CALCULATION, BOOKKEEPING, REPRESENTATION, AND EXPLANATION: A PARABLE

TIM MAUDLIN
New York University
Department of Philosophy
twm3@nyu.edu

SUMMARY: Our topic is a basic discussion of what it is to regard a bit of mathematics as referring to something physical, and the different options one has for explaining —and hence rendering non-miraculous— the empirical success of a mathematical formalism. In order to strip away as much extraneous distraction as possible, we will leave the particular case of the wavefunction in quantum theory aside and tell a simpler tale. A fairy tale, in fact. The aim is to foreground some possible relations between a mathematical formalism and the physical system it is used to represent.

KEY WORDS: wavefunction, ontology, mathematics, quantum, theory

RESUMEN: Nuestro tema es una discusión básica de lo que es considerar que una porción de matemáticas se refiera a algo físico y las diferentes opciones que uno tiene para explicar —y por lo tanto hacer que no sea milagroso— el éxito empírico de un formalismo matemático. Para eliminar tanta distracción externa como sea posible, dejaremos a un lado el caso de la función de onda en la teoría cuántica y contaremos una historia más simple. Un cuento de hadas, de hecho. El objetivo es poner en primer plano algunas posibles relaciones entre un formalismo matemático y el sistema físico al que representamos con el formalismo.

PALABRAS CLAVE: función de onda, ontología, matemáticas, cuántica, teoría

Prologue

The fundamental question of ontology is: *What exists?* Hence the fundamental question of physical ontology is: *What physical entities exist?* A full answer to that question would include all physical entities, including (for example) particular tables and chairs, but no one wants that sort of full answer. God may be entertained by enumerating the hairs on each person's head, but neither philosophers nor physicists would be enlightened by such an undertaking.

So the question of physical ontology gets restricted in two ways. First, there is a focus on *types* or *species* of physical entities: are there particles or fields or strings or space-time? Second, there is a focus on the more *fundamental* as opposed to the *derivative*. Water is a type of physical entity, but a derivative type: to be water is to

be a composed of H_2O molecules, and to be an H_2O molecule is to be a certain bound state of two hydrogen atoms and an oxygen atom (which can be further analyzed into bound states of quarks and electrons). In some decent sense, that is all that water is, so having accounted for all the quarks and electrons (whatever *they* are) automatically accounts for the water. Properties or characteristics or quantities can similarly be distinguished into the more fundamental and the more derivative. The physical quantity *temperature* is understood as a statistical characteristic of a system composed of many subsystems, so temperature is a derivative quantity while the more fundamental characteristics are those of the subsystems that display the statistical profile. Pursued in this way, the central question of physical ontology is a request for an account of the fundamental physical properties, entities and kinds.

Of course, no one claims to know right now what the fundamental physical kinds are, as that would require having the fundamental physical theory —the correct Theory of Everything— which we do not possess. Still, one can ask questions like: "What are the fundamental physical kinds *according to some theory*?". And even if a theory is offered explicitly as non-fundamental, even if it is offered as some sort of *effective* theory valid only in a limited domain, one can ask: What are the most fundamental sorts of entities postulated by this theory as presently formulated? This is the sort of question that we have in mind in the sequel.

As a primary example, consider a cluster of questions raised in discussions of the ontology of quantum theory:

1) Is the wavefunction real?

2) What is the meaning of the wavefunction?

3) Is the wavefunction ontic or epistemic?

4) Is the wavefunction just a bookkeeping device?

5) What, if anything, does the wavefunction represent?

6) What, if anything, does the wavefunction represent about the physical system to which it is ascribed?

All of these questions are attempts to raise the same question, although I will argue that the last formulation is the best. Let's sort through the list.

Questions of ontology can become hopelessly confused if one does not cleanly separate *representations* from *what they represent*. In mathematical physics, mathematical objects are used as representations of physical states of affairs. Of course, there are completely separate questions about the ontology of mathematics: in what sense, if any, do mathematical objects and mathematical facts exist? Those are fascinating and difficult questions but are orthogonal to questions of physical ontology, for whatever mathematical objects are they are not *per se* physical.[1] As far as our inquiries are concerned, we will take the mathematical objects as given, with a certain uncontroversial mathematical structure. The question of physical ontology is not about the ontological status of the mathematical objects *per se*.

Unfortunately, in mathematical physics there is often a pervasive ambiguity in language that arises from using the same term to refer to a mathematical object and to a (putative) physical entity that it might represent. And in quantum mechanics that ambiguity infects almost all discussions. In particular, what is referred to by the term "the wavefunction" in the six questions above?

Clearly, the intent in question 1 cannot be to refer to any *purely mathematical* entity. If it were, then the question would be about the ontology of mathematics rather than physics. But still, there is a mathematical entity that is employed in the formalism used by quantum mechanics, and that mathematical entity is typically called "the wavefunction of the system". For example, we might be told that the wavefunction of a collection of N spinless particles is a complex function on the configuration space of the system, or a ray in a Hilbert space, or a density operator. All of these are clearly mathematical objects, and asking whether *they* are real is asking a question that has nothing to do with physics. Rather, what is under discussion is not the "reality" or otherwise of the wavefunction itself, but the reality or otherwise of something *physical* that in some way would *correspond* to the mathematical wavefunction. I have taken to calling this putative physical object the *quantum state* of a system. By definition, a quantum state would be a physical characteristic of a system, whose existence would be independent of any mathematical objects used to represent it. In this sense quantum states might not exist at all, just as it turns out that "states of caloric" do not exist because there is no caloric.

Talk of the "meaning" of the wavefunction is also suboptimal. Primarily it is representations, such as sentences in an interpreted lan-

---

[1] For a contrary view, see Tegmark 2014.

guage, that have meanings, so question 2 does put the wavefunction in the category of representations. But "meaning" is so polysemous that the question immediately demands further clarification.

Question 3 also presumes that the wavefunction is a representation, and foregrounds the question of what sort of thing it represents. The term "ontic" is slightly tendentious in this context, and should be read as a contrast class to "epistemic". To say that the wavefunction is epistemic is to say that it represents the credal state (or perhaps the ideal credal state) of some cognitive agent. This is a view explicitly endorsed by the QBists. Wavefunctions, for the QBist, are to be assigned not to *physical systems per se* but to *cognitive agents*. They provide advice to the agent about how to set their subjective credences. In QBism, these credences are not even about how standard physical systems will behave, but rather about what the personal experiences of that particular agent will be. In two words, the position of the QBist is both instrumentalist and solipsistic. Wavefunctions are mere mathematical tools for forming expectations, and the class of expectations is not about objective physical systems, but about egocentric subjective states.

Instrumentalism is an old story in philosophy of science. If all one wants out of science in general is an effective way to make reasonably accurate predictions then that's all one wants: there is nothing more to be said. Many people, however, have higher ambitions. They want to actually *understand* the world, not merely accurately predict it (much less only predict the content of their own experience of it). For such a person the instrumental effectiveness of a mathematically formulated scheme is not the *end* of scientific inquiry but rather the beginning. One wants to know *why the scheme predicts so well*. If the QBist doesn't want this sort of explanation, then there hardly seems reason to have a dispute: they can go happily on their way satisfied with the acknowledged predictive accuracy of "standard" quantum theory. I want more. Indeed, if all that physics aimed to provide were a reliable and accurate way to predict my own experiences, I just would not be interested in the topic at all. My own experiences are a tiny and rather uninteresting sliver of what there is in the world. I am after bigger game.

The real issue raised by the so-called "epistemic" accounts of the wavefunction isn't what the wavefunction is but rather what the whole point of doing physics is. I myself wonder why an instrumentalist solipsist of this sort would spend much time studying fundamental physics: if I want to predict my experiences of the weather tomorrow

I will consult a meteorologist rather than a physicist. But *de gustibus non disputandum est.*

The physicist —or should we say natural philosopher?— wants not just an accurate prediction-making scheme but an understanding of why it works. Even more precisely, the natural philosopher wants an accurate account of the nature of the physical world and would not be terribly surprised if that account also allows one to make accurate and reliable predictions, although that is just a side benefit. It seems that something like a wavefunction will play a central role in the mathematical formalism used to make accurate empirical predictions. The question before us is what claims about the physical ontology of the world are implied or suggested by that fact. What, if anything, in the physical world does the wavefunction *represent*?

That brings us to questions 4, 5, and 6. Once one asks what, if anything, the wavefunction represents, there are only three possible sorts of answers: it represents *everything*; it represents *something but not everything*; and it represents *nothing*.[2]

The claim that it represents everything has long been discussed using the terminology introduced by Einstein, Podolsky and Rosen: is the "quantum-mechanical description of physical reality" "complete"?. EPR explicitly answered "no", and Bohr and Von Neumann explicitly answered "yes". In contemporary discussions, the Everettians defend the completeness of the wavefunction as a representation of the physical universe: all the (non-*de-se*) physical characteristics of the universe are captured —somehow— by the wavefunction. That still leaves room for disagreement about exactly how the wavefunction represents, which mathematical degrees of freedom correspond to physical degrees of freedom and which are merely "gauge", etc. Proponents of GRW-type collapse theories also tend to posit the representational completeness of the wavefunction.

The claim that the wavefunction represents something but not everything is central to the pilot-wave picture. That proposes a dualistic physical ontology: some local beables (which could be particles or fields or "flashes" or strings) and a "pilot wave" that plays a role in determining the dynamics of the local beables. The local beables —not being represented by the wavefunction— require some additional mathematical representation. Historically, this additional piece of mathematics has been called "hidden variables", but as Bell remarks that nomenclature is silly. There certainly are additional

---

[2] This formulation is due to Shelly Goldstein.

mathematical variables, as required by the additional physical ontol-
ogy, but they are the opposite of "hidden". It is the local beables, or
some subclass of them, that correspond to what we experience, what
we see when we open our eyes. It is the quantum state represented
by the wavefunction that is hidden, which is why some people deny
its existence altogether. Once again, accepting that both the wave-
function and some other mathematical object represent real physical
ontology does not yet settle all outstanding questions about physi-
cal vs. gauge degrees of freedom, but at least it is a clear starting
point for that discussion.

The claim that the wavefunction represents nothing physical at
all is the most surprising option of the three. We have already
mentioned the QBist form of this gambit: the wavefunction is merely
a description (or better: part of a prescription) about the degrees
of subjective credence an agent has or should have. That leaves the
question of physical ontology completely untouched. And it poses
the further question of why the wavefunction should be so effective
as part of a predictive mechanism if it represents nothing at all in
the physical world. This is indeed a good question, but not perhaps
as devastating as it might seem at first glance.

Hilary Putnam famously opined:

> The positive argument for realism is that it is the only philosophy that
> doesn't make the success of science a miracle. That terms in mature
> scientific theories typically refer (this formulation is due to Richard
> Boyd), that the theories accepted in a mature science are typically
> approximately true, that the same term can refer to the same thing
> even when it occurs in different theories —these statements are viewed
> by the scientific realist not as necessary truths but as part of the only
> scientific explanation of the success of science, and hence as part of any
> adequate scientific description of science and its relation to its objects.
> (1973, p. 73)

A monograph could be written unpacking and critiquing this passage.
Each of the myriad qualifications ("mature", "typically", "approxi-
mately true") designed to block immediate refutation would have
to be clarified, the conditions for being a "miracle" spelled out,
etc., etc. Ptolemaic astronomy, Newtonian gravitation, Maxwellian
electrodynamics were in some sense "mature" theories with some
notable empirical success that nonetheless may be argued to have
fundamentally misguided ontologies. But even taking such a view, in
retrospect their empirical success is hardly miraculous. Something or

other in Putnam's version of scientific realism has gone wrong. But this paper is not that monograph.

It is, perhaps, a prolegomenon to such a study. Our topic is narrower and more modest: a basic discussion of what it is to regard a bit of mathematics as referring to something physical, and the different options one has for explaining —and hence rendering non-miraculous— the empirical success of a mathematical formalism. In order to strip away as much extraneous distraction as possible, we will leave the particular case of the wavefunction in quantum theory aside and tell a simpler tale. A fairy tale, in fact. The aim is to foreground some possible relations between a mathematical formalism and the physical system it is used to represent. Like all good fairy tales, it begins...

Once Upon A Time...

...there was a King, who jealously guarded his wealth. In order to make sure that none of his golden Ducats was stolen or embezzled, he had a special vault made. The only way the vault could be opened was by a combination lock, with four numbers to the combination. Each of the King's four daughters chose a number and set the lock to it, and then the Princesses retired to their castles in four distant corners of the kingdom. In order to open the vault, they all had to be called back together, a lengthy and cumbersome process.

Of course, the King needed to sometimes add to his treasure and sometimes withdraw from it, so the vault had two access points: one Deposit slot and one Withdrawal slot. The Royal Accountant could put any number N Ducats in the Deposit Slot at the top of the vault and turn a crank. The coins would disappear into the vault, and a ticket would be printed out with the notation: "N Ducats, Deposit" followed by the date and time. Similarly, the Royal Accountant could slide a lever at the bottom of the vault N times, ejecting N coins and producing a ticket reading "N Ducats, Withdrawal" followed by the date and time. The Royal Accountant collected these tickets and carefully stored them.

When the vault was originally closed and locked, 1,000 Ducats were deposited within it. At the end of every month, the Royal Accountant would produce a report for the King, stating the balance in the treasury. In order to assist him in this task and to prevent underhanded accounting, a Royal Calculating Machine was built. This device was constructed of heavy metal cogs and cylinders, hand cranked. It could be set in either of two modes —addition and

subtraction— but switching between the two modes was a laborious operation requiring replacing one set of gears and cylinders with another. The machine printed out numbers on a long paper tape. When set in the addition mode, a number would be keyed in and the crank pulled. The machine would then print the sum of the initial number on the tape and the keyed-in number. When set in subtraction mode, the machine would print the result of subtracting the keyed-in number from the number on the tape.

The Royal Accountant was only required to calculate the contents of the treasury at the end of the month, so he did not try to keep a running total. Rather, he would do the calculation all at once. Since switching between the addition and subtraction modes on the Royal Calculating Machine was difficult, he would process all the deposits together and all the withdrawals together. And since it was easier to pull the lever when smaller numbers were being manipulated, he processed all of the withdrawals first, and then added in the deposits. Thus the method of royal bookkeeping.

At the end of one month, when he came to get the monthly report, the King glanced at the long tape coming out of the Royal Calculating Machine. The tape started with the contents of the treasury at the end of the last month, and then showed a long series of decreasing numbers. At one point —right after the number 4 had been printed— the ink used to print the digits changed from black to red, and the value of the numbers started increasing. At a later point, the numbers started decreasing again, approaching zero, then the digits switched back to black ink and increased until it reached the final total: 783 Ducats. This was the amount the Royal Accountant announced as the present contents of the treasury.

The King was puzzled, and requested an explanation. The Royal Accountant outlined the procedure, adding that the digits printed in red represented "negative numbers".

This greatly troubled the King. He had vaguely heard about negative numbers, but had dismissed them as purely abstract mathematical fictions with no possible application in the real world. After all, he reasoned, he understood perfectly well what it was to have 1,000 Ducats in the treasury, or 50 Ducats, and even what it meant to have 0 Ducats. That would mean he had no Ducats at all, the treasury was empty, and the kingdom was broke. But what could it even *mean* to say that there were negative 8 Ducats in the treasury? What could a "negative Ducat" *be*? How could one have less than nothing, and different amounts of less than nothing? What could all of this

mathematical nonsense possibly have to do with how rich he was? As the King asked these questions of the Royal Accountant, his glance kept alighting on the Royal Executiener. The Royal Accountant took note of that and felt a light sweat break out on the back of his neck.

The Royal Accountant started talking rather rapidly. The King should take no notice of the negative numbers that appeared in the course of the calculation: they were a mere *bookkeeping device*. They did not actually represent *anything at all* in the physical world or ever characterize the actual contents of the treasury. The number at the start of the calculation represented something physically real —the contents of the treasury at the start of the month— and the number at the end represented something physically real —the contents at the end of the month. But the intermediate numbers in the calculation often didn't represent anything at all. If the first transaction of the month was a deposit, then *none* of the intermediate numbers, whether black or red, might correspond to the actual contents of the treasury at any time. All of the intermediate numbers could be, as it were, fictional.

This talk of fictions being used in the process of accounting for the royal wealth did not assuage the King. Just the opposite: he looked toward the Royal Executiener even more often, and the Royal Accountant started to sweat even more prodigiously.

The Royal Accountant soldiered on. He took out the slips of paper that had accumulated during the month and explained that *these* numbers were not at all fictional: they all represented something with straightforward factual content. The slip reading "9 Ducats, Deposit, June 4, 8:45 PM", for example, represented the very real act of putting 9 Ducat coins into the vault at precisely that time, and "12 Ducats, Withdrawal, June 5, 7:23 AM" represented the equally real act of taking out 12 coins the next morning. *Those* numbers were not fictional in the least, and the totality of the slips of paper represented every single monetary change that happened to the vault during the month, what happened and when it happened. So the initial count of coins in the vault represented something real, and each slip of paper represented something real, and the collection of slips of paper provided a *complete* account of what happened with respect to the monetary contents of the vault through the whole month. The only fiction that arose was in the calculating process since the tickets were not tabulated in chronological order. If one really wanted a complete *history* of what happened one would have to add or subtract the numbers in the time order indicated on the tickets, not doing all

the withdrawals followed by all the deposits. And if one did the calculation *that* way, the numbers would always be printed in black ink: no "negative Ducats" would ever be required.

However, the Royal Accountant went on, it is a purely *mathematical fact*, proven by purely abstract mathematical means, that the results of doing the calculation in proper chronological order and doing first the subtractions and then the additions will be the same. In that proof the appeal to the mathematical properties of negative numbers is completely benign, because no one is claiming that those numbers represent anything actual or physical or monetary. Pure math is pure math, not physics. As long as the calculations are guaranteed mathematically to give the same result it is a matter of pure convenience how it is done, and it happens to be more convenient to do it out of chronological order and using negative numbers. At the end of the day, the Royal Accountant insisted, the number arrived at for the contents of the vault is correct. That number does represent the actual state of the vault.

The King listened to the explanation, but being of a rather skeptical and suspicious nature (especially as regards his own wealth) he commanded his daughters to come and open the vault. It took a full day for the message to be delivered and the royal processions to be arranged. The Royal Executioner spent the time sharpening his blade. The Royal Accountant, despite his mathematical acumen, felt distinctly queasy.

When the Princesses finally arrived, the vault was unlocked and the number of Ducats carefully counted. The result was... 783 Ducats. The Royal Accountant was visibly relieved and the Royal Executioner —behind his mask— seemed a bit crestfallen.

Pause for Commentary

The fable of the Royal Treasury is, so far, intentionally trivial and unsurprising. Everyone understands what is really going on, and why the calculational technique works, and why the intermediate stages of the calculation with "negative Ducats" is unproblematic. But there are several comments and observations apposite at this point.

One is that although in a certain sense the calculation done by the Royal Accountant passes through "unrealistic" mathematical steps —steps where the mathematical object radically fails to represent *anything* physically real or *any* actual state of the vault— the predictive success of the procedure is hardly a "miracle". Its success is

both a plain observable matter of fact and something completely explicable, despite the use of "unrealistic" mathematical elements. The explanation comprises two steps: first a *physical* argument, which deploys mathematical representations that are to be taken realistically (the contents of the vault calculated in proper chronological sequence), and then a *purely mathematical* argument that the outcome of a different calculation, done via different intermediate mathematical steps, will give the same (or effectively the same) result.

This same two-step explanation of empirical success is deployed throughout physics. Newtonian gravitational theory is extremely (although not perfectly) empirically successful. One *physical* account of gravitation, a variant of that deployed by Newton himself, postulates Newtonian Absolute Space and Absolute Time, as well as a physically real gravitational potential and forces created by that potential operating on bodies via $\mathbf{F} = m\mathbf{A}$. From the perspective of the General Theory of Relativity, the entirety of that proposed physical ontology is incorrect: there is no Absolute Space nor Absolute Time, there is no gravitational potential, and $\mathbf{F} = m\mathbf{A}$ is not a law of nature that plays a role in producing gravitational effects. Nonetheless, the General Relativist hardly regards the empirical success of Newton's theory as a *miracle*. Rather, the General Relativist proposes a completely different ontological account of what is going on in terms of a Lorentzian space-time manifold of variable curvature, and then proves a purely mathematical result showing that under certain defined conditions (relative speeds well below that of light, small curvatures or curvature densities) differences between the predictions of the General Theory and those of Newton's theory will be small (below empirically discriminable scale). From this standpoint, even though the central theoretical terms deployed by Newton fail to refer, the empirical success of Newton's theory is perfectly scientifically explicable.

In short, what Putnam characterizes as the "positive argument for realism" is nonsense.

Scientific realism with respect to some particular theory (there is no such thing as "scientific realism" with respect to "science in general") is the view that presently available empirical evidence warrants a high credence that the central terms in that theory refer to real items, and that the central principles or laws are at least approximately correct. What that belief amounts to in detail will vary from theory to theory. And the reasonableness of the high credence depends critically on the structure of the theory, the space of theoretical alternatives, and the quality of the evidence. In antiquity,

it was an open question whether water was homoiomerous (so that
every quantity of water could be subdivided into smaller quantities
of water) or atomic (in the sense of there being a smallest quantum of
water that cannot be subdivided into parts that are both water). That
question has been settled. Water, in this sense, is "atomic" and our
grounds for accepting that are empirical. But understanding in full
detail the scope and quality of empirical evidence in that case may
shed no light at all on other cases. What is the strength of presently
available evidence that quarks are really some sort of stringy object
and that space-time has multiple compactified dimensions? Whatever
answer one gives to that question, the considerations have little to no
overlap with the water example.

So I have no brief to defend "scientific realism" in general, what-
ever that might mean. I do regard any serious doubts about the claim
that water is H_2O and hence not homoiomerous as basically lunacy.
But —unlike Putnam— I think it can be completely non-miraculous
that a theory with non-referring central terms be empirically quite
successful. The successful calculation which uses "negative Ducats"
is just a particularly transparent example.

Although the appeal to negative Ducats may seem fanciful, the
example is just a slightly amplified version of the same situation as
used by Richard Feynman to soften us up for a discussion of negative
probabilities:

> It is usual to suppose that, since the probabilities of events must be pos-
> itive, a theory which gives negative numbers for such quantities must
> be absurd. I should show here how negative probabilities might be
> interpreted. A negative number, say of apples, seems like an absurdity.
> A man starting a day with five apples who gives away ten and is given
> eight during the day has three left. I can calculate this in two steps:
> $5 - 10 = -5$ and $-5 + 8 = 3$. The final answer is satisfactorily posi-
> tive and correct although in the intermediate steps of the calculation
> negative numbers appear. In the real situation there must be special
> limitations of the time in which the various apples are received and
> given since he never really has a negative number, yet the use of neg-
> ative numbers as an abstract calculation permits us freedom to do out
> mathematical calculations in any order simplifying the analysis enor-
> mously, and permitting us to disregard inessential details. The idea
> of negative numbers is an exceedingly fruitful mathematical invention.
> Today a person who balks at making a calculation in this way is con-
> sidered backward or ignorant, or to have some kind of mental block. It
> is the purpose of this paper to point out we have a similar strong block
> against negative probabilities. By discussing a number of examples, I

hope to show that they are entirely rational of course, and that their use simplifies calculation and thought in a number of applications in physics. (Hiley and Peat 1987, pp. 235–236)

It is worthy of note that Feynman makes the remark about the restriction on the times at which the man is given and gives away the apples. Once again, the reasoning is predicated on the existence of an accurate (with respect to time order) account of the entire transaction that never adverts to negative apples. It is once that is appreciated that the person who balks at the use of negative numbers as a calculational convenience becomes regarded as ignorant or backward.

Another familiar place in classical physics where a purely fictitious mathematical "bookkeeping device" is introduced for calculational convenience is Maxwellian electrodynamics. There, the fundamental ontology of the theory is postulated to be the electric and magnetic fields, represented by the mathematical vector fields $\mathbf{E}$ and $\mathbf{B}$. Given the mathematical structure of Maxwell's equations, it follows that one can introduce as mathematical conveniences the scalar and vector potentials ϕ and $\mathbf{A}$, so chosen that $\mathbf{B} = \nabla \times \mathbf{A}$ and $\mathbf{E} = -\nabla \phi - \frac{\partial A}{\partial t}$. It is again a purely mathematical fact that if $\mathbf{E}$ and $\mathbf{B}$ obey Maxwell's equations, such a pair of $\mathbf{A}$ and ϕ can be found. Indeed, an infinitude of such pairs exist, creating extra "gauge" degrees of freedom in the specification of $\mathbf{A}$ and ϕ that don't —from the point of view of Maxwell's ontology— correspond to physical degrees of freedom in the system itself.

Those non-physical degrees of freedom are precisely the source of the *convenience* of calculating in terms of $\mathbf{A}$ and ϕ rather than $\mathbf{E}$ and $\mathbf{B}$: the choice of gauge can be made to harmonize with symmetries of the problem, making calculation simpler. And while $\mathbf{A}$ and ϕ do provide a *complete* mathematical representation of the electromagnetic situation according to Maxwell's ontology (because they fix the fundamental ontology), the unphysical gauge degrees of freedom allow the potentials to behave in ways that would be worrisome if *all* of their mathematical degrees of freedom represented physical characteristics of the system. As John Bell put it in "The Theory of Local Beables":

The word 'beable' will also be used here to carry another distinction, that familiar already in classical theory between 'physical' and 'non-physical' quantities. In Maxwell's electromagnetic theory, for example, the fields $\mathbf{E}$ and $\mathbf{B}$ are 'physical' (beables, we will say) but the potentials $\mathbf{A}$ and ϕ are 'non-physical'. Because of gauge invariance the same

physical situation can be described by very different potentials. It does not matter that in Coulomb gauge the scalar potential propagates with infinite velocity. It is not really supposed to *be* there. It is just a mathematical convenience.

One of the apparent non-localities of quantum mechanics is the instantaneous, all-over-space, 'collapse of the wave function' on 'measurement'. But this does not bother us if we do not grant beable status to the wave function. We can regard it simply as a convenient but inessential mathematical device for formulating correlations between experimental procedures and experimental results, i.e. between one set of beables and another. Then its odd behavior is as acceptable as the funny behavior of the scalar potential in Coulomb gauge. (Bell 2004, pp. 52–53, changing Bell's '**H**' for '**B**'.)

Bell's way of describing the situation is a bit inexact in terms of cleanly separating the physical ontology from the mathematical representation of it (his use of '**E**' and '**B**' is ambiguous between the postulated physical fields and the mathematical representations of them), but the point is perfectly clear nonetheless. Maxwell postulates two physical fields as part of the ontology, which are "directly" represented by the mathematical vector fields **E** and **B**. They are only indirectly represented by the mathematical objects **A** and ϕ, which have non-physical degrees of freedom and are "convenient but inessential" in the sense that any calculation could, in principle, be carried out using only **E** and **B**. If **E** or **B** were to "propagate instantaneously" that would unavoidably (according to Maxwell) represent a real physical item doing so. Not so for **A** and ϕ. But of course, one would have to check that instantaneous changes in **A** or ϕ did not imply similar changes in the electric or magnetic fields.

To recap, we are not at all obliged to take every bit of mathematical structure used to formulate and calculate in a physical theory "seriously" as a more-or-less "direct" representation of a physical entity. Some of the mathematics can benignly be regarded, rather, as a "mere calculational convenience" or a "bookkeeping device" or "surplus structure". In the case of the scalar potential, regarding it as a mere calculational convenience relieves of us worries about its instantaneous propagation in Coulomb gauge. In the case of Feynman's negative apples or our negative Ducats, we are relieved of having to account for what these mathematical items *could possibly* represent. The notion of having less than zero apples need not be given any ontological significance at all. Still, talk of them can be convenient for calculational purposes.

Of course, all of this pleasant mathematical convenience must be *paid for* by honest ontological toil. It is not enough to merely deny "reality" or a role as a "direct representation" to the scalar potential or the negative apples and leave it at that. That sort of attitude would leave the empirical success of the calculational technique truly unexplained or miraculous. Rather, to complete the explanatory task of science, one must postulate some real ontology, governed by some real physical laws, and then show how calculations using the fictive mathematical items will yield (nearly) the same predictions as calculations using the "direct" representations of the beables. In the case of the apples and the Ducats and the electro-magnetic field, we know how to do this. The sense in which someone objecting to negative apples or superluminal scalar potentials is "ignorant" or "backward" or "has a mental block" is simply in not appreciating the force of these rather simple mathematical arguments.

But no matter how simple the justificatory arguments for the mathematical conveniences are, they must be made. And in order to be made, one needs some postulates about what the real ontology of the theory is, and which bits of the mathematical formalism play the role of "directly" representing them. The $64,000 question is exactly how one arrives at this base of more direct representations from which the legitimizing arguments for the mathematical conveniences can be launched. Given *just* the mathematical formalism, how does one sort out the physical-ontology-representing wheat from the mere-mathematical-convenience chaff?

The answer to that question is: "You can't". No amount of study of the mathematical formalism of Maxwellian electrodynamics will yield the conclusion that the **E** and **B** vector fields more directly represent the physical ontology (beables) and the **A** and ϕ fields do so only more indirectly, with additional gauge degrees of freedom. There is absolutely nothing in the mathematical formalism that prevents regarding **A** and ϕ as the more direct (and more complete) representations of the physical ontology while **E** and **B** would represent merely derivative and incomplete generic characteristics of the fundamental fields. Even adding a theory of what is empirically accessible according to the theory would not settle the issue: granting that **E** and **B** completely characterize the *observable* features of the situation, nothing prevents one from regarding the unobservable parts of **A** and ϕ as real but empirically inaccessible rather than fictitious and merely gauge.

In order to determine what —according to a physical theory— is physically real one needs more than just the mathematical formalism:

one needs *the entire theory*. A clearly articulated physical theory in
mathematical physics is a mathematical formalism accompanied by
what I have elsewhere called a *commentary* (Maudlin 2018). The
commentary specifies what the fundamental ontology of the theory
is, what the derivative ontology is, and how the mathematical items
are to be understood as representing each. The ontology of the theory
is not discovered by some abstruse mathematical analysis of the for-
malism but rather by a straightforward reading of the commentary.

Of course, actual physical theories may not be clearly articulated
in this way. They may, for example, consist in a mathematical for-
malism and a *partial interpretation*: a specification of some of the
ontology and how it is represented, of some mathematical degrees
of a freedom that are to be regarded as merely gauge and non-
representational, and a grey area where no firm commitment is made
one way or the other. As a simple example, a quantum theory may
use a standard wavefunction in its mathematical formalization, and be
committed to the physical existence of a quantum state that the wave-
function represents, but be agnostic about whether the physically
possible quantum states correspond to vectors in the Hilbert space of
wave functions or to rays in that space (i.e. to elements of projective
Hilbert space). A lot of physics can be done without resolving this
particular question, but in an ideally formulated theory a decision
will be made one way or the other. If the quantum states are in one-
to-one correspondence with the vectors (or the normalized vectors)
then there will be an unobservable physical feature corresponding
to the overall phase, but if they are in one-to-one correspondence
with the rays there will not.

Only with such a commentary in hand can one begin to address
the question of which mathematical objects are mere "bookkeeping
devices" and which are not, and the concomitant project of justify-
ing the pragmatic usefulness of the bookkeeping devices by showing
how appeal to them is inessential. And one and the same mathe-
matical apparatus conjoined with different commentaries can yield
different ontologies. As an example of how that might play out, let's
now return to our fable. . . .

Once More Upon A Time. . . . .

Although the King was somewhat mollified when the vault was
opened and the count came out correct, he admonished the Royal
Accountant not to use any "negative Ducats" in his calculations go-
ing forward. Since negative Ducats were not merely fictional but

conceptually incoherent, the King said, they had no place in any proper accounting of the royal treasury.

The Royal Accountant adjusted his procedure in the most convenient way to meet the King's demand. Instead of subtracting off all the withdrawals first, he added in all the deposits and then subtracted the withdrawals. The Royal Calculating Machine only had to be changed in mode of operation once, and by this method all of the printed figures on the tape would be black.

At the end of next month, the King again came for the account. The numbers on the tape began with 783 Ducats and rose monotonically up to 2,834 Ducats before falling monotonically to the final number of 634. At first the King was pleased when he saw only black numbers, but then he became puzzled again.

"2,834 Ducats? My goodness, I had no idea that the Royal Treasury ever got so well-stocked last month!"

The Royal Accountant explained that at no time did the vault ever contain 2,834 Ducats. That number, he said, was just as fictional as the negative Ducats that appeared in the old method of calculation. In fact, none of the numbers on the tape save the first and the last corresponded to the contents of the vault at any time. They were all, in that sense, fictional. Still, the result at the end was sure to be right.

The King again became perplexed. "What you call a fiction I call by its proper name: a lie. You are telling me that none of the steps in this calculation actually correspond to anything that ever happened, but nonetheless I should be confident that the end result is correct. This strikes me as a very suspicious way to proceed. I don't like the idea of my accounts being filled with lies. I want to see all of this done properly, never departing from the truth, and showing precisely what was the contents of the vault all through the last month."

The Royal Accountant was disheartened and a bit vexed by the King's obtuseness, but there is no arguing with a monarch so he set about the rather tedious task of putting all the slips of paper in proper chronological order, resetting the Royal Calculating Machine back to 783, and then processing the slips in order. He had to change the machine from addition mode to subtraction mode and back again dozens of times, and the entire operation was long and exhausting. But when he processed the last of the slips he triumphantly showed the King the final total: 634 Ducats.

His feeling of triumph, however, was short-lived. Scanning back over the long tape, the King noticed that right about the middle the

numbers were again printed in red. It was only a small stretch, and
soon returned to black, but there it was. The King regarded the
Royal Accountant with a look of suspicion mixed with anger, and
called the Royal Executioner to have him bound. Certain he was
somehow being cheated, the King again sent messages out to his
daughters to return to the palace.

The next day, the Princesses all arrived and a large crowd assem-
bled to witness the opening of the vault. The Royal Accountant, still
bound, was pale and sweaty, and the Royal Executioner passed the
time honing his blade. The vault was opened and the coins metic-
ulously counted. Result: 634 Ducats. There was a palpable sense of
disappointment in the crowd as the Royal Accountant was freed of
his bonds. The King called a meeting of all of his advisors.

The first commanded to speak was the Royal Accountant. Could
he explain just what was going on?

The Royal Accountant said that he was just as puzzled as everyone
else, but insisted that the results proved that he had not embezzled a
single Ducat. On reflection, he said, the problem certainly had to be
in the system that printed out the tickets. Something must have gone
wrong with the time stamp operation: the amounts of the deposits
and withdrawals were right, but the order must have somehow gotten
misrepresented. That mistake then accounted for the appearance of
the negative numbers in the calculation. That, he said, is the only
possible explanation. Whether the malfunction of the time stamp
was a mere blunder or rather intentional he could not say. That
problem was not his doing, but rather the responsibility of the Royal
Instrument Maker, who designed and built the device.

Next to speak was the Royal Instrument Maker. She asserted that
there was nothing at all wrong with the machine, and the time
indications and amounts of the deposits and withdrawals on the
tickets were accurate. The explanation, she ventured, was not that
any of the information on the tickets was inaccurate but that the
set of tickets was incomplete. At some point in the last month, the
Royal Accountant must have added some Ducats to the vault and
then later withdrawn the same number. The tickets recording these
extra deposits and withdrawals had then been lost, or more likely
intentionally destroyed. The Royal Accountant's motivation for this
behavior she would not speculate on.

As the Accountant and Instrument Maker glared at each other,
the Minister of Propaganda offered an opinion.

"Sire", he said, "I think you are getting yourself all worked up
over nothing. The whole point of the accounting system is to tell you

how many Ducats you have at the end of each month, and it has not failed in that task. When I sit down to write a speech for you, I know that my objective is to sway the audience one way or another. My only consideration in composing the speech is achieving that end: whether what I write is true or false, accurate or a fiction, is neither here nor there. The only grounds you could have to be disappointed with my work is if it fails to move listeners as you wish it should. Don't worry your head about *how* it works, just be satisfied *that* it works. Take that attitude and all of your worries here will dissolve."

The King looked askance at the Minister of Propaganda. "My good sir, I fully appreciate your attitude toward your job, and that the very question of truth versus fiction does not enter into it. But what we are talking about here is *my money*, and I find your insouciance about how much of it I have at any given time. . . disturbing." The King signaled with his eyes to the Royal Executioner, who escorted the Minister of Propaganda from the room. He did not return.

Next the King called on the Royal Oracle. She slowly stood and gazed into the middle distance for half a minute, and then intoned: "There are more things in heaven and earth, Sire, than are dreamt of in your philosophy." She resumed her seat.

The King muttered "I do wish you would be a bit less. . . oracular", but he knew from experience that he would get nothing more from her.

The last to speak was the Jester. "It's all very simple, King. When the vault is closed, and no one can see inside, *there are no Ducats at all*. When the door is opened, they suddenly pop into existence, and in the right numbers according to the accounting. The tickets are all correct, and they are all the tickets there are."

The rest of the company had a good laugh at the Jester's joke.

"Alright", said the King, "let's get down to serious business. Since there seems to be no money missing, I cannot in all propriety order the execution of the Royal Accountant. But we have to get to the bottom of this. First, we will have a second clock made against which the times stamped on the tickets can be checked, and the number of Ducats deposited or withdrawn in each case will also be verified against the printed ticket by a second person." At this the Royal Accountant smiled. "Next, both the vault and the Royal Accountant will be put under 24-hour watch to make sure no deposits or withdrawals are being made surreptitiously." At this the Royal Instrument Maker smiled. "Finally, instead of doing the accounting only at the end of the month, the Royal Accountant will

keep a running total after each transaction." At this, the Royal Oracle smiled. And so it was decided.

The next month started with 634 Ducats in the vault, and the number on the running tally fluctuated as Ducats were deposited and withdrawn. Since the taxes were generally collected toward the end of the month, withdrawals outnumbered deposits early on, and the total drifted ever lower and lower. Finally, after a particularly large withdrawal, the total on the tape printed out in red: negative 2 Ducats. The Royal Accountant froze in his tracks. The alarm was sounded.

The King and all his ministers gathered, looking from the red number on the tape to the vault and back again. After a brief consultation, they decided there was nothing to be done but to call back the four Princesses again and see what was inside the vault. The vault was impounded and the Royal Accountant and Royal Instrument Maker both put under house arrest for good measure.

The next day, the crowd assembled again. The Princesses carefully entered the combination, the King pulled at the door and. . . the door would not open. No matter how much force was tried, it remained firmly in place. Levers and compound pulleys were applied to no effect.

Eventually, in an act of sheer desperation, the Royal Accountant made a peculiar suggestion. He proposed depositing two Ducats into the vault and then trying to open it. Two Ducats were inserted into the slot, and forthwith the door became easy to open. But when opened, everyone was further shocked to find the vault completely empty.

Further experimentation revealed what one might now suspect. More Ducats could indeed be withdrawn from the vault than had been put in, but whenever the state of the vault was "negative Ducats" the door could not be opened. It would only open when enough coins had been deposited to bring the balance back to zero, at which point the vault would be found empty.

These "Ducat rules" were simple to state, but quite a lot of theorizing and speculation ensued. Some philosophers posited "anti-Ducats" which could combine with regular Ducats causing mutual annihilation. Specific methods for the annihilation were debated. Others insisted that no such detailed structure needed to exist: all that was required was the recognition of distinct "negative Ducat" states of the vault, together with rules about how those states interacted with the deposit and withdrawal slots. These philosophers considered the

"negative Ducat states" of the vault to be physically fundamental, admitting of no finer description. The "anti-Ducat" theorists found these explanations unsatisfyingly austere, while their rivals thought the postulation of anti-Ducats to be extravagant. Occasionally, someone would propose a theory of the internal goings-on in the vault that would entail some limits on the number of Ducats that could be withdrawn from the empty vault, or suggest a specific time lag between adding enough Ducats to bring the total to zero and being able to open the door. From time to time, these various hypotheses would be put to empirical test.

But despite the various theoretical disputes, several propositions were accepted by all. One was that the tickets printed out all accurately reflected the number of Ducats deposited or withdrawn and the time of the withdrawal. They agreed that there had been no surreptitious deposits or withdrawals the month before.

And they all agreed that there was more in heaven and earth than had been dreamt of in their philosophies.

Morals

There are several morals that can be drawn from our little tale. They all have application in discussions of the ontology underlying quantum phenomena.

The first is that the "reality" or otherwise of the wave function —the mathematical gadget used in the mathematical formalism of quantum theory— is not an interesting question if taken literally. The wave function is a mathematical item, and has the same sort of existence (or non-existence) as any other mathematical object. One can discuss the ontology of mathematics, but that has nothing at all to do with physical ontology, and so is completely beside the point.

The real question is what —if anything— in the physical world does a certain piece of mathematics *represent* and which features of the mathematical object are representational as opposed to "mere gauge". It is possible to maintain that strictly speaking a piece of mathematics represents *nothing physical at all* and is a "mere bookkeeping device". That is the original attitude of the Royal Accountant to the "negative Ducats" that appear in the course of his calculation. The Royal Accountant is initially not a "scientific realist" *at all* with respect to those particular mathematical items: they represent no actual, or even possible, physical state of the vault. Nonetheless, the Royal Accountant does not regard the accuracy of the calculations he does using those mathematical items as in the least "miraculous" or

"inexplicable" since he had (he believed) a perfectly clear mathematical explanation of why the calculations work, which explanation does not require the existence of "negative Ducat" states. Of course, upon further investigation that explanation fails in the face of the newly revealed phenomena. In order to account for those new phenomena he —and everyone else— is forced to take a quite different attitude to the "negative Ducat" states.

The example of the scalar and vector potentials in Maxwellian electrodynamics makes a different point. There is a sense in which those mathematical items represent something real and a sense in which they do not. Given Maxwell's ontology, they represent something real by indirectly representing the electric and magnetic fields, which are part of the physical ontology. But at the same time, there are gauge degrees of freedom in the representation, so not every aspect of the mathematical object corresponds to something physical. When Bell says that we need not be concerned with the instantaneous propagation of changes in the scalar potential in Coulomb gauge, he tacitly adverts to the fact that we can legitimately regard that as the instantaneous propagation of a merely gauge degree of mathematical freedom rather than a physical degree of freedom. But as with the insouciance about negative Ducat states, this attitude requires justification by reference to a clearly articulated physical ontology, in this case the electric and magnetic fields.

The justificatory argument could, for example, point out that in Lorenz gauge neither the scalar nor vector potential propagates faster than light. Since those potentials provide a *complete* representation of the electro-magnetic situation (relative to Maxwell's ontology), that means that nothing physical propagates faster than light. It follows —logically follows in this setting— that any mathematical degree of freedom that appears to propagate faster than light *must* be merely gauge. Or, one could prove the result more directly by doing all the calculations only using more direct representations $\mathbf{E}$ and $\mathbf{B}$ of the electric and magnetic fields, avoiding the scalar and vectors fields altogether, just as the Royal Accountant believed he could in principle always avoid mention of negative Ducat states.

These sorts of "defanging" arguments, though, can only be launched from the base of a proposed physical ontology. If that ontology itself is brought into question, then the entire situation has to be re-examined.

In the case of the scalar and vector potentials, such a re-examination was forced by the discovery of the Aharonov-Bohm effect. Under

two ontological/nomic assumptions —first, that the basic electromag-
netic ontology is the electric and magnetic (or electro-magnetic) field
and second that electrons cannot be affected by the electromagnetic
field in regions from which the electrons have been shielded (a lo-
cality assumption)— it follows that electromagnetic changes inside a
shielded solenoid cannot produce any change in behavior of electrons
restricted to remain outside. But just as the observable behavior of
the vault and Ducats in our story refute the naïve account of what is
going on, the observable Aharonov-Bohm effect refutes this pair of
postulates. Maxwell's ontology cannot be maintained.

What should replace Maxwell's ontology, however, is not immedi-
ately obvious. One can, of course, go whole hog and come to regard
the scalar and vector potentials as "direct" representations of the
ontology, and what used to be thought of as gauge transformations
as real physical changes of real degrees of freedom. As a slogan, that
would be to regard **A** and ϕ as "real" (i.e. as representations of real
physical ontology with fewer gauge degrees of freedom than had been
realized).

But there are other intermediate possibilities. For example, one
can regard the more "direct" mathematical representation of the
physical situation to be provided by the connection on a fiber bundle,
and can regard the **E** and **B** fields as representations of the *curvature*
of that bundle. Since there can be connections that differ outside
the solenoid without differing in their curvature (because the region
outside the solenoid is not simply connected) one is able to explain
the observable phenomena while maintaining the locality of both
the fundamental ontology and the laws. The **E** and **B** fields alone
underrepresent the physical situation where they are defined, and
the **A** and ϕ fields *overrepresent* it, still containing gauge degrees of
freedom. The connection of the bundle is the Goldilocks piece of
mathematics that most directly corresponds to the physical situation.
From the basis of this ontology, the shortcomings of both **E** and **B**
and of **A** and ϕ are comprehensible and their practical utility (and
limits of practical utility) explained. Once again, one can deny, in a
certain sense, being a realist about either pair without falling afoul
of any need to invoke miracles to explain their utility.

Until the discovery of the unexpected phenomenology in our fable,
the *calculational use* of "negative Ducat" states could be accepted
without any ontological qualms, and similarly for the use of the
potentials in Maxwell's theory before the discovery of the Aharonov-
Bohm effect. And in each case the new phenomenology demands a

new understanding of the situation, or at least a rejection of the old understanding. The precise representational character of the mathematical items becomes up for grabs.

The points illustrated by our fable —and by the concretely similar case of electromagnetic theory— may seem trivial and obvious. Perhaps so. But even the trivial and obvious sometimes needs to be explicitly pointed out in order to be appreciated. Without these basic points being universally acknowledged, any discussion of the representational character (or lack of it) of the wave function in quantum theory is bound to descend into chaos.

And as one last illustration of the importance of the conceptual points being made here, consider the case of *virtual* (or *off-shell*) particles. Some discussions of the ontology of quantum theory are rife with references to "virtual particles". We are told that they are constantly "popping in and out of existence"; that they form swarms surrounding "real" charged particles that shield the charge by "polarizing the vacuum"; that it is the "exchange" of virtual particles that mediates "forces" between real particles; that each Feynman diagram illustrates a possible (or actual?) sort of interaction among virtual and real particles. We are told that the spontaneous production of pairs of virtual particles in a vacuum plays an essential role in the production of Hawking radiation: when one of the pair falls through the event horizon it can produce enough "negative energy" to boost its outside partner up from being "virtual" to being "real". Hawking himself told exactly this story on multiple occasions, and it is regularly repeated as if it were being presented as a serious explanation of an (in principle) observable phenomenon. But on the other hand, Daniel Harlow remarks:

> Although I won't use it in what follows, there is a heuristic explanation of Hawking radiation that is occasionally brought up. The idea is that entangled pairs of particles are constantly jumping into existence near the horizon via vacuum fluctuations, and sometimes one of them falls in and one of them gets out. This cartoon has several problems if it is taken literally, among them that the "particles" involved have wavelengths comparable to the size of the black hole and that the Hawking process isn't really stochastic, and in my view it tends to create more confusion than it resolves. (2016, p. 15)

A pure instrumentalist about the theory will see nothing here to discuss: "sums of Feynman diagrams" has a clear enough *mathematical* meaning, and if a theory predicts Hawking radiation (as an empirical

phenomenon) then the only question is whether it in fact occurs, not the correct physical account of it. The pure instrumentalist is, of course, represented in our parable by the Minister of Propaganda, and we bade him farewell some time ago. The physicist as physicist, in contrast, should care. If Harlow is right, then Hawking's proposed explanation is wrong, and any sense of understanding one derives from reading it is an illusion. This is not a question about which one should feel neutral.

I do not mean to delve further into the issue of "virtual particles" and their relation to mathematical structure and ontology here. It is a complicated and controversial issue. My only point is that it is a real, substantial question of physical ontology. If, in a proper sense, there really are virtual particles, then there is more in heaven and earth that was dreamt of in even Bohr's philosophy. If there aren't, then we are owed an explanation of the mathematical fruitfulness of acting as if there were, and a clear account of the limitations of "explanations" that appear to advert to them. Either way there are ontological lessons to be learned, or at least ontological disputes and possibilities to be clarified.

Finally, let's return to the original question about the wavefunction as it appears in the quantum formalism. The best way to formulate the central ontological question concerning it was:

> What, if anything, does the wavefunction represent about the physical system to which it is ascribed?

The possibility that it represents nothing at all would be analogous to the Royal Accountant's initial attitude toward "negative Ducats" appearing in his calculational scheme: there simply are no corresponding physical states of the vault, but the mathematical objects can play a useful role in a calculational scheme nonetheless. That seems like a tremendously implausible attitude to take to the wavefunction. Just as there are physical phenomena that convince everyone in our fable that the "negative Ducat states" of the vault are real ontological states, so there are both phenomena and arguments indicating the physical reality of something whose characteristics are somehow encoded in the wavefunction. First among these, of course, are interference effects such as the two-slit interference pattern, that indicate the existence of something sensitive to the presence of both of the open slits. The use of a dynamical wave equation —the Schrödinger equation— in producing the predictions is also telling.

Unlike the initial appearance of "negative Ducat" states in the calculation done out of chronological order, the Schrödinger equation generates the calculational states in proper chronological order. They are therefore candidates for representations of the successive physical states of the system. Finally there is the PBR theorem, which shows (given the natural assumption of statistical independence between created states and later experiments done on them) that the wavefunction cannot be "merely epistemic" in the sense that the same ontological state can properly be assigned different wavefunctions.

The wavefunction represents *something* physically real about a system. But that still leaves us two questions.

> First, is it a *complete* representation of all the physical characteristics of the system?

> Second, how *direct* a representation is it? That is, how many of the mathematical degrees of freedom correspond to physical degrees of freedom and how many are merely gauge, as illustrated by the scalar and vector potentials in Maxwell's theory?

The completeness of the wavefunction is an open question. Both Objective Collapse and Many Worlds ontologies can maintain that the wavefunction is complete. Pilot Wave theories cannot. Insofar as these are live theoretical options, the question is unsettled.

Similarly, the separation of the wavefunction into physically-representing and merely-gauge degrees of freedom is an open question. It is widely accepted that a more direct mathematical representation of the quantum state —one with fewer gauge degrees of freedom— is a ray in a Hilbert space rather than a vector, i.e. that the set of possible physical states has the structure of a projective Hilbert space rather than a vector space. Nonetheless, vectors are more frequently used as mathematical representations of quantum systems than rays are. Dynamical equations like the Schrödinger equation are typically written for vectors rather than rays. Just as the discussion of the ontology of electrodynamics was invigorated by considering, in addition to the pair $\{\mathbf{E}, \mathbf{B}\}$, the pair $\{\mathbf{A}, \phi\}$, and the electromagnetic field tensor $F_{\mu\nu}$, as well as mathematical objects like the connection on a fiber bundle, so also the discussion of the ontology of the quantum state can benefit by a more open and extensive consideration of alternative mathematical structures that might be used to represent it.

But without a clear distinction between mathematical objects on the one hand and their various representational and non-representational uses on the other, such a clarification cannot even begin. Our

simple fictional tale of the negative Ducats can serve to remind us of the various possibilities that might arise. Including those we did not at all expect.

REFERENCES

Bell, John, 2004, *Speakable and Unspeakable in Quantum Mechanics*, 2nd Edition, Cambridge University Press, Cambridge.

Harlow, Daniel, 2016, "The Jerusalem Lectures on Black Holes and Quantum Information", *Review of Modern Physics*, vol. 88, no. 1, 15002.

Hiley, B.J. and F. David Peat (eds.), 1987, *Quantum Implications: Essays in Honour of David Bohm*, Routledge and Kegan Paul, London.

Maudlin, Tim, 2018, "Ontological Clarity via Canonical Presentation: Electromagnetism and the Aharonov–Bohm Effect", *Entropy*, vol. 20, no. 6, pp. 465–486.

Putnam, Hilary, 1975, *Mathematics, Matter and Method*, Cambridge University Press, Cambridge.

Tegmark, Max, 2014, *Our Mathematical Universe*, Knopf, New York.