

Bayesian quantum thermometry based on thermodynamic lengthMathias R. Jørgensen ^{1,*}, Jan Kołodziej ², Mohammad Mehboudi ³, Martí Perarnau-Llobet ³ and Jonatan B. Brask ¹¹*Department of Physics, Technical University of Denmark, 2800 Kongens Lyngby, Denmark*²*Centre for Quantum Optical Technologies, Centre of New Technologies, University of Warsaw, 02-097 Warsaw, Poland*³*Département de Physique Appliquée, Université de Genève, 1211 Geneva, Switzerland*

(Received 1 September 2021; revised 10 January 2022; accepted 8 March 2022; published 1 April 2022)

In this work we propose a theory of temperature estimation of quantum systems, which is applicable in the regime of non-negligible prior temperature uncertainty and limited measurement data. In this regime the problem of establishing a well-defined measure of estimation precision becomes nontrivial. Furthermore, the construction of a suitable criterion for optimal measurement design must be reexamined to account for the prior uncertainty. We propose a fully Bayesian approach to temperature estimation based on the concept of thermodynamic length, which solves both these problems. As an illustration of this framework, we consider thermal spin-1/2 particles and investigate the fundamental difference between two cases: on the one hand, when the spins are probing the temperature of a heat reservoir and, on the other, when the spins themselves constitute the sample.

DOI: [10.1103/PhysRevA.105.042601](https://doi.org/10.1103/PhysRevA.105.042601)**I. INTRODUCTION**

Measuring the temperature of a physical system is a fundamental task in science and technology. At the micro- and nanoscale in particular, highly precise temperature measurements are essential for a large number of current experiments. Examples include real-time monitoring of temperature profiles within living organisms, e.g., utilizing color centers in nanodiamonds [1–3], the preparation of ultracold atoms in optical lattices, phase diagrams and exploring transport phenomena [4–10], and studies of quantum thermodynamic phenomena in microelectronic devices [11–15]. Temperature is not a directly measurable property of a system, and in contrast to, e.g., interferometry, phase estimation, or electromagnetic-field sensing [16,17], thermometry is further complicated by the fact that temperature is also not a Hamiltonian-encoded parameter. Rather, the temperature of a system is an entropic quantity which must be estimated indirectly from the statistical behavior of a variable which can be observed directly. The purpose of the theory of quantum thermometry is both to guide the design of optimal measurement processes, i.e., building good thermometers in the quantum regime, and to optimally infer from the acquired measurement data the underlying temperature [18,19].

The majority of previous works on quantum thermometric theory, with the notable exception of the recent studies in [20–22], have focused on local point estimation [23], termed for short the local paradigm, in which measurements are designed to detect small variations around a known temperature value [18,19]. Within the local paradigm, the expected precision of a temperature estimate is typically quantified by the frequentist mean-square error, with the associated signal-to-noise ratio providing a meaningful notion of relative error [24]. Given that certain conditions are satisfied,

e.g., that the temperature estimate is unbiased, the frequentist mean-square error is lower bounded, as well as typically well approximated, by the so-called Cramér-Rao bound [25–27]. Furthermore, optimal measurements applicable in the asymptotic (large-data-set) regime can be identified via the local optimization of the Cramér-Rao bound.

The motivation for constructing a theory applicable beyond the local paradigm is twofold: (i) It is typically an unjustified assumption that the temperature to be estimated is known with sufficient precision *a priori* to justify working within the local paradigm and (ii) the optimal measurement protocol generally depends on the prior temperature information and cannot be identified via an optimization of the asymptotic Cramér-Rao bound. Avoiding the restrictions of the local paradigm, i.e., providing a general approach to quantifying thermometric performance, and designing optimal measurements, under conditions of non-negligible prior uncertainty, requires a fully Bayesian framework [28,29].

In this work we develop a theory of Bayesian quantum thermometry, applicable for any amount of prior information, which is based on the concept of a thermodynamic length [30–33]. The basic idea is that a meaningful measure of thermometric precision should be based on the ability to distinguish states at different temperatures, i.e., colder from hotter, and should be independent of the particular parametrization of the states, e.g., temperature. This can be naturally achieved by introducing a distance function between the thermal states of the sample system considered. Such a distance is exactly the thermodynamic length between thermal states [30–33] and we argue that this choice is singled out by the requirement that a well-defined distance should respect the invariance properties of the sample. An interesting implication of the proposed framework is that any meaningful definition of relative error must be given with respect to the specific sample system considered. This feature is illustrated in Fig. 1. In particular, we find that the standard noise-to-signal ratio, defined in terms of the frequentist mean-square error, is only

*matrj@fysik.dtu.dk

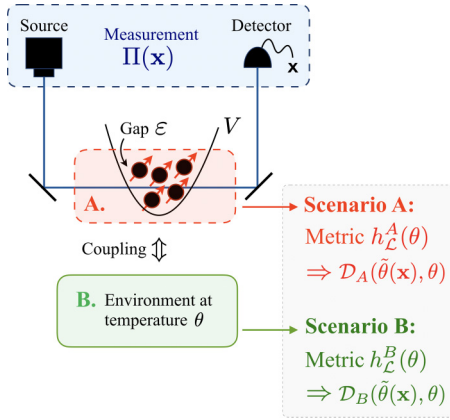


FIG. 1. Illustration of a hypothetical experiment. A collection of spin impurities (scenario A) interacts with a thermal environment (scenario B) at temperature θ ; the spin impurities are measured via an optical setup yielding measurement data \mathbf{x} . An essential message of this paper is that the sample system defines the appropriate distance function and that the distance is the basis of the estimation theory. To illustrate this point we consider two scenarios. In scenario A the spin impurities themselves constitute the sample system of interest, e.g., we are interested in the thermometry of the spin degrees of freedom themselves [8,34]; in this case we obtain a metric h_C^A and a distance \mathcal{D}_A . In scenario B the spin impurities are employed as equilibrium thermometers of the thermal environment [10,35]; in this case we obtain a different metric h_C^B and a different distance \mathcal{D}_B . As a consequence, the relevant figure of merit, as well as the associated optimal protocol, will be different in each sensing process.

recovered as a meaningful relative error, within the local regime, when the considered sample system can be effectively modeled as an ideal heat bath.

For the sake of illustration, we consider a thermometry setting that involves noninteracting spin-1/2 particles and compare the scenario in which the interest is in thermometry of the particles themselves to the scenario in which the particles are employed as probes of an underlying heat reservoir. Specifically, we look at simulated outcomes of projective energy measurements of thermal spin-1/2 particles and compare the computed temperature estimates and measures of precision in the two cases. The example illustrates that the rate of convergence to the local regime and the suitability of various precision bounds depend on the specific scenario considered.

II. BAYESIAN ESTIMATION THEORY

In this section we summarize the key concepts of Bayesian estimation theory [28,29]. The aim of estimation theory, at least when considering point estimation [23], is to provide a prescription for singling out an estimate, serving as a best guess of an unknown quantity of interest and for quantifying a measure of confidence in, or the precision of, the specified estimate. In this work we specifically consider a smooth manifold of quantum states, e.g., thermal Gibbs states, of a sample, depicted schematically as \mathcal{S} in Fig. 2. Our task is to estimate the true state of the sample, assuming that this belongs to the specified manifold.




FIG. 2. Illustration of a parameter space Θ and an alternative parametrization Φ , mapped onto a one-parameter curve of quantum states $\mathcal{S}_\Theta = \mathcal{S}_\Phi$. The distinguishability of parameter values is only meaningful when expressed as a distance \mathcal{D} between the corresponding quantum states, i.e., the state space induces a geometry on the parameter space. The induced geometry may be simpler for certain choices of parametrization. In particular, there may exist a Euclidean parametrization with a flat geometry. By considering the distance directly between quantum states, it follows that the distance is invariant with respect to the specific choice of parametrization of the states.

Within the Bayesian theory, the manifold of states is equipped with a probability distribution and this distribution is updated as measurement data are acquired. An estimate is computed, according to some prescription, based on the updated probability distribution. Providing a measure of confidence in an estimate requires a notion of distance between states on the manifold. Here we first argue for a metric structure on the manifold of states. A distance function can then be constructed as the geodesic length between states. As a result, the confidence in the computed estimate can then be unambiguously quantified via the mean-square distance over the updated probability distribution.

Having established the required notions of estimation theory, we provide a criterion for the optimal measurement design and show how the local estimation theory can be recovered as the asymptotic limit of the Bayesian framework. Moreover, we outline a number of known results on Bayesian Cramér-Rao bounds that are particularly useful. Finally, we discuss how to select an uninformative initial prior probability distribution.

A. Manifold of states and Bayesian updating

Physical systems are typically subject to a set of experimental conditions, e.g., some preparation procedure, which defines a smooth manifold of states [36–38]. To be specific, we consider a quantum system, the sample, and assume that the true, but unknown, state of the sample ρ_{true} belongs to a one-parameter family \mathcal{S}_Θ of quantum states in the manifold

$$\rho_{\text{true}} \in \mathcal{S}_\Theta := \{\tau_\Theta(\theta) \text{ for } \theta \in \Theta\}, \quad (1)$$

where $\tau_\Theta(\theta)$ labels a parametrized quantum state, i.e., a linear operator on the sample Hilbert space, and $\Theta \subseteq \mathbb{R}$ denotes the parameter space, e.g., Θ could be the space of temperatures. Our task is to identify the true sample state, which, given that the state belongs to a one-parameter family, can be formulated as a parameter estimation problem. Throughout, we focus on a

one-dimensional parameter space and leave the generalization to the multidimensional case for future work.

Within Bayesian estimation theory, we start from a prior probability density $p(\theta)$ on the parameter space Θ , which is subsequently updated as the sample state is probed via measurements. Quantum mechanically, any measurement can be represented by a positive-operator-valued measure (POVM). If we suppose that the sample is in the state $\tau_\Theta(\theta)$, then the likelihood $p(x|\theta)$ of observing a measurement outcome x is given by Born's rule

$$p(x|\theta) = \text{Tr}[\Pi(x)\tau_\Theta(\theta)], \quad (2)$$

where $\Pi(x)$ is the POVM element associated with the outcome $x \in X$, with X the outcome space of the measurement performed. The POVM elements must be positive semidefinite, i.e., $\Pi(x) \geq 0$, and satisfy the normalization condition $\int_X dx \Pi(x) = \mathbb{I}$, with \mathbb{I} the identity operator [39]. Conditioned on observing the specific outcome x , we can find the posterior probability distribution using Bayes's rule

$$p(\theta|x) = \frac{p(x|\theta)p(\theta)}{p(x)}, \quad (3)$$

where $p(x) := \int_\Theta d\theta p(x|\theta)p(\theta)$ is the marginal probability density on the space of outcomes. For later use we also define the joint density $p(\theta, x) = p(x|\theta)p(\theta)$. The posterior density represents our degree of belief in different parameter values given the available measurement data. Note that here we formulate Bayes's rule as a single-shot update. If we consider v independent measurements giving outcomes $\mathbf{x} = \{x_1, \dots, x_v\}$, then it is equally valid to compute the posterior conditioned on the full set of observations.

B. Parametrization invariance

Having introduced the manifold of quantum states and discussed probability densities on the parameter space, we must face the subtlety of parametrization invariance. In the above we worked with a parametrization Θ ; however, the manifold of quantum states itself is invariant with respect to the specific choice of parametrization. For example, the manifold of thermal Gibbs states is the same whether we parametrize it using the temperature or the inverse temperature. This fact is illustrated in Fig. 2. In general, we can express this invariance as follows: If we consider a one-to-one mapping $\phi : \Theta \rightarrow \Phi$, where $\Phi \subseteq \mathbb{R}$ is the image of the map, then the function ϕ provides an equally valid parametrization of the manifold of states, i.e., $\mathcal{S}_\Phi \equiv \mathcal{S}_\Theta$ with

$$\mathcal{S}_\Phi = \{\tau_\Phi(\phi) \text{ for } \phi \in \Phi\}, \quad (4)$$

and we explicitly indicate that each state is now given with respect to the parametrization Φ , while the invariance condition reads

$$\phi = \phi(\theta), \quad \tau_\Theta(\theta) = \tau_\Phi(\phi) \forall \phi \quad (5)$$

and expresses the parametrization invariance of every equivalent quantum state. Furthermore, the probability assigned to a given region of state space must be independent of the specific parametrization employed. Thus, to be consistent, the prior probability density must satisfy the invariance condition

[28,40]

$$\phi = \phi(\theta), \quad d\theta p_\Theta(\theta) = d\phi p_\Phi(\phi) \forall \phi. \quad (6)$$

Since the likelihood function $p(x|\theta)$ only depends on the state itself, it is inherently parametrization invariant. Hence, it thus follows that if the above invariance condition holds for the prior density, it will also hold for the posterior density. To simplify our notation, we will for the most part not indicate the parametrization explicitly in what follows. The choice of parametrization is implicitly indicated by the parameter value itself.

C. Metric structure and geodesic length

The essential ingredient required for the development of a parametrization-invariant estimation theory on a Riemannian manifold is the metric structure, i.e., an infinitesimal notion of length, on the manifold of quantum states [40–42]. A metric makes it possible to define the distance between states on the manifold and, more fundamentally, the idea of defining a probability density on a continuous parameter space is not well defined in the absence of such a metric structure [40]. Suppose for a moment that the manifold possesses a metric denoted by g . Then it is always possible to construct a well-defined distance function, as the geodesic length, which we refer to as the metric-based distance, defined by [32,40,42]

$$\mathcal{D}(\tau(\theta_0), \tau(\theta_1)) := \left| \int_{\theta_0}^{\theta_1} d\theta g_\Theta(\theta) \right|, \quad (7)$$

where the quantity $d\theta g_\Theta(\theta)$ provides a parametrization-invariant integration measure, i.e., $d\theta g_\Theta(\theta) = d\phi g_\Phi(\phi)$ for $\phi = \phi(\theta)$. The metric-based distance is a statistical distance between quantum states and it is a parametrization-invariant quantity; this is depicted in Fig. 2. To simplify our notation, in what follows we will write the metric-based distance as $\mathcal{D}(\theta_0, \theta_1)$, referring only to the parameter values. However, it should be kept in mind that this is a distance between quantum states. The form of the geodesic length above is valid for one-parameter problems. More generally, defining the geodesic length involves a minimization over paths [40].

The above way of defining a distance function on the parameter space may still seem ambiguous unless a particular choice of metric can be justified. However, if we consider a reference measurement of the sample with POVM elements $\mathcal{M}(y)$ for $y \in Y$, then a key insight of Bayesian information geometry [32,40,42,43] is that the likelihood associated with this measurement induces a metric of the form $g^2 = h_{\mathcal{M}}$, where $h_{\mathcal{M}}$ is the Fisher information associated with the \mathcal{M} measurement [25,40,42]

$$h_{\mathcal{M}}(\theta) = \int_Y dy p(y|\theta) [\partial_\theta \ln p(y|\theta)]^2. \quad (8)$$

From the point of view of Bayesian probability theory, the Fisher information metric with respect to the reference measurement is the unique Riemannian metric which satisfies both parametrization invariance and any other invariance property of the likelihood function [40,44]. Moreover, according to Chentsov's theorem, any other (monotonic) metric on the parameter space corresponds to the Fisher information metric,

with respect to some reference measurement, up to a multiplicative constant [44,45].

1. Quantum Fisher information metric

The choice of reference measurement represents a degree of freedom, i.e., we can define the metric-based distance and the estimation theory itself, relative to an arbitrary reference measurement of the sample. A natural choice would be to maximize the distance $\mathcal{D}(\theta_0, \theta_1)$ over all possible measurements; however, this procedure does not generally yield a unique reference measurement. To see this, we consider the problem of maximizing the metric at a specific parameter value. This local maximization problem can be solved analytically and yields a metric defined relative to a projective measurement of the symmetric logarithmic derivative $\mathcal{L}(\theta)$ associated with the manifold of states [27]. The symmetric logarithmic derivative (SLD) is defined implicitly by the relation [27]

$$\mathcal{L}(\theta)\tau(\theta) + \tau(\theta)\mathcal{L}(\theta) = 2\partial_\theta\tau(\theta) \quad (9)$$

and is in general a function of the parameter value. The problem alluded to above is then that the SLD defines a natural reference measurement only if the eigenbasis of the SLD is parameter independent, i.e., the projectors must be parameter independent. Note that this need not be the case for the SLD eigenvalues.

If we adopt the SLD reference measurement, then with this choice of reference we obtain the quantum Fisher information (QFI) metric, which gives the maximum-likelihood-induced metric-based distance between sample states. The QFI metric is given by [27]

$$\sqrt{h_{\mathcal{L}}(\theta)} = \text{Tr}[\mathcal{L}(\theta)^2\tau(\theta)]^{1/2}. \quad (10)$$

The resulting metric is equal to four times the Bures metric [25] and is thus directly related to the so-called fidelity between infinitesimally separated states in the manifold. If the family of states considered is a thermal state ensemble, then the metric-based distance is also called the thermodynamic length [30–33]. In the remainder of this paper we will always consider the SLD reference measurement, written \mathcal{L} , and the QFI metric following from that choice.

D. Choosing the correct Bayesian error

Having equipped the manifold with a metric structure and outlined how to construct a metric-based distance function, we can develop a parametrization-invariant estimation theory by considering the mean-square distance as a suitable Bayesian error function [40]. Before proceeding with this development, we want to stress the basic logic of the approach. For a specific thermal sample system, the QFI metric associated with the SLD reference measurement specifies the thermodynamic geometry, i.e., it specifies how to measure the distance between thermal sample states. Only once the states are defined whose temperature should be distinguished (the manifold of states) is it possible to define the correct corresponding Bayesian error function.

1. Mean-square distance within the $\lambda_{\mathcal{L}}$ parametrization

In this section we consider an implemented measurement Π giving measurement data x and the SLD reference measurement \mathcal{L} defining the thermodynamic length. Suppose we denote the estimator of the true parameter value constructed using the measurement data x by $\tilde{\theta}(x)$. Then we may define the mean-square distance (MSD) based on Eq. (7) as

$$\text{MSD}(\tilde{\theta}; x) := \int_{\Theta} d\theta p(\theta|x)\mathcal{D}(\tilde{\theta}(x), \theta)^2, \quad (11)$$

which provides a parametrization-invariant measure of the confidence assigned to the adopted parameter estimate. The mean-square distance is defined with respect to the likelihood function $p(x|\theta)$ associated with the implemented measurement Π and it is defined with respect to the distance function associated with the SLD reference measurement \mathcal{L} . Furthermore, note that the mean-square distance is conditional on the measurement data, i.e., it is a stochastic quantity.

For convenience, we define, implicitly, the function $\lambda_{\mathcal{L}}(\theta)$ as the inverse derivative of the QFI metric associated with the measurement \mathcal{L} as [40]

$$\partial_\theta\lambda_{\mathcal{L}}(\theta) := h_{\mathcal{L}}^{1/2}(\theta). \quad (12)$$

Since the QFI is non-negative, it follows that the function $\lambda_{\mathcal{L}}$ is monotonically increasing. Furthermore, if we consider a change of parametrization $\theta \rightarrow \phi$ and note that under such a transformation the QFI transforms as $h_{\mathcal{L}}(\theta) = [\partial_\theta\phi]^2 h_{\mathcal{L}}(\phi)$, then we see that it follows directly from the definition of $\lambda_{\mathcal{L}}$ that it is a parametrization-invariant quantity, i.e., $\lambda_{\mathcal{L},\Theta}(\theta) = \lambda_{\mathcal{L},\Phi}(\phi)$ for $\phi = \phi(\theta)$. Using the $\lambda_{\mathcal{L}}$ function, it is always possible to express the distance (7) defined based on the QFI metric in the Euclidean form, i.e.,

$$\mathcal{D}(\tilde{\theta}, \theta) = |\lambda_{\mathcal{L}}(\tilde{\theta}) - \lambda_{\mathcal{L}}(\theta)|. \quad (13)$$

This follows directly from an application of the fundamental theorem of calculus for $g_{\Theta}(\theta) = \sqrt{h_{\mathcal{L}}(\theta)}$.

If we constrain ourselves to reference measurements for which the QFI is nonvanishing (except perhaps at isolated points, e.g., the boundaries of the parameter domain), then $\lambda_{\mathcal{L}}$ itself constitutes a valid parametrization of the one-parameter family within the manifold. Referring to the associated parameter space as $\Lambda_{\mathcal{L}}$, it follows that when working in this parametrization the MSD takes the simple form

$$\text{MSD}(\tilde{\lambda}; x) = \int_{\Lambda_{\mathcal{L}}} d\lambda p(\lambda|x)[\tilde{\lambda}(x) - \lambda]^2, \quad (14)$$

where $\tilde{\lambda}(x) = \lambda(\tilde{\theta}(x))$ and we have made use of Eq. (6), i.e., the invariance property of the posterior probability density function. In the above we dropped the subscript \mathcal{L} when referring to the λ parameter itself. Furthermore, from here on we will not explicitly indicate the space which is integrated over, i.e., $\Lambda_{\mathcal{L}}$. Note that the $\Lambda_{\mathcal{L}}$ parametrization is special in that it is associated with QFI equal to unity, i.e., a flat metric. The form of the MSD is then simply the standard Euclidean distance with respect to the parametrization $\Lambda_{\mathcal{L}}$.

The remaining open question is how to choose a suitable estimation function. A natural choice, and the one primarily adopted here, is the mean of λ over the posterior probability

distribution

$$\bar{\lambda}(x) := \int d\lambda p(\lambda|x)\lambda. \quad (15)$$

It can be shown that this choice of estimation function is classically optimal, in the sense that it minimizes the mean-square distance [Eq. (14)] [24,29,40]. For this reason the posterior mean of λ is called the minimal mean-square distance (MMSD) estimator [40]. Notice that since the MSD is a parametrization-invariant quantity, we can then directly obtain the corresponding MMSD estimate of θ as $\bar{\theta}(x) = \lambda^{-1}\bar{\lambda}(x)$, where λ^{-1} denotes the inverse of the function λ .

2. Expected mean-square distance

Having introduced the mean-square distance as a measure of confidence, the next problem is to identify a suitable criterion for optimal measurement design. We immediately face the difficulty that the mean-square distance is a stochastic quantity, i.e., in identifying the optimal measurement one must consider the full ensemble

$$\Xi \equiv \{\text{MSD}(\bar{\lambda}; x) \text{ for } x \in X\}. \quad (16)$$

In an actual experiment we sample this ensemble according to the likelihood $p(x|\theta_{\text{true}})$ with respect to the true quantum state of the sample. Since the true state is unknown *a priori*, we cannot gauge the expected confidence in an estimate resulting from a measurement based on the true likelihood. Instead we must assume that the ensemble Ξ is sampled according to the marginal density $p(x)$, giving the likelihood averaged over the prior density.

In the simplest case we suppose that the ensemble Ξ is well captured by its average behavior, which constitutes a deterministic quantity, typically referred to as the Bayesian mean-squared error [24,29]. As in our work we interpret it as the average evaluated over the marginal density of the MSD (14), we refer to it as the expected mean-square distance (EMSD), i.e.,

$$\text{EMSD}[\bar{\lambda}] := \int dx p(x)\text{MSD}(\bar{\lambda}; x) \quad (17)$$

$$= \int d\lambda dx p(\lambda, x)[\bar{\lambda}(x) - \lambda]^2. \quad (18)$$

In (17) we use square brackets to denote the choice of estimator. Hence, a natural criterion for determining the optimal measurement design is to minimize the EMSD. Note that the optimal measurement will generally be a functional of the prior probability distribution; in other words, the optimal choice of measurement depends on the prior knowledge.

E. Connection to the local paradigm

We now consider the local limit of a narrow posterior distribution. In this limit the EMSD is lower bounded, as well as typically well approximated, by the expected Cramér-Rao bound (ECRB) [26,27]. To derive this bound, we first rewrite the EMSD as

$$\text{EMSD}[\bar{\lambda}] = \int d\lambda p(\lambda)\Delta^2\bar{\lambda}, \quad (19)$$

where we have decomposed the joint probability and defined the frequentist mean-square error (MSE) [24]

$$\Delta^2\bar{\lambda} := \int dx p(x|\lambda)[\bar{\lambda}(x) - \lambda]^2. \quad (20)$$

Adopting an asymptotically unbiased estimator implies that the MSE satisfies the Cramér-Rao bound [24], and upon substitution we obtain the ECRB

$$\begin{aligned} \text{EMSD}[\bar{\lambda}_{\text{ub}}] &\geq \text{ECRB} := \int d\lambda \frac{p(\lambda)}{h_{\Pi}(\lambda)} \\ &= \int d\theta p(\theta) \frac{h_{\mathcal{L}}(\theta)}{h_{\Pi}(\theta)}, \end{aligned} \quad (21)$$

where $\bar{\lambda}_{\text{ub}}$ denotes any unbiased estimator and we recall that h_{Π} is the Fisher information associated with the implemented measurement. The final equality follows from a change of parametrization and an application of Eq. (12). The ECRB is typically tight in the asymptotic limit, i.e., the limit of a large number of measurement repetitions [26].

The expression for the ECRB given here offers a generalization of the results reported in [20,22]; these correspond to $h_{\mathcal{L}}(\theta) \propto 1$ and $h_{\mathcal{L}}(\theta) \propto \theta^{-2}$, respectively. The ECRB shows that the EMSD converges to a generalized version of the relative error averaged over the prior probability, i.e., relative to the metric structure of the manifold. Note that the ECRB shows that the noise-to-signal ratio is only meaningful if the geometry of the problem is scale invariant, i.e., $h_{\mathcal{L}}(\theta) \propto \theta^{-2}$. For general geometries a physically meaningful notion of relative error is given by the mean-square distance. This insight alters the data analysis even in the asymptotic limit.

F. Bayesian Cramér-Rao bounds

As noted above, the expected thermometric performance is naturally quantified via the EMSD and the optimal measurement strategy is the one for which the EMSD is minimal. As an alternative to directly considering the EMSD, we can assess the expected performance by studying bounds on the attainable precision. In many cases this is computationally advantageous. The cost of considering lower bounds is that we do not know in general if the EMSD saturates them.

An advantage of working in the $\Lambda_{\mathcal{M}}$ parametrization is that the EMSD takes the form, as mentioned, of the BMSE. The EMSD can then be directly lower bounded by means of the so-called Van Trees inequalities [29], which are valid when $p(\lambda)$ and $\lambda p(\lambda)$ vanish at the boundaries of $\Lambda_{\mathcal{M}}$. We first consider the tightened Bayesian Cramér-Rao bound (TBCRB) [26,46]

$$\text{EMSD}[\bar{\lambda}] \geq \text{TBCRB} := \int dx p(x)\mathcal{Q}(x)^{-1}, \quad (22)$$

given in terms of the Bayesian information

$$\mathcal{Q}(x) = \int d\lambda p(\lambda|x)[\partial_{\lambda} \log_{10} p(\lambda|x)]^2. \quad (23)$$

The EMSD saturates the TBCRB, i.e., the bound is tight, when the posterior probability density takes the form of a Gaussian distribution in the $\Lambda_{\mathcal{M}}$ parametrization with inverse variance $\mathcal{Q}(x)$ [26]. According to the Laplace–Bernstein–von Mises theorem, this condition is satisfied in the asymptotic limit [46].

Evaluating the TBCRB is in general no less demanding than computing the EMSD, and our reason for starting with the TBCRB is that it comes with a condition for when it is tight. Often it is more convenient to consider the Bayesian Cramér-Rao bound (BCRB), which is obtained via an application of Jensen's inequality [26,47]

$$\begin{aligned} \text{TBCRB}^{-1} &\leq \text{BCRB}^{-1} \\ &:= Q_{\text{prior}} + \int d\theta p(\theta) \frac{h_{\Pi}(\theta)}{h_{\mathcal{M}}(\theta)}, \end{aligned} \quad (24)$$

where the prior information content is quantified by

$$Q_{\text{prior}} = \int d\lambda p(\lambda) [\partial_{\lambda} \log_{10} p(\lambda)]^2, \quad (25)$$

which is simply the Bayesian information evaluated for the initial prior. The TBCRB saturates the BCRB if the Bayesian information function is a constant independent of the measurement data. The advantage of working in the $\Lambda_{\mathcal{M}}$ parametrization is that the derived bound is automatically parametrization invariant.

Both the TBCRB and the BCRB are applicable at the single-shot level of the measurement protocol. In many cases the TBCRB is as difficult to compute as the EMSD; this leaves BCRB optimization as a viable choice of optimization strategy.

G. Specifying the initial prior

Within the Bayesian approach, prior densities are updated to posterior densities as measurement data are acquired. Applying the Bayesian framework requires specifying an initial prior and it is crucial to choose a prior which correctly represents ones prior knowledge or does not strongly effect the inference process. If we recall that the QFI metric encodes the thermodynamic geometry of the sample space, then it is convenient to express a general prior as a product

$$p(\theta) = h_{\mathcal{L}}^{1/2}(\theta) f(\theta), \quad (26)$$

where $f(\theta)$ is a parametrization-invariant density. If we are interested in applying an equal *a priori* probability postulate on a specific domain of parameter space, i.e., we want an uninformative prior over the specified domain, then this corresponds to having a constant, $f(\theta) = \text{const}$, prior density on this domain. For a constant density the prior is proportional to the QFI metric; in other words, the probability assigned to an infinitesimal parameter interval is proportional to the length of the interval. In this case we recover the so-called Jeffrey prior, which is the prior expressing complete ignorance given the metric structure of the parameter space [48]. For all simulations performed in the next section we adopt a smoothed density

$$f(\theta) = \frac{1}{\mathcal{N}} \left\{ \exp \left[\alpha \sin^2 \left(\pi \frac{\lambda(\theta) - \lambda_{\min}}{\lambda_{\max} - \lambda_{\min}} \right) \right] - 1 \right\} \quad (27)$$

with a normalization factor given by

$$\mathcal{N} \equiv (\lambda_{\max} - \lambda_{\min}) [\exp(\alpha/2) I_0(\alpha/2) - 1], \quad (28)$$

where $\theta \in [\theta_{\min}, \theta_{\max}]$, $\lambda_{\min, \max} = \lambda(\theta_{\min, \max})$, and I_0 is the modified Bessel function of the first kind. The resulting prior is essentially the one studied by Li *et al.* [46]; however, we

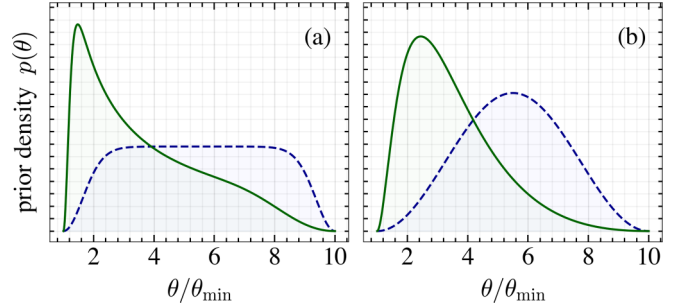


FIG. 3. Illustration of the smoothed Jeffrey prior [Eqs. (26) and (27)] for a flat reference $h_{\mathcal{L}}(\theta) \propto 1$ (blue dashed lines) and a curved metric $h_{\mathcal{L}}(\theta) \propto 1/\theta^2$ (green solid lines). (a) Plot for $\alpha = -10$. In this limit we approach an equal *a priori* probability postulate on the bounded interval $\theta \in [\theta_{\min}, \theta_{\max}]$. For a flat metric this gives a constant prior density with smoothed boundaries. For the curved metric the prior density is concentrated at small parameter values as the metric assigns the majority of metric length to this regime. (b) Plot for $\alpha = 0$. In this regime, the prior resembles a Gaussian distribution on a bounded interval.

specify the density with respect to the Euclidean parametrization. In the limit of large negative α , the density goes to a constant on the parameter domain and we thus recover Jeffrey's prior. The prior is illustrated in Fig. 3 for a flat reference metric $h_{\mathcal{L}}(\theta) \propto 1$ and a curved reference metric given by $h_{\mathcal{L}}(\theta) \propto \theta^{-2}$.

III. APPLICATIONS

In this section we consider some applications of the above estimation theory to the problem of thermometry. First, we argue for a suitable form of the reference metric when the sample system can be represented as a thermalizing channel and show that in this case the relative error is given by the standard noise-to-signal ratio. Second, we consider the case of thermal spin-1/2 particles and illustrate the difference between a thermometric scenario in which the spin-1/2 particles are considered as a thermometer for an underlying heat reservoir and a scenario in which the spin-1/2 particles themselves constitute the sample-system of interest. These two scenarios correspond to scenarios *B* and *A* illustrated in Fig. 1.

A. Preliminaries

We consider the case where the one-parameter family of sample-system states is the thermal Gibbs states with respect to a Hamiltonian operator H and take θ to denote the temperature. Recall that the Gibbs state takes the form [18,19]

$$\rho(\theta) = \frac{\exp(-H/k_B\theta)}{\text{Tr}[\exp(-H/k_B\theta)]}. \quad (29)$$

As our reference measurement we take a projective measurement of the SLD operator, which for Gibbs states takes the form [18,19,49]

$$\mathcal{L}(\theta) = \frac{H - \text{Tr}[\tau(\theta)H]}{2k_B\theta^2}, \quad (30)$$

where k_B is the Boltzmann constant. Measuring the SLD thus corresponds to a projective measurement of the sample-

system energy. Noting that the Hamiltonian operator is temperature independent and that this feature carries over to the eigenbasis of the SLD, it follows that a projective measurement of the SLD operator provides a valid reference measurement. For a projective energy measurement the associated Fisher information (FI), or equivalently the QFI, is then directly related to the sample-system heat capacity [18,19,49], i.e.,

$$h_{\mathcal{L}}(\theta) = \frac{\partial_{\theta} \text{Tr}[\tau(\theta)H]}{k_B \theta^2} =: \frac{C(\theta; H)}{k_B \theta^2}. \quad (31)$$

The square root of the QFI then provides a metric on the manifold of thermal Gibbs states. Recall that the associated metric-based distance is called the thermodynamic length [30–33]. Below we evaluate the thermodynamic length in the case of a bosonic mode and a spin-1/2 particle. First, however, we discuss how to represent a thermalizing channel.

B. Thermalizing channel

A common scenario in quantum thermometry is that of a quantum probe subject to a thermalizing channel. A thermalizing channel can in general be modeled as induced by a sample system which is effectively an infinitely large heat reservoir. Here we model such an ideal heat reservoir by a heat capacity which, either approximately or by definition, equals a constant value $k_B \mathcal{V}$ across the range of relevant temperatures, i.e., the sample energy is a linear function of the temperature. Note that a constant heat capacity corresponds to the QFI metric $h_{\mathcal{L}}(\theta) = k_B \mathcal{V} / \theta^2$. In this case we can evaluate the thermodynamic length analytically and find

$$\mathcal{D}(\theta_0, \theta_1) = |\log_{10}(\theta_1/\theta_0)|, \quad (32)$$

where for convenience we set $\mathcal{V} = 1$. The form of the MSD resulting from this distance function is called the mean-square logarithmic error (MSLE). The MSLE can be adopted whenever it can be assumed that the manifold of thermal states is generated via a weak coupling to an infinite heat reservoir. In practice, this assumption might break down at low temperatures, and more fundamentally there are cases where thermal behavior cannot be linked to an infinite heat reservoir, e.g., the case of subsystem thermalization described by the eigenstate thermalization hypothesis [38,50].

The MSLE was recently proposed by Rubio *et al.* [20] as a suitable measure of confidence, in the special case of a reference measurement for which the associated likelihood function satisfies the scale-invariance property

$$p(x|\theta) = \frac{g(x/\theta)}{\int dx g(x/\theta)}, \quad (33)$$

where x denotes the outcome from a projective measurement of the sample energy and $g(x/\theta)$ is a function only of the dimensionless ratio x/θ . This scale-invariance property of the likelihood function is satisfied for sample systems with a constant density of states or, equivalently, a constant heat capacity.

When considering the MSLE, it follows that the reference QFI takes the form $h_{\mathcal{L}}(\theta) = \theta^{-2}$, and we then find that the

associated ECRB for an unbiased estimator is given by

$$\text{ECRB} = \int d\theta \frac{p(\theta)}{\theta^2 h_{\mathcal{L}}(\theta)}. \quad (34)$$

The quantity $\theta^2 h_{\mathcal{L}}(\theta)$ provides an upper bound on the signal-to-noise ratio within the frequentist estimation paradigm [51,52]. This shows that when the sample can be modeled as an ideal heat reservoir, the standard notion of relative error is recovered in the local limit where $p(\theta)$ is sharply peaked. However, our analysis also points out that the standard relative error is not suitable unless the sample system has an approximately constant heat capacity across the range of relevant temperature. This condition typically breaks down at sufficiently low temperatures.

1. Thermodynamic length for a bosonic mode

A specific system approximately realizing the above assumptions on an ideal heat reservoir is a Bose gas at a fixed density well above the critical temperature [51]. As an illustration we consider a gas of bosonic modes with energy gap ε . In this case the QFI per mode is given by [18]

$$h_{\mathcal{L},\text{boson}}(\theta) = \frac{\varepsilon^2/k_B^2 \theta^4}{4 \sinh^2(\varepsilon/2k_B \theta)}. \quad (35)$$

For this QFI the associated λ function [see Eq. (12)] can be given an analytic expression

$$\lambda_{\text{boson}}(\theta) = -\log_{10}[\tanh(\varepsilon/4k_B \theta)], \quad (36)$$

which implies that $\Lambda_{\mathcal{L},\text{boson}} = [0, +\infty)$. The bosonic QFI and the associated λ function are shown in Figs. 4(a) and 4(b), together with the corresponding quantities for an ideal heat reservoir. We observe that in the limit where the temperature is large compared to the boson energy gap, the bosonic modes approximate the ideal heat reservoir, i.e., $h_{\mathcal{L},\text{boson}}(\theta) \rightarrow 1/\theta^2$. This suggests that we can generically represent an ideal thermalizing channel physically by a collection of low-frequency bosonic modes.

C. Noninteracting spin-1/2 particles

We now consider N noninteracting spin-1/2 particles, or qubits, with identical energy gaps ε . The spin-1/2 particles are in a thermal Gibbs state and as above we take the θ parametrization to be the temperature. We are going to compare and contrast two scenarios: scenario *A*, in which the spin-1/2 particles themselves constitute the sample system of interest, e.g., we are interested in thermometry of the spin degrees of freedom of the ultracold atoms themselves [8,34], and scenario *B*, where the spin-1/2 particles are employed as equilibrium thermometers of an underlying heat reservoir, e.g., the particles could model impurities within an ultracold gas, mapping motional information of gas atoms onto the quantum-spin state [10,35].

In scenario *B*, as argued above, the MSLE is the suitable measure of confidence and the QFI metric takes the form $h_{\mathcal{L},\text{res}}(\theta) = \theta^{-2}$. In scenario *A* we take the MSD resulting from adopting the QFI metric of a thermal spin-1/2 particle




FIG. 4. (a) Plot of the QFI for a spin-1/2 particle (blue dashed line), the QFI for a bosonic mode (red dash-dotted line), and the QFI for the ideal heat reservoir (green solid line). (b) Plot of the λ functions associated with the various QFI.

as our reference, i.e., [51]

$$h_{\mathcal{L},\text{spin}}(\theta) = \frac{\epsilon^2/k_B^2\theta^4}{4 \cosh^2(\epsilon/2k_B\theta)}, \quad (37)$$

which, we recall, corresponds to a projective energy measurement of a single spin-1/2 particle. When referring to the MSD below, we refer to the spin-1/2 particle QFI metric. For the spin-1/2 reference, the inverse derivative function $\lambda(\theta)$ can be obtained analytically; it takes the form

$$\lambda_{\text{spin}}(\theta) = \pi - 2 \arctan(e^{\epsilon/2k_B\theta}), \quad (38)$$

and thus $\Lambda_{\mathcal{L},\text{spin}} = [0, \pi/2]$. In Fig. 4(a) we plot the spin-1/2 QFI and compare it with the heat reservoir reference and in Fig. 4(b) we plot the associated λ functions. We observe that for temperature around $k_B\theta/\epsilon \approx 0.3$ the two λ functions exhibit a similar gradient; however, away from this temperature regime the specific thermometric scenario considered plays a role.

The estimation strategy employed consists of projectively measuring the energy of a subset of μ particles; this is written as the number of measurements in Fig. 5. For thermalized spin-1/2 particles, projective energy measurements maximize the associated FI for all temperatures [53], and it follows that the FI associated with this measurement can be expressed as

$$h_{\Pi}(\theta) = \mu h_{\mathcal{L},\text{spin}}(\theta). \quad (39)$$

In Figs. 5(a)–5(c) we show stochastic simulations of measurement trajectories sampled according to three different true temperatures. In all cases we plot the minimal MSLE (MMSLE) estimator and the MMSD estimator and note that

only negligible differences exist between these, despite different priors being used. This feature is to be expected as the respective λ functions rapidly become approximately constant across the posterior; when this is the case the estimator is simply the maximum *a posteriori* temperature, i.e., the temperature at which the posterior takes its maximum value, independently of the specific λ function.

Although the various temperature estimates are only negligibly different, the confidence assigned to the estimates depends on the thermometric scenario. In Figs. 5(e) and 5(f) we show the expected MSLE (EMSLE) and the EMSD, respectively, as a function of the subset size μ , or equivalently the number of independent measurements, and compare these with the associated BCRB and the associated ECRB. The EMSLE and the EMSD are evaluated using the corresponding smoothed Jeffrey priors shown in Fig. 5(d). In the case of the EMSLE, we do not observe convergence to the BCRB. This is to be expected since $h_{\mathcal{L},\text{res}}/h_{\Pi}$ is generally not constant across the domain of the employed prior. This feature also means that we do not observe a convergence to have an ECRB at the trajectory level, i.e., we observe fluctuations around the average. For the EMSD we observe rapid convergence to the BCRB and trajectory level convergence to the ECRB. This is due to the fact that $h_{\mathcal{L},\text{spin}}/h_{\Pi} = 1/\mu$ is a constant independent of the temperature.

When considering the spin-1/2 particles as thermometers of a underlying heat reservoir, it is sensible to consider adaptation of the energy gap ϵ to optimize the thermometric performance.¹ That is, we consider a protocol consisting of thermalizing a spin-1/2 particle and projectively measuring the energy. Based on the observed outcome, the energy gap of the particle employed in the following measurement is adapted. In Fig. 5(e) we also show the adaptive EMSLE, which is obtained via a BCRB optimization strategy of the spin-1/2 energy gap ϵ in each iteration of the protocol. By a BCRB optimization strategy we refer to a minimization of the BCRB with respect to the energy gap ϵ in each iteration. In this case we observe a rapid convergence to the optimized local result, i.e., the adaptive EMSLE $\sim \max_{\theta}[h_{\mathcal{L},\text{res}}/h_{\Pi}]$. Furthermore, the adaptive EMSLE converges to a lower bound on the set of MSLE trajectories. This is to be expected as the energy gap can be tuned to the optimal ratio of the energy gap to the temperature sampled from the prior.

IV. CONCLUSION

In this paper we have developed a general Bayesian approach to quantum thermometry based on the concept of thermodynamic length. The basic idea is that a meaningful measure of the precision of a temperature estimate should be based on the ability to distinguish states at different temperatures, i.e., colder from hotter, and should be independent of the particular parametrization of the states, e.g., temperature. The theory allows us to meaningfully quantify the thermometric performance of a given measurement strategy and to design

¹See the companion paper [54] focusing explicitly on the role of adaptivity in saturating fundamental precision bounds on Bayesian quantum thermometry via equilibrium quantum probes.

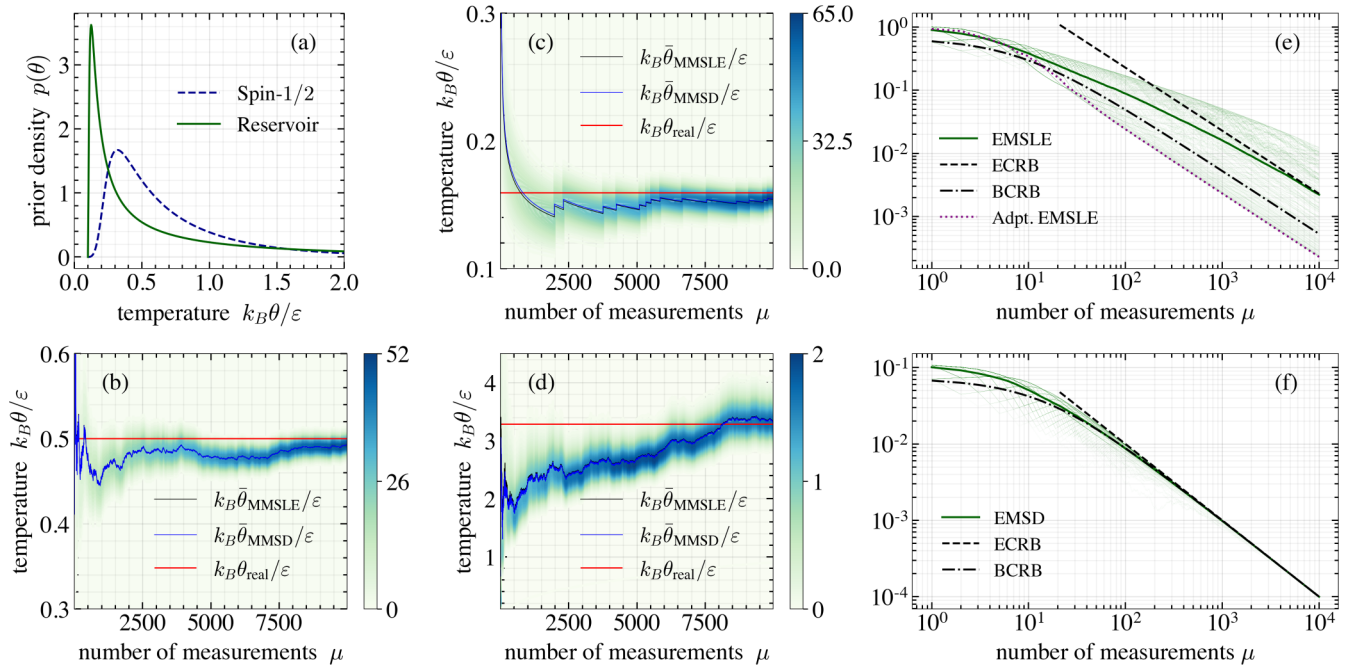


FIG. 5. (a) Smoothed Jeffrey prior [Eqs. (26) and (27)] for $\alpha = -4.5$ and $k_B \theta \in [\varepsilon/10, 10\varepsilon]$, corresponding to the spin-1/2 particle and the reservoir reference metric. (b)–(d) Simulated stochastic measurement trajectories sampled according to a specific true temperature (red horizontal line). The posterior probability density function is shown as a colormap. Both the MMSLE and the MMSD estimators are shown. (e) Simulation of the EMSLE (scenario B) for thermometry employing free spin-1/2 particles. The simulation is performed by averaging over 250 stochastic trajectories where the measurement data are generated with respect to a true temperature sampled according to the prior density. The EMSLE is compared with the ECRB, convergence is observed for $\mu \gtrsim 10^3$, and we note that the generated trajectories do not converge to the mean asymptotically. Furthermore, we plot the BCRB and observe that this is not a tight bound. Finally, we plot the adaptive EMSLE. This is obtained by adapting the energy gap ε in each repetition of the measurement. Here we employ a BCRB optimization strategy, which is described in more detail in the text. (f) Simulation of the EMSD (scenario A) for thermometry of free spin-1/2 particles. The simulation is performed by averaging over 250 stochastic trajectories where the measurement data are generated with respect to a true temperature sampled according to the prior density. The EMSD is compared with the BCRB and convergence is observed for $\mu \gtrsim 10^2$. Furthermore, we observe that the generated trajectories converge to the average.

optimal measurements under conditions of prior temperature uncertainty. Furthermore, the framework is applicable at the single-shot level, which allows us to consider adaptive measurement schemes. Insisting on a framework which satisfies both parametrization invariance and any other symmetry of the sample system considered leads us to gauge the thermometric performance using a distance function defined relative to the quantum Fisher information metric on the manifold of thermal states of the sample, i.e., the thermodynamic length. Formulating the problem in terms of the thermodynamic length then leads to the realization that any meaningful notion of relative error must be given with respect to the specific sample system considered. Furthermore, we demonstrate that the suitability of both Bayesian and local Cramér-Rao bounds is contingent on the specific sample system which is measured.

The generalization of the notion of relative error is particularly important in the low-temperature regime, where it shows that the standard signal-to-noise ratio is not a good

indicator of estimation precision, and provides an alternative point of view on the problems of low-temperature quantum thermometry [51,52]. In addition, the generalization is important at the microscale, where we are often not interested in thermometry of ideal heat baths, but rather in characterizing the thermodynamics of operating quantum devices such as quantum simulators.

ACKNOWLEDGMENTS

M.R.J. and J.B.B. acknowledge financial support from the Independent Research Fund Denmark. M.M. and M.P.-L. acknowledge financial support from the Swiss National Science Foundations (NCCR SwissMAP and Ambizione Grant No. PZ00P2-186067). J.K. acknowledges financial support from the Foundation for Polish Science within the “Quantum Optical Technologies” project carried out within the International Research Agendas program cofinanced by the European Union under the European Regional Development Fund.

[1] G. Kucsko, P. C. Maurer, N. Y. Yao, M. Kubo, H. J. Noh, P. K. Lo, H. Park, and M. D. Lukin, Nanometre-scale thermometry in a living cell, *Nature (London)* **500**, 54 (2013).

[2] M. Fujiwara, S. Sun, A. Dohms, Y. Nishimura, K. Suto, Y. Takezawa, K. Oshimi, L. Zhao, N. Sadzak, Y. Umehara, Y. Teki, N. Komatsu, O. Benson, Y. Shikano, and E. Kage-Nakadai,

- Real-time nanodiamond thermometry probing in vivo thermogenic responses, *Sci. Adv.* **6**, eaba9636 (2020).
- [3] E. Moreva, E. Bernardi, P. Traina, A. Sosso, S. D. Tchernij, J. Forneris, F. Picollo, G. Brida, Ž. Pastuović, I. P. Degiovanni, P. Olivero, and M. Genovese, Practical Applications of Quantum Sensing: A Simple Method to Enhance the Sensitivity of Nitrogen-Vacancy-Based Temperature Sensors, *Phys. Rev. Appl.* **13**, 054057 (2020).
- [4] C. Carcy, G. Hercé, A. Tenart, T. Roscilde, and D. Clément, Certifying the Adiabatic Preparation of Ultracold Lattice Bosons in the Vicinity of the Mott Transition, *Phys. Rev. Lett.* **126**, 045301 (2021).
- [5] D. C. McKay and B. DeMarco, Cooling in strongly correlated optical lattices: Prospects and challenges, *Rep. Prog. Phys.* **74**, 054401 (2011).
- [6] M. T. Mitchison, T. Fogarty, G. Guarnieri, S. Campbell, T. Busch, and J. Goold, *In Situ* Thermometry of a Cold Fermi Gas via Dephasing Impurities, *Phys. Rev. Lett.* **125**, 080402 (2020).
- [7] M. Mehboudi, A. Lampo, C. Charalambous, L. A. Correa, M. Á. García-March, and M. Lewenstein, Using Polarons for sub-nK Quantum Nondemolition Thermometry in a Bose-Einstein Condensate, *Phys. Rev. Lett.* **122**, 030403 (2019).
- [8] T. Hartke, B. Oreg, N. Jia, and M. Zwierlein, Doublon-Hole Correlations and Fluctuation Thermometry in a Fermi-Hubbard Gas, *Phys. Rev. Lett.* **125**, 113601 (2020).
- [9] J.-P. Brantut, C. Grenier, J. Meineke, D. Stadler, S. Krinner, C. Kollath, T. Esslinger, and A. Georges, A thermoelectric heat engine with ultracold atoms, *Science* **342**, 713 (2013).
- [10] Q. Bouton, J. Nettersheim, D. Adam, F. Schmidt, D. Mayer, T. Lausch, E. Tiemann, and A. Widera, Single-Atom Quantum Probes for Ultracold Gases Boosted by Nonequilibrium Spin Dynamics, *Phys. Rev. X* **10**, 011018 (2020).
- [11] S. Gasparinetti, K. L. Viisanen, O.-P. Saira, T. Faivre, M. Arzeo, M. Meschke, and J. P. Pekola, Fast Electron Thermometry for Ultrasensitive Calorimetric Detection, *Phys. Rev. Appl.* **3**, 014007 (2015).
- [12] M. Mecklenburg, W. A. Hubbard, E. R. White, R. Dhall, S. B. Cronin, S. Aloni, and B. C. Regan, Nanoscale temperature mapping in operating microelectronic devices, *Science* **347**, 629 (2015).
- [13] D. Halbertal, J. Cuppens, M. B. Shalom, L. Embon, N. Shadmi, Y. Anahory, H. R. Naren, J. Sarkar, A. Uri, Y. Ronen, Y. Myasoedov, L. S. Levitov, E. Joselevich, A. K. Geim, and E. Zeldov, Nanoscale thermal imaging of dissipation in quantum systems, *Nature (London)* **539**, 407 (2016).
- [14] F. Giazotto, T. T. Heikkilä, A. Luukanen, A. M. Savin, and J. P. Pekola, Opportunities for mesoscopics in thermometry and refrigeration: Physics and applications, *Rev. Mod. Phys.* **78**, 217 (2006).
- [15] B. Karimi, F. Brange, P. Samuelsson, and J. P. Pekola, Reaching the ultimate energy resolution of a quantum detector, *Nat. Commun.* **11**, 367 (2020).
- [16] V. Giovannetti, S. Lloyd, and L. Maccone, Advances in quantum metrology, *Nat. Photon.* **5**, 222 (2011).
- [17] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum Metrology, *Phys. Rev. Lett.* **96**, 010401 (2006).
- [18] M. Mehboudi, A. Sanpera, and L. A. Correa, Thermometry in the quantum regime: Recent theoretical progress, *J. Phys. A: Math. Theor.* **52**, 303001 (2019).
- [19] A. De Pasquale and T. M. Stace, in *Thermodynamics in the Quantum Regime: Fundamental Aspects and New Directions*, edited by F. Binder, L. A. Correa, C. Gogolin, J. Anders, and G. Adesso, Fundamental Theories of Physics Vol. 195 (Springer, Cham, 2018), p. 503.
- [20] J. Rubio, J. Anders, and L. A. Correa, Global Quantum Thermometry, *Phys. Rev. Lett.* **127**, 190402 (2021).
- [21] G. O. Alves and G. T. Landi, Bayesian estimation for collisional thermometry, *Phys. Rev. A* **105**, 012212 (2022).
- [22] W.-K. Mok, K. Bharti, L.-C. Kwek, and A. Bayat, Optimal probes for global quantum thermometry, *Commun. Phys.* **4**, 62 (2021).
- [23] E. L. Lehmann and G. Casella, *Theory of Point Estimation*, Springer Texts in Statistics Vol. 31 (Springer, New York, 1998).
- [24] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory* (Prentice Hall, Englewood Cliffs, 1993).
- [25] M. G. A. Paris, Quantum estimation for quantum technology, *Int. J. Quantum. Inf.* **07**, 125 (2009).
- [26] L. Bacharach, C. Fritsche, U. Orguner, and E. Chaumette, Some results on tighter Bayesian lower bounds on the mean-square error, [arXiv:1907.09509](https://arxiv.org/abs/1907.09509).
- [27] S. L. Braunstein and C. M. Caves, Statistical Distance and the Geometry of Quantum States, *Phys. Rev. Lett.* **72**, 3439 (1994).
- [28] U. von Toussaint, Bayesian inference in physics, *Rev. Mod. Phys.* **83**, 943 (2011).
- [29] *Bayesian Bounds for Parameter Estimation and Nonlinear Filtering/Tracking*, edited by H. L. Van Trees and K. L. Bell (Wiley, New York, 2007).
- [30] F. Weinhold, Metric geometry of equilibrium thermodynamics, *J. Chem. Phys.* **63**, 2479 (1975).
- [31] P. Salamon and R. S. Berry, Thermodynamic Length and Dissipated Availability, *Phys. Rev. Lett.* **51**, 1127 (1983).
- [32] G. E. Crooks, Measuring Thermodynamic Length, *Phys. Rev. Lett.* **99**, 100602 (2007).
- [33] M. Scandi and M. Perarnau-Llobet, Thermodynamic length in open quantum systems, *Quantum* **3**, 197 (2019).
- [34] D. M. Weld, P. Medley, H. Miyake, D. Hucul, D. E. Pritchard, and W. Ketterle, Spin Gradient Thermometry for Ultracold Atoms in Optical Lattices, *Phys. Rev. Lett.* **103**, 245301 (2009).
- [35] R. Olf, F. Fang, G. E. Marti, A. MacRae, and D. M. Stamper-Kurn, Thermometry and cooling of a Bose gas to 0.02 times the condensation temperature, *Nat. Phys.* **11**, 720 (2015).
- [36] G. Gallavotti, *Statistical Mechanics: A Short Treatise* (Springer, Berlin, 1999).
- [37] G. Guarnieri, G. T. Landi, S. R. Clark, and J. Goold, Thermodynamics of precision in quantum nonequilibrium steady states, *Phys. Rev. Res.* **1**, 033021 (2019).
- [38] M. Brenes, S. Pappalardi, J. Goold, and A. Silva, Multiparticle Entanglement Structure in the Eigenstate Thermalization Hypothesis, *Phys. Rev. Lett.* **124**, 040605 (2020).
- [39] L. Ballentine, *Quantum Mechanics: A Modern Development* (World Scientific, Singapore, 2014).
- [40] I. Jermy, Invariant Bayesian estimation on manifolds, *Ann. Stat.* **33**, 583 (2005).
- [41] S.-I. Amari and H. Nagaoka, *Methods of Information Geometry* (American Mathematical Society, Providence, 2007), Vol. 191.
- [42] S. Amari, *Differential-Geometrical Methods in Statistics* (Springer, Berlin, 1985).

- [43] H. Snoussi, in *Bayesian Information Geometry: Application to Prior Selection on Statistical Manifolds* (Elsevier, Amsterdam, 2007), pp. 163–207.
- [44] M. Jarzyna and J. Kołodyński, Geometric approach to quantum statistical inference, *IEEE J. Sel. Areas Inf. Theory* **1**, 367 (2020).
- [45] N. N. Chentsov, Algebraic foundation of mathematical statistics, *Math. Operationsforsch. Stat. Ser. Statist.* **9**, 267 (1978) [*Statistical Decision Rules and Optimal Inference*, Translation of Mathematical Monographs Vol. 53 (Providence American Mathematical Society, Providence, 1982)].
- [46] Y. Li, L. Pezzè, M. Gessner, Z. Ren, W. Li, and A. Smerzi, Frequentist and Bayesian quantum phase estimation, *Entropy* **20**, 628 (2018).
- [47] L. Bacharach, C. Fritsche, U. Orguner, and E. Chaumette, in *Proceedings of the 2019 IEEE International Conference on Acoustics, Speech, and Signal Processing, Brighton, 2019* (IEEE, Piscataway, 2019), pp. 5277–5281.
- [48] A. Caticha, Lectures on probability, entropy, and statistical physics, [arXiv:0808.0012](https://arxiv.org/abs/0808.0012).
- [49] L. A. Correa, M. Mehboudi, G. Adesso, and A. Sanpera, Individual Quantum Probes for Optimal Thermometry, *Phys. Rev. Lett.* **114**, 220405 (2015).
- [50] Y. Ashida, K. Saito, and M. Ueda, Thermalization and Heating Dynamics in Open Generic Many-Body Systems, *Phys. Rev. Lett.* **121**, 170402 (2018).
- [51] P. P. Potts, J. B. Brask, and N. Brunner, Fundamental limits on low-temperature quantum thermometry with finite resolution, *Quantum* **3**, 161 (2019).
- [52] M. R. Jørgensen, P. P. Potts, M. G. A. Paris, and J. B. Brask, Tight bound on finite-resolution quantum thermometry at low temperatures, *Phys. Rev. Res.* **2**, 033394 (2020).
- [53] M. G. A. Paris, Achieving the Landau bound to precision of quantum thermometry in systems with vanishing gap, *J. Phys. A: Math. Theor.* **49**, 03LT02 (2016).
- [54] M. Mehboudi, M. R. Jørgensen, S. Seah, J. B. Brask, J. Kołodyński, and M. Perarnau-Llobet, Fundamental Limits in Bayesian Thermometry and Attainability via Adaptive Strategies, *Phys. Rev. Lett.* **128**, 130502 (2022).