



CHICAGO JOURNALS



Statistical Mechanics and the Asymmetry of Counterfactual Dependence

Author(s): Adam Elga

Source: *Philosophy of Science*, Vol. 68, No. 3, Supplement: Proceedings of the 2000 Biennial Meeting of the Philosophy of Science Association. Part I: Contributed Papers (Sep., 2001), pp. S313-S324

Published by: [The University of Chicago Press](#) on behalf of the [Philosophy of Science Association](#)

Stable URL: <http://www.jstor.org/stable/3080955>

Accessed: 27/09/2013 13:37

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press and Philosophy of Science Association are collaborating with JSTOR to digitize, preserve and extend access to Philosophy of Science.

<http://www.jstor.org>

Statistical Mechanics and the Asymmetry of Counterfactual Dependence

Adam Elga^{†‡}

Massachusetts Institute of Technology

In “Counterfactual Dependence and Time’s Arrow”, David Lewis defends an analysis of counterfactuals intended to yield the *asymmetry of counterfactual dependence*: that later affairs depend counterfactually on earlier ones, and not the other way around. I argue that careful attention to the dynamical properties of thermodynamically irreversible processes shows that in many ordinary cases, Lewis’s analysis fails to yield this asymmetry. Furthermore, the analysis fails in an instructive way: it teaches us something about the connection between the asymmetry of overdetermination and the asymmetry of entropy.

1. Introduction. The *asymmetry of counterfactual dependence* is (roughly) the thesis that later affairs depend counterfactually on earlier ones, and not the other way around. This asymmetry seems to be a feature of our counterfactual talk (at least in many contexts), which is reason enough to seek an analysis of counterfactuals that reproduces it. There are other reasons. Many analyses of causation rely on counterfactuals that exhibit such an asymmetry (see, e.g., Lewis 1973; Ramachandran 1997). Causal decision theory does, too (see Gibbard and Harper 1978; Lewis 1981). And one might invoke the asymmetry of counterfactual dependence to

[†]Send requests for reprints to the author, Department of Linguistics and Philosophy, MIT, 77 Massachusetts Ave., Cambridge, MA 02139; email: adam@mit.edu.

[‡]Thanks to Ned Hall, Robert Stalnaker, Sarah McGrath, and Anthony Newman, to conference audiences at Princeton University, the University of Western Ontario, and the 2000 meeting of the Philosophy of Science Association, to attendees of the M.A.T.T.I. group at MIT, and to David Albert (for a great seminar on the direction of time). Thanks to Norm Margolus, both for helpful discussion and for kind assistance on using the CAM8 computing architecture. Thanks also to the Josephine de Kármán Fellowship Trust for research support.

Philosophy of Science, 68 (Proceedings) pp. S313–S324. 0031-8248/2001/68supp-0026\$0.00
Copyright 2001 by the Philosophy of Science Association. All rights reserved.

explain an apparent asymmetry of *openness*: the feeling that while the one true past is fixed, many different alternative futures lie ahead.

David Lewis (1986) defends an analysis of counterfactuals intended to yield this asymmetry. Lewis is loath to rule out backward causation and future-to-past counterfactual dependence a priori. So his analysis doesn't have any time asymmetry built in. Instead, it is designed to yield the desired asymmetry when combined with a contingent feature of the world he calls the *asymmetry of overdetermination*.

This paper applies some reasoning from statistical mechanics to Lewis's analysis. It turns out that in many cases that involve thermodynamically irreversible processes, Lewis's analysis fails. Furthermore, the analysis fails in an instructive way: it teaches us something about the connection between the asymmetry of overdetermination and the asymmetry of entropy.

2. Lewis's Analysis.

2.1 For present purposes we can take Lewis's analysis of counterfactuals to be the following:¹

The counterfactual "If *A* were true, then *C* would be true" is true if and only if *C* is true at the *A*-world that is most similar to the actual world. (An *A*-world is a world at which *A* is true.)

To make the discussion concrete, consider the following example. At 8:00, Gretta cracked open an egg onto a hot frying pan. According to the analysis, are the following counterfactuals true?

- (1) If Gretta hadn't cracked the egg, then at 8:05 there wouldn't have been a cooked egg on the pan.
- (2) If Gretta hadn't cracked the egg, then at 7:55 she wouldn't have taken an egg out of her refrigerator.

To answer, we must first ask, of the *no-crack worlds* (worlds in which Gretta doesn't crack the egg), which one is *closest* (i.e., which one is most similar to the actual world)? In order for the analysis to yield the asymmetry of counterfactual dependence for this choice of antecedent, it has to turn out that the closest no-crack world is one in which

- (A) history before 8:00 is almost exactly like actual history before 8:00, and
- (B) history after 8:00 differs significantly from actual history after 8:00.

1. The analysis given isn't Lewis's analysis exactly (Lewis complicates matters to deal with the case in which there is no most similar *A*-world), but the added complications don't matter for present purposes. So I'll use Stalnaker's (1968) analysis.

If the closest no-crack world meets these conditions, then counterfactuals such as (1)—i.e., those describing how matters after 8:00 would be different if matters at 8:00 were different—will often turn out true, but counterfactuals such as (2)—i.e., those describing how matters *before* 8:00 would be different if matters at 8:00 were different—will almost never turn out true.

So the crucial question is: Does the closest no-crack world meet conditions (A) and (B)?

2.2 Lewis stipulates that the following criteria² determine how similar a given world is to the actual world:

- (I) “It is of the first importance to avoid big, widespread, diverse violations of [actual] law.”
- (II) “It is of the second importance to maximize the spatio-temporal region throughout which perfect match of particular fact [with the actual world] prevails.” (Lewis 1986, 47)

Let us follow Lewis in assuming that the laws of nature are deterministic, in order to explore how the asymmetry of counterfactual dependence might arise even under such laws. In other words, let us assume that the state of the world at one time, together with the (fundamental dynamical) laws, determines the state of the world at all other times.

To see how criteria (I) and (II) work under deterministic laws, imagine that you are a god looking down at all of actual history. Your job is to perform the minimal modification that results in a world in which Gretta does not crack the egg.

One strategy is to make modifications so that in the resulting world, (i) Gretta doesn't crack the egg, and (ii) no actual laws are violated. Let W_1 be the resulting world. Since W_1 differs from the actual world at 8:00, and since no laws are violated at W_1 , it follows from the assumption of determinism that W_1 differs from the actual world at all times (see Figure 1).

Another strategy is to introduce a tiny *miracle* (violation of actual law) shortly before 8:00. The strategy is to leave everything before the miracle untouched but to have the state just after the miracle evolve (via the laws) into a future in which Gretta doesn't crack the egg. (Perhaps the miracle is that a few extra neurons fire in Gretta's brain, getting her to put the egg back in her refrigerator rather than crack it.) Let W_2 be the resulting world.

How do these first two strategies compare? W_1 and W_2 are on a par as far as (I) goes: neither contains big, diverse violations of actual law. But W_2 beats W_1 on criterion (II): while no spatio-temporal region of W_1

2. I list only the first two members of Lewis's list of criteria since the other two play no role in the following discussion.




Figure 1. Comparison of three no-crack worlds. The straight solid line labeled “@” represents history in the actual world. W_1 (thin solid line) contains no violations of actual law, and so differs from the actual world at all times. W_2 (shaded line) matches actual history until shortly before 8:00, at which time a small divergence miracle (marked by a star) occurs. Thereafter, history in W_2 diverges from actual history. W_3 (dotted line) differs from actual history up until shortly after 8:00, at which time a convergence miracle (marked by a many-pointed star) occurs. Thereafter, history in W_3 perfectly matches actual history.

matches the actual world, the whole region before the miracle in W_2 matches the actual world.

A third strategy is to introduce a miracle shortly after 8:00. The strategy is to leave everything after the miracle untouched but to have the state just before the miracle evolve backwards (via the laws) into a past in which Gretta doesn’t crack the egg. Let W_3 be the resulting world. Whether this third strategy is promising depends on how big of a miracle is required.

Lewis thinks that a very big, diverse, widespread miracle is required to implement the third strategy. Here’s his idea: Suppose that the miracle in W_3 occurs at 8:05. Then W_3 matches the actual world perfectly after 8:05. In the actual world after 8:05, there are many *traces* of Gretta’s having cracked an egg: Gretta has memories of cracking an egg, there are bits of cooked egg stuck to the pan, and so on. We may even suppose that Gretta’s voyeuristic neighbor videotaped the cracking of the egg. So after 8:05, W_3 also contains all of those varied traces.

That’s what W_3 looks like after 8:05. What about before 8:05? W_3 is a world in which Gretta *doesn’t* crack the egg. So in W_3 right before 8:05, there *aren’t* any traces of her having cracked the egg. (In §5 we’ll see that the argument commits a crucial error at this step.) Yet we just saw that in W_3 right after 8:05 there are *tons* of traces³ of her having cracked the

3. For convenience I use a non-factive sense of “trace.” That is, I use “trace” in such a way that from the fact that there are traces of an explosion it does not follow that there was an explosion. I use “memory” similarly.

egg. So the miracle in W_3 has to take care of *making* all of those (misleading) traces, and that requires doctoring Gretta's memories, the bits of egg stuck to the pan, the neighbor's videotape, and so on. That's enough doctoring to require a big, widespread, diverse miracle.

If all of this is right, then the second strategy (ensuring that Gretta doesn't crack the egg by putting in a miracle before 8:00) has a giant advantage over the third strategy (ensuring that Gretta doesn't crack the egg by putting in a miracle after 8:00). The purported advantage is an instance of an alleged *asymmetry of miracles*:

Ensuring that Gretta doesn't crack the egg by putting in a miracle *before* 8:00 requires only a *tiny* miracle, but ensuring that Gretta doesn't crack the egg by putting in a miracle *after* 8:00 requires a *huge* miracle.

If there is such an asymmetry, then W_2 counts as closer than W_3 because W_2 contains only a tiny miracle while W_3 contains a huge one. Granting for the sake of argument that there are no other likely candidates, it follows that the closest no-crack world is a world such as W_2 —one whose history perfectly matches actual history until shortly before 8:00.

Recall that this is just the result needed in order for the analysis to yield the asymmetry of counterfactual dependence. For then it will turn out that if Gretta hadn't cracked the egg, history after 8:00 would have been different (potentially very different) than it actually is. And it will turn out that if Gretta hadn't cracked the egg, history before 8:00 would have been exactly the same as it actually is, except for a small transition period immediately preceding 8:00.

Note that this whole account rests on the asymmetry of miracles. If the boxed statement above is false—if somehow the third strategy can be implemented with a small miracle—then there is no reason to think that W_2 is closer than W_3 , and hence no reason to think that the analysis yields the asymmetry of counterfactual dependence.

The boxed statement above *is* false.

The third strategy *can* be implemented with a small miracle.

It will take a little statistical mechanics to see why.

3. Sensitivity to Final Conditions.

3.1 To keep things simple, pretend that the laws of nature are the laws of Newtonian mechanics.⁴ Then to specify the state of the world at a time

4. The discussion remains in relevant respects the same if we consider more sophisticated (deterministic) laws. When such laws are in play, a more complicated transformation plays the role that velocity-reversal plays in the upcoming discussion.

it is sufficient to specify the positions and momenta of all of the particles that exist at that time. The set of physically possible states is called *phase space*. A specification of how the particles move over time corresponds to a trajectory through phase space. (Each point on the trajectory corresponds to the state the system is in at a particular time.)

Let S_0 be the state of the world at 8:00—a state in which Gretta is about to crack the egg into the pan. Over the course of five minutes, S_0 evolves into S_1 , a state in which the egg is sitting on the pan, cooked.

The S_0 -to- S_1 process The egg oozes out of the cracked shell and drops down towards the pan, where it splats on the pan, making a noise, slightly heating up the surrounding air, and setting up some vibrations in the pan. Then the egg cooks by absorbing heat from the pan.

To repeat: starting with S_0 and running the laws forwards for five minutes results in S_1 . We can also look at things in the other temporal direction: starting with S_1 and running the laws *backwards* for five minutes results in S_0 .

The aim of this section is to show that the process that gets from S_1 to S_0 by running the laws backwards is extremely sensitive to certain small changes in S_1 .

We've assumed that the laws are deterministic, and so assumed that any given state has a unique lawful past just as it has a unique lawful future. Nevertheless, it is easier to think about running the laws forwards than it is to think about running them backwards. So we will investigate the process of running the laws backwards to get from S_1 to S_0 indirectly, by investigating the following closely related process.

Let Z_1 be the *velocity-reverse* of S_1 —the result of reversing the velocities of all of the particles in S_1 . Like S_1 , Z_1 is a state in which the cooked egg sits on the pan. But Z_1 has an unusual future: the particle motions that result from starting with Z_1 and running the laws *forwards* are exactly the motions that result from starting with S_1 and running the laws *backwards*. In other words, the five minutes that lawfully follow Z_1 involve the egg uncooking and then jumping back up into its shell. The resulting state (at the end of the five minutes) is Z_0 , the velocity-reverse of S_0 .

The Z_1 -to- Z_0 process The cooked egg uncooks by giving up heat to the (already very hot) pan. Meanwhile, molecules in the pan start to coordinate to form a pattern of vibration converging on the center of the pan. Air molecules around the room begin to form a series of spherical waves that converge on the pan. Just as the egg finishes uncooking, the coordinated action of these inward-directed air waves and pan vibrations congeals the egg into a round shape and propels it vertically towards the waiting open shell, which then seals around it.

As far as positions of particles go, the process that gets from S_1 to S_0 by running the laws *backwards* is exactly like the process that gets from Z_1 to Z_0 by running the laws *forwards*. So in order to show that the first process is sensitive to certain small changes in S_1 , it is enough to show that the second process is sensitive to certain corresponding changes in Z_1 .

3.2 A process in which a cooked egg sits on a pan and gradually cools is unremarkable. In contrast, a process in which an egg spontaneously uncooks and jumps back up into its shell (such as the Z_1 -to- Z_0 process) is *amazing*. We would be shocked if such a process were to occur. Yet both processes are perfectly in accord with the (fundamental dynamical) laws.

Let COOKED be the set of states that are exactly like Z_1 with respect to coarse-grained macroscopic parameters (such as temperature and pressure distribution). All of the states in COOKED are ones in which a cooked egg sits on a pan; these states differ from each other only in microscopic respects.

Some of the states in COOKED (such as S_1) have futures in which the egg acts in ways that we would consider thermodynamically normal: for example, futures in which the egg just sits there and cools. The rest of the states in COOKED (such as Z_1) have futures in which the egg acts in ways that we would consider thermodynamically abnormal: for example, futures in which the egg uncooks and jumps into the air. Let AB be the set of members of COOKED with abnormal futures.⁵

There is a tradition dating back to Boltzmann of trying to explain, for example, the scarcity of instances of egg-uncooking-and-jumpings by appealing to the scarcity of states in COOKED with abnormal futures. More precisely, the tradition appeals to the following fact:

AB occupies only a tiny part of the phase space volume occupied by COOKED (on a natural way of measuring such volumes⁶).

For present purposes, it doesn't matter whether the fact that AB is so tiny can serve the purposes of the Boltzmannian tradition. What matters is simply that AB is so tiny (see, e.g., Penrose 1989; Price 1996; Sklar 1993). Even more important than the size of AB is its shape. AB does *not* consist of a single compact blob. Instead it consists of a dispersed pattern of miniscule specks and thin strands. Z_1 is a member of AB, and so it sits on one of those specks or strands. Since the specks are so miniscule and the

5. Here my terminology follows the terminology in Albert 1994.

6. For the details on how to measure the volumes of regions of phase space, see Huang 1963 or Sklar 1993.

strands so thin, almost all of the states near Z_1 in phase space are *not* members of AB.⁷

A small change in phase space location (for example, a change that gets from Z_1 to one of Z_1 's neighbors) might correspond to two sorts of changes to particles. It might correspond to slight changes in the states of many particles. Or it might correspond to slight changes to the states of just a small, localized bunch of particles. Call this second sort of change a *small-miracle change*. The purpose of the preceding discussion is to make plausible the following empirical claim:

Since so little of the phase-space neighborhood of Z_1 is within AB, some small-miracle change of Z_1 results in a point outside of AB.

This claim tells us something about the sensitivity of the Z_1 -to- Z_0 process. While Z_1 itself has an *abnormal* future (one in which the egg uncooks and jumps back into its shell), most of the states *near* Z_1 have *normal* futures—ones in which the egg just sits on the pan, cooling. And some of these neighboring states differ from Z_1 merely by a small-miracle change. We already knew that starting with Z_1 and running the laws forwards yields a process in which the egg uncooks and jumps back into its shell. What we just learned is that starting with Z_1 , *making the right sort of small-miracle change*, and *then* running the laws forwards yields a very different process—one in which the egg just sits on the pan and never jumps back into its shell.

3.3 It is worth making vivid the manner in which such a small difference in initial state can lead to such dramatically different processes. To do so we'll compare the Z_1 -to- Z_0 process with a modified process whose starting state differs from Z_1 by a small-miracle change. For concreteness, suppose that the starting state of the modified process can be gotten from Z_1 by slightly changing the positions of a small bunch of molecules in the pan. Suppose also that the starting state of the modified process has a normal future.

The two processes start out much the same: in both, a cooked egg begins to uncook and air molecules engage in complicated preliminary motions that in the Z_1 -to- Z_0 process will lead to them forming converging spherical waves. But there is a tiny difference between the processes in the motion of a few molecules of the pan.

In the modified process, the pan molecules whose positions were changed bump into neighboring molecules, making the trajectories of those neighbors differ from the trajectories of their counterparts in the Z_1 -to- Z_0 process. The complicated patterns that in the Z_1 -to- Z_0 process lead

7. David Albert puts this fact to different use in Albert 1994.

to the formation of inwardly directed vibrational waves are, in the modified process, disrupted by these changed trajectories. The disruption spreads: In the modified process, air molecules bump into the surface of the pan in slightly different ways than they do in the Z_1 -to- Z_0 process, making them move in ways that increasingly differ from the ways their counterparts move in the Z_1 -to- Z_0 process. These disrupted air molecules alter the trajectories of their neighbors, interfering with the coordinated motion needed to form inwardly directed airwaves.

The upshot is that in the modified process, the inwardly directed air waves and pan vibrations never form. So while in the Z_1 -to- Z_0 process the uncooked egg ends up being propelled back up into its shell, in the modified process the egg just sits on the pan.

3.4 The whole point of investigating the process that gets from Z_1 to Z_0 by running the laws forwards is to shed light on the process that gets from S_1 to S_0 by running the laws backwards. The main lesson—that the Z_1 -to- Z_0 process is very sensitive to certain small changes in Z_1 —leads immediately to a corresponding lesson about the S_1 -to- S_0 process.

Suppose that a small-miracle change gets from Z_1 to Z'_1 , a state with a future in which the egg just sits on the pan. Then a corresponding change gets from S_1 to S'_1 , a state with a *past* in which the egg just sits on the pan. In other words, the past history of S'_1 is one in which the egg was never cracked onto the pan.

4. There Is No Asymmetry of Miracles.

4.1 Go back to being a god looking at all of actual history. We know that you can guarantee that Gretta doesn't crack the egg by inserting a small miracle before 8:00, and then evolving that modified state forwards according to the laws. The crucial question is: Can you guarantee that Gretta doesn't crack the egg by inserting a small miracle *after* 8:00, and then evolving that modified state backwards according to the laws? At the end of §2, we saw that if the answer is yes, there is no reason to believe that Lewis's analysis yields the asymmetry of counterfactual dependence.

But because the S_1 -to- S_0 process is so sensitive to small changes in S_1 , the answer *is* yes. Take the actual state of the world at 8:05. Modify it by appropriately changing the positions of a few molecules in the pan, as was illustrated in the previous section (i.e., make the small-miracle change that gets from S_1 to S'_1). Now evolve this changed state backwards according to the laws. The result is a past history in which the egg never fell onto the pan, and hence is a past history in which Gretta never cracked the egg onto the pan.

Therefore, in this case there *is* no asymmetry of miracles, and hence in

this case Lewis's analysis fails to yield the asymmetry of counterfactual dependence.

4.2 The history of the actual world, from a time-reversed point of view, is quite amazing. Eggs uncook, congeal, and get propelled up into waiting shells. Coordinated waves in the earth conspire to eject meteorites out into space. Air and ground waves push rubble up the sides of a mountain in time-reversed landslides. The processes that look so amazing in reverse are the so-called thermodynamically irreversible processes—processes that are associated with increases in entropy.

In §3 we saw that the egg-uncooking-and-jumping process is *fragile*. If you are watching the history of the world backwards and see an egg starting to uncook, all it takes is a small miracle to disrupt the finely coordinated action required for the egg to completely uncook and be propelled upwards. But the point is general: many thermodynamically irreversible processes are fragile in this way.

Every thermodynamically irreversible process is sensitive to changes in its end state that correspond to small changes in phase space location. Whether such a process is sensitive to a change in its end state that corresponds to a *small-miracle* change depends on the degree of coupling between the parts of the system that undergoes the process. Make a small-miracle change to the end state of a process and run the laws backwards. Certainly the change disrupts the coordinated movement of the process in the neighborhood of the change. If the parts of the system are strongly coupled, then the region “infected” by the change (i.e., the region containing particles whose trajectories are greatly altered from what they would have been without the change) will grow rapidly.⁸

Many ordinary thermodynamically irreversible processes are strongly coupled in this way, and so are sensitive to small-miracle changes in their final conditions. (Examples include the processes of milk mixing into coffee, glasses shattering, water boiling, and balloons popping.) So the violation of the asymmetry of miracles described in §4 is no fluke—similar violations arise in many other mundane cases.

5. Traces.

5.1 The world that we've used to make trouble for Lewis's analysis is W_3 , a world gotten from the actual world by inserting an appropriate small miracle at 8:05 and evolving the changed state backwards according to the laws. W_3 makes trouble for the analysis because it is a no-crack world that (i) contains only a tiny violation of actual law and (ii) matches the actual world perfectly after 8:05.

8. I borrow “infected region” terminology from Tim Maudlin, who puts it to different use in Maudlin 1998.

In §2.2 I gave some (faulty) reasoning concluding that there is no such world. We're now in a position to see how that reasoning goes wrong. Here is the reasoning:

1. At 8:05 the actual world contains traces of Gretta's having cracked the egg.
2. So immediately after 8:05 W_3 also contains those traces.
3. But since W_3 is a world in which Gretta doesn't crack the egg, immediately before 8:05 W_3 does not contain those traces.
4. Therefore a large trace-manufacturing miracle occurs in W_3 at 8:05.

The error is in step 3. Though Gretta doesn't crack the egg in W_3 , at 8:05 W_3 is filled with false traces to the contrary.

Let's look at how those extremely misleading traces are formed in W_3 . We'll start by seeing what W_3 looks like from a future-to-past perspective. Then we'll take the time-reverse to see what W_3 looks like from an ordinary (past-to-future) perspective.

Start with the state of the actual world at 8:05. The cooked egg sits on the pan. Apply an appropriate small-miracle change by altering the positions of a small bunch of molecules of the pan. The resulting state is the state of W_3 immediately before 8:05. Now run time backwards. We already saw that the egg just sits on the pan (since the miracle interferes with the formation of the coordinated waves required to congeal the egg together and propel it upward). The history of the egg as we continue to run time backwards looks increasingly like the (past-to-future) history of an ordinary egg. The egg cools and eventually rots.

More generally, the situation is as follows. At 8:05, W_3 matches the actual world except for a tiny infected region (the region in which the miracle occurs). As we run time backwards, the infected region rapidly expands. Within that region, what we see looks (from our backwards-in-time vantage point) thermodynamically typical. (For example, eggs get more rotten as time gets earlier.) Outside of that region, events look thermodynamically *reversed*. (For example, eggs get less rotten as time gets earlier.)

Now look at W_3 in the ordinary (past-to-future) direction. In the distant past, the infected region is huge. Within that region are events that look thermodynamically reversed. Events outside the infected region look thermodynamically typical.

At 8:00 the infected region includes a cooked egg sitting on a pan in Gretta's kitchen. Over the next five minutes, the egg warms up, and by 8:05 is in a state that would suggest that it had been recently cracked on that pan and cooked. This suggestion is entirely misleading. The egg was never raw and was never cracked onto the pan. Long ago, the egg formed

in the pan as a puddle of rotten slime, and reached its 8:05 state by a process of reverse-rotting.

All of the traces in W_3 of the egg-cracking were formed in such ways. Each such trace was formed by an anti-thermodynamic process—by what would seem to be a finely tuned conspiracy in the motions of large numbers of microscopic particles.

5.2 In W_3 at 8:05, all signs point to Gretta's having cracked the egg. The cooked egg sits on the pan. Gretta remembers having cracked the egg onto the pan. The neighbor's videotape encodes images of Gretta cracking an egg. Offhand it might seem as though the only lawful way for all of those traces to have been formed is by Gretta's having cracked the egg at 8:00. More generally, one might agree with Lewis that there is an asymmetry of *overdetermination*—that “very many simultaneous disjoint combinations of traces of any present fact are determinants thereof; there is no lawful way for the combination to have come about in the absence of the fact” (Lewis 1986, 50).

But that's all wrong. It only takes a small miracle to make the difference between the actual world (a world with many veridical traces of the egg-cracking) and W_3 (a world in which those same traces are all highly misleading). In general, the existence of apparent traces of an event (together with the laws, and together with the absence of evidence that those traces have been faked) falls far short of entailing that the event occurred.⁹

REFERENCES

- Albert, David Z (1994), “Quantum Mechanics and the Approach to Thermodynamic Equilibrium”, *British Journal for the Philosophy of Science* 45: 669–677.
- (2001), *Time and Chance*. Boston: Harvard University Press, forthcoming.
- Gibbard, Allan and William Harper (1978), “Counterfactuals and Two Kinds of Expected Utility”, in J. J. Leach, C. A. Hooker, and E. F. McClennen (eds.), *Foundations and Applications of Decision Theory*. Dordrecht: Reidel, 125–162.
- Huang, Kerson (1963), *Statistical Mechanics*. New York: John Wiley & Sons.
- Lewis, David (1973), “Causation”, *Journal of Philosophy* 70: 556–567.
- (1981), “Causal Decision Theory”, *Australasian Journal of Philosophy* 59: 5–30.
- (1986), “Counterfactual Dependence and Time's Arrow”, in *Philosophical Papers*. Oxford: Oxford University Press, 32–66.
- Maudlin, Tim (1998), “A Modest Proposal Concerning Laws, Counterfactuals, and Explanations”, manuscript.
- Penrose, Roger (1989), *The Emperor's New Mind*. Oxford: Oxford University Press.
- Price, Huw (1996), *Time's Arrow and Archimedes' Point*. New York: Oxford University Press.
- Ramachandran, Murali (1997), “A Counterfactual Analysis of Causation”, *Mind* 106: 263–277.
- Sklar, Lawrence (1993), *Physics and Chance: Philosophical Issues in the Foundations of Statistical Mechanics*. New York: Cambridge University Press.
- Stalnaker, Robert (1968), “A Theory of Conditionals”, in Nicholas Rescher (ed.), *Studies in Logical Theory*. Oxford: Blackwell, 98–112.

9. David Albert (2001) makes a similar observation.