Determining dynamical equations is hard

Toby S. Cubitt,[1] Jens Eisert,[2] and Michael M. Wolf[3]
[1]*Department of Mathematics, University of Bristol University Walk, Bristol BS8 1TW, UK*
[2]*Institute for Physics and Astronomy, Potsdam University 14476 Potsdam, Germany*
[3]*Niels Bohr Institute, Blegdamsvej 17, 2100 Copenhagen, Denmark*

The behaviour of any physical system is governed by its underlying dynamical equations—the differential equations describing how the system evolves with time—and much of physics is ultimately concerned with discovering these dynamical equations and understanding their consequences. At the end of the day, any such dynamical law is identified by making measurements at different times, and computing the dynamical equation consistent with the acquired data. In this work, we show that, remarkably, this process is a provably computationally intractable problem (technically, it is NP-hard). That is, even for a moderately complex system, no matter how accurately we have specified the data, discovering its dynamical equations can take an infeasibly long time (unless P=NP [1]). As such, we find a complexity-theoretic solution to both the quantum and the classical embedding problems; the classical version is a long-standing open problem, dating from 1937 [2], which we finally lay to rest.

A significant part of physics is concerned with identifying the dynamical equations of physical systems and understanding their implications. But how do we deduce the dynamical equations from experimental observations? Whether deducing the laws of celestial mechanics from observations of the planets, determining economic laws from observing monetary parameters, or deducing quantum mechanical equations from observations of atoms, this task is clearly a fundamental part of physics and science in general. As we will see both for classical and for quantum systems, it is closely related to two long-standing open problems in mathematics (in the classical case, dating back over 70 years [2]).

Here, we give complexity-theoretic solutions to both these open problems which lead to a surprising conclusion: Regardless of how much information one has gained, deducing dynamical equations is in general an intractable problem—it is *NP-hard*. More precisely, the task of determining dynamical equations in general is equivalent to solving the (in)famous P versus NP problem [1]. If P≠NP, as is widely believed, then there *cannot* exist an efficient method of deducing dynamical equations.

On the positive side, our work leads to the first known algorithms for extracting dynamical equations from measurement data that are guaranteed to give the correct answer. For systems with few degrees of freedom, this is immediately applicable to many current experiments. And, indeed, the primary goal of many experiments is to characterize and understand the dynamics of a system [3–10].

In the classical setting, the problem of extracting dynamical models from experimental data has spawned an entire field known as *system identification* [11]. In the quantum case, interest in understanding the dynamics, especially externally-induced noise and decoherence, has been spurred on by efforts to develop quantum information processing technology [12, 13].

Let us make the task more concrete. Recall that in classical mechanics, the state of a system is described by a probability distribution $\boldsymbol{p}$ over its state space. Its evolution is then described by a *master equation*, whose form is determined by the system's Liouvillian, corresponding to a matrix L, as

$$\dot{\boldsymbol{p}} = L\boldsymbol{p}. \tag{1}$$

The Liouvillian expresses interactions, conservation laws, external noise etc., in short, it describes the underlying physics. We will throughout consider *open system dynamics* which takes external influences and noise into account. In order that the probabilities remain positive and sum to one, the elements $L_{i,j}$ necessarily obey two simple conditions: (i) $L_{i\neq j} \geq 0$, (ii) $\sum_i L_{i,j} = 0$.

In the quantum setting, the density matrix ρ plays the analogous role to that of the classical probability distribution, but the quantum master equations are still determined by a Liouvillian:

$$\dot{\rho} = \mathcal{L}(\rho). \tag{2}$$

In his seminal 1976 paper [14], Lindblad established the general form

$$\mathcal{L}(\rho) = i[\rho, H] + \sum_{\alpha,\beta} G_{\alpha,\beta}\Big(F_\alpha \rho F_\beta^\dagger - \frac{1}{2}\{F_\beta^\dagger F_\alpha, \rho\}_+\Big). \tag{3}$$

Here, H is the Hamiltonian of the system, G is a positive semi-definite matrix and, along with the matrices F_α, describes decoherence processes. ($[\cdot,\cdot]$ and $\{\cdot,\cdot\}_+$ denote respectively the commutator and anti-commutator.) Master equations of *Lindblad form* have become the mainstay of the dynamical theory of open quantum systems, and are crucial to the description of quantum mechanical experiments [15].

What is the best possible data that an experimentalist could conceivably gather about an evolving system? They could repeatedly prepare the system in any desired initial state, allow it to evolve for some period of time, and then perform any desired measurement. In fact, for a careful choice of initial states and measurements, it is possible to reconstruct a complete "snapshot" of the system dynamics at any particular time. In the quantum setting, this technique is known as *quantum process tomography* [13]. Full quantum process tomography is now routinely carried out in many different physical systems, from NMR [3–6] to trapped ions [7, 8], from photons [9] to solid-state devices [10].

A tomographic snapshot tells us *everything* there is to know about the evolution at the time t when the snapshot was taken. Each snapshot gives us a dynamical map $\mathcal{E}_t$, which describes how the initial state, $\boldsymbol{p}_0$ or ρ_0, is transformed into $\boldsymbol{p}(t) = \mathcal{E}_t(\boldsymbol{p}_0)$ or $\rho(t) = \mathcal{E}_t(\rho_0)$. Any measurement at time t can be viewed as an imperfect version of process tomography, since it gives partial information about the snapshot. Thus the most complete data that can be gathered about a system's dynamics consists of a set of snapshots. So, finally, we will allow our hypothetical experimentalist the ability to gather sufficient data to reconstruct as many complete snapshots of the dynamics as desired.

Let us concentrate first on the quantum case. The most general possible quantum dynamical map is described mathematically by a completely positive, trace-preserving (CPT) map. CPT maps are widely used in quantum information theory, where they are often called *quantum channels* [13].

The problem of deducing the dynamical equations from measurement data is then one of finding a Lindblad master equation (2) that accounts for the snapshots $\mathcal{E}_t$. This is essentially the converse problem to that considered by Lindblad [14, 16]. Given its relevance, it is not surprising that numerous heuristic numerical techniques have been applied to tackle this problem [5, 17–19]. But they give no guarantee as to whether a correct answer has been found.

Before tackling the problem of finding dynamical equations, let us start by considering an apparently much simpler question: given a *single* snapshot $\mathcal{E}$, does there even *exist* a Liouvillian $\mathcal{L}$ that could have generated it? Not every physically possible dynamical map can be generated by a master equation [20, 21], so the question of the existence of a Liouvillian is a well-posed problem. A dynamical map that *is* generated by some Liouvillian is said to be *Markovian*, so this problem is sometimes referred to as the *Markovianity problem*. Non-Markovian snapshots [22, 23] can arise if the environment carries a memory of the past. Then the system's evolution cannot be described by eqs. (2) and (3) since these treat the environment as a stationary bath.

Returning, then, to the Markovianity problem, it is convenient to represent a snapshot $\mathcal{E}$ by a matrix E,

$$E_{i,j;k,l} = \mathrm{Tr}\big[\mathcal{E}\big(|i\rangle\langle j|\big) \cdot |k\rangle\langle l|\big] \tag{4}$$

(the row- and column-indices of E are the double-indices i, j and k, l, respectively). Looked at this way, each measurement that is performed pins down the values of some of these matrix elements [13]. A snapshot of a Markovian evolution is then one with a Liouvillian $\mathcal{L}$ (represented in the same way by a matrix L) such that $E = e^L$ and, for all times $t \geq 0$, $E_t = e^{Lt}$ are also valid quantum evolutions. Since a snapshot can only ever be measured up to some experimental error, we should be satisfied if we can answer the Markovianity question for any approximation E' to the measured snapshot E, as long as the approximation is accurate up to the experimental error. This is known as a *weak membership* formulation of the problem.

The Markovianity problem can be transformed into an equivalent question about the Liouvillian. Inverting the relationship $E = e^L$, we have $L = \log E$. There are, however, infinitely many possible branches of the logarithm, since the phases of complex eigenvalues of E are only defined modulo $2\pi i$. The problem then becomes one of determining whether *any one of these* is a valid Liouvillian (i.e. of Lindblad form (3)). For the matrix L, this translates into the following necessary and sufficient conditions on L [21]:

(i). L^Γ is Hermitian, where the Γ operation $|i,j\rangle\langle k,l|^\Gamma = |i,k\rangle\langle j,l|$ is defined by its action on the basis elements.

(ii). L fulfils the normalisation $\langle\omega|\, L = 0$, where $|\omega\rangle = \sum_i |i,i\rangle / \sqrt{d}$ is maximally entangled.

(iii). L satisfies *conditional complete positivity* (ccp),

$$(\mathbb{1} - \omega)L^\Gamma(\mathbb{1} - \omega) \geq 0,\ \omega = |\omega\rangle\langle\omega|\,. \tag{5}$$

All branches L_m of the logarithm can be obtained by adding integer multiples of $2\pi i$ to the eigenvalues of the principle branch L_0, so we can parametrise all the possible branches by a set of integers m_c:

$$\begin{aligned} L_m = \log E &= L_0 + 2\pi i \sum_c m_c\big(\,|l_c\rangle\langle r_c| - \mathbb{F}(|l_c\rangle\langle r_c|)\big) \\ &= L_0 + \sum_c m_c A^{(c)}, \end{aligned} \tag{6}$$

$$A^{(c)} = 2\pi i\big(|l_c\rangle\langle r_c| - \mathbb{F}(|l_c\rangle\langle r_c|)\big), \tag{7}$$

with $|l_c\rangle$ and $\langle r_c|$ the left- and right-eigenvectors of E. $\mathbb{F}$ is the operation $\mathbb{F}(|i,j\rangle\langle k,l|) = |j,i\rangle\langle l,k|^*$, where * denotes the complex-conjugate.

To address this *Liouvillian problem*, we will require some basic concepts from complexity theory. Recall that P is the class of computational problems that can be solved efficiently on a classical computer. The notorious class NP instead only requires an efficient verification of solutions, and contains problems which we do not know how to solve efficiently, including such tough chestnuts as 3SAT, travelling salesman, and factoring. A problem is *NP-hard* if solving it efficiently would also lead to efficient solutions to *all* other NP problems. (A problem that is NP-hard and is also itself in the class NP is said to be *NP-complete*.)

We will prove that the Liouvillian problem is NP-hard, by showing how to encode any instance of the well-known NP-complete 3SAT problem into it. Recall that the task in 3SAT is to determine whether a given logical expression can be satisfied or not. The expression is made up of "clauses", all of which must be satisfied simultaneously. Each clause involves three boolean variables (variables with values "true" or "false") or their negations, and is satisfied if and only if at least one of the three is true.

Note that by representing the boolean variables by integers $m_c = 0, 1$, a 3SAT clause containing, say, the variables m_i, the negation of m_j, and m_k can be expressed as

$$m_i + (1 - m_j) + m_k \geq 1. \tag{8}$$

Compare this to condition (iii), above. If the matrices L_0 and $A^{(c)}$ were diagonal, and if we could arrange that, for each diagonal element, only three of the $A^{(c)}$ were non-zero, then condition (iii) would be a concise way of writing a set of inequalities of the form of eq. (8). If, furthermore, each $A^{(c)}$

contained some additional non-zero diagonal entries which were zero in all the other $A^{(c)}$, then we could also write

$$m_c \geq 0 \qquad \text{and} \qquad -m_c \geq -1, \tag{9}$$

which constrain the integers m_c to be either 0 or 1. condition (iii) can then be used to encode any 3SAT problem.

If condition (iii) were the only condition required for L_m to be a valid Liouvillian, and if L_0 and $A^{(c)}$ could be chosen independently, this would allow us to rewrite any 3SAT problem as a Liouvillian problem. In reality, diagonal L_0 and $A^{(c)}$ matrices will never satisfy the other conditions (i) and (ii), and the matrices certainly cannot be chosen independently, since they arise from eigenvectors and eigenvalues of E. However, a more elaborate construction *does* allow any 3SAT problem to be encoded in this way into valid L_0 and $A^{(c)}$ matrices, such that conditions (i) and (ii) are always satisfied, and condition (iii) is equivalent to the inequalities of eqs. (8) and (9) (for details see [24].) Since the Liouvillian problem is equivalent to the Markovianity problem, this proves that the Markovianity problem is itself NP-hard. The above setting now easily generalizes to the original question of *finding* which dynamical equations (if any) could have generated a given set of snapshots [24]: Any method of finding dynamical equations consistent with the data would allow us to solve all NP problems.

What of the classical setting? The classical analogue of the Markovianity problem is the so-called *embedding problem* for stochastic matrices, originally posed in 1937 [2]. Despite considerable effort [25] the general problem has, however, remained open until now [26]. Strictly speaking, the quantum result does not directly imply anything about the classical problem. Nevertheless, the arguments used in the quantum setting can readily be adapted to the classical embedding problem [24], proving that this is NP-hard, too.

These results show that determining the dynamical equations of a system is in general a computationally intractable problem. This also implies that various closely related problems, such as finding the dynamical equation that best approximates the data, or testing a dynamical model against experimental data, are also NP-hard, as solutions to these could easily be used to solve the Markovianity or embedding problem.

Finally, note that this work also leads to rigorous algorithms for extracting the underlying dynamical equations from experimental data. Although these algorithms are not efficient for systems with many degrees of freedom, the complexity-theoretic results show that they are nonetheless optimal. For systems with few effective degrees of freedom, such as all current quantum experiments [3–10], this gives the first practical and provably correct algorithm for this key task.

Acknowledgements The authors would like to thank J. I. Cirac, A. Winter, C. Goldschmidt, and J. Martin for valuable discussions. This work has been supported by a Leverhulme early career fellowship, by the European Commission (QAP, QESSENCE, MINOS, COMPAS, COQUIT, QUEVADIS), the EURYI, and the Danish Research Council (FNU).

[1] S. Cook, `www.claymath.org/millennium/P_vs_NP/pvsnp.pdf`(2000).
[2] G. Elfving, Acta Soc. Sei. Fennicae, n. Ser. A **2** (1937).
[3] M. Nielsen, E. Knill, and R. Laflamme, Nature **396**, 52 (1998).
[4] L. M. K. Vandersypen and I. L. Chuang, Rev. Mod. Phys. **76**, 1037 (2004).
[5] N. Boulant, T. Havel, M. Pravia, and D. Cory, Phys. Rev. A **67**, 042322 (2003).
[6] J. Emerson *et al.*, Science **317**, 1893 (2007).
[7] M. Riebe *et al.*, Phys. Rev. Lett. **97**, 220407 (2006).
[8] T. Monz *et al.*, Phys. Rev. Lett. **102**, 040501 (2009).
[9] J. L. O'Brien *et al.*, Phys. Rev. Lett. **93**, 080502 (2004).
[10] M. Howard *et al.*, New J. Phys. **8**, 33 (2006).
[11] L. Ljung, *System Identification: Theory for the User* (Prentice Hall, 1999).
[12] Nature Insight **453**, 1003 (2008).
[13] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (CUP, 2000).
[14] G. Lindblad, Commun. Math. Phys. **48**, 119 (1976).
[15] H. J. Carmichael, *Statistical Methods in Quantum Optics*, Vol. 1 (Springer, 2003).
[16] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan, J. Math. Phys. **17**, 821 (1976).
[17] M. Howard *et al.*, New Journ. Phys. **8**, 33 (2006).
[18] Y. S. Weinstein *et al.*, J. Chem. Phys. **121**, 6117 (2004).
[19] D. A. Lidar, Z. Bihary, and K. B. Whaley, Chem. Phys. **268**, 35 (2001).
[20] M. M. Wolf and J. I. Cirac, Commun. Math. Phys. **279**, 147 (2008).
[21] M. M. Wolf, J. Eisert, T. S. Cubitt, and J. I. Cirac, Phys. Rev. Lett. **101**, 150402 (2008).
[22] H.-P. Breuer, Phys. Rev. A **75**, 022103 (2007).
[23] M. Scala, B. Militello, A. Messina, J. Piilo, and S. Maniscalco, Phys. Rev. A **75**, 013811 (2007).
[24] T. S. Cubitt, J. Eisert, and M. M. Wolf, Supporting material (online).
[25] J. F. C. Kingman, Z. Wahrscheinlichkeitstheorie **1**, 14 (1962).
[26] A. Mukherjea, in *Matrix theory and applications*, edited by C. R. Johnson (American Mathematical Society, Providence, R.I., 1990)

SUPPORTING MATERIAL

Encoding 3SAT in a Liouvillian

Rather than considering 3SAT per se, it is more convenient to consider the equivalent 1-IN-3SAT problem, into which 3SAT can easily be transformed [1], and which is therefore also NP-complete. Like its close-cousin 3SAT, the task in 1-IN-3SAT is to determine whether or not a given boolean logical expression can be satisfied, where the expression consists of clauses, all of which must be satisfied simultaneously, and each of which involves three boolean variables. However, in 1-IN-3SAT, a clause is satisfied if and only if *exactly one* of the variables appearing in the clause is true (as opposed to 3SAT, in which *at least one* must be true), and no boolean negation is necessary. Note that, in terms of integer variables m_c, a 1-IN-3SAT clause containing variables m_i, m_j and m_k can be expressed as

$$m_i + m_j + m_k = 1, \tag{1a}$$
$$0 \leq m_i, m_j, m_k \leq 1. \tag{1b}$$

The key step in encoding these constraints in a quantum Liouvillian is to restrict our attention to matrices L_0 and $A^{(c)}$ with the following special form:

$$L_0 = 2\pi \sum_{i,j} Q_{i,j}\, |i,i\rangle\langle j,j| + 2\pi \sum_{i\neq j} P_{i,j}\, |i,j\rangle\langle i,j|\,, \tag{2}$$
$$A^{(c)} = 2\pi \sum_{i\neq j} B^{(c)}_{i,j}\, |i,i\rangle\langle j,j| \tag{3}$$

(the factors of 2π are included here for convenience), with coefficient matrices

$$\begin{aligned} Q = {} & \sum_r \boldsymbol{v}_r \boldsymbol{v}_r^T \otimes \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \otimes \begin{pmatrix} k+\lambda_r & \lambda_r \\ \lambda_r & k+\lambda_r \end{pmatrix} \\ & + \sum_c \boldsymbol{v}_c \boldsymbol{v}_c^T \otimes \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} \otimes \begin{pmatrix} k & -\frac{1}{3} \\ \frac{1}{3} & k \end{pmatrix} \\ & + \sum_{c'} \boldsymbol{v}_{c'} \boldsymbol{v}_{c'}^T \otimes \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} \otimes \begin{pmatrix} k & 0 \\ 0 & k \end{pmatrix}, \end{aligned} \tag{4}$$

$$B^{(c)} = \boldsymbol{v}_c \boldsymbol{v}_c^T \otimes \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} \otimes \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \tag{5}$$

The sets of real vectors $\{\boldsymbol{v}_r\}$ and $\{\boldsymbol{v}_c, \boldsymbol{v}_{c'}\}$ should each form an orthogonal basis, and the parameters k and λ_r are also real. The advantage of this restriction is that the action of the Γ operation on matrices of this form is somewhat easier to analyse, as can readily be seen from its definition (given in condition (i) on page 2 of the main text).

It is a simple matter to verify that the eigenvalues and eigenvectors of L_0 and $B^{(c)}$ do indeed parametrise the logarithms of a matrix E, and that the Hermiticity and normalisation conditions required for L to be a valid quantum Liouvillian (conditions (i) and (ii) on page 2 of the main text), are indeed satisfied by the forms given in eqs. (2) to (5), as long as $\boldsymbol{w}^T Q = 0$ and $\mathrm{diag}[P]^\Gamma$ is Hermitian (where for d–dimensional Q, $\boldsymbol{w} = (1,1,\ldots,1)^T/\sqrt{d}$, and $\mathrm{diag}[P]$ denotes the d^2–dimensional matrix with $P_{i,j}$ down its main diagonal). Furthermore, the ccp condition (condition (iii) on page 2 of the main text) reduces for this special form to the pair of conditions:

$$\sum_c B^{(c)}_{i,j}\, m_c + Q_{i,j} \geq 0 \qquad i \neq j, \tag{6a}$$
$$\left(\mathbb{1} - \boldsymbol{w}\boldsymbol{w}^T\right) [\,\mathrm{diag}\, Q + \mathrm{offdiag}\, P]\left(\mathbb{1} - \boldsymbol{w}\boldsymbol{w}^T\right) \geq 0, \tag{6b}$$

where $M = [\mathrm{diag}\, Q + \mathrm{offdiag}\, P]$ denotes the d–dimensional matrix with diagonal elements $M_{i,i} = Q_{i,i}$ and off-diagonal elements $M_{i\neq j} = P_{i,j}$.

We encode a 1-IN-3SAT problem into these matrices by writing the clauses into the vectors $\boldsymbol{v}_c$. Denote the total number of variables and clauses by V and C, respectively. For each clause n involving the i^{th}, j^{th} and k^{th} boolean variables, write a “1” in the n^{th} element of $\boldsymbol{v}_i$, $\boldsymbol{v}_j$ and $\boldsymbol{v}_k$, and write a “0” in the same element of all the other $\boldsymbol{v}_c$’s. Now, for each $\boldsymbol{v}_c$, write a “1” in its $C+c$’th element, writing a “0” in the corresponding element of all the other vectors. Finally, extend the vectors so that they are mutually orthogonal and have the same length, which can always be done. This produces vectors with at most $C + 2V$ elements.

This procedure encodes the coefficients for the 1-IN-3SAT inequalities into some of the on-diagonal 4×4 blocks of the $B^{(c)}$ matrices. Specifically, if we imagine colouring $B^{(c)}$ in a chess-board pattern (starting with a “white square” in the top-leftmost element), then the coefficients for one constraint from eqs. (1a) and (1b) are duplicated in all the “black squares” of one diagonal 4×4 block.

Colouring Q in the same chess-board pattern, the contribution to its “black squares” from the first term of eq. (4) is generated by the off-diagonal elements λ_r:

$$\sum_r \boldsymbol{v}_r \boldsymbol{v}_r^T \otimes \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \otimes \begin{pmatrix} \cdot & \lambda_r \\ \lambda_r & \cdot \end{pmatrix} = S \otimes \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \otimes \begin{pmatrix} \cdot & 1 \\ 1 & \cdot \end{pmatrix}. \tag{7}$$

Since $\boldsymbol{v}_r$ and λ_r can be chosen freely, the first tensor factor in this expression is just the eigenvalue decomposition of an arbitrary real, symmetric matrix S. If we choose the first C diagonal elements of S to be $1/2$, and choose the next V diagonal elements to be $5/6$, then it is straightforward to verify that the equations in the ccp condition of eq. (6a) corresponding to the “black squares” in on-diagonal 4×4 blocks are given by

$$\begin{aligned} m_i, m_j, m_k &\geq -\frac{1}{2}, & -m_i, m_j, m_k &\geq -\frac{7}{6}, \\ m_i + m_j + m_k &\geq \frac{1}{2}, & -m_i - m_j - m_k &\geq -\frac{3}{2}, \end{aligned} \tag{8}$$

for all m_i, m_j, m_k appearing together in a 1-IN-3SAT clause. Since the m_c are integers, these inequalities are exactly equivalent to the 1-IN-3SAT constraints of eqs. (1a) and (1b).

We have successfully encoded the correct coefficients and constants into certain matrix elements of $B^{(c)}$ and Q. But all the other elements of these matrices also generate inequalities via eq. (6a). To “filter out” these unwanted inequalities, we choose the remaining diagonal elements and all off-diagonal elements of the symmetric matrix S to be large and positive, thereby ensuring all unwanted inequalities are always trivially satisfied.

L_m, as constructed so far, will not satisfy the normalisation condition (condition (ii) on page 2 of the main text). For that, we need to ensure that $\boldsymbol{w}^T Q = 0$, i.e. that the columns of Q sum to zero. We use the "white squares" of Q, generated by the diagonal elements in the third tensor factors of eq. (4), to renormalise these column sums to zero. Recall that both $\{\boldsymbol{v}_r\}$ and $\{\boldsymbol{v}_c, \boldsymbol{v}_{c'}\}$ are complete sets of mutually orthogonal vectors. Rearranging eq. (4), Q is therefore given by

$$Q = k\mathbb{1} + S \otimes \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \otimes \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} + \sum_c \boldsymbol{v}_c \boldsymbol{v}_c^T \otimes \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} \otimes \begin{pmatrix} 0 & -\frac{1}{3} \\ \frac{1}{3} & 0 \end{pmatrix}, \quad (9)$$

where $\mathbb{1}$ is the identity matrix. Now, the only requirement on the off-diagonal elements of S is that they be sufficiently positive. Also, from the form of eq. (9), the columns in any individual 4×4 block of Q sum to the same value. Thus, by adjusting the elements of S, we can ensure that all columns of $Q - k\mathbb{1}$ sum to the *same* positive value, σ say. Choosing $k = -\sigma$, the negative on-diagonal element in each column generated by the $k\mathbb{1} = -\sigma\mathbb{1}$ term will cancel the positive contribution from the other terms, thereby satisfying the normalisation condition, as required.

Finally, we must ensure that the second ccp condition of eq. (6b) is always satisfied:

Lemma 1 *If $Q = -k\mathbb{1}$ is d-dimensional, then for any real k there exists a matrix P such that* $\operatorname{diag} P = 0$ *and* $(\mathbb{1} - \boldsymbol{w}\boldsymbol{w}^T)(Q + P)(\mathbb{1} - \boldsymbol{w}\boldsymbol{w}^T) \geq 0$*, where* $\boldsymbol{w} = (1, 1, \ldots, 1)^T/\sqrt{d}$.

Proof Choose $P = \alpha(\mathbb{1} - \boldsymbol{w}\boldsymbol{w}^T) + \alpha(1-d)\boldsymbol{w}\boldsymbol{w}^T$. Then the diagonal elements of P are

$$P_{i,i} = \alpha\left(1 - \frac{1}{d}\right) + \alpha(1-d)\frac{1}{d} = 0, \quad (10)$$

and

$$(\mathbb{1} - \boldsymbol{w}\boldsymbol{w}^T)(Q + P)(\mathbb{1} - \boldsymbol{w}\boldsymbol{w}^T) = (\alpha - k)(\mathbb{1} - \boldsymbol{w}\boldsymbol{w}^T), \quad (11)$$

which is positive semi-definite for $\alpha \geq k$. □

The coefficients $P_{i,j}$ in eq. (2) can be chosen freely, since they play no role in either the normalisation or in encoding 1-IN-3SAT, so the $[\operatorname{offdiag} P]$ term in the ccp condition of eq. (6b) can be chosen to be any matrix with zeros down its main diagonal. Also, from eq. (9), all diagonal elements of Q are equal to $k = -\sigma$. Thus eq. (6b) is exactly of the form given in lemma 1, and choosing P accordingly ensures that it is always satisfied.

We have constructed L_0 and $A^{(c)}$ such that there exists an L_m satisfying conditions (i) to (iii) on page 2 of the main text if (and only if) the original 1-IN-3SAT instance was satisfiable. But we have already shown that condition (iii), along with conditions (i) and (ii), are satisfied if (and only if) L_m is of Lindblad form, which in turn is equivalent to $E = e^{L_m} = e^{L_0}$ being Markovian.

Furthermore, the integer solutions of eqs. 8 are insensitive to small perturbations of the coefficients and constants, so any sufficiently good approximation E' will still be Markovian if E is, and vice versa, as long as we impose sufficient precision requirements. Indeed, it is natural to expect that if a snapshot E is close to being Markovian, it will have a generator L_m that is close to being of Lindblad form. Making this rigorous is less trivial, but follows from continuity properties of the matrix exponential [2] and logarithm [3]. The Markovianity problem is therefore equivalent to the problem of determining whether any L_m obeys the three conditions (i) to (iii), up to the necessary level of approximation. Thus we have successfully encoded 1-IN-3SAT into the Liouvillian problem, such that the corresponding snapshot E is Markovian if (and only if) the 1-IN-3SAT instance was satisfiable.

By applying standard perturbation theory results for eigenvalues and eigenvectors [4, 5], A careful analysis reveals that a precision of $O(V^{-1}(C + 2V)^{-3})$ is sufficient, which scales only polynomially with the number of degrees of freedom in the system (i.e. with the size of the Liouvillian matrix). Though a polynomial scaling is not strictly speaking necessary to prove NP-hardness, it does make the result more compelling, as it shows that the complexity does not depend on imposing unreasonable precision requirements. This is sometimes called *strong NP-hardness* of a weak-membership problem (cf. Refs [6]).

This weak-membership formulation allowing for approximate answers, notably, is necessary if the question is to be reasonable from an experimental perspective, and can only make the problem easier than requiring an exact answer. In fact, it is also necessary from a theoretical perspective; note that if E happened to be close to the boundary of the set of Markovian maps, then it would be close to both Markovian and non-Markovian maps, and an exact answer could require the matrix elements of E to be specified to infinite precision, which is not reasonable even theoretically.

Several snapshots

Clearly, if we can *find* a set of dynamical equations whenever they exist, we can also determine *whether* they exist. So finding the dynamical equations is at least as hard as answering the existence question. For a single snapshot, the latter is just the Markovianity problem again. But, having constructed L_0 and $A^{(c)}$ as described above, it is easy to generalise this to any number of snapshots $\mathcal{E}_t$: simply take $E_t = e^{L_0 t}$ for as many different times t as desired.

The classical setting

The analogue of the Markovianity problem in the classical setting is known as the *embedding problem*. Given a stochastic matrix, this asks whether it can be generated by any continuous-time Markov process (i.e. by dynamics obeying eq. (1) on page 1 of the main text). The quantum mechanical proof described above does not directly imply anything about the classical problem (nor vice versa). Nevertheless, it turns out that the arguments used in the quantum setting can readily be adapted to the classical embedding problem.

We can reduce the embedding problem to a question about the (classical) Liouvillian, in the same way as in the quantum case. Comparing the conditions for L to be a valid classical Liouvillian (conditions (i) and (ii) on page 1 of the main text) with the matrices Q and $B^{(c)}$ from eqs. (4) and (5), we see that $Q + \sum m_c B^{(c)}$ is a valid *classical* Liouvillian if and only if the 1-IN-3SAT problem was satisfiable. In other words, for the classical case, we simply need to use the matrices Q and $B^{(c)}$, rather than the full matrices L_0 and $A^{(c)}$ used in the quantum construction. The rest of the arguments proceed as in the quantum case, thereby proving that the embedding problem too is NP-hard.

[1] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness* (W. H. Freeman, 1979).

[2] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis* (CUP, 1994).

[3] J. Weilenmann, Journ. Func. Anal. **27**, 1 (1978).

[4] R. A. Horn and C. R. Johnson, *Matrix Analysis* (CUP, 1990).

[5] G. H. Golub and C. F. van Loan, *Matrix Computations*, 3rd ed. (Johns Hopkins University Press, 1996).

[6] S. Gharibian, arXiv:0810.4507(2008)