

PAPER • OPEN ACCESS

Fisher information of correlated stochastic processes

To cite this article: Marco Radaelli *et al* 2023 *New J. Phys.* **25** 053037

View the [article online](#) for updates and enhancements.

You may also like

- [Relationships among Global Climate Indices and Rainfall Pattern to Detect Impact of Climate Change in Yogyakarta Special Region, Indonesia](#)
B D A Nugroho, R K Nurrohman and P Sudira

- [A close look at market microstructure](#)
Giulia Iori

- [Review and comparative assessment of engineering correlations for the fire-induced air inflow rate through a compartment opening](#)
Dimitrios Pikiokos and Dionysios Kolaitis

New Journal of Physics

The open access journal at the forefront of physics

Deutsche Physikalische Gesellschaft 

IOP Institute of Physics

Published in partnership with: Deutsche Physikalische Gesellschaft and the Institute of Physics



PAPER

OPEN ACCESS

RECEIVED
20 June 2022

REVISED
1 March 2023

ACCEPTED FOR PUBLICATION
5 May 2023

PUBLISHED
1 June 2023

Original Content from
this work may be used
under the terms of the
Creative Commons
Attribution 4.0 licence.

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Fisher information of correlated stochastic processes

Marco Radaelli^{1,*}, Gabriel T Landi^{1,2,3,*}, Kavan Modi^{4,5} and Felix C Binder^{1,*}

¹ School of Physics, Trinity College Dublin, Dublin 2, Ireland

² Instituto de Física da Universidade de São Paulo, 05314-970 São Paulo, Brazil

³ Department of Physics and Astronomy, University of Rochester, Rochester, NY 14627, United States of America

⁴ School of Physics & Astronomy, Monash University, Clayton, VIC 3800, Australia

⁵ Centre for Quantum Technology, Transport for New South Wales, Sydney, NSW 2000, Australia

* Authors to whom any correspondence should be addressed.

E-mail: radaelli@tcd.ie, giland@gmail.com and quantum@felix-binder.net

Keywords: information theory, metrology, Fisher information, spin chains

Abstract

Many real-world tasks include some kind of parameter estimation, i.e. the determination of a parameter encoded in a probability distribution. Often, such probability distributions arise from stochastic processes. For a stationary stochastic process with temporal correlations, the random variables that constitute it are identically distributed but not independent. This is the case, for instance, for quantum continuous measurements. In this article, we derive the asymptotic Fisher information rate for a stationary process with finite Markov order. We give a precise expression for this rate which is determined by the process' conditional distribution up to its Markov order. Second, we demonstrate with suitable examples that correlations may both enhance or hamper the metrological precision. Indeed, unlike for entropic information quantities, in general nothing can be said about the sub- or super-additivity of the joint Fisher information in the presence of correlations. To illustrate our results, we apply them to thermometry on an Ising spin chain, considering nearest-neighbour and next-to-nearest neighbour coupling. In this case, the asymptotic Fisher information rate is directly connected to the specific heat capacity of the spin chain. We observe that the presence of correlations strongly enhances the estimation precision in an anti-ferromagnetic chain, while in a ferromagnetic chain this is not the case.

1. Introduction

In many experimentally meaningful scenarios, a quantity of interest cannot be accessed directly via a measurement, but only retrieved from a sample of data; this problem is commonly referred to as *estimation*. Examples arise in many disciplines: for instance, one might be interested in estimating financial parameters from the time series of stock prices, the lethality of epidemic diseases from real-world hospital data, or the Rabi frequency of a two level atom. Estimation theory [1–4] provides safe formal grounds for this estimation procedure, and gives ways to compute the maximum achievable precision that can be attained for a certain parameter, given the probability distribution from which the sample is extracted. The variance σ_θ^2 of an unbiased estimator (inverse to the precision) for the parameter θ is subject to the Cramér-Rao bound [5, 6],

$$\sigma_\theta^2 \geq \frac{1}{F_{1:N}(\theta)}, \quad (1)$$

where $F_{1:N}(\theta)$ is the Fisher information associated to N outcomes. In the case of i.i.d. random variables, the Cramér-Rao bound simply reduces to its most well-known version, given by $F_{1:N}(\theta) = NF_1(\theta)$. However, this is not true when, as often happens in the real world, processes exhibit correlations, where past values influence the future behaviour, possibly up to a fixed number of steps.

Being able to quantify how many past outcomes have to be recorded to retrieve an accurate description of the future statistics is of crucial importance for data compression purposes, but also as a way to understand

the structure and complexity of the underlying model [7–9]. In the context of discrete-time classical stochastic processes, the notion of memory length is expressed by the Markov order [10].

An important scenario featuring such correlated outcomes is that of quantum systems subject to continuous weak measurements [11, 12]. This is common place in various experimental platforms, from quantum optics [13–15] to mesoscopic conductors [16]. It involves a stream of outcomes whose correlations crucially depend on the underlying quantum properties of the model. A famous example is the emission of light from quantum dots, which is known to present intermittencies [17] (called *blinking*) with very long range correlations [18, 19]. Several papers have recently analysed the metrological properties of continuous measurements [20–28], aiming both at constructing useful estimation strategies, as well as establishing the ultimate bounds in precision.

The application of estimation theory to correlated processes is still widely unexplored, and a general framework is lacking, except for very specific parameter encoding, such as location parameters [29, 30]. In this work, we specifically address this issue, and prove through suitable examples that the joint Fisher information is not in general sub-additive with respect to the marginal Fisher information; this fact sets the Fisher information apart from entropy-like information measures, for which sub-additivity is a fundamental requirement.

The main result of our work, summarized in equation (30), is a general decomposition for the Fisher information for a stationary stochastic process of finite Markov order \mathcal{M} . This result showcases the interplay between the number of outcomes N , and the Markov order \mathcal{M} . As we discuss in detail, this leads to two fundamental consequences. First, it shows that when $N \gg \mathcal{M}$, the Fisher information will asymptotically always scale as

$$F_N(\theta) \propto Nf(\theta), \quad (2)$$

where $f(\theta)$ is the Fisher information rate, which is independent of N . We here report the precise form of $f(\theta)$, its conditional dependence on preceding samples, and the non-asymptotic residual. Second, equation (2) allows us to establish the sub- or super-additivity of $F_{1:N}$, by comparing f with the Fisher information F_1 of a single outcome. In fact, we show below by means of examples that in general $f \leq F_1$; that is, unlike entropic information quantities, the Fisher information does not satisfy sub-additivity. Correlations can therefore be both advantageous and deleterious. The impact of these correlations in concrete estimation scenarios will be discussed in detail below.

This article is organized as follows. In section 2, we review the basis of estimation theory, introduce the notion of unbiased estimator and the Fisher information. In section 3 we discuss the behaviour of the Fisher information in the presence of correlations, showing with suitable toy models that in general it can exhibit super-additive or sub-additive behaviour with respect to the individual Fisher information. In section 4 we introduce the notion of Markov order, and present our main result: a decomposition of the joint Fisher information in a finite-Markov order stochastic process. We discuss the consequences of the decomposition for the Cramér-Rao bound. To provide more intuition for our results, in section 4.2, we discuss how the sign of the correlations affects the behaviour of the Fisher information, and we show that, for the specific class of estimators based on the sample mean, anti-correlated processes usually yield higher Fisher information. Finally, in section 5 we apply our results to the task of estimating temperature on an Ising spin chain, and show a direct connection between our decomposition of the Fisher information and the heat capacity per unit length (specific heat capacity) of the chain.

2. Estimation theory

Consider the following problem. We are given a probability distribution $P_\theta(x_{1:N})$, describing the statistics of a sequence $X_{1:N} = X_1, \dots, X_N$ (left- and right-inclusive notation will be assumed throughout this work) of N random variables, and depending on a vector of real parameters $\theta = (\theta_1, \theta_2, \dots, \theta_d)$. Our aim is to estimate, in the most precise way possible, the vector θ , by sampling from the probability distribution.

In many scenarios, the parameters vector θ reduces to a scalar, such as in the case of thermometry we will deal with below. Other times, θ is a genuine vector (multi-parameter estimation), for example of Cartesian coordinates, velocities, rotation or torsion parameters. In the following, we will introduce the formalism of estimation theory in full generality in a multi-parameter fashion.

It is possible to show that the precision that can be achieved on the estimation of θ is upper-bounded by an intrinsic property of the probability distribution itself. An *estimator* T is a function which, given as an input a realization $x_{1:N} = \{x_1, x_2, \dots, x_N\}$, outputs an estimate for the parameter vector:

$$T : \{x_1, x_2, \dots, x_N\} \mapsto \tilde{\theta}. \quad (3)$$

The Fisher information matrix is defined as [2, 5, 31] $F_{1:N}(\theta)$ as:

$$[F_{1:N}(\theta)]_{ij} = \sum_{x_{1:N}} P_\theta(x_{1:N}) \left(\frac{\partial}{\partial \theta_i} \log P_\theta(x_{1:N}) \frac{\partial}{\partial \theta_j} \log P_\theta(x_{1:N}) \right). \quad (4)$$

That is, the covariance between the two first derivatives (with respect to the parameter vector components θ_i and θ_j) of the logarithms of the probability distribution (i.e. the *surprisal*).

The importance of the Fisher information is mainly due to its role in the expression of the Cramér-Rao bound, the ultimate limit to the precision of any estimate of the parameters. We define the *bias* of an estimator [3] as the average difference between the estimate of the parameter vector and the true value of the parameter vector itself:

$$b(T) = \langle T(X_{1:N}) - \theta \rangle. \quad (5)$$

An estimator is *unbiased* if its bias is equal to zero. Given an unbiased estimator T , we can define its *covariance matrix* σ^2 as the matrix [2]

$$\sigma_{ij}^2 = \sum_{x_{1:N}} P_\theta(x_{1:N}) ([T(x_{1:N})]_i - \theta_i) ([T(x_{1:N})]_j - \theta_j). \quad (6)$$

The Cramér-Rao bound then reads [6, 32]

$$\sigma^2 \geq \frac{1}{F_{1:N}(\theta)}, \quad (7)$$

where $\frac{1}{F_{1:N}(\theta)}$ is the inverse of the Fisher information matrix, and the inequality means that the difference $\sigma^2 - \frac{1}{F_{1:N}(\theta)}$ is a positive semi-definite matrix. Positive definiteness implies, in particular, that the variances of each θ_k will satisfy equation (1), with F replaced by the (k, k) th component of the Fisher information matrix.

3. Fisher information in presence of correlations

The simplest parameter estimation scenario is that of independent and identically distributed (i.i.d.) random variables, for which factorization of the joint probability distribution

$$P_\theta(X_1, \dots, X_N) = P_\theta(X_1) \dots P_\theta(X_N), \quad (8)$$

holds. In this case the Fisher information matrix behaves additively:

$$F_{1:N}(\theta) = N F_1(\theta), \quad (9)$$

as proven in appendix B, where $F_1(\theta)$ is the Fisher information matrix corresponding to the probability distribution of $P_\theta(X)$. The Cramér-Rao bound (equation (7)) in this case expresses the variance in θ scaling inversely with N :

$$\sigma_\theta^2 \geq \frac{1}{N F_1(\theta)}. \quad (10)$$

In the presence of correlations the factorization of the joint probability distribution no longer holds. Instead, we define the conditional Fisher information matrix between two random variables X_1 and X_2 , as

$$[F_{X_2|X_1}(\theta)]_{ij} = \sum_{x_1, x_2} P_\theta(x_1, x_2) \cdot \left[\frac{\partial}{\partial \theta_i} \log P_\theta(x_2|x_1) \frac{\partial}{\partial \theta_j} \log P_\theta(x_2|x_1) \right]. \quad (11)$$

It is possible to show [33] that the joint and the conditional Fisher information matrices between two random variables are connected by the relation

$$F_{X_2|X_1}(\theta) = F_{X_1, X_2}(\theta) - F_{X_1}(\theta), \quad (12)$$

which may serve as an alternative definition of $F_{X_2|X_1}(\theta)$. The single-parameter case was treated in [33]. A proof for the multi-parameter case can be found in appendix B (we are not aware of this general proof in the literature).

It directly follows that the joint Fisher information matrix of N outcomes can be decomposed as [34]

$$F_{1:N}(\theta) = F_1(\theta) + \sum_{k=2}^N F_{k|1:k-1}(\theta), \quad (13)$$

valid also in the case of multi-parametric estimation. If the variables are independent, this reduces back to equation (9).

Until this point, our treatment has considered the case of multi-parameter estimation in full generality. From now on, for the sake of clarity we will restrict our attention to single-parameter estimation, and we will therefore think of the Fisher information as a scalar quantity. Our results can be trivially generalized to a multi-parameter scenario.

3.1. Lack of sub-additivity for the Fisher information

In general, nothing can be said about the inequality

$$F_{X_1 \dots X_N}(\theta) \stackrel{?}{\leq} F_{X_1}(\theta) + \dots + F_{X_N}(\theta). \quad (14)$$

In this sense, the Fisher information does not behave as an entropy-like quantity, for which convexity is a fundamental requirement [3]. To illustrate this idea, we consider a simple example based on Gaussian random variables X, Y , with identical mean μ and a covariance matrix σ defined as

$$\sigma = \gamma_0 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}, \quad (15)$$

for $\rho \in [-1, 1]$ and $\gamma_0 > 0$. The task at hand is then the estimation of the mean value μ . The joint Fisher information for Gaussian random variables with mean vector $\vec{\mu}$ and covariance matrix σ is well known, and given by $F = (\partial_\theta \vec{\mu})^T \sigma^{-1} (\partial_\theta \vec{\mu}) + \frac{1}{2} \text{tr}\{\sigma^{-1} (\partial_\theta \sigma) \sigma^{-1} (\partial_\theta \sigma)\}$ [1]. Applying this to the estimation of μ itself yields

$$F_{X,Y}(\mu) = \frac{1}{\gamma_0} \frac{2}{1 + \rho}, \quad (16)$$

The marginal Fisher information, on the other hand, is given by:

$$F_X(\mu) = F_Y(\mu) = \frac{1}{\gamma_0}. \quad (17)$$

whence

$$\frac{F_{X,Y}(\mu)}{F_X(\mu) + F_Y(\mu)} = \frac{1}{1 + \rho}. \quad (18)$$

It is therefore clear that we obtain a super-additive behaviour ($F_{X,Y} > F_X + F_Y$) when the variables are anti-correlated, $\rho < 0$, and a sub-additive behaviour ($F_{X,Y} < F_X + F_Y$) when they are correlated, $\rho > 0$. For instance, if X and Y are perfectly correlated ($\rho = 1$), we get $F_{X,Y} = F_X$: We learn nothing from Y that we had not already learned from X . The fact that we can get super-additive behaviour for $\rho < 0$, though, is not at all intuitive. We will revisit this in section 4.2, where we show that it happens because for anti-correlated variables the errors tend to cancel out, while in the correlated case they tend to add up.

4. The Markov order

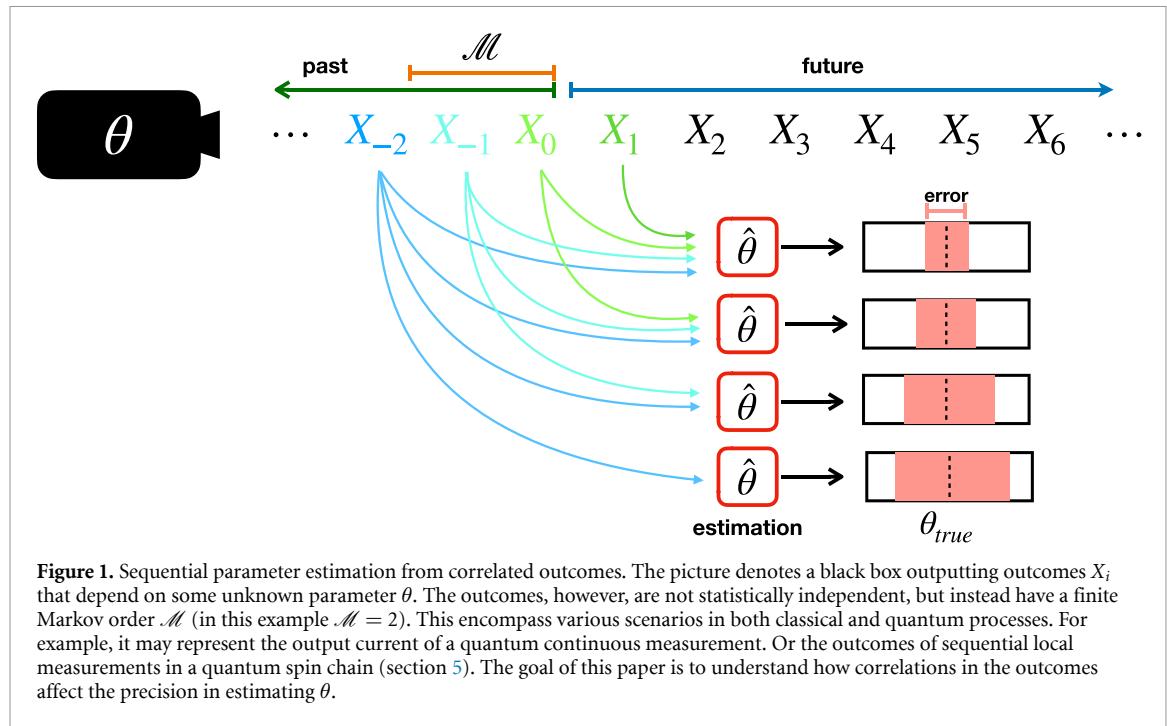
Let us consider a family of stochastic processes $\dots, X_0, X_1, \dots, X_N, \dots = \overleftrightarrow{X}$, fully represented by the overall joint parameter-dependent probability distribution $P_\theta(\overleftrightarrow{x})$. Here, we idealize stochastic processes as bi-infinite. A process is *stationary* [3] if any marginal of the overall joint probability distribution is translationally invariant, i.e.:

$$P_\theta(X_{k:m} = x_{k:m}) = P_\theta(X_{k+l:m+l} = x_{k:m}), \quad (19)$$

for any $k, m, l \in \mathbb{Z}, m \geq k$.

Given a stationary stochastic process, one can define its Markov order [10] as:

$$\mathcal{M} = \min \{\ell \text{ s.t. } P(X_1|X_{-\ell:0}) = P(X_1|X_{-\infty:0})\}, \quad (20)$$



where $X_{-\infty:0} = \overleftarrow{X}$ denotes the left-semi-infinite string of random variables in the stochastic process. The concept is illustrated in figure 1. The definition of the Markov order implies the possibility to truncate any conditioning in the probabilities to a fixed number of steps in the past, representing the *memory* of the system [35] (see also equation (28) below):

$$P(\overrightarrow{X} | \overleftarrow{X}) = P(\overrightarrow{X} | X_{-\mathcal{M}+1:0}), \quad (21)$$

where $\vec{X} = X_{1:+\infty}$. For example, in a process of Markov order $\mathcal{M} = 2$, all knowledge about the future is encoded in the two ‘most recent’ random variables in the past:

$$P(\overrightarrow{X} | \overleftarrow{X}) = P(\overrightarrow{X} | X_{-1}X_0). \quad (22)$$

Markov order $M = 0$ represents independent random variables

$$P(\overrightarrow{X} | \overleftarrow{X}) = P(\overrightarrow{X}). \quad (23)$$

Processes yielding Markov order $M \leq 1$ are customarily called *Markov processes*.

It is worth emphasizing that the Markov order only establishes conditional independence.

Unconditionally, *all* random variables in the chain are correlated (unless $\mathcal{M} = 0$). These correlations, however, decay exponentially with their distance for all processes with finite Markov order. For instance, in the case of Markov order $\mathcal{M} = 1$ the process forms a Markov chain specified by the transition probabilities $Q_{xy} = P(X_{i+1} = x | X_i = y)$. For any two random variables X_i and X_j , one will then have that (assuming $j > i$)

$$P(X_j|X_i) = (Q^{j-i})_{X_j, X_i}. \quad (24)$$

Since Q forms a stochastic matrix, its eigenvalues must have absolute value no larger than unity. Thus, besides possible deterministic periodicities, the correlations between X_i and X_j will decay exponentially with $|j - i|$. The same reasoning also holds for processes having $\mathcal{M} > 1$, provided that $|j - i| > \mathcal{M}$.

The interplay between conditional and unconditional independence is subtle, and often counterintuitive. Some non-trivial examples are discussed in appendix F. We also mention that, in general the Markov order is a *structural* property of a stochastic process (or *topological* in the language of computational mechanics [36]), independent of the precise choice of the parameters; other quantities like the (covariance-based) correlation length, on the contrary, usually exhibit a significant dependence on the parameters.

Alongside the probability-theoretic approach (20), it is equivalently possible to define the Markov order from an information-theoretic perspective as [37]

$$\mathcal{M} = \min \{ \ell \text{s.t. } H(X_{1:\ell}) = \ell h + E \}. \quad (25)$$

Here, H denotes the Shannon entropy of the distribution, while h is the *entropy rate*, defined as

$$\lim_{N \rightarrow \infty} \frac{1}{N} H(X_{1:N}). \quad (26)$$

Finally, the *excess entropy* E is given by the mutual information between past and future

$$E = \lim_{N \rightarrow \infty} I(X_{-N:0}; X_{1:N}). \quad (27)$$

It follows that, for a process of finite Markov Order \mathcal{M} , the memory renders past and future conditionally independent:

$$I(\vec{X}; X_{-\infty:-\mathcal{M}} | X_{-\mathcal{M}+1:0}) = 0, \quad (28)$$

meaning that no information about the future is retained in the past beyond the Markov order. The relation between the information-theoretic and the probability-theoretic versions of the Markov order is discussed in appendix A.

For a stationary process of finite Markov order \mathcal{M} we now consider the decomposition (13) for the Fisher information for a string of $N \geq \mathcal{M}$ outcomes $F_{1:N}$. In this case, all the conditional probabilities can be truncated at length \mathcal{M} , and the same will happen to the conditional Fisher information. We are therefore left with:

$$\begin{aligned} F_{1:N}(\theta) &= F_1(\theta) + \sum_{k=2}^N F_{k|1:k-1}(\theta) = F_1(\theta) + F_{2|1}(\theta) + F_{3|1:2}(\theta) + \dots + F_{\mathcal{M}|1:\mathcal{M}-1}(\theta) \\ &+ \sum_{k=1}^{N-\mathcal{M}} F_{\mathcal{M}+k|k-1:\mathcal{M}+k}(\theta), \end{aligned} \quad (29)$$

but, since the process is stationary, all the terms of the summation yield the same contribution, and, using (13), we can directly write:

$$F_{1:N}(\theta) = F_{1:\mathcal{M}}(\theta) + (N - \mathcal{M}) F_{\mathcal{M}+1|1:\mathcal{M}}(\theta). \quad (30)$$

This is our main result. In the remainder of this section, we showcase its significance by discussing some of the main consequences.

It is natural to write equation (30) in the spirit of (25), as

$$F_{1:N}(\theta) = Nf(\theta) + \epsilon, \quad (31)$$

defining the Fisher information rate $f(\theta)$ as

$$f(\theta) = \lim_{N \rightarrow \infty} \frac{1}{N} F_{1:N}(\theta) = F_{\mathcal{M}+1|1:\mathcal{M}}(\theta), \quad (32)$$

and the excess Fisher information $\epsilon(\theta)$ as

$$\epsilon(\theta) = F_{1:N}(\theta) - Nf(\theta) = F_{1:\mathcal{M}}(\theta) - \mathcal{M}f(\theta). \quad (33)$$

This expression leads to a clear analogy with (25); however, it is remarkable that, as already shown in the example of the two Gaussian random variables (in section 3), ϵ can be both positive or negative, while the excess entropy E is always positive.

To provide a concrete example, consider the case $\mathcal{M} = 1$. We then get

$$f = F_{2|1} = F_{12} - F_1, \quad (34)$$

$$\epsilon = F_1 - f = 2F_1 - F_{12}. \quad (35)$$

We must always have $F_{12} > F_1$, since extra information can only improve the estimation. However, ϵ does not have a well-defined sign because $F_{12} \leq 2F_1$ (or, what is equivalent $F_{2|1} \leq F_1$).

Exploiting the decomposition (30), the Cramér-Rao bound (7) can be expressed as

$$\sigma_\theta^2 \geq \frac{1}{F_{1:N}(\theta)} = \frac{1}{F_{1:\mathcal{M}}(\theta) + (N-\mathcal{M})F_{\mathcal{M}+1|1:\mathcal{M}}(\theta)} \sim \frac{1}{NF_{\mathcal{M}+1|1:\mathcal{M}}(\theta)}, \quad (36)$$

where the last approximation is valid in the large N limit. Hence: for a finite-Markov order process, the maximum precision of the measurement, given by the Cramér-Rao bound, grows linearly with the number of measurements, in the limit of a large number of measurements.

Since almost every stochastic process can always be assumed to have a finite (possibly large) Markov order, this means that for sufficiently large N the precision will always eventually scale linearly with N . The remaining question, then, is how does $F_{\mathcal{M}+1|1:\mathcal{M}}$ behave. More concretely, we can aim to compare it with F_1 , which would be the Fisher information that one would obtain if the outcomes were independent.

Equation (18), however, has already provided one illustration that nothing can in general be said about this. If $F_{\mathcal{M}+1|1:\mathcal{M}} > F_1$, then having correlations is advantageous for the estimation, while if $F_{\mathcal{M}+1|1:\mathcal{M}} < F_1$ they are deleterious. We therefore now turn to additional examples, which aim to pinpoint which ingredients might lead to $F_{\mathcal{M}+1|1:\mathcal{M}} \leq F_1$.

4.1. An example of super- and sub-additive behaviours

We illustrate the above results with minimal two-state models having Markov order $\mathcal{M} = 1$. For instance, a model having a sub-additive Fisher information is one given by the transition matrix:

$$P_\theta(X_{j+1}|X_j) = \begin{pmatrix} \theta & 1-\sqrt{\theta} \\ 1-\theta & \sqrt{\theta} \end{pmatrix}, \quad \theta \in [0, 1], \quad (37)$$

Conversely, a model having a super-additive Fisher information is that given by

$$P_\theta(X_{j+1}|X_j) = \begin{pmatrix} \theta & 1-e^{-\theta/3} \\ 1-\theta & e^{-\theta/3} \end{pmatrix}, \quad \theta \in [0, 1]. \quad (38)$$

The Fisher information sub- or super-additivity is spoiled by the difference $F_{1:2} - 2F_1$, or equivalently $f - F_1$, represented for both the processes in the left-most plot of figure 2 as a function of the parameter θ .

The relevance of the sub- or super-additive behaviour becomes apparent if we consider an actual estimation task for these two models. It is well-known [3, 38] that the Maximum Likelihood Estimator always asymptotically attains the Cramér-Rao precision bound (also in the case of correlated variables, under adequate regularity conditions, see appendix C); in figure 2 we compare the mean squared error (MSE) given by the maximum likelihood estimation (MLE) on the models with our asymptotic bound (36).

The plots show two relevant facts: first, the convergence of the MSE of the MLE estimations to the asymptotic Cramér-Rao bound (36) is verified; second, such a bound can be both much larger and much smaller than the i.i.d. Cramér-Rao bound for the specific case: once more, the Fisher information behaves sometimes super- and sometimes sub-additively.

The existence of processes yielding a sub-additive Fisher information could raise the question whether it could be possible to *beat the system* by employing some estimator which, by ignoring the correlations, can achieve the i.i.d. Cramér-Rao bound. It turns out that this is not the case, as discussed in more details in appendix D: the Fisher information (and consequently the Cramér-Rao bound) is an intrinsic property of the process, and cannot be circumvented by any choice of an estimator.

4.2. Influence of the sign of correlations in the sample mean

Equation (18) provided an example where the sign of the correlations $\rho \leq 0$ determined whether correlations are advantageous or not for the estimation process. When considering a general estimator, it unfortunately turns out that a direct link cannot be established; however, an interesting insight can be gained by restricting our attention to a specific subclass of estimators, namely those based on the sample mean. The employment of the sample mean usually represents the simplest choice when asked to estimate something based on a stochastic process.

Let us consider a θ -dependent stationary stochastic process X_1, \dots, X_N ; its sample mean is given by

$$Y = \frac{X_1 + \dots + X_N}{N}. \quad (39)$$

Let $\mu = \mathbb{E}[X_i]$ be the local expectation value, and $\sigma^2 = \text{Var}[X_i]$ the local variance. Due to the stationarity requirement, the covariance between two random variables X_i and X_j depends only on their distance, let $\text{Cov}[X_i, X_j] = C_{i-j}$. We refer to $C_i > 0$ as *positively correlated* and $C_i < 0$ as *negatively correlated*.

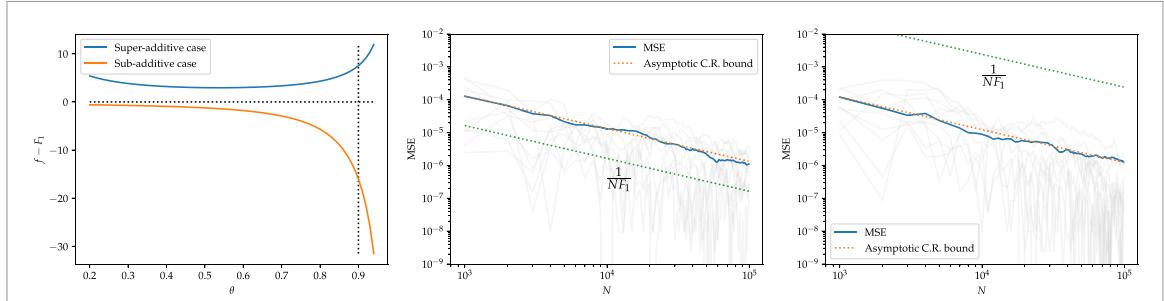


Figure 2. Two toy-model estimation tasks using the Maximum Likelihood Estimator show the convergence of the expectation value of the MSE, obtained by averaging over 50 iterations of the simulation, to the asymptotic behaviour given by (36). The expression of the conditional probabilities yielding these processes are in the main text, and the left-most plot shows the corresponding sub- or super-additive behaviours of the Fisher information. The semi-transparent lines represent the MSE of single runs of the simulation. Notice that the bound is achieved only asymptotically, hence it is possible to have, for finite N , an MSE lower than the bound.

Clearly, $\mathbb{E}[Y] = \mu$. The variance, on the other hand, reads

$$\text{Var}[Y] = \frac{1}{N}\sigma^2 + \frac{2}{N^2} \sum_{i=1}^{N-1} (N-i)C_i, \quad (40)$$

where we used the relation

$$\sum_{i,j>i} \text{Cov}[X_i, X_j] = \sum_{i=0}^{N-1} \sum_{j=i+1}^{N-1} C_{i-j} = \sum_{i=1}^{N-1} (N-i)C_i. \quad (41)$$

Under appropriate hypotheses (discussed in appendix E), a generalization of the Central Limit Theorem for correlated variables guarantees asymptotic normality of the distribution of Y , with average μ and variance ν , given by

$$\nu = \frac{1}{N} \left[\sigma^2 + 2 \sum_{i=1}^{N-1} C_i \right]. \quad (42)$$

Under this normality assumption, the Fisher information of the sample mean Y will be:

$$F_Y(\theta) = N \left[\frac{1}{\nu} \left(\frac{\partial}{\partial \theta} \mu \right)^2 + \frac{1}{2N} \frac{1}{\nu^2} \left(\frac{\partial}{\partial \theta} \nu \right)^2 \right], \quad (43)$$

where $\nu = N\nu$. Notice that the second term is negligible in the $N \rightarrow \infty$ limit, hence we finally get a Fisher information

$$F_Y(\theta) \sim N \left(\frac{\partial \mu}{\partial \theta} \right)^2 \frac{1}{\sigma^2 + 2 \sum_{i=1}^{N-1} C_i}. \quad (44)$$

This result, which also appears in [39] in the context of super-resolution imaging theory, clearly illustrates the role that the sign of C_i has on the Fisher information. Positively correlated variables ($C_i > 0$) increase the denominator and hence decrease F_Y , while negatively correlated C_i do the opposite. *For estimation, it is thus advantageous to have negatively correlated outcomes.* To understand the intuition behind this, consider equation (39) and write $X_i = \mu + \delta X_i$, where δX_i represent the fluctuations around the mean. Then

$$Y = \mu + \frac{\delta X_1 + \dots + \delta X_N}{N}. \quad (45)$$

If two variables δX_i and δX_j are negatively correlated, then an outcome $\delta X_i > 0$ implies there is a tendency to observe $\delta X_j < 0$. As a consequence, *errors tend to cancel out* in equation (45). Conversely, for positively correlated variables, the errors tend to add up.

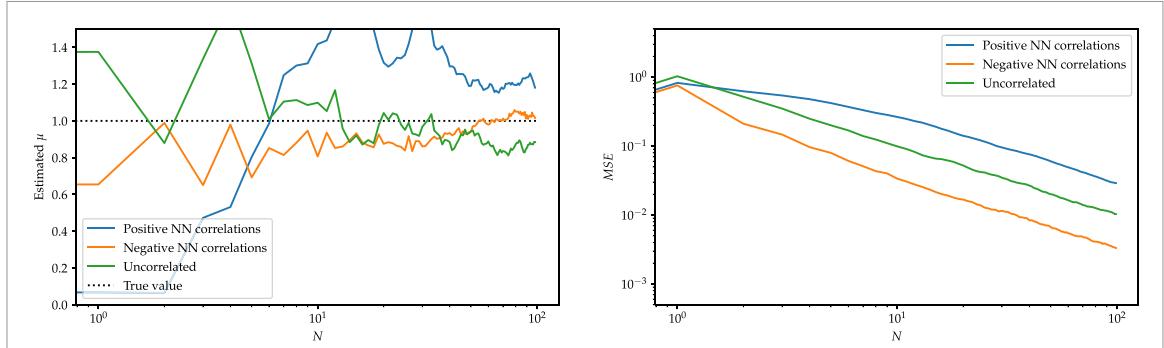


Figure 3. A simulated estimation task exploiting the sample mean on a Markov chain of Gaussian variables, for different values of the correlation parameter ρ . On the left, a single run of the estimation, on the right the MSE is averaged over 1000 iterations of the task. For details about the construction of this Gaussian Markov process, see appendix F.

Since F_Y is the Fisher information for a restricted set of estimators, it must necessarily lower bound the full Fisher information,

$$F_Y \leq F_{1:N}. \quad (46)$$

Using (30) we have therefore that, in the $N \rightarrow \infty$ limit, $F_Y \leq NF_{\mathcal{M}+1|1:\mathcal{M}}$. Hence,

$$F_{\mathcal{M}+1|1:\mathcal{M}}(\theta) \geq \left(\frac{\partial \mu}{\partial \theta} \right)^2 \frac{1}{\sigma^2 + 2 \sum_{i=1}^{N-1} C_i}. \quad (47)$$

This provides a fundamental bound for the asymptotic conditional Fisher information, clearly showcasing the effects of correlations.

These ideas are illustrated in figure 3, which analyses the estimation of the mean μ of a Gaussian process, with tunable nearest-neighbour correlations given by the parameter $\rho \in [0, 1]$. Details on how to construct the process are discussed in appendix F. The left panel shows a single stochastic realization using the sample mean as an estimator. As can be seen, positively and negatively correlated samples seem to converge at a somewhat similar rate to the true parameter, while uncorrelated samples seem to converge more slowly. The impression from this single realization, however, is deceiving. In the right panel of figure 3 we show the MSE averaged over multiple realizations. It is now quite clear that negatively correlated samples yield much more precise estimates when compared to the uncorrelated case (notice the log scale), in agreement with the predictions of equation (44). Conversely, positively correlated samples yield a clear and sizable disadvantage.

5. Application to thermometry of spin chains

In this section, we apply the above detailed methods to the task of temperature estimation on an Ising spin chain [40], consisting of N 1/2-spins, with Hamiltonian

$$\hat{H} = \sum_j -B\hat{\sigma}_j^z - \sum_k J_k \sum_j \hat{\sigma}_j^z \hat{\sigma}_{j+k}^z, \quad (48)$$

where σ_j^z are Pauli matrices, B is the external magnetic field, and J_k is a k -distance coupling strength. The system is prepared in a thermal state

$$\hat{\rho} = \frac{1}{Z} e^{-\hat{H}/T}, \quad (49)$$

whose temperature T we wish to estimate (i.e. it plays the role of the parameter θ), and Z is the partition function (we set the Boltzmann constant to $k_B = 1$).

Due to the structure of the Gibbs state ρ , the best strategy for thermometry is always to measure the system in the energy basis [41, 42]. Equation (48) is already diagonal, and thus a projective energy measurement is tantamount to N local measurements in the basis of σ_j^z . The outcome is a string $\mathbf{s} = (s_1, \dots, s_N)$, with $s_j = \pm 1$, which occurs with probability

$$p(\mathbf{s}) = \frac{1}{Z} \exp \left[-\frac{1}{T} H(\mathbf{s}) \right], \quad (50)$$

where

$$H(\mathbf{s}) = \sum_j -Bs_j - \sum_k J_k \sum_j s_j s_{j+k}. \quad (51)$$

Because of the interactions, the outcomes of each site are not statistically independent. Hence $p(\mathbf{s})$ forms a Markov chain, precisely of the form described in the previous section.

For intuition, one may imagine that the spins are measured sequentially, from left to right as in figure 1. In this way, subsequent measurements (each yielding the current state of the measured spin) will progressively form a stochastic process. To apply the definition of the Markov order (20) to the spin chain, one first has to be able to marginalize the chain's thermal state (50) in order to express probability distributions for configurations of finite number of spins. The procedure for doing so is based on the well-known transfer matrix method [43] and is detailed in [44, 45]. In appendix G, we briefly summarize the main results.

The problem of the Markov order of spin chains has been discussed in [10]. Let us first define the *interaction range* R of a spin chain as

$$R = \min \{k \text{ s.t. } J_s = 0 \ \forall s > k\}. \quad (52)$$

For instance, for nearest neighbour interactions $J_1 \neq 0$, while $J_s = 0$ for $s > 1$, which would lead to $R = 1$. Naively, one might expect that $R = \mathcal{M}$. However, this is not always true. In [44] it was implicitly proved that $R \geq \mathcal{M}$ for a completely general spin chain. Despite its intuitiveness, it is not clear whether the equality holds in all cases (as already noted in [45]). E.g. obvious exceptions appear for $T = 0$ and $1/T = 0$; however, these extremal, singular points bear no mathematical relevance in this context where T is the parameter to be estimated. We leave a general proof of the full conditions under which $R = \mathcal{M}$ as open question. A review of the discussion about the Markov order of classical spin chains can be found in appendix H.

Once we have determined the Markov order of a specific spin chain, for any number of consecutive spins larger than the Markov order the Fisher information scales according to (30).

If a system is in a thermal state, expressed by a Gibbs distribution (50), then the well-known relation between the thermal Fisher information and the heat capacity C holds [41]:

$$F(T) = \frac{C}{T^2}, \quad (53)$$

where C is defined as:

$$C = \frac{\partial \langle H \rangle}{\partial T}. \quad (54)$$

It is important to remark that this relation holds exactly only if we take into account the whole chain, since the marginals of a Gibbs distribution are not, in general, of Gibbs form. However, the Fisher information decomposition (30) allows us to relate directly the heat capacity per unit length (i.e. the *specific heat capacity*) to the Fisher information of suitable marginals.

Let us recall the decomposition of the Fisher information (equation (31)); defining the specific heat capacity as

$$c = \frac{C(N)}{N}, \quad (55)$$

where $C(N)$ denotes the heat capacity for a chain of length N , we have:

$$\frac{c}{T^2} = \frac{F_{1:N}(T)}{N} = \frac{1}{N} (Nf + \epsilon). \quad (56)$$

Taking the $N \rightarrow \infty$ limit, since neither f nor ϵ depend on N , we then obtain:

$$\lim_{N \rightarrow \infty} \frac{c}{T^2} = f = F_{\mathcal{M}+1|1:\mathcal{M}}(T). \quad (57)$$

Spin chains constitute an extremely interesting playground for our discussion: already in their simplest version, the nearest-neighbour Ising chain, they exhibit a surprising richness of features. This nearest-neighbour model requires a single parameter to encode the coupling strength; we will call it J instead of J_1 , for the sake of conciseness. Given the decomposition (36), for $\mathcal{M} = 1$ all the joint probabilities involving more than two spins are fully determined by the one-site and two-sites probabilities. Hence, it

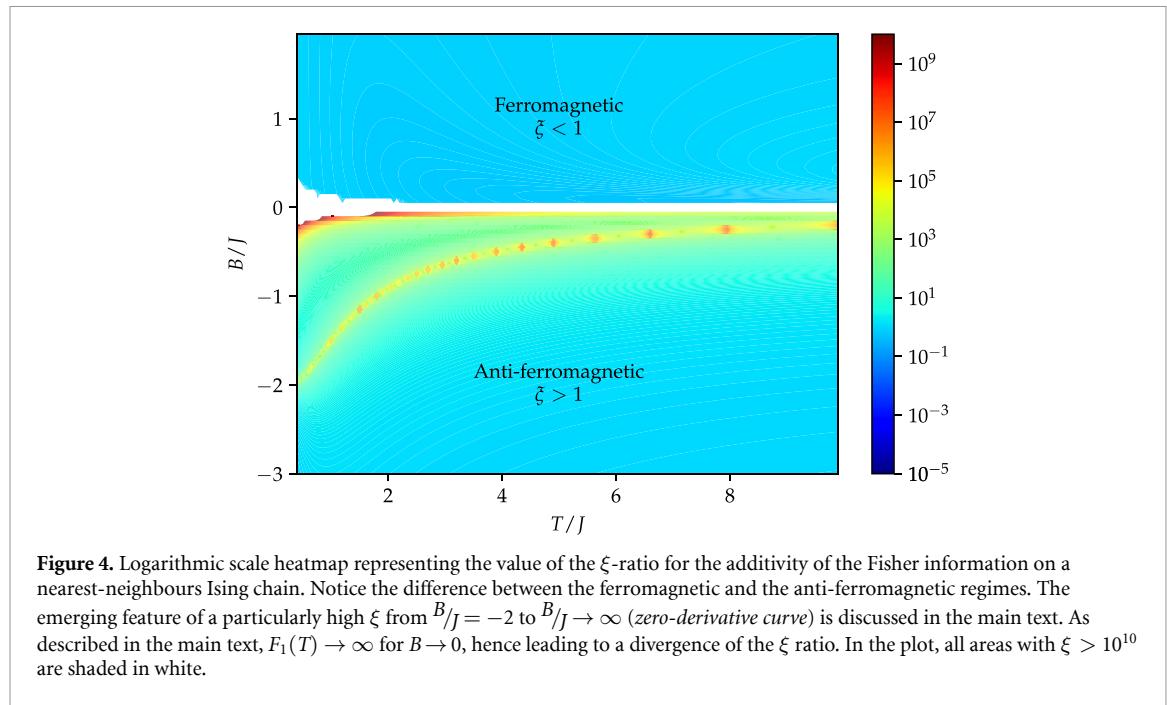


Figure 4. Logarithmic scale heatmap representing the value of the ξ -ratio for the additivity of the Fisher information on a nearest-neighbours Ising chain. Notice the difference between the ferromagnetic and the anti-ferromagnetic regimes. The emerging feature of a particularly high ξ from $B/J = -2$ to $B/J \rightarrow \infty$ (*zero-derivative curve*) is discussed in the main text. As described in the main text, $F_1(T) \rightarrow \infty$ for $B \rightarrow 0$, hence leading to a divergence of the ξ ratio. In the plot, all areas with $\xi > 10^{10}$ are shaded in white.

makes sense to discuss the relation between the 2-sites joint Fisher information $F_{1:2}(T)$ and the 1-site marginal Fisher information $F_1(T)$. In the same spirit of the Gaussian example in section 3, we can discuss the inequality

$$F_{1:2}(T) \stackrel{?}{\leqslant} 2F_1(T). \quad (58)$$

Or, what is equivalent, $F_{2|1} \stackrel{?}{\leqslant} F_1$ (since $F_{1:2} = F_1 + F_{2|1}$ in this case). Let us define the ratio

$$\xi = \frac{F_{1:2}(T)}{2F_1(T)}. \quad (59)$$

Super-additive behaviour corresponds to $\xi > 1$, and sub-additive to $\xi < 1$; independent random variables yield a perfectly additive behaviour with $\xi = 1$. The value of ξ for a wide choice of parameters is represented in figure 4.

It is apparent that there is a significant difference between the behaviours in the ferromagnetic and the anti-ferromagnetic regimes: the first gives $\xi \lesssim 1$ almost everywhere, corresponding to a slightly sub-additive Fisher information. In contrast, the anti-ferromagnetic regime sees significantly high ξ values, representing a strongly super-additive Fisher information.

Considering the shape of the different probabilities as a function of the temperature allows us to gain a precious intuition about the conditions under which each regime is attained, as illustrated in figure 5. Note how, in the $T \rightarrow 0$ limit, each curve reaches the probability of the specified configuration in the ground state of the chain: total anti-alignment in the anti-ferromagnetic case ($P(\uparrow\downarrow) = P(\downarrow\uparrow) = 1/2$) and total alignment in the ferromagnetic case ($P(\uparrow\uparrow) = P(\uparrow) = 1$).

The $\xi > 1$ condition in the anti-ferromagnetic case is mainly due to the strong suppression of the denominator $2F_1(T)$. The reason of the suppression is apparent if one considers how two two-sites joint probabilities add up to give a one-site marginal: the $P(\uparrow\uparrow)$ and $P(\uparrow\downarrow)$ probabilities in the anti-ferromagnetic case exhibit almost opposite derivatives with respect to T ; hence, their sum $P(\uparrow)$ will show a very mild dependence on the temperature, yielding a reduced Fisher information. This does not happen in the ferromagnetic regime, where the derivatives of $P(\uparrow\uparrow)$ and $P(\uparrow\downarrow)$ have significantly different absolute values.

This difference at the level of the derivatives can be easily thought in terms of the structure of the ground state of the chain, or, equivalently, of the thermal state for $T \rightarrow 0$. In the anti-ferromagnetic case, the ground state is two-fold degenerate (for values of B that are not too large): $P(\uparrow\downarrow)$ and $P(\downarrow\uparrow)$ both have $1/2$ probability, while $P(\uparrow\uparrow) = P(\downarrow\downarrow) = 0$. In the high- T limit, all the two-site probabilities will achieve a value of $1/4$. Hence, in the anti-ferromagnetic case $P(\uparrow\downarrow)$ has to lose $1/4$ probability between $T \rightarrow 0$ and $T \rightarrow \infty$; in the same temperature interval, $P(\uparrow\uparrow)$ has to gain $1/4$; the two derivatives can almost cancel one another.

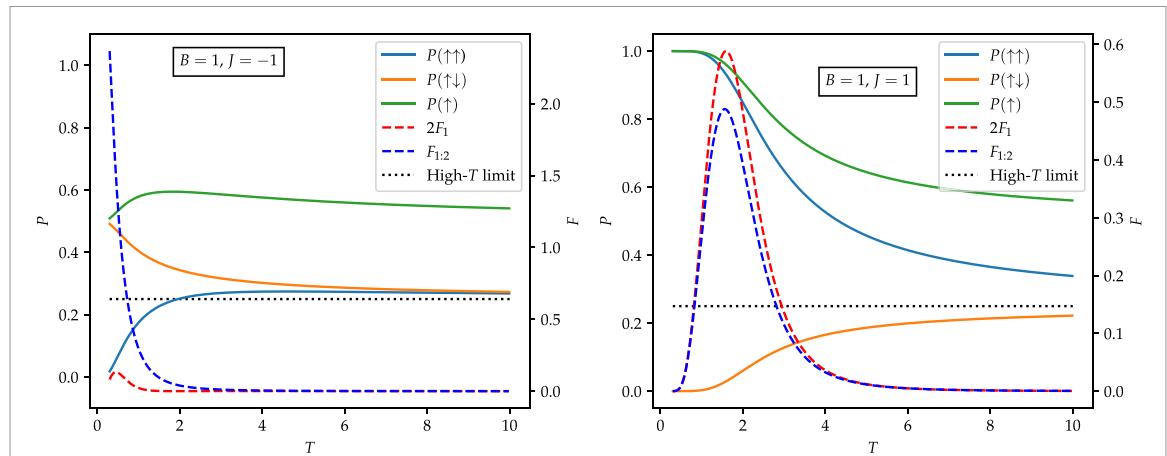


Figure 5. Relevant joint and marginal probabilities on the nearest-neighbours Ising spin chain, represented as a function of the temperature T , along with the one-site and two-sites Fisher information.

On the opposite side, the ferromagnetic chain has a non-degenerate ground state: $P(\uparrow\uparrow)$ has to lose $3/4$ of probability in the T -interval, while $P(\uparrow\downarrow)$ has only to gain $1/4$: the probabilities cannot exhibit opposite derivatives.

This perspective allows us also to interpret the emerging curve labeled as *zero-derivative curve* in figure 4 in the anti-ferromagnetic region. The curve corresponds to point of maximum or minimum of the one-site probability $P(\uparrow)$ with respect to the temperature; consequently, they have zero derivative and zero one-site thermal Fisher information, hence causing a divergence of the ξ ratio. Notice that the zero-derivative curve is only visible for a B/J ratio larger than -2 . This can be traced back to a fundamental regime change in the structure of the ground state, happening exactly at $B/J = -2$: the anti-ferromagnetic ground state $\dots \uparrow\downarrow\uparrow\downarrow\uparrow\downarrow \dots$ gives way to a ferromagnetic ground state $\dots \uparrow\uparrow\uparrow\uparrow\uparrow \dots$, due to the frustration of the anti-ferromagnetic interaction operated by the extremely strong external magnetic field. The one-site probability in this high- B regime is monotonous, thus leading to the absence of zero-derivative points.

Finally, the ξ ratio diverges at all temperatures for $B \rightarrow 0$; this is due to $F_1(T) \rightarrow 0$ in that limit. In fact, in the absence of an external magnetic field, it does not make any energetic difference for a single spin to be in state \uparrow or \downarrow ; no information about temperature can therefore be obtained from observations of single spins ($F_1(T) = 0$), whereas correlations can still give place to meaningful estimations of temperature ($F_{1:2}(T) \neq 0$).

We also considered how the introduction of a second-order coupling (J_2 in the language of equation (51)) affects the temperature estimation task. What appears is a much richer phenomenology, summarized in appendix I. The key points discussed above are preserved also in this more general case: chains with predominantly anti-ferromagnetic behaviour tend to give a large advantage, exploiting correlations for temperature estimation ($\xi \gg 1$), whereas chains dominated by ferromagnetic interactions tend to yield a disadvantage ($\xi < 1$).

6. Discussion and outlook

In this paper we studied the Fisher information of correlated stochastic processes. The central result is the decomposition of the joint Fisher information (30). It entails that the precision that may be achieved for parameter estimation from any process with finite Markov order obeys an asymptotic bound that scales linearly with the number of outcomes (36). A finite Markov order does not mean that the correlation is finite; on the contrary, even for $\mathcal{M} = 1$, all outcomes are correlated, except that these correlations decay exponentially. Thus, our results encompass an extremely large class of physical processes (a notable exception being critical systems, for which the correlation length diverges). Second, a central realization of our framework was that, in contrast to entropic information measures, the Fisher information does not necessarily obey subadditivity, and can be both super and sub-additive, depending on the problem. This means that correlations can sometimes help, or sometimes hamper the estimation process, in the sense that the convergence of the mean-squared error with N may be quicker or slower. As our results illustrate, generally speaking, anti-correlated outcomes tend to be advantageous, whereas correlated outcomes tend to result in reduced scaling.

Our focus here was on classical stochastic processes. It is natural to ask about a possible generalization of our main result, equation (30), to the quantum domain. In this regard, it is useful to distinguish between two scenarios. The first is that of continuous quantum measurements [12]. In this case, the process is quantum,

but the outcomes are classical and local. Indeed, quantum continuous measurements fall *precisely* within the paradigm of our framework, which is thus directly applicable. For instance, [22] studied the estimation of the Rabi frequency from photon counting experiments. Their results fall precisely within our framework with $\mathcal{M} = 1$. The same is true for [23], which studied homodyne detection.

A classical time series may be also be obtained from a quantum collision model [46–48]. For instance, we may consider a memory system which is inaccessible to direct observation. At each time step, a fresh probe systems is made to interact with the memory system at and then subject to a fixed, generalized, quantum measurement. Even for Markovian dynamics on the composite space, the resulting time series is in general non-Markovian, with any Markov order realizable. Our results apply directly to this scenario of indirectly probing a quantum system.

More generally, however, one may also consider scenarios requiring non-local measurements, which is often the case for systems exhibiting entanglement or other forms of quantum correlations. The generalization of the notion of Markov order to this case has proved to be rather cumbersome [35, 46]: when a completely general measurement is taken into account, it can be shown [35] that the quantum Markov order is either 0, 1 or infinite (one could instead try to apply the concept of approximate Markov order [49, 50]). In fact, it has been shown that the fundamental decomposition (13) generally does not hold for the Quantum Fisher information [33]; this fact is at the basis of the possibility of achieving higher precision by exploiting specifically quantum properties, such as entanglement. Furthermore, when dealing with multi-parameter estimation, one has to take into account the possible non-commutativity of the corresponding operators, thus leading to the necessity of more refined theoretical tools than the Fisher information matrix [4].

To give a simple example of a non-local measurement scenario, suppose one wished to do thermometry, exactly as in section 5, but considering instead the transverse field Ising model; i.e. where the field is now applied in the x direction. There has been considerable interest in thermometry of quantum systems (for reviews, see [51, 52]). And, as already discussed in section 5, in this case the optimal approach is to measure the system in the energy eigenbasis. In general, however, this basis is non-local and hence our main result, equation (30) is not necessarily applicable.

An extension to the fully-quantum case would be an interesting future avenue of research.

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgments

The authors thank Alessio Benavoli for helpful discussions on the Bayesian approach to parameter estimation. This research was supported by Grant Number FQXi-RFP-IPW-1910 from the Foundational Questions Institute and Fetzer Franklin Fund, a donor advised fund of Silicon Valley Community Foundation. G T L acknowledges the financial support of the São Paulo Funding Agency FAPESP (Grant No. 2019/14072-0), and the Brazilian funding agency CNPq (Grant No. INCT-IQ 246569/2014-0). M R acknowledges funding by the Irish Research Council under Government of Ireland Postgraduate Scheme Grant No. GOIPG/2022/2321. FCB acknowledges funding by the Irish Research Council under grant number IRCLA/2022/3922.

Appendix A. Information-theoretic and probability-theoretic definitions of the Markov order

In this appendix, we briefly recall the connection between the probability-theoretic and the information-theoretic definition of the Markov order [37]. In (20), we defined the Markov order as

$$\mathcal{M} = \min \{ \ell \text{ s.t. } P(X_0|X_{-\ell:0}) = P(X_0|X_{-\infty:0}) \}. \quad (\text{A1})$$

Let us now consider the N -site joint Shannon entropy $H(X_{1:N})$, for which a well-known chain rule holds [3]:

$$H(X_{1:N}) = \sum_{i=1}^N H(X_i|X_{i-1} \dots X_1). \quad (\text{A2})$$

Each conditional Shannon entropy can be written as:

$$\begin{aligned} H(X_i|X_{i-1}\dots X_1) &= - \sum_{x_1\dots x_i} P(x_1\dots x_i) \log P(x_i|x_{i-1}\dots x_1) \\ &= - \sum_{x_{i-\mathcal{M}}\dots x_i} P(x_{i-\mathcal{M}}\dots x_i) \log P(x_i|x_{i-1}\dots x_{i-\mathcal{M}}), \end{aligned} \quad (\text{A3})$$

where we used the probabilistic definition of the Markov order (20) in the last equality. Hence, the joint Fisher information can be reduced to:

$$H(X_{1:N}) = - \sum_{i=1}^N \sum_{x_{i-\mathcal{M}}\dots x_i} P(x_{i-\mathcal{M}}\dots x_i) \cdot \log P(x_i|x_{i-1}\dots x_{i-\mathcal{M}}) = H(X_{1:\mathcal{M}-1}) + (N-\mathcal{M})H(X_{1:\mathcal{M}}), \quad (\text{A4})$$

where in the last step the stationarity of the process has been exploited. With the definitions

$$\begin{aligned} h &= H(X_{1:\mathcal{M}}), \text{ and} \\ E &= H(X_{1:\mathcal{M}-1}) - \mathcal{M}H(X_{1:\mathcal{M}}), \end{aligned} \quad (\text{A5})$$

we obtain the structure of the information-theoretic definition (25).

Appendix B. Decomposition of the joint Fisher information

The proof of the decomposition of the joint Fisher information in this appendix is analogous to that in [33], but generalized to the case of multi-parameter dependence. Let us consider definition (4) of the Fisher information matrix for two random variables X_1, X_2 :

$$\begin{aligned} [F_{X_1 X_2}(\theta)]_{ij} &= \sum_{x_1 x_2} P_\theta(x_1, x_2) [\partial_i \log P_\theta(x_1, x_2) \partial_j \log P_\theta(x_1, x_2)] = \sum_{x_1 x_2} \frac{P_\theta(x_1)}{P_\theta(x_2|x_1)} \partial_{\theta_i} P_\theta(x_2|x_1) \partial_{\theta_j} P_\theta(x_2|x_1) \\ &+ \sum_{x_1 x_2} \partial_{\theta_i} P_\theta(x_2|x_1) \partial_{\theta_j} P_\theta(x_1) + \sum_{x_1 x_2} \partial_{\theta_i} P_\theta(x_1) \partial_{\theta_j} P_\theta(x_2|x_1) + \sum_{x_1 x_2} \frac{P_\theta(x_2|x_1)}{P_\theta(x_1)} \partial_{\theta_i} P_\theta(x_1) \partial_{\theta_j} P_\theta(x_2|x_1). \end{aligned} \quad (\text{B1})$$

Now, notice that the second and the third terms are zero. In fact, for the second:

$$\sum_{x_1 x_2} \partial_{\theta_i} P_\theta(x_2|x_1) \partial_{\theta_j} P_\theta(x_1) = \sum_{x_1} x_1 \partial_{\theta_j} P_\theta(x_1) \partial_{\theta_i} \sum_{x_2} P_\theta(x_2|x_1), \quad (\text{B2})$$

but by normalization of probabilities we have that

$$\sum_{x_2} P_\theta(x_2|x_1) = 1 \implies \partial_{\theta_i} \sum_{x_2} P_\theta(x_2|x_1) = 0. \quad (\text{B3})$$

In the same way, it is easy to prove that the third term vanishes. Hence, we have:

$$\begin{aligned} [F_{X_1 X_2}(\theta)]_{ij} &= \sum_{x_1 x_2} \frac{P_\theta(x_1)}{P_\theta(x_2|x_1)} \partial_{\theta_i} P_\theta(x_2|x_1) \partial_{\theta_j} P_\theta(x_2|x_1) + \sum_{x_1 x_2} \frac{P_\theta(x_2|x_1)}{P_\theta(x_1)} \partial_{\theta_i} P_\theta(x_1) \partial_{\theta_j} P_\theta(x_2|x_1) \\ &= [F_{X_1}(\theta)]_{ij} + [F_{X_2|X_1}(\theta)]_{ij}. \end{aligned} \quad (\text{B4})$$

The multi-partite case can be proven by iteratively following the same procedure. As a direct consequence, in case of i.i.d. multi-parametric estimation, we have simply:

$$F_{X_1 X_2}(\theta) = F_{X_1}(\theta) + F_{X_2}(\theta). \quad (\text{B5})$$

Appendix C. Maximum likelihood estimation

In this appendix, we present a brief review of MLE, following mainly the book by Kay [1], and give some indications about its application to parameter estimation on finite-Markov order processes on a finite sample space.

Let us consider a parameter-dependent probability distribution $P_\theta(X_{1:N})$, and define the function (of θ):

$$\ell(\theta) = P_\theta(X_{1:N}), \quad (C1)$$

which is called *likelihood* function. The value of the parameter θ proposed by the MLE estimator is given by the maximization of $\ell(\theta)$ over the entire space in which θ can take values. This corresponds to choosing the value of θ which makes the current outcomes as *likely* as possible.

Notice that the maximising θ does not change if a monotonous function is applied to $\ell(\theta)$; it is therefore common to maximize the *log-likelihood* $L(\theta) = \log \ell(\theta)$.

When dealing with processes with finite Markov order \mathcal{M} ($\mathcal{M} \ll N$) and a finite sample space, the computation of the MLE is greatly eased by making use of the following observations. First, exploiting the definition of the Markov order (20), we can rewrite the log-likelihood as:

$$L(\theta) = \sum_{k=\mathcal{M}+1}^N \log P_\theta(X_k|X_{k-\mathcal{M}:k-1}) + \log P_\theta(X_{\mathcal{M}}|X_{1:\mathcal{M}-1}) + \dots + \log P_\theta(X_2|X_1) + \log P_\theta(X_1). \quad (C2)$$

Notice that all the terms in the summation have the same form: the probability of a certain outcome in the sequence, conditioned on exactly the previous \mathcal{M} . The remaining terms, conditioned on a smaller number of random variables, give a negligible contribution in the $N \gg \mathcal{M}$ limit. Let us now group together the items in the sequence in groups of homogeneous size $\mathcal{M} + 1$, with a sliding window approach:

$$\eta_1 = (X_1, X_2, \dots, X_{\mathcal{M}}, X_{\mathcal{M}+1}); \eta_2 = (X_2, X_3, \dots, X_{\mathcal{M}+1}, X_{\mathcal{M}+2}) \dots \quad (C3)$$

Each of the η sequences is associated to a conditional probability in the form:

$$C_{\eta_j} = P_\theta(X_{j+\mathcal{M}}|X_{j:\mathcal{M}-1}). \quad (C4)$$

Notice that, if any of the variables X_j can take up to d different values, then there are only $z = (\mathcal{M} + 1)^d$ possible different groups η_j ; the summation can therefore be transferred to the different blocks:

$$L(\theta) = \sum_{q=1}^z f_q \log C(q) + \mathcal{R}, \quad (C5)$$

where f_q denotes the number of η groups having the same form, indexed by q , in the sequence, $C(q)$ is the associated conditional probability, and \mathcal{R} denotes the negligible remaining terms. In the common situation $(\mathcal{M} + 1)^d \ll N$, this allows us to compute the likelihood with a significantly enhanced efficiency with respect to a direct attempt.

We recall also some relevant properties of the MLE. Let us first consider a set $X_{1:N}$ of i.i.d. random variables, extracted from the same θ -dependent distribution $P_\theta(X)$. Under mild regularity conditions on P_θ , it can be proven that [38]:

- the MLE estimation is *consistent*, meaning that, in the limit $N \rightarrow \infty$, the MLE estimator $\hat{\theta}$ is such that $\hat{\theta} \rightarrow \theta^*$, where θ^* is the *true* value of the parameter;
- the MLE estimator is *asymptotically normal*, meaning that the random variable $\hat{\theta} - \theta^*$ (where randomness is given by the extraction of different datasets) tends to be distributed according to a Gaussian distribution, with variance given by the Cramér-Rao bound in the form (10).

These propositions were originally formulated for i.i.d. random variables. For the purposes of this paper, it is however fundamental to consider non-i.i.d. random variables. It has been proven [53] that, under appropriate regularity conditions, they can be extended to Markov chains, once the marginal Fisher information $F_1(\theta)$ is replaced by the conditional Fisher information $F_{2|1}(\theta)$. Furthermore, a process with a finite sample space and of finite Markov order \mathcal{M} can be directly translated into a Markov order 1 process by putting together groups of \mathcal{M} random variables, in a similar fashion as it is done in the transfer matrix approach, see appendix G. This allows the extension of the result to larger Markov orders.

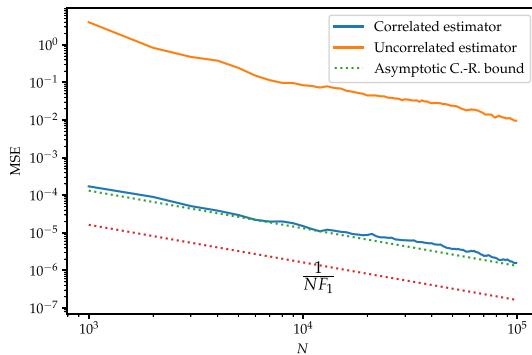


Figure 6. An estimation task for both the regular and the uncorrelated MLE, on the process defined by the transition matrix (37); the MSE is computed by averaging over 50 repetitions.

Appendix D. On the meaning of sub-additive Fisher information

When dealing with a process yielding a sub-additive Fisher information, it could seem that, by using some estimator able to ignore the correlations, the experimenter could be able to achieve a precision beyond the Cramér-Rao bound. We illustrate the situation with an example.

Consider the Markov process defined by the transition matrix (37); we know that it yields a sub-additive Fisher information for θ . Being a $\mathcal{M} = 1$ process, its genuine MLE has to take into account the correlations; the likelihood function for a sequence of outcomes $x_{0:N}$ has the form

$$\ell(\theta) = P_\theta(x_N|x_{N-1})P_\theta(x_{N-1}|x_{N-2})\dots P_\theta(x_1|x_0)P_\theta(x_0). \quad (\text{D1})$$

We could now define an *uncorrelated* MLE, taking into account only the marginal probabilities, as:

$$\ell'(\theta) = P_\theta(x_N)P_\theta(x_{N-1})\dots P_\theta(x_1)P_\theta(x_0). \quad (\text{D2})$$

Naively, one could think that this uncorrelated MLE should result in a higher precision, with an MSE converging to the i.i.d. version of the Cramér-Rao bound (10) instead of the correlated version (36). However, this is not verified, as shown in figure 6.

This gives an interesting perspective on the meaning of the Fisher information and of the Cramér-Rao bound: they are related to the nature of the *process*, not to the way in which the estimator is designed. In other terms, there is no way of taking advantage of the sub-additivity of the Fisher information to obtain a precision beyond the Cramér-Rao bound.

Appendix E. The Central Limit Theorem for correlated random variables

The usual formulations of the Central Limit Theorem [54] consider i.i.d. variables. Many different possible generalizations of the theorem, relaxing the independence hypothesis have been presented; in this appendix we discuss one that is directly applicable to most cases of physical interest.

Let X_1, X_2, \dots be a stationary stochastic process, and let us define the coefficients α_n by the inequality

$$|P(X_0 \cap X_n) - P(X_0)P(X_n)| \leq \alpha_n. \quad (\text{E1})$$

In most cases, we will have that $\lim_{n \rightarrow \infty} \alpha_n = 0$, meaning that two random variables in the stochastic process become increasingly independent for an increasing distance between them. If $\alpha_n \rightarrow 0$, the process is called α -mixing. It is possible to prove for instance that every Markov process is α -mixing, and that α_n decays exponentially on it.

Let us now consider an α -mixing process such that:

- $\alpha_n = \mathcal{O}(n^{-5})$
- the expectation value $\mathbb{E}[X_0^{12}]$ is finite (hence, due to stationarity, it is finite for all random variables composing the process);

then the results of the Central Limit Theorem can be recovered [54].

Notice that the first hypothesis is easily fulfilled by any process exhibiting exponential decay of the α_n coefficients, such as in the case of Markov processes; the second one is trivially true for any process on a finite

space, and also for processes on continuous sample spaces with an exponential depression of the distribution tails, such as normal distributions.

Appendix F. A Gaussian Markov process

In order to verify our claims about the role of the correlation sign in the scaling of the Fisher information, we considered in figure 3(a) Gaussian process with $\mathcal{M} = 1$ (i.e. a Gaussian Markov chain). In this section, we will detail its construction. Furthermore, we will use it to illustrate the lack of a relation between the conditional independence required by the definition of Markov order and the independence as defined from the covariances.

A stationary Markov process is fully determined by the joint probability $P(X_1 X_2)$, in terms of which all the relevant conditional and marginal probabilities can be written. We assume this two-sites joint distribution being Gaussian with the form:

$$P(x_1, x_2) \sim \exp \left[-\frac{1}{2} (\vec{x} - \vec{\mu})^T \sigma^{-1} (\vec{x} - \vec{\mu}) \right], \quad (\text{F1})$$

where $\vec{x} = (x_1, x_2)^T$, $\vec{\mu} = (\mu, \mu)^T$, and the normalization factors have been dropped. Here σ is the two-sites covariance matrix, written in terms of the two parameters $\gamma_0 \in \mathbb{R}_+$ and $\rho \in [-1, 1]$ as

$$\sigma = \gamma_0 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}. \quad (\text{F2})$$

Notice that the sign of the nearest-neighbour correlations is encoded in the sign of ρ . It is necessary to determine the conditional probability distribution $P(X_2|X_1)$, which we want to sample for the simulation. Like the joint probability, it is also Gaussian, with mean and variance given respectively by:

$$\begin{aligned} E[X_2|X_1] &= \mu + \rho(X_1 - \mu) \\ \text{Var}[X_2|X_1] &= \gamma_0 - \rho^2 \gamma_0. \end{aligned} \quad (\text{F3})$$

Using these two relations, we are able to iteratively simulate a Gaussian Markov process.

The process built in this way clearly has Markov order 1 by construction, because only the most recent element from the past plays a role in the probability of the future. The full covariance matrix for this process turns out to be

$$\sigma = \gamma_0 \begin{pmatrix} 1 & \rho & \rho^2 & \rho^3 & \dots \\ \rho & 1 & \rho & \rho^2 & \dots \\ \rho^2 & \rho & 1 & \rho & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \quad (\text{F4})$$

It is rather instructive to compare this behaviour with another possible definition of a Gaussian process; let us consider the Gaussian joint probability distribution with mean vector $\vec{\mu} = (\mu, \mu, \dots, \mu)$ and covariance matrix

$$\sigma = \gamma_0 \begin{pmatrix} 1 & \rho & 0 & 0 & \dots \\ \rho & 1 & \rho & 0 & \dots \\ 0 & \rho & 1 & \rho & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \quad (\text{F5})$$

In this case, the covariances for a distance larger than 1 are exactly zero; however, computing explicitly the conditional probabilities, we find that a probability of the form $P(X_3|X_2 X_1)$ depends non-trivially on X_1 , hence it is necessarily different from $P(X_3|X_2)$: despite its appearance, equation (F5) does not have Markov order 1. Instead, its Markov order is infinite. The process with actual Markov order 1 is that given by equation (F4).

Appendix G. Marginalization of spin chains

A detailed description of the method to obtain the marginals has been provided in [44]; here we only give a brief review, considering for the sake of simplicity a $R = 1$ spin chain with local dimension $d = 2$ (a spin- $1/2$ Ising chain).

The Gibbs state (50) of N spins can be conveniently rewritten in terms of an iterated product of components of the so-called *transfer matrix* V (note that periodic boundary conditions are assumed):

$$p(s_1, \dots, s_N) = \frac{1}{Z} V(s_1, s_2) V(s_2, s_3) \dots V(s_N, s_1), \quad (\text{G1})$$

where the transfer matrix is defined as [55]:

$$V = \begin{bmatrix} \exp[(J_1 + B)/T] & \exp[-J_1/T] \\ \exp[-J_1/T] & \exp[(J_1 - B)/T] \end{bmatrix}, \quad (\text{G2})$$

and it is well-known that the partition function Z of N spins can be written as:

$$Z = \text{Tr}[V^N]. \quad (\text{G3})$$

We would now marginalize $p(s_1, \dots, s_N)$ over all the spins but the first m , i.e. to compute:

$$p(s_1, \dots, s_m) = \sum_{s_{m+1}, \dots, s_N} p(s_1, \dots, s_N). \quad (\text{G4})$$

Note that, since the periodic boundary conditions grant translational invariance, the obtained expression is the general form of an m -spin marginal, and not specific to the first m spins. In terms of the transfer matrix expression (G1) we have therefore:

$$p(s_1, \dots, s_m) = \frac{1}{Z} V(s_1, s_2) \cdot \dots \cdot V(s_{m-1}, s_m) \cdot \sum_{\tilde{s}_{m+1}, \tilde{s}_N} V(s_m, \tilde{s}_{m+1}) V(\tilde{s}_{m+1}, \tilde{s}_{m+2}) \cdot \dots \cdot V(\tilde{s}_{N-1}, \tilde{s}_N) V(\tilde{s}_N, s_1), \quad (\text{G5})$$

but the summation over the spins $\tilde{s}_{m+1}, \dots, \tilde{s}_N$ is equivalent [56] to a matrix multiplication. Thus:

$$p(s_1, \dots, s_m) = \frac{1}{Z} V(s_1, s_2) \cdot \dots \cdot V(s_{m-1}, s_m) \cdot [V^{N-m+1}]_{s_1, s_m} = \frac{1}{Z} [V^{N-m+1}]_{s_1, s_m} \prod_{j=1}^{m-1} V(s_j, s_{j+1}). \quad (\text{G6})$$

This expression has been obtained without any approximation, and holds for any value of N . However, we are interested in the thermodynamic limit $N \rightarrow \infty$. It is noteworthy that only the dominant eigenvalue λ of the transfer matrix and the corresponding eigenvector u play a role in this limit, yielding the final expression:

$$p(s_1, \dots, s_m) = \frac{[u]_{s_1} [u]_{s_m}}{\lambda^{m-1}} \prod_{j=1}^{m-1} V(s_j, s_{j+1}). \quad (\text{G7})$$

A procedure to generalize the transfer matrix to higher- d or higher- R chains is detailed in [56], and involves the creation of a $d^R \times d^R$ -dimensional matrix, whose components are denoted by multi-indices built on R spins (here we will only put the relevant indices one after the other, like $s_i s_j$). The obtained transfer matrix is in general non-symmetric, and it is hence not possible to guarantee that it can be fully diagonalized. However, it can be proven [44, 57] that the dominant eigenvalue of V still dominates the trace (and the partition function, according to (G3)); the corresponding left and right eigenvectors (u^L and u^R respectively) have non-negative components.

It is important here to mention the fact that, unlike what happens in the $R = 1$ scenario, when dealing with longer ranges R the transfer matrix only permits the direct computation of marginals of $m = kR$, $k \in \mathbb{N}$ spins. Any marginal consisting of a number of spins which is not a multiple of R has to be computed by manual marginalization of a larger marginal. The general expression for a $m = kR$ marginal of a range- R spin chain is:

$$P_T(s_1, \dots, s_{kR}) = \frac{[u^L]_{s_1 \dots s_R} [u^R]_{s_{(k-1)R+1} \dots s_{kR}}}{\lambda^{k-1}} \cdot \prod_{j=0}^{k-2} V(s_{jR+1} \dots s_{(j+1)R}; s_{(j+1)R+1} \dots s_{(j+2)R}). \quad (\text{G8})$$

Appendix H. Markov order of spin chains

The proof of $\mathcal{M} \leq R$ is fully general, and has been first presented in [45]. Let R be the interaction range, and consider the calculation of a marginal over L spins, where $L = RL'$ for a suitable $L' \in \mathbb{N}$. Let us introduce a multi-index notation to describe the behaviour of a whole block of R spins as:

$$\eta_i = s_{Ri}s_{Ri+1}\dots s_{R(i+1)-1}. \quad (\text{H1})$$

The marginal probability distribution for a sequence of length L is then given by:

$$P(s_0, \dots, s_{L-1}) = \frac{u_{\eta_{L'-1}}^{\mathcal{R}} u_{\eta_0}^{\mathcal{L}}}{\lambda^{L'-1}} \prod_{i=1}^{L'-2} V(\eta_i, \eta_{i+1}). \quad (\text{H2})$$

Let us now consider two sequences, both of length L , centred in position 0, called $\overleftarrow{s}^L = s_{-L}s_{-L+1}\dots s_{-1}$ and $\overrightarrow{s}^L = s_0s_1\dots s_{L-1}$. We want to express the conditional probability:

$$P(\overrightarrow{s}^L | \overleftarrow{s}^L) = \frac{P(\overrightarrow{s}^L, \overleftarrow{s}^L)}{P(\overleftarrow{s}^L)}. \quad (\text{H3})$$

Using the marginal probability detailed before:

$$P(\overrightarrow{s}^L | \overleftarrow{s}^L) = \frac{u_{\eta_{L'-1}}^{\mathcal{R}}}{\lambda^{L'} u_{\eta_{-1}}^{\mathcal{R}}} \prod_{i=-1}^{L'-2} V(\eta_i, \eta_{i+1}), \quad (\text{H4})$$

from which it is apparent that of all the spin blocks in \overleftarrow{s}^L , the conditional probability only depends on the single spin block η_{-1} , since it is the only one which appears in the expression. One has finally:

$$P(\overrightarrow{s}^L | \overleftarrow{s}^L) = P(\overrightarrow{s}^L | \eta_{-1}). \quad (\text{H5})$$

The thesis follows, because η_{-1} has length R .

In this paper, we have only considered nearest-neighbours chains, with $R = 1$. It is rather easy to prove for such chains that, for $T \neq 0, \infty$, $\mathcal{M} = 1$. Notice that the condition $R = 1$ implies that $J_1 \neq 0$, by definition of R .

Proceeding by contradiction, let us assume that the Markov order is smaller than the interaction range, hence $\mathcal{M} = 0$, $R = 1$. Therefore the spins behave independently and the conditionals can be dropped:

$$P(s_2 | s_1) = P(s_1), \quad (\text{H6})$$

or, equivalently, the probability factorizes, as

$$P(s_1 s_2) = P(s_1)P(s_2). \quad (\text{H7})$$

The expression for both the two-sites joint probabilities $P(s_1 s_2)$ and the one-site marginals are analytically known (but lengthy). Plugging them into (H7) and solving for J , the only possible solution we obtain is $J = 0$, which contradicts the assumption $R = 1$. Hence, for the Ising chain we have, at finite temperature, $\mathcal{M} = R$.

Appendix I. Advantage and disadvantage on NNN spin chains

In this appendix, we extend the spin chain thermometry from the main body to a higher coupling order. In particular, we analyse a NNN spin chain with a Hamiltonian of the form

$$\hat{H} = -B \sum_j \hat{\sigma}_j^z - J \sum_j \hat{\sigma}_j^z \hat{\sigma}_{j+1}^z - \alpha J \sum_j \hat{\sigma}_j^z \hat{\sigma}_{j+2}^z, \quad (\text{H1})$$

where α represents the relative strength between the first and the second-order couplings. In this case, it is more instructive to discuss the advantage in terms of a difference instead of a ratio: let $\Delta F = F_{1:2} - 2F_1$. This choice is due to the fact that, for some of the plots represented, the value of F_1 is so small that the ratio ξ would become difficult to represent, even on a logarithmic scale. Metrological advantage corresponds to $\Delta F > 0$, disadvantage to $\Delta F < 0$. Figure 7 shows the behaviour of ΔF for different regimes of J and B .

We notice that, even though the NNN spin chain exhibits a richer phenomenology than the NN Ising chain, the relevant points are the ones already mentioned in the previous parts. Again, for $J > 0$ (primary ferromagnetic), one has disadvantage, for $J < 0$ (primary anti-ferromagnetic) advantage arises. In presence of a very strong external magnetic field, alignment is also enforced in the otherwise anti-ferromagnetic chain; in this case, the behaviour is similar to the ferromagnetic case, with $\Delta F < 0$. Overall, weaker coupling (relative to B) leads to weaker ΔF .

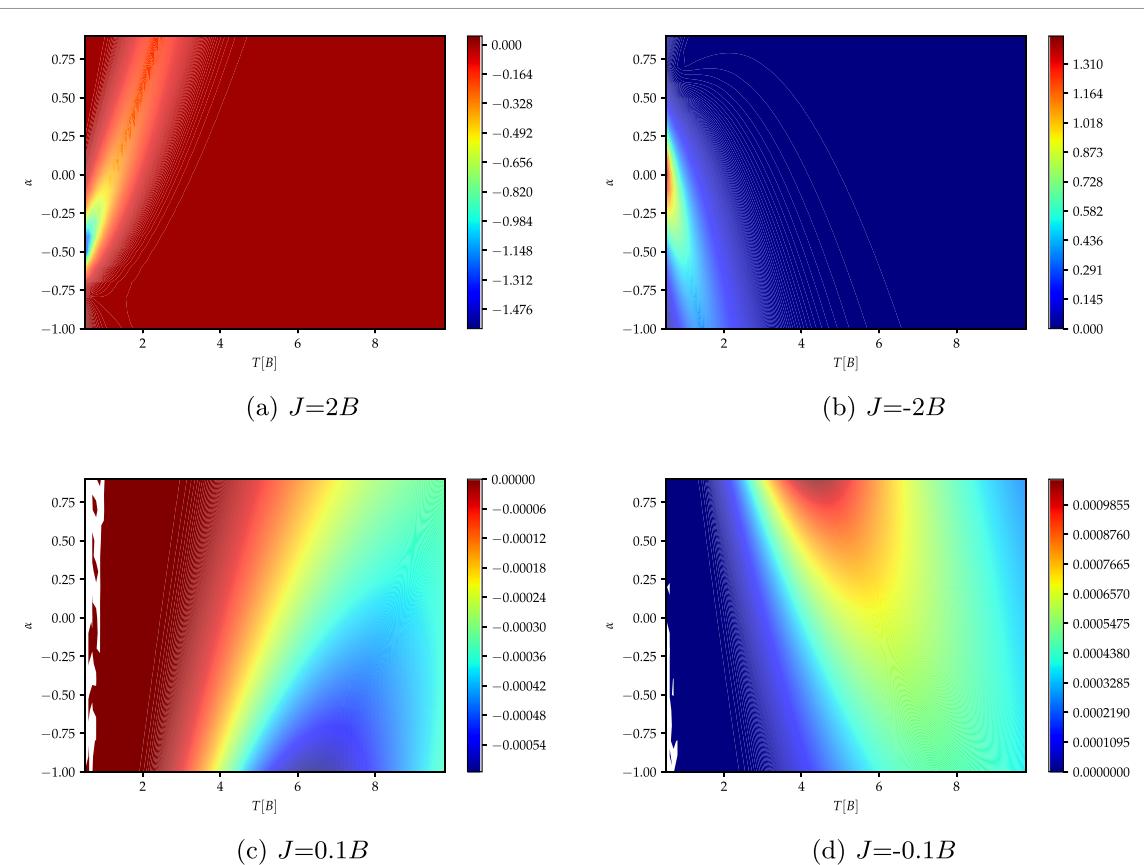


Figure 7. Behaviour of ΔF for next-to-nearest-neighbour for different values of the first-order coupling parameter J , as a function of temperature T and relative coupling strength α . Negative values (disadvantage) occur for all cases of positive correlation (panels (a) and (c)); positive values (advantage) occur in all cases with negative correlations (panels (b) and (d)).

ORCID iDs

Marco Radaelli  <https://orcid.org/0000-0002-9210-5779>

Gabriel T Landi  <https://orcid.org/0000-0001-8451-9712>

Kavan Modi  <https://orcid.org/0000-0002-2054-9901>

Felix C Binder  <https://orcid.org/0000-0003-4483-5643>

References

- [1] Kay S M 1993 *Fundamentals of Statistical Signal Processing: Estimation Theory* (Englewood Cliffs, NJ: Prentice-Hall)
- [2] Van Trees H L 2004 *Detection, Estimation and Modulation Theory, Part I: Detection, Estimation and Linear Modulation Theory* (New York: Wiley)
- [3] Cover T M 1999 *Elements of Information Theory* (New York: Wiley)
- [4] Demkowicz-Dobrzański R, Górecki W and Guță M 2020 *J. Phys. A: Math. Theor.* **53** 363001
- [5] Fisher R A 1922 *Phil. Trans. R. Soc. A* **222** 309
- [6] Cramér H 1946 *Mathematical Methods of Statistics* (Princeton, NJ: Princeton University Press)
- [7] Crutchfield J P and Young K 1989 *Phys. Rev. Lett.* **63** 105
- [8] Shalizi C R and Crutchfield J P 2001 *J. Stat. Phys.* **104** 817
- [9] Crutchfield J P 2012 *Nat. Phys.* **8** 17
- [10] James R G, Mahoney J R, Ellison C J and Crutchfield J P 2014 *Phys. Rev. E* **89** 042135
- [11] Jacobs K 2014 *Quantum Measurement Theory and Its Applications* (Cambridge: Cambridge University Press)
- [12] Wiseman H M and Milburn G J 2009 *Quantum Measurement and Control* (Cambridge: Cambridge University Press)
- [13] Zhou X, Dotsenko I, Peaudecerf B, Rybarczyk T, Sayrin C, Gleyzes S, Raimond J M, Brune M and Haroche S 2012 *Phys. Rev. Lett.* **108** 243602
- [14] Jasperse M, Kewming M J, Fischer S N, Pakkiam P, Anderson R P and Turner L D 2017 *Phys. Rev. A* **96** 063402
- [15] Anderson R P, Kewming M J and Turner L D 2018 *Phys. Rev. A* **97** 013408
- [16] Goan H-S, Milburn G J, Wiseman H M and Sun H B 2001 *Phys. Rev. B* **63** 125326
- [17] Nirmal M, Dabbousi B O, Bawendi M G, Macklin J J, Trautman J K, Harris T D and Brus L E 1996 *Nature* **383** 802
- [18] Yuan G, Gómez D E, Kirkwood N, Boldt K and Mulvaney P 2018 *ACS Nano* **12** 3397
- [19] Munoz R N, Frazer L, Yuan G, Mulvaney P, Pollock F A and Modi K 2021 arXiv:2106.12179v2 [cond-mat.mes-hall]
- [20] Gammelmark S and Mølmer K 2013 *Phys. Rev. A* **87** 032115
- [21] Gammelmark S and Mølmer K 2014 *Phys. Rev. Lett.* **112** 170401
- [22] Kiilerich A H and Mølmer K 2014 *Phys. Rev. A* **89** 052110

[23] Kiilerich A H and Mølmer K 2016 *Phys. Rev. A* **94** 032103

[24] Burgarth D, Giovannetti V, Kato A N and Yuasa K 2015 *New J. Phys.* **17** 113055

[25] Godley A and Guta M 2022 arXiv:2204.08964v1 [quant-ph]

[26] Ilias T, Yang D, Huelga S F and Plenio M B 2022 *PRX Quantum* **3** 010354

[27] Seah S, Nimmrichter S, Grimmer D, Santos J P, Scarani V and Landi G T 2019 *Phys. Rev. Lett.* **123** 180602

[28] Strasberg P 2020 *Quantum* **4** 240

[29] Carlen E A 1991 *J. Funct. Anal.* **101** 194

[30] Kagan A and Landsman Z 1997 *Stat. Probab. Lett.* **32** 175

[31] Ly A, Marsman M, Verhagen J, Grasman R and Wagenmakers E-J 2017 A tutorial on Fisher information *J. Math. Psychol.* **80** 40–55

[32] Rao C R 1973 *Linear Statistical Inference and Its Applications* (New York: Wiley)

[33] Micadei K, Rowlands D A, Pollock F A, Celeri L C, Serra R M and Modi K 2015 *New J. Phys.* **17** 023057

[34] Zegers P 2015 *Entropy* **17** 4918

[35] Taranto P, Pollock F A, Milz S, Tomamichel M and Modi K 2019 *Phys. Rev. Lett.* **122** 140401

[36] James R G, Mahoney J R, Ellison C J and Crutchfield J P 2014 The many roads to synchrony: natural time scales and their algorithms *Phys. Rev. E* **89** 042135

[37] Crutchfield J P and Feldman D P 2003 *Chaos* **13** 25

[38] Van der Vaart A W 2000 *Asymptotic Statistics* vol 3 (Cambridge: Cambridge University Press)

[39] Kurdziałek S and Demkowicz-Dobrzański R 2021 *J. Opt.* **23** 075701

[40] Ising E 1925 *Z. Phys.* **31** 253

[41] Correa L A, Mehboudi M, Adesso G and Sanpera A 2015 *Phys. Rev. Lett.* **114** 220405

[42] Hovhannisyan K V, Jørgensen M R, Landi G T, Alhambra A M, Brask J B and Perarnau-Llobet M 2021 *PRX Quantum* **2** 020322

[43] Kramers H A and Wannier G H 1941 *Phys. Rev.* **60** 252

[44] Feldman D P 1998 Computational mechanics of classical spin systems *PhD Thesis* University of California, Davis

[45] Feldman D P and Crutchfield J P 1998 Discovering noncritical organization: statistical mechanical, information theoretical, and computational views of patterns in one-dimensional spin systems *Working Paper* 98-04-026 (Santa Fe Institute) (available at: www.santafe.edu/research/results/working-papers/discovering-noncritical-organization-statistical-m)

[46] Taranto P, Milz S, Pollock F A and Modi K 2019 *Phys. Rev. A* **99** 042108

[47] Landi G T, Paternostro M and Belenchia A 2022 *PRX Quantum* **3** 1

[48] Ciccarello F, Lorenzo S, Giovannetti V and Palma G M 2022 *Phys. Rep.* **954** 1

[49] Fawzi O and Renner R 2015 *Commun. Math. Phys.* **340** 575

[50] Kato K and Brandão F G S L 2019 *Commun. Math. Phys.* **370** 117

[51] Mehboudi M, Sanpera A and Correa L A 2019 *J. Phys. A: Math. Theor.* **52** 303001

[52] Pasquale A D and Stace T M 2018 Quantum thermometry *Thermodynamics in the Quantum Regime: Fundamental Aspects and New Directions* Fundamental Theories of Physics 195 (Cham: Springer) pp 503–27

[53] Billingsley P 1961 *Statistical Inference for Markov Processes* (Chicago, IL: University of Chicago Press)

[54] Billingsley P 1995 *Probability and Measure* (New York: Wiley)

[55] Yeomans J M 1992 *Statistical Mechanics of Phase Transitions* (Oxford: Clarendon)

[56] Dobson J F 1969 *J. Math. Phys.* **10** 40

[57] Ruelle D 1999 *Statistical Mechanics: Rigorous Results* (Singapore: World Scientific)