ARTICLE

doi:10.1038/nature16059

# Undecidability of the spectral gap

Toby S. Cubitt[1,2], David Perez-Garcia[3,4] & Michael M. Wolf[5]

**The spectral gap—the energy difference between the ground state and first excited state of a system—is central to quantum many-body physics. Many challenging open problems, such as the Haldane conjecture, the question of the existence of gapped topological spin liquid phases, and the Yang–Mills gap conjecture, concern spectral gaps. These and other problems are particular cases of the general spectral gap problem: given the Hamiltonian of a quantum many-body system, is it gapped or gapless? Here we prove that this is an undecidable problem. Specifically, we construct families of quantum spin systems on a two-dimensional lattice with translationally invariant, nearest-neighbour interactions, for which the spectral gap problem is undecidable. This result extends to undecidability of other low-energy properties, such as the existence of algebraically decaying ground-state correlations. The proof combines Hamiltonian complexity techniques with aperiodic tilings, to construct a Hamiltonian whose ground state encodes the evolution of a quantum phase-estimation algorithm followed by a universal Turing machine. The spectral gap depends on the outcome of the corresponding 'halting problem'. Our result implies that there exists no algorithm to determine whether an arbitrary model is gapped or gapless, and that there exist models for which the presence or absence of a spectral gap is independent of the axioms of mathematics.**

The spectral gap is one of the most important physical properties of a quantum many-body system, determining much of its low-energy physics. Gapped systems exhibit non-critical behaviour (for example, massive excitations and short-range correlations), whereas phase transitions occur when the spectral gap vanishes and the system exhibits critical behaviour (for example, massless excitations and long-range correlations). Many seminal results in condensed matter theory prove that specific systems are gapped or gapless, for example, that the Heisenberg chain is gapless for half-integer spin[1] (later extended to higher dimensions[2]), or that the 1D AKLT (Affleck–Kennedy–Lieb–Tasaki) model is gapped[3]. Similarly, many famous and long-standing open problems in theoretical physics concern the presence or absence of a spectral gap. A paradigmatic example is the antiferromagnetic Heisenberg model in 1D with integer spins. The 'Haldane conjecture' that this model is gapped, first formulated in 1983[4], has yet to be rigorously proven despite strong supporting numerical evidence[5]. The same question in the case of 2D non-bipartite lattices such as the kagome lattice was posed in 1973[6]. Numerical evidence[7] strongly indicates that these systems may be topological spin liquids. This problem has attracted substantial attention[8] because materials such as herbertsmithite[9] have emerged whose interactions are well-approximated by the Heisenberg coupling. The presence of a spectral gap in these models remains one of the main unsolved questions concerning the long-sought topological spin liquid phase. In the related setting of quantum field theory, one of the most notorious open problems again concerns a spectral gap—the Yang–Mills mass gap problem[10]. Proving the existence of a gap in Yang–Mills theory could provide a full explanation of the phenomenon of quark confinement. Although there is strong supporting evidence of such a gap from numerical lattice quantum chromodynamics computations[11], the problem remains open.

All of these problems are specific instances of the general spectral gap problem: given a quantum many-body Hamiltonian, is the system it describes gapped or gapless? Our main result is to prove that the spectral gap problem is undecidable in general. This involves more than merely showing that the problem is computationally or mathematically hard. Although one may be able to solve the spectral gap problem in specific cases, our result implies that it is, in general, logically impossible to determine whether a system is gapped or gapless. This statement has two meanings, and we prove both.

(1) The spectral gap problem is algorithmically undecidable: there cannot exist any algorithm that, given a description of the local interactions, determines whether the resultant model is gapped or gapless. This is the same sense in which the halting problem is undecidable[12].

(2) The spectral gap problem is axiomatically independent: given any consistent recursive axiomatization of mathematics, there exist particular quantum many-body Hamiltonians for which the presence or absence of the spectral gap is not determined by these axioms. This is the form of undecidability encountered in Gödel's incompleteness theorem[13].

## Precise statement of results

It is important to be precise in what we mean by the spectral gap problem. To this end, we must first specify the systems we are considering. Because we are proving undecidability, the simpler the system, the stronger the result. We restrict ourselves to nearest-neighbour, translationally invariant spin lattice models on a 2D square lattice of size $L \times L$ (which we later take to $\infty$), with local Hilbert space dimension $d$. Any such Hamiltonian $H_L$ is completely specified by at most three finite-dimensional Hermitian matrices describing the local interactions of the system: two $d^2 \times d^2$ matrices $h_{\mathrm{row}}$ and $h_{\mathrm{col}}$ that specify the interactions along the rows and columns of the lattice, and a $d \times d$ matrix $h_1$ that specifies any on-site interaction. All matrix elements will be algebraic numbers, and we normalize the interaction strength such that $\max\{\|h_{\mathrm{row}}\|, \|h_{\mathrm{col}}\|, \|h_1\|\} = 1$.

We must also be precise in what we mean by 'gapped' and 'gapless' (see Fig. 1). Because quantum phase transitions occur in the thermodynamic limit of arbitrarily large system size, we are interested in the spectral gap $\Delta(H_L) = \lambda_1(H_L) - \lambda_0(H_L)$ as the system size $L \to \infty$ (where $\lambda_0$ and $\lambda_1$ are the eigenvalues of $H_L$ with the smallest and second-smallest magnitude). We take 'gapped' to mean that the system has a unique ground state and a constant lower bound on the spectral gap:

[1]Department of Computer Science, University College London, Gower Street, London WC1E 6BT, UK. [2]DAMTP, University of Cambridge, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WA, UK. [3]Departamento de Análisis Matemático and IMI, Facultad de CC Matemáticas, Universidad Complutense de Madrid, Plaza de Ciencias 3, 28040 Madrid, Spain. [4]ICMAT, C/Nicolás Cabrera, Campus de Cantoblanco, 28049 Madrid, Spain. [5]Department of Mathematics, Technische Universität München, 85748 Garching, Germany.



© 2015 Macmillan Publishers Limited. All rights reserved

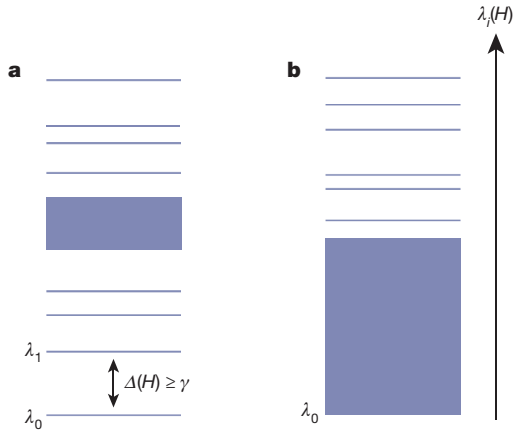


**Figure 1 | Gapped and gapless systems. a**, A gapped system has a unique ground state $\lambda_0(H)$ and a constant lower-bound $\gamma$ on the spectral gap $\Delta(H) = \lambda_1 - \lambda_0$ in the thermodynamic limit. **b**, A gapless system has continuous spectrum $\lambda_i(H)$ above the ground state in the thermodynamic limit.

$\Delta(H_L) \geq \gamma > 0$ for all sufficiently large $L$. We take 'gapless' to mean the system has continuous spectrum above the ground state in the thermodynamic limit.

Here gapped is not the negation of gapless; there are systems that fall into neither category. We adopt such strong definitions to deliberately exclude ambiguous cases, such as systems with degenerate ground states. A Hamiltonian that is gapped or gapless according to the above definitions is recognized as such throughout the literature. We show that the spectral gap problem is undecidable even given that the Hamiltonian either has a unique ground state and a spectral gap of magnitude one, or has continuous spectrum above the ground state.

We prove this by showing that the halting problem for Turing machines can be encoded in the spectral gap problem, implying that the latter is at least as hard as the former. A Turing machine is a simple, abstract model of computation in which a head reads and writes symbols from some finite alphabet on an infinite tape and moves left or right, following a finite set of rules. The halting problem asks: given an initial input written on the tape, does the Turing machine halt? Turing proved that this problem is undecidable[12]; we relate it to the spectral gap problem in the following way.

### Theorem 1

We can explicitly construct a dimension $d$, $d^2 \times d^2$ matrices $A$, $B$, $C$ and $D$, and a rational number $\beta > 0$, which can be chosen to be as small as desired, such that

(i) $A$ is Hermitian, with matrix elements in $\mathbb{Z} + \beta\mathbb{Z} + \frac{\beta}{\sqrt{2}}\mathbb{Z}$;
(ii) $B$ and $C$ have integer matrix elements; and
(iii) $D$ is Hermitian, with matrix elements in $\{0, 1, \beta\}$.

For each positive integer $n$, define the local interactions of a translationally invariant, nearest-neighbour Hamiltonian $H(n)$ on a 2D square lattice as

$$h_1(n) = \alpha(n)\Pi$$
$$h_{\mathrm{row}} = D$$
$$h_{\mathrm{col}} = A + \beta(\mathrm{e}^{i\pi\varphi(n)}B + \mathrm{e}^{-i\pi\varphi(n)}B^\dagger + \mathrm{e}^{i\pi 2^{-|\varphi(n)|}}C + \mathrm{e}^{-i\pi 2^{-|\varphi(n)|}}C^\dagger)$$

where $\varphi(n) = n/2^{|n|-1}$ is the rational number whose binary fraction expansion contains the binary digits of $n$ after the decimal point, $|\varphi(n)|$ denotes the number of digits in this expansion, $\alpha(n) \leq \beta$ is an algebraic number that is computable from $n$, $\Pi$ is a projector and the daggers denote Hermitian conjugation. Then

(i) the local interaction strength is $\leq 1$ (that is, $\|h_1(n)\|, \|h_{\mathrm{row}}\|, \|h_{\mathrm{col}}(n)\| \leq 1$);
(ii) if the universal Turing machine halts on input $n$, the Hamiltonian $H(n)$ is gapped with $\gamma \geq 1$; and
(iii) if the universal Turing machine does not halt on input $n$, the Hamiltonian $H(n)$ is gapless (that is, has continuous spectrum).

Theorem 1 implies that the spectral gap problem is algorithmically undecidable because the halting problem is. By a standard argument[14] algorithmic undecidability also implies axiomatic independence. Both forms of undecidability extend to other low-temperature properties of quantum systems, such as critical correlations in the ground state. In fact, our method allows us to prove undecidability of any physical property that distinguishes a Hamiltonian from a gapped system with unique, product ground state.

## Hamiltonian construction

We first relate undecidability of the spectral gap to undecidability of another important physical quantity, the ground state energy density, which, for a 2D lattice, is given by $E_\rho = \lim_{L\to\infty}[\lambda_0(H_L)/L^2]$. We then transform the halting problem into a question about ground state energy densities.

Reducing the ground state energy density problem to the spectral gap problem requires two ingredients.

(1) It requires a translationally invariant Hamiltonian $H_{\mathrm{u}}(\varphi)$ on a 2D square lattice with local interactions $h_{\mathrm{u}}(\varphi)$, whose ground state energy density is either strictly positive or tends to zero from below in the thermodynamic limit, depending on the value of an external parameter $\varphi$; however, determining which case holds should be undecidable. Constructing such a Hamiltonian constitutes the main technical work of our result. (These properties of $H_{\mathrm{u}}(\varphi)$ are unaffected if we multiply $h_{\mathrm{u}}(\varphi)$ by an arbitrary fixed rational number $\beta$, no matter how small.)

(2) It requires a gapless Hamiltonian $H_{\mathrm{d}}$ with translationally invariant local interactions $h_{\mathrm{d}}$ and a ground state energy of zero. (Recall that by 'gapless' we mean continuous spectrum above the ground state, not merely a vanishing spectral gap.) There are many well-known examples of such Hamiltonians, for example, that associated with the critical XY model[1].

Given Hamiltonians with these properties, we construct a new translationally invariant Hamiltonian, with local interactions $h(\varphi)$, that is gapped or gapless depending on the value of $\varphi$. The local Hilbert space of $h(\varphi)$ is the tensor product of those of $h_{\mathrm{u}}$ and $h_{\mathrm{d}}$ together with one additional energy level: $\mathcal{H} = |0\rangle \oplus \mathcal{H}_{\mathrm{u}} \otimes \mathcal{H}_{\mathrm{d}}$. We take the interaction $h^{(i,j)}$ between nearest-neighbour sites $i$ and $j$ to be

$$h(\varphi)^{(i,j)} = |0\rangle\langle 0|^{(i)} \otimes (\mathbb{1} - |0\rangle\langle 0|)^{(j)} + h_{\mathrm{u}}^{(i,j)}(\varphi) \otimes \mathbb{1}_{\mathrm{d}}^{(i,j)} + \mathbb{1}_{\mathrm{u}}^{(i,j)} \otimes h_{\mathrm{d}}^{(i,j)} \quad (1)$$

The spectrum of the new Hamiltonian $H$ is

$$\mathrm{spec}H = \{0\} \cup \{\mathrm{spec}H_{\mathrm{u}}(\varphi) + \mathrm{spec}H_{\mathrm{d}}\} \cup S \quad (2)$$

with $S \geq 1$ (see Supplementary Information for details). Recalling that we chose $H_{\mathrm{d}}$ to be gapless, we see immediately from equation (2) that if the ground state energy density of $H_{\mathrm{u}}$ tends to zero from below (so that $\lambda_0(H_{\mathrm{u}}) < 0$), then $H(\varphi)$ is gapless; if $H_{\mathrm{u}}$ has a strictly positive ground state energy density (so that $\lambda_0(H_{\mathrm{u}})$ diverges to $+\infty$), then it has a spectral gap $\geq 1$, as required (see Fig. 2).

This construction is rather general: by choosing different $h_{\mathrm{d}}$, we obtain undecidability of any physical property that distinguishes a Hamiltonian from a gapped system with a unique product ground state.

## Encoding computation in ground states

To construct the Hamiltonian $H_{\mathrm{u}}(\varphi)$, we encode the halting problem into the local interactions $h_{\mathrm{u}}(\varphi)$ of the Hamiltonian. The halting problem concerns the dynamics of a classical system—a Turing machine. To relate it to the ground state energy density—a static property of a quantum system—we construct a Hamiltonian whose ground state encodes the entire history of the computation carried out by the Turing


© 2015 Macmillan Publishers Limited. All rights reserved

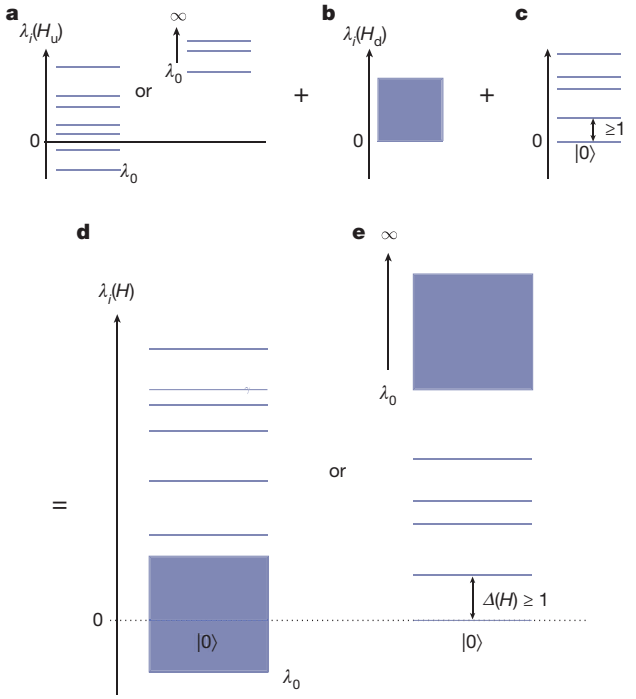


**Figure 2 | Relating ground state energy density to spectral gap. a–c**, To relate ground state energy density and spectral gap, we need a Hamiltonian $H_u(\varphi)$ whose ground state energy density is either strictly positive or tends to zero from below in the thermodynamic limit, but determining which is undecidable (**a**), and a gapless Hamiltonian $H_d$ with a ground state energy of zero (**b**). We combine $H_u(\varphi)$ and $H_d$ to form a new local interaction, $h(\varphi)$, in such a way that $H(\varphi)$ has an additional non-degenerate zero-energy eigenstate $|0\rangle$ (**c**), and that the continuous spectrum of $H_d$ is shifted immediately above the ground state energy of $H_u$. **d**, If the ground state energy density of $H_u(\varphi)$ tends to zero from below, then its ground state energy in the thermodynamic limit must be $\leq 0$, and $H(\varphi)$ is gapless. **e**, Alternatively, if the ground state energy density of $H_u(\varphi)$ is strictly positive, then its ground state energy in the thermodynamic limit must diverge to $+\infty$, and $H(\varphi)$ is gapped with gap $\Delta(H) \geq 1$.

machine in superposition[15]: if the state of the computation at time $t$ is represented by the state vector $|\psi_t\rangle$, and the computation runs until time $T$, then the ground state is the so-called 'computational history state' $\frac{1}{\sqrt{T}}\sum_{t=0}^{T-1}|t\rangle|\psi_t\rangle$. In the following, when we refer to the Turing machines encoded in the Hamiltonian 'running' on some input, we mean that the evolution produced by running the Turing machine on that input appears in the ground state as the corresponding computational history state.

If there are no other constraints, writing down a Hamiltonian whose ground state is the computational history state is straightforward. However, constructing such a Hamiltonian out of the local interactions of a many-body system is more involved. The construction method of ref. 15 was later substantially developed[16] and, after a long sequence of results[17–19], culminated in the construction for 1D spin chains with translationally invariant, nearest-neighbour interactions presented in ref. 20.

For any quantum Turing machine[21] (QTM), an interaction $h$ between neighbouring particles may be constructed[20] such that the ground state of the 1D translationally invariant Hamiltonian $H_{GI}=\sum_{i=1}^{N}h_{i,i+1}$ is of the form $\frac{1}{\sqrt{T}}\sum_{t=1}^{T}|t\rangle|\psi_t\rangle$, where the 'clock' part of the computational history state $|t\rangle \approx |1\rangle^{\otimes t}|0\rangle^{\otimes N-t}$ counts time in unary, and $|\psi_t\rangle$ represents the state of the QTM after $t$ time-steps. Moreover, the ground state energy may be taken equal to zero.

The translationally invariant Hamiltonians we are considering are completely specified by the finite number of matrix elements in the local interactions $h_{row}$, $h_{col}$ and $h_1$. To encode the halting problem in the Hamiltonian, we use quantum phase estimation[22] to encode any of the countably infinite possible inputs to the universal Turing machine (UTM) into these matrix elements.

## Quantum phase estimation

Given a unitary matrix $U$, the quantum phase estimation algorithm estimates an eigenvalue $e^{2\pi i\varphi}$ of $U$ to a given number of bits of precision (which must be chosen in advance). It is well-known[22] that if the number of bits of precision in the quantum phase estimation algorithm is greater than or equal to the number of digits in the binary fraction expansion of $\varphi$, then the quantum phase estimation algorithm, rather than estimating the phase approximately, will output all the digits of $\varphi$ (written as a binary fraction) exactly.

We use this property to construct a family of QTMs $P_n$, indexed by $n \in \mathbb{N}$, with the following properties: (i) the number of internal states and tape symbols of $P_n$ are independent of $n$; and (ii) given a number $N=2^x-1 \geq n$, with $x \in \mathbb{N}$, as input (written in binary), $P_n$ writes the binary expansion of $n$ on its tape and then halts deterministically. (The reason for having the input $N$ of this form will become clear later.) To construct $P_n$, we construct a QTM that uses the input $N$ to determine how many digits of precision to use, then runs the quantum phase estimation algorithm on the single-qubit gate $U=\begin{pmatrix}1 & 0\\ 0 & e^{2\pi i\varphi}\end{pmatrix}$. The phase $\varphi$ in $U$ is determined by the transition rules of the QTM[21]. Choosing $\varphi$ to be the rational number whose binary fraction expansion contains the digits of $n$ (expressed in binary) achieves the desired behaviour for $P_n$. By 'dovetailing' $P_n$ with a UTM (that is, running $P_n$ first, then running the UTM), the UTM runs on the input specified by $\varphi$.

The quantum computation carried out by $P_n$ followed by the UTM is encoded in the Hamiltonian using the history state construction described above. The phase $e^{2\pi i\varphi}$ being estimated then becomes one of the matrix elements of the Hamiltonian. The same happens with the $e^{i\pi 2^{-|\varphi|}}$ term that appears in the inverse quantum Fourier transform—the key ingredient of the quantum phase estimation algorithm.

Finally, we must ensure that the $|\psi_0\rangle$ component of the history state is correctly initialized to input of the form $N=2^x-1$ (written in binary) required by $P_n$. But $N=2^x-1$ in binary is simply a string of $N$ '1's, and it is easy to ensure that $|\psi_0\rangle$ is the state $|1\rangle^{\otimes N}$ using translationally invariant local interactions.

If we add an on-site interaction $h_1=|\top\rangle\langle\top|$ to the history-state Hamiltonian constructed above, which gives additional energy to the halting state $|\top\rangle$, then its ground state will pick up additional energy if and only if the UTM halts. However, the ground state energy still converges to zero as $L\to\infty$ in both cases (see Supplementary Information). The energy density therefore tends to zero in the thermodynamic limit, whether or not the UTM halts.

To remedy this, and amplify the difference between the halting and non-halting cases, we use the second spatial dimension and exploit Wang tilings.

## Quasi-periodic tilings

A Wang tile[23] is a square with markings along each edge. A tiling is then an arrangement of such tiles covering the whole plane, so that the markings on adjacent edges match. A tiling can easily be encoded in a ground state of a classical Hamiltonian on a 2D square lattice: by representing tile types by an orthogonal basis $\{|T\rangle\}$ for the local Hilbert space $\mathcal{H}_c$, and choosing local interaction terms $|T_i\rangle\langle T_i|\otimes|T_j\rangle\langle T_j|$ to give an energy penalty to all adjacent non-matching pairs of tiles $T_i$, $T_j$, a tiling of the plane is equivalent to a ground state with zero energy.

We prove, and subsequently exploit, very particular properties of the aperiodic Robinson tiling[24], and combine them with the history-state Hamiltonian. Although the pattern of tiles in the Robinson tiling extends infinitely in all directions, it never repeats. More precisely, it contains periodically repeating subpatterns that form squares with sizes given by $4^k$ for all $k \in \mathbb{N}$ (see Fig. 3). This periodicity allows us to encode



© 2015 Macmillan Publishers Limited. All rights reserved

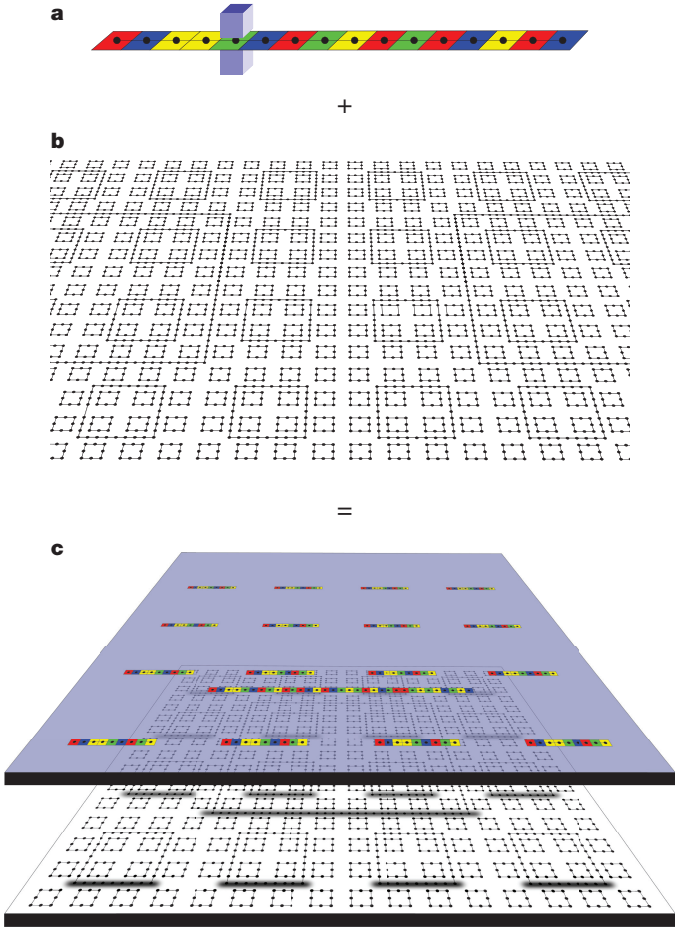


**Figure 3 | Complete Hamiltonian construction. a–c**, The Robinson tiles enforce a recursive pattern of interlocking squares, the sizes of which are given by $4^k$ for all $k \in \mathbb{N}$ (**b**). As with any Wang tiling, we can readily represent this tiling as a classical Hamiltonian whose ground state has the same quasi-periodic structure. Because the set of tiles is fixed, the local dimension of this Hamiltonian is constant. By adding a 'quantum layer' on top of the Robinson-tiling Hamiltonian and choosing a suitable translationally invariant coupling between the layers, we effectively place copies of the QTM encoded in a 1D history-state Hamiltonian (**a**) along one edge of all of the squares. The ground state of this Hamiltonian consists of the Robinson tiling configuration in the tiling layer, with computational history states in the quantum layer along one edge of each square in the tiling (**c**). Each of these encodes the evolution of the same quantum phase estimation algorithm and UTM. The effective tape length available for each QTM is determined by the size of the square it 'runs' on.

in the ground state many copies of the UTM running on the same input $\varphi$, with tapes of all possible finite lengths and for every possible finite run time (see Fig. 3).

This encoding is achieved by sandwiching the 1D quantum history-state Hamiltonian $h_q$ 'on top of' the Robinson-tiling Hamiltonian $h_c$ to form two 'layers', so that the local Hilbert space at each site is $\mathcal{H} = \mathcal{H}_c \otimes (\mathcal{H}_e \oplus \mathcal{H}_q)$ (where $\mathcal{H}_e = |0\rangle$ is an additional energy level). One can then construct a Hamiltonian (see Supplementary Information) whose ground state is of the form $|T\rangle_c \otimes |\psi\rangle_{eq}$, where $|T\rangle_c$ is a product state representing a classical configuration of the tiling layer and $|\psi\rangle_{eq}$ contains—in a tensor product structure—computational history states along one edge (a 'segment') of all of the squares appearing in the configuration given by $T$. These computational history states are essentially the only constituents of $|\psi\rangle_{eq}$ that contribute to the energy. The Hamiltonian also has an on-site interaction $h_1 = |\top\rangle\langle\top|$ that gives an additional energy to the halting state of the Turing machine $|\top\rangle$. Hence the ground state will pick up additional energy from all encoded Turing machines that halt. This energy still decreases as the relevant system size increases, which, however, is now the size of the corresponding segment in the Robinson tiling (see Fig. 3), not the overall system size.

We now consider the ground state energy. If the UTM does not halt on input $n$, then $|T\rangle_c$ is a valid tiling and for all segments that are larger than $|n|$, the ground state energy contribution is zero. The contribution for each segment smaller than $|n|$ is some algebraic computable number. If $\alpha(n)$ is the sum of the contributions of all segments smaller than $|n|$, then the addition of the constant energy shift $h_1 = -\alpha(n)\mathbb{1}$ to the Hamiltonian makes the ground state energy density negative (but tending to zero from below as $L \to \infty$) in the non-halting case (see Supplementary Information).

In the halting case, one of two things may happen. If $|T\rangle_c$ is a valid tiling, then the number of squares large enough for the encoded Turing machine to halt grows quadratically with system size, and each square contributes a small but non-zero energy. Because such a state also picks up the energy contribution from segments of size smaller than $|n|$, the energy diverges with lattice size even after adding $h_1 = -\alpha(n)\mathbb{1}$. Hence the ground state energy density is strictly positive in the halting case, as desired.

Alternatively, one could try to reduce the energy by introducing defects in the tiling, which effectively 'break' some of the Turing machines so that they do not halt. However, we prove that the Robinson tiling is robust to such defects: a tile mismatch only affects the pattern of squares in a finite region around the defect, and each defect contributes $\mathcal{O}(1)$ energy. We can choose the parameters (see Supplementary Information) to guarantee that introducing defects is energetically unfavourable. This completes the argument establishing our main result, Theorem 1.

Additional technical details can be found in the Supplementary Information.

## Discussion

We now discuss both the implications and the limitations of these results. This result is relevant to mathematical models of quantum many-body systems, as well as the behaviour of, and methods for treating, the thermodynamic limit. Moreover, it can also be seen as an indication of new physical phenomena.

An immediate consequence of the undecidability of the spectral gap is that there cannot exist an algorithm or a computable criterion that solves the spectral gap problem in general. Although algorithmic undecidability always concerns infinite families of systems, the axiomatic interpretation of the result also allows us to apply it to individual systems: there are particular Hamiltonians within these families for which one can neither prove nor disprove the presence of a gap, or of any other undecidable property. Unfortunately, our methods cannot pinpoint these particular cases, let alone prove that one of the aforementioned long-standing open problems is axiomatically undecidable.

A further consequence concerns the behaviour of the thermodynamic limit. In practice, we usually probe the idealized infinite thermodynamic limit by studying how the system behaves as we consider finite systems of increasing size. One often assumes that the systems, although finite, are so large that the asymptotic behaviour is already observed. In numerical simulations of condensed matter systems, one typically simulates finite systems of increasing size and extrapolates the asymptotic behaviour from the finite-size scaling[25]. Similarly, lattice quantum chromodynamics calculations simulate finite lattice spacings, and extrapolate the results to the continuum[11]. Renormalization group techniques accomplish something similar mathematically[26]; however, the undecidable quantum many-body models constructed in this work exhibit behaviour that defeats such approaches, in the following way. As the system size increases, the Hamiltonian will initially look like a gapless system, with the low-energy spectrum appearing to converge to a continuum. But at some threshold lattice size, a spectral gap of magnitude one will suddenly appear (or, vice versa, a gap will suddenly close[27]). Not only can the lattice size at which the system switches from


© 2015 Macmillan Publishers Limited. All rights reserved

gapless to gapped be arbitrarily large, the threshold at which this transition occurs is uncomputable. The analogous implication also holds for all other undecidable low-temperature properties. Thus, any method of extrapolating the asymptotic behaviour from finite system sizes must fail in general.

This conclusion leads us directly to new physical phenomena. First, it hints at a new type of 'phase transition', which is not driven by temperature or extrinsic local parameters, but by the size of the system. Some of the models constructed in the proof of Theorem 1 exhibit a drastic and abrupt change of properties when their size is increased beyond a certain scale. The scale at which this happens can be very large, and is not generally computable from the local description of the system. Second, our results show that certain quantum many-body models exhibit a radical form of instability. An arbitrarily small change in the parameters can cause the system to cross an arbitrary number of gapped/gapless transitions. In a sense, this phenomenon is the source of the undecidability in our models.

We finish with a closer look at some of the limitations of our results. First, all our results concern 2D (or higher-dimensional) systems. Although the majority of our construction is already 1D, we do not currently know whether the entire result holds in 1D as well. Second, although a theoretical model of a quantum many-body system is always an idealisation of the real physics, the models we construct in the proof of Theorem 1 are highly artificial. Whether the results can be extended to more natural models is yet to be determined. A related point is that we prove undecidabiltiy of the spectral gap (and other low-temperature properties) for Hamiltonians with a very particular form. We do not know how stable the results are to small deviations from this. This is a general issue with most many-body models; stability in this sense is not understood even for much simpler models such as the Ising model. Recent stability proofs only apply to certain types of frustration-free Hamiltonians[28,29]. Our results restrict the extent to which such stability results can be generalized. Similarly, we do not know whether the results hold for systems with low-dimensional local Hilbert spaces. Although the dimension $d$ in Theorem 1 is fixed and finite, providing an estimate for it would be cumbersome and certainly involve large exponentials. However, the steps in the proof described above are not tailored to minimizing this dimension. Whether there is a non-trivial bound on the dimension of the local Hilbert space below which the spectral gap problem becomes decidable is an intriguing open question.

**Received 25 March; accepted 21 September 2015.**


1. Schultz, T. D., Mattis, D. C. & Lieb, E. H. Two-dimensional Ising model as a soluble problem of many fermions. *Rev. Mod. Phys.* **36,** 856–871 (1964).
2. Hastings, M. B. Lieb-Schultz-Mattis in higher dimensions. *Phys. Rev. B* **69,** 104431 (2004).
3. Affleck, I., Kennedy, T., Lieb, E. H. & Tasaki, H. Valence bond ground states in isotropic quantum antiferromagnets. *Commun. Math. Phys.* **115,** 477–528 (1988).
4. Haldane, F. D. M. Nonlinear field theory of large-spin Heisenberg antiferromagnets: semiclassically quantized solitons of the one-dimensional easy-axis Neel state. *Phys. Rev. Lett.* **50,** 1153–1156 (1983).
5. Golinelli, O., Jolicoeur, Th. & Lacaze, R. Finite-lattice extrapolations for a Haldane-gap antiferromagnet. *Phys. Rev. B* **50,** 3037–3044 (1994).
6. Anderson, P. W. Resonating valence bonds: a new kind of insulator? *Mater. Res. Bull.* **8,** 153–160 (1973).
7. Yan, S., Huse, D. A. & White, S. R. Spin-liquid ground state of the $S = 1/2$ kagome Heisenberg antiferromagnet. *Science* **332,** 1173–1176 (2011).
8. Balents, L. Spin liquids in frustrated magnets. *Nature* **464,** 199–208 (2010).
9. Han, T.-H. *et al.* Fractionalized excitations in the spin-liquid state of a kagome-lattice antiferromagnet. *Nature* **492,** 406–410 (2012).
10. Jaffe, A. & Witten, E. in *The Millennium Prize Problems* (eds Carlson, J. A., Jaffe, A. & Wiles, A.) 129–152, http://www.claymath.org/sites/default/files/yangmills.pdf (Clay Mathematics Institute, American Mathematical Society, 2006).
11. Rebbi, C. *Lattice Gauge Theories and Monte Carlo Simulations* (World scientific, 1983).
12. Turing, A. M. On computable numbers, with an application to the Entscheidungsproblem. *Proc. Lond. Math. Soc.* **42,** 230–265 (1936).
13. Gödel, K. Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatsh. Math. Phys.* **38,** 173–198 (1931).
14. Poonen, B. in *Interpreting Gödel: Critical Essays* (ed. Kennedy, J.) 211–241 (Cambridge Univ. Press, 2014).
15. Feynman, R. P. Quantum mechanical computers. *Optics News* **11,** 11–20 (1985).
16. Kitaev, A.Y., Shen, A.H. & Vyalyi M.N. *Classical and Quantum Computation* Vol. 47 *of Graduate Studies in Mathematics* Ch. 14.4 (American Mathematical Society, 2002).
17. Kempe, J., Kitaev, A. & Regev, O. The complexity of the local Hamiltonian problem. *SIAM J. Comput.* **35,** 1070–1097 (2006).
18. Oliveira, R. & Terhal, B. M. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum Inf. Comput.* **8,** 900–924 (2008).
19. Aharonov, D., Gottesman, D., Irani, S. & Kempe, J. The power of quantum systems on a line. *Commun. Math. Phys.* **287,** 41–65 (2009).
20. Gottesman, D. & Irani, S. The quantum and classical complexity of translationally invariant tiling and Hamiltonian problems. In *Proc. 50th Annu. Symp. Found. Comput. Sci.* 95–104 (IEEE, 2009).
21. Bernstein, E. & Vazirani, U. Quantum complexity theory. *SIAM J. Comput.* **26,** 1411–1473 (1997).
22. Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* Ch. 5.2 (Cambridge Univ. Press, 2000).
23. Wang, H. Proving theorems by pattern recognition. *Bell Syst. Tech. J.* **40,** 1–41 (1961).
24. Robinson, R. M. Undecidability and nonperiodicity for tilings of the plane. *Invent. Math.* **12,** 177–209 (1971).
25. Cardy, J. (ed.) *Finite-size Scaling* (Elsevier, 2012).
26. Cardy, J. *Scaling and Renormalization in Statistical Physics* (Cambridge Univ. Press, 1996).
27. Cubitt, T. S., Perez-Garcia, D. & Wolf, M. M. Undecidability of the spectral gap. Preprint at http://arXiv.org/abs/1502.04573 (2015).
28. Bravyi, S., Hastings, M. & Michalakis, S. Topological quantum order: stability under local perturbations. *J. Math. Phys.* **51,** 093512 (2010).
29. Michalakis, S. & Zwolak, J. P. Stability of frustration-free Hamiltonians. *Commun. Math. Phys.* **322,** 277–302 (2013).


**Supplementary Information** is available in the online version of the paper.

**Acknowledgements** T.S.C. thanks IBM. T. J. Watson Laboratory for their hospitality, and C. Bennett in particular for discussions about this work. T.S.C., D.P.-G. and M.M.W. thank the Isaac Newton Institute for Mathematical Sciences, Cambridge for their hospitality during the programme "Mathematical Challenges in Quantum Information", where part of this work was carried out. T.S.C. is supported by the Royal Society. D.P.G. acknowledges support from MINECO (grant MTM2011-26912 and PRI-PIMCHI-2011-1071), Comunidad de Madrid (grant QUITEMAD+-CM, ref. S2013/ICE-2801) and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 648913). This work was made possible through the support of grant no. 48322 from the John Templeton Foundation. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

**Author Contributions** All authors contributed extensively to the paper.

**Author Information** Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to T.S.C. (t.cubitt@ucl.ac.uk).



© 2015 Macmillan Publishers Limited. All rights reserved