Pragmatism and the content of quantum mechanics

Peter J. Lewis

Draft – please don’t quote

Abstract

Pragmatism about quantum mechanics provides an attractive approach to the question of what quantum mechanics says. However, the conclusions reached by pragmatists concerning the content of quantum mechanics cannot be squared with the way that physicists use quantum mechanics to describe physical systems. In particular, attention to actual use results in ascribing content to claims about physical systems over a much wider range of contexts than countenanced by recent pragmatists. The resulting account of the content of quantum mechanics is much closer to quantum logic, and threatens the pragmatist conclusion that quantum mechanics requires no supplementation.

1. Introduction

Quantum mechanics is, notoriously, a theory in need of interpretation. But there is very little agreement on what kind of interpretation it needs. That is, there is very little agreement concerning what the foundational problems of quantum mechanics *are*, and without such agreement, there is little hope for a consensus concerning what an acceptable solution to the problems might look like.

Here is a way to divide up the territory. We can distinguish between *descriptive* and *normative* questions concerning quantum mechanics. Descriptive questions concern what quantum mechanics *says*—the *content* of the theory, as expressed in textbooks and used in labs. Normative questions concern what quantum mechanics *should* say—and in particular, whether it should say something different from what it actually does say.

All parties to the debates over the foundations of quantum mechanics would agree, I think, that there is a legitimate descriptive question concerning the content of quantum mechanics. Even those philosophers and physicists who think that quantum mechanics wears its interpretation on its sleeve at least feel the need to correct the mistaken impressions of *other* philosophers and physicists concerning what quantum mechanics says. The normative question presupposes an answer to the descriptive one: some think quantum mechanics is just fine the way it is, others contend that it needs to be replaced or supplemented with something radically different, and in large part this difference in attitude depends on prior differences concerning the answer to the descriptive question.

As an illustration, consider a fairly standard narrative concerning the descriptive and normative questions. Descriptively speaking, quantum mechanics depends on a distinction between measurements and non-measurements: measurements follow one dynamical law, the collapse dynamics, and non-measurements follow a different dynamical law, the Schrödinger dynamics. Since these two dynamical processes are incompatible, a precise formulation of quantum mechanics requires a precise dividing line between measurements and non-measurements. Quantum mechanics nowhere provides such a thing—and indeed, it seems highly unlikely that a term like “measurement” could be given a physically precise definition. So descriptively speaking, quantum mechanics is inadequate as a physical theory. On the basis of this measurement problem, Bell (2004, 213–231) recommends replacing quantum mechanics with either a pilot-wave theory or a spontaneous collapse theory. For similar reasons, Wallace (2012, 35) recommends replacing quantum mechanics with a many-worlds theory.[1]

But not everybody concurs. There are alternative narratives according to which quantum mechanics, descriptively speaking, is just fine as it is, and hence there is no normative pressure to supplement or replace it. One prominent version proceeds from the quantum logic of von Neumann (1936) and Putnam (1975) through to the quantum information theory of Bub (2016). According to this approach, quantum mechanics describes a non-classical event space—in terms of truth values, a non-Boolean algebra, and in terms of probability ascriptions, a non-simplex distribution. No-go theorems (arguably) show that it is impossible to construct a set of events obeying classical Boolean logic or classical Kolmogorov probability that reproduces the empirical predictions of quantum mechanics. The implication is that in quantum mechanics we have discovered something important about the fundamental event structure of the world. Seeking to replace or supplement quantum mechanics with a theory obeying classical logic and classical probability theory amounts to a quixotic attempt to impose a structure on the world that it manifestly does not have (Bub 2016, 222). The measurement problem, on this account, results from a mistaken demand for a dynamical explanation of the individual events in the quantum structure, when no such explanation is available (Bub 2016, 223)

---

[1] Wallace takes the many-worlds theory to be a precise statement of the content of quantum mechanics, rather than a replacement for it. I take up the question of whether the many-worlds structure is present in quantum mechanics as it stands in section 2.

This fundamental difference of opinion—between those who take the measurement problem seriously and those who regard it as a pseudo-problem—continues to divide the foundations of physics community today. Hence the descriptive question—the question of what quantum mechanics actually *says*—remains a pressing one. In this paper, I argue for a particular way of approaching the descriptive question. The methodology is the pragmatist one of Healey (2012; 2017) and Friederich (2015), but the answer to the descriptive question that results from following this methodology, I argue, differs in an important way from the answers that Healey and Friederich give. I conclude by assessing the consequences of this answer to the descriptive question for the normative question.

2. The descriptive question

So how should we approach the descriptive question? Consider a straightforward realist approach to the content of scientific theories. A theory, at least in physics, is typically expressed using a particular mathematical structure. The *state* of a physical system is generally identified with a mathematical entity that resides in a particular abstract space, and the *dynamics* of the theory tell us how that state evolves over time. So, for example, in many applications of classical mechanics, the state of a physical system can be represented by a set of vectors in a three-dimensional Euclidean space, and the dynamical laws of Newtonian mechanics tell us how the set of vectors evolves over time. The interpretation of the mathematics is fairly straightforward: the vectors represent the positions and momenta of point-like particles, and classical mechanics tells us how the properties of the particles change.

Such an approach can equally be applied to quantum mechanics (Albert 1996). According to quantum mechanics, the state of a physical system is identified with a complex-valued function defined on a configuration space—a space with three dimensions for each particle in the system. A dynamical law, the Schrödinger equation, tells us how this function, the wave-function, changes over time. Then by analogy with classical mechanics, the wave-function must be a representation of the physical properties of the quantum system as they change over time.

The continuity with classical mechanics in the above account is attractive, but there are surprising consequences. For an *N*-particle system, the wave-function is defined over a 3*N*-dimensional configuration space, and it cannot be represented without loss in a three-dimensional space. This has led some to conclude that a straightforward realist reading of quantum mechanics shows that the three-dimensionality of our physical world is illusory (Albert 1996). Furthermore, if we model a measurement using quantum mechanics, the wave-function ends up with components corresponding to each possible outcome of the measurement—not just one outcome, as is the case classically. This leads Everettians like Wallace (2012) to conclude that a straightforward realist reading of quantum mechanics shows that every possible outcome of a measurement actually occurs.

These conclusions might be right, but do they simply follow from close attention to the structure of quantum mechanics? There are reasons to be suspicious. As Healey (2017, 116) notes, conclusions of this kind depend on the assumption that the wave-function plays the same descriptive role in quantum mechanics as the position-momentum vectors play in classical mechanics. If this assumption is itself up for grabs in the interpretation of quantum mechanics, then neither of these conclusions is warranted. But how do we adjudicate the question of whether the wave-function describes physical systems or whether it has some other, non-descriptive role? Is there a metaphysically neutral methodology that could be used to answer this question? Healey (2012; 2017) and Friederich (2015) think that there is.

3. Pragmatism

Consider an analogy. “Stealing is bad” has the same grammatical structure as “Cherries are red”. But it is far from clear that both sentences should be taken as descriptive. In particular, badness, taken as a property of actions, seems like a queer kind of property, imperceptible and disconnected from the other properties of the action. Expressivists seek to dissolve the problem of the nature of badness by claiming that a sentence like “Stealing is bad” should be taken as expressive rather than descriptive—as expressing our attitude towards stealing. Pragmatists further coopt expressivism as a variety of pragmatism (Price 2011, 9). Pragmatists stress the variety of uses of language, noting that sentences with superficially similar form can be used in radically different ways. “Cherries are red” is used to describe a class of objects, whereas “Stealing is bad” is used to express our attitude towards a class of actions.

Pragmatism, then, enjoins us to pay close attention to how a sentence is *used* in order to find out what it means. Healey (2012; 2017) and Friederich (2015) each suggest that the pragmatist approach provides us with a metaphysically neutral methodology for probing the content of quantum mechanics. That is, we can look at how various quantum mechanical claims are used by physicists in order to determine what those claims mean. This strikes me as a welcome suggestion. In the rest of this section I present the conclusions of their pragmatist inquiries; in the next, I consider whether the language use of physicists actually supports those conclusions.

Healey (2012) distinguishes between *quantum claims* and *non-quantum magnitude claims*. The former explicitly mention quantum states, quantum probabilities, or other novel elements of the theory of quantum mechanics. The latter are claims about the magnitude of a physical quantity that do *not* involve quantum states, quantum probabilities etc. In keeping with the pragmatist methodology, Healey bases this distinction on the way the two kinds of claims are used. Non-quantum magnitude claims are used in a straightforwardly descriptive way. But quantum claims are used in a different way: they are used, not to *describe* a system, but to *prescribe* a user's degrees of belief in various non-quantum magnitude claims.

As an example, Healey appeals to the Interference experiments of Juffmann et al. (2009), in which C_{60} molecules are passed through an array of slits and then deposited on a silicon surface. To derive quantum mechanical predictions for this experimental arrangement, quantum states are ascribed to C_{60} molecules. That is, quantum claims of the form "The molecule has state $|\psi\rangle$" are used, via the Born rule, to ascribe probabilities to claims concerning the various possible locations of the molecules on the silicon surface. These latter claims—of the form "The molecule is located in region R"—are non-quantum magnitude claims. The job of the non-quantum magnitude claims is to describe the physical system, but the job of the quantum claims is to prescribe degrees of belief in the non-quantum magnitude claims for an appropriately situated observer. In this respect Healey's approach is like the expressivist's in ethics: claims that have superficially similar grammatical forms have very different functions.

Another important strand in the pragmatist approach concerns the role of decoherence. After the C_{60} molecule hits the silicon surface, complicated interactions with the surface mean that the state of the molecule-environment system becomes approximately diagonal when written as a density matrix in the position basis. This in turn insures that the probabilities ascribed by the Born rule to various claims about the molecule's position closely obey the probability axioms. But before the molecule encounters the silicon surface, its state is a coherent superposition—a state that is not even approximately diagonal, and for which the Born rule does not ascribe probabilities to location claims that closely obey the probability axioms. For such a state, the Born rule does not prescribe appropriate degrees of belief in the non-quantum location claims, and so assertion of such claims prior to decoherence is not *licensed* by quantum mechanics. Decoherence, then provides a demarcation between situations in which it is appropriate to have a well-defined degree of belief in a non-quantum magnitude claim, and situations in which it is not.

The central finding of the Healey-Friederich pragmatist approach is that attention to the use of quantum mechanical language shows that claims about the quantum state of a system are not used to describe that system. Hence, we should not think of the wave-function as a representation of the physical properties of the quantum system as they change over time. This perspective has the advantage that the measurement problem does not arise: if the wave-function doesn't represent the system, then we don't have to worry that the dynamical laws for wave-function evolution are different for measurements and non-measurements. In fact, if the quantum state is prescriptive, then the difference between measurements and non-measurements arises quite naturally: the results of measurements have a direct and obvious influence on what you should believe.

Hence the pragmatist approach provides a clear answer to the descriptive question: quantum mechanics, in itself, says *nothing* about the world. As Healey (2017, 12) puts it, "quantum theory has no physical ontology". Rather, quantum mechanics tells us what to believe about non-quantum ontology—about particles, or in the case of quantum field theory, about fields. Furthermore, this answer to the descriptive question suggests an answer to the normative question: since the measurement problem doesn't arise, there is no motivation for supplementing or replacing quantum mechanics with something else.

4. Actual use, counterfactual content

Thus far, I have said little about the evidence that backs up Healey's claims about how quantum claims and non-quantum magnitude claims are used. Indeed, direct evidence from the language use of physicists is likely to be unenlightening: that a claim is asserted in a given context provides no direct evidence concerning whether its content is descriptive or prescriptive.

To fill this gap, Healey appeals to an inferentialist account of the link between use and meaning derived from the work of Robert Brandom (2000): the meaning of a claim is identified with the set of material inferences it licenses. So by looking at the way a claim is used in licensing inferences, we can gain evidence about what it means. And here the distinction between prescriptive quantum claims and descriptive non-quantum magnitude claims seems to be well motivated. In the practice of physics, a claim about the quantum state of a system is used to infer Born probabilities, and nothing more. If Born probabilities are taken to be rational degrees of belief, then the prescriptive content of a quantum claim exhausts its meaning.

A non-quantum magnitude claim, on the other hand, can license a wide variety of inferences. From the claim that a C_{60} molecule is located in a particular region of the silicon surface, we can infer that an electron microscope will produce an image of the molecule if directed at that region (Juffmann et al. 2009, 2). We can infer that if the silicon surface is left untouched for two weeks, the C_{60} molecule will remain in the same place (Juffmann et al. 2009, 2). Under suitable conditions, we can infer that the C_{60} molecule will emits photons; under different conditions, that it will act as a nucleation core for molecular growth (Juffmann et al. 2009, 3). In other words, the inferences licensed by the non-quantum magnitude claim support the interpretation that the meaning of the claim is descriptive rather than merely prescriptive.[2]

So there is a good case to be made, I think, that actual use supports the distinction between prescriptive quantum claims and descriptive non-quantum magnitude claims. But there is a further strand to the Healey-Friederich interpretation, namely that non-quantum magnitude claims are only licensed after decoherence. This claim, I think, does not stand up so well to scrutiny.

Consider C_{60} interference again. After the molecule has adhered to the silicon surface, the state of the molecule is decoherent, and the claim that the molecule has a particular

---

[2] There is a sense in which the meaning of *any* claim is prescriptive according to the inferentialist program: the claim about the location of the molecule licenses an inference to a certain *degree of belief* that the electron microscope will produce an image of it. But still, there is a reasonable distinction here: the quantum claim licenses inferences only via the Born rule, whereas the non-quantum magnitude claim licenses inferences via a huge variety of schema typical of small physical objects. The latter is just what it is for a claim to be descriptive.

location is licensed—that is, it is appropriate to associate a particular degree of belief with the claim, and if that degree of belief is high enough, it is appropriate to assert the claim. But before the molecule has adhered to the silicon surface, the state of the molecule is coherent, and no claim about the location of the molecule is licensed—it is not appropriate to associate a degree of belief with such a claim, or to assert it. Similar considerations apply to properties other than location.

This seems to fly in the face of actual use. For example, in the description of the C_{60} interference experiment, Juffmann et al. (2009, 2) assert that “all transmitted particles arrive with the same speed,” and “about 110cm behind the source, the molecules encounter the first diffraction grating,” apparently ascribing both speed and location to C_{60} molecules prior to decoherence. This doesn’t seem to be an isolated incident: physicists routinely talk of preparing, selecting, spraying, shooting and trapping particles, ions and molecules, and this talk typically involves making claims about these objects prior to any eventual decoherence.

It is possible, of course, that this is just “loose talk”, or an indirect way of making claims about the quantum state of the systems concerned. But given the frequency of such claims, and given the reliance of the pragmatist methodology on *use*, this seems like a shaky game to play. It would be better, all things considered, if such claims could be accommodated within the pragmatist interpretation, rather than explained away as anomalies.

But there are obvious barriers to licensing non-quantum magnitude claims prior to decoherence. As Friederich (2015, 79) notes, the Born rule is only “reliable” when applied to decoherent states, in the sense that only for such states are the numbers it produces guaranteed to closely obey the probability axioms. Given some reasonable assumptions about rationality, it is plausible that numbers that do not closely obey the probability axioms could not be rational degrees of belief. Furthermore, Healey argues that asserting a non-quantum magnitude claim prior to decoherence is likely to be misleading. For example, suppose one asserts (with Juffmann et al.) that “about 110cm behind the source, the molecules encounter the first diffraction grating.” One might infer from this that each molecule passes through exactly one slit in the grating, and hence that the presence of the other slits is irrelevant, and hence that there is no possibility of interference (Healey 2012, 745).

So the pragmatist approach seems to face a dilemma: either it fails to accommodate the actual language use of physicists, or it licenses misleading assertions and irrational degrees of belief. Isn’t there another way? I think there is. Consider a mundane claim like “There is beer in the fridge.” In typical contexts, an assertion of this claim licenses the inference that if you were to go to the fridge and open the door, you could take a beer and drink it. Of course, you might not actually do this; maybe you don’t want a beer. That is, the inference here is a counterfactual one. A good deal of the inferential content of our assertions has this counterfactual character.

Now return to the quantum context. Consider again the claim that “about 110cm behind the source, the molecules encounter the first diffraction grating.” What content could that claim have? If we broaden the notion of inferential content to include counterfactual inferences, then the content seems fairly clear: if we were to replace the first diffraction grating with a detector taking up the same region of space, then the Born rule would ascribe a degree of belief close to 1 to detecting the molecules.

How does the inclusion of counterfactual content avoid the barriers to licensing non-quantum magnitude claims prior to decoherence? Note that the counterfactual content of the claim about the molecules involves a counterfactual intervention on the system—a counterfactual measurement. The counterfactual measurement induces counterfactual decoherence. The Born probabilities are conditional on this intervention and the associated decoherence, so the Born probabilities for various position claims concerning the molecules are not, after all, unreliable, in the sense of violating the probability axioms.

Neither should there be any danger of being misled by an assertion that the C_{60} molecules encounter the grating, because the counterfactual conditions implicit in the content of that assertion are distinct from the conditions that actually obtain in the apparatus. That you *could* detect the molecules at the diffraction grating, given a different experimental arrangement, doesn't license the inference that there *is* no interference, given the actual experimental arrangement. Admittedly, though, this amounts to a weakening of the content of position claims from the classical case, as spelled out in the next section.

5. A happy convergence?

I have argued that non-quantum magnitude claims have assertible content in a far wider range of contexts than countenanced by Healey or Friederich. If there is some counterfactual intervention on a system that would produce decoherence in the basis defined by a given observable, then claims about the values of that observable have content. And since counterfactual interventions only have to be realizable in principle, this means that claims about the value of an observable for a system *generally* have content, whether or not the system *actually* decoheres in the basis defined by that observable. This has the welcome consequence that the frequent assertions made by physicists about the properties of systems prior to decoherence are contentful.

A potential cost of such permissiveness about content is that the structure of this content is, in general, non-Boolean. Consider again a C_{60} molecule that is approaching the first diffraction grating, and consider an assertion of “The molecule passes through the leftmost slit”. This assertion has content, on the proposed view, because in principle there is an intervention on the system that would produce decoherence in a basis defined by an observable that distinguishes which slit the molecule passes through. Still, assertion of the claim would not be appropriate, simply because there are many slits in the grating, so the Born rule ascribes it a low probability. The same goes for every other slit in the grating. Nevertheless, the assertion that “The molecule passes through the leftmost slit, or the second to the left, or…” is assertible, since the Born rule ascribes it a probability close to 1. The disjunction is assertible, but none of the disjuncts is assertible. Since assertibility is a surrogate for truth in the pragmatist context, this is equivalent to saying that the disjunction is true, but none of the disjuncts is true.

One might take this to be unacceptable on the pragmatist view—especially if you endorse an inferentialist pragmatism, as Healey does. From a disjunctive claim you can straightforwardly infer that at least one of the disjuncts is true. If the content of a claim is identified with the inferences that it licenses, then part of the meaning of the disjunctive claim about the C_{60} molecule is that some assertion of the form “The molecule went through slit *x*” is true. Hence my proposal about content threatens to violate the inferentialist account of meaning. The pragmatist interpretation of Healey and Friederich avoids this problem by insisting that claims about systems have meaning only after suitable decoherence.

Of course, pragmatism is not necessarily tied to an inferentialist account of meaning. But even given inferentialism, there is arguably no real problem here. Physicists are *selective* in the inferences they draw: from the disjunctive claim, they don't infer that the C_{60} molecule goes through some particular slit, so they don't infer a lack of interference. But they do infer that the molecule will arrive at the silicon surface, that it might radiate a photon in flight, and so forth. That is, the inferences drawn by physicists from their claims about pre-decoherent systems suggest that the non-Boolean structure of those claims is already *built in* to the meanings associated with those claims and revealed in inference.

This suggests that close attention to the way non-quantum magnitude claims are actually used leads to a happy convergence between pragmatism and the quantum logical approach. Physicists assert claims about particles even when the state does not decohere, and such claims seem to be meaningful. But physicists are not inclined on that basis to draw all the inferences that a full Boolean structure to their claims would license. Quantum mechanics apparently weakens the meaning of many claims about pre-decoherent physical systems, but without rendering those claims meaningless.

6. The normative question

As a methodology for addressing the *descriptive* question of the content of quantum mechanics, the pragmatist approach seems entirely appropriate: look to the *use* of physicists to determine what the various claims involved in the theory mean. At the hands of Healey and

Friederich, this approach yields the important insight that while non-quantum magnitude claims are used to describe physical system, quantum claims are used to prescribe appropriate degrees of belief in non-quantum magnitude claims. But Healey and Friederich go further, in limiting the assertibility of non-quantum magnitude claims to contexts in which the quantum state is decoherent in the relevant basis. This, I have argued, cannot be squared with the actual use of such claims. I propose instead that non-quantum magnitude claims *generally* have well-defined content, understood in terms of a counterfactual intervention on the system. This change to the pragmatist approach means that it ends up looking a lot like the quantum logical approach that preceded it. Indeed, the pragmatist approach might be regarded as a *justification* for quantum logical claims concerning the content of quantum mechanics.

But where does all this leave the *normative* question concerning whether quantum mechanics is fine as it is, or whether it should be supplemented or replaced? Healey and Friederich argue that quantum mechanics is fine as it is: if quantum claims do not describe physical systems, then there can be no conflict between the way that quantum mechanics describes systems during measurements and the way it describes them during non-measurements. If there is no measurement problem, then there is no motivation to replace such a successful theory. If, as Healey (2017, 12) maintains, quantum theory “states no facts about physical objects or events,” then there can be no requirement that we come up with an *explanation* of quantum facts and events.

However, I have suggested that quantum theory has more content than the pragmatists countenance. In one sense, I agree that quantum theory states no facts: a quantum claim, such as the attribution of a quantum state to a system, is not a description. But in another sense, there are distinctive quantum facts, or at least facts with a distinctive quantum structure: non-quantum magnitude claims about pre-decoherent systems exhibit the non-Boolean structure characteristic of quantum mechanics. This is the sense in which quantum logic gets things right.

Notably, though, the proponents of quantum logic *also* often take the view that quantum logic dissolves the measurement problem (e.g. Putnam 1975, 186). But this dissolution is widely regarded to be a failure (e.g. Bacciagaluppi 2009, 65) Once one has admitted that the structure of true (i.e. assertible) claims for a quantum system is non-Boolean, the question of *how* the world manages to instantiate this structure becomes legitimate and pressing. A denial that any explanation is required looks suspiciously like instrumentalism. And since any answer to this question goes beyond quantum mechanics as it stands, the call for explanation involves a demand to supplement quantum mechanics, or to replace it with something more fundamental.

Of course, given the no-go theorems, the path to an explanation of the structure of quantum facts is by no means clear. But neither do the no-go theorems show that an explanation is *impossible* (Friederich 2015, 161).[3] If the foregoing is correct, then pragmatism is an excellent way to *expose* the foundational problems of quantum mechanics, but it is not a means to *dissolve* them.

References

---

[3] Interestingly, Friederich (2015, 161) suggests supplementing quantum mechanics with sharp values for all observables, even though this seems at odds with his therapeutic aim of dissolving the foundational problems of quantum mechanics rather than solving them (2015, 6).

Albert, David Z. (1996), "Elementary quantum metaphysics," in J. T. Cushing, A. Fine and S. Goldstein (eds.), *Bohmian Mechanics and Quantum Theory: An Appraisal*. Dordrecht: Springer, 277-284.

Bacciagaluppi, Guido (2009), "Is logic empirical?" in K. Engesser, D. M. Gabbay and D. Lehmann (eds.), *Handbook of Quantum Logic and Quantum Structures.* Amsterdam: North-Holland, 49-78.

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

Brandom, R. (2000), *Articulating Reasons: An Introduction to Inferentialism*. Cambridge, MA: Harvard University Press.

Bub, Jeffrey (2016), *Bananaworld: Quantum Mechanics for Primates*. Oxford: Oxford University Press.

Friederich, Simon (2015), *Interpreting Quantum Theory: A Therapeutic Approach*. Basingstoke: Palgrave Macmillan.

Healey, Richard (2012), "Quantum theory: a pragmatist approach," *British Journal for the Philosophy of Science* 63: 729-771.

Healey, Richard (2017), *The Quantum Revolution in Philosophy*. Oxford: Oxford University Press.

Juffmann, T., Truppe, S., Geyer, P., Major, A. G., Deachapunya, S., Ulbricht, H., and Arndt, M. (2009), "Wave and particle in molecular interference lithography," *Physical Review Letters* 103: 263601.

Price, Huw (2011), *Naturalism Without Mirrors*. Oxford: Oxford University Press.

Putnam, Hilary (1975), “The logic of quantum mechanics,” in *Mathematics, Matter and Method: Philosophical Papers Volume 1*. Cambridge: Cambridge University Press.

von Neumann, John (1932), *Mathematische Grundlagen der Quantenmechanik*. Berlin: Springer-Verlag.

Wallace, David (2012), *The Emergent Multiverse*. Oxford: Oxford University Press.