

# The Philosophy of Quantum Computing\*

Michael E. Cuffaro<sup>a</sup>

<sup>a</sup>Munich Center for Mathematical Philosophy, Ludwig-Maximilians-Universität München

*To appear in the book *Quantum Computing in the Arts and Humanities: An Introduction to Core Concepts, Theory and Applications*. E. R. Miranda (Ed.). Cham: Springer Nature, 202x. Pre-publication draft v.16 Mar 2021.*

## 1 Introduction

From the philosopher’s perspective, the interest in quantum computation stems primarily from the way that it combines fundamental concepts from two distinct sciences: physics (especially quantum mechanics) and computer science, each long a subject of philosophical speculation and analysis in its own right. Quantum computing combines both of these more traditional areas of inquiry into one wholly new (if not quite independent) science. There are philosophical questions that arise from this merger, and philosophical lessons to be learned. Over the course of this chapter we will be discussing what I take to be some of the most important.<sup>1</sup>

We begin, in Section 2, by introducing the fundamental concepts from physics and computation that will be essential for framing the further philosophical discussion that will follow. Section 2.1 and Section 2.2 introduce concepts from classical mechanics and the classical theory of computation, respectively. In Section 2.1 we discuss the concept of the state of a physical system as it is given in classical mechanics. We emphasise in particular

---

\*My work on this chapter benefited significantly from my interactions with students and other audience members during and after a series of lectures I gave at the University of Urbino’s twenty-third international summer school in philosophy of physics, held online in June 2020 in the midst of the first wave of the COVID-19 pandemic, as well as a further lecture I gave for Michel Janssen’s “The Age of Entanglement” honours seminar at the University of Minnesota in December 2020 as the second wave of the pandemic began in earnest. Thanks to Ari Duwell, Eduardo Miranda, Philip Papayannopoulos, and Lev Vaidman for comments on a previous draft of this chapter. I am also grateful for informal discussions, over the years, with Guido Bacciagaluppi, Jim Baggot, Michel Janssen, Christoph Lehner, Lev Vaidman, and David Wallace; my presentation of the Everett interpretation in Section 3, in particular, is significantly informed by what I take myself to have learned from these discussions, though I hold only myself responsible for any mistakes or misunderstandings in my presentation of the Everettian view. Section 2, on “Fundamental concepts” is heavily informed by my recent work on related topics with Stephan Hartmann, Michael Janas, Michel Janssen, and Markus Müller, as well as by informal correspondence with Jeffrey Bub and (the late) Bill Demopoulos; though here again I take sole responsibility for any mistakes. Finally, I gratefully acknowledge the generous financial support of the *Alexander von Humboldt Stiftung*.

<sup>1</sup>Space does not permit me to exhaustively survey all of the philosophical issues brought up by quantum computing. The interested reader can find a summary of other important issues in Hagar & Cuffaro (2019).

the way that we are invited, in classical mechanics, to think of the state of a physical system as a compact description of the properties possessed by it. These properties determine, in advance, the answers to the experimental questions that we can pose of a system. And considering these questions and their answers leads one to an alternative representation of a system's state that is useful for representing computation, a subject we then take up in more detail in Section 2.2.

In Section 2.2 we introduce the concept of a model of computation (or computational architecture), the concept of the cost of carrying out a computation under a given computational model, and explain what it means to solve a problem just as easily under one computational model as under another. After discussing some of the more important computational complexity classes, we then introduce two theses that can be used to ground the model-independence of computational cost: the so-called universality of Turing efficiency thesis (sometimes also referred to as the 'strong', or 'physical', or 'extended' Church-Turing thesis), and the invariance thesis. These theses are both called into question by the existence of quantum computers. And since some have taken the absolute model-independence guaranteed by them to be foundational for the science of computational complexity theory, the question arises of what to make of those foundations in light of the existence of quantum computation. We will discuss this question, in an abstract way, in Section 2.3, where I will argue that although the universality thesis must indeed be given up, this is not a great loss. The invariance thesis, by contrast, remains, but in a different form, transformed from a metaphysical to a methodological principle.

In Section 2.4 we turn to the physics behind quantum computation, and begin by introducing some of the more important concepts of quantum theory. Most importantly we introduce the concept of a quantum state. We emphasise that the way that a quantum state determines what the answers to experimental questions will be is fundamentally different than the way that they are determined by a classical state. We then turn to quantum computation proper in Section 3, where we review the basic concepts of quantum computing and consider what we can say regarding the physical explanation of the power of quantum computers. The many-worlds explanation of quantum computing—the idea that quantum computers outperform classical computers by running their computations in exponentially many physical universes—is then introduced. We note two major problems that arise for this interpretation. The first arises from the so-called preferred basis problem. This problem is a challenge (that is arguably surmountable) for the many-worlds view in the more general context of quantum mechanics. But we will see that it is especially problematic in the context of quantum computers. The second major problem arises from the fact that there are many different models of quantum computation, but the many-worlds explanation of quantum computing only seems motivated by one of them.

In Section 4 we consider the role of quantum entanglement in enabling quantum computers to outperform classical computers. We summarise an unsuccessful argument to the conclusion that quantum entanglement is insufficient to enable this quantum 'speedup' in Section 5, noting in Section 6 that reflecting on what it means to provide a classical computer simulation of a quantum phenomenon should convince us to reach the opposite conclusion. We continue the discussion of classically simulating quantum computers in Section 7, and reflect on general aspects of the computational perspective that the study of quantum computing provides on physical theory. We note that reflecting on quantum computing emphasises that there are important differences in the methodological presuppositions that lie at the basis of physics and computer science, respectively, and that conflating these can lead to confusion. We also note the emphasis that studying quantum computation places

on the fact that quantum mechanics and classical mechanics are each alternative universal languages for describing physical systems, and that the difference between quantum mechanics and classical mechanics lies, fundamentally, in the differences in the expressive power of these languages.<sup>2</sup> We reflect on this in Section 8.

## 2 Fundamental concepts

### 2.1 Classical states

Classical mechanics (in its various forms<sup>3</sup>) was, prior to quantum mechanics, our fundamental theoretical framework for describing the dynamics of physical systems (i.e., how physical systems change over time). What is meant by a physical system here is just one of the concrete objects that a particular physical theory describes. All such systems are described as evolving through time in accordance with the dynamical constraints that a theoretical framework applies universally to every physical system. In classical mechanics (and the same is true, as we will see later, in quantum mechanics) a physical system can be anything from a single particle to (in principle) the entire physical universe. Mathematically, though, what the dynamical laws of a theory actually govern are the relations between the possible state descriptions of systems. By the state description (or state specification) of a system is meant a description of the particular physical situation that it happens to be in at a given moment in time. For instance, at a given moment in time, a classical-mechanical system will have a particular kinetic energy, it will be accelerating (or not) in a particular direction, and so on.

It turns out that we do not have to explicitly specify each and every one of a system's dynamical properties to exactly and exhaustively specify the system's dynamics. In classical mechanics, given the dynamical laws, it is enough to specify values for a system's position and momentum. The value of any other dynamical property can then be calculated by relating these with each other and with the values of non-dynamical properties of the system such as, for instance, its charge or its mass.

Figure 1 illustrates the two different strategies described in the previous paragraph for specifying the state of a physical system, for instance a system composed of a particle attached to the end of a spring constrained to move in one direction (see Hughes, 1989, p. 73). On the left we explicitly specify values ( $v_1, v_2, \dots$ ) for each and every one of the physical system's dynamical parameters. On the right the values of momentum,  $p$ , and position,  $q$ , are specified, and all other dynamical quantities are calculated on their basis. In particular the total energy,  $H$ , of the system is defined as the sum of the kinetic energy,  $T$ , of the particle, and the potential energy,  $V$ , stored in the spring, which are in turn defined in terms of  $p$  and  $q$ , respectively. Note that  $m$ , the system's mass, is a *non*-dynamical parameter that does not change over the history of the system, and  $k$  is the force (the first derivative of momentum) per unit distance required for the spring to be displaced from its equilibrium point (i.e., the point at which it is neither stretched nor compressed). Other forces (e.g.,

---

<sup>2</sup>By 'quantum mechanics' I mean the fundamental theoretical framework shared in common by every specific quantum-mechanical theory (quantum field theories, for instance) of a particular class of systems; see Aaronson (2013b, pp. 110–111), Janas, Cuffaro, & Janssen (Forthcoming, ch. 1 and §6.3), Nielsen & Chuang (2000, p. 2), and Wallace (2019).

<sup>3</sup>These include Newtonian, Lagrangian, Hamiltonian, relativistic, and classical statistical mechanics. For a recent comparison and philosophical discussion of Lagrangian and Hamiltonian mechanics, see Curiel (2014).

$$\begin{array}{l|l}
 P_1 = v_1 & p = v_p \\
 P_2 = v_2 & q = v_q \\
 P_3 = v_3 & T = p^2/2m \\
 P_4 = v_4 & V = kq^2/2 \\
 P_5 = v_5 & H = V + T \\
 \dots & \dots
 \end{array}$$

**Figure 1:** Two different strategies for specifying the state of a physical system, for instance a system composed of a particle attached to the end of a spring constrained to move in one direction (see Hughes, 1989, p. 73). On the left we explicitly specify values ( $v_1, v_2, \dots$ ) for each and every one of the physical system’s dynamical properties. On the right the values of momentum,  $p$ , and position,  $q$ , are specified, and all other dynamical quantities are calculated on their basis.

gravitational forces) are neglected but in principle can be included (for further discussion, see Hughes, 1989, §3.2.5).

The upshot is that specifying the position and momentum of a system provides us with enough information, in classical mechanics, to completely characterise all of the other dynamical properties of that system. Accordingly, the dynamical state,  $\omega_t$ , of a system at some particular time  $t$  is given in classical mechanics by:

$$\omega_t = (\mathbf{q}_t, \mathbf{p}_t), \tag{2.1}$$

where  $\mathbf{q}_t$  and  $\mathbf{p}_t$  are vectors (in three dimensions) representing the system’s position and momentum, respectively, at  $t$ . Further, we can infer from the state of a system at  $t$ , and the classical-mechanical laws of motion, exactly what the state of the system will be in the next moment and at every other time both in the system’s past and in its future (Hughes, 1989, §2.6).

Classical-mechanical states have another feature. Imagine all of the possible (experimental) yes-or-no questions one might want to ask about the dynamical properties of a particular system at a particular time, questions like: *Is the value of the dynamical property A within the range  $\Delta$ ?* Completely specifying a system’s dynamical state, i.e., specifying its momentum,  $\mathbf{p}$ , and its position,  $\mathbf{q}$ , yields a simultaneous answer to all such questions irrespective of whether any question has actually been asked. Indeed, this is what actually justifies our having considered the values of  $\mathbf{p}$ ,  $\mathbf{q}$ , and the other quantities derived from them to be (observable) *properties* of the system in the first place, properties possessed by the system whether or not, and however, we enquire concerning them (Janas et al., Forthcoming, §6.3). The same goes for a system composed of many parts; for after all, any observable property  $A_1$  of a subsystem  $S_1$  of some larger system  $S$  is also an observable property of  $S$ . Thus, given the state specification for a particular system we can construct a sort of ‘truth table’ which could be used, in principle, to answer any conceivable yes-or-no question concerning the system and its subsystems (see Figure 2).

In Figure 2, I wrote  $Y$  and  $N$  (for “yes” and “no”), but I could have equally well used  $T$  (for true) and  $F$  (for false), or alternately the binary digits 0 and 1. Using binary digits, i.e., ‘bits’, is convenient because they are useful for representing numbers, which can be manipulated abstractly using logico-mathematical operations.<sup>4</sup> In particular, if we build physical systems whose states can reliably be used to represent binary numbers (in the

---

<sup>4</sup>I say ‘logico-mathematical’ because logical operations on bits can be thought of as modulo-2 arithmetical operations (see Boole, 1847).

$\mathbf{p}_1$	$\mathbf{q}_1$	$\mathbf{p}_2$	$\mathbf{q}_2$	$A$	$B$	...
$v_{p_1}^1$	$v_{q_1}^1$	$v_{p_2}^1$	$v_{q_2}^1$	$N$	$N$	
$v_{p_1}^2$	$v_{q_1}^2$	$v_{p_2}^2$	$v_{q_2}^2$	$N$	$Y$	
$v_{p_1}^3$	$v_{q_1}^3$	$v_{p_2}^3$	$v_{q_2}^3$	$N$	$Y$	
$v_{p_1}^4$	$v_{q_1}^4$	$v_{p_2}^4$	$v_{q_2}^4$	$Y$	$N$	
$v_{p_1}^5$	$v_{q_1}^5$	$v_{p_2}^5$	$v_{q_2}^5$	$N$	$Y$	
$v_{p_1}^6$	$v_{q_1}^6$	$v_{p_2}^6$	$v_{q_2}^6$	$Y$	$N$	
$v_{p_1}^7$	$v_{q_1}^7$	$v_{p_2}^7$	$v_{q_2}^7$	$Y$	$N$	
$v_{p_1}^8$	$v_{q_1}^8$	$v_{p_2}^8$	$v_{q_2}^8$	$N$	$N$	
...						

**Figure 2:** A ‘truth table’ for a classical system composed of two subsystems whose states, respectively, are  $(\mathbf{p}_1, \mathbf{q}_1)$  and  $(\mathbf{p}_2, \mathbf{q}_2)$ . Various combinations of values for the state parameters  $\mathbf{p}_i$  and  $\mathbf{q}_i$  are given on the left of the double-vertical line, with the superscript  $j$  in  $v^j$  indicating the  $j^{\text{th}}$  combination. Relative to a given combination of values, the answers to experimental questions concerning the values of derived observable quantities  $A, B, \dots$  on the right of the double-vertical line can be determined.




**Figure 3:** A physical system consisting of four light switches. For each switch, we can ask: “Is switch  $S$  turned on?” Using 0 to represent yes and 1 to represent no, the state of the overall system relative to this question can be represented using the bit-string 0110 (which, in base-ten, is the number 6). This is just one of the  $2^4 = 16$  possible states that a system like this one can be in. More generally, for a system made up of  $n$  2-dimensional subsystems, the number of possible states for the system as a whole is  $2^n$ .

sense that yes-or-no questions concerning their observable properties  $A, B, \dots$  can be laid out as in Figure 2, replacing 0 for  $Y$  and 1 for  $N$ ), and reliably evolve them in ways that mirror a small basic set of logico-mathematical operations, then we can (by combining these operations) use physical systems to carry out computations that can in principle be arbitrarily complex (see Figures 3–4).<sup>5</sup>

## 2.2 Classical computers

When the word ‘computer’ is used in popular discourse today, generally what is meant is a physical system like the one depicted in Figure 3—not, of course, literally a collection of light switches, but a physical system composed of subsystems that, like the switches, can reliably be interpreted as being either on or off, and that can be organised into registers, random access memory, hard drives, and other peripheral systems such as keyboards, ‘mice’, touchpads, and so on, that can be manipulated to reliably realise various logical operations like the ones depicted in Figure 4. The laptop I am writing this chapter with is an example of such a device. We turn such devices on when we need them, and turn them off or ‘put them

<sup>5</sup>For more general accounts of how physical systems can be used to represent computations, see Fletcher (2018); Horsman, Kendon, & Stepany (2018); Maroney & Timpson (2018).




**Figure 4:** Some of the logico-mathematical operations, or ‘logic gates’, that can be used to manipulate bits. OR and NOT together constitute a universal set of gates, in the sense that any more complicated logico-mathematical operation, such as the ‘if-then’ gate at the right, can be constructed using combinations of OR and NOT. There are other universal sets besides {OR, NOT}, for instance: {AND, NOT}, {NOR}, and {NAND}, where NAND is the ‘not-and’ gate and NOR is the ‘not-or’ gate.

to sleep’ when we do not. But the word ‘computer’ did not always signify a kind of machine. Prior to the 1930s, ‘computation’ generally meant an activity carried out by a human being. Even after Alan Turing’s introduction of what later came to be known as the ‘Turing machine’ in 1936—an abstract device that is itself modelled on human computation—‘computers’ were generally understood to be human beings until well into the 1960s.<sup>6</sup>

As for the Turing machine: this, as mentioned, is also not a concrete physical device (although it could, in principle, be instantiated by one) but an abstract mathematical model. Turing’s primary reason for introducing it had been to address the so-called *Entscheidungsproblem* (which is the German word for ‘decision problem’), an important but unsettled question that had arisen in the foundations of number theory. The *Entscheidungsproblem* concerned the existence of an effective procedure for deciding whether an arbitrarily given expression in first-order logic can be proven from the axioms of first-order logic. The answer, as Turing and, independently, Alonzo Church were able to demonstrate, is that there is no such procedure. Turing’s own proof relied on a penetrating philosophical analysis of the notion of effective (human) computation, and a corresponding argument that one could design an automatic machine (now called a *Turing machine*; see Figure 5) to carry out the essential activities of a human computer.

Turing (1936-7, pp. 249–51) argued that, for a computer to carry out a computation, it is essential that she have access to a notebook from which she can read and onto which she can write various symbols in the course of her work. These symbols need to be distinguishable by her from one another, on the basis of which Turing argued that the alphabet from which she chooses her symbols must be a finite alphabet. At any given moment during a computation, a computer may find herself in one of a finite number (again, because she must be able to distinguish them) of states of mind relevant to the task at hand which summarise her memory of the actions she has performed up until that point along with her awareness of what she must now do (pp. 253–4). The actions that are available to her are characterised by a finite set of elementary operations, such as ‘read the next symbol’ from the notebook, ‘write symbol *a*’ to the notebook, and so on. In an automatic machine, a *control unit* is constructed to encode the machine’s ‘state of mind’ (i.e., a logical representation of its current state, its current input, and its transition function), which in general changes after every operation of the *read-write head*. The latter moves back and forth along

<sup>6</sup>The role of human computers in the United States of America’s space program, for instance, has been documented in Shetterly (2016).



**Figure 5:** A version of what is now called a ‘Turing machine’, with a single bi-directional tape, a control unit encoding a particular set of instructions or program, and a read-write head. Such a machine is an idealised representation of the components involved in human computation. In a *universal* Turing machine, the computer’s program is read from the tape at the beginning of the computation. For a good introduction to the Turing machine and other related topics, see Martin (1997).

a *one-dimensional tape* (the machine’s ‘notebook’), from which it reads and onto which it writes various symbols from a finite alphabet, in conformity with a particular finite set of exactly specifiable rules. Turing was able to show that no automatic machine of this kind can be used to solve the *Entscheidungsproblem*.

Through the work of Turing, Church, and others, the modern science of computation, and in particular *computability theory*—the science of which problems can and which cannot be solved by a computer, i.e., by any thing or any person that can carry out (but is restricted to) the essential activities associated with human computation—was born.<sup>7</sup>

Irrespective of whether it is a machine or a human being that is doing the computing, the question of how much it actually costs to compute a solution to a computable problem is also a very important one. In particular we generally would like to know how much *time* and how much *space* are required (though time is usually regarded as the more important measure). The question of the cost of a given computation had become especially important as more and more problems came to be carried out by machines in the 1950s and 1960s. Even the earliest digital computers of the 1930s and 1940s performed their computations far faster than human beings could possibly perform them, and with improvements in design and in technology, machines could be made to run ever faster. Through the need to use these machines’ resources efficiently, a sub-field of computer science began to develop whose concern was with how efficiently a given problem can be solved and on which types of machine.

In the context of the Turing machine model, time and space are quantified in terms of the number of computational ‘steps’ (i.e., the number of basic operations) and the number of squares of tape, respectively, that are needed to carry out a given computation. Although it is in principle possible (assuming we have enough tape) to instantiate a Turing machine physically,<sup>8</sup> generally it makes more sense, from a practical perspective, to design a mechan-

<sup>7</sup>For more on the *Entscheidungsproblem* and the early history of computer science, see Copeland (2020); Davis (2000); Dawson Jr. (2007); Lupacchini (2018).

<sup>8</sup>A Turing machine’s tape does not need to be actually infinite in length. What is required is only that the tape be *indefinitely* long, so that, for a given (finite) computation, the machine can be supplied with enough tape to carry it out. To put it another way: What constitutes ‘enough tape’ to carry out a computation is not part of the general definition of a Turing machine. It is, rather, assumed, for the purposes of that definition, that enough tape to carry out a given computation can be supplied. That said, certain variations on the Turing

ical computer in accordance with some other model of computation. Particularly important is the von Neumann architecture, a type of stored-program computer,<sup>9</sup> that involves a central processing unit (CPU), an instruction register and program counter, internal memory used to store data (including the program being run), external memory used to store less frequently needed data, and finally an input and an output mechanism. A von Neumann machine can in principle compute all of the same problems as a universal Turing machine (see Figure 5), but a time step for a von Neumann machine is not concretely interpreted in the same way as it is for a Turing machine, just as ‘space’ for a von Neumann machine is not measured in squares of tape.

There are many other universal models of computation besides von Neumann machines and Turing machines. Nevertheless, at least when it comes to (reasonable) classical computational models, it is possible to talk about the computational cost of solving a particular problem (like, for instance, the problem of finding the prime factors of a given integer), in a model-independent way; i.e., without having to specify which machine model that cost is defined with respect to. There are two things involved in explicating what it means for the computational cost of solving a given problem to be model-independent in this sense. The first, mentioned just now in passing, is the idea of a reasonable machine architecture.

Roughly, a *reasonable* (universal) machine architecture or model is one that it would be possible to physically instantiate. Turing machines like the one described in Figure 5 and von Neumann machines are examples of reasonable machine models. A computational model that employs an unbounded amount of parallel processing, by contrast, is an example of an unreasonable model, or what Peter van Emde Boas (1990, p. 5) calls a model of the ‘second machine class’. There are legitimate theoretical reasons for studying the complexity-theoretic properties of unreasonable models. But no finite physical system could possibly instantiate an unreasonable architecture. There are thus good reasons to want to distinguish the reasonable from the unreasonable ones (see also Dean, 2016).

The second thing we need to do, in order explicate what it means for the computational cost of solving a given problem to be model-independent, is to make sense of what it means to solve a given problem *just as easily* with one computational architecture as with some other one. But before we can make sense of solving something just as easily we need to first try and make sense of what it means to solve something *easily* under one computational model. Consider, to begin with, a human computer working alone. For a human computer, the difference between a problem that could in principle require up to a thousand steps to solve, and one that could in principle require up to a million steps, is very great indeed. With some justification this person might call the former problem easy to solve, and the latter problem hard. Any criterion based on a fixed number of computational steps, however, would not be very useful in the context of machines, at any rate certainly not any criterion based on numbers this low. Even by the 1950s a mechanical computer could carry out a computation far faster than a human being. As of 2020, a typical laptop computer can perform hundreds of thousands of Millions of Instructions per Second (MIPS) (Wikipedia contributors, 2020), and computers are becoming more and more performant every year. In the context of machine computation, it is far more useful to ask how a given solution

---

machine model restrict the ways in which tape can be read by the control unit in various ways. For instance some variants employ separate tape(s) for the machine to write “rough work” on in addition to an output tape, some variants only allow the read-write head to move in one direction along the tape, and so on.

<sup>9</sup>Although the von Neumann architecture, or ‘von Neumann machine’ is only one of a number of various types of stored-program computer, the terms have today (inappropriately, from a historical point of view) come to be understood synonymously (Copeland, 2017).



to a computational problem *scales*; i.e., how the number of resources required to carry out a given solution to that problem grows with the size of the input to the problem. Take, for instance, the problem to sort a given set of numbers. All else equal, the more numbers there are to sort, the longer it will take. How much longer? The fastest sorting algorithms (one of which is MergeSort) will take on the order of  $n \log n$  steps to sort  $n$  numbers in the worst case (Mehlhorn & Sanders, 2008).

In modern computational complexity theory, we say that a problem is *easy* to solve, or alternately that it is *tractable*, if there exists an efficient algorithm to solve it. If no efficient algorithm exists, then we say that the problem is *hard*, or *intractable*. This definition, of course, merely trades in one informal notion for another, so let us state more precisely what we mean by ‘efficient’. By an *efficient algorithm* for solving a given problem, we mean an algorithm that will solve it in *polynomial time*; i.e., in a number of time steps that is no more than a polynomial function of the size of its input. Stated in mathematical terms, this means that for input size  $n$ , such an algorithm will take on the order of  $n^k$  steps, where  $k$  is some constant.<sup>10</sup> As for an algorithm that takes more than a polynomial number of steps, we consider it to be *inefficient*, even if it might be the most efficient algorithm known for a particular problem. An example of an inefficient algorithm is one that requires *exponential time*; i.e., on the order of  $k^n$  steps for input size  $n$  and some constant  $k$ .<sup>11</sup> Even harder problems take *factorial time*, i.e., on the order of  $n!$  steps, and so on.

The idea of using polynomial time as our criterion for distinguishing between efficient and inefficient algorithms—what we will henceforth call the *polynomial principle*—was introduced independently by Alan Cobham and Jack Edmonds in 1965.<sup>12</sup> There are a number of good reasons for adopting the polynomial principle. One reason is simply that the set of problems picked out by it has so far tended to correspond with those that we would (so to speak) pre-theoretically regard as efficiently solvable (Cuffaro, 2018b, §11.6).<sup>13</sup> Perhaps the most useful feature of the polynomial principle stems from the fact that polynomial functions compose (see Figure 6). That is, given an algorithm that runs in polynomial time, if we add to it any number of calls to polynomial-time subroutines, the total running time of the algorithm will still be polynomial (Arora & Barak, 2009, p. 27).<sup>14</sup>

Now that we have a handle on what it means for a problem to be solvable easily, we can state what it means for a problem to be solvable just as easily on one type of computational architecture as on another. We say that a problem is solvable *just as easily* on machine model  $\mathfrak{M}_1$  as on machine model  $\mathfrak{M}_2$  if and only if there is an algorithm for solving the problem on  $\mathfrak{M}_1$  that requires no more than a polynomial number of extra time steps as compared to  $\mathfrak{M}_2$ .<sup>15</sup> Another way of putting this is that anything that  $\mathfrak{M}_2$  can do is *efficiently simulable*

---

<sup>10</sup>On the order of’ is a technical term, usually symbolised in “big-oh notation” as  $O(T(n))$ . An algorithm is  $O(T(n))$  for some function  $T(n)$  if for every sufficiently large  $n$ , its actual running time  $t(n) \leq c \cdot T(n)$  for some constant  $c$ . For instance, an algorithm that never takes more than  $5n^3$  steps is  $O(n^3)$ .

<sup>11</sup>A famous example of a problem for which only exponential-time solutions are known is the Travelling Salesman Problem (Cook, 2012).

<sup>12</sup>In some literature it is referred to as the *Cobham-Edmonds thesis*. Kurt Gödel anticipated the principle, to some extent, in a private letter he wrote to John von Neumann in 1956. For further discussion, see Cuffaro (2018b).

<sup>13</sup>This correspondence is not perfect, but the usefulness of the polynomial principle is such that we appeal to it despite this (Cuffaro, 2018b, §11.6).

<sup>14</sup>It is easy to see this: Consider a program that consists of  $n^k$  calls of a subroutine that takes  $n^l$  steps, where  $n$  is the number of bits used to represent the input, and  $k$  and  $l$  are finite constants. The total number of steps needed to carry out this program will be  $n^{k+l}$ . If  $k$  and  $l$  are finite constants then so is  $k+l$ . In other words,  $n^{k+l}$  is still a polynomial.

<sup>15</sup>Note that it makes sense to talk about solving a given problem just as easily on  $\mathfrak{M}_1$  as on  $\mathfrak{M}_2$  even when

```

// A polynomial-time algorithm:
myProgram1(arg1, arg2) {
    statement1;
    statement2;
    ...
}

// Another polynomial-time algorithm:
myProgram2(arg1, arg2) {
    statement1;
    statement2;
    ...
    myPolynomialTimeSubroutine(arg1);
    ...
    myPolynomialTimeSubroutine(arg2);
}

```

**Figure 6:** Inserting any number of calls to polynomial-time subroutines within the body of a polynomial-time algorithm results in a polynomial-time algorithm.

by  $\mathfrak{M}_1$ .

Any problem that is solvable in polynomial time by a Turing machine is said to belong to the class **PTIME** (for ‘polynomial time’), though it is usually referred to simply as **P**. **P** is, of course, not the only complexity class. If a given problem is such that, for a given solution to it, there is a Turing machine that will *verify* that the solution is indeed a solution, then the problem is said to belong to the class **NP**, for ‘nondeterministic polynomial time’. The reason for the name is that, equivalently, this class of problems can be defined as those that can be solved in polynomial time on a nondeterministic Turing machine. A *nondeterministic Turing machine*, or *choice machine*, is such that at a given step of a given computation, an imagined external operator of the machine can choose to have it transition in one way rather than another. This is unlike a standard (deterministic) Turing machine for which every computational step is completely determined in advance given a particular input. Choice machines are interesting from a theoretical point of view, but given that they involve an external operator they are not really automatic machines in the sense that is relevant to our discussion here.

A variant of the choice machine that *is* automatic is a *probabilistic Turing machine*. Such a machine’s choices are not made by an external operator but by the machine itself, which we can imagine as equipped with the equivalent of a fair coin that it can flip a number of times to decide whether to transition one way or another whenever such a choice is open to it.<sup>16</sup> The class of problems solvable in polynomial time on a probabilistic Turing machine is called **BPP** (for ‘bounded error probabilistic polynomial’). Note that what it means to be able to solve a given problem is not the same for a probabilistic Turing machine as it is for a deterministic Turing machine. In particular a given ‘solution’ output by a probabilistic Turing machine is allowed to be wrong. We only demand that it be right with high enough probability so that, if we re-run the machine for the given input on the order of a polynomial number of further times, our confidence in the majority answer will approach certainty.

There are very many more complexity classes besides these ones (see Aaronson, 2016). But we will stop here as we now have the basic ingredients with which to state a couple of related theses that can be used to explicate what it means for the cost of solving a given problem to be model-independent. As we will see a little later, there are good reasons to prefer the second thesis, but for now we simply introduce both. First, the *universality of Turing efficiency thesis* (sometimes also called the ‘strong’, ‘physical’, or ‘extended’ Church-

---

the problem under consideration is actually intractable for both. For instance if some problem requires  $2^n$  steps to solve on  $\mathfrak{M}_1$  and  $2^n + n^3$  steps to solve on  $\mathfrak{M}_2$  then it is no harder, from the point of view of the polynomial principle, to solve it on  $\mathfrak{M}_2$  than on  $\mathfrak{M}_1$ .

<sup>16</sup>For more on probabilistic and nondeterministic Turing machines and how they compare to their deterministic counterparts, see Cuffaro (2018b, §11.3).

Turing thesis<sup>17</sup>) asserts that any problem that can be efficiently solved on *any* reasonable machine model  $\mathfrak{M}$  (von Neumann architecture, Harvard architecture, or whatever) can be efficiently solved on a probabilistic Turing machine, or more formally:

$$\bigcup \text{Poly}_{\mathfrak{M}} = \mathbf{BPP}. \quad (2.2)$$

In other words, the thesis implies that the set of problems solvable in polynomial time does not grow beyond **BPP** if we allow ourselves to vary the underlying model. Assuming the thesis is true, we do not need to worry about what model an algorithm is implemented on when discussing the computational complexity of various problems; we can simply focus on the abstract probabilistic Turing machine model and carry out our analyses in relation to it.<sup>18</sup>

The second thesis we will introduce is called the *invariance thesis* (van Emde Boas, 1990, p. 5),<sup>19</sup> which asserts that given any two reasonable machine models  $\mathfrak{M}_i$  and  $\mathfrak{M}_j$ ,  $\mathfrak{M}_i$  can efficiently simulate  $\mathfrak{M}_j$ ; i.e.  $\mathfrak{M}_i$  can solve any problem just as easily (in the sense explained above) as  $\mathfrak{M}_j$  can. More formally:

$$\forall_{i,j} \mathfrak{M}_i \stackrel{\text{poly}}{\sim} \mathfrak{M}_j. \quad (2.3)$$

Note that the invariance thesis implies the universality thesis, but not vice versa.

### 2.3 Physical perspectives on computer science

Neither the universality of Turing efficiency thesis, nor the invariance thesis, is a mathematical theorem. These statements can be true or false, and for a long time they were thought to be true, for none of the reasonable (classical) universal models of computation that had been developed since the time of Turing were found to be more efficient than the Turing model by more than a polynomial factor, and all of them had been shown to be able to simulate one another efficiently (van Emde Boas, 1990). Over the last three decades, however, evidence has been mounting against universality and invariance, primarily as a result of the advent of quantum computing (Aaronson, 2013b, chs. 10, 15).

Quantum mechanics, as we will see in the next section, is an irreducibly probabilistic theory. Thus a quantum computer is a type of probabilistic machine. Analogously to the way we defined **BPP** as the class of problems solvable (in the probabilistic sense) in polynomial time on a probabilistic Turing machine, we can also define the complexity class **BQP** (for ‘bounded error quantum polynomial’) as the class of problems solvable (in the same sense)

---

<sup>17</sup>See Timpson (2013, ch. 6) for discussion of a different, more general thesis, that is only indirectly relevant to computational complexity theory. For a discussion of how these theses relate, see Cuffaro (2018b, §11.4).

<sup>18</sup>We could have also expressed the thesis in terms of **P** rather than **BPP**. Although it was thought, for many years, that there are more problems efficiently solvable on a probabilistic Turing machine than on a standard Turing machine, a number of recent results have pointed in the opposite direction (e.g., Agrawal, Kayal, & Saxena, 2004), and it is now generally believed that classical probabilistic computation does not offer any performance advantage over classical deterministic computation (Arora & Barak, 2009, Ch. 20). In other words it is now widely believed that **P** = **BPP**, or that it is just as easy to solve a given problem on a deterministic Turing machine as it is on a probabilistic one. We have nevertheless stated the universality thesis in terms of **BPP** because this will prove convenient when it comes time to discuss the differences between classical and quantum computation. A quantum computer is, from one point of view, just another kind of probabilistic computer (that calculates probabilities differently), and it has the same success criterion as a classical probabilistic computer, i.e., we only demand that a given ‘solution’ be correct with ‘high enough’ probability.

<sup>19</sup>See also: Goldreich (2008, p. 33), who names it differently.

in polynomial time on a quantum computer. It is easy for a quantum computer to simulate a probabilistic Turing machine. Thus:

$$\mathbf{BPP} \subseteq \mathbf{BQP}, \tag{2.4}$$

i.e., the class of problems efficiently solvable on a quantum computer is at least as large as the class of problems efficiently solvable on a probabilistic Turing machine. What is still unknown, though it would be surprising if it were not the case, is whether

$$\mathbf{BPP} \subsetneq \mathbf{BQP}, \tag{2.5}$$

i.e., whether the class of problems efficiently solvable on a quantum computer (a physically realisable computational architecture) is *strictly larger than* the class of problems efficiently solvable on a probabilistic Turing machine. Note that if Eq. (2.5) is true, then both the universality and invariance theses are false.

There are some authors who view the consequences of the falsification of the universality thesis, in particular, to be profound. Bernstein & Vazirani (1997), for example, take it that computational complexity theory “rests upon” this thesis (p. 1411), and that the advent of quantum computers forces us to “re-examine the foundations” (p. 1412) of the theory. The sense in which complexity theory rests upon universality is expressed by Nielsen & Chuang (2000), who write that the falsity of the thesis implies that complexity theory cannot achieve an “elegant, model independent form” (p. 140). For Amit Hagar (2007), the failure of model-independence shakes, not only the foundations of complexity theory, but certain views in the philosophy of mind that depend on the model-independence of computational kinds (pp. 244–245).

If we actually examine what the universality thesis is saying, however, then it is not really clear, at least not *prima facie*, how it can ground the model-independence of complexity-theoretic concepts. The statement of the thesis is that any efficiently solvable problem is solvable efficiently by a probabilistic Turing machine. A probabilistic Turing machine is a particular model of computation, though. How can a thesis whose very definition makes reference to a particular computational model ground the model-independence of computational complexity theory? In fact there is a weak notion of model-independence being alluded to here. The point (see Nielsen & Chuang, 2000, p. 140) is that, for any assertion of the form: ‘problem  $P$  is efficiently solvable under computational model  $\mathfrak{M}$ ’, the qualification ‘under computational model  $\mathfrak{M}$ ’ can *always* (given the truth of the universality thesis) be substituted with ‘by a probabilistic Turing machine’ without changing the truth value of that sentence. Further, to show that such a sentence is true in general, it suffices to show that it is true for a probabilistic Turing machine. Finally, because ‘by a probabilistic Turing machine’ qualifies every such sentence we can leave it off and still expect to be understood. ‘Problem  $P$  is efficiently solvable by a probabilistic Turing machine’ is thus abbreviated to ‘problem  $P$  is efficiently solvable’. That we can do this is not insignificant, but this is arguably not a particularly deep sense of model-independence (see Cuffaro, 2018b, §11.6).

Far more interesting in relation to the question of model-independence is the invariance thesis. Unlike the universality thesis, model-independence is built right into the very statement of invariance. For after all, it amounts to the quite direct claim that the details of any particular reasonable machine model, since these can be efficiently simulated by any other reasonable model, are irrelevant for the purposes of characterising the complexity of a given problem.

Invariance, if taken to be true without qualification, clearly brings with it an absolute notion of model-independence (at least with respect to physically reasonable models). And if taken to be false (as it seems we should, given the existence of quantum computation) it clearly precludes such a notion. Arguably, however, what is (and always was) most valuable about the idea of invariance is not the absolute model-independence that is implied when it is taken to hold without qualification. As we will see later, the term ‘quantum computer’ does not refer to some single particular model of computation but is, rather, an umbrella term for a number of distinct computational models,<sup>20</sup> all of which have been shown to be computationally equivalent in the sense that they are all efficiently simulable by one another.<sup>21</sup> In other words, what we have learned from the study of quantum computing is that, in addition to the existence of an equivalence class of reasonable classical models that satisfy the invariance thesis with respect to one another, there is a further class of reasonable computational models, based on the principles of quantum mechanics, that satisfy the invariance thesis with respect to one another. Thus there are *two* distinct equivalence classes of physically reasonable computational models from the viewpoint of computational complexity theory. This is a discovery.

Invariance, thought of as a guiding rule or methodological principle, rather than as an absolute thesis, can be understood as grounding these investigations, and arguably this was the point all along (Cuffaro, 2018b, §11.6). Through the search for equivalence classes we carve out the structure of the space of computational models, yielding a notion of *relative* model-independence among the machine models comprising a particular equivalence class. And the fact that relative model-independence exists within the space of computational models at all arguably tells us something deep about how computer science connects up with the world, for the differences in computational power between these two reasonable classes of computational model are best understood by considering the *physics* needed to describe them. We discussed the physics of classical computers in Section 2.1. In the next section we turn to the physics of quantum computers.

## 2.4 Quantum states and operations


We saw earlier that in classical mechanics, assigning a particular state,  $\omega = (\mathbf{q}, \mathbf{p})$ , to a system fixes the answer to every conceivable yes-or-no experimental question that we can ask about it in advance, irrespective of whether we actually ask any questions or not. And we saw how to arrange the possible questions we can ask about a system, and their possible answers, into a truth-table-like structure like the one given in Figure 2. Note that specifying the values of the answers to the questions on the right-hand side of the double-vertical line in Figure 2 is another way of representing the state of a system like the one depicted in Figure 3. We saw how one can manipulate such a system in ways that can be represented abstractly as logico-mathematical operations on binary digits, where the binary digits are abstract representations of the properties of the system’s (two-level) subsystems.

The subject matter of *quantum* mechanics, just as it is for classical mechanics, is physical systems, where these can be anything from a single particle to the entire physical universe. In other words, quantum mechanics is not just a theory of the small. It is, rather, a new

---

<sup>20</sup>For instance, the quantum Turing model (Deutsch, 1985), the quantum circuit model (Deutsch, 1989), the cluster-state model (Briegel, Browne, Dür, Raussendorf, & den Nest, 2009), the adiabatic model (Farhi, Goldstone, Gutmann, & Sipser, 2000), and so on.

<sup>21</sup>See, for instance, Aharonov, van Dam, Kempe, Landau, Lloyd, & Regev (2007); Raussendorf & Briegel (2002); Nishimura & Ozawa (2009).



**Figure 7:** An ordinary vector in the Cartesian plane. It can be decomposed into  $x$  and  $y$  coordinates, as on the left, or into the alternative coordinates  $x'$  and  $y'$  as on the right. Similarly, a state vector, in quantum mechanics, can be decomposed into more than one Hilbert space basis. Two common bases (but not the only two) with which to decompose the state vector for a qubit are the  $\{|0\rangle, |1\rangle\}$  basis (known as the computational basis) and the  $\{|+\rangle, |-\rangle\}$  basis.

universal language for describing physical systems,<sup>22</sup> and it describes our experience with those systems better and more accurately than classical mechanics describes it. In quantum mechanics, just as in classical mechanics, a system can be, at a given moment in time, in any one of a number of possible physical states. These states are represented by *state vectors* (or ‘wave functions’). The state vector for a two-dimensional quantum system or ‘qubit’, for instance, is given in general by:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle. \quad (2.6)$$

When the complex numbers  $\alpha$  and  $\beta$  are both non-zero, a qubit is said to be in a *superposition* of the *basis vectors*  $|0\rangle$  and  $|1\rangle$ . Just as an ordinary vector in the Cartesian plane can be decomposed, mathematically, into  $x$  and  $y$  coordinates, a state vector for a two-dimensional quantum system can be decomposed, mathematically, into two basis vectors for the (Hilbert) space that it is situated in (see Figure 7). Two common bases (but not the only two) with which to decompose the state vector for a qubit are the  $\{|0\rangle, |1\rangle\}$  basis (known as the computational basis) and the  $\{|+\rangle, |-\rangle\}$  basis.

Associated with a given basis is a particular experimental question; for instance, “Is the qubit in the state  $|0\rangle$  (as opposed to the state  $|1\rangle$ )?” in the case of the computational basis, and in the case of the  $\{|+\rangle, |-\rangle\}$  basis, “Is the qubit in the state  $|+\rangle$  (as opposed to the state  $|-\rangle$ )?” When a system in the state  $\alpha|0\rangle + \beta|1\rangle$  is ‘asked’ whether it is in state  $|0\rangle$  as opposed to the state  $|1\rangle$ —i.e., when it is *measured* in the computational basis—then with probability  $|\alpha|^2$  the answer will come back as  $|0\rangle$  and with probability  $|\beta|^2$  it will come back as  $|1\rangle$ ; i.e.,

$$M(\alpha|0\rangle + \beta|1\rangle) \xrightarrow{\text{Pr} = |\alpha|^2} |0\rangle \quad (2.7)$$

$$M(\alpha|0\rangle + \beta|1\rangle) \xrightarrow{\text{Pr} = |\beta|^2} |1\rangle \quad (2.8)$$

where  $M$  is a *measurement operator* used to represent a computational basis measurement. Note that there is also a basis corresponding to the question: “Is the qubit in the state  $|\psi\rangle$  (as opposed to the state  $|\phi\rangle$ )?” Why not just ask this question? One reason is that measuring a quantum-mechanical system involves using a concrete physical device, and it turns out

<sup>22</sup>See note 2 above.

that it is technologically more feasible to construct measurement devices corresponding to some experimental questions than to others (Janas et al., Forthcoming, §6.5).

In terms of their dynamics, states of quantum systems evolve linearly:

$$U(\alpha|0\rangle + \beta|1\rangle) = \alpha U|0\rangle + \beta U|1\rangle \quad (2.9)$$

and unitarily:

$$UU^\dagger|\psi\rangle = U^\dagger U|\psi\rangle = |\psi\rangle \quad (2.10)$$

over time, where  $U$  is a *unitary operator* and  $U^\dagger$  is its *adjoint*.

Classical and quantum mechanics differ in the way that answers to experimental questions are determined. In contrast to classical mechanics, specifying a quantum system's state at a given moment in time does not fix in advance the answers to the experimental questions that can be asked of that system, or as Bub & Pitowsky (2010) put it, the quantum state is not a *truthmaker* in relation to those questions. It fails to be a truthmaker in two senses: First, a given state specification yields, in general, only the probability that the answer to a given experimental question will take on one or another value when asked. This in itself is not as much of a departure from classical mechanics as one might think, however, because conditional upon the selection of an experimental question, one can, in quantum mechanics, describe the observed probabilities as stemming from a prior classical probability distribution over the dynamical properties of the system that is determined in advance by the quantum state.<sup>23</sup>

This brings us to the second, more important, sense in which the state of a quantum-mechanical system fails to be a truthmaker in relation to questions that can be asked about that system's dynamical properties: The probability distributions associated with the answers to individual experimental questions cannot be embedded into a global prior probability distribution over all of the answers, as they can be for a classically describable system (see Figure 2). In quantum mechanics one can only say that *conditional* upon our inquiring about the observable  $A$ , there will be a particular probability for the answer to that question to take on a particular value. Thus in quantum mechanics, unlike in classical mechanics, we actually have to ask the system a question in order to get an answer from it.<sup>24</sup>

When quantum-mechanical systems are combined, they can sometimes become *entangled* with one another. Consider, by way of contrast, the simple illustration in Figure 3 of a classically describable system composed of many subsystems. This system is in a *product state*, which means that the state of the overall system can be expressed as a product of the states of each individual subsystem, i.e.,  $0_A 1_B 1_C 0_D$ , where the individual states of the subsystems are  $0_A$ ,  $1_B$ ,  $1_C$ , and  $0_D$ . In classical mechanics there is no other way to describe the combined state of a number of subsystems. Even when these subsystems are correlated with one another (as the state of one side of a coin is correlated with the state of the other side, for instance), as long as we include facts about those correlations in our description of the overall system, then its overall state can be factored into the individual states of its subsystems (Bell, 2004 [1981]). This is true even if the subsystems happen to be far apart, for instance if I print out two copies of a document and mail each of them to distant parts of the globe.

<sup>23</sup>This is analogous to the way we interpret probabilities in classical *statistical* mechanics.

<sup>24</sup>See also Janas et al. (Forthcoming, chs. 1 and 6), who call this (ironically) the 'small' measurement problem in contrast to the 'big' (this label is also ironic) problem described in the previous paragraph. These labels are originally due to Bub & Pitowsky (2010), who used them ironically too.

In quantum mechanics, in contrast, a system composed of multiple subsystems can sometimes be in a state like the following:

$$\frac{1}{\sqrt{2}}(|0\rangle_A|1\rangle_B - |1\rangle_A|0\rangle_B). \quad (2.11)$$

This state, describing a system composed of two subsystems  $A$  and  $B$ , is *not* a product of the individual states of its subsystems, for there is no way to factorise it into a state of the form:  $|\psi\rangle_A|\phi\rangle_B$  (the reader is invited to try). It is sometimes said that in an entangled quantum system the whole is somehow prior to the parts (Howard, 1989), or that the systems are *nonlocally correlated* (Maudlin, 2011), for if Alice measures her qubit in the computational basis and receives the result  $|0\rangle_A$ , then a measurement by Bob on his qubit (which in principle he may have taken to the opposite side of the universe) is thereby instantaneously determined to be  $|1\rangle_B$ , although there is no way for Bob to take advantage of this instantaneous determination for the purpose of predicting the outcome of a measurement on his subsystem, short of waiting for Alice to send him a classical signal (a text message or a phone call, for instance).<sup>25</sup>

Note that the correlation between the results of computational (i.e.,  $\{|0\rangle, |1\rangle\}$ ) basis measurements on an entangled system in this state do not, in themselves, represent as much of a departure as one might think from what one can describe using classical language. For with respect to this specific correlation, Bell (2004 [1964]) showed that it can be seen as arising from a classically describable system that has been prepared in a particular way. It is only when we consider other experiments on such a system that we see a departure from the possibilities inherent in classical description. In particular, because

$$|+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}, \quad |-\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}}, \quad (2.12)$$

Eq. (2.11) can be rewritten, in the  $\{|+\rangle, |-\rangle\}$  basis, as:

$$\frac{1}{\sqrt{2}}(|+\rangle_A|-\rangle_B - |-\rangle_A|+\rangle_B). \quad (2.13)$$

Indeed, we obtain the *same form* for the state of the system *regardless* of which basis we choose to express it in:

$$\frac{1}{\sqrt{2}}(|b_1\rangle_A|b_2\rangle_B - |b_2\rangle_A|b_1\rangle_B). \quad (2.14)$$

Eqs. (2.11), (2.13), and (2.14) all yield the same probabilities for the results of experiments.

If we collect a number of pairs of subsystems prepared in this entangled state, and then ask of each pair: “Is this pair in the state  $|0\rangle|1\rangle$  (as opposed to  $|1\rangle|0\rangle$ )?” then we will find that the answers ‘yes’ and ‘no’ will be obtained with equal probability. Conditional upon this question, we can imagine these answers as arising from a prior classical probability distribution over the properties of the subsystem pairs, half of which we imagine to be in the state  $|0\rangle|1\rangle$ , and half of which we imagine to be in the state  $|1\rangle|0\rangle$ . This prior probability distribution is *incompatible*, however, with the prior classical probability distribution that we might imagine to be responsible for the answers yielded from repeatedly asking the question: “Is this pair in the state  $|+\rangle|-\rangle$  (as opposed to  $|-\rangle|+\rangle$ )?” In the context of the

---

<sup>25</sup>This restriction on the use that can be made of nonlocal correlations in quantum mechanics is called ‘no signalling’. For discussion, see Bub (2016); Cuffaro (2020).



latter question, we can imagine the answers as arising from a prior classical probability distribution over the properties of the subsystem pairs, half of which we imagine to be in the state  $|+\rangle|-\rangle$ , and half of which we imagine to be in the state  $|-\rangle|+\rangle$ . Asking a question of an imagined ensemble of systems, half of which are in the state  $|0\rangle|1\rangle$ , and half of which are in the state  $|1\rangle|0\rangle$ , will yield different statistics, however, than the ones that would be yielded by asking the same question of an imagined ensemble of systems, half of which are in the state  $|+\rangle|-\rangle$ , and half of which are in the state  $|-\rangle|+\rangle$ . These probability distributions are, in this sense, incompatible. And yet *all* of these statistics are predicted by one and the same quantum state, the one expressed in the unitarily equivalent forms given in Eqs. (2.11–2.14). In quantum mechanics, unlike in classical mechanics, the questions we ask or do not ask actually matter for the ways that we can conceive of the underlying properties of a physical system, even when the state of that system is fully specified (Janas et al., Forthcoming, §6.5).

### 3 Quantum computation and parallel universes

There are problems that are hard for a classical computer, but that a quantum computer can solve easily (where ‘easy’ and ‘hard’ are meant in the computational sense defined above in Section 2.2).<sup>26</sup> The most famous example of such a problem is that of factoring large integers into primes. The best-known classical algorithm for factoring is the number field sieve (Lenstra, Lenstra Jr., Manasse, & Pollard, 1990), which takes on the order of  $2^{(\log N)^{1/3}}$  steps to factor a given integer  $N$ .<sup>27</sup> No one has yet proven that this is the best that a classical computer can do. All the same, encryption algorithms such as RSA (named for the paper by Rivest, Shamir, & Adleman in which it was introduced), that are today widely used for secure transactions on the internet, rely on the assumption that it is; i.e., that one cannot factor in polynomial time. In 1994, Peter Shor discovered a *quantum* algorithm—now known as Shor’s algorithm—that can factor integers in on the order of  $\log N$  steps for a given integer  $N$ —an exponential ‘speedup’ over the number field sieve.

What explains this phenomenon? What is it, exactly, that allows quantum computers to compute certain problems more efficiently than classical computers can? Surprising as it may seem, there is still no consensus regarding the answer, neither among the researchers working in the field, nor among the philosophers commenting on this work. By contrast, in the popular literature on the topic, one answer has tended to dominate all the others. This is the idea that quantum computers are actually *parallel* devices that perform exponentially many classical computations simultaneously, with each computation taking place in a parallel physical universe (or ‘world’), different in certain ways from our own but just as real as it is.

Perhaps the strongest proponent of this view, which we will call the *many-worlds explanation* of the power of quantum computers, is one of the fathers of the field, David Deutsch, discoverer of the first quantum algorithm and of the universal quantum Turing machine (Deutsch, 1985). The many-worlds explanation of quantum computing is, for Deutsch, not just a speculative conjecture. For Deutsch, many worlds are the only plausible explanation

---

<sup>26</sup>There are also problems for which a quantum computer, despite being unable to solve them easily, can nevertheless solve them significantly *more* easily than a classical computer can. An example is the problem to search an unstructured database, for which a quantum (“Grover’s”) algorithm can reduce the number of steps required by a quadratic factor over any known classical algorithm. See: Grover (1996), Bennett, Bernstein, Brassard, & Vazirani (1997), and for further discussion see Cuffaro (2018b, p. 269).

<sup>27</sup>For the meaning of ‘on the order of’ see fn. 10.

for how quantum computers work. As he puts it in his 1997 book, *The Fabric of Reality*: “[t]o those who still cling to a single-universe world-view, I issue this challenge: Explain how Shor’s algorithm works” (1997, p. 217).

Before we discuss what those who, like Deutsch, defend this view of quantum computation are really saying, let us review briefly what the view amounts to in more general terms in the context of quantum mechanics, for what the many-worlds explanation of quantum computing amounts to is an application of this more general philosophical interpretation of quantum mechanics to the case of quantum-mechanical computers.<sup>28</sup> Thus far, our examples of quantum-mechanical systems have been of qubits, either considered singly or in combination with other qubits. But quantum mechanics is not restricted, of course, to describing qubits. Quantum mechanics is, rather, a universal language for describing any physical system, be it a single particle or (in principle) the entire physical universe. From the point of view of quantum mechanics, you and I, the laptop I am writing this chapter with, the table it is sitting on, and everything else in the world is a quantum-mechanical system. And just as a qubit can be in a superposition of its basis states (see Eq. (2.6) and Figure 7), so too can *any* physical system. As Schrödinger famously remarked, according to quantum mechanics there are even ways of preparing a cat in a superposition of ‘dead’ and ‘alive’ states (Schrödinger, 1935a).

When we actually come to pose a question to the system; i.e., to measure the qubit, the cat, or what have you, we never actually find that it is in a superposition state. Quantum mechanics predicts that the cat will be found to be either dead or alive, the qubit to be either  $|0\rangle$  or  $|1\rangle$ , and so on, with a certain probability. Thus, unlike the physical state of a system as described by classical mechanics, we cannot simply read off from the quantum state description of a system what the outcome of an experiment on that system will be (except in special cases). This raises the question of how to interpret the quantum state. Are we to take superpositions literally, despite the fact that we never observe them? What would it mean if we did? The proponent of the many-worlds interpretation of quantum mechanics answers yes to the first question. As for the second question, the view of the many-worlds advocate is that each branch of a superposition represents a distinct, classically describable (from the point of view of that branch), physical universe.

Just as it is with a cat, or with any physical system, so it is with a quantum computer. What is different, arguably, about a quantum computer (according to the defender of the many-worlds explanation of quantum computing) is that in a quantum computer the resources made available by these multiple physical universes are harnessed to perform computations. This is a striking claim. But it can be said, in favour of the many-worlds explanation, that some quantum algorithms certainly do give us this impression. The following evolution is representative of a typical step in many quantum algorithms (note that normalisation factors have been omitted for simplicity):

$$\sum_{x=0}^{2^n-1} |x\rangle^n |0\rangle \rightarrow \sum_{x=0}^{2^n-1} |x\rangle^n |f(x)\rangle, \quad (3.1)$$

---

<sup>28</sup>The interpretation of quantum mechanics that we will be discussing in this section is one of a number of related interpretations of quantum mechanics that are collectively referred to as the ‘Everett interpretation’. These include but are not limited to Hugh Everett III’s original formulation (Everett III, 1956), the ‘Berlin Everettianism’ of Christoph Lehner (1997), Lev Vaidman’s version of Everett (Vaidman, 1998), so-called ‘many minds’ variants (Albert & Loewer, 1988), and finally the ‘many-worlds’ variants that are the direct inspiration for the many-worlds explanation of quantum computing. Belonging to the last group are Bryce DeWitt’s (1973 [1971]) view, as well as the ‘Oxford Everett’ interpretation (Deutsch, 1997; Saunders, 1995; Wallace, 2003, 2012) with which we will be mostly concerned here.

What is being depicted here is the state of a quantum computer, composed of  $n + 1$  qubits, where the first  $n$  qubits are called the *input qubits* of the computer, and the last qubit is called the computer's *output qubit*. The computer begins in the state described on the left-hand side of the equation, and then transitions to the state described on the right-hand side.

This notation is very compact, so let us unpack it. Considering, first, the left-hand side: If we take the  $n$  input qubits as comprising, together, one subsystem of the system, and the output qubit as comprising another subsystem, then we can say that the input and output subsystems of the computer begin in a product state. In other words we can write the left-hand side of the equation in product form as:

$$\text{LHS} = \left( \sum_{x=0}^{2^n-1} |x\rangle^n \right) |0\rangle. \quad (3.2)$$

Note that, in Eq. (3.2), the  $n$  input qubits are together in a superposition of all of the possible computational basis states for a system of  $n$  qubits. This is obscured somewhat, both by the summation notation as well as by the shorthand,  $|x\rangle^n$ , being used to represent each superposition term. To clarify this, recall that for a two-dimensional system, i.e., a single qubit, the computational basis states are, as we pointed out earlier (see Eq. (2.6) and Figure 7), the states  $|0\rangle$  and  $|1\rangle$ . As for a system composed of *two* qubits, it has four computational basis states, namely,  $|0\rangle|0\rangle$ ,  $|0\rangle|1\rangle$ ,  $|1\rangle|0\rangle$ , and  $|1\rangle|1\rangle$ . In general, for a system of  $n$  qubits there are  $2^n$  basis states:

$$\begin{aligned} &|0\rangle|0\rangle \dots |0\rangle|0\rangle, \\ &|0\rangle|0\rangle \dots |0\rangle|1\rangle, \\ &|0\rangle|0\rangle \dots |1\rangle|0\rangle, \\ &\quad \dots \\ &|1\rangle|1\rangle \dots |1\rangle|1\rangle, \end{aligned} \quad (3.3)$$

each of which can be thought of as representing a binary number. For short, we can represent these in decimal as  $\{|0\rangle, |1\rangle, |2\rangle, |3\rangle, \dots, |2^{n-1}\rangle\}$ . A superposition of all of these basis states is then given by:

$$|0\rangle + |1\rangle + |2\rangle + |3\rangle + \dots + |2^{n-1}\rangle = \sum_{x=0}^{2^n-1} |x\rangle^n, \quad (3.4)$$

exactly as we see in Eq. (3.2).

On the right-hand side of Eq. (3.1), the  $n + 1$  qubits of the quantum computer are no longer in a product state. The state of the computer has now transitioned to an *entangled* state: a superposition in which the state obtained for the output qubit, conditional upon measuring it, is correlated with the state that will be obtained for the input qubits, conditional upon measuring them. For instance, if we measure the input qubits and get the state  $|1\rangle$ , then the state of the output qubit will be  $|f(1)\rangle$ , for the given function  $f$ , and similarly for every other value of  $x$ . What is the function  $f$ ? That will depend on the particular problem that the algorithm is designed to solve. More importantly, from the point of view of our current discussion, notice that the right-hand side of the equation encodes exponentially many evaluations of that function in the state of the computer! We can, or so it seems, appeal to the results of all of these evaluations to help us with whatever problem we are trying to solve.

In reality things are not so easy, for when we actually come to read off the result of the computation, we will only ever find the computer's output qubit to be in *one* of its exponentially many superposition terms; i.e., in some state  $|f(a)\rangle$  for some one particular  $a$  in the domain of  $f$ . If we are tempted to read the presence of these exponentially many function evaluations in the description of the state of the computer literally, then the fact that only one of them can ever be accessed in a particular run should give us some pause (cf. Mermin, 2007, p. 38). That said, in any given run, there will be an, in general, non-zero probability of obtaining any one of them, and this, perhaps, speaks in favour of viewing them all as somehow literally there. As for the goal (which we should not lose sight of) of actually solving the problem at hand, what is just as important as achieving the form of the state of the computer on the right-hand side of Eq. (3.1) is the *next* step in a computation like this one, which requires that we manipulate the system cleverly enough so that one of the desired solutions is found with higher probability than the other, undesirable, solutions (Pitowsky, 2002, §3).

The many-worlds explanation of a quantum computational process enjoins us to take an evolution like the one given in Eq. (3.1) ontologically seriously. It affirms, in other words, that things are indeed as they seem: The computer most definitely *is* performing many simultaneous function evaluations in parallel when it is in a state like the one on the right-hand side of Eq. (3.1). Moreover the many-worlds explanation directly answers the question of *where* this parallel computation occurs, namely in distinct physical universes. For this reason it is also, arguably, one of the more intuitive of the physical explanations of quantum speedup. And it is certainly thought-provoking. It is no wonder, then, that is the one most often mentioned in the popular literature on quantum computation. And it has advocates in the more serious literature as well.<sup>29</sup>

So far we have highlighted one potential problem for the many-worlds explanation: Despite the appearance of parallel processing in an equation like Eq. (3.1), only one of these function evaluations is ever accessible on a given run of the computer. This is not to say that no story can be told from the point of view of the many-worlds explanation about why this is so, but at the very least this is enough to cast doubt on the claim that an evolution like the one given in Eq. (3.1) constitutes, all by itself, evidence for the many-worlds explanation of the power of quantum computers. We will return to this issue later. But for now we need to consider a somewhat deeper problem faced by the advocate of the many-worlds explanation. Recall that decomposing a qubit in the state  $|\psi\rangle$  into the computational basis states  $|0\rangle$  and  $|1\rangle$  (see Eq. (2.6) and Figure 7) is only one way to express this state. The same state can also be expressed in other bases besides the computational basis. This means that a system that is in the state

$$\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle, \quad (3.5)$$

from the point of view of the computational basis, is simply in the state:

$$|+\rangle \quad (3.6)$$

from the point of view of the  $\{|+\rangle, |-\rangle\}$  basis (see Eq. (2.12)). But now it is no longer clear just how many universes we should take this qubit to exist in. If we decompose its

---

<sup>29</sup>In addition to Deutsch's 1997 book, see Deutsch (2010), and see also Vaidman (2018 [2002], §7) and Wallace (2012, Ch. 10). The strongest and most in-depth defence of the many-worlds explanation of quantum computing that I am aware of is the one given by Hewitt-Horsman (2009).

state in the computational basis then it would seem that it exists in two worlds, and if we decompose it into the  $\{|+\rangle, |-\rangle\}$  basis then it would seem that it only exists in one. Which is it? If we decide that one of these decompositions is to be preferred to the other, then the challenge for the advocate of the many-worlds explanation is to give a compelling reason why, for *prima facie* there does not appear to be any reason for preferring one of these decompositions over the other.

This is known as the *preferred basis problem*, and it is a problem for the many-worlds view more generally, i.e., not just in the context of quantum computing. In this more general context, advocates of the many-worlds picture (see, for instance, Wallace, 2003) attempt to solve the preferred basis problem by appealing to the dynamical process of *decoherence*. In quantum mechanics, as we saw (see Eq. (2.6) and Figure 7), the state of a system at a particular moment in time is described by a state vector that evolves linearly and unitarily. However the evolution thereby described is the evolution of a system that is completely isolated from its external environment, which is an idealisation; in reality it is actually never possible to completely isolate a system from its environment,<sup>30</sup> unless, perhaps, we take our system of interest to be the universe in its entirety.<sup>31</sup> But barring the universe as a whole, when a system interacts with an external environment—for example with our measurement apparatus as we ask the system an experimental question—then the terms in the superposition describing its state begin to decohere (Zurek, 2003 [1991]) and come to achieve a kind of independence from one another (although they never decohere completely). On the many-worlds picture we are to think of such (approximately) decoherent terms as existing in independently evolving physical universes. And in each of these independently evolving physical universes there is in addition a different version of ourselves, all of whom receive a different answer to the experimental question that was asked. And with every additional question, the universe as we know it is branching, spawning exponentially more and more versions of the system, and more and more versions of ourselves along with it, and so on and on and on.

Fundamentally, however, decoherence is an approximate phenomenon, and some small amount of residual interference between worlds always remains despite it. Nevertheless decoherence tells us that, when the environment and system happen to be correlated in a particular way, then *a particular basis will emerge with respect to which* we can effectively describe the superposition terms expressed in the state of the system as evolving independently of one another. As David Wallace (2003, p. 90) puts it: “the basic idea is that dynamical processes cause a preferred basis to emerge rather than having to be specified a priori.” In this way, superposition terms that for all practical purposes evolve stably and independently over time with respect to the decoherence basis can be identified with different copies of measurement pointers, cats, experimenters, and whatever else is theoretically useful for us to include in our ontology: “the existence of a pattern as a real thing depends on the usefulness—in particular, the explanatory power and predictive reliability—of theories which admit that pattern in their ontology” (Wallace, 2003, p. 93). Whatever else one may think of the many-worlds view, decoherence, at least, does provide a principled way to identify worlds in the wave-function.

For the advocate of the many-worlds explanation of *quantum computation*, however, there is still a problem. Appealing to decoherence may solve the preferred basis problem for

---

<sup>30</sup>At the very least, the gravitational effects of other distant systems will not be able to be neglected.

<sup>31</sup>Some philosophers have questioned whether we should think of even the universe as a whole as a closed system (see, for instance, Cuffaro & Hartmann, 2021; Wallace, 2012, §10.5).

the purposes of describing the world of our everyday experience,<sup>32</sup> but the inner workings of a quantum computer are not part of that everyday experience. The problem is not just that qubits are too small to see. The problem is that the superpositions characteristic of quantum algorithms are *coherent* superpositions (Nielsen & Chuang 2000, p. 278; see also Aaronson 2013c; Cuffaro 2012). Thus the terms in the wave-function of a quantum computer do not seem to meet the criterion for world-identification advocated for by the many-worlds view.

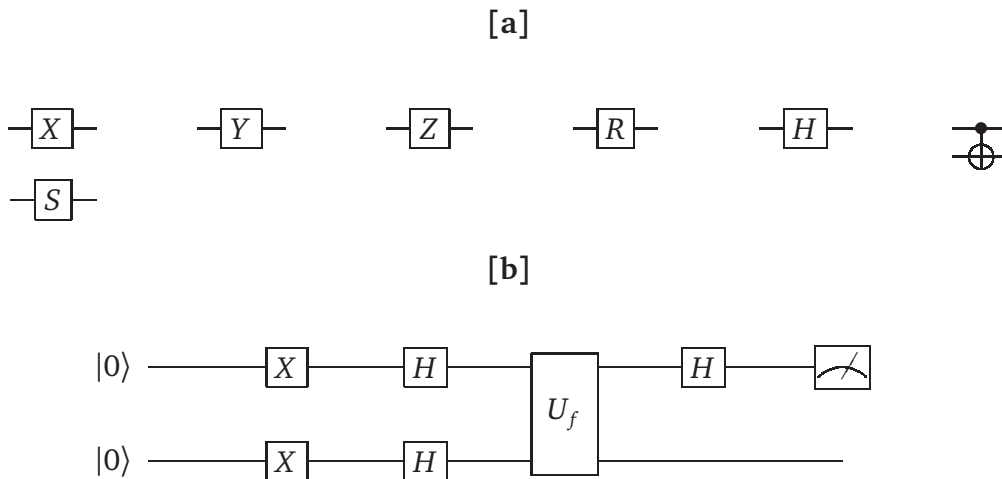
Now, to be fair, a similar thing can be said with regard to our everyday experience. So-called decoherent superpositions are (as I mentioned) not *really* decoherent, after all. But they are decoherent enough, and for long enough, that it is useful (or so says the many-worlds advocate) to think of the terms in such a superposition as independent, and the worlds that they describe as ontologically real. Likewise, although the superposition on the right-hand side of Eq. (3.1) is actually a coherent superposition, it may nevertheless be useful to think of the terms in that superposition as independent, at least for the short time that the quantum computer is in that state (Hewitt-Horsman, 2009, p. 876). The problem, however, is that it is the very fact that they are coherent that allows us to ‘cleverly’ extract desirable solutions from these superpositions with higher probability than undesirable solutions (Bub, 2010; Duwell, 2007).

Even if we grant that there is some heuristic value in describing a quantum computer in a superposition state as evaluating functions in exponentially many parallel worlds (I do not doubt that this was of some heuristic use to Deutsch, for instance, even if I question whether it is *necessary* to think of such superpositions in this way), it does not follow that this is enough to licence granting ontological status to those worlds. Wallace (2003, p. 93) mentions (as we saw) explanatory power and predictive reliability, for instance, and discusses the way that these and other ideas are applied in contemporary physics to support the many-worlds view outside of the context of quantum computing. It is not at all clear that these criteria are met in the context of quantum computing, however, and even Wallace admits that they for the most part are not: “There is no particular reason to assume that *all* or even *most* interesting quantum algorithms operate by any sort of ‘quantum parallelism’” (Wallace, 2010, p. 70, n. 17). Wallace goes on: “Shor’s algorithm, at least, does seem to operate in this way” (ibid.), but he does not describe how. Yet there are very plausible accounts of how Shor’s algorithm works that do not appeal to massive parallelism at all (see Bub, 2010). Far from it, on Jeffrey Bub’s account of Shor’s algorithm, the quantum algorithm is more efficient than known classical algorithms because it performs *fewer*, not more, computations (see also Bub, 2008).

The final reason that I will mention for being skeptical of the many-worlds explanation of quantum computing is that it only really seems to be useful in the context of one particular model of quantum computing. This is the so-called *quantum circuit model*, the model for which many of the first quantum algorithms were designed (Deutsch, 1989). This model is useful for abstract theoretical purposes, as well as for pedagogical purposes, as it borrows many familiar ideas from the classical circuit model of computation (see Figure 4). In the quantum circuit model, similarly to the classical circuit model, logical circuits are constructed out of various ‘quantum logic gates’. These instantiate unitary transformations of one or more qubits that are prepared beforehand in computational basis states (typically qubits begin a computation in the state  $|0\rangle$ ). The unitary gates transform the qubits’ states

---

<sup>32</sup>The preferred basis problem is not the only challenge that needs to be met by an advocate of the Everett interpretation of quantum mechanics. Another issue that has been much discussed in recent literature is the problem of how to account for probabilities on the Everettian view. For more on this issue see Adlam (2014), Dawid & Thébault (2015), Greaves & Myrvold (2010), Vaidman (1998, 2012), and Wallace (2007).



**Figure 8:** [a] A number of important one- and two-qubit quantum gates. The  $X$ ,  $Y$ ,  $Z$ ,  $R$ ,  $H$ , and CNOT gates together form the *Clifford group*. If we add the  $S$  gate to the Clifford group, they together form a *universal set*. [b] A quantum circuit diagram for Deutsch’s Algorithm, which determines whether a given function  $f$  on one bit is constant ( $f(0) = f(1)$ ) or balanced ( $f(0) \neq f(1)$ ).

into various superpositions of computational basis states, and at the end of the computation a measurement is performed, again in the computational basis, on (some of) the qubits and the results are read out.

Figure 8[a] depicts a number of important one- and two-qubit quantum gates. The  $X$  gate implements a qubit-flip operation; i.e., it takes  $|0\rangle \rightarrow |1\rangle$  and vice versa. The  $Y$  gate takes  $|0\rangle \rightarrow i|1\rangle$  and  $|1\rangle \rightarrow -i|0\rangle$ . The  $Z$  gate takes  $|0\rangle \rightarrow |0\rangle$  and  $|1\rangle \rightarrow -|1\rangle$ . The  $R$  gate takes  $|0\rangle \rightarrow |0\rangle$  and  $|1\rangle \rightarrow i|1\rangle$ . The  $H$  (or Hadamard) gate takes  $|0\rangle \rightarrow (|0\rangle + |1\rangle)/\sqrt{2}$  and  $|1\rangle \rightarrow (|0\rangle - |1\rangle)/\sqrt{2}$ . The  $S$  gate takes  $|0\rangle \rightarrow |0\rangle$  and  $|1\rangle \rightarrow e^{i\pi/4}|1\rangle$ . At the extreme right is the two-qubit CNOT (or controlled-not) gate. It leaves the topmost qubit unchanged. The bottom qubit is then assigned the output of taking the exclusive-or of both, i.e., this gate takes  $|0\rangle|0\rangle \rightarrow |0\rangle|0\rangle$ ,  $|0\rangle|1\rangle \rightarrow |0\rangle|1\rangle$ ,  $|1\rangle|0\rangle \rightarrow |1\rangle|1\rangle$ , and  $|1\rangle|1\rangle \rightarrow |1\rangle|0\rangle$ . The  $X$ ,  $Y$ ,  $Z$ ,  $R$ ,  $H$ , and CNOT gates together form the *Clifford group* of gates, which we will have more to say about later. If we add the  $S$  gate to the Clifford group, they together form a *universal set* of gates, i.e., any quantum circuit implementing any series of unitary transformations can be simulated to arbitrary accuracy using combinations of these seven gates.

Figure 8[b] depicts a quantum circuit diagram for Deutsch’s Algorithm, which determines whether a given function  $f$  on one bit is constant ( $f(0) = f(1)$ ) or balanced ( $f(0) \neq f(1)$ ): Two qubits are prepared in the product state  $|0\rangle|0\rangle$  and are each sent through an  $X$ -gate and a Hadamard gate, after which they are together input to the two-qubit entangling unitary gate  $U_f$ . The first qubit is then sent through a further Hadamard gate and finally measured to yield the answer (see Deutsch, 1989).

Just as there are various models of universal classical computation (for instance the various versions of the Turing machine, as well as the von Neumann architecture, and so on, that I mentioned above), there are various models of universal quantum computation. One computational model that presents a particularly difficult problem for those who would advocate for the many-worlds explanation is the *cluster-state* model of quantum computing, also known as *measurement-based* and *one-way* quantum computing (Raussendorf & Briegel, 2002; Raussendorf, Browne, & Briegel, 2003; Nielsen, 2006).<sup>33</sup> In the cluster-state

<sup>33</sup>For introductions to cluster-state quantum computing aimed at philosophers, see Cuffaro (2012, §4) and

model, the computer begins, not in a product state but in a highly entangled state, and measurements are performed not only at the end of the computation but throughout it. These measurements are *adaptive*, in the sense that each measurement is performed in a *different basis*, which depends on the random outcomes of whatever previous measurements have been performed.

Why is this a problem for the advocate of the many-worlds explanation of quantum computing? Because the fact that measurements are adaptive means, on the one hand, that there is no principled way to select a preferred basis *a priori* in the context of a given computation (Cuffaro, 2012, §4). Whichever basis we choose, few qubits will actually be measured in that basis in the course of the computation. On the other hand, there is no sense in which we can say that a preferred basis ‘emerges’ from the computational process. There is therefore no way to identify the worlds that the computation as a whole is supposed to be happening in.

As I alluded to above, both the cluster-state model and the circuit-model are universal models of quantum computing. Thus anything that the circuit model can do can be done (and, as it happens, just as efficiently) in the cluster-state model and vice versa. Perhaps, then, it could be argued that understanding the ‘true nature’ of algorithms in the cluster-state model requires that we first translate them into the language of the circuit model, though I can think of no reason why one should think so other than a desire to hold on to the many-worlds idea come what may. The proper response to anyone who would put forward such an argument is that a large part of what motivates those who adhere to the many-worlds explanation of quantum computing in the first place, is that it is useful for algorithm analysis and design to believe that a quantum computer is carrying out its computations in parallel worlds. This does not seem to be so for the cluster-state model. On the contrary, dogmatically holding on to the view that many worlds are, at root, physically responsible for the speedup evinced in the cluster-state model is at best useless, for it is of no help in designing algorithms for the cluster-state model. At worst, dogmatically holding on to the many-worlds idea could prove positively detrimental if it prevents us from exploiting the power of the cluster-state model or discovering other quantum computational models in the future.

## 4 Quantum computation and entanglement

The many-worlds explanation of the power of quantum computers is not the only explanation that has been proposed by philosophers. I have already mentioned the upshot of Bub’s analysis of Shor’s algorithm in the previous section. Bub’s view, more generally (2006; 2010), is that the key to understanding the power of quantum computers lies in understanding the way that they exploit the novel logico-probabilistic structure of quantum mechanics (Pitowsky, 1989). We saw above that in quantum mechanics, the (classical) probability distribution over the answers to a particular experimental question—the one that we infer, conditional upon our asking that question—cannot be embedded into a global prior probability distribution over all of the answers to all of the questions we might want to ask. But although these individual probability distributions do not logically fit together as they do for a classical system, there are logical relations between them nonetheless, that are compactly described by the quantum state.

For instance, as a general rule (which admits exceptions) one cannot use a quantum

---

Duwell (2018, §4.5).



system to compute a logical disjunction by computing the individual values of each disjunct. This is simply because in the general case both disjuncts will not be globally defined. However quantum mechanics' logical structure provides other ways to compute a logical disjunction, and these other ways (which are not available to a classical computer) are exploited by a quantum computer to compute certain problems more efficiently than a classical computer can compute them.

Another view is Armond Duwell's (2018; 2021), who in contrast, agrees with the proponent of the many-worlds explanation in identifying quantum parallelism as at least part of the source of the power of quantum computers. Duwell, however, resists the temptation to draw a metaphysical inference from parallelism to many computational worlds. For Duwell the source of the power of a quantum computer lies in the way that it can efficiently correlate multiple values of a function and use these correlations to efficiently extract global information about the function. The disagreement over whether a quantum computer performs more, or fewer, computations than a classical computer is one that Duwell views as arising from conflicting intuitions about how to appropriately describe the quantum systems that perform computational tasks.

In these and other candidate explanations of the power of quantum computers that one encounters in the philosophical literature on the topic, the fact that quantum-mechanical systems exhibit *entanglement* (see Section 2.4) invariably plays an important role. For Bub entanglement is absolutely central, as entanglement is a (direct) manifestation of the fact that the logical structures of quantum and classical mechanics differ. As for Duwell, he takes his quantum parallelism thesis to be completely compatible with the idea that entanglement plays a central role in quantum computing, even if his explanation emphasises the correlations between the values of the function being evaluated by a system rather than the logical structure of its underlying state space (Duwell, 2018, p. 101). The many worlds advocate, as well, views quantum entanglement to be indispensable in the analysis of the power of quantum computing (Hewitt-Horsman, 2009, p. 889), even if for the many-worlds advocate it does not, by itself, suffice as a philosophical explanation for it.

The debate over the interpretation of the phenomenon of quantum entanglement has historically been one of the central debates in the controversy over quantum theory's conceptual foundations. First emphasised by Albert Einstein, Boris Podolsky, and Nathan Rosen in their 1935 criticism of the orthodox interpretation of quantum mechanics, Erwin Schrödinger called it "*the characteristic trait of quantum mechanics, the one that enforces its entire departure from classical lines of thought*" (Schrödinger, 1935b, p. 555, emphasis in the original). It was only with the work of John Bell (2004 [1964]; 2004 [1966]), however, that its significance for our understanding of the break that quantum mechanics makes with classical physics was first made fully clear.

Consider Louise and Robert, two friends who live in the 15th and 11th arrondissements, on the left bank and the right bank of the Seine, respectively, in the city of Paris. Every morning, Louise and Robert each receive a letter in their respective mailboxes that consists of a single sheet of paper inscribed with either a large bass clef:



or a large treble clef



After awhile, Louise determines that it is equally likely, on any given day, that she will receive a bass clef as it is that she will receive a treble clef. After awhile Robert determines the same. But when they compare their notes, they find that whenever Robert receives a bass clef, so does Louise. Similarly, every time he receives a treble clef, she does too. In other words, there is a one-in-two chance, on any given day, that they both receive a bass clef, and a one-in-two chance that they both receive a treble clef. No other combinations are ever observed. In other words, their outcomes are *correlated*, such that the probability distribution over the possible combinations of outcomes is given by:<sup>34</sup>

$$\frac{1}{2}[\text{Bass}]_L[\text{Bass}]_R + \frac{1}{2}[\text{Treble}]_L[\text{Treble}]_R. \quad (4.1)$$

What explains this correlation? Well in this case it turns out that Louise and Robert both play in a jazz ensemble. The band's leader, Carsten, lives in the centre of the city (on the Île Saint-Louis). Every afternoon he flips a fair coin. Depending on whether the coin lands heads or tails, he either writes a large bass clef or a large treble clef on a sheet of paper, photocopies it, and sends one copy each to Robert and Louise. If Robert receives a treble clef from Carsten, then he will know to take his tenor horn to the jazz club that night. Otherwise he will take his trombone. As for Linda, if she receives a treble clef, she will know to bring her soprano clarinet. Otherwise she will bring her bassoon.

The result of Carsten's coin flip is called the *common cause* of Louise's and Robert's outcomes, and the story we tell about how Carsten's coin flip determines Louise's and Robert's outcomes is called a *causal model* for the correlations that they see (see Pearl, 2009). If Louise and Robert do not know how Carsten is determining what to send them, then they will wonder about the *hidden-variable* that explains their correlation. In this case the result of Carsten's coin flip is actually a *local* hidden-variable, since the process by which the outcome of the coin flip determines what is written on the letters is confined to the small localised region in the vicinity of the desk in Carsten's central office. He flips the coin, takes note of the outcome (the local hidden-variable), and writes the corresponding symbol.

Instead of simply flipping a coin while sitting at his desk, we can imagine a more complicated, spatially distributed, process by which Carsten determines what to write. For instance Carsten might begin by flipping his coin, and then, corresponding to heads or tails he might telephone Tilde or Bjarne, who live in the northern and southern suburbs of the city respectively, and ask whichever of them he calls to roll a six-sided die at exactly seven-o'clock in the evening, and to afterwards call him back at either 7:13PM, if it is Bjarne, or 7:18PM, if it is Tilde, to tell him the outcome. Then, if the result of the die roll is one, three, or six, Carsten will write a treble clef on the sheet of paper before photocopying it and sending copies to Robert and Louise, while if it is two, four, or five, he will write a bass clef. All of these actions together constitute, just like the coin flip in our simpler scenario, a locally causal model to explain Louise's and Robert's correlation. Why do we call it local even though it is spatially distributed? Because the physical processes (the coin flips, die rolls, and telephone calls) by which Carsten determines what to write propagate locally in the *physical sense*, i.e., at a finite speed that is less than the speed of light.

<sup>34</sup>If the outcomes were completely uncorrelated, the probability distribution would be

$$\frac{1}{4}[\text{Bass}]_L[\text{Bass}]_R + \frac{1}{4}[\text{Bass}]_L[\text{Treble}]_R + \frac{1}{4}[\text{Treble}]_L[\text{Bass}]_R + \frac{1}{4}[\text{Treble}]_L[\text{Treble}]_R.$$

Bell (2004 [1964]) showed, first, that there are certain probabilistic constraints—what are now called *Bell inequalities*—on the statistics arising from measurements on *any* locally correlated system. He then showed that these constraints are violated by statistics that are predicted to arise from certain experiments on a system composed of two qubits in the quantum-mechanical state:<sup>35</sup>

$$\frac{1}{\sqrt{2}}(|0\rangle_A|1\rangle_B - |1\rangle_A|0\rangle_B), \quad (4.2)$$

where the  $A$  and  $B$  subsystems can be as far apart as one likes. The proof of this violation is known as *Bell's theorem* and it, and its variants, have since been experimentally confirmed many times over (Genovese, 2016).

The predicted violation only occurs for certain measurements. If we measure both the  $A$  and  $B$  qubits in the computational basis (see Eq. (2.6) and Figure 7), then the predicted statistics will actually be compatible with the constraints imposed by local hidden-variable theories, as Bell himself showed (2004 [1964], p. 16). But as we rotate the measurement basis (see Figure 7) that we use for  $B$  away from the measurement basis that we use for  $A$ , quantum mechanics predicts that we will see a violation, that it will reach a maximum at a certain point, and then decrease again as the measurement basis for  $B$  begins to line up again with the measurement basis for  $A$ .

Another name for the computational basis is the  $Z$ -basis. We call it the  $Z$ -basis because the two computational basis vectors  $|0\rangle$  and  $|1\rangle$  correspond to the two possible outcomes of a  $Z$ -basis measurement on a qubit, where “measuring a qubit in the  $Z$ -basis” means sending it through a  $Z$ -gate (see Figure 8) and then recording its state (Janas et al., Forthcoming, §6.5). Similarly, the basis vectors  $|+\rangle$  and  $|-\rangle$  are the two possible outcomes of an  $X$ -basis measurement on a qubit, and  $|y^+\rangle$  and  $|y^-\rangle$  are the two possible outcomes of a  $Y$ -basis measurement. The  $X$ ,  $Y$ , and  $Z$  gates, together with the trivial  $I$  gate that leaves a qubit's state unchanged, are known as the *Pauli gates*.

I will have more to say about this family of gates later. For now I want to point out that, for a system of qubits in the state given by Eq. (4.2), as long as both the  $A$  and  $B$  qubit are measured in one of the Pauli bases, the predicted statistics arising from those measurements will not violate the constraints that Bell's inequality imposes on local hidden-variable theories. In other words, if all that we have access to are measurement devices that can measure a qubit in one of the Pauli bases, then there will be no way to disprove some local hidden-variable story a skeptic might cook up to explain the observed statistics. To experimentally disprove such a story, we will need to have measurement devices that can measure in other measurement bases besides these ones, the ones for which Bell showed that a violation of the Bell inequalities will occur.

Bell's theorem should convince us that no local hidden-variable theory can reproduce the correlations arising from such measurements (i.e., from measurements in bases other than  $X$ ,  $Y$ , and  $Z$ ), but this does not bar a skeptic from positing a *nonlocal* hidden-variable theory to recover them instead. An example of such a theory is one in which the outcome of a measurement on  $A$  depends on the measurement basis used to measure  $B$ . But since  $A$  and  $B$  could in principle be far apart in space, and since we require the measurements to be performed simultaneously, the causal influence from  $B$  to  $A$  in such a theory will have to be propagated faster than the speed of light. This is a hard pill to swallow, but one might be inclined to swallow it anyway in order to avoid the other alternatives.<sup>36</sup>

<sup>35</sup>This state is identical to the one given in Eq. (2.11) but we repeat it here for convenience.

<sup>36</sup>For further discussion, see Myrvold, Genovese, & Shimony (2020).

Coming back to our discussion of quantum computers, recall that when we discussed the many-worlds explanation of quantum computing, we noted that the computer's state as given on the right-hand side of Eq. (3.1) is entangled. The question arises, then, as to what role is played by entanglement more generally in enabling quantum computers to outperform their classical competitors. There are two ways to think of this question. We might ask, on the one hand, whether realising an entangled state is *necessary* to enable a quantum speedup. On the other hand we might ask whether it is *sufficient* to enable it.

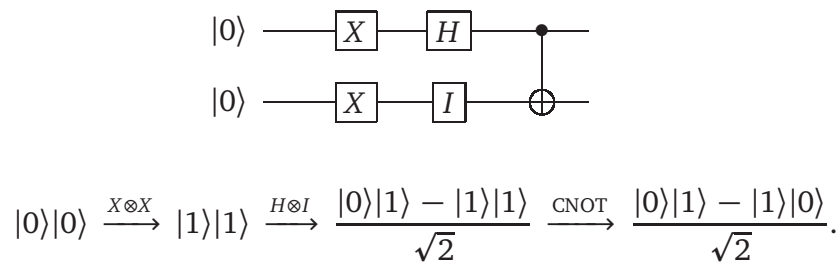
There has been some debate surrounding the first question. On the one hand, Jozsa & Linden (2003) have proven that realising an entangled state is necessary to enable speedup if a quantum computer can be assumed to be in a so-called 'pure state', i.e., in a state that represents a maximally specific description of the system from the point of view of the quantum-mechanical formalism. This is the case when we represent a system with a state vector as we have been doing up until now (and will continue to do). On the other hand it has been argued (Biham, Brassard, Kenigsberg, & Mor, 2004) that quantum computers that are in 'mixed states'—i.e., states which describe the computer in a less than maximally specific way, either because we are ignorant of the actual state of the computer, or because the computer is coupled to its environment—are in some cases capable of achieving a quantum speedup over their classical competitors without ever being entangled. There is insufficient space to discuss this (largely technical) debate here, but in Cuffaro (2013) I argue that the purported counter-examples to what I there call the "necessity of entanglement thesis" do not actually demonstrate what they purport to show, but instead clarify the necessary role that entanglement does play in quantum computation.

From the philosopher's point of view, the more interesting question in relation to the role of entanglement in quantum computing is the question regarding sufficiency, for as we will see, reflecting on this question sheds light on the tension between physical and computational ways of thinking that is at the heart of the science of quantum computing, the tension that is the primary source of the insight this new science brings into both physics and computation.

## 5 The Gottesman-Knill theorem

The main reason for being skeptical of the idea that entanglement suffices to enable quantum speedup is the Gottesman-Knill theorem (Gottesman, 1999). This theorem states that any quantum algorithm that employs (exclusively) some combination of the following operations (which together form what we will call the *Gottesman-Knill set*) is efficiently simulable by a classical computer: (i) The Clifford group of gates (see Figure 8), i.e., the  $X$ ,  $Y$ ,  $Z$ ,  $R$ ,  $H$ , and CNOT gates; (ii) Clifford group gates conditioned on the values of classical bits (indicating, e.g., the results of previous measurements); (iii) state preparation of a qubit in the computational basis (as one typically does for each qubit prior to the beginning of a computation); (iv) measurements in one of the Pauli bases (as one does at the end of a computation) (Nielsen & Chuang, 2000, §10.5.4).

The reason the theorem represents a challenge, to the view that entanglement suffices to enable a quantum computer to outperform a classical computer, is that by using only the operations given in (i)–(iv), which according to the theorem are all efficiently classically simulable, it is possible to generate an entangled state. Figure 9 depicts a quantum circuit to realise an entangled state using only operations from the Gottesman-Knill set. It begins with the preparation of two qubits in the computational basis state  $|0\rangle$ , then subjects them



**Figure 9:** A quantum circuit (above), written out explicitly in terms of the unitary transformations required to realise it (below), involving only operations from the Gottesman-Knill set. At the end of this sequence of operations the computer will be in an entangled state.

each to an  $X$ -gate, as a result of which they will both be in the state  $|1\rangle$ . The first qubit is then run through a Hadamard gate which transforms it into the superposition state  $(|0\rangle - |1\rangle)/\sqrt{2}$ , and following that both qubits are run through a CNOT gate. The resulting state of the computer is an entangled state.

It would be wrong to conclude (cf. Jozsa & Linden, 2003, pp. 2029–30), on the basis of this, that entanglement is not sufficient to enable a quantum computer to outperform a classical computer. Why? Well, let us reflect on what the Gottesman-Knill theorem is saying. The Gottesman-Knill theorem shows that there is a certain set of quantum operations that can be efficiently simulated on a classical computer. Let us then consider what we mean when we say that we have ‘efficiently classically simulated’ something. We can say, to start with, that in a classical computer simulation (efficient or otherwise), whatever computer is doing the simulating will be such that it can be given a classical description. As we saw in Section 2.1, a complete description of a system, in classical mechanics, is always factorisable into complete descriptions of all of its subsystems (see Figure 3). Besides this, classically describable systems and processes are locally causal, just like the convoluted procedure that Carsten uses to determine what to stuff in the envelopes he sends to Louise and Robert. The classical language of classical mechanics *constrains* all the descriptions of systems that can be given in that language to be this way. The behaviour of a classically describable system, in other words, is *always* expressible in terms of local causes and effects, and correlations between effects manifested by such a system’s subsystems can always be explained by appealing either to common or direct causes, whose influences propagate no faster than the speed of light. This is true no matter how large a system is conceived to be. For classical mechanics aims to be a universal language for describing *any* physical system, be it a single particle or the entire physical universe. The local causes of an effect will not always be known to us, but classical mechanics tells us that they are there, and that we can find them in principle.

Let us return to Carsten, our band leader who lives on the Île Saint-Louis in the centre of Paris. Carsten has lately become interested in quantum computers as a way of helping with composition,<sup>37</sup> and has acquired a shiny silver box, now sitting on his desk, which he claims is one. Quadeisha, a friend and fellow musician who has happened by, is skeptical, and asks him to show her how it works. Obliging, Carsten turns on his computer and has it run what he tells her is a calculation whose result depends on the prior generation of an entangled state. Quadeisha, still not satisfied, asks Carsten to describe the algorithm

<sup>37</sup>For a review of some of the uses envisioned for quantum computers in music, see Miranda (2021).

by which the computer arrives at its result. Upon discovering (after looking through the owner’s manual for the device) that the algorithm employs no operations outside of the Gottesman-Knill set, she refuses to believe that Carsten’s computer is quantum.

Quadeisha is right to be skeptical, *even if* the box on Carsten’s desk (which, let us assume, he does not know how to open) *actually* is a quantum computer.<sup>38</sup> Moreover she need not have heard of the Gottesman-Knill theorem to be justified in her skepticism. She need only be familiar with the work of John Bell. For it can be shown that the combined effect of any sequence of Gottesman-Knill operations, for any subsystem of the system to which they are applied, is equivalent to a measurement in one of the Pauli bases,  $X$ ,  $Y$ ,  $Z$ , on a system whose state is given by one of the basis vectors of a Pauli basis (Cuffaro, 2017, p. 115). Bell already showed us how to provide a locally causal model to reproduce the statistics of such measurements on a system in the state given in Eq. (4.2). Moreover his technique is straightforwardly extendable to other similarly entangled two-party systems (Cuffaro, 2017, §A.1), and further techniques have been devised for constructing efficient locally causal models to recover the statistics arising from Pauli-basis measurements on entangled systems composed of three or more parties as well (see Cuffaro, 2017, §A.2, which summarises a technique first described in Tessier, 2004).<sup>39</sup>

The upshot of the Gottesman-Knill theorem is that a certain number of quantum operations are efficiently classically simulable despite being capable of generating entangled states. And we have just seen that to say that some phenomenon is ‘efficiently classically simulable’ is to say that it can be efficiently redescribed in a locally causal way. Since any sequence of Gottesman-Knill operations is equivalent to a measurement in one of the Pauli bases,  $X$ ,  $Y$ ,  $Z$ , on a system whose state is given by one of the basis vectors of a Pauli basis, the Gottesman-Knill theorem is essentially telling us that there are some statistics associated with measurements on systems in entangled states that admit of a locally causal description.

We do not really need the Gottesman-Knill theorem to tell us this, for Bell’s and related inequalities amount to general constraints that any locally causal model of combined measurement statistics must satisfy, and we already know that in some cases, i.e., for measurements in one of the Pauli bases, the predictions of a locally causal model are compatible with quantum-mechanical predictions. It is therefore misleading to conclude, on the basis of the Gottesman-Knill theorem, that entanglement is not sufficient to enable quantum computational speedup. What the Gottesman-Knill theorem shows us is that one also needs to *use* this entanglement to its full potential, in particular by not restricting a quantum computer to use only a small proportion of the operations of which it is capable (i.e., by not restricting it to the Gottesman-Knill set).

In the context of more general discussions of quantum mechanics, there are those who claim that entanglement is *the* one and only distinguishing feature of the theory (Schrödinger, 1935b). This is a controversial claim, but it is not proved false by the fact that Bell (2004 [1964]) showed that some operations on some entangled states can be re-described in a locally causal way. Likewise, pointing out essentially the same thing in the context of quantum computation is not an objection to the claim that quantum entanglement constitutes a sufficient physical resource for realising speedup over classical computation. Indeed, the Gottesman-Knill theorem serves to highlight the role that is actually

---

<sup>38</sup>This is actually the attitude (and for similar reasons) that many theorists take toward certain claims by private corporations to have built a working quantum computer (see, e.g., Aaronson (2013a); Shin, Smith, Smolin, & Vazirani (2014)).

<sup>39</sup>The  $n$ -party case, for  $n \geq 3$ , introduces subtleties which we will discuss in the next section.

played by entanglement in a quantum computer and to clarify why and in what sense it is sufficient to preclude the computer's evolution from plausibly being classically simulated (for more on this, see Cuffaro, 2017).

## 6 Computer simulations and locally causal models of quantum phenomena

Above we mentioned that techniques have been developed for providing classical computer simulations, i.e., locally causal descriptions, to recover the statistics arising from Pauli-basis measurements on entangled systems whose subsystems are in superpositions of eigenstates of Pauli observables, for any number of parties. The classical computer simulations of three or more party systems are especially interesting in that they necessarily involve a small amount of communication. For three parties,  $A$ ,  $B$ , and  $C$ , one bit must be communicated between  $B$  and  $A$  in order to make the protocol work. For  $n$  parties,  $n - 2$  bits are required, i.e., the number of bits required scales linearly with  $n$  (Tessier, 2004). Note that this counts as being 'easy' in the complexity-theoretic sense we discussed in Section 2.2.

From a certain point of view, however, it may seem quite odd to call a classical computer simulation of quantum correlations, that involves a certain amount of communication between parties, a locally causal model of those correlations. Recall that, earlier, when we introduced Bell's theorem, we claimed that communication between distant parties in the context of a Bell experiment constitutes a *nonlocal* influence, i.e., one that would have to be communicated at a speed exceeding that of light.

We can start to resolve the tension here if we consider, first, just what a locally causal redescription of quantum measurement statistics is required to do. If a pair of qubits,  $A$  and  $B$ , are prepared in an entangled state and then spatially separated, quantum mechanics tells us that when, for instance, a  $Z$ -basis measurement is performed on  $B$ ,  $A$ 's state changes 'instantaneously', so that a  $Z$ -basis measurement on  $A$  will yield an outcome that is correlated with the result of the experiment on  $B$  in a particular way. Our challenge, however, is not to reproduce these apparent dynamics. Our challenge is to come up with an alternative description of the statistics that we *actually observe* in our experiments.

In order to assert that one has actually observed the results of a *combined* experiment on two or more spatially separated subsystems, it must be assumed that the results associated with the individual experiments on those subsystems have actually been combined in some way. That is, we must somehow gather together the results registered locally at the sites of the experiments on  $A$  and  $B$ . If Alice (a bassist) is assigned to measure  $A$ , for instance, and Bob (who plays the piano) to measure  $B$ , then once they record their individual results they need to communicate them, either to one another, or perhaps to Carsten who is sitting in his office on the Île Saint-Louis, where he is in the process of showing Quadeisha how his quantum computer is but one part of a network of spatially distributed quantum computers that are together performing a computation that depends on the prior generation of a system in an entangled state. While Alice and Bob are making their way toward Carsten and Quadeisha with their results, however, there is time for Alice and Bob to, surreptitiously, exchange further information with one another by local means (Bob might call Alice on his cell phone, for instance, as they are making their way toward Carsten's office) in order to coordinate the outcomes of their individual measurements and 'correct' them if necessary, unbeknownst to Quadeisha. And since this further information is propagated locally, i.e., at a speed no greater than that of light, it can be considered as *part of the common cause*

of the actually observed (combined) measurement event, and thus as part of a classical, locally causal, description of that measurement event (see Cuffaro, 2017, §4, fig. 1).

Now, given a particular combined calculational outcome at a time  $t$ , one can try to explain it in one of two ways: (i) Carsten's way, i.e., as the result of a *bona fide* quantum-mechanical process manifesting nonlocal behaviour, or (ii) Quadeisha's more skeptical way, i.e., as the result of a locally causal mechanism whose different computational components are communicating with one another at less than the speed of light. Without further evidence—if Carsten cannot or will not open his silver box—it is perfectly legitimate to side with Quadeisha here, and be a skeptic with respect to Carsten's quantum-mechanical description of the box's observed behaviour. Quadeisha's alternative account of the observed correlations is a perfectly good one in such a scenario, for her proposed classical alternative can solve the problem just as easily, in the complexity-theoretic sense, as Carsten's purported quantum computer.

Bell's and related inequalities specify *constraints* on what a locally causal description that accounts for a combined probability distribution must be like, but in and of themselves they are essentially little more than formal statements. Bell's (2004 [1964]) original inequality, for instance, is essentially just a theorem of probability (Bell 2004 [1981], Janas et al. Forthcoming, §3.4, Pitowsky 1989, Pitowsky 1994). This is no fault. But if we are to make a meaningful distinction between what is and is not ruled out by them, then we must consider the context within which they are being applied. Normally this does not need to be made explicit, for normally the context is what I have elsewhere called the *theoretical* context (Cuffaro, 2017, p. 107), wherein we are considering alternative theories of the natural world. In this context, Bell's inequalities help us to answer the question of whether there may be some deeper hidden-variable theory of the natural world underlying quantum mechanics, and what such a theory needs to be like.

They do not answer that question all by themselves, however. Any alternative hidden-variable theory of the natural world must do more than merely satisfy the constraints imposed by the Bell inequalities. It must also be *plausible*. Thus, besides reproducing the well-confirmed statistical predictions of quantum mechanics, any deeper candidate theory should also, for instance, be consistent with our other well-confirmed theories of physics such as special and general relativity (and if not, then a compelling reason will need to be given for why we should not worry about such contradictions). It is on the basis of such plausibility constraints, and not on the basis of the Bell inequalities, that we rule out many of the so-called 'loopholes' to Bell's theorem. Quadeisha's alternative account of the observed correlations in Carsten's office counts as a particularly implausible example of such a loophole in the theoretical context.<sup>40</sup>

But the theoretical context is not the only context that the Bell inequalities are relevant to. Nor is it the context appropriate to a discussion of the respective capabilities of quantum and classical computers. The context that is appropriate here is what I have elsewhere called the *practical* context (Cuffaro, 2017, p. 107). In the practical context what concerns us is not alternative theories of the natural world. What concerns us is what *we* are capable of *building* with the aim of reproducing the statistical predictions of quantum mechanics. The plausibility constraints that restrict us in the practical context are not the same as the ones that restrict us in the theoretical context. But there are plausibility constraints in the practical context nevertheless.

---

<sup>40</sup>Quadeisha's loophole is actually conceptually similar to the 'collapse locality' loophole at the heart of Adrian Kent's *causal quantum theory* (Kent, 2005). For discussion see Cuffaro (2017, pp. 104–106). For a more general discussion of the methodology of no-go theorems, see Dardashti (2021).



In particular, in the practical context, we should rule out, as implausible, alternative locally causal descriptions of systems that would be *too hard* (in the complexity-theoretic sense) for us to build. And we should rule in alternative locally causal descriptions of systems that can be built *easily* (again, in the complexity-theoretic sense). The complexity involved in the specification of such a system, in other words, needs to be *tractable*. If, for instance, Quadeisha's basis for being skeptical of Carsten's claim to own a quantum computer is that she can specify a locally causal model for his computer's calculation that requires a number of additional resources that scales exponentially, then Carsten will be justified in dismissing her skepticism, for he can with justice say that it is wildly implausible that anyone can have built a machine according to Beatrice's specifications, because simulating the quantum calculation would amount, for such a machine, to an intractable task (Cuffaro, 2017, §5).<sup>41</sup>

## 7 Computational perspectives on physics

In the previous section we discussed classical computer simulations of quantum-mechanical phenomena, and we made the point that such simulations can be thought of as locally causal alternative descriptions of those phenomena. This is not the way that they are normally framed in the literature on quantum information and computation. The way that it is normally put is that such simulations quantify the extent to which quantum mechanics differs from classical mechanics, by enumerating the number of resources required to account for quantum-mechanical phenomena using classical means (see, for instance Tessier 2004, Tessier, Caves, Deutsch, & Eastin 2005, Toner & Bacon 2003, Brassard, Cleve, & Tapp 1999, Rosset, Branciard, Gisin, & Liang 2013).

Framing their significance in this way is certainly not incorrect, and it can be very theoretically useful.<sup>42</sup> But from the philosopher's perspective there is value in, so to speak, calling a spade a spade. A description of a plausible classical computational model for efficiently simulating some of the operations that can be carried out by a quantum computer is the same kind of thing as a plausible alternative local hidden-variable account of some particular class of quantum phenomena, in the simple sense that in both cases these constitute plausible locally causal models of quantum phenomena. The difference between them lies in the way that one interprets the word 'plausible' in each case. In the practical context we assume that a system has been built by a rational agent to achieve a particular purpose. We do not presuppose this in the theoretical context.<sup>43</sup> Clearly, certain assumptions that would be regarded as plausible in the former context will not be in the latter and it is important to be on guard against conflating them. As we saw in the previous section, practical investigators attempting to isolate and/or quantify the computational resources provided by physical systems may be in danger of conceptual confusion if they are not cognisant of the differences between the practical and theoretical contexts. More positively, it must not be assumed by practical investigators that every 'no-go result', formulated in the

---

<sup>41</sup>Note that I am taking 'tractable' here in a relative sense. That is, the resources required by a classical computer to reproduce a particular effect should differ tractably from those required by a quantum computer. Or in other words: it should *not be essentially harder* for the classical system to produce the same effect as the quantum system.

<sup>42</sup>The general subject of classical simulations of quantum systems is an important and burgeoning area of modern physics (see, for example, Lee & Thomas 2002; Borges, Hor-Meyll, Huguenin, & Khoury 2010).

<sup>43</sup>This statement is not meant to express any sort of theological opinion. It is merely a statement about how science operates, at least in this century.

theoretical context, will constrain what they can accomplish in the practical context (Cufaro, 2018a). More generally it is important to think seriously about what one means by a ‘plausible locally causal model’ regardless of the context of investigation.

Calling a spade a spade gives us, moreover, insight into how the practical context can illuminate the investigations of the traditional philosopher of physics in the context of quantum mechanics. The silly story we told in the previous section about how Carsten, Alice, and Bob manage to reproduce, by communicating just a few classical bits, all of the correlational phenomena associated with Pauli-basis measurements on an  $n$ -party quantum system in an entangled state, besides representing just another silly classical story, emphasises perhaps the most fundamental difference that we can point to between classical and quantum mechanics. Classical and quantum mechanics are, fundamentally, alternative universal languages for describing physical systems (Janas et al., Forthcoming, §6.3). And by investigating the computational power inherent in those languages we gain insight into the respective logical structures of classical and quantum theory, into the logical structure of quantum mechanics that enables the efficient representation of correlational phenomena by quantum-mechanical systems, and the logical structure of classical mechanics which precludes this (Janas et al., Forthcoming, §4.2).

## 8 Concluding remarks

Reflecting on the lessons of quantum mechanics, Niels Bohr wrote that:

In representing a generalization of classical mechanics suited to allow for the existence of the quantum of action, Quantum mechanics offers a frame *sufficiently wide* to account for empirical regularities *which cannot* be comprised in the classical way of description (Bohr, 1948, p. 316, my emphasis).

This is a lesson with which the artist will be familiar.<sup>44</sup> From time to time, in literature, music, and in other forms of art, new methods of writing and composing emerge that allow one to express the subtleties and nuances of experience more easily than before. E. M. Forster (1942, p. 28) once said, about Virginia Woolf, that “She pushed the light of the English language a little further against the darkness.” In music, the (re-)introduction, into the Western musical tradition, of modality and chromaticism through Jazz, Blues, and other forms of popular music, and also through art music (see Vincent, 1951), has enabled modern composers in the Western tradition to explore musical landscapes and express aspects of our emotional and intellectual experience that composers limited to the major and minor scales, so dominant during Western music’s classical period, cannot easily capture. In his musings about quantum mechanics, Bohr himself was wont to appeal to an analogy with poetry. Werner Heisenberg recalls, in his memoirs, a conversation he had with Bohr in which the latter stated that:

We must be clear that, when it comes to atoms, language can be used only as poetry. The poet, too, is not nearly so concerned with describing facts as with creating images and establishing mental connections” (Heisenberg, 1971, p. 41).

---

<sup>44</sup>A similar point is made in Janas et al. (Forthcoming, sec. 6.3, note 22).

These and like analogies were not lost on actually practising poets and other artists of the same period (see Mairhofer, 2021). The subject of quantum mechanics has even been found to lend itself very naturally to the format of the comic book (see Bub & Bub, 2018, and the review by Cuffaro & Doyle, 2021, especially §4). As for the science of quantum computing: The science of quantum computing takes advantage, as we have seen throughout the course of this chapter, of the increased expressive power of the quantum-mechanical language. It shows us that using the new language of quantum mechanics, we can interact with what we take to be quantum-mechanical systems in ways that cannot be easily replicated classically.

In the course of this chapter we have reflected on the fundamentals of classical and quantum physics and computation, on the resources used and on the explanation of the power of quantum computers, and finally on the broader insights that can be gleaned from the science of quantum computing both for physics and for computation. The philosophical issues addressed in this chapter are not the only issues brought up by the science of quantum computing. But I hope to have convinced the reader, with this sample of some of the more central ones, of quantum computing's potential for illuminating the world that we are all participants in.

## References

- Aaronson, S. (2013a). D-Wave: Truth finally starts to emerge. Posted: June 05, 2013. Retrieved: August 11, 2014.  
URL [www.scottaaronson.com/blog/?p=1400](http://www.scottaaronson.com/blog/?p=1400)
- Aaronson, S. (2013b). *Quantum Computing Since Democritus*. New York: Cambridge University Press.
- Aaronson, S. (2013c). Why philosophers should care about computational complexity. In B. J. Copeland, C. J. Posy, & O. Shagrir (Eds.) *Computability: Turing, Gödel, Church, and Beyond*, (pp. 261–327). Cambridge, MA: MIT Press.
- Aaronson, S. (2016). Complexity zoo. [complexityzoo.uwaterloo.ca/](http://complexityzoo.uwaterloo.ca/)  
Complexity\_Zoo.
- Adlam, E. (2014). The problem of confirmation in the Everett interpretation. *Studies in History and Philosophy of Modern Physics*, 47, 21–32.
- Agrawal, M., Kayal, N., & Saxena, N. (2004). PRIMES is in P. *Annals of Mathematics*, 160, 781–793.
- Aharonov, D., van Dam, W., Kempe, J., Landau, Z., Lloyd, S., & Regev, O. (2007). Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM Journal on Computing*, 37, 166–194.
- Albert, D., & Loewer, B. (1988). Interpreting the many worlds interpretation. *Synthese*, 77, 195–213.
- Arora, S., & Barak, B. (2009). *Computational Complexity: A Modern Approach*. Cambridge: Cambridge University Press.
- Bell, J. S. (2004 [1964]). On the Einstein-Podolsky-Rosen paradox. In *Speakable and Un-speakable in Quantum Mechanics*, (pp. 14–21). Cambridge: Cambridge University Press.

- Bell, J. S. (2004 [1966]). On the problem of hidden variables in quantum mechanics. In *Speakable and Unsayable in Quantum Mechanics*, (pp. 1–13). Cambridge: Cambridge University Press.
- Bell, J. S. (2004 [1981]). Bertlmann’s socks and the nature of reality. In *Speakable and Unsayable in Quantum Mechanics*, (pp. 139–158). Cambridge: Cambridge University Press.
- Bell, M., & Gao, S. (Eds.) (2016). *Quantum Nonlocality and Reality*. Cambridge: Cambridge University Press.
- Bennett, C. H., Bernstein, E., Brassard, G., & Vazirani, U. (1997). Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26, 1510–1523.
- Bernstein, E., & Vazirani, U. (1997). Quantum complexity theory. *SIAM Journal on Computing*, 26, 1411–1473.
- Biham, E., Brassard, G., Kenigsberg, D., & Mor, T. (2004). Quantum computing without entanglement. *Theoretical Computer Science*, 320, 15–33.
- Bohr, N. (1948). On the notions of causality and complementarity. *Dialectica*, 2, 312–319.
- Boole, G. (1847). *The Mathematical Analysis of Logic*. New York: Philosophical Library. Reprinted: Bristol: Thoemmes Press, 1998.
- Borges, C. V. S., Hor-Meyll, M., Huguenin, J. A. O., & Khoury, A. Z. (2010). Bell-like inequality for the spin-orbit separability of a laser beam. *Physical Review A*, 82, 033833.
- Brassard, G., Cleve, R., & Tapp, A. (1999). Cost of exactly simulating quantum entanglement with classical communication. *Physical Review Letters*, 83, 1874–1877.
- Briegel, H. J., Browne, D. E., Dür, W., Raussendorf, R., & den Nest, M. V. (2009). Measurement-based quantum computation. *Nature Physics*, 5, 19–26.
- Bub, J. (2006). Quantum computation from a quantum logical perspective. arXiv:quant-ph/0605243v2.
- Bub, J. (2008). Quantum computation and pseudotelepathic games. *Philosophy of Science*, 75, 458–472.
- Bub, J. (2010). Quantum computation: Where does the speed-up come from? In A. Bokulich, & G. Jaeger (Eds.) *Philosophy of Quantum Information and Entanglement*, (pp. 231–246). Cambridge: Cambridge University Press.
- Bub, J. (2016). *Bananaworld, Quantum Mechanics for Primates*. Oxford: Oxford University Press.
- Bub, J., & Pitowsky, I. (2010). Two dogmas about quantum mechanics. In S. Saunders, J. Barrett, A. Kent, & D. Wallace (Eds.) *Many Worlds? Everett, Quantum Theory, and Reality*, (pp. 433–459). Oxford: Oxford University Press.
- Bub, T., & Bub, J. (2018). *Totally Random: Why Nobody Understands Quantum Mechanics*. Princeton: Princeton University Press.

- Church, A. (1936). An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58, 345–363.
- Cobham, A. (1965). The intrinsic computational difficulty of functions. In Y. Bar-Hillel (Ed.) *Logic, Methodology and Philosophy of Science: Proceedings of the 1964 International Congress*, (pp. 24–30). Amsterdam: North-Holland.
- Cook, W. J. (2012). *In Pursuit of the Traveling Salesman: Mathematics at the Limits of Computation*. Princeton: Princeton University Press.
- Copeland, J. B. (2017). The modern history of computing. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2017 ed.
- Copeland, J. B. (2020). The Church-Turing thesis. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2020 ed.
- Cuffaro, M. E. (2012). Many worlds, the cluster-state quantum computer, and the problem of the preferred basis. *Studies in History and Philosophy of Modern Physics*, 43, 35–42.
- Cuffaro, M. E. (2013). On the necessity of entanglement for the explanation of quantum speedup. arXiv:1112.1347v5.
- Cuffaro, M. E. (2017). On the significance of the Gottesman-Knill theorem. *The British Journal for the Philosophy of Science*, 68, 91–121.
- Cuffaro, M. E. (2018a). Reconsidering no-go-theorems from a practical perspective. *The British Journal for the Philosophy of Science*, 69, 633–655.
- Cuffaro, M. E. (2018b). Universality, invariance, and the foundations of computational complexity in the light of the quantum computer. In S. O. Hansson (Ed.) *Technology and Mathematics: Philosophical and Historical Investigations*. Cham: Springer.
- Cuffaro, M. E. (2020). Information causality, the Tsirelson bound, and the ‘being-thus’ of things. *Studies in History and Philosophy of Modern Physics*, 72, 266–277.
- Cuffaro, M. E., & Doyle, E. P. (2021). Essay review of Tanya and Jeffrey Bub’s *Totally Random: Why Nobody Understands Quantum Mechanics: A Serious Comic on Entanglement*. *Foundations of Physics*, 51, 28:1–28:16.
- Cuffaro, M. E., & Hartmann, S. (2021). The open systems view. In preparation.
- Curiel, E. (2014). Classical mechanics is Lagrangian; it is not Hamiltonian. *The British Journal for Philosophy of Science*, 65, 269–321.
- Dardashti, R. (2021). No-go theorems: What are they good for? *Studies in History and Philosophy of Science*, 86, 47–55.
- Davis, M. (2000). *The Universal Computer: The Road from Leibniz to Turing*. New York: W. W. Norton and Company.
- Dawid, R., & Thébault, K. P. Y. (2015). Many worlds: Decoherent or incoherent? *Synthese*, 192, 1559–1580.

- Dawson Jr., J. W. (2007). Classical logic's coming of age. In D. Jacquette (Ed.) *Philosophy of Logic*, (pp. 497–522). Amsterdam: Elsevier.
- Dean, W. (2016). Squeezing feasibility. In A. Beckmann, L. Bienvenu, & N. Jonoska (Eds.) *Pursuit of the Universal: Proceedings of the 12th Conference on Computability in Europe*, (pp. 78–88). Cham: Springer International Publishing.
- Deutsch, D. (1985). Quantum theory, the Church-Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 400, 97–117.
- Deutsch, D. (1989). Quantum computational networks. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 425, 73–90.
- Deutsch, D. (1997). *The Fabric of Reality*. New York: Penguin.
- Deutsch, D. (2010). Apart from universes. In S. Saunders, J. Barrett, A. Kent, & D. Wallace (Eds.) *Many Worlds? Everett, Quantum Theory, and Reality*, (pp. 542–552). Oxford: Oxford University Press.
- DeWitt, B., & Graham, N. (1973). *The Many-Worlds Interpretation of Quantum Mechanics*. Princeton: Princeton University Press.
- DeWitt, B. S. (1973 [1971]). The many-universes interpretation of quantum mechanics. In DeWitt & Graham (1973, pp. 167–218).
- Duwell, A. (2007). The many-worlds interpretation and quantum computation. *Philosophy of Science*, 74, 1007–1018.
- Duwell, A. (2018). How to make orthogonal positions parallel: Revisiting the quantum parallelism thesis. In M. E. Cuffaro, & S. C. Fletcher (Eds.) *Physical Perspectives on Computation, Computational Perspectives on Physics*, (pp. 83–102). Cambridge: Cambridge University Press.
- Duwell, A. (2021). *Computation and Physics*. Cambridge: Cambridge University Press. Forthcoming.
- Edmonds, J. (1965). Paths, trees, and flowers. *Canadian Journal of Mathematics*, 17, 449–467.
- Einstein, A., Podolsky, B., & Rosen, N. (1935). Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 47, 777–780.
- Everett III, H. (1956). The theory of the universal wave function. In DeWitt & Graham (1973, pp. 3–140).
- Farhi, E., Goldstone, J., Gutmann, S., & Sipser, M. (2000). Quantum computation by adiabatic evolution. Tech. Rep. MIT-CTP-2936, MIT. arXiv:quant-ph/0001106.
- Fletcher, S. C. (2018). Computers in abstraction / representation theory. *Minds & Machines*, 28, 445–463.
- Forster, E. M. (1942). *Virginia Woolf. The Rede Lecture 1941*. Cambridge: Cambridge University Press.

- Genovese, M. (2016). Experimental tests of Bell's inequalities. In Bell & Gao (2016, pp. 124–140).
- Gödel, K. (1956). Private letter to John von Neumann, 20 March 1956. Translated by Wensinger in: Sipser (1992).
- Goldreich, O. (2008). *Computational Complexity: A Conceptual Perspective*. Cambridge: Cambridge University Press.
- Gottesman, D. (1999). The Heisenberg representation of quantum computers. In S. P. Corneey, R. Delbourgo, & P. D. Jarvis (Eds.) *Group22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, (pp. 32–43). Cambridge, MA: International Press. Longer version available at: arXiv:quant-ph/9807006v1.
- Greaves, H., & Myrvold, W. (2010). Everett and evidence. In S. Saunders, J. Barrett, A. Kent, & D. Wallace (Eds.) *Many Worlds? Everett, Quantum Theory, and Reality*, (pp. 181–205). Oxford: Oxford University Press.
- Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC '96, (pp. 212–219). New York, NY, USA: Association for Computing Machinery.
- Hagar, A. (2007). Quantum algorithms: Philosophical lessons. *Minds & Machines*, 17, 233–247.
- Hagar, A., & Cuffaro, M. (2019). Quantum computing. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2019 ed.
- Heisenberg, W. (1971). *Physics and Beyond*. New York: Harper & Row.
- Hewitt-Horsman, C. (2009). An introduction to many worlds in quantum computation. *Foundations of Physics*, 39, 869–902.
- Horsman, D., Kendon, V., & Stepney, S. (2018). Abstraction/representation theory and the natural science of computation. In M. E. Cuffaro, & S. C. Fletcher (Eds.) *Physical Perspectives on Computation, Computational Perspectives on Physics*, (pp. 127–152). Cambridge: Cambridge University Press.
- Howard, D. (1989). Holism, separability, and the metaphysical implications of the Bell experiments. In J. T. Cushing, & E. McMullin (Eds.) *Philosophical Consequences of Quantum Theory*, (pp. 224–253). Notre Dame: University of Notre Dame Press.
- Hughes, R. I. G. (1989). *The Structure and Interpretation of Quantum Mechanics*. Cambridge, MA.: Harvard University Press.
- Janas, M., Cuffaro, M. E., & Janssen, M. (Forthcoming). *Understanding Quantum Raffles: Quantum Mechanics on an Information-Theoretic Approach: Structure and Interpretation*. Springer.
- Jozsa, R., & Linden, N. (2003). On the role of entanglement in quantum-computational speed-up. *Proceedings of the Royal Society of London. Series A. Mathematical, Physical and Engineering Sciences*, 459, 2011–2032.

- Kent, A. (2005). Causal quantum theory and the collapse locality loophole. *Physical Review A*, 72, 012107.
- Lee, K. F., & Thomas, J. E. (2002). Experimental simulation of two-particle quantum entanglement using classical fields. *Physical Review Letters*, 88, 097902.
- Lehner, C. (1997). What it feels like to be in a superposition. And why. *Synthese*, 110, 191–216.
- Lenstra, A. K., Lenstra Jr., H. W., Manasse, M. S., & Pollard, J. M. (1990). The number field sieve. In *Proceedings of the Twenty-second Annual ACM Symposium on Theory of Computing*, STOC '90, (pp. 564–572). New York, NY: Association for Computing Machinery.
- Lupacchini, R. (2018). Church's thesis, Turing's limits, and Deutsch's principle. In M. E. Cuffaro, & S. C. Fletcher (Eds.) *Physical Perspectives on Computation, Computational Perspectives on Physics*, (pp. 60–82). Cambridge: Cambridge University Press.
- Mairhofer, L. (2021). *Atom und Individuum: Bertolt Brechts Interferenz mit der Quantenphysik*. Berlin: De Gruyter.
- Maroney, O. J. E., & Timpson, C. G. (2018). How is there a physics of information? On characterising physical evolution as information processing. In M. E. Cuffaro, & S. C. Fletcher (Eds.) *Physical Perspectives on Computation, Computational Perspectives on Physics*, (pp. 103–126). Cambridge: Cambridge University Press.
- Martin, J. C. (1997). *Introduction to Languages and the Theory of Computation*. New York: McGraw-Hill, second ed.
- Maudlin, T. (2011). *Quantum Non-Locality and Relativity*. Cambridge, MA: Wiley-Blackwell, third ed.
- Mehlhorn, K., & Sanders, P. (2008). *Algorithms and Data Structures*. Berlin: Springer.
- Mermin, D. N. (2007). *Quantum Computer Science: An Introduction*. Cambridge University Press.
- Miranda, E. R. (2021). Quantum computer: Hello, music! In E. R. Miranda (Ed.) *Handbook of Artificial Intelligence for Music: Foundations, Advanced Approaches, and Developments for Creativity*. Cham: Springer. Forthcoming.
- Myrvold, W., Genovese, M., & Shimony, A. (2020). Bell's theorem. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2020 ed.
- Nielsen, M. A. (2006). Cluster-state quantum computation. *Reports on Mathematical Physics*, 57, 147–161.
- Nielsen, M. A., & Chuang, I. L. (2000). *Quantum Computation and Quantum Information*. Cambridge: Cambridge University Press.
- Nishimura, H., & Ozawa, M. (2009). Perfect computational equivalence between quantum turing machines and finitely generated uniform quantum circuit families. *Quantum Information Processing*, 8, 13–24.



- Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press, 2nd ed.
- Pitowsky, I. (1989). *Quantum Probability — Quantum Logic*. Hemsbach: Springer.
- Pitowsky, I. (1994). George Boole's 'conditions of possible experience' and the quantum puzzle. *British Journal for the Philosophy of Science*, 45, 99–125.
- Pitowsky, I. (2002). Quantum speed-up of computations. *Philosophy of Science*, 69, S168–S177.
- Raussendorf, R., & Briegel, H. J. (2002). Computational model underlying the one-way quantum computer. *Quantum Information and Computation*, 2, 443–486.
- Raussendorf, R., Browne, D. E., & Briegel, H. J. (2003). Measurement-based quantum computation on cluster states. *Physical Review A*, 68, 022312.
- Rivest, R. L., Shamir, A., & Adleman, L. (1978). A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21, 120–126.
- Rosset, D., Branciard, C., Gisin, N., & Liang, Y.-C. (2013). Entangled states cannot be classically simulated in generalized Bell experiments with quantum inputs. *New Journal of Physics*, 15, 053025.
- Saunders, S. (1995). Time, quantum mechanics, and decoherence. *Synthese*, 102, 235–266.
- Schrödinger, E. (1935a). Die gegenwärtige situation in der quantenmechanik. *Naturwissenschaften*, 23, 807–812; 823–828; 844–849. Translated in: John D. Trimmer (1980) in *Proceedings of the American Philosophical Society* 124, 323–338.
- Schrödinger, E. (1935b). Discussion of probability relations between separated systems. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31, 555–563.
- Shetterly, M. L. (2016). *Hidden Figures*. New York: HarperCollins.
- Shin, S. W., Smith, G., Smolin, J. A., & Vazirani, U. (2014). How "quantum" is the D-Wave machine? arXiv:1401.7087v2.
- Shor, P. W. (1994). Algorithms for quantum computation: Discrete logarithms and factoring. *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*, (pp. 124–134).
- Sipser, M. (1992). The history and status of the P versus NP question. In *Proceedings of the Twenty-fourth Annual ACM Symposium on Theory of Computing, STOC '92*, (pp. 603–618). New York, NY: Association for Computing Machinery.
- Tessier, T. E. (2004). *Complementarity and Entanglement in Quantum Information Theory*. Ph.D. thesis, The University of New Mexico, Albuquerque, New Mexico.
- Tessier, T. E., Caves, C. M., Deutsch, I. H., & Eastin, B. (2005). Optimal classical-communication-assisted local model of  $n$ -qubit Greenberger-Horne-Zeilinger correlations. *Physical Review A*, 72, 032305.

- Timpson, C. G. (2013). *Quantum Information Theory & the Foundations of Quantum Mechanics*. Oxford: Oxford University Press.
- Toner, B. F., & Bacon, D. (2003). Communication cost of simulating Bell correlations. *Physical Review Letters*, 91, 187904.
- Turing, A. M. (1936-7). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society. Second Series*, s2-42, 230–265.
- Vaidman, L. (1998). On schizophrenic experiences of the neutron or why we should believe in the many-worlds interpretation of quantum theory. *International Studies in the Philosophy of Science*, 12, 245–261.
- Vaidman, L. (2012). Probability in the many-worlds interpretation of quantum mechanics. In Y. Ben-Menahem, & M. Hemmo (Eds.) *Probability in Physics*, (pp. 299–311). Berlin: Springer.
- Vaidman, L. (2018 [2002]). Many-worlds interpretation of quantum mechanics. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2018 ed. First published: 2002.
- van Emde Boas, P. (1990). Machine models and simulations. In J. van Leeuwen (Ed.) *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity*, (pp. 1–66). Cambridge, MA: MIT Press/Elsevier.
- Vincent, J. (1951). *The Diatonic Modes in Modern Music*. New York: Mills Music. Page references to second edition (1974) published by Curlew Music Publishers, Hollywood, CA.
- Wallace, D. (2003). Everett and structure. *Studies in History and Philosophy of Modern Physics*, 34, 87–105.
- Wallace, D. (2007). Quantum probability from subjective likelihood: Improving on Deutsch’s proof of the probability rule. *Studies in History and Philosophy of Modern Physics*, 38, 311–332.
- Wallace, D. (2010). Decoherence and ontology. In S. Saunders, J. Barrett, A. Kent, & D. Wallace (Eds.) *Many Worlds? Everett, Quantum Theory, and Reality*, (pp. 53–72). Oxford: Oxford University Press.
- Wallace, D. (2012). *The Emergent Multiverse*. Oxford: Oxford University Press.
- Wallace, D. (2019). On the plurality of quantum theories: Quantum theory as a framework, and its implications for the quantum measurement problem. In S. French, & J. Saatsi (Eds.) *Realism and the Quantum*, (pp. 78–102). Oxford: Oxford University Press.
- Wikipedia contributors (2020). Instructions per second. In *Wikipedia, The Free Encyclopedia*. Posted on 15 September, 2020, 11:15.
- Zurek, W. H. (2003 [1991]). Decoherence and the transition from quantum to classical – revisited. arXiv:quant-ph/0306072v1.