

Supplementary Material for Axion Mass Prediction from adaptive mesh refinement cosmological lattice simulations

Joshua N. Benabou, Malte Buschmann, Joshua W. Foster, and Benjamin R. Safdi

This Supplementary Material provides additional details and results for the analyses discussed in the main Letter.

1. SUPPLEMENTARY FIGURES




Figure S1. (*Left*) The inverse expectation value $\langle H/k \rangle^{-1}$ (data points) of the axion emission spectrum F for our fiducial spectrum with $k_{\text{IR}}/H = 50$ (plotted in Fig. 3). Here we assume the spectral index is at its upper limit at 1σ , $q = 1.02$. The error bars are data-driven, obtained by fitting $\langle H/k \rangle^{-1} = \delta\sqrt{\xi}$, for a constant δ , across the indicated log data points (best fit shown in solid). We assume a Gaussian likelihood and treat the standard deviation σ as a nuisance parameter. We find $\delta = 318.2 \pm 2.1$. (*Right*) As in the left panel but assuming a conformal spectrum $q = 1$. In this case, as discussed in the main text, we expect $\delta = \delta_1 \text{Log}$; we measure $\delta_1 = 8.6 \pm 0.1$. In both panels, the UV cutoff is set by $\log(m_r/H_{\text{QCD}}) = 70$.

2. SIMULATION SETUP AND REFINEMENT CRITERIA

Our simulations are performed with the public code `sledgehamr` [53], which is based on the block-structured adaptive mesh refinement framework `AMReX` [79, 80]. Using `sledgehamr`, we model the dynamics of the PQ field, cosmic strings, and domain walls by evolving a complex field in a comoving volume with periodic boundary conditions in a radiation-dominated cosmological background using the method of lines. Our simulations, which we detail further here, are similar to but improve upon in several ways those implemented in [49].

In making our spatial discretization, we evaluate the Laplacian of fields using the fourth-order accurate 13-point finite-difference stencil. Our equations of motion are integrated in time using a third-order strong-stability preserving forward Runge-Kutta. The time-step used in our simulations is set by a Courant-Friedrichs-Lewy (CFL) condition of $\Delta\eta \approx \Delta x/3$ to sufficiently resolve the characteristic wave speed, which is 1 in our units. Note that because the grid spacing Δx depends on the refinement level, our time-step size varies from level to level as well. To coordinate the different time-step sizes we use a sub-cycling-in-time algorithm [53].

In our simulations, the adaptive meshing technique is determined via user-defined criteria for tagging lattice sites. Lattice sites that are tagged, as well as those in their vicinity, are upsampled by a nested sub-lattice with twice the spatial resolution as the tagged cells. Tagging and refinement proceed iteratively so that lattice sites tagged at the base level are upsampled to the first refinement level; then, if lattice sites on the first refinement level are tagged, they are in turn upsampled to a second refinement level, and so on. As this procedure is computationally expensive it is done only periodically every ten time-steps. This explains why not only tagged cells but also those in their vicinity are upsampled, as we need to ensure that no string will be able to leave the refined region before it is readjusted ten

time-steps later. Note that this also implies that higher refinement levels are necessarily readjusted more often. For details, see, *e.g.*, [53, 79, 80].

We implement two independent refinement criteria, with lattice sites tagged for up-sampling if either criterion is satisfied. The first and simplest criterion is based on the identification of strings, allowing us to manually enforce high resolution of regions that contain these intrinsically small-scale features. The second and more general criterion is based on a *self-shadow* method which automatically identifies areas of the simulation volume that must be maintained at high spatial resolution based on an estimate of the local absolute truncation error.

A. String Tagging

We localize strings in our simulation volume following an algorithm detailed in [54] to identify all plaquettes (square loops with vertices at the lattice sites of our simulation) that are pierced by a string. The lattice sites that form the vertices of pierced plaquettes at level ℓ are then tagged if strings would be resolved by less than 4 lattice sites as measured by if $m_r \Delta x < 1/4$ at the refinement level of interest. If strings are resolved by more than four lattice sites at level ℓ , then those sites are not tagged for continued adaptive meshing as the string resolution is considered sufficient. At a given refinement level ℓ with spatial resolution Δx_ℓ , we may estimate the number of lattice sites that resolve a string core by $1/(\eta m_r \Delta x_\ell)$, where the radial mode has mass $m_r = \sqrt{2}$ in our simulation units. Note that a factor of η appears in estimating the string resolution as the fixed physical size of our strings appears to shrink on our lattice defined in comoving coordinates. As a result, the number of levels of refinement required to maintain sufficient resolution of strings will grow over the course of simulations. Note also that because the strings shrink continuously while our resolution improves only in factors of 2, our string resolution will vary between four and eight lattice sites. Additionally, the number and location of string-plaquette piercings at the finest tagged level are tracked to determine the total string length ξ .

B. Self-Shadow Tagging

Our second, and more general, refinement criterion is based on the *self-shadow* technique, and its implementation is described in detail in [53]. In brief, the self-shadow technique assesses local absolute truncation error by comparing the simulation state at identical locations at refinement levels ℓ and $\ell + 1$ after both levels have been independently evolved for $\Delta \eta_\ell$ as part of the sub-cycling-in-time algorithm. If the difference from level ℓ and $\ell + 1$ exceeds some sufficiently small threshold ϵ_ℓ , then this indicates that although data at level $\ell + 1$ remains trustworthy, that truncation error may become appreciable, and so a new level $\ell + 2$ is generated to keep the simulation over-resolved at the finest level. We are free to choose truncation error thresholds which depend on the fields being compared and on the levels at which they are compared. We choose

$$\epsilon_\ell = 3 \times 10^{-4}, \quad \epsilon'_\ell = 2^\ell \times 3 \times 10^{-4}, \quad (\text{S1})$$

using ϵ_ℓ to compare ψ_1 and ψ_2 across levels $\ell - 1$ and ℓ , and similarly, using ϵ'_ℓ to compare ψ'_1 and ψ'_2 . While phenomenologically motivated, we found these provided effective data-driven tagging criteria sufficient for identifying domain walls, oscillons, and other localized small-scale features that arise over the course of our simulations without feature-specific tagging algorithms like those implemented for strings. Additionally, these criteria were previously found satisfactory in [24].

3. RECONSTRUCTING THE STRING AND DOMAIN WALL NETWORK

In this section, we review the algorithmic procedures by which we identify and track strings and domain walls in our simulations. Beyond just measuring the total string length ξ , we aim to more fully reconstruct the string and domain wall network in the interest of extracting additional properties of the network evolution, such as the local string curvatures. These properties help to confirm that the network behaves as expected and that we have indeed approached the scaling solution by the time we extract the spectral index q .

A. Identifying Individual Strings

We start by identifying string-plaquette piercings at the finest refinement level of a single simulation snapshot using the algorithm described in [54]. We then estimate the exact location of the string piercing within that plaquette using

the interpolation method outlined in [81]. After identifying the precise location at which each plaquette is pierced, we link nearby piercings together using a nearest-neighbor algorithm to identify continuous string loops.

String loop identification is computationally expensive finding nearest-neighbor pierced plaquettes involves evaluating their pairwise distances, and a single simulation snapshot may contain over 10^7 string-pierced plaquettes. To alleviate this issue we divide our simulation volume into $16^3 = 4096$ equally-size subvolumes to cut down the number of potential locations that need to be considered at any given point. In each subvolume, we start by picking a random location that has not yet been assigned to a string chain and use it to start a new chain. We compute the distance from this location to all other locations within that subvolume. We find the closest unassigned location and add it to the chain as long as its distance is less than $\sqrt{3}\Delta x$. If no such location can be found we stop. This happens whenever the current string segment leaves the subvolume or when the chain closes on itself. Afterward, if there are still unassigned string locations, we start a new chain by repeating the above steps. We follow this procedure in all subvolumes in parallel.

After identifying partial string chains, we link them by making pairwise connections of chain ends that are within a distance of $\sqrt{3}\Delta x$. Here, we must carefully take into account the periodicity of the simulation volume, and we manually confirm that after the linking procedure, all string chains form closed loops. Generally, the nearest-neighbor approach works well but can produce small errors in regions where strings are intersecting or where the string curvature is very large, *i.e.*, comparable to the inverse lattice spacing at the refinement level where string tagging is performed, as nearest-neighbor identification can lead to isolated plaquette-piercings which are not assigned to any chain. However, this occurs in fewer than 10^{-5} of the plaquettes so we neglect them in the following.

B. Extracting the String Curvature

After identifying continuous string loops, we can parametrize each string as a curve by $\gamma(s) = (x(s), y(s), z(s))$, where s is the arc length along the string. From $\gamma(s)$, we compute the local string curvature κ as

$$\kappa = \frac{|\gamma' \times \gamma''|}{|\gamma'|^3} = \frac{\sqrt{(x''y' - y''x')^2 + (x''z' - z''x')^2 + (y''z' - z''y')^2}}{(x'^2 + y'^2 + z'^2)^{3/2}}, \quad (\text{S2})$$

where primes indicate a derivative with respect to s . We compute these derivatives numerically using a finite difference stencil. Since all string locations are unevenly spaced in s , we need to use a non-uniform stencil. For the first derivative, we use the fourth-order stencil, which we derive to be

$$x'_i = \frac{A_0 + A_1 + A_2}{h_{-1}h_{+1}h_{-2}h_{+2}(h_{-1} + h_{+1})(h_{-2} + h_{+2})(h_{-1}h_{+1} - h_{-2}h_{+2})}, \quad (\text{S3})$$

with

$$\begin{aligned} A_0 &= (h_{-1} + h_{+1})(h_{-2} + h_{+2})(h_{-1}^2h_{+1}^2(h_{+2} - h_{-2}) + h_{-2}^2h_{+2}^2(h_{-1} - h_{+1}))x_i, \\ A_1 &= h_{-2}^2h_{+2}^2(h_{+2} + h_{-2})(h_{-1}^2x_{i+1} - h_{+1}^2x_{i-1}), \\ A_2 &= h_{-1}^2h_{+1}^2(h_{+1} + h_{-1})(h_{-2}^2x_{i+2} - h_{+2}^2x_{i-2}), \end{aligned} \quad (\text{S4})$$

where $h_{\pm k} = |s_i - s_{i\pm k}|$. Not shown here are the trivial modifications to account for the periodic boundary conditions. For the second derivative, we use the second-order stencil which we derive to be

$$x''_i = 2 \frac{h_{-1}x_{i+1} + h_{+1}x_{i-1} - (h_{+1} + h_{-1})x_i}{h_{-1}h_{+1}(h_{-1} + h_{+1})}. \quad (\text{S5})$$

We also tested the third-order stencil² for the second derivative but found it to be negatively affected by numerical noise at the grid level. We furthermore increase our stencil window to $s_{i\pm k} \rightarrow s_{i\pm 10k}$ as typically the curvature radius is much larger than the grid spacing, $R_c = 1/\kappa \gg \Delta x$. This increases the numerical stability of the calculation. We go on to use these string curvatures as part of our algorithm which tracks the motion of identified string loops in time.

² The non-uniform stencil involving $x_{i\pm 2}$ is only of third-order in contrast to the corresponding uniform stencil which is of fourth-

order. This is due to the lack of symmetry in s which allows higher order terms to cancel in the uniform case.

C. Tracking Strings Over Time

We track individual strings over time to confirm that loops collapse as expected in our QCD epoch simulation. To do so, we need to identify which string, parameterized as $\gamma(s, \eta)$ at time η , corresponds to which other string $\gamma(s, \eta + \Delta\eta)$ at time $\eta + \Delta\eta$. This is a non-trivial task due to frequent string reconnections and break-ups. Moreover, it is ambiguous to identify a string with a pre-existing one or as a new one after a major merger or break-up event.

Here, we choose to identify a string at time $\eta + \Delta\eta$ with a pre-existing string at η only if the difference in the string lengths between those times could have been realized by isolated string motion rather, *i.e.* characterized by a loop radius that grows at the speed of light. If the string length changes faster than this causal bound, it must be due to a major merger or break-up, and we identify the string at time $\eta + \Delta\eta$ as a new string. Specifically, while a circular loop can change its radius r by up to $r \rightarrow r \pm c\Delta\eta$ over a time interval of $\Delta\eta$, where c is the speed of light, our strings are not circular. However, using the local curvature radius R_c (see previous section) we can average over the relative radius change $\langle (R_c \pm c\Delta\eta)/R_c \rangle$ of each string segment. We use this average fractional change as an upper and lower bound on the allowed string length change. We find this criterion to work well in practice.

For a string at time η , we identify all string candidates at time $\eta + \Delta\eta$ within the allowed length range. To determine which candidate is the correct match, we pick up to 50 string segments at random and compute their minimum distance to each of the string candidates while ensuring proper treatment of periodic boundary conditions. If all minimum distances are closer than $c\Delta\eta$ it is considered a match. We only take a sample of 50 segments since it is sufficient to avoid false positives while not being too computationally demanding. If no match can be found we conclude the string must have been broken up significantly or merged with another string. Any string at time $\eta + \Delta\eta$ that has not been matched with another string is considered *new*. In our analysis, we only included strings with at least 20 string-plaquette piercings. Fig. S13 shows the resulting evolution of all such strings in our QCD epoch simulation, where we extract string locations every $\Delta\eta = 0.08$.

D. Measuring the Domain Wall Area

We extract the domain wall area analogously to the string length. This is done by counting links where the axion field value $\theta = \text{atan2}(\psi_2, \psi_1)$ wraps from π to $-\pi$ from one cell to another. The physical domain wall area is then $\mathcal{A}_{\text{dw}} = 2/3 N_{\text{link}}(\eta\Delta x)^2$, where the factor $2/3$ is the same over-counting correction factor that is also present when computing the string length [54]. Since domain walls are not necessarily refined to the finest level we perform this analysis self-consistently at the coarse level. We then define the domain wall area parameter analogously to the string length parameter ξ as

$$\xi_{\text{dw}} = \frac{\mathcal{A}_{\text{dw}} t}{\mathcal{V}}, \quad (\text{S6})$$

where $\mathcal{V} = (\eta L)^3$ is our physical simulation volume. An illustration of ξ_{dw} in our QCD epoch simulation is shown in Fig. S11.

4. SYSTEMATICS: IMPACT OF THE EMISSION SPECTRUM FITTING REGION

Our fiducial analysis, the analyses of [17, 47], and the analysis of [51] all differ in terms of the range of wavenumbers considered while fitting the instantaneous emission spectrum $F(k/H)$. In this section, we provide comprehensive results varying the lowest and highest wavenumbers in our fitting regions. We set the lowest wavenumber allowed in the fit by $k_{\text{IR}} = x_{\text{IR}}H$, where H is the physical Hubble rate. Similarly, we set the highest wavenumber allowed in the fit by $k_{\text{UV}} = x_{\text{UV}}m_r$. In our fiducial analysis, we take $x_{\text{IR}} = 50$ (as in Ref. [51]) and $x_{\text{UV}} = 1/16$, corresponding to an identical set of choices made in [49]. By comparison, [17, 47] consider x_{IR} as small as 30 and x_{UV} as large as $1/4$.

In Fig. S2, we inspect the impact of our choice of UV cutoff on our modeling of the emission spectrum index when all other details of the fiducial analysis are held fixed. Larger values of the UV cutoff are more aggressive as they extend the simple power-law modeling across a broader range, and, critically, to momenta near the radial mode mass, receiving nontrivial contributions from radial mode dynamics that could contaminate the spectrum. As we observe in Fig. S2, no appreciable evidence for a q growing linearly with $\log(m_r/H)$ is found until a UV cutoff at least as large as $m_r/8$. At larger UV cutoffs, the radial mode masking may in principle play a role in determining the spectrum, but its effect was found to be negligible in [49]. We inspect the emission spectra and associated best-fit indices for these large values of the UV cutoff in Fig. S3. We find increasing levels of systematic mismodeling that bias the best-fit index to a given instantaneous spectrum which grows in effect with increasing UV cutoff.




Figure S2. (Left) A comparison of the best-fit values for the index q of the instantaneous axion emission at $\log(m_r/H) = 9.5$ (black) and $\log(m_r/H) = 9$ (gray), corresponding to two of the latest times in our simulation. Shown are several choices of the UV cutoff, with the choice made in our fiducial analysis highlighted with a light red band. Smaller values of the UV cutoff are more conservative. (Right) Slope q_1 of the linear fit to the evolution of q . We observe no evidence for a linear increase of q with $\log(m_r/H)$ for a UV cutoff at or below $m_r/16$.




Figure S3. As in Fig. 3, but examining the instantaneous emission spectrum and time evolution of the emission spectrum index for large values of our UV cutoff. All other analysis choices are held fixed to those of our fiducial analysis. In the top panel, the data are downsampled by a factor of four for the purposes of visualization.

We note that our definition of the UV cutoff in an AMR context is somewhat more subtle than for a static lattice. At the end of our simulation, while the string width $\sqrt{2}m_r$ is resolved at our finest level of refinement by four lattice sites, it is under-resolved by a factor of eight at the base level where we measure the emission spectrum. As a result, the most aggressive choices of UV cutoff, like one at $m_r/4$, will admit wavenumbers up to nearly the Nyquist mode of our lattice, where discretization effects are certain to be important. To rectify this, we impose an additional constraint such that, independent of our UV cutoff, no momenta within a factor of four of the Nyquist momentum are included within our fitting range. This also ensures a more fair comparison of our $m_r/4$ analysis with that [17, 47], which takes a UV cutoff of $m_r/4$ but maintains at least unit resolution of the string core on their static lattice. We note that in both prior static lattice simulations and on the coarsest level of our adaptive mesh simulations, at late times, m_r is




Figure S4. As in Fig. S2, but comparing different choices of the IR cutoff.

Coefficient	$x_{\text{IR}} = 30$	$x_{\text{IR}} = \mathbf{50}$	$x_{\text{IR}} = 75$	$x_{\text{IR}} = 100$
q_1	0.04 ± 0.02	0.02 ± 0.02	0.00 ± 0.03	-0.03 ± 0.04
q_0	0.65 ± 0.16	0.87 ± 0.22	1.03 ± 0.01	1.31 ± 0.37
$q_0^{\text{const.}}$	1.01 ± 0.01	1.01 ± 0.01	1.02 ± 0.01	1.03 ± 0.01

Table S1. Results of the fits to the spectral evolution holding all our fiducial analysis choices fixed but for various IR cutoffs x_{IR} . We provide the fits and uncertainties for the q_1 and q_0 in the linearly growing index model and the best fit for $q_0^{\text{const.}}$ in the constant index model. Our fiducial choice of $x_{\text{IR}} = 50$ is shown in bold.

Coefficient	$x_{\text{UV}} = 1/4$	$x_{\text{UV}} = 1/6$	$x_{\text{UV}} = 1/8$	$x_{\text{UV}} = 1/12$	$x_{\text{UV}} = \mathbf{1/16}$	$x_{\text{UV}} = 1/24$	$x_{\text{UV}} = 1/28$	$x_{\text{UV}} = 1/32$
q_1	0.18 ± 0.01	0.11 ± 0.01	0.06 ± 0.01	0.04 ± 0.02	0.02 ± 0.02	0.03 ± 0.04	-0.01 ± 0.05	-0.02 ± 0.06
q_0	-0.34 ± 0.11	0.13 ± 0.11	0.49 ± 0.12	0.66 ± 0.16	0.87 ± 0.21	0.77 ± 0.33	1.10 ± 0.41	1.14 ± 0.50
$q_0^{\text{const.}}$	1.23 ± 0.01	1.13 ± 0.01	1.07 ± 0.01	1.03 ± 0.01	1.01 ± 0.01	1.00 ± 0.01	0.98 ± 0.02	0.98 ± 0.02

Table S2. As in Tab. S1, but for varying UV cutoff x_{UV} with all other parameters fixed to their fiducial values. Our fiducial choice of $x_{\text{UV}} = 16$ is shown in bold.

roughly at the Nyquist frequency, meaning that modes near the UV cutoff of $m_r/4$ are likely quite strongly affected by discretization choices. We do not expect these choices to strongly affect the data for our more conservative choices of UV cutoff.

Similarly, in Fig. S4, we study the results when all details of the fiducial analysis are held fixed with the exception of varying x_{IR} between $\{30, 50, 75, 100\}$. A lower choice of IR cutoff risks systematic mismodeling of the instantaneous emission spectrum by a simple power law model as low momentum axion emission is cut off by the interstring spacing. Additionally, the smaller density of states of low wavenumbers of our rectilinear lattice means lower wavenumber data is intrinsically noisier. Taken together, a larger value of the IR cutoff is a more conservative one. As we observe in Fig. S4, more aggressive smaller values of the IR cutoff result in estimated linear growth of the emission spectrum index, while more conservative larger values for the IR cutoff result in little to no evidence for a time-evolving index.

In Tab. S1 and Tab. S2, we provide the tabulated results from the systematic variations to the UV and IR cutoffs examined in this section. Other than the variations to those cutoffs, all details of the fiducial analysis are held fixed. We note that thanks to increased simulation volume and dynamical range relative to that of the previously state-of-the-art simulation in [49], fit uncertainties have been reduced at least threefold and, in the case of large x_{UV} , by as much as an order of magnitude.

We comment that Ref. [51] also finds that the spectral index extraction is sensitive to x_{UV} . In particular, Fig. [14] of that work illustrates that decreasing x_{UV} can somewhat flatten the log growth of the index and make it more consistent with the conformal hypothesis. On the other hand, while the general trend observed in [51] with changing x_{UV} is consistent with that we observe here, our results still disagree in detail with [51] (and also [47]).

5. SYSTEMATICS: PRE-EVOLUTION TO REMOVE TRANSIENTS

The results of this Letter differ appreciably from those of [17, 47, 51]. While none of these works implement AMR techniques to maintain the high resolution of strings, they also differ in their procedure of generating a string network for studying axion string emission. We consider the possibility here that these differences in the initialization procedure are the origin of the discrepant measurements of q across the works.

Our simulations begin with a thermal state before the breaking of the PQ symmetry (see [76] for procedure used to generate the thermal state), with strings realized dynamically as the shrinking thermal mass shifts the minimum of the scalar field potential. By contrast, [17, 47, 51] do not include a thermal mass for the PQ scalar, corresponding to starting their simulations deep in the broken phase of the theory well after strings should have first formed. It is unclear how to generate principled initial conditions for the complex scalar at this time, and so these works take an alternate approach of generating initial conditions appropriate for the PQ scalar in the unbroken phase and evolve the unphysical field dynamics under strong damping until strings form in the simulation and reach a target string length (typically corresponding to matching the claimed attractor solution of [17, 47] realized absent the external damping). This strategy is often referred to as *pre-evolution*.

We perform an AMR simulation initialized with pre-evolution and apply our analysis procedure to extract q , allowing us to directly assess the impact of the initialization procedure on the results presented in the main text. We perform our simulation in a box with size $\bar{L} = 96$ resolved by $2,048^3$ lattice sites at the base level; this enables a simulation out to $\log(m_r/H) = 8.5$. All other simulation parameters are the same as in our primary simulation. We generate an initial state of independent Gaussian noise at each lattice site in the ψ_1 and ψ_2 fields, and we evolve them with the equations of motion

$$\psi_i'' + \frac{3}{\eta}\psi_i' - \frac{\nabla\psi_i}{\eta^2\eta_i} + \frac{\eta_i^2}{\eta^2}\psi_i(|\psi|^2 - 1) = 0. \quad (\text{S7})$$

These equations of motion were developed in [24] to realize the desired additional damping by replacing $R(t) \propto t^{1/2}$ with $R(t) \propto t$ while also providing a tuneable parameter η_i such that the width of the string on the lattice as evolved under (S7) would match the width of the string on the lattice as evolved under the fiducial equations of motion at conformal time η_i . This then allows us to use the pre-evolved state at an arbitrary time η_{pre} as the initial condition for our standard simulations at time η_i . During this transition, we account for the change of physics by rescaling the field derivatives by η_{pre}/η_i .

We perform a simulation following this procedure, which we compare to our primary simulation performed without any pre-evolution. Note that this pre-evolved simulation uses a smaller box and has a more limited dynamical range, meaning that its individual statistical power is lesser. We choose $\eta_i \approx 3.77$, corresponding to $\log(m_r/H) = 3$, and we evolve under damping until the string length reaches its predicted value under the logarithmically growing attractor solution of $\xi = 0.3$ at this time [51]. After reaching the targeted string length, we transition to our standard equations of motion, taking the conformal time to be η_i . The pre-evolution phase ran for about 8 hours on 128 nodes of the Perlmutter GPU cluster, which provides 512 NVIDIA A100 (40 GB) GPUs, 8,192 CPU cores, and 66 TB of aggregate memory. The subsequent evolution ran on 256 Perlmutter GPU nodes for about 5 hours.

We first study the time-evolution of the string length in our pre-evolved simulation as compared with that found in our primary simulation and prior results in the literature [17, 47, 49, 51]. The results are shown in Fig. S5, where we find good consistency between the growth of the string network in our pre-evolved simulation and other works that also implement a pre-evolution procedure, namely [47, 51].

We go on to measure q and evaluate its possible time-evolving behavior in the pre-evolved simulation following the analysis procedure described in the main text. We compare the results of the linearly growing emission index shown in Fig. S6, finding results that are compatible at the level of the statistical errors. In general, we observe the same trend as in SM Sec. 4, with larger UV cutoffs for the fitting region resulting in more rapid growth of the emission spectrum index. Based on the good compatibility of the growing emission spectrum index between our primary simulation and our pre-evolved simulation, we conclude that our assessment of the systematic bias of the fit at large UV cutoffs is a robust one.

For completeness, in Tab. S3, we provide the best-fits and 1σ errors for the fit parameters to the constant and linearly growing emission index models for each of the variations of the UV cutoff considered here.

6. SYSTEMATICS: IMPACT OF RESOLUTION

By implementing AMR, our simulations realize a nearly constant (to within a factor of 2) resolution of strings, while simulations that use a static lattice are doomed to run out of resolution as the physical width of the string




Figure S5. Comparison of the string length per Hubble volume ξ between our main result (black, dash-dotted) and a simulation using a pre-evolved initial state tuned to be close to the scaling solution (data points). This plot illustrates consistency between our result and that of works from Gorghetto et al. [47] and Saikawa et al. [51], which also used a pre-evolved initial state.

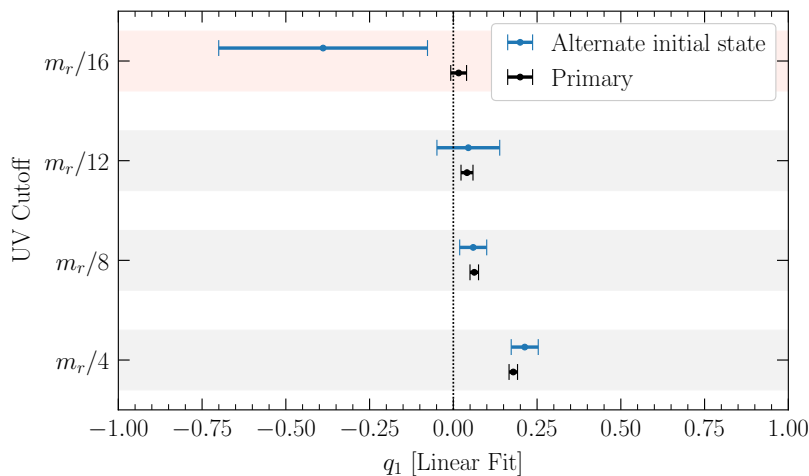


Figure S6. Comparison of spectrum linear fit slope parameter q_1 from our simulation using the alternate initial state procedure with pre-evolution (blue), to our primary large AMR simulation (black). Note that the simulation with an alternate state is the same as shown in Fig. S5. Here we assume the same IR cutoff as in our fiducial spectrum ($k_{\text{IR}}/H = 50$). The UV cutoff corresponding to our fiducial spectrum is shaded in light red. The numerical values of the fit indices are tabulated in Table S3.

Coefficient	$x_{\text{UV}} = 1/16$	$x_{\text{UV}} = 1/12$	$x_{\text{UV}} = 1/8$	$x_{\text{UV}} = 1/4$
q_1	0.05 ± 0.45 -0.39 ± 0.31	0.21 ± 0.09 0.04 ± 0.09	0.21 ± 0.05 0.06 ± 0.04	0.28 ± 0.03 0.21 ± 0.04
q_0	0.73 ± 3.58 4.09 ± 2.51	-0.56 ± 0.69 0.63 ± 0.74	-0.54 ± 0.35 0.53 ± 0.31	-0.95 ± 0.21 -0.52 ± 0.3
$q_0^{\text{const.}}$	1.15 ± 0.07 0.96 ± 0.04	1.05 ± 0.04 0.98 ± 0.03	1.07 ± 0.04 1.00 ± 0.01	1.13 ± 0.06 1.05 ± 0.05

Table S3. Fits and uncertainties for the q_1 and q_0 in the linearly growing index model and the best fit for $q_0^{\text{const.}}$ in the constant index model, for our simulations prepared via the alternate initial state procedure as described in SM Sec. 5 and SM Sec. 6. We perform these alternate simulations at two different grid resolutions, labeled ‘Low-res’ and ‘High-res’ in Fig. S9, corresponding to the top and bottom quantities in each cell, respectively. Note that these spectral fits are compared to our primary simulation in Figs. S6 and S9. Here we assume our fiducial choice of IR cutoff $k_{\text{IR}}/H = 50$ and vary the UV cutoff (with the fiducial choice bolded).

shrinks on the comoving lattice. For instance, in [17, 47, 51], the majority of simulations are run until the string

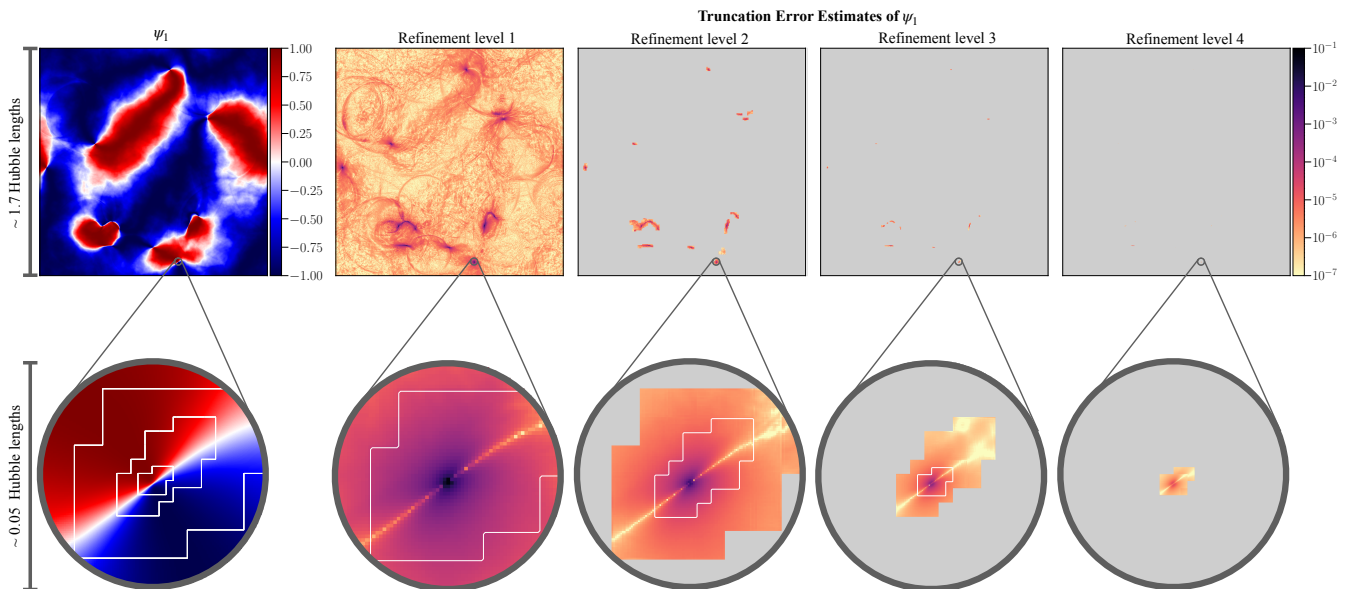


Figure S7. Illustration of truncation error estimates in the high-resolution simulation based on a pre-evolved initial state. The far-left panel is a slice through the ψ_1 field component with the other four panels containing the corresponding truncation error estimates at different refinement levels. The bottom row is a zoom-in centered around a string piercing the slice. Overlaid are coarse-fine boundaries (white lines). Areas in gray do not contain data as they are not covered by the respective refinement level.

width is resolved by a single lattice site. This corresponds to the η at which $1/(m_r \Delta x)$ is unity.

Not only does our AMR approach allow us to maintain resolution in a computationally efficient manner, but it also enables us to directly test the impact of resolution by loosening our string tagging criteria. We perform identical pre-evolved simulations to those studied in SM Sec. 5, but in our tagging criteria, we tag strings only at levels where they are expected to be resolved by one or fewer lattice sites. This means both simulations will have the same base resolution with identical initial states and only differ in the number of refinement levels. We refer to this simulation as our *low-resolution* simulation, while the previous simulation of SM Sec. 5 is referred to as our *high-resolution* simulation. Note that we chose to perform this comparative test between pre-evolved simulations to sharply inspect the resolution criteria in the existing literature. The low-resolution simulation ran around 25 times faster than the high-resolution simulation, requiring a total of 32 Perlmutter GPU nodes with 128 NVIDIA A100 (40 GB) GPUs for about 90 minutes.

We present truncation error estimates of the high-resolution simulation on a slice through the simulation volume in Fig. S7. Truncation error estimates correspond to the absolute difference between refinement level ℓ and $\ell - 1$ after evolving both levels independently for $\Delta\eta_{\ell-1} = 2\Delta\eta_\ell$. See SM Sec. 2B for a description of how truncation error estimates are computed and utilized in our simulation. From Fig. S7 it becomes clear that the majority of the simulation volume is captured much better at a given resolution than the string core, with errors several orders of magnitude smaller. At the first refinement level, the typical error does not exceed the 10^{-5} level away from strings, whereas the errors around the string core almost approach unity. For every extra refinement level, the error estimate drops by roughly an order of magnitude or more, eventually bringing it down to an acceptable level at the string core. This highlights how important resolution around the string core is and how AMR is the correct tool to simulate the network. Since we resolve the string core width by four grid sites in our primary simulation whereas other groups (e.g. [17, 47, 51]) often reduce this criterion to just one grid site per string core width, it implies that we effectively use two extra refinement levels. We can thus infer that our numerical error is roughly two orders of magnitude smaller.

We compare the time-evolving string length in the low-resolution and high-resolution simulation, with results shown in Fig. S8. We find modest discrepancies at the 2% level. The absolute difference is slowly increasing with time but with no clear trend in either direction.

As before, we measure q in the low-resolution simulation and compare our findings with those of the high-resolution simulation and our primary simulation. The results are depicted in Fig. S9 and provided in detail in Tab. S3. Strikingly, we find that low-resolution simulations systematically prefer a larger emission spectrum index in the case of the constant model and a more rapidly growing emission spectrum index in the logarithmically growing model.

One possible reason for the difference in the spectrum is illustrated in Fig. S15. Here, we performed two simulations




Figure S8. The relative difference between the high-resolution and low-resolution pre-evolved simulations. Only very small differences at the percent level are observed. See text for more details.




Figure S9. As in Fig. S6, but comparing two simulations at differing grid resolutions which are generated with the alternate initial state procedure (for details see text). The simulation labeled ‘High-res’ is the same as shown in Fig. S6. Note that we do not directly compare the constant fit parameter $q_0^{\text{const.}}$ between the primary simulation and those generated using the alternate initial state as the primary simulation differs in simulation volume from the latter.

of a single circular string in the QCD epoch, which is described in detail in Sec. 9. One simulation uses our high-resolution setup, the other is based on our low-resolution settings. Both simulations are otherwise identical. From Fig. S15 it becomes apparent that the low-resolution string is not sufficiently resolved to properly track its motion. The string collapses more slowly and loses its circular shape. Facets diagonally to the underlying grid are developing. This is due to inaccuracies in the finite-difference Laplacian caused by the limited resolution, which systematically reduces the speed of the string along the grid lines. The effect is large enough for visible kinks to develop that radiate extra energy. This extra artificial source of radiation is likely polluting the measurement of q . Even though this comparison was done in the QCD epoch, the same numerical arguments apply to PQ epoch simulations.

7. AXIONS THROUGH THE QCD PHASE TRANSITION

Understanding the fate of axions radiated by strings through the transient epoch during which the axion acquires its zero temperature mass is critical for accurately estimating the axion relic abundance. For example, it is possible that axions produced from the collapse of strings and domain walls during the QCD phase transition contribute appreciably to the DM relic abundance.




Figure S10. 3D \rightarrow 2D projection of the total axion energy density $\rho_a = \frac{1}{2}\dot{a}^2 + \frac{1}{2}(\nabla a)^2 + V(a)$ from a simulation of the QCD phase transition. From left to right: (1) Axion string network before the QCD phase transition. (2) The beginning of the phase transition. The axion mass starts growing and domain walls form. (3) The extra tension from the domain walls causes the network to collapse. (4) The QCD phase transition is complete. All strings have vanished and just a few oscillons are visible as bright point-like spots, though we do not investigate these features in this work. An animated version of this simulation can be found [here](#).

In this section, we evaluate the efficiency of strings and domain walls in producing axions at late times using analytic arguments calibrated against a simulation of a string-domain-wall network. The QCD-epoch simulation that we run for this analysis is described in App. C. The simulation takes an initial state containing axion strings with no axion mass and turns on a rapidly growing axion potential (with $N_{\text{dw}} = 1$), such that domain walls form and collapse the network. In Fig. S10 we illustrate a few snapshots of the simulation, showing the initial string network, the formation of domain walls bounded by these strings, and the subsequent collapse of the defect network under its tension. Oscillons, also referred to as axitons, appear in the final stages of the network collapse. We defer a more detailed study of them to future work.

It is useful to define a few important time periods during the epoch of the QCD phase transition to understand the dynamics of string-wall network collapse. First, we define t_* as the time at which the axion begins to oscillate when $m_a(t_*) = 3H(t_*)$. In our simulations, this time occurs at $\eta \approx 32$ or equivalently $\log(m_r/H) \approx 7.3$. Prior estimates of the axion abundance in *e.g.* [49] neglected axion emission by strings after this somewhat arbitrary time. More physically, strings will continue to emit axions, though possibly with reduced efficiency, until the total collapse of the defect network.

The next relevant time is t_{m_a} , which defines when the axion mass begins to play a role in determining the string emission spectrum. In particular, once $m_a \gg H$, the axion provides an IR cutoff to the string emission spectrum, suppressing emission at momenta $k < m_a$. To estimate this effect, let us assume a simple, analytic functional form for the instantaneous emission spectrum $F[k/H]$. In particular, note that for $t < t_*$ we may reproduce our fiducial result by approximating $F[k/H] \propto 1/(k/H)$ for $x_L < k/H < m_r/H$, where we take $x_L = \delta_1\sqrt{\xi}$, and zero otherwise. (Throughout this section we assume the strings radiate a conformal spectrum of axions; *i.e.*, we assume $q = 1$ for definiteness.) That is, we assume simply that $F[k/H]$ is a power-law for all k/H above some critical value x_L , which slowly moves to the UV like $\sqrt{\xi}$ as ξ grows with time. Now accounting for the time-dependent axion mass we simply use the above functional form but with $x_L = \max(\delta_1\sqrt{\xi}, m_a(t)/H)$. Then we have t_{m_a} defined as the time at which $m_a/H = \delta_1\sqrt{\xi}$. Note that from our primary simulation we measure $\delta_1 \approx 8.6$ and $\xi_* \in (11, 15)$. For concreteness in this section, we fix $\delta_1 = 8.6$ and $\xi_* = 13$, while $\xi_* \approx 0.8$ for the simulations which we study.

At later times, after t_{m_a} , the string network begins to collapse, and we define t_0 to be the time at which the string length parameter ξ begins to shrink. Finally, the collapse completes at t_{coll} . After this time, the axion field is (with the exception of some oscillons), to good approximation, linear, and no additional processes contribute to the axion abundance. In Fig. S11, we illustrate ξ as a function of time as measured in our QCD-epoch simulation, with the four critical times indicated. We also show the corresponding domain wall area parameter ξ_{dw} as defined in Sec. 3 D.

Given an axion mass parametrization $m_a(t)$ and string length parameter ξ , the times t_* and t_{m_a} are analytically calculable, whereas computing t_0 and t_{coll} is less straightforward. We begin by recognizing that, at t_* , in addition to abundant subhorizon-scale defects, there will also exist defects that are superhorizon-scale at this time. These superhorizon defects are comparatively less numerous but appear with nonzero probability due to the horizon-scale statistics associated with the axion field configuration via the Kibble mechanism, see, *e.g.* [82] for more details. Further evidence for the presence of these superhorizon defects in our simulations is presented in Sec. 9, particularly in Fig. S16 and the surrounding discussion.

The evolution of defects at times after t_* when the axion field becomes dynamic depends sharply on their size. A subhorizon string loop that bounds a domain wall will rapidly collapse (within roughly a Hubble time) under the joint




Figure S11. The string length parameter ξ and the domain wall area parameter ξ_{dW} during the QCD phase transition simulation. We indicate the time at which the axion mass becomes dynamical t_* , and the estimated time t_{m_a} when the IR cutoff of the emission spectrum becomes set by the axion mass rather than the inter-string separation distance. We further indicate the time t_0 when ξ begins to decrease due to the shrinking of string loops under domain wall tension, and when the network has collapsed completely, t_{coll} .

string and domain wall tension.³ By comparison, a string loop bounding a domain wall that is superhorizon at t_* will not begin to efficiently collapse until it enters the horizon, and in the intervening time, its physical size will grow due to the growth of the scale factor. As a result, in order to estimate axion production from the defect network around the time of the collapse, we treat subhorizon and superhorizon defects independently. Throughout this section, we make the simplifying assumption that string loops of physical length ℓ are circular and bound by a domain wall with radius $\ell/2\pi$. A defect is then taken to be superhorizon if its physical diameter d is greater than H^{-1} . We also assume, for simplicity, that domain walls within the horizon collapse at the speed of light, while domain walls outside the horizon are frozen in comoving coordinates.

A. Collapse of Subhorizon Defects

We begin with defects which are initially subhorizon at t_* . To put an upper bound on their axion production, we take ξ_* , the string length at t_* , to be fully accounted for by subhorizon defects. Our collapse ansatz specifies then that the physical radius of the defect evolves as

$$r(t) = r(t_*) - (t - t_*), \quad (\text{S8})$$

starting at t_* and ending at $t_{\text{end}} = r(t_*) + t_*$. We now consider the production efficiency of both these strings and their associated domain walls.

Emission from Subhorizon Strings

We assume that strings radiate axions with an emission spectrum $F(k, t) = 1/[k \log(k_{\text{UV}}/k_{\text{IR}}(t))]$ at k between an IR cutoff at $k_{\text{IR}}(t)$ and a UV cutoff of m_r . In our simplified collapse model, the rate at which a loop emits axions is given by

$$\frac{dE_a}{dt} = 2\pi^2 f_a^2 \log(m_r/H). \quad (\text{S9})$$

The instantaneous axion emission spectrum associated with a single loop is then given by

$$\frac{dN_a}{dt dk} = \frac{2\pi^2 f_a^2 \log(m_r/H)}{k \sqrt{k^2 + m_a(t)^2} \log(m_r/k_{\text{IR}}(t))} \theta(r(t)), \quad (\text{S10})$$

³ Note that this differs from the time-evolution of relatively large but still subhorizon string loops at early times, which do not

collapse rapidly absent the domain wall tension.

where θ is the Heaviside theta function. By integrating over k and t , we find the total number of axions emitted is

$$N_a(\ell) = \int_{t_*}^{t_* + \ell/2\pi} dt \int_{k_{\text{IR}}(t)}^{m_r} dk \frac{dN_a}{dt dk}, \quad (\text{S11})$$

where k_{IR} is set by either the inter-string spacing or the axion mass. That is,

$$k_{\text{IR}}(t) = \max \left[m_a(t), \delta_1 \sqrt{\xi_*} H(t) \right], \quad (\text{S12})$$

with turnover between the two cutoffs occurring at t_{m_a} .

Following *e.g.* [49], the number density of closed string loops was found to be $dn_\ell/d\ell \propto 1/\ell$; we expect this to still hold true at t_* when the axion mass has just entered the horizon but has not yet had time to affect the network dynamics. Then, subject to our assumption that all string loops at t_* are subhorizon (having a length less than $2\pi t_*$), the physical number density of strings at time t_* with length ℓ is given by

$$\frac{dn_\ell}{d\ell} = \frac{\xi_*}{2\pi t_*^3 \ell}. \quad (\text{S13})$$

We can evaluate the number density of subhorizon-string-emitted axions at $2t_*$, at which time all subhorizon defects have collapsed, by

$$\delta n_a^{\text{string,sub}}(2t_*) = \left[\frac{R(t_*)}{R(2t_*)} \right]^3 \int_0^{2\pi t_*} d\ell \frac{dn_\ell}{d\ell} N_a(\ell), \quad (\text{S14})$$

where the ratio of scale factors arises from redshifting the number density from t_* to $2t_*$. We compare this to the number density at $2t_*$ of axions emitted by strings at times before t_* , which is given by

$$n_a(2t_*) = \frac{8\pi f_a^2 \xi_*^{1/2} H(t_*)}{\delta_1} \left[\frac{R(t_*)}{R(2t_*)} \right]^3. \quad (\text{S15})$$

These number densities are useful, as from them we can compute the ratio of the energy density of axions produced by the collapse of subhorizon string loops to the number density of axions produced by string emission prior to t_* by

$$\frac{\delta \rho_a^{\text{string,sub}}}{\rho_a} \equiv \frac{\delta n_a^{\text{string,sub}}(2t_*)}{n_a(2t_*)} = \frac{\delta_1 \xi_*^{1/2}}{4t_*^2} \int_0^{2\pi t_*} \frac{d\ell}{\ell} \int_{t_*}^{t_* + \ell/2\pi} dt \frac{\log(m_r/H)}{\log(m_r/k_{\text{IR}})} \int_{k_{\text{IR}}}^{m_r} \frac{dk}{k \sqrt{k^2 + m_a^2}}. \quad (\text{S16})$$

In the physical hierarchy, this integral is easy to evaluate as all subhorizon loops will collapse by $2t_*$, while $t_{m_a} > t_*$. Hence, the axion mass never plays an important role in cutting off the string emission. In this case, we obtain

$$\frac{\delta \rho_a^{\text{string,sub}}}{\rho_a} = 0.65. \quad (\text{S17})$$

On the other hand, a similar case for the simulated hierarchy is somewhat more complicated as $t_{m_a} \approx 1.4t_*$, meaning that we must account for the axion mass in the IR cutoff of the string emission spectrum. Treating the cases of loop emission before and after t_{m_a} appropriately, we obtain

$$\frac{\delta \rho_a^{\text{string,sub}}}{\rho_a} = 0.75. \quad (\text{S18})$$

Hence, we see in both the simulated and physical hierarchies, we expect string emission after the time of axion oscillation to make an $\mathcal{O}(1)$ contribution to the late-time axion abundance.

Emission from Subhorizon Domain Walls

Next, we consider axion production from the collapse of domain walls that are subhorizon at time t_* . First, however, we describe the domain walls themselves in more detail. An infinitely wide and flat domain wall has the axion profile in the z -direction (perpendicular to the domain wall)

$$a(z) = 4f_a \tan^{-1} \exp(m_a z), \quad (\text{S19})$$

where the domain wall is at $z = 0$. The energy density in the domain wall is proportional to $|\partial_z a(z)|^2$. Let us now Fourier transform $\partial_z a(z)$ and take the magnitude squared of the result to gain insight into the k -modes available for the domain wall to produce upon its radiation. Denoting $\rho_{\text{dw}}(k)$ as the Fourier transform of the energy density in the domain wall in the z -direction, we estimate

$$\rho_{\text{dw}}(k) \propto \text{sech}^2\left(\frac{k\pi}{2m_a}\right). \quad (\text{S20})$$

By inspection, $\rho_{\text{dw}}(k)$ peaks at $k \lesssim m_a$, suggesting that most axions emitted from domain wall decay are nonrelativistic or mildly semi-relativistic, in which case their energies are approximately m_a . In the interest of placing an upper bound on axion production from domain wall collapse, we assume below that domain walls emit zero-momentum axions. In Sec. 8 we provide additional evidence that domain walls radiate semi-relativistically by simulating the collapse of a single circular string and domain wall. However, we emphasize that the assumption that the domain walls radiate nonrelativistically is a crucial one that should be checked more rigorously in future work, as it plays a key role in determining the DM abundance from domain wall decay.

Now, consider at time t_* a single string loop of physical length ℓ which bounds a domain wall of physical radius $r = \ell/2\pi$. The energy of the domain wall is given by

$$E(t) = 8\pi r(t)^2 m_a f_a^2. \quad (\text{S21})$$

Then, the rate at which some number N_a axions are emitted by the shrinking domain wall is given by

$$\frac{dN_a}{dt} = 16\pi r(t) f_a^2, \quad (\text{S22})$$

where we crucially assume all of the radiated axions are nonrelativistic with energy equal to m_a . Integrating this expression to the time when the loop radius has shrunk to zero, we find a total of $N_a(\ell) = 2\ell^2 f_a^2/\pi$ axions emitted by the shrinking domain wall. As in the case of strings, we integrate this emission over the loop distribution to evaluate the number density of these axions at time $2t_*$ as

$$\delta n_a^{\text{dw,sub}} = \left[\frac{R(t_*)}{R(2t_*)}\right]^3 \int_0^{2\pi t_*} d\ell \frac{dn_\ell}{d\ell} N_a(\ell). \quad (\text{S23})$$

This yields the result

$$\frac{\delta \rho_a^{\text{dw,sub}}}{\rho_a} = \frac{\delta_1}{2\pi} \xi_*^{1/2}. \quad (\text{S24})$$

Hence, we find that for our simulations, where $\xi_* \approx 0.8$ and $\delta_1 \approx 8.6$, we expect the collapse of subhorizon domain walls to contribute roughly equally to the string emission. In the physical case, the collapse of subhorizon domain walls may be as much as 5 times more efficient in producing axions as compared to string emission prior to t_* . (See Tab. S4.)

B. Collapse of Superhorizon Defects

We now proceed to consider the case of superhorizon defects. For simplicity, we treat all superhorizon defects as having a single characteristic size. While a spectrum of superhorizon defect sizes might be physically expected, larger and larger defects grow exponentially unlikely, and so we expect quantities associated with the spectrum of superhorizon defects to be dominated by a single typical scale.

Since these superhorizon domain walls are initially frozen out, their physical radius evolves as $r(t) \propto t^{1/2}$ until a time t_0 defined by when $r(t_0) = t_0$; after this time they are within the horizon and they collapse relativistically. We parametrize this physical radius as

$$r(t) \approx \begin{cases} r_{\text{dw}}\left(\frac{t}{t_*}\right)^{1/2} & \text{if } t_* < t < t_0 \\ r_{\text{dw}}\left(\frac{t_0}{t_*}\right)^{1/2} - (t - t_0) = 2t_0 - t & \text{if } t_0 < t < t_{\text{coll}} \\ 0 & \text{if } t_{\text{coll}} < t, \end{cases} \quad (\text{S25})$$

where r_{dw} is the physical radius of the domain wall at time t_* . The time t_0 is then defined by

$$r_{\text{dw}} \left(\frac{t_0}{t_*} \right)^{1/2} = t_0. \quad (\text{S26})$$

We determine r_{dw} via the observed ratio of t_0/t_* in our QCD-epoch simulation, finding $r_{\text{dw}} \approx 1.7t_*$. Moreover, we can determine the number density of these superhorizon defects from the value $\xi_0 \approx 1.2$, the string length parameter at time t_0 . The physical number density at t_0 , of these superhorizon defects is then

$$\frac{dn}{d\ell} = \frac{\xi_0}{2\pi t_0^3} \delta(\ell - 2\pi t_0). \quad (\text{S27})$$

Further, while we have heavily calibrated against our QCD-epoch simulation, we note that this model for superhorizon defects makes two nontrivial predictions. First, at times shortly before t_0 , the string length parameter ξ should grow. This is because the frozen-out superhorizon defects, which are coming to dominate the total string length, drive the growth of ξ linear in t . Additionally, our relativistic collapse model also predicts that these superhorizon defects collapse by $t_{\text{coll}} = 2t_0$. Both these predictions are seen in Fig. S11.

Emission from Superhorizon Strings

Now we consider the emission of axions from the collapse of initially superhorizon strings. Since $r_{\text{dw}} \approx 1.7t_*$, we have $t_0 \approx 2.9t_*$. At this time, whether in the reduced hierarchy of the simulation or in the physical scenario, the axion mass always plays the dominant role in setting the IR cutoff on the emission spectrum. Analogous to the case of subhorizon string emission, we can compute the number density of axions produced by emission from strings that are superhorizon at time t_* . These strings emit at times beginning at t_0 when they enter the horizon until $2t_0$ when the collapse has completed. We thus compute the axion number density at $2t_0$ as

$$\delta n_a^{\text{string,sup}}(2t_0) = \frac{\pi f_a^2 \xi_0}{t_0^3} \left[\frac{R(t_0)}{R(2t_0)} \right]^3 \int_{t_0}^{2t_0} \frac{\log(m_r/H)}{\log(m_r/k_{\text{IR}}(t))} \int_{m_a}^{m_r} \frac{dk}{k\sqrt{k^2 + m_a^2}}, \quad (\text{S28})$$

convolving the strings' emission spectrum with the delta function number density of strings. As before, we compare the abundance of these axions to the abundance of axions from emission prior to t_* , given by

$$n_a(2t_0) = \frac{8\pi f_a^2 \xi_*^{1/2} H(t_*)}{\delta_1} \left[\frac{R(t_*)}{R(2t_0)} \right]^3, \quad (\text{S29})$$

after redshifting to time $2t_0$. Taking $n = 6.68$, $\xi_* = 0.8$, and $\xi_0 = 1.2$ to calculate this quantity for our simulation, we find

$$\frac{\delta \rho_a^{\text{string,sup}}}{\rho_a} = 0.06, \quad (\text{S30})$$

while in the physical extrapolation, with $\xi_* \approx 13$ and $n = 8.16$, we have

$$\frac{\delta \rho_a^{\text{string,sup}}}{\rho_a} = 0.007. \quad (\text{S31})$$

These results encode that at late times, the rapidly growing axion mass renders string emission highly inefficient.

Emission from Superhorizon Domain Walls

We now determine the number of axions emitted by the collapse of the superhorizon domain walls. The calculation is identical to that of the subhorizon domain walls except that the collapse process starts at t_0 rather than t_* . We then have

$$N_a = 16\pi f_a^2 \int_{t_0}^{2t_0} r(t) = 8\pi f_a^2 t_0^2, \quad (\text{S32})$$

Scenario	t_{m_a}/t_*	t_0/t_*	t_{coll}/t_*	$\delta\rho^{\text{string,sub}}/\rho_a$	$\delta\rho^{\text{string,sup}}/\rho_a$	$\delta\rho^{\text{dw,sub}}/\rho_a$	$\delta\rho^{\text{dw,sup}}/\rho_a$	$\delta\rho^{\text{tot}}/\rho_a$
Simulation	1.4	2.9	5.8	0.75	0.06	1.2	6.3	8.3
Physical	2.16	2.9	5.8	0.65	0.007	4.9	1.6	7.2

Table S4. Summary of our analytic estimates for the energy densities in axions produced at times $t > t_*$ from strings and domain walls, relative to the energy density ρ_a in axions emitted at times $t < t_*$. We split these into contributions from string loops/walls that are subhorizon at the oscillation time t_* and those which are superhorizon (denoted by superscripts ‘sup’, ‘sub’, respectively). The parameters in the first (second) row correspond to our QCD epoch simulation (physical scenario). We measure $\xi_* = 0.8$ and $\log(m_r/H_*) = 7.34$ in our QCD epoch simulation and fix $\xi_* = 13$ and $\log(m_r/H_*) = 70$ in the physical scenario.

yielding

$$\frac{\delta\rho_a^{\text{dw,sup}}}{\rho_a} \approx 0.65\delta_1\xi_*^{-1/2} \quad (\text{S33})$$

when comparing both evaluated at $t = 2t_0$. In our simulations, where $\xi_* = 0.8$, this suggests the collapse of initially superhorizon domain walls outproduces string emission by a factor ~ 6.3 . For the physical scenario of $\xi_* \approx 13$, this falls to approximately 1.6. (See Tab. S4.)

C. Summary of axion production in the QCD epoch

We summarize our results for the efficiency of axion production at times after t_* in Tab. S4. In general, we find that domain walls may make an appreciable or even dominant contribution to the late-time axion abundance as compared to string emission. This is primarily because domain walls may radiate axions into low-momentum modes, though this assumption should be checked more carefully in future work (see also Sec. 8).

However, we qualify this in several ways. First, we have assumed domain walls radiate zero-momentum axions, while radiating semi-relativistic axions may reduce their production efficiency by an $\mathcal{O}(1)$ factor. Moreover, we have assumed string loops that bound domain walls are circular, which has the effect of maximizing the area at a fixed string length. Non-circular configurations would reduce the domain wall energy available for axion emission, leading to a further reduced production efficiency. As a result, the domain wall production efficiency we have estimated in this section should be considered an upper bound and represents an important direction for future study.

We go on to study the efficiency of axion-production after t_* by comparing the differential number density spectrum $\partial n_a/\partial k$ in our QCD-epoch simulation with the equivalent spectrum obtained from the string-only simulation of [49]. Since these two simulations are identical up to the inclusion of the QCD potential, this provides the opportunity to sharply and self-consistently examine the estimates made in this section. To extract the number density spectrum from our QCD simulations, we use

$$\frac{\partial\rho_a}{\partial k} = \frac{1}{2} \left[\frac{\partial}{\partial k} |\dot{\tilde{a}}(k)|^2 + \frac{1}{2} (k^2 + m_a(\eta)^2) \frac{\partial}{\partial k} |\tilde{a}(k)|^2 \right] \approx \frac{\partial}{\partial k} |\dot{\tilde{a}}(k)|^2, \quad (\text{S34})$$

which we relate to a physical number density by

$$\frac{\partial n_a}{\partial k} = \frac{1}{\sqrt{k^2 + m_a(\eta)^2}} \frac{\partial\rho_a}{\partial k}. \quad (\text{S35})$$

In practice, it is convenient to work with the differential spectrum of comoving density by comoving momentum, which is constant as a function of η when there is no active axion production.

In Fig. S12, we compare these differential comoving number densities between the two simulations. By integrating these comoving number densities, we find the comoving axion abundance at t_{coll} in our QCD simulation to be roughly 4.5 times larger than the axion abundance at t_* in our equivalent PQ simulation. Integrating only modes with $k_{\text{com}} > 10a(t_*)H(t_*)$ to avoid standard misalignment, which dominantly contributes to modes at or below the comoving horizon at t_* , given by $a(t_*)H(t_*)$, this ratio falls to 3.5.

By comparison, based on a scale-invariant emission spectrum, we expect the strings to radiate axions with an efficiency that enhances the axion production by only a total factor of ~ 1.8 , suggesting that domain walls are contributing a roughly equal amount to the late time abundance of axions as strings. This is less than our simple estimate of domain walls making a contribution as much as 7.5 times larger than string emissions prior to t_* under the assumption that they emit nonrelativistic axions. This suggests that domain walls are emitting somewhat relativistic axions with

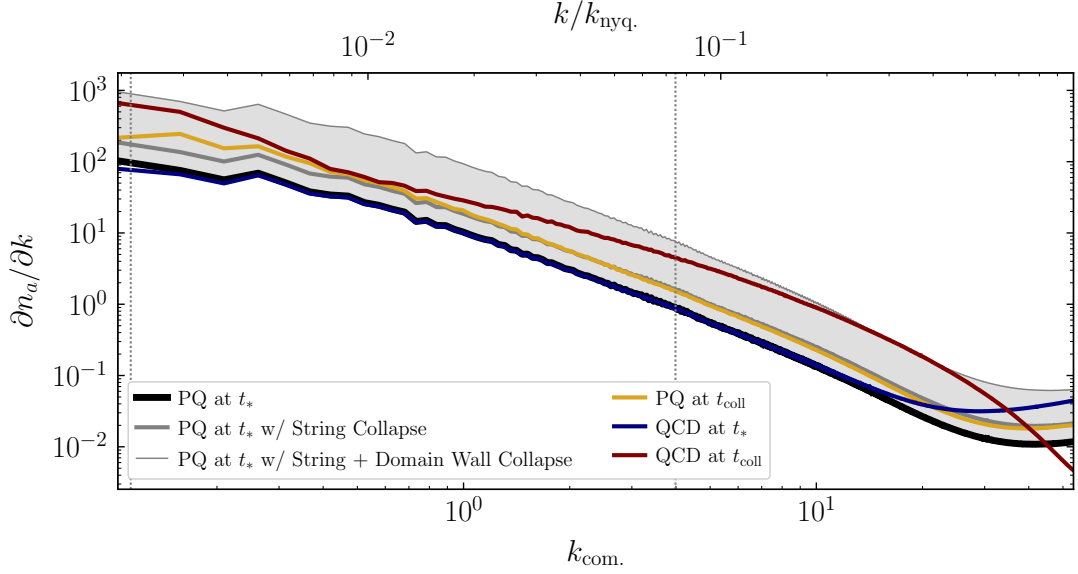


Figure S12. The number density spectra realized in our QCD epoch simulation and the otherwise identical string-only simulation of [49]. The left (right) vertical dotted gray lines indicate the axion mass at time t_* ($t_{\text{coll.}}$) in these comoving units, *i.e.*, they are the values $a(\eta)m_a(\eta)$. The heavy black line, corresponding to the string-only simulation of [49] at time t_* , and the navy line, corresponding to the QCD epoch simulation of this work at time t_* , are highly similar, with some mild discrepancy at large k . This is to be expected, as the axion mass has only recently begun to play an important role in the dynamics at time t_* . At t_{coll} , when all defects have collapsed and axion production has ceased, the number density spectrum realized by our simulation is shown in maroon. Based on our parametric estimates for axion production in this section, we present two additional spectra. The heavy gray line indicates the spectrum obtained by enhancing the spectrum from the PQ simulation at t_* by a factor of 1.81. This factor represents the estimated production of axions by the populations of subhorizon and superhorizon strings at times between t_* and t_{coll} . The lighter gray line indicates the spectrum obtained by enhancing the spectrum from the PQ simulation at t_* by a factor of 9.3. This factor represents the estimated production of axions by both strings and domain walls at times between t_* and t_{coll} . By comparison, the number density spectra at t_{coll} in our PQ simulation realized by undisturbed string emission is shown in yellow. Note that our calculation of the number of axions produced by domain walls assumes a primarily nonrelativistic spectrum, so applying this enhancement factor to the PQ spectrum is not strictly correct but is useful in determining where axion production is most efficient during the times between t_* and t_{coll} . The heavy black and light gray lines derived from the PQ-simulation spectrum at time t_* are used to define the lightly-shaded gray band.

$\mathcal{O}(1)$ boost factors. As we will see in later sections, we see further evidence that axions emitted by domain walls are at least somewhat relativistic. However, even if domain walls are emitting axions with $\mathcal{O}(1)$ boost factors, our estimates when extrapolated to the physical hierarchy still indicate that axion emission by domain walls will contribute nearly half of their late time abundance.

8. COLLAPSE OF A STRING-DOMAIN WALL LOOP

We study in further detail the collapse of a loop of string bounding a domain wall in an expanding background. Let us consider at time η , a string loop bounding a domain wall that has comoving radius r . We parametrize the spacetime location of a domain wall element with (η, R, ϕ) via the general map to the four-vector

$$\begin{pmatrix} \eta \\ R \\ \phi \end{pmatrix} \rightarrow \begin{pmatrix} \frac{\eta^2}{2f_a} \\ Rr(\eta) \cos \phi \\ Rr(\eta) \sin \phi \\ 0 \end{pmatrix}, \quad (\text{S36})$$

where ϕ is the azimuthal coordinate taking values in $[0, 2\pi)$, $r(\eta)$ is the maximum comoving radius of the domain wall at time η , and R is the relative location on the domain wall taking values in $[0, 1]$. (Note that R should not be confused with the scale factor.) This parametrization specifies an induced metric g_{dw} on the domain wall, from which

we derive the domain wall action

$$S_{\text{dw}} = - \int d\tau \sigma(\tau) \int_0^1 dR \int_0^{2\pi} d\phi \sqrt{|g|} = - \int d\tau [8\pi f_a^2 \eta^3 m_a(\eta) r(\eta)^2]. \quad (\text{S37})$$

The associated string parametrization is

$$\begin{pmatrix} \eta \\ \phi \end{pmatrix} \rightarrow \begin{pmatrix} \frac{\eta^2}{2f_a} \\ r(\eta) \cos \phi \\ r(\eta) \sin \phi \\ 0 \end{pmatrix}, \quad (\text{S38})$$

from which we obtain the string action

$$S_{\text{string}} = - \int d\tau \mu(\tau) \int_0^{2\pi} d\phi \sqrt{|g|} = - \int d\tau \left[2\pi^2 f_a^2 \eta^2 \log\left(\frac{m_r}{H}\right) r(\eta) \sqrt{1 - r'(\eta)^2} \right]. \quad (\text{S39})$$

For simplicity, we will approximate $\log(m_r/H)$ to be constant. From this action, we may derive the equation of motion for $r(\eta)$ with the Euler-Lagrange equations, obtaining

$$r'' = \frac{(\eta + 2rr') (r'^2 - 1)}{\eta r} - \frac{8\eta m_a(\eta) (1 - r'^2)^{3/2}}{\pi \log(m_r/H)}. \quad (\text{S40})$$

For a circular loop of string, the comoving radius $r(\eta)$ is related to the string length parameter by $\xi = \pi \eta^2 r(\eta) / 2H_1^2 L^3$ where L is the comoving box length. We compare the ξ associated with the numerical solution of (S40) to the string evolution measured in our QCD epoch simulation in Fig. S13. Specifically, we identify strings and track their individual contributions to ξ with the i^{th} string contributing ξ_i and the total string length parameter being $\xi = \sum_i \xi_i$.

We observe subhorizon strings collapsing on relatively short time scales whereas superhorizon strings decay very little until t_0 besides frequent reconnections. After t_0 , the superhorizon loops undergo a rapid decay driven by the domain wall tension. This quick decay is often accompanied by the shedding of smaller loops near kinks, which can be visually seen by the simultaneous emergence of many subhorizon loops. While these effects are not captured by (S40), the collapse of individual strings is fairly well described by our simple isolated loop model. See also the top left panel of Fig. S16.

We also note that a key feature of our analytic treatment of collapsing strings and domain walls at times was the assumption of defects that were superhorizon at t_* . These superhorizon defects are clearly visible in Fig. S16 as those string trajectories that lie above the dashed blue line defining the horizon diameter for circular string loops. These large loops demonstrate the expected $t^{1/2}$ growth until times as late as roughly t_0 , at which time they begin to collapse efficiently.

9. SIMULATION OF A SINGLE AXION STRING AND DOMAIN WALL

To improve our understanding of axion emission in the QCD epoch, in particular the emission from domain walls, we perform a simulation of a single axion string that comes to bound a collapsing domain wall. To do so we engineer an initial state that will result in a perfectly circular string by setting

$$\psi_1(\mathbf{x}) = \frac{2}{1 + e^{-(|\mathbf{x} - N/2| - N/4)/20}} - 1, \quad \psi_2(\mathbf{x}) = \sin\left(2\pi \frac{k}{N}\right), \quad (\text{S41})$$

with N the number of coarse-level cells in each spatial direction, and $\mathbf{x} = (i, j, k)^T$ the index space with $i, j, k = \{0, \dots, N-1\}$. The diameter of the string is exactly half of the simulation volume at the beginning of the simulation, and the domain wall forms on the inside of the string. The simulation starts as a normal PQ phase simulation with an identical setup to our primary simulation with $\bar{L} = 200$. Since we are dealing with an isolated string, however, we can reduce the coarse level to $1,024^3$ grid cells compensated by a larger number of refinement levels. To form a domain wall we adiabatically turn on the axion mass using the same parameterization as in our QCD simulation of (4) with $N = 19$, $\eta_* = 50$, and $n = 6.68$. The adiabatic period is described by a logistic function

$$m_a(\eta) \rightarrow \frac{1}{1 + e^{-f(\eta - \eta_0)}} m_a(\eta = 35), \quad (\text{S42})$$

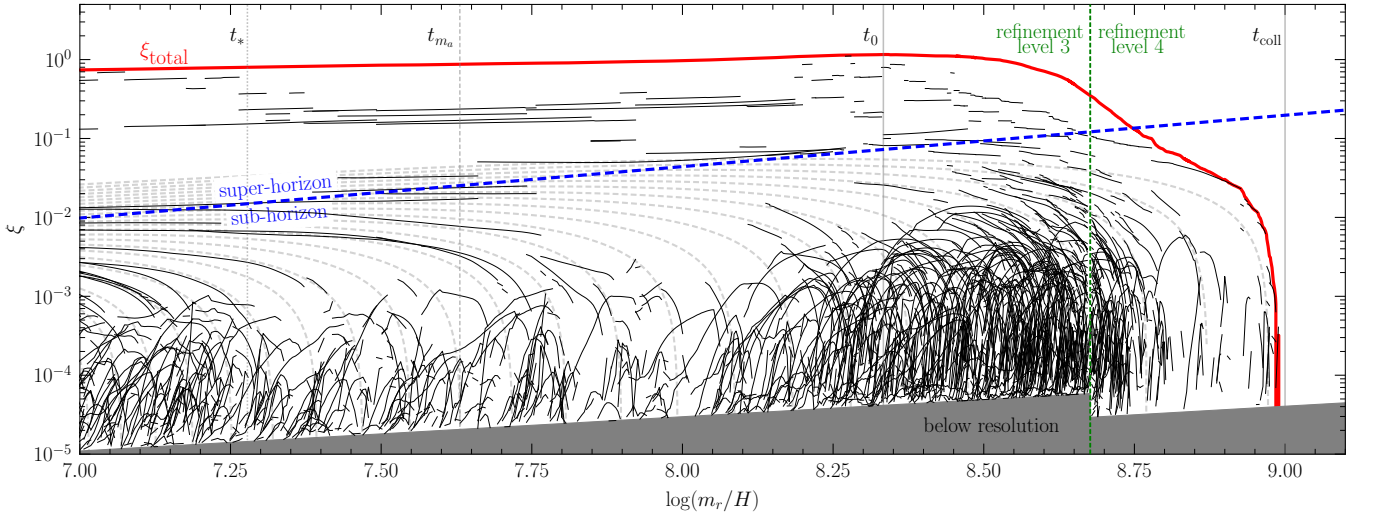


Figure S13. String length evolution in our QCD epoch simulation. **Black lines show the lengths of individual string segments, with the total string length indicated in red.** The black lines often stop or appear suddenly due to string reconnections and break-ups. The results are overlaid with gray dashed lines of a few examples of the string collapse model described in the text. The threshold between superhorizon and subhorizon loops is indicated by the dashed blue line and the simulation transitions from three to four refinement levels at the vertical green dashed line. Only strings with at least twenty string-plaquette piercings are taken into account as indicated by the grayed-out area at the bottom of the plot. The reference times t_* , t_{m_a} , t_0 , and t_{coll} are highlighted as well.

with $f = 3$ and $\eta_0 = 1$. This implies that beyond this brief adiabatic regime, the axion mass remains constant. We check that neither the exact initial state setup nor the exact parameterization of the adiabatic regime significantly affects the result. Due to its smaller size, we were able to run this simulation on GPUs. We used 128 nodes of the Perlmutter GPU cluster, which provides 512 NVIDIA A100 (40 GB) GPUs, 8,192 CPU cores, and 66 TB of aggregate memory. The total run time was around 3.5 hours.

We extract the axion emission spectra in the usual way but additionally mask out the domain wall by setting $\dot{a}^2 \rightarrow 0$ within the planar slab with width $6/m_a$ centered on the domain wall, which has characteristic width $1/m_a$. This masking region is illustrated in Fig. S14. A top-down projection is shown in Fig. S15 for two simulations with different resolution criteria. We use the high-resolution simulation in our analysis.

One means of examining the axion emission spectrum from domain walls is to compare the number density of axions realized in our simulation with the number density one would predict associated with the shrinking of the domain wall given an assumed emission spectrum. Consider at some time t_0 a domain wall of physical radius $r(t_0)$. Then the energy associated with that domain wall is

$$E_{\text{wall}} = 8\pi m_a f_a^2 r(t_0)^2. \quad (\text{S43})$$

If the wall were not contracting then its energy would grow like $a(t)^2 \propto (t/t_0)$. We define the free wall energy

$$E_{\text{wall}}^{\text{free}}(t|r(t_0), t_0) = 8\pi m_a f_a^2 r(t_0)^2 \left(\frac{t}{t_0} \right), \quad (\text{S44})$$

leading to an instantaneous rate of change in the domain wall energy of

$$\frac{\partial E_{\text{wall}}^{\text{free}}}{\partial t} = \frac{8\pi m_a f_a^2 r(t_0)^2}{t_0}. \quad (\text{S45})$$

On the other hand, from our simulation, we numerically extract $r(t)$ and $\dot{r}(t)$, so that the change in the wall energy is

$$\frac{\partial E_{\text{wall}}^{\text{sim}}}{\partial t} = 16\pi m_a f_a^2 r(t_0) \dot{r}(t_0). \quad (\text{S46})$$

The energy loss computed in our simulation is given by the difference of the emission into axions less the free expansion energy gain. Hence, from our simulation, we can compute the emission into axions as

$$\frac{\partial E_a(t)}{\partial t} = -\frac{\partial E_{\text{wall}}^{\text{free}}}{\partial t} - \frac{\partial E_{\text{wall}}^{\text{sim}}}{\partial t} = 16\pi m_a f_a^2 r(t) \left[\frac{r(t)}{2t} - \dot{r}(t) \right]. \quad (\text{S47})$$

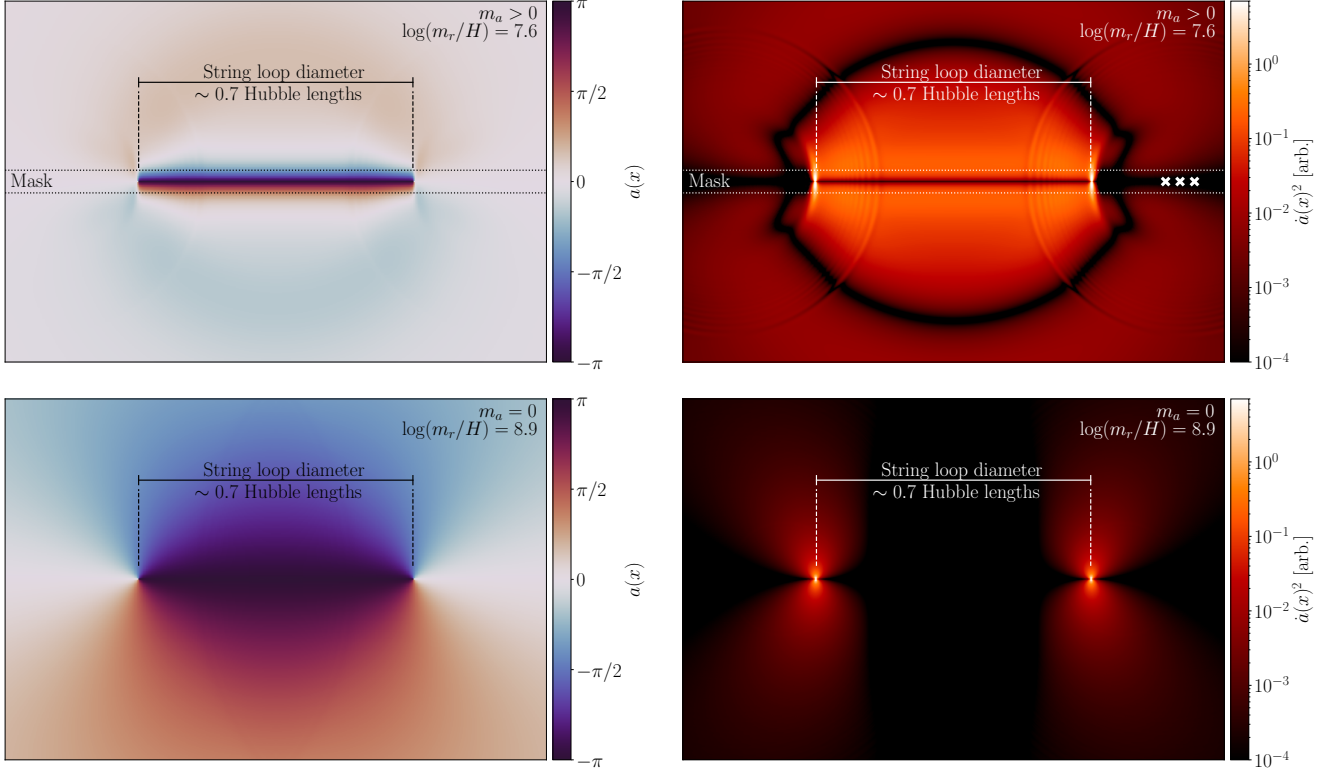


Figure S14. Cut through the axion field $a(x)$ (left column) and the axion kinetic energy $\dot{a}(x)^2$ (right column) of a single circular string in a simulation with constant non-zero m_a (top row) and $m_a = 0$ (bottom row). The string is piercing the slice twice at a distance of roughly 0.7 Hubble lengths. A cut through the domain wall is visible in the top row. The dotted horizontal lines indicate the mask encompassing the plane of the domain wall. Both simulations are based on the same initial state. Note that the string in the $m_a > 0$ case reaches a radius of 0.7 Hubble lengths more quickly due to the extra domain wall tension. As a result, the top row is a snapshot at an earlier time with $\log(m_r/H) \sim 7.6$ and the bottom row at $\log(m_r/H) \sim 8.9$.

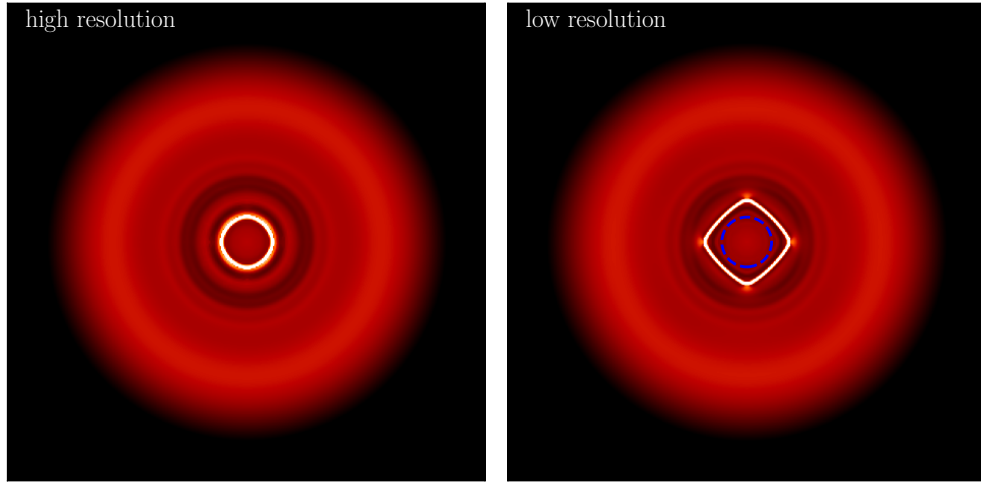


Figure S15. Comparison of a single circular string in the QCD epoch between a high-resolution setup (left) and a low-resolution setup (right) shortly before the strings final collapse at $\log(m_r/H) \sim 8.7$. The 3D \rightarrow 2D projection of $\dot{a}(x)^2$ on a logarithmic color scale is shown. Each panel spans roughly 0.9 Hubble lengths. The contour of the high-resolution string is overlaid in blue in the low-resolution panel. In the low-resolution simulation, the string collapses more slowly and loses its circular shape with facets developing diagonally to the underlying grid. This introduces artificial kinks that produce extra radiation.




Figure S16. (*Top Left*) The time evolution of the physical radius in our single string loop simulations. In solid black, we depict the physical radius of the string loop measured directly in the simulation, which we compare to the evolution of the string assuming free expansion with Hubble flow (gray) and inverse Hubble which defines the horizon size (dotted gray). When the radius of the string loop is much larger than the horizon, the size of the loop grows with η and only begins to contract efficiently once $r_{\text{phys}}H < 1$. We also compare the observed loop evolution with that expected by solving (S40) as shown in red, which demonstrates qualitatively good agreement. (*Top Right*) The instantaneous rate of axion emission per logarithmic interval was obtained via numerical evaluation of (S47). We see that the rate of axion emission is expected to peak around $\log(m_r/H) \approx 8.2$. (*Bottom*) A comparison of the total number density of axions in our string loop simulation (black) as compared to the prediction from integrating axion emission associated with a contracting string loop under the assumption of a nonrelativistic emission spectrum. By multiplying by a factor of η^3 , we compare comoving number density, which is constant at times when axions are not produced efficiently. The integrated emission number density exceeds that observed in our simulations, suggesting that the emission of axions from a shrinking domain wall is at least quasi-relativistic. See text for details.

Assuming that axions emitted by the wall collapse are nonrelativistic, the instantaneous rate of axion emission is

$$\frac{\partial N_a}{\partial t} = 8\pi f_a^2 r(t) \left[\frac{r(t)}{t} - 2\dot{r}(t) \right]. \quad (\text{S48})$$

We evaluate this numerically and divide it by the physical volume of our simulation to determine the physical number density n_a . We then compare this integrated axion emission to the number density realized in our simulations, with the results shown in Fig. S16.

It is clear from the results in Fig. S16 that the majority of the axions in our simulation volume were not produced by domain wall emission. Namely, the time-dependence of the number produced by domain wall emission does not match the time dependence of axions in the simulation, which saturates much earlier. The majority of axions produced in this simulation are likely generated by spurious modes realized in the initial conditions or the transient as the axion enters the horizon. On the other hand, the number of axions produced by domain wall collapse under the assumption of emission into nonrelativistic modes would nonetheless overproduce the observed axion number density. This suggests that the axions which are emitted from the domain wall must be at least somewhat relativistic as a harder emission spectrum would produce a smaller axion number density.

Another means of examining the axion production in the domain wall collapse is to study the typical momentum with which axions are emitted. To compute this, as in the main text, we evaluate $F(k/H)$ at each time in our wall collapse simulation and compute $\langle k/H \rangle = (\int dk(k/H)F(k/H))/(\int dkF(k/H))$. For comparison we also perform a string-only simulation, *i.e.*, with $m_a = 0$, using an identical initial condition, and similarly calculate $\langle k/H \rangle$. The results are shown in Fig. S17. Broadly, $\langle k/H \rangle$ is reduced in the domain wall simulation with $m_a > 0$ as compared to the string-only simulation with $m_a = 0$ and is typically less than the axion mass. This suggests that axions produced in our domain wall simulations are typically non- or semi-relativistic, though we caution that these results could be somewhat contaminated by unphysical modes produced from transient effects of the initial state.




Figure S17. A comparison of the typical momentum of axion emission computed for two simulations using an identical initial state, but with one simulating $m_a = 0$ (gray) while the other simulates $m_a > 0$ (black), such that a domain wall forms and collapses by the end of the run. In dotted gray, we show the ratio of the axion mass used in our $m_a \neq 0$ simulation to the Hubble expansion rate.

10. SIMULATING AXION NUMBER NON-CONSERVATION DURING QCD PHASE TRANSITION

The cosine potential experienced by the axion is an intrinsically nonlinear one that supports number density-changing processes. We consider the role of these processes on axions emitted by strings by performing an axion-only simulation with the equations of motion

$$\theta'' + \frac{2}{\eta}\theta' - \bar{\nabla}^2\theta + \frac{\eta^2 m_a(\eta)^2}{f_a^2} \sin\theta = 0, \quad (\text{S49})$$

where θ is a compact field varying between 0 and 2π related to the axion field a by $\theta \equiv a/f_a$. Note we have otherwise maintained a consistent set of simulation units and axion mass parametrization from other simulations considered in this work.

We generate initial conditions at $\eta \approx 33$ when the axion begins to oscillate in a box of comoving sidelength $\bar{L} \approx 66$ such that the simulation contains 8 Hubble volumes at the time of oscillation. We take the axion to oscillate at $\log(m_r/H_{\text{osc}}) \approx 65$, and assume an emission history between $\log(m_r/H) = 1$ and $\log(m_r/H_{\text{osc}})$ described by a normalized instantaneous emission spectrum of the form

$$F(k/H, t) \propto \frac{1}{k/H(t)}, \quad (\text{S50})$$

with support at k/H between $6\xi^{1/2}$ and m_r/H . We adopt our logarithmically growing fit to describe the time evolution of ξ . We then obtain the total energy in axions by integrating this emission history as

$$\frac{\partial \rho_a}{\partial k}(t) = \int dt' \frac{\Gamma(t')}{H(t')} \left(\frac{R(t')}{R(t)} \right)^3 F \left[\frac{kR(t)}{R(t')H(t')}, t' \right], \quad (\text{S51})$$

where $\Gamma(t) = 8\pi H(t)^3 f_a^2 \xi(t) \log(m_r/H(t))$ is the string decay rate. We may relate the energy density spectrum to the axion and axion derivative spectrum via (S34) assuming energy is shared equally between the kinetic term set by $|\dot{a}|^2$ and the gradient and mass energy term set by $(k^2 + m_a^2)|\bar{a}|$. However, we allow the phases of the $a(k)$ and $\dot{a}(k)$ to be random and uncorrelated.

We generate our initial conditions on a 512^3 lattice without adaptive meshing, zeroing out support in the momentum spectrum of $a(k)$ and $\dot{a}(k)$ for momentum within a factor of three of the Nyquist momentum. This ensures that our simulations are initially well-resolved. We then simulate from $\eta_i \approx 33$ to $\eta = 76$, corresponding to the end time of our full QCD simulation after defects have fully collapsed. Note that because we have generated initial conditions directly from the axion field, domain walls but not strings may be realized in our field configuration.

From our simulation, we calculate the comoving axion number density via (S35), including both field and field derivative terms in $\partial \rho_a / \partial k$. If number density is indeed conserved over these simulations, then we expect this




Figure S18. (*Left*) The time-evolution of the comoving axion number density relative to the initial axion number density in three simulations of an identical initial state at three choices of lattice resolution. Perfect number conservation would be expected to result in a constant line at 1, while decreases in the comoving number density are indicative of number density loss, associated with either discretization error or number-changing physical processes. For more details, see text. (*Right*) The inverse expectation value $\langle (H/k)^2 \rangle^{-1}$ (data points) of the axion emission spectrum F for our fiducial spectrum with $k_{\text{IR}}/H = 50$ (plotted in Fig. 3). Here we assume the spectral index is at its upper limit at 1σ , $q = 1.02$. We fit $\langle (H/k)^2 \rangle^{-1} = \beta\xi$ for a constant β (best fit model in solid), finding $\beta = 3226.2 \pm 65.2$. We fix $\log(m_r/H_{\text{QCD}}) = 70$.

comoving number density to be constant. In the left panel of Fig. S18, we depict the ratio of the comoving number density evaluated in the simulation to the initial comoving number density. The results using the 512^3 lattice are shown in gray, where we observe $\mathcal{O}(10\%)$ loss in the axion number by the end of our simulation. To test the origin of this loss, we upsample the initial state to a 1024^3 lattice, finding the red line, which exhibits only 1% loss in the number density. We then perform one additional simulation with a 2048^3 lattice, which achieves a resolution comparable to the coarse-level resolution of our network collapse simulation. The number density evolution is shown in black and reveals entirely negligible loss.

We conclude that the nonlinear potential is inefficient in changing the number density of axions emitted by strings. However, these simulations establish that appreciable spatial/temporal resolution is an important criterion for ensuring convergence, even in the axion-only scenario.

Following the method of Ref. [49] we also directly verify that nonlinearities in the axion field are small in our primary simulation. We fit $\langle (H/k)^2 \rangle^{-1} = \beta\xi$ for a constant β (see the right panel of Fig. S18) and estimate the nonlinearity of the axion field $\langle (a/f_a)^2 \rangle \approx (4\pi/\beta) \log(m_r/H) \leq 0.27$ for $\log(m_r/H) \leq 70$, consistent with our measurement of number density loss at the sub-percent level.