Local Beables and the Foundations of Physics

Tim Maudlin
New York University

Introduction: The Theory of Local Beables

John Bell's most celebrated contribution to the foundations of physics is his famous theorem. The theorem demonstrates that any physical theory capable of generating the predictions of the standard quantum-mechanical algorithm, in particular the prediction of violations of Bell's inequality for experiments done at space-like separation, cannot be local. The sense of "locality" used here is the same sense that Einstein had in mind when he pointed out that the standard interpretation of the quantum algorithm was committed to "spooky action at a distance". To this day, the import of Bell's theorem is not universally appreciated. I have written about this elsewhere (Maudlin, 2014), and others in this volume will take up that task. It is properly the main focus during this 50^{th} anniversary of that great achievement.

But it is also important to recall and celebrate Bell's other achievements. In many of his later writings, including "The theory of local beables", "Quantum mechanics for cosmologists", "On the impossible pilot wave", "Beables for quantum field theory", "Six possible worlds of quantum mechanics", "Are there quantum jumps?" and "Against 'measurement'" [1], Bell turned his attention to the more general problem of physically construing the mathematical formalism used to derive these predictions. This activity is often denominated "interpreting quantum theory", as if there were some precise physical theory that might somehow be supplemented with an "interpretation". Once framed this way it is easy to ask: But if I already have a theory in hand, what can be gained by supplementing it with an "interpretation"? Many physicists, at this juncture, are happy to conclude that "interpretations" are not a matter of physics at all—maybe they are only of interest to philosophers—and that therefore the whole enterprise of "interpreting quantum theory" is not within the purview of physics *per se*.

What then is in the purview of physics proper? One answer to this question goes under the banner "instrumentalism": all physics, as such, is concerned about is *predicting the outcomes of experiments*. In the service of making these predictions physicists may invent various mathematical formalisms, together with rules for their use as prediction-generating instruments. It is neither necessary, nor perhaps even desirable, to accompany these prediction-generating algorithms with any "picture" or "account" or "story" of what exists *beside* the instruments. Indeed, a common myth about quantum theory is that it is actually *impossible* to provide any such accompanying story, and that the progress of physics requires the positive renunciation of the desire for one. If this were correct then the desire for anything

[1] All of these are reproduced in Bell 2004, from which the page citations will be taken in this paper.

more than such a prediction-generating set of rules must arise from concerns outside of physics proper.

Bell rejected this account of physics root and branch. As usual, he expressed his dissatisfaction so clearly and elegantly that there is nothing to do but quote him:

> In the beginning, natural philosophers tried to understand the world around them. Trying to do that they hit upon the great idea of contriving artificially simple situations in which the number of factors involved is reduced to a minimum. Divide and conquer. Experimental science was born. But experiment is a tool. The aim remains: to understand the world. To restrict quantum mechanics to be exclusively about piddling laboratory operations is to betray the great enterprise. A serious formulation will not exclude the big world outside the laboratory. (Bell, 2004: 216-17)

Physics itself aims at more than just predicting the outcomes of experiments. What more is easily stated: physics aims at a complete and accurate account of the physical structure of the universe. Of course! And the different "interpretations of quantum theory" are really different physical theories, which happen to make exactly, or nearly, the same predictions as the standard quantum-mechanical algorithm. But what general features should such a physical theory have?

One of Bell's signal contributions to this problem is what he called the *theory of local beables*. There is a certain irony here. For while his most famous achievement was to show that the *non*-locality that Einstein long ago identified in the standard "interpretation" of the quantum formalism (the Copenhagen interpretation) could not be eliminated, his attention to local beables highlighted just the opposite problem: the standard story *fails* to be clear about what exists locally. So the standard account, if one tries to take it seriously, both *contains* a non-locality that was not acknowledged and *lacks* a different kind of locality that it requires. It is this second sort of locality I want to discuss here.

Any clearly formulated and articulated physical theory should contain an *ontology*, which is just a statement of what the theory postulates to exist. The word “ontology” can perhaps look a little intimidating, or overly “philosophical”, so Bell invented his own terminology for this: the “beables” of the theory. Stating what the beables of the theory are is nothing more nor less than stating what the theory postulates as being physically real. Once what the ontology *is* has been made clear, then (and only then) can one go on to ask what the ontology *does*, how it behaves. This question is answered by a *dynamics*: a mathematically precise characterization of how the beables change through time. The dynamics might be deterministic or might be stochastic. But according to the professional standards of mathematical physics, the dynamics ought to be precise. It should be specified in sharp equations relating the beables, rather than by using vague words (such as “measurement”).

Within the ontology (the beables) of the theory yet another distinction can be made. Some of the beables (but not necessarily all) are *local* beables. The local beables are

> those which (unlike for example the total energy) can be assigned to some bounded space-time region. For example, in Maxwell’s theory the beables local to a given region are just the fields **E** and **H** in that

> region, and all the functionals thereof. It is in terms of local beables that we can hope to formulate some notion of local causality. Of course, we may be obliged to develop theories in which there *are* no strictly local beables. That possibility will not be considered here.[2]

The obvious *non*-local beable that arises in the context of quantum theory is, of course, the "wavefunction" or "quantum state" of a system. My own preference is to use "wavefunction" for the mathematical representative of this piece of physical ontology and "quantum state" for the beable itself. Any theory that can properly be called a quantum theory must have some such part of its ontology. But it is not a *local* beable. As Bell wrote "[I]t makes no sense to ask for the amplitude or the phase or whatever of the wavefunction at a point in ordinary space. It has neither amplitude nor phase nor anything else until a multitude of points in ordinary space are specified".[3] So the wavefunction—or the quantum state it represents—is not a local beable. Any theory that only commits itself to the existence of the quantum state faces the challenge of providing an account of the physical universe without any local beables at all.

At the end of the quote above Bell mentions such a possibility, only to dismiss it from present consideration. And he never returned to discuss the possibility in any of his other writings. This is already *prima facie* evidence that he considered the local beables to play a central role in any physical theory. Our main challenge now is to clearly articulate what that role is, and what sorts of local beables might play it.

---

[2] From "The theory of local beables" (Bell, 2004: 53).
[3] From "Are there quantum jumps?" (Bell, 2004: 204-5).

Two Examples Proposed by Bell

Bell clarifies the methodological role he takes the local beables of a theory to play when discussing some possible choices for them. The most explicit discussion—which is characteristically quite compact—occurs in "Beables for quantum field theory". Since the topic is quantum field theory, one might well expect that the local beable will be some sort of *field*, i.e. something continuously distributed in space-time. But it is not. He first reminds us what beables are:

> In particular we will exclude the notion of 'observable' in favor of that of '*be*able'. The beables of a theory are those elements that might correspond to elements of reality, to things which exist. Their existence does not depend on 'observation'. Indeed, observation and observers must be made out of beables. (Bell, 2004: 174)

Of course, specifying how an observer such as a human being is "made out of beables" would be a monumental task, requiring detailed physiology and biology. But Bell makes clear that he does not (yet) demand this. What then is the criterion for an adequate choice of beable?

The following passage provides the key:

> Not all 'observables' can be given beable status, for they don't all have simultaneous eigenvalues, i.e. do not all commute. It is important to realize therefore that most of these 'observables' are entirely redundant. What is essential is to be able to define the

> positions of things, including the positions of pointers or (the modern equivalent) of ink of computer output. (Bell, 2004: 175)

Bell first considers a quantity that might well be field-like—the energy density—but rejects it on technical grounds.

> We fall back then on a second choice—fermion number density. The distribution of fermion number in the world certainly includes the positions of instruments, instrument pointers, ink on paper….and much, much more. (Bell, 2004: 175)

This last sentence holds the key to the method of local beables.

We began by insisting that a properly articulated physical theory must have an ontology, a set of entities that (according to the theory) exist. Further, some of these entities (but not all) must be local in the sense that they exist in delimited regions of space-time. Why should it be essential for a theory to postulate such things? Because, as the sentence above indicates, it is here that *the language of the theory makes contact with reports of the empirical data*.

What, after all, does a report from a lab, or indeed any observation report look like? It may describe the lab set-up using macroscopic language. It probably describes the outcome of the experiment in terms of macroscopic situations of things (pointers). Or if not pointers, the experimental apparatus certainly outputs its result as a certain disposition of ink on paper. So if a theory can accurately predict the behavior of the pointers and the distribution of ink on paper, then it correctly predicts the reported data.

And, of course, much, much more. A theory's specification of the fermion density in every region of the universe entails the distribution of matter at macroscopic scale. And if what it predicts at macroscopic scale matches everything we think we know about the world at macroscopic scale (including where the pointers ended up pointing, where the ink is on the paper, the shape of the earth, the dimensions of the Empire State building, etc., etc., etc.) then the theory is *empirically adequate* in any reasonable sense. There may be objections to such a theory, but they cannot rightfully be called empirical objections.

Why the emphasis here on macroscopic scale? Because although we take ourselves to know a lot about the microscopic structure of many things, the evidence for that knowledge always appeals to macroscopically observable facts. A theory that gets the macroscopics right accounts for all our evidence about microscopic structure *even if it also contradicts what we believe about microscopic structure*.

This is not just an idle possibility. We have seen that Bell made the somewhat surprising choice of fermion number density (rather than a continuous field quantity) for a local beable in quantum field theory. Even more striking was his choice for a local beable in his version of the GRW collapse theory. At a purely mathematical level (but not an *ontological* level) the theory is specified by just a wavefunction and its dynamics. Since the wavefunction is a continuous field-like mathematical object, one might again anticipate that any local beable postulated by the theory would also be field-like. But Bell's choice was not. Instead he proposed the completely novel "flash" ontology. Continuing the passage about the wavefunction cited above, Bell writes:

> However, the GRW jumps (which are part of the wavefunction, not something else) are well localized in ordinary space. Indeed, each is centered on a particular spacetime point ($\mathbf{x}$,t). So we can propose these events as the basis of the 'local beables' of the theory. These are the mathematical counterparts in the theory to real events at definite places and times in the real world (as distinct from the many purely mathematical constructions that occur in the working out of physical theories, as distinct from things that may be real but not localized, and as distinct from the 'observables' of other formulations of quantum mechanics, for which we have no use here). A piece of matter is then a galaxy of such events. (Bell, 2004: 205)

These localized point-events have come to be called "flashes".

To what extent does the distribution of flashes in space-time constitute an adequate representation of a piece of matter? A calculation of the density of these events in space-time, according to the GRW dynamics, reveals that they reflect very, very, very little of the detailed spatial structure that we attribute to things at microscopic scale. To take a simple example, and average DNA molecule contains about 2 x 10^{12} atoms. A generous rough estimate is therefore 10^{14} quarks. But each quark, according to the standard GRW dynamics, suffers a collapse only once every 10^{15} seconds. So each strand of DNA would manifest itself in space-time with less than a single point-event a second. No double-helix structure could be inferred from full information about these flashes without centuries of data (and even then one could not easily sort out which flashes belong to which molecules, how the molecules might have "moved" in the interim, etc.). The spatial structure of a cell would, according to this theory, be inexpressibly less detailed than we think it is.

Nonetheless, the motion of a whole human body at macroscale would come out just right. The whole body would be associated with over 10^{13} flashes per second: much more information than required to convey everything we think we can observe with the naked eye. And when we use, say, a scanning electron microscope the theory will make accurate predictions about the macroscopically observable images produced by the machinery. This is how the theory of local beables allows us to determine what the macroscopic predictions of a theory are, and hence whether the theory should be counted as consistent with all available data.

Note that we have, *en passant*, completely solved the "measurement problem" or the problem of Schrödinger's cat. The measurement problem is best formulated not as a problem about 'measurement' *per se*, or about laboratory operations, but about the *universal scope* of the quantum theory. As a proposal for the fundamental physical structure of everything that exists, quantum theory ought to be applicable to everything, including stars, planets, laboratory equipment, and cats. The macroscopic structure and behavior of all these things ought to be just the cumulative or aggregate structure and behavior of their microscopic parts. If the beables of a cat, for instance, are provided by the fermion density of particles in the cat, then it is easy to tell from that whether the cat is alive and kicking or has expired. Similarly for the distribution of flashes in the flash ontology. The postulation of local beables is exactly what allows for this result without appeal to any "shifty split" between a "classical" and "quantum" domain. The universe, according to these theories, is quantum-mechanical through and through. And the local beables, existing at microscopic scale, automatically determine also the macroscopic local spatio-temporal properties of things.

Pilot Wave Theory and Bell's Many Worlds

Beside inventing his own novel proposals for the local beables of quantum field theory and of the GRW collapse theory, Bell also expressed great admiration for another theory with a clear commitment to local beables: pilot wave theory. The local beables here are more familiar than flashes. They are plain old vanilla particles. It is simplest to consider them as point particles—giving them a finite volume would not contribute to the role they play in the theory—which always have definite positions and always move on continuous trajectories. Indeed, the concept of a point particle is so familiar that Bell expends no words at all on further elaboration. In the pilot wave picture matter is made of particles that move in precisely specified ways. There is, of course, also a real, physical non-local beable represented by the wavefunction. This non-local object determines the motions of the particles via the *guidance equation*

.

One of the virtues of Bell's presentation of the theory was to avoid the talk of a "quantum potential" and the appeal to Newtonian dynamics that sometimes obscures the fundamental architecture of the theory. Given an initial wavefunction, an initial configuration of particles, a dynamics for the wavefunction (e.g. Schrödinger's equation) and the guidance equation, one has a completed physics. Supplied with these four pieces of information, it becomes a matter of pure mathematical analysis to determine how the particles will move, and therefore what the macroscopic behavior of material objects will be. That can then be compared with the reported results of observation.

It cannot be overstated both how simple this physical theory is and how hard it is to get it across to the average physicist. Somehow, the very notion of a point particle moving around in some precisely specified way is taken to be incoherent, or impossible. Even if the basic ontology is understood, strange objections are raised like this: according to the theory, an electron in the ground state of an atom is *static*: it does not move. True. But why anyone takes this to be an *objection* to theory is completely obscure. If those atoms and their electrons (which are of course not directly visible) demonstrably move *en masse* so as to produce the macroscopic motion of matter that we think we see, then the theory has no empirical problems. All of this was so transparently obvious to Bell that he makes no comment about it.

Bell expressed his admiration for the pilot wave theory in many places. In "Six possible worlds of quantum mechanics" we find this:

> The last unromantic picture I will present is the 'pilot wave' picture. It is due to de Broglie (1925) and Bohm (1952). While the founding fathers agonized over the question
>
> 'particle' *or* 'wave'
>
> de Broglie in 1925 proposed the obvious answer
>
> 'particle' *and* 'wave'.
>
> Is it not clear from the smallness of the scintillation on the screen that we have to do with a particle? And is it not clear, from the diffraction and interference patterns, that the motion of the particle is directed by a wave? De Broglie showed in detail how the motion of a particle, passing through just one of two holes in the screen, could be influenced by waves propagating through both holes. And so influenced that the particle does not go where the waves cancel out, but is attracted to where they cooperate. This idea seems to me so natural and simple, to resolve the wave-particle dilemma in such a clear and ordinary way, that it is a great mystery that it was so generally ignored. (Bell, 2004: 191)

Even more telling, perhaps, is his account of the evolution of his own interest in foundational matters. As a young man, Bell had been taught that the observable quantum-mechanical phenomena (such as the double slit experiment) simply *could not* be accounted for by any theory that postulates beables existing independently of observation and measurement, and certainly not by any theory in which those beables are governed by deterministic equations. As he relates in "On the impossible pilot wave",

> But in 1952, I saw the impossible done. It was done in papers by David Bohm. Bohm showed explicitly how parameters could indeed be introduced, into nonrelativistic quantum mechanics, with the help of which the indeterministic description could be transformed into a deterministic one. More importantly, in my opinion, the subjectivity of the orthodox version, the necessary reference to 'the observer,' could be eliminated.
>
> Moreover, the essential idea was one that had been advanced already by de Broglie in 1927, in his 'pilot wave' picture.
>
> But why then had Born not told me of this 'pilot wave'? If only to point out what was wrong with it? Why did von Neumann not consider it? More extraordinarily, why did people go on producing 'impossibility' proofs, after 1952, and as recently as 1978? When even Pauli, Rosenfeld, and Heisenberg, could produce no more devastating criticism of Bohm's version than to brand it as 'metaphysical' and 'ideological'? Why is the pilot wave picture ignored in text books? Should it not be taught, not as the only way, but as an antidote to the prevailing complacency? To show that vagueness, subjectivity, and indeterminism are not forced on us by experimental facts, but by deliberate theoretical choice? (Bell, 2004: 160)

Bell's concern about fundamental interpretive problems of quantum theory is brilliantly expressed in this passage. It is clear that his concern was inspired, from the beginning, by reflection on what the pilot wave theory had accomplished. And it is equally clear that the main accomplishment for Bell was not the recovery of determinism. It was rather the recovery of *objectivity* and *clarity*. Reference to "the observer", and hence to "observation", "measurement" and so on does not occur in the articulation of the theory. There is rather a clear set of beables and a mathematically precise dynamics for them. *Extracting the empirical predictions from the theory is a matter of appreciating what it implies about how matter will move, nothing else*. And it is exactly the commitment to a precise set of local beables—the point particles—that enables all of the conceptual progress.

Indeed, Bell regarded the role of the particles in the pilot wave theory as so central to understanding its account of the physical world that he even inserted the same set of local beables into theories whose originators had no such thing in mind. This occurs in "Quantum theory for cosmologists". Bell there provides his account of the Many Worlds Interpretation, which he denominates "Everett(?)". The question mark is well deserved. Bell presents the Many Worlds idea in a completely idiosyncratic way, essentially as a version of the pilot wave theory (whose wavefunction, like that of the Many World approach, never collapses) but with a radically stochastic evolution law for the local beables. While the pilot wave particles move continuously through space as a consequence of the guidance equation, Bell's Many Worlds particles have their configuration chosen randomly at each moment, with the probability supplied by the wavefunction. The configuration of particles at any given time is completely uninfluenced by the configuration at any earlier time, and provides no information (beyond what can be extracted from the uncollapsed universal wavefunction alone) about what the future configuration might be. Here are two salient passages from that paper:

> Then it could be said that classical variables *x* do not appear in Everett and De Witt. However, it is taken for granted there that meaningful reference can be made to experiments having yielded one result rather than another. So instrument readings, or the numbers on computer output, and things like that, are the classical variables of the theory. We have argued already against the appearance of such vague quantities at a fundamental level. There is always some ambiguity about an instrument reading; the pointer has some thickness and is subject to Brownian motion. The ink can smudge in computer output and it is a matter of practical human judgment that one figure has been printed rather than another. These distinctions are unimportant in practice, but surely the theory should be more precise. It was for that reason that the hypothesis was made of fundamental variables *x*, from which instrument readings and so on can be constructed, so that only at the stage of this construction, of identifying what is of interest to gross creatures, does an inevitable and unimportant vagueness intrude. I suspect that Everett and De Witt wrote as if instrument readings were fundamental only in order to be intelligible to specialists in quantum measurement theory. (Bell, 2004: 134)

On Bell's reading, the *x* variable that appears in the wavefunction really is a variable representing the possible spatial locations of point particles, which are the fundamental local beables. Thus the Everett(?) theory solves the measurement problem by the very same resources as the pilot wave theory.

The difference between the two theories, on Bell's telling, is solely in the dynamics for the *x*s. In Everett(?) they jump around in the most astonishing fashion, from one configuration to another. Each configuration is drawn from a possible pilot wave theory history, but there is no consistency over time about which of these many possible histories is being sampled. There is, for example, almost no chance according to this theory that one's present apparent memories (which are encoded in the present configuration of particles) provide accurate information about anything that actually occurred in the past. The Everett(?) theory is a sort of physical blueprint for a Cartesian demon, but one where the subject is deceived about the past (and even her own past). As Bell puts it:

> Thus in our interpretation of the Everett theory there is no association of the particular present with any particular past. And the essential claim is that this does not matter at all. For we have no access to the past. We have only our 'memories' and 'records'. But these records and memories are in fact *present* phenomena. The instantaneous configuration of the *x*s can include clusters which are markings in notebooks, or in computer memories, or in human memories. These memories can be of the initial conditions in experiments, among other things, and of the results of those

> experiments. The theory should account for the present correlations between these present phenomena. And in this respect we have seen it to agree with ordinary quantum mechanics, insofar as the latter is unambiguous….
>
> Everett's replacement of the past by memories is a radical solipsism—extending to the temporal dimension the replacement of everything outside my head by my impressions, of ordinary solipsism or positivism. Solipsism cannot be refuted. But if such a theory were taken seriously it would hardly be possible to take anything else seriously. So much for the social implications. It is always interesting to find that solipsists and positivists, when they have children, have life insurance. (Bell, 2004: 134-35)

Bell's Everett(?) theory is certainly not a theory any contemporary Everettian would recognize as their own. Indeed, contemporary Everettians rather claim that they can see no important function for the particles in the plot wave picture! David Deutsch has famously characterized pilot wave theories as "parallel universes theories in a state of chronic denial" (Deutsch, 1996: 225). To borrow yet another phrase from Bell, misunderstanding could hardly be more complete. For Bell, the particles in the pilot-wave theory, the local beables, are exactly what provide the physical foundation of all claims about the macroscopic behavior of material objects in space. Eliminate the particles and you eliminate the heart of the theory. Deutsch and his fellow Everettians, in contrast, are convinced that the empirical consequences flow already just from the behavior of the quantum state, and the particles serve merely as a "pointer" to one or another branch of the wavefunction (Brown and Wallace, 2005: 527). But the wavefunction, as we have repeatedly mentioned, does not represent any sort of local beable at all. So a theory with only a wavefunction (or more precisely, with only a quantum state represented by a wavefunction) has no local beables. It is therefore unclear how such a theory could entail any claims about anything happening in ordinary space or space-time.

The standard contemporary Everettian response to this worry is to gesture at the idea of "functional definitions". A functional definition, if framed at a sufficiently high level of abstraction, is supposed to be the sort of thing that a quantum state (or parts of it), all on its own, can satisfy. But the exact details of how this works, beyond slogans such as "A tiger, instead, is to be understood as a pattern in the physical state" (Wallace, 2003: 92), are not easy to pin down. Specifying, even in the most broad-brush terms, the sort of "pattern" in *a non-spatial* object that would make a tiger is a conundrum. Contrast this with the corresponding question: what sort of behavior of microscopic beables in space-time (particles, say, or flashes) would correspond, at macroscale, to what we take ourselves to know about tigers? It would not be difficult to make passible answers to such a question, and to judge whether a proposed distribution of local beables could possibly be a tiger, without engaging in a speck of functional analysis of anything. Tigers have characteristic shapes and sizes and move in characteristic ways. All of this could be read off of the behavior of the local beables without further ado.

This methodological role of local beables as the part of the ontology of a theory whose behavior determines the theory's observable consequences underpins

Bell's concern. It is this role that explains why Bell never even considers a theory bereft of local beables. He simply would not know what to make of it as an account of the world we take ourselves to inhabit. But with the local beables in place the empirical consequences of a theory are easy, in principle, to identify.

The requirement that a theory posit some sort of local beables in "ordinary space" is not very stringent. As we have already seen by example, many different sorts of local beable are available: particles, fermion number densities and flashes can work, as could local fields (as in electromagnetic theory) or microscopic strings. There is a similar wealth of options for the exact specification of "ordinary space". Classical space-time works, of course, as do the space-times of Special and General Relativity. Extra compactified spatial dimensions, such as those postulated by string theory would pose no difficulty. Nor would a space-time structure that becomes discrete at microscopic scale.

The latitude in the choice of both local beables and the space-time they inhabit arises from the same source: the data against which the theory is tested are ultimately reported as the behavior of things at macroscopic scale. So the detailed, precise microscopic geometry and distribution of local beables according to the theory merely needs to *coarse grain* the right way to provide an adequate account of the data. It is obvious, and has been from antiquity, that both a fundamentally continuous and a fundamentally discrete space-time structure are consistent with the world as we experience it. Not for nothing is detail below a certain scale called "microscopic". Since it is too small to see or otherwise access directly, our evidence about it is always mediated by macroscale events and objects. Many different microscopic structures could coarse-grain into the four-dimensional space-time occupied by macroscopic objects to which we have direct perceptual access.

Bell and Bohr

The insistence that for a physical theory to be empirically acceptable it needs to account for the macroscopic structure and motions of material bodies is not only central to Bell's account of local beables, it was central to the Copenhagen interpretation. Bell's commentary about Bohr's position here is notable. He strongly approves of Bohr's insistence that there *must* be a "classical" part of the ontology of a physical theory if the theory is to be comprehensible. "Classical" here does not mean "obeying the laws of classical physics". Of course, if any physical theory is to be empirically adequate it must somehow entail the existence of macroscopic bodies that very nearly obey the laws of classical physics because the laws of classical physics, in many, many circumstances, describe the behavior of bodies to very, very high precision. But Bell's use of "classical" is essentially identical to his use of "local beable": the classical part of the ontology is the local stuff that is *just there* in space-time, independently of whether it is being "observed" or "measured". Observation and measurement themselves, as Bell insists, must be built out of the things that are just there on their own.

Bell's complaint about Bohr and Copenhagen, then, is not that the account fails to have any local beables at all. It is rather that the observation-independent, locally existing objects postulated by Bohr are all *fundamentally macroscopic*. That is, Bohr is committed to macroscopic local beables *while simultaneously denying the existence of any microscopic local beables.* Bell finds such a position conceptually unacceptable. One part of its unacceptability arises from the fact that "macroscopic" is a vague term: it has no precise meaning. But that does not really get to the heart of the matter. Even if one were to propose some precise criterion of "macroscopic", it is just not at all clear how macroscopic items can have any spatio-temporal characteristics apart from those they inherit in the obvious way from their microscopic local parts. If a tiger, as a whole, has no small parts that are situated in parts of space-time, it is extremely obscure how *it* can be so situated. But if a tiger has microscopic parts situated in space-time, then its spatio-temporal structure is nothing more than the collective spatio-temporal structure of those parts.

Regarded in this way the abstract structure of the Copenhagen theory, as presented by Bell, is identical to the abstract structure of the pilot wave theory, and the Evertt(?) theory, and to GRW flash theory. Each of these theories has a bipartite division of its ontology. One part, the non-local beable(s) of the theory, is the quantum state or quantum states. The other part is some sort of local beable. The non-local part must, of course, have some influence on the behavior of the local items since the observable consequences of the theory all flow from the behavior of the local items. In the pilot wave and Everett(?) theories the local ontology is a collection of particles, in the GRW flash theory it is a set of flashes, in the Copenhagen theory it is a collection of macroscopic "classical" objects including laboratory equipment. These theories may disagree about the dynamics of the non-local part (pilot wave and Everett(?) vs. GRW) and may disagree about how the behavior of the local items are constrained by the non-local part (pilot wave vs. GRW vs. Everett(?)). But these are disagreements of detail between theories with the same general architecture. And the articulation of these theories (pilot wave, Everett(?) and GRW flash theory) is uniformly clear and mathematically precise.

Bell understands Copenhagen to have exactly the same overall architecture, but implemented in a vague and imprecise way. Here is the description from his masterpiece of foundational discussion, "Against 'measurement'":

> Then came the Born interpretation. The wavefunction gives not the density of *stuff*, but rather (on squaring its modulus) the density of probability. Probability of *what* exactly? Not of the electron *being* there, but of the electron being *found* there, if its position is 'measured'.
>
> Why this aversion to 'being' and insistence on 'finding'? The founding fathers were unable to form a clear picture of things on the remote atomic scale. They became very aware of the intervening apparatus, and for the need for a 'classical' base from which to intervene on the quantum system. And so the shifty split.
>
> The kinematics of this world, on the orthodox picture, is given by a wavefunction (maybe more than one?) for the quantum part, and classical variables—variables which *have* values—for the classical part: $(\Psi(t, q \ldots), X(t), \ldots)$. The *X*s are somehow macroscopic. This is not spelled out very explicitly. The dynamics is not very precisely formulated either. It includes a Schrödinger equation for the quantum

> part, and some sort of classical mechanics for the classical part, and 'collapse' recipes for their interaction.
>
> It seems to me that the only hope of precision with the dual (Ψ, *x*) kinematics is to omit completely the shifty split, and let both Ψ and *x* refer to the world as a whole. Then the *x*s must not be confined to some vague macroscopic scale, but must extend to all scales. In the picture of de Broglie and Bohm, every particle is attributed a position *x*(*t*). Then instrument pointers—assemblies of particles—*have* positions, and experiments *have* results. (Bell, 2004: 228)

The foundational significance of Bell's insistence on the postulation of local beables in any precise theory cannot be given clearer or more concise expression.

Even "instrumentalists" who want to abjure all speculation about the microscopic are committed to the real existence of their laboratory apparatus. Any physical theory that purports to describe the physical world must be able to encompass the apparatus within its scope. The postulation of macroscopic localized bodies thereby justifies the search for the microscopic local beables of which they are composed.

Equally important in this discussion are the consequences for the status of the quantum state. For just as elimination of the "shifty split" drives the demand for local beables down to the microscopic scale, so too it drives the scope of the quantum state up to the universal or cosmological scale. The physical universe, as a whole, is a quantum system. And the entanglement of the wavefunction implies that the quantum state of the whole cannot be reduced to a collection of quantum states of the parts. This invites a seldom asked question: if, fundamentally, there is only one quantum state, the quantum state of the entire universe, how do we come to be able to ascribe separate quantum states to subsystems and use them so effectively for making predictions? Whereas the relation of the fundamental microscopic local beables to the local characteristics of macroscopic objects is trivial (the macroscopic object is just where its microscopic parts are), the relation of the wavefunction ascribed to small parts of the universe to the wavefunction of the universe as a whole is much more opaque. But that is the topic for another paper.

Conclusion

Bell's analysis of how the local beables postulated by a theory function in deriving its empirical consequences stands on its own as a contribution to the foundations of physics. From this perspective we can understand both what the "measurement problem" of quantum theory is and how it can be solved in a principled way. We can also see how theories that appear on the surface to be quite different—pilot wave theory, GRW, Everett(?), and even Copenhagen—all appeal to the same basic architectonic[4]. From this perspective, it is obvious why, for example, the decoherence of the quantum state, all on its own, could not solve the measurement problem. For it cannot, all on its own, bring any local beables into existence. It also highlights a challenge for the orthodox Everettian position.

[4] For elaboration on this theme, see Allori *et al*., 2008.

Orthodox Everettians would certainly renounce Everett(?) as their theory: they will have no truck with real particle configurations at all. But what, then, are the local beables that their theory is committed to? No answer is readily forthcoming. This is a completely different problem from the problems about probability and unique outcomes of experiments that are usually discussed. But it is a problem to be solved if the theory's relation to observational data is to be comprehensible.

Bell's theory of local beables is a signal accomplishment in the foundations of physics completely independent of his famous theorem. But also, at the end of the day, it is not irrelevant to understanding the theorem. Violations of Bell's inequality in experiments done at space-like separation indicate some sort of non-locality in the world, exactly the sort of non-locality that Einstein abhorred. There is a lively debate about how exactly to define the non-locality at issue. (It has nothing to do with signaling, for example, as is clear from Einstein's complaint about the standard theory.) But no clear and comprehensible account of any sort of locality or non-locality can even proceed without the postulation of some sort of local beables in space-time. If *nothing* definite ever happens *anywhere* in space-time, how could any question of locality even be posed?

Bell was aware of this connection as well. We have seen that in "The theory of local beables" he remarks:

> It is in terms of local beables that we can hope to formulate some notion of local causality. (Bell, 2004: 53)

Skeptics about the significance of Bell's theorem (of which there are still distressingly many) might try to seize on this remark. "So," they might argue, "if I *refuse* to recognize any local beables at all, no question of the holding (or not holding) of local causality in my theory can even arise". Is this then an escape route from the claim that the violations of Bell's inequality by experiments done at space-like separation show that that actual physics is non-local?

Only in the Pickwickian sense that such a blanket refusal to admit any local beables implies that *no such experiments were ever done at space-like separation at all*. For the results of the experiments—the disposition of the laboratory apparatus at the end of the day—is taken to be some local physical fact about what happened in the lab. Throw out all local beables and you throw out the labs and their results altogether. Bohr would not countenance such a thing, as his remarks about the necessity of a classical description of the laboratory operations shows. Metaphors about babies and bathwater spring to mind.

So not only does the theory of local beables stand on its own as a contribution to the foundations of physics, it also underpins all analyses of the significance of violations of Bell's inequality. Looking back at that epochal proof from the distance of half a century, the view is sharpened by the other great contributions John Bell made since then.

References

Allori, Goldstein, Tumulka, *et al*. (2008). On the common structure of Bohmian Mechanics and the Ghirardi-Rimini-Weber theory. *British Journal for the Philosophy of Science*, **59**(3), 353-389.

Bell, J. S. (2004). *Speakable and Unspeakable in Quantum Mechanics*, 2nd edn. Cambridge: Cambridge University Press.

Brown and Wallace (2004). Solving the measurement problem: de Broglie-Bohm loses out to Everett. *Foundations of* Physics, **35**, 517-40.

Deutsch, D. (1996). Comment on Lockwood. *British Journal for the Philosophy of Science*, **47**(2), 222-228.

Maudlin, T. (2014). What Bell Did. *J. Phys. A: Math. Theor.* **47** 424010.

Wallace, D. (2003). Everett and structure. *Studies in the History and Philosophy of Modern Physics*, **34**, 86-105.