On the Significance of the Gottesman-Knill Theorem[*]

Michael E. Cuffaro[1,2]

[1]The University of Western Ontario, Department of Philosophy
[2]Munich Center for Mathematical Philosophy, Ludwig Maximilians Universität München

August 13, 2013

1 Introduction

Toil in the field of quantum computation promises a bountiful harvest, both to the pragmatic-minded researcher seeking to develop new and efficient solutions to practical problems of immediate and transparent significance, as well as to those of us moved more by philosophical concerns: we who toil in the mud and black earth, ever desirous of those remote and yet more profound insights at the root of scientific inquiry. Some of us have seen in quantum computation the promise of a solution to the interpretational debates which have characterised the foundations and philosophy of quantum mechanics since its inception (e.g., Deutsch, 1997). Some of us have seen the prospects for a deeper understanding of the nature of computation as such (cf., Hagar & Korolev, 2007; Hagar, 2007). Others have seen quantum computation as potentially illuminating ongoing debates in epistemology and the philosophy of mind (e.g., Hameroff, 1998). These are only some of the topics to which quantum computation may be relevant.[1]

The investigation represented by this present piece should be understood to be of a kin with these more philosophical inquiries, though my goals are

[*] **This is a draft version. Please do not cite.**

[1] For a survey of some of the philosophical issues relevant to quantum computation, see Hagar (2011); Aaronson (2013).

perhaps more modest than some. Specifically, I aim to clarify the discussion surrounding the claim that the presence of a quantum system in a pure entangled state is sufficient to allow a quantum computer to realise a computational advantage (or 'quantum speedup') over a classical computer.[2] Call this claim the "sufficiency of entanglement thesis". In the literature on quantum computation, one often encounters the statement that the sufficiency of entanglement thesis is false. Motivating those who would deny the sufficiency of entanglement thesis is the Gottesman-Knill theorem (Gottesman, 1999). According to this theorem, any quantum algorithm which exclusively utilises the elements of a restricted set of quantum operations can be re-expressed using an alternative formalism which shows us how the algorithm can be efficiently simulated by classical means. Since some of the algorithms which exclusively utilise operations from this set involve the use of pure entangled states, it seems that entanglement cannot be sufficient to enable one to achieve a quantum speedup.

Two distinct ways of expressing the sufficiency of entanglement thesis should be distinguished here, however. The first is this: the mere presence of a pure entangled state is sufficient to realise quantum computational speedup.

[2] I will be focusing almost exclusively on the computational capabilities associated with pure states in this paper. A system in a mixed entangled state can be thought of as being in the presence of "noise", strong enough, in some cases, to prevent the system from being capable of realising more than a very small speedup (cf. Linden & Popescu, 2001). Our concern, however, is mainly with the issue of whether even an entangled system not in the presence of any noise whatsoever (i.e., a pure state) is sufficient to enable speedup. Although the analogy with noise suggests a particular ontological relationship between pure and mixed states, the present study should be of interest whether or not one interprets pure states as somehow ontologically prior, for there obviously are philosophically interesting conceptual relationships that obtain between pure and mixed states considered as distinct classes. The fact, for instance, and the different ways in which, every mixed state *can* be viewed as an imperfect perspective on some underlying pure state (or states) is surely of philosophical interest. It is of course possible that certain properties of pure states will not generalise to, or even contradict—as the direction of motion of a single particle may run contrary to the tendency of the fluid as a whole to which it belongs—the properties of mixed states (for a claim to this effect in the context of quantum computing, see see, for example, Biham et al. 2004). Thus a more general analysis of these topics awaits if we are to come to a full understanding of the situation before us. But it *must* await—or more properly put: it must attend upon and be responsive to the conclusions of the more special investigation, for whatever one thinks of the ontological priority of pure states, the fact that one *can* view the more general description as a generalisation over pure states is a property that the general description must satisfy and therefore conform itself to.

The Gottesman-Knill theorem shows, conclusively, that this claim is false. The second, arguably more interesting, way of expressing it is the following: quantum entanglement is a resource sufficient *to enable*, or *make possible*, quantum computational speedup; i.e., *no other physical resources are needed* to make quantum speedup possible if one begins with a quantum system in a pure entangled state. This claim, or so I will argue, is true. What the Gottesman-Knill theorem shows us is only that if we limit ourselves to the Gottesman-Knill operations, we will not have used the entanglement with which we have been provided to its full potential; for, as I will clarify, all of the Gottesman-Knill operations are such that their associated statistics (even when they involve entangled states) are reproducible in a local hidden variables theory.

In arguing for this latter claim, it will be necessary to clarify just what it is that we mean to express with the notion of a local hidden variables theory. In particular I will make the case that it can be conceptually illuminating to distinguish three different types of local hidden variables theory, all compatible with Bell's minimalist conditions on a local hidden variables theory, and yet characterised, respectively, by distinct plausibility constraints specifying just what we take to be allowed in such theories. Making such a conceptual division provides us, I will argue, with a natural way of framing the difference in the relative capabilities of classical and quantum computers, and helps to make clear the precise sense in which the sufficiency of entanglement thesis is true.

Tangentially, making this distinction also provides a fresh perspective from which to evaluate the relationship between various of the traditional "no-go" theorems of quantum theory. Two in particular that I will focus on here are the "GHZ" and "CHSH" inequalities. According to a commonly held opinion, the GHZ inequality is a more powerful refutation of local hidden variables theories than the CHSH inequality. Yet as I will argue, whether the GHZ refutation is more or less powerful than the CHSH inequality depends on the way in which we explicate what we mean by such a theory.

In Section 2 of this paper I will introduce the Gottesman-Knill theorem and motivate the assertion that the sufficiency of entanglement thesis is false. In Section 3 I will consider in more detail a version of this assertion due to Jozsa & Linden (2003). I will begin to formulate a response to this assertion by, at the end of the section, attempting to physically motivate the content of the Gottesman-Knill theorem on the basis of a consideration of the CHSH inequality. In Sections 4 and 5 I will broaden the discussion by including systems subject to the GHZ inequality. In Sections 5 and 6 I will consider the consequences of our discussion for our understanding of the sufficiency of entanglement thesis, and for the explication of the notion of a local hidden variables theory.

2 The Gottesman-Knill theorem

Call[3] an operator A a *stabiliser* of the state $|\psi\rangle$ if

$$A|\psi\rangle = |\psi\rangle. \tag{2.1}$$

For instance, consider the Bell state of two qubits:[4]

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|0\rangle \otimes |0\rangle + |1\rangle \otimes |1\rangle).$$

For this state we have

$$\begin{aligned}(X \otimes X)|\Phi^+\rangle &= \frac{1}{\sqrt{2}}(|1\rangle \otimes |1\rangle + |0\rangle \otimes |0\rangle)\\ &= \frac{1}{\sqrt{2}}(|0\rangle \otimes |0\rangle + |1\rangle \otimes |1\rangle) = |\Phi^+\rangle,\\ (Z \otimes Z)|\Phi^+\rangle &= \frac{1}{\sqrt{2}}(|0\rangle \otimes |0\rangle + (-|1\rangle \otimes -|1\rangle))\\ &= \frac{1}{\sqrt{2}}(|0\rangle \otimes |0\rangle + |1\rangle \otimes |1\rangle) = |\Phi^+\rangle.\end{aligned}$$

$X \otimes X$ and $Z \otimes Z$ are thus both stabilisers of the state $|\Phi^+\rangle$. Here, X and Z are the Pauli operators:

$$X \equiv \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Z \equiv \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

[3] In the following exegesis of the Gottesman-Knill theorem I have drawn substantially from Nielsen & Chuang (2000) and from Gottesman (1999).

[4] A qubit is the basic unit of quantum information, analogous to a classical bit. It can be physically realised by any two-level quantum mechanical system. Like a bit, it can be "on": $|0\rangle$, or "off": $|1\rangle$, but unlike a bit it can also be in a superposition of these values: $\alpha|0\rangle + \beta|1\rangle,\ \alpha, \beta \in \mathbb{C}$.

The remaining Pauli operators, I (the identity operator) and Y, are defined as:[5]

$$I \equiv \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad Y \equiv \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}.$$

The set P_n of n-fold tensor products of Pauli operators (plus those multiplied by $\alpha \in \{\pm 1, \pm i\}$) forms a group of operators, called a *Pauli Group*, which is closed under matrix multiplication. For example, for $n = 2$: $P_2 \equiv \{\alpha II, \alpha IX, \alpha IY, \alpha IZ, \alpha XI, \alpha XX, \alpha XY, ...\}$.[6,7]

Call the set V_S of states that are stabilised by every element in S, where S is some group of operators closed under matrix multiplication, the *vector space stabilised by S*. Consider a state $|\psi\rangle \in V_S$. For any $s \in S$ and any unitary operation U, we have

$$U|\psi\rangle = Us|\psi\rangle = UsU^\dagger U|\psi\rangle, \tag{2.2}$$

where the last equality follows from the definition of a unitary operator. Thus $UsU^\dagger$ stabilises $U|\psi\rangle$ and the vector space UV_S is stabilised by the group $USU^\dagger \equiv \{UsU^\dagger | s \in S\}$. Consider, for instance, the state $|0\rangle$, stabilised by the Z operator. To determine the stabiliser of this state after it has been subjected to the (unitary) Hadamard[8] transformation $H|0\rangle = |+\rangle$ we simply compute $HZH^\dagger$. Thus the stabiliser of $|+\rangle$ is X.

Now let $s_1, ..., s_m$ be elements of S. $s_1, ..., s_m$ are said to *generate* the group S if every element of S can be written as a product of elements from

[5] The I, X, Y, Z operators are also sometimes denominated as σ_I, σ_X, σ_Y, σ_Z, and sometimes as σ_0, σ_1, σ_2, σ_3.

[6] For brevity, from now on I will usually omit the tensor product symbol $\otimes$ from the expressions for joint operators; i.e., AB can be considered as a shorthand form of $A \otimes B$. Similarly for states: $|\alpha\beta\rangle$ and $|\alpha\rangle|\beta\rangle$ should be understood as shorthand forms of $|\alpha\rangle \otimes |\beta\rangle$. Additionally, all of the following should be taken to be equivalent:

$$A_1 \otimes A_2 \otimes ... \otimes A_n \equiv A^{\otimes n},$$
$$|\alpha\rangle_1 \otimes |\alpha\rangle_2 \otimes ... \otimes |\alpha\rangle_n \equiv |\alpha^n\rangle \equiv |\alpha\rangle^n \equiv |\alpha\rangle^{\otimes n}.$$

[7] The Pauli operators I, X, Y, Z are rare in that they are both unitary and Hermitian operators (it is because of the latter, of course, that they are also called the Pauli *observables*). When we generalise these operators to allow multiples $\alpha = \{\pm 1, \pm i\}$, however, this is sometimes no longer the case. For example, the operators iX and $-iX$, though unitary, are not Hermitian (they are anti-Hermitian).

[8] The H or Hadamard operator takes $|0\rangle$ to $\frac{|0\rangle + |1\rangle}{\sqrt{2}} \equiv |+\rangle$ and $|1\rangle$ to $\frac{|0\rangle - |1\rangle}{\sqrt{2}} \equiv |-\rangle$.

$s_1, ..., s_m$. For instance, the reader can verify that the subgroup, A, of P_3, defined by $A \equiv \{III, ZZI, IZZ, ZIZ\}$ can be generated by the elements $\{ZZI, IZZ\}$ (Nielsen & Chuang, 2000, §10.5.1). We may thus alternately express A in terms of its generators as follows: $A = \langle ZZI, IZZ \rangle$.

In order to compute the action of a unitary operator on a group S it suffices to compute the action of the unitary operator on the generators of S. For instance, $|0\rangle^{\otimes n}$ is the unique state stabilised by $\langle Z_1, Z_2, ..., Z_n \rangle$ (where the latter expression is a shorthand form of $\langle ZI^{\otimes n-1}, IZI^{\otimes n-2}, ..., I^{\otimes n-1}Z \rangle$). Consequently, the stabiliser of the state $H^{\otimes n}|0\rangle^{\otimes n}$ is $\langle X_1, X_2, ..., X_n \rangle$. Note that this state, expressed in the standard state vector formalism:

$$H^{\otimes n}|0\rangle^{\otimes n} = \frac{1}{2^{n/2}}(|0\rangle + |1\rangle)_1(|0\rangle + |1\rangle)_2 \dots (|0\rangle + |1\rangle)_n \tag{2.3}$$

$$= \frac{1}{2^{n/2}}(|00\dots00\rangle + |00\dots01\rangle + \dots + |11\dots10\rangle + |11\dots11\rangle, \tag{2.4}$$

specifies 2^n different amplitudes. Contrast this with the stabiliser description of the state in terms of its generators $\langle X_1, X_2, ..., X_n \rangle$, which is linear in n and thus capable of an efficient classical representation.[9]

Using the stabiliser formalism, it can be shown that all, as well as all combinations, of the following operations are capable of an efficient classical representation. (i) The *Clifford group* of gates; i.e., those unitary transformations which map elements of the Pauli group to other elements of the Pauli group.[10] These are the Pauli (I, X, Y, Z) gates, the Hadamard gate, the Phase gate (a $\pi/2$ rotation of the Bloch sphere[11] about the $\hat{z}$-axis), and the

[9] A basic distinction, in Computational Complexity Theory, is between those computational problems that are amenable to an *efficient* solution in terms of time and/or space resources, and those that are not. Easy (or 'tractable', 'feasible', 'efficiently solvable', etc.) problems are those for which solutions exist which involve resources bounded by a polynomial in the input size, n. Hard problems are those which are not easy, i.e., they are those whose solution requires resources that are 'exponential' in n, i.e., that grow faster than any polynomial in n (Nielsen & Chuang, 2000, p. 139). Note that the term 'exponential' is being used rather loosely here. Functions such as $n^{\log n}$ are called 'exponential' but do not grow as fast as a true exponential such as 2^n.

[10] A quantum 'gate' is just a unitary transformation. In a quantum computational circuit it plays a role analogous to a logic gate in a classical circuit.

[11] The Bloch sphere is a geometrical representation of the state space of a single qubit. States on the surface of the sphere represent pure states, while those in the interior represent mixed states (cf. Nielsen & Chuang, 2000).

controlled not ("CNOT") gate.[12,13] (ii) Clifford group gates conditioned on classical bits (indicating, e.g., the results of previous measurements). (iii) State preparation in the computational (i.e., $\{|0\rangle, |1\rangle\}$) basis. (iv) Measurements of observables in the Pauli group. This is the content of the *Gottesman-Knill theorem* (Nielsen & Chuang, 2000, §10.5.4).

What is especially notable about this theorem from the point of view of our discussion is that some of the states which may be realised through the operations in this set are actually entangled states. In particular, by combining a Hadamard and a CNOT gate, one can generate any one of the Bell states (which one is generated depends on the value assigned to the input qubits); i.e.,

$$\begin{aligned}
|0\rangle|0\rangle &\xrightarrow{H\otimes I} \frac{|0\rangle|0\rangle + |1\rangle|0\rangle}{\sqrt{2}} \xrightarrow{\text{CNOT}} \frac{|0\rangle|0\rangle + |1\rangle|1\rangle}{\sqrt{2}} = |\Phi^+\rangle,\\
|0\rangle|1\rangle &\xrightarrow{H\otimes I} \frac{|0\rangle|1\rangle + |1\rangle|1\rangle}{\sqrt{2}} \xrightarrow{\text{CNOT}} \frac{|0\rangle|1\rangle + |1\rangle|0\rangle}{\sqrt{2}} = |\Psi^+\rangle,\\
|1\rangle|0\rangle &\xrightarrow{H\otimes I} \frac{|0\rangle|0\rangle - |1\rangle|0\rangle}{\sqrt{2}} \xrightarrow{\text{CNOT}} \frac{|0\rangle|0\rangle - |1\rangle|1\rangle}{\sqrt{2}} = |\Phi^-\rangle,\\
|1\rangle|1\rangle &\xrightarrow{H\otimes I} \frac{|0\rangle|1\rangle - |1\rangle|1\rangle}{\sqrt{2}} \xrightarrow{\text{CNOT}} \frac{|0\rangle|1\rangle - |1\rangle|0\rangle}{\sqrt{2}} = |\Psi^-\rangle.
\end{aligned}$$

In fact many quantum algorithms utilise just such a combination of gates (e.g., teleportation; cf., Bennett et al. 1993). Recall that the sufficiency of entanglement thesis is the claim that the presence of a quantum system in a pure entangled state is sufficient to allow a quantum computer to realise a computational advantage over a classical computer. If all of operations from this set are efficiently classically simulable, however, then it appears as though the sufficiency of entanglement thesis must be false, for evidently there are quantum algorithms utilising pure entangled states that are efficiently simulable classically.

[12] The **CNOT** or controlled-not gate takes two qubits $|s\rangle|t\rangle$ to $|s\rangle|t \oplus s\rangle$, where $|s\rangle$ is the control, $|t\rangle$ the target qubit, and $\oplus$ is addition modulo 2 (i.e., 'exclusive-or'). Intuitively, the control qubit determines whether or not to apply a bit-flip operation (i.e., a NOT or X operation) to the target qubit.

[13] Note that the Hadamard, Phase, and CNOT gates by themselves suffice to generate the Clifford Group.

3 The significance of the Gottesman-Knill theorem

Reflecting on this circumstance in their influential (2003) article (in a section entitled *Is entanglement a key resource for computational power?*), Jozsa & Linden write:

> Recall that the significance of entanglement for pure-state computations is derived from the fact that unentangled pure states ... of n qubits have a description involving poly(n) parameters (in contrast to $O(2^n)$ parameters for a general pure state). But this special property of unentangled states (of having a 'small' descriptions [*sic.*]) is contingent on a particular mathematical description, as amplitudes in the computational basis. If we were to adopt some other choice of mathematical description for quantum states (and their evolution), then, although it will be mathematically equivalent to the amplitude description, there will be a different class of states which will now have a polynomially sized description; i.e. two formulations of a theory which are mathematically equivalent (and hence equally logically valid) need not have their corresponding mathematical descriptions of elements of the theory being [*sic.*] interconvertible by a *polynomially bounded* computation. With this in mind we see that the significance of entanglement as a resource for quantum computation is not an *intrinsic* property of quantum physics *itself*, but is tied to a particular additional (arbitrary) choice of mathematical formalism for the theory. ... An explicit example of an alternative formalism and its implications for the power of quantum computation is provided by the so-called stabilizer formalism and the Gottesman-Knill theorem ... Thus, in a fundamental sense, the power of quantum computation over classical computation ought to be derived simultaneously from *all* possible classical mathematical formalisms for representing quantum theory, not any single such formalism and associated quality (such as entanglement), ... (Jozsa & Linden, 2003, 2029-2030).[14]

[14] I will limit my discussion to Jozsa & Linden (2003), but Jozsa & Linden are not alone in concluding that pure state entanglement is insufficient to realise speedup. Similar

Rather than trying to make sense of the general notion of two equivalent mathematical representations of the same theoretical entity having ineliminably vastly differently sized descriptions, let us restrict ourselves to the particular case under consideration. It is easy to see, first of all, that even within the amplitude formalism one and the same system can admit of either a large or a small description. We have seen an example of this already. The state that results from applying $H^{\otimes n}$ to a system in the state $|0\rangle^{\otimes n}$ can be described as a superposition of 2^n states, as in Eq. (2.4). It can also be described as a product of n states, as in Eq. (2.3)—an exponentially smaller description. Indeed we can do much better than this, and can make do with one of the even more compact expressions:

$$\frac{1}{2^{n/2}}(|0\rangle + |1\rangle)^n \ , \qquad \frac{1}{2^{n/2}}\sum_{x}^{2^n-1}|x\rangle. \tag{3.1}$$

All of these descriptions are equivalent. In this case there is no mystery as to why. Facts about the underlying system are what make the alternative descriptions possible. For instance, the fact that the properties of each individual subsystem are maximally specifiable makes it possible to represent the superposition (2.4) as the product state (2.3). And since in this particular case each subsystem is in an identical state, we do not really need to single out any one of them, and thus we can use one of the descriptions given in (3.1).

This is not true in general. It is a quantum mechanical fact that subsystems of entangled systems are not maximally specifiable (i.e., their states are never pure). Thus entangled quantum systems cannot be given a product

considerations, presumably, lead Datta et al. to write: "the Gottesman-Knill theorem ... demonstrates that global entanglement is far from sufficient for exponential speedup." (2005, 1). Nielsen & Chuang (2000, ibid.) writing some years earlier, are, perhaps, more cautious: "The Gottesman-Knill theorem highlights how subtle is the power of quantum computation. It shows that some quantum computations involving highly entangled states may be simulated efficiently on classical computers. ... There is much more to quantum computation than just the power bestowed by quantum entanglement!" I say that this statement is more cautious because while Nielsen & Chuang correctly point out that an entangled quantum state will not, so to speak, yield a quantum speedup of its own accord, they (intentionally or not) decline to make the stronger claim, strongly suggested in my above quote of Jozsa & Linden, that further (or perhaps some other) *physical resources* besides entanglement (which are, according to Jozsa & Linden, hidden by the formalism) are required in order to make quantum speedup *possible*. I will discuss this distinction further as we proceed.

state representation. Descriptions of entangled states and of the transitions to and from them cannot therefore be compressed in the same way that (2.4) is compressible into (2.3). At least this is true in the standard amplitude formalism. Strangely, if we move from the amplitude to the stabiliser formalism it seems as though it *is* possible, somehow, to give more compact descriptions of these states and their transitions, despite the quantum mechanical fact just mentioned.

However let us persist, for a little while longer at least, in our conviction that it is facts about the underlying system and its transitions which make alternative (smaller) descriptions possible. And let us try and determine, if we can, just what these further facts may be. The Gottesman-Knill theorem does not say that *all* quantum state transitions are efficiently classically simulable when expressed in the stabiliser formalism, but only that a specific restricted subset of them are: the Clifford group of transformations (possibly conditioned on classical bits) which map the Pauli group into itself, measurements of observables in the Pauli group, and state preparation in the computational basis. Let us consider whether these transformations share anything in common apart from this abstract mathematical fact.

Consider, to start with, state preparation. This is, by hypothesis, the preparation of a product state in the computational basis. This means that each qubit will initially be in one of the states $|0\rangle$ or $|1\rangle$, stabilised by Z and $-Z$ respectively; i.e., each qubit will begin in a state equivalent to either $Z|0\rangle$ or $-Z|1\rangle$. The Pauli gates X, Y, and Z (I is just the trivial transformation) represent π rotations of the Bloch sphere about the x, y, and z axes respectively. Applied to Z they yield: $XZX^\dagger = -Z$, $YZY^\dagger = -Z$, $ZZZ^\dagger = Z$. Applied to X and Y they yield: $XXX^\dagger = X$, $YXY^\dagger = -X$, $ZXZ^\dagger = -X$, $XYX^\dagger = -Y$, $YYY^\dagger = Y$, $ZYZ^\dagger = -Y$. The Hadamard gate is a $\pi/2$-rotation about the y-axis, followed by a π-rotation about x. Applied to X, Y, and Z it yields $HXH^\dagger = Z$, $HYH^\dagger = -Y$, and $HZH^\dagger = X$. The Phase gate (R) is a $\pi/2$ rotation about the z-axis, with: $RXR^\dagger = Y$, $RYR^\dagger = X$, $RZR^\dagger = Z$. The CNOT gate is a two qubit gate but its result is either an X or I transformation applied to the target qubit.

From the foregoing it is readily seen that the combined effect of any of these operations, for any subsystem of the system, must be equivalent to the measurement of one of the Pauli observables $\pm X$, $\pm Y$, $\pm Z$ on one of the computational basis states $|0\rangle$, $|1\rangle$. The reader can easily verify that this fact continues to hold if we also include the generalisations of the Pauli operators $\pm iX$, $\pm iY$, and $\pm iZ$ among our allowed operations. Thus this fact holds true for all of the Gottesman-Knill operations.

Now since the Pauli observables (disregarding I) represent π rotations of the Bloch sphere about the x, y, and z axes, the respective orientations of different ends of an experimental apparatus set up to conduct an experiment on a combined system will never differ by anything other than an angle proportional to $\pi/2$. There is something very special about orientations of experimental devices which satisfy this property. These are precisely the orientations for which it is possible to provide a local hidden variables theory to reproduce the statistics associated with a Bell state.

To clarify, for a system in the singlet state ($|\Psi^-\rangle$), the expectation value for joint experiments on its subsystems is given by the following expression:

$$\langle \sigma_m \otimes \sigma_n \rangle = -\hat{m} \cdot \hat{n} = -\cos\theta. \tag{3.2}$$

Here σ_m, σ_n represent spin-m and spin-n experiments on the first (Alice's) and second (Bob's) subsystem, respectively, with $\hat{m}, \hat{n}$ the unit vectors representing the orientations of their two experimental devices, and θ the difference in these orientations. It is well known that it is not possible to provide a hidden variables theory to reproduce all of the predictions of quantum mechanics for this state if that theory makes the very reasonable assumption that the probabilities of local experiments on Alice's subsystem (and likewise Bob's) are completely determined by Alice's local experimental setup together with a shared hidden variable taken on by both subsystems at the time the joint state is prepared.

Consider the following expression relating different spin experiments on Alice's and Bob's respective subsystems for arbitrary directions $\hat{m}, \hat{m}', \hat{n}, \hat{n}'$:

$$|\langle \sigma_m \otimes \sigma_n \rangle + \langle \sigma_m \otimes \sigma_{n'} \rangle| + |\langle \sigma_{m'} \otimes \sigma_n \rangle - \langle \sigma_{m'} \otimes \sigma_{n'} \rangle|.$$

Let $A_\lambda(\hat{m}) \in \{\pm 1\}, B_\lambda(\hat{n}) \in \{\pm 1\}$ represent the results, given a specification of the hidden variable λ, of spin experiments on Alice's and Bob's subsystems. Given the aforementioned reasonable assumption, once λ is specified there are no further dependencies between Alice's and Bob's local experimental configurations. Thus we may substitute $\langle A_\lambda(\hat{m})B_\lambda(\hat{n})\rangle$ for $\langle \sigma_m \otimes \sigma_n \rangle$. This yields:

$$|\langle A_\lambda(\hat{m})B_\lambda(\hat{n})\rangle + \langle A_\lambda(\hat{m})B_\lambda(\hat{n}')\rangle| + |\langle A_\lambda(\hat{m}')B_\lambda(\hat{n})\rangle - \langle A_\lambda(\hat{m}')B_\lambda(\hat{n}')\rangle|,$$

which, as the reader can verify, must always be ≤ 2. This expression, a variant of the 'Bell inequality' (2004 [1964]), is known as the *Clauser-Horne-Shimony-Holt* (CHSH) inequality (cf., Clauser et al., 1969; Bell, 2004 [1981]).

Quantum mechanics violates the CHSH inequality for some experimental configurations. For example, let the system be in the singlet state and let the unit vectors $\hat{m}, \hat{m}', \hat{n}, \hat{n}'$ (taken to lie in the same plane) have the orientations $0, \pi/2, \pi/4, -\pi/4$ respectively. The differences, θ, between the different orientations (i.e., $\hat{m} - \hat{n}, \hat{m} - \hat{n}', \hat{m}' - \hat{n}$, and $\hat{m}' - \hat{n}'$) will all be in multiples of $\pi/4$ and we will have, from (3.2):

$$|\langle \sigma_m \otimes \sigma_n \rangle + \langle \sigma_m \otimes \sigma_{n'} \rangle| + |\langle \sigma_{m'} \otimes \sigma_n \rangle - \langle \sigma_{m'} \otimes \sigma_{n'} \rangle| = 2\sqrt{2} \nleq 2.$$

The predictions of quantum mechanics for arbitrary orientations $\hat{m}$, $\hat{m}'$, $\hat{n}$, $\hat{n}'$ cannot, therefore, be reproduced by a hidden variables theory in which all of the correlations between subsystems are due to a common parameter endowed to them at state preparation. *They can*, however, be reproduced by such a hidden variables theory for certain special cases. In particular, the inequality is *satisfied* (as the reader can verify) when $\hat{m}$ and $\hat{n}$, $\hat{m}$ and $\hat{n}'$, $\hat{m}'$ and $\hat{n}$, and $\hat{m}'$ and $\hat{n}'$ are all oriented at angles with respect to one another that are given in multiples of $\pi/2$.[15]

The Gottesman-Knill transformations are, therefore, precisely those for which the resultant quantum statistics for a system in a Bell state are reproducible by a local hidden variables theory; i.e., by an alternative theory in which the statistics for joint experiments are factorisable, and hence "compressible", into the products of the statistics of the (fully specified) individual subsystems. In light of this it seems unsurprising—i.e., we have given a physical motivation for the fact—that these transformations are efficiently simulable with a classical computer. *Also* unsurprising, since the quantum mechanical predictions are presumed to be consistent with a local hidden variables theory, is the fact that it is possible to give a "small" description (using the stabiliser formalism) of a system subjected to these operations *without leaving* the quantum mechanical conceptual framework.

4 The GHZ argument

But alas things are not quite so simple, for while the statistics associated with the Bell states for measurement angles which differ by an angle proportional to $\pi/2$ are reproducible by a local hidden variables theory, this is not *prima*

[15] These are the cases for which Eq. (3.2) predicts perfect correlation ($\theta = \pi$), perfect anti-correlation ($\theta = 0$), or no correlation ($\theta = \pi/2$).

facie true for every entangled state. In particular, one can use the so-called GHZ (a.k.a. "cat") state, a kind of generalisation of the Bell state for n qubits, to demonstrate, using measurements of Pauli observables *exclusively*, an incompatibility between the predictions of certain local hidden variables theories and the predictions of quantum mechanics. Moreover one can cast the demonstration in the form of a quantum circuit which uses only the Gottesman-Knill operations.

We begin[16] by considering three spatially separated spin-1/2 systems, a, b, c, which, having previously interacted, are now in the GHZ state:[17]

$$\frac{1}{\sqrt{2}}(|0\rangle_a|0\rangle_b|0\rangle_c + |1\rangle_a|1\rangle_b|1\rangle_c). \tag{4.1}$$

In this state, the eigenvalues associated with X and Y observables on individual subsystems are (as always) ± 1, while each of the tripartite observables,

$$X^a \otimes Y^b \otimes Y^c, \quad Y^a \otimes X^b \otimes Y^c, \quad Y^a \otimes Y^b \otimes X^c, \tag{4.2}$$

takes the eigenvalue -1.

Spin measurements on distinct particles are compatible with one another. However, the X and Y spin observables for any *single* system are incompatible (since they anticommute). Despite this, (4.2) is a set of compatible (i.e., commuting) operators, for the product of any two of the tripartite observables in (4.2) contains an even number of such anticommutations.

If we assume that the result of a combined measurement is factorisable, i.e., that it is just the product of the results of its constituent measurements, then we must have:

$$\begin{aligned} v(X^a \otimes Y^b \otimes Y^c) &= -1 = v(X^a) \cdot v(Y^b) \cdot v(Y^c), \\ v(Y^a \otimes X^b \otimes Y^c) &= -1 = v(Y^a) \cdot v(X^b) \cdot v(Y^c), \\ v(Y^a \otimes Y^b \otimes X^c) &= -1 = v(Y^a) \cdot v(Y^b) \cdot v(X^c), \end{aligned} \tag{4.3}$$

[16] The original version of the GHZ argument was published (1989) by Greenberger, Horne, & Zeilinger (hence the name "GHZ"). Our discussion will be based on the version of the argument subsequently given by Mermin (1990). Mermin indicates that his version is intended as a simplification of Clifton, Redhead, & Butterfield's (1991) generalised and refined version of GHZ's original proof. Note that GHZ also produced their own generalised and refined version in collaboration with Abner Shimony (1990).

[17] The label "GHZ" actually refers not just to this but to the family of states: $\frac{1}{\sqrt{2}}(|0\rangle^{\otimes n} \pm |1\rangle^{\otimes n})$.

where $v(A)$ is the result of measuring the observable A.

According to the EPR reality criterion: "If, without in any way disturbing a system, we can predict with certainty (i.e., with probability equal to unity) the value of a physical quantity, then there exists an element of physical reality corresponding to this physical quantity" (Einstein et al., 1935). With this in mind, note that, given the relationships (4.3), one can, without disturbing it, predict with certainty the result of measuring X and Y on any of the systems a, b, or c. For instance, to determinately predict $v(X^a)$ we first measure Y^b and Y^c: if $v(Y^b) \cdot v(Y^c) = +1$, then $v(X^a) = -1$, otherwise if $v(Y^b)$ and $v(Y^c)$ yield opposite results, then $v(X^a) = +1$. In a similar manner it is possible to determine all of the $v(X^\alpha)$ and $v(Y^\alpha)$. According to the EPR reality criterion we may, therefore, take these measurements to reveal six independently existing elements of physical reality corresponding to the x and y spin components of each of the three individual particles:

$$s_x^a,\ s_y^a,\ s_x^b,\ s_y^b,\ s_x^c,\ s_y^c, \tag{4.4}$$

where $s_x^a = v(X^a)$, $s_x^b = v(X^b)$, and so on. This is significant, for since spin measurements in distinct directions on a single system do not commute, orthodoxy would have it that a single system simply does not possess simultaneous values for both spin-x and spin-y. The apparent ability to ascribe elements of physical reality corresponding to both components of spin for each of the three particles flies very much in the face of the orthodox interpretation.

There is a problem with these elements of reality, however, for if we take the product of the commuting observables in (4.2), then since $XY = iZ = -YX$, $XZ = -iY = -ZX$, $YZ = iX = -ZY$, $XX = YY = ZZ = I$, this must yield:

$$\begin{aligned} & (X^a \otimes Y^b \otimes Y^c) \cdot (Y^a \otimes X^b \otimes Y^c) \cdot (Y^a \otimes Y^b \otimes X^c) \\ = \ & - X^a \otimes X^b \otimes X^c. \end{aligned} \tag{4.5}$$

This implies that

$$v(X^a \otimes X^b \otimes X^c) = 1. \tag{4.6}$$

Yet our description in terms of local hidden variables—the elements of reality (4.4)—has it, from (4.3), that

$$v(X^a \otimes X^b \otimes X^c) = v(X^a) \cdot v(X^b) \cdot v(X^c) = -1.$$

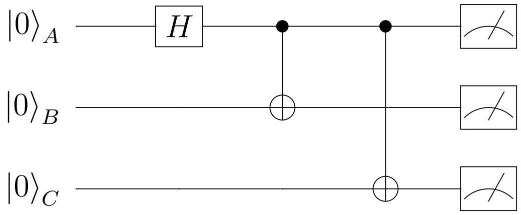

Figure 1: A quantum circuit diagram representation of the GHZ argument. The boxes at the far right indicate measurements (of X or Y observables in this case). Up to that point the diagram should be read as: $\mathrm{CNOT_{AC}CNOT_{AB}H_A}|0\rangle|0\rangle|0\rangle$. All of the operations depicted, including the final measurements, are examples of Gottesman-Knill operations.

We therefore conclude that it is impossible to give an independent specification of each of the elements of reality depicted in (4.4) in a way that is consistent with the predictions of quantum mechanics.

Now in light of the GHZ argument it seems that the physical motivation for the content of the Gottesman-Knill theorem which, in the previous section, we gleaned from the CHSH inequality cannot be maintained beyond that particular case. Indeed it is easy to represent the GHZ argument by means of a series of exclusively Gottesman-Knill operations (see Figure 1).

5 The sufficiency of entanglement thesis

And yet there is nevertheless and in spite of this still a sense in which the physical insight we gleaned from our consideration of the CHSH inequality can in fact be maintained. This will become increasingly clear as we proceed. To frame our discussion from here on, consider the following hidden variables theory (Tessier, 2004; Tessier et al., 2005) for reproducing the statistics associated with Pauli measurements on the three-qubit GHZ state

$(1/\sqrt{2})(|000\rangle + |111\rangle)$:

$$\begin{array}{c|ccc|} & q_A & q_B & q_C \\ \hline X & R_2R_3 & R_2 & R_3 \\ Y & iR_1R_2R_3 & iR_1R_2 & iR_1R_3 \\ Z & R_1 & R_1 & R_1 \\ I & 1 & 1 & 1 \end{array} . \tag{5.1}$$

q_A, q_B, and q_C are the constituent qubits of the system, offering themselves up to be measured by Alice, Bob, and Chris respectively. X, Y, Z, and I refer to measurements of Pauli observables. The R_k are random variables which return a value of ± 1 with equal probability. They are to be interpreted epistemically in the sense that they represent a determinate value of either $+1$ or -1 that is taken on by the system at state preparation. This value can *only* be revealed by measurement; i.e., nothing can be done at state preparation to *fix* the value of R_k: distinct systems subjected to identical state preparations will in general have different values for their R_k.[18] To determine the outcome of a particular measurement, we multiply the entries in the lookup table corresponding to the sub-measurements performed on each qubit, with $(R_k)^2 = 1$, discarding any lone straggling value of i that remains after calculating the final result. For example, $v(XXX) = R_2R_3R_2R_3 = 1$, $v(XYY) = R_2R_3iR_1R_2iR_1R_3 = -1$, $v(YYX) = iR_1R_2R_3iR_1R_2R_3 = -1$, $v(XYI) = R_2R_3iR_1R_2 = \pm i \Rightarrow \pm 1$, etc.

It can be verified that all of the predictions of quantum mechanics for joint Pauli experiments on the GHZ state are recovered by this hidden variables theory. Unfortunately the results of these experiments cannot be made consistent with one another under the assumption that each qubit's x and y spin component is an independently existing element of reality associated with the system. For instance, the outcome of the joint measurement XYY is $v(XYY) = R_2R_3iR_1R_2iR_1R_3 = -1$. Under the supposition that each qubit possesses independent values of both spin-x and spin-y, this joint measurement outcome must be consistent with the product of the outcomes of the individual measurements XII, IYI, and IIY. But it is not, for $v(XII) \times v(IYI) \times v(IIY) = (R_2R_3)(R_1R_2)(R_1R_3) = 1$.

This can, however, be compensated for. We can ensure consistency by adding just a wee bit of *signalling* to the model. For instance, Bob and Alice

[18] No such interpretation of the variables R_k is given in either Tessier (2004) or Tessier et al. (2005), but such an interpretation is implicit if one is to make sense of any of the claims made therein.

can agree that he will send her a single classical bit indicating whether or not he performed a Y measurement on his qubit. Upon receipt of this bit, Alice should flip the sign of her local outcome if either she or Bob (or if both of them) measured Y. Thus for the case above we will have $v(XII) \times v(IYI) \times v(IIY) = (-R_2R_3)(R_1R_2)(R_1R_3) = -1$, consistent both with the value obtained for the joint measurement XYY *and* with the individual results for the measurements XII, IYI, and IIY (each of the latter three produces a random outcome of ±1 with equal probability). In this manner it is possible to make the outcome of every joint measurement specifiable in the model consistent with the corresponding product of individual measurement outcomes.

The scheme generalises. For an n-qubit GHZ state

$$\frac{1}{\sqrt{2}}(|0\rangle^{\otimes n} \pm |1\rangle^{\otimes n}),$$

only $n-2$ bits of classical communication are required to accurately model the statistics associated with measurements of Pauli observables on the system. Indeed, unsurprisingly given the discussion of Section 3, this is true for any circuit consisting exclusively of Gottesman-Knill operations (details are given in Tessier 2004).

Tessier characterises the significance of these results in the following way:[19]

> Our results yield an alternative perspective on the GK theorem, and demonstrate that we may replace the nonlocal hidden variables represented by the stabilizer generators with LHVs and an amount of classical communication that scales efficiently with the size of the problem. This is a general feature of quantum circuits obeying the constraints of the GK theorem since, as our model illustrates, such circuits do not utilize the full capabilities of the available entanglement in the probability distributions that they generate (Tessier, 2004, 103).[20]

Conjecturing that the amount of classical communication required to simulate measurements of observables *outside* of the Pauli group will grow exponentially with the number of qubits in the system, Tessier concludes:

[19] See also Tessier et al. (2005).

[20] "GK" refers to Gottesman-Knill; "LHV" stands for local hidden variables.

> The success of our simulation provides strong evidence that the power of quantum computation arises not directly from entanglement, but rather from the nonexistence of an efficient, local realistic description of the computation, even when supplemented by an efficient amount of nonlocal, but classical communication (Tessier, 2004, 117).

While I am in agreement with the spirit of Tessier's conclusions, there are some points which require further clarification if confusion is to be avoided. In particular, the claim that "the power of quantum computation arises not directly from entanglement" requires some qualification. First, we must clarify precisely what is meant by "arises" in this context. Both the necessity and sufficiency of entanglement for enabling quantum speedup have been questioned on independent grounds, and indeed they are separate questions. For even if realising an entangled state is enough to enable a quantum speedup, it may be that there are alternative ways in which to achieve this.[21] On the other hand, from the fact that entanglement is required to enable quantum speedup, it does not follow that nothing else is.

The Gottesman-Knill theorem is most relevant to the question of *sufficiency*: the fact that a quantum algorithm utilising entanglement is efficiently simulable classically does not rule out that entanglement may nevertheless be *required* to enable quantum speedup. Yet it does seem to rule out that the realisation of an entangled state *is enough* to achieve a quantum speedup.

We have so far distinguished two senses in which the power of quantum computation may be said to "arise" from quantum entanglement, but this analysis is still too coarse for our purposes, for even the question of sufficiency may be given different significations. It will be helpful to illustrate this with an analogy. It may rightfully be said that a life vest is sufficient to keep me afloat in the event that I lean over too far and fall overboard while enjoying the Mediterranean sun on the deck of my neighbour's yacht. In saying this the intention is not to convey that the mere presence of a life vest in the water will be sufficient to keep me from drowning, for of course I must actually wear the vest if it is to perform its function. That, however, is not a fact about the life vest's capabilities, but about my choice of whether to use it or not. In seeking for an explanation for a physical process, it is usually helpful, conceptually, to distinguish facts about the capabilities of a system from

[21] For some examples purportedly demonstrating that entanglement is unnecessary for enabling speedup, see Biham et al. (2004); Datta et al. (2008).

facts pertaining to what is actually done with it in particular cases. Both of these questions are valuable but they should not be confused. Only the former question can be interpreted as an inquiry into the *resources available* in physical systems.

Thus if one asks, as do Jozsa & Linden (cf. §3 of this paper), whether entanglement is a sufficient resource to enable quantum speedup, then the question one is asking is whether any further physical resources are required to make quantum speedup possible once one has a system in an entangled state. The answer to this question is, I would argue, no. Consider the individual state spaces of two quantum mechanical systems, $\mathcal{H}_1^{d_1}$ and $\mathcal{H}_2^{d_2}$, where d_1 and d_2 are the dimensionality of the first and second system, respectively. In quantum mechanics, the overall state space of the combined system is given by the tensor product of the two systems, $\mathcal{H}_1^{d_1} \otimes \mathcal{H}_2^{d_2}$, with dimensionality $d_1 \cdot d_2$. Thus the state space of a combined system of n two-dimensional qubits is $\otimes^n \mathcal{H}^2$, with overall dimensionality 2^n. In classical mechanics, on the other hand, the total state space of two individual subsystems $\omega_1^{d_1}$, $\omega_2^{d_2}$ is given by the Cartesian product, $\omega_1^{d_1} \times \omega_2^{d_2}$, with dimensionality $d_1 + d_2$. Thus the dimensionality of the state space of a classical system of n two-dimensional subsystems is $2n$.

As both Ekert & Jozsa (1998) and Bub (2010) note, the possibility of entangled quantum systems is what is responsible for this difference in the allowable state space. To illustrate, consider how one would go about representing a general superposition of n two-dimensional values classically. It is possible to describe certain classical systems in terms of superpositions; for instance, the state of motion of a vibrating string can be characterised as a superposition of its two lowest energy modes, in the same way that the state of a qubit can be characterised as a superposition of the states $|0\rangle$ and $|1\rangle$. The joint state of a system of n strings, however, will always be a *product* state; *general* superpositions which include, in particular, values representable by entangled quantum states, cannot be physically represented using n classical systems in this way. Indeed I made essentially the same point in §3 when I there remarked that superpositions associated with entangled states are in general, 'incompressible'.

It is, of course, possible to classically represent a general superposition of n two-dimensional values in a more roundabout way; one may use, for instance, a single classical system which allows for the discrimination of 2^n resource levels within it. The cost of such a representation scales exponentially with n, however, either (if the spacing between resource levels is kept fixed) in terms of the total amount of resource required, or (if the total amount of the resource is kept fixed) in terms of the increasing precision required to discriminate the different resource levels (Ekert & Jozsa, 1998).

Quantum systems, in contrast, are not subject to this limitation; because of the possibility of entanglement, a superposition of n d-dimensional quantum systems can be used to represent a general superposition of n d-dimensional values *directly*; i.e., without incurring the cost associated with the roundabout classical method.[22] Quantum mechanical systems, therefore, allow us to efficiently exploit the full representational capacity of Hilbert space. Classical systems do not; they require exponentially more resources in order to do so. If we have an n-fold entangled quantum system, therefore, it follows straightforwardly that the possibilities for representation associated with such a system cannot, *in general*, be efficiently simulated classically.[23,24]

Evidently, it is possible to utilise only a small portion of the state space of a quantum system. This has no bearing on the nature of the actual physical resources that are *provided* by the quantum system, however. Thus *pace* Tessier, there is a sense (which we are still in the process of clarifying) in which the power of quantum computation can be said to arise directly from entanglement despite this.[25]

[22] Duwell (2004, Ch. 8) calls this 'well-adaptedness'.

[23] There is the caveat, of course, that a quantum computer will never be found, when experimented upon, to be in one of these 'extra', nonseparable, states, and thus the final 'readout' of a quantum computer will never be one of those states. Any problem, therefore, whose solution requires such a representation cannot be solved efficiently by a quantum computer. Nevertheless, such states represent a wealth of resources that are capable of being used as intermediaries in the calculation of a solution which is representable as a separable final state.

[24] The reader familiar with the literature on quantum computation will perhaps object that quantum speedup has not yet been conclusively proven. Thus although this is generally believed to be very unlikely, it may be the case that for every polynomial time quantum algorithm there is a polynomial time classical algorithm to achieve the same result (though note that because a classical computer is incapable of efficiently representing most quantum states and state transformations, such an algorithm would have to achieve its results using a method *very* different than the quantum one). Even in this unlikely event, however, the—still interesting—question remains as to which resources enable a quantum computer to solve certain computational problems in polynomial time, and the question of the source of quantum speedup can be recast accordingly. The answer to it will likely not change. I am indebted to Filippo Annovi for this observation.

[25] Note that my characterisation of entanglement as a physical resource is not motivated only by the conceptual arguments I have just given. To cite but one of many examples, one can show (Masanes, 2006) that for any non-separable state ρ, some other state σ is

Though this was worth mentioning, considered as a challenge to Tessier's conclusions it is arguably no more than a nitpicking one, for I am not fundamentally in disagreement with the spirit of Tessier's statements on this point. More problematic, however, is Tessier's claim that the Gottesman-Knill operations are recoverable using only "LHVs and an amount of classical communication that scales efficiently with the size of the problem." It is his characterisation of the theory (5.1) as *local*, in particular, which deserves our attention. Presumably Tessier calls the theory local because one can view the correlations associated with different spin measurements as arising from purely local interactions. Figure 2 illustrates this. All of the operations depicted there that are involved in the preparation of the GHZ state are localised in space and time. This includes the CNOT gates, whose implementation requires the qubits involved to be brought together in order to interact.[26] During these interactions, the variables R_k come to be shared amongst the qubits, and it is these variables which give rise to the correlations that are responsible for the measurement results that follow. These R_k are clearly local variables, and the way in which they are shared in the GHZ state is clearly locally explicable. They are also noncontextual: for a particular system (i.e., for any *one* particular system of an ensemble), each R_k takes on the same values regardless of the experiment performed. This is visualised in Figure 3.

Unfortunately the theory, as interpreted, is in fact nonlocal, for the variables R_k do not, by themselves, completely determine the results of spin measurements; the i terms are essential in predicting the outcome of a combined measurement like XYY, and yet the i terms associated with different qubits interfere nonlocally to prevent us from unambiguously assigning values to the x, y, and z components of each qubit's spin. The classical communication that is subsequently employed to compensate for this contextuality should not, therefore, be seen as a supplement to an otherwise local hidden variables theory. Rather it should be seen as a corrective to compensate for the nonlocal influences present in the model.

capable of having its teleportation fidelity (cf. Nielsen & Chuang, 2000, §9.2.2) enhanced by ρ's presence. It is also possible to quantify the amount of entanglement contained in a given state by means of so-called entanglement measures, the theory of which is surveyed in Plenio & Virmani (2007). Conceptual considerations aside, these uses legitimate, in this author's mind, the characterisation of entanglement as a physical resource.

[26] If this were not so, one could use a CNOT gate to signal faster than light (D'Hooghe, 2003).

$|000\rangle$:

	q_A	q_B	q_C
X	R_1	R_2	R_3
Y	$-iR_1$	iR_2	iR_3
Z	1	1	1
I	1	1	1

$\xrightarrow{H_A}$

	q_A	q_B	q_C
X	1	R_2	R_3
Y	iR_1	iR_2	iR_3
Z	R_1	1	1
I	1	1	1

$\xrightarrow{\text{CNOT}_{AB}}$

	q_A	q_B	q_C
X	R_2	R_2	R_3
Y	iR_1R_2	iR_1R_2	iR_3
Z	R_1	R_1	1
I	1	1	1

$\xrightarrow{\text{CNOT}_{AC}}$

	q_A	q_B	q_C
X	R_2R_3	R_2	R_3
Y	$iR_1R_2R_3$	iR_1R_2	iR_1R_3
Z	R_1	R_1	R_1
I	1	1	1

Figure 2: A series of hidden variables tables modelling the preparation of the state GHZ $= (|000\rangle + |111\rangle)/\sqrt{2}$ (Tessier, 2004). The update rules for the H and CNOT gates are: **H:** $X^f = Z^i$, $Y^f = -Y^i$, $Z^f = X^i$. **CNOT:** $X_s^f = X_s^i X_t^i$, $Y_s^f = Y_s^i X_t^i$, $Z_s^f = Z_s^i$, $X_t^f = X_t^i$, $Y_t^f = Z_s^i Y_t^i$, $Z_t^f = Z_s^i Z_t^i$, where P^i is the specification of P before the given transformation and P^f is its new specification. s and t refer to the control and target qubits, respectively, involved in a given CNOT operation.

6 The idea of a local hidden variables theory

Despite this it is, in fact, possible to cast the hidden variables theory (5.1) as local. Doing so, moreover, is conceptually illuminating for our understanding of the capabilities associated with a quantum system, for it brings to bear our discussion of the physical motivation for the Gottesman-Knill theorem from Section 3, as well as our comparison, from the last section, of the state spaces available to systems limited to product states versus the state spaces available to systems not so limited. Doing so, however, requires reinterpreting that theory. This reinterpretation, illustrated in Figure 4, is motivated by the idea that *actually observing* the result of a combined measurement outcome requires that one *actually combine* the results of the various individual sub-measurements. To put this another way: the point being made here is that if one wishes to verify that the sub-measurements associated with, for example, an XYY measurement yielded, say, -1, $+1$, and $+1$, or even only the statistics associated with a sequence of such measurements, one must somehow gather together the results registered locally by Alice, Bob, and Chris in order to examine them. There is simply no getting around this.

			“XYY”	“YXY”	“YYX”	“XXX”
R_1	R_2	R_3	$R_2R_3iR_1R_2iR_1R_3$	$iR_1R_2R_3R_2iR_1R_3$	$iR_1R_2R_3iR_1R_2R_3$	$R_2R_3R_2R_3$
+1	+1	+1	−1	−1	−1	+1
+1	+1	−1	−1	−1	−1	+1
+1	−1	+1	−1	−1	−1	+1
+1	−1	−1	−1	−1	−1	+1
−1	+1	+1	−1	−1	−1	+1
−1	+1	−1	−1	−1	−1	+1
−1	−1	+1	−1	−1	−1	+1
−1	−1	−1	−1	−1	−1	+1

			“XII”	“IYI”	“IIY”	. . .
R_1	R_2	R_3	R_2R_3	iR_1R_2	iR_1R_3	
+1	+1	+1	+1	+1	+1	
+1	+1	−1	−1	+1	−1	
+1	−1	+1	−1	−1	+1	
+1	−1	−1	+1	−1	−1	
−1	+1	+1	+1	−1	−1	
−1	+1	−1	−1	−1	+1	
−1	−1	+1	−1	+1	−1	
−1	−1	−1	+1	+1	+1	

Figure 3: A “Tractarian” (Wittgenstein, 2005 [1921], 5.15–5.156) summary of the statistics for Pauli measurements on the GHZ state, with the R_k representing the “atomic propositions.” Each row specifies the R_k of a particular system from an identically prepared ensemble. Probabilities for outcomes of experiments are determined by taking the ratio of the number of rows favourable to an outcome, to the total number of rows. Thus $Pr(XYY = -1) = Pr(YXY = -1) = Pr(YYX = -1) = Pr(XXX = +1) = 8/8 = 1$. For all X, Y, and Z measurements on a single qubit, we have $Pr(\cdot = +1) = 4/8 = 0.5$.

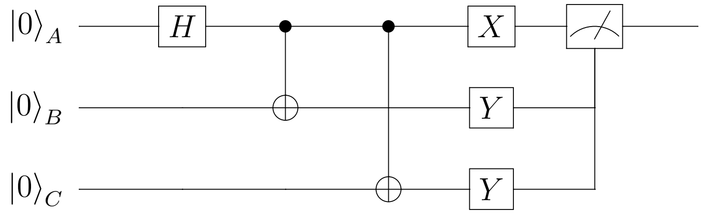

Figure 4: A reinterpretation of Tessier's (Tessier, 2004; Tessier et al., 2005) hidden variables model.

Whether the parties physically meet with one another to discuss the results over tea, or whether they physically transmit their results to a neutral party via telegraph, or use some other physical means, it is absolutely necessary that the results be collated together at some point, somehow, if the combined outcome is to be actually observed.

During this process of collating the results, there is time for Alice, Bob, and Chris (or more conspiratorially: for their systems) to signal classically to one another (at a velocity no greater than that of light) so as to coordinate the observed outcomes of their individual sub-measurements. If we now take our measurement event to consist in the act of actually observing the combined result, then all of this signalling activity will have taken place in the past light cone of that measurement event. Thus interpreted this way the theory (5.1) is indeed a local hidden variables theory.[27] And from this point of view we can see how to recover the physical motivation we appealed to in §3 for the content of the Gottesman-Knill theorem; i.e., it is available to us once again to argue that the Gottesman-Knill transformations are precisely those for which the resultant quantum measurement statistics of the system are reproducible by a local hidden variables theory; i.e., by an alternative theory in which the statistics for joint experiments are factorisable, and hence "compressible", into the products of the statistics of the fully specified individual subsystems.

The reader will likely balk at the suggestion that considering such a contrived local hidden variables theory can provide any physical motivation whatsoever, to the Gottesman-Knill theorem or to anything else. For surely,

[27] Note that, of course, no claim is being made here that (5.1) *fully* models the measurement statistics corresponding to the GHZ state; we are here speaking *only* of measurements of observables in the Pauli group.

it will be objected, the device of delaying the evaluation of the combined measurement result until the parties have found time to meet together over crumpets and tea is ad hoc and conspiratorial, almost comically so. In response, I ask the reader to consider the following: what is our goal when we provide an alternative hidden variables theory to reproduce the statistics associated with the quantum state? More succinctly: what question do we take a hidden variables theory to be answering?

One possibility is that we are answering what I will call a *theoretical* question; i.e., is there a (and if so, what is the) deeper underlying theory of the world in relation to which quantum mechanics is only an approximation? Any answer to this question will need to be very serious. It will not only need to reproduce the statistical predictions of quantum mechanics; it will also need to satisfy a number plausibility constraints. Such an answer, for instance, will need to be consistent with our other theories of physics (and if not it will need to provide a convincing reason why those should be modified), and with the body of our experiential knowledge in general. The de Broglie-Bohm family of theories is an example of a (nonlocal) answer that takes these particular constraints very seriously. Our reinterpretation of (5.1), on the other hand, is an example of an answer which does not.

A second possibility is that we are answering what I will call a *purely conceptual* question; i.e., what is logically possible and still consistent with the predictions of quantum mechanics? Answers to this question need not be very serious at all in the above sense; in fact they need satisfy no plausibility constraints whatsoever. The various toy theories are examples of such answers. Maudlin's (2011, 89-90) criticism of Howard's claim that outcome dependence implies separability, for instance, utilises a toy theory of this sort.

There is yet a third way of posing this question, which is neither theoretical nor purely conceptual. I call this the *practical* question. By this I do not intend anything to do with pragmatics, nor do I mean to imply that the question is any less profound or important. Rather I am using practical in its old signification as having to do with action and in particular human action; i.e., the practical question asks what *we* are capable of doing with the aim of reproducing the predictions of quantum mechanics. This is not merely the question of whether it is possible to provide a toy theory for the quantum state, for as we will see, answers to the practical question do, just as do answers to the theoretical question, have to satisfy plausibility constraints; and yet these are not the same constraints which must be satisfied by the answer to the theoretical question.

> Who do we think *we* are? *We* who can make 'measurements', *we* who can manipulate 'external fields', *we* who can 'signal' at all, even if not faster than light? Do *we* include chemists, or only physicists, plants, or only animals, pocket calculators, or only mainframe computers? (Bell, 2004 [1990], 245).

Indeed who *do* we think we are? For the purposes of answering the practical question I would argue that the appropriate thing to say is that it is we who can perform computations, we for whom this activity is modellable by Turing machine; indeed we whose calculational abilities Turing's analysis is specifically intended to capture (Turing, 1937, 1938). Turing computability is not the only constraint an answer to the practical question must satisfy, however. For although the constraint of Turing computability serves to delineate the space of answers that are *possible* for us (or rather for idealised versions of ourselves), what we seek is not bare possibility but *plausibility*; we seek for answers suitable to beings with finite time and space resources. Thus an answer to the practical question must not simply be modellable by Turing machine; it must also be *efficiently* modellable, in the sense that the quantum statistics are no harder (i.e., require additional resources that are at most polynomial in the input size n) for a system described by our practical local hidden variables theory to produce than they are to produce for a quantum mechanical system.

Any efficient classical computational simulation of quantum statistical predictions is *itself*, therefore, a local hidden variables theory that is neither a toy theory nor an attempt to describe what the world is actually like. Our reinterpretation of the theory (5.1) is thus, in this sense, a local hidden variables theory to recover the predictions associated with Pauli measurements on the state (4.1). And yet no-go theorems such as Bell's (and others) assure us that it is impossible to provide a local hidden variables theory to recover the predictions associated with *all* possible measurements on states like the Bell states, or states like (4.1). Thus for this reason if we limit ourselves to the Gottesman-Knill operations, we will not have used the entanglement with which we have been provided to its full potential.

Returning now to the GHZ argument, it is hard to overestimate the impact the GHZ proof has had upon the physical and philosophical community. For Mermin, for instance,

> This is an altogether more powerful refutation of the existence of elements of reality than the one provided by Bell's theorem for the two-particle EPR experiment. Bell showed that the elements of reality inferred from one group of measurements are incompatible with the *statistics* produced by a second group of measurements. Such a refutation cannot be accomplished in a single run, but is built up with increasing confidence as the number of runs increases [...] In the GHZ experiment, on the other hand, the elements of reality require a class of outcomes to occur *all* of the time, while quantum mechanics *never* allows them to occur. [...] I recently declared in writing that no set of experiments, real or *gedanken*, was known that could produce such an all-or-nothing demolition of the elements of reality. With a bow of admiration to Greenberger, Horne and Zeilinger, I hereby recant (Mermin, 1990).

Mermin's sentiment is widely shared, and yet it is misleading to claim that GHZ's argument is in this sense more 'powerful' than Bell's. For before one can say this one must first specify what one will admit within the class of possible hidden variables theories; i.e., one must be clear on the *context* in which the question is being asked. From a theoretical point of view it may be that the GHZ argument is more powerful than Bell's. From either a purely conceptual or a practical point of view, however, this is not the case at all, for the quantum mechanical predictions for measurements of Pauli-group observables on the GHZ state are fully reproducible by a local hidden variables theory like (our reinterpretation of) (5.1). Of course, if we allow measurements of observables outside of the Pauli group, then it will no longer be the case that the ensuing statistics can be modelled by a local hidden variables theory like (our reinterpretation of) (5.1). But then the contradiction with quantum mechanical predictions will be a *statistical* one—certainly it will *not* be an "all-or-nothing demolition of the elements of reality."

7 Conclusion

I have argued, in this paper, that there is an important sense in which entanglement may be said to provide sufficient physical resources to enable a quantum computer to achieve quantum computational speedup. In support of this conclusion, I have argued that claims to the contrary rest on a misunderstanding of the implications of the Gottesman-Knill theorem—that indeed, far from being a problem for the view that entanglement is a sufficient resource, the Gottesman-Knill theorem serves to *highlight* the role that is actually played by entanglement in the quantum computer and to clarify exactly *in what sense* it is sufficient.

Though I will not argue this here, I believe it unlikely that an investigation into quantum computation and quantum information theory will, by itself, resolve any of the interpretational debates at the heart of the foundations and philosophy of quantum mechanics. I also do not believe it likely that an investigation into the capabilities of quantum computers will, by itself, lead us to revise our understanding of the nature of computation as such, or to resolve any of the truly foundational questions at the centre of computational complexity theory. What I do believe likely, however, is that investigating the characteristics of quantum computers will furnish us with a fresh perspective from which to consider all of these old questions—a new opportunity to reconsider exactly what we mean in asking them. The foregoing essay is an attempt in this direction.

References

Aaronson, S. (2013). *Quantum Computing Since Democritus*. New York: Cambridge University Press.

Bell, J. S. (2004 [1964]). On the Einstein-Podolsky-Rosen paradox. In *Speakable and Unspeakable in Quantum Mechanics*, (pp. 14–21). Cambridge: Cambridge University Press.

Bell, J. S. (2004 [1981]). Bertlmann's socks and the nature of reality. In *Speakable and Unspeakable in Quantum Mechanics*, (pp. 139–158). Cambridge: Cambridge University Press.

Bell, J. S. (2004 [1990]). La nouvelle cuisine. In *Speakable and Unspeakable in Quantum Mechanics*, (pp. 232–248). Cambridge: Cambridge University Press.

Bennett, C. H., Brassard, G., Crépeau, C., Jozsa, R., Peres, A., & Wootters, W. K. (1993). Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters*, *70*, 1895–1899.

Biham, E., Brassard, G., Kenigsberg, D., & Mor, T. (2004). Quantum computing without entanglement. *Theoretical Computer Science*, *320*, 15–33.

Bub, J. (2010). Quantum entanglement and information. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Winter 2010 ed.

Clauser, J. F., Horne, M. A., Shimony, A., & Holt, R. A. (1969). Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, *23*, 880–884.

Clifton, R. K., Redhead, M. L. G., & Butterfield, J. N. (1991). Generalization of the Greenberger-Horne-Zeilinger algebraic proof of nonlocality. *Foundations of Physics*, *21*, 149–184.

Datta, A., Flammia, S. T., & Caves, C. M. (2005). Entanglement and the power of one qubit. *Physical Review A*, *72*, 042316.

Datta, A., Shaji, A., & Caves, C. M. (2008). Quantum discord and the power of one qubit. *Physical Review Letters*, *100*, 050502.

Deutsch, D. (1997). *The Fabric of Reality*. New York: Penguin.

D'Hooghe, B. (2003). Communication through measurements and unitary transformations. arXiv:quant-ph/0310185.

Duwell, A. (2004). *How to Teach an Old Dog New Tricks: Quantum Information, Quantum Computing, and the Philosophy of Physics*. Ph.D. thesis, University of Pittsburgh, Pittsburgh.

Einstein, A., Podolsky, B., & Rosen, N. (1935). Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, *47*, 777–780.

Ekert, A., & Jozsa, R. (1998). Quantum algorithms: Entanglement-enhanced information processing. *Philosophical Transactions of the Royal Society A*, *356*, 1769–1782.

Gottesman, D. (1999). The Heisenberg representation of quantum computers. In S. P. Corney, 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.

Greenberger, D. M., Horne, M. A., Shimony, A., & Zeilenger, A. (1990). Bell's theorem without inequalities. *American Journal of Physics*, *58*, 1131–1143.

Greenberger, D. M., Horne, M. A., & Zeilenger, A. (1989). Going beyond Bell's theorem. In M. Kafatos (Ed.) *Bell's Theorem, Quantum Theory and Conceptions of the Universe*, (pp. 69–72). Dordrecht: Kluwer Academic Publishers.

Hagar, A. (2007). Quantum algorithms: Philosophical lessons. *Minds & Machines*, *17*, 233–247.

Hagar, A. (2011). Quantum computing. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Spring 2011 ed.

Hagar, A., & Korolev, A. (2007). Quantum hypercomputation - hype or computation? *Philosophy of Science*, *74*, 347–363.

Hameroff, S. (1998). Quantum computation in brain microtubules? the Penrose-Hameroff 'Orch OR' model of consciousness. *Philosophical Transactions of the Royal Society A*, *356*, 1869–1896.

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.

Linden, N., & Popescu, S. (2001). Good dynamics versus bad kinematics: Is entanglement needed for quantum computation? *Physical Review Letters*, *87*, 047901.

Masanes, L. (2006). All bipartite entangled states are useful for information processing. *Physical Review Letters*, *96*, 150501.

Maudlin, T. (2011). *Quantum Non-Locality and Relativity*. Cambridge, MA: Wiley-Blackwell, third ed.

Mermin, N. D. (1990). What's wrong with these elements of reality? *Physics Today*, *43*, 9–11.

Nielsen, M. A., & Chuang, I. L. (2000). *Quantum Computation and Quantum Information*. Cambridge: Cambridge University Press.

Plenio, M. B., & Virmani, S. (2007). An introduction to entanglement measures. *Quantum Information & Computation*, *7*, 1–51.

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.

Turing, A. M. (1937). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society. Second Series*, *42*, 230–265.

Turing, A. M. (1938). On computable numbers, with an application to the Entscheidungsproblem. A correction. *Proceedings of the London Mathematical Society. Second Series*, *43*, 544–546.

Wittgenstein, L. (2005 [1921]). *Tractatus Logico-Philosophicus*. Trans. D. Pears, & B. McGuinness. London: Routledge.