occur as a result of having been perturbed by one of these successor particles. The probability might be mind-bogglingly minute! But we have infinitude to play with.

Consider a googol: 10^{100}. That's well over a trillion times as many particles as are estimated to exist in the observable portion of the universe. What a tiny number! A googolplex puts a googol to shame. Instead of ten raised merely to a hundred, it's ten raised to the googol: $10^{10^{100}}$. But this too is minuscule. We laugh at its smallness. How about a "power tower" of googolplexes – a googolplex raised to the googolplex raised to the googolplex raised to the googolplex… a googolplex times, or $10^{10^{100}} \uparrow\uparrow 10^{10^{100}}$, as it is sometimes notated.[22] Let's call this number a Vast. If the events discussed here happen once in a Vast years, that's eyeblink-

[22] Knuth 1976. Still not satisfied? How about a Vast ↑↑… [Vastly many arrows here] … ↑↑ Vast. Call that a Boggle – still tinier than almost all finite numbers. If we need to wait a Boggle years for an outcome, no sweat.

frequent compared to infinitude – or rather, of course, even more relatively frequent than that, if we're truly comparing to infinitude, which we are. These are the kinds of magnitudes we have in mind, not mere lifetime-of-the-galaxy magnitudes.

A successor particle from your hand-raising just now will eventually hit a system it will perturb in such a way that a person will live who would otherwise have died. At some point, a galaxy will fluctuate into existence containing an Earth-like planet populated with human-like people, containing a radio telescope which the successor particle strikes, causing a bit of information to appear on a sensitive device. This bit of information pushes the device over a threshold needed to trigger an alert to a waiting scientist, who now pauses to study the device rather than send the email she was composing. Because she didn't send that email, a certain fateful hiking trip is postponed and the scientist does not fall to her death, which she would have done but for your particle. However improbable all of this is, one improbability stacked on another stacked on another, there is no reason to think that any of this is less than finitely probable. Thus, given the assumptions above, it will occur, eventually, with virtually 100% probability. You saved her! Let's pause for a celebratory toast.

Of course, there is another scientist you killed. There are wars you started and peaces you precipitated. There are great acts of heroism you enabled, children you brought into existence, plagues you caused, great works of poetry that would never have been written but for your intervention, and so on. It would be bizarre to think you deserve any credit or blame for this. You didn't cause them in the sense of intending them or being what residents of those worlds would regard as among the primary causes worth describing in their history books. However, in another sense you did cause them. None of these events would have happened to the people they did in fact happen to, had it not been for the raising of your arm. And there is an unbroken chain of physical processes from the moment of your arm's going up to those various future events. Your arm raising isn't a proximal cause but rather a "distal" cause – very distal indeed – but a cause nonetheless.

If the goodness or badness of your actions is measured by their positive or negative effects, as in standard consequentialist ethics, then under the current set of cosmological assumptions the utility of every action you do will be $\infty + -\infty$.[23]

Our framework puts a few important limitations on the types of future events you will cause to occur. They must be finitely probable, less than galaxy-sized (though there's room for negotiation here), and not specified in a way that would make them unique. With those caveats, almost everything you do causes almost everything.

6. Signaling Across the Vastness.

The following will also almost certainly occur, given our assumptions so far: On some far distant post-heat-death counterpart of Earth will exist a counterpart of you – let's call that person *you-prime* – with the following properties: You-prime will think "right hand" after the ripple from the act of your raising your right hand arrives at their world, and you-prime will not have thought "right hand" had that ripple not arrived at their world. Maybe the ripple initiates a process that affects the weather which causes a slightly different growing season for grapes, which causes small nutritional differences in you-prime's diet, which causes one set of neurons to fire rather than another at some particular moment when you-prime happens to be thinking

[23] See Lenman 2000 for an argument against consequentialism on approximately these grounds, but in a finite, Earth-bound version, and Nye 2014 for similar concerns about constraints against harmful actions. See also Bostrom 2011 on infinite ethics. It might be that such effects would be symmetrical and canceling in the sense of Chapter 4, §10. See also Lenman 2000 vs Greaves 2016 on the issue of canceling.

about their hands. Likewise, there's a future you-prime who would have thought "A" if you, here on our Earth, had held up a sheet with that letter and not otherwise. Indeed, infinitely many future counterparts of you have that property. You can specify the message as precisely as you wish, within the bounds of what a counterpart of you could possibly think. Some you-prime will think, "Whoa! Infinite causation!" as a result of your having raised your hand and would not have done so otherwise.

These message recipients will mostly not believe that they have been signaled to. However, we can dispel their disbelief by choosing the fraction who, for whatever reason, are such that they believe they are receiving a signal if and only if they do in fact receive a signal. We can stipulate that we're interested in you-primes who share the property that when your signal arrives they think not only the content of the signal but also "Ah, finally that signal I've been waiting for from my earlier counterpart."[24]

There's a question of whether one of your future counterparts could *rationally* think such a thought. But maybe they could, if they had the right network of surrounding beliefs, and if those beliefs were themselves reasonably arrived at. We'll consider one such set of beliefs in Section 8.

7. Infinite Puppetry.

You needn't limit yourself to ordinary communicative signals. You can also control your future counterparts' actions. Consider future counterparts with the following property: They will raise their right hand if you raise your right hand, and they will not raise their right hand if you

[24] Compare this procedure with Sinhababu's 2008 procedure for writing love letters between possible worlds. One advantage of our method over Sinhababu's is that there actually is a causal connection.

do not. Exactly which counterparts have this feature will depend on exactly when you raise your hand and how, since that will affect which particles follow which trajectories when they are disturbed by your hand. But no matter. Whenever and however you raise your hand, such future counterparts exist.

Your counterparts' actions can be arbitrarily complex. There is a future you-prime who will, if you raise your hand, write an essay word-for-word identical with the chapter you are now reading and who will otherwise write nothing at all. Maybe that you-prime is considering whether to write some fanciful philosophy of cosmology, as their last hurrah in a failing career as a philosopher. They're leaning against. However, the arriving particle triggers a series of events that causes an internet outage that prevents them from pursuing an alternate plan, so they do write the essay after all. (A much greater proportion[25] of such future counterparts, of course, will write very different essays from this one, but we can focus on the tiny fraction of them who create word-for-word duplicates of this essay.)

Let's call someone a *puppet* if they perform action A as a consequence of your having performed an action (such as raising your hand) with the intention of eventually causing a future person to perform action A. (Admittedly, you might need to agree with the assumptions of this chapter to be able to successfully form such an intention.) You can now wave your hand around with any of a variety of intentions for your future counterparts' actions, and an infinite number of these future counterparts will act accordingly – puppets, in the just-defined sense.

We recommend that you intend for good things to happen. This might seem silly, since if the assumptions of this chapter are correct, almost every type of finitely probable, non-unique

[25] Here and throughout we bracket quibbles about ratios of infinitude by considering the limit of the ratio of counterparts with property A to counterparts with property B as the region of spacetime defined by your forward lightcone goes to infinity.

future event occurs, regardless of your benevolent or malevolent intent right now. Still, there is a type of good event that can occur as a result of your good intentions, which could not otherwise occur. That's the event of a good thing happening in the far distant future as a consequence of your raising your hand *with the intention of causing* that future good event. So let's choose benevolence, letting good future events be intentionally caused while bad future events are merely foreseen side effects.

A deeper kind of puppet mastery would involve influencing a person's actions through a sequence of moves over time and with some robustness to variations in the details of execution. This might not be possible on the current set of assumptions. Raising your right hand, you can trigger arbitrarily long sequences of actions in some future you-prime. But if you then raise your left hand, there's no guarantee that a ripple of particles from your left hand will also hit the same you-prime. Maybe all the ripples from your right hand head off toward regions A, B, and C of the future universe and all the ripples from your left hand head off toward regions D, E, and F. Similarly, if you raise your right hand like *this*, the ripples might head toward regions A, B, and C, and if you raise it instead like *that*, they head toward regions G, H, and I. So there might be no future counterparts of you who do what you intend if you raise your right hand now and then do what you intend when you raise your left hand later; and there might be no future counterparts who will do what you intend if you raise your right hand now, insensitively to the particular manner in which you raise it. In this way, there might be no sequencing and no implementational robustness to your puppetry.

Sequential and robust puppetry might only be reliably possible if we change one of the assumptions in this chapter. Suppose that although the universe endures infinitely in time, spatially it repeats – that is, it has a closed topology in the sense we described in Section 1 – so that any particle that travels far enough in one direction eventually returns to the spatial region from which it originated, as if traveling on the surface of a sphere. Suppose, further, that in this finite space, every ripple eventually intersects every other ripple infinitely often. Over the course of infinite time each ripple eventually traverses the whole of space infinitely many times; none get permanently stuck in regions or rhythms that prevent them from all repeatedly meeting each other. (If a few do get stuck, we can deal with them using the n^m strategy of Section 4. Also the rate of ripple stoppage would presumably increase with so much intersection, but hopefully again in a way that's manageable with the n^m strategy.) When you raise your right hand, the ripples initially head toward regions A, B, and C; when you raise your left hand, they initially head toward regions D, E, and F; but eventually those ripples meet.

With these changed assumptions, we can now find future counterparts who raise their right hands as a result of your raising your right hand and who then afterward raise their left hand as a result of your afterward raising your left hand. We simply look at the infinite series of systems that are perturbed by both ripples. Eventually some will contain counterparts of you who raise their right hands, then their left, as a result of that joint perturbation. In a similar way, we can find implementationally robust puppets: counterparts living in systems that are perturbed by your actual raising of your right hand (via the ripple that initially traversed regions A, B, and C) and which are also such that they *would have* been perturbed had you, counterfactually, raised your hand in a somewhat different way (via the ripple that would have initially traversed regions G, H, and I). Multiplying the minuscule-but-finite upon the miniscule-but-finite, we can now find puppets whose behavioral matching to yours is long and implementationally robust, within reasonable error tolerances.

8. We Might All Be Puppets.

So far, we have not assumed that anything existed before the Big Bang. But if the universe is infinite in duration, with infinitely many future sibling galaxies, it would be in a sense surprising if the Big Bang were the beginning. It would be surprising because it would make us amazingly special, in violation of the Copernican Principle of cosmology, which holds that our position in the cosmos is not special or unusual. We would be special in being *so close* to the beginning of the infinite cosmos. Within the first *14 billion years,* out of infinity! It's as though you had a lotto jar with infinitely many balls numbered 1, 2, 3… and you somehow managed to pull out a ball with the low, low number of 14 billion. If you don't like a strictly infinite lotto, consider instead a Vast one. The odds of pulling a number as low as 14 billion in a fair lottery from one to a Vastness are far less than one in a googolplex.[26]

Cosmologists don't ordinarily deny that there might have been something before the Big Bang. Plenty of theories posit that the Big Bang originated from something prior, though there's no consensus on these theories.[27] If we assume that somehow the Big Bang was brought into existence by a prior process, and that process in turn had something prior to it, and so on, then the Copernican lottery problem disappears. We're in the middle of a series, not at the beginning of one. Maybe Big Bangs can be seeded in one way or another. Heck, maybe the whole

[26] Our reasoning here resembles the reasoning in the "Doomsday argument", e.g., Gott 1993, according to which it's highly unlikely that we're very near the beginning of a huge run of cosmological observers. For a bit more detail, see Schwitzgebel 2022b. For another related perspective, see Huemer 2021.

[27] See notes 12 and 13 for references. A note on terminology: "Prior" sounds kind of like "earlier" but is also more general in that there's a sense in which one thing can be ontologically prior to another, or ground it, or give rise to it, even if they one doesn't temporally precede the other (e.g., an object is prior to its features, or noumena are prior to phenomena [see Chapter 5]). Possibly, temporal priority is a relationship that only holds among events within our post-Big Bang universe while whatever gave rise to the Big Bang stands in some broader priority relationship to us.

observable universe is a simulation nested in a whole different spatial reality (Chapters 4 and 5) or is itself a very large fluctuation from a prior heat-death state.

Suppose, then, that we are in the middle of an infinite series rather than at the beginning of one, conjoining Copernican mediocrity with an infinite future. If so, and if we can trace chains of causation or contingency infinitely backward up the line, and if a few other assumptions hold, then eventually we ought to find our puppeteers – entities who act with the intention of causing people to do what we are now doing and whose intentions are effective in the sense that had they not performed those actions, we would not be here doing those things. Suppose you are knitting your brow right now. Somewhere in the infinite past, there is a near-duplicate counterpart of you with the following properties: They are knitting their brow. They are doing so with the intention of initiating ripples that cause later counterparts of them to knit their brows. And you are just such a later counterpart, because among the events that led up to your knitting your brow, absent which you wouldn't have knit your brow, was a ripple from that past counterpart.

We the authors of this chapter – Eric and Jacob – can work ourselves into the mood of finding this probable. An infinite cosmos is simpler, more elegant, and more consistent with cosmological theory; if it's infinite, it's probably infinite in all directions; and if it's truly infinite in all directions, there will be bizarre consequences of that infinitude. Puppetry is one such consequence. We would not be so special as to be only puppeteers and never puppet. It seems only fair to our future puppets to acknowledge this.