THE UNIVERSITY OF CHICAGO

TOWARD A NEXT-GENERATION MEASUREMENT OF THE ELECTRON ELECTRIC DIPOLE MOMENT

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY
PEIRAN HU

CHICAGO, ILLINOIS
MARCH 2026

ORCID iD: https://orcid.org/0009-0004-1367-671X

To family and friends

"Blah, blah, blah. Did I miss anything?" — Jinx

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

This work would not have been possible without the support and guidance of many people.

First and foremost, I would like to NSF, the Gordon and Betty Moore Foundation, and the Alfred P. Sloan Foundation for funding this research.

I am deeply grateful to my advisor, Prof. David DeMille, for his free-range mentoring style. Dave is like a unicorn to me: he possesses every trait I could want in an advisor and physicist, and he appears less frequently than one might expect—but always at exactly the right moment. Despite his busy schedule, Dave always manages to quickly grasp the details of the experiment and provide insightful feedback. His influence can be seen in almost every aspect (at least the correct ones) of this thesis, even though I may not explicitly acknowledge it. Dave's trust and encouragement have allowed me to grow tremendously as both a scientist and a person over the years.

I would also like to thank the other two PIs of the ACME collaboration, Prof. John Doyle and Prof. Gerald Gabrielse. John showed me incredible hospitality when I was still working on the Harvard beamline. His engineering expertise and science leadership have been invaluable to the success of ACME III. Jerry also made tremendous contributions through his deep hardware knowledge and outstanding students. He always encouraged us to try really really hard and make really really sure that the apparatus would work smoothly and reliably, which is a critical aspect of a precision measurement experiment.

I am grateful for the incredible support and mentorship of my committee members, Prof. Carlos Wagner, Prof. David Miller, and Prof. Zoe Yan. I did not always make reasonable requests, but they were always patient and understanding.

Many thanks to other faculty and staff who have helped me along the way, including Prof. Zosia Krusberg, Putri Kusumo, Shashi Dyamenahalli, Amima Diagne, Caterina Mackowiak, Brenda Thomas, Samuel Roberts, Daphne Klemme, Salomon Rodriguez, Steven

Sansone, David Johnson, Jason Tresback, Erica Mantone, Joel Day, and Bonnie Currier.

I am indebted to the work from the previous ACME members. It was also my privilege to be able to work with many talented colleagues in the ACME collaboration over the years: Ayami Hiramoto, Bingjie Hao, Cole Meisenhelder, Collin Diver, Cristian Panda, Daniel Ang, James Chow, Koji Yoshimura, Maya Watts, Nick Hutzler, Noboru Sasao, Satoshi Uetake, Siyuan Liu, Takahiko Masuda, Xing Fan, Xing Wu, Zack Lasner, and Zhen Han. In particular, Xing Wu and Xing Fan acted as my unofficial mentors during our overlap. Xing Wu helped me get up to speed quickly when I first joined ACME and was often very effective at quickly generating ideas and feedback. Xing Fan acted as an effective PI after joining ACME. In addition to his many contributions in the lab, he provided important impedance matching between the students and the PIs.

Special thanks to Collin Diver and Zhen Han, with whom I worked most closely during my PhD. Collin Diver taught me physics, American culture, and Texas. He is quick to point out mistakes in my thinking—both in physics and in life—and even quicker to stay quiet once I become upset by those corrections. Zhen Han taught me a great deal about the experiment when I first joined ACME and later offered to help with thesis figures after hearing that I had started writing. His keen eye frequently allowed him to notice things that Collin and I missed. Despite frequent hallucinations, Zhen never ceases to make 99.999%-confident predictions that are either correct or hilariously wrong. The three of us worked together with remarkable efficiency, in a manner reminiscent of SpongeBob, Patrick, and Squidward.

It was a pleasure to learn from the other DeMille group members over the years: Olivier Grasdijk, Mangesh Bhattarai, Silviu-Marian Udrescu, Wesley Cassidy, Abby Kinney, Emma McClure, Pranav Chamakkad Muthukrishnan, Mohit Verma, Qian Wang, Emma McClure, Mohit Verma, Jianwei Wang, Qian Wang, Robert Xi, Shaozhen Yang, Yuanhang Yang, Geoffrey Zheng, Harish Ramachandran, Kevin Su, Akiva Davis, Michelle Hu, Juan Ignacio Prieto, José Manuel Escobar Rodriguez, Grace Simmons, Thomas Langin, Rohan Kapur,

Varun Jorapur, Jakob Kastelic, Jack Dewhurst, Oskari Timgren, Yuqi Zhu, Jack Roth, Maximilian Beyer, Trevor Wright, and Sid Cahn. In particular, Olivier, Oskari, Qian, and Mohit were always patient in answering any stupid question that I had.

I would also like to thank my non-ACME friends and colleagues in the Gabrielse group: Andra Ionescu, Banibrato Sinha, Benedict Sukra, HyunSeung Kim, Lily Soucy, Nathaniel McDonough, Tharon Morrison, Thomas Myers, Shungo Fukaya, and Alyssa Matthews. It was always fun to discuss science with Nathaniel and even more fun to complain about life with Ben.

I am also grateful to my friends and colleagues in the Doyle group: Yicheng Bao, Zack Lasner, Ben Augenbraun, Loïc Anderegg, Arian Jadbabaie, Sean Burchesky, Alex Frenett, Debayan Mitra, Nathaniel Vilas, Hiromitsu Sawaoka, Christian Hallas, Scarlett Yu, Abdullah Nasir, Annika Lunstad, Paige Robichaud, Mingda Li, and Grace Li. In particular, Yicheng and Loïc consistently offered helpful suggestions on experimental techniques during group meetings. Zack contributed hands-on work to ACME III in addition to his helpful thesis. The YbOH team also helped move ACME out of G14 and G16.

My journey with ACME began during my undergraduate years in Paul Hamilton's lab. Adam West was my unofficial mentor and taught me a great deal about experimental physics. I also learned a great deal from Randy Putnam, Elizabeth West, Rob Niederriter, Chandler Schlupf, Sami Khamis, and Kayla Rodriguez. Outside the lab, I received invaluable guidance from Per Kraus, Michael Gutperle, and Joshua Samani.

I must also thank my imaginary friends in the UChicago Physics Department and my real friends in the Yale Physics Department. It was my privilege to learn alongside of Zachary Gao Sun, Hui Liang Peng, Tianyu Yuan, Yuxuan Cheng, and Joseph Shomar.

Finally, I would like to express my deepest gratitude to my family. My father told me that it was the teacher's problem after I got a D on my first physics exam. My mother claims that she only became goofy after giving birth to me, evidence that I inherited whatever

intelligence she had. My grandmother never cared about science but always cared about my well-being. My grandfather married my grandmother. I think that was his greatest achievement in life.

ABSTRACT

Over the past decades, measurements of the electron electric dipole moment (eEDM, d_e) have emerged as a powerful probe of physics beyond the Standard Model. The current best limit, $|d_e| < 4.1 \times 10^{-30}$ e·cm, set by the JILA EDM experiment, improves upon the previous ACME II result by a factor of two. This thesis presents the progress and developments of the ACME III experiment, which aims to further improve the sensitivity to d_e by an additional order of magnitude. Such an improvement will enable probes of new physics at the 10 TeV scale and beyond using a tabletop experiment.

To achieve this target sensitivity, several major upgrades have been implemented in ACME III. These include engineering efforts to increase the molecular flux, extend the spin precession time, and improve the readout efficiency. In parallel, new engineering controls have been developed to mitigate known sources of systematic error. This thesis describes these upgrades in detail and documents the progress made toward understanding and controlling systematic effects.

CHAPTER 1

INTRODUCTION

1.1 The Standard Model and its limitations

The Standard Model (SM) of particle physics is undoubtedly one of the most successful theories in modern physics. From the prediction and observation of the Higgs boson [1, 2] to the precise calculation and measurement of the electron anomalous magnetic dipole moment [3, 4], the direct predictions of the Standard Model have so far survived all experimental tests. However, there is plenty of indirect evidence indicating that the Standard Model is not the ultimate theory of nature. The biggest mysteries can be briefly summarized as follows:

1. **Dark matter and dark energy:** The SM contains no viable candidate to explain the observed dark matter or dark energy in the universe [5].

2. **Neutrino mass and flavor oscillation:** Within the minimal SM, neutrinos are massless, which is inconsistent with the observed neutrino flavor oscillations [6].

3. **Matter–antimatter asymmetry:** The SM contains CP-violating sources that are too small to explain the observed matter–antimatter asymmetry in the universe [7–9].

4. **Hierarchy problem:** The SM lacks a natural explanation for the large separation between the electroweak scale and the Planck scale [10].

5. **Gravity:** The SM does not include gravity [11].

These mysteries strongly suggest the existence of new physics beyond the Standard Model (BSM). Due to the lack of direct evidence for BSM physics at high-energy colliders such as the LHC, there are few guiding principles beyond mathematical consistency and

naturalness for theorists to construct BSM models. Among the existing elegant frameworks, supersymmetry (SUSY) is perhaps the most well-known and extensively studied [12, 13].

SUSY introduces a symmetry that relates known fermions (bosons) to new bosonic (fermionic) superpartners. As a consequence of this symmetry, the contributions of superpartners can cancel large quantum corrections to the Higgs mass, thereby providing a natural solution to the hierarchy problem. Furthermore, potentially stable superpartners can act as Weakly Interacting Massive Particles (WIMPs) and explain the observed dark matter abundance [14]. In addition, SUSY generically introduces new CP-violating phases, providing additional sources of CP violation beyond the SM that may explain the matter–antimatter asymmetry [15]. Lastly, SUSY arises naturally in string theory, which is a leading candidate for a theory of quantum gravity [16].

The search for BSM physics such as SUSY was one of the main motivations for the construction of the LHC, but no direct evidence has been found to date [17, 18]. This thesis presents a BSM physics search effort complementary to high-energy collider experiments [19]; it is part of a broader class of non-collider experiments searching for new physics [20, 21]. More specifically, the ACME experiment searches for new CP-violating physics by probing the electron electric dipole moment (eEDM or d_e) [22, 23].

The Standard model predicts a value of d_e that is many orders of magnitude below the current experimental sensitivity [20, 23, 24]. Thus, any non-zero measurement of d_e would be a clear sign of BSM physics. On the other hand, a null measurement can be interpreted as a constraint on BSM physics. Quantitatively, we can make order-of-magnitude estimates for an n-loop Feynman diagram,

$$\Lambda^2 = e\,\frac{m_e}{\delta d_e}\left(\frac{\alpha}{4\pi}\right)^n \sin\phi_{\mathrm{CP}}, \tag{1.1}$$

where Λ is the energy scale of new CP-violating physics, m_e is the electron mass, δd_e is the uncertainty in the measured d_e, α is the fine-structure constant, and ϕ_{CP} is the new

CP-violating phase [25]. A null measurement rules out BSM physics at energy scales Λ or below. Assuming $\sin\phi_{\rm CP} \sim 1$ and $n \leq 2$ [26], the previous generation of the experiment was already ruling out BSM physics at the ~ 3 TeV scale[1] [23]. As we continue to improve the experimental sensitivity δd_e, we can either rule out BSM physics at progressively higher energy scales or eventually observe a non-zero d_e and provide crucial experimental guidance for constructing the ultimate theory of nature.

1.2 Overview of eEDM experiments

Given that we want to experimentally measure the eEDM, how should we do it? Without the benefit of hindsight, this is an extremely difficult problem. The interaction of an electron with its environment is dominated first by its electric monopole moment and then by its magnetic dipole moment; the effect of its electric dipole moment is vanishingly small in comparison. Moreover, an electron is difficult to trap in isolation and measure with high precision. The key conceptual breakthrough occurred when physicists realized that it is possible to measure the eEDM in atomic or molecular bound systems [28].

Before explaining this breakthrough, it is useful to first understand why measuring the eEDM in a bound system initially seemed impossible. In such a system, the eEDM Stark interaction can be characterized using an effective average electric field felt by the electron, $\vec{\mathcal{E}}_{\rm eff}$. Naively, this average electric field is expected to be zero; otherwise, the electron would rapidly redistribute its wavefunction until the average field vanished. This intuitive argument can be made fully rigorous and is known as Schiff's theorem [29]. However, Schiff's theorem relies on several assumptions that are not fully satisfied in real atomic or molecular systems. For our purposes, the most important assumption is that the electron moves non-relativistically. The key realization was that in atoms or molecules containing heavy nuclei, the valence electrons can move relativistically near the nucleus, thereby violating this

1. For a more general comparison, see Fan [27].

assumption of Schiff's theorem to an almost maximal extent [28, 30–32]. For example, in thorium monoxide, the molecule used by ACME, the effective electric field is $\mathcal{E}_{\text{eff}} \approx 78$ GV/cm [33, 34]. This is orders of magnitude larger than any electric field that can be applied in a laboratory setting.

Figure 1.1 shows the history of eEDM experimental limits in atomic and molecular systems. Initial rapid progress was made in atomic systems such as thallium (Tl) and cesium (Cs) [35–40], while molecular eEDM experiments have taken over in the last decade or two [20, 22, 23, 41–44]. Some additional discussion is required to explain why molecules provide superior platforms for eEDM measurements.

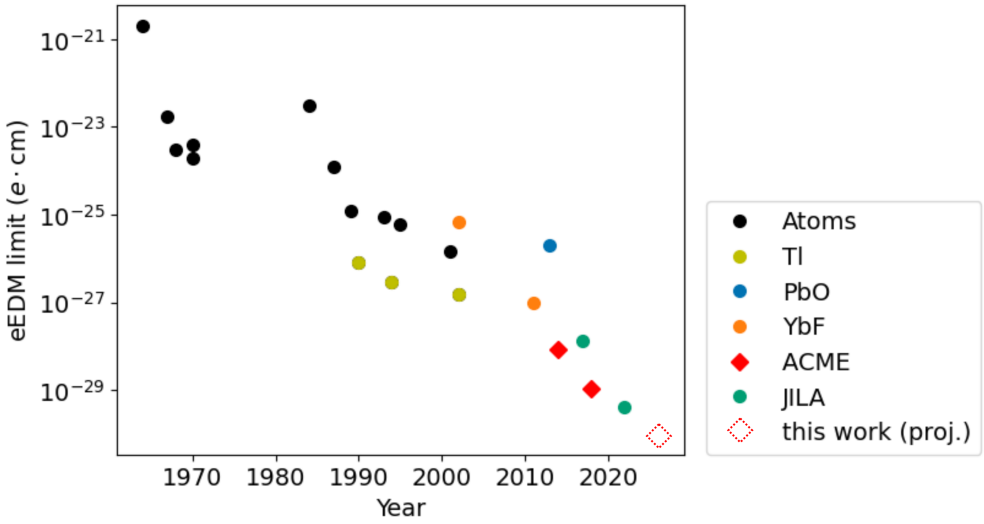

Figure 1.1: History of eEDM experimental limits. Tl: best limit for atomic eEDM [45]; YbF: first molecular eEDM experiment [42]; PbO: first molecular eEDM experiment using the so-called Ω-doublet structure [41]; JILA: current eEDM record using HfF^+ ions [20]; ACME (this work): ThO molecular eEDM experiment [23]. The triangle marker indicates the projected limit based on the current experimental sensitivity. Data points compiled by Xing Fan.

The electron electric dipole moment $\vec{d}_e$ is a vector, and an important question is: in which direction does this vector point? The answer is that it always points along the electron

spin $\vec{S}$, such that[2]

$$\vec{d}_e = d_e \hat{S}. \tag{1.2}$$

But why must this be the case? One common explanation invokes the Wigner–Eckart theorem, although this explanation can be somewhat misleading. While the Wigner–Eckart theorem implies that $\vec{d}_e$ must be colinear with $\vec{S}$, it does not determine discrete transformation properties. As a result, the most that can be concluded from this argument alone is $\vec{d}_e = \pm d_e \hat{S}$, with the relative sign potentially flipped under combinations of C, P, and T transformations. A more satisfactory explanation follows from the broader reasoning outlined above. The motivation for studying the eEDM is that most natural extensions of the SM predict its existence, and in all such predictions one finds $\vec{d}_e = d_e \hat{S}$ because they obey the fundamental principles of quantum field theory (QFT). From the perspective of symmetry and degrees of freedom, the only internal pseudovector associated with the electron is its spin[3]; any vector quantity not aligned with Eq. (1.2) would constitute an additional internal degree of freedom beyond those allowed in QFT.

A direct corollary of Eq. (1.2) is that the eEDM violates both P and T symmetries. This is readily seen because the electric field $\vec{\mathcal{E}}$ is a vector that changes sign under P but not under T, so the eEDM interaction Hamiltonian

$$H_{\text{eEDM}} = -\vec{d}_e \cdot \vec{\mathcal{E}} = -d_e \hat{S} \cdot \vec{\mathcal{E}} \tag{1.3}$$

is odd under both P and T transformations. By the CPT theorem[4], CP symmetry is also violated. In this sense, the electron behaves like an antenna for new CP-violating physics,

2. Here, d_e can be either positive or negative.

3. In contrast, composite particles like the ThO molecule have additional internal degrees of freedom, such as the molecular axis, in addition to their total angular momentum. This means that the molecular EDM does not have to lie along the direction of the total angular momentum.

4. This assumption is used throughout the thesis. It can be derived under reasonable conditions [46], and no experimental violation of CPT has been observed to date [47–49].

and the eEDM points along the direction of this antenna, $\hat{S}$, because there is no other internal direction available.

Given these discrete symmetry properties, we are now ready to explain why molecules are superior platforms for eEDM measurements. The Hamiltonian for an atomic or molecular bound system can be written as

$$H_{\text{total}} = H_{\text{bnd}} + H_{\text{eEDM}} + H_{\text{ext}}, \tag{1.4}$$

where H_{bnd} is the internal Hamiltonian of the bound system, H_{eEDM} is the eEDM interaction Hamiltonian, and H_{ext} describes interactions with external electric and magnetic fields. Ignoring H_{ext} for the moment, because H_{bnd} is even under parity while H_{eEDM} is odd, the two Hamiltonians cannot be simultaneously diagonalized; the already tiny H_{eEDM} interaction is therefore further suppressed to second order in perturbation theory by the energy gap between opposite-parity states. This suppression can be reduced, or even completely lifted, by applying an external electric field to mix parity eigenstates. In atomic systems this is difficult to achieve in practice because the energy gap between opposite-parity states is typically on the scale of eV. In molecular systems, however, opposite-parity states are separated by at most a few rotational constants, corresponding to an energy scale of 10^{-3} eV or smaller [50]. In certain diatomic molecules with Λ- or Ω-doublet structure, and more generally in polyatomic molecules, this separation can be further reduced due to additional internal degrees of freedom [51]. Once the parity states are fully mixed, no suppression of H_{eEDM} remains, and increasing the external electric field provides no additional benefit. There can, however, still be practical advantages to operating at smaller external electric fields, such as reduced sensitivity to systematic errors [43, 45].

1.3 ThO as an eEDM platform

As explained in the last section, it is advantageous to measure the electron electric dipole moment using molecules. The ACME experiment performs the eEDM measurement in a cold molecular beam of thorium monoxide (ThO). This molecule offers several advantages as an eEDM measurement platform (roughly in order of importance):

1. large effective internal electric field $\mathcal{E}_{\text{eff}} \approx 78$ GV/cm [33, 34], inducing a large eEDM energy shift despite the small expected value of d_e;

2. Ω-doublet structure in the $H\,^3\Delta_1$ science state with ms-scale lifetime[5] [52]. This allows full polarization of the molecule in a small external electric field and spectroscopic reversal of molecular orientation[6], fully mitigating several known systematic errors of atomic eEDM experiments [45];

3. it can be created in a slow, cold, and bright ($\sim 10^{13}$ molecules/sr/s) molecular beam using the technique of a cryogenic buffer-gas beam source [53–55];

4. convenient optical transitions, as defined by the laser technology standards and budgets available ~ 15 years ago [56];

5. no nuclear spin. Neither thorium-232 nor oxygen-16 has nuclear spin, simplifying the angular momentum coupling. This is helpful for interpreting a non-zero eEDM signal[7]

5. See Sec. 2.1 for more details

6. Also known as an internal co-magnetometer. Flipping the molecular axis flips the energy correction due to the eEDM without flipping the Zeeman shift, which is governed purely by angular momentum, thereby allowing us to separate the two effects.

7. There are multiple sources of T- or CP-violation expected in molecules. In the absence of nuclear spin, there is an additional eEDM-independent source of CP violation via the scalar lepton–nucleon interaction coefficient C_S [33, 34]. With nuclear spin, the heavy nucleus may contribute via its Schiff moment and via the tensor lepton–nucleon interaction coefficient C_T [57]. The C_S term turns out to be the dominant SM CP-violating effect in ThO [58]. We typically do not discuss this term because it is more complicated to explain. Similar to the eEDM, the C_S term is also expected to be small within the SM alone, and observation of a nonzero value would constitute clear evidence of BSM physics.

and is practically helpful for modeling systematic errors[8].

As we will explain in Chapter 2, the existence of a non-zero d_e causes a relative energy shift of size $2d_e\mathcal{E}_{\text{eff}}$ between two molecular states related by a time-reversal operation. By performing precision spectroscopy to measure this energy shift, we can infer the value of d_e with high precision through the large internal electric field $\mathcal{E}_{\text{eff}}$ and improved experimental sensitivity.

1.4 Other EDM searches

This thesis exclusively focuses on the measurement of the electron electric dipole moment using ThO molecules. However, other eEDM searches are also actively ongoing. Notably, but not exclusively, JILA is building their next-generation bucket-brigade experiment using a new ion species, ThF^+ [61–63]. Imperial College London is assembling their next-generation YbF molecular beam experiment [64]. A proposal to build a solid-state eEDM experiment is also being pursued [65].

Beyond eEDM searches, there are also active efforts targeting other types of EDMs. In particular, searches for nuclear EDMs have become increasingly prominent in recent years. The current best limit comes from the atomic mercury-199 experiment [66]. Another beam experiment using thallium fluoride (TlF) molecules is also ongoing within the DeMille group [67]. Due to the additional BSM signal amplification offered by nuclei with octupole deformation [68, 69], many groups are shifting to species containing these rare nuclei. An experiment assembling francium (Fr) atoms with silver (Ag) atoms to form FrAg molecules is being developed with rapid progress within the DeMille group [70, 71]. A cryogenic crystal experiment using europium (Eu) doped into yttrium orthosilicate (YSO) is being developed in Toronto [72, 73]. Xing Fan's group at Harvard is building a molecular ion trap experi-

8. This is primarily a practical advantage. On the other hand, work from the Hutzler and Doyle groups has demonstrated substantial gains offered by additional internal degrees of freedom [59, 60].

ment similar to JILA's, but targeting a nuclear EDM using ^{229}Th. Xing Wu's QuESTLAB at MSU is developing a thorium-containing trapped-molecule experiment using novel slowing and cooling techniques [74]. Hosted at Harvard, the RaX collaboration is building a trapped radium (Ra)-containing molecular experiment [75]. On the atomic side, a radium-225 experiment is being pursued both at Argonne and at USTC, following their recent success with ytterbium as a proxy [76, 77].

Moving beyond nuclear EDMs, an interesting nuclear MQM search using YbOH is being undertaken at Caltech [59]. Searches for neutron and muon EDMs are also ongoing, typically using experimental techniques significantly different from those employed in atomic and molecular systems [78–83].

1.5 Structure of this thesis

The ACME experiment has already been very well documented in previous theses [52, 53, 56, 84–93]. This thesis focuses on new progress and updates. Sufficient background is usually provided to make the new developments understandable, but the interested reader is encouraged to consult previous ACME publications for further details.

In Chapter 2, we provide an overview of the experimental methods used in ACME. ACME is an engineering-heavy experiment, and we include lessons learned that may be of general interest to other experimentalists. In Chapter 3, we explain the data analysis methods used in ACME, with a focus on the new methodologies adopted in this generation of the experiment. In Chapter 4, we present our current understanding of systematic errors and excess noise in the experiment. One of the systematic errors identified arises from decays of a metastable state. We believe this effect may be of general interest and present our modeling efforts in Chapter 5. Finally, in Chapter 6, we summarize the current status of the experiment with preliminary results.

CHAPTER 2

EXPERIMENTAL METHODS

2.1 Molecular structure

Numerous excellent presentations of the molecular structure of ThO can be found in previous ACME theses. I present my own approach in this section with the right mixture of hand-waving[1]. I do not intend to make any serious derivations, but rather to provide a qualitative understanding of the molecular structure and notation used in ACME. Relevant molecular constants will also be summarized for easy reference.

2.1.1 TL;DR

Thankfully, both thorium-232 and oxygen-16 are spinless isotopes and have $> 99\%$ natural abundance. This lack of nuclear spin simplifies the angular momentum coupling and provides a clean interpretation platform, with the main T-violating signals being the eEDM d_e and the electron–nucleon coupling C_S. The ground and lowest-lying metastable states of thorium monoxide are well described by the Hund's case (a) coupling scheme [50]. This includes the $H\ ^3\Delta_1$ science state used for the eEDM measurement; although the excited states are better described by Hund's case (c), we restrict most of the mathematical descriptions of the quantum states to the Hund's case (a) basis for simplicity.

In the Hund's case (a) coupling scheme, the molecular wavefunction consists of separable parts

$$|\Psi_{\text{total}}\rangle = |\psi_{\text{electronic}}\rangle\,|\psi_{\text{vibrational}}\rangle\,|\psi_{\text{rotational}}\rangle \tag{2.1}$$

$$= |\eta, \Lambda, \Sigma, \Omega\rangle\,|\nu\rangle\,|J, M, \Omega\rangle\,, \tag{2.2}$$

1. Judged entirely based on my personal preference.

where η is the quantum number for the electronic state; Λ is the projection of the electronic orbital angular momentum onto the nuclear axis; Σ is the projection of the electronic spin angular momentum onto the nuclear axis; $\Omega = \Lambda + \Sigma$ is the projection of the total electronic angular momentum onto the nuclear axis; ν is the quantum number for the vibrational state; J is the total molecular angular momentum; and M is the projection of the total angular momentum onto the lab Z axis. The separable parts are ordered in decreasing energy scale. There are four energy scales of interest:

$$E_{\text{electronic}} \gg E_{\text{spin-orbit}} \approx E_{\text{vibrational}} \gg E_{\text{rotational}}. \tag{2.3}$$

This relative energy scale justifies the Born–Oppenheimer approximation, and specifically the choice of the Hund's case (a) coupling scheme. By convention, the electronic and spin–orbit quantum numbers are labeled by a letter denoting the electronic state (e.g. $\eta = X, H, Q$) followed by the term symbol ${}^{2S+1}\Lambda_{|\Omega|}$, where Λ is represented by Greek letters (Σ, Π, Δ, etc.). For example, the H state used for the eEDM measurement is labeled $H^3\Delta_1$. This fully describes the energy contribution from the electronic and spin–orbit parts. For the list of electronic plus spin–orbit states used in ACME, see figure 2.1. The vibrational energy levels are labeled by $\nu = 0, 1, 2, \ldots$, starting from the lowest-energy level. For almost all of the ACME experiment, the molecules occupy the vibrational ground state $\nu = 0$, which we omit by default. The rotational energy levels are labeled by J, M, Ω as described above. Importantly, for a given electronic plus spin–orbit state such as $H^3\Delta_1$, the angular momentum projections Σ, Λ, Ω can only take two values of equal magnitude (e.g. $\Omega = \pm 1$), since different magnitudes correspond to different spin–orbit states. This leads to the Ω-doublet structure in ThO, which is crucial for the eEDM measurement, as explained later.

Ignoring the vibrational wavefunction, the full molecular wavefunction for a single electronic spin–orbit state can be written as $\eta^{2S+1}\Lambda_{|\Omega|}\,|J, M, \Omega\rangle$. For simplicity and for the purposes of ACME, we focus on the science-state H, $J = 1$ manifold, where the effective

Hamiltonian is given by

$$H_{\eta=H,J=1} = H_{\text{Stark}} + H_{\text{Zeeman}} + H_{\Omega\text{-doublet}} + H_{\text{eEDM}} \tag{2.4}$$

$$= -\vec{d}\cdot\vec{\mathcal{E}} - \vec{\mu}\cdot\vec{\mathcal{B}} + \frac{\Delta_\Omega}{2}\left|J=1,M,+\Omega\right\rangle\left\langle J=1,M,-\Omega\right| - d_e\mathcal{E}_{\text{eff}}\Omega, \tag{2.5}$$

where $\vec{d}$ is the molecular electric dipole moment operator, $\vec{\mathcal{E}}$ is the external laboratory electric field, $\vec{\mu}$ is the molecular magnetic dipole moment operator, $\vec{\mathcal{B}}$ is the external laboratory magnetic field, Δ_Ω is the Ω-doublet splitting, d_e is the electron electric dipole moment, and $\mathcal{E}_{\text{eff}}$ is the effective internal electric field in the H state.

The energy shifts corresponding to each term in Eq. 2.5 are explained pictorially in figure 2.2. There are a total of six states in this manifold with $|J=1,M,\Omega=\pm1\rangle$ and $M=-1,0,1$. For now, we assume that there are no external fields other than a static electric field $\vec{\mathcal{E}}=\mathcal{E}_Z\hat{Z}$ and a static magnetic field $\vec{\mathcal{B}}=\mathcal{B}_Z\hat{Z}$ applied along the lab Z axis. The Stark and Zeeman effect terms can be evaluated as diagonal matrix elements in this basis:

$$\left\langle J=1,M,\Omega\right|\left(-\vec{d}\cdot\vec{\mathcal{E}}\right)\left|J=1,M',\Omega'\right\rangle = -d_{H,J=1}\mathcal{E}_Z\Omega M\delta_{MM'}\delta_{\Omega\Omega'}, \tag{2.6}$$

$$\left\langle J=1,M,\Omega\right|\left(-\vec{\mu}\cdot\vec{\mathcal{B}}\right)\left|J=1,M',\Omega'\right\rangle = -\mu_{H,J=1}\mathcal{B}_Z M\delta_{MM'}\delta_{\Omega\Omega'}. \tag{2.7}$$

Here, $d_{H,J=1}$ is the magnitude of the molecular electric dipole moment in the $H, J=1$ state, and $\mu_{H,J=1}$ is the magnitude of the molecular magnetic dipole moment in the $H, J=1$ state. Generally, these dipole moments scale with J as $d_{H,J}=\frac{d_H}{J(J+1)}$ and $\mu_{H,J}=\frac{\mu_H}{J(J+1)}$, where d_H and μ_H are the J-independent dipole moments of the H state. Since the H state is very weakly magnetic, this J scaling does not apply to $\mu_{H,J}$ due to additional perturbations. Furthermore, in a large electric field there is an additional correction to the magnetic moment that depends on the orientation of the molecule relative to the electric-field direction.

The Stark effect term can be explained classically as the interaction between the per-

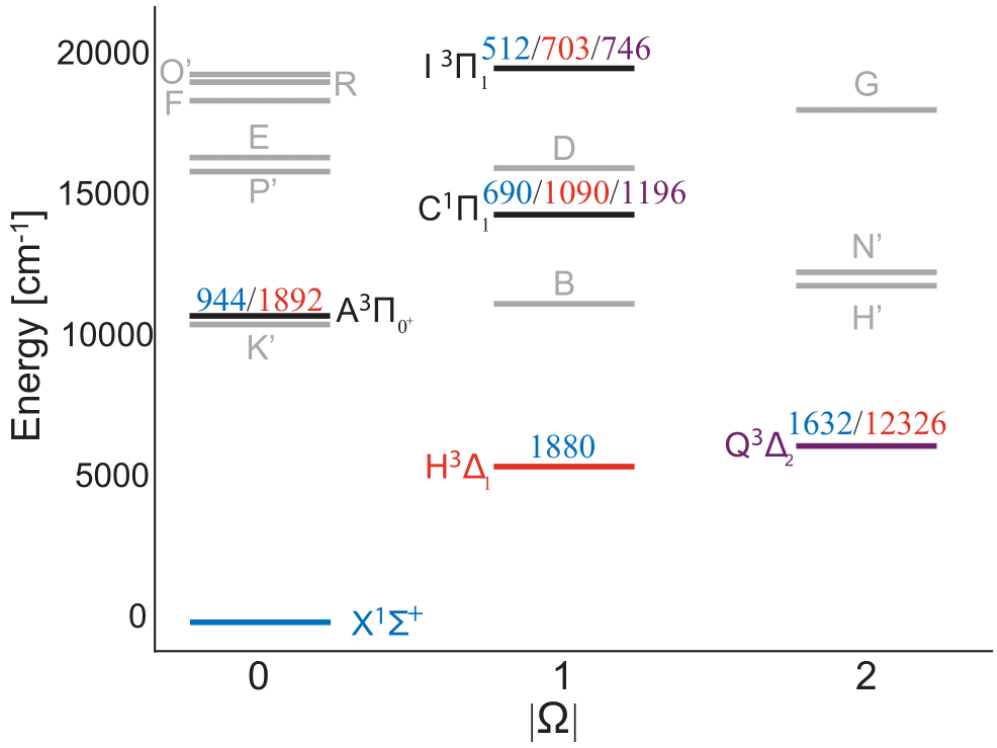

Figure 2.1: Energy level diagrams of ThO. Grey lines indicate levels considered to be not useful for ACME III. The lowest three levels X, H, and Q are colored blue, red, and purple, respectively. Transition wavelengths are given in the corresponding colors in nm [94].

manent molecular electric dipole moment and the external electric field. We define the orientation of the molecular dipole moment relative to the lab electric-field direction as $\mathcal{N} = \hat{\mathcal{E}}\Omega M = \pm 1$, where $\hat{\mathcal{E}} \equiv \mathrm{sgn}(\vec{\mathcal{E}} \cdot \hat{Z})$ denotes the orientation of the external electric field along the lab Z axis[2]. The Stark effect term in Eq. 2.6 now simplifies to

$$H_{\mathrm{Stark}} = -d_{H,J=1}|\mathcal{E}|\mathcal{N}. \tag{2.8}$$

In the absence of the Stark effect, the eEDM Hamiltonian term is suppressed to second order. This can be seen by examining the good quantum numbers. For the eEDM Hamiltonian to contribute without suppression, Ω must be a good quantum number. However, the

2. This is denoted as D in the JILA EDM experiment, which avoids possible confusion with the nuclear angular momentum N in the Hund's case (b) scheme. We will not use the Hund's case (b) symbol N in this work.

presence of the Ω-doublet interaction mixes opposite-Ω states at the energy scale Δ_Ω. This mixing in turn suppresses the eEDM interaction by a factor $\frac{H_{\text{eEDM}}}{H_{\Omega\text{-doublet}}} \lesssim \frac{100\ \mu\text{Hz}}{360\ \text{kHz}}$ in ThO. To prevent this suppression, an external interaction is required such that Ω becomes a good quantum number. The Stark interaction satisfies this criterion, as the good quantum numbers are $\mathcal{N}, M$, or equivalently Ω, M. For this reason, we measure the eEDM while applying an electric field capable of generating an energy shift ~ 100 MHz $\gg H_{\Omega\text{-doublet}} = 360$ kHz. More details about the Ω-doublet structure are given in Sec. 2.1.6.

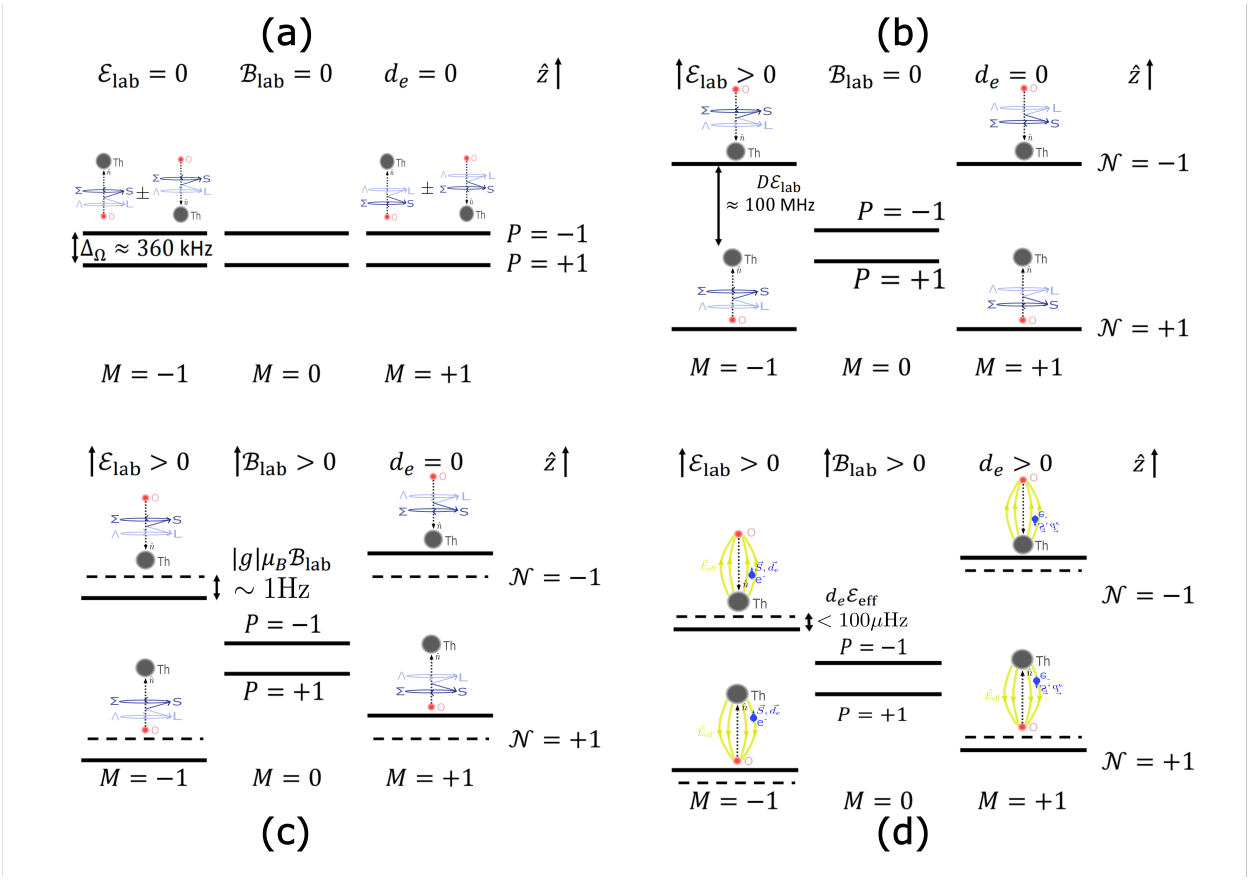

Figure 2.2: Classical picture of energy shifts in the science state H under external fields. (a) Zero-field energy levels with Ω-doublet structure. (b) The Stark effect polarizes the molecules, and the resulting $|M| = 1$ energies depend on the molecular orientation $\mathcal{N}$. (c) The Zeeman effect splits the $M = \pm 1$ levels. (d) The eEDM interaction shifts the levels depending on the relative orientation between the electron spin and the molecular dipole moment. Figure modified from Zack Lasner's thesis [91].

It remains to describe the effective Hamiltonian due to interactions with external AC

or DC fields. We only showed the diagonal elements in the case of fields applied along the lab Z direction earlier. In general, the interaction Hamiltonian can be written as

$$H_I = -\vec{d} \cdot \vec{\mathcal{E}} - \vec{\mu} \cdot \vec{\mathcal{B}}, \tag{2.9}$$

which looks identical to the Stark and Zeeman terms above. However, we are now interested in the off-diagonal matrix elements between different states or for fields applied along different directions. In Hund's case (a), the E1 transition dipole matrix element is given by

$$\langle \eta, J, M, \Omega; \Lambda, S, \Sigma \mid \vec{d} \cdot \vec{\mathcal{E}} \mid \eta', J', M', \Omega'; \Lambda', S', \Sigma' \rangle \tag{2.10}$$

$$= (-1)^{M'-\Omega} T^1_{M'-M}(\mathcal{E})\, \delta_{SS'}\, \delta_{\Sigma\Sigma'}\, \langle \eta, \Lambda | T^1_{\Omega-\Omega'}(d) | \eta', \Lambda' \rangle$$
$$\times \sqrt{(2J+1)(2J'+1)} \begin{pmatrix} J & 1 & J' \\ -\Omega & \Omega-\Omega' & \Omega' \end{pmatrix} \begin{pmatrix} J & 1 & J' \\ -M & M-M' & M' \end{pmatrix}, \tag{2.11}$$

where $T^1_k(\mathcal{E})$ and $T^1_k(d)$ are the k^{th} spherical tensor components of rank 1 for the electric field and electric dipole operators, respectively. The evaluation of the reduced matrix elements $\langle \eta, \Lambda | T^1_k(d) | \eta', \Lambda' \rangle$ relies on experimental measurements, and the measured values are summarized in Table 2.1. The physical meaning of these reduced matrix elements is that they quantify the molecular-frame E1 moment between electronic states. By convention, $d_{\eta,\parallel} \equiv \langle \eta, \Lambda | T^1_0(d) | \eta', \Lambda' \rangle$ is called the parallel transition dipole moment, while $d_{\eta,\perp} \equiv \langle \eta, \Lambda | T^1_{\pm 1}(d) | \eta', \Lambda' \rangle$ is called the perpendicular transition dipole moment. For a specific pair of transitions $\eta, \Omega \mapsto \eta', \Omega'$, it is clear from Eq. 2.10 that only $d_{\eta,\parallel}$ or $d_{\eta,\perp}$ contributes. Thus, there is no ambiguity in defining a table of reduced E1 matrix elements, as shown in Table 2.1.

	X	H	Q	A	C	E	I
X	1.09(1)			0.15(4)	0.51(2)		0.72
H		1.67(4)		0.31(9)	0.02	0.018(11)	0.27
Q			1.60		0.40		0.23
A	0.15(4)	0.31(9)					
C	0.51(2)	0.02	0.40		1.03		
E		0.018(11)				1.39(1)	
I	0.72	0.27	0.23				1.67(1)

Table 2.1: Reduced E1 matrix elements in units of ea_0 (1.28 MHz/(V/cm), 2.54 D). Values are based on previous measurements [52, 86, 87, 90, 91, 94–99]. For each transition, only one of $d_{\eta,\parallel}$ or $d_{\eta,\perp}$ is non-zero.

A similar expression exists for the M1 transition dipole matrix element:

$$\langle \eta, J, M, \Omega; \Lambda, S, \Sigma \mid \vec{\mu} \cdot \vec{\mathcal{B}} \mid \eta', J', M', \Omega'; \Lambda', S', \Sigma' \rangle \tag{2.12}$$

$$= (-1)^{M'-\Omega} T^1_{M'-M}(\mathcal{B}) \langle \eta, \Omega | T^1_{\Omega-\Omega'}(\mu) | \eta', \Omega' \rangle \sqrt{(2J+1)(2J'+1)} \tag{2.13}$$

$$\times \begin{pmatrix} J & 1 & J' \\ -\Omega & \Omega-\Omega' & \Omega' \end{pmatrix} \begin{pmatrix} J & 1 & J' \\ -M & M-M' & M' \end{pmatrix}, \tag{2.14}$$

where $T^1_k(\mathcal{B})$ and $T^1_k(\mu)$ are the k^{th} spherical tensor components of rank 1 for the magnetic field and magnetic dipole operators, respectively. The key difference between the E1 and M1 transition matrix elements is that we do not treat the spin as an independent quantum number in the M1 case, because the magnetic field couples directly to the spin; this explains the absence of the factors $\delta_{SS'}\delta_{\Sigma\Sigma'}$, and the spin dependence is implicitly absorbed into the reduced matrix element $\langle \eta, \Omega | T^1_k(\mu) | \eta', \Omega' \rangle$. By convention, we define $\mu_{\eta,\parallel} \equiv \langle \eta, \Omega | T^1_0(\mu) | \eta', \Omega' \rangle$ as the parallel magnetic transition dipole moment and $\mu_{\eta,\perp} \equiv \langle \eta, \Omega | T^1_{\pm 1}(\mu) | \eta', \Omega' \rangle$ as the perpendicular magnetic transition dipole moment. Again, although we have written the matrix element assuming both states are in Hund's case (a), the result is compatible with the upper state being in Hund's case (c). The measured and estimated values of these reduced M1 matrix elements are summarized in Table 2.2. Unlike the reduced E1 matrix element

$\langle \eta, \Lambda | T^1_k(d) | \eta', \Lambda' \rangle$, the reduced M1 matrix element $\langle \eta, \Omega | T^1_k(\mu) | \eta', \Omega' \rangle$ depends on the signs of Ω and Ω'. This is intuitive because the magnetic dipole moment depends on both the orbital and spin angular momenta, which have different signs for different Ω projections. This implies that the (non-reduced) E1 and M1 transition dipole moments differ by a sign of Ω or Ω', depending on the specific transition.

	X	H	Q	C	I
X	< 0.001			0.40	
H		-0.00784(4)		-0.16	0.08(5)
Q			2.07(11)		
C	0.40	-0.16		1.23(6)	
I		0.08(5)			0.5260(1)

Table 2.2: Directly measured molecular-frame diagonal matrix elements $\mu_{\eta,\parallel}$ in units of μ_B [52, 86, 87, 90, 91, 94–99]. Note that this does not include the rotational factor, which scales as $\mu_{\eta,J} \equiv \frac{\mu_{\eta,\parallel}\Omega}{J(J+1)}$. The off-diagonal matrix elements for X–C and C–H are based on theoretical calculations and scaling from measured E1 matrix elements [100]. The H–I matrix element is estimated by Brendon O'Leary in his thesis [87]. The H-state molecular-frame dipole moment was measured to be $-0.0088(1)$ in previous generations of ACME [101]. The new measurement, which we believe is more accurate, will be published in an upcoming paper that also describes the ACME III magnetic shield and field coil system.

State	Lifetime	Ω-doublet splitting
H, $J = 1$	4.2(5) ms	360 kHz
Q, $J = 2$	> 62 ms	< 10 kHz
C, $J = 1$	468(30) ns	51 MHz
I, $J = 1$	115(4) ns	91 MHz

Table 2.3: Measured lifetimes and Ω-doublet splittings [52, 87, 90, 95–97].

2.1.2 Born–Oppenheimer approximation

We begin the introduction of the Born–Oppenheimer approximation by writing out the spinless full Hamiltonian for a diatomic molecule and separating electronic and nuclear

motions.

$$H_{\text{mol}} = H_{\text{elc}} + H_{\text{nuc}}, \tag{2.15}$$

$$\begin{aligned} H_{\text{elc}} = &-\frac{\hbar^2}{2m}\sum_i \nabla_i^2 - \frac{\hbar^2}{2(M_1+M_2)}\sum_{i,j}\nabla_i\cdot\nabla_j \\ &+\sum_{i<j}\frac{e^2}{4\pi\epsilon_0 R_{ij}} - \sum_{\alpha,i}\frac{Z_\alpha e^2}{4\pi\epsilon_0 R_{i\alpha}}, \end{aligned} \tag{2.16}$$

$$H_{\text{nuc}} = -\frac{\hbar^2}{2\mu}\nabla_R^2 + \frac{Z_1 Z_2 e^2}{4\pi\epsilon_0 R}, \tag{2.17}$$

where m is the electron mass, M_1, M_2 are the nuclear masses, $\mu = \frac{M_1 M_2}{M_1+M_2}$ is the reduced mass of the nuclei, R_{ij} is the distance between electrons i and j, $R_{i\alpha}$ is the distance between electron i and nucleus α, and R is the internuclear distance. Treating the nuclear positions as fixed parameters, we can write the solutions of the electronic Hamiltonian as

$$H_{\text{elc}}\psi_{\text{e}}^{\eta}(\vec{r}_i; R) = E_{\text{e}}^{\eta}(R)\,\psi_{\text{e}}^{\eta}(\vec{r}_i; R), \tag{2.18}$$

where η is the electronic state label and $\vec{r}_i$ are the electronic coordinates in the molecular frame. We can do the same for the nuclear Hamiltonian:

$$H_{\text{nuc}}\psi_{\text{nuc}}^{\eta}(R,\theta,\phi) = E_{\text{nuc}}^{\eta}\,\psi_{\text{nuc}}^{\eta}(R,\theta,\phi), \tag{2.19}$$

where θ, ϕ are the angular coordinates of the internuclear axis in the lab frame. Note that the nuclear wavefunction depends on the electronic state η through the electronic potential energy surface $E_{\text{e}}^{\eta}(R)$, which acts as a potential for the nuclear motion, as expected since the electron density is the source of binding in molecules.

The basis set of the Born–Oppenheimer approximation is then given by product states

of the electronic and nuclear wavefunctions:

$$\Psi_{\text{total}} = \psi_{\text{e}}^{\eta}(\vec{r}_i; R)\, \psi_{\text{nuc}}^{\eta}(R, \theta, \phi). \tag{2.20}$$

This is not the exact solution of the full molecular Hamiltonian. However, it turns out to be a very good approximation because of the large mass difference between the electrons and the nuclei. To examine the quality of this approximation, we consider the off-diagonal matrix elements of the full Hamiltonian that couple different Born–Oppenheimer basis states:

$$\begin{aligned} C_{\eta\eta'} &\equiv \langle \psi_{\text{e}}^{\eta} \psi_{\text{nuc}}^{\eta} | H_{\text{elc}} + H_{\text{nuc}} | \psi_{\text{e}}^{\eta'} \psi_{\text{nuc}}^{\eta'} \rangle & (2.21)\\ &= \langle \psi_{\text{e}}^{\eta} \psi_{\text{nuc}}^{\eta} | H_{\text{nuc}} | \psi_{\text{e}}^{\eta'} \psi_{\text{nuc}}^{\eta'} \rangle & (2.22)\\ &= -\frac{\hbar^2}{2\mu} \langle \psi_{\text{nuc}}^{\eta} | \left(Q_{\eta,\eta'} + P_{\eta,\eta'} \cdot \nabla_R \right) | \psi_{\text{nuc}}^{\eta'} \rangle \,, & (2.23)\\ P_{\eta,\eta'} &\equiv 2 \langle \psi_{\text{e}}^{\eta} | \nabla_R | \psi_{\text{e}}^{\eta'} \rangle, & (2.24)\\ Q_{\eta,\eta'} &\equiv \langle \psi_{\text{e}}^{\eta} | \nabla_R^2 | \psi_{\text{e}}^{\eta'} \rangle, & (2.25) \end{aligned}$$

where we assume $\eta \neq \eta'$, so the cross terms of the electronic Hamiltonian vanish because it does not act on the nuclear wavefunction. The Born–Oppenheimer approximation is equivalent to setting $C_{\eta\eta'} = 0$ for all $\eta \neq \eta'$. Physically, these coefficients describe the change in the electronic wavefunction as the nuclear separation changes, and therefore quantify the inability of the electrons to instantaneously redistribute upon changes in nuclear positions. Setting them to zero corresponds to assuming that electrons instantaneously relax into the new energy minimum for any nuclear configuration. Using a simple back-of-the-envelope estimate, we can show that this approximation is very good for heavy molecules such as ThO. Using atomic units, the off-diagonal coupling terms can be estimated as

$$C_{\eta\eta'} \sim \frac{\hbar^2}{2\mu (2a_0)^2} \sim 10 \text{ GHz}, \tag{2.26}$$

where we use $\mu = 16$ amu for ThO and $2a_0$ as the typical bond length. This coupling is negligible compared to the typical electronic and spin–orbit energy scales, which are $\gtrsim$ 10 THz. This justifies the choice of the Born–Oppenheimer basis set in Eq. 2.20. With this basis, the Hamiltonian problem simplifies to

$$H_{\mathrm{mol}}\Psi_{\mathrm{total}} = (H_{\mathrm{elc}} + H_{\mathrm{nuc}})\psi_{\mathrm{e}}^{\eta}(\vec{r}_i; R)\,\psi_{\mathrm{nuc}}^{\eta}(R,\theta,\phi) \tag{2.27}$$

$$\approx \psi_{\mathrm{e}}^{\eta}(\vec{r}_i; R)\left[E_{\mathrm{e}}^{\eta}(R) + H_{\mathrm{nuc}}\right]\psi_{\mathrm{nuc}}^{\eta}(R,\theta,\phi), \tag{2.28}$$

i.e., by treating the electronic part as its own degree of freedom labeled by the quantum number η, the quantum many-body problem is reduced to a two-body problem involving only the two nuclei moving in the lab frame, with a binding potential $E_{\mathrm{e}}^{\eta}(R)$ provided by the electrons, plus the Coulomb repulsion between the nuclei contained in H_{nuc}.

2.1.3 Spin–orbit

To motivate the spin–orbit coupling in diatomic molecules, we begin by considering light diatomic molecules where the spin–orbit coupling is negligible compared to the rotational energy scale. The total molecular Hamiltonian is symmetric with respect to rotations of the molecular-frame coordinates about the internuclear axis. This leads to the conservation of the projection of the orbital angular momentum onto the internuclear axis, which we denote by Λ. Thus, Λ is a good quantum number. There are three energy scales of interest that determine which Hund's case coupling scheme is appropriate.

$$\langle H_{\mathrm{tot}}\rangle = \langle H_{\mathrm{elc}}\rangle + \langle H_{\mathrm{SO}}\rangle + \langle H_{\mathrm{rot}}\rangle \tag{2.29}$$

$$\approx \frac{\hbar^2}{2m_e(2a_0)^2} + \langle A\,\vec{L}\cdot\vec{S}\rangle + \langle B(\vec{J}-\vec{J}_e)^2\rangle\,, \tag{2.30}$$

where by convention A and B denote the spin–orbit and rotational constants, respectively, and $\vec{J}_e = \vec{L} + \vec{S}$ is the total electronic angular momentum. Because of the difference between the electron and nuclear masses, the electronic energy is generally greater than the rotational constant by a factor of order $\sim \frac{\mu}{m_e} \sim \frac{16 \text{ amu}}{m_e}$ where μ is the reduced mass of the ThO nuclei. The only remaining question is the size of the spin–orbit constant A. For light molecules, $A \ll B$, and terms such as $B\,\vec{J}\cdot\vec{S}$ prevent Σ from being a good quantum number even though S remains a good quantum number. This term is therefore called the spin-uncoupling term.

If instead $H_{\text{elc}} \gg A \gg B$, then Λ remains a good quantum number because the electronic energy dominates. The spin–orbit term can be written in terms of raising and lowering operators as

$$H_{\text{SO}} = A\,\vec{L}\cdot\vec{S} = A\left[L_0 S_0 + \frac{1}{2}(L_+ S_- + L_- S_+)\right]. \tag{2.31}$$

The $L_\pm$ terms change Λ and are therefore suppressed by the large electronic energy scale. Thus, only the $L_0 S_0$ term survives, and Σ remains a good quantum number, with the 1^{st} order spin–orbit coupling energy given by

$$\langle H_{\text{SO}} \rangle = A\,\Lambda\,\Sigma. \tag{2.32}$$

In this limit, both the orbital and spin angular momenta are said to be pinned to the internuclear axis, and this coupling scheme is called Hund's case (a). The classical picture of Hund's case (a) is shown in figure 2.3.

In the limit where $A \gtrsim \Delta\langle H_{\text{elc}}\rangle$, where $\Delta\langle H_{\text{elc}}\rangle$ is the energy spacing with the closest electronic state, the off-diagonal terms in Eq. 2.31 imply that neither Λ nor Σ are good quantum numbers. However, the projection of the total electronic angular momentum onto the internuclear axis, $\Omega = \Lambda + \Sigma$, remains a good quantum number because of the rotational symmetry about the internuclear axis. This coupling scheme is called Hund's case (c). Some of the excited, non-metastable states in ThO used in ACME are better described by Hund's

case (c). This can be seen from the density of states in figure 2.1, where $\Delta \langle H_{\mathrm{elc}} \rangle$ is expected to be much smaller in the $> 10\,000$ cm^{-1} region than in the $< 10\,000$ cm^{-1} region where the ground and metastable states reside. For simplicity, we describe all molecular states in the Hund's case (a) basis throughout this work.

For standardization of notation, we label the simultaneous energy eigenstates of the electronic and spin–orbit interaction by η, Λ, S, Σ as

$$\psi_e^{\eta}(\vec{r}_i; R) \equiv |\eta, \Lambda; S, \Sigma\rangle \,. \tag{2.33}$$

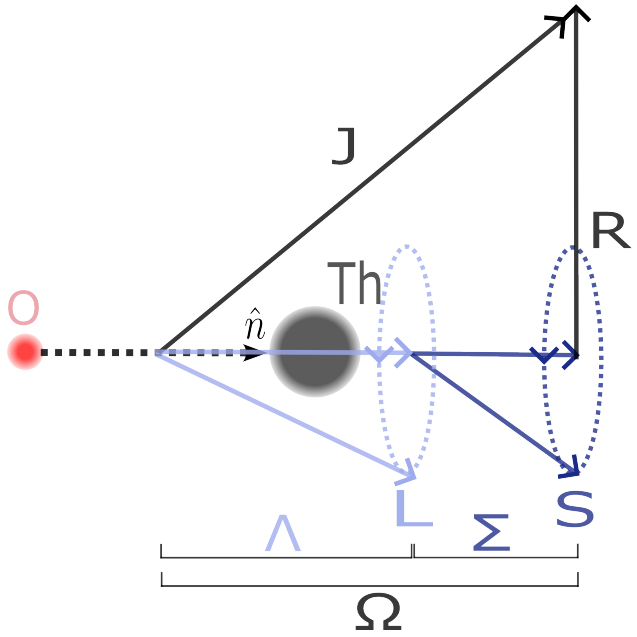

Figure 2.3: Hund's coupling scheme (a). **R** denotes the nuclear angular momentum.

2.1.4 Vibrational motion

The Born–Oppenheimer approximation greatly reduces the degrees of freedom in the molecular Hamiltonian by treating the electronic state as a separate degree of freedom. As shown in Eq. 2.28, the remaining problem is a two-body problem of the two nuclei moving in the lab frame with a binding potential provided by the electrons. The nuclear Hamiltonian

can be further separated into vibrational and rotational parts:

$$H_{\mathrm{ThO}} = H_{\mathrm{nuclear}} + H^{\eta}_{\mathrm{binding}} \tag{2.34}$$

$$= -\frac{\hbar^2}{2\mu}\nabla_R^2 + \frac{Z_{\mathrm{Th}} Z_{\mathrm{O}} e^2}{4\pi\epsilon_0 R} + E_{\mathrm{e}}^{\eta}(R) \tag{2.35}$$

$$= \left[-\frac{\hbar^2}{2\mu}\frac{\partial^2}{\partial R^2} + V_{\mathrm{eff}}^{\eta}(R)\right] + \frac{\mathbf{R}^2}{2\mu R^2}, \tag{2.36}$$

$$V_{\mathrm{eff}}^{\eta}(R) \equiv \frac{Z_{\mathrm{Th}} Z_{\mathrm{O}} e^2}{4\pi\epsilon_0 R} + E_{\mathrm{e}}^{\eta}(R), \tag{2.37}$$

where $\mathbf{R}$ is the nuclear angular momentum operator (not to be confused with the internuclear separation R). The effective potential $V_{\mathrm{eff}}^{\eta}(R)$ contains an attractive part from the electronic binding energy and a repulsive part from the nuclear Coulomb repulsion, resulting in an energy minimum at the equilibrium bond length $R = R_e$. Each electronic state plus spin–orbit state η has its own effective potential curve $V_{\mathrm{eff}}^{\eta}(R)$. The terms in brackets correspond to vibrational motion along the internuclear axis, while the last term corresponds to rotational motion of the nuclear axis, which is discussed in the next subsection.

In general, the effective potential curve $V_{\mathrm{eff}}^{\eta}(R)$ is anharmonic. A good approximation for its functional form is given by the Morse potential

$$V_{\mathrm{eff}}^{\eta}(R) \approx D_e \left(1 - e^{-\beta(R-R_e)}\right)^2, \tag{2.38}$$

where D_e is the dissociation energy, R_e is the equilibrium bond length, and β is a parameter related to the curvature of the potential at the minimum. The vibrational wavefunctions can be solved analytically together with the rotational wavefunctions. The energy spectrum contains contributions from both anharmonic vibrations and rotational motion, as well as their couplings. For simplicity, we do not discuss these effects here, since in ACME we mostly work with the lowest vibrational state $v = 0$. More details about this spectrum can be found in Brown and Carrington [50]. Near the minimum, this potential is well approximated as

harmonic, with an approximately factorizable vibrational wavefunction

$$\psi_{\text{vib}}(R) \equiv |\nu\rangle\,. \tag{2.39}$$

2.1.5 Rotational motion and branching ratios

After solving the vibrational motion, we now only need to deal with the rotational degree of freedom in Eq. 2.36. The rotational plus spin-orbit Hamiltonian is given by

$$H_{\text{rot + spin-orbit}} = \frac{\mathbf{R}^2}{2\mu R^2} + A\Lambda\Sigma \tag{2.40}$$

$$= B\mathbf{R}^2 + A\Lambda\Sigma \tag{2.41}$$

$$= B(\mathbf{J} - \mathbf{J}_e)^2 + A\Lambda\Sigma, \tag{2.42}$$

where $B = \frac{\hbar^2}{2\mu R_e^2}$ is the rotational constant, and $A\Lambda\Sigma$ is the spin-orbit contribution in Hund's case (a). This is the Hamiltonian of a rigid rotor plus an additional contribution from spin-orbit coupling. The angular momentum about the internuclear axis is conserved, and the total electronic angular momentum projection onto the internuclear axis $\mathbf{J}_e \cdot \hat{n} = \Omega = \Lambda + \Sigma$ is a good quantum number. The good quantum numbers are thus J, M, Ω. The rotational wavefunctions are given by the Wigner D-matrices

$$\langle \phi, \theta, \psi | J, M, \Omega \rangle = \sqrt{\frac{2J+1}{8\pi^2}} D^{(J)*}_{M,\Omega}(\phi, \theta, \psi) \tag{2.43}$$

$$= \sqrt{\frac{2J+1}{8\pi^2}} e^{-iM\phi} d^{(J)}_{M,\Omega}(\theta) e^{-i\Omega\psi}, \tag{2.44}$$

where ϕ, θ are the Euler angles relating the molecular axis frame to the lab frame, and ψ is the internal rotation angle about the internuclear axis.

The rotational energy without the spin-orbit contribution is given by

$$H_{\text{rot}} \left|J, M, \Omega\right\rangle \approx B \left[J(J+1) - \Omega^2\right] \left|J, M, \Omega\right\rangle, \tag{2.45}$$

where we have neglected the off-diagonal coupling term $-2B\mathbf{J} \cdot \mathbf{J}_e$ for now. This term is actually very important in generating the so-called Ω-doublet structure in heavy molecules like ThO. We will discuss this in more detail in the next subsection.

We are often interested in the branching ratios due to spontaneous emission from an excited electronic state $\eta', \Lambda', S', \Sigma'$ with vibrational state ν' and rotational state J', M', Ω'. Let the lower state be η, Λ, S, Σ with vibrational state ν and rotational state J, M, Ω. For E1 spontaneous emission, we need to compute the E1 transition dipole matrix element between these two states. We can organize this calculation by separating the electronic, vibrational, and rotational parts of the total wavefunction.

$$\left|\left\langle \eta, \Lambda; S, \Sigma\right| \left\langle \nu\right| \left\langle J, M, \Omega\right|\vec{d} \cdot \vec{\mathcal{E}}\left|J', M', \Omega'\right\rangle \left|\nu'\right\rangle \left|\eta', \Lambda'; S', \Sigma'\right\rangle\right|^2 \tag{2.46}$$

$$\propto \left|\left\langle \nu|\nu'\right\rangle\right|^2 |d|^2 \, \Xi(J, M, \Omega; J', M', \Omega') \, \delta_{SS'} \delta_{\Sigma\Sigma'}, \tag{2.47}$$

where we have assumed that the electronic transition dipole moment is approximately constant as a function of R, so we can factor out the vibrational overlap $\langle\nu|\nu'\rangle$. The squared modulus of this overlap is the Franck–Condon factor and quantifies the probability of an electronic transition between states with arbitrary vibrational quantum numbers.. The rotational factor $\Xi(J, M, \Omega; J', M', \Omega')$ captures the angular-momentum dependence and is related to the Hönl–London factor, which we derive next. The term $|d|$ is the reduced transition dipole matrix element in the molecular frame characterizing the strength of the transition.

To derive the Hönl–London factor, we first note the difficulty in evaluating this matrix element: the electric field $\vec{\mathcal{E}}$ is defined in the lab frame, while the electronic wavefunctions

such as $|\eta, \Lambda; S, \Sigma\rangle$ are defined in the molecular frame. To carry out the dot product, we rotate the molecular dipole operator from the molecular frame to the lab frame:

$$\langle \eta, \Lambda; S, \Sigma| \langle J, M, \Omega | \vec{d} \cdot \vec{\mathcal{E}} | J', M', \Omega' \rangle |\eta', \Lambda'; S', \Sigma'\rangle \tag{2.48}$$

$$= \sum_{p,q} (-1)^p \langle \eta, \Lambda; S, \Sigma| \langle J, M, \Omega | T^1_p(\vec{\mathcal{E}}) \mathcal{D}^{(1)*}_{-p,q} T^1_q(\vec{d}) | J', M', \Omega' \rangle |\eta', \Lambda'; S', \Sigma'\rangle \,, \tag{2.49}$$

where $\mathcal{D}$ is the Wigner D-matrix operator rotating from the molecular frame to the lab frame so that the dot product can be evaluated in the lab frame. The Wigner D-matrix acts only on the nuclear degrees of freedom, so we can factor out the electronic part of the matrix element.

$$\begin{aligned}(2.48) =& \sum_{p,q} (-1)^p \langle J, M, \Omega | T^1_p(\vec{\mathcal{E}}) \mathcal{D}^{(1)*}_{-p,q} | J', M', \Omega' \rangle \\ & \times \langle \eta, \Lambda; S, \Sigma| \, T^1_q(\vec{d}) \, |\eta', \Lambda'; S', \Sigma'\rangle \qquad (2.50) \\ =& \sum_{p,q} (-1)^p T^1_p(\vec{\mathcal{E}}) (-1)^{M-\Omega} \sqrt{(2J+1)(2J'+1)} \\ & \times \begin{pmatrix} J & 1 & J' \\ -M & p & M' \end{pmatrix} \begin{pmatrix} J & 1 & J' \\ -\Omega & -q & \Omega' \end{pmatrix} \delta_{SS'} \delta_{\Sigma\Sigma'} \langle \eta, \Lambda | T^1_q(\vec{d}) | \eta', \Lambda' \rangle \qquad (2.51) \\ =& (-1)^{M'-\Omega} T^1_{M'-M}(\vec{\mathcal{E}}) \sqrt{(2J+1)(2J'+1)} \\ & \times \begin{pmatrix} J & 1 & J' \\ -M & M-M' & M' \end{pmatrix} \begin{pmatrix} J & 1 & J' \\ -\Omega & \Omega-\Omega' & \Omega' \end{pmatrix} \delta_{SS'} \delta_{\Sigma\Sigma'} \langle \eta, \Lambda | T^1_{\Omega-\Omega'}(\vec{d}) | \eta', \Lambda' \rangle \,. \qquad (2.52)\end{aligned}$$

In the first step we simply factored out the electronic part. In the second step we used the standard formula for matrix elements of Wigner D-matrices [50] and the fact that the electric dipole operator does not act on spin degrees of freedom. In the last step we performed the sum over p, q using the properties of the $3j$ symbols. This result is identical to the E1

matrix element written earlier in eq. 2.10. The squared modulus of this transition matrix element has the angular-momentum dependence

$$\Xi(J, M, \Omega; J', M', \Omega') = (2J+1)(2J'+1) \begin{pmatrix} J & 1 & J' \\ -M & M - M' & M' \end{pmatrix}^2 \begin{pmatrix} J & 1 & J' \\ -\Omega & \Omega - \Omega' & \Omega' \end{pmatrix}^2 . \tag{2.53}$$

This rotational dependence, together with the Franck–Condon factor and the reduced dipole matrix element, completely determines the transition strength and thus the branching ratios to different quantum states. Sometimes we are interested in the total transition strength between two rotational levels J, Ω and J', Ω' summed over all initial and final M, M' states. This motivates the definition of the Hönl–London factor.

The Hönl–London factor $S(J, \Omega; J', \Omega')$ is defined as the rotational factor Ξ summed over all initial and final M, M' states. Using the properties of $3j$ symbols, we can evaluate this sum to obtain

$$S(J, \Omega; J', \Omega') \equiv \sum_{M, M'} \Xi(J, M, \Omega; J', M', \Omega') \tag{2.54}$$

$$= (2J'+1) \begin{pmatrix} J & 1 & J' \\ -\Omega & \Omega - \Omega' & \Omega' \end{pmatrix}^2 . \tag{2.55}$$

This Hönl–London factor quantifies the relative strength of rotational transitions between different J, Ω states for a given electronic and vibrational transition. In molecular spectroscopy, it is common to label rotational transitions using the branch notation $P(J)$, $Q(J)$, and $R(J)$ corresponding to $J' - J = -1, 0, 1$, respectively[3]. For example, the $R(0)$ line corresponds to a transition connecting $J = 0$ to $J' = 1$, while the $Q(1)$ line corresponds to

3. Note that J' is the upper state and J is the lower state, so a given transition has a fixed label regardless of direction.

a transition connecting $J = 1$ to $J' = 1$. Their relative strength can be computed using the ratios of Hönl–London factors $\frac{S(0,\Omega;1,\Omega')}{S(1,\Omega;1,\Omega')}$.

2.1.6 Ω-doubling

We now return to the off-diagonal coupling term $-2B\mathbf{J}\cdot\mathbf{J}_e$ in eq. 2.42, which we neglected in the last subsection. This term can be expanded as

$$-2B\mathbf{J}\cdot\mathbf{J}_e = -2B\left[J_0 J_{e0} + \frac{1}{2}(J_+ J_{e-} + J_- J_{e+})\right]. \tag{2.56}$$

The first term can contribute to the diagonal matrix elements and is already included in the rotational energy expression in the last subsection. The next two terms contribute off-diagonally, and couple states with $\Delta\Omega = \pm 1$. Consider, for example, the effect of these off-diagonal terms for a $|\Omega| = 1$ state such as the H or I state in ThO. These terms can couple the H or I state with $\Omega = \pm 1$ to states with $\Omega = 0$. The resulting energy shift can be estimated using second-order perturbation theory as

$$\Delta_\Omega \approx \frac{B^2}{E_{|\Omega|=1} - E_{\Omega=0}}. \tag{2.57}$$

There are two important differences here between the H and I states in ThO. The I,$^1\Pi$ state lies much closer to $\Omega = 0$ states than the H state, resulting in a much smaller denomenator and thus a much larger energy shift. In particular, because $\mathbf{J}\cdot\mathbf{J}_e$ conserves parity, the $|\Omega| = 1$ state can only couple to $\Omega = 0$ states of the same parity. This means only one of the two nominally degenerate parity states of the I state gets shifted via this coupling, while the other remains unshifted. The resulting energy eigenstates of parity in Hund's case (a) are written as

$$|J, M, \pm\rangle = \frac{|J, M, +\Omega\rangle \pm (-1)^{J-S}|J, M, -\Omega\rangle}{\sqrt{2}}. \tag{2.58}$$

These doublets are called parity doublets. Whether the upper or lower parity state corresponds to the symmetric or antisymmetric combination depends on whether the nearby $\Omega = 0$ state lies above or below in energy and on the parity of the $\Omega = 0$ state coupled, which in turn depends on the value of J and S and on the reflection symmetry of the electronic wavefunction about the nuclear axis if $\Lambda = 0$.

Unlike the I state, the H state is very far from any $\Omega = 0$ states, resulting in a large denomenator. In addition, with only 1 unit of angular momentum change, the H,$^3\Delta_1$ can only be coupled to $^3\Pi_0$ states. There are always two such states next to each other with opposite parity, resulting in a nearly equal amount of pushing for the two opposite parity states of the H state. These two factors together lead to a much smaller Ω doublet splitting for the H state. The presence of these closely spaced opposite-parity states is extremely useful for EDM measurements, as they allow large molecular polarization in relatively small applied electric fields. We will elaborate on this further in the next subsection.

The physical origin of Ω-doubling is often described as a Coriolis coupling between the nuclear rotation and the electronic motion. One intuitive picture is discussed in Nick Hutzler's thesis, where different parity states correspond to different effective moments of inertia for the nuclear rotation due to differences in the electronic motion [53]. However, this picture does not reliably produce quantitative estimates of the energy scale for most states, so we do not pursue it further here.

2.1.7 Interaction with external fields

An important aspect of molecular structure for EDM measurements is its interaction with external electric and magnetic fields. We are interested in E1 and M1 matrix elements

of the form

$$\langle \eta, \Lambda; S, \Sigma | \langle J, M, \Omega | \vec{d} \cdot \vec{\mathcal{E}} | J', M', \Omega' \rangle | \eta', \Lambda'; S', \Sigma' \rangle , \tag{2.59}$$

$$\langle \eta, \Lambda; S, \Sigma | \langle J, M, \Omega | \vec{\mu} \cdot \vec{\mathcal{B}} | J', M', \Omega' \rangle | \eta', \Lambda'; S', \Sigma' \rangle . \tag{2.60}$$

We derived the E1 matrix element in subsection 2.1.5. It is given in eq. 2.10 and is copied below for convenience:

$$\begin{aligned} &\langle \eta, \Lambda; S, \Sigma | \langle J, M, \Omega | \vec{d} \cdot \vec{\mathcal{E}} | J', M', \Omega' \rangle | \eta', \Lambda'; S', \Sigma' \rangle \\ =&(-1)^{M'-\Omega} T^1_{M'-M}(\vec{\mathcal{E}}) \sqrt{(2J+1)(2J'+1)} \begin{pmatrix} J & 1 & J' \\ -M & M-M' & M' \end{pmatrix} \\ &\times \begin{pmatrix} J & 1 & J' \\ -\Omega & \Omega-\Omega' & \Omega' \end{pmatrix} \delta_{SS'} \delta_{\Sigma\Sigma'} \langle \eta, \Lambda | T^1_{\Omega-\Omega'}(\vec{d}) | \eta', \Lambda' \rangle . \end{aligned} \tag{2.61}$$

Most of this derivation consists of rotating the dipole operator from the molecular frame to the lab frame using Wigner D-matrices. The M1 matrix element can be derived in an analogous way. The only difference arises in the step where the electronic part of the matrix element is factored out:

$$\text{E1:} \quad \langle \eta, \Lambda; S, \Sigma | \, T^1_q(\vec{d}) \, | \eta', \Lambda'; S', \Sigma' \rangle = \delta_{SS'} \delta_{\Sigma\Sigma'} \langle \eta, \Lambda | T^1_q(\vec{d}) | \eta', \Lambda' \rangle , \tag{2.62}$$

$$\text{M1:} \quad \langle \eta, \Lambda; S, \Sigma | \, T^1_q(\vec{\mu}) \, | \eta', \Lambda'; S', \Sigma' \rangle = \langle \eta, \Omega | T^1_q(\vec{\mu}) | \eta', \Omega' \rangle . \tag{2.63}$$

The M1 matrix element cannot be factorized into spin and orbital parts because the magnetic dipole operator acts on both spin and orbital degrees of freedom. Since the excited states in ThO are predominantly Hund's case (c) states, we express the M1 matrix element in terms

of Ω rather than Λ, S, Σ. The final M1 matrix element is therefore

$$\begin{aligned}
&\langle \eta, \Lambda; S, \Sigma | \langle J, M, \Omega | \vec{\mu} \cdot \vec{\mathcal{B}} | J', M', \Omega' \rangle | \eta', \Lambda'; S', \Sigma' \rangle \\
=&(-1)^{M'-\Omega} T^1_{M'-M}(\vec{\mathcal{B}}) \sqrt{(2J+1)(2J'+1)} \begin{pmatrix} J & 1 & J' \\ -M & M-M' & M' \end{pmatrix} \\
&\times \begin{pmatrix} J & 1 & J' \\ -\Omega & \Omega-\Omega' & \Omega' \end{pmatrix} \langle \eta, \Omega | T^1_{\Omega-\Omega'}(\vec{\mu}) | \eta', \Omega' \rangle \,. \qquad (2.64)
\end{aligned}$$

From this expression, it might appear that the Zeeman and Stark shifts always have the same sign. This is not the case. The diagonal matrix elements follow from the properties of the time-reversal operator $\mathcal{T}$:

$$\begin{aligned}
\langle \eta, \Omega | T^1_0(\vec{\mu}) | \eta, \Omega \rangle &= (\mathcal{T} \langle \eta, \Omega |)(\mathcal{T} T^1_0(\vec{\mu}) \mathcal{T})(\mathcal{T} | \eta, \Omega \rangle) \qquad (2.65) \\
&= - \langle \eta, -\Omega | T^1_0(\vec{\mu}) | \eta, -\Omega \rangle \,, \qquad (2.66) \\
\langle \eta, \Lambda | T^1_0(\vec{d}) | \eta, \Lambda \rangle &= (\mathcal{T} \langle \eta, \Lambda |)(\mathcal{T} T^1_0(\vec{d}) \mathcal{T})(\mathcal{T} | \eta, \Lambda \rangle) \qquad (2.67) \\
&= + \langle \eta, -\Lambda | T^1_0(\vec{d}) | \eta, -\Lambda \rangle \,. \qquad (2.68)
\end{aligned}$$

This makes intuitive sense because the molecular-frame electric dipole does not care about the direction in which the electrons are orbiting, while the magnetic dipole does. Therefore, the magnetic dipole moment changes sign with Ω, while the electric dipole moment does not. This leads to different signs for the Stark shift $\propto \Omega M \propto \mathcal{N}$ versus the Zeeman shift $\propto \Omega^2 M \propto M$. Assuming that the electric field $\vec{\mathcal{E}} = \mathcal{E}_Z \hat{Z}$ and magnetic field $\vec{\mathcal{B}} = \mathcal{B}_Z \hat{Z}$ are collinear with the lab Z direction, the diagonal energy shifts for E1 and M1 are

$$\text{E1:} \quad = -\mathcal{E}_Z d_{\eta,\parallel} \frac{\Omega M}{J(J+1)}, \tag{2.69}$$

$$\text{M1:} \quad = -\mathcal{B}_Z \mu_{\eta,\parallel} \frac{\Omega M}{J(J+1)}, \tag{2.70}$$

where $d_{\eta,\parallel} \equiv \langle \eta, \Lambda | T^1_0(\vec{d}) | \eta, \Lambda \rangle$ and $\mu_{\eta,\parallel} \equiv \langle \eta, \Omega | T^1_0(\vec{\mu}) | \eta, \Omega \rangle$ are the body-fixed-frame dipole moments along the internuclear axis. It is customary to absorb the J dependence into the definition of the dipole moments such that $d_{\eta,J} \equiv \frac{d_{\eta,\parallel}}{J(J+1)}$ and $\mu_{\eta,J} \equiv \mu_B g_{\eta,J} \equiv \frac{\mu_{\eta,\parallel}\Omega}{J(J+1)}$, such that[4]

$$\text{E1:} \quad = -\mathcal{E}_Z d_{\eta,J}\, \Omega M = -\mathcal{E}_Z d_{\eta,J}\, \mathcal{N}, \tag{2.71}$$

$$\text{M1:} \quad = -\mathcal{B}_Z \mu_{\eta,J} M = -\mathcal{B}_Z g_{\eta,J} \mu_B M, \tag{2.72}$$

where $g_{\eta,J}$ is the J-dependent molecular g factor and $\mathcal{N} \equiv \Omega M \,\mathrm{sgn}(\mathcal{E}_Z)$ is the molecular orientation quantum number [102]. This orientation quantum number can be understood classically as follows. Consider the science state $H, J = 1$, where $J = |\Omega| = |M| = 1$. This implies $J = J_e$. From the definitions of Ω and M, we have $\Omega = \vec{J}_e \cdot \hat{n} = \vec{J} \cdot \hat{n}$ and $M = \vec{J} \cdot \hat{Z}$. Therefore, $\Omega M = \hat{n} \cdot \hat{Z}$ gives the direction of the molecular axis relative to the lab Z axis. Thus, $\mathcal{N} = \Omega M \,\mathrm{sgn}(\mathcal{E}_Z)$ describes whether the molecular axis is aligned or anti-aligned with the external electric field direction. Note that with this definition we can equivalently write $\hat{n} \cdot \hat{Z} = \Omega M = \mathcal{N} \,\mathrm{sgn}(\mathcal{E}_Z)$. This will be important in the upcoming discussion of the EDM measurement switching sequence.

We now proceed to the discussion of the off-diagonal matrix elements, which are important when considering E1 and M1 transitions or mixing between different states. The

4. Note that $\mu_{\eta,\parallel}$ carries a $\mathrm{sgn}(\Omega)$ dependence while $d_{\eta,\parallel}$ does not, as explained earlier. This means that both $d_{\eta,J}$ and $\mu_{\eta,J}$ do not depend on the sign of Ω, which is a more natural definition for these dipole moments. Beware that a different convention, in which the Ω dependence is not absorbed into $\mu_{\eta,J}$, is sometimes used in the literature [97].

off-diagonal matrix elements can be read off directly from eqs. 2.61 and 2.64. For the ACME experiment, the most important property of these off-diagonal matrix elements is the relative sign between the E1 and M1 matrix elements connecting the same two states as a function of Ω. In other words, we are interested in

$$\mathrm{sgn}\left(\frac{\langle \eta, \Lambda | T^1_{\Omega-\Omega'}(\vec{d}) | \eta', \Lambda' \rangle}{\langle \eta, \Omega | T^1_{\Omega-\Omega'}(\vec{\mu}) | \eta', \Omega' \rangle}\right) \text{ vs. } \mathrm{sgn}\left(\frac{\langle \eta, -\Lambda | T^1_{-\Omega+\Omega'}(\vec{d}) | \eta', -\Lambda' \rangle}{\langle \eta, -\Omega | T^1_{-\Omega+\Omega'}(\vec{\mu}) | \eta', -\Omega' \rangle}\right), \tag{2.73}$$

assuming that the transition is spin-allowed so that $\delta_{SS'}\delta_{\Sigma\Sigma'} = 1$. Let π be the parity operator. Using the identity $(\langle \pi\alpha|)(\pi O \pi)(|\pi\beta\rangle) = \langle \alpha | O | \beta \rangle$, where O is any operator, it is straightforward to show that

$$\mathrm{sgn}\left(\frac{\langle \eta, \Lambda | T^1_{\Omega-\Omega'}(\vec{d}) | \eta', \Lambda' \rangle}{\langle \eta, \Omega | T^1_{\Omega-\Omega'}(\vec{\mu}) | \eta', \Omega' \rangle}\right) = -\mathrm{sgn}\left(\frac{\langle \eta, -\Lambda | T^1_{-\Omega+\Omega'}(\vec{d}) | \eta', -\Lambda' \rangle}{\langle \eta, -\Omega | T^1_{-\Omega+\Omega'}(\vec{\mu}) | \eta', -\Omega' \rangle}\right), \tag{2.74}$$

because $\pi T^1_q(\vec{d})\pi = T^1_q(\vec{d})$ while $\pi T^1_q(\vec{\mu})\pi = -T^1_q(\vec{\mu})$. In other words, upon flipping the signs of $\Omega, \Lambda, \Sigma, \Omega', \Lambda', \Sigma'$, the relative sign between the E1 and M1 matrix elements flips. Colloquially, we say that the E1 and M1 matrix elements differ by a sign of Ω [91].

So far the discussion has been somewhat abstract, but this sign difference has a very intuitive classical explanation, as shown in figure 2.4. Essentially, the internuclear axis $\hat{n}$ determines the interaction with electric fields, while the direction of electron rotation $\vec{J}_e$ determines the interaction with magnetic fields. The projection $\Omega = \vec{J}_e \cdot \hat{n}$ therefore determines the relative sign between the E1 and M1 matrix elements.

The Ω-dependent relative sign between the E1 and M1 matrix elements has very important consequences in ACME. From eq. 2.5, we see that the eEDM Hamiltonian also flips sign with Ω. This means that, in the presence of non-reversing fields, interference between the E1 and M1 transition matrix elements can closely resemble, or even exactly mimic, an eEDM signal. In ACME II and III, there is a very significant source of E1–M1-interference-induced systematic error due to AC fields, which we will discuss in more detail in the systematic error

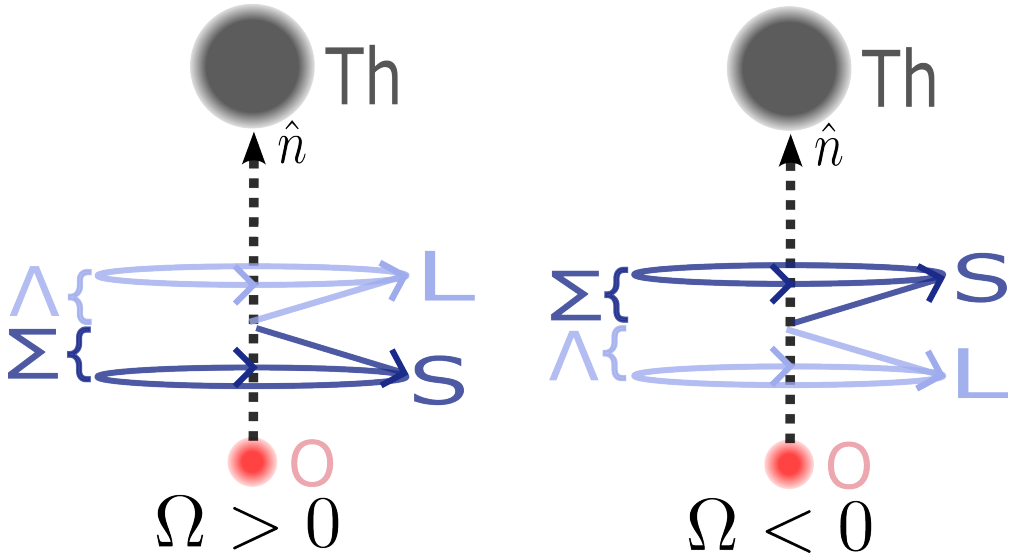

Figure 2.4: Classical pictures of molecular rotational states with different signs of Ω. The electrons rotate around the internuclear axis in opposite directions for opposite signs of Ω. This leads to opposite magnetic dipole moments while the electric dipole moments are the same, resulting in a sign difference between E1 and M1 matrix elements upon flipping the sign of Ω.

chapter. DC electric and magnetic fields also have the potential to generate large systematic errors from such interference, but we show here that this can be calibrated out by reversing the direction of the electric or magnetic field.

Consider a DC electric field $\vec{\mathcal{E}} = \mathcal{E}_Z \hat{Z}$ and a magnetic field $\vec{\mathcal{B}} = \mathcal{B}_Z \hat{Z}$. There is off-diagonal mixing between different J states. For example, the $H, J = 1$ state can mix with the $H, J = 2$ state through both E1 and M1 couplings. They are separated by $4B_H \sim 40$ GHz, and the second-order energy shift can be estimated as

$$\frac{|\langle J=1,\Omega|\,\mathrm{E1{+}M1}\,|J=2,\Omega\rangle|^2}{4B_H} \approx \frac{\left|\langle J=1,\Omega|\,\mathcal{E}_Z d_{H,\parallel} + \Omega\mathcal{B}_Z\mu_{H,\parallel}\,|J=2,\Omega\rangle\right|^2}{4B_H} \tag{2.75}$$

$$\approx \frac{\mathcal{E}_Z^2 d_{H,\parallel}^2 + 2\Omega\mathcal{E}_Z\mathcal{B}_Z d_{H,\parallel}\mu_{H,\parallel} + \mathcal{B}_Z^2\mu_{H,\parallel}^2}{4B_H}, \tag{2.76}$$

where we have approximated rotational factors to be of order unity for simplicity. The second term is the E1–M1 interference term, which flips sign with Ω. This looks exactly

like an eEDM signal, except that it depends on the applied electric and magnetic fields. In principle, it can be calibrated out by reversing the direction of the magnetic field $\mathcal{B}_Z \to -\mathcal{B}_Z$. However, in practice there is always a small residual magnetic field that does not reverse perfectly when performing the magnetic-field switch. Therefore, it is also important to reverse the electric field $\mathcal{E}_Z \to -\mathcal{E}_Z$ to calibrate out this systematic. The combination of these two switches ensures that it is possible to confirm any residual E1–M1 interference effect due to DC fields is below our desired sensitivity. As an additional precaution, we also vary the magnitude of the applied electric field to search for possible similar effects that we may not have anticipated.

The interference effect described above was first observed and understood in the PbO eEDM experiment [41]. By ACME convention, we call this effect the $\mathcal{N}$-correlated g factor for the following reasons. Writing its Ω-dependent energy correction using $\mathcal{N} = \Omega \operatorname{sgn}(\mathcal{E}_Z) M$, we obtain

$$\mathcal{B}_Z \left(\frac{\mathcal{E}_Z d_{H,\|} \mu_{H,\|}}{2B_H} \right) \mathcal{N} \operatorname{sgn}(\mathcal{E}_Z) M = \mathcal{B}_Z \left(\frac{|\mathcal{E}_Z| d_{H,\|} \mu_{H,\|}}{2B_H} \right) \mathcal{N} M, \tag{2.77}$$

which has exactly the form of a g factor correlated with $\mathcal{N}$ and linearly proportional to $|\mathcal{E}_Z|$, hence its name. It turns out that additional molecular perturbation effects also produce an energy correction with this same functional form; these are explained in detail in Nick Hutzler's thesis [53].

2.1.8 EDM interaction and the $^3\Delta_1$ state

The interaction between the electron EDM and the internal molecular electric field is best illustrated pictorially. The classical picture is shown in figure 2.5. The key property is that the electron EDM $\vec{d_e}$ is proportional to the electron spin $\vec{S}$, so they always flip directions together. Under the time-reversal operation $\mathcal{T}$, both the spin $\vec{S}$ and the electron EDM $\vec{d_e}$ flip direction. This is different from a classical electric dipole moment, which does not flip direction under time reversal. This difference results in a time reversal-odd Hamiltonian

term

$$H_{\mathrm{EDM}} = -\vec{d}_e \cdot \vec{\mathcal{E}}_{\mathrm{eff}} \tag{2.78}$$

$$= -(d_e \hat{S}) \cdot (-\mathcal{E}_{\mathrm{eff}} \hat{n}) \tag{2.79}$$

$$= d_e \mathcal{E}_{\mathrm{eff}} \Sigma \tag{2.80}$$

$$= -d_e \mathcal{E}_{\mathrm{eff}} \Omega, \tag{2.81}$$

where we have used the fact that in the science state $H\,^3\Delta_1$, the electron spin projection satisfies $\Sigma = -\Omega$. This is exactly the same term written earlier in eq. 2.5. In order for the

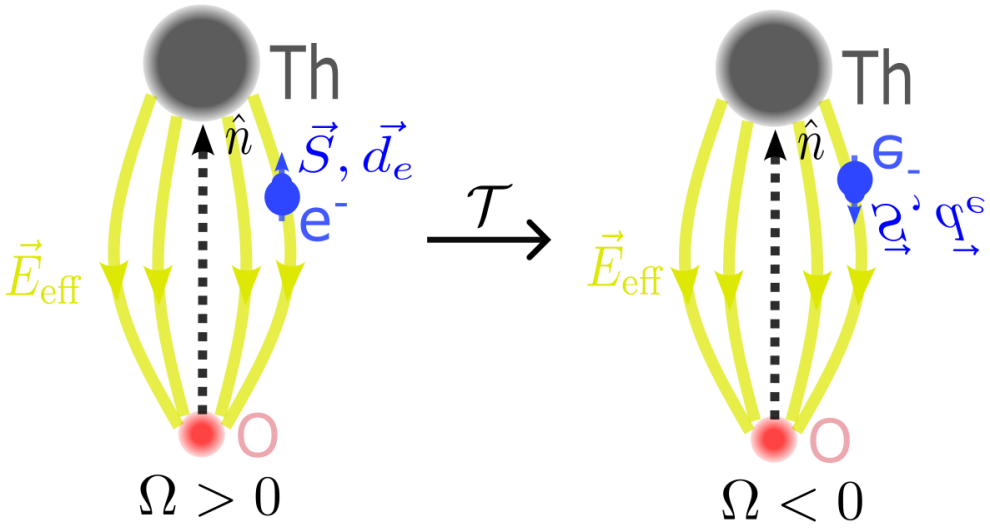

Figure 2.5: Classical picture of the interaction between the electron EDM and the effective internal molecular electric field. The molecular state on the left is related to the state on the right by the time-reversal operation. The electronic angular momentum projection $\Omega = \vec{J}_e \cdot \hat{n}$ flips under time reversal. In addition, because $\vec{d}_e \propto \vec{S}$, the electron EDM also flips direction under time reversal. The EDM interaction therefore changes sign between these two states, leading to an energy shift equal to $-d_e \mathcal{E}_{\mathrm{eff}}$ depending on the sign of Ω.

effective electric field $\mathcal{E}_{\mathrm{eff}}$ to be large, at least one of the valence electrons must occupy a σ orbital so that it can approach the heavy nucleus, where relativistic effects are large. This requires the other valence electron to occupy a δ orbital to produce the desired $^3\Delta_1$ state. This is indeed the case for ThO. The two valence electrons in ThO are expected to have

$6d\sigma$ and $7s\delta$ character [94]. The exact value of $\mathcal{E}_{\text{eff}}$ can be most accurately computed using relativistic many-body calculations, with the most up-to-date value being $\mathcal{E}_{\text{eff}} = 78$ GV/cm and an uncertainty of approximately 3% [33, 34]. The value of $\mathcal{E}_{\text{eff}}$ cannot be checked experimentally without knowing the value of d_e. However, these relativistic calculations also predict other molecular properties, such as the electronic state energies, molecular-frame dipole moment, and hyperfine constants, which can be measured experimentally and used for indirect validation.

More generally, there may be additional time-reversal-symmetry-violating interactions between the electrons and the molecule, such as the scalar–pseudoscalar nucleon–electron interaction $-W_S C_S$, giving the full time-reversal-symmetry-violating Hamiltonian

$$\langle H_{\mathcal{T}} \rangle = -d_e \mathcal{E}_{\text{eff}} \Omega - W_S C_S \Omega. \tag{2.82}$$

Notably, any time-reversal-symmetry-violating interaction within the science state H must be correlated with the sign of Ω, because Ω is also $\mathcal{T}$-odd. The ACME experiment thus probes multiple possible sources of time-reversal-symmetry violation simultaneously. It is mathematically impossible to disentangle different sources of time-reversal-symmetry violation using only a single molecular species. Therefore, experiments with different molecular species are required in order to distinguish among them once nonzero signals are observed.

The g factor of Hund's case (a) states can be easily estimated. We are interested in the matrix element of $\vec{\mu} = g_S \mu_B \vec{S} + g_L \mu_B \vec{L}$, where $g_S \approx 2$ and $g_L = 1$. Because both the spin and orbital angular momenta are pinned to the internuclear axis, we can approximate

$$\left| \mu_{\eta,\parallel} \right| \approx \mu_B (2\Sigma + \Lambda) \tag{2.83}$$

$$\approx 0 \quad \text{for the } {}^3\Delta_1 \text{ state.} \tag{2.84}$$

In other words, because the spin and orbital angular momenta are anti-aligned in the correct

proportions in the $^3\Delta_1$ state, their magnetic dipole moments nearly cancel. The classical picture of this cancellation is shown in figure 2.2. In reality, this cancellation is not exact because the electron g factor is not exactly 2 and because of additional contributions from spin-orbit mixing with other molecular states. This is worked out in detail in Nick Hutzler's thesis [53].

The $^3\Delta_1$ state is not unique in having a small g factor; other states such as $^2\Pi_{1/2}$ also exhibit similar cancellation effects. However, the $^2\Pi_{1/2}$ state is not well suited for an EDM measurement because its valence electron occupies a π orbital, which does not produce $\mathcal{E}_{\text{eff}}$ due to the centrifugal barrier suppressing relativistic enhancement. Thus, for Hund's case (a) states, at least two valence electrons are required to simultaneously achieve a very small molecular g factor and the strongest EDM sensitivity.

The class of two-valence-electron molecules in the $^3\Delta_1$ state has been the most popular choice for electron EDM measurements for the past 15 years. In recent years, however, there has been growing interest in molecules with a single valence electron, which generally allow much greater quantum control. If one additionally considers polyatomic molecules, it is possible to recover most of the features that make $^3\Delta_1$ states attractive for EDM measurements [51, 59]. Among these, polyatomic molecules with integer total angular momentum F appear particularly promising for beam experiments, as they have been shown to be compatible with ACME-style measurements [60].

2.1.9 Other effects

There are other small effects in molecules that can be treated as effective Hamiltonian terms. A common source of such terms comes from the expansion of the molecular-axis rotational Hamiltonian term $B(\vec{J}-\vec{J}_e)^2$. As shown in the spin-orbit coupling subsection 2.1.3

and the Ω-doubling subsection 2.1.6, there are terms such as

$$B\vec{J}\cdot\vec{S} \quad \text{and} \quad B\vec{J}\cdot\vec{J}_e, \tag{2.85}$$

where the first term describes the spin-uncoupling effect and the second term describes the Ω-doubling effect. They differ simply by an L-uncoupling term $\vec{J}\cdot\vec{L}$ and are therefore not independent. We chose the Ω-doubling term in subsection 2.1.6 because it is the more suitable choice for Hund's case (a) and (c) states, where Ω is a good quantum number.

Other than the Ω-doubling term, additional Hamiltonian terms are generally not necessarily small compared to the Stark splitting we apply in the lab (typically a few hundred MHz). However, they are almost always diagonal in the $|\Omega|, J, M$ basis and therefore do not cause observable effects in our measurements. Examples include the spin–spin coupling between different electrons and the spin–rotation coupling between the electron spin and the nuclear rotation, which are proportional to $\lambda\left(3\Sigma^2 - S(S+1)\right)$ and $\gamma\,\vec{S}\cdot(\vec{J}-\vec{J}_e)$, respectively [50].

It is obvious from a symmetry point of view why there is a lack of correction terms that are odd in M, which are the only kinds of Hamiltonian terms measured in ACME. Any such term would require the breaking of time-reversal symmetry, which is not present in the standard molecular Hamiltonian without external fields. To induce such M-odd terms, one would need time-odd effects such as a $\vec{\mathcal{B}}$ field, rotation of the quantization axis, or rotation of the coordinate system itself (e.g. Earth's rotation).

In the external-field interaction subsection 2.1.7, we discussed how an $\mathcal{N}$-correlated g factor can arise from E1–M1 interference effects. In fact, there is another contribution from third-order perturbation theory involving the Stark effect, the Zeeman effect, and the spin-uncoupling interaction. This was first worked out in an effort led by Nick Hutzler [53], and it was shown to be comparable in magnitude to the E1–M1 interference effect discussed in 2.1.7.

2.2 Spin precession measurement and switches

2.2.1 Ramsey measurement in $J = \frac{1}{2}$

We perform our EDM measurement using a Ramsey-type spin precession measurement. It is not exactly a Ramsey measurement, but for background we briefly review the standard Ramsey technique here. In a standard Ramsey measurement, one seeks to measure the energy difference between two states $|\uparrow\rangle$ and $|\downarrow\rangle$. The measurement sequence is as follows. First, one prepares a superposition state $\frac{|\uparrow\rangle+|\downarrow\rangle}{\sqrt{2}}$ using a $\frac{\pi}{2}$ pulse. Then, one lets the state freely precess for a time T under the Hamiltonian

$$H = \frac{\hbar\omega}{2}(|\uparrow\rangle\langle\uparrow| - |\downarrow\rangle\langle\downarrow|). \tag{2.86}$$

After time T, the state evolves to

$$|\psi(T)\rangle = \frac{|\uparrow\rangle + e^{-i\omega T}|\downarrow\rangle}{\sqrt{2}}. \tag{2.87}$$

Finally, one applies another $\frac{\pi}{2}$ pulse and performs state readout. The result is fluorescence proportional to

$$S = |\langle\uparrow|\psi(T)\rangle|^2 \tag{2.88}$$

$$= \frac{1}{2}\left(1 + \cos(\omega T)\right). \tag{2.89}$$

If both the up and down states are detected, one can define the asymmetry as

$$\mathcal{A} = \frac{S_\uparrow - S_\downarrow}{S_\uparrow + S_\downarrow} = \cos(\omega T). \tag{2.90}$$

By measuring the asymmetry $\mathcal{A}$, one can extract the precession frequency ω and thus the energy difference between the two states, given knowledge of the precession time T.

In general, it is convenient for visualization and mathematical simplicity to map any two-level system onto a spin-$\frac{1}{2}$ system. The quantum state is then described by two angular parameters θ and ϕ on the Bloch sphere. The state $|\uparrow\rangle$ corresponds to the north pole with $\theta = 0$, while the state $|\downarrow\rangle$ corresponds to the south pole with $\theta = \pi$. A general pure state is given by

$$|\psi\rangle = e^{-i\phi/2}\cos\left(\frac{\theta}{2}\right)|\uparrow\rangle + e^{i\phi/2}\sin\left(\frac{\theta}{2}\right)|\downarrow\rangle. \tag{2.91}$$

This quantum state can be represented by a vector $\vec{S}$ on the Bloch sphere with polar angle θ and azimuthal angle ϕ. Any non-dissipative evolution of this two-level system can be represented as a rotation of the Bloch vector $\vec{S}$ about some axis.

2.2.2 Ramsey-like measurement in $J = 1$ and correspondence with photon polarization

ACME's spin precession measurement is very similar to a Ramsey measurement, but there are some important differences. First, the two-level system we are working with is not a spin-$\frac{1}{2}$ system, but rather the $J = 1$ rotational manifold of the H state. In particular, we are only interested in the two $M = \pm 1$ sublevels of the $J = 1$ state, which we denote as $|M = +1\rangle$ and $|M = -1\rangle$. The $H, M = 0$ sublevel is not used in our measurement. This is similar to plane-wave photons, where only two polarization states are used, with the longitudinal polarization state being prohibited by Maxwell's equations. In fact, this similarity can be taken almost literally.

For the science H state in ACME, we only drive transitions between the $H, J = 1, M = \pm 1$ states and the $J = 1, M = 0$ state of an electronically excited state (either C or I); an example laser coupling is shown in figure 2.6. Clearly, there will always be one polarization bright state and one polarization dark state in this Λ-type three-level system. The bright

and dark states can be defined as

$$|B\rangle = e^{-i\theta}\sin\Theta|M=+1\rangle + e^{i\theta}\cos\Theta|M=-1\rangle, \tag{2.92}$$

$$|D\rangle = -e^{-i\theta}\cos\Theta|M=+1\rangle + e^{i\theta}\sin\Theta|M=-1\rangle. \tag{2.93}$$

If we change the polarization of the laser light, the bright and dark states will change as well. It turns out that the parameters θ and Θ defined here for the bright and dark states rotate exactly like the polarization of the laser.

Let us define the coordinate system using the readout laser. The readout light is linearly polarized along either the x or y direction[5]. This choice defines the local x and y axes and therefore fully specifies the polarization bright and dark states. In this coordinate system, the bright and dark states are given by

$$|B\rangle = \frac{|M=+1\rangle + |M=-1\rangle}{\sqrt{2}}, \tag{2.94}$$

$$|D\rangle = \frac{|M=+1\rangle - |M=-1\rangle}{\sqrt{2}}, \tag{2.95}$$

with $\theta = 0$ and $\Theta = \frac{\pi}{4}$. It turns out that if we change the laser polarization, the dark state will change exactly like the polarization state of a photon. For example, if we change the laser polarization to right circular polarization, the dark state becomes

$$|D\rangle = |M=-1\rangle, \tag{2.96}$$

which corresponds to a right-circularly polarized photon state. In general, θ and Θ of the bright and dark states, as defined in Eq. 2.94 and 2.95, correspond exactly to the polarization

5. These axes do not necessarily coincide with the laboratory X and Y directions; they instead define a local measurement frame set by the chosen linear polarization. Following ACME convention, we use x, y and X, Y interchangeably when no ambiguity arises.

parameters of a photon, where θ describes the orientation of the polarization ellipse and Θ describes the ellipticity. This correspondence is very useful for visualizing and understanding our spin precession measurement in ACME, and is the main reason we often do not directly map the quantum state to a spin-$\frac{1}{2}$ system but instead to a spin-1 system.

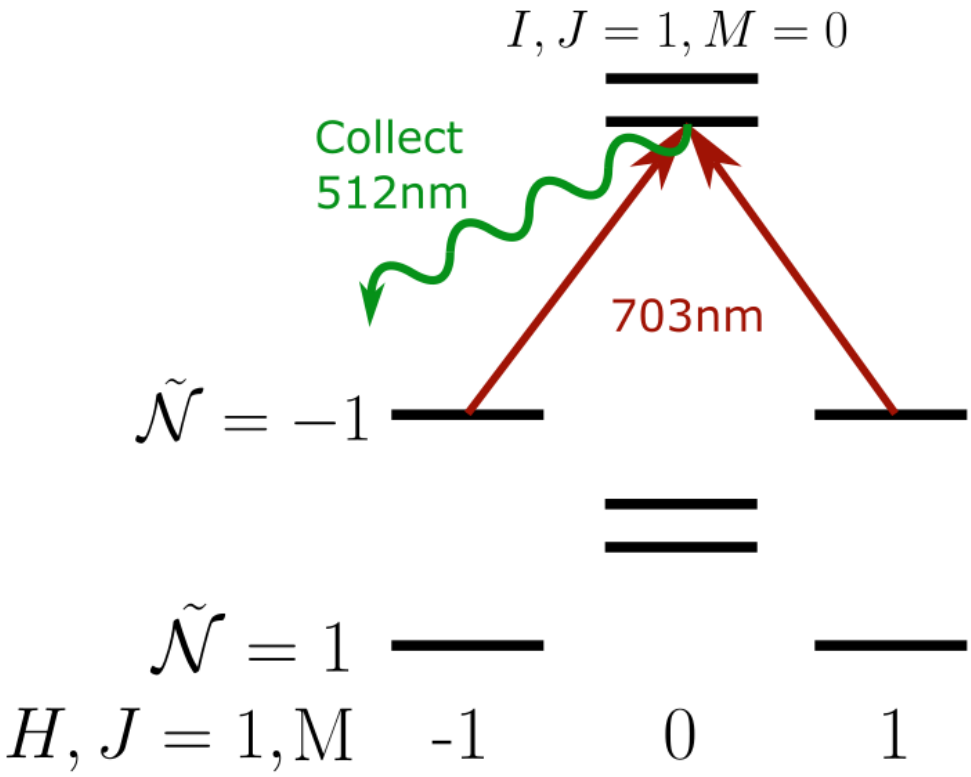

Figure 2.6: Laser coupling the $H, J = 1, M = \pm 1$ states to the $I, J = 1, M = 0$ state. The resulting 512 nm fluorescence is collected for state readout.

We are now ready to describe the ACME spin precession measurement sequence. First, we prepare the initial state by optically pumping into a dark state using diagonal linear polarization. The prepared state is thus

$$|\psi_i\rangle = |D\rangle = \frac{-e^{-i\pi/4}|M = +1\rangle + e^{i\pi/4}|M = -1\rangle}{\sqrt{2}}. \tag{2.97}$$

Next, we let the state freely precess for a time T, with the $M = \pm 1$ states separated in

energy by ω. After time T, the state evolves to

$$|\psi(T)\rangle = \frac{-e^{-i(\omega T+\pi/4)}|M=+1\rangle + e^{i(\omega T+\pi/4)}|M=-1\rangle}{\sqrt{2}}. \tag{2.98}$$

Finally, we perform state readout by alternating x and y linear polarization rapidly at around 178 kHz. This is fast compared to the molecule transit time through the detection laser beam, which is around 10 μs, so that each molecule experiences both polarizations during readout. The resulting fluorescence signals for each polarization are proportional to

$$S_x = |\langle B_x|\psi(T)\rangle|^2 = \sin^2(\omega T + \pi/4), \tag{2.99}$$

$$S_y = |\langle B_y|\psi(T)\rangle|^2 = \cos^2(\omega T + \pi/4), \tag{2.100}$$

where $|B_x\rangle$ and $|B_y\rangle$ are the bright states for x and y linear polarization, respectively. By defining the asymmetry as

$$\mathcal{A} = \frac{S_x - S_y}{S_x + S_y} = -\sin(2\omega T), \tag{2.101}$$

we see that, instead of the $\cos(\omega T)$ dependence in a standard Ramsey measurement, we obtain a $\sin(2\omega T)$ dependence. The $\pi/4$ phase offset arises from the diagonal linear polarization used in the initial state preparation, while the factor of 2 in front of ωT comes from the fact that the $J = 1$ system precesses twice as fast as a spin-$\frac{1}{2}$ system for the same energy splitting between the two levels. The asymmetry has the standard shot-noise-limited uncertainty $\delta\mathcal{A} = 1/\sqrt{N}$ for $\langle\mathcal{A}\rangle = 0$, where N is the total number of detected photons. The corresponding uncertainty in the precession phase $\phi = \omega T$ is therefore $\delta\phi = 1/(2\sqrt{N})$.

The ACME measurement scheme described above is not exactly a Ramsey measurement, not only because we are working with a $J = 1$ system instead of a spin-$\frac{1}{2}$ system, but also because we do not apply explicit $\frac{\pi}{2}$ pulses for state preparation and readout. Instead, we use optical pumping to prepare the initial dark state and optical detection for readout. Although

this is not a traditional Ramsey $\frac{\pi}{2}$ pulse sequence, the spirit and outcome are the same. The preparation creates a superposition state in the $M = \pm 1$ manifold, and the readout projects the final state onto two orthogonal states (here in the x-y plane of the Bloch sphere rather than along $\pm z$) to extract the precession phase. Therefore, we still sometimes refer to this as a Ramsey-type spin precession measurement.

2.2.3 Switches and parity channels

In a perfectly ideal experiment, one would have perfect contrast and know the precession time T exactly. The electric and magnetic fields would be perfectly controlled so that the only contribution to the precession frequency ω would come from the eEDM interaction. In this case, the simple Ramsey-type measurement described above would be sufficient to extract the eEDM signal. However, in reality there are many imperfections and additional phase contributions not related to the eEDM. Therefore, we perform a series of experimental switches to characterize the contrast and precession time and to isolate the eEDM signal from other contributions.

Most switches take binary values of ± 1. For a set of switches such as $\{\mathcal{N}, \mathcal{E}, \mathcal{B}, \mathcal{P}, \ldots\}$, we denote their values as $\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, \tilde{\mathcal{P}}, \ldots = \pm 1$. The full set of switches is described below.

1. **Polarization switch**: We rapidly switch the readout laser polarization between x and y linear polarization at around 178 kHz to perform state readout. The integrated fluorescence when the polarization is x/y is denoted S_x/S_y, respectively. We do not assign a switch label to the readout laser polarization. Instead, we use these signals to compute an asymmetry at time t_i given by

$$\mathcal{A}(t_i) = \frac{S_x(t_i) - S_y(t_i)}{S_x(t_i) + S_y(t_i)},$$

which is used to extract the accumulated precession phase. The relative timescales of

the switches are shown in figure 2.7. The molecular beam pulse is generated at 50 Hz; data from each pulse is called a "shot". We sum 25 shots to form a "trace". After collecting a trace, additional switches are applied before collecting the next trace.

2. **$\mathcal{N}$ switch (molecular orientation)**: This corresponds to the orientation of the molecule relative to the applied electric field. In particular, this switch is performed by changing the frequency of the state preparation laser to spectroscopically select molecular states with different orientations. $\tilde{\mathcal{N}} = \pm 1$ corresponds to the molecular axis aligned or anti-aligned with the applied electric field.

3. **$\mathcal{E}$ switch (electric field direction)**: This corresponds to reversing the direction of the applied electric field in the lab. This is performed by reversing the voltage applied to the electrodes. $\tilde{\mathcal{E}} = \pm 1$ corresponds to the field pointing along or against the lab Z direction.

4. **θ switch (polarization dither)**: Due to imperfections such as molecular beam velocity dispersion and polarization imperfections, the measured asymmetry depends on both the phase and the contrast,

$$\mathcal{A} = \mathcal{C} \sin\big(2(\phi - \theta)\big),$$

where $\phi \ll 1$ is the precession phase, $\mathcal{C}$ is the contrast, and $\theta = \pm 1^\circ$ is a small dither, or deliberate offset, of the readout laser linear polarization away from the nominal readout direction ($\theta = 0$). By dithering θ between $\pm 1^\circ$, we estimate the contrast in situ for each trace using

$$\mathcal{C} = -\frac{1}{2}\frac{\partial \mathcal{A}}{\partial \theta} \approx \frac{1}{2}\frac{\partial \mathcal{A}}{\partial \phi}. \tag{2.102}$$

The phase is then estimated as

$$\phi = \frac{\mathcal{A}}{2\mathcal{C}}. \tag{2.103}$$

These relations rely on small-angle approximations whose validity has been examined in previous generations of ACME [91]. The switch value $\tilde{\theta} = \pm 1$ corresponds to $\theta = \pm 1^\circ$.

5. **$\mathcal{B}$ switch (magnetic field direction)**: This corresponds to reversing the applied magnetic field. This is performed by reversing the applied current of the magnetic field coils. $\tilde{\mathcal{B}} = \pm 1$ corresponds to the field pointing along or against the lab Z direction[6]. The difference in precession phase between the two $\tilde{\mathcal{B}}$ states is used to extract the precession time using the calibrated magnetic moment of the H state and the applied field magnitude. The $\mathcal{B}$ switch is performed every 8 traces; the resulting 16 traces (both $\mathcal{B}$ states) are combined to form a block, which is the smallest unit used to extract the precession time.

   The $\mathcal{P}$, $\mathcal{R}$, $\mathcal{L}$, and $\mathscr{C}$ switches explained below are collectively referred to as superblock switches. There are $2^4 = 16$ possible combinations. The corresponding 16 blocks are combined to form a superblock, which is the smallest unit of data expected to be free of systematic errors and is therefore the smallest unit used to compute an eEDM result.

6. **$\mathcal{P}$ switch (readout parity)**: The readout laser couples the $H, J = 1, M = \pm 1$ states to an excited $I, J = 1, M = 0$ state. Due to Ω-doubling in the I state, there are two such $M = 0$ states with opposite parity seperated by around 91 MHz. The $\mathcal{P}$ switch is performed through frequency selection of the excited-state parity, with $\tilde{\mathcal{P}} = \pm 1$ corresponding to even or odd parity. Effectively, this operation rotates the readout laser polarization by 90° similar to the $\mathcal{R}$ switch.

7. **$\mathcal{R}$ switch (readout polarization rotation)**: By rotating the readout half wave

6. With the exception of the readout X,Y lasers, the coordinate system of this writing follows ACME convention where molecules fly toward $+X$ direction, gravity points in $+Y$ direction, and Z is defined by the right hand rule.

plate by 45°, the $\mathcal{R}$ switch rotates the readout laser polarization by 90°, effectively swapping the definitions of x and y linear polarization.

8. $\boldsymbol{\mathcal{L}}$ **switch (lead reversal)**: This switch changes the lead assignment for the voltage supplies connected to the field plates without changing the actual electric field direction. Both plate voltages are reversed simultaneously to keep the field direction fixed. This switch is used to diagnose imperfections associated with the voltage supplies.

9. $\mathscr{C}$ **switch (cleanup birefringence)**: This switch changes the birefringence seen by the preparation (cleanup) laser, which prepares the initial dark state through optical pumping. It is implemented by simultaneously rotating the cleanup laser polarization by 90° and switching the parity of the $M = 0$ excited state used for cleanup. The net effect is that the effective linear polarization seen by the molecules is unchanged while the ellipticity induced by birefringence flips sign. This is a new switch introduced in ACME III to control cleanup-related systematics.

10. **Slow switches**: In addition to the fast switches above, there are slow switches applied on timescales of hours to days, such as changing the electric field magnitude, readout laser power, and other experimental parameters. These are used to probe potential systematic effects. At the time of writing, the final set of slow switches for the ACME III data run has not yet been finalized.

After we collect a superblock of data, we compute what we call parity channels. For any measured quantity X (e.g. asymmetry $\mathcal{A}$ or precession phase ϕ), we can decompose it into different parity channels according to its dependence on the various switches. For example, the $\mathcal{NE}$-correlated channel of $\mathcal{A}$ is computed as

$$\mathcal{A}^{\mathcal{NE}} = \frac{1}{2^n} \sum_{\tilde{\mathcal{N}},\tilde{\mathcal{E}},\tilde{\mathcal{B}},\tilde{\mathcal{P}},\cdots} \tilde{\mathcal{N}}\tilde{\mathcal{E}}\, \mathcal{A}(\tilde{\mathcal{N}},\tilde{\mathcal{E}},\tilde{\mathcal{B}},\tilde{\mathcal{P}},\cdots), \tag{2.104}$$

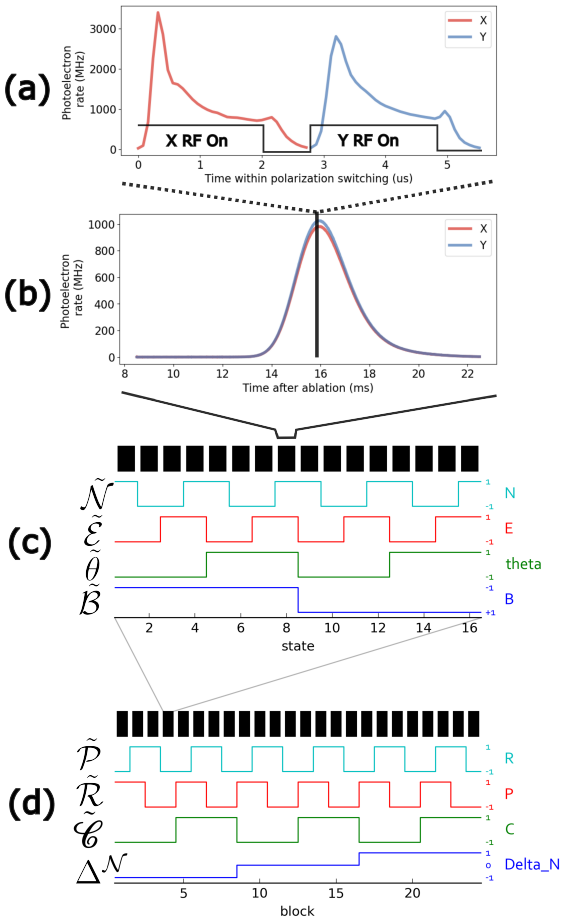

Figure 2.7: Example switching timescales used for checking systematic errors. (a) The fastest timescale is the polarization switch, which alternates between X and Y linear polarization every 2.8 μs. (b) Each molecular beam pulse (shot) is generated at 50 Hz, and only around 14 ms of data is collected from each pulse. We combine 25 shots to form a trace. (c) After every trace, one or more of $\mathcal{N}, \mathcal{E}, \theta, \mathcal{B}$ is switched. Here, $\mathcal{N}$ and $\mathcal{E}$ are switched every other trace so that $\mathcal{NE}$ is switched every trace. (d) After every block (16 traces), one or more of the superblock switches $\mathcal{P}, \mathcal{R}, \mathcal{L}, \mathscr{C}$ is switched. Sometimes $\mathcal{L}$ is not switched to speed up data collection. After collecting one superblock, we vary the knob of interest (e.g., $\Delta^{\mathcal{N}}$ or $\mathcal{N}$-correlated detuning). See 4.1.1 for details about systematics checks.

where n is the total number of switches in $\{\mathcal{N}, \mathcal{E}, \mathcal{B}, \mathcal{P}, \ldots\}$, and $\mathcal{A}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, \tilde{\mathcal{P}}, \cdots)$ is the measured asymmetry for a particular combination of switch values. Equivalently, we can write the measured quantity $\mathcal{A}$ as a sum over all possible parity channels:

$$\mathcal{A}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, \tilde{\mathcal{P}}, \cdots) = \mathcal{A}^{\mathrm{nr}} + \mathcal{A}^{\mathcal{N}}\tilde{\mathcal{N}} + \mathcal{A}^{\mathcal{E}}\tilde{\mathcal{E}} + \mathcal{A}^{\mathcal{NE}}\tilde{\mathcal{N}}\tilde{\mathcal{E}} + \mathcal{A}^{\mathcal{B}}\tilde{\mathcal{B}} + \cdots, \tag{2.105}$$

where $\mathcal{A}^{\mathrm{nr}}$ is the non-reversing channel that is not correlated with any switch, $\mathcal{A}^{\mathcal{N}}$ is the $\mathcal{N}$-correlated channel, and $\mathcal{A}^{\mathcal{NE}}$ is the channel correlated with the product of the $\mathcal{N}$ and $\mathcal{E}$ switch values. The ellipsis denotes all other possible combinations of switches.

In general, the set of values $\mathcal{A}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, \tilde{\mathcal{P}}, \cdots)$ forms a 2^n-dimensional vector space, and so does the set of parity channels $\{\mathcal{A}^{\mathcal{N}}, \mathcal{A}^{\mathcal{E}}, \mathcal{A}^{\mathcal{NE}}, \ldots\}$. We refer to the former as the *state basis* and the latter as the *parity basis*. These two vector spaces describe the same data and are isomorphic. By convention, we call the isomorphism that maps the state basis to the parity basis the *parity transformation*. Its explicit matrix representation depends on the ordering of the basis states but follows directly from the definitions above.

We record measurements naturally in the state basis. However, for data analysis it is often much more convenient to work in the parity basis because different channels isolate different physical effects. For example, the $\mathcal{B}$ channel is dominated by the Zeeman effect and is used to extract the precession phase. The eEDM contribution is expected to appear only in the $\mathcal{NE}$ channel. To see this, recall that the eEDM Hamiltonian term is proportional to $\Omega = \tilde{\mathcal{N}}\tilde{\mathcal{E}}\, M$, and the measured phase corresponds to the energy difference between the $M = \pm 1$ states. Physically, $\tilde{\mathcal{N}}$ defines the orientation of the molecule relative to the applied electric field, while $\tilde{\mathcal{E}}$ defines the direction of the applied electric field in the lab. Their product therefore specifies the orientation of the molecular axis in the lab frame, which in turn determines the direction of the effective internal electric field $\mathcal{E}_{\mathrm{eff}}$ and the sign of the eEDM energy shift. This switching protocol is crucial for isolating the eEDM signal from other contributions to the precession phase.

2.3 Experimental setup

2.3.1 Experimental schematic

The schematic of the ACME III experimental setup is shown in figure 2.8. The main components of the experiment include the molecular beam source, compact rotational cooling stage, molecular electrostatic lens for beam focusing, "Qleanup" stage for state preparation, and the interaction region for spin precession measurement. Each of these components is described in detail in the following subsections.

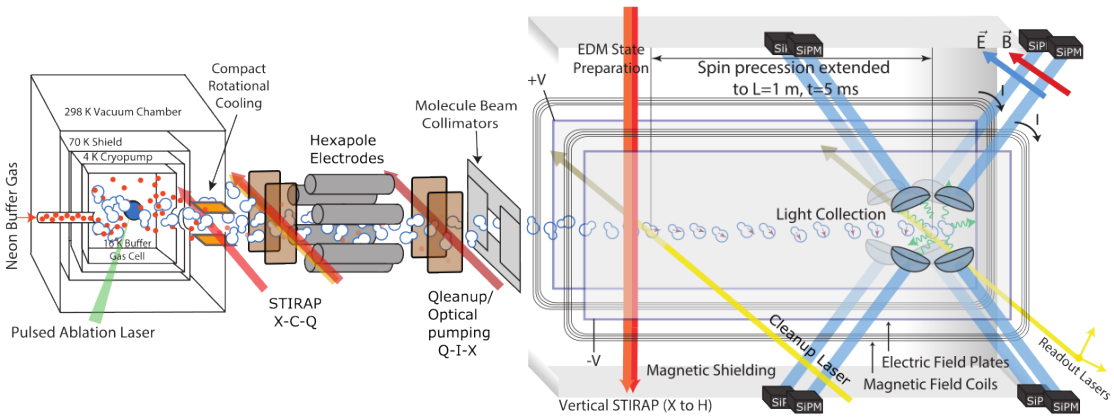

Figure 2.8: Experimental schematic of the ACME III electron EDM measurement apparatus. The ThO molecules are produced in a cryogenic buffer gas beam source, rotationally cooled to the ground $X, J = 0$ state using optical pumping, transferred to the $Q, J = 2$ lensing state, focused using an electrostatic lens, transferred back to the ground $X, J = 0$ state using the "Qleanup" optical pumping stage, prepared into the science $H, J = 1$ state using vertical STIRAP, with imperfections in the prepared states from this process subsequently cleaned up by the cleanup/preparation laser, and finally the spin precession measurement is performed in the interaction region using a Ramsey-type measurement. The resulting fluorescence from state readout is collected using a light collection system with SiPM detectors.

2.3.2 Beam source

We produce ThO molecules by ablating a ThO_2 ceramic target with ~ 15 mJ pulses from a Nd:YAG laser. ThO molecules are subsequently formed in the locally high-temperature region. The ThO molecules are then rapidly cooled through collisions with cold neon buffer

gas at around 20 K. The cold ThO molecules are extracted from the cell and undergo further cooling and collimation in a cryogenic buffer gas beam source [53, 54]. The resulting molecular beam has a forward velocity of around 200 m/s with a rotational temperature of approximately 4 K.

2.3.3 Compact rotational cooling

To gather the molecules into a single quantum state, we perform rotational cooling by optically pumping molecules from the $J = 1, 2$ states to the ground $X, J = 0$ state via the excited $C, J = 1$ state. The rotational cooling level scheme is shown in figure 2.9. The rotational cooling laser is multi-passed using an in-vacuum prism. The resulting rotational cooling is compact and efficient, with approximately a factor of 3 increase in the population of the ground $X, J = 0$ state. This compactness is required to reduce the distance between the beam source and the molecular electrostatic lens in order to minimize loss from beam divergence. The resulting distance between the source and the lens is around 30 cm [103].

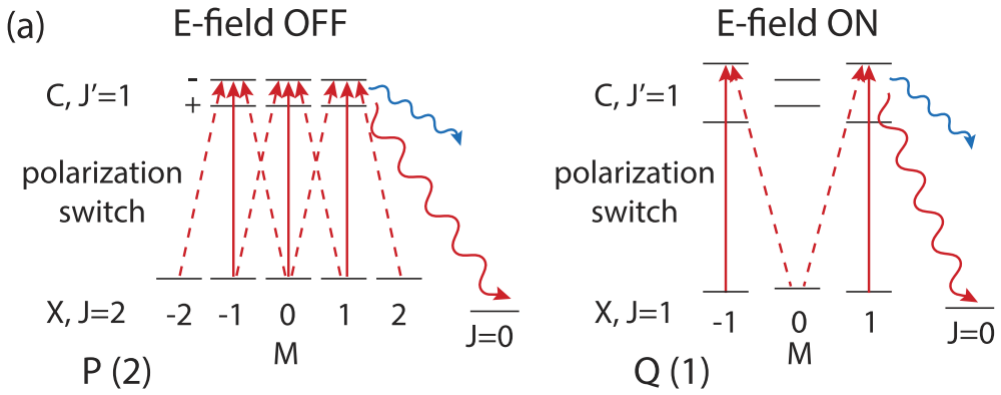

Figure 2.9: Rotational cooling level scheme. Molecules are optically excited from the $X, J = 1, 2$ states to the excited $C, J = 1$ state, which then decays to the $X, J = 0$ ground state with high branching ratio. Polarization switching is performed to destabilize dark states. The frequencies of the rotational cooling lasers are rapidly switched between the P(2) and Q(1) transitions to address both $X, J = 1, 2$ states. The Q(1) branch is parity-forbidden and therefore requires mixing of opposite-parity states in the C state using an electric field switched at the same frequency. Figure reproduced from Wu et al. [103].

2.3.4 Molecular electrostatic lens

After rotational cooling, we need to transfer the $X, J = 0$ molecules to the $Q, J = 2$ state. The $Q, J = 2$ state has a large Stark shift and small parity-doublet splitting, making it ideal for lensing at a modest applied electric field of 10.4 kV/cm. To transfer to the lensing state in $Q, J = 2$, we use stimulated Raman adiabatic passage (STIRAP) to coherently transfer population from the ground $X, J = 0$ state to the lensing $Q, J = 2$ state with near-unity efficiency via the intermediate $C, J = 1$ state [103].

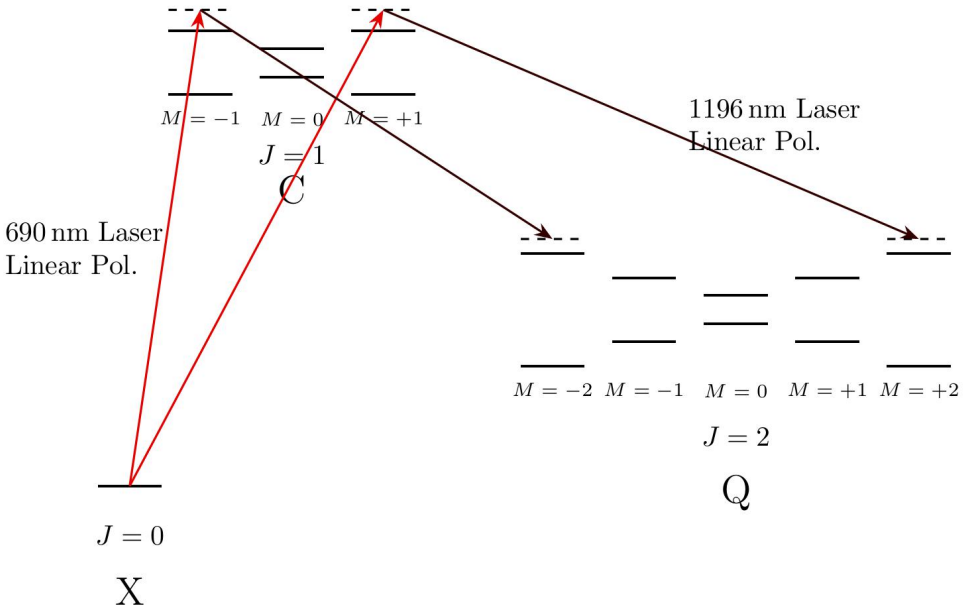

Figure 2.10: Level scheme for transferring molecules from the ground $X, J = 0$ state to the lensing $Q, J = 2$ state using STIRAP via the excited $C, J = 1$ state. The pump laser couples the $X, J = 0 \leftrightarrow C, J = 1$ transition, while the Stokes laser couples the $C, J = 1 \leftrightarrow Q, J = 2$ transition. The 690 nm pump laser and the 1196 nm Stokes laser are partially overlapped in space. The molecules first encounter the Stokes laser, whose intensity adiabatically ramps down while the pump laser intensity ramps up. The molecules remain in the dark state, which evolves from the initial $X, J = 0$ state to the final $Q, J = 2$ state as the molecules fly through the laser beams. Figure prepared by Zhen Han.

After the STIRAP transfer stage, the molecules are collimated in a hexapole electrostatic lens. The electrostatic lens consists of six cylindrical electrodes arranged in a hexapole configuration, as shown in figure 2.11. By applying alternating positive and negative voltages

to adjacent electrodes, we generate an electric field whose strength is proportional to the square of the distance r^2 from the center of the hexapole. The lensing state behaves like a weak-field-seeking classical dipole (due to its linear Stark shift) and therefore experiences a harmonic potential capable of collimating a point source.

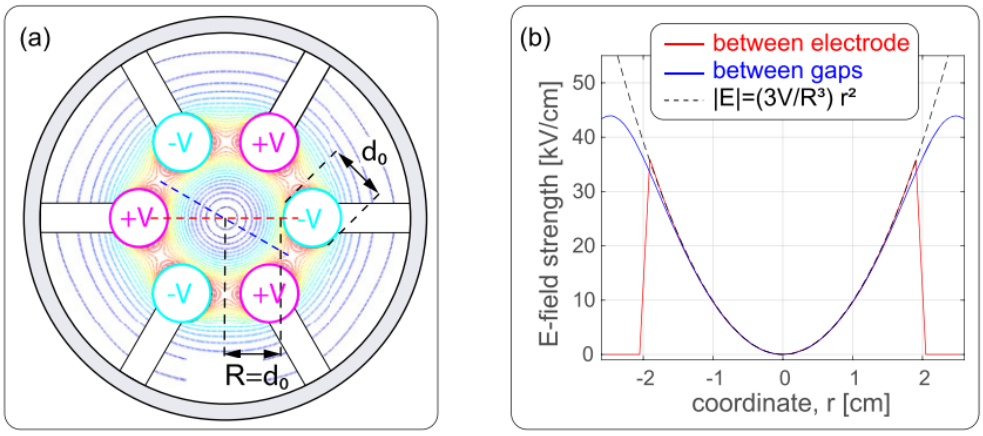

Figure 2.11: Electrostatic potential generated by the electrostatic lens. (a) Cross-sectional view of the hexapole with $R \approx 2$ cm. (b) Electric field, and hence electrostatic potential, seen by the molecules along two different cut lines. Figure reproduced from Wu et al. [103].

Typically, the need for a harmonic potential to collimate a point source is explained mathematically using the thick-lens approximation [91, 103]. For simplicity, we provide a harmonic-oscillator picture here. If the potential is harmonic, the transverse motion of the molecules in the lens can be described as simple harmonic motion. To show that collimation can be achieved, it suffices to show that when the molecules enter the lens, their oscillation phase is identical regardless of their initial transverse position and velocity; then, by choosing the appropriate lens length, all molecules exit the lens with zero transverse velocity, i.e. they are collimated. For a point source, the velocity–position relation is $v_{r,0} \propto r_0$, where r_0 and $v_{r,0}$ are the initial transverse position and velocity at the lens entrance. The maximal transverse displacement inside the lens is then $r_{\max} \propto \sqrt{E_0} \propto v_{r,0} \propto r_0$, where E_0 is the initial transverse kinetic energy. Importantly, since all molecules are traveling radially outward, their oscillation phase upon entering the lens is uniquely determined by $r_0/r_{\max}$,

which is independent of initial conditions. Thus, all molecules execute harmonic motion with the same phase but different amplitudes. By choosing the appropriate passage time, collimation is achieved.

2.3.5 Qleanup stage

After the electrostatic lens, the molecules are transferred back to the ground $X, J = 0$ state using optical pumping via the excited $I, J = 1$ state. This stage is called Qleanup because it cleans up almost all residual population from the lensing $Q, J = 2$ state after the STIRAP stage. Near-complete depletion of this state turns out to be crucial for reducing systematics. The exact level scheme for Qleanup is shown in figure 2.12.

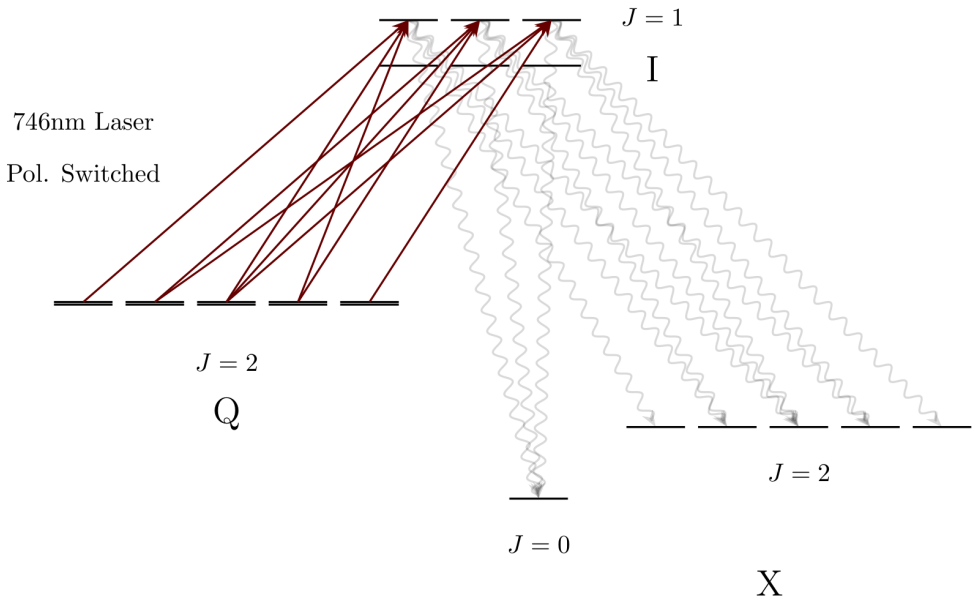

Figure 2.12: Qleanup level scheme. We pump out the $Q, J = 2$ state using optical pumping via the excited $I, J = 1$ state. Polarization switching is used to destabilize dark states. To evade parity selection rules, we apply an electric field at around 100 kHz to mix opposite-parity states in the Q state, which are separated by < 10 kHz. The molecules then decay to the ground $X, J = 0$ state with high branching ratio. Rotational branching to the $J = 2$ state is repumped back to the $J = 0$ state through the C state (not shown). Figure prepared by Zhen Han.

2.3.6 Interaction region for Ramsey measurement

After the Qleanup stage, the molecules enter a magnetically shielded region with residual magnetic fields below $10\,\mu$G after compensation. After an additional $\sim$ 10 cm of flight, the molecules enter a region of uniform electric and magnetic fields for the spin precession measurement. The electric field is generated by a pair of glass field plates with transparent conductive coating[7]. The magnetic field is generated by a set of self-shielded coils.

To prepare the molecules in the $H, J = 1$ science state, we use another STIRAP stage, with laser beams sent vertically through the apparatus to coherently transfer population from the ground $X, J = 0$ state to the $H, J = 1$ state via the excited $C, J = 1$ state. The level scheme for this "vertical STIRAP" stage is shown in figure 2.13. The prepared quantum state is largely set by the polarization of the 1090 nm Stokes beam, with some uncontrolled state-preparation imperfections that can lead to systematics in the EDM measurement.

After the vertical STIRAP stage, molecules next encounter a cleanup stage to remove systematics-inducing state imperfections from the state-transfer process via optical pumping. To maximize signal, the linear polarization direction of the cleanup laser is aligned to maximize overlap of the molecular state with the dark state of the cleanup laser. This stage is also sometimes called the preparation or refinement stage in the ACME literature. The level scheme for this cleanup stage is shown in figure 2.14.

After the cleanup stage, the molecules freely precess for $\approx$1 m, corresponding to about 5 ms of evolution time, before state readout is performed. This is done using 703 nm light to couple the $H, J = 1, M = \pm 1$ states to the excited $I, J = 1, M = 0$ state. The resulting 512 nm fluorescence is collected for state readout. The level scheme for state readout is shown in figure 2.15. The fluorescence is collected by eight large lens doublets, with the photons guided by lightpipes to silicon photomultipliers (SiPMs).

7. ITO coating. The transparency is needed to allow laser beams with polarizations orthogonal to the quantization axis defined by the electric field.

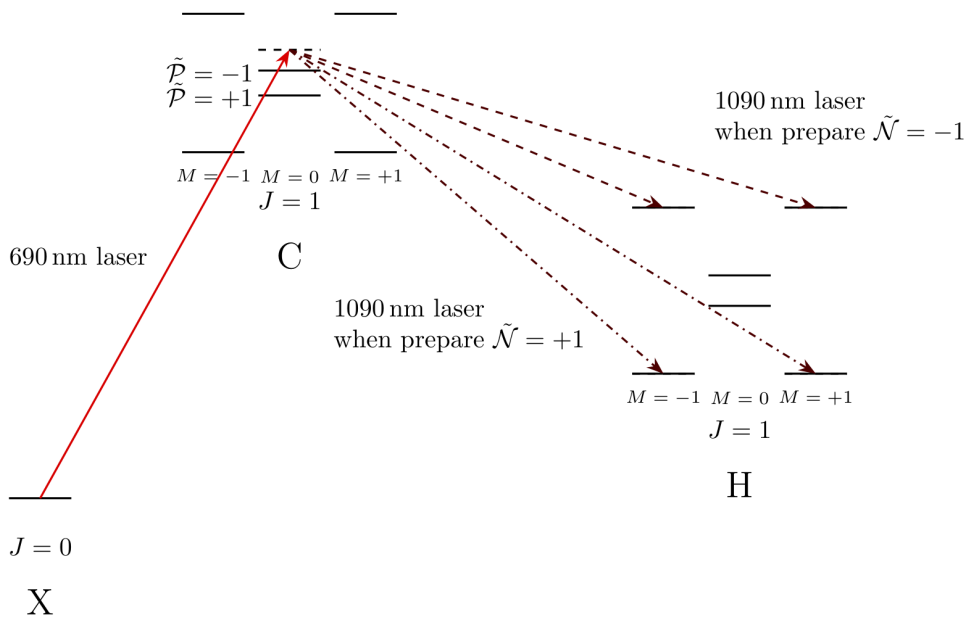

Figure 2.13: Level diagram for the vertical STIRAP stage used to transfer population from the ground $X, J = 0$ state to the science $H, J = 1$ state. The pump laser couples the $X, J = 0 \leftrightarrow C, J = 1$ transition, while the Stokes laser couples the $C, J = 1 \leftrightarrow H, J = 1$ transition. The 1090 nm frequency is changed to switch between different $\mathcal{N}$ states of H. Figure prepared by Zhen Han.

2.4 Improvements over the last generation

Table 2.4 summarizes the main statistical sensitivity improvements in ACME III over ACME II. With the exception of the technical noise improvement, all improvements have been demonstrated [52, 103–107]. To rigorously demonstrate the suppression of technical noise, we need to wait for the final data run to compare excess noise relative to shot noise.

Table 2.4: Sensitivity Improvement

Proposed Improvement	Signal Gain	τ Gain	EDM Sensitivity Gain
Longer precession time	0.3	5	2.7
Electrostatic lens	12	1	3.5
SiPM detector	2.7	1	1.6
Improved collection optics	1.7	1	1.3
Eliminated timing jitter noise	1	1	1.7
Load-lock target change	1.4	1	1.2
Total	23	5	39

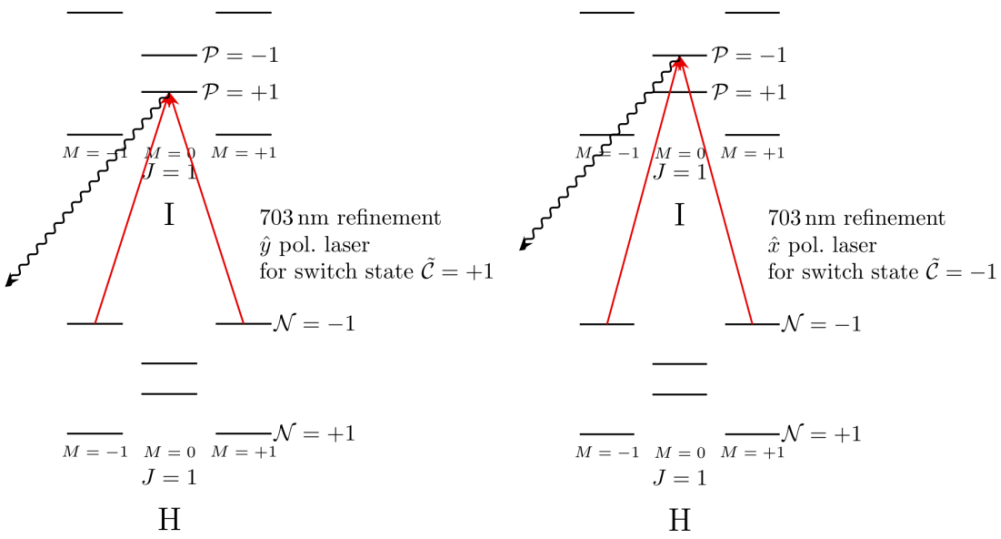

Figure 2.14: Level diagram for the cleanup stage after the vertical STIRAP, used to clean up systematics from the state-transfer process. When the $\mathscr{C}$ switch is changed, both the polarization of the cleanup laser and the parity of the excited $C, J = 1$ state used for cleanup are changed. This effectively flips the ellipticity seen by the molecules while keeping the linear polarization direction unchanged. Figure prepared by Zhen Han.

The longer precession time is made possible by the most recent measurements of the H state lifetime [52], which showed that the H state lifetime is around 4.2 ms instead of the previously believed ~ 1 ms. This allows us to increase the precession time from 1 ms in ACME II to 5 ms in ACME III. We cannot extend the precession time further due to signal loss from the decay of the H state [99].

The electrostatic lens prevents beam divergence loss, and the expected gain from this is around a factor of 12 in signal. The SiPM detectors have higher quantum efficiency compared to the PMTs used in ACME II, leading to a factor of 2.7 gain in signal. The improved collection optics, with larger numerical aperture, also lead to a factor of 1.7 gain in signal.

The ACME II experiment suffered from technical noise due to timing jitter in the data acquisition system [107]. After understanding the mechanism behind this noise, we redesigned the data acquisition system to eliminate this noise source in ACME III. The fix is to synchronize all clocks and ensure that the polarization switching and the data acquisition

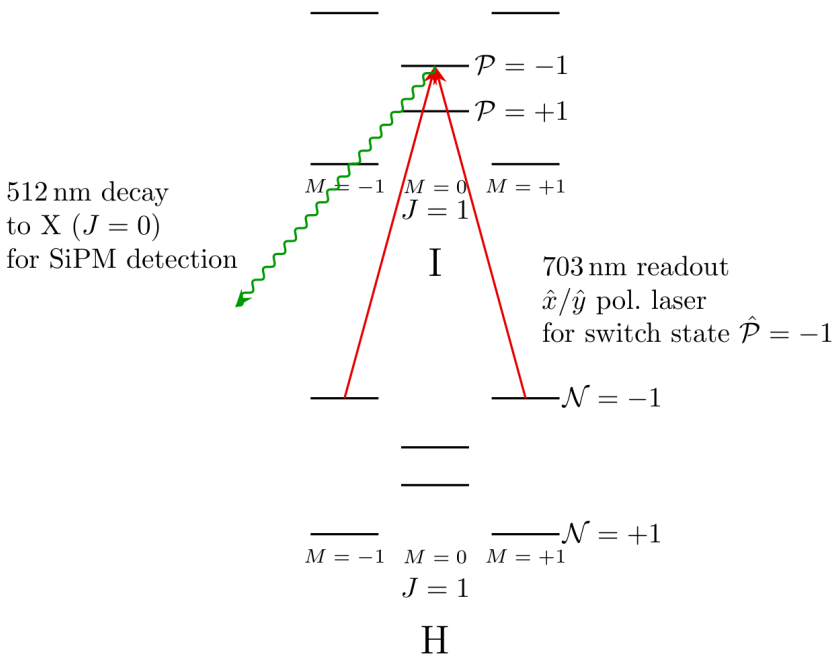

Figure 2.15: Level diagram for state readout. The readout laser couples the $H, J = 1, M = \pm 1$ states to an excited $I, J = 1, M = 0$ state. The resulting 512 nm fluorescence is collected for state readout. The polarization of the readout laser is rapidly switched between x and y linear polarization at around 178 kHz to perform state readout. The $\mathcal{P}$ switch changes the parity of the excited $I, J = 1, M = 0$ state used for readout. Figure prepared by Zhen Han.

are commensurate with the clock rate.

Finally, the load-lock target change system allows us to change the ThO_2 target without breaking vacuum. This allows us to reduce signal loss from target degradation by enabling rapid target swapping [106].

2.5 Hardware details or where to find them

This section provides a more detailed description of the various hardware components that I have worked on. For components that I have not worked on directly, I provide references to the relevant ACME writings where more details can be found.

2.5.1 Beam source

The beam source development was almost entirely led by Zhen Han. Zack Lasner contributed to the design drawings and helped with the initial construction. The 300 K vacuum chamber and the pulse tubes were recycled from the previous generation. The cryogenic 4 K and 40 K radiation shields and the 20 K buffer gas cell were newly designed and built for ACME III. Both the 4 K and 40 K radiation shields are made of gold-plated OFHC copper with a lightweight design to minimize thermal mass. It was found by the Doyle group that 40 K copper (instead of aluminum) radiation shields are better for minimizing cool-down time.

The most important improvement of the beam source is the addition of the vacuum load-lock system for changing the ThO_2 target without breaking vacuum. The load-lock system is described in detail in Han et al. [106].

2.5.2 Compact rotational cooling

The compact rotational cooling stage was designed and built by Xing Wu. After moving ACME from Harvard to Northwestern, Zhen Han constructed all the optics for the rotational cooling stage, and the setup is shown in figure 2.16. Although we initially planned to perform fast polarization switching of the rotational cooling light to destabilize dark states, Zhen Han realized that the polarization-switching Pockels cell did not noticeably increase the rotational cooling gain. This is possibly due to the fact that the laser beams traverse the molecules 13 times in the prism while the excited C state can scatter at most four photons. The birefringence inside the in-vacuum prism may already be sufficient to scramble dark states. Because of this discovery, we moved the Pockels cell elsewhere where it can be used more efficiently.

One possibly useful technical detail is that the in-vacuum prisms are mounted with the vacuum-compatible epoxy OPTOCAST 3553-LVUTF-HM. Xing Wu measured the out-

gassing rate to be $\leq 10^{-12}$ torr L s^{-1} cm^{-2}. I performed a similar measurement with the 3553-250K-HM variant and found a comparable outgassing rate. The part numbers LVUTF and 250K denote low viscosity and 250k cps viscosity, respectively. This suggests that the OPTOCAST 3553 epoxy family, across different viscosities, is vacuum compatible.

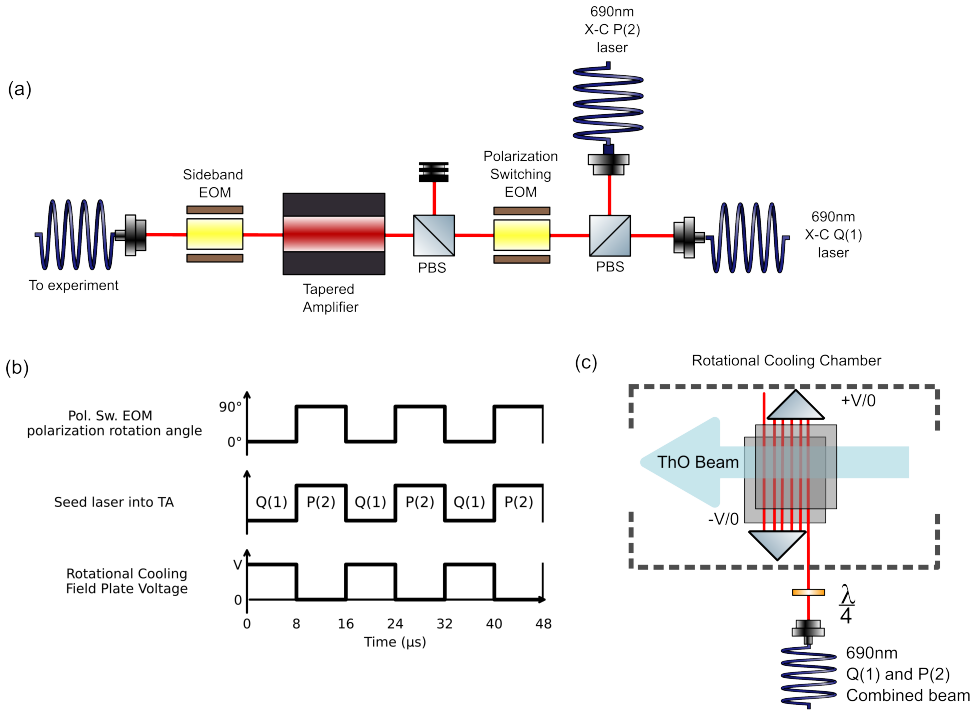

Figure 2.16: (a) Optics setup for rotational cooling. (b) Switching cycle. The polarization EOM controls whether the Q(1) or P(2) light enters the tapered amplifier for rotational cooling. The electric field is switched on and off accordingly to mix opposite-parity states in the C state, enabling excitation of the nominally parity-forbidden Q(1) transition. (c) Actual rotational cooling setup inside the vacuum chamber. Figure prepared by Zhen Han.

2.5.3 Molecular electrostatic lens

The electrostatic lens was designed and built by Xing Wu with some help from me. After moving from Harvard to Northwestern, Xing Wu, Zhen Han, Collin Diver, and I worked together to integrate the electrostatic lens with the new beamline. The $X-C-Q$

STIRAP for transferring to the lensing state and the Qleanup stage for pumping the lensing Q state back to X were built by Zhen Han.

The electrostatic lens design is described in detail in Wu et al. [103]. The previous generation of ACME attempted to build an electrostatic lens using the ground electronic state X for lensing, in an effort led by Adam West. However, due to the small Stark shift in the X state, the electric field required for lensing was too challenging from an engineering standpoint. It was also subsequently realized that the high-voltage lens generated X-rays due to uncontrolled discharges and bremsstrahlung. This is a significant health hazard and was one of the main reasons for abandoning the electrostatic lens in ACME II [108]. Nevertheless, significant progress was made in the last generation toward building an electrostatic lens, and the CeNTREX experiment built upon this work to make a working electrostatic lens without X-rays [57]. In this generation of the experiment, we borrowed heavily from their design in terms of CAD, high-voltage electronics, and manufacturing. The CeNTREX design already follows best practices for high-voltage applications, such as avoiding triple junctions [109]. More details can be found in Oskari Timgren's thesis [57].

After first assembling the electrostatic lens, we performed high-voltage conditioning up to $\pm 30\,\mathrm{kV}$ on each electrode. This is the maximum voltage that our high-voltage supply, a Spellman TOF3000, can provide. We also installed Geiger counters around the lens to monitor possible X-ray generation. No excess counts above background were observed during normal operation. One of the main improvements in ACME III is the use of the $Q, J = 2$ state for lensing instead of the X state. The required electric field for lensing in the Q state is only $14\,\mathrm{kV/cm}$ for the short beamline at Harvard and $10.4\,\mathrm{kV/cm}$ for the longer beamline at Northwestern. This is an exponentially easier problem, especially given our improved understanding of high-voltage engineering.

One major concern with the electrostatic lens is the possibility of overfocusing the molecular beam such that molecules deposit on the surface of the ITO-coated glass field

plates in the interaction region. If a significant number of molecules deposit on the field plates, they can potentially create uncontrolled patch electric fields that lead to systematic errors. To address this concern, we performed detailed simulations of the molecular beam trajectory through the lens using our best knowledge of the phase-space distribution of the molecular beam from the source. The simulation result is shown in figure 2.17. One useful lesson from the simulation is that it is more efficient to add beam collimators before the lens to block highly divergent molecules rather than blocking molecules after the lens. This is likely because highly divergent molecules tend to become overfocused and deposit on the field plates anyway. After optimization of the collimator size and position, we found that the fraction of molecules depositing on the field plates is negligible.

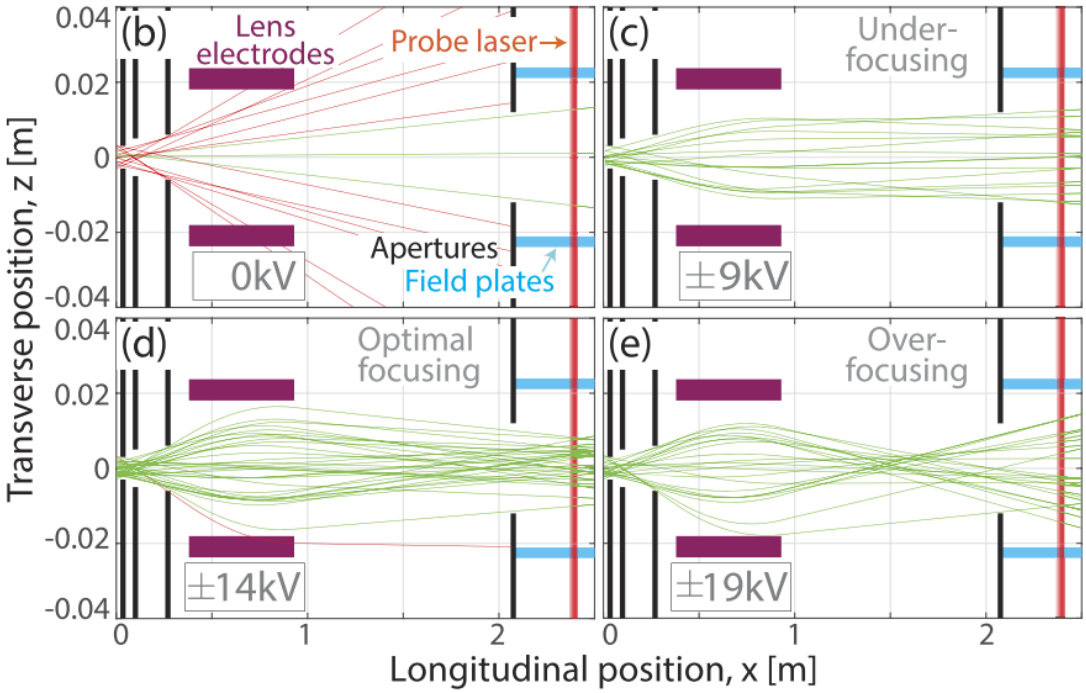

Figure 2.17: Simulation of molecular beam trajectories through the electrostatic lens at various applied voltages. The red lines indicate molecules that do not reach the detection region, while the green lines indicate molecules that successfully reach the detection region. Figure reproduced from Wu et al. [103].

2.5.4 Magnetic shields

The magnetic shields were the biggest engineering challenge in ACME III, requiring roughly two full PhD-years of effort until commissioning. The effort was led by Siyuan Liu, Maya Watts, and Xing Fan. Details of the magnetic shield design and construction can be found in [93]. As of the time of this writing, a paper summarizing all technical details of the shields and magnetic coils is in preparation for publication.

The main challenge with the magnetic shields is their sheer size and weight. The outer-layer shield is more than 2 m long, and there are a total of three layers of mu-metal shields. An additional difficulty arises from the fact that mu-metal is a delicate material, as deformation and mechanical stress can significantly degrade its magnetic permeability. There are good reasons to suspect that the magnetic field performance inside the shields degraded in the last generation of the experiment due to mechanical stress during assembly and disassembly [90, 91]. To minimize deformation-induced stress, the shields were designed to be rectangular rather than cylindrical. Nevertheless, smooth and non-destructive assembly and disassembly of the shields still required significant planning and engineering effort. Xing Fan was the primary driver of this aspect and substantially optimized the assembly and disassembly procedure.

When the shields were first assembled and degaussed, we observed unexpectedly large residual fields of approximately $50\,\mu$G. This was traced to a small DC offset current in the degaussing coils, which magnetized the shields during the degaussing process. After careful investigation led by Xing Fan, this offset was eliminated and the residual fields were reduced to below $10\,\mu$G. Another interesting observation during the degaussing investigation is that the DC shielding factor appears to be significantly higher than the AC shielding factor. In other words, the shields appear to self-magnetize after degaussing in a way that cancels external DC fields more effectively than time-varying fields. We modeled this behavior using standard hysteresis models, with details described in appendix A. Jakob from CeNTREX

independently observed similar behavior and performed related modeling [110].

2.5.5 Magnetic coils

The magnetic field coil design was initially led by Bingjie Hao. After she left the group, Daniel Ang took over the coil simulation work, and Xing Fan led the design and construction of the coils. Ayami Hiramoto and Siyuan Liu also contributed to the construction and characterization of the coils. I performed independent simulations of the field coils for verification and visualization, and I also helped with testing the coils after construction. Details of the magnetic coil design will be included in an upcoming magnetic shield paper that is currently in preparation.

For the purposes of this thesis, it suffices to give the field and field-gradient efficiencies of the coils, as well as their off-diagonal contributions, which were measured using commercial magnetometers[8] set up by Maya Watts and Ayami Hiramoto. These results are summarized in figures 2.18 and 2.19.

2.5.6 Interaction region vacuum chamber

The interaction region vacuum chamber was designed by Cole Meisenhelder and manufactured by Atlas Technologies. Details of the design can be found in Cole's thesis [92]. The initial testing and characterization of the chamber were done by me and Zhen Han.

After baking the chamber at 100°C for a week, the pressure was still surprisingly high, consistent with a leak rate of $10^{-5}\,\mathrm{mbar}\cdot\mathrm{L/s}$, despite passing the helium leak test. The final pressure is below $10^{-7}\,\mathrm{mbar}$ and is sufficient for the experiment, but the high apparent leak rate was puzzling and we decided to investigate further.

We initially suspected that there might be virtual leaks inside the chamber due to trapped volumes. However, John Doyle realized that this leak rate could be due to diffusion

8. QuSpin Total-Field Magnetometer (QTFM). Do not recommend.

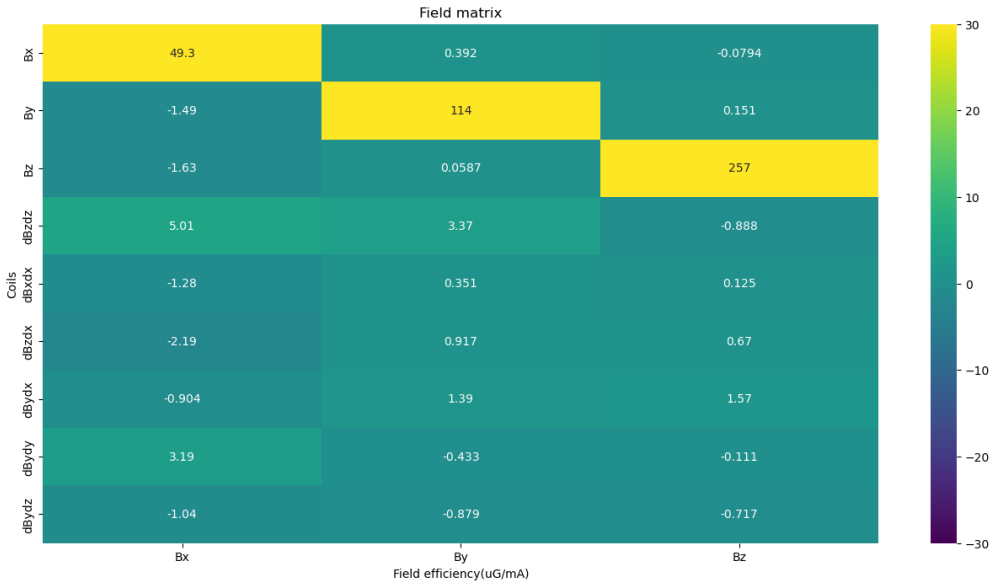

Figure 2.18: Magnetic field efficiency matrix from the coils. Vertical axis: different coil sets/modes. Horizontal axis: generated field direction. Color: coil efficiency in μG/mA. The fact that the off-diagonal elements are small compared to the diagonal elements indicates good orthogonality of the coil design.

of gases through the Viton O-rings used in the chamber. To test this hypothesis, we sealed the chamber in a bag and backed the turbo pump using the leak checker. We then filled the sealed bag with helium gas and observed the leak-checker reading rising very slowly. In fact, even after removal of the surrounding helium gas, the leak-checker reading remained high for hours, indicating that helium had diffused into the Viton O-rings and was slowly diffusing back out. We performed finite-element analysis of the diffusion process using literature values and found good agreement with the observed leak rate and time constant. We were satisfied with this explanation and decided to move on.

2.5.7 Magnetometer pockets

The magnetometer pockets are re-entrant aluminum insets that protrude into the interaction region vacuum chamber to allow insertion of magnetic field sensors close to the molecular beam path. The pockets were designed by Xing Fan and installed by me. During

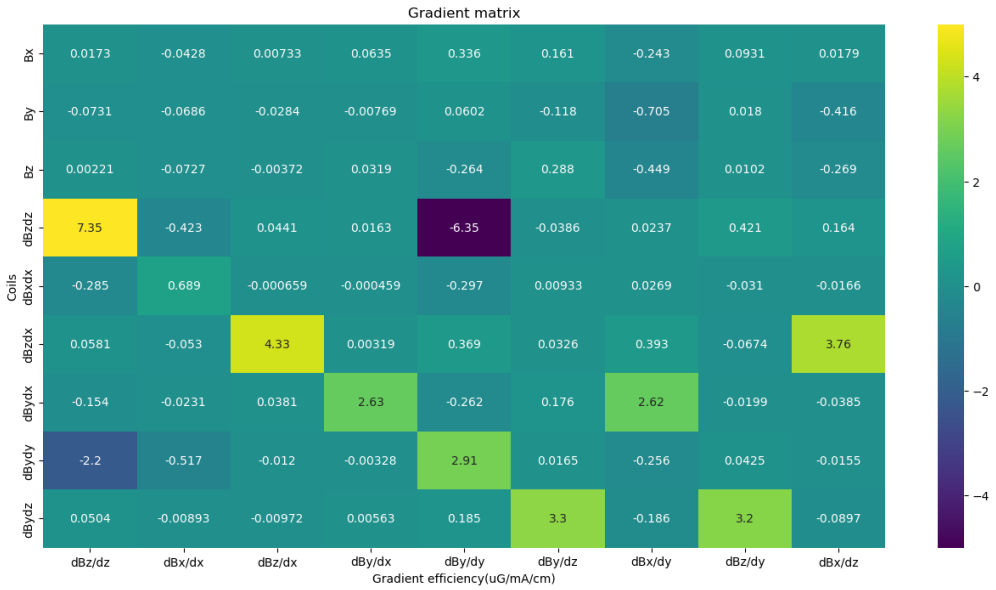

Figure 2.19: Magnetic field gradient efficiency matrix from the coils. Vertical axis: different coil sets/modes. Horizontal axis: generated field-gradient component. Color: coil efficiency in μG/cm/mA. Note that not all off-diagonal elements are small due to constraints imposed by Maxwell's equations.

initial testing, we found that the magnetometer pockets were likely magnetic. Xing Fan cut open one of the pockets and found conclusive evidence that the welding joints were magnetic, as shown in figure 2.20. He then ordered monolithic aluminum pockets without any welding joints for the final installation.

2.5.8 Glass electric field plates

The glass electric field plate system was the most open-ended engineering project of ACME III. No one had a concrete plan that was guaranteed to work. The main challenge was to build a large 1.3 m $\times$ 0.3 m glass plate with a very low birefringence gradient ($\leq 10^{-4}\,\mathrm{mm}^{-1}$) and good electric-field uniformity. The initial investigation was led by Cole Meisenhelder and Daniel Ang. There were two proposals for building a composite field plate that I continued to work on after joining the group. Eventually, both proposals were abandoned due to technical challenges, and we decided to implement a monolithic field plate.

Figure 2.20: Magnetometer pocket with two welding joints. Both were found to be magnetic. Figure prepared by Xing Fan.

Zhen Han built the polarimeter for testing the birefringence [111]. Xing Fan and Collin Diver helped with the mounting design and assembly.

The initial proposal was to use an ultra-low-birefringence glass from Schott called SF57HTULTRA. Roughly speaking, the ultra-low birefringence is achieved by mixing the right amount of lead oxide into the glass. The lead oxide pulls the stress-optic coefficient toward the negative side, canceling the positive stress-optic coefficient of regular glass [112]. The stress-optic coefficient crosses zero around 580 nm and remains small at our operating wavelength of 703 nm[9]. Zhen Han used his polarimeter to measure the birefringence gradient of a small test piece of SF57HTULTRA glass and found it to be $\sim 0.075 \times 10^{-12}\,\mathrm{Pa}^{-1}$, compared to $\sim 3 \times 10^{-12}\,\mathrm{Pa}^{-1}$ for typical glass.

The main challenge with using SF57HTULTRA glass is that it is not manufactured in large sizes. The largest size that Schott can produce is around 20 cm × 20 cm. To build a large field plate, we initially planned to glue multiple glass pieces together to form a composite field plate. One idea was to glue the glass pieces in a rectangular, LEGO-like

9. For optics experts, this is analogous to ULE glass, which has a zero crossing of the coefficient of thermal expansion at a particular temperature. ULE is a glass–ceramic, with the glass phase contributing a positive coefficient of thermal expansion and the ceramic phase contributing a negative one.

pattern. Another idea, proposed by Cole Meisenhelder and John Doyle, was to machine circular holes in a large piece of glass and insert SF57HTULTRA glass pieces into the holes. In this design, the specialized glass would only be used in regions where the laser beams pass through, while the non-laser regions could be made of regular soda-lime glass, which has a very similar coefficient of thermal expansion to SF57HTULTRA.

The composite field plate plan rendered the birefringence problem largely trivial. The key remaining engineering challenge was electric-field uniformity. The glass pieces needed to be chamfered at the edges, which meant that at the glue joints there would be surface flatness issues that could lead to electric-field inhomogeneity. One solution investigated by Daniel Ang was to add a shim electrode at every glue joint, allowing the electric-field uniformity to be tuned using these shim electrodes. We eventually decided that this approach was too complicated from a systematic-error diagnostics perspective and abandoned it. Another possible solution, proposed by Dave DeMille, was to fill the gaps with vacuum-compatible epoxy. We began with Master Bond EP30-2 epoxy, which has been tested by LIGO [113], and we independently verified that it has a low outgassing rate. This epoxy approach improved the joint flatness from $\sim 250\,\mu$m to $\sim 5\,\mu$m.

The main challenge I encountered with this approach was that the epoxy would not remain flat due to surface tension. I experimented with different epoxy viscosities and found that more viscous variants performed better, which is consistent with expectations for dynamic wetting. I also tried adding $\sim 1\,\mu$m copper powder to the epoxy to modify the contact angle. This yielded the best flatness of $\sim 5\,\mu$m at the joint. We also attempted to overfill the joint with epoxy and then grind and polish the surface flat. The main issue with this approach was that the epoxy tended to delaminate from the polished glass surface during grinding. Cracks then rapidly propagated at the delaminated joints, limiting the achievable surface flatness. Development ultimately stalled at around $\sim 5\,\mu$m flatness at the joints.

There were three main concerns with the epoxy-based approach. First, the ITO coating

we planned to use is only about 200 nm thick. If the surface flatness is much larger than this, there may be small patch areas that are not coated with ITO, leading to patch potentials and known systematic effects in the EDM measurement. Second, although we never observed systematic errors associated with electric-field gradients, poor surface flatness could introduce electric-field gradients significantly different from those in previous generations, potentially leading to new systematics. Finally, Jerry Gabrielse frequently reminded me in collaboration meetings that epoxy had caused major problems in his previous experiments. Due to these concerns, I eventually began investigating the use of monolithic field plates.

If we decide to build monolithic field plates, we cannot use SF57HTULTRA glass due to manufacturing size limitations. This means the field plates will have a large stress–optic coefficient, and a serious investigation is required to study the stress mechanics of such a large glass plate. In particular, we need to ensure that mounting stress, residual manufacturing stress, and thermal stress from laser heating do not produce large birefringence gradients.

To control mounting-stress-induced birefringence, I first investigated the mounting designs used in previous generations of ACME. The ACME II mounting design caused the field plates to bend after tightening, as shown in figure 2.21. It was believed that this bending stress was the main source of birefringence in ACME II. I simulated the birefringence from this bending in COMSOL and found that bending alone does not produce birefringence. In fact, this is well known in the glass manufacturing industry and is documented in the literature [114]. The detailed mathematical explanation is derived in Appendix B. The key takeaway is that all birefringence contributions arise from accidentally applied tangential stress normal to the clamping surface. Bending stresses average out to zero through the thickness and do not contribute to birefringence.

To minimize tangential clamping stress, I designed a mounting system using ball joints to ensure that the clamping force is always normal to the glass surface. The idea is that the ball can roll and relax any tangential force component. In reality, as the clamping force is

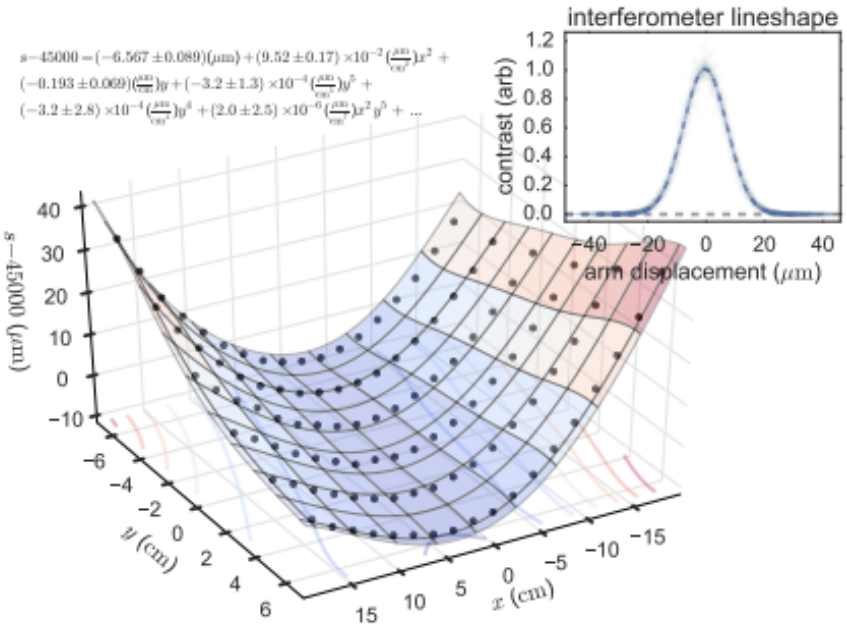

Figure 2.21: Bending of the electric field plate in ACME II due to mounting. Figure reproduced from O'Leary [87].

increased, the ball and/or contact surfaces deform slightly, leading to rolling resistance. This rolling resistance can generate a small tangential force component. In practice, I found that even when the rolling resistance is large enough to hold the glass against gravity using only six clamping points, the induced birefringence gradient remains well below our requirement. The rolling resistance is therefore a feature rather than a bug, as it helps hold the glass in place.

I first demonstrated this approach on the ACME II field plate assembly and then verified it with a test assembly capable of mounting a 1.3 m long soda-lime glass field plate. This gave us confidence that mounting-stress-induced birefringence can be kept under control. The details of the mounting design are shown in Figs. 2.22 and 2.23.

We also need to ensure that the residual manufacturing stress is small enough to avoid producing large birefringence gradients. The difficulty lies in the fact that residual stress is size dependent, and it would be prohibitively expensive, both in cost and time, to purchase a full-size test piece solely to characterize the residual birefringence. We therefore had to study

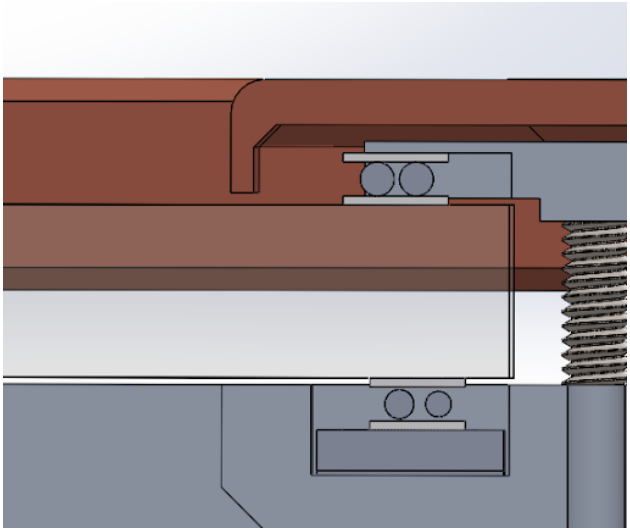

Figure 2.22: Mounting scheme of the field plates using ball joints to ensure a normal clamping force. Titanium spheres are held between two circular titanium disks. Between the titanium disks and the glass surface, a thin Kapton sheet is inserted to provide electrical insulation and to prevent direct glass–metal contact, which could lead to cracking of the glass.

and predict the residual stress indirectly. Fortunately, this turns out to be a manageable task. Residual stress in glass is mainly caused by thermal gradients during the annealing process [115]. It is essentially a differential thermal expansion problem. The key parameters are the cooling rate, thermal expansion coefficient, thermal conductivity, heat capacity, density, stress–optic coefficient, and the geometry of the glass piece. All of these, except the cooling rate, are known for the glass we selected, Corning 7980 fused silica. The cooling rate can be estimated from the annealing schedule, which is on the order of days. Using conservative estimates for the cooling rate, I performed finite-element analysis in COMSOL to simulate the thermal gradients during cooling and the resulting residual stress. To model the cutting process, I set the surfaces of the cut pieces to be traction-free. Multiple cutting geometries were simulated, and in all cases the residual stresses were found to be below our requirement. This gave us confidence that residual manufacturing stress would not be a problem. We also measured the residual stress of a small test piece of Corning 7980 fused silica using Zhen Han's polarimeter and extrapolated to the full-size plate using empirical scaling laws from

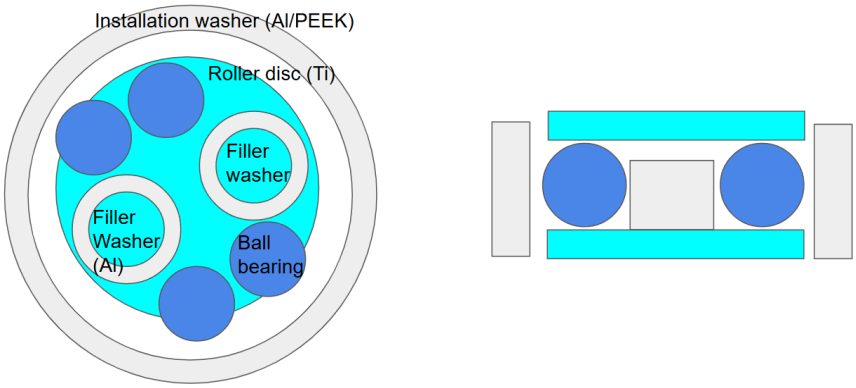

Figure 2.23: Ball bearing assembly for the field plates. The titanium spheres are held between two circular titanium disks. During assembly, an installation washer with a slightly larger inner diameter is used to hold the spheres in place. Its height is smaller than the clamping assembly so that it does not provide any clamping force. Because only 3–4 spheres are needed per assembly, filler washers are used to fill the remaining space and ensure maximal separation of the spheres.

the literature [115]. The agreement between this empirical approach and the simulations provided further confidence. After receiving the full-size fused silica plates, we measured the birefringence gradients and found them to be approximately 1.5×10^{-4}/mm in the regions traversed by the laser beams, which is below the ACME III design target of 1.7×10^{-4}/mm obtained based on scaling from birefringence gradients in ACME I.

Finally, we must ensure that thermal stress from laser heating is small enough to avoid producing large birefringence gradients. In ACME I, the preparation and readout laser operated at 1090 nm rather than 703 nm. The field plates were made of Schott Borofloat 33, which has a relatively large absorption of about 5% at 1090 nm. This resulted in significant thermally induced birefringence gradients and constituted a leading systematic error [53]. We performed polarimetry tests on both fused silica and thin ITO-coated samples to independently verify that absorption in either the glass substrate or the ITO coating is sufficiently small to avoid significant thermally induced birefringence. For a bulk material, the thermally induced birefringence at fixed laser power and wavelength scales as $\propto KE\beta\mu/\kappa$, where K is the stress–optic coefficient, E is Young's modulus, β is the thermal expansion

coefficient, μ is the absorption per unit length, and κ is the thermal conductivity. To measure this coefficient, it suffices to illuminate the material with a heating laser and measure the resulting change in birefringence a few beam waists away from the heating spot. In the absence of cooling, a Gaussian heating beam produces an approximately constant birefringence across the transverse plane of the glass at distances sufficiently far from the beam center. This follows from the exact solution of the biharmonic equation governing Airy stress functions [116, 117]. Therefore, measuring the birefringence a few beam waists away directly probes the material-dependent coefficient $\propto KE\beta\mu/\kappa$ relevant for thermal stress-induced birefringence. This method is fast and straightforward, and our measurements showed that thermally induced birefringence is not a concern for ACME III.

After we concluded that the stress-induced birefringence was under control, we proceeded to the manufacturing of the full-size field plates. This was difficult mainly because it was hard to find vendors willing to handle glass pieces of this size. The glass was manufactured by Corning using material 7980STD/0C. The dimensions are 1.3 m by 0.3 m by 15 mm. The subsequent grinding and polishing were performed by Displays & Optical Technologies, Inc., which was the best vendor I have ever worked with in terms of communication, quality, and timeliness. The final glass was specified to have a global surface flatness of 5 μm. After polishing, the glass plates were sent to Thin Film Devices Inc. for ITO coating. Although the company initially promised to coat the pieces at their California facility, they ultimately shipped the glass to their South Korea facility without notifying us. The surface resistance was also approximately 200 $\Omega/\square$ instead of the initially promised 2000 $\Omega/\square$. Nevertheless, after measuring uncontrolled electric fields between the field plates, we did not observe any significant difference compared to previous generations.

After receiving the coated glass plates, we proceeded with assembly. Initially, we observed significant bending of the field plates during mounting. Although bending itself does not cause birefringence, we were concerned that it could introduce electric field inhomogene-

ity. We determined that the bending was likely caused by deformation of the aluminum frame. The machining process had relieved a significant amount of internal stress in the aluminum, leading to distortion. We communicated this issue to the machining company, 3ERP, who proposed an improved process. Their solution was to perform a rough machining step, anneal the frame, and then carry out final precision machining[10]. This approach significantly reduced the bending. We then performed white-light interferometry measurements of the surface flatness of the assembled field plates [87], finding a flatness of approximately $100\,\mu$m. Because this was not fully satisfactory, we shimmed the frame to further reduce bending. The shimming process was lengthy and labor-intensive, as assembly and disassembly were complicated by the ball-bearing mounting design. I received substantial help from Collin Diver and Xing Fan during this process. After multiple iterations of shimming and measurement, we achieved a surface flatness of approximately $40\,\mu$m, matching the global flatness achieved in ACME II despite the much larger plate size. We measured the birefringence gradient after final assembly and confirmed that it was between 0.8×10^{-4}/mm to 1.5×10^{-4}/mm, which is below the ACME III design target of 1.7×10^{-4}/mm.

After finalizing the field plate assembly, we installed it into the interaction region chamber. A photograph of the completed assembly is shown in figure 2.24. After installation, we repeated both polarimetry and flatness measurements for confirmation. The polarimetry measurements verified that the birefringence gradient remained below the requirement. The locations used for the flatness measurements are shown in figure 2.25. The maximum separation deviation between the measured points is approximately $20\,\mu$m, confirming that the flatness of the field plate assembly was not significantly altered during installation.

10. The initial frame was made of 6061 aluminum, which has substantial residual stress. T651 aluminum is designed to have roughly ten times lower residual stress. We purchased a piece of T651 and measured its residual stress, finding results consistent with the literature. However, because we sought an improvement beyond a factor of ten, we adopted the solution proposed by 3ERP instead.

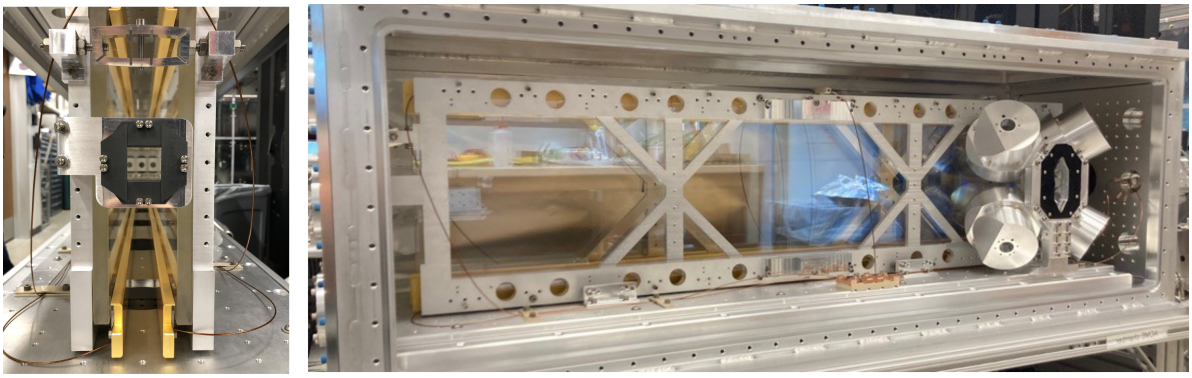

Figure 2.24: (a) Field plate assembly viewed from the beam box side. The collimating black razor blades are ceramic to prevent residual magnetism and are coated with graphite to prevent charge buildup. The gold-coated copper guard rails block dielectrics from direct line of sight of the molecular beam. By applying a bias voltage to them, we can also reduce electric field inhomogeneity. (b) Field plate assembly viewed from the side. One of the plates is fixed to the breadboard, while the other plate sits on two sapphire spheres for alignment. The kinematic design follows the ACME I design [56].

2.5.9 Low birefringence vacuum window

In addition to the field plates, we also needed to build low-birefringence vacuum windows for the interaction region chamber. Initially there were plans to make in-situ rotatable vacuum windows so that we could rotate the window birefringence axis to align with the laser polarization. This effort was led by Cole Meisenhelder and Zhen Han. Eventually, we realized that a regular vacuum window design with O-rings would be sufficient as long as we use the ultra-low-birefringence glass SF57HTULTRA. There were some initial concerns about vacuum-pressure-induced birefringence, but recall that bending induced by vacuum cannot cause birefringence on its own, as shown in appendix B. The dominant source of birefringence comes from the O-ring contact stress, which may contain tangential components.

We made a small test vacuum chamber with regular glass to show that even with regular glass, vacuum only induced a birefringence gradient of around $1.5{\cdot}10^{-4}$/mm, which is already below the ACME III design target of $1.7{\cdot}10^{-4}$/mm. If we use the SF57HTULTRA glass, the birefringence gradient induced by vacuum should be even smaller. This was confirmed by

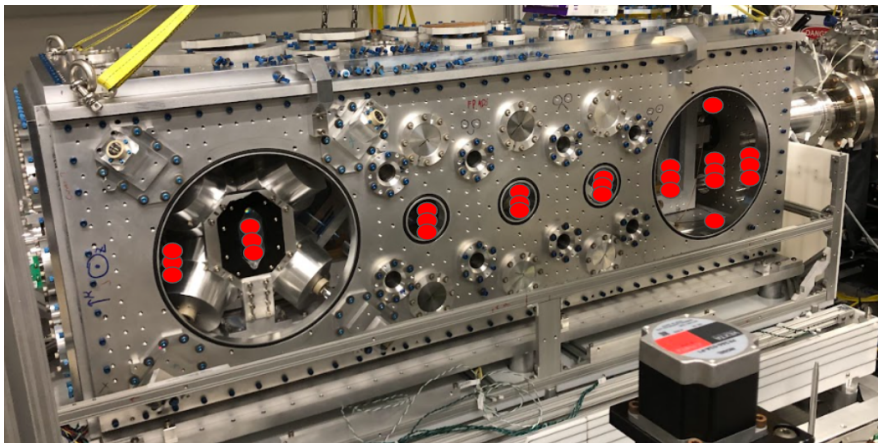

Figure 2.25: Red: points used for flatness measurement. The maximum separation deviation between the points is approximately 20 μm.

Takahiko Masuda at Okayama University using a separately built polarimeter. Interestingly, in the test setup with regular glass, vacuum induced around 2–3% birefringence on the glass, with the fast axis oriented about 20° away from the rectangular glass symmetry axis. We think this is due to randomness in the O-ring contact stress and not due to vacuum-induced bending, whose stress pattern would be symmetric.

The final design of the vacuum window was done by Cole Meisenhelder [92]. The assembly was done by me and Ayami Hiramoto.

2.5.10 Interaction region optics

There are three main optical subsystems in the interaction region: the vertical STIRAP optics, the cleanup optics, and the readout optics. The vertical STIRAP optics and the supporting structure were almost entirely built by Collin Diver. The optics used for the cleanup stage were built by me with significant input from Collin Diver and Zhen Han. The readout optics were built by Zhen Han and me.

The setup of the cleanup optics is shown in figure 2.26. The main purpose of the cleanup optics is to clean up systematics from the vertical STIRAP stage by optical pumping. To maximize signal, the linear polarization direction of the cleanup laser is aligned with the

molecule to maximize the overlap of the molecular state with the dark state of the cleanup laser. This is done using the transmitted beam by rotating the last Glan–Laser polarizer right before the vacuum chamber. Once this is done, the orientation of the polarizers is fixed, and we align the reflected beam to be colinear with the transmitted beam at the second PBS. By moving the motorized beam blocker to block either the reflected or transmitted beam, we obtain pure horizontal or vertical linear polarization output from the cleanup optics.

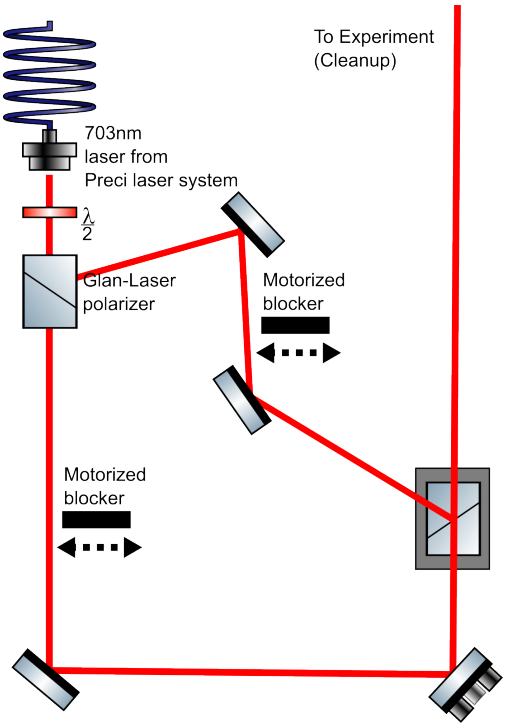

Figure 2.26: Cleanup optics setup. The reflected and transmitted beams of the first PBS are recombined to copropagate at the second PBS, forming the identity operator for laser polarization. By using the motorized beam blocker to block either the reflected or transmitted beam, we obtain a projection operator that projects the output fiber polarization into either horizontal or vertical linear polarization. When the $\mathscr{C}$ switch is changed, the frequency of the laser is also changed to address the opposite-parity $C, J = 1$ state. This frequency change is implemented using two AOMs combined with another PBS on a seperate breadboard.

We initially planned to perform the polarization rotation of the $\mathscr{C}$ switch with a half-wave plate. However, the birefringence gradients from the half-wave plate caused significant systematics, and we had to abandon this plan. We then tried rotating the PBS cube directly by 90°. This worked suboptimally. The main issue was that the PBS had limited parallelism

tolerance. When we rotated the PBS by 90°, there was a significant pointing change in the output laser beam, on the order of a millimeter at the molecule location, which is 1 m away. This resulted in imperfect cancellation of systematics from the $\mathscr{C}$ switch in many scenarios, and we decided to abandon this method as well. The final configuration shown in figure 2.26 no longer allows us to freely rotate the orientation of the cleanup polarization but provided the best $\mathscr{C}$ switch suppression, and we ultimately adopted it.

For the readout optics, we also went through several versions before settling on the final design shown in figure 2.29.

The first version is shown in figure 2.27 and is almost identical to the ACME II readout setup. The key issue with this setup is that the laser beams had to have a small waist to achieve sufficient AOM diffraction efficiency. This resulted in significant beam-shape differences between the X and Y readout beams because they have different path lengths. The beams also suffered from aberrations immediately after the fiber asphere and were not very Gaussian. To address this issue, Zhen Han proposed a second version of the readout optics shown in figure 2.28. The key change is that the AOM polarization switching is done before the fiber. The X and Y polarization beams are then coupled into two separate fibers. This removes any requirement on the beam waist after the fiber and resulted in much better beam quality. However, the two different fibers led to intensity drifts between the X and Y beams due to temperature changes, despite the fact that the fibers were nominally polarization maintaining. This intensity drift led to noise in the asymmetry channel, and we eventually decided to address the beam quality issue by using different AOMs.

The final version of the readout optics is shown in figure 2.29 and was built by Zhen Han. New AOMs with a larger active aperture were purchased to allow for a larger beam waist and improved beam quality. The AOMs were also water-cooled to improve thermal stability. The entire setup was enclosed in a box to prevent room light from entering the interaction chamber. Zhen Han further added water cooling to the breadboard, which should

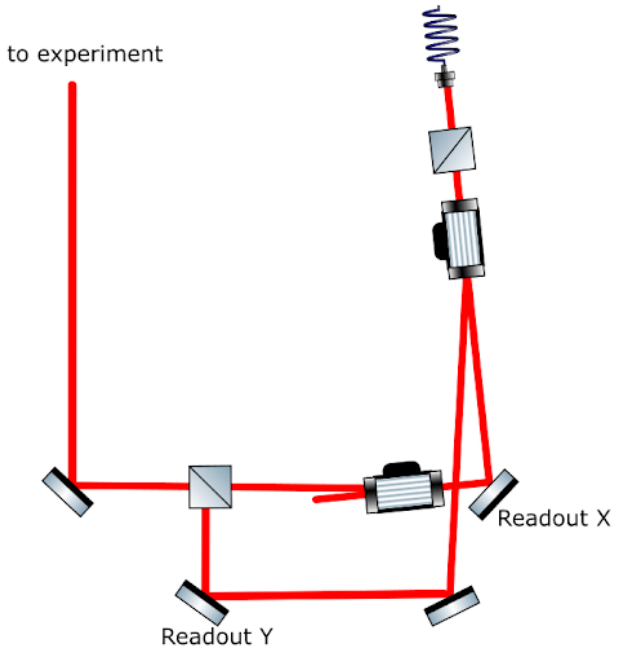

Figure 2.27: First version of the readout optics, almost identical to the ACME II readout setup. Other optics are removed for easy comparison with the second version shown in figure 2.28.

stabilize the index of refraction of the air inside the box. Two sets of quadrant photodiodes (QPDs) were also added to monitor the intensity and pointing stability of the readout beams. The QPD signals can also be used for active feedback if needed using commercial solutions from Thorlabs. With all these improvements, the final readout optics provided good beam quality, stability, and low noise.

To accurately align the readout X-polarization beam with the Y-polarization beam, we send them 3–4 m downstream to maximize spatial overlap. To ensure that they propagate along the same wavevector $\hat{k}$, we observe the interference fringes formed by the two beams. An additional PBS oriented at 45° with respect to both polarizations is added to allow interference between the two orthogonal polarizations. Polarization switching is also turned off so that the beams overlap in time. We additionally introduce a 1 Hz frequency difference between the X and Y polarization AOMs. This creates a 1 Hz differential phase shift between the two beams, and any misalignment in $\hat{k}$ shows up as moving interference fringes.

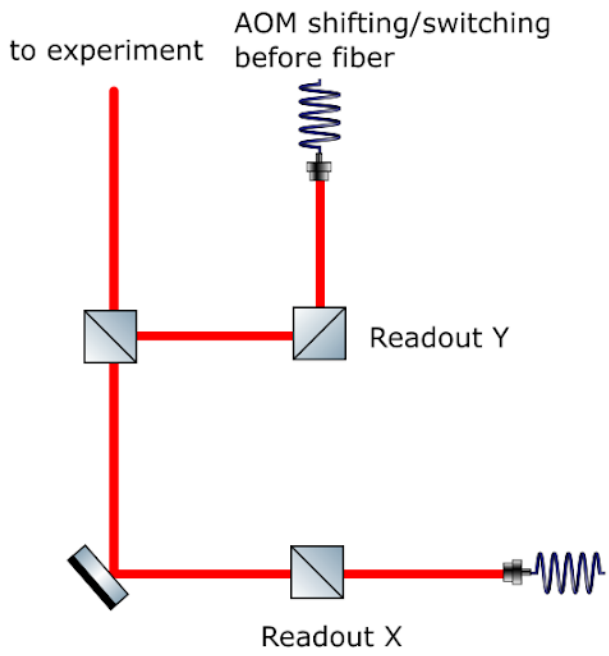

Figure 2.28: Second version of the readout optics proposed by Zhen Han. The main improvement is that both the X and Y readout beams have a larger waist and more similar beam sizes. Other optics are removed for easy comparison with the first version shown in figure 2.27.

By adjusting the mirrors to minimize fringe motion on a distant screen, we ensure that the two beams are well aligned in $\hat{k}$.

2.5.11 Light collection system with SiPM

The new light collection lenses were designed by Daniel Ang [99]. The design choice of using straight lightpipes instead of bent lightpipes is explained in Cole Meisenhelder's thesis [92]. The actual assembly in the experiment was done by Ayami Hiramoto, Xing Fan, Zhen Han, and me. The SiPM detectors were designed, tested, and integrated by Takahiko Masuda and Ayami Hiramoto, with assistance from Cole Meisenhelder and Daniel Ang [104, 105, 118].

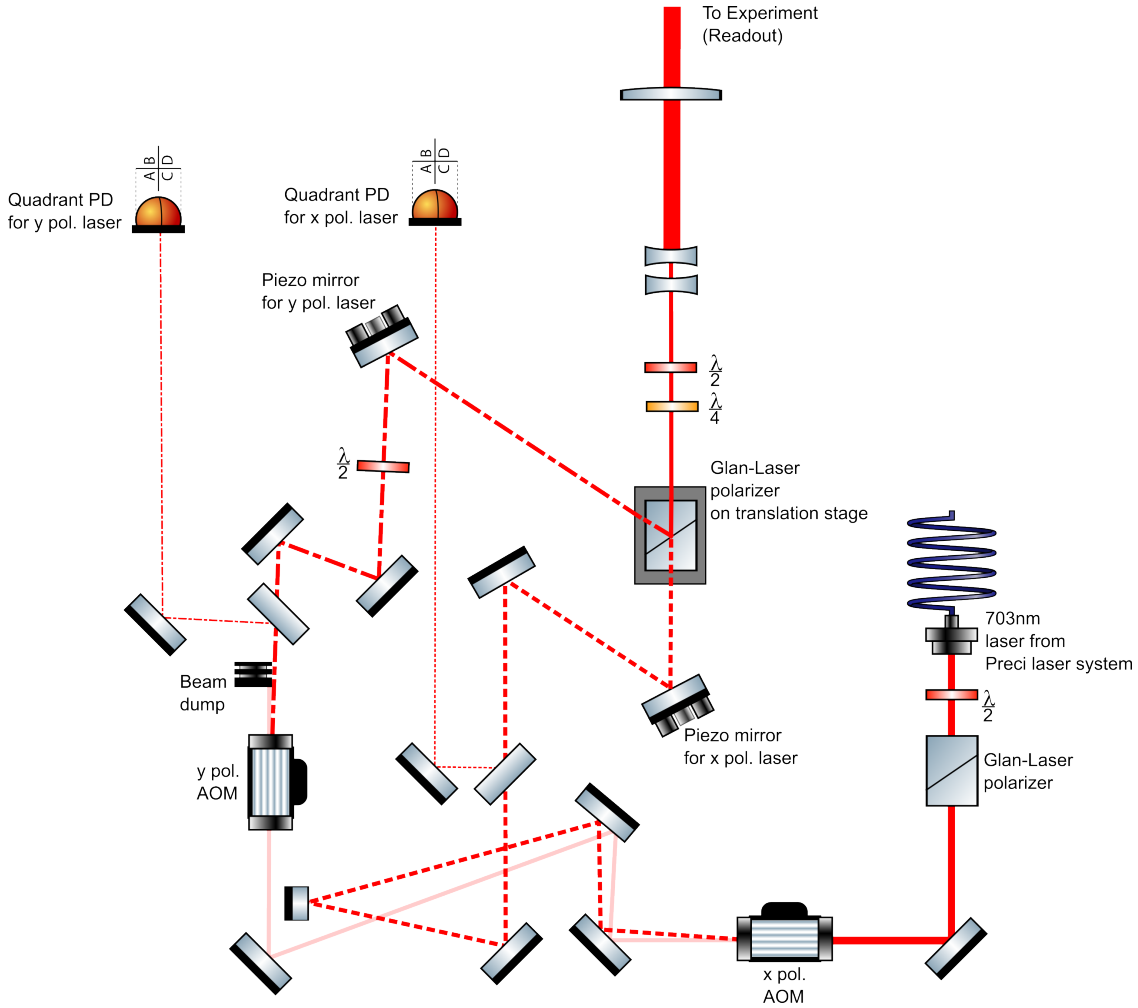

Figure 2.29: Final version of the readout optics currently used in the experiment. A triplet lens expansion system is used to make the final beam size continuously tunable. Figure prepared by Zhen Han.

2.5.12 Data acquisition system

The data acquisition system was designed and assembled by Daniel Ang [99]. Integration of the data acquisition system with the rest of the experiment was done by Zhen Han, Collin Diver, and me. The main improvement of the data acquisition system is higher data throughput. In the previous generation of the experiment, we only saved the signal from each trace by summing every 25 shots. It was later realized that it is potentially useful for debugging purposes to examine data from individual shots. The new data acquisition system

is capable of saving data from each individual shot at the full experimental repetition rate of 50 Hz. This allows more detailed analysis of the data and facilitates identification of sources of excess noise. For most data runs, we still save the summed data to reduce data size, but for the final data run we will fully utilize this shot-saving capability.

CHAPTER 3

DATA ANALYSIS

The ACME-style data analysis theoretical framework has been studied in great detail by Zack Lasner[91]. The hardware and technical aspects of the data analysis implementation were explained in Cris Panda's thesis [90]. Lessons learned from ACME II data analysis and additional technical changes have been summarized by Cole Meisenhelder and Daniel Ang [92, 99]. In this thesis, I mainly focus on aspects not covered in these previous works.

3.1 Pipeline to extract phase and precession frequency

We record the molecular fluorescence data from the eight SiPMs using an 8-channel FPGA digitizer (NI PXIe-5171). It is triggered by the YAG laser at 50 Hz. For each trigger, we record 14 ms of data every 20 ms. Digitized data from each channel contain 175,000 samples with 14-bit resolution on a 5 V_{pp} scale. Each sample has a duration of 80 ns. The X and Y readout laser beams are alternately switched on and off every 35 samples, i.e., 2.8 μs. This corresponds to a polarization switching frequency of approximately 178 kHz[1], rather than the originally planned 250 kHz, based on lessons learned during systematic studies. The TTL signal used to switch the laser polarization is generated by a high-precision delay generator (SRS DG645), which is phase-locked to the 10 MHz rubidium frequency standard (SRS FS725). The FPGA digitizer is also phase-locked to the same rubidium frequency standard.

The data throughput from the FPGA digitizer is near the limit of a 10 Gb fiber connection. This means that reading all saved data takes approximately as much time as acquiring it. Therefore, we also save summed trace data—defined as the sum of every 25 shots—into a 32 bit integer binary file. This summed data is sufficient for most analyses, and we some-

1. See 4.4.5 for details on the choice of this frequency.

times disable shot-by-shot saving to reduce data volume. At the time of writing, we have accumulated roughly 240 TB of data; had every shot been saved, the total data size would have exceeded 1 PB. For the remainder of this thesis, I focus exclusively on the analysis of the trace data.

Each binary file consists of one block of data containing 16 traces with distinct $\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\theta}, \tilde{\mathcal{B}}$ states, as explained in subsection 2.2.3. After reading the binary file, we obtain a $16 \times 8 \times 175000$ array of 32-bit integers. The first dimension corresponds to the 16 switch states, the second to the 8 SiPM channels, and the third to the time samples. This array is subsequently reshaped into two $16{\times}8{\times}2500{\times}35$ arrays corresponding to S_X and S_Y fluorescence. The axis with 2500 samples corresponds to time within a shot, while the 35-sample axis corresponds to time within a polarization bin. The average signal on these time scales is shown in figure 3.1.

Next, we sum the fluorescence signal within each polarization bin to obtain the total number of photoelectrons collected per bin. This yields two $16{\times}8{\times}2500$ arrays corresponding to X and Y fluorescence. We then perform background subtraction by subtracting the mean of the first 100 bins of each trace.[2] These first 100 bins contain molecules from the previous shot or trace. To improve estimation of the true electronic background, we added an RF switch to turn off both X and Y polarizations during the first 4 ms of collected data.

After background subtraction, we convert the 32-bit integers to photoelectrons by dividing by the single-photon response of each SiPM channel. We then apply a signal cut that discards data points with signal below 10% of the maximum value. This ensures that the denominator in the asymmetry calculation is not too small, which would otherwise lead to

2. ACME II performed independent background subtraction for the X and Y polarizations because scattered readout light was detected by the photon collection system. The new SiPM collection system preferentially detects 512 nm fluorescence, and we have not observed evidence of 703 nm scattering. Nevertheless, we follow the same background subtraction procedure.

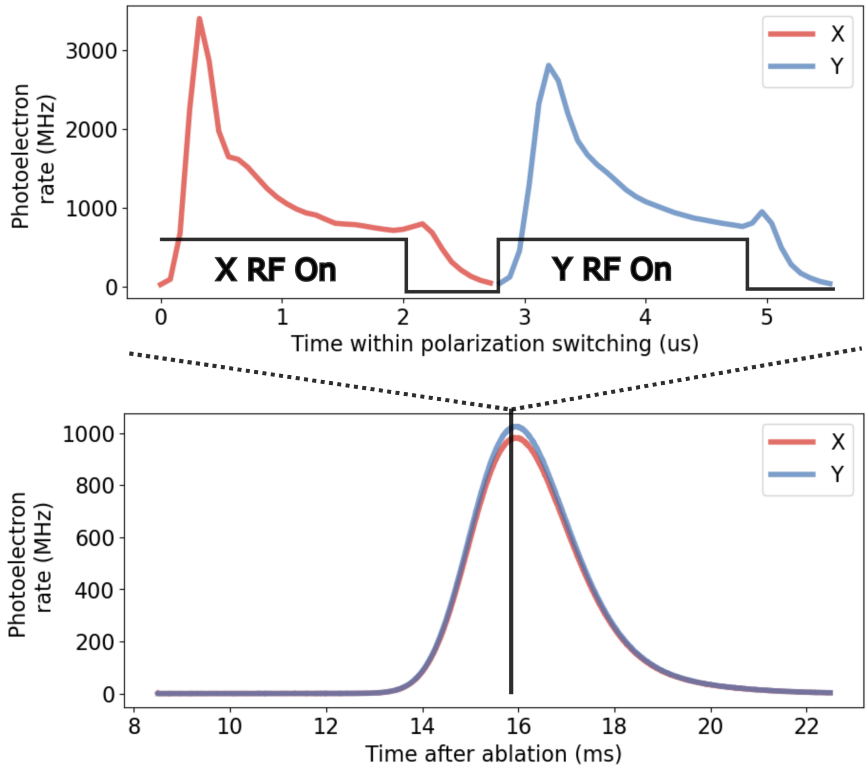

Figure 3.1: Fluorescence signal within a shot. The vertical axis is the instantaneous photoelectron rate, such that the area under the curve corresponds to the total number of photoelectrons collected. Top: time within a polarization switching cycle, which is 5.6 μs or 70 samples. When the laser light is turned on, the fluorescence starts to increase and reaches a maximum after roughly one Rabi oscillation period. The fluorescence then decays to a steady state whose level depends on newly arriving molecules entering the laser beam. After the laser light is turned off, the fluorescence decays at approximately the excited-state lifetime. Bottom: fluorescence signal within a shot, which is 14 ms or 175,000 samples. Each shot contains 2,500 polarization switching cycles.

non-Gaussian statistics [91]. We then compute the asymmetry and total signal:

$$\mathcal{A} = \frac{S_X - S_Y}{S_X + S_Y}, \tag{3.1}$$

$$S = S_X + S_Y. \tag{3.2}$$

We now have two $16 \times 8 \times 2500$ arrays corresponding to the asymmetry and total signal. Points removed by the signal cut are marked as NaN in both arrays. We then group the 2500

polarization-bin samples into groups of 20, forming 125 time bins. A linear fit is performed on the asymmetry versus time data to extract the mean asymmetry and the reduced chi-squared, which is used to check for possible excess noise in the asymmetry channel, as shown in figure 3.2. A linear fit is used instead of a constant fit because the asymmetry inherently has a time-dependent mean due to effects such as correlations between molecule arrival time and velocity, as well as the deterministic shape of the fluorescence signal within a shot. Typically, the reduced chi-squared from this procedure is around 1.15, which is not significantly higher than the shot-noise limit, and we do not investigate the origin of this small excess noise further.

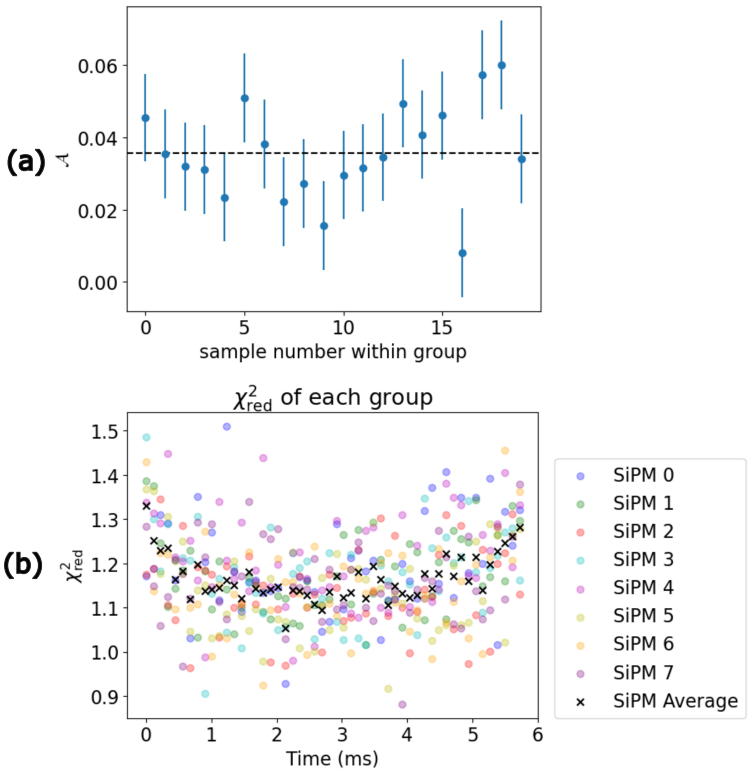

Figure 3.2: (a) Distribution of measured asymmetry within a group. (b) Distribution of reduced chi-squared values from different SiPM channels and the average reduced chi-squared within the signal-cut region. The average reduced chi-squared is around 1.15, slightly higher than the SiPM-limited noise floor given by the Fano factor of approximately 1.07 [104].

Next, we use the asymmetry data, now a $16 \times 8 \times 125$ array, to compute the contrast and

phase. Let ϕ denote the molecular precession phase, and let $\theta \pm \frac{\pi}{4}$ be the readout polarization angles corresponding to X and Y polarization. The measured asymmetry can be written as

$$\mathcal{A} = \mathcal{C} \sin[2(\phi - \theta)] \,. \tag{3.3}$$

We define $\mathcal{A}_+$ and $\mathcal{A}_-$ as two $16 \times 8 \times 125$ arrays, differing only in that $\mathcal{A}_+$ corresponds to $\tilde{\theta} = +1$ and $\mathcal{A}_-$ corresponds to $\tilde{\theta} = -1$. The θ-averaged asymmetry and contrast are then given by

$$\mathcal{A}_0 = \frac{\mathcal{A}_+ + \mathcal{A}_-}{2}, \tag{3.4}$$

$$\mathcal{C}_0 = -\frac{180^\circ}{2\pi\Delta\theta} \frac{\mathcal{A}_+ - \mathcal{A}_-}{2}, \tag{3.5}$$

where $\Delta\theta = 2^\circ$ is the total dithering angle. The contrast defined in this way can be either positive or negative, depending on which side of the Ramsey fringe the measurement is performed. This procedure is somewhat unique to ACME, as we never measure the full Ramsey fringe; instead, we operate on the region of the fringe with the largest slope to maximize phase sensitivity. Using the small-angle approximation, the precession phase can be extracted as

$$\phi \approx \frac{\mathcal{A}_0}{2\mathcal{C}_0}. \tag{3.6}$$

The validity of this small-angle approximation is discussed in detail in Lasner [91]. The primary concern is potential systematic error arising from imperfect contrast measurement. In ACME nomenclature, this class of systematics is referred to as "correlated contrast" and is discussed extensively in previous theses [90, 91, 99].

After extracting the phase, we can compute the precession frequency by examining the $\tilde{\mathcal{B}}$-correlated phase shift. The phase array is now an $8 \times 8 \times 125$ array. For each of the 125 time bins, we compute the difference between the mean phase at $\tilde{\mathcal{B}} = +1$ and $\tilde{\mathcal{B}} = -1$. This

difference, divided by twice the precession time, gives a one-dimensional array of 125 Zeeman precession phases. Using the applied magnetic field and the known magnetic moment of the H state, we convert this into a one-dimensional array of 125 precession times τ. As expected for a beam whose forward velocity is correlated with arrival time, the precession time τ is correlated with time within a trace, with faster molecules arriving earlier.

With the one-dimensional arrays of precession phase ϕ and precession time τ, we compute the precession frequency as

$$\omega = \frac{\phi}{\tau}, \tag{3.7}$$

which yields an $8 \times 8 \times 125$ array, just like ϕ. This completes the data analysis for each binary file, which contains a block of 16 traces. To extract meaningful eEDM information, we collect all blocks within a superblock (16 blocks from the $\mathcal{P}, \mathcal{R}, \mathcal{L}, \mathscr{C}$ switches). These are concatenated to form $128 \times 8 \times 125$ arrays of asymmetry, contrast, phase, and precession frequency. As explained in Sec. 2.2.3, these quantities are initially represented in the state basis. For physical interpretation, we transform them into the parity basis using Eq. 2.104. This produces another set of $128 \times 8 \times 125$ arrays, now expressed in the parity basis. The eEDM signal, for example, is contained in the $\mathcal{NE}$ channel of the precession-frequency array. This describes nearly all steps of the data analysis, except for two subtle but important procedures that are intentionally omitted here for simplicity. These are discussed in Secs. 3.2 and 3.3.

The contrast $\mathcal{C}_0$, as defined so far, corresponds to the slope of the Ramsey fringe and can take either positive or negative values. This is often inconvenient for data visualization. We therefore define the unsigned contrast and contrast-sign-adjusted asymmetry as

$$\mathcal{C} = |\mathcal{C}_0|, \tag{3.8}$$

$$\mathcal{A} = \mathrm{sgn}(\mathcal{C}_0)\, \mathcal{A}_0, \tag{3.9}$$

which we refer to simply as the contrast and asymmetry for the remainder of the thesis. When needed, $\mathcal{C}_0$ and $\mathcal{A}_0$ are referred to as the signed contrast and unadjusted asymmetry, respectively. Typically, the non-reversing contrast $\mathcal{C}^{\mathrm{nr}}$ is around 0.95 in ACME III, and the signed contrast $\mathcal{C}_0^{\mathcal{PR}}$ therefore also has a magnitude of approximately 0.95. The signed contrast predominantly appears in the $\mathcal{PR}$ channel because the $\mathcal{P}$ and $\mathcal{R}$ switches effectively swap the roles of the X and Y polarizations, thereby flipping the slope of the Ramsey fringe. This makes the $\mathcal{PR}$ parity channel particularly important for data interpretation.

3.2 Blinding the eEDM

3.2.1 The blinding procedure

The purpose of blinding is to prevent experimenter bias from affecting the final result. From discussions with other members of the EDM community, a common trend is that PIs tend to hope for a non-zero result, while students tend to hope for a null result. To prevent such bias from influencing the analysis, we blind the eEDM channel with a large offset. The final statistical uncertainty of the last ACME II data run was approximately 400 μrad/s. We therefore chose to blind the eEDM channel with a value drawn from a gaussian distribution with a standard deviation of 1 mrad/s. Given our expected final statistical sensitivity of ~ 10 μrad/s, this constitutes a comfortably large blind.

ACME II implemented blinding by directly adding the blind in the parity basis,

$$\omega^{\mathcal{NE}}_{\text{blinded}} = \omega^{\mathcal{NE}} + \omega_{\text{blind}}, \tag{3.10}$$

$$\phi^{\mathcal{NE}}_{\text{blinded}} = \phi^{\mathcal{NE}} + \omega_{\text{blind}}\,\tau, \tag{3.11}$$

$$\mathcal{A}^{\mathcal{NE}}_{\text{blinded}} = \mathcal{A}^{\mathcal{NE}} + 2\mathcal{C}\,\omega_{\text{blind}}\,\tau, \tag{3.12}$$

$$\omega^{p\neq\mathcal{NE}}_{\text{blinded}} = \omega^{p\neq\mathcal{NE}}, \tag{3.13}$$

$$\phi^{p\neq\mathcal{NE}}_{\text{blinded}} = \phi^{p\neq\mathcal{NE}}, \tag{3.14}$$

$$\mathcal{A}^{p\neq\mathcal{NE}}_{\text{blinded}} = \mathcal{A}^{p\neq\mathcal{NE}}, \tag{3.15}$$

where non-eEDM parity channels are left unblinded. While conceptually simple, this procedure is not strictly self-consistent. In practice, fluctuations in contrast and precession time lead to mixing between parity channels. As a result, adding a blind only to the $\mathcal{NE}$ channel inevitably leaks some of the blind into other parity channels. This is not a serious issue provided one is aware of it and avoids overly careful comparisons [91]. However, we find it equally simple to implement a fully self-consistent blinding scheme. The solution is to blind the data in the state basis rather than in the parity basis:

$$\omega_{\text{blinded,state}}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, \ldots) = \omega_{\text{state}} + \omega_{\text{blind}}\,\tilde{\mathcal{N}}\tilde{\mathcal{E}}, \tag{3.16}$$

$$\phi_{\text{blinded,state}}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, \ldots) = \phi_{\text{state}} + \omega_{\text{blind}}\,\tau\,\tilde{\mathcal{N}}\tilde{\mathcal{E}}, \tag{3.17}$$

$$\mathcal{A}_{\text{blinded,state}}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, \ldots) = \mathcal{A}_{\text{state}} + 2\mathcal{C}_{\text{state}}\,\omega_{\text{blind}}\,\tau\,\tilde{\mathcal{N}}\tilde{\mathcal{E}}. \tag{3.18}$$

This procedure is conceptually simple: it amounts to adding a $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-correlated offset directly to the state-basis observables. It is fully self-consistent because any mixing between parity channels arising from variations in contrast or precession time is automatically included. We do not view parity-channel mixing as problematic, since a true eEDM signal would mix channels in precisely the same way. This consideration does, however, place an upper bound

on the size of the blind: an excessively large blind could introduce so much mixing into other parity channels that data interpretation becomes difficult. Given the observed level of channel mixing in our system, a blind of 1 mrad/s is sufficiently small to avoid such issues.

In practice, this blinding procedure is implemented immediately after extracting the precession frequency from each block of 16 traces.

3.2.2 Blinding information leaking

In the previous generation of the experiment, the blind value happened to be around 2σ away from zero. This raised concern among some members of the collaboration, because such a value would be implausibly large for the true eEDM given previous experimental limits[3]. This implies that information about the value of the blind was leaked after inspecting the blinded results. We argue here that although information about the blind value itself may be leaked, no meaningful information about the true eEDM channel is revealed, which is the primary purpose of blinding.

Let us define the following random variables:

1. E: the true EDM value. Its distribution is subjective and reflects prior belief about the real EDM value. Let its probability density function (pdf) be $f(e)$, which may be informed by previous experiments or theoretical expectations (e.g. zero in the Standard Model).

2. B: the blind value. Its distribution is chosen by us, with pdf $g(b)$.

3. $Y \equiv E + B$: the blinded EDM value visible to the experimenters. Since E and B are independent, its pdf is given by $\int f(y-b)g(b)\,db$.

To quantify how much information is revealed by observing the blinded EDM, we con-

3. 2σ of the blind distribution typically corresponds to $> 4\sigma$ relative to previous measurements.

sider the conditional pdf of E given Y,

$$h_{E|Y}(e|y) = \frac{\mathrm{pdf}_{E,Y}(e,y)}{\mathrm{pdf}_Y(y)} = \frac{f(e)\,g(y-e)}{\int f(y-b)g(b)\,db}. \tag{3.19}$$

As an illustrative example, suppose both E and B follow Gaussian distributions,

$$f(e) = \frac{1}{\sqrt{2\pi}\,\sigma_e}\exp\left(-\frac{e^2}{2\sigma_e^2}\right), \tag{3.20}$$

$$g(b) = \frac{1}{\sqrt{2\pi}\,\sigma_b}\exp\left(-\frac{b^2}{2\sigma_b^2}\right). \tag{3.21}$$

In this case, the conditional distribution can be evaluated analytically,

$$h_{E|Y}(e|y) = \sqrt{\frac{1}{2\pi}\left(\frac{1}{\sigma_e^2}+\frac{1}{\sigma_b^2}\right)}\exp\left[-\frac{\left(y\sigma_e^2 - e(\sigma_b^2+\sigma_e^2)\right)^2}{2\sigma_b^2\sigma_e^2(\sigma_b^2+\sigma_e^2)}\right]. \tag{3.22}$$

This conditional pdf is Gaussian with mean

$$\mu_{E|Y} = \frac{y\,\sigma_e^2}{\sigma_b^2+\sigma_e^2}$$

and standard deviation

$$\sigma_{E|Y} = \sqrt{\frac{\sigma_b^2\sigma_e^2}{\sigma_b^2+\sigma_e^2}}.$$

Our blinding procedure is designed to operate in the regime $\sigma_b \gg \sigma_e$. In this limit,

$$\sigma_{E|Y} = \sqrt{\frac{\sigma_b^2\sigma_e^2}{\sigma_b^2+\sigma_e^2}} \approx \sigma_e. \tag{3.23}$$

This is the key result. As long as the blind width is much larger than the expected spread of the true EDM, inspection of the blinded EDM value does not allow us to update either our estimate or uncertainty of the true EDM. Nearly all information revealed by inspecting

the blinded result pertains to the blind value b, rather than to the true EDM value e.

3.3 Arrival time center of mass shifts

We first explain what we mean by arrival time center-of-mass (COM) shifts. Over timescales of minutes, the cryogenic buffer-gas beam is not a stable source in terms of its kinematic properties. As a result, the arrival-time distribution of the molecules at the detection region slowly drifts over time. In my analysis, I shift the center of mass of the arrival-time distribution of each block (roughly 10–20 s of data) so that each block has approximately the same arrival-time center of mass. This procedure is illustrated in figure 3.3. The shift is applied immediately after grouping the polarization-bin data, as described in Sec. 3.1. This procedure is unique to my analysis and is not performed by other members of the collaboration.

To motivate this procedure, we briefly review how data from different measurements are combined. Suppose we have two sets of measurements, $\mathcal{A}_1$ and $\mathcal{A}_2$, with corresponding signals S_1 and S_2. They are combined as

$$\mathcal{A}_{\text{combined}} = \frac{S_1\mathcal{A}_1 + S_2\mathcal{A}_2}{S_1 + S_2} \tag{3.24}$$

$$= \frac{(X_1+Y_1)\frac{X_1-Y_1}{X_1+Y_1} + (X_2+Y_2)\frac{X_2-Y_2}{X_2+Y_2}}{X_1+Y_1+X_2+Y_2} \tag{3.25}$$

$$= \frac{(X_1+X_2)-(Y_1+Y_2)}{(X_1+X_2)+(Y_1+Y_2)}, \tag{3.26}$$

$$S_{\text{combined}} = S_1 + S_2, \tag{3.27}$$

where X_i and Y_i are the X- and Y-polarization signals for measurement set i. This procedure is equivalent to directly summing the signals from different measurements separately for the X and Y polarizations.

If the two data sets have the same excess noise, this method is also equivalent to a

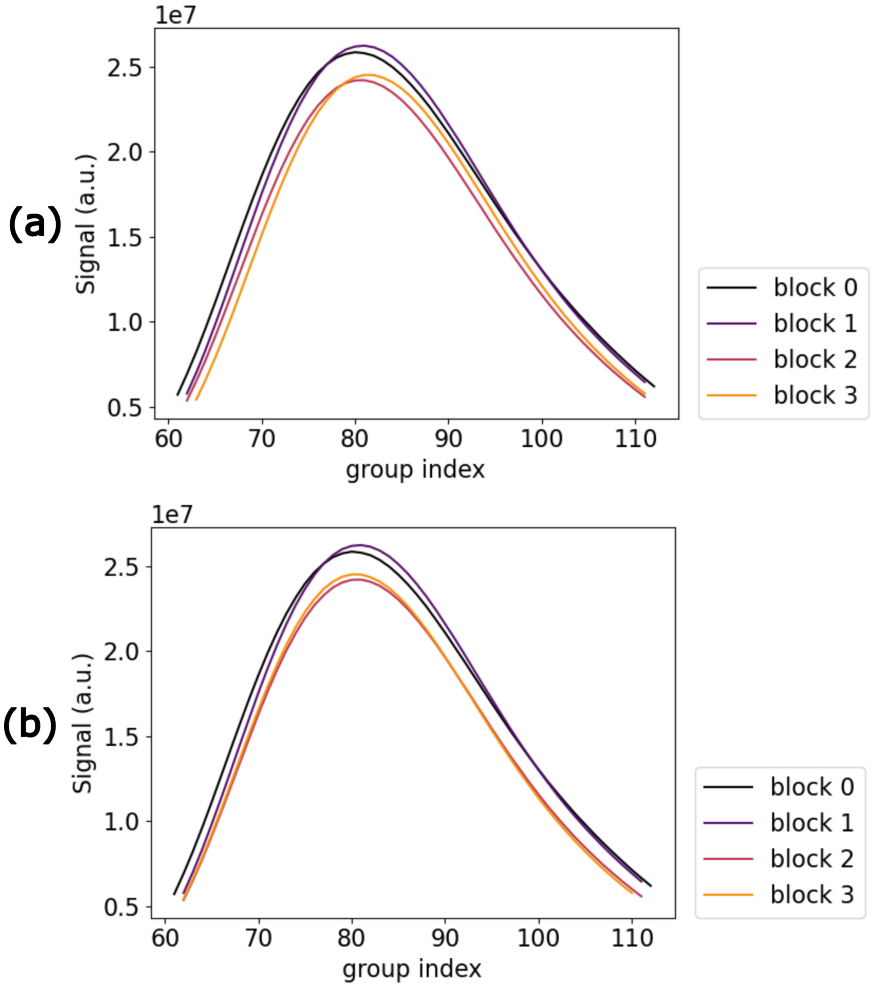

Figure 3.3: (a) Total molecular signal within a trace, with time labeled by group index, before center-of-mass shifting. (b) After center-of-mass shifting.

weighted average based on error estimates. Except for a few special cases, I do not combine data using rescaled errors; instead, all weighting is based on shot-noise uncertainty. A key advantage of this approach is that it is automatically commutative. As described in Sec. 3.1, quantities such as ϕ and ω are stored as higher-dimensional arrays. This combining procedure guarantees that the final result is independent of the order in which different axes are averaged, providing substantial conceptual simplicity[4].

4. Strictly speaking, this is not exact because of NaN entries in the arrays. In practice, summation is performed using the `nansum` function in NumPy. This technical detail is the primary motivation for applying the center-of-mass shift.

The parity transformation used to extract parity channels can be viewed as a generalized unweighted mean. The non-reversing channel (nr) is literally the unweighted mean over all switch states, while the other parity channels differ from unweighted means only by sign flips. As a result, the transformed quantities are not, in general, shot-noise limited compared to results obtained using an optimally weighted mean. This unweighted approach is unfortunately unavoidable, because introducing weights into the parity transformation leads to systematic errors [53].

To minimize the excess noise introduced by the parity transformation, it is advantageous for all switch states to have the same, or at least similar, signal levels. The arrival-time center-of-mass shifting procedure achieves exactly this. In most cases, this procedure does not significantly reduce the noise. However, there are situations—for example, when the ablation spot on the target is adjusted during a superblock—where the molecular arrival-time distribution can shift by as much as $\sim 20\%$ of its width. In such cases, applying the center-of-mass shift is especially beneficial.

An alternative approach adopted by other members of the ACME III collaboration is to first combine all group data before performing any further analysis. After this step, there are no remaining NaN values, and all subsequent averaging becomes fully commutative. This approach is also closer to the data analysis methods used in previous generations of the experiment. Its disadvantage is that it removes access to time-dependent effects within the molecular time-of-flight distribution. For systematic error studies, I find this time-resolved information to be valuable, and therefore I choose to retain the group axis until the final stage of the analysis.

Examples of the final $\phi^{\mathcal{NE}}$ and $\phi^{\mathcal{B}}$ as functions of arrival time after center-of-mass shifting are shown in Figs. 3.4 and 3.5, respectively. From the $\phi^{\mathcal{B}}$ versus arrival-time plot, for example, we observe that earlier-arriving molecules exhibit smaller Zeeman precession phases, consistent with faster molecules spending less time in the magnetic field.

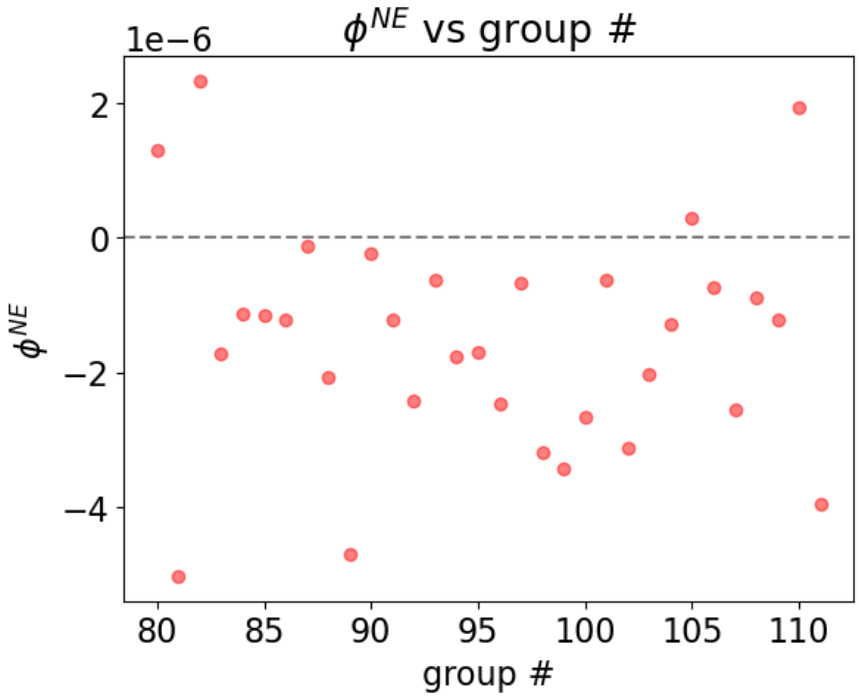

Figure 3.4: Example of $\phi^{\mathcal{NE}}$ versus arrival time (equivalently, group index). Data are taken from a systematic-study run, as the final data run has not yet begun.

In previous generations of the experiment, final analysis results from different analysis codes were combined to produce a single final result. Because the approach presented here is more liberal and may be more susceptible to false positives in systematic-error identification, the current plan is to use this analysis primarily for systematic studies, while relying on the more conservative analyses from other members of the collaboration for the final eEDM result. This is not a concern, as we already have sufficient independent analysis codes from the University of Chicago, Northwestern University, and Okayama University.

3.4 Simulated data

To validate the data analysis pipeline, we wrote a data simulator to generate fake data with parity-channel means set by us. The simulator is written in Python and is capable of generating data with statistical properties very similar to those of the real data. In

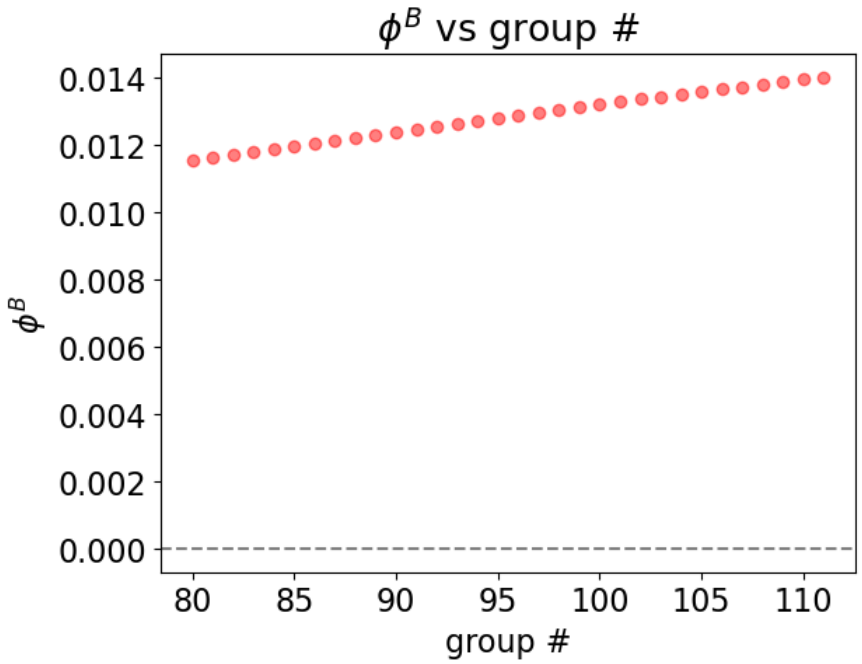

Figure 3.5: Example of $\phi^{\mathcal{B}}$ versus arrival time (equivalently, group index). Data are taken from a systematic-study run, as the final data run has not yet begun.

particular, molecular kinematic properties are drawn from realistic distributions, which then fully determine the precession times and phases. The fluorescence signal is drawn from a Poisson distribution with a mean determined by the realistic molecule number and detection efficiency.

One main difference between the simulated data and the real data is that Poisson sampling is performed only at the polarization-bin level in the simulated data. This means each bin contains a single independent sample instead of 35 samples. This choice is made to reduce computation time, as each Poisson sampling at our data volume scales as $O(1)$ [119]. Different SiPM channels receive independent Poisson samples, and the results are saved to a binary file in the same format as the real data after conversion using the single-photon response. The simulated data are then analyzed using exactly the same analysis pipeline as the real data.

We initially used the simulated data to cross-check the analysis pipeline without applying

blinding. During this process, we identified a systematic effect in the asymmetry, which was traced to the shape of the signal envelope within a trace. Let $S_{X,i}$ and $S_{Y,i}$ denote the X- and Y-polarization signals within polarization bin i. Because of the polarization-switching scheme, $S_{Y,i}$ is always measured after $S_{X,i}$. As a result, $S_{Y,i}$ tends to be smaller or larger than $S_{X,i}$ depending on whether the signal envelope is decreasing or increasing. This effect is also observed in the real data. We believe it is not problematic because it does not carry a $\mathcal{NE}$ correlation. In addition, its net contribution averages out provided that the signal levels at the beginning and end of the signal-cut region are similar.

We also used the simulated data to study potential noise sources in the eEDM channel. In particular, we suspected that electronic ground noise might contribute significantly after initial data taking began. To test this, we recorded electronic noise data without producing molecules and added this noise to the simulated data. The results showed that electronic ground noise can indeed contribute to the total noise. Efforts to reduce this noise were then undertaken by Ayami Hiramoto, Takahiko Masuda, and Zhen Han. After mitigation, we again recorded electronic noise data and added it to the simulated data. The results are shown in figure 3.6. For simulated signal levels of 3×10^6 counts per shot, electronic noise is not a significant contributor to the eEDM channel. However, when the signal level drops below $\sim 8 \times 10^5$ counts per shot, the noise contribution becomes significant and increases nonlinearly with decreasing signal. This behavior is consistent with observations from the real data, where substantial excess noise appears at low signal levels.

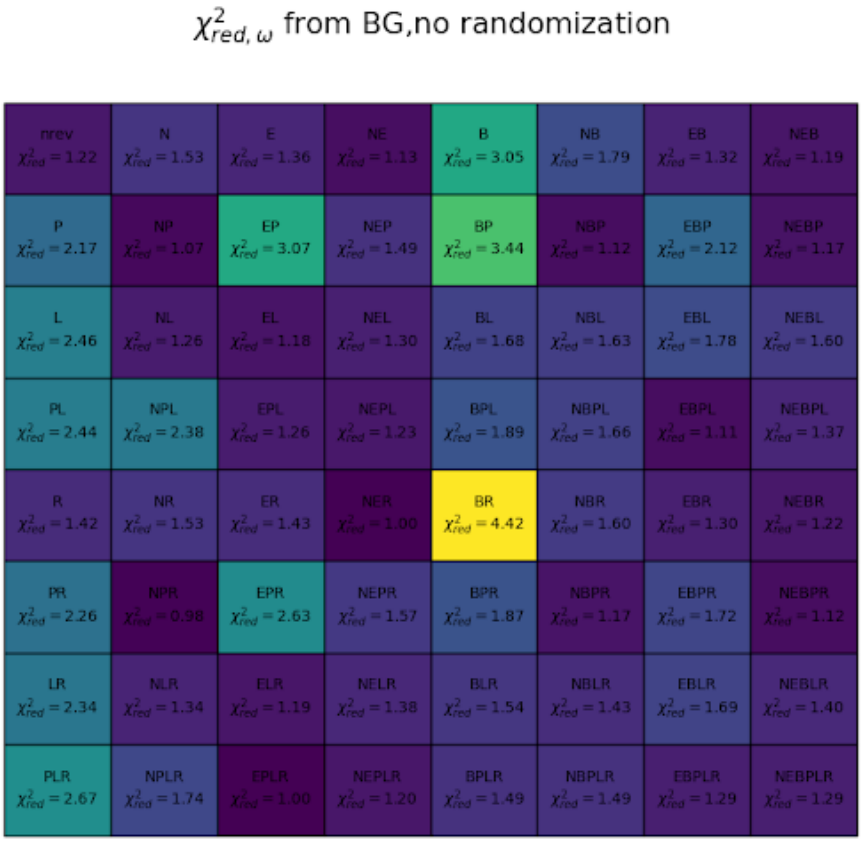

Figure 3.6: Electronic noise measured experimentally was added to the simulated data to examine noise propagation into each parity channel. For the simulated signal levels shown here, electronic noise is not a dominant contribution to the eEDM channel. At lower signal levels, the simulations reproduce the excess noise increase observed in real data. "No randomization" indicates that the electronic noise is added state by state without shuffling.

CHAPTER 4

SYSTEMATICS AND EXCESS NOISE

4.1 Search philosophy

The rigorous search for systematic errors is always a difficult task in precision measurement experiments. Historical results from other experiments show that systematic errors are often underestimated. The most problematic systematics typically arise from incomplete understanding of the apparatus, which is difficult to predict or model *a priori*.

In general, the search for systematic errors can be viewed as a Bayesian inference problem. We begin with prior beliefs about the properties of the measurement apparatus and perform a wide range of measurements while varying experimental parameters or control knobs. If the measurement results change in a manner inconsistent with these prior beliefs, we must update our understanding of the apparatus by investigating possible causes of the unexpected behavior. This process is iterated until a satisfactory understanding is achieved and a reasonable bound on systematic errors can be assigned.

In an ideal experiment, this procedure would continue until no hidden information about the apparatus remains. The resulting systematic uncertainty would then represent a true bound on possible systematic effects and would follow Gaussian statistics. In practice, however, precision measurements are constrained by practical considerations such as engineering limitations and the finite lifetime of graduate students. Consequently, the final assigned systematic uncertainty also reflects the judgement, experience, and risk tolerance of the experimenters, and need not follow Gaussian statistics. For this reason, it is important to document the entire systematic search process in detail so that future experimenters can independently assess the methodology and conclusions.

In ACME III, we varied approximately 60 experimental control parameters in the search for systematic effects. The eEDM signal resides in the $\mathcal{NE}$-correlated channel. We therefore

carefully monitor nearby channels that differ by one symmetry label, such as $\mathcal{N}$, $\mathcal{E}$, and $\mathcal{NEC}$, as well as the $\mathcal{NEPR}$ channel. Although the latter differs by two symmetry labels, it is particularly relevant because the signed contrast $\mathcal{C}_0$ carries a $\mathcal{PR}$ correlation.

4.1.1 Intentional parameter variation

The procedure in which we intentionally vary experimental parameters or control knobs is referred to as intentional parameter variation (IPV). Often, a parameter P has an ideal value $P = 0$ at which no systematic error is expected. The final experiment is performed near this ideal point, with mean μ_P and uncertainty σ_P; these quantities are typically characterized using auxiliary measurements.

We then deliberately vary P to an exaggerated non-ideal value $P = \Delta P \gg \sigma_P$ and measure the resulting change in the eEDM channel, $\Delta\omega_{\mathrm{IPV}}^{\mathcal{NE}}$, with statistical uncertainty $\delta\omega_{\mathrm{IPV}}^{\mathcal{NE}}$. From this, we estimate the slope

$$S_P = \frac{\Delta\omega_{\mathrm{IPV}}^{\mathcal{NE}}}{\Delta P}, \tag{4.1}$$

with uncertainty

$$\delta S_P = \frac{\delta\omega_{\mathrm{IPV}}^{\mathcal{NE}}}{\Delta P}. \tag{4.2}$$

The systematic uncertainty associated with parameter P in the final measurement is then

$$\delta\omega_{\mathrm{sys,P}}^{\mathcal{NE}} = \sqrt{(S_P\sigma_P)^2 + (\delta S_P\mu_P)^2} \tag{4.3}$$

$$= \sqrt{\left(\frac{\Delta\omega_{\mathrm{IPV}}^{\mathcal{NE}}}{\Delta P}\sigma_P\right)^2 + \left(\frac{\delta\omega_{\mathrm{IPV}}^{\mathcal{NE}}}{\Delta P}\mu_P\right)^2}. \tag{4.4}$$

For most parameters we vary, the measured slope is statistically consistent with zero. In such cases, $\Delta\omega_{\mathrm{IPV}}^{\mathcal{NE}} \sim \delta\omega_{\mathrm{IPV}}^{\mathcal{NE}}$. Similarly, the knob is zeroed for the final measurement such that $\mu_P \sim \sigma_P$. Under these conditions, the two terms in the systematic uncertainty are

comparable and

$$\delta\omega^{\mathcal{NE}}_{\mathrm{sys,P}} \sim \frac{\sigma_P}{\Delta P}\,\delta\omega^{\mathcal{NE}}_{\mathrm{IPV}}. \tag{4.5}$$

This uncertainty is proportional to the statistical uncertainty of the IPV measurement, $\delta\omega^{\mathcal{NE}}_{\mathrm{IPV}}$, as expected. Crucially, it is also inversely proportional to the ratio $\Delta P/\sigma_P$, which quantifies the degree of exaggeration used in the IPV measurement. Without this enhancement factor, the systematic uncertainty associated with each parameter would exceed the final statistical uncertainty of the eEDM measurement unless each IPV study were carried out for a duration comparable to the full eEDM data run, which is impractical. The IPV exaggeration ratio is therefore the key factor that enables a comprehensive systematic search within a reasonable amount of time.

The IPV procedure outlined above relies on two important assumptions. First, it assumes that the systematic effect vanishes at the ideal point $P = 0$. This assumption is often well motivated but rarely proven rigorously. Second, it assumes that the eEDM response is linear over the range $[0, \Delta P]$[1]. This assumption is typically valid because ΔP is chosen to be large compared to the residual imperfection σ_P but still small compared to the scale at which nonlinear effects become significant.

4.1.2 Slow switches

Not every parameter P has an ideal value at which zero systematic error is expected. For example, we perform the eEDM measurement at two different electric field magnitudes, $\mathcal{E} = 80$ and 140 V/cm. The molecule is fully polarized at both field strengths due to the small Ω-doublet splitting, and therefore there is no reason to expect either electric field value to be intrinsically ideal. In such cases, the only reasonably rigorous approach is to perform the eEDM measurement at both field values, spending approximately half of the total data-

1. This estimate is conservative for linear or superlinear scaling. If the scaling is sublinear, the systematic uncertainty would be underestimated.

taking time at each setting. The difference between the two measured eEDM values then provides an estimate of the systematic uncertainty associated with the choice of electric field strength. From this, it is immediately clear that the resulting systematic uncertainty cannot, in general, be smaller than the final statistical uncertainty. In ACME III, we refer to this approach as *slow switching*, because these parameters are switched on timescales of hours, much longer than the duration of a superblock (a few minutes).

In practice, the difficulty of the experiment tends to scale exponentially with the number of such slow switches. This is because slow switches are typically much harder to implement and often require manual intervention. If there are n slow switches and equal time is to be spent at each setting, approximately 2^n configurations must be measured. As an example, the $\hat{k}$ switch corresponds to reversing the propagation direction of the 703 nm readout laser. Ideally, this could be accomplished by installing optics on both sides of the vacuum chamber and switching the laser power between them. In reality, however, one side of the chamber is occupied by optics required for 746 nm co-magnetometry. As a result, switching $\hat{k}$ requires manual rearrangement of optics, which is time consuming. At the time of writing, the final list of slow switches to be implemented in the production data run is still under internal discussion.

4.2 Inclusion criteria for systematics

In principle, if we include systematic uncertainties from all parameters that were varied, the final systematic error bar would be very large. This is because not every parameter can be varied with a high IPV ratio, and the systematic uncertainty from some parameters may exceed the final statistical uncertainty. Including such parameters would unnecessarily inflate the final systematic error budget when no systematic slope is observed. To avoid this, we adopt a set of inclusion criteria to determine which effects enter the final systematic error budget. These criteria closely follow those used in previous generations of the experiment,

with only minimal modifications [22, 23], and are summarized below.

A. In an IPV measurement, if a non-zero slope is observed at greater than the 3σ confidence level, we initiate an investigation to understand the origin of the systematic effect. Unless the observed slope is clearly attributable to an accidental or correctable mistake in the setup, the corresponding systematic uncertainty is included in the final budget using the procedure described in Sec. 4.1.1.

   Similarly, if a discrepancy in the eEDM result is observed between two slow-switch settings at greater than the 3σ confidence level, a systematic uncertainty equal to the size of the discrepancy is included.

B. If a non-zero mean is observed at greater than the 5σ confidence level in a parity channel that is one symmetry label away from the eEDM channel $\mathcal{NE}$, or in a channel differing by a $\mathcal{PR}$ label, we investigate the origin of the unexpected mean. If no satisfactory explanation is found, we assign a systematic uncertainty due to possible correlation or leakage between that channel and the eEDM channel.

C. If a similar experiment has previously observed a systematic effect associated with a parameter P, but the underlying mechanism was not understood, we include the systematic uncertainty inferred from the IPV measurement of that parameter, even if no statistically significant slope is observed in our own IPV study.

In previous generations of the experiment, these were referred to as class A, B, and C systematics, respectively. Systematic effects that do not satisfy any of the above inclusion criteria are not included in the final systematic error budget.

4.3 Auxiliary measurements

As discussed in Sec. 4.1.1, estimating the systematic uncertainty associated with each parameter P requires knowledge of its mean value μ_P and uncertainty σ_P during the fi-

nal eEDM measurement. Some parameters, such as laser detuning, are straightforward to characterize. Others are both more difficult to measure and more critical for controlling systematics. These include the non-reversing electric field $\mathcal{E}^{\mathrm{nr}}$, magnetic field gradients, and laser beam ellipticity. We therefore perform a set of auxiliary measurements to accurately characterize these quantities.

4.3.1 Microwave electric field mapping

In the first generation of the experiment, the non-reversing electric field $\mathcal{E}^{\mathrm{nr}}$ was identified as a significant source of systematic error [22]. A microwave spectroscopy technique was subsequently developed to map $\mathcal{E}^{\mathrm{nr}}$ in the spin precession region. This technique, with only minor refinements, has been used in all three generations of the experiment.

The method relies on driving the microwave transition

$$|H, J = 1, M = \pm 1\rangle \rightarrow |H, J = 2, M = 0\rangle \,,$$

which occurs near 39 GHz and has a differential Stark shift of approximately 1 kHz/(mV/cm). By scanning the microwave frequency and measuring population transfer as a function of molecular flight time, we obtain a spatial map of the electric field in the spin precession region. A detailed description of the technique is given in Zack Lasner's thesis [91]. The microwave system was set up by Xing Fan and Ayami Hiramoto; here we report only the results.

The results of the microwave electric field mapping are shown in figure 4.1. We observe a non-reversing electric field $\mathcal{E}^{\mathrm{nr}}$ at the level of $\sim$ 5–10 mV/cm throughout the spin precession region. This magnitude is comparable to that observed in the previous generation of the experiment and is likely attributable to patch potentials on ITO coating of the field plates. The high-voltage supplies were independently calibrated to remove offsets, and the residual

lead-correlated field $\mathcal{E}^{\mathcal{L}}$ is consistent with zero within uncertainties.

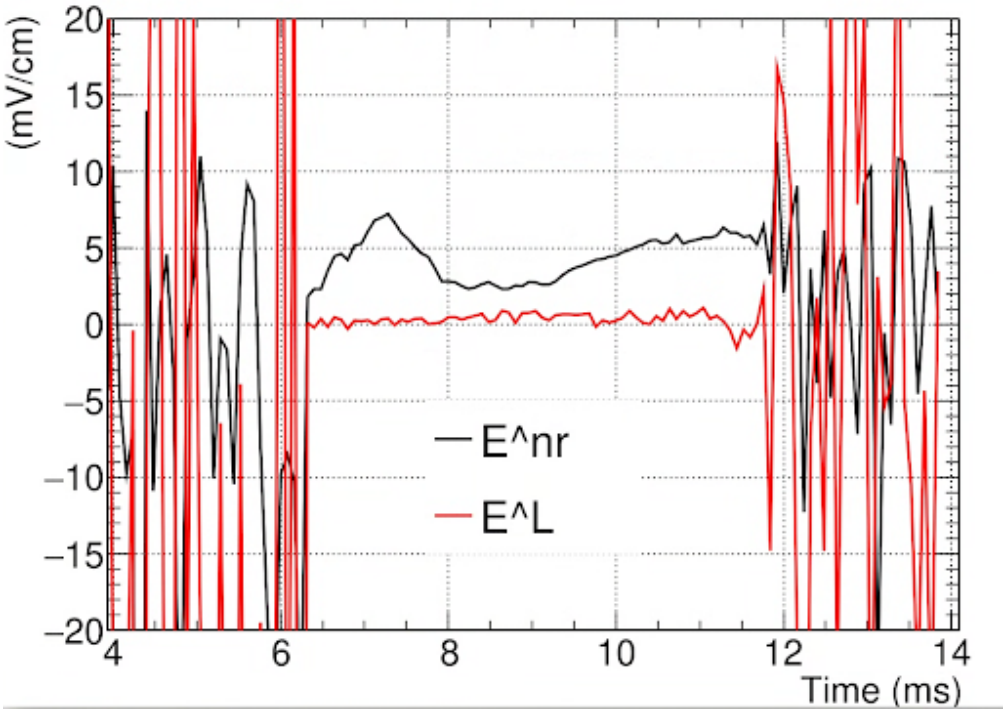

Figure 4.1: Electric field mapping obtained using microwave spectroscopy. Two physically relevant parity channels are shown. Black: non-reversing electric field $\mathcal{E}^{\mathrm{nr}}$ attributed to patch potentials on the ITO coating of the field plates. Red: lead-correlated electric field $\mathcal{E}^{\mathcal{L}}$, sensitive to possible voltage offsets in the electric-field supplies. Figure prepared by Ayami Hiramoto.

4.3.2 Magnetic field mapping using commercial magnetometers

Historically, magnetic fields have been a significant source of systematic errors in eEDM experiments. In ACME II, a mechanism was identified by which magnetic field gradients can generate eEDM-correlated systematics [90, 91]. It is therefore crucial to accurately characterize magnetic field gradients within the spin precession region. After the interaction region chamber was installed inside the magnetic shields, we observed that the residual magnetic field increased from approximately 10 μG to 50 μG. The origin of this increase is not fully understood. To investigate it, we performed detailed three-dimensional magnetic field mapping using commercial QuSpin magnetometers with a resolution of approximately

1 μG.

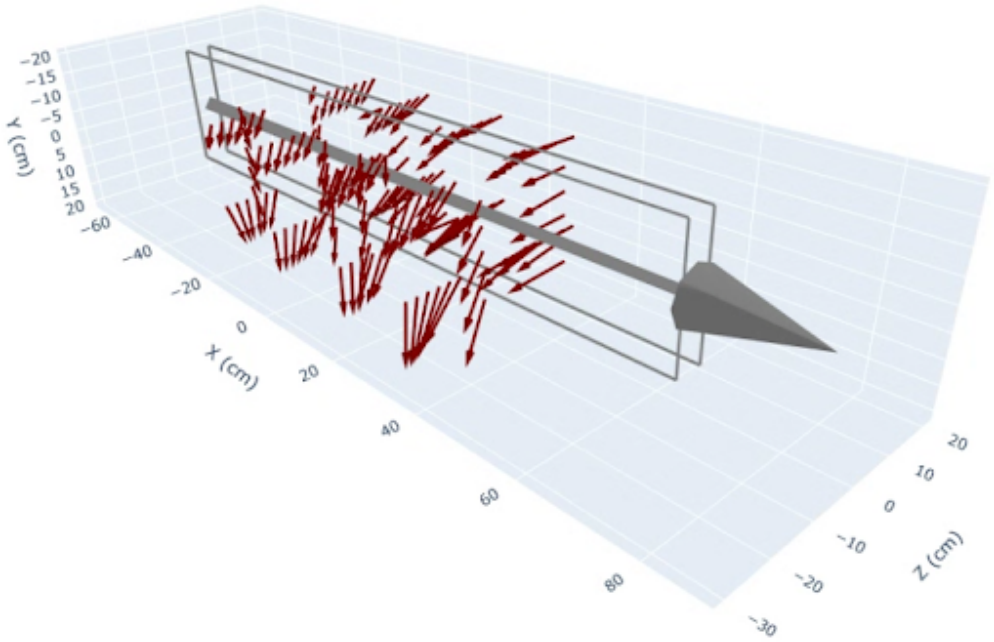

Figure 4.2: Magnetic field mapping using commercial QuSpin magnetometers. Gray arrow: molecular beam direction. Black square: field plates. Red arrows: direction of the magnetic field component being mapped. The maximum observed field strength is approximately 50–60 μG.

Further development of the magnetic field mapping was led by Maya Watts and Ayami Hiramoto. Approximately eight QuSpin magnetometers were mounted on a motorized rotation stage and inserted into the magnetometer pockets of the interaction region. The QuSpins are vector magnetometers capable of measuring all three components of the magnetic field. Controlled rotations were used to calibrate offsets in the magnetometers. Unlike the fluxgate magnetometers employed in earlier generations of the experiment, the QuSpin sensors are zero-field rubidium magnetometers that require only a small bias field for operation[2]. We did not observe any measurable increase in noise in the eEDM channel when the QuSpin magnetometers were operating.

Although it is in principle possible to perform the eEDM measurement with the QuSpins

2. The Rb magnetometers require a bias field comparable to the ambient background field, whereas fluxgate magnetometers require a field large enough to saturate a ferromagnetic core.

powered on, we are reconsidering this approach. When the QuSpins are operated during magnetic field reversals, the internal temperature of the sensors undergoes thermal cycles. This behavior was confirmed by independent measurements of the external case temperature performed by Ayami Hiramoto and Maya Watts. We suspect that repeated thermal cycling is responsible for the observed failure of several QuSpin units after approximately one year of operation. Ongoing efforts aim to mitigate this issue by actively air-cooling the QuSpins, although the effectiveness of this approach remains uncertain. In addition, the QuSpin sensors employ AC heating coils to regulate the temperature of the rubidium vapor cell. The associated AC magnetic fields may constitute an additional concern, as none of our available magnetometers can directly characterize these fields.

4.3.3 Q state co-magnetometer

In addition to mapping the magnetic fields using commercial magnetometers, we also developed a co-magnetometer based on the Q state of ThO. Unlike the H state used for the eEDM measurement, the Q state has a much larger g factor of around 0.7, making it ideal for measuring magnetic fields. The Q state is also already prepared in the lens STIRAP process. To utilize it for the co-magnetometer, we simply turn off the Qleanup beam, which normally depletes the Q-state population. The molecules in the Q state then enter the spin precession region and accumulate Zeeman phase just like the H-state molecules. The preparation and readout scheme of the Q state is exactly the same as that of the H state, except that we use a different laser frequency of around 746 nm.

The initial testing of the Q-state co-magnetometer was led by Xing Wu and me at the Harvard beamline. After we moved to Northwestern, Collin Diver took over the development with some help from Zhen Han and me. Because of the large Stark shift along the Z direction, we are only sensitive to the magnetic field component along Z. This means we can only measure the $\mathcal{B}_Z$ field and the gradients $\partial\mathcal{B}_Z/\partial X$, $\partial\mathcal{B}_Z/\partial Y$, and $\partial\mathcal{B}_Z/\partial Z$. As a quick

reminder, the Z direction is the direction of the applied electric field, X is the direction of the molecular beam, and Y is the lab vertical direction and the direction along which we use a cylindrical lens to expand the laser beams.

To measure the $\partial\mathcal{B}_Z/\partial X$ gradient, a longitudinal preparation beam that drives the $Q \mapsto C$ transition is applied. The magnetic field as a function of the X coordinate is then obtained using time-of-flight information.

To measure the $\partial\mathcal{B}_Z/\partial Y$ gradient, we block half of the readout beam using a razor blade. The change in the average measured magnetic field is then used to extract the gradient information.

To measure the $\partial\mathcal{B}_Z/\partial Z$ gradient, a $Q \mapsto I$ laser beam is sent vertically downward in the Qleanup region to selectively deplete molecules at different Z positions.

4.3.4 Molecular polarimetry

The idea for molecular polarimetry originated from the need to accurately measure the laser beam ellipticity inside the spin precession region. One concern about measuring the laser beam ellipticity outside the vacuum chamber is that only the net birefringence from both sides of the field plates and vacuum windows can be measured. In reality, the molecules only see the light after it passes through one side of the field plates and vacuum windows. To address this concern, we developed a molecular polarimetry technique to measure the laser beam ellipticity *in situ* using the molecules themselves. The basic idea utilizes the correspondence principle explained in 2.2.2. For example, to measure the ellipticity of the preparation laser, we note that the molecular dark state prepared by the preparation laser is defined by its polarization,

$$|D\rangle = -e^{-i\theta_{\text{prep}}} \cos\Theta_{\text{prep}} |M = +1\rangle + e^{i\theta_{\text{prep}}} \sin\Theta_{\text{prep}} |M = -1\rangle, \tag{4.6}$$

where $\Theta_{\text{prep}} = \frac{\pi}{4} + d\Theta_{\text{prep}}$ describes nearly linear light with a small ellipticity parameterized by $d\Theta_{\text{prep}}$. This ellipticity is also proportional to the population difference between the $M = \pm 1$ states, with

$$\sin^2 \Theta_{\text{prep}} - \cos^2 \Theta_{\text{prep}} = \sin(2d\Theta_{\text{prep}}) \tag{4.7}$$

$$\approx 2d\Theta_{\text{prep}}. \tag{4.8}$$

This means that to measure the ellipticity $d\Theta_{\text{prep}}$, we simply need to measure the population difference between the $M = \pm 1$ states after the preparation laser. This is done by switching the polarization of the readout laser between left- and right-circular polarization, which selectively measures the population in the $M = +1$ and $M = -1$ states, respectively, as shown in figure 4.3. The difference between these two measurements then gives the population difference and thus the ellipticity of the preparation laser. In a manner similar to the regular eEDM measurement, we define the asymmetry for molecular polarimetry as

$$\mathcal{A}_{\text{polarimetry}} = \frac{S_{\sigma+} - S_{\sigma-}}{S_{\sigma+} + S_{\sigma-}} \approx 2d\Theta_{\text{prep}}. \tag{4.9}$$

In terms of the Stokes parameters, the asymmetry defined here is equivalent to S/I, where S is the circular-polarization Stokes parameter and I is the total intensity. We therefore also use the notation S/I interchangeably to refer to the asymmetry measured in molecular polarimetry.

Molecular polarimetry can also be performed for the readout laser in a similar manner. We perform the polarimetry measurements with switches analogous to those used in the regular eEDM measurement to isolate various effects, including imperfections from the waveplates. Because this technique is newly introduced in ACME III, we include a more detailed discussion of the procedure and interpretation of the results in appendix C. One major benefit of this polarimetry technique is that it is almost trivial to set up: it only requires

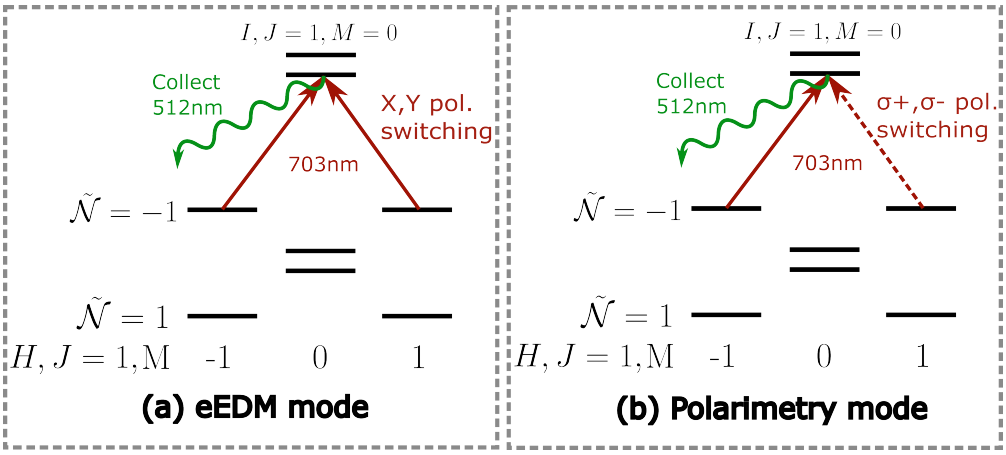

Figure 4.3: (a) Polarization switching used for the regular eEDM measurement, with X and Y linear polarizations alternated. (b) Polarization switching scheme used for molecular polarimetry, with left- and right-circular polarizations alternated.

inserting and aligning a quarter-wave plate in the optics. In the final eEDM measurement, however, the quarter-wave plate used for polarimetry will not be left in the optical path due to concerns about its birefringence.

4.4 Discovered systematics and mitigation

4.4.1 Structure of this section

The tradition in ACME theses is to discuss all systematic errors that meet the inclusion criteria defined in section 4.2. This includes systematic errors identified in previous generations of the experiment. I choose not to follow this tradition here for two reasons. First, at the time of writing, the systematic error studies are not yet complete, so a comprehensive summary is not possible. Second, we have discovered many new phenomena in this generation of the experiment, and I believe it is more fruitful to focus on these new discoveries rather than rehashing previously studied systematics. Readers interested in earlier systematic errors are referred to the theses of Cris Panda and Zack Lasner [90, 91].

There are several aspects of the discovered systematics that may be of general interest:

how they were discovered, how they were modeled, and how they were mitigated. The modeling aspect tends to be more technical and is therefore separated into its own section, section 4.5. For the systematic error arising from excited-state decay, the modeling has not been worked out in previous generations and therefore warrants its own chapter. This is discussed in detail in chapter 5.

In this section, we focus on the discovery process and the mitigation strategies.

4.4.2 AC Stark shift from reflection

We send the 703 nm cleanup and readout lasers through a series of optics before they encounter the molecular beam. Ideally, we align the laser beam via retroreflection from the electric field plates, which define the quantization axis and the lab Z direction. However, the vacuum windows are not aligned with the field plates at the mrad level. The two windows also appear to be misaligned with respect to each other at a similar level. This means we effectively have three different reflections of the laser beam: one from the front vacuum window, one from the back vacuum window, and one from the two field plates, which were aligned using white-light interferometry[3]. The front-window reflection cannot reach the molecules, but we theorize that if either of the other two reflections hits the molecules, it could cause systematic errors via AC Stark shift.

Recall that, due to our switching scheme, most AC Stark shift effects do not show up as systematic errors. For an AC Stark shift to generate a systematic error, it must be correlated with $\mathcal{N}\mathcal{E}$. In our experiment, we have a correlated detuning $\Delta^{\mathcal{N}\mathcal{E}}$ due to the presence of a non-reversing electric field $\mathcal{E}^{\mathrm{nr}}$, which does not reverse in the lab frame when we attempt to flip the applied electric field $\mathcal{E}$. This non-reversing field couples to the orientation of the molecular dipole $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ in the lab frame via the linear Stark shift, generating a correlated detuning $\Delta^{\mathcal{N}\mathcal{E}}$, as shown in figure 4.4.

3. We only noticed this after assembling the magnetic shields, so fixing it could take weeks of beam time.

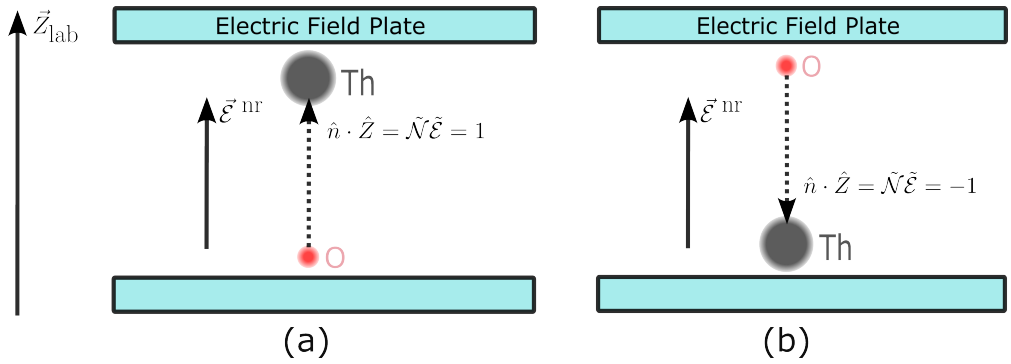

Figure 4.4: Origin of the correlated detuning $\Delta^{\mathcal{NE}}$ from the non-reversing electric field $\mathcal{E}^{\mathrm{nr}}$ via. linear Stark shift $-(d_{\mathrm{ThO}}\hat{n})\cdot\vec{\mathcal{E}}^{\mathrm{nr}} = -d_{\mathrm{ThO}}\mathcal{E}^{\mathrm{nr}}\tilde{\mathcal{N}}\tilde{\mathcal{E}}$.(a) When the molecular lab orientation $\mathcal{NE}$ is aligned with the non-reversing electric field $\mathcal{E}^{\mathrm{nr}}$, the energy is shifted down. (b) When the molecular lab orientation $\mathcal{NE}$ is anti-aligned with $\mathcal{E}^{\mathrm{nr}}$, the energy is shifted up. The differential Stark shift between different $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ configurations thus leads to a correlated detuning $\Delta^{\mathcal{NE}}$.

In addition, our laser beams acquire a small ellipticity $d\Theta$ from the birefringence of the optics they pass through. The effect of this ellipticity is that one of the $M = \pm 1$ states is coupled more strongly to the excited state than the other. This leads to a differential AC Stark shift between the two M states, as shown in figure 4.5. The amplitude of this differential AC Stark shift is proportional to both the detuning Δ and the ellipticity $d\Theta$ of the laser beam. Therefore, the presence of a correlated detuning $\Delta^{\mathcal{NE}}$ leads to a differential AC Stark shift that is correlated with $\mathcal{NE}$, which directly enters the eEDM channel as a systematic error.

The derivation of this systematic error model is shown in section 4.5.1, which uses time-independent perturbation theory. This treatment assumes that the laser beam is unsaturated, with a Rabi frequency much smaller than either the decay rate or the detuning. This assumption is only valid for the weak reflection beams. To mitigate this systematic error for the cleanup beam, we deliberately slightly misalign the cleanup beam such that the reflection intersects the molecular beam slightly upstream of the main beam. In this way, any systematic phase accumulated from the reflection beam is removed by the main beam through dissipation. A similar procedure is applied to the readout beam, with the reflection

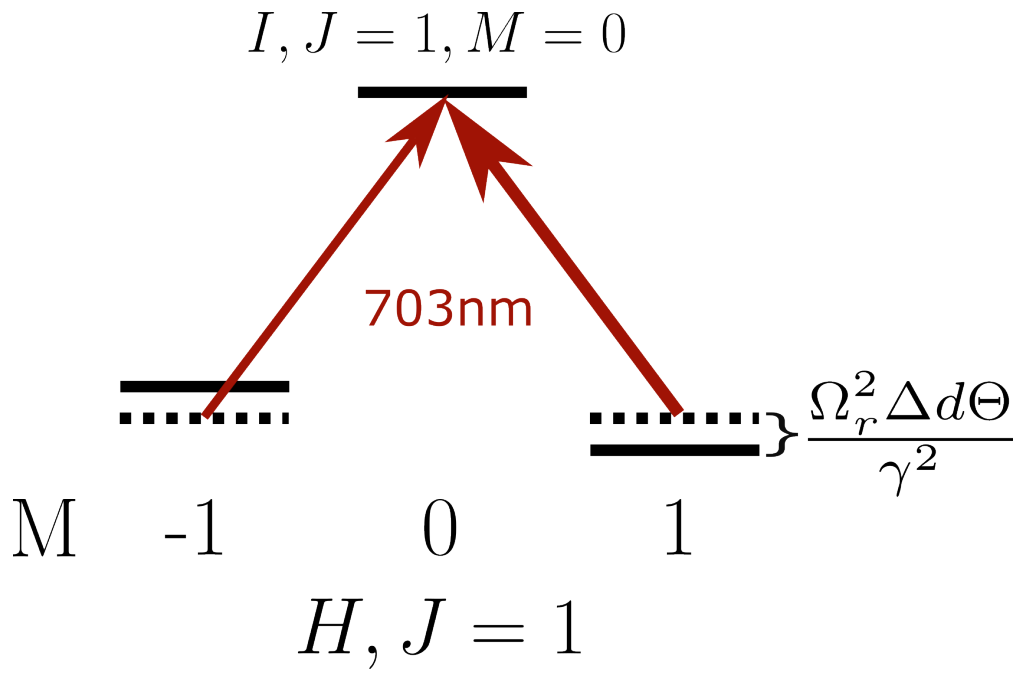

Figure 4.5: Coupling diagram illustrating the AC Stark shift systematic from an unsaturated laser beam, e.g. a reflection beam. Here, Δ is the detuning of the laser from resonance, Ω_r is the Rabi frequency of the laser, γ is the decay rate of the excited state, and $d\Theta$ is the ellipticity of the laser beam. The ellipticity causes different coupling strengths to the $M = \pm 1$ states, leading to a differential AC Stark shift that is proportional to both the detuning and the ellipticity.

aligned downstream.

We first observed this systematic error when attempting to measure the electric field using a two-photon Raman transition. We noticed that even with large one-photon and two-photon detunings, phase shifts were still observed in the molecules as we modulated the ellipticity of the Raman beams. This led us to the AC Stark shift model in the weakly driven regime. We have also observed evidence of this effect when implementing the $\mathscr{C}$ switch, which flips the ellipticity of the laser beam without necessarily preserving power or alignment. This switch should convert the AC Stark shift systematic from the $\mathcal{NE}$ channel to the $\mathcal{NEC}$ channel. When the conversion was imperfect, we were usually able to identify a downstream weak reflection beam that appeared only in one of the $\tilde{\mathscr{C}}$ states due to a $\mathscr{C}$-correlated misalignment. Correcting the alignment of this reflection beam typically restored

proper conversion. Further discussion of the $\mathscr{C}$ switch can be found in section 4.4.4.

4.4.3 AC Stark shift from the cleanup beam

A very similar AC Stark shift systematic error due to the correlated detuning $\Delta^{\mathcal{NE}}$ also exists for the main cleanup beam. However, because the cleanup beam is typically very strong and saturated, the perturbative model used for the weakly driven system in the previous subsection is no longer valid. Physically, the cleanup beam provides a strong dissipative interaction that removes any memory of the molecular state. This strong drive then ramps down on a timescale of around 10 μs as the molecules exit the laser beam. During the strong drive, any correlated detuning $\Delta^{\mathcal{NE}} \ll \Omega_r$ is not expected to cause significant systematic errors, where Ω_r is the maximum Rabi frequency of the cleanup beam. The systematic error only begins to accumulate during the ramp-down process. This therefore becomes a time-dependent perturbation problem, which is much more difficult to solve. Understanding this systematic error was one of the main projects in ACME I. We reproduce their derivations in section 4.5.2 and only summarize the main results here.

The functional form of the systematic error from the cleanup beam turns out to be very similar to that of the weak reflection beam. The main difference is that the proportionality constant is different and must be determined numerically. The systematic error from the cleanup beam can be expressed as

$$\phi^{\mathcal{NE}}_{\text{AC Stark}} \approx \frac{d\Theta}{dX}\sigma_X \cdot \frac{\Delta^{\mathcal{NE}}}{\sigma_\Pi}, \tag{4.10}$$

where $\frac{d\Theta}{dX}$ is the ellipticity gradient of the cleanup beam in the lab X direction, σ_X is the width of the laser beam in the X direction, and σ_Π is a numerically calculated constant that depends on the exact dynamics of the ramp-down process. For our experiment, the ellipticity gradient $\frac{d\Theta}{dX}$ is typically around 2×10^{-4}/mm, the beam width σ_X is around

1 mm, $\Delta^{\mathcal{NE}}$ is around $2\pi \times 5$ kHz, and σ_Π is around $2\pi \times 10$ MHz. Plugging in these values gives $\phi^{\mathcal{NE}}_{\text{AC Stark}} \sim 100$ nrad and $\omega^{\mathcal{NE}}_{\text{AC Stark}} \sim 20$ μrad/s. This is comparable to the final statistical uncertainty of the eEDM measurement.

Interestingly, the numerical constant σ_Π turns out to be mostly insensitive to the Rabi frequency. Significantly different Rabi frequencies have been used across different generations of ACME, yet this numerical constant remains roughly the same [85–87]. This was understood in ACME I as a consequence of adiabatic following during the ramp-down process. The detuning dependence only becomes important during the last Rabi oscillation cycle, when the Rabi frequency is no longer sufficient to pump all molecules out of the dark state. Changing the Rabi frequency therefore mainly shifts the timing of this last cycle, without substantially altering the overall dynamics.

Another difference between the cleanup-beam systematic and the weak-reflection-beam systematic is that the cleanup-beam effect is sensitive to the spatial gradient of the ellipticity rather than the absolute ellipticity. This is because, in the high-intensity region of the cleanup beam, the molecules begin in the instantaneous dark state defined by the local ellipticity. Additional pumping dynamics only occur if the dark state itself changes as the molecules move through the beam. This change is quantified by the spatial gradient of the ellipticity.

To mitigate this systematic error, we have carefully engineered the field plates and vacuum windows to be as free of birefringence gradients as possible, as explained in sections 2.5.8 and 2.5.9. For the optics outside of vacuum, we avoided the use of any waveplates after the final polarizer and used only large lenses for beam shaping. We find that large lenses tend to have smaller birefringence gradients than smaller lenses, likely because it is much harder to securely mount a small lens without introducing stress.

In addition to this careful engineering to reduce birefringence gradients, we developed a new switching scheme to further suppress this systematic error. This new switch, called the $\mathscr{C}$ switch, flips the ellipticity of the laser beam. In this way, any systematic error from the

cleanup beam is converted from the $\mathcal{NE}$ channel to the $\mathcal{NEC}$ channel. By averaging over both $\mathscr{C}$ states, the systematic error from the cleanup beam is suppressed. More discussion of the $\mathscr{C}$ switch can be found in section 4.4.4.

4.4.4 Addition of the $\mathscr{C}$ switch

In ACME I, there was a switch called the $\mathcal{G}$ switch, which performs a global rotation of the readout and cleanup polarizations by 90°. This rotation flips the ellipticity of the laser beams and therefore converts AC Stark shift systematic errors from the cleanup and readout beams from the $\mathcal{NE}$ channel to the $\mathcal{NEG}$ channel. However, this suppression mechanism was not well recognized at the time, and the $\mathcal{G}$ switch was not implemented in ACME II due to the difficulty of state preparation with the vertical STIRAP scheme.

In ACME III, the same technical difficulty with the vertical STIRAP scheme still exists, as the STIRAP beams cannot have any significant polarization component along the lab vertical direction. This means that if we naively rotate the cleanup and readout polarizations by 90°, the molecular state prepared by STIRAP will almost completely overlap with the bright state of the cleanup beam, leading to nearly complete loss of signal. To circumvent this problem, we developed a new switch called the $\mathscr{C}$ switch, which flips the ellipticity of the cleanup beam without significantly redefining the dark state. To see this, let the polarization of the cleanup beam be described by

$$\hat{\epsilon} = -e^{-i\theta}\cos\Theta\hat{\epsilon}_{+1} + e^{i\theta}\sin\Theta\hat{\epsilon}_{-1}. \tag{4.11}$$

This couples the science H state to the excited I state, as shown in figure 4.3. Recall that the $I, M = 0$ state remains a good parity eigenstate even in the presence of an electric field

and is given by

$$|I, M=0, \tilde{\mathcal{P}}\rangle = \frac{1}{\sqrt{2}}\left(|I, \Omega=+1, M=0\rangle + \tilde{\mathcal{P}}\,|I, \Omega=-1, M=0\rangle\right). \tag{4.12}$$

This $\tilde{\mathcal{P}}$-dependent sign has real consequences for the molecular dark state defined by the cleanup beam. The molecular dark state is given by the effective polarization

$$\hat{\epsilon}_{\text{eff}} = -e^{-i\theta - i(\tilde{\mathcal{P}}+1)\frac{\pi}{4}} \cos\Theta \hat{\epsilon}_{+1} + e^{i\theta + i(\tilde{\mathcal{P}}+1)\frac{\pi}{4}} \sin\Theta \hat{\epsilon}_{-1}, \tag{4.13}$$

which differs from the original polarization by a $\tilde{\mathcal{P}}$-dependent 90° rotation. From the perspective of the molecular dark state, rotating the linear polarization by 90° while changing the parity $\tilde{\mathcal{P}}$ of the excited state is therefore equivalent to a 180° rotation, which is the identity operation. However, from the perspective of birefringent optics, rotating the linear polarization by 90° flips the ellipticity[4]. This allows us to flip the ellipticity of the cleanup beam without significantly altering the molecular dark state by combining a 90° polarization rotation with a parity switch. We refer to this combined operation as the $\mathscr{C}$ switch.

We have demonstrated suppression of systematic errors using the $\mathscr{C}$ switch in several different settings. We first demonstrated the suppression by clamping a piece of acrylic with a vise, intentionally creating a large birefringence gradient. We later also demonstrated this effect unintentionally by inserting a small cylindrical window intended for intensity monitoring. This window was clamped by a standard retaining ring and exhibited a surprisingly large birefringence gradient. The results of this demonstration are shown in figure 4.6.

We expect the $\mathscr{C}$ switch to suppress AC Stark shift systematic errors from the main cleanup beam to well below the final statistical uncertainty, given the already low birefringence gradients achieved through engineering controls. However, there remain several AC Stark shift systematics that are not suppressed by the $\mathscr{C}$ switch.

4. If this is not obvious, we derive it in appendix C.

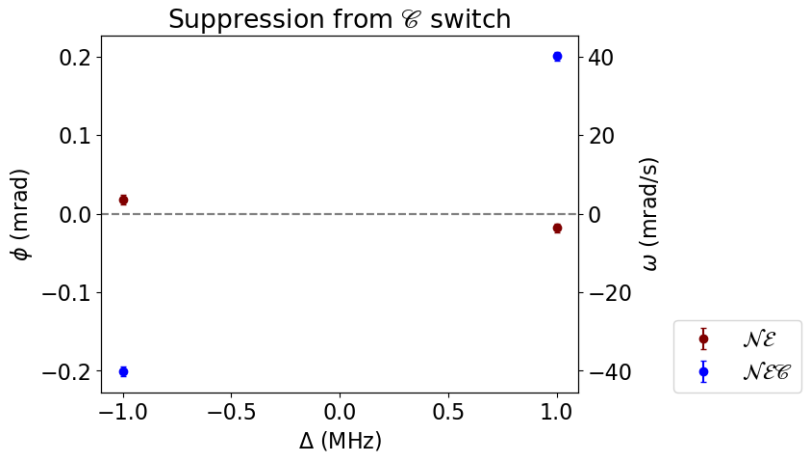

Figure 4.6: Demonstration of systematic error suppression from the $\mathscr{C}$ switch. Here, the dominant birefringence gradient comes from a small cylindrical window clamped very tightly. Without the $\mathscr{C}$ switch, almost all of the systematic error appears in the $\mathcal{NE}$ channel. With the $\mathscr{C}$ switch, the systematic error is converted to the $\mathcal{NEC}$ channel. The imperfect conversion is likely due to the readout beam, whose frequency was scanned together with the cleanup beam.

When we first implemented the $\mathscr{C}$ switch, the 90° polarization rotation was performed using a PBS with limited parallelism. As a result, the outgoing cleanup beam exhibited a $\mathscr{C}$-correlated pointing shift. This pointing shift can lead to $\mathscr{C}$-dependent stray reflections and generate unsuppressed systematic errors, as discussed in section 4.4.2. To mitigate this effect, we ultimately minimized the correlated pointing shift by pre-aligning the two $\mathscr{C}$ states and switching between them using motorized beam blockers. Details of this implementation are described in section 2.5.10.

Another source of unsuppressed AC Stark shift systematic error arises from the readout beam. The readout beam intrinsically includes $\mathscr{C}$ suppression [87], because the X and Y polarizations used for the eEDM measurement are already separated by 90°. However, while we can avoid half-wave plates in the cleanup beam path, we must use a half-wave plate to rotate the polarization of both the X and Y readout beams for the phase measurement scheme. Polarimetry measurements show that the waveplates we used exhibit significant first- and second-order birefringence gradients. As a result, even when the X and Y beam profiles

are well overlapped, their ellipticity gradients can differ sufficiently to generate a systematic error. Therefore, we will need to monitor the magnitude of this systematic throughout the final eEDM data run. Additional suppression techniques exist, such as rotating the half-wave plate by 90° and 180° to average over birefringence gradients. These methods have been experimentally demonstrated to be effective at suppressing readout-beam AC Stark systematics. At the time of writing, however, we have not decided whether to implement them in the final eEDM run, as they would increase the duration of a superblock, reducing the number of statistically independent data points and limiting the time available for slow switches.

4.4.5 AC Stark shift from pointing/differential phase

During the IPV process, we noticed that the measured eEDM depended on the pointing of the readout beam. To our surprise, this dependence was sinusoidal, with functional form

$$\phi_{\text{sys}} \propto (\text{XY overlap}) \cdot \sin(\alpha \cdot \Delta\theta_{XY}), \tag{4.14}$$

where $\Delta\theta_{XY}$ is the differential pointing angle between the two beams, and α is a proportionality constant. For our setup, $\alpha \approx \frac{2\pi}{40\ \mu\text{rad}}$. This sinusoidal dependence disappeared when the spatial overlap between the two beams was poor. This behavior puzzled us, and we began a series of careful two-dimensional scans to investigate how other experimental knobs affected this sinusoidal dependence. This was a month-long effort, during which we systematically went through all plausible knobs in the experiment. To perform this study at all, we first had to improve the pointing stability of the readout system to below the 10 μrad level. This was implemented by Zhen Han, who also first noticed the pointing dependence. The detailed implementation of the improved pointing stability system is described in section 2.5.10.

In these two-dimensional scans, we were unable to identify any knob that significantly

affected the sinusoidal dependence other than the spatial overlap. There was one important hint, however, that eventually led us in the right direction: the differential-pointing dependence was not isotropic. In particular, the sinusoidal dependence was much stronger along a direction approximately 45° from the lab X axis. Along the orthogonal direction, the dependence was much weaker. We suspected that this anisotropy could be used to identify the origin of the systematic error.

To test this, we installed a Dove prism in one of the readout beam paths. The Dove prism rotates the applied pointing by 90°, so naively we expected the sinusoidal dependence to rotate by 90° as well. To our surprise, the sinusoidal dependence was virtually unchanged, as shown in figure 4.7. This observation led to a shift in our interpretation: the relevant experimental parameter was not the differential pointing between the two beams itself, but rather another parameter that is modulated when the pointing of the readout beam is changed. Collin Diver quickly realized that one such parameter is the differential phase between the two readout polarizations. As the pointing of the readout beam is adjusted with a motorized mirror, the optical path length also changes slightly, at the scale of a few laser wavelengths, as illustrated in figure 4.8. This induces a differential phase shift between the X and Y polarizations. We tested this hypothesis by directly varying the differential phase, either by adjusting the AOM RF phase or by changing the optical path length with a motorized translation stage. Both tests confirmed that the differential phase between the two polarizations is indeed responsible for this systematic error.

How does the differential laser phase lead to systematic errors? It turns out that even though the RF TTL pulses used to switch the polarization of the AOMs are programmed to separate the two polarizations temporally, the extinction ratios of the AOMs are not perfect, and both polarizations are simultaneously present at a small level. This differential phase can then be converted into an ellipticity, which leads to AC Stark shift systematic errors through a mechanism similar to that explained in section 4.4.3.

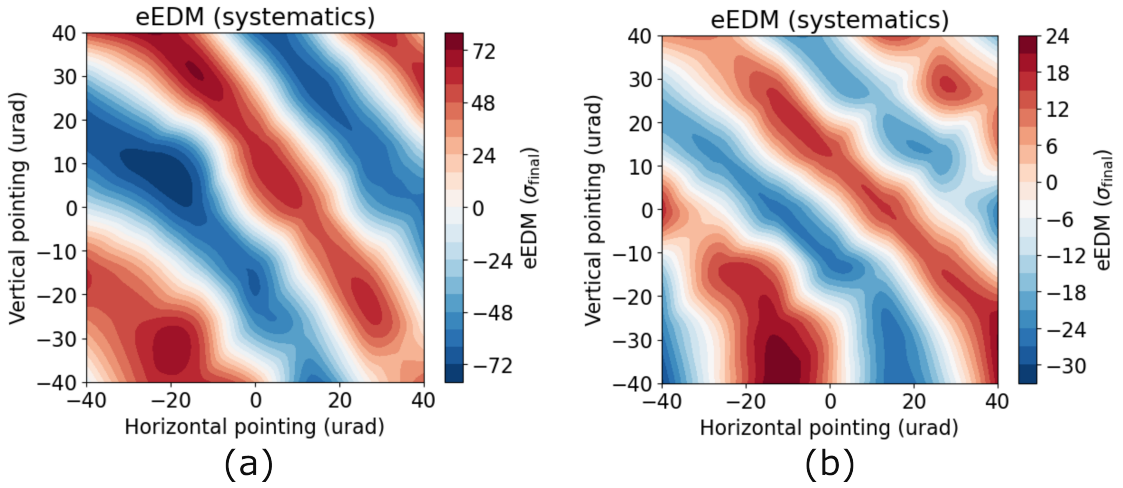

Figure 4.7: Two-dimensional scans of the pointing dependence, where the horizontal axis corresponds to horizontal pointing and the vertical axis corresponds to vertical pointing. Clear anisotropy is observed, as the sinusoidal dependence is much stronger along one diagonal direction. (a) After Dove prism rotation by 90°. (b) Before Dove prism rotation. The anisotropy direction remains unchanged even after the Dove prism rotation. The amplitude and phase change are likely due to slight misalignment of the Dove prism and fluctuations in the system between the two measurements. Color represents the size of the systematic error in multiples of our estimated final statistical uncertainty of 20 μrad/s.

We further investigated the origin of the finite extinction ratios of the AOMs. By performing an interferometric measurement of the AOM output beam, we discovered that the acoustic waves inside the AOM crystal were reflecting off the crystal surfaces, as shown in figure 4.9. Because our polarization switching cycle was not commensurate with the acoustic round-trip time, there was always some residual acoustic wave present when the other polarization was turned on.

To mitigate this systematic error, we first changed the polarization switching cycle to be close to the acoustic round-trip time of the AOM crystal. Due to concerns about timing jitter, we restricted the polarization switching cycle to be an integer multiple of one sample period of our data acquisition system (80 ns). Our final polarization switching period is 70 samples (5.6 μs), corresponding to a switching frequency of about 178 kHz. This procedure reduced the systematic error by approximately a factor of 20.

In addition, we added differential phase scrambling between the two polarizations by in-

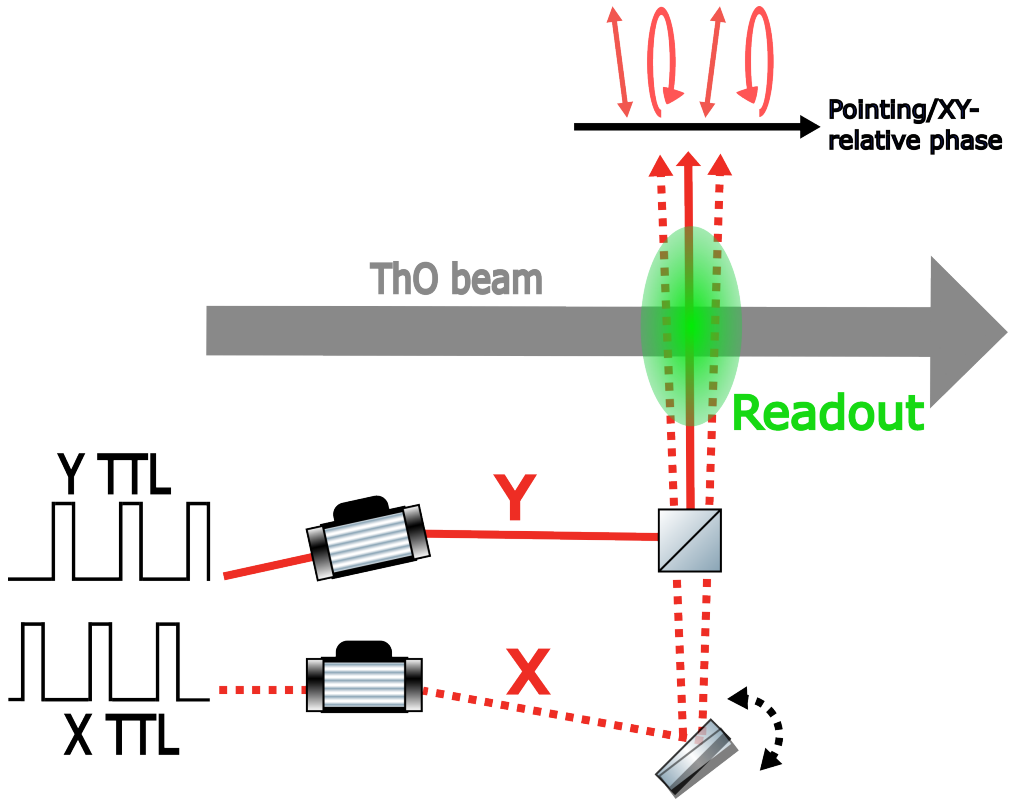

Figure 4.8: Origin of the so-called pointing systematics. As we modulate the pointing of the readout beam, the laser beam path length also changes slightly, at the scale of a few laser wavelengths. This leads to a differential phase between the X and Y polarizations. If, in addition, the extinction ratios of the AOMs used for polarization switching are not perfect, this differential phase can be converted into an ellipticity, which leads to AC Stark shift systematic errors.

troducing a 2 kHz frequency offset between the two AOM RF drives. This rapidly scrambles the differential phase between the two polarizations when combined with the random kinematic properties of the molecular beam. The combination of these two mitigation strategies reduced this systematic error to below a resolvable level.

We expect the 2 kHz frequency offset between the two AOM RF drives not to introduce additional systematic errors, because the frequency difference is small compared to the differential Doppler shift, which is limited by alignment ($\sim$ 10 kHz)[5]. We also performed

5. This suggests that the applied differential frequency is physically distinct from a differential Doppler shift. Although this may seem counterintuitive, the optical phase is given by $kx - \omega t$. In the laboratory frame, the Doppler shift arises from the k term via a phase gradient, while a frequency shift arises from the

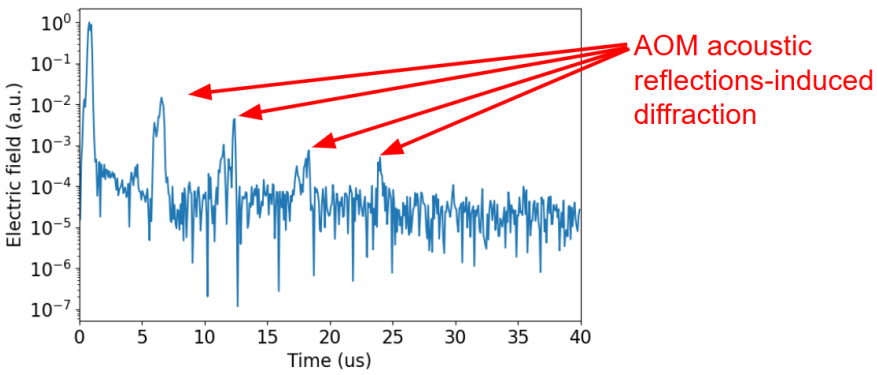

Figure 4.9: Interferometric measurement of the electric field amplitude of the AOM output beam as a function of time after one RF pulse. The multiple bumps in the electric field amplitude correspond to acoustic wave reflections inside the AOM crystal. The electric field is normalized to the first pulse, which is the main pulse used for the experiment. We estimated the round-trip time of the AOM using caliper measurements of the AOM crystal thickness and the speed of sound in the TeO_2 crystal. The round-trip time agrees with the time delay between the bumps in the electric field amplitude at the percent level.

IPV tests to confirm that this differential phase scrambling does not introduce additional systematic errors.

4.4.6 Q state decay

When we first started performing molecular polarimetry to check the polarization quality of our lasers, we noticed a puzzling phenomenon: there was an ellipticity component of the measured polarization that was correlated with $\mathcal{NE}$, $\mathcal{N}$, and $\mathcal{E}$. We denote these as $(S/I)^{\mathcal{NE}}$, $(S/I)^{\mathcal{N}}$, and $(S/I)^{\mathcal{E}}$ using Stokes-parameter notation. This was very concerning because new $\mathcal{NE}$-correlated imperfections in the system can lead to a new class of systematic errors.

To study this phenomenon, we performed a series of IPV tests to try to find a knob that would modulate these ellipticity components. After a month-long effort of systematically going through all possible knobs, we found that the voltage offset of the electric field plates can significantly modulate these ellipticity components. This was a surprise because the voltage offset nominally does not change the electric field inside the interaction region. However,

ω term. The two effects are therefore indeed physically distinct.

the voltage offset does change the electric field upstream of the glass electrodes, where the molecules fly through. This led us to suspect that the imperfections were generated during upstream state preparation.

We were quickly able to identify the problematic state preparation step: the lens STIRAP process, which prepares the molecules into the $Q, M = 2$ lensing state, was the main contributor to these ellipticity components. The smoking-gun evidence was that when we prepared the molecules into the $Q, M = -2$ lensing state instead, the sign of all these ellipticity components flipped.

How could the Q-state preparation affect what we measured in the H state? It turns out that the Q state decays into the H state via a ~ 12 μm transition. The $(S/I)^{\mathcal{NE}}$, $(S/I)^{\mathcal{N}}$, and $(S/I)^{\mathcal{E}}$ components we measured are essentially correlated population imbalances among the M sublevels, which can be partially preserved through the decay process[6]. We had considered this decay channel previously but rejected it because an order-unity ellipticity in the Q-state population would be required to explain the magnitude of the measured ellipticity components. This seemed unlikely because the molecules fly through a long region of low electric and magnetic fields, thereby losing a well-defined quantization axis. We therefore expected the Q state to be thoroughly mixed among all M sublevels before entering the interaction region.

It turns out that our intuition was completely wrong. We provide a brief qualitative explanation here and relegate the detailed modeling to chapter 5. To explain this phenomenon, we first need to convert the observed $(S/I)^{\mathcal{NE}}$ into a more transparent quantity in the molecular basis. Recall that we define S/I as the population in $M = +1$ minus the population in $M = -1$, divided by the total population within the same $\tilde{\mathcal{N}}$ state. Essentially, it is a measure of an M-correlated population imbalance. Therefore, the observed $(S/I)^{\mathcal{NE}}$ corresponds to a $\mathcal{NE}M = \Omega$-correlated population imbalance, where, as a reminder, Ω is the

6. The selection rule is $|\Delta M| \leq 1$ and Ω must be preserved.

projection of the total electronic angular momentum along the molecular axis. This recasts the observation into a more familiar language for AMO physicists.

Recall that in the lens STIRAP process, we prepare the molecules into the $M = 2$ sublevel of the lensing state. For a fixed orientation of the electric field in the STIRAP region, this fully specifies the Ω state of the molecules. Without loss of generality, let the molecules be prepared in the $\Omega = +2$ state. To explain the observed $(S/I)^{\mathcal{NE}}$, we therefore need to understand how the molecules can remain predominantly in the $\Omega = +2$ state throughout the entire flight path despite possible mixing from electric and magnetic fields.

The different mixing and coupling mechanisms are summarized in figure 4.10. The mixing from ambient fields depends on the adiabaticity of the system and is explained in detail in chapter 5. The key observation is that, regardless of the exact dynamics, ambient fields cannot efficiently mix different Ω states. This is intuitive because Ω is the electronic angular momentum about the molecular axis and behaves like a gyroscopic degree of freedom. An ideal gyroscope is insensitive to reorientation of its spin axis and maintains its spin angular momentum. In reality, non-ideal gyroscopes do respond to reorientation via Coriolis-type coupling. In diatomic molecules, this coupling is precisely the Ω-doublet splitting, quantified by Δ_Ω. In the Q state, this splitting satisfies $\Delta_\Omega < 10$ kHz, resulting in nearly ideal gyroscopic behavior[7]. Therefore, the molecules approximately preserve their initial Ω state throughout the entire flight path, leading to the population pattern shown in figure 4.10 and the observed $(S/I)^{\mathcal{NE}}$ measured in polarimetry mode. Some mixing between the $\Omega = \pm 2$ states does occur, leading to imperfect Ω conservation. This residual mixing turns out to explain the observed $(S/I)^{\mathcal{N}}$ and $(S/I)^{\mathcal{E}}$.

After we understood the mechanism of this decay pattern, we characterized the resulting systematic errors using an IPV-style measurement led by Collin Diver. The intentional parameter we varied was the initial population imbalance in the $Q, M = \pm 2$ state. By

7. This smallness is expected because the angular momentum would need to change by four units.

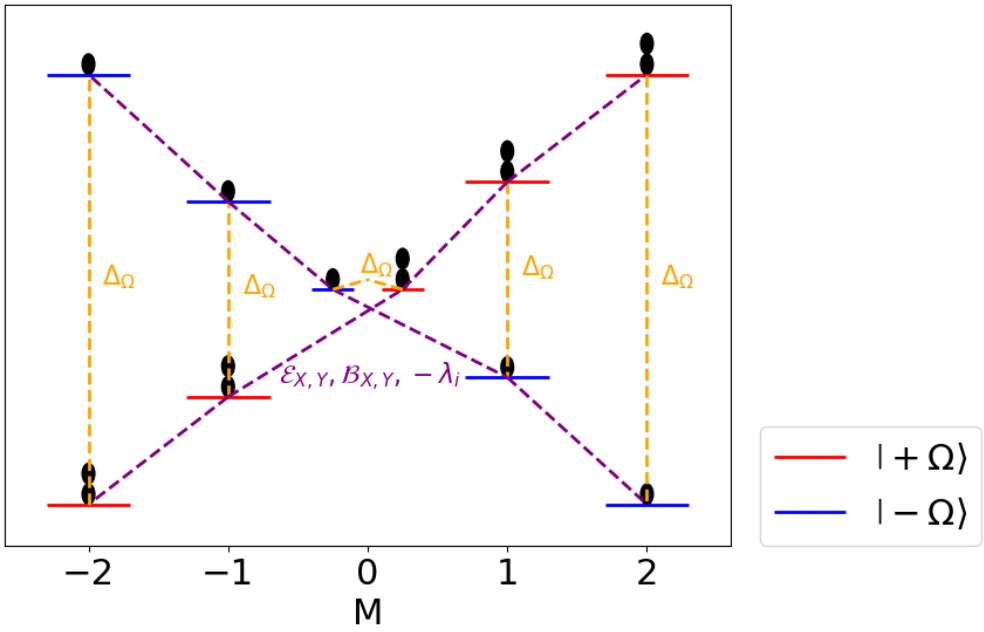

Figure 4.10: Molecular states in the $|J, \Omega, M\rangle$ basis. Interactions that mix Ω and M are shown in orange and purple, respectively: orange, Δ_Ω, the Ω-doublet splitting; purple, $\mathcal{E}_{X,Y}$ and $\mathcal{B}_{X,Y}$, the transverse electric and magnetic fields; also purple, $-\lambda_i$, the quantization-axis rotation rate describing the molecules' inability to adiabatically follow back to the same Z axis after traversing different experimental regions. The approximate final Q-state population distribution is indicated by black circles. See chapter 5 for a detailed explanation.

controlling the polarization of the STIRAP beams or by varying the magnetic field in the lens STIRAP region, we varied the initial population imbalance in $Q, M = \pm 2$ and measured the corresponding eEDM systematic error. The result showed that the lens STIRAP process was capable of generating a gigantic systematic error as large as 2 mrad/s, which is two orders of magnitude larger than our final statistical uncertainty.

To mitigate this systematic error, we implemented a new state-preparation step after the electrostatic lens to pump all residual Q-state population into the X state. Our original plan for this $Q \mapsto X$ step was to use another STIRAP. However, the residual Q-state population after STIRAP would be much higher than that obtained with an optical pumping step. This optical pumping step, which removes the residual Q-state population, is called the Qleanup step and is explained in detail in 2.3.5. In addition, instead of preparing the molecules in the

$Q, M = 2$ lensing state, we now prepare them in a balanced superposition of the $Q, M = \pm 2$ lensing states to further suppress possible systematic errors from this decay process. Lastly, we found that the decay systematics scale with the retardance error of the readout half-wave plate. We therefore ordered custom half-wave plates at our operating wavelength of 703 nm to further suppress this systematic error. Together, these mitigation strategies reduce the decay systematic error to below our final statistical uncertainty according to the IPV procedure.

The exact mechanism responsible for the systematic errors observed in the IPV measurement is still not fully understood. Because the size of the systematic error depends on the retardance error of the readout half-wave plate, we suspect that the effect does not arise from a true phase shift, but instead from an effective ellipticity imperfection in the readout process. The apparent phase that we extract contains a component that behaves like an ellipticity–ellipticity coupling, $d\Theta_{\mathrm{readout}}\, d\Theta_{\mathrm{ThO}}\, \tilde{\mathcal{P}}\tilde{\mathcal{R}}$, where $d\Theta_{\mathrm{readout}}$ is the ellipticity imperfection of the readout optics and $d\Theta_{\mathrm{ThO}}$ is the ellipticity imperfection of the ThO molecular state itself. We know that the lens STIRAP can generate a $d\Theta_{\mathrm{ThO}}^{\mathcal{NE}}$ component. Collin Diver has proposed that the readout process also contains an effective $d\Theta_{\mathrm{readout}}\tilde{\mathcal{P}}\tilde{\mathcal{R}}$ component, which can lead to an apparent phase $\phi_{\mathrm{sys}}^{\mathcal{NE}} \propto d\Theta_{\mathrm{readout}} \cdot d\Theta_{\mathrm{ThO}}^{\mathcal{NE}}$. However, the exact origin of this effective ellipticity in the readout process is still not fully understood. At the time of writing, we are prioritizing beam time for characterizing other systematic errors, as this particular effect has already been bounded by the IPV procedure.

Could the decay process directly give rise to phase systematic errors? This is, in principle, possible because the $Q, M = \pm 1 \mapsto H, M = \pm 1$ decay proceeds only via Z-polarized light. Therefore, the relative phase information between the $M = \pm 1$ states in the Q state is preserved through the decay process[8]. We give a simple example of how phase systematic

8. This phase preservation does not generally apply to all systems; the frequency splitting between the $M = \pm 1$ states must be smaller than the linewidth of the transition for phase preservation to hold. In our case, the interaction length is only 5 ms, and the resulting transit-time broadening enables phase preservation. Physically, the spread in decay time is not sufficient to fully scramble the relative phase between the two M

errors can arise from this decay process in chapter 5. However, we have not observed any evidence for phase systematic errors from this decay process in our IPV procedure.

The discovery of this decay systematic error could have important implications for future EDM experiments. If the decay lifetime is much longer than the interaction time, the decay can appear as a continuously accumulating signal that may be indistinguishable from a genuine precession frequency shift. Additionally, the extra internal degrees of freedom available in molecules make it difficult to achieve complete quantum control and to identify all possible metastable states. In our experiment, for example, we know that there are still residual decaying populations in vibrationally excited states of both the H and Q states. Due to practical constraints, we have not rigorously characterized these states, although we have good reasons to believe that their contributions are small. Future experiments will need to be aware of such decay channels as they push toward even higher precision.

4.4.7 STIRAP Stark Interference

The Stark interference systematic error occurs when we intend to drive an $E1$ transition but simultaneously drive an $M1$ transition due to constraints imposed by Maxwell's equations. The mechanism was first worked out by Brendon O'Leary in the context of ACME I [87]. It turns out that this particular systematic error is maximized when the laser is sent in a direction orthogonal to the lab Z quantization axis. This is exactly the geometry used in our vertical STIRAP setup. Therefore, we focus the discussion of this systematic error on the vertical STIRAP configuration.

We will not go through a detailed derivation of this systematic error, as it has already been presented in [87]. Instead, we provide a qualitative explanation based on molecular symmetry. Recall that the eEDM Hamiltonian term is an Ω-correlated energy shift. Therefore, any Ω-correlated phase shift can mimic an eEDM signal. Recall also that in the presence

states through Zeeman precession.

of non-zero electric and magnetic fields, molecular states with opposite signs of Ω respond differently, as shown in figure 2.4. Because of this difference in response, when $E1$ and $M1$ transitions are simultaneously driven, the interference between these two transition amplitudes can differ for the two Ω states, thereby preparing an Ω-correlated phase during the STIRAP process. This is the essence of the Stark interference systematic error.

The mechanism derived by Brendon O'Leary depends on the ellipticity of the 1090 nm Stokes beam. We indeed observe this behavior when we intentionally vary the ellipticity of the Stokes beam. However, we also scanned additional parameters, such as the linear polarization angles of the pump and Stokes beams, and found further dependences that are not explained by O'Leary's model. This scanning capability was newly implemented by Collin Diver, and therefore these additional dependences were not observed in ACME II.

To model these additional dependences, we built a full numerical simulation of the vertical STIRAP process including both $E1$ and $M1$ transitions. All 13 states in the X, C, and H manifolds involved in the STIRAP process were included. The $M1$ transition amplitudes are not experimentally known, so we used theoretical values from the literature [100]. The simulation successfully reproduced the additional dependences to within factors of order unity, which we attribute to uncertainties in the exact STIRAP parameters. The simulation allowed us to identify three additional pathways that can lead to Stark interference systematic errors, as shown in figure 4.11.

The mitigation strategies for this systematic error are similar to those used in ACME II. We use the cleanup beam to dissipatively remove any systematic phase generated by the vertical STIRAP process. In addition, we carefully align the STIRAP waveplates to minimize the systematic phase even in the absence of the cleanup. We will also monitor this systematic error throughout the entire final eEDM data run using our IPV procedure.

Why do the cleanup and readout lasers not suffer significantly from this Stark interference systematic error? This can again be explained by a symmetry argument. When the

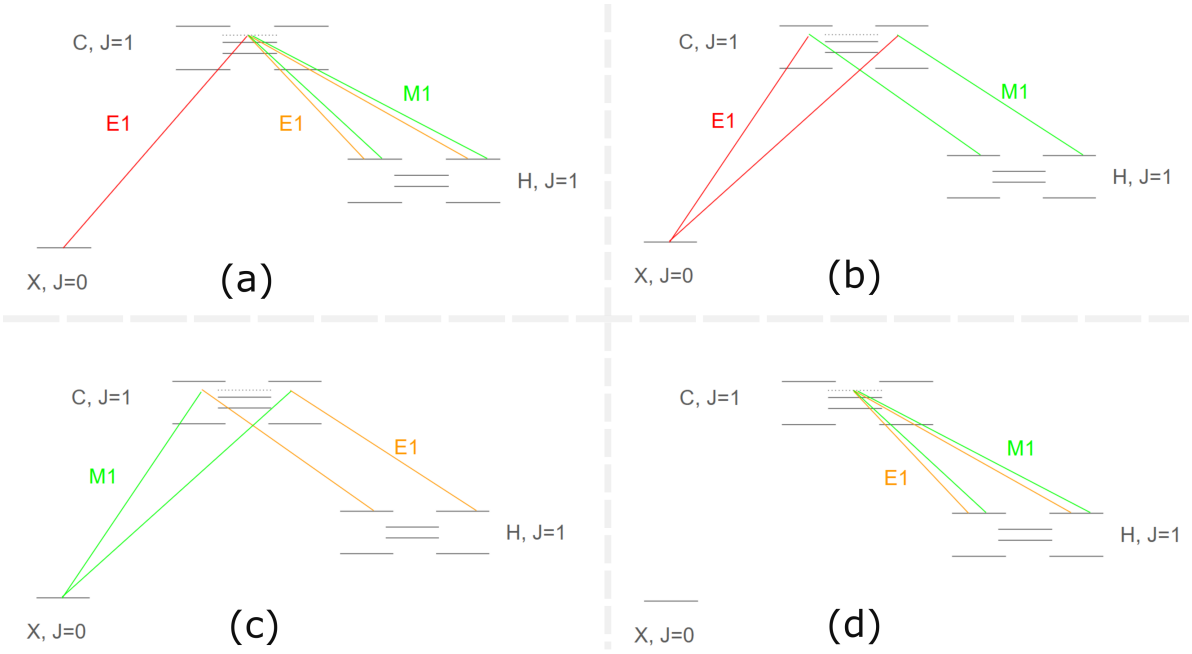

Figure 4.11: (a) Stark interference path worked out by Brendon O'Leary. The $E1$ and $M1$ are between the same pair of states and lead to different dark states depending on the sign of Ω. (b) additional pathway where the pump is driven by $E1$ and the Stokes is driven by $M1$. (c) additional pathway where the pump is driven by $M1$ and the Stokes is driven by $E1$. (d) additional pathway describing a two photon $E1-M1$ transition within H through a virtual intermediate state in C. While (a) is a bona fide STIRAP process, (b)-(d) are better described as Raman processes that prepare an Ω correlated phase. All these pathways can lead to Stark interference systematic errors.

laser propagates along the lab Z axis, without loss of generality, suppose the $E1$ transition is driven by X-polarized light and the corresponding $M1$ transition is driven by Y-polarized light. The $E1$ and $M1$ transitions differ by parity, and recall that a change in parity in the cleanup and readout processes effectively corresponds to a 90° rotation in polarization. As a result, the $M1$ transition attempts to drive the same transition as the $E1$ transition, which prevents a Stark interference phase from being generated.

Naively, this argument might suggest that Stark interference would instead produce an Ω- or $\mathcal{NE}$-correlated Rabi frequency in the cleanup and readout processes. However, this possibility was ruled out by Brendon O'Leary's detailed calculation [87]. Briefly, the key point is that the $E1$ and $M1$ matrix elements differ by a factor of i, analogous to the σ_x

and σ_y Pauli matrices. Their amplitudes therefore add in quadrature rather than linearly, preventing the generation of any correlated Rabi frequency $\Omega_r^{\mathcal{NE}}$. Equivalently, the $E1$ and $M1$ drives are 90° out of phase with each other.

4.4.8 Mysterious $\Omega_r^{\mathcal{NE}}$

In ACME I, there was evidence suggesting the presence of a mysterious $\mathcal{NE}$-correlated Rabi frequency $\Omega_r^{\mathcal{NE}}$ [53, 87]. After switching the readout transition from 1090 nm to 703 nm in ACME II, the evidence for this effect became much weaker [90, 91]. Although there was no evidence that this effect led to systematic errors, ACME II nevertheless used the IPV procedure to place a bound on possible systematic errors induced by it [23].

In ACME III, we further studied this effect using our improved statistical sensitivity. We defer the detailed mathematical description to sections 4.5.2 and 4.5.3. Briefly, if the cleanup detuning is not set exactly to zero, a correlated detuning $\Delta^{\mathcal{NE}}$ can be misinterpreted as a correlated Rabi frequency $\Omega_r^{\mathcal{NE}}$ within an AC Stark shift model.

Using microwave spectroscopy, we reduced the non-reversing electric field and hence $\Delta^{\mathcal{NE}}$ from around 5 kHz to below 1 kHz. We then measured $\Omega_r^{\mathcal{NE}}$ using a procedure very similar to that used in ACME II. We found that the measured $\Omega_r^{\mathcal{NE}}$ became consistent with zero within our statistical uncertainty. This suggests that we do not observe the same mysterious $\Omega_r^{\mathcal{NE}}$ effect as in ACME I. However, this effect is classified as a class C systematic error and satisfies our inclusion criteria. Therefore, we will still include it in the final systematic error budget using the same IPV procedure as in ACME II.

4.4.9 Driving transitions between unwanted states

Molecules offer a rich set of internal states that can be utilized for precision measurements. However, additional degrees of freedom also mean more possible unwanted states that can be accidentally driven by our lasers. One of the most obvious unwanted states is the

opposite-parity state for the readout transition. As a quick reminder, the readout transition excites the molecules to the $I, M = 0$ state with even or odd parity depending on the choice of switch state $\tilde{\mathcal{P}}$. These parity states are separated by only about 91 MHz, and whenever we excite one, we inadvertently excite the other off-resonantly at some small level. Possible AC Stark shift systematics from this opposite-parity state can arise and were first worked out by Brendon O'Leary [87]. The expected effect size is smaller than our final statistical uncertainty and is therefore not a source of concern. We include a detailed derivation in 4.5.4, as it was not included in previous ACME publications.

In ACME III, we have seen evidence that additional unwanted states are being driven by our lasers. At the time of writing, we are still actively investigating the exact mechanism, so we only provide a brief summary here. During the IPV tests, we noticed a non-zero $\phi^{\mathcal{N}}$ and an oriented, $\mathcal{N}\mathcal{E}$-correlated signal that increases as we increase the saturation of the readout laser. The non-zero $\phi^{\mathcal{N}}$ was later understood by Collin Diver to arise from unwanted transitions from $H, M = 0$ to $I, M = \pm 1$ states that are off-resonantly driven by the readout laser. The oriented $\mathcal{N}\mathcal{E}$-correlated signal is still not fully understood at the time of writing. However, there is evidence suggesting that, depending on the sign of $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$, scattered photons preferentially travel toward one Z direction compared to the other at the $\sim 0.1\%$ level. We suspect this could be due to interference effects in the unwanted transitions (e.g. $E1$ decay interfering with $M1$ decay), leading to a small directional bias in the fluorescence pattern.

Due to the presence of these unwanted transitions, we have decided to implement a slow switch on the readout laser power. By reducing the laser power by half, for example, we expect the effect from most unwanted transitions to be reduced by a similar factor. We believe this will provide a reliable way to bound the associated systematic errors using our slow-switch procedure.

4.5 Systematic error models

In this section, we present detailed models of various AC Stark shift systematic errors in ACME. We first derive the AC Stark shift in the unsaturated regime in Sec. 4.5.1. This treatment applies to weak reflection beams whose effects can be handled perturbatively. The non-perturbative case is more involved and is addressed in Sec. 4.5.2 using time-dependent perturbation theory. This latter analysis not only explains how a correlated detuning $\Delta^{\mathcal{NE}}$ can produce systematic errors, but also shows how the same effect can be misinterpreted as a correlated Rabi frequency $\Omega_r^{\mathcal{NE}}$, as discussed in Sec. 4.5.3. Finally, we derive the AC Stark shift systematic arising from driving the opposite-parity state in Sec. 4.5.4.

4.5.1 AC Stark shift from the reflection beam

How the presence of a correlated detuning $\Delta^{\mathcal{NE}}$ can lead to AC Stark shift systematics is best explained in the unsaturated regime, where time-independent perturbation theory can be applied. The laser coupling diagram is shown in figure 4.5. Let the ellipticity of the interaction laser be described by $\Theta = \frac{\pi}{4} + d\Theta$. The Hamiltonian of this three-level Λ system in the $\{|I, M=0\rangle, |H, M=+1\rangle, |H, M=-1\rangle\}$ basis is given by

$$H = \begin{pmatrix} \Delta - i\gamma/2 & \frac{\Omega_r}{2}\cos\Theta & \frac{\Omega_r}{2}\sin\Theta \\ \frac{\Omega_r}{2}\cos\Theta & 0 & 0 \\ \frac{\Omega_r}{2}\sin\Theta & 0 & 0 \end{pmatrix} \begin{matrix} |I, M=0\rangle \\ |H, M=+1\rangle \\ |H, M=-1\rangle \end{matrix} \tag{4.15}$$

$$\approx \begin{pmatrix} \Delta - i\gamma/2 & \frac{\Omega_r}{2\sqrt{2}}(1 - d\Theta) & \frac{\Omega_r}{2\sqrt{2}}(1 + d\Theta) \\ \frac{\Omega_r}{2\sqrt{2}}(1 - d\Theta) & 0 & 0 \\ \frac{\Omega_r}{2\sqrt{2}}(1 + d\Theta) & 0 & 0 \end{pmatrix}, \tag{4.16}$$

where Δ is the detuning of the laser from resonance, Ω_r is the Rabi frequency of the laser, and γ is the decay rate of the excited state. Using time-independent perturbation theory, the AC Stark shift of each of the $M = \pm 1$ states is given by

$$\Delta E_{M=\pm 1} = \frac{\Omega_r^2}{8(\Delta - i\gamma/2)}(1 \mp d\Theta)^2. \tag{4.17}$$

In particular, we are only interested in the differential AC Stark shift between the two M states, which is given by

$$\Delta E_{M=+1} - \Delta E_{M=-1} = \frac{\Omega_r^2}{8(\Delta - i\gamma/2)}\left[(1 - d\Theta)^2 - (1 + d\Theta)^2\right] \tag{4.18}$$

$$\approx -\frac{\Omega_r^2 d\Theta}{2(\Delta - i\gamma/2)} \tag{4.19}$$

$$\approx -\frac{\Omega_r^2 d\Theta}{-2i\gamma/2}\left(1 + \frac{\Delta}{i\gamma/2}\right) \tag{4.20}$$

$$\approx i\frac{\Omega_r^2 d\Theta}{\gamma} - \frac{2\Omega_r^2 d\Theta \Delta}{\gamma^2}, \tag{4.21}$$

where we have assumed that the detuning is small compared to the decay rate, $\Delta \ll \gamma$. The first term is imaginary and describes differential decay between the two M states, while the second term is real and describes the differential AC Stark shift. Ignoring the dissipative term, this perturbative calculation shows that the differential AC Stark shift is proportional to both the detuning Δ and the ellipticity $d\Theta$ of the laser beam. Therefore, if there is a correlated detuning $\Delta^{\mathcal{NE}}$ together with a non-zero ellipticity $d\Theta$, a $\mathcal{NE}$-correlated phase shift will be induced when the molecules interact with this laser beam. This calculation is only valid in the unsaturated regime where time-independent perturbation theory applies. In the actual experiment, this regime is only relevant for weak reflection beams. To model the systematics from the main cleanup beam, time-dependent perturbation theory is required.

4.5.2 AC Stark shift from the cleanup beam

The derivation of the AC Stark shift systematics from the cleanup beam is much more involved due to the time dependence of the problem. We follow the formulation laid out by Brendon O'Leary [87], attempting to fill in details that may have been obvious to him but not to the author. The basic idea is to apply time-dependent perturbation theory to the same three-level Λ system formed by the $\{|I, M = 0\rangle, |H, M = +1\rangle, |H, M = -1\rangle\}$ states. To do so, we first transform the Hamiltonian such that it is approximately diagonal. This is achieved through two transformations. The first transformation is to the bright/dark basis defined by the laser polarization, and the second transformation is to the dressed-state basis, or instantaneous eigenstate basis of the time-dependent Hamiltonian. The different bases are illustrated in figure 4.12.

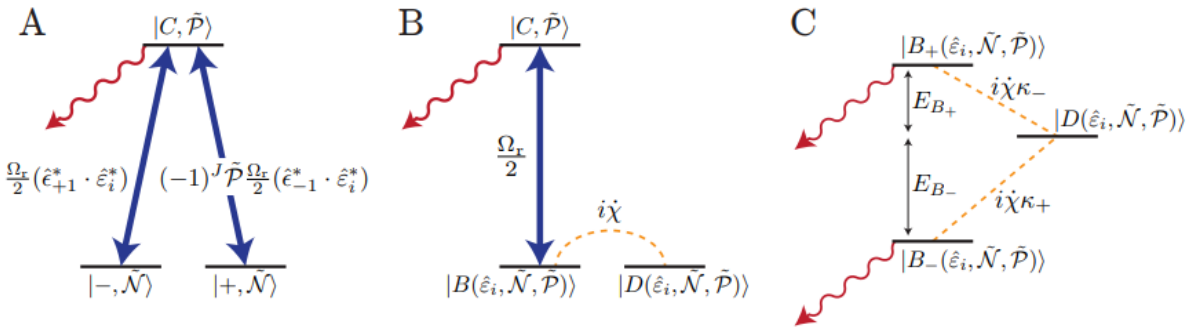

Figure 4.12: Energy diagram for three different bases used to understand the AC Stark shift systematics. A: $M = \pm 1$ basis. B: bright/dark basis defined by the laser polarization. $\dot{\chi}$ is the rate of change of the laser polarization. C: dressed-state basis, or instantaneous eigenstate basis of the time-dependent Hamiltonian. Figure reproduced from Brendon O'Leary's thesis [87]. For ACME I, the excited state used was C instead of I as in our case, but the structure is exactly the same. Additional state labels $\tilde{\mathcal{P}}, \tilde{\mathcal{N}}$ are added.

As a quick reminder, the laser polarization can be expressed in the spherical basis as

$$\hat{\epsilon} = -e^{-i\theta}\cos\Theta\hat{\epsilon}_{+1} + e^{i\theta}\sin\Theta\hat{\epsilon}_{-1}, \tag{4.22}$$

The initial Hamiltonian for this polarization is given by

$$\tilde{H} = \begin{pmatrix} \Delta_I & \frac{\Omega_r}{2}\left(\hat{\epsilon}_{+1}^* \cdot \hat{\epsilon}^*\right) & (-1)^J \tilde{\mathcal{P}} \frac{\Omega_r}{2}\left(\hat{\epsilon}_{-1}^* \cdot \hat{\epsilon}^*\right) \\ \frac{\Omega_r}{2}\left(\hat{\epsilon}_{+1} \cdot \hat{\epsilon}\right) & -\delta & 0 \\ (-1)^J \tilde{\mathcal{P}} \frac{\Omega_r}{2}\left(\hat{\epsilon}_{-1} \cdot \hat{\epsilon}\right) & 0 & \delta \end{pmatrix} \begin{matrix} |I, \tilde{\mathcal{P}}\rangle \\ |H, M=-1, \tilde{\mathcal{N}}\rangle, \\ |H, M=+1, \tilde{\mathcal{N}}\rangle \end{matrix} \tag{4.23}$$

where $\Delta_I = \Delta - i\gamma/2$ is a complex-valued detuning describing the combined effect of the single-photon detuning Δ and excited-state decay γ, δ is the two-photon detuning due to the Zeeman shift, and Ω_r is the Rabi frequency. As a reminder, $\hat{\epsilon}_{\pm 1}$ denote right- and left-handed circular polarizations. Following Brendon O'Leary, we add the switch-state labels $\tilde{\mathcal{P}}, \tilde{\mathcal{N}}$ to the states. However, they are cumbersome, and for notational simplicity we define $|I\rangle \equiv |I, \tilde{\mathcal{P}}\rangle$ and $|\pm\rangle \equiv |H, M = \pm 1, \tilde{\mathcal{N}}\rangle$ for the remainder of this section.

We first perform a basis transformation to the bright/dark basis defined by the laser polarization using the unitary transformation

$$U^\dagger = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \langle -|B\rangle & \langle -|D\rangle \\ 0 & \langle +|B\rangle & \langle +|D\rangle \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \dfrac{|\varepsilon_{+1}|}{\varepsilon_{+1}\sqrt{\dfrac{|\varepsilon_{-1}|^2}{|\varepsilon_{+1}|^2}+1}} & -\dfrac{\varepsilon_{-1}}{\varepsilon_{+1}\sqrt{\dfrac{|\varepsilon_{-1}|^2}{|\varepsilon_{+1}|^2}+1}} \\ 0 & \dfrac{\varepsilon_{-1}^*}{|\varepsilon_{+1}|\sqrt{\dfrac{|\varepsilon_{-1}|^2}{|\varepsilon_{+1}|^2}+1}} & \dfrac{1}{\sqrt{\dfrac{|\varepsilon_{-1}|^2}{|\varepsilon_{+1}|^2}+1}} \end{pmatrix}. \tag{4.24}$$

where, for simplicity, we have introduced two new variables

$$\varepsilon_{+1} \equiv \left(\hat{\epsilon}_{+1}^* \cdot \hat{\epsilon}^*\right), \tag{4.25}$$

$$\varepsilon_{-1} \equiv (-1)^J \tilde{\mathcal{P}}\left(\hat{\epsilon}_{-1}^* \cdot \hat{\epsilon}^*\right). \tag{4.26}$$

The transformed Hamiltonian via $\tilde{H} = UHU^\dagger - iU\dot{U}^\dagger$ is then given by

$$\tilde{H} = \begin{pmatrix} \Delta_I & \frac{1}{2}\Omega_\mathrm{r} & 0 \\ \frac{1}{2}\Omega_\mathrm{r} & 0 & i\dot{\chi} \\ 0 & -i\dot{\chi}^* & 0 \end{pmatrix} \begin{matrix} |I\rangle \\ |B(\hat{\epsilon})\rangle, \\ |D(\hat{\epsilon})\rangle \end{matrix} \tag{4.27}$$

where $|B(\hat{\epsilon})\rangle$ is the bright state defined by the polarization $\hat{\epsilon}$, $|D(\hat{\epsilon})\rangle$ is the dark state, and $\dot{\chi} \equiv \dot{\Theta} - i(\dot{\theta} + \delta)$ describes the rate of change of the laser polarization via equation 4.22. Physically, $\dot{\chi}$ describes the rotation of the bright and dark states into each other. For this rotation, the effect of the Zeeman interaction is equivalent to a rotation of the laser polarization.

The Hamiltonian in equation 4.27 is clearly still not diagonal. However, it is straightforward to approximately diagonalize it by transforming to the dressed-state basis, assuming that $\dot{\chi}$ is small compared to the other energy scales. The dressed-state basis, which defines the unitary transformation W, is given by

$$|B_\pm\rangle \equiv \pm\kappa_\pm|C\rangle + \kappa_\mp|B(\hat{\epsilon})\rangle, \qquad |D\rangle \equiv |D(\hat{\epsilon})\rangle, \tag{4.28}$$

where the mixing amplitudes are defined as

$$\kappa_\pm \equiv \frac{1}{\sqrt{2}}\sqrt{1 \pm \frac{\Delta_I}{\sqrt{\Delta_I^2 + \Omega_\mathrm{r}^2}}}. \tag{4.29}$$

The transformed Hamiltonian via $\tilde{H} = WHW^\dagger - iW\dot{W}^\dagger$ is then given by

$$\tilde{H} = \begin{pmatrix} 0 & -i\dot{\chi}^*\kappa_+ & -i\dot{\chi}^*\kappa_- \\ i\dot{\chi}\kappa_+ & E_{B-} & -\frac{i}{2}\frac{\dot{\Omega}_\mathrm{r}\Delta_I}{\Delta_I^2+\Omega_\mathrm{r}^2} \\ i\dot{\chi}\kappa_- & \frac{i}{2}\frac{\dot{\Omega}_\mathrm{r}\Delta_I}{\Delta_I^2+\Omega_\mathrm{r}^2} & E_{B+} \end{pmatrix} \begin{matrix} |D\rangle \\ |B_-\rangle, \\ |B_+\rangle \end{matrix} \tag{4.30}$$

with diagonal terms

$$E_{B\pm} = \frac{1}{2}\left(\Delta_I \pm \sqrt{\Delta_I^2 + \Omega_{\mathrm{r}}^2}\right), \qquad E_D = 0. \tag{4.31}$$

This Hamiltonian is now nearly diagonal, with only small off-diagonal terms. We are thus ready to apply time-dependent perturbation theory to solve for the time evolution of

$$|\psi(t)\rangle = c_D(t)|D\rangle + c_{B+}(t)\,|B_+\rangle + c_{B-}(t)\,|B_-\rangle\,. \tag{4.32}$$

The initial conditions are

$$c_D(0) = 1, \qquad c_{B+}(0) = c_{B-}(0) = 0, \tag{4.33}$$

because the cleanup removes essentially all previous history of the molecules, such that they start in the dark state. Using first-order time-dependent perturbation theory, the solution is

$$c_D(t) = c_D(0) - \sum_{\pm}\int_0^t \dot{\chi}^*\big(t'\big)\,\kappa_{\mp}\big(t'\big)\,e^{-iE_{B\pm}t'}c_{B\pm}(0)\,dt', \tag{4.34}$$

$$c_{B\pm}(t) = e^{-iE_{B\pm}t}c_{B\pm}(0) + e^{-iE_{B\pm}t}\int_0^t \dot{\chi}\big(t'\big)\,\kappa_{\mp}\big(t'\big)\,e^{iE_{B\pm}t'}c_D(0)\,dt'. \tag{4.35}$$

This solution is written in the dressed-state basis. For interpretation, we transform back to the bright/dark basis by applying the inverse of W,

$$c_B(t) = \sum_{\pm} e^{-iE_{B\pm}t}\int_0^t dt'\,\dot{\chi}\big(t'\big)\,e^{iE_{B\pm}t'}\,(\kappa_{\mp})^2 \tag{4.36}$$

$$= \frac{d\chi}{t}\sum_{\pm} e^{-iE_{B\pm}t}\int_0^t dt'\,t' e^{iE_{B\pm}t'}\,(\kappa_{\mp})^2 \tag{4.37}$$

$$= d\chi\,\Pi, \tag{4.38}$$

with

$$\Pi \equiv \sum_{\pm} (\kappa_{\mp})^2 \, e^{-iE_{B\pm}t/2} \, \mathrm{sinc}(E_{B\pm}t/2) \,, \tag{4.39}$$

where we have assumed that $\dot{\chi} = d\chi/t$ is constant over the laser profile and $d\chi$ is the total change in the polarization parameter χ. To summarize, after an interaction time t with the cleanup laser, the small admixture of the bright state acquired by the molecules is

$$c_B(t) = d\chi \, \Pi, \tag{4.40}$$

where Π is a complex optical pumping factor and $d\chi$ is the total change in laser polarization over the interaction time.

It turns out that the phase induced by this small admixture is simply given by $\mathrm{Im}(c_B)$, with the proof derived in appendix D.

$$\begin{aligned}
\phi_{\text{AC Stark}} &= \mathrm{Im}(c_B) && (4.41)\\
&= \mathrm{Im}(d\chi \, \Pi) && (4.42)\\
&= \mathrm{Im}[(d\Theta - i(\delta \cdot t + d\theta))(\mathrm{Re}\,\Pi + i\,\mathrm{Im}\,\Pi)] && (4.43)\\
&= d\Theta \,\mathrm{Im}\,\Pi + (\delta \cdot t + d\theta)\,\mathrm{Re}\,\Pi. && (4.44)
\end{aligned}$$

Recall that we set out to derive the systematic phase shift induced by a correlated detuning $\Delta^{\mathcal{NE}}$. If the optical pumping factor Π is completely even in detuning, then no $\phi^{\mathcal{NE}}_{\text{AC Stark}}$ can appear after the parity transformation. However, the imaginary part of Π is completely odd in detuning, while the real part is completely even. This means that we obtain a systematic

phase shift from the first term,

$$\phi^{\mathcal{NE}}_{\text{AC Stark}} = d\Theta \operatorname{Im}\Pi^{\mathcal{NE}} \tag{4.45}$$

$$\approx \frac{d\Theta}{dX}\sigma_X \cdot \left.\frac{\partial \operatorname{Im}\Pi}{\partial \Delta}\right|_{\Delta=0} \Delta^{\mathcal{NE}} \tag{4.46}$$

$$= \frac{d\Theta}{dX}\sigma_X \cdot \frac{\Delta^{\mathcal{NE}}}{\sigma_\Pi}, \tag{4.47}$$

$$\sigma_\Pi \equiv \left(\left.\frac{\partial \operatorname{Im}\Pi}{\partial \Delta}\right|_{\Delta=0}\right)^{-1}, \tag{4.48}$$

where $\frac{d\Theta}{dX}$ is the ellipticity gradient of the laser beam, σ_X is the characteristic length scale of the laser beam, and σ_Π is the characteristic detuning scale of the optical pumping factor.

This is the main takeaway. Recall that $d\Theta$ is the change in ellipticity over the laser profile. We thus obtain a functional form very similar to that of the unsaturated case discussed in 4.5.1, with the important distinction that it is the change in ellipticity over the laser profile that matters, rather than the absolute ellipticity. For the cleanup beam, this means that the birefringence gradient matters, not the absolute birefringence.

While we used first-order time-dependent perturbation theory to derive this result, series-expansion results for many laser profiles were worked out by Brendon O'Leary [87]. An exact proof of this functional dependence using generalized symmetry arguments is given in appendix E.

4.5.3 Correlated Rabi frequency

The formulation developed in the previous subsection can also be used to understand the effect of a correlated Rabi frequency $\Omega_r^{\mathcal{NE}}$. The systematic phase is given by

$$\phi_{\text{AC Stark}} = d\Theta \operatorname{Im}\Pi + (\delta \cdot t + d\theta) \operatorname{Re}\Pi, \tag{4.49}$$

where Π is the complex optical pumping factor defined previously, $d\Theta$ is the change in ellipticity over the laser profile, δ is the two-photon detuning due to the Zeeman shift, t is the characteristic interaction time with the laser, and $d\theta$ is the change in the linear polarization angle over the laser profile.

The real part of Π can be expanded as

$$\mathrm{Re}\,\Pi^{\mathcal{NE}} = \left.\frac{\partial\,\mathrm{Re}\,\Pi}{\partial\Omega_r}\right|_{\Omega_r=\Omega_r^0} \Omega_r^{\mathcal{NE}} + \left.\frac{\partial^2\,\mathrm{Re}\,\Pi}{\partial\Delta^2}\right|_{\Omega_r=\Omega_r^0} \Delta^{\mathrm{nr}}\Delta^{\mathcal{NE}} + \ldots, \tag{4.50}$$

which to first order mainly depends on the correlated Rabi frequency $\Omega_r^{\mathcal{NE}}$. If the one-photon non-reversing detuning Δ^{nr} is nonzero, there is also a second-order contribution from the correlated detuning $\Delta^{\mathcal{NE}}$.

This correlated optical pumping factor can couple with the reversing and non-reversing magnetic fields to produce $\mathcal{NEB}$ and $\mathcal{NE}$ contributions given by

$$\phi^{\mathcal{NEB}}_{\text{AC Stark}} = \delta^{\mathcal{B}} \cdot t \cdot \mathrm{Re}\,\Pi^{\mathcal{NE}}, \tag{4.51}$$

$$\phi^{\mathcal{NE}}_{\text{AC Stark}} = \delta^{\mathrm{nr}} \cdot t \cdot \mathrm{Re}\,\Pi^{\mathcal{NE}}. \tag{4.52}$$

This common dependence on the optical pumping factor $\mathrm{Re}\,\Pi$ allows us to bound the contribution to the $\mathcal{NE}$ channel using the measured $\mathcal{NEB}$ channel, together with measured or bounded values of the Zeeman shifts $\delta^{\mathcal{B}}$ and δ^{nr}.

Note, however, that $\Delta^{\mathcal{NE}}$ contributes to $\mathrm{Re}\,\Pi^{\mathcal{NE}}$ if the one-photon detuning is not zero. Based on our IPV measurements, we believe this to be the dominant contribution to $\mathrm{Re}\,\Pi^{\mathcal{NE}}$ in our experiment.

4.5.4 *Opposite parity state*

We would like to understand how the presence of an opposite-parity state with $M = 0$ can lead to possible systematics in the eEDM measurement. We therefore begin with the rotating-frame Hamiltonian of the laser–molecule interaction that includes both parities of the excited I state. We denote the parity of the on-resonant state as $\tilde{\mathcal{P}}$ and the parity of the off-resonant state as $-\tilde{\mathcal{P}}$.

$$H_{\text{rot}} = \begin{pmatrix} \Delta_* + \tilde{\mathcal{P}}\Delta_I & 0 & \frac{-ie^{-i\theta}\Omega_r \cos\Theta}{2} & \frac{ie^{-i\theta}\Omega_r \sin\Theta}{2} \\ 0 & \Delta_* & \frac{e^{-i\theta}\Omega_r \cos\Theta}{2} & \frac{e^{-i\theta}\Omega_r \sin\Theta}{2} \\ \frac{ie^{i\theta}\Omega_r \cos\Theta}{2} & \frac{e^{i\theta}\Omega_r \cos\Theta}{2} & 0 & 0 \\ \frac{-ie^{i\theta}\Omega_r \sin\Theta}{2} & \frac{e^{i\theta}\Omega_r \sin\Theta}{2} & 0 & 0 \end{pmatrix} \begin{matrix} |I, -\tilde{\mathcal{P}}, M = 0\rangle \\ |I, \tilde{\mathcal{P}}, M = 0\rangle \\ |H, M = +1\rangle \\ |H, M = -1\rangle \end{matrix}, \quad (4.53)$$

where $\Delta_* \equiv \Delta - i\gamma_I/2$ is the complex detuning including the actual detuning Δ and decay from the I state at rate γ_I, Δ_I is the Ω-doublet splitting of the I state, Ω_r is the Rabi frequency, and Θ, θ are the laser polarization parameters. We keep Δ_* until the end before separating its real and imaginary parts.

Next, we transform to the bright and dark state basis via

$$U = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & e^{-i\theta}\cos\Theta & e^{i\theta}\sin\Theta \\ 0 & 0 & -e^{-i\theta}\sin\Theta & e^{i\theta}\cos\Theta \end{pmatrix}. \quad (4.54)$$

The transformed Hamiltonian $UHU^\dagger$ is

$$H_{BD} = \begin{pmatrix} \Delta_* + \tilde{\mathcal{P}}\Delta_I & 0 & -\frac{i\Omega_r}{2}\cos 2\Theta & \frac{i\Omega_r}{2}\sin 2\Theta \\ 0 & \Delta_* & \frac{\Omega_r}{2} & 0 \\ \frac{i\Omega_r}{2}\cos 2\Theta & \frac{\Omega_r}{2} & 0 & 0 \\ -\frac{i\Omega_r}{2}\sin 2\Theta & 0 & 0 & 0 \end{pmatrix} \begin{matrix} |I, -\tilde{\mathcal{P}}, M=0\rangle \\ |I, \tilde{\mathcal{P}}, M=0\rangle \\ |H, B\rangle \\ |H, D\rangle \end{matrix} . \tag{4.55}$$

Notably, the dark state is no longer fully dark due to its coupling to the opposite $-\tilde{\mathcal{P}}$ parity state. In addition, if the light has nonzero ellipticity, i.e. $\Theta \neq \frac{\pi}{4}$, the opposite-parity state acquires an admixture of the bright state through the $\cos 2\Theta$ term. Since the dark state acquires an admixture of the opposite-parity state via off-resonant coupling, it therefore acquires a small admixture of the bright state through second-order perturbation.

To work out the exact second-order contribution, we first transform to the dressed-state basis via

$$U = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \kappa_+ & \kappa_- & 0 \\ 0 & -\kappa_- & \kappa_+ & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{4.56}$$

and the transformed Hamiltonian $UHU^\dagger$ is given by

$$H_{\text{dressed}} = \begin{pmatrix} \Delta_* + \tilde{\mathcal{P}}\Delta_I & -\frac{i\Omega_r\kappa_-\cos 2\Theta}{2} & -\frac{i\Omega_r\kappa_+\cos 2\Theta}{2} & \frac{i\Omega_r\sin 2\Theta}{2} \\ \frac{i\Omega_r\kappa_-\cos 2\Theta}{2} & E_{B_+} & 0 & 0 \\ \frac{i\Omega_r\kappa_+\cos 2\Theta}{2} & 0 & E_{B_-} & 0 \\ -\frac{i\Omega_r\sin 2\Theta}{2} & 0 & 0 & 0 \end{pmatrix} \begin{matrix} |I, -\tilde{\mathcal{P}}, M=0\rangle \\ |B_+\rangle \\ |B_-\rangle \\ |H, D\rangle \end{matrix} , \tag{4.57}$$

where $\kappa_\pm$ and $E_{B_\pm}$ are the dressed-state eigenvectors and eigenvalues, given by

$$
\begin{aligned}
|B_\pm\rangle &\equiv \pm\kappa_\pm |I, \tilde{\mathcal{P}}, M=0\rangle + \kappa_\mp |B\rangle\,, \\
E_{B\pm} &\equiv \frac{1}{2}\left(\Delta_* \pm \sqrt{\Delta_*^2 + \Omega_r^2}\right), \\
\kappa_\pm &\equiv \frac{1}{\sqrt{2}}\sqrt{1 \pm \frac{\Delta_*}{\sqrt{\Delta_*^2 + \Omega_r^2}}}.
\end{aligned}
$$

This is a nearly diagonal Hamiltonian, and hence we can apply second-order perturbation theory to compute the admixture of the bright state arising from coupling through the opposite (off-resonant) parity state.

$$
|D'\rangle = |D\rangle + \sum_\pm a_{B\pm} |B_\pm\rangle\,, \tag{4.58}
$$

where $|D'\rangle$ is the perturbed dark state and the admixture coefficient is

$$
a_{B_\pm} = \frac{\langle B_\pm | H_{\text{dressed}} | I, -\tilde{\mathcal{P}}, M=0\rangle \langle I, -\tilde{\mathcal{P}}, M=0 | H_{\text{dressed}} | H, D\rangle}{(E_D - E_{B\pm})(E_D - E_{I,-\tilde{\mathcal{P}}})}, \tag{4.59}
$$

where $E_D \equiv 0$ is the eigenvalue of the dark state and $E_{I,-\tilde{\mathcal{P}}} \equiv \Delta_* + \tilde{\mathcal{P}}\Delta_I$ is the eigenvalue of the off-resonant parity state. Plugging in the matrix elements yields

$$
a_{B_\pm} = -\frac{\Omega_r^2 \cos 2\Theta \sin 2\Theta\, \kappa_\mp}{4(\Delta_* + \tilde{\mathcal{P}}\Delta_I) E_{B\pm}}. \tag{4.60}
$$

We are ultimately interested in the admixture of the bright state in $|D'\rangle$, which is given

by the projection

$$\langle B|D'\rangle = \sum_{\pm} a_{B\pm} \langle B|B_{\pm}\rangle \tag{4.61}$$

$$= \sum_{\pm} a_{B\pm}\kappa_{\mp} \tag{4.62}$$

$$= -\frac{\Omega_r^2 \cos 2\Theta \sin 2\Theta}{4(\Delta_* + \tilde{\mathcal{P}}\Delta_I)} \sum_{\pm} \frac{\kappa_{\mp}^2}{E_{B\pm}} \tag{4.63}$$

$$= -\frac{\Omega_r^2 \cos 2\Theta \sin 2\Theta}{4(\Delta_* + \tilde{\mathcal{P}}\Delta_I)} \frac{4\Delta_*}{\Omega_r^2} \tag{4.64}$$

$$\approx -\frac{2d\Theta\, \Delta_*}{\Delta_* + \tilde{\mathcal{P}}\Delta_I} \tag{4.65}$$

$$= -\frac{2d\Theta(\Delta - i\gamma_I/2)}{\Delta - i\gamma_I/2 + \tilde{\mathcal{P}}\Delta_I} \tag{4.66}$$

$$= -\frac{2d\Theta\left(\tilde{\mathcal{P}}\Delta_I\Delta + \Delta^2 + \gamma_I^2/4\right)}{(\Delta + \tilde{\mathcal{P}}\Delta_I)^2 + (\gamma_I/2)^2} + i\frac{d\Theta\tilde{\mathcal{P}}\Delta_I\gamma_I}{(\Delta + \tilde{\mathcal{P}}\Delta_I)^2 + (\gamma_I/2)^2}. \tag{4.67}$$

This gives the admixture of the bright state. The phase induced by this admixture is given by the imaginary part of the bright-state amplitude, as shown in appendix D:

$$d\phi = \mathrm{Im}\big(\langle B|D'\rangle\big) = \frac{d\Theta\tilde{\mathcal{P}}\Delta_I\gamma_I}{(\Delta + \tilde{\mathcal{P}}\Delta_I)^2 + (\gamma_I/2)^2} \tag{4.68}$$

$$\approx \tilde{\mathcal{P}}\frac{d\Theta\Delta_I\gamma_I}{\Delta^2 + \Delta_I^2 + \gamma_I^2/4} - \frac{2d\Theta\Delta\Delta_I^2\gamma_I}{(\Delta^2 + \Delta_I^2 + \gamma_I^2/4)^2}, \tag{4.69}$$

where we have expanded the result into a $\tilde{\mathcal{P}}$-correlated part and a $\tilde{\mathcal{P}}$-independent part.

Plugging in $d\Theta = 0.01$, $\Delta^{\mathcal{NE}} = 5$ kHz, $\Delta_I = 91$ MHz, and $\gamma_I = 1.1 \times 10^7\,\mathrm{s}^{-1}$, we obtain $\phi^{\mathcal{NE}} \approx 20$ nrad, or equivalently $\omega^{\mathcal{NE}} \approx 4\ \mu$rad/s, which is below our final statistical sensitivity by a factor of order unity. Note that this estimate assumes an average ellipticity in the readout X, Y beams. Because the two readout polarizations are orthogonal, average ellipticity picked up from birefringent optics should be suppressed, so the actual systematic is likely even smaller and unlikely to be observable in our experiment.

4.6 Excess noise sources

4.6.1 Zeeman phase noise

In ACME II, we discovered that excess noise begins to appear when a magnetic field larger than about 2.6 mG is applied [90, 91]. Subsequent investigation showed that this excess noise was due to fluctuations in the mean forward velocity of the molecular beam. This caused trace-to-trace variations in the Zeeman phase accumulated by the molecules. The variation is random in nature, and the noise therefore leaks into the eEDM channel.

To mitigate this noise in ACME III, we shimmed the residual magnetic field down to below 10 μG and reduced the switching magnetic field to 100 μG. This reduced the Zeeman phase noise to a negligible level. However, excess noise does begin to dominate again if we increase the magnetic field to values much larger than 100 μG.

4.6.2 Power imbalance of readout beams

As mentioned in 2.5.10, one significant source of excess noise observed early on was due to limited saturation of the readout transition. This can be seen from the following simple argument. The lifetime of the excited I state is about 115 ns, while the X, Y polarizations are only turned on for approximately 1.2 μs. The depleted population can then be estimated as $1 - e^{-1.2\,\mu\mathrm{s}/(2\tau_I)} \approx 99.5\%$ if the transition is fully saturated. If the transition is not fully saturated and the laser intensity drifts, the population then fluctuates at a characteristic scale of $\sim 0.5\%$, corresponding to a phase noise of order 10^{-3}, which is enormous compared to our statistical sensitivity of $\sim 10^{-4}/\sqrt{\mathrm{Hz}}$. In reality, an accurate calculation of the saturation must account for spatial transport in a beam experiment, and this simple estimate just serves as a plausibility argument for the observed noise.

Why is the readout transition not deeply saturated? There are practical limits, such as limited power from the sum-frequency-generation laser system and stimulated Brillouin

scattering in the optical fiber discovered by Collin Diver. However, the dominant reason is physical and relates to systematics. In a diatomic molecule, the Ω-doublet structure provides closely spaced opposite-parity states that can be used to orient the molecules in a modest electric field. The same structure, however, also provides an ideal environment for driving unwanted transitions, with many states lying only ~ 100 MHz away from the main readout transition. There is substantial evidence in our systematic data that some of these unwanted transitions are being driven, although we do not observe a statistically significant shift in the eEDM channel. Out of caution, we therefore do not fully saturate the readout transition and instead modulate the readout laser power as a slow switch. We are continuing to investigate possible systematics from this effect, as such interference mechanisms have been shown to be significant sources of systematic error in other precision measurements [120, 121]. As explained in 2.5.10, our primary method to mitigate this noise is to deliver the readout power through a single fiber instead of two fibers. This makes the power drift in the X and Y polarizations highly correlated and thereby reduces noise in the asymmetry.

4.6.3 Miscellaneous sources

In addition to the two main sources of excess noise discussed above, we have also investigated several other possible sources of excess noise. These effects have the potential to increase the reduced chi-squared by 0.2–0.5, which is sufficient to explain the currently observed excess noise level of $\chi^2_{\text{red}} \approx 1.5$–2. We do not dwell on these sources in detail, as our primary focus is on understanding systematic errors.

We have observed increased excess noise in the parity channel when the asymmetry in any state exceeds approximately 0.3. This is likely due to the nonlinear relationship between asymmetry and phase. This observation motivated the use of a small dither angle of $\pm 1^\circ$ to keep the asymmetry small. We also perform daily to weekly zeroing of the non-reversing phase using either a magnetic field or the readout half-wave plate.

A 1° dither angle prevents us from saturating the Cramér–Rao bound on the phase. Our phase estimator is based on $\phi = \mathcal{A}/(2\mathcal{C})$. With a dither angle of 1°, the uncertainty in the contrast is no longer negligible and limits the best achievable excess noise ratio to about $\chi^2_{\text{red}} \approx 1.3$ relative to shot noise. Similarly, any signal-level mismatch between different parity channels further limits the minimum achievable excess noise. In my analysis, I compare the excess noise to the naive photon shot noise obtained by simply summing photon counts. Other members of the collaboration instead compare to the propagated photon shot noise, which accounts for the dither angle and signal-level mismatch. That approach typically yields a lower excess noise ratio of $\chi^2_{\text{red}} \approx 1.5$, whereas my analysis typically yields $\chi^2_{\text{red}} \approx 2$.

CHAPTER 5

ADIABATIC AND NON-ADIABATIC FOLLOWING OF ELECTRIC AND MAGNETIC FIELDS IN ACME III

5.1 Overview

One major difference between ACME III and previous generations is the use of molecules in the $J = 2$ state before the interaction region. Previously, the molecules were mostly in the $X, J = 0$ ground state. When we first set up ACME III, we planned to transfer the molecules back to the $X, J = 0$ state after the lens using a second STIRAP stage (STIRAP II). However, we discovered that there was a large population of molecules in the $Q, J = 2$ state after the STIRAP II stage, which nominally removes these molecules. These molecules continue to decay to the $H, J = 1$ state even after the $H \to I$ cleanup laser and therefore evade being properly cleaned up. For these reasons, it is especially important to understand the dynamics of these molecules in the context of ACME III.

In section 5.2, we write down the Hamiltonian in the $|J, \Omega, M\rangle$ basis. It turns out that the Hamiltonian can be conveniently separated into a 2×2 block matrix form by treating the $\pm\Omega$ submanifold as an almost isolated subspace. Of course, the $\pm\Omega$ subspaces are not strictly isolated but are instead coupled at a frequency given by the Ω-doublet splitting. In the Q state of ThO, however, the Ω-doublet splitting is sufficiently small that this approximation is often valid.

In section 5.3, we attempt to diagonalize the Hamiltonian and discuss properties of its solutions, including

1. Near conservation of population in each Ω state in the $Q, J = 2$ manifold in ACME III. This statement by itself is not surprising, since the Ω-doublet splitting in the $Q, J = 2$ state is very small. However, recall that $\Omega = \tilde{\mathcal{N}}\tilde{\mathcal{E}}M$. Consequently, conservation of population in Ω, combined with preparation in a well-defined Ω state, directly leads to

the observed $(S/I)^{\mathcal{NE}}$ signal.

2. Low co-magnetometer signal. Contrary to earlier expectations, stray fields actually mix the Q-state molecules rather efficiently (at least at Northwestern). However, this mixing occurs only within a single Ω state, and the molecules remain highly elliptical in each $\mathcal{N}$ state, leading to a reduced co-magnetometer signal.

3. Geometric phase. We show that expressions for the geometric phase can be straightforwardly derived using Wigner $\mathcal{D}$ matrices. We also estimate its approximate magnitude for ACME III, following a procedure similar to that used by Amar for ACME I [122, 123].

4. Potential systematics. We discuss possible systematic effects arising from this Hamiltonian and present two examples in which an Ω-correlated phase can occur for the $Q, J = 2, M = \pm 1$ states. Such phases can appear as EDM-mimicking signals in the H state via Z-polarized decay.

Various symbols are introduced throughout this chapter. For clarity, they are summarized in table 5.1.

5.2 Hamiltonian

Let us begin by choosing an ordering of the basis. This will dramatically simplify the math that follows. Let us say we are interested in the manifold where the total angular momentum is J and the projection of the total electronic angular momentum onto the molecular axis is $\pm|\Omega|$. We order the basis states $|J, \Omega, M\rangle$ as follows:

$$\underbrace{|J,|\Omega|,-J\rangle, |J,|\Omega|,-J+1\rangle, \ldots, |J,|\Omega|,J\rangle}_{|\Omega|\text{, increasing in } M}, \underbrace{|J,-|\Omega|,-J\rangle, |J,-|\Omega|,-J+1\rangle, \ldots, |J,-|\Omega|,J\rangle}_{-|\Omega|\text{, increasing in } M}. \tag{5.1}$$

Table 5.1: Definitions of symbols used in this chapter.

Symbol	Physical meaning
$\vec{\mathbb{E}}$	Electric field vector (in frequency units).
$\vec{\mathbb{B}}$	Magnetic field vector (in frequency units).
Δ_Ω	Splitting of the Ω doublets.
$\hat{n}_\pm$	Quantization axis for the $\pm\lvert\Omega\rvert$ manifolds.
$\phi_\pm$	First Euler (zyz) angle for the $\pm\lvert\Omega\rvert$ manifolds.
$\theta_\pm$	Second Euler (zyz) angle for the $\pm\lvert\Omega\rvert$ manifolds.
ψ_{meas}	Third Euler (zyz) angle; rotation needed to align with the lab measurement axes.
$R_{\text{Euler-zyz}}$	Rotation operator using the Euler zyz convention.
$U_\pm$	Euler rotation from the lab frame to the frame whose z axis is aligned with $\hat{n}_\pm$.
$m_\pm$	Projection of total angular momentum onto $\hat{n}_\pm$.
U_{Δ_Ω}	Differential rotation operator between the $\pm\lvert\Omega\rvert$ quantization axes.
M	Projection of total angular momentum onto the lab Z axis.
λ_i	Eigenvalues of $-iU_\pm\dot{U}_\pm^\dagger$ (rotation rates of $\hat{n}_\pm$).

We also need to choose a quantization axis. We define this to be the lab Z axis. In ACME, this is colinear with the propagation direction of the STIRAP beams and with the electric field direction of the interaction-region field plates. In this chapter, we are interested in the time evolution between STIRAP I ($X \to C \to Q$ lens STIRAP) and the end of the interaction region defined by the field plates. Let $t = 0$ be the time when the molecules leave STIRAP I, and let $t = T$ be the time when the molecules reach the end of the precession region and are read out. We are interested in the time evolution described by the time-dependent Hamiltonian $H(t)$.

Let us first look at a specific example of what the Hamiltonian looks like. This will help us write down the general time-dependent Hamiltonian later. Suppose that at $t = 0$, the electric and magnetic fields are both aligned with the lab Z axis. The Hamiltonian is then

given by

$$H(t=0) =$$

$$\left.\begin{pmatrix}
(\mathbb{E}_Z - \mathbb{B}_Z)(-J) & 0 & 0 & 0 & \frac{\Delta_\Omega}{2} & 0 & 0 & 0 \\
0 & (\mathbb{E}_Z - \mathbb{B}_Z)(-J+1) & 0 & 0 & 0 & \frac{\Delta_\Omega}{2} & 0 & 0 \\
0 & 0 & \dots & 0 & 0 & 0 & \frac{\Delta_\Omega}{2} & 0 \\
0 & 0 & 0 & (\mathbb{E}_Z - \mathbb{B}_Z)(J) & 0 & 0 & 0 & \frac{\Delta_\Omega}{2} \\
\frac{\Delta_\Omega}{2} & 0 & 0 & 0 & (-\mathbb{E}_Z - \mathbb{B}_Z)(-J) & 0 & 0 & 0 \\
0 & \frac{\Delta_\Omega}{2} & 0 & 0 & 0 & (-\mathbb{E}_Z - \mathbb{B}_Z)(-J+1) & 0 & 0 \\
0 & 0 & \frac{\Delta_\Omega}{2} & 0 & 0 & 0 & \dots & 0 \\
0 & 0 & 0 & \frac{\Delta_\Omega}{2} & 0 & 0 & 0 & (-\mathbb{E}_Z - \mathbb{B}_Z)(J)
\end{pmatrix}\right|_{t=0}
\begin{matrix}
|J,|\Omega|,-J\rangle \\
|J,|\Omega|,-J+1\rangle \\
\cdots \\
|J,|\Omega|,J\rangle \\
|J,-|\Omega|,-J\rangle \\
|J,-|\Omega|,-J+1\rangle \\
\cdots \\
|J,-|\Omega|,J\rangle
\end{matrix}$$

Here, $\mathbb{E}_Z(t) \equiv \frac{D_\parallel \mathcal{E}_Z(t)|\Omega|}{J(J+1)}$ is the magnitude of the Stark shift without the dependence on M and the sign of Ω, where $D_\parallel$ is the parallel dipole moment of the molecule in this manifold and $\mathcal{E}_Z$ is the projection of the electric field onto the lab Z axis. Similarly, $\mathbb{B}_Z(t) \equiv \frac{\mu_B g_\parallel \mathcal{B}_Z(t)|\Omega|}{J(J+1)}$ is the magnitude of the Zeeman shift without the dependence on M, where μ_B is the Bohr magneton, $g_\parallel$ is the parallel g factor, and $\mathcal{B}_Z$ is the projection of the magnetic field onto the lab Z axis. Finally, Δ_Ω is the frequency splitting of the Ω doublets in this manifold. In this special case of perfectly aligned electric and magnetic fields along Z, the Hamiltonian has a simple block form

$$H(t=0) = \left.\left(\begin{array}{c|c} (\mathbb{E}_Z - \mathbb{B}_Z)J_Z & \frac{\Delta_\Omega}{2} I \\ \hline \frac{\Delta_\Omega}{2} I & (-\mathbb{E}_Z - \mathbb{B}_Z)J_Z \end{array}\right)\right|_{t=0} \begin{matrix} |J,\Omega\rangle \\ \\ |J,-\Omega\rangle \end{matrix} \tag{5.2}$$

In general, the electric and magnetic fields are not fully aligned with the lab Z axis. We can therefore define $\vec{\mathbb{E}} \equiv (\mathbb{E}_X, \mathbb{E}_Y, \mathbb{E}_Z)$ and $\vec{\mathbb{B}} \equiv (\mathbb{B}_X, \mathbb{B}_Y, \mathbb{B}_Z)$ such that

$$H(t) = \left(\begin{array}{c|c} (\vec{\mathbb{E}}(t) - \vec{\mathbb{B}}(t)) \cdot \vec{J} & \frac{\Delta_\Omega}{2} I \\ \hline \frac{\Delta_\Omega}{2} I & (-\vec{\mathbb{E}}(t) - \vec{\mathbb{B}}(t)) \cdot \vec{J} \end{array}\right) \begin{matrix} |J,\Omega\rangle \\ \\ |J,-\Omega\rangle \end{matrix} . \tag{5.3}$$

Here we have dropped the explicit evaluation at $t = 0$ to be fully general. This is the most important step. The key point is that we have moved to a basis of even and odd superpositions of good-parity states, which rotates the off-diagonal Stark terms onto the diagonal. The Stark interaction now appears as a pure vector shift, which dramatically simplifies the mathematics.

We are now ready to fully state the Hamiltonian problem of interest.

Statement of the Hamiltonian problem:

Using the basis $|J, \Omega, M\rangle =$

$$\underbrace{|J, -|\Omega|, -J\rangle, |J, -|\Omega|, -J+1\rangle, \ldots, |J, -|\Omega|, J\rangle}_{-|\Omega|\text{, increasing in } M}, \underbrace{|J, |\Omega|, -J\rangle, |J, |\Omega|, -J+1\rangle, \ldots, |J, |\Omega|, J\rangle}_{|\Omega|\text{, increasing in } M} \tag{5.4}$$

we have a time-dependent Hamiltonian of the block form

$$H(t) = \left(\begin{array}{c|c} (\vec{\mathbb{E}}(t) - \vec{\mathbb{B}}(t)) \cdot \vec{J} & \frac{\Delta_\Omega}{2} I \\ \hline \frac{\Delta_\Omega}{2} I & (-\vec{\mathbb{E}}(t) - \vec{\mathbb{B}}(t)) \cdot \vec{J} \end{array}\right) \begin{array}{c} |J, |\Omega|\rangle \\ \\ |J, -|\Omega|\rangle \end{array} . \tag{5.5}$$

In later derivations, we will refer to the $(\pm\vec{\mathbb{E}}(t) - \vec{\mathbb{B}}(t)) \cdot \vec{J}$ blocks as the diagonal blocks and the $\frac{\Delta_\Omega}{2} I$ blocks as the off-diagonal blocks.

The system starts at $t = 0$ and ends at some $t = T$. What can we say about the dynamics described by this Hamiltonian?

Of course, the exact solution of the wavefunction cannot be obtained without exact knowledge of the fields as functions of time. This is impossible in ACME for many reasons. However, we will show that we can still say quite a few general things about this system

using properties of angular momentum.

5.3 Derivations

5.3.1 Setup and definitions

Let us begin by transforming the Hamiltonian to a near-diagonal form. Let us define the following two quantization axes[1]:

$$\hat{n}_{\pm}(t) \equiv \frac{\vec{\mathbb{E}}(t) \mp \vec{\mathbb{B}}(t)}{\left|\pm\vec{\mathbb{E}}(t) - \vec{\mathbb{B}}(t)\right|} \equiv (\sin\theta_{\pm}(t)\cos\phi_{\pm}(t), \sin\theta_{\pm}(t)\sin\phi_{\pm}(t), \cos\theta_{\pm}(t)) \,, \tag{5.6}$$

where $\theta_{\pm}(t), \phi_{\pm}(t)$ are the polar and azimuthal angles of the quantization axes $\hat{n}_{\pm}(t)$. Observe that the effect of the electric field is equal and opposite in the $\pm|\Omega|$ manifolds relative to the magnetic field. This means the $\pm|\Omega|$ manifolds effectively see two different quantization axes, as defined by $\hat{n}_{\pm}(t)$.

With $\hat{n}_{\pm}(t)$ defined, note that the diagonal block matrix elements are simply of the form $\hat{n}_{\pm}(t) \cdot \vec{J}$. This means that if we rotate to a local coordinate system such that the local z axis is aligned with $\hat{n}_{\pm}(t)$, the diagonal block matrices are diagonalized into the form J_z. This motivates the following definitions.

$$U_{\pm}(t) \equiv R_{\text{Euler-zyz}}(0, \theta_{\pm}(t), \phi_{\pm}(t)) = \left(\mathcal{D}^{J}_{m,m'}(0, \theta_{\pm}(t), \phi_{\pm}(t))\right) , \tag{5.7}$$

where the unitary operator is an Euler–zyz rotation operator and its matrix elements are given by the finite Wigner matrices $\mathcal{D}^{J}_{m,m'}$. At this point there is no preferred measurement x, y axes defined by the final z rotation, so we are free to set that rotation angle to zero. Each $U_{\pm}$ acts only on the $\pm|\Omega|$ subspace, so we define the full unitary operator on the entire

1. There are two possible sign choices here. We choose the convention where, if the electric field dominates, the two doublets have their quantization axes pointing in the same direction instead of opposite directions.

Hilbert space as

$$U(t) = \left(\begin{array}{c|c} U_+(t) & 0 \\ \hline 0 & U_-(t) \end{array}\right). \tag{5.8}$$

The Hamiltonian is then transformed according to

$$\tilde{H}(t) = U(t)H(t)U^\dagger(t) - iU(t)\dot{U}^\dagger(t) \tag{5.9}$$

$$= \left(\begin{array}{c|c} \left|\vec{\mathbb{E}}(t) - \vec{\mathbb{B}}(t)\right| J_z & \frac{\Delta_\Omega}{2} U_+(t)U_-^\dagger(t) \\ \hline \frac{\Delta_\Omega}{2} U_-(t)U_+^\dagger(t) & -\left|-\vec{\mathbb{E}}(t) - \vec{\mathbb{B}}(t)\right| J_z \end{array}\right) - i \left(\begin{array}{c|c} U_+(t)\dot{U}_+^\dagger(t) & 0 \\ \hline 0 & U_-(t)\dot{U}_-^\dagger(t) \end{array}\right) \begin{array}{l} |J,\Omega\rangle_{\hat{n}_+(t)} \\ |J,-\Omega\rangle_{\hat{n}_-(t)} \end{array} \tag{5.10}$$

$$= \left(\begin{array}{c|c} \left|\vec{\mathbb{E}}(t) - \vec{\mathbb{B}}(t)\right| J_z & \frac{\Delta_\Omega}{2} U_{\Delta_\Omega}(t) \\ \hline \frac{\Delta_\Omega}{2} U_{\Delta_\Omega}^\dagger(t) & -\left|-\vec{\mathbb{E}}(t) - \vec{\mathbb{B}}(t)\right| J_z \end{array}\right) - i \left(\begin{array}{c|c} U_+(t)\dot{U}_+^\dagger(t) & 0 \\ \hline 0 & U_-(t)\dot{U}_-^\dagger(t) \end{array}\right) \begin{array}{l} |J,\Omega\rangle_{\hat{n}_+(t)} \\ |J,-\Omega\rangle_{\hat{n}_-(t)} \end{array}, \tag{5.11}$$

where we have added the label $\hat{n}_\pm(t)$ in $|J,\pm|\Omega|\rangle_{\hat{n}_\pm(t)}$ to keep track of the fact that the quantization axes for the $\pm|\Omega|$ states are no longer the lab Z axis but instead $\hat{n}_\pm(t)$. The off-diagonal blocks

$$U_{\Delta_\Omega}(t) \equiv U_+(t)U_-^\dagger(t) \approx \left(\mathcal{D}^J_{m,m'}(0, \theta_+(t) - \theta_-(t), \phi_+(t) - \phi_-(t))\right)$$

represent the differential rotation between $U_\pm(t)$, and the last approximation assumes small angles. Usually, this is a very small rotation and $U_{\Delta_\Omega}(t) \approx I$. Only when the two quantization axes $\hat{n}_\pm(t)$ differ significantly in orientation does this differential rotation become nontrivial, which occurs if and only if the Stark and Zeeman shifts are comparable in magni-

tude and orthogonal in direction. This is not the case in the ACME spin-precession region.

With the above setup and definitions, we are ready to derive some properties of this time-dependent Hamiltonian problem. For simplicity, we will drop the explicit time dependence (t) when possible. We find it helpful to conceptually divide the Hamiltonian matrix elements as follows.

1. $\pm\left|\pm\vec{\mathbb{E}}-\vec{\mathbb{B}}\right| J_z$ is fully diagonal and proportional to $m_\pm$, the projection of total angular momentum onto $\hat{n}_\pm$, the local quantization axes[2]. It is the dominant term in most regions of ACME such as the STIRAP I, electrostatic lens, and interaction region. It is fully Ω- and $m_\pm$-conserving. When it dominates, Ω and $m_\pm$ are good quantum numbers. However, some Ω sublevels may be degenerate. In this case, other effects such as the Ω-doublet term come into play.

2. $\frac{\Delta_\Omega}{2}U_{\Delta_\Omega}$ is the only term in our Hamiltonian that couples the $\pm|\Omega|$ states. Its relative size compared to the energy splittings between eigenstates of $\pm\left|\pm\vec{\mathbb{E}}-\vec{\mathbb{B}}\right| J_z$ determines whether Ω is conserved. In most regions of ACME, the rotation $U_{\Delta_\Omega}=I$ is trivial, and thus it only couples states with $m_+=m_-$. However, in transition regions where $|\vec{\mathbb{B}}|\approx|\vec{\mathbb{E}}|$, this rotation may become nontrivial and can couple different $m_\pm$ states, e.g. $m_-=-m_+$.

3. The off-diagonal elements of $U_\pm\dot{U}_\pm^\dagger$ describe how fast the quantization axis rotates. This is the non-adiabatic term. When these elements are not small compared to the energy differences of the matrix elements of $\pm\left|\pm\vec{\mathbb{E}}-\vec{\mathbb{B}}\right| J_z$, $m_\pm$ is no longer a good quantum number and transitions between different $m_\pm$ sublevels occur. This term, however, never causes mixing between different Ω states.

4. The diagonal elements of $U_\pm\dot{U}_\pm^\dagger$ describe the phenomenon of geometric phase. When

2. Whenever we use the quantum number M, we are referring to the lab Z axis instead of the local z axis.

the quantization axis moves in space, the different $m_\pm$ sublevels acquire a Berry phase due to the geometric nature of parallel transport.

Next, we consider a few ACME-relevant regimes where

1. $\hat{\mathbb{E}} = \hat{\mathbb{B}} = \hat{Z}$, and $|\vec{\mathbb{E}}| \gg |\vec{\mathbb{B}}|, \Delta_\Omega, |\lambda_i|$, where λ_i is any eigenvalue of $-iU_+\dot{U}_+^\dagger$. Physically, this is equivalent to saying that the magnetic field, the Ω-doublet splitting, and the rotation rate of the local quantization axis are small. This is applicable to all STIRAP regions and the interaction region. It is not strictly true in the lens region because the electric field direction there is not uniform, but the qualitative properties are similar.

2. $\hat{\mathbb{B}} \approx \hat{Z}$, and $|\vec{\mathbb{B}}| \gg |\vec{\mathbb{E}}| \approx \Delta_\Omega \approx |\lambda_i|$. This is applicable to the bellow plus nipple region after the second STIRAP ($Q \to C \to X$) but before entering the magnetic shields.

3. $|\vec{\mathbb{E}}| \approx |\lambda_i| \approx \Delta_\Omega \gg |\vec{\mathbb{B}}|$. This is applicable to most non-interaction regions inside the magnetic shields.

4. $|\lambda_i| \gg |\vec{\mathbb{E}}|, |\vec{\mathbb{B}}|, \Delta_\Omega$ for some eigenvalue λ_i of $U_+\dot{U}_+^\dagger$. This is only applicable to a few small regions in the transition zones between the electrostatic lens and STIRAP II and is possibly applicable to some small regions near the fixed collimator right before the field plates. The fast rotation of the local quantization axis arises because the fields in these regions are effectively zero, so any tiny change can appear as a very rapid rotation of the quantization axis.

Regime 1: $\left|\vec{\mathbb{E}}\right| J_z$ dominates

This is the most common regime in ACME. In this case, $\hat{\mathbb{E}} = \hat{\mathbb{B}} = \hat{Z}$, and $|\vec{\mathbb{E}}| \gg |\vec{\mathbb{B}}|, \Delta_\Omega, |\lambda_i|$, where λ_i is any eigenvalue of $-iU_+\dot{U}_+^\dagger$. From the definition, $\hat{n}_+ = \hat{n}_- = \hat{Z}$, so the differential rotation operator $U_{\Delta_\Omega} = I$. Thus, the Ω-doublet matrix elements behave exactly as expected and couple states with the same M. By our assumption, these same M

states are well separated in energy because

$$\left|\left|\vec{\mathbb{E}} - \vec{\mathbb{B}}\right| M - \left(-\left|-\vec{\mathbb{E}} - \vec{\mathbb{B}}\right| M\right)\right| \approx 2\left|\vec{\mathbb{E}}\right| M \gg \Delta_\Omega \tag{5.12}$$

for all $M \neq 0$. Therefore, the off-diagonal term $\Delta_\Omega/2$ is strongly suppressed according to perturbation theory, because it is much smaller than the difference between the two diagonal elements it connects. The only exception is $M = 0$. As a result, the eigenstates in this regime are $|J, \Omega, M\rangle$ for $M \neq 0$ and $|J, \Omega, M\rangle \pm |J, -\Omega, M\rangle$ for $M = 0$. Thus, if molecules start in one of the $M \neq 0$ states, they remain there, as shown in figure 5.1.

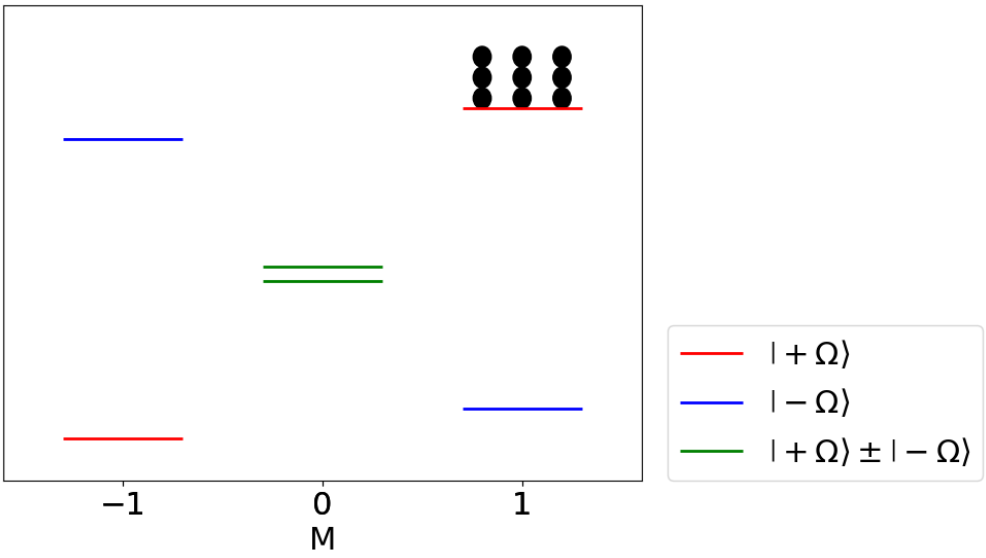

Figure 5.1: Energy levels in regime 1. Population in $|J, \pm|\Omega|, M \neq 0\rangle$ is conserved.

Regime 2: $\left|\vec{\mathbb{B}}\right| J_z$ dominates

In this regime, $\hat{\mathbb{B}} \approx \hat{Z}$ and $|\vec{\mathbb{B}}| \gg |\vec{\mathbb{E}}|, \Delta_\Omega, |\lambda_i|$, where λ_i is any eigenvalue of $-iU_+\dot{U}_+^\dagger$. This regime is applicable to the bellows-plus-nipple region after the second STIRAP ($Q \to C \to X$) but before entering the magnetic shields. To see why, we roughly estimate the energy scales of each interaction:

(a) **Magnetic field:** The Earth's magnetic field is around half a gauss. Using the g factor for the Q state with $M = 1$, $\Omega = 2$, and $J = 2$, this corresponds to an energy of order 300 kHz.

(b) **Electric field:** When no electric field is intentionally applied, stray fields are dominated by surface charges. Since all surfaces visible to the molecules in this region are metallic and relatively far away (> 5 cm), we estimate electric fields of a few mV/cm, corresponding to an energy scale of order 10 kHz.

(c) Ω **doublet:** For the Q state with $J = 2$, this splitting was measured to be less than 10 kHz [97].

(d) **Rate of quantization-axis rotation, λ_i:** We estimate this for regions where the electric and magnetic fields are not near zero. The characteristic length scale is $L \approx 2$ cm, set by the separation between the molecular beam and the nearest charge-carrying surface. With a beam velocity of 200 m/s, the maximal characteristic rotation rate is

$$\lambda_i \lesssim \frac{200\,\mathrm{m/s}}{2\,\mathrm{cm}} \approx 10\,\mathrm{kHz}.$$

Although the solution in this regime is simple, it is still useful to work through it using our formulation. The only subtlety is that the quantization axes $\hat{n}_\pm$ are no longer aligned. In this case, we can no longer use the same label M for both $\pm|\Omega|$ manifolds. Instead, we use $m_\pm$ to denote the projection along the axis $\hat{n}_\pm$ within the $\pm|\Omega|$ submanifold. Plugging in the definitions of the quantization-axis angles from equation 5.6, we find that the two

quantization axes point in opposite directions:

$$\hat{n}_+ = -\hat{n}_- \approx -\hat{Z}, \tag{5.13}$$

$$\theta_+ = \theta_- + \pi, \tag{5.14}$$

$$\phi_+ = \phi_-. \tag{5.15}$$

Hence, $U_{\Delta_\Omega} = R_{\text{Euler-zyz}}(0, \pi, 0)$, and its matrix elements are given by $\mathcal{D}^J_{m,m'}(0, \pi, 0)$. This connects states with $m_+ = -m_-$. This may seem counterintuitive because we are working in the local coordinate system. To clarify this, let us map back to the lab frame. Since $\hat{\mathbb{B}} \approx \hat{Z}$ by assumption, we have $\hat{n}_\pm \approx \mp\hat{Z}$. Thus, $M = \mp m_\pm$. In other words, in the lab frame, the Δ_Ω term simply connects states with the same M but opposite Ω, as shown in figure 5.2.

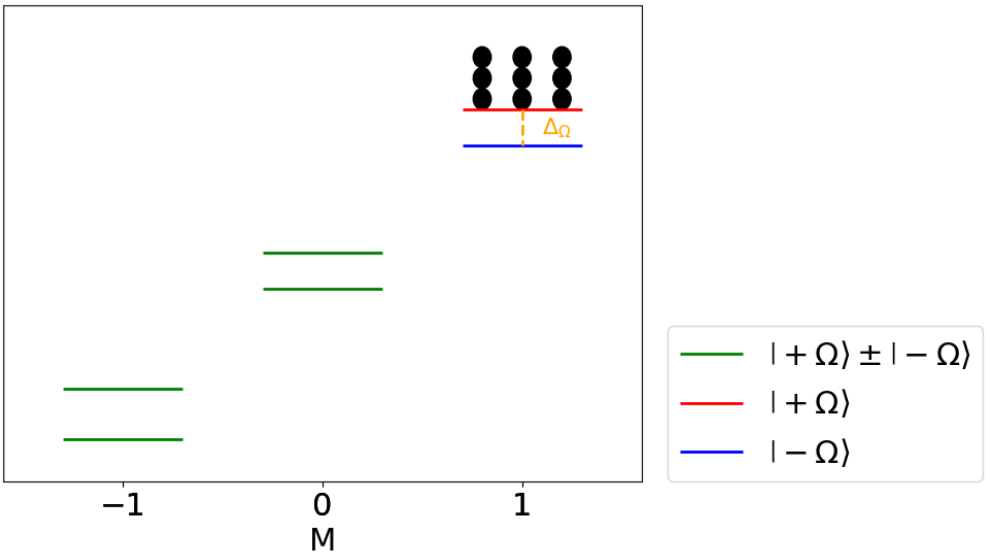

Figure 5.2: Energy levels in regime 2. Population in $|J, M\rangle$ is conserved, but Ω is not a good quantum number. If the molecules are prepared in a single $+\Omega$ state, they will begin to mix into $-\Omega$ at a rate set by the Ω-doublet splitting Δ_Ω, suppressed depending on the size of the Stark shift. In this figure, the Stark shift is assumed to be comparable to the doublet splitting.

Regime 3: no term dominates

In this regime, $|\vec{\mathbb{E}}| \approx |\lambda_i| \approx \Delta_\Omega \gg |\vec{\mathbb{B}}|$. This is applicable to most non-interaction regions inside the magnetic shields. Neither the electric nor the magnetic field is strong enough to clearly define a quantization axis. In this case, it is best not to use our formulation above, since near-zero fields may lead to ill-defined $\hat{n}_\pm$ and formally large λ_i near singularities. Instead, it is best to work in the lab Z basis.

In this regime, different M states within the same Ω sub-manifold are mixed by stray electric fields in the X and Y directions. The corresponding energy scale is likely at the level of a few kHz or less. At the same time, the $+\Omega$ states mix into the $-\Omega$ states within the same M sub-manifold at a rate set by Δ_Ω, as shown in figure 5.3.

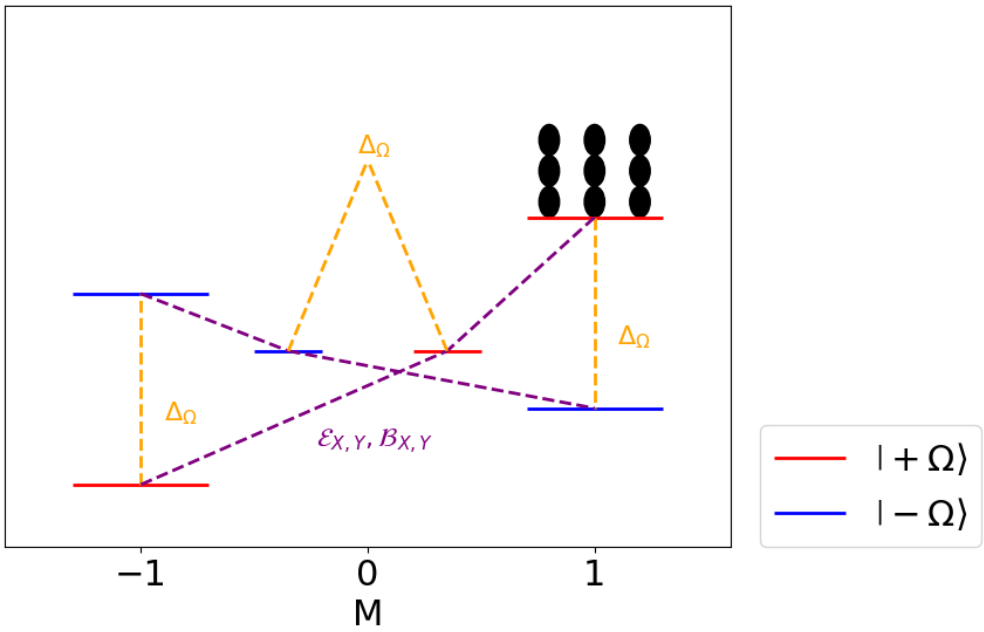

Figure 5.3: Energy levels in regime 3. Two types of mixing (purple and orange) are relevant. Transverse electric and magnetic fields mix states with different M but the same Ω. The Δ_Ω term mixes states with the same M but opposite Ω.

Regime 4: λ_i dominates

In this regime, $|\lambda_i| \gg |\vec{\mathbb{E}}|, |\vec{\mathbb{B}}|, \Delta_\Omega$ for some eigenvalue λ_i of $U_+ \dot{U}_+^\dagger$. This is only applicable to a few tiny regions in the transition between the electrostatic lens and STIRAP, and possibly to some small regions near the fixed collimator right before the field plates. In these regions, the electric field passes through zero, and the rate of rotation of the quantization axis, λ_i, becomes extremely large.

In this case, the time evolution is simply described by a net rotation across the region. The molecular states before entering and after leaving the region are related by a Wigner $\mathcal{D}$ matrix, as in equation 5.8. The rotation occurs strictly within a given $\pm|\Omega|$ sub-manifold, and Ω is conserved. Technically, in this regime there is no change in the molecular quantum state in the lab frame. The apparent state changes arise because the local coordinate system changes rapidly while the molecules are unable to adiabatically follow. In this sense, there is no intrinsic mixing between different M levels in this regime.

However, consider the example in which the molecules leave the electrostatic lens in the low-field-seeking $|m_- = 2, \Omega = -2\rangle$ state and enter a region where one of $\left|\pm\vec{\mathbb{E}} - \vec{\mathbb{B}}\right|$ passes through zero. This region is very small in volume, so only a small fraction of molecules traverse it. Most molecules therefore adiabatically follow the local quantization axis and end up in $M = 2$ again. In contrast, molecules that pass through the zero-field region generally do not return to the $M = 2$ state due to the large λ_i term when viewed in the local coordinate system.

As a result, when considering the net effect of this zero-field region, the λ_i term effectively appears as mixing between different M states, while still conserving Ω. This is illustrated in figure 5.4.

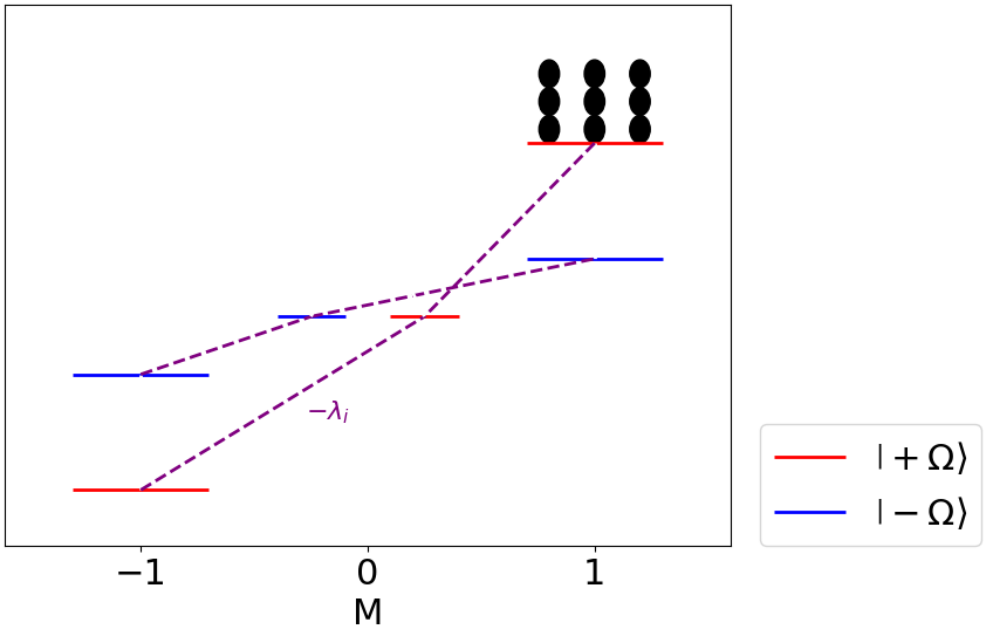

Figure 5.4: Energy levels in regime 4. The net effect of the fast rotation of the local quantization axis is that not all molecules adiabatically return to the same M state they started in. In other words, this effect mixes different M states while conserving Ω.

5.3.2 Conservation of M, Ω and detected $(S/I)^{\mathcal{NE}}$

Using the $|J, \Omega, M\rangle$ basis, we summarize the interactions that mix different Ω, M states in figure 4.10. Using the black circles, we also show the approximate Q-state population distribution detected in the interaction region after initially preparing the molecules in the $M = 2$ state. From the figure, it is clear that there is substantial mixing among different M states within the same $\pm|\Omega|$ sub-manifold, while the mixing between the $\pm|\Omega|$ manifolds is incomplete. This is not surprising, since the mixing among different M states is typically $\sim$ 10 kHz, whereas the mixing between the $\pm|\Omega|$ states is likely only $\lesssim$ 1 kHz. Although the matrix element Δ_Ω is likely larger than 1 kHz by an order-unity factor, it is always suppressed by stray electric fields by an order-unity factor unless $M = 0$. As a result, its effective contribution is probably $\lesssim$ 1 kHz. Given that the no-quantization-axis flight time is around 1 ms, we therefore expect imperfect mixing between the $\pm|\Omega|$ states.

The near conservation of Ω observed in the Q state also explains the observed $(S/I)^{\mathcal{NE}}$

in the H state. Note that E1 decay from $Q \to H$ strictly conserves Ω. Moreover, since

$$\Omega = \tilde{\mathcal{N}}\tilde{\mathcal{E}}M, \tag{5.16}$$

a Ω-correlated population distribution is exactly equivalent to a $\mathcal{NE}$-correlated population difference between the $M = \pm 1$ sub-levels, which is $(S/I)^{\mathcal{NE}}$. In this context, it is also straightforward to show that $(S/I)^{\mathcal{E}}$ and $(S/I)^{\mathcal{N}}$ simply reflects incomplete mixing among different M states within the same Ω sub-manifold. We do observe this effect, but it is smaller than $(S/I)^{\mathcal{NE}}$ by a factor of 2–20.

5.3.3 Low Q-state magnetometer signal level

Because the population is distributed as shown in figure 4.10, when we prepare the molecules in the superposition $|\Omega, M\rangle + |-\Omega, -M\rangle$ by optical pumping, we pump out most of the molecules. This effect was especially apparent at Harvard, where the detected population is around 0.1% of the total population. There, we likely have insufficient mixing among both M and Ω. The magnetometer signal is significantly improved at Northwestern, where the detected population is around 5%, possibly due to the longer flight path that allows additional mixing between the $\pm|\Omega|$ states and a different stray-field distribution that induces substantially better mixing among different M states.

5.3.4 Geometric phase and $\phi^{\mathcal{E}}$

In the limit of large electric and magnetic fields such that no states other than $m_\pm = 0$ are degenerate, almost all off-diagonal terms in equation 5.11 are strongly suppressed. The only first-order corrections remaining for the $m_\pm \neq 0$ levels are the diagonal terms of $U_\pm \dot{U}_\pm^\dagger$. As mentioned earlier, these terms correspond to the geometric phase. We now work out the

relevant expressions. As a reminder, the unitary transformation is defined as

$$U_{\pm}(t) \equiv R_{\text{Euler-zyz}}(0, \theta_{\pm}(t), \phi_{\pm}(t)) = \left(\mathcal{D}^{J}_{m,m'}(0, \theta_{\pm}(t), \phi_{\pm}(t))\right), \tag{5.7}$$

where the polar and azimuthal angles are defined by the two quantization axes of the $\pm|\Omega|$ sub-manifolds,

$$\hat{n}_{\pm}(t) \equiv \frac{\vec{\mathbb{E}}(t) \mp \vec{\mathbb{B}}(t)}{\left|\pm\vec{\mathbb{E}}(t) - \vec{\mathbb{B}}(t)\right|} \equiv (\sin\theta_{\pm}(t)\cos\phi_{\pm}(t), \sin\theta_{\pm}(t)\sin\phi_{\pm}(t), \cos\theta_{\pm}(t)). \tag{5.6}$$

In earlier sections, we set ψ, the final Z rotation of the Euler angles, to zero. This corresponds to choosing a particular XY coordinate system for measurement. Since none of the earlier discussion involved measuring phases between the $\pm M$ states, this choice was allowed for conceptual simplicity. We now restore this degree of freedom and modify equation 5.7 to

$$U_{\pm}(t) \equiv R_{\text{Euler-zyz}}(\psi_{\text{meas}}(t), \theta_{\pm}(t), \phi_{\pm}(t)) = \left(\mathcal{D}^{J}_{m,m'}(\psi_{\text{meas}}(t), \theta_{\pm}(t), \phi_{\pm}(t))\right). \tag{5.17}$$

Note that the physical meaning of $\psi_{\text{meas}}(t)$ is the angle required to rotate the local xy coordinates defined by equation 5.7 to the actual experimental XY coordinates, which are arbitrary by choice. Hence, there is no $\pm$ label on ψ_{meas} because we measure the phases of the $\pm|\Omega|$ states in the same XY coordinate system, as is the case in ACME.

As an example, in the ACME interaction region we have defined coordinate systems at the cleanup and at the readout. Let the angular offsets between these coordinate systems and the lab XY axes be $\psi(t_{\text{cleanup}})$ and $\psi(t_{\text{readout}})$. If these two angles differ, we would measure a corresponding phase

$$\psi(t_{\text{cleanup}}) - \psi(t_{\text{readout}}).$$

This measured phase is solely due to misalignment and not to any geometric-phase physics. This is one benefit of using the Euler zyz convention: it cleanly separates this arbitrary contribution from the true geometric phase.

Let us now examine the geometric-phase contributions from the diagonal elements. We begin with the following properties of the Wigner D matrix. Let

$$U_{mm'} \equiv \mathcal{D}^J_{mm'}(\psi, \theta, \phi) \equiv e^{-im\psi} e^{-im'\phi} d^J_{mm'}(\theta).$$

The diagonal elements satisfy

$$-i \left(U \dot{U}^\dagger \right)_{mm} = m\dot{\psi} + 0 \cdot \dot{\theta} + m \cos\theta \, \dot{\phi}. \tag{5.18}$$

The first contribution, $m\dot{\psi}$, from the final Euler z rotation, simply describes the phase arising from the angular offset between the lab XY coordinates and the local xy coordinates defined by the Euler rotations. It is completely history independent and depends only on the instantaneous angular offset, so it is not related to geometric phase.

The second contribution from $\dot{\theta}$ vanishes because geometric phase is associated with rotations about the initial z axis. From the definition of θ, we see that it carries no handedness about the z axis and therefore does not contribute. The final term, $m \cos\theta \, \dot{\phi}$, is the genuine geometric-phase contribution. The total geometric phase accumulated over time for a state with quantum number m is

$$\phi_{\text{geometric}} = m \int_{t_0}^{t_f} \cos\theta(t) \, \dot{\phi}(t) \, dt = m \int_{t_0}^{t_f} \big(\cos\theta(t) - 1\big) \dot{\phi}(t) \, dt, \tag{5.19}$$

where we have used the standard assumption that $\phi(t_f) - \phi(t_0) = 0 \mod 2\pi$ to obtain an expression proportional to a solid angle. In the context of ACME, we are free to add constant terms to the geometric phase because we are interested only in its $\mathcal{NE}$-correlated component

rather than its non-reversing component.

Due to the presence of Ω doublets, there are two sets of Euler angles, $\theta_\pm$ and $\phi_\pm$, corresponding to the two quantization axes $\hat{n}_\pm$. Therefore,

$$\phi_{\text{geometric},|\pm|\Omega|,m_\pm\rangle} = m_\pm \int_{t_0}^{t_f} [1-\cos\theta_\pm(t)]\,\dot{\phi}_\pm(t)\,dt. \tag{5.20}$$

Estimating geometric phase in ACME

Amar has estimated this previously, so we use similar numbers [123]. Suppose there is both a non-reversing and an $\mathcal{E}$-correlated transverse electric field of 10 kHz that rotates by a full 2π over the interaction region. Similarly, suppose there is both a non-reversing and an $\mathcal{B}$-correlated transverse magnetic field of 1 Hz pointing along the same direction as the transverse electric field. Then

$$\theta_\pm \approx \frac{\mathbb{E}_\perp \mp \mathbb{B}_\perp}{\mathbb{E}_Z} \approx \frac{(10+10\tilde{\mathcal{E}} \mp 0.001\tilde{\mathcal{B}})\,\text{kHz}}{100\ \text{MHz}} = \frac{(1+1\tilde{\mathcal{E}} \mp 10^{-4}\tilde{\mathcal{B}})}{10^4}, \tag{5.21}$$

where we have dropped the $\mathcal{B}^{\text{nr}}$ contribution since it is only a tiny correction. The total phase accumulated is approximately

$$\phi_{\text{geometric},|\pm|\Omega|,m_\pm\rangle} = (1-\cos\theta_\pm)\,2\pi \approx \theta_\pm^2\,\pi \approx \frac{2\pi}{10^8}\left(1+\tilde{\mathcal{E}} \mp 10^{-4}\tilde{\mathcal{B}} \mp 10^{-4}\tilde{\mathcal{E}}\tilde{\mathcal{B}}\right). \tag{5.22}$$

We now consider the relative phase between m_+ and $m_- = -m_+$ within the same $\mathcal{N}$ state. For notational convenience, let $|\mathcal{N}| = |\Omega| = |m_\pm| = 1$, as in our standard EDM measurement.

$$\frac{\phi_{\text{geometric},|\mathcal{N},m_+\rangle} - \phi_{\text{geometric},|\mathcal{N},m_-=-m_+\rangle}}{2} = \frac{\phi_{\text{geometric},|\Omega=\tilde{\mathcal{N}}\tilde{\mathcal{E}}m_+,m_+\rangle} - \phi_{\text{geometric},|-\Omega,m_-=-m_+\rangle}}{2} \tag{5.23}$$

$$\approx \frac{2\pi}{10^8}\left(1+\tilde{\mathcal{E}}\right), \tag{5.24}$$

where the magnetic-field contributions cancel because they are the same for the $\pm\Omega$ states, and because we have neglected the second-order term above. If retained, it would contribute at the level of $\sim 10^{-16}$ rad.

This estimate shows that the leading contribution is only ~ 10 μrad/s, consistent with Amar's estimate and barely detectable by ACME III. Importantly, as Amar first pointed out, this phase is only $\mathcal{E}$ correlated, and it is difficult to generate a $\mathcal{N}$ dependence. In principle, motional fields and a $\mathcal{N}$-correlated g factor could introduce a $\mathcal{NE}$ dependence, but it would be completely undetectable in ACME.

5.3.5 Systematics from non-adiabatic following or electric fields

One way to eliminate the Ω population imbalance, or $(S/I)^{\mathcal{NE}}$, is to prepare the initial Q-state population evenly in both Ω states. Although this sounds appealing, we must ensure that this does not merely hide the problem and lead to a worst-case scenario: rotating the quantum state to minimize the $\mathcal{NE}$ population imbalance while maximizing the $\mathcal{NE}$ phase.

Another scenario to consider is the region after $Q \to I \rightsquigarrow X$ pumping in the second STIRAP. The remaining molecules in the dark state have $M = 0$ and negative parity. This configuration would not produce any $\mathcal{NE}$ population imbalance, but it could still generate a $\mathcal{NE}$ phase. Below, we discuss possible mechanisms by which this can occur through adiabatic following. We begin with a brief summary:

1. If, in the interaction region, there is an Ω-correlated phase for the $Q, M = \pm 1$ states, this systematic phase will be inherited by the $H, M = \pm 1$ states when the $Q \to H$ decay occurs via Z-polarized light, producing a systematic phase.

2. In the first STIRAP ($X \to C \to Q$), if the molecules begin with an Ω-correlated phase and this phase survives up to the magnetic shields, it could appear as a systematic phase. We will show that this mechanism is suppressed and is likely not a concern.

3. In the "Qleanup" region after the lens, if the dark-state $M = 0$ molecules do not mix uniformly with the opposite-parity $M = 0$ state, they could interact with stray transverse electric fields inside the shields and generate systematics.

First STIRAP equal population in $M = \pm 2$

This is a contrived example, included only as an existence proof. Suppose the population after the first STIRAP begins in

$$\frac{1}{\sqrt{2}}\left(e^{-i\phi}\left|\Omega, M=2\right\rangle + e^{i\phi}\left|-\Omega, M=-2\right\rangle\right),$$

so that there is an initial relative phase between the $\pm\Omega$ states. Such a relative phase could arise, for example, if the linear polarization used for first STIRAP and that used for readout in the interaction region are not perfectly aligned, which is certainly true to some extent.

Next, suppose that through a combination of adiabatic and non-adiabatic following, the molecular state transforms according to $D^{J=2}_{m,m'}(0, \frac{\pi}{2}, 0)$ upon arriving at the interaction region. The final state is then

$$\begin{aligned}\left|\text{final}\right\rangle =& \frac{1}{2}\left(-e^{-i\phi}\left|\Omega, M=1\right\rangle - e^{-i\phi}\left|\Omega, M=-1\right\rangle + e^{i\phi}\left|-\Omega, M=1\right\rangle + e^{i\phi}\left|-\Omega, M=-1\right\rangle\right) \\ &+ (\text{other } M \text{ states}).\end{aligned}$$

We see that the relative phase ϕ between the $\pm\Omega$ states is retained. This provides an explicit example of an EDM-mimicking phase. However, this mechanism is strongly suppressed in practice because the molecules rapidly dephase in the electrostatic lens. As a result, the evolution cannot be described by a simple unitary rotation as assumed here. This rapid dephasing should erase any pre-existing phase coherence between the $\pm\Omega$ states. Because the lens acts as an effective dephasing stage, preparing the molecules in the state $\left|\Omega, M=2\right\rangle + \left|-\Omega, M=-2\right\rangle$ is actually a reasonable strategy and should not be particularly susceptible

to systematic errors.

Second STIRAP residual $M = 0$ dark state

The current plan to replace the second STIRAP is to instead perform optical pumping $Q \to I \rightsquigarrow X, J = 0$ via the near-cycling state I. Even with polarization switching, some population remains in

$$|D\rangle = \frac{1}{\sqrt{2}} \left(|\Omega = 2, M = 0\rangle + |\Omega = -2, M = 0\rangle \right). \tag{5.25}$$

First, we assume that this pure state is not completely remixed before entering the magnetic shields. This is possible because the lab magnetic field provides a quantization axis and suppresses mixing between different M states. As a result, some molecules remain in the $|D\rangle$ state upon entering the shields.

Next, suppose that there is a phase-dependent asymmetry from both the H state and the Q state (which eventually decays to H), which we denote as

$$A_H = \sin(2\phi_H) = \frac{D_H}{S_H} \equiv \frac{X_H - Y_H}{X_H + Y_H}, \tag{5.26}$$

$$A_Q = \sin(2\phi_Q) = \frac{D_Q}{S_Q} \equiv \frac{X_Q - Y_Q}{X_Q + Y_Q}, \tag{5.27}$$

where A_H, A_Q are the asymmetries measured from the H state and the Q state (via decay), ϕ_H, ϕ_Q are the corresponding molecular phases, X_H, Y_H, X_Q, Y_Q are the signals in the X, Y polarizations for the H and Q states, and D, S denote the difference and sum in X, Y, respectively.

For the total asymmetry after combining both signals, we obtain

$$A_{\text{total}} = \frac{D_H + D_Q}{S_H + S_Q} \tag{5.28}$$

$$= \frac{D_H + S_Q A_Q}{S_H + S_Q} \tag{5.29}$$

$$\approx \frac{D_H + S_Q A_Q}{S_H}, \tag{5.30}$$

where we have used the fact that the H-state signal S_H is much larger than the Q-state signal S_Q, which arises from decay into H.

Substituting the phase dependence of the asymmetries, we find

$$A_{\text{total}} \approx \frac{D_H + S_Q A_Q}{S_H} \tag{5.31}$$

$$= \frac{S_H \sin(2\phi_H) + S_Q \sin(2\phi_Q)}{S_H}, \tag{5.32}$$

$$\phi_{\text{total}} \approx \frac{A_{\text{total}}}{2} \approx \phi_H + \frac{S_Q}{2S_H}\sin(2\phi_Q) = \begin{cases} \phi_H + \dfrac{S_Q}{S_H}\phi_Q, & \text{if } \phi_Q \approx 0, \\ \phi_H + \dfrac{S_Q}{2S_H}, & \text{if } \phi_Q - \phi_H \approx \dfrac{\pi}{4} \pmod{\pi}, \\ \phi_H - \dfrac{S_Q}{2S_H}, & \text{if } \phi_Q - \phi_H \approx \dfrac{3\pi}{4} \pmod{\pi}, \end{cases} \tag{5.33}$$

where we have used the fact that the laser polarization is aligned such that $\phi_H \approx 0$.

If we are looking for $\mathcal{NE}$-correlated systematics, we are essentially looking for terms of

the form

$$\phi_{Q,\mathrm{sys}} \approx \left(\frac{D_Q}{2S_H}\right)^{\mathcal{N}\mathcal{E}} \approx \begin{cases} \sum_{P'} \left(\frac{S_Q}{S_H}\right)^{P'} \phi_Q^{P'\mathcal{N}\mathcal{E}}, & \text{if } \phi_Q \approx 0, \\ \pm\left(\frac{S_Q}{2S_H}\right)^{\mathcal{N}\mathcal{E}}, & \text{if } \phi_Q - \phi_H \approx \frac{\pi}{4} \bmod \frac{\pi}{2}, \\ \sum_{P'} \left(\frac{S_Q}{S_H}\right)^{P'} \left(\frac{\sin(2\phi_Q)}{2}\right)^{P'\mathcal{N}\mathcal{E}}, & \text{else.} \end{cases} \tag{5.34}$$

Here the sum is performed over all parity channels P'. The key takeaway is that we must worry not only about an actual Q-state phase correlated with $\mathcal{N}\mathcal{E}$, but also about a Q-state signal size correlated with $\mathcal{N}\mathcal{E}$, especially when the non-reversing phases of the H and Q states are not well aligned.

We are now ready to present the systematic error model for the dark state in the "Qleanup" step in the second STIRAP region. Between the field plates and the entrance of the magnetic shields, the electric and magnetic fields are small. As a result, there is no preferred reference frame, and we work entirely in the lab frame. For simplicity, we assume the magnetic field is zero. Let there be a stray non-reversing electric field $\mathcal{E}_Z$ along the Z direction and a non-reversing transverse electric field $\mathcal{E}_\perp$ in some direction in the XY plane. In combination with coupling from the Ω-doublet splitting, we show that this configuration can generate both a $\mathcal{N}\mathcal{E}$-correlated signal and a non-reversing phase, which together produce a $\mathcal{N}\mathcal{E}$-correlated phase according to equation 5.34.

The level diagram of this setup is shown in Fig. 5.5. Without loss of generality, consider starting in the upper doublet state with $M = 0$. Because the $M = 0$ states contain both $\pm\Omega$ components, they can in general couple to all four $|M| = 1$ states via transverse electric fields. However, for a fixed direction of $\mathcal{E}_\perp$, only two of the four $|M| = 1$ states are coupled. This is illustrated on the right side of the diagram, where the basis has been redefined using the angle $\phi_\perp$ between $\mathcal{E}_\perp$ and the lab X axis.

The transfer rate from the $M = 0$ state to the coupled $|M| = 1$ states can be estimated using perturbation theory as

$$\frac{\mathcal{E}_\perp^2}{\mathcal{E}_Z \pm \Delta_\Omega/2},$$

where the electric fields are expressed in frequency units. This rate is explicitly upper/lower-state dependent and therefore depends on ΩM. If $\mathcal{E}_Z$ and $\mathcal{E}_\perp$ do not reverse with the interaction-region electric field, this mechanism generates an $\mathcal{N}\mathcal{E}$-dependent population. The effect is maximized in the regime $\mathcal{E}_Z \sim \Delta_\Omega$, which is likely relevant in the chamber region upstream of the field plates. If the phase defined by $\mathcal{E}_\perp$ differs from the typical non-reversing phase ϕ^{nr} of the ThO molecules, this appears as a systematic according to Eq. 5.34.

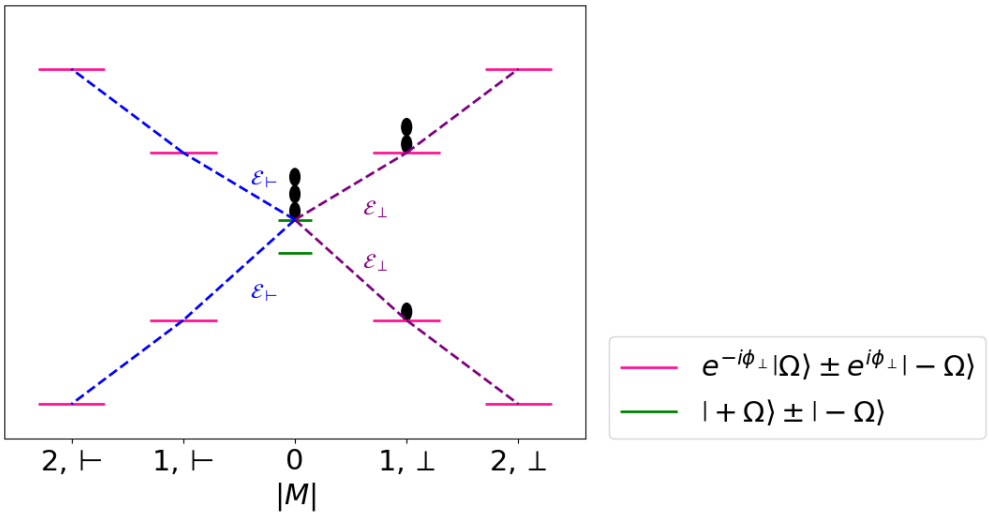

Figure 5.5: Level diagram for illustrating $NE = \Omega M$ dependent mixing. For simplicity, the only Hamiltonian terms included are a quantization field $\mathcal{E}_Z$, a transverse field $\mathcal{E}_\perp$, and a Δ_Ω splitting. The angle $\phi_\perp$ is the angle between the transverse field and the lab X axis.

In sum, we showed a simple scenario in which the time evolution of the STIRAP II dark state in stray electric fields can lead to EDM systematics. This phenomenon is in fact quite general when stray fields in other directions are included, because we operate in the regime $\mathcal{E}_{\mathrm{stray}} \sim \Delta_\Omega \sim 1/T_{\mathrm{flight}}$. In this regime, the ΩM dependence is maximized and the flight time is not long enough to ensure full mixing. This also makes estimating the possible size of this effect difficult. The time evolution between the first two STIRAP

chambers should be mostly adiabatic, and the population remaining in the dark state should be small, so this effect is expected to be small for ACME III. Indeed, tests performed by Collin Diver suggested that although we did observe systematics from the decay of the Q state, the dominant contribution likely comes from an ellipticity–ellipticity coupling, in which an effective ellipticity imperfection $d\Theta^{\mathcal{PR}}$ of the readout couples to $(S/I)^{\mathcal{NE}}$ of the Q state to produce an apparent systematic phase. This is not an effect of coherent phase-preserving decay, but rather a population effect.

CHAPTER 6

SUMMARY AND OUTLOOK

6.1 Summary

In this thesis, we have presented the progress of the ACME III experiment, which aims to measure the electron electric dipole moment (eEDM) with unprecedented sensitivity. In chapter 1, we motivated the eEDM search as a probe of physics beyond the Standard Model; molecules with large internal effective electric fields provide ideal platforms for such searches. In chapter 2, we described the ACME III experimental method and apparatus in detail, with emphasis on the changes and improvements relative to the previous ACME II experiment. In chapter 3, we explained our data analysis and blinding procedures, focusing on new protocols implemented in this generation of ACME. In chapter 4, we described our methods for identifying and controlling systematic errors in the experiment. Due to the increased sensitivity, many new effects have been observed and understood. One of the largest systematic effects arises from the metastable Q state decaying into the science H state. The detailed modeling of this effect in ACME is presented in chapter 5. Currently, the ACME experiment is in the final phase of systematic investigations before beginning the acquisition of final eEDM data. The next section summarizes our progress on systematic studies and the outlook moving forward.

6.2 Outlook

We began taking beamline data in early 2024 after finishing the major hardware construction. Since then, we have been carefully characterizing the system and studying hidden imperfections that we were previously unaware of. Around November 2025, we felt that most of the unexpected imperfections in the system had been identified and understood. We then began taking long IPV data sets while checking the statistical sensitivity of the

system. Figure 6.1 shows some preliminary data taken during this systematics investigation phase. We have demonstrated a factor of 30 improvement in statistical sensitivity per unit time compared to ACME II, sufficient to reach the ACME II statistical limit in around 40 minutes.

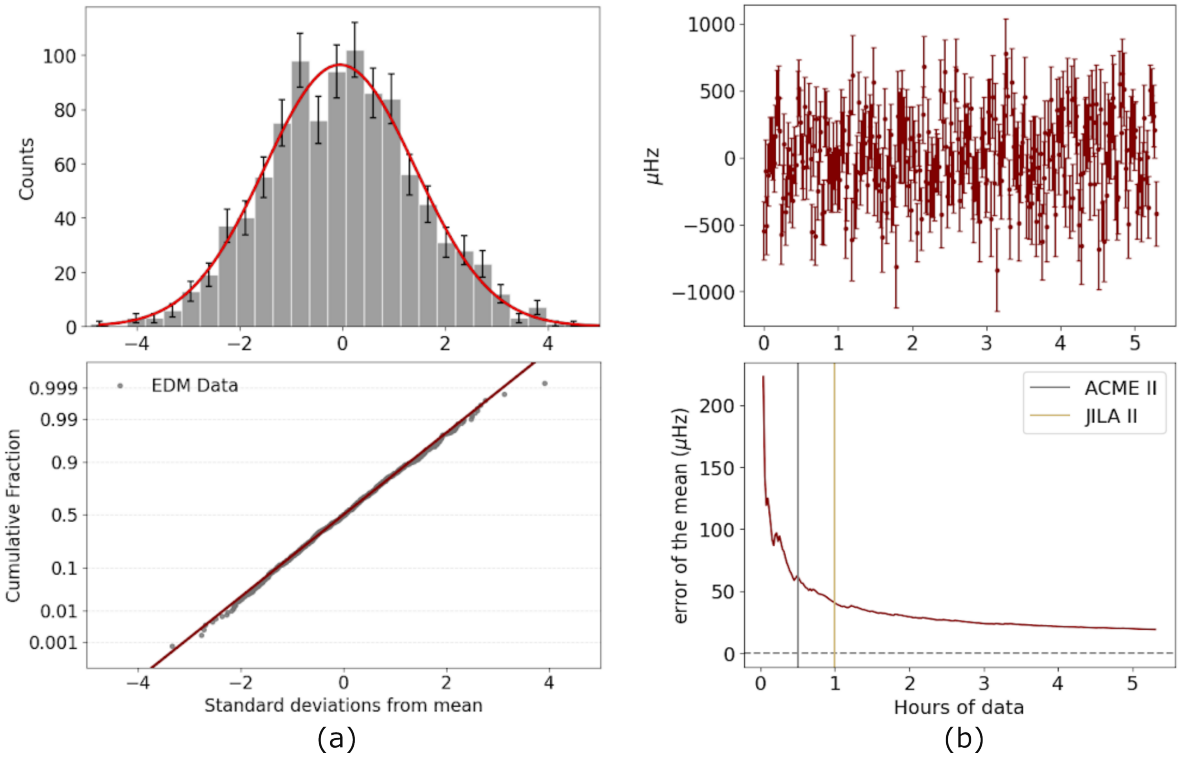

Figure 6.1: Preliminary data taken during the systematics investigation. (a) Top: histogram of centered and normalized eEDM measurements. Bottom: normal probability plot compared with an ideal normal distribution. (c) Top: mean-removed eEDM measurements collected as a function of experiment time. Bottom: error of the mean as a function of averaging time. It takes around 40 minutes to reach the final statistical sensitivity of ACME II and 60 minutes to reach that of JILA II.

We summarize our progress on the intentional parameter variation (IPV) systematics investigation in Table 6.1. In ACME II, over 40 parameters were intentionally varied to look for possible correlations with the eEDM channel. In ACME III, we have added around another 20 parameters to the IPV list due to the increased complexity of the system. So far, we have collected data for all but 3 parameters. It is important to point out that not all parameters have an ideal value for which we can confidently say that no systematic errors

can occur even if we average down for a month. This means we cannot rigorously bound possible systematic errors from these parameters, as explained in 4.1.1. Our approach is to try as hard as we reasonably can to measure a non-zero effect. For most of these parameters, we collected half a day to a full day of IPV data; if we do not see any statistically significant correlation, we consider the parameter to be not concerning.

At the time of this writing, we are working to understand the correlation between the eEDM channel and the correlated detuning $\Delta^{\mathcal{NE}}_{\mathrm{STIRAP}}$ of the 1090 nm STIRAP laser. This parameter was shown to produce systematic errors in the previous generation through correlated center-of-mass shifts of the molecular beam coupling to magnetic field gradients [23]. In this generation of the experiment, however, we have observed additional systematic effects even after zeroing the magnetic field gradients. Currently, we are actively investigating the root cause of this systematic effect and exploring possible mitigation strategies. If we do not successfully understand and mitigate this effect, we will have to conservatively assign a systematic uncertainty associated with this parameter, as done in other eEDM searches [43, 124]. This procedure would limit our final achievable sensitivity to approximately $1 \times 10^{-30}\ e \cdot \mathrm{cm}$, which is still a factor of 10 improvement over ACME II.

In sum, we have demonstrated a factor of 30 improvement in statistical sensitivity compared to the previous generation of ACME. We have completed most of the systematic error investigations and are working to understand what we believe to be the last problematic parameter, $\Delta^{\mathcal{NE}}_{\mathrm{STIRAP}}$. We expect to finish this investigation and start the final eEDM data collection in the next few months. Depending on the outcome of the $\Delta^{\mathcal{NE}}_{\mathrm{STIRAP}}$ investigation, the final data set will consist of approximately one to two months of data collection, and we expect to reach a final sensitivity of at least $1 \times 10^{-30}\ e \cdot \mathrm{cm}$, which is a factor of 4 improvement over the current best limit [20].

Parameters	Data collected	10× limit
Lasers: Pointing and position		
Differential X,Y translation	✓	×
Differential horizontal pointing	✓	✓
Differential vertical pointing	✓	✓
Precession time	✓	×
Laser propagation direction $\hat{k}$	✓	×
Cleanup pointing	✓	✓
Cleanup translation	✓	×
Lasers: Detuning		
X, Y differential detuning	✓	✓
Cleanup detuning	✓	✓
$\Delta^{\mathrm{nr}}, \Delta^{\mathcal{P}}, \Delta^{\mathcal{N}}, \Delta^{\mathcal{NE}}$ from 703 cleanup and readout	✓	✓
$\Delta^{\mathcal{NE}}$ from 1090 STIRAP	×	×
Δ^{nr} from 1196 lens STIRAP	✓	✓
Δ^{nr} from 690 lens STIRAP	✓	✓
Δ^{nr} from 746	✓	✓
746 91 MHz sideband on/off	✓	✓
Rotational cooling detuning	✓	✓
Lasers: Power		
Cleanup power, readout power	✓	×
X,Y power asymmetry	✓	✓
$\tilde{\mathcal{P}}$ state power asymmetry	✓	×
Lasers: Polarization		
Cleanup polarization rotation	✓	×
Cleanup ellipticity	✓	×
Probe ellipticity	×	×
Probe polarization rotation, dither angle	✓	×
Lens STIRAP HWP, QWP	✓	×
Electric Field		
$\mathcal{E}$ field magnitude $\lvert\mathcal{E}\rvert$	✓	×
V_{offset} on field plates	✓	✓
$\mathcal{E}^{\mathrm{nr}}$	✓	×
Magnetic Field		
$B_z^{\mathrm{nr}}, B_z^{\mathcal{B}}$	✓	✓
$B_x^{\mathrm{nr}}, B_y^{\mathrm{nr}}, B_x^{\mathcal{B}}, B_y^{\mathcal{B}}$	✓	✓
$\frac{\partial B_x}{\partial x}^{\mathrm{nr}}, \frac{\partial B_y}{\partial y}^{\mathrm{nr}}, \frac{\partial B_y}{\partial x}^{\mathrm{nr}}, \frac{\partial B_z}{\partial x}^{\mathrm{nr}}$	✓	✓
$\frac{\partial B_x}{\partial x}^{\mathcal{B}}, \frac{\partial B_y}{\partial x}^{\mathcal{B}}, \frac{\partial B_y}{\partial y}^{\mathcal{B}}, \frac{\partial B_y}{\partial z}^{\mathcal{B}}, \frac{\partial B_z}{\partial x}^{\mathcal{B}}, \frac{\partial B_z}{\partial z}^{\mathcal{B}}$	✓	✓
Other		
Q state decay	✓	✓
Block switches settling times	×	×
Polarization switching frequency	✓	×

Table 6.1: IPV progress as of January 2026. We have collected data for more than 50 experimental parameters. The last column indicates whether a limit can be placed that is 10 times smaller than the ACME II systematic error budget.

APPENDIX A

MODELING DEGAUSSING OF MAGNETIC SHIELDS

During the degaussing investigation led by Xing Fan, we observed several phenomena that were initially not understood. In particular, we observed that

1. the DC shielding factor is much higher than the AC shielding factor. In other words, after degaussing, the shield magnetizes itself in a way that cancels external fields;

2. a demagnetized shield can be magnetized by a very small DC current offset in the degaussing coils. Suppose this corresponds to an applied field of δH. If the same field is applied directly to the demagnetized shield without degaussing, no irreversible magnetization occurs. This indicates a strong nonlinearity in the magnetization process;

3. the residual magnetic field drift is smaller when the degaussing performance is better.

We argue that the first two observations can be explained using the two most common hysteresis models: the Jiles–Atherton model and the Preisach model. The third observation is more difficult to interpret, as the drift may be caused by a variety of effects, such as temperature fluctuations and mechanical stress changes.

In this appendix, we review the basics of these two hysteresis models as they apply to degaussing and use them to explain the observed phenomena. Readers not interested in the details may skip directly to the conclusion section, which is written to be understandable without reading the intervening sections.

A.1 The Jiles–Atherton hysteresis model

A.1.1 Pros

The main advantages of the Jiles–Atherton model include the following:

1. Simplicity: hysteresis is described by a relatively simple ordinary differential equation, which is a familiar and convenient framework for physicists.

2. Implementation in COMSOL: under the Ampère's law node, COMSOL includes a built-in implementation of the Jiles–Atherton model. In addition, the material library provides default Jiles–Atherton materials, allowing one to quickly begin simulating and exploring basic properties of magnetic hysteresis.

3. Low memory requirement: compared to the Preisach model, which requires a large number of hysterons for accurate results, the Jiles–Atherton model has significantly lower memory requirements.

A.1.2 Description of the model

Since COMSOL has already implemented the Jiles–Atherton model, we follow the notation used in the COMSOL documentation. Readers interested in additional details are referred to Jiles [125].

We are interested in predicting how a ferromagnetic material changes its magnetization in response to the magnetizing field $\mathbf{H}$. We define the effective field experienced by the material as

$$\mathbf{H}_e = \mathbf{H} + \alpha \mathbf{M}, \tag{A.1}$$

where α is the inter-domain coupling constant.

The total magnetization $\mathbf{M}$ is written as the sum of a reversible part and an irreversible part. The reversible part is described by the anhysteretic magnetization curve. The anhysteretic magnetization as a function of the effective magnetizing field $\mathbf{H}_e$ is given by

$$\mathbf{M}_{\text{an}} = M_s \, L\Big(\Big|a^{-1}\mathbf{H}_e\Big|\Big) \, \frac{\mathbf{H}_e}{|\mathbf{H}_e|}, \tag{A.2}$$

where $L(x) = \coth(x) - 1/x$ is the Langevin function, M_s is the saturation magnetization, and a characterizes the domain wall density.

We are now ready to describe the change in the total magnetization $d\mathbf{M}$ as the sum of an irreversible contribution and a reversible contribution:

$$d\mathbf{M} = \max(\boldsymbol{\chi} \cdot d\mathbf{H}_e,\, 0)\, \frac{\boldsymbol{\chi}}{|\boldsymbol{\chi}|} + c_r\, d\mathbf{M}_{\text{an}}, \tag{A.3}$$

where c_r is the magnetization reversibility. The vector

$$\boldsymbol{\chi} = k_p^{-1}\, (\mathbf{M}_{\text{an}} - \mathbf{M})$$

is referred to as the auxiliary vector. It measures the distance between the total magnetization $\mathbf{M}$ and the anhysteretic magnetization $\mathbf{M}_{\text{an}}$ in units of the pinning parameter k_p, which characterizes pinning losses. Equation (A.3) is the ordinary differential equation governing the evolution of the magnetization. Given the current magnetization $\mathbf{M}$, the applied field $\mathbf{H}$, and an incremental change $d\mathbf{H}$, the corresponding change in magnetization can be computed from Eq. (A.3).

In summary, this phenomenological model contains five free parameters:

1. magnetization reversibility: c_r (dimensionless);
2. saturation magnetization: M_s (units of A/m);
3. domain wall density parameter: a (units of A/m);
4. pinning loss parameter: k_p (units of A/m);
5. inter-domain coupling constant: α (dimensionless).

These parameters are scalar constants by default, but in COMSOL they can be generalized to tensors or made dependent on quantities such as magnetization and temperature. For

experimental methods used to determine these parameters, see Chapter 7 of Jiles's book [125].

A.1.3 Understanding degaussing through the Jiles–Atherton model

The action of degaussing is straightforward to interpret within the Jiles–Atherton framework. During degaussing, the evolution of the system is dominated by the irreversible contribution to the magnetization change, corresponding to the first term in Eq. (A.3). Note that the vector

$$\boldsymbol{\chi} = k_p^{-1} \left(\mathbf{M}_{\mathrm{an}} - \mathbf{M} \right)$$

points from the current magnetization toward the anhysteretic magnetization. Physically, this term represents an effective attraction of the system toward the anhysteretic state. In addition, the rate of approach to the anhysteretic curve decreases linearly with the distance from it through the factor $\boldsymbol{\chi} \cdot d\mathbf{H}_e$. Intuitively, this suggests that after a sufficiently well-designed degaussing waveform, the system will converge toward the center of attraction, namely $\mathbf{M}_{\mathrm{an}}$[1].

Note that the ratio between the AC and DC shielding factors is determined by the magnetization reversibility c_r. To see this, consider a system that has been fully degaussed so that $\boldsymbol{\chi} = 0$. In this case, the first term in Eq. (A.3) vanishes, and the magnetization change reduces to

$$d\mathbf{M} = c_r \, d\mathbf{M}_{\mathrm{an}}.$$

In contrast, when the system is not degaussed, one has $d\mathbf{M} = d\mathbf{M}_{\mathrm{an}}$. Therefore, the ratio of the AC to DC permeabilities, $\mu_{\mathrm{AC}}/\mu_{\mathrm{DC}}$, is set by the relative change in magnetization and is simply given by c_r. This behavior is confirmed by our numerical simulations in COMSOL.

1. A fully rigorous argument is complicated by the vector nature of COMSOL's implementation. For the scalar formulation, the proof is straightforward.

The Jiles–Atherton model also provides a sufficient condition for reversibility. From Eq. (A.3), the irreversible contribution vanishes when $\boldsymbol{\chi} = 0$, which occurs when the magnetization coincides with the anhysteretic magnetization, i.e., after perfect degaussing. This explains the second phenomenon described in the introduction. A fully degaussed shield is difficult to magnetize using fields much smaller than the pinning loss. However, if the same small field is applied as a DC offset during the degaussing process, the magnetization follows the anhysteretic curve and the shield acquires a magnetization increment $d\mathbf{M}_{\text{an}}$. After the DC offset is removed, a residual magnetization of $(1 - c_r)\, d\mathbf{M}_{\text{an}}$ remains.

A.2 The Preisach hysteresis model

A.2.1 Pros

The main advantages of the Preisach model include the following:

1. Generality: the Preisach model describes hysteresis using a distribution of bistable elements, or hysterons. By appropriately choosing this distribution, one can model a wide range of hysteretic phenomena, not limited to ferromagnetism.

2. Existing implementations: the Preisach model is already implemented in both MATLAB and Python.

3. Amenability to thermodynamic analysis: thermodynamic treatments based on the Preisach model have been developed [126]. In particular, under certain simplifying assumptions, degaussing can be shown to gradually drive the system toward a global free-energy minimum.

A.2.2 Description of the model

In this subsection, we follow the notation of Bertotti [126].

We begin by motivating the use of bistable elements. Consider a binary bistable switch, shown in Fig. A.1.

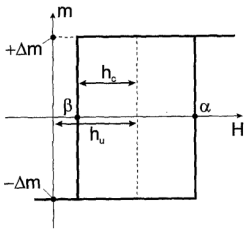

Figure A.1: Possible values m of a binary bistable switch.

This switch outputs a binary value ($\pm\Delta m$) in response to the input field H. Its hysteresis loop is centered at $H = h_u$. Outside the indeterminate region, defined by $h_u - h_c < H < h_u + h_c$, the switch has a unique output. Specifically, if H is sufficiently low ($H < h_u - h_c$), the switch outputs $-\Delta m$, whereas if H is sufficiently high ($H > h_u + h_c$), the switch outputs $+\Delta m$.

Within the indeterminate region, the switch can take either of the two values ($\pm\Delta m$), with the actual value determined by the history of the switch. In practice, the switch retains its previous value. For example, if the switch output is Δm before entering this region, it remains Δm as H varies within the indeterminate region. Formally, this behavior is described by

$$m(H(t),t) = \begin{cases} \Delta m, & \text{if } H(t) \geq h_u + h_c, \\ -\Delta m, & \text{if } H(t) \leq h_u - h_c, \\ m(H(t-dt),t-dt), & \text{if } h_u - h_c < H(t) < h_u + h_c. \end{cases} \tag{A.4}$$

The input–output relation of this bistable switch exhibits hysteresis, and the dependence on its history is explicit in Eq. (A.4). To model macroscopic hysteresis more accurately, a single switch is insufficient. Instead, we consider an ensemble of such switches spanning all possible values of h_u and h_c. This ensemble can be represented in the (h_c, h_u) plane, commonly referred to as the Preisach plane, shown in Fig. A.2.

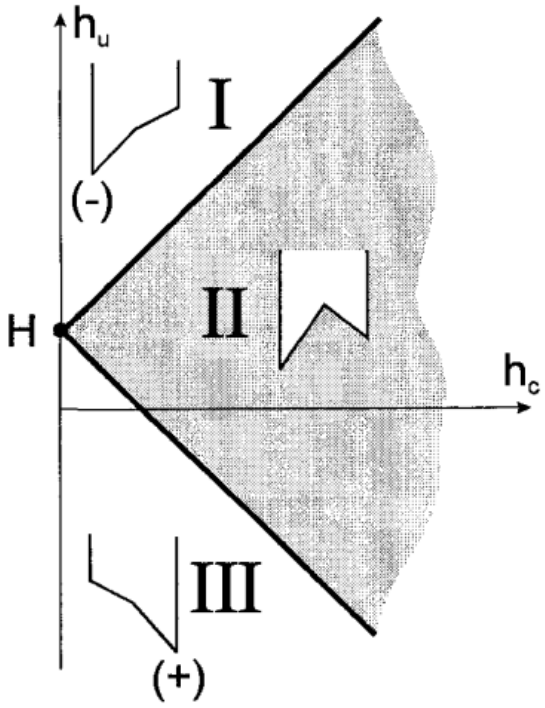

Figure A.2: The right half of the (h_c, h_u) plane (the Preisach plane) contains all possible bistable switches defined in Eq. (A.4). For a given value of H, the plane is divided into three regions according to Eq. (A.4). In regions I and III, the switches output $-\Delta m$ and $+\Delta m$, respectively. In region II, the switch states are determined by their history. The physical picture can be understood by considering the sum of a symmetric double-well potential and a linear potential proportional to H, as illustrated in the figure. In regions I and III, the linear potential dominates, resulting in a single minimum and hence a single-valued output. In the indeterminate region II, the field H is insufficient to overcome the double-well potential, so two minima exist, and a classical particle in this potential remains in its previous minimum.

For each possible pair of switch parameters (h_u, h_c), the density of switches may differ. We denote this density by $p(h_c, h_u)$, which is referred to as the Preisach distribution. By symmetry considerations, the Preisach distribution must be even in h_u, i.e.,

$$p(h_c, h_u) = p(h_c, -h_u). \tag{A.5}$$

Next, we introduce the concept of the Preisach state line. Suppose the system starts in a completely magnetized (saturated) state. In this case, the switches can be partitioned

by a single zigzag line, called the Preisach state line $b(h_c)$, as illustrated in Fig. A.3. All switches below $b(h_c)$ take the value $+\Delta m$, while all switches above it take the value $-\Delta m$.

To determine $b(h_c)$, one applies the definition in Eq. (A.4) sequentially as the field H evolves. At first glance, this may appear impractical, since the definition would need to be applied an infinite number of times. However, it is sufficient to apply the update rule only at the turning points of H, as well as at the initial and final values of H. A graphical illustration of one step in this procedure is shown in Fig. A.2. An explicit example of constructing the state line $b(h_c)$ is shown in Fig. A.4.

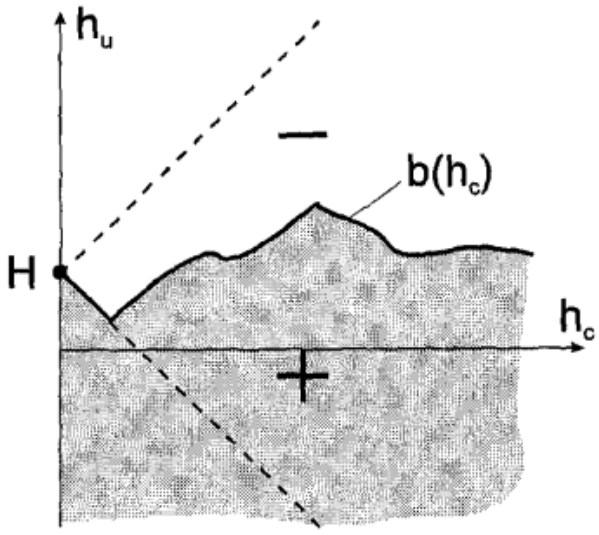

Figure A.3: Example Preisach state line.

Together with the Preisach distribution, the switch states in the Preisach plane completely determine the macroscopic magnetization, which can be written as[2]

$$M = \Delta m \int_0^\infty dh_c \left[\int_{-\infty}^{b(h_c)} p(h_c, h_u)\, dh_u - \int_{b(h_c)}^{+\infty} p(h_c, h_u)\, dh_u \right] \tag{A.6}$$

$$= 2\Delta m \int_0^\infty dh_c \int_0^{b(h_c)} p(h_c, h_u)\, dh_u, \tag{A.7}$$

2. We deviate from the normalization conventions used by Bertotti by constant factors for simplicity. See Chapter 13 of Ref. [126] for details.

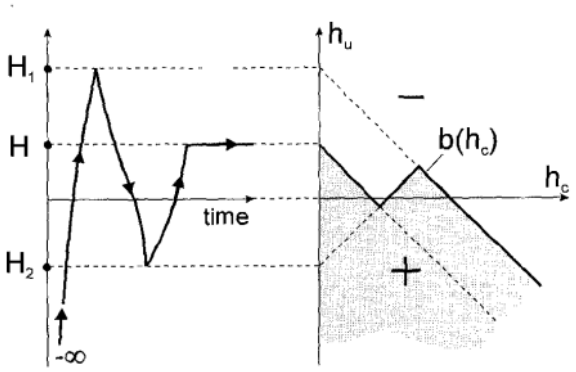

Figure A.4: An example of a Preisach state line $b(h_c)$ (right) generated by the magnetizing field shown on the left. The state line is completely determined by the turning points of $H(t)$, including the initial and final values. One can verify that $b(h_c)$ is generated according to the definition given in Eq. (A.4).

where the second line follows from the symmetry of the Preisach distribution given in Eq. (A.5). In other words, the magnetization M is proportional to the weighted area under the Preisach state line $b(h_c)$, with the weighting determined by $p(h_c, h_u)$.

Finally, we consider the Preisach state line generated by a degaussing waveform, shown in Fig. A.5. In general, the state line produced by a degaussing waveform superposed with a DC offset field H_a has a sawtooth-like structure centered at H_a. The size of the sawteeth is determined by the difference between successive extrema of the degaussing field. As this difference approaches zero, the state line converges to $b(h_c) \to H_a$.

A.2.3 Understanding degaussing through the Preisach model

In the previous subsection, we noted that the Preisach state line $b(h_c)$ produced by an ideal degaussing process is approximately $b(h_c) = 0$. It follows immediately from Eq. (A.7) that this corresponds to a state of zero net magnetization. Thus, although the Preisach framework requires some setup, the description of degaussing within this model is straightforward.

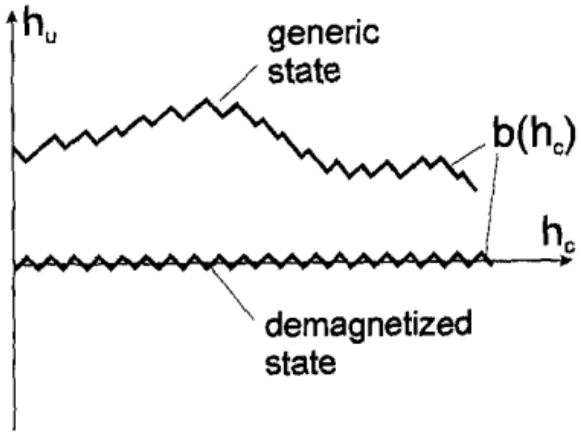

Figure A.5: Additional examples of Preisach state lines. The line oscillating about the h_c axis is generated by a degaussing waveform. The outermost sawteeth correspond to the largest applied fields. As the field amplitude decreases while oscillating, additional sawteeth are generated progressively from right to left. A degaussing waveform with an added DC bias H_a produces an identical state line, shifted vertically by H_a.

The Preisach model also provides a clear explanation for the enhancement of the effective permeability along the anhysteresis curve. Consider the example waveforms shown in Fig. A.6. To demonstrate this enhancement, it suffices to show that the magnetization produced by a biased degaussing waveform is larger than that produced by a simple linear ramp[3]. This is immediately evident from the corresponding Preisach planes shown in Fig. A.7. The weighted area under the state line $b(h_c)$ is clearly larger for the biased degaussing waveform, leading to a larger magnetization.

Within the Preisach model, the enhancement of the effective permeability arises naturally from the collective behavior of bistable switches. The increase in μ results from the progressive alignment of switches that require increasingly large fields to change state. Without degaussing, the system generally becomes trapped in local energy minima with only partial alignment, leading to a smaller magnetization. In contrast, the Jiles–Atherton model incorporates the effective permeability enhancement phenomenologically through the

3. As discussed earlier, only the turning points of the field are relevant, so the detailed shape of the ramp is unimportant. We also implicitly assume that the magnetization M depends linearly on H both before and after degaussing; discussion of an effective permeability is only meaningful under this assumption.

constant c_r. From this perspective, the Preisach model provides a more transparent physical interpretation of anhysteresis.

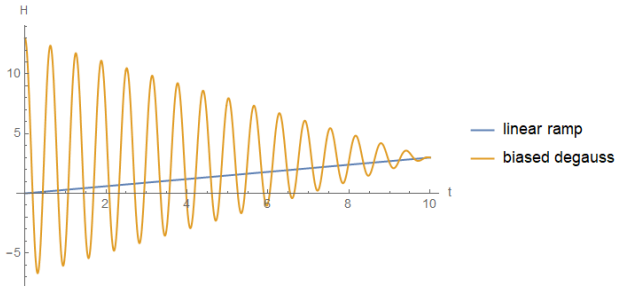

Figure A.6: Comparison between a linear ramp to $H = 3$ and a degaussing waveform with a DC offset of 3.

A.2.4 Free energy minimization

Studying the magnetization of ferromagnetic materials is challenging due to the strong internal couplings present in these systems. Such couplings give rise to many local minima and metastable states in the free-energy landscape. Within the Preisach framework, this metastability is modeled through an ensemble of bistable switches. In this picture, each Preisach state line $b(h_c)$ corresponds to a particular metastable state of the ferromagnetic material. Finding the global energy minimum of the system is therefore equivalent to identifying the state line $b(h_c)$ that minimizes the free energy.

This problem is treated in detail by Bertotti [126]; here we summarize only the main conclusions. When the mean-field energy contributions (such as magnetostatic and domain-wall energies) are small compared to the coupling to the external magnetic field, the free energy is minimized after degaussing. However, when the mean-field terms are non-negligible, degaussing does not necessarily drive the system to the global energy minimum. This result is somewhat counterintuitive and highlights the complexity of the magnetic energy landscape. Readers interested in a more detailed discussion are referred to Chapters 13 and 14

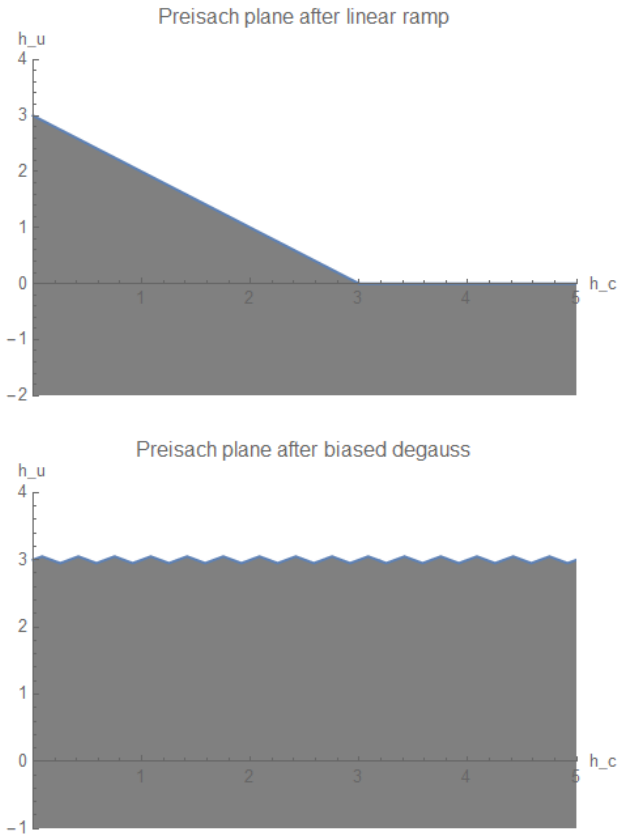

Figure A.7: Resulting Preisach planes corresponding to the waveforms shown in Fig. A.6.

of Bertotti[4][126].

A.3 Conclusion

In summary, we believe we have a reasonable understanding of the origins of the higher DC shielding factor and the nonlinear response to degaussing offsets described in the introduction.

The enhanced DC shielding factor is a well-known phenomenon in the literature. After degaussing, the magnetization follows the so-called anhysteresis curve, as illustrated in Fig. A.8. The effective permeability of this curve is generally higher than that of the initial

4. Strictly speaking, because the temperature is finite, the system is never guaranteed to reach the global minimum. Nevertheless, the conclusion remains somewhat surprising.

magnetization curve, resulting in an increased DC shielding factor. A clear discussion of this behavior is given by Yashchuk *et al.* [127]. The underlying physical mechanism is most transparently explained within the Preisach model. The magnetic system possesses many local energy minima in the space of possible magnetization states. When an external field is applied, the system may become trapped in one of these local minima before reaching the global minimum magnetization $M_{\text{global min}}$, leading to a reduced magnetization and hence a smaller effective permeability. Degaussing effectively provides the activation needed to help the system escape local minima and approach the global minimum.

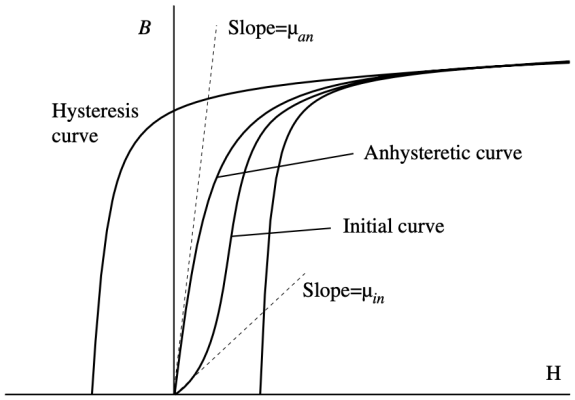

Figure A.8: Various magnetization curves that a ferromagnetic material may follow. If the shield is initially demagnetized, then upon application of a magnetizing field H it follows the initial magnetization curve. If a degaussing sequence is performed while the same magnetizing field H is applied, the system instead follows the anhysteresis curve. In general, the effective permeability μ_{an} of the anhysteresis curve is larger than the permeability μ_{in} of the initial curve.

The degaussing offset nonlinearity can also be understood in terms of anhysteresis, together with an additional statement about reversibility. This mechanism is most transparent in the Jiles–Atherton model. Once the system is sufficiently degaussed[5], the magnetization response becomes reversible. Consider the application of a small external field δH. If the shields are initially demagnetized, the reversibility condition implies that δH alone is insufficient to produce any irreversible magnetization.

5. As noted earlier, the criterion for sufficient degaussing is set by the pinning loss.

In contrast, if the same field δH is applied as a DC offset during the degaussing process, the magnetization follows the anhysteresis curve and acquires a magnetization $\mu_{\mathrm{an}}\,\delta H$, where μ_{an} is the effective relative permeability of the anhysteresis curve. When the DC offset field δH is subsequently removed, the magnetization decreases by $\mu_{\mathrm{in}}\,\delta H$, where μ_{in} is the relative permeability of the initial magnetization curve near $H = 0$, as illustrated in Fig. A.8. Since both the Jiles–Atherton and Preisach models predict that $\mu_{\mathrm{an}} > \mu_{\mathrm{in}}$, a residual magnetization remains. Thus, a DC offset applied during degaussing can produce a net magnetization, even though the same offset field applied by itself is incapable of magnetizing the shield.

APPENDIX B

FIELD PLATE BENDING STRESS MECHANICS

B.1 Background

In ACME II, it was observed that the field plates were bent. We suspected that this bending might be responsible for the observed birefringence in the field plates. Here we argue that this is not the case. A simple (quadratic) bending profile, such as the one shown in Fig. 2.21, cannot produce birefringence. Any birefringence must instead arise from applied tangential forces near the edges of the glass field plates. In addition, we present a set of design targets that, if satisfied, will result in zero birefringence from clamping-induced stress.

B.2 Boundary value problem

For problems in linear elasticity, the governing equations consist of the equilibrium equations

$$\sigma_{ij,j} = 0, \tag{B.1}$$

which express force balance at mechanical equilibrium. This equation states that an infinitesimal volume element of the solid is not accelerating. For simplicity, we neglect body forces such as gravity. In addition to the equilibrium equations, the stress and strain fields must satisfy the stress compatibility equations,

$$\epsilon_{ijk}\,\epsilon_{pqr}\,\frac{\partial^2 e_{jq}}{\partial x_k \partial x_r} = 0,$$

where e_{jq} is an element of the infinitesimal strain tensor, describing how an infinitesimal volume deforms under stress. The strain tensor is related to the stress tensor through the

stress–strain relation (Hooke's law),

$$\sigma_{ij} = \lambda\, \delta_{ij}\, e_{mm} + 2\mu\, e_{ij},$$

where λ and μ are material constants determined by Young's modulus and Poisson's ratio. We do not assign equation numbers to these relations, as they are not central to the discussion that follows and are included primarily for completeness.

Given this set of partial differential equations, appropriate boundary conditions are required to ensure existence and uniqueness of a solution. Consider solving for the stress tensor within a body bounded by a closed surface S. Intuitively, if the forces acting on the surface are known, or if the deformation of the surface from its natural state is specified, the stress field within the body is fully determined. Formally, let the surface be partitioned into two disjoint parts, S_T and S_u. A sufficient condition for existence and uniqueness is to specify the traction $\mathbf{T} = \boldsymbol{\sigma} \cdot \hat{\mathbf{n}}$ on S_T and the displacement $\mathbf{u}$ on S_u. The traction represents the force per unit area acting on the surface in the direction of the outward normal $\hat{\mathbf{n}}$, while the displacement describes the deviation of the surface from its undeformed, natural configuration.

B.3 Solving the boundary value problem

We have formulated the boundary value problem in a somewhat abstract way by allowing, at each point on the closed surface, either the traction or the displacement to be specified. This formulation is, however, well suited to the ACME II field plates. On the main faces (normal to ACME Z direction), the displacement is known[1]. At the edges (normal to the ACME X or Y directions), forces are applied by the clamping mechanism. If these

1. Strictly speaking, the interferometric measurement only determines differences in displacement. However, if the measured displacement difference has a quadratic dependence on x, it is difficult to envision a scenario in which each individual field plate does not also exhibit a quadratic profile in x. For the purposes of this discussion, we therefore assume that each field plate has a quadratic displacement profile.

forces were known exactly, the stress tensor could be uniquely determined. Although the edge forces are not known precisely, we can nevertheless establish a useful result under the following assumptions:

1. The rectangular field plate occupies the region $-L_x \le x \le L_x$, $-L_y \le y \le L_y$, and $-L_z \le z \le L_z$. Suppose that $\sigma_{xz} = \sigma_{yz} = 0$ on the main faces at $z = \pm L_z$. This assumption is valid for ACME because these faces are exposed to vacuum and must therefore be traction-free, i.e., $\boldsymbol{\sigma} \cdot \hat{\mathbf{z}} = 0$.

2. Suppose that
$$\int_{-L_z}^{L_z} \sigma_{xx}\, dz = \int_{-L_z}^{L_z} \sigma_{xy}\, dz = 0 \tag{B.2}$$
on the surfaces $x = \pm L_x$, and that
$$\int_{-L_z}^{L_z} \sigma_{yy}\, dz = \int_{-L_z}^{L_z} \sigma_{xy}\, dz = 0 \tag{B.3}$$
on the surfaces $y = \pm L_y$. Physically, this corresponds to a clamping mechanism that does not apply net tangential forces in the $\hat{z}$ direction. The integrals account for the total stress across the thickness of the field plates.

3. Assume that the stresses are approximately uniform across the main faces of the field plates. In particular, suppose that
$$\int_{-L_i}^{L_i} \sigma_{jk}\, di = 2L_i\, \sigma_{jk},$$
for $i, j, k \in \{x, y\}$. This assumption is only approximately satisfied for the ACME II field plates; deviations from it are discussed later.

With these assumptions, we show that the stress components capable of producing birefringence must average to zero when integrated across the thickness of the plate. In

other words,

$$\int_{-L_z}^{L_z} \sigma_{ij}\, dz = 0 \quad \text{for } i,j \in \{x,y\}.$$

This implies that a laser propagating along the Z direction does not accumulate any net birefringence after passing through the field plates.

Proof: We begin by explicitly writing the first equilibrium equation from Eq. (B.1),

$$\sigma_{xx,x} + \sigma_{xy,y} + \sigma_{xz,z} = 0. \tag{B.4}$$

We integrate this equation over an entire cross section normal to the x direction:

$$\begin{aligned}
0 &= \int_{-L_z}^{L_z}\int_{-L_y}^{L_y} \left(\sigma_{xx,x} + \sigma_{xy,y} + \sigma_{xz,z}\right)\, dy\, dz && \text{(B.5)}\\
&= \int_{-L_z}^{L_z}\int_{-L_y}^{L_y} \sigma_{xx,x}\, dy\, dz + \int_{-L_z}^{L_z} \sigma_{xy}\big|_{-L_y}^{L_y}\, dz + \int_{-L_y}^{L_y} \sigma_{xz}\big|_{-L_z}^{L_z}\, dy && \text{(B.6)}\\
&= \int_{-L_z}^{L_z}\int_{-L_y}^{L_y} \sigma_{xx,x}\, dy\, dz, && \text{(B.7)}
\end{aligned}$$

where the last equality follows from Assumptions 1 and 2, since σ_{xy} is evaluated at $y = \pm L_y$ and σ_{xz} at $z = \pm L_z$. Equation (B.7) allows us to perform the following integration:

$$\begin{aligned}
\int_{-L_z}^{L_z}\int_{-L_y}^{L_y} \sigma_{xx}(x_0,y,z)\, dy\, dz &= \int_{-L_z}^{L_z}\int_{-L_y}^{L_y} \sigma_{xx}(-L_x,y,z)\, dy\, dz\\
&\quad + \int_{-L_x}^{x_0}\int_{-L_z}^{L_z}\int_{-L_y}^{L_y} \sigma_{xx,x}\, dy\, dz\, dx, && \text{(B.8)}\\
&= \int_{-L_z}^{L_z}\int_{-L_y}^{L_y} \sigma_{xx}(-L_x,y,z)\, dy\, dz, && \text{(B.9)}\\
&= 0. && \text{(B.10)}
\end{aligned}$$

The step from Eq. (B.8) to Eq. (B.9) follows directly from Eq. (B.7), while the final equality follows from Assumption 2. Physically, Assumption 2 states that no net forces tangential

to $\hat{z}$ are applied at the edges of the field plates. Equation (B.7) implies that, for any cross section normal to x, the total force in the x direction obtained by integrating the stress over that cross section is independent of x. Since this force vanishes at the edges, it must vanish for all cross sections normal to x.

We have not yet made use of Assumption 3. This assumption states that σ_{xx} is uniform across y, which allows us to write

$$\int_{-L_z}^{L_z} \sigma_{xx}\, dz = \frac{1}{2L_y} \int_{-L_y}^{L_y} \int_{-L_z}^{L_z} \sigma_{xx}\, dz\, dy = 0. \tag{B.11}$$

Using essentially the same argument, with appropriately chosen cross sections, one can similarly show that

$$\int_{-L_z}^{L_z} \sigma_{xx}\, dz = \int_{-L_z}^{L_z} \sigma_{xy}\, dz = \int_{-L_z}^{L_z} \sigma_{yy}\, dz = 0. \tag{B.12}$$

This completes the proof that the stress components capable of producing birefringence must average to zero across the thickness of the field plates.

B.4 Implications for field plate clamping design

In the previous section, we identified a set of conditions that, if satisfied, guarantee zero birefringence arising from clamping stress. We now discuss these conditions in the context of ACME II and ACME III. Condition 1 is always satisfied. Condition 2 must be violated in ACME II. In the previous section, we showed that if Condition 2 were satisfied, the birefringence averaged over a cross section normal to y would vanish. However, the ACME II field plates exhibited an approximately constant $\sim 2\%$ birefringence offset along such cross sections [88]. This observation implies the presence of a nonzero clamping force with a tangential component along the $\hat{z}$ direction.

With improved mechanical design, Condition 2 can be approximately satisfied. Condition 3 is only approximately satisfied in ACME II, as the stress-induced birefringence varied

by roughly 20% across the field plates. Note, however, that the stress component responsible for bending should be nearly uniform. This follows from the fact that the observed bending profile is close to quadratic, corresponding to an approximately constant radius of curvature of about 40 m across the plates. A nearly constant radius of curvature implies a nearly uniform bending stress. The observed nonuniformity is therefore most likely due to spatial variations in the clamping forces.

In ACME III, we aim to better satisfy Condition 3 by distributing the clamping force over as many contact points as is practical. Ultimately, we experimentally verified that Conditions 2 and 3 are sufficiently well satisfied by constructing and testing a dedicated clamping setup.

APPENDIX C

MOLECULAR POLARIMETRY PARITY CHANNELS

C.1 Overview

In this appendix, we work out the contributions to the various parity channels arising from the polarimetry switches $\mathcal{N}$, $\mathcal{E}$, $\mathcal{Q}$, $\mathcal{B}$, $\mathcal{P}$, and $\mathcal{R}$. A new switch, $\mathcal{Q}$, is introduced; it rotates the quarter-wave plate (QWP) used in polarimetry by 90°, thereby flipping the sign of the measured asymmetry. This provides a means to separate true ellipticity from possible systematic effects.

We perform polarimetry in two modes. In the preparation polarimetry mode, the readout light is fully circularly polarized. In the readout polarimetry mode, the preparation light is fully circularly polarized.

To determine the contributions to the different parity channels, we first analyze the imperfections generated when light propagates through birefringent materials in Sec. C.2. We then determine the total imperfections introduced by the full polarimetry optics in Secs. C.3 and C.4. Using these results, we compute the measured asymmetry in Sec. C.5. Finally, we summarize the contributions to the different parity channels in Sec. C.6.

We work in the circular polarization basis,

$$\hat{\epsilon}_i = -e^{-i\theta_i}\cos\Theta_i\,\hat{\epsilon}_{+1} + e^{i\theta_i}\sin\Theta_i\,\hat{\epsilon}_{-1}, \tag{C.1}$$

and also in the linear polarization basis,

$$\hat{\epsilon}_i = -e^{-i\psi_i}\cos\Psi_i\,\hat{\epsilon}_x + e^{i\psi_i}\sin\Psi_i\,\hat{\epsilon}_y. \tag{C.2}$$

The transformation from the circular basis to the linear basis is given by

$$U = \begin{pmatrix} -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{i}{\sqrt{2}} & -\frac{i}{\sqrt{2}} \end{pmatrix}. \tag{C.3}$$

In practice, we work in the vicinity of $\Theta, \Psi = \pi/4$, $\theta = 0, \pi/2$, and $\psi = \pm\pi/4$. In this limit, $d\theta$ corresponds to a small admixture of diagonal linear polarization, $d\Theta$ to a small admixture of circular polarization, $d\Psi$ to a small admixture of vertical or horizontal polarization, and $d\psi$ to a small admixture of diagonal linear polarization. As a convenient reference, the imperfections generated by different wave plates are summarized in Table C.1.

$\Theta_{\rm in}$	$\theta_{\rm in}$	Wave plate	$\Theta_{\rm out,0}$	$\theta_{\rm out,0}$	$d\Theta_{\rm out}$	$d\theta_{\rm out}$
$\frac{\pi}{4}$	0	HWP @ $\theta_{\rm HWP}$	$\frac{\pi}{4}$	$2\theta_{\rm HWP}$	$\frac{d\delta_{\rm HWP}}{2}\sin(2\theta_{\rm HWP})$	$-\frac{1}{8}d\delta^2_{\rm HWP}\sin(4\theta_{\rm HWP})$
$\frac{\pi}{4}$	0	QWP @ $\frac{\pi}{4} + d\theta_{\rm QWP}$	0	NA	$\frac{1}{2}\sqrt{d\delta^2_{\rm QWP} + 4d\theta^2_{\rm QWP}}$	NA
$\frac{\pi}{4}$	0	QWP @ $d\theta_{\rm QWP}$	$\frac{\pi}{4}$	0	$-d\theta_{\rm QWP}$	$d\theta_{\rm QWP}(1 + d\delta_{\rm QWP})$
$\frac{\pi}{4}$	0	ZWP @ $\theta_{\rm ZWP}$	$\frac{\pi}{4}$	0	$-\frac{d\delta_{\rm ZWP}}{2}\sin(2\theta_{\rm ZWP})$	$\frac{1}{8}d\delta^2_{\rm ZWP}\sin(4\theta_{\rm ZWP})$
$\frac{\pi}{2}$	0	HWP @ $\theta_{\rm HWP}$	0	NA	$\frac{d\delta_{\rm HWP}}{2}$	NA
$\frac{\pi}{2}$	0	QWP @ $\theta_{\rm QWP}$	$\frac{\pi}{4}$	$\theta_{\rm QWP} + \frac{\pi}{4}$	$-\frac{d\delta_{\rm QWP}}{2}$	0
$\frac{\pi}{2}$	0	ZWP @ $\theta_{\rm ZWP}$	$\frac{\pi}{2}$	NA	$-\frac{d\delta_{\rm ZWP}}{2}$	NA

Table C.1: Polarization imperfections generated after propagation through different wave plates. Here HWP denotes a half-wave plate, QWP a quarter-wave plate, and ZWP a zero-order wave plate (or, more generally, a weakly birefringent element). All orientation angles θ are defined with respect to a common laboratory axis. The retardance error $d\delta$ is defined as the deviation from the ideal retardance of the HWP, QWP, or ZWP. The birefringent axes of the HWP, QWP, and ZWP are denoted by $\theta_{\rm HWP}$, $\theta_{\rm QWP}$, and $\theta_{\rm ZWP}$, respectively.

C.2 Polarization imperfections from different wave plates

In this section, we derive the outgoing polarization imperfections after propagation through different wave plates for various incident polarizations. All results are expanded to second order in small imperfections.

C.2.1 Linear input into HWP

Wave plate parameters: $\theta = \theta_{\rm HWP}$, $\delta = \pi + \kappa\, d\delta_{\rm HWP}$.

Incident polarization:

$$\Theta = \frac{\pi}{4} + \kappa\, d\Theta_{\rm in}, \tag{C.4}$$

$$\theta = \theta_0 + \kappa\, d\theta_{\rm in}. \tag{C.5}$$

Outgoing polarization:

$$\Theta = \frac{\pi}{4} + \kappa\left[\frac{1}{2}\sin\!\big(2\theta_0 - 2\theta_{\rm HWP}\big)\, d\delta_{\rm HWP} - d\Theta_{\rm in}\right] + \kappa^2 \cos\!\big(2\theta_0 - 2\theta_{\rm HWP}\big)\, d\delta_{\rm HWP}\, d\theta_{\rm in}, \tag{C.6}$$

$$\theta = -\theta_0 + 2\theta_{\rm HWP} - \kappa\, d\theta_{\rm in} + \kappa^2 \frac{d\delta_{\rm HWP}}{8}\Big[d\delta_{\rm HWP}\sin\!\big(4(\theta_0 - \theta_{\rm HWP})\big) - 8\, d\Theta_{\rm in}\cos\!\big(2(\theta_0 - \theta_{\rm HWP})\big)\Big]. \tag{C.7}$$

C.2.2 Linear input into QWP, circular output

Wave plate parameters: $\theta = \tilde{\mathcal{Q}}\frac{\pi}{4} + \kappa\, d\theta_{\rm QWP}$, $\delta = \frac{\pi}{2} + \kappa\, d\delta_{\rm QWP}$, where $\mathcal{Q}$ is the polarimetry switch.

Incident polarization:

$$\Theta = \frac{\pi}{4} \pm \kappa\, d\Theta_{\rm in}, \tag{C.8}$$

$$\theta = \frac{\pi}{4} \mp \frac{\pi}{4} + \kappa\, d\theta_{\rm in}. \tag{C.9}$$

Outgoing polarization:

$$\Psi = \frac{\pi}{4} \pm \kappa\left(\tilde{\mathcal{Q}}\, d\Theta_{\rm in} + \frac{d\delta_{\rm QWP}}{2}\right) \pm \kappa^2\, 2(d\theta_{\rm in} - d\theta_{\rm QWP})\, d\theta_{\rm QWP}, \tag{C.10}$$

$$\psi = \pm\frac{\pi}{4}\tilde{\mathcal{Q}} + \kappa\,\tilde{\mathcal{Q}}(d\theta_{\rm in} - d\theta_{\rm QWP}) - \kappa^2 d\theta_{\rm QWP}\big(\tilde{\mathcal{Q}}\, d\delta_{\rm QWP} + 2d\Theta_{\rm in}\big). \tag{C.11}$$

C.2.3 Linear input into QWP, linear output

Wave plate parameters: $\theta = \kappa\, d\theta_{\mathrm{QWP}}$, $\delta = \frac{\pi}{2} + \kappa\, d\delta_{\mathrm{QWP}}$.

Incident polarization:

$$\Theta = \frac{\pi}{4} \pm \kappa\, d\Theta_{\mathrm{in}}, \tag{C.12}$$

$$\theta = \frac{\pi}{4} \mp \frac{\pi}{4} + \kappa\, d\theta_{\mathrm{in}}. \tag{C.13}$$

Outgoing polarization:

$$\Theta = \frac{\pi}{4} \pm \kappa(d\theta_{\mathrm{QWP}} - d\theta_{\mathrm{in}}) \mp \kappa^2 d\delta_{\mathrm{QWP}}\, d\Theta_{\mathrm{in}}, \tag{C.14}$$

$$\theta = \frac{\pi}{4} \mp \frac{\pi}{4} + \kappa(d\theta_{\mathrm{QWP}} + d\Theta_{\mathrm{in}}) + \kappa^2 d\delta_{\mathrm{QWP}}(d\theta_{\mathrm{QWP}} - d\theta_{\mathrm{in}}). \tag{C.15}$$

C.2.4 Linear input into ZWP

Wave plate parameters: $\theta = \theta_{\mathrm{ZWP}}$, $\delta = \kappa\, d\delta_{\mathrm{ZWP}}$.

Incident polarization:

$$\Theta = \frac{\pi}{4} \pm \kappa\, d\Theta_{\mathrm{in}}, \tag{C.16}$$

$$\theta = \frac{\pi}{4} \mp \frac{\pi}{4} + \kappa\, d\theta_{\mathrm{in}}. \tag{C.17}$$

Outgoing polarization:

$$\Theta = \frac{\pi}{4} \pm \kappa\left(d\Theta_{\mathrm{in}} + \frac{d\delta_{\mathrm{ZWP}}}{2}\sin\big(2\theta_{\mathrm{ZWP}}\big)\right) \mp \kappa^2 d\delta_{\mathrm{ZWP}}\, d\theta_{\mathrm{in}} \cos\big(2\theta_{\mathrm{ZWP}}\big), \tag{C.18}$$

$$\theta = \frac{\pi}{4} \mp \frac{\pi}{4} + \kappa\, d\theta_{\mathrm{in}} + \kappa^2 d\delta_{\mathrm{ZWP}}\left[\frac{d\delta_{\mathrm{ZWP}}}{8}\sin\big(4\theta_{\mathrm{ZWP}}\big) + d\Theta_{\mathrm{in}}\cos\big(2\theta_{\mathrm{ZWP}}\big)\right]. \tag{C.19}$$

C.2.5 Circular input into HWP, circular output

Wave plate parameters: $\theta = \theta_{\mathrm{HWP}}$, $\delta = \pi + \kappa\, d\delta_{\mathrm{HWP}}$.

Incident polarization:

$$\Psi = \frac{\pi}{4} \pm \kappa\, d\Psi_{\mathrm{in}}, \tag{C.20}$$

$$\psi = \pm\frac{\pi}{4} + \kappa\, d\psi_{\mathrm{in}}. \tag{C.21}$$

Outgoing polarization:

$$\Psi = \frac{\pi}{4} \pm \kappa\left[d\Psi_{\mathrm{in}} \cos\!\big(4\theta_{\mathrm{HWP}}\big) - d\psi_{\mathrm{in}} \sin\!\big(4\theta_{\mathrm{HWP}}\big) - \frac{d\delta_{\mathrm{HWP}}}{2} \sin\!\big(2\theta_{\mathrm{HWP}}\big)\right], \tag{C.22}$$

$$\psi = \pm\frac{\pi}{4} + \kappa\left[d\Psi_{\mathrm{in}} \sin\!\big(4\theta_{\mathrm{HWP}}\big) + d\psi_{\mathrm{in}} \cos\!\big(4\theta_{\mathrm{HWP}}\big) + \frac{d\delta_{\mathrm{HWP}}}{2} \cos\!\big(2\theta_{\mathrm{HWP}}\big)\right]. \tag{C.23}$$

C.2.6 Circular input into ZWP

Wave plate parameters: $\theta = \theta_{\mathrm{ZWP}}$, $\delta = \kappa\, d\delta_{\mathrm{ZWP}}$.

Incident polarization:

$$\Psi = \frac{\pi}{4} \pm \kappa\, d\Psi_{\mathrm{in}}, \tag{C.24}$$

$$\psi = \pm\frac{\pi}{4} + \kappa\, d\psi_{\mathrm{in}}. \tag{C.25}$$

Outgoing polarization:

$$\Psi = \frac{\pi}{4} \pm \kappa\left(d\Psi_{\mathrm{in}} + \frac{d\delta_{\mathrm{ZWP}}}{2} \sin\!\big(2\theta_{\mathrm{ZWP}}\big)\right), \tag{C.26}$$

$$\psi = \pm\frac{\pi}{4} + \kappa\left(d\psi_{\mathrm{in}} + \frac{d\delta_{\mathrm{ZWP}}}{2} \cos\!\big(2\theta_{\mathrm{ZWP}}\big)\right). \tag{C.27}$$

C.3 Combined imperfections, readout polarimetry mode

In this section, we derive the total polarization imperfections after the light propagates through all optical elements in the readout polarimetry mode.

C.3.1 Preparation side

Incident polarization:

$$\Theta = \frac{\pi}{4} + \kappa\, d\Theta_{\mathrm{in}}, \tag{C.28}$$

$$\theta = \kappa\, d\theta_{\mathrm{in}}. \tag{C.29}$$

Outgoing polarization:

$$\Psi = \frac{\pi}{4} + \kappa\left(\tilde{\mathcal{Q}}\, d\Theta_{\mathrm{in}} + \frac{d\delta_{\mathrm{QWP}}}{2} + \tilde{\mathcal{Q}}\,\frac{d\delta_{\mathrm{ZWP}}}{2}\sin\!\big(2\theta_{\mathrm{ZWP}}\big)\right) + \kappa^2\, 2(d\theta_{\mathrm{in}} - d\theta_{\mathrm{QWP}})\, d\theta_{\mathrm{QWP}}, \tag{C.30}$$

$$\psi = \frac{\pi}{4}\tilde{\mathcal{Q}} + \kappa\left(\tilde{\mathcal{Q}}\, d\theta_{\mathrm{in}} - \tilde{\mathcal{Q}}\, d\theta_{\mathrm{QWP}} + \frac{d\delta_{\mathrm{ZWP}}}{2}\cos\!\big(2\theta_{\mathrm{ZWP}}\big)\right) - \kappa^2 d\theta_{\mathrm{QWP}}\big(\tilde{\mathcal{Q}}\, d\delta_{\mathrm{QWP}} + 2d\Theta_{\mathrm{in}}\big). \tag{C.31}$$

C.3.2 Readout side

Incident polarization:

$$\Theta = \frac{\pi}{4}, \tag{C.32}$$

$$\theta = \frac{\pi}{4} \mp \frac{\pi}{4}. \tag{C.33}$$

Outgoing polarization:

$$\Theta = \frac{\pi}{4} \pm \frac{\kappa}{2}\Big[-d\delta_{\mathrm{HWP}}\sin\big(2\theta_{\mathrm{HWP}}\big) - d\delta_{\mathrm{ZWP}}\sin\big(4\theta_{\mathrm{HWP}} - 2\theta_{\mathrm{ZWP}}\big) - 2d\theta_{\mathrm{QWP}}\Big] \\ \pm \kappa^2 d\theta_{\mathrm{QWP}}\Big[d\delta_{\mathrm{HWP}}\cos\big(2\theta_{\mathrm{HWP}}\big) + d\delta_{\mathrm{ZWP}}\cos\big(4\theta_{\mathrm{HWP}} - 2\theta_{\mathrm{ZWP}}\big)\Big], \tag{C.34}$$

$$\theta = -\frac{\pi}{4} \pm \frac{\pi}{4} - \kappa\, d\theta_{\mathrm{QWP}} \\ + \frac{\kappa^2}{8}\Big[-8d\theta_{\mathrm{QWP}}\big(d\delta_{\mathrm{HWP}}\cos\big(2\theta_{\mathrm{HWP}}\big) + d\delta_{\mathrm{QWP}}\big) - 2d\delta_{\mathrm{ZWP}}\cos\big(4\theta_{\mathrm{HWP}} - 2\theta_{\mathrm{ZWP}}\big) \\ \times \big(2d\delta_{\mathrm{HWP}}\sin\big(2\theta_{\mathrm{HWP}}\big) + d\delta_{\mathrm{ZWP}}\sin\big(4\theta_{\mathrm{HWP}} - 2\theta_{\mathrm{ZWP}}\big)\big) + 4d\theta_{\mathrm{QWP}} - d\delta_{\mathrm{HWP}}^2\sin\big(4\theta_{\mathrm{HWP}}\big)\Big]. \tag{C.35}$$

C.4 Combined imperfections, prep polarimetry mode

In this section, we derive the total polarization imperfections after the light propagates through all optical elements in the preparation polarimetry mode.

C.4.1 Preparation side

Incident polarization:

$$\Theta = \frac{\pi}{4}, \tag{C.36}$$

$$\theta = 0. \tag{C.37}$$

Outgoing polarization:

$$\Theta = \frac{\pi}{4} + \kappa\left(\frac{d\delta_{\mathrm{ZWP}}}{2}\sin\big(2\theta_{\mathrm{ZWP}}\big) + d\theta_{\mathrm{QWP}}\right) - \kappa^2 d\delta_{\mathrm{ZWP}}\, d\theta_{\mathrm{QWP}}\cos\big(2\theta_{\mathrm{ZWP}}\big), \tag{C.38}$$

$$\theta = \kappa\, d\theta_{\mathrm{QWP}} + \kappa^2\left[d\theta_{\mathrm{QWP}}\big(d\delta_{\mathrm{QWP}} + d\delta_{\mathrm{ZWP}}\cos\big(2\theta_{\mathrm{ZWP}}\big)\big) + \frac{d\delta_{\mathrm{ZWP}}^2}{8}\sin\big(4\theta_{\mathrm{ZWP}}\big)\right]. \tag{C.39}$$

C.4.2 Readout side

Incident polarization:

$$\Theta = \frac{\pi}{4}, \tag{C.40}$$

$$\theta = \frac{\pi}{4} \mp \frac{\pi}{4}. \tag{C.41}$$

Outgoing polarization:

$$\begin{aligned}\Psi = \frac{\pi}{4} \pm \frac{\kappa}{2}\Big[\tilde{\mathcal{Q}}\Big(-d\delta_{\mathrm{HWP}}\sin\big(2\theta_{\mathrm{HWP}}\big) - d\delta_{\mathrm{ZWP}}\sin\big(2\theta_{\mathrm{ZWP}}\big) + 2\tilde{\mathcal{Q}}\,d\theta_{\mathrm{QWP}}\sin\big(4\theta_{\mathrm{HWP}}\big)\Big) - d\delta_{\mathrm{QWP}}\cos\big(4\theta_{\mathrm{HWP}}\big)\Big] \\ \pm \kappa^2 d\theta_{\mathrm{QWP}}\Big[d\delta_{\mathrm{QWP}}\sin\big(4\theta_{\mathrm{HWP}}\big) + 2d\theta_{\mathrm{QWP}}\cos\big(4\theta_{\mathrm{HWP}}\big)\Big],\end{aligned} \tag{C.42}$$

$$\begin{aligned}\psi = \mp\tilde{\mathcal{Q}}\frac{\pi}{4} + \frac{\kappa}{2}\Big[d\delta_{\mathrm{HWP}}\cos\big(2\theta_{\mathrm{HWP}}\big) - \tilde{\mathcal{Q}}\,d\delta_{\mathrm{QWP}}\sin\big(4\theta_{\mathrm{HWP}}\big) + d\delta_{\mathrm{ZWP}}\cos\big(2\theta_{\mathrm{ZWP}}\big) \\ - 2\tilde{\mathcal{Q}}\,d\theta_{\mathrm{QWP}}\cos\big(4\theta_{\mathrm{HWP}}\big)\Big] + \kappa^2\tilde{\mathcal{Q}}\,d\theta_{\mathrm{QWP}}\Big[2d\theta_{\mathrm{QWP}}\sin\big(4\theta_{\mathrm{HWP}}\big) - d\delta_{\mathrm{QWP}}\cos\big(4\theta_{\mathrm{HWP}}\big)\Big].\end{aligned} \tag{C.43}$$

C.5 Measured asymmetry in the state basis

In this section, we derive the measured asymmetry in both the readout polarimetry mode and the preparation polarimetry mode, using the total imperfections obtained in Secs. C.3 and C.4, respectively. In both cases, we include a small spin-precession phase ϕ and expand to second order in small imperfections.

C.5.1 Readout polarimetry

In this mode, the preparation light is circularly polarized and the readout light is linearly polarized, as parameterized below. A small precession phase ϕ is included.

$$\psi_{\mathrm{prep}} = \tilde{\mathcal{Q}}\frac{\pi}{4} + \kappa\, d\psi_{\mathrm{prep}}, \tag{C.44}$$

$$\Psi_{\mathrm{prep}} = \frac{\pi}{4} + \kappa\, d\Psi_{\mathrm{prep}}, \tag{C.45}$$

$$\theta_X = 2\theta_{\mathrm{HWP}} + \tilde{\mathcal{R}}\frac{\pi}{4}, \tag{C.46}$$

$$\theta_Y = \frac{\pi}{2} + 2\theta_{\mathrm{HWP}} + \tilde{\mathcal{R}}\frac{\pi}{4}, \tag{C.47}$$

$$\Theta_X = \frac{\pi}{4} + \kappa\, d\Theta_{\mathrm{read}}, \tag{C.48}$$

$$\Theta_Y = \frac{\pi}{4} - \kappa\, d\Theta_{\mathrm{read}}, \tag{C.49}$$

$$\phi = \kappa\, d\phi. \tag{C.50}$$

The measured asymmetry is

$$\begin{aligned} A = \kappa\, 2\Big[& d\Theta_{\mathrm{read}}\, \tilde{\mathcal{Q}} - \tilde{\mathcal{P}}\, d\Psi_{\mathrm{prep}} \cos\left(\frac{\pi\tilde{\mathcal{R}}}{2} + 4\theta_{\mathrm{HWP}}\right) + \tilde{\mathcal{P}}\tilde{\mathcal{Q}}\, d\psi_{\mathrm{prep}} \sin\left(\frac{\pi\tilde{\mathcal{R}}}{2} + 4\theta_{\mathrm{HWP}}\right)\Big] \\ - \kappa^2\, 4d\phi\Big[& \tilde{\mathcal{P}}\, d\Psi_{\mathrm{prep}} \sin\left(\frac{\pi\tilde{\mathcal{R}}}{2} + 4\theta_{\mathrm{HWP}}\right) + \tilde{\mathcal{P}}\tilde{\mathcal{Q}}\, d\psi_{\mathrm{prep}} \cos\left(\frac{\pi\tilde{\mathcal{R}}}{2} + 4\theta_{\mathrm{HWP}}\right)\Big]. \end{aligned} \tag{C.51}$$

Here, the term $d\Theta_{\mathrm{read}}\,\tilde{\mathcal{Q}}$ is the primary quantity of interest. The remaining terms arise from linear–linear polarization couplings, which are modulated by the wave-plate angles and therefore depend on the switches $\tilde{\mathcal{Q}}$ and $\tilde{\mathcal{R}}$.

C.5.2 Preparation polarimetry

In this mode, the preparation light is linearly polarized and the readout light is circularly polarized, as parameterized below. A small precession phase ϕ is included.

$$\theta_{\text{prep}} = \kappa\, d\theta_{\text{prep}}, \tag{C.52}$$

$$\Theta_{\text{prep}} = \frac{\pi}{4} + \kappa\, d\Theta_{\text{prep}}, \tag{C.53}$$

$$\psi_X = \tilde{\mathcal{Q}}\frac{\pi}{4} + \kappa\, d\psi_{\text{read}}, \tag{C.54}$$

$$\psi_Y = -\tilde{\mathcal{Q}}\frac{\pi}{4} + \kappa\, d\psi_{\text{read}}, \tag{C.55}$$

$$\Psi_X = \frac{\pi}{4} + \kappa\, d\Psi_{\text{read}}, \tag{C.56}$$

$$\Psi_Y = \frac{\pi}{4} - \kappa\, d\Psi_{\text{read}}, \tag{C.57}$$

$$\phi = \kappa\, d\phi. \tag{C.58}$$

The measured asymmetry is

$$A = \kappa\left(2\tilde{\mathcal{Q}}\, d\Theta_{\text{prep}} - 2\tilde{\mathcal{P}}\, d\Psi_{\text{read}}\right) + 4\kappa^2\tilde{\mathcal{P}}\tilde{\mathcal{Q}}\, d\psi_{\text{read}}\left(d\theta_{\text{prep}} + d\phi\right). \tag{C.59}$$

This expression appears relatively simple because the detailed structure of the imperfections has not yet been substituted; doing so would introduce additional dependence on the switches $\tilde{\mathcal{Q}}$ and $\tilde{\mathcal{R}}$. The first-order terms arise from circular–circular coupling, which depends on $\tilde{\mathcal{Q}}$, and vertical–vertical coupling, which depends on $\tilde{\mathcal{P}}$. The second-order term corresponds to diagonal–diagonal coupling and is suppressed by one additional order, since both the preparation and readout fields contain only zeroth-order vertical or circular polarization components.

C.6 Polarimetry mode parity channels

In this section, we use the detailed forms of the imperfections derived in Secs. C.3 and C.4 to determine the contributions to the various parity channels. These results form

the basis for interpreting the measured asymmetries.

C.6.1 Readout polarimetry

Writing out the full expression for the asymmetry is prohibitively cumbersome. Instead, we list the contributions to each parity channel explicitly:

$$
\begin{aligned}
Q : \kappa\Big[&2d\Theta_{\mathrm{read}} - d\delta_{\mathrm{HWP,read}}\sin\big(2\theta_{\mathrm{HWP,read}}\big) - d\delta_{\mathrm{ZWP,read}}\sin\big(4\theta_{\mathrm{HWP,read}} - 2\theta_{\mathrm{ZWP,read}}\big) - 2d\theta_{\mathrm{QWP,read}}\Big] \\
&+ \kappa^2\, 2d\theta_{\mathrm{QWP,read}}\Big[d\delta_{\mathrm{HWP,read}}\cos\big(2\theta_{\mathrm{HWP,read}}\big) + d\delta_{\mathrm{ZWP,read}}\cos\big(4\theta_{\mathrm{HWP,read}} - 2\theta_{\mathrm{ZWP,read}}\big)\Big], \\
QPR : \kappa\, &d\delta_{\mathrm{ZWP,prep}}\cos\big(4\theta_{\mathrm{HWP,read}} - 2\theta_{\mathrm{ZWP,prep}}\big) \\
&+ \kappa^2\, 2d\delta_{\mathrm{ZWP,prep}}(d\phi + d\theta_{\mathrm{QWP,read}})\sin\big(4\theta_{\mathrm{HWP,read}} - 2\theta_{\mathrm{ZWP,prep}}\big), \\
PR : -\kappa\Big[&d\delta_{\mathrm{QWP,prep}}\sin\big(4\theta_{\mathrm{HWP,read}}\big) - 2d\theta_{\mathrm{QWP,prep}}\cos\big(4\theta_{\mathrm{HWP,read}}\big)\Big] \\
&+ \kappa^2(d\phi - d\theta_{\mathrm{QWP,prep}} + d\theta_{\mathrm{QWP,read}})\Big[2d\delta_{\mathrm{QWP,prep}}\cos\big(4\theta_{\mathrm{HWP,read}}\big) - 4d\theta_{\mathrm{QWP,prep}}\sin\big(4\theta_{\mathrm{HWP,read}}\big)\Big].
\end{aligned}
$$

Interpretation.

Q : First-order terms correspond to the readout ellipticity and its contributions from wave-plate imperfections. Second-order terms arise from linear rotation induced by $d\theta_{\mathrm{QWP,read}}$.

QPR : Linear–linear coupling from imperfectly circular preparation due to $d\delta_{\mathrm{ZWP,prep}}$.

PR : Linear–linear coupling from imperfectly circular preparation due to $d\delta_{\mathrm{QWP,prep}}$.

QR : Imperfect R via Q, potentially arising from angle-dependent $d\delta_{\mathrm{HWP,read}}$.

QP : Imperfect R via QPR; angle-dependent $d\delta_{\mathrm{HWP,read}}$ does not contribute here.

P : Imperfect R via PR; again, angle-dependent $d\delta_{\mathrm{HWP,read}}$ does not contribute. Possible additional contribution from imperfect half-wave-plate rotation.

$QBPR$: Leakage from QPR due to A^B.

BPR : Leakage from PR due to A^B.

θPR : Leakage from PR due to A^θ.

C.6.2 Preparation polarimetry

$$
\begin{aligned}
Q :\ & \kappa\Big[2d\Theta_{\mathrm{prep}} + 2d\theta_{\mathrm{QWP,prep}} + d\delta_{\mathrm{ZWP,prep}}\sin\big(2\theta_{\mathrm{ZWP,prep}}\big)\Big] \\
& - \kappa^2\, 4d\delta_{\mathrm{ZWP,prep}} d\theta_{\mathrm{QWP,prep}}\cos\big(2\theta_{\mathrm{ZWP,prep}}\big), \\
QP :\ & \kappa\, d\delta_{\mathrm{ZWP,read}}\sin\big(2\theta_{\mathrm{ZWP,read}}\big) \\
& + \kappa^2\, 2(d\phi + d\theta_{\mathrm{QWP,prep}}) d\delta_{\mathrm{ZWP,read}}\cos\big(2\theta_{\mathrm{ZWP,read}}\big), \\
QP\ \&\ QPR :\ & \kappa\, d\delta_{\mathrm{HWP,read}}\sin\big(2\theta_{\mathrm{HWP,read}}\big) \\
& + \kappa^2\, 2(d\phi + d\theta_{\mathrm{QWP,prep}}) d\delta_{\mathrm{HWP,read}}\cos\big(2\theta_{\mathrm{HWP,read}}\big), \\
PR :\ & -\kappa\Big[d\delta_{\mathrm{QWP,read}}\sin\big(4\theta_{\mathrm{HWP,read}}\big) + 2d\theta_{\mathrm{QWP,read}}\cos\big(4\theta_{\mathrm{HWP,read}}\big)\Big] \\
& - \kappa^2(d\phi + d\theta_{\mathrm{QWP,prep}} + d\theta_{\mathrm{QWP,read}})\Big[2d\delta_{\mathrm{QWP,read}}\cos\big(4\theta_{\mathrm{HWP,read}}\big) - 4d\theta_{\mathrm{QWP,read}}\sin\big(4\theta_{\mathrm{HWP,read}}\big)\Big].
\end{aligned}
$$

Interpretation.

Q : Preparation ellipticity and its contributions from the preparation QWP and ZWP.

QP & QPR : Linear–linear coupling arising from $d\delta_{\mathrm{HWP,read}}$ and $d\delta_{\mathrm{ZWP,read}}$.

This channel may also contain leakage from imperfect R, e.g., due to angle-dependent $d\delta_{\mathrm{HWP,read}}$.

PR : Linear–linear coupling from readout imperfections ($d\delta_{\mathrm{QWP,read}}, d\theta_{\mathrm{QWP,read}}$).

P : Leakage from the QP channel due to imperfect PR, not related to $d\delta_{\mathrm{HWP,read}}$.

QBP : Leakage from the QP channel via A^B.

BPR : Leakage from the PR channel via A^B.

θPR : Leakage from the PR channel via A^θ.

APPENDIX D

SYSTEMATIC PHASE FROM BRIGHT-STATE ADMIXTURE

A lot of potential systematic errors in ACME come from our state preparation scheme where the molecules are optically pumped into the dark state. Imperfections in this optical pumping can lead to a small admixture of the bright state, which can in turn lead to systematic shifts in the measured precession phase. In this appendix, we derive the relationship between bright-state admixture and the resulting systematic phase shift.

We start with the standard parameterization of the molecular dark and bright states:

$$|D\rangle = -e^{-i\theta}\cos\Theta\ |M=+1\rangle + e^{i\theta}\sin\Theta\ |M=-1\rangle\,, \tag{D.1}$$

$$|B\rangle = e^{-i\theta}\sin\Theta\ |M=+1\rangle + e^{i\theta}\cos\Theta\ |M=-1\rangle\,, \tag{D.2}$$

where the nominal value $\Theta = \pi/4$ corresponds to purely linearly polarized molecules.

We now consider two scenarios:

1. the molecules acquire a small precession phase $\kappa\, d\phi$, where κ is a bookkeeping parameter for perturbative expansions. For completeness, we also allow the molecules to acquire a small ellipticity $\kappa\, d\Theta$;

2. the molecules acquire a small admixture of the nominal bright state, $\kappa\, d\Lambda$.

We are interested in the mapping between these two descriptions. The second description is often more convenient for solving Hamiltonian problems, while the first is the physically relevant parameterization for extracting the electron EDM.

To obtain this relationship, we equate the two descriptions perturbatively:

$$\begin{aligned}
|D\rangle + \kappa d\Lambda\, |B\rangle &= -e^{-i\theta - i\kappa d\phi} \cos(\Theta + \kappa d\Theta)\, |M=+1\rangle + e^{i\theta + i\kappa d\phi} \sin(\Theta - \kappa d\Theta)\, |M=-1\rangle + \mathcal{O}(\kappa^2) \\
&= \kappa d\Theta\, |B\rangle + \kappa i d\phi \left(e^{-i\theta} \cos\Theta\, |M=+1\rangle + e^{i\theta} \sin\Theta\, |M=-1\rangle \right) + \mathcal{O}(\kappa^2) \\
&= \kappa d\Theta\, |B\rangle + \kappa i d\phi \left(e^{-i\theta} \sin\Theta\, |M=+1\rangle + e^{i\theta} \cos\Theta\, |M=-1\rangle \right) + \mathcal{O}(\kappa^2) \\
&= \kappa(d\Theta + i d\phi)\, |B\rangle + \mathcal{O}(\kappa^2),
\end{aligned}$$

where we have used the nominal value $\Theta = \pi/4$, for which $\cos\Theta = \sin\Theta$.

This establishes the important correspondence

$$d\phi = \mathrm{Im}(d\Lambda), \tag{D.3}$$

$$d\Theta = \mathrm{Re}(d\Lambda). \tag{D.4}$$

Notably, this relationship holds independently of the initial angle θ, and therefore applies for all wave-plate angles and for either parity of the excited I state.

APPENDIX E

AC STARK SHIFT SYSTEMATICS WITHOUT PERTURBATION

In Sec. 4.5.2, we derived the detuning dependence of the ellipticity-induced AC Stark shift systematic using perturbation theory. Here, we present a more general and rigorous derivation based on symmetry arguments, without relying on perturbation theory.

Before introducing the symmetry argument, it is helpful to clarify the statement we wish to prove. We begin with the Hamiltonian given in Eq. (4.27),

$$\tilde{H}(\Delta) = \begin{pmatrix} \Delta - i\gamma/2 & \frac{1}{2}\Omega_{\mathrm{r}} & 0 \\ \frac{1}{2}\Omega_{\mathrm{r}} & 0 & i\dot{\chi} \\ 0 & -i\dot{\chi}^* & 0 \end{pmatrix} \begin{matrix} |I\rangle \\ |B\rangle, \\ |D\rangle \end{matrix} \tag{E.1}$$

and study the properties of its solutions. The key feature of interest is the behavior of the solution under a change of the detuning sign, $\Delta \to -\Delta$. To emphasize this dependence, we explicitly write the Hamiltonian as $\tilde{H}(\Delta)$ rather than simply $\tilde{H}$.

We write the corresponding solution as

$$|\psi(t,\Delta)\rangle = c_D(t,\Delta)|D\rangle + c_B(t,\Delta)|B\rangle + c_I(t,\Delta)|I\rangle. \tag{E.2}$$

In general, reversing the detuning in the Hamiltonian leads to different solutions, so the coefficients $c_\alpha(t,\pm\Delta)$, with $\alpha = D, B, I$, are not equal.

As shown in Appendix D, the AC Stark shift phase is given by

$$\delta\Phi(\Delta) = \mathrm{Im}\, c_B(t,\Delta). \tag{E.3}$$

This quantity represents the total AC Stark shift phase. To isolate the AC Stark shift

systematic, we take the part that is odd in detuning. Explicitly,

$$\delta\Phi(\Delta) = \frac{\delta\Phi(\Delta) + \delta\Phi(-\Delta)}{2} + \frac{\delta\Phi(\Delta) - \delta\Phi(-\Delta)}{2}, \tag{E.4}$$

where the first term is even in Δ and the second term is odd. Only the detuning-odd term contributes to the AC Stark shift systematic.

Our goal is to show that when the ellipticity changes sign, $d\Theta \mapsto -d\Theta$, the systematic phase also changes sign:

$$\frac{\delta\Phi(\Delta) - \delta\Phi(-\Delta)}{2} = \frac{\operatorname{Im} c_B(t,\Delta) - \operatorname{Im} c_B(t,-\Delta)}{2} \;\mapsto\; -\frac{\operatorname{Im} c_B(t,\Delta) - \operatorname{Im} c_B(t,-\Delta)}{2}. \tag{E.5}$$

This shows that the Δ-odd AC Stark shift phase is also odd in ellipticity, which is the desired result. Having clarified our goal, we are now ready to present the symmetry argument. To do so, we first introduce an elementary but useful lemma.

Lemma: Suppose we have two systems of differential equations

$$\partial_t \vec{\psi} = \mathbf{J}(t)\vec{\psi}, \qquad \partial_t \vec{\phi} = \mathbf{K}(t)\vec{\phi}, \tag{E.6}$$

subject to the initial condition

$$\vec{\psi}(0) = \vec{\phi}(0) = \vec{c}_0. \tag{E.7}$$

Suppose in addition that there exists a time-independent invertible matrix $\mathbf{R}$ such that

$$\mathbf{R}^{-1}\mathbf{K}(t)\mathbf{R} = \mathbf{J}(t), \tag{E.8}$$

and that the initial condition is invariant under $\mathbf{R}$,

$$\mathbf{R}\vec{c}_0 = \vec{c}_0. \tag{E.9}$$

Then, by substituting $\vec{\phi} = \mathbf{R}\vec{\psi}$, we see that $\vec{\phi}$ solves $\partial_t\vec{\phi} = \mathbf{K}(t)\vec{\phi}$.

This lemma is proved in one line and requires little further discussion; it is essentially a generalized (not necessarily unitary) change of basis.

To apply the lemma, we rewrite the Schrödinger equation in the form of Eq. (E.6). From Eq. (E.5), we need to understand how $\mathrm{Im}\, c_B(t, \pm\Delta)$ transforms for both signs of detuning. Because both $+\Delta$ and $-\Delta$ solutions enter, our enlarged system must include both $\tilde{H}(\Delta)$ and $\tilde{H}(-\Delta)$. Each Hamiltonian has three complex-valued amplitudes (c_I, c_B, c_D), which we decompose into real and imaginary parts, giving $2 \times 3 \times 2 = 12$ real variables. This formulation allows more generalized transformations that mix solutions associated with $\tilde{H}(\Delta)$ and $\tilde{H}(-\Delta)$.

We order these variables as

$$\vec{\psi}^T = (I_{r+}, I_{i+}, I_{r-}, I_{i-}, B_{r+}, B_{i+}, B_{r-}, B_{i-}, D_{r+}, D_{i+}, D_{r-}, D_{i-})\,. \tag{E.10}$$

where we write $c_I(t)$, $c_B(t)$, and $c_D(t)$ as I, B, and D for brevity; the subscripts r and i denote real and imaginary parts; and the subscripts $+$ and $-$ denote solutions corresponding to $\tilde{H}(\Delta)$ and $\tilde{H}(-\Delta)$, respectively.

In this basis, it can be shown using the Hamiltonian $\tilde{H}(\pm\Delta)$ that the corresponding

time-evolution matrix $\mathbf{J}$ is

$$\mathbf{J} = \begin{pmatrix}
-\frac{\gamma}{2} & \Delta & 0 & 0 & 0 & \frac{\Omega}{2} & 0 & 0 & 0 & 0 & 0 & 0 \\
-\Delta & -\frac{\gamma}{2} & 0 & 0 & -\frac{\Omega}{2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & -\frac{\gamma}{2} & -\Delta & 0 & 0 & 0 & \frac{\Omega}{2} & 0 & 0 & 0 & 0 \\
0 & 0 & \Delta & -\frac{\gamma}{2} & 0 & 0 & -\frac{\Omega}{2} & 0 & 0 & 0 & 0 & 0 \\
0 & \frac{\Omega}{2} & 0 & 0 & 0 & 0 & 0 & 0 & \dot{\Theta} & \dot{\theta} & 0 & 0 \\
-\frac{\Omega}{2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\dot{\theta} & \dot{\Theta} & 0 & 0 \\
0 & 0 & 0 & \frac{\Omega}{2} & 0 & 0 & 0 & 0 & 0 & 0 & \dot{\Theta} & \dot{\theta} \\
0 & 0 & -\frac{\Omega}{2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\dot{\theta} & \dot{\Theta} \\
0 & 0 & 0 & 0 & -\dot{\Theta} & \dot{\theta} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & -\dot{\theta} & -\dot{\Theta} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & -\dot{\Theta} & \dot{\theta} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & -\dot{\theta} & -\dot{\Theta} & 0 & 0 & 0 & 0
\end{pmatrix}. \tag{E.11}$$

To reduce redundancy, we do not write out $\mathbf{K}$ explicitly; it is obtained from $\mathbf{J}$ by the

replacement $\dot{\Theta} \mapsto -\dot{\Theta}$. We find that with the transformation

$$\mathbf{R} = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \end{pmatrix}, \tag{E.12}$$

we satisfy $\mathbf{R}^{-1}\mathbf{K}(t)\mathbf{R} = \mathbf{J}(t)$. In addition, the initial condition $I, B, D = 0, 0, 1$ is invariant under this transformation. This means that $\mathbf{R}\vec{\psi}$ solves the differential equation for $-\dot{\Theta}$ while $\vec{\psi}$ solves the differential equation for $+\dot{\Theta}$. To obtain the systematic shift, we return to Eq. (E.5). In our present notation, the systematic contribution is

$$\frac{\operatorname{Im} c_B(t, \Delta) - \operatorname{Im} c_B(t, -\Delta)}{2} = \frac{B_{i+} - B_{i-}}{2}. \tag{E.13}$$

From the transformation $\mathbf{R}$ we see that $B_{i+} \mapsto B_{i-}$ while $B_{i-} \mapsto B_{i+}$. This obviously reverses the systematic shift. $\square$

REFERENCES

[1] Benjamin W. Lee, C. Quigg, and H. B. Thacker. Weak interactions at very high energies: The role of the Higgs-boson mass. *Physical Review D*, 16(5):1519–1531, September 1977. doi:10.1103/PhysRevD.16.1519. URL https://link.aps.org/doi/10.1103/PhysRevD.16.1519.

[2] ATLAS collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, 716 (1):1–29, September 2012. ISSN 0370-2693. doi:10.1016/j.physletb.2012.08.020. URL https://www.sciencedirect.com/science/article/pii/S037026931200857X.

[3] Tatsumi Aoyama, Toichiro Kinoshita, and Makiko Nio. Revised and Improved Value of the QED Tenth-Order Electron Anomalous Magnetic Moment. *Physical Review D*, 97(3):036001, February 2018. ISSN 2470-0010, 2470-0029. doi:10.1103/PhysRevD.97.036001.

[4] X. Fan, T. G. Myers, B. A. D. Sukra, and G. Gabrielse. Measurement of the Electron Magnetic Moment. *Physical Review Letters*, 130(7):071801, February 2023. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.130.071801.

[5] Gianfranco Bertone, Dan Hooper, and Joseph Silk. Particle dark matter: evidence, candidates and constraints. *Physics Reports*, 405(5):279–390, January 2005. ISSN 0370-1573. doi:10.1016/j.physrep.2004.08.031. URL https://www.sciencedirect.com/science/article/pii/S0370157304003515.

[6] M. C. Gonzalez-Garcia and Yosef Nir. Neutrino masses and mixing: evidence and implications. *Reviews of Modern Physics*, 75(2):345–402, March 2003. ISSN 0034-6861, 1539-0756. doi:10.1103/RevModPhys.75.345. URL https://link.aps.org/doi/10.1103/RevModPhys.75.345.

[7] Michael Dine and Alexander Kusenko. The Origin of the Matter-Antimatter Asymmetry. *Reviews of Modern Physics*, 76(1):1–30, December 2003. ISSN 0034-6861, 1539-0756. doi:10.1103/RevModPhys.76.1. _eprint: hep-ph/0303065.

[8] Patrick Huet. Electroweak Baryogenesis and the Standard Model, June 1994. URL http://arxiv.org/abs/hep-ph/9406301. arXiv:hep-ph/9406301.

[9] M. B. Gavela, P Hernandez, J Orloff, and O Pene. Standard Model CP-violation and Baryon Asymmetry. *Modern Physics Letters A*, 09(09):795–809, 1994.

[10] G. F. Giudice. Naturally Speaking: The Naturalness Criterion and Physics at the LHC. pages 155–178. June 2008. doi:10.1142/9789812779762_0010. URL http://arxiv.org/abs/0801.2562. arXiv:0801.2562 [hep-ph].

[11] S. Carlip. Quantum Gravity: a Progress Report. *Reports on Progress in Physics*, 64(8): 885–942, August 2001. ISSN 0034-4885, 1361-6633. doi:10.1088/0034-4885/64/8/301. URL http://arxiv.org/abs/gr-qc/0108040. arXiv:gr-qc/0108040.

[12] Stephen P. Martin. A Supersymmetry Primer. volume 18, pages 1–98. July 1998. doi:10.1142/9789812839657_0001. URL http://arxiv.org/abs/hep-ph/9709356. arXiv:hep-ph/9709356.

[13] Anthony Zee. Quantum Field Theory in a Nutshell | Princeton University Press, February 2010. URL https://press.princeton.edu/books/hardcover/9780691140346/quantum-field-theory-in-a-nutshell. ISBN: 9780691140346.

[14] G. Jungman, M. Kamionkowski, and K. Griest. Supersymmetric Dark Matter. *Physics Reports*, 267(5-6):195–373, March 1996. ISSN 03701573. doi:10.1016/0370-1573(95)00058-5. URL http://arxiv.org/abs/hep-ph/9506380. arXiv:hep-ph/9506380.

[15] Patrick Huet and Ann E. Nelson. Electroweak baryogenesis in supersymmetric models. *Physical Review D*, 53(8):4578–4597, April 1996. doi:10.1103/PhysRevD.53.4578. URL https://link.aps.org/doi/10.1103/PhysRevD.53.4578.

[16] Michael B. Green, John H. Schwarz, and Edward Witten. *Superstring Theory: 25th Anniversary Edition: Volume 1: Introduction*, volume 1 of *Cambridge Monographs on Mathematical Physics*. Cambridge University Press, Cambridge, 2012. ISBN 978-1-107-02911-8. doi:10.1017/CBO9781139248563. URL https://www.cambridge.org/core/books/superstring-theory/35152940E2D9C3B0BD327EB7FF74DBBE.

[17] ATLAS Collaboration. Search for the electroweak production of supersymmetric particles in s = 8 tev p p collisions with the atlas detector. *Physical Review D*, 93(5):052002, March 2016. ISSN 2470-0010, 2470-0029. doi:10.1103/PhysRevD.93.052002. URL https://link.aps.org/doi/10.1103/PhysRevD.93.052002.

[18] ATLAS collaboration. Summary of the ATLAS experiment's sensitivity to supersymmetry after LHC Run 1 - interpreted in the phenomenological MSSM. *JHEP*, 10:134, 2015. doi:10.1007/JHEP10(2015)134. _eprint: 1508.06608.

[19] Jonathan L. Feng. Naturalness and the Status of Supersymmetry. *Annual Review of Nuclear and Particle Science*, 63(1):351–382, October 2013. ISSN 0163-8998, 1545-4134. doi:10.1146/annurev-nucl-102010-130447.

[20] Tanya S. Roussy, Luke Caldwell, Trevor Wright, William B. Cairncross, Yuval Shagam, Kia Boon Ng, Noah Schlossberger, Sun Yool Park, Anzhou Wang, Jun Ye, and Eric A. Cornell. A New Bound on the Electron's Electric Dipole Moment. *arXiv*, 2022. doi:10.48550/ARXIV.2212.11841.

[21] M. G. Aartsen and others. Improved limits on dark matter annihilation in the Sun with the 79-string IceCube detector and implications for supersymmetry. *JCAP*, 04:022, 2016. doi:10.1088/1475-7516/2016/04/022. _eprint: 1601.00653.

[22] The ACME Collaboration, J. Baron, W. C. Campbell, D. DeMille, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, E. Kirilov, I. Kozyryev,

B. R. O'Leary, C. D. Panda, M. F. Parsons, E. S. Petrik, B. Spaun, A. C. Vutha, and A. D. West. Order of Magnitude Smaller Limit on the Electric Dipole Moment of the Electron. *Science*, 343(6168):269–272, January 2014. ISSN 0036-8075, 1095-9203. doi:10.1126/science.1248213.

[23] The ACME Collaboration, V. Wirthl, D. G. Ang, D. DeMille, J. M. Doyle, G. Gabrielse, J. Haefner, N. R. Hutzler, Z. Lasner, C. Meisenhelder, B. R. O'Leary, C. D. Panda, A. D. West, E. P. West, and X. Wu. Improved Limit on the Electric Dipole Moment of the Electron. *Nature*, 562(7727):355–360, October 2018. ISSN 0028-0836, 1476-4687. doi:10.1038/s41586-018-0599-8.

[24] Maxim Pospelov and Adam Ritz. CKM Benchmarks for Electron Electric Dipole Moment Experiments. *Physical Review D*, 89(5):056006, March 2014. ISSN 1550-7998, 1550-2368. doi:10.1103/PhysRevD.89.056006.

[25] Jonathan Engel, Michael J. Ramsey-Musolf, and U. van Kolck. Electric Dipole Moments of Nucleons, Nuclei, and Atoms: The Standard Model and Beyond. *Progress in Particle and Nuclear Physics*, 71:21–74, July 2013. ISSN 0146-6410. doi:10.1016/j.ppnp.2013.03.003.

[26] Maxim Pospelov and Adam Ritz. Electric Dipole Moments as Probes of New Physics. *Annals of Physics*, 318(1):119–169, July 2005. ISSN 00034916. doi:10.1016/j.aop.2005.04.002.

[27] Xing Fan. EDM Limit Calculator. URL `https://xfa9626.github.io/BSM/index.html`.

[28] P.G.H. Sandars. The Electric Dipole Moment of an Atom. *Physics Letters*, 14(3): 194–196, February 1965. ISSN 00319163. doi:10.1016/0031-9163(65)90583-4.

[29] L. I. Schiff. Measurability of Nuclear Electric Dipole Moments. *Physical Review*, 132 (5):2194–2200, December 1963. ISSN 0031-899X. doi:10.1103/PhysRev.132.2194.

[30] Eugene D. Commins, J. D. Jackson, and David P. DeMille. The Electric Dipole Moment of the Electron: An Intuitive Explanation for the Evasion of Schiff's Theorem. *American Journal of Physics*, 75(6):532–536, June 2007. ISSN 0002-9505, 1943-2909. doi:10.1119/1.2710486.

[31] Eugene D Commins and David P DeMille. The Electric Dipole Moment of the Electron.

[32] Iosif B. Khriplovich. *Cp Violation without Strangeness: Electric Dipole Moments of Particles, Atoms, and Molecules*. Springer, Place of publication not identified, 2012. ISBN 978-3-642-60838-4.

[33] Malika Denis and Timo Fleig. In Search of Discrete Symmetry Violations beyond the Standard Model: Thorium Monoxide Reloaded. *The Journal of Chemical Physics*, 145 (21):214307, December 2016. ISSN 0021-9606, 1089-7690. doi:10.1063/1.4968597.

[34] L. V. Skripnikov. Combined 4-Component and Relativistic Pseudopotential Study of ThO for the Electron Electric Dipole Moment Search. *The Journal of Chemical Physics*, 145(21):214301, December 2016. ISSN 0021-9606, 1089-7690. doi:10.1063/1.4968229.

[35] M. C. Weisskopf, J. P. Carrico, H. Gould, E. Lipworth, and T. S. Stein. Electric Dipole Moment of the Cesium Atom. A New Upper Limit to the Electric Dipole Moment of the Electron. *Physical Review Letters*, 21(24):1645–1648, December 1968. ISSN 0031-9007. doi:10.1103/PhysRevLett.21.1645.

[36] K. Abdullah, C. Carlberg, E. D. Commins, Harvey Gould, and Stephen B. Ross. New Experimental Limit on the Electron Electric Dipole Moment. *Physical Review Letters*, 65(19):2347–2350, November 1990. ISSN 0031-9007. doi:10.1103/PhysRevLett.65.2347.

[37] Eugene D. Commins, Stephen B. Ross, David DeMille, and B. C. Regan. Improved Experimental Limit on the Electric Dipole Moment of the Electron. *Physical Review A*, 50(4):2960–2977, October 1994. ISSN 1050-2947, 1094-1622. doi:10.1103/PhysRevA.50.2960.

[38] S. A. Murthy, D. Krause, Z. L. Li, and L. R. Hunter. New Limits on the Electron Electric Dipole Moment from Cesium. *Physical Review Letters*, 63(9):965–968, August 1989. ISSN 0031-9007. doi:10.1103/PhysRevLett.63.965.

[39] M A Player and P G H Sandars. An Experiment to Search for an Electric Dipole Moment in the \textsuperscript3 P \textsubscript2 Metastable State of Xenon. *Journal of Physics B: Atomic and Molecular Physics*, 3(12):1620–1635, December 1970. ISSN 0022-3700. doi:10.1088/0022-3700/3/12/007.

[40] P. G. H. Sandars and E. Lipworth. Electric Dipole Moment of the Cesium Atom. A New Upper Limit to the Electric Dipole Moment of the Free Electron. *Physical Review Letters*, 13(24):718–720, December 1964. ISSN 0031-9007. doi:10.1103/PhysRevLett.13.718.

[41] S. Bickman, P. Hamilton, Y. Jiang, and D. DeMille. Preparation and Detection of States with Simultaneous Spin Alignment and Selectable Molecular Orientation in PbO. *Physical Review A*, 80(2):023418, August 2009. ISSN 1050-2947, 1094-1622. doi:10.1103/PhysRevA.80.023418.

[42] J. J. Hudson, B. E. Sauer, M. R. Tarbutt, and E. A. Hinds. Measurement of the Electron Electric Dipole Moment Using YbF Molecules. *Physical Review Letters*, 89(2): 023003, June 2002. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.89.023003.

[43] J. J. Hudson, D. M. Kara, I. J. Smallman, B. E. Sauer, M. R. Tarbutt, and E. A. Hinds. Improved Measurement of the Shape of the Electron. *Nature*, 473(7348):493–496, May 2011. ISSN 0028-0836, 1476-4687. doi:10.1038/nature10104.

[44] William B. Cairncross, Daniel N. Gresh, Matt Grau, Kevin C. Cossel, Tanya S. Roussy, Yiqi Ni, Yan Zhou, Jun Ye, and Eric A. Cornell. Precision Measurement of the Electron's Electric Dipole Moment Using Trapped Molecular Ions. *Physical Review Letters*, 119(15):153001, October 2017. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.119.153001.

[45] B. C. Regan, Eugene D. Commins, Christian J. Schmidt, and David DeMille. New Limit on the Electron Electric Dipole Moment. *Physical Review Letters*, 88(7):071805, February 2002. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.88.071805.

[46] Gerhart Lüders. Proof of the TCP theorem. *Annals of Physics*, 2(1):1–15, July 1957. ISSN 0003-4916. doi:10.1016/0003-4916(57)90032-5. URL `https://www.sciencedirect.com/science/article/pii/0003491657900325`.

[47] G. Gabrielse, S. Fogwell Hoogerheide, J. Dorr, and E. Novitski. Precise Matter and Antimatter Tests of the Standard Model with E\textsuperscript-, E\textsuperscript+, p, \bar\textrmp, and \bar\textrmH. In Wolfgang Quint and Manuel Vogel, editors, *Fundamental Physics in Particle Traps*, volume 256, pages 1–40. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014. ISBN 978-3-642-45200-0 978-3-642-45201-7. doi:10.1007/978-3-642-45201-7_1.

[48] J. DiSciacca, M. Marshall, K. Marable, G. Gabrielse, S. Ettenauer, E. Tardiff, R. Kalra, D. W. Fitzakerley, M. C. George, E. A. Hessels, C. H. Storry, M. Weel, D. Grzonka, W. Oelert, T. Sefzick, and ATRAP Collaboration. One-Particle Measurement of the Antiproton Magnetic Moment. *Physical Review Letters*, 110(13):130801, March 2013. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.110.130801.

[49] F. Ambrosino, A. Antonelli, M. Antonelli, C. Bacci, P. Beltrame, G. Bencivenni, S. Bertolucci, C. Bini, C. Bloise, S. Bocchetta, V. Bocci, F. Bossi, D. Bowring, P. Branchini, R. Caloi, P. Campana, G. Capon, T. Capussela, Filippo Ceradini, S. Chi, G. Chiefari, P. Ciambrone, S. Conetti, E. De Lucia, A. De Santis, P. De Simone, G. De Zorzi, S. Dell'Agnello, A. Denig, A. Di Domenico, C. Di Donato, S. Di Falco, Biagio Di Micco, A. Doria, M. Dreucci, G. Felici, A. Ferrari, M. L. Ferrer, G. Finocchiaro, S. Fiore, C. Forti, P. Franzini, C. Gatti, P. Gauzzi, S. Giovannella, E. Gorini, E. Graziani, M. Incagli, W. Kluge, V. Kulikov, F. Lacava, G. Lanfranchi, J. Lee Franzini, D. Leone, M. Martini, P. Massarotti, W. Mei, S. Meola, S. Miscetti, M. Moulson, S. Muller, F. Murtas, M. Napolitano, F. Nguyen, M. Palutan, E. Pasqualucci, A. Passeri, V. Patera, F. Perfetto, L. Pontecorvo, M. Primavera, P. Santangelo, E. Santovetti, G. Saracino, B. Sciascia, A. Sciubba, F. Scuri, I. Sfiligoi, A. Sibidanov, T. Spadaro, M. Testa, L. Tortora, P. Valente, B. Valeriani, G. Venanzoni, S. Veneziano, A. Ventura, R. Versaci, and G. Xu. First observation of quantum interference in the process phi -> KSKL -> pi(+)pi(-)pi(+)pi(-): A test of quantum mechanics and C P T symmetry. 2006. doi:10.1016/j.physletb.2006.09.046. URL `https://iris.uniroma3.it/handle/11590/121465`. Accepted: 2015-09-24T13:15:42Z.

[50] John M. Brown and Alan Carrington. *Rotational Spectroscopy of Diatomic Molecules*. Cambridge Molecular Science Series. Cambridge University Press, Cambridge ; New York, 2003. ISBN 978-0-521-81009-8 978-0-521-53078-1.

[51] Ivan Kozyryev and Nicholas R. Hutzler. Precision Measurement of Time-Reversal Symmetry Violation with Laser-Cooled Polyatomic Molecules. *Physical Review Letters*, 119(13):133002, September 2017. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.119.133002.

[52] D. G. Ang, C. Meisenhelder, C. D. Panda, X. Wu, D. DeMille, J. M. Doyle, and G. Gabrielse. Measurement of the H \textsuperscript3Δ\textsubscript1 Radiative Lifetime in ThO. *Physical Review A*, 106(2):022808, August 2022. ISSN 2469-9926, 2469-9934. doi:10.1103/PhysRevA.106.022808.

[53] Nicholas Richard Hutzler. *A New Limit on the Electron Electric Dipole Moment: Beam Production, Data Interpretation, and Systematics*. PhD Thesis, Harvard Univeristy, 2014.

[54] Nicholas R. Hutzler, Hsin-I Lu, and John M. Doyle. The Buffer Gas Beam: An Intense, Cold, and Slow Source for Atoms and Molecules. *Chemical Reviews*, 112(9):4803–4827, September 2012. ISSN 0009-2665, 1520-6890. doi:10.1021/cr200362u.

[55] Nicholas R. Hutzler, Maxwell F. Parsons, Yulia V. Gurevich, Paul W. Hess, Elizabeth Petrik, Ben Spaun, Amar C. Vutha, David DeMille, Gerald Gabrielse, and John M. Doyle. A Cryogenic Beam of Refractory, Chemically Reactive Molecules with Expansion Cooling. *Physical Chemistry Chemical Physics*, 13(42):18976, 2011. ISSN 1463-9076, 1463-9084. doi:10.1039/c1cp20901a.

[56] Amar Chandra Vutha. *A Search for the Electric Dipole Moment of the Electron Using Thorium Monoxide*. PhD Thesis, Yale University, 2011.

[57] Oskari Timgren. Progress towards a measurement of time-reversal symmetry violation in thallium fluoride.

[58] Yohei Ema, Ting Gao, and Maxim Pospelov. Standard Model Prediction for Paramagnetic EDMs. *Physical Review Letters*, 129(23):231801, November 2022. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.129.231801. _eprint: 2202.10524.

[59] Yuiki Takahashi, Chi Zhang, Arian Jadbabaie, and Nicholas R. Hutzler. Engineering Field-Insensitive Molecular Clock Transitions for Symmetry Violation Searches. *Physical Review Letters*, 131(18):183003, October 2023. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.131.183003. URL `https://link.aps.org/doi/10.1103/PhysRevLett.131.183003`.

[60] Loïc Anderegg, Nathaniel B. Vilas, Christian Hallas, Paige Robichaud, Arian Jadbabaie, John M. Doyle, and Nicholas R. Hutzler. Quantum control of trapped polyatomic molecules for eEDM searches. *Science*, 382(6671):665–668, November 2023.

ISSN 0036-8075, 1095-9203. doi:10.1126/science.adg8155. URL https://www.science.org/doi/10.1126/science.adg8155.

[61] Kia Boon Ng, Sun Yool Park, Anzhou Wang, Addison Hartman, Patricia Hector Hernandez, Rohan Kompella, Lan Cheng, Stephan Malbrunot-Ettenauer, Jun Ye, and Eric A. Cornell. High-efficiency quantum-state detection of ThF + with resonance-enhanced multiphoton asymmetric dissociation. *Physical Review A*, 112(5): 053104, November 2025. ISSN 2469-9926, 2469-9934. doi:10.1103/mzd6-9kml. URL https://link.aps.org/doi/10.1103/mzd6-9kml.

[62] Sun Yool Park, Anzhou Wang, Addison Hartman, Patricia Hector Hernandez, Kia Boon Ng, Rohan Kompella, Jun Ye, and Eric Cornell. The JILA Gen III Experiment on electron's Electric Dipole Moment (eEDM): Achieving long coherence time. *Bulletin of the American Physical Society*, 2025. URL https://meetings.aps.org/Meeting/DAMOP25/Session/E09.2.

[63] Anzhou Wang, Sun Yool Park, Addison Hartman, Patricia Hector Hernandez, Jun Ye, and Eric Cornell. Towards continuous load and readout of the JILA Generation III eEDM experiment. *Bulletin of the American Physical Society*, 2025. URL https://meetings.aps.org/Meeting/DAMOP25/Session/D00.133.

[64] F. J. Collings, N. J. Fitch, R. A. Jenkins, J. M. Dyne, E. Wursten, M. T. Ziemba, X. S. Zheng, F. Castellini, J. Lim, and B. E. Sauer. Low-noise environment for probing fundamental symmetries. *New Journal of Physics*, 27(10):105004, 2025. URL https://iopscience.iop.org/article/10.1088/1367-2630/ae0ea7/meta.

[65] A. Marsman, D. Heinrich, M. Horbatsch, E. A. Hessels, and EDM3 Collaboration. Large optical forces on a barium monofluoride molecule using laser pulses for absorption and stimulated emission: A full density-matrix simulation. *Physical Review A*, 107(3): 032811, March 2023. ISSN 2469-9926, 2469-9934. doi:10.1103/PhysRevA.107.032811. URL https://link.aps.org/doi/10.1103/PhysRevA.107.032811.

[66] B. Graner, Y. Chen, E. G. Lindahl, and B. R. Heckel. Reduced Limit on the Permanent Electric Dipole Moment of Hg 199. *Physical Review Letters*, 116(16):161601, April 2016. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.116.161601.

[67] Olivier Grasdijk, Oskari Timgren, Jakob Kastelic, Trevor Wright, Steve Lamoreaux, David DeMille, Konrad Wenz, Michael Aitken, Tanya Zelevinsky, and Tristan Winick. CeNTREX: a new search for time-reversal symmetry violation in the 205Tl nucleus. *Quantum Science and Technology*, 6(4):044007, 2021. URL https://iopscience.iop.org/article/10.1088/2058-9565/abdca3/meta.

[68] V. V. Flambaum and V. A. Dzuba. Electric dipole moments of atoms and molecules produced by enhanced nuclear Schiff moments. *Physical Review A*, 101(4):042504, April 2020. doi:10.1103/PhysRevA.101.042504. URL https://link.aps.org/doi/10.1103/PhysRevA.101.042504.

[69] V. V. Flambaum and H. Feldmeier. Enhanced nuclear Schiff moment in stable and metastable nuclei. *Physical Review C*, 101(1):015502, January 2020. doi:10.1103/PhysRevC.101.015502. URL https://link.aps.org/doi/10.1103/PhysRevC.101.015502.

[70] Mohit Verma, Shaozhen Yang, Alan Jamison, and David DeMille. A new search for flavor-conserving hadronic CP violation using ultracold assembled 223 FrAg molecules. *Bulletin of the American Physical Society*, 2023. URL https://meetings.aps.org/Meeting/HAW23/Session/C14.7.

[71] Shaozhen Yang, Mohit Verma, Wesley Cassidy, Harish Ramachandran, Abby Kinney, Alan O. Jamison, and David P. DeMille. Searching for hadronic CP violation using ultracold assembled FrAg molecules. In *DNP 2025*. APS, 2025. URL https://archive.aps.org/dnp/2025/n08/8/.

[72] Bassam Nima, Mingyu Fan, Aleksandar Radak, Andrew M. Jayich, and Amar Vutha. Precision comagnetometry for T -violation searches in crystals. *Physical Review A*, 112(3):L030801, September 2025. ISSN 2469-9926, 2469-9934. doi:10.1103/jykv-rsd1. URL https://link.aps.org/doi/10.1103/jykv-rsd1.

[73] H. D. Ramachandran and A. C. Vutha. Nuclear T -violation search using octopole-deformed nuclei in a crystal. *Physical Review A*, 108(1):012819, July 2023. ISSN 2469-9926, 2469-9934. doi:10.1103/PhysRevA.108.012819. URL https://link.aps.org/doi/10.1103/PhysRevA.108.012819.

[74] Xing Wu, Sebastian Miki-Silva, Nathan Czopp, Alexander J. Frenett, Maggie Tseng, Monika Fouad, Nicholas Emtage, and Dorothy Gan. Towards a magnetic centrifuge decelerator for testing fundamental symmetries. In *DNP 2025*. APS, 2025. URL https://archive.aps.org/dnp/2025/p08/1/.

[75] Ronald Fernando Garcia Ruiz. Towards Laser Cooling of Radium Containing Molecules. *Bulletin of the American Physical Society*, 2025. URL https://meetings.aps.org/Meeting/DAMOP25/Session/B07.1.

[76] T. A. Zheng, Y. A. Yang, S.-Z. Wang, J. T. Singh, Z.-X. Xiong, T. Xia, and Z.-T. Lu. Measurement of the Electric Dipole Moment of Yb 171 Atoms in an Optical Dipole Trap. *Physical Review Letters*, 129(8):083001, August 2022. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.129.083001. URL https://link.aps.org/doi/10.1103/PhysRevLett.129.083001.

[77] R. H. Parker, M. R. Dietrich, M. R. Kalita, N. D. Lemke, K. G. Bailey, M. Bishof, J. P. Greene, R. J. Holt, W. Korsch, Z.-T. Lu, P. Mueller, T. P. O'Connor, and J. T. Singh. First Measurement of the Atomic Electric Dipole Moment of Ra 225. *Physical Review Letters*, 114(23):233002, June 2015. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.114.233002.

[78] C. Abel, S. Afach, N. J. Ayres, C. A. Baker, G. Ban, G. Bison, K. Bodek, V. Bondar, M. Burghoff, E. Chanel, Z. Chowdhuri, P.-J. Chiu, B. Clement, C. B. Crawford, M. Daum, S. Emmenegger, L. Ferraris-Bouchez, M. Fertl, P. Flaux, B. Franke, A. Fratangelo, P. Geltenbort, K. Green, W. C. Griffith, M. van der Grinten, Z. D. Grujić, P. G. Harris, L. Hayen, W. Heil, R. Henneck, V. Hélaine, N. Hild, Z. Hodge, M. Horras, P. Iaydjiev, S. N. Ivanov, M. Kasprzak, Y. Kermaidic, K. Kirch, A. Knecht, P. Knowles, H.-C. Koch, P. A. Koss, S. Komposch, A. Kozela, A. Kraft, J. Krempel, M. Kuźniak, B. Lauss, T. Lefort, Y. Lemière, A. Leredde, P. Mohanmurthy, A. Mtchedlishvili, M. Musgrave, O. Naviliat-Cuncic, D. Pais, F. M. Piegsa, E. Pierre, G. Pignol, C. Plonka-Spehr, P. N. Prashanth, G. Quéméner, M. Rawlik, D. Rebreyend, I. Rienäcker, D. Ries, S. Roccia, G. Rogel, D. Rozpedzik, A. Schnabel, P. Schmidt-Wellenburg, N. Severijns, D. Shiers, R. Tavakoli Dinani, J. A. Thorne, R. Virot, J. Voigt, A. Weis, E. Wursten, G. Wyszynski, J. Zejma, J. Zenner, and G. Zsigmond. Measurement of the Permanent Electric Dipole Moment of the Neutron. *Physical Review Letters*, 124(8):081803, February 2020. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.124.081803. URL `https://link.aps.org/doi/10.1103/PhysRevLett.124.081803`.

[79] Ryohei Matsumiya, Hiroaki Akatsuka, Chris P. Bidinosti, Charles A. Davis, Beatrice Franke, Derek Fujimoto, Michael T. W. Gericke, Pietro Giampa, Robert Golub, Sean Hansen-Romu, Kichiji Hatanaka, Tomohiro Hayamizu, Takashi Higuchi, Go Ichikawa, Sohei Imajo, Blair Jamieson, Shinsuke Kawasaki, Masaaki Kitaguchi, Wolfgang Klassen, Emma Klemets, Akira Konaka, Elie Korkmaz, Ekaterina Korobkina, Florian Kuchler, Maedeh Lavvaf, Larry Lee, Thomas Lindner, Kirk W. Madison, Yasuhiro Makida, Russell Mammei, Juliette Mammei, Jeffery W. Martin, Mark McCrea, Eric Miller, Kenji Mishima, Takamasa Momose, Takahiro Okamura, Hooi Jin Ong, Ruediger Picker, William D. Ramsay, Wolfgang Schreyer, Hirohiko M. Shimizu, Steve Sidhu, Shawn Stargardter, Isao Tanihata, Sean Vanbergen, Willem T. H. van Oers, and Yutaka Watanabe. The Precision nEDM Measurement with UltraCold Neutrons at TRIUMF. In *Proceedings of the 24th International Spin Symposium (SPIN2021)*, December 2022. doi:10.7566/JPSCP.37.020701. URL `http://arxiv.org/abs/2207.09880`. arXiv:2207.09880 [physics].

[80] C. A. Baker, D. D. Doyle, P. Geltenbort, K. Green, M. G. D. van der Grinten, P. G. Harris, P. Iaydjiev, S. N. Ivanov, D. J. R. May, J. M. Pendlebury, J. D. Richardson, D. Shiers, and K. F. Smith. Improved Experimental Limit on the Electric Dipole Moment of the Neutron. *Physical Review Letters*, 97(13):131801, September 2006. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.97.131801.

[81] J. M. Pendlebury, S. Afach, N. J. Ayres, C. A. Baker, G. Ban, G. Bison, K. Bodek, M. Burghoff, P. Geltenbort, K. Green, W. C. Griffith, M. van der Grinten, Z. D. Grujić, P. G. Harris, V. Hélaine, P. Iaydjiev, S. N. Ivanov, M. Kasprzak, Y. Kermaidic, K. Kirch, H.-C. Koch, S. Komposch, A. Kozela, J. Krempel, B. Lauss, T. Lefort, Y. Lemière, D. J. R. May, M. Musgrave, O. Naviliat-Cuncic, F. M. Piegsa, G. Pignol, P. N. Prashanth, G. Quéméner, M. Rawlik, D. Rebreyend, J. D. Richardson,

D. Ries, S. Roccia, D. Rozpedzik, A. Schnabel, P. Schmidt-Wellenburg, N. Severijns, D. Shiers, J. A. Thorne, A. Weis, O. J. Winston, E. Wursten, J. Zejma, and G. Zsigmond. Revised Experimental Upper Limit on the Electric Dipole Moment of the Neutron. *Physical Review D*, 92(9):092003, November 2015. ISSN 1550-7998, 1550-2368. doi:10.1103/PhysRevD.92.092003.

[82] Riccardo Schmid. *New Search for the Neutron Electric Dipole Moment Using Ultracold Neutrons at the Spallation Neutron Source.* PhD Thesis, California Institute of Technology, 2014.

[83] Rebecca Chislett. The Muon EDM in the G-2 Experiment at Fermilab. *EPJ Web of Conferences*, 118:01005, 2016. ISSN 2100-014X. doi:10.1051/epjconf/201611801005.

[84] Yulia Vsevolodovna Gurevich. *Preliminary Measurements for an Electron EDM Experiment in ThO.* PhD Thesis, Harvard Univeristy, 2012.

[85] Paul William Hess. *Improving the Limit on the Electron EDM: Data Acquisition and Systematics Studies in the ACME Experiment.* PhD Thesis, Harvard Univeristy, 2014.

[86] Benjamin Norman Spaun. *A Ten-Fold Improvement to the Limit of the Electron Electric Dipole Moment.* PhD Thesis, Harvard Univeristy, 2014.

[87] Brendon R O'Leary. *In Search of the Electron's Electric Dipole Moment in Thorium Monoxide: An Improved Upper Limit, Systematic Error Models, and Apparatus Upgrades.* PhD Thesis, Yale University, 2016.

[88] V. Wirthl. Polarimetry on the Advanced Cold Molecule Electron Electric Dipole Moment Experiment. Master's thesis, Technische Universität München, 2016.

[89] Elizabeth P. West. *A Thermochemical Cryogenic Buffer Gas Beam Source of ThO for Measuring the Electric Dipole Moment of the Electron.* PhD Thesis, Harvard Univeristy, 2017.

[90] Cristian D. Panda. *Order of Magnitude Improved Limit on the Electric Dipole Moment of the Electron.* PhD Thesis, Harvard Univeristy, 2018.

[91] Zack Lasner. *Order-of-Magnitude-Tighter Bound on the Electron Electric Dipole Moment.* PhD Thesis, Yale University, 2019.

[92] Cole Meisenhelder. Advances in the Measurement of the Electron Electric Dipole Moment, 2023. URL `https://cfp.physics.northwestern.edu/documents/2023-Meisenhelder_Thesis.pdf`.

[93] Siyuan Liu. MITIGATING MAGNETIC FIELD IMPERFECTIONS FOR PRECISION MEASUREMENT OF THE ELECTRON ELECTRIC DIPOLE MOMENT, 2024. URL `http://doylegroup.harvard.edu/wiki/images/9/93/2024-Siyuan_Liu_PhD_Thesis.pdf`.

[94] Leonid A. Kaledin, Alexey L. Kaledin, and Michael C. Heaven. The electronic structure of thorium monoxide: Ligand field assignment of states in the range 0-5 eV. *Journal of Computational Chemistry*, 40(2):430–446, January 2019. ISSN 1096-987X. doi:10.1002/jcc.25710.

[95] Damian L. Kokkin, Timothy C. Steimle, and David DeMille. Branching Ratios and Radiative Lifetimes of the U , L , and I States of Thorium Oxide. *Physical Review A*, 90(6):062503, December 2014. ISSN 1050-2947, 1094-1622. doi:10.1103/PhysRevA.90.062503.

[96] Damian L. Kokkin, Timothy C. Steimle, and David DeMille. Characterization of the I (|Ømega|=1)–X\textsuperscript1 Σ\textsuperscript+ (0 , 0) Band of Thorium Oxide. *Physical Review A*, 91(4):042508, April 2015. ISSN 1050-2947, 1094-1622. doi:10.1103/PhysRevA.91.042508.

[97] X Wu, Z Han, J Chow, D G Ang, C Meisenhelder, C D Panda, E P West, G Gabrielse, J M Doyle, and D DeMille. The Metastable Q \textsuperscript3Δ\textsubscript2 State of ThO: A New Resource for the ACME Electron EDM Search. *New Journal of Physics*, 22(2):023013, February 2020. ISSN 1367-2630. doi:10.1088/1367-2630/ab6a3a.

[98] Fang Wang, Anh Le, Timothy C. Steimle, and Michael. C. Heaven. Communication: The permanent electric dipole moment of thorium monoxide, ThO. *The Journal of Chemical Physics*, 134(3):031102, January 2011. ISSN 0021-9606. doi:10.1063/1.3533274. URL `https://doi.org/10.1063/1.3533274`.

[99] Daniel G. Ang. *Progress towards an Improved Measurement of the Electric Dipole Moment of the Electron*. PhD Thesis, Harvard University, 2023.

[100] A. N. Petrov and L. V. Skripnikov. Interference between the E1 and M1 Amplitudes of the Transition from the H State to C of a ThO Molecule. *Optics and Spectroscopy*, 126(4):331–335, April 2019. ISSN 1562-6911. doi:10.1134/S0030400X19040192. URL `https://doi.org/10.1134/S0030400X19040192`.

[101] E. Kirilov, W. C. Campbell, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, B. R. O'Leary, E. Petrik, B. Spaun, A. C. Vutha, and D. DeMille. Shot-Noise-Limited Spin Measurements in a Pulsed Molecular Beam. *Physical Review A*, 88 (1):013844, July 2013. ISSN 1050-2947, 1094-1622. doi:10.1103/PhysRevA.88.013844.

[102] Charles Townes and Arthur Schawlow. *Microwave Spectroscopy*. 2013. URL `https://store.doverpublications.com/products/9780486162317`.

[103] X Wu, P Hu, Z Han, D G Ang, C Meisenhelder, G Gabrielse, J M Doyle, and D DeMille. Electrostatic Focusing of Cold and Heavy Molecules for the ACME Electron EDM Search. *New Journal of Physics*, 24(7):073043, July 2022. ISSN 1367-2630. doi:10.1088/1367-2630/ac8014.

[104] Takahiko Masuda, Ayami Hiramoto, Daniel G. Ang, Cole Meisenhelder, Cristian D. Panda, Noboru Sasao, Satoshi Uetake, Xing Wu, David P. DeMille, John M. Doyle, Gerald Gabrielse, and Koji Yoshimura. High-Sensitivity Low-Noise Photodetector Using a Large-Area Silicon Photomultiplier. *Optics Express*, 31(2):1943, January 2023. ISSN 1094-4087. doi:10.1364/OE.475109.

[105] Takahiko Masuda, Daniel G. Ang, Nicholas R. Hutzler, Cole Meisenhelder, Noboru Sasao, Satoshi Uetake, Xing Wu, David DeMille, Gerald Gabrielse, John M. Doyle, and Koji Yoshimura. Suppression of the Optical Crosstalk in a Multi-Channel Silicon Photomultiplier Array. *Optics Express*, 29(11):16914, May 2021. ISSN 1094-4087. doi:10.1364/OE.424460.

[106] Zhen Han, Zack Lasner, Collin Diver, Peiran Hu, Takahiko Masuda, Xing Wu, Ayami Hiramoto, Maya Watts, Satoshi Uetake, Koji Yoshimura, Xing Fan, Gerald Gabrielse, John M. Doyle, and David DeMille. A cryogenic buffer gas beam source with in-situ ablation target replacement, May 2025. URL `http://arxiv.org/abs/2505.11647`. arXiv:2505.11647 [physics].

[107] C D Panda, C Meisenhelder, M Verma, D G Ang, J Chow, Z Lasner, X Wu, D DeMille, J M Doyle, and G Gabrielse. Attaining the Shot-Noise-Limit in the ACME Measurement of the Electron Electric Dipole Moment. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 52(23):235003, December 2019. ISSN 0953-4075, 1361-6455. doi:10.1088/1361-6455/ab4a61.

[108] Adam D. West, Zack Lasner, David DeMille, Elizabeth P. West, Cristian D. Panda, John M. Doyle, Gerald Gabrielse, Adam Kryskow, and Corinne Mitchell. An Underappreciated Radiation Hazard from High Voltage Electrodes in Vacuum. *Health Physics*, 112(1):33–41, January 2017. ISSN 0017-9078. doi:10.1097/HP.0000000000000593.

[109] W. A. Stygar, J. A. Lott, T. C. Wagoner, V. Anaya, H. C. Harjes, H. C. Ives, Z. R. Wallace, G. R. Mowrer, R. W. Shoup, J. P. Corley, R. A. Anderson, G. E. Vogtlin, M. E. Savage, J. M. Elizondo, B. S. Stoltzfus, D. M. Andercyk, D. L. Fehl, T. F. Jaramillo, D. L. Johnson, D. H. McDaniel, D. A. Muirhead, J. M. Radman, J. J. Ramirez, L. E. Ramirez, R. B. Spielman, K. W. Struve, D. E. Walsh, E. D. Walsh, and M. D. Walsh. Improved design of a high-voltage vacuum-insulator interface. *Physical Review Special Topics - Accelerators and Beams*, 8(5):050401, May 2005. ISSN 1098-4402. doi:10.1103/PhysRevSTAB.8.050401. URL `https://link.aps.org/doi/10.1103/PhysRevSTAB.8.050401`.

[110] Jakob Kastelic. Search for Time-Reversal-Symmetry Violation in Thallium Fluoride Using a Cryogenic Buffer-Gas Beam Source.

[111] V. Wirthl, C. D. Panda, P. W. Hess, B. Spaun, and G. Gabrielse. A Self-Calibrating Polarimeter to Measure Stokes Parameters, February 2021. Issue: arXiv:1703.00963 _eprint: 1703.00963.

[112] M. Guignard, L. Albrecht, and J. W. Zwanziger. Zero-Stress Optic Glass without Lead. *Chemistry of Materials*, 19(2):286–290, January 2007. ISSN 0897-4756. doi:10.1021/cm062208a. URL https://doi.org/10.1021/cm062208a.

[113] LIGO. LIGO-E1000386-v1: Material Qualification RGA Test Results: MasterBond EP30-2 epoxy. URL https://dcc.ligo.org/LIGO-E1000386/public.

[114] Hillar Aben and Claude Guillemet. *Photoelasticity of Glass*. Springer Science & Business Media, December 2012. ISBN 978-3-642-50071-8. Google-Books-ID: 1VfrCAAAQBAJ.

[115] Schott. Schott-Tie-27-Stress-in-Optical-Glass-Eng.Pdf, July 2004. URL https://www.schott.com/shop/medias/schott-tie-27-stress-in-optical-glass-eng.pdf.

[116] J. R. Barber. *Elasticity*. Number v. 172 in Solid Mechanics and Its Applications. Springer, Dordrecht ; New York, 3rd rev. ed edition, 2010. ISBN 978-90-481-3808-1 978-90-481-3809-8.

[117] A. R. Lamichhane and C. S. Chen. The closed-form particular solutions for Laplace and biharmonic operators using a Gaussian function. *Applied Mathematics Letters*, 46:50–56, August 2015. ISSN 0893-9659. doi:10.1016/j.aml.2015.02.004. URL https://www.sciencedirect.com/science/article/pii/S0893965915000531.

[118] A. Hiramoto, T. Masuda, D.G. Ang, C. Meisenhelder, C. Panda, N. Sasao, S. Uetake, X. Wu, D. Demille, J.M. Doyle, G. Gabrielse, and K. Yoshimura. SiPM Module for the ACME III Electron EDM Search. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1045:167513, January 2023. ISSN 01689002. doi:10.1016/j.nima.2022.167513.

[119] Luc Devroye. *Non-Uniform Random Variate Generation*. Springer, New York, NY, 1986. ISBN 978-1-4613-8645-2 978-1-4613-8643-8. doi:10.1007/978-1-4613-8643-8. URL http://link.springer.com/10.1007/978-1-4613-8643-8.

[120] A. Marsman, E. A. Hessels, and M. Horbatsch. Shifts due to quantum-mechanical interference from distant neighboring resonances for saturated fluorescence spectroscopy of the $2{\phantom{\rule{0.16em}{0ex}}}^{3}S$ to $2{\phantom{\rule{0.16em}{0ex}}}^{3}P$ intervals of helium. *Physical Review A*, 89(4): 043403, April 2014. doi:10.1103/PhysRevA.89.043403. URL https://link.aps.org/doi/10.1103/PhysRevA.89.043403.

[121] A. Marsman, M. Horbatsch, and E. A. Hessels. The Effect of Quantum-Mechanical Interference on Precise Measurements of the n = 2 Triplet P Fine Structure of Helium. *Journal of Physical and Chemical Reference Data*, 44(3):031207, July 2015. ISSN 0047-2689. doi:10.1063/1.4922796. URL https://doi.org/10.1063/1.4922796.

[122] Amar Vutha and David DeMille. Geometric phases without geometry, December 2009. URL http://arxiv.org/abs/0907.5116. arXiv:0907.5116 [quant-ph].

[123] A C Vutha, W C Campbell, Y V Gurevich, N R Hutzler, M Parsons, D Patterson, E Petrik, B Spaun, J M Doyle, G Gabrielse, and D DeMille. Search for the Electric Dipole Moment of the Electron with Thorium Monoxide. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 43(7):074007, April 2010. ISSN 0953-4075, 1361-6455. doi:10.1088/0953-4075/43/7/074007.

[124] S. Eckel, P. Hamilton, E. Kirilov, H. W. Smith, and D. DeMille. Search for the Electron Electric Dipole Moment Using Ømega -Doublet Levels in PbO. *Physical Review A*, 87 (5):052130, May 2013. ISSN 1050-2947, 1094-1622. doi:10.1103/PhysRevA.87.052130.

[125] David Jiles. *Introduction to Magnetism and Magnetic Materials*. Springer Netherlands, 1991. ISBN 978-0-412-38630-5. Google-Books-ID: FrUPAQAAMAAJ.

[126] Giorgio Bertotti. *Hysteresis in Magnetism*. May 1998. ISBN 978-0-12-093270-2. URL `https://shop.elsevier.com/books/hysteresis-in-magnetism/bertotti/978-0-12-093270-2`.

[127] Dmitry Budker and Derek F. Jackson Kimball, editors. *Optical Magnetometry*. Cambridge University Press, Cambridge, 2013. ISBN 978-1-107-01035-2. doi:10.1017/CBO9780511846380. URL `https://www.cambridge.org/core/books/optical-magnetometry/33E638960A4E6B159DF83F82BBDDC5E1`.