

# 1

## Introduction

Under normal conditions the research scientist is not an innovator but a solver of puzzles, and the puzzles upon which he concentrates are just those which he believes can be both stated and solved within the existing scientific tradition.

– *Thomas Kuhn, The Essential Tension, 1977.*

Quantum theory has been puzzling physicists and philosophers since its birth in the early 20th century. However, starting in the 1980s, rather than asking why quantum theory is so weird, many people started to ask the question:

*What can we do with quantum weirdness?*


In this book we not only embrace this perspective shift, but challenge the quantum icons even more. We contend that one should not only change the kinds of questions we ask about quantum theory, but also:

*change the very language we use to discuss it!*

Before meeting this challenge head-on, we will tell a short tale to demonstrate how the quantum world defies conventional intuitions ...

### 1.1 The Penguins and the Polar Bear

Quantum theory is about very special kinds of physical systems – often very small systems – and the ways in which their behaviour differs from what we observe in everyday life. Typical examples of physical systems obeying quantum theory are microscopic particles such as photons and electrons. We will ignore these for the moment, and begin by considering a more ‘feathered’ quantum system. This is Dave:



He's a dodo. Not your typical run-of-the-mill dodo, but a *quantum dodo*. We will assume that Dave behaves in the same manner as the smallest non-trivial quantum system, a two-level system, which these days gets referred to as a quantum bit, or *qubit*. Let's compare Dave's state to the state of his classical counterpart, the *bit*. Bits form the building blocks of classical computers, whereas (we will see that) qubits form the building blocks of quantum computers. A bit:

1. admits two states, which we tend to label 0 and 1,
2. can be subjected to any function, and
3. can be freely read.

Here, 'can be subjected to any function' means that we can apply any function on a bit to change its state. For example, we can apply the 'NOT' function to a bit, which interchanges the states 0 and 1, or the 'constant 0' function which sends any state to 0. What we mean by 'can be freely read' is that we can read the state of any bit in a computer's memory without any kind of obstruction and without changing that state.

The fact that we even mention all of this may sound a bit odd...until we compare this to the quantum analogue. A qubit:

1. admits an entire sphere of states,
2. can only be subjected to rotations of the sphere, and
3. can only be accessed by special processes called *quantum measurements*, which only provide limited access, and are moreover extremely invasive.

The set of states a system can occupy is called the *state space* of that system. For classical bits, this state space contains just two states, whereas a qubit can be in infinitely-many states, which we can visualise as a sphere. In the context of quantum theory, this state space is called the *Bloch sphere*. For the sake of explanation, any sphere will do, so we'll just take the Earth. There's plenty of space on Earth for two states of a bit, so put 0 on the North Pole and 1 on the South Pole:

### 1.1 The Penguins and the Polar Bear

3



The particular choice of North Pole/South Pole is not important, but it is important that they are *antipodal* points on the sphere.

Since we can only apply rotations to the sphere of qubit-states, we cannot map both 0 and 1 to 0 (as we could with classical bits), simply because there is no rotation that does that. On the other hand, there are lots of ways to interchange 0 and 1, since there are many (different!) rotations that will turn a sphere upside-down.

So what are quantum measurements? Just like when we read a normal bit, measuring a qubit will produce one of two answers (e.g. 0 or 1, hence the name qubit). However, this act of ‘measuring’ is not quite as innocent as simply reading a bit to get its value. To get a feel for this, we return to Dave. Since qubits can live anywhere in the world, Dave – like one particularly famous (classical) dodo – lives in Oxford:



Now, suppose we wish to ascertain where in the world certain animals live, subject to the following assumptions:

1. we are only allowed to ask whether an animal lives at a specific location on Earth or its antipodal location;
2. all animals can talk and will always answer ‘correctly’; and
3. predatory animals will refrain from eating the questioner.

If we ask a polar bear whether she lives at the North Pole or the South Pole, then she'll say 'the North Pole'. If we ask again, she'll say 'the North Pole' again, because that's just where polar bears are from. Similarly, if we ask a penguin, he'll keep saying 'the South Pole', as long as we keep asking.

On the other hand, what will Dave say if we ask him whether he lives at the North Pole or the South Pole? Now, Dave doesn't really understand the question, but since dodos are a bit thick, he'll give an answer anyway. However, assumption 2 was that all animals will answer correctly. Consequently, as soon as Dave says 'the North Pole', his statement is correct: he actually is at the North Pole!



Now, if we ask him again, he'll say 'the North Pole' again, and he'll keep answering thus until he's eaten by a polar bear (Fig. 1.1). Alternatively, if he had initially said 'the South Pole', he would immediately have been at the South Pole.






Figure 1.1 A polar bear attempting a ‘demolition measurement’ on Dave.

So, no matter what answer Dave gives, his state has changed. The fact that he was originally in Oxford is permanently lost. This phenomenon, known as the *collapse* of the quantum state, happens for almost all questions (i.e. measurements) we might perform. Crucially, this collapse is almost always *non-deterministic*. We almost never know until we measure Dave whether he’ll be at the North Pole or the South Pole. We say ‘almost’, because there is one exception: if we ask whether Dave is in Oxford or the Antipodes Islands, he’ll say ‘Oxford’ and stay put.

While quantum theory cannot predict with certainty the fate of Dave, what it does provide are the *probabilities* for Dave to either collapse to the North Pole or to the South Pole. In this case, quantum theory will tell us that Dave is more likely to go to the North Pole and get eaten by a polar bear than to go to the South Pole and chill with some penguins. The dodo is extinct for a reason after all ...

## 1.2 So What's New?

Almost a century has passed since Dave’s unfortunate travels to the North Pole. In particular, the past two decades have seen a humongous surge in new kinds of research surrounding quantum theory, ranging from re-considering basic concepts (Fig. 1.2) to envisioning radically new technologies. A paradigmatic example is *quantum teleportation*, whereby the non-local features of quantum theory are exploited to send a quantum state across (sometimes) great distances, using nothing but a little bit (actually two little bits ...) of classical communication. Quantum teleportation exposes a delicate interaction between quantum theory and the structure of spacetime at the most fundamental level. At the same time, it is also a template for an important quantum computational model (measurement-based quantum computing), as well as a component in many quantum communication protocols.




Figure 1.2 The paper by Einstein, Podolsky, and Rosen, which was the first to identify quantum non-locality, has enjoyed a huge surge in citations over the past two decades according to Google Scholar, now making it Albert Einstein’s most cited paper. And considering the competition, that’s saying something.

Quantum theory as we now know it – that is to say, its formulation in terms of *Hilbert spaces* – first saw daylight in 1932 with John von Neumann’s book *Mathematische Grundlagen der Quantenmechanik*. On the other hand, quantum teleportation was only discovered in 1992. Hence the question:

*Why did it take 60 years for quantum teleportation to be discovered?*

A first explanation is that within the tradition of physics research during those 60 years, the question of whether something like quantum teleportation would be possible was simply never asked. It only became apparent when researchers stepped outside the existing scientific tradition and asked a seemingly bizarre question:

*What are the information processing features of quantum theory?*

However, one could go a step further and ask why it was even necessary to first pose such a question for teleportation to be discovered. Why wasn’t it plainly obvious that quantum theory allowed for quantum teleportation, in the same way that it is plainly obvious that hammers are capable of hitting nails? Our answer to this question is that the traditional language of Hilbert spaces just isn’t very good at exposing many of the features of quantum theory, and in particular, those features such as teleportation that involve the interaction of multiple systems across time and space. Thus, we pose a new question:

*What is the most appropriate language to reason about quantum theory?*

The answer to this question is what this book is all about. The reader will learn about many important new quantum features that rose to prominence within the emerging fields of quantum computation, quantum information, and quantum technologies, and how these developments went hand-in-hand with a revival of research into the foundations of quantum theory. All of this will be done by using a novel presentation of quantum theory in a purely diagrammatic manner. This not only consists of developing a two-dimensional notation for describing and reasoning about quantum processes, but also of a unique methodology that treats quantum processes, and most importantly *compositions* of processes, as first-class citizens.

### 1.2.1 A New Attitude to Quantum Theory: ‘Features’

Since its inception, many prominent thinkers were deeply unsettled by quantum theory. A great deal of effort and ingenious mathematics in the early twentieth century went

into demonstrating the *bugs* in quantum theory, starting with the now famous EPR paper by Einstein, Podolsky, and Rosen (EPR) in 1935, which claimed that the quantum state provided an ‘incomplete description’ of physical reality. Roughly speaking, they claimed that something must be missing in order to make sense of quantum theory in a manner compatible with our conventional intuitions. However, John Bell showed in 1964 that any attempt to ‘complete’ quantum theory to EPR’s standards was doomed to failure and thereby binned our conventional intuitions as far as quantum theory is concerned. Bell showed that quantum theory contains at its heart a fundamental, irreducible non-locality (Fig. 1.3).

While relativity theory led Einstein to a beautiful and elegant description of the universe in-the-large, quantum theory seemed to muddy the waters. And this more or less characterises how most scientists perceived quantum theory. There were essentially two ways of dealing with this discomfort with ‘quantum weirdness’. One way is to simply ignore any conceptual considerations. This has been the main attitude within the particle physics community, who exemplify the motto ‘shut up and calculate’. Alternatively, one can be obsessively concerned with the conceptual problems surrounding

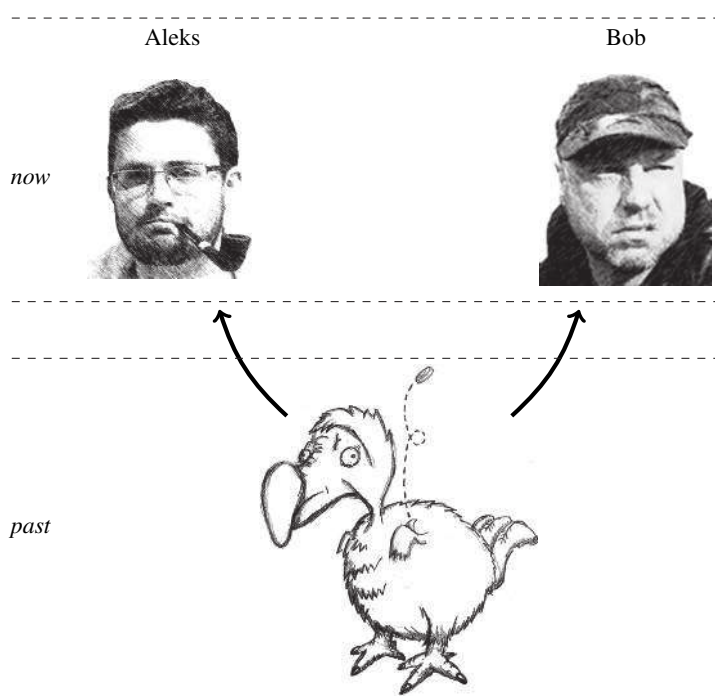


Figure 1.3 Non-locality of quantum theory means that quantum features cannot be explained by means of a classical probabilistic model. In other words, there are situations (unlike the one above) where distantly located observers can experience statistical correlations when they make quantum measurements that cannot be explained by a common cause.




Figure 1.4 Alice queries Dave about Aleks having been partnered with Bob.

quantum theory, sacrificing most of one's life (not to mention sanity) trying to 'fix' them.

Then, starting in the early 1980s, there was an important attitude change, which could be summed up in a simple question:

*What if the purported bugs of quantum theory are actually features?*

In other words, people began to realise that there was much to be gained by embracing quantum theory as it is and trying to figure out how one can actually exploit 'quantum weirdness'. One may even hope that by doing so, we will become more acquainted with quantumness, get more comfortable with its quirkiness, and maybe the resulting less conventional intuitions might even start to make a lot of sense.

And indeed, quantum non-locality, once perceived by Einstein as some unwanted 'spooky action at a distance', suddenly became a key resource. In fact, decades before software developers started using the motto above to excuse their lazy debugging practices ('It's not a bug, it's a feature!'), Richard Feynman had already pointed out that there was at least one thing that quantum systems were really good at: simulating quantum systems! As it turns out, this problem is pretty difficult using a normal, classical computer. Over the next few decades, scientists discovered lots of weird and wonderful things that quantum systems can do: send secure messages, teleport physical systems, and efficiently factor large numbers.

The new focus on quantum features gave birth to several new fields: quantum computing, which studies how quantum systems can be used to compute; quantum information theory, which studies the implications of incorporating quantum phenomena into gathering and sharing information; and quantum technologies, which concerns the actual business of building devices that exploit quantum effects to make our lives better.



```

0 0 1 1 0 1 0 0 1 1 1 0 1 0
0 1 0 1 0 1 0 0 1 0 0 1 0 0
1 0 1 0 0 1 0 1 0 0 0 1 0 0
0 1 1 1 0 0 1 1 1 0 0 0 1 1
1 0 0 0 1 0 0 0 0 1 1 0 1 0
0 1 1 0 0 1 1 1 0 1 0 0 1 0
0 1 1 1 0 1 0 0 1 0 1 0 1 0
1 1 0 1 0 0 0 0 1 0 1 0 1 0
0 1 0 1 0 1 0 0 1 0 1 0 1 0
0 1 0 0 1 1 0 1 0 0 1 1 0 1
1 0 0 1 1 0 1 0 1 0 1 0 0 1 1
0 1 0 0 1 0 0 1 1 0 0 0 1 1
1 1 0 0 0 1 1 0 1 0 1 1 1 0
0 1 0 1 1 1 0 1 1 0 1 0 0 0

```

vs.




Figure 1.5 Contrasting a low-level and a high-level representation of the digital data that one may find within a computational device.

### 1.2.2 A New Form of Mathematics: ‘Diagrams’

It should be emphasised that discovering these new quantum features wasn’t trivial and involved some very smart people. Our bold claim is that when one adopts the appropriate language for quantum theory, these features jump right off the page. Conversely, the traditional, Hilbert space–based language of quantum theory forms a major obstruction to discovering such features. To give some idea of why this is the case, we will make use of some simple metaphors.

Imagine that you were trying to determine what was happening in a video just by looking at its digital encoding (Fig. 1.5). Obviously this is a more or less impossible task. While digital data, i.e. strings of 0s and 1s, is the workhorse of digital technology, and while it is possible to understand ‘in principle’ how they encode all of the media stored on your hard drive, asking a person to decode a particular string of binary by hand is more suitable for punishing greedy bankers and corrupt politicians than solving interesting problems.

Of course, even skilled computer programmers wouldn’t be expected to interact directly with binary data. Somewhere along the way to modern computer programming came the advent of assembly language, which gives a (somewhat) human-readable translation for individual instructions sent to a computer processor. While this made it more practical to write programs to drive computers, it still takes a lot of head-scratching to figure out what any particular piece of assembly code does. Using *low-level languages* such as assembly language creates an artificial barrier between programs and the concepts that they represent and places practical limits on the complexity of problems those programs can solve. For this reason, virtually every programmer today uses *high-level languages* in their day-to-day work (Fig. 1.6).

Similarly, ‘detecting new quantum features’ in terms of the traditional (i.e. low-level) language for quantum theory, namely ‘strings of complex numbers’ (rather than ‘strings of 0s and 1s’), isn’t that easy either. This could explain why it took six highly esteemed researchers to discover quantum teleportation, some 60 years since the actual birth of the quantum theoretical formalism. By contrast, the diagrammatic language we use in this book is a *high-level* language for exploring quantum features (Fig. 1.7). We will soon see

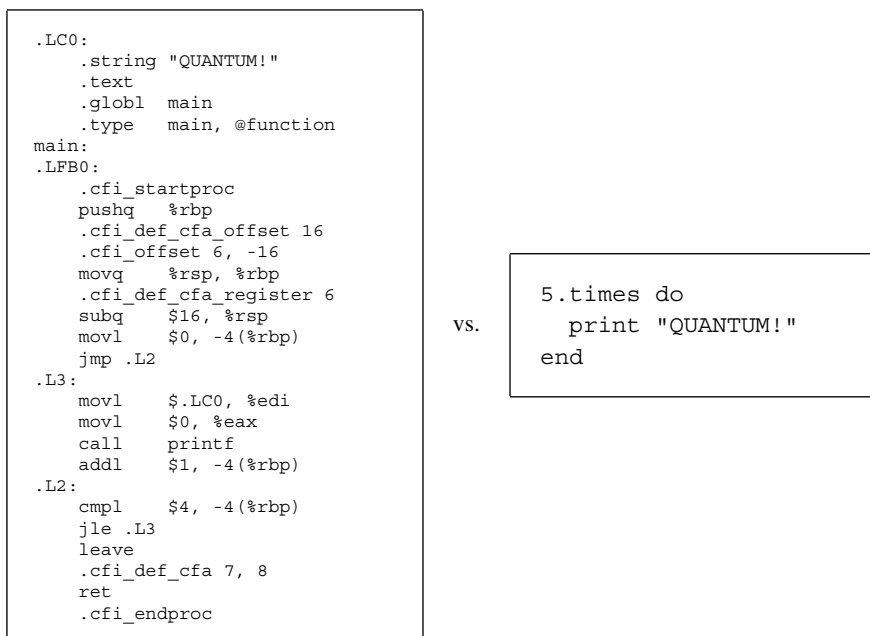


Figure 1.6 Contrasting a low-level and a high-level language for computer programs. The programs on the left and right perform the same task, but one is written in the low-level x86 assembly language and one in the high-level language Ruby.




Figure 1.7 Contrasting a low-level and a high-level language for quantum processes, just like we contrasted the low-level and a high-level representation for digital data in Fig. 1.5 and a low-level and a high-level programming language in Fig. 1.6.

that by embracing the diagrammatic language for quantum theory, features like quantum teleportation are pretty much staring you in the face!

Although it goes beyond the scope of this book, it is worth mentioning that the diagrammatic language we use has found applications in other areas as well, such as modelling meaning in natural language (Fig. 1.8), doing proofs in formal logic, control theory, and modelling electrical circuits.

Diagrams are also becoming increasingly important in some fancy research areas of pure mathematics, such as knot theory, representation theory, and algebraic topology. By using diagrams we eliminate a huge amount of redundant syntactic garbage in representing