# Measurement-induced phase transitions in quantum inference problems

With Curt von Keyserlingk (KCL), Austen Lamacraft (Cambridge)

**arXiv:2504.08888**

**Sun Woo P. Kim (KCL)**

**25-09-12, Christ's College, Cambridge**

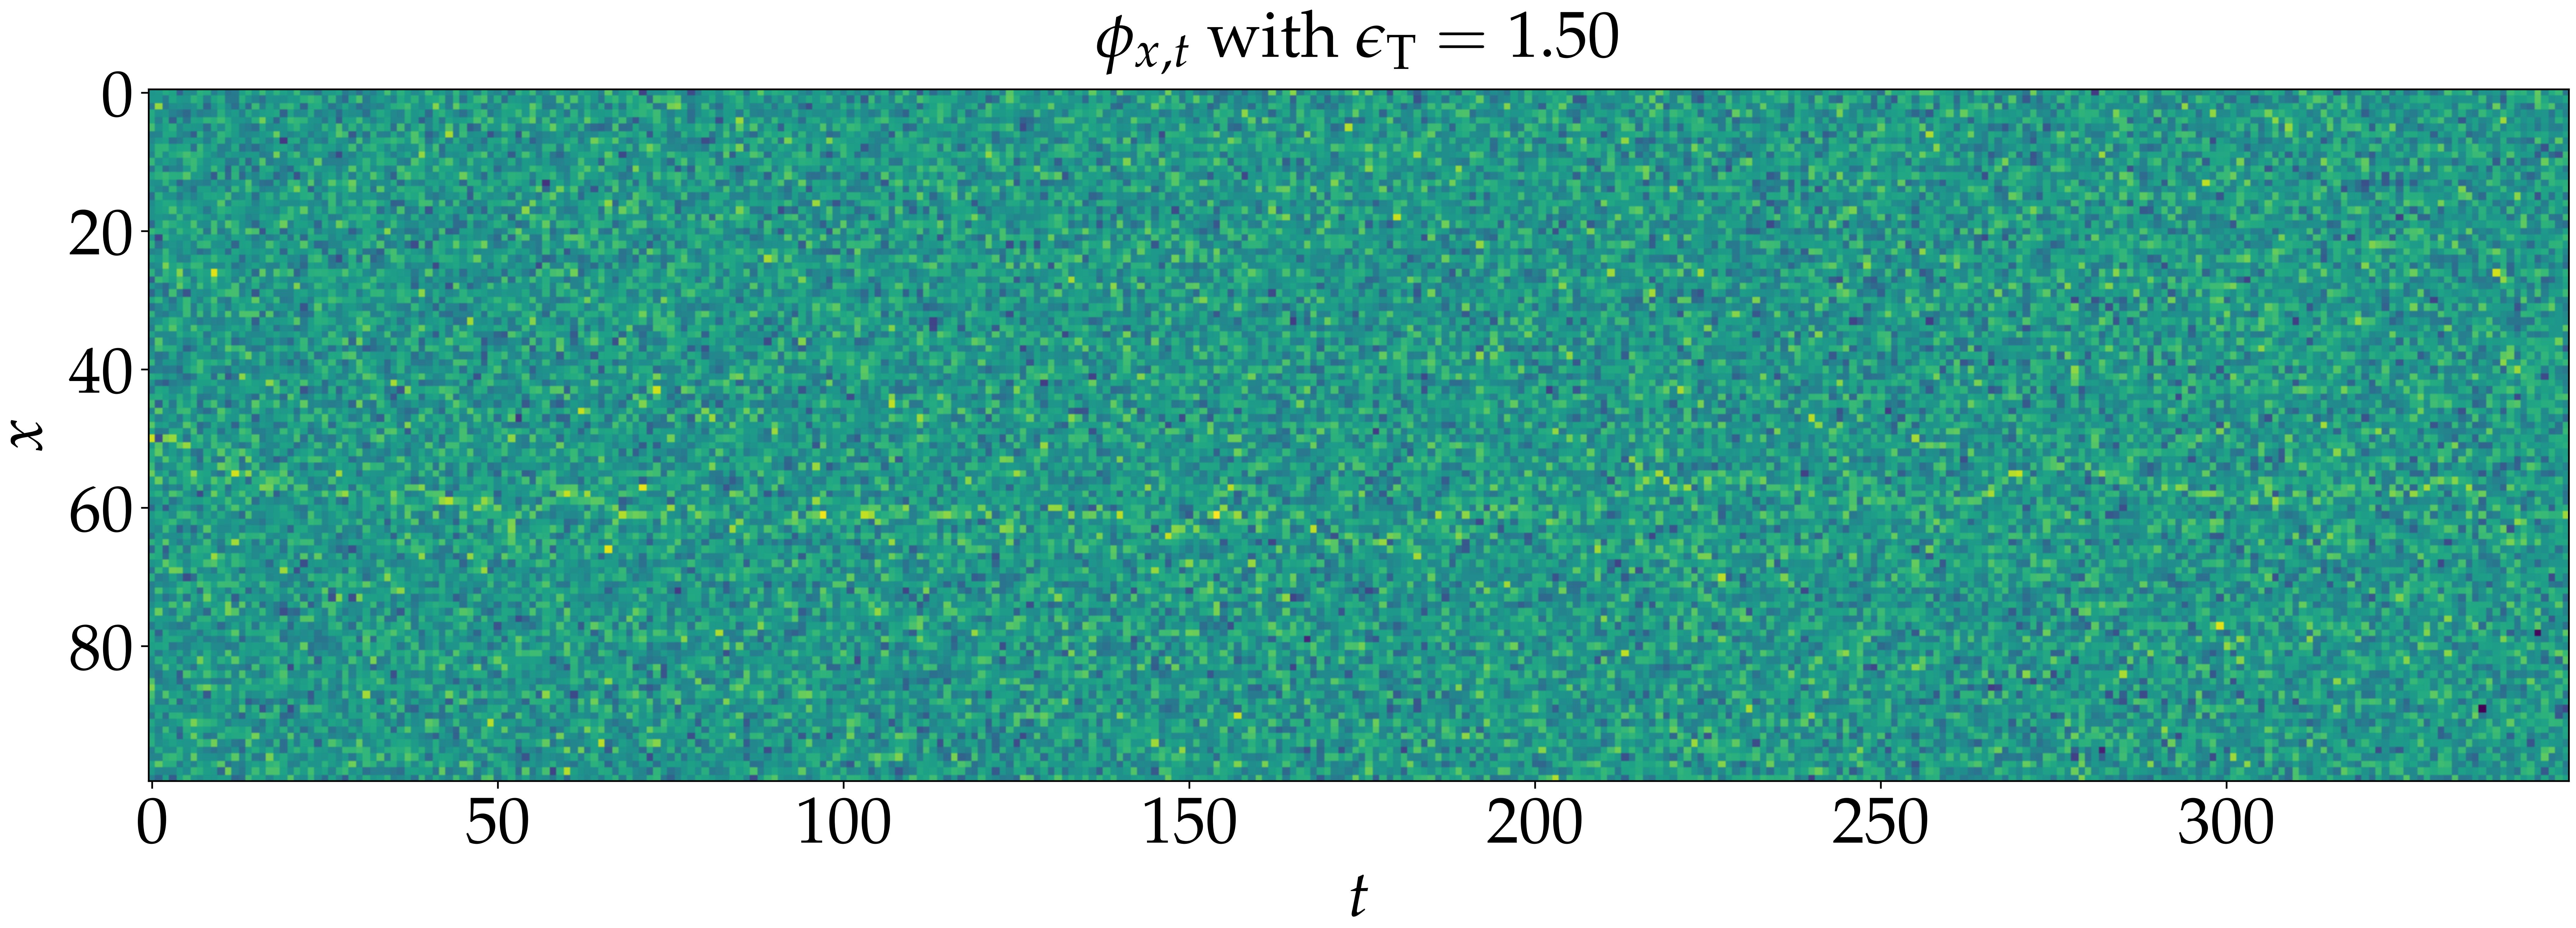
$\phi_{x,t}$ with $\epsilon_{\mathrm{T}} = 1.50$
0
20
40
60
80
$x$
0
50
100
150
200
250
300
$t$

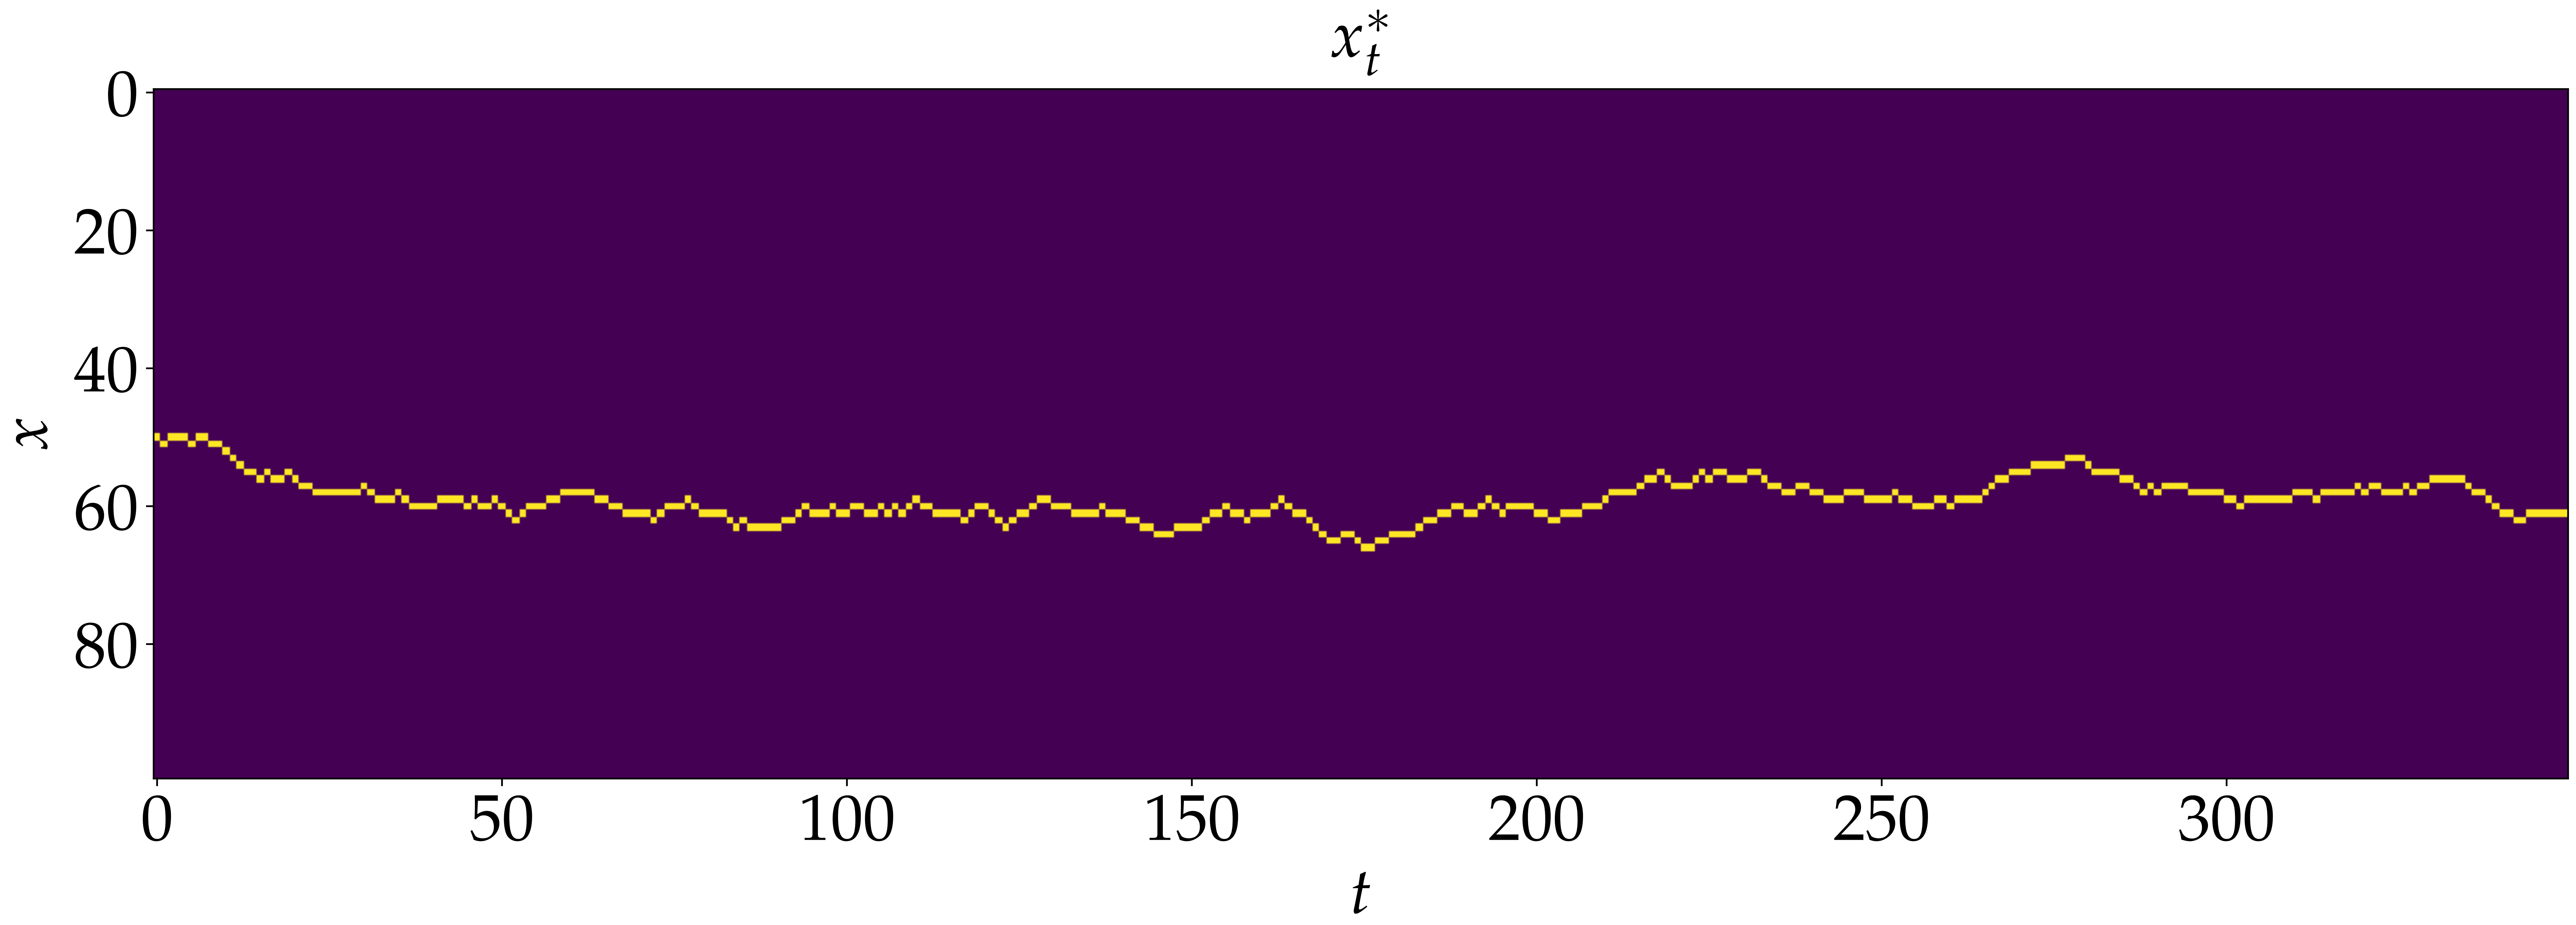

$x_t^*$
0
20
40
60
80
$x$
0
50
100
150
200
250
300
$t$

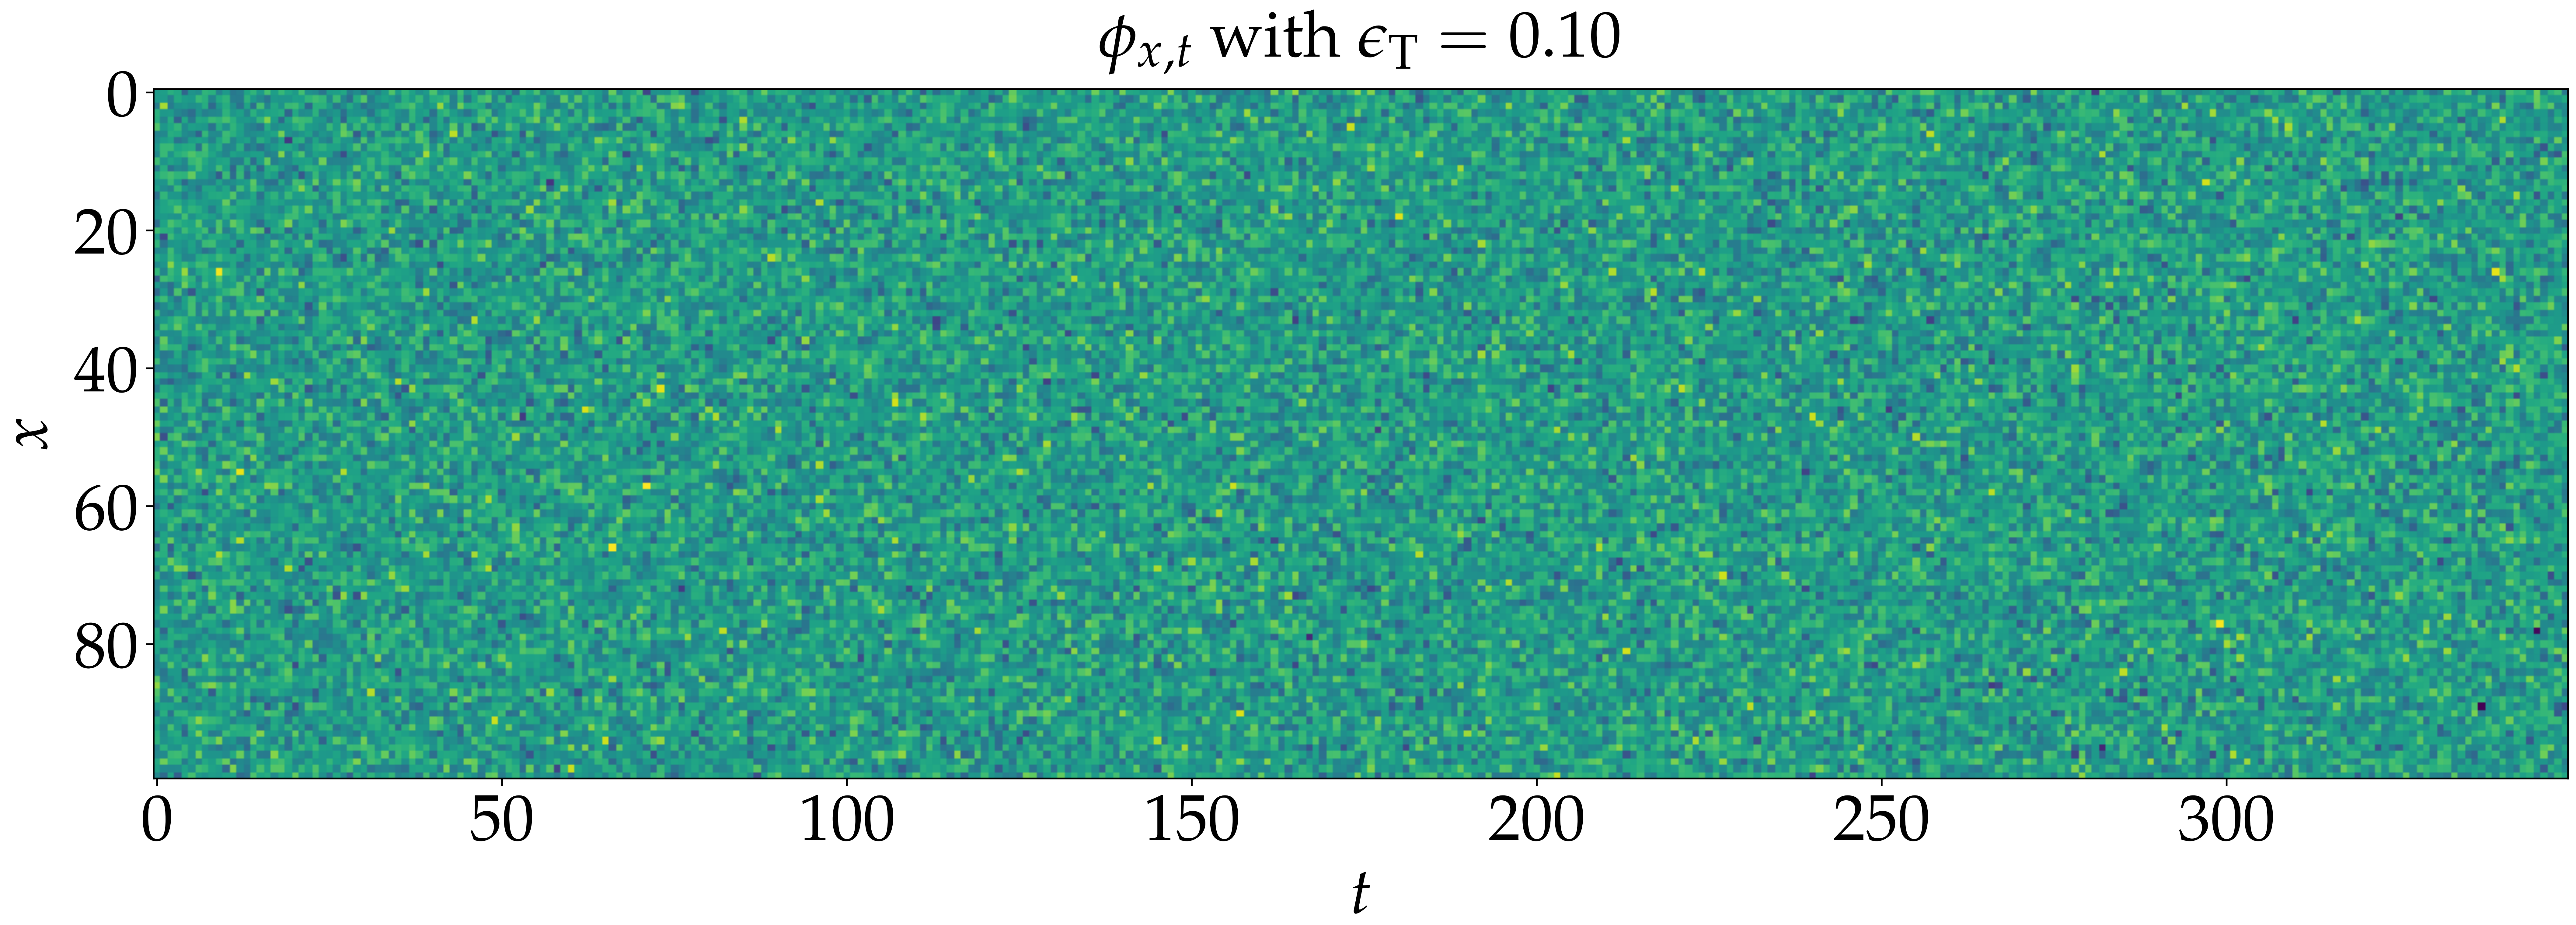
$\phi_{x,t}$ with $\epsilon_T = 0.10$
0
20
40
60
80
$x$
0
50
100
150
200
250
300
$t$

# Outline

1. **Bayesian inference and hidden Markov models**

2. Teacher-student scenario and Bayes optimality

3. General quantum inference problem and quantum hidden Markov models

4. Quantum error correction and the random-bond Ising model

5. The Planted SSEP and the planted XOR

6. Discussion of the phase diagrams

# Bayesian inference

Infer posterior distribution for state $X$ given data $Y$

$$p(X|Y) = \frac{p(Y|X)p(X)}{p(Y)}$$

$p(Y|X)$: "likelihood"/measurement model

$p(X)$: "prior"

$p(Y) = \sum_X p(Y|X)p(X)$: "evidence"/normalisation

In general $X$ is high-dimensional

# Hidden Markov models (HMMs)

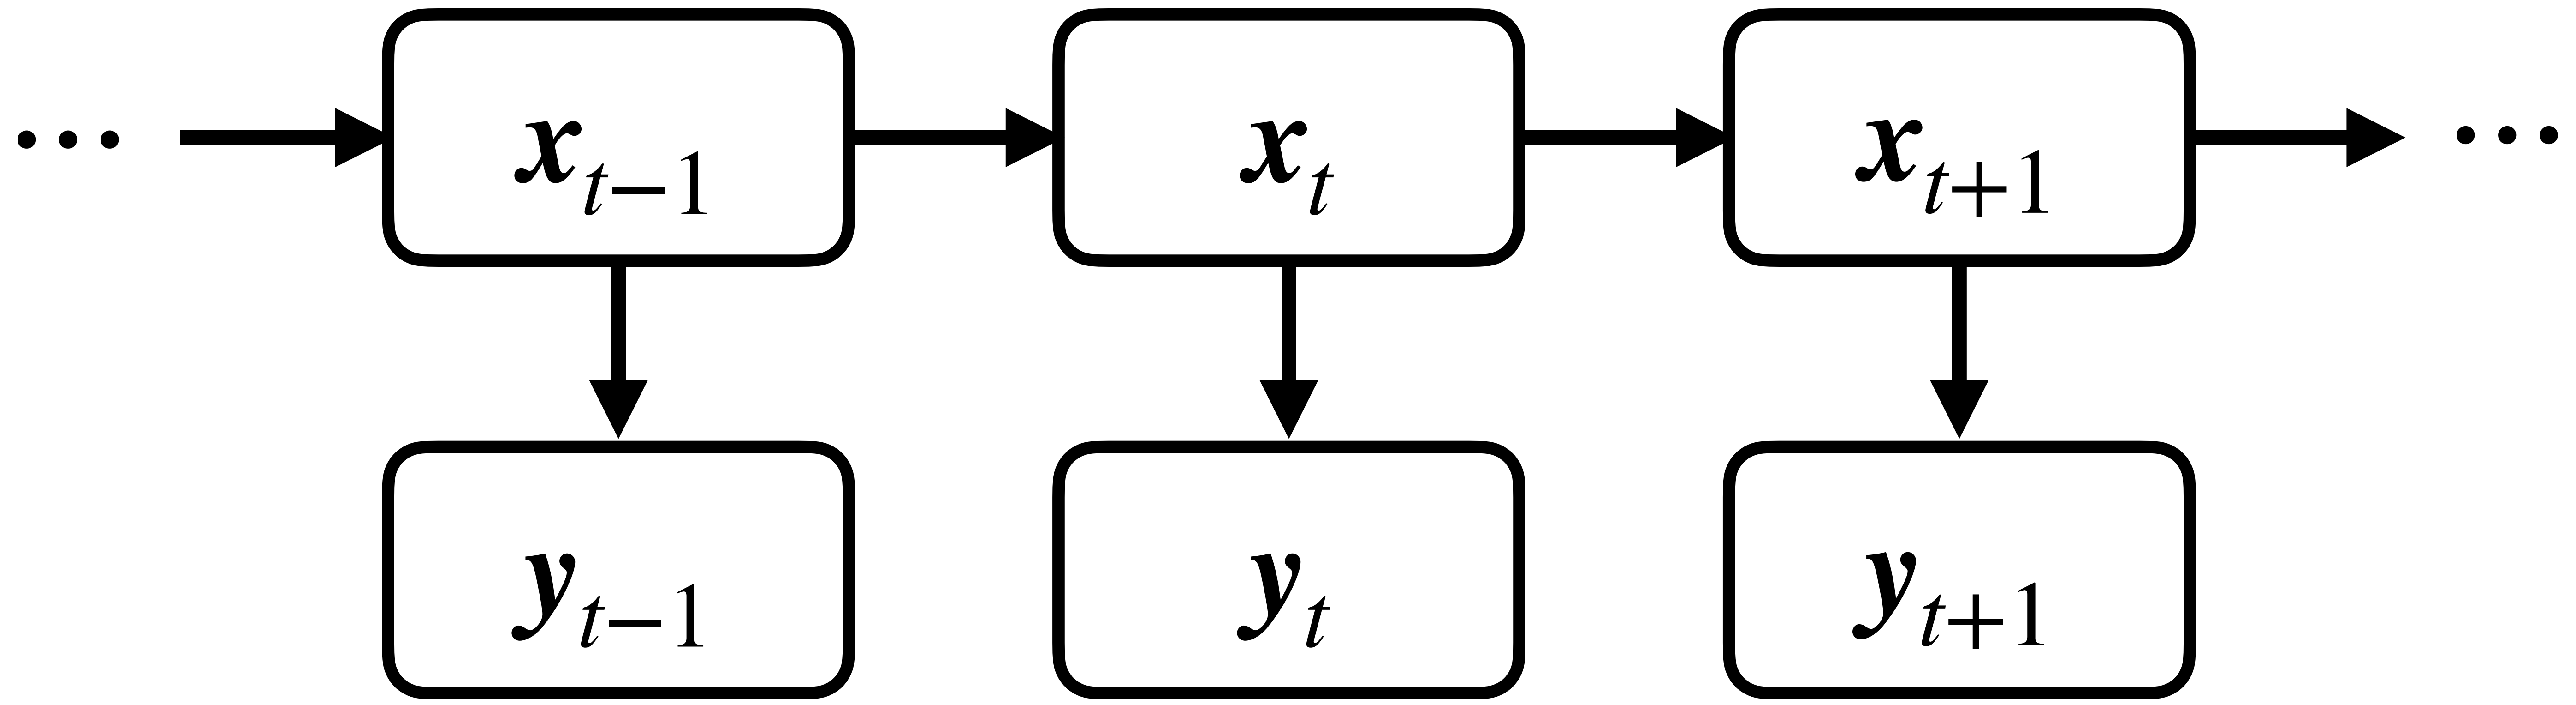


$X := \boldsymbol{x}_{1:t}$: entire trajectory $\implies p(X) = p(\boldsymbol{x}_t | \boldsymbol{x}_{t-1}) \cdots p(\boldsymbol{x}_2 | \boldsymbol{x}_1) p(\boldsymbol{x}_1)$

$Y := \boldsymbol{y}_{1:t}$: measurements over all timesteps $\implies p(Y | X) = \prod_\tau p(\boldsymbol{y}_\tau | \boldsymbol{x}_\tau)$

Filtering: posterior for current state given history of measurements $p(\boldsymbol{x}_t | \boldsymbol{y}_{1:t})$

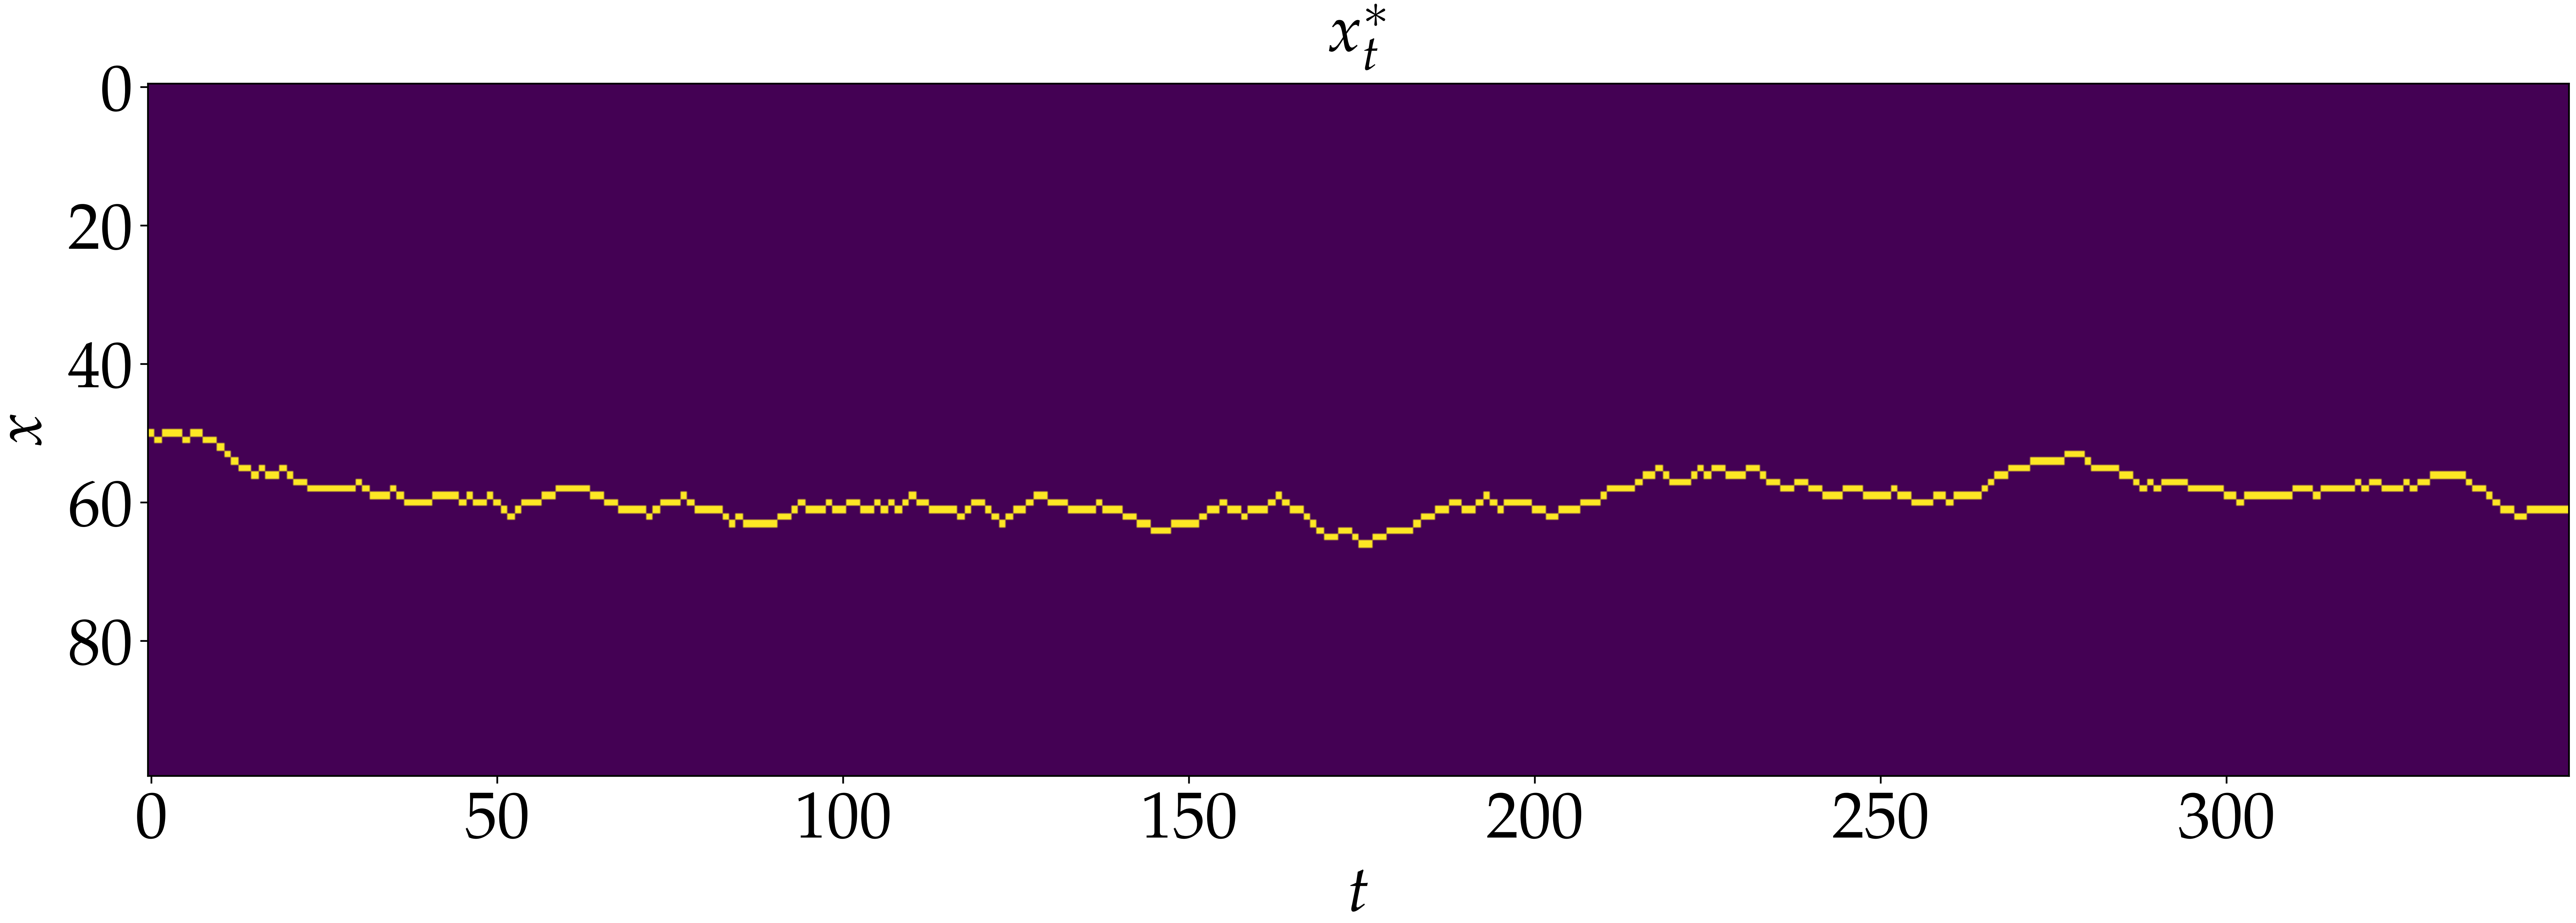

$x_t^*$
0
20
40
60
80
$x$
0
50
100
150
200
250
300
$t$

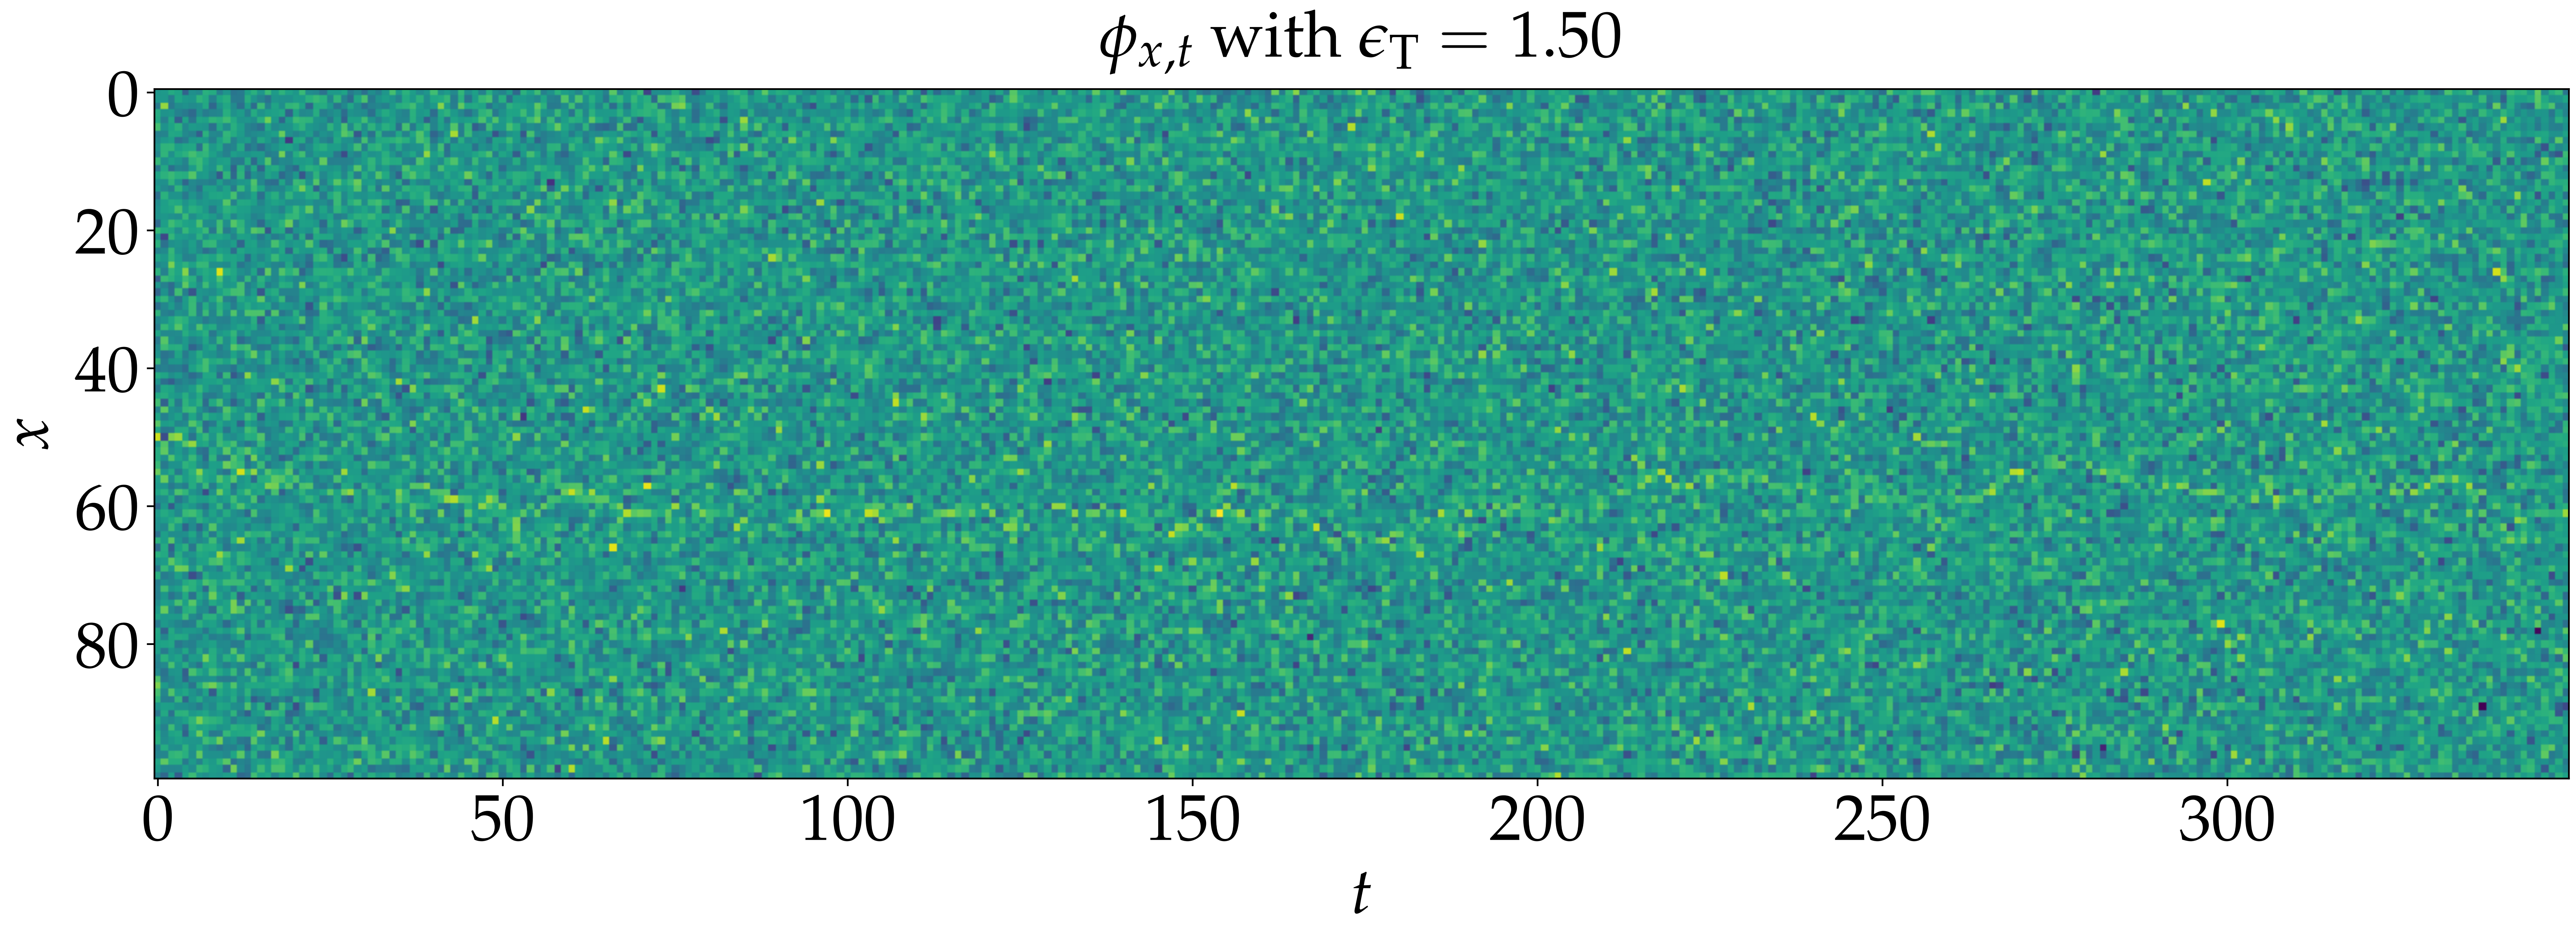

$\phi_{x,t}$ with $\epsilon_T = 1.50$
0
20
40
60
80
$x$
0
50
100
150
200
250
300
$t$

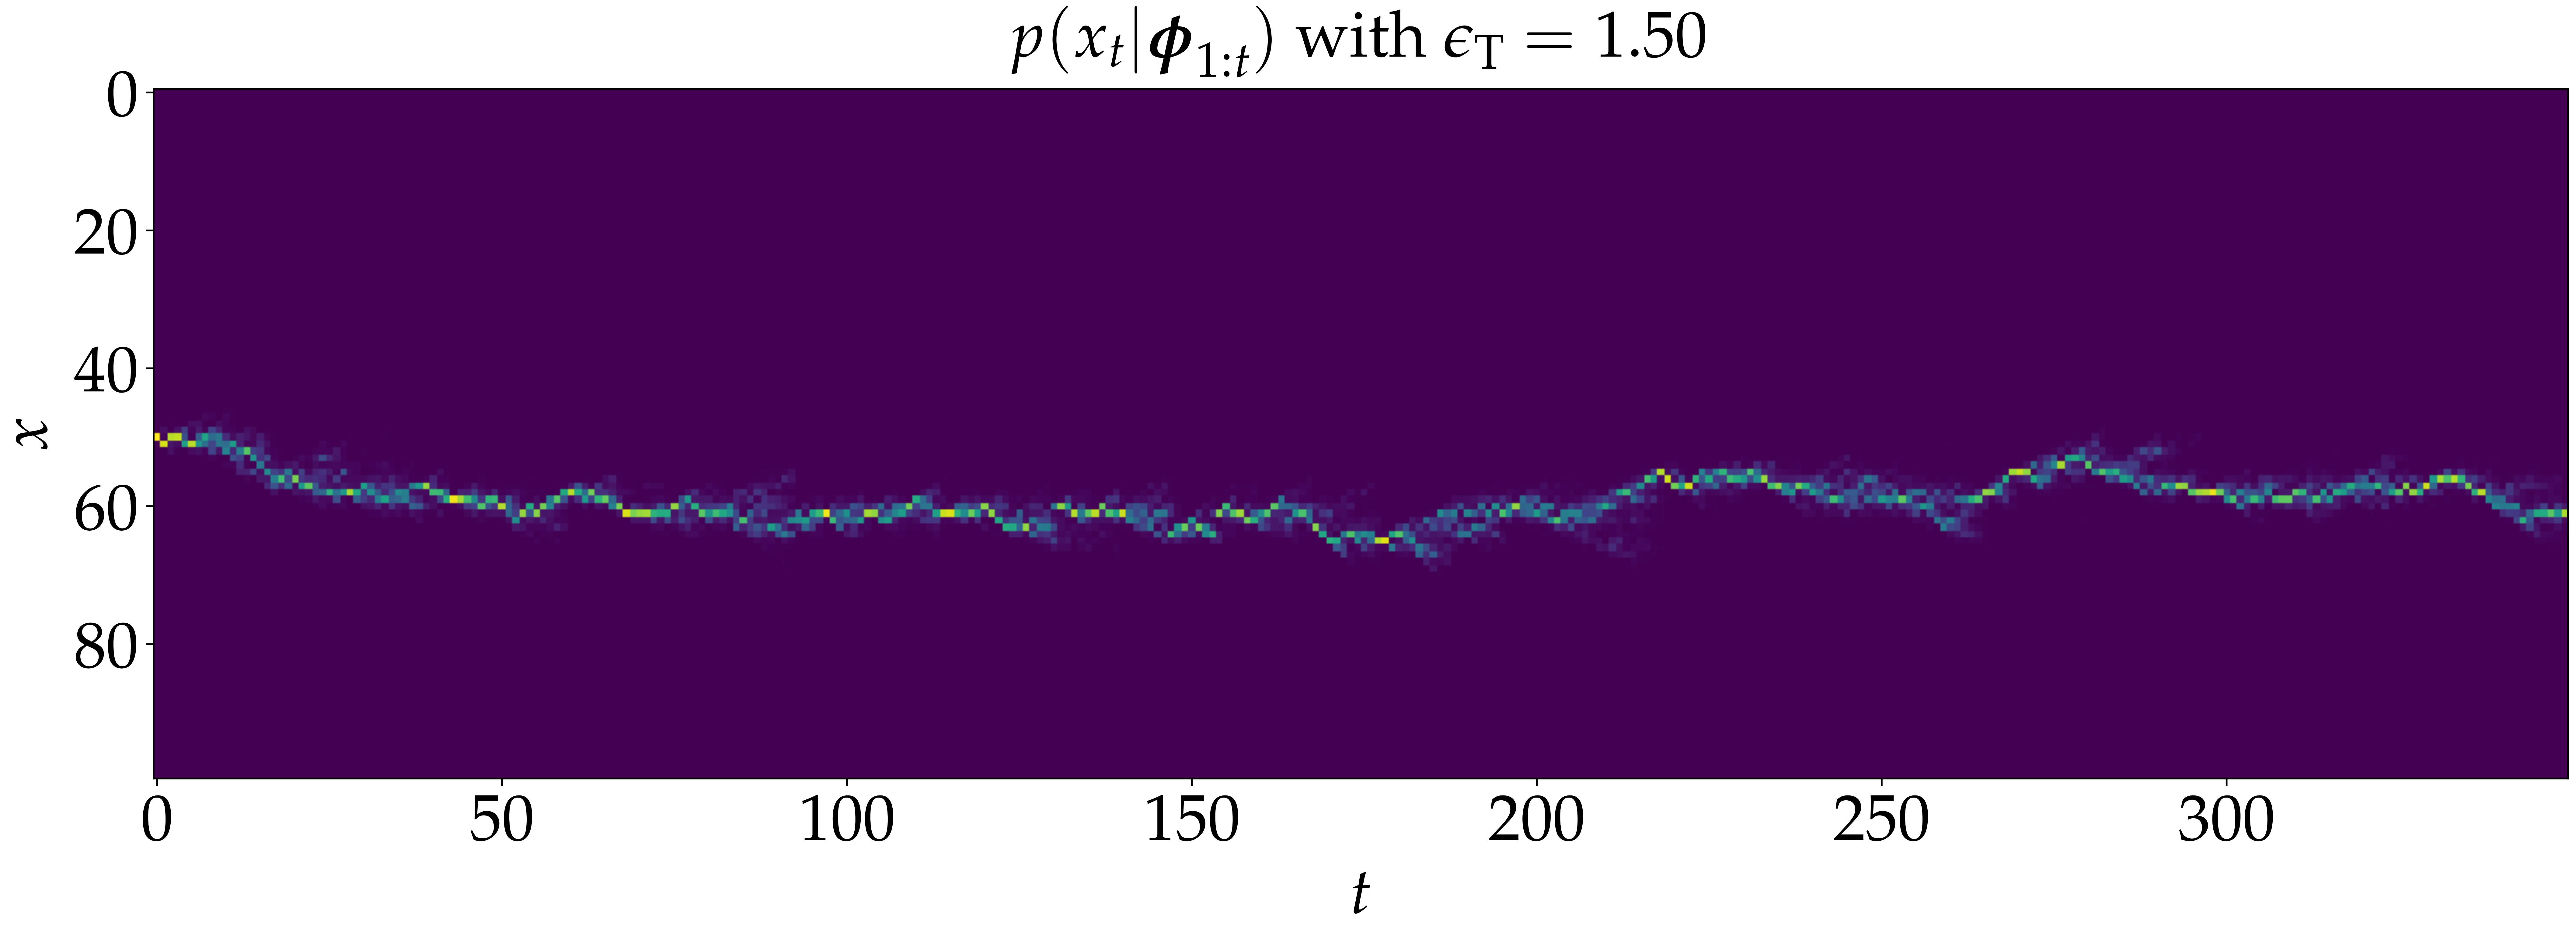
$p(x_t|\boldsymbol{\phi}_{1:t})$ with $\epsilon_{\mathrm{T}} = 1.50$
0
20
40
60
80
$x$
0
50
100
150
200
250
300
$t$

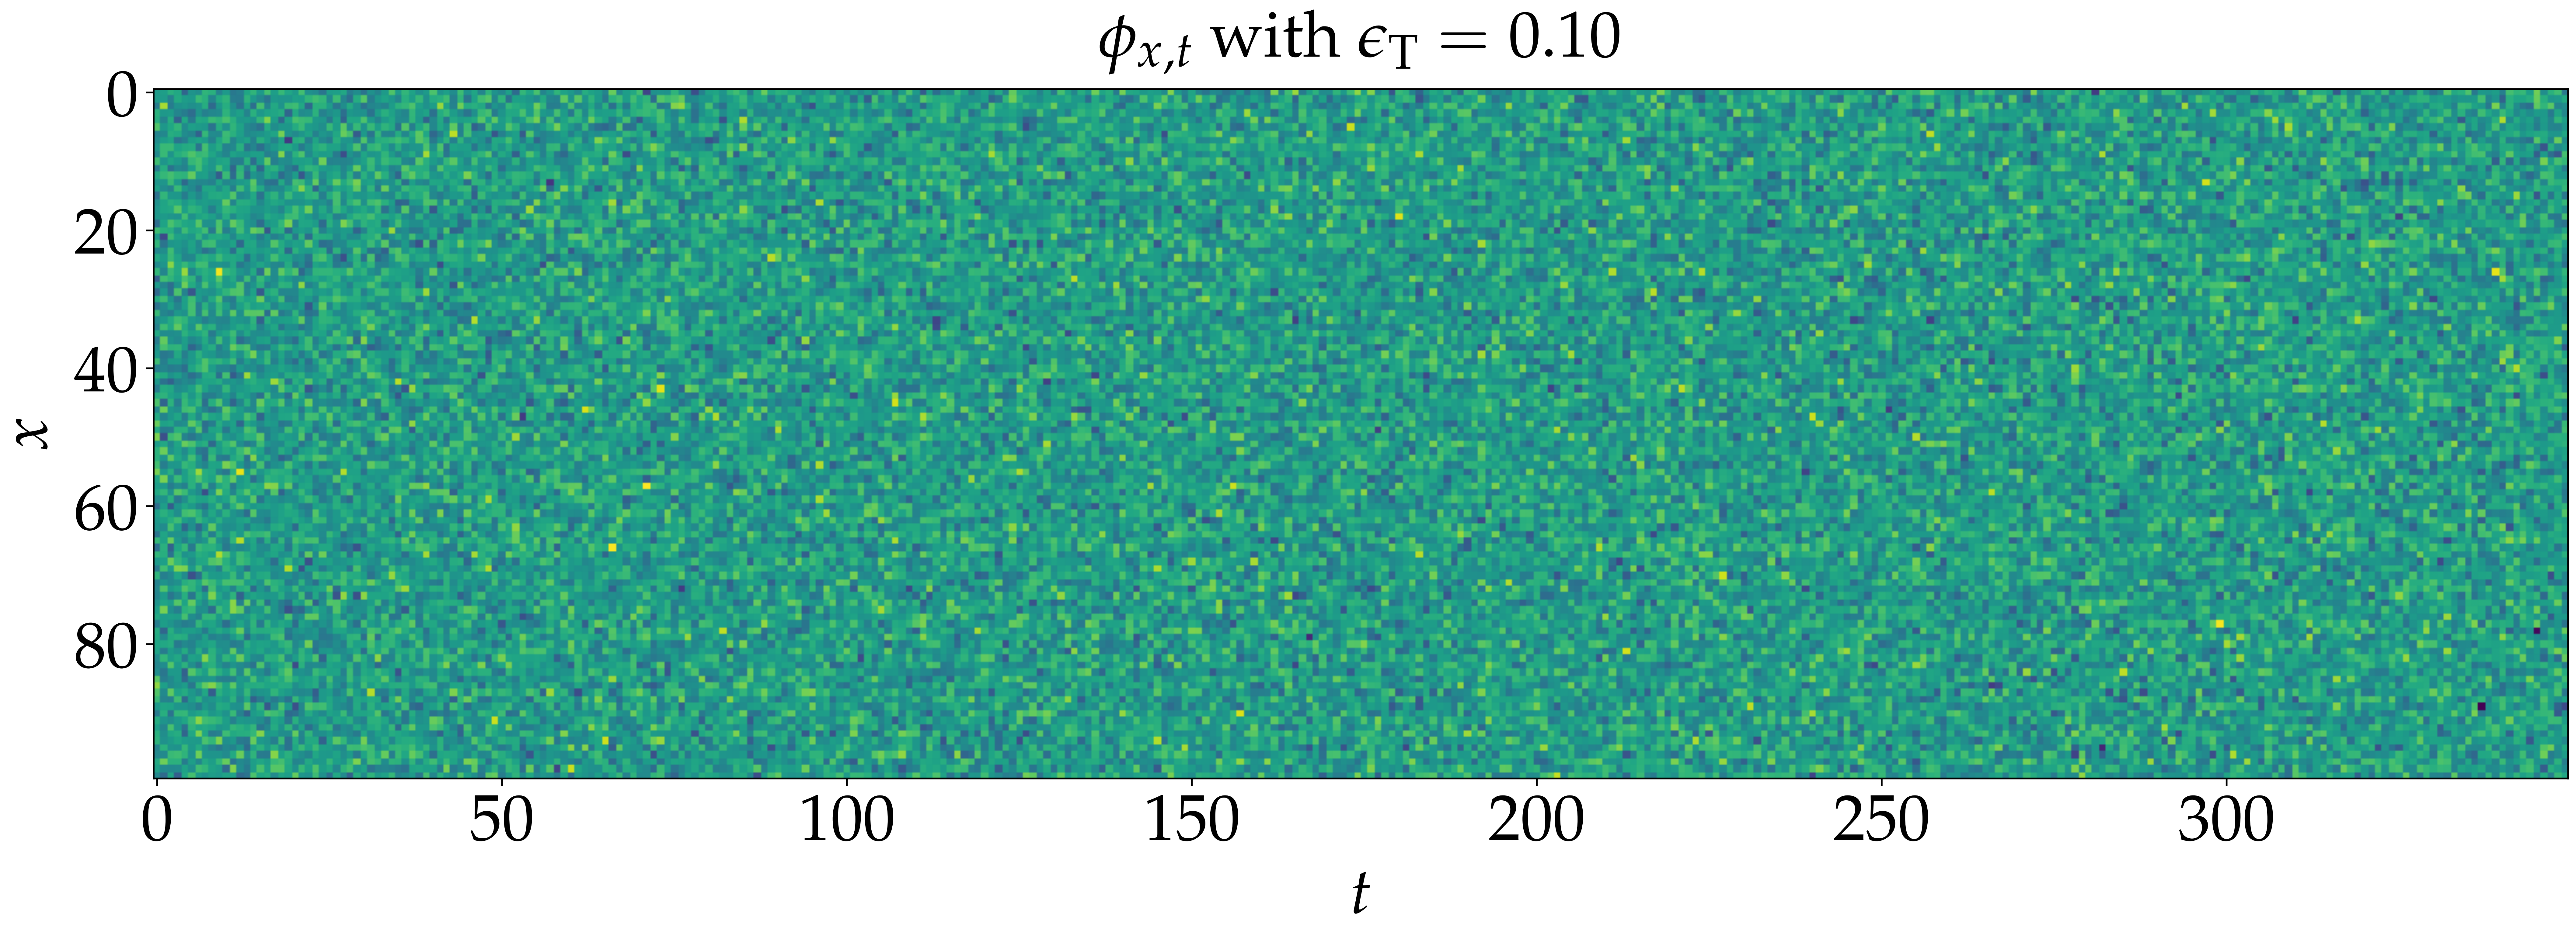
$\phi_{x,t}$ with $\epsilon_T = 0.10$
0
20
40
60
80
x
0
50
100
150
200
250
300
t

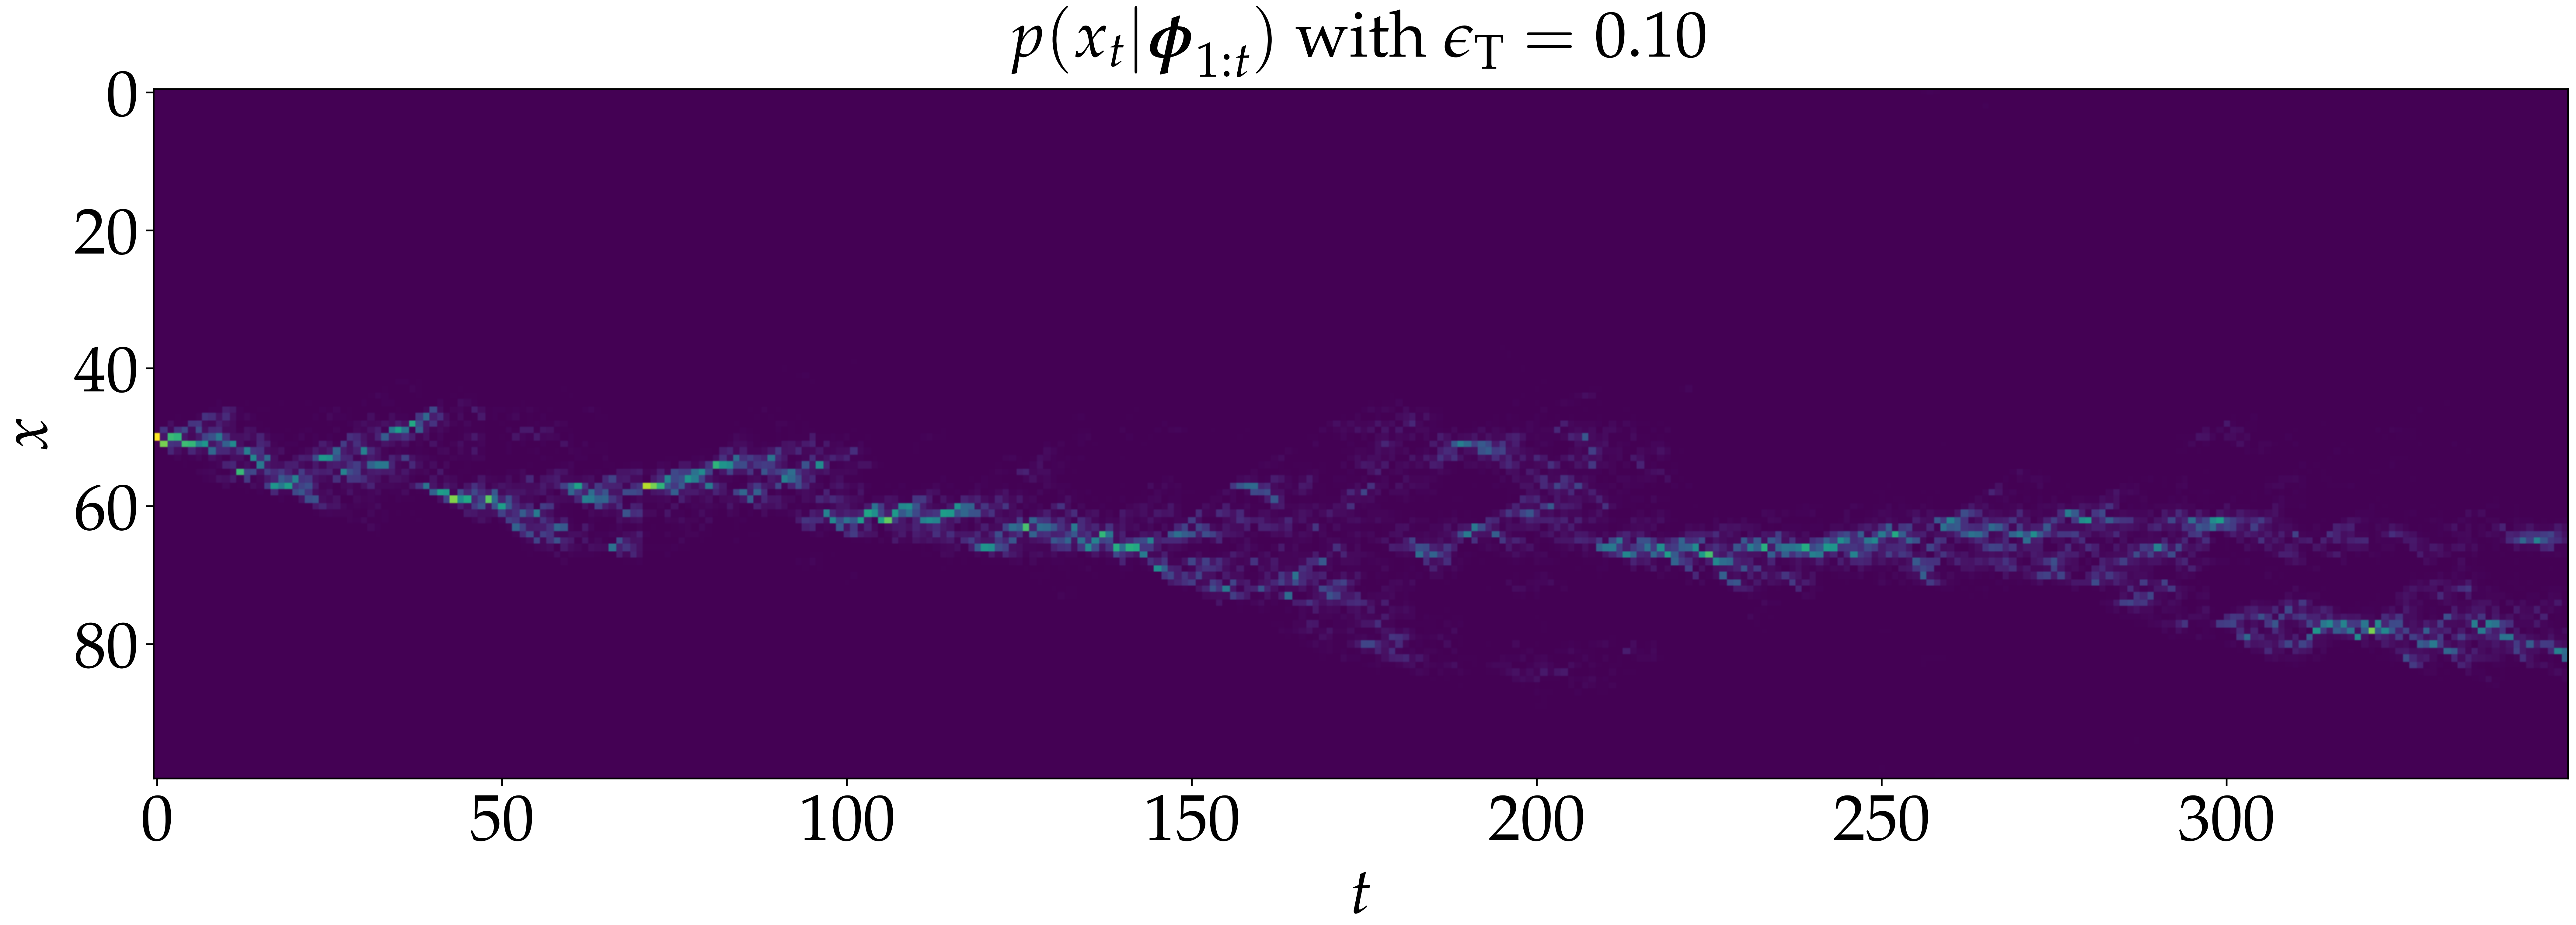

$p(x_t|\boldsymbol{\phi}_{1:t})$ with $\epsilon_{\mathrm{T}} = 0.10$
0
20
40
60
80
x
0
50
100
150
200
250
300
t

# Outline

1. Bayesian inference and hidden Markov models
2. **Teacher-student scenario and Bayes optimality**
3. General quantum inference problem and quantum hidden Markov models
4. Quantum error correction and the random-bond Ising model
5. The Planted SSEP and the planted XOR
6. Discussion of the phase diagrams

# Teacher-student scenario ([Zdeborová & Krzakala, 2016])

- Teacher generates true state $X^* \sim p_*(X^*)$ then data $Y \sim p_*(Y|X^*)$ with teacher's parameters $\theta_*$
- Student receives only $Y$ and conducts Bayesian inference assuming student parameters $\theta$ to generate posterior for inferred state $X$, $p_s(X|Y)$
- Joint distribution is known as the **planted ensemble**

$$p(X, Y, X^*) = p_s(X|Y)p_*(Y|X^*)p_*(X^*) = \frac{p_s(Y|X)p_s(X)p_*(Y|X^*)p_*(X^*)}{p_s(Y)}$$

- At Bayes optimality (Nishimori condition): $s = *$ , $X$ distributed identically to $X^*$
- **Note:** Even with full knowledge of teacher's parameters, perfect inference is not possible in general as data is still generated randomly

# Planting

- Student's posterior can be thought of as a Gibbs probability:

$$p_s(X|Y) = \frac{p_s(Y|X)p_s(X)}{p_s(Y)} \leftrightarrow \frac{e^{-\beta H(X|Y)}}{Z_s(Y)}$$

- Then observations $Y$ can be thought of as a "disorder field":

$$p_*(Y) = \sum_{X^*} p_*(Y|X^*)p_*(X^*)$$

- However, $Y = (y_{x,t})_{x\in 1:L, t\in 1:T}$ are correlated

**cf.** traditional disordered systems where $Y$ is iid

The true configuration is "planted" in the disorder

# Bayes optimality ( $* = s$ )

Why is it called Bayes optimal?

For $X \in \mathcal{X}$, consider classifier $h(Y) \in \mathcal{X}$

Classification performance $\mathbb{E}_{Y,X*}\left[\delta_{h(Y),X*}\right]$

**Theorem**: Optimal classifier $h_{\text{optimal}}(Y) = \underset{X}{\text{argmax}}\, p_*(X \mid Y)$

Consider Estimator $\tilde{O}(Y) \in \mathbb{R}$ for real-valued observables $O(X) \in \mathbb{R}$

Average mean-squared error of estimator $\mathbb{E}_{Y,X*}\left[(\tilde{O}(Y) - O(X*))^2\right]$

**Theorem**: Optimal estimator $\tilde{O}(Y) = \mathbb{E}_{X \sim p_*(X|Y)}\left[O(X)\right]$

# Observables

- Mean-squared error of the mean $\mathrm{MSEM} = \mathbb{E}_{Y,X*}\left[\left(\langle O(X)\rangle_{\mathrm{s}} - O(X*)\right)^2\right]$
- Mean-squared error $\mathrm{MSE} = \mathbb{E}_{X,Y,X*}\left[\left(O(X) - O(X*)\right)^2\right]$
- Observable variance $\delta O_{\mathrm{s}}^2 = \mathbb{E}_{Y,X*}\left[\langle O(X)^2\rangle_{\mathrm{s}} - \langle O(X)\rangle_{\mathrm{s}}^2\right]$
- $\langle \cdot \rangle_{\mathrm{s}} = \mathbb{E}_{X\sim p_{\mathrm{s}}(X|Y)}[\, \cdot \,]$

At Bayes optimality, $\mathrm{MSEM} = \frac{1}{2}\mathrm{MSE} = \delta O_{\mathrm{s}}^2$!

# Outline

1. Bayesian inference and hidden Markov models
2. Teacher-student scenario and Bayes optimality
3. **General quantum inference problem and quantum hidden Markov models**
4. Quantum error correction and the random-bond Ising model
5. The Planted SSEP and the planted XOR
6. Discussion of the phase diagrams

# General quantum inference problem

- Teacher prepares state $\rho_{*,0}$
- Applies channel $\mathscr{C}_*$ which records outcomes in a **classical register**

$$\mathscr{C}_* : \rho_{*,0} \to \sum_{O*,Y} \rho_*(Y, O*) \otimes |Y\rangle\langle Y| \otimes |O*\rangle\langle O*|$$

- $Y$, $O*$ could be measurement outcomes, or record of randomly chosen subchannel
- $Y$ is revealed to the student. $O*$ is hidden from the student.
- Goal of student: infer $O*$!

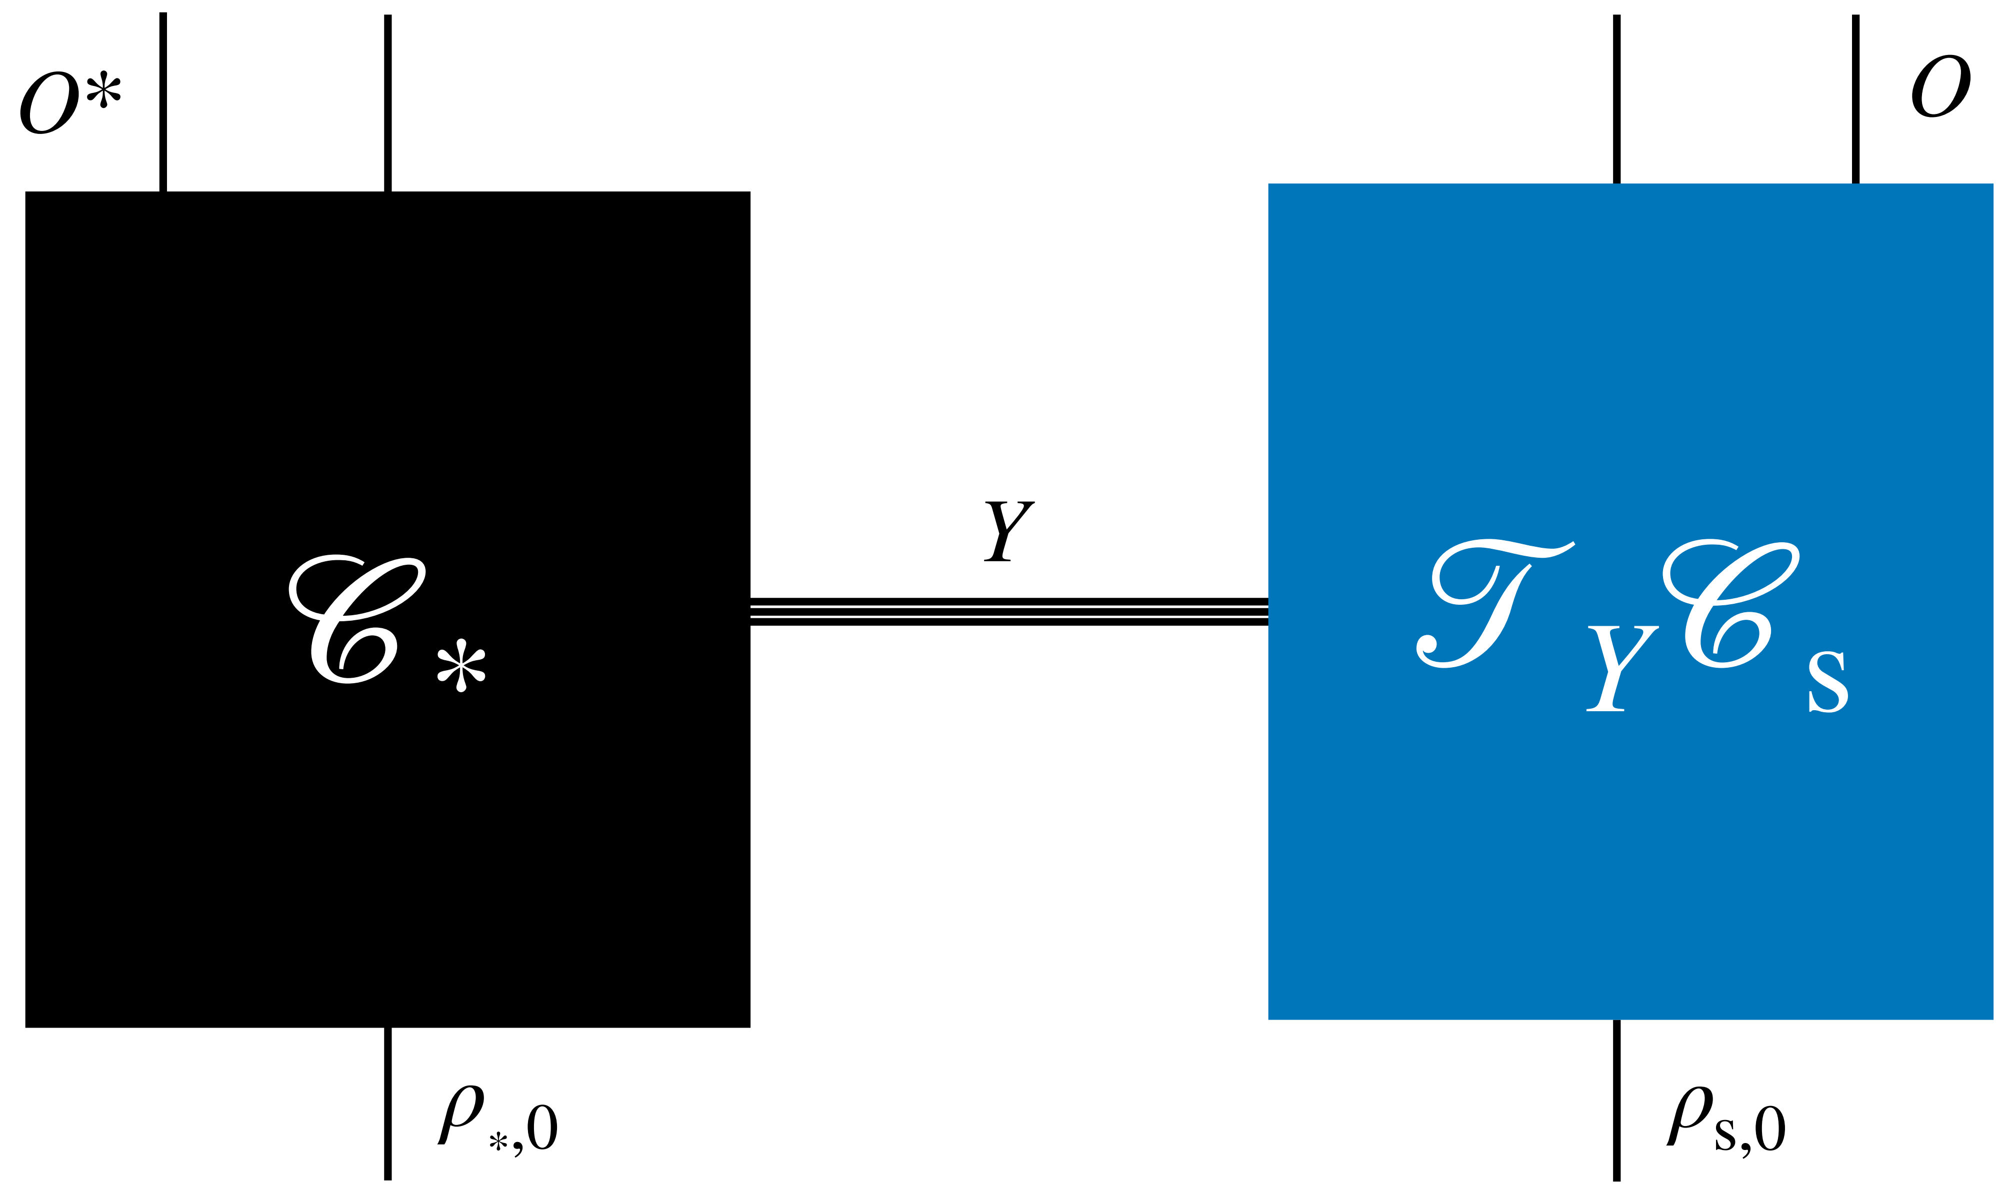

# General quantum inference problem

- What's the best that the student can do, assuming model $\mathscr{C}_\mathrm{s}$?

- → *Simulate* a density matrix, and **condition** it on measurement outcomes $Y$

$$p_\mathrm{s}(O\,|\,Y) = \frac{\mathrm{tr}[\mathbb{P}_O \mathscr{T}_Y \mathscr{C}_\mathrm{s} \rho_{\mathrm{s},0}]}{\mathrm{tr}[\mathscr{T}_Y \mathscr{C}_\mathrm{s}(\rho_{\mathrm{s},0})]}$$

- $\mathscr{T}_Y = \mathrm{tr}_{\mathscr{Y}} \mathscr{P}_Y$ (project then trace out)

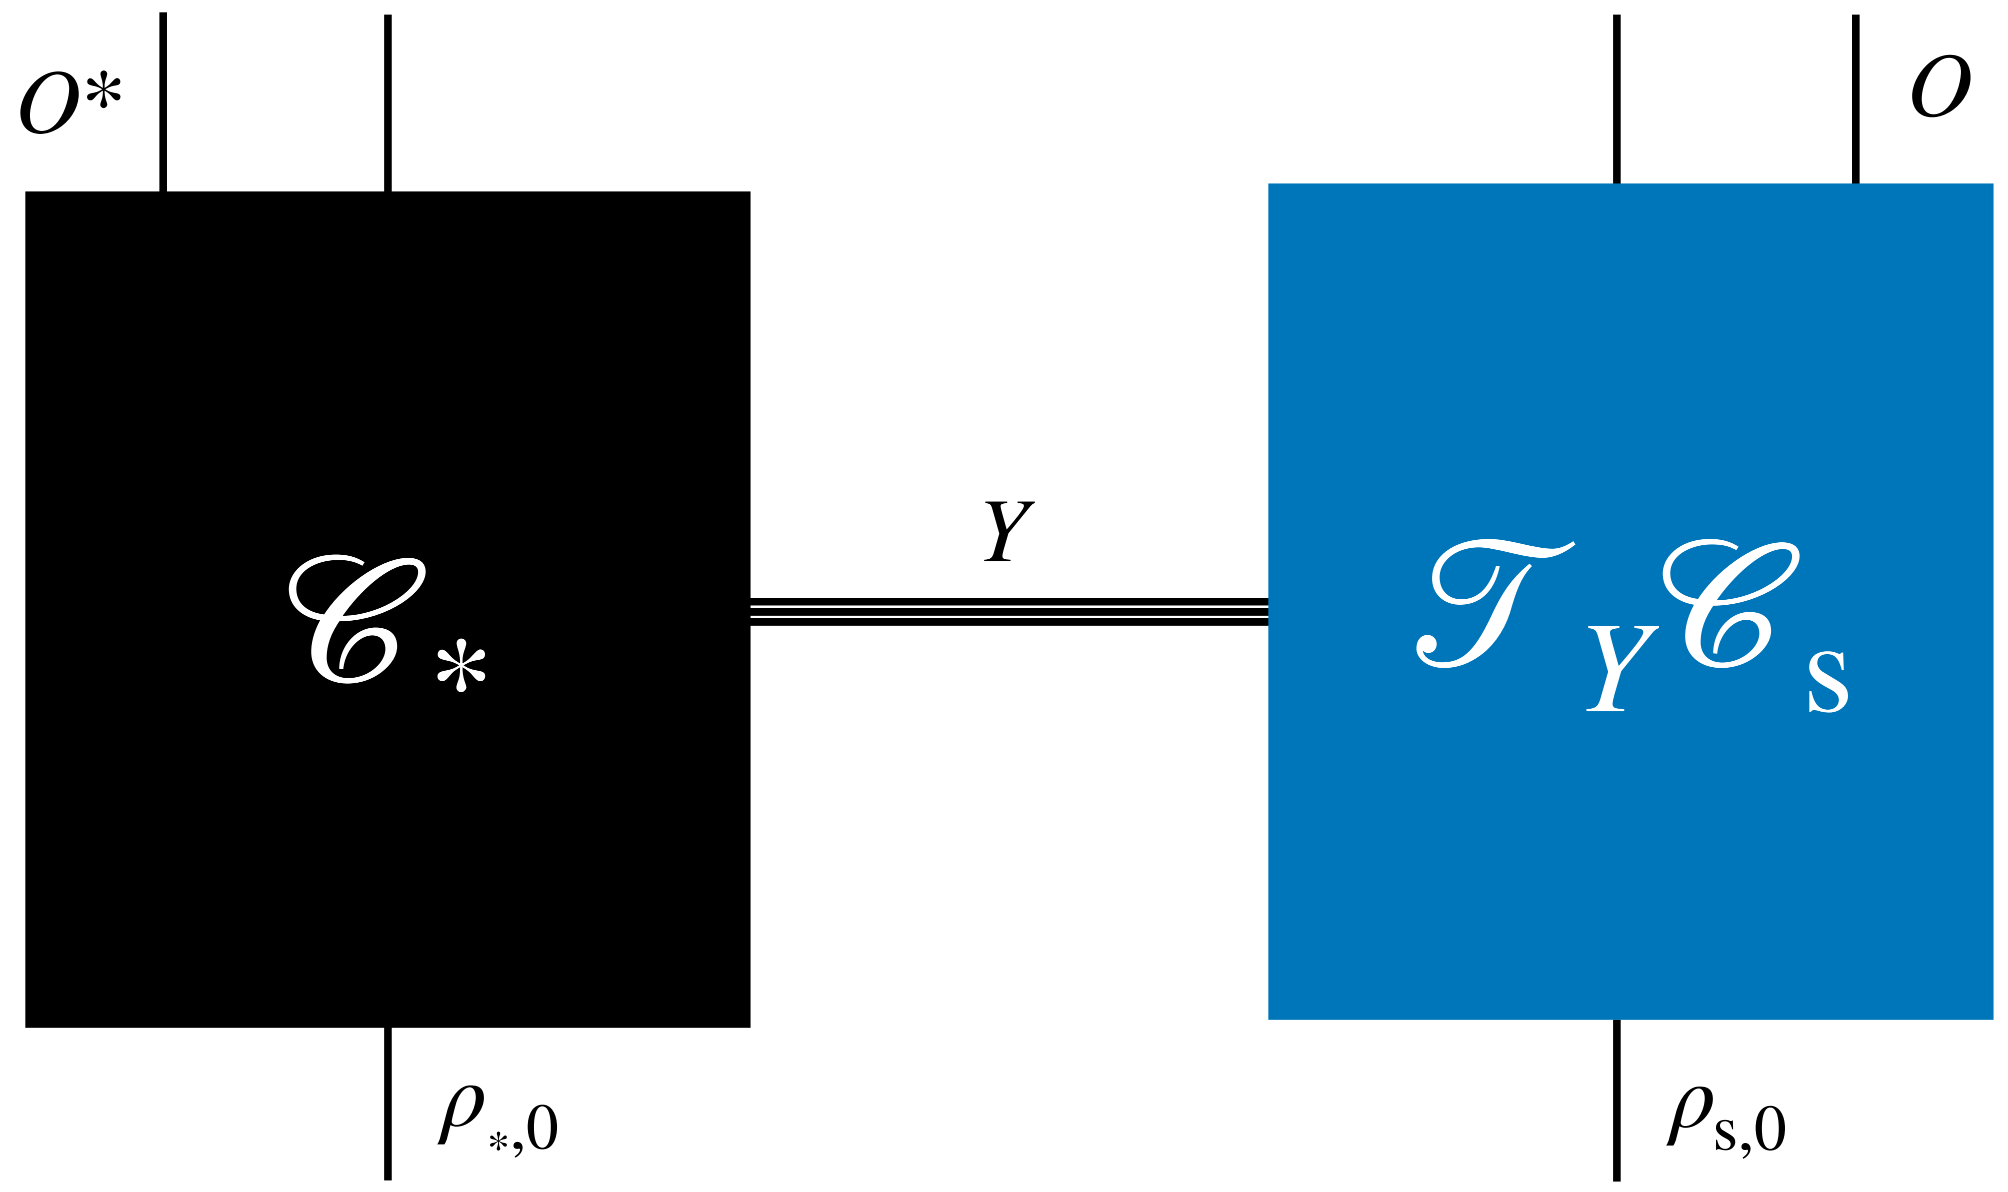

# General quantum inference problem

- Think of student's simulated density matrix as an actual density matrix:

$$\rho = \sum_Y \frac{\rho_{\mathrm{s}}(Y)}{\mathrm{tr}[\rho_{\mathrm{s}}(Y)]} \otimes |Y\rangle\langle Y| \otimes \rho_*(Y)$$

- Equivalent of the planted ensemble in the classical setting
- Is symmetric under $* \leftrightarrow \mathrm{s}$ at Bayes optimality

- Can study observable sharpening of the teacher's system:

$$\delta O_*^2 = \sum_Y \left( \frac{\mathrm{tr}[O_*^2 \mathscr{P}_Y \rho]}{\mathrm{tr}[\mathscr{P}_Y \rho]} - \frac{\mathrm{tr}[O_* \mathscr{P}_Y \rho]^2}{\mathrm{tr}[\mathscr{P}_Y \rho]^2} \right)$$

- At Bayes optimality $\delta O_*^2 = \delta O_{\mathrm{s}}^2 = \mathrm{MSEM}$! 'Sharpening' = 'learnability'

# Quantum hidden Markov models (qHMMs)

- Break down whole channel into subchannels

- Some channels record outcomes:

$$\mathscr{R}_* : \rho_* \to \sum_{\boldsymbol{a}} K_{*,\boldsymbol{a}} \rho_* K_{*,\boldsymbol{a}}^\dagger \otimes |\boldsymbol{a}\rangle\langle\boldsymbol{a}|$$

- $Y = (\boldsymbol{y}_t)_{t\in 1:T}$ are revealed, while $O^*$ are hidden

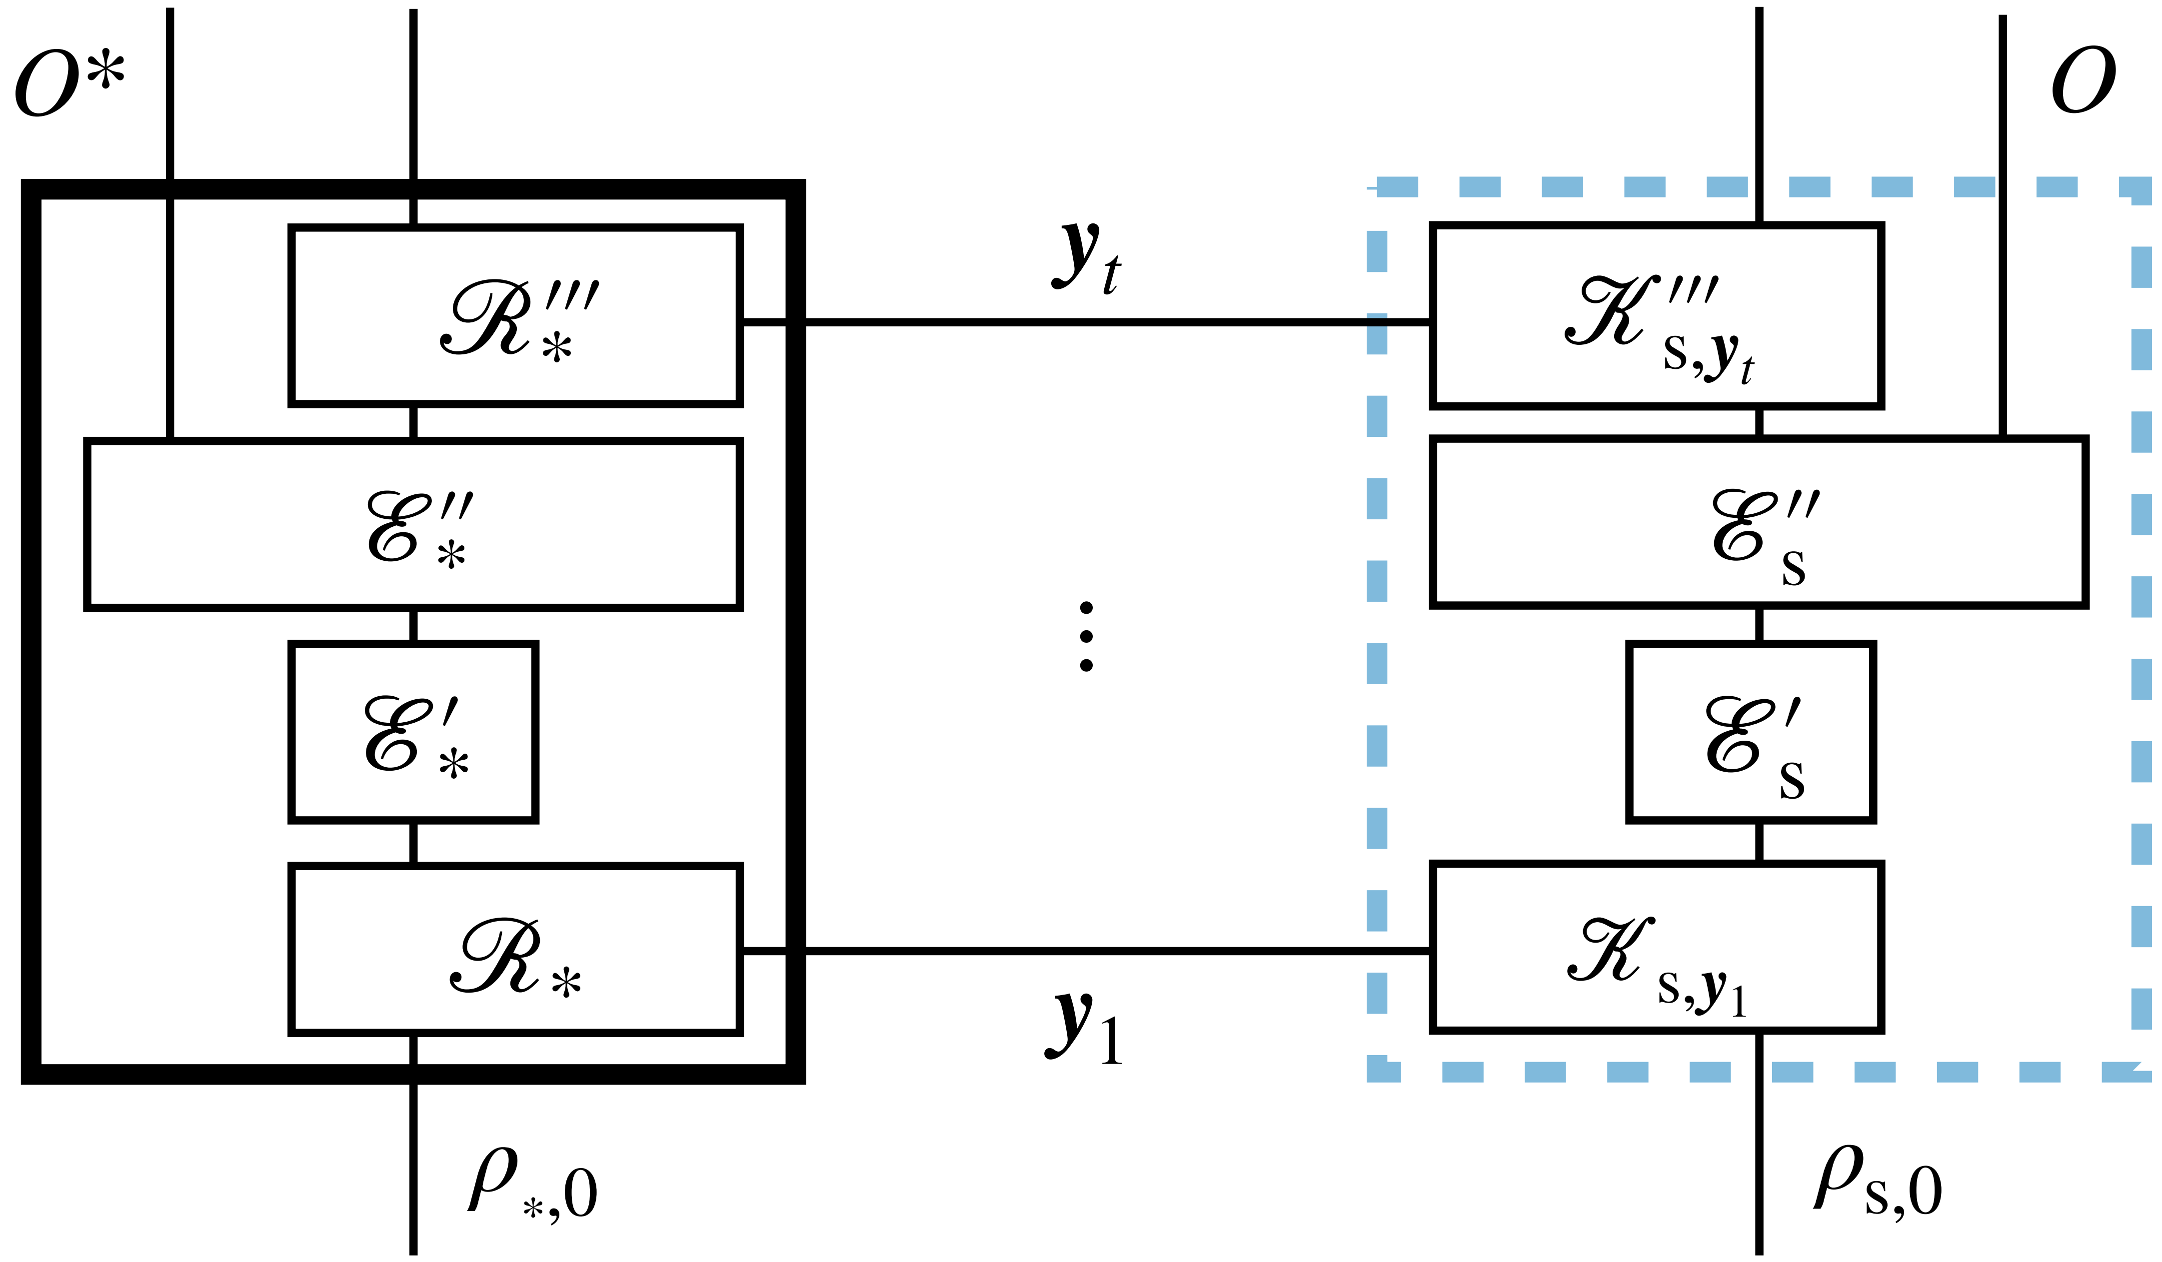

# Quantum hidden Markov models (qHMMs)

- Assume that channel applied at each step **does not depend on** any of the previously revealed hidden registers

- Then, can condition channel at each timestep separately:

$$\mathcal{T}_{\boldsymbol{y}_t}\mathcal{R}_{\mathrm{s}} = \mathcal{K}_{\mathrm{s},\boldsymbol{y}_t} : \rho_{\mathrm{s}} \to K_{\mathrm{s},\boldsymbol{y}_t}\rho_{\mathrm{s}}K^{\dagger}_{\mathrm{s},\boldsymbol{y}_t}$$

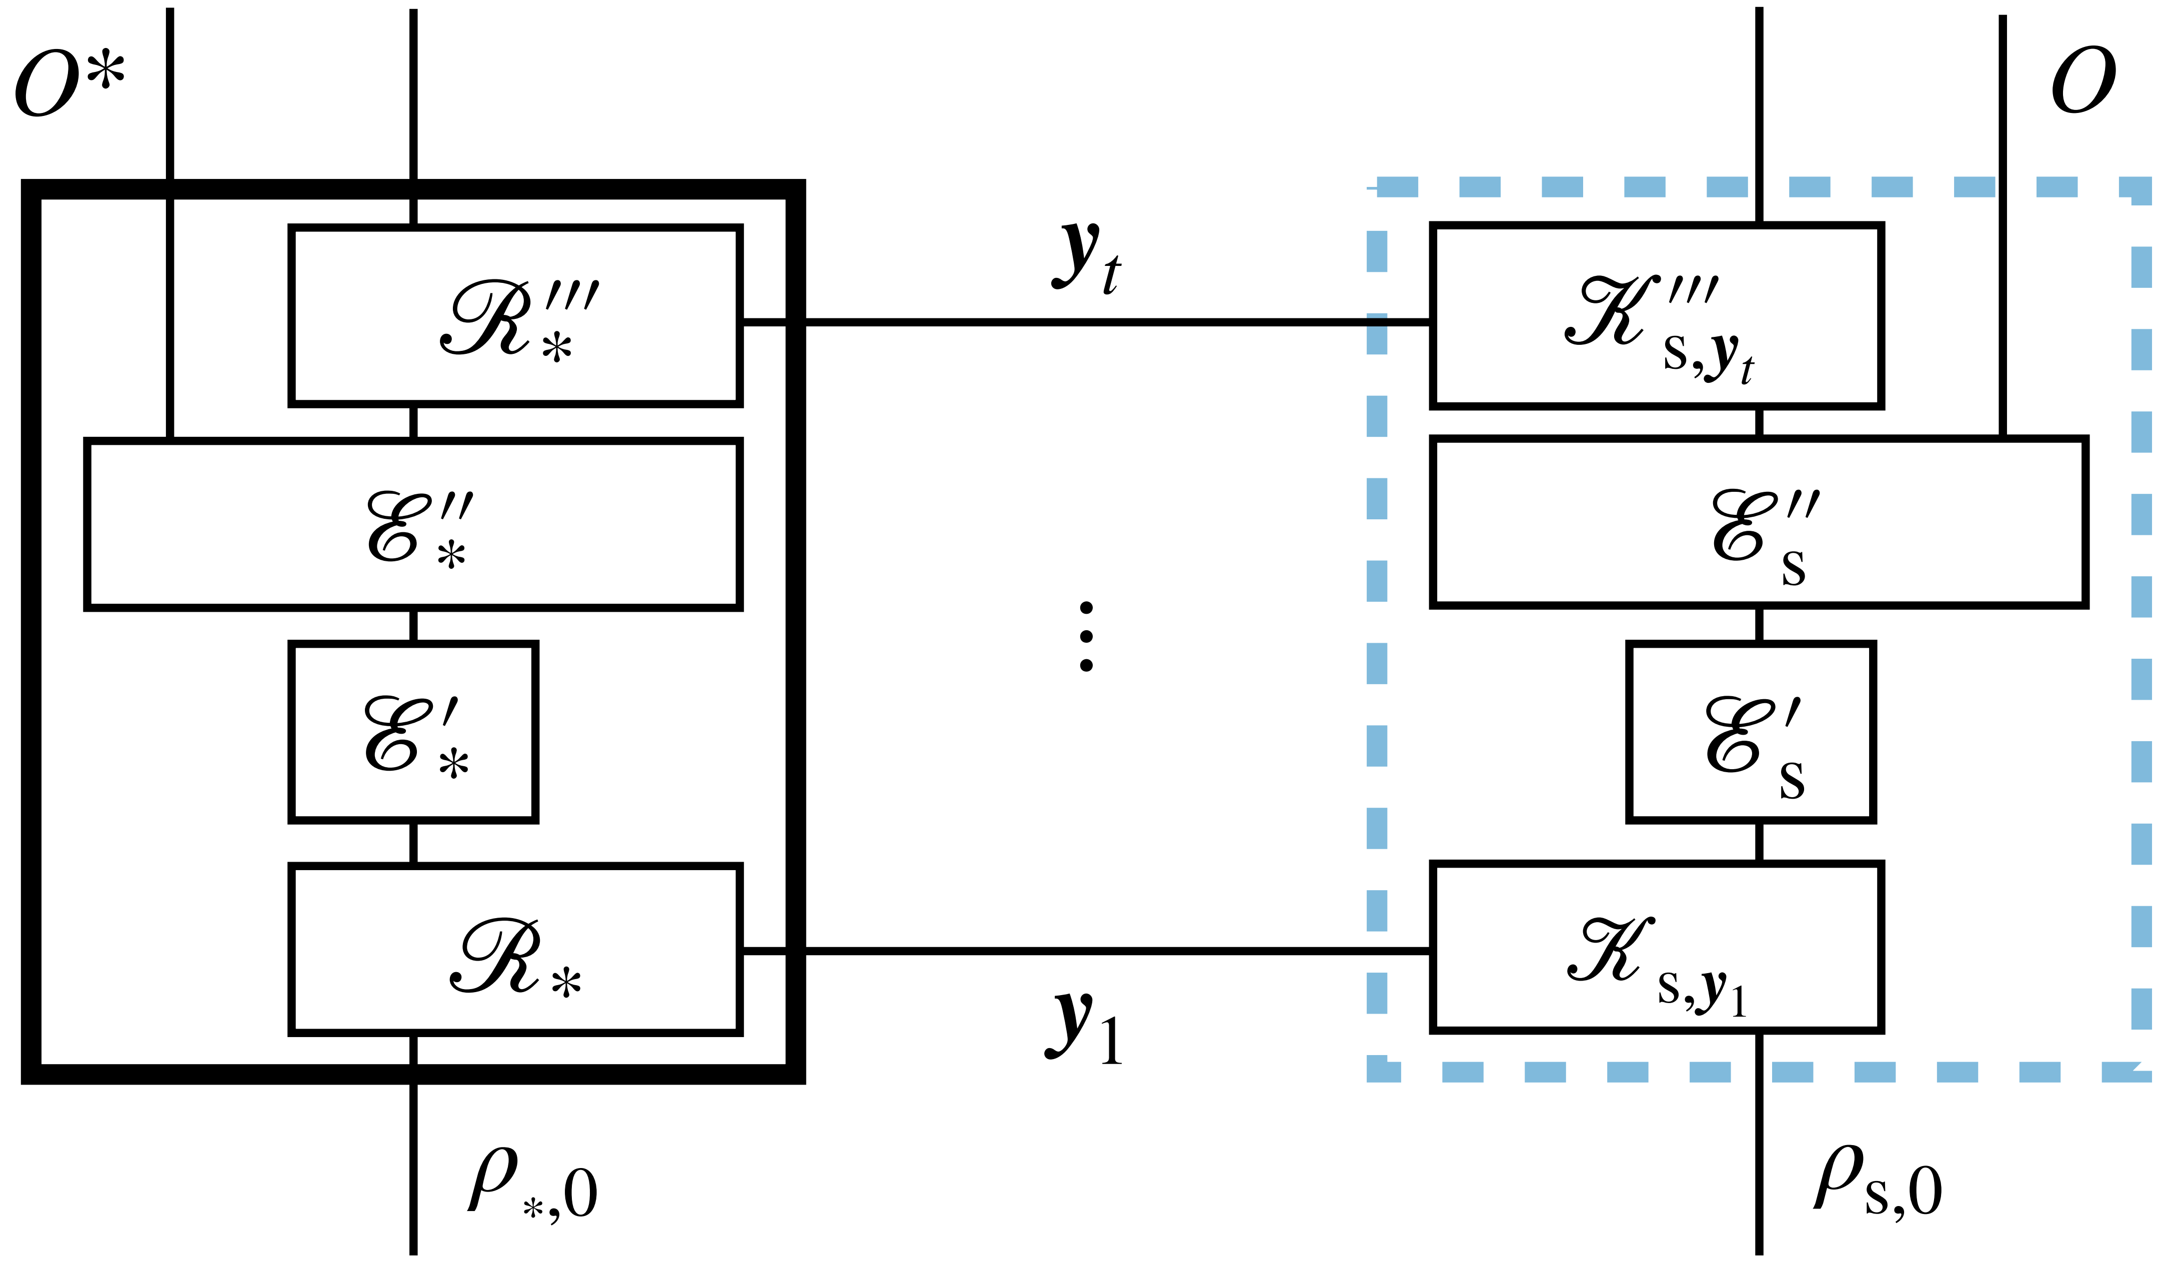

# Connection with HMMs

- For HMMs, the (*unnormalised*) posterior distribution for the **filtering** task evolves according to "the forward algorithm"

$$q_s(\boldsymbol{x}_t | \boldsymbol{y}_{1:t}) = \sum_{\boldsymbol{x}_{t-1}} p_s(\boldsymbol{y}_t | \boldsymbol{x}_t) p_s(\boldsymbol{x}_t | \boldsymbol{x}_{t-1}) q_s(\boldsymbol{x}_{t-1} | \boldsymbol{y}_{1:t-1})$$

- Using vector notation for posterior ex. $q_s(\boldsymbol{x}_t | \boldsymbol{y}_{1:t}) = (\boldsymbol{x}_t | q_s(\boldsymbol{y}_{1:t}))$,

$$| q_s(\boldsymbol{y}_{1:t})) = \mathsf{K}_{s,\boldsymbol{y}_{1:t}} \mathsf{E}_{s,t} | q_s(\boldsymbol{y}_{1:t-1}))$$

- $\mathsf{K}_{s,\boldsymbol{y}_t}$: diagonal matrix; measurement model $(\boldsymbol{x}_t | \mathsf{K}_{s,\boldsymbol{y}_t} | \boldsymbol{x}_t) = p_s(\boldsymbol{y}_t | \boldsymbol{x}_t)$
- $\mathsf{E}_{s,t}$: prior Markov kernel $(\boldsymbol{x}_t | \mathsf{E}_{s,t} | \boldsymbol{x}_{t-1}) = p_s(\boldsymbol{x}_t | \boldsymbol{x}_{t-1})$

# Connection with HMMs

- for qHMMs, the student's density matrix evolves as:

$$\rho_{\mathrm{s}}(\boldsymbol{y}_{1:t}) = \mathscr{K}_{\mathrm{s},\boldsymbol{y}_t}\mathscr{E}_{\mathrm{s},t}\rho_{\mathrm{s}}(\boldsymbol{y}_{1:t-1})$$

- Group all channels between revealed registers as $\mathscr{E}_t$
- $\mathscr{K}_{\mathrm{s},\boldsymbol{y}_t}$ is the Kraus operator from conditioning on current measurement outcome $\boldsymbol{y}_t$

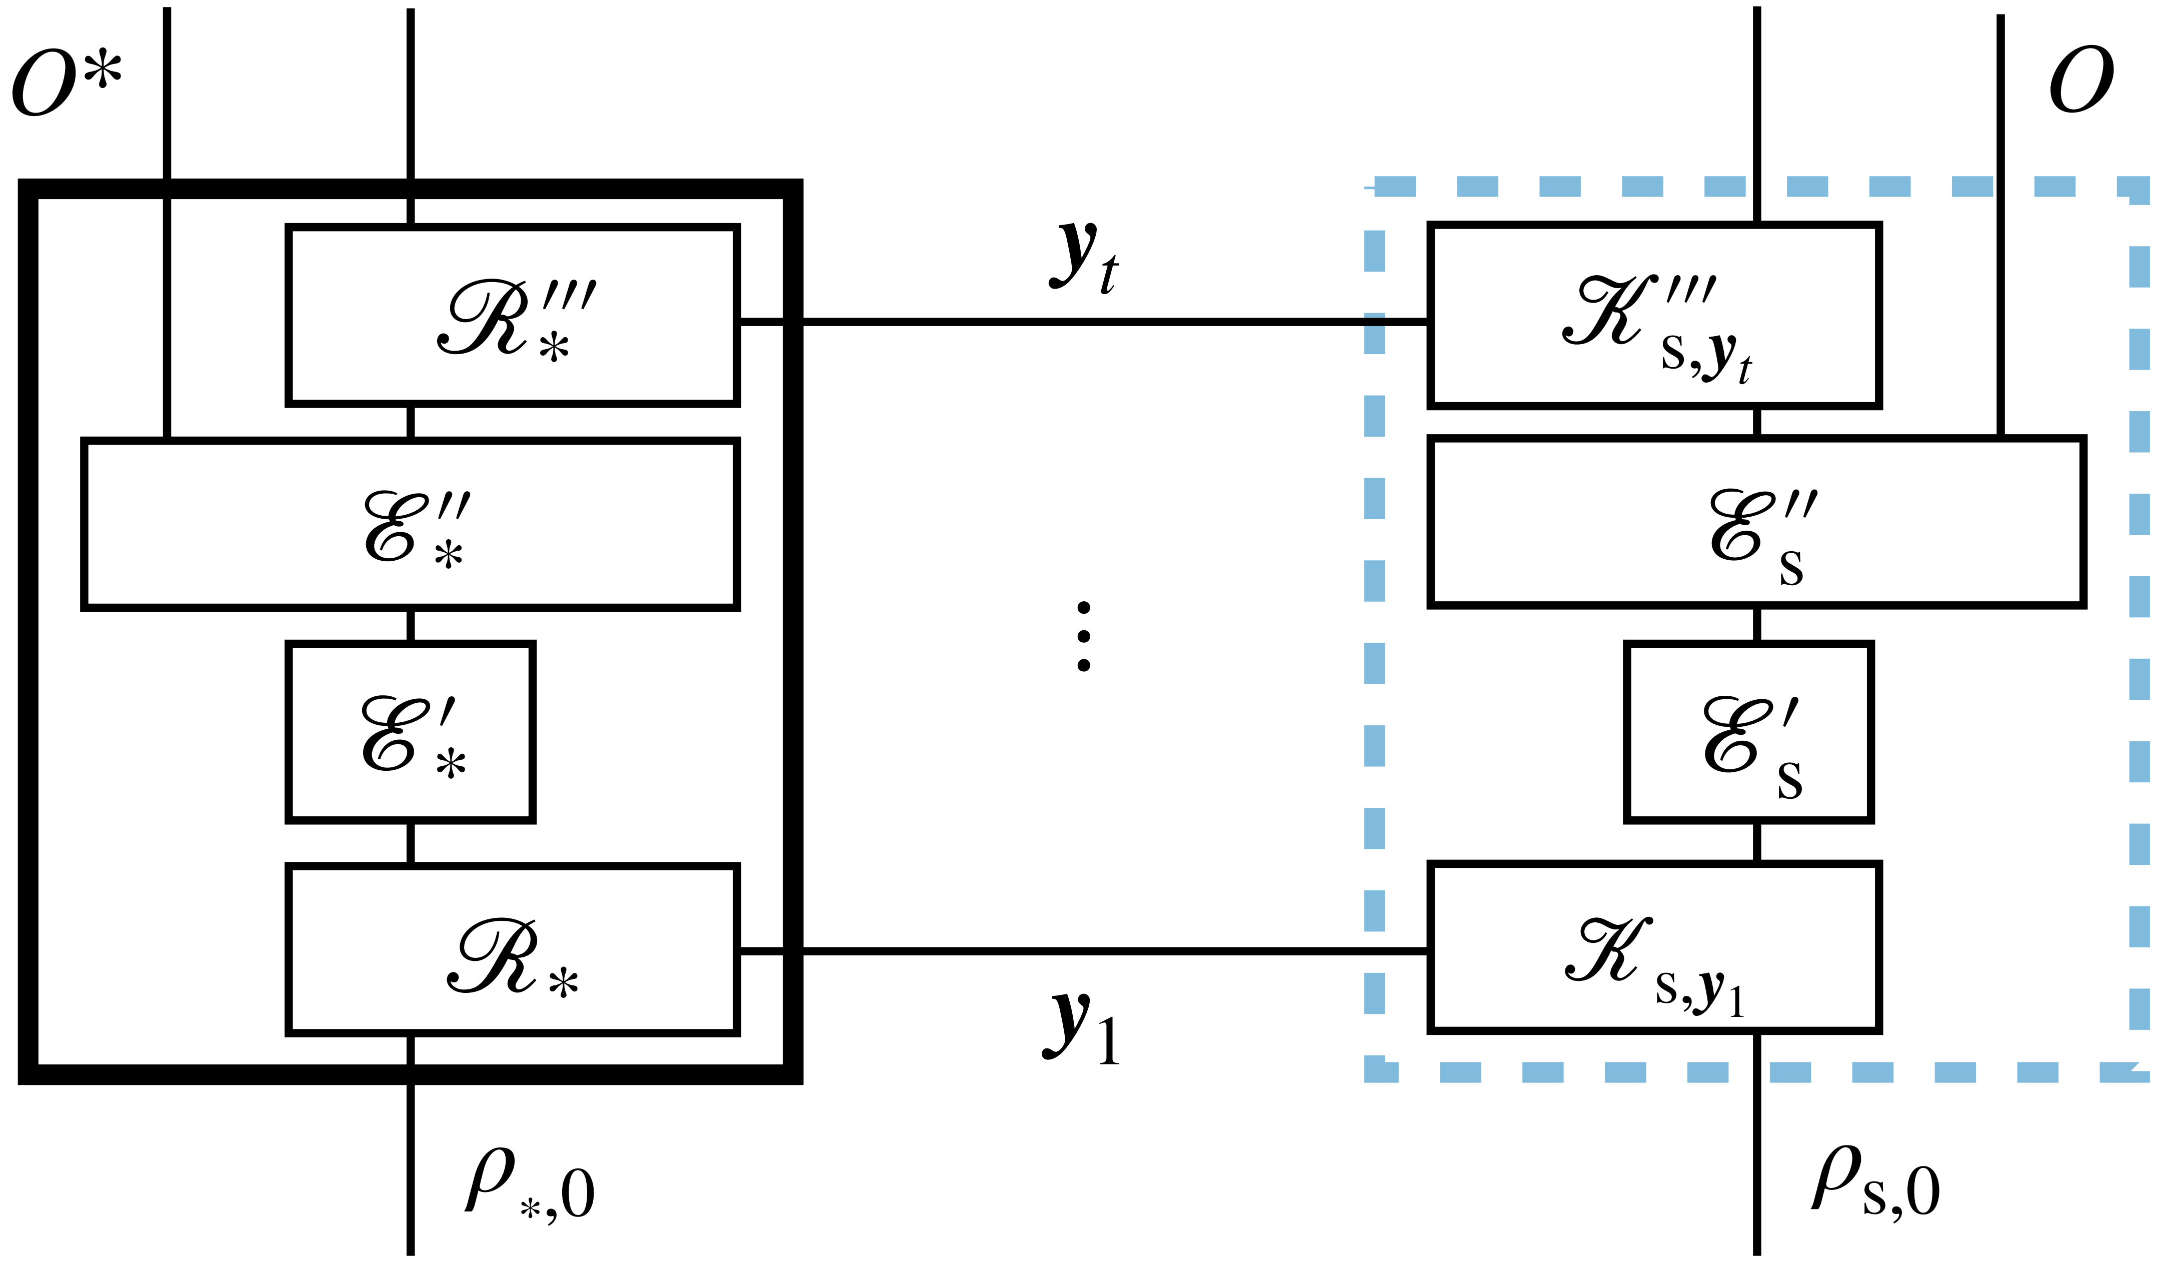


- If (reduced) density matrix is diagonal, then retrieve HMMs $\Rightarrow$ subset of qHMMs

# Haar-random unitaries

- Haar-random unitary gate $u$ can be thought of as a channel

$$\rho \to \int d\mu(u)\ u\rho u^\dagger \otimes |u\rangle\langle u|$$

- Can have block diagonal structure

$$\text{ex. } u_{\mathrm{U}(1)} = u^{(-1)} \oplus u^{(0)} \oplus u^{(1)} = \begin{pmatrix} u^{(-1)} & & \\ & u^{(0)} & \\ & & u^{(1)} \end{pmatrix}$$

# Haar-random unitaries

- If the choice of unitary gate is **not** revealed to the student, then can trace over classical registers, which is completely depolarising for each block

$$\rho \to \int d\mu(u)\ u\rho u^\dagger = (\mathscr{D} \oplus \cdots \oplus \mathscr{D})[\rho]$$

- ex. for $u_{\mathrm{U}(1)}$, we get $\mathscr{E}_{\mathrm{ssep}}$ which acts on diagonal elements of the density matrix as a "simple-symmetric exclusion process",

$$\begin{pmatrix} 1 & & & \\ & 1/2 & 1/2 & \\ & 1/2 & 1/2 & \\ & & & 1 \end{pmatrix}.$$

# Haar-random unitaries

- For a setup with **hidden** (block-diagonal) Haar-random unitaries and measurements on observables diagonal in this basis, the student can do optimal inference using a classical HMM

- Consider local Hilbert space of qubits and qu**d**its $\mathscr{H}_{\mathrm{loc}} = \mathbb{C}^2 \otimes \mathbb{C}^d$.
  We prove that even with **revealed** unitary gates, in the $d \gg \mathrm{e}^{LT}$ limit, the student can still do optimal inference with a classical HMM

- Haar unitary $\leftrightarrow$ Markov kernel
- Quantum measurement $\leftrightarrow$ classical likelihood/measurement model

# Outline

1. Bayesian inference and hidden Markov models
2. Teacher-student scenario and Bayes optimality
3. General quantum inference problem and quantum hidden Markov models
4. **Quantum error correction and the random-bond Ising model**
5. The Planted SSEP and the planted XOR
6. Discussion of the phase diagrams

# Quantum error correction (QEC) and random-bond Ising model (RBIM)

- Teacher applies error channel $\mathscr{E}$ applies bit-flips with some rate $\pi_* = e^{-\beta_*}/2\cosh\beta_*$

- Record current bit-flip status in 'environment' register $|\boldsymbol{f}_*\rangle\langle\boldsymbol{f}_*|, f_l = \pm 1$

$$\rho_* \to \pi_*(X \otimes X_{\text{env}})\rho_*(X \otimes X_{\text{env}}) + (1-\pi_*)\rho_*$$

# Quantum error correction (QEC) and random-bond Ising model (RBIM)

- Measure syndrome $s^*$ and record measurement outcome

$$\rho_* \to \sum_{s^*} \mathbb{P}_{s^*} \rho_* \mathbb{P}_{s^*} \otimes | s^* \rangle \langle s^* |$$

- Possibly bit-flip errors on the syndrome measurement $s^* \to s$

- Goal of the student: Infer $\boldsymbol{f}_*$ from $S = (s_{l,t})_{l \in 1:L}^{t \in 1:T}$

- Student assumes $\pi = e^{-\beta} / 2 \cosh \beta$

# Quantum error correction (QEC) and random-bond Ising model (RBIM)

- Natural order parameter is $\mathbb{E}[(f_l - f^*_{l,t})^2]$

- Writing the student's guess as $f_l = \sigma_v \sigma_{v'} f^*_l$, above becomes a ferromagnetic order parameter $\mathbb{E}[\langle \sigma_v \sigma_{v'} \rangle]$ of an RBIM

- $\beta \leftrightarrow 1/T, p = e^{-\beta_*}/2\cosh\beta_*$

- Sometimes known as the 'planted Ising model'

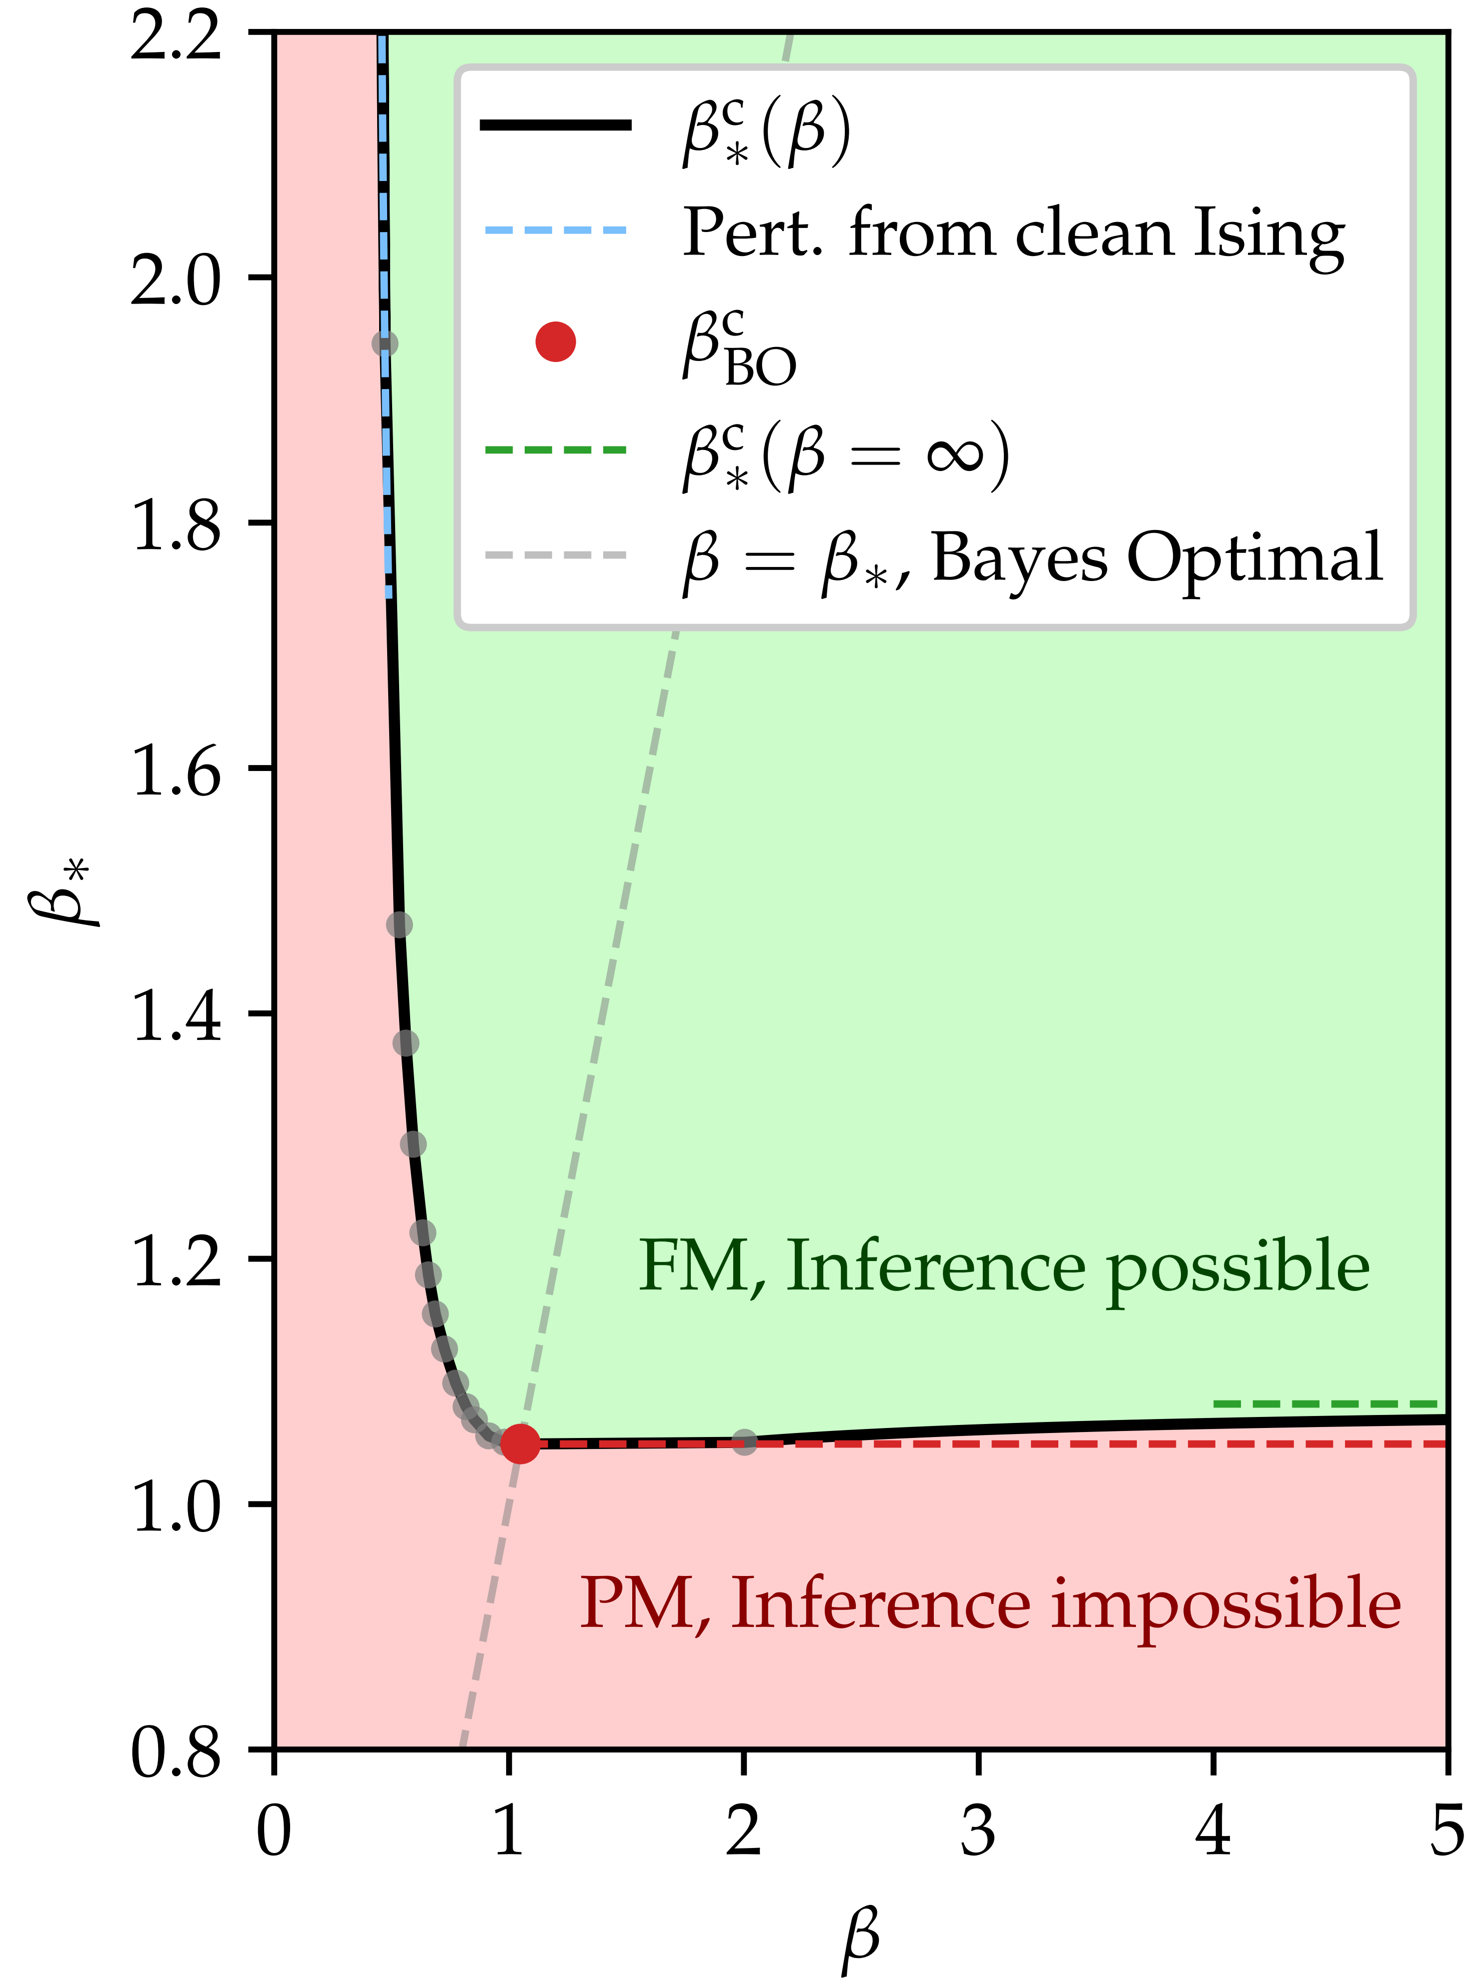

# Quantum error correction (QEC) and random-bond Ising model (RBIM)

- Plot in $\beta - \beta_*$ space → Nishimori condition is diagonal line
- $\beta_* \to \infty$: clean Ising model
- $\beta \to \infty$: Minimum weight perfect matching (MWPM)
- Not pictured: student's uncertainty in their guess: $\delta f_l^2 = 1 - \mathbb{E}[\langle \sigma_v \sigma_{v'} \rangle^2]$. Confident in their **wrong** answer → Spin-glass phase

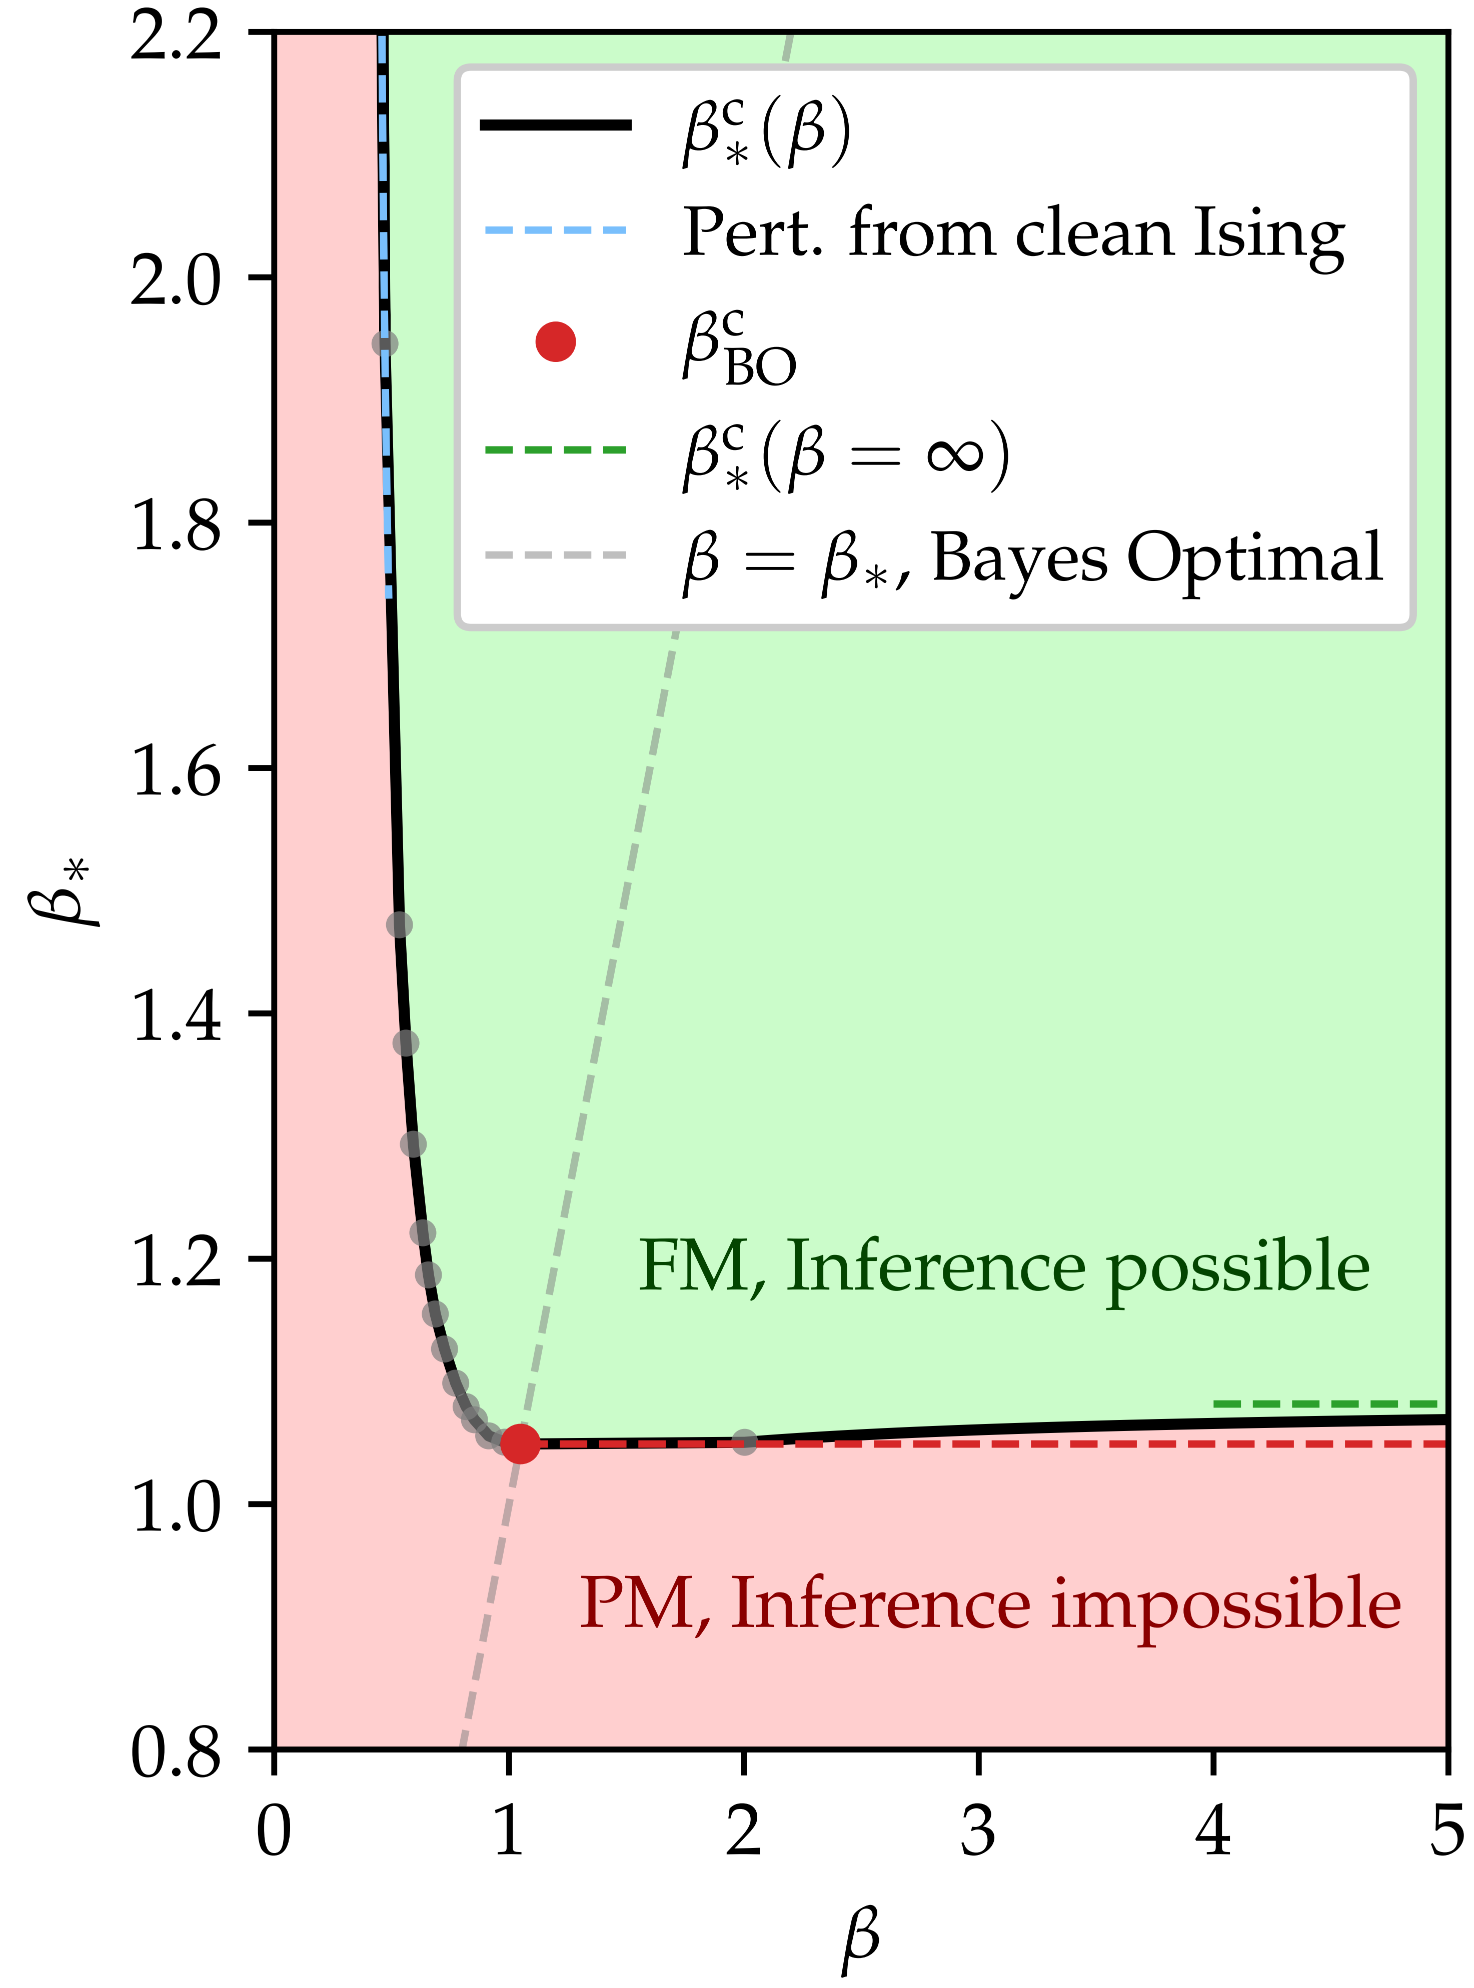

# Outline

1. Bayesian inference and hidden Markov models
2. Teacher-student scenario and Bayes optimality
3. General quantum inference problem and quantum hidden Markov models
4. Quantum error correction and the random-bond Ising model
5. **The Planted SSEP and the planted XOR**
6. Discussion of the phase diagrams

# (Quantum) Planted SSEP (Agrawal et al., 2022)

- Local Hilbert space of qubits and qudits $\mathbb{C}^2 \otimes \mathbb{C}^d$
- Teacher evolves with Haar-random $U(1)$-symmetric unitary gates in brickwork fashion
- Teacher performs on-site weak measurements at every timestep

$$Q(y) = \frac{1}{\sqrt[4]{2\pi}} \exp\left( -\frac{(y - \epsilon_* Z)^2}{4} \right)$$

- Teacher measure total charge $C_*$ at the end
- As a standalone circuit: charge-sharpening $\delta C_*^2$
- Goal of the student: infer $C$ from $Y = (y_{x,t})_{x \in 1:L}^{t \in 1:T} \Rightarrow \mathrm{MSE}(C)$

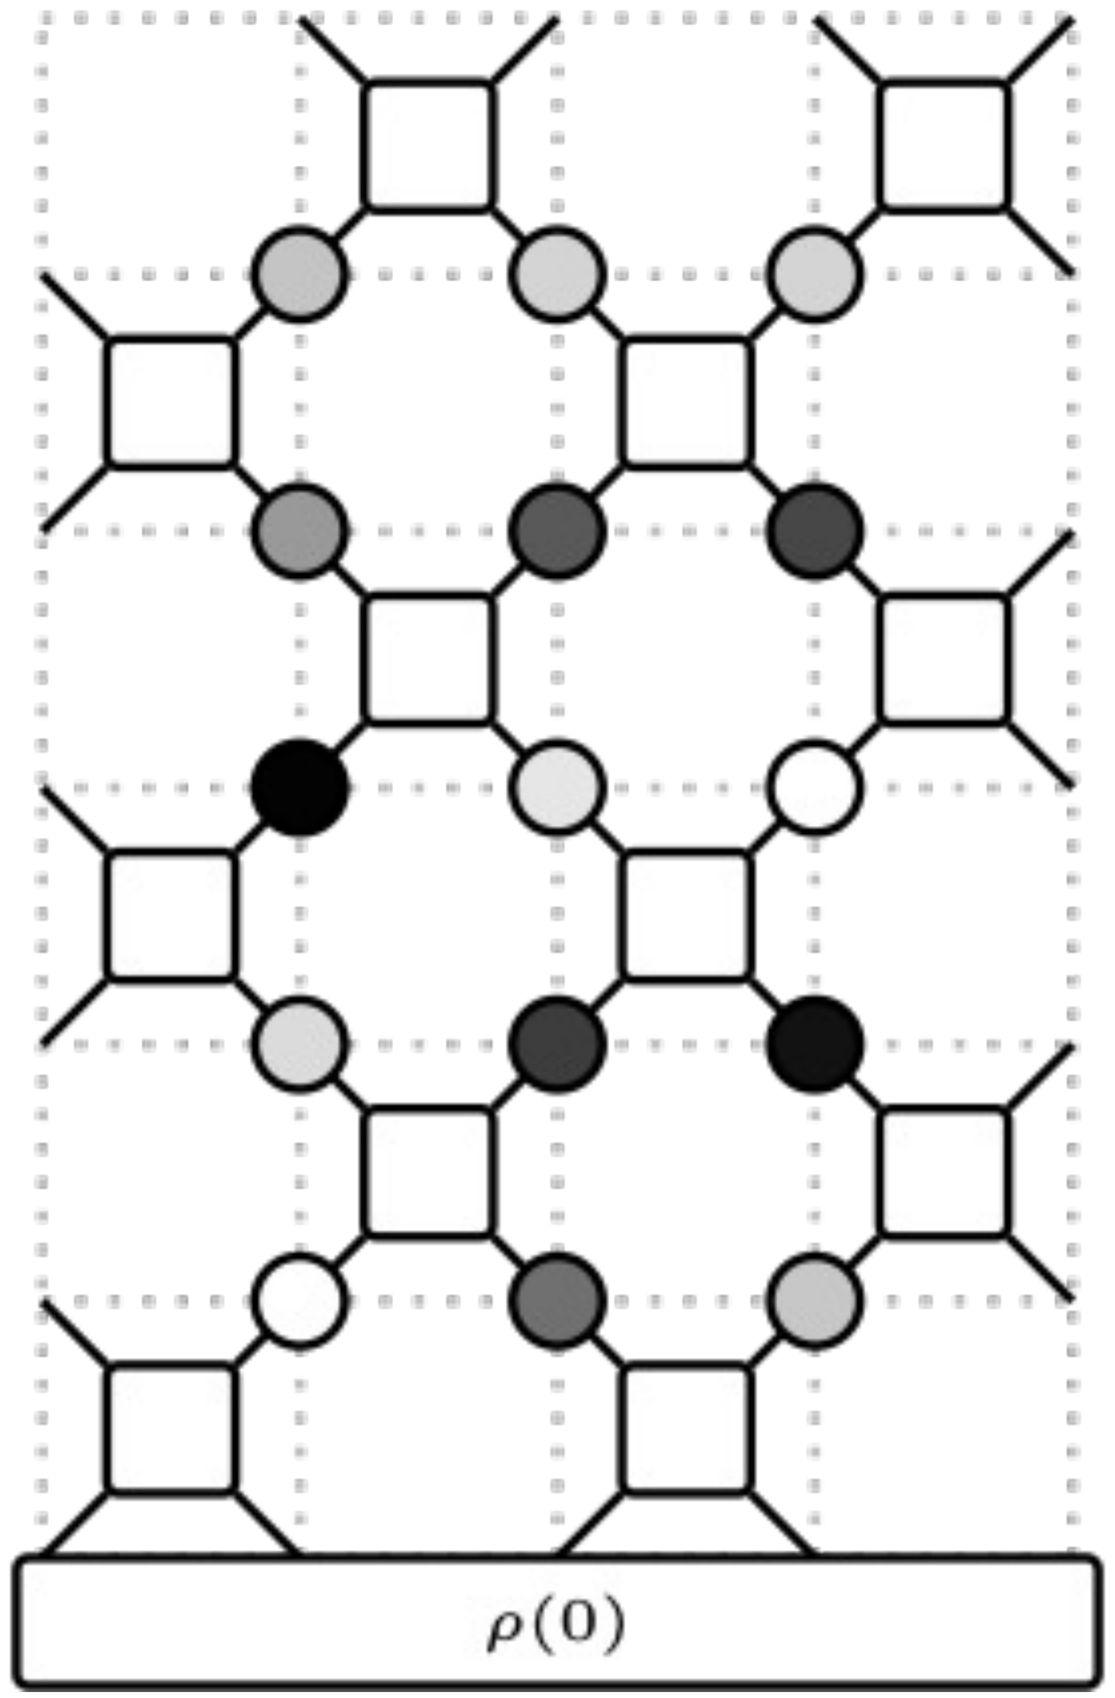

# Planted SSEP

- In the limit $d \gg \mathrm{e}^{LT}$, student can conduct optimal inference with assuming a classical HMM of inferring a SSEP from noisy images:

$$p(s_t, s_t' | s_{t-1}, s_{t-1}) = (s_t, s_t' | \begin{pmatrix} 1 & & & \\ & 1/2 & 1/2 & \\ & 1/2 & 1/2 & \\ & & & 1 \end{pmatrix} | s_{t-1}, s_{t-1}')$$

$$p_*(y_{x,t} | s_{x,t}^*) = \frac{1}{\sqrt{2\pi}} \exp\left( -\frac{(y_{x,t} - \epsilon_* s_{x,t}^*)^2}{2} \right)$$

- Generalisation of inferring a random walk (SWPK, Lamacraft, 2022)

# Planted SSEP

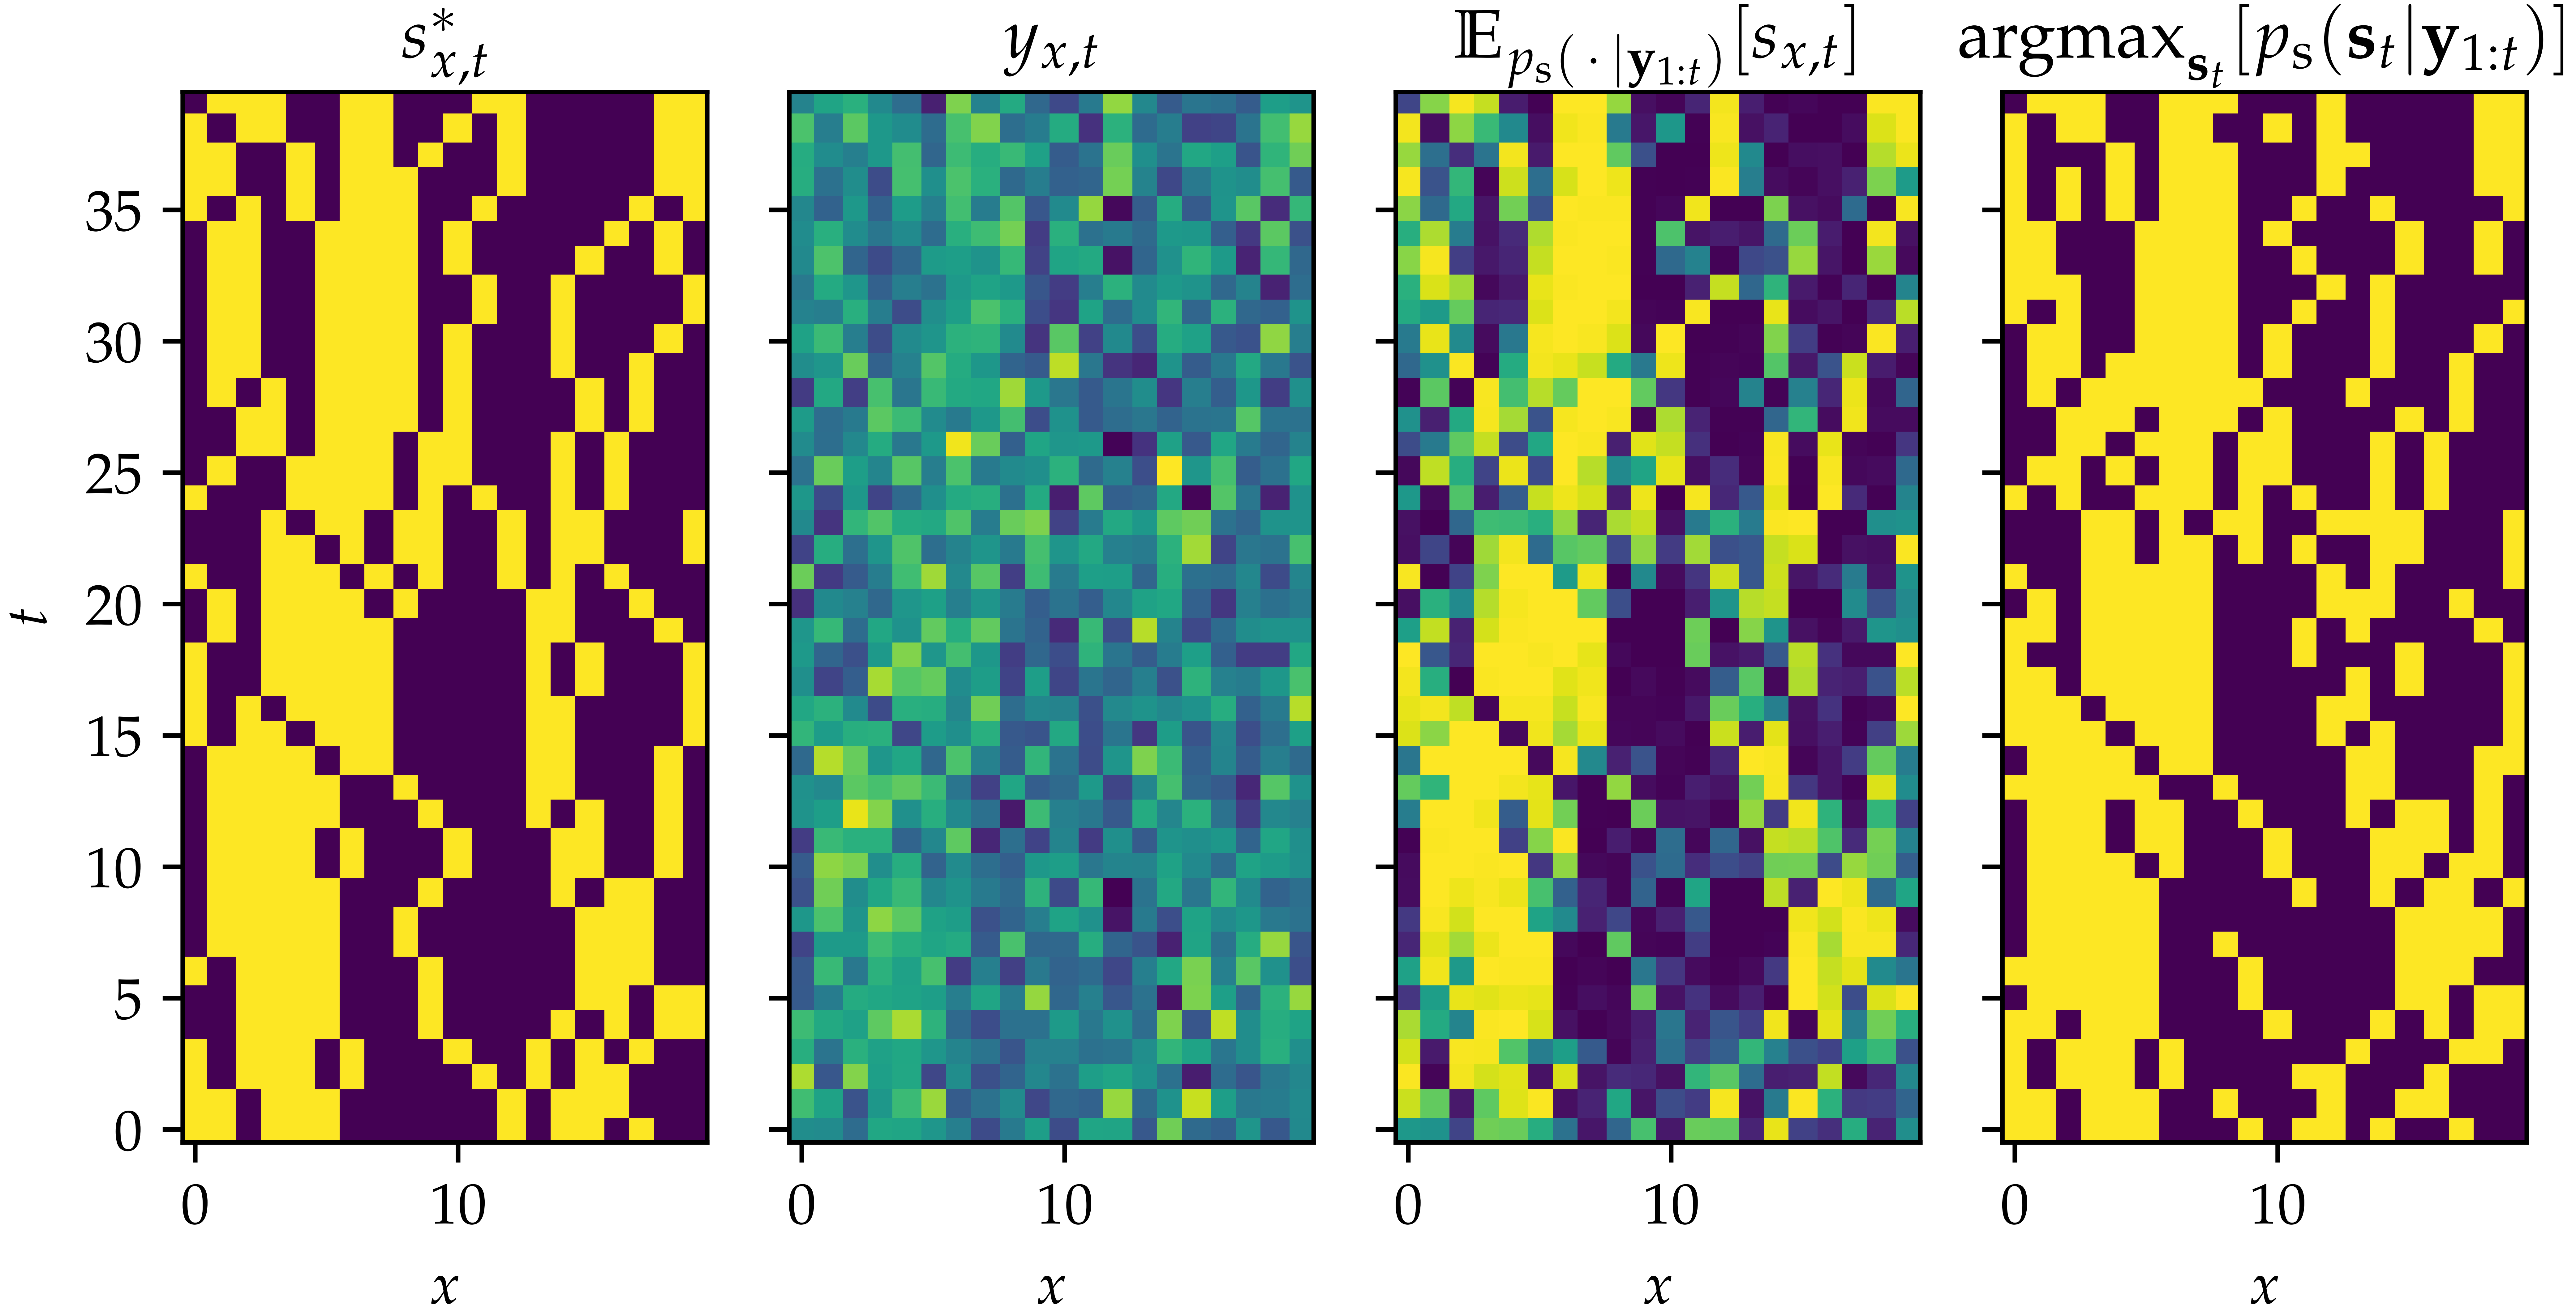

# Phase diagram of the planted SSEP

- Go beyond existing work by expanding the phase diagram into the $\epsilon - \epsilon_*$ parameter space
- Follow (Barratt et al. 2022) to develop replica field theory, which predicts the same universality class of phase boundary
- Sharpening and learnability phases coincide for this model

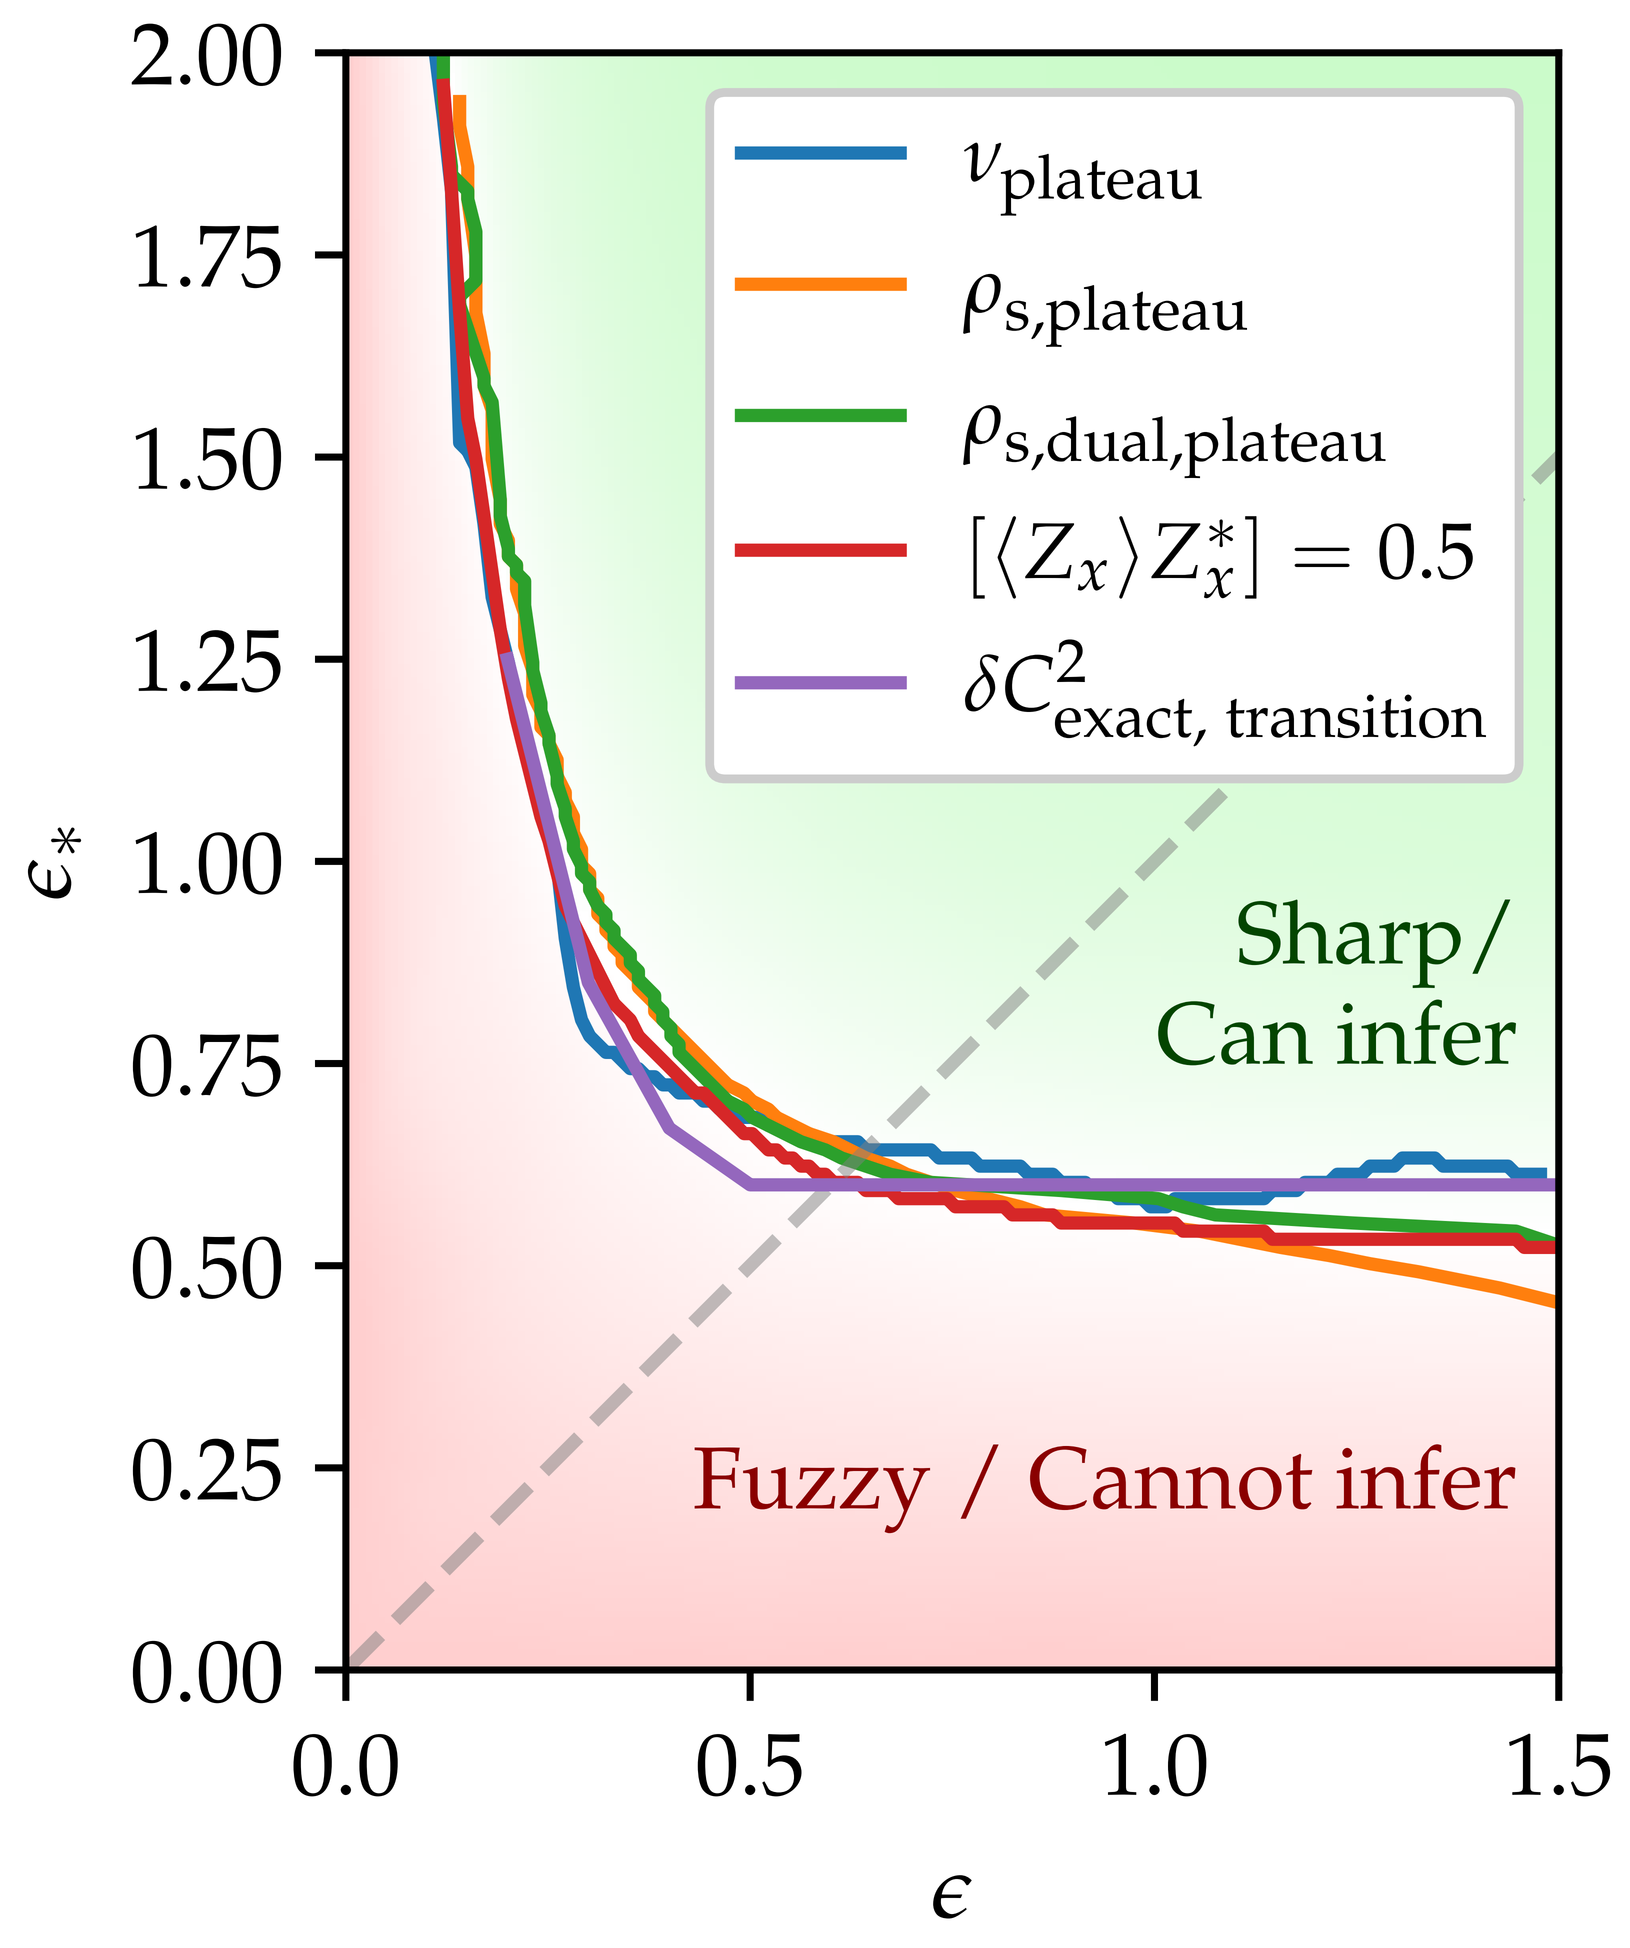

# Phase diagram of the planted SSEP

- Both replica field theory and perturbative expansion predict that the disconnected correlator should change sign across the Bayes optimal line in the fuzzy phase

$$\mathbb{E}[\langle s_x s_0 \rangle] \sim \left( \frac{\epsilon_*^2}{\epsilon^2} - 1 \right)$$

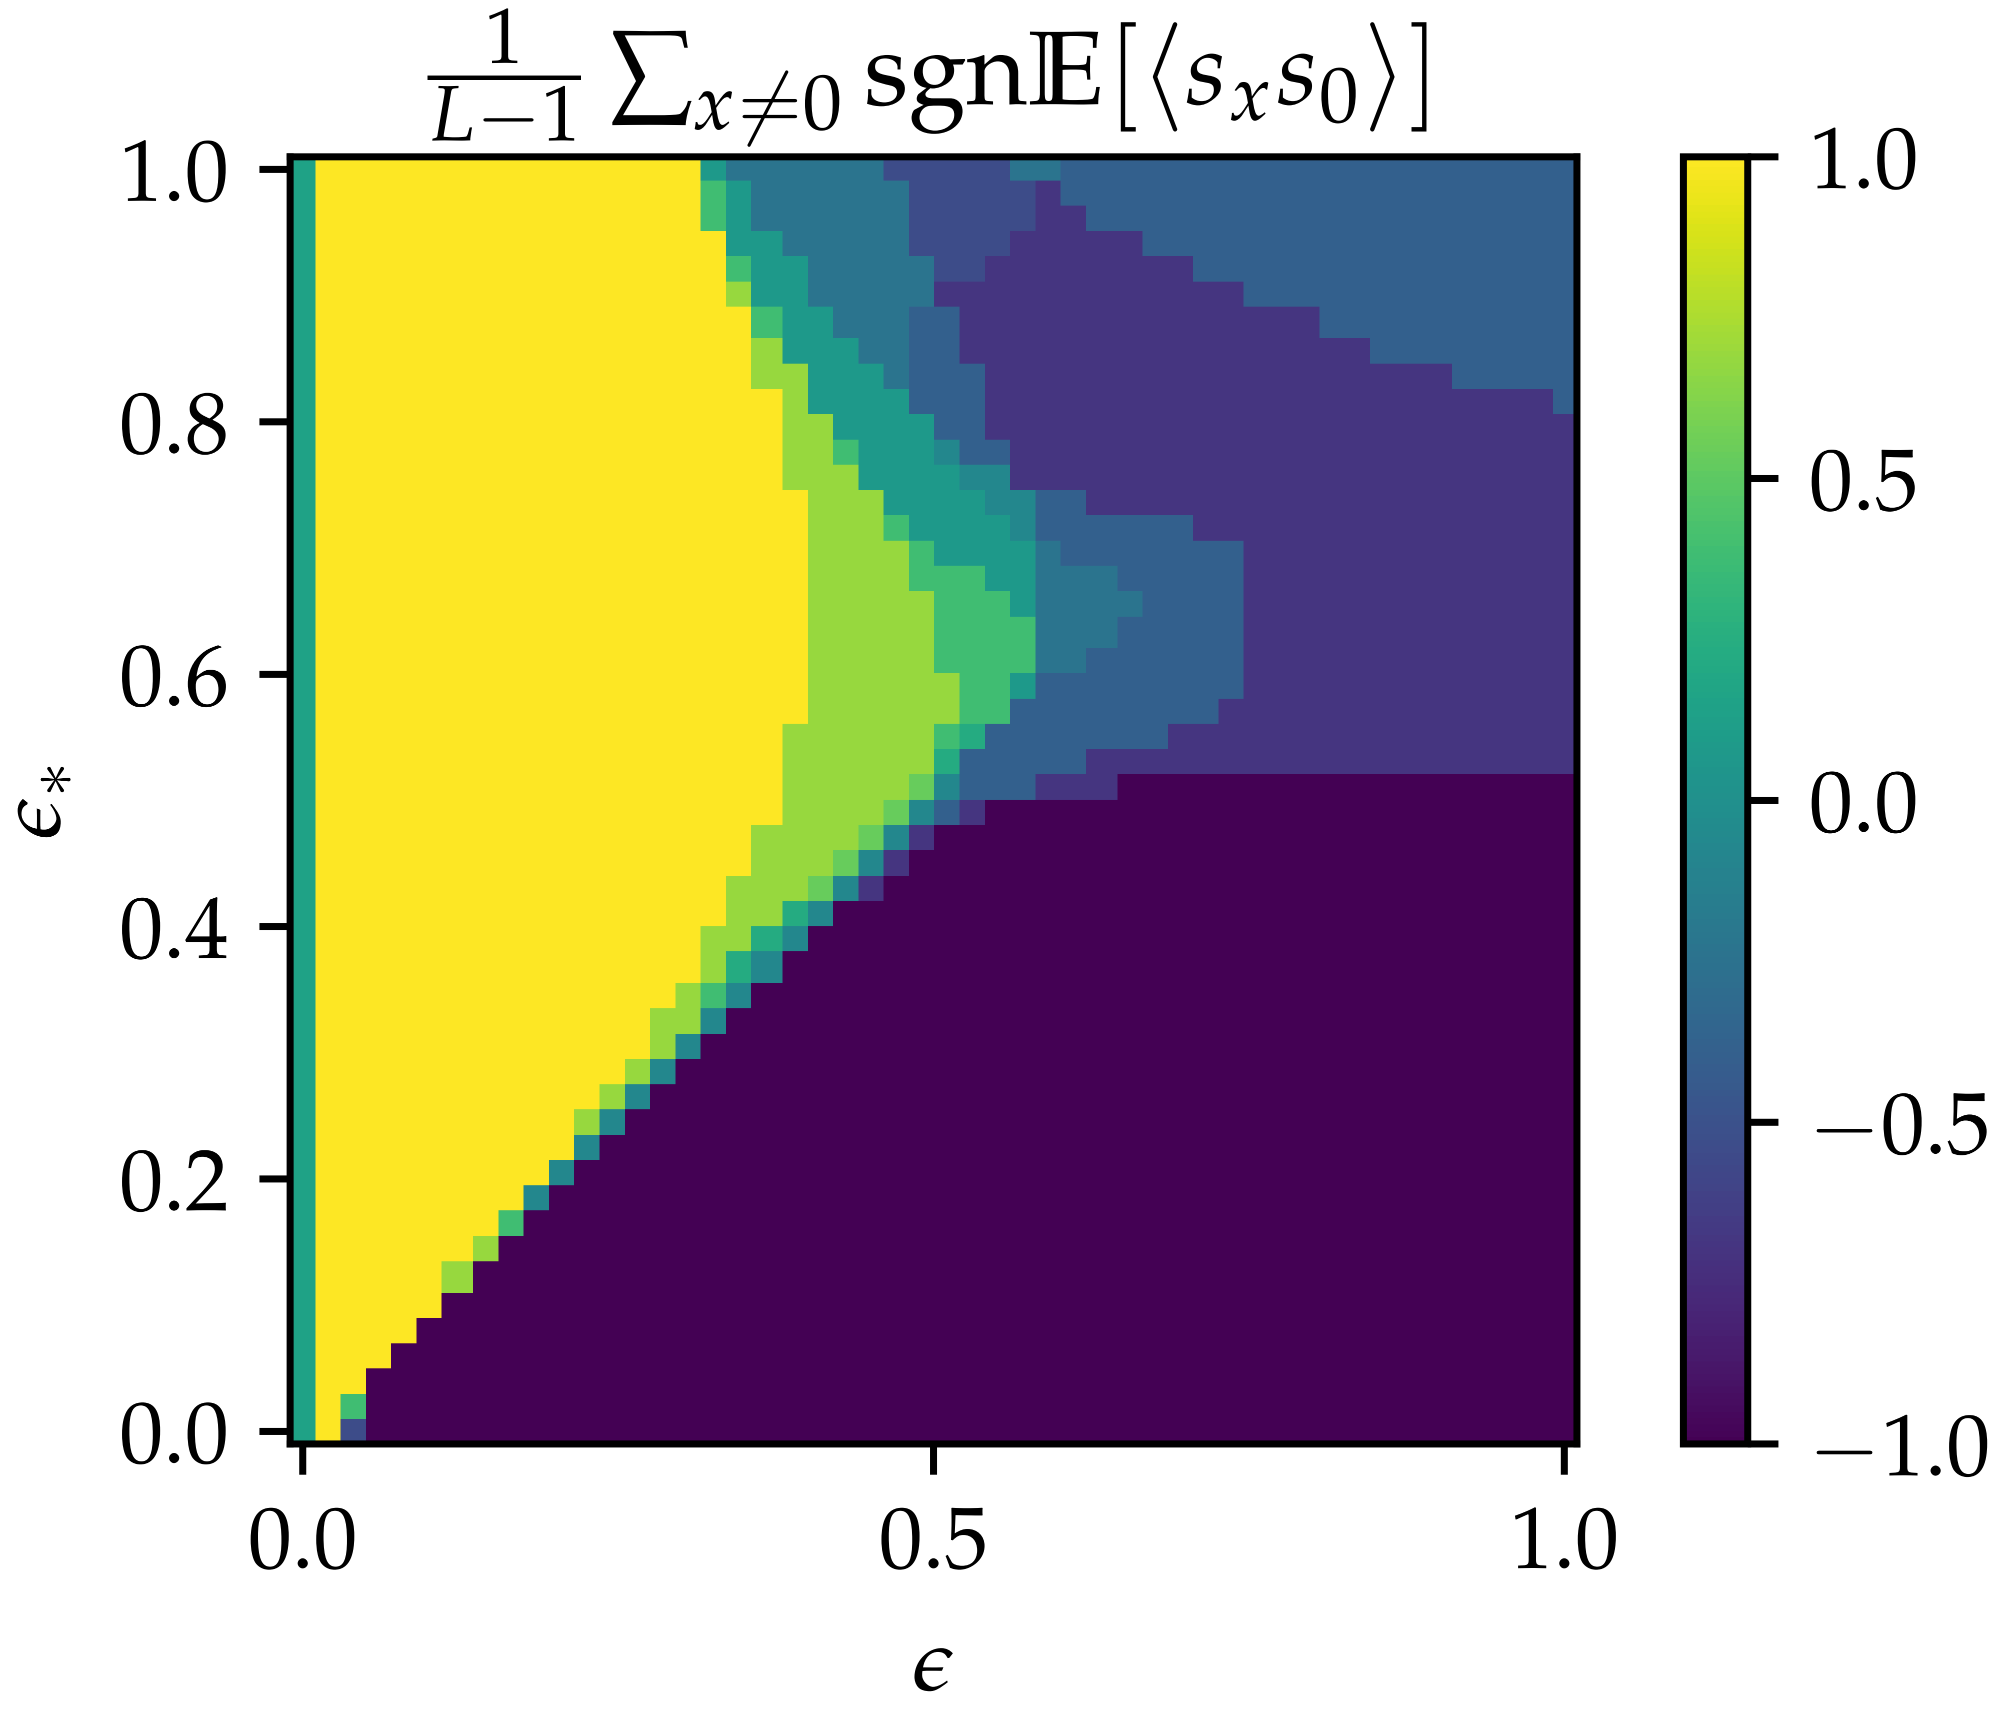

# Planted XOR

- Teacher randomly picks Booleans $b_{1:2^N,1}$ then performs XOR operations
- Produces "images" of the Booleans $y \sim \mathcal{N}(\epsilon_* b, \sigma_*^2)$
- Goal of the student: deduce the final bit

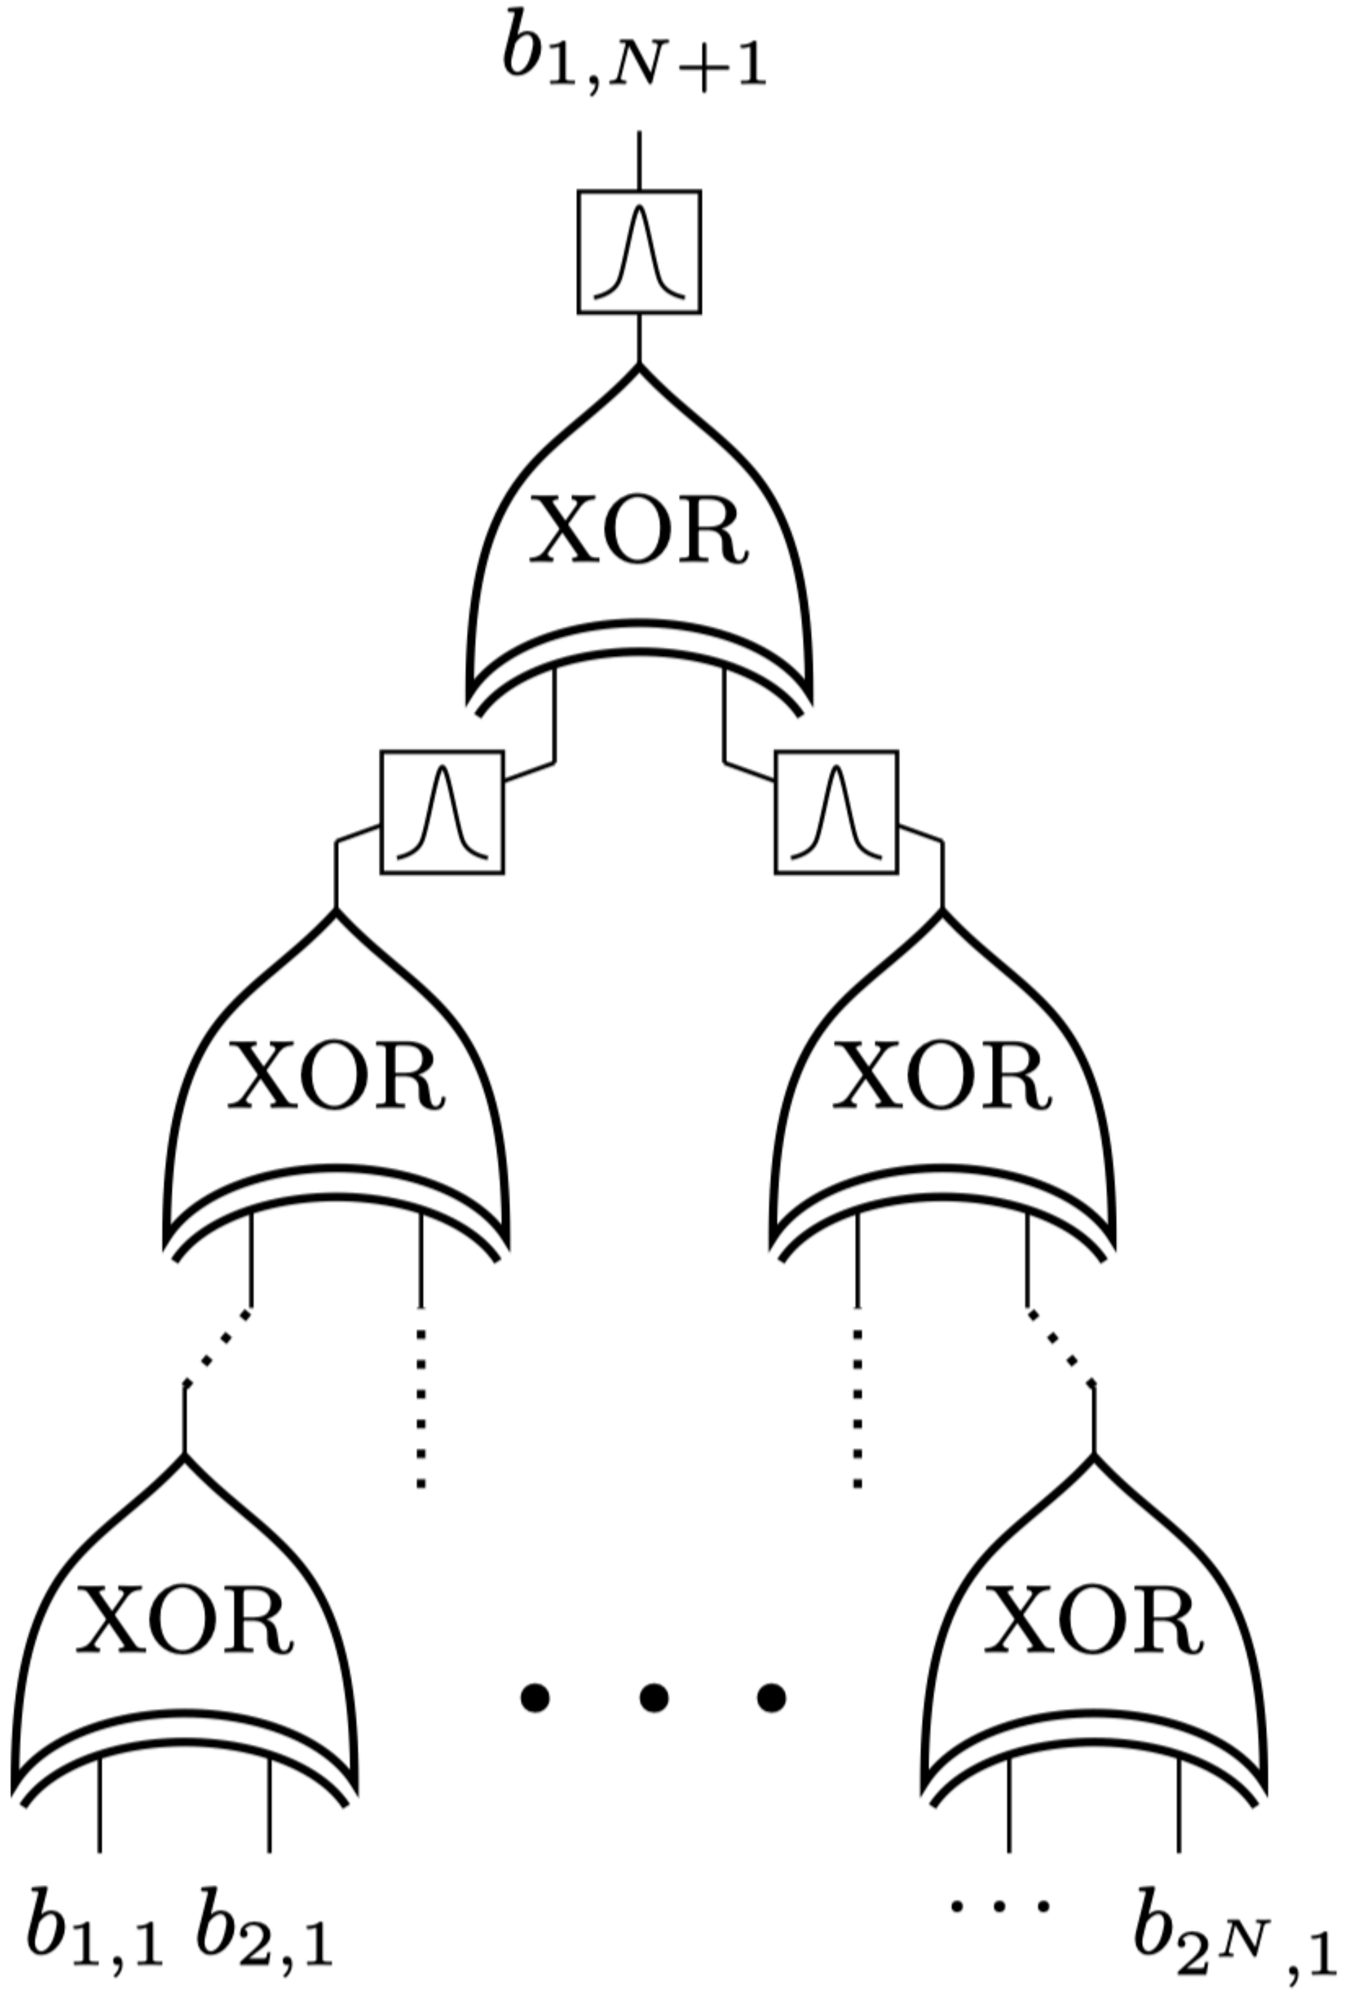

# (Quantum) Planted XOR

- Note that applying a SSEP gate + measuring one of the outputs is like picking XOR or NOTXOR and telling the student which one you picked

- Produce a quantum model by replacing the XOR gates with $u_{\mathrm{U}(1)}$ with one projective measurement, and Gaussian measurements with corresponding weak measurement

- Teacher picks from an ensemble of computational basis states

- Qudit dimension $d = 1$ is "maximally quantum"

- Dilute projective measurement version studied in (Feng et al. 2024)

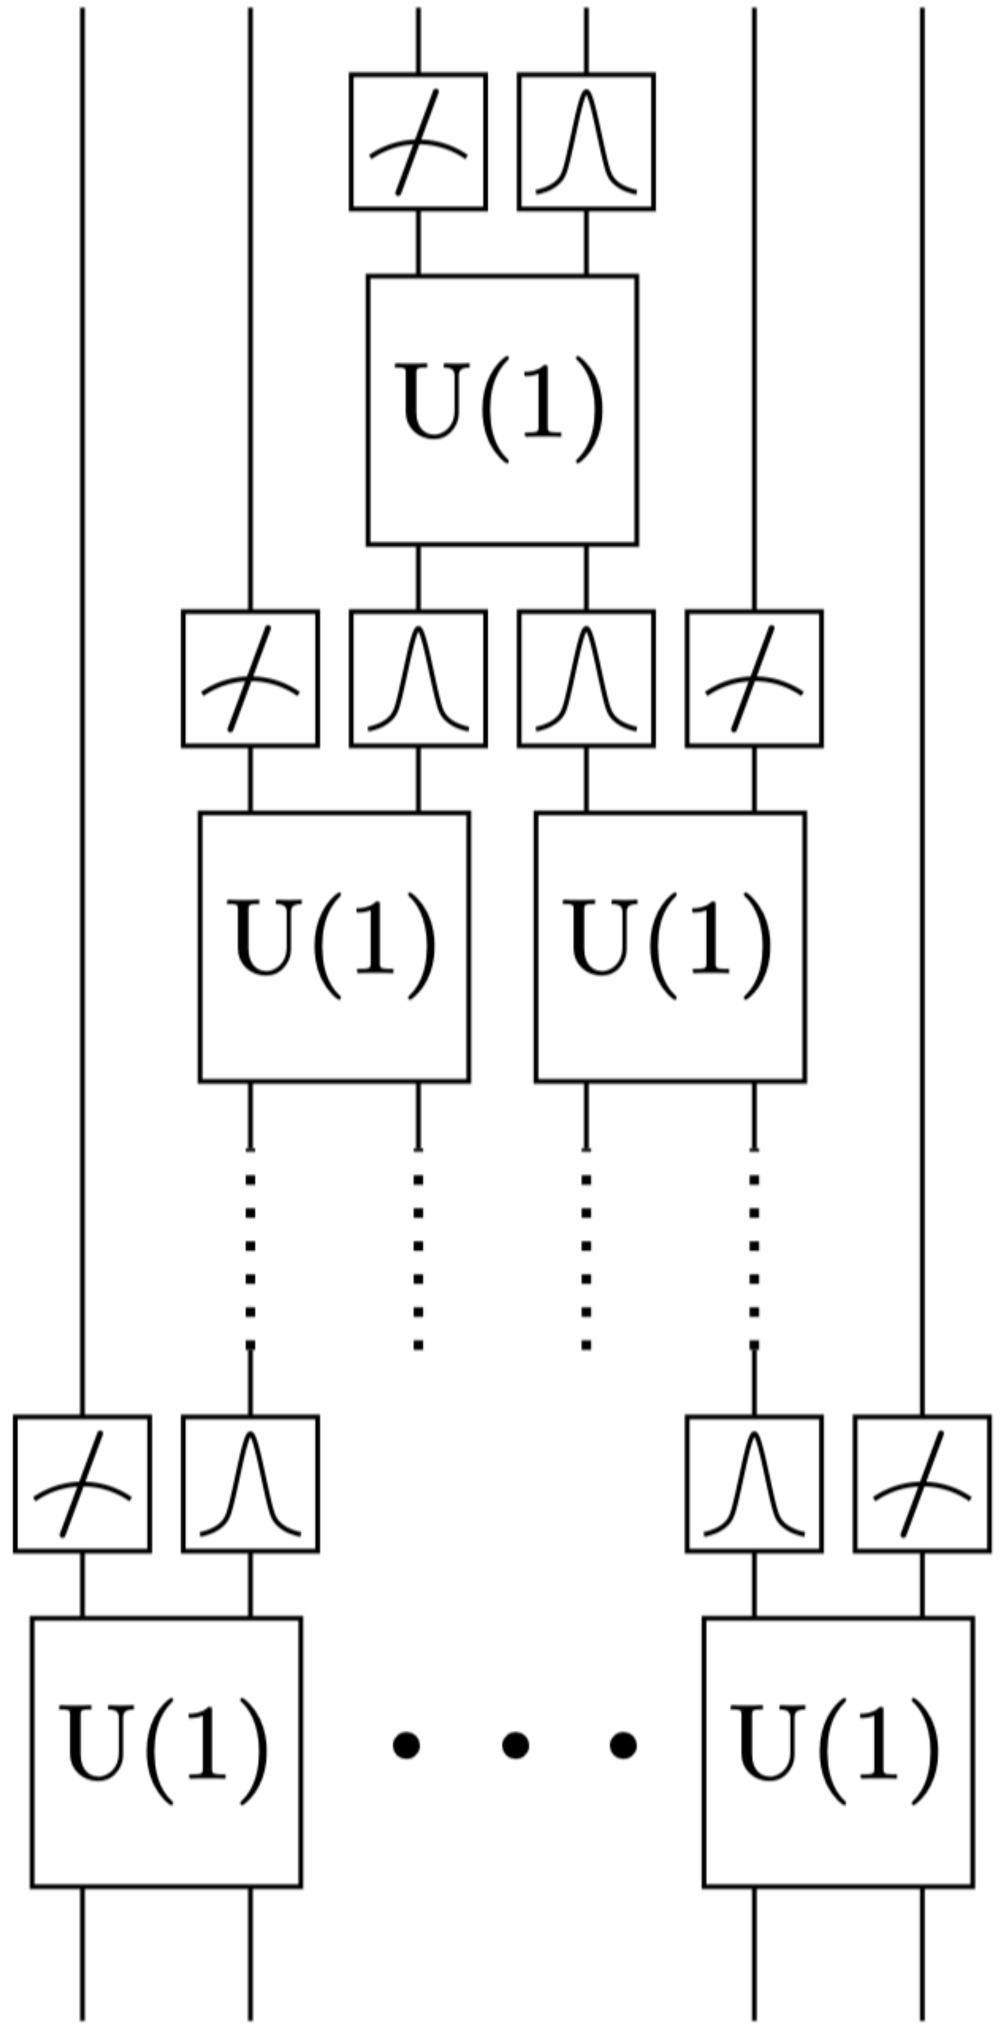

# Solving the planted XOR

- $\rho_t$ only depends on $\rho_{t-1}^{\mathrm{left}}$ and $\rho_{t-1}^{\mathrm{right}}$ which are independent of each other

- Exploit travelling wave approach originally used to study the directed polymer (Derrida, Spohn, 1988) and that near the transition, can linearise the evolution of the posterior/density matrix (Feng, Nahum, Skinner 2022)

# Phase diagram of quantum and classical planted XOR models

- Again, fuzzy/sharp phases coincide with inference possible/impossible phases

- Classical model has a larger fuzzy phase
  → *recall that this is same as the teacher hiding the sampled unitary gates, therefore less information*

- Reentrance in the phase diagram, cf. 2D RBIM

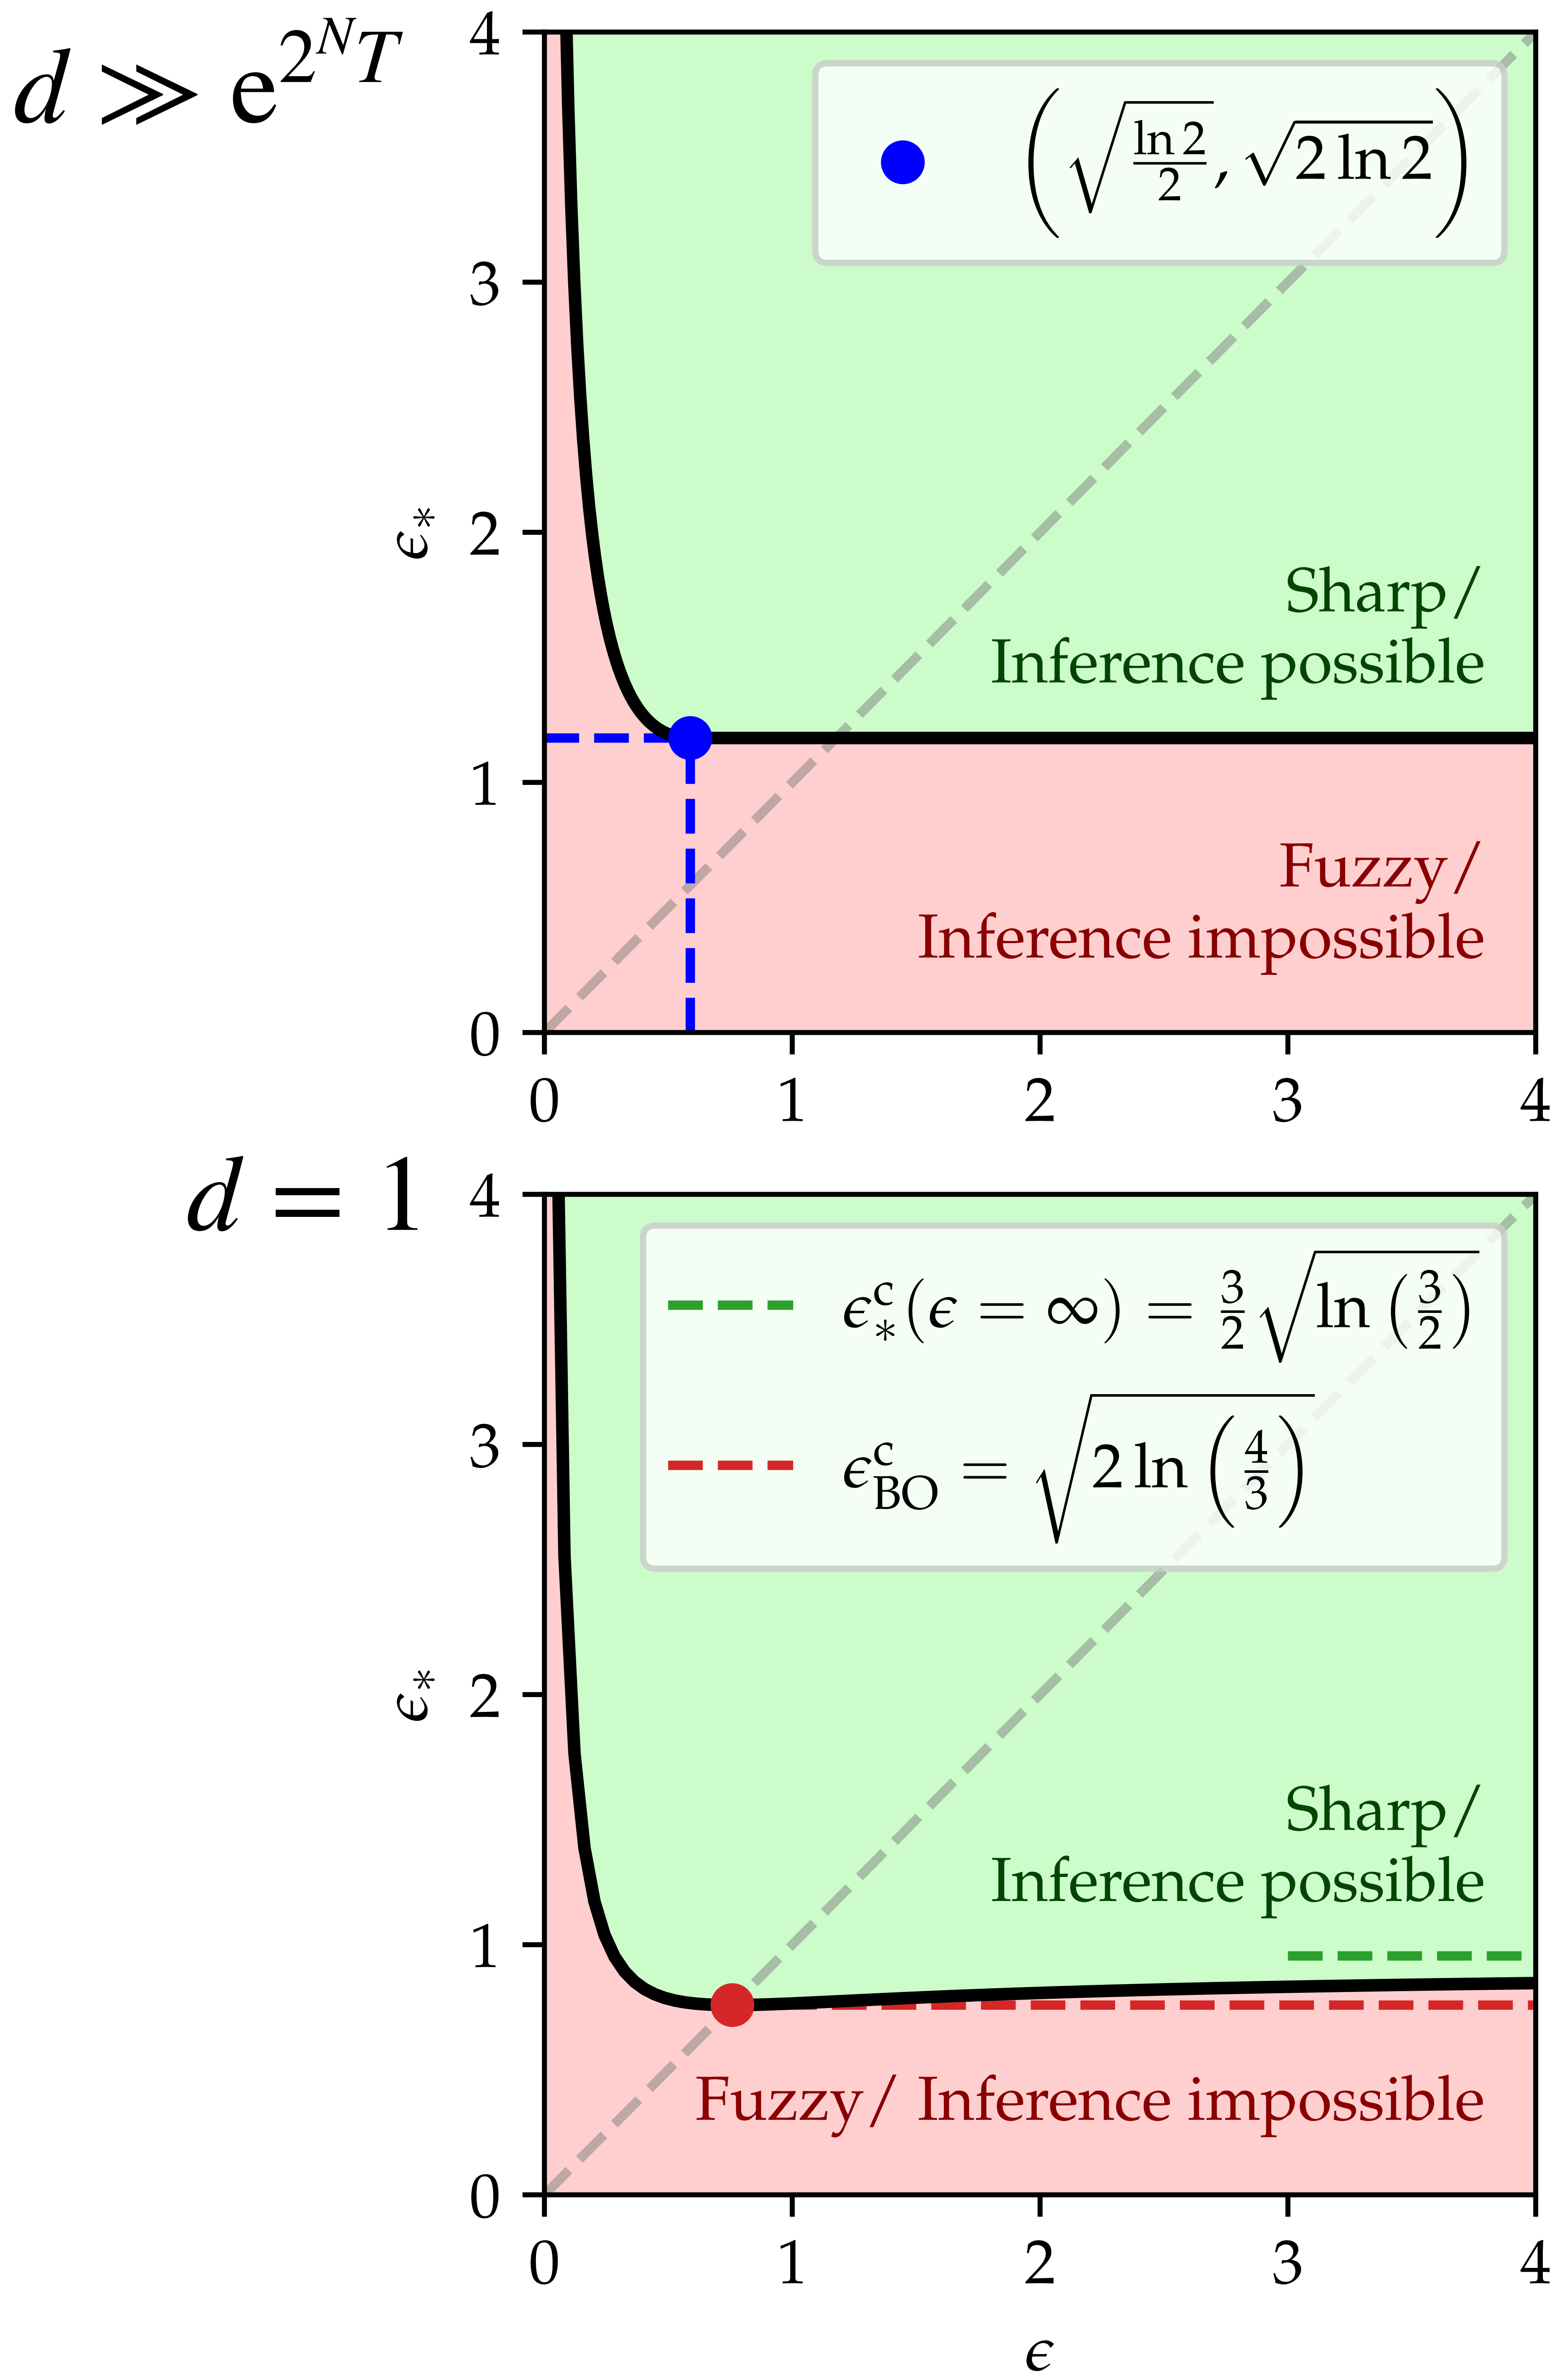

# Outline

1. Bayesian inference and hidden Markov models
2. Teacher-student scenario and Bayes optimality
3. General quantum inference problem and quantum hidden Markov models
4. Quantum error correction and the random-bond Ising model
5. The Planted SSEP and the planted XOR
6. **Discussion of the phase diagrams**

# Phase diagrams

- All models we studied had same phase boundary for fuzzy / sharp and inference impossible / possible
- This does not have to be so!
  $\delta C_{\mathrm{s}}^2 = \mathrm{MSEM} = \mathrm{MSE}/2$ **only** on the Bayes optimal line

- $\mathrm{MSE}$ is like the FM order parameter
- $\delta C_{\mathrm{s}}^2$ is like the Edwards-Anderson order parameter

- Small $\mathrm{MSE} \Rightarrow$ small $\delta C_{\mathrm{s}}^2 \leftrightarrow$ Ferromagnetic phase
- Large $\mathrm{MSE}$ but small $\delta C_{\mathrm{s}}^2 \leftrightarrow$ Spin Glass phase
- Large $\mathrm{MSE}$ and large $\delta C^2\mathrm{s} \leftrightarrow$ Paramagnetic phase

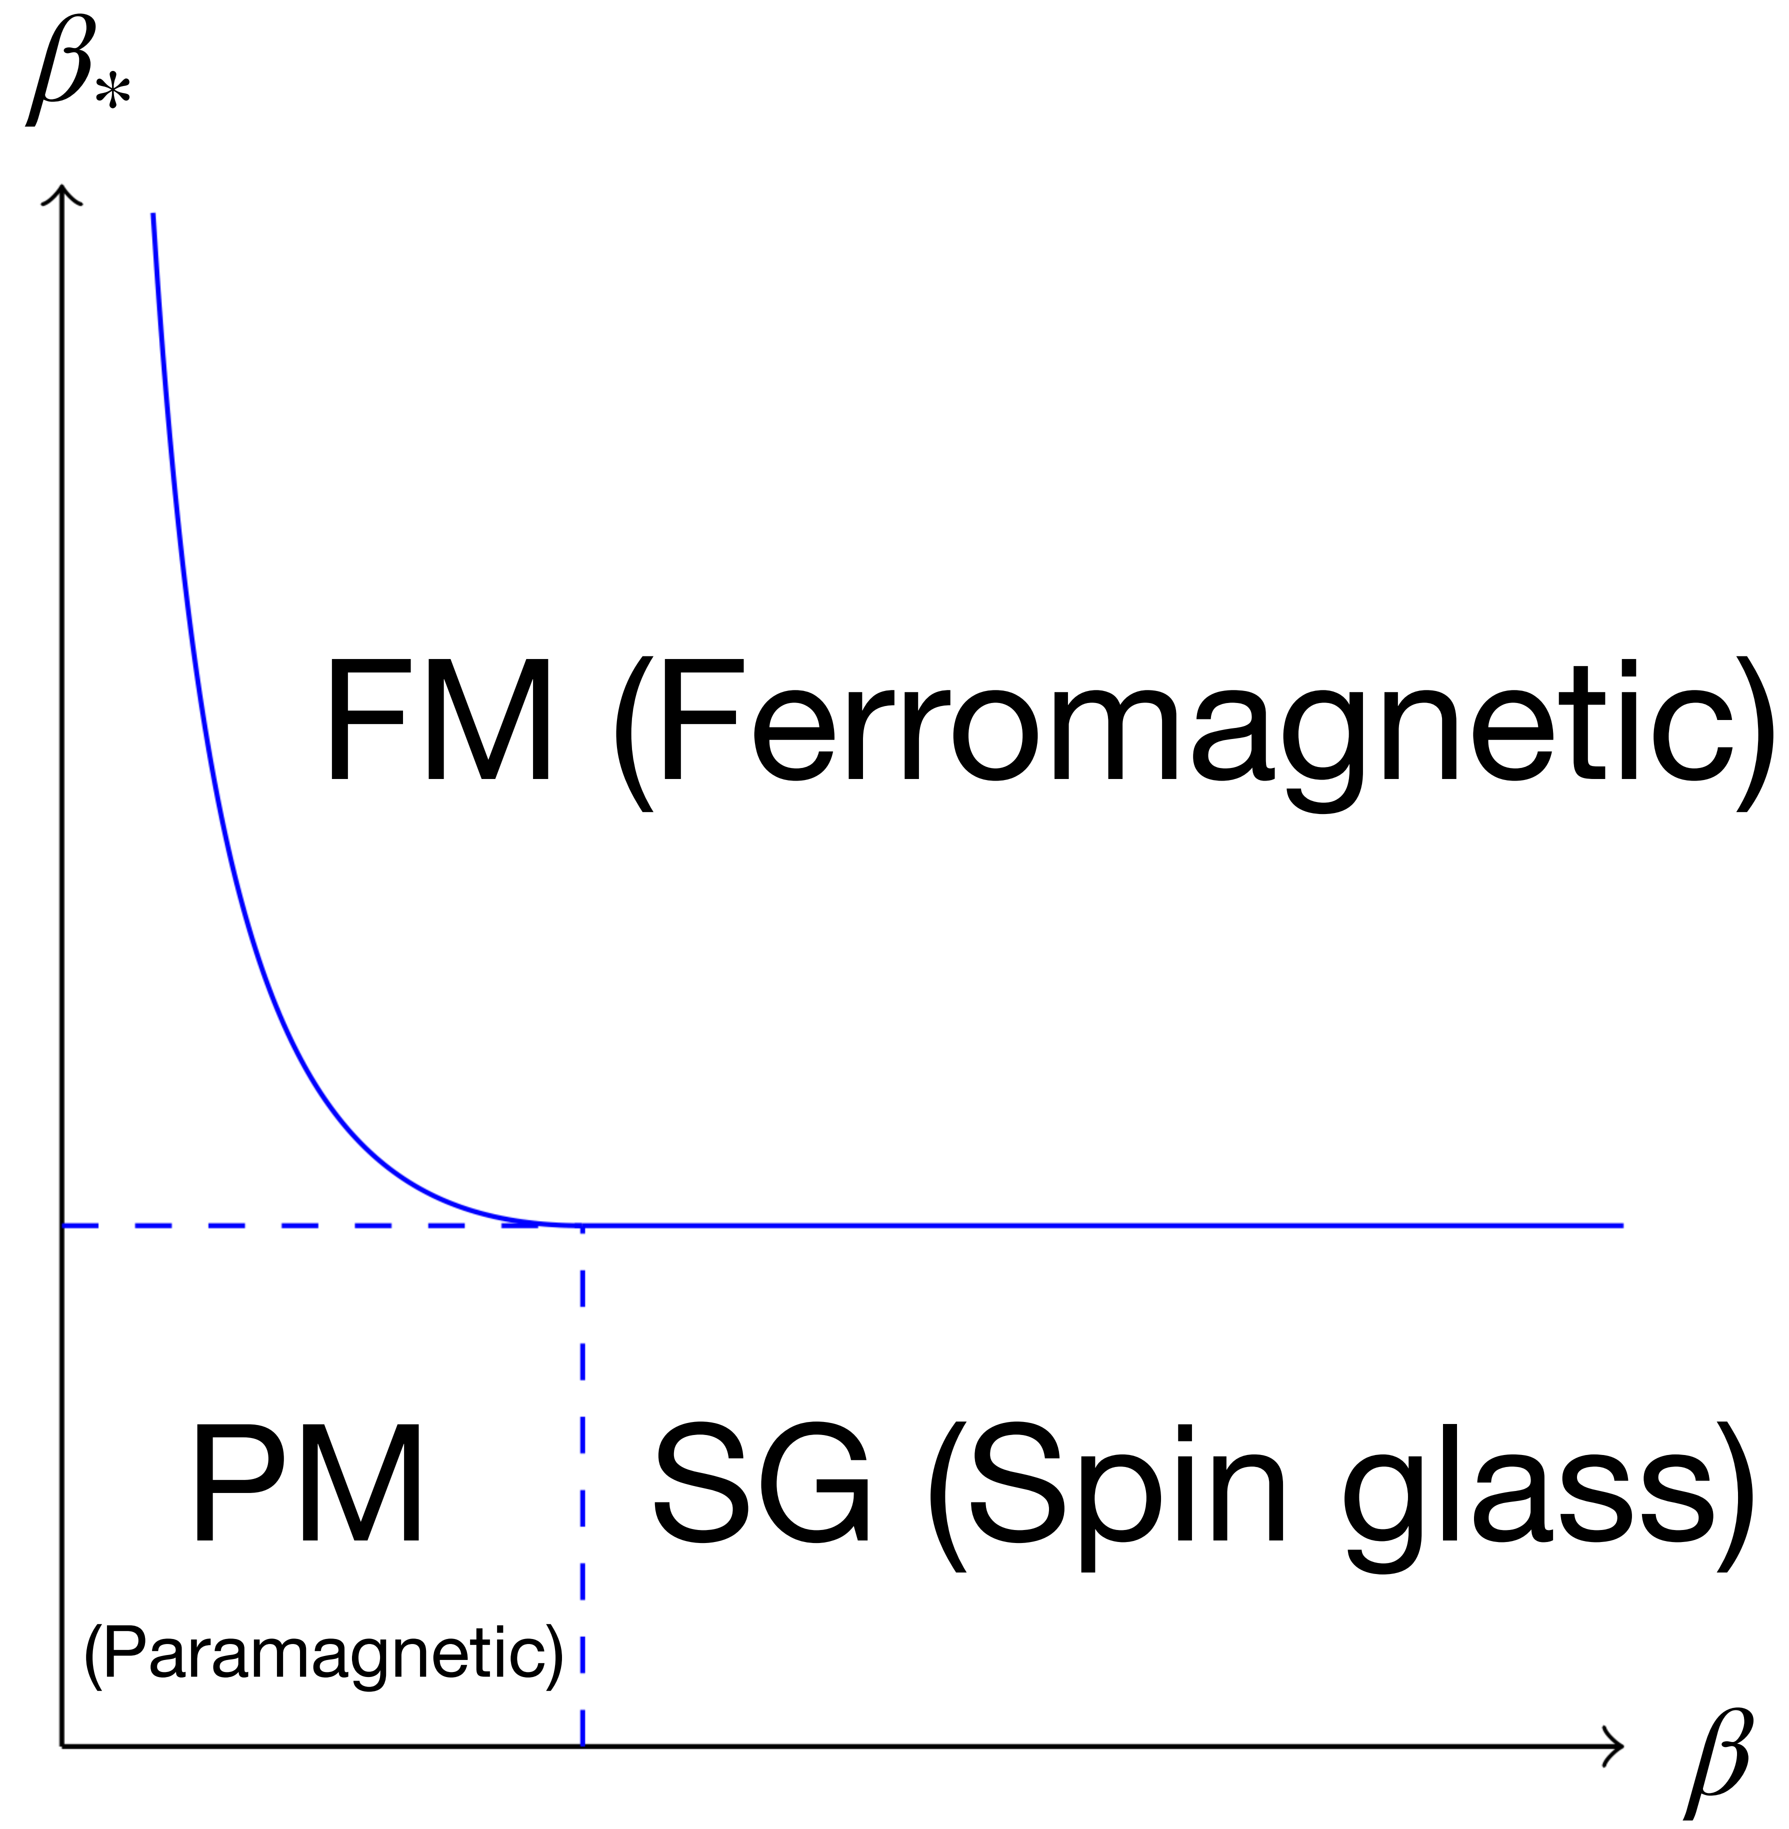

# Phase diagrams

- **Theorem**: Phase diagram for $\mathrm{MSEM}$ cannot curve down, due to the Bayes optimal estimator theorem
- For a fixed $\epsilon_*$, the best the student can do is to set $\epsilon = \epsilon_*$
- Therefore can only escape the inference possible phase to inference impossible phase, but never better

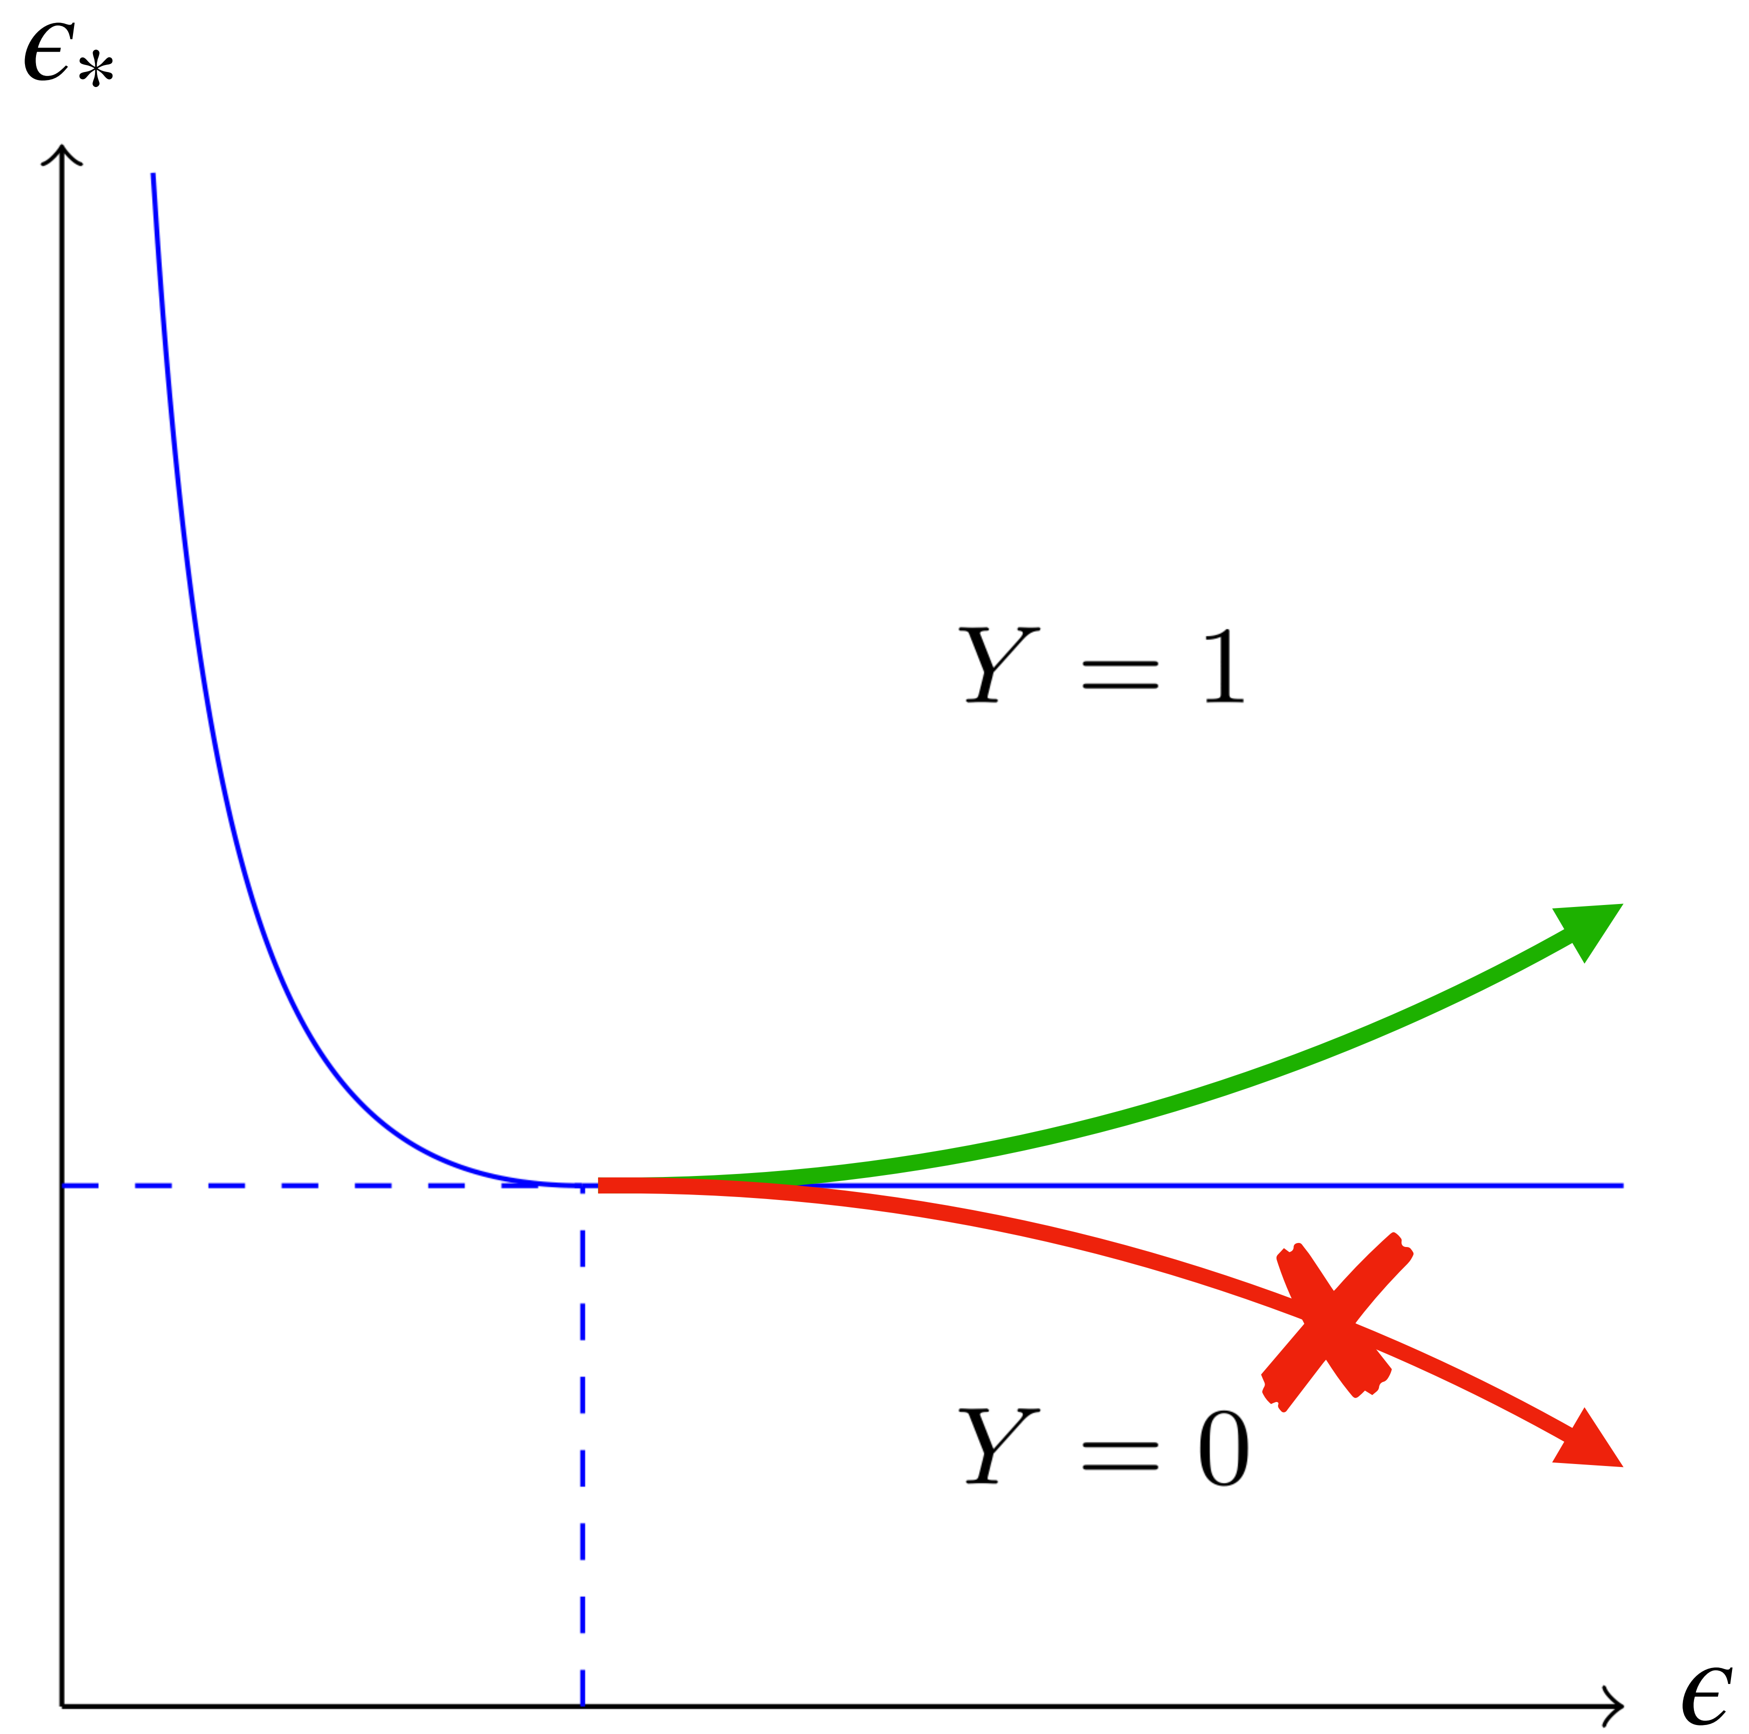

# Quantum state preparation perspective

Underestimating measurement strength

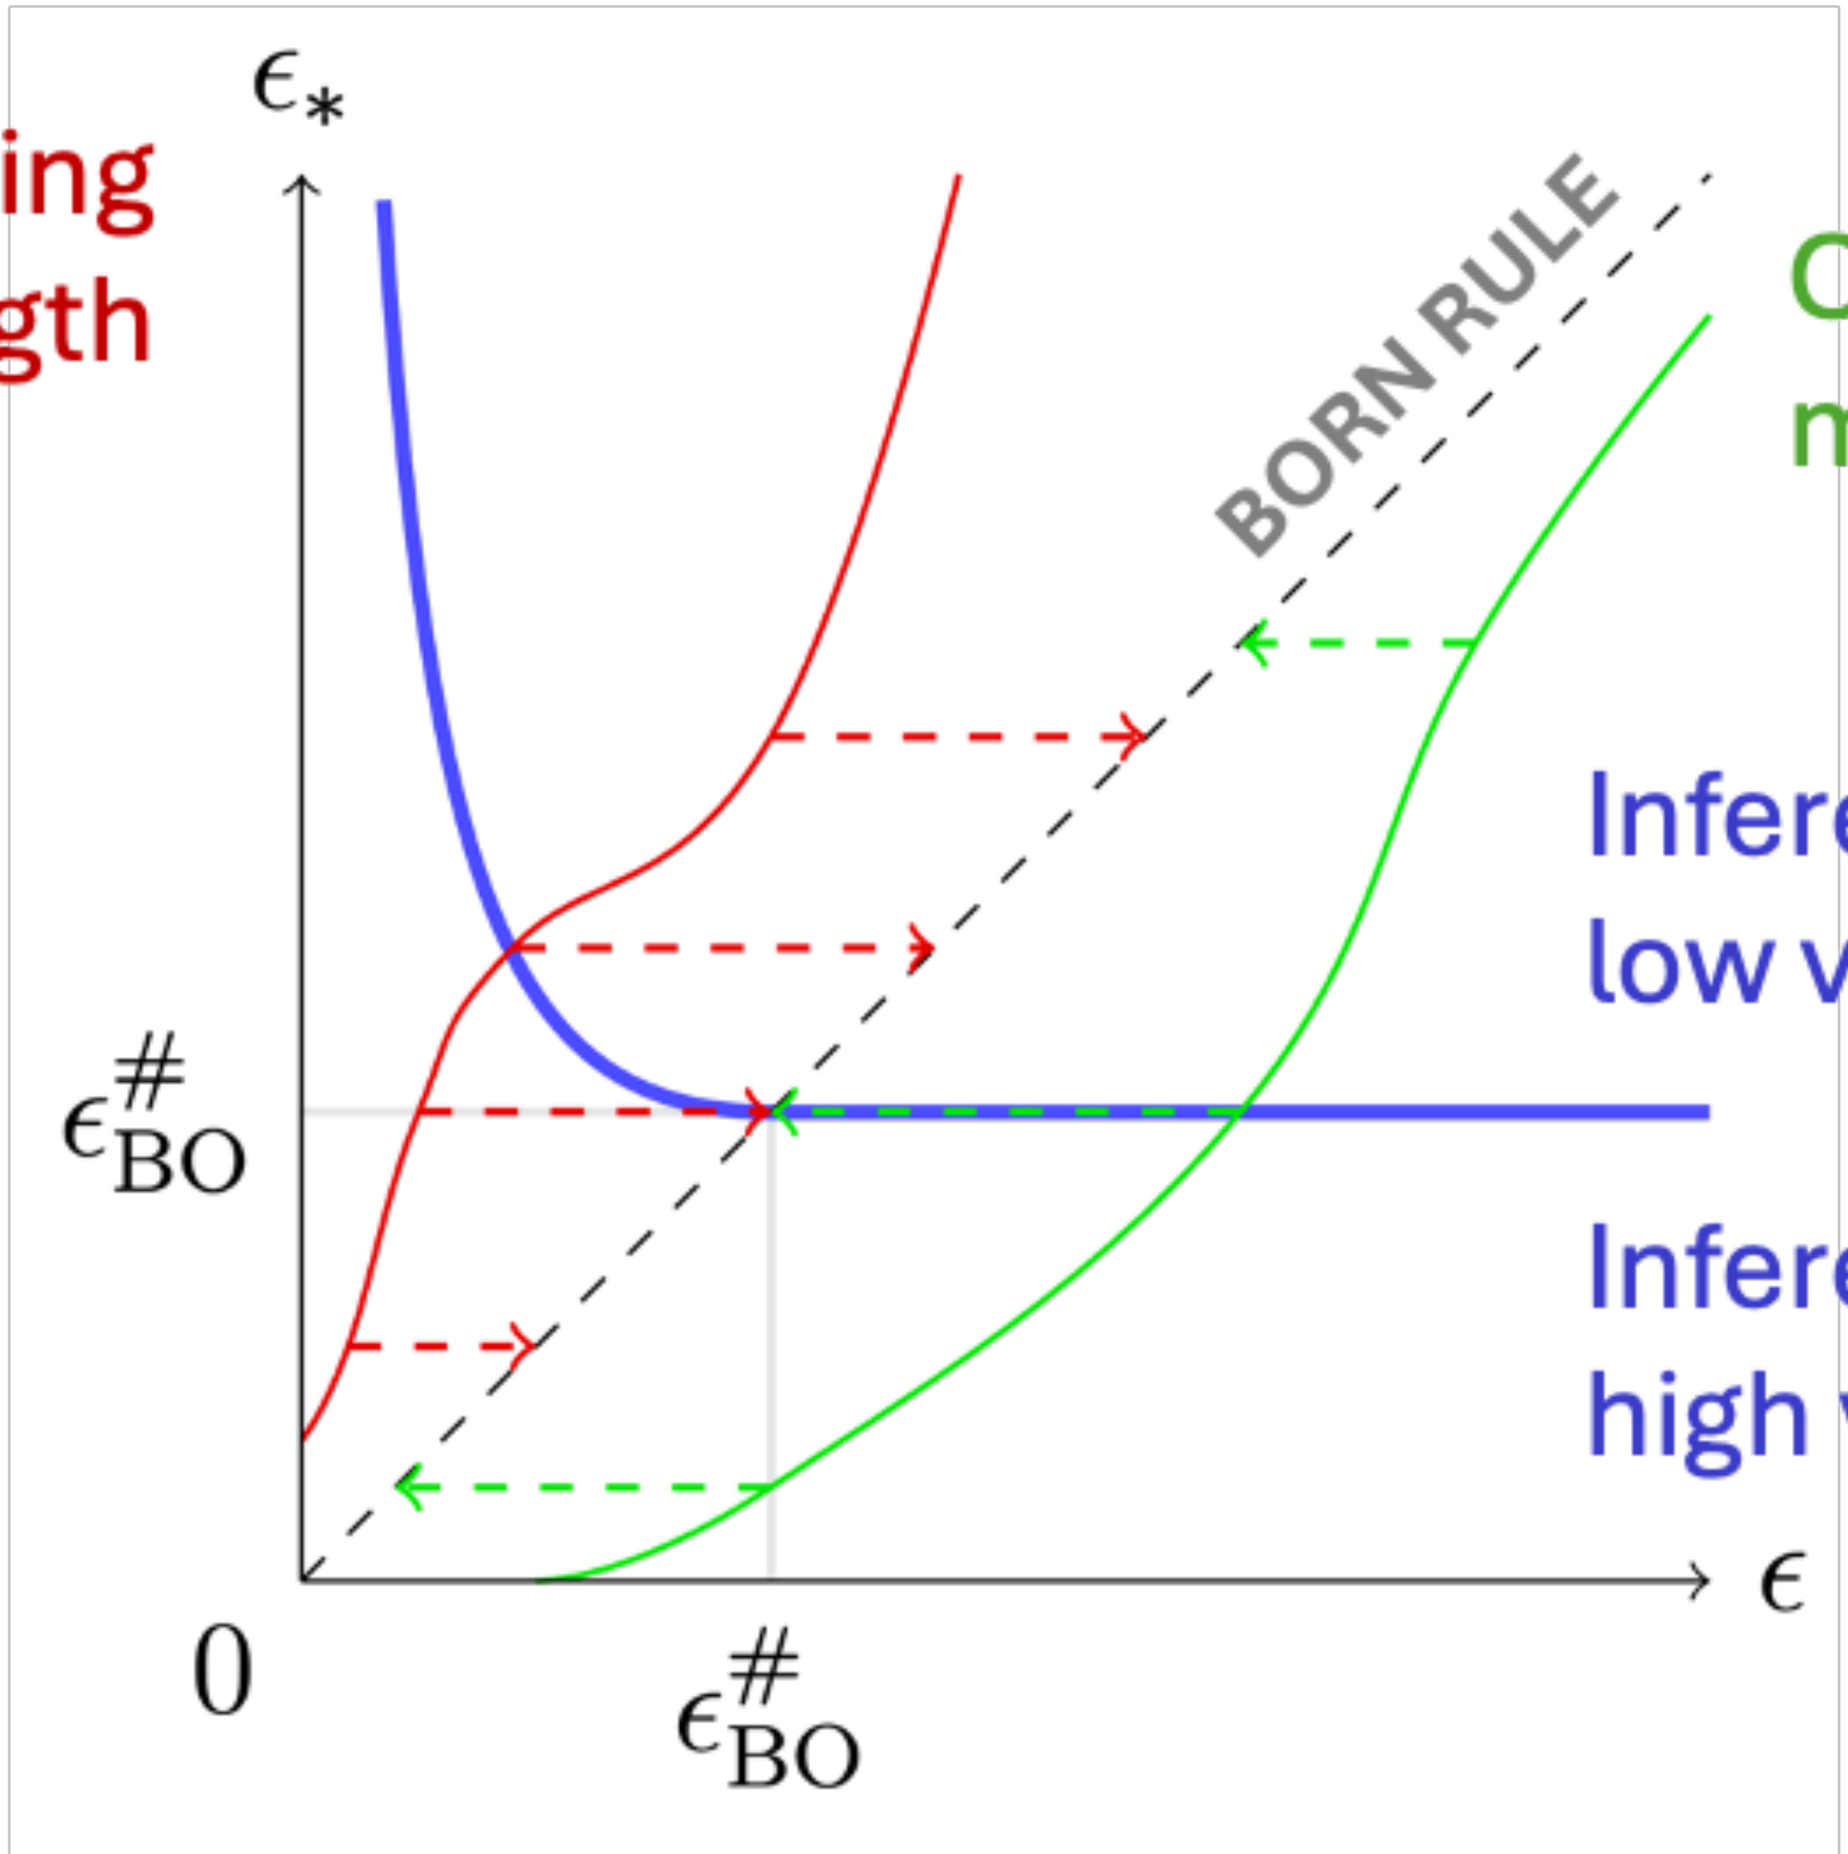


Overestimating measurement strength

Inference easy/ low variance

Inference hard/ high variance

# Outlook

- Can we use quantum fluctuations to get different universality class to the classical HMM models? (ex. quantum planted directed polymer)
- General field-theoretic RG analysis off of the Bayes-optimal line (cf. Nahum, Jacobson 2025; Gopalakrishnan, McCulloch, Vasseur 2025)
- Non-Markovian inference problems (ex. errors with memory in QEC)