Quantum Metaphysics

Richard Healey

April 22nd 2003

Abstract

Should we allow quantum mechanics to reshape our view of the world and ourselves? What view emerges if we do? Philosophers as well as physicists have answered these questions in many different ways for 80 years without reaching consensus, but at least we now have a better idea of which answers to take seriously. I will offer a rough guide to the possible implications of quantum mechanics, real and merely imagined.

- What is quantum mechanics?
- What is metaphysics?
- Some quantum phenomena and how quantum mechanics accounts for them
- Is quantum mechanics complete: Copenhagen (Bohr & Dirac) vs. Einstein
- Why quantum probabilities concern measurements
- The measurement problem
- Collapse interpretations: GRW
- No-collapse interpretations 1: Bohm
- No-collapse interpretations 2: Everett
- Modal interpretations

–Where do we stand today?
–Can we reach a consensus on how to interpret quantum mechanics just by further careful thought, or will that require new physics?

Quantum mechanics is a theory of physics. It initially arose as a way of accounting for the structure of atoms and the way they emit and absorb light. It has since been successfully applied to a much wider range of physical systems of various sizes, from atoms and their component particles to devices at the much larger scales that have been the focus of this workshop–large molecules, and superconducting devices the thickness of a human hair–to metals, crystals, semi- and super-conductors, lasers, superfluids, and even (more speculatively) to the entire universe in the early stages of its evolution.

A new theory was needed because of the failure of classical physics correctly to account for the behavior of atoms. This predicted, for example, that an atom of hydrogen would be unstable, that it would collapse in a burst of light in about a millionth of a second! Before encountering such failures at the turn of the 20th century, Newton's classical mechanics, and Maxwell's theory of electricity and magnetism, had provided wonderfully successful accounts of all kinds of phenomena. Indeed, these theories are still usefully applied to a wide range of so-called macroscopic phenomena (that is, phenomena involving systems of a size typically encountered in daily life). But that is only because classical physics is good enough to enable us to predict and control these phenomena. It is not, we believe, because its theories correctly describe the underlying processes in any but a rough and ready way. The new theory, quantum mechanics, explains why classical physics works as well as it does, and also succeeds where it fails.

Quantum mechanics has now been supported by a wealth of experimental tests and successfully employed to account for phenomena that were quite unknown when it was first proposed. It is a pillar on which contemporary physics rests. As we have seen at this workshop, it promises to ground amazing 21st century technology. But it is also a deeply puzzling theory. Niels Bohr, one of its founders, said that anyone who is not shocked by the theory has not understood it. The late, great Richard Feynman once claimed that, by contrast with relativity, no-one understands quantum mechanics. This is not because of the mathematical difficulty of the theory, nor is it because of any special problems in applying it to derive predictions. The real difficulty arises when one asks what the world could possibly be like so that it conforms to these predictions. An answer to that question would be an interpretation of quantum mechanics. Many have been offered, but none has won, or deserved, universal acceptance.

Many working physicists have little patience for this question: they often dismiss worries about the foundations of quantum mechanics as “mere philosophy”. I hope to convince you that while pursuit of this question quickly leads one into typically philosophical concerns, there is nothing “mere” about them. Almost any would-be interpretation of quantum mechanics forces one to question deeply held views about the world we experience and the character of our experience of that world. You will have seen several illustrations by the end of my talk. These views are appropriately termed metaphysical, metaphysics being the branch of philosophy devoted to the investigation of the most general categories of being. But before getting down to details I must face some initial skeptical challenges to the project.

Metaphysics suffers from a bad reputation, in both popular and academic circles. The volumes you will find in the “Metaphysics” section of your local book store amply justify this, and some of them these days even have the word “quantum” in their titles! The great philosopher Immanuel Kant set out to confine metaphysical investigation to what he took to be limits imposed by the nature of the human mind. But the progress of science has convinced us that many metaphysical “truths” he claimed to find within these limits are simply empirical falsehoods. Quantum mechanics may provide a salient example: on a widely accepted interpretation, this theory allows events to occur, even though nothing that happened before determined that they would.

The 20th century logical positivists declared all metaphysical assertions meaningless because they were unverifiable. Like Kant, they were profoundly influenced by the science of their day, including quantum mechanics as well as relativity. And their attitude, if not their doctrines, have attracted adherents among physicists from Werner Heisenberg to Stephen Hawking. But the positivists’ views of meaning did not hold up to philosophers’ critical scrutiny, and metaphysics has emerged once again as a flourishing area of philosophical investigation. Some of the best work in this area is practically indistinguishable from that of theoretical physicists, and even appears in the same journals.

But even if metaphysics can be an intellectually respectable discipline, why should its practitioners take the pronouncements of contemporary physics seriously? An influential school of thought in the history and sociology of science holds that they should not. Social constructivists regard the creation of physical theories as an enterprise constrained more by the historical and social contingencies of scientific communities than by their tenuous connections to the very indirect observations made in the highly controlled artificial environments of modern experimental laboratories. In his provocatively titled book Constructing Quarks, Andy Pickering concluded that no-one interested in constructing a world-view should pay any attention to the theories of contemporary physics. But other historians, including Peter Galison in his How Experiments End have provided detailed studies of how physicists carefully weigh the evidence provided by experiments and modify their views more in response to the data these provide than in response to their own interests and prejudices or those of their colleagues. Historical contingencies may influence the fortunes of a physical theory for a while, but what makes it good science is precisely the role of observation as final arbiter of its worth.

But if observation is so important to science, is it not scientifically irresponsible to form beliefs about things that are not observable? Even after the demise of logical positivism, empiricism remains an influential school of thought in contemporary philosophy of science. Its most notable contemporary standard bearer is Bas van Fraassen. He maintains that the goal of science is empirical adequacy–being right about all observable matters–and that all that a scientist should be committed to in accepting a theory is that it is empirically adequate. He denies that it is or should be any part of the scientist’s goal to come up with theories that truly describe a world that lies behind our observations. So while a theory may make claims about unobservable structures, it is not unreasonable for scientists to dispute or ignore these claims, as long as the theory is right about everything that we can observe. Quarks may exist, or they may not: it doesn’t matter to science whether they do or don’t, as long as the observable predictions of theories that posit quarks are invariably correct. Of course, such a view requires a distinction between what is and what is not observable. In this workshop we have seen electron microscope pictures of SQUIDs and heard of observations of tiny currents and magnetic fluxes through these devices, while our unaided senses present us with a SQUID as at most a tiny dot at the limits of vision. We have seen complex 3-dimensional representations of manganese 12 molecules, microscopic images of domain walls in ferromagnetic crystals, and other structures about which sophisticated instruments give us detailed information even though none of this is manifest to our senses. If our powers of observation extend beyond our unaided senses, then why should any structure postulated by a successful theory and detectable by instruments based on it count as unobservable?

But suppose one can draw some distinction between what a theory says about observable and unobservable things. Should this matter to us–to the laymen who finance the scientists, and to the philosophers who try to come to terms with their findings? Van Fraassen's answer is interesting. He takes a theory's assertions about unobservable things to be part of its metaphysical, rather than purely scientific, content. But he certainly does not dismiss this as meaningless. Indeed, in his book Quantum Mechanics: an Empiricist View he has himself set out an interpretation of quantum mechanics, which is precisely an account of how the unobservable as well as observable world might be if that theory is true. He believes that provision of such an interpretation contributes to our understanding of the theory, and that understanding may be further increased by addition of alternative, incompatible, interpretations. As to which, if any, of these is correct, he maintains a deliberate silence. Science cannot answer that question, and nor can a responsible metaphysics, whose ambitions must be restricted to the task of providing us with a menu of options. For him, any choice among these would indeed be purely a matter of taste.

His rival, the scientific realist, disagrees. The scientific realist maintains, on the contrary, that the evidence for quantum mechanics can provide us with reasons to adopt one interpretation rather than another–to draw justified conclusions about the nature of the quantum world that lies behind our observations of it. But contemporary empiricists and scientific realists can join hands in the enterprise of seeking to interpret quantum mechanics. It is to that enterprise that I now turn.

Quantum Mechanics and How it Handles Some Phenomena

Let's start with something everyone agrees on. Quantum mechanics represents the state of a system (for example, a hydrogen atom all by itself) by what is called a wave-function. In the case of our hydrogen atom, the value of this function at a particular place at some time would fix the chance of finding the atom's electron there if one were to look for it then. As long as this chance remains the same, the atom will always be found to have a definite energy. There are many such states, of different energies. But not all energies are permitted. The energies available to a hydrogen atom are said to be quantized, and that is why it is appropriate to call the theory that predicts the available energies quantum mechanics. The state of lowest energy, called the ground state, is stable. States of higher definite energies are called excited. They are not stable. A hydrogen atom in an excited state may later be found in the ground state after emitting light of frequency proportional to the difference between the energies of the excited state and the ground state. Whether, and if so when, this happens is completely unpredictable, though the theory does predict the chance of its happening during any period of time. Quantum mechanics thus accounted for the stability of the hydrogen atom, as well as predicting the frequencies and relative intensities of the light it would emit when excited. As Dirac noted, it also imposed an absolute scale on the world by showing why atoms were just the size they are: Classical mechanics had no such explanation to offer. So far so good.

But other states of a system may also be represented by sums of wave-functions, each of which represents a possible state of that system. In that case, the system is said to be in a superposition of those states. For example, the state of a hydrogen atom may be represented by the sum of its ground-state wave-function and the wave-function of one of its excited states. This state has properties that are not intermediate between, but generally quite different from, the properties of the states represented by the wave-functions that were added to get its wave-function. If a hydrogen atom is in a state that is a superposition of ground state and excited state then the chance of finding its electron in a particular place varies with time. And if one were to measure its energy instead, then one would find either the ground state energy or the excited state energy, but nothing in between. While it is quite unpredictable which energy one would find at (almost) any time, that chance changes from moment to moment in a regular fashion. Moreover, the theory determines what that chance is at any time by applying a rule to the superposed wave-function then: This is just the *same* rule that was applied to the ground-state wave-function to determine that there is never any chance of finding a stable hydrogen atom with any energy other than that of its ground state.

More generally, we can say that quantum mechanics is a theory that applies a general rule to the wave-function to predict the probability that a precisely specified observation on any part of the physical world, large or small, will have one rather than another of various possible results. The rule is known as the Born rule, after a co-discoverer of quantum mechanics. Ironically, it was added in a footnote to his paper at the proof stage, as a correction to the mistaken rule appearing in the body of the paper! The basic reason we don't need to apply quantum mechanics to everyday situations in which we observe large-scale phenomena is that the chance of getting any result except the one we expect from classical mechanics turns out to be unimaginably tiny.

I sketched the way in which quantum mechanics predicts the chance of observing a particular value for the energy of a hydrogen atom in a state represented by a superposed wave-function. But what does the theory say about the value of its energy when it is not observed? Here we leave the safety of common ground and enter into the disputed territory of interpretation! In exploring this territory it will be helpful to have before us another illustration which highlights the role of observation in quantum mechanics.

This is a beautiful device called a neutron interferometer, which has now been in operation in various laboratories for 20 years. At the heart of the device is a large (1 meter) rectangular, perfect crystal of silicon which has had two rectangular "bites" cut out of it so it looks like a letter 'E' on its side. A beam of neutrons (the neutral particles contained in the nuclei of all atoms except hydrogen) is prepared from a nuclear reactor so that all the neutrons are heading in the same direction toward the first face of the device at the same speed. Some of them are observed at the other end of the crystal by detectors that are placed to intercept them as they emerge in particular directions. There turn out to be just two such directions; and the relative numbers of neutrons detected at each depend on what, if anything, the neutrons encountered in the gaps within the crystal.

If nothing is placed in the gaps, then all the emerging neutrons are detected going in the same direction (upward, *U*, say) that they went in at.

If a barrier is inserted at *B* in what we naturally think of as the top path, then only half as many neutrons are detected as before: But now as many are detected going *downward* as upward.

The same thing happens if a barrier is inserted at *C* in the bottom path.

This is surprising if each neutron takes either the top path or the bottom path: How can *blocking* a path a neutron *doesn't* take make it end up somewhere it is never found when that path is open?

Nothing is detected if barriers are placed both at *B* and at *C*.

Now suppose that instead of a barrier, a device is inserted at *B* that detects any neutrons that pass by, without in any way disturbing their passage. Then the detectors placed after the crystal detect as many neutrons as when this device is absent; but now half of them are detected by the lower detector!

Conclusions from the phenomena(?): A neutron follows one path or the other through the crystal if, but only if, it is observed to do so. We cannot say that an unobserved neutron follows one path or the other, though it does seem reasonable to conclude that no unobserved neutron follows some *third* path. It is tempting to conclude that each unobserved neutron in some sense follows *both* paths. But what could this mean? Besides, each *observed* neutron follows just *one* of the two paths.

What does QM say about these phenomena?

1) The wave-function of the neutrons correctly predicts their probability of detection at *U* or *L*

2) This is generally intermediate between 0 and 1, but in special cases may equal 0 or 1.

3) The wave-function does *not* say how any particular neutron passes through the device.

4) Any *observation* of how a neutron passes through the device alters these probabilities.

In the light of such phenomena and their quantum mechanical explanation, it is not altogether surprising that Niels Bohr, one of the founders of quantum mechanics and perhaps its most influential early thinker once remarked:

"There is no quantum world. There is only an abstract quantum physical description. It is wrong to think that the task of physics is to find out how nature *is*. Physics concerns what we can *say* about nature."

If Bohr is right, and the quantum world is *not* just (as he often said) the *atomic* world, then QM simply has nothing to say about what the world is like: it's just a tool for helping us to predict what we are likely to observe at the everyday level (Leggett's minimal statistical "interpretation")

Besides being irredeemably vague, the minimal statistical interpretation fails to realize the goals of physics: physics becomes merely a blunt tool for helping each of to us get around in a world on which his thought and experience give him essentially no purchase

Einstein stated the goals of physics as follows:

"Physics is an attempt conceptually to grasp reality as it is thought independently of its being observed. In this sense one speaks of 'physical reality' ".

The minimal statistical interpretation is irremediably vague because while it requires a distinction between systems that are treated quantum mechanically by means of a wave-function and other systems that are described classically, it offers no clear way of making that distinction. This vagueness is becoming more urgent as quantum mechanics comes to be applied to the larger scale and more complex systems that are the focus of this workshop.

But any attempt to do better must acknowledge that observation/measurement doesn't *simply* reveal the world to us– it also *affects* the world in significant ways (cf. the neutron interferometer).

Then *how* does measurement affect things?

Any answer to that question must include a description of what properties things have before and after a measurement. How can such a description be extracted from quantum mechanics?

–*Dirac/von Neumann property rule*: a system has a property if and only if the wave-function assigns that property probability 1 (so that the Born rule says it is certain to be found on observation).

Measurement affects a system by discontinuously "collapsing" its wave-function.

This changes its properties in accordance with the DvN property rule.

Born's probabilities must be understood as probabilities *of measurement results*, not of how things are whether or not they are observed.

But exactly when does the wave-function collapse?

A development of the minimal statistical interpretation answers that *it doesn't matter* exactly where this happens along the chain from observed system (e.g. the neutron) to observer (e.g. a physicist becoming aware of the result of

the experiment by viewing a computer screen on which this is displayed).
As Bell put it, just apply QM without collapse to enough of this chain before introducing a measurement collapse so that treating further links quantum mechanically wouldn't make any difference to the predicted probabilities.
This is fine for all practical purposes (FAPP). But it didn't satisfy Bell ("FAPP-trap"): nor did it satisfy Einstein!
Einstein argued that the DvN property rule is wrong, for two reasons.
1) "Spooky" action at-a-distance (for "entangled" systems, a measurement on one system here instantaneously affects the properties of another system there through no intervening force)
2) The measurement problem (Schrodinger's cat): the measuring device typically records no result!
He introduced a character (Physicist *A*) who held the following naively realistic views:
–Systems have a full compliment of properties even when unobserved
–Observation reveals whether or not a system has a property (though not all properties can be observed at once)
–Born's rule gives the relative numbers of systems that *have* various properties, whether or not these are observed.
But, by an extension of one of Einstein's own arguments, Bell showed that if these properties are *not* subject to "Spooky" action at-a-distance, then physicist *A*'s naively realistic views are *false*.
The Born rule gives probabilities *of observing a property*, not of a system's *having that property*.
Quantum probabilities concern measurements.
But then we are forced to face the measurement problem!

Maybe we can use QM to model the effects of measurement/observation without collapse? (Decoherence)
Look again at the neutron interferometer experiment.
If an additional detecting device is present at *B* say, it's interaction with the neutrons can be modeled by QM as an ordinary physical interaction in the usual way.
Just as without the detector at *B*, we had
|*B*>+ |*C*> ✡|*U*+*L*> + |*U*✐*L*> = |*U*>
So now we get instead
|*B*, detected at *B*> + |*C*, not detected at *B*>✡ |*U*+*L*, detected at *B*> + |*U*✐*L*, not detected at *B*>
Note that we still have a superposition of two components: there has been no collapse.
But the neutrons' wave-function has been coupled to that of the device, so that when the two components overlap again at the third "ear" of the crystal, the lower component no longer cancels.
That explains why a neutron is now just as likely to be observed to go down as up, whether or not it was detected at *B*. And it explains it without assuming the detector collapsed the wave-function.
Note that we didn't need to actually observe which path a neutron took: just the physical interaction with the detector at B was sufficient. And this didn't disturb the state of the neutrons at all–it only affected the state of the detector.
So far, so good. But of course this still doesn't explain how a detector ends up in a definite condition, either of detecting a neutron, or not detecting a neutron!
Consider a detector at *U* with a barrier at *B*.
We get |*C*> ✡|*U*✐*L*>✡|*U,* detected> ✐ |L, not detected>.
But the detector is itself strongly interacting with the rest of its environment *E* in a way that couples its state to that of the environment. So we get
|*U,* detected> ✐ |L, not detected> ✡ |*U,* detected, E_d> ✐ |L, not detected, E_n>
And it turns out that this coupling is extremely quick, and the resulting coupled state has the following remarkable property:
If a bunch of systems are all in that state, then the probability of getting any result in a measurement either on the neutrons, or on the detectors, or on both at once, is just the same as if every neutron went up or down and was reliably detected as having done so!
So interaction with the environment (environmental decoherence) quickly makes the state of this bunch experimentally indistinguishable from that of a bunch of neutrons, each of which has gone up or down and been reliably detected as having done so. For all practical purposes a reliable observation has occurred.
But the problem is not practical, but one of consistency. For the Dirac/von Neumann rule still implies that in the final state after interaction with the environment no neutron has gone up or down, and no detector has detected a neutron having gone up or down!

So it seems that not only do we need to understand the Born rules as yielding probabilities *for measurement results*, but also at some point in the chain of interactions between a system on which a measurement is performed and the recording of its result there must be some interaction that induces a real physical collapse. The experiments Professor Leggett discussed in his talk at the start of the workshop may help investigate this issue at least in one context. As he noted, in recent years specific and detailed models of such a collapse have been proposed (GRWP). From one perspective, these models are not so much interpretations of quantum mechanics as rival theories to it. I shall have nothing more to say about them, except to note that at least some such theories face interesting interpretative problems of their own arising from the rules by which they attribute properties to systems--problems quite similar to those arising in interpretations of quantum mechanics that I'll come to later (so-called modal interpretations).

J. S. Bell was one of the most original and influential thinkers about the foundations of quantum mechanics. He once said that *either* the story quantum mechanics tells of a world modeled by continuously changing wave-functions is not right, *or* it is not everything. Those who appeal to decoherence to "solve" the measurement problem try, unsuccessfully in my opinion, to evade Bell's dilemma. Collapse "interpretations" grasp the first horn, and thereby effectively give up on the task of interpreting quantum mechanics. Let's now look more closely at what happens if we grasp the other horn. As you'll see, we will rapidly find ourselves involved in some fundamental metaphysical issues.

Bohm's View (developed recently by Holland, Goldstein etc.)
Perhaps the most straightforward supplement to the QM story is due to David Bohm.
In 1952, Bohm presented a theory which can be regarded as an interpretation of quantum mechanics. On this view, a particle like an electron always has a precise position. Changes in its position are governed by the wave-function of the entire system of which the particle is a part. All measurements are thought of ultimately as measurements of position, and their results are recorded in the position of some part of the apparatus. Each particular result is completely determined by the wave-function together with the prior positions of all the particles concerned. But attempts to find out the relevant positions disturb them and inevitably fail to provide more predictive information than the probabilities specified by the Born rule. So while the world is deterministic, the probabilistic predictions of quantum mechanics are the best we can hope to come up with.

It is striking the extent to which Bohm's view has been ignored or disparaged by physicists. It has been rejected on the basis of a variety of inconclusive and sometimes conflicting arguments.
This constitutes a fascinating episode well worth further investigation by historians and sociologists of science–we might call it the social destruction of reality!

Everettian Interpretations (developed recently by Deutsch, Gell-Mann and Hartle, Coleman)
In 1957 Hugh Everett developed what he called the relative state formulation of quantum mechanics, according to which the wave-function never collapses.
On his view, observation/measurement is a QM interaction internal to a quantum system whereby the state of one part gets correlated to that of another part.
These states are *"relative states"* of the parts.
A measurement does not have one result rather than another: The occurrence of each possible result is represented by the corresponding relative state of the system.
E.g. if the state of neutron and detector is |*U,* detected> ✐ |L, not detected>, the neutron is detected relative to its going up, but not detected relative to its going down.
There are two "branches", or "worlds". In one the neutron is detected and goes up, in the other it is not detected and goes down.
Problems
1) What is a "world", and how do "worlds" diverge?
2) The preferred basis problem: why is it *those* relative states that define the "worlds"
3) What does probability mean?

"World" = space + its physical contents?
"World" = the physical contents of space?
"World" = a mind and/or its mental content?
"World" = how things are recorded in a memory state of some device?

"Everett's" Interpretation come in many Versions(see Barrett, The Quantum Mechanics of Minds and Worlds)
Everett
DeWitt, Deutsch
Bell's Everett(?) Interpretation
Many Minds
Decoherent histories
relative facts

Everettian Metaphysics
the role of probability and the status of determinism
subjective/objective and the mind/body problem
ontological proliferation and its skeptical consequences
fission of observers and Parfit's views on personal identity
Deutsch vs. David Lewis

Is it worth the metaphysical cost?
alleged advangtages:
–save locality (?)
–solve(?) the measurement problem
–preserve the symmetries (and so readily extended to relativistic domain)

Modal Interpretations (van Fraassen, Bub, Kochen, Dieks, Healey)
Everett without the crazy metaphysics!
What additional properties are there, when are they possessed, and how do they change?
Bohm's view as a modal interpretation
Bub's interpretation in *Interpreting the Quantum World* as adding further properties
van Fraassen's interpretation as adding just enough properties to solve the measurement problem in theory, though not, perhaps, in practice: but leading to skeptical problems
Kochen-style interpretations as letting QM pick out the additional properties:
but they face problems
–imperfect measurement: solved by decoherence? (No: as Guido Bacciagaluppi has shown!)
–strange features of property assignments (failure of composition(?), property intersection)
–incompatible with Lorentz invariance: solved by making all properties relational?

<u>*Conclusions*</u>
Where do we stand today?
Can we reach a consensus on how to interpret quantum mechanics just by further careful thinking, or do we need new physics?
So how, if at all, should we allow quantum mechanics to reshape our view of the world and ourselves?

My hunch is that we will need new physics.
In the absence of new physics we may assess the implications of quantum mechanics as it is.
Proponents of rival interpretations have made fervent proclamations, such as Deutsch's proclamation in *The Fabric of Reality* that quantum phenomena themselves prove the existence of many unobservable worlds just as real as this one but containing more or less altered copies of everything and everyone in it. But at the present stage of our understanding there is no sufficient reason to accept any such proclamations. This contrasts with other results of scientific investigation that do and should influence one's world-view, such as the progressive displacement of the position of the earth from the center of the universe since Copernicus, mankind's continuity with the rest of life on earth that we have known about since Darwin, and the recent cosmological finding that we have no idea what 90% of the matter of the universe is made of.
Philosophers, on the other hand have a lot to learn from interpretations of quantum mechanics; not so much about how the world is as about how very different it may turn out to be from the way they had assumed it was. In this way reflection on physics can help to free contemporary metaphysics from its blinkers and lead to a deeper and more probing analysis of such fundamental concepts as causation, probability, realism and observation.