Book Symposium: David Albert, *After Physics*

90th Annual Meeting of the Pacific Division of the American Philosophical Association
April 1, 2016

Craig Callender, University of California, San Diego
Jenann Ismael, University of Arizona
Wayne Myrvold, The University of Western Ontario
David Albert, Columbia University

On April 1, 2016, at the Annual Meeting of the Pacific Division of the American Philosophical Association, a book symposium, organized by Alyssa Ney, was held in honor of David Albert's *After Physics* (Harvard University Press, 2015). All participants agreed that it was a valuable and enlightening session. We have decided that it would be useful, for those who weren't present, to make our remarks publicly available. Please bear in mind that what follows are remarks prepared for the session, and that on some points participants may have changed their minds in light of the ensuing discussion.

W.M.

Release Your Inner Humean: Comments on *After Physics*[1]

Craig Callender

David is the author of two books that significantly shaped my intellectual life, *Quantum Mechanics and Experience* and *Time and Chance*. Both broke new ground in philosophy of physics, fundamentally reshaping the discussion of the measurement problem in the first, and laying out a bold understanding of thermodynamics and the temporal asymmetries in the second. Unlike these two books, the present one is a collection of essays and has the character of a welcome sequel. We're all familiar with the sadness one feels when finishing a particularly good book. I was crushed when I finished David Mitchell's *Clock Bones*, for instance, because I knew I would miss joining the strange compelling universe portrayed there. I've just learned that Mitchell wrote a brief sequel and I can't wait to jump back into that world. *After Physics* inspires a similar feeling, where one is delighted to rejoin a wonderful conversation that ones dearly misses. It doesn't disappoint, for it is no less packed with original, fun, deep, and dense argumentation than the first two books. The big difference in the two sequels is that whereas Mitchell's genre-busting fantasy is fiction, David's book purports to be nonfiction. :)

Barry Loewer once said to me that whereas he and David are fundamentalists first, Humeans second, I am Humean first, fundamentalist second. There is something to that. What I'll try to do is sketch two places in David's overall picture where this may matter. Even if one isn't inclined to a spare Humean picture of the world, the discussion may be of interest to anyone because the resulting pictures are at odds with David's and (I think) independent of their genesis in Humean motivations.[2]

Reduction, Special Sciences, and Statistical Mechanical Imperialism

While philosophers of science have been busy developing and bouncing from one form of non-reductionism after another, David proposes the most audacious form of reductionism seen since Hempel and Oppenheim. His motivation is the special science of thermodynamics. Assume a kind of physicalism or supervenience of the macro upon the micro (what David calls "translation"). If David is right, the projectible macroscopic generalizations of the science of heat reduce to microphysics via the following package: a past hypothesis (that the entropy of initial

---

[1] Please do not quote. These are quickly typed up comments for an Author-Meets-Critics Session at the Pac APA, San Francisco, 2016. The Author is David Albert and the book is *After Physics*. I now stand by about two-thirds of the comments.

[2] I feel bad focusing on this theme because it means I won't discuss some of the best parts of the book. Let me just register here my opinion that the chapter on the "technique of significables" is almost unspeakably cool and interesting. Using a technique involving the locations of golf balls, David plugs holes in proofs of no-superluminal signaling in quantum mechanics. And his work on a failure of narratability could have fundamental repercussions in our understanding of quantum mechanics and is the subject of ongoing research.

macrostate M(0) is extremely low), a uniform probability distribution over the microstates that realize M(0), the present macrostate M(t), an the dynamical laws of the microlevel. Take the special science generalization that entropy increases. We “derive” this from the lower level of physics by making a case that

> (1) P(S increases/M(t) & M(0) & laws & uniform probability over M(0)) = high

Boltzmann, Gibbs and many others make a case for this in statistical mechanics, although it’s worth bearing in mind that they do so rigorously only in ideal cases. Still, it strikes many, including myself, as physically plausible. A lot could be said on this topic, but that will take us away from my main point. If correct, we have an honest-to-goodness reduction of a special science law — the second law of thermodynamics — from the bottom. Grant this.

What does this imply about *other* special science generalizations, or really, any kind of stable counterfactual supporting generalization about the macrostate of the world? Does it somehow follow that bunnies reproduce at the rates ecology says, that 80% of people flashed smiley faces themselves smile, that Smith stably prefers blue ties to red ties, etc?

One thing we might consider is *copying* this strategy in ecology, psychology and so on. Sheldon Smith and John Earman consider this when discussing ceteris paribus laws. Absent a Boltzmann or Gibbs relating bunny reproduction to quarks, a basic state space, measure, etc, we’re really in cross-your-fingers-and-hope territory. But it could be that that the other special sciences get the same kind of treatment as thermodynamics, except that here the intervening science is too hard to actually produce.[3] David seems to want something much stronger, something more unified, and something where the other special sciences *follow* from the PH.

I have trouble squaring this aspiration with has always seemed to me a plausible observation of Schrodinger’s. Every now and then (recently in the book, *Into the Cool)* scientists propose that non equilibrium thermodynamics explains pretty much everything, e.g., life, economics. In *What is Life?* Schrödinger tiptoes to the edge of saying that NET explains life, but backs away:

> But F. Simon has very pertinently pointed out to me that my simple thermodynamical considerations cannot account for our having to feed on matter 'in the extremely well ordered state of more or less complicated organic compounds' rather than on charcoal or diamond pulp. He is right.

The point is that NET may be necessary for life and economic markets, but it doesn't seem sufficient. Because we're thermodynamic organisms and far from equilibrium, physics says that we need low entropy intake. Simon points out that thermodynamics doesn't provide further discrimination about what the form of low entropy intake must be, yet there are generalizations

---

[3]We could claim that, for any generalization G over macrostates we like, that

> (2) P(G/M(t) & Mg(0) & uniform probability over Mg(0)) = high

where Mg is the initial state such that uniform probability over it makes G likely. Repeat for G1, G2, G3… Logically, that works, but of course we sacrifice any support we ever had from Boltzmann, Gibbs, etc. And then we don’t have a unified reducing package. We just have a bunch of hypothetical separate reductions.

about this in biology. In accepting Simon's point, Schrödinger is seeing that some biological facts don't follow from thermodynamics.

Put in the above framework, the claim then is that if we replace “S increases” with claims in a different vocabulary, in a different science, such as the claim about reacting to smiley faces, then

(3) P(S smiles../M(t) & M(0) & uniform probability over M(0)) =/= high

Looking only through a thermodynamic lens at the past state doesn’t tighten the likely paths to psychological ones. Perhaps the point can best be put visually. Consider a state space parametrized by the fundamental variables and an actual history through that space. Each special science can then be viewed as offering a way of partitioning this space into sets. Thermodynamics does it one way, and in the figure you can see that the actual history progresses into bigger and bigger sets, corresponding to larger and larger entropy. Other sciences will carve up this set differently. They are all different windows onto the same world. And it’s perfectly compatible with this picture that ecology and psychology and the rest carve it up differently. Moreover, the trajectory compatible with entropy increase is compatible with inequivalent and more fine-grained carvings that say a lot more than that entropy increases. To be specific, consider entropic region R. Suppose psychology carves up R into two pieces, R1 and R2. It might be a psychological law that the state typically heads towards R1, not R2. That is compatible with entropy increase, but it is conveying more info than thermodynamics does.

David could make evolution into R1 likely by conditionaling upon an even more special past state, the state M’(0) that makes (2) true for all G’s one cares about. Then he is basically positing the special initial condition that gets it all right, not the initial condition picked out by thermodynamic parameters.

Is there a problem with this scenario as I’ve sketched it? From the perspective of statistical mechanics, and David, I think, it’s a *coincidence* that the state always ends up in R1 and not R2. Suppose, by the carving done by thermodynamics, R1 is the same size as R2, and that there is nowhere else to go. Then states should go into R1 50% of the time. Instead they go in 80% of the time,or so says (eg) psychology. As an analogy, consider a sequence of fair coin flips generated by a truly stochastic device: HTHHHTHTTTHTTHHH. Now invent some macropredicate, say, number of T’s in a chunk of 12 tosses. It’s compatible with the underlying randomness that one finds regularity among these macropredicates, say, for a few billion years. Special science laws could be like this, “accidents” frozen into the world.

David’s objection, I suspect, is that this makes ecology just a matter of dumb luck, a series of coincidences bound to run out. This point about coincidences is the lynchpin here, I think, for it’s what makes crossing the Schrodinger divide between a necessary condition to a sufficient condition tempting.

And here is where Humeanism might enter. There is no divine coin tosser, nor is there a level where causes *really* cause, laws *really* govern, etc. George Bush said he was the Decider. On David’s view, the basic fundamental chances are the Deciders of whether any pattern is a coincidence. By contrast, on a less reductionistic picture, physics, like ecology, is just another window onto the world — and all the modal features arise in our making sense of patterns, wherever they may be found. On this picture, coincidences are the norm, seen whenever one translates between different systems. Heading into R1 may be highly likely according to the carving and measure provided by ecology, but only 50% likely according to statistical mechanics. Equally, thermodynamics may see differences that ecology doesn't, and then again we'd have coincidences from the other perspective.

With this kind of picture in mind, Jonathan Cohen and I elsewhere developed a quite different picture than David's. We agree on a lot: the assumption of physicalism or translation, a Humean theory of laws, and the rest. But we try to give voice to the idea that the special sciences are in some sense autonomous. Instead of *one* best system competition, we imagine multiple ones, one for each vocabulary of interest and science. Scientists figure out notions of simplicity, strength and balance in field specific ways, using metrics designed by and implemented in their fields. They needn't look over their shoulder at physics. In terms of my earlier figure, there is a best system for each partition of the world's basic state space. Systems provide laws, chances and counterfactuals. Multiple systems give you differences amongst these. So the theory is ontologically reductionistic — bunnies are really quarks, etc, behaving in fuzzy hopping ways — but it is anti-reductionistic about projectibility.

My first question, then, is to ask David to react to Schrodinger's point, and my second, is to ask him to react to this multiple-systems perspective.

Lost in Space

"Descartes held that the mind and the body inhabit different "spaces." The mind lives in an extensionless mental realm whereas the body lives in the extended realm of objects. His contemporaries, Pierre Gassendi and Princess Elizabeth of Bohemia, famously asked how these two objects communicated, that is, how causal influences surmounted the lack of a common space. Wouldn't ripples triggered by an event in one realm fail to reach the shore of the other? How does a thought in the mental realm trigger changes in the pineal gland (Descartes' preferred location for the mind–body nexus)? Philosophers know this problem as the "interaction problem" for Cartesian dualism. While it is one of the more famous problems in the philosophical canon, it is far from decisive against Cartesianism. Better objections are probably methodological in nature, arising from the lack of the mental realm's predictive power and its assault on parsimony. Still, different spaces bring awkward questions. How does the interaction work? Where are we? Is motion conserved in each space? And more.

A similar set of problems afflict a naive reading of [what Einstein called] "pea" theories. Because it is the most developed "pea" theory, we'll focus on the deBroglie–Bohm interpretation of quantum mechanics; however, I expect that most of what I say will be true of others, too. deBroglie–Bohm (Bohm 1952) is an attempt to solve the notorious measurement problem. According to this class of theories, the quantum wavefunction evolves unitarily according to a linear wave equation which guides the motion of Einstein's peas (e.g., particles, classical fields, fermion number density). In the familiar non-relativistic case, the ontology appears to be a wave evolving according to the Schrödinger equation and a particle configuration evolving with a velocity given by the conserved current divided by the probability density. No less an authority than J.S. Bell tells us that " ...no one can understand this [Bohm] theory until he is willing to think of ψ as a real objective field rather than just a " probability amplitude" (2004, p. 128). Hence the apparent ontological dualism: existence of both ψ and particles.

The quantum state is a ray in a high-dimensional Hilbert space, H = L2 (R3N ,d3N q), where N is the number of particles in the universe. Since the Bohmian privileges position, it is more natural to work with the field the wavefunction defines (see below) on a configuration space R3N than Hilbert space. The peas, by contrast, live in a low 3-dimensional space, R3. Unless N = 1, the field on configuration space and the primitive ontology then live in different spaces, e.g., R3N/= R3 (Fig. 1).
Different spaces again invite awkward questions. The field in R3N "guides" an N-particle system in R3. How does this puppet-master operate the puppet with- out any strings? 1 And where do we live, in R3 or R3N? Is motion conserved in the two spaces? These problems are not threats to the logical coherence or empir- ical adequacy of the theory, and some may even be said to be badly motivated (Callender and Weingard 1997); still, if not blemishes on the theory they raise the hope that something better might be possible. Call this the problem of being lost in space.

The philosophical literature has directed attention to the "two-space" problem (see Ney and Albert 2013). Some respond by trying to stuff the waves and particles into the same space. Two directions are possible (Fig. 2). Albert (1996) champions boosting the particle configuration upstairs into R3N . Here the beable must be a single Bohmian "world" particle. People, planets, particles and all the rest somehow emerge from this world particle. Alternatively, one can kick the field on configuration space downstairs into R3, providing each particle with its own wave to surf. Doing so is highly non- trivial, yet recently Norsen (2010) has gone some distance toward showing how this might be accomplished. Each view has its merits and demerits. However these stack up, it's fair to say that either is sufficiently radical that having a third option seems worthwhile.

Here we can be guided by the mind–body debate. There one eliminates the inter- action problem not by retaining dualism and putting the different kinds of entity in the same space, as Albert and Norsen do, but by eliminating one half of the ontology. Berkeley famously eliminates the corporeal and retains only the mental entities in the mental realm. Materialists and physicalists instead eliminate the mental entities and keep only the material. Here we have two options as well. One is to retain only the wavefunction. Doing so commits one to an Everettian interpretation and all of its attendant challenges. The other is to simply posit the non-wavefunction beables, Einstein's peas, and nothing else.

The best response to the above set of concerns is to eliminate them altogether by denying that the quantum state's status is ontological and asserting that it is instead nomological (Dürr et al. 1997; Goldstein and Zanghì 2013). The idea is that the quantum state does not represent an entity in the world, a piece of ontology, but rather is part of the representational structure of the laws of physics. The quantum state is therefore akin to the Hamiltonian function. Like the quantum state, the classical Hamiltonian generates the motion of the beables and also their statistics. It doesn't "live" in ordinary three-dimensional space, for it is defined on an even higher dimensional space than our field on configuration space, namely, phase space, R3N xR3N . Yet no one worries about where the Hamiltonian lives because it is viewed as part of the laws, not the physical world itself. There is simply no expectation that laws be functions over three-space, nor that they be decomposable into functions on three-space.

This conceptual reorientation results in a very satisfying picture. In non-relativistic theory, the world consists simply of a bunch of Bohmian particles or GRW flashes or other forms of Einstein's peas—and that's it. In the Bohmian case, the Schrödinger and guidance equations are then the laws, existing wherever laws exist. We have two spaces, unlike in Albert and Norsen, but no dualistic ontology and therefore no interaction problem. This resolution eliminates various complaints against beable theories, as we'll see below, and it has a number of other attractive features, parsimony high among them.

David thinks this position is “crazy” (126). Why? Here he makes the same kind of points made elsewhere in the literature (Brown and Wallace, Belot, and others). The wavefunction evolves, and therefore is time-dependent, unlike the laws of nature. It also varies by system. Together, these two features make the wavefunction “contingent-looking” (126). Albert adds in a footnote: “It is the very essence of stuff to be, in general, a mess—and it is the very essence of laws (on the other hand) to be clean and simple.”

There are reasons to think that such intuition-mongering shouldn't be decisive. Neither the notion of a law of nature nor our understanding of the quantum state are solid enough to ground objections based on analogies and disanalogies. On the law of nature side, note that some philosophers and physicists believe that “evolving” laws of nature make sense (e.g., Smolin 2013). For them an inference from a quantum state's time-dependence to its not being a law is not sound. Stepping back, one may also point out that our intuitions about what is nomological are presumably formed by comparison with classical physics. Yet one can reasonably ask why we should take lessons about ontology from a theory that we know is wrong. Perhaps quantum theory is telling us about the nature of the nomological.

However the burden falls, note that for beable theories, one needs to approach these disanalogies by first distinguishing universal and effective wavefunctions. The universal wavefunction is the wavefunction of the entire universe. The wavefunctions associated with systems in laboratories, people, planets and other subsystems, by contrast, are effective wavefunctions. Let Q_t be the actual configuration of particles in the universe at a time. In a composite system $Q_t = (X_t, Y_t)$, where X is the actual subsystem of interest and Y is the actual environment. Then a natural definition for a subsystem's wavefunction is the conditional wavefunction $\psi_t(x) = t(x, Y_t)$, where we calculate the universal wavefunction in the actual configuration of the envi- ronment (Dürr et al. 1992). The conditional wavefunction will not in general evolve according to the Schrödinger equation, but when it does—which it will if the univer- sal wavefunction evolves into a wide separation of components in the configuration space of the entire system—then we call the conditional wavefunction an effective wavefunction. Effective wavefunctions correspond to the wavefunctions discussed in quantum textbooks, labs, and so on.

When this distinction is appreciated, we see that the above disanalogies must be approached with caution. Just as offspring can differ from their parents, so too can effective wavefunctions differ from the universal wavefunction. You may be tall and your parents short. Similarly, the universal wavefunction may be simple and the effective one for some subsystem complicated. One may be real and the other complex. One may not vary by system (universal) whereas the other (effective) may vary with system. And perhaps most surprisingly, the universal wavefunction may be time independent and an effective one time-dependent. Many Bohmians have explicitly shown how this might arise in quantum gravity (e.g., Callender and Weingard 1994; Goldstein and Teufel 2001), but the possibility exists even in ordinary non-relativistic theory (Esfeld et al. 2013). Given this distinction, it's not clear that we should infer from the effective wavefunction's variability that the universal wavefunction is variable in any way. The universal wavefunction is clearly not variable by subsystem, since it applies to everything, and it's an open question whether the universe as a whole could have different universal wavefunctions. Moreover, the universal wavefunction might not even be time-dependent. The true form of the universal wavefunction is simply a matter of speculation, and therefore, so are the claimed analogies and disanalogies.

Two more points…

First, an *ad hominen*. Intuitions about what's makes for a simple law turn out to be decisive for David here, but they don't earlier in the book. There the Past Hypothesis is regarded as a law of nature, and certainly, expressed in the macro language it sounds simple ("entropy is low"). But that's not fair, as it's purportedly a law in the best system of basic physics. Hence it ought to be translated into these variables, and when it is, I bet it's a mess.

Second, surely this position wins against David's on any cost-benefit analysis. Look at how attractive it is! There is just good old-fashioned three-space with a bunch of particles on it. The very same ontology, if you want (non-relativistically, maybe relativistically), as classical mechanics. Immunity to the Redundancy and Action-Reaction objections to Bohm. A clear and gorgeous spare ontology worthy of Quine. By contrast, David posits a kind of intellectual monstrosity, a Bohmian world particle living in a massively high-dimensional space. After the worry of lack of supervenience is discharged, as Elizabeth Miller and I do in our papers, the only thing standing in the way of adopting this beautiful world picture over the ungainly one is a bunch of intuitions we have, based on a false theory and probably not applying to the universal wavefunction, about laws and messes.

In sum, David, listen more closely to your inner Humean! I sympathize: on the East Coast you're surrounded by metaphysicians and it's natural to keep the inner Humean muzzled slightly. But you're on the West Coast today, a place where you're allowed to express yourself freely. Later, while doing yoga, try to pay attention to the inner Hume. When freed from distractions, you'll hear that no scientific level is the Decider, and hence no theory's chance functions uniquely determines coincidences. Not only is this plausible, but it will make you more popular—scientists everywhere will be freed from the shackles of your physics. And heeding that inner Humean voice will also allow you to see that wavefunctions are simply aspects of laws of nature, summaries of the real stuff in the real and beautiful low dimensional world.

Thanks for writing my favorite type of book, one that is simultaneously challenging, educational and fun.

Comments for APA Symposium on *After Physics*

Jenann Ismael
jtismael@u.arizona.edu

Introduction

I would like to David Albert for a beautiful book. I'm both happy and honored to be able to comment on it. The book is full of insights that are characteristically gem-like in their clarity and beauty and permanence.

Although not written systematically, the pieces fit together into a distinctive program that gives a stratigraphic view of the architecture of the physical world. Starting with the familiar world of everyday sense, it peels back the layers revealing the physical structures that underlie the macroscopic surface of the world. Lower levels generate structures that scaffold higher levels, and as we descend to lower levels the structures become increasingly austere and increasingly unfamiliar. The temporal asymmetries arise at an intermediate level, scaffolded by the classical microscopic laws and low entropy past. Below the classical level, quantum mechanics gives us glimpses of even deeper, non-spatiotemporal structures.

The book is enormously ambitious. It comes with a broad epistemology rooted in the neo-Boltzmannian account of the foundations of statistical mechanics presented in *Time and Chance* that sheds light on central features of human experience. It takes its departure from well-worn topics in philosophy of physics (e.g., accounting for thermodynamic irreversibility), but strides into areas of philosophy well outside the traditional purview of foundations of physics[1]: questions having to do with human being, her place of human action in nature, and the status of agency. This brings philosophy of physics into contact with areas of philosophy that it doesn't normally speak to.

The Neo-Boltzmannian account of the origin of the macroscopic asymmetries is familiar from *Time and Chance*.[2] It has been scrutinized and well picked over in the 16 years since its publication. *After Physics* provides a new, and extended treatment of one of the most philosophically significant and controversial aspects of that story: the account of the asymmetry of influence or causation. As Albert says:

"What seems to have struck so many people as impossible to swallow in the account of the time asymmetry of counterfactual dependence that was presented in *Time and Chance* is that – on that account – that asymmetry turns out not to be absolute. On that account, the difference between our capacity to influence the future and our capacity to influence the past is apparently a matter of degree"[3]

The asymmetry is the subject of "The Difference between the Past and the Future", written in the form of a

[1] Against the advice of Callender. See his contributions to the symposium.
[2] Harvard University Press, 2000.
[3] P. 46.

dialogue between "Huckleberry" (a stand-in for Albert) and "Jedediah," (who acts forcefully and quite effectively as the voice-peace for objections). The epistemic and practical asymmetries both have their origin in the thermodynamic arrow. And they are asymmetries that apply to the human being, which we can understand for these purposes as a physical system that is coupled to the world in such a way that it picks up macroscopic information about its surroundings and intervenes locally. The epistemic asymmetry is an asymmetry in the kind of information about the past and future that is available to a system like that in the here and now. The practical asymmetry is an asymmetry in the ability of a system like that to influence the past and future by actions undertaken in the present.

The Neo-Boltzmannian account of the thermodynamic asymmetry offered in *Time and Chance* has three components:

1. The Newtonian laws of Motion
2. A statistical postulate (SP), which gives the canonical statistical mechanical probability distribution that a system in a specified macrostate M, is in a corresponding microstate m_n.
3. The Past Hypothesis (PH), which says that the universe started in a state of low enough entropy to make thermodynamic generalizations applicable for the roughly 15 billion years we think these generalizations have held.

Albert's new name for the probability distribution that follows from SP, HP, and the classical dynamical laws in the Mentaculus.[4] Here's how the practical asymmetry gets recovered from the Mentaculus. The agent starts with what he calls the 'fiction of agency', assuming unmediated control over certain features of the world (her 'actions'). Influence is understood in terms of counterfactual dependence on those features. Albert's claim is that on almost any reasonable criterion for assessing truth-values of counterfactuals, it comes out that the agent's actions make a difference to the future, but not the past. And the reason traces to the same macroscopic asymmetries that underwrite her differing knowledge of past and future. The recipe that he himself proposes is to look in phase space for the possible world in which the past hypothesis is true, whose macrohistory is not abnormal, and which is closest to the actual one at the time of the antecedent. Evolve it backward or forward (as the case may be) in accord with the equations of motion, and see whether it satisfies the consequent.[5]

Following the recipe gives us some good news and some bad news. The good news is that if we look at the effects of local macroscopic changes in the world of the kind that we normally associate with human actions we don't find particular identifiable counterfactual dependencies in the macroscopic past that provide strategic routes for bringing about past events. The bad news is that we do find some esoteric counterfactual dependencies of the past on our present actions. It follows that on Albert's account of the practical

[4] The name is taken from taken from the Coen brothers movie "A Serious Man", in which it is used to refer to what is supposed to be a probability map of the universe.
[5] If the consequent precedes the antecedent, we evolve the state backward in time. If the consequence comes after the antecedent, we evolve the state forward.

asymmetry, our inability to influence the past turns out to be only an approximate, as-far-as-we-can-tell, for-practical-purposes inability rather than a true physical impossibility. Examples of these kinds of dependencies were marshaled as objections to the account of the causal asymmetry offered in *Time and Chance*.[6] Jedidiah raises two of these (one due to Elga, one due to Frisch) and Huckleberry responds with explanations of why the dependencies are not practically utilizable as ways of influencing the past. Albert refers to these dependencies as 'causal handles'. He shows again why it follows from the Mentaculus that *if* the set of features of the world over which we think of ourselves as having immediate control is small and localized in space (e.g., the movements of a certain body), *then* (i) causal handles on the future are vastly more plentiful than causal handles on the past, and (ii) it will be out of our ken to discover what causal handles there are or put them to practical use.

Up to this point, Albert is merely rehearsing the account of the asymmetry that was offered in *Time and Chance*. He then adds a new observation: a strict epistemic difference between measurement and retrodiction that is intended to explain the sharpness of the intuitive, or pre-theoretic, distinction between past and future. The new observation is intended to explain why we tend to think of the future as wide open and the past as not just *approximately-and-for-practical-purposes*, but *strictly* fixed. The reason is illustrated with billiard balls: Albert's all purpose stand-in for a classical system. He suggests that the position of a particular billiard ball at some particular future time presents itself to us as open or as unfixed or as susceptible of being influenced or as amenable to our control over the interval between now and then because we can have no empirical access to the condition of that ball at that time except by way of the story of what's going to get done to it over that interval. In the case of the past position, we can eliminate uncertainty now about where it was at some earlier time without knowing about what gets done to it between now and then by finding a record of its position: a photograph, a witness, a scrapbook in which someone wrote it down. So, Albert says,

"The fact that our experience of the world offers us such vivid and plentiful examples of this *epistemic* independence (of the past on the present) very naturally brings with it the feeling of a *causal* and *counterfactual* independence as well."[7]

We can marshal all kinds of concrete, local evidence to confirm or deny beliefs about earlier events, without knowing anything of what has happened in the interim. Nothing of the kind can be done to confirm or deny beliefs about the future. There are no localized records, no concrete traces, no documents or witnesses that could bear testimony of what is to come. Nothing short of full knowledge of the events along a cross section of the past light cone of even the smallest and most mundane object could tell us where it will be, and what it will be doing at some future time.

Discussion

---

[6]What I'm calling the 'practical asymmetry', Albert calls the 'causal asymmetry'. I prefer the former because it brings the relational character of the asymmetry (and in particular, its relationship to human action) into more explicit focus.

[7] p. 57

This part of the stratigraphic vision that comes through in *After Physics* is not etched as finely as the other parts, including the discussion of the epistemic asymmetry. Albert is non-committal about what features of the world an agent views herself as having control over (mental acts like decision or volition, brain states, bodily movements) and he is non-committal about the correct semantics for counterfactuals. The neutrality is meant to avoid taking a stand on contested details, but it has the effect of leaving it somewhat out of focus. It is also not as illuminating as the account of the epistemic asymmetry. One may dispute the details, but one feels that one understands how knowledge gathering is both possible and successful in a world with a thermodynamic gradient. In the case of the practical asymmetry, it is never clear whether Albert thinks that his account under*writes* or under*mines* the belief that we influence the future. His references to the '*fiction* of agency' and the language about epistemic dependence bringing with it the '*feeling* of causal and counterfactual dependence' leave one with an uneasy feeling that agency – on his account – turns out to be an illusion, and indeed an illusion of a very traditional kind. It seems that Albert is suggesting that information about what is going to happen in the future is *there in the world*, but is simply distributed so widely that we don't have access to it. So, the 'feeling' of counterfactual dependence turns out to be a product of ignorance. Although he is following a distinguished tradition in explicating the practical asymmetry in counterfactual terms, moreover, the fit between the *content* of counterfactuals and their *functional role* seems wrong. The ordinary semantics for counterfactuals makes beliefs about counterfactuals beliefs about possible, non-actual worlds. But it is unclear why beliefs about *other possible worlds* should play any role in guiding one's decisions about what to do in the actual world, here and now.[8] So there are reasons that this part of the account has encountered a lot of pushback. It's not only the intrinsic interest of the question. It is also that it is less persuasive.

I think Albert has the material to do a little better. First, I see no good reason not to be concrete and specific about the features of the macroscopic environment that the agent takes to be under her immediate control. They are the bodily movements that she learns early on in development that she can control by will. There is no mystery here. In psychological terms, perceptual feedback on inner efforts of will that teaches an infant which changes in the environment are under voluntary control. She kicks her legs and waves her arms, eventually getting them under control and spends the rest of her life learning to use her movements to effect changes in the wider world. That chases the question of whether agency is 'real' or fictional back to the question of whether our volitions are under our control, but that is where those questions belong, and I think Albert should be happy to leave them in the hands of the moral psychologists.

[8] And for that matter, why *these* worlds, rather than *those*. The ordinary semantics for counterfactuals might do a good job of systematizing our ordinary judgments about their truth, but it does little to illuminate their role in practical reasoning. As Albert pointed out (in discussion), the Lewisian semantics for counterfactuals, combined with his best systems analysis of laws, yields the conclusion that counterfactuals turn out to be about the actual world in a more roundabout way. I'm not sure whether the connection is too roundabout and buried too deeply in Lewis's account to capture the content of counterfactuals for everyday users. In any case, the connection could be displayed more perspicuously, without the detour through possible worlds and a reductive analysis of modality. I'm sure that there is more to be said here, and the issue could use some clarification.

Second, I that think that maybe counterfactuals, with the attendant talk of possible worlds and metrics of possibility, are unhelpful. I want to suggest a way of doing things that brings the causal asymmetry more into line with the epistemic asymmetry, and lets Albert appeal to the Mentaculus in a straightforward way. The idea is to exploit Pearl's insight that causal thinking is to practical reasoning as probabilistic thinking is to epistemic reasoning. Instead of talking about counterfactuals (and being pushed into talk global, non-actual possibilities), we can think about hypothetical conditionals of the form 'What would happen in the actual world if I A'd rather than B'd? or B'd rather than C'd?', where A'ing, B'ing, or C'ing are things that the agent thinks of as coming under her voluntary control. These beliefs are clearly beliefs about our world, and their role in practical reasoning is easy to discern.[9] Then we appeal to the interventionist account of causal claims. Causal claims, on that account, are claims about the results of hypothetical interventions. Interventions are 'surgical' changes in the value of a variable. If A is a variable in a probability space S, an intervention on A is a change in the value of A that comes from outside S and leaves the connections between A and other variables in S intact. In the paradigmatic setting, an agent deliberating about what to do is thinking about what changes to effect in the macroscopic environment. She treats her actions as interventions, and her causal beliefs will tell her how differences in her actions will ripple through the space. If she were representing the world in full physical detail, her actions would appear as bodily movements. But we don't typically represent the world in full physical detail. More typically, we choose a set of variables of interest, leave background conditions tacit, and represent our actions, by whatever elements in the space we take to be most immediately under voluntary control. If we are deciding between airplane seats, marriage proposals, or dinner entrees, for example, our actions will be sitting here rather than there, marrying Jim rather than Joe, ordering this rather than that. The framework itself is generic. Causal beliefs are generically beliefs about how interventions on the value of one variable in a chosen space (against sometimes only tacitly specified background conditions)[10] ripple through the space.[11]

Causal beliefs matter in practical reasoning because they carry information about the effects of hypothetical actions, and an agent *uses those beliefs* to guide her decisions about what to do.[12] Let's see what the Mentaculus tells an agent about the likely effect of some local macroscopic change to the value of a variable **A** in the environment. She starts with information about the present, surveyable macroscopic state,

[9]Counterfactuals and future conditionals are both species of hypothetical. Science deals generically with hypotheticals, and although hypotheticals give us the logical resource to define counterfactuals, it is the future conditionals that have the most basic cognitive function. This is what Alison Gopnik meant when she said: "past counterfactuals are the price we pay for future conditionals." (*The Philosophical Baby: What Children's Minds Tell Us About Truth, Love, and the Meaning of Life* Farrar, Straus and Giroux).

[10] Or, if you prefer, how interventions on the value of one variable in a network affect other variables in the network.

[11] The probability space can be widened or narrowed and background conditions can be added or dropped. The setting in which the probability space is widened to include the agent's own decisions is the one in which the question of whether she is justified in treating her own decisions as under her control arises.

[12] Causal beliefs are beliefs about the results of hypothetical interventions, where an intervention is an action from outside the probability space that sets the value of a variable. Intuitively, what causal beliefs add to probabilistic beliefs is a separation of correlations due to information that A carries in virtue of common causes in the past and correlations due to A's own influence. It is because it makes this distinction that causal information goes beyond mere probability, and it is because it makes this distinction that causal beliefs identify strategic routes to bringing about ends.

conditionalizes the Mentaculus on that. This leaves her with many specific beliefs about the past, and few such secure beliefs about the future. Now, she conditionalizes on the change in A, and what she sees is the probabilistic effect of that change rippling into the local future along routes that she can identify, leaving the past mostly untouched. If she jumps into a swimming pool, for example, that makes it vastly more likely that she will be wet, that the water will be disturbed, that she won't be eating a hot dog, and so on. What all of this means is that the causal landscape relevant to human action will have influence running into the future in more or less the way that common sense supposes.

This doesn't change anything in Albert's account. It simply brings the asymmetry of influence a little more transparently into alignment with the epistemic asymmetry. By understanding the position of the agent, we understand why she raises the question that she does. By understanding the physics, we understand why she gets the answers that she does. From a thermodynamic point of view, when an agent acts as she does in the here and now, we can think of her as putting in place future records of her present action. The present macroscopic effects of her A'ing are the future macroscopic records of her having A'd. So, although there is no intrinsic direction of influence built into the fundamental fabric of the world, if we look just for those lines of influence that the human being who intervenes locally in space, can harness as strategic routes for effecting parts of the landscape at which she is not located, those lines of influence will run from past to future. Our ordinary, everyday causal ideas, then, have a built-in direction because they carry information about those lines of influence. We don't have to start with any 'fiction of agency', and we don't have to take the detour through counterfactuals and possible worlds. Causal beliefs are beliefs about the results of hypothetical interventions. When those interventions are local macroscopic changes in the environment of the kind that we take to be under voluntary control, the causal asymmetry comes out very naturally of the thermodynamic gradient as the other side of the epistemic asymmetry.

I mentioned that people have balked at the conclusion that the causal asymmetry is not strict. On Albert's account, it is not true that nothing about the present state of the world can influence the past. It is only that *we* cannot influence the past by the means at our disposal because of limitations on what we have direct and unmediated control over. The practical impossibility of influencing the past strikes many as far too weak. The interventionist account of causal relations has resources that can be marshaled to generate a stricter prohibition on influencing the past. It is designed to distinguish correlations due to causes from correlations due to merely evidential relationships, and in the two cases that Jedidiah raises as difficulties for Albert's analysis, the problematic dependencies of past on present turn out to be evidential, rather than causal. In the example due to Frisch, a practiced pianist playing a familiar piece is running through passage that is repeated twice, but followed by two different endings. He doesn't remember whether he's played the first ending, but he trusts his instincts and decides to play the second ending if it feels right. The decision to play the second ending is correlated with having already played the first ending because it provides evidence of having played the first ending, but one wants to say that the second ending doesn't cause or bring it about that the first ending was played. Deciding to play the second ending doesn't serve as a strategic route for bringing it about that the first ending had already been played.

On the interventionist account, to establish a *causal* connection on the interventionist account, the probabilistic correlation should persist when we intervene on the decision. Ordinarily we cut ties with *pre*-decision events, intervene on the decision to play the second ending, and see how the decision to play the second ending influences *later* events. But here we are looking at whether the decision to play the second ending causes *earlier* events, so we take the temporal reverse of the ordinary case, by cutting ties between the *post*-decision events and intervening on the decision. So we hold fixed post-decision events and twiddle the decision, and see whether the correlation between the decision to play the second ending and the prior playing of the first ending persists. It does not. There are lots of post-decision records of the first playing that remain in place when we twiddle the decision. Since these carry the information that the first ending had already been played, the correlation between the decision to play the second ending and the prior occurrence of the first ending is broken, and this fails to count as a causal link. The other case Jedidiah considers is drawn from Elga, and it fails to pass the test for a causal link for a similar reason. In Elga's case, there is a weak correlation between the existence of Atlantis and my not (for example) snapping my fingers right now because all of the familiar kinds of concrete traces that we ordinarily count as evidence for the existence of Altantis has been dissipated into the landscape and disappeared from view. And so it seems that by snapping my fingers now, I can influence whether Atlantis existed. I think that the thing to say here is that so long as the Mentaculus entails that the correlation *would* be broken by the appearance of any post-finger-snapping evidence of the existence of Atlantis, it does not support the existence of a backward causal link, on an interventionist account. The fact that no such evidence may appear only tells us that the evidence underdetermines whether the link is causal or evidential (a familiar situation to be in, in causal modeling). Only a correlation between the snapping of my fingers and the existence of Atlantis that survives the potential discovery of later evidence for the existence of Atlantis qualifies as a causal link on an interventionist analysis, and only such a correlation would provide a strategic route for bringing it about that Atlantis existed.

A question for Albert about imaginative difficulty

All of this, however, might be beside the point if what the examples of influencing the past were really expressing is a feeling that there is something about the whole vision that is wrong, something bordering on unintelligible, something that has less to do with the fact that the asymmetry isn't *strict*, with the fact that it is not *fundamental*. Albert's account is an account of how the world *is* that recovers how things *look* through the eyes of an agent who is subject to certain epistemic and practical limitations. It presents a model of the way the world is (or might be) in which we recover notions of possibility and openness that are part of the surface structure of a human life by seeing how the world appears through the eyes of an information gathering and utilizing creature characterized by certain epistemic and practical limitations. But surely (one might think) we ought to be able to imagine what the world would be like if those limitations were removed. Surely (that is to say) we ought to be able to answer questions like: What does the world look like to microscopic agents who can intervene microscopically and see the effects of their interventions? What would *we* see if our eyes were suddenly enhanced so that we could see right down to the fundamental fabric of the world? What if technology got good enough that we could harness strategic routes of influence into the past?[13] The imagination balks a little at these kinds of questions. One feels somewhat at a loss to deliver answers. They arise for Albert because his program entails that the temporal asymmetries that allow for the *gathering* and *utilization* of information (and hence make cognition and agency possible) are emergent.[14]

I think that it is a feature not a bug, but it is difficult to think clearly about a world in which the asymmetries that seem so essential to experience and agency disappear on the fundamental level. So I wanted to ask what Albert thinks here: does he think that without the well-defined temporal asymmetries that make it possible to *collect and store information* in one temporal direction and *use* it to guide how things go along the other, it doesn't make sense to talk about knowledge or influence, and so there simply is no answer to questions about what things would *look* like? Or does he think there is more to say here?

A Metaphilosophical Question

All of this raises a general question about how restrained (with respect to common sense) philosophers of physics should be in the philosophical conclusions they draw, or, to put it a little differently, the depth of the things that we might learn about the world (and about ourselves) from physics. On this question we are beginning to see an emerging divide among philosophers of physics between conservatives and progressives. Conservatives side with common sense and allege that philosophy has gone seriously off the rails when we accept deeply revisionary conclusions from physics. Progressives give common sense little authority except as a repository of folk wisdom, and follow the science where it seems to lead.

For what it is worth, I think that Albert is on the side of the angels here. The epistemology implicit in *After Physics* comes out more clearly than in the earlier work, and underwrites the progressivism that we see also in other parts of the program (e.g., the emergence of space). Unreflectively, common sense always seems to claim the intuitive weight, but one might ask by what right? The possibility that our worldviews might take us beyond imaginative intelligibility in the second sense was always implicit if physics is treated seriously as

[13] On the question 'when we act, would we see influence running in both directions, into the microscopic future and the microscopic past?', my instinct is to say yes. We should see our own actions as making microscopic differences to the past that simply fall below the threshold of macroscopic resolution. That wouldn't quite make knowledge gathering and utilization impossible for reasons that come from statistical mechanics. It would mean that we couldn't form reliable, stable beliefs about the *micro*scopic past, because the microscopic past would be influenced by our own decisions. But those microscopic differences would not affect the past *macro*scopic state, so we could still stabilize beliefs about the macroscopic past, and use that information to guide behavior. The information that our enhanced microscopic vision provides us with, moreover, would not do us much good in the way of bringing about better macroscopic outcomes. Statistical mechanics tells us that better *micro*scopic knowledge of the present will not typically lead to better *macro*scopic knowledge of the future than we get by conditionalizing on the present macrostate. The only way in which the additional microscopic information could lead to different or better *macro*scopic predictions is if we were on an entropy-decreasing trajectory.

[14] They are also contingent, if one takes 'contingent' in its ordinary meaning, as meaning, compatible with the dynamical laws. Indeed, the microscopic laws are not only *compatible* with absence of an entropy gradient; equilibrium is really the natural state of the universe. In the larger scheme of Albert's metaphysics, there is a proposal to treat the Past Hypothesis as a *law*, but the notion of 'law' there is metaphysically thin. Being a law is a logical status, not a metaphysical one.

telling us how we fit into the architecture of nature. It is commonplace that physics proceeds downward from our immediate grasp of the world's coarse-grained toward a more fine- grained view. So it should not be surprising that some of the structure that is basic to our experience will not turn out to be present at the fundamental level. This is so even for some of the structures that make experience itself possible. This is arguably the point at which physics can be most revealing. If the vision of our place in the order of nature that Albert presents in *After Physics* succeeds in recovering all of the structure and local asymmetries that are part of our everyday, untutored experience of the world, but grounds them in emergent features of our world, how could that be an objection to the view, rather than a discovery, and one that could only have been delivered by the methods of physics?

Addendum:

In his response to symposiasts, David points out that there is an asymmetry in the presumption that A lies in the future (rather than past). So, for example, the presumption was that I fix the present state of the world and jump into the pool two seconds in the future, rather than 2 seconds in the past. And he worries that it was the presumption that A lies in the future rather than that past that introduced the temporal asymmetry in the effects of A. I don't think that he is right to worry. It doesn't matter when A happens relative to the present moment, its causal effects are assessed by looking at the effects of hypothetical interventions on A into its own past and future. He was right, however, that my example left things very badly and confusingly underspecified. What I said was that we fix the present macroscopic state of the environment, and assess the effects of some intervention on the past and future. What I *should have* said was that we fix the expected macroscopic state of the environment at the time of the proposed intervention, then assess the effects of the intervention, first on *its own* future, and then on it's own past. There is a difference between holding fixed the present macroscopic state and the macroscopic state at the time of the proposed intervention, if the intervention lies some time in the future. That difference doesn't make a difference to the probabilities if the expected macrostate is derived by conditionalizing present macrostate on the Mentaculus. But it matters to removing the appearance of an asymmetry built into the recipe for assessing the effects of interventions.

So, to assess the effects of a surgical change in the value of a variable A at some time t on the future, we fix the expected macroscopic state of the world at t, break connections with pre-t variables and look at how the effects of change ripple into post-t history. To assess the effects of such a change on the past, we fix the expected macroscopic state of the world at t, break connections with post-t variables and look at how the effects of change ripple into pre-t history. And the claim was that the effects of an intervention on a local, macroscopic variable of the kind typical of human action, ripple into the future, leaving the past in all discernible respects untouched.

Physics and Spacetime Structure: for Symposium in Honor of David Albert's *After Physics*

Pacific APA, April 1, 2016

Wayne C. Myrvold
Department of Philosophy
The University of Western Ontario

1 Introduction

I first want to say what it an honor it is to be participating in this symposium, and to thank Alyssa Ney for organizing it. Like the other participants, my thinking about many things has been very much shaped by reading David's work. To name just one example: my introduction to how Bohm's theory works was David's *Scientific American* article of 1994. There is much that is insightful in *After Physics* and much that we agree on. If, in the following, I focus on points where I don't follow David, it's because we have an opportunity here, in face-to-face conversation, to clear up some matters about which I'm not sure I understand his reasoning.

A recurring theme in David Albert's writings on the metaphysics of quantum mechanics is that quantum theory gives us reason to think that our familiar, three-dimensional space (or four-dimensional spacetime) is derivative, emergent, or in some other sense ontologically inferior to some other space of extremely high dimension (perhaps infinite-dimensional, if we take quantum field theory into account).[1] I do not follow him to this conclusion. This talk will be my attempt to get clear about what it is that we do agree about, what we don't agree about, and why I don't follow him to that conclusion. The reasoning that leads to that conclusion, as far as I can see, has a gap, and I don't see how to fill it. At very least, it seems that there is some assumption or premise that has not been made sufficiently explicit, and I hope that the outcome of this exchange will be to fill it in.

[1]See Albert 1996, 2013; also Chs. 6 and 7 of Albert 2015.

2 What is geometry?

First, a point of agreement. Physical geometry (or, perhaps, chronogeometry, to bring time into it), the subject that treats of the distances between things in space and the durations between events, cannot be divorced from consideration of how things *behave*; that is, it cannot be divorced from dynamics. There are several passages in *After Physics* in which David expresses something along these lines. Here's one.

> what it is to be a table or a chair or a building or a person is—at the end of the day—*to occupy a certain location in the causal map of the world*. The thing to keep in mind is that the production of geometrical appearance is—at the end of the day—a matter of *dynamics* (p. 127).

I find myself *mostly* in agreement with this, with the exception of the word "appearance." The causal map of which David speaks is real, and the geometrical relations evinced in this causal map are not *mere* appearances; they are the ones that matter for physics: if there were a background space with a different metric that played no dynamical role whatsoever, it would have no claim to the title of physical geometry.[2]

If we consider nonrelativistic quantum mechanics, with Galilean invariant potentials, the dynamics of the theory depends on distances in three-dimensional space; in "Elementary Quantum Metaphysics" these were called *interactive distances.* I think that we are in agreement that interactive distances are what occupy the role of spatial distances in physical theory; there aren't any other sort of spatial distances.

3 What is a quantum theory?

Suppose we want to construct a quantum theory of some system. It might, for example, be a large and complicated molecule. What we first do is to identify the dynamical degrees of freedom of interest. For a system of n particles, this will be their positions and spins, if they are not spin-0 particles. We associate with these degrees of freedom operators obeying the appropriate canonical commutation relations; this gives us an algebra of operators associated with the system, the self-adjoint members of which represent "observables." A quantum state is a positive linear functional on this algebra. A pure quantum state can be represented in a number of mathematically equivalent ways, from which we may choose the most convenient for whatever purpose is at hand: as a vector in a Hilbert Space on which those operators act, as a one-dimensional subspace of the Hilbert space, as a projector onto that subspace, ...

When the system is a system of point particles, we can represent pure states via wave functions. For a single, spinless particle, this can be an assignment of complex numbers to points in physical space, whose magnitude squared yields a probability density function for

[2] This is a point made in David's talk at the Tucson Philosophy of Science Workshop, February 27, 2016.

position; alternatively, but equivalently, we can write down momentum-space wave functions, which assign complex numbers to points in a three-dimensional momentum space. One could transform to other variables, or even consider (though it is not clear what the utility of this would be) mixed wave functions, that are functions of (say) one position and two momentum variables.

For a particle of spin s, the spin-space is a Hilbert space of dimension $2s+1$, and a wave function associates a spinor, that is, a vector in this spin space, with each point of its domain (be it position-space or momentum-space or something else). For a system consisting of n particles, we construct the appropriate total spin space that is the tensor product of all of the component spin spaces. A configuration-space wave function associates vectors in this n-particle spin space with n-tuples of points in space, a momentum-space wave function, with n-tuples of points in momentum space.

4 Quantum state realism

A vexed question, which is going on 90 years old, is whether we should think of quantum states as features of the world, rather than, say, of our knowledge of the world or our beliefs about it. Or, better: should we take an ascription of a quantum state to say something about the system itself? I will take any affirmative answer to this question to be an affirmation of quantum state realism. This is not the same as quantum state *fundamentalism*; it is meant to be open to the thought that one might have an ontology that, at the fundamental level, doesn't evolve anything like quantum states, on which quantum states supervene.

Some motivation for realism about the quantum state stems from the eigenstate-eigenvalue link. When an atom or molecule is ascribed an eigenstate of some dynamical variable, this is traditionally thought to be an ascription of a definite value of the variable to the system, and possession of this value is take to be an intrinsic property of the system. The familiar problem, inaptly called the *measurement problem*, is that this doesn't get us very far. The quantum state of any bounded spatial region will almost always be entangled with that of its environment, in such a way that it won't be an eigenstate of any nontrivial local observable. The quantum state, equipped with the eigenstate-eigenvalue link, does not, by itself, give us a world containing cars and bars and wars and experimental outcomes. For our purposes, it suffices to note that the major avenues of approach to this issue—Bohmian theories, collapse theories, and Everettian theories—are all realist about the quantum state. And it is this that, allegedly, leads one to infer the existence of a fundamental space of extraordinarily large dimension. Let's look into the reasoning.

5 The criterion of nonseparability

Prima facie, if a quantum state is meant to represent something physically real, it won't be like anything in classical physics. This is not, of course, an objection, any more than it would be an objection to someone who regards electromagnetic fields as free-standing entities (rather than disturbances in some aetherial fluid) to point out that this, at the time it was introduced, was a novel conception.

On the other hand, there is a line of thought that makes quantum states seem more like something in classical physics. As mentioned, for a single, spinless particle, the quantum state can be represented as an assignment of complex numbers to points in physical space, or to points in momentum space, both of which are three-dimensional. That these numbers are complex marks a difference between such a field and a classical field, but this might be regarded (rightly, in my view) as a minor difference. For a single particle with nonzero spin, the quantum state can be represented as an assignment of a spinor to each point in space. This is less familiar, but might be regarded as a reasonable extension of the notion of a classical field—after all, we are used, in general relativity, to field tensors as a generalization of the traditional notion of a field. What these notions share is that there is an assignment of *something*, be it scalar or vector or tensor or spinor, to each point of spacetime, and each of these assignments can be thought of as a *local* state of affairs at that point, consistent with arbitrary states of affairs outside arbitrarily small neighbourhoods of that point.

For more than one particle, this feature, called *separability*, is lost (at least *prima facie*), as we have to deal with assignments of scalars or spinors in the appropriate high-dimensional spinor space, not to points in space, but to *n-tuples* of points in space. One reaction—which, for a while at least, seemed to be the consensus of philosophers who thought about such things—is to accept that accepting quantum mechanics means accepting nonseparability. This is not David's way. One can retain separability if one takes the quantum state of n particles to be an assignment of local somethings to points in a $3n$-dimensional space (be it configuration space, or momentum space, or some other space). It is this, as far as I can tell, that motivates the idea that this high-dimensional space is more fundamental than ordinary space, as it is only in a high-dimensional space that one can satisfy the demand of separability.

Textual support for this can be found in several of David's essays on this topic. In "Elementary Quantum Metaphysics," we have

> The sorts of physical objects that wave functions *are*, on this way of thinking, are (plainly) *fields*—which is to say that they are the sorts of objects whose states one specifies by specifying the values of some set of numbers at every point in the space where they live ... (Albert 1996, p. 278; 2013, p. 53).

Therein David contrasts the space whose distances are interactive distances, that is, the space that we are used to thinking of as physical space, with what he regards as a more fundamental idea of space, namely,

> a space . . . in which a specification of the local conditions at every address at some particular time (but not at any proper subset of them) amounts to a complete specification of the physical situation of the world, on that theory, at that time (1996, pp. 282–283).

In “Quantum Mechanics and Everyday Life” (Ch. 6 of *After Physics*), we have

> people like Bell have famously announced that the Bohmian-mechanical wave function is “a physically real field, as real here as Maxwell’s fields were for Maxwell.” And it goes without saying that part of what it *is* to be the sort of field that Bell is talking about is to take on values at points in a real, fundamental, freestanding, concrete physical space. And so, if all this is right, the project of getting rid of the higher-dimensional space, or of somehow demoting it to lesser or inferior or secondary ontological status, isn’t going to work (p. 127).

And in “Primitive Ontology” (Ch. 7 of *After Physics*), we have

> one way . . . of *imposing* geometrical relations between the wave function and the mass density . . . is to deny that the $3N$-dimensional space in which the wave fuction undulates is a concrete, free-standing, fundamental space at all, and to regard it (instead) as literally constituted out of N-tuples of points of the *three-dimensional* space. This (indeed) is precisely what proponents of GRW_M typically have in mind—and this will of course make it perfectly clear which particular parts of the wave function are responsible for determining which particular parts of the mass distribution. But this is also going to transform the wave function back into a shadowy and mysterious and traditionally quantum-mechanical sort of thing—it will no longer be the sort of thing (that is) that takes on values at individual points in any free-standing fundamental physical space, and it will no longer be susceptible of being thought about as anything along the lines of concrete physical *stuff*, and we are going to be saddled, once again, with the old-fashioned and unwelcome quantum-mechanical weirdness of nonseparability (p. 149).

The general precept at work here seems to be (and, of course, if this is unfair or uncharitable or a misreading, I’d be pleased to learn that):

> Faced with a physical theory that, taken at face value, seems to violate the condition of separability, we are to find (or construct) another space, such that states of the theory can be regarded as assignments of local quantities to points in that space, and to take some space of this sort as the *fundamental space* of the theory.

If this is the precept in operation, then I have two sorts of question about it. One is: in what sense is a space of this sort (assuming that we can find one and, out of possible candidates, pick *one* out in a principled sort of way), more fundamental than the space in which we live and move, the space whose geometrical structure is given by the interaction-distance metric, that is, the space in which one’s location is picked out as the place in whose vicinity lie the things one most immediately interacts with?

The other question (the answer to which may entail an answer to the first) is: What is the significance of the separability requirement? Why should it be a requirement (if that's what it's taken to be) on an acceptable account of the world that there be some space in which separability is satisfied? What would be objectionable about taking as part of the furniture of the world, for example, a wave-function that takes as its arguments n-tuples of points in 3-D space, and whose role it is to provide dynamics for either Bohmian corpuscles in 3-D space or for mass densities on 3-D space?

I take it that the rhetoric of unwelcome weirdness at end of the above-quoted passage from *After Physics* is not meant to be the actual argument for rejecting nonseparability; we would not want to adopt a general principle of inference that licenses an inference from "X is unwelcome and weird" to "X is not the way the world is."[3] So, we need to look elsewhere for an argument in favour of separability as a requirement on an acceptable theory of fundamental physics.

One well-known advocate of separability was Einstein. In a comment on the last chapter ("Metaphysical Conclusions") of Max Born's book, *Natural Philosophy of Cause and Chance*, Einstein wrote,

> But whatever we regard as existing (real) should somehow be localised in time and space. That is, the real in part of space A should (in theory) somehow 'exist' independently of what is thought of as real in space B. When a system in physics extends over the parts of space A and B, then that which exists in B should somehow exist independently of that which exists in A. ... if one abandons the assumption that what exists in different parts of space has its own, independent, real existence, then I simply cannot see what it is that physics is meant to describe. For what is thought to be a 'system' is, after all, just a convention, and I cannot see how one could divide the world objectively in such a way that one could make statements about parts of it. (In Born 1971, pp. 164–165; this was sent to Born 18 March 1948).

These remarks were echoed in Einstein's *Dialectica* article (1948), a copy of which he sent to Born a few weeks later.

This argument for separability as a requirement on a physical theory strikes me as an argument for a weaker conclusion, namely, an argument against physical theories that lack any local beables, rather than an argument that *all* beables be local beables. If there are local beables in addition to nonlocal ones, one can talk objectively about parts of the world and make statements about the contents of bounded regions of space; one just has to bear in mind that, if we divide the world up into little parts and specify the contents of all these parts, on a nonseparable theory we have not thereby said everything that there is to say about the world.

[3] If reason to hold such a principle in suspicion were needed, the daily news from the Republican primaries provides a good illustration of the rashness of adopting such a mode of inference.

Statements echoing Einstein's have become a mainstay of presentations of quantum field theory. One finds, for instance, in Weinberg's *The Quantum Theory of Fields*,

> It is one of the fundamental principles of physics (indeed, of all science) that experiments that are sufficiently separated in space have unrelated results. The probabilities for various collisions measured at Fermilab should not depend on what sort of experiments are being done at CERN at the same time. If this principle were not valid, then we could never make any predictions about any experiment without knowing everything about the universe (Weinberg 1995, p. 177).

Though this sounds like an argument for strict separability and locality, rhetoric like this is actually used to justify something far weaker, namely, the *Cluster Decomposition Condition*, which mandates that factorizability of the scattering matrix between two spatial regions be approached as the distance between the regions is increased without limit.

If these sorts of considerations don't support support separability as a requirement on viability of physical theories, is there anything that can take their place?

One consideration might stem from the doctrine of Humean Supervenience, which, in Lewis' formulation, is committed to separability.

> Humean supervenience ... is the doctrine that all there is to the world is a vast mosaic of local matters of particular fact, just one little thing and then another ... We have geometry: a system of external relations of spatiotemporal points. Maybe points of spacetime itself, maybe point-sized bits of matter or aether or fields, maybe both. And at those points we have local quantities: perfectly natural intrinsic properties which need nothing bigger than a point at which to be instantiated (Lewis, 1986, pp. ix–x).

If one took this as an *a priori* commitment, then one might be motivated to recast a theory that *prima facie* violates separability into a form in which it can be thought of as satisfying it.

But this is an odd sort of commitment to impose *a priori*. For Lewis, the picture was supposed to be inspired by physics (albeit classical physics) (Lewis, 1994, p. 474), and he took it to be a "contingent, therefore an empirical, issue" (1986, p. x). It is thus not a condition to be imposed *a priori* as a requirement every physical theory must satisfy.

It seems to me that what is *not* at issue is Humeanism about the laws of nature. Earman and Roberts (2005) have suggested, rightly I think, that Humeanism as a doctrine about the laws of nature can be divorced from Lewis' particular commitment to the structure of the Humean base.

The attempt to recast quantum mechanics into a form that satisfies a craving for separability seems to be me to get things exactly backwards; we should be looking to physics for clues about the underlying metaphysics, instead of trying to force physics into some metaphysical straightjacket. In this I am in agreement with Tim Maudlin, in “Why Be Humean?”

> Ontology is the general account of what there is, and our knowledge of what there is is grounded in empirical science, not in a priori speculation or prejudice. . . . As we have seen, contemporary physics posits physical facts that are Non-Separable. What grounding could a preference for Separability have to suggest that we ought to warp either the physics itself, or our account of space to accommodate Separability? (Maudlin 2007, p. 77)

6 Conclusion

The view that quantum mechanics gives us incentive to think that there is a space, more fundamental than the familiar one whose geometrical structure is tied to dynamics, is based on the requirement that a wave function be thought of as a *field* assigning local quantities to points of an underlying space. I can see no reason for this requirement other than a commitment to separability at the fundamental level. But as to the reason we have for taking separability to be a genuine feature of the world, I find myself at a loss. Help!

7 Epilogue

In his responses to my comments at the Pacific APA, David suggested a generalization of his notion of a fundamental space: one can consider, for any x, a space such that a complete specification of the way the world is consists of a specification of conditions at all x-tuples of addresses in that space. David opts for $x = 1$, and suggested that I was opting for $x = n$, where n is the total number of elementary particles in the world.

That was helpful, as it threw into sharper relief the differences between the ways in which David and I were approaching things. I don’t think that a notion of a space of this sort, for any x, is required for doing physics. (In contrast, everyone should regard the 4-dimensional spacetime of interactive distances as physically real, whether fundamental or emergent from something else.)

Consider a theory of the following form: a quantum field theory on some 4-dimensional spacetime, perhaps Minkowski spacetime, supplemented by local beables, such as Bohmian particles or a mass-energy density. The ontology of such a theory consists of the local beables at various addresses in spacetime, together with the quantum state, which, in all its glory, has no address smaller than the whole of spacetime: though we can (and must, in order to apply the theory) talk of the quantum states of bounded spacetime regions, specifying the states of all bounded regions of any fixed finite size won’t tell us everything there is to say about the quantum state. The fact that the ontology of a theory like that contains something, the quantum state, that has no local address, is not an obstacle to it being regarded as a fundamental theory.[4] The notion of "fundamental space" David is talking about is not required to do physics, and, I for one, don't see the utility of trying to find a space of this sort.

References

Albert, D. Z. (1996). Elementary quantum metaphysics. In J. T. Cushing, A. Fine, and S. Goldstein (Eds.), *Bohmian Mechanics and Quantum Mechanics: An Appraisal*, pp. 277–284. Dordrecht: Kluwer.

Albert, D. Z. (2013). Wave function realism. In A. Ney and D. Z. Albert (Eds.), *The Wave Function*, pp. 52–57. Oxford: Oxford University Press.

Albert, D. Z. (2015). *After Physics*. Cambridge, MA: Harvard University Press.

Born, M. (Ed.) (1971). *The Born-Einstein Letters*. Macmillan. Tr. Irene Born.

Earman, J. and J. T. Roberts (2005). Contact with the nomic: A challenge for deniers of Humean Supervenience about laws of nature. Part I: Humean Supervenience. *Philosophy and Phenomenological Research 71*, 1–22.

Einstein, A. (1948). Quanten-Mechanik und Wirklichkeit. *Dialectica 2*, 320–324.

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

Lewis, D. (1994). Humean supervenience debugged. *Mind 103*, 473–490.

Maudlin, T. (2007). Why be Humean? In *The Metaphysics Within Physics*, pp. 50–77. Oxford: Oxford University Press.

Weinberg, S. (1995). *The Quantum Theory of Fields, Vol. I*. Cambridge University Press.

[4]We have good reason to believe that no theory of that sort is a fundamental theory of our world, because it omits gravity, but that's another matter.

DAVID ALBERT'S RESPONSE TO HIS CRITICS

Let me start by saying that I am astonished, and that I am enormously grateful, that people for whom I have such profound admiration have been willing to take the trouble of schlepping themselves up to San Francisco to talk about this very flawed little book. Thank you all so much — and thank you, especially, Alyssa, for the work of conceiving and organizing all this.

The comments are all very rich, and very thoughtful — and it seems to me that the most productive way of reacting to them, in the short time we have left here, is to try to open them up, to try to start conversations about them, to try (that is) not so much to be right, as to be provocative.

*

Let me begin, more or less at random, with some of Craig's remarks. Craig says that the best way to think about the quantum-mechanical wave-function is to deny it any serious ontological status at all, and to assert that it is, instead, *nomological*. The idea is that the wave-function does not represent an entity in the world, a piece of ontology, but rather is part — as Craig puts it — of the representational structure of the laws of physics. The quantum state is therefore akin to the *Hamiltonian* function. Like the quantum state, the classical Hamiltonian generates the motion of the beables and also their statistics. It doesn't "live" in ordinary three-dimensional space — and the space where it *does* live has (as a matter of fact) even more dimensions configuration space. And yet — as Craig rightly says — no one worries about where the Hamiltonian lives, because it is viewed as part of the laws, not the physical world — not (that is) the Humean mosaic — itself. There is simply no *expectation* that laws be functions over three-space, nor that they be *decomposable* into functions on three-space.

And Craig correctly reports that I think that this view is crazy. But I don't think that he entirely does justice — and this is at least in part because I don't think I *myself*, in the book of mine that we're supposed to be talking about here, entirely did justice — to my reasons for thinking that. Let me see if I can do a little better here.

Think for a minute about how it is that we decide what ought to be counted as ontologically serious *stuff*.

Here's an instructive historical example. At the end of the 19^{th} century, with the advent of Maxwell's theory of electrodynamics, people became convinced that the category of serious ontological stuff was larger than they had ever previously imagined — that it consisted not only of material particles but of electromagnetic *fields* as well. How did they decide that? They might easily have taken the sort of route that Craig is suggesting we take with the quantum-mechanical wave-function — they might (that is) have insisted that all there really are, notwithstanding Maxwell's equations, are material particles, and that electromagnetic fields should properly be thought of as merely "part of the representational structure of the laws" of the *motions* of those particles. One way to do that would be to stipulate that the fact that the electromagnetic field has the specific spatio-temporal configuration that it does counts as a law. That (it seems to me) would be a disaster, from the outset, for all sorts of reasons — because the laws would then be no less complex than the world they are supposed to be part of an attempt at *summarizing*, and because the laws would then be no less elaborately *time-dependent* than the world they are supposed to be part of an attempt at summarizing, and because the set of physically possible motions of the particles would then come out to be *infinitely* smaller than we have good reason to think it is (that is: the set of physically possible such motions would be limited to those that happen to be compatible with the particular initial conditions of the electromagnetic field in question), and so on. A more attractive option, I should think, would be to stipulate that a certain set of motions of all of the material particles in the world is physically possible if and only if those motions form a part of at least *one* solution of the unified equations of the electrodynamics of material particles and electromagnetic *fields*. Note that there's *nothing* complicated or time-dependent or otherwise ungainly about a law like *that* — it's simply the familiar and concise and beautiful Maxwellian law of electrodynamics re-phrased in such a way as to make it clear that electromagnetic fields are to be understood, again, as a part of the representational structure of that law, and *not* any part of what the law is actually a law *about*. Why was it, then, that people didn't, in fact, *go* that way? Why was it that everybody at the time, and everybody since, seems to have concluded that Maxwell's equations leave us with no reasonable choice but to admit an entirely new category of fundamental, ontologically serious, physical *stuff*?

Here's an even simpler (more abstract, less historical) example. Suppose that the Hamiltonian of a certain world is the Hamiltonian of a collection of N Newtonian billiard-balls that interact with one another exclusively by means of elastic collisions. And suppose that we were to somehow get it into our heads that only the first M of those balls are actually parts of the ontologically serious fundamental furniture of the universe — and that the remaining N–M of them are to be regarded merely as a part of the representational structure of the law of the motions of those first M balls. Not a problem. We can immediately write down a neat and simple and clear and time-invariant law of the motions of those first M balls, which looks like this: A certain motion of the first M balls is physically possible if and only if there is at least one solution of the N-ball equations of motion of which the M-ball motion in question forms a part.

But this, I take it, would be a manifestly stupid thing to do. And it will be useful to try to be explicit about exactly *why* — and to consider what lessons that might hold for the more complicated cases of electromagnetic fields and quantum-mechanical wave-functions.

To begin with, the strategy we are imagining here — the strategy (that is) of stripping billiard-balls M through N of their ontological status — makes profound metaphysical distinctions between items that the *mathematical structure* of the theory is clearly treating as entirely *on a par* with one another. And there's another thing, or maybe another way of saying more or less the same thing, which feels even more significant: The world we are imagining here — the one (again) in which balls *M* through *N* have been stripped of their ontological status — is wildly indeterministic. Indeed, the entire history of a world like that up to an arbitrary time *t* is in general going to license *no predictions whatsoever* — neither deterministic predictions nor probabilistic ones — about what's going to happen even so little as an instant *later* than *t*. And none of that is going to change so much as an iota if we decide to rehabilitate one or two or three of the nomic balls back into the fundamental ontic furniture of the world — unless and until (that is) we rehabilitate *every last one of them* — at which point, all of a sudden, we have a deterministic world with a well-posed initial value problem. It's as if the mathematics is positively *screaming* at us that *there were always N real balls* — and that reason that the world looked so odd to us when we tried to insist that there were only *M* of them was for no more complicated reason than that *we were idiotically failing to take account of what was simply and concretely and ontologically **there***.

And I think that the reasons why everybody has always taken the electromagnetic field ontologically seriously, and the reasons why everybody *ought* to take the quantum-mechanical *wave-function* ontologically seriously, are going to turn out to be very much of a piece with the reasons we have just been through for taking billiard-balls *N*–*M* ontologically seriously. The argument is going to have to be elaborated somewhat further, and with more care, and in more detail, in those more complicated cases — but the crux of it, I think, is going to turn out to be very much the same.

*

Wayne has an altogether different take on this stuff. He's squarely on *my* side of the question about the ontological status of the wave-function. He is never in any doubt, in so far as I can see, that the wave-function counts as full-fledged, first rank, concrete physical *stuff*. But once *that's* granted, the next question that comes up is the question of what sort of a *fundamental space* is required in order to *house* stuff like that. And everything is now going to depend on exactly what one means by this notion of *housing*.

Here's the way I was thinking about it: the fundamental space of the world is the totality of opportunities for things, at any given time, to be one way or another — the fundamental space of the world is the manifold of points, or (rather) the *smallest* manifold of points, such that a specification of what's going on at every one of those points at every time amounts to a complete specification of the physical history of the world. And it follows right away, of course, that the fundamental space of a non-relativistic N-particle quantum-mechanical world is going to have 3N dimensions.

Now, the conception of what it is to be the fundamental space of the world that I just described can be viewed as one particular member of a *family* of such conceptions — parameterized by a variable X — as follows: The fundamental space of the world is that manifold of points, or maybe the smallest manifold of points, such that a specification of what's going on at every *X-tuple* of those points at every time amounts to a complete specification of the physical history of the world.

My own conception, then, would correspond to the choice X=1. This choice — as Wayne says —amounts to a demand that the physical situation of the world be separable in the fundamental space of the world — and it will entail that the fundamental space of a non-relativistic N-particle quantum-mechanical world has $3N$ dimensions.

Wayne, on the other hand, suggests that we think of the wave-function as taking on values at *N-tuples* of points in a *3-dimensional* fundamental space. So his conception corresponds to the choice X=N.

And of course there will be myriad other choices available as well. If we choose X=$3N$, then the fundamental space of a non-relativistic N-particle quantum-mechanical world will turn out to have 1 dimension. And if we choose X=$3N/2$ (supposing that N is even) then the fundamental space of a non-relativistic N-particle quantum-mechanical world will turn out to have 2 dimensions. And so on.

And it goes without saying that we could play exactly the same sort of game with classical theories as well. What has always been taken for granted, in theories like Newtonian Mechanics or Maxwellian Electrodynamics, is that the fundamental space of the world is the one that corresponds to the choice X=1 — what has always been taken for granted (that is) is that the fundamental space of the world is the smallest manifold of points such that a specification of what's going on at every one of those points at every time amounts to a complete specification of the physical history of the world. But if we decide that things like mass-densities and electromagnetic fields are in fact the sorts of things that can take as their arguments more general X-tuples of points in the fundamental space of the world, then nothing is going to stand in the way of setting X at (say) 3 — In which case the fundamental space of the world of classical physics will be 1-dimensional.

And the question is what sorts of principled arguments there might be for settling on this or that particular value of X.

The arguments for X=1 have to do with the sorts of intuitions I alluded to above — intuitions to the effect that what it is to be the fundamental space of the world is to be the sort of manifold within which the entirety of the physical furniture of the world can be *housed*, the sort of manifold into which the entirety of the physical furniture of the world can be *embedded*, the sort of manifold whose elements constitute the entirety of opportunities for things to be one way or another. And note that these sorts of arguments are going to apply in exactly the same way, and for exactly the same reasons, to any physical theory, classical or quantum-mechanical, whatsoever.

But I'm not sure what sorts of principled reasons there could be for suspecting that *X* ought to be equal to (say) the number of particles in the world. One could say, of course, that in the case of non-relativistic quantum mechanics one *stipulates* that *X* is equal to the number of particles in the world just in order to obtain the result that the fundamental space of the world is 3-dimensional — but that would seem to me to be grotesquely begging the question. I mean — what we're supposed to be doing here, in so far as I understand it, is not figuring out how to *impose* this or that particular fundamental space of the world on the physics — but (rather) to figuring out what it is that the physics is trying to *tell* us about the fundamental space of the world! Wayne must favor some way of figuring that out — Wayne must favor some general *principle* for figuring that out — which differs from the *X*=1 principle I sketched out above. But I'm not sure I understand — and I would be very eager to hear more about — what sort of a principle that is.

*

Ok. Back to Craig — but now on the topic of the relationship between statistical mechanics and the special sciences.

There are one or two things that Craig says that I just don't understand. I mean, I don't know what the deal was with *Schrödinger*, but nobody *I* know wants to derive biology or economics from non-equilibrium thermodynamics — they want to derive it from *statistical mechanics*. There are all sorts of questions on which non-equilibrium thermodynamics is simply *agnostic*, but statistical mechanics — as Craig himself, a bit later on in his remarks, is at pains to point out — is agnostic about absolutely *nothing*. But never mind. Let me go directly to what seems to me to be the heart of the matter.

Craig wants us to consider an image of the sciences which, while it is adamantly physicalist, nonetheless '*gives voice'* — as he puts it — to the idea that the special sciences are in some sense autonomous. The thought is that the statistical *regularities* of the special sciences might represent statistical *flukes* from the standpoint of fundamental physics, and (conversely) that the statistical regularities of *fundamental physics* might represent statistical flukes from the standpoint of this or that particular *special science*, and that there is nothing further or more absolute to say than that about what's *really* a fluke and what isn't. And so, although (say) eco-systems are made of atoms, it isn't going to follow that the *Humean regularities* of ecology are going to be *derivable* from the Humean regularities of atoms — or that either one of them is somehow going to take *precedence* over the other. On this sort of a picture, says Craig, "physics, like ecology, is just another window onto the world — and all the modal features arise in our making sense of patterns, wherever they may be found." "On this picture, coincidences (which is to say: flukes) are the *norm*, seen whenever one translates between different systems."

And this is indeed a very different picture than the much more militantly reductionist one presented in *After Physics*. And I won't have time to do anything more, at the moment, than just to mention two things that puzzle me about it. One is a formal point and the other is maybe a more practical or methodological one. The formal point is just the trivial remark that there is, unavoidably, ineluctably, whether you like it or not, something special about physics — which is that physics is, by definition, the most *comprehensive* of the sciences. And so — although it seems perfectly intelligible to say that ecology turns out to be a fluke from the standpoint of physics (because physics assigns a determinate probability to *every* formulable proposition whatsoever about the physical history of the world) — it isn't at all clear what it might mean to say that physics turns out to be a fluke from the standpoint of ecology (because ecology is presumably agnostic with respect to the question of whether this or that is the best Humean system for the totality of the fundamental physical mosaic). And the methodological point is that it seems absolutely clear to me, and I put it to Craig that it is also absolutely clear to him, that if the statistical regularities of ecology were to turn out to be a fluke from the standpoint of fundamental physics, if it were to turn out (for example) that radioactive nuclei actually decay at different rates in swamps than they do in intergalactic space, then the scientific community would be thrown into crisis – and very rightly so — and it isn't going to calm anybody down, under such circumstances, to be reminded that (after all) "physics, like ecology, is just another window onto the world — and all the modal features arise in our making sense of patterns, wherever they may be found".

*

And let me, finally, turn to Jenann's remarks. The first thing I have to say is that I am just so grateful, and so full of admiration, for the way Jenann does philosophy — both in general and in the particular case of her contribution here today. She's the real thing. She doesn't waste her time with anything other than the largest and deepest and most fundamental questions. She has no interest at all in being clever or in scoring points. She somehow makes every conversation of which she is a part into a brave and exhilarating intellectual adventure — and she ennobles the activity of philosophical discourse. And one of the consequences of that is that I am going to be mostly out of my depth in trying to respond to her here. But let me try, at least, to keep the conversation going.

Jenann raises a whole bunch of really really interesting questions here — and I'm not going to have enough time even to so much as mention them all. But let me start with this:

Let's suppose — at least for the sake of this particular conversation — that Jenann is right that it was silly of me to be so pointedly non-committal about what it is about the world over which we take ourselves to have direct and unmediated control. Let's suppose (that is) that it isn't brain states, and it isn't muscular contractions, and what it obviously and ineluctably *is*, exactly as Jenann says, and for exactly the *reasons* Jenann says, is our own *bodily movements*. I think there's actually more that needs talking about here — but let's leave it to one side for the moment.

I don't know what to think of the claim that those movements are exercises of a genuine libertarian kind of freedom. That's deeper than I have any aspiration to dive in the sort of work that we've been talking about here. Let's leave that to one side too. What *I* have been aiming to do, and *all* I have been aiming to do, in my discussions of what Jenann calls the 'practical' asymmetry — is to understand the fact that the sort of influence that such movements can have on the future is very different than the sort of influence that they can have on the past — that the sort of *control* that one can exert, by means of such movements, over the future, is much *larger*, and much more *varied*, and much more *significant*, than the sort of control one can exert, by means of such movements, over the past.

And Jenann seems to think that counterfactuals are the wrong tool with which to attempt an understanding like that. And I'm not altogether sure I understand why she thinks that — and to the extent that I *do* understand it I'm not sure I agree. But what follows is, again, and at best, only the very beginning of a much much longer conversation.

Jenann says that counterfactuals are about worlds other than the actual one — but certainly, on anything like a Lewisian metaphysics, that's just not true. On anything like a Lewisian metaphysics counterfactuals are claims about the actual world in exactly the same way, and to exactly the same degree that claims about the laws of nature are claims about the actual world. And there is (moreover, and more importantly) a sense in which the indicative conditionals that Jenann proposes as a replacement don't seem to me to be up to the job.

Jenann suggests that the right way for an agent to assess the likely effect of some local macroscopic change *A* in the environment — the raising of her arm, say, or the pressing of a button — is as follows: "She starts with information about the present, surveyable macroscopic state, conditionalizes the Mentaculus on that. This leaves her with many specific beliefs about the past, and few such secure beliefs about the future. Now, she conditionalizes on *A*, and what she sees is the probabilistic effect of *A* rippling into the local future along routes that she can identify, leaving the past mostly untouched. If she jumps into a swimming pool, for example, that makes it vastly more likely that she will be wet, that the water will be disturbed, that she won't be eating a hot dog, and so on. What all of this means is that the causal landscape relevant to human action will have influence running into the future in more or less the way that common sense supposes. This doesn't change anything in Albert's account. It simply brings the asymmetry of influence a little more transparently into alignment with the epistemic asymmetry." And Jenann is quite right that if we evaluate the effect of performing *A* in this way, the evaluation will require nothing along the lines of a metric of distances between possible worlds, and that it will be possible to read the result of that evaluation directly off of the Mentaculus, and that everything is going to work (in that respect) exactly as it does in the case of the asymmetry of epistemic access. But what *I* have always aspired to do — in my various discussions of the practical asymmetry — is to show how that asymmetry (that is: the *practical* one) can be understood as a consequence of the fundamental laws of physics, to show how that asymmetry can be understood (like the thermodynamic and epistemic asymmetries) as a *mechanical phenomenon of nature*. And what worries me about the way Jenann has set things up in the passage I just quoted is that *the way that she raises the question* of the likely effect of some local macroscopic change *A already has a time-asymmetry* ***built into it***. Jenann's suggestion (remember) is that the agent *first* conditionalize the Mentaculus on the present, surveyable, macroscopic state of the world, and *then* conditionalize *that* on her *future* performance of *A*. And the worry is that it is the fact that the conditionalization on the surveyable macroscopic state of the world refers to an *earlier* time than the conditionalization on the performance of *A* — rather than anything in the fundamental laws of nature — that ends up producing the time-asymmetry. Or some of it. Or something like that. And my thought was that the only reliable way *avoiding* such worries — the only way (that is) of assuring ourselves that whatever temporal asymmetries we discover in our capacities to influence events are not mere by-products of some such asymmetry that is somehow hidden in the way we use words like 'influence', but are (rather) rooted in the fundamental objective structure of the world — is to adopt a method of evaluating the effects of the performance of A which is fully and explicitly time-reversal symmetric. And *that* — I think — is going to require that we evaluate the difference between performing A and not performing it *now*, rather than some time in the future. And the evaluation of a difference like *that* is the evaluation of a *counterfactual*.

It may be that all this is a symptom of some deep and genuine disagreement. I bet it isn't — but I'm not altogether sure. Jenann says at one point in her remarks that "there is no intrinsic direction of influence built into the fundamental fabric of the world" — but it's supposed to be *precisely* the lesson of *my own* work on this stuff that there *is*.

Ok. I know I'm running overtime here — so let me finish up, very quickly, by just trying to provoke one or two other conversations. Jenann asks — near the end of her remarks: "What does the world look like to microscopic agents who can intervene microscopically and see the effects of their interventions? What would *we* see if our eyes were suddenly enhanced so that we could see right down to the fundamental fabric of the world? What if technology got good enough that we could harness strategic routes of influence into the past?" And I just want to make two very preliminary remarks about that. The first is that I, too, find it hard to think about what the world would look like, I find it hard (in particular) to think about whether anything would survive of our ideas of *agency*, if we could know and predict the exact and complete microcondition of the world. But I *have* been thinking of myself as arguing, in what I hope has been a very general way, that *no* imaginable advances in technology could make it possible for *anything* that it makes any sense to think of as an *agent* to exert volitional control over the past in the way that we familiarly do over the future —and if I have failed to make that clear, then I have indeed expressed myself poorly.