Philosophy and
Phenomenological Research

Philosophy and Phenomenological Research
Vol. LXXVIII No. 2, March 2009

Must the Fundamental Laws of Physics be Complete?

MARC LANGE
University of North Carolina, Chapel Hill

> The beauty of electricity, or of any other force, is not that the power is mysterious and unexpected, touching every sense at unawares in turn, but that it is under *law*…
>
> Michael Faraday, *Wheatstone's Electric Telegraph's Relation to Science (being an argument in favour of the full recognition of Science as a branch of Education)*, 1854

1. Introduction

Faraday made this remark even before the laws of electricity had all been discovered. Nevertheless, he was utterly confident that all electric phenomena are covered by laws and that all other forces are too—indeed, apparently, that the laws cover every kind of situation that every possible kind of thing can manage to get into. The laws are "complete;" there are no gaps in their coverage.

If the laws are in fact complete, then is this merely a peculiarity of the actual laws? Or is the laws's completeness metaphysically compulsory? That is the question I shall investigate.

In section 2, I shall offer a provisional characterization of what it would be for the laws to be complete. In section 3, I will try to capture some reasons for supposing that the laws must be complete. I will also argue that neither David Lewis's best-system account nor David Armstrong's relations-among-universals account of laws entails that the laws must be complete. If the laws's completeness is an important part of the idea that the laws govern the universe, then an important (but heretofore neglected) criterion of adequacy for any metaphysical analysis of natural law is that it account for the laws's completeness. Since standard analyses of natural law fail to do so, there is an opportunity for some other analysis to do better. In section 4, I shall sketch an alternative account of laws (which I have developed at greater length elsewhere), and in section 5, I shall extend this account so that it is equipped to entail that the laws must be complete. I shall derive this corollary in section 6.

By this route, I shall argue for the metaphysical necessity of the completeness of the laws—or, more precisely, of the laws of fundamental physics. I do not maintain that the laws of some special science (if there be any such laws) or that the laws of some branch of physics (such as low-temperature physics or thermodynamics) must be complete. That I am concerned only with the laws of fundamental physics will ultimately play a role in my argument, but not for a while. Until then, I will refer simply to the "laws."

2. What Would it Be for the Laws to be Complete?

Intuitively, everything that happens is not only logically consistent with the laws, but also covered by the laws. The laws govern the outcome of every process; no entity or behavior falls beyond the scope of their sovereignty. Nothing operates outside the laws or above the laws. That is to say, the laws are "complete." But beyond all of the metaphors, what does the laws's completeness really amount to?

For the laws to "cover" a fact, it is not necessary that the fact follow logically from the laws. After all, Hempel's D-N model of scientific explanation is a covering-law model, and yet the fact being explained may be entailed not by the laws alone, but only by the laws together with some accidental truth (an "initial condition"). For that matter, in covering-law models of statistical explanation, the fact being explained is not even entailed by the laws and the initial conditions; only its objective chance is so entailed. Nevertheless, the fact being explained is thereby covered by the laws.

What would a *gap* in the laws's coverage be? For the sake of having a definite, simple example to examine, let's suppose that the laws specify (i) that everything consists entirely of elementary particles of certain kinds (A-ons, B-ons, etc.); (ii) how elementary particles behave when they are not undergoing any interaction; (iii) the chances that various kinds of particles in various circumstances will interact and, if they do, the chances of various results—and nothing more. In particular, suppose that the laws entail nothing about the result of an A-on's interacting with a B-on when the two particles are separated by between 1 nm and 2 nm. No law prohibits A-ons from being 1-2 nm from B-ons. Indeed, the laws specify the chance of two such particles's interacting—just nothing about what might result (not even the chances of various outcomes). To accentuate this gap in the laws's coverage, we might even add that the laws do specify the chances of various results of an A-on interacting with a B-on at distances less than 1 nm or greater than 2 nm. The only gap is between 1 and 2 nm. There is likewise a gap if the laws specify that any A-B interaction must yield one (α) or the other (β) of two possible results that are logically exclusive but not logically exhaustive—but fail to specify anything further when the two particles are 1-2 nm apart (such as either result's chance under various conditions).

In each of these cases, the laws are incomplete because they fail to cover the interaction of an A-on with a B-on at a distance of 1-2 nm. The laws would not explain the outcome of any such interaction. Whether any such interaction ever in fact occurs is irrelevant to the laws's incompleteness; it suffices that such an interaction is physically possible.

This example of a gap in the laws's coverage suggests what it would be for the laws to be complete. Let's begin by distinguishing the "bare facts"—bare, that is, of any reference to which of them are laws and which are accidents. These are the facts that could be governed by laws but do not concern which facts are (or aren't) the laws: the "sub-nomic" facts. They include that all emeralds are green (a law) as well as that all gold cubes are smaller than a cubic mile (an accident). But the fact that it is *a law* that all emeralds are green is not sub-nomic; neither is the fact that it is *not a law* that all gold cubes are smaller than a cubic mile. The sub-nomic facts include facts about single-case objective chances but not facts about which of those facts hold as a matter of law and which are accidental. For example, it is a sub-nomic fact that every atom of Polonium-210, at each moment it exists, has a 50% chance of surviving for the next 138.39 days (the isotope's half-life), but that this is a law is not a sub-nomic fact. I shall henceforth reserve lower-case English letters (such as *p*) for claims that, if true, state sub-nomic facts.

Let Λ be the set of facts consisting of every truth *m* where it is a law that *m*—that is, the set containing all and only the laws governing sub-nomic facts.[1] The laws are complete if and only if Λ "covers" every sub-nomic fact. The idea behind "coverage" here seems to be the covering-law conception of scientific explanation. Accordingly, let's say

[1] Hence, some of the natural laws may not belong to Λ. A law that governs *other laws* (rather than sub-nomic facts), such as the law that all laws are time-displacement symmetric, is excluded from Λ. I presume that the logical, conceptual, mathematical, and metaphysical necessities are included "by courtesy" in Λ and that Λ is logically closed. Note: When I say that Λ is true, that Λ entails *p*, etc., I shall mean that every member of Λ is true, that Λ's members (taken together) entail *p*, etc.

that the laws are complete if and only if in any history allowed by the laws, every event has a covering-law explanation. Here is the way that I will provisionally cash out this requirement. (I will reconsider this scheme in the final section of the paper.) The laws are "complete" if and only if

> given any hypothetical world-history allowed by the laws (i.e., where Λ is true) and
>
> given any hypothetical event E (letting e be that E occurs),

that history contains certain events

> not involving chances and
>
> all occurring at or before some moment T preceding the time with which E is concerned[2]

(letting h be that those events occur) such that one of the following is true:

> (h & Λ) logically entails e,
>
> (h & Λ) logically entails $\sim e$, or
>
> there is some N such that (h & Λ) logically entails $\mathrm{ch}_T(e) = \mathrm{N}$ (i.e., that at T, E's chance of occurring is N).[3]

Roughly speaking, the laws are complete exactly when for any physically possible history, every hypothetical event E's occurrence (or non-occurrence) in that history is explained by the laws, where the explanation involves the laws together with certain initial conditions entailing E's occurrence (or non-occurrence) or at least E's chance at T. The laws are complete exactly when every actual event is "covered" by them and this broad coverage is no accident. That the laws are complete might be considered a weakening of determinism.

The gap in the laws's coverage in my A-B example precludes the laws there from qualifying as complete. Although h & Λ may entail

2 This notion of completeness is suited only to a universe with absolute time. No matter: I shall suggest later that different law-governed universes have different completeness principles.

3 I shall presume that E is not vague. Otherwise $\mathrm{ch}_T(e)$ might be vague. The demands of "completeness" might be extended to accommodate this possibility.

that an A-on will be within 1-2 nm of a B-on at a certain moment, or even also that they will interact at that moment, nevertheless h & Λ neither entails whether α will result in this case nor assigns some chance at a given moment to α's doing so.

That the natural laws are complete should not be confused with

> Lawful-magnitude principle (LMP): For any time T and any proposition e that some event will occur at a subsequent time, if it is the case that $ch_T(e) = N$, then the laws of nature plus the history of instantiations of categorical properties at and before T logically entail that $ch_T(e) = N$ (Lewis 1994: 230-1, Schaffer 2003: 36-7).

Whereas LMP starts with whatever chances there may be at a given moment and demands that they be fixed by the laws and the non-chancy history up to and including that moment, completeness demands that there be chances in the first place: for every event at a given moment, the laws plus the history through some time preceding that moment must fix its chance at that time (or, in the extreme case, that the event will come to pass). That every chance there is must be fixed in accordance with LMP does not ensure completeness, since there may be too few chances for the laws to be complete. However, if in every possible world allowed by the laws, LMP holds and there is a well-defined $ch_T(p)$ for any proposition p and any time T, then completeness holds.[4]

Here is a final way to understand the laws's "completeness" as I have just defined it. According to a famous quip, on an English conception of (civil and criminal) law, everything is permitted that isn't expressly forbidden, whereas on a Prussian conception, everything is forbidden that isn't expressly permitted.[5] The two conceptions of law involve different defaults. But if the laws of nature are complete, then everything (as far as sub-nomic facts are concerned) is either expressly

4 (a) Completeness does not entail LMP since it suffices for completeness that every event have a chance at some earlier moment that is fixed by the laws and the non-chancy history through that moment. There may be additional chances as well that are not so fixed, violating LMP. (b) It is not clear to me whether Lewis is inclined to grant that for any proposition p and any time T, there is a well-defined $ch_T(p)$. He writes, "It is only caution, not any definite reason to think otherwise, that stops me from assuming that chance of truth applies to any proposition whatever" (1986: 91). However, Lewis may be saying merely that it makes sense to speak of $ch_T(p)$ for any p, rather than that there actually exists a $ch_T(p)$ for any p and T. As I will explain shortly, it seems to me that Lewis's "Best-System Account" of laws and chances permits there to be p's and T's for which there is no well-defined $ch_T(p)$.

5 Van Fraassen (1989: 171) credits this aphorism to Oliver Wendell Holmes.

forbidden or expressly permitted. There is no default; the laws plus some history through a given moment expressly categorize every hypothetical *E*. If the laws are complete, then they are both English and Prussian.

3. Must the Laws be Complete?

Science appears to presume that the laws are complete. When scientists discover a new phenomenon, they try to find some laws that explain it by covering it. That the phenomenon has no covering-law explanation is not a hypothesis that is taken seriously; that various proffered covering-law explanations have failed is not regarded as confirming that the phenomenon is governed by no laws at all. For example, before Newtonian natural philosophers had discovered any of the actual laws governing electric and magnetic interactions, they presumed that there were such laws rather than that such interactions were ungoverned by any laws.[6] Likewise, as Feynman (1967: 151) notes, if ESP were verified, then since ESP is not a consequence of the known laws, its discovery would show that physics is incomplete and lead physicists to seek the laws governing ESP. But, again, the discovery of ESP would not lead physicists to take seriously the possibility that there are no laws covering it.

If the actual laws are in fact complete, then we must consider whether their completeness is just a notable feature of the actual universe or metaphysically compulsory. Of course, we could ask the same question regarding the existence of some natural laws rather than none at all. Let's keep these two questions separate by asking: *If* there are laws, must they be complete? (A possible world where there are no laws would obviously be one where the laws are incomplete—but of an uninteresting kind.) Is the laws's completeness a corollary of what it *is* to be a law?

The laws of nature are sometimes characterized as the rules of the "game" that is played by the world's various inhabitants (particles, fields, or whatever).[7] Of course, the rules of a game typically are not so complete as to dictate the move that a player makes (or specify the chance that the player will make a certain move) in a given situation. That would make for a boring game. But the rules of a game are supposed to govern play in that they are supposed to specify, for any

---

6 Today physicists seek a "theory of everything" (TOE). That sounds like it would be complete!

7 "The chess-board is the world, the pieces are the phenomena of the universe, the rules of the game are what we call the laws of Nature." (Huxley 1893: III, 82). See also Feynman 1967: 36, 59.

circumstances that might arise, which moves are allowed in that circumstance and which are not. Imagine a game with pieces that are moved around on a board. Perhaps one rule of the game says that the "fortress" can move diagonally when it is next to a "cardinal," another rule says that the number of spaces the fortress must move in a given turn is equal to the number on which the die lands on that turn, and so forth. Suppose, however, that through a sequence of lawful moves from the arrangement of pieces at the start of the game as required by the rules, it is possible for a fortress to find itself no longer next to a cardinal, but the rules specify nothing about the directions it can move in that situation. The game is then fundamentally flawed, its rules incomplete. (Playground strife sometimes ensues when my son and his friends discover that the rules of the game they had concocted a few minutes earlier are incomplete.) But the natural laws, as nature's rules, are not supposed to be flawed in this way. They are supposed to cover every possible eventuality that might arise.[8]

Likewise, the transition rules of a cellular automaton (such as Conway's "game of life") are frequently characterized as its "laws of nature" (e.g., Gardner 1970). Such rules are complete: for each of the $2^9 = 512$ possible patterns of occupation of a 3x3 grid, the rules specify whether or not the central square is occupied at the next time step.

A related metaphor understands the natural laws as "the software of the universe," directing the functioning of the hardware (particles, fields, or whatever).[9] Incomplete laws would be analogous to a computer program afflicted with a "bug:" it calls upon a subroutine that isn't there. In our earlier example with a gap for A-ons interacting with B-ons 1-2 nm away, the cosmic software contains code (in BASIC!) something like the following:

8 On the other hand, if a game's rules are incomplete but no cases ever fall into a gap, then their incompleteness might never lead to problems. Games where the rules are made up as the players go along may typically have incomplete rules and nevertheless work adequately. Perhaps an *ideal* game has complete rules, and the laws of nature, as the rules of nature's game, are supposed to be ideal. Furthermore, even if an ideal game's rules must cover "every possible eventuality," this range may be narrower than the circumstances that are logically consistent with the rules. For instance, the rules of a ball game may be complete yet fail to cover a case where a ball turns into a bird and flies away. (This example was suggested to me by Bill Lycan, who believed it was Wittgenstein's.) That is not one of the "possible eventualities" because the game presupposes certain background conditions that are not entailed by its rules. Presumably, any such "background conditions" for nature's game are among the laws of nature (understood as per note 1).

9 See, for example, Davies 1995: 256 and Dorato 2005.

```
⋮

540 REM I = 0 MEANS NO INTERACTION, R IS THE
    PARTICLES'S SEPARATION IN NANOMETERS

550 IF I = 0 GOTO 1200

560 IF R ≤ 1 GOTO 1300

570 IF R > 1 AND R < 2 GOTO 1400

⋮
```

But there is no step 1400. The program is incomplete. Unlike a computer program, the universe cannot "crash," so the cosmic software cannot contain such a "bug." The laws must be complete.

A widespread belief that the laws must be complete may also play a role in the way that science approaches cases where the "laws" break down. Consider a singularity, i.e., a situation where a physical quantity figuring in the putative laws fails to be well-defined, preventing the "laws" from yielding physically meaningful predictions (even statistical ones) regarding that situation. For example, according to classical electromagnetic theory, it is a law that the electric field at a given point P at a given moment T is equal to the vector sum of contributions from all of the charges in the universe. A charge's contribution is proportional to its magnitude and (in the simplest case) inversely proportional to the square of the distance from P to that charge's "retarded position" (i.e., its location at the moment when light leaving the charge's location at that moment would arrive at P at T). The contribution's direction is the direction from P away from the charge's retarded position. Hence, if there is a charged point body at P at T, then its contribution to the electric field there then is infinite (since at T the charge's distance from P is zero, and the charge divided by zero squared is infinity) and has no well-defined direction (since every direction from P is away from P). Hence, the "law" I just gave breaks down for point charges. It has a gap there.

However, physicists generally do not regard singularities (such as those arising when general relativity is applied to a black hole, where there is infinite density, pressure, and spacetime curvature) as indicating that the actual laws of nature are really incomplete (i.e., really have nothing to say about certain physically possible conditions). Rather, the alleged "laws" (whether of classical electromagnetism or general relativity) are generally held not to be quite accurate; the genuine laws contain no singularities because they are complete. (For example, perhaps quantum mechanics or its successor is needed to deal adequately with the contribution of a point charge to the electric field at its location, or perhaps quantum mechanics or its successor precludes a finite quantity of charge or mass from being packed into a point.)

This attitude toward singularities seems to have been Einstein's:

> It seems Einstein always was of the opinion that singularities in a classical field theory are intolerable. They are intolerable from the point of view of classical field theory because a singular region represents a breakdown of the postulated laws of nature. I think that one can turn this argument around and say that a theory that involves singularities and involves them unavoidably, moreover, carries within itself the seeds of its own destruction … (Bergmann 1980: 156)

That is, the theory does not give the genuine laws of nature. This attitude is quite common[10] and is surely one motive for Penrose's and Hawking's "cosmic censorship" hypothesis, a version of which is roughly that there are no naked singularities (i.e., no singularities that are able to affect the outside universe because they are neither at the end of time nor safely hidden behind event horizons) other than perhaps at the beginning of time. The laws of physics say nothing about what a naked singularity would spew forth into the rest of the universe, nothing even about its chances of emitting various things. As Earman remarks:

> The principles of classical GTR [general theory of relativity] do not tell us whether a naked singularity will passively absorb whatever falls into it or will regurgitate helter-skelter TV sets, green slime, or God only knows what. (Earman 1995: 94; cf. 65-6)

Clearly, then, the physical possibility of a naked singularity would make the laws incomplete.[11] Here's Earman again:

> Perhaps one can also argue that violations of cosmic censorship would show that classical GTR is incomplete in a stronger sense. The premise required is not that determinism holds but the weaker premise that all physical processes be law governed. The argument would be completed by showing that classical GTR places no constraints, not even statistical ones, on what can emerge from a naked singularity. (Earman 1995: 225)

---

10 Earman (1995) contains many passages evincing this attitude (including the remark I quoted from Bergmann). Earman, however, regards it as "a pious hope that some quantum theory of gravity, yet to be formulated, will contain mechanisms for the avoidance of singularities" (224), though he says that it remains an open question whether GTR contains built-in mechanisms for avoiding naked singularities (225).

11 Indeed, even the possibility of a clothed singularity would reveal the laws to be incomplete, since events occurring in the spacetime region within the event horizon would fail to be covered by the laws.

With the laws's "completeness," I have tried to capture the premise that all physical processes are "law governed" and to give some intuitive motivations for believing it metaphysically compulsory (if there are laws at all).[12] (Don't worry: You haven't heard the last about green slime!)

However, familiar philosophical accounts of natural law do not entail that the laws must be complete. Consider, for example, Lewis's (1994) "Best System Account," according to which the laws are the generalizations in the deductive system of truths and history-to-chance conditionals with the optimal combination of simplicity, informativeness, and fit to the Humean mosaic, and the chances at a given moment are exactly those entailed by that "Best System" coupled with the history of categorical-property instantiations through that moment. There is no reason why the Best System has to be rich enough to be complete. Although a system that was complete would be mighty informative (and could fit the Humean mosaic well), it need not be very simple. An incomplete system would be less informative but could be simpler—so much so as to make it better overall than any complete system. There would then be no profit in filling its gaps. For example, suppose there exist many bodies of "gas," each body characterized by two fundamental quantities, P and V. Suppose that over the world's history, there are many bodies of gas where $P \leq 500$ or $V \leq 500$ (in some units), and for each, $P = V$. There are only a few bodies of gas where $P > 500$ and $V > 500$, but they do not obey $P = V$. Rather, they fall into no simple pattern. A complicated polynomial curve could be fit through them, but it would presumably be far better to include

12 There are now some wonderful arguments (Earman 1986, Laraudogoitia 1996, Norton forthcoming) showing that surprisingly, a world "governed" by all and only the laws of Newtonian physics violates determinism: the full set of initial conditions and laws allow certain events to occur (such as "space invaders" swooping in from infinity and particles suddenly starting to move) but also leave room for them not to occur. Nor are these events assigned any specific chances by the laws and initial conditions. Thus, if (as I shall argue) the laws must be complete, then what these wonderful arguments show is not that a world governed by exactly the Newtonian laws is indeterministic, but rather that no such world is possible. It might be objected that a Newtonian world certainly *seems* possible. I agree; it does seem possible—until the arguments from Earman *et al.* reveal that in such a world, the laws unexpectedly have gaps. Of course, as I mentioned earlier, there are other well-known singularities in Newtonian physics; they might already have suggested that Newtonian "laws" could not have been all of the laws. The arguments from Earman *et al.* that Newtonian physics is not deterministic reveal additional sorts of events uncovered by the "laws" that could occur in a world supposedly governed by Newtonian laws. I regard a Newtonian world as weirder even than a world governed by statistical laws: it is metaphysically impossible because its laws are incomplete. (See also the penultimate paragraph of the paper.)

"P = V for P ≤ 500 or V ≤ 500" in the system and to leave unfilled the gap above this threshold.

So the Best System Account does not make the laws's completeness mandatory.[13] Armstrong's (1983) account of laws as relations of "nomic necessitation" among universals likewise fails to entail that the laws must be complete. These relations are contingent; even given which universals exist, it is not metaphysically compulsory that the nomic-necessitation relations among them be rich enough to make the laws complete.[14]

That the laws are complete seems to be an important part of the idea that the laws govern the universe. Since standard analyses of natural law fail to account for the laws's completeness, I shall now propose an analysis that seems better equipped to do so.

4. The Stability of the Laws

In science, the natural laws are called upon to tell us what would have been the case, had things been different in some physically possible way. For example, had Uranus's axis not been so nearly aligned with its orbital plane, then conditions on Uranus would have been quite different, but the laws of nature would still have been laws (which is *why* conditions on Uranus would have been so different). Scientific practice thereby suggests a principle that I shall call "Nomic Preservation" (NP):

> NP: for any counterfactual supposition *q*, the laws would still have been laws, had *q* been the case—as long as *q* is physically possible, i.e. logically consistent with Λ.

(Recall that I have reserved lower-case English letters for claims that, if true, state "sub-nomic" facts—that is, facts that could be governed by laws but do not concern which facts are the laws and which facts are not. Λ is the set of facts consisting of every truth *m* where it is a law that *m*.) Although the truth-values of counterfactual conditionals are

13 Of course, Lewis's account could be amended to require that any system eligible for the competition for Best be complete. Or Lewis's account could be amended so that completeness joins informativeness, simplicity, and fit as among the desiderata an optimal satisfaction of which makes a system "Best." Loewer (2004: 1118) may have something like this in mind. But it would be far better for the laws's completeness to fall out nicely as a corollary of some integral part of the account of laws rather than to be inserted expressly into the account "by hand."

14 In a personal communication, David Armstrong kindly acknowledged that on his account, the laws do not have to be complete. He added, however, that he "should think worse of the world if there actually is 'incompleteness'."

notoriously context-sensitive, NP is intended to hold in all contexts, since it purports to capture the *logical* relation between laws and counterfactuals, and logic is not context-sensitive. Principles roughly like NP have been defended by Bennett (1984), Carroll (1994), Chisholm (1946), Goodman (1983), Horwich (1987), Jackson (1977), and many others.

If lawhood is not context sensitive, then NP entails that if it is a law that *m*, then (in any conversational context) *m* would still have been true, had *q* been the case—for any *q* that is logically consistent with Λ. No accident would still have been true under every such *q*, since ∼*m* is physically possible if *m* is an accident (and obviously *m* is not preserved under the counterfactual supposition that ∼*m*). However, this principle

> It is a law that *m* if and only if (in any context) *m* would still have been true, had *q* been the case (i.e., $q \mathrel{\Box\!\!\rightarrow} m$), for every *q* that is logically consistent with Λ

cannot reveal what it is to be a law, since the laws appear on both sides of the "if and only if." That is, the above principle uses the laws to pick out the range of counterfactual suppositions invariance under which sets the laws apart. It would be circular to distinguish the laws as the truths *m* that would still have held under all counterfactual suppositions that are logically consistent with *the laws*!

Elsewhere (Lange 1999, 2000, 2002, 2005a) I have suggested a way to avoid this circularity. The range of counterfactual suppositions considered by NP (namely, every physically possible *q*) is designed expressly to suit the laws. What if any logically closed set of truths was allowed to pick out for itself a convenient range of counterfactual suppositions: those with which the set is logically consistent? Let us call the set "stable" exactly when the set's members would all still have held (whatever the context), under every such counterfactual supposition.

More precisely: Consider a non-empty, logically closed set Γ of subnomic facts. Define

> Γ is *stable* exactly when for any member *m* of Γ and any *q* where Γ ∪{*q*} is logically consistent, the subjunctive conditionals (which will be counterfactuals if *q* is false) $q \mathrel{\Box\!\!\rightarrow} m$ hold in any context.

The intuitions behind NP (manifested in scientific practice) suggest that Λ is stable. In contrast, the logical closure of (e.g.) Reichenbach's favorite accident, "All gold cubes are smaller than a cubic mile," is unstable, since had Bill Gates wanted to build a gold cube exceeding a cubic mile, I dare say there would have been such a cube.

It is *nearly* true, I have argued (Lange 2000, 2005a), that stability distinguishes Λ from *any* set of sub-nomic truths containing accidents in that no set containing an accidental truth is stable. (Perhaps there is one exception: the set of *all* sub-nomic facts, which is *trivially* stable if counterfactuals obey "Centering" (the principle that $(p \Box\!\!\rightarrow q)$ holds if p and q hold), since no counterfactual supposition q is logically consistent with all sub-nomic facts.[15]) For example, consider the accident g: whenever the gas pedal of a certain car is depressed by x inches and the car is on a dry, flat road, then the car's acceleration is $f(x)$. Had the gas pedal on a certain occasion been pressed a little farther, g would still have held. However, a set containing g is unstable unless it also includes a description of the car's engine, since had the engine contained six cylinders instead of four, $\sim g$ might have held. (If the set includes a description of the car's engine, then the counterfactual supposition positing six cylinders is logically inconsistent with the set, and so the set does not have to be invariant under that supposition in order to qualify as stable.) But now to be stable, the set must also include a description of the engine factory, since had the factory been different, the engine might have been different. (With a description of the factory in the set, a counterfactual supposing the factory to have been different in some respect is logically inconsistent with the set.) By packing more and more into the set, will we ever arrive at a stable set containing g before we have reached the set containing *all* sub-nomic facts? The prospects seem dim indeed.

Consider a logically closed set containing g but omitting the fact that I am not wearing an orange shirt. Here is a counterfactual supposition that is logically consistent with every member of the set: "Had it been the case that either $\sim g$ or I wear an orange shirt." What would the world have been like then? Of course, as I have mentioned, counterfactual conditionals are notoriously context sensitive. In Quine's famous example, it is correct in some conversational contexts that had Caesar been in command in the Korean War, he would have used the atomic bomb. In other conversational contexts, it is correct that he would have used catapults. What about "Had either $\sim g$ or I wear an orange shirt"? Would g still have held (and so I have worn an orange shirt)? "No" is the correct answer in at least some conversational

15 Lewis-Stalnaker semantics for counterfactuals endorses Centering. I will reject Centering near the close of this paper, but my denial of Centering plays no role in the present argument.

contexts. In those contexts, it is *not* the case that the truth *in* the set we are talking about would still have held, had either it or an arbitrary truth *out* of that set been false. That is enough to make the set unstable. But if a set containing *g* must include even an arbitrary sub-nomic fact in order to be stable, then presumably the set must include *all* sub-nomic facts in order to be stable.

On this view, what makes the laws special, as far as their range of invariance under counterfactual suppositions is concerned, is that they are stable: all of the laws would still have held under every counterfactual supposition under which they *could* all still have held (i.e., every supposition with which they are all logically consistent). No set containing an accident can make that boast non-trivially. A stable set is *maximally* resilient under counterfactual perturbations; it has as much invariance under counterfactual suppositions as it could logically possibly have.

That lawhood consists of membership in a stable set that is not the set of all sub-nomic truths (i.e., a nonmaximal stable set) suggests why there is a variety of *necessity* that the laws alone possess. Intuitively, "necessity" involves an especially strong sort of persistence under counterfactual perturbations. But not every fact that would still have held, under even a wide range of counterfactual perturbations, qualifies as possessing some species of "necessity." Being necessary is supposed to be *qualitatively* different from merely being invariant under a wide range of counterfactual suppositions. Because the set of laws is *maximally* resilient—as resilient as it could logically possibly be—its members possess a variety of necessity.

Here is another argument for this analysis of necessity. Suppose that *q* is possible and that *p* would have held, had *q* been the case. Then intuitively, *p* must be possible: whatever would have happened, had something possible happened, must also qualify as possible. Now suppose that the necessities of some particular variety (such as the physical necessities) are exactly the members of some particular logically closed set of truths. What must that set be like in order to respect the above principle? It says that if *q* is possible (that is to say, logically consistent with every member of the relevant set) and if *p* would have held, had *q* been the case, then *p* must be possible (that is, logically consistent with every member of that set). That is immediately guaranteed if the set is stable. (If *q* is logically consistent with every member of a given stable set, then under the counterfactual supposition that *q* holds, every member of that set would still have held, and so anything else that would *also* have been the case must join the members of that set and therefore must be logically consistent with them.)

However, look what happens if a logically closed but *unstable* set of truths contains exactly the necessities of some variety. Because the set is unstable, there is a counterfactual supposition q that is logically consistent with every member of the set but where some member m of the set would not still have held under this supposition. That is to say, m's negation might have held. But m, being a member of the set, is supposed to be necessary, so m's negation is an impossibility. Therefore, if an unstable set contains exactly the necessities (of some variety), then had a certain possibility (of that variety) come to pass, something impossible might have happened. This result conflicts with a principle slightly broader than the one we were just looking at—namely, that whatever *might* have happened, had something possible happened, must also qualify as possible.

In short, if an unstable set contains exactly the necessities (of some variety), then though some q-world is possible, the closest q-world (or, at least, one of the optimally close q-worlds) is impossible. This conflicts with the intuition that *any possible* q-world is closer to the actual world than is *every impossible* q-world. Hence, if a logically closed set of truths contains exactly the necessities (of some variety), then that set must be *stable*.

Thus, that lawhood consists of membership in a nonmaximal stable set not only allows us to break out of the circle afflicting the earlier proposal (that the laws are the truths that would still have held under every counterfactual supposition that is logically consistent with the laws), but also accounts for the laws's possession of a species of necessity. However, I shall now suggest that this proposal requires a slight modification.

To begin with, it fails to account for some of the counterfactual suppositions under which the laws would still have held true—namely, nested counterfactual suppositions. For example, we believe not only that had we tried to accelerate a body from rest to beyond the speed of light, we would have failed, but also that had we access to 23rd century technology, then had we tried to accelerate a body from rest to superluminal speed, we would have failed. As defined above, the laws's "stability" does not ensure that the laws are preserved under such nested counterfactuals, since stability requires the truth of various counterfactuals of the form $q \Box\!\!\rightarrow m$, not the truth of any counterfactuals of the form $p \Box\!\!\rightarrow (q \Box\!\!\rightarrow m)$.

Initially, nested counterfactuals may appear remote from actual scientific practice. But in fact, scientists routinely employ them ("Had the chamber been completely evacuated, then had it contained a few CO_2 molecules, they would have had a long mean free path;" "Had gravity declined with the cube of the distance, then a solar system, had it begun with many planets, would not long have so remained; they would have soon escaped or spiraled into the sun''). As the first of these examples illustrates, $p \; \Box\!\!\rightarrow \; (q \; \Box\!\!\rightarrow \; m)$ is not in general logically equivalent to $(p \;\&\; q) \; \Box\!\!\rightarrow \; m$.[16]

Another reason for including nested counterfactuals in the definition of ''stability'' is to foreclose further the possibility of a nonmaximal stable set that contains accidents. If there is such a set under the definition of ''stability'' I gave earlier, then I suggest that its invariance under all of those counterfactual suppositions is just a fluke. That is, although each of the set's members m would still have held under every q that is logically consistent with the set, that *invariance* is *not* likewise invariant: although every $q \; \Box\!\!\rightarrow \; m$ holds, there is some r for which $r \; \Box\!\!\rightarrow \; (q \; \Box\!\!\rightarrow \; m)$ does not hold. (Or if every $r \; \Box\!\!\rightarrow \; (q \; \Box\!\!\rightarrow \; m)$ holds, then *that* invariance is not invariant under further iterated counterfactuals ....)

Furthermore, the principle ''Whatever would (or might) have happened, had something possible happened, must also qualify as possible,'' to which I appealed earlier, generalizes to nested counterfactuals: ''Had something possible happened, then whatever would (or might) have happened, had something possible happened, must also qualify as possible.''

Accordingly, I suggest amending the definition of ''stability'' as follows:

> Consider a non-empty set Γ of sub-nomic truths containing every sub-nomic logical consequence of its members. Γ is *stable* exactly when for any member m of Γ (and in every conversational context),
>
> $p \; \Box\!\!\rightarrow \; m$,
>
> $q \; \Box\!\!\rightarrow \; (p \; \Box\!\!\rightarrow \; m)$,
>
> $r \; \Box\!\!\rightarrow \; (q \; \Box\!\!\rightarrow \; (p \; \Box\!\!\rightarrow \; m))$, ...
>
> obtain for any p, q, r, ... where $\Gamma \cup \{p\}$ is logically consistent, $\Gamma \cup \{q\}$ is logically consistent, $\Gamma \cup \{r\}$ is logically consistent ....

[16] The same applies even when p is logically consistent with q. For example, suppose that you and I run a race, I win, and I would always win were I to try. Had you won, then had I tried, I would have won. This nested counterfactual is plainly not equivalent to ''Had you won and I tried, then I would have won.'' (See my 2000: 290–1)

A final, important reason for including nested counterfactuals in the definition of "stability" is that it allows us to capture part of NP that was not covered by the original definition of "stability." Recall that NP demands that under any physically possible circumstance p, the laws not only would still have been true, but also would still have been laws. (That is *why* they would still have been true.) This result follows from identifying the laws as the members of a nonmaximal stable set—but only after "stability" has been amended to include the nested counterfactuals. Suppose that m is a member of Γ, a stable set, and that q, r, s … are all logically consistent with Γ. Then by the amended definition of "stability," $q \Box\!\!\rightarrow (r \Box\!\!\rightarrow m)$, $q \Box\!\!\rightarrow (s \Box\!\!\rightarrow m)$, $q \Box\!\!\rightarrow (r \Box\!\!\rightarrow (s \Box\!\!\rightarrow m))$, etc. So in the closest q-world, these counterfactuals hold: $r \Box\!\!\rightarrow m$, $s \Box\!\!\rightarrow m$, $r \Box\!\!\rightarrow (s \Box\!\!\rightarrow m)$, etc.—exactly the counterfactuals needed for Γ to be stable in the closest q-world. Hence, if Λ is actually stable, then had q been the case, Λ would still have been stable, and so its members would still have been laws. We thereby save the intuition that had Jones missed his bus to work this morning, then the actual laws would still have been laws—and so (here comes a nested counterfactual) Jones would not have gotten to work on time had he simply clicked his heels and made a wish to get there.[17]

5. Stability Further Amended

We are almost ready to show that the laws must be complete. But to do so, we must amend our definition of "stability" one last time. Just as we had to include counterfactuals having counterfactuals as their consequents (such as $q \Box\!\!\rightarrow (r \Box\!\!\rightarrow m)$), so also we must include

17 Lewis rejects this intuition (even though it appears to reflect scientific practice). If we insist that the laws would have been no different, had Jones missed his bus, then (he argues) we must say (if the world is deterministic) that the world's state billions of years ago would have been different, had Jones missed his bus. That sounds counterintuitive. Although I cannot discuss this issue here (see my 2000), I am inclined to think that $q \Box\!\!\rightarrow m$ says (roughly speaking) not that m is true in the closest q-world, but that m is true in the relevant fragment of the closest q-world. In a context where we should not 'backtrack' in assessing counterfactuals, the relevant fragment (or 'nonmaximal situation') does not concern the events responsible for bringing q about (Lewis's 'small miracle'). Therefore, an actual law would still have been true had q, since the miracle (which violates the actual law) is 'offstage'. The world's state billions of years ago likewise stands outside of the relevant fragment of the closest q-world. So when we occupy a non-backtracking context and consider what would have happened had Jones missed his bus, we are not interested in whether the world's state billions of years ago would have been different. If we focus our attention upon q's past light cone (e.g., in discussing what a remarkable doctrine determinism is), then we enter a (backtracking) context where it is true that the world's state billions of years ago would have been different, had Jones missed his bus.

counterfactuals having counterfactuals as their antecedents, such as $(p \,\Box\!\!\rightarrow\, q) \,\Box\!\!\rightarrow\, m$.

One of our reasons for amending stability to include counterfactuals having counterfactuals as their consequents was to account for some of the counterfactual suppositions under which the laws would still have been true. The same rationale applies to amending stability to include counterfactuals having counterfactuals as their antecedents. Let's say I am holding a well-made, oxygenated, dry match. The match would have lit, had it been struck. But had it been the case that the match would *not* have lit, had it been struck, then … what? It would (have to) have been wet or incorrectly made or starved of oxygen. The actual laws of nature would still have been laws, had it been the case that the match would not have lit had it been struck. (The *reason* that the match would (have to) have been wet or incorrectly made or starved of oxygen is precisely because the actual laws of nature would still have been laws.) Thus, for Λ's stability to capture the laws's characteristic resilience under counterfactual suppositions, Λ's stability needs to require Λ's invariance under certain counterfactual suppositions that posit the truth of certain counterfactual conditionals.

While such counterfactuals might initially sound exotic, they actually are not. Consider counterfactuals with antecedents involving dispositions, such as "Had the box contained a fragile vase, then I would have taken great care not to drop it, since the vase might well have broken had the box been dropped." Even if the ascription of a disposition is not logically equivalent to some counterfactual conditional, dispositions seem to involve "threats and promises" (Goodman 1983: 40) and hence to have a counterfactual flavor. A counterfactual antecedent involving a counterfactual conditional seems no more exotic than "Had the box contained a fragile vase."

Moreover, counterfactuals with antecedents involving counterfactual conditionals are not out-of-the-ordinary in science. Consider: "Had the boiling point of the liquid in our test tube been under 300K (under standard pressure), then it would already have boiled by now." This counterfactual conditional is perfectly innocuous, yet its antecedent is (roughly) "Had it been the case that the liquid in our test tube would boil were it heated to 300K (under standard pressure).''[18]

Of course, Λ would *not* still have held had *this* counterfactual conditional been true: Had I sneezed a moment ago, then a body would have been accelerated from rest to beyond the speed of light! Of course, this counterfactual conditional is logically inconsistent with one of the counterfactual conditionals required for Λ to count as ''stable'' by the definition given in the previous section. Thus, in order for its ''stability'' to distinguish Λ from sets containing accidents, a given set Γ's ''stability'' had better not require Γ's preservation under a counterfactual supposition positing the truth of a counterfactual conditional that is logically inconsistent with one of the counterfactual conditionals

[18] Bennett (2003: 167-8) offers this example: ''Jones was not careless when he threw the lighted match onto the leaves. He knew that the leaves were too damp to ignite. If it had been the case that if he were to throw the match onto the leaves a forest fire would ensue, then he would have known this was the case and not thrown the match onto the leaves.'' In (Lange 2005b), I argue that a fact about a body's instantaneous velocity at t is a fact about what that body's trajectory would be like, were the body to remain in existence after t. Hence, an ordinary-looking counterfactual ''Had the marble's speed at t been 10 cm/s …'' actually has a counterfactual hidden in its antecedent.

Here's an example of reasoning with counterfactuals having counterfactual conditionals as their antecedents. Consider a match that is wet (so had it been struck, it would not have lit) but otherwise in propitious conditions (e.g., oxygenated). Had ''If the match had been struck, it would have lit'' been true, the match would have been dry. Had the match been dry, then ''If the match had been struck, it would have lit'' would have been true and the actual laws of nature would still have been laws. Therefore, had ''If the match had been struck, it would have lit'' been true, then the actual laws of nature would still have been laws.

Of course, the truth-values of counterfactuals with antecedents involving counterfactuals are context-sensitive just as the truth-values of other counterfactuals are. In one conversational context, it might be accurate to say ''If you would get a million dollars if you touched Jason's head, then everyone would be chasing Jason,'' whereas in another conversational context, it might be accurate to say ''If you would get a million dollars if you touched Jason's head, then everyone would be touching everyone else's head to see if that would work, too.''

Question: How is $(p \mathrel{\Box\!\!\rightarrow} q) \mathrel{\Box\!\!\rightarrow} m$ to be understood in a non-backtracking context if $(p \mathrel{\Box\!\!\rightarrow} q)$ is a backwards-directed counterfactual? What would it be for $(p \mathrel{\Box\!\!\rightarrow} q)$ to be true in a non-backtracking context? For instance, if $(p \mathrel{\Box\!\!\rightarrow} q)$ is ''Had I worn an orange shirt this morning, then Lincoln wouldn't have been assassinated,'' then it is false in a non-backtracking context, and so how can we—in a non-backtracking context—entertain the counterfactual supposition that it is true?

Answer: Easy. In a non-backtracking context, ''Had I worn an orange shirt this morning, then Lincoln wouldn't have been assassinated'' is true if Lincoln wasn't assassinated. If $(p \mathrel{\Box\!\!\rightarrow} q)$ is backwards-directed and q is actually false, then there is a way for $(p \mathrel{\Box\!\!\rightarrow} q)$ to be true in a non-backtracking context—namely, for q to be true. In a possible world where q is true, $(p \mathrel{\Box\!\!\rightarrow} q)$ holds in a non-backtracking context. So the counterfactual ''Had the counterfactual conditional 'Had I worn an orange shirt this morning, then Lincoln wouldn't have been assassinated' been true, then … ,'' entertained in a non-backtracking context, may amount simply to ''Had Lincoln not been assassinated, then … .''

required by Γ's "stability" according to the definition given in the previous section.

Hence, to capture the laws's characteristic resilience under counterfactual suppositions (and to preclude a nonmaximal set containing accidents from qualifying as "stable" on a fluke), I shall amend the definition of "stability" one last time. Let lower-case Greek letters represent claims that can be constructed exclusively out of sub-nomic claims, "∼", and "□→ ". Consider a non-empty set Γ of claims ω containing every logical consequence ς of its members. Now define:

> Γ is *stable* exactly when for any member ω of Γ, φ □→ ω holds (in any conversational context) for every φ where $\Gamma \cup \{\varphi\}$ is logically consistent (and there are such φ).

In short, Γ is stable exactly when its members are preserved together under every subjunctive or counterfactual supposition (constructible out of sub-nomic claims, "∼", and "□→ ") under which they *could* logically possibly be preserved together. Because a stable Γ is *maximally* resilient—as resilient as it could logically possibly be—its members possess a variety of necessity.[19]

We have already seen good reason to believe that every nonmaximal set containing accidents is unstable. Let's now see what sort of set would qualify as stable. Suppose it is a law that *n*. Let φ be that *n* would not still have held, had there obtained some arbitrary accident – say, had Jones missed his bus to work this morning. Of course, φ is actually false; since *n* is a law, *n* would still have held under some accident. Accordingly, had φ obtained, then *n* would not still have been a law or the laws would have been different in some other respect (so as to render φ's antecedent physically impossible). Had φ, then the laws might not still have been true. So for a set Γ containing every *m* where it is a law that *m* to be stable, Γ must also contain the counterfactual (Jones missed his bus to work this morning) □→ *n*, so that $\Gamma \cup \{\varphi\}$ is not logically consistent, and hence the laws do not have to be preserved under φ in order for Γ to be stable. Hence, for Γ to be stable, Γ must contain all of the counterfactual conditionals *p* □→ *n* for every *p* where $\Gamma \cup \{p\}$ is logically consistent. Now we can take the argument we just gave concerning Γ's member *n* and apply it instead to Γ's member (*p* □→ *n*). Let φ be that (*p* □→ *n*) would not still have held, had

[19] It is superfluous to add to our new definition that a "stable" set's members all be true: to be stable, according to the new definition, Γ must be logically consistent (since otherwise there is no φ where $\Gamma \cup \{\varphi\}$ is logically consistent) and so must be preserved (in any context) under any logical truth *p*, which precludes Γ from containing falsehoods.

there obtained a given arbitrary accident q. Had φ, then the laws might not still have been true. So to be stable, Γ must also contain the counterfactual $q \,\Box\!\!\rightarrow\, (p \,\Box\!\!\rightarrow\, n)$, so that $\Gamma \cup \{\varphi\}$ is not logically consistent, and hence the laws do not have to be preserved under φ in order for Γ to be stable. Therefore to be stable, Γ must contain all of the counterfactual conditionals $q \,\Box\!\!\rightarrow\, (p \,\Box\!\!\rightarrow\, n)$, for every p where $\Gamma \cup \{p\}$ is logically consistent. And so on for the entire cascade of multiply nested counterfactuals.

Thus, if we are building a stable set Γ by starting with the laws m, we must add all of the counterfactual conditionals in the cascade

$$p \,\Box\!\!\rightarrow\, m,$$

$$q \,\Box\!\!\rightarrow\, (p \,\Box\!\!\rightarrow\, m),$$

$$r \,\Box\!\!\rightarrow\, (q \,\Box\!\!\rightarrow\, (p \,\Box\!\!\rightarrow\, m)), \ldots$$

for every member m of Γ and every p, q, r, … where $\Gamma \cup \{p\}$ is logically consistent, $\Gamma \cup \{q\}$ is logically consistent, $\Gamma \cup \{r\}$ is logically consistent …. These are exactly the counterfactual conditionals required for stability under our previous definition.

But that's not all. We have added various subjunctive conditionals to Γ, but all of them have had sub-nomic claims as their antecedents. What about conditionals having subjunctive conditionals in their antecedents? We should also include in Γ all of the conditionals like this one concerning a dry, oxygenated match: "Had it been the case that the match would not have lit, had it been struck, then Coulomb's law would still have held" ($\varphi \,\Box\!\!\rightarrow\, m$). This conditional is true; had it been the case that the match would not have lit, had it been struck, then the match would have been wet or deoxygenated or etc., but the laws would still have held. So had $\varphi \,\Box\!\!\rightarrow\, m$ been false, then the laws would have been different; Coulomb's law would perhaps not still have held. Therefore, the stable set must include $\varphi \,\Box\!\!\rightarrow\, m$, so that the set's stability does not require its invariance under the counterfactual supposition that ($\varphi \,\Box\!\!\rightarrow\, m$) is false.

What, then, is the stable set that includes all of the laws m? It must include not only the above cascade of conditionals that had to be true for Λ to qualify as stable under our earlier definition, but also every $\varphi \,\Box\!\!\rightarrow\, m$ where φ is logically consistent with that cascade, and every $\varphi' \,\Box\!\!\rightarrow\, (\varphi \,\Box\!\!\rightarrow\, m)$ where φ is logically consistent with the cascade and φ' is too, and so forth. Let's call that set Λ^*. I suggest that it is a law that m if and only if m belongs to a stable set that does not include all of the sub-nomic truths, and that the members of Λ* have a characteristic species of necessity in virtue of Λ*'s stability.[20]

I have not suggested that Λ* is the *only* stable set. The set of broadly logical truths (together with various counterfactuals such as "Had the match been struck, then it would have been the case that all circles are round") is stable, and there is a variety of necessity that all and only its members possess, just as there is a variety of necessity possessed by all and only the members of Λ*. Furthermore, I have argued elsewhere (Lange 1999, 2000, 2005a, forthcoming) that certain other proper subsets of Λ* are also stable. For example, consider the set generated by the fundamental dynamical law(s) (which, in classical physics, is $F = ma$), the law of the composition of forces, and the conservation laws—but without the various force laws. The laws in this set would still have held, even if there had been different forces. For example, according to classical physics, if gravity had been (replaced by) an inverse-cubed force, the relation between force and acceleration would still have been $F = ma$. Thus, the laws Λ governing the sub-nomic facts may come in several strata. In any event, I am not concerned with any of Λ's proper subsets—only with Λ's completeness.

6. Why the Laws must be Complete

Now, at last, for the argument that the laws must be complete. Let's start with a quick-and-dirty version. Suppose, for the sake of *reductio*, that there is a gap in the laws's coverage. Let's return to my example from section 2, where the laws fail to cover the interaction of an A-on with a B-on at a distance of 1-2 nm, though an A-B interaction at that distance is logically consistent with the laws and the laws do specify the chances of various results of A-B interactions at distances less than 1 nm or greater than 2 nm. Suppose that in a given spatiotemporal region L, no A-B interactions actually occur at a distance of 1-2 nm. Let φ be the following counterfactual conditional: had an A-B interaction occurred in L at a distance of 1-2 nm, then all such interactions would have turned the interacting particles into green slime. Since the laws Λ are silent about the possible result of an A-B interaction at 1-2 nm, φ is logically consistent with Λ*. Hence, for Λ* to be stable, Λ*'s members must be preserved under φ.

---

[20] Since the stable set contains the cascade $p \,\Box\!\!\rightarrow\, m$, $q \,\Box\!\!\rightarrow\, (p \,\Box\!\!\rightarrow\, m)$, etc., we do not need to add to our new definition of stability that a stable set's members be preserved under *nested* counterfactuals; that the member $(p \,\Box\!\!\rightarrow\, m)$ is preserved under the non-nested supposition q automatically ensures that the member m is preserved under nested suppositions of q and p, i.e., $q \,\Box\!\!\rightarrow\, (p \,\Box\!\!\rightarrow\, m)$.

However, it does not seem that they would be. Suppose that green slime is *wildly* different from what the laws entail to result from A-B interactions at distances less than 1 nm or greater than 2 nm. Had it been the case that φ, then presumably the laws governing A-B interactions at other distances would (or at least might) have been different—in particular, have assigned chances to something like green slime being produced at those distances. Had it been the case that green slime would have been produced by all A-B interactions in L at 1-2 nm had some such interactions occurred, then green slime would perhaps have also been produced by some A-B interactions in other spacetime regions or at other distances. That is:

[(An A-B interaction occurs in L at 1-2 nm) □→

(Green slime is produced by all such interactions)]

◇→

(Some A-B interactions at other distances also produce green slime).

But the laws would have been violated, had some A-B interactions at other distances produced green slime.

Of course, A-B interactions at 1-2 nm *could* be very different from A-B interactions at other distances; there is no metaphysical obligation that the laws vary "smoothly" with distance. It suffices for my argument that in certain contexts, the laws governing A-B interactions at other distances *might* have been different, had it been the case that green slime would have been produced by all A-B interactions in L at 1-2 nm, had there been any such interactions.

This seems very plausible. After all, consider this counterfactual conditional:

[(An A-B interaction occurs at greater than 2 nm) □→

(Green slime is produced by all such interactions)]

◇→

(Some A-B interactions at other distances also produce green slime).

This counterfactual conditional seems true. Admittedly, its truth does not threaten Λ^*'s stability; its antecedent posits the truth of a counterfactual conditional that conflicts with the counterfactuals in Λ^*, and so

Λ* does not need to be preserved under this counterfactual supposition in order to be stable. But if Λ* would not have been preserved under this counterfactual supposition (involving A-B interactions beyond 2 nm yielding green slime), why would Λ* have been preserved had it been the case that green slime would have been produced by all A-B interactions in L at 1-2 nm, had there been any such interactions?

To summarize: Suppose (for *reductio*) that the laws Λ are incomplete. Consider a counterfactual conditional φ specifying that something wild would have happened, had there been a case falling into the gap. Had φ, would the laws still have held? I have suggested that the answer is not "Yes." But the answer must be "Yes" for Λ* to be stable.[21] Hence, Λ* is not stable, and so (by my earlier argument) Λ is not the set of laws. Contradiction. So the laws must be complete.

That was the quick-and-dirty argument. Now for a slightly more careful version. I will break the argument into two steps.

First step: According to Goodman (1983), if a counterfactual conditional $p \;\Box\!\!\rightarrow\; q$ holds, where p is logically consistent with the actual laws Λ, then one of two options must hold:

(i) (p & Λ) logically entails q, or

(ii) (p & Λ) does not logically entail q, but there is some actual sub-nomic fact f that is not a law and that the context implicitly invokes where (p & Λ & f) logically entails q.

To take Goodman's own example, if $p \;\Box\!\!\rightarrow\; q$ is that had I struck the match, it would have lit, then in a typical case where this counterfactual conditional is true, f is that the match is dry, well-made, surrounded by oxygen etc. On the second "Goodman option," it may even be that q is logically entailed by (p & f)—without the aid of Λ. For example, suppose that whether a given radioactive atom decays is governed only by irreducibly statistical laws. Suppose that the atom actually does decay, but I made a bet that it wouldn't and so lost. Had I bet that it would decay, then I would have won. (This counterfactual conditional is true in at least some contexts.) Here typically a suitable f is that the atom decays, and q is logically entailed by (p & f). Likewise, in

21 For two reasons, each of which is sufficient: (i) Suppose it is a law that m. Since m is a member of Λ* and φ is logically consistent with Λ*, m must be preserved under φ for Λ* to be stable—so ($\varphi \;\Box\!\!\rightarrow\; m$) must be true. (ii) To be stable, Λ* must consist exclusively of truths (see note 19), and ($\varphi \;\Box\!\!\rightarrow\; m$) is a member of Λ* (since φ is logically consistent with the cascade of conditionals that had to be true for Λ to qualify as stable under the definition from section 5)—so ($\varphi \;\Box\!\!\rightarrow\; m$) must be true.

Goodman's match example, various features of the match's state (such as the fact that the match is surrounded by oxygen) are invoked by the context, and so if $p \,\Box\!\!\rightarrow\, q$ is that had the match been struck, it would still have been surrounded by oxygen, then q is logically entailed by (p & f), without the aid of Λ, since q is f . Not every counterfactual conditional that is contingently true must be "covered" by a law.[22]

Goodman's point applies more broadly: not only must one of the two Goodman options hold if $p \,\Box\!\!\rightarrow\, q$ is in fact true, but also one of them *would* have held *had* $p \,\Box\!\!\rightarrow\, q$ been true. For example, suppose that in fact, the match is wet, so "struck $\Box\!\!\rightarrow$ lit" is false. Had "struck $\Box\!\!\rightarrow$ lit" been true (and the match not been struck), then (typically) there would then have been laws Γ and a salient accidental truth f (that the match is dry, oxygenated, etc.) such that "The match lights" is entailed by "The match is struck" together with f and Γ. Accordingly, in general: Had p been false and $p \,\Box\!\!\rightarrow\, q$ held, where p is logically consistent with whatever the laws Γ would have been had p been false and $p \,\Box\!\!\rightarrow\, q$ held, then one of these two "Goodman options" would have held:

(i) (p & Γ) logically entails q,

(ii) (p & Γ) does not logically entail q, but there is some subnomic claim f that would have held accidentally, had $\sim p$ and $p \,\Box\!\!\rightarrow\, q$ held, and that the context implicitly invokes, where (p & Γ & f) logically entails q.

Now which option would have applied to φ, had φ held?

If the first Goodman option would have applied, then had φ held, there would have been some or another "green-slime law" (such as "Every A-B interaction at 1-2 nm produces green slime"). That is: $\varphi \,\Box\!\!\rightarrow$ green-slime law.

If the second Goodman option would have applied, then had φ held, φ would have held partly by the grace of some f. But suppose we amend φ to be something like "Had an A-B interaction occurred in L at a distance of 1-2 nm *prior to which there was nothing in the universe's entire history except for an A-on and a B-on approaching each other*, then all such interactions would have turned the interacting particles

[22] Unlike Goodman, I am not presuming that what *makes* it true that $p \,\Box\!\!\rightarrow\, q$ is that one of these two options holds. I am presuming only the fact that if the counterfactual conditional is true, then one of these two options holds. I am also not suggesting (in using this fact to show that the laws must be complete) that this fact is the *reason why* the laws must be complete. Indeed, the order of explanation might run in the opposite direction.

into green slime." We might even imagine adding further to the italicized portion of the antecedent, making it completely describe the posited universe's sub-nomic history until the A-B interaction takes place. This counterfactual's antecedent (unlike "Had the match been struck") leaves no room for some *f* to supplement it in entailing its consequent. Therefore, had this φ held, then the first Goodman option would have applied to it: (φ □→ green slime law).[23]

Second step: Is it the case that had φ held and a green-slime law obtained, Λ would still have held? I suggest not (or, at least, that it is not the case that in every context, Λ would still have held). Here is the only place in my argument where it matters that we are considering the fundamental laws of physics. If Λ were just the laws of ecology (population-growth laws, laws governing the relations over time between predator and prey populations ...), for example, then a green-slime law (at least for inanimate slime!) could perhaps have obtained without disrupting them, just as they would not have been perturbed by other modest tinkerings with the fundamental microphysical laws (Lange 2002). Indeed, as I mentioned at the end of the previous section, there are even stable sets that are generated by proper subsets of the laws of fundamental physics, and a green-slime law could have obtained without disrupting them. Once again, consider the set generated by the fundamental dynamical law(s) (in classical physics: $F = ma$), the law of the composition of forces, and the conservation laws—but without the force laws. The members of this set would still have held, even if there had been a green-slime law, just as they would still have held even if gravity had been (replaced by) an inverse-cube force.

However, the laws of fundamental physics are different from the laws of ecology. Although some systems are not ecological systems, every system is a fundamental physical system. That is not because as it happens, every system is made of matter (or whatever). Rather, it is because no system falls outside of the interests of fundamental physics.

---

23 Let me emphasize why I have considered a counterfactual conditional with a counterfactual conditional φ in its antecedent rather than a more straightforward counterfactual conditional, such as "Had an A-B interaction occurred in L at 1-2 nm and produced green slime, then ... ." To suggest that under the latter counterfactual supposition, there would have been a green-slime law requires presupposing that any event must be covered by a law. To avoid begging the question in this way, I have instead appealed to the counterfactual conditional having φ in its antecedent. I do not presuppose that every counterfactual conditional that is contingently true (where the antecedent is logically consistent with the laws) must be covered by a law. As I mentioned in the main text, if a counterfactual conditional p □→ q holds, where p is logically consistent with the laws Λ, then there may be certain actual facts *f* that are not laws and that the context invokes where (p & *f*) suffices to logically entail q. (I do not even presume that a world where non-trivial counterfactual conditionals obtain must have laws.)

Whereas the laws of a special science could be excused for their failure to cover A-B interactions by the fact that A-B interactions do not fall within the science's scope, the same does not apply to the laws of fundamental physics. Hence, only the laws of fundamental physics, not the laws of some special science, are intuitively expected to be complete.

Even if the laws governing the relation of forces to motions (e.g., $F = ma$) would have withstood the addition of further forces, the actual force laws need not have done. As I argued earlier, it seems very implausible that in every context, the laws governing A-B interactions at other distances would still have held, had there been a green-slime law for A-B interactions at 1-2 nm. There are plenty of examples where the actual laws would (or, at least, might well) have failed still to be true, had there been an additional law with which they are nevertheless logically consistent. For instance, had there been a law prohibiting gold cubes larger than a cubic mile, wouldn't the force laws perhaps have been different? The actual force laws *could* still have held (the laws could merely have imposed a further constraint on, e.g., the universe's possible initial conditions) but I see no reason to insist that the actual force laws *would* still have held, had the gold-cubes generalization been a law.

Likewise, consider σ: it is a law that when two bodies of unequal mass collide, the less massive body disappears and the more massive body goes on its way as if the other body had not been there. The actual laws are logically consistent with σ: In a possible world where it is a law that that there is always just a single particle with constant mass moving uniformly forever, σ holds (vacuously) and the rest of the familiar laws of classical physics (e.g., the force laws, energy and momentum conservation) also still hold. But (in some contexts, at least, it is true that) had σ obtained, then the familiar conservation and force laws would not all still have held. Rather, there would still have been many bodies, but they would have behaved differently than bodies actually do.[24]

Similarly, consider τ: it is a law that the sum of each body's $(mv)^{1/2}$ is a conserved quantity. Again, the actual laws are logically consistent with τ: In a universe where it is a law that there is nothing but a single point body of constant mass moving uniformly forever, it is a law that $\Sigma(mv)^{1/2}$ is conserved and all of the actual conservation and dynamical laws are still laws too. But (in at least some conversational contexts, it is true that) had τ held, then some of the actual conservation laws would not still have held; the momentum conservation law or energy conservation law would presumably have been replaced by the law of

24 The failure of (σ $\Box\!\!\rightarrow$ Λ) is no threat to Λ^*'s stability since σ is not constructed exclusively out of sub-nomic claims and "$\Box\!\!\rightarrow$." Rather, σ refers to a law.

the conservation of $\Sigma(mv)^{1/2}$. In all of these examples, had an additional law held, the actual laws would have been disobliged.

Let's now gather the fruits of the two steps of this argument. Suppose (for *reductio*) that the fundamental physical laws Λ are incomplete (using the A-B example). Then

$\varphi \, \Box\!\rightarrow$ green-slime law [first step]

(φ & green-slime law) $\Diamond\!\rightarrow \; \sim\Lambda$ (holds in some context) [second step]

Therefore, by the kind of transitivity to which counterfactuals adhere,

$\varphi \, \Diamond\!\rightarrow \; \sim\Lambda$ (holds in some context)

contrary to Λ^*'s stability. *Reductio* achieved.

Let's consider an objection. I supposed that in L, no A-on actually interacts with a B-on at a distance of 1-2 nm. But this supposition seems to play no role in the argument, so let's drop it. In fact, let's suppose that there are many A-B interactions in L at 1-2 nm and all of them produce green slime—which is nevertheless quite different from what the laws for A-B interactions at other distances call for. So the subjunctive conditional φ (Were there an A-B interaction in L at 1-2 nm, then all such interactions would produce green slime) is true (since the actual world is the closest world where there is an A-B interaction in L at 1-2 nm), and its truth obviously fails to undermine any of the laws. Of course, we might select a different candidate for φ, such as "Were there an A-B interaction in L at 1-2 nm, then all such interactions would transform the interacting particles into TV sets" (where TV sets are also wildly different from the results of A-B interactions at other distances). Now φ is false. But had φ obtained, the laws would have been no different; if the laws governing A-B interactions at other distances are undaunted by green slime resulting from A-B interactions at 1-2 nm, despite green slime being wildly different from the results of A-B interactions at other distances, then why should the laws governing A-B interactions at other distances be undermined by TV sets resulting from A-B interactions at 1-2 nm? (End of objection.)

In response, I deny that in all conversational contexts, the truth of "Were there an A-B interaction in L at 1-2 nm, then all such interactions would produce green slime" is ensured simply by the fact that there is an A-B interaction in L at 1-2 nm and every such interaction produces green slime. In other words, I deny "Centering:" that

$p \,\Box\!\!\rightarrow\, q$ is true if p and q are true. Here is a counterexample to it. Suppose there is an indeterministic process (a "coin flip") having a 50% chance of yielding "heads" and a 50% chance of yielding "tails." Suppose this process will occur here sometime during the next 5 seconds. Then contrary to Centering, there is a familiar context in which it is not true that were the coin flipped here sometime during the next 5 seconds, it would land heads, and it is also not true that were it flipped here sometime during the next 5 seconds, it would land tails. Rather, were it flipped, then (just before the toss) it would have a 50% chance of landing heads and a 50% chance of landing tails, so although it might land heads, it might just as well land tails. In this context, the flip's actual outcome is irrelevant to the conditional's truth.[25]

For it to be true (in that context) that the coin would land heads were it flipped, the outcome of the coin flip must be (as it were) preordained—"fated." If the law governing the coin flip's outcome is indeterministic, then its outcome is not fated. I am not suggesting that such a "fate" can be supplied only by a law; a counterfactual conditional that is contingently true need not be "covered" by a law. However, if the coin flip falls into a gap in the laws, then the flip's outcome can be at least as unfated as it would be if the coin flip is governed by indeterministic laws. Therefore, even if there are A-B interactions in L at 1-2 nm and all of them produce green slime, there is a perfectly natural context in which "Were there an A-B interaction in L at 1-2 nm, all such interactions would produce green slime" (φ) is false. In that context, then, had φ been true, some of the laws Λ might not still have held—contrary to Λ^*'s stability, giving us the *reductio*. So the laws must be complete.

Let's look at this example in one final way. Suppose there is a law m specifying that any "coin flip" has a 50% chance of yielding heads and a 50% chance of yielding tails. Had it been the case (in the context I was discussing) that (flip in L $\Box\!\!\rightarrow$ heads) holds, then m would not still have held.[26] Indeed, even coin flips outside of L would (or at least

[25] There is a different, familiar context in which it is true that had Jones bet on "heads," then Jones would have won, since the outcome would have been no different from what it actually was: heads (let's say). In that context, a world's sharing the flip's actual outcome is thereby rendered more similar to the actual world.

[26] This counterfactual conditional's truth is logically consistent with Λ^*'s stability since (flip $\Box\!\!\rightarrow$ heads) is logically inconsistent with (flip $\Box\!\!\rightarrow$ 50% chance heads). Hence, Λ^* (including (flip $\Box\!\!\rightarrow$ 50% chance heads)) does not need to be preserved under the supposition that (flip $\Box\!\!\rightarrow$ heads) in order for Λ^* to be stable. One way to show that (flip $\Box\!\!\rightarrow$ heads) is logically inconsistent with (flip $\Box\!\!\rightarrow$ 50% chance heads), even though (heads) and (50% chance heads) are logically consistent, is to point out that (flip $\Box\!\!\rightarrow$ 50% chance heads) logically entails "Had the coin been flipped, it might have landed tails", which contradicts (flip $\Box\!\!\rightarrow$ heads). See (Lange 2006a).

might) not still have had 50% chances of yielding tails. Just as the supposition that (flip in L $\Box\!\!\rightarrow$ heads) posits that such a flip's outcome is preordained, so likewise the supposition that "Were there an A-B interaction at 1-2 nm, all such interactions would produce green slime" (φ) posits that such an interaction's outcome is preordained. Just as "(flip in L $\Box\!\!\rightarrow$ heads) $\Box\!\!\rightarrow$ (coin flips outside L have 50% chance of yielding tails)" is false, so likewise had φ, then the actual laws about A-B interactions at other distances would (or at least might) not still have held.

7. Conclusion

I have argued that for the laws to have a gap for A-B interactions at 1-2 nm, it must be that the laws would all still have been true, had it been the case that had there been A-B interactions at 1-2 nm, they would all have produced green slime. But, I have suggested, it is not the case that the laws would then still have been true. To conflict with Λ*'s stability (where Λ has a gap), it is not necessary for "(A-B interaction at 1-2 nm $\Box\!\!\rightarrow$ green slime) $\Box\!\!\rightarrow$ $\sim\Lambda$" (that is, "φ $\Box\!\!\rightarrow$ $\sim\Lambda$") to be true in every context. It suffices that in some context, "φ $\Diamond\!\!\rightarrow$ $\sim\Lambda$" is true. This seems modest enough.

Admittedly, one might simply dig in one's heels and say: If the gappy Λ does indeed contain all and only the laws, then it must be true that Λ would still have been true, had φ. Although this counterfactual seems false to me, I don't have a knock-down argument against it; it involves no outright contradiction. However, I would ask someone who digs in her heels which of the two key counterfactual conditionals in my argument she believes false (in every context): "φ $\Box\!\!\rightarrow$ green-slime law" or "φ & green-slime law $\Diamond\!\!\rightarrow$ $\sim\Lambda$." (Or both.) Our precise point of disagreement might then be identified.

I have concluded that the laws (if there are any) must be complete. "What colossal presumption," you may say, "for a philosopher reasoning *a priori* to purport to ascertain such a contingent fact about the universe!" You might press the point along these lines:

> Surely it should be left for empirical science to figure out what the laws happen to be like. In particular, we should not prejudge the ways that the laws might constrain the future given the past. The above requirement that the laws be "complete" audaciously presumes that the only constraints that the laws could impose on some hypothetical future event *E* given past events are by entailing *e*, entailing $\sim e$, or entailing *E*'s chance. But there is no *a priori* limit to discovery in science. We might someday discover additional ways for laws to explain sub-nomic facts—ways that make the laws "incomplete" in the sense formulated above, but do not intuitively involve a gap in the laws's coverage. After all, had this paper been written before quantum mechanics was discovered, it

> might have posited a "completeness requirement" that left no room for irreducibly statistical explanations and ontologically primitive chances in fundamental physics. How can we be sure that any purported "completeness requirement" has allowed for *all* of the ways that laws *could* (as a matter of physical possibility, let alone metaphysical possibility) 'cover' an event?"

I agree with this objection. I suggest that the "completeness requirement" differs in different universes. In a universe where none of the fundamental laws ascribes chances, a gap in the laws's coverage involves a violation of determinism. In such a universe, the above argument for completeness goes through; the requisite counterfactuals hold. In contrast, in a universe where there are also statistical laws, the completeness requirement is the one I gave in section 2. If the laws say nothing about the outcomes of A-B interactions at 1-2 nm, not even ascribing chances to them even though such interactions are physically possible, then the laws are incomplete by either standard.

We might even imagine a universe where there are chance processes that have chances rather than non-chancy events as their outcomes. Such a process has various chances of yielding various outcomes consisting exclusively of various different chances of *e*. Until a given process has run its course (at *T'*), there is no N such that ch(*e*) = N; there is (at *T*) only some chance N that when the process yields an outcome (at *T'*), *e*'s chance will then be M: $ch_T(ch_{T'}(e) = M) = N$, an irreducibly second-order chance.[27] In such a universe, it suffices to satisfy the completeness requirement that the laws and various events not involving chances and occurring at or before some moment *T* preceding the time with which *E* is concerned entail that some N is the chance at *T* that at some later moment *T'*, *E*'s chance is M—i.e., $ch_T(ch_{T'}(e) = M) = N$. If the laws say nothing about the outcomes of A-B interactions at 1-2 nm, even about the chance of green slime's having a given chance of resulting, then the laws are incomplete. The laws cannot contain such a gap, on pain of Λ^*'s instability.

While it is metaphysically compulsory that some or another "completeness requirement" hold, no particular "completeness requirement" is metaphysically compulsory. Rather, a particular "completeness requirement" holds in a given possible world as a *meta-law*—a law governing the "first-order laws" Λ (i.e., the laws governing the subnomic facts). The meta-law expressing the completeness requirement specifies the manner in which every event must be "covered" by first-order laws.

27 Elsewhere (Lange 2006b) I have argued for the metaphysical possibility of irreducible second-order chances.

Other meta-laws are widely accepted in physics. For instance, a spacetime symmetry principle in physics is standardly construed as a meta-law—as "a superprinciple which is in a similar relation to the laws of nature as these are to the events" (Wigner 1972:10) and "as laws which the laws of nature have to obey" (Wigner 1985: 700; cf. Feynman 1967: 59). Time-displacement symmetry, for instance, requires roughly that the first-order laws fail to privilege any particular moment in time. I suggest that in any possible world where there are laws, there must be an appropriate completeness principle that holds with the modal force of a meta-law, constraining the laws in the same manner as symmetry principles are thought to do.[28]

Imagine, for example, a possible world where the first-order laws include Newton's second law of motion (F = ma, relating the net force on a body to its mass and acceleration) and various force laws—a world along the lines contemplated in classical mechanics. As a meta-law, time-displacement symmetry would impose restrictions on the kinds of fundamental forces there could be in such a world. For example, it would preclude a fundamental force law demanding that all bodies feel a component force in a given direction that is zero until a certain time T, and then a constant non-zero strength thereafter. In the same way, a completeness meta-law in such a world would restrict the kinds of fundamental forces there could be. For example, it would rule out a fundamental force on a body varying in a given direction as the square-root of the body's speed in that direction. Such a force (if permitted to act in isolation on a body at rest) generates from Newton's second law an equation of motion that is satisfied by more than one trajectory: by the body's sitting still for *any* span of time, and then beginning to move (Hutchison 1993: 320). The body's launching into motion (or remaining at rest) is not covered by the laws; they do not even ascribe chances (given the prior history) to these events. Like a symmetry meta-law, a completeness meta-law explains why the first-order laws have a certain feature.

By allowing different possible law-governed worlds to have different completeness meta-laws, this view neither forecloses the conceptual innovations and empirical discoveries open to future science nor imposes *a priori* limits on the laws's ingenuity in "covering" events—in finding ways to constrain the future given the past. I have argued only that there will be *some* completeness meta-law suited to the actual laws; it is up to empirical science to discover what it is. Admittedly, this view does purport to give an *a priori* argument that a certain kind of

[28] I elaborate that variety of constraint (in terms of "stability" generalized to a higher-level of law) in (Lange forthcoming).

meta-law obtains—a law demanding that the first-order laws be complete. But what good is a philosophical view if it tells us nothing that we didn't already know?

References

Armstrong, David 1983. *What is a law of nature?* Cambridge: Cambridge University Press.

Bennett, Jonathan 1984. "Counterfactuals and temporal direction," *Philosophical Review* 93: 57–91.

—— 2003. *A philosophical guide to conditionals*. Oxford: Clarendon.

Bergmann, P.G. 1980. Comment in "Open discussion following papers by S. Hawking and W.G. Unruh," in Woolf, Harry (ed.), *Some strangeness in the proportion*, Reading, MA: Addison-Wesley, pp. 156–8.

Carroll, John 1994. *Laws of nature*. Cambridge: Cambridge University Press.

Chisholm, Roderick 1946. "The contrary-to-fact conditional," *Mind* 55: 289–307.

Davies, Paul 1995. "Algorithmic compressibility, fundamental and phenomenological laws," in Weinert, Friedel (ed.), *Laws of nature: essays on the philosophical, scientific, and historical dimensions*, Berlin: de Gruyter, pp. 248–67.

Dorato, Mauro 2005. *The software of the universe*. Aldershot: Ashgate.

Earman, John 1986. *A primer on determinism*. Dordrecht: Reidel.

—— 1995. *Bangs, crunches, whimpers, and shrieks*. New York: Oxford University Press.

Feynman, Richard 1967. *The character of physical law*. Cambridge, MA: MIT Press.

Gardner, Martin 1970. "Mathematical games: The fantastic combinations of Conway's new solitaire game 'life'," *Scientific American* 223 (October): 120–3.

Goodman, Nelson 1983. *Fact, fiction, and forecast*. Cambridge, MA: Harvard University Press.

Horwich, Paul 1987. *Asymmetries in time*. Cambridge, MA: MIT Press.

Hutchison, Keith 1993. "Is classical physics really time-reversible and deterministic?", *British Journal for the Philosophy of Science* 44: 307–23.

Huxley, Thomas Henry 1893. *Science and education: T.H. Huxley's collected essays, volume III*. New York: Appleton.

Jackson, Frank 1977. "A causal theory of counterfactuals," *Australasian Journal of Philosophy* 55: 3–21.

Lange, Marc 1999. "Laws, counterfactuals, stability, and degrees of lawhood," *Philosophy of Science* 66 (June 1999), 243–267.
—— 2000. *Natural laws in scientific practice*. New York: Oxford University Press.
—— 2002. "Who's afraid of *ceteris-paribus* laws? Or: how I learned to stop worrying and love them," *Erkenntnis* 57: 407–23.
—— 2005a. "Laws and their stability," *Synthese* 144: 415–32.
—— 2005b. "How can instantaneous velocity fulfill its causal role?" *Philosophical Review* 114: 433–68.
—— 2006a. "How to account for the relation between chancy facts and deterministic laws," *Mind* 115: 917–46.
—— 2006b. "Do chances receive equal treatment under the laws? Or: must chances be probabilities?", *British Journal for the Philosophy of Science* 57: 383–403.
—— forthcoming. "Laws and meta-laws of nature: Conservation laws and symmetries," *Studies in History and Philosophy of Modern Physics*.
Laraudogoitia, Jon Perez 1996. "A beautiful supertask," *Mind* 105: 81–3.
Lewis, David 1986. "A subjectivist's guide to objective chance," *Philosophical papers, volume 2*, New York: Oxford University Press, pp. 84–132.
—— 1994. "Humean supervenience debugged," *Mind* 103: 473–90.
Loewer, Barry 2004. "David Lewis's Humean theory of objective chance," *Philosophy of Science* 71: 1115–25.
Norton, John forthcoming. "Causation as folk science," Philosophers' Imprint Vol. 3, No. 4, http://www.philosophersimprint.org/003004/, to be reprinted in Price, H. and Corry, R. (eds.), *Causation and the constitution of reality*, Oxford: Oxford University Press.
Schaffer, Jonathan 2003. "Principled chances," *British Journal for the Philosophy of Science* 54: 27–41.
van Fraassen, Bas 1989. *Laws and symmetry*. Oxford: Clarendon.
Wigner, Eugene 1972. "Events, laws of nature, and invariance principles," in *Nobel lectures: Physics 1963-1970*, Amsterdam: Elsevier, pp. 6–19.
—— 1985. "Events, laws of nature, and invariance principles," in Zichichi, A. (ed.), *How far are we from the gauge forces – Proceedings of the 21st course of the international school of subnuclear physics, Aug 3-14, 1983, Enrice Sicily*, New York and London: Plenum, pp. 699–708.