

Measurement-driven Langevin modeling of superparamagnetic tunnel junctions

Liam A. Pocher^{1,2}, Temitayo N. Adeyeye^{1,2}, Sidra Gibeault^{1,2}, Philippe Talatchian³,
Ursula Ebels³, Daniel P. Lathrop¹, Jabez J. McClelland⁴, Mark D. Stiles⁴, Advait Madhavan^{1,2}
and Matthew W. Daniels^{4,*}

¹*Institute for Research in Electronics and Applied Physics, University of Maryland, College Park, Maryland, USA*

²*Associate, Physical Measurement Laboratory, National Institute of Standards and Technology, Gaithersburg, Maryland, USA*

³*Univ. Grenoble Alpes, CEA, CNRS, Grenoble INP, SPINTEC, 38000 Grenoble, France*

⁴*Physical Measurement Laboratory, National Institute of Standards and Technology, Gaithersburg, Maryland, USA*



(Received 18 March 2024; revised 17 May 2024; accepted 28 June 2024; published 23 July 2024)

Superparamagnetic tunnel junctions are important devices for a range of emerging technologies, but most existing compact models capture only their mean switching rates. Capturing qualitatively accurate analog dynamics of these devices will be important as the technology scales up. Here we present results using a one-dimensional overdamped Langevin equation that captures statistical properties of measured time traces, including voltage histograms, drift and diffusion characteristics as measured with Kramers-Moyal coefficients, and dwell-time distributions. While common macrospin models are more physically motivated magnetic models than the Langevin model, we show that for the device measured here, they capture even fewer of the measured experimental behaviors.

DOI: [10.1103/PhysRevApplied.22.014057](https://doi.org/10.1103/PhysRevApplied.22.014057)

I. INTRODUCTION

Magnetic tunnel junctions (MTJs) are versatile devices with many modes of operation, several of which are now being recognized for their value to neuromorphic and alternative computing schemes [1–5]. Their superparamagnetic regime, in particular, holds great promise not only for the acceleration of traditional scientific computing workloads [6–9] but also for stochastic computing [10–12], other neural networks [13–17], and combinatorial optimization accelerators [18–21]. Their technological relevance, and the attendant need to design scaled-up circuits that include these devices, has made the task of finding appropriate dynamical models an engineering priority.

Currently, three classes of models are used for different levels of granularity in representing superparamagnetic tunnel junction (SMTJ) physics. The simplest is the Néel-Brown model [22–24], a two-state Markov model where the transition rates are exponential functions of the current and field applied to a device. These models fit transition rates of the Markov model to experiment [25], and can sometimes be aligned with the measurable parameters of the system in a physics-driven way [26,27]. Their discrete state spaces, however, hide any intermediate analog dynamics, and experiments to date indicate that

this analog magnetic behavior is visibly non-negligible [8,25,28–31]. Néel-Brown models implicitly assume an exponential distribution of the dwell times.

The most detailed form of modeling is achieved through micromagnetic simulations. Some micromagnetics packages can explicitly solve the stochastic Landau-Lifshitz-Gilbert (sLLG) equation over large magnetic domains using finite-element methods [32]. The simulation of devices in this manner can lead to physically realistic results, but the required computational timescale renders such models inappropriate for single elements in large circuit simulations. To wit, the authors of Ref. [33] show that micromagnetic simulations qualitatively reproduce measured two-level fluctuations but are not able to collect enough statistics with the full simulations to make a robust statistical comparison.

The intermediate level between these approaches is that of analog compact modeling. Compact models—relatively low-dimensional differential equations that capture the essential analog behaviors of a device without full physical realism—have been developed for many nanodevices [34–37], including MTJs [38–40]. When SMTJs are incorporated into integrated circuits, modeling the interaction between the SMTJs and the transistors will require a statistically accurate description of the SMTJ dynamics on the timescale of switching events in both, but that description cannot be significantly slower than the modeling of the

*Contact author: matthew.daniels@nist.gov

integrated circuit. Those restrictions force the development and use of a compact model that accurately describes the transition dynamics. A Markov model, while fast, does not accurately capture the dynamics of the transition, whereas a micromagnetic model would capture that dynamics, but is far too slow.

Ideas for computing systems that use SMTJs are now being proposed that assume the existence and even interaction of thousands to millions of individual devices [10,13,41,42]. The engineering work of design, test, and verification of large-scale designs like those will require mature models that faithfully capture device physics. Such models already exist for memory-class MTJs used in magnetic random access memory designs and have enabled extraordinary progress in that field [43]. Our present goal, then, is to facilitate a move beyond demonstrating the mere viability of stochastic and probabilistic computing proposals based on SMTJs and toward full-stack-engineered system design based on SMTJ device physics.

For SMTJs in particular, a common compact modeling methodology in the literature is the use of a macrospin (single-domain) model [44]. Macrospin modeling has been applied as a theoretical approach to understanding the physics of magnetic systems [45] and has been used to demonstrate the viability of SMTJ-based alternative computing schemes [10,18]. Yet it is generally understood that real magnetic devices rarely operate in the single-domain regime, calling into question whether macrospin models provide an appropriate simulation framework for the engineering context. Even in devices that are expected to be near-single-domain, the authors of Ref. [46] show explicitly that a simple macrospin model cannot reproduce their experimental results because details of the dynamics change in nontrivial ways during the reversal process. They theorize that the dissipative mechanisms of the model behave outside of what can be captured by the usual sLLG equation, which has a diffusive part that is pre-determined [47] and cannot be tuned except in overall amplitude (through the Gilbert damping). The authors of Ref. [48] show that entropic effects due to degrees of freedom neglected in the macrospin approximation significantly affect both the energetics and the dynamics of the reversal process.

From a practical standpoint, numerical integration of the sLLG equation often requires careful attention [49]; spherical coordinate approaches can depend in subtle ways on the integration scheme [50], and integrating the more stable Euclidean equations requires renormalization of the magnetization [51,52], which slows down the integration, or carefully chosen solvers [53] that an end-user may not be able to access in an engineering context (e.g., when using a commercial circuit simulator).

Facing these limitations of simple macrospin models, we propose a data-driven method to capture the dynamics of interest in circuit simulations. We turn to generic

overdamped one-dimensional (1D) Langevin models for the voltage across the device and fit such models to experimental measurements. Similar data-driven approaches have been applied in other fields of physics to learn the underlying physics of stochastic processes through a variety of methods [54–56]. We show that our approach leads to high-fidelity matching of the voltage histograms and the drift and diffusion coefficients of the device between model and experiment while maintaining the analog nature of macrospin approaches. We also show that our model correctly predicts the drift dynamics and dwell-time distributions of the experiment without explicitly encoding these in the simulation, confirming a level of self-consistency between the model and the underlying physics.

We organize the remainder of the paper as follows. We give an overview of the approach we use to fit a time trace measured for an SMTJ in Sec. II so that we can reproduce a statistically identical time trace in simulation. In Sec. III we describe the different statistical measures of such voltage-time traces that provide reductions of the data to relevant dynamical quantities. Section IV introduces an overdamped, 1D Langevin model and describes the method that we use to fit the experimental data. Section V compares the Langevin modeling with experiment and shows that the model reproduces the experimental data used to determine the model and exhibits self-consistency with the assumed Fokker-Planck dynamics. In Sec. VI, we analyze standard macrospin models and use a comparison of their allowed probability distributions and drift-diffusion statistics with those of the experimental data to show that they cannot reasonably provide a quantitative model for the data. Finally, in Sec. VII, we discuss possible extensions to our compact modeling approach, orienting future research in ways we view as most useful for supporting the device-circuit-system codesign practices needed to realize large-scale computational systems based on SMTJs.

Experimental details are given in the appendices. Appendix A briefly introduces the fabrication of our experimental device. Appendix B describes the direct current (dc) characterization of the measured device. Appendix C describes the circuitry used to make the high-speed measurements of the superparamagnetic behavior. Appendix D describes some alternate fitting schemes. Appendix E describes the extraction of mean dwell times from time traces.

II. MODEL CONSTRUCTION ROADMAP

Our goal is to develop a data-driven, compact, semi-analytic model that produces dynamics that are statistically identical to those of the voltage-time trace. We refer to the model as data-driven because we use some of the statistical properties of the data to produce the model. In contrast to other, physically motivated models, such as a




FIG. 1. Visual outline of the model generation process developed in this paper. Top row: we start with an experimental device from which a voltage-time trace is extracted under fixed experimental conditions. The histogram and the voltage-dependent diffusion coefficient (second Kramers-Moyal coefficient D_2) are extracted numerically from the time trace. The drift coefficient and dwell-time distributions are also extracted for later use. Middle row: we fit drift and diffusion characteristics for our model to the histogram and diffusion characteristics found in experiment, in a way that compensates for the high-frequency cutoff of the experimental data. Bottom row: we then use these fitted drift and diffusion characteristics to simulate a Langevin equation from which we extract a voltage-time trace and its attendant statistics. Far right: the suite of statistics from the model is validated against the experiment. In future work with more complex fits, one may need to introduce a self-consistent fitting procedure that uses observed error between theory and experiment to inform iterative refinements of the model (dashed line).

macrospin approximation, our models are not derived from approximations based on device physics. Not being tied to approximations of the physics can be an advantage, as it allows the model to capture behavior that is neglected by physically motivated models that may be oversimplified compared with experimental reality.

Figure 1 shows the modeling flowchart for this paper. From the experiment with the SMTJ device we obtain a voltage-time trace. The statistics of that data are used to determine our model, a 1D equation of (directed, overdamped Brownian) motion. This equation of motion describes the state of the system with drift (deterministic) and diffusion (stochastic) terms, with parameterizations that can be calculated directly from the experimental voltage-time trace. We integrate the model forward in time to produce a simulated time trace, do the same statistical analysis that we had performed for the experiment, and compare the experimental and simulated statistics.

Figure 2 shows a representative segment of a voltage-time trace of our experimental device described in Appendix A. The models we use to reproduce this behavior are 1D stochastic differential equations whose drift and diffusion parts are determined by both short-time statistics and long-time statistics. The resulting equation is the Langevin equation

$$\dot{\Phi} = \underbrace{f(\Phi)}_{\text{drift}} + \underbrace{g(\Phi)\eta_{\Phi}(t)}_{\text{diffusion}}, \quad (1)$$

where Φ is the dependant variable (voltage), t the independent variable (time), $f(\Phi)$ the function characterizing the drift, $g(\Phi)$ the function characterizing the diffusion, and

$\eta_{\Phi}(t)$ a Gaussian white noise term [57] with the properties $\langle \eta_{\Phi}(t)\eta_{\Phi}(t') \rangle_t = 2\delta(t' - t)$ and $\langle \eta_{\Phi}(t) \rangle_t = 0$. The drift and diffusion terms of the Langevin equation are related to each other through the voltage-time trace histogram, which gives the long-time average of the stationary state of the system for a fixed current flowing through the SMTJ; see Sec. IV for further details.

The 1D Langevin models are completely determined by fitting the histogram and either the first, $n = 1$, or second,




FIG. 2. Typical bistable voltage fluctuations of an SMTJ in our experimental setup, measured at a sampling rate of $\tau_{\text{exp}} = 50$ ps. Shown is a small subset (a 100-ns window) of the full-time 1 ms trace captured in the experiment. The total number of transitions between the two states in the full voltage-time trace is of the order of 3×10^5 . The dashed line indicates the boundary between the parallel and antiparallel wells. That value is chosen from the local minimum in the distribution of measured voltages given in Fig. 3.

$n = 2$, Kramers-Moyal coefficient

$$D_n(\Phi) = \frac{1}{n!} \lim_{\tau \rightarrow 0} \frac{\langle [\Phi(t + \tau) - \Phi(t)]^n \rangle_t}{\tau}, \quad (2)$$

with τ the sampling time, the angle brackets denoting a time average over the voltage-time trace, and $[\Phi(t + \tau) - \Phi(t)]$ a τ -delayed difference in voltage. The first and second Kramers-Moyal coefficients encode short-time dynamics into our model. The D_1 coefficients capture the drift behavior of the system while D_2 captures the diffusive behavior. Fitting the histogram and one of the Kramers-Moyal coefficients then determines the functional coefficients in the Langevin model [Eq. (1)], fully characterizing it and allowing us to perform simulations that mimic the experimental measurement.

As will be shown in Sec. V and discussed in Sec. VII, this data-driven approach captures multiple statistical and dynamical metrics with high fidelity. Specifically, it captures the Kramers-Moyal coefficients, the histogram, the dwell-time distributions, and the power spectral density (discussed in Appendix C). The agreement for the dwell-time distributions is perhaps the most significant validation of the model because the experimental input to the model does not directly encode these timescales. This ability to recover higher-order statistics is necessary so that we can use the model to design scaled-up circuits in engineering applications.

III. STATISTICS OF TIME TRACES

In this section, we describe in more detail the statistical properties that we use to set up the Langevin model. These properties and others are compared with those of simulations of the resulting model in Sec. V.

The first statistical reduction of the voltage-time trace we use is the histogram of the device state over the entire 1-ms measurement window; see Fig. 3. We assume that the system giving rise to the voltage-time trace is in quasiequilibrium, so the histogram primarily depends on the effective energy and entropy of the system at each voltage and less on the short-time dynamics of the system. The binning resolution used for the histogram and other statistical properties described below along the voltage axis impacts the fit quality and the observed characteristics. If the bins are too small, the statistical uncertainty prevents good fits; if they are too large, the details of the dynamics are obscured. The bin size we use for the histogram is a compromise between these factors at approximately $200 \mu\text{V}$. Each bin captures approximately 50 quantized voltage signals out of the approximately 2^{12} unique levels reported by the oscilloscope.

Though the histograms we observed above can (together with an assumption of Boltzmann statistics) tell us about the effective energy landscape of the system, the detailed




FIG. 3. Histogram of SMTJ state. The black circles give the experimental results for the entire voltage-time trace of 160 ms, part of which is shown in Fig. 2, using bins $200 \mu\text{V}$ wide. The blue rectangles give the results of a simulation, discussed in Sec. V, using a data-driven Langevin model. Note the excellent recovery of the experimental histogram's values in the simulation, demonstrating the capacity of the first-order Langevin model to capture coarse-grained statistics. Statistical error bars are smaller than the symbols. The inset plot shows the ratio between the simulation and experiment histogram values. Note the near-unity agreement for the entire domain save where both histogram's values approach zero. The local minimum near -1 mV is taken as the boundary between the parallel and antiparallel wells.

short-time dynamics generated by thermally induced magnetic fluctuations—which are by construction separate from the conservative forces on the system—remain hidden. To determine their behaviors, we compute conditional moments [58] of the experimental voltage-time trace. The n th-order conditional moment M_n is

$$M_n(\Phi, \tau) = \langle [\Phi(t + \tau) - \Phi(t)]^n \rangle_t, \quad (3)$$

with the angle brackets denoting a time average over an entire trajectory, and $\Phi(t + \tau) - \Phi(t)$ a τ -delayed difference in the system's state. We capture these conditional moments with the same bins as we use for the histogram.

At very small τ , Eq. (3) can be used to approximate the n th-order Kramers-Moyal coefficient [59–61] of the system. Formally, these Kramers-Moyal coefficients are connected to the conditional moments as

$$D_n(\Phi) = \frac{1}{n!} \lim_{\tau \rightarrow 0} \frac{M_n(\Phi, \tau)}{\tau}. \quad (4)$$

The D_1 and D_2 terms, which we refer to as the drift and diffusion terms, are the specific Kramers-Moyal coefficients used in our first-order Langevin model and are the only nonzero Kramers-Moyal coefficients required to describe systems obeying the Fokker-Planck equation [62].

Due to the finite resolution inherent in experimental data, we cannot immediately take the limit as $\tau \rightarrow 0$ in

our calculation of the Kramers-Moyal coefficients. We also find in practice that our experimental timestep was not quite small enough to approximate this limit; subsampling our data to compute the M_n at slightly longer timesteps indicates that we are not in the converged regime, and strongly so for M_2 . As the τ increases from zero, however, finite-time corrections can be used to relate the conditional, time-delayed moments M_n to the underlying Kramers-Moyal coefficients that are needed to construct an analytic Langevin model [63,64]. These corrections up to second order in the time delay τ are

$$M_1 = \tau D_1 + \frac{\tau^2}{2} (D_1 D_1' + D_2 D_1'') + \mathcal{O}(\tau^3), \quad (5)$$

$$M_2 = 2\tau D_2 + \tau^2 (D_1^2 + D_1 D_2' + D_2 D_2'' + 2D_2 D_1') + \mathcal{O}(\tau^3), \quad (6)$$

with a prime denoting differentiation with respect to Φ , and each Kramers-Moyal coefficient and time-delayed moment understood to be a function of Φ . Equations (5) and (6) describe a second-order differential system for the Kramers-Moyal coefficients in terms of calculated M_1 and M_2 ; these equations cannot be solved analytically in general. Our approach to extracting the Kramers-Moyal coefficients, which we require to parameterize our Langevin model, is to choose a parameterized ansatz for D_2 ; that is, we choose some functional form with a number of free parameters. Combining this ansatz with an analytic fit to the histogram induces an ansatz on D_1 via the stationary solution to the Fokker-Planck equation. We then choose the free parameters [see Eq. (10) below] by fitting the right-hand sides of Eqs. (6) [using our ansatz on the right-hand sides of Eq. (5) and Eq. (6)] to the M_2 extracted from our experimentally measured voltage-time trace. We could fit both M_1 and M_2 simultaneously, but finding agreement between M_1 and the predictions of the model when only M_2 is fitted argues for the appropriateness of using the 1D Fokker-Planck equation.

IV. ONE-DIMENSIONAL LANGEVIN MODEL

We note in Sec. VIB that macrospin models fail to capture qualitative metrics—let alone statistical metrics—associated with the experimental data. Yet models that capture the statistics of experimental devices will be required for high-fidelity circuit simulations of SMTJ devices as their attendant technological applications scale. To address this discrepancy, we introduce a first-order Langevin model inspired by existing works on the modeling of fluctuating bistable processes [65,66]. This approach may not be capable of predicting the behavior of uncharacterized devices, but we anticipate it could be used to characterize the devices from a particular manufacturing process and those fit results could be used for circuit simulations of those devices.

We take a data-driven approach to determining the details of the 1D Langevin model. The first step is to find an analytic expression for the experimentally determined probability density ρ_0 . We require an analytic expression because we will need to take its derivative to infer the deterministic forces in the system. To guarantee a positive-definite fit, we assume a Boltzmann distribution and then fit not to the histogram itself but to a dimensionless effective energy $U_{\text{eff}}(\Phi)$ defined so that $\rho_0(\Phi) = (1/Z) \exp(-U_{\text{eff}}(\Phi))$, where Z ensures that $\int d\Phi \rho_0(\Phi) = 1$. A judicious choice of basis function is required to capture the distribution with high fidelity; we scale our data appropriately and then use the Chebyshev polynomials of the first kind $T_n(\Phi)$ which form an orthonormal, complete basis on $[-1, 1]$ and are each bounded between -1 and 1 over this range [67]. Figure 4 shows progressively higher order fits to U_{eff} ; for the rest of the work presented in this paper, we use the $n = 20$ fit for the effective energy and the stationary distribution.




FIG. 4. Fit to the effective energy. (a) Residues (normalized l^2 norm) between the Chebyshev polynomial fits to the data as a function of even order n and (b) sample fits— $n = 4, 10, 20$ —to the experimental effective energy U_{eff} . Error bars on experimental data indicate statistical uncertainties due to the number of counts per histogram bin; however, these uncertainties are smaller than the symbols.

For practical applications the choice of the fit order would be a balance between fidelity to the data and speed of calculation.

The effective energy in Fig. 4, derived from the stationary distribution for the experimental data, exhibits several features that are inconsistent with a macrospin model as we show later in Sec. VI. These include asymmetry for the interior/exterior sides of the parallel (P)/antiparallel (AP) states, disparate P/AP well widths, and exterior boundaries that are significantly rounded compared with the macrospin results. We believe that these discrepancies may be related to a common cause: the way the macrospin model neglects many degrees of freedom that may play an important role in the dynamics. The asymmetry between the parallel and antiparallel alignments could be due to asymmetry between the configurations due to the nonuniform fringing fields from the fixed layer and synthetic antiferromagnet. The tails in the distribution could have several explanations. At a basic level, it could be that the states with the minimum and maximum resistances are not the lowest energy states. One may also note that for an open system (the SMTJ) in contact with a thermal reservoir it is the free energy, rather than internal energy, that is the relevant quantity [48] for making thermodynamic predictions like the probability distributions. The measured distribution integrates over the unmeasured degrees of freedom (those that do not lie along the fixed layer magnetization) so, even if the strictly parallel and strictly antiparallel states *were* energy minima, slightly higher energy states may be significantly more prolific so that the peaks in the marginal distribution along the voltage, Φ , may be shifted from the energy minima.

The desire to capture these features of the histograms and other statistical measures is what motivates us to develop a compact model beyond the macrospin approach. As we have detailed above, our approach is to assume a 1D overdamped Langevin equation driven by white noise. The simplicity of this model makes our task tractable, and in fact, we show in Sec. V that the model gives very good agreement with experiment.

To determine the parameters for the Langevin model, we start with the Fokker-Planck equation, which describes the evolution of the probability density $\rho(\Phi, t)$ as

$$\frac{\partial \rho}{\partial t} = \frac{\partial}{\partial \Phi} \left[-D_1(\Phi)\rho(\Phi, t) + \frac{\partial}{\partial \Phi} D_2(\Phi)\rho(\Phi, t) \right], \quad (7)$$

with Φ a generic state variable (voltage, in the present paper), t the time, $D_1(\Phi, t)$ the drift, and $D_2(\Phi, t)$ the diffusion. The terms $D_1(\Phi, t)$ and $D_2(\Phi, t)$ are precisely the first- and second-order Kramers-Moyal coefficients from Eq. (4). The Kramers-Moyal coefficients may be functions of time in general, but the measured probability distribution ρ_0 is stationary on the timescale of the experiment and thus we also assume $D_n(\Phi, t)$ are not functions of time. In

the steady-state limit $\partial \rho / \partial t = 0$ we have

$$D_1(\Phi) = D'_2(\Phi) - D_2(\Phi) \frac{dU_{\text{eff}}(\Phi)}{d\Phi}, \quad (8)$$

so that D_1 is uniquely determined given the diffusion coefficient and the steady-state distribution of the system. In passing from Eq. (7) to Eq. (8), we used the fact that the right-hand side of Eq. (7) is the divergence of the probability current J . In the steady-state limit, $\partial_\Phi J = 0$, and thus the stationary current J_0 is constant. For the system to be physically bounded, the constant value of J_0 must in fact be zero, which allows us to uniquely solve for $D_1(\Phi)$ as a function of $D_2(\Phi)$ and $\rho_0(\Phi)$.

In a system with additive noise (i.e., $D'_2(\Phi) = 0$) the drift term would be determined through the derivative of the effective energy, in which case D_1 coincides with a typical conservative force, up to prefactors. A system with multiplicative noise (where $D'_2(\Phi) \neq 0$) exhibits mixing between the stochastic terms and the deterministic part; here the nonzero diffusion gradient introduces a stochastic drift term.

The Langevin equation derived from the Fokker-Planck equation is

$$\dot{\Phi} = f(\Phi) + g(\Phi)\eta_\Phi(t) \quad (9)$$

with Φ the voltage state, $f(\Phi)$ the deterministic drift term, $g(\Phi)$ the stochastic diffusion term, and η_Φ is white noise with $\langle \eta_\Phi(t)\eta_\Phi(t') \rangle_t = 2\delta(t' - t)$ and $\langle \eta_\Phi(t) \rangle_t = 0$. We interpret this equation in the Stratonovich sense [68], as one would for the typical construction of the sLLG equation. In the Stratonovich interpretation, the Langevin equation is related to the Fokker-Planck equation's Kramers-Moyal coefficients by $f(\Phi) = D_1(\Phi) - D'_2(\Phi)/2$ and $g(\Phi) = \sqrt{D_2(\Phi)}$. The second term on the right-hand side of $f(\Phi)$ comes from gradients in the diffusion which give rise to an effective drift term independent of the usual drift coefficient $D_1(\Phi)$.

The first step in determining a Langevin model that can reproduce aspects of the experimental voltage-time trace in Fig. 2 is to compute the conditional moments in Eq. (3). These moments are shown in Fig. 5. With these experimental conditional moments, we determine the Kramers-Moyal coefficients by making an ansatz $\tilde{D}_2$ for the diffusion part of the Fokker-Planck equation, which when combined with the stationary distribution is used to obtain an induced form for D_1 via Eq. (8). Measurements of M_2 , shown in Fig. 5, inform our choice for $\tilde{D}_2$. The observation that the data show larger diffusion in the AP well than in the P well leads us to propose

$$\tilde{D}_2(\Phi; \mu) = m\Phi + b, \quad (10)$$

with $\mu = (m, b)$ the fit parameters. In Appendix D we show that fitting with a constant $\tilde{D}_2$ also provides an adequate fit, and introducing more fitting parameters can give




FIG. 5. Time-delayed drift coefficient (a) and diffusion coefficient (b), showing experiment (black circles), Langevin simulation (blue squares), computed at a sampling time $\tau_{\text{samp}} = 50$ ps equivalent to the experimental measurement time, and the analytic, fitted results that were inserted into the model, D_n (solid green line) and $M_n/n\tau$ (dashed red line), where $n = 1$ for the drift term and $n = 2$ for the diffusion term. Error bars on the experimental indicate single-standard-deviation uncertainties in the mean. They are smaller than the plot symbols except at the extreme voltages.

an even better fit. To ensure that D_2 is positive definite we take $D_2 = \lambda \log(1 + \exp(\tilde{D}_2)/\lambda)$, where we take $\lambda = 1 \text{ V}^2/\text{s}$, which does not significantly affect $D_2 > 0$. Combining this ansatz for D_2 with the analytical expressions for D_1 and the stationary distribution ρ_0 completely specifies the model. In the next section, we use this model to run stochastic Langevin equation simulations, that is, numerical integrations of Eq. (9).

We fit the parameters of the ansatz in Eq. (10) following the procedure outlined in Sec. IV. The fit parameters in Eq. (10) are determined by first computing an analytic form for D_1 from the ansatz for D_2 , the fit to the histogram, and Eq. (8). Then an analytic form for M_2 is determined from Eq. (6). The parameters of this last expression are adjusted to fit the experimental data. The agreement between the model and the experimental data is shown in Fig. 5. The error in the drift and diffusion coefficients was calculated from the standard deviation $\sigma_i/\sqrt{N_i}$ of the underlying $\delta\Phi^n(\Phi)$ distribution as a function for

the binned Φ_i levels shown in the figure, where σ is the standard deviation and N_i is the number of counts within the i th Φ_i level. The binning level used for Φ_i here was approximately $295 \mu\text{V}$ to have exactly 64 total bins over the sample. The maximum error shown in the plots is observed on the exterior well boundaries where the number of N_i counts is of the order of hundreds of data points, instead of millions nearer the well's interior.

Both drift and diffusion coefficients (i.e., the induced ansatzes for M_n), agree well with the conditional moments extracted from experiment except at the extreme voltage values where the statistical certainty is poor. Figure 5 shows the functional forms determined for the model. The model is constructed so that the $M_2/2\tau$ is fitted to the experimental data; the consequent agreement between M_1/τ and the experimental data speaks to the appropriateness of the underlying Fokker-Planck model as a description of this system.

V. RESULTS

Our ultimate goal is to develop a stochastic differential equation of motion for SMTJs for use in circuit simulators. Though our current model is restricted to a particular biasing condition, we can still explore how well our model matches the device. We simulate the Langevin equation [Eq. (9)] and compare simulation results both with the experimental data used to fit the model and to selected statistical metrics of the experimental data.

After obtaining the Kramers-Moyal coefficients D_1 and D_2 , we insert them into Eq. (9). We set the integration timestep to a fifth of the experimental sampling rate, $dt = \tau_{\text{exp}}/5 = 10$ ps. Integration is performed using the Euler-Maruyama method. Simulation histogram results are compared with the experimental histogram in Fig. 3, confirming that the first-order, data-driven Langevin model captures this aspect of the experiment well. The histogram displays asymmetrical energy-well probabilities and widths, greater-than-exponential decay for the probability density at exterior data boundaries, and a large energy barrier between the two well peaks. We assert that the reproduction of statistical features such as these is key to developing an engineering-appropriate model for the device. For purely analog circuits in particular—such as those proposed in Refs. [69] and [42]—correct modeling of these distributions may be crucial for capturing emergent statistics of scaled-up circuitry based on superparamagnetic tunnel junctions.

Comparing the conditional moment extracted from experiment and simulation demonstrates that the first-order Langevin model gives dynamics similar to those seen in experiment. Figure 5 confirms excellent agreement for $M_1(\Phi, \tau)$ and $M_2(\Phi, \tau)$ calculated for a time delay equal to the experimental sampling time ($\tau_{\text{samp}} = \tau_{\text{exp}} = 50$ ps).

The statistics that have received perhaps the most attention in the literature are the mean dwell times of the parallel and antiparallel states. These are captured directly by construction in Néel-Brown models, and macrospin models are generally fitted to ensure these dwell times are accurate [40]. In circuits and systems where the SMTJ state is binarized into a two-level telegraph signal, the mean dwell times together with the assumption that the switching events are independent of the past history (the Markov property) completely determine the circuit output. With a voltage threshold chosen to demark the boundary between the P and AP states, the dwell times imbue the data we have already captured in the histogram with a characteristic timescale. Crucially, this timescale information is carried into the model through our fitted ansatz on D_2 . For the purposes of this work we consider a transition between the P and AP states to occur when the voltage crosses the barrier location determined by the local maximum of the energy landscape in Fig. 4. We define a dwell time as the amount of time the system spends in the positive or negative half of state space before crossing that $\partial_\phi U_{\text{eff}} = 0$ threshold into the other half of state space.

The main panels in Fig. 7 show simulation and experimental dwell times for a sampling rate of $\tau_{\text{samp}} = \tau_{\text{exp}}$ on log-log and semilog plots for the two distinct P and AP states. These figures show that the simulation captures the full distribution of dwell times for both states. This agreement for the dwell-time distributions is perhaps the most significant validation of the model because the experimental input to the model does not directly encode these timescales. The model uses just the short-time dynamics captured by the Kramers-Moyal coefficients and the stationary limit of the time-averaged time series in the form of the histogram.

One unusual feature of both experimental and simulated distributions is the crossover from the familiar exponential behavior to approximately $t^{-3/2}$ power-law behavior at short times. We attribute this behavior to fluctuations around the threshold we defined between the states. Recall that we chose the threshold between states to correspond to the local maximum of the effective energy landscape; linearizing about this threshold therefore reveals a locally flat energy landscape where the dynamics are driven entirely by stochastic fluctuations. In other words, the dynamics in this small neighborhood around the threshold are given by a random walk. It is well known that the dwell-time (return-time) distribution of a random walk is proportional to $\sqrt{\tau/t^3}$, where τ is the timestep of the walk [70]. As τ decreases, more and more probability mass accumulates at smaller and smaller t , dominating the dwell-time distribution. This suggests that the transition between the power-law and exponentially distributed dwell times is controlled by the curvature of the energy barrier: the flatter the energy near the threshold, the longer the system can behave as an unbiased random walk.

Intuitively, we can imagine zooming in to a threshold-crossing trajectory sampled at some frequency. If we “zoom in” by increasing our sampling frequency on the same underlying trajectory, the newly sampled points have the potential to introduce additional threshold crossings. Since the spectrum of the white noise process contains all frequencies, we expect to introduce more and more crossing events as we sample at higher and higher frequencies; as the sampling time goes to zero, the number of crossings will diverge.

To see this visually, consider Fig. 6. In the top panel, we plot in gray a simulated trajectory of an Ornstein-Uhlenbeck process with unit mean and volatility, spring constant $\theta = 2$, and initial condition $y = -1$. As it traverses from the initial condition toward its energy minimum at $y = 1$, it crosses the $y = 0$ line, which we regard as a threshold. The question is how many times it crosses this threshold. This question is easily addressed if we have the full trajectory data. But suppose the number of points we have access to on the trajectory is much smaller than the “true” number of points. In Fig. 6, we denote subsampled trajectories with a blue line.




FIG. 6. Progressively increased zoom and sampling rates on a random trajectory generated by an arbitrarily simulated Ornstein-Uhlenbeck process. Top panel: the full-resolution trajectory which has 10^7 sampling points is given in gray. The blue trajectory, subsampled to have only 10^2 points, represents what might one collect experimentally. Red vertical bars denote zero-crossings of the blue trajectory. We zoom in on the gray region of the top panel in the center panel. Center panel: here we locally resample with 10^2 points, an order of magnitude faster sampling than in the top panel. The resampling reveals more zero-crossings than were present at that larger scale. Bottom panel: zooming and resampling by factors of 10 again reveals yet more crossings. Note the qualitative self-similarity of the blue curves in each panel, suggesting that (with infinite underlying resolution) we would continue to reveal more and more crossings with higher and higher sampling rates.

In the top panel of Fig. 6, we observe three crossings of the subsampled trajectory across the threshold. But if we zoom in on the gray region and resample the underlying trajectory there (middle panel), we find that far more crossings become apparent, approximately 11. We zoom in and resample once again on a subdomain of this trajectory to find an even finer resolution (bottom panel), which has yet more transitions. Note that in each case, the size scale of the fluctuations (i.e., the range of the vertical axis) gets smaller and smaller. But the thresholding process, by definition, ignores this amplitude information, and since the white noise driving the system has support at all frequencies, we will always find more transitions by going to higher and higher sampling rates.

In the physical SMTJ system, we expect some physical mechanism to impose an ultraviolet cutoff—but whatever this cutoff may be clearly resides at a higher frequency than our experiment can access. Since this power-law behavior is simply a property of stochastic dynamics around a

threshold, we expect it to be found in other stochastic magnetic systems as well. In Sec. VI, we will find that the behavior is also reproduced in the macrospin model, but note that it is impossible to capture in a simple Néel-Brown model, which is based on a two-state Markov model that assumes exponential distributions of dwell times; such a system has no notion of dynamics in the barrier, which is the essential cause of the power-law effect.

Traditionally, mean dwell times are extracted from experimental dwell-time histograms by examining the slope of the exponential dwell-time distributions on a semilog scale (wherein the distribution appears linear, as in Fig. 11). Because of the power-law behavior at small times in our data, these inverse slopes—which we will call *characteristic dwell times*—are clearly very different than the literal mean dwell times. Yet merely isolating the slope of the curve where it does appear exponential in Fig. 7 does not lead to a physically invariant characteristic dwell time; this slope can change significantly under minor subsampling of the signal, because this changes the probability-dominating power-law part. We elaborate more on the topic of characteristic dwell times in Appendix E.

Finally, we compare the power spectral densities of the experimental signal with those produced by the model. Accurate power spectra are crucial from an electrical engineering context. Unlike the diffusion characteristics and histograms, power spectra are known to exhibit the correct qualitative (Lorentzian) form in macrospin models [32] and in Néel-Brown models [71]. In Appendix C, we show that the power spectral density of the simulation agrees with that of the experiment and both follow the expected behavior at high frequencies.




FIG. 7. Dwell-time distributions as extracted from experiment (black) and simulation (blue) in (a) a log-log plot and (b) a semilog plot. The log-log plot uses a sampling time bin size of 100 ps, while the semilog plot uses a sampling time bin size of 1 ns. The experimental data and simulation both have of the order of 300,000 transitions. The straight-line behavior of both the simulation and experiment for short times in the upper panel indicates a power-law distribution and for long times in the lower panel an exponential distribution.

VI. ANALYSIS OF MACROSPIN MODELING

In the previous sections, we described the modeling of an experimental dataset collected from a device in the laboratory. In service of our engineering-focused modeling goals, our objective was to create compact models that faithfully capture the relevant physics of *that particular device*. A common approach in the literature to modeling SMTJs is through the use of a macrospin model. In the present section, we consider the macrospin model for SMTJs in two separate contexts. First, we show that it cannot reproduce important aspects of the measured properties of our device. Second, we use the macrospin model as experimental data to test the generality of our approach using the Langevin model to fit data. In this context, we show that our Langevin modeling approach can capture the dynamics of an easy-axis macrospin. However, for other devices of current interest, referred to as low-barrier devices [28,31], in which there is a strong easy-plane anisotropy axis and a weaker in-plane easy axis, the Langevin approach can model the drift and diffusion characteristics but is unable to capture the correct dwell-time

distributions of the device. This failure is due to the inertial dynamics that are neglected in the overdamped Langevin approach. Treating such systems will require additional developments as discussed at the end of this section.

A. Macrospin model

We consider here a typical macrospin model with anisotropy. The dynamics are generated by the sLLG equation for dimensionless magnetization $\mathbf{m}$,

$$\dot{\mathbf{m}} = -\tilde{\gamma}_0 \mathbf{m} \times \mathbf{h}_{\text{eff}} - \lambda \mathbf{m} \times (\mathbf{m} \times \mathbf{h}_{\text{eff}}), \quad (11)$$

with $\mathbf{h}_{\text{eff}} = \xi - \partial_{\mathbf{m}} \mathcal{E}/(\mu_0 M_s)$ the effective field with thermal part ξ and M_s the saturation magnetization, $\tilde{\gamma}_0 = \gamma_0/(1 + \alpha^2)$ the renormalized gyromagnetic ratio, $\lambda = \gamma_0 \alpha/(1 + \alpha^2)$ with α the Gilbert damping, and $\gamma_0 = \mu_0 g \mu_B / \hbar$, with μ_0 the vacuum magnetic permeability, g the electron g -factor, μ_B the Bohr magneton, and $\hbar$ the reduced Planck constant. We take ξ to be a spherically symmetric Gaussian noise term with variance $2\Gamma = 2\alpha kT/(\mu_0 \gamma_0 M_s \mathcal{V})$ [45], where $\mathcal{V}$ is the volume, k the Boltzmann constant, and T the temperature. Finally, $E = -K_z m_z^2 - K_x m_x^2$ is the energy density. The first term, with $K_z > 0$, is an easy-axis anisotropy along the direction of the fixed layer magnetization, that is, the component of the magnetization that determines the magnetoresistance. We take m_z as a proxy for the voltage signal across an MTJ at fixed bias; in Fig. 8, for instance, m_z is the vertical axis (compare with Fig. 2). The second term, with $K_x \leq 0$, is a (potential) easy-plane anisotropy. Other contributions, like external fields, other anisotropies, or current-driven torques, could be included but would not change the conclusions we draw below. We assume the Stratonovich interpretation here as we have for all other stochastic differential equations throughout the paper [68].

B. Capturing macrospin physics in a first-order Langevin model

Figure 8 shows typical time traces for a macrospin where $K_z = 5$ kT; in the top panel, $K_x = 0$, and in the bottom panel, $K_x = -10$ kT. Compare with Fig. 2 (or experimentally measured voltage-time trace signals); macrospin simulations show much harder walls at the $m_z = \pm 1$ state than we see in experiment. We will soon see that the diffusion characteristics also look quite different than in the experiment. Nevertheless, in this section, we attempt to capture the physics of a macrospin within an overdamped Langevin model with a single degree of freedom, m_z .

Our experimental results indicate that our real device has two relatively well-defined states, which suggests an easy-axis description in the macrospin limit. If the system has only easy-axis anisotropy in the energy density, we can isolate the equations of motion for the z -component of a macrospin. Rewriting Eq. (11) in terms of m_z and ϕ gives




FIG. 8. Typical magnetization-time traces from the sLLG simulation of the easy-axis (top) and hard-axis with in-plane easy-axis (bottom) macrospin model. In both panels, the simulation was conducted with a timestep of 10 fs, but the trajectories are plotted by subsampling the raw data with a 5 ps timestep (corresponding to the subsampling used in the experimental section). Compare with Fig. 2: we have harder walls at the extremal states in the vertical direction.

an equation for $\dot{\phi}$ that depends on m_z and one for $\dot{m}_z$ that is independent of ϕ ,

$$\dot{m}_z = \frac{\lambda K_z}{\mu_0 M_s \mathcal{V}} (m_z - m_z^3) - \tilde{\gamma}_0 (L_\phi + \alpha L_\theta) \sqrt{1 - m_z^2}, \quad (12)$$

where the two independent white noise processes, L_ϕ and L_θ , each have variance 2Γ and produce thermal fields in the $\hat{\phi}$ and $\hat{\theta}$ directions. We combine these and regard them as a combined white noise process L with variance $2\tilde{\Gamma} = 2\Gamma(1 + \alpha^2) = 2\lambda kT/(\mu_0 M_s \mathcal{V})$. Then it is straightforward to derive the drift and diffusion coefficients,

$$D_1(m_z) = \frac{\lambda m_z}{\mu_0 M_s \mathcal{V}} (K_z(1 - m_z^2) - kT), \quad (13a)$$

$$D_2(m_z) = \frac{\lambda kT}{\mu_0 M_s \mathcal{V}} (1 - m_z^2). \quad (13b)$$

The critical point $K_z = kT$ is a pitchfork bifurcation in the drift coefficient where $m_z = 0$ changes stability and two additional zeros of $D_1(m_z)$ appear. Note that the value

of $D_1(\pm 1)$ is exactly $\mp \tilde{\Gamma}$; the conservative forces do not contribute there.

Equipped with the drift and diffusion coefficients of this process in m_z coordinates, it would be theoretically valid to write down a Langevin equation

$$dm_z = \left[D_1(m_z) - \frac{1}{2} D_2'(m_z) \right] dt + \sqrt{2D_2(m_z)} dW, \quad (14)$$

where W is a Wiener process [72] with unit variance and zero mean. Given Eqs. (13), however, it is clear that this equation would be difficult to simulate numerically. Any finite-timestep integration of this process is liable to step beyond the valid domain $m_z \in [-1, 1]$, leading both to imaginary diffusion and to a drift term that would push the system even further away from the valid regime.

There may well be methods to properly integrate Eq. (14), but for the purposes of this paper we will instead consider our data-driven approach. First, we simulate the system according to the original description [Eq. (11)]. By examination of the conditional moments $M_1 = \langle m_z(t + \tau) - m_z(t) \rangle$ and $M_2 = \langle [m_z(t + \tau) - m_z(t)]^2 \rangle$ of those “experimental data,” we once again choose an ansatz for the D_2 of our synthetic model and fit it to the data via Eq. (10). The histogram of the data again defines an energy function $E(m_z) = -\log \rho(m_z)$ which, together with D_2 , determines D_1 through the equilibrium condition of the Fokker-Planck equation,

$$D_1(m_z) = D_2'(m_z) - D_2(m_z)E'(m_z), \quad (15)$$

just as in Sec. IV. Then we simulate Eq. (14) using these fitted D_1, D_2 which will be chosen to have good behavior at the boundaries.

We numerically integrate the sLLG equation that will provide our dataset in two different ways: first, using only an easy-axis anisotropy $K_z = 5$ kT; second, using a hard-axis anisotropy $K_x = -10$ kT with a weak easy axis $K_z = 5$ kT. In both cases, we take $\alpha = 0.01$, $M_s = 1.5$ MA/m, and $T = 300$ K. We assume a free layer thickness of 1 nm and a diameter of 20 nm, which approximates our experimental device.

We then take these data and apply our procedure from Sec. IV to extract an effective first-order Langevin system. As the structure of the macrospin dynamics is much simpler than the experimental device, we use only a fourth-order polynomial to fit the effective energy function, though we also include two sharp exponential functions $A_{\pm} \exp[\pm 100(x \mp 1)]$ in order to capture the hard-wall boundaries at $m_z = \pm 1$. We then simulate this effective model and compare its histogram with the macrospin simulation data in Fig. 9. Both simulations used a timestep of 10 fs and ran for 10^{11} timesteps (a total simulated time of 1 ms). While the agreement is good in the easy-axis simulation, our models naturally exhibit “leak” of




FIG. 9. Comparison of histograms between our sLLG simulation of the easy-axis macrospin model and the first-order Langevin model that was fitted to the sLLG simulation statistics, together with the analytic solution for these distributions (dashed line). The top panel shows the easy-axis macrospin, the bottom panel a hard-axis macrospin with an in-plane easy axis.

probability beyond $[-1, 1]$. Both models have analytically tractable histograms with which the macrospin and Langevin simulations well agree.

Theoretically, the structure of the multiplicative noise in the sLLG equation (which manifests here as D_2) arises entirely from the geometric considerations of $\mathbf{m}$'s normalized amplitude [47]; the thermally driven part of the sLLG equation in Cartesian coordinates has coefficients $\epsilon_{ijk}m_k + \alpha(\delta_{ij} - m_i m_j)$ irrespective of the effective field. The agreement between M_2 from all simulations with the known $1 - m_z^2$ form of D_2 in the macrospin model is excellent; they are essentially indistinguishable except for very tiny deviations at $m_z = \pm 1$. We conclude that a simple macrospin model cannot capture the complex diffusion characteristics such as those observed in Fig. 5(b), since the functional form of D_2 is fixed by construction. Figure 10 shows the first conditional moments for these sLLG simulations compared with their fitted Langevin models. These also have nice agreement, though the convergence of the noise is slower than for M_2 and the boundary conditions more artificially imposed.




FIG. 10. Comparisons of $M_1(\tau)/\tau$ between our sLLG simulation of the easy-axis macrospin model and the first-order Langevin model that was fitted to the sLLG simulation statistics. The top panel shows the easy-axis macrospin, the bottom panel a hard-axis macrospin with an in-plane easy axis. In both cases, $\tau = 10$ fs.

In both the easy-axis and hard-axis models, the *literal* mean dwell times agree very well between macrospin and Langevin models: the easy-axis mean dwell times are 0.624 ns (macrospin) and 0.623 ns (Langevin); the hard-axis mean dwell times are 0.730 ns (macrospin) and 0.731 ns (Langevin). We also consider the distribution of dwell times, that is, the probability that any particular dwell event has a given dwell time. These agree very well between macrospin and Langevin models in the easy-axis system (Fig. 11, top panel). Although we have elaborated above and in Appendix E on the potential unphysicality of characteristic dwell times extracted from the slopes of these curves, they can still offer a comparison between macrospin and Langevin models. Using a linear fit to extract the slope of the exponential distribution from the semilog plot, we obtain “characteristic dwell times” of 121 ns and 120 ns for both the macrospin and Langevin simulations, respectively. In the hard-axis system, despite the remarkable agreement of the mean dwell times, the dwell-time distributions are considerably different (bottom panel). The “characteristic dwell time” of the Langevin model is 293 ns, while the “characteristic dwell time” of




FIG. 11. Long-time behavior of the dwell-time histograms for the easy-axis macrospin (top) and hard-axis macrospin with an in-plane easy axis (bottom). As in Fig. 12, the model fit to the hard-axis macrospin has poor agreement, but the easy-axis system matches well. Note that these data have been rebinned to a 1-GHz sampling frequency compared with Fig. 12; on this timescale, the coherent switching dynamics of the easy-plane oscillator are washed out.

the macrospin model is only 27 ns—an order-of-magnitude difference.

The top panel of Fig. 12 zooms into the dwell-time distribution of the hard-axis macrospin simulation for very short times. Both the easy- and hard-axis macrospin models exhibit the same $t^{-3/2}$ power-law behavior [73] at the smallest time steps (not shown) as was observed in the experimental device [Fig. 7(a)]. At intermediate but subnanosecond dwell times, however, the hard-axis system shows a distinctive set of peaks. We associate these peaks with coherent precession events around multiple cycles of the easy plane. The bottom panel of Fig. 12 shows an example of a “transition” we found in a trajectory generated from the hard-axis macrospin simulations; the ringing here artificially creates “switching events” around twice some resonant frequency of the macrospin. Such dynamics rely crucially on second-order behavior, with energy being stored in the out-of-plane angle serving as a source of inertia. Since these dynamics strongly affect the switching times but cannot possibly be modeled in an 1D overdamped Langevin equation, it is unsurprising that our model fails to capture these particular statistics.




FIG. 12. Top: short-time behavior of the dwell-time histogram for the hard-axis macrospin with an in-plane easy axis. The central peak in the distribution is not captured by the fitted Langevin model; abnormally frequent dwell times around 250 ps (and, to a lesser degree, around 750 ps) are attributed to coherent in-plane switching cycles. Bottom: a high-resolution (10-fs-timestep) time trace of a transition in the hard-axis macrospin model that exhibits coherent ringing effects on subnanosecond timescales leading to approximately 0.2-ns dwell events, commensurate with the central peak in the top panel.

The inclusion of memory and higher-order differential terms would, therefore, seem a necessary next step to capture higher-fidelity models, especially given the high interest in low-barrier easy-plane devices. Representation of such systems as multiple coupled first-order Langevin models seems the natural direction of study, but this will come hand-in-hand with a quadratic increase in the number of free parameters. Fitting the increased set of parameters will require extracting more statistical information from the time traces. In the simple case we consider here, the first-order Kramers-Moyal coefficient is not needed, nor are the dwell-time distributions. In addition to these, higher-order or time-lagged Kramers-Moyal coefficients could be extracted and used. It is not implausible that the inverse problem (i.e., divining the model directly from the experimental data without having to run simulations to compare) may become intractable. One solution

might be to use an iterative, self-consistent forward-fitting approach, where trial models are refined by comparing the statistics extracted from simulations with those of experiments and refining the parameters until the comparison is satisfactory. While finding the appropriate model will become much more complicated, using the resulting model will not. This type of approach is represented by the dashed arrow in Fig. 1, and may be a topic of future research.

VII. DISCUSSION

While the sLLG equation for a macrospin is physically motivated and contains terms directly identifiable with distinct physics, it cannot capture certain statistical metrics associated with our experimental device. This is partly due to the sLLG model's drift and diffusion coefficients; fitting the former to the observed drift would require including sufficiently "data-driven" terms as to lose the simple plausibility of the model, while directly fitting the diffusion is simply impossible in general. Indeed, even single-domain MTJs may show diffusion properties that do not generally agree with the macrospin. Theory indicates that the Gilbert damping (an overall prefactor for D_2) in common magnetic materials has strong directional dependence [74], while recent experiments on SMTJ switching suggests that the Suhl instability may create dynamically variable damping mechanisms during single-domain reversal [46]. These considerations suggest that we must move beyond the simple macrospin model to accurately capture SMTJ physics in compact models in general, and the data-driven approach explored in this paper is our initial attempt at this enterprise.

We show that our data-driven approach captures multiple statistical and dynamical metrics with high fidelity in an experimentally measured SMTJ device with perpendicular anisotropy. Specifically, we were able to capture the Kramers-Moyal coefficients (Fig. 5), the stationary distribution ρ_0 by construction (Fig. 3), and the dwell-time distributions (Fig. 7). The first-order Langevin model outperforms traditional sLLG macrospin simulations in capturing coarse-grained statistics of the experimental data. We note that perpendicular devices of this type are of strong applied interest, as the fabrication of their thermally stable kin are already integrated into commercial foundry processes, and challenges with device-to-device variation are expected to be easy to address in this geometry compared with in-plane magnets (so-called low-barrier magnets) [75].

We have also simulated macrospin models to demonstrate that the behaviors they exhibit are not commensurate with the experimental observations that we capture with our Langevin model. Of course, the macrospin model *is*, ultimately, a Langevin equation. So long as one is willing

to accept a macrospin's finite domain, it should be possible in principle to deduce complicated, potentially pathological forms of the effective field and (state-dependent) Gilbert damping functions that would indeed mimic the behavior we observe in experiment. But to do so would be to lose the physical motivation of the macrospin model; we contend that achieving this would amount to using a data-driven approach similar to what we present here, simply in spherical rather than Euclidean geometry.

There are several refinements of our approach that are necessary to successfully design scaled-up circuits that operate over ranges of input parameters. One, which is referred to at the end of the previous section, is to extend the model to include memory effects and incorporate additional statistical properties of the device into the fitting procedure. A second is to extend the model to varying currents through the device. In the present results, the current through the device was kept constant, but for a model to be useful in simulation it must extrapolate to an extensive current range. Additionally, it is likely that the devices will have a range of properties and cannot all be described with a single model. If the distributions of device properties are wide enough, it will be necessary to fit a distribution of device properties with a distribution of model parameters. Ultimately, the needs of compact models will be dictated by advancements in circuits and systems that use SMTJs as critical computational elements. Progress on the front of applied physics and engineering in this regard must occur in tandem with physics-level device modeling to ensure that large-scale circuit simulation remains rooted in physical reality.

ACKNOWLEDGMENTS

The authors are indebted to Di Xiao for his useful notes on stochastic processes and to Steve Moxim for his help with the experimental setup. We also thank William Rippard and Robert McMichael for their invaluable comments on the manuscript, and we again thank Robert McMichael for his advice on microwave theory. This work was funded by the National Institute of Standards and Technology, National Science Foundation, and Agence nationale de la recherche. S.G., T.A., L.P., D.P.L., and A.M. acknowledge support under NSF Grant No. CCF-CISE-ANR-FET-2121957. A.M. acknowledges support under the NIST Cooperative Research Agreement Award No. 70NANB14H209 through the University of Maryland. P.T. and U.E. acknowledge support under the ANR StochNet Project Award No. ANR-21-CE94-0002-01. The authors acknowledge J. Langer and J. Wrona from Singulus Technologies for the MTJ stack deposition and N. Lamard, R. Sousa, L. Prejbeanu, and the Upstream Technological platform PTA, Grenoble, France for the device nanofabrication.

APPENDIX A: EXPERIMENTAL SAMPLES

To serve as a basis for our modeling efforts we collected experimental time-voltage traces from an SMTJ. The device's material stack is Si (base)/SiO₂/TaN/[Co(0.5)/Pt(0.2)]₆/Ru(0.8)/[Co(0.6)/Pt(0.2)]₃/Ta(0.2)/Co(0.9)/W(0.25)/CoFeB(1)/MgO(0.8)/CoFeB(1.4)/W(0.3)/CoFeB(0.5)/MgO(0.75)/Ta(150)/Ru(8), where the numbers in parentheses refer to layer thicknesses in nanometers and the subscripts on square brackets show bilayer repetitions. The device is roughly circular in cross section and has a diameter of about 20 nm [76,77]. The materials in this stack produce a perpendicular anisotropy [78] that largely cancels the demagnetization field, leaving only a small energy barrier in the final device. The device thus shows superparamagnetic behavior at a small range of applied currents and applied field. The biasing scheme in this paper is described in Appendix C. An external magnetic field of 124.5 mT was applied at an 18° angle to the device's normal axis. This tilt accelerates the occurrence of stochastic reversals and allows for operation at lower bias voltages and currents.

APPENDIX B: EXPERIMENTAL DIRECT CURRENT CHARACTERIZATION

We first characterized the dc properties of the SMTJ by performing static dc measurements using a varied dc applied voltage and magnetic field. This initial measurement provides the operating voltage and field conditions needed to find superparamagnetic regimes in the fabricated devices.

To accurately measure the dc resistance of the fabricated SMTJs, a constant dc voltage V_0 was applied to the SMTJ through a static series resistance $R_0 = 3 \text{ k}\Omega$. For real-time monitoring and precise measurement of the device's dc resistance, an oscilloscope with an internal impedance of $Z_0 = 1 \text{ M}\Omega$ was connected in parallel to the SMTJ. The oscilloscope, used as a dc voltmeter, outputs an average voltage that is used to evaluate the dc resistance of the SMTJ. At the $V_0 = 505 \text{ mV}$ operating voltage used to measure the time trace analyzed in this paper, the parallel (P) and antiparallel (AP) resistances are 2500Ω and 3400Ω , respectively.

We measured the evolution of the dc resistance in response to a perpendicular magnetic field, which was swept back and forth while maintaining the constant applied voltage during the field sweep. The obtained magnetoresistance loops (resistance versus field curves) are summarized in Fig. 13, presenting them as a dc resistance phase diagram.

As shown in the diagram, positive (negative) fields tend to stabilize the P (AP) state. Similarly, positive (negative) voltages increase (decrease) the stability range of the P (AP) state. For an intermediate set of fields and voltages, a “bistable” region corresponding to configurations where




FIG. 13. Direct current resistance phase diagram of the investigated SMTJ device. The chart summarizes the dc resistance versus field (x axis) curves obtained at different SMTJ voltages (y axis). The red (blue) region corresponds to antiparallel (AP) [parallel (P)] configurations. The central green region illustrates a hysteretic bistable configuration where both AP and P states are possible depending on the variation direction of the applied field. A reduction in the dimension of the bistable region at large voltages coincides with the emergence of a superparamagnetic regime in the device.

both AP and P states can be stabilized is depicted in green. In a traditional resistance versus field loop, that region can be understood as the opening of the hysteresis. For small voltage values, roughly close to zero, the bistable region occupies a roughly -150 mT hysteresis opening, which coincides with the behavior of a thermally stable MTJ, typically suitable for nonvolatile memory applications (e.g., magnetic random access memory).

The size of the bistable region can be reduced by applying positive (negative) voltages, which in our case is interpreted as a stabilization of the P (AP) state through spin-transfer torque. In this line of thinking, the slope of the top-right and bottom-left boundaries of the bistable region can be interpreted as the efficiency of the spin-transfer torque. For large voltages, the boundaries present a curvature (quadratic or higher-order contribution in voltage), which could be a partial consequence of Joule heating effects in perpendicular MTJs [79].

For conditions close to the zero field, the bistable region seems to progressively disappear with voltage, at least for negative voltages larger than -400 mV, leading to a superparamagnetic regime. In our case, such a set of dc field and voltage conditions is typically used to observe superparamagnetic behavior, which is confirmed by voltage-time trace measurements in real time presenting random telegraph noise.

APPENDIX C: HIGH-BANDWIDTH MEASUREMENT CIRCUIT

To precisely capture random telegraph noise on the nanosecond scale, we engineered a radio frequency measurement setup optimized for identifying transitions with dwell times below $10 \mu\text{s}$ as shown in Fig. 14(a). The setup uses a wideband bias-tee in order to isolate the high-frequency noise from other electrical components interfering with our signal while allowing the fast, alternating current (ac) fluctuations of our devices to be captured effectively. It is specified to have an operating frequency range from 80 kHz to 26 GHz . The dc-only terminal of the bias-tee is connected to a voltage source through a $3 \text{ k}\Omega$ source resistor, while the rf+dc terminal is connected to the SMTJ. The AC-only terminal is connected to a 5-cm -long coaxial cable which is terminated into a wide-bandwidth amplifier. The characteristic impedance of the cable and the termination into the amplifier is 50Ω . The amplifier is specified to have a gain of 4 and an operating frequency range from 50 kHz to 40 GHz , and is powered by a filtered source to minimize the effects of noise. The output of the amplifier is impedance-matched to a $50\text{-}\Omega$ coaxial cable




FIG. 14. (a) A schematic diagram of the electrical circuit designed for the detection of resistance fluctuations between parallel (P) and antiparallel (AP) states on submicrosecond timescales. This setup is specifically tailored for measuring the nanosecond-scale random telegraph noise generated by SMTJs. (b) Schematic illustration of the trajectory-averaging process we use to analyze transition dynamics. The gray lines represent a small subset of (the initial timesteps of) the positive voltage sub-trajectories. Error bars on the blue mean curve indicate a standard deviation of uncertainty on the mean.

which terminates into the 50- Ω input of the oscilloscope. The oscilloscope has a frequency range of 10 GHz and a sampling rate of 20 billion samples per second.

The rf path of the measurement setup has a nominal complex impedance of 50 Ω , which helps in achieving impedance matching with the oscilloscope, thus ensuring the efficient transmission of rf signals from the SMTJ. This is effective even when the cumulative impedance of the SMTJ and the static resistor diverge from the 50 Ω target. Within such a setup, SMTJ resistance changes lead to a fluctuating current, which acts as the primary stimulus for the entire rf circuitry. Nonetheless, this configuration significantly reduces the voltage amplitude of the outgoing rf signal observed on the oscilloscope.

1. Temporal sensitivity

After implementing the measurement circuit outlined above, the effective RC time constant of the circuit is decreased to below 0.5 ns, as illustrated in Fig. 14(b). This adjustment facilitates the electrical detection of phenomena within the SMTJs occurring on timescales from submicroseconds to nanoseconds. Of particular interest is how the electronic setup affects our measurement of the transitions between P and AP states. This is difficult to interrogate for single trajectories due to the intrinsic noise of the system, so we use the following procedure to compute averaged switching events. We consider all positive-voltage (negative-voltage) subtrajectories of our time trace, representing stretches of time that the device spent purely in the AP (P) state. We discard subtrajectories shorter than 20 ns to ensure that we are studying events that reliably transition, sit in a metastable state, and then transition back. This filtering leaves us with 3263 positive-voltage subtrajectories and 5550 negative-voltage subtrajectories.

From each of these subtrajectories, we extract their initial (final) 10 ns stretches and average these across all subtrajectories, creating averaged trajectories into (out of) the two states. The averaging process on the initial-stretch trajectories for the P and AP states is depicted in Fig. 14(b) to give a sense of the underlying distributions. We repeat this entire averaging procedure for the voltage-time trace generated by our Langevin simulation from Sec. V so that we may compare theory and experiment.

We glue the averaged state-entering and state-exiting trajectories together to get a sense of average transition events. We zoom in on the 2-ns window around barrier crossings to examine crossing dynamics in Fig. 15. There are some apparently coherent oscillations in the experimental trajectories before and after barrier crossings. These are not captured by the Langevin model in simulation—an unsurprising result given that the Langevin model is inertia-free by construction. It is not straightforward at this point to tell whether these oscillations




FIG. 15. Blue curves show the experimental subtrajectory averages around a transition; red curves show the simulation subtrajectory averages. Note that at 1 ns the simulation curve is larger in amplitude than the experimental curve, but by the end of the trajectory (i.e., at -1 ns) the experimental amplitude exceeds the simulated amplitude. The dashed black line, which approximately captures the simulation behavior, is $V_0 t / \sqrt{\tau^2 + (0.25 \text{ ns})^2}$ for $V_0 = 4$ mV. The early rise times and asymptotic behavior of the experiment are captured by the dashed gray line, $V_0 \tanh[t/(0.15 \text{ ns})]$.

represent magnetic physics or artifacts of the experimental setup. We note that the simulation curve was difficult to fit with functions of exponential character (like hyperbolic tangent); here we fit it with a curve of the form $t/\sqrt{\tau^2 + t^2}$. We speculate that, since the Langevin model is given by diffusion motion *per se* available to make a sensible exponential-type fit, but again, more investigation is needed. The timescales reported in Fig. 15 are close to the limits of the oscilloscope. Accurately observing even faster dynamics necessitates further enhancements in the electrical circuit design, including the use of oscilloscopes and circuits that support a broader frequency bandwidth.

Finally, in Fig. 16, we zoom out and show these averaged trajectories for a span of 20 ns. We note that at about 6 ns post transition in the experimental data, there is a notable voltage shift in both the P and AP wells (though it is more visible in the P well). Since these curves are averaged over thousands of subtrajectories, it seems unlikely that these correspond to device physics. We attribute them to signal reflection in the cable connecting our oscilloscope to the amplifier.

2. Frequency response

This paper is focused on a temporal analysis of the behavior of the device. Analysis in the frequency domain is complicated by the dependence of the behavior on the state of the SMTJ as it crosses the barrier. Nonetheless, it




FIG. 16. Zoomed-out version of Fig. 12. A notable artifact in the signal is present at 6 ns post transition. Note that at long times (i.e., the left half of the plot) the simulated signal stays flat at around ± 4 mV while the experiment signal rises to almost ± 4.5 mV.

can be useful to look in the frequency domain in the additive noise limit (where no state dependence of the diffusion processes exists) to understand the effect of measurement noise on the analysis. First, we note that measurements of electronic noise in this measurement circuit indicate that it only makes minor contributions to the voltage-space histograms. The electrical noise is characterized by replacing the SMTJ with a 3-k Ω resistor giving measured noise with a standard deviation of (256 ± 1) μ V. Meanwhile, the full widths at half maximum of the P and AP wells found in our data are approximately 2 mV and 4 mV, respectively. Since these widths add in quadrature, we expect the effects of electrical noise to contribute to less than 5% of the distributional full widths at half maximum around each well. In what follows, we use a voltage-time trace collected with this 3-k Ω resistor to represent the effective noise floor of the measurement apparatus.

In Fig. 17, we note that power spectral density [80] of our model simulations are well matched to the experiment to which it was fitted. These spectra are consistent with the expected Lorentzian behavior of an asymmetric random telegraph noise model [71]. Figure 17 shows the 10-GHz cutoff frequency of the oscilloscope. Note that the slight upward curvature of the simulation curve at high frequency is an expected deviation from the continuous-time behavior of any discrete-timestep simulation of the Langevin equation (i.e., any autoregressive process of order 1).

The power spectral densities make two things clear. First, although we do measure slightly above the cutoff frequency of the scope, the loss of total signal power resulting from this choice is minimal. Because of the $1/f^2$ behavior, the spectral contributions here would be negligible compared with where we do have meaningful signal.




FIG. 17. Estimated power spectral density (PSD) of the experimental time trace compared with the noise floor of the measurement apparatus and the output of the simulated Langevin model described in Sec. V. The dashed line shows the slope of $1/f^2$ behavior, which would be typical of the spectral density expected from asymmetric random telegraph noise (i.e., a Néel-Brown model) beyond its Lorentzian low-frequency cutoff.

Estimating from the plot (and assuming a continuation of the Lorentzian tail), the total missing signal power from 10 GHz out to infinity is approximately 3×10^{-8} V². The integrated power of such a Lorentzian integrated from zero frequency up to the cutoff is three orders of magnitude larger, approximately 2×10^{-5} V².

Of course, the quantity we extract from the data that is most sensitive to this noise power is M_2 , which is our direct characterization of the noise from experiment. Together with the histograms, M_2 is a central quantity on which all our fitting is based, so its fidelity is crucial to our process. One may ask how not only the high-frequency truncation but also the *low-frequency* cutoff from the bias-tee may affect this characterization. As a first approximation, we can consider the case of additive noise (in Appendix D we show that, though modeling our device with additive noise loses some of the features present in the data, it still does a reasonably good job). In this case, one can expand the signal $\Phi(t)$ in the definition of M_2 [Eq. (3)] in its Fourier components, and after some simplification one finds that

$$M_2 = \frac{2}{\pi} \int_0^{\pi/\tau} S(\omega)[1 - \cos(\omega\tau)] d\omega, \quad (C1)$$

where $S(\omega)$ is the power spectral density of the signal measured with timestep τ ; note that the upper integration limit is the (angular) Nyquist frequency. This expression for M_2 suggests that while the maximal contribution is nominally from the Nyquist frequency itself, there are still strong contributions at half the Nyquist frequency (where $\cos \omega\tau = 0$) and lower.

The infrared cutoff of about 0.1 MHz to 1 MHz imposed by the bias-tee corresponds to $(1 - \cos \omega\tau)$ of about $5 \times$




FIG. 18. The integrand of Eq. (C1) for the experimental data and for the noise analysis data of the measurement setup without the SMTJ. At all frequencies, the experimental curve equals or exceeds the diffusion density found in the noise-floor characterization.

10^{-11} to 5×10^{-9} , respectively. This suppression is sufficiently strong that any signal power lost from the sub-megahertz range would never meaningfully contribute to M_2 even if we could capture it.

More problematic, however, is the ultraviolet cutoff. Figure 18 plots the integrand of Eq. (C1) for the experimental data and the noise-floor characterization. Because of the convergence of the noise floor with the signal at the cutoff frequency (Fig. 17), the contribution of the electrical noise to M_2 becomes significant there; all in all, the total M_2 from the noise floor derived from integrating Eq. (C1) amounts to 20% of the same integral for the experiment. Unlike the histograms, these quantities do not add in quadrature. While the majority of our M_2 signal is indeed device physics, it is also clear that the effect of the electrical noise is far from negligible.

From a theoretical standpoint, however, accounting for the electrical noise may be straightforward. Assuming the magnetic fluctuations have zero mean and are independent of the electrical noise, Eq. (3) factors to give $M_2 = M_2^{\text{device}} + M_2^{\text{electrical}}$. The measured M_2 for the noise characterization could then be subtracted out directly before beginning the fitting process for the theoretical model. Verifying this model against the experiment would then involve the additional step of injecting modeled electrical noise on top of the signal from the Langevin simulation in order to make a fair comparison. Though this procedure could be used to mitigate the effect of experimental artifacts on the model, we omit such a procedure from the main text in the interest of clarity.

APPENDIX D: OTHER PARAMETERIZATIONS OF $\tilde{D}_2$

The ansatz chosen in Sec. IV only has two fit parameters that can be optimized to fit the experiment's calculated

$M_2/2\tau$ Kramers-Moyal coefficient. This section describes the results using a higher and lower number of fitting parameters for the $\tilde{D}_2$ ansatz.

The higher-order fit differs from the linear fit by including a Gaussian that seeks to capture the difference between the Langevin model's calculated $M_2/2\tau$ and the experiment's calculated $M_2/2\tau$ seen in Fig. 5. The analytic form for this fit is $\tilde{D}_2(\Phi; \mu) = m\Phi + b + a e^{(\Phi - \mu_0)/2\sigma^2}$. Note that we ensure D_2 remains positive, as discussed in the main text. The lower-order fit assumes a constant $D_2 \rightarrow D_2(\Phi; \mu) = b$.

Figure 19 compares the calculated $M_2/2\tau$ coefficients from experiment and simulations using the Gaussian plus linear fit, the linear fit, and the constant fit for $\tilde{D}_2$. With progressively higher-order fits, the simulation fidelity increases, indicating that the actual noise term in the SMTJ device has a more complicated structure than these simple analytic models. However, model fidelity to the empirical




FIG. 19. Experimental and Langevin simulation results for the time-delayed, drift panel (a) and diffusion panel (b) coefficients computed at a sampling time $\tau_{\text{samp}} = 50$ ps equivalent to the experimental measurement time. The various symbols are as follows: black circles show experimental data; blue squares show first-order, data-driven Langevin simulations with a Gaussian plus linear $\tilde{D}_2$; yellow circles show first-order, data-driven Langevin simulations with a linear $\tilde{D}_2$ as reproduced from the main paper body; green triangles show first-order, data-driven Langevin simulations with a constant $\tilde{D}_2$.

data is remarkable and demonstrates the flexibility and robustness of our data-driven method to capture gross statistics of the SMTJ device.

In addition to comparing calculated Kramers-Moyal coefficients, we also compare the dwell-time distributions of each model. Figure 20 shows the dwell-time distributions in the P state for the experiment and simulations measured using 2×10^6 points. This number of points is smaller than used in Fig. 7, which is why there is a difference between the experimental curves in the two figures. This was done to ensure magnitude agreement in the dwell-time distribution between the experiment and simulations,




FIG. 20. Log-log panel (a) and semilog panel (b) dwell-time distribution for the P state τ_P as calculated in experiment and produced by simulation: black circles show experimental data; the blue dashed line shows the first-order, data-driven Langevin simulations with a Gaussian plus linear $\tilde{D}_2$; the yellow dotted line shows the first-order, data-driven Langevin simulations with a linear $\tilde{D}_2$ as reproduced from the main paper body; and the green dot-dashed line shows the first-order, data-driven Langevin simulations with a constant $\tilde{D}_2$. The linelike behavior of the simulation and experiment in the upper plot at short times indicated a power-law dwell-time distribution in that regime. In the lower plot, the linelike behavior of the simulation and experiment for longer times indicates an exponential dwell-time distribution in that regime. The same sampling time τ_{samp} is used as in Fig. 7, but the overall dwell-time distribution is only measured over the first 2,000,000 measurement points of the experiment.

which were limited in total length because of computational expense. The dwell-time distribution is unchanged for these different assumed fits for D_2 .

APPENDIX E: CHARACTERISTIC DWELL TIMES

In the discussion of Sec. V, we described how the short-time power-law behavior of the dwell-time histograms in both experiment its fitted simulations [Fig. 7] and, later, in simulations of the macrospin model [Fig. 11] can be attributed to the scale-free, fractallike nature of random walks around the switching threshold at high sampling frequencies. They are not mere artifacts of the experiment, but additional structure that is usually neglected in dwell-time analysis of magnetic devices, and structure that we expect to similarly arise in any magnetic nanodevice.

Calculating the mean dwell times from the entire distributions gives very small dwell times dominated by these short-time power-law behaviors. Figure 21 shows good agreement between the mean dwell times (τ) extracted from the simulation and from the experiment. Even when subsampled, the simulation and experimental average dwell times agree moderately well. Recall that D_2 , which is fitted to M_2 via finite-time corrections, holds all this timescale information in the model. The mean dwell time increases as the subsampling time increases because more and more short-time events are eliminated from the distribution.

Generally speaking, one expects that if a dwell time is relevant for a particular application, it is likely the characteristic dwell time of the exponential distribution that is most meaningful. If we ignore the power-law behavior of the very short dwell times, we can extract these from the slope of the exponential in the semilog plot in Fig. 7(b). The resulting fits give $\tau_{P,\text{exp}} = 7.10$ ns, $\tau_{AP,\text{exp}} = 11.5$ ns, $\tau_{P,\text{sim}} = 5.95$ ns, and $\tau_{AP,\text{sim}} = 9.51$ ns. The one-standard-deviation ranges for these slopes are $\tau_{P,\text{exp}} \pm \sigma =$




FIG. 21. Mean dwell times of experiment and simulation as a function of sampling time.

[6.82, 7.39] ns, $\tau_{\text{AP,exp}} \pm \sigma = [11.23, 11.85]$ ns, $\tau_{\text{P,sim}} \pm \sigma = [5.79, 6.12]$ ns, and $\tau_{\text{AP,sim}} \pm \sigma = [9.33, 9.69]$ ns. The agreement is good even though both simulation dwell-time slopes slightly underestimate the corresponding experimental values.

It is important to note that different fitting schemes will give different results that lie outside these ranges, which are strictly determined by the statistics of the points that have been fitted. Indeed, it is not clear that these extracted times represent anything of experimental relevance. By a fine-tuned subsampling or RC filtering of the data, we can artificially flatten the power-law behavior so that the resulting distribution is exponential. Using this procedure to obtain a plausibly exponential distribution gives $\tau_{\text{P,sim}} \approx 25$ ns and $\tau_{\text{AP,sim}} \approx 30$ ns. Going back to Fig. 2 qualitatively illustrates this effect. Keeping all short-time events gives on the order of 35 crossings of the zero line, which would correspond to mean dwell times of approximately 3 ns. On the other hand, if the rapid fluctuations were filtered out, there would be approximately six transitions for a mean dwell time of about 17 ns.

-
- [1] J. Grollier, D. Querlioz, and M. D. Stiles, Spintronic nanodevices for bioinspired computing, *Proc. IEEE* **104**, 2024 (2016).
- [2] K. Roy, A. Sengupta, and Y. Shim, Perspective: Stochastic magnetic devices for cognitive computing, *J. Appl. Phys.* **123**, 210901 (2018).
- [3] J. Grollier, D. Querlioz, K. Y. Camsari, K. Everschor-Sitte, S. Fukami, and M. D. Stiles, Neuromorphic spintronics, *Nat. Electron.* **3**, 360 (2020).
- [4] T. Blachowicz and A. Ehrmann, Magnetic elements for neuromorphic computing, *Molecules* **25**, 2550 (2020).
- [5] B. R. Zink, Y. Lv, and J.-P. Wang, Review of magnetic tunnel junctions for stochastic computing, *IEEE J. Explor. Solid-State Comput. Devices Circuits* **8**, 173 (2022).
- [6] A. Fukushima, T. Seki, K. Yakushiji, H. Kubota, H. Imamura, S. Yuasa, and K. Ando, Spin dice: A scalable truly random number generator based on spintronics, *Appl. Phys. Express* **7**, 083001 (2014).
- [7] D. Vodenicarevic, N. Locatelli, A. Mizrahi, J. S. Friedman, A. F. Vincent, M. Romera, A. Fukushima, K. Yakushiji, H. Kubota, S. Yuasa, S. Tiwari, J. Grollier, and D. Querlioz, Low-energy truly random number generation with superparamagnetic tunnel junctions for unconventional computing, *Phys. Rev. Appl.* **8**, 054045 (2017).
- [8] B. Parks, M. Bapna, J. Igboke, H. Almasi, W. Wang, and S. A. Majetich, Superparamagnetic perpendicular magnetic tunnel junctions for true random number generators, *AIP Adv.* **8**, 055903 (2018).
- [9] Y. Lv, B. R. Zink, and J.-P. Wang, Bipolar random spike and bipolar random number generation by two magnetic tunnel junctions, *IEEE Trans. Electron Devices* **69**, 1582 (2022).
- [10] M. W. Daniels, A. Madhavan, P. Talatchian, A. Mizrahi, and M. D. Stiles, Energy-efficient stochastic computing with superparamagnetic tunnel junctions, *Phys. Rev. Appl.* **13**, 034016 (2020).
- [11] B. R. Zink, Y. Lv, M. Zabihi, H. Cilasun, S. S. Sapatnekar, U. R. Karpuzcu, M. D. Riedel, and J.-P. Wang, A stochastic computing scheme of embedding random bit generation and processing in computational random access memory (SC-CRAM), *IEEE J. Explor. Solid-State Comput. Devices Circuits* **9**, 29 (2023).
- [12] B. Cai, Y. He, Y. Xin, Z. Yuan, X. Zhang, Z. Zhu, and G. Liang, Unconventional computing based on magnetic tunnel junction, *Appl. Phys. A* **129**, 236 (2023).
- [13] A. Mizrahi, T. Hirtzlin, A. Fukushima, H. Kubota, S. Yuasa, J. Grollier, and D. Querlioz, Neural-like computing with populations of superparamagnetic basis functions, *Nat. Commun.* **9**, 1533 (2018).
- [14] A. Mizrahi, J. Grollier, D. Querlioz, and M. D. Stiles, Overcoming device unreliability with continuous learning in a population coding based computing system, *J. Appl. Phys.* **124**, 152111 (2018).
- [15] D. Rajasekharan, N. Rangarajan, S. Patnaik, O. Sinanoglu, and Y. S. Chauhan, Scanet: Securing the weights with superparamagnetic-mtj crossbar array networks, *IEEE Trans. Neural Netw. Learn. Syst.* **34**, 5693 (2021).
- [16] C. M. Liyanagedera, A. Sengupta, A. Jaiswal, and K. Roy, Stochastic spiking neural networks enabled by magnetic tunnel junctions: From nontelegraphic to telegraphic switching regimes, *Phys. Rev. Appl.* **8**, 064017 (2017).
- [17] J. Cai, B. Fang, L. Zhang, W. Lv, B. Zhang, T. Zhou, G. Finocchio, and Z. Zeng, Voltage-controlled spintronic stochastic neuron based on a magnetic tunnel junction, *Phys. Rev. Appl.* **11**, 034015 (2019).
- [18] K. Y. Camsari, S. Salahuddin, and S. Datta, Implementing p-bits with embedded MTJ, *IEEE Electron Device Lett.* **38**, 1767 (2017).
- [19] W. A. Borders, A. Z. Pervaiz, S. Fukami, K. Y. Camsari, H. Ohno, and S. Datta, Integer factorization using stochastic magnetic tunnel junctions, *Nature* **573**, 390 (2019).
- [20] Y. Lv, R. P. Bloom, and J.-P. Wang, Experimental demonstration of probabilistic spin logic by magnetic tunnel junctions, *IEEE Magn. Lett.* **10**, 1 (2019).
- [21] S. Gibeault, T. N. Adeyeye, L. A. Pocher, D. P. Lathrop, M. W. Daniels, M. D. Stiles, J. J. McClelland, W. A. Borders, J. T. Ryan, P. Talatchian *et al.*, Programmable electrical coupling between stochastic magnetic tunnel junctions, *Phys. Rev. Appl.* **21**, 034064 (2024).
- [22] W. F. Brown, Thermal fluctuations of a single-domain particle, *Phys. Rev.* **130**, 1677 (1963).
- [23] A. Mizrahi, N. Locatelli, R. Matsumoto, A. Fukushima, H. Kubota, S. Yuasa, V. Cros, J.-V. Kim, J. Grollier, and D. Querlioz, Magnetic stochastic oscillators: Noise-induced synchronization to underthreshold excitation and comprehensive compact model, *IEEE Trans. Magn.* **51**, 1 (2015).
- [24] E. Becle, P. Talatchian, G. Prenat, L. Anghel, and I.-L. Prebeanu, in *ESSDERC 2021-IEEE 51st European Solid-State Device Research Conference (ESSDERC)* (IEEE, Grenoble, France, 2021), p. 259.
- [25] P. Talatchian, M. W. Daniels, A. Madhavan, M. R. Pufall, E. Jué, W. H. Rippard, J. J. McClelland, and M. D. Stiles, Mutual control of stochastic switching for two electrically

- coupled superparamagnetic tunnel junctions, *Phys. Rev. B* **104**, 054427 (2021).
- [26] R. Carboni, E. Vernocchi, M. Siddik, J. Harms, A. Lyle, G. Sandhu, and D. Ielmini, A physics-based compact model of stochastic switching in spin-transfer torque magnetic memory, *IEEE Trans. Electron Devices* **66**, 4176 (2019).
- [27] X. Yang, Y. Zhang, Y. Zhang, and P. Wang, A universal compact model for spin-transfer torque-driven magnetization switching in magnetic tunnel junction, *IEEE Trans. Electron Devices* **69**, 6453 (2022).
- [28] L. Schnitzspan, M. Kläui, and G. Jakob, Nanosecond true-random-number generation with superparamagnetic tunnel junctions: Identification of Joule heating and spin-transfer-torque effects, *Phys. Rev. Appl.* **20**, 024002 (2023).
- [29] C. Safranski, J. Kaiser, P. Trouilloud, P. Hashemi, G. Hu, and J. Z. Sun, Demonstration of nanosecond operation in stochastic magnetic tunnel junctions, *Nano Lett.* **21**, 2040 (2021).
- [30] M. Bapna and S. A. Majetich, Current control of time-averaged magnetization in superparamagnetic tunnel junctions, *Appl. Phys. Lett.* **111**, 243107 (2017).
- [31] K. Hayakawa, S. Kanai, T. Funatsu, J. Igarashi, B. Jinnai, W. Borders, H. Ohno, and S. Fukami, Nanosecond random telegraph noise in in-plane magnetic tunnel junctions, *Phys. Rev. Lett.* **126**, 117202 (2021).
- [32] J. Leliaert, J. Mulkers, J. De Clercq, A. Coene, M. Dvornik, and B. Van Waeyenberge, Adaptively time stepping the stochastic Landau-Lifshitz-Gilbert equation at nonzero temperature: Implementation and validation in MuMax3, *AIP Adv.* **7**, 125010 (2017).
- [33] D. E. Endean, C. Weigelt, R. Vitoria, and E. Dan Dahlberg, Tunable random telegraph noise in individual square permalloy dots, *Appl. Phys. Lett.* **104**, 252408 (2014).
- [34] S. Shin, K. Kim, and S.-M. Kang, Compact models for memristors based on charge-flux constitutive relationships, *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.* **29**, 590 (2010).
- [35] F. Ding, B. Peng, X. Li, L. Zhang, R. Wang, Z. Song, and R. Huang, A review of compact modeling for phase change memory, *J. Semicond.* **43**, 023101 (2022).
- [36] J. Jiménez-León, L. A. Sarmiento-Reyes, and P. Rosales-Quintero, A compact modeling methodology for experimental memristive devices, *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.* **41**, 4851 (2021).
- [37] F. A. Araujo, C. Chopin, and S. De Wergifosse, Data-driven thiele equation approach for solving the full nonlinear spin-torque vortex oscillator dynamics, *arXiv:2206.13596*.
- [38] L.-B. Faber, W. Zhao, J.-O. Klein, T. Devolder, and C. Chappert, in *2009 4th International Conference on Design & Technology of Integrated Systems in Nanoscale Era*, edited by H. Eltahawy (IEEE, Cairo, Egypt, 2009), p. 130.
- [39] W. Guo, G. Prenat, V. Javerliac, M. El Baraji, N. De Mestier, C. Baraduc, and B. Dieny, Spice modelling of magnetic tunnel junctions written by spin-transfer torque, *J. Phys. D: Appl. Phys.* **43**, 215001 (2010).
- [40] M. M. Torunbalci, P. Upadhyaya, S. A. Bhawe, and K. Y. Camsari, Modular compact modeling of MTJ devices, *IEEE Trans. Electron Devices* **65**, 4628 (2018).
- [41] S. Chowdhury, A. Grimaldi, N. A. Aadit, S. Niazi, M. Mohseni, S. Kanai, H. Ohno, S. Fukami, L. Theogarajan, G. Finocchio, S. Datta, and K. Y. Camsari, A full-stack view of probabilistic computing with p-bits: Devices, architectures and algorithms, *IEEE Journal on Exploratory Solid-State Computational Devices and Circuits* **9**, 1 (2023).
- [42] S. Bunaiyan, S. Datta, and K. Y. Camsari, Heisenberg machines with programmable spin-circuits, *Phys. Rev. Appl.* **22**, 014014 (2024).
- [43] S. Jung, H. Lee, S. Myung, H. Kim, S. K. Yoon, S.-W. Kwon, Y. Ju, M. Kim, W. Yi, S. Han, B. Kwon, B. Seo, K. Lee, G.-H. Koh, K. Lee, Y. Song, C. Choi, D. Ham, and S. J. Kim, A crossbar array of magnetoresistive memory devices for in-memory computing, *Nature* **601**, 211 (2022).
- [44] K. Y. Camsari, S. Salahuddin, and S. Datta, Implementing p-bits with embedded MTJ, *IEEE Electron Device Lett.* **38**, 1767 (2017).
- [45] J. Xiao, A. Zangwill, and M. D. Stiles, Macrospin models of spin transfer dynamics, *Phys. Rev. B* **72**, 014446 (2005).
- [46] S. Kanai, K. Hayakawa, M. Elyasi, K. Kobayashi, J. Igarashi, B. Jinnai, W. A. Borders, G. E. W. Bauer, H. Ohno, and S. Fukami, Stochastic switching time constant and instability in nanomagnets, (to be published).
- [47] D. Berkov and N. Gorn, Thermally activated processes in magnetic systems consisting of rigid dipoles: Equivalence of the Ito and Stratonovich stochastic calculus, *J. Phys.: Condens. Matter* **14**, L281 (2002).
- [48] L. Desplat and J.-V. Kim, Entropy-reduced retention times in magnetic memory elements: A case of the meyer-neldel compensation rule, *Phys. Rev. Lett.* **125**, 107201 (2020).
- [49] I. Cimrák, A survey on the numerics and computations for the Landau-Lifshitz equation of micromagnetism, *Arch. Comput. Methods Eng.* **15**, 1 (2007).
- [50] C. Aron, D. G. Barci, L. F. Cugliandolo, Z. G. Arenas, and G. S. Lozano, Magnetization dynamics: Path-integral formalism for the stochastic Landau-Lifshitz-Gilbert equation, *J. Stat. Mech: Theory Exp.* **2014**, P09008 (2014).
- [51] F. Romá, L. F. Cugliandolo, and G. S. Lozano, Numerical integration of the stochastic Landau-Lifshitz-Gilbert equation in generic time-discretization schemes, *Phys. Rev. E* **90**, 023203 (2014).
- [52] J. L. García-Palacios and F. J. Lázaro, Langevin-dynamics study of the dynamical properties of small magnetic particles, *Phys. Rev. B* **58**, 14937 (1998).
- [53] S. Ament, N. Rangarajan, A. Parthasarathy, and S. Rakheja, Solving the stochastic Landau-Lifshitz-Gilbert-Slonczewski equation for monodomain nanomagnets: A survey and analysis of numerical techniques, *arXiv:1607.04596*.
- [54] L. Boninsegna, F. Nüske, and C. Clementi, Sparse learning of stochastic dynamical equations, *J. Chem. Phys.* **148**, 241723 (2018).
- [55] J. L. Callahan, J.-C. Loiseau, G. Rigas, and S. L. Brunton, Nonlinear stochastic modelling with Langevin regression, *Proc. R. Soc. A* **477**, 20210092 (2021).
- [56] L. R. Gorrão, J. Heysel, K. Lehnertz, and M. R. R. Tabar, Analysis and data-driven reconstruction of bivariate jump-diffusion processes, *Phys. Rev. E* **100**, 062127 (2019).
- [57] We note that in reality the driving noise terms likely possess additional structure beyond mere white noise. However, as we show in Sec. V, the white noise assumption does very well in reproducing the statistics observed in experiment.

- Physically, the choice of white noise is the unique choice that ensures the dissipative mechanisms of the system are local in time [81].
- [58] These conditional moments are reminiscent of Kolmogorov's structure functions in turbulence. However, structure functions are conditional moments that use a spatial separation and spatial ensemble average to obtain a scalar value for fluid velocity, whereas the Kramers-Moyal coefficients here are evaluated from a voltage-time trace (which by the ergodic hypothesis is equivalent to an ensemble average) that integrates solely over the probability distribution.
 - [59] H. Risken and H. Risken, *The Fokker-Planck equation* (Springer, Berlin, 1996).
 - [60] C. Honisch and R. Friedrich, Estimation of Kramers-Moyal coefficients at low sampling rates, *Phys. Rev. E* **83**, 066701 (2011).
 - [61] R. Friedrich, J. Peinke, and C. Renner, How to quantify deterministic and random influences on the statistics of the foreign exchange market, *Phys. Rev. Lett.* **84**, 5224 (2000).
 - [62] R. F. Pawula, Approximation of the linear Boltzmann equation by the Fokker-Planck equation, *Phys. Rev.* **162**, 186 (1967).
 - [63] J. Gottschall and J. Peinke, On the definition and handling of different drift and diffusion estimates, *New J. Phys.* **10**, 083034 (2008).
 - [64] L. Rydin Gorjão, D. Witthaut, K. Lehnertz, and P. G. Lind, Arbitrary-order finite-time corrections for the Kramers-Moyal operator, *Entropy* **23**, 517 (2021).
 - [65] F. Sicard, V. Koskin, A. Annibale, and E. Rosta, Position dependent diffusion from biased simulations and Markov state model analysis, *J. Chem. Theory. Comput.* **17**, 2022 (2021).
 - [66] N. Vercauteren and C. Hartmann, Ph.D. thesis, École Polytechnique Fédérale de Lausanne. 2005.
 - [67] Legendre polynomials were also considered because of their similar properties, but we chose Chebyshev polynomials because they obtained more accurate results with fewer total coefficients.
 - [68] N. G. Van Kampen, Itô versus Stratonovich, *J. Stat. Phys.* **24**, 175 (1981).
 - [69] M. W. Daniels, W. A. Borders, N. Prasad, A. Madhavan, S. Gibeault, T. Adeyeye, L. Pocher, L. Wan, M. Tran, J. A. Katine *et al.*, in *Spintronics XVI*, edited by J.-E. Wegrowe, J. S. Friedman, and M. Razeghi (SPIE, San Diego, CA, USA, 2023), Vol. 12656, p. 84.
 - [70] S. Kostinski and A. Amir, An elementary derivation of first and last return times of 1D random walks, *Am. J. Phys.* **84**, 57 (2016).
 - [71] R. Fitzhugh, Statistical properties of the asymmetric random telegraph signal, with applications to single-channel analysis, *Math. Biosci.* **64**, 75 (1983).
 - [72] D. S. Lemons and P. Langevin, *An Introduction to Stochastic Processes in Physics* (JHU Press, Baltimore, 2002).
 - [73] We remark that, unlike what we would expect in a real magnetic tunnel junction, there is no physically imposed ultraviolet cutoff for macrospin dynamics. This is not particularly problematic, though, since any good circuit simulation involving such a device would impose its own high-frequency cutoffs through the finite bandwidth of realistic circuit components.
 - [74] K. Gilmore, M. D. Stiles, J. Seib, D. Steiauf, and M. Fähnle, Anisotropic damping of the magnetization dynamics in Ni, Co, and Fe, *Phys. Rev. B* **81**, 174414 (2010).
 - [75] L. Soumah, L. Desplat, N.-T. Phan, A. S. E. Valli, A. Madhavan, F. Disdier, S. Auffret, R. Sousa, U. Ebels, and P. Talatchian, Nanosecond stochastic operation in perpendicular superparamagnetic tunnel junctions, [arXiv:2402.03452](https://arxiv.org/abs/2402.03452).
 - [76] A. Chavent, V. Iurchuk, L. Tillie, Y. Bel, N. Lamard, L. Vila, U. Ebels, R. C. Sousa, B. Dieny, G. Di Pendina *et al.*, A multifunctional standardized magnetic tunnel junction stack embedding sensor, memory and oscillator functionality, *J. Magn. Magn. Mater.* **505**, 166647 (2020).
 - [77] R. Ma, A. Sidi El Valli, M. Kreibitz, G. Di Pendina, F. Protze, U. Ebels, G. Prenat, A. Chavent, V. Iurchuk, R. Sousa *et al.*, Microwave functionality of spintronic devices implemented in a hybrid complementary metal oxide semiconductor and magnetic tunnel junction technology, *Electron. Lett.* **57**, 264 (2021).
 - [78] B. Dieny and M. Chshiev, Perpendicular magnetic anisotropy at transition metal/oxide interfaces and applications, *Rev. Mod. Phys.* **89**, 025008 (2017).
 - [79] N. Strelkov, A. Chavent, A. Timopheev, R. Sousa, I. Prejbeanu, L. Buda-Prejbeanu, and B. Dieny, Impact of Joule heating on the stability phase diagrams of perpendicular magnetic tunnel junctions, *Phys. Rev. B* **98**, 214410 (2018).
 - [80] Numerical experimental power spectral densities were estimated using Welch's method [82], with a Hann window.
 - [81] R. Kubo, The fluctuation-dissipation theorem, *Rep. Prog. Phys.* **29**, 255 (1966).
 - [82] P. Welch, The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms, *IEEE Trans. Audio Electroacoust.* **15**, 70 (1967).