

Emergent equilibrium in many-body optical bistabilityM. Foss-Feig,^{1,2,3} P. Niroula,^{2,4} J. T. Young,² M. Hafezi,^{2,5} A. V. Gorshkov,^{2,3} R. M. Wilson,⁶ and M. F. Maghrebi^{2,3,7}¹*United States Army Research Laboratory, Adelphi, Maryland 20783, USA*²*Joint Quantum Institute, NIST and University of Maryland, College Park, Maryland 20742, USA*³*Joint Center for Quantum Information and Computer Science, NIST and University of Maryland, College Park, Maryland 20742, USA*⁴*Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*⁵*Department of Electrical and Computer Engineering and Institute for Research in Electronics and Applied Physics, University of Maryland, College Park, Maryland 20742, USA*⁶*Department of Physics, United States Naval Academy, Annapolis, Maryland 21402, USA*⁷*Department of Physics and Astronomy, Michigan State University, East Lansing, Michigan 48824, USA*

(Received 3 November 2016; revised manuscript received 10 January 2017; published 17 April 2017)

Many-body systems constructed of quantum-optical building blocks can now be realized in experimental platforms ranging from exciton-polariton fluids to ultracold Rydberg gases, establishing a fascinating interface between traditional many-body physics and the driven-dissipative, nonequilibrium setting of cavity QED. At this interface, the standard techniques and intuitions of both fields are called into question, obscuring issues as fundamental as the role of fluctuations, dimensionality, and symmetry on the nature of collective behavior and phase transitions. Here, we study the driven-dissipative Bose-Hubbard model, a minimal description of numerous atomic, optical, and solid-state systems in which particle loss is countered by coherent driving. Despite being a lattice version of optical bistability, a foundational and patently *nonequilibrium* model of cavity QED, the steady state possesses an emergent *equilibrium* description in terms of a classical Ising model. We establish this picture by making new connections between traditional techniques from many-body physics (functional integrals) and quantum optics (the system-size expansion). To lowest order in a controlled expansion—organized around the experimentally relevant limit of weak interactions—the full quantum dynamics reduces to nonequilibrium Langevin equations, which support a phase transition described by model A of the Hohenberg-Halperin classification. Numerical simulations of the Langevin equations corroborate this picture, revealing that canonical behavior associated with the Ising model manifests readily in simple experimental observables.

DOI: [10.1103/PhysRevA.95.043826](https://doi.org/10.1103/PhysRevA.95.043826)**I. INTRODUCTION**

While systems described by cavity quantum electrodynamics (QED) often contain many interacting degrees of freedom, they are unconventional from the standpoint of traditional many-body physics for two primary reasons. First, the mediation of interactions through a small number of delocalized cavity modes generally leads to extremely long-ranged interactions [1], which suppress the role of fluctuations and often enable accurate mean-field descriptions. In this sense they are simpler than conventional solid-state realizations of many-body physics, in which short-range interactions promote both quantum and thermal fluctuations to an important role, especially in low spatial dimensions [2,3]. Second, cavity-QED systems are typically driven and dissipative; as a result, even if they reach a time-independent steady state they will generally not be in thermal equilibrium [4]. In this sense they are more complicated than conventional solid-state realizations of many-body physics, in which coupling to a thermal reservoir is typically assumed and well justified, leaving the system in thermal equilibrium and enabling the powerful tools of statistical mechanics to be employed [5].

In recent years, experimental advances in quantum optics have begun to blur the first of these distinctions [6–9], with platforms including exciton-polariton fluids in semiconductor quantum wells [10–15], circuit QED [16–20], optical fibers, waveguides, and photonic crystals [21–26], small-mode-volume optical resonators [27,28], and Rydberg ensembles [29–32] all making progress towards realizing large-scale arrays of short-range coupled quantum-optical building blocks.

These developments have led many researchers to revisit fundamental questions surrounding the fate of nonequilibrium quantum-optical systems in situations where, due to the importance of either dissipative or quantum fluctuations, a mean-field description is insufficient [7,15,33–44]. The primary goal of this paper is to elucidate the physics of a canonical many-body model made relevant by these developments—the driven-dissipative Bose-Hubbard model [45–48]—which furnishes a minimal description of, e.g., coherently driven exciton-polariton fluids confined in coupled microcavities or other patterned semiconductor devices [14,15,49–51]. In that context, the weak-coupling limit of the model is by far the most experimentally relevant. However, as will be discussed, neither perturbation theory nor mean-field theory are sufficient for capturing the weak-coupling physics, even qualitatively, nor can it be inferred from the well-studied equilibrium physics of the Bose-Hubbard model.

Nevertheless, by exploiting connections between ideas from many-body physics (functional-integral treatments of nonequilibrium field theory) and quantum optics (phase-space techniques and the system-size expansion), we identify a *quantitatively accurate* mapping of the weak-coupling limit of the driven-dissipative Bose-Hubbard model onto a field theory governing the *equilibrium* physics of the finite-temperature classical Ising model. In this way (cf. our main results in Sec. V), we are able to make quantitative predictions for the steady-state phases and phase transitions of the model in a parameter regime that is highly relevant to ongoing experiments with exciton-polariton fluids [14,15].

Much of the previous work on the driven-dissipative Bose-Hubbard model has grown out of early proposals to simulate the equilibrium Bose-Hubbard model in photonic systems, either in the transient regime of very *weakly* dissipative systems [45,52], or through clever strategies to mitigate the effects of particle loss [53–55]. In this context, the driven-dissipative model has been considered in an attempt to understand the corruption of equilibrium physics by nonvanishing dissipation in realistic systems, and to identify qualitative signatures of equilibrium Bose-Hubbard physics—e.g., fermionization for strong interactions [46,56] or the incompressibility of the zero-temperature Mott-insulating phase [57–59]—that survive in steady state. The general spirit of this approach is to start from the intuitions and expectations appropriate for the equilibrium Bose-Hubbard model, and to build outward toward an understanding of the driven-dissipative dynamics; numerous interesting connections to the equilibrium physics of the Bose-Hubbard model [34,46,60], as well as a variety of surprising and genuinely nonequilibrium effects [17,46,61–64], have been discovered in this manner. But there are many reasons to expect that the search for *universal* features of the driven-dissipative model benefits from, and perhaps even requires, a fundamentally different approach. For example, the ground-state and thermal phase transitions of the Bose-Hubbard model are intimately related to U(1) symmetry and the associated particle-number conservation [65]. While the former can be preserved in a driven-dissipative context by pumping the cavities incoherently [6], the latter remains absent [66], calling into question whether—in the sense of phase transitions and universality—*any* properties of the Bose-Hubbard model can survive the presence of driving and dissipation [42].

Here, we instead pursue an understanding of the driven-dissipative Bose-Hubbard model from the ground up, starting from the well-understood nonequilibrium physics of a single cavity, namely optical bistability [67], and adopting a functional-integral formalism that is well suited to extending the essential single-cavity physics to a many-body setting [7,34,39,68]. As we show, the breaking of conservation laws and symmetries at the microscopic level leads to universal properties of the steady state that bear essentially no resemblance to the equilibrium physics of the Bose-Hubbard model [57], but neither do they retain the fundamentally nonequilibrium character of optical bistability. Instead, the steady state of the driven-dissipative Bose-Hubbard model admits an emergent equilibrium description in terms of a finite-temperature classical Ising model [69]. Specifically (see Fig. 1), two collective mean-field steady states are inherited from the optical bistability of the individual cavities; they play the role of the two local minima in the Ising model’s mean-field free energy, while dissipation (i.e., vacuum fluctuations) plays the role of thermal fluctuations, setting the effective temperature. By explicitly and quantitatively connecting two canonical and minimal models of many-body physics—one a cornerstone of nonequilibrium quantum optics and one a cornerstone of traditional equilibrium many-body physics—this paper provides a particularly simple and concrete example of the way in which equilibrium can emerge very naturally from an *a priori* nonequilibrium many-body problem, even when (a) mean-field theory fails and (b) the model with respect

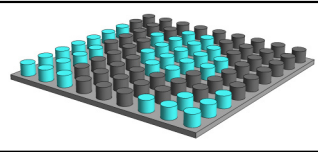
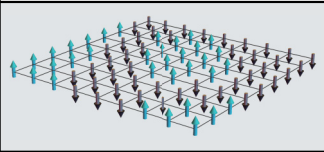
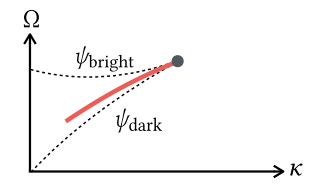
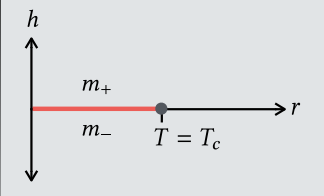
Optical bistability	Ising model
	
Degrees of freedom	
$\psi = \langle \hat{a} \rangle$	$m = \langle \sigma \rangle$ ($\sigma = \pm 1$)
Mean-field theory	
$i\dot{\psi} = -(\mu + i\kappa/2 + U \psi ^2)\psi + \Omega$	$\mathcal{F}(m) = rm^2 + gm^4 + hm$
Important control parameters	
(κ, Ω)	$(r \sim T - T_c, h)$
Mean-field solution	
$(\dot{\psi} \rightarrow 0) \psi = \psi_{\text{dark}}, \psi_{\text{bright}}$	$(\mathcal{F} \rightarrow \min) m = m_+, m_-$
Phase diagram	
	
Temperature	
$\sim \kappa/\mathcal{N}$	T

FIG. 1. Summary of the correspondence between many-body optical bistability and the classical Ising model that serves as its effective equilibrium description. The two possible magnetizations of the Ising model correspond to the bright and dark mean-field steady states of the optically bistable cavities. All parameters and variables are defined in the manuscript. In the bottom row, \mathcal{N} is a parameter controlling the density scale, i.e., $|\psi|^2 \sim \mathcal{N}$. Hence the low-temperature limit of the Ising model corresponds to a semiclassical (large density) limit of optical bistability.

to which equilibrium emerges is not simply connected to the equilibrium physics of the underlying Hamiltonian.

We note that our conclusions are enabled by a unification of ideas from nonequilibrium field theory (originally applied to a similar model in Ref. [69]) and quantum optics. In particular, we identify a small parameter $1/\mathcal{N}$ related to the inverse “system size” (in the sense of the system-size expansion often employed in quantum-optics studies of systems with only a few degrees of freedom), which controls the overall scale of fluctuations, and thus the effective temperature. In the limit of weak fluctuations, the qualitative predictions of Ref. [69] can be justified, made quantitative, and even verified numerically. In this way, we not only identify the model with respect to which an effective thermal description emerges, but also semiquantitatively obtain the phase boundaries, effective temperature, and near-critical dynamics in terms of microscopic parameters. Crucially, the limit in which these methods are accurate—weak coupling and large density—coincides with the limit of the model that is most accessible experimentally. Thus, in addition to establishing a universal perspective on

the physics of the driven-dissipative Bose-Hubbard model, the techniques established here can be used to make quantitative predictions for ongoing experiments with exciton-polariton fluids.

Before proceeding, we caution that the emergence of an effective equilibrium description as detailed in this paper, while potentially reasonably generic, should not be taken for granted; other more genuinely nonequilibrium situations can and do arise in other models [35,39,41,70–72]. Ultimately, the goal of this paper is not only to provide a detailed view into the mechanisms by which thermal equilibrium can emerge from the microscopically nonequilibrium setting of many-body cavity QED, nor by any means to insist that one must emerge, but also to establish and clarify deep connections between many-body physics and quantum optics that may elucidate more unusual behaviors made possible by the strong-coupling regime of quantum optics. It is natural to expect that analyses of prototypical situations in which equilibrium *does* emerge will play an important role in anticipating more exotic situations in which it does *not*.

After presenting the model and reviewing the well-known solution of the single-cavity case in Sec. II, our general strategy for the many-body problem is laid out in Sec. III. Our approach is based on a well-established exact mapping of the master equation onto a functional integral. Although it imposes some additional notational burden, the functional-integral formalism has the virtue of (1) being a convenient starting point for the identification of approximation schemes, including controlled strategies for going beyond mean-field theory [68], and (2) enabling powerful techniques such as the renormalization group to be applied [7]. Here, we use the functional integral to quickly identify an exactly solvable limit of the problem, around which a semiclassical expansion (related to the system-size expansion of quantum optics) can be made. To leading nontrivial order in this expansion, we obtain a quantitatively accurate mapping of the many-body quantum master equation onto classical nonequilibrium Langevin equations, with a small parameter controlling the strength of the noise. In Sec. IV we analyze the mean-field equations of motion near the mean-field critical point, which possess an emergent Z_2 symmetry in the spirit of Ref. [73]. We show how the complex order parameter decomposes into two real components, one of which stays massive at the critical point and one of which does not. By adiabatically eliminating the massive component, we arrive at a time-dependent Landau-Ginsburg equation for a scalar field, which supports two different homogeneous solutions within the bistable region. Near the critical point and inside the bistable region, we are able to analytically obtain the profile and velocity of domain walls separating domains of these two different phases, and the vanishing of the domain-wall velocity gives a zeroth-order approximation to the location of a true (first-order) phase transition in more than one spatial dimension. In Sec. V we consider the effects of fluctuations in both one and two spatial dimensions by (a) arguing that—near the critical point and for weak noise—the nonlinear Langevin equations become equivalent to model A of the Hohenberg-Halperin classification, and (b) solving the nonequilibrium Langevin equations numerically, which is valid even away from the critical point. As our earlier analysis would suggest, the numerical results are qualitatively consistent with the

expected equilibrium physics of a classical Ising model in a longitudinal field. In one dimension, domains are seeded by fluctuations, and the dynamics of their unbound domain walls smooths the mean-field transition into a crossover. In two dimensions the domain walls exhibit a surface tension, enabling a line of true first-order phase transitions terminating at a critical point.

II. MODEL

The model we consider can arise in a variety of contexts, but for concreteness we consider either a one-dimensional (1D) chain or a two-dimensional (2D) rectangular array of semiconductor microcavities supporting exciton-polaritons (see, e.g., Ref. [14]). We assume that the on-site energies of exciton-polaritons are spatially uniform and equal to ω_0 , and that the cavities are driven coherently and in phase by a laser with frequency ω_L . Upon making a unitary transformation to remove the time dependence of the driving, we obtain the Hamiltonian

$$\hat{H} = -J \sum_{\langle j,k \rangle} \hat{a}_j^\dagger \hat{a}_k - \delta \sum_j \hat{a}_j^\dagger \hat{a}_j + \frac{U}{2} \sum_j \hat{a}_j^\dagger \hat{a}_j^\dagger \hat{a}_j \hat{a}_j + \Omega \sum_j (\hat{a}_j + \hat{a}_j^\dagger). \quad (1)$$

Here, \hat{a}_j^\dagger (\hat{a}_j) creates (annihilates) an exciton-polariton in the j th cavity, J parametrizes the strength of a coherent coupling of exciton-polaritons between neighboring cavities, $\delta = \omega_L - \omega_0$ is the detuning of the laser from cavity resonance, U sets the two-body interaction energy for exciton-polaritons confined in the same cavity, and Ω is the amplitude of the coherent driving. The notation $\langle j,k \rangle$ implies that the sum should be taken over all nearest-neighbor pairs of sites j and k . The driving is necessary to stabilize a nontrivial steady state in the presence of particle loss out of the cavities at a rate κ . If the loss of exciton-polaritons is treated in the Born-Markov approximation, the dynamics of the combined unitary evolution under \hat{H} and loss is described by a Markovian master equation [4],

$$\frac{d\hat{\rho}}{dt} = -i[\hat{H}, \hat{\rho}] + \frac{\kappa}{2} \sum_j (2\hat{a}_j \hat{\rho} \hat{a}_j^\dagger - \hat{\rho} \hat{a}_j^\dagger \hat{a}_j - \hat{a}_j^\dagger \hat{a}_j \hat{\rho}). \quad (2)$$

More generally, Eqs. (1) and (2) provide a natural (though certainly not unique) generalization of the Bose-Hubbard model to the driven (Ω) and dissipative (κ) setting of quantum optics.

The case of a single cavity has been thoroughly studied in the quantum-optics literature, where it serves as a minimal model for dispersive optical bistability [67,74]. A mean-field description of the problem can be obtained by writing down the equation of motion for $\psi \equiv \langle \hat{a} \rangle$ and assuming that expectation values of normal-ordered operator products factorize (i.e., making the replacement $\langle \hat{a}^\dagger \hat{a} \hat{a} \rangle \rightarrow |\psi|^2 \psi$), giving

$$i\dot{\psi} = -(\delta + i\kappa/2)\psi + U|\psi|^2\psi + \Omega. \quad (3)$$

The steady-state equation $\dot{\psi} = 0$ can be recast as a cubic equation for the mean-field density $n = |\psi|^2$,

$$n[(\delta - Un)^2 + \kappa^2/4] = \Omega^2. \quad (4)$$

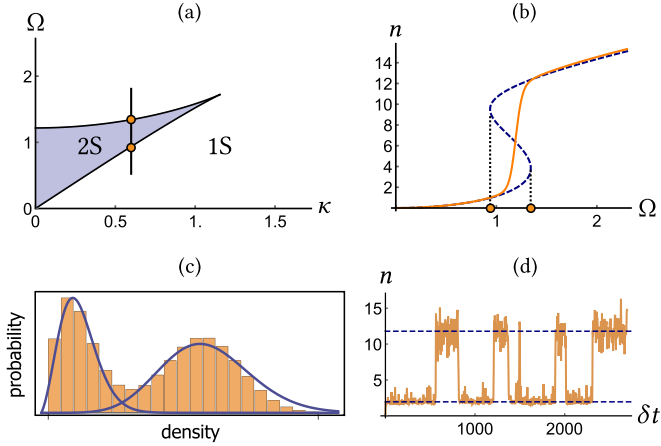


FIG. 2. (a) The mean-field phase diagram for a single cavity is divided into regions that support either one or two dynamically stable solutions. For this plot and those that follow, all parameter values are given in units of δ , and $U = 0.1$. (b) $\kappa = 0.6$: A cut through the bistable region of the mean-field phase diagram, showing both the mean-field (dashed blue curve) and exact (solid orange curve) solution for the density. (c) $\kappa = 0.6$ and $\Omega = 1.2$: Full counting statistics of the exact steady-state density matrix, together with that of both mean-field solutions (i.e., coherent-state distributions, with the relative normalization used as a fitting parameter). (d) $\kappa = 0.6$ and $\Omega = 1.2$: One trajectory obtained from a quantum-trajectories simulation of Eq. (2) for a single cavity, showing switching between two mean-field-like states; the densities associated with the two dynamically stable mean-field solution are shown as dashed lines.

This equation has either one or two solutions that are dynamically stable to small perturbations, leading to the mean-field phase diagram shown in Fig. 2(a).

For a single cavity, the full quantum solution of the master equation can be obtained efficiently in a variety of ways, for example by direct numerical integration of Eq. (2) within a truncated Hilbert space. Steady-state expectation values can even be obtained analytically by mapping the single-cavity version of Eq. (2) onto phase-space equations in the complex-P representation [67,75]. Reference [67] provides a comprehensive discussion of the solution, and here we simply summarize its main features, focusing primarily on the relationship between the exact and mean-field solutions. While the mean-field equations of motion can support two dynamically stable steady states, the exact steady-state density matrix of Eq. (2) is always unique, as are all observables calculated from it [for example, see Fig. 2(b)]. While there are never two truly stable steady states, two important signatures of mean-field bistability do survive in the limit of large cavity occupancy [76,77]: (1) The full-counting statistics of the exact solution exhibits a bimodal structure within the parameter regime yielding mean-field bistability, with the probability of observing different photon numbers clustering around the two mean-field stable values of the density [Fig. 2(c)]. Outside of the bistable region, this bimodality disappears and the exact counting statistics becomes similar to that corresponding to the one stable mean-field solution; in this sense, the exact solution interpolates between the two mean-field steady states within the bistable region. (2) If the system is initialized in

one of the mean-field steady states, it will only explore the phase space in the vicinity of that solution on the natural time scales of the problem (i.e., those associated with energy scales appearing explicitly in the master equation), and will only sample the phase space in the vicinity of the other mean-field solution on much longer time scales [Fig. 2(d)]. This slow time-scale for switching between mean-field-like steady states is associated with a small gap of the exact quantum Liouvillian, which vanishes inside the bistable region in the limit of large photon occupancy [48,78,79]. In this limit, which plays a role analogous to the “thermodynamic limit” of a spatially extended system [80], mean-field bistability can therefore be identified with the existence of a (zero-dimensional) dissipative phase transition.

III. FUNCTIONAL-INTEGRAL FORMULATION AND THE SYSTEM-SIZE EXPANSION

To analyze the steady-state behavior of many coupled cavities, it is convenient to recast the master equation in terms of an equivalent functional integral [81],

$$\mathcal{Z} = \int \mathcal{D}\psi(t) \mathcal{D}\varphi(t) \mathcal{W}(\psi_0, t_0) e^{i\mathcal{S}}, \quad (5)$$

with action

$$\begin{aligned} \mathcal{S} = & 2i \sum_j \int_{t_0}^{\infty} dt (\bar{\varphi} \partial_t \psi - \varphi \partial_t \bar{\psi}) \\ & - \sum_j \int_{t_0}^{\infty} dt (\mathcal{H}_w(\psi + \varphi) - \mathcal{H}_w(\psi - \varphi)) \\ & + i\kappa \sum_j \int_{t_0}^{\infty} dt (2\bar{\varphi}\varphi - \varphi\bar{\psi} + \bar{\varphi}\psi). \end{aligned} \quad (6)$$

The functional integral in Eq. (5) is over all unconstrained paths for the variables $\psi_j(t)$ and $\varphi_j(t)$, the spatiotemporal dependence of which has been suppressed in Eqs. (5) and (6) [82]. The factor $\mathcal{W}(\psi_0, t_0)$ in Eq. (5) is the Wigner function at the initial time t_0 , with ψ_0 being shorthand for the set of field variables $\psi_j(0)$ at the initial time. In this paper we only concern ourselves with steady-state properties, and so we can safely set $t_0 = -\infty$ in the integral limits of Eq. (6) and exclude the dependence on $\mathcal{W}(\psi_0, -\infty)$ from \mathcal{Z} . From here forward the integral limits of $\pm\infty$ in the action are implied and not shown.

The classical Hamiltonian \mathcal{H}_w in Eq. (6) is the Weyl symbol of the Hamiltonian in Eq. (1) [83,84]; this should be contrasted with the usual appearance of the Q symbol in the Keldysh functional integral. The appearance of the Weyl symbol is a consequence of the way in which Eqs. (5) and (6) are derived. In particular, we do not apply the usual Keldysh rotation to a coherent-state path integral, but rather provide a direct construction of \mathcal{Z} in terms of the variables φ and ψ (see Appendix A for a detailed derivation). This constructive approach has the benefit of elucidating the deep connections between the functional-integral treatment of open quantum systems—which is widely employed, and reviewed in Ref. [7]—and more traditional phase-space methods. These differences notwithstanding, the present approach is perfectly consistent with the appearance of the Q function in the usual

Keldysh approach; the difference is compensated by a slightly different set of rules for computing equal-time observables.

While the master equation can be used to compute arbitrary averages of time-dependent system operators via the quantum-regression formula [4,85], the functional integral in Eq. (5) is suited to calculating products of operators that are ordered along a closed-time (Keldysh) contour [86,87]. It is worth noting, however, that the restriction to calculating a particular class of time-ordered operators is not required in order to formulate the problem via a functional integral; rather, it enables the functional integral to be formulated along a particularly simple (Keldysh) time contour. The simplicity afforded by Keldysh time ordering also manifests itself in more traditional quantum-optics approaches to computing time-dependent observables: The observables made accessible via a functional integral formulated on the Keldysh contour are precisely the same as those computable via the quantum-regression formula without evolving the density matrix backwards in time.

Writing the so-called classical (ψ) and quantum (φ) fields as $(\psi^1, \psi^2) \equiv (\bar{\psi}, \psi)$ and $(\varphi^1, \varphi^2) \equiv (\bar{\varphi}, \varphi)$, and writing creation and annihilation operators as $(\hat{a}^1, \hat{a}^2) \equiv (\hat{a}^\dagger, \hat{a})$, Keldysh-ordered correlation functions can be computed as (restoring spatiotemporal indices and defining $\mu = 1, 2$)

$$\langle \mathcal{T}_K(\dots \hat{a}_j^\mu(t^\pm) \dots) \rangle = \langle \dots (\psi_j^\mu(t) \pm \varphi_j^\mu(t)) \dots \rangle_{\mathcal{Z}}. \quad (7)$$

On the left-hand side of Eq. (7) the expectation value is taken with respect to the initial density matrix, and the operators evolve in the Heisenberg picture [88]. The symbol \mathcal{T}_K time-orders all operators whose time arguments have a “+” superscript, and anti-time-orders those with time arguments that have a “−” superscript, placing all of the latter to the left of all of the former. On the right-hand side of Eq. (7) the expectation value is taken with respect to the functional integral \mathcal{Z} ; i.e., it is computed by inserting the relevant fields into the integrand of Eq. (5) (\mathcal{Z} is normalized to unity by construction, as it expresses the trace of the density matrix). In the calculations that follow, we exploit a semiclassical limit in which the quantum field is, in a sense to be made precise, parametrically smaller than the classical field; thus we are primarily interested in correlations of the classical field alone, which can be converted into operator expectation values by inverting Eq. (7),

$$\langle \psi_{j_1}^{\mu_1}(t_1) \dots \psi_{j_n}^{\mu_n}(t_n) \rangle_{\mathcal{Z}} = \frac{1}{2^n} \sum_{\sigma=\pm} \langle \mathcal{T}_K(\hat{a}_{j_1}^{\mu_1}(t_1^{\sigma_1}) \dots \hat{a}_{j_n}^{\mu_n}(t_n^{\sigma_n})) \rangle. \quad (8)$$

While it may appear that such correlation functions can be discontinuous at coinciding times due to the associated change of operator ordering on the right-hand side of Eq. (8), it is straightforward to show that this is not actually the case. Instead, when the times approach each other ($t_1, \dots, t_n \rightarrow t$), the limit of an arbitrary n -point correlation function of the classical field ψ smoothly approaches the equal-time value

$$\langle \psi_{j_1}^{\mu_1}(t) \dots \psi_{j_n}^{\mu_n}(t) \rangle_{\mathcal{Z}} = \langle (\hat{a}_{j_1}^{\mu_1}(t) \dots \hat{a}_{j_n}^{\mu_n}(t))_s \rangle, \quad (9)$$

where $(\dots)_s$ symmetrizes (i.e., Weyl-orders) products of creation and annihilation operators [4]. In other words, equal-time correlation functions of the classical field reproduce

the average of Weyl-ordered operator products. For example, the density can be computed from the two-point correlation function

$$\begin{aligned} \langle \bar{\psi}_j(t) \psi_j(t) \rangle_{\mathcal{Z}} &= \langle (\hat{a}_j^\dagger(t) \hat{a}_j(t))_s \rangle = \frac{1}{2} \langle \hat{a}_j^\dagger(t) \hat{a}_j(t) + \hat{a}_j(t) \hat{a}_j^\dagger(t) \rangle \\ &= \langle \hat{n}(t) \rangle + \frac{1}{2}. \end{aligned} \quad (10)$$

The Weyl symbol of \hat{H} is given by (ignoring additive constants, which do not affect correlation functions)

$$\begin{aligned} \mathcal{H}_w(\alpha) &= \sum_j \left(-\bar{\alpha}_j (J \nabla^2 + \mu + U) \alpha_j \right. \\ &\quad \left. + \frac{U}{2} |\alpha_j|^4 + \Omega (\bar{\alpha}_j + \alpha_j) \right). \end{aligned} \quad (11)$$

Here $\nabla^2 \alpha_j \equiv -z \alpha_j + \sum_{\langle k, j \rangle} \alpha_k$ is the discrete Laplacian, the lattice coordination number $z = 2D$ in D dimensions, and $\mu = \delta + zJ$. Inserting Eq. (11) into Eq. (6) yields the action

$$\begin{aligned} S &= 2 \sum_j \int dt \bar{\varphi} \left(i \partial_t \psi + \left(J \nabla^2 + \mu + i \frac{\kappa}{2} \right) \psi \right. \\ &\quad \left. - \Omega - U |\psi|^2 \psi \right) + \text{c.c.} + 2i\kappa \sum_j \int dt \bar{\varphi} \varphi \\ &\quad + 2U \sum_j \int dt (\bar{\varphi} \varphi + 1) (\psi \bar{\varphi} + \bar{\psi} \varphi). \end{aligned} \quad (12)$$

Because the functional integral is Gaussian for $U = 0$, and because optical bistability at the mean-field level can occur at arbitrarily small values of U , one might hope that some aspects of the relevant physics can be captured by doing perturbation theory in U . However, this is not the case; while optical bistability can indeed occur for small U , it always occurs when the typical interaction energy seen by a particle, $U|\psi|^2$, is comparable to the other energy scales of the problem. Nevertheless, the action can still be organized around a small parameter that enables a controlled approximation. To this end, we define rescaled fields and parameters

$$\Phi \equiv \varphi \sqrt{\mathcal{N}}, \quad \Psi \equiv \psi / \sqrt{\mathcal{N}}, \quad \omega \equiv \Omega / \sqrt{\mathcal{N}}, \quad u \equiv U \mathcal{N}, \quad (13)$$

in terms of which the action can be rewritten

$$\begin{aligned} S &= 2 \sum_j \int dt \bar{\Phi} \left(i \partial_t \Psi + \left(J \nabla^2 + \mu + i \frac{\kappa}{2} \right) \Psi \right. \\ &\quad \left. - \omega - u |\Psi|^2 \Psi \right) + \text{c.c.} + \frac{1}{\mathcal{N}} \sum_j \int dt (2i\kappa \bar{\Phi} \Phi) \\ &\quad + \frac{1}{\mathcal{N}^2} \sum_j \int dt 2u (\bar{\Phi} \Phi + \mathcal{N}) (\Psi \bar{\Phi} + \bar{\Psi} \Phi). \end{aligned} \quad (14)$$

The dimensionless parameter \mathcal{N} implicitly identifies a one-dimensional family of actions at fixed values of ω , u , κ , μ , and J , one limit of which (large \mathcal{N}) will be shown to admit a tractable analysis. Note that the limit $\mathcal{N} \rightarrow \infty$ at fixed u and ω is not the same thing as the limit $U \rightarrow 0$, even though the coupling U does become small in this limit. Rather, increasing \mathcal{N} amounts to increasing the drive strength Ω while

simultaneously decreasing the coupling U in such a way that the typical interaction energy per particle, $U|\psi|^2$, remains constant.

To see this, we first evaluate the functional integral in the limit $\mathcal{N} \rightarrow \infty$, in which only the first term in the action survives. The functional integral over Φ can be carried out and yields a functional δ function of the term inside parentheses, thereby enforcing the mean-field equation of motion

$$i\partial_t\Psi = -J\nabla^2\Psi - (\mu + i\kappa/2)\Psi + \omega + u|\Psi|^2\Psi. \quad (15)$$

Note that at this level of approximation, varying \mathcal{N} with ω and u held fixed leaves the equation of motion for Ψ invariant. Therefore, as discussed above, the typical interaction energy seen by each particle, $U|\psi|^2 = u|\Psi|^2$, stays fixed. The only consequence is that the actual density, $|\psi|^2 = \mathcal{N}|\Psi|^2$, is enhanced by a factor of \mathcal{N} , which therefore sets the overall density scale [79].

For \mathcal{N} large but finite, the final term on the second line of Eq. (14) suppresses contributions to the functional integral unless $\Phi \lesssim \mathcal{N}^{1/2}$. The final term can therefore be estimated as $\Phi^3\mathcal{N}^{-2} + \Phi\mathcal{N}^{-1} \lesssim \mathcal{N}^{-1/2}$ and can be safely ignored in the large- \mathcal{N} limit. At this level of approximation, the functional integral can no longer be solved exactly, but it can be mapped onto stochastic classical equations by standard techniques. Decoupling the term that is quadratic in Φ with a Hubbard-Stratonovich transformation,

$$e^{-(2\kappa/\mathcal{N})|\Phi|^2} = \frac{2\mathcal{N}}{\kappa\pi} \int d^2\zeta e^{2i(\Phi\bar{\zeta} + \bar{\Phi}\zeta)} e^{-2\mathcal{N}|\zeta|^2/\kappa},$$

the action again becomes linear in Φ . This time, for fixed ζ , the functional integral over Φ enforces the equation of motion

$$i\partial_t\Psi = -J\nabla^2\Psi - (\mu + i\kappa/2)\Psi + \omega + u|\Psi|^2\Psi + \zeta. \quad (16)$$

The remaining functional integral over ζ with a Gaussian weight $\exp(-2\mathcal{N}|\zeta|^2/\kappa)$ indicates that we should interpret Eq. (16) as a stochastic differential equation, with ζ being complex, Gaussian white noise of variance (restoring spatial and temporal indices)

$$\overline{\zeta_j(t_1)\zeta_k(t_2)} = \frac{1}{\mathcal{N}} \frac{\kappa}{2} \delta_{j,k} \delta(t_1 - t_2). \quad (17)$$

Hence the dynamics of the rescaled classical field Ψ , to this order in $1/\mathcal{N}$, is governed by a stochastic and dissipative Gross-Pitaevskii equation with parametrically weak noise [89]. Equation (16) can also be derived using phase-space techniques; in this context, it would arise as the so-called truncated Wigner approximation, an approximation to the (otherwise-exact) multidimensional partial-differential equation governing the time evolution of the Wigner function. However, the functional-integral approach used here makes the identification and justification of this approximation very transparent and has the advantage that one can assess the consequences of the approximations that lead to Eq. (16) within the framework of the renormalization group. In particular, as discussed in Ref. [69], Eq. (16) should reproduce the correct critical exponents for the phase transition exhibited by the exact steady state of Eq. (2). Unlike in Ref. [69], however, here we have explicitly identified a limit (large \mathcal{N}) in which Eq. (16) yields asymptotically exact results for microscopic observables and thus can be used to make

quantitative predictions about the behavior of correlation functions at the lattice scale (rather than just qualitative predictions about their long-distance asymptotics). Moreover, the existence of this limit furnishes a more formal justification for perturbative renormalization-group analyses.

It is important to realize that, even for a single cavity, Eqs. (15) and (16) must be interpreted with some care in order to correctly extract steady-state properties in the large- \mathcal{N} limit. In particular, there is a sense in which mean-field theory makes incorrect predictions about the steady-state even in the limit of large \mathcal{N} . The difficulty can be seen by returning to the full functional integral; there, it can be shown that the limits $\mathcal{N} \rightarrow \infty$ and $t \rightarrow \infty$ do not commute when the parameters are tuned to be inside the mean-field bistable region [90]. Indeed, if we take the limit $\mathcal{N} \rightarrow \infty$ first, Eq. (15) is exact at all times, and we are led to conclude (on the same basis as the analysis in Sec. II) that there are two stable steady states. If we instead take the large- t limit first, we would find [based on the analysis of $1/\mathcal{N}$ corrections contained in Eq. (16)] that there is a unique steady state at any finite value of \mathcal{N} , which is a fluctuation-induced admixture of the two stable steady states computed by reversing the order of limits. If we now take the large- \mathcal{N} limit, one of those steady states is generally preferred over the other, in the sense that it alone determines all steady-state expectation values. Thus we encounter a sudden (first-order) phase transition between a bright and a dark state when traversing through the mean-field bistable region.

IV. MEAN-FIELD THEORY

From the above considerations it is clear that, at least in the large- \mathcal{N} limit, the steady state can be understood by solving the mean-field equations of motion in the presence of parametrically weak noise. Thus we expect a detailed understanding of the mean-field dynamics in the absence of noise to form a useful starting point for understanding the dynamics of Eq. (16). At the level of Eq. (15), and assuming that only uniform steady states exist, the steady-state phase diagram is identical to that of single-cavity optical bistability, up to the replacement of δ by μ . In terms of the rescaled field Ψ and a rescaled density $N \equiv |\Psi|^2 = \mathcal{N}^{-1}|\psi|^2$, Eq. (4) becomes

$$N((\mu - uN)^2 + \kappa^2/4) = \omega^2. \quad (18)$$

Straightforward analysis of Eq. (18) shows that upon entering the bistable region from outside of it, the additional solution does not in general emerge continuously from the existing one. However, if one enters the bistable region through the cusp located at (see Fig. 3)

$$\{\kappa_c, \omega_c\} = \mu\{(4/3)^{1/2}, (2/3)^{3/2}(\mu/u)^{1/2}\}, \quad (19)$$

the two solutions do emerge continuously from a single solution, $\Psi_c = e^{-i\pi/3} \sqrt{2\mu/3u}$ (with critical density $N_c = |\Psi_c|^2 = 2\mu/3u$). Therefore, we can identify the cusp of the bistable region as a mean-field critical point locating a continuous phase transition from one to two steady states.

It is straightforward to show that one can only enter the bistable region through the critical point along the line $(\kappa - \kappa_c) = \sqrt{8u/\mu}(\omega - \omega_c)$. It is convenient in what follows to define new coordinates in the κ - ω parameter space that naturally parametrize deviations from the mean-field critical




FIG. 3. (a) Coordinates used to parametrize passage through the mean-field critical point and into the bistable region; r and h control deviations in a nonorthogonal coordinate system spanned by the dashed arrow. (b) Setting $h = 0$ and scanning r from positive to negative causes the order parameter to undergo a cusp bifurcation, at which the solution outside the bistable region goes unstable (dashed line) and two new dynamically stable steady states (solid lines) emerge continuously. In both plots $u = \mu$, and energies are given as dimensionless ratios with μ .

point along and away from this line,

$$r = \frac{1}{2}(\kappa - \kappa_c), \quad h = \frac{4}{\sqrt{3}}(\omega - \omega_c) - \sqrt{\frac{2\mu}{3u}}(\kappa - \kappa_c). \quad (20)$$

Note that this parameter transformation is designed so that only r varies as we enter the bistable region through the critical point—the overall normalization of r and h is arbitrary, and chosen to make the formulas that follow simpler. These new variables can be visualized as parametrizing deviations from the mean-field critical point in the nonorthogonal coordinate system shown in Fig. 3(a). For $h = 0$, moving from $r > 0$ to $r < 0$ causes the mean-field solution outside of the bistable region to undergo a cusp bifurcation [Fig. 3(b)]. For $r < 0$, sweeping h from negative to positive traverses the bistable region in such a way that the system goes from supporting only a low-density solution, to having coexisting low-density and high-density solutions, and then eventually to supporting only a high-density solution. This behavior is in close analogy to that of an Ising model: If the dark and bright solutions are identified with the up-down-symmetry-related free-energy minima, then r plays the role of the reduced temperature and h plays the role of a symmetry-breaking (longitudinal) field, causing one to be preferred over the other. The remainder of Sec. IV formalizes this analogy, and in Sec. V we argue that it continues to hold even when fluctuations are included.

A. Near-critical dynamics

We are particularly interested in the effects of fluctuations in the vicinity of the mean-field critical point, which requires that we first understand the mean-field response when (a) the parameters are tuned close to the mean-field critical point (both r and h are small) and (b) the order parameter Ψ is perturbed weakly from its steady-state value. First working directly at the critical point ($r = h = 0$, with steady-state solution $\Psi = \Psi_c$) and assuming that $\Psi = \Psi_c + \delta\Psi$ is uniform and close to the critical value, we expand Eq. (15) to first order in $\delta\Psi$ to obtain

$$\partial_t \delta\Psi = -\frac{\mu}{2}[(\sqrt{3} + i)\delta\Psi + (\sqrt{3} - i)\delta\Psi]. \quad (21)$$

The right-hand side (rhs) of Eq. (21) is purely real, and thus only the real part of $\delta\Psi$ decays; the dynamics stops when the rhs vanishes, i.e., when $\arg(\delta\Psi) = \pi/2 - \arg(\sqrt{3} + i) = \pi/3$. This observation motivates the following decomposition of the complex-valued $\delta\Psi$ into two real components,

$$\delta\Psi = \varrho + e^{i\pi/3}\sigma, \quad (22)$$

with the expectation that ϱ and σ will relax quickly and slowly, respectively, in the vicinity of the mean-field critical point. Inserting this decomposition into Eq. (15) but now keeping all orders in ϱ and σ , we obtain coupled nonlinear differential equations for ϱ and σ (see Appendix B for a detailed discussion). The fast variable ϱ can be adiabatically eliminated by solving $\partial_t \varrho = 0$ for ϱ (perturbatively in r and h) and inserting the solution into the equation of motion for σ . In this way, to lowest nontrivial order in r we obtain

$$\partial_t \sigma = \frac{J}{\sqrt{3}}\nabla^2 \sigma - r\sigma - \frac{u}{\sqrt{3}}\sigma^3 - \frac{h}{2}. \quad (23)$$

Note that we have also dropped higher-order derivative terms for σ that arise from the adiabatic elimination of ϱ ; this omission turns out to be justified (to lowest nontrivial order in r) near the mean-field critical point, where the fields vary slowly in space, even when J is not small compared to the other energy scales in Eq. (23) (the perturbative adiabatic elimination of ϱ is explained in detail in Appendix B). Restoring spatial indices and defining parameters

$$K = J/\sqrt{3}, \quad g = u/\sqrt{3}, \quad (24)$$

Eq. (23) can be rewritten as

$$\partial_t \sigma_j = -\frac{\partial \mathcal{H}(\sigma)}{\partial \sigma_j}, \quad (25)$$

where the effective Hamiltonian $\mathcal{H}(\sigma)$ is defined as

$$\mathcal{H}(\sigma) = \frac{1}{2} \sum_j \left(K |\nabla \sigma_j|^2 + r\sigma_j^2 + \frac{1}{2}g\sigma_j^4 + h\sigma_j \right). \quad (26)$$

Note that, as anticipated, $\mathcal{H}(\sigma)$ is precisely the energy functional defining the Landau theory of a classical Ising model, with σ playing the role of the magnetization. Equation (25) indicates that the dynamics of the slow field in the vicinity of the critical point is purely relaxational, evolving towards the minimum of the effective potential $\mathcal{H}(\sigma)$.

B. Domain walls

At the level of mean-field theory there are two truly stable homogeneous solutions within the bistable region. However, it is clear that if we place the system in one of the two mean-field steady states, the inclusion of fluctuations will seed defects of the other steady state; whether these defects shrink or grow will depend on the dynamics of the domain wall separating them from the bulk, and will determine which of the two mean-field steady states is preferred over the other. Thus the identification of a point in the bistable region where the mean-field velocity of a domain wall vanishes gives a first approximation to the location of the phase transition when fluctuations are included.

It is difficult to analytically extract domain-wall dynamics directly from Eq. (15), so to proceed we make three

assumptions: (1) the parameters are tuned to be inside the bistable region and close to the mean-field critical point, (2) the domains are smooth, such that a continuum approximation is justified, and (3) if $D > 1$, the domains are large and thus have vanishing curvature. The first assumption justifies the use of Eq. (25) to calculate the dynamics. The second assumption requires that J is large compared to the local energy scales of the problem, e.g., to the characteristic time scale associated with dynamics in the potential part of the Hamiltonian

$$U(\sigma) = \frac{1}{2}(r\sigma^2 + \frac{1}{2}g\sigma^4 + h\sigma). \quad (27)$$

Note that near criticality, this only requires that J is large compared to r and h , and not that J is large compared to the energy scales ω and γ . The third assumption is made because the phase that is favored in the limit of weak fluctuations is the one in which asymptotically large defects of the opposite phase are disfavored (i.e., tend to shrink).

The dynamics of a flat domain wall is effectively one-dimensional and can be ascertained from a one-dimensional continuum version of Eq. (25),

$$\partial_t \sigma(x,t) = K \partial_x^2 \sigma(x,t) - r\sigma(x,t) - g\sigma(x,t)^3 - \frac{h}{2}. \quad (28)$$

When $h = 0$, the symmetry of Eq. (28) under inversions $\sigma \rightarrow -\sigma$ implies that domain walls must be stationary; both uniform phases have the same effective potential, and relaxational dynamics cannot prefer one over the other. Therefore, the line $h = 0$ provides a first approximation to the dividing line between parts of the bistable region in which the bright phase is more stable and parts in which the dark phase is more stable.

For h small but nonzero, the domain-wall velocity can be estimated in the following manner [91]. Making a traveling wave ansatz $\sigma(x,t) = \sigma(\tau)$, with $\tau \equiv x - vt$, and denoting derivatives with respect to τ by dots, Eq. (28) becomes

$$K\ddot{\sigma} = -v\dot{\sigma} + \frac{\partial U(\sigma)}{\partial \sigma}. \quad (29)$$

Equation (29) can be interpreted as Newton's equation for a particle with position σ and mass K , moving in the inverted potential $-U(\sigma)$ and subject to a linear drag with friction coefficient v . Inside the mean-field bistable region ($r < 0$) there are two stationary solutions of Eq. (29) associated with the two local maxima of the potential energy $-U(\sigma)$ (Fig. 4); these correspond to the two spatially uniform mean-field steady states. To zeroth order in h , these solutions are located at




FIG. 4. (a) Domain-wall dynamics near the critical point. Since the dynamics is relaxational, the domain wall moves in such a direction that the lower-energy domain increases in size. (b) This dynamics can be mapped onto the motion of a fictitious particle in an inverted potential.

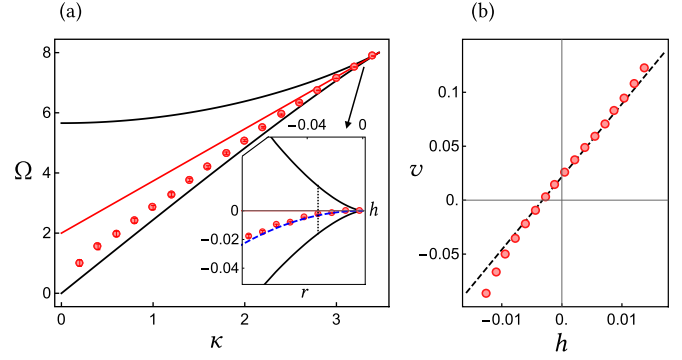


FIG. 5. (a) Numerically determined location of the domain-wall velocity zeros (red disks), together with the analytical estimate of the zero-velocity line, $h = 0$, valid near the critical point (red line). The size of the disks reflects the largest expected uncertainty in the numerical determination of these points. Inset: Exploded view of the main plot near the critical point, now in terms of the parameters h and r . The dashed blue line is an improved estimate of the zero-velocity line obtained by extending Eq. (23) to next leading (quadratic) order in r . (b) Numerically extracted domain-wall velocity (red disks) as a function of h [taken along the black dotted line shown in the inset of (a)], compared with the estimate in Eq. (30) (black dashed line). Note that the black dashed line does not vanish at $h = 0$. While its slope is taken from Eq. (30), it has been shifted by an amount that we determine by extending Eq. (23) to next-to-leading order in the deviations from the critical point [i.e., the same correction used to produce the blue dashed line in the inset of (a)], which clearly agrees well with the numerically calculated shift of the zero-velocity point.

$\sigma_{\pm} = \pm\sigma_0$, with $\sigma_0 = \sqrt{|r|/g}$. In addition to the two stationary solutions, a solution can be found that interpolates from the higher local maximum to the lower one, which for $h > 0$ is located at $\sigma = +\sigma_0$. The friction coefficient v must be determined self-consistently such that the particle comes to rest at the lower local maximum. Standard analysis of the solutions of Eq. (29) based on conservation of energy yields, to first order in h (see Appendix C for details),

$$v \approx h \frac{3}{2} \sqrt{\frac{Kg}{2r^2}}. \quad (30)$$

The analysis above is corroborated by brute-force numerical integration of Eq. (15). The true zero-velocity line can be determined numerically by solving Eq. (15) with a domain wall inserted at $t = 0$, and agrees with the $h = 0$ line near the critical point [Fig. 5(a)]. Also, as shown in Fig. 5(b), Eq. (30) agrees well with the numerically extracted domain-wall velocity.

V. LANGEVIN EQUATIONS

Mean-field theory suggests that the steady state of the master equation in Eq. (2) undergoes an Ising-like phase transition in sufficiently high spatial dimensions. However, in order to understand the detailed nature of this phase transition, and to determine its lower critical dimension, fluctuations must be taken into account. As discussed in Sec. III, for large \mathcal{N} the dominant fluctuations are captured by working with the stochastic and dissipative Gross-Pitaevskii equation (GPE) in

Eq. (16), reproduced here for clarity:

$$i\partial_t\Psi = -J\nabla^2\Psi - (\mu + i\kappa/2)\Psi + \omega + u|\Psi|^2\Psi + \zeta. \quad (31)$$

At this level of approximation, expectation values of the classical field are obtained by averaging the solution of Eq. (31) over realizations of the noise ζ , $\langle \cdots \rangle_{\mathcal{Z}} \approx \langle \cdots \rangle_{\text{SGPE}}$. Note that, unlike in mean-field treatments, Eq. (31) gives access to approximate correlation functions of operators at unequal points in space and time via the correspondence in Eqs. (8) and (9).

Though we cannot solve Eq. (31) analytically, simple arguments can be made to explain many features of the steady state quantitatively near the mean-field critical point, and qualitatively even away from it. As before, the near-critical dynamics is simplified by decomposing the field as $\Psi = \Psi_c + (\varrho + e^{i\pi/3}\sigma)$. Adiabatic elimination of ϱ can again be performed perturbatively in h and r ; the only subtlety is that fluctuations cause ϱ to undergo a lattice version of the Ornstein-Uhlenbeck process [92], which feeds back into the equation of motion for σ as non- δ -correlated noise. However, it is straightforward to show that near the critical point the correlation time of this additional noise is short compared to the dynamical time scales of σ , and it can be incorporated as a perturbative renormalization of the δ -correlated noise acting directly on σ . Details of the calculation are reported in Appendix B, and here we simply quote the final result,

$$\partial_t\sigma_j = -\frac{\partial\mathcal{H}(\sigma)}{\partial\sigma_j} + \xi_j(t). \quad (32)$$

Here, $\mathcal{H}(\sigma)$ is the same energy functional given in Eq. (26), and $\xi_j(t)$ is (real) Gaussian white noise with variance

$$\overline{\xi_j(t_1)\xi_k(t_2)} = \frac{\kappa}{3\mathcal{N}}\delta_{j,k}\delta(t_1 - t_2). \quad (33)$$

Equation (32) is a spatially discretized version of model A in the Hohenberg-Halperin classification [93], suggesting that the steady-state phase transition associated with optical bistability in the driven-dissipative Bose-Hubbard model is, as anticipated, in the universality class of the finite-temperature classical Ising model. In particular, steady-state and static observables generated by the stochastic dynamics in Eq. (32) can be computed with respect to a Boltzmann weight,

$$\mathcal{P}(\sigma) = \mathcal{Z}^{-1}e^{-\mathcal{H}(\sigma)/T_{\text{eff}}}, \quad \mathcal{Z} = \int \prod_j d\sigma_j e^{-\mathcal{H}(\sigma)/T_{\text{eff}}}, \quad (34)$$

with an effective temperature given by

$$T_{\text{eff}} = \kappa/3\mathcal{N}. \quad (35)$$

The large- \mathcal{N} limit was designed to suppress fluctuations in the microscopic action and so, unsurprisingly, it corresponds to a low-temperature limit of the effective equilibrium description of the phase transition. Returning to the underlying microscopic degrees of freedom, it is straightforward to see that the dynamics of this effective theory is imprinted on experimentally measurable observables. The connection is particularly simple near the mean-field critical point. For example, working to lowest nontrivial order in the deviations of the fields from their mean-field critical values, straightforward algebra yields

the equal-time connected density-density correlation function

$$\begin{aligned} \mathcal{C}_{jk} &= \langle \hat{n}_j(t)\hat{n}_k(t) \rangle - \langle \hat{n}_j(t) \rangle \langle \hat{n}_k(t) \rangle \\ &\propto \langle \sigma_j(t)\sigma_k(t) \rangle_{\text{SGPE}} - \langle \sigma_j(t) \rangle_{\text{SGPE}} \langle \sigma_k(t) \rangle_{\text{SGPE}}. \end{aligned} \quad (36)$$

The critical properties of the finite-temperature Ising model should, therefore, control the critical fluctuations of the intensity of light emitted from a coherently driven array of exciton-polariton microcavities.

Before considering what happens away from the critical point, we first briefly summarize a qualitative picture of model A dynamics and its connection to the Ising model. Suppose that the system is seeded in a locally random initial configuration: We would like to know what happens to it in steady state. At short times and for $r < 0$, we expect the system to form domains of both (locally stable) phases, separated by domain walls. In the absence of fluctuations (i.e., at $T_{\text{eff}} = 0$) the preferred steady state of the system can be understood by simple domain-wall dynamics; for $h \neq 0$ one phase is preferred over the other, and the system will eventually order in that phase. If fluctuations are now turned on, domains of the less favored phase will be seeded, and the consequence of these defects depends crucially on the dimensionality. In one dimension, the domain walls enclosing these defects move independently of each other when they are sufficiently far apart, undergoing a biased random walk. As a result, when $h \rightarrow 0$ and the dynamics becomes unbiased, defects proliferate and the system will be disordered at any finite temperature. In two or more spatial dimensions, defects of the less favored phase will still be seeded by fluctuations, but small defects contract aggressively even as $h \rightarrow 0$ due to a surface tension. Therefore, at least at sufficiently small temperature, as $h \rightarrow 0$ the system remains ordered in a phase that depends on whether h approaches zero from below or above, indicating a first-order phase transition.

Since the above argument relies very little on the dynamics being relaxational, and primarily on the existence of domain walls that—in the absence of fluctuations and for asymptotically large domains—have a directional preference that changes as we move through the bistable region, it is reasonable to expect the qualitative picture described above to be valid even away from the critical point. Nevertheless, because the dynamics generated by Eq. (16) does not induce an equilibrium steady-state distribution far away from the critical point, it is important to verify this picture numerically.

To this end, we carry out a brute-force numerical integration of Eq. (31) in both one and two dimensions using a fixed-time-step first-order Euler-Mayurama method. After a burn-in time, the equations are integrated until statistical error bars (1σ) of fractional size 0.01 are achieved for the density. Temporal autocorrelations on time scales of 1/10 the total integration time are also required to fall below a similar threshold to ensure that the averaging time is long compared to all dynamical time scales, which can become anomalously large near the crossover or phase transition. Exemplary results of these numerics in one dimension are shown in Fig. 6 and reflect the spatiotemporal dynamics of output light intensity that would be observed if the model were realized in an array of exciton-polariton microcavities. As expected, by sweeping vertically through the mean-field bistable region, we change from a dominantly dark steady state with small domains of




FIG. 6. Real-time dynamics (after burn in) in a 1D system with 128 sites and periodic boundary conditions, showing domain proliferation in the vicinity of the crossover in one dimension. In all three plots, $(J, U, \kappa) \approx (0.1\mu, 0.2\mu, 0.3\mu)$. The left-hand panel is just on the dark side of the crossover, the middle panel is roughly in the middle of the crossover, and the right-hand panel is just on the bright side of the crossover. The color indicates the density.

the bright phase to a predominantly bright steady state with small domains of the dark phase. Because the domain walls are unbound, this change manifests as a smooth crossover rather than a true phase transition, as confirmed in the 1D phase diagram shown in Fig. 7(a) [in particular, see the cross section plotted in Fig. 7(c)]. Near the mean-field critical point, the characteristic domain size at the crossover point (see, for example, the central panel of Fig. 6) reflects the effective temperature of the model; as κ shrinks the domains grow in size, and the crossover becomes sharper. However, the extent to which this behavior persists as $\kappa \rightarrow 0$, and whether this limit is strictly analogous to the $T \rightarrow 0$ limit of the Ising model, is difficult to say. A careful numerical analysis of density-density correlation functions in one dimension reveals that they always decay exponentially at the crossover point, with a correlation length that grows monotonically with decreasing κ . However, as κ decreases, eventually the domain size becomes so large—and the dynamics of domain-wall diffusion becomes so slow—that we are unable to obtain statistically converged results. This computational limitation imposes the lower limit on κ in the 1D phase diagram reported in Fig. 7(a).

In two dimensions [Figs. 7(b) and 7(d)], the dynamics is qualitatively different. In sweeping from small to large drives at sufficiently small κ , one encounters a clear first-order phase transition between the bright and dark phases, consistent with the expected equilibrium physics of the 2D Ising model. The size of the discontinuity increases with decreasing κ (and thus with decreasing effective temperature). The initial state is a random admixture of the two mean-field steady states, which plays the role of an infinite-temperature state. Thus for κ below the critical point and Ω chosen close to the first-order phase transition, the initial dynamics can be viewed (in the language of equilibrium physics) as a quench from an infinite-temperature phase to a final temperature below the ordering temperature. The short-time dynamics therefore shows the expected coarsening of small domains. Eventually the more favored phase wins out unless Ω is tuned precisely

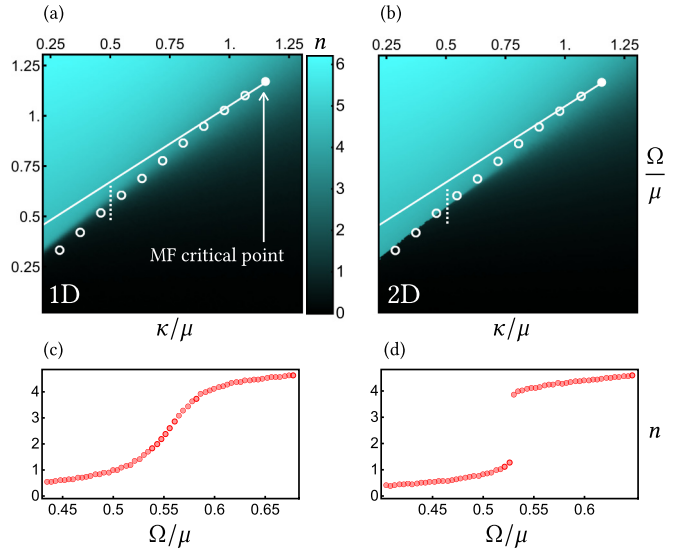


FIG. 7. Phase diagrams obtained by numerically solving Eq. (16) on (a) a 1D chain with 128 sites and (b) a 2D square lattice with 32×32 sites (in both cases periodic boundary conditions were used). For both plots, the parameters used are $(J, U) \approx (0.1\mu, 0.2\mu)$. The solid white line locates the (near-critical) condition for a vanishing domain-wall velocity, $h = 0$, while the white circles indicate numerically obtained velocity zeros. The dashed white lines indicate the parameter regime used for plots (c) and (d). Note that the phase diagrams are cut off at small κ , or equivalently low effective temperature, because statistically converged numerical solutions of Eq. (16) require prohibitively long integration times as fluctuations become weaker. (c) and (d) Cuts through the phase diagrams indicated by dashed white lines in (a) and (b). In one dimension (c) the mean-field phase transition is smoothed out into a crossover, while in two dimensions (d) bistability leads to a true first-order phase transition. The error bars in (c) and (d) are smaller than the size of the plot markers.

to the phase transition, but this claim takes progressively more averaging to establish reliably as one moves closer to the phase-transition line. Once again, decreasing the effective temperature by decreasing κ leads to a slowing down of the dynamics, making it difficult to obtain statistically converged results near the first-order phase transition and imposing a lower limit on κ in the 2D phase diagram reported in Fig. 7(b).

VI. DISCUSSION

By bringing together a number of ideas from both quantum optics and condensed-matter physics, we have identified a limit of the driven-dissipative Bose-Hubbard model in which the dominant fluctuations are captured by nonequilibrium Langevin equations, enabling a quantitatively accurate and computationally efficient determination of steady-state properties. Near the critical point, these fluctuations are thermal and lead to an effective equilibrium description. However, we emphasize that the Langevin description of the problem should be asymptotically exact in the limit of $\mathcal{N} \rightarrow \infty$ even away from the critical point, where an equilibrium description is *not* valid. Numerically, we find that in two dimensions the first-order phase transition expected from the mapping onto an

Ising model remains intact far from the critical point, where this mapping is not strictly valid. In addition, the absence of a phase transition in one dimension is the result of the same domain-wall phenomenology that prevents ordering of the 1D Ising model at finite temperatures. In this way, the numerical results reinforce and extend the assertion that the steady-state behavior of the driven-dissipative Bose-Hubbard model possesses an emergent description in terms of the equilibrium physics of a finite-temperature classical Ising model. These conclusions have direct consequences for a range of experiments in which particle loss is countered by coherent driving, for example a coherently driven exciton-polariton fluid in an array of semiconductor microcavities. Here, by tuning the laser driving strength through the mean-field bistable regime, one should be able to observe domain growth, hysteresis, critical fluctuations, and other generic features of the Ising model in a longitudinal field, all by simple correlation measurements on the output intensities of the cavities.

We caution that any claims about the universality class of the phase transition require more than just a microscopically accurate treatment of fluctuations—it is also important to identify the relevance of any ignored fluctuations in the sense of the renormalization group, even if they are parametrically small. In other words, there is no guarantee that small ($1/\mathcal{N}$) quantitative errors at the scale of the lattice spacing will not qualitatively affect the nature of the phase transition. Renormalization-group arguments (i.e., canonical power counting) supporting the Ising universality class as the correct critical theory at small but finite \mathcal{N} can be inferred from the results of Ref. [69]. Moreover, we note that because Eq. (32) provides an increasingly accurate approximation to the microscopic dynamics for increasing \mathcal{N} , Ising-like critical behavior should manifest itself at least as an intermediate length-scale crossover phenomenon, regardless of the true universality class of the phase transition. We also emphasize that the precise nature of the first-order phase transition far away from the critical point is less clear, in particular its fate as $\kappa \rightarrow 0$. It would be worthwhile to extend the numerical approach taken in Sec. V to confirm the universal aspects of both the critical point and the first-order phase transition.

It would also be worthwhile to compare some of the results in one dimension with numerically exact calculations based on the density-matrix renormalization group [94] in order to better understand the importance of higher-order (in $1/\mathcal{N}$) corrections that are not captured by the Langevin description. In particular, at large U (small \mathcal{N}), mean-field arguments suggest that the inclusion of fluctuations ignored to leading order in $1/\mathcal{N}$ may lead to richer steady-state behaviors [47,58], including phases that spontaneously break discrete spatial-translation symmetry [47,48]. In two dimensions, the formalism described here could be used to compute other dynamical aspects of the system near the first-order phase transition; for example, it should be possible to calculate the lifetime of the metastable phase via an instanton approach.

ACKNOWLEDGMENTS

We thank Cristiano Ciuti, Sebastian Diehl, Howard Carmichael, Sarang Gopalakrishnan, Victor Gurarie, Ana Maria Rey, Murray Holland, Anzi Hu, Michael Fleischhauer,

and Chih-Wei Lai for helpful discussions. M.F.M., J.T.Y., and A.V.G. acknowledge support by ARL CDQI, ARO MURI, NSF QIS, ARO, NSF PFC at JQI, and AFOSR. R.M.W. acknowledges partial support from the National Science Foundation under Grant No. PHYS-1516421. M.H. acknowledges support by AFOSR-MURI, ONR, and the Sloan Foundation.

APPENDIX A: FUNCTIONAL-INTEGRAL FORMALISM

The functional integral presented in Sec. III is closely related to the usual Keldysh functional-integral formalism, for which there are many good references (see, for example, Ref. [87]). However, there is a subtle difference between the formalism used here and that typically employed in the condensed-matter community, and it therefore seems worthwhile to provide an explicit derivation of Eqs. (5)–(9). For simplicity, we treat only the single-cavity case, but the generalization to many cavities that yields Eqs. (5)–(9) follows immediately.

In any functional-integral formulation of quantum mechanics, operators must be traded in for classical variables. The usual way to do this, as is the case in the standard approach to the Keldysh function integral, is to repeatedly insert coherent-state resolutions of identity during the time evolution. Operators get sandwiched between coherent states, and if they are normal ordered they turn into functions of phase-space variables. In the language of quantum optics, operators are exchanged for their Q symbols. However, there are many different ways to associate operators with functions over phase space, and thus many ways to formulate a functional integral. In the following derivation, we replace operators with classical variables by working in the Weyl representation (see, for example, Ref. [81]). In our opinion, even though this strategy entails some additional overhead in phase-space formalism, it is both more direct and conceptually simpler than the usual approach to the Keldysh functional integral. In particular, the canonical approach described in Refs. [86,87] relies on the construction of a formal continuous-time notation that—together with simple rules for computing equal-time correlation functions—correctly reproduces the continuous-time limit of the Green's functions of a noninteracting Bose field. Interactions are then included in a self-consistent fashion by ensuring that the rules for Gaussian integration produce correct results for the interacting theory at all orders of perturbation theory. In the approach taken here, the functional integral is derived constructively in such a way that an unambiguous continuous-time notation emerges naturally from a properly defined (i.e., discretized) functional integral.

1. Functional representation of the Wigner function

The Weyl symbol of an arbitrary operator \hat{A} , denoted $\mathcal{A}_w(\psi)$, can be defined via the relation

$$\mathcal{A}_w(\psi) = \text{Tr}[\delta_w(\psi - \hat{a})\hat{A}]. \quad (\text{A1})$$

Here, the Weyl-ordered (and operator-valued) delta function is defined by

$$\delta_w(\psi - \hat{a}) = \frac{1}{\pi^2} \int d^2\varphi \exp[i\bar{\varphi}(\psi - \hat{a}) - \varphi(\bar{\psi} - \hat{a}^\dagger)]. \quad (\text{A2})$$

When convenient, the correspondence between operators and their Weyl symbols is indicated below with the notation $\hat{A} \leftrightarrow \mathcal{A}_w(\psi)$. Given the special role played by the density operator $\hat{\rho}$, it is traditional to use a special notation for its Weyl symbol, $\mathcal{W}(\psi, t)$, which is also called the Wigner function; the explicit time dependence is included because we will work in the Schrödinger picture, where the density matrix (and therefore the Wigner function) evolves in time.

The Weyl representation is intimately related to symmetrically ordered operator products; if an arbitrary operator \hat{A} is expanded in terms of symmetrically ordered operator products,

$$\hat{A} = \sum_{p,q} \mathcal{A}_{pq} [\hat{a}^p (\hat{a}^\dagger)^q]_s, \quad (\text{A3})$$

then the coefficients in this expansion determine the Weyl symbol $\mathcal{A}_w(\psi)$ in a particularly natural way:

$$\mathcal{A}_w(\psi) = \sum_{p,q} \mathcal{A}_{pq} \psi^p \bar{\psi}^q. \quad (\text{A4})$$

Given the Wigner function at an initial time t_0 , we would like to understand how it has changed a short time δt later due to the evolution of the density matrix by the master equation. During this time, the density matrix evolves according to $\hat{\rho}(t_0 + \delta t) = V_{\delta t}(\hat{\rho}(t_0))$, where the infinitesimal time-evolution superoperator $V_{\delta t}$ satisfies

$$\begin{aligned} V_{\delta t}(\star) &= 1 - i\delta t [\hat{H}, \star] + \delta t \frac{\kappa}{2} \sum_j (2\hat{a}_j \star \hat{a}_j^\dagger \\ &\quad - \star \hat{a}_j^\dagger \hat{a}_j - \hat{a}_j^\dagger \hat{a}_j \star) + O(\delta t^2). \end{aligned} \quad (\text{A5})$$

This transformation induces a corresponding evolution of the Wigner function, which for now we write formally as

$$\mathcal{W}(\psi_1, t_0 + \delta t) = \int d^2\psi_0 \mathcal{V}(\psi_1, \psi_0) \mathcal{W}(\psi_0, t_0), \quad (\text{A6})$$

thereby implicitly defining the infinitesimal phase-space propagator for the Wigner function, \mathcal{V} .

From the structure of Eqs. (A5) and (A6), it is clear that finding the explicit form of \mathcal{V} requires us to compute the Weyl symbol of products of $\hat{\rho}$ with creation and annihilation operators. To this end we define an operator-valued generating function

$$\hat{G} = e^{\eta\hat{a} + \bar{\eta}\hat{a}^\dagger} \hat{\rho}(t_0), \quad (\text{A7})$$

which can be differentiated to produce symmetrically ordered operator products

$$\partial_\eta^p \partial_{\bar{\eta}}^q \hat{G} \Big|_{\eta=0} = [\hat{a}^p (\hat{a}^\dagger)^q]_s \hat{\rho}(t_0). \quad (\text{A8})$$

Using Eqs. (A3) and (A8), we can expand the product of an arbitrary operator with the density matrix as

$$\hat{A} \hat{\rho}(t_0) = \sum_{p,q} \mathcal{A}_{pq} \partial_\eta^p \partial_{\bar{\eta}}^q \hat{G} \Big|_{\eta=0}, \quad (\text{A9})$$

giving us a prescription to compute the Weyl symbol of $\hat{A} \hat{\rho}(t_0)$ from the Weyl symbol of \hat{G} ,

$$\hat{A} \hat{\rho}(t_0) \leftrightarrow \sum_{p,q} \mathcal{A}_{pq} \partial_\eta^p \partial_{\bar{\eta}}^q \mathcal{G}_w(\psi_1) \Big|_{\eta=0}. \quad (\text{A10})$$

Making use of the standard operator phase-space correspondences [4],

$$\begin{aligned} \hat{a} \hat{\rho} &\leftrightarrow (\psi + \tfrac{1}{2} \partial_{\bar{\psi}}) \mathcal{W}(\psi), & \hat{a}^\dagger \hat{\rho} &\leftrightarrow (\bar{\psi} - \tfrac{1}{2} \partial_\psi) \mathcal{W}(\psi), \\ \hat{\rho} \hat{a} &\leftrightarrow (\psi - \tfrac{1}{2} \partial_{\bar{\psi}}) \mathcal{W}(\psi), & \hat{\rho} \hat{a}^\dagger &\leftrightarrow (\bar{\psi} + \tfrac{1}{2} \partial_\psi) \mathcal{W}(\psi), \end{aligned}$$

the commutation relations $[\partial_\psi, \psi] = 1$, $[\partial_\psi, \bar{\psi}] = 0$, and the Baker-Campbell-Hausdorff formula, the Weyl symbol of \hat{G} can be written

$$\mathcal{G}_w(\psi_1) = e^{\frac{1}{2}\eta\partial_{\bar{\psi}_1} - \frac{1}{2}\bar{\eta}\partial_{\psi_1}} e^{\eta\psi_1 + \bar{\eta}\bar{\psi}_1} \mathcal{W}(\psi_1, t_0). \quad (\text{A11})$$

Inserting a standard representation of the δ function and integrating by parts, we obtain

$$\begin{aligned} \mathcal{G}_w(\psi_1) &= \frac{4}{\pi^2} \int d^2\varphi_0 d^2\psi_0 e^{2\varphi_0(\bar{\psi}_1 - \bar{\psi}_0) - 2\bar{\varphi}_0(\psi_1 - \psi_0)} \\ &\quad \times e^{\eta(\varphi_0 + \psi_0) + \bar{\eta}(\bar{\varphi}_0 + \bar{\psi}_0)} \mathcal{W}(\psi_0, t_0). \end{aligned} \quad (\text{A12})$$

Note that the choice of a generating function that produced symmetrically ordered operator products also led to all of the derivatives appearing on the left in Eq. (A11), which enabled the integration by parts to proceed in a particularly simple manner to obtain Eq. (A12). Inserting Eq. (A12) into Eq. (A10) and then using Eq. (A4), we obtain

$$\begin{aligned} \hat{A} \hat{\rho} &\leftrightarrow \frac{4}{\pi^2} \int d^2\varphi_0 d^2\psi_0 e^{2\varphi_0(\bar{\psi}_1 - \bar{\psi}_0) - 2\bar{\varphi}_0(\psi_1 - \psi_0)} \\ &\quad \times \mathcal{A}_w(\psi_0 + \varphi_0) \mathcal{W}(\psi_0, t_0). \end{aligned} \quad (\text{A13})$$

Given Eq. (A13), we can now deduce Eq. (A6) from Eq. (A5). To first order in δt we find

$$\mathcal{W}(\psi_1, t_0 + \delta t) = \frac{4}{\pi^2} \int d^2\psi_0 d^2\varphi_0 e^{i\delta t \mathcal{L}(\psi_1, \varphi_1; \psi_0, \varphi_0)} \mathcal{W}(\psi_0, t_0), \quad (\text{A14})$$

where

$$\begin{aligned} \mathcal{L}(\psi_1, \varphi_1; \psi_0, \varphi_0) &= 2i\bar{\varphi}_0(\psi_1 - \psi_0)/\delta t - 2i\varphi_0(\bar{\psi}_1 - \bar{\psi}_0)/\delta t \\ &\quad - \mathcal{H}_w(\psi_0 + \varphi_0) + \mathcal{H}_w(\psi_0 - \varphi_0) \\ &\quad + i\kappa(2\bar{\varphi}_0\varphi_0 - \varphi_0\bar{\psi}_0 + \bar{\varphi}_0\psi_0). \end{aligned} \quad (\text{A15})$$

The Wigner function at a general time t can be obtained from the Wigner function at time t_0 by iteration of Eq. (A14). Breaking the interval $[t_0, t]$ into N segments of size $\delta t = (t - t_0)/N$, we obtain

$$\mathcal{W}(\psi_N, t) = \int \prod_{j=0}^{N-1} \left(\frac{4}{\pi^2} d^2\psi_j d^2\varphi_j \right) e^{iS} \mathcal{W}(\psi_0, t_0). \quad (\text{A16})$$

Here we have defined the discretized action

$$S = \sum_{j=0}^{N-1} \delta t \mathcal{L}(\psi_{j+1}, \varphi_{j+1}; \psi_j, \varphi_j). \quad (\text{A17})$$

Defining functional-integration measures that include the fields ψ_N, φ_N at the final time,

$$\mathcal{D}\psi = \prod_{j=0}^N d^2\psi_j, \quad \mathcal{D}\varphi = \prod_{j=0}^N \frac{4}{\pi^2} d^2\varphi_j, \quad (\text{A18})$$




FIG. 8. Time ordering of operators in Eq. (A20).

the trace of the Wigner function at time t can now be written

$$\mathcal{Z} \equiv \int \mathcal{D}\psi \mathcal{D}\varphi e^{i\mathcal{S}} \mathcal{W}(\psi_0, t_0) = 1. \quad (\text{A19})$$

The continuous-time limit ($N \rightarrow \infty, \delta t \rightarrow 0$) of Eqs. (A17)–(A19), generalized to many coherently coupled bosonic modes, yields Eqs. (5) and (6) of the main text.

2. Expectation values

The functional-integral representation of the Wigner function lends itself naturally to calculating correlation functions that are time ordered along the Keldysh contour (Fig. 8), e.g.,

$$\mathcal{C} = \langle \mathcal{T}_K(\hat{B}^1(t_1^-) \cdots \hat{B}^n(t_n^-) \hat{A}^1(t_1^+) \cdots \hat{A}^m(t_m^+)) \rangle. \quad (\text{A20})$$

Here, \mathcal{C} is written in the Heisenberg picture, and \mathcal{T}_K orders operators at times with a “+” superscript such that the times increase from right to left, and orders operators at times with a “−” superscript such that the times increase from left to right. Note that the Heisenberg picture is to be interpreted *before* the adiabatic elimination of the reservoir, which we assume proceeds within the Born-Markov approximation. Assuming without loss of generality that $t_j^\pm > t_{j-1}^\pm$, insisting that the operators be inserted at the discretized time slices chosen above ($t_j^\pm = \delta t \times r_j^\pm$, with r_j^\pm an integer between zero and N), and making the notational change $\hat{A}(\delta t r_j^\pm) \rightarrow \hat{A}(r_j^\pm)$, \mathcal{C} can be written

$$\mathcal{C} = \text{Tr}(\hat{A}^m(r_m^+) \cdots \hat{A}^1(r_1^+) \hat{\rho}(t_0) \hat{B}^1(r_1^-) \cdots \hat{B}^n(r_n^-)). \quad (\text{A21})$$

For our purposes we need to rewrite this correlation function in the Schrödinger picture, which can be accomplished with the help of the quantum-regression formula. For example, if $t_1^+ < t_1^-$, we have [4]

$$\mathcal{C} = \text{Tr}(\cdots U_{\delta t(r_1^- - r_1^+)}(\hat{A}^1 U_{\delta t r_1^+}(\hat{\rho}(t_0))) \hat{B}^1 \cdots), \quad (\text{A22})$$

where $U_{\delta t j} = V_{\delta t}^j$ is the time-evolution superoperator. The choice of Keldysh ordering is necessary and sufficient to guarantee that, in the quantum-regression formula, it will never be necessary to evolve backwards in time. Using the functional-integral expression for the Wigner function at time t , together with the phase-space correspondences for operators \hat{A} and \hat{B} , we find

$$\mathcal{C} = \int \mathcal{D}\psi \mathcal{D}\varphi e^{-i\mathcal{S}} \mathcal{W}(\psi_0, t_0) \quad (\text{A23})$$

$$\times (\cdots \mathcal{A}_w^1(\psi_{r_1^+} + \varphi_{r_1^+}) \mathcal{B}_w^1(\psi_{r_1^-} - \varphi_{r_1^-}) \cdots) \quad (\text{A24})$$

$$\equiv \langle \cdots \mathcal{A}_w^1(\psi_{r_1^+} + \varphi_{r_1^+}) \mathcal{B}_w^1(\psi_{r_1^-} - \varphi_{r_1^-}) \cdots \rangle_{\mathcal{Z}}. \quad (\text{A25})$$

Such correlation functions can be conveniently computed with respect to a generating functional by adding source terms to

action,

$$\mathcal{C} = \left(\cdots \mathcal{A}_w^1 \left(\frac{\partial}{\delta t \partial J_{r_1^+}} + \frac{\partial}{\delta t \partial K_{r_1^+}} \right) \times \mathcal{B}_w^1 \left(\frac{\partial}{\delta t \partial J_{r_1^-}} - \frac{\partial}{\delta t \partial K_{r_1^-}} \right) \cdots \right) \mathcal{Z}(\mathbf{J}, \mathbf{K}) \Big|_{\mathbf{J}, \mathbf{K} = 0}.$$

Here

$$\mathcal{Z}(\mathbf{J}, \mathbf{K}) = \int \mathcal{D}\psi \mathcal{D}\varphi e^{i\mathcal{S}(\mathbf{J}, \mathbf{K})} \mathcal{W}(\psi_0, t_0), \quad (\text{A26})$$

and

$$\mathcal{S}(\mathbf{J}, \mathbf{K}) = \mathcal{S} + \sum_{j=0}^N \delta t (\mathbf{J}_j \cdot \boldsymbol{\psi}_j + \mathbf{K}_j \cdot \boldsymbol{\varphi}_j), \quad (\text{A27})$$

with

$$\mathbf{J}_j = \begin{bmatrix} J_j \\ \bar{J}_j \end{bmatrix}, \quad \mathbf{K}_j = \begin{bmatrix} K_j \\ \bar{K}_j \end{bmatrix}, \quad \boldsymbol{\psi}_j = \begin{bmatrix} \psi_j \\ \bar{\psi}_j \end{bmatrix}, \quad \boldsymbol{\varphi}_j = \begin{bmatrix} \varphi_j \\ \bar{\varphi}_j \end{bmatrix}. \quad (\text{A28})$$

Restricting to the special case when the operators \hat{A} and \hat{B} are creation and annihilation operators, defining

$$\frac{1}{\delta t} \frac{\partial}{\partial J_j} = \frac{\delta}{\delta J(t_j)}, \quad \frac{1}{\delta t} \frac{\partial}{\partial K_j} = \frac{\delta}{\delta K(t_j)}, \quad (\text{A29})$$

and taking the continuum limit $\delta t \rightarrow 0$, we recover the expressions for correlation functions in Sec. III of the main text.

The correspondence given above can also be reversed in such a way to rewrite expectation values of the fields ψ and φ in terms of operator averages. To this end, we rearrange the operator phase-space correspondences as

$$\{\hat{a}, \hat{\rho}\} \leftrightarrow 2\psi \mathcal{W}, \quad \{\hat{a}^\dagger, \hat{\rho}\} \leftrightarrow 2\bar{\psi} \mathcal{W},$$

$$[\hat{a}, \hat{\rho}] \leftrightarrow \partial_{\bar{\psi}} \mathcal{W}, \quad [\hat{a}^\dagger, \hat{\rho}] \leftrightarrow -\partial_{\psi} \mathcal{W}.$$

The substitutions $\partial_{\psi} \leftrightarrow -2\bar{\varphi}$ and $\partial_{\bar{\psi}} \leftrightarrow 2\varphi$ are valid under the functional-integration sign, so long as the derivative sits to the left of any other instances of the field ψ evaluated at the same time (not including those occurring in the action). Thus we are led to the identifications

$$\psi \mathcal{W} \leftrightarrow \frac{1}{2} \{\hat{a}, \hat{\rho}\}, \quad \varphi \mathcal{W} \leftrightarrow \frac{1}{2} [\hat{a}, \hat{\rho}], \quad (\text{A30})$$

with the understanding that when products of the fields ψ and φ at the same time arise, we should take the commutators *after* the anticommutators in the corresponding operator expectation values. This identification leads very directly to correlation functions of the classical field. For example, assuming $r_3 > r_2 > r_1$ we have

$$\langle \bar{\psi}_{r_1} \psi_{r_2} \bar{\psi}_{r_3} \rangle_{\mathcal{Z}} = \text{Tr}(\{\hat{a}^\dagger(r_3), \{\hat{a}(r_2), \{\hat{a}^\dagger(r_1), \hat{\rho}(t_0)\}\}\}). \quad (\text{A31})$$

Note that if $r_1 = r_2 = r_3 \equiv r$, this correlation function simplifies to

$$\langle \bar{\psi}_r \psi_r \bar{\psi}_r \rangle_{\mathcal{Z}} = \text{Tr}([\hat{a}^\dagger(r) \hat{a}(r) \hat{a}^\dagger(r)]_{\mathcal{S}}). \quad (\text{A32})$$

Since we just have a few fields in the correlation function, this result can be worked out by direct comparison of the right-hand sides of Eqs. (A31) and (A32) (using the commutation relation

$[\hat{a}, \hat{a}^\dagger] = 1$). More generally, it also follows by using the phase-space correspondence $[\hat{a}^\dagger \hat{a} \hat{a}^\dagger]_s \leftrightarrow (\bar{\psi} + \bar{\varphi})^2 (\psi + \varphi)$. Going back to operator expectation values by using Eq. (A30), and remembering the rule that commutators are taken after anticommutators when the fields φ and ψ are evaluated at the same times, terms with one or more power of φ result in an outer commutator that gets killed by the trace, so only the $\bar{\psi}^2 \psi$ term survives. The generalization of Eq. (A32) to arbitrary equal-time correlation functions of the classical field ψ leads, in the continuum-time limit and for more than one site, to Eq. (9) of the main text.

APPENDIX B: ADIABATIC ELIMINATION OF THE MASSIVE FIELD

Here we describe the perturbative adiabatic elimination of the field ϱ near the mean-field critical point. For simplicity, we first treat the case with no fluctuations ($\mathcal{N} \rightarrow \infty$), and afterwards we consider the effect of weak fluctuations.

1. Mean-field theory

Substitution of Eq. (22) into Eq. (15) yields coupled equations for ϱ and σ that are fully equivalent to the mean-field dynamics of Ψ . To simplify the following expressions we convert all energy and time scales into dimensionless ratios with the chemical potential μ , and in a slight abuse of notation we do not change any of the associated symbols—the μ dependence can be unambiguously restored by insisting on dimensional consistency. After some algebra, we find

$$\dot{\varrho} = -\frac{2}{\sqrt{3}}(\varrho - \varrho_0) - r\varrho - \frac{J\nabla^2}{\sqrt{3}}(\varrho + 2\sigma) + \frac{u}{\sqrt{3}}(2\sigma^3 + 3\sigma\varrho^2 + 3\sigma\sigma^2 + \varrho^3 - \sigma^2\sqrt{6/u}), \quad (\text{B1})$$

$$\dot{\sigma} = -\frac{h}{2} - r\sigma + \frac{J\nabla^2}{\sqrt{3}}(\sigma + 2\varrho) - \frac{u}{\sqrt{3}}(\varrho^2\sqrt{6/u} + 2\varrho^3 + 3\sigma\varrho^2 + 3\sigma\sigma^2 + \sigma^3), \quad (\text{B2})$$

where

$$\varrho_0 = \frac{h\sqrt{3}}{8} - \frac{r}{\sqrt{8u}}. \quad (\text{B3})$$

Several simplifications can now be made. First, because we are interested in dynamics near the critical point and after the field has nearly relaxed, h , r , ϱ , and σ can all be treated as small parameters. Though we do not know *a priori* how small the fields ϱ and σ are relative to the parameters r and h , it is perfectly consistent to keep, at any particular order in one of the parameters, only the lowest nontrivial order in any other parameter. Moreover, while we do not know how small spatial derivatives of the field are, we do expect them to be small compared to the fields themselves, which should vary slowly in space near the critical point, and thus we formally treat $J\nabla^2$ as an additional small parameter (one can check that this assumption is self-consistent at the end of the calculation, where it is seen that $J\nabla^2 \sim r$). Following this

logic, and introducing the rescaled parameters $K = J/\sqrt{3}$ and $g = u/\sqrt{3}$ used in the main text, we arrive at the simplified equations

$$\dot{\varrho} = -\frac{2}{\sqrt{3}}(\varrho - \varrho_0) - 2K\nabla^2\sigma + O(\sigma^2), \quad (\text{B4})$$

$$\dot{\sigma} = -\frac{h}{2} - r\sigma + K\nabla^2(\sigma + 2\varrho) - g\sigma^3 + O(\varrho^2) + O(\varrho\sigma^2). \quad (\text{B5})$$

Terms that are kept with the O notation are there to remind us that we do not know for sure whether they are parametrically small compared to other terms that are kept—they turn out to be unimportant for reasons explained below, which is why we do not keep track of the exact coefficients.

The justification for adiabatically eliminating ϱ near the critical point is now clear: As $r \rightarrow 0$, the term proportional to ϱ in Eq. (B4) stays finite, indicating that ϱ relaxes to zero exponentially in time even at the critical point (once μ is restored, we see that it decays on a time scale $\sim 1/\mu$). On the other hand, the term linear in σ in Eq. (B5) vanishes as $r \rightarrow 0$, indicating a divergent time scale for relaxation of σ (which relaxes algebraically precisely at the critical point, $r = 0$). To adiabatically eliminate ϱ we set the rhs of Eq. (B4) to zero, obtaining

$$\varrho = \varrho_0 - \sqrt{3}K\nabla^2\sigma + O(\sigma^2), \quad (\text{B6})$$

and then substitute this result into Eq. (B5). Many derivative terms are generated, but working to lowest order in $K\nabla^2$ we find

$$\dot{\sigma} = K\nabla^2\sigma - r\sigma - g\sigma^3 - \frac{h}{2} + O(r^2) + O(r\sigma^2). \quad (\text{B7})$$

Here we have implicitly assumed that $|h| < |r|$ to write $\varrho_0 = O(r)$, which poses no important restriction on what follows. If we ignore the final two terms and solve Eq. (B7) at $h = 0$ and $r < 0$ (inside the bistable region), we find two uniform solutions at $\sigma = \pm\sqrt{|r|/g}$, which sets the scale of σ in the relevant near-critical dynamics, $\sigma \sim \sqrt{r}$. From this scaling, it is easily seen that the final two terms in Eq. (B7) are parametrically smaller than the others; thus we were justified in dropping them, which results in Eq. (23) in the main text.

2. Fluctuations

Substitution of Eq. (22) into Eq. (16) yields coupled stochastic equations for ϱ and σ that are fully equivalent to the nonequilibrium Langevin equation for Ψ . In the limit of weak noise, an expansion in ϱ and σ is still justified. Moreover, since arbitrarily weak noise will induce arbitrarily small excursions away from the mean-field stationary state, much of the analysis that led to an effective relaxational description of σ (that is approximately decoupled from ϱ) should remain valid, as it assumed nothing more than being close to both the critical point and the steady state. Working near the mean-field critical point, the same assumptions that lead from Eqs. (B1) and (B2) to Eqs. (B4) and (B5) remain justified and yield (for now

keeping all gradient terms)

$$\dot{\varrho} = -\frac{2}{\sqrt{3}}(\varrho - \varrho_0) - K\nabla^2(\varrho + 2\sigma) + O(\sigma^2) + \eta(\tau), \quad (\text{B8})$$

$$\begin{aligned} \dot{\sigma} = & -\frac{h}{2} - r\sigma + K\nabla^2(\sigma + 2\varrho) - g\sigma^3 + O(\varrho^2) \\ & + O(\varrho\sigma^2) + \xi(\tau). \end{aligned} \quad (\text{B9})$$

In order to avoid confusion regarding the noise variances, here we have chosen a new symbol for the dimensionless time, $t\mu \equiv \tau$ (so $\dot{\sigma} = d\sigma/d\tau$, etc.), even though we continue to use the same symbols for the (now dimensionless) energies r , h , K , and g . The real noises η and ξ are defined as

$$\eta(\tau) = \frac{\zeta_I(t)}{\mu} - \frac{\zeta_R(t)}{\mu\sqrt{3}}, \quad \xi(\tau) = \frac{2\zeta_I(t)}{\mu\sqrt{3}}, \quad (\text{B10})$$

where ζ_R and ζ_I are the real and imaginary components, respectively, of the complex Gaussian white noise $\zeta(t)$ in Eq. (16). Thus from the variances of $\zeta(t)$ we have

$$\overline{\eta_j(\tau_1)\eta_k(\tau_2)} = \overline{\xi_j(\tau_1)\xi_k(\tau_2)} = \frac{\kappa}{3\mathcal{N}}\delta_{j,k}\delta(\tau_1 - \tau_2), \quad (\text{B11})$$

where one factor of $1/\mu$ has been absorbed into the now dimensionless κ and one is used to change variables from t to τ in the δ function.

We would like to eliminate the terms in Eq. (B9) that involve ϱ , but, strictly speaking, the adiabatic elimination of ϱ by setting $\dot{\varrho} = 0$ in Eq. (B8) is no longer justified. Nevertheless, for weak noise and near the critical point, there will be a separation of time scales, length scales, and typical sizes of the fluctuations of σ and ϱ . Because ϱ remains massive at the mean-field critical point while σ does not, we expect that in the presence of noise the scale of typical fluctuations for ϱ will be small compared to the scale of typical fluctuations for σ . Likewise, σ will relax more slowly than ϱ near the critical point, and will exhibit fluctuations on a longer length scale, i.e., in the presence of fluctuations it will have a longer autocorrelation time than ϱ , and will be roughly spatially homogeneous over the length scale on which ϱ is correlated. On the grounds of the latter statement, it is justified to solve Eq. (B8) at *fixed* σ and within a local-density approximation (i.e., we assume that ϱ is relaxing in a locally homogeneous environment set by the slowly varying value of σ), in which case ϱ simply undergoes a lattice version of Ornstein-Uhlenbeck relaxation to the mean value

$$\bar{\varrho} = \varrho_0 - \sqrt{3}K\nabla^2\sigma + O(\sigma^2). \quad (\text{B12})$$

Straightforward analysis reveals that fluctuations of ϱ around its mean value, $\vartheta(\tau) = \varrho(\tau) - \bar{\varrho}(t)$, obey

$$\overline{\vartheta_j(\tau_1)\vartheta_k(\tau_2)} \sim \frac{\kappa}{\mathcal{N}}e^{-|j-k|/\ell_\varrho}e^{-|\tau_2-\tau_1|/\tau_\varrho}, \quad (\text{B13})$$

where $\ell_\varrho \sim \sqrt{K}$ and $\tau_\varrho \sim 1$ are the correlation length and correlation time of the massive field ϱ , respectively. Inserting the solution $\varrho = \bar{\varrho} + \vartheta$ into Eq. (B9), and keeping for now all

terms that involve the fluctuations ϑ , we obtain

$$\begin{aligned} \dot{\sigma} = & -\frac{h}{2} - r\sigma + K\nabla^2\sigma - g\sigma^3 + \xi(\tau) \\ & + 2K\nabla^2\bar{\varrho} + 2K\nabla^2\vartheta + O(\bar{\varrho}^2) + O(\bar{\varrho})\vartheta + O(\vartheta^2). \end{aligned} \quad (\text{B14})$$

The first and third terms on the second line, $2K\nabla^2\bar{\varrho}$ and $O(\bar{\varrho}^2)$, contain terms that are either higher order in σ , in r , or in gradients than other terms on the first line, and thus can be ignored. The remaining terms depend on ϑ ; because the dynamics of σ is slow and dominated by long-wavelength fluctuations, we can approximate ϑ as spatially uncorrelated white noise with variance

$$\overline{\vartheta_j(\tau_1)\vartheta_k(\tau_2)} \sim \frac{\kappa}{\mathcal{N}}\delta_{j,k}\delta(\tau_1 - \tau_2). \quad (\text{B15})$$

In light of this approximation, the term $K\nabla^2\vartheta$ acts as an additional source of additive white noise—it can in principle be taken into account as a (K -dependent) renormalization of the noise $\xi(\tau)$ that is already present, but if we work to lowest order in K this renormalization can be ignored. The term $O(\bar{\varrho})\vartheta$ acts as a source of multiplicative noise; since $\bar{\varrho}$ is itself a small parameter, this noise can also be ignored in comparison to the already present white noise $\xi(\tau)$. The term $O(\vartheta^2)$ should be interpreted by writing $\vartheta(\tau)^2 = f(\tau) + \chi(\tau)$, where f is the average of ϑ^2 and χ is its fluctuations. The average $f \sim \kappa/\mathcal{N}$ can be ignored for weak noise, while the fluctuations χ obey (using the fact that ϑ is a Gaussian variable)

$$\overline{\chi_j(\tau_1)\chi_k(\tau_2)} \sim (\kappa/\mathcal{N})^2 e^{-2|j-k|/\ell_\varrho} e^{-2|\tau_2-\tau_1|/\tau_\varrho}. \quad (\text{B16})$$

As with ϑ , χ can be interpreted as spatially uncorrelated Gaussian white noise as far as the slow dynamics of σ is concerned. Since it has a variance that is parametrically smaller than that of ξ , it can once again be ignored. With all of the terms on the second line of Eq. (B14) dropped, we recover Eq. (32) of the main text.

APPENDIX C: DOMAIN-WALL VELOCITY

As described in the main text, the domain-wall velocity can be determined by solving for the dynamics of a fictitious particle obeying the equation of motion

$$K\ddot{\sigma} = -v\dot{\sigma} + \frac{\partial U(\sigma)}{\partial \sigma}, \quad (\text{C1})$$

where

$$U(\sigma) = \frac{1}{2}(r\sigma^2 + \frac{1}{2}g\sigma^4 + h\sigma). \quad (\text{C2})$$

Here the domain-wall velocity v plays the role of a velocity-dependent friction coefficient. We seek solutions of Eq. (C1) for which the particle starts (with $\dot{\sigma} = 0$) at the higher local maximum of $-U(\sigma)$ and comes to rest at the lower one. When h is small the two local maxima have similar energies, and the friction coefficient v must also be small for the particle to reach the top of the lower potential maximum. Thus the trajectory is similar to that in the case of $h = 0$, for which (from conservation of energy) $K\dot{\sigma}^2/2 = U(\sigma) - U(\sigma_0)$, such that

$$\dot{\sigma} \approx \sqrt{\frac{2[U(\sigma) - U(\sigma_0)]}{K}}. \quad (\text{C3})$$

The work done along this trajectory by the friction must be equal to the change in potential energy,

$$-v \int_{\sigma_-}^{\sigma_+} \dot{\sigma} d\sigma = U(\sigma_+) - U(\sigma_-) \approx h\sigma_0. \quad (\text{C4})$$

Inserting Eq. (C3) into Eq. (C4), taking the integral, and using $U(\sigma_0) = r^2/4g$, we obtain the domain-wall velocity given in

the main text,

$$v \approx h \frac{3}{2} \sqrt{\frac{Kg}{2r^2}}. \quad (\text{C5})$$

-
- [1] R. H. Dicke, Coherence in spontaneous radiation processes, *Phys. Rev.* **93**, 99 (1954).
- [2] Assa Auerbach, *Interacting Electrons and Quantum Magnetism* (Springer-Verlag, New York, 1994).
- [3] Thierry Giamarchi, *Quantum Physics in One Dimension* (Oxford University Press, New York, 2004).
- [4] Howard Carmichael, *Statistical Methods in Quantum Optics I* (Springer-Verlag, Berlin, 2010).
- [5] Mehran Kardar, *Statistical Physics of Fields* (Cambridge University Press, Cambridge, UK, 2007).
- [6] Iacopo Carusotto and Cristiano Ciuti, Quantum fluids of light, *Rev. Mod. Phys.* **85**, 299 (2013).
- [7] L. M. Sieberer, M. Buchhold, and S. Diehl, Keldysh field theory for driven open quantum systems, *Rep. Prog. Phys.* **79**, 096001 (2016).
- [8] Changsuk Noh and Dimitris G Angelakis, Quantum simulations and many-body physics with light, *Rep. Prog. Phys.* **80**, 016401 (2017).
- [9] Michael J. Hartmann, Quantum simulation with interacting photons, *J. Opt.* **18**, 104005 (2016).
- [10] Hui Deng, Gregor Weihs, Charles Santori, Jacqueline Bloch, and Yoshihisa Yamamoto, Condensation of semiconductor microcavity exciton polaritons, *Science* **298**, 199 (2002).
- [11] J. Kasprzak, M. Richard, S. Kundermann, A. Baas, P. Jeambrun, J. M. J. Keeling, F. M. Marchetti, M. H. Szymanska, R. Andre, J. L. Staehli, V. Savona, P. B. Littlewood, B. Deveaud, and Le Si Dang, Bose-Einstein condensation of exciton polaritons, *Nature (London)* **443**, 409 (2006).
- [12] M. H. Szymańska, J. Keeling, and P. B. Littlewood, Nonequilibrium Quantum Condensation in an Incoherently Pumped Dissipative System, *Phys. Rev. Lett.* **96**, 230602 (2006).
- [13] Tim Byrnes, Na Young Kim, and Yoshihisa Yamamoto, Exciton-polariton condensates, *Nat. Phys.* **10**, 803 (2014).
- [14] S. R. K. Rodriguez, A. Amo, I. Sagnes, L. Le Gratiet, E. Galopin, A. Lemaître, and J. Bloch, Interaction-induced hopping phase in driven-dissipative coupled photonic microcavities, *Nat. Commun.* **7**, 11887 (2016).
- [15] S. R. K. Rodriguez, W. Casteels, F. Storme, I. Sagnes, L. Le Gratiet, E. Galopin, A. Lemaître, A. Amo, C. Ciuti, and J. Bloch, Dynamic optical hysteresis in the quantum regime, [arXiv:1608.00260](https://arxiv.org/abs/1608.00260).
- [16] Andrew A. Houck, Hakan E. Tureci, and Jens Koch, On-chip quantum simulation with superconducting circuits, *Nat. Phys.* **8**, 292 (2012).
- [17] Felix Nissen, Sebastian Schmidt, Matteo Biondi, Gianni Blatter, Hakan E. Türeci, and Jonathan Keeling, Nonequilibrium Dynamics of Coupled Qubit-Cavity Arrays, *Phys. Rev. Lett.* **108**, 233603 (2012).
- [18] J. Raftery, D. Sadri, S. Schmidt, H. E. Türeci, and A. A. Houck, Observation of a Dissipation-Induced Classical to Quantum Transition, *Phys. Rev. X* **4**, 031043 (2014).
- [19] R. Barends, L. Lamata, J. Kelly, L. García-Álvarez, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Yu Chen, Z. Chen, B. Chiaro, A. Dunsworth, I. C. Hoi, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, E. Solano, and John M. Martinis, Digital quantum simulation of fermionic models with a superconducting circuit, *Nat. Commun.* **6**, 7654 (2015).
- [20] Mattias Fitzpatrick, Neereja M. Sundaresan, Andy C. Y. Li, Jens Koch, and Andrew A. Houck, Observation of a Dissipative Phase Transition in a One-Dimensional Circuit QED Lattice, *Phys. Rev. X* **7**, 011016 (2017).
- [21] Andrew D. Greentree, Charles Tahan, Jared H. Cole, and Lloyd C. L. Hollenberg, Quantum phase transitions of light, *Nat. Phys.* **2**, 856 (2006).
- [22] S. M. Spillane, G. S. Pati, K. Salit, M. Hall, P. Kumar, R. G. Beausoleil, and M. S. Shahriar, Observation of Nonlinear Optical Interactions of Ultralow Levels of Light in a Tapered Optical Nanofiber Embedded in a Hot Rubidium Vapor, *Phys. Rev. Lett.* **100**, 233602 (2008).
- [23] E. Vetsch, D. Reitz, G. Sagué, R. Schmidt, S. T. Dawkins, and A. Rauschenbeutel, Optical Interface Created by Laser-Cooled Atoms Trapped in the Evanescent Field Surrounding an Optical Nanofiber, *Phys. Rev. Lett.* **104**, 203603 (2010).
- [24] T. G. Tiecke, J. D. Thompson, N. P. de Leon, L. R. Liu, V. Vuletić, and M. D. Lukin, Nanophotonic quantum phase switch with a single atom, *Nature (London)* **508**, 241 (2014).
- [25] J. D. Thompson, T. G. Tiecke, N. P. de Leon, J. Feist, A. V. Akimov, M. Gullans, A. S. Zibrov, V. Vuletić, and M. D. Lukin, Coupling a single trapped atom to a nanoscale optical cavity, *Science* **340**, 1202 (2013).
- [26] A. Goban, C. L. Hung, S. P. Yu, J. D. Hood, J. A. Muniz, J. H. Lee, M. J. Martin, A. C. McClung, K. S. Choi, D. E. Chang, O. Painter, and H. J. Kimble, Atom-light interactions in photonic crystals, *Nat. Commun.* **5**, 3808 (2014).
- [27] K. Hennessy, A. Badolato, M. Winger, D. Gerace, M. Atature, S. Gulde, S. Falt, E. L. Hu, and A. Imamoglu, Quantum nature of a strongly coupled single quantum dot-cavity system, *Nature (London)* **445**, 896 (2007).
- [28] Takao Aoki, Barak Dayan, E. Wilcut, W. P. Bowen, A. S. Parkins, T. J. Kippenberg, K. J. Vahala, and H. J. Kimble, Observation of strong coupling between one atom and a monolithic microresonator, *Nature (London)* **443**, 671 (2006).
- [29] Alexey V. Gorshkov, Johannes Otterbach, Michael Fleischhauer, Thomas Pohl, and Mikhail D. Lukin, Photon-Photon

- Interactions via Rydberg Blockade, *Phys. Rev. Lett.* **107**, 133602 (2011).
- [30] Thibault Peyronel, Ofer Firstenberg, Qi-Yu Liang, Sebastian Hofferberth, Alexey V. Gorshkov, Thomas Pohl, Mikhail D. Lukin, and Vladan Vuletic, Quantum nonlinear optics with single photons enabled by strongly interacting atoms, *Nature (London)* **488**, 57 (2012).
- [31] C. Carr, R. Ritter, C. G. Wade, C. S. Adams, and K. J. Weatherill, Nonequilibrium Phase Transition in a Dilute Rydberg Ensemble, *Phys. Rev. Lett.* **111**, 113901 (2013).
- [32] N. Malossi, M. M. Valado, S. Scotto, P. Huillery, P. Pillet, D. Ciampini, E. Arimondo, and O. Morsch, Full Counting Statistics and Phase Diagram of a Dissipative Rydberg Gas, *Phys. Rev. Lett.* **113**, 023006 (2014).
- [33] Sarang Gopalakrishnan, Benjamin L. Lev, and Paul M. Goldbart, Emergent crystallinity and frustration with Bose-Einstein condensates in multimode cavities, *Nat. Phys.* **5**, 845 (2009).
- [34] Sarang Gopalakrishnan, Benjamin L. Lev, and Paul M. Goldbart, Atom-light crystallization of Bose-Einstein condensates in multimode cavities: Nonequilibrium classical and quantum phase transitions, emergent lattices, supersolidity, and frustration, *Phys. Rev. A* **82**, 043612 (2010).
- [35] Tony E. Lee, Sarang Gopalakrishnan, and Mikhail D. Lukin, Unconventional Magnetism via Optical Pumping of Interacting Spin Systems, *Phys. Rev. Lett.* **110**, 257204 (2013).
- [36] Michael Hoening, Wildan Abdussalam, Michael Fleischhauer, and Thomas Pohl, Antiferromagnetic long-range order in dissipative Rydberg lattices, *Phys. Rev. A* **90**, 021603 (2014).
- [37] K. Macieszczak, M. Guta, I. Lesanovsky, and J. P. Garrahan, Towards a Theory of Metastability in Open Quantum Dynamics, *Phys. Rev. Lett.* **116**, 240404 (2016).
- [38] Dominic C. Rose, Katarzyna Macieszczak, Igor Lesanovsky, and Juan P. Garrahan, Metastability in an open quantum Ising model, *Phys. Rev. E* **94**, 052132 (2016).
- [39] L. M. Sieberer, S. D. Huber, E. Altman, and S. Diehl, Dynamical Critical Phenomena in Driven-Dissipative Systems, *Phys. Rev. Lett.* **110**, 195301 (2013).
- [40] Jiasen Jin, Alberto Biella, Oscar Viyuela, Leonardo Mazza, Jonathan Keeling, Rosario Fazio, and Davide Rossini, Cluster Mean-Field Approach to the Steady-State Phase Diagram of Dissipative Spin Systems, *Phys. Rev. X* **6**, 031011 (2016).
- [41] Ehud Altman, Lukas M. Sieberer, Leiming Chen, Sebastian Diehl, and John Toner, Two-Dimensional Superfluidity of Exciton Polaritons Requires Strong Anisotropy, *Phys. Rev. X* **5**, 011017 (2015).
- [42] Hendrik Weimer, Variational Principle for Steady States of Dissipative Quantum Many-Body Systems, *Phys. Rev. Lett.* **114**, 040402 (2015).
- [43] V. R. Overbeck, M. F. Maghrebi, A. V. Gorshkov, and H. Weimer, Multicritical behavior in dissipative Ising models, [arXiv:1606.08863](https://arxiv.org/abs/1606.08863).
- [44] W. Casteels, F. Storme, A. Le Boité, and C. Ciuti, Power laws in the dynamic hysteresis of quantum nonlinear photonic resonators, *Phys. Rev. A* **93**, 033824 (2016).
- [45] Michael J. Hartmann, Fernando G. S. L. Brandao, and Martin B. Plenio, Strongly interacting polaritons in coupled arrays of cavities, *Nat. Phys.* **2**, 849 (2006).
- [46] I. Carusotto, D. Gerace, H. E. Türeci, S. De Liberato, C. Ciuti, and A. Imamoglu, Fermionized Photons in an Array of Driven Dissipative Nonlinear Cavities, *Phys. Rev. Lett.* **103**, 033601 (2009).
- [47] Alexandre Le Boité, Giuliano Orso, and Cristiano Ciuti, Steady-State Phases and Tunneling-Induced Instabilities in the Driven Dissipative Bose-Hubbard Model, *Phys. Rev. Lett.* **110**, 233601 (2013).
- [48] Ryan M. Wilson, Khan W. Mahmud, Anzi Hu, Alexey V. Gorshkov, Mohammad Hafezi, and Michael Foss-Feig, Collective phases of strongly interacting cavity photons, *Phys. Rev. A* **94**, 033801 (2016).
- [49] C. W. Lai, N. Y. Kim, S. Utsunomiya, G. Roumpos, H. Deng, M. D. Fraser, T. Byrnes, P. Recher, N. Kumada, T. Fujisawa, and Y. Yamamoto, Coherent zero-state and π -state in an exciton-polariton condensate array, *Nature (London)* **450**, 529 (2007).
- [50] T. Jacqmin, I. Carusotto, I. Sagnes, M. Abbarchi, D. D. Solnyshkov, G. Malpuech, E. Galopin, A. Lemaître, J. Bloch, and A. Amo, Direct Observation of Dirac Cones and a Flatband in a Honeycomb Lattice for Polaritons, *Phys. Rev. Lett.* **112**, 116402 (2014).
- [51] P. Hamel, S. Haddadi, F. Raineri, P. Monnier, G. Beaudoin, Isabelle Sagnes, Ariel Levenson, and Alejandro M. Yacomotti, Spontaneous mirror-symmetry breaking in coupled photonic-crystal nanolasers, *Nat. Photon.* **9**, 311 (2015).
- [52] D. G. Angelakis, M. F. Santos, and S. Bose, Photon-blockade-induced Mott transitions and xy spin models in coupled cavity arrays, *Phys. Rev. A* **76**, 031805 (2007).
- [53] M. Schiró, M. Bordyuh, B. Öztop, and H. E. Türeci, Phase Transition of Light in Cavity QED Lattices, *Phys. Rev. Lett.* **109**, 053601 (2012).
- [54] M. Hafezi, P. Adhikari, and J. M. Taylor, Chemical potential for light by parametric coupling, *Phys. Rev. B* **92**, 174305 (2015).
- [55] E. Kapit, M. Hafezi, and S. H. Simon, Induced Self-Stabilization in Fractional Quantum Hall States of Light, *Phys. Rev. X* **4**, 031039 (2014).
- [56] Elliott H. Lieb and W. Liniger, Exact analysis of an interacting Bose gas. I. The general solution and the ground state, *Phys. Rev.* **130**, 1605 (1963).
- [57] Matthew P. A. Fisher, Peter B. Weichman, G. Grinstein, and Daniel S. Fisher, Boson localization and the superfluid-insulator transition, *Phys. Rev. B* **40**, 546 (1989).
- [58] Alexandre Le Boité, Giuliano Orso, and Cristiano Ciuti, Bose-Hubbard model: Relation between driven-dissipative steady states and equilibrium quantum phases, *Phys. Rev. A* **90**, 063821 (2014).
- [59] José Lebreuilly, Iacopo Carusotto, and Michiel Wouters, Towards strongly correlated photons in arrays of dissipative nonlinear cavities under a frequency-dependent incoherent pumping, *C. R. Phys.* **17**, 836 (2016).
- [60] R. O. Umucalılar and I. Carusotto, Fractional Quantum Hall States of Photons in an Array of Dissipative Coupled Cavities, *Phys. Rev. Lett.* **108**, 206809 (2012).
- [61] Dario Gerace, Hakan E. Türeci, Atac Imamoglu, Vittorio Giovannetti, and Rosario Fazio, The quantum-optical Josephson interferometer, *Nat. Phys.* **5**, 281 (2009).
- [62] Michael J. Hartmann, Polariton Crystallization in Driven Arrays of Lossy Nonlinear Resonators, *Phys. Rev. Lett.* **104**, 113601 (2010).
- [63] T. Grujic, S. R. Clark, D. Jaksch, and D. G. Angelakis, Repulsively induced photon superbunching in driven resonator arrays, *Phys. Rev. A* **87**, 053846 (2013).

- [64] M. Biondi, G. Blatter, H. E. Türeci, and S. Schmidt, Nonequilibrium phase diagram of the driven-dissipative photonic lattice, [arXiv:1611.00697](#).
- [65] Subir Sachdev, *Quantum Phase Transitions* (Cambridge University Press, Cambridge, UK, 2011).
- [66] See Ref. [7] for a detailed discussion of the subtle connection between symmetries and conservation laws in driven-dissipative systems.
- [67] P. D. Drummond and D. F. Walls, Quantum theory of optical bistability. I. Nonlinear polarisability model, *J. Phys. A: Math. Gen.* **13**, 725 (1980).
- [68] Emanuele G. Dalla Torre, S. Diehl, M. D. Lukin, S. Sachdev, and P. Strack, Keldysh approach for nonequilibrium phase transitions in quantum optics: Beyond the Dicke model in optical cavities, *Phys. Rev. A* **87**, 023831 (2013).
- [69] Mohammad F. Maghrebi and Alexey V. Gorshkov, Nonequilibrium many-body steady states via Keldysh formalism, *Phys. Rev. B* **93**, 014307 (2016).
- [70] Ching-Kit Chan, Tony E. Lee, and Sarang Gopalakrishnan, Limit-cycle phase in driven-dissipative spin systems, *Phys. Rev. A* **91**, 051601 (2015).
- [71] Matteo Marcuzzi, Michael Buchhold, Sebastian Diehl, and Igor Lesanovsky, Absorbing State Phase Transition with Competing Quantum and Classical Fluctuations, *Phys. Rev. Lett.* **116**, 245701 (2016).
- [72] Jamir Marino and Sebastian Diehl, Driven Markovian Quantum Criticality, *Phys. Rev. Lett.* **116**, 070407 (2016).
- [73] Matteo Marcuzzi, Emanuele Levi, Sebastian Diehl, Juan P. Garrahan, and Igor Lesanovsky, Universal Nonequilibrium Properties of Dissipative Rydberg Gases, *Phys. Rev. Lett.* **113**, 210401 (2014).
- [74] J. H. Marburger and F. S. Felber, Theory of a lossless nonlinear Fabry-Perot interferometer, *Phys. Rev. A* **17**, 335 (1978).
- [75] P. D. Drummond and C. W. Gardiner, Generalised p-representations in quantum optics, *J. Phys. A: Math. Gen.* **13**, 2353 (1980).
- [76] H. J. Carmichael, Breakdown of Photon Blockade: A Dissipative Quantum Phase Transition in Zero Dimensions, *Phys. Rev. X* **5**, 031028 (2015).
- [77] J. M. Fink, A. Dombi, A. Vukics, A. Wallraff, and P. Domokos, Observation of the Photon-Blockade Breakdown Phase Transition, *Phys. Rev. X* **7**, 011012 (2017).
- [78] P. Kinsler and P. D. Drummond, Quantum dynamics of the parametric oscillator, *Phys. Rev. A* **43**, 6194 (1991).
- [79] Wim Casteels, Rosario Fazio, and Christiano Ciuti, Critical scaling of the Liouvillian gap for a nonlinear driven-dissipative resonator, *Phys. Rev. A* **95**, 012128 (2017).
- [80] The large-density limit on a single site is closely related to the thermodynamic limit (large system-size limit at fixed density) of a model with infinite-range interactions, for which the interaction energy is also proportional to the square of the number of particles.
- [81] B. Berg, L. I. Plimak, A. Polkovnikov, M. K. Olsen, M. Fleischhauer, and W. P. Schleich, Commuting Heisenberg operators as the quantum response problem: Time-normal averages in the truncated Wigner representation, *Phys. Rev. A* **80**, 033624 (2009).
- [82] Many of the usual caveats about continuum-time notation for functional integrals apply here, and the interested reader should consult Appendix A for a proper discrete-time regularization of \mathcal{Z} .
- [83] Anatoli Polkovnikov, Phase space representation of quantum dynamics, *Ann. Phys.* **325**, 1790 (2010).
- [84] Anatoli Polkovnikov, Quantum corrections to the dynamics of interacting bosons: Beyond the truncated Wigner approximation, *Phys. Rev. A* **68**, 053604 (2003).
- [85] Melvin Lax, Formal theory of quantum fluctuations from a driven state, *Phys. Rev.* **129**, 2342 (1963).
- [86] A. Kamenev and A. Levchenko, Keldysh technique and nonlinear σ -model: Basic principles and applications, *Adv. Phys.* **58**, 197 (2009).
- [87] Alex Kamenev, *Field Theory of Non-Equilibrium Systems* (Cambridge University Press, New York, 2011).
- [88] Note that the Heisenberg picture is to be interpreted with respect to the Hamiltonian describing the system-bath pair *before* the adiabatic elimination of the bath, which we assume proceeds within the Born-Markov approximation.
- [89] Equation (16) is closely related to the truncated Wigner approximation, though there is a subtle difference; the truncated Wigner approximation would result in an interaction term $u(|\Psi|^2 - 1/\mathcal{N})\Psi$ rather than $u|\Psi|^2\Psi$ owing to the Weyl ordering of the Hamiltonian. The additional term, however, should not be retained when making a consistent expansion in $1/\mathcal{N}$.
- [90] F. Piazza and P. Strack, Quantum kinetics of ultracold fermions coupled to an optical resonator, *Phys. Rev. A* **90**, 043823 (2014).
- [91] J. S. Langer, Theory of the condensation point, *Ann. Phys.* **41**, 108 (1967).
- [92] G. E. Uhlenbeck and L. S. Ornstein, On the theory of the Brownian motion, *Phys. Rev.* **36**, 823 (1930).
- [93] P. C. Hohenberg and B. I. Halperin, Theory of dynamic critical phenomena, *Rev. Mod. Phys.* **49**, 435 (1977).
- [94] Andrew J. Daley, Quantum trajectories and open many-body quantum systems, *Adv. Phys.* **63**, 77 (2014).