

Dimension and Illusion

Peter J. Lewis

September 28, 2011

And even as we, who are now in Space, look down on Flatland and see the insides of all things, so of a certainty there is yet above us some higher, purer region, whither though dost surely purpose to lead me— O my Priest, Philosopher, and Friend—some yet more spacious Space, some more dimensionable Dimensionality, from the vantage-ground of which we shall look down together upon the revealed insides of Solid things, and where thine own intestines, and those of thy kindred Spheres, will lie exposed to the view of the poor wandering exile from Flatland, to whom so much has already been vouchsafed. (Abbott 1952, 88)

Could the three-dimensionality of the world be an illusion? That is, could the world *appear* to have a different number of spatial dimensions than it *actually* has? If the question concerns mere possibility, then the answer is clearly yes—there are many ways this might be. The most straightforward, perhaps, is the scenario explored in *Flatland* (Abbott 1952). Consider a race of creatures whose movements and sensory experiences are confined to a plane within a three-dimensional world. Such a world appears to its inhabitants to be two-dimensional. However, the illusory nature of this appearance can be revealed to them by interaction with an ordinary three-dimensional being like us, via such phenomena as appearing “from nowhere” within a locked room, or describing the current state of a Flatlanders intestines. Similarly, we can imagine a race of creatures whose movements and sensory experiences are confined to a three-dimensional space within a four-dimensional world. Such a world appears to its inhabitants to be three-dimensional, and again the illusory nature of this appearance could be revealed by interaction with four-dimensional beings via the kinds of phenomena just mentioned.

So the three-dimensionality of the world *could* be an illusion. But is there any reason to think that it *is* an illusion? Certainly creatures don't materialize inside locked rooms, and if there are beings who can describe the current state of our intestines to us, they don't do so by hovering over us in a fourth dimension (unfortunately). However, modern physics has led to various kinds of claim that the dimensionality of the world is greater than three. First, special relativity arguably entails that the time dimension is not distinct from the three spatial dimensions, and hence that the world is four-dimensional. While there are interesting questions here about the sense in which the time dimension can be treated as "spacelike", I restrict attention for present purposes to spatial dimensions, excluding time. Second, string theory postulates that there are nine (or perhaps ten) spatial dimensions, but that all but three are "curled up" so small that they are irrelevant to our experience. A case might be made that in such a case the three-dimensionality of the world is an illusion, but since there is no direct evidence for string theory, I set this possibility aside as speculation. Finally, quantum mechanics represents the state of the world via a $3N$ -dimensional wavefunction, where N is the number of particles in the universe. Quantum mechanics is not just speculation; it is a well-confirmed theory. So perhaps there really is evidence that the three-dimensionality of the world is illusory. But to see whether this surprising conclusion it is warranted, we first need to understand why quantum mechanics requires a $3N$ -dimensional wavefunction.

1 The wavefunction

Consider a very simple world, consisting of just two particles moving around in three-dimensional space. We can pick an origin and three mutually orthogonal directions, and hence impose a coordinate system on the space; the state of the world at a time can then be represented as two points in this space, say $(1, 4, 3)$ and $(9, 2, 6)$. Alternatively, we can represent the same state as a single point in a six-dimensional space, namely $(1, 4, 3, 9, 2, 6)$. The latter is called a *configuration space* representation, since each point in the space represents a configuration of particles in ordinary three-dimensional space. Pretty clearly, the two representations are equivalent—any arrangement of particles in three-dimensional space corresponds to a point in configuration space, and conversely any point in configuration space corresponds to an arrangement of particles in three-dimensional space.

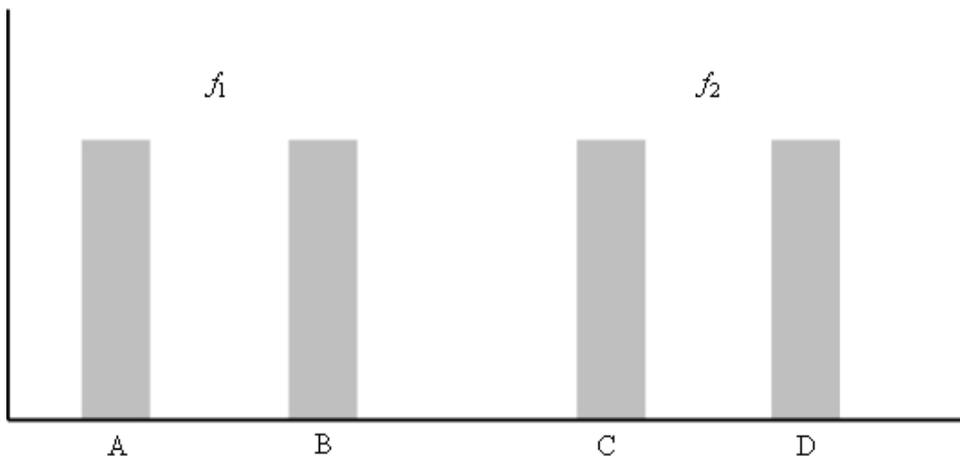


Figure 1: Field intensity in 3D space

Now suppose that, instead of particles, our toy world contains two *fields*—continuous distributions of some quantity over three-dimensional space, rather than discrete particles. That is, the three-dimensional space contains two objects—a function $f_1(x, y, z)$ and a second function $f_2(x, y, z)$, each representing the intensity of its respective field as a function of the three spatial coordinates. Again, we can represent the same state in a six-dimensional space as a function $F(x_1, y_1, z_1, x_2, y_2, z_2)$, where the first three coordinates are those of the field f_1 and the second three are those of f_2 . For example, suppose field f_1 has non-zero intensity only in two regions A and B, and f_2 has non-zero intensity only in two regions C and D, as shown (schematically) in Fig. 1. Here the horizontal axis represents the three spatial dimensions, reduced to one for ease of representation.

Alternatively, we can choose to represent the state of the system in the six-dimensional space shown (schematically) in Fig. 2(a). Here the horizontal axis represents the coordinates of f_1 , the vertical axis represents the coordinates of f_2 , and the shaded areas represent non-zero field intensity. In the f_1 coordinates, the field has non-zero intensity only in regions A and B, and in the f_2 coordinates, the field has non-zero intensity only in regions C and D, as in the three-dimensional representation. But note that exactly the same is true of the field distributions shown in Fig. 2(b) and 2(c). That is, any one of these six-dimensional field distributions (and many others besides) can

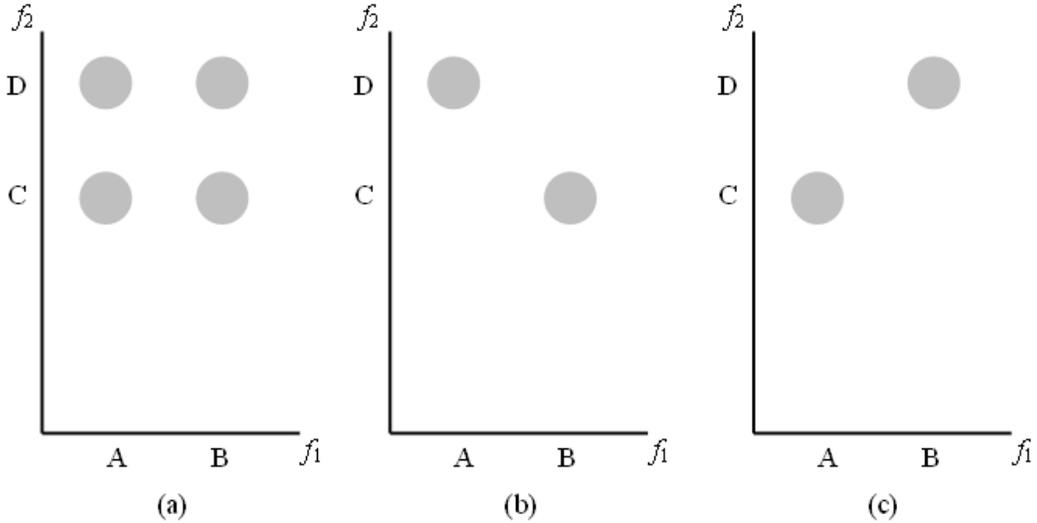


Figure 2: Field intensity in 6D space

adequately represent the three-dimensional field distribution of Fig. 1.

So any two fields in three-dimensional space can be represented by a field in six-dimensional space—and in fact there is considerable freedom in choosing the latter, since many such representations contain all the information in the three-dimensional representation. But by the same token, a field in a six-dimensional space cannot in general be represented without loss as two fields in three-dimensional space, since the six-dimensional field contains information that is not present in the corresponding pair of three-dimensional fields. This is reflected in the fact that the distinct six-dimensional fields of Figs. 2(a), 2(b) and 2(c) correspond to one and the same three-dimensional representation—that of Fig. 1.

In quantum mechanics, the state of a system of N particles can be expressed a function of $3N$ coordinates, 3 for each particle. This is the wavefunction, written $\psi(x_1, y_1, z_1, \dots, x_N, y_N, z_N)$. As this expression suggests, the wavefunction is most naturally represented in a $3N$ -dimensional space; in fact the wavefunction cannot be expressed as a set of N fields in three-dimensional space, because the empirical predictions of the theory depend on the information that is lost in the three-dimensional representation. For example, for a two-particle system, the three wavefunctions depicted in Figs. 2(a), 2(b)

and 2(c) would result in different predictions for the measured positions of the particles; Fig. 2(c) represents a state in which particle 2 is at C if and only if particle 1 is at A, Fig. 2(b) represents a state in which particle 2 is at C if and only if particle 1 is at B, and Fig. 2(a) represents a state in which there is no such correlation between the locations of the particles.

The foundations of quantum mechanics are notoriously contested, and the status of the wavefunction is no exception. But there is a *prima facie* case, at least, that the wavefunction should be regarded as the fundamental entity of quantum mechanics. It is the evolution of the wavefunction that is governed by the basic dynamical law of the theory, and it is the wavefunction at the end of an experiment that generates its empirical predictions. As J. S. Bell famously put it, “No one can understand this theory until he is willing to think of ψ as a real objective field . . . Even though it propagates not in 3-space but in $3N$ -space” (1987, 128). That is, realism in the context of quantum mechanics arguably commits you to the existence of an entity corresponding to the wavefunction of the world—a $3N$ -dimensional field, where N is the total number of particles in the universe. This is a radical violation of the intuition that there are three spatial dimensions—it is not that there are four, or ten, but that there are at least 10^{80} spatial dimensions.

2 Keeping up appearances

So why does it look to us as if there are three spatial dimensions, if in fact there are not? David Albert (1996) suggests that the answer lies in the dynamical law by which the quantum state evolves; an element of the dynamical law called the Hamiltonian takes a particularly simple form if the $3N$ spatial coordinates are grouped into N sets of 3 (rather than $3N/2$ sets of 2, or $3N/4$ sets of 4, etc.). The Hamiltonian is a representation of the energy properties of a physical system. In classical physics it is a function from the positions and momenta of the particles in the system to a real number representing the total energy of the system—kinetic energy plus potential energy. The use of the Hamiltonian in classical mechanics is somewhat arcane—it’s not the high school version of the theory—but it allows for a particularly succinct expression of the dynamical laws by which the state of the system evolves over time. In quantum mechanics, on the other hand, the Hamiltonian formulation of the theory is the canonical one. Properties are represented as operators rather than functions in quantum mechanics; the

Hamiltonian is a differential operator acting on the wavefunction. But as in classical mechanics, the Hamiltonian acting on a wavefunction represents the energy properties of the system with that wavefunction, and as in classical mechanics, this representation allows for a very succinct expression of the dynamical law by which the wavefunction evolves. This dynamical law—the Schrödinger equation—says that (up to a constant) the rate of change of the wavefunction is equal to the result of applying the Hamiltonian operator to the wavefunction. In brief, for both classical and quantum mechanics, the Hamiltonian tells you how the total energy varies across possible states of a system, and plays a key role in one formulation of the dynamics for the system.

With this understanding of the Hamiltonian in mind, let us return to Albert’s argument. Note first that there is a sense in which classical behavior emerges from quantum mechanical behavior in the macroscopic limit—that is, as systems become large and complicated. The sense is that while microscopic systems must typically be represented by a spread-out wavefunction in configuration space, macroscopic systems can always be represented to a good degree of approximation by a *point* in configuration space. And a point in $3N$ -dimensional configuration space, as shown above, can equally well be represented as N points in three-dimensional space. Of course, it can also be represented as $3N/2$ points in a two-dimensional space, and so on. But if we choose to group the coordinates into threes, then the Hamiltonian takes a particularly neat form; the potential energy term depends only on the distances between the N “particles”. On the other hand, if we choose to group the coordinates into two or fours or sevens, then the potential energy term will bear no straightforward relation to the “interparticle” distances so produced.

The key point is that the form of the Hamiltonian one obtains by grouping into threes corresponds to our classical world-view; potential energy, be it gravitational, electrostatic or some other, is taken to be generated by forces that depend only on the distances between the objects involved. Indeed, elementary classical mechanics takes these interparticle force laws (e.g. the inverse-square laws of gravitation and electrostatics) to be fundamental, and the Hamiltonian (if one thinks of it at all) is taken to be a construction based on these fundamental laws. If Albert is right we have to give up this picture, since there are no particles moving in three-dimensional space and hence no laws governing their motion; it is the wavefunction and Hamiltonian in $3N$ dimensions that are fundamental. But his claim is that we can nevertheless

explain why our intuitive picture of the world has three dimensions, since it is the obvious (though false) interpretation of the behavior of macroscopic objects. That is, even though there is nothing in fundamental reality corresponding to our choice of coordinate grouping, if we choose to group the coordinates in threes a particular description of the behavior of medium-sized everyday objects becomes available to us, namely the *classical* description of objects moving in three-dimensional space, subject to forces that depend on the distances between them. Hence “quantum-mechanical worlds are going to appear (falsely!) to their inhabitants, if they don’t look too closely, to have the same number of *spatial dimensions* as their *classical counterparts* do”, namely three (Albert 1996, 282). The reason for the caveat, of course, is that if we *do* look closely—if we perform experiments that reveal the underlying quantum-mechanical nature of microscopic reality—then we convince ourselves that the world can’t really be three-dimensional via the arguments of the previous section.

If Albert is right, then the world is really $3N$ -dimensional, yet appears three-dimensional to us. But is he right? There are reasons to be skeptical (Monton 2002; Lewis 2004). First, note that the above story about grouping the coordinates into threes is somewhat simplified; more is required to recover the coordinate system in which the Hamiltonian takes the familiar three-dimensional form. Consider a classical two-particle system, represented by a point in a six-dimensional space. The choice of origin and the order in which we list the coordinates is conventional; we can choose any system we like. But starting from such a coordinate system, we cannot recover the three-dimensional coordinates in which the standard force laws hold simply by grouping the coordinates into threes. First, we have to make sure that the axes line up—that the two x -axes, for example, point in the “same” direction. Second, we have to make sure that the coordinates are ordered in the “right” way, such that the first and fourth are the x -coordinates of the (fictional) particles, the second and fifth are the y -coordinates, and the third and sixth are the z -coordinates. Finally, we have to make sure that the origins for the two x -coordinates “coincide”, and similarly for the y - and z -coordinates. Note all the scare quotes here. The two x -axes can’t really point in the same direction, since they are orthogonal in configuration space, but nevertheless one particular relative orientation is required for the Hamiltonian to take the right form. Similarly for the ordering and the origin; no coordinates are really x -coordinates, so the origins of the two x -coordinates cannot really coincide. If there isn’t really a three-dimensional space, it seems that there’s no

explanation for these requirements other than that they work to produce the familiar Hamiltonian. Indeed, one might start to suspect that requirements like this are part of what it takes for a space to be three-dimensional.

On a related note, our picture of force laws in three dimensions is more than just an expendable intuition; it is the explanation of the form of the Hamiltonian. It is not just high school physics that takes force laws between three-dimensional objects as fundamental; when quantum mechanics is applied to a physical system, the potential energy term in the Hamiltonian is constructed by appeal to such laws. Under Albert's way of looking at things, we would have to forego this explanation of what we are doing. Albert is free to postulate that the Hamiltonian is unexplained and fundamental, but then it becomes mysterious why the method of constructing the Hamiltonian from the three-dimensional configuration of our system works at all.

Finally, although grouping the coordinates into threes yields a particularly simple Hamiltonian, this simplicity cannot be parlayed into an explanation of three-dimensional appearances. It is true that a particular choice of coordinates sometimes plays a role in the explanation of appearances. For example, I might choose coordinates in which my office is at rest (rather than ones in which the sun is at rest) in order to explain my experience of thunderclouds building behind the skyscrapers of downtown Miami. But the dependence of the explanation on the choice of coordinates here is superficial; it is simply a matter of calculational convenience. Even if I used the sun-centered coordinates, it would still be the case that the clouds, the buildings and my office lie in a straight line. The coordinate-dependence in Albert's account is much more thoroughgoing—to an extent that makes it problematic. For one thing, three points that lie on a straight line under one grouping of $3N$ -dimensional coordinates will not, in general, lie in a straight line under another grouping. For another, the patterns of points that we identify as buildings and clouds will not, in general, be present under another grouping. But patterns that appear only under one arbitrary choice of coordinates are generally regarded as artifacts of that choice rather than facts about the world, since the facts about the world presumably don't depend on our representational choices. If there are no buildings and clouds and offices that lie in a straight line, then the explanation of my experience evaporates.

But now we seem to have painted ourselves into a corner. The state of a quantum mechanical system can't be represented without loss in three dimensions; it has to be represented in a $3N$ -dimensional configuration space. But the evolution of the wavefunction in $3N$ dimensions can't explain our

three-dimensional appearances. So either quantum mechanics is wrong, or we are radically deceived about the nature of our own experience. Neither of these looks like an attractive option.

3 Configuration space

But all is not lost, since there is another option available to us—namely that we are wrong about the structure of space. A rather flatfooted solution along these lines would be to propose that there are in fact two spaces—the $3N$ -dimensional space in which the wavefunction evolves, and the three-dimensional space in which the objects of our experience move around—neither of which can be reduced to the other. But this proposal introduces as many problems as it solves; how does the motion of the wavefunction in one space give rise to the motion of objects in a completely separate space (Monton 2006)? A more subtle solution is that the three-dimensional structure is already present in the configuration space. Thus far, we have been treating the space in which the wavefunction evolves as $3N$ -dimensional in the same sense that ordinary space is three-dimensional—essentially, as a space spanned by $3N$ mutually orthogonal vectors and having no other structure. But perhaps there's more to configuration space than initially meets the eye.

Consider ordinary three-dimensional space. To impose coordinates on this space, you choose an origin, a length scale, and three mutually orthogonal axes. We have been supposing thus far that the coordinatization of configuration space is just like this; to impose coordinates on a $3N$ -dimensional configuration space, you choose an origin, a length scale, and $3N$ mutually orthogonal axes. It is precisely this freedom in choosing coordinates that leads to the problems facing Albert's argument, since most such choices do not yield the simple Hamiltonian or the objects of ordinary experience.

But suppose instead that we take seriously the idea of a configuration space as a space of configurations—that is, a space which is *intrinsically* structured as N sets of three-dimensional coordinates. Mathematically, this is not hard to do. Instead of modeling the space as an ordered $3N$ -tuple of parameters, $\langle x_1, x_2, \dots, x_{3N} \rangle$, we model it as an ordered N -tuple of ordered triples:

$$\langle \langle x_1, y_1, z_1 \rangle, \langle x_2, y_2, z_2 \rangle, \dots, \langle x_N, y_N, z_N \rangle \rangle.$$

And rather than specifying the coordinates by choosing $3N$ axes, we choose

3—the x , y and z axis, which are the same for each triple. That is, x_1 through x_N pick out points on the *same* axis, and similarly for y and z . Then the wavefunction can be regarded as a function of *these* parameters—as a mathematical entity inhabiting a $(3 \times N)$ -dimensional configuration space, rather than a $3N$ -dimensional plain space. And the basic thesis of wavefunction realism is that the world has *this* structure—the structure of a function on $(3 \times N)$ -dimensional configuration space.

Given that configuration space has this structure, then an Albert-style appeal to dynamics to generate three-dimensional appearance is impossible, but it is also unnecessary. It is impossible because the Hamiltonian takes exactly the same form under every choice of coordinates, so no choice makes it especially simple. But it is unnecessary because the outcome of that argument—that the coordinates are naturally grouped into threes—is built into the structure of reality, and hence doesn't need to be generated as a mere appearance based on the simplicity of the dynamics.

4 Dimension without illusion

The proposal of the previous section is designed to allow us to say that the world of quantum mechanics *really is* three-dimensional, and hence the three-dimensionality of appearances doesn't have to be generated as any kind of illusion. Our appearances are veridical. But does the proposal really allow us to do that? The wavefunction is still a function of $3N$ independent parameters, even if those parameters have some internal structure. Isn't that a prima facie reason to say that the wavefunction—and hence the world—is $3N$ -dimensional?

Previously, I took the lesson here to be that the term “dimension” is ambiguous in the quantum-mechanical world (Lewis 2004). One can take it to refer to the number of independent parameters required to specify a point in the space in which the quantum state evolves, or one can take it to refer to the number of independent axes required to impose coordinates on the space. In the classical case, these two coincide; a point in the space in which the classical state evolves is specified by three parameters, and imposing coordinates on the space requires three axes. But in the quantum case, the two dimension concepts come apart; it takes $3N$ parameters to specify a point in the space in which the wavefunction evolves, but only three axes are required to impose coordinates on this space. Hence if quantum mechanics

is true, there is a sense in which the world is $3N$ -dimensional and a sense in which it is three-dimensional. It is the latter sense under which the world *truly* appears three-dimensional.

My 2004 position is essentially the same as the one Albert adopts in the coda to his 1996 paper (although our reasons for adopting the position are different). Here Albert restates his thesis “a bit more diplomatically” (1996, 282). The diplomatic version of his thesis is that there are “two ideas we’re accustomed to having in mind when we think of ‘physical space’”. The first is “the space of possible *interactive distances*”—that is, the space in which the distances between points are the “interparticle distances” appearing in the Hamiltonian. This space is three-dimensional—the dynamical laws are laws of a three-dimensional world. The second is the space “in which a specification of the local conditions at every address at some particular time (but not at any proper *subset* of them) amounts to a complete specification of the physical situation of the world, on that theory, at that time”. This space is $3N$ -dimensional; it takes a specification of the wavefunction amplitude at every point in a $3N$ -dimensional space to completely specify the quantum state of the world. So, diplomatically speaking, it is not that the three-dimensionality of the world is an illusion; rather, the world really is three-dimensional under one reading of “dimension”, and it really is $3N$ -dimensional under the other reading.

This position still seems tenable to me, but there is an alternative position that I now find more attractive. Albert initially claimed that the world is $3N$ -dimensional, and the appearance that the world is three-dimensional (when we don’t look too closely) is illusory. Later he claimed (and I concurred) that the three-dimensional appearances (when we don’t look closely) and the $3N$ -dimensional appearances (when we do) can both be taken as veridical, since “dimension” is ambiguous. However, there is a third possibility—namely that the world is three-dimensional, and the $3N$ -dimensional appearance of the world when we look closely is illusory. That is, our everyday impression that the world is three-dimensional is correct, but the impression we get from our quantum-mechanical experiments that the world is $3N$ -dimensional is misleading. This third position seems to have something to be said for it, and in what follows I explore and defend it.

5 Dimensions and parameters

The claim that the quantum world is $3N$ -dimensional is based on the fact that the wavefunction is a function of $3N$ independent parameters. This is a perfectly good characterization of “dimension” in some general sense of the concept. But the number of independent parameters is not always a good way to characterize the *spatial* dimensionality of a system. Sometimes a parameter is not a spatial parameter at all. For example, in evolutionary game theory, the state of a population of organisms can be represented as a function of n parameters, one for each organism, where each parameter represents the continuum of possible strategies the organism can adopt in interacting with the others, and the function represents the probability distribution over the strategy space—i.e. the chance that organisms will use particular strategies in their interactions with each other. Clearly nothing should be inferred from this model about the (literal, rather than figurative) space in which the organisms live.

But of course in the quantum case the parameters *are* spatial; each is a position coordinate for a particle. Still, it doesn’t follow that the representation is of a spatially $3N$ -dimensional world. Consider again the case of evolutionary game theory. Insofar as such a model is intended to be realistic, it is intended to be about n individual organisms, each with its own probability distribution for adopting a particular strategy. But there may be information in the full n -dimensional representation that is lost when the population is represented as n single-organism states. We can reinterpret Figs. 1 and 2 to demonstrate this. Suppose Fig. 1 represents the strategy distributions for two organisms; organism 1 adopts strategies of type A and type B with equal probabilities, and organism 2 adopts strategies C and D with equal probabilities. But the single-organism properties of Fig. 1 are compatible with each of the three n -organism states represented in Fig. 2. Fig. 2(a) represents a state in which the strategy adopted by each organism is independent of that adopted by the other organism, and Figs. 2(b) and 2(c) represent states in which the strategy adopted by each organism depends on what the other does. Clearly the future evolution of the system depends on which n -organism state the population has. But this wouldn’t warrant the inference that reality consists of an n -organism entity rather than n individual organisms. The organisms just happen to have complicated conditional properties; their strategies depend on those of the organisms with which they interact. The most convenient way to represent these conditional

properties—perhaps the only convenient way—is via the n -dimensional state, but nothing ontological should be read into this.

One can adopt a similar position with regard to the quantum mechanical wavefunction. That is, the fact that the state of a quantum system can be represented without loss as a single wavefunction in a $3N$ -dimensional configuration space but not as N single-particle wavefunctions in three-dimensional space does not entail that the world is spatially $3N$ -dimensional. Rather, one could interpret the situation as one in which the particles have complicated conditional connections; the position properties of one particle, as encoded in the wavefunction, depend on those of the other particles (c.f. Monton 2012). The $3N$ -dimensional wavefunction may be the only convenient way to represent these properties, but it does not follow that it literally represents the spatial structure of the world. This is not to give up the assumption of wavefunction realism; the structure of the wavefunction accurately reflects the structure of the world, but some of that structure consists in correlations between spatial positions. The wavefunction is a function of $3N$ parameters, but we need not interpret each parameter as an independent spatial direction.

6 Spatial phenomena

So far, I have argued that we *need* not interpret the quantum world as spatially $3N$ -dimensional; we can't simply read the spatial structure of the world off the mathematical structure of our representation. But this leaves open the question of how we *should* interpret the wavefunction. The question is whether the $3N$ parameters of configuration space deserve to be called spatial dimensions. And the answer, I suggest, hangs on the connections between those parameters and spatial phenomena.

The claim that the quantum world is *three*-dimensional is based on a fairly direct correspondence between the structure of configuration space and the structure of our spatial experience. I can stick my arm out in some arbitrary direction and stipulate that every third configuration-space parameter refers to *that* direction, and similarly pick two directions orthogonal to my arm to correspond to the remaining configuration-space parameters. Hence the sense in which the world is three-dimensional is straightforwardly a spatial sense. But no such direct correspondence to experience is available for the claim that the quantum world is $3N$ -dimensional; if the first and fourth configuration space parameters are in fact orthogonal spatial directions in

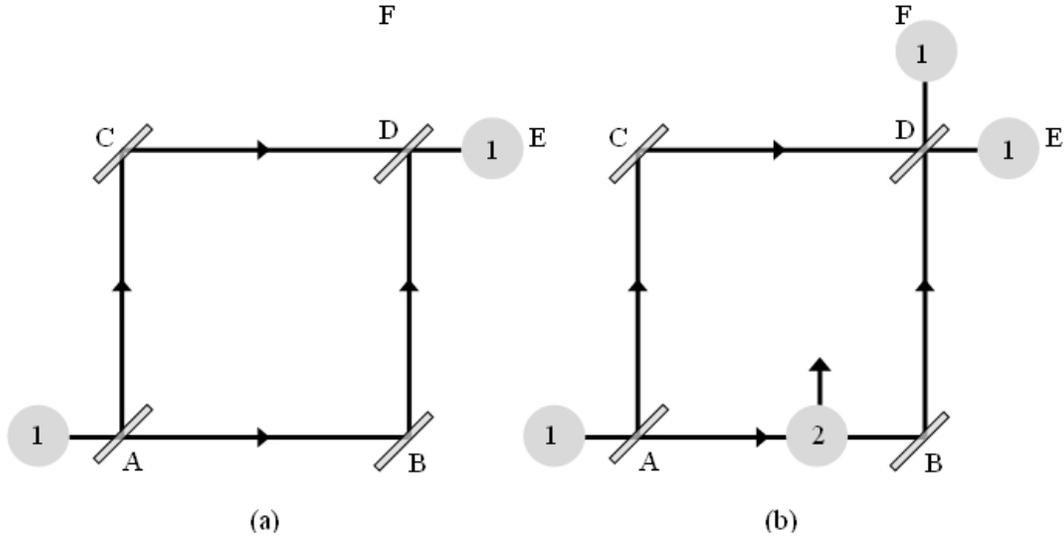


Figure 3: Mach-Zehnder interferometer

some sense, this is no part of my experience. Hence the doubt that the $3N$ parameters should be characterized as spatial.

Still, an indirect correspondence to experience might be available. Consider again the Flatland scenario that we started with. The two-dimensional Flatlander suspects that he is really living in a three-dimensional world because various phenomena he observes are explicable in terms of three dimensions, but not in terms of two. Similarly, various phenomena in the quantum world are explicable on the assumption that the quantum state evolves in a $3N$ -dimensional space, but arguably not under the assumption that it evolves in a three-dimensional space. Consider, for example, the Mach-Zehnder interferometer shown in Fig. 3(a). A single-particle wave packet (labelled 1 in the diagram) enters the device at the bottom left and is split into two equal wave packets by the half-silvered mirror A. The two packets are reflected by mirrors at B and C respectively, and both arrive at the second half-silvered mirror D. If the path lengths ABD and ACD are exactly equal, then the wave components emerging towards detector F are exactly out of phase—the troughs of one wave coincide with the peaks of the other—and the two waves exactly cancel out. By the same token, the wave components emerging towards detector E are exactly in phase—the peaks coincide and the troughs

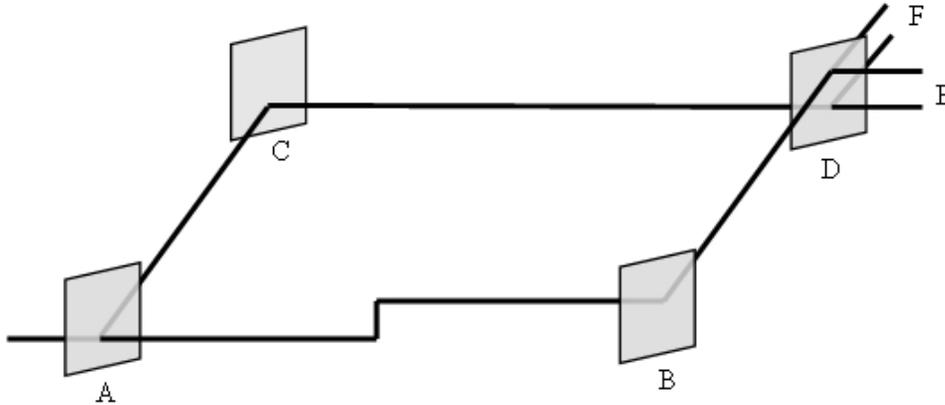


Figure 4: Configuration space representation

coincide—and the two waves add in intensity. The result is a single wave packet, of exactly the intensity of the one input at the bottom left, emerging towards detector E. That is, the two wave packets travelling via B and C collide and interact at D, producing a single wave packet travelling towards E.

Now suppose that the wave packet travelling via B interacts with a wave packet associated with another particle (labelled 2 in the diagram), and the wave packet travelling via C does not, as shown in Fig. 3(b). What happens now is that the two wave packets arriving at D no longer interact, resulting in wave packets travelling onwards towards E and F. It is as if the wave packets pass by each other rather than colliding. Indeed, if one adopts a configuration space representation, this is exactly what happens; the packet travelling via B undergoes a shift in the coordinates of the second particle, whereas the packet travelling via C undergoes no such shift. I have attempted a configuration space representation in Fig. 4. The full configuration space has six dimensions; I have chosen three to include in the diagram, the two horizontal dimensions representing the plane of the experiment in the coordinates of particle 1, and the vertical dimension representing one of the coordinates of particle 2. A wave packet representing both particles enters at the bottom left and is split in the coordinates of particle 1 by mirror A. The vertical jog on the lower path represents the collision of the particles; particle 2 moves, so the trajectory shifts in the coordinates of particle 2. But now the two wave

packets don't hit the same spot in the (five-dimensional) plane of mirror D, since they differ in the coordinates of particle 2. Hence no interference occurs at D, and each packet is again equally split into two further packets in the coordinates of particle 1.

Here, then, we apparently have indirect evidence that the quantum world is $3N$ -dimensional. The extra dimensions seem necessary to explain the interactions, or lack of them, between wave packets. But this evidence needs to be treated with considerable caution. Note in particular that wave packets corresponding to *different* particles cannot pass by each other in the way just described; if two wave packets corresponding to two different particles approach the same region of three-dimensional space, then they interact. Consider the interaction between two particles depicted in Fig. 5. Fig. 5(a) is the ordinary three-dimensional representation; the wave packets approach and bounce off each other, and this is the case even if particle 1 or 2 interacts with some other particle on its way to the collision point. Fig 5(b) depicts the same interaction in configuration space, with the vertical axis representing the x -coordinate of particle 1 and the horizontal axis the x -coordinate of particle 2. The two-particle wave packet travels up to the diagonal line where the coordinates of the two particles coincide, and here the trajectory suddenly changes direction. The trajectory in configuration space acts as if it bounces off a reflective wall, a wall that extends indefinitely in the coordinates of any other particles involved. But of course there is no physical wall; this behavior is purely an expression of the dynamical laws. Pictured in configuration space, these laws look mysterious; why should the trajectory suddenly change direction just there? But pictured in three-dimensional space, the reason for the sudden turn is obvious; the particles collide. The phenomena are not what one would expect of objects moving in a $3N$ -dimensional space, but they are precisely what one would expect of objects moving in a three-dimensional space.

Of course, Albert might reply here that I have switched back to the sense in which configuration space is three-dimensional—namely that the dynamical laws are those of a three-dimensional space. Quite right—but my contention here is precisely that the usual sense of the term “spatial” is intimately connected with dynamical laws. Spatial phenomena concern whether or not objects bounce off one another or pass by one another—and these are dynamical concepts. When we are dealing with wave packets for distinct particles, three-dimensional space, not $3N$ -dimensional space, is the arena in which spatial phenomena play out. What's more, the interactions of

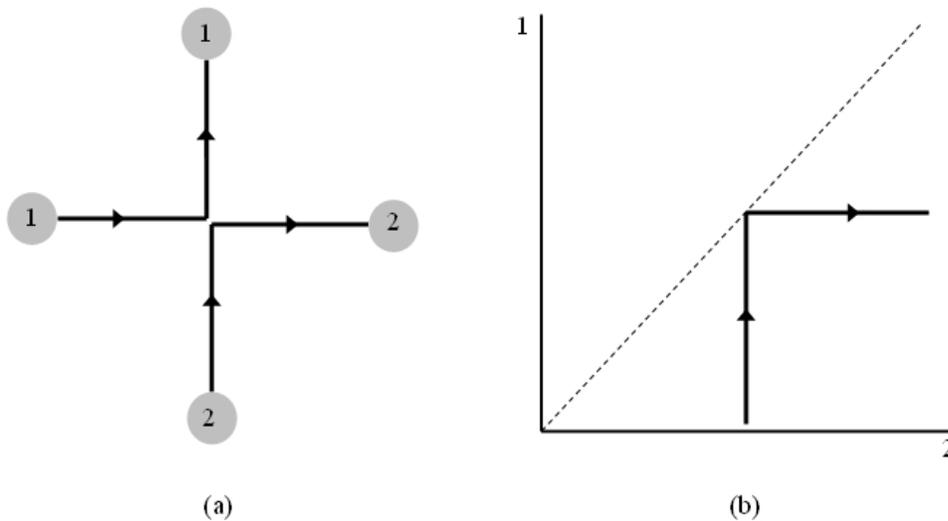


Figure 5: Interaction between two particles

ordinary macroscopic objects play out in three-dimensional space, and it is this kind of behavior that ultimately underlies our ordinary use of the term “spatial”.

Admittedly, consideration of the example with which we started this section complicates this story. Two wave packets that are components of the state of one and the same particle sometimes interact and sometimes pass by each other when their three-dimensional coordinates coincide. Doesn’t this require the existence of extra dimensions in which the passing-by can take place? Well, the lesson of the previous section is that such phenomena don’t *require* the existence of extra spatial dimensions. Certainly one needs *parameters* in the theory, the values of which determine whether or not the packets interact. And in the quantum case, the parameters in question refer to the coordinates of the *other* particles in the system—i.e. they encode how the wave packet for the particle we are following is correlated with the wave packets for the other particles. The question is whether we should regard this structure of parameters as spatial. And we can now see clearly the costs and benefits of each answer. If we think of spatial phenomena as $3N$ -dimensional, we get a nice literal interpretation of the operation of the interferometer (Fig. 4), but the behavior of everyday objects becomes mysterious (Fig. 5(b)). If we think of spatial phenomena as three-dimensional,

then we can go on thinking of everyday phenomena in terms of collisions between objects (Fig. 5(a)), but the operation of the interferometer has to be taken as involving correlation rather than spatial location, since the spatial locations don't determine the particle's behavior (Fig. 3). To my eye, the choice seems clear; the former choice does considerable violence to our ordinary spatial notions, whereas the latter just requires a slightly less intuitive understanding of interference phenomena. Even in the latter case, it's not as if we can't *use* the $3N$ -dimensional representation to picture interference to ourselves; it's just that we can't take that picture as a direct *spatial* representation.

7 The pragmatic dimension

The world looks three-dimensional unless one looks closely, when it looks $3N$ -dimensional. But which appearance is veridical, and which the illusion? Albert initially contends that the three-dimensionality of the everyday world is illusory, and that $3N$ -dimensional wavefunction one discerns in quantum phenomena is the reality behind the illusion. What I have tried to do here is to argue for the converse of Albert's initial position; the world really is three dimensional, and the $3N$ -dimensional appearance of quantum phenomena is the theoretical analog of an illusion; we represent quantum reality to ourselves as $3N$ -dimensional in order to more readily visualize the correlations between wave packets. The basic reason for thinking of things this way round is that the sense in which the quantum world is three-dimensional corresponds directly to the way we already use the term "spatial", whereas the sense in which it is $3N$ -dimensional does not.

Still, even if it would be a stretch to apply the adjective "spatial" to the $3N$ dimensions of configuration space, there is nothing to stop us doing so; it would not be the first time that ordinary usage has shifted as a result of scientific advances. This is ultimately a pragmatic matter. But I find it hard to see any motivation for such a shift, since the phenomena in question are so far removed from everyday life. We modified our concepts of rest and motion to accommodate the idea that the Earth is moving, but doing so also gave us the tools to better describe everyday phenomena. If a shift in our concept of spatial dimension would reap similar benefits, then no doubt such a shift should and will occur. But until then, we can say with a straight face that the world is three-dimensional.

Acknowledgments

I would like to thank Alyssa Ney and David Albert for very helpful comments on an earlier draft of this paper.

References

- Abbott, Edwin A. (1952), *Flatland*. New York: Dover.
- Albert, David Z. (1996), “Elementary Quantum Metaphysics”, in J. Cushing, A. Fine and S. Goldstein (eds.), *Bohmian Mechanics and Quantum Theory: An Appraisal*. Dordrecht: Kluwer, pp. 277–84.
- Bell, John S. (1987), *Speakable and Unspeakable in Quantum Mechanics*. Cambridge: Cambridge University Press.
- Lewis, Peter J. (2004), “Life in Configuration Space”, *British Journal for the Philosophy of Science* 55: 713–29.
- Monton, Bradley (2002), “Wavefunction Ontology”, *Synthese* 130: 265–77.
- (2006), “Quantum Mechanics and $3N$ -Dimensional Space”, *Philosophy of Science* 73: 778–89.
- (2012), “Against $3N$ -Dimensional Space”, this volume.