This item is the archived peer-reviewed author-version of:

Stochastic gradient ascent outperforms gamers in the Quantum Moves game

Reference:

Sels Dries.- Stochastic gradient ascent outperforms gamers in the Quantum Moves game
Physical review A - ISSN 2469-9926 - 97:4(2018), 040302(R)
Full text (Publisher's DOI): <https://doi.org/10.1103/PHYSREVA.97.040302>

Stochastic gradient ascend outperforms gamers in Quantum Moves

Dries Sels

Department of Physics, Boston University, Boston, MA 02215, USA

*Department of Physics, Harvard University, 17 Oxford st., Cambridge, MA 02138, USA and
Theory of quantum and complex systems, Universiteit Antwerpen, B-2610 Antwerpen, Belgium*

(Dated: March 22, 2018)

In a recent work on quantum state preparation, Sørensen and colleagues[1] explore the possibility of using video games to help design quantum control protocols. The authors present a game called “Quantum Moves”[2] in which gamers have to move an atom from A to B by means of optical tweezers. They report that, *players succeed where purely numerical optimization fails*[1]. Moreover, by harnessing the player strategies they can *outperform the most prominent established numerical methods*[1]. The aim of this manuscript is to analyze the problem in detail and show that those claims are untenable. In fact, without any prior knowledge and starting from a random initial seed, a simple stochastic local optimization method finds near optimal solutions which outperform all players. Counter-diabatic driving can even be used to generate protocols without resorting to numeric optimization^a. The analysis results in an accurate analytic estimate of the quantum speed limit which, apart from zero-point motion, is shown to be entirely classical in nature. The latter might explain why gamers are reasonably good at the game. A simple modification of the BringHomeWater challenge is proposed to test this hypothesis.

The entire gamification research program—the use of games to channel human brain-power to solve research problems—is based on the sole premise that “*Humans are better than computers at certain tasks because of their intuition and superior visual processing*” [3]. Previous examples of citizen-science games such as Foldit [4] and Eyewire [5] are 3-D puzzle games. Their success is attributed to our superior visual insight in 3D shapes and our capability to explore different strategies [4]. But what if our visual processing is impaired by our inability to imagine anything in more than 3 dimensions? Or our intuition about the problem is simply wrong. Are our heuristics any better than machines at that point? Recent progress in machine learning seems to indicate they are not. With almost super-human performance in Go [6], computers have mastered to most complex of classical games using a combination of reinforcement learning and supervised learning. Deep convolutional neural networks have resulted in unseen machine performance in the visual domain [7] and deep reinforcement learning [8] has produced a computer agent which rivals professional players in many Atari 2600 games. In light of these results one might wonder how players of Quantum Moves (Fig. 1) were able to outperform purely numerical optimization methods designed for quantum optimal control problems. Given the bizarre and counterintuitive nature of quantum mechanics one could even wonder how players came up with a reasonable strategy to begin with.

The goal of this work is to elucidate this tension. The manuscript starts by introducing the concept of counter-diabatic driving [9–12]. The latter formalizes the intuitive notion of players to move as fast as possible without spilling [3]. The idea can readily be applied to a single op-

Figure 1. **Quantum moves interface** Screenshot of the BringHomeWater challenge developed by Sørensen et al. [1]. In the game, players move an optical tweezer (blue circle is the cursor to move the tweezer) from an initial point on the left towards an atom trapped in a potential well on the right. The atom is depicted as the purple liquid filling up the potential-well. Players have to find a route to bring this liquid to the target (purple shaded region on the left) as fast as possible without spilling it.

tical tweezer and it is shown how this completely predicts the quantum speed limit. A recent result by Jarzynski et al. [13] is used to generalize it to the double tweezer situation. Protocols are compared to those obtained from purely numeric optimization and the convergence of the numeric algorithm is discussed. In short, all methods are shown to outperform human gamers.

^a As E. Wigner once said: *It is nice to know that the computer understands the problem. But I would like to understand it too.*

Figure 2. **Effective potential for moving objects** The potential of a static optical tweezer is depicted by the dashed blue line. The atom however feels a different potential once the tweezer is moving. By going to the co-moving frame the atom picks up an additional pseudo-force if the tweezer accelerates, just like a person in an elevator. For a tweezer of finite depth, it unfortunately allows atoms to escape if the acceleration is too large.

I. QUANTUM STATE PREPARATION

Let's consider first what happens to an atom in the ground state of just a single optical tweezer, i.e. the system's Hamiltonian is

$$H = \frac{p^2}{2m} + V(x - x_0(t)),$$

where $V(x)$ denotes the tweezer potential (Fig. 2) and $x_0(t)$ the tweezer's position. The goal is to move the tweezer without exciting the system. Since the entire system is simply being translated in space, the adiabatic gauge potential [12] is the momentum operator, i.e. the generator of translations. Consequently, the counter-diabatic Hamiltonian becomes

$$H_{CD} = \frac{p^2}{2m} + V(x - x_0(t)) + \dot{x}_0 p$$

For a more detailed discussion on counter-diabatic driving I refer to ref. [14]. Note that we can gauge transform this into an equivalent Hamiltonian which accomplishes the same task as long as we demand that the velocity $\dot{x}_0$ vanishes in the beginning and at the end of the protocol. Just like in electromagnetism, we turn the curl-free vector potential into a scalar potential, which results in

$$H'_{CD} = \frac{p^2}{2m} + V(x - x_0(t)) - m\ddot{x}_0 x \quad (1)$$

As expected the atom can be kept in the ground state even when the tweezer is moving as long as we cancel out the *pseudo-force* due to the acceleration of the system (Fig. 2). This is not completely within our reach since we can only tune the depth and the position of the tweezer. The best we can do is, shift the tweezer a bit such that it effectively generates the counter-diabatic force on the atom. Since the trap is not infinitely deep, the maximal

force it can generate is limited and the above procedure can not be implemented once the pseudoforce exceeds this maximal value. Our procedure can only work as long as

$$\max_x |\partial_x V(x)| > ma,$$

where a is the maximal acceleration. For a Gaussian beam with amplitude $\mathcal{A}$, and standard deviation σ

$$a < \frac{\mathcal{A}}{m\sigma\sqrt{e}}.$$

The smaller the acceleration, the better. One way to do so is of course to go to the adiabatic limit but in this case we want to do it in fixed duration T . We thus want a protocol $x_0(t)$ that covers a distance L , let's say from $x_i = L/2$ to $x_f = -L/2$, with the smallest magnitude of the acceleration. This would of course be a straight line or geodesic, which has zero acceleration. Recall however, that the velocity had to vanish at the beginning and the end of the protocol as well. Since we can not make the acceleration zero all the time let us simply minimize the average magnitude of the acceleration over the entire protocol. The latter just results in a cubic polynomial:

$$x_0(t) = L \left(2 \left(\frac{t}{T} \right)^3 - 3 \left(\frac{t}{T} \right)^2 + \frac{1}{2} \right). \quad (2)$$

This protocol has maximal acceleration at the beginning and the end, namely $a = 6L/T^2$. Combining this with the above restriction on the acceleration we find a constraint on the protocol duration,

$$T > T_{\text{CSL}} = \left(\frac{6mL\sigma\sqrt{e}}{\mathcal{A}} \right)^{1/2}$$

Plugging in the actual parameters used in Ref. [1], being $\mathcal{A} = 130$ to 160 , $m = 1$, $\sigma = 1/8$, $L = 1.1$, we find $T_{\text{CSL}} = 0.092$ to 0.102 . To find the approximate counter diabatic protocol one can simply approximate the tweezer potential by a harmonic potential with frequency $m\omega^2 = \mathcal{A}/\sigma^2$. For the harmonic oscillator the counter-diabatic protocol is simply given by [12]

$$x_{\text{CD}}(t) = x_0(t) + \frac{\ddot{x}_0(t)}{\omega^2}. \quad (3)$$

Moreover, by eliminating the tweezer amplitude in favor of the effective trap frequency, the classical speed limit becomes

$$T_{\text{CSL}} \approx \frac{\pi}{\omega} \left(\frac{L}{\sigma} \right)^{1/2}$$

Up to this point the derivation was completely classical, which has resulted in the complete absence of Planck's constant in the minimal required duration. One might wonder, how quantum mechanics will change these

Figure 3. **Fidelity for different protocol durations** The green shaded region indicates the classical speed limit for moving a particle in a finite depth potential. The fidelity of the associated quantum counter-diabatic driving (dashed green) protocol shows a sharp rise of fidelity near T_{CSL} . For the double tweezer, the protocol is adapted to account for an additional geometric contribution to the gauge potential. Up to T_{CSL} , its performance (blue line) is comparable to that obtained from numeric optimization (black dots). At later times, its performance starts to wane because it is derived from an approximate adiabatic gauge potential. Details on the numeric optimization algorithm can be found in section III. Gamers performance and KASS results are shown for comparison and taken from Ref. [1]; for that reason we adopt the same units $\hbar = m = 1$.

results. On the short timescale set by T_{CSL} , the system is still in the semi-classical limit. It's therefore reasonable to say that the dominant quantum contribution is the zero-point fluctuations in the initial state. Instead of having a sharp transition from perfectly controllable to completely uncontrollable there will be a smooth crossover around T_{CSL} . Numerical simulation of the quantum dynamics of the single tweezer counter-diabatic protocol confirms this argument (Fig. 3). The green dashed line depicts the fidelity obtained by applying (4), and the cubic trial (2). The fidelity reaches 1/2 at $T = 0.1$.

So far we have only been concerned with a single tweezer. While numerical results indicate that the fidelity for the single tweezer is comparable to that of the actual two tweezer problem, it remains a challenge to find the actual protocol. In the two tweezer problem, the original tweezer in which the atom is trapped can not be moved, and a second movable tweezer has to transport the atom (Fig. 1). It would be unwise to simply use (4) as it hinges on the fact that the adiabatic gauge potential is simply the momentum operator. The adiabatic gauge potential for the two tweezer problem is much more complicated and in general there is no hope that we can turn it into an effective scalar potential like before. For the ground-state of the system, it is however possible to design an approximate counter-diabatic gauge field that captures almost all of the physics and constitutes only a small (but important) correction to single tweezer result:

$$\begin{aligned} \text{single tweezer} &\Leftrightarrow \text{double tweezer} \\ \dot{x}_0 p &\quad \dot{x}_0 \sqrt{g(x_0)} p, \end{aligned}$$

where $g(x_0)$ serves as an effective metric tensor. The exact details are given in section III, but the result has a very intuitive explanation. The additional correction $\sqrt{g}$ is caused by the fact that the equilibrium position of an atom in the double tweezer potential does not always coincide with the minimum of the potential of the moving tweezer. In other words g effectively contracts and dilates space to account for the discrepancy between tweezer's and the atom's position. Since the gauge potential is the generator of adiabatic transformations, it translates the atom to its new equilibrium position as the tweezer is moved. When the two tweezers are close together the equilibrium position shifts slower than the tweezer, such that $g < 1$. For two far separated tweezers the single tweezer result should hold, implying $g \rightarrow 1$. In the middle, however, there is a region where the minimum of the potential quickly shifts between the two tweezers. The latter results in a significant peak in the metric (Fig. 4). One can thus avoid most non-adiabatic effects by moving the tweezer on a geodesic of the metric g . Just like for the single tweezer one has to make sure that the velocity goes to zero in the beginning and the end of the protocol. At present, this is done by simply adding a region of constant acceleration at the beginning and the end. Just like for the single tweezer we find an effective protocol by correcting the trial protocol for its deviation from the geodesic:

$$x_{CD}(t) = x_0(t) + \frac{1}{\omega^2} \frac{\partial}{\partial t} \left(\sqrt{g(x_0)} \frac{\partial x_0(t)}{\partial t} \right). \quad (4)$$

The performance of these protocols is summarized in Fig. 3. Again there is a sharp rise of the fidelity near the classical speed limit. For longer protocol durations, its performance starts to falter. This suboptimal performance is entirely caused by the approximate nature of the adiabatic gauge potential. At time $T > T_{CLS}$ there starts to be a significant discrepancy between the numerical best result (see Fig. 3) and the counter-diabatic protocol. However, it still grossly outperforms any Quantum Moves player and so does the numerical optimal control result. The details on the numerics can be found in section III. By taking piecewise constant protocols and only allowing a discrete set of amplitudes it becomes possible to efficiently implement a local stochastic ascend method. The results in Fig. 3 show the outcome of a typical stochastic run starting from a random initial seed. In contrast to the KASS results in [1], the method has no problem finding nearly optimal solutions.

II. DISCUSSION

We conclude that Quantum Moves gamers did not devise strategies that *go beyond the classical laws of physics* [3]; they also didn't need to. Gamers simply came

Figure 4. **Geometry and geodesic protocol** The left panel shows the effective adiabatic metric g for moving the tweezer. In the middle there is a region where a small change in the position of a tweezer can cause a large change in the shape of the potential energy, at which point the metric becomes much larger than one. The resulting geodesic is shown in right panel; note how it slows down where the metric is large. All quantities are expressed in units of $\hbar = m = 1$.

up with the expected intuitive strategy where you try to move the tweezer without spilling the “water”. They do not outperform any reasonable numerical method and at short times perform much worse than a theoretical designed counter-diabatic protocol. This doesn’t only falsify the thesis of Sørensen and colleagues [1], it also poses serious questions about the generalizability of the result. Both the initial and target state in the BringHomeWater challenge are localized states in position-space, which is also the space in which the tweezer can move. It thus allows to display everything in position-space. It’s been argued that this makes the game clever designed [15], but it also makes the problem easy. Quantum mechanics itself is basis-invariant and any physics can only be established by the trinity of Hamiltonian, observable and quantum state; in this case control, target and initial state. To make the problem quantum mechanically interesting, at least one of them must be local in a basis that is sufficiently different from the others. In that respect, the only quantum mechanically interesting target state in Quantum Moves is the superposition of the atom in the tweezer on the left and on the right, with a well defined phase difference between the two. So here is the real challenge, can gamers even come up with any reasonable strategy (see Fig. 5) to this problem?

III. METHODS

State specific counter-diabatic driving

Counter-diabatic driving is a method to speed up quantum state preparation for adiabatically connected states by adding an additional driving field to the Hamiltonian. In general, the additional Hamiltonian is hard to find. For chaotic systems, the exact field is even ill defined [12, 17]. Recently, some progress has been made to find approximate counter-diabatic drives for complex systems [14]. At present, the situation is slightly simpler

Figure 5. **Fidelity of superposition** Instead of having to move the atom from one tweezer to the other, the goal is to make an equal superposition of the ground state in the tweezer at $L/2$ and at $-L/2$. The green dots show the final outcome of the stochastic ascend method for 200 random initial seeds and the black line shows the best out of that set. For comparison, the fidelity for the original BringHomeWater challenge is also shown in grey. In the BringHomeWater challenge the difference between the worst and best fidelity, obtained from 200 random initial seeds, was only about 1%. In the present challenge there is an intermediate regime where the spread is more than 20%, indicating that the problem is indeed much harder (see also [16]). All quantities are expressed in units of $\hbar = m = 1$.

since we simply deal with ground states of a 1-D single particle problem. A recent idea by Jarzynski and colleagues [13] is particularly suited for this problem. The idea is very simple, i.e. consider a particle described by the wave function $\psi(x, t)$ and subject to the Hamiltonian:

$$H_1(t) = \frac{1}{2} (v(x, t)p + pv(x, t)).$$

From the Schrödinger equation one readily derives the continuity equation for the density $n(x, t) = |\psi(x, t)|^2$,

$$\partial_t n(x, t) = -\partial_x (v(x, t)n(x, t)). \quad (5)$$

Given a target time-dependent density, e.g. an adiabatic ground state, we are trying to find the velocity-field $v(x, t)$ such that we always recover this density. Instead of solving the continuity equation for the density we thus solve for the velocity-field, which results in

$$v(x, t) = -\frac{1}{n(x, t)} \int_{-\infty}^x d\xi \partial_t n(\xi, t) = -\frac{\partial_t I(x, t)}{\partial_x I(x, t)}, \quad (6)$$

where $I(x, t)$ is the cumulative distribution function (cf. expression (12) in [13]). If the density is that of an eigenstate of a time dependent Hamiltonian $H_0(t)$ one can of course add the bare Hamiltonian to $H_1(t)$ to get the full counter-diabatic drive $H_{CD} = H_0 + H_1$. While this will result in perfect state transfer, it will require access to arbitrarily complicated potential energies. In the end, we can only shift the tweezers and so we would like to restrict to spatially homogenous $v(x, t) = v(t)$, as they will result in a linear potential. To get the approximate

velocity field one can simply solve eq. (5) in the least square sense, i.e.

$$v(t) = \underset{w}{\operatorname{argmin}} \int dx (\partial_t n + w \partial_x n)^2, \quad (7)$$

which yields

$$\begin{aligned} v(t) &= -\frac{\int dx (\partial_t n)(\partial_x n)}{\int dx (\partial_x n)^2} = \dot{x}_0 \frac{\int dx (-\partial_{x_0} n)(\partial_x n)}{\int dx (\partial_x n)^2} \\ &\equiv \dot{x}_0 \sqrt{g(x_0)}, \end{aligned}$$

where we have used the fact that the ground-state density is fully parametrized by the position of the tweezer x_0 . This moreover results in a robust numerical scheme to extract $v(t)$. Note that, for a single tweezer $\partial_{x_0} n = -\partial_x n$ such that $v(t) = \dot{x}_0$. For the latter the velocity field is exact, for any other situation it is only an approximation. Finally, one readily identifies g with an effective adiabatic metric tensor, which dilates and contracts space to account for the shift between the tweezer position and the equilibrium position of the atom. Indeed, no pseudo-force will be generated whenever $v(t)$ is constant, a condition which exactly yields the geodesic equation:

$$\ddot{x}_0 + \frac{1}{2} (g^{-1} \partial_{x_0} g) \dot{x}_0^2 = 0.$$

The metric tensor and geodesic for the present — two tweezer problem — can be found in Fig. 4. The density is computed from the ground state of the Hamiltonian:

$$H(x_0) = \frac{p^2}{2m} - \mathcal{A} \exp\left(-\frac{(x - x_0)^2}{2\sigma^2}\right) - \mathcal{B} \exp\left(-\frac{x^2}{2\sigma^2}\right),$$

where $\mathcal{A} = 160$, $\mathcal{B} = 130$, $m = 1$ and $\sigma = 1/8$, in agreement with the parameters reported in [1].

Stochastic ascend optimizer (SGA)

In order to store and optimize the protocols we have to reduce them to a finite number of degrees of freedom. This is done by considering quasi-continuous protocols of duration $\Delta t = T/N$. Moreover, we also quantize the allowed values of the tweezer position with a spacing $\Delta x = \sigma/8$ and allow it to move in the interval $[-1, 1]$. This results in $M = 128$ different tweezer positions x^k . For each of those positions we have a Hamiltonian H_k and an associated unitary

$$U_k = \exp(-iH_k \Delta t / \hbar). \quad (8)$$

All those unitaries can be precomputed and stored, this only has to be done once. The total time evolution operator $U(T)$ for a protocol thus consists of a multiplication of a subset of these unitaries

$$U(T) = U_{k_N} U_{k_{N-1}} \dots U_{k_2} U_{k_1} = \prod_{i=1}^N U_{k_i}, \quad (9)$$

and the fidelity for reaching the target state $\langle \phi |$ from the initial state $|\psi\rangle$ is given by $F = |\langle \phi | U | \psi \rangle|^2$. To optimize the fidelity we now simply do local stochastic ascend optimization, which works as follows:

- 0 Draw a random starting protocol $x = (x_N, x_{N-1}, \dots, x_2, x_1)$
- 1 Randomly permute all integers $1 : N$, let's call this vector I . This will ensure we run over all the timesteps in a random fashion but sweep over the entire protocol before we change a local setting again. Consider the current integer to be W
- 2 Compute $\langle \phi_{W+1} | = \langle \phi | \prod_{i=W+1}^N U_{k_i}$ and $|\psi_{W-1}\rangle = \prod_{i=1}^{W-1} U_{k_i} |\psi\rangle$
- 3 Compute the fidelity $F_k = |\langle \phi_{W+1} | U_k | \psi_{W-1} \rangle|^2$ for all k .
- 4 Select the optimal k and replace the corresponding element in x with x^k , i.e. $x_W = x^k$
- 5 Repeat step 2,3,4 for all the elements of I . If done, we have swept over the entire protocol once and forced the protocol to the best local settings.
- 6 Sweep over the protocol again by going back to step 1. Stop when tolerance is reached or when no element was changed in the previous sweep. The latter implies a local optimum is reached.

Note that the method is relatively efficient, since we can compute the local change in the fidelity (step 3) for all tweezer positions without having to recompute the entire time-evolution (step 2). In stark contrast to the KASS method used in [1]. Since the latter works in Fourier space one must recompute the entire time-evolution for every small change to the parameters. The present method can however not escape out of local optima in the fidelity landscape and we should run it for multiple initial random seeds. A typical set of optimization trajectories is show in Fig. 6. While some initial seeds converge to a minimum very quickly others are substantially slower. Moreover, none of the final protocols are exactly the same but the corresponding fidelities are very similar. In fact all fidelities are somewhere between 0.5311 and 0.5364, hence the best fidelity is only 1% better than the worst. This behavior was also observed in [16] for controlling a single qubit.

ACKNOWLEDGMENTS

I acknowledge financial support of the Research Foundation Flanders—FWO and the BU CMTV program. I thank F. Brosens and A. Polkovnikov for detailed remarks on the manuscript. The work benefited from many discussion with M. Bukov, A. Day, P. Mehta on machine learning.

Figure 6. **Stochastic fidelity traces for $T=0.1$** Figure shows how 100 initial random protocols move through the fidelity landscape. There are $M = 128$ settings and $N = 40$ timesteps, such that a single sweep over the lattice results in 5120 fidelity evaluations. A large fraction of protocols has converged in about 20 000 steps which is as little as 4 sweeps. The slowest protocol took 22 sweeps to end up in an optimum. A single sweep runs in just a few seconds on a laptop.

-
- [1] J. J. W. H. Sørensen, M. K. Pedersen, M. Munch, P. Haikka, J. H. Jensen, T. Planke, M. G. Andreassen, M. Gajdacz, K. Mølmer, A. Lieberoth, and J. F. Sherson, *Nature* **532**, 210 (2016).
 - [2] https://www.scienceathome.org/games/quantum_moves/, (2016).
 - [3] S. Maniscalco, *Nature* **532**, 184 (2016).
 - [4] S. Cooper, F. Khatib, A. Treuille, J. Barbero, J. Lee, M. Beenen, A. Leaver-Fay, D. Baker, Z. Popovic, and F. players, *Nature* **466**, 756 (2010).
 - [5] J. S. Kim, M. J. Greene, A. Zlateski, K. Lee, M. Richardson, S. C. Turaga, M. Purcaro, M. Balkam, A. Robinson, B. F. Behabadi, M. Campos, W. Denk, H. S. Seung, and the EyeWriters, *Nature* **509**, 331 (2014).
 - [6] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, *Nature* **529**, 484 (2016).
 - [7] A. Krizhevsky, I. Sutskever, and G. E. Hinton, in *Advances in Neural Information Processing Systems 25*, edited by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Curran Associates, Inc., 2012) pp. 1097–1105.
 - [8] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, *Nature* **518**, 529 (2015).
 - [9] M. V. Berry, *Journal of Physics A: Mathematical and Theoretical* **42**, 365303 (2009).
 - [10] S. An, D. Lv, A. del Campo, and K. Kim, *Nature Communications* **7**, 12999 EP (2016).
 - [11] S. Deffner, C. Jarzynski, and A. del Campo, *Phys. Rev. X* **4**, 021013 (2014).
 - [12] M. Kolodrubetz, D. Sels, P. Mehta, and A. Polkovnikov, *Physics Reports* (2017).
 - [13] C. Jarzynski, S. Deffner, A. Patra, and S. Yiğit, *Phys. Rev. E* **95**, 032122 (2017).
 - [14] D. Sels and A. Polkovnikov, *Proceedings of the National Academy of Sciences* **114**, E3909 (2017).
 - [15] E. Gibney, *Nature* **532**, 160161 (2016).
 - [16] M. Bukov, A. G. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, arXiv:1705.00565 (2017).
 - [17] C. Jarzynski, *Phys. Rev. Lett.* **74**, 1732 (1995).