

Some “solutions” of the quantum measurement paradox

It is convenient to classify alleged solutions to the quantum measurement (or realization) “paradox” into two major categories, “exotic” and “non-exotic”, according to whether they do or do not require us to abandon certain rather deeply-held general prejudices about the everyday world, or to invoke considerations from areas, such as psychology, that are not normally regarded as the province of physical theory. What I called in the last lecture the “orthodox” solution clearly belongs to the “non-exotic” category, and in this lecture I will first discuss other solutions in this class.*

A. “Non-exotic” solutions

1. The Copenhagen interpretation

Most variants of the Copenhagen interpretation rest on making a firm distinction between the microworld, which has to be described by QM, and the macroworld of sources, measuring devices, etc., which has (allegedly) to be described in purely classical language. Thus, a really consistent advocate of the CI would probably not even concede the assumption, implicit in our setting-up of the cat paradox, that it is a legitimate demand that we are able to give an account of the behavior of the macroscopic apparatus in QM terms at all. Rather, such an advocate would maintain that a purely classical description of the apparatus is a necessary precondition for setting up the problems that require QM for their solution in the first place.

Obviously, the major problem that this point of view confronts is why, if QM is not only the most fundamental theory of the physical world that we have now, but also (as would almost certainly be maintained by most Copenhagen advocates) the most fundamental we are ever likely to get, should it be illegitimate to ask for an account of the working of a measuring apparatus (or other macroscopic system) in QM terms? After all, the atoms that compose the gas in a proportional counter and contribute to its working are presumably not specially marked “To be used only in measuring devices”, and in any other context it would seem not only possible but natural to describe them in the language of QM! Indeed, in any context other than the measurement paradox, most physicists, be they of the Copenhagen persuasion or not, are only too happy to assume that a consistent application of the principles of QM to the atoms, etc., composing a macrosystem could in principle explain all its properties, even though in practice this is virtually never done (on that, more at the end of this course).

It seems to me that none of the leading advocates of the CI have ever really faced up to this question or given a coherent answer to it, except perhaps to observe that it does not matter in practice where one draws the line (“Schnitt”) between the micro- and macro- levels, as long as one draws it somewhere – a remark that is no doubt true, but in my opinion does nothing to resolve the fundamental issue. It is worth remarking that discussion of this subject tends to be confused by the fact that Bohr did comment

*More precisely, that subclass of solutions that does not attempt to modify the formalism of quantum mechanics at any level. For that which does, see part C.

on a superficially similar, but in reality quite different, question: namely, whether in, for example, a Young’s double-slit experiment, the diffracting screen should, or could, be treated quantum mechanically. (In the case of polarization experiments, a similar question could be raised about the crystals constituting the polarizers – do we, for example, have to describe the epoxy, etc., that binds them to them to the laboratory bench in QM terms?) What Bohr in effect demonstrated in this analysis is that while under certain experimentally unrealistic conditions a QM description of the diffracting screen (or polarizers, etc.) may lead to results different from the “textbook” one, under any conditions that are likely to correspond even approximately to a real-life experiment, the application of QM to this part of the apparatus yields results identical to those obtained by describing it classically (which is what we have implicitly done, and most textbooks do, in obtaining the standard “Young’s slits” or polarization results). But this observation simply does not bite on the real difficulty – it is not the QM description of the polarizers, but that of the detectors, that gives rise to the quantum measurement paradox! (The detector must undergo a macroscopic change of state in order to work at all; the polarizers in general need not and should not.)

2. The “statistical” interpretation

This has in common with the Copenhagen interpretation that it refuses to “interpret” the quantum formalism at the microlevel, but goes beyond it in implicitly denying, under certain conditions, that there even is a classical level at all. The statistical interpretation, as its name implies, emphasizes the point that the quantum formalism should be applied only to ensembles, whereas experimental outcomes are observed on specific systems. In its extreme form (which I personally believe is the only internally consistent one), it takes as basic only the “raw data” – that is, macroscopic events such as clicks in counters, pointer positions, etc.: things that can be and are directly observed by human observers with the natural eye. The whole formalism for quantum mechanics is then taken as, literally, nothing more than a calculation recipe for the prediction of the probability that a system drawn from a particular ensemble will lead to a particular macroscopic event (e.g., a particular position of the pointer set up to read it). Since the formalism is explicitly recognized as nothing more than a recipe, the question of “interpreting” it simply does not arise. In particular, the exotic-looking quantum description of the final states of Schrödinger’s cat, namely

$$\Psi = a|\Psi_1\rangle + b|\Psi_2\rangle$$

where $|\Psi_1\rangle$, $|\Psi_2\rangle$ correspond to macroscopically distinct states, should be regarded as telling us that the probability of observing the macroscopic behavior corresponding to $|\Psi_1\rangle$ ($|\Psi_2\rangle$) is a^2 (b^2) and *nothing more*.

It is interesting to look at the statistical interpretation in the context of Reichenbach’s approach to quantum mechanics (lecture 17). Reichenbach, we recall, took the view that at the macroscopic level, it is possible to give either a “restrictive” or “exhaustive” interpretation of “raw data”. The first corresponds to talking only about the data

themselves, so that the question of the ontological states of the unobserved objects does not arise; the second (obviously much more convenient and universally used in practice!) is essentially our usual classical language, which presupposes, e.g., that the tree continues to exist when I have turned my back on it. Reichenbach then points out that, while at the macroscopic level we have this choice, the experimental data in the atomic realm have shown us that at this level no such choice is available: no “exhaustive” interpretation that is “normal” in his sense exists, and we must therefore either relax the condition of normality, or content ourselves with a restrictive interpretation, i.e., one that talks only in terms of the behavior of macroscopic devices, etc. (which, however, can still itself be described exhaustively). Thus, according to Reichenbach, the point at which an exhaustive (“objective”) interpretation becomes feasible is essentially the borderline between the atomic and macroscopic worlds, and for him, as for Copenhagen advocates, a quantum description is confined to the atomic side of the boundary. The advocates of the statistical interpretation, by contrast, recognize that there is no point at which quantum mechanics can be arbitrarily “turned off” and therefore are forced to push the borderline right back to the raw data themselves. Thus, in the extreme (and only consistent) version of this view, there is simply no level at which an exhaustive interpretation is guaranteed to be possible.[†]

The statistical interpretation seems to me to have one major virtue, namely internal consistency; unlike the Copenhagen interpretation, it does not need to invoke some vague and arbitrarily-defined borderline at which the quantum description simply stops. However, it has two features that, while they certainly cannot be used to refute its logic, should probably give most physicists pause. In the first place, it simply does not permit us to ask, let alone answer, questions like “What was the state of this particular cat before I (or some other human agent) looked in the box and observed it?” It is conceivable that we shall eventually get used to the idea that in a quantum-mechanical world, such questions are meaningless or a waste of time, but this will certainly take us some time and mental agony! Secondly, it embodies an extreme form of empiricism: no meaning is to be attributed to the “probability amplitudes” (wave function) at any intermediate stage, despite the fact that they obviously play an essential role within the formalism in giving rise to the phenomena of interference, etc. This extreme empiricist attitude may be compared to that of Mach toward the atoms in terms of which Maxwell and others interpreted the thermodynamic properties of gases: Mach held that, in effect, the atoms were mere “calculational conveniences” and that no physical reality should be ascribed to them. It is probably as well for the development of physics that this view did not prevail, as, for example, it is unlikely anyone would have thought of doing an experiment of the Stern-Gerlach type with a beam of calculational conveniences! More generally, nature seems to have a habit of showing us that quantities that we have introduced into our theories to satisfy certain formal requirements may actually have unexpected consequences and thereby establish, for most working physicists, their

[†]Of course, this is not to say that there may not be many specific situations in which such an interpretation is possible (roughly speaking, all those situations where application of quantum mechanics predicts a single macroscopic outcome).

“physical reality” (whatever that means): cf., e.g., the case of the electromagnetic field in free space (lecture 9). It is tempting, in fact, to formulate a sort of inverse Einstein principle, to the effect that any quantity that appears in the formal theory, or at least any that plays an essential role therein, must have a counterpart in physical reality. (We would, of course, all agree that the quantities like, say, the partition function of statistical mechanics do not correspond to anything obvious in the physical world; but then the partition function is a derived, rather than a primary, element in the theory.) The statistical interpretation would appear to violate this inverse Einstein principle, since the probability amplitudes, which are generally held to be central to the theory, are held to possess no counterpart at all in the physical world.

2a. Consistent histories

As we have just seen, an obvious difficulty with the statistical interpretation is that it is at first sight problematic to talk about the behavior of objects, even macroscopic ones, when they are not subject to direct observation. In particular, it seems problematic to assign *histories* to such objects (“did the electron come through slit 1 or slit 2?”). A line of argument that has seen much development in the last few years is to the effect that because of the phenomenon of “decoherence” (see stage I of the “orthodox” solution, last lecture), the construction and consideration of such “histories” is in fact legitimate, in the sense that it will not normally lead to conclusions that contradict experiment. (Roughly speaking, whereas the assumption that the electron either went through slit 1 or slit 2 may lead to experimentally incorrect predictions, the hypothesis that the counter either clicked or did not click before we looked at it does not, because these two states, being macroscopically distinct, cannot in practice mutually interfere [“decoherence”].) With this elaboration, the statistical interpretation begins to look, formally, very like the “orthodox” one: however, it differs from it in emphasizing explicitly that even at the macroscopic level the quantum-mechanical wave function is not a description of anything at all in the real world.

3. “Solutions” based on irreversibility

It is sometimes claimed that the solution to the measurement paradox is to be sought in the fact that QM is meant to describe only isolated systems, and that the moment the system is put into contact with a complex environment in such a way that it can interact irreversibly with it, the inferences we have been drawing no longer apply. However, it is important to distinguish two different meanings that can be attached to this claim.

The weaker sense of the claim, which is not controversial, is a technical point about the structure of QM itself, namely that, where a system is interacting (or has interacted in the past; cf. the EPR thought experiment) with any other quantum system (“environment”), in general it will not be possible to assign it a quantum state in its own right – all we can do is to assign a state to the “universe” composed of the system plus its environment. This remark is an important ingredient in the claim made at stage I of the “orthodox” solution that “decoherence” will in practice prevent the observation of in-

terference between macroscopically distinct states, and hence is relevant to the paradox insofar, and only insofar, as one accepts the general logic of that solution.

The stronger sense of the claim is that irreversibility is in some sense a qualitatively new ingredient in the problem, and that therefore, its treatment requires us to go outside the usual laws (Newton, Schrödinger. . .) that we apply to isolated objects. If that is so, the whole idea of applying standard QM to (for example) the interaction of a microscopic system with a macroscopic measuring device, which as we have seen normally involves a high degree of irreversibility, might be inappropriate – and thus the whole basis for the measurement paradox might vanish.

It seems to me that this point of view is simply unjustifiably pessimistic. There seems no compelling a priori reason to doubt that the “irreversibility” of which such a mystery is made is nothing other than the transfer of energy (and the associated information) from one or few macroscopic degrees of freedom to a myriad of microscopic (and not directly observed) ones in such a way that it is in practice unrecoverable (cf. the last part of the course); and that the processes involved can in principle be perfectly well described by Newton’s laws (for a classical system) or by Schrödinger’s equation (for a QM one). In fact, by now there exist a number of calculations which quite straightforwardly apply standard QM to the coupled system–environment complex and obtain the characteristic phenomena associated with irreversibility. Thus, there seems no good a priori reason to treat irreversibility as something fundamentally mysterious; and while no doubt future research *might* some day show that there is more to the question than meets the eye, the principle of Occam’s razor would suggest that we should not assume this ahead of time.

4. The Bohm–de Broglie “pilot wave” interpretation (theory)

We already encountered the ideas of Bohm briefly in the context of hidden-variable theories. Depending on the auxiliary assumptions made, the so-called “pilot-wave” scenario can be either an “interpretation” of QM or an alternative theory (i.e., something that under certain circumstances predicts different experimental results). In the recent (post-1990) revival of interest in Bohm’s ideas, most work has tended to favor the former point of view, so it seems appropriate to discuss it at this point; in any case, considerations concerning the measurement paradox are apparently insensitive to this choice (at least for the versions of the theory that have been explicitly considered so far).

Following earlier work of de Broglie, Bohm associates with a given microscopic entity such as an electron both a physically real “particle” and a physically real “wave”. The latter is simply the quantum-mechanical wave function $\Psi(x, t)$, which is not now interpreted (primarily) as a probability amplitude (see below); as to the particle, it has at all times a real position $x(t)$ whose dynamics are governed by ordinary classical (Newtonian) mechanics with one major difference: the potential $V(x, t)$ in which it moves contains, apart from all the usual classical terms, also a “quantum potential” governed

by the accompanying wave, of the form

$$-\frac{\hbar^2}{2m} \frac{1}{\Psi(x,t)} \frac{d^2\Psi}{dx^2}$$

In the original (one-particle) Bohm theory, the relation

$$\text{probability of finding particle near } x \propto \Psi^2(x)$$

is not an identity as it is in the standard interpretation of QM, and in principle need not be true; however, remarkably, Bohm was able to show that if this condition is initially fulfilled and no measurement is made, then the statement remains true subsequently.[‡] In earlier versions of the scenario, the possibility was considered that this relation might not always be true, e.g., immediately subsequent to a measurement or filtering process (cf. lecture 19); however, more recent versions, including some that attempt to apply the scenario to the Universe as a whole, have tended to try to guarantee that it will be valid always and everywhere, thereby guaranteeing agreement with the predictions of QM under all possible circumstances and converting the scenario from a “theory” into an “interpretation”. One might, by the way, ask how the Bohm theory evades the consequences of Bell’s theorem. The answer is straightforward: the theory turns out to be intrinsically *nonlocal*, and in particular the “quantum potential” felt by one particle may depend on the instantaneous state of a distant partner (but, of course, in a way that is guaranteed not to permit superluminal signaling, just as standard QM does not).

How does the Bohm interpretation deal with Schrödinger’s cat? There does not seem to have been a great deal of discussion of this question in the literature, but presumably it would say that in the absence of any “measurement”, the following statement is true for any particular cat: (a) she is *either* alive *or* dead, not in a quantum superposition, and yet (b) the corresponding “pilot wave” (obviously at this level, an entity in a highly abstract space!) has finite and possibly large components corresponding to *both* possibilities, alive and dead. The conceptual problem then clearly relates to the ontological status of that component that does *not* correspond to the “real state of affairs”: we know in particular that if the experimental results predicted are to be identical to those given by QM, then as soon as this particular cat is observed to be (say) alive, the part of the wave function corresponding to “cat dead” must either vanish or at least be guaranteed to leave no trace in any subsequent measurement. (This is because in QM, any cats found to be alive must immediately be reassigned to a new subensemble characterized entirely by Ψ_{alive} with no Ψ_{dead} component; and any subsequent predictions must be made using this state.) This seems, at least, a counterintuitive state of affairs. . .

[‡]It should be strongly emphasized, however, that the position variable is privileged in this respect – no corresponding statement can be made about probabilities for momentum, etc. For this reason, most adherents of the Bohm approach tend to consider the position variable as logically primary.

B. “Exotic” solutions

5. Schulman’s “fatalistic” solution

An intriguing, if decidedly exotic, solution to the measurement problem has been proposed by Schulman. He points out that since the development of the (universe) wave function in QM is strictly deterministic, then we can in principle start from a final state of the universe (system - plus - apparatus - plus - environment...) that corresponds to a definite macrostate of the apparatus and work back to the unique initial state[§] that would have produced it. In general, these states will, of course, be horribly complicated ones involving delicate superpositions of different (microscopic) states of the apparatus correlated with those of the environment; however, in principle they are guaranteed to exist. Schulman then postulates that in the real world, only those special initial states actually occur; from which it follows that every time a measurement is carried out, the final state of the apparatus as predicted by QM is a definite macrostate and a “Schrödinger’s cat”-type final state never arises.

At first sight this solution might seem technically invalid, in that we know that the measuring apparatus, to play its assigned role, must end up in the macroscopic state k *if and only if* the microsystem started in the corresponding state Ψ_k with the appropriate value a_k of A . Thus, at first sight, only such states are allowed for the initial state of the microsystem, whereas the whole point of the paradox is our ability to prepare this microsystem (or an ensemble thereof) in a linear superposition of the form $\sum_k c_k \Psi_k$. Schulman’s response to this difficulty is to postulate that whenever we indeed start with such a superposition, the initial state of the environment is exactly such as to convert it, by straight application of the laws of QM, into the appropriate state Ψ_k before the interaction with the measuring device even starts.

One immediately sees one feature of this “solution” that is very peculiar indeed: since for any *given* initial state of the environment the superposition principle still applies, it follows that the experimenter can never have the choice of preparing his ensemble in three different states Ψ_i , Ψ_j and $2^{-1/2}(\Psi_i + \Psi_j)$, where Ψ_i and Ψ_j correspond to different values of A . For, if Ψ_i and Ψ_j are guaranteed (starting from the particular initial state of the environment that is realized) to lead to a definite macrostate of the apparatus, then by the principle of superposition, the third state $2^{-1/2}(\Psi_i + \Psi_j)$ cannot![¶] However, this objection assumes that the “decision” as to which of the three states to prepare is effectively instantaneous. If we assume, more realistically, that it takes a finite time for the experimenter to decide which states to prepare, and to set up the necessary preparation device, then the possibility arises that the environment could “sense” this and adjust its initial state accordingly. It is clear that the issues here are somewhat related to those encountered in the context of Bell’s theorem, and we will return to them below.

[§]Or more precisely, a unique set of states, since states of the apparatus that are microscopically different, and hence not distinguished by us, will “retrodict” different initial states.

[¶]And conversely, if Ψ_j and $2^{-1/2}(\Psi_i + \Psi_j)$ both lead to definite (different) macrostates, then Ψ_i cannot, etc.

6. The “mentalistic” solution

A proposal that was originally contemplated by von Neumann, advocated explicitly by London and Bauer, and embraced, with varying degrees of tongue in cheek, by Wigner^{||} and a number of other distinguished physicists, attempts to involve human consciousness in the solution of the measurement problem. The idea is essentially that so long as no human agent is involved, the linear laws of QM apply and in particular the final state of the “universe” is a macroscopic superposition of the “Schrödinger’s cat” type. However, according to this view, human consciousness lies beyond the laws of physics as usually understood, and therefore, the linear laws of QM no longer apply; observation by a human agent is what “reduces” or “collapses” the wave function, so that a specific macroscopic outcome is always observed.

It is difficult to say much that is useful about this proposal, since it essentially relies, for the solution of something we understand badly (the quantum measurement paradox), on something we understand even worse (the phenomenon of consciousness). One might perhaps merely quote the remark of Margenau that it would be nice if, before calling on the discipline of psychology to solve its problems for it, physics were to display a little more competence in the reverse direction!

If the absurdity is justified... by the observation that processes in the human brain do matter and that ego is to be introduced into the scheme somehow, the only significant reply is... QM does not yet pretend to be a psychological theory. As such it would have to show a little more competence in the purely psychological realm.**

At any rate, it is clear that the mentalistic solution affords no opportunity for an experimental test in the near or probably even the remote future, since any such test, to be meaningful, would presumably have to involve the possibility of observing the lack of interference between different macroscopic states of a human observer under the condition that QM, extrapolated to cover the human being, etc., predicts such interference – a condition many orders of magnitude more stringent than demanded by even the most speculative tests contemplated today.

7. The “many-worlds” interpretation

This interpretation (also known, less provocatively, as the “relative state” or “Everett-Wheeler” interpretation) has the dubious distinction of being probably the most exotic solution of the measurement paradox conceived to date. The technical underpinning of the interpretation is a series of formal theorems of quantum measurement theory; in reading the literature on this subject, one should beware of the tendency of the defenders of MWI to switch without comment between discussions of these formal theorems, which are noncontroversial, and the metaphysical interpretation of them, which is controversial in the extreme.

^{||}Who however seems to have abandoned it before the end of his life.

^{**}Quoted by M. Jammer, *The Philosophy of Quantum Mechanics*, p. 487.

In its simplest and most extreme form, the MWI holds that our belief that we obtain definite results of our measurements, or more generally that the Universe is at any given time in a definite macroscopic state, is simply an illusion: the real state of the Universe is that which follows from a consistent application of the formalism of quantum mechanics without “collapse”, that is, it is a superposition of different macroscopically distinct states. The more pictorial versions describe the Universe as “splitting” into many different worlds, each of which is said to be “equally real”; when challenged as to why we are not conscious of this multiplicity, the advocates of the MWI reply that this is to be expected, since not only the correlation between my own observations at different times but also the correlation between my experiences and yours are guaranteed by the QM formalism itself to be such that (e.g.) the probability of my “observing” world I today and you “observing” world II tomorrow is guaranteed by the formal theorems to be zero.

I find it difficult to take the MWI seriously. To quote Ballentine: “Rather than deny [that] a state vector can be a complete model of the real world, Everett and de Witt choose to redefine the ‘real world’ so that a state vector [of the macro-superposition type] can be a model of it.” What *meaning* can one attach to the (ostensibly English) sentence: “The non-observed worlds are as real as the observed ones?”

C. Attempts to modify QM

All of the above alleged “solutions” of the quantum measurement paradox have accepted the assumption, implicit in its formulation, that the linear formalism of QM does in fact apply to the whole of the physical world. However, this assumption is not self-evident, and in recent years there has been increasing interest in the possibility of modifying the formalism in such a way as to resolve the measurement paradox; of such attempts, by far the most extensively developed is the model of GRWP (Ghirardi, Rimini, Weber and Pearle), which I will briefly describe.

Any theory of this type needs to satisfy, at least, the following constraints:

1. At the *microscopic* level it should reproduce, either exactly or with extremely small error, the experimental predictions of QM (since these have been very widely confirmed). In particular, it should reproduce the characteristic QM interference effects, e.g., in the Young’s slits experiment.
2. At the *macroscopic* level, by contrast, it should effect the “realization” of one alternative or another within a time that is, at a minimum, smaller than the minimum time needed for conscious human observation (so as to explain why we always “see” the pointer in a definite position, etc.).
3. Its statistical predictions for the frequencies of observed outcomes when measurement of the (general) physical quantity A is made on the ensemble in the (general) state Ψ should correspond with those of quantum mechanics. I.e., if the expansion

of Ψ is

$$\Psi = \sum_k c_k \Psi_k$$

then the probability of observing the value a_k of A should be c_k^2 .

4. Further, most people would regard it as highly desirable (though not perhaps essential) that any such modification should continue to respect our rather deeply-held general beliefs about the physical world, such as, e.g., that total energy is conserved, or that superluminal communication is impossible.

The GRWP theory fulfills criteria 1-3; it fails in some ways to satisfy criterion 4 in its present form, and it is unclear whether future extensions may be able to overcome this problem. In essence, GRWP start with the standard QM description and add to it a new effect: they postulate that there exists a universal “noise” background, not itself described by QM, whose effect is to localize the position of a body – i.e., to reduce the QM uncertainty in position. (As for Bohm, in the [present version of] the GRWP theory, the position variable has a privileged role.) It emerges naturally from the basic formalism (and does not have to be put in by hand) that the “realization (localization) time” τ_{loc} is inversely-proportional to the total number of (cooperating) particles N in the body; in fact, it can be written $\tau_{\text{loc}} = \tau_0/N$, where τ_0 is universal characteristic time that is an input to the theory. (The theory also has a characteristic length scale a [the scale, roughly speaking, on which “localization” is achieved], which in the present version is usually taken to be $\approx 10^{-5}$ cm.) Thus, if, for example, we choose $\tau_0 \approx 10^{16}$ secs ($\approx 10^8$ years), then we find that for photons ($N = 1$) arriving from all but the most distant quasars, any original quantum-mechanical superposition of different “positions” should be preserved (so that we expect to see the usual effects in Michelson interferometry, etc.); while for Schrödinger’s cat or a typical counter ($N \sim 10^{23}$), the “realization” of a definite macrostate takes place within about $\sim 10^{-7}$ secs – certainly fast enough to avoid any paradoxical consequences. From the details of its construction, the GRWP theory also satisfies criterion 3 (the selection of one out of the possible alternatives is stochastic in nature and of such a kind that the statistical predictions of QM are obeyed).