The Distribution Postulate in Bohm's Theory

Jeffrey A. Barrett

ABSTRACT. On Bohm's formulation of quantum mechanics particles always have determinate positions and follow continuous trajectories. Bohm's theory, however, requires a postulate that says that particles are initially distributed in a special way: particles are randomly distributed so that the probability of their positions being represented by a point in any region R in configuration space is equal to the square of the wave-function integrated over R. If the distribution postulate were false, then the theory would generally fail to make the right statistical predictions. Further, if it were false, then there would at least in principle be situations where a particle would approach an eigenstate of having one position but in fact always be somewhere very different. Indeed, we will see how this might happen even if the distribution postulate were true. This will help to show how loose the connection is between the wave-function and the positions of particles in Bohm's theory and what the precise role of the distribution postulate is. Finally, we will briefly consider two attempts to formulate a version of Bohm's theory without the distribution postulate.

1. Bohm's theory

Bohm's theory has long stood as a counterexample to many of the deep philosophical conclusions that people have tried to draw from quantum mechanics.[1] It is a deterministic hidden-variable theory that makes exactly the same predictions as the standard theory of quantum mechanics whenever the standard theory makes coherent predictions.[2] Unlike the standard theory, however, Bohm's theory treats measurements the same way it treats any other physical interaction. Since it is deterministic, probabilities are purely epistemic, but there is also a nice explanation for why there is nonetheless no way to do better than the standard statistical predictions of quantum mechanics.[3] Measurements have determinate outcomes because every particle always has a determinate position and it is assumed that every measurement is ultimately a measurement of position. So while Bohm's theory is referred to as a hidden-variable theory, where the positions of particles are the hidden variables, this is widely recognized as misleading. As Bohm and Hiley put it,

> . . . our variables are not actually hidden. For example, we introduce the concept that the electron *is* a particle with a well-defined position and momentum that is, however, profoundly affected by a wave that always accompanies it . . . . Far from being hidden, this particle is generally what is most directly manifested in an observation (1993, p. 2).

Indeed, since Bohm's theory requires every measurement to ultimately be a measurement of the position of some particle, one might say that the positions of particles are the only variables that are *not* hidden.

As Bohm's theory is usually understood, the world consists of a wave and a collection of particles that always have determinate positions and follow continuous trajectories. The complete physical state at a time is given by the ordered pair $\langle\psi, \S\rangle$, where ψ is the usual quantum-mechanical wave-function and § is the point in configuration space that represents the current positions of all of the particles. The state of the physical world at all times is determined by its complete state $\langle\psi, \S\rangle$ at any particular time and the Hamiltonian $\hat{H}$, which is determined by the energy properties of the world.

Suppose that there are N particles. The evolution of the wave-function $\psi(\mathbf{x}, t)$ is always described by the deterministic linear dynamics. That is,

$$(1) \quad \frac{\partial}{\partial t}\psi(\mathbf{x}, t) = -i\hat{H}\psi(\mathbf{x}, t)$$

where $\mathbf{x}$ is a variable representing a point in $3N$-dimensional configuration space and $\hat{H}$ is the Hamiltonian. Each particle moves in a deterministic way that depends on the wave-function and the position of every particle. The velocity of particle P is given by

$$(2) \quad \mathbf{v}_P = \frac{\mathrm{Im}\, \psi^*(\mathbf{x}, t)\, (\partial/\partial \mathbf{x}_P)\psi(\mathbf{x}, t)}{m_n|\psi(\mathbf{x}, t)|^2}$$

where m_P is the mass of the particle and $\mathbf{x}_P$ represents the three coordinates of configuration space that deter-

Topoi 14: 45–54, 1995.

mine its position. In order to calculate the velocity of P at some time t_1, evaluate the right-hand side of (2) for each component of v_P using $\psi(\mathbf{x}, t_1)$, then substitute the position of every particle at time t_1, $\S(t_1)$, into the resulting expressions.

As ψ evolves according to (1), $|\psi|^2$ evolves just as one would expect if it represented the density of a fluid. On the standard formulation of quantum mechanics $|\psi|^2$ is called the probability density, its gradient is called the probability current, and the "fluid" is called the probability. It is convenient to use the same terminology here even though the final picture is very different from that of the standard theory. Given the notion of a fluid associated with $|\psi|^2$, one might understand the way the two dynamical laws work for a single particle by picturing ψ evolving in the usual linear way in position space and the particle being carried along by the probability current associated with the change in $|\psi|^2$ much as a particle might be carried by an ordinary fluid. And, for many particles, one might picture the point in configuration space that represents the positions of all the particles being carried along by the probability current associated with the change in $|\psi|^2$ as ψ evolves in configuration space. While this many-particle picture may at first seem counterintuitive, it is very useful.[4] It shows how the motion of a particle generally depends not only on its own position but also on the current position of every other particle.

As the probability $|\psi|^2$ flows into a region R of configuration space, the probability of the point representing the particle configuration being carried into R increases – consequently, the probability of finding the particle configuration in R increases. We presumably want the probability of finding the particle configuration in R to be *equal* to $|\psi|^2$ integrated over R. If we have this, then Bohm's theory will make the same statistical predictions as the standard theory for the positions of the particles. If the particles are randomly distributed at some time t_0 in such a way that the probability of the particle configuration being represented by a point in any region R of configuration space is

$$(3) \quad P(R, t_0) = \int_R |\psi(\mathbf{x}, t)|^2 d\mathbf{x}$$

then it follows from the dynamics that the particles always were and always will be similarly distributed with respect to the wave-function, that the probability of finding the particle configuration in any region always was and always will be equal to the integral of $|\psi|^2$ over the region, which is exactly what the standard theory of quantum mechanics predicts whenever it makes a coherent prediction. Since the standard theory has proven itself empirically adequate over the measurements that we have made so far, if every actual measurement really amounts to a measurement of position, then Bohm's theory is also empirically adequate.[5] Conversely, if particles are distributed in some other way, then the probability of finding the particle configuration in a particular region on Bohm's theory would be different from the standard statistical prediction. To the extent that one is convinced that the standard statistical predictions are right, one would expect Bohm's theory to make the wrong predictions if particles are not at some time randomly distributed with respect to $|\psi|^2$. Let's call the claim that the positions of particles at some time were in fact randomly distributed with respect to $|\psi|^2$ the *distribution postulate.* The distribution postulate is usually simply taken to be a part of Bohm's theory.[6] It places a very strong constraint on the relationship between the wave-function ψ and the particle configuration § and plays an essential role in making the theory empirically adequate. Let's consider the role the distribution postulate plays in a bit more detail, then we will consider two proposals for constructing a Bohm-like theory without such a strong constraint on the relationship between ψ and §.

2. Measurement in Bohm's theory

One might measure the z-spin of a spin-1/2 particle P by passing it through a Stern-Gerlach device and observing which way it is deflected. Suppose that P's motion is effectively determined by the single-particle wave-function

$$(4) \quad \psi_P = \frac{1}{\sqrt{2}}(\psi^\uparrow + \psi^\downarrow)$$

where $\psi^\uparrow$ and $\psi^\downarrow$ are the z-spin up and z-spin down components of ψ_P, respectively. Suppose also that $|\psi_P|^2$ is constant in some spherical region R and that it is zero everywhere outside R. Finally, suppose that the distribution postulate is true, that the initial position of P is randomly distributed by the probability distribution $|\psi_P|^2$. Given these assumptions, there is a 50–50 chance that S will initially be in the top half of R and the same chance that it will be in the bottom half of R.

The magnetic field of the Stern-Gerlach measuring device will separate the two z-spin components in position space and deflect $\psi^{\uparrow}$ up and $\psi^{\downarrow}$ down. Suppose that the components are separated in a symmetric way so that P will be deflected up if it begins in the top half of R and down if it begins in the bottom half of R. Note that how P behaves in future interactions will depend on which way it is deflected since the way it is deflected determines which component of ψ_P carries P. P's motion will be completely determined by the probability current associated with $|\psi^{\uparrow}|^2$ if P is deflected up and by the current associated with $|\psi^{\downarrow}|^2$ if it is deflected down. Suppose that P is deflected up. Since $\psi^{\uparrow}$ would be deflected in the z-spin up direction on any subsequent z-spin measurement, so would P. More generally, P would behave just as one would expect a z-spin up particle to behave as long as its motion is determined by $\psi^{\uparrow}$, and its motion will be determined by $\psi^{\uparrow}$ until $\psi^{\uparrow}$ and $\psi^{\downarrow}$ overlap at P's position or until $\psi^{\uparrow}$ is itself split into *its* spin components. Similarly, if P were deflected down, it would behave just as one would expect a z-spin down particle to behave as long as its motion is determined by $\psi^{\downarrow}$. Consequently, it is natural to take the direction that P is deflected to indicate its z-spin.[7]

One might construct a slightly more sophisticated z-spin measuring device where the result of the measurement is determined by observing the position of a pointer particle Q rather than by observing the position of P. Suppose again that P begins in some middle position and that it is either deflected up or down by the measuring device. Suppose further that the measuring device is constructed so that Q might occupy any of three positions 0, -1, and $+1$ corresponding to Q reporting that the measuring device is ready to make a measurement, that the result is z-spin down, and that the result is z-spin up, respectively. Now, how does the position of Q end up accurately representing the direction that P was deflected?

Consider the motion of the point in configuration space that represents the positions of *both* P and Q. This point begins in a region where P is in some undeflected position and Q is at position 0. The measuring device splits the wave-function into two components in configuration space. Whether the position of Q ends up appropriately correlated with the position of P depends on how the measuring device does this, and this is determined by the Hamiltonian. An ideal measuring device would be constructed so that one component of the wave-function (the z-spin up component) is nonzero only in the region of configuration space where P is deflected up and Q is at position $+1$ and the other component of the wave-function (the z-spin down component) is nonzero only in the region where P is deflected down and Q is at position -1. If a measuring device does this, then the positions of P and Q will end up perfectly correlated in the appropriate way. One might think of the point representing the positions of the two particles being carried along by the probability current, and here all of the probability flows into regions where the two particles are correlated in the appropriate way. The point representing the positions of the two particles will not be carried into a region where P is deflected down and Q ends up at position $+1$, for example, because none of the probability flows in that direction in configuration space. Consequently, if the initial positions of the two particles are such that P is deflected up, then Q will end up at position $+1$; and if the initial positions are such that P is deflected down, then Q will end up at position -1.

3. What if the distribution postulate is false?

Suppose that the distribution postulate is false. Suppose, for example, that the particles were initially distributed so that each particle was in the top half of its effective wave-function. In this case, one would get *up* as the result of every z-spin measurement. More generally, if the particles were distributed in any way other than randomly with respect to $|\psi|^2$, then one would expect to get a different statistical distribution of measurement results than that predicted by the standard theory of quantum mechanics. The problem, of course, is that we have very good inductive evidence that the standard theory makes the right statistical predictions whenever it makes coherent predictions.

Another consequence of the distribution postulate being false is that it would generally be possible for a particle to be in an eigenstate of being at one position but in fact be somewhere very different. In order to prevent this from happening, one might stipulate that particles must initially be in a configuration where ψ is nonzero. As we shall see, however, even this would not necessarily prevent particle configuration from eventually ending up in a region of configuration where ψ is zero. Indeed, even if the distribution postulate were true, it would generally be possible for a particle to be in an eigenstate of one position but actually be somewhere else. This is the sense in which the relation between $|\psi|^2$ and the probability of finding a particle in a region is contingent on Bohm's theory.

Suppose that a perfect measuring device M measures the z-spin of each of a series of particles $P_1, P_2, \ldots$, and suppose that particle P_n is initially associated with the single-particle wave-function

$$\text{(5)} \quad \psi_{P_n} = \frac{1}{\sqrt{2}}(\psi^{\uparrow}{}_{P_n} + \psi^{\downarrow}{}_{P_n})$$

where ψ_{P_n} is zero outside region R_n and ψ_{P_n} is such that if P_n were in the top half of R_n, then it would be deflected up, and if it were in the bottom half of R_n, then P_n would be deflected down.

Now consider the evolution of the wave-function associated with the composite system $M + P_1 + P_2 + \ldots$ as M makes its z-spin measurements. Let $\psi[r_1, r_2, \ldots]$ represent the wave-function of the composite system $M + P_1 + P_2 + \ldots$ when M is in an eigenstate of reporting that the first z-spin result was r_1, that the second z-spin result was r_2, etc. It follows from the linear dynamics and the assumption that M is a perfect measuring device that after the first measurement the wave-function would be

$$\text{(6)} \quad \frac{1}{\sqrt{2}}(\psi[\uparrow] + \psi[\downarrow])$$

and after the second measurement the wave-function would be

$$\text{(7)} \quad (1/\sqrt{2})^2(\psi[\uparrow,\uparrow] + \psi[\uparrow,\downarrow] + \psi[\downarrow,\uparrow] + \psi[\downarrow,\downarrow])$$

Generally, after n measurements the wave-function would have 2^n terms, and the coefficient on each term would be $(1/\sqrt{2})^n$ when written in this basis.

One can easily show that the wave-function describing the composite system approaches an eigenstate of M reporting that the z-spin results were randomly distributed as the number of measurements n gets large for any notion of randomness satisfying two plausible-sounding conditions. Every length-n string of possible measurement results will be represented by some term in the wave-function describing $M + P_1 + P_2 + \ldots$ after n measurements, so if one assumes (a) that there are generally many more random finite strings than nonrandom finite strings of a particular length and (b) that nonrandom strings form a measure-zero subset (in a nontrivial product measure) of the set of infinite-length strings, then it immediately follows that the sum of the squared coefficients on the terms that describe M as getting a random sequence of results goes to one as the number of measurements n gets large. Which means that the wave-function describing the composite system approaches an eigenstate of reporting that the z-spin results were randomly distributed as n gets large.

One can also show that the wave-function associated with the composite system approaches an eigenstate of M reporting that the relative frequency of z-spin up results was within ε of 1/2 as n gets large for any $\varepsilon > 0$. For a given n, let $F(n)$ be the sum of the squared coefficients of each of the terms where the ratio of z-spin up results to n is within ε of 1/2. As n gets large, the ratio of the terms that describe a sequence of measurements results where the ratio of z-spin up to n is within ε of 1/2 to the total number of terms goes to one. The reason is that most length-n binary sequences have about the same proportion of $\uparrow$'s as $\downarrow$'s. After an even number of measurements, there are n choose $n/2$ terms where the ratio of $\uparrow$-results to n is exactly 1/2, $2[n$ choose $[(n/2) - 1]]$ where it is within $1/n$ of 1/2, etc. So as n gets large, an ever greater proportion of the terms have ratios of $\uparrow$-results to n ever closer to 1/2. Which means that, as n gets large, $F(n)$ goes to one for any $\varepsilon > 0$ – that is, the sum of the squared coefficients on the terms that describe M as getting a sequence of results where the frequency of $\uparrow$-results was within ε of 1/2 goes to one as n gets large for any $\varepsilon > 0$. And this means that the wave-function describing the composite system approaches an eigenstate of M reporting that the relative frequency of z-spin up results was within ε of 1/2 as n gets large.

Note that these are simply conclusions concerning the evolution of the wave-function of the composite system – now let's consider what an ideal measuring device M would *report* here according to Bohm's theory. Suppose that M has two pointers: one that points to either *random* or *nonrandom* and the other that points to a number between zero and one that represents the relative frequency of z-spin up results. These pointers are made of particles and, since particles always have determinate positions on Bohm's theory, these pointers always point in determinate directions. Which direction they point will generally depend on the position of every particle and the evolution of the wave-function. Suppose that it is possible to construct a measuring device where the positions of these pointers accurately represent the statistical properties of the particles measured.

Suppose that the distribution postulate is true of the measured particles, that particle P_n is initially randomly distributed according to $|\psi_{P_n}|^2$. This means that P_n will be equally likely to be in the top half and bottom half of region R_n, which means that it will be equally likely for P_n to be deflected up and be taken as a z-spin up particle as it will be for it to be deflected down and be taken as an z-spin down particle. In the limit as the number of measurements n gets large, with probability one, the first pointer will end up pointing at *random* and the second pointer will end up pointing at 1/2. So in this case, the wave-function of the composite system will approach an eigenstate of M reporting that its results were randomly distributed with the usual relative frequencies, and, with probability one, M will in fact report that its results were randomly distributed with the usual relative frequencies.

But now suppose that the distribution postulate is false of the measured particles – suppose that every particle is in the top half of its effective wave-function. In this case, the randomness pointer on the ideal measuring device will always point at *nonrandom* while the relative frequency pointer will always point at 1 regardless of how many measurements M makes, but the wave-function of the composite system will still approach an eigenstate of the first pointer pointing at *random* and the second pointer pointing at 1/2. In other words, while the composite system would be in an eigenstate of giving one report, it would in fact give a logically incompatible report – while the composite system would approach a state where each pointer is in an eigenstate of being at one position, each pointer would in fact always be somewhere very different.

On the standard interpretation of quantum mechanical states a system determinately has a property if and only if it is in a eigenstate of having the property. One direction of this biconditional is false on any hidden-variable theory that always makes a particular physical quantity determinate – on Bohm's theory, for example, a particle always has a determinate position regardless of whether the quantum-mechanical state is an eigenstate of position. One might suppose, however, that it would still be true on Bohm's theory that, if a particle is in an eigenstate of being at a particular position, then it is at that position. But here we see that even this is in part contingent on the truth of the distribution postulate. If the distribution postulate is false, then each of M's pointers *will approach an eigenstate of being at one position but in fact always be somewhere very different*. So if the distribution postulate is false, then Bohm's theory tells us that the wave-function ψ is in general entirely irrelevant to the positions of particles, and since every measurement is ultimately a measurement of position. This would mean that ψ is generally entirely irrelevant to the results of our observations.

The relationship between the wave-function and measurement results on Bohm's theory is even weaker than this last result might suggest since the theory allows for the possibility of a system approaching an eigenstate of reporting one thing but actually always reporting something else *even if the distribution postulate is true*. Suppose that the distribution postulate is true of the particles $P_1, P_2, \ldots$ but that *by chance* the observer only chooses to measure the z-spin of those particles that are in the top half of their wave-functions. As the number of measurements gets large, the composite system made up of the measuring device and the particles actually measured will approach an eigenstate of the randomness pointer pointing at *random* and the relative frequency pointer pointing at 1/2; but again, if the measuring device is reliable, then the first pointer will in fact always point at *nonrandom* and the second pointer will always point at 1. This means that regardless of the truth of the distribution postulate, there are experiments that would at least in principle distinguish between Bohm's theory and any formulation of quantum mechanics where a system is *guaranteed* to exhibit a property if the system is in an eigenstate of having the property. If the distribution postulate were false, then such an experiment would almost always distinguish between the theories; if it were true, then such an experiment would almost never distinguish between the theories but would still be possible *in principle*.

The distribution postulate then helps to make the wave-function relevant to the results of our measurements. While it does tell us that particles will in fact be distributed just as we have found them, it does ensure that the epistemic probability of finding the particle configuration in region R is equal to the integral of $|\psi|^2$ over R, and this means that one would expect to find particles randomly distributed with respect to $|\psi|^2$ with probability one. If the distribution postulate is false, then one would expect that the statistical properties of our measurement results would not be correctly predicted by the standard theory of quantum mechanics. Indeed, we have just seen that there are situations where the standard theory of quantum mechanics would predict one result with certainty and Bohm's theory would predict an incompatible result with certainty. This means that if one is committed to the dynamics of Bohm's theory and its interpretation of states and if one is committed to saving as many of the usual predictions of quantum mechanics as possible, then one is also committed to assuming that the distribution postulate is true.

4. Two theories without the distribution postulate

The distribution postulate places a very strong constraint on the relationship between the wave-function and the actual positions of particles. Further, as we have seen, the distribution postulate plays an essential role in Bohm's theory. If it is false, then one would not expect the theory to be empirically adequate; indeed, one would generally not expect any correlation whatsoever between the wave-function and the positions of particles, and, without an appropriate correlation, Bohm's theory would fail to be a serious candidate for a coherent formulation of quantum mechanics. Given all this, the distribution postulate looks like something cooked up just to provide the right relationship between the wave-function and the positions of the particles, or at least as close a relationship as the theory's dynamics and its interpretation of states will allow. Bohm himself seems to have been worried about the distribution postulate's ad hoc flavor, and he spent considerable effort arguing that his theory could be formulated without it. Bohm and Hiley argued that in typical experimental situations one might reasonably expect the distribution postulate to be at least approximately true:

> In our interpretation, the primary significance of the wave-function is that it is a quantum field . . . . However, we also say that the wave-function determines the probability density in a statistical ensemble through the relationship $P = |\psi|^2$. But . . . this is regarded as a secondary significance of the wave-function. In principle there is no reason why the probability could not be different from $|\psi|^2$, even though it is equal to $|\psi|^2$ in all cases that we have encountered so far . . . . What we have to explain then is why P should tend to approach $|\psi|^2$ in typical situations that are currently treated in physics (i.e. situations in which the quantum laws are valid) (Bohm and Hiley, 1993, p. 181).

We will look at this a little differently. Let's take the distribution postulate to be *a part* of Bohm's *original theory*. Given this, Bohm and Hiley argue that the original theory might be replaced by a new theory that makes almost the same empirical predictions but which looks less ad hoc and thus more plausible. They give two alternative formulations: let's call the first the *chaotic theory* and the second the *stochastic theory*. If one of these new theories does look less ad hoc than the original theory and if its predictions are close enough to the standard quantum-mechanical predictions that we are convinced that it is empirically adequate and will remain empirically adequate over our future experiments, then one might choose one of these over the original theory.

Let $P(t)$ be the epistemic probability distribution of the particle configuration at time t – that is, the integral of $P(t)$ over a region R represents the probability of the actual particle configuration being in R at t. Again, one might picture the actual positions of the particles being randomly determined by $P(t_0)$ at some initial time t_0, then think of the particles moving along continuous trajectories according to the dynamical laws from these initial positions. As the particles move, the epistemic probability distribution $P(t)$ changes. Bohm and Hiley want to describe conditions where $P(t)$ would approach $|\psi(t)|^2$, then claim that these conditions actually obtain.

The chaotic theory is based on the idea that even Bohm's deterministic dynamics allows for the possibility of a probability distribution $P(t)$ that is initially different from $|\psi(t)|^2$ approaching $|\psi(t)|^2$ over time. Some Hamiltonians, initial probability distributions, and wave-functions would have this property and some would not. The chaotic theory is simply Bohm's original theory where the distribution postulate is replaced with a postulate that says that the actual Hamiltonian, probability distribution $P(t_0)$, and wave-function $\psi(t_0)$ were such that $P(t)$ and $|\psi(t)|^2$ were indistinguishable by the time we started making quantum-mechanical measurements given the measurements we have in fact made. If this postulate is plausible enough, then the chaotic theory will look less ad hoc than the original theory. But is this postulate significantly more plausible than assuming that the probability distribution has always been given by $|\psi|^2$?

The argument that Bohm and Hiley give for the plausibility of this new assumption goes something like this. Consider a particle in a box. It is easy to imagine the particle moving in a very simple way. If, for example, the initial wave-function was an eigenstate of energy or if the particle was initially placed somewhere in the box where the wave-function was always zero, then the particle would not move at all.[8] In general, however, one would expect the wave-function to bounce around in the box and look increasingly complicated as the various reflected components interfere with each other; and, since the motion of the particle is directly determined by the wave-function, one would consequently expect the particle's motion to be increasingly chaotic. Since the evolution of the wave-function is determined by the Hamiltonian, the particle is more likely to exhibit chaotic behavior the more complicated the initial wave-function and the Hamiltonian are (a complicated Hamiltonian here might correspond to the box having very irregular walls); further, if the particle is ever going to exhibit chaotic behavior, then the more complicated the wave-function and Hamiltonian are, the sooner this will happen. So whether the particle eventually ends up behaving chaotically and how long it takes for this to happen, if it is ever going to happen, generally depend on the particle's initial position, the initial wave-function, and the Hamiltonian, and at least intuitively one would eventually expect to see chaotic behavior for most combinations of these three parameters.

Similarly, in a many-particle system one would eventually expect to see the point representing the positions of the particles move in a chaotic way in configuration space for most combinations of initial positions, wave-functions, and Hamiltonians – that is, one would eventually expect the point representing the positions of the particles to get arbitrarily close to every accessible point in configuration space (where a point is accessible if there is a path from the point representing the current positions of the particles to the point that does not pass through a region where the wave-function is always zero) and one would expect the distance between the trajectories corresponding to slightly different initial configurations to diverge relatively quickly. Consequently, given an arbitrary initial wave-function and Hamiltonian, one would expect even a very localized epistemic probability distribution to eventually spread throughout all of the accessible configuration space. Here $|\psi|^2$ and P might roughly be thought of as two fluids in the same container, something like a cup of cold coffee with a blob of cream in it, where the cream P is being mixed in by the sloshing about of the cold coffee $|\psi|^2$. One would expect P to form a complicated thread-like structure, and one would expect that more of P would be found in regions where more of $|\psi|^2$ has flowed. So what is $P(t)$ in the limit as t gets large? The dynamics implies the conservation equation

$$(8) \quad \frac{\partial \rho}{\partial t} + \sum_n \nabla_n \cdot \rho v_n = 0$$

where $\rho = |\psi|^2$ and v_n is the velocity of particle n. Given this conservation equation and the assumption that the conditions are such that the motions of particles are chaotic, Bohm and Hiley conclude that $P(t)$ will approach $|\psi|^2$ (1993, pp. 182–184).

While their argument is more intuitive than rigorous, let's suppose that the loose ends can be tied up. Where does this leave us? If chaotic motions are enough for one to expect that $P(t)$ would eventually approach $|\psi(t)|^2$, then if one had a good argument that the actual Hamiltonian, probability density $P(t_0)$, and wave-function $\psi(t_0)$ were in fact such that the particle motions of systems that we have measured were chaotic for long enough to get a sufficiently good correlation between $P(t)$ and $|\psi(t)|^2$ to account for our measurement results, then the chaotic theory would presumably look more plausible than the original theory. Bohm and Hiley argue

> There are a wide range of conditions in which the sort of process described above may take place. These include metals, gases, plasmas etc. so that not only in the laboratory, but also in the stars and interstellar space, we might reasonably expect that quantum theory itself contains processes that tend to produce distributions near $P = |\psi|^2$ (1993, p. 184).

It might be added that our best theories of cosmology suggest that the early motions of most particles were chaotic. There are, however, reasons for not liking the chaotic theory. First, the chaotic theory still requires a fairly strong boundary condition – since it is easy to specify conditions where $P(t)$ would not approach $|\psi(t)|^2$, the chaotic theory needs to make some assumption concerning former physical states in order for ψ to be at all relevant to measurement results. This new assumption may look more plausible than the distribution postulate, but how plausible it looks depends on the details of our best cosmology, on what we believe concerning the histories of the systems we have measured. If a system has had a relatively peaceful past, then one would expect P to be relatively independent of $|\psi|^2$. Further, no matter how long one waits, the actual probability density P will *never* be identical to $|\psi|^2$ if it did not start out identical to $|\psi|^2$. Since equation (8) is time-reversal symmetric, if it is ever the case that $P = |\psi|^2$, then not only will this always be true in the future but it always *was* true in the past, which means that the chaotic theory will always make statistical predictions that differ from the probabilities given by $|\psi|^2$. Finally, the closer one wants the predictions of the chaotic theory to approximate the standard quantum-mechanical probabilities, the more restrictive the conditions on the past physical state will have to be. Consequently, the strongest conclusion that Bohm and Hiley can make here is that if sufficiently chaotic initial conditions obtain for a system, then "in the sense of a coarse grained average" and "in any process that does not depend on the complex and chaotic fine details of the motion . . . we may say that the probability is effectively $\rho = |\psi|^2$" (1993, p. 184). To the extent that one is committed to the standard quantum-mechanical probabilities being *exact*, one would presumably not like the chaotic theory.

The stochastic theory differs from Bohm's original theory in a more striking way than the chaotic theory does since the stochastic theory has different dynamical laws. According to the stochastic theory, the velocity of particle *n* is given by

$$(9)\quad v_n = \frac{\mathrm{Im}\, \psi^*(x, t)\, (\partial/\partial x_n)\psi(x, t)}{m_n|\psi(x, t)|^2} + \xi_n(x, t)$$

where $\xi_n(\mathbf{x}, t)$ represents a random contribution to the velocity of particle *n*. But not just any random contribution will do what Bohm and Hiley have in mind:

> . . . our aim . . . is to determine the properties of the random motions described by [$\xi_n(\mathbf{x}, t)$] so that an arbitrary initial ensemble with a probability distribution $P(\mathbf{x}, t)$ will approach an ensemble with a probability distribution $P(\mathbf{x}, t) = \rho(\mathbf{x}, t) = |\psi(\mathbf{x}, t)|^2$ after a suitable interval of time. We then suppose that under typical conditions obtaining in quantum mechanical measurements, the interval of time will have been long enough so that the probability distribution will be given by $|\psi(\mathbf{x}, t)|^2$. In this way we understand the fact that while $P(\mathbf{x}, t)$ and $|\psi(\mathbf{x}, t)|^2$ are basically independent concepts their numerical values will experimentally turn out to be equal in quantum mechanical experiments (1993, p. 195).

Bohm and Hiley suggest that $\xi_n(\mathbf{x}, t)$ might arise from some "deeper level" of the physical world just as small objects exhibit Brownian motion as the result of "impacts originating at a finer molecular level" (1993, p. 194). They end up characterizing this random component of motion by an osmotic velocity

$$(10)\quad \mathbf{u} = \frac{D\nabla\rho}{\rho}$$

and a diffusion current

$$(11)\quad \mathbf{j}^{(d)} = -D\nabla P$$

so that the total current is

$$(12)\quad \mathbf{j} = \frac{P\nabla S}{m} + \frac{D\nabla\rho}{\rho} - D\nabla P$$

where D is the diffusion coefficient, $\rho = |\psi|^2$, P is the epistemic probability, and $\nabla S/m$ represents the velocities of the particles on the original deterministic dynamics. The osmotic velocity and diffusion current are chosen so that, if $P = \rho$, then the osmotic velocity would be balanced by the diffusion current and the mean velocities of the particles would consequently be just what the original theory would predict. Further, near a zero of $|\psi|^2$ the osmotic velocity approaches infinity in the opposite direction, which keeps particles out of regions where $|\psi|^2$ is zero. More generally, the osmotic velocity pushes particles away from regions where $|\psi|^2$ is small and toward regions where $|\psi|^2$ is large. The upshot of all this is that an arbitrary initial distribution $P(t)$ will always approach $|\psi(t)|^2$ as long as $P(t_0)$ is zero wherever $|\psi(t_0)|^2$ is zero and points where $P(t) \neq |\psi(t)|^2$ are mutually accessible. Near this equilibrium the *mean motions* of the particles will be very close to what the original theory would predict, which would presumably make the stochastic theory empirically adequate over the experiments we have performed so far.

The most serious problem with the stochastic theory is that, without an independent argument for why the motions of particles ought to contain a random component, its dynamical laws look at least as ad hoc as the original theory's distribution postulate. In order to guarantee that an arbitrary distribution P always approaches $|\psi|^2$, it is necessary to suppose that $\xi_n(\mathbf{x}, t)$ is contingent on both $|\psi|^2$ and P in a very special way. Bohm and Hiley justify their choice of $\xi_n(\mathbf{x}, t)$ by noting that

> without assuming an osmotic velocity field of this kind, there would be no way of explaining [why no particle ever reaches a point where the wave-function is zero and why most particles are found near the maxima of the wave-function]. As a result of random motions, for example, a particle just undergoing a random process on its own would have no way of "knowing" that it should avoid the zeros of the wave-function (1993, p. 200).

Of course, finding a $\xi_n(\mathbf{x}, t)$ that would guarantee that "an arbitrary initial ensemble with a probability distribution $P(\mathbf{x}, t)$ will approach an ensemble with a probability distribution $P(\mathbf{x}, t) = \rho(\mathbf{x}, t) = |\psi(\mathbf{x}, t)|^2$ after a suitable interval of time" was the strategy from the start. But if the particle dynamics of the stochastic theory is explicitly constructed just to get P to approach $|\psi|^2$, then we have gained nothing by changing theories – if one is worried about Bohm's original theory being ad hoc, then the stochastic theory with the justification given here for its dynamical laws can be no better off. Moreover, it seems that the original theory comes out looking *better* than the stochastic theory since introducing a special boundary condition is presumably less objectionable than cooking up a special dynamics with the sole aim of being able to handle a broader range of boundary conditions. After all, we have a precedent for the former strategy in classical statistical mechanics when we assume that the initial state of a system was a low entropy state with random particle motions in order to explain why entropy is low now and tends to increase over time. If we had good independent reasons for supposing that the stochastic dynamics is true, if we ended up with good reason to believe, for example, that there really is a "deeper level" of the physical world that generates random motions, then the stochastic theory might look better than it does. But again, in order for this particular stochastic formulation to be true, these random contributions to the motion of particles would have to be correlated to the wave-function and the current probability distribution in just the way that Bohm and Hiley stipulate, and at least to me this sounds implausible.

5. Conclusion

We have seen that the distribution postulate plays an essential role in Bohm's original theory. If it were false, then the theory would generally fail to make the right statistical predictions. And if making the wrong predictions were not bad enough, if the distribution postulate were false, then there would be situations where a particle would approach an eigenstate of having one position but in fact always be somewhere very different, so there would generally be no relationship whatsoever between the actual quantum-mechanical wave-function and the results of an observer's measurements. But we also saw that on Bohm's theory whether a particle will be at a position whenever it is in an eigenstate of being at the position is contingent on what systems an observer chooses to measure *even if the distribution postulate is true*, so even the truth of the distribution postulate is not a sufficient condition for the wave-function being relevant to the positions of particles and the results of measurements. This helps to show how loose the connection is between the wave-function and the positions of particles on Bohm's theory.

One might argue that the distribution postulate makes Bohm's original theory look at least somewhat ad hoc since it looks like an assumption cooked up just to ensure that with probability one the theory makes the right statistical predictions. Consequently, one might want to look for a new formulation of Bohm's theory that makes the same empirical predictions (or very nearly the same empirical predictions) but looks more plausible. But each of the two proposals we have considered here have unattractive features.

Rather than assuming the distribution postulate, the chaotic theory is based on the assumption that the Hamiltonian, initial particle distribution $P(t_0)$, and initial wave-function $|\psi(t_0)|^2$ were such that $P(t)$ would approach $|\psi(t)|^2$ and be very close, closer than our actual experiments could distinguish, by now. Since this would not always be the case, the chaotic theory must still make *some* assumption concerning the past. While this assumption is arguably more plausible than the distribution postulate, it is difficult to say just how much more plausible it is because among other things its plausibility is contingent on what we believe about the histories of the systems that we measure. Further, the plausibility of the chaotic theory as a whole depends on how much one worries about the fact that it makes different statistical predictions than the standard formulation of quantum mechanics.

The stochastic theory has dynamical laws that are explicitly formulated so that virtually any initial distribution would eventually approach $|\psi|^2$. All that is required is that $P(t_0)$ is zero wherever $|\psi(t_0)|^2$ is zero, that points where $P(t) \neq |\psi(t)|^2$ are mutually accessible, and that the initial particle configuration is always somewhere where ψ is nonzero. In addition to sharing with the chaotic theory the feature that $P(t)$ is almost never exactly equal to $|\psi(t)|^2$, if one formulates the dynamical laws with the sole aim of making the distribution postulate approximately true over time, then the resultant theory cannot be more plausible than Bohm's original theory.[9]

All things considered, I like Bohm's original theory better than the chaotic theory, which I like better than the stochastic theory. Indeed, a hidden-variable theory constructed along similar lines may well end up as our best formulation of quantum mechanics. If something like this happens, then presumably either we will have figured out some way to do without the distribution postulate or we will have convinced ourselves that it is not such a bad thing after all.

Notes

[1] A very short history of Bohm's theory: de Broglie had the idea of interpreting the wave-function as something that physically guides particles, which always have determinate positions (1930). Bohm wrote down a theory based on this idea (1952), and Bell clarified and defended the theory (1987).

[2] The standard theory is the nonrelativistic version of quantum mechanics presented, for example, by von Neumann (1932). On the standard theory the quantum-mechanical state evolves linearly and deterministically until a measurement is made. When a measurement is made, the state randomly collapses to an eigenstate of the observable being measured with the probabilities determined by Born's rule: the probability of the post-measurement state of a system initially in $|\phi\rangle$ ending up in $|\psi\rangle$ is $|\langle\phi|\psi\rangle|^2$.

[3] See Albert (1992, pp. 134–79) for a very clear discussion of how this works.

[4] This was probably what Bell had in mind when he said that "*no one can understand [Bohm's theory] until he is willing to think of* ψ *as a real objective field rather than just a 'probability amplitude' . . . even though it propagates not in* 3*-space but in* 3*N-space*" (Bell, 1987, 128; his italics). This is a very useful picture, but when it comes to telling a causal story concerning the motions of the particles, one might reasonably wonder how a field that propagates in one space would influence the motions of particles in another space. One might, of course, picture N interdependent fields – one field for each projection of ψ onto a subspace representing the position of a particle – each evolving in 3-space and only affecting the motion of one particle.

[5] It is not necessarily true, however, that every actual measurement is ultimately a measurement of position; that is, it may be that making only positions determinate fails to make the results of every actual measurement determinate. One might easily imagine an observer who records his measurement results in terms of some physical parameter other than position – energy, for example. Making position determinate here would not make the observer's measurement results determinate. So whether every measurement is ultimately a measurement of position is contingent on the physiology of sentient beings. If a theory fails to make an observer's measurement results determinate, then it is nonsense to claim that it is empirically adequate.

[6] This is the fifth principle in Bohm and Hiley's (1993, pp. 29–30) description of the single-particle theory, for example.

[7] There is a sense, however, in which P's only real property is its position. z-spin, for example, is not a property of P alone; it is rather the orientation of the magnetic field, the position of P, and P's wave-function that determine the result. Consequently, one might say that z-spin is only a *contextual property* on Bohm's theory.

[8] See Bohm and Hiley (1993, pp. 42–9) for an explanation of why a particle does not move in a stationary state but does move in a linear combination of stationary states.

[9] I would like to thank David Albert, Frank Arntzenius, and Wayne Aitken for comments on an earlier version of this paper.

References

Albert, D. Z.: 1992, *Quantum Mechanics and Experience*, Harvard University Press, Cambridge, MA.

Bell, J. S.: 1981, 'Quantum mechanics for cosmologists', reprinted in Bell, J. S. (1987, pp. 117–38).

Bell, J. S.: 1980, 'De Broglie-Bohm, delayed-choice double-slit experiment, and density matrix', reprinted in Bell, J. S. (1987, pp. 111–6).

Bell, J. S.: 1987, *Speakable and Unspeakable in Quantum Theory*, Cambridge University Press, Cambridge.

Bohm, D.: 1952, 'A suggested interpretation of the quantum theory in terms of "hidden" variables, I and II', *Physical Review* **85**, 166–79 and 180–93.

Bohm, D. and Hiley, B. J.: 1993, *The Divided Universe: An Ontological Interpretation of Quantum Theory*, Routledge, London.

de Broglie, L.: 1930, *An Introduction to the Study of Wave Mechanics*, E. P. Dutton and Co., New York.

DeWitt, B. S. and Graham, Neill (eds.): 1973, *The Many-Worlds Interpretation of Quantum Mechanics*, Princeton University Press, Princeton.

Everett, H. III.: 1957, '"Relative state" formulation of quantum mechanics', *Reviews of Modern Physics* **29**, 454–62, reprinted in DeWitt and Graham (eds.) (1973).

Everett, H. III.: 1973, 'Theory of the universal wave-function', in DeWitt and Graham (eds.) (1973).

Hartle, J.: 1968, 'Quantum mechanics of individual systems', *American Journal of Physics* **36**, 704–12.

von Neumann, J.: 1932, *Mathematische Grundlagen der Quantenmechanik*, Springer, Berlin; reprinted and translated as *Mathematical Foundations of Quantum Mechanics*, in R. Beyer (trans.) Princeton University Press, Princeton, 1955.

University of California, Irvine, USA