Designing optimal protocols in Bayesian quantum parameter estimation with higher-order operations

Jessica Bavaresco ●,[1,*] Patryk Lipka-Bartosik ●,[1,†] Pavel Sekatski,[1,‡] and Mohammad Mehboudi ●[2,§]
[1]*Department of Applied Physics, University of Geneva, Geneva, Switzerland*
[2]*Technische Universität Wien, 1020 Vienna, Austria*

Using quantum systems as sensors or probes has been shown to greatly improve the precision of parameter estimation by exploiting unique quantum features such as entanglement. A major task in quantum sensing is to design the optimal protocol, i.e., the most precise one. It has been solved for some specific instances of the problem, but in general even numerical methods are not known. Here, we focus on the single-shot Bayesian setting, where the goal is to find the optimal initial state of the probe (which can be entangled with an auxiliary system), the optimal measurement, and the optimal estimator function. We leverage the formalism of higher-order operations to develop a method based on semidefinite programming that finds a protocol that is close to the optimal one with arbitrary precision. Crucially, our method is not restricted to any specific quantum evolution, cost function, or prior distribution, and thus can be applied to any estimation problem. Moreover, it can be applied to both single-parameter or multiparameter estimation tasks. We demonstrate our method with three examples, consisting of unitary phase estimation, thermometry in a bosonic bath, and multiparameter estimation of an SU(2) transformation. Exploiting our methods, we extend several results from the literature. For example, in the thermometry case, we find the optimal protocol at any finite time and quantify the usefulness of entanglement.

I. INTRODUCTION

Quantum parameter estimation, also known as quantum metrology or quantum sensing, is at the heart of quantum technologies [1]. The quantitative assessment of some properties of a system, such as magnetic field amplitude, length, temperature or chemical potential, to name a few, is a key task for science and industry. A sensor is a device which manipulates probes interacting with the system of interest in order to readout its properties. Loosely speaking, the sensing becomes quantum, whenever the manipulation of the probes and their interaction with the measured system is governed by quantum physics. Quantum metrology has been very successful in advancing technological frontiers as showcased in several experiments, namely, the detection of gravitational waves [2,3], thermometry [4,5], magnetometry [6,7], and phase estimation in optical platforms [8].

The theory of quantum metrology aims at developing protocols that use optimally the probes and other *metrological resources*, such as quantum correlations, coherence, and

---

*Contact author: jessica.bavaresco@unige.ch
†Contact author: patryk.lipka.bartosik@gmail.com
‡Contact author: pavel.sekatski@gmail.com
§Contact author: mohammad.mehboudi@tuwien.ac.at

measurement time, in order to estimate the parameter with minimal error [9–13], and uncovers ultimate limits on the achievable estimation precision [14–19]. These limits are usually expressed as a bound on the quantum Fisher information (matrix) that must hold in a certain context and are related to the mean-squared error (MSE) via a Cramér-Rao type bound [20–25]. In the single-parameter case, such bounds are often saturable in the regime where the protocol is repeated many times [26,27].

However, in the limit of small data, such bounds are not generally saturable, and furthermore the MSE, addressed by the Carmér-Rao bound, may not be the best quantifier of the estimation precision. Such problems can be attacked from the perspective of the full Bayesian framework. In the Bayesian approach, one starts with a prior distribution (belief) of the parameter and updates it through the protocol based on the observed measurements results. Crucially, the choices of prior distribution and the cost (or reward) function have a substantial impact on the optimal protocol. Finding such optimal Bayesian protocols is one of the key problems in metrology. This is a nontrivial task even in the case of single-shot scenarios, where the protocol is described by the combination of the initial state, the final measurement, and the estimator function. Optimal protocols are only known for a few highly symmetric specific cases (see Refs. [28,29] for a review), and general effective numerical methods for finding them are lacking.

We therefore dedicate this work to address the shortcomings of quantum metrology within the single-shot Bayesian framework. Namely, we exploit the formalism of higher-order operations [30–33] to combine two pivotal aspects of the estimation protocol, the quantum state and the

measurement, referred to as the quantum strategy, into a single and equivalent higher-order transformation, called quantum tester [30,34,35]. While the standard approach to metrology typically involves the optimization over state and measurement individually [36–38], often in a nonefficient, heuristic manner, quantum testers allow us to optimize over the quantum strategy altogether, finding the optimal state and measurement efficiently with a single instance of a semidefinite program (SDP). Originally a tool applied to tasks such as channel discrimination [34,39,40], the higher-order operations formalism was recently extended to quantum parameter estimation problem, both in the frequentist setting in order to maximize the Fisher information of a protocol [41,42] and in the Bayesian setting in order to maximize the probability of a fixed-width credible interval [38]. In this work we focus on single-shot Bayesian setting and show how to leverage the properties of higher-order operations in order to efficiently optimize the estimation protocol with respect to any reward function.

We propose three different methods to integrate the optimization of the quantum strategy, i.e., state and measurement, with the optimization of the estimators, therefore finding the optimal overall protocol within arbitrary precision. Our methods take into account both numerical and practical limitations, finding application in a wide range of realistic scenarios. It is furthermore appropriate to any estimation problem regardless of prior distribution, reward, or cost function, or the type of quantum evolution. Moreover, we show how these methods can be straightforwardly adapted to multiparameter estimation problems.

To demonstrate the merit of our approach, we present three case studies where we apply our methods to relevant parameter estimation problems: phase estimation, thermometry, and SU(2) estimation. These examples cover single and multiparameter problems, both unitary and nonunitary evolution, reward, or cost functions of varying nature (e.g., fidelity and MSE), and different prior distributions (e.g., uniform and Gaussian). Moreover, we use one of our case studies, thermometry, to show how our approach can be adapted to approximate quantum strategies that do not permit for entanglement between the probe and an auxiliary system for their implementation. This allows us to demonstrate that entanglement provides an advantage over no-entanglement strategies in a finite-time temperature estimation task. Our techniques can be similarly used to answer whether entanglement can be useful in other estimation tasks, and put a lower bound on the usefulness of entanglement. In the thermometry problem, we also find the optimal protocol in finite time, which was previously only known in the frequentist regime [43], and show that the estimation precision only decreases with $t \to \infty$.

II. BACKGROUND

A. Bayesian parameter estimation

In a standard metrology problem, one is interested in estimating an unknown parameter θ by encoding it into the quantum state of a probe. The encoding process can be described by a quantum channel, a completely positive and trace-preserving map, which we denote by $\mathcal{E}_\theta : \mathcal{L}(\mathcal{H}^I) \to \mathcal{L}(\mathcal{H}^O)$ where $\mathcal{H}^I$ and $\mathcal{H}^O$ are the Hilbert spaces of the input
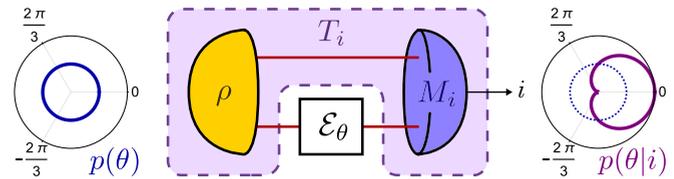
FIG. 1. *Strategy for Bayesian parameter estimation.* The left panel represents the prior probability distribution of the parameter θ encoded in the channel $\mathcal{E}_\theta$. The center panel shows a single-shot strategy of parameter estimation in which part of a quantum state ρ is sent through the channel $\mathcal{E}_\theta$ and then measured by POVM $\{M_i\}$, yielding a classical outcome i. The right panel then represents the posterior probability distribution of the parameter θ, conditioned on the obtained measurement outcome i.

and output systems of the channel, respectively. When probing the channel, it is in general more advantageous to also use an auxiliary system which is initially entangled to the probe, but does not go through the channel, as sketched in Fig. 1. In other words, one considers the extended channel $\mathcal{E}_\theta \otimes \text{id}$, where "id" is the identity channel acting on the auxiliary system. The chosen global input state, given by the density operator $\rho \in \mathcal{L}(\mathcal{H}^I \otimes \mathcal{H}^{\text{aux}})$, is then mapped to a global output state $\rho_\theta := (\mathcal{E}_\theta \otimes \text{id})[\rho]$ by the extended channel. In order to extract the information about the parameter θ encoded in this state, one performs a joint measurement $M = \{M_i\}_{i=1}^{N_O}$, $M_i \in \mathcal{L}(\mathcal{H}^{\text{aux}} \otimes \mathcal{H}^O)$ in the auxiliary system and the output state of the channel. Finally, in the considered setting, one designs an *estimator* $\hat{\theta}$ that assigns an estimate $\hat{\theta}_i$ to the true value of the parameter θ, conditioned on each measurement outcome i. The quality of the estimation can be then quantified by setting some score (cost) function, evaluating the closeness (deviation) of the estimator to the true parameter value. Indeed, the score should depend on the *protocol*; i.e., the triplet of the initial state, the measurement, and the estimator $\{\rho, \{M_i\}_i, \{\hat{\theta}_i\}_i\}$. A central problem in quantum metrology is finding the optimal protocol.

In the Bayesian approach, one starts with a prior belief in the parameter value given by a probability distribution $p(\theta)$. After the measurement, described by the Born rule

$$p(i|\theta) = \text{tr}(\rho_\theta M_i) = \text{tr}[(\mathcal{E}_\theta \otimes \text{id})[\rho] M_i], \quad (1)$$

one uses the Bayes' rule to update the distribution of the parameter based on the observed outcome i

$$p(\theta|i) = \frac{p(i|\theta)p(\theta)}{p(i)}, \quad (2)$$

where the normalization factor is defined as $p(i) := \int d\theta \, p(i|\theta)p(\theta)$.

The performance of the estimation strategy can be quantified according to a *score*. Generally, this can be cast as

$$\mathcal{S} := \sum_i p(i) \int d\theta \, p(\theta|i) \, r(\theta, \hat{\theta}_i) \quad (3)$$

$$= \int d\theta \sum_i p(\theta) r(\theta, \hat{\theta}_i) \, \text{tr}[(\mathcal{E}_\theta \otimes \text{id})[\rho] M_i], \quad (4)$$

where $r(\theta, \hat{\theta}_i)$ is a reward or cost function that quantifies the difference between the parameter θ and each estimate $\hat{\theta}_i$. A

particular choice of cost function is the MSE $r_{\mathrm{MSE}}(\theta, \hat{\theta}_i) = (\theta - \hat{\theta}_i)^2$. In light of this definition, it becomes clear that the optimal protocol will be the one that either maximizes or minimizes the score $\mathcal{S}$, depending on whether $r(\theta, \hat{\theta}_i)$ is a reward or a cost function, respectively.

As previously mentioned, this problem does not have a known analytical solution in general, and efficient numerical methods have only been proposed for a few special single-parameter problems [44,45]. In this work we provide an efficient algorithm that approximates the solution with arbitrary precision, and works for all cost functions and number of parameters.

B. Quantum testers: The quantum strategy as a higher-order operation

A typically cumbersome part of metrology and estimation problems is the optimization of the *quantum strategy*, i.e., of the state and measurement that is used to probe the channel that encodes the parameter to be estimated. Here, we apply techniques from the formalism of higher-order operations [30–33] to fully characterize the set of quantum strategies applicable to a given estimation task. We then use this reformulation to efficiently optimize over quantum strategies using semidefinite programming [46–48]. In particular, we exploit the connection between the states and measurements and an object of the higher-order formalism called a *quantum tester*.

While quantum maps describes transformations of quantum states, higher-order operations (also called supermaps) describe transformations of quantum maps themselves. The equivalent of a positive operator value measure (POVM) in this formalism is a quantum tester, the most general higher-order transformation that maps quantum channels to a probability distribution, effectively "measuring" a quantum channel and yielding a classical outcome with some probability. As illustrated in Fig. 1, a tester T is equivalent to the concatenation of a state ρ and a measurement M. Nevertheless, as we now explain, the Born rule in Eq. (1) becomes linear in the tester variable, which is characterized by simple SDP constraints. We then exploit these two properties to efficiently optimize over the quantum strategies.

In order to express an estimation problem in terms of testers, we start by restating the problem using the Choi-Jamiołkowski isomorphism [49,50]. In this representation, a map $\mathcal{E}_\theta : \mathcal{L}(\mathcal{H}^I) \to \mathcal{L}(\mathcal{H}^O)$ can be equivalently expressed as an operator $C_\theta \in \mathcal{L}(\mathcal{H}^I \otimes \mathcal{H}^O)$, given by $C_\theta = \sum_{ij} |i\rangle\langle j| \otimes \mathcal{E}_\theta[|i\rangle\langle j|]$, called the Choi operator.

Using the Choi operator, the output state of the probe can be expressed as

$$\rho_\theta = (\mathcal{E}_\theta \otimes \mathrm{id})[\rho] = \mathrm{tr}_I((C_\theta \otimes \mathbb{1}^{\mathrm{aux}})(\rho^{\mathrm{T}_I} \otimes \mathbb{1}^O)), \quad (5)$$

where $(\cdot)^{\mathrm{T}_I}$ denotes the partial transposition over the input space $\mathcal{H}^I$. Then, the probability of obtaining outcome i, as in Eq. (1), can be equivalently written as

$$p(i|\theta) = \mathrm{tr}[\mathrm{tr}_I((C_\theta \otimes \mathbb{1}^{\mathrm{aux}})(\rho^{\mathrm{T}_I} \otimes \mathbb{1}^O))M_i] \quad (6)$$

$$= \mathrm{tr}[C_\theta \, \mathrm{tr}_{\mathrm{aux}}((\rho^{\mathrm{T}_I} \otimes \mathbb{1}^O)(\mathbb{1}^I \otimes M_i))]. \quad (7)$$

We can now group the objects that constitute the quantum strategy, that is the state and the measurement, into a

single object called the quantum tester [30,34,35]. A tester $T = \{T_i\}_{i=1}^{N_O}$, $T_i \in \mathcal{L}(\mathcal{H}^I \otimes \mathcal{H}^O)$, is a set of N_O (standing for the "number of outcomes") operators defined as

$$T_i := \mathrm{tr}_{\mathrm{aux}}((\rho^{\mathrm{T}_I} \otimes \mathbb{1}^O)(\mathbb{1}^I \otimes M_i)), \quad (8)$$

which allows one to rewrite the probability of obtaining outcome i, in Eq. (7), as simply

$$p(i|\theta) = \mathrm{tr}(C_\theta \, T_i). \quad (9)$$

The usefulness of this representation comes from the fact that, as shown in Refs. [30,34,35], testers have a simple mathematical characterization. More specifically, they obey the following set of necessary and sufficient conditions:

$$T_i \geqslant 0, \quad \forall \, i \quad (10)$$

$$\sum_i T_i = \sigma \otimes \mathbb{1}^O, \quad (11)$$

where $\sigma \in \mathcal{L}(\mathcal{H}^I)$, $\sigma \geqslant 0$, and $\mathrm{tr}(\sigma) = 1$. It is straightforward to see that every set of operators T that satisfy Eq. (8) also satisfy Eqs. (10) and (11). The converse is also true. Given any set of operators T that satisfy Eqs. (10) and (11), one can define a state ρ and measurement $M = \{M_i\}_{i=1}^{N_O}$ according to

$$\rho := (\mathbb{1}^I \otimes \sqrt{\sigma}) \sum_{ij} |ii\rangle\langle jj| (\mathbb{1}^I \otimes \sqrt{\sigma})^\dagger, \quad (12)$$

$$M_i := (\sqrt{\sigma}^{-1} \otimes \mathbb{1}^O) T_i (\sqrt{\sigma}^{-1} \otimes \mathbb{1}^O)^\dagger, \quad \forall \, i \quad (13)$$

such that

$$\mathrm{tr}_{\mathrm{aux}}((\rho^{\mathrm{T}_I} \otimes \mathbb{1}^O)(\mathbb{1}^I \otimes M_i)) = T_i, \quad \forall \, i. \quad (14)$$

The state ρ and measurement M are called a *quantum realization* of the tester T. This realization is not unique, as different sets of states and measurements can lead to the same tester. However, crucially, different states and measurements that lead to the same tester will also yield the same probability distribution $\{p(i|\theta)\}_i$ in Eq. (1), and have the same performance in an estimation task.

Hence, the optimization of any linear function of $p(i|\theta)$ in Eq. (9) over a tester $T = \{T_i\}$ that satisfies Eqs. (10) and (11) is a semidefinite program, and its optimal tester is guaranteed to have a quantum realization in terms of a quantum state and measurement. Importantly, once the optimal quantum strategy (i.e., tester) is found, the corresponding optimal state and measurement can be easily determined using Eqs. (12) and (13).

Notice that, while a tester is a set of operators that act only on the input and output space of the channel C_θ, its quantum realization may require an auxiliary system. This implies that the optimal quantum strategy may require entanglement between the target and auxiliary systems, and a global measurement that acts on both of these systems. The dimension of the auxiliary space is bounded to be at most the dimension of $\mathcal{H}^I$, as established by the explicit construction of ρ in Eq. (12). The auxiliary system can also be interpreted as a (quantum) memory. Hence, by optimizing over testers, one is effectively optimizing over all possible quantum strategies, including those that may require memory or entanglement for their implementation.

However, certain experimental limitations might induce a situation in which it is necessary to design a quantum strategy that does not require entanglement for its implementation, or a means to certify whether entanglement is indeed advantageous in a given estimation task. In Appendix A we provide details on how quantum strategies that do not require entanglement can be approximated with SDPs. Moreover, in Sec. V B we provide an example of a temperature estimation problem in which our methods demonstrate a clear gap between the performance of strategies operating with and without entanglement.

III. OPTIMAL TESTER FOR METROLOGY VIA SEMIDEFINITE PROGRAMMING

Using quantum testers, we can now rewrite the score of an estimation problem in Eq. (4) as

$$\mathcal{S} = \sum_{i=1}^{N_O} \int d\theta \, p(\theta) \, r(\theta, \hat{\theta}_i) \, \mathrm{tr}(C_\theta \, T_i). \qquad (15)$$

Now, to find the optimal score of a given estimation task, the optimization of $\mathcal{S}$ over the triplet $\{\rho, \{M_i\}, \{\hat{\theta}_i\}_i\}$ can be substituted for an optimization over the pair $\{\{T_i\}, \{\hat{\theta}_i\}_i\}$.

We may express all dependencies of the score $\mathcal{S}$ on the estimates $\{\hat{\theta}_i\}_{i=1}^{N_O}$ with a set of operators $\{X(\hat{\theta}_i)\}_{i=1}^{N_O}$, $X(\hat{\theta}_i) \in \mathcal{L}(\mathcal{H}^I \otimes \mathcal{H}^O)$, which are given by an integral over the parameter θ, defined as

$$X(\hat{\theta}_i) := \int d\theta \, p(\theta) \, r(\theta, \hat{\theta}_i) \, C_\theta, \quad \forall \, i. \qquad (16)$$

These operators encompass all the given information about the task (prior distribution, cost function, and channels in which the parameter is encoded) that does not depend on the quantum strategy. Expressed in terms of these operators, the score is simply

$$\mathcal{S} = \sum_{i=1}^{N_O} \mathrm{tr}(X(\hat{\theta}_i) \, T_i). \qquad (17)$$

For any given set of fixed estimates $\{\hat{\theta}_i\}$, the optimization of the score is given by either a maximization or minimization (depending on the character of the cost function) of $\mathcal{S}$ over all testers T. Taking maximization for instance, we have that

$$\max_{\{T_i\}} \mathcal{S} = \max_{\{T_i\}} \sum_{i=1}^{N_O} \mathrm{tr}(X(\hat{\theta}_i) \, T_i) \qquad (18)$$

is the optimal score. The optimization over testers includes the constraints of Eqs. (10) and (11). Since testers $T = \{T_i\}$ are sets of positive-semidefinite operators characterized by linear constraints, the above optimal score can be efficiently computed using SDP. Once again, the optimal tester is guaranteed to have a quantum realization, hence for any optimal solution of T that the SDP should return, there exist a probe state ρ and measurement M that can realize it; they constitute the optimal quantum strategy for the given estimators.

Notice that this can be straightforwardly generalized to the multiparameter regime as well. In Appendix C we provide more details on this case, while in Sec. V C we present an

example of the application of our methods to the multiparameter problem of SU(2) estimation.

It is now clear that given the knowledge of the estimator values $\hat{\theta}_i$ and the operators $X(\hat{\theta}_i)$ one can find the optimal tester T efficiently. The remaining difficulties thus are as follows:

(i) Finding the optimal estimators $\{\hat{\theta}_i^*\}$ leading to the optimal score

$$\mathcal{S}^* := \max_{\{T_i\}, \{\hat{\theta}_i\}} \mathcal{S}, \quad \{\{T_i^*\}, \{\hat{\theta}_i^*\}\} := \arg\max_{\{T_i\}, \{\hat{\theta}_i\}} \mathcal{S}. \qquad (19)$$

(ii) Computing the integral in Eq. (16).

In the following, we construct three different approaches to tackle both of these problems.

IV. PARAMETER DISCRETIZATION AND ESTIMATOR OPTIMIZATION

In situations where the optimal estimators are unknown, or the integral in Eq. (16) cannot be calculated exactly, an approximation of the optimal score in Eq. (18) can still be computed with SDP. This can be achieved by first discretizing the parameter θ to a finite number of hypotheses, thereby mapping the original parameter estimation task onto one closely resembling channel discrimination.

Concretely, let us choose a discretization of θ such that $\theta \mapsto \{\theta_k\}_{k=1}^{N_H}$, where N_H (standing for the "number of hypotheses") is the total number of different values assigned to θ. We can then define a prior distribution over the new hypotheses as

$$p(\theta_k) := \frac{p(\theta = \theta_k)}{\sum_{k=1}^{N_H} p(\theta = \theta_k)}, \qquad (20)$$

which is computationally straightforward and has the advantage of giving a valid probability distribution.

Now, let us define the discrete equivalent of the operators in Eq. (16) as $\{\widetilde{X}(\hat{\theta}_i)\}_{i=1}^{N_O}$, where

$$\widetilde{X}(\hat{\theta}_i) := \sum_{k=1}^{N_H} p(\theta_k) \, r(\theta_k, \hat{\theta}_i) \, C_{\theta_k}, \quad \forall \, i. \qquad (21)$$

Hence, the approximate score $\widetilde{\mathcal{S}}$ can be expressed as

$$\widetilde{\mathcal{S}} := \sum_{i=1}^{N_O} \mathrm{tr}(\widetilde{X}(\hat{\theta}_i) \, T_i). \qquad (22)$$

The value of $\widetilde{\mathcal{S}}$ will depend on a chosen discretization $\{\theta_k\}$ of the continuous parameter θ: the finer the discretization, the better the approximation. Hence, for a given discretization $\{\theta_k\}$, the optimum score is given by either maximizing or minimizing $\widetilde{\mathcal{S}}$ again over the pair $\{\{\hat{\theta}_i\}_i, \{T_i\}_i\}$ of estimates and testers.

In what follows we propose three different methods, all based on semidefinite programming, with which this approximation can be computed.

A. Method 1: Approximating metrology with channel discrimination

The first approach we propose is heavily based on the problem of channel discrimination [51]. Its starting point is

the realization that, without loss of generality, we may restrict ourselves to testers with as many measurement outcomes as there are hypotheses to be distinguished. In the context of our discretized parameter estimation problem, this amounts to setting $N_O = N_H$; essentially, there is no advantage in increasing the number of measurement outcomes beyond the number of different values in the discretization of θ. The second simplification is to choose the values of the estimates $\{\hat{\theta}_i\}$ to be the same as the values in the discretization $\{\theta_k\}$, in such a way that each measurement outcome i is directly associated to a value $\hat{\theta}_i = \theta_i$. Choosing the values of the potential estimates of the parameter to correspond to the values in the discretization of the parameter reduces the estimation problem to a discrimination problem. In this case, the task can be interpreted as determining the "classical" label k that is encoded via the values of θ_k in the channel C_{θ_k}. In this case, the set of operators $\{\widetilde{X}(\hat{\theta}_i)\}$ becomes

$$\widetilde{X}(\hat{\theta}_i) = \widetilde{X}(\theta_i) = \sum_{k=1}^{N} p(\theta_k)\, r(\theta_k, \theta_i)\, C_{\theta_k}, \quad \forall\, i \tag{23}$$

where $N = N_O = N_H$, and the approximate score $\widetilde{\mathcal{S}}$ becomes

$$\widetilde{\mathcal{S}} = \sum_{i=1}^{N} \mathrm{tr}(\widetilde{X}(\theta_i)\, T_i). \tag{24}$$

For this fixed value of the discretization $\{\theta_i\}_{i=1}^{N}$, the optimum value of $\widetilde{\mathcal{S}}$ over all testers T is an SDP.

This approach circumvents problem (i), of finding the optimal estimators, by setting them to be the same values used in the discretization of the continuous parameter θ; and problem (ii), of computing the integral in Eq. (16), by discretizing it. In principle, the higher the number of values in the discretization of θ, the closer the estimates are to the optimal estimator. The advantage then is that the optimal score can be found with a single SDP that needs to optimize only over the quantum strategies. The drawback, on the other hand, is that to achieve a good approximation of the optimal estimator, a high number N of values in the discretization are necessary, and since this number is directly associated to the number of measurement outcomes in the quantum strategy, the problem can eventually become intractable numerically and experimentally. In practice, however, as demonstrated in our examples in Sec. V, this method yields very good results with a value of N that can still be straightforwardly handled numerically.

Nevertheless, our next approach is designed to overcome this problem as well.

B. Method 2: Parameter discretization with optimal estimator

One possible way to overcome computational challenges is to fix the number of measurement outcomes N_O, and hence the number of tester elements, to a value that is computationally (and experimentally) tractable and increase the value of N_H far beyond that. Since, in this case, the complexity of the problem does not depend on N_H, the discretization of θ can be arbitrarily fine. However, because the number of values in the discretization of θ can far surpass the number of measurement outcomes, the association of one estimate $\hat{\theta}_i$ to each

discretization value θ_k is no longer possible. Therefore, the problem of choosing a "good" set of estimates $\{\hat{\theta}_i\}$ is crucial.

Let us start by assuming that the optimal estimator is known to be $\{\hat{\theta}_i^*\}$. Then, the operators $\{\widetilde{X}(\hat{\theta}_i^*)\}$ amount to

$$\widetilde{X}(\hat{\theta}_i^*) := \sum_{k=1}^{N_H} p(\theta_k)\, r(\theta_k, \hat{\theta}_i^*)\, C_{\theta_k}, \quad \forall\, i \tag{25}$$

for a fixed discretization $\{\theta_k\}_{k=1}^{N_H}$, which can now in principle contain an arbitrarily high number of values $N_H \gg N_O$. The approximate score $\widetilde{\mathcal{S}}$ then becomes

$$\widetilde{\mathcal{S}} = \sum_{i=1}^{N_O} \mathrm{tr}(\widetilde{X}(\hat{\theta}_i^*)\, T_i). \tag{26}$$

The optimization of $\widetilde{\mathcal{S}}$ over the quantum strategy is then given by an SDP.

This approach essentially takes care of problem (ii), of numerically computing the integral in Eq. (16), by discretizing the parameter θ in an arbitrarily fine manner, while maintaining the number of measurements low enough to decrease the computational demand of the SDP. Hence, it is better suited for a situation in which the optimal estimator is known. It can nevertheless also be applied to a problem in which only a good guess for the optimal estimator is known, in which case the solution will be an approximation of the optimal score. Otherwise, to overcome problem (i) of finding the optimal estimator in the first place, we combine this approach with an estimator optimization in a seesaw algorithm, detailed in the following.

C. Method 3: Parameter discretization with estimator optimization

This final approach consists of a seesaw between two optimization problems, which are not necessarily SDPs, that will approximate an optimization over both the quantum strategy and the values of the estimates.

A seesaw is an iterative method that alternates between two optimization problems, using the solution of one as the input of the other. In our case, the first optimization problem is the SDP of the previous approach (method 2). Namely, given $\{\hat{\theta}_i\}_{i=1}^{N_O}$,

$$\max_{\{T_i\}}\ \sum_{i=1}^{N_O} \mathrm{tr}(\widetilde{X}(\hat{\theta}_i)\, T_i), \tag{27}$$

where $\{T_i\}$ is an N_O-outcome tester.

The second optimization problem will then be one that, for a fixed tester $\{T_i\}$, taken to be the optimal tester of the previous SDP, optimizes over the values $\{\hat{\theta}_i\}$ of the estimates. Namely, given $\{T_i\}_{i=1}^{N_O}$,

$$\max_{\{\hat{\theta}_i\}}\ \sum_{i=1}^{N_O} \mathrm{tr}(\widetilde{X}(\hat{\theta}_i)\, T_i), \tag{28}$$

where $\{\hat{\theta}_i\}$ are N_O possible values of θ.

Whether the problem in Eq. (28) is an SDP will depend on whether the reward function $r(\theta_k, \hat{\theta}_i)$ is linear on $\{\hat{\theta}_i\}$. In practice, this will often not be the case. Nevertheless, in some

cases this problem can be solved analytically, depending on the form of the reward function, the optimal estimator may be known or it may be found by standard Lagrangian optimization methods. In other cases, heuristic optimization methods may be applied.

This iterative method, although even for a fixed discretization $\{\theta_k\}$ does not necessarily converge to the optimal value of $\widetilde{S}$, in practice leads to very good approximations. A relevant point here is that, assuming a situation where the seesaw does converge to the optimal estimator, one may restrict themselves without loss of generality to a maximum number of outcomes N_O that are related to the extremality properties of the tester. In principle, since (i) the set of testers $T = \{T_i\}$ is convex and (ii) the function $\widetilde{S}$ is linear on each tester element T_i, the maximum (or minimum) of $\widetilde{S}$ will be achieved by an extremal tester. Analogously to extremal POVMs [52], extremal testers have at most d^2 (nonzero) elements, where d is the dimension of the space upon which the tester (or POVM) elements act. Hence, the number of outcomes N_O in the seesaw can be fixed to be at most $N_O \leqslant (d_I \times d_O)^2$, where $d_{I(O)} := \dim(\mathcal{H}^{I(O)})$ since, for optimal estimators, there is no advantage in optimizing over nonextremal testers. This fact also holds for method 2 if one is guaranteed to know the optimal estimator.

Furthermore, if the cost function is the mean-squared error r_{MSE}, then the optimal measurement will be projective (see Appendix A in Ref. [45]), and hence the optimal tester will have at most $(d_I \times d_O)$ outcomes. We present a case study in Sec. V B, which concerns the problem of thermometry, that precisely falls in this case.

D. Convergence of the methods

In all three methods above, we encounter some error due to discretization of the integral in finite hypotheses, as well as suboptimality due to our choice of estimators. The discretization error is expected to vanish as N_H increases since all three methods are based on approximating an integral with a Riemannian sum with an error that vanishes as $1/N_H$. As for the suboptimality, let us define the best approximate score $\widetilde{S}^*$ similarly to Eq. (19) but for the approximate score defined in Eq. (22). For large enough N_H, and when the cost function is supposed to be maximized, it means that $\widetilde{S}^* \leqslant S^*$ while for cost functions that are supposed to be minimized, it means that $\widetilde{S}^* \geqslant S^*$. The suboptimality roots from the fact that, none of the methods simultaneously optimize over both $\{\hat{\theta}_i\}$ and $\{T_i\}$. Each of the three methods, however, deals with suboptimality differently. In all three methods one has

$$|\widetilde{S}^* - S^*| = O\left(\frac{1}{N_O}\right) \tag{29}$$

and thus one can guarantee convergence by choosing $N_O \gg 1$. In Appendix B we rigorously derive the convergence for arbitrary cost functions and furthermore show that for certain cost functions that we will later use in the case studies (Examples 1 and 2) the convergence is even faster, i.e., $|\widetilde{S}^* - S^*| = O(1/N_O^2)$. When we cannot arbitrarily increase N_O, methods 2 and 3 come to the rescue. In particular, if *a priori* we know what the optimal estimators are, then method 2 allows to find the optimal testers in one shot. However, it is rarely the case that we do know the optimal estimators *a priori*. Nonetheless,

as we see in the examples below, method 2 typically finds suboptimal solutions that are very close to the optimum. Method 3, on the other hand, adds a powerful layer of optimization based on a seesaw between the estimators and testers, and therefore has a higher chance of finding the optimal protocol even with $N_O = (d_I \times d_O)^2$.

V. CASE STUDIES

Our methodology for solving the Bayesian parameter estimation problem using higher-order operations offers numerous advantages over conventional techniques in quantum metrology. By proposing to optimize over the input state and measurement with a single SDP, and combining this with effective heuristics for the joint optimization of the quantum strategy and estimator, we overcome the longstanding challenges of the Bayesian approach. Our approach provides a comprehensive and versatile set of techniques that can be applied to any Bayesian estimation problem, setting it apart from most existing methods in the literature.

The key strength of our approach lies in its ability to handle a wide range of estimation problems, without being limited to specific error quantifiers. This universality allows our method to be seamlessly applied to any estimation scenario. Moreover, the techniques we described here are equally effective for single-parameter and multiparameter estimation tasks. Finally, unlike most techniques in the Bayesian approach, our methods are not bound by the type of dynamics used to encode the parameter. Whether the parameter is encoded via a unitary evolution (e.g., phase estimation), or a more complex open system dynamics resulting from the probe's interaction with a thermal environment (e.g., quantum thermometry), our approach can be systematically applied and, as we show in the following, delivers consistent results which are very close to the optimal values.

We now delve into the practical application of our methods and explore how they can be applied to determine the optimal estimation strategy for various scenarios encountered in quantum metrology. The examples are deliberately chosen to be cover a wide range of different problems to demonstrate the versatility of our methods.

A. Example 1. Paradigmatic example: Local phase estimation

We start with a paradigmatic task in quantum metrology, namely, the single-parameter unitary phase estimation. In this problem a single parameter $\theta \in [0, 2\pi)$ is encoded in an n-qubit quantum system via a local unitary channel $\mathcal{E}_\theta[\cdot] = U_\theta(\cdot)U_\theta^\dagger$ with

$$U_\theta = e^{-i\theta S_z}, \tag{30}$$

where S_z is a collective spin operator $S_z = \frac{1}{2}\sum_{i=1}^n \sigma_z^{(i)}$ and $\sigma_z^{(i)}$ is the Pauli-Z matrix of the ith qubit. Due to the intrinsic symmetry of the problem, every state of the n-qubit system can be effectively described using an $(n+1)$-dimensional Hilbert space, i.e., the symmetric subspace [53]. Consequently, the n-qubit phase estimation problem can be equivalently mapped into a phase estimation problem of a d-dimensional system with $d = n + 1$. In this representation the generator

of the dynamics S_z expressed in the computational basis $\{|0\rangle, |1\rangle, \ldots, |n\rangle\}$ is given by $S_z = \sum_{i=0}^{n} i|i\rangle\langle i|$ [54].

For this example, we take a typical reward function in phase estimation, which takes into account the cyclicity of its parameter space, given by

$$r(\theta, \hat{\theta}_i) = \cos^2\left(\frac{\theta - \hat{\theta}_i}{2}\right). \quad (31)$$

We also choose two different priors, one given by a uniform distribution according to

$$p(\theta) = \frac{1}{\theta_{\max} - \theta_{\min}}, \quad (32)$$

where $\theta_{\min} = 0$ and $\theta_{\max} = 2\pi$ are, respectively, the minimal and maximal values of the parameter. The other distribution is given by a Gaussian distribution, according to

$$p(\theta) = \frac{1}{\mathcal{N}} \exp \frac{(\theta - \mu)^2}{(2\sigma^2)}, \quad (33)$$

where $\mathcal{N}$ is the normalization factor. We set the mean $\mu = \pi$ and the deviation $\sigma = 1$. For the discretization of the parameter θ and initial value of the estimators, we fix

$$\theta_k = \frac{\theta_{\max} - \theta_{\min}}{N_H}(k-1), \quad \forall\, k \in \{1, \ldots, N_H\} \quad (34)$$

and

$$\hat{\theta}_i = \frac{\theta_{\max} - \theta_{\min}}{N_O}(i-1), \quad \forall\, i \in \{1, \ldots, N_O\} \quad (35)$$

respectively.

We now discuss how to apply each of our methods to infer the optimal protocol in this case and present the results obtained for the problem of $(n = 2)$-qubit phase estimation, plotted in Fig. 2.

Method 1. To apply Method 1, we simply set $N_H = N_O = N$, implying $\hat{\theta}_i = \theta_i$. For the results plotted in Fig. 2, we take $N \in \{2, \ldots, 10\}$.

Method 2. Here we fix the number of outcomes $N_O \in \{2, \ldots, 10\}$ and set the number of hypotheses to be $N_H = 1000 \gg N_O$. We then discretize the parameter and set the estimators according to Eqs. (34) and (35). For the case of a uniform prior [Eq. (32), Fig. 2(a)], these are expected to be the optimal estimators. For the case of a Gaussian prior [Eq. (33), Fig. 2(b)], these estimators are not expected to be optimal, but are nonetheless used in method 2, serving as a starting point for the estimator optimization in method 3.

Method 3. For this method, we take the solution for the testers found using method 2 for each N_O as a starting point, and then optimize over the estimator. We prove in Appendix B that in this case the optimal estimators are given as a function of the optimal tester, according to

$$\hat{\theta}_i^* = \begin{cases} \arctan\left(\frac{\langle\sin(\theta_k)\rangle^{(i)}}{\langle\cos(\theta_k)\rangle^{(i)}}\right), & \langle\cos(\theta_k)\rangle^{(i)} \geqslant 0 \\ \arctan\left(\frac{\langle\sin(\theta_k)\rangle^{(i)}}{\langle\cos(\theta_k)\rangle^{(i)}}\right) + \pi, & \langle\cos(\theta_k)\rangle^{(i)} < 0 \end{cases} \quad (36)$$

where

$$\langle f(\theta_k)\rangle^{(i)} = \sum_k p(\theta_k|i) f(\theta_k) = \sum_k \frac{p(\theta_k)\mathrm{tr}(C_{\theta_k} T_i)}{p(i)} f(\theta_k) \quad (37)$$

(a)

(b)
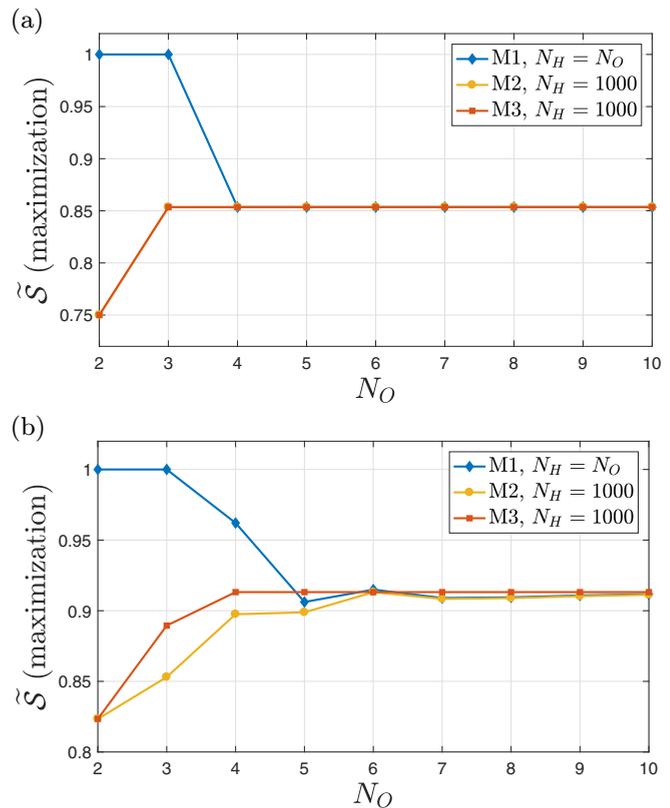
FIG. 2. *Local phase estimation (Example 1).* The maximum approximate score $\widetilde{S}$ in a local $n = 2$ qubit phase estimation problem. Each panel shows the scores corresponding to methods M1, M2, and M3 as a function of the number of outcomes $N_O \in \{2, \ldots, 10\}$ for different prior distributions of the local phase: (a) corresponds to the case of uniform prior, while (b) corresponds to a Gaussian prior. The phase parameter ranges from $\theta_{\min} = 0$ to $\theta_{\max} = 2\pi$. The considered cost function is the cosine squared in Eq. (31).

with $p(i) = \sum_k p(\theta_k)\mathrm{tr}(C_{\theta_k} T_i)$. Note that with the definition of Eq. (36) the range of the estimator is $[-\frac{\pi}{2}, \frac{3\pi}{2}]$, instead of $[0, 2\pi]$ as defined initially, this has no effect on the expected reward and can be resolved by adding 2π if $\hat{\theta}_i^* < 0$. Hence, the second step in the seesaw is solved analytically, and in each round of the seesaw we update the value of the estimators according to the expression above, as a function of the testers found by the SDP in the first step. Here and in the following examples, we iterate these two steps until the gap between the value of the score in subsequent rounds is smaller than 10^{-6}.

Results. In Fig. 2 we plot the maximal approximate scores $\widetilde{S}$ obtained via the three methods outlined above. In the case of the uniform prior [Fig. 2(a)] we observe that the approximate score very quickly reaches the optimal one, i.e., $\widetilde{S} \approx S = \frac{1}{2}[1 + \cos(\pi/4)]$ which was formerly obtained with alternative methods [28,55]. All three methods converge very quickly to this solution, already for $N_O = 3$ outcomes, using methods 2 and 3, and for $N_O = 4$ outcomes, using method 1. Notice also that, since we start already at the optimal estimator, we observe that for all N_O there is no advantage of applying the seesaw in method 3 since method 2 already returns the optimal solution. For the case of the Gaussian prior [Fig. 2(b)], we see that again method 3 converges to a

stable value of $\widetilde{S}$ with $N_O = 4$ outcomes. Here we can see an initial difference between methods 1 and 2, which take a fixed estimator, and method 3, which optimizes over the estimator. Nevertheless, all methods quickly converge to approximately the same value, at $N_O = 10$.

B. Example 2. Nonunitary evolution: Thermometry

Let us now discuss a different instance of single-parameter estimation, namely, thermometry [56,57]. In this case, the unknown parameter is the temperature θ of a sample (or a thermal bath) that is resting at thermal equilibrium, and it is encoded in the probe using a nonunitary quantum channel. We consider the probe to be a two-level system (qubit) which is potentially entangled with an auxiliary system, that does not undergo the nonunitary dynamics. At the initial time $t = 0$, the probe and the sample which are initially uncorrelated start to interact. After some fixed time t, the probe and the auxiliary system will be jointly measured to infer the temperature of the bath. The probe's reduced state $\rho_\theta^p = \mathrm{tr}_A[\rho_\theta]$ evolves according to a standard Markovian quantum master equation [58–61], i.e.,

$$\dot{\rho}_\theta^p(t) = -i[H, \rho_\theta^p(t)] + \Gamma_{\mathrm{in}}\mathcal{D}[\sigma_+]\rho_\theta^p(t) + \Gamma_{\mathrm{out}}\mathcal{D}[\sigma_-]\rho_\theta^p(t),$$
(38)

where $H = \epsilon|1\rangle\langle 1|$ is the Hamiltonian of the probe, $\sigma_- = |0\rangle\langle 1|$ and $\sigma_+ = |1\rangle\langle 0|$ are the jump operators, and $\mathcal{D}[A]\rho_\theta^p = A\rho_\theta^p A^\dagger - \frac{1}{2}\{A^\dagger A, \rho_\theta^p\}$ is the dissipator superoperator which captures the effect of the environment on the probe. The dissipation rates Γ_{in} and Γ_{out} are the only temperature-dependent parts of the dynamics and are responsible for encoding the parameter. For a bosonic and fermionic environment, we have $\Gamma_{\mathrm{in}} = J(\epsilon)N_{B/F}$ and $\Gamma_{\mathrm{out}} = J(\epsilon)(1 \pm N_{B/F})$, where minus sign should be used for fermions, and positive sign for bosons, with $J(\epsilon)$ being the bath spectral density while $N_{B/F}$ is the occupation number for the bosonic or fermionic bath, defined as $N_B = (e^{\epsilon/\theta} - 1)^{-1}$ and $N_F = (e^{\epsilon/\theta} + 1)^{-1}$, respectively. In what follows, we focus on the bosonic bath, however, our methods can be applied to the fermionic case as well.

The evolution specified by Eq. (38) generates an effective quantum channel $\mathcal{E}_\theta(t)$ that imprints the temperature into the probe's state (see Appendix D for the explicit expression). Note that in our notation we keep the time dependence because we are also interested in the optimal protocol at different times.

As for the cost function, we use the MSE

$$r(\theta, \hat{\theta}_i) = (\theta - \hat{\theta}_i)^2,$$
(39)

while the prior distribution $p(\theta)$ is uniform and given by Eq. (32), where we set $\theta_{\mathrm{min}} = 0.1$ and $\theta_{\mathrm{max}} = 2$ as the minimum and maximum values of the temperature.

We discretize the temperature parameter θ and fix the estimators according to Eqs. (34) and (35), respectively. We evaluate the thermometry problem for 100 different time steps, evenly distributed between $t = 0$ and 1.

Let us now discuss how to approach this problem using each of the three methods presented in this work.
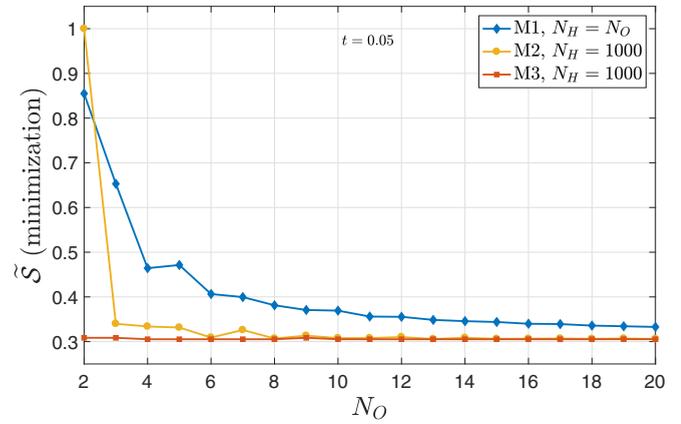
FIG. 3. *Thermometry (Example 2).* The minimum approximate score $\widetilde{S}$ in the finite-time temperature estimation problem $\widetilde{S}$, renormalized by the maximum value of $\widetilde{S}$ in the plot. The temperature θ is encoded via a qubit nonunitary evolution specified by Eq. (38) acting for an amount of time $t < \infty$, here shown for a fixed time $t = 0.05$. The plot shows the different scores corresponding to methods M1, M2, and M3 as a function of the number of outcomes $N_O \in \{2, \ldots, 10\}$ for a uniformly distributed prior in a temperature parameter range of $\theta_{\mathrm{min}} = 0.1, \theta_{\mathrm{max}} = 2$. The considered cost function is the MSE in Eq. (39). The remaining parameters chosen are $\epsilon = 0.1$ and $J(\epsilon) = 2$.

Method 1. To apply method 1, we again simply set $N_H = N_O = N$, implying $\hat{\theta}_i = \theta_i$. For the results presented here we take $N \in \{2, \ldots, 20\}$.

Method 2. Here we fix the number of outcomes $N_O \in \{2, \ldots, 20\}$ and set the number of hypotheses to be $N_H = 1000 \gg N_O$. These values for the estimator are not expected to be optimal but are nevertheless used in method 2, serving as a starting point for the estimator optimization in method 3.

Method 3. To apply the seesaw in method 3, we again begin with the solution provided by method 2 for each N_O and t as a starting point. For the thermometry problem, as was the case for the phase estimation problem, we can analytically express the optimal estimator as a function of the quantum strategy (tester). Since the score is quantified using the MSE, the optimal estimator is simply the mean over the posterior

$$\hat{\theta}_i = \langle\theta_k\rangle,$$
(40)

where $\langle\theta_k\rangle$ is given by Eq. (37). Once again, in this example the first step of the seesaw consists in an optimization over the testers while the second step consists in reassigning a value to the estimators as a function of the testers found in the previous step, according to the above expression.

Results. In Fig. 3 we compare the performance of the three methods outlined above as a function of the number of outcomes N_O for a fixed time $t = 0.05$. We observe that methods 1 and 2 start from a relatively large $\widetilde{S}$, which is now being minimized, that gradually decreases with increasing N_O, while method 3 already starts at a value of $\widetilde{S}$ close to where it will converge. While method 2 quickly converges to the same values of $\widetilde{S}$ of method 3 with increasing N_O, at $N_O = 20$ the approximate score predicted by method 1 is still somewhat above the corresponding one obtained using method 3. This is a result of the error in approximating the operators $X(\hat{\theta}_i)$ using
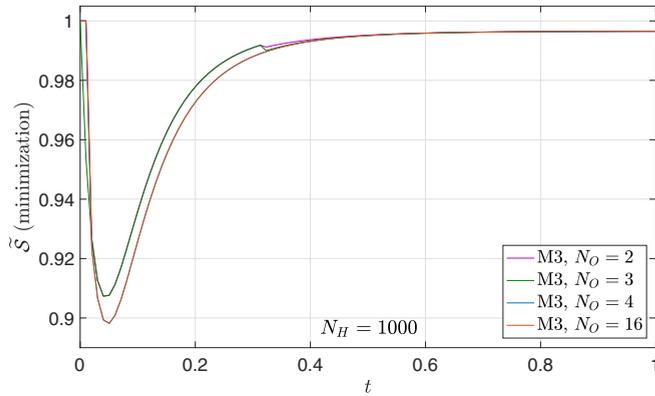
FIG. 4. *Thermometry (Example 2): Optimality of projective measurements.* The approximate score $\widetilde{S}$ is computed using method 3 as a function of time t for different values of N_O. The parameters are chosen as in Fig. 3. The inset plot is a log-log plot of the same curves. All values are renormalized by the maximum value of $\widetilde{S}$ in the plot. Since the cost function is the MSE, in this case projective measurements (which have at most $d_I \times d_O$ outcomes) are optimal. Indeed, we observe that increasing N_O beyond 4 does not change the value of the score.

a Riemannian sum with N_H elements. Indeed, it is guaranteed that only in the limit $N_H \to \infty$ the approximate score in method 1 converges to the true optimal score. Finally, we observe that method 3 saturates around its optimal value already for $N_O = d_I \times d_O = 4$ outcomes. This is expected since, for a MSE cost function, projective measurements are optimal [45]. In Fig. 4, we focus on this point, by comparing method 3 for some fixed values of N_O as a function of time t, for the whole interval of time evaluated in this problem. Here we can see that while there is an improvement in increasing N_O up to 4, the curve for $N_O = 16$ lies on top of that of $N_O = 4$, demonstrating that there is indeed no advantage in increasing the number of outcomes beyond $N_O = d_I \times d_O$.

Another interesting point to make here is that the score clearly depends on time of evolution. In particular, in Fig. 4 we observe that there exist times $t < \infty$ where the score is much better (recall that for a MSE cost function, the optimal score is being minimized) than at $t \to \infty$. This means that measuring the probe in the transient regime can be advantageous over estimation performed after reaching the steady state. Indeed, this effect has been observed in thermometry previously [62]. For the steady state, the optimal measurement strategy is known. In particular, in this case the auxiliary system is useless and the optimal measurement is a projector-valued measure (PVM) in the basis of the probe Hamiltonian [63]. However, as we discuss in the next paragraph, here we show that this is not the case for the transient regime, where entanglement with the probe leads to a more precise estimation. Our techniques therefore allow to determine the optimal probe and memory state, as well as the measurements and estimators in this difficult regime.

Finally, in order to investigate the role of entanglement in the transient regime of temperature estimation, we compare two different measurement strategies: a general strategy where the memory can be initially entangled with the probe, and the scenario where the initial state of the memory and probe is
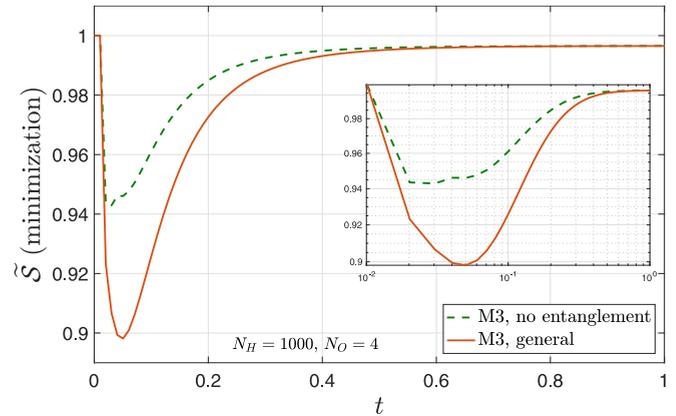
FIG. 5. *Thermometry (Example 2): Advantage of entanglement in the transient regime.* The main panel shows the approximate score $\widetilde{S}$ (computed via method 3) as a function of the evolution time t for a fixed number of outcomes $N_O = 4$. The parameters chosen are as in Fig. 3. All values are renormalized by the maximum value of $\widetilde{S}$ in the plot. The inset shows the same curves plotted in a log-log scale. We see that there are times t for which the precision of estimation is better than in the steady state $t \to \infty$. This can be understood as an advantage arising from having entanglement with the probe: the entanglement allows the transfer of the information about the parameter into the memory system which is itself not subject to the dephasing dynamics of the master equation. As a consequence, measuring the entangled probe and memory systems before the joint system thermalizes provides a significant advantage.

separable. Figure 5 highlights the importance of the entangled auxiliary qubit. To this aim, we have focused again on method 3, and depict the approximate score as a function of time, for strategies with and without entanglement. As one expects, at the limits of very short time or very long time the two kinds of strategies perform equally. While in the former this is because there has not been enough time to collect and add new information to the prior, in the latter case it is because after a long time, the system reaches a steady state regardless of the input state, namely, of it being entangled or not. However, at the transient regime, we observe that an auxiliary system entangled with the probe can significantly improve the score.

Let us emphasize that very often the parameter estimation problem described above cannot be solved analytically and is very difficult to solve numerically. In general, the effective evolution of the probe may result from a complicated master equation, which has to be evaluated many times. In our approach there is no need to evaluate the evolution for each potential state of the probe, as the only thing we need are the Choi operators associated with the effective channels. In this sense, our methods only require to solve the dynamics on a finite grid of parameter values, and thus makes finding the solution more tractable numerically.

Finally, for the method 3 we investigate the optimal initial probe-ancilla state ρ found by our algorithm for different times $t \in [0, 1]$. We consider the case with four outcomes $N_O = 4$, which is found to give the optimal precision. Here the tester has four elements $\{T_1, \ldots, T_4\}$, and the initial state can be computed with the help of Eqs. (11) and (12). Note that the state $\rho = |\Psi\rangle\langle\Psi|_{IA}$ is pure by construction, and its
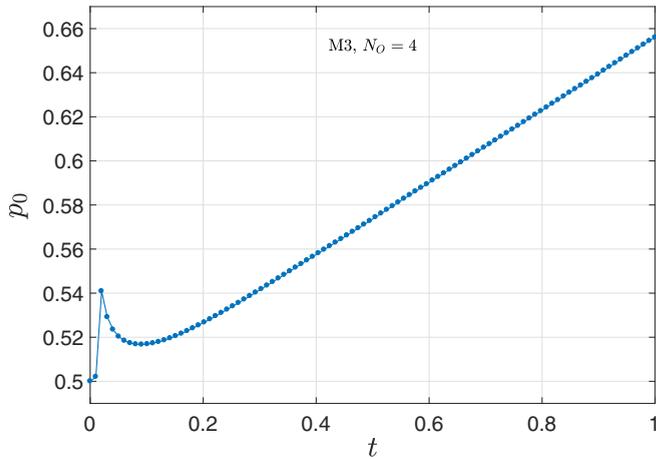
FIG. 6. *Thermometry (Example 2): Optimal state.* The entanglement of the optimal initial state $\rho = |\Psi\rangle\langle\Psi|_{IA}$ in Eq. (41) as function of time found by method 3 for $N_O = 4$. The corresponding score $\widetilde{S}$ is given in Fig. 4, the physical parameters are given in Fig. 3.

Schmidt diagonal form reads as (for $p_0 \geqslant p_1$)

$$|\Psi\rangle_{IA} = \sqrt{p_0}|\boldsymbol{n}\rangle_I|0\rangle_A + \sqrt{p_1}|-\boldsymbol{n}\rangle_I|1\rangle_A, \qquad (41)$$

where $\boldsymbol{n}$ is the Bloch vector corresponding to the state $|\boldsymbol{n}\rangle\langle\boldsymbol{n}| = \frac{1}{2}(\mathbb{1} + \boldsymbol{n}^T\boldsymbol{\sigma})$ and basis choice of the ancillary qubit plays no physical role. We find that for all times the optimal state is always Schmidt diagonal in the computational basis for the probe, and the Schmidt state corresponding the the larger value p_0 is always the ground state, i.e., $|\boldsymbol{n}\rangle_I = |0\rangle_I$. In contrast, the amount of entanglement in the optimal state depends on the interaction time t, as shown in Fig. 6. In particular, for the two first times $t = 0$ and 0.01 the state is closed to be maximally entangled. At $t = 0$ any state encodes no information on the parameter. At $t = 0.01$ this is due to a numerical error, as the score $\widetilde{S}$ is still found to be maximal by the algorithm (see the orange line Fig. 4). Then we find that the entanglement in $|\Psi\rangle$, as captured by the value p_0 changes smoothly with t. Asymptotically $t \to \infty$, we know that all initial states do equally well as the dynamics maps prepares as the steady state, which is product with the auxiliary system's state independent of the parameter. Moreover, from Fig. 5 it is clear that the presence of entanglement in the initial state does not give any substantial improvement for t close to one. We have also considered the optimal measurements $\{M_1, \ldots, M_4\}$ in Eq. (13) found by the algorithm. We see that the POVM elements are given by rank-1 projectors. The states corresponding to the projectors are also Schmidt diagonal in the computational basis (for the probe), and the amount of entanglement first increases and then decreases for $t \geqslant 0.06$ (similar to Fig. 6).

Lastly, as we have readily pointed out our methods can be simply adopted to other reward functions. As an example, in Appendix D we address the thermometry problem with the mean-square logarithmic error as the reward function, which has gained attention in recent years due to respecting scale-invariance properties [64–66].

C. Example 3. Multiparameter estimation: SU(2) gates

For our final example, we will consider a more complex metrology problem which involves multiple parameters. This is the problem of estimating any qubit unitary: the group SU(2).

As a first observation, note that any qubit unitary operator can be parametrized in terms of three independent parameters $\theta := (\theta^x, \theta^y, \theta^z)$, with $0 \leqslant \theta^i < 2\pi$ for all $i \in \{x, y, z\}$, as

$$U_\theta = e^{-i(\theta^x\sigma_x + \theta^y\sigma_y + \theta^z\sigma_z)}. \qquad (42)$$

Here, σ_i for $i \in \{x, y, z\}$ are the three Pauli operators. Since these generators do not commute, the estimation of the unitary U_θ (or equivalently the parameter vector θ) is a multiparameter estimation problem. The unitary channel that acts on the probe system and encodes the parameter θ is then given simply by $\mathcal{E}_\theta[\cdot] = U_\theta(\cdot)U_\theta^\dagger$, and will have a Choi operator C_θ associated to it.

We take a natural reward function that captures how close the estimated unitary is from the actual one, which is the fidelity, i.e.,

$$r(\theta, \hat{\theta}_i) = \frac{1}{d^2}\text{tr}(C_\theta\, C_{\hat{\theta}_i}), \qquad (43)$$

where in this example $d = 2$. Here $C_{\hat{\theta}_i}$, and for later reference C_{θ_k}, are defined analogously to C_θ, for a vector of estimator values $\hat{\theta}_i = (\hat{\theta}_a^x, \hat{\theta}_b^y, \hat{\theta}_c^z)$ and for a vector of discretization values $\theta_k = (\theta_a^x, \theta_b^y, \theta_c^z)$.

Here, we again analyze the cases of two different prior distributions, a uniform prior, as in Eq. (32), and a Gaussian prior, as in Eq. (33).

The parameter vector $\theta = (\theta^x, \theta^y, \theta^z)$ is discretized into values $\theta_k = (\theta_a^x, \theta_b^y, \theta_c^z)$, where each of the three elements follow the discretization in Eq. (34), with $\theta_{\min} = -\pi$ and $\theta_{\max} = \pi$, and all with the same value of $N_H = n_H$. Notice that this will amount to a final number of different discretization values of $N_H = n_H^3$.

The initial set of estimators $\hat{\theta}_i = (\hat{\theta}_a^x, \hat{\theta}_b^y, \hat{\theta}_c^z)$ is also set according to Eq. (35) for each parameter estimator, using the same values of $\theta_{\min} = -\pi$ and $\theta_{\max} = \pi$, and all with the same value of $N_O = n_O$. Here again this amounts to a total number of outcomes equal to $N_O = n_O^3$, analogously to the indexation of the parameter discretization.

We now discuss the application of our methods to this specific problem, plotting our results in Fig. 7.

Method 1. To apply method 1, we again set $n_H = n_O = n$ for each of the three parameters, amounting to a total number of discretization values and of outcomes equal to $N = n^3$. For the results presented here, we look at values of $n = N^{1/3} \in \{2, \ldots, 10\}$.

Method 2. Here we fix the number of outcomes $n_O = N_O^{1/3} \in \{2, \ldots, 10\}$ and set the number of hypotheses to be $N_H = n_H^3 = 10^3$. Notice that while the total number of hypotheses is 1000, each of the three parameters is discretized in only 10 different values. These values for the estimator are not expected to be optimal but are nevertheless used in method 2, serving as a starting point for the estimator optimization in method 3.

Method 3. To apply the seesaw in method 3, we again begin with the solution provided by method 2 for each N_O as a

(a)

(b)
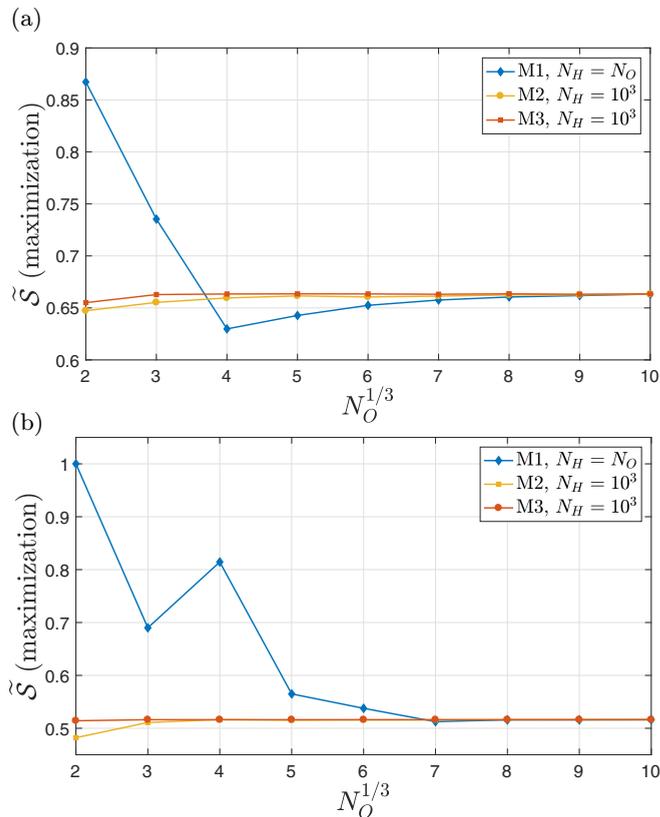
FIG. 7. *SU(2) estimation (Example 3)*. The maximum approximate score $\widetilde{S}$ for an SU(2) multiparameter estimation problem. Both panels show the scores corresponding to methods M1, M2, and M3 as a function of the cubic root of the total number of outcomes $N_O^{1/3} \in \{2, \ldots, 10\}$, for different prior distributions of the phase parameters $(\theta^x, \theta^y, \theta^z)$. (a) Shows the case of a uniform prior while (b) corresponds to a Gaussian prior. Each of the three parameters range from $\theta_{\min} = -\pi$ to $\theta_{\max} = \pi$. The considered cost function is the fidelity in Eq. (43).

starting point. In this case, we optimize over the estimators using standard gradient descent techniques. Therefore, the first step in the seesaw is the SDP in method 2, while the second step is a heuristic search over the estimators for a fixed tester.

Results. In Fig. 7, we plot the maximal approximate scores $\widetilde{S}$ obtained via the three methods outlined above. Figure 7(a) concerns the case of a uniform prior distribution and Fig. 7(b) that of a Gaussian prior distribution. In both cases we observe method 3 converging to its final value of $\widetilde{S}$ with $N_O^{1/3} = 3$, i.e., total number of outcomes $N_O = 27$. This is consistent with the fact that, in this case, extremal testers have at most $(d_I \times d_O)^2 = 16$ outcomes, and hence, for $N_O^{1/3} = 2$ (total number of outcomes $N_O = 8$) we are not yet optimizing over all possible extremal testers. This is an interesting case where extremal nonprojective POVMs with $(d_I \times d_O)^2$ outcomes show improvement over $(d_I \times d_O)$-outcome PVMs. Method 2 quickly approaches this same value, while method 1 requires higher values of $N_O^{1/3}$. Nevertheless, as expected, for a larger number of outcomes, namely, $N_O^{1/3} = 10$, all methods yield the same result.

VI. CONCLUSION AND OUTLOOK

We introduced a set of tools for addressing Bayesian parameter estimation problems, applying techniques from the formalism of higher-order operations and drawing inspiration from the problem of channel discrimination. The key insight that we exploit consists of describing the quantum strategy, i.e., the state preparation and the measurement (see Fig. 1), as a single operation called a quantum tester. The latter is characterized by SDP constraints, and can thus be optimized efficiently. We developed three methods for determining the state of the probe, the measurement, and the estimators in any parameter estimation problem, regardless of prior distribution, reward function, or description of the quantum evolution.

The first method exploits the connection between the Bayesian approach to parameter estimation and quantum channel discrimination. By discretizing the parameter to a finite set of values, and by furthermore associating each value of the estimator to a value of the discretized parameter, one can directly map a parameter estimation task onto a channel discrimination one, albeit with a reward function which inherits the geometry of the original parameter set. We leveraged this connection to create a general method for approximating the optimal solution of the estimation problem within any arbitrary precision. We also proved that our approximation converges to the optimal score. Although this method is conceptually simple and comes with a convergence guarantee, it may nonetheless be computationally demanding since the size of the optimization variables increases with the fineness of parameter discretization. Our second method computes an approximation of the optimal quantum strategy for a fixed set of estimators. This method is less computationally demanding in general, but it relies on a good guess for optimal estimators, which is not always available. To address this drawback, our third method iteratively combines an optimization over the quantum strategy and over the estimators, and hence does not require any previous knowledge over the estimator.

A key advantage of our methods is their universal applicability. They can be used for any parameter estimation problem, regardless of the nature of parameter encoding or the number of parameters to be estimated. To showcase this wide-ranging applicability, we examined three distinct case studies of high practical importance: local phase estimation, thermometry, and SU(2) estimation.

We also developed tools to bound the performance of estimation strategies that do not require entanglement and used them to show that, in the thermometry problem, probe states that are entangled with an auxiliary system lead to a more precise estimation of the temperature parameter, particularly at finite times.

Our work provides a starting platform for the application of higher-order operations to the problem of Bayesian parameter estimation. We conclude by summarizing further research directions that could draw further benefits from this approach.

Generalization to many shots. The quantum strategies explored here concern a situation in which, at each independent realization of the experiment, one is given access to a single call, or copy, of the channel that encodes the parameter θ. It is also the case that, in more general scenarios, where one has access to multiple calls at once, the optimization of the quantum

strategy can be done with SDP as well. Multiple-copy testers can take different forms, describing different classes of quantum strategies, such as parallel (nonadaptive) and sequential (adaptive), or even those involving an indefinite causal order. Such testers have been defined in Refs. [39,40] and explored in a frequentist approach to metrology [42]. These techniques can also be applied to multiple-copy Bayesian estimation protocols, and be exploited to investigate, for example, whether different classes of estimation strategies can lead to higher precision in the parameter estimation. Similarly, strategies with an indefinite causal order could lead to establishing new types of metrological resources.

While in the multishot scenario one can find the global optimal protocol as explained above, it can be hard to implement in practice, due to the exponential growth of of Hilbert space dimension. Alternatively, one can seek greedy optimal algorithms [67]. In such a strategy, one would (i) perform the optimization protocol as subscribed in our work, (ii) update the prior distribution to the posterior distribution based on the outcome, and repeat steps (i) and (ii) until all shots are consumed. Despite not being necessarily globally optimal, this strategy can be very strong and has shown to be asymptotically optimal in some cases [68]. An example of our interest is in Bayesian equilibrium thermometry [66]. Regardless of global optimality, it is practically easier to implement such multishot strategies since the required operations do not involve exponentially increasing Hilbert spaces.

Applications to complex noise models. One of the key advantages of our approach is its versatility in handling various types of parameter-encoding dynamics. Often, sensing methods are limited to specific parameter-encoding channels, however, our approach can effectively model and accommodate any type of dynamics and address different types of noise. Specifically, if one has a good description of the noise appearing in the measurement process, one can simply incorporate this noise into the encoding channel and compute optimal testers and optimal estimators according to any one of the three methods. Therefore, applying these techniques to real (noisy) experimental settings to infer their actual performance would be an interesting next step.

Quantum metrology techniques for asymptotic quantum channel discrimination. In this study, we utilized higher-order operations, a technique previously used to study channel discrimination, to provide a nearly optimal solution for the quantum parameter estimation problem. It would be interesting to explore whether the reverse approach could also yield novel insights into the field of channel discrimination. Specifically, one could investigate whether leveraging asymptotic theoretical results from quantum metrology, such as the Heisenberg scaling, can contribute to the investigation of asymptotic quantum channel discrimination. This direction holds promise for gaining a deeper understanding of the relationship between channel discrimination and quantum metrology.

Connections with the multihypothesis testing problem. The discretization of the parameter space that we perform suggests that the Bayesian estimation problem can be connected with a multihypothesis testing problem [69]. However, it should be noted that this connection is only partial. Indeed, our work exploits the fact that Bayesian estimation can be seen as a multihypothesis testing problem with (i) a continuous set of hypotheses and with (ii) a specific geometry on the "hypothesis space" as captured by the cost function. Still, we believe that the methods developed here (method 1 in particular) could be potentially useful for determining bounds on the errors for the multihypothesis testing problem. Another interesting question is whether some of the bounds on error probabilities arising in the multihypothesis testing scenario could be also applicable in the Bayesian setting.

All code developed for this work is freely available in our online repository [70].

APPENDIX A: APPROXIMATIONS FOR STRATEGIES WITHOUT ENTANGLEMENT

If the particular problem at hand does not allow for the use of auxiliary systems that may be entangled with the target system, the estimation problem in Eq. (1) above reduces to the following:

$$p(i|\theta) = \text{tr}(\mathcal{E}_\theta[\rho] M_i), \tag{A1}$$

where now $\rho \in \mathcal{L}(\mathcal{H}^I)$ and $M_i \in \mathcal{L}(\mathcal{H}^O)$ or, in the Choi representation

$$p(i|\theta) = \text{tr}(C_\theta (\rho^T \otimes M_i)). \tag{A2}$$

In this case, the optimization over the quantum strategy (state and measurement) of any linear function of $p(i|\theta)$ is no longer an SDP. The tester in this case, which does not require an auxiliary system for implementation, is given by

$$T_i = \rho^T \otimes M_i, \tag{A3}$$

where $\rho \in \mathcal{H}^I$ and $M_i \in \mathcal{H}^O$. Unfortunately, this kind of tester does not admit a simple mathematical characterization. In fact, the problem of characterizing this kind of tester is mathematically very similar to the problem of characterizing separable states, except in this case we are interested in the "separability" between the operators acting on the input and output spaces of the tester.

For this reason, techniques used to approximate the set of separable states can be applied here to approximate the set of testers that do not require an auxiliary system, or entanglement, for its implementation. Outer approximations of the set of testers that do not require entanglement are useful to determine whether or not entanglement or memory is advantageous in an estimation task; if the score achieved by a general tester is better than the best score achieved by a tester in the outer approximation, then entanglement is useful. On the other hand, inner approximations are useful to determine

a lower bound on how well a tester without entanglement can perform in a given task.

An example of outer approximation is the positive partial transpose (PPT) condition [71]. The set of all testers whose elements have a positive partial transpose, i.e., the set of all $T = \{T_i\}$ that satisfy

$$T_i \geqslant 0, \quad \forall\, i \tag{A4}$$

$$\sum_i T_i = \sigma \otimes \mathbb{1}^O, \tag{A5}$$

$$(T_i)^{\mathrm{T}_I} \geqslant 0, \quad \forall\, i \tag{A6}$$

is an outer approximation of the set of testers that satisfy Eq. (A3). Similarly, the imposition of the condition of k-symmetric extensibility [72] on the tester elements defines an outer approximation for the set of testers that require entanglement. Any other SDP method of approximating the set of separable states can also be applied to approximate the set of testers without entanglement, although the resulting approximation may not be tight or may not converge to the actual set.

Inner approximations can, for example, be achieved by applying a seesaw method of optimizing over states and measurements separately over many iterations. This method is not guaranteed to converge, but will yield a bound.

APPENDIX B: CONVERGENCE OF THE APPROXIMATIONS

The cornerstone of our results is the discretization of the estimators and the hypotheses. The intuition suggests that as the discretization is made finer, the approximation becomes more precise and converges to the exact value. Here, we make this statement more rigorous. We focus on a score function that needs to be maximized; for those that require minimization a similar argument holds.

First, let us denote the optimal protocol by $\{\{T_i^*\}, \{\hat{\theta}_i^*\}\}$, it maximizes the score in Eq. (17) to its optimal value $\mathcal{S}^*$. We know that the optimal protocol has at most $D := (d_I \times d_O)^2$ elements, therefore,

$$\{\{T_i^*\}, \{\hat{\theta}_i^*\}\} := \arg\max_{\{T_i\},\{\hat{\theta}_i\}} \sum_{i=1}^{D} \mathrm{tr}(X(\hat{\theta}_i)\, T_i), \tag{B1}$$

$$\mathcal{S}^* := \max_{\{T_i\},\{\hat{\theta}_i\}} \sum_{i=1}^{D} \mathrm{tr}(X(\hat{\theta}_i)\, T_i) = \sum_{i=1}^{D} \mathrm{tr}(X(\hat{\theta}_i^*)\, T_i^*). \tag{B2}$$

Now consider another protocol in which the estimators are fixed to $\{\hat{\theta}_i\}_{i=1}^{N_O}$, and the tester $\{\tilde{T}_i\}_{i=1}^{N_O}$ is the solution of the SDP maximizing the score for the given estimators

$$\{\tilde{T}_i\} = \arg\max_{\{T_i\}} \sum_{i=1}^{N_O} \mathrm{tr}(X(\hat{\theta}_i)\, T_i). \tag{B3}$$

The protocol achieves a certain score

$$\widetilde{\mathcal{S}}^* = \max_{\{T_i\}} \sum_{i=1}^{N_O} \mathrm{tr}(X(\hat{\theta}_i)\, T_i) = \sum_{i=1}^{N_O} \mathrm{tr}(X(\hat{\theta}_i)\, \tilde{T}_i) \leqslant \mathcal{S}^*. \tag{B4}$$

Next, let us consider the two sets of estimators $\{\hat{\theta}_i^*\}_{i=1}^{D}$ and $\{\hat{\theta}_i\}_{i=1}^{N_O}$ and for each $\hat{\theta}_i^*$ define $\bar{\theta}_i^*$ to be the closest value in the second set, i.e.,

$$\bar{\theta}_i^* := \arg\min_{\hat{\theta}_k | 1 \leqslant k \leqslant N_O} |\hat{\theta}_k - \hat{\theta}_i^*|, \tag{B5}$$

$$\epsilon_i := \hat{\theta}_i^* - \bar{\theta}_i^*, \tag{B6}$$

where ϵ_i quantifies how different these values are. For simplicity we also introduce $\epsilon := \max_i |\epsilon_i|$. Note that for concreteness we here used the absolute value of the difference $|\hat{\theta}_i - \hat{\theta}_k^*|$ as a distance between the estimated values, however, any other distance $d(\hat{\theta}_i, \hat{\theta}_k^*)$ could be used here and below to define $\bar{\theta}_k^*$ and ϵ instead (e.g., in the multiparameter case). The new values $\{\bar{\theta}_i^*\}$ allow us to define the protocol $\{\{T_i^*\}, \{\bar{\theta}_i^*\}\}$, where the testers are taken from the optimal protocol but the estimators have been modified. It achieves a certain score

$$\sum_{i=1}^{D} \mathrm{tr}[X(\bar{\theta}_i^*)\, T_i^*] \leqslant \max_{\{T_i\}_{i=1}^{D}} \sum_{i=1}^{D} \mathrm{tr}[X(\bar{\theta}_i^*)\, T_i]$$

$$\leqslant \max_{\{T_i\}_{i=1}^{N_O}} \sum_{i=1}^{N_O} \mathrm{tr}[X(\hat{\theta}_i)\, T_i] = \widetilde{\mathcal{S}}^*. \tag{B7}$$

Here, we used the fact that by construction $\{\bar{\theta}_k^*\}_{i=1}^{D}$ form a subset of $\{\hat{\theta}_i\}_{i=1}^{N_O}$, therefore, the maximization over the tester $\{T_i\}_{i=1}^{N_O}$ includes the maximization over the tester $\{T_i\}_{i=1}^{D}$ (some of the tester elements can be identically zero).

Our next goal is to bound the deviation between $\sum_{i=1}^{D} \mathrm{tr}[X(\bar{\theta}_i^*)\, T_i^*]$ and the optimal score $\mathcal{S}^*$, which are obtained with the same tester. To do so we recall their Bayesian interpretation in terms of posterior parameter distribution in Eq. (4):

$$\sum_{i=1}^{D} \mathrm{tr}[X(\hat{\theta}_i^*)\, T_i^*] = \sum_{i=1}^{D} p(i) \int \mathrm{d}\theta\, p(\theta|i)\, r(\theta, \hat{\theta}_i^*)$$

$$= \sum_{i=1}^{D} p(i)\, \mathbb{E}^{(i)}[r(\theta, \hat{\theta}_i^*)], \tag{B8}$$

where each expected value $\mathbb{E}^{(i)}[\cdot]$ is taken with respect to the probability distribution $p(\theta|i)$. This allows us to write

$$\mathcal{S}^* - \widetilde{\mathcal{S}}^* \leqslant \mathcal{S}^* - \sum_{i=1}^{D} \mathrm{tr}[X(\bar{\theta}_i^*)\, T_i^*]$$

$$= \sum_{i=1}^{D} p(i) \{\mathbb{E}^{(i)}[r(\theta, \hat{\theta}_i^*)] - \mathbb{E}^{(i)}[r(\theta, \bar{\theta}_i^*)]\}$$

$$= \sum_{i=1}^{D} p(i)\, \mathbb{E}^{(i)}[r(\theta, \hat{\theta}_i^*) - r(\theta, \bar{\theta}_i^*)]. \tag{B9}$$

Here, it is intuitively clear that for nearby values $\hat{\theta}_i^*$ and $\bar{\theta}_i^*$ the expected values $\mathbb{E}^{(i)}[r(\theta, \hat{\theta}_i^*)]$ and $\mathbb{E}^{(i)}[r(\theta, \bar{\theta}_i^*)]$ will also be close, provided that the reward function r is regular enough. For simplicity let us now assume that it is Lipschitz continuous, i.e.,

$$|r(\theta, \hat{\theta}_i^*) - r(\theta, \bar{\theta}_i^*)| \leqslant K_r\, \epsilon \quad \text{if} \quad |\hat{\theta}_i^* - \bar{\theta}_i^*| \leqslant \epsilon, \tag{B10}$$

for any small enough $\epsilon \leqslant \delta$ (here K_r might depend on delta) which directly implies $\widetilde{\mathcal{S}}^* \geqslant \mathcal{S}^* - K_r\, \epsilon$. Finally, for a scalar

parameter the N_O estimators $\{\hat{\theta}_i\}$ can be chosen such that $\epsilon \leqslant \frac{L}{N_O}$, where L is some constant depending on the prior. This would guarantee the convergence to the optimal score with

$$0 \leqslant \mathcal{S}^* - \widetilde{\mathcal{S}}^* \leqslant \frac{K_r L}{N_O}. \tag{B11}$$

1. Case of reward functions that are not Lipschitz continuous

Notably, Lipschitz continuity of the reward function is not necessary to guarantee the convergence of the score $\widetilde{\mathcal{S}}^* \to \mathcal{S}^*$. However, in such cases it seems difficult to make a general statement, which might furthermore require to assume some regularity of the prior. Nevertheless, for illustration let us consider a piecewise constant reward function

$$r(\theta, \hat{\theta}) = \begin{cases} 1, & |\theta - \hat{\theta}| \leqslant \Delta \\ 0, & \text{otherwise} \end{cases} \tag{B12}$$

that can be used to define a confidence interval for the parameter. This reward function coincides with the recent proposal in Ref. [38]. This function is manifestly discontinuous, with $r(\theta, \hat{\theta}_i^*) - r(\theta, \bar{\theta}_i^*)$ taking the value $+1$ on an interval $\theta \in I_+^i$ of width $|\hat{\theta}_i^* - \bar{\theta}_i^*| \leqslant \epsilon$, the value -1 on another interval of the same width, and is otherwise zero. From Eq. (B9) we then find

$$\mathcal{S}^* - \widetilde{\mathcal{S}}^* \leqslant \sum_{i=1}^D p(i) \mathbb{E}^{(i)}[r(\theta, \hat{\theta}_i^*) - r(\theta, \bar{\theta}_i^*)]$$

$$\leqslant \sum_{i=1}^D p(i) \mathrm{Pr}^{(i)}[\theta \in I_+^i], \tag{B13}$$

where each probability $\mathrm{Pr}^{(i)}[f(\theta)] = \int d\theta\, f(\theta) p(\theta|i)$ is taken over the conditional distribution $p(\theta|i)$. Defining the union of all the intervals $I_+ = \cup_{i=1}^D I_+^i$ we can further upper bound the score difference with

$$\mathcal{S}^* - \mathcal{S} \leqslant \sum_{i=1}^D p(i) \mathrm{Pr}^{(i)}[\theta \in I_+^i] \leqslant \sum_{i=1}^D p(i) \mathrm{Pr}^{(i)}[\theta \in I_+]$$

$$= \mathrm{Pr}[\theta \in I_+], \tag{B14}$$

where in the last term the probability is taken over the prior distribution $p(\theta) = \sum_i p(i) p(\theta|i)$, and we used

$$\sum_{i=1}^D p(i) \mathrm{Pr}^{(i)}[\theta \in I_+] = \int d\theta \sum_{i=1}^D p(i) p(\theta|i) \chi_{I_+}(\theta)$$

$$= \int d\theta\, p(\theta) \chi_{I_+}(\theta) \tag{B15}$$

for the indicator function $\chi_{I_+} = \{ \begin{smallmatrix} 1, & \theta \in I_+ \\ 0, & \theta \notin I_+ \end{smallmatrix}$. Here, the set I_+ is of measure at most $D\epsilon$. Thus, for any regular prior $p(\theta)$ supported $[p(\theta) \neq 0]$ on a set of finite measure, the probability $\mathrm{Pr}[\theta \in I_+]$ to find a value inside I_+ converges to zero with ϵ. In particular, for any prior with a bounded density $p(\theta) \leqslant p_*$
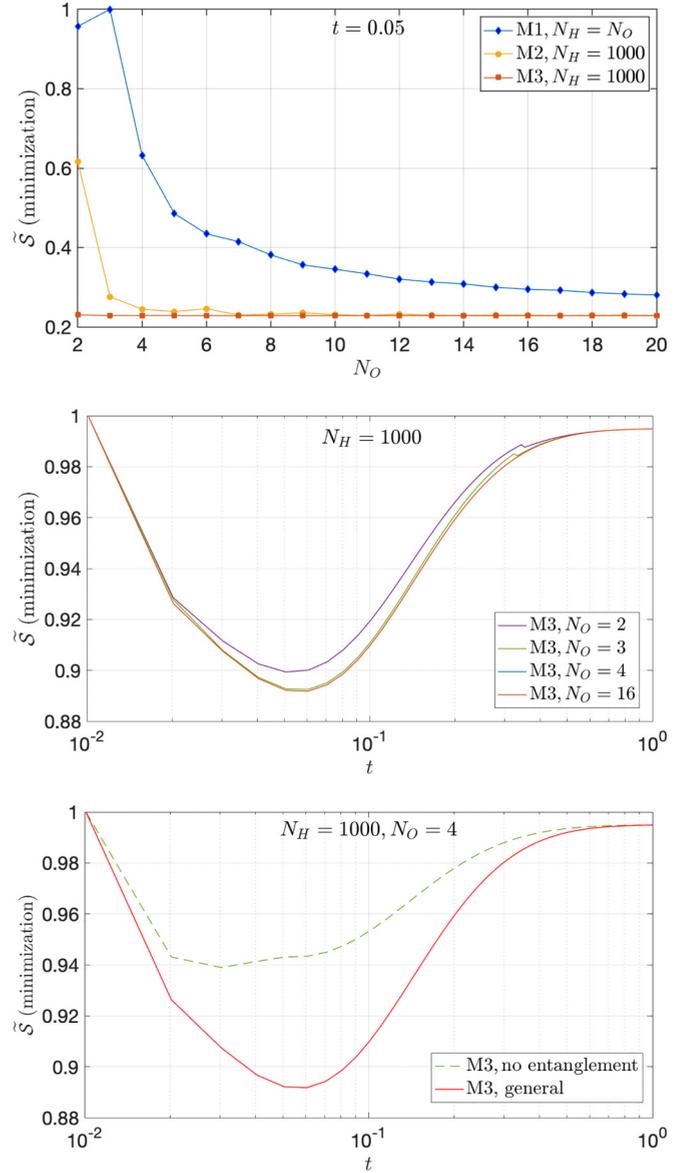
FIG. 8. The thermometry problem seen from the perspective of the EMSLE as the cost function. The top, middle, and low panels correspond to Figs. 3–5 of the main text, respectively; note the logarithmic scaling in the middle and bottom figures. All other parameters are kept the same as the corresponding graphs in the main text.

we find

$$\mathcal{S}^* - \mathcal{S} \leqslant \mathrm{Pr}[\theta \in I_+] = \int d\theta\, p(\theta) \chi_{I_+}(\theta)$$

$$\leqslant \int d\theta\, p_* \chi_{I_+}(\theta) \leqslant D\epsilon\, p_*. \tag{B16}$$

2. Convergence for the MSE

As a first example, let us have a look into the MSE that we used for the thermometry problem. In this case the score function is a cost which has to be minimized, so to match to the notation with the previous section we consider maximization

of $r(\theta, \hat{\theta}) = -(\theta - \hat{\theta})^2$. We have

$$r(\theta, \hat{\theta}_i^*) - r(\theta, \bar{\theta}_i^*) = -(\theta - \hat{\theta}_i^*)^2 + -(\theta - \bar{\theta}_i^*)^2 \quad \text{(B17)}$$

$$= -(\theta - \hat{\theta}_i^*)^2 + (\theta - \hat{\theta}_i^* + \epsilon_i)^2 \quad \text{(B18)}$$

$$= \epsilon_i^2 + 2\epsilon_i(\theta - \hat{\theta}_i^*), \quad \text{(B19)}$$

where ϵ_i is defined in Eq. (B6). Plugging this in Eq. (B9) one gets

$$\mathcal{S}^* - \widetilde{\mathcal{S}}^* \leqslant \sum_{i=1}^{D} p(i)\mathbb{E}^{(i)}[r(\theta, \hat{\theta}_i^*) - r(\theta, \bar{\theta}_i^*)]$$

$$= \sum_{i=1}^{D} p(i)\left(\epsilon_i^2 + 2\epsilon_i\mathbb{E}^{(i)}[(\theta - \hat{\theta}_i^*)]\right). \quad \text{(B20)}$$

But for the MSE we know that the optimal estimator is the mean, i.e., $\hat{\theta}_i^* = \mathbb{E}^{(i)}[\theta] = \int d\theta\, \theta\, p(\theta|i)$. Therefore, the second term is zero and we find

$$\mathcal{S}^* - \widetilde{\mathcal{S}}^* \leqslant \sum_{i=1}^{D} p(i)\epsilon_i^2 \leqslant \epsilon^2 \quad \text{(B21)}$$

with $\epsilon = \max_i |\epsilon_i|$.

3. Convergence for the cos² reward function

In the phase estimation problem we considered a reward function that reads as $r(\theta, \hat{\theta}_i^*) = \cos^2(\frac{\theta - \hat{\theta}_i^*}{2})$. First of all, note that in this case the optimal estimator can be found in a closed form. To do so, we first rewrite the score for the posterior distribution $p(\theta|i)$ as

$$\mathbb{E}^{(i)}\left[\cos^2\left(\frac{\theta - \hat{\theta}_i^*}{2}\right)\right] = \mathbb{E}^{(i)}\left[\frac{\cos(\theta - \hat{\theta}_i^*) + 1}{2}\right]$$

$$= \frac{1 + \mathbb{E}^{(i)}[\cos(\theta)]\cos(\hat{\theta}_i^*) + \mathbb{E}^{(i)}[\sin(\theta)]\sin(\hat{\theta}_i^*)}{2}.$$

$$\text{(B22)}$$

Impose that the derivative of the score with respect to the estimator $\hat{\theta}_i^*$ is zero:

$$\frac{\partial}{\partial\hat{\theta}_i^*}\mathbb{E}^{(i)}\left[\cos^2\left(\frac{\theta - \hat{\theta}_i^*}{2}\right)\right] = -\frac{1}{2}\mathbb{E}^{(i)}[\sin(\theta - \hat{\theta}_i^*)]$$

$$= \frac{1}{2}\{-\mathbb{E}^{(i)}[\cos(\theta)]\sin(\hat{\theta}_i^*) + \mathbb{E}^{(i)}[\sin(\theta)]\cos(\hat{\theta}_i^*)\} = 0,$$

$$\text{(B23)}$$

which is equivalent to

$$\tan(\hat{\theta}_i^*) = \frac{\mathbb{E}^{(i)}[\sin(\theta)]}{\mathbb{E}^{(i)}[\cos(\theta)]} \quad \text{(B24)}$$

and admits two solutions

$$\hat{\theta}_i^* = \arctan\left(\frac{\mathbb{E}^{(i)}[\sin(\theta)]}{\mathbb{E}^{(i)}[\cos(\theta)]}\right) \quad \text{or}$$

$$\hat{\theta}_i^* = \arctan\left(\frac{\mathbb{E}^{(i)}[\sin(\theta)]}{\mathbb{E}^{(i)}[\cos(\theta)]}\right) + \pi, \quad \text{(B25)}$$

with the notation from the main text. We then need to pick the value which gives the highest contribution to the reward in Eq. (B22). In fact, up to a constant the reward is the scalar product between the vectors $(\mathbb{E}^{(i)}[\cos(\theta)], \mathbb{E}^{(i)}[\sin(\theta)])$ and $(\cos(\hat{\theta}_i^*), \sin(\hat{\theta}_i^*))$, so its maximum is attained when the two vectors are in the same half of the disk. Since the range of $\arctan \in [-\frac{\pi}{2}, \frac{\pi}{2}]$ corresponds to positive cosine, the choice of the optimal estimator solution depends on the sign of $\langle\cos(\theta)\rangle^{(i)}$:

$$\hat{\theta}_i^* = \begin{cases} \arctan\left(\frac{\mathbb{E}^{(i)}[\sin(\theta)]}{\mathbb{E}^{(i)}[\cos(\theta)]}\right), & \mathbb{E}^{(i)}[\cos(\theta)] \geqslant 0 \\ \arctan\left(\frac{\mathbb{E}^{(i)}[\sin(\theta)]}{\mathbb{E}^{(i)}[\cos(\theta)]}\right) + \pi, & \text{otherwise.} \end{cases} \quad \text{(B26)}$$

Therefore, we can see that

$$\mathcal{S}^* - \widetilde{\mathcal{S}}^* \leqslant \sum_{i=1}^{D} p(i)\mathbb{E}^{(i)}[r(\theta, \hat{\theta}_i^*) - r(\theta, \bar{\theta}_i^*)]$$

$$= \sum_{i=1}^{D} p(i)\mathbb{E}^{(i)}\left[\cos^2\left(\frac{\theta - \hat{\theta}_i^*}{2}\right) - \cos^2\left(\frac{\theta - \hat{\theta}_i^*}{2}\right)\cos^2\left(\frac{\epsilon_i}{2}\right) + \frac{1}{2}\sin(\theta - \hat{\theta}_i^*)\sin(\epsilon_i) - \sin^2\left(\frac{\theta - \hat{\theta}_i^*}{2}\right)\sin^2\left(\frac{\epsilon_i}{2}\right)\right]$$

$$= \sum_{i=1}^{D} p(i)\mathbb{E}^{(i)}\left[\cos(\theta - \hat{\theta}_i^*)\sin^2\left(\frac{\epsilon_i}{2}\right) + \frac{1}{2}\sin(\theta - \hat{\theta}_i^*)\sin(\epsilon_i)\right]$$

$$= \sum_{i=1}^{D} p(i)\mathbb{E}^{(i)}\left[\cos(\theta - \hat{\theta}_i^*)\sin^2\left(\frac{\epsilon_i}{2}\right)\right] + \frac{1}{2}\sum_{i=1}^{D} p(i)\mathbb{E}^{(i)}[\sin(\theta - \hat{\theta}_i^*)\sin(\epsilon_i)]$$

$$= \sum_{i=1}^{D} p(i)\mathbb{E}^{(i)}\left[\cos(\theta - \hat{\theta}_i^*)\sin^2\left(\frac{\epsilon_i}{2}\right)\right] \leqslant \frac{1}{4}\sum_{i=1}^{D} p(i)\epsilon_i^2\mathbb{E}^{(i)}[\cos[2(\theta - \hat{\theta}_i^*)]] \leqslant \frac{1}{4}\sum_{i=1}^{D} p(i)\epsilon_i^2 \leqslant \frac{\epsilon^2}{4}, \quad \text{(B27)}$$

where in the penultimate line we use the optimality criterion (B23), and used the definition of ϵ_i in Eq. (B6), and $\epsilon = \max_i |\epsilon_i|$.

APPENDIX C: EXTENSION TO MULTIPARAMETER ESTIMATION

Take the two-parameter estimation example, where the parameters θ^1 and θ^2 are encoded via a channel C_{θ^1, θ^2}. The estimates of these parameters are given, respectively, by $\{\hat{\theta}^1_{i_1}\}_{i_1=1}^{N_{O_1}}$ and $\{\hat{\theta}^2_{i_2}\}_{i_2=1}^{N_{O_2}}$. The continuous version of the problem can be expressed as

$$
S = \sum_{i_1=1}^{N_{O_1}} \sum_{i_2=1}^{N_{O_2}} \int d\theta^1 \int d\theta^2 \, p(\theta^1, \theta^2) \, r(\theta^1, \theta^2, \hat{\theta}^1_{i_1}, \hat{\theta}^2_{i_1})
$$
$$
\times \mathrm{tr}\big(C_{\theta^1, \theta^2} \, T_{i_1, i_2}\big). \tag{C1}
$$

We then choose a discretization of θ^1 given by $\{\theta^1_{k_1}\}_{k_1=1}^{N_{H_1}}$ and of θ^2 given by $\{\theta^2_{k_2}\}_{k_2=1}^{N_{H_2}}$, leading to the approximation

$$
\widetilde{S} = \sum_{i_1=1}^{N_{O_1}} \sum_{i_2=1}^{N_{O_2}} \sum_{k_1=1}^{N_{H_1}} \sum_{k_2=1}^{N_{H_2}} p\big(\theta^1_{k_1}, \theta^2_{k_2}\big) \, r\big(\theta^1_{k_1}, \theta^2_{k_2}, \hat{\theta}^1_{i_1}, \hat{\theta}^2_{i_2}\big)
$$
$$
\times \mathrm{tr}(C_{\theta^1_{k_1}, \theta^2_{k_2}} T_{i_1, i_2}). \tag{C2}
$$

By mapping the indices $(i_1, i_2) \mapsto i$, where $i \in \{1, \ldots, N_O\}$, $N_O = N_{O_1} N_{O_2}$, and $(k_1, k_2) \mapsto k$, where $k \in \{1, \ldots, N_H\}$, $N_H = N_{H_1} N_{H_2}$, and furthermore defining $\vec{\theta}_k := (\theta^1_{k_1}, \theta^2_{k_2})$ and $\vec{\hat{\theta}}_i := (\hat{\theta}^1_{i_1}, \hat{\theta}^2_{i_2})$, S can be rewritten as

$$
\widetilde{S} = \sum_{i=1}^{N_O} \sum_{k=1}^{N_H} p(\vec{\theta}_k) \, r(\vec{\theta}_k, \vec{\hat{\theta}}_i) \, \mathrm{tr}(C_{\vec{\theta}_k} \, T_i), \tag{C3}
$$

which is equivalent to the single-parameter problem. Hence, the techniques presented here are applicable to multiparameter estimation problems as well.

APPENDIX D: DETAILS OF THERMOMETRY (EXAMPLE 2)

Here, we provide some technical details for the Example 2 where we want to estimate the temperature of a bosonic bath. A qubit that is initially prepared in the state $\rho^p(0) = \begin{pmatrix} \rho_{11} & \rho_{12} \\ \rho_{21} & 1 - \rho_{11} \end{pmatrix}$. Under the evolution (38) the probe evolves to

$$
\mathcal{E}_\theta(t)[\rho^p(0)] = \rho^p_\theta(t) = \begin{pmatrix} \frac{e^{-\Gamma t}\left[\rho_{11}(2N_{B/F}+1)+(e^{\Gamma t}-1)(N_{B/F}+1)\right]}{2N_{B/F}+1} & \rho_{12}\,e^{-\frac{\Gamma t+2i\epsilon t}{2}} \\ \rho_{21}\,e^{-\frac{\Gamma t-2i\epsilon t}{2}} & \frac{e^{-\Gamma t}\left[-\rho_{11}(2N_{B/F}+1)+N_{B/F}(1+e^{\Gamma t})+1\right]}{2N_{B/F}+1} \end{pmatrix}, \tag{D1}
$$

where

$$
\Gamma = \Gamma_{\mathrm{in}} + \Gamma_{\mathrm{out}} = J(\epsilon)\,(2N_{B/F}+1).
$$

The Choi operator of the thermalization channel is then given by

$$
\sum_{ij} \mathcal{E}_\theta(t)[|i\rangle\langle j|] \otimes |i\rangle\langle j| = \begin{pmatrix} \frac{N_{B/F}+N_{B/F}\,e^{-\Gamma t}+1}{(2N_{B/F}+1)} & 0 & 0 & e^{-\frac{\Gamma t+2i\epsilon t}{2}} \\ 0 & \frac{e^{-\Gamma t}\left(e^{\Gamma t}-1\right)\left(N_{B/F}+1\right)}{(2N_{B/F}+1)} & 0 & 0 \\ 0 & 0 & \frac{N_{B/F}-N_{B/F}\,e^{-\Gamma t}}{(2N_{B/F}+1)} & 0 \\ e^{-\frac{\Gamma t-2i\epsilon t}{2}} & 0 & 0 & \frac{e^{-\Gamma t}\left(N_{B/F}+N_{B/F}\,e^{\Gamma t}+1\right)}{(2N_{B/F}+1)} \end{pmatrix}. \tag{D2}
$$

Remark. The only temperature dependence comes from $N_{B/F}$. In particular, the Hamiltonian term is independent of (θ) and thus can be ignored. Then the optimal solution for this problem should be rotated with the same Hamiltonian in order to compensate for it. As such, we can ignore the phases in the off-diagonal terms above.

Using the expected mean logarithmic error as a cost function

In the main text, we took the MSE as our figure of merit. However, in recent years, an alternative cost function has been put forward for thermometry, which is motivated by scale invariance [65]. This is the so-called expected mean square logarithmic error (EMSLE) at the kernel of which lies the following reward function:

$$
r(\theta, \hat{\theta}^*_i) = \log^2(\hat{\theta}_i/\theta), \tag{D3}
$$

which can be analytically solved to find the optimal estimator as [65]

$$
\hat{\theta}^*_i = \exp \int d\theta \, p(\theta|i) \ln(\theta). \tag{D4}
$$

Interestingly, for this cost function, one can also prove that the optimal POVM is in fact a PVM [64]. Our results straightforwardly apply to such figure of merit. We showcase this by reproducing our Figs. 3–5. These are depicted here in the three panels of Fig. 8, respectively, from top to bottom. The fact that PVMs are optimal is reflected in the middle panel, where our method M3 is optimal with only $N_O = 4$ outcomes.

[1] A. Acín, I. Bloch, H. Buhrman, T. Calarco, C. Eichler, J. Eisert, D. Esteve, N. Gisin, S. J. Glaser, F. Jelezko, S. Kuhr, M. Lewenstein, M. F. Riedel, P. O. Schmidt, R. Thew, A. Wallraff, I. Walmsley, and F. K. Wilhelm, The quantum technologies roadmap: A European community view, New J. Phys. **20**, 080201 (2018).

[2] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Observation of gravitational waves from a binary black hole merger, Phys. Rev. Lett. **116**, 061102 (2016).

[3] R. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Gwtc-2: Compact binary coalescences observed by ligo and virgo during the first half of the third observing run, Phys. Rev. X **11**, 021053 (2021).

[4] 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).

[5] 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).

[6] D. Budker and M. Romalis, Optical magnetometry, Nat. Phys. **3**, 227 (2007).

[7] W. Wasilewski, K. Jensen, H. Krauter, J. J. Renema, M. V. Balabas, and E. S. Polzik, Quantum noise limited and entanglement-assisted magnetometry, Phys. Rev. Lett. **104**, 133601 (2010).

[8] M. W. Mitchell, J. S. Lundeen, and A. M. Steinberg, Super-resolving phase measurements with a multiphoton entangled state, Nature (London) **429**, 161 (2004).

[9] V. Giovannetti, S. Lloyd, and L. Maccone, Advances in quantum metrology, Nat. Photonics **5**, 222 (2011).

[10] C. L. Degen, F. Reinhard, and P. Cappellaro, Quantum sensing, Rev. Mod. Phys. **89**, 035002 (2017).

[11] G. Tóth and I. Apellaniz, Quantum metrology from a quantum information science perspective, J. Phys. A: Math. Theor. **47**, 424006 (2014).

[12] M. G. A. Paris, Quantum estimation for quantum technology, Int. J. Quantum Inf. **7**, 125 (2009).

[13] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum metrology, Phys. Rev. Lett. **96**, 010401 (2006).

[14] A. Fujiwara and H. Imai, A fibre bundle over manifolds of quantum channels and its application to quantum statistics, J. Phys. A: Math. Theor. **41**, 255304 (2008).

[15] B. Escher, R. L. de Matos Filho, and L. Davidovich, General framework for estimating the ultimate precision limit in noisy quantum-enhanced metrology, Nat. Phys. **7**, 406 (2011).

[16] R. Demkowicz-Dobrzański, J. Kołodyński, and M. Guţă, The elusive heisenberg limit in quantum-enhanced metrology, Nat. Commun. **3**, 1063 (2012).

[17] P. Sekatski, M. Skotiniotis, J. Kołodyński, and W. Dür, Quantum metrology with full and fast quantum control, Quantum **1**, 27 (2017).

[18] R. Demkowicz-Dobrzański, J. Czajkowski, and P. Sekatski, Adaptive quantum metrology under general markovian noise, Phys. Rev. X **7**, 041009 (2017).

[19] S. Zhou, M. Zhang, J. Preskill, and L. Jiang, Achieving the Heisenberg limit in quantum metrology using quantum error correction, Nat. Commun. **9**, 78 (2018).

[20] H. Cramer, *Mathematical Methods of Statistics*, Princeton Mathematical Series (Princeton University Press, Princeton, NJ, 1954).

[21] C. R. Rao, Information and the accuracy attainable in the estimation of statistical parameters, in *Breakthroughs in Statistics: Foundations and Basic Theory*, edited by S. Kotz and N. L. Johnson (Springer, New York, 1992), pp. 235–247.

[22] R. D. Gill and B. Y. Levit, Applications of the van Trees Inequality: A Bayesian Cramér-Rao Bound, Bernoulli **1**, 59 (1995).

[23] S. L. Braunstein and C. M. Caves, Statistical distance and the geometry of quantum states, Phys. Rev. Lett. **72**, 3439 (1994).

[24] C. W. Helstrom, Quantum detection and estimation theory, J. Stat. Phys. **1**, 231 (1969).

[25] Y. Li, L. Pezze, M. Gessner, Z. Ren, W. Li, and A. Smerzi, Frequentist and bayesian quantum phase estimation, Entropy **20**, 628 (2018).

[26] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory* (Prentice-Hall, Englewood Cliffs, NJ, 1993).

[27] J. Kolodynski, Precision bounds in noisy quantum metrology, Ph.D. thesis, University of Warsaw, 2014.

[28] S. D. Bartlett, T. Rudolph, and R. W. Spekkens, Reference frames, superselection rules, and quantum information, Rev. Mod. Phys. **79**, 555 (2007).

[29] R. Demkowicz-Dobrzański, W. Górecki, and M. Guţă, Multiparameter estimation beyond quantum fisher information, J. Phys. A: Math. Theor. **53**, 363001 (2020).

[30] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Quantum circuit architecture, Phys. Rev. Lett. **101**, 060401 (2008).

[31] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Transforming quantum operations: Quantum supermaps, Europhys. Lett. **83**, 30004 (2008).

[32] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Theoretical framework for quantum networks, Phys. Rev. A **80**, 022339 (2009).

[33] S. Milz and M. T. Quintino, Transformations between arbitrary (quantum) objects and the emergence of indefinite causality, arXiv:2305.01247.

[34] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Memory effects in quantum channel discrimination, Phys. Rev. Lett. **101**, 180501 (2008).

[35] M. Ziman, Process positive-operator-valued measure: A mathematical framework for the description of process tomography experiments, Phys. Rev. A **77**, 062112 (2008).

[36] M. Zhang, H.-M. Yu, H. Yuan, X. Wang, R. Demkowicz-Dobrzański, and J. Liu, Quanestimation: An open-source toolkit for quantum parameter estimation, Phys. Rev. Res. **4**, 043057 (2022).

[37] R. Kaubruegger, A. Shankar, D. V. Vasilyev, and P. Zoller, Optimal and variational multiparameter quantum metrology and vector-field sensing, PRX Quantum **4**, 020333 (2023).

[38] J. J. Meyer, S. Khatri, D. S. França, J. Eisert, and P. Faist, Quantum metrology in the finite-sample regime, arXiv:2307.06370.

[39] J. Bavaresco, M. Murao, and M. T. Quintino, Strict hierarchy between parallel, sequential, and indefinite-causal-order strategies for channel discrimination, Phys. Rev. Lett. **127**, 200504 (2021).

[40] J. Bavaresco, M. Murao, and M. T. Quintino, Unitary channel discrimination beyond group structures: Advantages of sequential and indefinite-causal-order strategies, J. Math. Phys. **63**, 042203 (2022).

[41] A. Altherr and Y. Yang, Quantum metrology for non-Markovian processes, Phys. Rev. Lett. **127**, 060501 (2021).

[42] Q. Liu, Z. Hu, H. Yuan, and Y. Yang, Optimal strategies of quantum metrology with a strict hierarchy, Phys. Rev. Lett. **130**, 070803 (2023).

[43] P. Sekatski and M. Perarnau-Llobet, Optimal nonequilibrium thermometry in Markovian environments, Quantum **6**, 869 (2022).

[44] R. Demkowicz-Dobrzański, Optimal phase estimation with arbitrary a priori knowledge, Phys. Rev. A **83**, 061802(R) (2011).

[45] K. Macieszczak, M. Fraas, and R. Demkowicz-Dobrzański, Bayesian quantum frequency estimation in presence of collective dephasing, New J. Phys. **16**, 113002 (2014).

[46] L. Vandenberghe and S. Boyd, Semidefinite programming, SIAM Rev. **38**, 49 (1996).

[47] P. Skrzypczyk and D. Cavalcanti, *Semidefinite Programming in Quantum Information Science* (IOPP, Bristol, 2023), pp. 2053-2563.

[48] H. Wolkowicz, R. Saigal, and L. Vandenberghe, *Handbook of Semidefinite Programming: Theory, Algorithms, and Applications* (Springer, New York, 2012), Vol. 27.

[49] A. Jamiołkowski, Linear transformations which preserve trace and positive semidefiniteness of operators, Rep. Math. Phys. **3**, 275 (1972).

[50] M.-D. Choi, Completely positive linear maps on complex matrices, Linear Algebra Applicat. **10**, 285 (1975).

[51] A. Y. Kitaev, Quantum computations: Algorithms and error correction, Russ. Math. Surveys **52**, 1191 (1997).

[52] G. M. D'Ariano, P. Lo Presti, and P. Perinotti, Classical randomness in quantum measurements, J. Phys. A: Math. Gen. **38**, 5979 (2005).

[53] A. W. Harrow, The church of the symmetric subspace, arXiv:1308.6595.

[54] S. I. Knysh, E. H. Chen, and G. A. Durkin, True limits to precision via unique quantum probe, arXiv:1402.0495.

[55] D. W. Berry and H. M. Wiseman, Optimal states and almost optimal adaptive measurements for quantum interferometry, Phys. Rev. Lett. **85**, 5098 (2000).

[56] M. Mehboudi, A. Sanpera, and L. A. Correa, Thermometry in the quantum regime: Recent theoretical progress, J. Phys. A: Math. Theor. **52**, 303001 (2019).

[57] 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 (Springer, Cham, 2018), pp. 503–527.

[58] A. Kossakowski, On quantum statistical mechanics of non-Hamiltonian systems, Rep. Math. Phys. **3**, 247 (1972).

[59] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan, Completely positive dynamical semigroups of *N*-level systems, J. Math. Phys. **17**, 821 (1976).

[60] G. Lindblad, On the generators of quantum dynamical semigroups, Commun. Math. Phys. **48**, 119 (1976).

[61] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, Oxford, 2007).

[62] S. Jevtic, D. Newman, T. Rudolph, and T. M. Stace, Single-qubit thermometry, Phys. Rev. A **91**, 012331 (2015).

[63] L. A. Correa, M. Mehboudi, G. Adesso, and A. Sanpera, Individual quantum probes for optimal thermometry, Phys. Rev. Lett. **114**, 220405 (2015).

[64] J. Rubio, Quantum scale estimation, Quantum Sci. Technol. **8**, 015009 (2023).

[65] J. Rubio, J. Anders, and L. A. Correa, Global quantum thermometry, Phys. Rev. Lett. **127**, 190402 (2021).

[66] 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).

[67] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity* (Dover, New York, 1998).

[68] U. Feige, A threshold of ln n for approximating set cover, J. ACM **45**, 634 (1998).

[69] Rupert G. Miller Jr., *Simultaneous Statistical Inference* (Springer, New York, 1981).

[70] https://github.com/jessicabavaresco/singleshot-bayesian-estimation.

[71] A. Peres, Separability criterion for density matrices, Phys. Rev. Lett. **77**, 1413 (1996).

[72] A. C. Doherty, P. A. Parrilo, and F. M. Spedalieri, Distinguishing separable and entangled states, Phys. Rev. Lett. **88**, 187904 (2002).