NOTES ON "SAVING SCIENCE FROM QUANTUM MECHANICS", ADLAM 2025

Jeffrey M Epstein, September 2025

CHAPTER 1 (PROLOGUE)

Adlam groups all current approaches to the measurement problem into two groups. (1) Those that "potentially undermine the methods of enquiry we have used to arrive at quantum mechanics as well as many of our other scientific beliefs". These are epistemically problematic. (2) The non-unitary approaches, which do not pose the same kinds of epistemological problems, but which seem unlikely to be extendable to QFT, and so are likely empirically inadequate. Thus she argues that at present we have no viable approach to the measurement problem. An overarching theme of the book will be to "gain a deeper understanding of the epistemology of science" by analyzing the measurement problem, while also using epistemology to "help us understand what a successful solution to the measurement problem would really look like".

CHAPTER 2 (WHAT IS THE MEASUREMENT PROBLEM?)

Here Adlam begins to make the case that we ought to view the measurement problem through an epistemological lens. As she says, "measurement...is nothing more or less than the ultimate contact point between our experience and physical reality" and thus (citing van Fraassen) "has both 'physical and intentional aspects'". As a result, problems in how we understand measurement lead to problems in how we understand the process by which we arrive at scientific beliefs in the first place. This is not only an issue in quantum mechanics, and indeed Mach and many other pre-quantum physicists and philosophers discussed the "coordination problem", i.e. the apparent chicken-and-egg nature of theory development. Demonstrating that a theory is empirically adequate requires trust in the reliability of measurement procedures, which trust in turn requires a belief in the adequacy of the theory (at least in the setting of fundamental physics). How then does science get off the ground? (Here I'll get a bit Continental with it, and just vaguely say that I want to think later about how Gadamer's arguments in Truth and Method about the methodological independence of the human sciences might be fruitfully juxtaposed here, perhaps in a way that undermines his possibly naive view of natural scientific method. Anyway, just a note to self.)

In any case, what we really need is not a methodological way to choose an immutable starting point for theory development, but rather a criterion of coherence and stability - a theory of fundamental physics ought to be capable (at least in principle) of explaining what measurements are, and why they provide reliable information about the structure of the theory. Quantum mechanics introduces the further problem that the intuitive notions we tend to have about measurement (that they reveal pre-existing values, that they have single outcomes, that records and measurements give accurate information about past measurement results, etc.) are likely unable all to survive in an empirically adequate interpretation (e.g. see the Kochen-Spekker theorem), though different interpretations preserve different subsets of these intuitions. Thus it becomes difficult to tell a satisfactory story about the epistemology of science. This becomes a statement of purpose: "this is the lens through which we will consider the measurement problem in this book - the central challenge is to find an empirically adequate solution which is also capable of accounting in a coherent, reasonable way for the epistemology of science".

Adlam also makes the point that the measurement problem, despite her framing in terms of apparently quite abstract philosophical concerns, is indeed a truly scientific problem that should be of concern to people who are concerned with such things but may turn up their noses at epistemology. To do so, she invokes the family of Extended Wigner's Friend thought experiments. In her view, these improve over the original Wigner's Friend thought experiment by demonstrating that different solutions to the measurement problem (here interpreted as where to place the Heisenberg cut) result in different empirical predictions, or "empirical scientific knowledge". In Wigner's original proposal, the distinctions are on the level of "theoretical scientific knowledge" (concerning unobservable parts of reality, in this case perhaps propositions about whether or not the Friend has observed an outcome at all before communicating with Wigner). Because these experiments seem possible in principle, many would agree that there ought to be a fact of the matter about what would happen if they were performed, and thus the measurement problem is not merely a "matter of idle intellectual curiosity. It is a concrete scientific question" as much as (she does not explicitly make this point) any question about quantum gravity or standard particle physics at energies far beyond what we will plausibly achieve in our lifetimes.

I think a fruitful way of precising/formalizing this very interesting discussion would be something along the lines of the formalism I use in (sorry) my paper on the Frauchiger-Renner thought experiment. In other words, perhaps we can view interpretations of quantum mechanics as methods for assigning probabilities of measurement outcomes to any one agent participating in a protocol defined in terms of a syntax of directed graphs of interacting systems, some of which are labeled as agents, and some of whose interactions are labeled as measurements. In particular, different interpretations differ in the extra structure that they may assign to handle these special interactions, and thus may in cases such as Extended Wigner's Friend settings make different predictions. The measurement problem is then to find a method for assigning these probabilities that (1) agrees with the empirical data we have for the subset of relatively simple protocols which have been actually executed, and (2) allows us to preserve enough of our pre-quantum intuitions about measurement so that scientific epistemology may be salvaged.

CHAPTER 3 (SOME SOLUTIONS)

Adlam turns to some specific proposals, or families of proposals, to the measurement problem, focusing on unitary-only approaches. These are those strategies for addressing the measurement problem in which unitary quantum mechanics (i.e. quantum mechanics without collapse) is taken to be complete and universal. Completeness is the condition that all physical properties are reflected in the wavefunction (so no hidden variables or other extra structure may be appended to the theory to explain these properties), while universality is the condition that all isolated systems (or rather, their wavefunctions) evolve unitarily, in any special relativistic inertial frame. The goal of this chapter is to introduce these approaches and show that they all present problems for scientific epistemology.

First, the Everett interpretation. Widely called the "many-worlds" interpretation, this is either an wildly extravagant or an admirably parsimonious approach, depending on your ontological sensibilities, and so considerations of theoretical simplicity don't seem to provide much objective guidance. On the other hand, epistemologically speaking, concrete difficulties appear. In order for an observer to believe that her observations provide empirical confirmation of quantum mechanics, she must hold the Everettian Statistical Assumption "that the relative frequencies she herself observes in experiments are, after a sufficiently large number of experiments, very close to the mod-squared-amplitudes for the corresponding outcomes". But in the Everett interpretation, all sequences of outcomes (or at least those with non-zero amplitude) do in fact occur, so how can an observer be confident that she is in one of the branches where the ESA holds? Thus without justification of the ESA, the Everett interpretation is "empirically incoherent" - it provides a picture of the world in which there is no reason to think that observations provide support for the theory. In a later chapter, options for justifying the ESA will be explored.

Second, the observer-relative interpretations. These are characterized by three postulates: (1) unitary QM is complete and universal, (2) all measurements have unique outcomes with probabilities given by the Born rule, and (3) "all quantum mechanical descriptions are relative to individual observers". To be honest, I don't understand why (1) and (2) are compatible. This seems to require that we don't take the unique measurement outcomes to be genuine physical properties as invoked in the completeness condition. I think postulate (3) is supposed to save us from this difficulty, perhaps by motivating this very non-identification of outcomes as physical properties. Later chapters go into more depth on observer-relative interpretations and on first- and third-person views of science, so I'm hoping to get more clarity on this.

In any case, the epistemological problem that arises here can be illustrated by considering a Wigner's Friend scenario. Suppose that a qubit is prepared in a plus state, and measured in the Z basis by the Friend, and then Wigner asks the Friend what he observed (this too is a measurement). Now by postulate (2), there is a fact of the matter about what the Friend observed, and this is a random variable. By postulate (1), after the Friend's measurement he and the qubit are in an entangled state, and so by postulate (2) Wigner's observation, which he takes to be a record of his Friend's observation, is also a random variable uncorrelated with the actual Friend's observation. We can also take the Friend to be in fact just a past version of Wigner himself, and so "we cannot expect to obtain any veridical information from reports obtained from other observers or past versions of ourselves...making it impossible to perform any empirical confirmation" and leading, again, to empirical incoherence. Put another way, "if records of the past don't reflect what has really happened, then the theory we confirm won't be a theory of measurement results as they actually occur: it will simply be a description of the contents of records at some particular point in time". I guess the crux here is that quantum mechanics is a theory of evolution in time, with these interpretations sitting on top. So if the interpretations tell us that in fact our records have nothing to do with the actual past, or don't yet give us a way to make this correspondence evident, then there is no reason to take our records to have the potential to validate the theory.

Finally, Adlam introduces the consistent histories family of approaches, which suffer from one or the other of the problems created by the previous two, either the problem of motivating the ESA or the unreliability of records, depending on whether or not the approach provides a way to select a single consistent set and a single history within this set.

In light of the fact that the unitary-only approaches seem at present to be the only ones capable of reproducing the experimental evidence for quantum mechanics in both relativistic and non-relativistic settings, we find ourselves in the untenable position that "if we take seriously what quantum mechanics is telling us about reality, right now there is a sense in which our empirical evidence seems to be telling us that empirical evidence cannot in general be relied upon!" Adlam reiterates her point from the previous chapter that the measurement problem is a bona fide scientific problem, and the book will argue that confronting the epistemic version of this problem can be a useful way of making progress in future research on quantum gravity, cosmology, and other related fields.

CHAPTER 4 (THE EPISTEMOLOGY OF MEASUREMENT)

A common, and naive, perspective on empirical confirmation is that it can be understood as "a direct relation between a mathematical hypothesis and a set of empirical experiences which directly confirm the hypothesis". From this (mistaken) point of view, Adlam says, all unitary-only approaches to the measurement problem seem have the same degree of empirical confirmation, as in all cases the mathematical models are identical. But this simple conception of scientific method, associated for example with the logical positivists, has long been discredited: "confirmation is not and cannot be a straightforward map from hypothesis to evidence: it is really a three-place relation, in which a hypothesis is confirmed by some evidence relative to our other background beliefs". In principle, this implies a "radical holism", which obliges us always to evaluate the confirmation of our entire set of beliefs all at once. In practice, we can often make a clean division between a set of background beliefs, underlying assumptions of our modeling and hypothesizing activities that are not up for reevaluation, and a set of beliefs we are open to revising or rejecting in light of empirical data. But this fails if empirical data leads us to form hypotheses that contradict these background beliefs, as in the case of quantum mechanics. So for example, "prior to the advent of quantum mechanics it was generally accepted that we are located in a single world, performing experiments that produce a single outcome". In the Everettian view, this is likely incorrect. We are free to reject this background belief, but in doing so we may lose the power of our empirical observations to confirm the mathematical hypothesis of unitary quantum mechanics. In order to have confidence in our ability to use observations to confirm quantum mechanics, we need "a measurement model and schematisation of the observer". I think I would actually argue that this model and schematisation ought ideally to be mathematical hypotheses in their own right (sets of observers, schemes for selecting consistent sets of histories, etc.) and so in fact adequately-formalized unitary-only approaches to the measurement problem ought to differ in terms of the degree of empirical confirmation they take empirical observations to confer to the theory, which is ideally a theory not just of the wavefunction, but also of the observers and their experiences (whatever these turn out to be). Again, I'm hoping for much more on this in the later chapter on first- and third-person views of science.

CHAPTER 5 (STRATEGY)

Adlam discusses some of the other framings of the measurement problem and desiderata for proposed solutions that appear in the discourse.

(1) Many discussions of the measurement problem reflect a desire to maintain classical intuitions, including but not limited to those about measurement. There are two problems here. First, such discussions often come down to a matter of taste about which of our intuitions are most sacrosanct. Second, even if an approach to the measurement problem restores classical intuitions at the microscopic level, it must be able

to simulate the violation of these intuitions at what perhaps we can now think of as a mesoscopic level, and typically will need to be supplied with constraints that prevent the appearance of these apparent violations at the macroscopic level. Thus for example, while there are approaches that restore locality at the microscopic level, macroscopic locality may not really be good evidence for supposing this feature, and such an approach will need to explain how quantum nonlocality can be simulated at the mesoscopic level.

(2) Compatibility with special or general relativity is often touted as a crucial feature of an approach, but what this actually means is up for debate. Some take it to mean compatibility at the empirical level, others to mean compatibility at the level of mathematical modeling, for example forbidding preferred inertial frames. Adlam agrees that empirical compatibility is required.

(3) Claims are sometimes made that the measurement problem refers to the issue of providing an ontology for quantum mechanics, i.e. clarifying what exactly it is to which the wavefunction refers. An adequate ontology would clearly be useful, but is not necessary to have a working account of scientific epistemology.

(4) The measurement problem is sometimes framed as the problem of figuring out how to excise from our fundamental physical theories the anthropocentric and imprecise notions "observer" and "measurement". The epistemic view that Adlam advocates clarifies the truly problematic aspects of involving such concepts in the basic structure of the theory.

Adlam notes that many discussions of the measurement problem (and of quantum mechanics more generally) center on various sorts of "weirdness". She divides these into two sorts. "Benign weirdness" is conflict with classical intuitions that does not lead to empirical incoherence. These are thus epistemically rationally acceptable. As I like to say about features of quantum mechanics, not liking them is a personal problem. Then there is "epistemic weirdness". These are features of approaches to the measurement problem that do lead to empirical incoherence, thus disrupting the process by which quantum mechanics could be confirmed in the first place, and undermining themselves in the process. Such features are unacceptable.

This distinction is neither totally precise nor necessarily obvious in any particular case. For instance, Einstein saw nonlocality as epistemic weirdness, as it violates classical assumptions about mereology and measurement. Nevertheless, the restricted form of nonlocality allowed within quantum mechanics does not in fact cause problems for scientific epistemology, so this feature is merely benignly weird.

CHAPTER 6 (THE EVERETT INTERPRETATION)

From here on the book becomes more technical, and my already unwieldy thread will become a sparser sampling of topics.

This chapter dives deeper into the epistemic issues arising in the Everett interpretation. To begin, Adlam considers the question of whether observations can provide any useful information at all in the service of empirical confirmation. The problem is that the information gained upon making an observation seems to be of a particular kind known as self-locating information. This is "information which tells you about where or when you are located, but does not tell you anything new about reality as a whole". (Thus it is information that locates the self who knows it, not information that locates itself.) Indeed, in the Everettian picture, all possible observations do occur, so finding yourself in one branch or another tells you nothing about the world as a whole. The argument now goes as follows: (1) The information an observer gains upon making an observation is purely self-locating. (2) Self-locating beliefs cannot rationally be used to update non-self-locating beliefs. (3) The belief that quantum mechanics accurately describes the world is non-self-locating. (4) Thus observations cannot be used to perform empirical confirmation in a world governed by Everettian quantum mechanics, which is therefore empirically incoherent. Point (2), also known as the "Relevance-Limiting Thesis" is supported by many philosophical arguments, and I suppose that in various formal epistemic logics it may be proven, although Adlam does not discuss such formal deductive systems.

Even if you accept that you can learn something relevant to empirical confirmation by making observations within an Everettian world, the question remains why you should update your beliefs in a way that corresponds to the Born rule. This is related to the Everettian Statistical Assumption introduced in Chapter 3. The discussion is quite extensive, but I'll just say that Adlam divides the arguments that have been made on this issue into three categories: ones based on decision theory, ones based directly on formal accounts of epistemic rationality, and ones that turn on some notion of branch-counting or weighting. All of these suffer from problems of various sorts (of course discussed at length in the book), but I'm more interested at the moment in the discussion of self-locating beliefs. Adlam closes by acknowledging that the Everettian picture, including “an assumption like ‘you should believe you are in a branch with high mod-squared amplitude’” indeed solves the measurement problem, if (and it’s a big if) you’re willing to accept a fundamental theory featuring “ineliminable references to vague, agent-centric notions”.

To add a bit to the discussion, I think it’s interesting to compare to the (deterministic) classical setting and ask why the same problem of empirical incoherence does not appear there. Consider asking the question “Is Newtonian mechanics empirically coherent?” Clearly in order to make sense of the question, we must think about observers making observations, which is to say having experiences. In the classical setting, we typically make the assumption that experiences correspond in some deterministic way to particular configurations of matter. Thus we can tell a story which is at least not clearly incoherent about how it is that in a world governed by Newtonian mechanics, observers could end up devising and empirically confirming Newtonian mechanics. But this relies on an implicit schematization of the observer! So the problem in the quantum setting is just that we do not have this schematization, and the tacking on the ESA or a similar assumption is an (unsatisfying) way to provide one.

Personally, I think that this sort of discussion shows that there’s a lot to be said about the interaction between the measurement problem and the problems of consciousness. For example, although Adlam doesn’t mention this, the classical example she gives to illustrate the notion of self-locating beliefs (involving several identical clones with differently colored hair) comes right up against what is known as the vertiginous problem in philosophy of mind. Moreover, there seems to be a lot to be said about the connection between self-locating and phenomenal knowledge, as both are intrinsically first-person rather than third-person. While it’s of course dangerous to say too much about quantum mechanics and consciousness, there is a long tradition of facing up to the connection between these twin mysteries. I wrote informally about this in the blog post in my pinned tweet (sorry again).

CHAPTER 7 (NO WORSE OFF)

Here Adlam responds to a family of potential objections to her claim that the quantum mechanical measurement problem, as it appears in the Everettian approach, is a particularly pressing issue in the epistemology of science. These are objections that point to epistemic issues in other theories to argue that Everettian quantum mechanics is at least “no worse off”. For me, the most interesting case is that of theories featuring objective chance. In a fundamentally probabilistic theory it seems that one has to rely on the Probabilistic Statistical Assumption, i.e. the assumption that you are observing a “typical” sequence of observations. Since this has the same basic form as the Everettian Statistical Assumption, Everettians might claim that the ESA is equally ad hoc, so if you allow that probabilistic theories can be empirically confirmed, Everettian quantum mechanics is no worse off, and you ought to accept the Everettian approach as a solution to the measurement problem.

In response, Adlam points out that when faced with a probabilistic theory, it makes sense to generalize the notion of epistemic rationality, so that we require only that the standard methods of empirical enquiry lead us to forming true beliefs with high probability, rather than with certainty. Here the so-called Principal Principle, due to Lewis, can be invoked: “the (epistemically) rational credence to assign to the occurrence of event C, conditional on the knowledge that the objective chance of C is P(C), is simply P(C)”. This can be viewed as essentially a definition of objective chance. I think a way to make sense of this is to return to the notion of “schematizing the observer”. In a (classical) probabilistic theory, all physical facts are the same for all observers, so that we can just identify the probability of an event happening for a given observer with the objective probability of that event happening. Thus for a given observer, the PSA is true with high probability.

The question then arises whether an Everettian can simply posit a Quantum Principal Principle, “which might say that the rational credence to assign to the occurrence of event C, conditional on the knowledge that the mod-squared amplitude associated with C is A(C), is A(C)”. But the issue here is that the amplitude already plays a dynamical role in quantum theory, and so we aren’t free to simply define it in this way, as we were in the case of objective chances.

The chapter also discusses Boltzmann brains and inflationary cosmology, and asks whether the Past Hypothesis or landscapes of island universes pose problems for scientific epistemology at the same level of severity as the measurement problem, concluding that they do not.

CHAPTER 8 (BAYESIANISM)

In this chapter, Adlam argues that two approaches to the measurement problem, the Everett interpretation and QBism, rely on a Bayesian conception of scientific confirmation, which she argues does not work to justify their empirical coherence. Thus even if it were possible to establish the validity of Bayesian approaches to the updating of credences (which she argued previously has not been sufficiently demonstrated), that would be insufficient to solve the epistemic version of the measurement problem present in these approaches. The main point seems to be that Bayesianism is naturally suited to the very narrow setting in which background assumptions do not change, so that it is reasonable to assume a relatively stable and well-defined set of conditional probabilities. In such a setting, convergence theorems guarantee that a Bayesian agent's particular prior is not too significant. But this guarantee doesn't help us understand situations in which those background assumptions, such as the notion that we live in a single world with definite measurement outcomes, do change in light of empirical evidence. In such cases, it is often not clear that there is any reasonable way to assign priors at all, or that such assignments would be meaningful (what would it mean to say "The likelihood that I am a brain in a vat is .02?).

To be honest, I don't find this chapter especially convincing. It seems clearly true that the process of theory development and confirmation is complex and canot be fully captured by the simple mechanistic procedures of Bayesian updating. Thus Adlam is obviously right that simply demonstrating the validity of Bayesian updating within Everettian quantum mechanics is insufficient for demonstrating the empirical coherence of that approach. However, the problems that she identifies with attempting to formalize scientific progress within a Bayesian setting don't seem to me to be particularly tied to quantum mechanics. For example, she discusses the problem of using past evidence to confirm a new theory, and provides as an illustration the confirmation of general relativity by the perihelion of Mercury, a piece of observational knowledge already known at the time GR was developed. It is unclear, Adlam points out, how one would go about "subtracting" that particular piece of evidence from one's beliefs in order to obtain a prior over a set of other beliefs that could then be used in a Bayesian model of empirical confirmation of GR.

But it's not surprising that such a simple model would fail to capture the methodological richness of the scientific approach. As Adlam says, "[we] need to remain aware of the bigger picture in which the goal is to arrive at a coherent system of beliefs which includes a satisfactory explanation of how we could possibly have come to know the things we take ourselves to know." It is still not clear to me that something along the lines of the Everettian Statistical Assumption, i.e. a schematization of the observer that justifies assuming that we will continue to observe frequencies roughly consistent with mod-squared amplitudes, doesn't give this to us. After all, in an Everettian universe augmented with such an assumption, we can tell more or less the story that actually happened: a randomly-selected observer (say, Hugh Everett) will find himself in a situation where what he has personally observed in the past - and his records of the entire observational history of the scientific community - is consistent with the predictions of microscopic quantum mechanics. Then he might reflect on the relative merits of adding in an ad-hoc Heisenberg cut to unitary quantum mechanics and, through a process that seems futile to try to model via Bayesian updating, decide that it seems more "likely" (in a purely colloquial, non-quantitative sense for which I lack a better word) that in fact we are in a multiverse. So oughtn't we to believe that in such a universe a randomly selected, or "typical" observer come to believe exactly what Everettians do believe?

Two comments here: First, I haven't talked at all about the discussion of QBism in this chapter. This is on one hand because at the moment I'm sort of following the Everettian thread through the book, and on the other because I don't fully understand what the claims of the QBist approach actually are. Adlam discusses these of course, and I'll have to go back and do a closer reading. Second, there's obviously something wrong with my reasoning here, and I need to come back and understand why my objections fail on a more careful reading of this chapter, but I'm recording them for my own future use.

CHAPTER 9 (OBSERVER-RELATIVE INTERPRETATIONS)

These are approaches in which measurements are taken to have unique outcomes, which are "relative to individual observers". In some sense, these outcomes are private, first-personal facts (not Adlam's characterization, and I don't know if she would agree with this terminology). Such approaches generically feature what Adlam calls Type-III disaccord – what one observer hears upon asking another what outcome they saw when making some measurement does not necessarily agree with what the latter experiences himself saying. Thus in such approaches, there can be no empirical evidence for regularities predicted by quantum mecahnics persisting across all perspectives. Even worse, as there is no clear way to identify observers at different times, if such approaches are on the right track then our own memories become unreliable, and we cannot rely on any sort of temporally-extended reasoning. This seems to create a problem for the standard methods of scientific epistemology.

Now, even without quantum mechanics, we face the problem that the goings-on in another mind are unknowable. Do all scientific theories then suffer from the epistemic difficulties arising from the failure of intersubjective communication? Adlam says that they do not, pointing out that in classical physics, scientific knowledge may be taken to be knowledge about "regularities occurring in some mind-independent external world", whereas in observer-relative approaches such an external world simply doesn't exist, so there is nothing mind-independent for knowledge to be about other than the perspectives of other observers. Moreover, in classical theories, an assumption of reliable intersubjective communication is not as ad hoc as it would be in quantum theories, precisely because there is a mechanism for this sharing of experiences to arise, namely the causal relationship between this external reality and the experiences of the various observers.

This distinction between the classical and quantum settings is not quite convincing to me – for example, in the Everettian picture, within each branch there is indeed a mechanism for reliable intersubjective communication, namely whatever mechanism/regularity/law we suppose accounts for the supervenience of mental phenomena upon physical phenomena in the classical setting, applied branch-wise. Thus it seems to me that there is some equivocation going on here between two distinct schematizations of the observer. If the fact of there being an observer having a particular experience at a point in time is identical to the fact of there being a particular set of physical facts obtaining at that time, then it seems that we can say that at all times the universal wavefunction should verify the proposition (projection) that all observers agree. If on the other hand there are additional facts about experiences, over and above the physical facts (as I guess there must be in an observer-relative approach), then it seems that we should also worry that this is the case in classical theories. In any case, I'm very confused at the moment, so I'm just thinking out loud (silently) here. I'll certainly have to return to this very rich chapter and think much more slowly about the points Adlam raises.

Adlam also discusses several "non-absoluteness theorems", which demonstrate that in various extended Wigner's Friend scenarios, it is impossible to assign probability distributions over variables representing Absolute Observed Events in a way that satisfies some other intuitive axioms, e.g. locality, no-superdeterminism, no-retrocausality, and the universality of unitary quantum mechanics. Such theorems are often taken to provide support for observer-relative approaches, but of course one of the other assumptions could fail instead. For example, retrocausality is certainly counterintuitive, but Adlam points out that the sorts of backwards causation that would need to occur in the EWF settings relevant to the non-absoluteness theorems involve supermeasurements, i.e. one observer making a measurement of another observer and all systems that second observer has become entangled with since performing a previous measurement. Such measurements would be very far from anything that has actually been done in the process of developing and confirming quantum mechanics, so if retrocausality shows up in that context, it wouldn't necessarily disrupt the standard methods of scientific confirmation. Thus retrocausality of the sort needed in the EWF scenarios is in fact benignly, rather than epistemically weird.

CHAPTER 10 (FIRST AND THIRD PERSON VIEWS OF SCIENCE)

In the previous chapter, Adlam made the case that the intrinsically first-personal nature of measurement outcomes in observer-relative approaches makes trouble for scientific epistemology because it prohibits reliable intersubjective communication. She goes further here, claiming that "even in the non-observer-relative case, it is still the case that scientific knowledge...would be less robust if I were the only knower of it and all the rest of you were just automata...[In] the process of doing science we are not just relying on other observers to report measurement outcomes to us like automata – we are also relying on them to exercise judgment and to report conclusions that we take seriously because we regard other observers as our epistemic peers". Doing away with intersubjectivity risks erasing van Fraassen's distinction between public hallucinations (rainbows, images seen through microscopes, and any observations related in systematic and predictable ways) and private hallucinations (dreams, visions, etc.). Without this distinction, it seems that we lose a major concept involved in claims of scientific objectivity and a distinctive subject matter for physical science.

Something I think should be addressed here is the distinction between the functional and phenomenal aspects of the "judgement" of other observers. Indeed, Adlam refers to the necessity of "[believing] that they have minds which are just as capable as our own of producing good judgement about scientific evidence". But the notion of production seems to require only the functional aspects of judgment. Thus it is unclear why we should not be satisfied with these aspects alone. Would one of Chalmers's zombies not be a valid member of our epistemic communities? And isn't this effectively how Alice should conceive of Bob, in a world in which what she hears when she asks Bob about an observation he has made isn't the same as what Bob (the "real", phenomenally-alive Bob) experiences himself saying? With respect to the public/private hallucination distinction, it seems to me that there might be a way to salvage at least some of this distinction, even in a fully observer-relative world, via a notion of coherence with the full collection of perceptions available to an observer. Perhaps such a move would require that we view someone having a perfectly coherent hallucination of an alternative physics, including hallucinated colleagues reporting hallucinated data from hallucinated experiments, is doing proper science – but it isn't clear to me that this is necessarily wrong.

Adlam next turns to an argument against the view that observer-relative approaches are the "natural endpoint of the progressive relativisation of more and more physical quantities", as exemplified by the progression from Galilean relativity to special relativity to general relativity. In fact, she argues, the proper way to understand this progression is as a process of making the observer's perspective less central to fundamental physical theory, by "removing from our fundamentla physical descriptions quantities that we now understand to be observer-dependent, in order that we can better understand the absolute, invariant reality that grounds the relational quantities and the transformations between them". Thus denying the existence of this objective third person reality, as in observer-relative approaches, may actually be directly opposed to this historical trend.

Finally, Adlam returns to the Everett approach. Because measurement outcomes in this picture are "irreducibly indexical information", the Everettian must conclude "that a theory which has been arrived at and confirmed by the effort of the epistemic community as a whole can nonetheless only be properly understood from the first person point of view". Adlam claims that the presence of such "irreducibly indexical facts" poses a much more serious challenge to the standard story of scientific epistemology than the presence in more standard thinking of phenomenal facts like "the experience of seeing redness, which we usually imagine can be ignored in the formulation of a physical theory", since the outcomes of quantum mechanical measurements, which are for the Everettian irreducibly indexical, are exactly what we use to confirm the theory. But this isn't quite right, I think. What we use are reports of measurement outcomes, and we do indeed use reports of color experiences to confirm theories of (for example) color perception. Thus it isn't clear to me why we can't take a similar heterophenomenological approach (to invoke Dennett) here. Of course, in the setting of a color perception theory developed within a classical worldview, we (may) make the assumption that there is a unique phenomenal experience supervening upon a unique physical situation in a lawlike way, whereas in the Everettian picture we have to make this assumption branchwise, but it's not obvious to me that this can't be made to work.

CHAPTER 11 (PRIMITIVE ONTOLOGY, FUNDAMENTALITY, AND SCALE)

Primitive ontology approaches are those that posit some classical ontology living in a three dimensional space, and providing the building blocks of everything else. Examples are the de Broglie-Bohm approach, which takes for its ontology a collection of classical particles characterized fully by their positions, and various spontaneous collapse approaches, which take mass distributions or pointlike events. Such approaches solve the epistemic version of the measurement problem, since "we need only specify that our conscious experiences, and/or our macroscopic reality, supervene directly on this classical ontology". Since the elements of this classical ontology cannot enter into superpositions, neither can our experiences, and we are not faced with the epistemic problems Adlam has argued appear in unitary-only approaches. Certainly there are still nonintuitive aspects of these approaches, including in what they say about measurements. For example, in the de Broglie-Bohm framework, outcomes of spin measurements don't reveal a preexisting quantity of a particle, but rather depend in some complicated way on the positions of that particle and the particles making up the measurement device. Nonetheless, these failures of classical intuition don't pose epistemic problems, so this is merely benign weirdness.

Unfortunately, it seems difficult or impossible to produce a primitive ontology approach that is consistent with QFT. The basic problem is that, because of the complexity of scaling transformations in QFT, it seems unlikely that there is a relatively simple set of microscopic variables that gives rise to a relatively simple set of macroscopic variables on which for experience to supervene. In particular, "it is very challenging to define any relatively simple structures at small scales in QFT which would reduce to something like de Broglie-Bohm particles or spontaneous collapses in the position basis in just one branch of the wavefunction in the regimes relevant to non-relativistic quantum mechanics". Thus, either the predictive success of QFT is a coincidence, and the emergence of non-relativistic theory from a more fundamental theory is not mediated via QFT, or it is unlikely that the way in which these primitive ontology approaches resolve the measurement problem (by eliminating superposition/branching) can be maintained in a modified version of the approach that accommodates more fundamental physics.

Adlam closes the chapter with a discussion of notions of "autonomy of scales" (more fittingly autonomies of scales). It's widely understood that in order to do physics (the discipline) at one scale, we must be able to a large extent to ignore physics (the actual workings of the world) at much smaller scales. Following Franklin, Adlam introduces two versions of this notion of autonomy. "Autonomy from microstates" is the condition that the dynamics at the level of macrostates (sets of microstates) doesn't depend much on which particular microstate a system is in. Clearly this is necessary for empirical knowledge, since without it we wouldn't have any macroscopic regularities whatsoever. "Autonomy from microlaws" is the condition that the dynamics at the level of macrostates doesn't depend too much on the details of the micro-physics, for example that it is not too sensitive to small changes in parameters of the microscopic dynamical equations. Thus this latter autonomy concerns counterfactual changes rather than actual ones. Nevertheless, failure of autonomy from microlaws can also pose epistemic problems, because we rely on macroscopic theories to interpret measurement outcomes. Thus if changes to the microscopic theory radically alter our picture of the macroscopic world, this undermines what we have taken to be established scientific knowledge. This, Adlam says, is what happens in the case of the Everett approach, where the microscopic physics is taken to tell us that the macroscopic world is in fact a branching multiverse.

CHAPTER 12 (SUPERDETERMINISM AND NATURALNESS)

Superdeterminism is a family of approaches to addressing Bell inequalities that rejects non-correlation of measurement settings with the random ontic states corresponding to entangled pair preparation. A standard implementation of this idea is what Adlam calls "superdeterminism 1.0", in which the relevant correlations are built into the initial state of the universe. A possible epistemic concern that arises in such scenarios is that "the ability to make independent interventions is a crucial part of the scientific method". However, Adlam argues, while superdeterminism may cause problems for causal inference, it needn't necessarily disturb our usual methods of empirical knowledge production, as long as these correlations are "fairly mild and subtle". This leads into a discussion of naturalness and fine-tuning, in the context of initial state selection and choice of fundamental constants. The problem of initial state selection does, Adlam argues, lead to epistemic problems for superdeterminism 1.0 that typical anthropic approaches cannot resolve.

The chapter also discusses retrocausality approaches, but I'm not familiar with the concrete mathematical instantiations of these programs, so I'll have to come back to this part another time.

CHAPTER 13 (WHERE TO FROM HERE?)

Adlam discusses some potential directions for making progress on the epistemic version of the measurement problem in the event that neither unitary-only approaches nor primitive ontology approaches can be made to work, as she has argued they are unlikely to. I'll just touch on one of these.

Adlam herself, together with Rovelli, has worked on an approach known as relational quantum mechanics with cross-perspectival links. The core idea here is to add a new postulate to the framework of relational quantum mechanics. This requires that "[in] a scenario where some observer Chidi measures a variable V of a system S, then provided that Chidi does not undergo any interactions which destroy the information about V stored in Chidi's physical variables, if Alice subsequently measures Chidi in a basis corresponding to the physical variables representing Chidi's information about the variable V, then Alice's measurement result will match Chidi's measurement result". Adlam points out two issues that need addressing in this program. First, quantum field theory does not have well-defined subsystems (except possibly for spatial regions, which raises more problems), it is unclear how to formulate the cross-perspectival links postulate in the QFT setting. Second, the postulate itself is not clearly justified. I certainly can't say this any better than Adlam does:

"One might be tempted to offer a transcendental argument to the effect that a condition like CPL is necessary for the possibility of objective scientific knowledge, so we are entitled to simply assume CPL without any particular empirical or theoretical support for it. And indeed, I think we certainly must begin the process of scientific enquiry by assuming something like CPL. However, as discussed throughout this book, in general we expect that our scientific enquiries will provide empirical and/or theoretical support for the basic assumptions that we need in order to get the process of enquiry started, so our initially unjustified assumptions about the reliability of our means of enquiry can be retrospectively justified in a progressive coherentist sense."

The failure of this retrospective justification, and therefore of the "progressive coherentist" version of theory development Adlam advocates, is the consequence of the various forms of "epistemic weirdness" argued to be present in the various unitary-only approaches to the measurement problem. Adlam suggests, but leaves open for further investigation, the possibility that RQM+CPL fares better.

I do wish that the additional structure of this approach were more explicitly set out here (as it presumably is in the relevant papers). It seems that to the standard equipment of non-relativistic quantum mechanics we need to add a set of phenomenal state spaces, indexed by a set of observers, and some prescription for picking out a single phenomenal state for each observer given a universal wavefunction. The CPL postulate then becomes a condition on this prescription. The framework begins then to look like the type of dualism that Wigner assumes in "Remarks on the Mind-Body Question" (as I argued in that pinned blog post). Whereas Wigner resolves the technical problem by positing that consciousness causes wavefunction collapse, Adlam refers to RQM+CPL as a "modified unitary-only" approach, and no actual collapse of the wavefunction is supposed to occur. Thus presumably it is still meaningful in the latter picture to say that the physical variables representing an observer's experiences participate in superpositions, even if the actual experiences do not. CPL+RQM thus seems to imply a relatively complicated and indirect form of supervenience of the mental on the physical.

Adlam closes out the chapter by discussing connections to quantum gravity. She points out that Healey has argued that the problem of time raises similar epistemic concerns to those involved in the measurement problem, because "measurement is a process which takes place in time, and which by definition involves a change in the state of knowledge of the observer performing the measurement". Thus a theory of quantum gravity without time may be empirically incoherent: the "unreality of time is a form of epistemic weirdness rather than merely benign weirdness, meaning that it is not something we can just choose to accept". Perhaps there is a way to make sense of measurement in a timeless universe, but the problem must be addressed. The technical problems of quantum gravity are intimately related to the fact that we don't know how to handle the incorporation of observation and measurement, and thus the epistemic version of the quantum measurement problem could plausibly provide guidance in the development of a coherent theory of quantum gravity.

CHAPTER 14 (CONCLUSION)

Adlam reiterates her argument that empirical coherence is a productive criterion for scientific progress, and expresses optimism that this new perspective on the measurement problem provides real guidance for physics. Having reached the end of the book, I think I agree this is a real possibity! As expected, and hoped, I now have a ton more reading to do. In particular, I think my main stumbling block is that I'm not fully convinced that Everett + ESA can't be made to work, so following up on that argument is going to be my next step.