

Optimal sampling of dynamical large deviations via matrix product states

Luke Causer,^{1,2} Mari Carmen Bañuls,^{3,4} and Juan P. Garrahan^{1,2,5}

¹*School of Physics and Astronomy, University of Nottingham, Nottingham, NG7 2RD, UK*

²*Centre for the Mathematics and Theoretical Physics of Quantum Non-Equilibrium Systems, University of Nottingham, Nottingham, NG7 2RD, UK*

³*Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Str. 1, D-85748 Garching, Germany*

⁴*Munich Center for Quantum Science and Technology (MCQST), Schellingstr. 4, D-80799 München*

⁵*All Souls College, Oxford, UK*

The large deviation (LD) statistics of dynamical observables is encoded in the spectral properties of deformed Markov generators. Recent works have shown that tensor network methods are well suited to compute the relevant leading eigenvalues and eigenvectors accurately. However, the efficient generation of the corresponding rare trajectories is a harder task. Here we show how to exploit the MPS approximation of the dominant eigenvector to implement an efficient sampling scheme which closely resembles the optimal (so-called “Doob”) dynamics that realises the rare events. We demonstrate our approach on three well-studied lattice models, the Fredrickson-Andersen and East kinetically constrained models (KCMs), and the symmetric simple exclusion process (SSEP). We discuss how to generalise our approach to higher dimensions.

I. INTRODUCTION

The complex behaviour of the non-equilibrium dynamics of stochastic systems can be characterised by studying trajectory ensembles, that is, the set of all possible trajectories alongside the probability that they occur under the evolution defined via a stochastic master operator. This is analogous to standard thermodynamics, where static properties are entirely determined by the equilibrium ensemble of all microstates and the probabilities [1]. Often dynamical behaviour of interest is dominated not by trajectories that are typical under the dynamics, but by “rare events” which are exponentially (in time and in system size) scarce. Studying these rare events is made possible by using the framework of large deviations (LDs) [2–7], where in large time limits time-extensive dynamical observables obey a LD principle, and their statistics is encoded in functions which play for dynamics the role that thermodynamic potentials play for statics (see below for definitions).

LD functions can be obtained in principle from a deformation or *tilting* of the dynamical generator (in the case of continuous-time dynamics) or the Markov matrix (in the case of discrete-time dynamics), through its largest eigenvalue. Obtaining this eigenvalue is not always an easy - or even possible - task, and often one needs to resort to numerical methods. Methods to overcome this difficulty often include techniques based on population dynamics, namely cloning or splitting [8–11], and importance sampling [12–16] which provide information about the configurations frequently visited by the rare events. Notice that even if one manages to diagonalise the tilted generator (or the Markov matrix), the generation of rare trajectories is non-trivial: while rare trajectories are “generated” by the tilted operator, this is not a proper stochastic operator and these trajectories cannot be directly sampled.

The efficient sampling of rare events can be achieved by searching for another stochastic dynamics which gen-

erates trajectories with desirable probabilities that are the same as (or a close approximation to) those of the tilted generator (with any small discrepancy corrected via importance sampling techniques). Methods for doing so currently include optimal control [17, 18] and machine learning approaches, where one attempts to “learn” this convenient sampling dynamics [19–21]. The most optimal choice for a reference dynamics is the so-called *generalised Doob dynamics* [22–25], which generates trajectories with the exact tilting corresponding to the deformed generator. The Doob dynamics thus produces rare trajectories of the original dynamics “on demand”. To construct such optimal dynamics, however, requires knowledge of the leading eigenvector of the tilted generator.

Variational tensor network (TN) techniques [26–32], originally devised as a tool to study quantum many-body systems, have recently been shown to be useful in the context of classical stochastic models, specifically for studying LD statistics [33–36]. In particular, it is often both possible and easy to approximate the leading eigenstate of the tilted generator of a one-dimensional stochastic lattice system using a matrix product state (MPS) ansatz, even those with dynamical (i.e. LD) phase transitions. Recent works have made use of this eigenstate to determine the statistical properties of the dynamics [33–36]. To our knowledge, however, such TN approach has not been exploited yet to sample efficiently rare trajectories. This is what we do in this paper. We present a scheme to use the MPS approximation to the leading eigenvalue of the tilted generator to construct a new dynamics which very closely resembles the optimal Doob dynamics, and we show how we can use this new dynamics to efficiently sample rare events.

We focus on three paradigmatic models. The first two correspond to kinetically constrained models (KCMs) [6, 37, 38], specifically the Fredrickson-Andersen (FA) and the East [39] model, two well-studied models known for their connection to structural glasses [40, 41]. The

third model is the symmetric simple exclusion processes (SSEP) [42, 43]. All these models have interesting LD statistics, including trajectory phase transitions controlled by their activities and/or currents (in the case of the SSEP) [3, 5, 44–47].

The paper is organised as follows. In Sec. II, we review continuous time Markov dynamics and LDs. We also recap how one can apply an MPS ansatz to study KCMs. In Sec. III, we define the Doob dynamics and introduce a scheme to approximate it with a reference dynamics, constructed using an MPS approximation to the leading eigenstate of the tilted generator. In Sec. IV we present the numerical results from our method applied to the three models. We show how our approach can effectively be used to accurately measure the statistics of time-extensive observables. We provide an outlook on possible generalisations and our conclusions in Sec. V.

II. LARGE DEVIATIONS AND MATRIX PRODUCT STATES

In this section we introduce continuous-time Markov dynamics, giving specific examples in the context of kinetically constrained models (KCMs) and exclusion processes. We then also review the framework of large deviations (LDs) and how variational matrix product states (MPS) can be used to determine the LD statistics.

A. Continuous time Markov dynamics for KCMs and exclusion processes

We consider stochastic Markov dynamics which evolves continuously in time. Suppose we have some system with the set of configurations $\{x_1, x_2, \dots, x_M\}$ where M is the size of the configuration space. The probability that the system is in some configuration x at the time t is encoded in the probability vector $|P(t)\rangle = \sum_x P_x(t) |x\rangle$ which evolves under the stochastic master equation

$$\frac{d}{dt} |P(t)\rangle = \mathbb{W} |P(t)\rangle. \quad (1)$$

Here the generator of the dynamics \mathbb{W} is given by

$$\mathbb{W} = \sum_{x, x' \neq x} w_{x \rightarrow x'} |x'\rangle \langle x| - \sum_x R_x |x\rangle \langle x|, \quad (2)$$

where $w_{x \rightarrow x'}$ are the transition rates from configuration x to x' and $R_x = \sum_{x' \neq x} w_{x \rightarrow x'}$ is the escape rate from x . The largest eigenvalue of the generator is zero, with left eigenvector the flat state $\langle -| = \sum_x \langle x|$, and right eigenvector the steady state $|\text{ss}\rangle = \sum_x P_x^{\text{ss}} |x\rangle$, which describes the probability of finding any configuration at equilibrium. If our system obeys detailed balance, then we are guaranteed that any initial state will eventually relax to some equilibrium state given enough time. Here we assume this to be the case.

We will focus on two broad areas of 1D constrained systems. The first is KCMs (for reviews see [6, 37, 38]), for which configuration changes are governed by a kinetic constraint which is explicitly encoded in the generator. For concreteness, we focus on the 1D spin facilitation Fredrickson-Andersen (FA) [48] and East [39] models. Both models are defined on a 1D lattice of N binary variables (spins) $n_j = 0, 1$ for $j = 1, \dots, N$, and configuration changes are only allowed via single-spin flips. The Markovian generators for both models are given by

$$\mathbb{W}^{\text{East/FA}} = \sum_{i=1}^N \mathbb{P}_i^{\text{East/FA}} [c\sigma_i^+ + (1-c)\sigma_i^- - c(1-n_i) - (1-c)n_i] \quad (3)$$

where σ_i^\pm are the Pauli raising/lowering operators acting on site i and $c \in (0, 0.5]$ controls the rates at which spins flip, given they satisfy the kinetic constraints

$$\mathbb{P}_i^{\text{FA}} = n_{i-1} + n_{i+1}, \quad \mathbb{P}_i^{\text{East}} = n_{i-1}, \quad (4)$$

where the first only allows a transition if the spin attempting to flip has a neighbouring excitation, and the second only if the neighbouring spin to the left is excited. (For the FA model the constraint is sometimes defined as the projector $n_{i-1} + n_{i+1} - n_{i-1}n_{i+1}$, but in practice it makes little difference with the definition above.)

The second area we consider are exclusion processes [42, 43] - particles hopping around sites on a lattice, with a hardcore exclusion such that we can have at most one particle per site. We focus on the 1D symmetric simple exclusion process (SSEP), adopting the lattice notation we used for KCMs, where now $n_j = 1(0)$ implies the site is occupied (empty). In the SSEP, a particle can hop left or right to its neighbouring sites, both with the same rate ($\gamma = 1/2$) if the neighbouring site is not already occupied. The generator for the dynamics is

$$\mathbb{W}^{\text{SSEP}} = \frac{1}{4} \sum_{i=1}^N \sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y + \sigma_i^z \sigma_{i+1}^z - 1 \quad (5)$$

where σ_i^a are the Pauli operators acting on site i .

For the entirety of this paper, we will assume open boundary conditions (OBC), which will later reduce the computational cost of tensor network contractions. This formally means that we set $n_0 = n_{N+1} = 0$. Furthermore, we impose certain restrictions on the state space. For the FA model, we simply exclude the disconnected zero state $n_i = 0, \forall i$. On the other-hand, we set $n_1 = 1$ for the East model which ensures the state space remains fully connected on each dynamical site $i > 1$. Finally, we restrict SSEP such that the total number of particles $N_p = \sum_i n_i$ is fixed, with particle density $n_p = N_p/N$ which will be assumed to be $n_p = 1/2$.

B. Trajectories and large deviations

Consider some general trajectory $\omega_t = \{x_0 \rightarrow x_{t_1} \rightarrow \dots \rightarrow x_{t_K}\}$ where the system moves into the configura-

tion x_{t_i} at time t_i and has the total time $t > t_K$. The dynamical activity \hat{K} [3–6, 49] is a trajectory observable which measures the number of configuration changes for a given trajectory. The probability of observing some activity K can then be calculated as the sum over all trajectories with K configuration changes, and the probability they occur,

$$P_t(K) = \sum_{\omega_t} \pi(\omega_t) \delta[\hat{K}(\omega_t) - K], \quad (6)$$

where $\pi(\omega_t)$ is the probability of observing ω_t . For large times, this obeys the large deviation (LD) principle [2–5]

$$P_t(K) \sim e^{t\varphi(K/t)}, \quad (7)$$

where $\varphi(K/t)$ is called the LD rate function and plays the role of entropy density for trajectories. Alternatively, one can consider the moment generating function (MGF) [2]

$$Z_t(s) = \sum_K P_t(K) e^{-sK} = \sum_{\omega_t} \pi(\omega_t) e^{-s\hat{K}(\omega_t)}, \quad (8)$$

which contains equivalent information to Eq. (7) and can be considered the partition function. From Eq. (8), we see that the weighting of each trajectory is the probability that the trajectory occurs, exponentially re-weighted by its dynamical activity. The MGF also obeys a LD principle,

$$Z_t(s) \sim e^{t\theta(s)}, \quad (9)$$

where $\theta(s)$ is the scaled cumulant generating function (SCGF), whose derivatives evaluated at $s = 0$ give the cumulants of K scaled by time. The SCGF plays the role of the thermodynamical free energy of trajectories and is related to the LD rate function by a Legendre transform $\theta(s) = -\min_k (sk + \varphi(k))$ [2].

The MGF Eq. (8) can be expressed as

$$Z_t(s) = \langle -|e^{t\mathbb{W}_s}|\text{in}\rangle, \quad (10)$$

where $|\text{in}\rangle$ is some initial probability vector and \mathbb{W}_s is a new operator which we name the *tilted generator*, and is a deformed version of Eq. (2) where we tilt with respect to the dynamical observable of interest [2–5]. For the case of the dynamical activity [3–5], we simply tilt the off-diagonals of \mathbb{W} with the same factor to obtain

$$\mathbb{W}_s = \sum_{x, x' \neq x} e^{-s w_{x \rightarrow x'}} |x'\rangle \langle x| - \sum_x R_x |x\rangle \langle x|. \quad (11)$$

The largest eigenvalue of \mathbb{W}_s is the SCGF $\theta(s)$, with associated left and right eigenvectors $\langle l_s|$ and $|r_s\rangle$. Since $\langle l_s|$ in general is not the flat state, \mathbb{W}_s is not a proper stochastic generator for $s \neq 0$ [3–5]. If one could exactly diagonalise Eq. (11) to find its leading eigenvalue and eigenvectors, then they would entirely unravel the LD statistics. We now briefly recap how this can be achieved using numerical TN techniques [33–36].

C. Variational matrix product states

A matrix product state (MPS) is an ansatz for describing vector states of many-body systems [26–28, 50, 51],

$$|\Psi\rangle = \sum_{i_1, \dots, i_N} \text{Tr}(A_1^{i_1} A_2^{i_2} \dots A_N^{i_N}) |i_1 i_2 \dots i_N\rangle, \quad (12)$$

where each subsystem k has its own rank-3 tensor A_k with the dimensions $d \times D \times D$. The allowed entanglement within the state is controlled by the *bond dimension* D [30]. It is often convenient to represent tensor networks in a diagrammatic form using shapes to represent tensorial objects, and (connecting) lines to represent contractions over tensors. For example, the corresponding diagram for an MPS is

$$|\Psi\rangle = \begin{array}{c} | \\ \bullet \\ \text{---} \\ \bullet \\ \text{---} \\ \bullet \\ \text{---} \\ \bullet \\ | \end{array}, \quad (13)$$

where each circle corresponds to one of the tensors A_k . Similarly, one can also attempt to write some operator \hat{O} as a matrix product operator (MPO) [31, 32, 52–55]. Operators which act locally on the sub-systems, such as Eqs. (3–5), can be efficiently described as a MPO. That is to say we can represent them exactly in MPO form with only a small constant bond dimension. The diagrammatic representation for MPOs is

$$\hat{O} = \begin{array}{c} | \\ \blacksquare \\ \text{---} \\ \blacksquare \\ \text{---} \\ \blacksquare \\ \text{---} \\ \blacksquare \\ | \end{array}. \quad (14)$$

MPS allow for the easy and efficient implementation of the widely used density matrix renormalization group (DMRG) method [56, 57], an algorithm designed to iteratively minimize the energy of a state E_Ψ with respect to some Hamiltonian \hat{H} . In the language of MPS [28], we start with some guess at some fixed bond dimension, and sweep through each tensor applying local optimizations with all other tensors fixed. This is done until we reach convergence, which is usually when the change in energy of the state per sweep is small. At the end of the routine, one can efficiently calculate the variance of the state with respect to the Hamiltonian

$$\delta E_\Psi^2 = \text{var}_{\hat{H}}(\Psi) = \langle \hat{H}^2 \rangle_\Psi - \langle \hat{H} \rangle_\Psi^2 \quad (15)$$

where $\langle \cdot \rangle_\Psi = \langle \Psi | \cdot | \Psi \rangle$ denotes an expectation value. We check to see if it has fallen below some desired value, ϵ ; if not, we run the algorithm with an increased bond dimension, where we typically use the state from the previous run as an initial guess. For more details on the workings of variational MPS (vMPS) algorithms, see the reviews [28, 58].

Many recent works have shown that vMPS algorithms are very effective for studying the LD statistics of classically constrained systems which obey detailed balance

[33, 34, 36]. In particular, if we write the titled generator in a way such that it is Hermitian then the state we are searching for is the ground state. This guarantees each update is an improvement upon the last. For dynamics obeying detailed balance, the activity-tilted generator can be brought to a Hermitian form using a similarity transformation that is independent of s [5],

$$\mathbb{H}_s = -Q^{-1}\mathbb{W}_sQ, \quad (16)$$

For the case of the East/FA models [5], the diagonal operator Q is given by

$$Q^{\text{FA/East}} = [\sqrt{1-c}|0\rangle\langle 0| + \sqrt{c}|1\rangle\langle 1|]^{\otimes N}, \quad (17)$$

and for the SSEP by $Q^{\text{SSEP}} = \mathbb{I}$. The new Hamiltonian \mathbb{H}_s has the ground state $|\psi_s\rangle$ with energy $-\theta(s)$. The ground state is related to the left and right eigenvectors of \mathbb{W}_s in the following way [33],

$$|\psi_s\rangle = Q^{-1}|r_s\rangle, \quad (18)$$

$$\langle\psi_s| = \langle l_s|Q, \quad (19)$$

$$|\psi_s\rangle = \sum_x \sqrt{l_s(x)r_s(x)}|x\rangle, \quad (20)$$

where $l_s(x) = \langle l_s|x\rangle$ and $r_s(x) = \langle x|r_s\rangle$.

III. DOOB TRANSFORMATION AND OPTIMAL SAMPLING

We now define the so-called generalised Doob transformation [19, 22, 23, 25, 59], and show how one can use our MPS solution to Eq. (16) to construct a reference dynamics which closely resembles the true Doob dynamics. We then present a method to optimally sample the rare events of our toy models using these new dynamics.

A. Generalised Doob dynamics

The goal is to find a proper stochastic generator which generates trajectories with the same probabilities as those in the tilted dynamics \mathbb{W}_s , cf. Eq. (11). This can be achieved using the (long-time) generalised Doob transformation [19, 22, 23, 25, 59], defined as

$$\mathbb{W}_s^{\text{Doob}} = \mathbb{L}[\mathbb{W}_s - \theta(s)\mathbb{I}]\mathbb{L}^{-1}, \quad (21)$$

where $\mathbb{L} = \text{diag}(\langle l_s|)$ is the left eigenvector $\langle l_s|$ as a diagonal matrix. It is easy to check that Eq. (21) is annihilated by the flat state $\langle -|$, which means that $\mathbb{W}_s^{\text{Doob}}$ is a stochastic operator. Its stationary state is

$$|\text{ss}\rangle_s^{\text{Doob}} = \sum_x l_s(x)r_s(x)|x\rangle. \quad (22)$$

The generator $\mathbb{W}_s^{\text{Doob}}$ can also be expressed as a sum of its diagonal and off-diagonal elements

$$\begin{aligned} \mathbb{W}_s^{\text{Doob}} = & \sum_{x,x'\neq x} \frac{l_s(x')}{l_s(x)} e^{-s w_{x\rightarrow x'}} |x'\rangle\langle x| \\ & - \sum_x (R_x + \theta(s)) |x\rangle\langle x|. \end{aligned} \quad (23)$$

Thus our new dynamics has the transition rates and escape rates

$$\tilde{w}_{x\rightarrow x'} = \frac{l_s(x')}{l_s(x)} e^{-s w_{x\rightarrow x'}} \quad (24)$$

$$\tilde{R}_x = R_x + \theta(s) \quad (25)$$

respectively. That is to say the transition rates are reweighted by e^{-s} and by some ratio $l_s(x')/l_s(x)$ which depends on the structure of the configurations, and the escape rate is shifted by $\theta(s)$.

We now consider some general time-dependent observable \hat{A} , and ask what is the expectation value in the tilted dynamics,

$$\langle\hat{A}\rangle_s \equiv \frac{\langle\hat{A}e^{-sK}\rangle}{\langle e^{-sK}\rangle} = Z_t(s)^{-1} \sum_{\omega_t} \pi(\omega_t) \hat{A}(\omega_t) e^{-s\hat{K}(\omega_t)}. \quad (26)$$

One can now apply importance sampling to arrive at

$$\begin{aligned} \langle\hat{A}\rangle_s &= Z_t(s)^{-1} \sum_{\omega_t} \tilde{\pi}(\omega_t) \frac{\pi(\omega_t)}{\tilde{\pi}(\omega_t)} \hat{A}(\omega_t) e^{-s\hat{K}(\omega_t)}, \\ &= Z_t(s)^{-1} \left\langle \frac{\pi}{\tilde{\pi}} \hat{A} e^{-s\hat{K}} \right\rangle_{\text{Doob}} \end{aligned} \quad (27)$$

where $\tilde{\pi}(\omega_t)$ is the probability of observing ω_t in the dynamics generated by $\mathbb{W}_s^{\text{Doob}}$ and $\langle\cdot\rangle_{\text{Doob}}$ denotes an expectation value with respect to trajectories with probabilities from the Doob dynamics. At a first glance, it might look that we have not gained much from expressing the expectation of A using the Doob generator $\mathbb{W}_s^{\text{Doob}}$. However, if one calculates the ratio of probabilities in Eq. (27) then the power of this expression becomes apparent.

Let us first consider the original dynamics described by Eq. (2). If we have some system in configuration x , then the probability it flips to some other state x' at the time Δt is

$$P_{x\rightarrow x'}(\Delta t) = w_{x\rightarrow x'} e^{-R_x \Delta t}. \quad (28)$$

It then follows that the trajectory ω_t occurs with probability

$$\pi(\omega_t) = P(x_0) e^{-R_{x_K}(t-t_{x_K})} \prod_{i=1}^K w_{x_{i-1}\rightarrow x_i} e^{-R_{x_{i-1}}(t_{x_i}-t_{x_{i-1}})}, \quad (29)$$

where we have also accounted for the fact that the system must remain in the same state after the final flip for the remainder of the time, and $P(x_0)$ is the probability of the initial configuration. The probability of the trajectory under the Doob dynamics has a similar form, with the substitutions $w_{x \rightarrow x'} \rightarrow \tilde{w}_{x \rightarrow x'}$, $R_x \rightarrow \tilde{R}_x$, and $P(x_0) \rightarrow \tilde{P}(x_0)$,

$$\tilde{\pi}(\omega_t) = \tilde{P}(x_0) e^{-sK} e^{-t\theta(s)} \frac{l_s(x_K)}{l_s(x_0)} e^{-R_{x_K}(t-t_K)} \quad (30)$$

$$\times \prod_{i=1}^K w_{x_{i-1} \rightarrow x_i} e^{-R_{x_{i-1}}(t_{x_i} - t_{x_{i-1}})},$$

where all but the endpoint factors of $l_s(x)$ cancel out telescopically. The ratio of probabilities then goes as

$$\frac{\pi(\omega_t)}{\tilde{\pi}(\omega_t)} = \frac{e^{sK} e^{t\theta(s)}}{l_s(x_0) l_s(x_K)}, \quad (31)$$

where we have used $\tilde{P}(x_0) = P(x_0) l_s(x_0)^2$. Substituting Eq. (31) back into Eq. (27) cancels out the exponential tilting e^{sK} . Furthermore, for large times, $Z_t(s)^{-1} \approx e^{-t\theta(s)}$ giving the final result

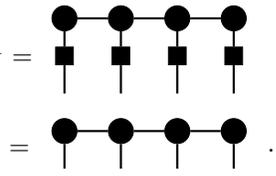
$$\langle \hat{A} \rangle_s = \left\langle \frac{1}{l_s(x_0) l_s(x_K)} \hat{A} \right\rangle_{\text{Doob}}. \quad (32)$$

And so it follows that one can exactly sample the expectation value of a trajectory observable in the tilted ensemble defined by the non-stochastic tilted generator, by sampling it directly from trajectories generated by the stochastic Doob dynamics Eq. (21), up to factors at the endpoints of each trajectory (which become negligible in the long time limit).

B. Reference Dynamics

While the above shows how to optimally sample if one has access to the Doob generator, which is obtained from the exact minimisation of the tilted generator, we now consider how to approximate it efficiently.

Suppose we have an MPS approximation $|\psi_s^{\text{ref}}\rangle$ to the ground state of the Hermitian operator \mathbb{H}_s , where our choice of bond dimension D controls the error. By applying the operator Q^{-1} to $|\psi_s^{\text{ref}}\rangle$, as is done in Eq. (19), one can also retrieve an approximation to the left eigenvector. This is easily done as an MPS-MPO product,

$$\begin{aligned} \langle l_s^{\text{ref}} | &= \langle \psi_s^{\text{ref}} | Q^{-1} = \\ &= \end{aligned}$$


$$(33)$$

We then construct the generator of the so-called *reference dynamics*, which goes as Eq. (2) with the transition rates

and escape rates given by

$$w_{x \rightarrow x'}^{\text{ref}} = \frac{l_s^{\text{ref}}(x')}{l_s^{\text{ref}}(x)} e^{-s} w_{x \rightarrow x'}, \quad (34)$$

$$R_x^{\text{ref}} = \sum_{x' \neq x} w_{x \rightarrow x'}^{\text{ref}}, \quad (35)$$

respectively. Note that here we have not used Eq. (25) for the escape rates, as these reference dynamics only act as an approximation to the Doob dynamics, and thus would not give a true stochastic dynamics. In appendix A, we show the steady-state solution to the reference dynamics is given by

$$|\text{ss}\rangle_s^{\text{ref}} = \sum_x \psi_s^{\text{ref}}(x)^2 |x\rangle, \quad (36)$$

where $\psi_s^{\text{ref}}(x) = \langle x | \psi_s^{\text{ref}} \rangle$.

If we repeat the steps between Eqs. (27) and (32) but for the reference dynamics, the expectation Eq. (27) looks like

$$\langle \hat{A} \rangle_s = \left\langle \frac{1}{l_s^{\text{ref}}(x_0) l_s^{\text{ref}}(x_K)} e^{-t\theta(s) + \int dt \Delta \hat{R}} \hat{A} \right\rangle_{\text{ref}}, \quad (37)$$

where $\int dt \Delta \hat{R}$ is the time integral of the difference of escape rates between the reference dynamics and the original dynamics, with $\Delta \hat{R}_x = R_x^{\text{ref}} - R_x$. We can estimate a sampling error when using Eq. (37) in the following way [59]. First, let us assume the effects of the time-edge factors is negligible (as they are not exponential in time) and try to sample the quantity

$$\langle e^{-s\hat{K}} \rangle = \langle \mathcal{R} e^{-s\hat{K}} \rangle_{\text{ref}} \approx \frac{1}{N_{\text{sp}}} \sum_{\alpha=1}^{N_{\text{sp}}} \mathcal{R}(\omega^\alpha) e^{-s\hat{K}(\omega^\alpha)}, \quad (38)$$

where $\mathcal{R}(\omega^\alpha) = e^{s\hat{K}(\omega^\alpha) + \int dt \Delta \hat{R}(\omega^\alpha)}$ is the umbrella which compensates for change in sampling dynamics and we estimate for a fixed number of samples, N_{sp} . The variance of Eq. (38) gives a way to quantify the sampling error,

$$\begin{aligned} \epsilon_{\text{ref}}^2 &= \frac{\text{Var}_{\text{ref}} \left(\frac{1}{N_{\text{sp}}} \sum_{\alpha=1}^{N_{\text{sp}}} \mathcal{R}(\omega^\alpha) e^{-s\hat{K}(\omega^\alpha)} \right)}{\left\langle \frac{1}{N_{\text{sp}}} \sum_{\alpha=1}^{N_{\text{sp}}} \mathcal{R}(\omega^\alpha) e^{-s\hat{K}(\omega^\alpha)} \right\rangle_{\text{ref}}^2} \\ &= \frac{1}{N_{\text{sp}}} \left[\frac{\langle \mathcal{R}^2 e^{-2s\hat{K}} \rangle_{\text{ref}}}{\langle \mathcal{R} e^{-s\hat{K}} \rangle_{\text{ref}}^2} - 1 \right] \\ &= \frac{1}{N_{\text{sp}}} \left[\frac{\langle e^{2 \int dt' \Delta \hat{R}} \rangle_{\text{ref}}}{\langle e^{\int dt' \Delta \hat{R}} \rangle_{\text{ref}}^2} - 1 \right]. \end{aligned} \quad (39)$$

In appendix B we show

$$\epsilon_{\text{ref}}^2 \approx \frac{e^{t^2 \delta E^2} - 1}{N_{\text{sp}}} \approx \frac{t^2 \delta E^2}{N_{\text{sp}}} \quad (40)$$

for small $t^2 \delta E^2 \ll 1$, where δE^2 is the calculated variance on our MPS approximation of the leading eigenvector.

C. Simulating trajectories

We are now in a position to efficiently simulate trajectories from our reference dynamics. The sampling of trajectories from a classical generator is usually achieved using a continuous time Monte Carlo (CTMC, otherwise known as the BKL algorithm) [60]. Given that our system is in some configuration x at time t' , we need to calculate the next jump in the trajectory. That is, we need to decide the next configuration the system will move into, and the time it does so. Calculating this can be split into five separate steps:

1. Find each configuration x' the system can move into from x .
2. Calculate the transition rates $w_{x \rightarrow x'}$ for each x' .
3. Calculate the escape rate R_x as the sum of all transition rates.
4. Randomly choose one x' , each with the probability $w_{x \rightarrow x'} / R_x$.
5. Randomly choose the jump time Δt with probability $P(\Delta t) = R_x e^{-R_x \Delta t}$.

By starting at a configuration sampled from equilibrium (which in the case of the reference dynamics can be efficiently done using the MPS $|\psi_s^{\text{ref}}\rangle$ [61, 62]), or otherwise, one can simply repeat this procedure until some total time t has elapsed.

We can use this method for our reference dynamics, where the only step that needs slight adjustment is the second. While one must still calculate the transition rates of the original dynamics in the usual way, we must also calculate the left vector components $l_s(x)$ and $l_s(x')$. Let us assume the former is carried over from the previous jump in the algorithm. Then all one needs to do is calculate each $l_s(x')$. We start by noting that any configuration x can be written in MPS form with bond dimension 1,

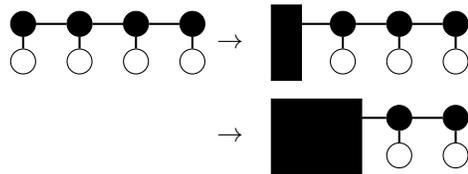
$$\begin{aligned} |x\rangle &= |n_1\rangle \otimes |n_2\rangle \otimes \cdots \otimes |n_N\rangle \\ &= \begin{array}{cccc} \text{---} & \text{---} & \text{---} & \text{---} \\ | & | & | & | \\ \bigcirc & \bigcirc & \bigcirc & \bigcirc \end{array} \end{aligned} \quad (41)$$

and then we can simply calculate the left component as a MPS-MPS contraction

$$l_s^{\text{ref}}(x) = |\langle l_s^{\text{ref}} | x \rangle| = \begin{array}{cccc} \bullet & \bullet & \bullet & \bullet \\ \text{---} & \text{---} & \text{---} & \text{---} \\ | & | & | & | \\ \bigcirc & \bigcirc & \bigcirc & \bigcirc \end{array}. \quad (42)$$

The transition rates for the reference dynamics are then calculated using Eq. (34), and the method proceeds as before. The total computational cost for calculating each $l_s^{\text{ref}}(x)$ is $O(D^2 N)$, and thus the total cost of each Monte Carlo (MC) step is $O(D^2 N N_F)$, where N_F is the total number of configurations x' for a given step.

Let us now consider our KCMs where we have single-spin flip dynamics. We first note that the number of possible configuration changes from x is bounded by the number of sites, that is, $1 \leq N_F \leq N$. Using the method described above, the computational cost for each step is at worst, quadratic in the system size. However, by realising that the tensor network contractions $\langle l_s^{\text{ref}} | x \rangle$ and $\langle l_s^{\text{ref}} | x' \rangle$ are identical apart from just one tensor (corresponding to the spin which would flip), we can reduce the computational cost by recycling partial contractions from the edges. We first need to identify the first and last sites on the lattice which are able to flip, which we label i_l and i_r respectively. In a similar fashion to variational algorithms, we then contract from the left edge of tensor network $\langle l_s^{\text{ref}} | x \rangle$ up to $i_r - 1$, and saving each tensor block along the way.



We do the same but from the right and up to $i_l + 1$. This initialization of partial contractions has a one-time cost of

$$O(D^2(N + i_r - i_l - 2)) < O(2D^2 N). \quad (43)$$

Calculating each $l_s^{\text{ref}}(x')$ at site j is then easy, we just perform a contraction over the left block up to site $j - 1$, the right block up to site $j + 1$ and the remaining tensors,

$$l_s^{\text{ref}}(x') = \begin{array}{ccc} \text{---} & \bullet & \text{---} \\ | & | & | \\ \bigcirc & \bigcirc & \bigcirc \end{array}$$

where of course we now choose the remaining tensor at site j from the MPS $|x'\rangle$. This is done for each possible site which can flip, and thus entails a computational cost $O(D^2 N_F)$. Once a choice is made for which site to flip, which we will label i , we must update the blocks of partial contractions up to (the now possibly different) i_l and i_r . Note that this time we do not have to start from the edges of the MPS, but just from site i as the previous partial contractions that come before do not change. The total cost of updating the partial contractions is

$$O(D^2[(i_r - i) + (i - i_l)]) = O(D^2(i_r - i_l)). \quad (44)$$

The total computation cost for each MC step is the sum of the cost for calculating each $l_s^{\text{ref}}(x)$ and updating the partial blocks after a choice is made,

$$O(D^2(N_F + i_r - i_l)) \leq O(2D^2 N). \quad (45)$$

Consequently, the cost of each MC step is reduced to one which is at most linear in system size.

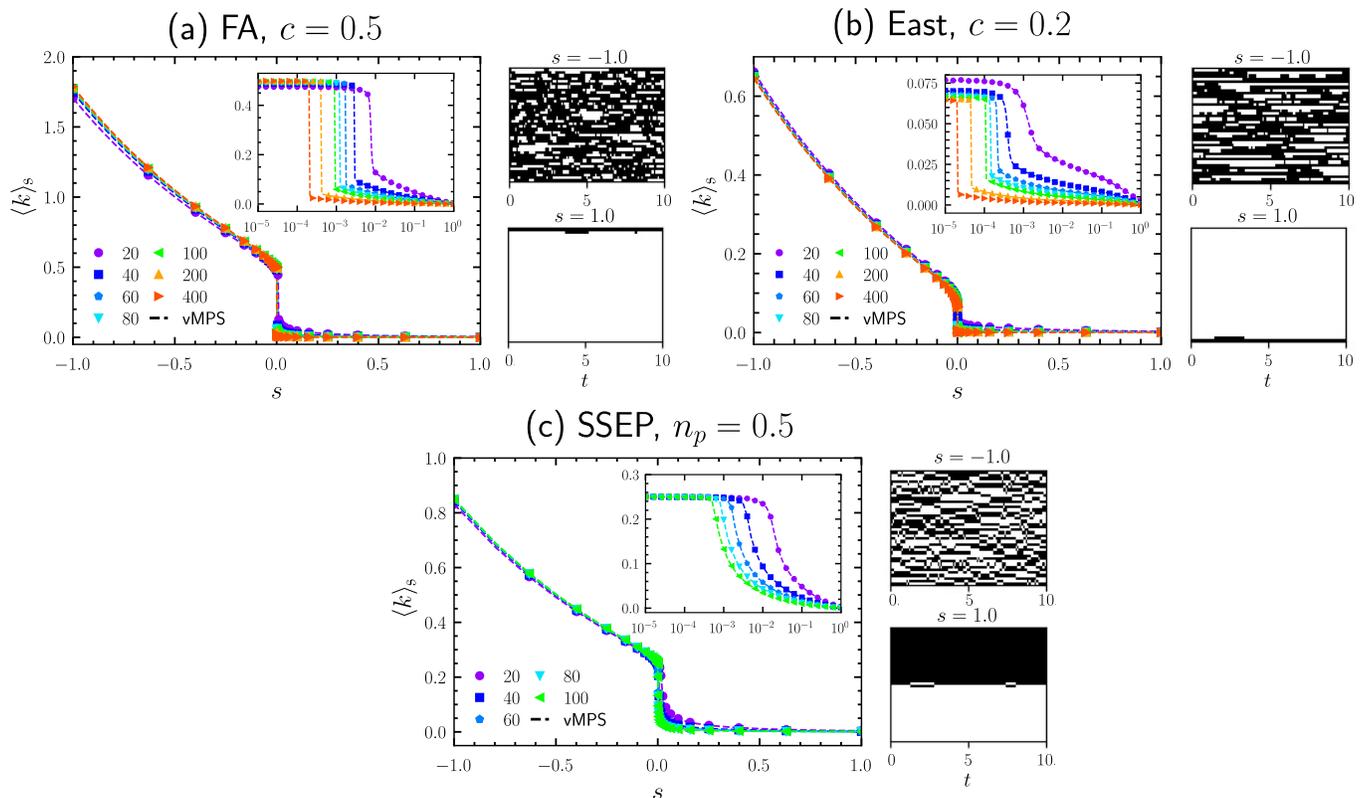


FIG. 1. **The dynamical activity from brute-force Monte Carlo.** We show the dynamical activity measured for (a) the FA model with $c = 0.5$, (b) the East model with $c = 0.2$ and (c) SSEP with $n_p = 0.5$. A variety of system sizes $N \in [20, 400]$ are shown for each. The dashed lines shown expected activity calculated directly from the MPS $\tilde{k}(s)$, whereas the markers show the activity measured via CTMC $\langle k \rangle_s$ with a time $t = 100/\tilde{k}(s)$. The inactive phase is shown with a log s scale in the insets. We also show representative trajectories at $s = -1, 1$ for each.

IV. NUMERICAL RESULTS

A. Approximating the Doob Dynamics

We put to the test the general method presented above by approximating the Doob dynamics of each model defined in Sec. II. **We show that the Doob dynamics is well estimated using the MPS reference dynamics**, and can even be well approximated with truncated MPS.

Each of the three models is known to exhibit a trajectory phase transition (when tilted against the activity) for long times and in the thermodynamic limit $N \rightarrow \infty$, manifested in the SCGF $\theta(s)$ at $s = 0$ with a discontinuous drop in the dynamical activity $\hat{K}(s) = -\theta'(s)/N$ [3, 5, 36, 45, 47]. We call the dynamical phase for $s < 0$ the *active phase*, and that for $s > 0$, the *inactive phase*. One is able to do a detailed investigation of this first-order phase transition by considering the finite-size scaling of the model [33, 36, 46, 63, 64]. We can estimate a critical point $s_c(N) \gtrsim 0$ by finding the peak of the dynamical susceptibility $\chi(s) = \theta''(s)$, which shows a drastic change in a small region around the transition point.

We start by taking the usual approach of approxim-

ing the ground states $|\Psi_s\rangle$ using vMPS. That is, we run the algorithm allowing the bond dimension to increase until the variance of the energy (with respect to the Hamiltonian) falls sufficiently, cf. Eq. (15). The resultant MPS is then used to construct the reference dynamics, which approximates the Doob dynamics to a high accuracy, as explained in the previous section. Note that because the vMPS tries to keep entanglement as low as possible, for $s > s_c(N)$ the approximated ground state exhibits localisation at just one edge of the system [33]. While for the East case this corresponds to the structure of the ground state in the sector with fixed occupation 1 in the leftmost site, the FA and SSEP models have reflection symmetry, spontaneously broken for $s > 0$ and large N . Thus, in order to maintain the symmetry in the latter two cases, we construct an MPS which is a superposition of the result from vMPS and its spatially reflected state to construct our dynamics in the inactive phase.

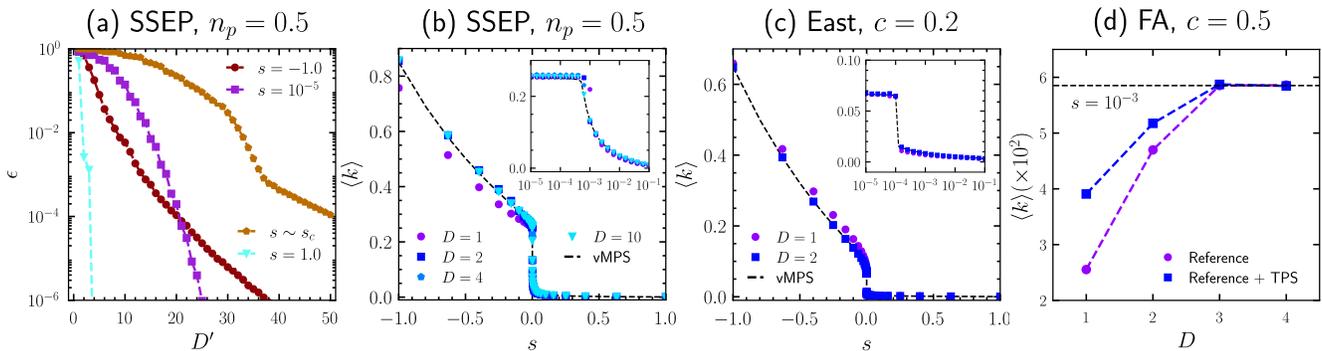


FIG. 2. **Reference dynamics from a truncated MPS.** All data is for the system size $N = 100$. (a) The truncation error $\epsilon = 1 - \langle \psi_D | \psi_{D'} \rangle^2$ as a function of truncated bond dimension D' for SSEP at various values of s . (b) The measured average dynamical activity with a reference dynamics constructed with truncated MPS. The dashed line shows the expected value obtained through vMPS with a large ($D \geq 50$) bond dimension, and the inset shows the same but on a log-scale around the critical point. (c) The same but for the East model. (d) The measured dynamical activity $\langle k \rangle$ as a function of D for $s = 10^{-3}$ close to the critical point s_c . The purple circles show the values measured using the reference dynamics alone, whereas the blue squares show values obtained using the reference dynamics and TPS to incorporate umbrella sampling. The dashed line shows the expected value obtained from vMPS. Each point is done for a trajectory time of $t = 100$ and $N_{sp} = 10^6$ trajectories.

1. Direct sampling with the reference dynamics but without re-weighting

We first check that the CTMC algorithm with our MPS reference dynamics gives the expected results. We do this *without using* the trajectory re-weighting, cf. Eq. (37). This amounts to only considering infinite-time dynamics, and assuming that our approximation is actually exact. Despite this strong assumption, we find that it produces excellent results as shown in Fig. 1. The expected dynamical activity (per unit site and time, dashed lines) can be calculated as a TN contraction over our MPS and MPO,

$$\tilde{k}(s) = \frac{1}{N} \left\langle \Psi_s \left| \frac{d\mathbb{H}_s}{ds} \right| \Psi_s \right\rangle. \quad (46)$$

The same quantity can be calculated on a trajectory level (symbols) by counting the total number of configuration changes, $\langle K \rangle$ and taking its time (and spatial) average,

$$\langle k \rangle_s = \frac{\langle K \rangle}{Nt}, \quad (47)$$

where t is the run time for each trajectory. We show results for each model, for a range of system sizes of $N \in [20, 400]$. The expected and measured results have excellent agreement. This simplified algorithm struggles most around the transition point, $s_c(N)$, due to the required large bond dimension (see Refs. [33, 36]).

We also show representative trajectories for the active ($s = -1$) and inactive phases ($s = 1$). Each model excellently contrasts the difference in dynamics between the two phases. The active phase displays very rapid changes with structures that allow for unconstrained dynamics. For the FA and East models this means having a large number of excitations, while SSEP requires particles to be spaced apart. Conversely, this inactive phase has just

few configuration changes with highly constrained dynamics. This means minimizing the number of excitations for the FA and East resulting in the dynamics responsible for the so-called “space-time bubble” in local regions of space [5, 40, 65], while for SSEP we restrict the activity by clustering the particles [47, 66]. To our knowledge, direct dynamical sampling of trajectories for these systems sizes and values of $s \neq 0$ is unprecedented for these three models.

2. Reference dynamics with truncated bond dimensions

While in the extreme active/inactive limits we can achieve a good MPS description with just a bond dimension of $O(10)$, one may need a bond dimension of $O(100)$ for the more difficult regions such as around $s = 0$ [33, 36]. One reason for the necessity of this high bond dimension could be that the state has longer-ranged spatial correlations. Another could be that when one runs the vMPS, we run it against some constraint in the state space. For the FA model, this is the weak constraint that restricts to the connected component of all configuration but the one with $n_i = 0$ for all i . For the SSEP, we have the stronger constraint that we are within the state space with fixed N_p particles.

The goal is to look for a state with a smaller bond dimension than we currently have which still contains all the necessary interactions, but if necessary, discards the information which enforces the constraint. Then, by starting our CTMC algorithm in a state which satisfies the constraint, we will automatically enforce it for the rest of the trajectory, as the dynamics keeps the system in the constrained subspace.

Bond dimensions of TNs are typically reduced via truncation. That is, we move through each bond within the

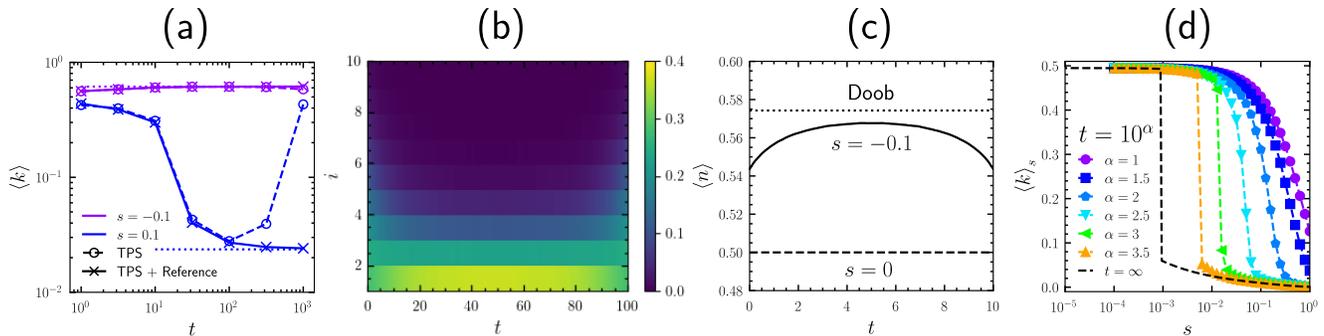


FIG. 3. **Sampling finite time trajectories.** All results are done for the FA model with $c = 0.5$. (a) The measured dynamical activity $\langle k \rangle$ as a function of time t for $s = -0.1$ (top) and $s = 0.1$ (bottom) for $N_{sp} = 10^6$ trajectories and system size $N = 40$. The circles show the values obtained via TPS with the normal dynamics, and crosses TPS with the reference dynamics. The dotted lines show the expected value at infinite times. (b) The local occupations $\langle n_i \rangle$ as a function of time in the inactive regime $s = 0.1$ and $N = 40$. (c) The average excitation density $\langle n \rangle$ as a function of time in the active regime, $s = -0.1$. The value approaches the expected value in the Doob dynamics (dotted line) in the bulk, but moves towards the equilibrium value ($s = 0$, dashed line) at the edges and $N = 40$. (d) The dynamical activity as a function of s and time t . The data for $t = \infty$ is obtained directly from the MPS, whereas finite t is obtained using TPS. Note that the sharp drop in activity shifts with time. The dynamics are run at the system size $N = 100$.

TN and apply a singular value decomposition over it, keeping only the D' states with the largest singular values, where $D' < D$ is our new desired bond dimension. We show this in Fig. 2(a) for SSEP (as this typically requires the largest bond dimension), where we run the vMPS to at least (but higher if required) $D = 50$ to find the state $|\psi_D\rangle$, and then truncate to $|\psi_{D'}\rangle$ with the bond dimension $D' < D$. We measure the truncation error $\varepsilon = 1 - |\langle \psi_D | \psi_{D'} \rangle|^2$ between the two states, where we assume both are normalised. We find that when far from the critical point, we can describe the original state to a high accuracy with bond dimensions as small as $D' \sim 20$. Conversely, we cannot attain the same level of accuracy for $s \sim s_c$, where the state exhibits larger amounts of entanglement.

There are multiple reasons that one may want to find a state with a truncated bond dimension. The first is that our Monte Carlo algorithm scales quadratically with the bond dimension - this could hinder the convergence of time-dependent observables at large times, which can require a large sample size to be determined with sufficient accuracy. For such situations, reducing the scaling of the algorithm would be desired. Another reasoning could be that we want to investigate a system which requires a higher complexity of TN, such as 2D system with projected entangled pair states (PEPS) [58, 67]. Not only would the scaling of our CTMC algorithm increase, but so would the scaling of the variational algorithm used to find the reference dynamics. In this case, one may not be able to reach a bond dimension large enough to give a desirable variance.

We show the measured dynamical activity for SSEP and the East model (symbols), with a reference dynamics constructed from states with a truncated bond dimension in Fig. 2(b, c), and compare to the expected

result from the non-truncated MPS (dashed line). Surprisingly, we find that for the most part, even for bond dimensions as small as $D = 2$, we can accurately reproduce the correct dynamical activity for each of the models. As expected, the truncation struggles mostly around the transition point. Nevertheless, we can achieve good results for the FA (not shown) and East with a truncated bond dimension of $D = 4$, and $D = 10$ for SSEP.

The calculations done thus far have been with a reference dynamics constructed using a truncated bond dimension without any trajectory re-weighting. In principle, Eq. (37) is exact and thus allows for further improvements by using the umbrella

$$g(\omega) = e^{-t\theta(s) + \int dt \Delta \hat{R}(\omega)}. \quad (48)$$

We implement this re-weighting via transition path sampling (TPS) with the *shifting* method, see Refs. [12, 59] for further details. Figure 2(d) shows the results of this umbrella sampling for the FA model with a s value close to the critical point, s_c . It is here the discrepancy is the largest, and we can do a more detailed analysis by looking at a larger range of bond dimensions. As was the case for (a), we see a significant improvement when using the re-weighting factor Eq. (48). It might be that we could see further improvements with more TPS iterations.

The main point to take from this is that we are able to achieve accurate results for the dynamical activity (the observable we are tilting) and some local observables with a relatively small bond dimension. This of course comes at a cost however, as when we truncate we discard some of the information that accounts for the long-ranged spatial correlations. For the case of SSEP, even though it seems as if we are discarding a large amount of information when truncating (cf. Fig. 2(a)), it seems that we

keep the relevant information needed to reproduce the correct dynamics, but at the cost of not maintaining the conservation law. We note however that it is possible to explicitly implement the conservation laws in the MPS [68], but it is not clear how this will affect the quality of the reference dynamics in the CTMC algorithm.

B. Sampling rare events of finite times

For the previous results, we disregard finite-time effects by considering our sampled trajectories to be a “slice” of a larger infinite-time trajectory. We now look to incorporate these effects back into our sampling by considering the full re-weighting factor

$$g(\omega) = \frac{e^{-t\theta(s) + \int dt \Delta \hat{R}(\omega)}}{l_s^{\text{ref}}(x_0) l_s^{\text{ref}}(x_K)}. \quad (49)$$

Note that previously, for a large bond dimension, the part of Eq. (49) which accounts for the difference in escape rate had a negligible effect, and could be ignored. This is not always the case here, as the umbrella sampling at the time-edges of the trajectory causes the system to visit configurations which are atypical in the Doob dynamics, and not well described by our MPS approximation.

As a proof of principle, we start by comparing results from TPS with the original dynamics against TPS with the reference dynamics for a small system size $N = 40$, and a variety of times, as is shown for the FA model at $s = \pm 0.1$ in Fig. 3(a). For small times, both show excellent agreement. For large times however, the normal dynamics struggles to correctly account for the expected activity shown by the dotted lines, a result of the exponential time-dependence in Eq. (26) (as \hat{K} is time extensive). While sampling with our reference dynamics reduces the exponential cost in time, the time-edges still suffer from an exponential sampling difficulty in the system size. This is most noticeable for the inactive phase, where each model exhibits an exponential localization [33, 36, 69] at the spatial edge(s) of the system. This causes the $l_s^{\text{ref}}(x)$ values to exponentially vary. Nevertheless, it is still a significant improvement on the previously exponential cost in space, time and s . [We offer a solution to ease this sampling difficulty in Appendix C, where we construct a reference dynamics using our approximation to the Doob dynamics, in addition to a small perturbation of the stationary state of the generator \$\mathbb{W}\$.](#)

The average occupations $\langle n_i \rangle$ (at site i) for $s = 0.1$ and $t = 100$ is shown in Fig. 3(b), while Fig. 3(c) shows the average excitation density $\langle n \rangle = N^{-1} \sum_i \langle n_i \rangle$ for $s = -0.1$ and $t = 10$. It is here the time-edge effects become obvious; we start at a state which lies somewhere between the expected $s = 0$ (dashed line) dynamics and the expected long-time dynamics, which depends on the whole spectrum of \mathbb{W}_s , as well as the total trajectory time. The system quickly evolves and resembles the Doob dynamics. Note that at the end of trajectory, it is

again described by the original probability vector, as is expected due to the time-symmetry in Eq. (37).

Finally, Fig. 3(d) shows the average dynamical activity as a function of s and time, t . [We show the expected activity in the infinite time limit \$t = \infty\$ as a black dashed line, and the measured activity for finite times as symbols.](#) Notice that as time decreases, the drop in activity becomes less sharp. Furthermore, the transition from the active to inactive phase happens at decreasing s . While the methods presented here could allow for a detailed investigation into the temporal scaling of the critical point, doing so for desirable system sizes would be at a large computational cost. We hope to investigate this more extensively using time-evolution methods (see e.g. [28, 58, 70]).

V. CONCLUSIONS

We have expanded on previous applications of TNs to classical constrained models [33–36], using the MPS solution to the ground state of a tilted stochastic generator in 1D to construct a reference dynamics which well approximates the exact Doob dynamics. This allows us to (nearly) optimally sample the rare events of 1D constrained systems with just a polynomial cost in both space and time - rather than the exponential cost of most sampling methods. We have demonstrated here the efficiency of this approach by generating tilted trajectory ensembles for the FA and East KCMs and the symmetric simple exclusion process. Our simulations are for sizes and times unprecedented for such large deviation studies.

Furthermore, our results show that it is possible to obtain an accurate dynamics away from the dynamical transitions of the models we studied with a truncated bond dimension, which enables close to optimal sampling simulations at little cost. Further extensions of our work includes generalising our methods to higher dimensions, for example by using two-dimensional variational tensor network techniques, such as PEPS [58, 67] to approximate the leading eigenvectors of 2D classical generators, as is done in Ref. [35]. From the associated leading eigenvectors, as we have shown here, we can in turn to construct a reference dynamics which is nearly optimal for sampling rare trajectories. While PEPS algorithms do not currently allow for bond dimensions comparable to vMPS, they remain a fruitful area of research which is constantly being improved on [71–79]. [Recent works \[80\] have shown the effectiveness of using recurrent neural networks \(RNN\) to approximate the leading eigenstates of tilted generators in two dimensions. The methods presented here could be easily generalized to RNN to allow for the efficient sampling of 2D rare events.](#)

Another area that deserves exploration is to apply similar TN methods to systems which do not obey detailed balance, and for which their generators cannot be brought to a Hermitian form. While this would dampen

the effectiveness of variational algorithms, approaches based on time evolution may offer a promising solution (see e.g.[28, 58, 70]). Such approaches could also offer further insights into intermediate time rare events, where both usual sampling methods and large deviation methods fall short. We hope to report on such studies in the near future.

ACKNOWLEDGMENTS

We acknowledge financial support from EPSRC Grant no. EP/R04421X/1 and the Leverhulme Trust Grant

No. RPG-2018-181. M.C.B. acknowledges support from Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC-2111 – 390814868. J.P.G. is grateful to All Souls College, Oxford, for support through a Visiting Fellowship. We acknowledge access to the University of Nottingham Augusta HPC service. Much of the numerical data during the later stages of this work was acquired using the ITensor library [81].

-
- [1] D. Chandler, “Introduction to Modern Statistical Mechanics,” (1987).
- [2] H. Touchette, *Phys. Rep.* **478**, 1 (2009).
- [3] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, *Phys. Rev. Lett.* **98**, 195702 (2007).
- [4] V. Lecomte, C. Appert-Rolland, and F. van Wijland, *J. Stat. Phys.* **127**, 51 (2007).
- [5] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, *J. Phys. A* **42**, 75007 (2009).
- [6] J. P. Garrahan, *Physica A* **504**, 130 (2018).
- [7] R. L. Jack, *Eur. Phys. J. B* **93**, 74 (2020).
- [8] C. Giardinà, J. Kurchan, and L. Peliti, *Phys. Rev. Lett.* **96**, 120603 (2006).
- [9] V. Lecomte and J. Tailleur, *J. Stat. Mech.* **2007**, P03004 (2007).
- [10] F. Cérou, A. Guyader, and M. Rousset, *Chaos* **29**, 43108 (2019).
- [11] C. Pérez-Espigares and P. I. Hurtado, *Chaos* **29**, 83106 (2019).
- [12] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, *Annu. Rev. Phys. Chem.* **53**, 291 (2002).
- [13] U. Ray, G. K.-L. Chan, and D. T. Limmer, *J. Chem. Phys.* **148**, 124120 (2018).
- [14] U. Ray, G. K.-L. Chan, and D. T. Limmer, *Phys. Rev. Lett.* **120**, 210602 (2018).
- [15] K. Klymko, P. L. Geissler, J. P. Garrahan, and S. Whiteman, *Phys. Rev. E* **97**, 32123 (2018).
- [16] A. Guyader and H. Touchette, *J. Stat. Phys.* **181**, 551 (2020).
- [17] T. Nemoto, F. Bouchet, R. L. Jack, and V. Lecomte, *Phys. Rev. E* **93**, 62123 (2016).
- [18] G. Ferré and H. Touchette, *J. Stat. Phys.* **172**, 1525 (2018).
- [19] T. H. E. Oakes, A. Moss, and J. P. Garrahan, *Mach. Learn.: Sci. Technol.* **1**, 035004 (2020).
- [20] D. C. Rose, J. F. Mair, and J. P. Garrahan, *New J. Phys.* **23**, 013013 (2021).
- [21] E. Gillman, D. C. Rose, and J. P. Garrahan, *arXiv* (2020), [arXiv:2002.05185](https://arxiv.org/abs/2002.05185).
- [22] R. L. Jack and P. Sollich, *Prog. Theor. Phys. Supp.* **184**, 304 (2010).
- [23] R. Chetrite and H. Touchette, *Ann. Henri Poincaré* **16**, 2005 (2015).
- [24] J. P. Garrahan, *J. Stat. Mech.: Theory Exp* **2016**, 73208 (2016).
- [25] F. Carollo, J. P. Garrahan, I. Lesanovsky, and C. Pérez-Espigares, *Phys. Rev. A* **98**, 10103 (2018).
- [26] G. Vidal, *Phys. Rev. Lett.* **91**, 147902 (2003).
- [27] F. Verstraete, D. Porras, and J. I. Cirac, *Phys. Rev. Lett.* **93**, 227205 (2004).
- [28] U. Schollwöck, *Ann. Phys.* **326**, 96 (2011).
- [29] R. Orús, *Nat. Rev. Phys.* **1**, 538 (2019).
- [30] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, *Phys. Rev. Lett.* **100**, 30504 (2008).
- [31] I. P. McCulloch, *J. Stat. Mech.* **2007**, P10014 (2007).
- [32] B. Pirvu, V. Murg, J. I. Cirac, and F. Verstraete, *New J. Phys.* **12**, 25012 (2010).
- [33] M. C. Bañuls and J. P. Garrahan, *Phys. Rev. Lett.* **123**, 200601 (2019).
- [34] P. Helms, U. Ray, and G. K.-L. Chan, *Phys. Rev. E* **100**, 022101 (2019).
- [35] P. Helms and G. K.-L. Chan, *Phys. Rev. Lett.* **125**, 140601 (2020).
- [36] L. Causer, I. Lesanovsky, M. C. Bañuls, and J. P. Garrahan, *Phys. Rev. E* **102**, 052132 (2020).
- [37] F. Ritort and P. Sollich, *Adv. Phys.* **52**, 219 (2003).
- [38] J. P. Garrahan, P. Sollich, and C. Toninelli, in *Dynamical Heterogeneities in Glasses, Colloids, and Granular Media*, International Series of Monographs on Physics, edited by L. Berthier, G. Biroli, J.-P. Bouchaud, L. Cipelletti, and W. van Saarloos (Oxford University Press, Oxford, UK, 2011).
- [39] J. Jäckle and S. Eisinger, *Z. für Phys. B* **84**, 115 (1991).
- [40] D. Chandler and J. P. Garrahan, *Annu. Rev. Phys. Chem.* **61**, 191 (2010).
- [41] G. Biroli and J. P. Garrahan, *J. Chem. Phys.* **138**, 12A301 (2013).
- [42] K. Mallick, *Physica A* **418**, 17 (2015).
- [43] R. A. Blythe and M. R. Evans, *J. Phys. A* **40**, R333 (2007).
- [44] T. Bodineau and B. Derrida, *C. R. Acad. Sci.* **8**, 540 (2007).
- [45] C. Appert-Rolland, B. Derrida, V. Lecomte, and F. van Wijland, *Phys. Rev. E* **78**, 21122 (2008).
- [46] T. Bodineau, V. Lecomte, and C. Toninelli, *J. Stat. Phys.* **147**, 1 (2012).
- [47] R. L. Jack, I. R. Thompson, and P. Sollich, *Phys. Rev. Lett.* **114**, 60601 (2015).

- [48] G. H. Fredrickson and H. C. Andersen, *Phys. Rev. Lett.* **53**, 1244 (1984).
- [49] C. Maes, *Phys. Rep.* **850**, 1 (2020).
- [50] M. Fannes, B. Nachtergaele, and R. F. Werner, *Comm. Math. Phys.* **144**, 443 (1992).
- [51] S. Östlund Stellanstlund and S. Rommer, *Phys. Rev. Lett.* **75**, 3537 (1995).
- [52] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, *Phys. Rev. Lett.* **93**, 207204 (2004).
- [53] M. Zwolak and G. Vidal, *Phys. Rev. Lett.* **93**, 207205 (2004).
- [54] C. Hubig, I. P. McCulloch, and U. Schollwöck, *Phys. Rev. B* **95** (2017).
- [55] D. E. Parker, X. Cao, and M. P. Zaletel, *Phys. Rev. B* **102**, 35147 (2020).
- [56] S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [57] U. Schollwöck, *Rev. Mod. Phys.* **77**, 259 (2005).
- [58] F. Verstraete, V. Murg, and J. I. Cirac, *Adv. Phys.* **57**, 143 (2008).
- [59] T. Oakes, S. Powell, C. Castelnovo, A. Lamacraft, and J. P. Garrahan, *Phys. Rev. B* **98**, 64302 (2018).
- [60] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, *J. Comput. Phys.* **17**, 10 (1975).
- [61] A. J. Ferris and G. Vidal, *Phys. Rev. B* **85**, 165146 (2012).
- [62] S. Iblisdir, *New J. Phys.* **16**, 103022 (2014).
- [63] T. Bodineau and C. Toninelli, *Comm. Math. Phys.* **311**, 357 (2012).
- [64] T. Nemoto, R. L. Jack, and V. Lecomte, *Phys. Rev. Lett.* **118**, 115702 (2017).
- [65] J. P. Garrahan and D. Chandler, *Phys. Rev. Lett.* **89** (2002).
- [66] V. Lecomte, J. P. Garrahan, and F. van Wijland, *J. Phys. A* **45**, 175001 (2012).
- [67] H. N. Phien, J. A. Bengua, H. D. Tuan, P. Corboz, and R. Orús, *Phys. Rev. B* **92**, 35142 (2015).
- [68] C. Hubig, *SciPost Phys.* **5**, 47 (2018).
- [69] N. Pancotti, G. Giudice, J. I. Cirac, J. P. Garrahan, and M. C. Bañuls, *Phys. Rev. X* **10**, 21051 (2020).
- [70] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, *Ann. Phys.* **411**, 167998 (2019).
- [71] P. Corboz, *Phys. Rev. B* **93**, 45116 (2016).
- [72] P. Corboz, *Phys. Rev. B* **94**, 35133 (2016).
- [73] L. Vanderstraeten, J. Haegeman, P. Corboz, and F. Verstraete, *Phys. Rev. B* **94**, 155123 (2016).
- [74] P. Corboz, P. Czarnik, G. Kapteijns, and L. Tagliacozzo, *Phys. Rev. X* **8**, 31031 (2018).
- [75] M. Rader and A. M. Läuchli, *Phys. Rev. X* **8**, 31030 (2018).
- [76] L. Vanderstraeten, J. Haegeman, and F. Verstraete, *Phys. Rev. B* **99**, 165121 (2019).
- [77] P. Czarnik, J. Dziarmaga, and P. Corboz, *Phys. Rev. B* **99**, 35115 (2019).
- [78] R. Haghshenas, M. J. O’Rourke, and G. K.-L. Chan, *Phys. Rev. B* **100**, 54404 (2019).
- [79] M. J. O’Rourke and G. K.-L. Chan, *Phys. Rev. B* **101**, 205142 (2020).
- [80] C. Casert, T. Vieijra, S. Whitelam, and I. Tamblyn, “Dynamical large deviations of two-dimensional kinetically constrained models using a neural-network state ansatz,” (2020), arXiv:2011.08657.
- [81] M. Fishman, S. R. White, and E. M. Stoudenmire, “The ITensor Software Library for Tensor Network Calculations,” (2020), arXiv:2007.14822.

APPENDIX

A. Steady-state solution in the reference dynamics

The generator of the reference dynamics defined by Eqs. (34-35) can be written as

$$\mathbb{W}_s^{\text{ref}} = \sum_{x, x' \neq x} e^{-s} w_{x \rightarrow x'} \frac{l_s^{\text{ref}}(x')}{l_s^{\text{ref}}(x)} \left[|x'\rangle \langle x| - |x\rangle \langle x'| \right]. \quad (50)$$

By definition, the stationary state $|\text{ss}\rangle_s^{\text{ref}} = \sum_z P_s^{\text{ref}}(z) |z\rangle$ is annihilated by Eq. (50). It follows that

$$\begin{aligned} \mathbb{W}_s^{\text{ref}} |\text{ss}\rangle_s^{\text{ref}} &= \sum_{x, x' \neq x} e^{-s} w_{x \rightarrow x'} \frac{l_s^{\text{ref}}(x')}{l_s^{\text{ref}}(x)} P_s^{\text{ref}}(x) \left[|x'\rangle - |x\rangle \right] \\ &= \sum_{x, x' \neq x} e^{-s} \left[w_{x \rightarrow x'} \frac{l_s^{\text{ref}}(x')}{l_s^{\text{ref}}(x)} P_s^{\text{ref}}(x) \right. \\ &\quad \left. - w_{x' \rightarrow x} \frac{l_s^{\text{ref}}(x)}{l_s^{\text{ref}}(x')} P_s^{\text{ref}}(x') \right] |x'\rangle \\ &= 0, \end{aligned} \quad (51)$$

where we have used a change of variables in the second and third line. Let us assume our original dynamics obeys detailed balance, and that the state space is fully connected. Then it follows that if $w_{x \rightarrow x'} = 0$, so does $w_{x' \rightarrow x}$, in which Eq. (51) is satisfied. Otherwise, we must have that

$$\frac{P_s^{\text{ref}}(x)}{P_s^{\text{ref}}(x')} = \frac{w_{x' \rightarrow x} l_s^{\text{ref}}(x)^2}{w_{x \rightarrow x'} l_s^{\text{ref}}(x')^2}. \quad (52)$$

Now let us recall that our original dynamics obeys detailed balance. This allows us to use a similarity transformation to write the generator as a Hermitian one, cf. Eq. (16). In particular, let us define the diagonal transformation matrix as

$$Q = \sum_z Q(z) |z\rangle \langle z|. \quad (53)$$

One can easily show that for \mathbb{H} to be Hermitian, we must have

$$\frac{Q(x)^2}{Q(x')^2} = \frac{w_{x' \rightarrow x}}{w_{x \rightarrow x'}}. \quad (54)$$

Substituting this back into Eq. (52), we find

$$\frac{P_s^{\text{ref}}(x)}{P_s^{\text{ref}}(x')} = \frac{Q(x)^2 l_s^{\text{ref}}(x)^2}{Q(x')^2 l_s^{\text{ref}}(x')^2}. \quad (55)$$

and it follows the stationary state is given by

$$|\text{ss}\rangle_s^{\text{ref}} = \sum_x l_s^{\text{ref}}(x)^2 Q(x)^2 |x\rangle. \quad (56)$$

For the case of our MPS dynamics, we defined $\langle l_s| = \langle \psi_s| Q^{-1}$ of our solution $\langle \psi_s|$. It follows that Eq. (56) can be written as $|\text{ss}\rangle_s^{\text{ref}} = \sum_x \psi_s^{\text{ref}}(x)^2 |x\rangle$, where $\psi_s^{\text{ref}}(x) = \langle x| \psi_s \rangle$.

B. Sampling variance in the reference dynamics

We start by assuming that we are always at the stationary state of the dynamics. This allows us to calculate the trajectory ensemble average of some observable (per unit time) as the average over all configurations with respect to the stationary state,

$$\langle \hat{O} \rangle_{\text{ref}} \equiv \frac{1}{t} \left\langle \int_0^t dt' \hat{O}(t') \right\rangle_{\text{ref}} = \langle l_s^{\text{ref}} | \hat{O} | r_s^{\text{ref}} \rangle. \quad (57)$$

The aim is to calculate the expectation value and the variance of the time-integrated difference in escape rates (c.f. Eq. (39)). Using Eq. (57), we can write

$$\begin{aligned} \langle \Delta \hat{R} \rangle_{\text{ref}} &= \sum_{x, y \neq x} l_x \left(\frac{l_y}{l_x} e^{-s} \omega_{x \rightarrow y} - \omega_{x \rightarrow y} \right) r_x \\ &= \sum_{x, y \neq x} l_y e^{-s} \omega_{x \rightarrow y} r_x - l_x \omega_{x \rightarrow y} r_x \\ &= \langle l_s^{\text{ref}} | \mathbb{W}_s | r_s^{\text{ref}} \rangle = \theta^{\text{ref}}(s), \end{aligned} \quad (58)$$

where we have written $l_x \equiv l_s^{\text{ref}}(x)$ and $r_x \equiv r_s^{\text{ref}}(x)$ for brevity. Performing the same calculation for $\Delta \hat{R}^2$, we find

$$\langle \Delta \hat{R}^2 \rangle_{\text{ref}} = \langle l_s^{\text{ref}} | \mathbb{W}_s^2 | r_s^{\text{ref}} \rangle, \quad (59)$$

giving the variance

$$\text{Var}_{\text{ref}} \Delta \hat{R} \equiv \langle \Delta \hat{R}^2 \rangle_{\text{ref}} - \langle \Delta \hat{R} \rangle_{\text{ref}}^2 = \delta E^2, \quad (60)$$

where δE^2 is the measured variance of the MPS used to construct the reference dynamics with respect to the tilted generator (or tilted Hamiltonian).

We are now in a position to estimate the sampling error Eq. (39). The time integrated difference in escape rate can be written as it's average plus some random noise,

$$\int_0^t dt' \Delta \hat{R}(t') = t \langle \Delta \hat{R} \rangle + t \delta \hat{R}, \quad (61)$$

with $\delta \hat{R}$ having zero mean and δE^2 variance. For simplicity, we also assume $\delta \hat{R}$ to be normally distributed, allowing us to calculate Eq. (39),

$$\epsilon_{\text{ref}}^2 = \frac{1}{N_{\text{sp}}} \left[\frac{\langle e^{2t \delta \hat{R}} \rangle_{\text{ref}}}{\langle e^{t \delta \hat{R}} \rangle_{\text{ref}}^2} - 1 \right] \approx \frac{e^{t^2 \delta E^2} - 1}{N_{\text{sp}}}. \quad (62)$$

[LC: We need some justification to why this is a fine assumption.]

C. Reference dynamics in the inactive phase

The exponential cost of the time-edge dependence in Eq. (49) leads to a slow initial convergence in TPS. We

aim to ease this cost by constructing a reference dynamics which depends on the spectrum of the titled generator \mathbb{W}_s . As the biggest sampling difficulty is close to (but after) the critical point $s_c(N)$, we will construct a reference dynamics using just the approximations of two eigenstates. The first is of course the leading eigenvector we find using vMPS, $|\psi_s^{\text{ref}}\rangle$. The second is the stationary state of the equilibrium dynamics, $|\psi_{s=0}\rangle$, which closely approximates an eigenstate of \mathbb{W}_s for small s . Together, they give the complete reference dynamics

$$|\tilde{\psi}_s^{\text{ref}}\rangle = |\psi_s^{\text{ref}}\rangle + \kappa |\psi_{s=0}\rangle, \quad (63)$$

up to a renormalisation factor. To make an appropriate choice for κ , one can consider the spectrum of \mathbb{W}_s for both states. We first calculate their expectation values

$$\theta^{\text{ref}}(s) = -\langle \psi_s^{\text{ref}} | \mathbb{H}_s | \psi_s^{\text{ref}} \rangle, \quad (64)$$

$$\theta^{\text{ss}}(s) = -\langle \psi_{s=0} | \mathbb{H}_s | \psi_{s=0} \rangle. \quad (65)$$

Note that the dynamics also evolves exponentially in time under the generator $e^{t\mathbb{W}_s}$. We use this to justify our choice

$$\kappa = e^{-t[\theta^{\text{ref}}(s) - \theta^0(s)]}. \quad (66)$$

We stress that this should only be considered a perturbation to the original choice of the dynamics, where the perturbation serves as a connection to the original dynamics with the purpose of aiding the convergence close to the critical point. Its exponential time dependence ensures that we retrieve the Doob approximation for large times. As we move further from the critical point, the gap $\theta^{\text{ref}}(s) - \theta^{\text{ss}}(s)$ grows and thus the the perturbation is again exponentially suppressed.