Further insights into the thermodynamics of the Kitaev honeycomb model

Kexin Feng, Natalia B. Perkins ●, and F. J. Burnell

School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA

Here we revisit the thermodynamics of the Kitaev quantum spin liquid realized on the honeycomb lattice. We address two main questions: First, we investigate whether there are observable thermodynamic signatures of the topological Majorana boundary modes of the Kitaev honeycomb model. We argue that for the time-reversal invariant case the residual low-temperature entropy is the primary thermodynamic signature of these Majorana edge modes. Using large-scale Monte Carlo simulations, we verify that this residual entropy is present in the full Kitaev model. When time-reversal symmetry is broken, the Majorana edge modes are potentially observable in more direct thermodynamic measurements such as the specific heat, though only at temperatures well below the bulk gap. Second, we study the energetics, and the corresponding thermodynamic signatures, of the flux excitations in the Kitaev model. Specifically, we study the flux interactions on both cylinder and torus geometries numerically and quantify their impact on the thermodynamics of the Kitaev spin liquid by using a polynomial fit for the average flux energy as a function of flux density and extrapolating it to the thermodynamic limit. By comparing this model to Monte Carlo simulations, we find that flux interactions have a significant quantitative impact on the shape and the position of the low-temperature peak in the specific heat.

I. INTRODUCTION

The Kitaev model on the honeycomb lattice [1] and related models have recently attracted much interest in both theoretical and experimental communities [2–5]. Strikingly, these 2D and 3D Kitaev models have exact quantum spin liquid (QSL) ground states [1] and can be potentially realized in Mott insulating magnets on tricoordinated two- (2D) and three-dimensional (3D) lattices with the strong spin-orbit coupling [3,4,6–9].

One hallmark of a quantum spin liquid is fractionalization. Spin excitations in the Kitaev model are fractionalized into two types of quasiparticles: itinerant spinonlike excitations, which are described by the Majorana fermions which are gapless or gapped depending on the coupling parameters, and localized gapped $\mathbb{Z}_2$ fluxes, also referred to as visons [1]. Much effort has been devoted to searching for traces of such fractionalization in the spin dynamics of α-RuCl$_3$ [10–14] and H$_3$LiIr$_2$O$_6$ [15], which are believed to be proximate to the Kitaev QSL [6]; evidence suggestive of Majorana fermions has been found in these compounds at temperatures up to 100 K. Though most of these materials have magnetically ordered ground states, and hence are not spin liquids, the idea is that even if residual long range magnetic order sets in below a certain temperature, the fractionalized quasiparticles of the nearby QSL phase may still lead to characteristic signatures in the dynamical response spectrum reminiscent of the nearby QSL [2,5,16]. One promising route to look for such signatures is by using various dynamical probes, such as inelastic neutron scattering [17–20], Raman scattering with visible light [21–30], resonant inelastic x-ray scattering [31–33], and through the phonon dynamics [34,35].

Another important probe of fractionalization in the Kitaev QSL is thermodynamics [5,36–40]. Not surprisingly, the emergent fractionalized quasiparticles of the Kitaev QSL reveal themselves in the thermodynamic behavior in a peculiar manner. In the 2D Kitaev QSL, two characteristic crossovers are seen in the specific heat, indicating a two-stage release of magnetic entropy. The first is associated with itinerant fermionic excitations, the second with the localized $\mathbb{Z}_2$ fluxes [36,37]. In 3D Kitaev QSLs, flux freezeout is associated with a phase transition in which topological order is lost. Under certain conditions in three-dimensional tricoordinated lattices, coexistence between the low-temperature chiral QSL phase and crystalline ordering of the $\mathbb{Z}_2$ fluxes has also been observed numerically [39].

The goal of the present work is to revisit the thermodynamics of the Kitaev model on the honeycomb lattice with two main foci. First, we scrutinize the impact of the boundary geometry on the specific heat. In the presence of a boundary, one distinctive signature of the Kitaev spin liquid phase is the existence of topologically protected Majorana boundary modes [26,27,41–44]. Because these protected boundary modes are associated with fractionalized Majorana fermion excitations, they are more difficult to detect than the boundary modes characteristic of topological insulators and superconductors, and thermodynamic measurements have been proposed [43] as one approach. Our second focus is to study the impact of flux interactions on the thermodynamics of the Kitaev model and present a quantitative model of the low-temperature specific heat that we argue captures the main features of the crossover in the thermodynamic limit.

For the first question, our main results are as follows. We show that in principle, the Majorana boundary modes can be

observed in the thermodynamics. Specifically, we find that when time-reversal symmetry is broken, or with unbroken time-reversal symmetry on very narrow nanoribbons, the Majorana boundary modes lead to a low-temperature power law in the specific heat that differs from the $C(T) \sim T^2$ power-law behavior of the bulk. When time-reversal symmetry is unbroken, however, the temperatures at which these power laws occur in currently available materials are well below the range accessed by current experiments. In this case, the Majorana boundary modes can be detected only indirectly, from their contribution to the residual low-temperature entropy. We expect that this residual entropy can be detected using methods similar to those employed in spin ice, where the macroscopic ground state degeneracy contributes to a finite entropy density at ultralow temperatures [45–47], which has been detected using highly accurate measurements of the specific heat [46,48,49]. For Kitaev materials, such experiments offer an alternative to resonant Raman scattering, which can also be used to detect the fractionalized Kitaev boundary modes under certain conditions [25,27].

For the second focus, we study in detail several aspects of the energetics of flux excitations in the Kitaev model that have not been scrutinized in the literature. First, we show that on a lattice with open boundaries, the energy cost of a single flux is significantly reduced near the boundary relative to its bulk value; this has a noticeable impact on the specific heat, as it leads to a lower onset temperature for flux excitations on the cylinder relative to the torus. Second, we study the energetics of fluxes at finite density. We use the results to construct a model of the flux-only contribution to the specific heat in the thermodynamic limit and show that this has good quantitative agreement with the low-temperature peak in the specific heat obtained from the Monte Carlo (MC) simulations in finite systems. This affirms the conclusion of previous numerical work [37] that in the 2D Kitaev model there is no finite-temperature phase transition but rather a crossover from a low-temperature region with vanishingly small flux density to an intermediate temperature region where fluxes (but not fermionic excitations) have proliferated.

Our model is based on an analysis of the distribution of the energies of different flux configurations at fixed flux density. We show that for sufficiently large torus lattices, this distribution is sharply peaked and is approximately independent of lattice size. Therefore, we can model flux thermodynamics by numerically fitting the average energy as a function of flux density. A similar fit is obtained on the cylinder, by separately accounting for flux densities in the bulk and on the boundary. These models shed light on the flux's contribution to the low-temperature specific heat and help identify the role of the boundary fluxes therein. Since the best-fit energy depends only on the flux density, we refer to it as to the pseudopotential energy (PPE), by analogy with the local density approximation (LDA) in density functional theory [50,51]. In Sec. VII, we show that some features of the multiflux interactions can be understood by looking into the microscopics of the two-flux interactions.

The paper is organized as follows. In Sec. II we briefly review the exact solution of the Kitaev model [1]. In Sec. III, we analyze the energy spectrum of the fermionic boundary modes in the flux-free sector and identify the signatures of
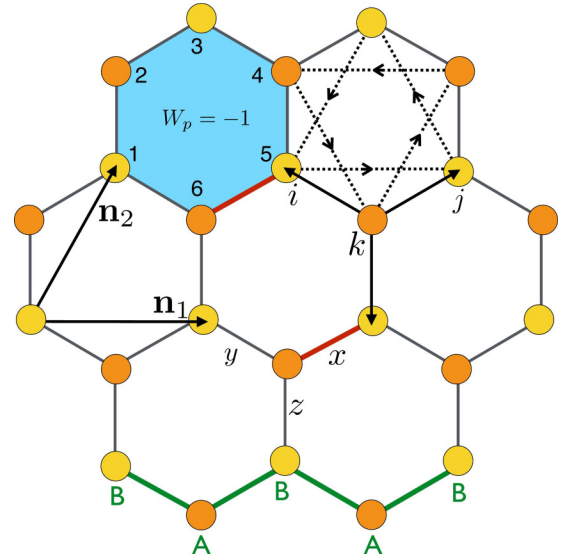

FIG. 1. Kitaev model on the honeycomb lattice. The unit vectors $\mathbf{n}_1 = (1, 0)$, $\mathbf{n}_2 = (\frac{1}{2}, \frac{\sqrt{3}}{2})$. The blue hexagon shows a π flux, i.e., the eigenvalue of $\hat{W}_p = \sigma_1^x \sigma_2^y \sigma_3^z \sigma_4^x \sigma_5^y \sigma_6^z$ is equal to -1. The convention for the sign of the NN couplings on the bonds labeled by x, y, and z is as follows: An arrow pointing from site k to i means u_{ki} on the corresponding bond (k, i) in Eq. (2) is positive. The sites (i, k, j) are an example of the NNN triplet used in Eq. (2). The solid green line shows the zigzag edge of the finite size system.

these modes in the low-temperature scaling behavior of the specific heat. In Sec. IV, we show that the residual entropy from the fermionic boundary modes can be observed in MC simulations of the Kitaev model. In Sec. V, we study the flux energetics in the time-reversal symmetric case. Using best-fit polynomials to describe the PPE, we propose phenomenological flux models for both torus and cylinder lattices, and use these to describe the flux thermodynamics in the thermodynamic limit. Next, Sec. VI focuses on the flux energetics when time-reversal symmetry is broken. We analyze how the specific heat changes with varying magnitude of the time-reversal breaking term. Finally, to better understand the resulting flux PPE models, in Sec. VII we examine the two-flux interactions and show that they capture the essence of the flux energetics in the multiflux systems. Finally, in Sec. VIII, we summarize the main results of this paper.

II. THE MODEL

The extended Kitaev model on a honeycomb lattice is given by the Hamiltonian [1]:

$$H = -\sum_{\langle i,j \rangle_\alpha} J^\alpha \sigma_i^\alpha \sigma_j^\alpha - \kappa \sum_{\langle i,j,k \rangle} \sigma_i^\alpha \sigma_j^\beta \sigma_k^\gamma, \qquad (1)$$

where J^α in the first term denotes the nearest neighbor (NN) Kitaev interaction on the corresponding bond of type $\alpha = x, y, z$ (see Fig. 1), and $\sigma_\mathbf{r}^\alpha$ are the Pauli matrices. The second term is a three-spin interaction on the three adjacent sites (see Fig. 1) of strength κ, which breaks time-reversal symmetry (TRS). It mimics the effect of a magnetic field but preserves exact solubility of the model [1].

At the heart of exact solvability is the macroscopic number of local symmetries in the plaquette operators $[H, \hat{W}_p] = 0$, where $\hat{W}_p$ are defined as $\hat{W}_p = \sigma_1^x \sigma_2^y \sigma_3^z \sigma_4^x \sigma_5^y \sigma_6^z$ (see Fig. 1). These operators commute with the Hamiltonian and take eigenvalues of ± 1. Thus the Hilbert space can be divided into eigenspaces of $\hat{W}_p$, and the ground state is the one with all W_p equal to one, which is also referred to as the flux-free sector.

Using Kitaev's representation of spins in terms of Majorana fermions [1], $\sigma_j^\alpha = i b_j^\alpha c_j$, we rewrite the spin Hamiltonian (1) as

$$H = -\sum_{\langle ij \rangle_\alpha} J^{\alpha_{ij}} i \hat{u}_{ij} c_i c_j - \kappa \sum_{\langle\langle ij \rangle\rangle} i \hat{u}_{ik} \hat{u}_{jk} c_i c_j, \qquad (2)$$

where the bond operators are defined as $\hat{u}_{ij} = i b_i^\alpha b_j^\alpha$, and $\langle ij \rangle$ and $\langle\langle ij \rangle\rangle$ denote NN and next NN (NNN) bonds, respectively. When $\kappa \neq 0$ and time reversal symmetry is broken, Eq. (2) is closely related to the Haldane model for the anomalous quantum Hall effect [52].

The Hilbert space of the fermionic model is larger than that of the spin model; the latter is recovered when we impose the constraint $b_j^x b_j^y b_j^z c_j |\Psi\rangle = |\psi\rangle$ at each site j of our lattice. This constraint commutes with the Hamiltonian, as well as with the bond operators $\hat{u}_{ij}$. Thus the eigenvalues $u_{ij} = \pm 1$ of the operators $\hat{u}_{ij}$ are constants of motion of the model (2) and can be understood as $\mathbb{Z}_2$ gauge fields. This picture captures the fact that not all choices of $\{u_{ij}\}$ correspond to distinct physical states of the spin model, and only those that are gauge inequivalent should be treated as distinct. In the fermionic representation, $\hat{W}_p = \prod_{(i,j) \in \text{edge}(p)} \hat{u}_{ij}$, and $\hat{W}_p$ can be understood as a gauge invariant Wilson loop operator around a single plaquette p, with the eigenvalue $W_p = -1$ corresponding to a (gapped) π-flux excitation on the plaquette in question. This is why in the following we will also use notation $\{\phi_p\}$ to denote a particular flux sector, described (in a given gauge) by choosing a particular configuration of the eigenvalues $\{u_{ij}\}$.

Topological band structure of the Kitaev spin liquid

One interesting feature of the Kitaev spin liquid is that, with u_{ij} chosen such that the hopping matrix element from the A sublattice to the B sublattice is $+i$, for $\kappa = 0$ Eq. (2) describes a band structure in the symmetry class BDI whose band crossings are topologically protected [43], leading to a protected, zero energy band of states at the system's boundary in the thermodynamic limit [42]. These topologically protected boundary (or edge) modes are characteristic of the Kitaev QSL *phase*, rather than the specific Kitaev Hamiltonian. When time-reversal symmetry is broken for $\kappa \neq 0$, the bulk band structure becomes gapped with a nonvanishing Chern number, leading to topologically protected chiral edge modes. We briefly review the nature and origins of these Majorana edge modes, as this understanding is crucial for the remaining discussion.

To understand the topology of the band structure, it is helpful to write the Majorana fermion Hamiltonian Eq. (2) in the zero-flux sector as $H = \sum_\mathbf{k} \mathbf{c}_\mathbf{k}^T H_\mathbf{k} \mathbf{c}_\mathbf{k}$, where $\mathbf{c}_\mathbf{k}^T = (c_{A,\mathbf{k}}, c_{B,\mathbf{k}})$, and

$$H_\mathbf{k} = \mathbf{d}_\mathbf{k} \cdot \boldsymbol{\sigma}, \qquad (3)$$
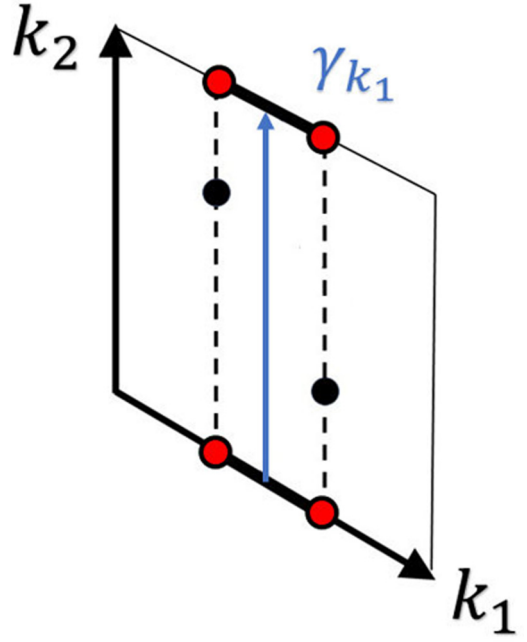

FIG. 2. The 2D Brillouin zone with two Dirac points. γ_{k_1} (blue vertical line) is the integration loop path along which the winding number is calculated. The bold segment on the $\mathbf{k_1}$ axis denotes the region where the gapless edge mode exists (in the limit of large system size) when the lattice is made open in the $\mathbf{n}_1$ direction. It is referred to as a topologically nontrivial region.

where $\boldsymbol{\sigma}$ is a Pauli matrix over the two sublattices of the honeycomb lattice [27,44]. Explicitly, for $\kappa = 0$, we have

$$
\begin{aligned}
d_\mathbf{k}^x &= -J^x \sin k_2 - J^y \sin(k_2 - k_1), \\
d_\mathbf{k}^y &= -J^z - J^x \cos k_2 - J^y \cos(k_2 - k_1), \\
d_\mathbf{k}^z &= 0,
\end{aligned}
\qquad (4)
$$

where $k_i \equiv \mathbf{k} \cdot \mathbf{n}_i$. The Fermi surface then consists of the pair of Dirac points obtained by taking $d_\mathbf{k}^x = 0$ and $d_\mathbf{k}^y = 0$. Time-reversal symmetry and particle-hole symmetry constrain the $\mathbf{d}_\mathbf{k}$ vector to lie in the plane of the 2D honeycomb lattice, and the Hamiltonian (3) can be used to define the topological winding number [43,44,53]:

$$
\begin{aligned}
\nu[\gamma_{k_1}] &= \frac{1}{4\pi i} \int dk_2 \, \text{Tr}\left[H_\mathbf{k}^{-1} \sigma^z \partial_{k_2} H_\mathbf{k}\right] \\
&= \frac{1}{2\pi} \int dk_2 \, \text{Im}\left[\frac{\partial_{k_2} \Gamma_\mathbf{k}}{\Gamma_\mathbf{k}}\right],
\end{aligned}
\qquad (5)
$$

where γ_{k_1} is a path that traverses the Brillouin zone in the k_2 direction at fixed k_1 (see Fig. 2). For $J^x = J^y = J^z = J$, the winding number $\nu[\gamma_{k_1}] = 1$ if $k_1 \in [\frac{2\pi}{3}, \frac{4\pi}{3}]$ and is 0 elsewhere.

On a cylinder with open boundaries in the $\mathbf{n}_2$ direction and periodic boundaries in the $\mathbf{n}_1$ direction, the effective Hamiltonian at a fixed value of the conserved momentum k_1 describes a 1D Majorana chain with open boundary conditions. The topologically nontrivial winding implies that each 1D Majorana chain at fixed $k_1 \in [\frac{2\pi}{3}, \frac{4\pi}{3}]$ hosts zero-energy topologically protected Majorana boundary modes [27,44,54,55].

Breaking time-reversal symmetry by taking $\kappa \neq 0$ introduces a diagonal term $d_{\mathbf{k}}^z \sigma_z$ into the bulk Hamiltonian (3) with

$$d_{\mathbf{k}}^z = 2\kappa(\sin(k_2) + \sin(k_1 - k_2) - \sin(k_1)). \quad (6)$$

The bulk band structure becomes gapped and acquires nonvanishing Chern number [1]. This destroys the boundary flat bands described above and leads instead to chiral edge modes and a corresponding thermal Hall conductance.

III. THERMODYNAMIC SIGNATURES OF MAJORANA EDGE MODES

The topological boundary modes of the Kitaev spin liquid are a distinctive signature of the spin liquid phase at temperatures well below the flux gap, where flux excitations are exponentially suppressed [56]. In this section, we discuss under what conditions these gapless edge modes can be experimentally detected with low-temperature thermodynamic probes by studying the fermionic contribution to the specific heat—i.e., by studying the specific heat of the model (2) in the absence of fluxes. We expect this to be a good description of the actual specific heat of the Kitaev QSL at temperatures well below the flux gap. Throughout this section, we work on a cylindrical lattice with open boundaries in the $\mathbf{n}_2$ direction and periodic boundaries in the $\mathbf{n}_1$ direction (see Fig. 1). We let N_2 denote the number of unit cells in the $\mathbf{n}_2$ direction. Since the system is translationally invariant in the $\mathbf{n}_1$ direction, we set the distance a between two A sites to $a = 1$ and diagonalize the Hamiltonian (3) for each k_1 to obtain the dispersion.

A. Specific heat of fermions

First, we briefly review how the low-temperature fermionic specific heat is calculated. Suppose that the fermionic dispersion in an energy window $\epsilon_{\min} \leqslant \epsilon \leqslant \epsilon_{\max}$ is well approximated by a power law of the form $\epsilon - \epsilon_{\min} \sim k^\alpha$ ($\alpha \geqslant 1$). The density of states in this energy window is given by

$$D(\epsilon) = \mathcal{A} \int \delta(\epsilon - \epsilon_k) dk$$

$$= \mathcal{A}(\epsilon - \epsilon_{\min})^{-(1 - \frac{1}{\alpha})}, \quad \alpha \geqslant 1, \quad (7)$$

where $\mathcal{A}$ is a normalization constant, which in general will depend on the choice of the interval $(\epsilon_{\min}, \epsilon_{\max})$. The corresponding contribution to the fermionic specific heat can be evaluated via:

$$C = \mathcal{A} \frac{1}{T^2} \int_{\epsilon_{\min}}^{\epsilon_{\max}} D(\epsilon) \frac{\epsilon^2 e^{\beta\epsilon}}{(e^{\beta\epsilon} + 1)^2} d\epsilon \quad (8)$$

$$= \mathcal{A} T^{\frac{1}{\alpha}} \int_{x_0}^{\beta\epsilon_{\max}} (x - x_0)^{-(1 - \frac{1}{\alpha})} \frac{x^2 e^x}{(e^x + 1)^2} dx, \quad (9)$$

where $\beta = 1/T$ and $x_0 = \beta\epsilon_{\min}$. Thus we find

$$C = \mathcal{A} I T^{\frac{1}{\alpha}}, \quad (10)$$

where

$$I = \int_{\epsilon_{\min}}^{\epsilon_{\max}} (x - \beta\epsilon_{\min})^{-(1 - \frac{1}{\alpha})} \frac{x^2 e^x}{(e^x + 1)^2} dx. \quad (11)$$

TABLE I. The quantities k_c, r_c, and Δ^{fs} for a variety of cylinder widths. Here N_2 is the number of unit cells along $\mathbf{n}_2$ direction on a cylinder lattice. The momentum k_c, which indicates the boundary between regions (1) and (2), is given in units of π. r_c is the ratio of the length of the region (2) to the length of the topologically nontrivial region $[2\pi/3, 4\pi/3]$. Δ^{fs} is the fermionic energy at the projected Dirac point.

N_2	5	10	25	50	100	500	1000
k_c/π	1.011	1.06	1.166	1.235	1.277	1.319	1.326
r_c	0.967	0.82	0.502	0.296	0.169	0.042	0.021
Δ^{fs}	0.569	0.299	0.123	0.062	0.031	0.006	0.003

To obtain the low temperature limit, we take $\epsilon_{\min} \ll T$, in which case we can safely treat $\epsilon_{\min}$ as 0. In this case, for $T \ll \beta\epsilon_{\max}$, we have

$$I \approx \int_0^\infty x^{1 + \frac{1}{\alpha}} \frac{e^x}{(e^x + 1)^2} dx$$

$$= \sum_{l=0}^\infty (-1)^l G\left(\frac{1}{\alpha} + 2, l + 1\right). \quad (12)$$

Here we define $G(\nu, l) = \int_0^\infty x^{\nu-1} (\cosh x)^{-l} dx$, which is converging for $\nu > 0$, $l \in \mathbb{N}$. When the temperature increases, $\beta\Delta^{\mathrm{fs}}$ becomes finite and the integral in Eq. (9) needs to be computed more accurately.

Our calculation gives the contribution of states within the energy window $(\epsilon_{\min}, \epsilon_{\max})$ to the specific heat, for any temperature. However, at temperatures comparable to $\epsilon_{\max}$, other higher energy states also begin to contribute to the specific heat via $C = \frac{\mathcal{A}}{T^2} \int_{\epsilon_{\max}}^\infty D(\epsilon) \frac{\epsilon^2 e^{\beta\epsilon}}{(e^{\beta\epsilon} + 1)^2} d\epsilon$. In what follows, we will discuss the appropriate energy window, and corresponding values of α and $n_{\epsilon_{\max}}$, that describe the topological boundary modes of both the nodal topological Kitaev QSL with $\kappa = 0$, and the gapped topological Kitaev QSL obtained for $\kappa > 0$.

B. Time-reversal symmetric Kitaev QSL ($\kappa = 0$)

We first consider $\kappa = 0$, where the edge state consists of a boundary Majorana flat band in the thermodynamic limit. Clearly the entropy associated with this flat band never freezes out; in the thermodynamic limit its thermodynamic signature is a residual zero-temperature entropy, which we discuss in detail in Sec. IV A. However, when the lattice size is finite, the edge modes localized at opposite boundaries have finite overlap, and the boundary bands have finite-size splitting everywhere except at $k_1 = \pi$. Here we determine when this finite-size splitting is large enough to be observed in thermodynamic measurements. We quote all energies in units of J and set the Boltzmann constant $k_B = 1$.

As described above, in order to calculate the specific heat, we first need to understand the low-energy spectrum of our model. To this end, Fig. 3 illustrates the effects of finite size on the fermionic band structure, which is particularly clear from the comparison of the fermion energy spectrum for cylinder lattices with $N_2 = 10$ and $N_2 = 100$ shown in Figs. 3(a) and 3(c), respectively. The finite N_2 manifests itself in two main
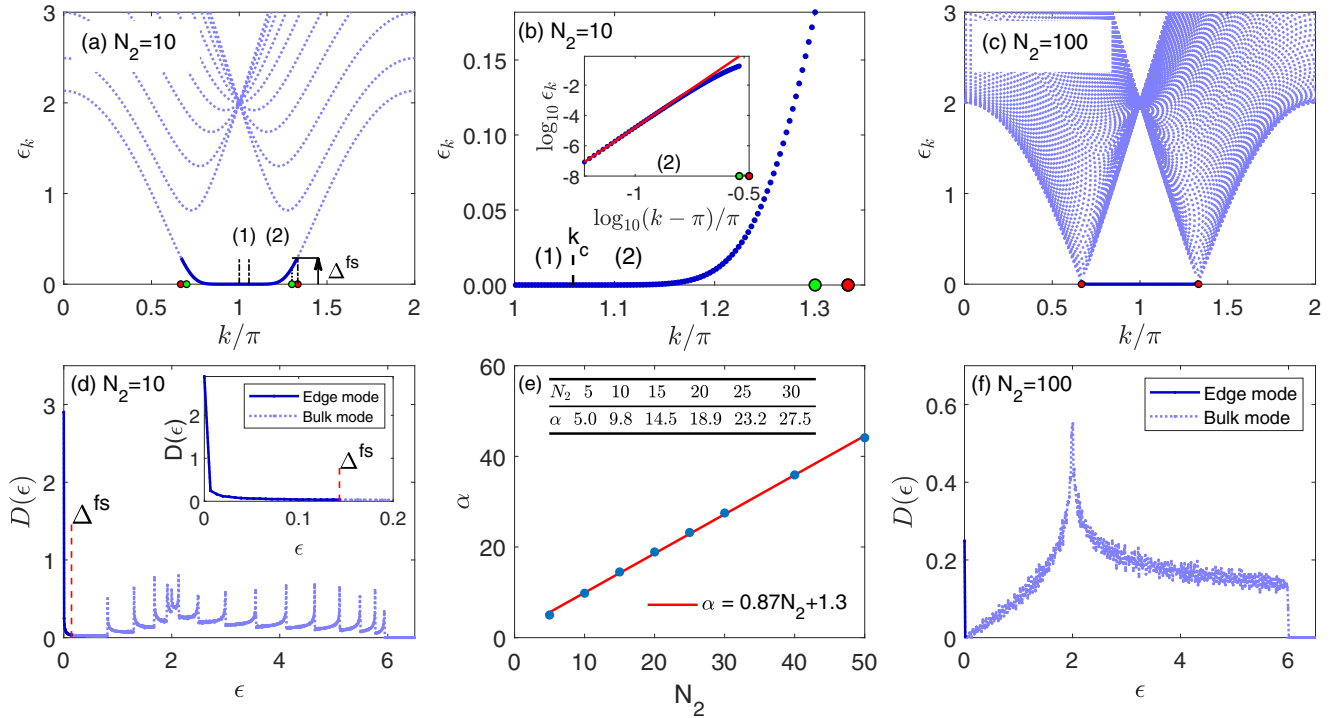
FIG. 3. Fermionic spectrum and density of states (DOS) of the Majorana fermion Hamiltonian (2) with $\kappa = 0$ obtained on the cylinder lattices with $N_2 = 10$ [(a) and (d)] and with $N_2 = 100$ [(c) and (f)]. The nearly-flat lowest energy band in (a) and (c) appears as a near-zero energy peak in the DOS. Here red dots indicate the position of the bulk Dirac points projected to k_1 (i.e., $k_1 = \frac{2\pi}{3}, \frac{4\pi}{3}$), and Δ^{fs} corresponds to the energy of the lowest band at these projected Dirac points. Green dots indicate the analytical solution [57] for the boundary of the k_1 region inside which the states in the lowest band are exponentially localized to the edges. In (a) and (b), the regions (1) and (2) denote $[\pi, k_c]$ and $[k_c, 4\pi/3]$, respectively. Δ^{fs} and k_c for different values of N_2 are given in Table I. (b) A zoomed-in view of the lowest band obtained on the cylinder lattice with $N_2 = 10$. The inset shows the zoomed-in edge mode spectrum in the log-log scale (blue dots), with the red line showing an approximate power-law fit to the dispersion in the leftmost one fifth of region (2). We use this to approximate the dispersion throughout region (2) by a power law, noting however that this deviates substantially from the exact dispersion before the boundary of the region containing exponentially localized edge modes (green dot). (e) The N_2 dependence of the scaling powers of the edge mode dispersion, obtained as shown in (b). Blue dots show the exact values (also quoted in Table I in the inset). The red line $\alpha = 0.87 N_2 + 1.3$ is the linear fit of these data points.

ways. First, the Dirac cones are lifted, and the purely bulk bands (i.e., the second and higher subbands) acquire a finite-size gap, which decreases as $1/N_2$. As can be seen in Fig. 3(c), for $N_2 = 100$ the resulting bulk gap is extremely small compared to the bandwidth. Second, states in the lowest subband localized at opposite edges of the system have finite overlap, leading to a finite-size splitting that falls off exponentially in N_2. In Fig. 3(b), in order to present the details of this finite-size effect, we show the zoomed-in dispersion of the lowest band for $N_2 = 10$.

We numerically introduce the scales ϵ_{max} and ϵ_{min} for this band structure as follows. We take ϵ_{max} to be the energy of the lowest subband at the projected Dirac points $k_1 = \frac{2\pi}{3}, \frac{4\pi}{3}$ [shown by red dots in Figs. 3(a)–3(c)], which is nonzero due to finite-size effects. Thus, we denote $\epsilon_{\mathrm{max}} = \Delta^{\mathrm{fs}}$. Note that states in the lowest subband become delocalized in the bulk at a slightly lower energy than this, at the momentum indicated by the green dot in Fig. 3 (see Ref. [57] for exact expressions); however this difference vanishes as N_2 increases, and is already small for $N_2 = 10$. We take $\epsilon_{\mathrm{min}} = 10^{-7} J$ to be a low-temperature cutoff below which thermodynamic measurements cannot be performed. For the value $J \approx 100$ K, this corresponds to a temperature scale of 10^{-5} K.

The corresponding momentum k_c depends on the value of N_2 (see Table I), approaching $4\pi/3$ as N_2 increases.

To understand the contribution of modes in the energy window $(\epsilon_{\mathrm{min}}, \epsilon_{\mathrm{max}})$ to the specific heat, we approximate the dispersion in the lowest subband in the region $k_c \leqslant k_1 \leqslant 4\pi/3$ [region (2) in Figs. 3(a) and 3(b)] with a power law: $\epsilon_{\mathrm{edge}}^{\mathrm{fs}} \sim k^\alpha$, $\alpha \geqslant 1$. Figure 3(b) (inset) shows this fit on log-log scale for $N_2 = 10$; the best-fit power α for different values of N_2 is presented in Fig. 3(e). We emphasize that because the dispersion is not an exact power law (see Ref. [57]), the precise value we obtain for α depends on the range of momenta that we include in our scaling. Here, we obtain α by fitting only the leftmost fifth of region (2), meaning that our power law is a good approximation for the dispersion at the lowest energies but becomes less accurate as we approach Δ^{fs}. Irrespective of this choice, however, we find that α is large compared to 1, and grows linearly with N_2, reflecting the fact that for $\pi \leqslant k_1 < 4\pi/3$, the lowest subband becomes flatter as N_2 increases. For the range of momenta that we fit to, we find $\alpha \approx 0.87 N_2 + 1.3$, with exact values shown in the inset of Fig. 3(e).

We may now use Eq. (10) to evaluate the contribution of the boundary modes to the specific heat. The total
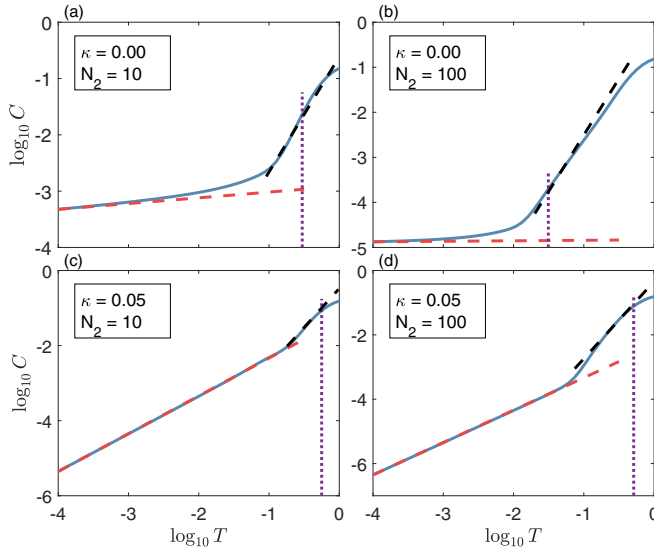
FIG. 4. The total fermionic specific heat, $C_{tot} = C_{edge} + C_{bulk}$, in the flux-free sector for different κ and N_2. The inverse of the slope of the red dashed line gives the power of the low-temperature scaling behavior of the specific heat: (a) $\alpha = 9.8$, (b) $\alpha = 88.3$, (c) $\alpha = 1$, and (d) $\alpha = 1$. The black dashed line shows the expected bulk fermionic contribution to the specific heat with $\alpha = 2$. The purple dotted line marks the temperature corresponding to Δ^{fs}, which is virtually indistinguishable from ϵ_{max} on a logarithmic scale.

number of states in region (2) is approximately $\frac{1}{3}N_1 r_c$, where N_1 is the number of unit cells in the $\mathbf{n}_1$ direction. Here $r_c = (4 - 3k_c/\pi)$ is the length of region (2) divided by $\pi/3$ (which is half the length of the topologically nontrivial region), and $\frac{1}{3}N_1$ gives the number of states in the topologically nontrivial region.

The remaining boundary states, contained in region (1) in Figs. 3(a) and 3(b), do not contribute to the specific heat in the temperature range of interest (where T is large compared to ϵ_{min}); instead they contribute to the zero-temperature entropy. Therefore, for $\epsilon_{min} \ll T \ll \epsilon_{max}$ the specific heat obtained in Eq. (10) becomes

$$C_{edge} = \frac{1}{3}r_c N_1 I T^{\frac{1}{\alpha}}, \tag{13}$$

where I is given by Eq. (12). Note that as $N_2 \to \infty$, r_c [and hence the length of region (2)] approaches 0, as shown in Table I, while $T^{1/\alpha}$ approaches unity. Thus for fixed N_1, the edge's contribution to the specific heat vanishes as N_2 increases.

At temperature scales on the order of Δ^{fs}, the low-temperature fermionic specific heat also has a contribution from the bulk states, which in practice dominates over the contribution from the edge. This is shown in Figs. 4(a) and 4(b), which plot the total fermionic specific heat as a function of temperature for $\kappa = 0$ and $N_2 = 10$ and 100, respectively. The solid blue line is obtained from the exact fermionic density of states in the absence of fluxes, shown for $N_2 = 10$ and 100 in Figs. 3(d) and 3(f), respectively. At temperature scales larger than Δ^{fs}, the specific heat is approximately proportional to T^2 (black dashed line), which is the usual bulk power law from the Dirac cone. We see a crossover from this bulk power

law to the boundary power law in Eq. (13), at a temperature scale of approximately $0.1\epsilon_{max}$, below which the specific heat changes much more slowly with temperature. The red dashed line indicates our prediction of the low temperature specific heat based on the power α shown in Fig. 3(e).

C. System with broken time-reversal symmetry ($\kappa \neq 0$)

In the extended Kitaev model (1), time reversal symmetry is broken by the κ term, which introduces the diagonal term (6) into the bulk Hamiltonian (3). This diagonal term opens a bulk energy gap equal to $\Delta_{bulk} = 6\sqrt{3}\kappa$ in the thermodynamic limit [1]. It also introduces energy dispersion to the edge modes. For small κ, the dispersion of the corresponding edge modes in the thermodynamic limit can be obtained by perturbation theory in κ/J [1,27]:

$$\epsilon_k^{edge} = 12\kappa|\sin k|, \quad k \in [2\pi/3, 4\pi/3]. \tag{14}$$

We will refer to this dispersion as the magnetic dispersion. Note that near the gapless point $k = \pi$, ϵ_k^{edge} is linear in k, with a velocity proportional to κ. At the projected Dirac point, $\epsilon_{4\pi/3}^{edge} = 6\sqrt{3}\kappa$, this agrees with the bulk energy gap Δ_{bulk}, and in the thermodynamic limit, the boundary mode merges with the bulk bands at this point.

However, Eq. (14) does not account for finite-size corrections that modify both the bulk gap and the dispersion of the boundary modes. In order to illustrate how the magnitude of these corrections depends on the magnitude of the time-reversal symmetry breaking term and the size of the system, in Fig. 5 we plot the fermionic energy spectrum for different κ and N_2, where the numerical band structures are shown by solid blue lines and the magnetic dispersions computed from Eq. (14) are shown by red dashed lines. Two observations are in order here. First, we see that for all values κ and N_2, the magnetic dispersion is essentially indistinguishable from the true dispersion over much of the range $2\pi/3 < k_1 < 4\pi/3$, with deviations most noticeable in the immediate vicinity of the projected Dirac points. Second, as expected, the finite-size effects decrease with increasing N_2 [Fig. 5(b)], however they also decrease with increasing κ [Figs. 5(c) and 5(d)].

As for the $\kappa = 0$ case, it is convenient to divide the interval $k_1 \in [\pi, 4\pi/3]$ into two regions, as shown in Fig. 5(a). Region (1) consists of the interval $[\pi, k_c]$, on which the boundary dispersion Eq. (14) deviates from the numerical dispersion of the lowest subband by less than 1%; here finite-size effects are negligible and the gap of the edge mode is dominated by κ. In region (2), consisting of the interval $[k_c, 4\pi/3]$, finite-size effects cannot be neglected, as they lead to a correction of 1% or more relative to Eq. (14). In this region, the dispersion of the edge mode can still be approximated by the power law, $\epsilon_k^{edge(2)} \sim k^{\alpha_2}$; numerically, we find $\alpha_2 = (2.14, 2.02, 1.94, 1.7)$ for the system sizes $N_2 = (10, 25, 30, 50)$. Table II shows the extent of these finite-size effects for a variety of system sizes with $\kappa = 0.01$. Specifically, it lists the momentum k_c at which the deviation from Eq. (14) exceeds 1%, the corresponding energy ϵ_{min}, the ratio $r_1 = 3(k_c/\pi - 1)$ of the length of region (1) to the length of the interval $[\pi, 4\pi/3]$, and the minimum energy Δ^{fs} at the projected bulk Dirac point. We see that by $N_2 = 100$, finite-size corrections in all of these quantities are on the order of
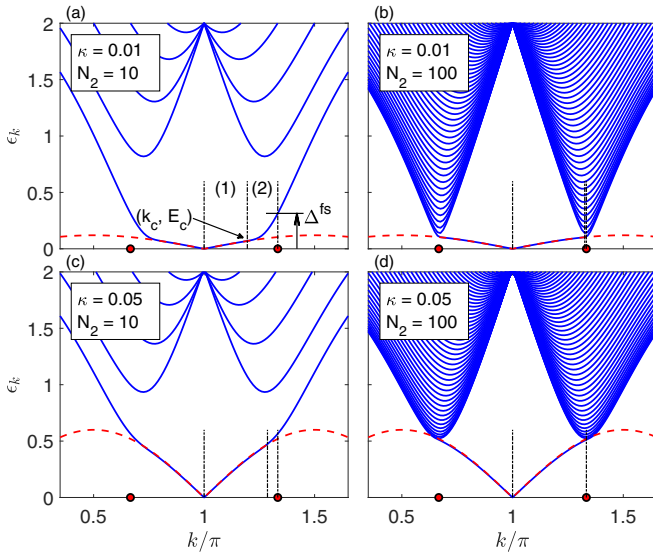
FIG. 5. The fermionic energy spectrum for different κ and N_2. The spectrum obtained by numerically diagonalizing the Majorana fermion Hamiltonian (2) with $\kappa \neq 0$ is shown with blue lines. The magnetic dispersion of the edge mode obtained from the perturbative analytical expression Eq. (14) is shown by the red dashed line. The right half of the topologically nontrivial region $k_1 \in [2\pi/3, 4\pi/3]$ is divided into a region (1) spanning the interval $[\pi, k_c]$ and a region (2), containing $[k_c, 4\pi/3]$. The division is such that in region (1) the deviation between two dispersions is less than 1%, and in region (2) the deviation between two dispersions is more than 1%, so the finite-size effects are essential. Δ^{fs} is the energy of the lowest subband at the projected Dirac point (red dots). Δ^{fs} and k_c for different values of N_2 are given in Table II.

a few percent, and by $N_2 = 500$ they are less than a fraction of a percent. At larger values of κ, the finite-size corrections will be further reduced relative to those shown in Table II, as is clear from Fig. 5(c).

When the finite-size corrections can be neglected, the energy dispersion is well approximated by $\epsilon_k^{\text{edge}(1)} \approx 12\kappa|k - \pi|$, and the DOS is given by $D_1 \sim \kappa^{-1}$ leading to the scaling behavior of the specific heat given by

$$C_{\text{edge}}^{(1)} \sim \tfrac{1}{3}N_1 \kappa^{-1} T, \quad T < \epsilon_{\text{min}}, \tag{15}$$

TABLE II. k_c, ϵ_{min}, r_1, and Δ^{fs} computed for different N_2 and taking $\kappa = 0.01$. k_c corresponds to the point where the energy calculated by numerical diagonalization and the energy from the perturbative analysis are different by 1%, as shown in Fig. 5. r_1 is the ratio of the number of states in region (1) to the number of states in the whole topologically nontrivial region. Δ^{fs} is the fermionic energy at the projected Dirac point.

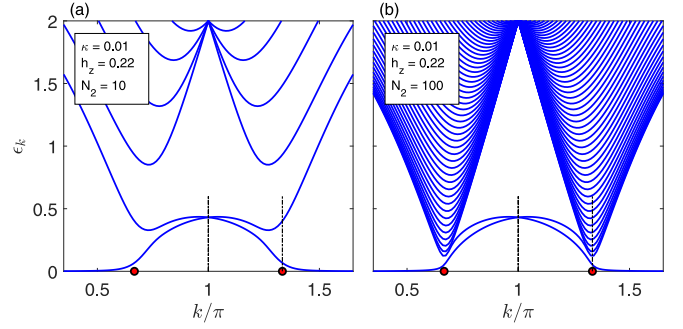| N_2	5	10	25	50	100	500
k_c/π	1.101	1.201	1.285	1.315	1.328	1.333
ϵ_{min}	0.038	0.072	0.095	0.101	0.104	0.104
r_1	0.302	0.604	0.856	0.944	0.984	1.000
Δ^{fs}	0.576	0.314	0.159	0.12	0.108	0.104

FIG. 6. The fermionic energy spectrum for various values of $\kappa = 0.01$ and (a) $N_2 = 10$ and (b) $N_2 = 100$ computed in the presence of the Zeeman coupling h between the dangling edge fermions and the itinerant fermions. Here we set the external filed components $h_x = h_y = h_z = \kappa^{1/3}$. The red dots show the projected Dirac points.

in agreement with Ref. [27]. Therefore, a low-temperature specific heat $C/T \sim \kappa^{-1}$ at scales below the bulk gap is a signature of the edge mode in the presence of time reversal symmetry breaking. Unlike the time-reversal symmetric case, this signature is robust in the thermodynamic limit.

We note that for very small N_2, finite-size effects can lead to observable deviations from the scaling (15) for temperatures below the bulk gap but above the energy scale ϵ_{min}. In region (2), where $\epsilon_{\text{min}} \leqslant \epsilon_k^{\text{edge}(2)} \leqslant \epsilon_{\text{max}} = 6\sqrt{3}\kappa$, the corresponding DOS scales as $D_2 \sim (\epsilon - \epsilon_{\text{min}})^{1/\alpha_2 - 1}$ and contributes a term to the specific heat that scales as:

$$C_{\text{edge}}^{(2)} = \tfrac{1}{3}N_1(1 - r_1)IT^{1/\alpha_2}, \tag{16}$$

where the constant I is given by Eq. (12), with $\epsilon_{\text{max}} = 6\sqrt{3}\kappa$.

Figures 4(c) and 4(d) show the total fermionic specific heat for $\kappa = 0.01$. At temperatures well below the bulk gap, the specific heat fits well to the analytical result in Eq. (15) for both $N_2 = 10$ [panel (c)] and $N_2 = 100$ [panel (d)]. At temperatures on the order of 10% of the bulk gap $\Delta_{\text{bulk}} = 6\sqrt{3}\kappa$, bulk states begin to dominate the fermionic specific heat, which fits reasonably well to the T^2 dependence expected for a 2D Dirac cone.

In the pure Kitaev model, a further correction to the edge spectrum at finite magnetic field comes from the Zeeman coupling to spins on the boundary [1]. In the absence of magnetic field, these edge spins are associated with a nontopological zero-energy boundary flat band of "dangling" Majorana fermions, which happen to be completely decoupled from the bulk of the system (see Sec. IV A). Since the resulting flat band is nontopological, we expect that in real materials, terms that perturb the bulk Hamiltonian away from the Kitaev limit will gap out these states; hence we have ignored them in our analysis thus far. However, it is worth noting that if we consider the Zeeman term as the *only* coupling of our system to these dangling boundary fermions, the edge spectrum is modified significantly, since the zero-energy trivial flat band hybridizes strongly with the chiral boundary modes that we have focused on here. This is shown in Fig 6, where we include a coupling $h_z \sum_i c_i b_i^z$ between the dispersing Majorana c_i and the dangling Majorana b_i^z on each site i on the edge, via a Zeeman field $h_z = \kappa^{1/3}$. We see that the velocity

of the chiral edge mode is drastically reduced through hybridization with the boundary flat band and that the gapless point is moved from the center to the edge of our Brillouin zone. Though this does not alter the qualitative nature of the low-energy power law arising due to the chiral boundary modes, it does significantly alter the energy scale at which it sets in.

D. Specific heat as a probe of topological boundary modes

We now turn to the question of whether the Majorana edge modes can be detected experimentally using realistic specific heat measurements. As observed above, for $\kappa = 0$ this requires measurements to be made at temperatures below the finite-size gap Δ^{fs}. The best chance of detection is therefore in nanoribbons with the smallest possible N_2, which increases both the temperature scale at which the specific heat becomes dominated by C_{edge} and the magnitude of the edge contribution at the higher end of this range. To the best of our knowledge, the smallest α-RuCl$_3$ monolayer nanosheet has a size of around 10 μm × 10 μm × 2 nm. Given that the spacing of the ions in α-RuCl$_3$ is around 10 Å [58,59], the number of the unit cells of the nanosheet is estimated to be $N_1 \times N_2 \sim 10^4 \times 10^4$. The corresponding high temperature cutoff is set by $\Delta^{\mathrm{fs}} \sim 2 \times 10^{-4}$ which is around 2×10^{-2} K, given that $J \approx 100$ K [17,22]. To the best of our knowledge, the lowest accessible temperature for the specific heat measurements is on a scale of 0.1 K [38,60], which is three orders of magnitude larger than Δ^{fs}. We conclude that direct observation of the Majorana boundary modes in the specific in α-RuCl$_3$ without time-reversal symmetry breaking is likely out of reach for current experiments.

For $\kappa > 0$, on the other hand, the linear-in T specific heat predicted in Eq. (15) up to a temperature scale corresponding to approximately $0.1\Delta_{\mathrm{bulk}}$, followed by a crossover to the quadratic temperature dependence of the bulk, could potentially be observed even for realistic sample sizes. For example, with $N_2 \sim 10^4$ and $\kappa = 0.01$, we have $\Delta_{\mathrm{bulk}} = 6\sqrt{3}\kappa = 0.104J$, which gives a bulk gap of 10 K and a temperature scale of 1 K for T-linear specific heat to be observed α-RuCl$_3$. We caution, however, that in real experiments time-reversal symmetry is typically broken by applying a magnetic field; in this situation our analysis applies only for κ much smaller than the bulk flux gap.

Note that in realistic samples, the edges will not be of the perfect zigzag type considered here, but rather will include a degree of boundary disorder. While disorder only at the boundary cannot eliminate the gapless chiral boundary modes found at $\kappa \neq 0$, the fate of the boundary states of the nodal topological phase for $\kappa = 0$ is much less well understood. However, a few remarks on its likely impact are in order here. Disorder can include both random steps (or more generally, segments of armchair edge) interspersed with segments of zigzag edges, as well as random vacancies along the zigzag edge. In the former case, the armchair edge segments do not contribute to the boundary states [57,61]. If the zigzag segments are long on the scale of the lattice constant, we expect that the system can be heuristically understood as a series of short, disconnected zigzag ribbons,

such that much of the phenomenology described here will be qualitatively unchanged at energy scales that are larger than the typical segment length. If the zigzag segments are short, then we expect the effect of disorder to be qualitatively similar to adding a high density of vacancies on the edge. The latter can be viewed as a particular type of (strong) hopping disorder at the boundary of the zigzag nanoribbon. The random-hopping problem has been much studied in 1D systems with linear dispersion and with the same symmetries as the Kitaev model (see, e.g., Refs. [62–64]). It is known that in this case, the disorder does not gap out the low-energy modes, but instead leads to a low-energy divergence in the density of states, which can be either power-law or Dyson-like ($\rho(\epsilon) \sim \epsilon^{-1}|\ln^3 \epsilon|$). Thus we expect that the system will retain a divergent low-energy density of states and a corresponding contribution to the zero-temperature entropy; however, the low-temperature specific heat, which is sensitive to the precise nature of this divergence, will evidently be sensitive to such disorder. To fully understand this dependence requires further systematic analysis, similar to that performed in Refs. [65,66], where the finite-size effects were properly accounted for.

IV. RESIDUAL ENTROPY AND GAPLESS EDGE MODES IN THE TIME-REVERSAL SYMMETRIC KITAEV MODEL

In the previous section we discussed the fermionic contribution to the specific heat, which we expect to dominate at temperatures well below the temperature scale associated with the flux gap. In the time-reversal invariant case, we argued that the boundary modes do not give a measurable contribution to the specific heat for realistic sample sizes; instead their primary thermodynamic signature is a residual low-temperature entropy S_0. Here we derive the value of S_0 in the absence of fluxes and compare this to S_0 for the full Kitaev spin liquid. To this end, we will compute the specific heat by using the MC method for the Kitaev spin liquid developed by Nasu, Udagawa, and Motome [36]. The entropy is then given by

$$S(T) = S_\infty - \int_T^{T_{\mathrm{max}}} C \, d \ln T, \quad (17)$$

where S_∞ is the entropy in the limit $T \to \infty$, i.e., the maximum entropy of our system, and we choose $T_{\mathrm{max}} = 10^{1.5}$, as above this the specific heat is vanishingly small. Recall that throughout we work in units where $J = 1$. In the limit $T \to 0$, the residual low-temperature entropy is estimated numerically using $S_0 = S(T_{\mathrm{min}})$, where $T_{\mathrm{min}} = 10^{-2}$ is the lowest temperature simulated. In order to extract an approximate finite-size scaling of S_0, we compare numerical results on tori (with no boundary modes) and cylinders (with boundary modes) of different sizes.

A. S_∞ in the system with open boundaries

At this point a note on S_∞ is in order. For a system of N_s spin-1/2 degrees of freedom, $S_\infty = \ln \Omega = N_s \ln 2$, where Ω is the number of states. However, on a lattice with open boundary conditions, this does not correspond exactly to the infinite-temperature entropy of our flux-fermion model. This

is because the Kitaev spin model (1) contains extra static degrees of freedom at the boundary [1] that are not included in the flux-fermion model (2).

To see this, consider a zigzag boundary with A sites on the boundary (such a zigzag edge is shown by the solid green line in Fig. 1). For a given site on the boundary, the Hamiltonian includes only two terms:

$$H_{\text{edge}} = \sum_i \sigma_{i,A}^x \sigma_{i,B}^x + \sigma_{i,A}^y \sigma_{i-1,B}^y. \quad (18)$$

Thus in addition to the conserved plaquette operators W_p, there is a subextensive set of operators

$$W_i^{(b)} = \sigma_{i,A}^y \sigma_{i,B}^z \sigma_{i+1,A}^x, \quad (19)$$

which commute with the Hamiltonian (1) for every site i on the boundary. In the Majorana fermion representation, by applying the constraint operator $\hat{D}_j = b_j^x b_j^y b_j^z c_j$, $W_i^{(b)}$ operator can be expressed as $W_i^{(b)} = -i\hat{u}_{iA,iB}\hat{u}_{(i+1)A,iB}b_{i,A}^z b_{i+1,A}^z$. It is always possible to work in a gauge where all $\hat{u}_{ij}$ on the boundary are $+1$, where $W_i^{(b)} = -ib_{i,A}^z b_{i+1,A}^z$. Since the "dangling" fermions $b_{i,A}^z$ (for unit cells i on the edge) do not enter the Hamiltonian at all, the two possible values of $W_i^{(b)}$ correspond to states with the same energy. This leads to a twofold degeneracy in the spectrum for every pair of dangling fermions (i.e., every pair of unit cells) on the boundary, i.e., a total of 2^{N_1} zero-energy states from both boundaries combined.

These zero-energy states, which we refer to as "dangling fermion states," are not included in the model (2) but must be accounted for in calculating S_∞. Specifically, each pair of unit cells on the zigzag boundary contributes an additional $\ln 2$ to the entropy of the spin model, which we must subtract to obtain S_∞. Thus, on a cylinder with N_1 unit cells on each of the two zigzag boundaries we find that

$$S_\infty = (N_s - N_1) \ln 2. \quad (20)$$

B. Results of Monte Carlo simulations

The thermodynamics of the Kitaev honeycomb model has been studied previously by Nasu, Udagawa, and Motome [36,37]. In particular, using MC simulations they have determined the basic structure of the specific heat of the Kitaev honeycomb model on finite-size systems with $L \times L$ unit cells, using both periodic and open boundary conditions. It was found that the specific heat $C(T)$ has a two-peak structure, with the low-temperature peak at $T = T_L$ associated with the freezeout of flux excitations, and a high-temperature peak from the fermionic excitations at $T = T_H$. Since this is important for our discussion, we review the derivations of the relevant thermodynamic quantities for the Kitaev model in Appendix B.

Here we extend the analysis of Refs. [36,37] and present a detailed study of the specific heat and entropy of the Kitaev model based on extensive MC simulations on various finite-size lattices with periodic (tori) and semiperiodic (cylinders) boundary conditions. Our main focus is on exploring the contribution of the Majorana fermion boundary modes to the specific heat and the entropy. As such, we focus on systems with $N_1 = 4N_2$, leading to a long boundary in the cylinder's case, thereby maximizing effects of the edge states. A short

description of the implementation of the MC algorithm [36] is outlined in Appendix C.

The MC results for the specific heat are presented in Fig. 7(a). As in Refs. [36,37], we observe a two-peak structure for the specific heat on both torus and cylindrical lattices, with a high-temperature peak that is basically insensitive to both the system size and boundary conditions. This is expected since the bulk band structure for these high-energy modes is qualitatively similar in all cases. In contrast, the low-temperature peak (T_L) is markedly different between the torus and the cylindrical lattices. Notably, the onset temperature for this lower peak on a cylindrical lattice is lower than on a comparably sized torus lattice. This indicates a smaller average flux gap due to the presence of open boundaries. Finite-size effects are also significant for the lower peak, primarily due to interactions between fluxes that are enhanced by small system size. We will discuss both of these effects in detail in Sec. V. Figure 7(b) presents the corresponding entropy, computed by using Eq. (17). The two peaks in $C(T)$ lead to a two-step entropy release, with entropy release near T_L predominantly due to the proliferation of flux degrees of freedom, while that near T_H stems from the proliferation of high-energy itinerant fermionic degrees of freedom.

Figure 7(b) also clearly shows the difference between torus and cylindrical lattices. On torus lattices, the entropy decreases to nearly zero at the lowest temperatures. This is expected: In the thermodynamic limit all the physical degrees of freedom except the states very near the gapless Dirac points are frozen out at temperatures small compared to the flux gap. On finite-sized tori the freezeout is more pronounced, since the Dirac points acquire a gap on the order of $1/N_2$ unless N_1 and N_2 are both divisible by 3 (see Appendix D), leading to a freezeout of all degrees of freedom below this scale. In contrast, the residual entropy on cylindrical lattices is nonzero due to the band of low-energy fermionic edge modes. Specifically, with a low-temperature cutoff at a scale $T_{\min}$, the residual entropy S_0 is approximately given by the number of states with energy less than $T_{\min}$. For the model (2), and $T_{\min} \sim \Delta^{\text{fs}}$, this is approximately:

$$S_0 = \ln \Omega_0, \quad (21)$$

where Ω_0 is the number of states associated with the zero-energy edge mode, $\Omega_0 \approx 2^{\frac{1}{3}N_1}$. Note that for the small values of N_2 simulated here, $\Delta^{\text{fs}} > T_{\min} = 10^{-2}$, and we find $S_0 \lesssim \frac{1}{3}N_1 \ln 2$. To obtain the residual entropy associated with the original Kitaev spin model (1), we must add the contribution from dangling gauge fermions discussed in Sec. IV A giving

$$S_0^{\text{spin}} = \ln \Omega_0 + N_1 \ln 2. \quad (22)$$

A quantitative comparison between these analytical estimates and the MC results is given in Table III, which compares various ways of evaluating $s_0^{\text{spin}} \equiv S_0^{\text{spin}}/(N_s \ln 2)$. In the table, $S_0^{\text{spin}}(\text{I}) = \frac{4}{3}N_1 \ln 2$ is the residual entropy of the spin model in the thermodynamic limit, i.e., assuming that all edge states contribute to S_0; $S_0^{\text{spin}}(\text{II})$ is the residual entropy obtained from Eq. (21) by taking Ω_0 equal to the number of fermionic states with energy less than $T_{\min}$; $S_0^{\text{spin}}(\text{III})$ is the residual entropy obtained from the MC results using Eq. (17). In both $S_0^{\text{spin}}(\text{I})$ and $S_0^{\text{spin}}(\text{II})$ we have added the
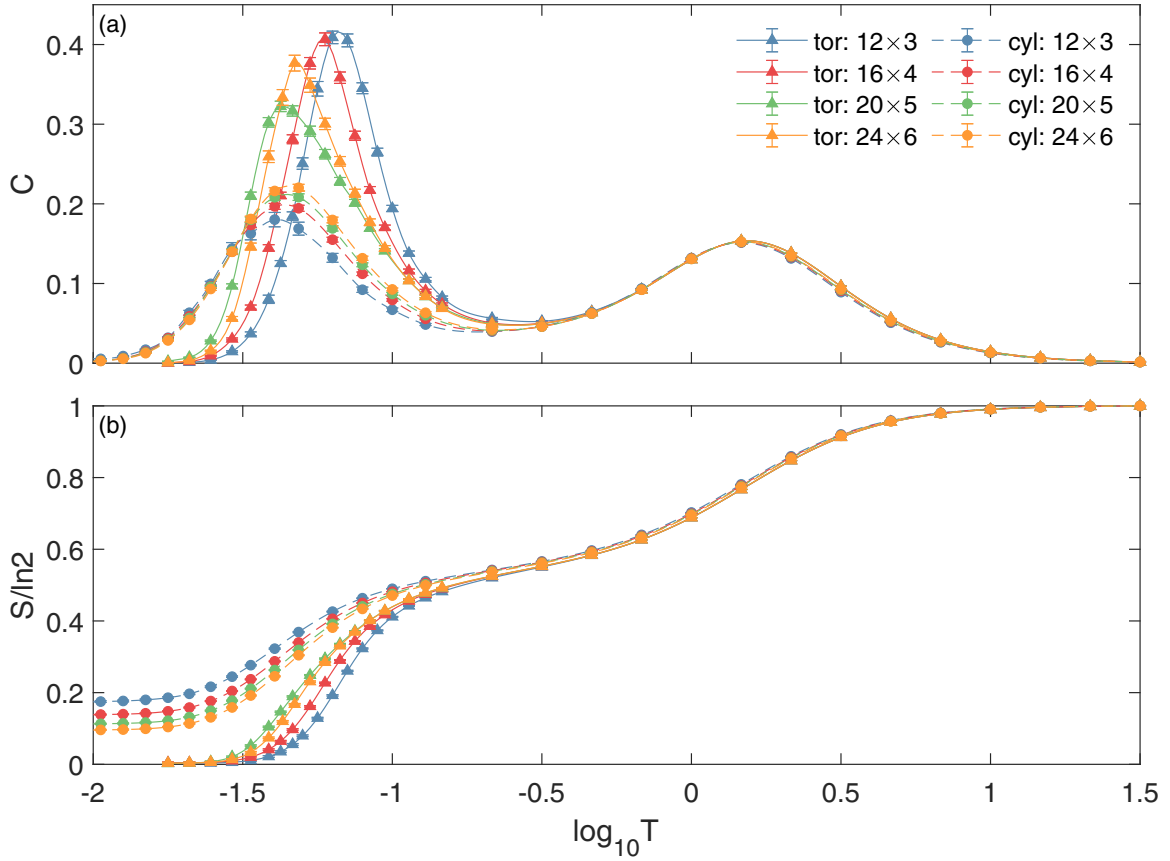
FIG. 7. Specific heat (a) and entropy (b) on torus and cylinder lattices with different $N_1 \times N_2$ sizes. For cylinder lattices, $\mathbf{n}_1$ is periodic direction. We choose the aspect ratio to be 4 to have longer open edges and more edge modes.

contribution $N_1 \ln 2$ due to dangling fermions, to match the residual entropy of the spin model. As expected we find excellent agreement between $S_0^{\text{spin}}(\text{II})$ and $S_0^{\text{spin}}(\text{III})$, even when finite size effects are significant. This illustrates how, in principle, a measurement of the residual entropy S_0^{spin} for different sample geometries can be used to establish the presence of the topological boundary modes characteristic of the Kitaev spin liquid.

We emphasize that because the boundary flat band is topologically protected provided that time-reversal symmetry is

TABLE III. Residual entropy per site divided by $\ln 2$, i.e., $s_0^{\text{spin}} = S_0^{\text{spin}}/(N_s \ln 2)$, computed for cylindrical lattices with open boundaries in the $\mathbf{n}_2$ direction. $S_0^{\text{spin}}(\text{I}) = \frac{4}{3}N_1 \ln 2$ is obtained by neglecting finite size effects and assuming all states in the lowest subband on the interval $k_1 \in [2\pi/3, 4\pi/3]$ contribute to the residual entropy. $S_0^{\text{spin}}(\text{II})$ is obtained from Eq. (22) by numerically counting the number Ω_0 of fermionic states below the energy scale of ϵ_{min}. $S_0^{\text{spin}}(\text{III})$ is obtained by integrating the specific heat from high temperature T_{max} down to T_{min}, according to Eq. (17).

$N_1 \times N_2$	$s_0^{\text{spin}}(\text{I})$	$s_0^{\text{spin}}(\text{II})$	$s_0^{\text{spin}}(\text{III})$
12×3	0.222	0.176	0.174 ± 0.002
16×4	0.167	0.136	0.138 ± 0.001
20×5	0.133	0.112	0.113 ± 0.001
24×6	0.111	0.095	0.096 ± 0.002

unbroken, the residual low-temperature entropy is a feature of the entire gapless Kitaev spin liquid phase and thus is a robust experimental signature of the gapless Kitaev spin liquid. Reference [67] verified this explicitly, showing that the residual entropy survives in the presence of non-Kitaev interactions that are finite but not too strong.

V. FLUX ENERGETICS IN THE TIME-REVERSAL SYMMETRIC KITAEV MODEL ($\kappa = 0$)

The MC results discussed in the previous section reveal significant differences in the low-temperature specific heat peak associated with flux freezeout on the cylinder relative to the torus. In this section, we give a physical interpretation for these differences, by studying the energetics of flux excitations in both torus and cylinder lattices. We use this to model how flux energetics impact thermodynamic quantities in both geometries, for both finite and infinite systems. Here we will focus on the time-reversal symmetric Kitaev model, $\kappa = 0$, and consider the effects of finite $\kappa \neq 0$ in Sec. VI.

A. One-flux gap energy in torus and cylinder

We start by computing the one-flux energy gap Δ_ϕ in toroidal and cylindrical geometries. On the torus, the one-flux energy gap was analyzed by Kitaev [1], who associated it with a half of the two-flux energy $\Delta_\phi = \frac{1}{2}E_{2\phi}^{(0)}$, where $E_{2\phi}^{(0)}$ is defined to be the energy of the flux configuration $\{\phi_p\} = 2\phi$
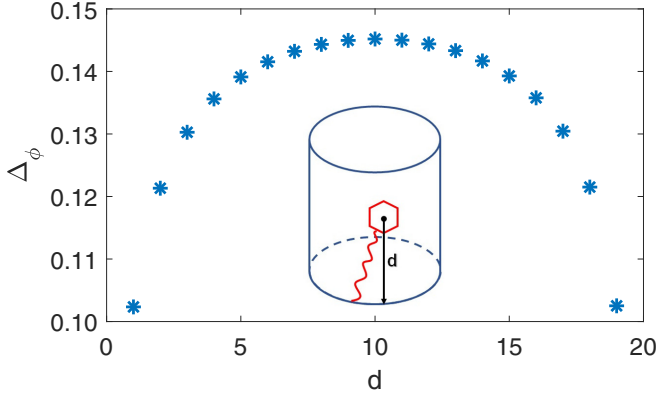
FIG. 8. Position dependence of the one-flux gap on a $N_1 \times N_2 = 40 \times 20$ cylinder. Here d indicates the distance between the flux and the edge.

with the maximum separation between the two fluxes [see Eq. (B1)]. He also showed that $E_{2\phi}^{(0)}$ displays oscillatory period-3 modulations as the lattice size increases, due to the fact that on a $L \times L$ lattice, for L divisible by 3, the Dirac cone is located on the reciprocal lattice point, while for other values of L it is not (see Appendix D). By extrapolating $E_{2\phi}^{(0)}$ to the thermodynamic limit, Kitaev obtained an estimate for the flux gap $\Delta_\phi^\infty = 0.154$.

We follow the same method and compute the one-flux gap in cylinder lattices. In this case, however, the open boundary allows us to create just one flux and study how the one-flux energy depends on the distance of the flux to the open boundary. Figure 8 shows the one-flux energy gap at different distances from the boundary computed on a $N_1 \times N_2 = 40 \times 20$ cylinder. We can clearly see that the flux gap decreases as the flux approaches the open boundary and is significantly lower (by over 30%) on the boundary than deep in the bulk. This decrease in energy can be seen as an effective attractive interaction between the fluxes and the boundary, due to the Majorana boundary modes [68]. To see the finite size scaling of the one-flux gap, in Fig. 9 we plot the flux gap when it is created on the plaquette on the edge (a) and in the middle of the cylinder (b) calculated for different system sizes L. By extrapolation to thermodynamic limit separately for $L = 3k$, $L = 3k + 1$ and $L = 3k + 2$ (see details in Appendix D), we get the edge gap equal to $\Delta_{\phi,e}^\infty = 0.102$ and the bulk gap equal to $\Delta_{\phi,b}^\infty = 0.154$. The latter, as expected, is the same as on the torus lattices.

To summarize, our numerical analysis suggests that, on average, the one-flux gap in cylinder lattices is smaller than that on comparable torus lattices and decreases as the flux approaches the boundary. This finding is in agreement with our MC results of the specific heat (see Fig. 7) which shows that the onset temperature for flux excitations on cylinders are lower than those on tori.

B. Phenomenological flux energy models

To understand the energetics of systems with many fluxes, we analyze the distribution of the flux energy $E_{\phi_p}^{(0)}$ [defined as the lowest-energy state with a given flux configuration ϕ_p—see Eq. (B1)] for different flux configurations $\{\phi_p\}$ at
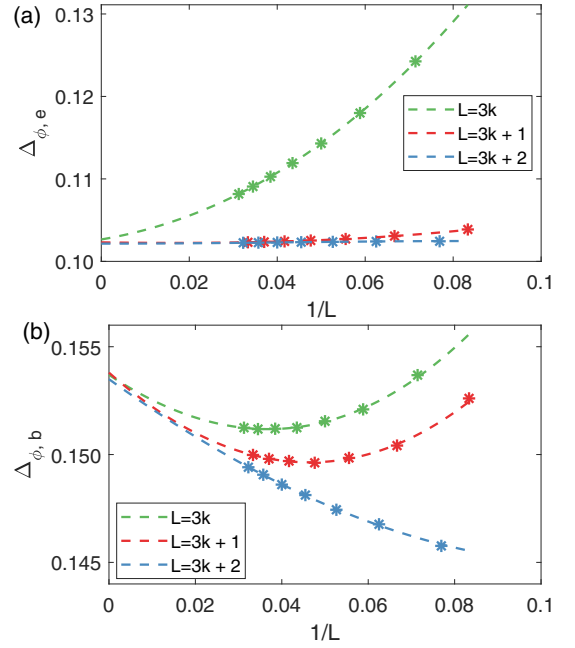

FIG. 9. One-flux energy as a function of the inverse lattice size $1/L$ computed when the flux is located (a) on a boundary plaquette and (b) on a plaquette maximally distant from the edge, for different cylinder systems with $N_1 = N_2 = L$. Due to the periodic dependence on system size modulo 3, we use separate polynomial fits to the series $L = 3k$ (green), $L = 3k + 1$ (red), and $L = 3k + 2$ (blue), and extrapolate these in the limit $1/L \to 0$ to estimate the gap for large system sizes. We find that the energy gaps for the edge and bulk fluxes extrapolate to (a) $\Delta_{\phi,e}^\infty = 0.102$ and (b) $\Delta_{\phi,b}^\infty = 0.154$, respectively.

fixed flux density on torus and cylinder lattices. We perform a polynomial fit to the average energy obtained over many random flux configurations at each flux density. We call this fit the flux pseudopotential energy (PPE). By comparing the resulting curves on lattices of different sizes and geometries, we argue that in the limit of large system size, the PPE on the torus describes a universal dependence of flux energy density on flux density. On the cylinder the dependence on geometry is not universal, but we find evidence of a universal dependence of total flux energy to total flux density on cylinders of fixed width in the thermodynamic limit.

1. Flux pseudopotential energy on a torus

For a given flux configuration, the flux energy density is given by $E_{\phi_p}^{(0)}/N_p$, where N_p is the number of plaquettes and N_ϕ is the number of fluxes. For fixed flux density $n_\phi = N_\phi/N_p$, we define $\mathcal{E}(n_\phi)$ to be the distribution of $E_{\phi_p}^{(0)}/N_p$ for different flux configurations with the same flux number N_ϕ. We begin by studying how $\mathcal{E}(n_\phi)$ depends on n_ϕ for torus lattices. An example is shown in Fig. 10(a), which plots $\mathcal{E}(n_\phi)$ as a function of n_ϕ on the $N_1 \times N_2 = 24 \times 6$ torus lattice. To generate this distribution we sample 100 random flux configurations at each flux density n_ϕ; each blue dot represents the energy of one such configuration. The solid red line represents a fifth-order polynomial best fit to the average $\bar{\mathcal{E}}(n_\phi)$ at each flux density n_ϕ, given by $\bar{\mathcal{E}}(n_\phi) \simeq -0.143n_\phi^5 + 0.274n_\phi^4 - 0.181n_\phi^3 - $
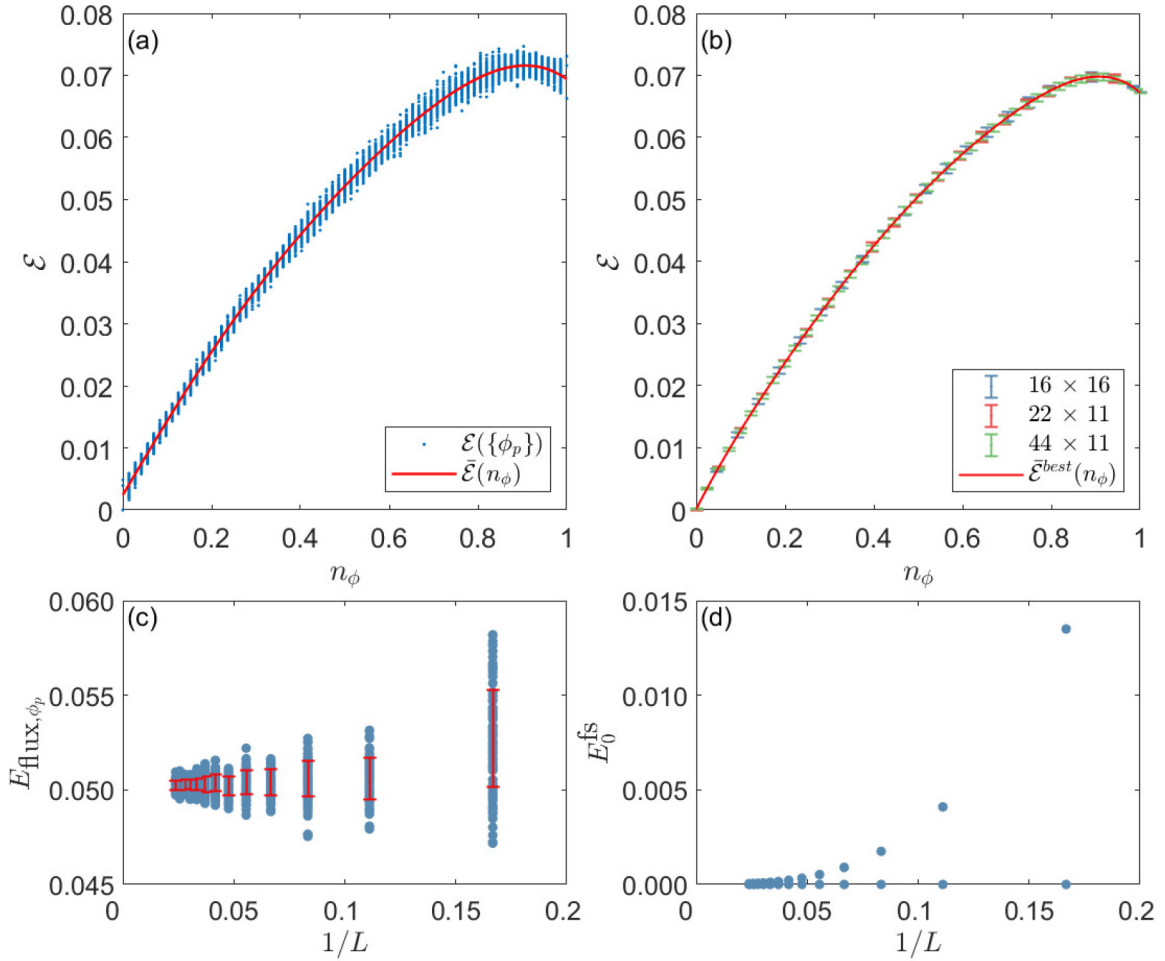
FIG. 10. (a) Distribution of the flux energy density $\mathcal{E}(n_\phi)$ for different flux densities n_ϕ on the torus lattices with $N_1 \times N_2 = 24 \times 6$, generated by taking 100 randomly generated flux configurations at each flux density. The best-fit polynomial $\bar{\mathcal{E}}(n_\phi) \simeq -0.143n_\phi^5 + 0.274n_\phi^4 - 0.181n_\phi^3 - 0.004n_\phi^2 + 0.121n_\phi + 0.002$ is shown by the red curve. The standard deviation of the residual errors is 0.14. (b) The flux energy densities on the lattices of $N_1 \times N_2 = 16 \times 16$ (blue), 22×11 (green), and 44×11 (yellow) from a random sampling of 60 flux configurations. Vertical bars at each flux number show one standard deviation of energy. The best-fit polynomial that simultaneously fits the three sets of data is $\bar{\mathcal{E}}^{best}(n_\phi) \simeq -0.321n_\phi^6 + 0.850n_\phi^5 - 0.898n_\phi^4 + 0.475n_\phi^3 - 0.180n_\phi^2 + 0.141n_\phi$. (c) Distribution of flux energies for various lattices sizes L for a fixed flux density $n_\phi = 0.5$ from a sampling of 60 flux configurations. The red error bars show one standard deviation of the flux energies. (d) The finite-size energy splitting of the ground state energy E_0^{fs} as a function of the inverse lattice size $1/L$, where $N_1 = N_2 = L$.

$0.004n_\phi^2 + 0.121n_\phi + 0.002$. We call this polynomial fit the flux pseudopotential energy (PPE). We choose to fit to fifth-order polynomials as this is the lowest order that consistently fits the data, based on the distribution of residual errors.

The PPE curves for smaller tori, where the finite-size splitting of topological sectors is large, are not universal and can differ significantly depending on the system size and aspect ratio. To understand why, recall that in the absence of fluxes the Kitaev model has four topologically distinct sectors on the torus [1]. These sectors are energetically degenerate in the thermodynamic limit; however on finite-size tori the splitting between them can be appreciable, as is evident from the distribution of energies for the four configurations with $n_\phi = 0$ in Fig. 10(a). This splitting is shown for $L \times L$ tori in Fig. 10(d), as a function of $1/L$. For comparison, Fig. 10(c) shows how the width of the distribution $\mathcal{E}(n_\phi = 0.5)$ depends on the system size L for an $L \times L$ torus. We see that for $L \lesssim 20$, the differences between the energies of these four sectors rapidly

becomes comparable to the width of the distribution of flux energies at $n_\phi = 0.5$, suggesting that the difference between topological sectors accounts for a significant part of the finite-size broadening of this distribution at small L.

For $L \gtrsim 20$, in contrast, differences between the energies of these four sectors are small compared to the width of the distribution of flux energies. Though the distribution narrows with increasing L in this range, it does so only very slowly. Indeed we expect this distribution to reach a plateau of finite width as $L \to \infty$, since the interactions between fluxes lead to variations in the energies of different flux configurations at fixed n_ϕ. Our numerics suggest that for $L \gtrsim 20$ the distribution of energies converges to a universal function, independent of both system size and aspect ratio, which captures the main features of the flux energetics as a function of flux density in the thermodynamic limit. In order to find this universal function, we consider tori with $N_1 \times N_2 = 16 \times 16$, 22×11, and 44×11 (for which the aspect ratios N_1/N_2 are 1, 2, and 4,

respectively) and for each of the systems obtain $\mathcal{E}(n_\phi)$ [shown in Fig. 10(b)] generated from 60 random flux configurations. The height of the bars at each value of n_ϕ represents the standard deviation of the corresponding energy distribution. We see that the overlap between the three curves is large, indicating that the dependence on aspect ratio (as well as system size, not shown here) is negligible. All three curves are well described by the best-fit polynomial (shown by the red line)

$$\bar{\mathcal{E}}^{\text{best}}(n_\phi) \simeq -0.321n_\phi^6 + 0.850n_\phi^5 - 0.898n_\phi^4$$
$$+ 0.475n_\phi^3 - 0.180n_\phi^2 + 0.141n_\phi . \quad (23)$$

We find that the sixth-order polynomial gives better fit for the distribution of energies shown in Fig. 10(b) than the fifth-order polynomial. We thus use the functional form (23) as an approximate universal PPE fit, which describes the average flux energy density in the thermodynamic limit.

The PPE polynomial Eq. (23) gives us a useful tool to obtain insights into the nature of the flux interactions on the torus. First, the concave shape of the flux PPE as a function of the flux density indicates that the flux interactions are, on average, attractive; our analysis suggests this remains true at all flux densities. Second, it is natural to try to identify the coefficient of the linear term in the best-fit curve with an average flux gap $\bar{\Delta}_\phi$, since on a large torus $2\bar{\Delta}_\phi$ should give a good approximation of the average energy of a configuration of two fluxes. The value obtained from the universal PPE in Eq. (23) is 0.14, which is smaller than $\Delta_\phi^\infty = 0.154$. We attribute this to the attractive interactions between fluxes: Δ_ϕ^∞ is measured by putting the two fluxes at the maximum separation, while $\bar{\Delta}_\phi$ is obtained from the 2-flux energy averaged over all separations. As interactions between the fluxes are attractive on average, the latter is smaller.

2. Flux pseudopotential energy on a cylinder

In order to study how the flux energies are distributed on the cylinder lattices, we again sample different flux configurations for fixed flux densities. However, since the energy of a single flux significantly depends on its proximity to the edge of the lattice (see Fig. 8), we now do this for a pair of two independent flux densities: one for the edge, $n_\phi^e = N_\phi^e/N_p^e$, and one for the bulk, $n_\phi^b = N_\phi^b/N_p^b$. Here N_ϕ^e is the edge flux number, counting the fluxes on the two outermost layers of plaquettes, and $N_\phi^b = N_\phi - N_\phi^e$ is the bulk flux number. N_p^b and N_p^e count the number of plaquettes on the edges and in the bulk, respectively. The energy density is defined as $\mathcal{E}(n_\phi^e, n_\phi^b) = E_{\phi_p}^{(0)}(N_\phi^e, N_\phi^b)/N_p$, where $E_{\phi_p}^{(0)}(N_\phi^e, N_\phi^b)$ is the flux energy of the configuration with N_ϕ^e edge fluxes and N_ϕ^b bulk fluxes.

On the cylinder, we do not expect the flux PPE to be universal even for large system sizes, since it depends on the fraction of plaquettes that lie on the boundary. However, for fixed N_2 (and hence a fixed boundary-to-area ratio), we do find good agreement to a universal PPE, as shown in Fig. 11. The distributions of flux energy densities $\mathcal{E}(n_\phi^e, n_\phi^b)$ computed on cylindrical lattices with $N_1 \times N_2 = 28 \times 6$, 24×6, 20×6, and 16×6 are shown in Fig. 11(a), with 30 different flux configurations sampled randomly for each pair (n_ϕ^e, n_ϕ^b).
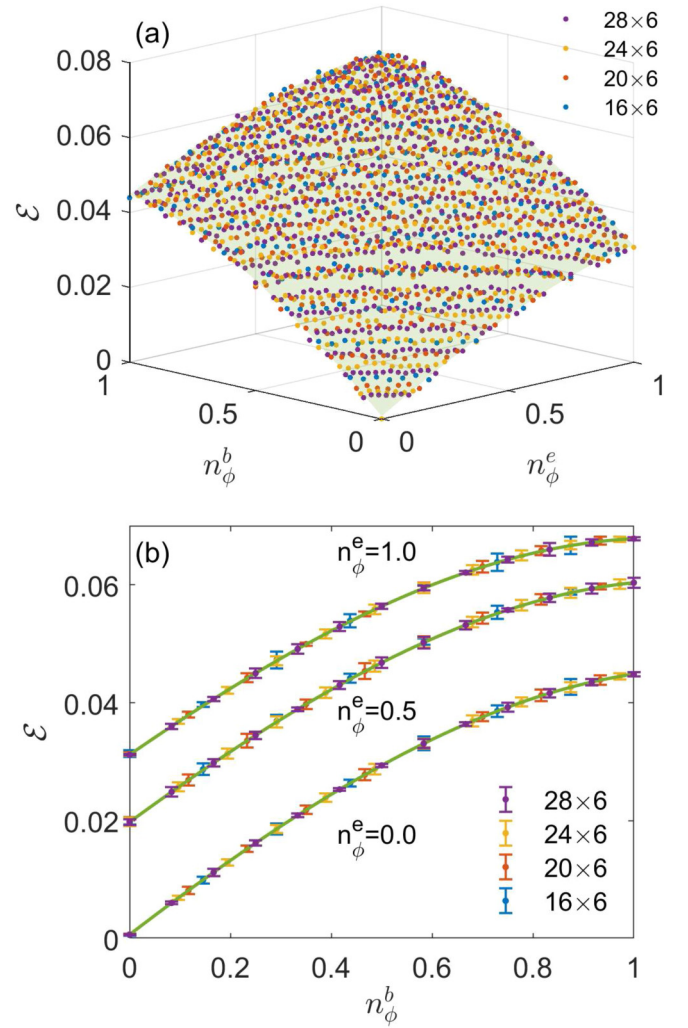

FIG. 11. (a) Distribution of the flux energy density $\mathcal{E}$ as a function of edge ($n_\phi^e = N_\phi^e/N_p^e$) and bulk ($n_\phi^b = N_\phi^b/N_p^b$) flux densities computed on cylindrical lattices with $N_1 \times N_2 = 28 \times 6$ (purple), 24×6 (yellow), 20×6 (orange), and 16×6 (blue). The best-fit polynomial surface is shown. (b) The line cuts of the best-fit surface obtained for fixed edge flux densities $n_\phi^e = 0, 0.5, 1.0$. Vertical bars indicate one standard deviation in energies.

As for the torus lattices, due to the attraction (on average) between fluxes, the average energy density $\bar{\mathcal{E}}(n_\phi^e, n_\phi^b)$ lies on a concave-down surface.

To obtain the flux PPE surface appropriate to cylinders with $N_2 = 6$, we perform a two-variable polynomial fit to the average of all of the energy distributions shown in Fig. 11(a). This gives the best-fit polynomial surface shown in Fig. 11(a), described by the equation:

$$\bar{\mathcal{E}}^{\text{best}}(n_\phi^e, n_\phi^b) = 0.001 + 0.042n_\phi^e + 0.066n_\phi^b$$
$$- 0.005(n_\phi^e)^2 - 0.006n_\phi^e n_\phi^b - 0.011(n_\phi^b)^2$$
$$- 0.007(n_\phi^e)^3 - 0.001(n_\phi^e)^2 n_\phi^b - 0.001n_\phi^e(n_\phi^b)^2$$
$$- 0.010(n_\phi^b)^3 . \quad (24)$$

Using this polynomial, we can also determine the edge and bulk flux gaps as

$$\Delta_{\phi,e} = \frac{\partial \bar{\mathcal{E}}^{\text{best}}(n_\phi^e, n_\phi^b)}{\partial n_\phi^e}\bigg|_{n_\phi^e, n_\phi^b = 0} \cdot \frac{N_p}{N_p^e} = 0.105, \qquad (25)$$

$$\Delta_{\phi,b} = \frac{\partial \bar{\mathcal{E}}^{\text{best}}(n_\phi^e, n_\phi^b)}{\partial n_\phi^b}\bigg|_{n_\phi^e, n_\phi^b = 0} \cdot \frac{N_p}{N_p^b} = 0.110. \qquad (26)$$

While the estimate of the edge flux gap is close to the extrapolated result shown in Fig. 9(a), the bulk gap energy is lower than the energy obtained from the extrapolation shown in Fig. 9(b). As for the torus lattice, one contribution to this difference is that the two-flux interactions are attractive on average. In addition, however, for the small N_2 shown here many of our "bulk" fluxes are in fact quite close to the edge, leading to a significant further reduction.

The line cuts of the best-fit surface obtained for fixed edge flux densities $n_\phi^e = 0, 0.5, 1.0$ are shown in Fig. 11(b), where the height of the bars at each value of n_ϕ^b represents the standard deviation of the corresponding energy distribution for a given value of n_ϕ^e. All system sizes shown are well fit within error bars by the PPE given in Eq. (24).

C. Flux contribution to specific heat ($\kappa = 0$)

We now leverage our phenomenological flux models to predict the contribution of fluxes to the specific heat in the thermodynamic limit of the time-reversal invariant Kitaev model. In particular, since the low-temperature peak in the specific heat is mainly due to the flux excitations, its position and shape are closely related to the flux energetics, which we can now describe in the thermodynamic limit using our numerically obtained best-fit polynomials. We will also compare the predictions of the specific heat based on the PPE polynomials with the results of the MC simulations in analogous boundary conditions and analyze how well they capture the position and shape of the low-temperature peak in the specific heat.

1. Torus lattices

We calculate thermodynamic quantities on the torus lattices by a saddle point approximation of the effective free energy of the flux degrees of freedom,

$$F_\phi = N_p \mathcal{E}(n_\phi) - T S_\phi, \qquad (27)$$

where $\mathcal{E}(n_\phi)$ from now on denotes the average flux energy density and S_ϕ is the flux entropy. On a torus with N_p plaquettes, there are $\binom{N_p}{N_\phi} = \frac{N_p!}{N_\phi!(N_p - N_\phi)!}$ possible configurations with a total of N_ϕ fluxes. Hence the configurational entropy is equal to $S_\phi = \ln\binom{N_p}{N_\phi}$. Using the Stirling's formula, for large N_p and N_ϕ we find:

$$S_\phi = N_p[-n_\phi \ln n_\phi - (1 - n_\phi) \ln(1 - n_\phi)]. \qquad (28)$$

In the thermodynamic limit, the value of the flux density n_ϕ can be obtained by minimizing the free energy:

$$\frac{1}{N_p} \frac{dF_\phi}{dn_\phi} = \mathcal{E}'(n_\phi) - T \ln \frac{1 - n_\phi}{n_\phi} = 0, \qquad (29)$$

where $\mathcal{E}'(n_\phi) = \partial \mathcal{E}(n_\phi)/\partial n_\phi$. This yields the flux density:

$$n_\phi(T) = \frac{1}{e^{\mathcal{E}'(n_\phi)/T} + 1}. \qquad (30)$$

Since the energy density depends on T only through its dependence on n_ϕ, the specific heat per site arising from solely flux degrees of freedom can be expressed as:

$$C_\phi = \frac{1}{2} \mathcal{E}'(n_\phi) \frac{dn_\phi}{dT}. \qquad (31)$$

Explicitly,

$$\frac{\partial n_\phi(T)}{\partial T} = -\frac{e^{\mathcal{E}'(n_\phi)/T}}{(e^{\mathcal{E}'(n_\phi)/T} + 1)^2}\left(-\frac{\mathcal{E}'(n_\phi)}{T^2} + \frac{1}{T}\frac{\partial \mathcal{E}'(n_\phi)}{\partial T}\right)$$

$$= n_\phi(1 - n_\phi)\left(\frac{\mathcal{E}'(n_\phi)}{T^2} - \frac{1}{T}\frac{\partial \mathcal{E}'(n_\phi)}{\partial n_\phi}\frac{\partial n_\phi(T)}{\partial T}\right). \qquad (32)$$

Solving the above for $\partial n_\phi(T)/\partial T$ gives the specific heat per site:

$$C_\phi = \frac{1}{2}\left(\frac{\mathcal{E}'(n_\phi)}{T}\right)^2 \frac{1}{\frac{1}{n_\phi(1-n_\phi)} + \frac{\mathcal{E}''(n_\phi)}{T}}, \qquad (33)$$

where $\mathcal{E}''(n_\phi) = \partial^2 \mathcal{E}(n_\phi)/\partial n_\phi^2$. The denominator of Eq. (33) shows two distinct contributions to the specific heat. The first, from the term $\frac{1}{n_\phi(1-n_\phi)}$, is entropic in origin and comes from the curvature of the configurational entropy, $S_\phi''/N_p = -\frac{1}{n_\phi(1-n_\phi)}$.

The second contribution is due to the curvature of the average flux energy density $\mathcal{E}(n_\phi)$. For an attractive interaction, such as the one we found for our flux PPE in Sec. VI A, $\mathcal{E}''(n_\phi) < 0$, and this term increases the value of the specific heat. Intuitively, this is because the attractive interaction between the fluxes lowers the energy required to excite additional fluxes as the flux density increases, increasing the number of thermal flux excitations at a given temperature.

We evaluate the flux contribution to the specific heat by using the best-fit polynomial (23) to describe the average flux energy density $\mathcal{E}(n_\phi)$ in Eq. (33). This gives us the specific heat of the interacting flux model (IFT) on the torus. For comparison, we also compute the specific heat Eq. (33) for the noninteracting flux model on the torus (NIFT), for which the energy density is simply $\mathcal{E}(n_\phi) = \Delta_\phi n_\phi$, with $\Delta_\phi = \overline{\Delta}_\phi$, and $\mathcal{E}''(n_\phi) = 0$. In Fig. 12(a) we plot the specific heat for both the IFT (dark blue) and NIFT (light blue) models. Clearly, flux interactions have an important effect on the shape of the peak in specific heat: The IFT model has a higher and narrower peak than the NIFT model due to the attractive flux interactions. The difference is on the order of 30%, indicating that interaction effects are quantitatively important. We also find that the precise shape and height of the specific heat curve is sensitive to small changes in the flux PPE, resulting in significant error bars to our predictions for C in spite of the relatively small error of our PPE fits. (See Appendix F for details.)

Figure 12(a) also compares these flux-only specific heat curves $C_\phi(T)$ to those obtained by the MC simulations on different-size torus lattices, which describe both flux and
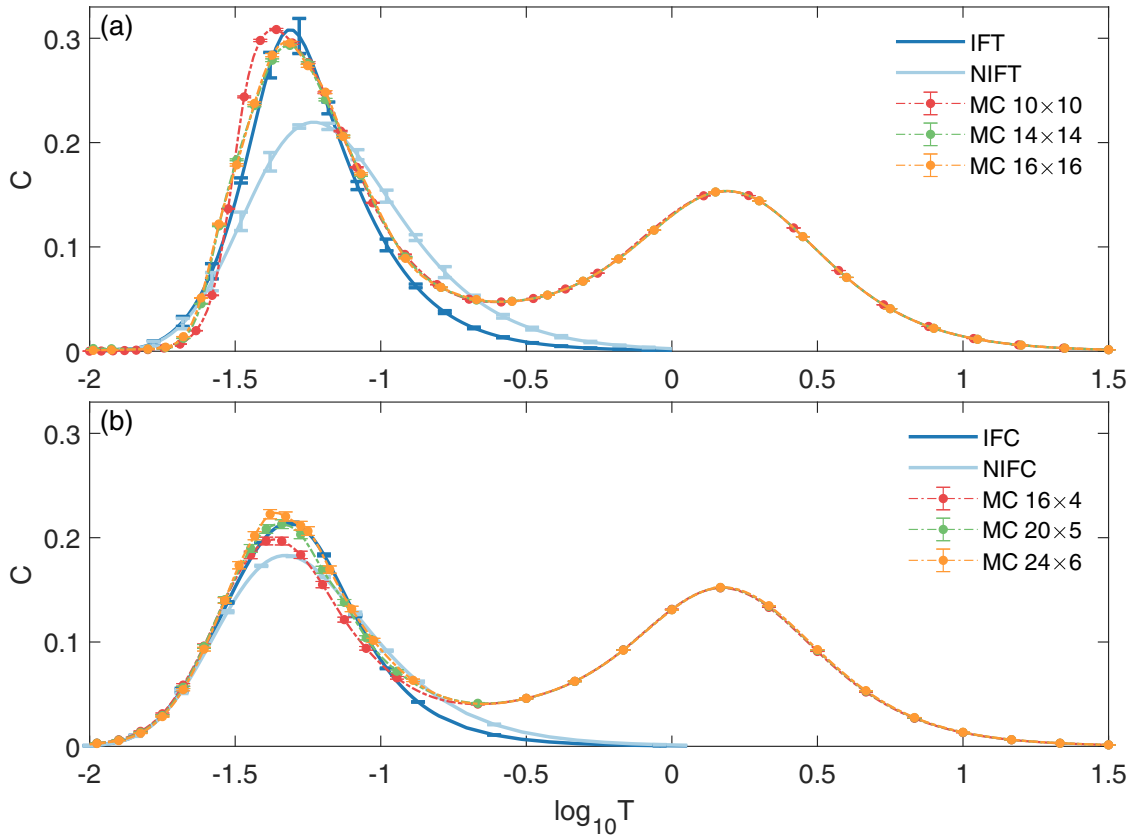
FIG. 12. Comparison between IFT, NIFT, and MC predictions of the specific heat. (a) The specific heat per site with IFT and NIFT models, computed according to Eq. (33), compared to Monte Carlo simulations on various $L \times L$ tori. (b) The specific heat per site with IFC and NIFC models, computed according to Eq. (37), compared to Monte Carlo simulations. Note that our NIFC fit is specific to cylinders with $N_2 = 6$. Here the uncertainty in the specific heat of the PPE models are estimated from an ensemble of fit curves created by bootstrapping the regression data [69], as described in Appendix F.

fermionic contributions to the specific heat $C(T)$. We find reasonable agreement between the specific heat of the IFT model and that of the MC simulations over most of the lower peak in $C(T)$, particularly for the largest lattice sizes simulated. The two models do exhibit significant differences at low temperatures, with the specific heat falling off more quickly in the MC simulations than in our IFT model. We conjecture that this is primarily due to finite-size effects, which decrease the average distance (and thus increase the average interaction) between fluxes at low flux numbers in our MC simulations. Thus, as expected, we find that the low-temperature peak in $C(T)$ can be well accounted for by a model that includes only flux excitations and ignores the dispersing fermions.

Moreover, the IFT model based on the PPE polynomial fit suggests that the low-temperature peak in specific heat is of finite width even in the thermodynamic limit. This confirms the expectation [37] that the gapless phase of the 2D Kitaev model does not have a finite-temperature phase transition, but rather a crossover from a state of vanishingly small flux density at temperatures well below the flux gap, to one with flux density $n_\phi \approx 0.5$ at temperatures above the flux gap. To see why, we expand the exact expression for the free energy at large N_p and small N_ϕ to find $F_\phi \approx (\mathcal{E}'(0) - T \ln N_p)N_\phi$. This is minimized by $N_\phi > 0$ at the flux onset temperature

$$T_{\text{onset}} \approx \bar{\Delta}_\phi / \ln N_p, \tag{34}$$

where $2\bar{\Delta}_\phi$ is the average energy cost of inserting a pair of fluxes. Thus in the thermodynamic limit, for any $T > 0$ there will be some (finite) number of fluxes in the system. Although for $T \ll \bar{\Delta}_\phi$ the flux *density* is effectively 0, these few fluxes are sufficient to destroy the topological order, such that there are no singularities in $C(T)$. Our numerics suggest that including the low-energy fermionic excitations in this analysis will not substantively alter T_{onset}, or the nature of the crossover. We emphasize that T_{onset} does not correspond to the position of the low-temperature peak in $C(T)$, which always occurs at a temperature scale set by $\bar{\Delta}_\phi$, irrespective of the system size. On the finite-sized systems studied here, however, where the flux density can never be lower than $2/N_p$, T_{onset} is a good proxy for where this lower peak begins.

In Fig. 12(a), the specific heat curves of the NIFT and IFT models merge together at low temperatures, where the flux density is small. This is because at low flux densities the energy density of the IFT model is dominated by the linear term, $\mathcal{E}(n_\phi) \approx \mathcal{E}'(0)n_\phi$. Thus taking the flux gap of the NIFC model to be $\bar{\Delta}_\phi = \mathcal{E}'(0)$, the two models exhibit very similar

behavior for $C(T)$ at temperatures well below $\bar{\Delta}_\phi$, where the flux density is small.

2. Cylinder lattices

We now discuss the analogs of both the noninteracting (NIFC) and the interacting (IFC) flux models appropriate for cylinder lattices. With the 2-variable model of the flux energy described in Sec. V B 2, the approach taken above to calculate $C(T)$ becomes cumbersome. Instead, we start from the expression for the specific heat in Eq. (B5) and exploit the fact that the last term vanishes in the absence of fermionic excitations, i.e.,

$$C_\phi(T) = \frac{d\langle E_{\phi_p}^{(0)} \rangle}{dT} = \frac{1}{T^2}\left(\langle (E_{\phi_p}^{(0)})^2 \rangle - \langle E_{\phi_p}^{(0)} \rangle^2\right). \quad (35)$$

Here the expectation $\langle \ldots \rangle$ is taken over different flux configurations characterized by the pairs of the flux numbers (N_ϕ^e, N_ϕ^b), weighted by the probability

$$p(N_\phi^e, N_\phi^b) = e^{-\beta E_{\phi_p}^{(0)} + \ln\binom{N_p^e}{N_\phi^e} + \ln\binom{N_p^b}{N_\phi^b}}. \quad (36)$$

On a finite-size cylinder, this allows us to compute the specific heat exactly by summing over the possible values of N_ϕ^e, N_ϕ^b:

$$C_\phi(T) = \frac{1}{N_s T^2} \sum_{N_\phi^e, N_\phi^b} \left(E_{\phi_p}^{(0)} - \langle E_{\phi_p}^{(0)} \rangle\right)^2 p(N_\phi^e, N_\phi^b), \quad (37)$$

where the sum over N_ϕ^e (N_ϕ^b) runs from 1 to N_p^e (1 to N_p^b), where $N_p^e(N_p^b)$ is the number of plaquettes on the edge (in the bulk). Using Eq. (37), we compute the specific heat using both the best-fit PPE surface for the cylinders (IFC model) and the noninteracting flux model (NIFC), in which $E_{\phi_p}^{(0)} = N_\phi^e \Delta_{\phi,e} + N_\phi^b \Delta_{\phi,b}$, with the gaps $\Delta_{\phi,e}$ and $\Delta_{\phi,b}$ given by Eqs. (25) and (26).

Figure 12(b) shows the specific heat of the NIFC and IFC models evaluated on a 24×6 lattice along with the ones obtained by the MC simulations on 16×4, 20×5, and 24×6 lattices. (Recall that the fit used in our IFC model is specific to cylinders with $N_2 = 6$.) Again, we see that flux interactions, which are attractive on average, play an important role: The specific heat in the NIFC model has a lower and broader peak than that in the IFC model, with a difference in peak heights on the order of 20%. We also find good agreement between the specific heat in the IFC and our MC simulations in the vicinity of the low-temperature peak, particularly for the two largest system sizes. Notice that there is still a small disagreement between the IFC and the MC results on a 24×6 lattice: The peak temperature in the IFC model is slightly higher than that obtained with the MC simulation. This is because the IFC model assigns the same energy to any flux not on the edge. Though this leads to an underestimate for the gap of fluxes deep in the bulk of the system, it also overestimates the gap of fluxes close to, but not on, the boundary (see Fig. 8). Since these lower energy excitations contribute disproportionately to the specific heat at lower temperatures, our IFC model actually overestimates the peak position. Still, Fig. 12(b) shows that the IFC model captures the specific heat due to thermal flux excitations relatively well.
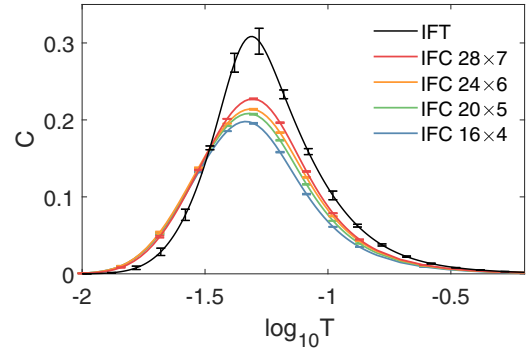

FIG. 13. Comparison of the specific heat computed using an IFC model based on fits for various lattice sizes, with the specific heat computed using the IFT model in the thermodynamic limit. Error estimates are obtained using the same method as in Fig. 12.

3. Comparison between the IFT and IFC specific heats

To illustrate the effect of the boundary fluxes on the thermodynamics, here we directly compare the specific heat on torus and cylinder lattices. Figure 13 presents the specific heat of the IFC model computed on a series of finite-size cylinder lattices and that of the universal IFT model on the torus lattice. It shows that the peak begins at a lower temperature on the cylinder than on the torus, predominantly due to the smaller single-flux gap near the edge of the cylinder, which broadens the temperature range over which the flux entropy is released. Because the total entropy released is the same in both cases, the peak height for the cylinder lattice is correspondingly lower than that on the torus lattice.

Additionally, as the lattice size increases, the specific heat peak of the cylinders moves slightly towards the right, with the right half of the peak appearing to tend towards the IFT curve. This is because for an $aL \times L$ lattice ($a > 1$), as the lattice size L increases, the fraction of the plaquettes on the boundary falls off as $N_p^e/N_p = a/L$. Thus any effects associated with the edge diminish in importance in the thermodynamic limit.

VI. FLUX ENERGETICS IN THE EXTENDED KITAEV MODEL ($\kappa \neq 0$)

A. Flux pseudopotential energy on a torus at $\kappa \neq 0$

We now turn to the time-reversal symmetry broken Kitaev model and study the energetics of the many-flux problem at finite κ. As we discussed in Sec. III C, the κ term introduces both an energy dispersion to the edge modes and a bulk energy band gap. Collectively, these lead to a slightly weaker dependence of the single flux energy on its proximity to the edge of the lattice; however, the main effects of the boundary on the flux energetics remain the same as in the time-reversal invariant case. Here we therefore only consider the flux energetics on torus lattices, for which we numerically obtain best-fit PPE polynomials to the average flux energy as a function of the flux density n_ϕ and the strength of κ [see Fig. 14(a)]. Note that we treat the three-spin interaction as an independent time-reversal symmetry breaking term rather than as a perturbative effect of the magnetic field [1], and thus for curiosity consider values of κ up to 1.
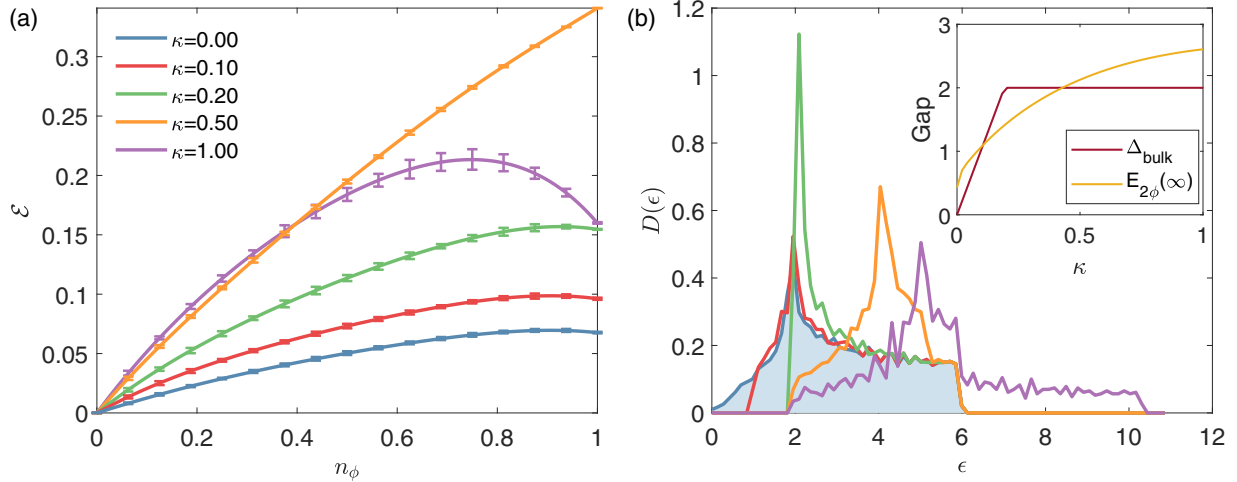
FIG. 14. (a) The flux PPE for different κ on the $N_1 \times N_2 = 16 \times 16$ torus. The corresponding best-fit polynomials are given in Appendix E. For $\kappa = 0.5$, the nearly straight band and small standard deviation indicates that flux interactions are very weak. (b) The fermionic DOS, $D(\epsilon)$, in the flux-free sector obtained by numerically diagonalizing the Bloch Hamiltonian $H_\mathbf{k}$ defined in Eq. (3) on the $N_1 \times N_2 = 30 \times 30$ torus. The inset in (b) shows the two-flux gap $E_{2\phi}(\infty)$ (dark yellow line) and the fermionic gap Δ_{bulk} (dark red line). Here $E_{2\phi}(\infty)$ is defined as the flux energy when two fluxes are placed at the maximum separation.

The value of κ affects the flux energetics through its impact on the fermionic spectrum. In Fig. 14(b) we plot the fermionic DOS in the flux-free sector for various values of κ. Recall that in the extended Kitaev model, Majorana fermions can hop to nearest neighbor sites with an amplitude of J and to second neighbor sites with an amplitude of κ. Two main effects of finite κ should be noticed: First, as expected, κ leads to a bulk gap Δ_{bulk} in the fermionic spectrum, which grows linearly with κ up to $\kappa \simeq 0.2$. The inset in (b) shows that both this fermionic energy gap (dark red line) and the flux gap $E_{2\phi}(\infty)$ (dark yellow line, defined as the flux energy with a pair of maximally separated fluxes on the torus), grow monotonically with κ. Second, for $\kappa > 1/\sqrt{3}$, the maximum fermionic energy is greater than its $\kappa = 0$ value of $6J$ and increases linearly with κ, leading to the increased maximum fermionic energy that is apparent for $\kappa = 1$.

Figure 14(a) shows the phenomenological models of flux energetics in the time-reversal symmetry broken case computed on the $N_1 \times N_2 = 16 \times 16$ torus. The PPE polynomials are shown using solid lines, while the vertical bars show the mean and standard deviation of 60 random flux configurations at each flux density, with colors indicating the value of κ in each case. As before, the best-fit polynomials (given explicitly in Appendix E) are obtained by fitting the mean energy as a function of n_ϕ for each value of κ. We see that the slope of the PPE at low flux densities increases monotonically with κ, due to the increase in the flux energy gap. The explicit dependence of $\bar{\Delta}_\phi$ on the strength of κ is shown in Fig. 22 and is consistent with the trend for $E_{2\phi}(\infty)$ shown in the inset of Fig. 14(b). For $\kappa \lesssim 0.5$ this difference in slope is clearly the most significant effect of increasing κ. We also observe differences in the curvature of the PPE for different κ, indicating differences in the average importance of flux interactions, which we will attempt to understand with the minimal two-flux interaction model in Sec. VII.

B. Flux contribution to the specific heat at $\kappa \neq 0$

The change of the flux energetics due to κ is clearly observable in the behavior of the specific heat. This is shown in Fig. 15, which compares the predictions of our flux-PPE model $C_\phi(T)$ with MC results $C(T)$ on 16×16 tori. The solid lines show the specific heat for various values of κ as a function of temperature using the best-fit flux PPE polynomials explicitly given in Eq. (E1) and shown in Fig. 14(a). Comparing these with the results for $C(T)$ obtained with MC simulations on the 16×16 torus shows that the specific heat peak positions are well approximated by the flux PPE model: In Fig. 15 we clearly see a rightward shift in the position of the low-temperature peak with increasing κ in both the flux PPE models and the MC simulation results. This is due to the fact that the flux gap increases with κ.

The peak's height and shape are more sensitive to the details of the PPE fit. In particular, small changes in the values of the polynomial coefficients of the flux PPE can lead to large differences in the peak height and shape; this is apparent when comparing the fit at $\kappa = 0$ in Fig. 15 (obtained from a PPE derived from a single system size) to that in Fig. 12, obtained from the universal best-fit PPE in Eq. (23). As a result, as discussed in detail in Appendix F, the small sampling error in the PPE fits translates to significant error bars near the peak of $C_\phi(T)$. Moreover, the flux PPE model predicts $C_\phi(T)$ based on the average flux energy for a given flux density and does not account for fluctuations in this energy over different configurations with the same flux number, which can also impact the specific shape of the peak. These factors, together with finite-size effects in the MC data, lead to differences between the PPE and MC predictions for the peak height and shape. These are particularly apparent when comparing the PPE results for the specific heat computed for $\kappa = 0.1$ with corresponding MC result.

For $\kappa \gtrsim 0.2$, our MC simulations also show changes in the high-temperature (fermionic) peak in $C(T)$, which shifts
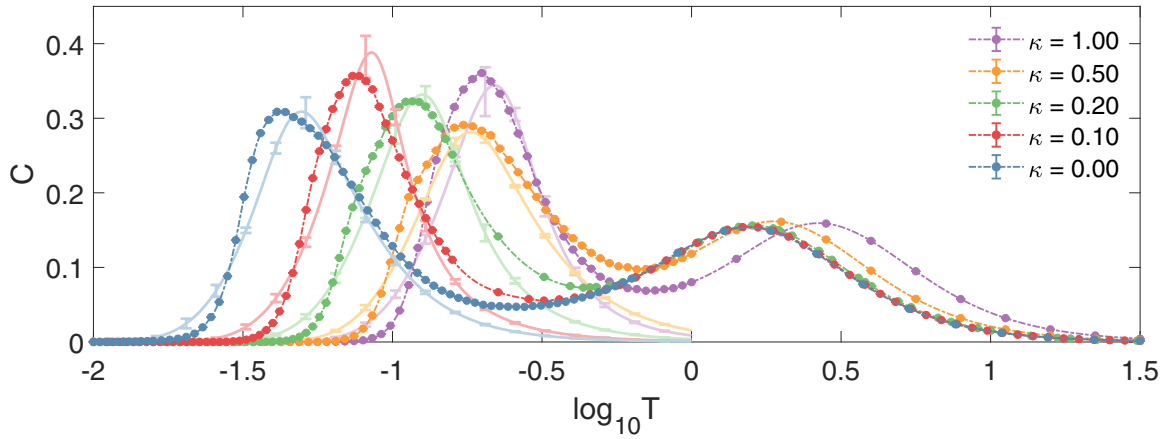
FIG. 15. Comparison of the specific heat per site for various values of κ on a 16×16 torus lattice, computed by the MC simulations (the dots) and from the flux PPE model via Eq. (33) (the solid lines). The uncertainty of the PPE model specific heat is estimated the same way as described in Fig. 12.

to noticeably higher temperatures for $\kappa = 1$. This reflects the changes in the high-energy fermionic DOS for larger values of κ [see Fig. 14(b)].

VII. TWO-FLUX INTERACTIONS

In the previous sections, we presented phenomenological models of the average flux energy $\bar{\mathcal{E}}(n_\phi)$ as a function of flux density n_ϕ that we argued were universal on sufficiently large tori and universal for a fixed value of N_2 on the cylinder. We found that $\bar{\mathcal{E}}(n_\phi)$ is a concave function of n_ϕ, indicating that on average, interactions between the fluxes are attractive. For sufficiently large system sizes, we showed that the width of the distribution of energies around the best-fit PPE curve was narrow, indicating that at given flux density the flux energy depends only weakly on the specific flux configuration. In this section, we examine the microscopics of the flux interactions and argue that the essential features of our flux PPE models can be understood by studying the interaction between just two fluxes. We will also show how these interactions can be tuned by the time-reversal symmetry breaking term κ.

At finite κ, the interaction between fluxes in the Kitaev model originates from the Majorana zero mode localized near each flux [1], leading to flux interactions that are similar to those of vortices in chiral p-wave superconductors [70]. The resulting fermion energetics was studied numerically by Ref. [68], who obtained a phenomenological fit to the Majorana fermion energy as a function of the separation between the two fluxes. Subsequently, similar effective Hamiltonians describing the interactions between Majorana zero modes have been used to study the Majorana zero mode spectrum in the presence of various flux lattices, often called vison crystals [71–74].

Specifically, a pair of Majorana zero modes corresponds to two possible states, which can be characterized by their fermion parity. Thus due to the bound Majorana zero modes, each pair of fluxes can have either even or odd fermion parity. As the separation d between the two fluxes approaches ∞, the energy splitting $2\epsilon_0(d)$ between these states approaches 0, and the two possibilities are energetically degenerate, hence the name "zero mode." At finite separation, however, the wave

functions of the Majorana zero modes hybridize, leading to a nonvanishing value of $\epsilon_0(d)$, which can be quantitatively described by [68]

$$\epsilon_0(d) = \Delta_{\text{bulk}} \cos\left(\frac{2\pi d}{\lambda}\right) e^{-\frac{d}{\xi}}, \qquad (38)$$

where λ is a characteristic wavelength and Δ_{bulk} is the complex fermion bulk energy gap. The coherence length ξ is inversely proportional to the bulk fermion gap, i.e., $\xi \sim \Delta_{\text{bulk}}^{-1}$, indicating that the length scale over which the flux interactions are appreciable decreases with increasing κ. In addition to the exponential decrease with separation, Eq. (38) shows that $\epsilon_0(d)$ oscillates as a function of the separation d.

To show the relationship between Majorana zero modes and the two-flux interactions, we follow Ref. [68] and plot the fermionic energy spectrum on the torus as a function of the two-flux separation d along the $\mathbf{n}_1$ direction, for both $\kappa = 0$ [Fig. 16(a)] and $\kappa = 0.05$ [Fig. 16(b)]. In Fig. 16 we show the full particle-hole symmetric Majorana fermion spectrum; the positive energy corresponding to a complex fermion at momentum $\mathbf{k}$ is given by twice the energy of the upper band at $\mathbf{k}$. Note that throughout this paper, we use the symbol ϵ to denote this complex fermion energy. We identify the two Majorana states with energies closest to 0 (red lines) as those associated with the Majorana zero mode pair, with energy $\pm\frac{1}{2}\epsilon_0(d)$. The two-flux interaction energy $V_{\text{int}}(d)$ is given by the difference between the energy of the flux pair in the lower energy state at separation d [see Eq. (B1)] and the energy of the flux pair at infinite separation:

$$V_{\text{int}}(d) = -\frac{1}{2}\epsilon_0(d) - \frac{1}{2}\sum_{i>0}(\epsilon_i(d) - \epsilon_i(\infty)), \qquad (39)$$

where we have used the fact that $\epsilon_0(\infty) = 0$, and the second sum runs over all positive-energy fermionic states except the lowest-energy one. It is clear that due to the oscillating behavior of $\epsilon_0(d)$, $V_{\text{int}}(d)$ also exhibits short-range modulations.

Figure 16(b) shows that for $\kappa = 0.05$, all fermionic modes except the lowest one (shown in red) are almost independent of d. This is true for all nonzero κ, thus at finite κ the two-flux interaction energy is due almost entirely to the hybridization
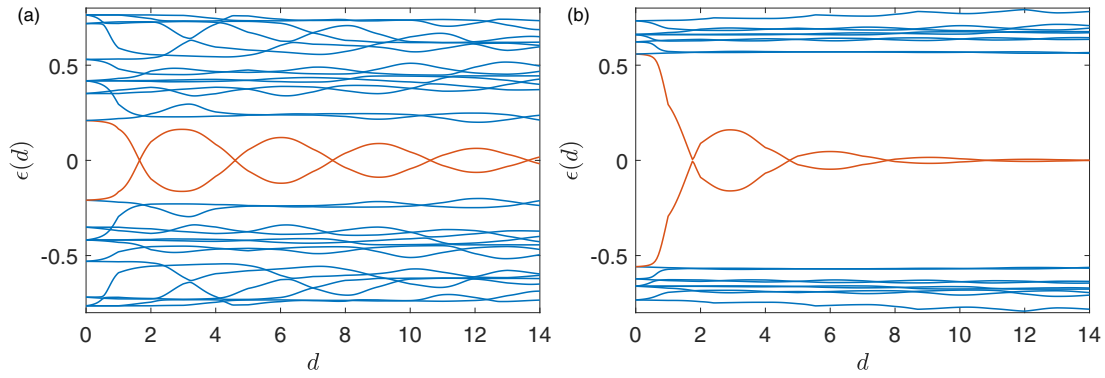
FIG. 16. The fermion energy spectrum on $N_1 \times N_2 = 30 \times 20$ torus as a function of the two-flux separation d along the $\mathbf{n}_1$ direction computed with (a) $\kappa = 0$ and (b) $\kappa = 0.05$. The plot shows the 40 lowest energy bands. The lowest-energy fermionic band is associated with the Majorana zero modes bound to the flux pair, with the corresponding energy $\epsilon_0(d)$ shown by the red line. In both cases $\epsilon_0(d)$ displays decaying oscillations as a function of d. For $\kappa = 0.05$ the rest of the fermionic modes are almost independent of d for $d > 1$; for $\kappa = 0$ all modes exhibit d dependence.

of the Majorana zero modes, which determines $\epsilon_0(d)$. On the other hand at $\kappa = 0$ [see Fig. 16(a)], the entire fermionic spectrum varies with the two-flux separation d, and we can no longer study $V_{\text{int}}(d)$ merely by examining $\epsilon_0(d)$. Moreover at $\kappa = 0$ the expression (38) no longer holds in the thermodynamic limit, where the bulk fermion gap $\Delta_{\text{bulk}} = 0$, and ξ is infinite. In this case we expect a power-law decay to replace the exponential envelope; for finite-size systems, we expect ξ to be on the order of the system size. This is apparent if we compare Figs. 16(a) and 16(b), which show a clear contrast in the length scale over which $\epsilon_0(d)$ falls off. Thus $V_{\text{int}}(d)$ is effectively long ranged for $\kappa = 0$ but short ranged (for sufficiently large system sizes) when $\kappa > 0$.

To demonstrate the evolution of the two-flux interaction with increasing κ, Fig. 17 shows spatial maps of $V_{\text{int}}(d)$ on a 30×20 torus for κ ranging from 0 [panel (a)] to 1 [panel (f)]. For $\kappa = 0$, where the coherence length ξ is comparable to the system size, we see appreciable interactions on virtually all plaquettes in the system. Moreover, they can be either attractive (negative) or repulsive (positive), depending on the relative positions of the two plaquettes. In addition, at the system sizes studied here, the spatial pattern of flux interactions depends strongly on the topological sector. This is clearly seen in Fig. 18, which shows the two-flux interaction maps for the four topological sectors. Here fixing the topological sector effectively determines which of the possible configurations of bonds with $\hat{u}_{ij} = -1$ connect our flux pair (see Appendix C for our conventions); these differences lead to a significant anisotropy in the pattern of spatial interactions, as seen in Fig. 18. As a consequence, on finite-size tori a fixed flux configuration corresponds to multiple distinct energies, depending on the underlying topological sector; this is responsible for a significant fraction of the variance in the energy at a given flux sector observed in Fig. 10(a). We emphasize that the anisotropy in the pattern of spatial interactions seen in Fig. 18 is due to finite-size effects: Choosing a different topological sector in the $\mathbf{n}_1$ or $\mathbf{n}_2$ direction corresponds to interchanging periodic and antiperiodic boundary conditions in $\mathbf{n}_1$ or $\mathbf{n}_2$, respectively, for the itinerant fermions. On small tori, where the fermion momenta are discrete, such changes in boundary conditions can lead to significant changes in the energy. In the

thermodynamic limit, however, these differences vanish with the inverse system size, and the four topological sectors are energetically degenerate. This highlights the fact that finite-size effects due to these topological sectors are present in all of the system sizes shown in Fig. 10(b).

When κ increases, the coherence length ξ decreases, and already for $\kappa = 0.05$, we see that ξ is less than the maximal separation between fluxes on the 30×20 lattice shown in Fig. 17. Consequently, the anisotropy corresponding to different topological sectors for $\kappa = 0$ is no longer present, indicating that the impact of finite-size effects on flux energetics is already small for this system size.

Figure 17 also shows how the two-flux interaction depends on the strength of κ. For κ up to 0.2, we see that ξ decreases with increasing κ, as anticipated above. For larger values of κ [see Figs. 17(d)–17(f)], however, we can see an interesting transition. As was discussed in Ref. [68], for $\kappa \gtrsim 0.2$, the coherence length ξ is on the order of a single plaquette and cannot decrease further with increasing κ; this corresponds to the point at which the perturbative treatment used to obtain $\xi \sim \Delta_{\text{bulk}}^{-1}$ is no longer valid. Instead, by $\kappa \sim 0.5$, we find that the coherence length has begun to *increase* with κ, with a flux interaction pattern that is qualitatively different from the one at small κ.

We now connect our findings about the pairwise flux interactions to the character of the flux energetics when the flux density is finite. For $\kappa = 0$ the interactions are longer ranged and hence stronger in magnitude on average than for finite κ; however from Fig. 18 we see that they can be either attractive or repulsive, depending on the relative positions of the two plaquettes. Thus, we do not see a higher curvature for our flux PPE (which averages over all positions) for $\kappa = 0$ than for $\kappa = 0.1, 0.2$ shown in Fig. 14. As noted above, somewhere between $\kappa = 0.2$ and $\kappa = 0.5$ the coherence length begins to increase with κ. Since interactions in this regime are uniformly attractive, this leads to an increase in curvature of the flux PPE, at least at low flux densities. However this does not carry over to high flux density: Near $n_\phi = 1$ the flux PPE for $\kappa = 0.5$ shows the lowest curvature of any of our PPE curves and is relatively close to the straight line expected for noninteracting fluxes. This indicates that the two-flux interaction
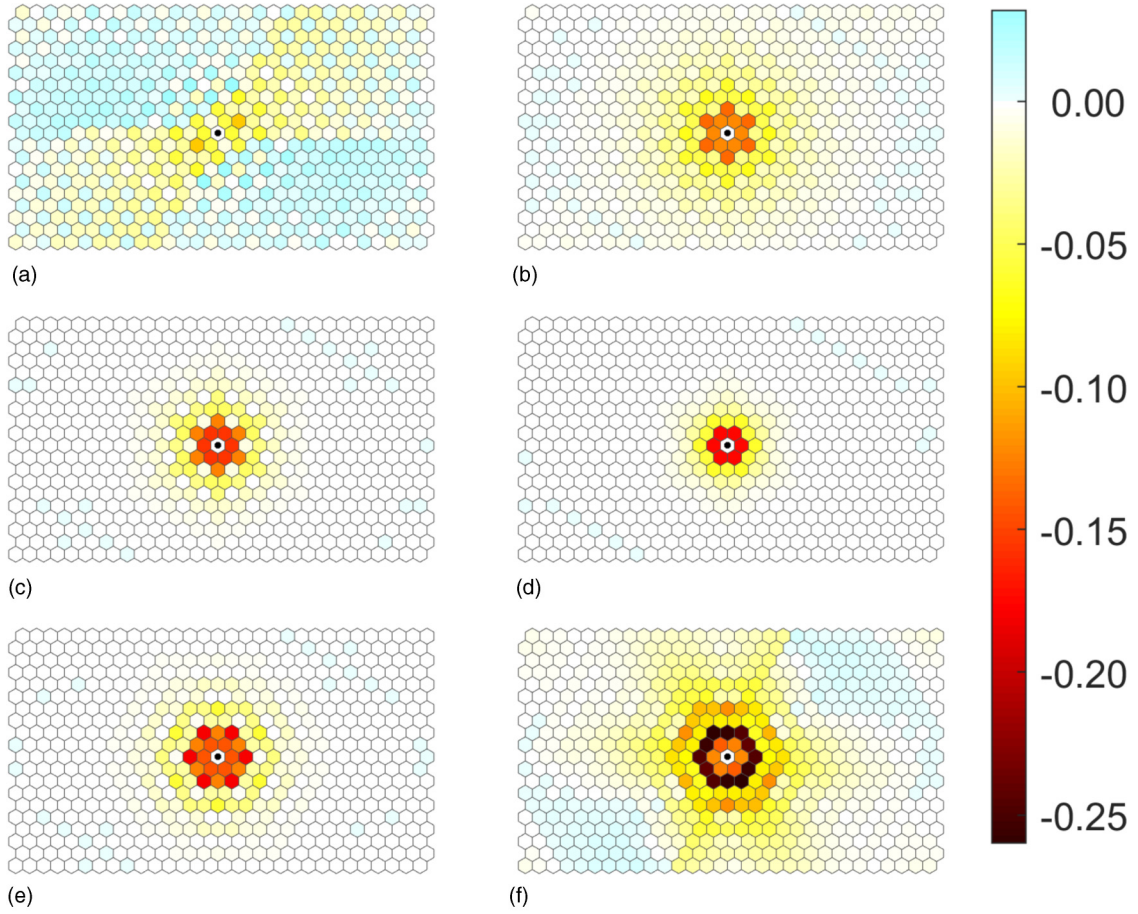
FIG. 17. The two-flux interaction map on a 30×20 torus for various values of κ: (a) $\kappa = 0$, (b) $\kappa = 0.05$, (c) $\kappa = 0.1$, (d) $\kappa = 0.2$, (e) $\kappa = 0.5$, (f) $\kappa = 1$. $E_{\phi,2}(\infty)$ is taken as the energy where separation of the two fluxes is 15 plaquettes.
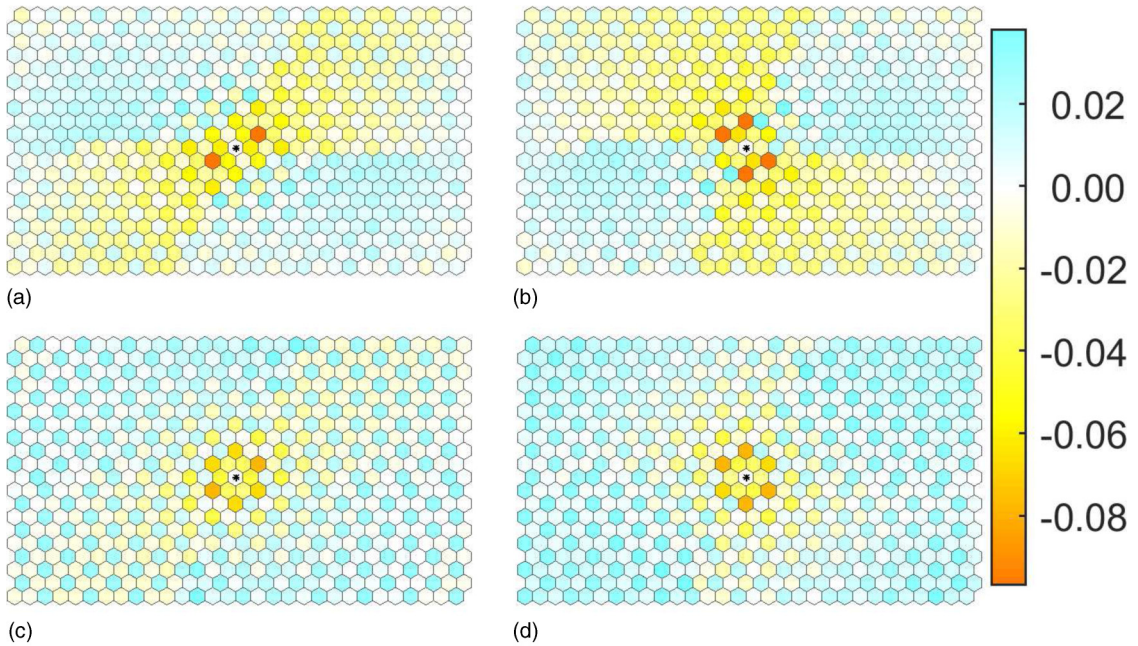

FIG. 18. The two-flux interaction map on a 30×20 lattice for the four topological sectors at $\kappa = 0$ obtained by: (a) is generated by taking bonds along the path described in Fig. 19 to be -1, with all other bonds equal to $+1$. Bonds in (b) and (c) are obtained from bonds in (a) by flipping the z bonds along a loop in the $\mathbf{n}_1$ direction, and the y bonds along a loop in the $\mathbf{n}_2$ direction, respectively. In (d) we have flipped the appropriate bonds along both $\mathbf{n}_1$ and $\mathbf{n}_2$ directions. $E_{\phi,2}(\infty)$ is taken as the energy where separation of the two fluxes is 15 plaquettes.

does not fully capture the flux energetics at high densities. In contrast, for $\kappa = 1$ we observe substantial curvature, both at low flux densities [in agreement with the longer coherence length observed in Fig. 17(f)] and at high flux densities. This is consistent with the 2-flux interaction map, which indicates significantly longer ranged, predominantly attractive interactions at this value of κ.

VIII. SUMMARY

In this work, we address the question of whether the physics characteristic of the boundary of the Kitaev honeycomb spin liquid has observable thermodynamic signatures. We consider two aspects of this question: First, in the ultralow temperature regime, we consider possible thermodynamic signatures of the gapless boundary fermion modes associated with the topological Majorana band structure. Second, at temperatures on the order of the bulk flux gap, we study the impact of boundary conditions, as well as time-reversal symmetry breaking, on flux energetics of the Kitaev model, and discuss the resulting quantitative impact on the specific heat.

We find that in realistic experiments on currently available sample sizes, the topological boundary flat band cannot be seen in the specific heat; rather, it can only be detected indirectly through a contribution to the residual entropy. Because the boundary flat band is topological in nature, this residual entropy is a robust feature of the gapless spin liquid phase—though at the Kitaev point, there is an additional contribution due to boundary operators that commute with the Kitaev Hamiltonian. The chiral edge modes that arise in the presence of time-reversal symmetry breaking, on the other hand, are in principle observable in the specific heat, though the temperature scale for this is set by the parameter κ. In most experiments this term would be generated by a magnetic field, which must be small compared to the flux gap in order for our analysis to be valid. We note that though in the pure Kitaev model the temperature scale of the flux gap is $\Delta_\phi \sim 15$ K for $J \sim 100$ K, in real materials we expect this value to be smaller.

We also find that boundary conditions have a significant impact on flux excitations in the Kitaev model. We quantify this impact on both cylinder and torus geometries by using a polynomial fit (the PPE) for the average flux energy as a function of flux density. We show that a single universal curve provides a good fit for tori of various sizes and aspect ratios for both the time-reversal symmetric and time-reversal broken cases. On the cylinder the best fit depends on both the density of fluxes in the bulk and on the edges and can be described by a single universal surface. Finally, we use these universal best fits to compute the flux contribution to the low-temperature specific heat in the thermodynamic limit and show that these agree well with the specific heat obtained in Monte Carlo simulations at finite system sizes. We show that the flux interactions captured by the PPE models have important quantitative fingerprints on the shapes and positions of the low-temperature peaks in the specific heat computed in different geometries. In particular, in the cylinder, the lower average energy of fluxes near the boundary leads to a broadening of the specific heat peak. Finally, we analyze how a nonzero value of

the time-reversal symmetry breaking strength κ tunes the flux energetics and thus affects the specific heat, finding that the flux energetics are qualitatively different between large and small values of κ.

APPENDIX A: MAJORANA FERMION SPECTRUM IN A GIVEN FLUX SECTOR

Within a given flux sector, the $\hat{u}_{ij}$ operators in the Hamiltonian (2) are replaced by the corresponding eigenvalues u_{ij}, so the Hamiltonian (2) becomes quadratic in the Majorana fermion operators. Exploiting the bipartite nature of the honeycomb lattice, and noting that each unit cell l has two sites $\mathbf{r}_{A,l}$ and $\mathbf{r}_{B,l}$ in the two sublattices A and B, the Hamiltonian (2) can be written as

$$H = \sum_{\langle l,l' \rangle} iM_{ll'}c_{A,l}c_{B,l'} + \sum_{\langle\langle l,l' \rangle\rangle} i\tilde{M}_{ll'}(c_{A,l}c_{A,l'} + c_{B,l}c_{B,l'}),$$

(A1)

where the first term describes the nearest neighbor hopping of the Majorana fermions with $M_{ll'} = -J^\alpha u_{AB,ll'}$ if $\mathbf{r}_{A,l}$ and $\mathbf{r}_{B,l'}$ are connected by α bond and $M_{ll'} = 0$ otherwise. The second term describes the second neighbor hopping between two sites of the A (or B) sublattice with $\tilde{M}_{ll'} = -\kappa u_{AB,ll''}u_{AB,l'l''}$. Using the singular-value decomposition $M = U \cdot \Lambda \cdot V^T$, the resulting free-fermion Hamiltonian can be written in the canonical form as

$$H = \sum_n \varepsilon_n(\psi_n^\dagger \psi_n - 1/2),$$

(A2)

where the fermions $\psi_n = (\gamma_{A,n} + i\gamma_{B,n})/2$ are expressed in terms of Majorana fermions $\gamma_{A,n} = \sum_l U_{ln}c_{A,l}$ and $\gamma_{B,n} = \sum_l V_{ln}c_{B,l}$ on sublattices A and B, respectively, and $\varepsilon_n = 2\Lambda_{nn}$ are their energies. Since M is a real matrix, U and V are real orthogonal matrices, while Λ is a diagonal matrix with non-negative (real) entries.

APPENDIX B: STATISTICAL MECHANICS OF THE KITAEV MODEL

Here we briefly outline the distinctive aspects of the thermodynamics of the Kitaev model [5,24,37]. Given the exact solution of the model [1], both $\mathbb{Z}_2$ fluxes and fermionic fractional excitations contribute to the thermodynamic behavior of the system. In a given flux configuration, ϕ_p, with a given fermionic occupation number configuration, $\{n_i\}$, the corresponding energy of the system is $\sum_i \epsilon_{i,\phi_p}(n_i - \frac{1}{2})$, where the fermionic energy levels ϵ_{i,ϕ_p} are obtained by diagonalizing the Majorana fermion Hamiltonian in a given flux configuration

ϕ_p (see details in Appendix A). The energy of the lowest-energy state in a given flux configuration,

$$E_{\phi_p}^{(0)} \equiv -\frac{1}{2} \sum_i \epsilon_{i,\phi_p}, \tag{B1}$$

which corresponds to all unoccupied fermionic states, is now associated with the energy of a corresponding flux sector.

The partition function of the system is given by:

$$Z = \sum_{\phi_p, n_i} e^{-\beta \sum_i \epsilon_{i,\phi_p}(n_i - \frac{1}{2})} = \sum_{\phi_p} Z_{\phi_p}, \tag{B2}$$

with

$$Z_{\phi_p} = e^{-\beta E_{\phi_p}^{(0)}} \prod_i (1 + e^{-\beta \epsilon_{i,\phi_p}}), \tag{B3}$$

where we denote $\beta = 1/T$. The expectation value of the total energy at a given temperature is then given by

$$\langle E \rangle = \frac{1}{Z} \sum_{\phi_p, n_i} \sum_i \epsilon_{i,\phi_p} \left(n_i - \frac{1}{2} \right) e^{-\beta \sum_i \epsilon_{i,\phi_p}(n_i - \frac{1}{2})}$$

$$= \frac{1}{Z} \sum_{\phi_p} E_{\phi_p} Z_{\phi_p}, \tag{B4}$$

where $E_{\phi_p} = \sum_i \epsilon_{i,\phi_p} n_F(\beta \epsilon_{i,\phi_p}) + E_{\phi_p}^{(0)}$ and $n_F(\beta \epsilon_{i,\phi_p}) = 1/(e^{\beta \epsilon_{i,\phi_p}} + 1)$ is the Fermi-Dirac distribution function. The specific heat of the Kitaev model is thus given by:

$$C = \frac{d\langle E \rangle}{dT}$$

$$= -\frac{1}{T^2} \sum_{\phi_p} \left(E_{\phi_p} \frac{\partial(Z_{\phi_p}/Z)}{\partial \beta} + \frac{Z_\phi}{Z} \frac{\partial E_{\phi_p}}{\partial \beta} \right)$$

$$= \frac{1}{T^2} \left(\langle E_{\phi_p}^2 \rangle - \langle E_{\phi_p} \rangle^2 - \left\langle \frac{\partial E_{\phi_p}}{\partial \beta} \right\rangle \right), \tag{B5}$$

where we have simplified the first term in the second line via:

$$\frac{1}{Z} \sum_{\phi_p} E_{\phi_p} \frac{\partial Z_{\phi_p}}{\partial \beta} - \frac{1}{Z^2} \frac{\partial Z}{\partial \beta} \sum_{\phi_p} E_{\phi_p} Z_{\phi_p} \tag{B6}$$

$$= -\frac{1}{Z} \sum_{\phi_p} E_{\phi_p}^2 Z_{\phi_p} + \langle E_{\phi_p} \rangle \frac{1}{Z} \sum_{\phi_p} E_{\phi_p} Z_{\phi_p} \tag{B7}$$

$$= -\langle E_{\phi_p}^2 \rangle + \langle E_{\phi_p} \rangle^2, \tag{B8}$$

i.e., it is simply the variance of the fermionic energy.

APPENDIX C: MONTE CARLO METHOD FOR THE FLUX-FERMION MODEL

In this Appendix, we discuss details of the implementation of the MC algorithm [36], which has been used for the computation of the specific heat of the flux-fermion model (B5). The basic idea of this MC algorithm is that we can perform sampling over flux configurations $\{\phi_p\}$ classically by exploiting the fact that the energy of each flux configuration, $E_{\phi_p}^{(0)}$, can be computed exactly by diagonalizing the quadratic Majorana Hamiltonian (A2). From Eqs. (B2) and (B3), the
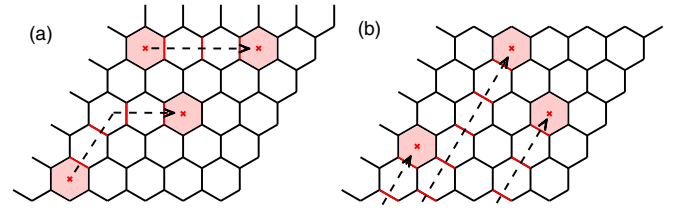

FIG. 19. Gauge choice used to generate a given flux configuration on (a) torus and (b) cylinder in our MC algorithm. The red bonds correspond to $u_{ij} = -1$. The fluxes are located on the red plaquettes.

probability distribution function for fluxes is defined as

$$p(\phi_p) = \frac{Z_{\phi_p}}{Z} = \frac{1}{Z} e^{-\beta E_{\phi_p}^{(0)}} \prod_i (1 + e^{-\beta \epsilon_i}). \tag{C1}$$

In practice, to implement a given flux configuration ϕ_p, we must choose one of the many possible bond configurations $\{u_{ij}\}$ that lead to fluxes on the desired plaquettes. Figure 19 shows the convention that we use for our MC simulations. On a torus lattice, a pair of fluxes is generated as follows. (i) Choose a pair of plaquettes, which are separated by $a\mathbf{n}_1 + b\mathbf{n}_2$. (ii) Choose a path on the dual lattice that connects these two plaquettes by first crossing a bonds in the $+\mathbf{n}_1$ direction, and then b bonds in the $+\mathbf{n}_2$ direction, and flip the sign of u_{ij} on all bonds that this path crosses. The same convention can be used to move a flux between two plaquettes. To create a more general flux configuration, steps (i) and (ii) can be repeated. On a cylinder lattice, we can create a single flux on the plaquette p by choosing a path along the $+\mathbf{n}_2$ direction connecting p to the lower boundary of the cylinder and flipping the sign of u_{ij} on all of the bonds crossed by this path. Flux annihilation is carried out using the same conventions. On the cylinder, we move a flux from plaquette p to plaquette p' by annihilating the flux on plaquette p and creating a new flux on p'.

All the simulations for the specific heat reported in the main text were done in the following way. The simulations were performed on finite-size lattices with spanning vectors $(\mathbf{R}_1, \mathbf{R}_2) = (N_1\mathbf{n}_1, N_2\mathbf{n}_2)$ with either periodic boundary conditions, which we refer to as torus lattices, or with open boundary conditions in the $\mathbf{n}_2$ direction and periodic boundary conditions in the $\mathbf{n}_1$ direction, which we refer to as cylinder lattices. We initially perform 1000 MC steps for thermalization and then 100 000 steps for measurement on a lattice of up to $2 \times 24 \times 6 = 288$ sites.

Three kinds of MC moves are implemented: (1) Randomly shuffle all the fluxes while keeping the total flux number fixed. (2) Count the current total flux number, change the total flux number by two, and then randomly place the updated number of fluxes on the plaquettes. (3) Flip all the bonds crossed by a path encircling the system in (one of) the periodic directions, which will change the topological sector while keeping the flux configuration fixed. The transition (or acceptance) probability, satisfying the detailed balance condition, is then defined as

$$T(x \to y) = \min\left(1, \frac{p(y)n(y)}{p(x)n(x)}\right), \tag{C2}$$
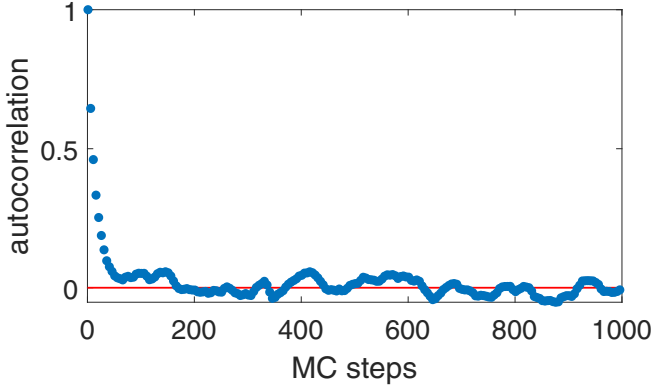
FIG. 20. The autocorrelation as a function of the number of MC steps on an $N_1 \times N_2 = 16 \times 16$ torus lattice, at the temperature of $10^{-1.3}$. This is approximately the temperature of the low-temperature peak in the specific heat shown in Fig. 7.

where x represents a certain flux configuration $\{\phi_p\}$ and $n(x)$ is the number of the candidate flux configurations in a proposal (see Ref. [75] for details). Since these MC moves are all global, our algorithm does not struggle with local minima, and we find good convergence without requiring parallel tempering. In order to demonstrate the quality of convergence, in Fig. 20 we plot the autocorrelation function of our MC algorithm for $N_1 \times N_2 = 16 \times 16$ torus lattice at temperature equal to $10^{-1.3} \approx 0.05$, a typical temperature near the low-temperature peak of specific heat. (We have found that this is the temperature regime where MC convergence is most difficult). The autocorrelation function reaches zero in around 200 MC update steps, indicating that our MC moves lead to thermalization on a relatively rapid timescale.

A note of caution is in order here. For lattices with periodic boundaries, there is a parity constraint [76] on the total fermion parity $\prod_i (-1)^{n_i}$. Specifically, for a given lattice shape with fixed boundary conditions and gauge field (bond operator) configurations, we find that the physical states have fixed fermion parity; states in the flux-fermion model with opposite fermion parity are unphysical, in the sense that they do not correspond to states in the spin Hilbert space [65]. In the thermodynamic limit, this constraint has little relevance for the model's spectrum and hence its thermodynamics; however it can lead to substantial finite-size corrections to the energy spectrum, which in turn can be relevant for thermodynamic quantities such as the specific heat.

One way to circumvent the parity constraint is to remove a single bond from the lattice [36]. Exactly as for the open cylinder boundary discussed in the main text, this creates a pair of dangling Majorana fermions and a corresponding zero-energy state. This ensures that for every state with odd fermion parity, there is a state of identical energy with even fermion parity; in this case we do not need to know which of the two states is physical to compute thermodynamic quantities exactly. (We do, however, need to account for the impact of this zero mode on the entropy S_∞.) Though the numerical results shown here are evaluated with all bonds present in the lattice, we have checked that the differences caused by removing the single bond are within the statistical error of our Monte Carlo simulations.
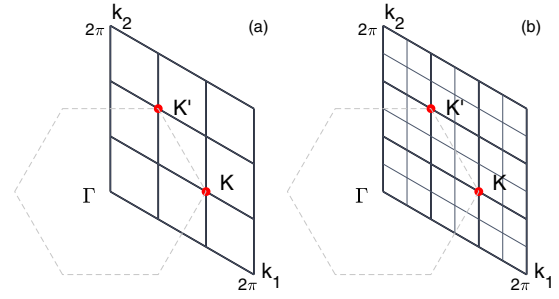

FIG. 21. The reciprocal lattices that correspond to two real-space lattice sizes: (a) $L = 3$, (b) $L = 6$, where $N_1 = N_2 = L$. The dashed honeycomb is the Brillouin zone. The two red points K, K' are the two nonequivalent Dirac points. This shows that only when $L = 3k$ do the Dirac points sit on the sites of reciprocal lattice.

APPENDIX D: EXTRAPOLATING THE FLUX GAP ENERGY TO THERMODYNAMIC LIMIT

As we discussed in Sec. V A of the main text, when extrapolating the flux gap energy to infinitely large lattice size, generally the energies display period-3 oscillatory finite-size behavior [1]. In order to understand this behavior, we consider the flux-free sector, where the 2D reciprocal space is well defined. Figure 21 shows the reciprocal lattices corresponding to two real-space lattice sizes: $L = 3$ and $L = 6$. We see that for $L = 3k$ the two Dirac points K, K' can sit on reciprocal lattice sites, while for $L = 3k + 1$ and $L = 3k + 2$, the Dirac point is off the reciprocal lattice by a distance of order $1/L$. This difference in the fermionic spectrum $\{\epsilon_i\}$ leads to a different series of flux gaps $E_{\phi_p}^{(0)}$ for $L = 3k, 3k + 1$, and $3k + 2$, respectively.

APPENDIX E: THE BEST-FIT FLUX PPE POLYNOMIALS FOR DIFFERENT κ ON THE $N_1 \times N_2 = 16 \times 16$ torus

The equation of the best-fit polynomials $\mathcal{E}(n_\phi)$ for various κ shown in Fig. 14 are:

$$\kappa = 0.0, \quad \mathcal{E}(n_\phi) = -0.309 n_\phi^6 + 0.814 n_\phi^5 - 0.859 n_\phi^4$$
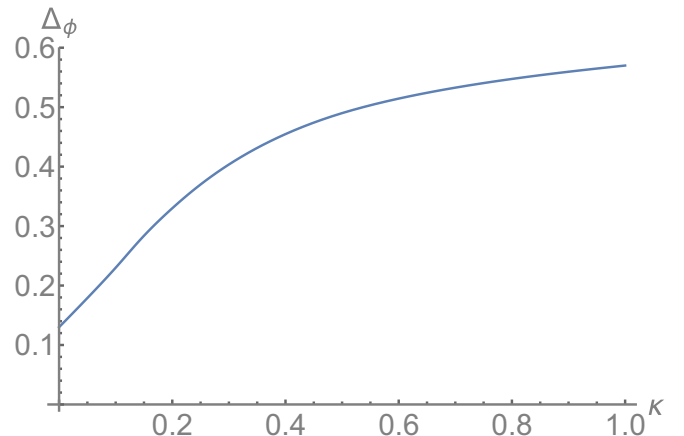$$+ 0.456 n_\phi^3 - 0.176 n_\phi^2 + 0.141 n_\phi$$

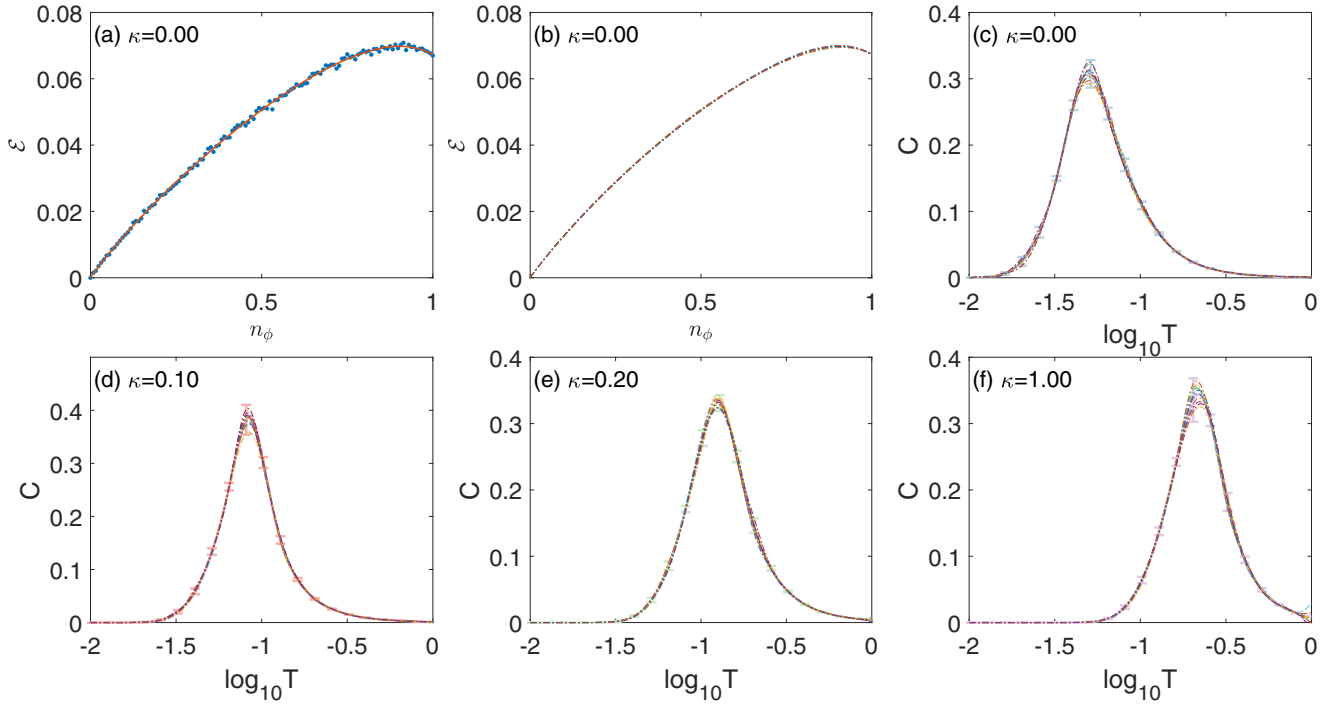FIG. 22. The dependence of the flux gap Δ_ϕ on κ.

FIG. 23. Error estimates for specific heat curves from flux PPE fits. (a) One random sample drawn from the data, shown together with the corresponding flux PPE curve. (b) An ensemble of 10 selected flux PPE curves generated from these random samples. (c)–(f) Plots of the corresponding specific heat, calculated using Eq. (33). We see that very small variations in the flux PPE correspond to relatively large variations in the precise shape of the specific heat peak.

$$
\begin{aligned}
\kappa = 0.1, \quad \mathcal{E}(n_\phi) &= -0.024 n_\phi^5 - 0.115 n_\phi^4 \\
&\quad + 0.286 n_\phi^3 - 0.282 n_\phi^2 + 0.232 n_\phi \\
\kappa = 0.2, \quad \mathcal{E}(n_\phi) &= -0.055 n_\phi^5 - 0.075 n_\phi^4 \\
&\quad + 0.270 n_\phi^3 - 0.316 n_\phi^2 + 0.330 n_\phi \\
\kappa = 0.5, \quad \mathcal{E}(n_\phi) &= 0.239 n_\phi^5 - 0.652 n_\phi^4 \\
&\quad + 0.679 n_\phi^3 - 0.403 n_\phi^2 + 0.489 n_\phi \\
\kappa = 1.0, \quad \mathcal{E}(n_\phi) &= -0.157 n_\phi^5 - 0.062 n_\phi^4 \\
&\quad + 0.348 n_\phi^3 - 0.536 n_\phi^2 + 0.566 n_\phi.
\end{aligned}
\tag{E1}
$$

We find that while sixth-order polynomials give the best fit to C for $\kappa = 0$, for nonzero kappa a better fit is achieved with fifth-order polynomials. We note, however, that it is difficult to distinguish the fifth- and sixth-order fits based on residual errors, which are comparable in both cases. This reflects the extreme sensitivity of C to the details of the PPE curve used. The dependence of the flux gap Δ_ϕ computed by the first derivatives of the PPE polynomials on the strength of κ is shown in Fig. 22.

APPENDIX F: THE COMPUTATION OF THE ERROR BARS FOR PPE CURVES

To obtain the error estimates in Figs. 12, 13, and 15, we generate 60 PPE-like curves by randomly sampling from the data used to generate the initial PPE curve, choosing one value of the energy density for each value of the flux density. For each set of data, we generate a PPE curve using the best-fit polynomial obtained from the mean of the data. Then, from each such PPE curve, we generate a specific heat curve using Eq. (33) [or, for cylinder lattices, Eq. (37)]. The error bars at each point show one standard deviation of the resulting curves. Some results are shown in Fig. 23.

[1] A. Kitaev, Ann. Phys. **321**, 2 (2006).

[2] M. Hermanns, I. Kimchi, and J. Knolle, Annu. Rev. Condens. Matter Phys. **9**, 17 (2018).

[3] S. Trebst, arXiv:1701.07056.

[4] H. Takagi, T. Takayama, G. Jackeli, G. Khaliullin, and S. E. Nagler, Nat. Rev. Phys. **1**, 264 (2019).

[5] Y. Motome and J. Nasu, J. Phys. Soc. Jpn **89**, 012002 (2020).

[6] G. Jackeli and G. Khaliullin, Phys. Rev. Lett. **102**, 017205 (2009).

[7] J. Chaloupka, G. Jackeli, and G. Khaliullin, Phys. Rev. Lett. **105**, 027204 (2010).

[8] J. G. Rau, E. K.-H. Lee, and H.-Y. Kee, Ann. Rev. Cond. Matt. Phys. **7**, 195 (2016).

[9] S. M. Winter, A. A. Tsirlin, M. Daghofer, J. van den Brink, Y. Singh, P. Gegenwart, and R. Valentí, J. Phys.: Condens. Matter **29**, 493002 (2017).

[10] K. Plumb, J. Clancy, L. Sandilands, V. V. Shankar, Y. Hu, K. Burch, H.-Y. Kee, and Y.-J. Kim, Phys. Rev. B **90**, 041112(R) (2014).

[11] M. Majumder, M. Schmidt, H. Rosner, A. A. Tsirlin, H. Yasuoka, and M. Baenitz, Phys. Rev. B **91**, 180401(R) (2015).

[12] R. D. Johnson, S. C. Williams, A. A. Haghighirad, J. Singleton, V. Zapf, P. Manuel, I. I. Mazin, Y. Li, H. O. Jeschke, R. Valentí, and R. Coldea, Phys. Rev. B **92**, 235119 (2015).

[13] J. A. Sears, M. Songvilay, K. W. Plumb, J. P. Clancy, Y. Qiu, Y. Zhao, D. Parshall, and Y.-J. Kim, Phys. Rev. B **91**, 144420 (2015).

[14] Y. Kasahara, T. Ohnishi, Y. Mizukami, O. Tanaka, S. Ma, K. Sugii, N. Kurita, H. Tanaka, J. Nasu, Y. Motome *et al.*, Nature (London) **559**, 227 (2018).

[15] K. Kitagawa, T. Takayama, Y. Matsumoto, A. Kato, R. Takano, Y. Kishimoto, S. Bette, R. Dinnebier, G. Jackeli, and H. Takagi, Nature (London) **554**, 341 (2018).

[16] I. Rousochatzakis, S. Kourtis, J. Knolle, R. Moessner, and N. B. Perkins, Phys. Rev. B **100**, 045117 (2019).

[17] A. Banerjee, C. A. Bridges, J.-Q. Yan, A. A. Aczel, L. Li, M. B. Stone, G. E. Granroth, M. D. Lumsden, Y. Yiu, J. Knolle, S. Bhattacharjee, D. L. Kovrizhin, R. Moessner, D. A. Tennant, D. G. Mandrus, and S. E. Nagler, Nat. Mater. **15**, 733 (2016).

[18] J. Knolle, D. L. Kovrizhin, J. T. Chalker, and R. Moessner, Phys. Rev. Lett. **112**, 207203 (2014).

[19] J. Knolle, D. L. Kovrizhin, J. T. Chalker, and R. Moessner, Phys. Rev. B **92**, 115127 (2015).

[20] M. Udagawa and R. Moessner, arXiv:1912.01545.

[21] J. Knolle, G.-W. Chern, D. L. Kovrizhin, R. Moessner, and N. B. Perkins, Phys. Rev. Lett. **113**, 187201 (2014).

[22] L. J. Sandilands, Y. Tian, K. W. Plumb, Y.-J. Kim, and K. S. Burch, Phys. Rev. Lett. **114**, 147201 (2015).

[23] L. J. Sandilands, Y. Tian, A. A. Reijnders, H.-S. Kim, K. W. Plumb, Y.-J. Kim, H.-Y. Kee, and K. S. Burch, Phys. Rev. B **93**, 075144 (2016).

[24] J. Nasu, J. Knolle, D. L. Kovrizhin, Y. Motome, and R. Moessner, Nat. Phys. **12**, 912 (2016).

[25] B. Perreault, J. Knolle, N. B. Perkins, and F. J. Burnell, Phys. Rev. B **92**, 094439 (2015).

[26] B. Perreault, J. Knolle, N. B. Perkins, and F. J. Burnell, Phys. Rev. B **94**, 060408(R) (2016).

[27] B. Perreault, J. Knolle, N. B. Perkins, and F. J. Burnell, Phys. Rev. B **94**, 104427 (2016).

[28] Y. Wang, G. B. Osterhoudt, Y. Tian, P. Lampen-Kelley, A. Banerjee, T. Goldstein, J. Yan, J. Knolle, H. Ji, R. J. Cava, J. Nasu, Y. Motome, S. E. Nagler, D. Mandrus, and K. S. Burch, npj Quantum Mater. **5**, 14 (2020).

[29] A. Sahasrabudhe, D. A. S. Kaib, S. Reschke, R. German, T. C. Koethe, J. Buhot, D. Kamenskyi, C. Hickey, P. Becker, V. Tsurkan, A. Loidl, S. H. Do, K. Y. Choi, M. Grüninger, S. M. Winter, Z. Wang, R. Valentí, and P. H. M. van Loosdrecht, Phys. Rev. B **101**, 140410(R) (2020).

[30] D. Wulferding, Y. Choi, S.-H. Do, C. H. Lee, P. Lemmens, C. Faugeras, Y. Gallais, and K.-Y. Choi, Nat. Commun. **11**, 1603 (2020).

[31] G. B. Halász, N. B. Perkins, and J. van den Brink, Phys. Rev. Lett. **117**, 127203 (2016).

[32] G. B. Halász, B. Perreault, and N. B. Perkins, Phys. Rev. Lett. **119**, 097202 (2017).

[33] G. B. Halász, S. Kourtis, J. Knolle, and N. B. Perkins, Phys. Rev. B **99**, 184417 (2019).

[34] M. Ye, R. M. Fernandes, and N. B. Perkins, Phys. Rev. Res. **2**, 033180 (2020).

[35] A. Metavitsiadis and W. Brenig, Phys. Rev. B **101**, 035103 (2020).

[36] J. Nasu, M. Udagawa, and Y. Motome, Phys. Rev. Lett. **113**, 197205 (2014).

[37] J. Nasu, M. Udagawa, and Y. Motome, Phys. Rev. B **92**, 115122 (2015).

[38] S.-H. Do, S.-Y. Park, J. Yoshitake, J. Nasu, Y. Motome, Y. S. Kwon, D. T. Adroja, D. J. Voneshen, K. Kim, T.-H. Jang, J.-H. Park, K.-Y. Choi, and S. Ji, Nat. Phys. **13**, 1079 (2017).

[39] Y. Kato, Y. Kamiya, J. Nasu, and Y. Motome, Phys. Rev. B **96**, 174409 (2017).

[40] T. Eschmann, P. A. Mishchenko, T. A. Bojesen, Y. Kato, M. Hermanns, Y. Motome, and S. Trebst, Phys. Rev. Research **1**, 032011 (2019).

[41] G. Kells and J. Vala, Phys. Rev. B **82**, 125122 (2010).

[42] S. Matsuura, P.-Y. Chang, A. P. Schnyder, and S. Ryu, New J. Phys. **15**, 065001 (2013).

[43] K. O'Brien, M. Hermanns, and S. Trebst, Phys. Rev. B **93**, 085101 (2016).

[44] R. Schaffer, Eric Kin-Ho Lee, Y.-M. Lu, and Y. B. Kim, Phys. Rev. Lett. **114**, 116803 (2015).

[45] R. Wang, C. Nisoli, R. S. d. Freitas, J. Li, W. McConville, B. Cooley, M. Lund, N. Samarth, C. Leighton, V. Crespi *et al.*, Nature (London) **439**, 303 (2006).

[46] A. P. Ramirez, A. Hayashi, R. a. Cava, R. Siddharthan, and B. Shastry, Nature (London) **399**, 333 (1999).

[47] S. T. Bramwell and M. J. Gingras, Science **294**, 1495 (2001).

[48] G. Lau, R. S. d. Freitas, B. Ueland, B. Muegge, E. Duncan, P. Schiffer, and R. J. Cava, Nat. Phys. **2**, 249 (2006).

[49] R. Higashinaka, H. Fukazawa, and Y. Maeno, Phys. Rev. B **68**, 014415 (2003).

[50] P. Hohenberg and W. Kohn, Phys. Rev. **136**, B864 (1964).

[51] R. O. Jones and O. Gunnarsson, Rev. Mod. Phys. **61**, 689 (1989).

[52] F. D. M. Haldane, Phys. Rev. Lett. **61**, 2015 (1988).

[53] Y. Zhao and Z. Wang, Phys. Rev. Lett. **110**, 240404 (2013).

[54] S. Ryu, A. P. Schnyder, A. Furusaki, and A. W. Ludwig, New J. Phys. **12**, 065010 (2010).

[55] B. A. Bernevig and T. L. Hughes, *Topological Insulators and Topological Superconductors* (Princeton University Press, Princeton, NJ, 2013).

[56] Though strictly speaking the classification of nodal topological phases requires translation invariance, it is easy to check numerically that zero-energy states localized at the boundary remain a feature of the spectrum with weak disorder—and in fact persist up to strong disorder.

[57] K. Wakabayashi, K.-i. Sasaki, T. Nakanishi, and T. Enoki, Sci. Tech. Adv. Mater. **11**, 054504 (2010).

[58] G. Bastien, M. Roslova, M. H. Haghighi, K. Mehlawat, J. Hunger, A. Isaeva, T. Doert, M. Vojta, B. Büchner, and A. U. B. Wolter, Phys. Rev. B **99**, 214410 (2019).

[59] M. Roslova, J. Hunger, G. Bastien, D. Pohl, H. M. Haghighi, A. U. Wolter, A. Isaeva, U. Schwarz, B. Rellinghaus, K. Nielsch *et al.*, Inorg. Chem. **58**, 6659 (2019).

[60] A. U. B. Wolter, L. T. Corredor, L. Janssen, K. Nenkov, S. Schönecker, S.-H. Do, K.-Y. Choi, R. Albrecht, J. Hunger, T. Doert, M. Vojta, and B. Buchner, Phys. Rev. B **96**, 041405(R) (2017).

[61] K. Wakabayashi, Y. Takane, M. Yamamoto, and M. Sigrist, New J. Phys. **11**, 095016 (2009).

[62] F. J. Dyson, Phys. Rev. **92**, 1331 (1953).

[63] O. Motrunich, K. Damle, and D. A. Huse, Phys. Rev. B **63**, 224204 (2001).

[64] I. A. Gruzberg, N. Read, and S. Vishveshwara, Phys. Rev. B **71**, 245124 (2005).

[65] F. Zschocke and M. Vojta, Phys. Rev. B **92**, 014403 (2015).

[66] W.-H. Kao, J. Knolle, G. B. Halász, R. Moessner, and N. B. Perkins, arXiv:2007.11637.

[67] H. Li, D.-W. Qu, H.-K. Zhang, Y.-Z. Jia, S.-S. Gong, Y. Qi, and W. Li, Phys. Rev. Research **2**, 043015 (2020).

[68] V. Lahtinen, New J. Phys. **13**, 075009 (2011).

[69] B. Efron, SIAM Rev. **21**, 460 (1979).

[70] M. Cheng, R. M. Lutchyn, V. Galitski, and S. Das Sarma, Phys. Rev. Lett. **103**, 107001 (2009).

[71] V. Lahtinen, A. W. W. Ludwig, J. K. Pachos, and S. Trebst, Phys. Rev. B **86**, 075115 (2012).

[72] A. W. Ludwig, D. Poilblanc, S. Trebst, and M. Troyer, New J. Phys. **13**, 045014 (2011).

[73] V. Lahtinen, A. W. W. Ludwig, and S. Trebst, Phys. Rev. B **89**, 085121 (2014).

[74] S.-S. Zhang, Z. Wang, G. B. Halász, and C. D. Batista, Phys. Rev. Lett. **123**, 057201 (2019).

[75] C. J. Geyer, Stat. Sci. **7**, 473 (1992).

[76] F. L. Pedrocchi, S. Chesi, and D. Loss, Phys. Rev. B **84**, 165414 (2011).