Fundamental Limits in Bayesian Thermometry and Attainability via Adaptive Strategies

Mohammad Mehboudi,[1,*] Mathias R. Jørgensen,[2,†] Stella Seah,[1] Jonatan B. Brask,[2] Jan Kołodyński,[3] and Martí Perarnau-Llobet[1,‡]

[1]*Département de Physique Appliquée, Université de Genève, 1211 Geneva, Switzerland*
[2]*Department of Physics, Technical University of Denmark, 2800 Kongens Lyngby, Denmark*
[3]*Centre for Quantum Optical Technologies, Centre of New Technologies, University of Warsaw, 02-097 Warsaw, Poland*

(Received 1 September 2021; revised 30 December 2021; accepted 7 March 2022; published 1 April 2022)

We investigate the limits of thermometry using quantum probes at thermal equilibrium within the Bayesian approach. We consider the possibility of engineering interactions between the probes in order to enhance their sensitivity, as well as feedback during the measurement process, i.e., adaptive protocols. On the one hand, we obtain an ultimate bound on thermometry precision in the Bayesian setting, valid for arbitrary interactions and measurement schemes, which lower bounds the error with a quadratic (Heisenberg-like) scaling with the number of probes. We develop a simple adaptive strategy that can saturate this limit. On the other hand, we derive a no-go theorem for nonadaptive protocols that does not allow for better than linear (shot-noise-like) scaling even if one has unlimited control over the probes, namely, access to arbitrary many-body interactions.

DOI: 10.1103/PhysRevLett.128.130502

Introduction.—Preparing quantum systems at low temperatures is an essential task for development of quantum technologies [1–3]. Measuring temperature precisely is necessary to validate cooling and ensure the performance of quantum protocols, and has been demonstrated in cutting-edge experiments [4–12]; it is however challenging. Due to the scarcity of thermal fluctuations at such low temperatures, the relative error on thermometry can be enormous. Moreover, the fragility of quantum systems requires additional forward planning to minimise disturbance while maximising the information obtained. The theory of quantum thermometry is built to address these pivotal challenges [13,14].

Quantum thermometry finds fundamental limits on precision [15–18] and designs protocols to achieve them in different platforms [19–22], and improve them thanks to quantum correlations [23,24], coherence [25,26], many-body interactions and criticality [27–32] or other resources [33,34]. To date, such enhancements have been developed in the context of local thermometry, aiming at designing a thermometer that detects the smallest temperature variations around a known temperature [13,14]. In many practical situations, however, one might not know the temperature accurately beforehand. Rather, one has limited prior knowledge about the temperature of the sample. Under such circumstances, Bayesian estimation is a more suitable approach, and has been the subject of a few recent studies [35,36].

The goal of this Letter is to set the ultimate bounds of Bayesian equilibrium thermometry, and to develop adaptive strategies to saturate them. It is insightful to first recall analogous results in the local approach to equilibrium thermometry [13,14]. Within such a framework—contrary to dynamical approaches where the probe evolves according to some predefined model parametrized by the temperature [37,38], e.g., a superconducting qubit in radiometry [39]—the probe always thermalizes to the temperature of the sample whose value is known *a priori*. In that case, for any unbiased estimator $\tilde{\theta}$ of the temperature θ_0, the mean square error is inversely proportional to the heat capacity of the probe: $\Delta\tilde{\theta} \propto 1/C$ [15,16,29,40]. For n-body probes, C can scale superextensively with n in the vicinity of a critical point, with the ultimate bound $C \approx n^2/4$ [15,41]—a quadratic scaling with the number of resources reminiscent of the Heisenberg scaling in quantum metrology [42]. Here, we show that similar bounds hold in the Bayesian approach, but adaptive strategies are needed to saturate them, contrary to the local case. In fact, we prove that any nonadaptive strategy necessarily leads to $\Delta\tilde{\theta} \propto 1/n$ for sufficiently large n—i.e., a shot-noise-like scaling [42]—a no-go result that holds even when arbitrary control over the n-body probe Hamiltonian is allowed. Thus, adaptive measurement strategies are a crucial ingredient for optimal thermometry whenever the temperature value is *a priori* not perfectly known.

Preliminaries and setup.—We consider estimation of the temperature θ_0 of a (possibly macroscopic) sample given some prior distribution $p(\theta)$ reflecting our initial knowledge on θ_0. We assume we have at our disposal N copies of a d-dimensional system that we use as probes, which are much smaller than the sample. When put in contact with the sample, we assume that the probes eventually reach thermal equilibrium at temperature θ_0. By measuring them we infer

θ_0. This corresponds to the framework of equilibrium thermometry, which is by nature robust [13,14]. In order to establish fundamental bounds, we assume full control on the Hamiltonian of the probes, and in particular the ability to make them interact. Therefore, alternatively one can think of a d^N-dimensional probe, which constitutes our resource.

The thermometry process is divided into m rounds, each involving $n = N/m$ probes. Every round consists of (I) preparation of the n-body probe, (II) interaction with the sample and thermalization, (III) measurement or data acquisition, and (IV) data analysis (see Fig. 1). In the first round, we start by engineering the Hamiltonian $H_n^{(1)}$ of the n-body probe into any desired configuration based on the prior distribution $p(\theta)$. That is, we arrange the energy distribution of the n-body probe to become most sensitive to the relevant temperature range. Next, in step (II), this n-body system is put in contact with the sample, and reaches thermal equilibrium with it. Therefore, it can be described by the Gibbs state $\omega_{\theta_0}(H_n^{(1)}) \coloneqq \exp[-H_n^{(1)}/\theta_0]/Z$, with $Z = \mathrm{Tr}(\exp[-H_n^{(1)}/\theta_0])$ the partition function. Then, in step III, a measurement is performed that yields an outcome x_1. We focus on energy measurements since they are optimal as the Gibbs state is diagonal in the energy basis. In the data analysis (step IV), the posterior distribution is obtained through Bayes' rule:

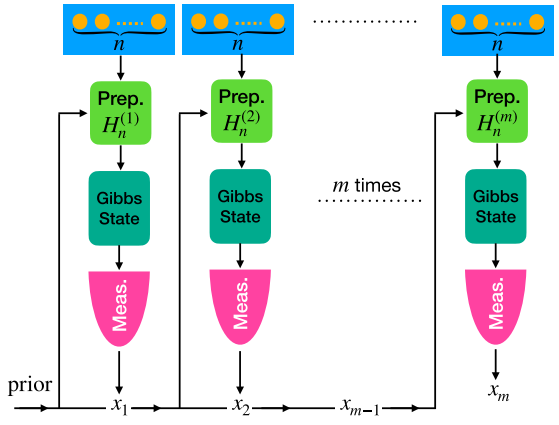

FIG. 1. Schematic representation of the adaptive scenario. A total of N probes are used in groups of n to estimate the temperature of the sample, θ_0. Initially, our prior temperature distribution is given by $p(\theta)$, according to which we choose the Hamiltonian of the first n probes to be $H_n^{(1)}$ that minimizes the expected mean square logarithmic error. The probes interact and thermalize with the sample followed by an energy measurement, yielding an outcome, say x_1. Our knowledge about the temperature will be reflected in the posterior distribution $p(\theta|x_1)$. This will be used as the prior for the second round—in order to find the optimal Hamiltonian $H_n^{(2)}$. This process is repeated $m = N/n$ times. In contrast, in the nonadaptive scenario the Hamiltonian is fixed $H_n^{(k)} = H_n \ \forall \ k$.

$$p(\theta|x_1) = \frac{p(x_1|\theta)p(\theta)}{p(x_1)}, \tag{1}$$

where $p(x|\theta)$ is the likelihood function (which depends on the temperature *and* the Hamiltonian), $p(\theta)$ is the prior distribution on θ, and $p(x) = \int \mathrm{d}\theta p(\theta)p(x|\theta)$ is the outcome probability. The next round proceeds in an analogous way, but replacing the prior $p(\theta)$ by $p(\theta|x_1)$ and $H_n^{(1)}$ by $H_n^{(2)}$. Likewise, in round $k > 1$, $p(\theta)$ is replaced by $p(\theta|\mathbf{x}_{k-1})$ with $\mathbf{x}_{k-1} \equiv \{x_{k-1}, \ldots, x_2, x_1\}$ and $H_n^{(1)}$ is replaced by $H_n^{(k)}$. Such a strategy is adaptive since $H_n^{(k)}$ depends on $\mathbf{x}_{k-1}$. In contrast, a nonadaptive strategy satisfies $H_n^{(k)} = H_n \ \forall \ k$, where H_n is chosen according to the initial prior $p(\theta)$ only. At the end of the thermometry process (round m), the final estimate $\tilde{\theta}(\mathbf{x}_m)$ of θ_0 is computed.

In order to gauge the quality of the estimator, we need to introduce an error quantifier that describes how far $\tilde{\theta}$ is from θ_0, on average. A natural measure which is suitable for equilibrium probes is the expected mean square logarithmic error (EMSLE) (see Ref. [35] for justification and the accompanying paper [43] for a deeper analysis and generalization)

$$\mathrm{EMSLE} \coloneqq \int d\theta p(\theta) \int d\mathbf{x}_m \, p(\mathbf{x}_m|\theta) \ln^2\left[\frac{\tilde{\theta}(\mathbf{x}_m)}{\theta}\right], \tag{2}$$

with $d\mathbf{x}_m \coloneqq dx_m \ldots dx_1$. Moreover,

$$\tilde{\theta}(\mathbf{x}_m) = \exp\left[\int d\theta \frac{p(\theta)p(\mathbf{x}_m|\theta)}{p(\mathbf{x}_m)} \ln\theta\right], \tag{3}$$

is the optimal temperature estimator, i.e., it minimizes EMSLE [35].

We wish to find lower bounds for EMSLE, as well as optimal strategies to saturate them, for both adaptive and nonadaptive measurements. More precisely, our aim is to minimize EMSLE as a function of the number N of probes, with $N = mn$. We will pay particular attention to the relevant case where $m \gg 1$ is large (asymptotic regime) but n is limited due to, e.g., experimental limitations on the amount of probes that can be collectively processed. In this case, we will focus on the scaling of EMSLE with n for a fixed but large m.

Main results.—Our main results are (i) an ultimate precision limit for Bayesian thermometry that holds for both adaptive and nonadaptive strategies, which in principle allows for a quadratic (Heisenberg-like) scaling with n, (ii) a no-go theorem that forbids superextensive scaling in any nonadaptive scenario, and (iii) an adaptive strategy that reaches the ultimate limit. These results are derived in what follows (technical details are given in the Supplemental Material [44]).

Given the prior $p(\theta)$, and by utilizing the Van Trees inequality [46,47] we construct a lower bound on the estimation error after m rounds

$$\mathrm{EMSLE}^{-1} \leqslant Q[p(\theta)] + \sum_{k=1}^{m} \int d\mathbf{x}_{k-1} p(\mathbf{x}_{k-1}) \times \int d\theta p(\theta|\mathbf{x}_{k-1}) C(\theta; H_n^{(k)}), \tag{4}$$

where $p(\theta|\mathbf{x}_0) = p(\theta)$, $p(\mathbf{x}_0) = p(\mathbf{x}_0|\theta) = 1$, and $\int d\mathbf{x}_0 = 1$ are introduced to compress our notation. Here, $Q[p(\theta)]$ quantifies the prior information and reads

$$Q[p(\theta)] \coloneqq \int d\theta\, p(\theta)[1 + \theta \partial_\theta \log p(\theta)]^2. \tag{5}$$

The second term quantifies the information acquired through all measurements. It also establishes a connection to the quantum Fisher information through its proportionality to the heat capacity [15]. The heat capacity of the probe at round k of the measurement is denoted $C(\theta; H_n^{(k)})$, with the Hamiltonian $H_n^{(k)}$ designed according to the prior *and* the information acquired so far. Recall that, by definition, $C(\theta; H_n) \coloneqq \partial_\theta E(\theta; H_n)$, where $E(\theta; H) = \mathrm{Tr}[H\omega_\theta(H)]$ is the energy of the probe at thermal equilibrium. To bound Eq. (4), we first define the maximum of the integrand over $\{H_n^{(k)}\}_{k=1}^m$ for a specific trajectory $\mathbf{x}_m$:

$$\Gamma(\mathbf{x}_m) \coloneqq \max_{\{H_n^{(k)}\}_k} \sum_{k=1}^{m} \int d\theta\, p(\theta|\mathbf{x}_{k-1}) C(\theta; H_n^{(k)}) \leqslant \sum_{k=1}^{m} \int d\theta\, p(\theta|\mathbf{x}_{k-1}) C_D = m C_D, \tag{6}$$

where $C_D \coloneqq \max_{H_n} C(\theta; H_n)$, i.e., the maximum heat capacity of an n-body probe. In the last line we used that C_D is independent of θ (see Refs. [15] or [44] for the explicit expression of C_D). Furthermore, we have $C_D \approx \frac{n^2}{4} \log^2 d$, for large enough n. Putting everything together, we obtain from Eq. (4)

$$\mathrm{EMSLE}^{-1} \leqslant Q[p(\theta)] + m C_D \overset{\eta \gg 1}{\approx} Q[p(\theta)] + m \frac{n^2}{4} \log^2 d. \tag{7}$$

This gives an ultimate bound on Bayesian thermometry [result (i)], which both adaptive and nonadaptive strategies should respect. This bound implies that any Bayesian thermometry protocol is ultimately limited by a quadratic Heisenberg-like scaling.

The ultimate bound (7) becomes tight and can be saturated by adaptive strategies in the regime $m \gg 1$ (see results below). However, nonadaptive strategies fail to saturate it, and in fact EMSLE^{-1} can increase at most linearly with n [result (ii)]

$$\mathrm{EMSLE}^{-1} \overset{\text{nonadaptive}}{\leqslant} Q[p(\theta)] + f[p(\theta)] mn \log d, \tag{8}$$

where $f[p(\theta)] = \int_{\mathcal{R}} d\theta[-\partial_\theta p(\theta)]\theta$ is a functional of only the prior distribution, and $\mathcal{R}$ is the temperature domain where $\partial_\theta p(\theta) \leqslant 0$. This result is rigorously proven in the Supplemental Material [44], but let us provide some intuition. It is already noted in the literature that engineered probes for thermometry show enhanced sensitivity only in a small temperature range Δ [13,15,48–50]. Finite-size scaling theory hints that if $C \propto n^{1+\alpha}$, then $\Delta \propto n^{-\gamma}$ with $\gamma \geqslant \alpha$ in order to ensure that the energy density of an equilibrium state remains finite [51]. This implies that, for any $p(\theta)$ with a finite width (independent of n), the term $\int \mathrm{d}\theta\, p(\theta) C(\theta)$ in Eq. (4) grows at most linearly with n for sufficiently large n. In other words, optimal n-body probes require priors with a width smaller than $\mathcal{O}(1/n)$ to obtain superlinear scaling, and conversely a finite width in $p(\theta)$ will eventually kill any superlinear scaling. The no-go result (8) makes this intuition rigorous.

The above reasoning also explains why adaptive protocols can potentially saturate (7). By updating the prior $p(\theta)$ to the posterior $p(\theta|\mathbf{x}_{k-1})$ in each step of the process $(k = 1, \ldots, m)$, it can stay inside the optimal region for sufficiently large m, thus enabling superlinear precision. This also suggests using optimal probes for local thermometry as an *ansatz* for the Bayesian thermometry with adaptive strategies. The optimal thermometer in the local scenario is an effective two-level system with $d^n - 1$-fold degeneracy in the excited state [15]. Although this Hamiltonian is useful to obtain fundamental bounds [15] it involves n-body interactions and is hence highly complex for $n \gg 1$. Nonetheless, it can be well approximated through two-body interactions by the method developed in Ref. [52] and, furthermore, it can be effectively realized with a few-fermionic mixture confined in a one-dimensional harmonic trap [41]. Motivated by this progress, at the kth round we restrict to the class of Hamiltonians $H_n^{(k)}$ with the aforementioned two-level structure, and tune the energy gap to minimize the EMSLE (2). As we show in the example below, we can achieve a quadratic scaling with n and saturate (7) using this strategy [result (iii)].

Case study.—The results presented here are valid for a broad class of priors, but in what follows we stick to a specific choice in order to illustrate their usage. In any relevant application of thermometry, the temperature is known *a priori* to lie within a certain range, i.e., $\theta_{\min} \leqslant \theta_0 \leqslant \theta_{\max}$. We use a family of probability distributions that are suitable in this case and were proposed in Ref. [53]:

$$p(\theta) = \frac{1}{k_\alpha(\theta_{\max} - \theta_{\min})}\left[e^{\alpha \sin^2\left(\pi \frac{\theta-\theta_{\min}}{\theta_{\max}-\theta_{\min}}\right)} - 1\right] \tag{9}$$

with

$$k_\alpha := e^{\alpha/2} I_0(\alpha/2) - 1, \tag{10}$$

where I_0 is the modified Bessel function of the first kind. In the limit $\alpha \to -\infty$ the above prior becomes a constant, while in the limit $\alpha \to 0$ we have $p(\theta) \propto \sin^2(2\theta)$.

The adaptive strategy works as follows. We consider as a resource N qubits, which are divided in m groups of n qubits. In each group, the n-qubit Hamiltonian is engineered to become a two-level system with degeneracy $(2^n - 1)$ and with a tunable gap ϵ. In the first round, we tune the gap to $\epsilon^{(1)}$ to minimize the single shot EMSLE, that is we set $m = 1$ in Eq. (2). Then, we measure the energy of the system. Given the outcome x_1 is observed, we update the prior to $p(\theta) \to p(\theta|x_1)$, and implement the same procedure to choose $\epsilon^{(2)}$ in the second round [i.e., we minimize Eq. (2) replacing $p(\theta) \to p(\theta|x_1)$]. This process is repeated until all probes are used.

In our simulations, we apply the adaptive process for a given θ_0 sampled from $p(\theta)$, which yields a trajectory as illustrated in the left panel of Fig. 2. We see that the prior peaks around the true temperature as k increases, and the estimated temperature gets closer to the true temperature, i.e., $\tilde{\theta}/\theta_0 \to 1$. The average over a large amount of trajectories enables us to compute EMSLE in Eq. (2) with high accuracy (in the numerical simulations, we consider $\mathcal{O}(1000/m)$ trajectories, which ensures convergence). In the right panel of Fig. 3 we plot EMSLE in the adaptive scenario for various values of n, benchmarked against the no-go bound for nonadaptive scenarios—only the shaded area can be accessed by nonadaptive strategies given any $n \leqslant N$. We see that as n increases the error gets smaller for large enough N. In particular, there exist some thresholds n for which one can beat the no-go bound via adaptive strategies. As an example, given $N = 10^3$ and $\theta_{\max}/\theta_{\min} = 10$ in Eq. (9) (with $\alpha = 1$), adaptive strategies using $n \approx 10$ interacting qubits outperform arbitrary nonadaptive strategies.

Next, we ask whether the adaptive strategy can reach the Heisenberg-like scaling, $\text{EMSLE}^{-1} \propto mn^2$. To this aim, we study the behavior of the error with the resources n for a sufficiently large number of repetitions m. The results are depicted in Fig. 3, where we see Eq. (7) is saturated and therefore the proposed adaptive scheme reaches the ultimate bound on thermometry.

Finally, we note that although the optimal protocol requires a very idealized Hamiltonian for the probe [a $(2^n - 1)$-degenerate two-level system], adaptive protocols already become useful for small n. Namely, for $n = 1, 2$, they decrease the error more than 60% and 80%, respectively, compared to the nonadaptive protocols [44]. For larger n, a realistic method to obtain a scaling of the EMSLE beyond the SNL would be to combine the adaptive method derived here with thermal phase transitions [51].

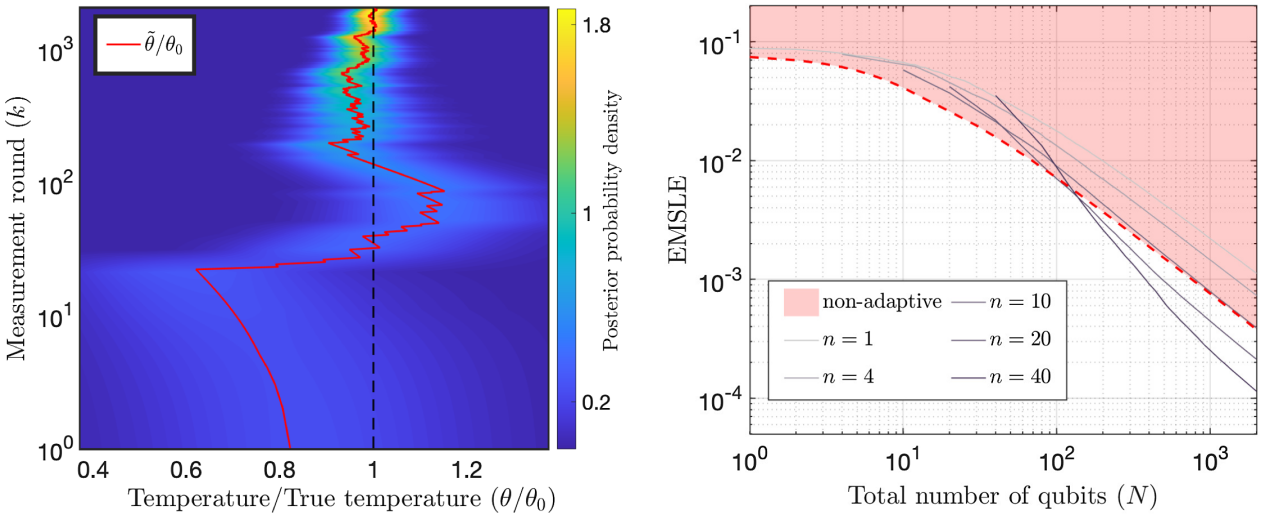

FIG. 2. Left—Contour plot of the prior versus the measurement round $k \in \{1, \ldots, m\}$ (logarithmic scale), and temperature normalized to its true value θ/θ_0. The red trajectory shows the ratio between the estimated temperature and the true temperature $\tilde{\theta}/\theta_0$. As k increases, the prior sharpens around the true temperature, and $\tilde{\theta}/\theta_0$ approaches one. Here, we have set $n = 1$, $\alpha = 1$, $\theta_{\min} = 1$, and $\theta_{\max} = 10$ in arbitrary units. Right—Loglog plot of the expected mean square logarithmic error attained by the adaptive strategy vs the total number of qubits N. Dark solid lines represent different values of n. They show that, for sufficiently large N, the bigger n is the smaller the error can get. The red-dashed line is the (not necessarily tight) bound on nonadaptive strategies: only the shaded area can be achieved using nonadaptive protocols. One can cross the border with adaptive strategies for $n > 10$.

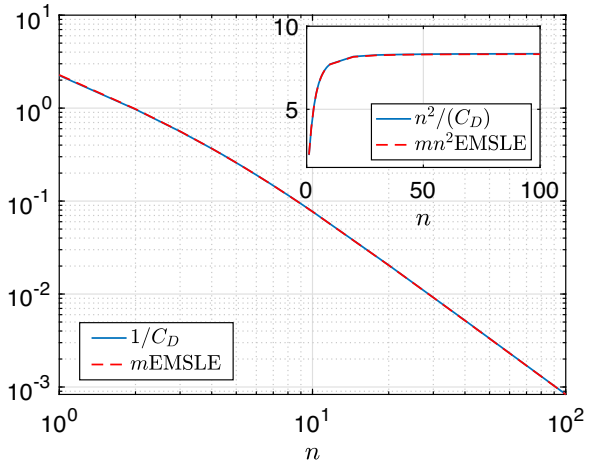

FIG. 3. (Dashed red) Log-log plot of the normalized expected mean square logarithmic error after m rounds of the adaptive scheme—for sufficiently large m, here $m = 2 \times 10^3$—vs n. This shows that for large enough n the error vanishes quadratically with n, which can be better seen from the inset. (Blue) The minimum achievable EMSLE given by the right-hand side of Eq. (7). The perfect agreement shows the efficiency of the proposed adaptive protocol.

Conclusions and future directions.—We derived fundamental limitations of the Bayesian approach to equilibrium thermometry, which shows a Heisenberg-like quadratic scaling with the number of probes. We showed nonadaptive strategies cannot saturate this bound and are limited to shot-noise-like scaling whenever the initial prior is not sharp. We also constructed an adaptive protocol that saturates the ultimate bound, thus highlighting the crucial role of adaptivity in quantum thermometry. This is importantly different to Bayesian phase-estimation protocols [54], where the Heisenberg limit that applies to most general adaptive protocols [55] can be attained by resorting only to measurements being adaptively varied in between the phase-encoding channel uses [56]. In contrast, in equilibrium thermometry the form of probe states (Gibbs) and measurement (energy-basis) is fixed, and it is the probe Hamiltonian that must be adaptively adjusted for the quadratic scaling to become reachable.

While here we considered the total number of probes N as our resource, future works could include time as an extra resource. This naturally leads to non-equilibrium thermometry, where the probe is measured before reaching thermalization. While considerable progress in this framework has been obtained within the frequentist approach [13,25,37,38,57,58], adaptive protocols could be developed following the Bayesian approach pursued here. Lastly, exploiting adaptive schemes for other metrological tasks involving criticality and quantum phase transitions [59], or restrictions such as limited measurement resolution [17,18,60], can be a subject of future work.

We gratefully thank J. Rubio and L. A. Correa for fruitful discussions in an early stage of this work. M. M. and M. P-L. acknowledge financial support from the Swiss National Science Foundation (NCCR SwissMAP and Ambizione Grant No. PZ00P2-186067). J. K. acknowledges the Foundation for Polish Science within the "Quantum Optical Technologies" project carried out within the International Research Agendas programme cofinanced by the European Union under the European Regional Development Fund. M. R. J. and J. B. B. acknowledge support by the Independent Research Fund Denmark.

*mohammad.mehboudi@unige.ch
†matrj@fysik.dtu.dk
‡marti.perarnaullobet@unige.ch

[1] Alessio Celi, Anna Sanpera, Veronica Ahufinger, and Maciej Lewenstein, Quantum optics and frontiers of physics: The third quantum revolution, Phys. Scr. **92**, 013003 (2016).
[2] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger, Many-body physics with ultracold gases, Rev. Mod. Phys. **80**, 885 (2008).
[3] Immanuel Bloch, Jean Dalibard, and Sylvain Nascimbene, Quantum simulations with ultracold quantum gases, Nat. Phys. **8**, 267 (2012).
[4] A. E. Leanhardt, T. A. Pasquini, Michele Saba, A. Schirotzek, Y. Shin, David Kielpinski, D. E. Pritchard, and W. Ketterle, Cooling bose-einstein condensates below 500 picokelvin, Science **301**, 1513 (2003).
[5] Quentin Bouton, Jens Nettersheim, Daniel Adam, Felix Schmidt, Daniel Mayer, Tobias Lausch, Eberhard Tiemann, and Artur Widera, Single-Atom Quantum Probes for Ultracold Gases Boosted by Nonequilibrium Spin Dynamics, Phys. Rev. X **10**, 011018 (2020).
[6] Marco Scigliuzzo, Andreas Bengtsson, Jean-Claude Besse, Andreas Wallraff, Per Delsing, and Simone Gasparinetti, Primary Thermometry of Propagating Microwaves in the Quantum Regime, Phys. Rev. X **10**, 041054 (2020).
[7] Ryan Olf, Fang Fang, G. Edward Marti, Andrew MacRae, and Dan M. Stamper-Kurn, Thermometry and cooling of a bose gas to 0.02 times the condensation temperature, Nat. Phys. **11**, 720 (2015).
[8] Rudolf Gati, Börge Hemmerling, Jonas Fölling, Michael Albiez, and Markus K. Oberthaler, Noise Thermometry with Two Weakly Coupled Bose-Einstein Condensates, Phys. Rev. Lett. **96**, 130404 (2006).
[9] Alberto Ronzani, Bayan Karimi, Jorden Senior, Yu-Cheng Chang, Joonas T. Peltonen, ChiiDong Chen, and Jukka P. Pekola, Tunable photonic heat transport in a quantum heat valve, Nat. Phys. **14**, 991 (2018).
[10] F. M. Spiegelhalder, A. Trenkwalder, D. Naik, G. Hendl, F. Schreck, and R. Grimm, Collisional Stability of ^{40}K Immersed in a Strongly Interacting Fermi Gas of ^{6}Li, Phys. Rev. Lett. **103**, 223203 (2009).
[11] Kuan Yen Tan, Matti Partanen, Russell E. Lake, Joonas Govenius, Shumpei Masuda, and Mikko Möttönen, Quantum-circuit refrigerator, Nat. Commun. **8**, 1 (2017).

[12] Daniel Adam, Quentin Bouton, Jens Nettersheim, Sabrina Burgardt, and Artur Widera, Coherent and dephasing spectroscopy for single-impurity probing of an ultracold bath, arXiv:2105.03331.
[13] Mohammad Mehboudi, Anna Sanpera, and Luis A. Correa, Thermometry in the quantum regime: Recent theoretical progress, J. Phys. A **52**, 303001 (2019).
[14] Antonella De Pasquale and Thomas M. Stace, Quantum thermometry, in *Thermodynamics in the Quantum Regime* (Springer, New York, 2018), pp. 503–527.
[15] Luis A. Correa, Mohammad Mehboudi, Gerardo Adesso, and Anna Sanpera, Individual Quantum Probes for Optimal Thermometry, Phys. Rev. Lett. **114**, 220405 (2015).
[16] Matteo G. A. Paris, Achieving the landau bound to precision of quantum thermometry in systems with vanishing gap, J. Phys. A **49**, 03LT02 (2015).
[17] Patrick P. Potts, Jonatan Bohr Brask, and Nicolas Brunner, Fundamental limits on low-temperature quantum thermometry with finite resolution, Quantum **3**, 161 (2019).
[18] Mathias R. Jørgensen, Patrick P. Potts, Matteo G. A. Paris, and Jonatan B. Brask, Tight bound on finite-resolution quantum thermometry at low temperatures, Phys. Rev. Research **2**, 033394 (2020).
[19] Mohammad Mehboudi, Aniello Lampo, Christos Charalambous, Luis A. Correa, Miguel Ángel García-March, and Maciej Lewenstein, Using Polarons for sub-nK Quantum Nondemolition Thermometry in a Bose-Einstein Condensate, Phys. Rev. Lett. **122**, 030403 (2019).
[20] Gerardo A. Paz-Silva, Leigh M. Norris, and Lorenza Viola, Multiqubit spectroscopy of gaussian quantum noise, Phys. Rev. A **95**, 022121 (2017).
[21] J. Ruostekoski, C. J. Foot, and A. B. Deb, Light Scattering for Thermometry of Fermionic Atoms in an Optical Lattice, Phys. Rev. Lett. **103**, 170404 (2009).
[22] Mark T. Mitchison, Thomás Fogarty, Giacomo Guarnieri, Steve Campbell, Thomas Busch, and John Goold, In Situ Thermometry of a Cold Fermi Gas via Dephasing Impurities, Phys. Rev. Lett. **125**, 080402 (2020).
[23] Stella Seah, Stefan Nimmrichter, Daniel Grimmer, Jader P. Santos, Valerio Scarani, and Gabriel T. Landi, Collisional Quantum Thermometry, Phys. Rev. Lett. **123**, 180602 (2019).
[24] Steve Campbell, Mohammad Mehboudi, Gabriele De Chiara, and Mauro Paternostro, Global and local thermometry schemes in coupled quantum systems, New J. Phys. **19**, 103003 (2017).
[25] Sania Jevtic, David Newman, Terry Rudolph, and T. M. Stace, Single-qubit thermometry, Phys. Rev. A **91**, 012331 (2015).
[26] Sholeh Razavian, Claudia Benedetti, Matteo Bina, Yahya Akbari-Kourbolagh, and Matteo G. A. Paris, Quantum thermometry by single-qubit dephasing, Eur. Phys. J. Plus **134**, 284 (2019).
[27] Antonella De Pasquale, Davide Rossini, Rosario Fazio, and Vittorio Giovannetti, Local quantum thermal susceptibility, Nat. Commun. **7**, 12782 (2016).
[28] Karen V. Hovhannisyan and Luis A. Correa, Measuring the temperature of cold many-body quantum systems, Phys. Rev. B **98**, 045101 (2018).
[29] M. Mehboudi, M. Moreno-Cardoner, G. De Chiara, and A. Sanpera, Thermometry precision in strongly correlated ultracold lattice gases, New J. Phys. **17**, 055020 (2015).
[30] S. S. Mirkhalaf, D. Benedicto Orenes, M. W. Mitchell, and E. Witkowska, Criticality-enhanced quantum sensing in ferromagnetic Bose-Einstein condensates: Role of readout measurement and detection noise, Phys. Rev. A **103**, 023317 (2021).
[31] C. L. Latune, I. Sinayskiy, and F. Petruccione, Collective heat capacity for quantum thermometry and quantum engine enhancements, New J. Phys. **22**, 083049 (2020).
[32] Guim Planella, Marina F. B. Cenni, Antonio Acin, and Mohammad Mehboudi, Bath-Induced Correlations Enhance Thermometry Precision at Low Temperatures, Phys. Rev. Lett. **128**, 040502 (2022).
[33] Luis A. Correa, Martí Perarnau-Llobet, Karen V. Hovhannisyan, Senaida Hernández-Santana, Mohammad Mehboudi, and Anna Sanpera, Enhancement of low-temperature thermometry by strong coupling, Phys. Rev. A **96**, 062103 (2017).
[34] Patrick P. Hofer, Jonatan Bohr Brask, Martí Perarnau-Llobet, and Nicolas Brunner, Quantum Thermal Machine as a Thermometer, Phys. Rev. Lett. **119**, 090603 (2017).
[35] Jesús Rubio, Janet Anders, and Luis A. Correa, Global Quantum Thermometry, Phys. Rev. Lett. **127**, 190402 (2021).
[36] Gabriel O. Alves and Gabriel T. Landi, Bayesian estimation for collisional thermometry, Phys. Rev. A **105**, 012212 (2022).
[37] Alexander Holm Kiilerich, Antonella De Pasquale, and Vittorio Giovannetti, Dynamical approach to ancilla-assisted quantum thermometry, Phys. Rev. A **98**, 042124 (2018).
[38] Stella Seah, Stefan Nimmrichter, Daniel Grimmer, Jader P. Santos, Valerio Scarani, and Gabriel T. Landi, Collisional Quantum Thermometry, Phys. Rev. Lett. **123**, 180602 (2019).
[39] Zhixin Wang, Mingrui Xu, Xu Han, Wei Fu, Shruti Puri, S. M. Girvin, Hong X. Tang, S. Shankar, and M. H. Devoret, Quantum Microwave Radiometry with a Superconducting Qubit, Phys. Rev. Lett. **126**, 180501 (2021).
[40] T. Jahnke, S. Lanéry, and G. Mahler, Operational approach to fluctuations of thermodynamic variables in finite quantum systems, Phys. Rev. E **83**, 011109 (2011).
[41] M. Płodzień, R. Demkowicz-Dobrzański, and T. Sowiński, Few-fermion thermometry, Phys. Rev. A **97**, 063619 (2018).
[42] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone, Quantum-enhanced measurements: Beating the standard quantum limit, Science **306**, 1330 (2004).
[43] Mathias R. Jørgensen, Jan Kołodyński, Mohammad Mehboudi, Martí Perarnau Llobet, and Jonata Bohr Brask, companion paper, Bayesian quantum thermometry based on the thermodynamic length, Phys. Rev. A **105**, 042601 (2022).
[44] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.128.130502, which includes Ref. [45], for detailed derivation of the bounds, as well as nonasymptotic behavior of the error.
[45] M. Ronen (to be published).

[46] Harry L Van Trees, *Detection, Estimation, and Modulation Theory, Part I: Detection, Estimation, and Linear Modulation Theory* (John Wiley & Sons, New York, 2004).

[47] Harry L. Van Trees and Kristine L. Bell, *Bayesian Bounds for Parameter Estimation and Nonlinear Filtering/Tracking* (Wiley-IEEE Press, New York, 2007).

[48] Steve Campbell, Marco G. Genoni, and Sebastian Deffner, Precision thermometry and the quantum speed limit, Quantum Sci. Technol. **3**, 025002 (2018).

[49] Ricardo Román-Ancheyta, Barış Çakmak, and Özgür E. Müstecaplıoğlu, Spectral signatures of non-thermal baths in quantum thermalization, Quantum Sci. Technol. **5**, 015003 (2019).

[50] Wai-Keong Mok, Kishor Bharti, Leong-Chuan Kwek, and Abolfazl Bayat, Optimal probes for global quantum thermometry, Commun. Phys. **4**, 62 (2021).

[51] Kerson Huang, *Introduction to Statistical Physics* (Chapman and Hall/CRC, Boca Raton, 2009).

[52] Nicholas Chancellor, Stefan Zohren, Paul A. Warburton, Simon C. Benjamin, and Stephen Roberts, A direct mapping of max k-sat and high order parity checks to a chimera graph, Sci. Rep. **6**, 1 (2016).

[53] Yan Li, Luca Pezzè, Manuel Gessner, Zhihong Ren, Weidong Li, and Augusto Smerzi, Frequentist and bayesian quantum phase estimation, Entropy **20**, 628 (2018).

[54] D. W. Berry and H. M. Wiseman, Optimal States and Almost Optimal Adaptive Measurements for Quantum Interferometry, Phys. Rev. Lett. **85**, 5098 (2000).

[55] W. Górecki, R. Demkowicz-Dobrzański, H. M. Wiseman, and D. W. Berry, π-Corrected Heisenberg Limit, Phys. Rev. Lett. **124**, 030501 (2020).

[56] H. M. Wiseman, D. W. Berry, S. D. Bartlett, B. L. Higgins, and G. J. Pryde, Adaptive measurements in the optical quantum information laboratory, IEEE J. Sel. Top. Quantum Electron. **15**, 1661 (2009).

[57] Vasco Cavina, Luca Mancino, Antonella De Pasquale, Ilaria Gianani, Marco Sbroscia, Robert I. Booth, Emanuele Roccia, Roberto Raimondi, Vittorio Giovannetti, and Marco Barbieri, Bridging thermodynamics and metrology in non-equilibrium quantum thermometry, Phys. Rev. A **98**, 050101(R) (2018).

[58] M. M. Rams, P. Sierant, O. Dutta, P. Horodecki, and J. Zakrzewski, At the Limits of Criticality-Based Quantum Metrology: Apparent Super-Heisenberg Scaling Revisited, Phys. Rev. X **8**, 021022 (2018).

[59] Irénée Frérot and Tommaso Roscilde, Quantum Critical Metrology, Phys. Rev. Lett. **121**, 020402 (2018).

[60] Karen V. Hovhannisyan, Mathias R. Jørgensen, Gabriel T. Landi, Álvaro M. Alhambra, Jonatan B. Brask, and Martí Perarnau-Llobet, Optimal quantum thermometry with coarse-grained measurements, PRX Quantum **2**, 020322 (2021).