6-10-2022

Symmetry and Control in Thermodynamics

Emily Adlam
Chapman University, adlam@chapman.edu

L. Uribarri
Etxea LLC

N. Allen
University of Maryland

Recommended Citation

E. Adlam, L. Uribarri, N. Allen; Symmetry and control in thermodynamics. *AVS Quantum Sci.* 1 June 2022; 4 (2): 022001. https://doi.org/10.1116/5.0065442

Symmetry and Control in Thermodynamics

Comments

This article was originally published in *AVS Quantum Science*, volume 4, issue 2, in 2022. https://doi.org/10.1116/5.0065442

Copyright

RESEARCH ARTICLE | JUNE 10 2022

Symmetry and control in thermodynamics

Special Collection: Quantum Thermodynamics

E. Adlam ; L. Uribarri; N. Allen

AVS Quantum Sci. 4, 022001 (2022)
https://doi.org/10.1116/5.0065442

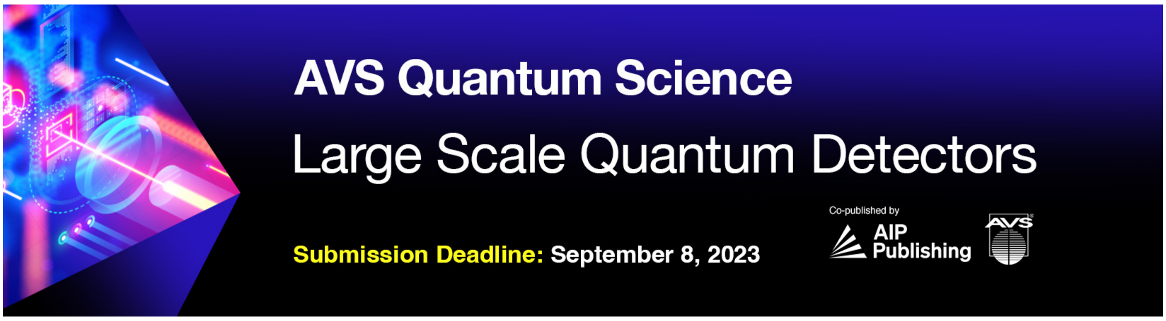

Symmetry and control in thermodynamics

Cite as: AVS Quantum Sci. **4**, 022001 (2022); doi: 10.1116/5.0065442
Submitted: 2 August 2021 · Accepted: 10 May 2022 ·
Published Online: 10 June 2022

View Online Export Citation CrossMark

E. Adlam,[1,2,a)] L. Uribarri,[3] and N. Allen[4]

AFFILIATIONS

[1]Basic Research Community for Physics, Leipzig, Germany
[2]Western University, London, Ontario, Canada, N6A 3K7
[3]Etxea LLC, Arlington, Texas 76010, USA
[4]University of Maryland, College Park, Maryland 20742, USA and University of Queensland, Brisbane, Australia

Note: This paper is a part of the Special Topic Collection on Quantum Thermodynamics.
[a)]Electronic mail: eadlam90@gmail.com

ABSTRACT

We explore the relationship between symmetry and entropy, distinguishing between symmetries of state and dynamical symmetries, and in the context of quantum thermodynamics between symmetries of pure and mixed states. Ultimately, we will argue that symmetry in thermodynamics is best understood as a means of control within the control theory paradigm, and we will describe an interesting technological application of symmetry-based control in the context of a quantum coherence capacitor. Symmetry, the concept from which Noether derived the conservation laws of physics, is one of the most important guiding principles of modern physics. Moreover, symmetry is often regarded as a form of order, and entropy is sometimes regarded as a measure of disorder, so it is natural to suppose that symmetry and entropy are related in some way. In this article, we will explore the relationship between symmetry and entropy, demonstrating that this relationship is by no means a simple one: in particular, it is important to distinguish between symmetries of state and dynamical symmetries, and in the context of quantum thermodynamics to distinguish between symmetries of pure and mixed states. Ultimately, we will argue that symmetry in thermodynamics is best understood as a means of control within the control theory paradigm, and we will describe an interesting technological application of symmetry-based control in the context of a quantum coherence capacitor.

I. THERMODYNAMICS ASACONTROL THEORY

The role of symmetry in thermodynamics has engendered considerable discussion. In particular, links have often been made between Curie's principle (the idea that whatever symmetry or asymmetry is present in an effect must also have been present in its cause) and thermodynamics; we will return to this idea later in the article. It has also been argued that symmetry considerations determine the correct choice of thermodynamic variables—"A very few variables meet the criteria of extensivity and time independence required for macroscopic measurability. Those which qualify do so by virtue of one or another symmetry consideration,"[1] and of course, it is common to invoke symmetry considerations to justify the use of uniform phase space distributions in discussions of the origins of thermodynamics in statistical mechanics. Symmetries also play a major role in the study of phase transitions.[2] However, in this article we will argue for a slightly different conceptualization of the relation between symmetry and thermodynamics—we will argue that symmetries can be regarded as a locus of control within the control theory approach to thermodynamics. This approach also has important practical consequences, as there are many possible technological applications of symmetry-based thermodynamical control.

There is a longstanding tradition of understanding thermodynamics not as a true theory of dynamics but as a practical description of the ways in which agents can manipulate thermodynamical systems. The role of agents in thermodynamics was famously highlighted by Maxwell, who observed that if a diminutive "sentient being" could control things on the level of fundamental particles, perhaps we could escape some of the more burdensome consequences of the second law. The idea that thermodynamics is a theory about the abilities of agents to manipulate systems was subsequently championed by Jaynes[3] and later by Wallace[4] and Myrvold.[5]

Myrvold's approach reinforces the fact that our thermodynamic descriptions typically assume that the certain controllable external parameters are *exogenous*, meaning that we take it that these variables have an effect on the system we are studying, but we disregard any effects of the system we are studying on the variables. For example, when we place an external heat bath next to a thermodynamical system, we take it that the heat bath is infinite, so the thermodynamic system will change in temperature to reach the temperature of the adjacent bath, but this does not have any impact on the temperature of the bath. Myrvold notes that this is always an idealization made for the purpose of calculation; in many real situations, it is a good idealization since the effect of the system on the external variables is negligible, but in cases where the system has a non-negligible effect on the controllable variables this assumption may lead to errors in the long term. We will see an example of such a case in Sec. III C.

In Wallace's account, it is shown that "by treating [any system's dynamics, e.g.,] thermodynamics as a theory of transitions between states resulting from simple controls, one can recover [the system's dynamics, e.g.,] thermodynamics if we do not include measurements which affect subsequent choices of operations [i.e., feedback]." The control operations Wallace considers consist of "doing things to the external parameters," like volume or temperature, such as "(a) smooth modifications to the external parameters over some finite interval of time; and (b) leaving the controlled object alone" for a long enough time for it to reach a new equilibrium. A common realization of these operations in thermodynamics is to bring two control systems into physical contact (through, e.g., a diathermic membrane). Wallace demonstrates that in the absence of feedback, physically possible control processes are limited to those which induce transitions that do not lower Gibbs entropy.

While quantum thermodynamics differs from classical thermodynamics in a number of ways, it turns out that the control theory paradigm is useful here too, although in the context of quantum thermodynamics, the control theory paradigm is usually described as a "resource theory."[6–8] In formulating a resource theory, we identify certain operations which are regarded as "free" in the sense that they can be performed without using external resources, and then, we identify the set of all states which can be prepared using only those operations as free,—all other states are then regarded as a "resource." For example, in the resource theory of quantum thermodynamics, which applies in the regime of a small number of particles interacting with a heat bath, we might identify the energy-conserving operations as free, and thus for a heat bath held at some fixed temperature, the free states consist of those that are in thermal equilibrium with the bath. The resource theory then characterizes which state transitions are possible using only the free, energy-conserving operations. Again, it is clear that this description is to some degree agent-relative, since the set of free operations is typically supposed to coincide with those operations that the relevant class of agents can easily perform. The resource theory formulation then enables us to formulate an analog of the second law for microscopic quantum systems: it turns out that we obtain not just a single constraint on entropy changes in possible transformations but an entire family of constraints, each of which can be regarded as a variant of the second law;[9] and in this case, these "second laws" are explicitly derived not just from the dynamics of the system itself but from the way in which the resource theory limits possible operations.

A valuable approach to making these insights quantitative is afforded by the thermodynamic length, which measures the distance between equilibrium thermodynamic states, which defines a metric on the equilibrium states.[10] Thus, rather than simply asking questions about transitions that thermodynamical systems will spontaneously undergo, we can consider exerting finer control by driving the system along a desired path through the space of equilibrium space. The metric defined by the thermodynamical length can then be used to answer questions about the optimality of various such paths—for example, it can be shown that the protocols with minimal dissipation are geodesics of this metric,[11] that dissipation is inversely proportional to the duration of the protocol, and that the optimal protocol is independent of duration.[12] This methodology can be employed in both classical and quantum thermodynamics, thus allowing us to determine more effective ways of exerting control over thermodynamical systems.

Once we recognize that thermodynamical descriptions are relativized to a choice of macroscopic parameters and/or free operations, it should not be surprising that interesting new thermodynamics can arise when we identify and learn to control some new sort of macroscopic parameter and/or add a new operation to our set of free operations. In this paper, we will argue that symmetries can be used in this way. In Sec. III, we will explore aspects of the role of symmetry in thermodynamics by studying some relationships between symmetry and entropy, and we will argue that symmetry is best understood as a means of control within the control theory paradigm. In Sec. IV, we will see an example where symmetries are used as a way of accessing and controlling phase properties of a thermodynamic system via quantum symmetry-protected states.

II. SYMMETRY-PROTECTED STATES

In this section, we briefly review the physics of symmetry-protected states, which will play an important role in our discussion. The study of quantum entanglement has given rise to a rich understanding of the structure of several classes of entangled states. We typically distinguish between short-range entanglement (SRE),[13] which refers to entangled states which can be transformed into direct product states using only local unitary transformations, and long-range entanglement (LRE), which refers to entangled states that cannot be transformed into direct product states unless we act globally on the whole system. In a gapped quantum system with no symmetries, all SRE states can be transformed into one another using reversible transformations with no energy cost, so we say that they all belong to the same phase, but in a gapped quantum system with symmetry, there are different classes of SRE states which cannot be transformed into one another using local unitary transformations; thus, they belong to different phases. Any Hamiltonian in a symmetry-protected phase can be reversibly transformed into a Hamiltonian whose ground state is a product state, but only if the relevant symmetry can be broken during this transition. Both the symmetry-protected state and the product state respect the symmetry, but they are separated by a "symmetry barrier"—i.e., a wall of states that do not respect the symmetry. We have to pass through the symmetry barrier to get to the product state, so as long as the symmetry cannot be broken, the entanglement will remain robust. Thus, we refer to these states as "symmetry protected topological states" (SPTS). Well-known examples of SPTS include the 1D Haldane phases for a chain of spin-1 systems,[14] and topological insulators.[15] A symmetry barrier of this type may be compared to the "energy barrier" that exists in a classical meta-stable state, as with excited nuclear isomer states of certain radioisotopes;[16] however, in the case of the energy barrier it is necessary to input energy to overcome the barrier, whereas in the case of the symmetry barrier, in principle the symmetry may be removed adiabatically so it may be possible to induce the relevant transition without any input of energy.

For example, in a quantum spin Hall insulator (or "topological insulator"), the edge states are composed of pairs of states with equal energy and opposite spin (Kramers pairs):[17] the up-spin electrons flow in one direction and the down-spin electrons flow in the other direction. Because the states come in matching pairs, elastic back-scattering of a left-moving state to a right-moving state or vice versa would break time-reversal symmetry. Thus, no time-reversal invariant interaction can cause back-scattering, and therefore, if we assume that the electrons are exposed only to time-reversal invariant interactions, no back-scattering will occur and therefore electrons will flow unimpeded, like a current subject to no electrical resistance. Thus, the state is "protected" by the time reversal symmetry.

However, if a "bath" of nuclear spins is present near the edges of the system, the nuclear spins are coupled to the edge electrons by the hyperfine interaction. These hyperfine interactions do not satisfy time-reversal symmetry, and thus, it becomes possible to have some elastic back-scattering where a left-moving electron becomes a right-moving electron or vice versa; every time this happens, a nucleus spin is flipped. Under these circumstances, the symmetry barrier can be bypassed, and thus, the state is no longer protected by the symmetry. In previous experiments on this topic,[18] the spin bath has been provided by impurities within the topological insulator itself, so it is not subject to external control. However, since the edge states occur at the physical edge of the system, one could also consider providing an external spin bath which can be brought into diathermal contact with the outside edge of the topological insulator and subsequently removed; provided the systems are in close enough contact, then the spins should become coupled to the edge electrons by the hyperfine interaction and thus the time-reversal symmetry will be broken. Thus, by manipulating the spin bath, we should be able to switch back-scattering on and off. Of course, in reality most topological insulators will have at least a few impurities, and thus, there will still be some violations of time-reversal invariance even when the external spin bath is not present, but by purifying the insulator as much as possible and then providing a spin-rich material as an external spin bath, one could nonetheless switch between a high degree of time-reversal invariance and a significantly lower degree of time-reversal invariance. We will assume that adding or removing the spin bath can be performed adiabatically (i.e., no exchange of energy occurs during this process). In general, adiabatic transformations cannot lower the entropy of a state, and indeed, we will see that this is the case in this instance.

Similarly, it has been shown[19] that topological states can be produced by shining monochromatic circularly polarized light on a semiconductor quantum well. By tuning the form of the periodic modulation of the light, we can control the symmetries present, switching from an insulating state protected from back-scattering by a time-reversal symmetry to an unprotected state where back-scattering can occur. Thus, by controlling the symmetries we also control the spectral properties of the edge state. This method has some practical advantages over the spin bath approach, since photon polarization is easier to manage than nuclear spin. It is, however, less clear that shining light on an insulator can be achieved adiabatically, but we will assume that energy transfer is minimal, so the process is still effectively adiabatic. Based on these examples, we think it is reasonable to treat the time-reversal symmetry of systems of this kind as a parameter which is subject to external manipulations; let us now see what kinds of role symmetry-based control might play in the thermodynamics of these systems.

III. SYMMETRY IN THERMODYNAMICS

We have argued that symmetries may play a role in the thermodynamic description of a system; in this section, we will demonstrate this by exploring the relationship between symmetry and entropy. For classical cases, we will employ the Boltzmann entropy, which is given by the formula $-k_B \ln(\Theta)$, where Θ is the number of microstates which are compatible with the known macrostate of the system,[3] where a "microstate" of a system is defined by giving the position and state of each of the individual particles in the system (in most classical cases, the "state" will simply be the momentum, but other properties like spin may also be relevant), and a "macrostate" is defined by the values of the chosen macroscopic parameters, so in general, a single macrostate is associated with many possible microstates. For example, putting the same quantity of gas in a larger box will increase the entropy since there are more possible positions for the gas molecules compatible with the same macrostate. For quantum cases, we will employ the von Neumann entropy, which is given by the formula $-k_B Tr(\rho \ln(\rho))$, where ρ is the quantum state of the system. The definition uses the trace function because trace is an invariant property of a matrix of this class; note that this entails that the von Neumann entropy of a pure quantum state is always zero. Technically speaking, the von Neumann entropy is the quantum analog of the Gibbs entropy rather than the Boltzmann entropy,[15] but we employ the Boltzmann entropy here because the connection between symmetries and Boltzmann entropies is conceptually clearer and more easily illustrated. Since the Gibbs and Boltzmann entropies are equivalent in a wide variety of important cases (including whenever the systems in question are in equilibrium[20]), our conclusions are likely to generalize to the Gibbs entropy in most circumstances, although no doubt there will be a few interesting exceptions.

It will be important for us to distinguish between symmetries of states and dynamical symmetries.[21,22] We use the term "symmetry of state" to refer to the group of invertible automorphisms on the space of microstates of a system which do not change the macrostate of the system (where an "automorphism" is a structure-preserving mapping from an object to itself—in this case, from the set of microstates to the set of microstates). Symmetries of state are, therefore, instantaneous properties of an individual macrostate, and we will say that for a given system, one macrostate is more symmetric than another if the associated symmetry group is larger. We will use the term "dynamical symmetry" to refer to the group of transformations which map solutions of the equations of motion of a system into other solutions: for example, in the topological insulator without spin bath, time reversal is a symmetry of the system since a time-reversed solution is still a solution. Dynamical symmetries are thus properties of the laws by which the system evolves and not of individual states; when a system is subject to a dynamical symmetry, we know that the only possible evolutions are evolutions which obey the relevant symmetry, so for example, in the topological insulator case all allowed dynamical evolutions must be time-reversal invariant. We will now consider each of these types of symmetry in turn and discuss how they relate to classical and quantum entropy measures.

A. Symmetries of state

Because entropy is a function of state, it is natural to suppose that symmetries of state will have a direct impact on the entropy associated with a state. In particular, symmetry may in some cases be regarded as a macroscopic parameter of the system or may have a direct quantitative relationship with some macroscopic parameter, and thus, including symmetry properties in the definition of the macrostate will often have an effect on the quantitative value of the entropy.

1. The classical case

Suppose we have a set of n particles whose macrostate is fully defined by the net polarization p, which is equal to the number of spin up particles minus the number of spin down particles, i.e., $u = n/2 + p/2$ where u is the number of spin up particles. Clearly, the number of automorphisms in the symmetry group for the macrostate of polarization p is given by the number of microstates compatible with the macrostate of polarization p, which is equal to

$$\frac{n!}{(n/2 + p/2)!(n/2 - p/2)}.$$

Moreover, the Boltzmann entropy of the system at polarization p is equal to the logarithm of this number, and therefore, states with a larger symmetry group are also states of higher entropy. In particular, the "most symmetric" state is the one where the number of spin up particles is equal to the number of spin down particles, which is also the highest entropy state. Moreover, this relationship between symmetry and entropy is not merely a quirk of one particular case, because it follows directly from the combinatoric features of any problem involving assigning properties to individuals that the number of automorphisms in the symmetry group associated with a given macrostate is equal to the number of different possible microstates compatible with that macrostate, so in general, states which are more symmetric in this sense will have higher Boltzmann entropy, and the highest entropy state will be the most symmetric state.

This makes it clear that the "order" involved in symmetries of state and the order which entropy is supposed to measure the absence of are in a sense inverses of one another: the entropy measure takes highly non-uniform states to be "ordered," whereas from the point of view of the symmetry group description, highly uniform states are more ordered. The control theory paradigm gives a crisp way of understanding these different perspectives on the nature of order, because in the control theory approach, entropy is intended to the absence of the type of order which we are able to *make use of* to implement various thermodynamical transformations. In general, the usefulness of order in thermodynamics lies in the fact that it allows us to control and accumulate microscopic fluctuations in order to do useful macroscopic work, and heuristically speaking, accumulating fluctuations involves arranging that the fluctuations should all "point in the same direction" so their effects sum rather than canceling, which in general requires *decreasing* the size of the symmetry automorphism group. For example, in the simple example above, we can use the spins in the system to achieve macroscopic effects only if a significant number of spins are pointing in the same direction, so that a net macroscopic polarization is produced. Similarly, in the example of the topological insulator, we achieve a net macroscopic current at the edge only if the number of left-movers is significantly greater than the number of right-movers or vice versa. So from a thermodynamic point of view, it is *asymmetry* which is generally of practical value to us: the more balanced a system is, the more difficult to extract work from the energy of its microscopic fluctuations.

Moreover, the observation that states with a large symmetry group have high entropy provides an interesting link between Curie's principle and the second law of thermodynamics. Curie's principle states that an asymmetry in an effect must be produced by an asymmetry in the cause.[23] This principle remains the subject of debate, and there are a few putative counterexamples,[24] but nonetheless, it is a good heuristic which is true in many circumstances. In fact, the content of Curie's principle is quite similar to the second law of thermodynamics: it tells us that asymmetry cannot increase over time, or correspondingly, symmetry cannot decrease over time, which is exactly what we would expect based on the observation that, in general, states with higher entropy are also more symmetric. Thus, as noted by Rosen,[25] for certain sorts of symmetries (i.e., those which are associated with a group of automorphisms on microstates) the fact that symmetry is in general non-decreasing can be understood as a consequence of the fact that entropy is in general non-decreasing, and thus, the relationship between entropy and symmetry provides a quantitative justification for Curie's principle for these cases.[39]

That said, it is also important to observe that there is an alternative point of view in which symmetries *decrease* entropy. Consider another type of case where we take a system and impose some symmetry as an external constraint on that system. In general, this will have the effect of *decreasing* the Boltzmann entropy, since the set of states obeying the symmetry is a proper subset of the whole set of microstates S. For example, in the case of our spin system of n particles, if we originally have no information about the value of p and we then apply an external constraint which resets the polarization to zero, the Boltzmann entropy will actually decrease, since the number of microstates compatible with the more constrained macrostate is smaller. So in fact one and the same symmetry of state can have the effect of increasing or decreasing entropy, depending on whether we are (1) moving from a specific set of non-symmetric microstates to a disjoint set of symmetric microstates, in which case entropy will typically increase, or (2) we are moving from an entire state space to a proper subset of symmetric microstates, in which case entropy will typically decrease. However, it should be reinforced that in the latter case the entropy decreases not because of any specific properties of the symmetry, but simply because we have imposed some constraint which narrows down the state space—moving from an entire state space to a particular subset of non-symmetric microstates will have the same effect.

2. The quantum case

The von Neumann entropy and the Boltzmann entropy are equivalent in many circumstances,[26] and thus the conclusions we have drawn above also apply to many quantum systems. For example, cases where we have classical uncertainty are typically translated to *mixed* quantum states, with classical probabilities transformed into coefficients in a convex decomposition: so, for example, for our n-particle spin system with polarization p, the appropriate quantum analog might become an equally weighted sum of all of the $n!/(n/2 + p/2)!(n/2 - p/2)$ pure states in which exactly $n/2 + p/2$ spins are up and the remainder are down. The von Neumann entropy of this arrangement would then be equal to the Boltzmann entropy of the original classical system, and thus, we will see the same relationships between entropy and symmetry as we saw in the classical case.

However, this representation assumes that the n particles are all distinguishable (for example, because they are labeled by their position in some sort of lattice). An alternative quantum representation of this case would regard the particles as indistinguishable, which means that rather than having $n!/(n/2+p/2)!(n/2-p/2)$ distinct pure states, we would be working with a single pure state which consists of $n/2+p/2$ up spins and $n/2-p/2$ down spins, with nothing further to be said about which particular spins are up or down. In this setting, each polarization is associated with only a single possible (pure) microstate, and therefore, all the polarization macrostates have entropy equal to zero, so symmetry or asymmetry is irrelevant. Thus, the conclusions we have drawn about the relationship between entropy and symmetry still hold in quantum thermodynamics, but only insofar as the symmetry or asymmetry properties are understood as constraints on *classical* uncertainty, rather than constraints on a pure quantum state. Put simply, the reason symmetries of state are associated with high entropy is that they typically engender high uncertainty about which particles have which properties, but when we are dealing with a pure quantum state, there are in general no objective facts about which particles have which properties, and this severs the link between symmetry and uncertainty.

B. Dynamical symmetries

Let us now consider the effect of dynamical symmetries on entropic descriptions. Since entropy is a function of state which is blind to dynamics, one might think that dynamical symmetries would be irrelevant to the entropy of a system. However, there is a relationship between symmetries of state and dynamical symmetries, because dynamical symmetries are frequently linked with variational symmetries, which are groups of infinitesimal transformations of variables under which the action of the system is invariant.[40] From Noether's first theorem, if a system obeys a continuous global variational symmetry, then it must have a conserved quantity,[27] and this conserved quantity may sometimes function to preserve a symmetry of state. For example, symmetry under translations leads to the conservation of momentum, so if a state exhibits some symmetry with regard to momentum (e.g., the total momentum of system A is equal and opposite to the total momentum of system B, which is not interacting with A), the dynamical symmetry works to ensure that the relevant symmetry of state continues to hold.

In the case of the topological insulator, there is no conserved Noether quantity corresponding to the time reversal symmetry, because time reversal symmetry is not a continuous symmetry (there are only two possibilities for the direction of time, so there is no continuous parametrization) and therefore Noether's theorem does not apply. However, it is nonetheless true that in this case the dynamical symmetry is related to a symmetry of state, because the dynamical symmetry prevents back-scattering and thus if we start out with the number of left-movers equal to the number of right-movers, these numbers will continue to be equal throughout the dynamical evolution. We will now examine these effects in greater detail for the classical and quantum case.

1. The classical case

We have seen that, in general, the effect of imposing a dynamical symmetry on a system is to conserve some features of the original state by limiting the possible evolutions that the system can potentially undergo so that its microstate is limited to some region R of state space. Thus, removing a symmetry will allow access to a region R' of state space which was previously forbidden. The effect of this process on the long-run entropy will depend on the nature of the regions in question: for example, if a system is prepared in such a way that it exhibits some symmetry of state, and it is subject to some dynamical symmetry which conserves that symmetry of state, removing the dynamical symmetry may allow the system to reach less symmetrical states which have lower entropy. Thus, in this instance removing the dynamical symmetry may be the precursor to a *decrease* in entropy. For example, in the case of the topological insulator, the system is able to move from the higher-entropy symmetrical states to the lower-entropy states with the number of left-movers not equal to the number of right-movers. (Of course in this scenario it is also necessary to input some energy to effect this transformation, and this is likely to be the case in similar examples, since transitions to lower entropy states do not typically happen spontaneously.)

On the other hand, if a system is prepared in a low entropy state, the dynamical symmetry may instead function to preserve the system in the low entropy state until such a time as the dynamical symmetry is removed and the system can evolve toward the higher entropy state. This effect could also be achieved with the topological insulator if it were possible to remove the spin bath and thus restore the time-symmetry after some back-scattering has occurred: the low entropy state with the number of left-movers not equal to the number of right-movers would in principle be preserved until time-symmetry could be broken again.

Note that the preceding discussion assumed that the region of state space R' made accessible by removing the dynamical symmetry is macroscopically distinct from the region R in which the state is constrained to lie for as long as the symmetry holds, so these two regions of state space are associated with distinct macrostates. In the topological insulator, the states in the region R' made accessible by back-scattering have nonzero edge current and polarization, so they are macroscopically distinct from states in region R which all have zero edge current and polarization. However, it was noted in Sec. I that, in general, we have some freedom of choice about which macroscopic parameters are considered relevant, so we can also consider a similar case where R' is regarded as having the same values of the macroscopic parameters as R. For example, in the topological insulator case we could declare that polarization and edge current should be left out of the macroscopic description of the state (e.g., because in our lab we have no instruments which can measure them). Thus, in this case, removing the dynamical symmetry and then allowing the system to undergo free evolution has the effect of increasing the number of microstates compatible with the known macrostate, so the effect of breaking the symmetry is always to *increase* the entropy of the system, regardless of the nature of the states in R and R'.

It is important to reinforce that dynamical symmetries forbid *transitions*, not states in-and-of themselves, and thus in the case we have just considered, a dynamical symmetry can forbid the state from moving from R to R', but the symmetry in-and-of-itself does not rule out any of the states in R and R'. So if the information available to us about the system consists only of the values of the relevant macroscopic parameters and the fact that the relevant dynamical symmetry currently applies, all of the states in R and R' are compatible with the known macrostate, and the entropy will not increase when we remove the symmetry (removing the symmetry widens the range of possible evolutions but not the number of microstates compatible with our information). This demonstrates that the increase in entropy induced by removing a symmetry is in part a function of our knowledge of the history of a system; in the topological insulator case we are confident, based on theoretical arguments, that the original state has equal numbers of left-movers and right-movers, even though we have not given ourselves access to macroscopic parameters which could confirm that, so we are able to say that removing the symmetry strictly increases the number of microstates compatible with our information about the system, even though the new microstates still have the same values of the macroscopic parameters.

It is also natural to ask whether adding the symmetry back in *decreases* the Boltzmann entropy. This would be surprising, because typically we expect that adding and removing a symmetry can be achieved adiabatically (e.g., adding and removing a spin bath is an adiabatic process), and adiabatic transformations cannot decrease entropy in classical thermodynamics. However, as a matter of fact removing the symmetry does not change the Boltzmann entropy. This is because prior to removing the symmetry, the system could be anywhere in Ror R', and after removing the symmetry, it could also be in either R or R', under the assumption that we have no access to macroscopic parameters which would distinguish these two regions. So the entropy actually remains the same, because all of the same states remain compatible with the macrostate; we are just in the slightly unusual situation where we know that the system is confined by the symmetry in some region of the state space which is smaller than the region which is compatible with the current macrostate. The reason the entropy increases when we remove the symmetry but simply stays the same when we put the symmetry back in is not really to do with the dynamical or thermodynamical properties of the system at all, but is simply a function of our knowledge about the history of the system. In fact it is often argued that this sort of effect is to some extent responsible for the increase in entropy over time:[28] in general, we know the most about the state of a system when we have deliberately prepared it in a particular state, and when we subsequently allow it to evolve, we lose information as it may evolve away from the region we initially selected. Removing the dynamical symmetry is in this sense precisely analogous to removing the partition in the textbook example where a gas is confined in one half of a box and we then remove the partition: as we allow the gas to disperse through the box, we lose information about its state and thus the entropy increases.

This is also an interesting example of a case where the Boltzmann entropy fails to fully capture the information that we have about the microstate of a system, because it does not distinguish between the case where the system can evolve freely between the microstates compatible with the current macrostate and the case where we can partition this set of microstates into subsets such that we know the system is confined to one of these subsets but we cannot find out which subset it is confined to. Clearly, there is a sense in which we have much more information in the latter case than the former, since we can rule out all the trajectories for the microstate which do not remain within a single subset throughout the time of observation; thus, the possible microscopic evolutions of the system are much more limited. However, the Boltzmann entropy is not capable of quantifying the extra information we have in the latter case, as it characterizes only our (lack of) knowledge of the instantaneous microstate of the system, and not our (lack of) knowledge of trajectories though the state space. This is an important reminder that the entropic description does not necessarily exhaust all the information that we have about the system—it is simply intended to capture all the information that is useful from the point of view of thermodynamics, i.e., information that we can use to control the system in useful ways. A similar effect can be achieved in the case of the gas in a box by removing the partition, allowing the system to equilibrate, and then closing the partition again: now we know that the ratio of the volume of gas on one side to the volume of gas on the other side must remain constant over time, but we have no new information about what that ratio is, and thus the insertion of the partition does not change the entropy of the system even though in a sense we do have extra information. The gas in the box case has been well explored in the literature[29,30] (it was first formulated in the 1870s[31]), and thus, we can say with some confidence that the extra information available in this case does not allow the construction of a Maxwell's demon or other violations of the law of thermodynamics, and presumably the same is true in the case of the dynamical symmetry, though it might nonetheless be interesting to consider if this special effect could be leveraged for some new technological applications.

2. The quantum case

Of course, the topological insulator is in fact a quantum system, so the preceding analysis should really be performed in terms of the von Neumann entropy. Much of the preceding analysis will also carry over directly to the quantum case. The von Neumann entropy is a function only of the instantaneous state, so it is clear that changes in dynamical symmetries will not instantaneously alter the von Neumann entropy—for example, simply adding a spin bath to our topological insulator does not instantaneously change the quantum state (and in general no significant changes will occur until we add a charging current), and thus, the von Neumann entropy remains the same. However, by changing the set of allowed evolutions, we do change the set of states which can be reached, and thus when the state space is extended over spacetime, changes in dynamical symmetries will have an effect on the von Neumann entropy if some of the states that have become accessible have different von Neumann entropies from the original state. As before, if we are dealing with a case where classical uncertainty over different possible microstates is translated into a convex decomposition over pure states, then the von Neumann entropies will be equal to the Boltzmann entropies and thus the same conclusions will apply.

On the other hand, if classical uncertainty is translated to a single pure state, we will get different results. In the topological insulator, we typically regard the edge particles as indistinguishable, so states can be labeled by the difference in the number of left and right movers, but there is nothing further to be said about which particles in particular are moving left or right. When the time symmetry is removed the edge states become coupled with the spins in the spin bath and thus although the state of the system as a whole remains pure, the reduced state of the edge particles becomes more mixed and therefore the entropy of the edge particles actually increases in this process, even though the state is becoming "less symmetric." This is again a consequence of the fact that the indistinguishability of particles in pure quantum states severs the link between symmetry and uncertainty—since there is no fact of the matter about which edge particles are going which way in the original edge state, the high degree of symmetry of this state does not lead to high uncertainty, and thus, the symmetry of this state makes no contribution to its entropy.

C. Symmetries as a means of control

The preceding discussion makes it clear that neither symmetries of state nor dynamical symmetries are equivalent to order as measured by either the Boltzmann or von Neumann negentropy. Indeed, symmetries of state are in some sense exactly the opposite of the order measured by negentropy, and thus, highly symmetrical states are often undesirable in thermodynamics as they make it more difficult to accumulate microscopic fluctuations. However, we have seen that manipulating symmetries of state and/or dynamical symmetries can be used as a way to either decrease or increase entropy, depending on the circumstances and the set of relevant macroscopic parameters. This leads us to an important moral: symmetries cannot typically be regarded as a thermodynamical property akin to temperature, pressure, heat capacity and so on. Rather the most perspicacious way to understand the role of symmetries in thermodynamics is to see them as a controllable feature within the control-theory paradigm. Symmetries—and in particular *dynamical symmetries*—are relevant to a thermodynamical description insofar as they can be used to preserve and/or control the features that *are* relevant to thermodynamical descriptions, like the other sorts of control operations that Wallace describes in his presentation of the control theory approach to thermodynamics.[4] For example, in a topological insulator, ensuring that there is no spin bath present entails that the system will remain in the high-entropy state with $n = 0$ and will not gain an edge current or polarization even if we attempt to charge it up.

As with the other "control operations" employed in control theory approaches to thermodynamics, dynamical symmetries are a property of the external constraints on a system, not of the system itself, since changing the symmetries of a system does not instantaneously change the state. Indeed, it is usually reasonable to regard dynamical symmetries or the lack thereof as exogenous variables. In the topological insulator case, if the size of the spin bath is effectively infinite relative to the number of electrons in the edge state, then the time-reversal symmetry can be regarded as an exogenous variable: adding the spin bath has an effect on the edge state (by permitting back-scattering) but we do not need to consider the effect of the edge state's behavior on the spin bath. However, if the spin bath is finite some care is required because each instance of backscattering uses up one of the spins, and therefore, after a certain amount of back-scattering we will run out of spins and thus back-scattering will again be forbidden. So in the finite case, it is not simply a matter of switching the symmetry on and off: rather we exert control by providing a finite resource which can be used up. This is in fact directly analogous to the standard thermodynamical use of heat baths: when we provide a heat bath to change the temperature of the system being studied, we typically assume that the bath is effectively infinite so its temperature remains the same while it provides heat to the system or absorbs heat from the system; but of course in real life heat baths are finite and their temperature will eventually change as they provide or absorb heat. The difference is that in the case of the topological insulator, only the spins adjacent to the edge can affect the symmetry, so there is in fact a strong limit on the size of the bath, making the idealization of infinite size inaccurate in certain circumstances.

We have seen that the precise relationship between symmetry and entropy depends on a variety of factors, including the type of symmetry, the type of entropy, and crucially, which regions of state space are regarded as macroscopically distinct and what we know about the history of the system. This is to be expected, because we have argued here that both entropy and symmetry have epistemic aspects, since they are features not only of systems in and of themselves, but also of our *knowledge* of those systems—for example, dynamical symmetries serve to delineate a set of possible transitions, but of course a system in a fixed microstate at a fixed time will actually undergo exactly one transition, so the relevance of the *set* is in large part to characterize what observers who do not know the exact microstate can infer about the evolution of the system. Thus, we should not expect to be able to write down a straightforward mathematical relationship between the symmetry increasing/decreasing aspects of a manipulation and its thermodynamic properties: such a relation would also have to take account all the various possibilities for an observer's knowledge of a system and its history, and would therefore become a very complex three-way relationship between the system, the external constraints, and the knowledge of the observer, but the fact that the relation between symmetry and thermodynamics is complex does not undermine our observation that the thermodynamic relevance of symmetry is primarily as a means of control. This observation is particularly important given growing interest in the control theory paradigm of thermodynamics—symmetries are not mentioned as a possible means of control in any of the articles that we have cited on the subject, and yet we have seen here that they offer interesting possibilities in this regard, so we consider that symmetries may be a promising area for future research in this field.

Moreover, we see great potential for novel technological applications once it is recognized that symmetries provide us with new ways of exerting control over thermodynamical systems. For example, symmetry-protected states are a particularly useful way of making use of the order inherent in quantum entanglement, because they provide a convenient middle ground between simple entangled states and topologically ordered states. Ordinary entangled states like two-particle Bell states are very unstable because they can be disrupted by arbitrary local perturbations, and thus they are not sufficiently reliable for use in thermodynamical applications. By contrast, topologically ordered states[32] (which involve a stronger form of order than symmetry-protected topological states) are *too* stable for many thermodynamical applications—the topological phase cannot be changed by any local perturbation whatsoever, which makes it difficult to access and make use of the order inherent in the system in any thermodynamical context, but with symmetry-protected topological states, we can change the phase with a local perturbation if the relevant symmetry is turned off, so we can access and make use of the system's order by controlling the symmetries of the system, but nonetheless, the state can be held in a meta-stable state for long periods of time when the symmetry remains switched on. Moreover, because symmetry-protected topological states require only short-range entanglement, they are easier to produce than full topologically ordered states, and therefore, technology based on such states is much more accessible today.

Some steps have already been taken toward implementing symmetry-based control of thermodynamic systems. For example, Ref. 33 has experimentally demonstrated the decrease in entropy associated with the removal of a symmetry by measuring the heat dissipated when a Brownian particle transitions from a single-well to a double-well potential; this article focuses on spontaneous symmetry breaking, but in fact in practice the manipulation of an optical trap is employed to bring the symmetry-breaking under the control of the experimenters, thus providing a clear good demonstration of the possibilities for symmetry-based control within the control-theory paradigm. Meanwhile, Ref. 34 has shown theoretically that controlling the symmetries of thermoelectric materials can make it more efficient. In particular, they introduce an "asymmetry parameter" quantifying the degree of symmetry-breaking, and demonstrate that as the asymmetry parameter becomes it is possible to overcome the Curzon–Ahlborn limit within linear response and to reach the Carnot efficiency even if the usual figure of merit for thermoelectric materials is small. These projects provide a proof-of-principle that symmetry-based control can have practically useful applications, but there is a much wider space of possibilities yet to be explored, and thus, we anticipate that symmetry-protected states will be a fruitful domain for future investigations into new technologies employing quantum thermodynamics.

IV. APPLICATION: COHERENCE CAPACITOR USING SYMMETRY-PROTECTED STATES

In this section, we describe a concrete application of symmetry-based control using symmetry-protected topological states—specifically, we show how to use symmetries to construct a coherence capacitor (CC). Acting in concert with a quantum heat engine (QHE), a coherence capacitor is a source of power (such as electrical power on a macroscopic level or kinetic power on a microscopic level) whose efficiency exceeds classical thermodynamic limits and whose storage capacity can greatly exceed electrochemical or mechanical batteries of identical mass. It is also a safe source of power, being immune to unintentional energy releases such as frequently occur in chemical storage systems such as lithium-ion batteries.

To understand the significance of coherence capacitor concepts, we first start with an abstraction revealing the importance of superlinear scaling and then discuss how "low -quality" energy from the environment is converted to work using the information[35] in the coherence capacitor as a resource.[6] Finally, we discuss a practical implementations of the coherence capacitor which makes use of symmetry as a means of control.

A. Information content of a graph

A coherence capacitor in its ideally abstracted form can be conceived as a graph, a set of nodes connected by edges. The nodes of the graph represent physical entities, such as atoms, molecules, or qubits of some sort; the edges represent our resource, connections between these entities, such as a quantum correlation only possible in a coherent system. The nodes, therefore, are massive; the edges are massless but contain information (the absence or presence of an edge between two nodes is analogous to a bit in a state 0 or 1).

The number of edges in a fully connected or complete graph of n nodes is

$$\frac{n(n-1)}{2},$$

which scales as n^2 in the limit of large n. This *superlinear* scaling in number of edges represents a similarly superlinear scaling in the information content of the graph, so that the specific information (information per unit mass) contained in the graph is not constant but increases with the graph size.

It is nearly certain that in practical graph implementations, distance or shielding effects will disallow to some extent the ability of a sufficiently large graph to be fully connected; but even on graphs where such a limit is active, it is likely that the scaling will be superlinear in some fashion (n^β, where $\beta > 1$). As will be shown in Sec. IV B, it is superlinear scaling of information content that makes a coherence capacitor interesting as a quantum thermodynamic system.

B. The coherence capacitor as a source of work

The combination of a store of information or negentropy (the coherence capacitor) and an engine capable of converting that information to useful work (the Quantum Heat Engine) represents a powerful method of harvesting ambient energy that is inaccessible to conventional heat engines. Whereas classically the "quality" of the heat processed by an engine depends on the ratio of the temperatures of the "cold" and "hot" reservoirs (giving rise to the Carnot efficiency $\eta_c = 1 - T_c/T_h$), the CC+QHE combination can avoid this metric altogether by introducing negentropy as a consumable resource, rendering all heat of equal quality. This allows the QHE to exceed the Carnot efficiency when also consuming information;[36] in the limiting case, the CC+QHE combination can generate work from a single temperature reservoir.

The work generated by a QHE in excess of that allowed by the Carnot limit is work undiscoverable classically and arises as a result of the information stored in the CC. To distinguish this additional quantity of work, and to allow a statement of the utility of the CC, we call it "quantum-accessible" work; that is to say, it is work that is only accessible through a proper quantum-mechanical mechanism. In the limit of a CC+QHE operating on a single heat bath, the entire work output of the quantum thermodynamic process consists of quantum-accessible work extracted from the maximal-entropy heat of the bath.

The quantity of quantum-accessible work associated with a quantity of information is given by Landauer's principle,[37]

$$L = k_b T \ln 2,$$

in which T is the temperature of the reservoir, k_b is Boltzmann's constant, and L is the quantum-accessible work associated with a single bit of information. Greater work can, therefore, be extracted when the QHE is operating in a higher-temperature environment, as should be expected; and the quantum-accessible work accessible by the information in a CC is not constant but is environment-dependent. (This contrasts with other methods of energy storage and retrieval, such as electrochemical batteries, which have environment-independent output characteristics—within operational limits.) Also, with constant quantum-accessible work per bit of information ℓ, the benefit of the superlinear scaling of Sec. IV A is clear: the information content and the quantum-accessible work both increase more quickly than the mass of the physical realization of the CC. This also contrasts positively with electrochemical energy storage devices, which have a constant energy per unit mass ϵ and therefore which increase in energy content proportionately to an increase in mass.

C. Physical realizations and their scaling

There are two salient features of a CC+QHE system that should be realized in practice in order to make a practical power source. The first, and *sine qua non*, is the ability to extract work from the environment which beats the efficiency performance of classical counterparts by utilizing information as a resource. This means that these systems should appear to exceed the Carnot efficiency as traditionally defined using equilibrium temperatures. It is possible to carefully define thermodynamic empirical temperatures for the working fluid so as to come up with a generalized quantum Carnot bound which will not be exceeded by the quantum device, but the important point is that the quantum device should exceed the *classical* Carnot bound as it would traditionally be defined for a classical counterpart. The second condition, highly desirable but not necessary, is that the scaling of the extractable quantum-accessible work be superlinear with increasing CC mass. The first condition ensures that the CC has technological benefit to offer that is unique when compared with electrochemical batteries; the second allows the CC+QHE to perhaps outclass electrochemistry as the size of the system increases.

An example realization of the first criterion is given by Ref. 37, in which the CC and QHE are integrated into a single device. This device uses a topological insulator as described in Sec. II together with a spin bath which breaks the time-reversal invariance, as depicted in Fig. 1. As an example, suppose the electrons moving right on the top edge of the topological insulator are spin down and the electrons on the top edge moving left are spin up. If there is no spin-bath present, then time-reversal symmetry prevents back-scattering, but by including a spin bath which removes this dynamical symmetry, we increase the accessible region of phase space and so it is now possible for back-scattering to occur. On the top edge, an electron scattering from right to left will produce a spin down nucleus, while scattering from left to right will produce a spin up nucleus. Suppose the current is flowing to the left; then there will be more right-moving electrons than left-moving electrons, so we will get more back-scattering from right to left, and therefore, the material near the top edge will end up with more down than up nuclear spins and will thus take an overall nuclear polarization. Thus, the presence or absence of the dynamical time-reversal symmetry produces observable changes in the macroscopic polarization parameter. Heat is dissipated during this process due to the back-scattering of the electrons. Note that all of this can be achieved using nuclear spins which are degenerate, meaning that up and down spins have the same energy, and therefore, the down-polarized state has the same overall energy as the unpolarized state, and thus, there is no energy stored in the nuclear spins—this is a non-energetic memory. However, the process does decrease the classical Boltzmann entropy of the system, since we move to a non-zero n state, and this decrease in entropy is the reason why we need to dissipate heat during the polarization process, so that the increase in entropy of the environment will balance the decrease in entropy of the system.

When we remove the applied current, the imbalance in left and right movers must be preserved, since it is recorded in the polarization of the nuclear spins. Therefore, we now have more left-moving than right-moving electrons, so a current flows in the opposite direction from the charging current. As the current flows the nuclear spins flip in the opposite direction and the polarization dissipates, so eventually the number of left-movers and right-movers will become balanced again and there will no longer be any net current. If we connect a non-zero bias to the circuit while the current is still flowing, the current will extract energy from the surrounding thermal reservoirs in order to do work. The device memory is thus used to convert heat directly to work. However, this does not violate the second law of thermodynamics, as the entropy of the device increases as the polarization is lost and thus this device cannot be operated continuously in a cycle without going through the polarization phase again.

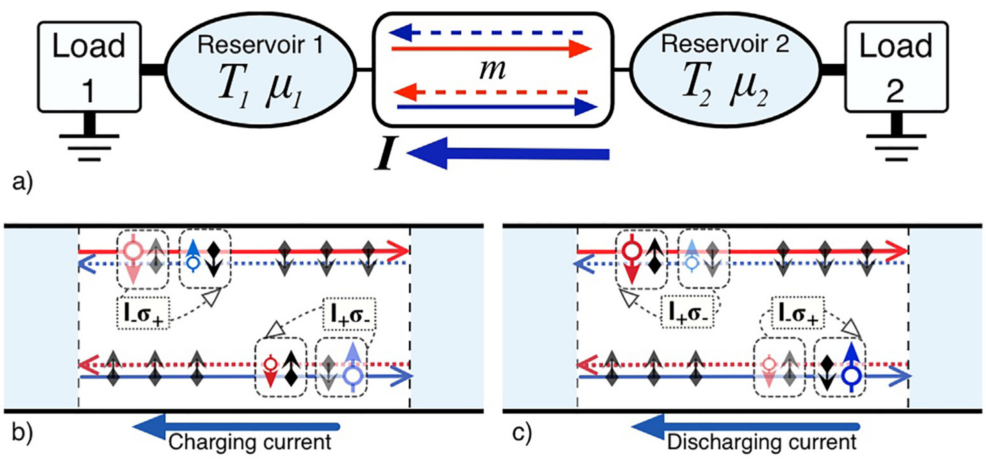

Fig. 1. Schematic diagram of (a) the CC+QHE system using a topological insulator, (b) the charging phase, and (c) the discharging phase. In the charging phase, an applied bias current increases the number of right-movers (solid lines) at the edges relative to left-movers. This excess creates a net nuclear spin polarization with opposite values in each edge. In the discharging phase, even without external bias, the net polarization of the nuclear spins increases the number of right-movers, driving a net discharging current to the left. Reprinted with permission from Bozkurt *et al.*, Phys. Rev. B **97**, 245414 (2018). Copyright 2018 American Physical Society.[39]

Note that the storage capability of this device depends on the strength of the hyperfine interactions that permit the symmetry breaking. If the interaction is strong, the nuclear polarization will be produced quickly during charging, but will also dissipate quickly, so it will not be possible to use the device for storage (but it might be suitable as a capacitor). If the interaction is weak, it will take longer to produce a nuclear polarization but the polarization will not dissipate quickly, so we will be able to use the device as a kind of "battery." In the limit where the hyperfine interaction goes to zero, in principle the polarization cannot dissipate at all and thus, the imbalanced edge state will be preserved. Different materials exhibit different strength of hyperfine interaction—for example, thin film flakes of 3D topological QSHI Bi2Te2Se (BTS221) have a fairly weak hyperfine interaction and thus implementations of this kind require a very strong charging current but have been experimentally demonstrated to retain a polarization for several days, whereas InAs/GaSB quantum wells have a much stronger hyperfine interaction, and therefore, implementations of this kind are easy to charge but also depolarize quickly.[18] Moreover, it is possible to change the concentration of nuclear spins in InAs/GaSB quantum wells using magnetic impurity dosing, which offers further possibilities for controlling the level of symmetry breaking. Thus, by making appropriate choices for the material of the spin bath, we are able to exert thermodynamical control over the rate of charging and discharging of the device, so the choice of material acts as a macroscopic parameter which alters the values of thermodynamic quantities like entropy.

This observation suggests possibilities for a modified version of the topological coherence capacitor where symmetry is actively manipulated in order to control the system's behavior. Let us suppose that it is possible to add or remove an external spin bath; in principle this would allow us to turn time-reversal invariance symmetry on and off, thus changing the rate of charging and discharging. Note that when the system is unpolarized, adding the spin bath is not by itself enough to prompt the system to evolve to a polarized state: this transition decreases the entropy of the system and thus we must still provide an external current to induce it. A more interesting case occurs if the symmetry barrier can be removed while the system is in the higher entropy state, because then the removal of the barrier would be expected to lead to spontaneous evolution into the lower entropy state without any further external manipulation. For example, suppose we add the spin bath and create a polarization, then remove the spin bath to preserve the polarization, and then at some later time we add the spin bath again, at which point the system is expected to spontaneously depolarize (and produce a current if there is a load attached). Assuming that this is technologically possible, it would be an interesting example of a case where we can induce a spontaneous evolution simply by (adiabatically) altering a symmetry of the system and making no other changes.

In a weak coupling and short-edge limit, the quantum-accessible work available from the topological coherence capacitor scales as N^2, where N is the number of nuclear spins; this is the ideal scaling described earlier for a fully connected graph. However, there is no identifiable graph-like structure to this implementation, which suggests that the N^2 scaling is coincidental and may disappear outside the limits under which it was derived. In particular, there would seem to be a significant likelihood that in the thermodynamic (large N) limit, the scaling will asymptote to a linear behavior. The open question that remains, therefore, is whether the superlinear scaling can be accessed in devices large enough to be useful in practical applications, such as nanoelectronics. The most likely reason for the failure of the topological coherence capacitor to manifest superlinear scaling at macroscopic scales is that the coherences involved are spin coherences rather than quantum coherences, and the spin coherences are not long-range interactions and are therefore unlikely to correspond to edges in the fully connected graph analogy. A theoretical manifestation of the CC+QHE that clearly manifests the N^2 scaling law is based on heat exchange coherences (HEC) in a system of qubits.[38] The HECs, joined by displacement and squeezing coherences, scale quickly with increasing numbers of qubits in the system and mimic the complete graph model. The extraction of work using these HECs is proposed to be accomplished through super-radiance in superconducting circuit QED schemes. While this approach is promising—particularly because it directly addresses the desirable superlinear scaling feature—to date no experiment has been performed to demonstrate that this theory is practicable.

V. CONCLUSION

In this article, we have explored the rich and interesting relationship between symmetry and entropy. We noted that in many cases, highly symmetric states also have high entropy. This illustrates some of the limitations of regarding entropy as a measure of "disorder," because a state that looks ordered from the symmetry point of view may also be highly "disordered" from the entropy point of view; the control theory paradigm provides a better way to conceptualize this relationship in terms of order, which can actually be used to produce macroscopic work. The relationship between entropy and symmetries of state then allows us to provide a quantitative justification for some instances of Curie's principle: asymmetry typically does not increase because asymmetry is the inverse of entropy and entropy does not decrease. However, we also noted that when symmetries are treated as constraints on a state-space they may decrease rather than increase entropy, so much hangs on the specific way in which symmetries are being employed. We also saw that these symmetry-related effects provide an elegant demonstration of the differences between the probabilities of statistical mechanics and the probabilities of quantum mechanics; for the von Neumann entropy of a pure quantum state is completely independent of the symmetry properties of the state, since the probabilities associated with that state are not to be interpreted as probability distributions over genuinely different microstates.

In the case of dynamical symmetries, we saw that the effect of removing a dynamical symmetry may be to allow transitions that were previously forbidden, and this may either increase or decrease entropy depending on the nature of the newly accessible states. We also observed that in some cases, a purely entropic description may fail to fully characterize our knowledge of a system's microstate, since it is possible to construct examples where the Boltzmann entropy is blind to our knowledge about the ways in which dynamical evolution is constrained by symmetries.

Finally, we argued that the thermodynamic role of symmetries is best understood in the context of the control theory approach to thermodynamics, and we explored how symmetry-based thermodynamic control can be used to construct useful energy storage and harvesting systems, such as a "negentropy battery" or coherence capacitor. The coherence capacitor is inspired by Maxwell's original vision, using quantum thermodynamics and symmetry-based effects as to get around some of the restrictions of the second law while still living within the law's constraints. We proposed a possible extension of the coherence capacitor concept that would make more active use of controllable symmetries; we believe symmetry-based control of this nature may provide interesting possibilities for further advances in technologies based on quantum thermodynamics.

ACKNOWLEDGMENTS

We would like to acknowledge significant discussions and the critical role in developing these ideas played by the staff, advisors, and research contractors of the Chief Scientist's Office at Lockheed Martin's Corporate Engineering and Technology Office, especially Dani Couger and Anthony W. Jacomb-Hood (Lockheed Martin), Joan Vaccaro and Erik Streed (Griffith University), InancAdagideli, Ahmet Bozkurt, and BarisPekerten (Sabanci University), Özgür Müstecaplioğlu (Koç University), Alexander Brinkman (University of Twente), Seth Lloyd (MIT), Yong Chen (Perdue University), and Lan Wang (RMIT).

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

DATA AVAILABILITY

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

REFERENCES

[1]H. Callen, Found. Phys. **4**(4), 423 (1974).
[2]H. E. Stanley, *Introduction to Phase Transitions and Critical Phenomena*, International Series of Monographs on Physics (Oxford University Press, Oxford, 1971).
[3]E. T. Jaynes, Am. J. Phys. **33**(5), 391–398 (1965).
[4]D. Wallace, "Thermodynamics as a Control Theory" Entropy **16**, 699–725 (2013); available at https://philpapers.org/rec/WALTAC-6.
[5]W. C. Myrvold, Found. Phys. **50**(10), 1219–1251 (2020).
[6]E. Chitambar and G. Gour, Rev. Mod. Phys. **91**(2), 025001 (2019).
[7]N. Huei, Y. Ng, and M. P. Woods, "Resource theory of quantum thermodynamics: Thermal operations and second laws," in *Thermodynamics in the Quantum Regime* (Springer International Publishing, 2018), pp. 625–650.
[8]M. Lostaglio, Rep. Prog. Phys. **82**(11), 114001 (2019).
[9]F. Brandão, M. Horodecki, N. Ng, J. Oppenheim, and S. Wehner, Proc. Natl. Acad. Sci. **112**(11), 3275–3279 (2015).
[10]G. E. Crooks, Phys. Rev. Lett. **99**(10), 100602 (2007).
[11]P. Salamon and R. S. Berry, Phys. Rev. Lett. **51**, 1127 (1983).
[12]D. A. Sivak and G. E. Crooks, Phys. Rev. Lett. **108**, 190602 (2012).
[13]X. Chen, Z.-C. Gu, Z.-X. Liu, and X.-G. Wen, Phys. Rev. B **87**, 155114 (2013).
[14]K. Wierschem and P. Sengupta, Mod. Phys. Lett. B **28**(32), 1430017 (2014).
[15]J. K. Asbóth, L. Oroszlány, and A. Pályi, *A Short Course on Topological Insulators*, Lecture Notes in Physics (Springer International Publishing, 2016).
[16]P. Walker and J. Carroll, Nucl. Phys. News **17**, 11–15 (2007).
[17]M. König, S. Wiedmann, C. Brüne, A. Roth, H. Buhmann, L. W. Molenkamp, X.-L. Qi, and S.-C. Zhang, Science **318**(5851), 766 (2007).
[18]A. M. Bozkurt, B. Pekerten, and Î. Adagideli, Phys. Rev. B **97**(24), 245414 (2018).
[19]N. H. Lindner, G. Refael, and V. Galitski, Nat. Phys. **7**(6), 490 (2011).
[20]S. Goldstein, J. L. Lebowitz, R. Tumulka, and N. Zanghì, "Gibbs and Boltzmann entropy in classical and quantum mechanics," in *Statistical Mechanics and Scientific Explanation: Determinism, Indeterminism and Laws of Nature* (World Scientific, 2020), Chap. 14, pp. 519–581.
[21]S. De Haro and J. Butterfield, "On symmetry and duality," arXiv:1905.05966 (2019).
[22]B. Dóra, J. Cayssol, F. Simon, and R. Moessner, Phys. Rev. Lett. **108**(5), 056602 (2012).
[23]E. Castellani and J. Ismael, Philos. Sci. **83**(5), 1002 (2016).
[24]B. W. Roberts, "The simple failure of Curie's principle: How to get out what hasn't gone in," LSE Philosophy Blog (2014); available at https://www.lse.ac.uk/philosophy/blog/2014/12/22/the-simple-failure-of-curies-principle-how-to-get-out-what-hasnt-gone-in/.
[25]J. Rosen, Entropy **7**(4), 308 (2005).
[26]D. Wallace, *The Necessity of Gibbsian Statistical Mechanics* (American Journal of Science, 1878); available at URL: http://philsci-archive.pitt.edu/15290/.
[27]H. R. Brown and P. Holland, Mol. Phys.**102**(11–12),1133 (2004).
[28]T. L. Duncan and J. S. Semura, Found. Phys. **37**(12), 1767 (2007).
[29]E. Jaynes, "The Gibbs paradox," in *Maximum Entropy and Bayesian Methods* (Springer, Dordrecht, 1992).
[30]P. M. Ainsworth, Philos. Sci.**79**(4), 542 (2012).
[31]J. W. Gibbs, *On the Equilibrium of Heterogeneous Substances* (Transactions of the Connecticut Academy of Arts and Sciences, Connecticut, 1874).
[32]X. G. Wen, Int. J. Mod. Phys. B **4**(2), 239 (1990).
[33]É. Roldán, I. A. Martínez, J. M. R. Parrondo, and D. Petrov, Nat. Phys. **10**(6), 457(2014).
[34]G. Benenti, K. Saito, and G. Casati, Phys. Rev. Lett. **106**(23), 230602 (2011).
[35]R. Balian, Stud. Hist. Philos. Mod. Phys. **36**(2), 323 (2005).
[36]M. O. Scully, M. S. Zubairy, G. S. Agarwal, and H. Walther, Science **299**(5608), 862 (2003).
[37]R. Landauer, IBM J. Res. Dev. **5**(3), 183 (1961).
[38]A. Tuncer and Ö. E. Müstecaplioğlu, "Quantum thermodynamics and quantum coherence engines," Turk. J. Phys. **44**, 404 (2020); available at https://journals.tubitak.gov.tr/physics/vol44/iss5/1/.
[39]Of course, Curie's principle may also be applied to symmetries which are not associated with groups of automorphisms on microstates in this way, and in such cases, we require a different sort of explanation, such as Chalmers' derivation from the invariance properties of deterministic physical laws, Ref. 26.
[40]Though as noted in Ref. 28, there certainly exist some dynamical symmetry transformations which are not so linked, either because the transformation employed by the dynamical symmetry fails to leave the action invariant, as for rescaling transformations, or because the dynamical system is not amenable to a Lagrangian description at all, and therefore, it has no variational symmetries.