

Data blinding for the nEDM experiment at PSI

N. J. Ayres^{1,2}, G. Ban³, G. Bison⁴, K. Bodek⁵, V. Bondar^{1,4,6}, E. Chanel⁷, P.-J. Chiu⁴, C. Crawford⁸, M. Daum⁴, S. Emmenegger¹, L. Ferraris-Bouchez⁹, P. Flaux³, Z. Grujić^{10,†}, P.G Harris², N. Hild⁴, J. Hommet³, M. Kasprzak^{4,6,10}, Y. Kermaidic⁹, K. Kirch^{1,4}, S. Komposch^{1,4}, A. Kozela¹¹, J. Krempel^{*,1}, B. Lauss⁴, T. Lefort³, Y. Lemiere³, A. Leredde⁹, P. Mohanmurthy^{1,4}, A. Mtchedlishvili⁴, O. Naviliat-Cuncic³, D. Pais⁴, F. M. Piegsa⁷, G. Pignol⁹, M. Rawlik¹, D. Rebreyend⁹, I. Rienäcker⁴, D. Ries¹², S. Rocca^{13,14}, D. Rozpedzik⁵, P. Schmidt-Wellenburg⁴, A. Schnabel¹⁵, R. Viro⁹, A. Weis¹⁰, E. Wursten^{6,‡}, J. Zejma⁵, and G. Zsigmond⁴

¹Institute for Particle Physics and Astrophysics, ETH Zürich, Zürich, Switzerland

²Department of Physics and Astronomy, University of Sussex, Falmer, Brighton, UK

³LPC Caen, ENSICAEN, Université de Caen, CNRS/IN2P3, Caen, France

⁴Paul Scherrer Institute, Villigen, Switzerland

⁵M. Smoluchowski Institute of Physics, Jagiellonian University in Krakow, Poland

⁶Instituut voor Kern- en Stralingsfysica, Katholieke Universiteit Leuven, Leuven, Belgium

⁷University of Bern, Albert Einstein Center for Fundamental Physics, Laboratory for High Energy Physics, Bern, Switzerland

⁸Department of Physics and Astronomy, University of Kentucky, Lexington, USA

⁹Univ. Grenoble Alpes, CNRS, Grenoble INP, LPSC-IN2P3, Grenoble, France

¹⁰Physics Department, University of Fribourg, Fribourg, Switzerland

¹¹H. Niewodniczanski Institute of Nuclear Physics, Polish Academy of Sciences, Krakow, Poland

¹²Department of Chemistry - TRIGA site, Johannes Gutenberg University Mainz, Mainz, Germany

¹³CSNSM, Université Paris Sud, CNRS/IN2P3, Université Paris Saclay, Orsay-Campus, France

¹⁴Institut Laue-Langevin, Grenoble, France

¹⁵Physikalisch Technische Bundesanstalt, Berlin, Germany

[†]Present address: Institute of Physics Belgrade, Belgrade, Serbia

[‡]Present address: CERN, Geneva, Switzerland

*Corresponding author: Jochen.Krempel@phys.ethz.ch

Abstract

Psychological bias towards, or away from, a prior measurement or a theory prediction is an intrinsic threat to any data analysis. While various methods can be used to avoid the bias, *e.g.* actively not looking at the result, only data blinding is a traceable and thus trustworthy method to circumvent the bias and to convince a public audience that there is not even an accidental psychological bias.

Data blinding is nowadays a standard practice in particle physics, but it is particularly difficult for experiments searching for the neutron electric dipole moment (nEDM), as several cross measurements, in particular of the magnetic field, create a self-consistent network into which it is hard to inject a fake signal.

We present an algorithm that modifies the data without influencing the experiment. Results of an automated analysis of the data are used to change the recorded spin state of a few neutrons of each measurement cycle.

The flexible algorithm is applied twice to the data, to provide different data to various analysis teams. This gives us the option to sequentially apply various blinding offsets for separate analysis steps with independent teams. The subtle modification of the data allows us to modify the algorithm and to produce a re-blinded data set without revealing the blinding secret. The method was designed for the 2015/2016 measurement campaign of the nEDM experiment at the Paul Scherrer Institute. However, it can be re-used with minor modification for the follow-up experiment n2EDM, and may be suitable for comparable efforts.

1 Introduction

The electric dipole moment (EDM) of the neutron is a fundamental observable in particle physics that may directly relate to the observed dominance of matter over antimatter in the Universe. It has been sought after for almost seven decades, with ever-improving sensitivity, but to date all experimental results have been compatible with zero [1, 2, 3, 4, 5]. However, many beyond the Standard Model theories naturally predict values that are close to current experimental sensitivities [6, 7, 8]. Thus, depending upon their outlook, scientists analysing the data from EDM experiments may be biased unintentionally towards a result that favours their own expectations of either seeing, or not seeing, a statistically significant signal. Data blinding removes this psychological bias and, if applied properly, does not introduce any other bias. In experimental particle physics, blinding has been used quite commonly for many years [9], but to date it has not been applied to any neutron EDM measurement.

In general, at least two different types of blinding can be distinguished:

1. Data corresponding to a region of interest is withheld from the analysis team, or, correspondingly, “fake” events can be added to obscure the signal. This is often the case in discovery experiments. See, *e.g.* (not the latest but representative) searches for rare decays [10], dark matter [11] or gravitational waves [12].
2. For precision experiments the observable of interest can be scaled by an unknown factor [13], or in some cases, an unknown offset can be added to the observable [14].




Figure 1: Measured neutron counts plotted versus spin-flip frequency. Both are corrected for fluctuations as indicated in the axis labels. A sine with offset is fitted to the data points. $\bar{\nu}_{\text{Hg}}$ is the average reading of the mercury co-magnetometer. Both averages used in this plot are calculated from the measurements shown in this graph.

The latter is obviously applicable to EDM experiments, and it is the approach that we have adopted for the nEDM experiment at the Paul Scherrer Institute (PSI) [15]. In deciding to modify the observable, one can choose to do so either by changing an aspect of the experiment itself, or by modifying the data *post hoc*. The latter has the advantage that it does not change or corrupt the experiment, and a hidden set of original data can be stored securely. Thus, if the blinding were to affect the data quality in any way, *e.g.* by reducing the sensitivity, or by introducing a new bias, the original data can still be used knowing that the final result is unaffected by any systematic effects that may have been introduced through blinding.

2 Experimental overview

In nEDM experiments the target observable is the dependence of the neutrons' Larmor precession frequency upon an applied static electric field [15]. A set of ultracold neutrons (UCN) is polarized and then stored in a volume within a stable and highly uniform magnetic field. In most experiments to date the frequency measurement is based on Ramsey's technique of separated oscillatory fields. Two spin-flip pulses which induce each a $\pi/2$ flip through transverse oscillating magnetic fields are applied. The two pulses must have a well-known phase relation and were kept in phase at all times during our data taking. Between these pulses the neutron spins can precess freely. If the spin-flip frequency is perfectly in resonance with the Larmor frequency of the neutrons, they will have undergone a π spin-flip by the end of the procedure. If not, the accumulated phase difference, a highly sensitive measure of the difference between the Larmor and reference frequencies, will result in a partial spin-flip. The neutrons are then counted in a spin-sensitive detector. By repeating such measurements while scanning the reference frequency, and plotting the final neutron spin state versus that frequency, a Ramsey fringe pattern emerges as plotted in fig. 1. For a non-zero EDM value the pattern will shift when the electric field direction is reversed, which is done periodically.

Typically an entrance spin-flipper is present, in our case called SF1, which will

invert the initial neutron spin orientation if employed. Alternating its state is used to investigate the influence of systematic effects. Regular changes of the magnetic field orientation and a variation of the magnetic field gradient serve the same purpose.

During the 2015/16 data taking campaign at PSI [15] we detected up to 20 000 UCN after having stored them in a volume of 22 litres within a highly uniform magnetic field of approximately 1 μ T. Each of such single measurements, called “cycle”, was composed of two $\pi/2$ spin-flip pulses of a frequency of $\nu_{F,i} \approx 30$ Hz applied for a duration of $t = 2$ s each and a free precession period of $T = 180$ s before the UCN were counted. The detector simultaneously measured spin-up/down neutrons in two branches. Each consisted of a controllable spin-flipper, a magnetized spin-analysing foil, and ${}^6\text{Li}$ based neutron detectors that were read out via photomultiplier tubes [16, 17]. Timestamp, integrated charge and detector channel of every event were recorded in the data files. A set of consecutive cycles carried out with a stable magnetic field configuration, but with variation of the applied electric field, was called a “run”. During a run, lasting for up to several days, and typically consisting of several hundred cycles, the configuration of the spin-flippers in the detector branches was reversed every four cycles, and the entrance spin-flipper status (the spin orientation of the neutrons let in) was changed every 112 cycles. Within those runs, eight cycles with no electric field were followed by 48 cycles with high voltage of a given polarity. Thus, both electric field polarities were applied between each change of the entrance spin-flipper state.

2.1 Formal description

The $\pi/2$ spin-flips of frequency $\nu_{F,i}$, that are applied in a particular cycle i , cause a change of polarization described by the phase

$$\phi_i = \frac{(\nu_{F,i} - \nu_L)}{\Delta\nu} \pi, \quad (1)$$

where ν_L is the Larmor frequency and the fringe width $\Delta\nu$ is

$$\Delta\nu = \frac{1}{2(T + 4t/\pi)}, \quad (2)$$

with T and t being respectively the free-precession time and the duration of each spin-flip pulse.

The true numbers of neutrons $N'_{\uparrow,i}$ and $N'_{\downarrow,i}$ for the particular spin states that are in the precession chamber right before guiding them to the detector are

$$N'_{\uparrow,i} = \frac{N'_i}{2} (1 - \alpha' \cos \phi_i) \quad (3)$$

$$N'_{\downarrow,i} = \frac{N'_i}{2} (1 + \alpha' \cos \phi_i), \quad (4)$$

where N'_i is the total number of neutrons in the chamber after the precession and α' describes the true visibility after the precession, but has a negative sign in case SF1 is enabled.

The neutrons will then fall through a polarisation analyser with spin selectivities p_A and p_B to the detectors that operate with efficiencies ϵ_A and ϵ_B . Thus, the

numbers of measured neutrons are

$$N_{\uparrow,i} = (N'_{\uparrow,i} p_A + N'_{\downarrow,i} (1 - p_A)) \epsilon_A \quad (5)$$

$$N_{\downarrow,i} = (N'_{\downarrow,i} p_B + N'_{\uparrow,i} (1 - p_B)) \epsilon_B . \quad (6)$$

In this model the efficiency of the spin-flippers is neglected since it is very close to 1.

“Spin-up” (\uparrow) refers to neutrons with the spin polarisation anti-parallel to the magnetic field B_0 , and therefore with the magnetic moment parallel to the field. They are also known as “high field seekers”. When SF1 is off, this is the state in which they enter the bottle and is thus their state before the Ramsey sequence is applied.

A fit of eqs. (1)–(6) to the data yields the Larmor frequency. The steepest part of the slope, *i.e.* where $\phi_i \approx \pm 90^\circ$, is most sensitive to variations in frequency. Thus the spin-flip frequencies were configured to operate sequentially at four working points (see points in fig. 2) close to this condition with a detuning of about 5% of the fringe width in order to have some sensitivity also to further experimental parameters such as the visibility and the asymmetry of the detector efficiency.

In the presence of an nEDM d and an applied electric field \vec{E} collinear to the magnetic field \vec{B} , the resonant frequency ν_L shifts by

$$\delta\nu = 2d\vec{E} \cdot \frac{\vec{B}}{|\vec{B}|} / h , \quad (7)$$

with h being the Planck constant. Note that the $\vec{B}/|\vec{B}|$ term is required only to obtain the proper sign.

Unfortunately, any change of the amplitude of the magnetic field causes a corresponding change of the Larmor frequency. We used a mercury co-magnetometer to correct for magnetic-field fluctuations by taking the ratio of the measured frequencies $\mathcal{R} = \nu_n/\nu_{\text{Hg}}$ [18]. Correspondingly, eq. (1) is altered as shown in eq. (12). However, although the (thermal) mercury atoms populate the storage cell rather uniformly, the UCN have such low kinetic energies that they sag a little under gravity. Any vertical gradient of the magnetic field therefore results in a different average value of the magnetic field for the two species. This in turn leads to a small shift in the mercury-corrected neutron Larmor frequency. For a given vertical gradient, this shift is in opposite directions for the two different orientations (up vs. down) of the main magnetic field. Furthermore, there is a systematic effect leading to a significant false-EDM arising from a conjunction of the vertical magnetic-field gradient and the relativistic motional magnetic field seen by the mercury atoms (in particular) as they move through the electric field [19]. We therefore needed to interpolate our measured nEDM results to zero vertical magnetic-field gradient. As we did not have an absolute gradiometer, we intentionally applied small magnetic-field gradients using trim coils in order to determine the situation at zero gradient from the intersection of the two curves arising from the two magnetic-field directions [15, 19, 20]. Hence, it is important to state that a blinding offset, in contrast to some kind of blinding factor, does not interfere with the interpolation of the curves.

3 Data blinding

3.1 Blinding concept

Any blinding method for an nEDM experiment must shift the measured Larmor frequency in correlation to the electric field, ideally while leaving all other observables unaltered. The following blinding procedures were briefly considered by our collaboration:

- Apply a modified spin-flip frequency with respect to the recorded value during the experiment. However, this would modify the experiment in an insidious manner as the change in actual physical conditions applied would be correlated to the electric field changes. This could therefore potentially introduce systematic effects, and, additionally, it would be irreversible. So, one would have no possibility to investigate it *a posteriori*.
- Register a shifted spin-flip frequency with respect to the one actually applied. This would require also subtle adjustments of all other magnetometer readings in order not to immediately reveal the shift by comparison. In our case this would have meant to adjust consistently a total of 16 magnetometer readings (1 mercury and 15 caesium) [15, 21] — A daunting task.

Both of these rejected techniques would mimic an EDM according to the usual analysis strategies. There are other observables that can be modified, but they would not have exactly the same appearance as an EDM signal, and thus it would ultimately be fairly trivial to identify them as fake signals. Note, that manipulating the value of the electric field cannot be used to introduce a blinding offset.

The remaining variable that can be modified is the neutron counts. The primary difficulty in that case is that, since the size of the required shift depends upon the neutron counts themselves, a partial but automatic analysis of the data must be done. And, in order to hand out as early as possible blinded data to the analysis teams, this must be done during data taking in a way that is fully defined before starting the actual data-taking campaign. This was the approach that we have adopted for the nEDM experiment, and its implementation will be described in detail in the following sections.

3.2 Algorithm

The blinding algorithm operates stepwise. First the necessary parameters are extracted from a full run (sect. 3.2.2). Then the position of each cycle on the Ramsey curve, the so-called working point, is determined (sect. 3.2.3), before the number of neutrons in each cycle can be modified (sect. 3.2.4).

3.2.1 Calculation of the number of neutrons to be transferred

In order to generate an \vec{E} -field dependent frequency shift a small number of spin-up neutrons have to be reclassified as spin-down, or *vice versa*, see fig. 2.

We follow eqs. (1), (5), (6) and (7) as well as the first-order Taylor expansion

$\delta N = \left(\frac{d}{d\phi} N\right) \left(\frac{d}{d\nu_L} \phi\right) \delta\nu$ to find the number of neutrons that need to change state:

$$\delta N_{\uparrow,i} = \epsilon_A \frac{N'_i}{2} \alpha' (2p_A - 1) (\sin \phi_i) \left(\frac{d}{d\nu_L} \phi\right) \delta\nu \quad (8)$$

$$= -\epsilon_A \frac{N'_i}{2} \alpha' (2p_A - 1) (\sin \phi_i) \frac{\pi}{\Delta\nu} \frac{2 d \vec{E} \cdot \vec{B} / |\vec{B}|}{h} \quad (9)$$

$$\delta N_{\downarrow,i} = +\epsilon_B \frac{N'_i}{2} \alpha' (2p_B - 1) (\sin \phi_i) \frac{\pi}{\Delta\nu} \frac{2 d \vec{E} \cdot \vec{B} / |\vec{B}|}{h}. \quad (10)$$

Note that \vec{E} and \vec{B} have to be parallel or antiparallel, and that the sign of α' can be negative depending on the state of SF1.

At this point, we make the convenient approximation that $\epsilon_A = \epsilon_B$, which is reasonable in light of the performance of our detectors [16, 22]. This allows us to introduce the total number of measured counts per cycle $N_i = \frac{\epsilon_A + \epsilon_B}{2} N'_i$ and its average over a run $\bar{N} = \langle N_i \rangle$. We assume further that the performance of our spin analysers is equal for the two spin states, that is $p_A = p_B$, which is, again, reasonable and supported by the data. Thus, we can introduce $\alpha = \alpha' (2p_A - 1)$ which describes the measured visibility, but again, is negative in case SF1 is enabled. Therefore,

$$\delta N_{\uparrow,\downarrow,i} = \mp N_i \frac{\pi \alpha}{\Delta\nu} \frac{d \vec{E} \cdot \vec{B} / |\vec{B}|}{h} \sin \phi_i. \quad (11)$$

We will discuss below the implications of removing the assumptions of $\epsilon_A = \epsilon_B$ and $p_A = p_B$.

Typical values for the nEDM experiment are $N_i = 15\,000$, $|\sin \phi| = 0.99$, $\alpha = 0.75$, $T = 180$ s, $t = 2$ s, and $E = 11$ kV/cm. Thus an EDM offset of 1.0×10^{-25} e cm would require a shift of about 3.39 neutrons in each cycle. In terms of signs, bearing in mind that the neutron has a negative magnetic moment, if \vec{B} and \vec{E} are parallel, a positive nEDM would *reduce* the precession frequency. This would shift the Ramsey curves towards smaller frequencies, which means that neutrons measured at a working point above the resonant frequency shift from the spin-down detector arm to the spin-up. Neutrons that are measured below the working point shift from down to up, correspondingly. Figure 2 illustrates this reclassification and the resulting shift.

Obviously, it is impossible to shift a non-integer number of neutrons in a single cycle. One could simply round the number, but this would effectively cause a granularity of $\sim 3 \times 10^{-26}$ e cm in the available blinding offsets. However, we can add to δN a small random number with a normal distribution, and round the sum to the nearest integer number. The choice of the width of this normal distribution was driven by two facts. On the one hand, a small width does not smooth the granularity sufficiently. On the other hand, a large width adds noise to the neutron counts and thus to the blinded nEDM value. We found a suitable compromise using a standard deviation of 2 counts. In this case the granularity is sufficiently suppressed so that the result differs from a flat distribution by less than 10^{-7} . An improved method will be suggested in sect. 6.

As mentioned above, this algorithm assumes the same \bar{N} and α for each of the two spin states. If this were not to be the case, a direct transfer of neutrons from one spin state to the other would not be appropriate. Instead, one would have to




Figure 2: Simulated neutron counts plotted versus applied spin-flip frequency ν_F . The transfer of a small number of neutrons (green arrows) from their initially recorded state, *e.g.* counts N_\uparrow (blue circles), corresponding to the original Larmor frequency ν_L , to the other spin state creates the blinded data points (orange squares). If this is done systematically and proportionally to the electric field, one can extract from the resulting dashed orange line a different Larmor frequency $\nu_d = \nu_L + \delta\nu$. The frequency shift by $\delta\nu$ (violet arrow) represents a false EDM signal d given by eq. (7). For the detector arm counting the opposite spin state, *e.g.* N_\downarrow , the corresponding shift leads from the solid red to the dotted magenta curve. This yields the same false EDM signal. In case SF1 is on, all points and lines must be mirrored at a horizontal line at $N = 7500$. For clarity the strongly exaggerated values $|\sin\phi| = 0.951$ and $d = 3 \times 10^{-23}$ e cm have been used here.

analyse and treat the two states separately, and neutrons would have to be added to or deleted from the spin-up and spin-down arrays as required. While this is trivial if the neutron data merely consists of a simple sum of counts per cycle, it is a substantial effort for a more detailed data format such as ours, which lists charge and time per event.

3.2.2 Determination of α and detector asymmetry

Before the data can be blinded one has to determine α and ν_L . While α is sufficiently constant throughout an entire run, ν_L might change from cycle to cycle and must be corrected with the field values recorded by the mercury co-magnetometer. We therefore refer to it as $\nu_{L,i}$ and write

$$\nu_{L,i} = \left| \frac{\gamma_n}{\gamma_{\text{Hg}}} \right| \nu_{\text{Hg},i} - \frac{\Phi}{\pi} \Delta\nu, \quad (12)$$

where γ_n and γ_{Hg} are the gyromagnetic ratios of the neutron and mercury respectively, $\nu_{\text{Hg},i}$ the frequency obtained from the mercury co-magnetometer, and the phase Φ accommodates any difference in the average magnetic field sampled by the two species. We had measured the ratio of gyromagnetic ratios in a previous experiment [23]. For the blinding algorithm we used a fixed value of $\left| \frac{\gamma_n}{\gamma_{\text{Hg}}} \right| = 3.8424574$. As magnetic field gradients were not changed during a run, Φ kept the same value throughout all cycles of the run.

We rewrite eqs. (5) and (6) as

$$\frac{N_{\uparrow,i} - N_{\downarrow,i}}{N_{\uparrow,i} + N_{\downarrow,i}} = \frac{\epsilon_A - \epsilon_B - \alpha' (\epsilon_A(2p_A - 1) + \epsilon_B(2p_B - 1)) \cos \phi_i}{\epsilon_A + \epsilon_B + \alpha' (-\epsilon_A(2p_A - 1) + \epsilon_B(2p_B - 1)) \cos \phi_i} \quad (13)$$

$$\approx \frac{\epsilon_A - \epsilon_B}{\epsilon_A + \epsilon_B} - \alpha \cos \phi_i, \quad (14)$$

where we made the approximations $\epsilon_A \approx \epsilon_B$ and $p_A \approx p_B$. Furthermore, we used that $\epsilon_A - \epsilon_B \ll \epsilon_A + \epsilon_B$ which is also justified by the data.

We define $A_m = \frac{\epsilon_A - \epsilon_B}{\epsilon_A + \epsilon_B}$ which is the detector asymmetry originating from the slightly different efficiencies of the two detector arms counting the two spin states. We use eqs. (12) and (1) to rewrite eq. (14) as

$$\frac{N_{\uparrow,i} - N_{\downarrow,i}}{N_{\uparrow,i} + N_{\downarrow,i}} = f \left(\nu_{F,i} - \left| \frac{\gamma_n}{\gamma_{\text{Hg}}} \right| \nu_{\text{Hg},i} \right), \quad (15)$$

where we have defined a function as:

$$f(x) = A_m - \alpha \cos \left(\frac{\pi}{\Delta\nu} x + \Phi \right). \quad (16)$$

The independent variable x is beneficial for the fit algorithm, since it can be calculated from the observables for each cycle. It represents the frequency difference between applied spin-flip and neutron resonance. The parameter estimation of A_m , α and Φ is done by fitting the data of a full run to eq. 15. Every four cycles we inverted the spin-flipper states on both detector arms by activating and deactivating spin-flippers that are mounted inside the arms [16, 22, 17]. This results in a “Normal” and an “Inverted” configuration, with asymmetries A_N and A_I respectively.

Both values are almost constant throughout a run. We retained them as fit parameters to accommodate for long term changes. Consequently, the data contains two collated subsets and the fit must be conducted as a simultaneous fit, within which α and Φ are common parameters while A_N and A_I apply to the respective partial data sets only.

3.2.3 Determination of the Ramsey phase ϕ_i

After having carried out the fit on the full run we can use eq. (11) to calculate the number of neutrons to be transferred for each cycle. However, we still need to determine ϕ_i . This may be done either via eq. (14):

$$\cos \phi_i = \frac{1}{\alpha} \left(\frac{N_{\uparrow,i} - N_{\downarrow,i}}{N_{\uparrow,i} + N_{\downarrow,i}} - A_m \right), \quad (17)$$

or via eq. (1):

$$\phi_i = \frac{\nu_{F,i} - \left| \frac{\gamma_B}{\gamma_{Hg}} \right| \nu_{Hg,i}}{\Delta\nu} \pi + \Phi. \quad (18)$$

We have implemented the first variant, as it is more robust in instances where in a single cycle the co-magnetometer provides a bad reading. Note that also this variant uses eq. (18) to determine the sign of ϕ_i .

3.2.4 Transferring neutrons

The data files are an event-driven list where each entry consists of a time stamp, the integrated charge recorded at the time, and the identification number of the photomultiplier tube (PMT) that observed the event [17]; if the charge exceeds a certain threshold then the event is classified as a neutron detection. Each of the two detector arms, one per spin state, consists of a set of nine PMTs, which are sequentially numbered from 0 to 17 with 0–8 in the first arm and 9–17 in the second. In order to reclassify the spin of a neutron it is therefore sufficient to take the PMT number of that event, add 9 and do a modulo 18 operation. A neutron that is to be transferred is chosen by randomly selecting an event from the list counted in the correct detector arm, and then checking whether it is suitable to be moved simply and cleanly across: the requirement is that there must be a minimum separation in time between the event in question and the previous and subsequent events. We apply this condition to both the source and the recipient channel. The reason is to avoid the transfer of events for which the charge is split between neighbouring PMTs, or where the baseline correction algorithm has or would have to modify the charge [24]. If the event is not suitable, another randomly chosen event is tested until an appropriate one is found.

3.3 Choice of the blinding offset

Obviously, the value of the blinding offset must be kept secret from the analysis teams. In order not to provide any indirect psychological bias as to what it might be it should be randomly chosen from a distribution that allows a wide range of values. It is convenient for its modulus to be larger than the known upper limit of the nEDM, since this allows a “sanity check” of having a sign that can be confirmed for consistency prior to publication of results (see sect. 5).




Figure 3: Probability density function for the choice of the blinding offset created with 10^6 samples. The dashed vertical lines indicate the $\pm 1\sigma$ sensitivity of the data accumulated in 2015 and 2016 assuming a mean value of 0. For psychological reasons we keep the probability that an offset in this range is selected very small but non-zero (integrated probability $\approx 2 \times 10^{-4}$).

At the same time, it should be sufficiently small in order to guarantee that the working points are not shifted away from the steep slope of the Ramsey pattern such that we maintain the sensitivity and the Taylor expansion used in eqs. (8)–(10) holds. Any error in the calculation of the number of neutrons that are shifted by the blinding process will add noise to the EDM signal and therefore make it more difficult to look for effects and correlations that might indicate possible systematic effects such as the motional-field effect described above.

For the nEDM experiment we chose a combination of four Heaviside step functions that together define a range of $\pm 15 \times 10^{-26}$ ecm and exclude a modulus smaller than 5×10^{-26} ecm. We then blurred this function with a Gaussian of width $\pm 1.5 \times 10^{-26}$ ecm. We also explicitly excluded the extremely unlikely possibility that the tail of the Gaussian would extend beyond $\pm 1 \times 10^{-24}$ ecm, in order to ensure that we stay within the linear region of the original Ramsey fringe. One could argue that this latter step represents a small psychological bias, but — notwithstanding the previously existing world limit — a one-day measurement without blinding leads to the certain conclusion that the true nEDM value must be smaller. Finally, we also excluded a modulus of $< 1 \times 10^{-28}$ ecm for technical reasons, since when communicating between different programs we use a value of exactly zero for cycles that should not be blinded at all, *e.g.* those with no applied electric field. Figure 3 shows the probability distribution of the blinding offset.

3.4 Secondary blinding and reblinding

The nEDM collaboration decided before data taking to have the analysis carried out by two independent teams, referred to as eastern and western, loosely reflecting the geographic distributions of the involved institutions. In order to allow them to communicate without introducing a bias in case of any discrepancy over the mean value of the nEDM, it was decided that in addition to the first-stage “primary” blinding we would use exactly the same algorithm to apply a separate “secondary” blinding that was distinct for each group, *i.e.* with a different offset. Figure 4




Figure 4: Illustration of primary and secondary blinding. Each analysis group has access only to their respective blinded data set, “eastern” or “western”.

illustrates this process.

During the early days of data taking some concern was expressed that the automatic fitting algorithm might not work properly in all cases, or that some important properties of the data might be hidden as a result of the blinding, or that some other similarly unexpected events might make it necessary to significantly change the blinding algorithm. In order to provide a consistent data set in any of those cases, it would be necessary to run a modified blinding program again from scratch on the raw data. However, since the first set of blinded data would by then already be available to the analysis teams, it would be trivial for them to compare two versions of the same data file, and by leaving out all mismatching events they would have an unblinded data set with a statistical significance close to the original data set. In order to avoid such a scenario we made sure that our pseudo-random number generator delivers reproducible numbers, and that the neutrons that are transferred are selected reproducibly. Thus, if *e.g.* one version of the blinding algorithm shifts seven neutrons and the other eight within a given cycle, the two resulting files would only differ by one neutron for that cycle. Therefore, a reblinding using the same or similar offset and a slightly modified algorithm can be carried out without danger of inadvertent unblinding. It shall be mentioned that reblinding with an offset of opposite sign would immediately reveal both offsets.

In addition to transferring the neutrons between spin states, the blinding algorithm also marks each blinded data file with the date of blinding and the version number of the blinding code in order to ensure that those otherwise very similar files remain clearly distinguishable.

After all, the reblinding feature was not used, since in our data analysis no large discrepancies occurred.

3.5 Pseudo-random number generator (PRNG)

In principle, the random numbers used should meet the same strict requirements as those for strong cryptography regarding the prediction of numbers, the correlation between them and the uniformity of their distribution. However, the quantity of random numbers that we need is very small — typically a dozen per cycle for about 50000 cycles, where each cycle gets a new seed. Therefore, a prediction attack to reveal the blinding offset would be extremely unlikely to succeed even if the random number generator is not of highest quality. In contrast, the quality of the generator is important in terms of non-correlation and uniformity in order to avoid the danger of introducing noise or a systematic bias to the blinded data. The standard PRNG

of many computer languages, the linear congruential generator, therefore may be not suitable. Furthermore, for the reblinding it is absolutely necessary that the same algorithm should remain available for a significant number of years. Thus, we avoided any sort of library that may vary either over time or between different computers. We decided to use WELL1024a [25]. The Box-Muller transform [26] was used to convert uniformly distributed to normally distributed random numbers where necessary.

The random seed for each cycle must be reproducible over years, and it must be secret after blinding. Our data format, which, as noted above, is an event list of particle detection per channel, includes a periodic counter of accumulated events in every channel. This led to the choice of using a 1024-bit checksum over the last 130 kByte of the unblinded file. Note that if the data files were not to include such a counter, the blinded data file would be very similar to the unblinded one and the seed would not be secret. In such cases an alternative approach would be to use the noise in the detector for the seed creation, *e.g.* from gamma events. For the secondary blinding, the original unblinded data were used for the seed calculation. This would help if a reblinding at both primary and secondary level were ever to become necessary.

3.6 Online blinding

In order to calculate the phase of the actual working point ϕ_i the blinding process requires the knowledge of Φ and α . Yet, this information is available only after a full run has been recorded, since it results from the overall Ramsey fit. Consequently, no blinded data are available before the end of each run. However, it is absolutely necessary to have some live data available for quality checks of the ongoing measurement. An intuitive thought would be to publish a rounded version of the neutron counts in order to disguise the blinding offset. It turns out though, that, in order to make this disguise effective, the rounding must be so coarse that the obtained number would be useless as a quality check. As a solution to this problem, we devised an online blinding mechanism. For this, an additional blinding offset was created randomly for each run. The range for these random numbers with uniform probability distribution was $\pm 1 \times 10^{-23}$ ecm and thus about a hundred times larger than the range of the regular blinding offset. The list of online blinding offsets used was stored in a location with restricted access, and has not been used for any other purpose than debugging the program. The online blinding algorithm does not provide a Φ parameter, that means it assumes a zero magnetic field gradient for the calculation of the number of neutrons to be shifted. Furthermore, it assumed perfectly symmetric detector efficiencies and uses \bar{N} instead of N_i . In all other aspects the algorithm was identical to that used for the regular blinding. The frequency shift introduced through the online blinding was of the same order of magnitude as the naturally occurring magnetic field gradients, either intentionally introduced by the trim coils, or by external fluctuations. With the online blinding system in place, it was possible to make neutron counts available to the user immediately after each cycle without affecting the quality of the data, as far as online checks are concerned. However, these data can obviously not be used meaningfully for any further analysis.

It is important that cycles without electric field are not blinded at all, since they are used by the DAQ to choose the working points to have symmetric neutron

count rates even in the presence of magnetic field gradients.

3.7 Technical details

The supervisory control and data acquisition system of the experiment was partly modular, with all processes and file handling concerning neutron counting being hosted on a dedicated computer running Linux. Time was synchronised among all computers, and control communication was done via Ethernet (TCP/IP). Thus, with simple user permissions provided by the file system we were able to restrict access to the binary code of the blinding program, which contains the blinding offset, as well as to the raw data files. It was particularly beneficial that the computer could be started with a common unprivileged account. The DAQ program and thus the blinding process were given different permissions via the setuid bit. Consequently, the blinding process had access to the secret blinding offset and could write data files that standard users could not read.

A typical run of several days generated about a dozen gigabytes of data. With files of this size, the blinding process took several minutes. We obviously desired immediate feedback about the blinding and any potential problems, but we did not wish to block our DAQ system for as long as that. We therefore split the blinding process into two parts. The first part was to select data and to do the fit of α which was reported to the main DAQ and thus to the user. This could be carried out within a second. The process would then fork itself, on the one hand quitting to make the system available for the following run, while, on the other hand, simultaneously carrying out the intensive work of transferring neutron data between the two detector arms.

During data taking the blinding program was supposed to run autonomously and without intervention. This meant it had to handle some irregular conditions:

- Data that did not contain EDM information must not be blinded; they were instead revealed immediately. These were typically runs without applied high voltage, or runs with cycles that did not have two spin-flips. Such measurements occurred fairly frequently in order to characterise the UCN source [27, 28], the detector, or the background.
- The fitting process ignored single cycles with a low neutron count rate. We chose 1000 counts as a threshold, as such a low count rate would not be used for nEDM analysis.
- Cycles with an unphysically high count rate were not blinded, since these could effectively disclose the blinding offset. We chose a threshold of 50000 neutrons per cycle, ten times more than during commissioning and nearly three times more than the maximum observed.
- Blinding was only automatically applied if the quality of the Ramsey fit was sufficiently good, we chose $\chi_{\text{red}}^2 < 3$.

In case of doubt the blinding process neither blinded nor revealed data, but rather made a request via E-mail for human intervention.

3.7.1 Manual interventions

We took great care during the design of the blinding algorithm to minimize the need for human intervention during data taking. This required automatic handling of unusual circumstances with respect to data quality or malfunctions of parts of the apparatus. Inevitably, due to the complexity of the experiment, some manual interventions during data taking were necessary. In these circumstances the data were assessed by the blinding coordinator in order to either reject bad cycles, and to apply the blinding on the remaining cycles of the run, or divide a run into pieces between magnetic field jumps and apply the blinding on these parts. During data taking in 2015 and 2016, 1072 runs with data files from the neutron detector were recorded. Of these, 113 runs were automatically blinded, 14 runs needed manual blinding and 20 runs needed manual revealing. All of these contained information on the EDM. The other 925 runs were revealed promptly. This is important as these were special characterization measurements that were required immediately for detailed analysis.

3.7.2 Secrecy

The scenarios we want to protect against are the following:

- During the data analysis process somebody might be inclined to do a test that would be simpler to run on the unblinded raw data.
- If forbidden data exists some human beings are tempted to try accessing them, only to prove that they can.
- Others may seek the challenge that there is something which is claimed to be impossible, namely to decrypt the data or to apply statistical attacks on them.

Any of them might be without malicious intent, but it may lead to accidental display of the blinding secret which is a simple number.

The blinding offset was stored using asymmetric encryption with the public part of a RSA-key pair directly after it had been created. The blinding offset together with some metadata only amounts to 192 bits, thus a simple asymmetric encryption is possible. The private key to decrypt the blinding offset was injected into the executable of the blinding program at compile time. Access to the executable program is restricted by file system permissions. The original private key was stored with password RSA encryption using OpenSSL and thus only available to the blinding coordinator. Access to data files was restricted by file system permissions.

These cryptographic and organisational measures were deemed reasonable in order to prevent accidental unblinding of the data. They were easy to implement and did not impact any permitted workflow. Although fairly robust, they are certainly not sufficient to protect against either physical theft of hard drives or manipulation of software with malicious intent. Any further protection would require the restriction of physical access to the DAQ computer or its boot process. Encryption of the operation system via a Trusted Platform Module chip is nowadays available and would suffice for this task. However, this would have been a potential impact on the maintainability of the system, especially in case of hardware problems, hence we considered the existing hurdles high enough for our case.

3.8 Effects of noise and asymmetry

The blinding algorithm manipulates the data. This includes the use of random numbers and fit results. This necessary procedure naturally introduces some noise. In this section we discuss the level of this noise and the resulting consequences.

3.8.1 Noise from fractional neutron numbers

In sect. 3.2.1 we described how a random number (normal distribution with $\sigma = 2$) is added to the fractional number of neutrons to be transferred before rounding to an integer value. Solving eq. (11) for d allows one to calculate how much noise is added to the final nEDM result due to this additional random process. Using the average number of neutrons per cycle $\overline{N} = 11400$, the average visibility $|\overline{\alpha}| = 0.75$, and the applied electric field $E = 11 \text{ kV/cm}$ the additional noise amounts to $7.7 \times 10^{-26} \text{ e cm}$ per cycle. The additional statistical uncertainty for the mean of all 54068 cycles is $3.3 \times 10^{-28} \text{ e cm}$, which is about 3% of the uncertainty due to counting statistics.

3.8.2 Noise from detector asymmetry

In sect. 3.2.2 we described the determination of α and A_m through fitting. These quantities each have their own statistical uncertainty. The mean of the fit value of the visibility $|\alpha|$ was 0.75, and the mean of its uncertainty was 0.003. The mean values of the detector asymmetry were $A_N = 0.032$ and $A_I = -0.036$ in 2015; both with a standard deviation of 0.002. In 2016 the mean was $|A_m| = 0.004$ with a standard deviation of 0.001. The mean of the individual uncertainties within each run was always below 0.001. Thus in 2015 there was a significant asymmetry. The number of neutrons to be transferred is calculated from these numbers via eqs. (17) and (11). At our working points the result of $\sin(\arccos(x))$ lies between 0.98 and 0.99 for any x . Thus no matter how large the fluctuations of A_m may have been, the resulting noise on δN is less than 1% and thus much smaller than the noise arising from the integer rounding described in sect. 3.8.1, and is thus also negligible.

3.8.3 Noise from visibility

The parameter α enters directly in eq. (11), but as the observed relative uncertainty $\frac{0.003}{0.74} = 0.004$ is also very small, the same argument applies once again.

3.8.4 Noise from neutron number per cycle

In eq. (11) also the measured quantity N_i enters. Despite being a noisy observable, it does not contribute to any noise in the blinding, since it is the exact value of the number of neutrons for this particular cycle.

3.8.5 Verification on test data

For a test of the blinding process with real data, we took advantage of the very first data taken before September 13th in 2015. This part of the data was made available to the analysis teams with and without a blinding offset, though the blinded one was not used in the analysis presented in ref. [5]. It was blinded with $d^* = 1.951 \times 10^{-25} \text{ e cm}$ to test the blinding algorithm. These 24 runs have each an irreducible statistical sensitivity that ranges from $0.9 \times 10^{-25} \text{ e cm}$ to $2.4 \times 10^{-25} \text{ e cm}$.

They accumulate to 3.2×10^{-26} e cm. The data was analysed in two stages, once in September 2015, right before the decision was taken to continue with the full implementation of the blinding, at this time obviously with a still quite rudimentary data analysis. A second time this test was done with an almost final analysis. Both tests showed that the blinding algorithm increased the uncertainty by 2×10^{-28} e cm corresponding to 0.5% of the statistical sensitivity of the data set. The blinding offset predicted by the analysis matched the applied within 0.2×10^{-26} e cm which was a tenth of the uncertainty of the analysis. This comparison was done before any unblinding in order to have a better estimate which outcome of the secondary (relative) unblinding should be considered as successful. After unblinding, this test was repeated with the full dataset as described in the next section.

4 Unblinding

Each data analysis team worked on a twice (primary + secondary) blinded data set, and ultimately extracted their own estimator for the blinded nEDM value and its uncertainty. Once the collaboration was convinced that these analyses were complete, a comparison based on appropriate parameters and distributions was done. One comparator was for example the nEDM uncertainty. Moreover, since the data taking was organized in sequences, it was possible to check, sequence by sequence, the difference between the extracted nEDM and its mean value (averaged over all sequences). This difference was useful to check that the two analysis results showed the same correlations with respect to external parameters.

The decision to proceed to the first un-blinding step, which consisted of removing the secondary blinding offsets, was taken based on the agreement of all comparators.

After this first unblinding it was possible to cross-check the two analyses with respect to the secondary blinding offset, see table 1.

Thereafter we compared directly the nEDM values obtained by the two teams. If any discrepancy would have been found, a longer and detailed comparison would have had to be carried out at this point. Among other things, we have discussed to run both analysis codes on a common subset of data and converge parameters and code, *i.e.* cut criteria and methods, until the results match. Another option would have been to produce a new set of secondary blinded data, although this would have been of limited use since by then both analysis teams would implicitly know their offsets. For this reason we rather had considered to produce an alternative blinded data set directly from the original raw data with a new unknown random offset.

Since the two analysis teams were in agreement we proceeded to the removal of the primary blinding after the full completion of the evaluation of all systematic effects. The offset itself was revealed and subtracted, to yield a true nEDM estimator. In addition though, the same analysis codes together with the same settings, *e.g.* for cuts, were applied to the original, non-blinded data set, which had been kept hidden up to that point. The result of the direct analysis of this never-blinded data set should match with that emerging from the analysis of the blinded set minus the applied blinding offset. From theoretical estimation as well as from our experience with the early data taken without blinding we expected this agreement to be on the 10^{-27} e cm level. In the posterior comparison, given in table 1, this was perfectly confirmed. Figure 5 shows a comparison between the injected blinding offset and the one predicted by the analysis from the western team. The non-zero width of the

nEDM estimator	Western analysis team	Eastern analysis team
Twice blinded $\tilde{\tilde{d}}$	$15.39 \times 10^{-26} \text{ e cm}$	$3.80 \times 10^{-26} \text{ e cm}$
Single blinded \tilde{d}	$5.97 \times 10^{-26} \text{ e cm}$	$6.15 \times 10^{-26} \text{ e cm}$
non-blinded d	$-0.02 \times 10^{-26} \text{ e cm}$	$0.16 \times 10^{-26} \text{ e cm}$
$\tilde{\tilde{d}} - \tilde{d}$	$9.55 \times 10^{-26} \text{ e cm}$	$-2.24 \times 10^{-26} \text{ e cm}$
Input offset d''	$9.48 \times 10^{-26} \text{ e cm}$	$-2.33 \times 10^{-26} \text{ e cm}$
Difference $\tilde{\tilde{d}} - \tilde{d} - d''$	$-0.05 \times 10^{-26} \text{ e cm}$	$-0.02 \times 10^{-26} \text{ e cm}$
$\tilde{d} - d$	$5.99 \times 10^{-26} \text{ e cm}$	$5.99 \times 10^{-26} \text{ e cm}$
Input offset d'	$6.02 \times 10^{-26} \text{ e cm}$	$6.02 \times 10^{-26} \text{ e cm}$
Difference $\tilde{d} - d - d'$	$-0.03 \times 10^{-26} \text{ e cm}$	$-0.03 \times 10^{-26} \text{ e cm}$

Table 1: Estimators of the neutron EDM by the two analysis teams. $\tilde{\tilde{d}}$ is the estimator of the twice blinded data. \tilde{d} is the estimator of the primarily blinded data. The input offset d'' is the value of the secondary blinding offset which was de-encrypted during the first, relative unblinding on October 23rd 2019. The input offset d' is the value of the primary blinding offset which was de-encrypted during the second unblinding on November 28th 2019. All analysis results in this table comprise only data taken after September 13th 2019, data prior to this was not blinded with the same offsets and thus cannot be compared. Consequently, the value d listed here differs slightly from the final result [5]. The statistical uncertainty of the never-blinded data shown here was determined by one analysis team to be $1.07 \times 10^{-26} \text{ e cm}$ with $\chi^2/N_{\text{dof}} = 92.5/86$.

peaks indicates that the blinding algorithm does inject some noise into individual sequences or cycles. The widths of the Gaussian fitted to the distribution were $0.31(5) \times 10^{-26} \text{ e cm}$ and $0.41(4) \times 10^{-26} \text{ e cm}$ for secondary and primary blinding, respectively. And indeed, the blinding injects some noise in each cycle as explained in sect. 3.8.1. The sequences consisted of 514 cycles on average. Thus, the observed widths are compatible with the expected uncertainty of the mean due to the noise of $0.34 \times 10^{-26} \text{ e cm}$.

The agreement of the difference ‘analysis of blinded data’ minus ‘analysis of never-blinded data’ with the blinded offset is marvellous.

5 Costs and benefits

As discussed in ref. [29], blinding does not come for free. For the method presented here, the costs were primarily the manpower required for design, implementation and study of the technique. As noted above, our blinding introduces a small amount of statistical noise into the blinded data. This tiny, additional noise was only present in the blinded datasets; the final analysis was never affected by this. Ultimately, we did not suffer from various costs that are typical for other blinding techniques. For example, in our case all analysis channels were immediately available and no signals or features, other than the true nEDM itself, was hidden. Notably, the blinding permitted the analysis of a periodically changing nEDM [30] without revealing the unblinded result of the static signal.




Figure 5: Difference of analysis on blinded and unblinded data sets and the corresponding offset for the two blinding steps. The bin width is 10^{-27} ecm. Both peaks are centred well within 10^{-27} ecm. Only data of the western analysis and after September 13th 2015 is shown. The eastern analysis yields similar results.

Most importantly, the blinding provided a very substantial benefit to the nEDM analysis, not only in that it eliminated the effects of an unconscious bias. Since in the past the true nEDM result has always been indistinguishable from zero, it was sign insensitive, and as such also insensitive to potential sign errors in the analysis. However, once blinded, the signal in the data had a value significantly away from zero, and thus included a clearly identifiable sign. This sign showed up in the various analysis channels, *e.g.* with its dependence upon magnetic-field gradients, and as such at one point it actually revealed a mistake in an early version of our data analysis when the code was tested with a known fake nEDM.

6 Possible Improvements

In order to handle the non-integer number of neutrons to be transferred in each single cycle we used a normally distributed random number of width 2 as described in sect. 3.2.1. Future implementations will use a rectangular probability density function of width 1. This will provide perfect linearity and will reduce the introduced noise by a factor of two, down to the intrinsic minimum.

7 Summary & Conclusion

We have developed and applied for the first time for an nEDM measurement a blinding technique. Our algorithm only modifies a copy of already recorded data and saves the original data on a hidden and secret disk. Secondary blinding and the re-blinding option are innovations which reduce risks that are often associated with blinding. The artificial increase in noise in the blinded datasets through the process was negligibly small and disappears automatically in the final result.

8 Acknowledgement

We gratefully acknowledge financial support from the Swiss National Science Foundation through project grants 149211, 162574 and 172639 for the ETHZ group, grant 181996 for the Bern group and grants 117696, 144473, 137664, 163413, 169596, and 171626 for the PSI group. This work was funded in part by the United Kingdom Science and Technology Facilities Council (STFC) through grants ST/N000307/1 and ST/M503836/1, as well as by the School of Mathematical and Physical Sciences at the University of Sussex. One of the authors (P. M.) would like to acknowledge support from the Swiss State Secretariat for Education, Research and Innovation (SERI) - Federal Commission for Scholarships (FCS) award #2015.0594. We gratefully acknowledge financial support from the Polish National Science Center, grants UMO-2015/18/M/ST2/00056, UMO-2016/23/D/ST2/00715 and UMO-2018/30/M/ST2/00319. We acknowledge the excellent support provided by the PSI technical groups and by various services of the collaborating universities and research laboratories.

References

- [1] S. K. Lamoreaux and R. Golub. Experimental searches for the neutron electric dipole moment. *J. Phys. G*, 36(10):104002, September 2009. DOI: 10.1088/0954-3899/36/10/104002.
- [2] C. A. Baker, D. D. Doyle, P. Geltenbort, et al. Improved experimental limit on the electric dipole moment of the neutron. *Phys. Rev. Lett.*, 97:131801, September 2006. DOI: 10.1103/PhysRevLett.97.131801.
- [3] J. M. Pendlebury, S. Afach, N. J. Ayres, et al. Revised experimental upper limit on the electric dipole moment of the neutron. *Phys. Rev. D*, 92:092003, November 2015. DOI: 10.1103/PhysRevD.92.092003.
- [4] A. P. Serebrov, E. A. Kolomenskiy, A. N. Pirozhkov, et al. PNPI differential EDM spectrometer and latest results of measurements of the neutron electric dipole moment. *Phys. At. Nucl.*, 78:1601–1605, 2015. DOI: 10.1134/S1063778815130293.
- [5] C. Abel, S. Afach, N. J. Ayres, et al. Measurement of the permanent electric dipole moment of the neutron. *Phys. Rev. Lett.*, 124:081803, February 2020. DOI: 10.1103/PhysRevLett.124.081803.
- [6] J. Ellis. Theory of the neutron electric dipole moment. *Nucl. Instrum. Methods Phys. Res. A*, 284(1):33–39, 1989. DOI: 10.1016/0168-9002(89)90243-X.
- [7] D. E. Morrissey and M. J. Ramsey-Musolf. Electroweak baryogenesis. *New J. Phys.*, 14:125003, 2012. DOI: 10.1088/1367-2630/14/12/125003.
- [8] J. Engel, M. J. Ramsey-Musolf, and U. van Kolck. Electric dipole moments of nucleons, nuclei, and atoms: The standard model and beyond. *Prog. Part. Nucl. Phys.*, 71:21–74, 2013. DOI: 10.1016/j.ppnp.2013.03.003.

- [9] J. R. Klein and A. Roodman. Blind analysis in nuclear and particle physics. *Annu. Rev. Nucl. Part. S.*, 55(1):141–163, 2005. DOI: 10.1146/annurev.nucl.55.090704.151521.
- [10] J. Adam, X. Bai, A. M. Baldini, et al. New limit on the lepton-flavor-violating decay $\mu^+ \rightarrow e^+ \gamma$. *Phys. Rev. Lett.*, 107:171801, October 2011. DOI: 10.1103/PhysRevLett.107.171801.
- [11] E. Aprile, M. Alfonsi, K. Arisaka, et al. Dark matter results from 225 live days of XENON100 data. *Phys. Rev. Lett.*, 109:181301, November 2012. DOI: 10.1103/PhysRevLett.109.181301.
- [12] J. Abadie, B. P. Abbott, R. Abbott, et al. All-sky search for gravitational-wave bursts in the first joint LIGO-GEO-Virgo run. *Phys. Rev. D*, 81:102001, May 2010. DOI: 10.1103/PhysRevD.81.102001.
- [13] D. B. Chitwood, T. I. Banks, M. J. Barnes, et al. Improved measurement of the positive-muon lifetime and determination of the Fermi constant. *Phys. Rev. Lett.*, 99:032001, July 2007. DOI: 10.1103/PhysRevLett.99.032001.
- [14] A. Vacheret and D. Lhuillier. A precise measurement of $\sin^2 \theta_W$ at low Q^2 in Møller scattering. *Eur. Phys. J. A*, 24(2):151–154, February 2005. DOI: 10.1140/epjad/s2005-04-038-6.
- [15] C. Abel, N. J. Ayres, G. Ban, et al. nEDM experiment at PSI: Data-taking strategy and sensitivity of the dataset. *EPJ Web Conf.*, 219:02001, 2019. DOI: 10.1051/epjconf/201921902001.
- [16] S. Afach, G. Ban, G. Bison, et al. A device for simultaneous spin analysis of ultracold neutrons. *Eur. Phys. J. A*, 51(11):143, 2015. DOI: 10.1140/epja/i2015-15143-7.
- [17] G. Ban, G. Bison, K. Bodek, et al. Ultracold neutron detection with ^6Li -doped glass scintillators. NANOSC: A fast ultracold neutron detector for the nEDM experiment at the Paul Scherrer Institute. *Eur. Phys. J. A*, 52:326, October 2016. DOI: 10.1140/epja/i2016-16326-4.
- [18] C. A. Baker, Y. Chibane, M. Chouder, et al. Apparatus for measurement of the electric dipole moment of the neutron using a cohabiting atomic-mercury magnetometer. *Nucl. Instrum. Methods Phys. Res. A*, 736:184–203, 2014. DOI: 10.1016/j.nima.2013.10.005.
- [19] J. M. Pendlebury, W. Heil, Yu. Sobolev, et al. Geometric-phase-induced false electric dipole moment signals for particles in traps. *Phys. Rev. A*, 70:032102, September 2004. DOI: 10.1103/PhysRevA.70.032102.
- [20] S. Afach, N. J. Ayres, C. A. Baker, et al. Gravitational depolarization of ultracold neutrons: Comparison with data. *Phys. Rev. D*, 92:052008, September 2015. DOI: 10.1103/PhysRevD.92.052008.
- [21] C. Abel, S. Afach, N. J. Ayres, et al. Optically pumped Cs magnetometers enabling a high-sensitivity search for the neutron electric dipole moment. *Phys. Rev. A*, 101:053419, May 2020. DOI: 10.1103/PhysRevA.101.053419.

- [22] V. Helaine. *Neutron Electric Dipole Moment measurement: simultaneous spin analysis and preliminary data analysis*. Thesis, Université de Caen, September 2014. URL <https://tel.archives-ouvertes.fr/tel-01063399>.
- [23] S. Afach, C. A. Baker, G. Ban, et al. A measurement of the neutron to ^{199}Hg magnetic moment ratio. *Phys. Lett. B*, 739:128–132, 2014. DOI: 10.1016/j.physletb.2014.10.046.
- [24] LPC Caen. FASTER Fast Acquisition System for nuclEAR Research, 2013. <http://faster.in2p3.fr/>.
- [25] F. Panneton, P. L’Ecuyer, and M. Matsumoto. Improved long-period generators based on linear recurrences modulo 2. *ACM Trans. Math. Softw.*, 32(1): 1–16, March 2006. DOI: 10.1145/1132973.1132974.
- [26] G. E. P. Box and M. E. Muller. A note on the generation of random normal deviates. *Ann. Math. Statist.*, 29(2):610–611, June 1958. DOI: 10.1214/aoms/1177706645.
- [27] G. Bison, B. Blau, M. Daum, et al. Neutron optics of the PSI ultracold-neutron source: characterization and simulation. *Eur. Phys. J. A*, 56:33, 2020. DOI: 10.1140/epja/s10050-020-00027-w.
- [28] B. Lauss. Ultracold Neutron Production at the Second Spallation Target of the Paul Scherrer Institute. *Phys. Proc.*, 51:98–101, 2014. DOI: 10.1016/j.phpro.2013.12.022.
- [29] R. Golub. Blind analysis in physics experiments: Is this trip necessary? In S. J. Seestrom, editor, *Next Generation Experiments to Measure the Neutron Lifetime*, pages 97–102, March 2014. DOI: 10.1142/9789814571678_0010.
- [30] C. Abel, N. J. Ayres, G. Ban, et al. Search for axionlike dark matter through nuclear spin precession in electric and magnetic fields. *Phys. Rev. X*, 7:041034, November 2017. DOI: 10.1103/PhysRevX.7.041034.