

THE PHILOSOPHY OF QUANTUM INFORMATION

Chris Timpson

Quantum information theory is an exciting and still youthfully vigorous area of research which lies at the intersection of quantum physics, communication theory, and the theory of computation. Its point of departure is to seize upon the very marked and peculiar ways in which the quantum world differs from our classical conceptions of physics, and to see in these differences *opportunities* for new forms of communication and computation. The oddity of the quantum world, in this approach, is not seen as a potentially troublesome conundrum, best avoided or treated in as classical-like a way as possible, but rather as something which we may be able to harness to our advantage, to do things which we would not otherwise be able to do, or to do things in ways in which we would not otherwise be able to do them. The defining, strikingly non-classical, and conceptually puzzling quantum features such as *superposition*, *entanglement*, and *non-commutativity* (we shall see more of each of these notions in due course) are positively embraced, and put to concrete work.

The field primarily began to emerge in the early- to mid- 1980s, with the work of Deutsch and others (Benioff 1980, Feynman 1982, Deutsch 1985) on the concept of quantum computers – a distinctively *quantum* approach to information processing. In the 1990s, the concept of quantum information proper was introduced by Schumacher (1995) in his development and extension to the quantum realm of Shannon's paradigm of information theory [CROSSREF] (this was Schumacher's *quantum noiseless coding theorem*, see also Barnum *et al.* 1996). In addition, the first protocols were developed which unequivocally involved the transmission of quantum information proper (Bennett *et al.* 1993), whilst Shor (1994) showed that there exists an efficient quantum algorithm for an important computational problem – factoring large numbers – for which it is universally believed that there is no efficient classical algorithm: thus the claim that quantum computers are 'exponentially more powerful' than classical ones. By 2000 and the publication of the canonical textbook of Nielsen and Chuang, quantum information theory had reached a mature stage. (Nielsen and Chuang is now in its second edition (2010) and is still canonical.) Even so, however, the general view is that in some ways, quantum information scientists have only scratched the surface so far. There is still much to be done to reach a general understanding of the ways in which the quantum world differs from the classical world, to understand how these differences might be harnessed for interesting information-processing and communication ends, and to settle how these quantum-classical differences should be understood at the fundamental conceptual level.

The *philosophy* of quantum information – the topic of this chapter – has three main components or tasks. The first task is that of seeking to understand the nature and content of quantum information theory as it currently stands, and of seeking to highlight or to resolve any conceptual puzzles internal to the theory. The second task is that of reflecting on how the existence and striking success of quantum information theory affects our understanding of information - or *theories* of information - more generally: what does the advent of quantum information mean for the concept of information? The third task is that of exploring whether and how developments in quantum information theory affect the traditional, well-worn, conceptual problems in the foundations of quantum theory itself, for many have seen suggestive avenues here.

In this chapter I shall focus on the first two of these questions (for the third, see Timpson 2013). But first we need to begin by getting some grip on the characteristic features of quantum theory which have such an important role to play.

Features of Quantum Theory

The starting place is with the concept of *superposition*. In classical physics it is always the case that, for whatever physical property one considers, a given system either determinately does, or determinately does not, have a particular value of that property. We might think of energy, position, momentum, angular momentum, and so-on: the *dynamical* variables – those that are susceptible to change over time – as opposed to the fixed properties such as charge or mass. We can put this in terms of the features of the *states* of the system: the state of a system at a given time fixes *numerically exact* and *ontologically determinate* values for all the physical properties of the system at that time. Conversely, a specification of all of the values of the physical properties uniquely determines the state. We might be *ignorant* of what true state a system is in, in which case we might work with a probability distribution over the possible states, but there is always a fact about what the true state of a system is, and this true state is one in which all the physical values are determinate. In quantum physics, by contrast, this is not so.

Let us imagine a single particle for which there are only two possibilities of where it can be located. Either in a box located at the origin of some coordinate system set-up in the lab (call it 'Box 0' for 0 on the x-axis), or in a box one unit of length away along the x-axis (call it 'Box 1' for 1 on the x-axis). (Suppose the boxes to be of relatively small extent compared to the distance separating them.) Then in classical physics – and, one might be inclined to add, in common sense – at any time, the particle will either determinately be in Box 0, or it will determinately be in Box 1 – it has definite value 0 or definite value 1 for position on the x-axis. Certainly, whenever we open a box to look (whenever we 'perform a position measurement') we will find the particle to be in one or other location. But classical physics (and common sense) in addition insists that when we open a box we are merely revealing a position which was already determinate. If there was to begin with some probability of the particle being in Box 0 and some probability of it being in Box 1, then that was just down to our ignorance. Really, it was (must!) be in one or the other.

In quantum mechanics, however, as well as the two possibilities of the particle's determinately being in Box 0 and its determinately being in Box 1, it is also perfectly possible that it should be *neither* determinately in 0, *nor* determinately in 1: there can simply be no fact of the matter either way. (This is the example of Schrödinger's cat, where the cat is neither alive nor dead, but in some sense is both.) Yet even so, if we look, we will only ever find the particle to be in one or other box, so it's not as if the particle might actually be determinately in some third box instead, or that it might have split in two. We need more states in the theory, to correspond to these kinds of cases. Label the state corresponding to definitely being in Box 0, ' $|0\rangle$ ', and label the state corresponding to definitely being in Box 1, ' $|1\rangle$ '. Then there is a further family of states which correspond to not having either position determinately, and in these circumstances, the system is said to be in a *superposition* with respect to spatial location. States of this kind are given by a linear combination of the two definite position states, of the form:

$$a|0\rangle + b|1\rangle.$$

The coefficients a and b are complex numbers, whose moduli squared sum to one ($|a|^2 + |b|^2 = 1$). The number $|a|^2$ gives the probability that the particle will be found in Box 0 when we look, the number $|b|^2$ gives the probability that the particle will be found in Box 1 when we look. With probability one, therefore, the particle will be found to be in one or other box when we look. Crucially, however, for states of this kind, these probabilities are not understood as stemming from ignorance of a pre-determined fact of where the particle was located: in these states, there are no such facts, according to quantum mechanics.

Now imagine that we introduce a further pair of boxes: Box 2 sitting at two units along the x-axis, and Box 3 sitting at three units along the x-axis. Suppose that our first particle remains constrained only to be found in Box 0 or 1, and that we introduce a second particle, which will only be found in Box 2 or Box 3. We can now think about correlations between the positions of our two particles. Suppose that our boxes are prepared in such a way that whenever the first particle is found to have x-position 0, the second particle is found to have x-position 2; and that whenever the first particle is found to have x-position 1, the second particle is found to have x-position 3. In other words, the two particles are always found two units of position apart. There are two importantly different scenarios in which correlations such as this might arise.

In the first scenario, there is a definite value for x-position for particle 1 (it is either definitely in Box 0 or definitely in Box 1) and there is a definite x-position for particle 2 (it is either definitely in Box 2 or definitely in Box 3): the correlations arise simply because these definite values have been arranged to coincide in a particular manner. Either the true state of the pair of particles is $|0\rangle|2\rangle$, i.e., the first particle is sitting in Box 0 and the second particle is sitting in Box 2; or the true state of the pair is $|1\rangle|3\rangle$, i.e., the first particle is sitting in Box 1 and the second in Box 3. We may not know which particle is where – in which case we will introduce a probability distribution over the two states – but each particle is somewhere, and wherever it is, it is two units away from the other particle.

Much more interesting is the scenario in which there is no definite value for x-position for either particle: each is in a superposition with respect to position. Remarkably, it is still possible in this case for the stated correlations to exist. In a state like

$$a|0\rangle|2\rangle + b|1\rangle|3\rangle$$

we have superposed the two states considered in the previous paragraph. (Take neither a nor b to be zero.) Now there is no fact about where particle 1 is, and no fact about where particle 2 is, yet *it is still the case* – it is a fact - *that the particles are a distance of two units apart!* This is an extraordinary state of affairs. The two particles can definitely be two units apart, without either even *having* a position itself. States of this kind are called *entangled* states, and they all possess this striking feature of holism. When one's systems of interest are in an entangled state, there are facts about the properties of the whole joint system which are not reducible to facts about properties possessed by the individual systems making up the whole.

So far we have only considered the simplest form of measurement procedure on our particles, namely, opening a box to see whether a particle is inside. There also exist more sophisticated possibilities. First, notice a very important feature of the two states we started with, $|0\rangle$ and $|1\rangle$. It is possible to distinguish between these two states perfectly in a single-shot measurement. That is, if a particle is definitely in one or other state to begin with, but we don't know which, then we have a measurement procedure (opening one or other box, or both) which, in one go, will allow us to tell what that state was. Compare this with the pair of states $|0\rangle$ and $a|0\rangle + b|1\rangle$. If our particle was definitely prepared in one or other of *these* states to begin with, then there is no measurement procedure which applied on a single occasion is guaranteed to determine which state the particle was actually prepared in. For example, if we opened Box 0 and found the particle there, then this could have arisen either because the initial state was $|0\rangle$, or because it was $a|0\rangle + b|1\rangle$. Both options give a non-zero probability to this measurement outcome, so we can't distinguish between the two possibilities from this piece of data. If we had opened Box 0 and found instead that the particle *wasn't* there, then this would tell us that the initial state was certainly $a|0\rangle + b|1\rangle$ and not $|0\rangle$, but we are not guaranteed, when we perform the measurement once, that this will be the outcome we see: it only occurs with some probability less than one.

States, like $|0\rangle$ and $|1\rangle$, between which it is possible to distinguish perfectly (i.e., with probability 1) in a single-shot measurement are called *orthogonal*. States like $|0\rangle$ and $a|0\rangle + b|1\rangle$, where this is not possible, are called *non-orthogonal*.

It turns out that whilst superposition states of the form $a|0\rangle + b|1\rangle$ are (for a, b non-zero) not orthogonal to $|0\rangle$, nor indeed to $|1\rangle$, there is, for each of them, *some other* superposition state to which it *is* orthogonal. In other words, there is some measurement scheme which will perfectly distinguish between suitably chosen pairs of these superposition states (such schemes will need to be more sophisticated than just opening a box and looking). Sets of mutually orthogonal states are said to correspond to *observable quantities* in quantum mechanics. (*Position on the x-axis* is the observable quantity which goes with $|0\rangle$ and $|1\rangle$, of course.) Two observable quantities are said to *commute* when each member of the set of orthogonal states associated with one quantity is either identical with or orthogonal to any member of the set of orthogonal states associated with the other quantity. If two quantities do not commute, then when a system is in a state which gives it a definite value of one quantity, it will be in a superposition with respect to, and thus lack a definite value for, the other quantity. Furthermore, when two quantities do not commute, they cannot both be measured at the same time.

The summary is that the space of possible states in quantum theory is much richer than the space of states in a classical theory. In a classical theory of our particle in a box, the only states we have available are 'position is 0', 'position is 1' and statistical mixtures of the two (probability distributions representing our ignorance of what the actual state is). In a quantum theory for our particle in a box, we not only have the two definite position states, but all the superposition states as well - and statistical mixtures of all these too if we are ignorant of what actual state may have been prepared - whilst to all orthogonal pairs of states there corresponds an observable quantity, which quantities cannot all take definite values at the same time. Finally, when we imagine combining two systems, quantum theory, unlike a classical theory, allows there to be states of the whole which are not just products of the states of the parts. In other words, there are properties of the whole which are not determined by (do not supervene on) the properties of the parts.

Before moving on we should say something briefly about dynamics – the rules for evolution over time of the states of the system. In quantum mechanics there are (or there are apparently!) two rules for evolution over time. For closed systems, the evolution of the state is given by the so-called *unitary* evolution derived from the Schrödinger equation: it is deterministic and continuous. For systems when they are observed – when measurements are made – the evolution is indeterministic and discontinuous. When you open a box, the system jumps from $a|0\rangle + b|1\rangle$ to either $|0\rangle$ or $|1\rangle$, with corresponding probabilities $|a|^2$ and $|b|^2$, respectively. How these two apparently incompatible rules for time-evolution are to be reconciled – or if they can be - is perhaps (once we have got used to the idea of superposition) the deepest of the traditional conceptual problems of quantum mechanics.ⁱ

Quantum Information Theory

At least one way of thinking about what an information theory in general is, is as a theory describing what kinds of transformations – what processes, or what protocols – are possible given a certain defined range of resources.

Quantum information theory is a development of classical Shannon information theory which introduces new communication primitives – the quantum bit or *qubit*, and shared entanglement;

which explores what can be done with these new primitives; and which generalises Shannon's notions of information source and information channel to the quantum realm.

A *classical* bit is any two-state (classical) physical system, such as a classical ball placed in one or other of our boxes 0 or 1, or two distinct voltage levels in an electric circuit. Generically we label the two possible states of a bit with the Boolean logical values 0 or 1. We also use talk of bits to characterise the *quantity* of Shannon information produced by a (classical) source: the amount of (Shannon) information produced by a source is defined to be the number of bits that would be required to encode the output of the source optimally (Shannon 1948). The maximum amount of information that one physical bit can contain is therefore one bit's worth. The reason we are interested in a count in bits is because we are interested in the number of distinct possible messages that a Shannon information source will (over a suitably long run) produce with non-negligible probability, and we are therefore interested in the number of distinct states of some encoding medium or some transmission medium that would be required in order to allow us to differentiate between which of these possible messages was actually produced by the source on a given occasion. Suppose there are m such distinct messages for a given source, then m distinguishable states of the transmission or encoding medium will be required. If $m=2^n$ (or approximately so) then we can readily count the number of states in terms of the number of some standard physical resources: the number of two-state systems. The number required in this example is of course n bits.ⁱⁱ

A *quantum* bit – a *qubit* – is, in perfect parallel with the case of the classical bit, any two state *quantum* system, i.e., any system possessing two distinguishable (i.e., orthogonal) quantum states. For example – a quantum particle in our boxes 0 and 1 from earlier, or an atom with two particular energy states of interest, or a photon with horizontal or vertical polarisation. The two states are conventionally labelled (as earlier) $|0\rangle$ and $|1\rangle$, and are often termed the *computational basis states*. They are chosen to be states of the system which can readily be prepared (it is relatively easy to put the system controllably into one or other of the states) and which can readily be measured. However, in marked contrast to the case of the classical bit, there will of course be *very many more* distinct states that a qubit could be in than just $|0\rangle$ or $|1\rangle$. All of the superposition states $a|0\rangle + b|1\rangle$ are available too. Since the coefficients a and b can vary continuously, there are in fact, even for the humble qubit, the simplest possible instance of a quantum system, an infinite number of states which it can occupy.

An infinite number of distinct states in the humble qubit! This might seem to give us straight away an astonishing advantage over the classical case, information-theoretically. Cannot we now encode an infinite amount of information into a single one of these simplest of systems, if we want to? Well, no: this is too quick. For although a qubit possesses an infinite number of *distinct* states, it only ever has at most *two* states which can be *distinguished* from one another at a given time, as we have already seenⁱⁱⁱ. True, I could prepare a qubit in any one of its possible states, choosing one at random from the infinite set, but this would not count as encoding an infinite amount of (classical) information into the qubit, for there is no way of *decoding* that information: finding out what the state prepared was. This is an instance of an extremely important principle, which we can call the *Inaccessibility of Unknown State Identity*, or perhaps the *Hiddenness of Quantum Information*:

Given a single copy of a system prepared in an unknown quantum state, it is impossible to determine the identity of that state by measurement on, or interaction with, the system alone.

Only if the state of a system has been selected as one from a known *orthogonal set* of states is it possible to determine the identity of the state otherwise than just by asking whomever prepared it what its state is.

These facts force us to recognise an important conceptual distinction which is easy to miss in the classical case. This is the distinction between *specification information* – the amount of (classical) information required to specify the state of a given system; and *accessible information* (Schumacher 1995) – the amount of (classical) information which can successfully be encoded into a given system. (Encoding only counts as successful if *decoding* is in principle possible.) These two quantities coincide in the classical case, but differ in the quantum, as we have just seen: the specification information for the unknown state of a qubit may be infinite (two continuous parameters required), but the accessible information – the most that can be encoded – is just one bit, since by definition a qubit has only two states in any mutually orthogonal set of its states. An important result known as the *Holevo Bound* (Holevo 1973) shows quite generally that the maximum accessible information for a quantum system with d mutually orthogonal states available to it is $\log_2 d$.

There is no particular advantage, then, in straightforwardly encoding *classical* information into qubits: the main interest of quantum information theory therefore lies elsewhere, particularly when one begins to focus on quantum states themselves being the items of information-theoretic interest, and when one looks at the distinctive possibilities that arise when entanglement is used as a resource.

Quantum information proper

In the Shannon paradigm, the aim of a communication protocol is the reproduction at some location of a message – a sequence of distinguishable states – selected elsewhere. The apparatus of Shannon's theory is then used to characterise the resources required achieve this, particularly in the presence of noise. But the theory is largely silent on any particular reason why one might want to reproduce a given sequence of states: ultimately this will be down to the interests of whomever is seeking to design and build the communication system in question. (Very often, reproducing a sequence at a destination will enable the user to achieve something else that they are interested in.)

It was Schumacher's (1995) insight that the general framework which Shannon developed for thinking about information sources could be applied also in the case of quantum mechanics. Instead of thinking of an information source as producing a sequence of distinct and distinguishable (i.e., classical) states as its messages, each state appearing in the sequence with some fixed probability, we can think of a source as producing a sequence of systems being in particular *quantum* states, where a given quantum state will occur with a fixed probability in the sequence (call this a *quantum information source*). So considering, for example, a set $\{|a_i\rangle\}$ of non-orthogonal quantum states, we can contemplate a source which would produce a sequence of quantum systems, one after the other, where the probability that a given system in the sequence will be in the particular state $|a_i\rangle$ is $p(a_i)$. Then the message produced from a particular run of the source might be a sequence like $|a_2\rangle|a_5\rangle|a_1\rangle|a_4\rangle|a_4\rangle|a_4\rangle|a_7\rangle\dots$, and so on: a specific sequence of non-orthogonal states. The task of communication, following the Shannon paradigm, would then be to reproduce this sequence of states at the destination. Alternatively, it might be that one's source produces a sequence of quantum systems each of which is entangled in some way with some *other* quantum systems. Then the task of communication will be that of producing at the destination a new sequence of systems which are entangled in *exactly the same way* to this other set of systems as those in the sequence initially produced by the source were. Either way, Schumacher observed that the resources required to achieve these tasks could be quantified. He derived, using an extension of Shannon's techniques, the minimum number of qubits that would be required to encode the output of a quantum information source in such a way that the output sequence of states could be reproduced at the intended destination as required, including any entanglement which might have existed between systems in the output sequence and other systems. This was Schumacher's quantum noiseless coding theorem. In perfect parallel to the classical case, the minimum number of *qubits* required to encode the output of a *quantum* information source defines the notion of the *amount of*

quantum information that a given quantum source produces. Again in parallel to the classical case, the maximum amount of quantum information that a qubit can contain is simply one qubit's worth^{iv}.

A moment ago we noted that Shannon's theory had little to say on the general question of what *purpose* there might be in trying to reproduce elsewhere a sequences of states generated by some source. This question clearly also arises in the quantum case: why should we wish to be able to reproduce at location *B* a sequence of quantum states, or a pattern of entanglement, produced at location *A*? Well, the specific reasons could be many and various, but as before, the most general and fundamental answer is that this will ultimately just come down to the interests of whomever is setting-up the communication protocol. But given the introduction of the notion of a quantum information source, there is a further specific question which we might now press, and to which we can give some substantive answers, namely: why would anyone be interested in reproducing at a destination a sequence of quantum states *as opposed to* a sequence of distinguishable – classical – states?

One immediate very important example is given by the case of quantum computation. A useful way to think about quantum computation is in terms of the so-called *quantum circuit* architecture. Here we imagine a (possibly large, but finite) register of qubits, each of which can be acted on individually and jointly by various *quantum logic gates*, where each quantum gate implements some unitary quantum dynamics on the qubits it acts on. The register begins with each qubit in the standard $|0\rangle$ computational basis state, then, analogously to a sequence of logic gates in a classical logic circuit, some particular sequence of quantum gates is applied to the register, corresponding to a particular quantum computation being performed. This will typically leave the register in an entangled state. Readout is then performed by measurement in the computational basis - the 0 or 1 boxes for each qubit are opened - leaving one with a string of classical bit values as the result of the computation. It turns out that a relatively small set of one- and two- qubit quantum gates are sufficient, when suitably combined, to produce the result of any possible unitary dynamics on the qubits in the register, so we can think of a (suitably sized) quantum computer as a good model of any quantum-mechanical system whatsoever; and moreover, we recover the feature of computational universality familiar from the classical Turing machine model. Any possible quantum computation can be performed (given enough time, and given a large enough register) by a machine of finite specification, i.e., one which is able to apply each of the gates in the small 'universal set' of gates to its register of qubits, as needed.

Now, whilst the set of functions which a quantum computer can compute is in fact the same as the set of functions a classical Turing machine can compute, the important point is that quantum computers can perform some computations much more quickly than any classical computer. However, in order to achieve this, we will need to ensure that a computation 'remains quantum' from beginning to end. There is a very strong tendency for noise – unwanted interaction with uncontrolled or external degrees of freedom – to make quantum superposition states such as $a|0\rangle + b|1\rangle$ start behaving as if they were just an ordinary classical statistical mixture of definite classical bit values 0 and 1, with a probability $|a|^2$ of the definite value being 0 and a probability $|b|^2$ of the definite value being 1. If this occurs on too wide a scale in our quantum computer, it will just start behaving exactly as a classical computer does, and the quantum speed-up will be lost. Therefore it is important when designing practical quantum computers - which will typically involve various components located in various different spatial locations - that the quantum states involved in the computation can reliably be stored, and also reliably and accurately moved from one location to another, whilst preserving any necessary entanglement structure. In other words, it will be necessary that sequences of quantum states produced in one location can be *reproduced* at another location, accurately, and preserving any necessary entanglement to other systems. Moreover, we will be interested in the minimum resources required to achieve such transmission, in order to build the optimum computer. Thus we require the notions of quantum information source, of quantum compression (coding), and of successful message

reproduction that Schumacher introduced, in order to design and construct practical quantum computers.

Speaking more generally, one will be interested in reproducing quantum states at a destination whenever having a system, or systems, with a well-defined (even if perhaps unknown) quantum state in one's possession will confer an advantage in some task one is trying to achieve.

A further example of gaining an advantage is provided by the special kind of non-classical, holistic, correlations that obtain between systems in entangled states. Entanglement proves to be an extremely important non-classical communication resource. When two spatially separated parties – traditionally called Alice and Bob – share some entanglement, each possessing one half of a pair of qubits in an entangled state, say, they can do remarkable new things that they would not otherwise be able to do. The two paradigm cases of such *entanglement-assisted communication* are *superdense coding* (Bennett and Weisner 1982) and *quantum teleportation* (Bennett *et al.* 1993).

In superdense coding, shared entanglement is deployed to improve classical information transfer. When Alice and Bob share a state like: $1/\sqrt{2} (|0\rangle|2\rangle + |1\rangle|3\rangle)$, i.e., Alice has boxes 0 and 1 close by her (her qubit), and Bob - far away - has boxes 2 and 3 close by him (his qubit), Alice can manage to transmit to Bob two bits of classical information whilst only sending him a single qubit. At first sight this looks like a violation of the Holevo bound – Alice has somehow managed to stuff two classical bits into one qubit! She pulls the trick off by, first, applying one of four unitary gates to her qubit, corresponding to a choice of two classical bit values on her part, then, second, sending her qubit to Bob. When Bob receives it, he performs a measurement on both qubits taken together (this is called performing a joint measurement), and he can infer what gate Alice applied, thus infer the bit values she wishes to transmit. Formally, therefore, there is no violation of the Holevo bound, since two qubits are involved in the protocol. But what is remarkable is the time-ordering in the procedure: Alice is able to encode two classical bit values into these two qubits *when she only has access to one of them*. At the time Alice chooses her bit values, Bob *already has* his qubit - his half of the entangled pair! Since we are ruling out superluminal information transmission (quantum theory enforces the impossibility of signalling faster than light, which is healthy, for consistency with special relativity) it is remarkable that she can manage this.

By contrast with superdense coding, quantum teleportation uses shared entanglement to transmit *quantum information* in a remarkable way. In fact, it was the very first protocol explicitly to involve the transmission of quantum information properly speaking: the reproduction at one location of a quantum state produced at another location.

Again, we begin with Alice and Bob widely separated, but sharing the entangled state $1/\sqrt{2} (|0\rangle|2\rangle + |1\rangle|3\rangle)$. Alice is presented with a qubit in some unknown (to her) quantum state $|\psi\rangle$. As we know, $|\psi\rangle$ could be any one of an infinite number of (non-orthogonal) states. We can think of this state as the output of some quantum information source. Alice's task is now to bring it about that Bob should come to have a copy of the state $|\psi\rangle$, whatever it is. How could she do this? We stipulate that she is not allowed simply to package up the qubit which is in the state $|\psi\rangle$ and send it to Bob, and we stipulate more strongly that she is not allowed to send *any* quantum systems *at all* to Bob (to stop her swapping the state $|\psi\rangle$ onto some other qubit and sending that to Bob instead). However, she is permitted to send him a small number of classical bits if she wishes. Can she achieve the task?

The answer is *yes*, but only by making use of the shared entanglement. The procedure is as follows. Alice begins by making a particular fixed joint measurement on the qubit in the state $|\psi\rangle$ and her half of the entangled pair. (Which measurement she performs is, and must be, independent of the identity of the state $|\psi\rangle$, not least since Alice has no idea what the state is, and cannot find out.) This

measurement will have one of four outcomes, each with equal probability, and thus independent of the identity of $|\psi\rangle$. Alice records the outcome in two classical bits, and then sends these to Bob. Conditional on these bit values, Bob should then perform one from a previously chosen set of four unitary gates on his half of the entangled pair. The result is that his half of the entangled pair will now be guaranteed to be in the state $|\psi\rangle$, whilst following Alice's measurement, the states of her qubits were both completely scrambled – they have no definite value for any observable quantity. Thus the state $|\psi\rangle$ has disappeared from Alice's location, and reappeared at Bob's location, whilst nothing that bears any relation to the identity of $|\psi\rangle$ has passed between them! This is a truly remarkable phenomenon, and warrants the 'teleportation' label.

Dwelling on this point a little further: it is true that it is not *matter* - colloquially speaking, 'the stuff of Alice's initial qubit' - that has been transported in this protocol. Rather, there is a physical system at A and there is a distinct physical system at B , and the latter has been made to have the state which the former used to have, but has no longer.^v Moreover, the process is not instantaneous, it can only be completed when the classical message Alice sends reaches Bob. But even so, the manner in which the result is achieved is striking. The information characterising Alice's initial qubit has in some sense been completely disembodied during the protocol: to repeat, *nothing that bears any relation to the identity of $|\psi\rangle$ passes between Alice and Bob*. The information seems to disappear from Alice's location and to reappear at Bob's, without having been anywhere in between during the process. Furthermore, this process seems to be extraordinarily efficient: it takes an infinite number of bits to specify the state of Alice's initial qubit, yet she has managed to transport it to Bob whilst only sending him a measly two bits!

Quantum teleportation is both conceptually and information-theoretically striking, whilst it brings together in one package a number of *characteristic* features of quantum information. First, the impossibility of determining an unknown state, the *hiddenness* of quantum information – Alice can't find out what $|\psi\rangle$ is and then hope to send a classical description of it to Bob. Second, the fact that shared entanglement can be used as a resource. Third, the *impossibility of cloning quantum information*.

The significance of this last is as follows. A key feature of the teleportation protocol is that Alice's copy of $|\psi\rangle$ disappears, and a new copy appears with Bob. The important *no-cloning theorem* (Dieks 1982, Wootters and Zurek 1983) tells us that this couldn't have been otherwise. The theorem states that given only a single instance of a quantum system in an unknown state $|\psi\rangle$, it is impossible to generate any further systems in the same state, i.e., it is impossible start with a qubit in the state $|\psi\rangle$ and a register of n qubits each in some standard state, say $|0\rangle$, and to end up with more than one system in the state $|\psi\rangle$. Thus if Bob is to end up with a copy of (i.e., a system in) $|\psi\rangle$, Alice cannot be left with a copy, and this is indeed what we see in teleportation.

No-cloning represents an extremely important difference between quantum and classical information, for classical information can of course be cloned – for example, we can readily measure a classical bit, and then prepare many, many new bits with the same value as the measured one. Clearly we could not do the same thing with quantum states, for measurement cannot tell us what unknown state we have before us. It follows that unknown quantum states will be very precious: if you start with only one copy of $|\psi\rangle$, you had better look after it, for you will not be able to generate more.

In teleportation and superdense coding, shared entanglement is used as a communication resource. It is also *used up* in the course of these protocols. That is, Alice and Bob start with some shared entanglement, but finish without any. Since entanglement is useful, and is used-up, it becomes natural to seek to quantify the *amount* of entanglement one has when one has a group of systems in some

entangled state. An *ebit* is the term for a basic unit of entanglement, and this is the amount of entanglement used in the teleportation or in the superdense coding protocols above. It turns out that one ebit is the *maximum* amount of entanglement there can be in a pair of qubits. If one has a large number of less entangled systems, it proves to be possible to distil their entanglement into a smaller number of maximally entangled systems, leaving-over a number of unentangled systems. The maximally entangled systems would then be useful for teleportation, or what have you. The cornerstone of the quantitative theory of entanglement is the proposition that it is not possible for two separated parties, Alice and Bob, to increase the amount of entanglement (if any) that they share by anything that either could do locally to the systems in their possession, nor by any classical communication between them, nor by any combination of the two. In other words: no increase of entanglement by local operations and classical communication. Strikingly, it is possible to teleport entanglement: suppose Cynthia hands Alice a qubit to teleport to Bob, which Alice obliges by doing. If the qubit Cynthia initially gave Alice was in fact half of an entangled pair, the other half of which Cynthia is still holding, then Bob's qubit will now be entangled with Cynthia's one. This is called *entanglement swapping*: there was entanglement between *C* and *A* and between *A* and *B*, but this is turned into entanglement between *C* and *B*. Two quantum systems which were initially independent of one another, and have never directly interacted with one another, can nevertheless be made to enter into these special quantum-correlated states.

Internal puzzles of the theory

Quantum information theory presents a number of fairly immediate conceptual puzzles internal to the theory. For example, an important question to ask is – just what is it that powers quantum computation? Where does the speed-up come from?

An immediate attractive thought is that it just comes from the possibility of superposition. Suppose I have some quantum gate that, when I prepare my n -qubit register in the computational basis in some sequence of $|0\rangle$ s and $|1\rangle$ s, evaluates some particular function of the string of 0s and 1s. If I now prepare, as I can, each qubit in the register in an equal superposition of $|0\rangle$ and $|1\rangle$, I will have as my input to the quantum gate, a state in which all 2^n strings of 0s and 1s are equally superposed. After the quantum gate operates on this superposed input, the output will be a superposition of each value of the function for all 2^n inputs. In other words, in one computational step, all possible values of the function have been computed. And this looks like a massive parallelism speed-up.

Unfortunately, things aren't so simple. In order to read any result out, we need to do a measurement of the register in the computational basis, to get a sequence of 0s and 1s which we can actually read. Thus we cannot access all 2^n values of the function: in fact we will get one of them, at random. So this is no better than classical, in terms of what we have epistemic access to. Indeed, one might argue more strongly that a computation doesn't count as being performed at all, properly speaking, unless it is in principle possible to read its outputs out. In which case, in the example just given, it wasn't even the case that the 2^n values of the function were really evaluated in the first place.

Another reason to think that it cannot be superposition, or just superposition, which is at the heart of things, is that it is not only quantum systems which can be in superpositions: classical waves, such as waves on a string, for example, can be superposed too. But there is no computational speed-up available here.

A more plausible thought is therefore that it must be *entanglement* which is responsible, for no classical system, even one which supports superposition, can support entanglement; and indeed, in those quantum algorithms, such as Shor's (1994), which give a speed-up there seems to be both parallelism *and* entanglement at play. Yet this suggestion is not entirely straightforward either. The

Gottesman-Knill theorem (Gottesman 1998) implies that if you restrict the gates available to your quantum computer to a particular sub-set of gates, but one which includes the possibility of making as much entanglement as you like, your computer cannot be more powerful than a classical one. Thus entanglement on its own is not enough. The question therefore remains open exactly what needs to be added, or exactly what is responsible for the speed-up (Jozsa and Linden 2003, Aaronson 2013, Cuffaro 2015). Significant questions have also been asked (Steane 2003) about the role of parallelism in quantum computational speed-up given that an important alternative architecture for quantum computation – *measurement-based* or *one-way* computation (Raussendorf and Briegel 2001) – does not have anything that looks like the parallel-processing of the circuit model in it. Here there is no unitary, superposition-supporting, evolution over time, instead there are only sequences of *measurements* made on a network of pre-entangled quantum systems. A final thought might be, following Bub (2007), that the key to speed-up is not that the quantum computer evaluates *all* the values of a function at the same time, but rather that it manages to produce the result of a computation without actually evaluating *any* of the values in the intermediary steps. This perhaps suggests a re-orientation of the question: rather than ask *why* quantum computers are so quick for some tasks, perhaps we should be asking *why* classical computers are so slow (Timpson 2009).

Another set of puzzles cluster around the question of what exactly is going-on in teleportation. The key questions here are 1) How does *so much* information get from Alice to Bob in teleportation, when she only sends him two bits? And 2) Just *how exactly* does the information *get* from Alice to Bob? Perhaps the view which has been most tempting for many in regard to (2) is that the answer is: backwards in time! (Jozsa 2004, Penrose 1998) The thought is that there must be some physical, spatio-temporally continuous, route by which the information characterising $|\psi\rangle$ gets to Bob, yet it cannot be via the two classical bits that Alice sends him, since a) there is not enough room in these bits to carry all the information, and anyway, b) their values are simply independent of the identity of the teleported state. On the other hand, in order for Alice and Bob to share an entangled pair, it must have been created at some point in the past, and then split into two parts, one which was sent to Alice and one which was sent to Bob. Thus there is another spatio-temporally continuous path connecting Alice and Bob apart from the path of the two classical bits: there is the path that goes from Alice's measurement at the beginning of the protocol, backwards in time to the event of the creation of the entangled pair, and then forwards in time along the path of the half of the pair which is sent to Bob. So the conclusion reached is that it must be through *this* path that the information characterising $|\psi\rangle$ reached Bob, there is no other alternative. We have reached the startling implication that quantum information is a *new kind of information* which can travel backwards in time!

For some, understandably, this is too much. Deutsch and Hayden (2000) view the conclusion as untenable and instead develop an alternative approach to quantum theory which allows them to find, after all, an always-forward-in-time route for the quantum information travelling between Alice and Bob. They argue that the information is, after all, actually carried by the two classical bits: it was just hidden away in them in such a way that it couldn't be revealed by looking at them. Yet a further response, however, is to reject both the backwards-in-time and the Deutsch-Hayden answers to the puzzle as being predicated on a mistaken picture of how we should conceive of quantum information and questions of its location and travel in protocols such as these (Timpson 2013), of which more in a moment.

This brings us to perhaps the key question of all: Just what *is* quantum information?

Quantum information theory and the nature of information

There are two aspects to the question ‘what is quantum information?’, the first conceptual, the second, ontological. Both shed light on how we should think about information theory and its relation to physics and to the physical world.

On the conceptual side, the core questions are i) is the concept of quantum information *primitive* or *defined* and ii) if defined, what relation does it bear to the concept of Shannon information? On the ontological side, the core questions are iii) what kind of existence or being does quantum information have (if any), and iv) what is the relation between quantum information and our traditional conception of the physical world as being constituted by material particles and fields?

The prevailing view on the conceptual side seems to be that quantum information is *sui generis* and a primitive concept, for example:

‘Quantum information, though *not precisely defined*, is a fundamental concept of quantum information theory’ (Horodecki et al. 2006)

‘ $|\psi\rangle$ may be viewed as a carrier of “quantum information” which...we leave...*undefined in more fundamental terms*. Quantum information is a new concept *with no classical analogue....*’ (Jozsa 2004)

However, claims to this effect seem to be based on an erroneous, or at least, an unnecessary, conception of the content of the classical Shannon theory of information. If an appropriate view is taken of the Shannon theory then it is possible to see quantum and Shannon information as species of a single genus (Timpson 2013). If one approaches the Shannon theory as fundamentally concerning inference and the reduction of uncertainty then it will be impossible to set the quantum and the classical theories together, for uncertainty and inference give us no model to understand what quantum information could be. We can’t find out what an unknown state $|\psi\rangle$ is; and if we could, and did, then we would no longer be dealing with some quantum information that a system instantiates, but rather with some classical information *about* a system, and these are very different things. But we need not and, arguably, *should not* think of the Shannon theory this way. The core content of the theory is not about uncertainty and inference (epistemic notions) but is rather about the (mechanical) production and reproduction of messages of a certain type, and the resources required to achieve this. When put like this we can now see quantum and classical information as different species of a single genus, falling within a general paradigm that Shannon introduced; for just the same is true of the quantum theory. All that differs between the two is the nature of the message which is produced and which is required to be reproduced. In the classical case it is a sequence of distinguishable states produced by a classical source; in the quantum case it is a sequence of potentially non-orthogonal quantum states produced by a quantum source, including any entanglement correlations.

On the ontological side of the question, it is tempting to conceive of quantum information as something new which is *postulated* by quantum information theory. (This thought naturally goes along with *sui generis* primitivism on the conceptual side.) Quantum information then constitutes the subject matter of the theory - the thing which the theory is about – and the theory goes on to formulate various laws about how this subject matter behaves: how it may be produced, compressed, lost, exploited, transported from *A* to *B*, changed from one form to another, and so on. This mind-set tends to force the concept of quantum information into being that of some sort of novel substance (particular) or stuff. According to temperament one might then be inclined to think of this as some novel material - or perhaps quasi-material - substance or stuff; or one might be inclined to think of it as some sort of immaterial substance or stuff. Once conceived on the “substance or stuff” model it will then be natural to ask questions such as whether quantum information is something which exists above and beyond, in addition to, the familiar material notions of particles and fields, or whether it is

perhaps more fundamental than these familiar notions, and if so, whether these latter familiar items might be reducible to quantum information. Might the world be made of (quantum) information? Some notable physicists think so:

“An alternative view is gaining popularity: a view in which *information* is regarded as the primary entity from which physical reality is built” (Davies 2010)

“The universe is the biggest thing there is and the bit is the smallest possible chunk of information. The universe is made from bits.” (Lloyd 2006)

“What is the message of the quantum?...I suggest that...the distinction between reality and our knowledge of reality, between reality and information cannot be made” (Zeilinger 2005)

“It from bit symbolizes the idea that every item of the physical world has at bottom...an immaterial source and explanation;...all things physical are information-theoretic in origin...” (Wheeler 1990)

But it is quite wrong to think of quantum information as something which is postulated by quantum information theory (Timpson 2013). Quantum information theory is not that kind of theory: it is not in the business of postulating any *thing*, whether material, quasi-material, immaterial, substance, or stuff. Rather it is in the business of uncovering and describing *what can be done* with the various things and stuffs already postulated in the world by our familiar (quantum) physical theory. Moreover, when we recognise this, we release a great deal of tension which we felt about understanding quantum teleportation. It is only on the “substance or stuff” model that one will feel forced to think of information as something which always must have a spatiotemporal location and for which a spatiotemporal path must be traced from *A* to *B*. It was this pressure which gave rise to the backwards-in-time model and the “hidden in the classical bits” model for the information flow. But we do not need to trace a path for the information in teleportation, because there is no thing “the information” to trace a path for. The only job to do is to describe the processes which are involved in completing the teleportation protocol, and in the quantum case, when entanglement is involved, there need be no locally defined properties which have a dependence on the identity of the unknown state instantiated between *A* and *B* in order for the state to be transmitted from *A* to *B*.^{vi}

Returning to our main question: What, then, is the ontological status of quantum information? We must distinguish two cases. In the first case we say that quantum information is what is produced by a quantum information source that is required to be reproduced at the destination. What is produced? A sequence of quantum states. What is a sequence of quantum states? In the helpful philosophers’ jargon of the type/token distinction, it is a type – a particular pattern of properties and relations – of which there can be various tokens, located in various places. The tokens are the *concrete* things which instantiate the type, an *abstract* thing. A group of qubits, each being in the appropriate quantum state, will be the concrete physical objects which constitute the token of a particular type. This type which they instantiate - the abstract thing – will be the piece of quantum information – the message - that the quantum source produced.

In the second case we say that quantum information is a property of a source: ‘the quantum information of a source’, i.e., its compressibility. As with any property, quantum information so conceived will be abstract rather than concrete. It will be a *physical property* since it is a property formulated in a physical theory and defined on the basis of various other physical properties; but just as with any property, physical or no, it will not be part of the spatiotemporal contents of the world.

In either case, therefore, quantum information is ontologically abstract, rather than concrete.

Conclusion

A prominent and enticing slogan of quantum information theory is that “Information is Physical” (Landauer 1996). From what we have seen, this slogan cannot be taken to express a substantive ontological claim. If we take it to refer to pieces of quantum information – what is produced by the source that needs to be reproduced at the destination – then it is a category mistake. It is the tokens which instantiate the type which are physical, not the type itself. If we take it to refer to the information of the source – the amount of quantum information the source produces – then it is trivial. For it is trivial that a physically defined quantity is physical. The correct interpretation of this slogan, I submit, is as a misleadingly-stated *methodological* claim, not *ontological* one. It is the claim that it is a very productive and salutary business to try to discover what information-processing capacities or opportunities may be hidden away in our most fundamental physical theories. *Qua* methodological claim, this is exactly right, and important, as evidenced by the impressively rude health, richness, and interest of quantum information theory. But it is unfortunate that this key methodological commitment of the discipline is sometimes given a misleading ontological cast.

We may adopt the following useful mnemonic: rather than think of quantum information theory as a theory of some enticingly new stuff – *quantum information* – we should think of it as a *quantum* information theory. A theory of what one can do by way of computation and communication with the distinctively non-classical features of quantum theory. Thus it is all a question of bracketing:

Not: (quantum information) theory, but: quantum (information theory).

Since quantum information theory is itself therefore completely neutral on the question of the nature of the physical world - it is about what can be done with quantum physical resources - we can see that no novel information-based immaterialism or other reductionism regarding the physical world can gain the least support from the rich successes of quantum information theory.

References

- Aaronson, S 2013, *Quantum Computing to Since Democritus*, Cambridge University Press
- Benioff, P 1980, ‘The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines’, *Journal of Statistical Physics*, 22(5), 563-591.
- Bennett, C., Brassard, G., Crépeau C., Jozsa, R., Peres, A., Wootters, W. 1993, ‘Teleporting an unknown state via dual classical and EPR channels’ *Physical Review Letters*, 70, 1895 - 1899.
- Bennett, C. and Weisner, S 1982, ‘Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states’, *Physical Review Letters*, 69(20), 2881-2884.
- Bub, J 2007, ‘Quantum Computation from a Quantum Logical Perspective’, *Quantum Information and Computation*, 7(4), 281-296.
- Cuffaro, M 2015, ‘On the Significance of the Gottesman-Knill Theorem’ arxiv:quant-ph/1310.0938. Fothcoming in the *British Journal for the Philosophy of Science*.
- Davies, P 2010, ‘Universe from Bit’, in Davies and Gregersen (eds.) *Information and the Nature of Reality*, Chpt. 4, pp.65-88.
- Deutsch, D 1985, ‘Quantum theory, the Church-Turing Principle and the universal quantum

- computer', *Proceedings of the Royal Society of London A*, 400, 97-117.
- Deutsch, D. and Hayden, P 2000, 'Information flow in entangled quantum systems' *Proceedings of the Royal Society of London A*, 456, 1759-1774.
- Dieks, D 1982, 'Communication by EPR devices', *Physics Letters A*, 92(6), 271-272.
- Feynman, R 1982, 'Simulating physics with computers' *International Journal of Theoretical Physics*, 21(6/7), 467-488.
- Gottesman, D 1998, 'The Heisenberg Representation of Quantum Computers', arxiv:quant-ph/9807006.
- Holevo, A 1973, 'Information theoretical aspects of quantum measurement' *Problems of Information Transmission (USSR)*, 9, 177-183.
- Horodecki, M., Horodecki, P., Horodecki, R., and Piani, M (2006). 'Quantumness of ensemble from no-broadcasting principle', *International Journal of Quantum Information*, 4(1), 105-118.
- Jozsa, R 2004, 'Illustrating the concept of quantum information', *IBM Journal of Research and Development*, 4(1), 79-85.
- Landauer, R 1996, 'The physical nature of information', *Physics Letters A*, 217, 188-193.
- Lloyd, S 2006 *Programming the Universe*, Vintage.
- Jozsa, R. and Linden, N 2003, 'On the role of entanglement in quantum-computational speed-up', *Proceedings of the Royal Society of London A*, 459(2036), 2011-2032.
- Nielsen, M. and Chuang, I 2010, *Quantum Computation and Quantum Information*. Cambridge University Press.
- Penrose, R 1998, 'Quantum computation, entanglement and state reduction', *Philosophical Transactions of the Royal Society of London A*, 356, 1927-1939.
- Raussendorf, R. and Briegel, H 2001, 'A one-way quantum computer', *Physical Review Letters* 86, 5188.
- Schumacher, B 1995, 'Quantum coding', *Physical Review A*, 51(4), 2738.
- Shannon, C 1948, 'The mathematical theory of communication', *Bell Systems Technical Journal*, 27, 379-423, 623-656.
- Shor, P 1994, 'Algorithms for quantum computation: Discrete logarithms and factoring' *Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science* pp.124-134.
- Steane, A 2003, 'A quantum computer only needs one universe', *Studies in History and Philosophy of Modern Physics*, 34(3), 469-478.
- Timpson, C 2009, 'Philosophical aspects of quantum information theory', in D. Rickles (ed.) *The Ashgate Companion to Contemporary Philosophy of Physics*, pp. 197-261. Ashgate.
- Timpson, C 2013, *Quantum Information Theory and the Foundations of Quantum Mechanics*, Oxford University Press.
- Wheeler, J 1990, 'Information, physics, quantum: The search for links', W Zurek, ed., *Complexity, Entropy and the Physics of Information*, pp. 3-28, Addison-Wesley.

Wootters, W. and Zurek, W 1983, 'A single quantum cannot be cloned' *Nature*, 299, 802-803.

Zeilinger, A 2005, 'The message of the quantum', *Nature*, 438, 743.

ⁱThis problem is usually called the *Measurement Problem*. For an up-to-date framing of the issue, and review of approaches to it, see Wallace (2009).

ⁱⁱNothing important really hangs on the choice of two-state systems, as opposed to three or four, or d -state systems. One is really just interested in counting the number of distinguishable states that one's medium will need to have, and there are many ways in which one can do that.

ⁱⁱⁱAlthough a qubit may have many *pairs* of orthogonal states (for any a and b , given a state $a|0\rangle + b|1\rangle$ there is some other state of the qubit orthogonal to it), no set of *mutually* orthogonal states of the qubit has more than two members. The states of a quantum system in fact form a *vector space*. The states $|0\rangle$, $|1\rangle$ of a qubit form a *basis* for the vector space of its states, which is therefore a two-dimensional vector space. Any other pair of mutually orthogonal states of the qubit can just be thought of as *rotation* of the $|0\rangle$, $|1\rangle$ basis states, and thus amount to another choice of basis for the space.

^{iv} Whether a given qubit does in fact contain a qubit's worth of quantum information will depend on whether it has been used to encode the output of a quantum source, and on how much (quantum) information that source actually produces, if so.

^v From the perspective of quantum field theory however – our current fundamental theory of matter and forces – there isn't a very strong notion of 'physical system located at a particular position'. Rather, there are just various quantum fields which permeate all of space and evolve over time. Particular features in a particular region of space will then be determined by what quantum state the field(s) local to that region have. In which case teleporting the state might amount to teleporting the matter, in so far as the latter is a well-defined or interesting notion.

^{vi} However, the detailed story one will tell about the exact physical process involved will depend on one's interpretation of quantum mechanics. See Timpson (2013) chpt.4 for details.