A Dialogue on the Meaning of Quantum Mechanics

May 8, 2006

Recorded in 2002 by Florian Marquardt [incomplete draft]. Note: I wrote this while still a graduate student at Basel University, and never got around to finish it completely. Since, however, it is 80% complete, I thought it might be useful to make it available nonetheless. - FM, 2006
(Florian.Marquardt@physik.lmu.de)

Contents

1 Trajectories

QM: So you want to *understand* quantum mechanics, right?

LR: Yes. They have taught me everything about the wave function of the electron and how to use the Schrödinger equation to calculate its time evolution. I have learned how to find the energies of the hydrogen atom and how its orbital wave functions look like. I know how to solve scattering problems, when the wave is partially reflected and partially goes through a barrier. And I was told

that experimental predictions can be obtained by squaring the wave function, because this gives you the probability of finding the electron here or there. But I want to understand what *really* goes on. I want to grasp the physics behind the formalism, so to speak!

QM: Well, what exactly do you hope for when you say you want to "understand" it? Is it not enough to be able to make precise predictions for every conceivable experiment? Do you expect to find that, after all, the electron's motion can be described in the same way as that of macroscopic bodies - say, a stone being thrown or a satellite moving in orbit?

LR: No, of course not. It must be something very tricky, indeed.

QM: What about Maxwell's equations of electrodynamics? Do you "understand" them? Is it not also quite mysterious to have electric and magnetic fields propagating through vacuum, without being attached to anything material? But I guess you would *not* try to uncover some deeper reason behind those equations, would you?

LR: No, I am satisfied with Maxwell's equations as they stand now. But I claim that they are different from quantum mechanics: The electric and magnetic fields can always be *measured*, at every point and at every time, by looking at the deflection of little test charges. So they are very real, in contrast to the mysterious wave function! And I can *imagine* how these fields propagate as time runs on. There is a definite physical picture behind it.

QM: Still, when these equations had been only a few decades old, people were not at all content with that picture. They tried to "understand" them - I guess much in the same way you are trying to understand the physics behind quantum mechanics. They knew Maxwell's equations are similar to other wave equations, for example those describing the propagation of sound waves. So it was entirely natural to try to reason by analogy: "air is to sound waves as ... is to light waves". They even introduced a name for "...", it was the "ether" in which light was believed to propagate. But it could never be observed directly. With time, the properties one had to attribute to the "ether" in order not to contradict the known facts became more and more mysterious. Nowadays, we can do without the "ether" and any deeper physical explanation of electrodynamics. It is just a matter of getting used to it.

LR: Hold on. You have left out something important. I agree that we can do very well without the ether. But consider this: If they had not looked for the ether, they would not have found special relativity! So it paid off not to ignore these questions, although the final answer turned out to be quite different from that expected in the beginning.

QM: Well, I see you do not give up easily. So perhaps you should tell me in more detail what exactly are the points you find so difficult to grasp in quantum mechanics. Let us start with something basic: You mentioned the "probability of a particle being here or there" - so you agree that the wave function does not describe a smeared-out particle?

LR: Yes. Of course, that would have been the simplest possibility: to picture an electron as a small wave packet moving in space or as a standing wave around the nucleus of the atom. Waves would have been truly fundamental while "particles" would have been the word for localized disturbances in the wave field. In that interpretation, the square of the wave function would have given directly the charge density of the extended electron...

QM: Which is the initial interpretation proposed by Schrödinger himself.

LR: ...In fact, it works nicely for a number of purposes: For example, the charge density belonging to the ground state wave function of the atom is static - a fuzzy, spherically symmetric cloud around the nucleus. So it is natural that no electromagnetic field is radiated. But if you have excited the atom a little bit, then the wave function consists of an admixture of the excited states and the ground state, and the charge density starts to vibrate. The center of the negative electronic charge density does not coincide any longer with the positive nucleus during the vibration. So the atom acquires an oscillating electric dipole moment, just like a small antenna: Therefore, it radiates an electromagnetic wave, and both the frequency and also the intensity of that wave can be understood in this picture. The picture of a smeared-out charge density is also nice and useful for visualizing molecules, so it has always been applied in chemistry.

QM: But you know that it fails badly in other circumstances.

LR: Yes. For example, if we have an electron moving in free space, it may be described by a wave packet. Still, even if the extent of that packet is very small at first (say, only a few nanometers), it will necessarily spread further. After a while, it will cover macroscopic distances (say, millimeters). But in every experiment only localized, point-like electrons are detected. So the wave function cannot really describe an extended electron.

QM: Right. In contrast, the probability interpretation has worked perfectly for every experiment done so far. The square of the wave function at a particular point in space is simply the probability of finding the electron near that point in a measurement. But besides those problems with spreading wave packets, the picture of the charge density cannot work for another reason: If you describe a system of more than one electron, you have to use a wave which is a function of the coordinates of all the electrons. So this wave function does no longer live in threedimensional space but, instead, in a $3 \times N$ dimensional space, if you have N electrons. Obviously, it makes no sense to talk of a charge density in such a high-dimensional abstract space. On the other hand, it is perfectly reasonable to ask for the probability of finding the electrons in some particular configuration at a certain time: "What is the likelihood to find particle one at this point, and, at the same time, particle two at the other point over there?" This is done by inserting the prescribed coordinates of all the electrons into the many-particle wave function and squaring it. Anyway, this has not been a new concept at all: The same is done in classical statistical mechanics,

where one also deals with a probability distribution that depends on the coordinates (and, maybe, momenta) of all the particles simultaneously. It gives the probability of finding the set of particles in this or that configuration, when the total system is in thermal equilibrium. The major difference is, of course, that the quantum mechanical probability distribution will often show interference patterns, because it is derived from a wave function.

LR: Agreed. We cannot easily interpret everything in terms of waves. And the probabilities calculated using the wave function offer precisely the right description of experiments. But what I want to know is: How do those probabilities come about? You have mentioned statistical mechanics: We know that it is just an efficient way of describing the complicated behaviour of a large number of particles that excert forces on each other and move according to Newton's laws of motion. Since we do not know the precise initial positions and velocities of all the particles, we have to resort to a probabilistic description. It is also practically infeasible to calculate in detail all the trajectories, since they depend very sensitively on the initial conditions. Besides, we would not want to know all those details anyway. We are particularly interested in the average behaviour of the particles. And it is this average behaviour in the thermal equilibrium state which is very well described using the comparatively simple "Boltzmann distribution". Now, given the probability interpretation of quantum mechanics, is it not inevitable to wonder what is the microscopic physics behind the efficient description provided by the wave function and the Schrödinger equation? I could put it like this: "Newton's laws of motion are to classical statistical mechanics as ... is to quantum mechanics - fill in the missing word!"

QM: Well, of course that is what people tried at first, even before they had the final formalism of quantum mechanics in their hands. Perhaps it is best to discuss your ideas on the "physics behind quantum mechanics" for some concrete example. Let us take the double slit experiment, that should be simple enough.

LR: Fine. In fact, there are even two different points that I find disturbing about the usual description of this experiment. The first one is this: the wave spreads continuously over all of the screen and develops the pattern that stems from the interference of the waves coming from both slits. But, of course, we know that there is no such thing as a "smeared-out" electron. Indeed, in the end the electron hits the screen *either* here *or* there, at a well-localized spot. But what to do with the extended wave function in that event? Well, here is what people say and what sounds to my mind like a cheap excuse: They say the wave function "collapses instantaneously to the point where the electron has been detected". But, tell me, how does the wave function that is spread out over a larger portion of space "know" that the electron has been detected at a particular point, so it should immediately vanish everywhere except at that point? Don't we have the theory of special relativity which forbids influences to travel faster than light? Is this not a mysterious "action at a distance"? Imagine something similar for the classical electromagnetic field: If radio waves encounter

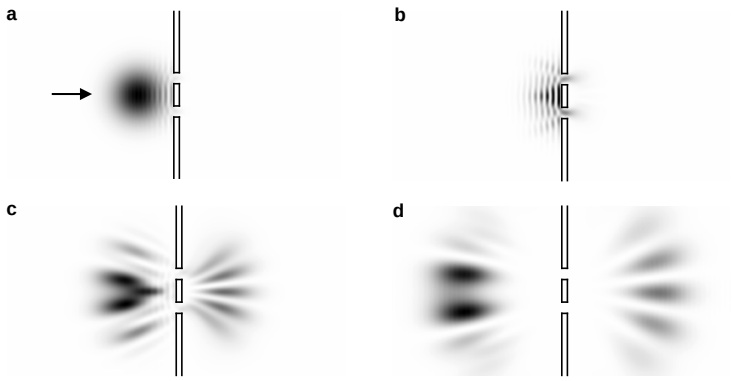

Figure 1: The double slit experiment: In the example shown here, a spherical wave packet moves towards the two slits from the left. The probability density is plotted in a grayscale plot, which means the electron is most likely to be found in the darkest regions. As the wave approaches the slits, a large part of it is reflected, such that a standing wave pattern develops to the left of the slits (a,b). Only a smaller fraction passes through the slits. The waves emanating from the two slits overlap and form an interference pattern (c,d), which might be observed on a screen to the right. Note that the reflected wave to the left also shows interference.

an antenna somewhere, then they may excite a current in it. This draws energy out of the field and the excited antenna emits radiation, which is superimposed onto the incoming field, canceling it partly. So the field is changed, but not instantaneously over arbitrarily large distances. In contrast, the "collapse of the wave function" is effective regardless of the distances between the different parts of the wave. For example, the wave could be split into two wave packets at a partially transparent barrier. Then one of the wave packets could travel to Jupiter, the other one to Mercury and in the case of detection on either planet, the other part of the wave would vanish at the same instant!

QM: True. But isn't this quite natural in a probabilistic description? At first you do not know whether the electron is here or there, so you ascribe a finite probability to both possibilities. As soon as you have detected it at a certain place, you can update your probabilities, because you have gained some knowledge. You can do further predictions on the basis of this refined information. This is true in the same way for classical statistical mechanics or any classical reasoning based on probabilities. No mysterious "action at a distance" is involved. What changes suddenly is just your knowledge about the actual situation.

LR: That sounds plausible. In that case the "collapse of the wave function" isn't really something to worry about. And it makes me confident that we can find a more detailed "common-sense" theory underlying quantum mechanics just as we know Newton's law of motion is the microscopic theory underlying

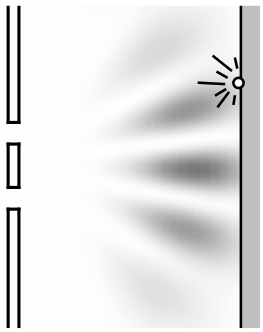

Figure 2: "Collapse of the wave function": The wave function is extended before the electron is detected at some random point on the screen. After the event, the wave function is to be replaced by a packet localized around that point.

classical statistical mechanics. After all, what you say amounts to the following: There always *is* some definite place for the electron, only we do not know it!

QM: Uhm, well, you see... Actually that is not *quite* the point I wanted to make. But I agree that what I have said sounds like that. In fact, things are more complicated: In order to really understand the "measurement problem", one should deal not only with the single electron but also with all the particles in the measuring apparatus - for example, those making up the screen in our example. So it becomes a many-particle system and the wave function lives in a higher dimensional space. Maybe that gets a bit too complicated at this point. We should discuss these things later on. You mentioned you had a second problem with the double slit experiment?

LR: Yes, it is this: Often, people say "the electron goes through both slits at once and interferes with itself" - or something like that. That sounds very vague and imprecise. Is it not entirely reasonable to assume that it goes *either* through this slit *or* through the other one?

QM: But certainly you should have been taught in your quantum mechanics course that this can never give you the correct interference pattern! Because if the electron only goes through a single slit, how should it know about the presence of the other one? And if it does not know about the presence of the other one, you will just end up with a distribution of electrons on the screen which is the sum of the distributions you obtain for each slit alone - no interference fringes!

LR: Of course I know this story. But there is a very naive classical reasoning behind it - maybe the electron *does* know about the presence of the other slit,

even though it only goes through a single one. There may be some long-range forces involved, or perhaps it emits a photon which goes through the other slit and is reabsorbed later, or maybe there is a mysterious wave field which pervades everything and influences the motion of the particle and knows about the presence of both of the slits. Actually, I like that possibility: Imagine water waves passing between some rocks. The ripples will form an interference pattern, similar to that in the double-slit experiment. Now, if a small piece of wood rides on the waves, it will be influenced by that pattern, although it has not mysteriously traveled "all paths at once". In any case, I am sure there *is* a definite position of the electron at each and every moment, even when we do not look at it. And therefore, the electron will follow some definite trajectory, even though we do not know which one.

QM: Come on, you know there are no particle trajectories in quantum mechanics!

LR: That's what people say. But I have some good reasons for believing in trajectories: First of all, you can detect the position of an electron instantaneously - or nearly so - by scattering light off it. And when you repeat that observation a short time later, you will find the electron at a position which is not too far from the first one. So, not only does it have some precisely defined position whenever we choose to observe it, but there are also no sudden jumps in its position. And, of course, there should be no jumps, because that would mean the particle could move faster than light. As far as I know, there is not a single experimental result that would force us to conclude that such faster-than-light motion can happen in quantum mechanics. So it is natural to assume the particle follows a continuous trajectory.

QM: Maybe...

LR: And secondly, it would be hard to imagine the particle being "everywhere at once" in a kind of mysterious superposition until we measure it. After all, it is not just a matter of a few nanometers: the different parts of the wave may be separated by macroscopic distances; remember my example of a wave packet going to Jupiter and the other one to Mercury. And it is not just about electrons; interference experiments are done with all kinds of other particles: That includes whole atoms or ions or even molecules composed of many individual atoms. You can trap those particles and manipulate them individually. You can even make pictures of them, for example with a scanning tunneling microscope and they appear to me very "real". So it is no longer like in the early days of quantum mechanics, when you could only do experiments on large ensembles of particles which you could not resolve individually. There, a merely statistical interpretation may have been enough, but now it is not.

QM: I agree that it is hard to imagine how a whole "buckeyball" molecule can be in a superposition of different places. But you must acknowledge that there have been equally good reasons for ruling out the existence of trajectories: In quantum mechanics, there is simply no way to have, at the same time, a precisely defined momentum and a precisely defined position of a particle. This is

Figure 3: Trajectories are changed when one tries to measure them too accurately: A series of position measurements, with a precision given by the gray circles, is carried out in constant time intervals. When the spatial resolution is increased (from left to right), the random deflections in momentum at each step are also enlarged, which finally leads to something like a random walk.

forbidden by Heisenberg's uncertainty relation, which is a direct consequence of quantum mechanics itself.

LR: I know: $\Delta p \Delta x \geq h/4\pi$. If the spread in position Δx is given, then the minimum spread in momentum Δp can be calculated using this inequality. And it means that the velocity of the electron will be defined with a precision no better than $100\,000$ meters per second, for a position uncertainty of a nanometer. If the spread in position is reduced by a factor of ten, then the minimum spread in velocity must grow by the same factor. And, of course, a tenth of a nanometer is just the diameter of a hydrogen atom, so the single electron inside that atom must have typical velocities on the order of at least a million meters per second, according to Heisenberg. Actually, this is indeed the correct magnitude. For a heavier particle, of course, the spread in momentum would be the same, only the spread in velocity would be reduced. That is all very fine - but why should it rule out the existence of trajectories?

QM: Well, if you try to follow the trajectory of the particle, you have to carry out successive position measurements. But, depending on how precisely these measurements pinpoint the electron's location, each of them will induce a smaller or larger spread in its momentum. Thus, the trajectory is changed. If you want to follow it very closely, measuring extremely often with high precision, you will observe just a random walk. And the position that the electron reaches at the end (on a screen, for example) will have nothing to do any more with the position it *would* have reached if you had let it alone. So, maybe, you may even choose to talk of a trajectory, but once you try to measure it, it is heavily perturbed and becomes a different one.

In fact, that is not only demanded formally by Heisenberg's uncertainty relation. It is an inevitable consequence of the physical measurement process. You may know the example given by Heisenberg himself: if you want to observe the position of the electron under a microscope, you have to scatter light off it. If the position is to be determined accurately, you will have to choose light of a short wave length. This means the photon which is scattered off the electron carries a high momentum, because wavelength and momentum are inversely

proportional, according to de Broglie: $\lambda = h/p$. A part of the momentum is passed to the electron in the "collision" between photon and electron. And since the lens of your microscope will collect photons flying in various directions, you can never determine precisely the momentum change of the photon. Thus, the momentum transferred to the electron is not only random but unknown, to an amount which turns out to be precisely that demanded by the uncertainty relation! Everything fits together - the experiment cannot determine anything more precisely than the theory of quantum mechanics is able to describe it and vice versa. Otherwise, there could be contradictions or at least something would be missing in quantum mechanics.

LR: Hm, what happens if I do *not* use a lens?

QM: Then you are able to determine exactly which direction the photon went after the scattering event, because you do not collect light going into all the different angles. So, if you know which way it came, you can calculate the momentum that must have been transferred to the electron, from momentum conservation. Therefore, although the momentum of the electron still has been changed by a random amount, you know its new value - provided you knew the initial momentum. Of course, you have to pay a price: Without a lens, you cannot any longer find out at which point exactly the photon was scattered, so you don't know where the electron is, you only know its momentum. You simply cannot determine both position and momentum exactly at the same time.

LR: I admit, it is impressive how quantum mechanics avoids all contradictions. But still: If we know that the physical reason behind the uncertainty relation is the inevitable perturbation during the measurement, it is all the more natural to assume that, in reality, the electron follows some trajectory. Of course, we can never find out which one without perturbing it, as you have said. But it would be helpful nevertheless to be able to *imagine* a definite trajectory, at least.

QM: Now assume you measure the position of the electron following such a trajectory. Why do you get a statistical distribution of positions? After all, quantum mechanics tells us that you do get such a distribution if you repeat the experiment several times, with a series of electrons being injected into your apparatus, one after the other, each prepared in the same way. And you cannot claim that it is the measurement which is responsible for the random component: That would be true only for subsequent position measurements of the same electron, because these are affected by the perturbation introduced by the first measurement.

LR: Oh, the trajectories are random, of course. You take care to prepare the wave function of each electron to be the same, with suitable velocity filters and slits and so on. But this is not sufficient: quantum mechanics forbids you to determine precisely both position and momentum. Both of them will have some spread, depending on the wave function. So at least the initial conditions of the trajectories are different and therefore you get a distribution of trajectories and a distribution of arrival points on the screen. It is as simple as that.

QM: I fear it is too simple! Remember, you could also choose to measure the momentum of the particle, instead of the position.

LR: Then I find a distribution of momenta. What is the problem?

QM: So you claim you can make the distributions of momenta *and* positions obtained from your trajectories coincide simultaneously with the quantum mechanical results? This is not possible, again due to Heisenberg's uncertainty relation. Because, what you would like to have is a kind of classical probability density that depends *both* on position and momentum simultaneously. But remember, the square of the wave function only gives you a probability distribution of the position alone. And if you want to get the distribution of momenta, you have to Fourier transform the wave function and afterwards take its square. That is, you have to change to the momentum basis. You can never get both at the same time. In fact, the nearest thing to what you would like to have that exists in quantum mechanics is the so-called Wigner density, which is a function of both position and momentum. It fulfills everything you would expect from a classical probability density, except that, unfortunately, it turns out to be partly *negative*, for most wave functions. So there you have "negative probabilities" - I guess you could not realize that with your random trajectories!

LR: Well, I don't know enough about the Wigner density you are talking about. But I can easily demonstrate to you that you must be wrong: If quantum mechanics tells me that the probability distribution of position is $|\Psi(x)|^2$ and that of momentum is $\left|\hat{\Psi}(p)\right|^2$, then I can set up a probability distribution that depends both on x and p, gives back these probabilities when you only ask for the position or for the momentum alone and is, on top of that, never negative. Just take $\rho(x,p) = |\Psi(x)|^2 \cdot \left|\hat{\Psi}(p)\right|^2$. Of course, I still do not know how to find a law of motion for my trajectories that gives this distribution at every instant of time. But at least it is not impossible from the outset.

QM: Oops, you are right. I forgot to mention one requirement: You should not only be able to obtain the correct probability distributions for position or momentum but also for every combination of position and momentum, in the form $Ax + Bp$. This is because you could choose to measure not in the position basis, not in the momentum basis, but in a basis with respect to such a mixture of both x and p.

LR: It seems you invent new rules whenever I have a good idea... Seriously, why should anyone measure in such a basis? And how could that be done physically, anyway? After all, I would be content to reproduce the real *physical* observations, not every formal quirk of quantum mechanics.

QM: It is not something merely formal. Assume that the particle starts out at position x and has momentum p. Then, if you measure its position after some time t, this will be $x + vt = x + pt/m$ in terms of the original position and mo-

mentum. So it just corresponds to a measurement of such a linear combination of x and p.

LR: Well, maybe you are right. But if I cannot reproduce all these probability distributions at once, I decide in favor of the position distribution: $|\Psi(x)|^2$. This is what should be observed if one measures the positions of such an ensemble of electrons following the random trajectories.

QM: Then the experimental results your theory predicts will no longer coincide with those of quantum mechanics, because momentum measurements will come out wrong!

LR: Oh yes, they will coincide! And you pointed out why they still can coincide!

QM: How come?

LR: You gave the example of a position measurement at a time t which is really, in part, a measurement of the momentum at an earlier time. If you think of it, the momentum is always measured indirectly in such a way: You may run the particle through velocity filters or scatter it off a grating, where the scattering angle depends on its momentum. But in the end you observe where it hits a screen, so you actually carry out a position measurement!

QM: Hm, at least it is an excuse that I cannot immediately prove wrong. However, it spoils the beautiful symmetry between position and momentum, which is basic for the theory of quantum mechanics.

LR: That does not matter to me. You could also claim that this symmetry is basic for classical mechanics, at least in Hamilton's formulation. But it is a merely formal symmetry. Physically, momentum and position are very different: The position can be observed directly, momentum only indirectly, by measuring at least two positions at different times. The position can have no jumps, because there is no motion faster than light. The momentum can easily jump arbitrarily fast, if a large force is applied suddenly. So they are not really similar.

QM: Agreed. But how do you think your trajectories look like? I mean, can you give any example that appears to be plausible physically?

LR: Of course, it is very hard to make up *any* example that works...
Maybe the trajectories are just random walks, where the particle takes a step in a random direction during each small time interval?

QM: Well, at least that would be a very simple possibility, if it worked...

LR: In fact, as I think about it, that is a particularly nice possibility, for the following reason: Random walks look "fractal", which means if you zoom into the details, they look similar to the larger picture. It also means that the "instantaneous" velocity is not defined for a random walk, in contrast to a smooth classical particle trajectory - the position $x(t)$ as a function of time is nowhere differentiable. This already fits nicely with the fact that, in quantum mechanics,

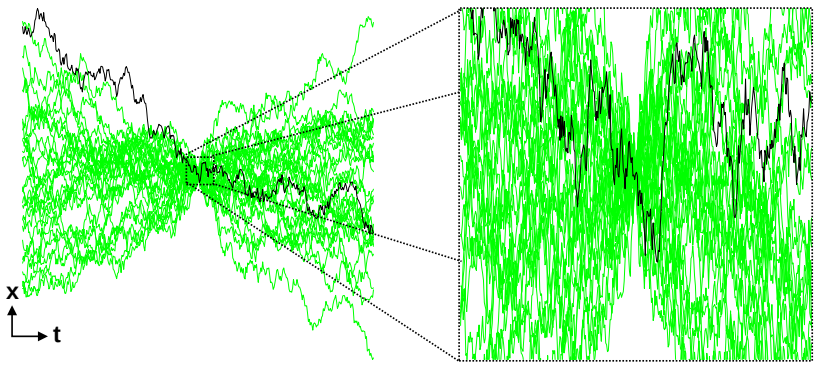

Figure 4: A couple of random walks that have been chosen to pass through the same point at a particular time. The space coordinate x is plotted vs. time t. A zoom-in is displayed to the right, which looks similar to the overall picture ("self-similarity", "fractal nature" of random walks).

the momentum is completely uncertain if you measure a definite position of the particle at a given time. But there is a much closer correspondence: Imagine one tries to find out the velocity - and therefore the momentum - by measuring the position twice, with a time interval t in between the measurements. Of course, because this is a random walk, there will be a spread in measured momenta...

QM: And the mean momentum is *always* zero! This is because the random walk has a preferred frame of reference - the average velocity vanishes. That rules out your theory.

LR: No. I just imagine there may be an average "drift" velocity - maybe a different one for each trajectory. Then, if one sets out to measure that velocity, one simply divides the change in position by the time interval. The position will have a component growing linearly with time, which is just vt, and an additional random component. The spread Δx of the random component grows only as the square root of time, as for any random walk. In order to obtain the velocity you divide by t; thus the spread in momentum will actually decrease with time, since you are able to measure the velocity ever more accurately. You see, we already have obtained what Heisenberg's uncertainty relation tells us: If the spread in momentum decreases, the spread in position increases.

QM: But the uncertainty relation involves Planck's constant, $\hbar$. Where does that enter your theory?

LR: Well, I simply have to choose the diffusion constant D correctly. Let's see: For a random walk, Δx is $\sqrt{Dt}$. We have $\Delta p = m\Delta v$, which equals $m\Delta x/t$, so if I multiply $\Delta x \cdot \Delta p$, I get mD. According to Heisenberg, I should have obtained something like $\hbar$. So I must have $D = \hbar/m$, up to some numerical

factor. The larger Planck's constant, the larger the diffusion constant, which means the random component grows faster. That is reasonable, because the classical limit must be reached for $\hbar = 0$ and then there would indeed be no random component at all.

QM: Diffusion constant equals $\hbar/m$! What a strange law!

LR: Oh, it is not strange at all! I can give you a very plausible interpretation: We have been taught that there is a minimum length scale down to which the usual quantum mechanics of an electron can be trusted, it is the Compton wavelength of the electron, $\lambda_C = h/mc$. That is about one hundreth the diameter of the hydrogen atom. If you try to localize an electron in a region that is smaller than this length, Heisenberg's uncertainty relation tells you that the velocity spread approaches the speed of light and the spread in kinetic energy becomes larger than the energy contained in the mass of the electron, mc^2. Then electron-positron pairs will be created and peculiar things can happen. Anyway, there is a time scale associated to the Compton wavelength, which is the time it takes light to travel the length λ_C, namely $\tau_C = h/mc^2$. That is about $1/10000$ of the time it takes the electron to go around the nucleus once in the hydrogen atom, so it is a very small time scale. Now in a usual random walk, if I know that during time τ the average distance traveled is l, the diffusion constant is $D = l^2/\tau$. Here I would claim that the random walk of the electron takes random steps of size λ_C during times of size τ_C. If I insert that, I get $D = \lambda_C^2/\tau_C = h/m$. There you are! It is interesting that the speed of light plays an important role at first but cancels in the end. This is even necessary, because the Schrödinger equation itself, whose results we want to reproduce, does not depend on c. Of course, the random motion during these short time intervals would proceed at the speed of light, but that is even to be expected and nothing to worry about. It seems we are up to something...

QM: Hm, I would call that yet another unmerited success obtained with dimensional analysis! You see, the way you have introduced Δx and Δp is just not honest; it does not correspond to the physical situation we talk about in quantum mechanics. Imagine a wave packet, part of which goes through a single slit of width Δx. Then we will observe that, after the slit, the wave will spread in the direction perpendicular to the initial velocity, because a momentum spread Δp has been introduced by the constriction of size Δx. But in your theory, you would have to start out with an ensemble of trajectories whose drift velocities all coincide more or less with the average velocity of the wave packet before the slit. Then, a few trajectories would pass the slit by chance, while most would be reflected. But there is no way that this selection of a few trajectories would produce the correct spread of the average drift velocities of the trajectories which have passed through the slit. In your model, they would not change at all. And your observation that the random component Δx of a random walk grows with time while the spread of the velocity decreases is just completely irrelevant.

LR: I admit that it is not helpful here. Obviously, I would have to assume that

the borders of the slit excert some forces on the particles. In fact, that is even true in quantum mechanics: For when I start out with a wave packet which is smaller than the slit, it will pass through the slit unhindered and will *not* acquire an extra momentum spread after the slit. So it is not affected if it does not touch the borders of the slit. The wave packet must feel the potential that defines the slit. And if the particles are indeed reflected there, the average drift velocity of the trajectories might change as well. Of course, I cannot immediately give you the right laws of motion...

QM: You would face still larger difficulties with the double slit experiment. I suggest we do not try to figure out how the detailed laws of motion might look like, at least for the moment. Let us just ask about some general features which the trajectories should have in order to fit the experimental observations predicted by quantum mechanics.

LR: OK. This is probably the most efficient way to proceed. The most important requirement is that, at every instant of time, the probability distribution of our imagined ensemble of electrons should be the one prescribed by quantum mechanics, $|\Psi(x)|^2$. In fact, I think this is the *only* requirement.

QM: Yes - it is the only requirement if we believe your assumption that, in the end, every measurement is a measurement of position. Let us begin our discussion with some very simple situations. For example, take a wave packet moving along in free space. Then all the particles of our imagined ensemble should better have the same velocity - at least on the average. Otherwise, the probability distribution would not move according to quantum mechanics.

LR: Maybe they just move along smooth trajectories, like classical free particles, with their velocity equal to the velocity of the wave packet as a whole. But we cannot be sure. Perhaps, on top of that motion, they also perform some random walk. Then, a particle that starts out in the middle of the wave packet would not remain there all the time but it could wander around.

QM: If we have several wave packets that do not overlap, we can say something about the average velocities of all the particles inside each wave packet separately. Besides, we could also look at the spreading of a wave packet. Then, we might learn something about the spread of velocities of the particles inside the wave packet.

LR: Wait a moment - we are always talking of velocities. But what if we really have to deal with random walks? Then the instantaneous velocity is not even defined.

QM: That is true. For these cases, one must be more careful about the calculation. We should find that the average *drift* velocity of all the particles equals the velocity of the wave packet. And the spreading of the wave packet would not only depend on the spread of drift velocities but also on the diffusion constant. Maybe the diffusion constant could even be different for particles at different positions and with different drift velocities. Who knows, maybe it is not even

a random walk but something even more complicated. In the usual random walk, the direction of each random step is independent of the previous history. Perhaps this is not true for the trajectories of our particles. There might be correlations in time, such that the random step the particle takes now depends on what it did a long time ago.

LR: There seems to be an abundance of possibilities. By the way, how does it look like for the double slit experiment that we started to discuss above? At first, we have a wave packet moving towards the slits, so all particles move with the same average velocity. But when the packet hits the slits, a complicated pattern of standing waves develops. Afterwards, some particles must have gone through the slits and others will have been reflected. And those that have passed through the slits will be inside one of the several wave packets that form the interference pattern and move on towards the screen. The average velocities of these packets are different, they move along different directions.

QM: Of course, again we cannot tell what the particles do inside each packet, apart from giving their average velocity. But one fact is very clear: They do not cross the boundaries between the wave packets, provided the wave function really vanishes there. The places with $\Psi = 0$ are forbidden for the particles.

LR: All right. However, that does not give us much information about the trajectories.

QM: Oh, I believe it can already tell us a lot, in particular if we have a situation where $\Psi = 0$ holds on a full plane dividing space in two halves. Then this plane can never be crossed by particles. Admittedly, the double slit setup is still a bit too complicated. The simplest situation would be this: Imagine two wave packets of equal size crossing each other. As long as they overlap, they form an interference pattern and we might have a plane with $\Psi = 0$.

LR: Of course this may only hold at a certain instant of time. Then the particles would not be allowed to cross that plane, but only at that particular time. Again, it would not be very helpful to know that.

QM: That is the most common situation. However, I believe we can set up the gedanken experiment in such a way that there is always a plane with $\Psi = 0$. First of all, the packets should be identical, apart from moving in exactly opposite directions. Otherwise, a full cancellation is unlikely, since the waves do not have the same magnitude. Secondly, the interference pattern depends on the relative phase between the two wave packets, since this decides where we get destructive rather than constructive interference. Now, if we set up the two waves to have opposite signs, we will get a plane $\Psi = 0$ exactly in the middle between the two packets. This persists all the time.

LR: OK. So we know the trajectories do never cross that plane...

QM: You should admit that this result is extraordinary from the usual point of view. After all, any particle in the wave packet coming from the left *must* be reflected to go back where it came from. Likewise for the particles in the packet

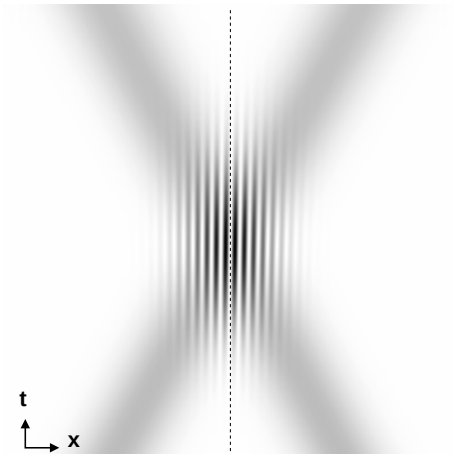

Figure 5: A space-time plot showing the evolution of the probability density $|\Psi(x,t)|^2$ for two wave packets crossing each other (in one space dimension). The dashed line belongs to $\Psi = 0$, which is the point (or, in three dimensions, the plane) that cannot be crossed by any trajectory.

coming from the right. This is in stark contrast to what one would reasonably expect. After all, the particles inside a single wave packet in free space move along with the packet, more or less as expected from a classical point of view. Maybe they also execute some random walk, but that is not very surprising. However, when the two packets cross, the trajectories do *not* cross, as they certainly would if this were a classical situation. Instead, they are reflected!

LR: This is indeed very strange. After all, there are no forces involved! Everything takes place in free space. Besides, after the wave packets have crossed, they proceed just with the same velocities as before. In the end, the particle will be detected either to the left or to the right. Judging from that result alone, we would conclude that, just as in the purely classical situation, nothing has happened to the particle on its way to the detector. But your argument with the plane $\Psi = 0$ indicates that, in fact, the particle has to be reflected on its way, such that a particle coming from the left also ends up going left.

QM: It is even more surprising if one remembers that the trajectories would not have been reflected if the packets had actually missed each other - I mean, if we think of a fully threedimensional situation, where each packet is a spherical cloud. Classically, it could never make a difference for the trajectories whether the packets happen to meet or do not meet, since they only represent the probability of having a *single* particle moving along this way or that way.
We could even imagine a situation where we have a choice of whether the two wave packets meet and therefore whether the trajectories get reflected: After the two wave packets have been separated at a beam splitter, and before they

are recombined, one decides whether one of them should be blocked. For example, one may put a detector in its path. Alternatively, one could simply remove the mirror that is normally used to reflect the packet, such that it will not meet the other one any more. Then, we would find that a particle taking the *other* way, where nothing is changed, is reflected or is not reflected at the crossing point depending on what we decided to do at the other end of the whole interference setup!

LR: Wait a moment. Actually, I begin to have some doubts about your assertion that trajectories can never cross a plane with $\Psi = 0$. Of course, the square of the wave function is used to calculate the probability of finding the particle in a small volume around a given point. Therefore, our theory should correctly reproduce this quantum mechanical prediction if we do not want to risk contradictions with the measurement results in a real experiment. But $\Psi = 0$ only holds *exactly* on a certain plane in space. If, in a measurement, one finds the particle slightly removed from that plane, there is no contradiction. And remember, *real* measurements of the position cannot have unlimited precision. At the very least, the precision probably will be limited by the Compton wave length. If you did any measurement with a precision higher than that, using light of very small wave length, electron-positron pairs would be created and you could no longer tell which particle has been your original electron. Perhaps this resolves the whole mystery - trajectories could pass through a $\Psi = 0$ plane without producing a conflict with real measurements.

QM: Again, I see you are not easily convinced. I will have to make up a rigorous argument.
Imagine every particle actually crosses that plane, as we would expect it to do classically. Then each trajectory has to spend some time Δt near that plane - say, inside a volume of extent Δx in the direction perpendicular to the plane. Of course, Δt will vary and sometimes a trajectory may even cross that region several times. Then Δt should be considered the total time spent inside the region. I want to show that Δt has to be so small that the velocities of particles going through that region come out larger than the speed of light!
Now let us ask: What is the average Δt? If you think about it, it is the space-integral of $|\Psi|^2$ taken over that region, and afterwards integrated over all times! Even if you do not understand it immediately, you can see that it is plausible: If Δt is larger, it means the trajectories spend more time inside the region, which must mean that either the probability of finding the particle inside the region is relatively large at certain times or that the time it takes for the two packets to cross is quite large or both. Now, if we know the average Δt, we can conclude that the maximum velocities reached by some particle trajectories cannot be less than $\Delta x/\Delta t$. If some relatively slow trajectories spend more than Δt inside the region, others have to be faster than that. You can then easily estimate this maximum velocity, for a situation where the wave packets move with velocity v: It turns out to be $v\,\lambda^2/\Delta x^2$, where $\lambda = h/mv$ is the wave length of the packets. The wave length is important, because it determines how steep the interference minimum at $\Psi = 0$ really is and how large the integral over

$|\Psi|^2$ is. In any case, if we could make Δx tend to zero, the maximum velocities would become infinite! But even if I accept your argument that the precision of measurements is limited by the Compton wave length, so $\Delta x = \lambda_C = h/mc$, I still get maximum velocities that are larger than the speed of light, namely by a large factor c/v.

LR: I agree, this does not work. So it seems we have to face the fact that, regardless of how sophisticated our trajectories may be, they will not cross the $\Psi = 0$ plane and they will be reflected. As you have said, that may even depend on changes made to the experimental setup at places that the given trajectory never comes close to.

QM: Of course, everything also depends on the relative phase between the two wave packets. If we did not have the opposite sign, but the same sign of the wave - or some different phase in between - then we could still get $\Psi = 0$ planes at certain locations and certain times, but they would not persist all the time and you could no longer demonstrate that the trajectories should be reflected.

LR: What determines the relative phase?

QM: It is simply the length difference between the two paths that are followed by the wave packets. If we move one of the mirrors slightly, then the path length changes and so does the relative phase. This is another example of a "nonlocal" influence. After all, moving the mirror will also affect trajectories that may actually have passed the *other* mirror, not the one we move. They will feel the dependence of the interference pattern on the relative phase once they go through the region where the wave packets cross.

LR: Does this influence go faster than light??

QM: No, at least we cannot claim that to be the case. This is because the change in relative phase, or the change produced by blocking one of the wave packets, will become important only at the time when the wave packets cross. But since the packets move slower than light, there would have been plenty of time for a light signal emanating from the mirror to reach that region before the trajectories are actually changed.

LR: I am glad to hear that.

QM: Well, just wait...
But for the moment, I just want to point out that there is another, more elegant and striking way of changing the relative phase between the paths. It is done by placing a coil somewhere between the two paths. If an electric current flows through the coil, it produces a magnetic field inside the coil, but not on the outside. Then, the phase difference will start to depend on the magnetic flux inside the coil. This is the famous "Aharonov-Bohm" effect.

LR: That is even more suprising, because the wave packets do not pass near the coil that contains the magnetic flux! They never enter the region inside the coil where the magnetic field could exert a force on the particle. It occurs to me that *this* is a situation where the influence may travel faster than light. After all,

the phase difference depends instantaneously on the flux inside the coil, which may be far away.

QM: No, it is not really *instantaneous*. If you think about it, during the change of the magnetic flux an electric field is generated around the circumference of the coil. This electric field pulse travels outward, at the speed of light. Any influence on the wave packets will not propagate faster than that. If they are reached by the electric field pulse before they overlap, their momenta will be changed but the details depend on their positions and directions of flight. There is a unique correspondence between the phase shift and the flux only if the magnetic flux remains the same during the time that the packets are separated.

2 David Bohm's pilot wave

LR: I must admit that the trajectories of quantum-mechanical particles behave in a strange way. For example, you have demonstrated to me in that situation with the interfering wave packets that the particle will always be reflected at the point where the waves meet - although there should be no forces in free space. And this holds completely generally, whatever the microscopic laws of motion might turn out to be! I hate to say it, but given this sort of behaviour, perhaps one could go further and prove that there is no way of having such trajectories at all? Perhaps there *are* no consistent microscopic laws of motion? Up to now, we have only listed a few features which they would have to have *if* they existed.

QM: Well, they do exist. It is time to tell you about Bohm's theory. He made up a microscopic law of motion that works.

LR: I am very curious. Does it involve random walks, as I expected?

QM: No, it is simpler than that. In fact, you could call it the simplest possible "hidden variable" theory.

LR: What is "hidden" in "hidden variable" theories?

QM: It is the trajectories that are hidden, because you cannot really measure them without disturbing them. Remember the discussion on Heisenberg's gamma-ray microscope. You will never be able to verify whether your predictions about the trajectories are correct. Which means, of course, you could as well do without them. It also means that there is more than one possibility and you cannot decide in favor of one or the other. But you wanted to have trajectories, so here they are.

LR: So how does it work?

QM: We have to fulfill two requirements: First of all, we want to have continuous trajectories. And secondly, we want to reproduce the predictions of quantum mechanics. If we assume that it is enough to reproduce correctly the

outcome of position measurements, then this means: if we consider a situation that quantum mechanics describes using a particular wave function, our "hidden variable" theory should always give the correct probability of finding the particle around a certain point, namely the square of the wave function: $|\Psi|^2$. It is not self-evident that the two requirements are compatible. Imagine we have two wave packets that are separated by a large distance. If $|\Psi|^2$ *could* change suddenly, such that the probability of finding the particle in the first packet decreases by a certain amount and this is compensated for by an increase for the other packet, then we could not invent continuous trajectories that did the job. The particle would have to jump instantaneously from one packet to the other, with some probability. Fortunately, this does not happen in quantum mechanics. The probability density $|\Psi|^2$ always changes continuously. Moreover, its increase at a certain point is given completely by the influx from neighboring points.

LR: Yes, I know. It is similar to liquids or gases. There, you have some density ρ and a current density $\vec{j}$ that describes the streaming of particles. Both are functions of position and time and they are coupled by the "equation of continuity"

$$\frac{\partial \rho(\vec{r},t)}{\partial t} = -div\, \vec{j}(\vec{r},t)$$

which means the density ρ grows or shrinks where the vector field $\vec{j}$ has its sinks or sources. So the change in ρ is indeed determined by the influx from (or outflux to) neighboring points, because the latter is given by the divergence of $\vec{j}$.

QM: The same holds for quantum mechanics. There, we have $\rho = |\Psi|^2$, describing the probability density. Obviously, the physical situation is different, we are only dealing with a single particle and the probability to find it here or there. But the mathematics is the same. We have a probability current density $\vec{j}$, such that the equation of continuity holds. By the way, I hope you do remember how the current density for a liquid can be related to the density, apart from the equation of continuity?

LR: Why, yes, of course. In the simplest case, the streaming of particles may be described by a velocity field $\vec{v}(\vec{r},t)$: If a particle happens to be at $\vec{r}$ at time t, then its velocity is given by $\vec{v}(\vec{r},t)$. In that case, the current density $\vec{j}$ is larger wherever the particles are streaming faster or there are more of them, it is simply the product of ρ and $\vec{v}$:

$$\vec{j}(\vec{r},t) = \rho(\vec{r},t)\vec{v}(\vec{r},t)$$

We can also use $\vec{v}$ to calculate the full trajectories $\vec{r}(t)$ of fluid particles, simply by setting

$$\frac{d\vec{r}(t)}{dt} = \vec{v}(\vec{r}(t),t)$$

All of this works perfectly for a liquid. But for a gas, $\vec{v}$ would only be the *average* velocity of particles at point $\vec{r}$ and time t. So if you pick a certain particle, it will, in general, have a random velocity, distributed according to the Boltzmann distribution of statistical mechanics, with the average equal to $\vec{v}(\vec{r},t)$. In that case, you could not obtain the full particle trajectory by knowing only $\vec{v}$. Besides, for that situation there is also another contribution to the current density $\vec{j}$, describing the diffusion of particles.

QM: Right. Now, the simple idea of Bohm's theory is to imagine an ensemble of particles, whose density is described by $\rho = |\Psi|^2$ and that are moving like the particles inside a liquid. So we obtain the trajectories just like you have explained it for the liquid! Of course, for this we need a velocity field $\vec{v}(\vec{r},t)$. Quantum mechanics does *not* provide us with any. But it gives us the density $\rho = |\Psi|^2$. Then, if we also know the current density $\vec{j}$, we can simply set $\vec{v} = \vec{j}/\rho$. You know the quantum-mechanical expression for $\vec{j}$?

LR: Yes. Given the wave function Ψ, we can write it as a magnitude times a phase factor, $\Psi = |\Psi|\, e^{i\varphi}$. Then the current density points along the gradient of the phase φ:

$$\vec{j}(\vec{r},t) = |\Psi(\vec{r},t)|^2 \, \frac{\hbar\vec{\nabla}\varphi(\vec{r},t)}{m}$$

In fact, for a plane wave $\Psi \propto \exp(i(\vec{k}\vec{r} - \omega t))$, we just get $\vec{j} = |\Psi|^2 \, \hbar\vec{k}/m$. This is very reasonable, since $\hbar\vec{k}$ is equal to the momentum $\vec{p}$ of the particle, so $\vec{j} = |\Psi|^2 \, \vec{p}/m$.

QM: This means: if we are to *define* our velocity field $\vec{v}(\vec{r},t)$ by setting

$$\vec{v}(\vec{r},t) = \frac{\vec{j}(\vec{r},t)}{|\Psi(\vec{r},t)|^2} = \frac{\hbar\vec{\nabla}\varphi(\vec{r},t)}{m}$$

we get the most reasonable result for a plane wave of momentum $\vec{p}$: $\vec{v}(\vec{r},t) = \vec{p}/m$. The same would be true for a wave packet: The particles in our imagined ensemble would move, more or less, with the average velocity of the packet. No extra random motion is involved.

LR: What will happen in the ground state of the hydrogen atom? There, the phase does not depend on position. So, according to Bohm's theory, the velocity is zero?

QM: Right. The electron sits here or there, but it does not move.

LR: Isn't this quite strange? Wouldn't we expect the electron to move around the nucleus with a typical velocity of around 1 percent of the speed of light? At least that's what Bohr's old quantum mechanics tells us. And it is also what one expects if one calculates the probability distribution of the momentum from the ground state wave function.

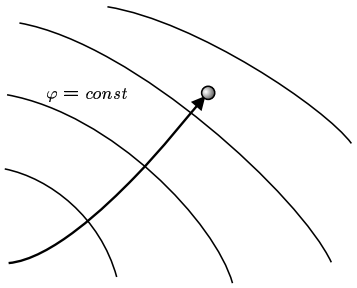

Figure 6: According to Bohm's theory, the particle moves along the gradient of the phase, perpendicular to the "wavefronts" of constant phase.

QM: You are right. In fact, one could even try to argue that this could lead to a conflict of Bohm's hidden variable theory with experiment! The argument goes like this: There are hydrogen atoms that contain a muon instead of an electron. The muon is unstable and may decay after a certain time. However, if it is in the ground state of the hydrogen atom, its decay time turns out to be longer than usual. One can explain it by noting that it moves at a high velocity, such that the relativistic time dilation will make the decay appear slowed down. In this way, the muon decay acts as a kind of clock that tells us about the velocity of the muon. But you cannot really demand Bohm's theory to reproduce this. After all, it builds on the nonrelativistic Schrödinger equation, while here we are talking about a relativistic effect that reveals the nonzero velocity of the particle in the ground state of the hydrogen atom. Still, a finite velocity would look somewhat more reasonable, I agree.

LR: What is the prediction of Bohm's theory for the trajectories in that situation we discussed before, with two interfering wave packets? I would like to see how the particles are actually reflected.

QM: You can easily calculate that: First you need to know the wave function $\Psi(\vec{r}, t)$, which you get by solving the Schrödinger equation. Then, you obtain the current density and from that the velocity field and the trajectories. Of course, in every run of the experiment, the single particle follows only one of these paths. But when you repeat the experiment, the initial condition will be different and the trajectory will be another one. There is, actually, a further assumption involved in the theory: At some time the probability distribution must be given correctly by $|\Psi|^2$. If that holds for one particular point in time, it will hold always, thanks to the equation of continuity, and because we have set up the trajectories such that they give the correct $\vec{j}$.

LR: I see. The particles seem to go faster in the interference minima and spend more time in the maxima. That is how the probability of finding a particle at the maxima is enhanced. All of the particles are reflected. I guess it *must* be

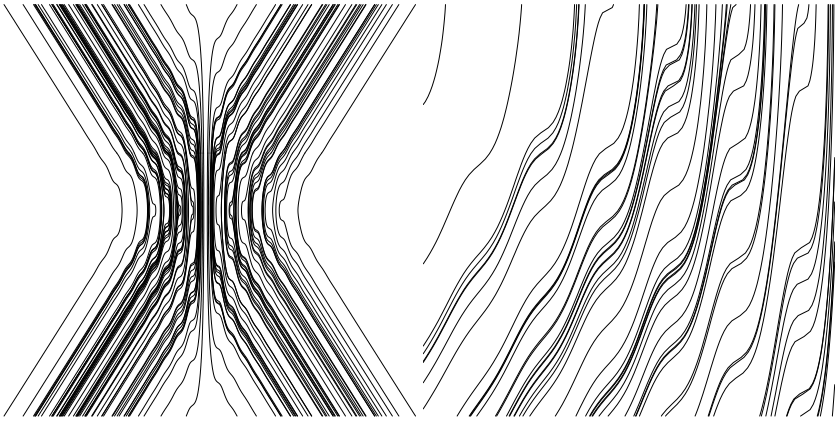

Figure 7: Space-time diagram showing particle trajectories according to Bohm's theory for a situation with two wave packets crossing each other (in one space dimension; time runs upward). The initial positions have been picked at random, according to the probability density $|\Psi|^2$. See Fig. 5 for a plot of $|\Psi(x,t)|^2$ for this situation. Magnification to the right: Velocities are highest (slope $\Delta t/\Delta x$ smallest) whenever trajectories cross regions of low probability density.

like that for the simple reaspon that the trajectories never cross.

QM: Yes, in a one-dimensional situation, the Bohmian trajectories never cross, because the velocity is a unique function of the position. So, for Bohm's theory we could have proved the reflection of trajectories much easier, at least in one dimension. Alternatively, we could simply have calculated the trajectories without further thinking about why they should come out the way they do. But it is important to know that we would necessarily end up with the same behaviour even for completely different microscopic laws of motion.

LR: It seems the particles can move very fast in the interference minima. Is it possible that they may exceed the speed of light?

QM: You are right. One can prove that in this situation, the particle velocities calculated from the current density can become larger than the speed of light at the minima! This is perhaps astonishing, since the wave packets themselves may move at small velocities. However, it happens only in a very narrow range around such a minimum - in fact, in a range that is necessarily on the order of the Compton wavelength λ_C. Now we have already argued that we cannot trust a nonrelativistic theory on such a small length scale. Thus, it is probably an artefact that does not really speak against Bohm's theory.

LR: Are there other situations where Bohm's theory gives particle velocities exceeding the speed of light?

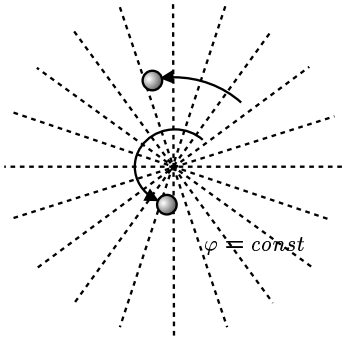

Figure 8: According to Bohm's theory, a particle in a p-orbital will move faster if it is nearer to the nucleus, since the phase gradient diverges when approaching the origin.

QM: Yes: Take an excited state of the hydrogen atom that has a finite component of the angular momentum - for example, a p-orbital. Then, the phase is given by the angle in polar coordinates. If you calculate the gradient of φ, you will see that its magnitude grows without bounds as you approach the origin. Thus, a particle close to the origin may move faster than light. But again, this happens only at distances of the order of λ_C where we should not be surprised to find some artefacts in a nonrelativistic theory.

LR: I believe I know a situation where the particles move faster than light but not only on a scale of λ_C. Consider the following: If you have, say, a p_z orbital of the hydrogen atom, then space is divided into two halves by a plane where the wave function vanishes. Now, if we are able to rotate that plane with a certain angular frequency, the following will have to happen: Since the particles cannot cross the plane, they will have to move. After half a revolution every particle that had been in the right half-space will have to be in the left one and vice versa. It is easy for the particles near the origin to switch sides, but those that are sufficiently far from the origin will necessarily have to go faster than light! This argument is completely general, it does not really depend on Bohm's theory, it will work for any theory that reproduces the quantum mechanical probability density at each instant of time.
Of course, you may ask: how do we rotate that plane? The answer is simple - apply a magnetic field that points along the plane, say in x direction. Then, you may consider the initial p_z orbital to be a superposition of $m = +1$ and $m = -1$ states for the x-quantization direction. Since the magnetic field lifts the degeneracy of the energies of those two states, the relative phase in this superposition will change in time and this just gives you the desired rotation of the plane where $\Psi = 0$. Voila!

QM: I agree that in Bohm's theory the particles will indeed move faster than

light for your example, if they are sufficiently far from the origin. But consider how far they have to be: If the angular frequency of rotation is not larger than the typical transition frequency between different atomic orbitals, then they should be more than 100 atomic radii away from the nucleus. At that distance, the wave function is extremely small, such that the probability of finding a particle out there is almost zero. This means that subtle relativistic corrections may apply that are not important for the bulk of the wave function. Indeed, the picture of a simple change in energies of the $m = +1$ and $m = -1$ states is no longer correct at these distances, because the contribution of the magnetic field to the Hamiltonian is no longer a small perturbation. Anyway, if the current density yields a speed larger than that of light, it seems we are overtaxing the nonrelativistic theory in any case.

LR: OK. So is there a relativistic extension to Bohm's theory? Taking Dirac's equation rather than Schrödinger's equation as a starting point?

INSERT DISCUSSION ON RELATIVISTIC THEORIES

LR: Are there modifications to Bohm's theory? How much freedom do we have in choosing the trajectories?

QM: That is a good question. Indeed, there are plenty of possibilities to modify the theory. For example, we can change the current density $\vec{j}$ as long as its divergence stays the same. Then the equation of continuity will still give the correct time evolution of the probability density, that is compatible with the quantum mechanical evolution. So we can add any divergence-free vector field to $\vec{j}$ and use the new current density to calculate the velocity of the particle. The resulting trajectories will be different but the coincidence with the quantum-mechanical results will remain intact.

LR: OK. In this way we could make the electron move around the nucleus in the ground state by adding some vector field that winds around it. But is it also possible to add some diffusion?

QM: Indeed, we can do that. Just look at the situation in a gas, where the diffusion of particles may be as important as their drift motion. There, the current density contains an extra term,

$$\vec{j} = \rho\vec{v} - D\vec{\nabla}\rho$$

which is proportional to the diffusion constant D and to the gradient of the density ρ. It will lead to the diffusive spreading of the density. Now we have to be careful, since the divergence of this extra term does not vanish, so it could change the time-evolution of the density. We must modify the definition of the drift velocity, such that the current density remains the same - equal to the quantum mechanical result. But this is simple enough. Just use the equation above as the new *definition* of the velocity:

$$\vec{v} \equiv (\vec{j} + D\vec{\nabla}\rho)/\rho = \frac{\hbar\vec{\nabla}\varphi}{m} + D\frac{\vec{\nabla}\rho}{\rho}$$

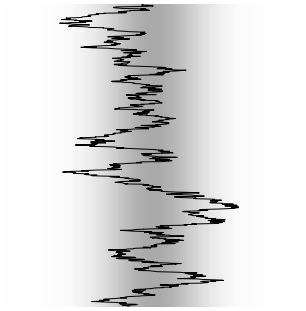

Figure 9: A particle trajectory according to Nelson's theory for the case of the harmonic oscillator, in a space-time diagram. The shading gives the stationary probability density of the ground state. This kind of random walk, with a drift velocity proportional to the distance to the origin, is called an "Ornstein-Uhlenbeck process".

Of course, now this is only the *average* velocity of a given particle. There is an extra random walk on top of that, with a diffusion constant equal to D.

LR: So we can set the value of D equal to $\hbar/m$? Then we are back to our previous idea about random walks that derive from a jitter motion taking place on the length-scale of the Compton wave length!

QM: Yes. But you can as well pick *any* other value of D! If you choose the special value $D = 0$, there is no random walk and you are back to Bohm's theory. There is an infinity of possibilities.

LR: I see. But the choice $D = \hbar/m$ gives physical results that seem to be the most reasonable: In particular, the order of magnitude of the velocities comes out as expected. The ground state of the hydrogen atom has an extent a, so the magnitude of the density gradient is roughly $\vec{\nabla}\rho \sim \rho/a$. Therefore, we will get a random walk with a typical drift velocity of about $D\vec{\nabla}\rho/\rho \sim \hbar/ma$. This is just what you would estimate from the momentum spread $\hbar/a$ obtained from Heisenberg's uncertainty principle.

QM: Still, there is no particular reason why you shouldn't choose $D = 2\hbar/m$, for example. And as I have told you before: There is no experimental way to distinguish between these possibilities, since they are all compatible with quantum mechanics. If there were any such experiment, quantum mechanics would be "incomplete", i.e. it couldn't make a prediction for something which can be measured.

LR: Perhaps one could turn this wealth of possibilities into a virtue: We could put other demands on our theory, instead of just having the correct probability density at each instant of time. For example, we could demand to have the right momentum distribution as well.

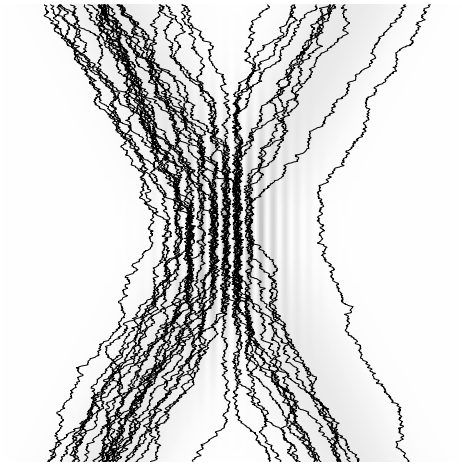

Figure 10: Particle trajectories according to Nelson's theory for the case of two interfering wave packets (compare Fig. 5), for some randomly chosen initial conditions. In contrast to Bohm's theory (compare Fig. 7), different trajectories can cross, but still they never cross the middle of this picture, where $\Psi = 0$.

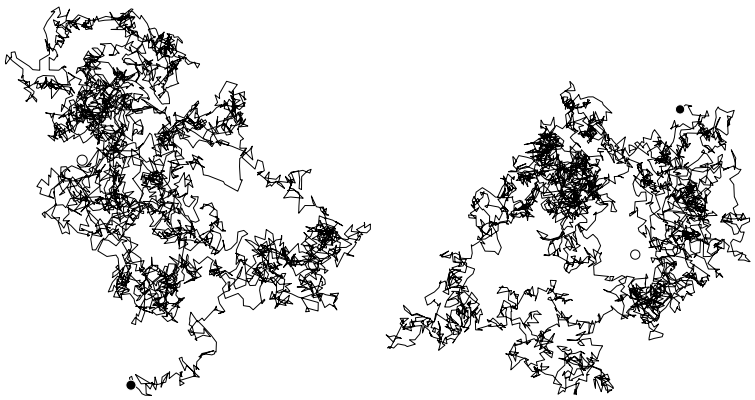

Figure 11: Left: A trajectory according to Nelson's theory for the hydrogen atom ground state, projected onto the xy-plane. The white dot marks the nucleus, the black dot the final position of the electron. Right: A trajectory for an excited state with angular momentum (p-orbital, $m = +1$). The electron has some tendency to move counter-clockwise around the nucleus, not easily discernible in the picture.

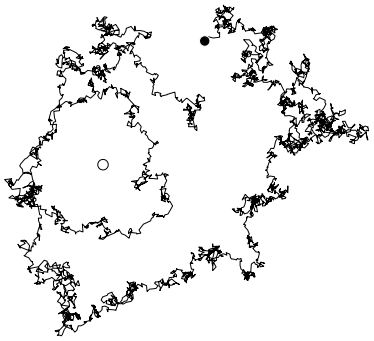

Figure 12: Another particle trajectory for the hydrogen p state with finite angular momentum projection along the z-axis. In this case, the value of the diffusion constant has been chosen smaller, such that the drift along the gradient of the phase, associated with the circular motion, becomes visible more easily. The resulting probability density is still equal to $|\Psi|^2$.

QM: ...or perhaps the right correlation function of the position!
In classical statistical mechanics, when we are dealing with a particle whose position x is statistically distributed around its average value, say around zero, then we can calculate $\langle x^2 \rangle$ to learn something about the spread of position x. We can extend this to learn something about the time evolution of our system, by calculating the correlation function $\langle x(t)x(0) \rangle$. Its magnitude tells us whether the value of x at time t is still statistically correlated with that at time 0 and its sign tells us whether the value of x at time t has a tendency to be opposite to the initial value, perhaps because we are dealing with an oscillator that swings back and forth between positive and negative values of x. After half an oscillation period, the sign of x would be opposite to the initial value, whatever that has been.

LR: I see. It would be nice to have a theory where the trajectories also give the correct quantum-mechanical result for the correlation function of the position, not only the correct distribution of the position at each single time point. Then the theory would capture more of the real dynamics of the system. And, I guess, this would rule out Bohm's theory, because quantum-mechanics will not predict the electron to sit at rest, if it is in the hydrogen atom's ground state!

QM: You are right. Let us test these things for the simple harmonic oscillator. The quantum-mechanical correlation function of the position looks very reasonable: At $t = 0$ it is of course equal to the spread $\langle x^2 \rangle$ in the ground state, and after half a period it is again of the same magnitude, but with an opposite sign. This just means the position $x(t)$ will be minus the value of $x(0)$ after half a period. Besides, the correlation function is periodic, since the motion itself is periodic, just as for the classical harmonic oscillator. In fact, the quantum

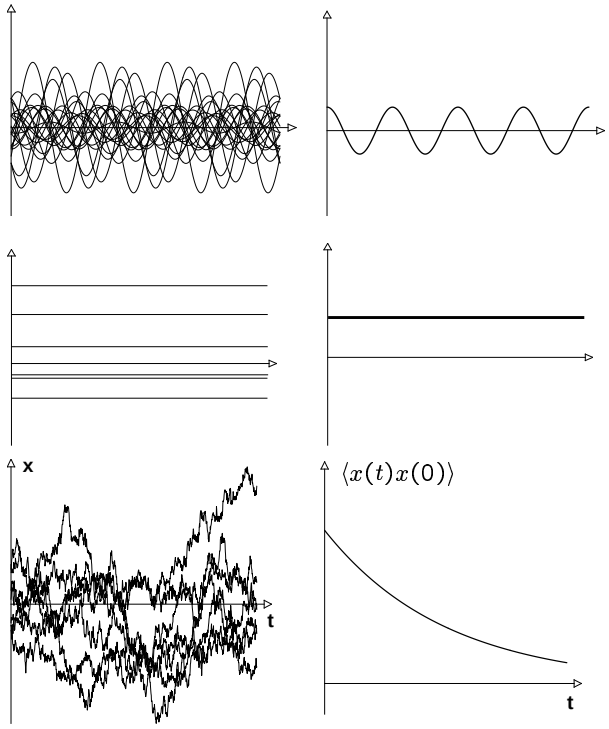

Figure 13: Correlation function of the coordinate x of the harmonic oscillator (right) and corresponding ensemble of trajectories (left), for three different descriptions: Statistical ensemble of particles moving according to classical equations of motion (top), ensemble of Bohmian trajectories (middle), and ensemble of trajectories according to Nelson's theory (bottom).

mechanical and the classical oscillator are practically the same, if you impose the correct probability distribution; I will come to that in a minute.
Now look at the prediction of Bohm's theory: In the ground state, the particle always sits at the same place, so $x(t)$ will not change over time, which also means the correlation function does not change. This is in complete disagreement with the oscillations. For the theory including diffusion, it is different, but not necessarily better: There, if you start out with $x(0)$, the average position will have a tendency to drift towards the bottom of the potential, and this is reflected in the decay of the correlation function. Of course, there will be an extra random walk on top of that, but since it is equally likely to lead to higher or lower values of the position, it averages away when calculating the correlator.

LR: So we have ruled out both of these hidden variable theories?

QM: No! The fact that their trajectories do not reproduce the quantum-mechanical correlation function does not mean that they get into conflict with experiments. The point is: There is no experimental procedure of measuring the quantum-mechanical correlation function! This is because, if you were to measure the position x at time 0, you would severely disturb the system and the particle would obtain a momentum spread which would even be infinite if the position measurement were infinitely accurate. So the dynamics would be altered completely and the distribution of x at time t would be quite different from what it would have been if you had not measured at time 0. You cannot measure both $x(0)$ and $x(t)$ in order to multiply their values and average over many runs, or at least if you try this, the results would severely depend on your measurement precision and have nothing to do with the quantum mechanical correlator.

LR: From what you say I would rather conclude that this correlator is quite a useless and unphysical quantity!

QM: No. First of all, if you go towards superpositions of higher excited states, such that the dynamics becomes more and more classical, the quantum-mechanical correlator goes over into the classical version. And besides, in the quantum-regime, the correlator is directly connected to the linear response of the system, via the "fluctuation-dissipation theorem", similarly to the connection in classical statistical mechanics. And the linear response of the system to a small external perturbation is certainly measurable experimentally.

LR: Fine. So we should rather demand the theories to reproduce the linear response.

QM: But they do that anyway, without difficulties. After all, linear response means applying a small force initially and then measuring, for example, the average change in the position after some time t. Since the dynamics of our hidden variable theories are derived from the Schrödinger equation itself and since they reproduce the correct probability density for the position at every instant of time, they give just the correct answer. Of course, if you were to ask for the average momentum, they could not give directly the proper answer. You

would have to remember that finally every measurement boils down to a measurement of position, so you just have to include the momentum measurement apparatus in your description, as we have discussed before.

LR: You have mentioned some similarity between the classical and the quantum mechanical harmonic oscillator?

QM: Yes. The harmonic oscillator is very special in many respects. In particular, you can reproduce an important part of the quantum mechanical results by using the classical equations of motion. If the quantum oscillator is in the ground state, its position and its momentum have a certain spread. You can set up a statistical ensemble of particles inside a classical harmonic oscillator that have the same distributions of position and momentum. Then, these distributions are stationary, i.e. they do not change over time if you let the particles oscillate according to the classical equation of motion! If you tried to set up the same for the ground state of a quantum-mechanical particle in a box, you would fail.

LR: What about the correlation function?

QM: It is reproduced entirely within this classical statistical picture! And there is more: You also get right the response to an arbitrary external time-dependent force. This is remarkable, since you do not use the Schrödinger equation, in contrast to Bohm's theory or similar hidden variable theories. It just comes out correctly from the classical dynamics. You can even apply an arbitrary time-dependent variation of the oscillator frequency, that is you change the steepness of the oscillator parabola. This will make the width of the particle distribution oscillate in time, something which is called "squeezed states" in quantum mechanics. But you do not need quantum mechanics to understand these, classical mechanics of such a statistical ensemble of particles will give you the right answers.

LR: This is indeed astonishing. Tell me: where, in the classical description, do appear the energies and wave functions of the excited states of the harmonic oscillator?

QM: Well, nowhere! You never need to know about them in the classical description. For example, if you apply a time-dependent force to the oscillator, the Gaussian particle distribution will start to oscillate in the potential well. Quantum-mechanically speaking, the wave function becomes a superposition of excited states, more precisely a so-called "coherent state". In the classical picture, Newton's equation of motion automatically leads to the same results. The same applies to collections of oscillators, coupled together by linear forces, for example a vibrating molecule. If there were only oscillators, we really would not need quantum mechanics.

LR: Is *everything* reproduced by the classical picture in the case of oscillators?

QM: Almost everything. Of course, if you drive with a force that depends on position in a nonlinear way, then your equations of motion are no longer

linear and you are leaving the regime where quantum mechanics and classical mechanics yield the same answers. And then, there are some formal aspects: For example, quantum-mechanics tells you that the energy does *not* fluctuate if you are in the ground-state of the oscillator, while classical mechanics gives you a finite spread of energy. However, this discrepancy does not affect the dynamics nor any results for measurements of position and momentum.

LR: Isn't it very strange that it just works perfectly for harmonic oscillators?

QM: Well, linear equations of motion are always special. In particular, you can, say, add some high-frequency fluctuations on top of a trajectory without disturbing this low-frequency average motion. As soon as you go towards nonlinear equations, the fluctuations would influence the average motion. So even from the viewpoint of classical statistical mechanics, it is not unreasonable that the linear case may be deceptively simple.

LR: Up to now, you have told me about several hidden variable theories that work, although it does not seem to be easy to fulfill requirements that go further than just having the correct probability density. And, of course, we have learned that the trajectories sometimes have to behave in a strange way. It is nice to know that such theories are possible, but they seem to be a bit arbitrary. Above all, they do not give any clear physical picture of why the Schrödinger equation looks the way it does. Shouldn't one be able to do better? Is there *any* hidden variable theory which starts from some physical principles to *derive* the Schrödinger equation?

QM: No, I fear there is no such theory. It is true that having such a theory would be very gratifying, even if there would be no experimental procedure by which it could be distinguished from usual quantum mechanics. I believe that a theory like that would be adopted as the new "true" interpretation of quantum mechanics. However, it is hard to imagine that any theory that necessarily has to predict trajectories as strange as those we encountered for the two interfering wave packets could be based on a sufficiently simple physical picture. And, on the other hand, it is also hard to imagine that any theory *without* those continuous trajectories could appear physically "natural", given that there are no experiments that indicate particle jumps.

3 The measurement process

LR: You promised to tell me more about the "collapse of the wavefunction", that is, about measurements.

QM: OK. The important lesson to learn is this: The measurement process can be described completely *within* quantum mechanics, you do not need any extra postulate, no mysterious prescription for "collapsing the wave function"! However, you must pay a prize: You have to include the detector into the description. This means you have to deal with a many-particle wave function

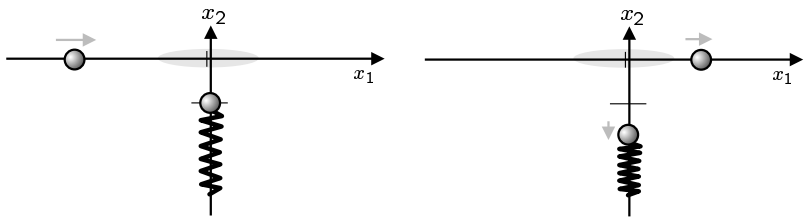

Figure 14: The "detector" in this simple example is a mass attached to a spring that feels a force when the particle moves through the interaction region (gray oval). Afterwards it will oscillate.

that depends on all the coordinates of the particles in the detector plus the coordinates of the particle you want to measure.

LR: Sounds pretty complicated.

QM: I will illustrate it in a very simple example, where the "detector" is just a mass attached to a spring, or, in any case, a particle that can oscillate around an equilibrium position. It will act as a detector for another particle because we say there is an appreciable force between the two particles whenever they come close to each other. Put simply, the "detector particle" will get a kick and start to oscillate when the other particle moves past.

LR: That sounds very classical...

QM: Of course, in the quantum-mechanical description, we have to solve the Schrödinger equation for the wave function of the two particles together. Let's keep things simple and assume that each of them can only move along one direction, so the wave function just depends on two coordinates - $\Psi(x_1, x_2)$, where x_1 belongs to the particle to be measured and x_2 is the coordinate of the mass attached to the detector spring, counted from its equilibrium position.
The detector is a harmonic oscillator, the particle itself is moving freely and there is an interaction potential $V(x_1, x_2)$ describing the mutual influence. We want this potential to describe some force that acts on particle 2, the detector, but only if particle 1 is near the origin, where the detector is located. So it should be of the form $V(x_1, x_2) = f(x_1)x_2$ with a function $f(x_1)$ that is nonzero only near the origin. You can imagine the detector spring to be placed below the track on which the particle is moving, so that the mass on the spring will feel a downward force if the particle moves nearby, provided f is positive.

LR: You claim that the "measurement process" somehow appears automatically once we solve the Schrödinger equation for this situation? However, we still need an initial condition for the time-evolution of the wave function!

QM: Yes. In the classical counterpart of this situation we would start with the detector being exactly in its equilibrium position and the particle moving

towards the detector. In quantum mechanics, we should demand the detector to be in its ground state, since that is the state it will relax into, if left alone. As for the particle itself, we will try an initial superposition of wave packets, one moving to the left and the other one to the right, towards the detector. Using this, we should be able to see how the initial superposition actually gets "collapsed" onto one state or the other by the measurement.
Thus, our initial wave function should be a product of two wave functions, one for the particle and one for the detector:

$$\Psi(x_1, x_2, t = 0) = \frac{1}{\sqrt{2}}(\psi_R(x_1) + \psi_L(x_1))\phi_0(x_2)$$

Here ψ_R and ψ_L are the wave functions of the packets moving to the left and to the right, which are added to form the superposition, and ϕ_0 is the harmonic oscillator ground state wave function, describing the detector. The prefactor is for normalization, so that the integral of the probability density $|\Psi(x1, x2)|^2$ over x_1 and x_2 is equal to one, if we assume the same for the integrals of $|\psi_R(x_1)|^2$ and $|\psi_L(x_1)|^2$ over x_1.
This form of Ψ means the particle and the detector are statistically independent at first, because the probability of finding the particle at x_1 and, at the same time, the detector at x_2, is $|\Psi(x_1, x_2)|^2$, which will be simply the product of the respective individual probabilities, calculated from the wave functions of x_1 and x_2 separately.

LR: Before things become too complicated, let's try to figure out what is the *classical* behaviour in this situation. As you have said, the particle moving past the detector will give it a kick, so it starts to oscillate. Of course, that means energy has been transferred to the spring, so the particle should be somewhat slower afterwards. However, maybe it is not really a sudden kick. If the function f is smooth, then the force acts for a while, becoming stronger gradually and then fading away. If this process is slow enough, we may imagine the spring to become compressed slowly and then to go back to its old equilibrium position, once the particle has passed. In that case, there would be no detection, since afterwards there are no oscillations, so you cannot recognize that the particle has been there if you look only at the detector. It seems that it depends crucially on the speed of the particle whether we get one or the other situation. Only if it is fast enough, we will have a "working detector". If the particle is *very* slow, it might even be reflected by the repulsive force!

QM: You are right. All of this remains true in the quantum-mechanical situation, apart from the quantitative details. In particular, if the particle is too slow, the detector will just be influenced only as long as the particle is moving past and there will be no trace left in the detector afterwards. That is essentially the reason why, as a rule of thumb, it becomes harder to detect particles or radiation when they do not carry much energy.
Let's look at the actual quantum-mechanical evolution of the wave function. Remember that we have to plot the probability density as a function of both x_1

and x_2 in the two-dimensional plane, so that the product of a wave packet in x_1 and a Gaussian oscillator wave function in x_2 gives some fuzzy blob. Since there are two wave packets making up the particle's wave function, $\psi_L + \psi_R$, there are also two of them for the total wave function $\Psi(x_1, x_2)$.
If the wave packet describing the particle is too slow, we will have essentially the same situation before and after the particle passes nearby the detector. The two-particle wave function $\Psi(x_1, x_2)$ will be distorted by the interaction only temporarily, and no trace will be left in the detector.

On the other hand, if the particle is fast enough, then the detector starts oscillating afterwards. In the picture of $|\Psi(x_1, x_2)|^2$, we see the fuzzy blob moving up and down in x_2-direction, while it still proceeds to the right in the x_1 direction, corresponding to the free motion of the particle. Of course, this applies only to the particle wave packet moving to the right, that has passed the detector. The other one remains undisturbed. Now the detector carries information about whether the particle has been there.
In this respect, there are not much qualitative differences to the classical behaviour. The most important one, perhaps, is that the energy transfer between particle and detector is quantized. So the particle will lose some random integer number of energy quanta $\hbar\omega$, if ω is the oscillator's frequency. This is the same physics that is behind the Franck-Hertz experiment, where fast-moving electrons could transfer part of their kinetic energy to gas atoms with which they collide, but only in discrete portions. All of this is contained in $|\Psi(x_1, x_2)|^2$ and can be observed especially at late times, when the particle will have moved a larger or a smaller distance depending on the loss of kinetic energy. But, of course, this is rather unimportant for the detection process. The ideal detector would have a very small frequency ω, so that the loss of kinetic energy does not really matter.

LR: I do not quite understand. Now we still have two wave packets. Where's the collapse?

QM: Yes, the two-particle wave function still describes a superposition of two possibilities: Either the particle is moving to the right *and* the detector has been excited, or it is moving to the left *and* the detector has not been excited. Before, either the particle was moving to the right or to the left, but the detector was in its ground state regardless of what the particle was doing. So the important point to notice is that the wave function $\Psi(x_1, x_2)$ is *qualitatively* different from before, since the state of the particle and that of the detector have become correlated: If you are to measure the detector's position, x_2, and you find a large elongation, this must mean that the particle has passed the detector. Large amplitudes of x_2 are a clear indication for "particle was (and is) moving to the right". This is what we would require from a detector.

LR: True. But still: Where is the "collapse" of the wave function? Or has this been a completely wrong concept all along?

QM: No, it is a useful concept. But now we can tell much more precisely than before what it really means: In the beginning, before the particle had a chance

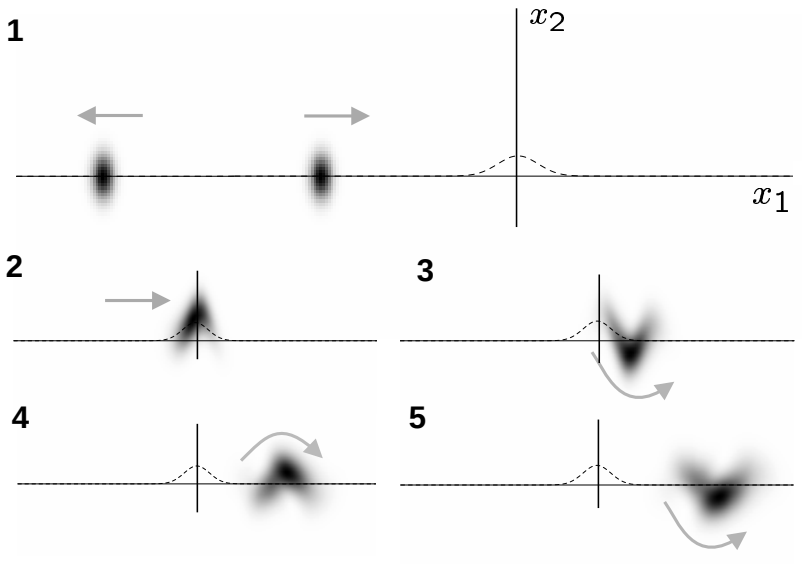

Figure 15: Plots of the probability density $|\Psi(x_1, x_2)|^2$ to find the particle and the detector mass at x_1 and x_2, respectively. At first (1) the two wave packets of the superposition start to move to the left and the right, while the oscillator (vertical component, x_2) is in its ground state: $|\Psi(x_1, x_2)|^2$ is a product. Only the right-moving packet enters the interaction region (indicated by the bump in the dashed line, that gives the minimum of the oscillator potential for x_2, as a function of x_1). The detector gets deflected (2) and oscillates afterwards (3,4,5), after the particle has already left the region. The left-moving packet (not shown in 2-5) propagates freely, i.e. the detector is not excited.

of reaching the detector, we could as well have used a single-particle wave function for our description of the particle, depending only on x_1 and describing the superposition of wave packets, $\psi_L(x_1) + \psi_R(x_1)$. This is because the particle and the detector are independent before they interact: The total wave function $\Psi(x_1, x_2)$ is simply given by the product of the individual wave functions, each evolving according to its own single-particle Schrödinger equation. We have used Ψ from the outset only because we anticipated we would need it later. Of course, *during* the interaction, we *must* use the two-particle wave function $\Psi(x_1, x_2)$, in order to describe the effect of the interaction potential $V(x_1, x_2)$. But even *afterwards*, we have to use Ψ, provided the detector is working. After all, even after the particle has left the region where the interaction takes place, there are correlations between particle and detector which means Ψ is no longer simply the product of two single-particle wave functions. That is to say, Ψ is now called "entangled". Therefore, the wave function of the particle, depending only on x_1, ceases to be useful. In fact, it no longer exists independently. The interaction with the detector forces us to consider the total wave function Ψ. That is why we speak of the "collapse" of the wave function, which means the collapse of the *single-particle* wave function.

LR: That sounds reasonable. But I expected only *one* out of the two wave packets to be picked at random by the measurement process. Then, afterwards we would only have to deal with that particular wave packet and could follow its further time-evolution without caring about the other packet any more.

QM: Indeed, that is essentially what has happened. However, it is the right packet *or* the left packet which has been detected, each with a probability of one half. That is the simple reason why, even after the measurement process, Ψ still must contain the two possibilities.

LR: That sounds very convincing in a situation involving classical statistics. But usually the wave function contains more than classical probabilities. Why can you argue that, after the measurement, we should regard it in such a classical way, while, before the measurement, we have been talking of superpositions?

QM: One way to distinguish between quantum-mechanical superpositions and mere classical probabilities is to ask about interference effects. We could imagine the two wave packets to be reflected, so that they overlap later on. Of course, they should not hit the detector a second time, so let us assume it has been removed in the meantime, or the packets follow another route that does not pass nearby the detector. We imagine to measure the position of the particle at the time when the two wave packets should overlap again. Then, if no detection had taken place previously, we would observe the usual interference fringes that we have discussed before. That means: If we can observe interference, then certainly we are not allowed to talk simply of classical probabilities to have either one or the other possibility. However, if the detector has been working, we do not observe any interference pattern at all! Now it is really the probability densities rather than the wave functions which are superimposed,

in a completely classical manner.
That is already an indication that the "excuse" for still having two possibilities contained in Ψ is correct. The interpretation is that we are averaging over many runs of the same experiment, including those where the particle has been detected to the right and those where it has not been detected, which means it must have been moving to the left. In each of these runs, we could imagine that actually only one packet is left after the detection event, so there cannot be any further interference in the zone where usually we would expect the two wave packets to overlap. However, because of the averaging over many runs, we still don't always observe the particle traveling only along one direction. At least, this is a consistent interpretation of the experimental results in this "gedanken experiment".

LR: Wait a moment. You have talked about checking for interference by another *measurement*. Certainly, this should be described by including also that additional detector into the description, or rather even a whole detector array in the region where the interference is expected?

QM: Of course. For that purpose, we would consider a total wave function $\Psi(x_1, x_2, x_3, \ldots)$ that is a function of all those detector coordinates.

LR: Another question comes to my mind: You said that the state of the particle and that of the detector become correlated, such that, by "looking" at the detector, one could learn something about whether the particle is moving to the right or to the left. Certainly, this read-out of the detector state also involves another measurement?

QM: Yes.

LR: So when, actually, are you allowed finally to take $|\Psi|^2$ to calculate the probabilities? It seems that this goes on forever, without any definite result. In the end, you will have to deal with the computer that registers the detector signals, and with the researcher that looks at the screen of the computer...

QM: It is true that, in principle, you cannot stop at any point of this "hierarchy" of measurements without introducing some extra postulate like "collapse of the wave function". But that does not mean that this whole line of description is useless. After you have described even the first measurement stage, you may already stop. It is enough for all practical purposes, in the following sense: If you apply the standard rule about "collapsing" the wave function of the particle to one or the other state, and assume the respective probabilities to be given by the magnitudes squared of the amplitudes, then this gives you a consistent description. As I have explained just before, it describes correctly the vanishing of the interference effect, and it also gives the relative weights with which you have to superimpose the probability densities of the two wave packets to obtain the probability of detecting the particle here or there after the measurement. Of course, you can go further and include also another detector that measures the state of the first detector or the one which carries out a further measurement on the particle after the initial detection event. The important point is that this will change nothing in the statistical predictions for

the behaviour of the particle. You add more realism and more complications, but the results are consistent with that simple rule. On the other hand, it is a good way to verify that the rule is correct.

LR: So the message is that, although there is such a hierarchy, the first measurement process is something special, since it really transforms a quantum-mechanical superposition into two classical possibilities - while the other measurements do not change the picture any more?

QM: That is exactly the point. In this sense, the quantum measurement process is described completely within quantum mechanics, without introducing any postulate about "collapse of the wave function". Rather, this postulate arises automatically from the description, as a convenient rule that summarizes the effects of the measurement and tells you that you do not need to incorporate all the further measurements into your description, although you could do it and it would not hurt, so to speak.

LR: If that is true, not only the total probability of finding the particle somewhere after the measurement should come out correctly using the classical superposition of the two possibilities. It is also the conditional probabilities that should be right. For example, if we ask for the probability of a second detector being excited by the particle, given that the first detector has been excited, then the classical expectation is that this should be zero if the second detector is placed somewhere to the left of the first one, so it cannot be reached by a particle moving to the right. And it should be one if the second detector is placed to the right of the first one and is working well, so it doesn't let the particle go unnoticed.

QM: That is exactly what we find: If you set up such a situation with two detectors, and you look at the time-evolution of the total wave function $\Psi(x_1, x_2, x_3)$, in the end the probability of having large elongations of x_2 and x_3 simultaneously vanishes, if you consider the first of the two possibilities you mentioned. Thus, quantum-mechanics tells you that the simple classical picture of conditional probabilities for different sequences of detection events is reliable, after the first measurement has turned the quantum superposition into a classical one.

LR: I see, this approach is self-consistent, at least in this example. But how can we be sure that there cannot be any interference effects after the detection has taken place? After all, the total wave function still contains both possibilities, and it is not so clear to me that they are always to be considered mere *classical* possibilities.

QM: You are asking about cases where one would make wrong predictions by using merely classical probabilities after the measurement. Is it really enough to start from the probability densities for the left- and right-moving wave packets separately and add them with weights given by the classical probabilities of having measured one or the other possibility?
Fortunately, there is a straightforward proof that demonstrates why this works: After the detection, the total wave function will be something like

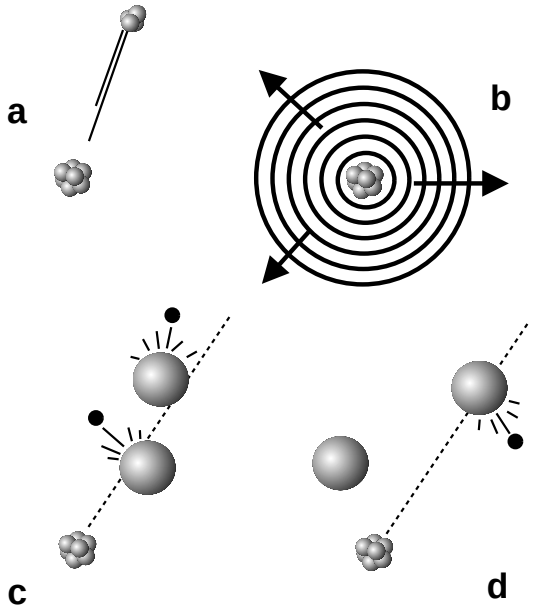

Figure 16: Mott's gedanken experiment concerning conditional detection: How is the classical particle picture of an α-particle emitted from a nucleus into a particular direction (a) compatible with the quantum-mechanical wave spreading out into all directions (b)? If two atoms are placed roughly in a line, as seen from the nucleus, then the fully quantum-mechanical calculation (treating the α and the electrons inside the atoms) predicts that in the end both atoms may be found in an ionized state (c). However, if they are placed in different directions, the many-particle wave function will vanish for such doubly-ionized configurations: At most one of the atoms can be ionized (d). This works only because the α-particle's kinetic energy is so large that it is not deflected much when kicking an electron out of the atom (the momentum uncertainty acquired by the α is related to the resolution of the position measurement, given roughly by the extent of the atom - which is obviously not drawn to scale with respect to the nucleus.). Qualitatively, the correct prediction is also obtained by properly invoking collapse of the wave function at the first ionized atom.

$$\Psi(x_1, x_2) = \frac{1}{\sqrt{2}}\{\psi_L(x_1)\phi_0(x_2) + \psi_R(x_1)\phi(x_2)\}$$

This is the entangled state, where the correlation between detector and particle is apparent: ϕ_0 is still the detector's ground state, because the left-moving particle has never hit the detector. On the other hand, ϕ is an excited state of the detector. Now, if the detector has been working properly, the overlap between the quantum-mechanical states ϕ and ϕ_0 is *zero*:

$$\langle\phi|\phi_0\rangle = \int \phi^*(x)\phi_0(x)\,dx = 0$$

Actually, that can be considered the *definition* of a perfect measurement. If the particle had been relatively slow, the detector would have been excited only with a small probability and the overlap would not be zero. In that case, by looking at the state of the detector, you could not determine unambiguously whether the particle has been there or not. But we will assume that the particle is fast enough.

LR: Are you sure that, in the case of excitation, the wave function is simply a product $\psi_R(x_1)\phi(x_1)$? Didn't you tell me that, with certain probabilities, the particle would transfer a larger or smaller number of energy quanta $\hbar\omega$ to the detector? Then, I would expect each of the different possible excited states of the harmonic oscillator to appear, with corresponding different states for the particle's wave function, since it will have different kinetic energies depending on the energy transfer.

QM: Actually, you are right. In principle, we would have a more complicated state, something like $a_1\psi_{R1}(x_1)\phi_1(x_1) + a_2\psi_{R2}(x_1)\phi_2(x_2) + \ldots$. Each of these contributions would correspond to $1, 2, \ldots$ transferred quanta and the probabilities would be $|a_1|^2$, $|a_2|^2, \ldots$. But the point here is to discuss the ideal case first, where the detector only distinguishes between left- and right-moving particles. We can set up a nearly ideal detector by having a very small frequency, as discussed above. Then we may arrange parameters such that the particle loses almost no kinetic energy although the harmonic oscillator will still go into a superposition of excited states, which we call ϕ. In that case, for our purposes, it is a good approximation to deal only with $\psi_R(x_1)\phi(x_2)$.
Then, if we want to calculate the probability of finding the particle itself around a point x_1, we get that by integrating the total probability density over the unimportant detector coordinate x_2:

$$p(x_1) = \int |\Psi(x_1, x_2)|^2\,dx_2$$

Now it becomes important that ϕ and ϕ_0 are orthogonal, because it means that p just turns out to be the classical superposition of probability densities of the

two wave packets. We just have to insert the expression for Ψ from above and carry out the integration over x_2:

$$p(x_1) = \frac{1}{2}\int |\psi_L(x_1)|^2 \, |\phi_0(x_2)|^2 + |\psi_R(x_1)|^2 \, |\phi(x_2)|^2 + \\ \psi_L^*(x_1)\psi_R(x_1)\phi_0^*(x_2)\phi(x_2) + c.c.\, dx_2$$

In the second line, the integration over x_2 gives the overlap integral from above, which is zero. In the first line, we may use the normalization of the detector wave functions, which gives 1 for the x_2 integral:

$$p(x_1) = \frac{1}{2}\,|\psi_L(x_1)|^2 + \frac{1}{2}\,|\psi_R(x_1)|^2$$

That is just the classical superposition, as announced. The prefactors of $1/2$ are the classical probabilities. Of course, they would be equal to $|a_L|^2$ and $|a_R|^2$ if we had set up the initial particle wave function to be $a_L\psi_L(x_1) + a_R\psi_R(x_1)$.

LR: I see how it works out. But this is only right immediately after the detection event. Couldn't it happen that the interference reappears later on? After all, if the detector is realized as a harmonic oscillator, like in our example, it swings back and forth and also comes back to the origin at certain times. Maybe then the interference term in $p(x_1)$ does not drop out any more?

QM: No, it vanishes at all times, in spite of the dynamics of the excited state in the detector's harmonic oscillator potential well. This is because the overlap still remains zero. And it must remain zero because both ϕ_0 and ϕ evolve according to the same Schrödinger equation, that is according to the same Hamiltonian operator. It is a general rule that the overlap of two states, when they evolve according to the same Hamiltonian, stays the same for all times. Of course it is true that ϕ describes an oscillator wave packet that swings back to the origin, so you might think it could have a nonzero overlap with the ground state ϕ_0 that just sits there, at the origin. However, the overlap still remains zero, because ϕ has acquired a finite average velocity, while ϕ_0 is at rest: In momentum space, they are displaced with respect to each other. Therefore, you could still distinguish perfectly between these two states - not by a position measurement of the detector, but by a momentum measurement. The interference cannot reappear.

LR: Is there really no way of making the interference between the two wave packets reappear? What about this: If we apply a carefully tuned force pulse to the oscillator, we may stop its oscillation, such that it goes into the ground state ϕ_0 again. Then, the overlap of the two detector states would be reset to unity and the interference pattern in the particle's probability density would reappear!

QM: I do not doubt that, with a proper force pulse, you could turn the state ϕ into the ground state ϕ_0. However, you cannot be sure that the detector has actually been excited into ϕ, without measuring its state. Maybe the particle has been traveling to the left, such that the detector already *is* in the ground state ϕ_0, before you apply the force. In that case, you will excite it to another state! That means your whole procedure is fruitless: Afterwards, the overlap between the two detector states, belonging to left- and right-moving wave packets will still be zero, only the states have changed: the excited state has been turned into the ground state and vice versa.

LR: Well, of course I assumed that, first of all, I look at the detector. Only when I find it to be in its excited state, I will apply the force in order to reset it to the ground state.

QM: This will not help. By "looking" at the detector, you have measured it, that is you have to include another measuring device in the description. Let us assume that, whenever this device finds the detector in the excited state, it will automatically trigger a force that resets the detector, so you do not have to bother about it and we don't need to include *you* in the description. Such a setup is certainly possible. However, afterwards, you will have an entangled state of the following form:

$$\Psi(x_1, x_2, x_3) = \frac{1}{\sqrt{2}}\{\psi_L(x_1)\phi_0(x_2)\chi_0(x_3) + \psi_R(x_1)\phi_0(x_2)\chi(x_3)\}$$

Here χ_0 is the ground state of the second measuring device, and χ an excited state. You see: Entanglement has only been transferred to the second detector. The interference terms will not reappear, simply because now it is the overlap between χ_0 and χ which is zero.
To put it in a nutshell: After you have read out the detector, you have no chance of "undoing" the measurement.

LR: I see. Obviously the force pulse fails because it changes the state of the detector in *any* case, even if the detector had been in its ground state. And, as you have said, I cannot make sure that the detector is in its excited state because that would mean another measurement. But suppose there is some dissipation, such that the detector relaxes from its excited state into the ground state after some time, without any intervention. On the other hand, if it already is in the ground state, it will simply stay there. I know that this works at least for excited states of an atom: it will decay to the ground state by spontaneous emission of a photon. In that way, the detector will definitely go into its ground state, without any further measurement, and without any carefully tailored force pulse.

QM: Yes, the detector will relax to its ground state. But, no, the interference will *not* reappear. That is because, by emitting a photon, the entanglement has been transferred again, now into the electromagnetic field. It is just the case we discussed before, only that the second measuring device has been replaced

by the field. Although *you* do not know which state the detector had been in, you could find out *in principle* by detecting the photon afterwards. This means, from the point of view of quantum mechanics, it is the same story again.

LR: So there is no chance of "undoing" the measurement?

QM: There is, but only if the particle returns to the detector, before the state of the detector has been read out. Only then the overlap between the two detector states can be reset to unity, by employing the interaction between particle and detector. This is possible because then the two detector states do *not* evolve under the action of the same Hamiltonian. One of them is influenced by the interaction with the particle, the other one is not. Imagine a setup where the right-moving particle is reflected, goes back to the detector and reaches it just at the right point in time in order to stop its oscillation when it is about to swing through its equilibrium position. This could be done, with a proper tuning of the path lengths and so on. On the other hand, if the particle had been moving to the left originally, it will *not* be reflected and *not* disturb the detector. In either case, the detector will end up in its ground state. Afterwards, you will have a total wave function that is no longer entangled, and, consequently, any subsequent interference experiment will show interference fringes again. The modern buzz-word for this is "quantum eraser": The measurement record has been erased again. You have to make sure, that, in the end, the particle leaves no trace anywhere, even though it had become entangled with the detector temporarily.

LR: Well, it seems a little bit difficult to realize in this case, when you need just the proper timing and so on. Are there situations where one can observe in detail how this trick is performed?

QM: Of course, it gets easier if both the system and the detector are as simple as possible. A good example is a spin measurement in a Stern-Gerlach device. I assume you know how that works?

LR: Yes: You measure the spin by sending the particle through a magnetic field whose strength varies with position. The spin has a magnetic moment attached to it, which gives a different potential energy depending on the direction of the spin with respect to the field and depending on the field strength. Therefore, this energy will also vary with position. That means the particle experiences a force that depends on its spin direction. If the spin is that of a single unpaired electron inside an atom, it can assume only two quantum-mechanical states: along the direction of the magnetic field or opposite to it. So the particle will get deflected upwards or downwards, and you can find out the spin direction by detecting the particle on a screen.

QM: Right. In that case, the "system" to be measured is the spin and the "detector" is the particle's coordinate, because it gets deflected according to the state of the spin. Only later on, that kind of "detector" is "read out" when the particle hits the screen. Let's assume we start out with an arbitrary superposition of the two spin directions $\sigma = \pm\frac{1}{2}$:

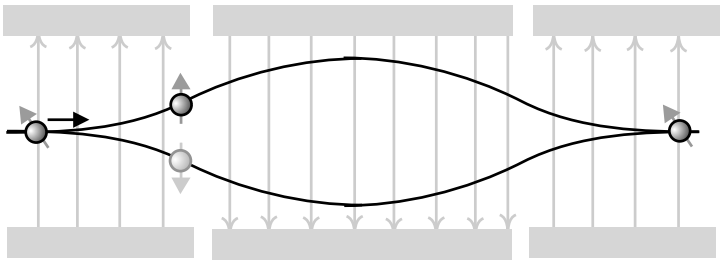

Figure 17: Recombining the beams in a Stern-Gerlach apparatus before they impinge on the screen "undoes" the measurement.

$$\Psi(\sigma) = a\chi_\uparrow(\sigma) + b\chi_\downarrow(\sigma)$$

Here I have written the spin-dependent wave function similar to a position-dependent wave function. $\chi_\uparrow$ is the eigenstate for "spin up", it gives 1 if $\sigma = +\frac{1}{2}$ and 0 otherwise. Similarly, $\chi_\downarrow$ belongs to "spin down". The "quantization axis" that defines the meaning of "up" and "down" is given by the magnetic field. Through the deflection by the magnetic field gradient, the position of the particle will start to depend on the spin. That is, just as before, we are now forced to consider the total wave function, depending on both spin and coordinate. It will evolve into an entangled state:

$$\Psi(\sigma, x) = a\chi_\uparrow(\sigma)\psi_\uparrow(x) + b\chi_\downarrow(\sigma)\psi_\downarrow(x)$$

As you can guess, $\psi_\uparrow(x)$ describes a wave packet that is deflected upwards, for example, and $\psi_\downarrow(x)$ is deflected downwards. Of course, the actual deflection will depend on the direction of the magnetic field gradient.

LR: I see, it is the same story again. But perhaps it is a little bit more striking than before. If we start out with a spin pointing along the y direction, then we will always end up either with "spin up in z direction" or "spin down in z direction", if the magnetic field is pointing along the z direction. Both will occur with equal probability, but it is *not* like in the corresponding classical situation, where you would expect a spin not to be deflected at all if it points in a direction orthogonal to the magnetic field. Here, only *on the average* there is no deflection. And in order to describe that behaviour properly, we have to think of the spin pointing along y as a superposition of those two possibilities, which amounts to a change of basis states. It becomes more obvious than in the previous example that the measurement device chooses a preferred basis of quantum-mechanical states.

QM: Yes. Now back to your initial question of "undoing" the measurement. Here it can be done relatively easily: You simply have to consider an extended apparatus, where a field gradient of opposite direction applies a force of equal strength but different sign to the particle. Actually, you have to apply it twice as long, such that the wave packet that initially moved upwards is now moving downwards, and vice versa for the other packet. Afterwards you still have to stop the particle by applying a field gradient along the initial direction. In this way, the two wave packets, corresponding to the two different spin directions, will be recombined. Then, the total wave function is "dis-entangled" again, and the measurement has been undone. Now, if you had a spin pointing along y initially, you will also end up again with a spin pointing along y. On the other hand, assume you had chosen to measure the spin state of the particle in the "upper" wave packet before they had been recombined, using another measurement device. If that device had been set up to measure the spin in the y-direction, you would have found "up" or "down" with equal probability, because these are the results for a spin pointing along z direction. In this intermediate situation, the correct description has been: The spin is either "up" or "down" along z, with *classical* probabilities given by $|a|^2$ or $|b|^2$. However, the fact that one is able to undo the measurement shows that the information about the initial spin state still has been around in the total wave function. But, just as before, in order to "erase" the measurement record and restore the initial state, we needed to employ again the interaction between "system" and "detector". Here, it was the additional field gradient which was able to reset the entangled wave function to the original state. If the particle had not been allowed to interact with any magnetic field afterwards, we could not have achieved our goal, because then every deflection would have been spin-independent, similar to the force pulse that did not help to undo the measurement either.

LR: Doesn't that mean that it is somewhat dangerous in general to rely only on the simple picture of classical probabilities after the interaction with the detector?

QM: Well, in practice it is much harder to "undo" the measurement, if the measurement device is really "macroscopic", i.e. if the detection excites a lot of degrees of freedom. We can observe this very nicely in the previous example where the detector has been represented by an oscillator. Imagine that this oscillator, in turn, is coupled to many other particles that form a whole crystal. Then, the energy transferred to the "detector" will spread out through the crystal in the form of vibrations, i.e. sound waves. Now, if you want to undo the measurement, you would have to time the return of the particle in such a way that, when it gives the detector a second kick, all the sound waves have just reconverged and the whole vibrational energy is contained, once again, only in the detector particle. Even classically, you can imagine that this is an extremely improbable event, especially if the crystal is large and especially if it is not shaped extremely regularly. In this way, the first detection event has become irreversible, for all practical purposes. That's why you usually do not have to bother about the possibility of the interference terms showing up later

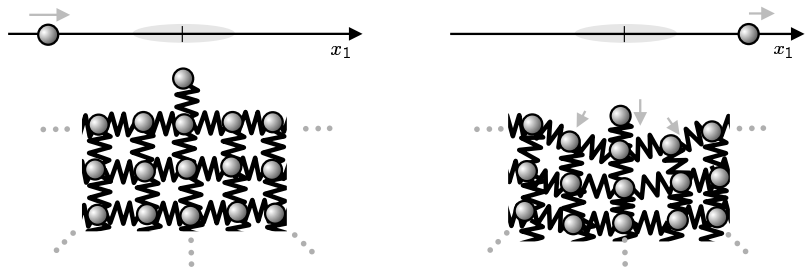

Figure 18: If the "detector" is actually part of a larger crystal, then sound waves will spread through the crystal and it becomes practically impossible to "undo" the measurement later on. This is a very simple example of "irreversibility" introduced by a macroscopic measurement device.

on, because, even if you let the particle interact with the detector again, it is very unlikely that this will "erase the measurement record".

LR: I see: for all practical purposes we may treat the measurement as being irrevocable, after the first detection has taken place, apart from examples like the Stern-Gerlach apparatus, where not many degrees of freedom are involved in the first stage of the measurement. And subsequent measurements give results that can be readily understood within the context of classical probability theory, nothing mysterious.

QM: Yes, but remember that the quantum-mechanical measurement is still very different from merely detecting a statistically distributed pre-defined property: If you want to measure different kinds of observables, for example momentum as opposed to position, you have to change the whole setup and cannot measure them both at the same time, because that is not allowed by Heisenberg's uncertainty relation.

LR: You told me about Bohm's theory. There, we *did* have a "pre-defined property", some "hidden variable", namely the position of the particle. Of course, you have pointed out that, in order to measure the momentum instead, we would have to consider a different setup, where, ultimately, different momenta are mapped onto different positions. But I wonder how the measurement process is described there - or whether we need any special description at all when the position is assumed to be real all the time anyway!

QM: The important idea behind Bohm's theory is to stick very closely to the Schrödinger equation, so that all the statistical results will remain completely compatible with the "usual" quantum mechanics. This means that the measurement process is also described just as we have learned it, with many-particle wave functions depending on the coordinates of particle and detectors together. The only difference is the following: now we *imagine* that, at

every instant of time, the particle itself and also all the particles in the detectors are at certain positions and they move along continuous trajectories as time progresses. The trajectories are deterministic and can be calculated once the time-evolution of the complete wave function is known. The statistical element enters only because the initial positions of all particles are distributed randomly. In our example of the measurement process, sometimes you will happen to start out with a particle moving to the left, sometimes with a particle moving to the right. In the first case, the detector does not get excited: the detector mass will always remain at rest, since that is what a particle does in Bohm's theory when it is in the ground state. In the other case, when the particle is actually moving to the right, the detector mass will start to oscillate after the interaction between particle and detector has taken place.

LR: I remember that, in Bohm's theory, the velocity of the particle is given by the gradient of the phase of the wave function, evaluated at the current position of the particle. But now the wave function depends on many coordinates, so what should I insert for all the other coordinates? Should I average over them?

QM: No, you simply have to insert the coordinates of all the other particles, with the values they happen to have in the current configuration of the complete system. So the velocity of particle number 1 will be given by

$$\vec{v}_1 = \frac{\hbar\vec{\nabla}_1\varphi(\vec{r}_1, \vec{r}_2, \vec{r}_3, \ldots, t)}{m}\,.$$

The phase φ is the phase of the many-particle wave function: $\Psi(\vec{r}_1, \vec{r}_2, \ldots) = |\Psi(\vec{r}_1, \vec{r}_2, \ldots)| \exp[i\varphi(\vec{r}_1, \vec{r}_2, \ldots)]$. The gradient of the phase is only taken with respect to $\vec{r}_1$.

LR: This is remarkable. The velocity of particle 1 depends on the current positions of all the other particles as well.

QM: That's true. Of course, there are important exceptions: If the total wave function is a product of wave functions belonging to the different particles, the phase becomes a sum of phases depending on the individual particle coordinates separately, $\varphi(\vec{r}_1, \vec{r}_2, \ldots) = \varphi_1(\vec{r}_1) + \varphi_2(\vec{r}_2) + \ldots$. Then the velocity $\vec{v}_1$ will only depend on $\vec{r}_1$, as in the single-particle case. This is reasonable: If the particles have never interacted, then their trajectories are independent.

LR: And obviously we shouldn't be surprised that their trajectories start to depend on each other once the particles are interacting.

QM: Yes. However, even *after* they have interacted, this dependence does not cease. The wave function will in general have become entangled: it is not a product of individual particle wave functions any longer and the phase is not a sum of individual phases. Then the particles "influence" each other's motion in Bohm's theory although we would not expect any such influence on classical grounds. It is not just the same as some statistical correlation appearing in the positions of the particles after they have left the zone of interaction. Of course,

something like that is to be expected and it will also lead to correlations in the further time evolution of the trajectories. But it is more than that: If one of the particles experiences some force, the whole many-particle wave function will evolve in a different way and that will also change the trajectories of the other particles which may be far away.

LR: Doesn't that contradict the special theory of relativity? Couldn't we use it to send messages over large distances, faster than the speed of light?

QM: No, we cannot. We could only do that if we were able to actually observe those trajectories. But we can't, without disturbing them. The first measurement of the particle's position will disturb all future measurements. And we know that Bohm's theory will yield just the same probability distribution for this measurement as quantum mechanics predicts. What is most important: This probability distribution is *independent* of what you do to the other particle far away, once the interaction has ceased. That can be proven in a straightforward way, using essentially the same arguments as those which tell us that a measurement cannot be undone unless particle and detector interact again. So the statistical feature of quantum mechanics prevents us from sending faster-than-light messages by this nonlocal influence. Of course, within the standard interpretation of quantum mechanics, this suggestion would never come up anyway, since there we do not have any of those "hidden variables" that could experience such influences.

LR: When we discussed the behaviour of trajectories in general, we came across a peculiar effect: In a situation with two wave packets crossing each other, all the trajectories had to be reflected in the zone of overlap, if they were assumed to be continuous and no faster-than-light motion was allowed. This was also the case for Bohm's theory, in particular. Now you have told me that a measurement will destroy interference effects. If we place a detector in one of the arms of the interferometer, sometimes we will detect a particle there and sometimes the particle will go the other way. No interference effects will show up any more in the zone where the wave packets overlap - even if the particle is not absorbed by the detector but instead allowed to move on after it has been detected. My question is the following: Are the trajectories still reflected or do they cross each other? I cannot rule out either of these possibilities, since they could lead to the same probability density, which is the only quantity that has to be reproduced correctly by Bohm's theory or any other hidden variable theory of this kind.

QM: You are right, in principle both alternatives are possible. However, when we actually describe the complete situation, including the detector, using Bohm's theory, we find the following: The trajectories are not reflected anymore! Remember that previously, at least for the one-dimensional situation, the trajectories simply had to be reflected, because they are not allowed to cross each other, according to Bohm's theory. But now we are dealing effectively with a higher-dimensional situation, because the trajectories also have to describe the time-evolution of the detector's coordinate. When the particle moves nearby

the detector and is detected, the detector's coordinate will change. More precisely, the wave packet describing both detector and particle will get displaced in configuration space, along the dimension belonging to the detecor. The "system point" of Bohm's theory will follow suit, if it had been located within that wave packet previously, i.e. if it had described a situation with the particle actually moving along the arm of the interferometer where the detector has been placed. Afterwards, when the wave packets would usually overlap, they just miss each other in the higher-dimensional space we are dealing with. Therefore, no interference pattern will show up in the probability density of finding the particle here or there. And, at the same time, the trajectories living in that space will neither cross each other nor be reflected, they will simply pass by each other. If you are interested in the trajectory of the particle alone, you have to project the higher-dimensional trajectory onto the dimension of the particle's coordinate. Then you may, and will, get a crossing of those trajectories in the zone were they would be reflected in the setup without a detector.

4 Local realistic theories

LR: I have read about a hidden variable theory which seems to be much more natural and physical than Bohm's theory. It is called "Stochastic Electrodynamics".

QM: Tell me more about that theory - does it involve continuous trajectories as well?

LR: Yes, it even starts from the classical equations of motion! And it avoids any "nonlocal influences".

QM: Well, I doubt that is possible... But go on.

LR: Perhaps it is most natural to introduce this theory in relation to the stability of atoms. You know that the classical model of an atom, with an electron tracing its orbit around the nucleus like a planet revolving around the sun, is unstable. The electron represents an accelerated charge that will radiate away its energy and will spiral into the nucleus, after about a million revolutions or so. Therefore, one might imagine that in a universe dominated by classical laws, all the atoms would collapse within a mere nanosecond. But now try to go one step further in this story: The energy radiated away by the atoms is not lost. Instead, empty space is now filled with electromagnetic energy in the form of fluctuations of the electric and magnetic field. This means the electron in any given atom will experience some fluctuating force that tends to increase its energy. In other words, the possibility arises of a balance between energy input due to fluctuations and energy dissipation due to the radiation of the accelerated electron. In this way, the atoms could become stable.

QM: So you want to tell me that we settle into a thermodynamic equilibrium belonging to some finite temperature?

LR: No. The idea would be that these fluctuations in the electromagnetic field are not just classical heat radiation. In fact, for this effect to be fundamental, we would require the fluctuations to be Lorentz-invariant. Heat radiation is not Lorentz-invariant, because it prefers a certain frame of reference, the frame where the hot material objects, that have equilibrated with the radiation field, are at rest. We should ask: Is there any classical fluctuating electromagnetic field which is statistically Lorentz-invariant? For example, the average energy contained in a given volume and a given frequency interval is called the power spectrum, when you consider it as a function of frequency. Can we set up fluctuations such that their power spectrum will be the same in every frame of reference? The answer is yes. And what is more, it is found that this requirement of Lorentz-invariance fixes the power-spectrum of the fluctuations uniquely, up to a constant factor that determines the overall strength of the fluctuations. The power spectrum *has* to rise like the third power of the frequency, ω^3. Only then the transition to another frame of reference will reproduce the same spectrum, due to a delicate cancellation of Doppler shifts and Lorentz contraction effects.

QM: So you are talking about the zero-point fluctuations of the electromagnetic field in vacuum?

LR: Yes, the ω^3 spectrum is exactly the power spectrum which one obtains from the quantum theory of the electromagnetic field. But there, it is obtained in quite a different fashion: The field is formally equivalent to a collection of harmonic oscillators, one for each mode of the field. At zero temperature, each of those oscillators still has some finite average energy, the so-called zero-point energy $\hbar\omega/2$, which is proportional to both Planck's constant and the frequency. Now in order to obtain the power spectrum, one still has to multiply by the number of modes contained in a given frequency interval, which is proportional to ω^2. Then one arrives at a spectrum proportional to ω^3, with a prefactor involving Planck's constant. When we compare with the derivation based upon Lorentz-invariance, we learn two important facts: First of all, the fact that in quantum mechanics a harmonic oscillator has an energy proportional to its frequency is absolutely necessary to obtain a Lorentz-invariant power spectrum for the zero-point fluctuations. This is amazing, since the quantum-mechanical derivation of this harmonic oscillator ground state energy seems to be completely independent of special relativity. In a way, one could claim that the requirement of Lorentz-invariance is an alternative derivation of $E \propto \omega$ for harmonic oscillators in quantum theory! And secondly, we see that the overall strength of the fluctuations is proportional to Planck's constant $\hbar$. That is also very reasonable, since in a completely classical world, $\hbar$ would be zero and therefore the zero-point fluctuations would vanish.

QM: So how does $\hbar$ appear in the theory of "Stochastic Electrodynamics"?

LR: It does not have to be introduced at all. On the contrary: We start with fluctuations whose spectrum goes like ω^3, and then the prefactor describing the overall strength of these fluctuations will serve to *define* Planck's constant.

That is the principal idea of Stochastic Electrodynamics: Nothing is changed about classical mechanics and electromagnetism, apart from the initial conditions for the electromagnetic field in vacuum. Instead of having both electric and magnetic fields equal to zero in free space, the world is assumed to be actually described by a solution of Maxwell's equations that consists of a random superposition of plane waves, with a power spectrum rising like ω^3. This then will induce stability of atoms and it should lead to all the other effects known from quantum mechanics.

QM: I agree it sounds interesting. But does it work?

LR: Well, of course it is very hard to actually solve the equations of motion of an electron subject to these fluctuations and to the Coulomb potential of the atom's nucleus. It is so difficult because the force fluctuations are correlated over large time spans, they are not just simple "white noise" as it is encountered in the theory of Brownian motion. But in the case of the harmonic oscillator, it can be done relatively easily, and it works indeed! One has to take into account both the fluctuations and the radiation damping force. Then one obtains an equilibrium distribution for the particle's position which is of Gaussian shape and *exactly* equal to the quantum-mechanical expectation. Provided, of course, that one inserts the correct power-spectrum of the zero-point fluctuations.
Qualitatively, you can also describe many other things in quantum-mechanics using such fluctuations. For example, tunneling through a barrier, which is classically impossible, could be made possible by some temporary energy input due to the fluctuations. I have also read that people have been able to derive the correct Planck spectrum of heat radiation within Stochastic Electrodynamics.

QM: That is interesting. However, I am a little bit skeptical of that example of the harmonic oscillator. It is well known that systems with *linear* equations of motion, that is collections of oscillators, show a very close correspondence between classical and quantum-mechanical results - I have told you about that earlier. But in the case of nonlinear systems, like the atom, you will probably run into trouble: If I am not mistaken, the fluctuations will be able to ionize the atom, and then the electron will not return. This is in contrast to quantum mechanics, where an atom in its ground state is stable against ionization. But of course, as you have said, this situation is extremely difficult to analyze in detail.
Therefore, let us look at a simpler example: A potential well with infinite walls. It is hard to imagine that, in this case, you will be able to obtain the correct ground state probability density of quantum mechanics. This is because the fluctuations are not very strong: You yourself mentioned that, in the case of the hydrogen atom, radiation damping would considerably affect the orbit roughly after one million revolutions. Since this effect is to be compensated by the fluctuations, we know that they will not severely alter the dynamics of a system during one period of revolution, or one period of oscillation in the

case of the potential well. This fact is also confirmed explicitly by the calculation in the case of the harmonic oscillator. Now we know that the quantum-mechanical ground state distribution for the potential well with infinite walls is concentrated in the middle of the well. However, this distribution is not stable under the time-evolution according to the classical equations of motion, regardless of which momentum distribution you prescribe. And since the fluctuations are so relatively weak, I do not think they could stabilize this distribution.
But even worse, if we go away from these bound systems and just consider motion in free space and interference effects of overlapping wave packets: How do you think those phenomena could be described by a theory that essentially only introduces fluctuations?

LR: Well, frankly speaking I do not know. But one could speculate that, perhaps, the interference pattern is actually first in the electromagnetic field, which then acts as a sort of "guiding wave" or "pilot wave" for the electron. This would be similar to Bohm's theory, only the "guiding wave" would have a direct physical meaning. Or one may speculate that the high-frequency fluctuations of the electromagnetic field, around the Compton frequency, become important. Perhaps this would not only lead to some jitter motion, as we discussed it earlier, but it would also have a relatively large effect on the motion of the electron, because the power spectrum grows towards high frequencies. Then we could hope to get the correct ground state distribution even for such examples like the potential well. Even better, I have read that one can actually get the de-Broglie relation $\lambda = h/mv$ out of these high-frequency modes, in a way. Consider standing waves formed out of superpositions of electromagnetic plane waves at the Compton frequency: the interference pattern of these plane waves will develop a beating period if one changes to a different frame of reference via a Lorentz transformation. And this spatial period will be just h/mv, if one moves at velocity v with respect to the original frame of reference! Therefore, somehow this may then govern the motion of the electrons, which would open the possibility of *deriving* Schrödinger's equation from physical principles!

QM: That gets too speculative for my taste. Of course you can always keep me busy by introducing yet another hard-to-analyze feature into the theory.
But I think there is a simple gedanken experiment that demonstrates why every theory that uses just the vacuum electromagnetic field as a sort of "guiding wave" for the electron must fail. The point is not to analyze something as complicated as a bound state or motion through a double slit, where one cannot really carry through the analysis. Instead, let's imagine an electron that has been prepared in a wave packet which is spread over a rather large region of space, such that its momentum uncertainty is relatively small. This we could check experimentally by repeatedly preparing an electron in the same wave function and then measuring its position at some variable instant of time. In this way, we could observe the probability density of the wave packet and we would find that it expands only rather slowly, because of the small momen-

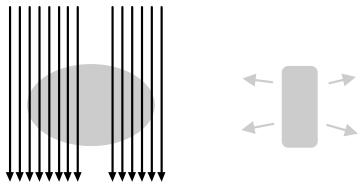

Figure 19: "Cut-away" gedanken experiment demonstrating failure of certain theories using the vacuum electromagnetic fluctuations as a "guiding field" for the electron. Left: A pulse of gamma rays will hit the electron, unless it happens to be in the central region of the wave packet. Right: If it has not been hit, it is described by a narrow wave packet that expands fast.

tum uncertainty. Now imagine an intense pulse of gamma rays is fired at the region where the electron must reside, so that it will receive a substantial impulse and will fly towards an array of detectors. However, the gamma pulse is shaped such that a small region in the center of the wave packet is spared. That means, in some trials the electron will not be registered by the detectors, simply because it has been missed by the gamma rays. Quantum-mechanically speaking, after the "position measurement" performed by the gamma rays has failed to detect the electron, it is described by a much smaller wave packet, corresponding to the central region. *This* will have a considerable momentum uncertainty and will expand fast in space. That can be confirmed experimentally by subsequent position measurements.
Now try to explain this using a theory like "Stochastic Electrodynamics". The gamma rays are just part of the electromagnetic field, they corrrespond to an intense additional field superimposed onto the vacuum fluctuations. It is easy to understand that they give a kick to the electron if it is located in the region that is exposed to the rays, but they should not influence at all an electron that happens to be in the central region. Therefore, this electron could by no means acquire a sufficiently large momentum to explain the fast expansion of the probability density in the case where the electron has been missed by the gamma pulse. And, on the other hand, it cannot have had this large momentum even before the pulse, because that contradicts the observations made when no pulse is applied. The only way out seems to be either to postulate some nonlocal influence or at least some nonlinear interaction within the electromagnetic field, such that the field will be changed by the passage of the pulse. In any case, that would destroy the simple structure of a theory like this. Of course, physically the same thing happens already when a wave packet passes through a narrow slit and starts to expand afterwards in the transverse direction because of the momentum spread acquired according to Heisenberg's uncertainty relation. But for our purposes this experiment would not be suitable, because you could always argue that the matter making up the constriction will influence the vacuum electromagnetic field in the region of the slit

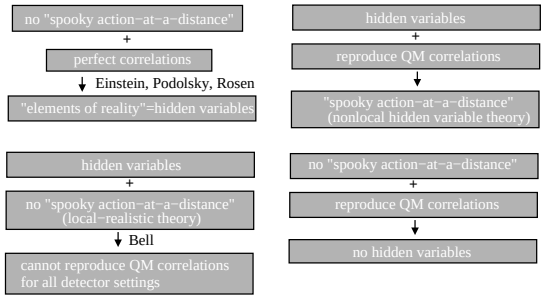

Figure 20: EPR's original reasoning, Bell's result, and its consequences

and this could indeed alter the trajectory of the electron.

5 John Bell's inequality

LR: I agree that it seems to be difficult to build a theory that gives the same results as quantum mechanics, like Bohm's theory, and which, at the same time, only involves local influences, like Stochastic Electrodynamics. On the other hand, there are so many complicated phenomena in the field of classical nonlinear dynamics and chaos. I would guess there must be *some* such description, perhaps involving fluctuating fields with nonlinear field equations. The particles might be some sort of singularities in the field, like the "soliton" pulses known from nonlinear optics, or some kind of vortices...

QM: Certainly it would be hard to invent a crucial "gedanken experiment" for each and every such theory in order to demonstrate that it gives results that are incompatible with quantum mechanics. And even if you succeeded in each case, still you could never be sure that there is not a working local hidden variable theory waiting around the corner, so to speak.
But fortunately, we do not have to analyze every such theory anew. There is a general theorem by John Bell which rules out *all* local hidden variable theories at once!

LR: How is that possible? We cannot enlist all of those theories.

QM: The idea is to use only the general features common to all of these theories to show that in a certain kind of experiment they must necessarily give results which deviate from those of quantum mechanics. This experiment is a variation on a gedanken experiment due to Einstein, Podolsky and Rosen (EPR), whose intention had been to prove that there *should* be local hidden variable theories!
Now let me explain the EPR experiment, as it was proposed originally. This was before Bohm invented his hidden variable theory, so it was not even clear

at the time whether any hidden variable theory existed at all! They wanted to demonstrate that one should indeed expect there to be such a "deeper" description underlying quantum mechanics, just as classical mechanics underlies statistical physics.
Of course, quantum mechanics tells us, for example, that it is impossible to measure precisely both the position and the momentum at the same time, in contrast to classical mechanics. But this does not necessarily mean that one could not at least *think* of them as having some well defined values simultaneously, *before* the measurement. After all, Heisenberg himself has explained how the measurement of the position will invariably disturb the momentum, and vice versa, in his "gamma ray microscope" gedanken experiment. Therefore, one might conclude that the particle does have a well-defined momentum *and* a well-defined position as long as we do not start to make a measurement. The problem with this "realistic" interpretation is simply that you could never hope to prove it in an experiment, because of this unavoidable disturbance. It is the same story as with the trajectories in Bohm's theory, which cannot be observed. Therefore, a skeptic might argue that you should not attribute reality to both position and momentum, because that exceeds what actually can be measured. It seems to become a matter of taste whether you want to believe in position and momentum taking on some definite values before the measurement.
However, EPR invented a clever scheme involving two particles, where the measurement on one of the particles will tell us something about the momentum or position of the *other* particle as well, because the properties of the particles are correlated. If they are far apart, this should *not* involve any disturbance of the other particle, so the usual objection is avoided. In order to introduce the desired correlation between the positions and the momenta of the two particles, they are assumed to be prepared in a wave function like this:

$$\Psi(x1, x2) \propto \delta(x_1 - x_2 - x)$$

This is an entangled state, since it cannot be written as a product of two functions that only depend on x_1 and x_2 separately. Obviously, the *relative* coordinate $x_1 - x_2$ takes on a precisely defined value, namely x. This necessarily means the relative momentum $p_1 - p_2$ can take on every possible value, with equal probability, due to Heisenberg's uncertainty relation. On the other hand, the wave function does not depend at all on the center-of-mass position $(x_1 + x_2)/2$. That position can therefore take on any possible value. This is consistent with the fact that the center-of-mass momentum $(p_1 + p_2)/2$ is precisely defined, it is equal to zero.
Now the argument of EPR goes like this: if one were to measure the position x_1 of the first particle, one could immediately conclude that the position of the other particle must be $x_2 = x_1 - x$, even before any measurement. If one actually *does* the position measurement on the second particle, one will obtain precisely that value. On the other hand, one could still choose to measure instead

the *momentum* p_2 of the second particle. Since the first measurement cannot possibly disturb the far-away second particle, one may be sure that p_2 is still equal to its "original" value. Similarly, if one measures p_1, one can immediately deduce the value of the momentum $p_2 = -p_1$, but one can still choose to measure x_2. In this way, it is possible to find out both the values of the position and the momentum of the second particle. Therefore, it is not a matter of taste any more whether to think of them as having some definite values prior to the measurement. Since quantum mechanics only makes statistical predictions and does not say anything in detail about the behaviour of the position and momentum of a given particle, it is "incomplete". It is not the full truth, so to speak. That was the conclusion of EPR, although, at that time, they could not yet present a working hidden variable theory to bolster their conclusion.
There is one important assumption in the argument, which EPR stated right at the beginning: It is assumed that, in accordance with the theory of special relativity, no faster-than-light influences can travel from one particle to the other. Therefore, if the particles are sufficiently far apart and the measurements are carried out in quick succession, there is no time for any signal to travel from one to the other particle, which could bring about a disturbance similar to that in the "gammy ray microscope". In other words, EPR *assumed locality* to show that the results which quantum mechanics gives for such a gedanken experiment prove the existence of "independent elements of reality", i.e. a hidden variable theory.

LR: Did they explain how a state like this could be prepared in the first place?

QM: No. But at least in principle this is a valid quantum-mechanical two-particle state. Of course, the wave function given here is symmetric under interchange of the particles, so it cannot describe fermions of equal spin direction, since then it would have to be antisymmetric. However, it is valid for the case of two atoms of spin zero, since they are bosons. In addition, one definitely needs some interaction between the two particles to produce this state. An indirect interaction may be enough for this purpose: For example, both of them could be scattered off the same object whose momentum would be measured before and after the event. Then momentum conservation would automatically introduce a correlation between the momenta p_1 and p_2, although it would still require some effort to produce exactly *that* state.
In fact, nowadays there are even approximate realizations of this original EPR state, in the field of quantum optics, where the coordinates x_1 and x_2 may be thought of as the electric field amplitudes at two different points in space, rather than the positions of two particles. However, there the field oscillates, which does not correspond to free motion in the original problem.

LR: Suppose you have been able to produce this state. Isn't it true that it will only describe the situation at a single instant of time? After all, if we think of a particle that is prepared in a localized wave packet, it will soon start to spread out, because the momentum uncertainty is necessarily large. The same should happen here, in the case of the two particles that are prepared in a state which is localized at $x_1 - x_2 = x$ with respect to their relative coordinate.

Then, maybe, the whole procedure of measuring either position or momentum of each particle becomes problematic. Each measurement will take some time to proceed and the state may change significantly. At least it is not obvious that everything can be carried out in this idealized fashion or that the deviations will turn out to be unimportant for the argument.

QM: It is true that the two-particle wave function given above will spread, at least as long as the particles are moving freely.
However, we should not bother about any of these concerns. This is because there are other variants of the original EPR experiment that do not have such problems. In order to demonstrate the essential features of the EPR argument, we only need the following: First of all, two observables that cannot be measured simultaneously, such as position and momentum. Secondly, two particles which are in an entangled state, such that their properties become correlated with respect to these observables. And thirdly, a way to separate them and to measure either one or the other observable at each of the particles.
The most important of these EPR variants has been proposed by David Bohm. He considered a correlation between the *spins* of particles. The different spin components S_x, S_y and S_z cannot be measured simultaneously: If you send a particle through a Stern-Gerlach magnet that measures the z-component S_z, then you will end up with the particle whose spin points along $+z$ or $-z$, depending on whether it is found in the upper or lower beam. In both cases, any *subsequent* measurement of the x-component will only give $50/50$ spin up or down in x-direction, which doesn't reveal anything about what you *would* have obtained if you had chosen to measure the x-component in the first place. So the spin components replace position and momentum. That has the nice effect that now the state does not evolve with time, at least as far as the entangled spin part of the total wave function is considered and as long as there is no magnetic field.
The state that is assumed to describe the correlations between the spins of the two particles is the singlet state:

$$\frac{1}{\sqrt{2}}(|\uparrow\rangle_1 |\downarrow\rangle_2 - |\downarrow\rangle_1 |\uparrow\rangle_2)$$

This state is a superposition of two possibilities, each of which occurs with the same probability of $1/2$: either the spin of particle 1 points in $+z$ direction and, at the same time, that of particle 2 points in $-z$ direction, or it is the other way round. If we measure the z-component of the spin of each particle, we will always obtain results that are opposite to each other. This is analogous to the perfect anticorrelation between momenta in the original EPR proposal. However, there we could also choose to measure the position instead of the momentum. In that case, again a perfect correlation was found, namely $x_1 - x_2 = x$. Likewise, in the spin version of the EPR experiment, we want to find out what happens if we measure the spins in x-direction, instead of the z-direction.
Therefore, we should ask how the singlet state looks like if we write it in spin

states that refer to the x-axis. Do you remember the connection between z- and x- basis for a spin $1/2$?

LR: Of course: If the states describing spin up and down in z-direction are denoted as $|\uparrow_z\rangle$ and $|\downarrow_z\rangle$, then the analogous basis vectors for the x quantization axis are given by symmetric and antisymmetric superpositions:

$$|\uparrow_x\rangle = \frac{1}{\sqrt{2}}(|\uparrow_z\rangle + |\downarrow_z\rangle)$$

$$|\downarrow_x\rangle = \frac{1}{\sqrt{2}}(|\uparrow_z\rangle - |\downarrow_z\rangle)$$

QM: Now, if you use these equations to express $|\uparrow_z\rangle$, $|\downarrow_z\rangle$ in terms of $|\uparrow_x\rangle$, $|\downarrow_x\rangle$ and insert the results into the singlet state, you will find that many terms cancel. We end up with the *same* form of the singlet state again, although now it is written with respect to the x axis:

$$\frac{1}{\sqrt{2}}(|\uparrow_z\rangle_1 |\downarrow_z\rangle_2 - |\downarrow_z\rangle_1 |\uparrow_z\rangle_2) = \frac{1}{\sqrt{2}}(|\uparrow_x\rangle_1 |\downarrow_x\rangle_2 - |\downarrow_x\rangle_1 |\uparrow_x\rangle_2)$$

This only works because of the minus sign in front of the second contribution. If you were to transform to the y-basis, you would still recover the same form, and that holds for *any* other direction as well! The singlet state is *isotropic*, there is no preferred direction. This feature will become very convenient in our discussion of the correlations between the two particles' spins, although it is not a *necessary* condition for obtaining those striking results.

LR: I agree that there are perfect anticorrelations between the spin directions, regardless of the direction in which you measure the spin components. But this is not all too surprising. I can even make up a similar completely classical model. Suppose you have two asteroids orbiting around each other. Each of them may rotate around some axis, so they have intrinsic angular momentum in addition to the angular momentum associated with their orbital motion. They are similar to two particles with spin that interact with each other. In the course of their motion, they may exchange angular momentum, due to tidal forces. At some point, they might fly apart, with some random orientations of their angular momenta. But if you have been able to measure the total angular momentum in the beginning, you will find that, in the end, their angular momenta are perfectly correlated, since they always add up to the value of the total angular momentum which has been conserved. Of course, some part of it may have gone into the orbital angular momentum, so it is not as nice as with the two particles in the singlet state. Nevertheless, it demonstrates that even perfect correlations are not very surprising, in particular if conservation laws are involved. Why should such a gedanken experiment serve to rule out local hidden variable theories?

QM: Of course, it is very easy to imagine that perfect correlations are observed because they have been there already at the time at which the two particles separated. If you take this viewpoint, it even appears to support the assumption that there are hidden variables. After all, it is most natural to assume that the correlations between the measurement results arise because of such hidden variables that have become correlated at the time of interaction.
But there is an important difference between the classical angular momentum of a rotating rigid body and the spin: You cannot measure the different spin components simultaneously. You have to *choose* a single particular direction along which to measure the projection of the spin, and still you will always find only the two maximum values of "spin up" or "spin down" along that direction, not any intermediate results. Therefore, measuring the spin does not mean you obtain a vector, as in the classical case. Rather, you provide a direction and get either of two possibilities as a result.

LR: OK, so I have to start from spin $1/2$ particles rather than classical rotating bodies. Why not simply say: When the particles fly apart, their spins are pointing into some definite, but opposite, directions. However, in order to explain the isotropy, we postulate that the initial directions are random for each experimental run.

QM: If I understand you correctly, your initial states are of the type

$$|\uparrow_{\hat{n}}\rangle\,|\downarrow_{\hat{n}}\rangle$$

where the spin is written in a basis that belongs to some direction $\hat{n}$, which is chosen at random for each pair. Unfortunately, this does not work. Although the measurement results will be isotropic indeed, and there are obviously some correlations between the spin directions, these correlations will not be *perfect*: Suppose you choose to measure along x but the direction $\hat{n}$ of the current pair happens to be z. Then the results will be $50/50$ up/down, *independently* at each measurement apparatus. Therefore, sometimes you will get *both* spins up, for example. This is in contrast to what is expected for the singlet state.

LR: I see. The trick is to have isotropy, perfect anticorrelations and only up/down results, all at the same time...

QM: In fact, there *is* a working local hidden variable model that is able to reproduce these results. It has been invented by John Bell. In principle, it is just like the model which you proposed: randomly oriented initial spin directions, which are always exactly opposite for the two particles. However, these "spins" are modeled in a manner that goes beyond usual quantum mechanics. One has to do this in order to avoid the problem of having less than perfect correlations if the measurement direction is different from the direction $\hat{n}$ of the current pair.
The model can be visualized very easily by imagining the particle to be a globe that is painted black and white on its northern and southern hemisphere, respectively. The spin of the particle points along an axis that goes through the

poles. The rule for the spin measurement along an arbitrary direction $\hat{d}$ is the following: If the vector $\hat{d}$ points onto the northern hemisphere, the result is "up", otherwise it is "down". For any given pair, characterized by $\hat{n}$, and any measurement direction, described by $\hat{d}$, we get perfect anticorrelation between the results: If the direction $\hat{d}$ points onto the northern (black) hemisphere for particle 1, it will necessarily point onto the southern (white) hemisphere for particle 2, and vice versa. Since the directions $\hat{n}$ are distributed randomly, we also have isotropy.

LR: Fine. Now we have at least one working local hidden variable model for this situation. Of course, it sounds a bit unrealistic, but certainly we will be able to find some modifications to improve it. Where, then, is the *contradiction* with local hidden variable theories?

QM: There is no contradiction as long as you choose to measure both spins along the *same* direction. However, as soon as the two measurement directions are tilted with respect to each other, you find quantitative deviations between quantum mechanics and our toy model. Of course, *qualitatively*, they both give the same answer, namely a decrease of correlations. For example, in our model, it becomes possible that the measurement direction $\hat{d}_1$ points onto the northern hemisphere of particle 1 and $\hat{d}_2$ *also* points onto the northern hemisphere of particle 2, even though particle 2 is colored in exactly the reverse manner. This is because $\hat{d}_1$ and $\hat{d}_2$ do not point at the same place on the globe any longer. Therefore, there is some small percentage of cases where the perfect anticorrelations are not observed. As you can imagine, this percentage grows if you increase the angle between $\hat{d}_1$ and $\hat{d}_2$ even further. In fact, you can find that percentage if you imagine the two globes to be tilted with respect to each other by the angle between $\hat{d}_1$ and $\hat{d}_2$. Then the fraction of "wrong" cases is proportional to the area on globe 1 that has a different color from globe 2. It grows linearly with the tilt angle.

LR: Let me see: If I want to quantify these correlations, I could attribute a value of "$+1$" to spin up and "-1" to spin down. Then, if the measurement directions are the same, one will always get $+1, -1$ or $-1, +1$, but never $+1, +1$ or $-1, -1$ for the two spin measurements at the two particles. So the product of the results is always -1. These are the perfect anticorrelations. Now if we tilt $\hat{d}_1$ versus $\hat{d}_2$ by an angle θ, we will sometimes get the "wrong" results $+1, +1$ or $-1, -1$. This happens in a fraction θ/π of all cases: If $\theta = \pi$, the measurement directions $\hat{d}_1$ and $\hat{d}_2$ are opposite to each other and we *only* get the wrong results. If we take the average over the product, which is -1 for the right cases and $+1$ for the wrong ones, we have:

$$\begin{aligned}(-1)P(right) + (+1)P(wrong) = \\ (-1)(1 - \frac{\theta}{\pi}) + (+1)\frac{\theta}{\pi} = \quad \tfrac{2}{\pi}\theta - 1\end{aligned}$$

When the measurement directions are at right angles, at $\theta = \pi/2$, we get exactly zero: Then, there are as many "right" results as there are "wrong" ones and overall there are no correlations at all. If θ is negative, the fraction of the "wrong" results will be $|\theta|/\pi$, so we should use $|\theta|$ instead of θ.

QM: Now we would like to compare that result with what quantum mechanics has to tell us. We can understand quite easily what happens at $\theta = \pi/2$: Suppose the measurement at particle 1 is carried out along the z-direction and gives "spin down". This means the total singlet state is projected onto that part which is compatible with this measurement result:

$$\frac{1}{\sqrt{2}}(|\uparrow_z\rangle_1 |\downarrow_z\rangle_2 - |\downarrow_z\rangle_1 |\uparrow_z\rangle_2) \mapsto |\downarrow_z\rangle_1 |\uparrow_z\rangle_2$$

Now particle 2 is definitely "spin up" in z direction, so we would observe perfect anticorrelations if we also measured in that direction. However, the measurement at particle 2 is carried out in a direction perpendicular to z, let's take the x direction. We know that we will obtain $50/50$ spin up/down along x if we start from a spin up in z direction. Therefore, we get the same result as in our toy model: no correlations at all for $\theta = \pi/2$.

LR: This has been a particular choice of measurement directions. What happens for others?

QM: As long as they are perpendicular, we will get the same result, because the singlet is isotropic. But suppose we do not measure particle 2 along x, but along a direction which is tilted with respect to the z axis by an arbitrary angle θ. Then our task is to find the probability of having "spin up" in this more general case. That is easy if you know the expression for the state of "spin up" pointing along an arbitrary direction, written down in the z-basis. In any case, here is the result for the probability:

$$(\cos(\frac{\theta}{2}))^2$$

This is reasonable: At $\theta = 0$, we definitely get "up", at $\theta = \pi$, we have "spin down", so we get 0 for the probability of observing "up", and at $\theta = \pi/2$ both possibilities are equally likely, so we obtain $1/2$. Now this probability for "spin up" is equal to the fraction of "right" results, since we had measured "down" at particle 1. If we had measured "up" at particle 1, the same reasoning would go through and we would still get the same probability of having a "right" result, with perfect anticorrelation. Overall, quantum mechanics tells us that the *correct* expression for the correlation function is

$$\begin{aligned}(-1)P(right) + (+1)P(wrong) = \\ (-1)(\cos(\frac{\theta}{2}))^2 + (+1)(1 - (\cos(\frac{\theta}{2}))^2) = \quad -\cos(\theta)\end{aligned}$$

This also starts at -1, goes through 0 at $\theta = \pi/2$ and ends up with $+1$ at $\theta = \pi$. But for angles smaller than $\pi/2$, it gives *stronger* anticorrelations than our toy model.

This is the essential point of Bell's argument: For the singlet state, quantum mechanics predicts *stronger* correlations than any possible local hidden variable model!

Index