

Sequential analysis in a continuous spin-noise quantum sensor

E. Roda-Salichs ^{1,*} G. Gasbarri ^{1,†} A. Alou ^{1,2,‡} M. Skotiniotis ^{1,§} A. Sierant ³ D. Méndez-Avalos ³
M. W. Mitchell ^{3,4,||} and J. Calsamiglia ^{4,¶}¹*Física Teòrica: Informació i Fenòmens Quàntics, Departament de Física, Universitat Autònoma de Barcelona, 08193 Bellaterra, Barcelona, Spain*²*Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Avda Complutense 22, 28040 Madrid, Spain*³*ICFO—Institut de Ciències Fotòniques, The Barcelona Institute of Science and Technology, 08860 Castelldefels, Barcelona, Spain*⁴*ICREA—Institució Catalana de Recerca i Estudis Avançats, 08010 Barcelona, Spain*

(Received 3 October 2025; accepted 10 June 2026; published 9 July 2026)

Many control and detection applications require real-time analysis of signals from sensors, in order to quickly and accurately act upon events revealed by the sensors. Such signal analysis benefits from statistical models of signal and sensor behavior. This creates a need for data analysis methods that are simultaneously model based, computationally efficient, and causal, in the sense that they employ only sensor data available prior to a specific point in time. In this work, we implement sequential data analysis techniques on a spin-noise-based quantum sensor to perform two key tasks: hypothesis testing and quickest change-point detection. These online protocols allow us to detect weak magnetic fields by adaptively collecting measurement data until a predefined confidence threshold is reached. We demonstrate these methods in a realistic experimental setting and derive performance bounds for the achievable precision and response time. Our approach has potential utility when detecting small perturbations to the magnetic field, in both applied and fundamental contexts including biomagnetism, geophysical surveys, detection of concealed materials, searches for dark matter candidates, and exotic spin interactions. Our results demonstrate that sequential techniques enable faster and more sensitive detection, making them a powerful tool for quantum sensing.

DOI: [10.1103/z7zj-864q](https://doi.org/10.1103/z7zj-864q)

I. INTRODUCTION

Continuously monitored quantum systems, such as atomic magnetometers and optomechanical sensors, offer a promising avenue for ultraprecise quantum sensing [1–11]. These platforms provide opportunities for real-time statistical inference in applications ranging from biomedical diagnostics to materials characterization [5,10,12]. Their ability to operate in real time makes them particularly suitable for detecting and tracking stochastic time-varying signals. As such, they

play a key role in applications requiring online and efficient detection of abrupt transitions or sudden changes in dynamics.

Despite remarkable theoretical and experimental progress [13–29], many aspects of real-time quantum sensing protocols remain largely unexplored.

In most standard approaches to inference tasks, particularly in hypothesis testing, signal information is analyzed only after data collection or after multiple experimental repetitions, thereby limiting the possibility of real-time decision-making. Sequential analysis offers a fundamentally different paradigm: Decisions are made progressively as data arrives, enabling real-time inference while maintaining rigorous statistical guarantees. Sequential analysis has deep roots in classical statistics dating back to Wald's pioneering work in the 1940s [30]. The field has experienced significant renewed interest across multiple scientific domains. Modern developments include algorithms for quickest detection of abrupt changes in data streams [31], and theoretical advances in safe anytime-valid inference and e-processes [32] that address fundamental challenges of continuous monitoring, optional stopping, and decision-making under finite data constraints. These methods have found increasing practical application in large-scale hypothesis-testing scenarios, from clinical trial monitoring to high-throughput experimental platforms.

A few seminal works have laid the foundation for quantum sequential analysis in settings where identically prepared copies of a quantum system are supplied on demand. In Refs. [33,34], the quantum analog of Wald's [30] problem of

*Contact author: elisabet.roda@uab.cat

†Present address: Naturwissenschaftlich-Technische Fakultät, Universität Siegen, Walter-Flex-Straße 3, 57068 Siegen, Germany.

‡Present address: Port d'Informació Científica (PIC), Campus UAB, C. Albareda s/n, 08193 Bellaterra, Spain.

§Present address: Quantum Thermodynamics and Computation Group, Department of Electromagnetism and Condensed Matter, Universidad de Granada, 18071 Granada, Spain and Instituto Carlos I de Física Teórica y Computacional, Universidad de Granada, 18071 Granada, Spain.

||Contact author: morgan.mitchell@icfo.eu¶Contact author: john.calsamiglia@uab.cat

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

sequential hypothesis testing was developed. In this setting, one aims to determine whether a quantum source is emitting states ρ_0 or ρ_1 , using measurements on successive, identically prepared systems, and stopping the process once sufficient evidence has been accumulated to make a reliable decision within a given confidence level. Similarly, Lorden's [35] classical problem of *quickest change-point detection* (QCPD), which seeks to identify as promptly as possible the moment when the statistical properties of a sequence change, was extended to the quantum domain in Ref. [36]. There, the system emits quantum states initially prepared as ρ_0 , which at an unknown time switch to ρ_1 , and the task is to detect this change with minimal delay while keeping false alarms under control. These works have established fundamental bounds on performance in such independent and identically distributed (IID) quantum scenarios, showing that the relevant asymptotic figure of merit is the quantum relative entropy $D(\rho_0\|\rho_1)$, just as the Kullback-Leibler divergence plays a central role in classical statistics. More recently, quantum-enhanced strategies for quickest change-point detection have been applied in the context of quantum communication [37–39], indicating their relevance for real-time decision-making across a broad range of quantum technologies. Related inference or stopping rule problems have appeared implicitly in early studies of continuous quantum error correction and quantum filtering, where random error events were inferred from noisy continuous measurement records and corrected via real-time feedback [40,41].

This paper considers a distinct quantum sensing paradigm in which, instead of performing a sequence of measurements on an increasing number of independently prepared copies, a continuous sequence of measurements is carried out on the *same* quantum system. The goal is to discriminate between alternative internal dynamics governing the system. This approach, explored theoretically in the context of binary hypothesis testing using sequential strategies [42], offers notable advantages: It enables real-time analysis of individual trajectories without requiring ensemble averages, and it can significantly reduce measurement time by allowing data collection to stop as soon as a hypothesis is identified with the desired level of confidence.

In this work, we report what is, to our knowledge, the first experimental implementation of sequential analysis methodologies in continuously monitored quantum sensors. We demonstrate this framework using two foundational sequential analysis primitives, sequential hypothesis testing and quickest change-point detection, applied to atomic spin-noise magnetometry. Atomic magnetometry is among the most mature and widely adopted in quantum technologies, with demonstrated relevance in diverse fields [24]. The sensing platform of optically pumped and probed atomic vapors is applied in medicine [43,44], materials science [45], geophysics [46], space science [47], and fundamental physics [48,49]. The statistical properties of spin-noise dynamics has been studied extensively [50–52], including quantum-enhanced spectroscopy with squeezed light [53,54] and measurement-induced spin squeezing [55], and found to have excellent agreement with gaussian statistical models [22,56], including in high-density regimes dominated by spin-exchange interactions [57,58].

This work is structured as follows: In Sec. II, we present hypothesis testing, both deterministic and sequential strategies, and quickest change-point-detection scenarios. Here, and throughout this work, we use the term *deterministic* to refer to standard approaches that rely on a predetermined, fixed number of samples, in contrast to sequential strategies, where the number of samples is stochastic, as it depends on previous measurement outcomes. In Sec. III, we explain the physical model, the dynamics behind the spin-noise quantum sensor, and the probability distribution that the observations follow. We then proceed in Sec. IV to show the experimental data and how we tackle the problem of low-frequency spurious noise in the system. With this at hand, in Sec. V we present the results of sequential analysis in the spin-noise system. Finally, in Sec. VI we derive the optimal asymptotic performance limits for the studied inference strategies.

II. SEQUENTIAL ANALYSIS: BASIC CONCEPTS AND METHODS

To introduce the class of sequential strategies considered in this work, we begin with sequences of n observations denoted by $\mathbf{R}_n = (R_1, \dots, R_n)^T$, where n may grow as new data become available. In hypothesis testing, we are given two known models labeled $h = 0$ and $h = 1$, with the assumption that the sequence $\mathbf{R}_n$ is generated by one of them. The goal is to sequentially evaluate the plausibility of these hypotheses as more observations are collected, terminating when the evidence reaches a predetermined confidence level.

The change-point-detection scenario, on the other hand, assumes that the observations initially follow model $h = 0$, but may transition to model $h = 1$ at some unknown point in the sequence. Here, the task becomes monitoring the stream of observations to detect such a change as quickly as possible while controlling the rate of false alarm.

As we will shortly see, in both sequential hypothesis testing and quickest change-point detection, the central quantity that must be computed and monitored is the log-likelihood ratio (LLR) statistic:

$$L_{R,n} := \ln \frac{p_1(\mathbf{R}_n)}{p_0(\mathbf{R}_n)}. \quad (1)$$

Here, $p_i(\mathbf{R}_n) := p(\mathbf{R}_n|h_i)$, for $i = 0, 1$, denotes the probability density function of $\mathbf{R}_n$ under h_i . Tracking this single scalar statistic is sufficient to implement our optimal stopping rules and inference decisions.

A. Binary hypothesis testing

In binary hypothesis testing, the goal is to decide which of two mutually exclusive hypotheses, labeled $h = 0$ and $h = 1$, correctly describes a system or process based on observed data. Each hypothesis corresponds to a distinct probability distribution for the observations. Since decisions are made from finite noisy data, errors are unavoidable. These errors are quantified by two standard metrics: the type-I error (false positive), α_1 , which occurs when $h = 0$ is rejected even though it is true, and the type-II error (false negative), α_0 , which occurs when $h = 0$ is accepted while the true hypothesis is $h = 1$. Broadly, the goal of hypothesis testing is to design decision

strategies that minimize these errors using the least amount of resources, such as the minimum number of measurements or samples.

1. Hypothesis testing with deterministic time

In the standard approach, a sequence of measurement data of predetermined length n is collected, after which a definitive decision regarding the true hypothesis is made. Such decision faces an inherent trade-off between the two error probabilities α_0 and α_1 , which generally cannot be minimized simultaneously. Different testing paradigms emerge depending on the optimization criteria. An asymmetric test strategy minimizes one error probability while constraining the other below a specified threshold, which is particularly relevant when one error type carries greater consequences than the other. In contrast, symmetric testing naturally appears in Bayesian frameworks where prior probabilities $\{\pi_0, \pi_1\}$ are assigned to the different hypotheses. Here, the optimal strategy is the one that minimizes the average error probability

$$P_{\text{error}} = \pi_0 \alpha_1 + \pi_1 \alpha_0. \quad (2)$$

In order to perform the optimal decision in these problems, it is sufficient to compute the LLR statistic $L_n(\mathbf{R}_n)$ from Eq. (1), choose an appropriate threshold a , and carry out the *likelihood ratio test* [59]:

- (1) accept $h = 1$, if $L_n > a$,
 - (2) accept $h = 0$, if $L_n \leq a$,
- which renders the error probabilities,

$$\begin{aligned} \alpha_0 &= P_1(L_n < a) = \int_{-\infty}^a p_1(\ell) d\ell, \\ \alpha_1 &= P_0(L_n \geq a) = \int_a^{\infty} p_0(\ell) d\ell, \end{aligned} \quad (3)$$

where $p_h(\ell)$ is the probability density function of the test statistic L_n under hypothesis $h = 0, 1$. The threshold a controls the trade-off between type-I and type-II errors: Decreasing a reduces α_1 at the cost of increasing α_0 , and vice versa. For equal prior probabilities ($\pi_0 = \pi_1$), the optimal decision rule is the *maximum-likelihood estimator*, which corresponds to setting $a = 0$.

For a large number n of IID observations, the error probabilities decay exponentially with n , with rates characterized by fundamental information-theoretic quantities [59]. In particular, the Chernoff information

$$C(p_0, p_1) := - \min_{0 \leq s \leq 1} \ln \int p_0(x)^s p_1(x)^{1-s} dx \quad (4)$$

determines the optimal error exponent for symmetric hypothesis testing, satisfying

$$- \lim_{n \rightarrow \infty} \frac{1}{n} \ln P_{\text{err}} = C(p_0, p_1). \quad (5)$$

In asymmetric hypothesis testing, the Kullback-Leibler divergence (relative entropy)

$$D(p_0 \| p_1) := \int p_0(x) \ln \frac{p_0(x)}{p_1(x)} dx \quad (6)$$

governs the optimal type-II error exponent. Specifically, when constraining the type-I error probability $\alpha_1 < \epsilon$ for some con-

stant $\epsilon < 1/2$, we have

$$- \lim_{n \rightarrow \infty} \frac{1}{n} \ln \alpha_0 = D(p_0 \| p_1). \quad (7)$$

This represents the fastest achievable error rate, as guaranteed by the converse of Stein's lemma [59]: Any attempt to achieve a better error exponent would necessarily drive the type-I error probability to 1, resulting in complete loss of discrimination power.

Equivalent asymptotic expressions can also be derived for (non-IID) Gaussian noise processes, like those arising in our setup (see Refs. [16,60] and Sec. VI).

2. Sequential hypothesis testing

In the context of continuously monitored sensors, rather than fixing the duration of the experiment and then assessing the error probability, it is more natural to perform inference *on the fly*, continuing measurements until a desired error threshold is reached. In this setting, the inference strategy involves not only a decision rule but also a stopping time—a criterion that dictates whether to continue or stop based on the observations. Note that since the observations are stochastic, the stopping time is itself a random variable.

In the Bayesian framework, such a stopping rule directly follows from the operational requirement that one may stop only when, based on the current measurement record, the probability of correctly identifying hypothesis $h = 0$ is at least $s_0 = 1 - \epsilon_0$, and that of $h = 1$ is at least $s_1 = 1 - \epsilon_1$, i.e.,

$$p(h_k | \mathbf{R}_n) = \frac{p(\mathbf{R}_n | h_k) \pi_k}{p(\mathbf{R}_n | h_0) \pi_0 + p(\mathbf{R}_n | h_1) \pi_1} \geq 1 - \epsilon_k, \quad (8)$$

where $k = 0, 1$ and we have made use of Bayes' rule. The above inequalities can be conveniently expressed in terms of the log-likelihood ratio:

- (1) Stop and accept $h = 0$ if $L_n \leq -a_0 := -\ln(\frac{\pi_0}{\pi_1} \frac{1-\epsilon_0}{\epsilon_0})$.
- (2) Stop and accept $h = 1$ if $L_n \geq a_1 := -\ln(\frac{\pi_1}{\pi_0} \frac{1-\epsilon_1}{\epsilon_1})$.
- (3) Continue if $-a_0 < L_n < a_1$.

This protocol is known as the *sequential probability ratio test* (SPRT) [30] and its stopping time is given by

$$N := \inf\{n : L_n(\mathbf{R}_n) \notin (-a_0, a_1)\}. \quad (9)$$

Since the decision to stop depends only on data observed up to time $n \leq N$, it defines a valid stopping time in the sense of stochastic decision theory. Notice that for $\epsilon_0 = \epsilon_1 = \epsilon$, the error probability from Eq. (2) satisfies $P_{\text{err}} \leq \epsilon$.

Wald [30,61] showed that the thresholds a_0, a_1 in the SPRT algorithm are also in one-to-one correspondence with the type-I and type-II errors, whose definitions do not involve Bayesian priors. Specifically, for a process that stops either at $L_N = -a_0$ or at $L_N = a_1$,

$$\alpha_0 = \frac{1 - e^{-a_1}}{e^{-a_0} - e^{-a_1}}, \quad \alpha_1 = \frac{1 - e^{-a_0}}{e^{a_1} - e^{-a_0}}. \quad (10)$$

Wald and Wolfowitz [62] showed that, among all strategies constrained by given type-I and type-II error bounds, the SPRT minimizes the expected stopping time under each hypothesis. Analogously to the deterministic error rates [Eqs. (4) and (6)] for sequential strategies, an asymptotic relation connects the expected number of observations to the error

bounds:

$$\mathbb{E}_0[N] \sim -\frac{\ln \alpha_0}{D(p_0 \| p_1)} \quad \text{and} \quad \mathbb{E}_1[N] \sim -\frac{\ln \alpha_1}{D(p_1 \| p_0)} \quad (11)$$

in the limit of small error bounds $\alpha_0, \alpha_1 \ll 1$. Clearly, sequential strategies achieve the same error probabilities with far fewer observations *on average*. In particular, the SPRT achieves the fastest possible exponential decay of both error probabilities, each governed by the corresponding Kullback-Leibler divergence. This is in contrast to fixed-length (deterministic) strategies that can attain such a decay rate for one error, but only at the expense of keeping the other error probability finite.

In non-IID and quantum scenarios, recent works [33,34,42] have also demonstrated the advantages of the SPRT, further extending its relevance beyond classical settings.

B. Quickest change-point detection

Detecting abrupt changes in continuously monitored quantum systems is a fundamental task, as it captures the general challenge of identifying sudden shifts in the underlying dynamics while balancing detection speed and reliability. Beyond its conceptual importance, this problem has direct practical relevance in contexts such as early detection of fires or hazardous substances, or other environmental and medical monitoring, where timely responses to unexpected events are crucial.

In this subsection, we review the problem of quickest change-point detection, which aims at identifying a sudden change in the underlying quantum dynamics while minimizing both false alarms and detection delay. The cumulative sum (CUSUM) algorithm first proposed by Page [63] is one of the most widely used methods for quickest change-point detection. Its popularity stems from two main features: its computational efficiency, which allows real-time implementation, and optimality under certain mild conditions. In the IID scenario, Lorden [35] showed that CUSUM minimizes the worst mean detection delay among the strategies with a given bound on the mean false alarm time. Lorden's asymptotic optimality proof was later shown to hold also for finite times [64,65]. The CUSUM algorithm was extended to non-iid settings and its asymptotic optimality was also proven under some reasonable conditions by Lai [66].

In the IID quickest change-point-detection problem, the observations are assumed to be independent and distributed according to a prechange distribution $p_0(R_n)$ before the unknown change time $n < \nu$ and a postchange distribution $p_1(R_n)$ for $n \geq \nu$. In more general, non-IID settings, the statistical model is specified through the conditional distribution of each observation given the past: The probability density function of R_n , conditional on $R_1 \cdots R_{n-1}$, is given by

$$p_0(R_n | R_1 \cdots R_{n-1}), \quad \text{for } n < \nu \quad (\text{before change}), \quad (12)$$

$$p_1(R_n | R_1 \cdots R_{n-1}), \quad \text{for } n \geq \nu \quad (\text{postchange}). \quad (13)$$

The CUSUM algorithms can be described in terms of the likelihood ratio between the change-point occurring at time

$1 \leq k \leq n$ and no change occurring in that time:

$$L_n^{(k)} := \ln \frac{p_1(R_{1:n})}{p_0(R_{1:n})} \quad (14)$$

$$\begin{aligned} &= \ln \frac{\prod_{i=1}^{k-1} p_0(R_i | R_{1:i-1}) \prod_{i=k}^n p_1(R_i | R_{1:i-1})}{\prod_{i=1}^n p_0(R_i | R_{1:i-1})} \\ &= \ln \prod_{i=k}^n \frac{p_1(R_i | R_{1:i-1})}{p_0(R_i | R_{1:i-1})} = \sum_{i=k}^n \Delta L_i, \end{aligned} \quad (15)$$

where the likelihood ratio increment at step i is defined as

$$\Delta L_i = \ln \frac{p_1(R_i | R_{1:i-1})}{p_0(R_i | R_{1:i-1})}, \quad (16)$$

which is positive whenever the event at i is more likely to be caused under the changed hypothesis than under the default one.

The crucial statistic is the cumulative sum:

$$M_n = \max_{1 \leq k \leq n} \sum_{i=k}^n L_i. \quad (17)$$

The CUSUM algorithm outputs change-point detection whenever this quantity exceeds a given threshold a and the corresponding stopping, or detection, time is given by

$$T = \inf\{n : M_n \geq a\}. \quad (18)$$

The threshold a is chosen so that the mean false alarm time is guaranteed to exceed a given positive number:

$$T_{\text{FA}} := \mathbb{E}_0(T) \geq \ln a. \quad (19)$$

Under such condition on the false alarm time, Lai showed that the stopping time T minimizes the worst-worst case mean delay [31,66]

$$\bar{\tau}^* := \sup_{\nu \geq 0} \text{ess sup}_{\nu} \mathbb{E}_\nu(T - \nu | T > \nu, R_{1:\nu}). \quad (20)$$

This figure of merit considers the worst delay over all possible locations of the change point and over (almost) all measurement outcomes before the change. A less strict way of assessing the performance of the protocol is the maximal conditional average delay of detection:

$$\bar{\tau} := \sup_{\nu \geq 0} \mathbb{E}_\nu(T - \nu | T > \nu) \leq \bar{\tau}^*. \quad (21)$$

We note that, as a valuable by-product, whenever the CUSUM algorithm signals a detection, i.e., $M_n > a$, the value of k that maximizes Eq. (17) provides a very good guess of where the change actually occurred. Indeed, from the definition in Eq. (14) it is clear that it corresponds to the most likely hypothesis $1 \leq k \leq n$. That is, the CUSUM algorithm effectively checks at every time step if the current measurement record is a given factor more likely to be given due to a change at a previous time than due to the default hypothesis.

In the IID setting, one can give a closed-form formula that quantifies the trade-off between the false alarm time and the mean detection time:

$$\bar{\tau} \sim \frac{\ln T_{\text{FA}}}{D(p_0 \| p_1)}. \quad (22)$$

As we will see in Secs. VB and VI, in our setting a similar relation holds where the regularized expectation value




FIG. 1. Scheme of the atomic sensor: An ensemble of ^{87}Rb atoms precesses at the Larmor frequency ω_L under the action of an external magnetic field $\mathbf{B}$. A probe laser beam monitors this precession via the Faraday effect, generating a photocurrent proportional to J_z , the z component of the collective spin, with added shot noise.

of the likelihood ratio plays the role of the relative entropy $D(p_0 \| p_1)$.

The same notation for T , T_{FA} , ν , $\bar{\tau}^*$, $\bar{\tau}$ will be used for both index values in a time series and actual times (discrete or continuous). The intended meaning will be clear from the context or explicitly specified when needed.

III. SPIN-NOISE QUANTUM SENSOR

This section outlines the experimental setup of the quantum sensor (Sec. III A) and the corresponding probability distribution of the measured signal (Sec. III B). These elements provide the foundation for calculating the L_n statistic in order to apply the sequential protocols described in Sec. II. However, since computing L_n directly from this distribution in real time is impractical, in Sec. III C we present a recursive method to obtain it using the so-called Kalman filter. Finally, in Sec. III D to enable an analytical and asymptotic analysis of the log-likelihood ratio in further sections, we examine the signal in the frequency domain and treat the dynamics of the system as a linear time-invariant filter.

Our experimental setup for the spin-noise sensor is shown schematically in Fig. 1. It consists of an ensemble of ^{87}Rb atoms, whose joint state is well described by the collective spin operators $\mathbf{J}_t$, with components $J_i(t)$ for $i = x, y, z$. These operators capture the total spin of the ensemble at time t and govern its dynamical evolution. Under the action of an external magnetic field $\mathbf{B} = B\mathbf{e}_x$, the atomic spin $\mathbf{J}_t$ undergoes a precession around the direction of $\mathbf{B}$, at the so-called Larmor frequency ω_L , which is proportional to the magnitude of $\mathbf{B}$. This precession can be continuously measured via the magneto-optical Faraday effect, using an off-resonant light beam propagating along the z axis and linearly polarized in the x direction. As the spin precesses, the plane of polarization of the probe beam rotates proportionally to J_z , and this rotation is detected by a polarimeter that outputs a photocurrent signal carrying information about the spin dynamics, superimposed with optical shot noise.

The measured photocurrent provides a continuous time trace whose statistical properties are directly determined by the underlying spin dynamics and, consequently, by the value of the magnetic field through the Larmor frequency ω_L . Atomic magnetometers based on this principle are among the most sensitive magnetic-field sensors available, achieving high precision without requiring cryogenic cooling, large superconducting magnets, or other complex infrastructure. Their relatively low cost, portability, and ability to operate under ambient conditions make them attractive for a wide range of applications, including biomagnetic sensing, geophysical prospecting, navigation, and the detection of concealed metallic objects. In this work, we focus on hypothesis-testing and change-detection scenarios in which the magnetic field can take one of two possible values, corresponding to different Larmor frequencies. Such situations arise, for example, when the presence of a nearby object perturbs a background magnetic field or when a physiological signal induces a transient magnetic variation. Our goal is to determine, from the continuously monitored photocurrent, which magnetic-field hypothesis is realized and, in particular, to detect abrupt transitions between them as quickly as possible.

A. Experimental details

Isotopically enriched ^{87}Rb atoms are contained in a 3-cm-long vapor cell filled with 100 torr of N_2 buffer gas (pressure specified at room temperature). The cell is enclosed in four layers of mu-metal shielding to suppress environmental magnetic fields. The atoms are heated inside a ceramic oven to 102°C using intermittent Joule heating, yielding an atomic number density of 8.2×10^{12} atoms/cm³. Induction coils driven by a waveform generator (Keysight 33500B) followed by a low-noise current source (TwinLeaf CSUA300) generate DC magnetic fields along the x axis with strengths around $B \approx 7.22 \mu\text{T}$, corresponding to a Larmor frequency of $\omega_L \approx 2\pi \times 50.5 \text{ kHz}$. The bias magnetic-field step was implemented via a pulse function from the waveform generator and synchronized with the data acquisition system.

The polarization rotation of a linearly polarized probe beam ($830 \mu\text{W}$) is measured with a shot-noise-limited polarimeter consisting of a half-wave plate, a Wollaston prism, and a balanced photodetector (Thorlabs PDB450A). The probe is derived from a tunable diode laser (Toptica DL100) operating at 795 nm, frequency stabilized via a fiber interferometer [67] to be 27 GHz blue detuned from the D^1 line. The detector output is digitized with a National Instruments PCI-4462 data acquisition card with 200 kSa/s sampling rate. Data were collected in 200 repeated samples, each acquisition lasting 2 s.

B. Physical model and underlying probability distribution of the signal

The stochastic model that describes the dynamics of the components of spin and the measured photocurrent I is described in detail in Ref. [22]. Relying on the Itô formulation for stochastic differential equations [68], the dynamics of the spin components in the transverse yz plane, $\mathbf{J}_t := (J_y(t), J_z(t))^T$, are described by the Ornstein-Uhlenbeck

process:

$$d\mathbf{J}_t = \hat{\mathbf{A}} \cdot \mathbf{J}_t dt + \hat{\mathbf{N}} d\mathbf{W}, \quad \text{with} \quad \hat{\mathbf{A}} = \begin{pmatrix} -\gamma & \omega_L \\ -\omega_L & -\gamma \end{pmatrix}, \quad (23)$$

$\hat{\mathbf{N}} = \sqrt{Q}\mathbb{1}_2$, and $d\mathbf{W} = (dW_x, dW_y)^T$ with dW_x and dW_y two independent Wiener noises. The matrix $\hat{\mathbf{A}}$ encapsulates coherent Larmor precession at frequency ω_L and transverse relaxation at rate γ . Both γ and the noise term with strength $\hat{\mathbf{N}}$ originate from different noise processes including atomic collisions, optical depolarization, and transit-time broadening. Together, these parameters define a linear Gaussian model for the spin dynamics under continuous monitoring.

The photocurrent I generated by the optical detection process is sampled at finite-time intervals $t_k = k\Delta$, where $k \in \mathbb{N}$ and Δ denotes the sampling period. The photocurrent captures information about the collective spin while being affected by optical shot noise $\xi(t_k)$, which is assumed to be Gaussian and statistically independent of $\mathbf{J}_t$. The sensor output is thus described by the discrete-time stochastic equation:

$$I_k = \hat{\mathbf{C}} \mathbf{J}_{t_k} + \xi(t_k) = g_D J_z(t_k) + \xi(t_k), \quad (24)$$

where $\hat{\mathbf{C}} = g_D(0, 1)$ defines the measurement coupling, with g_D the transduction constant. The measurement noise $\xi(t_k)$ is zero-mean Gaussian white noise $\xi(t_k) \sim \mathcal{N}(0, G_k)$, with variance $G_k = S_{\text{ph}}/\Delta \forall k \in \mathbb{N}$, where S_{ph} is the spectral density of the photon shot noise.

In the steady state, the probability distribution of n photocurrent observations, $\mathbf{I}_n = (I_1, \dots, I_n)^T$, follows a multivariate normal distribution $\mathbf{I}_n \sim \mathcal{N}(\boldsymbol{\mu}_n, \hat{\mathbf{T}}_n)$, with mean $\boldsymbol{\mu}_n = \mathbf{0}$ and covariance matrix $\hat{\mathbf{T}}_n$ with elements (see Appendix A):

$$T_{i,j} = K_{i-j},$$

with

$$K_k := g_D^2 \frac{Q}{2\gamma} e^{-\gamma|k|\Delta} \cos(k\omega_L\Delta) + \delta_{k,0} \frac{S_{\text{ph}}}{\Delta}. \quad (25)$$

Explicitly,

$$p(I_1, \dots, I_n) = \frac{\exp\left(-\frac{1}{2} \mathbf{I}_n^T \hat{\mathbf{T}}_n^{-1} \mathbf{I}_n\right)}{\sqrt{(2\pi)^n \det \hat{\mathbf{T}}_n}}. \quad (26)$$

From Eq. (26), one can see that the log-likelihood ratio $L_{I,n}$ in Eq. (1) is given by

$$L_{I,n} = \frac{1}{2} \ln \left(\frac{\det \hat{\mathbf{T}}_{0,n}}{\det \hat{\mathbf{T}}_{1,n}} \right) + \frac{1}{2} \mathbf{I}_n^T (\hat{\mathbf{T}}_{0,n}^{-1} - \hat{\mathbf{T}}_{1,n}^{-1}) \mathbf{I}_n, \quad (27)$$

where $\hat{\mathbf{T}}_{h,n}$ stands for the covariance matrix under hypothesis h after n measurements. Notice that, since the process is stationary, each descending diagonal (from left to right) of the covariance matrix is constant; i.e., $\hat{\mathbf{T}}_n$ is a Toeplitz matrix. We will exploit this structure, as the spectral properties of Toeplitz matrices and their associated functionals are well studied (see Appendix B). For instance, the mean of $L_{I,n}$ can be immediately read off Eq. (27), by noting that $\mathbb{E}_h[\mathbf{I}_n \mathbf{I}_n^T] = \hat{\mathbf{T}}_{h,n}$. Similarly, the characteristic function of $L_{I,n}$ under hypothesis h , which completely determines its probability density func-

tion after n observations, can be computed as

$$\phi_{L_I}(u) := \mathbb{E}_h[e^{iuL_I}] = \frac{\eta^{\frac{iu}{2}}}{\sqrt{\det(\mathbb{1} - iu\hat{\mathbf{D}}\hat{\mathbf{T}}_{h,n})}}, \quad (28)$$

where we defined $\hat{\mathbf{D}} := \hat{\mathbf{T}}_{0,n}^{-1} - \hat{\mathbf{T}}_{1,n}^{-1}$ and $\eta = \frac{\det \hat{\mathbf{T}}_{0,n}}{\det \hat{\mathbf{T}}_{1,n}}$. From here, we can easily obtain the m th cumulant $\kappa_{m,n}$ of the LLR at time step n :

$$\begin{aligned} \kappa_{m,n}^h &= i^{-m} \frac{d^m \ln \phi(u)}{du^m} \Big|_{u=0} \\ &= \frac{1}{2} (m-1)! \text{Tr}((\hat{\mathbf{D}}\hat{\mathbf{T}}_{h,n})^m) + \delta_{m,1} \frac{1}{2} \ln \eta. \end{aligned} \quad (29)$$

In the limit of large n , for Toeplitz matrices arising from short-memory stationary processes, such as those defined in Eq. (27), the operator traces appearing in Eq. (29) can be computed (see Appendix B):

$$\begin{aligned} \kappa_{m,n}^h &= \frac{t}{2\pi} (m-1)! \int_0^\infty \left(\frac{\bar{S}_{I,h}(\omega)}{\bar{S}_{I,0}(\omega)} - \frac{\bar{S}_{I,h}(\omega)}{\bar{S}_{I,1}(\omega)} \right)^m d\omega \\ &\quad + \frac{t}{2\pi} \delta_{m,1} \int_0^\infty \ln \left(\frac{\bar{S}_{I,0}(\omega)}{\bar{S}_{I,1}(\omega)} \right) d\omega + \mathcal{O}(1), \end{aligned} \quad (30)$$

where $t = n\Delta$ and $\bar{S}_{I,h}(\omega)$ is the power spectral density (PSD) of the signal, given by the Fourier transform of the autocorrelation function, i.e.,

$$\begin{aligned} \bar{S}_{I,h}(\omega) &:= \int_{-\infty}^\infty K_{h,t} e^{-i\omega t} dt \\ &= S_{\text{at}} \left(\frac{\gamma^2}{\gamma^2 + (\omega - \omega_L)^2} + \frac{\gamma^2}{\gamma^2 + (\omega + \omega_L)^2} \right) + S_{\text{ph}}, \end{aligned} \quad (31)$$

with $S_{\text{at}} := \frac{g_D^2 Q}{2\gamma^2}$. Here, and in the following, the subleading $\mathcal{O}(\Delta)$ terms from the continuum approximation are omitted for clarity. Also, with a slight abuse of notation, we write K_t for the continuous-time analog of the discrete K_k of Eq. (25) in the limit $\Delta \rightarrow 0$ and similarly generalize this notation to other quantities.

Equation (30) shows that the mean, variance, and higher-order cumulants of the log-likelihood ratio L_I grow linearly with time (up to subleading constant corrections), indicating that the distribution of L_I becomes increasingly peaked and concentrates around its mean.

C. Hybrid Kalman filter for probability distribution of the signal

The Kalman Filter provides a recursive Bayesian algorithm to infer the hidden state of linear Gaussian systems. It is widely used to estimate time-varying signals, as it yields the optimal estimator minimizing the mean-squared error (MSE) [69,70], e.g., in Ref. [22]. In our case, we use the KF not to estimate the hidden signal itself, but to recursively compute the probability distribution in Eq. (26).

The main idea is that the probability of the full measurement record $p(I_1, \dots, I_n)$ can be built iteratively via the chain rule. The recursion starts from a prior $p(\mathbf{J}_{t_0})$, which we take as the stationary distribution of the hidden process under the

hypothesis being tested. The state is then updated recursively through prediction and update steps as new observations become available.

At step $k = 1$, after observing I_1 , we update the prior via Bayes' rule:

$$p(\mathbf{J}_{t_1} | I_1) = \frac{p(I_1 | \mathbf{J}_{t_1}) p(\mathbf{J}_{t_1})}{\int p(I_1 | \mathbf{J}_{t_1}) p(\mathbf{J}_{t_1}) d\mathbf{J}_{t_1}}. \quad (32)$$

For $k > 1$, the prediction step propagates the posterior via the state-space transition model:

$$p(\mathbf{J}_{t_k} | I_{1:k-1}) = \int p(\mathbf{J}_{t_k} | \mathbf{J}_{t_{k-1}}) p(\mathbf{J}_{t_{k-1}} | I_{1:k-1}) d\mathbf{J}_{t_{k-1}},$$

followed by the update

$$p(\mathbf{J}_{t_k} | I_{1:k}) = \frac{p(I_k | \mathbf{J}_{t_k}) p(\mathbf{J}_{t_k} | I_{1:k-1})}{p(I_k | I_{1:k-1})}, \quad (33)$$

with normalization

$$p(I_k | I_{1:k-1}) = \int p(I_k | \mathbf{J}_{t_k}) p(\mathbf{J}_{t_k} | I_{1:k-1}) d\mathbf{J}_{t_k}, \quad (34)$$

where we have used the shorthand notation $I_{1:k} \equiv (I_1, \dots, I_k)$.

Since both the dynamical model and the measurement model are Gaussian, all these integrals can be computed analytically, yielding Gaussian distributions at each step [71]. In particular, $p(I_k | I_{1:k-1}) = \mathcal{N}(\mu_k, \sigma_k^2)$, where the mean μ_k and variance σ_k^2 are explicit symbolic functions of the model parameters and the previous data, computed recursively through the Kalman filter equations.

By iterating these steps and using the chain rule, we obtain the full likelihood of the observed data as

$$p(I_{1:n}) = \prod_{k=1}^n p(I_k | I_{1:k-1}) = \prod_{k=1}^n \frac{e^{-\frac{(I_k - \mu_k)^2}{2\sigma_k^2}}}{\sqrt{2\pi\sigma_k^2}}. \quad (35)$$

When comparing two hypotheses $h = 0, 1$, where each hypothesis specifies different system parameters and therefore distinct Kalman estimators $(\mu_{k,h}, \sigma_{k,h}^2)$, the log-likelihood ratio reads

$$L_{I,n} = \sum_{k=1}^n \left(\ln \frac{\sigma_{k,0}}{\sigma_{k,1}} + \frac{1}{2\sigma_{k,0}^2} (I_k - \mu_{k,0})^2 - \frac{1}{2\sigma_{k,1}^2} (I_k - \mu_{k,1})^2 \right). \quad (36)$$

Notice that Eq. (36) is the discrete-time counterpart of the log-likelihood-ratio process derived in Ref. [42] for general continuously monitored Gaussian systems. To account for continuous dynamics with discretely sampled measurements, we employ the continuous-discrete (hybrid) Kalman filter. The explicit algorithm is detailed in Appendix C.

D. Process as a filter and probability distribution in the asymptotic regime

As mentioned in the beginning of Sec. III, the dynamics of observations comprised in Eqs. (23) and (24) can be understood as a linear time-invariant filter that inputs Wiener noise and outputs the detected photocurrent.

A complete characterization of such continuous-time systems is provided by the transfer matrix function, which describes the system in the Laplace domain and fully specifies the response of the system to any input. By taking the Laplace transform, denoted by $\mathcal{L}$, of Eq. (23), we obtain

$$\mathcal{L}(\mathbf{J}_t) = \hat{H}_D(s) \mathcal{L}(\mathbf{W}_t), \quad (37)$$

where the transfer matrix is $\hat{H}_D(s) = (s\mathbb{1} - \hat{A})^{-1}$.

Restricting the Laplace variable to the imaginary axis, $s = i\omega$, yields the Fourier transform, $\mathcal{F}$, of the system:

$$\mathcal{F}(\mathbf{J}_t) = \hat{H}_D(i\omega) \mathcal{F}(\mathbf{W}_t). \quad (38)$$

This expression allows us to understand the process as a simple input-output relation in frequency domain in the long-time limit. The output inherits the essential statistical features of the input, shaped by the filter's frequency response. In particular, when $\mathbf{W}_t$ is a stationary Gaussian process,¹ $\mathbf{J}_t$ remains stationary and Gaussian. By multiplying Eq. (38) by its conjugate transpose, one can get a connection between their asymptotic multivariate periodograms,

$$S_J(\omega) = |\hat{H}_D(i\omega)|^2 S_W(\omega). \quad (39)$$

Evaluating $|\hat{H}_D(i\omega)|^2$ and taking the white noise input $\mathbb{E}[S_W(\omega)] = Q\mathbb{1}_2$, one recovers the well-known Lorentzian line shape centered at the Larmor frequency:

$$\mathbb{E}[S_J(\omega)] = \frac{Q/2}{\gamma^2 + (\omega - \omega_L)^2} + \frac{Q/2}{\gamma^2 + (\omega + \omega_L)^2}, \quad (40)$$

defined over $\omega \in (-\infty, \infty)$. Note that, since the signal is real valued, the Fourier components at positive and negative frequencies are always complex conjugates of each other. In addition, in the regime where $\omega_L \gg \gamma$, the Lorentzian peak centered at $-\omega_L$ contributes negligibly to the spectrum in the positive-frequency domain. It is therefore sufficient to consider only the positive-frequency branch:

$$\mathbb{E}[S_J(\omega)] \approx \frac{Q/2}{\gamma^2 + (\omega - \omega_L)^2}, \quad \omega \in (0, \infty). \quad (41)$$

To obtain the output signal from Eq. (24), we need to take into account that J_z enters in the photocurrent description sampled at periods Δ and therefore an analogous analysis in the discrete case must be carried out. In this scenario, the equivalence of the s -domain representation is the so-called z domain. The discrete version of the transfer function for our system is

$$\hat{H}_D^d(z) = (\mathbb{1}_2 - z^{-1}(\mathbb{1}_2 + \hat{A}\Delta))^{-1}, \quad (42)$$

and the frequency response, i.e., restricting $z = e^{-i\omega\Delta}$,

$$\hat{H}_D^d(e^{i\omega\Delta}) = (\mathbb{1}_2 - e^{-i\omega\Delta}(\mathbb{1}_2 + \hat{A}\Delta))^{-1}, \quad (43)$$

¹For continuous-time stochastic processes (e.g., Wiener noise appearing here) that are non-square-integrable, the Fourier transform $\mathcal{F}(X_t)$ must be understood as a finite-time, normalized limit: $\mathcal{F}(J_t)(\omega) := \lim_{T \rightarrow \infty} \frac{1}{\sqrt{T}} \int_0^T J_t e^{-i\omega t} dt$, which ensures the transform is well defined and properly captures the long-time spectral content of the process [72].

with $\omega \in (-\frac{\pi}{\Delta}, \frac{\pi}{\Delta})$. Notice that now the frequency resolution is bounded by the sampling frequency $\frac{1}{\Delta}$. Finally, for a small Δ , one can relate the continuous and discrete versions:

$$|\hat{H}_D^d(e^{i\omega\Delta})|^2 = \frac{|\hat{H}_D(i\omega)|^2}{\Delta} + O(1). \quad (44)$$

Equipped with these results, and noting that the detected signal, I in Eq. (24), is linearly coupled to $\mathbf{J}$ plus some added independent background noise, we can derive its PSD in the limit of high sampling rate, i.e., small Δ , and recover the Lorentzian profile found in Eq. (31).

Following the structure of the previous two subsections, we now present the asymptotic probability distribution of the signal in the continuum, long-time limit.

As discussed above, the statistical properties of the output signal I_t are inherited from those of the input white noises. In particular, stationarity implies that for any time shift τ , the process $I_{t+\tau}$ has the same statistics as I_t . Consequently, the Fourier components $\mathcal{F}_I(\omega)$ and $e^{-i\omega\tau}\mathcal{F}_I(\omega)$ must be identically distributed, which implies that the phase of $\mathcal{F}_I(\omega)$ is uniformly distributed over $[0, 2\pi)$. Since $\mathcal{F}_I(\omega)$ is a linear functional of a Gaussian process, it is itself a complex Gaussian variable. The uniformity of the phase then ensures that the real and imaginary parts of $\mathcal{F}_I(\omega)$ are independent, identically distributed Gaussian variables.

Moreover, stationarity also implies that the Fourier components $\mathcal{F}_I(\omega)$ and $\mathcal{F}_I(\omega')$ are uncorrelated for $\omega \neq \omega'$ (excluding the case $\omega' = -\omega$, which is perfectly correlated for real-valued signals, as noted above), and since they are jointly Gaussian, they are statistically independent. Therefore, in the frequency domain, $\mathcal{F}_I(\omega)$ consists of independent complex Gaussian variables across frequencies, each with independent real and imaginary parts.

As a result, the power spectral density $S_I(\omega) = |\mathcal{F}_I(\omega)|^2$ is the sum of the squares of two independent, identically distributed Gaussian variables; i.e., it follows a χ^2 distribution with two degrees of freedom or, equivalently, an exponential distribution:

$$p(S_I(\omega)) = \frac{1}{\bar{S}_I(\omega)} e^{-S_I(\omega)/\bar{S}_I(\omega)}, \quad (45)$$

where $\bar{S}_I(\omega) = \mathbb{E}[S_I(\omega)]$ is the mean power spectrum at frequency ω from Eq. (31) (see also Refs. [53,54]).

The joint probability distribution of the signal across all frequencies factorizes into a product of exponentials given by Eq. (45) at each frequency. In the continuum limit, this leads to the following expression for the asymptotic log-likelihood ratio:

$$L_{S_I} = \frac{t}{2\pi} \int_0^\infty \ln \left(\frac{\bar{S}_{I,0}(\omega)}{\bar{S}_{I,1}(\omega)} \right) d\omega + \frac{t}{2\pi} \int_0^\infty \left(\frac{1}{\bar{S}_{I,0}(\omega)} - \frac{1}{\bar{S}_{I,1}(\omega)} \right) S_I(\omega) d\omega. \quad (46)$$

That is, the LLR is given by a weighted sum of independent exponential random variables and yields cumulants matching the leading-order terms in Eq. (30), as expected in the large- n limit. Indeed, these results are fully consistent with the asymptotic diagonalization of Toeplitz covariance matrices discussed in Sec. III B, where the Fourier basis asymptoti-




FIG. 2. Representative time- and frequency-domain traces from two datasets, corresponding to hypotheses differing in the magnitude of the magnetic field, i.e., differing in our model [Eqs. (23) and (24)], only through the Larmor frequencies, labeled $\omega_{L,0}$ and $\omega_{L,1}$. The values of ω_L and other model parameters are estimated in Sec. V A. Top: example time traces before (left) and after (right) high-pass filtering. Bottom: PSDs averaged over 200×2 s traces, before (green, orange) and after (red, blue) filtering. The frequency response of the filter is indicated on the right-hand y axis, and the inset (same units as main) provides a close-up into 40–60 kHz near the Larmor peaks.

cally diagonalizes such matrices and the PSDs appear as their eigenvalue spectra [73,74]. There exists a well-established body of results (see Sec. IV B), which provides precise asymptotic expressions for functionals of Toeplitz matrices together with rigorous results on their convergence rates. This framework will allow us to systematically connect the infinite-size limit with the corrections that arise at finite n .

IV. EXPERIMENTAL DATA, HIGH-PASS FILTER, AND LOG-LIKELIHOOD RATIO

A. Experimental data and notch filter

The detection strategies considered in this work rely on the assumption that the observed data align with the underlying theoretical model. However, the current physical model does not account for low-frequency spurious noise, which is particularly prominent near zero frequency and is ubiquitous in laboratory environments. Figure 2 illustrates examples of random traces corresponding to two different high Larmor frequencies, where the presence of such low-frequency noise can be clearly observed.

A complete characterization of this spurious noise that could be incorporated into our model is futile, due to its diverse and unpredictable origins. Instead, we address the issue by applying a linear time-invariant causal filter to the observations. The filter must satisfy three essential properties:

first, it must be implementable in real time; second, it must effectively suppress the unwanted noise; and third, its impact on the relevant information in the data, specifically around the Lorentzian peak, must be minimal. This last point will be verified by quantifying the influence of the filter on our primary statistic, L_I , as will be shown in the next subsection.

For this purpose, we apply to the signal $\mathbf{I}_n$ the following infinite impulse response high-pass filter, designed to eliminate the zero-frequency component and attenuate low-frequency noise. Denoting the output by $\mathbf{R}_n = (R_1, \dots, R_n)^T$, the resulting signal is

$$R_k = \alpha R_{k-1} + \sqrt{\alpha}(I_k - I_{k-1}), \quad (47)$$

where the parameter $0 < \alpha < 1$ regulates the strength of the filter.

Figure 2 shows the result of applying the high-pass filter to the time traces. The filtered signals are now centered around zero, with low-frequency components effectively suppressed, yielding data that better match the model predictions of Eqs. (23) and (24).

To understand the impact of the high-pass filter, we start by noting that the transformation from $\mathbf{I}_n$ to $\mathbf{R}_n$ is affine, and therefore the Gaussianity of the observations is preserved. Now, $\mathbf{R}_n \sim \mathcal{N}(0, \hat{\Sigma}_n)$, with $\mathbf{R}_n = \hat{\mathbf{B}}_n \mathbf{I}_n$ and $\hat{\Sigma}_n = \hat{\mathbf{B}}_n \hat{\mathbf{T}}_n \hat{\mathbf{B}}_n^T$, where $\hat{\mathbf{T}}_n$ and $\hat{\mathbf{B}}_n$ are $n \times n$ Toeplitz matrices with matrix elements, respectively, given by Eq. (25) and $B_{i,j} = \sqrt{\alpha}(\alpha - 1)\alpha^{|i-j|} + \sqrt{\alpha}(2 - \alpha)\delta_{i,j}$. After the filter stabilizes, $\hat{\Sigma}_n$ also becomes a Toeplitz matrix with diagonal entries Σ_k that for high Larmor frequencies, i.e., far from spurious noises,

$$\Sigma_k \approx \left(1 - \frac{b^2}{\omega_L^2}\right) T_k - b \left(\frac{S_{\text{ph}}}{2} + \frac{g_D^2 Q \gamma}{2\omega_L^2}\right) e^{-b|k|\Delta} + \mathcal{O}\left(\frac{1}{\omega_L^3}\right), \quad (48)$$

where b is implicitly defined via $\alpha = e^{-b\Delta}$ to facilitate taking the continuous-time limit. Notice that Σ_k is essentially proportional to T_k —the correlations that will give rise to the Lorentzian centered at ω_L —since the term b^2/ω_L^2 is small in this regime. The other term in Eq. (48) is a contribution at zero frequency that will be negligible by the functional in the LLR expression.

In the frequency domain, the asymptotic transformation from $S_I(\omega)$ to $S_R(\omega)$ is given by the transfer function $H_R^d(z) = \sqrt{\alpha} \frac{1-z^{-1}}{1-\alpha z^{-1}}$ evaluated at $z = e^{-i\omega\Delta}$. Again, taking into account $\alpha = e^{-b\Delta}$, the square modulus of the frequency response satisfies

$$\begin{aligned} |H_R^d(e^{i\omega\Delta})|^2 &= 2\alpha \frac{1 - \cos(\omega\Delta)}{1 + \alpha^2 - 2\alpha \cos(\omega\Delta)} \\ &= \frac{\omega^2}{b^2 + \omega^2} + \mathcal{O}(\Delta), \quad \omega \in \left(-\frac{\pi}{\Delta}, \frac{\pi}{\Delta}\right), \end{aligned} \quad (49)$$

where $|H_R(i\omega)|^2 = \frac{\omega^2}{b^2 + \omega^2}$ is the continuous version. Now, one can see how this filter has the effect of eliminating the zero-frequency component, reducing the low-frequency noise, and at the same time leaving the high-frequency components largely untouched. Other filters, such as the moving average filter or the Kalman filter, could reduce zero-frequency spurious noise, but it would not have the desired effect of




FIG. 3. Block diagram of the process represented with the transfer matrices.

annihilating it. Indeed, for large frequencies,

$$S_R(\omega) = \frac{\omega^2}{b^2 + \omega^2} S_I(\omega) = \left(1 - \frac{b^2}{\omega^2}\right) S_I(\omega) + \mathcal{O}\left(\frac{1}{\omega^3}\right). \quad (50)$$

For high ω_L , the high-pass filter has negligible effect in the region of interest, i.e., near the Lorentzian peak. Figure 2 shows the PSD averaged over 200 experimental traces, with and without filtering. The spurious low-frequency noise exceeds the Lorentzian peak by nearly 10 orders of magnitude, justifying the use of the high-pass filter. Figure 3 presents a diagram of the complete process, interpreted as an effective filter.

B. Effect of the high-pass filter on the log-likelihood ratio

The previous section demonstrates, via frequency-domain analysis, that the filter preserves relevant signal information in the long-time limit while suppressing spurious low-frequency noise. However, the filtered time trace $\mathbf{R}_n = \hat{\mathbf{B}}_n \mathbf{I}_n$ differs substantially from the raw input $\mathbf{I}_n$ at every time step. Since the Kalman filter computes L_I recursively from filtered data, finite-time effects could in principle degrade performance or compromise the theoretical guarantees. We address this concern by quantifying how filtering modifies the Kalman filter output and verifying that these corrections remain negligible at timescales relevant to our protocol.

Since the log-likelihood ratio depends only on the observed data, we define

$$L_R(\mathbf{I}_n) := L_I(\mathbf{R}_n),$$

which has the same functional form as L_I given in Eq. (27), but evaluated on the filtered observations. Explicitly, this corresponds to

$$L_{R,n} = \frac{1}{2} \ln \frac{\det \hat{\mathbf{T}}_{0,n}}{\det \hat{\mathbf{T}}_{1,n}} + \frac{1}{2} \mathbf{I}_n^T \hat{\mathbf{B}}^T (\hat{\mathbf{T}}_{0,n}^{-1} - \hat{\mathbf{T}}_{1,n}^{-1}) \hat{\mathbf{B}} \mathbf{I}_n. \quad (51)$$

We quantify the error induced by the filter by

$$\epsilon_n := L_{I,n} - L_{R,n} = \frac{1}{2} \mathbf{I}_n^T \hat{\mathbf{M}} \mathbf{I}_n, \quad (52)$$

where we combined Eqs. (27) and (51) and defined $\hat{\mathbf{M}} := \hat{\mathbf{D}} - \hat{\mathbf{B}}^T \hat{\mathbf{D}} \hat{\mathbf{B}}$ with $\hat{\mathbf{D}} = \hat{\mathbf{T}}_{0,n}^{-1} - \hat{\mathbf{T}}_{1,n}^{-1}$.

We expect that the log-likelihood ratio is primarily dominated by spectral components near the Larmor frequencies [see Eq. (50)], the difference introduced by filtering should be small. Specifically, we anticipate that the error term $\epsilon_n \sim \frac{b^2}{\omega_c^2} L_{S_I}$, where $\omega_c := \frac{1}{2}(\omega_0 + \omega_1)$, is the high frequency around which both hypotheses $h = 0, 1$ are centered.

To verify that the effect of the low-frequency filtering can be safely neglected and that the finite-time corrections are

TABLE I. Results for the five estimated parameters. We present the inferred value, the analytical standard deviation σ obtained from the Cramér Rao bound, and the empirical MSE of the standard deviation simulating the experiments with a total number of 1000 batches.

Parameter	Unit	Estimated Value	σ	MSE
γ	Hz	330.90	0.90	0.98
$\omega_{L,0}$	Hz	50 114.03	0.97	0.98
$\omega_{L,1}$	Hz	50 550.88	0.97	1.02
S_{at}	$\mu\text{V}^2/\text{Hz}$	31.768	0.087	0.094
S_{ph}	$\mu\text{V}^2/\text{Hz}$	13.0457	0.0038	0.0040

small, we compute the mean and variance of ϵ_n . Note that Eq. (52) defines ϵ_n as a quadratic functional of the signal, and as in Sec. III B, a closed-form expression for the m th cumulant under hypothesis h can be written as

$$\tilde{\kappa}_{m,n}^h = \frac{(m-1)!}{2} \text{Tr}((\hat{M}\hat{T}_{h,n})^m) = \mathcal{O}(1) + \frac{t(m-1)!}{2\pi} \times \int_0^\infty \left(\frac{b^2}{\omega^2 + b^2} \left(\frac{\tilde{S}_{L,h}(\omega)}{\tilde{S}_{L,0}(\omega)} - \frac{\tilde{S}_{L,h}(\omega)}{\tilde{S}_{L,1}(\omega)} \right) \right)^m d\omega. \quad (53)$$

In particular, the mean and variance can be explicitly computed, yielding

$$\begin{aligned} \mathbb{E}_h[\epsilon_n] &= \frac{b^2}{\omega_c^2} \mathbb{E}_h[L_{S_I}] + \mathcal{O}\left(\frac{1}{\omega_c^3}\right) + \mathcal{O}(1), \\ \text{Var}_h(\epsilon_n) &= \frac{b^4}{\omega_c^4} \text{Var}_h(L_{S_I}) + \mathcal{O}\left(\frac{1}{\omega_c^5}\right) + \mathcal{O}(1), \end{aligned} \quad (54)$$

which formally shows that indeed the filter has a very limited effect on the LLR, and hence on the performance of our protocols.

Finally, to validate the procedure, we simulate 10^4 measurement trajectories using the parameters listed in Table I under hypothesis $h = 0$, with a sampling interval $\Delta = 5 \mu\text{s}$, matching our experimental conditions. The simulated data are processed identically to the experimental traces: we apply a high-pass filter with parameter $\alpha = 0.91$ and compute the LLR sequentially using the Kalman filter formalism described in Sec. III C, which is formally equivalent to Eq. (51). We then compare the resulting LLR to that obtained from the unfiltered data to compute the error ϵ_n .

Figure 4 shows the mean and variance of the time-normalized error ϵ_t/t alongside theoretical asymptotic predictions from Eq. (54). With $b^2/\omega_c^2 = 3.6 \times 10^{-3}$ in our setting, the filter's relative contribution is minimal. The inset displays representative simulated traces, demonstrating negligible filter effects even at short times.

V. SEQUENTIAL STRATEGIES ON THE SPIN-NOISE ATOMIC MAGNETOMETER

We are now in a position to demonstrate sequential strategies in a spin-noise quantum system. Specifically, we consider a scenario where all system parameters are fixed except for the Larmor frequency, ω_L . We set $\omega_{L,h} = \omega_c + (-1)^h \frac{1}{2} \Delta\omega$, with ω_c the frequency around which both hypotheses $h = 0, 1$




FIG. 4. Comparison between analytical and simulated results for the time-normalized error ϵ_t/t introduced by the high-pass filter. The mean (blue) and variance (red) are estimated from 10^4 simulated trajectories and are compared with the theoretical predictions. The asymptotic values are given in terms of the regularized mean and variance of L_{S_I} , i.e., of $\rho := \lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E}_0[L_{S_I}] = -0.207 \text{ ms}^{-1}$ and $v := \lim_{t \rightarrow \infty} \frac{1}{t} \text{Var}_0[L_{S_I}] = 0.477 \text{ ms}^{-1}$ [see Eq. (54)].

are centered. We have in mind applications such as detecting biomagnetic signals, magnetic anomalies in geophysical surveys, or concealed metallic objects, where a small magnetic field perturbs a baseline field, inducing a shift $\Delta\omega$ in ω_L . Details of $\Delta\omega$ values and the rest of parameters are given below.

While our setting assumes fully characterized models for both hypotheses, real devices typically require calibration before each use. This is done by operating the sensor under controlled conditions where only the baseline field is present. Therefore, before proceeding, we briefly outline how we estimate model parameters from finite datasets.

A. Calibrating the sensor: Estimation of model parameters

We employ the maximum-likelihood estimator (MLE) to infer the system parameters, building on the thorough analysis of the likelihood statistic's properties presented above. Indeed, for long observation times, the PSD components become independent and follow exponential distributions, making the MLE analytically tractable and statistically efficient. That is in the sense that it achieves the Cramér-Rao bound, ensuring optimal precision in parameter estimation given sufficiently large datasets [75].

For a given time trace $\mathbf{I}_n$, we estimate the parameters $\Theta_h = (S_{\text{ph}}, S_{\text{at}}, \omega_h, \gamma)^T$ via maximum likelihood. Recall that the DFT of $\mathbf{I}_n$ is unitary and preserves the full information content of the signal, that for stationary processes the Fourier phases carry no information, and that for sufficiently long traces the different frequency components become independent. As a result, given a time trace $\mathbf{I}_n$ or equivalently its periodogram $\mathbf{S}_I$, the MLE can be obtained by

maximizing

$$-\ln p(S_I|\Theta_h) = \sum_k \ln (\mathbb{E}_{\Theta_h}[S_I(\omega_k)]) + \frac{S_I(\omega_k)}{\mathbb{E}_{\Theta_h}[S_I(\omega_k)]}, \quad (55)$$

where k runs over the first $n/2$ frequencies, since the signal is real valued and the negative frequencies are fully determined by their positive counterparts.

For the expression of $\mathbb{E}_{\Theta_h}[S_I(\omega_k)]$, one has two options. Employ directly the asymptotic Lorentzian given in Eq. (31), evaluated at the discrete frequencies $\omega_k = \frac{2\pi k}{n\Delta}$, yielding the so-called Whittle log-likelihood [76], or use the expected expression of the periodogram, which is known as the debiased. The regime in which we will perform the estimation is long enough to neglect the finite size effects and therefore we will directly employ the Lorentzian approximation.

Notice that, due to the linearity of the log-likelihood in the periodogram components, $S_I(\omega_k)$, if we have N independent traces collected under identical conditions, the total log-likelihood takes the same form as Eq. (55), with $S_I(\omega_k)$ replaced by the ensemble average $\tilde{S}_I(\omega_k) = \frac{1}{N} \sum_{j=1}^N S_I^j(\omega_k)$.

In this work, we aim to illustrate the use of sequential methodologies and empirically verify their performance in terms of mean detection times and error probabilities. To this end, we generate multiple time traces under each hypothesis, forming two datasets that enable us to estimate these quantities and compare the performance of sequential and deterministic strategies.

We use these two labeled datasets, corresponding to hypotheses h_0 and h_1 , to infer the relevant system parameters for this batch of experiments. As the data are independently distributed and differ solely in the value of the Larmor frequency, we can perform maximum likelihood estimation (MLE) by optimizing the full joint likelihood over both datasets:

$$\sum_k \left(\ln (\mathbb{E}_{\Theta_0}[S_I(\omega_k)]) + \frac{\tilde{S}_{I,0}(\omega_k)}{\mathbb{E}_{\Theta_0}[S_I(\omega_k)]} + \ln (\mathbb{E}_{\Theta_1}[S_I(\omega_k)]) + \frac{\tilde{S}_{I,1}(\omega_k)}{\mathbb{E}_{\Theta_1}[S_I(\omega_k)]} \right). \quad (56)$$

Since the MLE is known to saturate the Cramér-Rao bound, we can assess the expected estimation errors by computing the inverse of the Fisher information matrix for our model.

Again, we validate this prediction by simulating the complete estimation procedure. Specifically, we generate synthetic datasets using the parameter values Θ listed in Table I, matching the number of time traces used in the experimental dataset. This process is repeated over $N_B = 1000$ independent batches, yielding an estimate $\hat{\Theta}_i$ for each batch. From these, we compute the empirical covariance matrix of the estimator,

$$\hat{\Gamma} = \frac{1}{N_B} \sum_{i=1}^{N_B} (\hat{\Theta}_i - \Theta)(\hat{\Theta}_i - \Theta)^T. \quad (57)$$

The simulated MSE values reported in the last column of Table I correspond to the diagonal entries of this matrix and

show good agreement with the theoretical values predicted by the Cramér-Rao bound.

For the estimation, we cut the region of frequencies where the Lorentzian is centered. The Larmor frequencies of both hypotheses are near 50 kHz. Therefore, we consider the region of frequencies between 40 and 60 kHz. The experimental datasets for each hypothesis contain 200 traces of 2 s. The results are detailed in Table I.

Notice that the error in the estimation of all parameters is at least 2 orders of magnitude smaller than the separation between the two hypotheses, since the difference between the two Larmor frequencies is $\Delta\omega = \omega_{L,0} - \omega_{L,1} = 436.85$ Hz. Therefore, we are in a regime where the inferred parameters can be reliably used for the implementation of the sequential strategies. While the reported errors correspond to theoretical estimates, either from the Cramér-Rao bound or from Monte Carlo simulations, we will later show that our experimental datasets exhibit parameter fluctuations beyond these estimates, particularly in the parameters that govern the relaxation processes. In Sec. VB and Appendix E, we provide a qualitative model that explains these fluctuations and quantify its effects.

B. Implementation of sequential strategies

For real-time implementation, both the high-pass and Kalman filters can be efficiently implemented on microcontrollers or field-programmable gate arrays (FPGAs) [77–80] capable of real-time processing, using simple arithmetic operations (see the sequential block diagrams in Appendix D). At each time step n , the device computes the filtered signal R_n , evaluates the LLR, and applies the SPRT or CUSUM thresholds for sequential hypothesis testing or quickest change-point detection, respectively. In this work, we implement these strategies on a prerecorded dataset. This approach allows us to assess the performance of the protocols, such as average detection times and error probabilities, by repeating the analysis over many independent realizations under controlled conditions.

In the hypothesis-testing scenario, we segment the recorded trajectories into 10^4 independent traces of duration 40 ms for each hypothesis in order to increase statistical sampling. The high-pass filter is applied with strength parameter $\alpha = 0.91$, and the acquisition sampling interval is $\Delta = 5$ μ s.

In Fig. 5, we compare the average time required by the sequential probability ratio test to reach a predefined error level with that of the optimal deterministic strategy achieving the same error. To ensure a fair comparison, we consider a symmetric hypothesis-testing scenario (see Sec. II A), where both types of errors are weighted equally, and define the average error as $\epsilon = \frac{1}{2}\epsilon_0 + \frac{1}{2}\epsilon_1$. For the deterministic strategy, we use time traces of fixed length t and apply the optimal decision rule described in Sec. II A 1: decide $h = 1$ if $L_t > 0$, and $h = 0$ otherwise. The empirical error probability is then estimated as the fraction of misidentified traces. For the SPRT, we set equal error thresholds $\epsilon_0 = \epsilon_1 = \epsilon$. According to Eq. (8), this determines the upper and lower thresholds for the SPRT: sampling stops at the first random time t when $L_{R,t} \notin (-a_0, a_1)$, followed by a decision: $h = 1$ if $L_t \geq a_1$ and $h = 0$ if $L_{R,t} \leq -a_0$, with $a_0 = a_1 = a$. The empirical error is computed as




FIG. 5. Relation between error probability and required experiment duration for sequential (dotted) and deterministic (solid) strategies. Each point is obtained by averaging over 10^4 experimental traces per hypothesis. The horizontal axis shows $\ln \epsilon$, where ϵ is the empirical error probability (i.e., fraction of incorrectly identified hypotheses), and the vertical axis shows the required experiment duration. For the sequential strategy (SPRT), we vary the stopping threshold $a = a_0 = a_1$, compute the error probability and the average stopping time over all runs, and plot $(\ln \epsilon, t)$. For the deterministic strategy, we fix the experiment duration t to various values, and for each value we compute the average error probability ϵ over all runs, and plot $(\ln \epsilon, t)$. Experimental results (blue) are plotted alongside simulated results (red) for comparison.

before, although it is worth emphasizing that, by construction, the SPRT guarantees a conditional error probability ϵ for every trajectory.

To validate our model, results are shown for both experimental data (blue) and simulated data (red), generated under identical conditions with matched parameters and the same high-pass filter. In both datasets, the SPRT consistently reaches the target error level in less time than the deterministic strategy, while also providing rigorous error guarantees. It is worth noting that, while sequential strategies are clearly more time and resource efficient on average, they come with a trade-off: the duration of each experiment is inherently random. Although its fluctuations are typically bounded within a characteristic range, the distribution can exhibit long tails, making the exact stopping time unpredictable [42]. Indeed, in our experimental implementation, we limited traces to 40 ms to improve statistics; as a result, a tiny fraction of trajectories (55 for $h = 0$ and 53 for $h = 1$ out of 10^4 each) did not reach a stopping decision. This negligible portion of the dataset was excluded from the averages shown in Fig. 5.

When comparing experimental and simulated results in Fig. 5, we observe good agreement at moderate measurement times. As the acquisition time increases, discrepancies appear, with experimental data requiring longer times to reach the same estimation error. This reflects that, at long times, the information gathered about the underlying process becomes increasingly precise, and small system imperfections, such as parameter drifts, stability issues, or unmodeled dynamics, become more significant. In Appendix E, we model the effect of small run-to-run variations in the system parameters,




FIG. 6. Quickest change-point detection applied to experimental traces with an abrupt magnetic-field shift. Each trajectory starts under hypothesis h_0 and undergoes a programmed magnetic-field change of magnitude $\Delta B = \Delta\omega/\gamma_G \approx 62.3$ nT at $t = 1$ s (indicated by a vertical line), where $\gamma_G \approx 7$ Hz/nT is the gyromagnetic ratio. For visualization purposes, the detection threshold is set to the large value $a = 50$. Circles mark the detection times produced by the algorithm for each run, while squares indicate the corresponding maximum-likelihood estimates of the actual change-point times. Note that the programmed change point does not coincide exactly with the true physical change point, nor with the estimated change points, as discussed in the main text.

particularly those governing spin relaxation, which qualitatively explains the statistics of the LLR.

To rigorously address scenarios with significant parameter fluctuations, our techniques must be extended to the framework of composite hypothesis testing [31,72], where each hypothesis h corresponds to a family of possible models defined by parameters in a given set, $\theta_h \in \mathcal{S}_h$. In this setting, different forms of the LLR are employed depending on whether the parameters are partially or completely unknown, and the stopping rules are adapted to select the most likely hypothesis.

Finally, we demonstrate the implementation of the QCPD algorithm using 2-s measurement traces that initially follow hypothesis h_0 and exhibit a sudden transition to hypothesis h_1 near $t = 1$ s. For the calibration of the detection model, we employ the parameters previously estimated from the hypothesis-testing datasets, as the measurements are acquired under identical conditions, including the application of the same high-pass filter. Representative results for different trajectories are presented in Fig. 6.

Notice the marked change in the behavior of the LLR after the change point. We use the CUSUM stopping rule defined in Eq. (18) with a fixed threshold $a = 50$.

Colored circles along the LLR traces indicate the corresponding stopping times, i.e., the points at which the decision criterion in Eq. (17) triggers a change-point detection alarm. In addition, as discussed in Sec. II B, the minimum of the LLR curve provides the maximum-likelihood estimate of the actual change-point time. Across all runs, this minimum (marked with colored squares) is consistently shifted with respect to the vertical line indicating the nominally programmed change.




FIG. 7. Same as Fig. 6, but for a different set of hypotheses with parameters $\gamma = 418.48$ Hz, $\omega_{L,0} = 42\,577.43$ Hz, $\omega_{L,1} = 43\,640.10$ Hz, $S_{\text{at}} = 0.770\,90\,\mu\text{V}^2/\text{Hz}$, and $S_{\text{ph}} = 1.4277\,\mu\text{V}^2/\text{Hz}$. Unlike Fig. 6, the magnetic-field sequence and data acquisition are synchronized here, with $t = 1$ s marking the programmed magnetic-field change. Circles mark detection times and squares indicate the corresponding maximum-likelihood estimates, which closely match the programmed change point.

This systematic offset revealed the presence of an unaccounted technical delay caused by a timing mismatch between the magnetic-field sequence and the data acquisition, resulting in a constant delay of approximately 60 ms.

To confirm this diagnosis, we performed a second batch of experiments (Fig. 7) in which the magnetic-field changes and data acquisition were synchronized by recording the coil current and atomic response on two channels of the same acquisition system. The estimated change points in this configuration show excellent agreement with the programmed times, confirming the source of the initial offset. This example demonstrates the robustness of sequential analysis: our algorithm correctly identified the change in both cases, revealing even subtle instrumental effects.

This new set of hypotheses is less distinguishable than the former set, as indicated by a smaller regularized Kullback-Leibler divergence in Eq. (66). This is also evident in Fig. 7, which shows that achieving the same detection threshold (false alarm rate) requires a detection time 2–3 times longer, consistent with the theoretical factor of 2.65 given by the ratio of the slopes in Eq. (66).

The threshold $a = 50$ considered so far was chosen to be rather large to clearly visualize the detection delay in Fig. 6 and corresponds to an astronomical false alarm time of $\sim 10^9$ years. We now turn to a more realistic scenario and consider a more modest bound on mean time of false alarm of $T_{\text{FA}} = 0.5$ s, which corresponds to a decision threshold of $a = \ln(T_{\text{FA}}/\Delta) \approx 11.5$, and a much quicker detection. We now use a widely adopted version of the CUSUM algorithm: sequentially compute

$$L_n^+ := \sum_{i=1}^n \Delta L_i^+, \quad \text{with} \quad \Delta L_i^+ = \max(0, \Delta L_i), \quad (58)$$




FIG. 8. Results from 200 experimental runs starting at $t_0 = 0.25$ s with a fixed initial magnetic field. A step change of $\Delta B = 62.3$ nT is programmed at $t = 1$ s (blue vertical line). The plot shows the four trajectories that triggered a false alarm, i.e., a detection before the programmed change. Triangles indicate the false alarm times T_{FA} , as determined by the algorithm with a threshold of $a = 11.5$ (black horizontal line) corresponding to $T_{\text{FA}} = 0.5$ s.

and signal a detection alarm as soon as $L_n^+ \geq a$. This algorithm is equivalent to the procedure defined in Eq. (17), but accumulates only positive log-likelihood contributions, effectively resetting the reference level to zero whenever the instantaneous evidence for change becomes negative.

Figure 8 shows the application of this procedure to a dataset of 200 experimental trajectories, each lasting 2 s. We show the four traces in which a false alarm was triggered; that is, a change was detected before the programmed nominal change point at $t = 1$ s. The sequential analysis begins at $t_0 = 0.25$ s. In each instance of a false alarm, instead of halting the measurement, we reset the statistic $S_n = 0$ and resumed data acquisition, effectively restarting the detection process from an uninformed state.

Finally, for trajectories that correctly detected the actual change, the average detection delay was found to be approximately 60 ms, in agreement with the theoretical value $\bar{\tau} = 56$ ms predicted from Eq. (69) using the mean log-likelihood ratio from Eq. (30).

VI. OPTIMAL ASYMPTOTIC LIMITS

In this section, we provide closed-form expressions for key information-theoretic quantities that determine the asymptotic performance bounds at long observation times, both for deterministic and sequential hypothesis testing, as well as for QCPD.

In the *standard deterministic strategy* where the observation time is fixed in advance, for a sequence of n observations $\mathbf{I}_n = (I_1, \dots, I_n)^T$, the probability of error is given by

$$P_{\text{err}}(n) = \int_{\mathbb{R}^n} \min\{\pi_0 p_0(\mathbf{I}_n), \pi_1 p_1(\mathbf{I}_n)\} d\mathbf{I}_n. \quad (59)$$

Evaluating Eq. (59) is generally intractable. However, analogously to the i.i.d. case [see Eq. (5)], one can derive asymptotic bounds in the long-time limit.




FIG. 9. Asymptotic performance metrics for the sequential probability ratio test (dotted) and the deterministic probability ratio test (solid), obtained by averaging over 4×10^4 simulated trajectories per hypothesis. The horizontal axis shows $\ln \epsilon$, with ϵ the error probability, and the vertical axis shows the experiment duration. These quantities are computed as explained in Fig. 5. Asymptotic theoretical rates are also shown: The Chernoff information in Eq. (61) (solid black) and the regularized Kullback-Leibler information in Eq. (66) (dotted black) give asymptotic linear scaling of time with $-\ln \epsilon$, with slopes 0.2055 and 0.0496 ms^{-1} , respectively. This corresponds to more than a fourfold efficiency gain for the SPRT, $\rho_{\text{seq}}/\rho_{\text{det}} \simeq 4.14$.

By applying the inequality $\min(a, b) \leq a^s b^{1-s}$ for all $s \in (0, 1)$ to Eq. (59), one obtains the Chernoff bound:

$$P_{\text{err}}(n) \leq \min_{s \in (0,1)} \pi_0^s \pi_1^{1-s} \int_{\mathbb{R}^n} p_0(\mathbf{I}_n)^s p_1(\mathbf{I}_n)^{1-s} d\mathbf{I}_n. \quad (60)$$

This yields a lower bound on the error exponent:

$$-\lim_{t \rightarrow \infty} \frac{1}{t} \ln P_{\text{err}}(t) \geq \rho_{\text{det}} := \lim_{t \rightarrow \infty} \frac{1}{t} \max_s C(s),$$

with

$$C(s) = -\ln \int_{\mathbb{R}^n} p_0(\mathbf{I}_n)^s p_1(\mathbf{I}_n)^{1-s} d\mathbf{I}_n. \quad (61)$$

As in the IID case, this bound is known to be attainable in the asymptotic limit for a wide range of stochastic processes [72] (see also Fig. 9), in which case the regularized Chernoff information provides a lower bound on the asymptotic rate at which the error probability decays,

$$-\ln \epsilon_{\text{det}} \sim \rho_{\text{det}} t. \quad (62)$$

In the continuum limit, the Chernoff exponent function in Eq. (61) has a closed form in terms of the Kalman estimators, introduced in Sec. III C [72]:

$$C(s) = \frac{1}{2G} \int_0^t ((1-s)\sigma_1(t') + s\sigma_0(t') - \sigma_c(t', s)) dt', \quad (63)$$

where G is the variance of the shot noise and σ_h is the continuous version of the innovation variance under hypothesis h . Moreover, $\sigma_c(s, t')$ is understood as the variance of the innovations associated with the artificial signal $I_c = \sqrt{s}I_0 + \sqrt{1-s}I_1$, where I_h for $h = 0, 1$ is the signal under hypothesis h . Equation (63) provides a useful bound of P_{err} valid for any

observation time and predicts a linear scaling of $C(s)$ with t , once the variances reach their stationary values.

Alternatively, the Chernoff function in Eq. (63) admits an equivalent representation in the frequency domain. Here, the signal is characterized by its periodogram $S_{n/2} = (S_1, \dots, S_{n/2})^T$, where each S_j follows an independent exponential distribution. This leads to

$$\begin{aligned} C(s) &= \prod_{j=0}^{n/2} \int_0^\infty p_0(S_j)^s p_1(S_j)^{1-s} dS_j \\ &= \prod_{j=0}^{n/2} \frac{\bar{S}_{I,0}(\omega_j)^{1-s} \bar{S}_{I,1}(\omega_j)^s}{(1-s)\bar{S}_{I,0}(\omega_j) + s\bar{S}_{I,1}(\omega_j)}, \end{aligned} \quad (64)$$

which, in the continuum limit, becomes [60]

$$C(s) = \frac{t}{2\pi} \int_0^\infty \ln \frac{\bar{S}_{I,0}(\omega)^{1-s} \bar{S}_{I,1}(\omega)^s}{(1-s)\bar{S}_{I,0}(\omega) + s\bar{S}_{I,1}(\omega)} d\omega. \quad (65)$$

The performance metric for sequential hypothesis testing, and QCPD, is given by the slope of the mean LLR:

$$\rho_{\text{seq}} := \lim_{t \rightarrow \infty} \frac{|\mathbb{E}[L_t]_h|}{t}. \quad (66)$$

This quantity corresponds to the time-regularized Kullback-Leibler divergence and determines the exponential rate of error decay and the average stopping time in sequential strategies. In the time domain, using Kalman filter estimators the expected LLR is given by [42]

$$\begin{aligned} \mathbb{E}_h[L_t] &= \int_0^t \frac{1}{S_{\text{ph}}} \mathbb{E}_h[dI'_t(\mu_1(t') - \mu_0(t'))] \\ &\quad + \int_0^t \frac{dt'}{2S_{\text{ph}}} \mathbb{E}_h[\mu_0(t')^2 - \mu_1(t')^2], \end{aligned} \quad (67)$$

while in frequency domain, we recall Eq. (46),

$$\begin{aligned} \mathbb{E}_h[L_t] &= \frac{t}{2\pi} \int_0^\infty \ln \frac{\bar{S}_{I,0}(\omega)}{\bar{S}_{I,1}(\omega)} d\omega \\ &\quad + \frac{t}{2\pi} \int_0^\infty \left(\frac{\bar{S}_{I,h}(\omega)}{\bar{S}_{I,0}(\omega)} - \frac{\bar{S}_{I,h}(\omega)}{\bar{S}_{I,1}(\omega)} \right) d\omega + \mathcal{O}(1), \end{aligned} \quad (68)$$

which allows for an analytical computation of ρ_{seq} . Note that the choice of hypothesis $h = 0, 1$ only affects the sign of the expected LLR, while its magnitude, and thus ρ_{seq} , remains the same.

The mean detection time in sequential tests then scales as

$$\bar{\tau} \sim \frac{a}{\rho_{\text{seq}}}, \quad (69)$$

where $a = -\ln \alpha_h$ for a Wald test under hypothesis h and $a = \ln(T_{\text{FA}}/\Delta)$ in the case of QCPD with false alarm constraint T_{FA} .

In Fig. 9, we present the results for sequential (dotted) and deterministic (solid) hypothesis testing. To estimate the performance in the asymptotic regime of long observation times, corresponding to small error probabilities, a large number of trials is required. For this reason, we rely on simulated data. Analytical asymptotic error rates are also plotted for




FIG. 10. Mean detection time (blue, 5×10^3 traces left axis) and mean false alarm (orange, 10^3 traces, right axis) for simulated data. The red line shows the asymptotic detection time rate, $\rho_{\text{seq}}^{-1} = 4.88$ ms from Eq. (69), and the green line indicates the expected linear relation between threshold h and $\ln(T_{\text{FA}}/\Delta)$, with slope 1. Both slopes match the simulation results.

reference. Notably, the numerical results approach the theoretical Chernoff bound, confirming its attainability in this setting. Furthermore, we note that the sequential strategy outperforms the deterministic one, exhibiting more than a fourfold improvement in time required to achieve a given error probability.

Similarly, Fig. 10 illustrates the asymptotic results for the quickest change-point-detection problem. In order to compute these asymptotic results, we rely again on simulated data. Notice how the simulated results are in agreement with both slopes validating numerically the theoretical predictions.

In Appendix F, we show that, in our setting where the two models differ only in the Larmor frequency $\omega_{L,h} = \omega_c + (-1)^h \frac{\Delta\omega}{2}$ and large central frequency ω_c , assuming the Chernoff bound is attainable, the sequential strategies outperform deterministic ones by at least a factor of 4.

To conclude this section, we show that when the two hypotheses become arbitrarily close, regardless of the specific way in which they differ, the Chernoff and Kullback-Leibler information, which determine the deterministic and sequential performances, respectively, differ by a factor of 4.

In order to see this, we can, without loss of generality, choose a parametrization such that the two hypotheses differ only in a single parameter. Specifically, we take $\theta = \theta_0$ under h_0 and $\theta_1 = \theta_0 + \delta\theta$ under h_1 . Then, $\bar{S}_{I,\theta_0}(\omega) := \bar{S}_{I,0}(\omega)$ and $\bar{S}_{I,1}(\omega) = \bar{S}_{I,0}(\omega) + \delta\theta \frac{\partial \bar{S}_{I,0}(\omega)}{\partial \theta}|_{\theta=\theta_0} + \mathcal{O}(\delta\theta^2)$.

Under these assumptions, and using Eq. (65), we readily obtain

$$\begin{aligned} \rho_{\text{det}} &= \min_{s \in (0,1)} \frac{s(s-1)\delta\theta^2}{4\pi} \int_0^\infty d\omega (\partial_\theta \ln \bar{S}_{I,\theta}(\omega)|_{\theta=\theta_0})^2 + \mathcal{O}(\delta\theta^3) \\ &= \frac{\delta\theta^2}{16\pi} \int_0^\infty d\omega (\partial_\theta \ln \bar{S}_{I,0}(\omega)|_{\theta=\theta_0})^2 + \mathcal{O}(\delta\theta^3) \\ &= \frac{\delta\theta^2}{8} \mathcal{F}(\theta_0) + \mathcal{O}(\delta\theta^3), \end{aligned} \quad (70)$$

where in the last equality we have identified the regularized Fisher information of the process with respect to parameter θ :

$$\begin{aligned} \mathcal{F}(\theta) &:= \lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E}_\theta [(\partial_\theta \ln p_\theta(\mathbf{I}_t))^2] \\ &= \frac{1}{2\pi} \int_0^\infty d\omega (\partial_\theta \ln \bar{S}_{I,\theta}(\omega))^2. \end{aligned} \quad (71)$$

Similarly, for nearing hypothesis, the regularized relative entropy becomes

$$\rho_{\text{seq}} = \frac{\delta\theta^2}{2} \mathcal{F}(\theta_0) + \mathcal{O}(\delta\theta^3), \quad (72)$$

which, by direct comparison with Eq. (70), implies that $\rho_{\text{seq}} = 4\rho_{\text{det}}$. That is, sequential strategies require up to 4 times less observation time to achieve the same error performance as deterministic ones.

VII. DISCUSSION AND OUTLOOK

In this work, we have presented what is, to our knowledge, the first experimental implementation of sequential data analysis methods in continuously monitored quantum sensors. We realize two cornerstone primitives, namely, sequential hypothesis testing and quickest change-point detection, in an atomic spin-noise magnetometer. This bridges a gap between quantum sensing theory and realistic, real-time sequential protocols. Our methods enable adaptive decision-making, allowing detection of weak magnetic perturbations with minimal delay and a quantifiable confidence threshold. This represents progress toward making quantum sensors more autonomous, efficient in data use, and compatible with real-world constraints.

We have presented the first comprehensive demonstration that sequential decision protocols can be applied to data from continuously monitored quantum sensors, further exploiting their intrinsic sensitivity. Although decisions were computed postacquisition, all components of the analysis were designed with real-time feasibility in mind. Using a linear Gaussian state-space model of spin-noise magnetometry and processing measurement records with a hybrid continuous-discrete Kalman filter, we evaluated the log-likelihood ratio via the Sequential Probability Ratio Test for binary discrimination and the CUSUM algorithm for change-point detection.

Beyond validating the feasibility of these strategies in a state-of-the-art experimental platform, our results open several promising directions for future research. From a methodological standpoint, the sequential framework can be extended to tackle multiple and composite hypothesis testing, signal segmentation (involving the detection of multiple abrupt changes), and anomaly detection, critical in contexts such as rare signal searches. Introducing adaptive feedback, open and closed loop and in both classical and quantum forms, and quantum-enhanced resources, such as spin-squeezed states, could further enhance detection efficiency or sensitivity. This synergy between quantum sensing and sequential decision theory has yet to be fully explored and could unlock new performance frontiers.

These results confirm that sequential decision-making protocols can be efficiently integrated with continuously monitored quantum systems. While we focused on atomic

spin-noise magnetometry, the methodology, combining state-space modeling, real-time filtering, and sequential decision rules, is broadly applicable across other quantum sensing platforms such as optomechanical resonators [10], superconducting circuits [28], NV centers in diamond [81], and cavity-based architectures [82], provided suitable measurement models exist.

Finally, an important open direction lies in model-free approaches, where the log-likelihood ratio is inferred directly from measurement data using machine-learning techniques. Such methods could bypass the need for full model knowledge, making sequential analysis more flexible and robust in complex, real-world environments.

ACKNOWLEDGMENTS

We are grateful to Charikleia Troullinou for insightful discussions during the early stages of this project. This work was supported by the QuantERA grant C'MON-QSENS! and Spanish Agencia Estatal de Investigación, Project No. PID2022-141283NB-I00 with the support of FEDER funds, by the Spanish MCIN with funding from European Union NextGenerationEU (Grant No. PRTR-C17.I1) and the Generalitat de Catalunya, as well as the Spanish MTDFF through the QUANTUM ENIA project call—Quantum Spain project, and by the European Union through the Recovery, Transformation and Resilience Plan—NextGenerationEU within the framework of the Digital Spain 2026 Agenda. E.R.-S was supported by FI AGAUR grant Joan Oró (Ref. 2024 FI-2 00410). J.C. also acknowledges support from the ICREA Academia award. A.S. and M.W.M. acknowledge European Commission projects Field-SEER (ERC 101097313), OPMMEG (101099379), and QUANTIFY (101135931); Spanish Ministry of Science MCIN project SAPONARIA (PID2021-123813NB-I00) and SALVIA (PID2024-158479NB-I00), “NextGenerationEU/PRTR,” (Grant No. FJC2021-047840-I) and “Severo Ochoa” Center of Excellence CEX2019-000910-S; Generalitat de Catalunya through the CERCA program, DURSI Grant No. 2021 SGR 01453 and QSENSE (GOV/51/2022); Fundació Privada Cellex and Fundació Mir-Puig. M.S. acknowledges support from Ayuda Ramón y Cajal 2021 (RYC2021-032032-I, MICIU/AEI/10.13039/501100011033, ESF+) as well as Project FEDER C-EXP-256-UGR23 Consejería de Universidad, Investigación e Innovación y UE Programa FEDER Andalucía 2021-2027. G.G. acknowledges financial support by the DAAD, the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation, Projects No. 447948357 and No. 440958198), the SinoGerman Center

for Research Promotion (Project No. M-0294), the German Ministry of Education and Research (Project QuKuK, BMBF Grant No. 16KIS1618K), and the Stiftung Innovation in der Hochschullehre.

DATA AVAILABILITY

The experimental data that support the findings of this article are openly available [83].

APPENDIX A: PROBABILITY DISTRIBUTION OF THE MEASURED SIGNAL IN THE STEADY STATE

The formal solution of the Ornstein-Uhlenbeck process given in Eq. (23) is well known [68] and can be readily obtained to be

$$\mathbf{J}_t = e^{\hat{A}t} \mathbf{J}_0 + \int_0^t e^{\hat{A}(t-s)} \sqrt{Q} d\mathbf{W} = e^{\hat{A}t} \mathbf{J}_0 + \mathbf{N}_t, \quad (\text{A1})$$

with $\mathbf{J}_0$ the initial condition and $\mathbf{J}_t$ being a Markovian process. The first term in Eq. (A1) is the deterministic solution, whereas the second is the stochastic contribution due to Wiener noise, defined by $\mathbf{N}_t$. Since Eq. (23) is a linear system with additive Gaussian noise, $\mathbf{J}_t$ is a stochastic Gaussian variable, and is therefore fully determined by its mean, $\mathbb{E}[\mathbf{J}_t]$, and covariance matrix $\text{Cov}(\mathbf{J}_t, \mathbf{J}_{t'})$. Averaging Eq. (23), and recalling that $\mathbb{E}[d\mathbf{W}_i] = 0$, we obtain the differential equation $d\mathbb{E}[\mathbf{J}_t] = \hat{A}\mathbb{E}[\mathbf{J}_t]dt$ whose solution can be verified to be

$$\mathbb{E}[\mathbf{J}_t] = e^{\hat{A}t} \mathbf{J}_0,$$

with

$$e^{\hat{A}t} = e^{-\gamma t} \begin{pmatrix} \cos(\omega_L t) & \sin(\omega_L t) \\ -\sin(\omega_L t) & \cos(\omega_L t) \end{pmatrix}. \quad (\text{A2})$$

Notice that the asymptotic expected value achieved in the steady state is $\mathbb{E}[\mathbf{J}_t]_A := \lim_{t \rightarrow \infty} \mathbb{E}[\mathbf{J}_t] = 0$ regardless of $\mathbf{J}_0$.

From Eq. (A1), it is straightforward to check that the noise properties of $\mathbf{J}_t$ are inherited by those of its stochastic component $\mathbf{N}_t$. Using that $\mathbb{E}[\mathbf{N}_t] = 0$ since $\mathbf{N}_t$ is the infinite sum of Wiener noises:

$$\begin{aligned} \text{Cov}(\mathbf{J}_t, \mathbf{J}_{t'}) &= \mathbb{E}[\mathbf{J}_t \cdot \mathbf{J}_{t'}^T] - \mathbb{E}[\mathbf{J}_t] \mathbb{E}[\mathbf{J}_{t'}^T] \\ &= \mathbb{E}[\mathbf{N}_t \mathbf{N}_{t'}^T] = \text{Cov}(\mathbf{N}_t, \mathbf{N}_{t'}). \end{aligned} \quad (\text{A3})$$

In fact, due to the properties of the independent Wiener noise processes, i.e., $\mathbb{E}[d\mathbf{W}_s^i d\mathbf{W}_{s'}^j] = \delta_{ij} \delta(s - s') ds ds'$ for $i = y, z$ and $j = y, z$, one can see that when $t' < t$,

$$\begin{aligned} \mathbb{E}[\mathbf{N}_t \mathbf{N}_{t'}^T] &= Q \int_0^t \int_0^{t'} e^{\hat{A}(t-s)} \mathbb{E}[d\mathbf{W}_s d\mathbf{W}_{s'}^T] (e^{\hat{A}(t'-s')})^T = Q \int_0^t \int_0^{t'} e^{\hat{A}(t-s)} \delta(s - s') ds ds' \mathbb{1}(e^{\hat{A}(t'-s')})^T \\ &= Q \int_0^{t'} e^{\hat{A}(t-s')} (e^{\hat{A}(t'-s')})^T ds' = \frac{Q}{2\gamma} (e^{-\gamma(t-t')} - e^{-\gamma(t+t')}) \begin{pmatrix} \cos(\omega_L(t-t')) & \sin(\omega_L(t-t')) \\ -\sin(\omega_L(t-t')) & \cos(\omega_L(t-t')) \end{pmatrix}. \end{aligned} \quad (\text{A4})$$

Repeating the same procedure for $t' > t$, we conclude that

$$\text{Cov}(\mathbf{J}_t, \mathbf{J}_{t'}) = \frac{Q}{2\gamma} (e^{-\gamma|t-t'|} - e^{-\gamma(t+t')}) \begin{pmatrix} \cos(\omega_L(t-t')) & \sin(\omega_L(t-t')) \\ -\sin(\omega_L(t-t')) & \cos(\omega_L(t-t')) \end{pmatrix}. \quad (\text{A5})$$

Note that for $t = t'$, $\text{Cov}(\mathbf{J}_t, \mathbf{J}_t) = \frac{Q}{2\gamma}(1 - e^{-2\gamma t})\mathbb{1}$, and at $t = t' = 0$, the variance is zero since $\mathbf{J}_0$ is a fixed initial condition. Moreover, the asymptotic variance is $\text{Cov}(\mathbf{J}_t, \mathbf{J}_t)_A := \lim_{t \rightarrow \infty} \text{Cov}(\mathbf{J}_t, \mathbf{J}_t) = \frac{Q}{2\gamma}\mathbb{1}$. For a process with unknown $\mathbf{J}_0$ but that starts in the steady state, i.e., $\mathbb{E}[\mathbf{J}_0] = 0$ and $\text{Var}(\mathbf{J}_0, \mathbf{J}_0) = \frac{Q}{2\gamma}\mathbb{1}$, the first two moments are $\mathbb{E}[\mathbf{J}_t]_{\text{steady}} = 0$ and

$$\begin{aligned} \text{Cov}(\mathbf{J}_t, \mathbf{J}_{t'})_{\text{steady}} \\ = \frac{Q}{2\gamma} e^{-\gamma|t-t'|} \begin{pmatrix} \cos(\omega_L(t-t')) & \sin(\omega_L(t-t')) \\ -\sin(\omega_L(t-t')) & \cos(\omega_L(t-t')) \end{pmatrix}. \end{aligned} \quad (\text{A6})$$

Observe that $\text{Cov}(\mathbf{J}_t, \mathbf{J}_{t'})_{\text{steady}}$ is translational invariant since it can be expressed as a function of the difference $t - t'$. Regarding the properties of the photocurrent I , since $\mathbf{J}$ and ξ are independent normal variables, I is also a Gaussian. At the discretized harvesting times $t_m = m\Delta$, the mean and covariance matrix elements of I in the steady state are $\mathbb{E}[I_m] = g_D \mathbb{E}[J_z(t_m)] = 0$ and

$$\begin{aligned} \text{Cov}(I_m, I_{m-k}) &= g_D^2 \text{Cov}(J_z(t_m), J_z(t_{m-k})) \\ &\quad + \text{Cov}(\xi(t_m), \xi(t_{m-k})) \\ &= g_D^2 \frac{Q}{2\gamma} e^{-\gamma|k|\Delta} \cos(\omega_L k \Delta) + \delta_{0k} \frac{S_{ph}}{\Delta} := T_k, \end{aligned} \quad (\text{A7})$$

where again the second moments of I are translational invariant. With this at hand, we present the normal probability distribution of a sequence of n photocurrent observations in the steady state, $\mathbf{I}_n = \{I_1, \dots, I_n\}$, i.e., $\mathbf{I}_n \sim \mathcal{N}(\boldsymbol{\mu}_n, \hat{T}_n)$ with mean and Toeplitz covariance matrix:

$$\begin{aligned} \boldsymbol{\mu}_n &= \mathbb{E}[\mathbf{I}_n] = \mathbf{0}, \\ \hat{T}_n &= \begin{pmatrix} T_0 & T_1 & T_2 & \cdots & T_n \\ T_1 & T_0 & T_1 & & T_{n-1} \\ T_2 & T_1 & T_0 & \ddots & \vdots \\ \vdots & & \ddots & \ddots & T_1 \\ T_n & T_{n-1} & \cdots & T_1 & T_0 \end{pmatrix}. \end{aligned} \quad (\text{A8})$$

APPENDIX B: TOEPLITZ MATRICES

In this Appendix, we overview the main definitions and theorems that we use in this work regarding Toeplitz operators, first introduced in Sec. III B.

A Toeplitz matrix $T_n(f)$ of size $n \times n$ is defined by a symbol $f(\theta)$, a 2π -periodic function $f \in L^\infty([0, 2\pi))$, through its entries [84]:

$$(T_n(f))_{ij} = C_{i-j},$$

where

$$C_k = \frac{1}{2\pi} \int_0^{2\pi} f(\theta) e^{-ik\theta} d\theta, \quad 1 \leq i, j \leq n. \quad (\text{B1})$$

Here, $L^\infty([0, 2\pi))$ is the space of all essentially bounded, Lebesgue-measurable functions on the interval $[0, 2\pi)$.

The symbol generates the matrix via Fourier coefficients, with $f(\theta) = \sum_{k=-\infty}^{\infty} C_k e^{ik\theta}$. For large n , the spectral

properties of $T_n(f)$ are governed by harmonic analysis on the unit circle.

For the process under study, recalling Eq. (25),

$$C_k = T_{i,i+k} = \gamma S_{\text{at}} e^{-\gamma|k|\Delta} \cos(k\omega_L \Delta) + \delta_{k,0} \frac{S_{ph}}{\Delta}, \quad (\text{B2})$$

with $S_{\text{at}} = g_D^2 \frac{Q}{2\gamma^2}$. The corresponding symbol $f(\theta)$ is

$$\begin{aligned} f(\theta) &= \frac{1}{2} \frac{\gamma S_{\text{at}} \sinh(\gamma \Delta)}{\cosh(\gamma \Delta) - \cos(\theta - \Delta \omega_L)} \\ &\quad + \frac{1}{2} \frac{\gamma S_{\text{at}} \sinh(\gamma \Delta)}{\cosh(\gamma \Delta) - \cos(\theta + \Delta \omega_L)} + \frac{S_{ph}}{\Delta}, \end{aligned} \quad (\text{B3})$$

which, by taking $\theta = \omega \Delta$, at leading order in $\Delta^{-1} \ll 1$ leads to Lorentzian profile found in Eq. (31).

Szegő's theorems and spectral properties

Theorem B1. First Szegő Theorem: For $f > 0$ with $\ln f \in L^1$, the determinant satisfies

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln \det T_n(f) = \frac{1}{2\pi} \int_0^{2\pi} \ln f(\theta) d\theta \equiv \ln G(f), \quad (\text{B4})$$

where $G(f)$ is the geometric mean of f .

Theorem B2. Strong Szegő Theorem: If f is sufficiently smooth (e.g., $f \in C^\infty$) and positive,

$$\det T_n(f) \sim G(f)^n \cdot E(f),$$

$$E(f) = \exp \left(\sum_{k=1}^{\infty} k (\ln f)_k (\ln f)_{-k} \right), \quad (\text{B5})$$

where $(\cdot)_k$ denote Fourier coefficients. The eigenvalues $\{\lambda_j\}_{j=1}^n$ satisfy $\lambda_j \rightarrow f(2\pi j/n)$ in distribution.

Traces and inverses for smooth $f > 0$:

Traces of powers:

$$\text{Tr}(T_n(f)^k) = \frac{n}{2\pi} \int_0^{2\pi} f(\theta)^k d\theta + C_k(f) + o(1).$$

Inverses:

$$T_n(f)^{-1} \approx T_n(1/f),$$

$$\text{Tr}(T_n(f)^{-1}) = \frac{n}{2\pi} \int_0^{2\pi} \frac{d\theta}{f(\theta)} + O(1). \quad (\text{B6})$$

Here, $C_k(f) = O(1)$ depends on the symbol's smoothness. The Fourier basis $\{e^{ij\theta}\}_{j=1}^n$ asymptotically diagonalizes $T_n(f)$.

Generalized traces: For products involving inverses of Toeplitz matrices with smooth positive symbols,

$$\text{Tr} \left(\prod_{j=1}^m T_n(f_j(\theta))^{\epsilon_j} \right) = \frac{n}{2\pi} \int_0^{2\pi} \prod_{j=1}^m f_j(\theta)^{\epsilon_j} d\theta + O(1), \quad (\text{B7})$$

where $\epsilon_j = \pm 1$.

Proof. We provide a sketch of the proof using some of the results in Ref. [84]. A Toeplitz matrix $T_n(f)$ can be understood as $T_n(f) = P_n T(f) P_n$, where $T(f)$ is the infinite-dimensional Toeplitz matrix (bounded operator) and $P_n : \{x_1, x_2, x_3, \dots\} \rightarrow \{x_1, x_2, \dots, x_n, 0, 0, \dots\}$ is the projector on

the n th dimensional subspace. The main ingredients for the proof are the following:

(1) By Proposition 2.12 (Widom), p. 40, if $a, b \in L^\infty$, then

$$T_n(ab) = T_n(a)T_n(b) + P_n H(a)H(\tilde{b})P_n + W_n H(\tilde{a})H(b)W_n, \quad (\text{B8})$$

where H denotes the so-called Hankel matrix, $W_n : \{x_1, x_2, x_3, \dots\} \rightarrow \{x_n, x_{n-1}, \dots, x_1, 0, 0, \dots\}$ is a projection, and $\tilde{a} := a(1/t)$ with $t = e^{i\theta}$.

(2) By Theorem 2.14 (Widom), p. 42, if $T(a)$ is invertible, for all sufficiently large n ,

$$T_n(a)^{-1} = T_n(a^{-1}) + P_n K(a)P_n + W_n K(\tilde{a})W_n + C_n, \quad (\text{B9})$$

where $K(a) = H(a^{-1})H(\tilde{a})T^{-1}(a)$, $K(\tilde{a}) = H(\tilde{a}^{-1})H(a)T^{-1}(\tilde{a})$, and $\|C_n\| \rightarrow 0$ as $n \rightarrow \infty$. Here, $\|\cdot\|$ is the operator norm.

(3) By Theorem 2.15 (p. 43), if $T(a)$ is invertible and $\sum_{n \in \mathbb{Z}} |n|^\alpha |a_n| < \infty$ for some $\alpha > 0$, then $\|C_n\| = o(n^{-\alpha})$.

(4) Let W be the Wiener algebra and $W \cap B^{1/2} = \{a \in W : \sum_{n \in \mathbb{Z}} (|n| + 1)|a_n|^2 < \infty\}$. From pp. 122–123, if $a \in W \cap B^{1/2}$ and has no zeros, then $a^{-1} \in W \cap B^{1/2}$. Moreover, if $a, b \in W \cap B^{1/2}$ (i.e., a, b sufficiently smooth), then $H(a)H(\tilde{b})$ is a trace class operator, i.e., $\|H(a)H(\tilde{b})\|_1 < \infty$, where $\|\cdot\|_1$ is the trace norm.

Notice that from (1) and (4),

$$\begin{aligned} & |\text{Tr}(T_n(ab) - T_n(a)T_n(b))| \\ &= |\text{Tr}(P_n H(a)H(\tilde{b})P_n + W_n H(\tilde{a})H(b)W_n)| \\ &\leq \|P_n H(a)H(\tilde{b})P_n + W_n H(\tilde{a})H(b)W_n\|_1 \\ &\leq \|P_n H(a)H(\tilde{b})P_n\|_1 + \|W_n H(\tilde{a})H(b)W_n\|_1 < c, \end{aligned} \quad (\text{B10})$$

where c is a constant since $\|P_n H(a)H(\tilde{b})P_n\|_1 \leq \|P_n\|^2 \|H(a)H(\tilde{b})\|_1 = \|H(a)H(\tilde{b})\|_1 < \infty$ and the same for $\|W_n H(\tilde{a})H(b)W_n\|_1$ (a trace class operator multiplied by a bounded operator is still trace class).

For traces involving inverses and $\alpha \geq 1$, from (2)–(4),

$$\begin{aligned} & |\text{Tr}(T_n(a)^{-1} - T_n(a^{-1}))| \\ &+ \|P_n H(a^{-1})H(\tilde{a})T^{-1}(a)P_n\|_1 \\ &+ \|W_n H(\tilde{a}^{-1})H(a)T^{-1}(\tilde{a})W_n\|_1 + \|C_n\|_1 < c', \end{aligned} \quad (\text{B11})$$

since $P_n H(a^{-1})H(\tilde{a})T^{-1}(a)P_n$ and $W_n H(\tilde{a}^{-1})H(a)T^{-1}(\tilde{a})W_n$ are trace class due to the bounded property of $T^{-1}(a)$ and $T^{-1}(\tilde{a})$. Moreover, $\|C_n\|_1 < n\|C_n\| = o(n^{1-\alpha}) < o(1)$ for $\alpha \geq 1$.

From Eq. (B11),

$$\begin{aligned} & |\text{Tr}(T_n(b)T_n(a)^{-1} - T_n(b)T_n(a^{-1}))| \\ &< \|T_n(b)\| \|T_n(a)^{-1} - T_n(a^{-1})\|_1 < c'', \end{aligned} \quad (\text{B12})$$

and therefore combining Eqs. (B10) and (B12),

$$\text{Tr}(T_n(b)T_n(a)^{-1}) = \text{Tr}(T_n(ba^{-1})) + O(1). \quad (\text{B13})$$

These processes can be repeated recursively, when increasing the number of multiplying Toeplitz inside the trace, to obtain

$$\text{Tr} \left(\prod_{j=1}^m T_n(f_j(\theta))^{\epsilon_j} \right) = \text{Tr} \left(\prod_{j=1}^m T_n(f_j(\theta)^{\epsilon_j}) \right) + O(1), \quad (\text{B14})$$

and by the Szegő theorems arrive to the result in Eq. (B7). ■

The generalized trace property in Eq. (B7) is defined for the symbol $f(\theta)$ with θ in the unit circle, i.e., $\theta \in (0, 2\pi)$. When we move to the continuous limit, i.e., $\Delta \rightarrow 0$, we have that ω is defined in the interval $(-\infty, \infty)$. Moreover, $f_j(\omega\Delta) = g_j(\omega) + O(\Delta)$ and for sufficiently smooth symbols,

$$\begin{aligned} & \text{Tr} \left(\prod_{j=1}^m T_n(f_j(\theta))^{\epsilon_j} \right) \\ &= \frac{n\Delta}{2\pi} \int_0^{2\pi/\Delta} \prod_{j=1}^m f_j(\omega\Delta)^{\epsilon_j} d\omega + O(1) \\ &= \frac{t}{2\pi} \int_{-\infty}^{\infty} \prod_{j=1}^m g_j(\omega)^{\epsilon_j} d\omega + tO(\Delta) + O(1). \end{aligned}$$

APPENDIX C: HYBRID KALMAN FILTER ALGORITHM

In the two-step procedure of the KF algorithm, we denote by $\tilde{\mathbf{j}}_{k|k-1}$ the predicted state of the system at time t_k , and the corresponding estimators that construct the mean and covariance of the normal distribution as $\tilde{\mathbf{j}}_{k|k-1}$ and $\hat{\Sigma}_{k|k-1}$. In addition, for the computation of the next predictions, we denote the updated estimators by $\tilde{\mathbf{j}}_{k|k}$ and $\hat{\Sigma}_{k|k}$. Then, the explicit procedure is given as follows [72,85].

Prediction ($\tilde{\mathbf{j}}_{k-1|k-1} \rightarrow \tilde{\mathbf{j}}_{k|k-1}$, $\hat{\Sigma}_{k-1|k-1} \rightarrow \hat{\Sigma}_{k|k-1}$). We define the transition matrix $\hat{\Phi}$ as the solution of the deterministic part of the state vector. When $\hat{A}$ is time independent, for any Δ interval, $\hat{\Phi} := \hat{\Phi}_\Delta = e^{\hat{A}\Delta}$. Then, we can construct $\tilde{\mathbf{j}}_{k|k-1}$ and $\hat{\Sigma}_{k|k-1}$ over the interval $[t_{k-1}, t_k]$, as

$$\begin{aligned} \tilde{\mathbf{j}}_{k|k-1} &= \hat{\Phi} \tilde{\mathbf{j}}_{k-1|k-1}, \\ \hat{\Sigma}_{k|k-1} &= \hat{\Phi} \hat{\Sigma}_{k-1|k-1} \hat{\Phi} + \hat{\mathcal{Q}}_k^\Delta, \end{aligned} \quad (\text{C1})$$

where $\hat{\mathcal{Q}}_k^\Delta := \int_{t_{k-1}}^{t_k} \hat{\Phi}_{t_k, \tau} \hat{\mathcal{Q}}_\tau \hat{\Phi}_{t_k, \tau}^T d\tau$ is the effective covariance matrix of the system noise.

Update ($\tilde{\mathbf{j}}_{k|k-1} \rightarrow \tilde{\mathbf{j}}_{k|k}$, $\hat{\Sigma}_{k|k-1} \rightarrow \hat{\Sigma}_{k|k}$). To take into account the k th outcome $\mathbf{I}_k$, one computes

$$\begin{aligned} \tilde{\mathbf{j}}_{k|k} &= \tilde{\mathbf{j}}_{k|k-1} + \hat{K}_k \tilde{\mathbf{y}}_k, \\ \hat{\Sigma}_{k|k} &= (\mathbb{1} - \hat{K}_k \hat{C}) \hat{\Sigma}_{k|k-1}, \end{aligned} \quad (\text{C2})$$

where the so-called innovation and the Kalman gain are, respectively, $\tilde{\mathbf{y}}_k = \mathbf{I}_k - \mu_k$ and $\hat{K}_k = \hat{\Sigma}_{k|k-1} \hat{C}^T \hat{S}_k^{-1}$. Here, $\mu_k = \hat{C} \tilde{\mathbf{j}}_{k|k-1}$ is the prediction of the k th outcome value and $\sigma_k^2 = \mathbb{E}[\tilde{\mathbf{y}}_k \tilde{\mathbf{y}}_k^T] = \hat{G}_k^\Delta + \hat{C} \hat{\Sigma}_{k|k-1} \hat{C}^T$ is the covariance matrix for the k th innovation $\tilde{\mathbf{y}}_k$.

Finally, to initialize the filter at time t_0 one chooses an appropriate initial distribution, $p(\mathbf{J}_{0|0}) \sim \mathcal{N}(\tilde{\mathbf{J}}_{0|0}, \hat{\Sigma}_{0|0})$, representing the knowledge about the state vector. In our case, we




FIG. 11. Block diagram of Eq. (36) (see main text) that describes the log-likelihood ratio in terms of the Kalman estimators. Here, n represents the time step and Z^{-1} the time delay. The rest of the elements are part of the hybrid Kalman filter algorithm (see Appendix C).

start with the steady-state solution derived in Sec. A, $\tilde{\mathbf{J}}_{0|0} = \mathbf{0}$ and $\hat{\Sigma}_{0|0} = \frac{Q}{2\gamma} \mathbf{1}$.

APPENDIX D: BLOCK DIAGRAMS FOR SEQUENTIAL STRATEGIES

In this Appendix, we present the block diagrams of the high-pass filter and the log-likelihood ratio that can be deployed on an Field-Programmable Gate Array (FPGA) board for real-time implementation. As explained in the main text, these two diagrams represent a chain process; i.e., at each time step n , the log-likelihood ratio is evaluated (Fig. 11) after the filtered signal R_n is first calculated (Fig. 12). Finally, either the SPRT or CUSUM thresholds are applied for sequential hypothesis testing or quickest change-point detection, respectively.

APPENDIX E: MODEL FOR STOCHASTIC SPIN RELAXATION PARAMETER AND EXPERIMENTAL DATA BEHAVIOR

As discussed in Sec. VB, at long observation times, we observe discrepancies between the simulated and experimental results in the hypothesis-testing scenario. To investigate the origin of these differences, we analyze the LLR as a function of time. In Fig. 13, we present the mean and variance of the statistic obtained from the 10^4 experimental traces, each

lasting $t = 40$ ms per hypothesis, and compare them with the simulated results and theoretical predictions. The exact result for the mean is obtained in a recursive way via the Kalman estimators and the analytical slope for the variance is directly the asymptotic solution.

The LLR is computed using both the Kalman filter (shown in blue) and the Whittle log-likelihood method introduced in Sec. V (shown in red). For the latter, we compute the raw periodogram (no high-pass filter applied) and perform a frequency-domain cut, considering only the contributions from the range near the Lorentzian peaks—specifically from 40 to 60 kHz. This range is chosen to be broad enough to retain relevant information while excluding less informative regions and the spurious low-frequency noise. Then, the Whittle LLR is

$$\sum_{k=m}^l \ln \frac{\bar{S}_1(\omega_k)}{\bar{S}_0(\omega_k)} + \left(\frac{1}{\bar{S}_0(\omega_k)} - \frac{1}{\bar{S}_1(\omega_k)} \right) S_l(\omega_k), \quad (\text{E1})$$

where $\omega_k = 2\pi k/t$, and m and l are the integers corresponding to the edges of the region of interest. Observe that both approaches converge as time t increases, and for both hypotheses, the variance grows faster than linearly with time. In contrast, the simulated traces (shown in green) closely follow the expected behavior of the first two moments of the LLR, as indicated by the theoretical curves (shown in black).

To understand the behavior of the LLR, especially with respect to the variance, we computed the covariance matrix across different frequencies of the experimental periodogram. Figure 14 displays a two-dimensional (2D) plot of these frequency correlations near the Larmor frequency $\omega_{L,h}$ and shows an unexpected structure emerging around this frequency, taking both positive and negative values. To better visualize this structure, we exclude the diagonal elements (i.e., the variance), which are orders of magnitude larger and obscure the finer details. Furthermore, we have applied a 5×5 2D binning to reduce noise and enhance the clarity of the correlation patterns.




FIG. 12. Block diagram of Eq. (47) (see main text) that describes the high-pass filter. Here, n represents the time step and Z^{-1} the time delay. Recall that α is the strength of the filter.




FIG. 13. Regularized mean and variance results of the LLR. Top plots for $h = 0$ and bottom plots for $h = 1$. Left and right show, respectively, the first and second moments. In all the cases, in blue the computation of the experimental LLR through the Kalman filter is presented and in red the Whittle version. The simulated results are shown in green and the black curves represent the theoretical predictions for constant parameters.

For more details about its shape, we present some horizontal cuts of the 2D plots in Fig. 15 and observe that near the Larmor frequency, the correlations show a sinc function structure.

Notice that this profile is not predicted by the physical model in Eqs. (23) and (24). At long timescales, it is expected to follow an exponential distribution and exhibit no correlations across different frequencies. This theoretical




FIG. 14. PSD covariance matrix at different frequencies near the Larmor frequency where the diagonal elements are removed. Left and right: results for the experimental dataset under $h = 0$ and $h = 1$, respectively.




FIG. 15. Horizontal cut of the PSD covariance matrix at different frequencies. Each color represents a fixed frequency. Again, the correlation at same frequencies is removed. Left and right: results for experimental dataset under $h = 0$ and $h = 1$, respectively.

behavior is consistent with the corresponding plots for the simulated data, as shown in Fig. 16. In these simulations, only noise near zero is observed, and no well-defined pattern emerges.

The structure of correlations followed by the experimental datasets can be explained qualitatively by a stochastic relaxation time. In Eq. (23), γ is assumed to be constant but it seems to fluctuate run to run, during the experimental acquisition time. Let us now briefly explain how run-to-run fluctuations in the relaxation constant can give rise to artificial correlations between different frequency components of the measured signal. In essence, when a parameter like γ fluctuates slightly from run to run, each realization of the signal is effectively drawn from a different process, each with a slightly different frequency response. So even though the frequency components remain uncorrelated in any single realization, the relative amplitudes of all spectral components are jointly influenced by the overall shape of the spectrum, giving rise to apparent correlations when computing ensemble averages.

We model the fluctuating parameter as $\tilde{\gamma} = \gamma + \delta\gamma$, where γ is a constant central value and $\delta\gamma$ is a random variable with $\mathbb{E}[\delta\gamma] = 0$ and standard deviation $\sigma_{\delta\gamma}$. Moreover, we assume that the total energy of the system is constant, and do not vary for different values of γ . The PSD function used is

$$\tilde{S}_{h,\gamma}(\omega) = \tilde{S}_{\text{at}} \frac{\gamma}{\gamma^2 + (\omega - \omega_{L,h})^2} + S_{\text{ph}}, \quad (\text{E2})$$

where $\tilde{S}_{\text{at}} = S_{\text{at}}\gamma = \frac{\tilde{g}_D Q}{2\gamma}$. Notice that $\int_0^\infty \tilde{S}_{h,\gamma}(\omega) d\omega = \frac{\tilde{g}_D \pi Q}{2\gamma} + \mathcal{O}(\frac{1}{\omega_L})$ and for the total area to be independent of γ , the strength of the Wiener noise has to be proportional to γ , i.e., $Q := Q'\gamma$. Then, $\tilde{S}_{\text{at}}$ is independent of γ and expanding around small $\delta\omega$, we obtain $\tilde{S}_{h,\gamma+\delta\gamma}(\omega) = \tilde{S}_{h,\gamma}(\omega) + \tilde{g}(\omega)\delta\gamma + \mathcal{O}(\delta\gamma^2)$, with

$$\tilde{g}(\omega) = \frac{\tilde{S}_{\text{at}}(\gamma + \omega - \omega_L)(\gamma - \omega + \omega_L)}{(\gamma^2 + (\omega - \omega_L)^2)^2}. \quad (\text{E3})$$




FIG. 16. Horizontal cut of the PSD covariance matrix at different frequencies. Each color represents a fixed frequency. Again, the correlation at same frequencies is removed. Left and right: results for simulated dataset under $h = 0$ and $h = 1$, respectively.




FIG. 17. Horizontal cut of the PSD covariance matrix at different frequencies. Each color represents a fixed frequency. Again, the correlation at same frequencies is removed. Left and right: results for the theoretical predictions under $h = 0$ and $h = 1$, respectively. The value for $\sigma_{\delta\gamma}$ used is 60 Hz.

It can easily be seen that the correlations under $\delta\gamma$ between two independent PSDs at different frequencies is

$$\text{Cov}((S(\omega), S(\omega'))_{\delta\gamma} = \tilde{g}(\omega)\tilde{g}(\omega')\text{Var}(\delta\gamma). \quad (\text{E4})$$

In Fig. 17, we show the predictions of this stochastic spin relaxation model, and see that they qualitatively reproduce the structures observed in Figs. 14 and 15.

Observe that it is crucial that the total energy of the PSD is independent of the fluctuations. If we do not take into account this fact and use

$$S_{h,\gamma}(\omega) = S_{\text{at}} \frac{\gamma^2}{\gamma^2 + (\omega - \omega_{L,h})^2} + S_{\text{ph}}, \quad (\text{E5})$$

where $S_{\text{at}} = \frac{g_D^2 Q}{2\gamma^2}$ and Q is independent of γ , the results after expanding around small $\delta\omega$ are $S_{h,\gamma+\delta\gamma}(\omega) = S_{h,\gamma}(\omega) + g(\omega)\delta\gamma + \mathcal{O}(\delta\gamma^2)$, with

$$g(\omega) = \frac{2S_{\text{at}}\gamma(\omega - \omega_L)^2}{(\gamma^2 + (\omega - \omega_L)^2)^2} > 0, \quad \forall \omega, \quad (\text{E6})$$

and the negative valleys in $\text{Cov}((S(\omega), S(\omega'))_{\delta\gamma}$ cannot be theoretically explained.

With all this information, we are in a position to study the effect on the log-likelihood ratio $\tilde{L}_t$ under the proposed model for stochastic fluctuations. Combining Eqs. (46) and (E4), the statistic asymptotically presents these first two




FIG. 18. Sum of correlations at different frequencies weighed with the LLR functional in dotted curves. In stripped curves, the fitting outcomes are represented. Red and blue show the results for $h = 0, 1$ datasets, respectively. In the plots and fits, the correlations at same frequencies are excluded.




FIG. 19. LLR results as a function of time for the simulated data with fluctuating spin relaxation (red). The experimental results are shown in blue.

moments:

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{\mathbb{E}[\tilde{L}_t]_{h, \delta\gamma}}{t} &= \frac{1}{2\pi} \int_0^\infty \ln \left(\frac{\bar{S}_0(\omega)}{\bar{S}_1(\omega)} \right) + \left(\frac{1}{\bar{S}_0(\omega)} - \frac{1}{\bar{S}_1(\omega)} \right) \bar{S}_h(\omega) d\omega, \\ \lim_{t \rightarrow \infty} \frac{\text{Var}(\tilde{L}_t)_{h, \delta\gamma}}{t^2} &= \frac{1}{4\pi^2} \int_0^\infty \int_0^\infty \left(\frac{1}{\bar{S}_0(\omega)} - \frac{1}{\bar{S}_1(\omega)} \right) \left(\frac{1}{\bar{S}_0(\omega')} - \frac{1}{\bar{S}_1(\omega')} \right) \text{Cov}(S(\omega), S(\omega'))_{\delta\gamma} d\omega d\omega' \\ &= \text{Var}(\delta\gamma) \frac{1}{4\pi^2} \int_0^\infty \int_0^\infty \left(\frac{1}{\bar{S}_0(\omega)} - \frac{1}{\bar{S}_1(\omega)} \right) \left(\frac{1}{\bar{S}_0(\omega')} - \frac{1}{\bar{S}_1(\omega')} \right) \tilde{g}(\omega) \tilde{g}(\omega') d\omega d\omega'. \end{aligned} \quad (\text{E7})$$

Notice that the leading term of $\text{Var}(L_t)_{h, \delta\gamma}$ scales quadratically with time t , unlike $\text{Var}(L_t)_h$, which corresponds to the case with constant parameters and grows linearly with time. This linear behavior can be derived using the same reasoning as in Sec. IV B. Finally, we can use the expression of $\text{Var}(L_t)_{h, \delta\gamma}$ in Eq. (E7) to estimate $\text{Var}(\delta\gamma)$ and automatically obtain the standard deviation $\sigma_{\delta\gamma}$. In Fig. 18, we show the contribution of the correlations between different frequencies in the LLR, for both experimental and simulated datasets, computed via the Whittle log-likelihood ratio for the range 40–60 kHz. The left plot shows the results for the simulated traces, where, as expected, the contribution remains constant over time. In contrast, the right plot displays the experimental data, where artificial fluctuations increase with time. The fitted curves are shown in black, from which we extract standard deviations of

$\sigma_{0, \delta\gamma} = 57.3$ Hz and $\sigma_{1, \delta\gamma} = 50.7$ Hz for the datasets of $\omega_{L,0}$ and $\omega_{L,1}$, respectively.

Finally, we take the largest value for $\sigma_{\delta\gamma}$, i.e., $\sigma_{0, \delta\gamma}$, and generate time traces under the conditions that this model predicts. In Fig. 19, we present the results of the LLR for simulated data with fluctuating spin relaxation (red) and compare them with the experimental results (blue).


APPENDIX F: ULTIMATE ERROR RATES

In this Appendix, we provide analytical asymptotic expressions for the sequential and Chernoff rates. Recall that we are in a scenario where the choice of hypothesis $h = 0, 1$ only affects the sign of the LLR and the sequential rate can be obtained from Eq. (68). For high ω_c ,

$$\rho_{\text{seq}} = \frac{S_{\text{at}} \gamma \Delta \omega^2 (S_{\text{at}} (-3\sqrt{S_{\text{ph}}} + \sqrt{S_{\text{at}} + S_{\text{ph}}}) \gamma^2 + S_{\text{ph}} (-\sqrt{S_{\text{ph}}} + \sqrt{S_{\text{at}} + S_{\text{ph}}}) (4\gamma^2 + \Delta \omega^2))}{2\sqrt{S_{\text{at}} + S_{\text{ph}}} (S_{\text{at}}^2 \gamma^4 + 2S_{\text{at}} S_{\text{ph}} \gamma^2 \Delta \omega^2 + S_{\text{ph}}^2 \Delta \omega^2 (4\gamma^2 + \Delta \omega^2))} - \frac{S_{\text{at}}^2 \gamma^4 \Delta \omega^2}{5\pi S_{\text{ph}} \omega_c^5} + \mathcal{O}\left(\frac{1}{\omega_c^6}\right). \quad (\text{F1})$$

On the other hand, for the Chernoff rate computation we rely on Eq. (65). To obtain the value of s that provides the tightest bound in our experimental scenario, i.e., $C(s^*) := \max_s C(s)$, assuming one can compute the following quantity,

$$\begin{aligned} \frac{\partial}{\partial s} \lim_{t \rightarrow \infty} \left(\frac{C(s)}{t} \right) &= \int_0^\infty \frac{d\omega}{2\pi} \frac{\partial}{\partial s} \ln \frac{\bar{S}_{I,0}(\omega_j)^{1-s} \bar{S}_{I,1}(\omega_j)^s}{(1-s) \bar{S}_{I,0}(\omega_j) + s \bar{S}_{I,1}(\omega_j)} = \frac{(1-2s) S_{\text{at}}^2 \gamma^4 \Delta \omega^2}{5\pi S_{\text{ph}}^2 \omega_c^5} \\ &+ \frac{((-1+2s) S_{\text{at}}^2 \gamma^4 \Delta \omega^2 (2(S_{\text{at}} + 2S_{\text{ph}}) \gamma^2 - S_{\text{ph}} \Delta \omega^2))}{7\pi S_{\text{ph}}^3 \omega_c^7} + \frac{(1+6(-1+s)s) S_{\text{at}}^3 \gamma^6 \Delta \omega^3}{12\pi S_{\text{ph}}^3 \omega_c^8} + \mathcal{O}\left(\frac{1}{\omega_c^9}\right), \end{aligned} \quad (\text{F2})$$

FIG. 20. Heatmap of $r(c_a, c_b)$ for different values of $c_a := \Delta\omega/\gamma$ and $c_b := S_{\text{at}}/S_{\text{ph}}$ that range up to 10.

and set it equal to zero to find the relative extreme. Observe that for large ω_c , the first two terms of the expansion are zero at $s^* = 1/2$. Then, evaluating Eq. (65) at $s^* = 1/2$, we obtain the following closed form at large ω_c :

$$\begin{aligned} \rho_c := \lim_{t \rightarrow \infty} \frac{C(s^*)}{t} &= \frac{S_{\text{at}}^2 \gamma^4 \Delta\omega^2}{20\pi S_{\text{ph}}^2 \omega_c^5} + \left(1 + \sqrt{\frac{S_{\text{at}} + S_{\text{ph}}}{S_{\text{ph}}}}\right) \gamma \\ &- \frac{1}{2\sqrt{S_{\text{ph}}}} \left(\sqrt{2S_{\text{at}}\gamma^2 + 4S_{\text{ph}}\gamma^2 - S_{\text{ph}}\Delta\omega^2 - 2\gamma\sqrt{S_{\text{at}}^2\gamma^2 - 2S_{\text{ph}}(S_{\text{at}} + 2S_{\text{ph}})\Delta\omega^2}} \right) \\ &- \frac{1}{2\sqrt{S_{\text{ph}}}} \left(\sqrt{2S_{\text{at}}\gamma^2 + 4S_{\text{ph}}\gamma^2 - S_{\text{ph}}\Delta\omega^2 + 2\gamma\sqrt{S_{\text{at}}^2\gamma^2 - 2S_{\text{ph}}(S_{\text{at}} + 2S_{\text{ph}})\Delta\omega^2}} \right) + \mathcal{O}\left(\frac{1}{\omega_c^6}\right). \end{aligned} \quad (\text{F3})$$

The leading terms of the sequential and Chernoff rates in Eqs. (F1) and (F3) are four-parameter dependent, i.e., $\Delta\omega, \gamma, S_{\text{at}}$ and S_{ph} (ω_c no longer plays a role). Therefore, one can express these two rates in terms of the following two ratios that provide an idea of the shape and relative distance between the two hypotheses under testing:

$$c_a := \frac{\Delta\omega}{\gamma} > 0, \quad c_b := \frac{S_{\text{at}}}{S_{\text{ph}}} > 0. \quad (\text{F4})$$

In the first case, c_a is an indicator of the overlap between the two Lorentzians since it takes into account their width and distance in between. In the second case, c_b stands for the ratio between the Lorentzian peak, and the background noise. With this change of variable, it turns out that both slopes become a two-parameter function. Then, with the aim to compare both the sequential and Chernoff rates we define the

quantity

$$r(c_a, c_b) := \frac{\rho_{\text{seq}}}{\rho_c} \quad (\text{F5})$$

and study its behavior at different parameter regimes, i.e., $\forall c_a, c_b > 0$. It can be symbolically checked with computational software, like, for instance, Wolfram Mathematica, that $r(c_a, c_b) \geq 4, \forall c_a, c_b > 0$. Therefore, for processes where the Chernoff bound is asymptotically attainable, the sequential strategy outperforms the deterministic one by at least a factor of 4 in terms of less observation time to achieve the same error performance.

In Fig. 20, we show a heatmap of $r(c_a, c_b)$ for different values of c_a, c_b that range up to 10. Notice that for small c_a, c_b , i.e., in the regimes where the two hypothesis are more difficult to distinguish, $r(c_a, c_b)$ saturates the minimum value of 4. Moreover, in scenarios where the two Lorentzians are widely separated or the shot noise is very small compared to the spin signal, $r(c_a, c_b)$ increases up to values bigger than 5.5.

[1] H. M. Wiseman and G. J. Milburn, *Quantum Measurement and Control* (Cambridge University Press, Cambridge, 2009).

[2] J. Kitching, S. Knappe, and E. A. Donley, Atomic sensors—A review, *IEEE Sens. J.* **11**, 1749 (2011).

- [3] A. Serafini, Feedback control in quantum optics: An overview of experimental breakthroughs and areas of application, *Int. Scholarly Res. Not.* **2012**, 275016 (2012).
- [4] M. Aspelmeyer, T. J. Kippenberg, and F. Marquardt, Cavity optomechanics, *Rev. Mod. Phys.* **86**, 1391 (2014).
- [5] B. Rogers, N. Lo Gullo, G. De Chiara, G. M. Palma, and M. Paternostro, Hybrid optomechanics for quantum technologies, *Quantum Meas. Quantum Metrol.* **2**, 000010247820140002 (2014).
- [6] W. Yu, W. C. Jiang, Q. Lin, and T. Lu, Cavity optomechanical spring sensing of single molecules, *Nat. Commun.* **7**, 12311 (2016).
- [7] J. Millen, T. S. Monteiro, R. Pettit, and A. N. Vamivakas, Optomechanics with levitated particles, *Rep. Prog. Phys.* **83**, 026401 (2020).
- [8] P. Kumar, T. Biswas, K. Feliz, R. Kanamoto, M.-S. Chang, A. K. Jha, and M. Bhattacharya, Cavity optomechanical sensing and manipulation of an atomic persistent current, *Phys. Rev. Lett.* **127**, 113601 (2021).
- [9] B.-B. Li, L. Ou, Y. Lei, and Y.-C. Liu, Cavity optomechanical sensing, *Nanophotonics* **10**, 2799 (2021).
- [10] S. Barzanjeh, A. Xuereb, S. Gröblacher, M. Paternostro, C. A. Regal, and E. M. Weig, Optomechanics for quantum technologies, *Nat. Phys.* **18**, 15 (2022).
- [11] S. Z. Ang, G. I. Harris, W. P. Bowen, and M. Tsang, Optomechanical parameter estimation, *New J. Phys.* **15**, 103028 (2013).
- [12] D. Budker and M. Romalis, Optical magnetometry, *Nat. Phys.* **3**, 227 (2007).
- [13] H. M. Wiseman and G. J. Milburn, Quantum theory of field-quadrature measurements, *Phys. Rev. A* **47**, 642 (1993).
- [14] J. Kitching, Chip-scale atomic devices, *Appl. Phys. Rev.* **5**, 031302 (2018).
- [15] S. Forstner, J. Knittel, E. Sheridan, J. D. Swaim, H. Rubinsztein-Dunlop, and W. P. Bowen, Sensitivity and performance of cavity optomechanical field sensors, *Photon. Sens.* **2**, 259 (2012).
- [16] M. Tsang, Continuous quantum hypothesis testing, *Phys. Rev. Lett.* **108**, 170502 (2012).
- [17] K. Jacobs, *Quantum Measurement Theory and its Applications* (Cambridge University Press, Cambridge, 2014).
- [18] K. Mølmer, Hypothesis testing with open quantum systems, *Phys. Rev. Lett.* **114**, 040401 (2015).
- [19] F. Albarelli, M. A. C. Rossi, M. G. A. Paris, and M. G. Genoni, Ultimate limits for quantum magnetometry via time-continuous measurements, *New J. Phys.* **19**, 123011 (2017).
- [20] A. H. Kiilerich and K. Mølmer, Hypothesis testing with a continuously monitored quantum system, *Phys. Rev. A* **98**, 022103 (2018).
- [21] J. F. Ralph, M. Toroš, S. Maskell, K. Jacobs, M. Rashid, A. J. Setter, and H. Ulbricht, Dynamical model selection near the quantum-classical boundary, *Phys. Rev. A* **98**, 010102(R) (2018).
- [22] R. Jiménez-Martínez, J. Kołodyński, C. Troullinou, V. G. Lucivero, J. Kong, and M. W. Mitchell, Signal tracking beyond the time resolution of an atomic sensor by Kalman filtering, *Phys. Rev. Lett.* **120**, 040503 (2018).
- [23] J. Liu, H. Yuan, X.-M. Lu, and X. Wang, Quantum Fisher information matrix and multiparameter estimation, *J. Phys. A: Math. Theor.* **53**, 023001 (2020).
- [24] M. W. Mitchell and S. Palacios Alvarez, *Colloquium: Quantum limits to the energy resolution of magnetic field sensors*, *Rev. Mod. Phys.* **92**, 021001 (2020).
- [25] J. Amorós-Binefa and J. Kołodyński, Noisy atomic magnetometry in real time, *New J. Phys.* **23**, 123030 (2021).
- [26] M. M. Marchese, A. Belenchia, and M. Paternostro, Optomechanics-based quantum estimation theory for collapse models, *Entropy* **25**, 500 (2023).
- [27] M. Wang, D. J. Perez-Morelo, G. Ramer, G. Pavlidis, J. J. Schwartz, L. Yu, R. Ilic, A. Centrone, and V. A. Aksyuk, Beating thermal noise in a dynamic signal measurement by a nanofabricated cavity optomechanical sensor, *Sci. Adv.* **9**, eadf7595 (2023).
- [28] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, Circuit quantum electrodynamics, *Rev. Mod. Phys.* **93**, 025005 (2021).
- [29] J. Amorós-Binefa and J. Kołodyński, Noisy atomic magnetometry with Kalman filtering and measurement-based feedback, *PRX Quantum* **6**, 030331 (2025).
- [30] A. Wald, Sequential tests of statistical hypotheses, *Ann. Math. Stat.* **16**, 117 (1945).
- [31] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Changepoint Detection* (CRC Press, New York, 2014).
- [32] A. Ramdas, P. Grünwald, V. Vovk, and G. Shafer, Game-theoretic statistics and safe anytime-valid inference, *Stat. Sci.* **38**, 576 (2023).
- [33] E. Martínez Vargas, C. Hirche, G. Sentís, M. Skotiniotis, M. Carrizo, R. Muñoz Tapia, and J. Calsamiglia, Quantum sequential hypothesis testing, *Phys. Rev. Lett.* **126**, 180502 (2021).
- [34] Y. Li, V. Y. F. Tan, and M. Tomamichel, Optimal adaptive strategies for sequential quantum hypothesis testing, *Commun. Math. Phys.* **392**, 993 (2022).
- [35] G. Lorden, Procedures for reacting to a change in distribution, *Ann. Math. Stat.* **42**, 1897 (1971).
- [36] M. Fanizza, C. Hirche, and J. Calsamiglia, Ultimate limits for quickest quantum change-point detection, *Phys. Rev. Lett.* **131**, 020602 (2023).
- [37] M. Reichert, Q. Zhuang, J. H. Shapiro, and R. Di Candia, Quantum illumination with a hetero-homodyne receiver and sequential detection, *Phys. Rev. Appl.* **20**, 014030 (2023).
- [38] Z. Gong and S. Guha, Quantum-enhanced change detection and joint communication-detection, *Phys. Rev. A* **112**, 032604 (2025).
- [39] S. Guha, T. C. John, and P. Basu, Quantum-enhanced quickest change detection of transmission loss, *Phys. Rev. Lett.* **135**, 210801 (2025).
- [40] C. Ahn, A. C. Doherty, and A. J. Landahl, Continuous quantum error correction via quantum feedback control, *Phys. Rev. A* **65**, 042301 (2002).
- [41] W. P. Livingston, M. S. Blok, E. Flurin, J. Dressel, A. N. Jordan, and I. Siddiqi, Experimental demonstration of continuous quantum error correction, *Nat. Commun.* **13**, 2307 (2022).
- [42] G. Gasbarri, M. Bilkis, E. Roda-Salichs, and J. Calsamiglia, Sequential hypothesis testing for continuously-monitored quantum systems, *Quantum* **8**, 1289 (2024).
- [43] E. Boto, S. S. Meyer, V. Shah, O. Alem, S. Knappe, P. Kruger, T. M. Fromhold, M. Lim, P. M. Glover, P. G. Morris, R. Bowtell, G. R. Barnes, and M. J. Brookes, A new generation of magnetoencephalography: Room temperature measurements

- using optically-pumped magnetometers, *NeuroImage* **149**, 404 (2017).
- [44] E. Boto, N. Holmes, J. Leggett, G. Roberts, V. Shah, S. S. Meyer, L. D. Muñoz, K. J. Mullinger, T. M. Tierney, S. Bestmann, G. R. Barnes, R. Bowtell, and M. J. Brookes, Moving magnetoencephalography towards real-world applications with a wearable system, *Nature (London)* **555**, 657 (2018).
- [45] S. Knappe, Y. Okada, and S. Krzyzewski, *U.S. Patent No. 12,019,127* (U.S. Patent and Trademark Office, Washington, DC, 2024).
- [46] G. Oelsner, R. IJsselsteijn, T. Scholtes, A. Krüger, V. Schultze, G. Seyffert, G. Werner, M. Jäger, A. Chwala, and R. Stolz, Integrated optically pumped magnetometer for measurements within Earth's magnetic field, *Phys. Rev. Appl.* **17**, 024034 (2022).
- [47] M. S. Mrozowski, A. S. Bell, P. F. Griffin, D. Hunter, D. Burt, J. P. McGilligan, E. Riis, C. D. Beggan, and S. J. Ingleby, Distributed network of optically pumped magnetometers for space weather monitoring, *Sci. Rep.* **14**, 28229 (2024).
- [48] L. Hunter, J. Gordon, S. Peck, D. Ang, and J.-F. Lin, Using the earth as a polarized electron source to search for long-range spin-spin interactions, *Science* **339**, 928 (2013).
- [49] M. S. Safronova, D. Budker, D. DeMille, D. F. J. Kimball, A. Derevianko, and C. W. Clark, Search for new physics with atoms and molecules, *Rev. Mod. Phys.* **90**, 025008 (2018).
- [50] N. A. Sinitsyn and Y. V. Pershin, The theory of spin noise spectroscopy: A review, *Rep. Prog. Phys.* **79**, 106501 (2016).
- [51] V. S. Zapasskii, Spin-noise spectroscopy: From proof of principle to applications, *Adv. Opt. Photon.* **5**, 131 (2013).
- [52] S. A. Crooker, D. G. Rickel, A. V. Balatsky, and D. L. Smith, Spectroscopy of spontaneous spin noise as a probe of spin dynamics and magnetic resonance, *Nature (London)* **431**, 49 (2004).
- [53] V. G. Lucivero, R. Jiménez-Martínez, J. Kong, and M. W. Mitchell, Squeezed-light spin noise spectroscopy, *Phys. Rev. A* **93**, 053802 (2016).
- [54] V. G. Lucivero, A. Dimic, J. Kong, R. Jiménez-Martínez, and M. W. Mitchell, Sensitivity, quantum limits, and quantum enhancement of noise spectroscopies, *Phys. Rev. A* **95**, 041803(R) (2017).
- [55] J. Kong, R. Jiménez-Martínez, C. Troullinou, V. G. Lucivero, G. Tóth, and M. W. Mitchell, Measurement-induced, spatially-extended entanglement in a hot, strongly-interacting atomic system, *Nat. Commun.* **11**, 2415 (2020).
- [56] K. Mouloudakis, F. Vouzinas, A. Margaritakis, A. Koutsimpela, G. Mouloudakis, V. Koutrouli, M. Skotiniotis, G. P. Tsironis, M. Loulakis, M. W. Mitchell, G. Vasilakis, and I. K. Kominis, Interspecies spin-noise correlations in hot atomic vapors, *Phys. Rev. A* **108**, 052822 (2023).
- [57] K. Mouloudakis, G. Vasilakis, V. G. Lucivero, J. Kong, I. K. Kominis, and M. W. Mitchell, Effects of spin-exchange collisions on the fluctuation spectra of hot alkali-metal vapors, *Phys. Rev. A* **106**, 023112 (2022).
- [58] K. Mouloudakis, J. Kong, A. Sierant, E. Arkin, M. Hernández Ruiz, R. Jiménez-Martínez, and M. W. Mitchell, Anomalous noise spectra in a spin-exchange-relaxation-free alkali-metal vapor, *Phys. Rev. A* **109**, L040802 (2024).
- [59] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, Wiley Series in Telecommunications and Signal Processing (Wiley-Interscience, New York, 2006).
- [60] D. Kazakos and P. Papantoni-Kazakos, Spectral distance measures between Gaussian processes, *IEEE Trans. Autom. Control* **25**, 950 (1980).
- [61] A. Wald, *Sequential Analysis* (Courier Corporation, New York, 2004).
- [62] A. Wald and J. Wolfowitz, Optimum character of the sequential probability ratio test, *Ann. Math. Stat.* **19**, 326 (1948).
- [63] E. S. Page, Continuous inspection schemes, *Biometrika* **41**, 100 (1954).
- [64] G. V. Moustakides, Optimal stopping times for detecting changes in distributions, *Ann. Stat.* **14**, 1379 (1986).
- [65] Y. Ritov, Decision theoretic optimality of the cusum procedure, *Ann. Stat.* **18**, 1464 (1990).
- [66] T. L. Lai, Information bounds and quick detection of parameter changes in stochastic systems, *IEEE Trans. Inf. Theory* **44**, 2917 (1998).
- [67] J. Kong, V. G. Lucivero, R. Jiménez-Martínez, and M. W. Mitchell, Long-term laser frequency stabilization using fiber interferometers, *Rev. Sci. Instrum.* **86**, 073104 (2015).
- [68] K. Jacobs, *Stochastic Processes for Physicists: Understanding Noisy Systems* (Cambridge University Press, Cambridge, 2010).
- [69] R. E. Kalman, A new approach to linear filtering and prediction problems, *J. Basic Eng.* **82**, 35 (1960).
- [70] R. E. Kalman and R. S. Bucy, New results in linear filtering and prediction theory, *J. Basic Eng.* **83**, 95 (1961).
- [71] R. Gurajala, P. B. Choppala, J. S. Meka, and P. D. Teal, Derivation of the Kalman filter in a Bayesian filtering perspective, in *Proceedings of the 2021 2nd International Conference on Range Technology (ICORT)* (IEEE, Piscataway, NJ, USA, 2021), pp. 1–5.
- [72] V. Trees and L. Harry, *Detection, Estimation, and Modulation Theory—Part I—Detection, Estimation, and Linear Modulation Theory* (John Wiley & Sons, New York, NY, 2001).
- [73] R. Gray, On the asymptotic eigenvalue distribution of Toeplitz matrices, *IEEE Trans. Inf. Theory* **18**, 725 (1972).
- [74] R. M. Gray *et al.*, Toeplitz and circulant matrices: A review, *Found. Trends® Commun. Inf. Theory* **2**, 155 (2006).
- [75] M. Levin, Power spectrum parameter estimation, *IEEE Trans. Inf. Theory* **11**, 100 (2003).
- [76] P. Whittle, *Hypothesis Testing in Time Series Analysis, Statistics* (Uppsala Universitet, Statistiska Institutionen, Almqvist & Wiksells, Uppsala, 1951), Vol. 4.
- [77] L. Magrini, P. Rosenzweig, C. Bach, A. Deutschmann-Olek, S. G. Hofer, S. Hong, N. Kiesel, A. Kugi, and M. Aspelmeyer, Real-time optimal quantum control of mechanical motion at room temperature, *Nature (London)* **595**, 373 (2021).
- [78] J. Liao *et al.*, FPGA implementation of a Kalman-based motion estimator for levitated nanoparticles, *IEEE Trans. Instrum. Meas.* **68**, 2374 (2018).
- [79] F. Auger, M. Hilairet, J. M. Guerrero, E. Monmasson, T. Orlowska-Kowalska, and S. Katsura, Industrial applications of the Kalman filter: A review, *IEEE Trans. Ind. Electron.* **60**, 5458 (2013).
- [80] J. Ruiz-Rosero, G. Ramirez-Gonzalez, and R. Khanna, Field programmable gate array applications—A scientometric review, *Computation* **7**, 63 (2019).
- [81] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. L. Hollenberg, The nitrogen-vacancy colour centre in diamond, *Phys. Rep.* **528**, 1 (2013).

- [82] M. H. Bitarafan and R. G. DeCorby, On-chip high-finesse Fabry-Perot microcavities for optical sensing and quantum information, *Sensors* **17**, 1748 (2017).
- [83] Zenodo's physical servers and technical infrastructure: CERN (the European Organization for Nuclear Research) in Geneva, Switzerland, <https://doi.org/10.5281/zenodo.17259292>.
- [84] A. Böttcher and B. Silbermann, *Introduction to Large Truncated Toeplitz Matrices* (Springer Science & Business Media, 2012).
- [85] A. Laub, A Schur method for solving algebraic Riccati equations, *IEEE Trans. Autom. Control* **24**, 913 (1979).