

A new limit on the electron's electric dipole moment

by

Tanya S. Roussy

BASc., University of British Columbia, 2014

MASc., University of British Columbia, 2016

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics

2022

Committee Members:

Eric A. Cornell, Chair

Prof. Jun Ye

Prof. John Bohn

Prof. Keith Ulmer

Prof. Zach Berta-Thompson

Roussy, Tanya S. (Ph.D., Physics)

A new limit on the electron's electric dipole moment

Thesis directed by Prof. Eric A. Cornell

Physicists have been able to describe our entire universe using just two mathematical models: General Relativity, for gravity, and the Standard Model, for everything else. While the Standard Model has been tested to exquisite precision and is generally successful, it has some serious inconsistencies: it fails, for example, to explain why matter dominates over antimatter in our universe. New models aiming to eliminate these inconsistencies often introduce additional symmetry violation, which in turn results in symmetry-violating properties in fundamental particles like electrons. In particular, many new physics theories predict that the electron will have a non-zero electric dipole moment (EDM) aligned with the spin axis, which violates time-reversal symmetry. Placing limits on the magnitude of the electron EDM directly constrains new physics theories at energies beyond the reach of the Large Hadron Collider. Here we present a new experimental limit on the electron's electric dipole moment using trapped HfF^+ in rotating bias fields: $|d_e| < 4.4 \times 10^{-30}$ e-cm. Our measurement is a factor of 2.5 better than the previous best limit.

Dedication

To everyone fighting for a better world.

"It is our duty to fight for our freedom.

It is our duty to win.

We must love each other and support each other.

We have nothing to lose but our chains".

- Assata Shakur

Acknowledgements

For my parents. Ok, ok, I'll get a real job

Contents

Chapter

1	Introduction	1
1.1	Motivation	1
1.2	JILA's eEDM experiment(s): a brief history	4
1.3	A non-exhaustive list of other searches for the eEDM	9
1.4	Overview of this thesis	10
2	How the experiment works	11
2.1	Measuring the electron's electric dipole moment: the basic recipe	11
2.1.1	The basic recipe: put a dipole in a field - twice	11
2.1.2	Key ingredient 1: sensitivity	12
2.1.3	Key ingredient 2: bias magnetic and electric fields	14
2.1.4	Key ingredient 3: Ω doubling (aka clever use of molecular energy structure)	16
2.2	Measuring the electron's electric dipole moment: theoretical details	20
2.2.1	$^3\Delta_1, v = 0, J = 1$ Hamiltonian	20
2.2.2	Measurement of $\delta g_F/g_F$	22
2.2.3	Molecular structure	22
2.2.4	Experimental noise vs the quantum projection noise limit	23
2.3	Experimental apparatus	28
2.3.1	Overview	28

2.3.2	Source chamber	28
2.3.3	Ion trap	30
2.3.4	Computer control	31
2.3.5	Applied magnetic fields	31
2.3.6	Applied electric fields	34
2.3.7	Lasers	37
2.3.8	Chops, switch states, blocks, and superblocks	38
2.3.9	Trap characterization	40
2.3.10	Experimental sequence	43
2.4	Data analysis	50
2.4.1	Hf counting and imaging	50
2.4.2	Spin asymmetry	52
2.4.3	Fitting	54
2.4.4	Frequency channels	55
2.4.5	Blinding	58
2.4.6	Data cuts	59
3	Accuracy evaluation	60
3.1	Accuracy evaluation of the HfF^+ electron electric dipole moment experiment	60
3.2	Introduction	60
3.3	Experiment	61
3.3.1	Lasers	64
3.3.2	Molecular beam and ionisation	65
3.3.3	Ion trap	65
3.3.4	Magnetic fields	66
3.3.5	State preparation	67
3.3.6	Ramsey sequence	69

3.3.7	Measurement	69
3.3.8	Noise	71
3.3.9	Experimental protocol and switch states	72
3.3.10	Images to determine doublet positions	74
3.4	Data Analysis	74
3.4.1	Ion counting and asymmetry	74
3.4.2	Least squares fitting	76
3.4.3	Switch parity channels	77
3.4.4	Blinding	78
3.4.5	Data cuts	78
3.5	Effective Hamiltonian for doublets	78
3.6	Accuracy evaluation	81
3.7	Frequency shifts	82
3.7.1	Magnetic effects	82
3.7.2	Berry's phase effects	95
3.7.3	Residual rotation induced doublet mixing	107
3.7.4	Oddities & miscellany	109
3.8	Phase Shifts	115
3.8.1	State preparation	115
3.8.2	Internal state measurement	120
3.9	Interpretation	130
3.10	Supplemental Information / Appendices to the paper	130
3.10.1	Overview of data	130
3.10.2	Overall comments	131
3.10.3	Frequencies	134
3.10.4	Contrasts, and contrast decays	135
3.10.5	Phases	136

3.10.6	Offsets	137
3.11	Not Systematics	137
3.11.1	Shifts vs z position	137
3.11.2	Ion number vs f^0	146
3.11.3	f^R	147
3.11.4	Shifts vs x-y position	147
4	A new limit on the electron's electric dipole moment	149
4.1	Summary of the eEDM result	149
4.2	Experimental Overview	151
4.3	Accuracy evaluation	154
4.4	Results	156
4.5	Conclusions	158
4.6	Methods	159
4.6.1	Bias fields	159
4.6.2	HfF ⁺³ Δ_1 details	159
4.6.3	Quantum state preparation	160
4.6.4	Ramsey spectroscopy	160
4.6.5	Final state readout	161
4.6.6	d_e signal extraction	162
5	Dark Matter Exclusion	164
5.1	Introduction	164
5.2	Experimental Constraint on Axionlike Particles over Seven Orders of Magnitude in Mass	168
5.3	Experimental Constraint on Axionlike Particles over Seven Orders of Magnitude in Mass - Supplemental Material	179
5.3.1	Frequency Ranges	179

5.3.2	Generating Least-Squares Spectra	181
5.3.3	Timestamp Shuffling	183
5.3.4	Estimating the no-ALP distribution	186
5.3.5	Sources of Attenuation	187
5.3.6	Parameterizing the ALP-gluon coupling	190
5.3.7	Aggregate Prior Update	191
6	Conclusion	194
	Bibliography	198
	Appendix	
A	Lasers	206
A.1	Neutral HfF production / ablation: "Abby" 532 nm	206
A.2	Photionization 1+2 pump: "Mia" 532 nm	207
A.3	Photoionization 1 seed: "Tuna" 1237.55598 nm	208
A.4	Photoionization 1 (first photon): "Harambe" 309.388 nm	209
A.5	Photoionization 2 (second photon): "Bertha" 735.464 nm	210
A.6	State transfer: "Toptica" 961.43495 nm	211
A.7	Cleanup: "Vibational Cleanup" 818.372 nm	212
A.8	mF pumping: "DBR/1083" 1082.4137 nm	213
A.9	Depletion: "814/Depletion" 814.508 nm	213
A.10	Photodissociation 1 (first photon): "Cora" 736.989/2 nm or 736.986/2 nm	214
A.11	Photodissociation 2 (second photon) + Cora pump: "Olga" 266 nm	215
B	One shot asymmetry	216
B.0.1	One-shot asymmetry and measured fringe contrast	216

C	A list of useful transitions	219
D	An approximation to our shot noise limit	221
E	All parity channel results	226

Tables

Table

2.1	Magnetic field coil parameters	36
2.2	Photons used in our experimental sequence	37
3.1	Constants used throughout this document. The total magnetic g-factor of $J = 1, F = 3/2$ states $g_F \equiv (G_{\parallel} + g_N \mu_N / \mu_B) / 3$ results from the combination of nuclear and electronic Zeeman effects.	62
3.2	Example experimental parameters and associated derived parameters from our 2022 data.	62
3.3	Photons used in our experimental sequence	64
3.4	Magnetic-field expansion	85
3.5	Constraints on electric fields from higher-harmonic voltages on radial electrodes and endcaps.	104
3.6	Constraints on Berry's phase frequency contributions to f^B or f^{BR} due to possible combinations of radial and axial field imperfections from harmonic distortion in amplifiers used to drive radial electrodes. All entries are in μHz	106
3.7	Systematic error budget	129
4.1	Systematic error budget	155
A.1	Minilite II specs from manual	207

A.2	Quanta-Ray-Pro-250-10 specs from manual	207
A.3	A few Tuna (Toptica TA Pro) specs from the data sheet	208
A.4	Select Harambe (Radiant Narrow Amp) specs	210
A.5	Select Harambe (Radiant Narrow Amp) specs	211
A.6	A few Toptica/Transfer (Toptica TA Pro) specs	211
A.7	A few Vibrational Cleanup (Homebuilt ECDL) specs	212
A.8	A few Vibrational Cleanup (Homebuilt ECDL) specs	213
A.9	A few Depletion (Toptica TA Pro) specs	214
A.10	Select Cora (Syrah Cobra Stretch) specs	215
A.11	Quanta-Ray-Lab-170-10H specs from manual. Note the manual specified the pulse width in nm but I'm pretty sure that's a typo, especially given footnote 1 on table 3-2.	215
E.1	All contrast parity channels, disaggregated over distinct f^0 values and aggregated over all 3 f^0 values.	227
E.2	All decay rate parity channels, disaggregated over distinct f^0 values and aggre- gated over all 3 f^0 values.	228
E.3	All frequency parity channels, disaggregated over distinct f^0 values and aggre- gated over all 3 f^0 values.	229
E.4	All phase parity channels, disaggregated over distinct f^0 values and aggregated over all 3 f^0 values.	230
E.5	All offset parity channels, disaggregated over distinct f^0 values and aggregated over all 3 f^0 values.	231
E.6	All frequency parity channels (but with 3 digits after the decimal), disaggregated over distinct f^0 values and aggregated over all 3 f^0 values.	232

Figures

Figure

1.1	A comic I drew about the baryon asymmetry	2
1.2	A cartoon depicting the time-reversal symmetry violation of an electron electric dipole moment. The spin is indicated with two arrows and the electric dipole moment is represented by a colour gradient and positive and negative ‘poles’. When one flips the arrow of time, the spin flips but the electric dipole moment does not.	3
1.3	A comparison between fringes taken in 2013 (upper left) and much more recently, in 2020 (remaining panels). Note the time axis goes to 25 ms in 2013, while it goes to 1.5 seconds in 2020.	8
1.4	A cartoon depicting the current status of eEDM searches. The pink region indicates where most new physics theories predict the eEDM should be, and the blue arrows indicate recent exclusions (at the 90% confidence level). The exclusions measured zero up to the uncertainty indicated, which means that they exclude eEDMs larger than their uncertainty. The brown region indicates the standard model prediction for the eEDM. One can see that the region we are probing is effectively ‘background free’ in terms of the standard model prediction, and that most new physics theories predict EDMs right in the range we can measure - so our measurements directly constrain those theories.	10

- 2.1 Conceptual cartoon depicting the basic measurement principle. In each panel, the ket indicates both the direction of the magnetic moment (spin) and the electric dipole moment (EDM). In the first panel, a simple spin flip measurement is performed. In the second panel, an electric field is added, pushing the energy levels apart for that particular relative orientation of the fields. In the third panel the electric field is flipped, pushing the levels together. The difference between the last two measurements cancels out the magnetic dipole contribution (to first order) leaving only the electric dipole contribution. 12
- 2.2 Conceptual cartoon depicting how three key parameters affect our sensitivity. In the first panel, the separation between peaks scales linearly with the effective electric field, E_{eff} . In the second panel, having a higher count rate N reduces the fractional error in the determination of the peak centers by $\sqrt{1/N}$. In the third panel, the width of the peaks scales like the inverse coherence time, or $1/\tau$ 13
- 2.3 A cartoon depicting a polarized HfF^+ molecule 13
- 2.4 A cartoon depicting how a radial quadrupole magnetic field amounts to a static, quantization-axis aligned magnetic field in the rotating frame. In each panel, the magnetic field gradient is indicated with light blue arrows. A circular orbit for a molecule is indicated near the top, and pink arrows depict the vector decomposition of the magnetic field which the ions see as they are rotating. You can see that as the molecule rotates, the vector decomposition of the magnetic field has two components: one which is always pointing "up", and one which always aligns (or anti-aligns) with the molecular axis. 15

2.5 A cartoon depicting the vectors for Hund's case a. The electronic spin and orbital angular momenta, \mathbf{S} and \mathbf{L} , are strongly coupled to the internuclear axis by spin-orbit and electrostatic interactions respectively. Hund's case a arises when spin-orbit interactions are much weaker than the electrostatic interactions which split different electronic states. This means that widely separated electronic levels are split into a set of spin-orbit manifolds, each of which contains many closely spaced rotational levels. $\Lambda = \mathbf{L} \cdot \mathbf{n}$ is the projection of the electronic orbital angular momentum \mathbf{L} on the internuclear axis \mathbf{n} , $S = |\mathbf{S}|$ is the total electronic spin angular momentum, $\Sigma = \mathbf{S} \cdot \mathbf{n}$ is the electron spin projection on the internuclear axis, $J = |\mathbf{J}| = |\mathbf{L} + \mathbf{S} + \mathbf{R}|$ is the electronic plus rotational angular momentum, $\Omega = \Lambda + \Sigma = \mathbf{J} \cdot \mathbf{n}$ is the projection of the electronic angular momentum onto the internuclear axis, $I = |\mathbf{I}|$ is the ^{19}F nuclear spin, $F = |\mathbf{F}| = |\mathbf{J} + \mathbf{I}|$ is the total angular momentum of the molecule, and $m_F = \mathbf{F} \cdot \hat{z}$ is the projection of \mathbf{F} on the rotating quantization axis [1] [2]. We take the internuclear axis \mathbf{n} to be directed from the ^{19}F nucleus to the ^{180}Hf nucleus [2]. 17

2.6 In $\text{HfF}^+ \text{s } ^3\Delta_1 J = 1$ state we have two hyperfine states with total angular momentum $F = 3/2$ and $F = 1/2$. Each of these has two degenerate manifolds made up of the different allowed projections $m_F = \mathbf{F} \cdot \hat{z}$ onto the quantization axis. The two sets of degenerate states correspond to the eigenstates of the parity operator: $|e\rangle = (|\Omega = +1\rangle - |\Omega = -1\rangle)/\sqrt{2}$ and $|f\rangle = (|\Omega = +1\rangle + |\Omega = -1\rangle)/\sqrt{2}$. These degenerate manifolds are separated by the Ω -doublet splitting, which we call ω_{ef} . 18

- 2.7 Focusing only on the $F=3/2$ manifold of the ${}^3\Delta_1$ $J = 1$ state of HfF^+ . $\Sigma = \mathbf{S} \cdot \mathbf{n}$ is the electron spin projection on the internuclear axis, and the internal electric field experienced by the valence electron is indicated as \vec{E}_{eff} . Left: an electric field polarizes the molecule along or against the field, which we call the “upper” and “lower” “doublets”. The energy levels without the added E field are indicated with dashed lines, and the energy levels with the added E field are indicated with solid lines. Center: adding a B field further lifts the degeneracy, when the spin is aligned with the B field the level moves down, and when it is anti-aligned it moves up. Again, the energy levels without the added B field are indicated with dashed lines and the energy levels with the added B field are indicated with solid lines. Right: finally, the addition of the eEDM itself further breaks the degeneracy, we assume the eEDM points in the same direction as the spin; when the eEDM is aligned with the internal electric field \vec{E}_{eff} the energy level moves further down and vice versa. As before, the energy levels without the effect of the eEDM are indicated with dashed lines and the energy levels with the effect of the eEDM are indicated with solid lines. Nothing in this cartoon is to scale. 19
- 2.8 HfF^+ 's ${}^3\Delta_1, J = 1, F = 3/2$ state complete details. $\Sigma = \mathbf{S} \cdot \mathbf{n}$ is the electron spin projection on the internuclear axis, and the internal electric field experienced by the valence electron is indicated as \vec{E}_{eff} . The energy levels without the effect of the eEDM are indicated with grey lines, while the ones with the effect of the eEDM are indicated with black lines. Dashed lines are for visual aid only. The effect of the eEDM has the opposite sign in the upper vs lower doublet. It is this energy difference that we aim to measure with our experiment. 20
- 2.9 Measurement of $\delta g_F/g_F$ by varying the mean fringe frequency f^0 , all data is taken at $\mathcal{E}_{\text{rot}} = 57.9\text{V/cm}$ 23
- 2.10 Measurement of $\delta g_F/g_F$ by adding a non-reversing axial magnetic field gradient, leading to large f^B . Note that f^{BD} is still blinded in this measurement. 23

2.11	CAD rendering of the entire apparatus, excluding lasers. In purple are the X shim coils, blue are the Y shim coils, and yellow are the Z shim coils. The orange coils generate the axial magnetic gradient that we call \mathcal{B}_{rot}	29
2.12	Top down schematic of our experiment's main chamber	29
2.13	CAD rendering of the main experimental chamber sliced in half for a view of the ion trap. To give a sense of scale, you can fit a coffee cup between the electrodes.	31
2.14	Confining and orienting field amplitudes throughout the experimental sequence, in Volts/cm, for a 5 ms Ramsey free-evolution time. The $\pi/2$ pulses are easily visible in the \mathcal{E}_{rot} panel.	32
2.15	Confining and orienting field amplitudes throughout the experimental sequence, in Volts, for a 5 ms Ramsey free-evolution time. The $\pi/2$ pulses are easily visible in the \mathcal{E}_{rot} panel.	33
2.16	The variation of the computed secular frequencies throughout the experimental sequence for a 5 ms Ramsey free-evolution time.	34
2.17	From \mathcal{E}_{rot} , computed values of \mathcal{B}_{rot} and r_{rot} throughout the experimental sequence for a 5 ms Ramsey free-evolution time.	35
2.18	Cartoon depicting the transitions used during our state preparation.	37
2.19	Measurement of the 'ambient' f^B	40
2.20	Measurement of the 'ambient' f^B	41
2.21	The time of flight trace of the HfF^+ cloud with a Gaussian overlaid	41
2.22	Using our HfF/voltage calibration, we compute the HfF in each slice of the TOF trace	42
2.23	Constructing the approximate HfF ellipsoid from our imaging and TOF data. In the Y-Z imaging data, the lines demarcate the approximate edge of the cloud.	42
2.24	Approximate HfF density along X, and approximate Mean Field Parameter for our cloud over a given temperature range	43

2.25	Some typical images showing how the two doublets appear on our phosphor screen: on the left, we depleted one doublet and imaged the one remaining. In the center we did the same with the opposite doublet, and on the right we depleted neither (so we are imaging both doublets at the same time). The numbers indicate pixels on the screen, the images are in false colour, and the images are an average of 20 separate runs of each experiment. The red line indicates our programmatic determination of the 'center' of the two blobs.	48
2.26	An overlay of 20 images from our MCP assembly, where each doublet was fully depleted 10 times. The center is indicated with a black dashed line.	51
2.27	The integrated trace of Figure 2.26, showing the upper and lower doublets separately.	51
2.28	Example Ramsey interferometry data	53
3.1	CAD rendering of the entire apparatus, excluding lasers. On the left is the source chamber, where we produce neutral molecules. To the right of the source chamber is the differential pumping chamber, where the beam of neutral molecules is supersonically cooled to a rotational temperature of roughly $\sim 10\text{K}$ and collimated by skimmers. Further right is the main experimental chamber, which contains our 8-electrode Paul trap. Surrounding the main chamber are the magnetic field coils: the X shim coils in purple, the Y shim coils in blue, and the Z shim coils in yellow; the coils which generate the quadrupole magnetic field \vec{B} are indicated in orange.	63
3.2	Top down schematic of our experiment's main chamber. The neutral molecules fly into the center of the trap where they are ionized and trapped. We apply the rotating electric bias field \vec{E}_{rot} and the quadrupole magnetic field \vec{B} . The ions undergo uniform circular motion, either aligned or anti-aligned with \vec{E}_{rot} , and their unpaired spins precess about $\vec{E}_{\text{rot}} \cdot \vec{B}$	63
3.3	Cartoon depicting the transitions used during our state preparation.	67

3.4	Ramsey spectroscopy in HfF^+ . Top, level structure in the eEDM-sensitive $^3\Delta_1$, $J = 1$, $F = 3/2$ manifold in external electric \mathcal{E}_{rot} and magnetic \mathcal{B}_{rot} bias fields (not to scale). The oriented molecules are either aligned or anti-aligned with \mathcal{E}_{rot} , which we call the lower and upper Stark doublets respectively. The component of the energy splitting due to the eEDM, $2d_e\mathcal{E}_{\text{eff}}$, flips sign between doublets. Bottom, example Ramsey fringes from our dataset. The fringes from the two doublets are collected simultaneously.	70
3.5	Results as we shift the threshold for applying cuts to the data. In each plot, only one cut is applied (the other cut is not applied). The cuts do not shift the resultant value of f^{BD}	79
3.6	Measurement of $\frac{\delta g_F}{g_F}$, by varying the mean fringe frequency f^0 and measuring the differential fringe frequency f^D . Data taken at $\mathcal{E}_{\text{rot}} \sim 50$ V/cm.	87
3.7	Measurement of $\frac{\delta g_F}{g_F}$, obtained by introducing a large $\delta\mathcal{B}_{nr}$ and measuring the resultant values of f^{BD} and f^B . Note that f^{BD} is still blinded in this measurement. Data taken at $\mathcal{E}_{\text{rot}} \sim 50$ V/cm.	88
3.8	Example axial oscillation-induced phase modulation data, for a both large kick (color) and a small kick (grey).	99
3.9	Phase modulation amplitude for various sizes of z kick. Upper doublet data is in blue and lower doublet data is in pink. Both fits have an intercept of $-0.0127(0.0108)$	99
3.10	Example data collected with $\mathcal{E}_{2fZ} \sim 4$ V m $^{-1}$ and an ellipticity $\mathcal{E}_e/\mathcal{E}_{\text{rot}} \sim 5 \times 10^{-3}$. As we vary the phase of the axial modulation, f^B and f^{BR} vary sinusoidally with a $\pi/2$ phase offset, as expected.	101
3.11	Toy model to explain ϕ^{DBI} . (i) uncoupled basis, (ii) coupled basis.	116
3.12	Example data. Plotting the slope of f^D when the early-time data is taken at varying offsets from the zero crossing of the difference fringe. The x -intercept gives us the true value of our imaging contrast, which we then use in our fitting program.	122

3.13	Varying the swatch width while correcting for the doublet contrast. We see no meaningful shifts in our science channel and a small residual shift in f^D . In the lower panel we have averaged over the 3 distinct values of f^D that we operated at.	124
3.14	The effect of moving the center of the swatch on our science channel.	126
3.15	Moving the center of the swatch has a measurable effect in several parity channels. Note all panels show the entire aggregated dataset except for the f^D panel, which shows a subset of measurements performed at a single f^0	127
3.16	Imaging data collated over 25 blocks. The lower doublet data is in pink while the upper doublet data is in blue.	128
3.17	129
3.18	Observed frequency shifts in z	138
3.19	Observed frequency shifts in z	138
3.20	Fitting $f(z)$	139
3.21	Comparing frequency shifts in z with the magnetic field shimming Z coil on or off	140
3.22	Comparing frequency shifts in z with the magnetic field shimming Z coil on or off	140
3.23	Comparing observed frequency shifts in z with 2 different values for the axial confinement	141
3.24	Comparing observed frequency shifts in z with 2 different values for the axial confinement	141
3.25	Fitting the measured frequency vs slosh amplitude in the non-displaced trap	145
3.26	Comparing the prediction for f^{BR} from our model of the anharmonicity to the empirical results. We see that this model does not account for the observed shifts	146
3.27	Averaging f^R over our entire dataset, per primary fringe frequency f^0 . f^R scales with f^0 with a slope of 0.00018, which is very close to $1/1024 = 0.00098$ or a single step in a 10-bit DAC.	147
3.28	Comparing observed frequency shifts with x translation	148
3.29	Comparing observed frequency shifts with x translation	148

- 4.1 Ramsey spectroscopy in HfF^+ . Top, level structure in the eEDM-sensitive $^3\Delta_1$, $J = 1$, $F = 3/2$ manifold in external electric \mathcal{E}_{rot} and magnetic \mathcal{B}_{rot} bias fields (not to scale). The oriented molecules are either aligned or anti-aligned with \mathcal{E}_{rot} , which we call the lower and upper Stark doublets respectively. The component of the energy splitting due to the eEDM, $2d_e\mathcal{E}_{\text{eff}}$, flips sign between doublets. Bottom, example Ramsey fringes from our dataset. The fringes from the two doublets are collected simultaneously. 152
- 4.2 Top down schematic of our experiment's main chamber. The neutral ions fly into the center of the trap where they are ionized and trapped. We apply the rotating electric bias field $\vec{\mathcal{E}}_{\text{rot}}$ and the quadrupole magnetic field gradient $\vec{\mathcal{B}}$. The ions undergo uniform circular motion, either aligned or anti-aligned with $\vec{\mathcal{E}}_{\text{rot}}$, and their unpaired spins precess about $\vec{\mathcal{E}}_{\text{rot}} \cdot \vec{\mathcal{B}}$ 152
- 4.3 Systematic shift in the eEDM channel f^{BD} due to non-reversing axial magnetic field, which shows up in the f^B channel. 156
- 4.4 Results as we change the threshold for cutting blocks from our dataset. In each plot, only one cut is applied (the other cut is not applied). Applying cuts does not systematically shift our result. 156
- 4.5 Summary of our dataset after cuts have been applied and scaling δf by $\sqrt{\chi^2}$ to account for over-scatter. Top left, normal probability plot of f^{BD} showing the data are consistent with a normal distribution. Bottom left, histogram of normalized and centered f^{BD} data. Overlaid in blue is a normal distribution. Right, superblocks taken under different experimental parameters compared to the overall average value of f^{BD} . Here N is the average number of trapped HfF^+ ions per experimental trail. 157
- 4.6 Cartoon depicting the transitions used during our state preparation. 161
- 5.1 Generation 1 EDM data vs time. Could there be an oscillatory signal buried in here? 164

5.2	Least-squares spectral analysis of our EDM data.	172
5.3	Bayesian power measured analysis of the EDM data. The grayscale matrix corresponds to a sub-aggregation of the prior updates U over 100 logarithmically-spaced frequency bins. Darker shading on the logarithmic color scale corresponds to increased prior update, which can be interpreted as an increased belief in an ALP of given frequency ν at ALP-gluon coupling $C_G/(f_a m_a)$. We use a blue line to indicate the 95% exclusion (which corresponds to the sub-aggregated update dropping to 0.05 [3]) for sub-aggregation over each decade. The corresponding greyscale matrix for this sub-aggregation is not shown. Finally, a dotted line corresponds to the aggregated exclusion over the full analysis range. The full un-aggregated matrix of U values is available upon request.	176
5.4	Astrophysical and laboratory limits on the ALP-gluon interaction. All limits assume that ALPs saturate the local dark matter content and we have taken the mean local dark matter density to be $0.4 \text{ GeV}/\text{cm}^3$ [4]. All constraints correspond to a 95% confidence level. The purple region embodies the constraint that the virial de-Broglie wavelength is smaller than the size of a dwarf galaxy [5]. The blue region is the constraint derived from looking at neutron EDM oscillation [6]. The green region represents constraints from big bang nucleosynthesis [7], and the yellow region represents constraints from supernova energy loss bounds [8, 9]. The pink region is the constraint described in this work.	178
5.5	(Top) Individual C_5 measurements plotted as a function of their acquisition time since August 16 2016 (the temporal ‘center’ of our dataset). (Bottom) Our most sensitive C_5 measurements were collected over 11 days in 2017. For the low frequency analysis we used the entire dataset (top), and for the high frequency analysis we used only the most sensitive 11 days (bottom).	180

5.6	Attenuation factor due to ALP field decoherence, calculated over the high frequency analysis range. Grey-filled circles are the calculated values, and the blue curve is the fit.	189
5.7	Attenuation due to finite sample time over the entire analysis range. The blue curve shows the piecewise function generated by the fit to the grey circles.	191
5.8	Aggregate prior update \mathcal{U} (equivalently, exclusion) as a function of ALP-gluon coupling $C_G/(f_a m_a)$. The horizontal dotted line indicates 95% exclusion (where the aggregate prior update drops to 0.05), while the vertical dotted line indicates the lower limit of $C_G/(f_a m_a)$ which is excluded at 95% confidence over the entire analysis band. Had our data presented us with strong affirmative evidence for the existence of an ALP in our frequency range, the curve of \mathcal{U} vs $C_G/(f_a m_a)$ might have peaked not at $\mathcal{U} = 2.5$ but for instance at $\mathcal{U} = 1,000$ or higher.	193
C.1	The infamous 'laser diagram'	220

Chapter 1

Introduction

1.1 Motivation

It's honestly hilariously ambitious of physicists to imagine that the entire universe - from the dawn of time until its inevitable heat death, from our galactic core out to infinity - can be described with a single mathematical equation. But that does seem to be the project for a sizeable portion of the discipline, and this thesis is one very small contribution to that project. Physicists have agreed for now that the language of our universe is mathematics, so the goal is to find the Grand Unified Theory of Everything, which is funnier if you read it in Douglas Adam's voice. In practice, this means taking a fairly hefty equation, like the Standard Model (SM), and testing it and refining it and testing it and refining it and so on until we have pulled the blindfolds off our eyes and can stare into the face of god as we walk backwards into hell¹.

The Standard Model, while a darling in most physics circles, is problematic². For example, the SM does not account for dark matter or dark energy³, the unification of the SM with general relativity remains elusive, and there are other issues with opaque names like 'the hierarchy problem' [11]. Our experiment is motivated by a problem called the 'baryon asymmetry', which

¹@dril wrote: "IF THE ZOO BANS ME FOR HOLLERING AT THE ANIMALS I WILL FACE GOD AND WALK BACKWARDS INTO HELL" Twitter, 2012-05-22

²The Standard Model: our problematic fave.

³How embarrassing - for example, dark matter makes up 84% of our universe's cosmic matter density [10]

describes the preponderance of matter (baryons) over antimatter (antibaryons) in our universe. This asymmetry explains why we actually *exist*, but this too is a problem in the SM. The SM predicts that our universe would have been born with virtually equal parts matter and antimatter, who would have met and (excitingly! explosively!) annihilated a long time ago. The result would be a present-day universe made entirely of light⁴, which sounds poetic - but a universe without matter would be terribly boring. Our very existence, then, poses a problem for the "let's describe literally everything with an equation" people.

The baryon asymmetry can be explained by modifying the SM with the addition of some extra 'asymmetry' in the math, in particular some time-reversal (T) violation (and by the CPT theorem, charge-parity (CP) violation). This was all worked out by a physicist named Andrei Sakharov [12], whose papers I do not understand in the least. For a slightly more technical explanation of all this I'd recommend you read the introduction to Will Cairncross' thesis [13], and for an even less technical explanation please enjoy this comic I drew (Figure 1.1).




Figure 1.1: A comic I drew about the baryon asymmetry

The tricky bit, as far as I understand it, is that we can't just add some T violation willy-nilly into the standard model, because there's a multitude of so-called *sectors* it can be added to and we need to figure out which one is the right one. Put a different way, there's a lot of terms in the SM which could be tweaked to accommodate the extra symmetry violation, and we need to figure out which terms are the right ones to adjust. Our measurement aims to check the *leptonic*

⁴no stuff, just vibes




Figure 1.2: A cartoon depicting the time-reversal symmetry violation of an electron electric dipole moment. The spin is indicated with two arrows and the electric dipole moment is represented by a colour gradient and positive and negative ‘poles’. When one flips the arrow of time, the spin flips but the electric dipole moment does not.

sector by measuring the permanent electric dipole moment (EDM) of the electron (eEDM). An eEDM violates time-reversal symmetry because flipping the arrow of time flips the spin but not the EDM (see Figure 1.2).

Finally, we can state the goal of this thesis: to measure the EDM of the electron (eEDM), if it exists. We will cover the details of how one goes about measuring the eEDM in the following chapters. For now, dear reader, please understand that the rest of this thesis is really about doing a spectroscopic experiment (with a quick detour into a dark matter search - see chapter 5), and not quantum field theory or complicated calculations.

There is a second motivation for our particular experiment. The addition of more T violation in the SM is, by some strange magic I don’t understand, the same as adding new high energy particles to the SM. Because of this, sorting out the details of the T violation is related to detecting new high-energy particles and measuring their properties. Traditionally this is what particle accelerators are built for - but the new particles we are searching for are beyond the reach of any existing accelerator, and (barring some incredible innovation) this fact is unlikely to change for decades to come. Because of this, we really need to look outside the accelerator paradigm to continue to search for new physics. Luckily for the field, precision AMO (atomic and molecular optics) experiments can do just that: our current experiment is probing new physics at a TeV

range which is already beyond the reach of the Large Hadron Collider (LHC) [32].

1.2 JILA's eEDM experiment(s): a brief history

Before Eric started this experiment, nobody had done an eEDM search with trapped ions. To this day we are still the only trapped ion eEDM experiment in the whole world. The relative *uniqueness* of this experiment is important to acknowledge off the top, because it means that this experiment was quite difficult to realize from scratch. Eric describes the early days of thinking and planning⁵:

"There were several years there when I was thinking of how do a really good AMO-meets-particle-physics experiment. I had BEC⁶, which seemed like a good starting point. Then (as now) I was conscious of hailing from a line of AMO or AMO-adjacent physicists who were interested in that sort of thing. I was in Florence, Italy, in 2000 and went to an ICAP where Dave Demille had a poster on his PbO, vapor cell, $\Omega = 1$ vision. That was my first exposure to the really cool possibilities of $\Omega = 1$ molecular states for precision metrology, and it blew my mind. I'm glad that four years went by before I definitively set out to steal that idea. And I'm glad that everything about my experiment except for $\Omega = 1$ was entirely different from Dave's EDM vapor cell project. It can be a little rude (although not "illegal") to go to a talk or poster and then just say "cool experiment. I am going to try to do it too, just do it faster." In the 2000s through summer 2003 I was thinking a lot about "what I could do for an encore to BEC". Starting in fall 2002 with Immanuel Bloch's paper on the Mott transition [14] which basically set everyone's hair on fire, and increasingly for the next couple years, this idea of doing quantum computing with individual neutral atoms in adjacent lattice sites was like all anyone wanted to talk about. The whole scheme didn't appeal to me, or maybe it was just sour grapes that it was not my idea, so that pushed me to think even more about precision metrology as an alternative to quantum information, and as a future direction for cool BEC experiments. Meanwhile, Debbie just down the hall was doing these mind-blowing things with fermions, so I could sort of enjoy cutting-edge ultra-cold physics vicariously through watching her progress. Before too long I was so set on "doing a cool precision measurement" that I was no longer committed to "doing precision measurement using BEC". I was pretty sure I didn't want to do a beam-line experiment, ala Rabi then Ramsey then Kleppner then Commins. I wanted long coherence times, so trapped or at least hobbled atoms or molecules. I spent time reading journal articles on all sorts of crazy ideas, e.g. "trapping" or at least immersing molecules in liquid

⁵Extensive quotes from E. Cornell extracted from a February 2022 email

⁶Bose Einstein Condensate

helium. I looked into exotic traps like acoustic-wave Paul traps and microwave cavity traps."

At this point it's helpful to point out that a key step in any EDM experiment is the addition of a very large electric field to polarize the species under investigation. And since Eric wanted to do the experiment with molecules to realize the " $\Omega = 1$ vision", but he also wanted long coherence/interrogation times, he needed to figure out how to both trap *and* electrically polarize molecules. At the time nobody had developed the technology to trap EDM-friendly *neutral* molecules. In fact, considerable effort is still being expended on getting EDM-friendly neutrals trapped, and the best neutral limits still come from beam experiments (which are severely limited in interrogation time). Of course, the technology for trapping *charged* species was very mature, but using trapped ions for an EDM experiment was widely considered impossible: in a normal ion trap the application of a polarizing DC electric field would immediately kick the ions right out of the trap. But here's where Eric had a fantastic idea: use a *rotating* electric field to polarize the molecules without kicking them out of the trap. Later chapters will discuss this in detail.

"The idea of trap with rotating field, I am not sure exactly when I came up with it but it was much on my mind in summer of 2003, and by fall 2003 I decided I wanted to try it. I had like 10 different molecules in mind, but HfF^+ was not one of them⁷. The idea of using photodissociation as read out technique was there from early on as well. I started trying to hire my first student in late fall of 2003. Russ (Russel Stutz) joined in January 2004."

The first recorded proposal for the experiment that I could find comes from 2004, from the 35th Annual Meeting of the APS AMO Division [15]. The abstract reads:

Search for the electron EDM using trapped molecular ions
The current limit on the electron electric dipole moment ($d_e < 1.6 \times 10^{-27} e\cdot\text{cm}$) was set using an atomic beam of Tl [16]. We propose a new experiment using molecular ions trapped in an RF quadrupole and cooled to cryogenic temperatures using a Helium buffer gas. This new experiment would benefit from the large effective electric fields experienced by an electron in polar molecules which can be orders of magnitude larger than in heavy atoms [17]. We also hope to

⁷Ultimately, HfF^+ was the molecule we used in the experiment

achieve much longer spin coherent times than previous experiments, resulting in narrow magnetic resonance lines. Utilizing the near degenerate Λ -doublet present in some molecular ions, we expect to achieve a highly polarized sample of ions in a relatively weak (10 's of V/cm) rotating electric field. This will lead to high sensitivities to de while minimizing systematic effects due to laboratory electric fields. Systematic effects should also be reduced by the low ion velocities compared to atomic beam experiments.

It's worth pausing for a moment here to give Russ the credit he deserves for being the very first grad student to sign on to a never-been-done-before experiment as a new grad student. Eric recalled a time when he gave a talk about the proposed experiment at NIST, with Russ in tow, and afterwards one of the NIST scientists said, with Russ listening, "cool idea, but I'd hate to be a grad student working on that project!"

By the spring of 2004 Laura Sinclair had also signed up as a grad student on the eEDM project, and Aaron Leanhardt had agreed to be a postdoc. All 4 of them met at DAMOP in 2004 for the first group eEDM meeting, even though half the group hadn't even been to Boulder yet! Around this time, Eugene Commins' experiment was the one to beat. He had published the leading paper on eEDM in 2002 [16] using an atomic beam of Thallium, and this limit stood for about a decade before being bested in 2011 by the first eEDM experiment done with a paramagnetic polar molecule: the YbF experiment from the Hinds group at Imperial College London [18].

By Fall of 2004 Russ, Laura, and Aaron were all in the lab, and John Bohn and his student Ed Meyer were thinking hard about the experiment in general and molecular species in particular. But just as things were getting off the ground, Eric got very, very sick with necrotizing fasciitis. In his words:

"It's hard to describe how very very sick I was, how close to death, for two months, and how for at least a full 12 months I was either not around at all or else around, but physically and emotionally kind of shattered, needing to work on physical therapy, pain issues, and something at least vaguely like ptsd. I had definitely lost my swagger. In the meantime, Russ and co. plodded on, and by around February 2005 Eric realized that "some combination of Aaron, John and

Ed had identified a really promising molecule, HfF^+ . I have no idea how they did that!!!!"

Obviously, settling on a species was a huge milestone for the experiment. After this the team needed to figure out how to make it, trap it, map its energy landscape, control the internal quantum state, do Ramsey spectroscopy, and read out the final quantum state. A tall order indeed! Eric describes the exploratory work that followed:

"We wanted to do "practice photodissociation" on a molecule whose photodissociation was well understood, and we wasted like four person-years on that effort. Also it took us too long to figure out we really could not do survey spectroscopy in our ion trap. We needed a neutral beam line, and we simply had to bite the bullet and learn to do our own survey spectroscopy. Enter Jun Ye!!!!"

Though the group did not publish a lot at this time, a metric tonne of work got done. They figured out how to make and trap HfF^+ , they did spectroscopy and saw their first of many assignable rotation bands, they found excited states which coupled strongly to both $^1\Sigma^+$ and $^3\Delta_1$, they measured the spectroscopic signature for Δ^8 with no applied B field, and of course they finally saw Ramsey fringes. A big deal for our experiment was when the team figured out rotational-state-selective resonance-enhanced-multiphoton-dissociation (REMPD), which was eventually published in 2014 [19]. This was so critical because the efficiency of the other readout option, light induced fluorescence (LIF), was so low that it was clear that they would never get a high precision result with it. Eric said it was the "much better than expected REMPD" that made him think "yes. We will by golly be able to beat Commins by a factor of ten. We may not do that before ACME does, but we will eventually. This whole idea IS NOT CRAZY!!!!"

Eventually we published our first generation result [2], and we managed to almost tie ACME's world-leading result [20]. It was pretty amazing given we were counting about 5 ions on the side of a fringe at the time (our entire advantage was that we were measuring out to $\sim 750\text{ms}$ compared to ACME's $\sim 5\text{ms}$). After that we tore down the first generation apparatus

⁸Defined in Chapter 2

and rebuilt everything while we re-invented almost every step of the experiment. We now count ~ 250 ions on the side of a fringe and we measure out to 1.5 – 2.5 seconds, allowing us to put a new limit on the eEDM.

Looking back at all the theses and papers published over the last decade and a half it is remarkable what a labour of love our eEDM experiment is. I owe a huge debt of gratitude to all the grad students, undergrads, post-docs, and PI's who worked incredibly hard to made this not-so-crazy idea actually come to life. I am exceedingly lucky to have walked into the experiment around the time the first measurement was happening, and that I've been able to see the second measurement through from beginning to end.

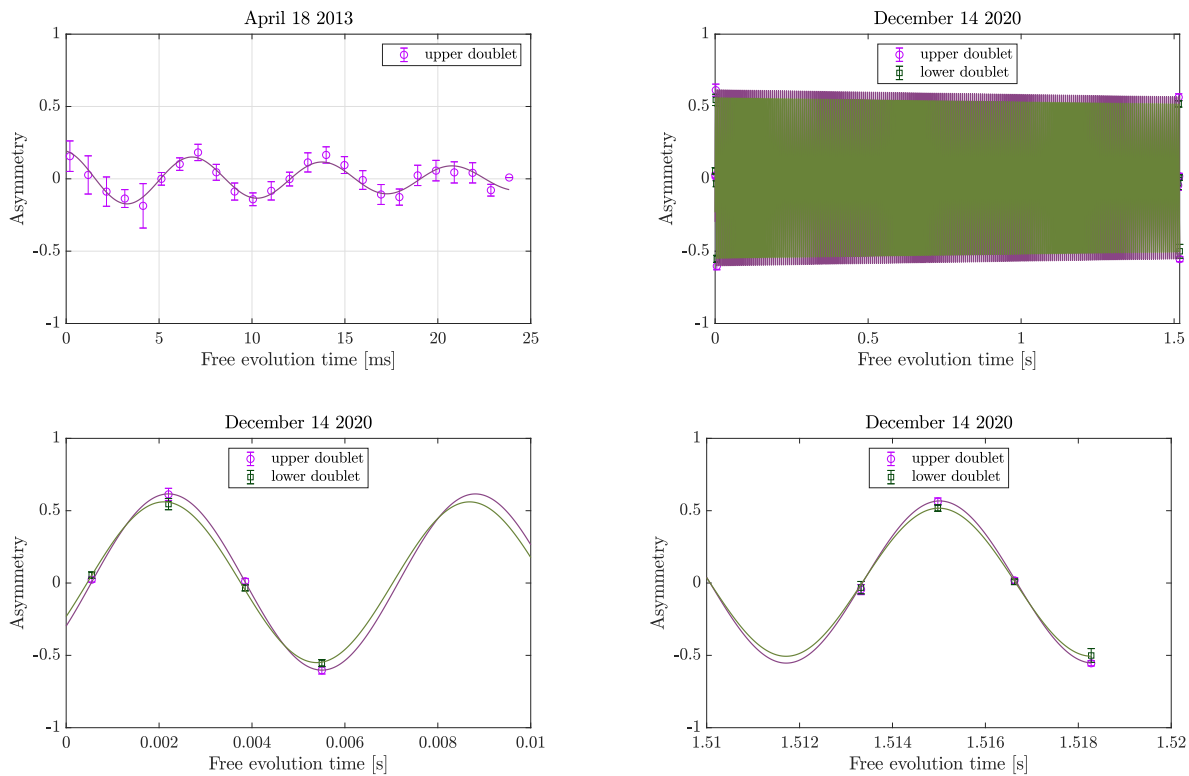


Figure 1.3: A comparison between fringes taken in 2013 (upper left) and much more recently, in 2020 (remaining panels). Note the time axis goes to 25 ms in 2013, while it goes to 1.5 seconds in 2020.

1.3 A non-exhaustive list of other searches for the eEDM

So far I've only given a few examples of the larger eEDM landscape, but these measurements have been happening for 70 years so it's a relatively mature field. Purcell and Ramsey were the first to propose such a search in 1950 [21]. The first limits were set via scattering [22], Lamb shift [23], and electron g-factor measurements in the late 1950's and early 1960's [24]. Limits were then set with paramagnetic atoms such as Rb, Cs, and Tl (see, for example, [25],[26],[27]) in the 1960's through 1970's. To give you a sense of the progress, the 1964 Cs result [26] constrained $|d_e| < 2 \times 10^{-21}$ e·cm, while the measurement detailed in this thesis constrains $|d_e| < 4.4 \times 10^{-30}$ e·cm.

In the early 1960's measurements were done in diamagnetic ^{129}Xe [28] and then later in the late 1980's with Hg [29]. From the late 1980's through the early 2000's Tl was consistently setting the most stringent limits ([30][31][16]) - by 2002 the limit was $|d_e| < 1.6 \times 10^{-27}$ e·cm. In the last 20 years, molecular experiments took over: since the first limit was set using YbF by the Hinds group in 2011 [18] ($|d_e| < 1.05 \times 10^{-27}$ e·cm), all subsequent limits have been set using paramagnetic polar molecules. In 2014 the ACME collaboration used a beam of ThO [20] to constrain $|d_e| < 8.7 \times 10^{-29}$ e·cm, and in 2017 our group published a similar limit using trapped HfF⁺ ions [2] ($|d_e| < 1.3 \times 10^{-28}$ e·cm). In 2018 ACME improved the limit by roughly an order of magnitude, again using a beam of ThO [32] ($|d_e| < 1.1 \times 10^{-29}$ e·cm). As mentioned above, the measurement in this thesis constrains $|d_e| < 4.4 \times 10^{-30}$ e·cm.

At the time of this writing, we are still the only eEDM experiment which actually *traps* molecules. There are a number of groups who are on the verge of trapping EDM-friendly neutrals, whether they have Ω -doubling structure [33] or not (which does not preclude sensitivity to BSM physics but requires much larger laboratory fields, often results in incomplete polarization, and offers no internal co-magnetometry) [34][35][36][37], a technique has been suggested for using already-trappable polar molecules [38], and some folks are even working on solid state

eEDM searches [39]. I even heard a rumor at DAMOP that another group is going to start using rotating fields [40].

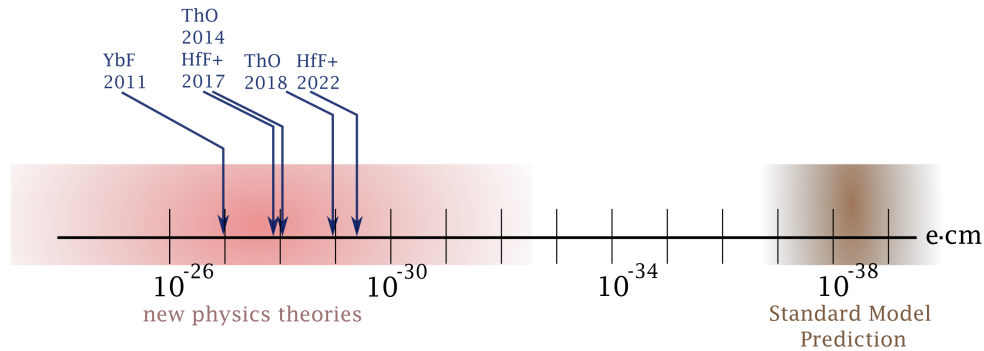


Figure 1.4: A cartoon depicting the current status of eEDM searches. The pink region indicates where most new physics theories predict the eEDM should be, and the blue arrows indicate recent exclusions (at the 90% confidence level). The exclusions measured zero up to the uncertainty indicated, which means that they exclude eEDMs larger than their uncertainty. The brown region indicates the standard model prediction for the eEDM. One can see that the region we are probing is effectively ‘background free’ in terms of the standard model prediction, and that most new physics theories predict EDMs right in the range we can measure - so our measurements directly constrain those theories.

1.4 Overview of this thesis

This thesis is focused entirely on JILA’s eEDM search, and almost completely on the second generation of the search, referred to as “Generation 2” or “Gen 2”. Chapter 2 details how one actually goes about measuring an eEDM and the specific details of our experiment, including details about the data analysis and the mathematical and algorithmic specifics we use to generate everything from Ramsey fringes (from raw ion counts) to linear combinations of data streams to get “frequency channels”. Chapter 3 is the systematics chapter, which is pretty self explanatory. Chapter 4 contains the final results from the Generation 2 measurement. Chapter 5 is about the dark matter exclusion we did with Generation 1’s data. Chapter 6 contains my conclusions about this work. Enjoy!

Chapter 2

How the experiment works

This chapter details how our experiment actually works, both theoretically and in practice. We begin this chapter with a section providing a generic overview of how one might measure an eEDM. Later sections will detail the theoretical foundations as well as the practical details of our experiment.

2.1 Measuring the electron's electric dipole moment: the basic recipe

In the most general terms, the goal of this section is to answer the question: how do you actually measure an electron EDM? We will cover the core concepts and basic ingredients here.

2.1.1 The basic recipe: put a dipole in a field - twice

Let's begin with something we understand¹: a simple spin-flip experiment. Let's say we have a spin in a magnetic field, and we measure the energy splitting between the two states (spin aligned or spin anti-aligned with the field, see Figure 2.1, left). Now we can add in the EDM: imagine it is pointing in the same direction as the spin, so the ket in Figure 2.1 represents both the direction of the spin *and* the EDM. If we also add an electric field, then the energy levels will get pushed apart for one relative orientation of the fields (see Figure 2.1, center). If we flip the direction of the electric field, the levels get pushed together (see Figure 2.1, right). If we take the

¹hopefully

difference between these two measurements, the magnetic term cancels out² and we measure an energy splitting directly proportional to the EDM. So you put a dipole in a field - twice.




Figure 2.1: Conceptual cartoon depicting the basic measurement principle. In each panel, the ket indicates both the direction of the magnetic moment (spin) and the electric dipole moment (EDM). In the first panel, a simple spin flip measurement is performed. In the second panel, an electric field is added, pushing the energy levels apart for that particular relative orientation of the fields. In the third panel the electric field is flipped, pushing the levels together. The difference between the last two measurements cancels out the magnetic dipole contribution (to first order) leaving only the electric dipole contribution.

2.1.2 Key ingredient 1: sensitivity

To keep pushing the limit on the eEDM we must keep improving our sensitivity (i.e. reducing our measurement uncertainty). To do this we need to consider three primary things: the electric field, the count rate, and the coherence/interrogation time. First, the larger the electric field the larger the energy level splitting will be (see Figure 2.2, left). Second, a large count rate will allow us to determine the centers of the peaks with more precision (see Figure 2.2, center). Third, a long coherence will make the lines more narrow (see Figure 2.2, right). All together, our sensitivity improves as the product of the effective electric field, the coherence time, and the square root of the counts. To optimize our sensitivity (minimize our uncertainty) we want to make all three parameters as large as possible.

$$\delta d_e \sim \frac{1}{|E_{\text{eff}}| \tau \sqrt{N}} \quad (2.1)$$

²to first order




Figure 2.2: Conceptual cartoon depicting how three key parameters affect our sensitivity. In the first panel, the separation between peaks scales linearly with the effective electric field, E_{eff} . In the second panel, having a higher count rate N reduces the fractional error in the determination of the peak centers by $\sqrt{1/N}$. In the third panel, the width of the peaks scales like the inverse coherence time, or $1/\tau$

To maximize the effective electric field we use paramagnetic polar molecules: the unpaired electron spin experiences a massively enhanced electric field inside the polarized molecule, requiring only modest laboratory fields to produce [41]. For example, in our experiment we apply a field of ~ 50 V/cm, which results in an effective electric field (experienced by the unpaired electron) of ~ 20 GV/cm. Of course, this comes with a catch: molecules have a much more complex energy structure compared to atoms. This can be both a blessing and a curse: it's a blessing because we are able to exploit certain molecular states to improve our measurement, but it's a curse because it makes the vast majority of molecules that would be sensitive to the eEDM very difficult to trap and cool. Molecules being *generically* difficult to trap and cool³ means that until recently all the leading eEDM measurements were done with molecular beams, which are constrained to \sim ms coherence/interrogation times.




Figure 2.3: A cartoon depicting a polarized HfF^+ molecule

To maximize the coherence time one needs to be able to trap the molecules - with a catch. To

³Of course, there are molecules amenable to trapping and cooling, but so far none of those candidates appear sensitive to the eEDM.

perform an eEDM experiment, you also need to *orient* the molecules in an external electric field. Doing both at the same time is the difficult bit. Lots of groups are working on the technology to trap eEDM-sensitive neutral molecules, and once these are trapped orienting them is (nominally) quite easy. Trapping eEDM-friendly neutrals is not easy, however - being sensitive to the eEDM usually means an energy level structure which is not very amenable to cooling and trapping. The other option is to work with charged species, which are very easy to trap. But until our experiment came along nobody could figure out how to orient trapped ions: any application of a DC electric field would push the ions right out of the trap. Eric's insight was to apply a *rotating* electric field in the $x - y$ plane, $\mathcal{E}_{\text{lab}} = \mathcal{E}_{\text{rot}}(\cos(\omega_{\text{rot}}t)\hat{x} + \sin(\omega_{\text{rot}}t)\hat{y})$ which the ions can adiabatically follow while staying trapped. With the rotating field trick, we enjoy coherence times three orders of magnitude larger than beam experiments [2] [32]. It's a really brilliant idea, and we are still the only people in the world using rotating fields to orient trapped ions.

Finally, to maximize the count rate you need to trap/cool as many molecules as possible, and detect them with maximum efficiency. This is where ion traps underperform compared to beam experiments: we can only trap so many ions at a time due to coulomb repulsion, and the experimental rep rate is limited by the interrogation time. In our experiment we count about 250 ions on the side of a Ramsey fringe (roughly per second), which seems really pathetic compared to the $\sim 10^6$ molecules per second counted in a beam experiment. But these numbers go under a square root, while the 3 orders of magnitude improvement we have in interrogation time compared to beam experiments does not. This is one of the ways we stay competitive.

2.1.3 Key ingredient 2: bias magnetic and electric fields

To perform the spin-flip experiment described above we need to apply a polarizing electric field as well as a bias magnetic field. The two fields need to be aligned then anti-aligned to get the two core measurements we need. The previous section described our polarizing electric field. The rotating field has two primary effects: first, it drives the ions in uniform circular motion,

where their orbits have radius $r = \frac{e\mathcal{E}_{\text{rot}}}{m\omega_{\text{rot}}^2}$ ($r \sim 0.5$ mm for our typical operating parameters). Second, the ions are accelerated into a rotating frame relative to the lab, which introduces a coupling between near-degenerate m_F levels, which we denote as Δ for $m_F = \pm 3/2$. Δ^D denotes the difference in coupling between Stark doublets due to differing proximity to the $F = 1/2$ manifold.

The rotating electric field and resultant circular orbits of the molecules leaves us with one more hurdle: how can we apply a bias magnetic field aligned (or anti-aligned) with a rotating quantization axis?⁴ The answer is to use a quadrupole magnetic field. The circular orbits (in the rotating frame) see a static magnetic field which is aligned or anti-aligned with the quantization axis⁵ (see Figure 2.4). The primary component of the spin precession frequency that we measure is due to this so-called rotating magnetic field.




Figure 2.4: A cartoon depicting how a radial quadrupole magnetic field amounts to a static, quantization-axis aligned magnetic field in the rotating frame. In each panel, the magnetic field gradient is indicated with light blue arrows. A circular orbit for a molecule is indicated near the top, and pink arrows depict the vector decomposition of the magnetic field which the ions see as they are rotating. You can see that as the molecule rotates, the vector decomposition of the magnetic field has two components: one which is always pointing “up”, and one which always aligns (or anti-aligns) with the molecular axis.

⁴Why have a magnetic field at all? We want to bias away from an avoided crossing between the two energy levels we measure, and to control the spin precession frequency.

⁵Technically, the quantization axis is the vector sum of the electric and magnetic fields, but the Zeeman interaction is so much smaller than the Stark interaction in our chosen states that we can, to an excellent approximation, neglect the contribution of the magnetic field to the direction of the quantization axis.

2.1.4 Key ingredient 3: Ω doubling (aka clever use of molecular energy structure)

We have almost everything queued up for a great experiment. The main remaining problem is that the shift we want to measure is either zero (up to our sensitivity) or very small, fractionally, compared to the primary Zeeman-induced precession frequency we will measure. To measure something that is fractionally tiny without polluting the result is difficult. One must take care to identify and, whenever possible, eliminate systematic shifts which might mimic the effect of a real eEDM. This is where Ω -doubling comes in! I'll provide a brief review of Ω -doubling here as well as a short discussion on why it is so helpful for systematic rejection. If you would like more details I'd recommend reading [42] and references therein, as well as [43]. The idea to exploit Ω -doubling to reduce systematics comes from Dave DeMille [44].

Ω -doubling is the name for a type of energetic substructure present in some molecules. Ω -doublet substructure generally occurs in molecular states with internal electronic angular momentum $J \geq 1$, and often in molecules with one heavy nucleus, two valence electrons in a triplet state, and at least one valence electron in a σ or $\pi_{1/2}$ orbital [42]. For the rest of this discussion we will discuss the $J = 1$ manifold of the $\text{HfF}^+ \ ^3\Delta_1$ state. HfF^+ 's $\ ^3\Delta_1$ state is well described by Hund's case a) basis states with coupled nuclear spin:

$$|\Lambda = \pm 2, S = 1, \Sigma = \mp 1, J, \Omega = \pm 1, I = 1/2, F, M_F\rangle, \quad (2.2)$$

where $\Lambda = \mathbf{L} \cdot \mathbf{n}$ is the projection of the electronic orbital angular momentum \mathbf{L} on the internuclear axis \mathbf{n} , $S = |\mathbf{S}|$ is the total electronic spin angular momentum, $\Sigma = \mathbf{S} \cdot \mathbf{n}$ is the electron spin projection on the internuclear axis, $J = |\mathbf{J}| = |\mathbf{L} + \mathbf{S} + \mathbf{R}|$ is the electronic plus rotational angular momentum, $\Omega = \Lambda + \Sigma = \mathbf{J} \cdot \mathbf{n}$ is the projection of the electronic angular momentum onto the internuclear axis, $I = |\mathbf{I}|$ is the ^{19}F nuclear spin, $F = |\mathbf{F}| = |\mathbf{J} + \mathbf{I}|$ is the total angular momentum of the molecule, and $m_F = \mathbf{F} \cdot \hat{\mathbf{z}}$ is the projection of \mathbf{F} on the rotating quantization axis [1] [2]. We take the internuclear axis \mathbf{n} to be directed from the ^{19}F nucleus to

the ^{180}Hf nucleus [2]. The above might be easier to digest in pictorial form; if so see Figure 2.5.




Figure 2.5: A cartoon depicting the vectors for Hund's case a. The electronic spin and orbital angular momenta, \mathbf{S} and \mathbf{L} , are strongly coupled to the internuclear axis by spin-orbit and electrostatic interactions respectively. Hund's case a arises when spin-orbit interactions are much weaker than the electrostatic interactions which split different electronic states. This means that widely separated electronic levels are split into a set of spin-orbit manifolds, each of which contains many closely spaced rotational levels. $\Lambda = \mathbf{L} \cdot \mathbf{n}$ is the projection of the electronic orbital angular momentum \mathbf{L} on the internuclear axis \mathbf{n} , $S = |\mathbf{S}|$ is the total electronic spin angular momentum, $\Sigma = \mathbf{S} \cdot \mathbf{n}$ is the electron spin projection on the internuclear axis, $J = |\mathbf{J}| = |\mathbf{L} + \mathbf{S} + \mathbf{R}|$ is the electronic plus rotational angular momentum, $\Omega = \Lambda + \Sigma = \mathbf{J} \cdot \mathbf{n}$ is the projection of the electronic angular momentum onto the internuclear axis, $I = |\mathbf{I}|$ is the ^{19}F nuclear spin, $F = |\mathbf{F}| = |\mathbf{J} + \mathbf{I}|$ is the total angular momentum of the molecule, and $m_F = \mathbf{F} \cdot \hat{\mathbf{z}}$ is the projection of \mathbf{F} on the rotating quantization axis [1] [2]. We take the internuclear axis \mathbf{n} to be directed from the ^{19}F nucleus to the ^{180}Hf nucleus [2].

In the $^3\Delta_1$ $J = 1$ state of HfF^+ we have two hyperfine states with total angular momentum $F = 3/2$ and $F = 1/2$. Each of these has a pair of manifolds made up of the two allowed projections $m_F = \mathbf{F} \cdot \hat{\mathbf{z}}$ onto the quantization axis. The two sets of states correspond to the eigenstates of the parity operator: $|e\rangle = (|\Omega = +1\rangle - |\Omega = -1\rangle)/\sqrt{2}$ and $|f\rangle = (|\Omega = +1\rangle + |\Omega = -1\rangle)/\sqrt{2}$. The degeneracy of these two manifolds is lifted by the Ω -doublet splitting, which we call ω_{ef} (see Figure 2.6).




Figure 2.6: In $\text{HfF}^{+'}s^3\Delta_1 J = 1$ state we have two hyperfine states with total angular momentum $F = 3/2$ and $F = 1/2$. Each of these has two degenerate manifolds made up of the different allowed projections $m_F = \mathbf{F} \cdot \hat{z}$ onto the quantization axis. The two sets of degenerate states correspond to the eigenstates of the parity operator: $|e\rangle = (|\Omega = +1\rangle - |\Omega = -1\rangle)/\sqrt{2}$ and $|f\rangle = (|\Omega = +1\rangle + |\Omega = -1\rangle)/\sqrt{2}$. These degenerate manifolds are separated by the Ω -doublet splitting, which we call ω_{ef}

The molecule has a molecule frame dipole moment⁶, and in the presence of an electric field states of good parity are mixed to become states of good orientation. The orientation states are either aligned or anti-aligned with the field, and we refer to these states as *Stark doublets*, or doublets for short. We refer to the aligned orientation as the “lower doublet” and the anti-aligned orientation as the “upper doublet” (see Figure 2.7, left). The addition of a magnetic field lifts the degeneracy further, as the spin is either aligned or anti-aligned with the external field (see Figure 2.7, center). Finally the effect of the eEDM further shifts the levels, as the eEDM (which we assume is aligned with the spin) is either aligned or anti-aligned with the internal electric field of the molecule (indicated with \vec{E}_{eff} , see Figure 2.7, right). A summary of all the shifts discussed above is depicted in Figure 2.8.

The effect of the eEDM has the opposite sign in the upper vs lower doublet, so this allows us to make a differential measurement between doublets to measure the energy shift due to the eEDM. A small complication is the magnetic g-factor is different in the upper vs lower doublet (indicated as g_f^u and g_F^l) due to electric-field-induced mixing with the $F = 1/2$ hyperfine level

⁶The molecule frame dipole moment points from negative to positive charge




Figure 2.7: Focusing only on the $F=3/2$ manifold of the ${}^3\Delta_1 J=1$ state of HfF^+ . $\Sigma = \mathbf{S} \cdot \mathbf{n}$ is the electron spin projection on the internuclear axis, and the internal electric field experienced by the valence electron is indicated as \vec{E}_{eff} . Left: an electric field polarizes the molecule along or against the field, which we call the “upper” and “lower” “doublets”. The energy levels without the added E field are indicated with dashed lines, and the energy levels with the added E field are indicated with solid lines. Center: adding a B field further lifts the degeneracy, when the spin is aligned with the B field the level moves down, and when it is anti-aligned it moves up. Again, the energy levels without the added B field are indicated with dashed lines and the energy levels with the added B field are indicated with solid lines. Right: finally, the addition of the eEDM itself further breaks the degeneracy, we assume the eEDM points in the same direction as the spin; when the eEDM is aligned with the internal electric field \vec{E}_{eff} the energy level moves further down and vice versa. As before, the energy levels without the effect of the eEDM are indicated with dashed lines and the energy levels with the effect of the eEDM are indicated with solid lines. Nothing in this cartoon is to scale.

and with the $J=2$ states.

Now here’s the clever bit: this energy level structure allows for robust systematic rejection due to misaligned fields: a transverse magnetic field, to first order, will couple m_F and $m'_F = m_F \pm 1$. Because we measure only the energy difference between stretched states, this first order shift won’t affect our measurement. To second order, the shift will affect the difference between the stretched states, but *this shift is common-mode between doublets*⁷ - so taking a differential measurement between doublets eliminates this (up to time-variations, which can be dealt with by measuring both doublets simultaneously, as we do) [42].

⁷Up to the difference in g-factor between doublets, see section 2.2.1

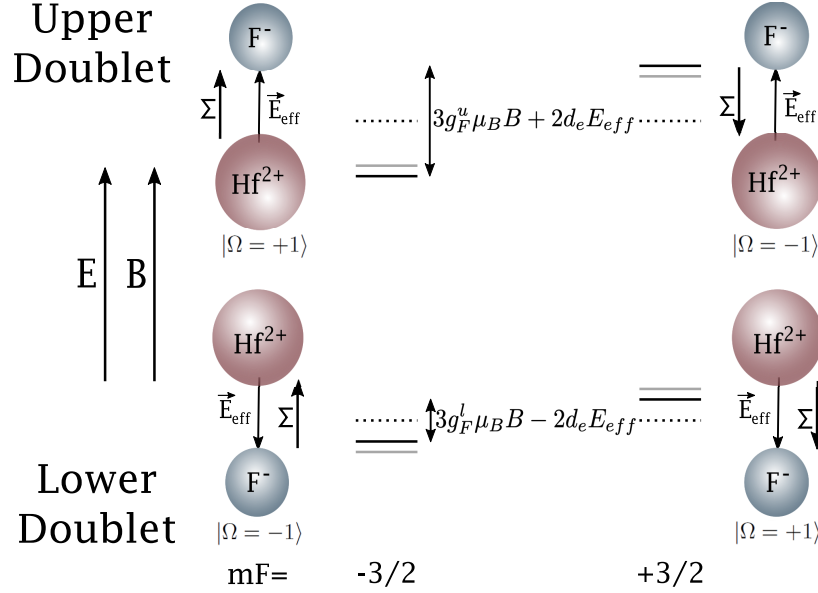


Figure 2.8: HfF^+ 's ${}^3\Delta_1, J = 1, F = 3/2$ state complete details. $\Sigma = \mathbf{S} \cdot \mathbf{n}$ is the electron spin projection on the internuclear axis, and the internal electric field experienced by the valence electron is indicated as \vec{E}_{eff} . The energy levels without the effect of the eEDM are indicated with grey lines, while the ones with the effect of the eEDM are indicated with black lines. Dashed lines are for visual aid only. The effect of the eEDM has the opposite sign in the upper vs lower doublet. It is this energy difference that we aim to measure with our experiment.

2.2 Measuring the electron's electric dipole moment: theoretical details

2.2.1 ${}^3\Delta_1, v = 0, J = 1$ Hamiltonian

We perform Ramsey spectroscopy on the stretched ($m_F = \pm 3/2$) states in HfF^+ 's ${}^3\Delta_1, J = 1, F = 3/2$ manifold. As mentioned in 2.1.4, the ${}^3\Delta_1$ state is well-described by Hund's case a) basis states with coupled nuclear spin. In the ${}^3\Delta_1, J = 1, F = 3/2$ manifold $|\Sigma| = 1$, $|\Lambda| = 2$, and $|\Omega| = 1$, and, again as mentioned in 2.1.4, both hyperfine (due to the fluorine nucleus' spin $I = 1/2$) and Ω -doubling structure are present. The hyperfine splitting is $-46.6(1)$ MHz between the $F = 1/2$ and $F = 3/2$ manifolds, and the Ω -doubling splitting is $740(40)$ kHz between $\Omega = 1$ and $\Omega = -1$ states. Treating each Stark doublet as a two-level system, we include the Zeeman, frame-rotation, Ω -doubling, and eEDM terms as primary contributors to the energy shift between the stretched states [2]. In the $F = 3/2, m_F = \pm 3/2$ basis, the two-level perturbative Hamiltonian has the form:

$$H_{\text{eff}}^{u/l} = \frac{1}{2} \begin{pmatrix} -3(g_F \pm \delta g_F)\mu_B(\mathcal{B}_{\text{rot}} + \delta\mathcal{B}_{nr} + \delta\mathcal{B}_i) \pm 2d_e|\mathcal{E}_{\text{eff}}| - 3\alpha h f_{\text{rot}} + \dots & h(\Delta \pm \Delta^D) \\ h(\Delta \pm \Delta^D) & 3(g_F \pm \delta g_F)\mu_B(\mathcal{B}_{\text{rot}} + \delta\mathcal{B}_{nr} + \delta\mathcal{B}_i) \mp 2d_e|\mathcal{E}_{\text{eff}}| + 3\alpha h f_{\text{rot}} + \dots \end{pmatrix}, \quad (2.3)$$

where the u/l superscript and upper or lower signs correspond to the upper or lower Stark doublet, $\alpha \equiv \mathcal{E}_Z/\mathcal{E}_{\text{rot}}$ is the tilt of the electric field out of the XY plane (in other words, the axial tilt), $\delta\mathcal{B}$ represents imperfect switching of the rotating magnetic field, and Δ is a rotation induced mixing of the angular momentum projections in the laboratory frame⁸ with a Stark doublet-odd contribution Δ^D (due to different levels of perturbation from the $F = 1/2$ manifold, the upper doublet being closer and therefore more perturbed). The subscripts on $\delta\mathcal{B}$, nr and i , indicate an axial magnetic field gradient $\delta\mathcal{B}_{nr}$ that does not switch with our applied switching axial field gradient (which ultimately comes down to the limit of our ability with the shim coil), and an actually-rotating-in-the-lab frame $\delta\mathcal{B}_i$ due to the charging current i in our 8-electrode trap. Here the magnetic g -factor g_F results from the combined nuclear and electronic magnetic moments, which are coupled by the nuclear hyperfine Hamiltonian. There is a difference in magnetic g -factors between the upper and lower Stark doublets, denoted as δg_F , which is due to Stark-induced mixing of $J = 1$ with $J = 2$ and \mathcal{E} -field induced mixing of adjacent magnetic sublevels at second order in perturbation theory. δg_F has the form

$$\frac{\delta g_F}{g_F} \approx -\frac{3d_{\text{mf}}\mathcal{E}_{\text{rot}}}{20B_e} \left(1 - \frac{g_N\mu_N}{3g_F\mu_B}\right) + \frac{3\hbar^2\omega_{\text{rot}}^2}{d_{\text{mf}}\mathcal{E}_{\text{rot}}A_{\parallel}} \left(1 - \frac{2g_N\mu_N}{3g_F\mu_B}\right) \approx -1 \times 10^{-3}. \quad (2.4)$$

For our experimental parameters the first term, which is linear in \mathcal{E}_{rot} , dominates. The coupling Δ first arises at fourth order in perturbation theory from the combined perturbations of rotation and Ω -doubling, and breaks the degeneracy of the $|m_F| = 3/2$ states in either Stark doublet at $\mathcal{B}_{\text{rot}} = 0$ [43, 45]. Δ and Δ^D are given by

⁸I think of it is like this: if there is a projection that is initially aligned with the quantization axis, then the axis rotates, the projection goes from being, eg., $+z$ to some mixture of $\pm z$. So, with our rotating bias field, adjacent m_F levels become mixed.

$$h\Delta = \frac{3\hbar\omega_{ef}}{2} \left(\frac{\hbar\omega_{rot}}{d_{mf}\mathcal{E}_{rot}} \right)^3 \left(\frac{18A_{\parallel}^2 - 19d_{mf}^2\mathcal{E}_{rot}^2}{A_{\parallel}^2 - d_{mf}^2\mathcal{E}_{rot}^2} \right), \quad h\Delta^D = \frac{3\hbar\omega_{ef}}{2} \left(\frac{\hbar^3\omega_{rot}^3}{d_{mf}^2\mathcal{E}_{rot}^2A_{\parallel}} \right) \left(\frac{9A_{\parallel}^2 - 8d_{mf}^2\mathcal{E}_{rot}^2}{A_{\parallel}^2 - d_{mf}^2\mathcal{E}_{rot}^2} \right). \quad (2.5)$$

These expressions are valid as long as $d_{mf}\mathcal{E}_{rot} \gg \hbar\omega_{ef}$ and $d_{mf}\mathcal{E}_{rot} \gg \hbar\omega_{rot}$ or, in other words, if the molecule is fully polarized. The strong scaling of Δ with \mathcal{E}_{rot} allows us to perform off-resonant $\pi/2$ pulses by modulating the magnitude of \mathcal{E}_{rot} .

The geometric frequency shift $3\alpha f_{rot}$ is due to a non-zero axial electric field, \mathcal{E}_Z . When \mathcal{E}_Z is nonzero, the total electric field tilts out of the $x - y$ plane, and it sweeps out a solid angle different from 2π . See [46] for a detailed discussion of this effect. $2d_e|\mathcal{E}_{eff}|$ is the EDM term, where d_e is the electron edm and \mathcal{E}_{eff} is the effective electric field experienced by the valence electron. Higher order contributions to $H_{eff}^{u/l}$, denoted by ellipses in Eq. (2.3), are due to higher-order combinations of perturbing Hamiltonian components, and primarily result in small corrections to the terms already discussed [2].

2.2.2 Measurement of $\delta g_F/g_F$

We re-measured the differential g-factor ratio $\delta g_F/g_F$ for improved accuracy compared to the Generation 1 measurement. There are two different ways to measure $\delta g_F/g_F$, one can look at the slope in f^D as one varies f^0 (by, for example, varying the magnitude of \mathcal{B}_{rot}), or the slope in f^{BD} as one purposely applies a non-reversing axial magnetic field gradient, generating an f_{nr}^B . We did both, see Figures 3.6 and 2.10.

2.2.3 Molecular structure

Put bluntly, HfF^+ is a complex molecule, and entire PHDs are done calculating various molecular parameters. Luckily for me, that is not what this thesis is about. For a nice detailed discussion of the energetic structure of HfF^+ , I'd recommend reading Chapter 2 of Will's thesis




Figure 2.9: Measurement of $\delta g_F/g_F$ by varying the mean fringe frequency f^0 , all data is taken at $\mathcal{E}_{\text{rot}} = 57.9\text{V/cm}$.




Figure 2.10: Measurement of $\delta g_F/g_F$ by adding a non-reversing axial magnetic field gradient, leading to large f^B . Note that f^{BD} is still blinded in this measurement.

[13]. This section is here simply to point you to the right place if you are looking for HfF⁺ details :)

2.2.4 Experimental noise vs the quantum projection noise limit

We've already seen a heuristic expression for the sensitivity in section 2.1.2. Now let's dig into the details, and try to figure out how far above the quantum projection noise limit we are.

The problem with really deriving the correct expression is that we never really know how many ions each experimental run starts with. In certain specialized experiments, the total number of spin-1/2 particles maybe known precisely (for instance, in chains of trapped ions, or in certain optical lattice experiments, or in situations where many photons can be scattered from each particle.) Often times, however, the total number of spins N_{tot} contained within a sample can be known to at best only $\pm \bar{N}_{\text{tot}}^{1/2}$. Where \bar{N}_{tot} is the average value for N_{tot} observed over many samples. One can imagine this limit arises because we have somehow randomly down-selected a small fraction much larger sample (say of neutral molecules) to get our sample spins (perhaps they are ionized molecules). In any case, when one knows the total number of spins only to the ?shot-noise? on N_{tot} , $\bar{N}_{\text{tot}}^{1/2}$, and when one then performs some coherent spin manipulation on the sample, and then measures a projection, say, the number of up spins N_{\uparrow} , one finds that the best one can do is to measure N_{\uparrow} to a one-sigma uncertainty of $\bar{N}_{\uparrow}^{1/2}$ where $\bar{N}_{\uparrow}^{1/2}$ is the average number of up spins measured over may samples. If one succeeds in realizing this precision, one is at the shot-noise limit. This situation can be contrasted to the situation in which one knows N_{tot} precisely, in which case the corresponding precision on measuring up-spins can be somewhat better:

$$\sigma = (\bar{N}_{\uparrow}(1 - \bar{N}_{\downarrow}/N_{\text{tot}}))^{1/2}. \quad (2.6)$$

Note that in this “known N_{tot} ” limit, variance vanished as \bar{N}_{\uparrow} goes to \bar{N}_{tot} whereas in the other limit the variance on N_{\uparrow} continues to be the naive shot-noise limit of $\bar{N}_{\uparrow}^{1/2}$ even as \bar{N}_{\uparrow} goes to \bar{N}_{tot} .

We are very much in the “imprecisely known N_{tot} ” limit, and worse yet, we actually have considerable technical noise in the production of our molecules, so noise from sample-to-sample on N_{tot} is actually often factor of 4 or 5 worse even than $\bar{N}_{\text{tot}}^{1/2}$. We use a differential technique, measuring the N_{\uparrow} population simultaneously in both stark doublets, so as to cancel the technical noise and thus on the side of the fringe we can measure the difference between the spin up

populations in the two stark doublets to the precision you would anticipate if each individual spin-up population were measured to within its shot noise limit [47].

I derived the expected shot noise limit in the case where we *do* know N_{tot} , you can find it in Appendix D.

For now, let's not worry about the stark doublet and think of a single fringe and see what we anticipate for the error in our spin asymmetry if we were able to measure individual spin projections at the shot noise limit. During our Ramsey experiment the spin *precesses*, so the initial state $(|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$ evolves such that at a later time t_R after a second $\pi/2$ pulse the state vector can be written as $\sim \cos(\omega t_R)|\uparrow\rangle + \sin(\omega t_R)|\downarrow\rangle$. So now the probabilities of measuring a given basis vector become sinusoidally time-varying: $p_{\uparrow} = |c_{\uparrow}|^2 = |\cos(\omega t_R)|^2$, which goes right into the binomial distribution of equation D.1 and the corresponding variance (eq D.2):

$$\sigma^2(t_R) = N_{\text{tot}}|\cos(\omega t_R)|^2(1 - |\cos(\omega t_R)|^2). \quad (2.7)$$

This is the quantum projection noise limited uncertainty for any measurement we will make of an ensemble of independent ions. In other words, this is the minimum uncertainty we can expect in any measurement of the number of ions in a given stretched state. How will this uncertainty affect our *total* uncertainty?

To compute the quantum projection noise limited total uncertainty, we need to go from an uncertainty in a single spin asymmetry measurement to an uncertainty in the precession frequency, then from that to a measurement of d_e . I'll sketch out a heuristic argument using only the sides of the fringe, I expect the rigorous argument is more complex than we need. Let's step through a perfect, quantum-projection-noise-limited version of our experiment and keep track of the uncertainties:

First, we prepare an ensemble of ions N_{tot} in a superposition state $(|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$, let it evolve, apply a second $\pi/2$ pulse, then measure the number of ions in the up state, N_{\uparrow} . The

probability of measuring a given value of N_\uparrow (on the side of a fringe, we expect $N_\uparrow = N_{\text{tot}}/2$) is given by the binomial distribution with variance $\sigma_{\uparrow/\downarrow}^2 = N_{\text{tot}}p_\uparrow(1-p_\uparrow) = N_{\text{tot}}/4$, since $p_\uparrow = 1/2$ on the side of a fringe.

Then we repeat the same process to measure N_\downarrow and again $\sigma_{\uparrow/\downarrow}^2 = N_{\text{tot}}/2$ since we are on the side of a fringe ($p_\uparrow = p_\downarrow = 1/2$).

Now we take these two measurements and form the fringe asymmetry: $\mathcal{A} = (N_\downarrow - N_\uparrow)/(N_\downarrow + N_\uparrow)$. Combining the errors we have a new uncertainty: the numerator and denominator have the same variance $\delta\mathcal{A}_{n/d}^2 = N_{\text{tot}}/2$, and the quotient of that has variance

$$\begin{aligned}\delta\mathcal{A}^2 &= \left(\frac{N_\downarrow - N_\uparrow}{N_\downarrow + N_\uparrow}\right)^2 \left(\frac{N_{\text{tot}}/2}{(N_\downarrow - N_\uparrow)^2} + \frac{N_{\text{tot}}/2}{(N_\downarrow + N_\uparrow)^2}\right) \\ &= \frac{N_{\text{tot}}/2}{(N_\downarrow + N_\uparrow)^2} + \frac{N_{\text{tot}}/2(N_\downarrow - N_\uparrow)^2}{(N_\downarrow + N_\uparrow)^4} \\ &= \frac{1}{2N_{\text{tot}}}\end{aligned}\tag{2.8}$$

where in the last line we used $N_\downarrow + N_\uparrow = N_{\text{tot}}$

Now we repeat this procedure r times, and we get $r/2$ asymmetry measurements. We average these measurements to get a better estimate of the spin asymmetry, and the sampling distribution of the mean⁹ is usually approximated by a normal distribution, so the quantum-projection-noise-limited uncertainty on an asymmetry measurement is $\delta\mathcal{A}_{\text{mean}}^2 = \delta\mathcal{A}^2/r$:

$$\delta\mathcal{A}_{\text{QPN}}^2 = \frac{1}{N_{\text{tot}}f_{\text{rep}}T_{\text{tot}}},\tag{2.9}$$

f_{rep} is the rep rate of the experiment, and T_{tot} is the total time the experiment was running, and $r = f_{\text{rep}}T_{\text{tot}}/2$.

On the side (zero-crossing) of a sinusoidally-varying asymmetry measurement, the slope (in units 1/s) is equal to ω ¹⁰, which we approximate as a line with zero offset: $y = \omega t$. If we are

⁹For $N_{\text{tot}} \gg 1$

¹⁰The side of the fringe $\cos(\omega t)$ is at $\omega t = \pi/2$, the derivative is $d/dt \cos(\omega t) = -\omega \sin(\omega t)$, plug in $\omega t = \pi/2$

fitting a line to a point with errorbar σ' (and $\sigma' \ll 1$), then the line can be moved up or down σ' , so the fitted line can range from $y = \omega t$ to $y = \omega t \pm \sigma'$. The horizontal distance (in this case, the time shift) between the "true line" $y = \omega t$ and the "line inferred due to the errorbar" $y = \omega t \pm \sigma'$ is given by finding the point on the x -axis where $y = \omega t = \sigma'$, so $t = \sigma'/\omega$. In a cycle, the time for a single degree is $t_{deg} = 1/360f = T/360$, where f is the frequency and T is the period. The time uncertainty due to our errorbar divided by the time for a single degree is the uncertainty in degrees: $\sigma'/\omega/1/360f = \sigma'/2\pi f/1/360f = \sigma'$ (remembering that the 2π is radians and 360 is degrees). So our phase error, in radians on the side of a fringe, is the same as the errorbar: σ' . Now if we say that our frequency measurement is given by $(\phi_L - \phi_E)/2\pi t_R$, where ϕ_L is the late-time phase, ϕ_E is the early-time phase and t_R is the ramsey free evolution time, then the frequency uncertainty is $\delta f = \frac{\sqrt{(\sigma')^2 + (\sigma')^2}}{2\pi t_R} = \frac{\sqrt{2}\sigma'}{2\pi t_R}$. This is all for a fringe with perfect contrast, for a non-unity contrast C we can scale it $\delta f = \frac{\sqrt{2}\sigma'}{C2\pi t_R}$. Finally, to get from a frequency uncertainty to a d_e uncertainty, we divide by $-2|\mathcal{E}_{eff}|\text{sgn}(g_F)/\hbar$ [2], where $|\mathcal{E}_{eff}|$ is the effective field experienced by the valence electron and $\text{sgn}(g_F)$ is the sign on the g -factor. Thus,

$$\delta d_e = \frac{\sqrt{2}\sigma'\hbar}{2|\mathcal{E}_{eff}|C2\pi t_R} \quad (2.10)$$

is the uncertainty in d_e for a single frequency measurement.

Finally, we will take the mean of all our frequency measurements, and the distribution of frequency measurements is again normal (we go to some effort to quantify this, see Chapter 4.) So our new uncertainty is just equation D.6 divided by the square root of the number of measurements:

$$\delta d_e = \frac{\sqrt{2}\sigma'\hbar}{2|\mathcal{E}_{eff}|C2\pi t_R \sqrt{f_{rep} T_{tot}/2}}. \quad (2.11)$$

Let's clean this up by taking $\hbar = \hbar/2\pi$ and cleaning up the factors of 2:

$$\delta d_e = \frac{\sigma'\hbar}{|\mathcal{E}_{eff}|C t_R \sqrt{f_{rep} T_{tot}}}. \quad (2.12)$$

Now all that remains is plug in the quantum-projection noise limited errorbar on a single frequency measurement, $\sigma' = 1/\sqrt{N_{\text{tot}}f_{\text{rep}}T_{\text{tot}}}$. Finally, we have the quantum-projection-noise-limited expression for our uncertainty in d_e :

$$\delta d_e = \frac{\hbar}{|\mathcal{E}_{\text{eff}}|Ct_Rf_{\text{rep}}T_{\text{tot}}\sqrt{N_{\text{tot}}}}. \quad (2.13)$$

The expression above assumes that in each experimental cycle we read out an identical number of ions, N_{tot} . In reality, N varies in every cycle due to fluctuations in the spatial mode and power of the lasers that ablate, ionize, and dissociate, as well as other experimental noise sources. We avoid most of that noise by performing a differential measurement between the upper and lower doublets: in a single shot of the experiment we measure both the lower and upper Stark doublets, and we perform our measurements when the fringes for both are in phase, so the proportional experimental noise is effectively cancelled out. Details of this scheme are in [47] and [48].

2.3 Experimental apparatus

2.3.1 Overview

Our experiment is a spectroscopy-on-trapped-ions experiment, so we have the usual stuff: lasers, a vacuum chamber, a source chamber, some oxford commas, and an ion trap. The general aesthetic is 'home built spaceship'.

2.3.2 Source chamber

Our source chamber is where we make neutral HfF. It contains our piezo valve, which opens to pulse SF₆ into the chamber. Details of the disk-piezo¹¹ valve can be found in Kevin Cossel's thesis [49]. The source is connected to the main chamber via a differential pumping chamber which contains two skimmers (detailed in Huanquian's thesis [50]).

¹¹Physik Instrumente P-286.23




Figure 2.11: CAD rendering of the entire apparatus, excluding lasers. In purple are the X shim coils, blue are the Y shim coils, and yellow are the Z shim coils. The orange coils generate the axial magnetic gradient that we call B_{rot}




Figure 2.12: Top down schematic of our experiment's main chamber

2.3.3 Ion trap

Our 10-electrode linear Paul trap has 8 radial electrodes and 2 axial electrodes. The 8 radial electrodes are driven in a quadrupole configuration to provide RF radial confinement (50 kHz), while they are operated independently to generate a rotating electric field (375 kHz) or DC shims. The 2 endcap electrodes provide DC axial confinement.

Our second-generation ion trap benefitted enormously from lessons learned in the first generation of the experiment [2]. The 8 radial electrodes have a shape carefully optimized to provide a drastically improved rotating field uniformity (over a several centimeter diameter area), increased trap depth, increased trap harmonicity, and highly suppressed electric field inhomogeneities. The result is a larger trapping volume which allows us to trap about 10 times more ions at the same or lower density. New electronics allowed us to generate about $2\times$ larger \mathcal{E}_{rot} , which directly suppresses collisional decoherence. The driving electronics were also improved to suppress ground loops and high frequency noise, all of which can cause general ‘heating’. Our trap is very unusual, in part for the rotating field but also because it is relatively *big*. To give you a sense of scale, you could put a coffee mug between the electrodes (see Figure 2.13). The details of the design of the Generation 2 ion trap and electronics can be found in Will Cairncross’ thesis [13].

The amplitudes of the confining, rotating, and shim fields vary throughout the experimental sequence, as pictured in Figures 2.17-2.16 for a 5 ms free evolution time Ramsey sequence.

To get from a voltage on the electrodes to a value of V/cm in the center of the trap, we need to account for the geometry of the electrodes. For the radial electrodes: take the electrode voltage and multiply by 0.1748 to get \mathcal{E}_{rot} or RF amplitudes in V/cm in the center of the trap, and for the endcap (axial) electrodes multiply the amplitude in voltage by 0.0146.




Figure 2.13: CAD rendering of the main experimental chamber sliced in half for a view of the ion trap. To give a sense of scale, you can fit a coffee cup between the electrodes.

2.3.4 Computer control

While almost every line of Matlab code got rewritten since Generation 1, and the LabView state machine was largely reprogrammed, the structure of the computer control is still about the same as detailed in Section 2.3 of Matt Grau's thesis [51] and Section 4.4 of Will Cairncorss' thesis [13].

2.3.5 Applied magnetic fields

Our experiment has 4 pairs of magnetic field coils¹² and no magnetic shielding. All coils are on the outside of the main vacuum chamber, see Figure 3.1.

Three of the coil pairs are set up along our lab-frame X , Y , Z axes in Helmholtz configuration for either shimming out the earth's magnetic field or for generating large transverse fields (for systematics studies). The magnetic field at the center of the trap is measured by an array

¹²*Slaps top of experiment* this baby can fit so many field coils in it

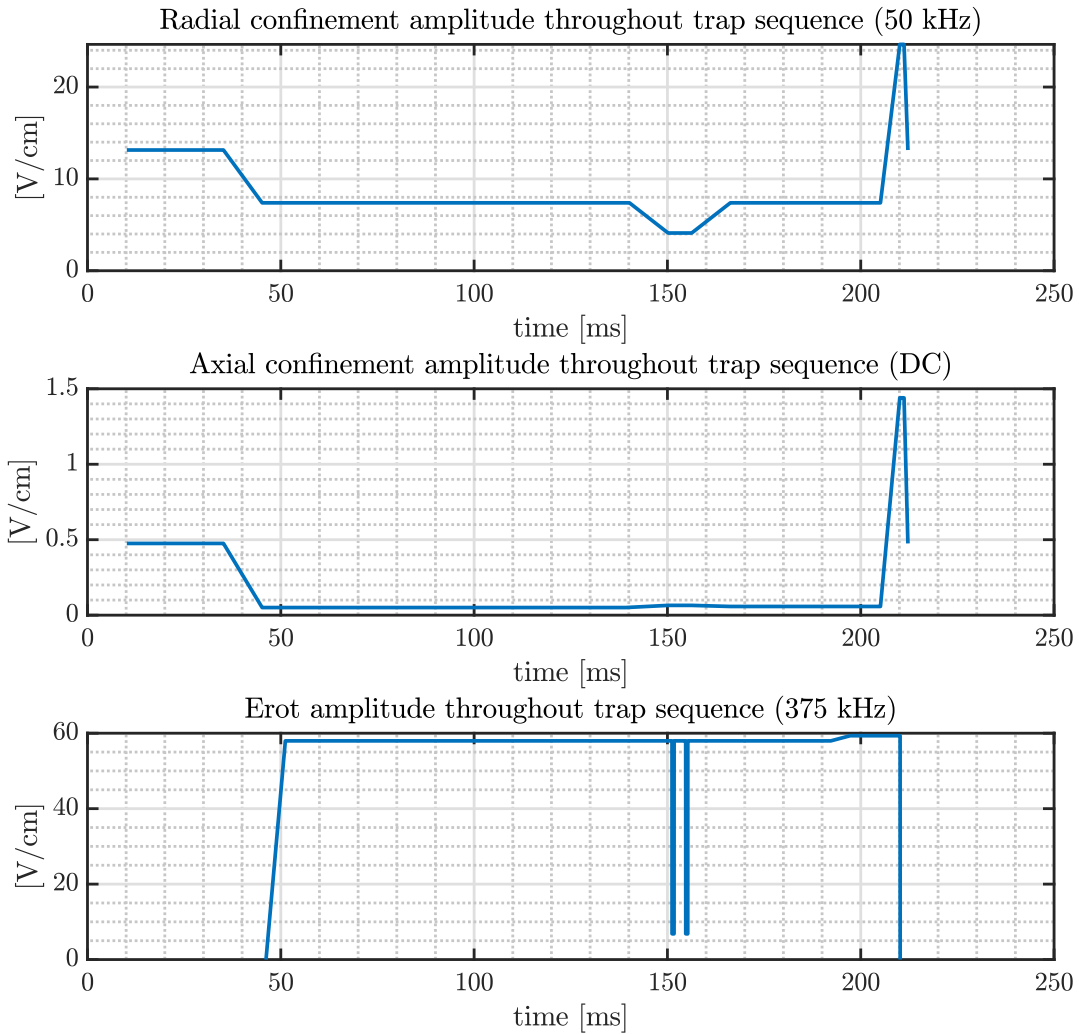


Figure 2.14: Confining and orienting field amplitudes throughout the experimental sequence, in Volts/cm, for a 5 ms Ramsey free-evolution time. The $\pi/2$ pulses are easily visible in the \mathcal{E}_{rot} panel.

of fluxgate magnetometers¹³ bolted to the main experimental chamber. The magnetometers are sensitive enough to detect whether the elevator across the hall is on the same floor as us or not.

The final coil pair, which we colloquially call “B axgrad coils” generate our so-called ‘rotating’ magnetic field. They are aligned with the lab-frame Z axis and produce the axial magnetic

¹³From Bartington Instruments, custom-built Mag612U’s with a 500 μ T range and a \sim 1kHz bandwidth

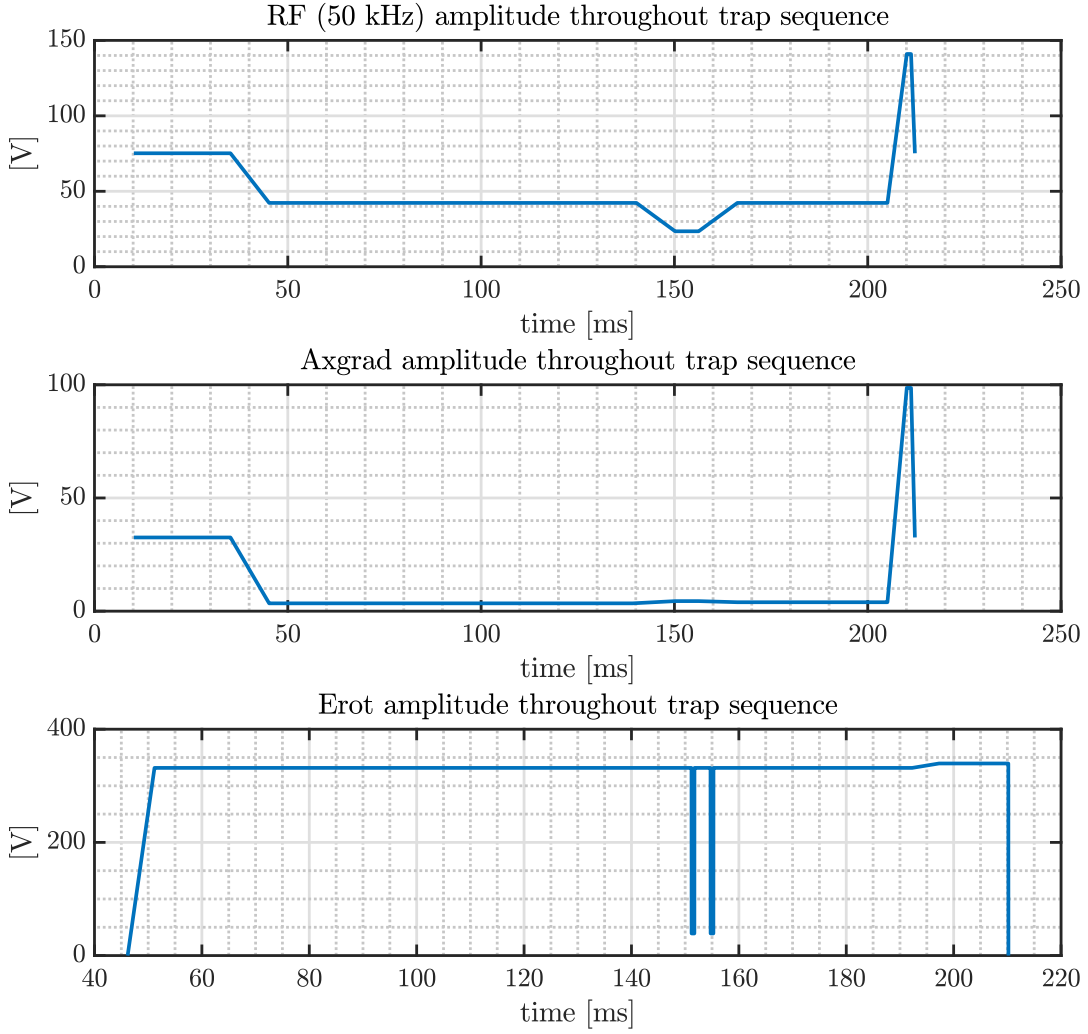


Figure 2.15: Confining and orienting field amplitudes throughout the experimental sequence, in Volts, for a 5 ms Ramsey free-evolution time. The $\pi/2$ pulses are easily visible in the \mathcal{E}_{rot} panel.

gradient $\mathcal{B} = \mathcal{B}'_{\text{axgrad}}(2Z - X - Y)$ (where, for example, $|\mathcal{B}'_{\text{axgrad}}| \approx 11.89$ G/m for an applied current of 0.55A). In the rotating frame of the ions this appears as a magnetic bias field

$$\mathcal{B}_{\text{rot}} \equiv \langle \mathcal{B} \cdot \mathcal{E}_{\text{rot}} / \mathcal{E}_{\text{rot}} \rangle \simeq \mathcal{B}'_{\text{axgrad}} r_{\text{rot}}.$$

Physical coil parameters are detailed in Table 2.1.

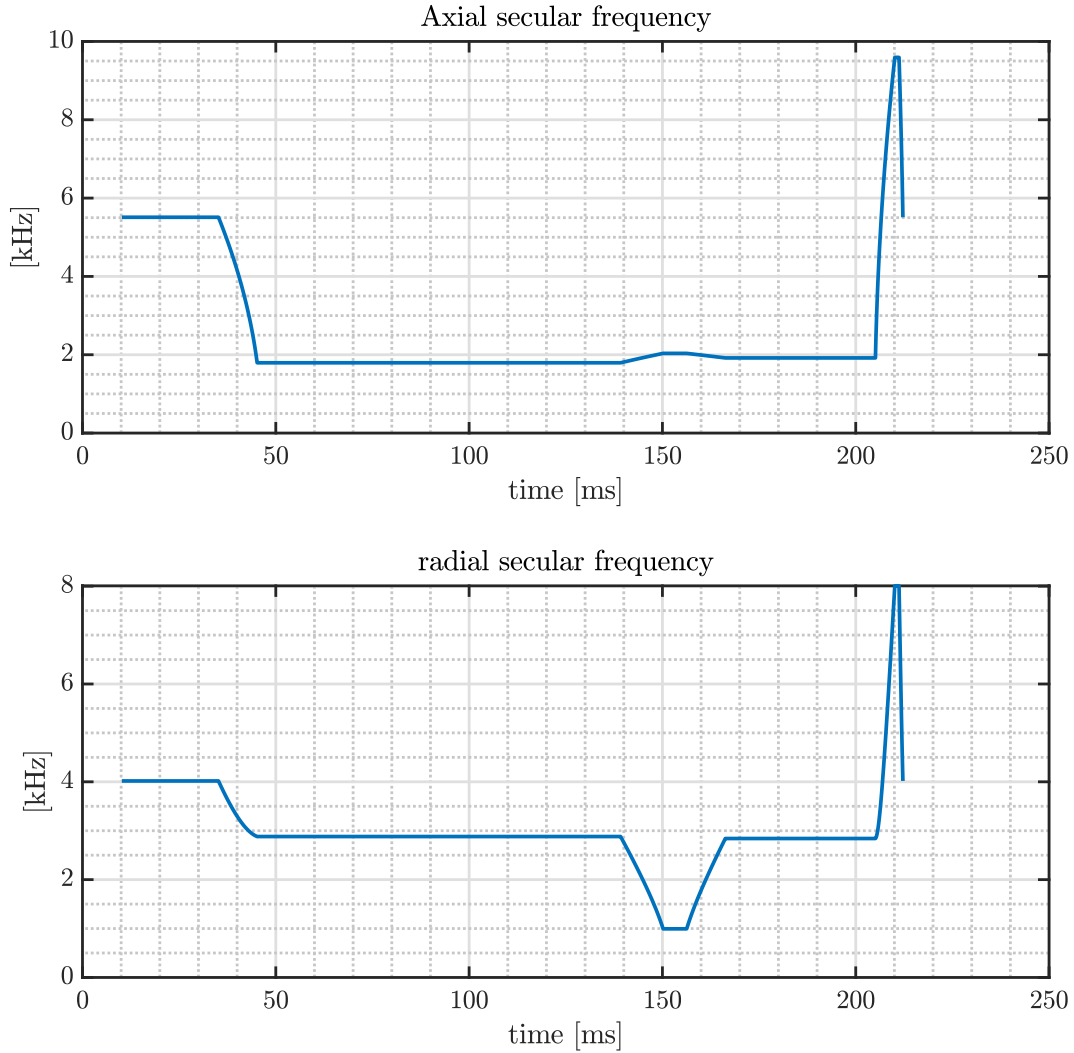


Figure 2.16: The variation of the computed secular frequencies throughout the experimental sequence for a 5 ms Ramsey free-evolution time.

2.3.6 Applied electric fields

As mentioned in 2.3.3, our 10-electrode ion trap has 8 radial electrodes and 2 axial electrodes. We can apply the RF, \mathcal{E}_{rot} or DC shim/confining potentials to any electrode, but in our normal operation we only apply a DC confining field to the endcap electrodes, and everything else is confined to the radial electrodes. In an idealized experiment, in the laboratory frame, we would only have 3 applied fields:

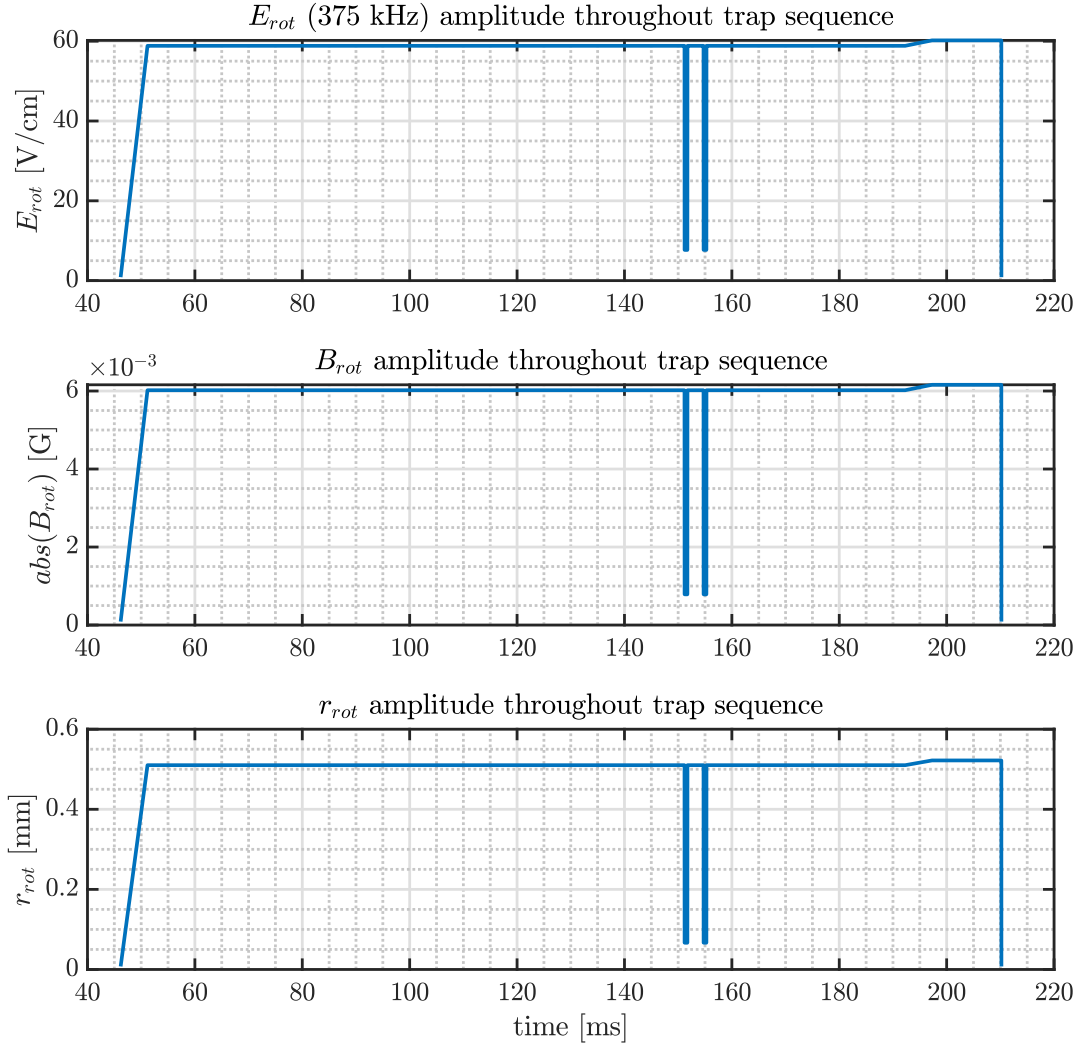


Figure 2.17: From \mathcal{E}_{rot} , computed values of \mathcal{B}_{rot} and r_{rot} throughout the experimental sequence for a 5 ms Ramsey free-evolution time.

First, the radial RF confining fields:

$$\mathcal{E}_{RF}(\mathbf{R}, t) = \frac{V_{rf}}{R_0^2} \cos(\omega_{RF}t)(\mathbf{X} - \mathbf{Y}), \quad (2.14)$$

where $\omega_{RF} \equiv 2\pi f_{rf} = 50$ kHz and $R_0 \sim 2.4$ cm is the effective radius the RF trap.

Table 2.1: Magnetic field coil parameters

Coil	# turns	Radius [m]	Separation [m]
X	80	0.217	0.216
Y	85	0.198	0.197
Z	57	0.171	0.172
Axgrad	85	0.135	0.192

Second, the axial DC confining fie

$$\mathcal{E}_{\text{DC}}(\mathbf{R}, t) = \frac{V_{\text{dc}}}{Z_0^2}(\mathbf{X} + \mathbf{Y} - 2\mathbf{Z}) \quad (2.15)$$

where $Z_0 \sim 8.3$ cm is the effective height of the RF trap.

And third, the radial rotating electric field:

$$\mathcal{E}_{\text{rot}}(\mathbf{R}, t) = \mathcal{E}_{\text{rot}} [\hat{X} \cos(\omega_{\text{rot}}t) - \tilde{R} \hat{Y} \sin(\omega_{\text{rot}}t)], \quad (2.16)$$

where $\omega_{\text{rot}} \equiv 2\pi f_{\text{rot}}$ and \tilde{R} is the rotation switch sign.

In our actual experiment, we see higher harmonics of both time-varying fields, though it is the higher harmonics of \mathcal{E}_{rot} which cause us trouble. We also see leakage of \mathcal{E}_{rot} and its higher harmonics on to the endcap, which is itself problematic (we will discuss this in Chapter 3). The second harmonic of \mathcal{E}_{rot} is of particular concern:

$$\mathcal{E}_{2h}(\mathbf{R}, t) = \mathcal{E}_{2h} \left(\frac{X\hat{X} - Y\hat{Y}}{R_0} \cos(2\omega_{\text{rot}}t + \phi_{2h}) + \frac{X\hat{Y} + Y\hat{X}}{R_0} \sin(2\omega_{\text{rot}}t + \phi_{2h}) \right), \quad (2.17)$$

where R_0 is the radius of the ion trap, and this is again in the laboratory frame [13]. These four fields comprise the main electric fields we are concerned with in our experiment.

For the radial electrodes: take the electrode voltage and multiply by 0.1784 to get \mathcal{E}_{rot} or RF amplitudes in V/cm in the center of the trap, and for the endcap (axial) electrodes multiply the amplitude in voltage by 0.0146.

2.3.7 Lasers

Below is a table detailing the laser beams used directly in our experiment. The generation of these beams takes a somewhat larger number of actual lasers. Complete details of all the lasers used in our experiment can be found in Appendix A.

Table 2.2: Photons used in our experimental sequence

Name	Transition	Power	λ (nm)	Polarization	Pulse Width
Ablation	n/a	10 mJ	532	linear	10 ns
Photoionization 1	$\Omega = 3/2 \leftarrow X^2\Delta_{3/2}$	30 μ J	309.388	linear	10 ns
Photoionization 2	Rydberg $\leftarrow \Omega = 3/2$	1.3 mJ	367.732	linear	10 ns
Transfer	$P(1) \ ^3\Pi_{0^+} \leftarrow X^1\Sigma^+$	600 mW	961.43495	linear	CW
mF Pumping	$P(1) \ ^3\Pi_{0^-} \leftarrow ^3\Delta_1$	21 mW	1082.4137	circular	CW, strobed
mF Depletion	$Q(1) \ ^3\Sigma_{0^+}^- \leftarrow ^3\Delta_1$	550 mW	814.508	circular	CW, strobed
Vibrational Cleanup	$P(1) \ ^3\Sigma_{0^+}^- \leftarrow ^3\Delta_1$	30 mW	818.37198	linear	CW
Dissociation 1	$\Omega = 2 \leftarrow ^3\Delta_1$	1.6 mJ	368.494	circular	10 ns
Dissociation 2	$? \leftarrow \Omega = 2$	25 mJ	266	circular	10 ns

All the lasers used in our experiment are locked to wavemeters using a simple servo loop. Note that for the first dissociation photon the wavelength is actually 368.4945nm or 368.493nm depending on the direction of rotation, to adjust for doppler broadening.




Figure 2.18: Cartoon depicting the transitions used during our state preparation.

2.3.8 Chops, switch states, blocks, and superblocks

The core of this measurement is taking the difference between measured transition energies (Ramsey fringe frequencies) when the electric field flips direction relative to the electron spin, as explained in Section 2.1.1. We also like to measure the difference and average of the fringe frequencies when we change other experimental parameters, like the \mathcal{E}_{rot} rotation direction. Ultimately, we measure the fringe frequency in 16 different experimental states, and we form linear combinations of those measurements to diagnose systematics and isolate our eEDM signal, as explained in Section 2.4.4. In addition to that, we prepare one of 2 stretched states and read out one of 2 stretched states, and we chop between those possible preparation and readout choices. Finally, we sometimes change a single parameter and take a set of measurements with that new parameter. All of this is what we call switch states, blocks, chops, and superblocks. Let me explain.

A *chop* is the fastest change we make to our experiment. Between single experimental runs, we change which stretched state we prepare and which we readout. For example in one run we might prepare $m_F = +3/2$ and readout $m_F = -3/2$, then in the next run we might prepare $m_F = -3/2$ and readout $m_F = -3/2$. We do this by changing the circularly polarized pumping and depletion lasers (see Section 2.3.7 and Section 2.3.10) to be σ^+ or σ^- in the ion's frame. The pumping and depletion lasers are circularly polarized by passing through $\lambda/4$ waveplates right before they enter the vacuum chamber. These waveplates stay fixed (they are not rotated between experimental chops). The pumping and depletion lasers are strobed at f_{rot} on a 49% duty cycle, and the relative direction of \vec{k} and $\vec{\mathcal{E}}_{\text{rot}}$, plus the handedness in the lab frame, determines if the light is σ^+ or σ^- . So, in practice, the *chop* consists strobing the lasers *in phase* or *out of phase* with \mathcal{E}_{rot} . There are 4 possible combinations: in/in, in/out, out/in, and out/out.

A *switch* is the next fastest change we make to our experiment. When we take measurements we do them in one of $2^4 = 16$ possible experimental states, which we call *switch states*.

We switch the sign of the magnetic field (*in* or *out*), the \mathcal{E}_{rot} rotation direction (again *in* or *out*), the Stark doublet (*upper* or *lower*), and the side of the screen we read a given Stark doublet out on (*left* or *right*). Note, we measure both doublets **simultaneously** (so the Stark doublet switch comes for ‘free’), but **separately** (on either side of the screen, so we still use it as a switch in the analysis - see Section 2.4.4). We indicate which *switch state* we are in with the letters $\tilde{\mathcal{B}}$, $\tilde{\mathcal{R}}$, $\tilde{\mathcal{D}}$, and $\tilde{\mathcal{I}}$, i.e. $\tilde{\mathcal{B}} = \pm 1$, $\tilde{\mathcal{R}} = \pm 1$, $\tilde{\mathcal{D}} = \pm 1$, and $\tilde{\mathcal{I}} = \pm 1$. Because the $\tilde{\mathcal{D}}$ switch comes for free, the actual number of switches we physically make is $2^3 = 8$.

In practice, there’s a few extra details to keep track of here, but we don’t have special names for them. The dissociation transitions need to be between stretched states to ensure the dissociation products inherit the correct momentum and the doublets separate on the screen. So if the chop is in/in or out/in (in both cases the readout state is “in”), then the dissociation beams, which are circularly polarized, need to be fired when they are also ‘in phase’ with \mathcal{E}_{rot} - in other words, if the depletion transition is σ^+ then the dissociation transition also need to be σ^+ , and vice-versa. In addition to this, the first dissociation photon will be red or blue shifted depending on the rotation direction, so we shift its wavelength depending on which combination of $\tilde{\mathcal{R}}$ and $\sigma^{+/-}$ we have.

A *Block* is what we call a set of Ramsey fringe measurements in all possible switch states - again since the $\tilde{\mathcal{D}}$ switch comes for free this means a block is comprised of 8 ramsey fringes.

Finally, a *superblock* is when we make a change to the apparatus then measure a set of blocks. For example, the $\lambda/4$ waveplates are all fixed in our normal operation, but they can each be rotated (and the corresponding logic changed to ensure we always dissociate stretched-stretched etc) and we can collect a set of blocks with them in their new position. The example given should have no effect on our science signal, but we do things like that just in case we do see something - in which case we investigate thoroughly!

2.3.9 Trap characterization

2.3.9.1 Secular frequencies

Because we change the RF and axial confinement throughout the experimental sequence, the secular frequencies are not constant. Their computed values for the entire sequence are shown in Figure 2.16, but those computed values do not account for any experimental imperfections. We can easily measure the secular frequencies by giving the ions a little kick to start them sloshing, then sending them on to the MCPs as a function of trapping time. Measuring the secular frequencies during the free-evolution time, we find that $f_X = 1$ kHz, $f_Y = 1.3$ kHz, and $f_Z = 1.7$ kHz.

2.3.9.2 Ambient f^B measurement

We collected data with I_{axgrad} set to 1.08 A in both the in and out direction and shifted the ions in z to see if we could get a handle on the 'ambient' f^B . The data is in `X:\2021\210901\systematic1` and `X:\2021\210831\systematic1` `X:\2021\210831\systematic2`. We observed quadratic shifts in both the f^B and f^{BR} channels. We explore this in more detail in Chapter 3.




Figure 2.19: Measurement of the 'ambient' f^B

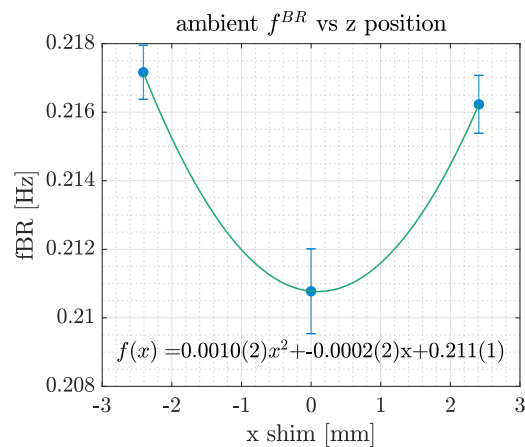


Figure 2.20: Measurement of the 'ambient' f^B

2.3.9.3 Mean field

We have approximately 20,000 ions in our trap on any given day. Coulomb repulsion causes mean field effects that we can observe. For example, in time of flight a thermal cloud should fit a Gaussian profile, but our cloud clearly deviates from this (see Figure 2.21).

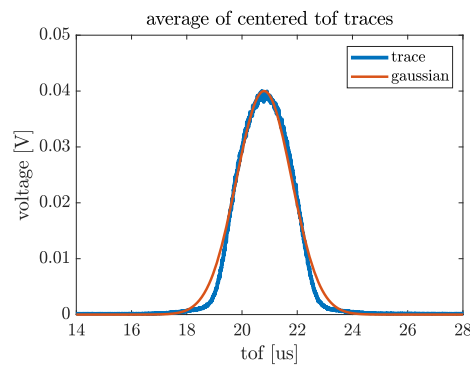


Figure 2.21: The time of flight trace of the HfF^+ cloud with a Gaussian overlaid

Generally, there is always going to be tension when we try to define things like temperature and size for an ion cloud with mean field effects. In particular, a lot of the following analysis relies on assumptions like expansion during time of flight is ballistic, so I just want to put a disclaimer for the rest of this section that there are some serious self-consistency issues. The part which does not rely on assumptions is the breathe, which we did measure, and which does show

the presence of mean field effects (see the very end of this section).

We can chop the trace in Fig 2.21 into slices and integrate each slice to compute the HfF number in each slice (see Fig 2.22).

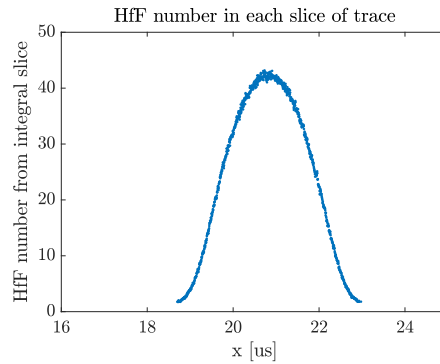


Figure 2.22: Using our HfF/voltage calibration, we compute the HfF in each slice of the TOF trace

The TOF data along with the Y-Z image data allows us to construct an approximate ellipsoid which contains the majority of the ions (see Fig 2.23).




Figure 2.23: Constructing the approximate HfF ellipsoid from our imaging and TOF data. In the Y-Z imaging data, the lines demarcate the approximate edge of the cloud.

Now we can use the area in each slice in the X direction along with how many ions are in each slice to determine the HfF density:

Using this, we can compute the density weighted density (8.9 for this dataset) and the




Figure 2.24: Approximate HfF density along X , and approximate Mean Field Parameter for our cloud over a given temperature range

approximate mean field parameter (see Fig 2.24).

If we estimate that our cloud is about 10 K then our mean field parameter is roughly 6. This means that we are definitely not a gas but clearly not in the coulomb crystal regime. The ion cloud behaviour should be somewhere between a pure gas and a liquid.

We can also give the cloud kicks and measure the breathe and slosh frequencies. For a thermal cloud the ratio of the breathe over the slosh frequency should be 2. For our cloud we measure $\sim 1.4^{14}$.

2.3.10 Experimental sequence

2.3.10.1 Neutral HfF production (~ 3 ms)

We make neutral HfF by ablating a Hafnium rod in the presence of SF_6 . First, a disk-piezo¹⁵ based valve¹⁶ opens for $\sim 65 \mu\text{s}$, and Ar(99%)/ SF_6 (1%) gas at ~ 120 psi is injected into the source chamber. The gas exits through a 1mm diameter nozzle, just downstream of which a Hf rod is positioned (such that it is on the edge of the expanding beam). About $1.21 \mu\text{s}$ after the start

¹⁴Analysis and plotting scripts are in X:\2020\200403\analyze_cloud_density .

¹⁵Physik Instrumente P-286.23

¹⁶See Kevin Cossel's thesis for details of the valve assembly [49]

of the gas pulse, an Nd:YAG pulsed laser (**Abby**, 532 nm, 10ns, ~ 10 mJ pulses) ablates the Hf rod. The Hf molecules from the ablation are entrained in the gas pulse and through the magic of chemistry, HfF molecules are produced.

The jet of HfF molecules (at this point both neutral and charged) is supersonically expanded, which cools the molecular beam to rotational temperature of ~ 10 K. As the beam traverses a differentially pumped chamber on its way to the main chamber it is collimated with two skimmers: the first (3 mm diameter) is located about 11 mm from the source and is biased at 200 V to remove any ions which were produced in the ablation; and the second (2 mm diameter) is about 9.6 mm further downstream and is not biased.

2.3.10.2 Photoionization (10ms)

When the neutral HfF molecules enter the trapping region, still entrained in the Argon buffer gas, they are excited via an optical-optical double resonance to a Rydberg state 54 cm^{-1} above the ionization threshold, and they autoionize (the two ionization photons come from **Harambe**, dye amplifier, 309.388 nm with $\sim 30 \mu\text{J}$ per pulse, and **Bertha**, dye laser, 368.351 nm with ~ 2 mJ per pulse, both pumped by **Mia**, YAG 532 nm, 10 ns, 10 Hz, and Harambe is seeded by **Tuna**, CW ECDL, 60 mW, $2 \times 309.388 \text{ nm}$)¹⁷. The ions inherit the 600 m/s center-of-mass velocity of the neutral molecules, and 'kicks' (short pulses) are applied to the electrodes to stop the ion beam in the center of the trap. The trap turns on immediately after the kicks are applied, and this RF quadrupole trap confines the ions. At the same time, the rotating electric field is turned on. We usually trap around 20,000 ions.

The ionization process is state-selective: the Rydberg state is energetically forced to decay to a single $^1\Sigma_0^+, \nu = 0$ vibronic ground state of HfF⁺. By choosing the correct parity of the intermediate state, as much as 50% of the ions end up in the $^1\Sigma_0^+, m_J = 0, J = 0$ state. Complete

¹⁷Harambe provides the 1st ionization photon and Bertha provides the 2nd ionization photon. Harambe is seeded by Tuna but both are pumped by Mia.

details of our ionization procedure can be found in Huanquian Loh's thesis [50].

We take some care when stopping and catching the ion cloud to minimize secular motion, which will ultimately dephase and result in a hotter cloud. We do this by applying small DC shims to the trap to shift the center of the RF trap to best overlap with the cloud (we call this 'minimizing the slosh'). After the trap is turned on and the ions are trapped, we slowly (over 10 ms) ramp the shim values down to zero, which we call 'centering the trap'. I've quoted the timing for photoionization as 10 ms because I have included the time taken to center the trap.

Note that the linewidth of the first ionization laser (Harambe) is narrow enough to resolve transitions of the various HfF isotopologues, allowing us to address transitions of only the $^{199}\text{Hf}^{19}\text{F}$ isotopologue¹⁸ [50]. Also note that ionization itself only takes tens of nanoseconds, but the YAG which pumps both ionization lasers operates at 10 Hz, so our experimental rep rate is limited to at most 10 Hz.

2.3.10.3 State transfer, cleanup, and m_F pumping (120 ms)

At this point we have trapped HfF⁺ in the ground electronic and vibrational state ($^1\Sigma^+$, $v = 0$), spread over the lowest 4 rotational levels. We connect $J = 0 \rightarrow 3$ using microwaves and perform incoherent transfer to our 'science' state: $^3\Delta_1$, $v = 0$, $J = 1$. We turn the CW transfer laser (**Topica**, CW ECDL, 961.434 nm, 600 mW) on as soon as the ions are trapped and leave it on for 120 ms. The transfer laser is resonant with $^3\Pi_{0+}$, $J = 0$ which decays preferentially to $^3\Delta_1$. You might notice that this seems to violate the transition selection rule that the total spin cannot change, but the $^3\Pi_{0+}$ state has an admixture of singlet and triplet character, allowing us to transfer from an $S = 0$ state to an $S = 1$ state.

Using a resonant transition and its resultant decay means that we are left with ions in $^3\Delta_1$

¹⁸The isotopologue of a chemical species has at least one atom which is a different isotope (has a different number of neutrons).

spread over several vibrational levels, primarily (until vibrational decay) in the lowest rotational level (thanks to coming from $J = 0$ upstairs). This presents two problems: first, ions which are in the lowest vibrational and rotational level are still spread over the zeeman sublevels, and we need to prepare a pure spin state. And second, ions in the upper vibrational levels will decay throughout our spin precession experiment, polluting any pure spin state that we might prepare in the lowest vibrational level.

We deal with the first issue by illuminating the cloud with low power CW lasers which resonantly connect ${}^3\Delta_1, v = 1, J = 1 \rightarrow {}^3\Sigma_{0+}^-, v = 1, J = 0$ or ${}^3\Delta_1, v = 1, J = 2 \rightarrow {}^3\Sigma_{0+}^-, v = 1, J = 1$ or ${}^3\Delta_1, v = 2, J = 1 \rightarrow {}^3\Sigma_{0+}^-, v = 2, J = 0$. The ${}^3\Sigma_{0+}^-$ preferentially decays back to ${}^1\Sigma^+$. These CW lasers are linearly polarized and are brought through the top of the trap, and have about 10 mW of power. We only actually use one laser at the moment (${}^3\Delta_1, v = 1, J = 1 \rightarrow {}^3\Sigma_{0+}^-, v = 1, J = 0$, which we call **vibrational cleanup**, home-built ECDL, 10 mW, 818.3719 nm) as this seems to maximize the contrast improvement we observe. We still have poor fringe contrast overall (only 50%) from non science state ions getting dissociated, we are not sure what the root cause is. I have been suspicious that the skimmer is no longer attached to the electrical lead and is no longer actually biased, so our background might be due to random state ions created during the ablation process being co-trapped with our state-prepared ions.

We perform our Ramsey spectroscopy between the $m_F = \pm 3/2$ levels in either the upper or lower Stark doublet, so we must prepare a pure spin state in anticipation of preparing a superposition. 40 ms after transfer is initiated we turn on a strobed, sigma polarized, CW laser (1083, DBR laser, 1082.4137 nm, strobed at f_{rot} , 49% duty cycle, 21 mW) which is brought in from the side of the trap and connects ${}^3\Delta_1, v = 0, J = 1 \rightarrow {}^3\Pi_{0-}, v = 0, J = 0$. The strobing is synchronous with \mathcal{E}_{rot} such that the wavevector is either parallel or anti-parallel to \mathcal{E}_{rot} , which drives a $\sigma^{+/-}$ transition to an $F' = 3/2$ manifold upstairs, eventually leaving only one stretched state populated in ${}^3\Delta_1$. This pumping laser is on for 80 ms.

2.3.10.4 First depletion (7 ms)

At this point we have nominally prepared a pure spin state, but just to be sure we illuminate the ion cloud with our CW depletion laser (**814**, CW ECDL, 814.5077nm, 540 mW at 50% duty cycle, strobed at f_{rot} such that it is parallel/antiparallel to \mathcal{E}_{rot}) which drives a $\sigma^{+/-}$ transition to ${}^3\Sigma_{0+}^-$. Our depletion laser connects ${}^3\Delta_1, v = 0, J = 1 \rightarrow {}^3\Sigma_{0+}^-, v = 0, J = 0$, which preferentially decays to ${}^1\Sigma^+$ by a 10:1 ratio, having weaker coupling to the ${}^3\Delta_1$ state. This additional step ensures that any ions in the wrong stretched state due to imperfect pumping are eliminated from the ${}^3\Delta_1$ manifold. Figure 4.6 depicts the lasers and transitions used so far (excluding ionization).

2.3.10.5 $\pi/2$ pulses and free evolution (varying timescale)

Now we have a pure spin state. We apply a $\pi/2$ pulse to create a coherent superposition of the $m_F = +3/2$ and $m_F = -3/2$ states. We apply the $\pi/2$ pulse by temporarily ramping down the magnitude of \mathcal{E}_{rot} , which increases a rotation-induced coupling $\Delta^{u/l}$ between $m_F = \pm 3/2$ states. The pure spin state will evolve into a superposition in ~ 1 ms [43] [50]. We ramp \mathcal{E}_{rot} back up, and allow the superposition to evolve for a variable amount of time. After the desired precession time has elapsed, we apply a second $\pi/2$ pulse to map the relative phase onto a population difference between the $m_F = \pm 3/2$ states.

2.3.10.6 Final state projection (second depletion, 25ms)

To project the ions into their final state, we again apply the depletion laser we used in Section 2.3.10.4, which selects which final state we will read out ($m_F = \pm 3/2$).

2.3.10.7 Dissociation and imaging

To detect and count the number of ions in our science state, we employ resonance-enhanced multiphoton dissociation [19] using two UV pulsed lasers (**Dissociation 1**, also known as Cora, a pulsed dye laser, fired such that its k-vector is parallel/antiparallel to \mathcal{E}_{rot} , and **Dissociation**

2, also known as 266, a pulsed YAG, fired such that its k-vector is parallel/antiparallel to \mathcal{E}_{rot}). The dissociation is partially state-selective, it should only dissociate molecules which are in the ${}^3\Delta_1, v = 0, J = 1, 2$ manifold (but we do have a large background of dissociated ions, we do not know where they are from. I suspect the skimmer biasing is broken). Because our ions are oriented in the trap, we dissociate them when they are parallel to our detector, and as long as we obey the angular momentum requirements detailed in Will Cairncross' thesis [13] (importantly, that both dissociation photons drive $\sigma^{+/-}$ transitions), the Hafnium atoms from the different doublets will fly to opposite sides of the phosphor screen. Chapter 6 of Will Cairncross' thesis covers our oriented photodissociation procedure in detail [13].

As mentioned in Section 2.3.8, if the chop is in/in or out/in (in both cases the readout state is "in"), then the dissociation beams need to be fired when they are also 'in phase' with \mathcal{E}_{rot} - in other words, if the depletion transition is σ^+ then the dissociation transitions also need to be σ^+ , and vice-versa. In addition to this, the first dissociation photon will be red or blue shifted depending on the rotation direction, so we shift its wavelength depending on which combination of $\tilde{\mathcal{R}}$ and $\sigma^{+/-}$ we have.




Figure 2.25: Some typical images showing how the two doublets appear on our phosphor screen: on the left, we depleted one doublet and imaged the one remaining. In the center we did the same with the opposite doublet, and on the right we depleted neither (so we are imaging both doublets at the same time). The numbers indicate pixels on the screen, the images are in false colour, and the images are an average of 20 separate runs of each experiment. The red line indicates our programmatic determination of the 'center' of the two blobs.

We gate the phosphor screen such that we only image the Hf^+ atoms, the un-dissociated HfF^+ 'spectator' ions (ions which were presumably in the wrong state) would otherwise ruin the

image. We detect both the Hf^+ and the HfF^+ spectator ions in time of flight. The HfF^+ signal allows to diagnose experimental issues. We do not normalize the counted Hf^+ by the integrated HfF^+ signal, we found it was less noisy to perform a differential measurement between the simultaneously detected orientation states [47]. The technical details of our imaging and counting system are detailed in [48] and [47]. The photons used in our experimental sequence are summarized in Table 3.3 and Figure 4.6.

For diagnostic purposes (see Section 2.4.1), we can also use the depletion laser tuned to be on resonance with one Stark doublet only (and sent in, linearly polarized, through the top of the trap) to deplete one Stark doublet at a time, and to image this doublet alone.

2.3.10.8 Data collection and basic processing

Our experimental signal is dissociated Hf^+ ions read out via phosphorescence on an imaging microchannel plate (MCP), as well as the time of flight scope trace from the MCP. The image is processed asynchronously such that instead of storing the whole image, which would take an enormous amount of disk space, we save a file which contains the locations and brightness of each bright spot which was determined to be an ion according to our smoothing and noise-reducing processing code. At the same time, we store the entire experimental state (which means, in practice, the state of every object in our LabView state machine). These two items, the ion counts and locations as well as the state of the experiment, form our base data.

Once we have this file, we use our determined 'center' of the image to count how many ions are on each side of the screen, and our switch state to determine which side of the screen corresponds to which doublet. We also know which stretch state we are reading out, so ultimately we end up with a count of the $m_F = \pm 3/2$ population in each doublet.

2.4 Data analysis

2.4.1 Hf counting and imaging

As mentioned in the previous section, our experimental signal is dissociated Hf⁺ ions read out via phosphorescence on an imaging microchannel plate (MCP), and we store a list of the locations and brightness of each ion in the image, as well as the precise state of the experiment.

Before we take any Ramsey fringes, we take a set of images in each switch state to determine where the *center line* between the two blobs of ions should be for each switch state (in case, say, the rotation direction shifts the center of the image). In total we take 3 types of image: one where we depleted the lower Stark doublet, one where we depleted the upper Stark doublet, and one where we deplete neither (see 2.25). We use these three images to determine the center line between the two blobs for each switch state. Then we define a *swatch*, which is a rectangular area centered on the center line whose ion counts are discarded. We do this because the doublets are not entirely separated on the screen, so in this area we can't be sure that we will assign the correct bright spots to the correct doublet.

Using Figures 2.26 and 2.27, we can get a heuristic understanding of the doublet/imaging contrast. The doublet contrast essentially quantifies the fraction of ions from a given doublet which land on the wrong side of the screen. In our experimental data, we see a blob for each doublet (like in Figure 2.26). If we integrate along one dimension, we will get two gaussian-like functions (Figure 2.27). If those traces are completely separated, then we can say the doublet contrast is 1. Say the 1D function resulting from integrating in one dimension is called f^L for the lower doublet and f^U for the upper doublet. Define the center of the blobs as y_c (indicated by the dashed black line in Figure 2.27). We can define the doublet contrast as

$$DC = 1 - \frac{2(\int_{y_c}^{1200} f^U(y)dy + \int_0^{y_c} f^L(y)dy)}{\int_0^{1200} f^U(y) + f^L(y)dy} \quad (2.18)$$

The doublet contrast defined above does not account for any background, meaning ions




Figure 2.26: An overlay of 20 images from our MCP assembly, where each doublet was fully depleted 10 times. The center is indicated with a black dashed line.




Figure 2.27: The integrated trace of Figure 2.26, showing the upper and lower doublets separately.

who were dissociated but who weren't from the science state. Our experimental data has a relatively high background (order 50 out of 350 ions in each doublet), so it's important to consider the effect of a background signal. The background ions arrive as a blob in the center of the image (we verify this by depleting all of $^3\Delta_1$ before dissociating), and call its integrated trace $b(y)$. In that case,

$$DC = 1 - \frac{2(\int_{y_c}^{1200} f^U(y)dy + \int_0^{y_c} f^L(y)dy) + \int_0^{1200} b(y)dy}{\int_0^{1200} f^U(y) + f^L(y)dy + b(y)dy} \quad (2.19)$$

Let's re-write this extremely cursed equation to make it more readable. Let $T = \int_0^{1200} (f^U + f^L)dy$, $B = \int_0^{1200} b(y)dy$ and $N = \int_{y_c}^{1200} f^U(y)dy + \int_0^{y_c} f^L(y)dy$ then our measured doublet contrast is

$$DC_{\text{apparent}} = 1 - \frac{2N}{T + B} - \frac{2B}{T + B} \quad (2.20)$$

So background ions reduce our apparent doublet contrast (for example, in the image above). However, these background ions get effectively subtracted away in our difference asymmetry (see section 2.4.2), so our *effective* doublet contrast is just

$$DC_{\text{effective}} = 1 - \frac{2N}{T}. \quad (2.21)$$

In either case, now that we have our center line and swatch for each switch state we can properly count Hf^+ ions in our images and assign them to the correct doublet and spin polarization. For every image (which corresponds to a single run of the experiment), we end up with a spin polarization assignment (in-in or in-anti etc), number of ions in the upper doublet, number of ions in the lower doublet, and complete experimental state including the switch state.

2.4.2 Spin asymmetry

For every switch state, we will take data in both the in-in and anti-in (or the anti-anti and in-anti) state-*prep* and readout combinations. If the preparation and readout phase are the same (ie in-in and anti-anti), then the fringe formed as we measure different free evolution times will have a 180 degree phase shift from the fringe formed when they are different (anti-in or in-anti). You can reason this through easily if you just imagine what happens on the Bloch sphere at zero time: if you prepare 'in' (say, $m_F = +3/2$) and readout 'in' with zero time evolution, you

should measure the maximum $N_{\text{in}} = N_{\text{max}}$, while if you prepare 'out' (in this example this is now $m_F = -3/2$) and readout 'in' you should measure the minimum $N_{\text{in}} = 0$.

Now let's call in-in and anti-anti "In" and in-anti and anti-in "Anti"¹⁹. Now we can form our asymmetry:

$$\mathcal{A}_{u/l} = \frac{N_{\text{In}} - N_{\text{Anti}}}{N_{\text{In}} + N_{\text{Anti}}}, \quad (2.22)$$

where the subscript u/l indicates the upper or lower doublet, respectively. As you can see, it takes at least two separate experimental runs to form a single asymmetry measurement. However, thanks to our imaging technique, we get an asymmetry measurement in both doublets simultaneously.

An typical ramsey fringe looks like Figure 2.28. The fringes from the upper and lower doublets have slightly different frequencies, mostly due to the difference in g -factor between the doublets.

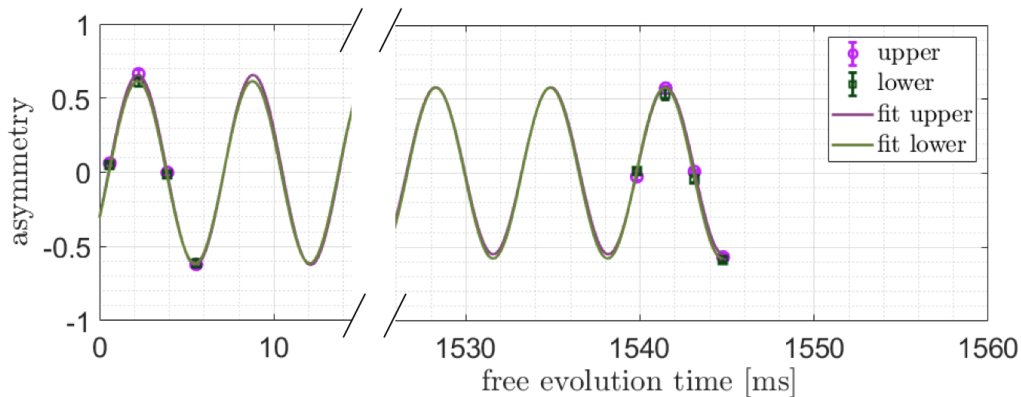


Figure 2.28: Example Ramsey interferometry data

We can also form two 'meta' asymmetries by taking the difference or the sum of the upper

¹⁹Sorry

and lower asymmetries:

$$\begin{aligned}\mathcal{A}_D &= \mathcal{A}_u - \mathcal{A}_l \\ \mathcal{A}_S &= \mathcal{A}_u + \mathcal{A}_l\end{aligned}\tag{2.23}$$

2.4.3 Fitting

As mentioned previously, we perform our Ramsey experiment simultaneously on both doublets and use their oppositely-signed angular momentum at the time of dissociation to read them out of opposing sides on the imaging MCP screen (see section 2.3.10.7). Unfortunately, the doublets are not fully separated on the screen, so we must be careful with how we fit our data. Below is a summary of our fitting routine.

For a simple Ramsey fringe with no leakage from the other doublet, we can define a decaying sine wave for the asymmetry:

$$g_{u/l}(t_R) = C_{u/l} e^{-\gamma_{u/l} t_R} \cos(2\pi f_{u/l} t_R + \phi_{u/l}) + O_{u/l}\tag{2.24}$$

where C is the initial fringe contrast, γ is the contrast decay rate, f is the fringe frequency, t_R is the free evolution time, ϕ is the initial phase, O is the offset, and the u/l subscript indicates the upper or lower Stark doublet. In our fitting routine, we initially fit each fringe separately to this function. From the fit parameters we define the mean and difference parameters as

$$\begin{aligned}\alpha_m &= \frac{\alpha_u + \alpha_l}{2} \\ \alpha_d &= \frac{\alpha_u - \alpha_l}{2},\end{aligned}\tag{2.25}$$

where α represents C , γ , f , ϕ , or O . We can redefine our upper and lower fit parameters in terms of the mean and difference parameters:

$$\begin{aligned}\alpha_u &= \alpha_m + \alpha_d \\ \alpha_l &= \alpha_m - \alpha_d.\end{aligned}\tag{2.26}$$

Due to imperfect imaging/doublet contrast, in reality each doublet's sine wave has a contribution from the other doublet, which leads to beating. If we let DC be our doublet contrast, then our

measured asymmetries will go like

$$\begin{aligned} A_u &= \left(\frac{1+DC}{2}\right)g_u + \left(\frac{1-DC}{2}\right)g_l \\ A_l &= \left(\frac{1-DC}{2}\right)g_u + \left(\frac{1+DC}{2}\right)g_l, \end{aligned} \quad (2.27)$$

Now our sum and difference asymmetries are

$$\begin{aligned} A_u + A_l &= g_u + g_l \\ A_u - A_l &= DC(g_u - g_l), \end{aligned} \quad (2.28)$$

which we can express in terms of the mean and difference fitting parameters:

$$\begin{aligned} A_u + A_l &= (C_m + C_d)e^{-(\gamma_u + \gamma_l)t_R} \cos(2\pi(f_m + f_d)t_R + (\phi_m + \phi_d)) + (O_m + O_d)... \\ &\quad + (C_m - C_d)e^{-(\gamma_u - \gamma_l)t_R} \cos(2\pi(f_m - f_d)t_R + (\phi_m - \phi_d)) + (O_m - O_d) \\ A_u - A_l &= DC((C_m + C_d)e^{-(\gamma_u + \gamma_l)t_R} \cos(2\pi(f_m + f_d)t_R + (\phi_m + \phi_d)) + (O_m + O_d)... \\ &\quad - (C_m - C_d)e^{-(\gamma_u - \gamma_l)t_R} \cos(2\pi(f_m - f_d)t_R + (\phi_m - \phi_d)) - (O_m - O_d)), \end{aligned} \quad (2.29)$$

We use the two functions above to simultaneously fit the sum and difference asymmetries. The fit uncertainty on f_d and ϕ_d are smaller than the fit uncertainty on f_m and ϕ_m thanks to our differential data collection and fitting routine.

2.4.4 Frequency channels

When we take measurements we do them in one of 16 possible experimental states, where we *switch* the sign of the magnetic field (*in* or *out*), the rotation direction (again *in* or *out*), the Stark doublet (*upper* or *lower*), and the side of the screen/image we read a given Stark doublet out on (*left* or *right*). Note, we measure both doublets **simultaneously**, but **separately** (on either side of the screen). We will indicate which *switch state* we are in with the letters \tilde{B} , \tilde{R} , \tilde{D} , and \tilde{I} , i.e. $\tilde{B} = \pm 1$, $\tilde{R} = \pm 1$, $\tilde{D} = \pm 1$, and $\tilde{I} = \pm 1$. With our switch state notation we can re-write our two-level Hamiltonian²⁰:

²⁰The \tilde{I} switch has no physical mechanism by which I can reasonably insert it into the hamiltonian, it's more of a sanity check than anything.

$$H_{\text{eff}}^{u/l} = \frac{1}{2} \begin{pmatrix} -3(g_F + \tilde{D}\delta g_F)\mu_B(\tilde{\mathcal{B}}\mathcal{B}_{\text{rot}} + \delta\mathcal{B}_{nr} + \tilde{\mathcal{R}}\delta\mathcal{B}_i) + \tilde{D}2d_e|\mathcal{E}_{\text{eff}}| - 3\alpha h\tilde{\mathcal{R}}f_{\text{rot}} + \dots & h(\Delta + \tilde{D}\Delta^D) \\ h(\Delta + \tilde{D}\Delta^D) & 3(g_F + \tilde{D}\delta g_F)\mu_B(\tilde{\mathcal{B}}\mathcal{B}_{\text{rot}} + \delta\mathcal{B}_{nr} + \tilde{\mathcal{R}}\delta\mathcal{B}_i) - \tilde{D}2d_e|\mathcal{E}_{\text{eff}}| + 3\alpha h\tilde{\mathcal{R}}f_{\text{rot}} + \dots \end{pmatrix}. \quad (2.30)$$

This two-level system has the familiar form

$$H = \begin{pmatrix} E_\alpha & \Delta \\ \Delta & -E_\alpha \end{pmatrix}, \quad (2.31)$$

which has eigenvalues $E = \pm 1/2\sqrt{\Delta^2 + E_\alpha^2}$. The full expression for the energy of this system is

$$\Delta E^{u/l} = \sqrt{[3(g_F + \tilde{D}\delta g_F)\mu_B(\tilde{\mathcal{B}}\mathcal{B}_{\text{rot}} + \delta\mathcal{B}_{nr} + \tilde{\mathcal{R}}\delta\mathcal{B}_i) + 3\alpha h\tilde{\mathcal{R}}f_{\text{rot}} - \tilde{D}2d_e|\mathcal{E}_{\text{eff}}|]^2 + [h(\Delta + \tilde{D}\Delta^D)]^2}. \quad (2.32)$$

To simplify this, we can think of it like $\sqrt{(A+a)^2 + x^2}$, where $A \gg a$ and $A \gg x$ since $d_e|\mathcal{E}_{\text{eff}}|$ is (obviously) tiny and (when \mathcal{E}_{rot} is large enough) $\Delta \pm \Delta^D$ is too. We expand our expression in a Taylor series: $\sqrt{(A+a)^2 + x^2} \sim |A+a| + \frac{x^2}{2|A+a|} \sim |A+a| + \frac{x^2}{2|A|}$. Dividing by h we get the measured frequency:

$$f = \frac{1}{h} |3(g_F + \tilde{D}\delta g_F)\mu_B(\tilde{\mathcal{B}}\mathcal{B}_{\text{rot}} + \delta\mathcal{B}_{nr} + \tilde{\mathcal{R}}\delta\mathcal{B}_i) + 3\alpha h\tilde{\mathcal{R}}f_{\text{rot}} - \tilde{D}2d_e|\mathcal{E}_{\text{eff}}|| + \frac{h(\Delta + \tilde{D}\Delta^D)^2}{2|3g_F\mu_B\mathcal{B}_{\text{rot}}|}. \quad (2.33)$$

You need to be careful when taking the absolute value: $3g_F\mu_B\mathcal{B}_{\text{rot}}$ is much bigger than all the other terms, so taking the absolute value flips \mathcal{B} 's sign onto the other terms (it can help to draw this out on the number line to convince yourself of this). So after taking the absolute value one is left with the theoretical measured frequency:

$$f(\tilde{\mathcal{B}}, \tilde{D}, \tilde{\mathcal{R}}) = \frac{1}{h} (3|g_F|\mu_B\mathcal{B}_{\text{rot}} + 3(|g_F|\mu_B(\delta\mathcal{B}_{nr} - \tilde{\mathcal{R}}\delta\mathcal{B}_i)\tilde{\mathcal{B}} - 3\delta g_F\mu_B\mathcal{B}_{\text{rot}}\tilde{D} - 3\delta g_F\mu_B(\delta\mathcal{B}_{nr} + \tilde{\mathcal{R}}\delta\mathcal{B}_i)\tilde{\mathcal{B}}\tilde{D} - 3\alpha h f_{\text{rot}}\tilde{\mathcal{B}}\tilde{\mathcal{R}} + 2d_e\mathcal{E}_{\text{eff}}\tilde{\mathcal{B}}\tilde{D} + \frac{h(\Delta + \tilde{D}\Delta^D)^2}{6g_F\mu_B\mathcal{B}_{\text{rot}}}). \quad (2.34)$$

We take the measurement in all possible switch states and we form linear combinations of that set of measurements to create what we call *frequency channels* or *parity states*. One measurement in all of the switch states is called a *block*. For example, (ignoring the $\tilde{\mathcal{I}}$ switch for now), we

form linear combinations of frequency measurements $f(\tilde{B}, \tilde{D}, \tilde{R})$ to form the frequency channels $f^{\tilde{B}\tilde{D}\tilde{R}}$ via the transformation

$$\begin{pmatrix} f^0 \\ f^B \\ f^D \\ f^{BD} \\ f^R \\ f^{BR} \\ f^{DR} \\ f^{BDR} \end{pmatrix} = \frac{1}{8} \begin{pmatrix} + & + & + & + & + & + & + & + \\ + & - & + & - & + & - & + & - \\ + & + & - & - & + & + & - & - \\ + & - & - & + & + & - & - & + \\ + & + & + & + & - & - & - & - \\ + & - & + & - & - & + & - & + \\ + & + & - & - & - & - & + & + \\ + & - & - & + & - & + & + & - \end{pmatrix} \begin{pmatrix} f^{++++} \\ f^{-+++} \\ f^{+--+} \\ f^{----} \\ f^{+--+} \\ f^{-+-} \\ f^{+--} \\ f^{----} \end{pmatrix}, \quad (2.35)$$

where \pm signs represent ± 1 . Combining the errors on each fringe we end up with one standard error for the D -odd channels and one for the D -even channels, given by

$$\delta f = \frac{1}{8} \sqrt{\sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm} (\delta f^{\tilde{B}\tilde{D}\tilde{R}})^2}, \quad (2.36)$$

where $\delta f^{\tilde{B}\tilde{D}\tilde{R}}$ are the standard error estimates obtained from the fitting routine described in Section 2.4.3. Note: given the matrix above you might expect the demoninator in equation 2.36 should be 4 since we average the D -odd channels separate from the D -even channels, but in the matrix above I ignored the \tilde{I} switch, so the correct number is 8.

We follow the same prescription to obtain parity states for other measured quantities like the contrast, phase, decoherence rate, and fringe offset.

To understand why we do this, consider the leading order terms of some of the frequency channels. $f^0 \sim 3|g_F|\mu_B\mathcal{B}_{\text{rot}}/h$ is the average fringe frequency. On the other hand, $f^B \sim 3|g_F|\mu_B\delta\mathcal{B}/h \sim f^0(\frac{\delta\mathcal{B}}{\mathcal{B}_{\text{rot}}})$ measures how accurately the magnetic field is reversed, you can think of this as the average fringe frequency suppressed by the magnetic field chop. f^D is the difference in frequency between the two Stark doublets: $f^D \sim -3\delta|g_F|\mu_B\mathcal{B}_{\text{rot}}/h$, the leading term is the fringe frequency suppressed by the difference in g-factors between the two doublets. If the

g-factor was identical for each doublet and the rotation-induced coupling terms were identical, the leading order f^D term would vanish. And so on for all the other channels. Listing the leading terms for all frequency channels we have:

$$\begin{aligned}
f^0 &\sim 3g_F\mu_B\mathcal{B}_{\text{rot}}/h + \frac{h(\Delta + \Delta^D)^2}{6g_F\mu_B\mathcal{B}_{\text{rot}}} \\
f^B &\sim 3g_F\mu_B\delta\mathcal{B}_{nr}/h \\
f^D &\sim -3\delta g_F\mu_B\mathcal{B}_{\text{rot}}/h + \frac{h\Delta\Delta^D}{3|g_F|\mu_B\mathcal{B}_{\text{rot}}} \\
f^{BD} &\sim 2d_e\mathcal{E}_{\text{eff}}/h - 3\delta|g_F|\mu_B/h(\delta\mathcal{B}_{nr} + \delta\mathcal{B}_i) + \dots \\
f^R &\sim 0 \\
f^{BR} &\sim -3\langle\alpha\rangle f_{\text{rot}} + 3g_F\mu_B\delta\mathcal{B}_i/h \\
f^{DR} &\sim 0 \\
f^{BDR} &\sim 3\delta g_F\mu_B\delta\mathcal{B}_i/h
\end{aligned} \tag{2.37}$$

Where again the subscripts on $\delta\mathcal{B}$, nr and i , indicate an axial magnetic field gradient $\delta\mathcal{B}_{nr}$ that does not switch with our applied switching axial field gradient (which ultimately comes down to the limit of our abilities with the shim coil), and an actually-rotating-in-the-lab frame $\delta\mathcal{B}_i$ due to the charging current i in our 8-electrode trap. The frequency channel we are most interested in is the one containing the eEDM: f^{BD} . f^{BD} is the antisymmetric combination of chopping both the doublet and the magnetic field. In general, forming frequency channels allows us to diagnose experimental issues and get a handle on most systematics.

2.4.5 Blinding

We performed our measurement *blind*: the frequency channel which contains d_e , f^{DB} , was blinded by the computer before we got to see it. In practice, this meant that as soon as the frequency channels were computed, a random number was added to the f^{DB} value. The random number was chosen once from a uniform distribution between 5 and 20 mHz, and it is stored in a coded file which the computer opens, reads the blind, adds to f^{DB} , then closes. When we

unblinded we subtracted the blind from the mean value of f^{BD} .

2.4.6 Data cuts

Please see chapter 4 for details of our data cuts.

Chapter 3

Accuracy evaluation

This chapter was co-written by Luke Caldwell and myself.

Welcome to the systematics chapter, comrade. This chapter begins with a reproduction of a manuscript in preparation, titled 'Accuracy evaluation of the HfF^+ electron electric dipole moment experiment'. After the manuscript draft follows a section titled 'Not Systematics' which contains everything that explicitly did not lead to systematic shifts in our measurement channel but which contributed to our understanding of our system.

The following sections are a reproduction of a manuscript in preparation

3.1 Accuracy evaluation of the HfF^+ electron electric dipole moment experiment

We have completed a new precision measurement of the electron's electric dipole moment using trapped HfF^+ in rotating bias fields. We report on the accuracy evaluation of this measurement, describing the mechanisms behind our systematic shifts. Our systematic uncertainty is reduced by a factor of 33 compared to the first generation of this measurement [2].

3.2 Introduction

Symmetries are fundamental in physics and exploring them has been vital to revealing much of what we understand about nature today. In 1967 Sakharov showed that violation of

combined charge and parity (CP) symmetry is required to explain the observed baryon asymmetry of the universe - an imbalance between the amount of matter and anti-matter. CP symmetry is not a good symmetry of nature but is only weakly violated by the Standard Model. Explaining the baryon asymmetry requires new physics with additional sources of CP violation and many extensions to the Standard Model have been proposed to do this.

Electric dipole moments of fundamental particles like the electron violate time-reversal (T) symmetry, equivalent to CP-violation assuming CPT invariance. The Standard Model predicts an electron EDM (eEDM) which is well below current experimental sensitivity. However many proposed extensions predict eEDMs which are several orders of magnitude larger, bringing its observation within experimental reach. Measurements of the eEDM thus constitute sensitive probes for physics beyond the Standard Model.

We recently completed the most precise measurement of the eEDM to date, constraining its size at $|d_e| < 4.4 \times 10^{-30}$ e·cm (90% confidence). This paper explains the details of our measurement. Section 4.2 describes our apparatus and experimental sequence, Section 3.4 the protocol used for data analysis, and Section 3.5 details the effective Hamiltonian we use to model the results. In sections 3.6–3.8 we report the measures we have taken to identify, characterize and mitigate sources of systematic error.

3.3 Experiment

Our experiment uses HfF^+ ions, confined in an ion trap and prepared in the metastable $^3\Delta_1$ state. In this state, the valence electrons are subject to a large intramolecular effective electric field $\mathcal{E}_{\text{eff}} = 23 \text{ GV cm}^{-1}$, along the internuclear axis of the molecule. We orient this molecular axis in the lab frame by applying an external electric field which rotates to maintain confinement of the ions. We then prepare the electron spin of the molecule in a coherent superposition of states parallel and antiparallel to \mathcal{E}_{eff} , and use Ramsey spectroscopy to measure the precession frequency. The electron EDM will give a contribution to this precession frequency proportional

Table 3.1: Constants used throughout this document. The total magnetic g-factor of $J = 1, F = 3/2$ states $g_F \equiv (G_{\parallel} + g_N \mu_N / \mu_B) / 3$ results from the combination of nuclear and electronic Zeeman effects.

Constant	Value	Description	Reference
B_e/h	8.983(1) GHz	Rotational constant of science state	[52]
A_{\parallel}/h	-62.0(2) MHz	Hyperfine constant	[2]
d_{mf}/h	1.79(1) MHz V ⁻¹ cm	Molecule-frame electric dipole moment	[2]
$\omega_{ef}/(2\pi)$	0.74(4) MHz	Ω -doubling constant	[52]
g_F	-0.0031(1)	$F = 3/2$ state g-factor	[50]
g_N	5.25774(2)	Nuclear magnetic g-factor of ¹⁹ F	[53]
$ \mathcal{E}_{\text{eff}} /h$	5.63×10^{24} Hz e ⁻¹ cm ⁻¹	Effective electric field	[54]

Table 3.2: Example experimental parameters and associated derived parameters from our 2022 data.

Parameter	Value	Description
\mathcal{E}_{rot}	58 V cm ⁻¹	Magnitude of rotating electric field during free evolution
ω_{rot}	$2\pi \times 375$ kHz	Rotation rate of \mathcal{E}_{rot}
\mathcal{B}_{rot}	11.9 mG	Computed from a measurement of B_{axgrad} and calculation of r_{rot}
$ \vec{\mathcal{B}}^0 $	23.5 G m ⁻¹	Measured using the magnetometers, for an I_{axgrad} of roughly 1.08 A
r_{rot}	0.5 mm	Radius of HfF molecule circular motion
$\frac{\delta g_F}{g_F}$	-0.002 146(2)	Stark doublet-odd magnetic g-factor ratio (see Figure 3.6)
Δ	~ 1 Hz	Electric field-induced m_F coupling
Δ^D	~ -0.6 Hz	Doublet-odd correction to Δ
RF	23.5 V	RF radial confinement amplitude during free evolution
RF	4.1 V cm ⁻¹	RF radial confinement amplitude during free evolution
RF	50 kHz	RF radial confinement frequency during free evolution
Axial confinement	3.7 V	DC axial confinement amplitude during free evolution
Axial confinement	0.05 V cm ⁻¹	DC axial confinement amplitude during free evolution
ω_x	$2\pi \times 0.95$ kHz	X secular frequency during free evolution
ω_y	$2\pi \times 1.51$ kHz	Y secular frequency during free evolution
ω_z	$2\pi \times 1.60$ kHz	Z secular frequency during free evolution




Figure 3.1: CAD rendering of the entire apparatus, excluding lasers. On the left is the source chamber, where we produce neutral molecules. To the right of the source chamber is the differential pumping chamber, where the beam of neutral molecules is supersonically cooled to a rotational temperature of roughly $\sim 10\text{K}$ and collimated by skimmers. Further right is the main experimental chamber, which contains our 8-electrode Paul trap. Surrounding the main chamber are the magnetic field coils: the X shim coils in purple, the Y shim coils in blue, and the Z shim coils in yellow; the coils which generate the quadrupole magnetic field \vec{B} are indicated in orange.




Figure 3.2: Top down schematic of our experiment's main chamber. The neutral molecules fly into the center of the trap where they are ionized and trapped. We apply the rotating electric bias field \vec{E}_{rot} and the quadrupole magnetic field \vec{B} . The ions undergo uniform circular motion, either aligned or anti-aligned with \vec{E}_{rot} , and their unpaired spins precess about $\vec{E}_{\text{rot}} \cdot \vec{B}$.

Table 3.3: Photons used in our experimental sequence

Name	Symbol	Transition	Power	λ (nm)	Polarization	Pulse Width
Ablation		n/a	10 mJ	532	linear	10 ns
Photoionization 1		$\Omega = 3/2 \leftarrow X^2\Delta_{3/2}$	30 μ J	309.388	linear	10 ns
Photoionization 2		Rydberg $\leftarrow \Omega = 3/2$	1.3 mJ	367.732	linear	10 ns
Transfer	$\mathcal{L}_{\text{trans}}^{961}$	$P(1) \ ^3\Pi_{0^+} \leftarrow X^1\Sigma^+$	600 mW	961.43495	linear	CW
m_F Pumping	$\mathcal{L}_{\text{op}}^{1082}$	$P(1) \ ^3\Pi_{0^-} \leftarrow ^3\Delta_1$	21 mW	1082.4137	circular	CW, strobed
m_F Depletion	$\mathcal{L}_{\text{depl}}^{814}$	$Q(1) \ ^3\Sigma_{0^+}^- \leftarrow ^3\Delta_1$	550 mW	814.508	circular	CW, strobed
Vibrational Cleanup	$\mathcal{L}_{\text{vc}}^{818}$	$P(1) \ ^3\Sigma_{0^+}^- \leftarrow ^3\Delta_1$	30 mW	818.37198	linear	CW
Dissociation 1		$\Omega = 2 \leftarrow ^3\Delta_1$	1.6 mJ	368.494	circular	10 ns
Dissociation 2		$n/a \leftarrow \Omega = 2$	25 mJ	266	circular	10 ns

to $d_e\mathcal{E}_{\text{eff}}$. To reject other unwanted contributions, we perform this measurement simultaneously on two spatially overlapping clouds of ions with their molecular axes aligned and anti-aligned with the externally applied field. The difference between the precession frequencies measured in each case is our science signal, proportional to $2d_e\mathcal{E}_{\text{eff}}$.

This section describes the apparatus, and each of the steps used in state preparation and measurement of the ions. A summary of typical parameters used in the experiment is given in Table 3.2.

3.3.1 Lasers

The experiment uses a total of 9 lasers¹; 5 pulsed UV lasers used for ablation, ionisation and photo dissociation, and 4 CW lasers—which we denote $\mathcal{L}_{\text{trans}}^{961}$, $\mathcal{L}_{\text{vc}}^{818}$, $\mathcal{L}_{\text{op}}^{1082}$, $\mathcal{L}_{\text{depl}}^{814}$ —used for state-preparation and readout. A summary is given in Table 3.3 and Fig. 4.6, and each is described in detail in the following sections. All lasers are locked to wavemeters using simple servo loops with a bandwidth of ~ 1 Hz.

¹There’s actually more lasers than this; the UV lasers need pumping and seed lasers, and we calibrate our wavemeter using an additional laser locked to a Rb transition. Our actual experiment had 13 ± 2 lasers going at any given time.

3.3.2 Molecular beam and ionisation

Our experiment begins with a pulsed beam of neutral molecules. We use a pulsed Nd:YAG laser to ablate a solid Hf rod into a pulsed supersonic expansion of Ar and SF₆. Chemical reactions between the Hf plasma and the SF₆ produce neutral HfF which are entrained in the supersonic expansion and rovibrationally cooled by collisions with the Ar to a temperature of ~ 10 K. When they arrive in our main chamber, ~ 50 cm away, a pair of pulsed UV lasers at 309 nm and 368 nm excite a two-photon transition to a Rydberg state 54 cm^{-1} above the ionization threshold, from which they autoionize [55, 50]. The molecular ions are created in the first few rotational levels of $^1\Sigma^+$, the electronic and vibrational ground state of the molecule. The ions are stopped at the center of our RF ion trap by pulsed voltages on the radial trap electrodes, after which the confining potentials are immediately turned on. We typically trap $\sim 20,000$ HfF⁺ ions with a lifetime of ~ 10 seconds. The trap is described in detail in the next section.

3.3.3 Ion trap

Our linear Paul trap has 8 radial electrodes and 2 endcaps. The radial confinement is provided by driving the radial electrodes in a quadrupole configuration producing a field,

$$\vec{\mathcal{E}}_{\text{RF}}(\vec{r}, t) = \frac{V_{\text{RF}}}{R_0^2} \cos(\omega_{\text{RF}} t) (\vec{x} - \vec{y}), \quad (3.1)$$

where $\omega_{\text{RF}} = 2\pi \times 50$ kHz, V_{RF} is the voltage applied on each electrodes, $R_0 \sim 4.8$ cm is the effective radius the RF trap, and \vec{x}, \vec{y} are the radial position coordinates of the ions in the laboratory frame. Axial confinement is provided by DC voltages V_{DC} on a pair of endcaps, producing a field

$$\vec{\mathcal{E}}_{\text{DC}}(\vec{r}, t) = \frac{V_{\text{DC}}}{Z_0^2} (\vec{x} + \vec{y} - 2\vec{z}) \quad (3.2)$$

where $Z_0 \sim 16$ cm is the effective height of the RF trap. We choose the values of V_{RF} and V_{DC} immediately after ionisation to best phase match with the size of the ion cloud, giving trap frequencies of around 5 kHz in all directions. We then linearly ramp the trapping voltages down over 10 ms to expand and cool the ion cloud. The ramp takes the trap frequencies to ~ 2.8 kHz

and ~ 2.0 kHz in the radial and axial directions respectively.

In addition to the RF confinement, we also apply a rotating electric field $\vec{\mathcal{E}}_{\text{rot}}$,

$$\vec{\mathcal{E}}_{\text{rot}}(t) = \mathcal{E}_{\text{rot}} [\hat{x} \cos(\omega_{\text{rot}}t) - \tilde{R}\hat{y} \sin(\omega_{\text{rot}}t)], \quad (3.3)$$

where $\omega_{\text{rot}} = 2\pi \times 375$ kHz, $\tilde{R} = \pm 1$ indicates the rotation direction and $\mathcal{E}_{\text{rot}} = |\vec{\mathcal{E}}_{\text{rot}}|$ is typically ~ 58 V cm $^{-1}$. This field serves to orientate the molecular axis, and thus the effective electric field, of the ions and we do our spectroscopy in this rotating frame. $\vec{\mathcal{E}}_{\text{rot}}$ causes an additional micromotion of the ions,

$$\delta\vec{R} = -\frac{e^2}{2m}\vec{\mathcal{E}}_{\text{rot}} = -r_{\text{rot}}\hat{\mathcal{E}}_{\text{rot}}, \quad (3.4)$$

where $r_{\text{rot}} \sim 0.5$ mm. The shape of the radial electrodes are optimized to minimize inhomogeneities in $\vec{\mathcal{E}}_{\text{rot}}$ across the ion cloud.

3.3.4 Magnetic fields

Measuring the electron EDM also requires orienting the electron spin of the molecules which we do with an applied magnetic field $\vec{\mathcal{B}}^0$. In order for the unpaired electrons to experience a time-averaged interaction with the intramolecular effective electric field, this magnetic field must corotate with $\vec{\mathcal{E}}_{\text{rot}}$. We achieve this by using a pair of coils in anti-helmholtz configuration, and aligned along the axial direction, giving

$$\vec{\mathcal{B}}^0 = \tilde{B}\mathcal{B}_{2,0}(2\hat{z} - \hat{x} - \hat{y}), \quad (3.5)$$

where $\mathcal{B}_{2,0}$ is typically 230 mG cm $^{-1}$ and $\tilde{B} = \pm 1$ indicates the direction of the current in the coils, explained in more detail in Sec. WHAT. In the rotating and co-moving frame of the ions, this quadrupole magnetic field appears as a time-averaged magnetic bias,

$$\mathcal{B}_{\text{rot}} = \langle \vec{\mathcal{B}}^0 \cdot \vec{\mathcal{E}}_{\text{rot}} \rangle = \mathcal{B}_{2,0}r_{\text{rot}}. \quad (3.6)$$

The coil pair is driven by a precision current source (Keysight B2902A). We refer to the pair of coils that produces this field as the $\vec{\mathcal{B}}^0$ -coils.

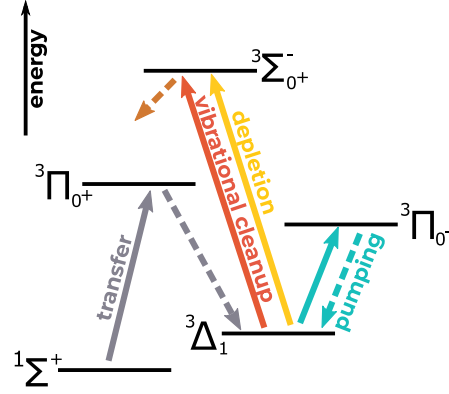


Figure 3.3: Cartoon depicting the transitions used during our state preparation.

The apparatus also includes three pairs of coils setup along the lab frame $\hat{x}, \hat{y}, \hat{z}$ axes in Helmholtz configuration for tuning the magnetic field at the position of the ions. The z coil is driven by the second channel of the precision current supply used for the \vec{B}^0 -coils, the x and y coils are driven by a lower precision current supply (Keithley 2230). The magnetic field around the periphery of the trap is measured by an array of 8 fluxgate magnetometers (Bartington Instruments, custom-built Mag612U's) bolted to the main experimental chamber. We use these measurements to infer the magnetic field at the center of the trap. In contrast to other modern eEDM experiments, the apparatus includes no magnetic shielding.

3.3.5 State preparation

Immediately after ionization, the HfF^+ ions are in the ground electronic and vibrational state ($1\Sigma^+(v=0)$), primarily distributed over the lowest 4 rotational levels $J=0-3$. We connect these rotational levels using microwaves and perform incoherent transfer to the eEDM-sensitive $3\Delta_1(v=0, J=1)$ state by using light from $\mathcal{L}_{\text{trans}}^{961}$ to drive the $3\Pi_{0^+}(J=0) \leftarrow 1\Sigma^+$ transition, the excited state of which decays preferentially to $3\Delta_1$. This light enters the chamber along the z -axis and is on for 80 ms beginning immediately after ionisation. The decay from $3\Pi_{0^+}$ puts population in several vibrational levels in $3\Delta_1$, which can decay into the $v=0$ science state if left untreated. We remove the population in higher vibrational levels by illuminating the cloud with

$\mathcal{L}_{\text{vc}}^{818}$ light which connects ${}^3\Sigma_{0+}^-(v = 1, J = 0) \leftarrow {}^3\Delta_1(v = 1, J = 1)$ at ~ 818 nm, preferentially decaying back to ${}^1\Sigma^+$. The $\mathcal{L}_{\text{vc}}^{818}$ laser also enters the chamber along the z axis and remains on for the duration of the experiment.

After transferring the ions to the science state, we ramp on $\vec{\mathcal{E}}_{\text{rot}}$ in 5 ms. In this field, the stretched states of ${}^3\Delta_1(J = 1, F = 3/2)$, can be considered states of good orientation, forming two pairs of levels—which we call doublets—with their molecular dipole, and thus \mathcal{E}_{eff} , either aligned or anti-aligned with \mathcal{E}_{rot} as shown in Fig. 4.1. Each doublet consists of one state with $m_F = 3/2$ and one with $m_F = -3/2$.

We polarize the molecules in the rotating frame by optically pumping them using $\mathcal{L}_{\text{op}}^{1082}$ light addressing the ${}^3\Pi_{0-}(v = 0, J = 0) \leftarrow {}^3\Delta_1(v = 0, J = 1)$ at 1082 nm. The light is circularly polarized with its k -vector in the plane of $\vec{\mathcal{E}}_{\text{rot}}$. We use an AOM to strobe $\mathcal{L}_{\text{op}}^{1082}$ synchronously with the rotation of $\vec{\mathcal{E}}_{\text{rot}}$ on a 50% duty cycle such that it drives either σ^+ or σ^- transitions to an $F' = 3/2$ manifold in the excited state. This eventually leaves population only in either the $m_F = 3/2$ or $m_F = -3/2$ states of ${}^3\Delta_1(v = 0, J = 1)$. The *preparation phase* of the experiment refers to the orientation of $\vec{\mathcal{E}}_{\text{rot}}$ relative to the k -vector of the light when the light is on; *in* when $\vec{\mathcal{E}}_{\text{rot}}$ is parallel, and *anti* when it is anti-parallel. This preparation phase can be changed by adjusting the timing of the strobing cycle. $\mathcal{L}_{\text{op}}^{1082}$ is turned on for 80 ms, starting 40 ms after trapping.

The final step of state preparation is applying $\mathcal{L}_{\text{depl}}^{814}$ light at 814 nm. This laser is tuned to address the ${}^3\Sigma_{0+}^-(v = 0, J = 0) \leftarrow {}^3\Delta_1(v = 0, J = 1)$, which preferentially decays to ${}^1\Sigma^+$ by a 10:1 ratio, having weaker coupling to the ${}^3\Delta_1$ state. $\mathcal{L}_{\text{depl}}^{814}$ is circularly polarized with the same handedness and k -vector as $\mathcal{L}_{\text{op}}^{1082}$. It is again strobed so as to only address and remove any residual population left over in other m_F states after $\mathcal{L}_{\text{op}}^{1082}$ is turned off. $\mathcal{L}_{\text{depl}}^{814}$ light is on for 7 ms, 3 ms after $\mathcal{L}_{\text{op}}^{1082}$ is turned off.

These steps leave the population in an incoherent mixture of the stretched states of the two doublets. The key difference from our previous measurement [2] is that the experiment proceeds

on *both* doublets simultaneously. Our detection scheme [47], described in Sec. 3.3.7, allows us to read out each independently, allowing us to take advantage of common-mode noise cancellation.

3.3.6 Ramsey sequence

Immediately prior to the Ramsey sequence, we ramp the radial confinement of the ions down further, to a trap frequency of ~ 1 kHz. This reduces the density of the cloud and improves the coherence time.

We apply a $\pi/2$ pulse to the ensemble of ions by temporarily ramping down the magnitude of $\vec{\mathcal{E}}_{\text{rot}}$ from 58 V cm^{-1} to 7 V cm^{-1} in $16 \mu\text{s}$, holding it there for 1 ms and then ramping back up in a further $16 \mu\text{s}$. Reducing \mathcal{E}_{rot} increases a rotation-induced coupling between $m_F = \pm 3/2$ states in a doublet (see Sec. 3.5), causing the pure spin states in each doublet to evolve into a coherent superposition. We allow this superposition to evolve for a variable amount of time t_R —up to 3 s —and then apply a second $\pi/2$ pulse to map the relative phase onto a population difference between the two states in a doublet.

3.3.7 Measurement

We project the ions into their final state by applying $\mathcal{L}_{\text{depl}}^{814}$ again to remove population from one of the stretched states in each doublet. The *readout phase* is defined in the same way as the *preparation phase*; *in* for $\vec{\mathcal{E}}_{\text{rot}}$ is parallel with the k -vector of the light when it is on and *anti* for antiparallel.

Finally, we detect and count the number of ions in the remaining stretched states via resonance-enhanced multiphoton dissociation [19], driven by two pulsed UV lasers at 368 nm and 266 nm . Immediately prior to the dissociation pulse, we ramp up both radial and axial confinement to compress the cloud and improve the dissociation efficiency. The dissociation pulse is timed so that $\vec{\mathcal{E}}_{\text{rot}}$ is along $\tilde{I}\hat{y}$, in the plane of a microchannel plate (MCP) and phosphor screen assembly. Here $\tilde{I} = \pm 1$ and \hat{y} defined by Fig. 4.2. Because the dissociation lasers enter at an

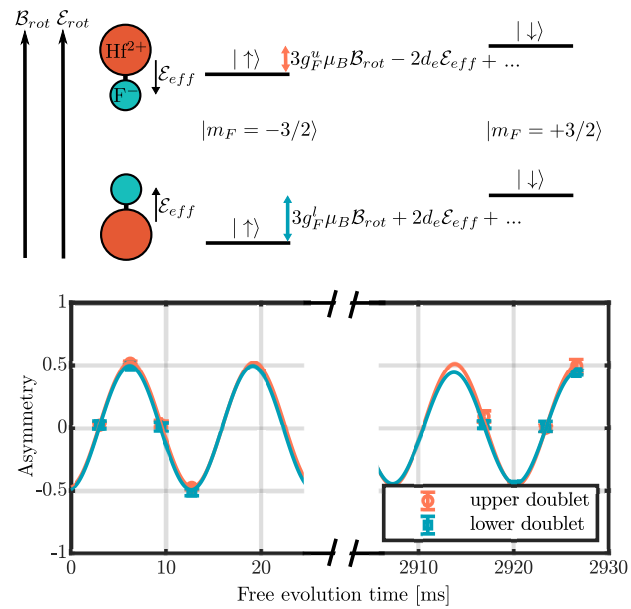


Figure 3.4: Ramsey spectroscopy in HfF^+ . Top, level structure in the eEDM-sensitive $^3\Delta_1$, $J = 1$, $F = 3/2$ manifold in external electric \mathcal{E}_{rot} and magnetic \mathcal{B}_{rot} bias fields (not to scale). The oriented molecules are either aligned or anti-aligned with \mathcal{E}_{rot} , which we call the lower and upper Stark doublets respectively. The component of the energy splitting due to the eEDM, $2d_e \mathcal{E}_{\text{eff}}$, flips sign between doublets. Bottom, example Ramsey fringes from our dataset. The fringes from the two doublets are collected simultaneously.

angle, there is considerable Doppler shift from the micromotion of the ions at 45° to the k -vector of the light. To account for this we adjust the wavelength of the 368 nm light depending both on the product $\tilde{R}\tilde{I}$ which gives the sign of the Doppler shift.

Each of the lasers is circularly polarized to drive transitions which preserve the orientation of the molecules during dissociation [13]. In this way the resultant Hf^+ ions from each doublet are ejected in opposite directions. The handedness of the dissociation lasers \tilde{P} is determined by a $\lambda/2$ waveplate which can be moved into or out of the beam path on a motorised mount. Immediately after dissociation, we turn off the RF confinement and apply pulsed voltage on the radial electrodes to kick the ions towards the MCP. The Hf^+ ions from each doublet are imaged on opposite sides of the phosphor screen; which side each doublet is imaged on is set by the value of \tilde{I} . We time-gate the phosphor screen such that we only image the dissociated Hf^+ and not any remaining HfF^+ ions. We detect both Hf^+ and HfF^+ ions in time of flight. Technical details of our imaging and counting system are detailed in [48] and [47].

We typically detect ~ 550 Hf^+ ions on each side of the screen early time, and ~ 120 after a 3 s interrogation time. The later is principally limited by the finite lifetime of the $^3\Delta_1$ state but with a contribution from ions being heated out of our shallow trapping potential during t_R .

3.3.8 Noise

Fluctuations in the power of the pulsed lasers used for ablation, ionization and photo-dissociation means that the fluctuations in the number of Hf^+ ions detected at the end of each shot are $\sim 30\%$, roughly $3\times$ the quantum projection noise limit for 120 ions. However these sources of noise, and many others, are common mode; the exact same laser pulses address the ions which end up in the upper and lower doublets. If we measure the ion number when the Ramsey oscillations of the two doublets are in phase with one another (or close) then we can take advantage of excellent noise cancellation in the number difference which we use to extract our science signal (see Sec 3.4). The two doublets oscillate at slightly different frequencies owing to

a small difference in their magnetic moments of about a part in 230 and so we deliberately take our data at the ‘beat’, where the two doublets are in phase. We take our early time data when the two doublets are in phase² and our late time data ~ 230 oscillations later when they come back into phase again. The time of the second beat can be controlled by varying oscillation frequency via the \vec{B}^0 -coils. In our final dataset, the noise in the science signal is roughly 10% above shot noise.

3.3.9 Experimental protocol and switch states

In each shot of the experiment we can choose the preparation phase to be either *in* or *anti*. For a given choice of \vec{I} —the direction of $\vec{\mathcal{E}}_{\text{rot}}$ at the moment of dissociation, and \vec{P} —the handedness of the dissociation laser polarization, the readout phase is constrained by the need to drive stretched-to-stretched transitions which preserve molecule orientation. We label each shot of the experiment with each of these phases. For example *in-anti* labels a shot where the cleanup (and optical pumping) laser are parallel to $\vec{\mathcal{E}}_{\text{rot}}$ during state preparation but anti-parallel during readout.

To record a Ramsey fringe, we repeat our measurement at different free evolution times. For a given fringe, the phase of readout is kept fixed. At each Ramsey time, we take an even number of shots with each pair consisting of one shot with each phase of state preparation. This set of shots is called a point and a Ramsey fringe consists of 8 of these points taken at different t_R ; we take 4 points at early Ramsey time, and 4 points at late Ramsey time, each consisting of two points on the sides of fringes and one point each on the top and bottom as shown in Fig. 4.1. The points on the sides of the fringes consist of 20 shots, while the points on top and bottom consist of 8 shots.

We record our data in ‘blocks’. Each block is constructed from a set of $2^3 = 8$ fringes

²Due to the finite length of the $\pi/2$ pulses, the doublets are already slightly out of phase at the earliest Ramsey times accessible to us. Systematic effects associated with this imperfection are discussed in Sec 3.8.

recorded in each possible combination of 3 experimental switches. Each switch corresponds to an experimental parameter whose sign can be reversed: \tilde{B} the direction of the current in the \vec{B}^0 -coils, \tilde{R} the direction of rotation of $\vec{\mathcal{E}}_{\text{rot}}$, and \tilde{I} the direction of $\vec{\mathcal{E}}_{\text{rot}}$ relative to the y axis at the time of dissociation, corresponding to which side of the phosphor screen each of the doublets are imaged onto. Note that in our implementation of the \tilde{I} switch, the direction of $\vec{\mathcal{E}}_{\text{rot}}$ is reversed at all points in time so that the opposite switch is prepared and read out. A fourth experimental switch, \tilde{P} the polarization of the dissociation light, is alternated every block.

We refer to each experimental configuration with $\{\tilde{B}, \tilde{R}, \tilde{I}, \tilde{P}\} = \pm 1$ as a switch state. In each block, the first Ramsey time is recorded for all switch states before moving onto the second Ramsey time for each switch state etc. The order of the switch states at each point is $\{\tilde{B}, \tilde{R}, \tilde{I}\} = \{1, 1, 1\}, \{-1, 1, 1\}, \{-1, 1, -1\}, \{-1, -1, 1\}, \{1, 1, -1\}, \{-1, 1, -1\}, \{1, -1, -1\}, \{-1, -1, -1\}$. Every other block, the order of switch states is reversed. In each switch state, we simultaneously collect data for molecules in each doublets, corresponding to orientation of the molecular axis with respect to the applied electric field which we represent by another switch $\tilde{D} = \pm 1$. The Ramsey times for each switch state are adjusted independently based on the data from the previous block to ensure that the 20 shot points are as close as possible to both the sides of the fringes, where our sensitivity is highest, and to the beat, where our noise cancellation is best and where various systematic shifts are minimized.

For the eEDM dataset, we collected 1370 blocks or ~ 600 hours of data over a ~ 2 month period of April–June 2022. During the data run, we took data with 3 different values of the $\mathcal{B}_{2,0}$, corresponding to fringe frequencies of $\sim 75, 105$ and 151 Hz. About halfway through the dataset, we changed the waveplates of $\mathcal{L}_{\text{op}}^{1082}$, $\mathcal{L}_{\text{depl}}^{814}$ and the dissociation lasers to reverse the handedness of the light from each.

3.3.10 Images to determine doublet positions

To determine where each doublet falls on the phosphor screen in each switch state, we take a series of images with no Ramsey sequence. For these images we prepare the stretched states as described in Sec. 3.8.1 and then remove the population in one of the doublets using low-power $\mathcal{L}_{\text{depl}}^{814}$ light, tuned to resonance with the doublet to be depleted, and propagating along the z direction to avoid micromotion-induced Doppler shifts. We take 3 types of image per switch state: one where we deplete the lower Stark doublet, one where we deplete the upper Stark doublet, and one where we deplete neither. We use these three images to determine the center line between the two blobs for each switch state as described in Section 3.4.1. We repeated this imaging routine roughly every 20 blocks during the dataset.

3.4 Data Analysis

3.4.1 Ion counting and asymmetry

Our experimental signal is dissociated Hf^+ ions read out via phosphorescence on an imaging microchannel plate (MCP). The images are processed asynchronously and we save a file which contains the locations and brightness of each bright spot which was determined to be an ion according to our smoothing and noise-reducing processing code.

We use this same code to analyze the test images described in Sec. 3.3.10 and find a *center line* for each switch state. We use this center line when analyzing the Ramsey data to define a *swatch* which is a rectangular area, of fixed width, at the center of the image whose ions counts are discarded. We do this because the doublets are not entirely separated on the screen, so in this area we can't be sure that we will assign ions to the correct doublet. The extent to which we are able to isolate the two doublets is given by the imaging contrast C_I . As the swatch width increases, the imaging contrast improves but total ion number N decreases as we throw out more ions. For the final analysis of the dataset, we used a swatch of 90 pixels which roughly maximizes

$C_I\sqrt{N}$.

Once we have our center line and swatch for each switch state we can properly count Hf^+ ions in our Ramsey data images and assign them to the correct doublet. For every image (which corresponds to a single run of the experiment), we end up with a number of ions in the upper doublet N^u , number of ions in the lower doublet N^l , and complete experimental state including the switch state, preparation phase and readout phase.

For every switch state, we take data in both the in-in and anti-in (or the anti-anti and in-anti) combinations of preparation and readout phase. If the preparation and readout phase are the same (i.e. in-in and anti-anti), then the fringe formed as we measure different free evolution times will have a 180 degree phase shift from the fringe formed when they are different (anti-in or in-anti). We will refer to in-in and anti-anti "In" and in-anti and anti-in "Anti". Now for each pair of shots, we can form our spin asymmetry,

$$\mathcal{A}_{u/l} = \frac{N_{\text{In}}^{u/l} - N_{\text{Anti}}^{u/l}}{N_{\text{In}}^{u/l} + N_{\text{Anti}}^{u/l}}, \quad (3.7)$$

where the subscript refers to the preparation and readout phase combination. For each Ramsey time and switch state we take n shots and so can form $n/2$ asymmetries for each of the upper and lower doublet. From these $n/2$ asymmetries we construct a mean $\mathcal{A}_{u/l}^m$ and a standard error $\delta\mathcal{A}_{u/l}$ which we use for fitting.

We can also form two 'meta' asymmetries by taking the difference (D) and sum (S) of the upper and lower asymmetries,

$$\begin{aligned} \mathcal{A}_D &= \mathcal{A}_u - \mathcal{A}_l, \\ \mathcal{A}_S &= \mathcal{A}_u + \mathcal{A}_l, \end{aligned} \quad (3.8)$$

with means $\mathcal{A}_D^m, \mathcal{A}_S^m$ and standard errors $\delta\mathcal{A}_D, \delta\mathcal{A}_S$. The $\delta\mathcal{A}_D$ are reduced compared to $\delta\mathcal{A}_S$ due to common-mode cancellation of many sources of noise.

3.4.2 Least squares fitting

As mentioned previously, we perform our Ramsey experiment simultaneously on both doublets and use their opposing orientations at the time of dissociation to read them out of opposing sides on the imaging MCP screen. Because the data are acquired simultaneously, the difference asymmetry allows us to cancel much of the common-mode noise, leaving us with doublet-odd data with less scatter than the raw data. Unfortunately, the doublets are not fully separated on the screen, so we must be careful with how we fit our data.

For a simple Ramsey fringe with no leakage from the other doublet, we can define a decaying sine wave for the asymmetry,

$$g_{u/l}(t_R) = C_{u/l} e^{-\gamma_{u/l} t_R} \cos(2\pi f_{u/l} t_R + \phi_{u/l}) + O_{u/l}. \quad (3.9)$$

Here C is the initial fringe contrast, γ is the contrast decay rate, f is the fringe frequency, t_R is the free evolution time, ϕ is the initial phase, O is the offset, and the subscripts indicate the upper or lower Stark doublet. In our fitting routine, we initially fit each fringe separately to this function. From the fit parameters we define the mean and difference parameters as

$$\begin{aligned} \alpha_m &= \frac{\alpha_u + \alpha_l}{2}, \\ \alpha_d &= \frac{\alpha_u - \alpha_l}{2}, \end{aligned} \quad (3.10)$$

with $\alpha \in \{C, \gamma, f, \phi, O\}$. Due to imperfect imaging contrast C_I , in reality each doublet's signal has a contribution from the other doublet. In this case, the measured asymmetries are

$$\begin{aligned} A_u &= \left(\frac{1+C_I}{2}\right)g_u + \left(\frac{1-C_I}{2}\right)g_l \\ A_l &= \left(\frac{1-C_I}{2}\right)g_u + \left(\frac{1+C_I}{2}\right)g_l \end{aligned} \quad (3.11)$$

Now our sum and difference asymmetries are

$$\begin{aligned} A_S &= g_u + g_l \\ A_D &= C_I(g_u - g_l), \end{aligned} \quad (3.12)$$

which we can express in terms of the mean and difference fitting parameters,

$$\begin{aligned}
A_S &= (C_m + C_d)e^{-2\gamma_m t_R} \cos(2\pi(f_m + f_d)t_R + (\phi_m + \phi_d)) + (O_m + O_d)... \\
&\quad + (C_m - C_d)e^{-2\gamma_d t_R} \cos(2\pi(f_m - f_d)t_R + (\phi_m - \phi_d)) + (O_m - O_d) \\
A_D &= C_I((C_m + C_d)e^{-2\gamma_m t_R} \cos(2\pi(f_m + f_d)t_R + (\phi_m + \phi_d)) + (O_m + O_d)... \\
&\quad - (C_m - C_d)e^{-2\gamma_d t_R} \cos(2\pi(f_m - f_d)t_R + (\phi_m - \phi_d)) - (O_m - O_d)),
\end{aligned} \tag{3.13}$$

We use these two expressions to perform a simultaneous least-squares fit for the sum and difference asymmetries in each experimental switch state. The value of C_I , the imaging contrast, is fixed at 0.9 for this fit—determined as described in Sec. 3.8.2.1. The parameter uncertainties are extracted based solely on the standard errors of the asymmetries used in the fits. The resultant uncertainties on the fitted values of f_d and ϕ_d are close to the shot noise limit and much smaller than those on f_m and ϕ_m thanks to our simultaneous data collection and fitting routine, which cancels most of the common-mode noise. The outputs of these simultaneous fits are used for all further analysis.

3.4.3 Switch parity channels

After fitting to each Ramsey fringe in a block to extract the 10 fitting parameters, we use the resulting 16 values of each parameter to form 16 linear combinations which we call switch-parity channels. The switch-parity channel for the mean and difference parameters α_m and α_d which are odd under the product of switches $[\tilde{\mathcal{S}}_a \tilde{\mathcal{S}}_b \dots]$ are given by

$$\begin{aligned}
\alpha^{\mathcal{S}_a \mathcal{S}_b \dots} &= \frac{1}{16} \sum_{\tilde{B}, \tilde{R}, \tilde{I} = \pm 1} [\tilde{\mathcal{S}}_a \tilde{\mathcal{S}}_b \dots] \alpha_m(\tilde{B}, \tilde{R}, \tilde{I}), \\
\alpha^{D\mathcal{S}_a \mathcal{S}_b \dots} &= \frac{1}{16} \sum_{\tilde{B}, \tilde{R}, \tilde{I} = \pm 1} [\tilde{\mathcal{S}}_a \tilde{\mathcal{S}}_b \dots] \alpha_d(\tilde{B}, \tilde{R}, \tilde{I}).
\end{aligned} \tag{3.14}$$

So for example, f^{BR} is formed by adding together the f_m measured in all switch states for which the product $\tilde{B}\tilde{R}$ is even, subtracting the f_m measured in all switch states for which the product $\tilde{B}\tilde{R}$ is odd and dividing by the total number of switch states. C^{DBRI} is formed from adding together the measured C_d in all switch states for which the product $\tilde{B}\tilde{R}\tilde{I}$ is even, subtracting all switch

states for which the product $\tilde{B}\tilde{R}\tilde{I}$ is odd and again, dividing by the number of switch states. The science signal is f^{BD} . All other parity channels allow us to diagnose experimental issues and identify sources of systematic error.

3.4.4 Blinding

We blinded the dataset by programming the fitting routine to add an unknown constant offset to the f^{BD} channel. This offset was stored in an encrypted file and was not removed until all statistical and systematic checks on the dataset had been completed, and the uncertainties finalised.

3.4.5 Data cuts

After completing the dataset, we applied cuts to the blinded data based on non-EDM channels indicating signal quality. Blocks with late-time contrast below 0.2 were cut due to low signal to noise. Figure 4.4 shows how the center value and error bar of the science channel vary as a function this cut. Blocks were also cut if they contained a fringe with a fitted difference frequency f_d in any switch state which was over 3.5σ different from the mean fringe frequency for that switch state. The mean fringe frequencies were calculated over all blocks which had the same value of \vec{B}^0 and were constructed from the linear combinations, including the blinded offset on f^{BD} . In total, the two cuts removed 42 blocks leaving 1328 blocks in the final dataset.

3.5 Effective Hamiltonian for doublets

To an excellent approximation, we can model the evolution of each of the Stark doublets as a two-level system. The effective Hamiltonian for each pair can be parameterized

$$H_{\text{eff}} = \frac{h}{2} \begin{pmatrix} f_0 & \Delta \\ \Delta & -f_0 \end{pmatrix}. \quad (3.15)$$

The diagonal f_0 components contain all terms which directly shift the energies of the two states relative to one another, whilst the off diagonal Δ components contain all terms which mix the

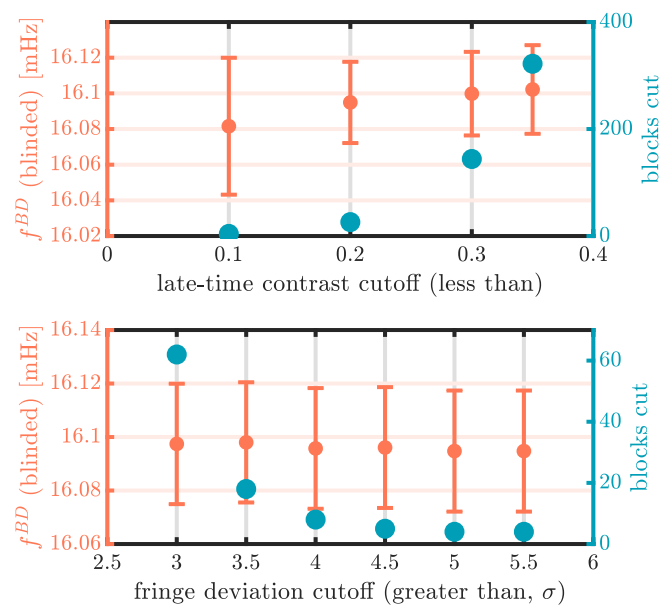


Figure 3.5: Results as we shift the threshold for applying cuts to the data. In each plot, only one cut is applied (the other cut is not applied). The cuts do not shift the resultant value of f^{BD} .

two states. We can expand both f_0 and Δ in terms of their leading contributions,

$$f_0 = \tilde{B}f_0^0 + \delta\mathcal{F}_0, \quad (3.16)$$

$$\Delta = \tilde{R}\Delta^0 + \tilde{R}\tilde{D}\Delta^D + \delta\mathcal{D}. \quad (3.17)$$

Here the quantities with tildes are equal to ± 1 and are determined by the experimental switch state, discussed in Sec. 3.3.9. The principle contribution to f_0 is the Zeeman splitting $f_0^0 = 3g_F\mu_B\mathcal{B}_{\text{rot}} \sim 100$ Hz. The off-diagonal component is dominated by two terms with similar magnitude, but different switch state dependence; Δ^0 and $\Delta^D \sim 1$ Hz arise at fourth order in perturbation theory in the full molecular Hamiltonian from the combined perturbations of rotation and Ω -doubling, breaking the degeneracy of the $m_F = \pm 3/2$ states in either Stark doublet at $\mathcal{B}_{\text{rot}} = 0$ [43, 45]. Δ and Δ^D are given by

$$\begin{aligned} h\Delta &= \frac{3\hbar\omega_{ef}}{2} \left(\frac{\hbar\omega_{\text{rot}}}{d_{\text{mf}}\mathcal{E}_{\text{rot}}} \right)^3 \left(\frac{18A_{\parallel}^2 - 19d_{\text{mf}}^2\mathcal{E}_{\text{rot}}^2}{A_{\parallel}^2 - d_{\text{mf}}^2\mathcal{E}_{\text{rot}}^2} \right) \\ h\Delta^D &= \frac{3\hbar\omega_{ef}}{2} \left(\frac{\hbar^3\omega_{\text{rot}}^3}{d_{\text{mf}}^2\mathcal{E}_{\text{rot}}^2 A_{\parallel}} \right) \left(\frac{9A_{\parallel}^2 - 8d_{\text{mf}}^2\mathcal{E}_{\text{rot}}^2}{A_{\parallel}^2 - d_{\text{mf}}^2\mathcal{E}_{\text{rot}}^2} \right). \end{aligned} \quad (3.18)$$

These expressions are valid as long as $d_{\text{mf}}\mathcal{E}_{\text{rot}} \gg \hbar\omega_{ef}$ and $d_{\text{mf}}\mathcal{E}_{\text{rot}} \gg \hbar\omega_{\text{rot}}$ or, in other words, if the molecule is fully polarized. The strong scaling of Δ with \mathcal{E}_{rot} allows us to perform off-resonant $\pi/2$ pulses by modulating the magnitude of \mathcal{E}_{rot} as described in Sec. 3.3.6.

The additional perturbations are given by

$$\delta\mathcal{F}_0 = \sum_{\tilde{s} \in \mathcal{W}} \tilde{s}\delta f_0^s, \quad (3.19)$$

$$\delta\mathcal{D} = \sum_{\tilde{s} \in \mathcal{W}} \tilde{s}\delta\Delta^s \quad (3.20)$$

where both summations are over \mathcal{W} , the set of all possible products of $\{\tilde{B}, \tilde{D}, \tilde{R}, \tilde{I}\}$, and the superscript s on the δf_0^s and $\delta\Delta^s$ denote the switch state dependence of the perturbation *relative* to the largest term in each matrix element, f_0^0 and Δ^0 respectively. For our purposes, the most important perturbation is that due to the electron EDM which contributes a diagonal term, $\tilde{D}\delta f_0^{BD} = 2\tilde{D}d_e\mathcal{E}_{\text{eff}}$. Some of largest diagonal perturbations are listed in Table WHAT. Those which are important for our determination of d_e are discussed in detail in Sec. 3.7.

For each experimental switch state and doublet, we measure a frequency $f(\tilde{B}, \tilde{R}, \tilde{I}, \tilde{D})$ corresponding to the energy difference between the two eigenstates, which we define to be always positive. For typical experimental parameters, f_0^0 is roughly two orders of magnitude larger than any other term in the Hamiltonian, and so we can expand f about f_0^0 ,

$$\begin{aligned}
f(\tilde{B}, \tilde{R}, \tilde{I}, \tilde{D}) &= \left| \tilde{B}f_0^0 + \delta\mathcal{F}_0 + \frac{\Delta^{0^2} + \Delta^{D^2} + 2\tilde{D}\Delta^0\Delta^D + 2\tilde{R}(\Delta^0 + \tilde{D}\Delta^D)\delta\mathcal{D} + \delta\mathcal{D}^2}{\tilde{B}f_0^0} \right. \\
&\quad \left. - \delta\mathcal{F}_0 \frac{\Delta^{0^2} + \Delta^{D^2} + 2\tilde{D}\Delta^0\Delta^D + 2\tilde{R}(\Delta^0 + \tilde{D}\Delta^D)\delta\mathcal{D} + \delta\mathcal{D}^2}{f_0^{0^2}} + \dots \right| \\
&= |f_0^0| + \tilde{B}\delta\mathcal{F}_0 + \frac{\Delta^{0^2} + \Delta^{D^2} + 2\tilde{D}\Delta^0\Delta^D + 2\tilde{R}(\Delta^0 + \tilde{D}\Delta^D)\delta\mathcal{D} + \delta\mathcal{D}^2}{|f_0^0|} \\
&\quad - \tilde{B}\delta\mathcal{F}_0 \frac{\Delta^{0^2} + \Delta^{D^2} + 2\tilde{D}\Delta^0\Delta^D + 2\tilde{R}(\Delta^0 + \tilde{D}\Delta^D)\delta\mathcal{D} + \delta\mathcal{D}^2}{|f_0^0|^2} + \dots
\end{aligned} \tag{3.21}$$

3.6 Accuracy evaluation

Accurate determination of the eEDM-induced energy shift in HfF^+ requires a careful evaluation of all possible systematic shifts. There are several methods one can employ to search for and understand systematics. First, one can change experimental parameters in a controlled way to look for shifts in the science channel. To see shifts with this method, one typically needs to change the parameter in question by a factor of 100, which is not always possible. Section 3.7.1.1 is a good example of doing this. Second, one can “break” the data analysis. This can be very precise since it can be performed on our entire dataset before unblinding, and it runs no risk of actually breaking the physical experiment. Section 3.8.2 contains several examples of this approach. Third, we can use other parity channels for insight into the eEDM channel as they can at times magnify certain effects. Section 3.7.2.4 is a good example of this approach. Finally, one can perform detailed analytical calculations or simulations to understand how experimental parameters might affect our measurement. Section 3.7.1 is an example of this approach.

Our systematics themselves can be divided into 2 categories: shifts in the actual precession frequency of the spin, and shifts in the phase at early or late free evolution time. Frequency shifts

are the dominant class of systematics in our experiment, and are covered in Section 3.7. Phase shifts tend to be smaller as they are suppressed due to our long coherence times and our ability to measure both the early-time and late-time phase (which is not possible in beam experiments). Phase shifts are covered in Section 3.8.

We include effects as systematics if they meet one of three criteria: first, if we saw a non-zero correlation between the eEDM channel and some parameter; second, if we saw a signal we thought would be important to understand but could not understand; and third, if we saw a correlation with the non-eEDM channel and some parameter but we know of a mechanism for that to leak into the eEDM channel. We set limits on systematics either by varying parameters over a large dynamic range and setting direct limits using the experiment, or by making conservative estimates using simulations.

3.7 Frequency shifts

3.7.1 Magnetic effects

Many experiments to measure EDMs are plagued by the interaction of the atom or molecule with stray magnetic fields. In our experiment, three features make us comparatively immune to these effects. First, the $^3\Delta_1$ state we use has an extremely small magnetic moment, roughly 300 times smaller than the magnetic moment of a bare electron. Second, our measurement is a differential one taken simultaneously between two pairs of states with almost identical magnetic moments—equal to roughly 1 part in 500. This allows us to measure magnetic effects by looking at the f^B channel whilst they cancel to a high degree in our measurement channel f^{BD} . Finally, the rotating quantisation axis in our experiment means that most effects considered in other experiments tend to average to zero over an integer number of rotation cycles. However, there are still some important magnetic effects which can shift our measured value of f^{BD} and we consider those in this section.

The interaction of the science states with a magnetic field \vec{B} is well described by the effective Zeeman Hamiltonian,

$$\begin{aligned} H_Z &= (\vec{\mu}_0 \pm \vec{\mu}_D) \cdot \vec{B}, \\ &= m_F \mu_B (g_F \pm \delta g_F) \hat{\mathcal{E}}_{\text{rot}} \cdot \vec{B}, \end{aligned} \quad (3.22)$$

where the \pm corresponds to the upper and lower doublet respectively. Here $\vec{\mu}_0$ is the part of the magnetic moment which is common to both doublets and $\vec{\mu}_D$ is differential part, typically 500 times smaller for our choice of experimental values. In the second line, we have made explicit the fact that both magnetic moments track the quantisation axis defined by the unit vector pointing along \mathcal{E}_{rot} . We choose to write the magnetic field as a sum of two parts, $\vec{B} = \tilde{B}\vec{B}^0 + \delta\vec{B}$ where the first part is due to the idealized applied quadrupole magnetic field gradient which reverses perfectly with the \tilde{B} switch, and the second part $\delta\vec{B}$ represents any additional magnetic field experienced by the molecules. In general, $\delta\vec{B}$ can be composed of many components with different dependence on the switch state and can be written

$$\delta\vec{B} = \sum_{\tilde{S} \in \mathcal{W}} \tilde{S} \delta\vec{B}^s, \quad (3.23)$$

where the summation is over \mathcal{W} , the set of all possible products of $\{\tilde{B}, \tilde{D}, \tilde{R}, \tilde{I}\}$, and the superscript s denotes the switch state dependence of the field relative to \vec{B}^0 . For example, the largest magnetic shifts measured in the experiment are caused by charging currents induced by the oscillating voltages on the radial electrodes. The charging currents produce a uniform magnetic field which rotates with $\vec{\mathcal{E}}_{\text{rot}}$ and whose sign depends on the rotation direction,

$$\tilde{R} \delta\vec{B}^{BR} = \tilde{R} \frac{\mu_0 \omega_{\text{rot}} C_{\text{eff}} V_{\text{rot}}}{2\pi R_{\text{eff}}} \hat{\mathcal{E}}_{\text{rot}}. \quad (3.24)$$

Here V_{rot} is the amplitude of the oscillating voltages on each of the fins, C_{eff} is the effective capacitance, and R_{eff} the effective radius of the trap. In our experiment $|\delta\vec{B}^{BR}| \sim 15 \mu\text{G}$.

We will be particularly interested in \vec{B}^B , the magnetic field which does not chop sign with

\vec{B} , causing shifts in f^B and f^{BD} . To first order³ these are

$$\begin{aligned} h\delta f^B &= 3g_F\mu_B\hat{\mathcal{E}} \cdot \delta\vec{\mathcal{B}}^B, \\ h\delta f^{BD} &= 3\delta g_F\mu_B\hat{\mathcal{E}} \cdot \delta\vec{\mathcal{B}}^B, \end{aligned} \quad (3.25)$$

where $\hat{\mathcal{E}}$ is the unit vector pointing along the total electric field $\vec{\mathcal{E}}$.

The differential magnetic moment $\vec{\mu}_D$ arises principally from two effects—mixing of neighbouring m_F levels in the rotating frame, and mixing of the $J = 1$ levels with $J = 2$ induced by the electric field. For the experimental parameters used throughout our dataset, the latter dominates and so we neglect the former here. The differential g-factor δg_F can then be calculated from second-order perturbation theory as,

$$\delta g_F \simeq \frac{3D_{\parallel}G_{\parallel}|\vec{\mathcal{E}}|}{20B}. \quad (3.26)$$

We see that, as expected, the mixing is linear in the size of the electric field. We can use this, in combination with Eq. 3.25 to give expressions for the time-averaged frequency shifts we measure in the experiment,

$$\begin{aligned} h\langle\delta f^B\rangle &= 3g_F\mu_B\langle\hat{\mathcal{E}} \cdot \delta\vec{\mathcal{B}}^B\rangle, \\ h\langle\delta f^{BD}\rangle &= 3\frac{\delta g_F}{\mathcal{E}_{\text{rot}}}\mu_B\langle\vec{\mathcal{E}} \cdot \delta\vec{\mathcal{B}}^B\rangle. \end{aligned} \quad (3.27)$$

We see that the common mode part of the Zeeman interaction is proportional to $\langle\hat{\mathcal{E}} \cdot \delta\vec{\mathcal{B}}^B\rangle$ whilst the differential part is proportional to $\langle\vec{\mathcal{E}} \cdot \delta\vec{\mathcal{B}}^B\rangle$.

Any possible $\vec{\mathcal{B}}^B$ can be expanded as

$$\delta\vec{\mathcal{B}}^B(\vec{x}, \vec{y}, \vec{z}) = \sum_{l=1,2,\dots} \sum_{m=-l}^l \mathcal{B}_{l,m} \vec{\nabla} r^l Y_{l,m}^* \quad (3.28)$$

³We note that there are higher order corrections to these shifts arising from rotation-induced mixing of the two states in a doublet. These effects are well understood but, under the experimental parameters used in the dataset, contribute less than 10% corrections to the shifts discussed and so are not included in our systematic uncertainty budget.

Table 3.4: Magnetic-field expansion

Coefficient	$Y_{l,m}^*$
Uniform fields	
$B_{1,-1}$	\hat{y}
$B_{1,0}$	\hat{z}
$B_{1,1}$	\hat{x}
First-order gradients	
$B_{2,-2}$	$\sqrt{3}(y\hat{x} + x\hat{y})$
$B_{2,-1}$	$\sqrt{3}(z\hat{y} + y\hat{z})$
$B_{2,0}$	$-x\hat{x} - y\hat{y} + 2z\hat{z}$
$B_{2,1}$	$\sqrt{3}(z\hat{x} + x\hat{z})$
$B_{2,2}$	$\sqrt{3}(x\hat{x} - y\hat{y})$
Second-order gradients	
$B_{3,-3}$	$\frac{3}{2}\sqrt{\frac{5}{2}}(2xy\hat{x} + (x - y)(x + y)\hat{y})$
$B_{3,-2}$	$\sqrt{15}(yz\hat{x} + xz\hat{y} + xy\hat{z})$
$B_{3,-1}$	$\sqrt{\frac{3}{2}}(-xy\hat{x} + \frac{1}{2}(-x^2 - 3y^2 + 4z^2)\hat{y} + 4yz\hat{z})$
$B_{3,0}$	$-3(xz\hat{x} + yz\hat{y} + \frac{1}{2}(x^2 + y^2 - 2z^2)\hat{z})$
$B_{3,1}$	$\sqrt{\frac{3}{2}}(\frac{1}{2}(-3x^2 - y^2 + 4z^2)\hat{x} - xy\hat{y} + 4xz\hat{z})$
$B_{3,2}$	$\sqrt{15}(xz\hat{x} - yz\hat{y} + \frac{1}{2}(x - y)(x + y)\hat{z})$
$B_{3,3}$	$\frac{3}{2}\sqrt{\frac{5}{2}}((x - y)(x + y)\hat{x} - 2xy\hat{y})$

where $\vec{x}, \vec{y}, \vec{z}$ are the ions position relative to the trap center, $\mathcal{B}_{l,m}$ is the coefficient of each component and $Y_{l,m}^*$ are the semi-normalized real spherical harmonics

$$Y_{l,m}^* = \sqrt{\frac{4\pi}{2l+1}} \times \begin{cases} \frac{i}{\sqrt{2}}(Y_l^m - (-1)^m Y_l^{-m}) & \text{if } m < 0 \\ Y_l^0 & \text{if } m = 0 \\ \frac{1}{\sqrt{2}}(Y_l^{-m} + (-1)^m Y_l^m) & \text{if } m > 0 \end{cases} \quad (3.29)$$

The $Y_{l,m}^*$ for $l \leq 3$ are given in Table 3.4.

3.7.1.1 Non-reversing $\mathcal{B}'_{\text{axgrad}}$

In an idealised version of the experiment where the ions are subject only to a perfect rotating electric field $\vec{\mathcal{E}}_{\text{rot}} = \mathcal{E}_{\text{rot}}(\cos(\omega t)\hat{x} + \sin(\omega t)\hat{y})$, the micromotion of the ions is $-\frac{e}{m\omega^2}\vec{\mathcal{E}}_{\text{rot}} = -r_{\text{rot}}\hat{\mathcal{E}}_{\text{rot}}$. In this case the only component of the magnetic field with $l \leq 3$ which causes a non-zero time-averaged shift in any frequency channel is the quadrupole magnetic field $\delta\vec{\mathcal{B}}^B = \mathcal{B}_{2,0}(-x\hat{x} - y\hat{y} + 2z\hat{z})$. As described in Section 4.2, we deliberately apply such a quadrupole magnetic field $\mathcal{B}'_{\text{axgrad}}$ whose sign reverses in the \tilde{B} chop which is the main contribution to f . However there can be an additional contribution which does not reverse sign with \tilde{B} , arising either from a background field present in the lab, or from imperfect reversal of the applied field. We have

$$\begin{aligned} h\langle\delta f^B\rangle &= -3g_F\mu_B\mathcal{B}_{2,0}r_{\text{rot}}, \\ h\langle\delta f^{BD}\rangle &= -3\delta g_F\mu_B\mathcal{B}_{2,0}r_{\text{rot}}. \end{aligned} \quad (3.30)$$

Both shifts are proportional to $\mathcal{B}_{2,0}$ but with the \tilde{D} -even shift being ~ 500 times larger. There are two possible approaches to removing the effect of this systematic from our measurement, the first is to make a correction to the measured value of f^{BD} based the measured value of f^B ,

$$\delta f_{\text{corr}}^{BD} = f^B \frac{\delta g_F}{g_F}. \quad (3.31)$$

The second is to shim out $\mathcal{B}_{2,0}$ by deliberately applying slightly different currents through the $\mathcal{B}'_{\text{axgrad}}$ coil in each direction. Both approaches are equivalent, we choose to shim f^B to zero




Figure 3.6: Measurement of $\frac{\delta g_F}{g_F}$, by varying the mean fringe frequency f^0 and measuring the differential fringe frequency f^D . Data taken at $\mathcal{E}_{\text{rot}} \sim 50$ V/cm.

on each block based on the measured value of f^B in the previous block. We then correct the measured value of f^{BD} on each block by the measured value of f^B on that block. The quantity $\frac{\delta g_F}{g_F}$ in Eq. 3.31 is determined directly from experiment as shown in Figure 3.6. The largest correction on any single block is 75 μHz , while the median correction size is 5 μHz . The average correction over the whole dataset is just 0.2 μHz .

This approach is effective in reducing the contribution of this systematic shift to levels well below our statistical error bar, but only holds if there are no other shifts in f^B that scale with f^{BD} with a constant of proportionality different from $\frac{\delta g_F}{g_F}$ —or even no shift in f^{BD} at all. The remainder of this section explores and places limits on such shifts arising from stray magnetic fields. Section 3.7.2 explores possible contributions from Berry’s phase effects.

3.7.1.2 General principles in determining other magnetic effects

Although no other magnetic field components couple directly to an idealised rotating electric field to give significant shifts, a number of potentially problematic shifts can arise when other contributions to the electric field experienced by the ions are taken into account. In particular, when the magnitude of the total electric field experienced by the ions is not constant in time,




Figure 3.7: Measurement of $\frac{\delta g_F}{g_F}$, obtained by introducing a large $\delta \mathcal{B}_{nr}$ and measuring the resultant values of f^{BD} and f^B . Note that f^{BD} is still blinded in this measurement. Data taken at $\mathcal{E}_{\text{rot}} \sim 50$ V/cm.

there can be effects for which $\langle \vec{\mathcal{E}} \cdot \delta \vec{\mathcal{B}}^B \rangle \neq \langle |\vec{\mathcal{E}}| \rangle \langle \hat{\mathcal{E}} \cdot \delta \vec{\mathcal{B}}^B \rangle$.

To enumerate possible shifts in the experiment, we analytically calculate the micromotion of a charged particle subject to an electric field $\vec{\mathcal{E}} = \vec{\mathcal{E}}_{\text{rot}} + \kappa \delta \vec{\mathcal{E}}$ where $\vec{\mathcal{E}}_{\text{rot}}$ is the electric field of the idealised rotating field, and $\delta \vec{\mathcal{E}}$ represents perturbations to this, discussed in the sections below. This micromotion allows us to write down the time dependent magnetic field experienced by a molecule moving through magnetic field imperfections $\kappa \delta \vec{\mathcal{B}}$. Finally we find the frequency shifts $\langle f^B \rangle$ and $\langle f^{BD} \rangle$ by calculating $\langle \hat{\mathcal{E}} \cdot \delta \vec{\mathcal{B}} \rangle$ and $\langle \vec{\mathcal{E}} \cdot \delta \vec{\mathcal{B}} \rangle$. Here the time average is over an integer number of periods of all relevant frequencies and we keep terms up to $\mathcal{O}(\kappa^2)$.

For those effects which give non-zero shifts in either channel, we constrain the maximum size of the shifts by direct measurement using the ions, or by auxiliary measurements of the size of the imperfections. We note that this analytical approach, by necessity, does not include the effects of spatially dependent electric fields and so misses the effects of ponderomotive forces such as those from the RF confining fields. To validate this approximation, we performed comprehensive numerical simulations including the full motion of the ions in the trap whilst varying the size of the imperfections listed over ranges well above their expected size in the experiment and observed no additional effects.

3.7.1.3 Second harmonic of \mathcal{E}_{rot} and transverse magnetic field

The rotating electric field in our experiment is generated by sinusoidally varying voltages on each of the radial electrodes. These voltages are driven by op-amps which inevitably suffer from harmonic distortion, adding higher harmonics of the input signal. Consider the effect of an additional electric field oscillating at the second-harmonic frequency. We have

$$\vec{\mathcal{E}} = \mathcal{E}_{\text{rot}} \begin{pmatrix} \cos(\omega t) \\ \sin(\omega t) \\ 0 \end{pmatrix} + \begin{pmatrix} \mathcal{E}_{2\text{hX}} \cos(2\omega t + \phi_{2\text{hX}}) \\ \mathcal{E}_{2\text{hY}} \cos(2\omega t + \phi_{2\text{hY}}) \\ 0 \end{pmatrix} \quad (3.32)$$

$$\begin{aligned} \hat{\mathcal{E}} &\simeq \begin{pmatrix} \cos(\omega t) \\ \sin(\omega t) \\ 0 \end{pmatrix} + \frac{1}{2\mathcal{E}_{\text{rot}}} \begin{pmatrix} \mathcal{E}_{2\text{hX}} \cos(2\omega t + \phi_{2\text{hX}}) \\ \mathcal{E}_{2\text{hY}} \cos(2\omega t + \phi_{2\text{hY}}) \\ 0 \end{pmatrix} \\ &+ \frac{\mathcal{E}_{2\text{hX}}}{4\mathcal{E}_{\text{rot}}} \begin{pmatrix} -\cos(\phi_{2\text{hX}}) \\ \sin(\phi_{2\text{hX}}) \\ 0 \end{pmatrix} + \frac{\mathcal{E}_{2\text{hY}}}{4\mathcal{E}_{\text{rot}}} \begin{pmatrix} \sin(\phi_{2\text{hY}}) \\ \cos(\phi_{2\text{hY}}) \\ 0 \end{pmatrix} \\ &+ \frac{\mathcal{E}_{2\text{hX}}}{4\mathcal{E}_{\text{rot}}} \begin{pmatrix} -\cos(4\omega t + \phi_{2\text{hX}}) \\ -\sin(4\omega t + \phi_{2\text{hX}}) \\ 0 \end{pmatrix} \\ &+ \frac{\mathcal{E}_{2\text{hY}}}{4\mathcal{E}_{\text{rot}}} \begin{pmatrix} -\sin(4\omega t + \phi_{2\text{hY}}) \\ \cos(4\omega t + \phi_{2\text{hY}}) \\ 0 \end{pmatrix} \end{aligned} \quad (3.33)$$

where $\hat{\mathcal{E}}$ has been expanded to first order in $1/\mathcal{E}_{\text{rot}}$. Alongside the components oscillating at ω , and the corrections oscillating at 2ω , $\hat{\mathcal{E}}$ also has terms which are time independent, and terms that oscillate at 4ω . The time-independent terms can couple to a uniform transverse magnetic

field $\delta\vec{\mathcal{B}}^B = \mathcal{B}_{1,1}\hat{x} + \mathcal{B}_{1,-1}\hat{y}$ to give time averaged shifts,

$$\begin{aligned} h\langle\delta f^B\rangle &= \frac{3g_F\mu_B}{4\mathcal{E}_{\text{rot}}}\left(-\mathcal{B}_{1,1}\mathcal{E}_{2hX}\cos(\phi_{2hX})\right. \\ &\quad + \mathcal{B}_{1,-1}\mathcal{E}_{2hY}\cos(\phi_{2hY}) \\ &\quad + \mathcal{B}_{1,-1}\mathcal{E}_{2hX}\sin(\phi_{2hX}) \\ &\quad \left.+ \mathcal{B}_{1,1}\mathcal{E}_{2hY}\sin(\phi_{2hY})\right), \end{aligned} \tag{3.34}$$

$$h\langle\delta f^{BD}\rangle = 0. \tag{3.35}$$

We note that depending on the details of the second harmonic, the phases ϕ_{2hX} and ϕ_{2hY} can change sign under rotation. In this way some of this shift can (and usually does) show up in f^{BR} rather than f^B . We have verified these shifts experimentally by intentionally applying a large second harmonic to our electrodes along with large transverse magnetic fields. Due to the correction applied described in Section 3.7.1.1, any shift in f^B not due to a non-reversing quadrupole magnetic field would cause a systematic $\delta f^{BD} = \frac{\delta g_F}{g_F}\langle\delta f^B\rangle$.

During the datarun, we measure the transverse field using an array of magnetometers positioned on the vacuum chamber and tune it to zero using three pairs of shim coils arranged around the apparatus. The magnetometers can be subject to small offsets and are calibrated at the beginning of the data run by intentionally applying large second harmonic electric fields and measuring shifts in f^B and f^{BR} . We used the same technique at regular (roughly weekly) intervals throughout the dataset to check the size of the magnetic fields and guard against any drifts in the magnetometer offsets. Based on these measurements, we conservatively estimate the maximum size of the transverse fields to be < 10 mG in x and y , limited principally by the precision of the current supplies used to drive the shim coils.

To reduce the size of the second harmonic electric fields present in the experiment, we intentionally apply a secondary signal at 2ω with the opposite phase, cancelling any signal at 2ω down to about a part in 40,000 of the signal at ω . In the worst case, where the relative phases of the residual 2ω signal are such that the corresponding electric field at the position of the ions is

maximised, this corresponds to residual shifts in f^B of about 80 mHz G^{-1} . After this procedure, we deliberately apply large magnetic fields to the ions and observe the residual shifts in f^B to be $103(1) \text{ mHz G}^{-1}$ and $46(9) \text{ mHz G}^{-1}$ for fields along x and y respectively, in good agreement with our direct voltage measurements. The amplitude of the second harmonic was checked periodically throughout the dataset to guard against any drifts, for example in the characteristics of the power operation amplifiers. Combined with the uncertainty on the transverse fields above, these measured gradients result in systematic uncertainty in f^B of 1 mHz with a corresponding systematic uncertainty in f^{BD} of $2 \text{ } \mu\text{Hz}$.

3.7.1.4 Higher-harmonics of \mathcal{E}_{rot}

We can generalize the arguments of the previous section to higher harmonics of \mathcal{E}_{rot} . For the n th harmonic, the expansion of $\hat{\mathcal{E}}$ in powers of $1/\mathcal{E}_{\text{rot}}$ will contain components which oscillate at $(n-2)\omega$ and $(n+2)\omega$. The $\vec{\mathcal{E}}_{\text{rot}}$ -induced micromotion of the ions combined with an m th order magnetic field gradient, produces a magnetic field in the frame of the ions which oscillates at $m\omega$. So in general an n th harmonic of \mathcal{E}_{rot} can couple to $(n-2)$ th and $(n+2)$ th order magnetic field gradients ($l = n-1$ or $n+3$) to give non-zero time-averaged frequency shifts. After the feedforward is applied to reduce the second-harmonic as described above, the next largest shift from harmonic distortion is due to 3ω , with an amplitude on each fin of about a part in 1500 of the amplitude at ω . The 3rd harmonic can couple to first-order field gradients of the form $\delta\vec{\mathcal{B}}^B = \mathcal{B}_{2,-2}\sqrt{3}(y\hat{x} + x\hat{y}) + \mathcal{B}_{2,2}\sqrt{3}(x\hat{x} - y\hat{y})$ to give shifts,

$$\begin{aligned} h\langle\delta f^B\rangle &= \frac{3\sqrt{3}g_F\mu_B e}{4m\omega^2} (\mathcal{B}_{2,2}\mathcal{E}_{3hX} \cos(\phi_{3hX}) \\ &\quad - \mathcal{B}_{2,-2}\mathcal{E}_{3hY} \cos(\phi_{3hY}) \\ &\quad - \mathcal{B}_{2,-2}\mathcal{E}_{3hX} \cos(\phi_{3hX}) \\ &\quad - \mathcal{B}_{2,2}\mathcal{E}_{3hY} \cos(\phi_{3hY})), \end{aligned} \tag{3.36}$$

$$h\langle\delta f^{BD}\rangle = 0. \tag{3.37}$$

We intentionally applied a large third harmonic electric field with various phases and observed shifts of ~ 20 mHz. The applied third harmonic field was 28 times larger than that present in the dataset and so we conservatively set a limit on the maximum size of any possible shift in f^B at 1 mHz, with a corresponding systematic uncertainty in f^{BD} of 2 μ Hz.

Higher harmonics are comparable to, or lower in amplitude (the fourth harmonic is down about a factor of 2/3 compared to the third, whilst the fifth is up a factor of about 4/3, sixth and higher harmonics are all smaller) and couple to higher-order magnetic field gradients. To set limits on their possible contributions to shifts in f^B we made measurements of the magnetic field gradients. These measurements were made outside of, but within a few cm of, the vacuum chamber. The electrodes in the experiment are machined from non-magnetic titanium and there are no other sources of magnetic fields closer to the ions than the steel vacuum chamber which is ~ 10 cm away. Because of this, we expect that the field gradients outside the chamber are similar to, or greater than, those inside the chamber. From these measurements we estimate the maximum size of first order field gradients to be 10 mG cm $^{-1}$ and second-order gradients to be 10 mG cm $^{-2}$. To determine the size of the magnetic field oscillating at $n\omega$, the n th order magnetic field gradient should be multiplied by r_{rot}^n where $r_{\text{rot}} = 0.05$ cm. This means we expect any higher-order harmonic contributions to shifts in f^B to be reduced by a factor of ~ 20 from those for the third harmonic, corresponding to systematic uncertainties in $f^{BD} < 100$ nHz. Therefore, we do not include any higher-order effects in our systematic uncertainty budget.

3.7.1.5 Ellipticity of \mathcal{E}_{rot}

Another possible electric field imperfection is an ellipticity of \mathcal{E}_{rot} (i.e. the oscillating voltage on different electrodes in the trap having slightly different amplitudes). A general rotating electric field having an ellipticity with its major axis orientated at angle θ to the x axis can be

written

$$\vec{\mathcal{E}} = \mathcal{E}_{\text{rot}} \begin{pmatrix} \cos(\omega t) \\ \sin(\omega t) \\ 0 \end{pmatrix} + \mathcal{E}_\epsilon \begin{pmatrix} \cos(2\theta - \omega t) \\ \sin(2\theta - \omega t) \\ 0 \end{pmatrix}. \quad (3.38)$$

The corresponding $\hat{\mathcal{E}}$ is

$$\begin{aligned} \hat{\mathcal{E}} = & \begin{pmatrix} \cos(\omega t) \\ \sin(\omega t) \\ 0 \end{pmatrix} + \frac{\mathcal{E}_\epsilon}{2\mathcal{E}_{\text{rot}}} \begin{pmatrix} \cos(2\theta - \omega t) \\ \sin(2\theta - \omega t) \\ 0 \end{pmatrix} \\ & + \frac{\mathcal{E}_\epsilon}{2\mathcal{E}_{\text{rot}}} \begin{pmatrix} -\cos(2\theta - 3\omega t) \\ \sin(2\theta - 3\omega t) \\ 0 \end{pmatrix}. \end{aligned} \quad (3.39)$$

This field can once again couple to first order magnetic field gradients of the form $\delta\vec{\mathcal{B}}^B = \mathcal{B}_{2,-2}\sqrt{3}(y\hat{x} + x\hat{y}) + \mathcal{B}_{2,2}\sqrt{3}(x\hat{x} - y\hat{y})$. After taking into account the modification to the ion micromotion caused by this perturbation, the resulting shifts are

$$h\langle\delta f^B\rangle = -\frac{9\sqrt{3}g_F\mu_{Be}\mathcal{E}_\epsilon}{2m\omega^2}(\mathcal{B}_{2,2}\cos(2\theta) + \mathcal{B}_{2,-2}\sin(2\theta)), \quad (3.40)$$

$$\begin{aligned} h\langle\delta f^{BD}\rangle = & -\frac{6\sqrt{3}\delta g_F\mu_{Be}\mathcal{E}_\epsilon}{m\omega^2}(\mathcal{B}_{2,2}\cos(2\theta) \\ & + \mathcal{B}_{2,-2}\sin(2\theta)). \end{aligned} \quad (3.41)$$

Note that in contrast to the effects of higher-harmonics discussed above, the ellipticity does cause a shift in f^{BD} , with $\delta f^{BD}/\delta f^B = \frac{4}{3}\frac{\delta g_F}{g_F}$. Due to the correction described in Sec. 3.7.1.1, which assumes $\delta f^{BD}/\delta f^B = \frac{\delta g_F}{g_F}$, the systematic associated with the shift described in this section will be 1/4 of the total shift induced in f^{BD} . The size of the ellipticity in our experiment was measured as described in Sec. 3.7.2. To set limits on the size of the shift, we applied an ellipticity 5 times larger than this, limited by our shallow ion trap during the Ramsey sequence, and varied its angle. The one sigma upper limit on the shifts seen in this way was 12 mHz and so we set an upper limit on the size of any shift in f^B due to this effect of 2.4 mHz, corresponding to a systematic uncertainty in f^{BD} of 1 μ Hz.

3.7.1.6 RF micromotion

In addition to the rotating electric field the ions are subject to the RF confining field oscillating at ω_{rf} . This RF field induces its own micromotion of the ions and, through $\mathcal{B}_{2,0}$, a magnetic field oscillating at ω_{rf} . The corresponding time-averaged shifts are

$$h\langle\delta f^B\rangle = 3g_F\mu_B \frac{e\mathcal{E}_{\text{rf}}^2\mathcal{B}_{2,0}}{4\mathcal{E}_{\text{rot}}m\omega_{\text{rf}}^2}, \quad (3.42)$$

$$h\langle\delta f^{BD}\rangle = 3\delta g_F\mu_B \frac{e\mathcal{E}_{\text{rf}}^2\mathcal{B}_{2,0}}{2\mathcal{E}_{\text{rot}}m\omega_{\text{rf}}^2}. \quad (3.43)$$

In this case we have $\langle\delta f^B\rangle/\langle\delta f^{BD}\rangle = 2\delta g_F/g_F$ and so there is potential for a systematic, but both shifts are proportional to $\mathcal{B}_{2,0}$ which we are shimming to zero. The ratio of the magnitude of the shift caused by the RF field to that caused by \mathcal{E}_{rot} is $\langle\mathcal{E}_{\text{rf}}^2\rangle\omega^2/\mathcal{E}_{\text{rot}}^2\omega_{\text{rf}}^2 \sim 10^{-2}$ so that this shift will just provide a $\sim 1\%$ correction to the—already very small—systematic uncertainty in Sec. 3.7.1.1. In addition, because our measurement of $\delta g_F/g_F$ was performed with the RF fields present, it already includes this correction. Therefore, we do not include any contribution from this effect in our systematic uncertainty budget. Other electric fields which oscillate at frequencies other than integer multiples of ω (e.g. time-averaged electric field due to secular motion) behave similarly but with even smaller contributions.

3.7.1.7 Out of plane electric fields

All the electric field imperfections discussed thus far are additional fields contained within the x, y -plane, the plane in which \mathcal{E}_{rot} rotates. However, it's also possible for the ions to experience oscillating fields in the z direction. This can occur, for example, due to thermal motion of the ions or patch charges displacing the center of the ion cloud from the geometrical centre of the trapping electrodes. The largest shifts from fields in the z direction are those due to $\mathcal{B}_{2,0}$ coupling to their own micromotion. The shifts have $\delta f^{BD}/\delta f^B = \delta g_F/g_F$ and so produce no systematic effects not already accounted for by the approach described in Sec. 3.7.1.1.

3.7.2 Berry's phase effects

The rotating electric bias field \mathcal{E}_{rot} defines our quantisation axis, and we do our spectroscopy in the rotating frame. Working in this rotating frame has two important effects: a mixing of the eigenstates of the system, which depends on the rotation rate; and an additional phase accumulation between eigenstates with different total angular momentum projections onto the quantization axis, which does not. This latter part, a geometric or Berry's phase—depending only on the path traced out by the quantisation axis in time—is discussed in this section, whilst the former is discussed in section 3.7.3.

As the electric field vector in the experiment rotates, it traces out closed loops in phase space. In each loop, the two states in a doublet accrue a differential geometric phase given by

$$\phi_{\text{geo}} = \Delta m_F \Omega = 3\Omega, \quad (3.44)$$

where Δm_F is the difference in the angular momentum projection of the two states and Ω is the solid angle traced out by the electric field. In the idealised experiment, the electric field rotates strictly in the X, Y -plane, subtending a solid angle of 2π every $T_{\text{rot}} = 1/f_{\text{rot}}$. Phase shifts of integer multiples of 2π are indistinguishable from zero and so not observable in the experiment. However effects which tilt the quantisation axis out the X, Y -plane can cause a non-zero frequency shift. For small tilts, the shift is given by

$$\delta f = -\frac{3f_{\text{rot}}}{2\pi} \int_0^{T_{\text{rot}}} \alpha(t) \dot{\phi}(t) dt \quad (3.45)$$

where $\alpha(t)$ is the tilt of the quantization axis out of the equatorial plane, $\dot{\phi}(t)$ is its azimuthal angular velocity, and the integral is over one period of the rotating field.

This frequency shift affects both doublets identically but is independent of the \tilde{B} chop and so appears in \tilde{B} -odd frequency channels. We note that the main contribution to $\dot{\phi}(t) \sim \tilde{R}\omega$ is rotation odd and so the most natural channel for the shift to appear in is f^{BR} , however rotation-odd components of $\alpha(t)$ can cause effects in f^B as well. Whilst the shifts do not appear directly

in our measurement channel, any shifts in f^B will cause a systematic shift—suppressed by a factor of $\delta g/g$ —due to the correction we make based on the measured value of f^B , described in Sec. 3.7.1.1.

The most straightforward way to tilt the quantisation axis is to add a time-independent electric field in the axial direction. Our measurement using trapped ions is well protected from this possibility; any electric field in the axial direction which does *not* time average to zero, *and* which is not balanced by another force on the ions, will cause the mean position of the ions to shift to where the time-averaged electric field *is* zero. However, there are a number of mechanisms which can cause non-zero time-averaged geometric phases in the experiment. In this section, we discuss these mechanisms and place constraints on their possible size during our measurement.

3.7.2.1 Gravity

The gravitational force on the ions causes them to sit at a position in the trap where the axial electric field is slightly different from zero. This field is $\mathcal{E}_{\text{grav}} = mg/e \sim 20 \mu\text{V m}^{-1}$, causing a shift $\delta f = 3f_{\text{rot}}\mathcal{E}_{\text{grav}}/\mathcal{E}_{\text{rot}} \sim 4 \text{ mHz}$. This shift appears rigorously in f^{BR} which we do not use for correction of f^{BD} and so we do not include any contribution from this effect in our systematic uncertainty budget.

3.7.2.2 AC Stark shift from vibrational cleanup light

The vibrational cleanup light is tuned to resonance with the ${}^3\Delta_1(v=1) \rightarrow {}^3\Sigma_{0+}(v=0)$ transition to remove any population in ${}^3\Delta_1(v=1)$ and prevent it decaying into the science state. The light is left on throughout the free evolution time and propagates in the axial direction through the chamber. To estimate the effect of this light, we assume that the polarisability of the ${}^3\Delta_1(v=0)$ state is dominated by interaction with ${}^3\Sigma_{0+}(v=0)$, from which the laser is detuned $\delta_L \sim 1.7 \text{ THz}$, and the nearby ${}^3\Sigma_1(v=0)$, a further $\sim 30 \text{ THz}$ away. Conservatively assuming transition matrix elements of $d_{\text{trans}} \sim 0.1 ea_0$ for each, we obtain a polarisabil-

ity of $d_{\text{trans}}^2 / (h^2 \epsilon_0 c \delta_L) \sim 2 \times 10^{-4} \text{ Hz W}^{-1} \text{ m}^2$. The light intensity at the position of the ions is $\sim 300 \text{ W m}^{-2}$, giving an AC Stark shift of order $\sim 1 \text{ Hz}$.

The scalar AC Stark shift affects all states in ${}^3\Delta_1(v=0)$ equally and so has no effect on our measurement. However, a gradient of the intensity could exert a force on the ions, pushing them to a position where the electric field is non-zero in a similar way to the gravitational effect described above. The light is roughly collimated through the chamber and we conservatively constrain the intensity gradient to be $< 500 \text{ W m}^{-3}$, corresponding to a force more than 10 orders of magnitude smaller than that due to gravity discussed in section 3.7.2.1.

Effects from possible vector or tensor AC stark shifts depend on the polarization of the light which is nominally linear in the X, Y -plane in our experiment. The tensor shift causes an AC stark shift which is differential between doublets, but common mode within a doublet, but which oscillates at $2f_{\text{rot}}$, imitating an ellipticity of \mathcal{E}_{rot} . Conservatively assuming the tensor polarisability is of similar size to the scalar, this effective ellipticity is $h \sim 1 \text{ Hz} / \mathcal{E}_{\text{rot}} d_{\text{mf}} \sim 10^{-8}$, much smaller than our best limit on the true ellipticity in the trap.

For linearly polarised light, the vector Stark shift is zero but we can consider the effect of a residual handedness. In this case, the vector shift can tilt the quantisation axis out of the X, Y -plane. Conservatively assuming vector polarisability similar size to the scalar, and a 2% residual handedness of the light, we find a geometric phase of $\sim 1 \text{ mHz}$. Like the gravitational effect, this shift is in f^{BR} and so has no systematic effect on our measurement. We took some auxiliary blocks with the vibrational cleanup light turned off during the free evolution time and observed no shifts, constraining any shift in f^{BR} to less than 9 mHz.

3.7.2.3 Phase modulation due to axial secular motion

Ions displaced from the center of the trap experience a non-zero axial confining field. Any axial motion due to the axial confining potential is harmonic, so the axial component of the

electric field will average out to zero over many cycles. *Instantaneously*, however, the ions will accrue a Berry's phase which we could then alias, leading to a perceived phase shift.

We investigated by pulsing the endcap electrode to 'kick' the ion cloud in the axial direction right before the first $\pi/2$ pulse; causing temporary coherent axial oscillations, then collecting a Ramsey fringe. Example data is shown in Fig. 3.8. We repeated this procedure over a range of kick amplitudes. For each Ramsey fringe we fit the data to the model

$$A \cos(2\pi ft + \phi + \delta \cos(2\pi f_z t + \phi_z)). \quad (3.46)$$

Figure 3.9 shows the amplitude of the phase modulation δ , plotted against the amplitude of the applied kick. The data constrain $\delta < 0.02$ radians at the 1σ confidence level. Assuming the worst case scenario where we take our data at the maximum of the phase modulation, this corresponds to an apparent frequency shift of < 1.4 mHz for the effective average late Ramsey time (~ 2.3 seconds) used. This shift is likely to be in f^{BR} but could leak into f^B if the phase of the axial oscillation were rotation odd. We did not monitor this during our dataset and so include an systematic uncertainty of $3 \mu\text{Hz}$ in our uncertainty budget.

While this phase shift is negligible at late time—we find any slosh decoheres in 7 ms—the integral of the phase shift could be non-zero. We explored this numerically and found a frequency shift of f^{BD} on the order of 100 pHz, which is completely negligible.

3.7.2.4 Axial 2nd harmonic + radial ellipticity

In addition to static tilts of the quantisation axis and aliasing of time-varying tilts, perturbations to the radial electric field can cause variations in the magnitude of the rotating field and its azimuthal angular velocity which can couple to axial displacements to give geometric phase effects which do not time-average to zero. An important example of this type of effect is an ellipticity of \mathcal{E}_{rot} combined with a second-harmonic field in the axial direction. Such an axial field could result either from leakage of the signals that drive the radial electrodes onto the endcaps,




Figure 3.8: Example axial oscillation-induced phase modulation data, for a both large kick (color) and a small kick (grey).




Figure 3.9: Phase modulation amplitude for various sizes of z kick. Upper doublet data is in blue and lower doublet data is in pink. Both fits have an intercept of $-0.0127(0.0108)$.

or from axial displacement of the ions from the geometric center of the trap combined with fields generated by the radial electrodes themselves.

We can express a radial ellipticity in \mathcal{E}_{rot} as we did before (Eq. 3.38) but now explicitly including rotational dependence,

$$\begin{aligned} \vec{\mathcal{E}} = & \mathcal{E}_{\text{rot}} \begin{pmatrix} \cos(\omega t + \phi_R + \tilde{R}\phi_0) \\ \tilde{R} \sin(\omega t + \phi_R + \tilde{R}\phi_0) \\ 0 \end{pmatrix} \\ & + \mathcal{E}_\epsilon \begin{pmatrix} \cos(2\tilde{R}\theta - \omega t - \phi_R - \tilde{R}\phi_0) \\ \tilde{R} \sin(2\tilde{R}\theta - \omega t - \phi_R - \tilde{R}\phi_0) \\ 0 \end{pmatrix}. \end{aligned} \quad (3.47)$$

Here $\phi_0 + \tilde{R}\phi_R$ is the angle of the radial electric field to the X -axis at $t = 0$ and θ is the angle between the major axis of the ellipse and the X -axis. An axial field oscillating at the second harmonic is given by

$$\vec{\mathcal{E}}_{2hZ} = \mathcal{E}_{2hZ} \begin{pmatrix} 0 \\ 0 \\ \cos(2\omega t + \phi_{2hZ}) \end{pmatrix}. \quad (3.48)$$

With these electric fields we find the tilt angle and azimuthal angular velocity are

$$\begin{aligned} \alpha(t) \simeq & \frac{\mathcal{E}_{2hZ}}{\mathcal{E}_{\text{rot}}} \cos(2\omega t + 2\phi_{2fZ}) \\ & \times \left(1 - \frac{\mathcal{E}_\epsilon}{\mathcal{E}_{\text{rot}}} \cos(2\omega t + 2\phi_R + 2\tilde{R}\phi_0 - 2\tilde{R}\theta) \right), \end{aligned} \quad (3.49)$$

$$\dot{\phi}(t) \simeq \tilde{R}\omega - \frac{2\tilde{R}\mathcal{E}_\epsilon\omega}{\mathcal{E}_{\text{rot}}} \cos(2\omega t + 2\phi_R + 2\tilde{R}\phi_0 - 2\tilde{R}\theta), \quad (3.50)$$

where in each case we have expanded to first order in $\mathcal{E}_\epsilon/\mathcal{E}_{\text{rot}}$. Using Eq. 3.45, we find

$$\delta f \simeq \frac{9\tilde{R}\mathcal{E}_{2hZ}\mathcal{E}_\epsilon f_{\text{rot}}}{2\mathcal{E}_{\text{rot}}^2} \cos(\phi_{2fZ} + 2\tilde{R}\theta - 2\tilde{R}\phi_0 - 2\phi_R). \quad (3.51)$$

Finally we can calculate the expected shifts in f^B and f^{BR} by taking the sum and difference of




Figure 3.10: Example data collected with $\mathcal{E}_{2fZ} \sim 4 \text{ V m}^{-1}$ and an ellipticity $\mathcal{E}_e/\mathcal{E}_{\text{rot}} \sim 5 \times 10^{-3}$. As we vary the phase of the axial modulation, f^B and f^{BR} vary sinusoidally with a $\pi/2$ phase offset, as expected.

the shifts in the two rotation directions,

$$\delta f^B = -\frac{9\mathcal{E}_{2hZ}\mathcal{E}_e f_{\text{rot}}}{2\mathcal{E}_{\text{rot}}^2} \sin(2\theta - 2\phi_0) \sin(\phi_{2fZ} - 2\phi_R), \quad (3.52)$$

$$\delta f^{BR} = \frac{9\mathcal{E}_{2hZ}\mathcal{E}_e f_{\text{rot}}}{2\mathcal{E}_{\text{rot}}^2} \cos(2\theta - 2\phi_0) \cos(\phi_{2fZ} - 2\phi_R). \quad (3.53)$$

We see that we expect frequency shifts which vary sinusoidally as a function of either θ or ϕ_{2fZ} , with a $\pi/2$ phase shift between f^B and f^{BR} . Figure 3.10 shows example data for deliberately applied ellipticity $\mathcal{E}_e/\mathcal{E}_{\text{rot}} \sim 5 \times 10^{-3}$ and an axial second-harmonic $\mathcal{E}_{2fZ} \sim 4 \text{ V m}^{-1}$.

We used data like this to manually shim out any residual ellipticity by applying \mathcal{E}_{2fZ} , and adjusting the \mathcal{E}_{rot} multiplying DAC on each electrode until we saw the effect was minimized. We then measured the residual ellipticity in the same way and found $\mathcal{E}_e/\mathcal{E}_{\text{rot}} = 3 \times 10^{-4}$, suppressed by a factor of 35. We constrain the axial second harmonic field by deliberately applying an ellipticity and varying θ to constrain $|\mathcal{E}_{2fZ}| < 10 \text{ mV m}^{-1}$ at the 1σ confidence level, suppressed by a factor 400. Combining these suppression factors allows us to place a limit on the maximum size of any shift in f^B , $\delta f^B < 0.8 \text{ mHz}$. We include a corresponding contribution to our systematic uncertainty of $2 \mu\text{Hz}$.

3.7.2.5 Effects of higher harmonics

We can generalize the discussion of the previous section to a pair of perturbations $\vec{\delta\mathcal{E}}_{X,Y}$ and $\vec{\delta\mathcal{E}}_Z$ which act in radial and axial direction respectively, and which are both phase locked to \mathcal{E}_{rot} . We can write general expressions for these perturbations,

$$\begin{aligned}\vec{\delta\mathcal{E}}_{X,Y} &= \mathcal{E}_{nhX} \cos(n\omega t + \phi_{nhX}) \hat{X} \\ &\quad + \mathcal{E}_{nhY} \cos(n\omega t + \phi_{nhY}) \hat{Y},\end{aligned}\tag{3.54}$$

$$\vec{\delta\mathcal{E}}_Z = \mathcal{E}_{mhZ} \cos(m\omega t + \phi_{mhZ})\tag{3.55}$$

where m, n are both integer. The combination of these perturbations will result in a non-zero time-averaged frequency shift at first order provided m and n differ by ± 1 . The shift is given by

$$\begin{aligned}\delta f &= \mp \frac{3\mathcal{E}_{mhZ}f_{\text{rot}}(n \pm 2)}{4\mathcal{E}_{\text{rot}}^2} (\mathcal{E}_{nhX} \cos(\phi_{nhX} - \phi_{mhZ}) \\ &\quad \pm \mathcal{E}_{nhY} \sin(\phi_{nhY} - \phi_{mhZ})),\end{aligned}\tag{3.56}$$

where the \pm correspond to $m = n \pm 1$. Depending on the \vec{R} dependence of the various phases, this shift could appear in either f^B or f^{BR} . We note that the frequency shift corresponding to the first harmonic on the endcap and second harmonic radially, $m = 1, n = 2$, is zero, and so the next largest shifts are expected to come from effects involving the third harmonic or higher.

To place constraints on the possible size of these effects, we measured the Fourier transform of the signal on each of the radial electrodes and the two endcaps. The measured amplitudes of each harmonic are shown in the second column of Table 3.5. We first consider the possible radial electric fields caused by these signals. The radial electric fields at the center of the trap generated by an $n\omega$ voltage signal on each electrode are

$$\delta\mathcal{E}_{Xnh} = - \sum_{k=1}^8 \frac{V_{nhk}}{R_{X,Y}} \cos\left(\left(9 - 2k\right)\frac{\pi}{8}\right) \cos(n\omega t + \phi_{nhk}),\tag{3.57}$$

$$\delta\mathcal{E}_{Ynh} = \sum_{k=1}^8 \frac{V_{nhk}}{R_{X,Y}} \sin\left(\left(9 - 2k\right)\frac{\pi}{8}\right) \cos(n\omega t + \phi_{nhk}),\tag{3.58}$$

where V_{nhk} and ϕ_{nhk} are amplitude and phase of the voltage on each of the 8 electrodes and $R_{X,Y} \simeq 23$ cm is constant relating to the geometry of the trap. The amplitude of these fields is

given by

$$|\delta\mathcal{E}_{Xnh}| = \left[\left(\sum_{k=1}^8 \frac{V_{nhk}}{R_{X,Y}} \cos\left((9-2k)\frac{\pi}{8}\right) \cos(\phi_{nhk}) \right)^2 + \left(\sum_{k=1}^8 \frac{V_{nhk}}{R_{X,Y}} \cos\left((9-2k)\frac{\pi}{8}\right) \sin(\phi_{nhk}) \right)^2 \right]^{1/2}, \quad (3.59)$$

$$|\delta\mathcal{E}_{Ynh}| = \left[\left(\sum_{k=1}^8 \frac{V_{nhk}}{R_{X,Y}} \sin\left((9-2k)\frac{\pi}{8}\right) \cos(\phi_{nhk}) \right)^2 + \left(\sum_{k=1}^8 \frac{V_{nhk}}{R_{X,Y}} \sin\left((9-2k)\frac{\pi}{8}\right) \sin(\phi_{nhk}) \right)^2 \right]^{1/2}. \quad (3.60)$$

The amplifiers driving each of the electrodes are nominally identical and so the harmonic distortion is approximately equal on each electrode. As a first order approximation, we take the amplitude of the signal on each electrode to be equal, $V_{nhk} = V_{nh}$, and the phases to be locked to the fundamental with a small offset that can be different on each electrode, $\phi_{nhk} = \phi_{nh0} + n(9-2k)\frac{\pi}{8} + \delta\phi_{nhk}$.

We first consider $\delta\phi_{nhk} = 0$. In this case all the radial fields from all electrodes cancel one another unless n differs from a multiple of 8 by ± 1 . When n is one greater than a multiple of 8, the relative phases on neighbouring electrodes are the same as for the fundamental and so the harmonic field corotates with \mathcal{E}_{rot} . For n one less than a multiple of 8, the relative phases on neighbouring electrodes are flipped compared to the fundamental and the resulting field counterrotates. In both cases the amplitude is $4V_{nh}/R_{X,Y}$ and we use this as an estimate of the field size for these values of n in the third column of Table 3.5. For the other harmonics we can estimate their size by looking at the variance of their amplitude due to the variance of the amplifier phases,

$$\langle |\delta\mathcal{E}_{Xnh}|^2 \rangle = \sum_{k=1}^8 \left(\frac{\partial |\delta\mathcal{E}_{Xnh}|}{\partial \delta\phi_{nhk}} \right)^2 \langle \delta\phi_{nhk}^2 \rangle, \quad (3.61)$$

$$\langle |\delta\mathcal{E}_{Ynh}|^2 \rangle = \sum_{k=1}^8 \left(\frac{\partial |\delta\mathcal{E}_{Ynh}|}{\partial \delta\phi_{nhk}} \right)^2 \langle \delta\phi_{nhk}^2 \rangle, \quad (3.62)$$

where the derivatives in brackets are evaluated at $\delta\phi_{nhk} = 0$. Since in these cases the phase relationship between the fields is random, we will be interested in their quadrature sum. After

Table 3.5: Constraints on electric fields from higher-harmonic voltages on radial electrodes and endcaps.

n	Amplitudes (mV)		Fields (mV m ⁻¹)	
	Radial	Axial	$ \delta\mathcal{E}_{X,Ynh} $	$ \delta\mathcal{E}_{Znh} $
2	11	0.5	250	1.0
3	310	0.3	1296	2.1
4	190	0.1	649	1.8
5	410	0.2	1715	2.7
6	100	0.4	341	0.8
7	230	0.5	4000	1.7
8	45	0.7	154	3.1
9	170	1.0	2957	1.8
10	37	1.4	126	2.0
11	140	1.8	586	2.7
12	42	2.1	143	3.1

some lengthy algebra, we find

$$\begin{aligned}
\langle |\delta\mathcal{E}_{X,Ynh}|^2 \rangle^{1/2} &= (\langle |\delta\mathcal{E}_{Xnh}|^2 \rangle + \langle |\delta\mathcal{E}_{Ynh}|^2 \rangle)^{1/2} \\
&= \frac{V_{nh}}{R_{X,Y}} \left[4 + (-1)^n \sqrt{2} \cos\left(\frac{n\pi}{4}\right) \right. \\
&\quad \left. - (-1)^n \sqrt{2} \cos\left(\frac{3n\pi}{4}\right) \right]^{1/2} \langle \delta\phi_{nh}^2 \rangle^{1/2},
\end{aligned} \tag{3.63}$$

where we have set the variances of the phases on all amplifiers to be equal, $\langle \delta\phi_{nhk}^2 \rangle = \langle \delta\phi_{nh}^2 \rangle$. From measurements of the signals on the electrodes, we conservatively estimate $\langle \delta\phi_{nh}^2 \rangle = (\pi/8)^2$. The third column of Table 3.5 shows the estimated size of the radial electric fields from higher harmonics calculated using this expression. Note that, because we shim the 2nd harmonic, we can make no such claims about relative phases on electrodes and so we assume the worst case of all electrodes on one side of the trap in phase and exactly out of phase with the electrodes on the other side, giving $\delta\mathcal{E}_{X,Y2h} = 4\sqrt{1 + \frac{1}{\sqrt{2}}V_{2h}/R_{X,Y}}$.

We can take a similar approach to estimating the possible higher harmonic fields in the Z direction. The most direct way of generating axial electric fields with angular frequency $n\omega$ is for them to appear directly on the endcaps of our trap where they generate a field V_{nhZ}/R_Z^{ax} ,

where $R_Z^{ax} \simeq 137$ cm. Our measurements of possible higher-harmonic signals onto the endcaps are shown in the third column of Table 3.5; for $n > 3$ these are upper limits, constrained by the frequency-dependent noise floor of our FFT measurements. We find that in most cases the fields generated by these voltages are dominated by another contribution, signals on the radial electrodes combined with axial displacement of the ions from the center of the trap. We can write the axial electric field due to $n\omega$ signals on the radial electrodes at axial displacement Z_0 as

$$\delta\mathcal{E}_{Znh} = \sum_{k=1}^8 \frac{V_{nhk}Z_0}{R_Z^2} \cos(n\omega t + \phi_{nhk}), \quad (3.64)$$

with amplitude

$$|\delta\mathcal{E}_{Znh}| = \frac{Z_0}{R_Z^2} \left(\left(\sum_{k=1}^8 V_{nhk} \cos(\phi_{nhk}) \right)^2 + \left(\sum_{k=1}^8 V_{nhk} \sin(\phi_{nhk}) \right)^2 \right)^{1/2}. \quad (3.65)$$

The constant $R_Z \simeq 27$ cm is fixed by the trap geometry. We measure the axial displacement by deliberately applying an ellipticity plus a large second harmonic in phase to all radial electrodes and find $Z_0 = 0.6$ mm.

Again, we make the substitutions $V_{nhk} = V_{nh}$ and $\phi_{nhk} = \phi_{nh0} + n(9 - 2k)\frac{\pi}{8} + \delta\phi_{nhk}$ and start by assuming $\delta\phi_{nhk} = 0$,

$$\delta\mathcal{E}_{Znh} = \frac{V_{nh}Z_0}{R_Z^2} \frac{\sin(n\pi)}{\sin\left(\frac{n\pi}{8}\right)} \cos(n\omega t + \phi_{nh0}). \quad (3.66)$$

The contribution of the 8 electrodes cancels out except for when n is a multiple of 8 when they are all in phase with one another and we have $|\delta\mathcal{E}_{Znh}| = 8V_{nh}Z_0/R_Z^2$. For these values of n , we use this as an estimate of the axial field amplitude from the radial electrodes. For the other possible harmonics we can estimate their amplitudes in the same way we did for the radial fields. The variance of the amplitude of the field is given by

$$\langle |\delta\mathcal{E}_{Znh}|^2 \rangle = \sum_{k=1}^8 \left(\frac{\partial |\delta\mathcal{E}_{Znh}|}{\partial \delta\phi_{nhk}} \right)^2 \langle \delta\phi_{nhk}^2 \rangle. \quad (3.67)$$

Table 3.6: Constraints on Berry’s phase frequency contributions to f^B or f^{BR} due to possible combinations of radial and axial field imperfections from harmonic distortion in amplifiers used to drive radial electrodes. All entries are in μHz .

$n_{X,Y}$	n_Z										
	2	3	4	5	6	7	8	9	10	11	12
2	0	17	0	0	0	0	0	0	0	0	0
3	11	0	95	0	0	0	0	0	0	0	0
4	0	22	0	87	0	0	0	0	0	0	0
5	0	0	75	0	83	0	0	0	0	0	0
6	0	0	0	31	0	38	0	0	0	0	0
7	0	0	0	0	139	0	947	0	0	0	0
8	0	0	0	0	0	13	0	24	0	0	0
9	0	0	0	0	0	0	545	0	556	0	0
10	0	0	0	0	0	0	0	16	0	35	0
11	0	0	0	0	0	0	0	0	90	0	200
12	0	0	0	0	0	0	0	0	0	33	0

where again, the derivative in brackets is evaluated at $\delta\phi_{nhk} = 0$. After some more lengthy algebra, we find

$$\langle |\delta\mathcal{E}_{Znh}|^2 \rangle^{1/2} = \frac{V_{nh}Z_0}{R_Z^2} \left(4 - \frac{\sin(n\pi)}{\sin(\frac{n\pi}{4})} \right)^{1/2} \langle \delta\phi_{nh}^2 \rangle^{1/2}. \quad (3.68)$$

We use this expression with $\langle \delta\phi_{nh}^2 \rangle^{1/2} = \pi/8$ to calculate the field estimates for all n which are not a multiple of 8. Note that, again due to our shimming of the second harmonic, we can’t make any claims about relative phases on the electrodes and so we assume the worst case of all in phase giving $|\delta\mathcal{E}_{Z2h}| = 8V_{2h}Z_0/R_Z^2$. The fourth column of Table 3.5 shows the sum in quadrature of our estimates for the axial field from the radial electrodes with those from the endcaps.

In Table 3.6, we combine these estimates of the radial and axial field amplitudes with Eq. 3.56 to set limits on the Berry’s phase from each possible combination of higher-harmonics. Summing all the contributions in quadrature we get an uncertainty on the Berry’s phase of 1.4 mHz, corresponding to a systematic uncertainty in f^{BD} of 3 μHz which we include in our systematic uncertainty budget.

3.7.3 Residual rotation induced doublet mixing

The systematic effects we have considered so far are all concerned with the second term in Eq. 3.21, various sources of δf_0^s . However, the third and fourth terms in the expression, containing the off-diagonal components of the effective Hamiltonian, can also potentially cause systematic shifts. The third term concerns shifts directly generated by the off-diagonal mixing components, while the fourth term concerns the leaking of frequency shifts into other channels due to this mixing. In this section we discuss possible contributions to each and place constraints on their size during the measurement.

3.7.3.1 Leaking of f^{BR}

The largest leak of any frequency channel into f^{BD} is expected to come from the frequency channel with the largest signal. This is f^{BR} where a signal of ~ 200 mHz is caused by the charging currents in our electrodes. This could potentially leak into f^{BD} when combined with a non-zero values of $\delta\Delta^{DR}$ or $\delta\Delta^R$,

$$\delta f^{BD} = -f^{BR} \frac{2\Delta^0 \delta\Delta^{DR} + 2\Delta^D \delta\Delta^R}{|f_0^0|^2}. \quad (3.69)$$

Non-zero values of $\delta\Delta^{DR}$ and $\delta\Delta^R$ can come from an axial magnetic field through a 4th order coupling similar to Δ and Δ^D but where one of the rotation matrix elements is replaced with a magnetic field matrix element. Any effect would also result in a contribution to f^{DR} ,

$$\delta f^{DR} = \frac{2\Delta^0 \delta\Delta^{DR} + 2\Delta^D \delta\Delta^R}{|f_0^0|}. \quad (3.70)$$

So we can rewrite Eq. 3.69 as

$$\delta f^{BD} \simeq -\frac{f^{BR}}{f_0^0} \left(f^{DR} - \frac{\delta g_F}{g_F} f^R \right). \quad (3.71)$$

Our data places a limit on this contribution to f^{BD} of < 170 nHz and so we include no contribution in our systematic uncertainty budget.

On similar grounds, we could also expect non-zero $\delta\Delta^{DR}$ or $\delta\Delta^R$ to cause f^{BR} to leak into

f^B ,

$$\delta f^B = -f^{BR} \frac{2\Delta^0 \delta\Delta^R + 2\Delta^D \delta\Delta^{DR}}{|f_0^0|^2}, \quad (3.72)$$

causing a systematic via the correction described in Sec. 3.7.1.1. Since this expression involves the same terms as Eq. 3.69, barring unlikely cancellations, the contribution can be expected to be of similar size. The corresponding systematic in f^{BD} is suppressed by a further factor of $\delta g_F/g_F$ and so we again include no contribution in our systematic uncertainty budget.

3.7.3.2 Axial magnetic fields

Analogously to how a static axial magnetic field can generate $\delta\Delta^{DR}$ and $\delta\Delta^R$, a \tilde{B} -odd axial magnetic field \mathcal{B}_Z^B can cause $\delta\Delta^{BDR}$ and $\delta\Delta^{BR}$. Such a field could be generated from the $\mathcal{B}'_{\text{axgrad}}$ coils if the ions are situated slightly away from their geometric center. The combination of these two effects can cause a shift in f^{BD} via the third term in Eq. 3.21,

$$\delta f^{BD} = \frac{2\delta\Delta^{DR}\delta\Delta^{BR} + 2\delta\Delta^R\delta\Delta^{BDR}}{|f_0^0|^2}. \quad (3.73)$$

An order of magnitude estimate for these \tilde{R} -odd mixing elements can be obtained from taking the measured values of Δ and Δ^D and multiplying them by the ratio of the magnetic-field matrix element to the rotational matrix element, $g_F\mu_B\mathcal{B}_Z/\hbar\omega_{\text{rot}} \sim 0.01 \text{ G}^{-1}$. Since Δ and Δ^D are $\sim 1 \text{ Hz}$, we can expect effects of 10 mHz G^{-1} . We took data with large axial magnetic fields and saw shifts in f^{DR} of $170(30) \mu\text{Hz G}^{-1}$ for $f^0 = 151 \text{ Hz}$, confirming the order of magnitude of this approximation. Applying these approximations to Eq. 3.73, we can then expect a systematic shift in f^{BD} of $\sim 2 \mu\text{Hz G}^{-2} \times \mathcal{B}_Z^B \mathcal{B}_Z$. Looking at $f^{DR} - \frac{\delta g_F}{g_F} f^R$ and $f^{BDR} - \frac{\delta g_F}{g_F} f^{BR}$ for our dataset puts limits on the fields of $\mathcal{B}_Z < 0.5 \text{ G}$ and $\mathcal{B}_Z^B < 0.4 \text{ G}$. We note that the same magnetometers used to shim out \mathcal{B}_Z are able to shim \mathcal{B}_X and \mathcal{B}_Y to $\sim 10 \text{ mG}$ but since we have not confirmed this for \mathcal{B}_Z , we use this more conservative limit. Finally we constrain any systematic shift in f^{BD} to $< 400 \text{ nHz}$.

3.7.4 Oddities & miscellany

3.7.4.1 Ion-ion interactions

The dominant ion-ion interaction is from the monopole charge on each ion. If we qualitatively model the cloud as 20,000 singly charged ions uniformly distributed across a sphere 1 cm in radius, we find that the resulting mean-field electric field is much smaller than either $\vec{\mathcal{E}}_{\text{rot}}$ or the peak $\vec{\mathcal{E}}_{\text{RF}}$, but comparable to the time-averaged value of $\vec{\mathcal{E}}_{\text{RF}}$. That is to say, the mean-field self-repulsion of the cloud is only modestly smaller than the effective electric fields providing secular confinement. We see evidence for this for instance in the frequencies of the breathing modes of the cloud. The ions therefore experience a net confinement that is somewhat anharmonic. None of the arguments for the limits of the size of systematic errors on d_e hinge on the confinement being particularly harmonic. Each trapped ion must perforce experience a time-averaged axial electric total field (whether external or from other ions) extremely close to zero, and thus to a high precision the rotation of $\vec{\mathcal{E}}_{\text{rot}}$ causes no Berry's phase frequency shift. It is the case, however, that the mean-field repulsion causes the ion cloud to be larger than it otherwise would be at a given temperature. In the presence of various field inhomogeneities, changes in the radius of the cloud can change both the decoherence rate and the average fringe frequency f^0 . We see shifts in the mean Ramsey frequency f^0 that correlate with the number of ions in the trap. These do not appear in the D-odd frequency channels. The mean-field coulombic repulsion does not break the various symmetries that keep frequency shifts out of the science signal.

On a microscopic level, the coulomb potential between two atoms is roughly $10^{-3}kT$, and thus ions are far from the ion-crystal regime. When small-impact-parameter ion-ion collisions occur, the ion-ion electric field can briefly spike to a magnitude which is not infinitesimal compared to \mathcal{E}_{rot} . This can cause the tip of the $\vec{\mathcal{E}}_{\text{rot}}$ to briefly wobble in a way that encloses solid angle, and there can be a resulting random Berry's phase shift that degrades the coherence between the $m_F = 3/2$ and $m_F = -3/2$ states, but does not bias the central value of the Ramsey

fringe frequency. Adiabatic relaxation of the confining potential of the trap dramatically reduces this source of decoherence, at the cost of increasing the decoherence from spatial inhomogeneity mentioned in this section. We empirically reoptimize the compromise value of the ramped-down confinement several times during the course of a long data collection run.

Fields from the molecule-frame electric dipole moment are down from the monopole-generated fields by a factor or more.

As for magnetic interactions, each ion carries a magnetic dipole moment $(3/2)g_F\mu_B$. Approximating the distribution of 20,000 ions just as described above, the field arising from a uniform magnetization within the ball of ions could cause frequency shifts of order pHz, and is thus neglected.

3.7.4.2 Effect of ion-cloud spatial distributions in which the Stark doublets are not perfectly overlapping

Much of the ultimate accuracy and precision our measurement of $f^{BD} \sim d_e \mathcal{E}_{\text{eff}}$ comes from the fact that we are measuring a resonance in two samples of ions which overlap perfectly in space and in time, but for which \mathcal{E}_{eff} points in opposite directions. If for some reason the two samples do not perfectly overlap, then spatial variation in \tilde{B} -odd frequencies can cause systematic shifts in our science signal. Generically, time-averaged \tilde{D} -odd displacements $\langle r_i \rangle^D$ or sizes of the ion cloud $\langle r_i^2 \rangle^D$ can couple to first- and second-order gradients in f^B to give systematic shifts,

$$\delta f^{BD} = \sum_{i=x,y,z} \langle r_i \rangle^D \frac{\partial f^B}{\partial r_i} + \langle r_i^2 \rangle^D \frac{\partial^2 f^B}{\partial r_i^2}. \quad (3.74)$$

By moving the ions around in the trap, we measure $\frac{\partial f^B}{\partial r_i} \sim 40 \text{ mHz cm}^{-1}$ and $\frac{\partial^2 f^B}{\partial r_i^2} \sim 10 \text{ mHz cm}^{-2}$ and so we are potentially interested in $\langle r_i \rangle^D \sim 1 \times 10^{-4} \text{ cm}$ and $\langle r_i^2 \rangle^D \sim 5 \times 10^{-4} \text{ cm}^2$.

We have good a priori reason to believe that the spatial distribution of ions in the two doublets are identical to very high precision. The two Stark doublets are initially populated by

optical pumping from the $^1\Sigma$ ground state via a $^3\Pi_0$ state, both with $\Omega = 0$. To an excellent approximation, both these states have well defined parity in the modest electric fields used in our experiment and so the pumping process is completely independent of doublet. Any \tilde{D} -odd spatial distributions must be imprinted subsequently, either by the lasers used to prepare or readout the states of the molecules, or by \tilde{D} -odd forces on the ions.

We first discuss possible effects of the lasers used for state-preparation. The next laser to interact with the ions is the 1082 nm which polarizes the ions by optically pumping them into stretched m_F states. Although this process again proceeds via an excited state with $\Omega = 0$, and so each photon scatter is equally likely to populate either doublet, in this case the laser excites out of the $^3\Delta_1$ state, and so can potentially cause a \tilde{D} -odd population difference. However, we operate in the strongly saturated regime where all ions have sufficient time to interact with the laser and so any effects of spatial intensity variation are strongly suppressed. Due to this suppression, we expect any effects to be smaller than those discussed in the next paragraph.

The other CW laser addressing molecules in the science state is the 814 nm laser used to remove population in unwanted m_F states. This laser is on for 7 ms before the Ramsey sequence, to clean up any population not successfully optically pumped, and for 25 ms after the Ramsey sequence to remove population in one stretched state. Using measurements of ion number as a function of laser interaction time, we estimate the time taken to remove all molecules in unwanted m_F levels is about 1 ms and so both interaction periods are strongly saturated, again greatly suppressing any effects of spatial variation. However for this laser, there is a second relevant time scale; because the electric field is not perfectly aligned with the k -vector of the light, molecules in the desired stretched m_F state can also occasionally scatter photons, removing a fraction β_{dep} from the state we detect. This effect is largest for the time after the Ramsey sequence, where interaction time is longest and we conservatively estimate that $\beta_{\text{dep}} < 0.1$. Although this step takes place after the Ramsey sequence, if the ions removed have some \tilde{D} -odd spatial dependence, that will result in similar dependence being imprinted onto the remaining ions which we detect.

The rotating $\vec{\mathcal{E}}_{\text{rot}}$ causes the ions to move in small circles at speeds of $\sim 1000 \text{ m s}^{-1}$. This motion in the X-Y plane causes a sinusoidally varying Doppler shift oscillating at ω_{rot} with amplitude $\Delta_{\text{Dopp}} \sim 1000 \text{ MHz}$. This Doppler shift, combined with the Stark splitting of $\Delta_{\text{St}} \sim (d_{\text{mf}}\mathcal{E}_{\text{rot}} - A_{\parallel}/2)/h \sim 140 \text{ MHz}$ between the two doublets means that the laser comes into resonance with each at a slightly different time on each rotation. We have identified two possible mechanisms related to this effect which can produce \tilde{D} -odd spatial structure in the cloud.

The first is that the interaction with each doublet takes place at slightly different spatial locations, separated by $\delta x_{\text{dep}} \sim 60 \mu\text{m}$. This can cause an initial offset between the two clouds, $r_{i0}^D \lesssim \beta_{\text{dep}}\delta x_{\text{dep}} \sim 6 \mu\text{m}$. Initial offsets in the center position of the two doublets can cause a non-zero $\langle r_i \rangle^D$ but, given that both clouds oscillate about the trap center with trap frequencies $\omega_{\text{sec}} \sim 2\pi \times 1 \text{ kHz}$, the time-averaged effect is heavily suppressed by the long Ramsey times, $t_R \sim 2 \text{ s}$ used in our experiment, $\langle r_i \rangle^D \sim \langle r_{i0} \rangle \omega_{\text{sec}} / 2\pi t_R \sim 30 \text{ nm}$. This effect is $\tilde{D}\tilde{R}$ -odd and so produces a systematic when coupled to gradients in f^{BR} . The f^{BR} gradients we measure in the radial directions are $< 20 \text{ MHz cm}^{-1}$, giving a systematic shift of $< 100 \text{ nHz}$. In reality we expect any effect to be significantly smaller still due to the depletion laser being on for a time which is long compared to one trap period so that any r_{i0}^D is greatly reduced.

If the laser beam has non-zero displacement from the center of the micromotion, x_{dep} , the interaction can selectively remove more of the hotter ions from one doublet than from the other, producing an effective $\langle r_i^2 \rangle^D$ in those ions remaining in ${}^3\Delta_1$. In the worst-case limit where the depletion laser is much smaller than the ion cloud, the mean square position of the removed ions is

$$(x_{\text{dep}} + \frac{\delta x_{\text{dep}}}{2})^2 - (x_{\text{dep}} - \frac{\delta x_{\text{dep}}}{2})^2 = 2x_{\text{dep}}\delta x_{\text{dep}}. \quad (3.75)$$

The ions oscillate backwards and forward in the trap and so, in the limit of no ion-ion collisions, their time-averaged mean square position during the Ramsey time is reduced by a factor of 2. Simulations of more realistic ratios of laser to cloud size reduce this by a further factor of ~ 3 . We

estimate $x_{\text{dep}} < 2 \text{ mm}$ and so estimate the difference in mean square position of the remaining ions as $\langle r_i^2 \rangle^D \sim \beta_{\text{dep}} x_{\text{dep}} \delta x_{\text{dep}} / 3 \sim 4 \times 10^{-5} \text{ cm}^2$. The effect is again $\tilde{D}\tilde{R}$ -odd and so can produce a systematic when accompanied by curvature in f^{BR} . The largest such curvature we measure is 70 MHz cm^{-2} , producing a systematic of $\sim 3 \mu\text{Hz}$. In reality, ion-ion collisions redistribute the velocities and positions of ions in the trap so that any initial difference in the size of the clouds is rapidly scrambled. To set an upper limit on the timescale of this thermalisation, we note that we measure a curvature in f of 1200 MHz cm^{-2} , caused by the spatial variation of the magnetic field from the \vec{B}^0 coils. For a 1 cm gaussian cloud with no ion-ion collisions, this would cause decoherence of $\sim 20\%$ of the contrast in $\sim 50 \text{ ms}$. This is roughly equal to our total loss of contrast over a 3 s Ramsey time and so $\sim 50 \text{ ms}$ sets a rough upper limit for the timescale of thermalisation. For the shortest Ramsey times used in the dataset of $\sim 1.5 \text{ s}$, the possible systematic shift is reduced by a factor of $50 \text{ ms} / 1.5 \text{ s} \sim 30$. As such we estimate the size of any residual systematic to be $< 100 \text{ nHz}$ and so do not include its contribution in our systematic error budget.

The second mechanism is that the electric field vector defining the quantisation axis is, on average, pointing in slightly different directions when the laser interacts with each doublet. These angles are given by $\theta_u = \theta_0/2 + \delta\theta$, $\theta_l = -\theta_0/2 + \delta\theta$ where $\theta_0 \sim \Delta_{\text{St}}/\Delta_{\text{Dopp}} \sim 0.14$ is the difference in the angle of \mathcal{E}_{rot} when it addresses each of the two doublets and $\delta\theta \sim \delta_{\text{depl}}/\Delta_{\text{Dopp}}$ is half the difference in the magnitude of those angles caused by imperfect laser detuning $\delta_{\text{depl}} \lesssim 20 \text{ MHz}$. In the unsaturated limit, the scattering rate with which each doublet interacts with the laser is then different by a factor $\cos\theta_u - \cos\theta_l \sim -2\Delta_{\text{St}}\delta_{\text{depl}}/\Delta_{\text{Dopp}}^2$. Again, the laser interacts with a fraction ~ 0.1 of the molecules and so we estimate the possible difference in mean square position as $\langle r_i^2 \rangle^D \sim \langle r_i^2 \rangle 0.1(2\Delta_{\text{St}}\delta_{\text{depl}}/\Delta_{\text{Dopp}}^2) \sim 6 \times 10^{-4} \text{ cm}^2$. This effect is \tilde{R} -even and so can produce a systematic when accompanied by curvature in f^B . The largest such curvature we measure is 10 MHz cm^{-2} , producing a systematic of $< 6 \mu\text{Hz}$. The same thermalization arguments discussed in the previous paragraph reduce this to below 200 nHz and so we include

no contribution in our systematic error budget.

Finally, we consider the first dissociation laser. The photodissociation beam is a doubled, pulsed dye laser with a 10 ns pulse width. The central frequency of the laser pulse is stabilized to a wave meter, with the set point chosen to be a compromise between the peaks of the two transitions for dissociating the upper and lower doublets. Over the course of the entire data run we find that on average we have see \tilde{D} -odd term of 2% in the number of Hafnium ions we create. A residual differences in efficiency, in combination with intensity-driven saturation and spatial structure on the beam can generate a D -odd contribution to the mean-square spatial extent of the molecules that participate in the Ramsey fringes. This in turn can combine with curvature in f^B , which can be as large as 10 MHz cm^{-2} , to give a systematic error on f^{BD} . As with possible \tilde{D} -odd effects arising from the optical pumping beams, the size of the resulting systematic error is strongly suppressed by ion-ion thermalization that occurs over the course of the t_R and washes out size differences. The pulsed lasers width in frequency space is large compared to the transform-limited value for a 10 ns pulse width, and frequency structure within the pulse is not well characterized, which makes modeling the effect of the laser beam a little uncertain. Just before entering the chamber, however, the dissociation beam passes through an iris with 1 cm diameter. We can make the simplifying and very pessimistic assumption that the extra 2% \tilde{D} -odd change in Hafnium ions is due entirely to molecules dissociated just at the outer edge of the laser beam. In this way we can set a conservative limit that the size of this effect on f^{BD} must be less than $3.5 \mu\text{Hz}$.

In addition to these effects from the lasers, we have also considered possible effects of \tilde{D} -odd forces on the ions due to electric field gradients, and \tilde{D} -odd heating due to photon scattering. We find each to be significantly smaller than our limits on those from the lasers.

3.8 Phase Shifts

In this section we explore possible systematic effects caused by shifts in the phase of the ions, or our measurement of that phase. Using trapped ions for our measurement makes us relatively insensitive to these effects when compared to other similar experiments using beams of atoms or molecules. The first reason is that we can vary the free evolution time of our measurement; each block of our dataset consists of data taken at short and long Ramsey times. Most phase shifts caused by state-preparation or measurement are common mode and so cancel out in our science signal which is only sensitive to the differential phase evolution between the early and late time data. The second reason is that any residual shifts—from phase effects which are different for the early and late time data—get divided by the difference in free evolution time between the two measurements. The coherence time of the ions in our experiment is roughly three orders of magnitude longer than comparable experiments using beams and so any effects are reduced by a similar factor. In Sec. 3.8.2, we consider errors caused by imperfections in our measurement of those accumulated phases, and in Sec. 3.8.1, we consider phase errors caused by imperfections in state-preparation.

3.8.1 State preparation

3.8.1.1 phases due to $\pi/2$ pulses

During preliminary data runs, we measured a small, and unexpected, contribution to the initial phase (roughly 10 mrad) which is odd under \tilde{D} , \tilde{B} and \tilde{I} . We studied the dependence of this effect on many experimental parameters. In particular we found: (a) the magnitude of the effect depended strongly on the parameters used for the $\pi/2$ pulses—we could greatly suppress the effect by implementing $\pi/2$ pulses with higher values of \mathcal{E}_{rot} during them and compensating by increasing their duration; (b) the sign of the effect changed when we projected into the opposite doublet; and (c) contrary to what we had anticipated for an effect which is odd under \tilde{I} , there is no dependence on the side of the phosphor screen that each doublet is imaged onto.




Figure 3.11: Toy model to explain ϕ^{DBI} . (i) uncoupled basis, (ii) coupled basis.

We have settled on a probable physical mechanism for the effect which qualitatively matches our observations by including the effects of other states ${}^3\Delta_1(v = 0, J = 1)$ outside the two doublets used for spectroscopy. The important features of this mechanism can be illustrated with the toy model system shown in Fig. 3.11(i). The model includes just 4 of the 12 levels in ${}^3\Delta_1(J = 1)$: the two $F = 3/2, |m_F| = 3/2$ states of a given doublet, labelled $|a\rangle$ and $|b\rangle$, and the two $F = 3/2, |m_F| = 1/2$ levels which lie closest to them in energy, labelled $|c\rangle$ and $|d\rangle$. In an electric field, and in the absence of any magnetic field or rotation, the two pairs of states with $|m_F| = 3/2$ and $|m_F| = 1/2$ are each degenerate. The Stark splitting between the two pairs is κ . When rotation is included, there is a direct coupling mixing $|a\rangle$ and $|c\rangle$, and $|b\rangle$ and $|d\rangle$. In the full 12-level model there is also a second-order coupling between $|c\rangle$ and $|d\rangle$ —involving a single matrix element associated with each of rotation and Ω -doubling—and a fourth-order coupling between $|a\rangle$ and $|b\rangle$ —involving three matrix elements associated with rotation and a matrix element associated with Ω -doubling. Although these couplings are much smaller than the direct coupling, they are important because they act in degenerate subspaces so we include their effects in our toy model through couplings Δ and Δ_2 . We will not include the effect of a magnetic field because we are interested in what happens during a $\pi/2$ -pulse where the other effects are all

much larger. The Hamiltonian for our toy system in this basis is then

$$H = \begin{pmatrix} \kappa & \delta & 0 & \Delta \\ \delta & 0 & \Delta_2 & 0 \\ 0 & \Delta_2 & 0 & \delta \\ \Delta & 0 & \delta & \kappa \end{pmatrix}, \quad (3.76)$$

where the states are in order of m_F : $|a\rangle, |c\rangle, |d\rangle, |b\rangle$.

We can simplify our thinking by changing to a different basis. Consider the case where we somehow independently tune δ to zero; we now have two uncoupled systems, $|a\rangle$ and $|b\rangle$, and $|c\rangle$ and $|d\rangle$. We can diagonalize each system to give eigenstates which are fully mixed as shown in Fig. 3.11(ii): $|3/2, +\rangle = \frac{1}{\sqrt{2}}(|a\rangle + |b\rangle)$ and $|3/2, -\rangle = \frac{1}{\sqrt{2}}(|a\rangle - |b\rangle)$, and $|1/2, +\rangle = \frac{1}{\sqrt{2}}(|c\rangle + |d\rangle)$ and $|1/2, -\rangle = \frac{1}{\sqrt{2}}(|c\rangle - |d\rangle)$. Rewriting the full toy-model Hamiltonian in this basis we have

$$H = \begin{pmatrix} \Delta_2 & \delta & 0 & 0 \\ \delta & \kappa + \Delta & 0 & 0 \\ 0 & 0 & -\Delta_2 & \delta \\ 0 & 0 & \delta & \kappa - \Delta \end{pmatrix}, \quad (3.77)$$

where the states are in the order $|1/2, +\rangle, |3/2, +\rangle, |1/2, -\rangle, |3/2, -\rangle$. The Hamiltonian is now block diagonal. During the second $\pi/2$ pulse, the mixing between $|a\rangle$ and $|b\rangle$ causes population to oscillate between them (as intended), but it also causes a weak oscillation between the $|m_F| = 3/2$ states and the $|m_F| = 1/2$ states. This oscillation depends on the initial phase difference between the $|a\rangle$ and $|b\rangle$ states. The two most important phases to consider are the sides of the fringe; where we are most sensitive to phase shifts, and thus where we take most of our data. On the side of the fringe, the state immediately before the second $\pi/2$ pulse is $|3/2, +\rangle$ or $|3/2, -\rangle$. It can be seen from Eq. 3.77 that when $\delta = 0$, these states are eigenstates of the Hamiltonian and the $\pi/2$ -pulse does nothing. When δ is non-zero, the coupling causes population to be transferred from $|3/2, +\rangle$ to $|1/2, +\rangle$ and from $|3/2, -\rangle$ to $|1/2, -\rangle$ with Rabi frequencies Ω_+

and Ω_- respectively,

$$\Omega_+ = \sqrt{4\delta^2 + (\kappa + \Delta - \Delta_2)^2} \sim \kappa + \Delta - \Delta_2, \quad (3.78)$$

$$\Omega_- = \sqrt{4\delta^2 + (\kappa - \Delta + \Delta_2)^2} \sim \kappa - \Delta + \Delta_2. \quad (3.79)$$

The two Rabi frequencies are different because the energy gap is different in each case. The amplitudes of the oscillations in each case are then

$$A_+ = \frac{\delta^2}{4\delta^2 + (\kappa + \Delta - \Delta_2)^2} \sim \frac{\delta^2}{\kappa^2}, \quad (3.80)$$

$$A_- = \frac{\delta^2}{4\delta^2 + (\kappa - \Delta + \Delta_2)^2} \sim \frac{\delta^2}{\kappa^2}. \quad (3.81)$$

The approximate expressions in both cases assume $\kappa \gg \delta, \Delta, \Delta_2$. This difference in Rabi frequencies results in a different amount of the population being transferred to $|c\rangle$ and $|d\rangle$, dependent on the initial phase. We measure only the population in either $|a\rangle$ or $|b\rangle$ and never the population in $|c\rangle$ or $|d\rangle$, and so this population transfer appears as an apparent phase shift. The population oscillations are happening very fast—the frequencies are of order the energy gap between the $3/2$ states and the $1/2$ states, $\kappa \sim \text{MHz}$ —and so the size of the apparent phase shift can depend very sensitively on the exact parameters of the $\pi/2$ -pulse. The maximum size of the effect goes as $1/\kappa^2$ and so depends strongly on the size of \mathcal{E}_{rot} during the $\pi/2$ pulse.

We now examine the dependence of this effect on the experimental switch state:

- (i) \tilde{I} : in our implementation of the \tilde{I} switch, the m_F state we project into, and read out of, is \tilde{I} odd. This corresponds to measuring, for example, $|a\rangle$ instead of $|b\rangle$. On the side of the fringe where the population in $|a\rangle$ is decreasing with Ramsey time, the population in $|b\rangle$ is increasing with Ramsey time and as a result the apparent phase shift—a relative change in population which has same sign for both states—is \tilde{I} odd.
- (ii) \tilde{B} : the phase shift is completely independent \vec{B}^0 and so is \tilde{B} -odd.
- (iii) \tilde{D} : in the opposing doublet, the sign of κ is flipped; the $|m_F| = 1/2$ states are above the

$|m_F| = 3/2$ states in energy. This means that the relative rate of transfer between the two sides of the fringe are flipped and the effect is \tilde{D} odd.

- (iv) \tilde{R} : δ , Δ and Δ_2 all involve odd numbers of matrix elements associated with rotation and so flip sign with rotation direction. This flips the sign of the effect so that the relative rates of population transfer out of the two states are reversed. However, the flipping sign of Δ means that the initial superposition produced by the first $\pi/2$ pulse is π out of phase compared with the opposite rotation direction so the phase shift at a given Ramsey time ends up being \tilde{R} even.

In the experiment it is very difficult to control the parameters of the $\pi/2$ pulses at the required level to quantitatively verify this mechanism is at play. In addition small differences in the electric field experienced by different ions, and coupling to other states in the full Hamiltonian, cause the population oscillations to dephase rapidly, and so we were never able to observe them. However, we believe the qualitative agreement with what we observe to be convincing.

We took three steps to mitigate the effects of ϕ^{DBI} . We reduced the size of the effect dramatically by increasing the value of \mathcal{E}_{rot} that we ramp down to during the $\pi/2$ -pulse. We also added the \tilde{P} chop to our experimental sequence for the eEDM dataset; in every other block, we read out of the opposite state, so that any remaining effect would change sign. Finally, midway through the dataset, we reversed the polarization of all the optical pumping, depletion and dissociation lasers, again flipping the sign of any effect. We can still check for any residual effect by looking at the difference in ϕ^{DBIP} before and after the waveplate change; in the eEDM dataset its average value was 1.1(1) mrad. Such phase shifts could potentially leak into frequency channels if, for example, the second $\pi/2$ pulse depended on the Ramsey time, perhaps because of heating in the amplifiers or similar. We note that during our exploration of the effect we were able to increase its size to 170 mrad, and saw no shift in f^{DBI} at the 1σ confidence level. Using this data we can set a limit on any shift caused by the phase at $-3(4) \mu\text{Hz mrad}^{-1}$, increasing

confidence in our ability to reject systematic effects associated with phases caused by the $\pi/2$ pulses. In the eEDM dataset, none of the other \tilde{D} -odd phases (except ϕ^D) exceed $750 \mu\text{rad}$.

3.8.2 Internal state measurement

3.8.2.1 Improperly characterized imaging contrast

Our imaging process allows us to differentiate between Hf ions originating from the upper and lower doublets by projecting them onto different sides of the imaging MCP. However, due to the initial size of the cloud and its finite temperature the two clouds are not perfectly separated and they overlap slightly at the center of the screen. This overlap can cause our measurement of the difference frequency to be pulled and, if improperly characterised, can potentially cause a systematic error. To see how this works, consider how the number of ions detected on each side of the MCP, and in each phase, depends on the phase evolution of the upper and lower doublets Φ_u and Φ_l ,

$$\begin{aligned}
 N_u^{\text{anti}} &= \frac{N}{2}(1 - \sin(\Phi_u))(1 - \epsilon) + \frac{N}{2}(1 - \sin(\Phi_l))(\epsilon) \\
 N_u^{\text{in}} &= \frac{N}{2}(1 + \sin(\Phi_u))(1 - \epsilon) + \frac{N}{2}(1 + \sin(\Phi_l))(\epsilon) \\
 N_l^{\text{anti}} &= \frac{N}{2}(1 - \sin(\Phi_l))(1 - \epsilon) + \frac{N}{2}(1 - \sin(\Phi_u))(\epsilon) \\
 N_l^{\text{in}} &= \frac{N}{2}(1 + \sin(\Phi_l))(1 - \epsilon) + \frac{N}{2}(1 + \sin(\Phi_u))(\epsilon),
 \end{aligned} \tag{3.82}$$

where N is the mean number of ions measured on each side of the MCP, ϵ is a parameter which characterises the amount of leakage of each cloud onto the other side of the MCP screen, and we have assumed perfect contrast. We can now form asymmetries as we do in our analysis,

$$\begin{aligned}
 A_u &= \frac{N_u^{\text{in}} - N_u^{\text{anti}}}{N_u^{\text{in}} + N_u^{\text{anti}}} \\
 &= (1 - \epsilon) \sin \Phi_u + \epsilon \sin \Phi_l, \\
 A_l &= \frac{N_l^{\text{in}} - N_l^{\text{anti}}}{N_l^{\text{in}} + N_l^{\text{anti}}} \\
 &= (1 - \epsilon) \sin \Phi_u + \epsilon \sin \Phi_l.
 \end{aligned} \tag{3.83}$$

We can simplify our thinking by assuming that we are taking data close to the side of the fringe such that $\Phi_u \simeq 2n\pi + \delta_u$ and $\Phi_l \simeq 2m\pi + \delta_l$ for n and m integer. In this limit we have

$$\begin{aligned} A_u &\simeq \delta_u - \epsilon(\delta_u - \delta_l) \\ A_l &\simeq \delta_l + \epsilon(\delta_u - \delta_l). \end{aligned} \quad (3.84)$$

The apparent phase of each of the doublets are pulled towards each other by an amount that depends on the leakage ϵ and their difference in true phase $\delta_u - \delta_l$. We account for this in our analysis by including a parameter C_I , the imaging contrast. By comparing Eq. 3.83 with Eq. 3.11 we can identify $\epsilon = (1 - C_I)/2$. Now suppose we mischaracterize the imaging contrast C_I , assigning it an incorrect value C'_I . The resultant fitted values for the phases of the upper and lower doublet δ'_u and δ'_l will then satisfy the following expressions

$$\begin{aligned} \delta'_u - \frac{1 - C'_I}{2}(\delta'_u - \delta'_l) &= \delta_u - \frac{1 - C_I}{2}(\delta_u - \delta_l), \\ \delta'_l + \frac{1 - C'_I}{2}(\delta'_u - \delta'_l) &= \delta_l + \frac{1 - C_I}{2}(\delta_u - \delta_l). \end{aligned} \quad (3.85)$$

We find that the fitted difference phase $\Delta' = \delta'_u - \delta'_l$ is different from the true difference phase $\Delta = \delta_u - \delta_l$ by an amount

$$\Delta' - \Delta = \frac{(C_I - C'_I)\Delta}{C'_I}. \quad (3.86)$$

Both the early and late time data in our Ramsey are fringes are potentially susceptible to this phase pulling effect and so it is important that we properly characterize and understand C_I . To do so, we took fringes where the data was collected deliberately offset from the early- or late-time zero crossings of the difference fringe. Taking the late time data removed from the beat by a number of periods n causes the fitted f^D to be pulled by an amount

$$\delta f^D = \frac{\Delta' - \Delta}{2\pi t_R} = \frac{n}{t_R} \frac{(C_I - C'_I)}{C'_I} \frac{f^D}{f}. \quad (3.87)$$

Moving the early time data has the same size effect but with the opposite sign. For a given C'_I and swatch size, we can fit the extracted f^D vs the number of periods we are offset from the zero crossing of the difference fringe. Figure 3.12 shows how this slope varies with the value of C'_I for the swatch size of 90 pixels used in our measurement, chosen to maximize $C\sqrt{N}$. The correct




Figure 3.12: Example data. Plotting the slope of f^D when the early-time data is taken at varying offsets from the zero crossing of the difference fringe. The x -intercept gives us the true value of our imaging contrast, which we then use in our fitting program.

value of C'_I is the one for which $\frac{\partial f^D}{\partial n} = 0$, here $C'_I = 0.90(1)$, corresponding to $\epsilon = 0.050(5)$, which we use for analysis of our whole dataset. Whilst we can measure C_I very precisely, its value can wander slightly over time. We have repeated this measurement many times over the course a year and find that the early and late time C_I are consistent with one another, and between measurements, to ± 0.05 . We conservatively estimate the largest possible deviation in C_I , averaged over the dataset to be $\delta C_I \sim 0.05$.

For this effect to leak into f^{BD} and cause a systematic shift in our science signal requires either the amount we miss the beat by to \tilde{B} -odd, or the imaging contrast itself to be \tilde{B} -odd. We consider the \tilde{B} -odd phase first. Because the mean value of ϕ^D is negative, the early time beat happens before zero, and so we are forced to take data with nonzero Δ . In this case a non-zero value of ϕ^{BD} will mean that the early time data is taken at a different Δ depending on the \tilde{B} switch. The value of ϕ^{BD} over the dataset is $-30(140) \mu\text{rad}$. Combining this with the uncertainty in C_I above leads to a systematic uncertainty $\frac{\phi^{DB} \delta C_I}{2\pi t_R C_I} \sim 0.7 \mu\text{Hz}$.

A systematic shift can also potentially be caused by ϕ^D combined with a \tilde{B} -odd imaging contrast C_I^B . In order for this shift to reach the $5 \mu\text{Hz}$ level, would require $C_I^B \sim 0.005$, or equivalently $\epsilon^B \sim 2.5 \times 10^{-3}$, or about 5% of the total leakage. We know of no mechanism which

can cause such an effect and so include no contribution in our systematic uncertainty budget.

3.8.2.2 Swatch position

As described in Section 4.2, the position of the center line and corresponding swatch for each switch state is determined by an algorithm based on images of the cloud taken without any $\pi/2$ -pulses. In this section we explore the consequences of a possible systematic error in the center position dividing line between the two doublets.

We reanalysed the dataset, with the center line in each image displaced by different number of pixels left and right. The largest observed shifts in the frequency and phase channels are shown in Fig. 3.15. We note that the signal to noise in these shifts is much better than indicated by the error bars shown because we are reanalysing the same data. We can explain the observed shifts using ideas from the previous section. Let's use Eq. 3.85 as before but this time with the phase pulling on each of the doublets separated,

$$\delta'_u - \delta_u = \frac{(\delta_u - \delta_l)(\epsilon - \epsilon')}{2\epsilon' - 1}, \quad (3.88)$$

$$\delta'_l - \delta_l = -\frac{(\delta_u - \delta_l)(\epsilon - \epsilon')}{2\epsilon' - 1}. \quad (3.89)$$

We can combine both of these equations as

$$\delta\phi = \tilde{D} \frac{\Delta(\epsilon - \epsilon')}{2\epsilon' - 1} \quad (3.90)$$

where $\delta\phi$ is the phase pulling of a single fringe and $\Delta = \delta_u - \delta_l$ is the amount we are missing the beat by. We can express the effect of moving the swatch center around by modifying ϵ , the amount each doublet leaks into the other

$$\epsilon = \epsilon_0 + \kappa \tilde{D} \tilde{I} Y_c - \nu (\tilde{D} \tilde{I} Y_c)^2. \quad (3.91)$$

This equation expresses the fact that as we move the swatch over to one side, there is more leakage into one of the doublets and less into the other. Which way around they are depends on the I switch. The first order effect, $\propto \kappa$, is linear but for larger swatch displacements, the




Figure 3.13: Varying the swatch width while correcting for the doublet contrast. We see no meaningful shifts in our science channel and a small residual shift in f^D . In the lower panel we have averaged over the 3 distinct values of f^D that we operated at.

second order effect, $\propto \nu$, starts to become important too. Substituting this back into Eq. 3.90 and assuming for simplicity that we have done a good job of picking the doublet contrast in the first place so that $\epsilon' = \epsilon_0$, we get

$$\delta\phi = \frac{\Delta}{2\epsilon' - 1} (\tilde{I}\kappa Y_c - \tilde{D}\nu Y_c^2). \quad (3.92)$$

The last factor we need to include to explain the data is to make the approximation that the main place we are missing the beat is at early time (due to the beat being before $T = 0$), so that $\Delta \sim \phi^D$. Here we have included the largest (by far) contribution to the D -odd phase only. Finally we have

$$\delta\phi = \frac{1}{2\epsilon' - 1} (\tilde{I}\kappa\phi^D Y_c - \tilde{D}\nu\phi^D Y_c^2). \quad (3.93)$$

So the principle effects we expect are a linear shift in ϕ^I and a smaller quadratic effects in ϕ^D . Because (almost) all of the pulling is happening at early time, each of these could be expected to have an echo in the corresponding f channel, with opposite sign and smaller by a factor of $2\pi \times T \sim 15$ s rad. Each of these can be seen in Fig. 3.15 from which we infer $\kappa = -1.89(6) \times 10^{-3}$ per pixel and $\nu = -7.0(2) \times 10^{-6}$ per pixel squared. The shifts seen in ϕ^{DI} and f^{DI} (about a factor of 5 smaller than those in ϕ^I and f^I) could be caused by a \tilde{I} -odd systematic error in the swatch position Y^I . Including this in Eq. 3.93 yields an additional contribution $\frac{\phi^D}{2\epsilon' - 1} (\kappa Y^I - \tilde{D}\nu Y^I{}^2 - 2\tilde{D}\tilde{I}\nu Y_c Y^I)$. Comparing the term linear in Y_c to the observed gradient in ϕ^{DI} gives $Y^I = 25(1)$ pixels. Other possible systematic shifts in swatch position can be constrained by the data in the same way to less than 5 pixels.

The largest gradient in a frequency channel seen in this analysis is $\partial f^I / \partial Y_c \sim 2$ μ Hz per pixel. It is important to realize that the swatch displacement that is carried out here is just one of many possible ways of moving the swatch. Because we moved the swatch the same direction in each switch state, Y_c shows up in Eq. 3.91 with $\tilde{D}\tilde{I}$. There are 7 other possible ways to move the swatch, always \tilde{D} -odd but with all other possible switch-state dependence. We can infer the effects of these possible systematic effects from those shown here. We are interested in any possible shifts in ϕ^{DB} and f^{DB} which can only show up quadratically in any systematic error




Figure 3.14: The effect of moving the center of the swatch on our science channel.

in the swatch position. To see an effect requires either missing the beat by a \tilde{B} -odd amount, or a combination of two systematic shifts in swatch position, one \tilde{B} -odd and one \tilde{B} -even. As discussed in the previous section, the former can happen due to a combination of our non-zero ϕ^D and a non-zero ϕ^{DB} . In our dataset, $\phi^{DB} = -30(140) \mu\text{rad}$ which, combined with Y^I gives a systematic $|\frac{\phi^{DB} \nu Y^{I^2}}{2\pi T(1-ee')}| < 70 \text{ nHz}$. The largest contribution to the latter is from Y^I combined with Y^{BI} , constrained by looking at the gradient of ϕ^{DBI} in our analysis to be 1(2) pixels. This gives a systematic of $|\frac{2\phi^D \nu Y^I Y^{BI}}{2\pi T(1-ee')}| < 1.2 \mu\text{Hz}$. We include the quadrature sum of these two in our systematic uncertainty budget.

3.8.2.3 Counting saturation

Our imaging system has a limited capacity, and if that were somehow magnetic field and doublet odd it could be a concern. It's hard to conceive of a method for this to be both magnetic field and doublet odd, however, because even if one side of the screen saturated faster than the other, that asymmetry is heavily suppressed by the fact that the \tilde{I} chop swaps which side of the screen we read each doublet out on.

Collating the imaging data over 25 blocks we can see that our imaging does indeed saturate (the image intensity does not scale linearly with the number of Hf^+ ions we count, see Figure




Figure 3.15: Moving the center of the switch has a measurable effect in several parity channels. Note all panels show the entire aggregated dataset except for the f^D panel, which shows a subset of measurements performed at a single f^0




Figure 3.16: Imaging data collated over 25 blocks. The lower doublet data is in pink while the upper doublet data is in blue.

3.16).

We can fit the data in Figure 3.16 to the function $I = \alpha \text{Hf}_{\text{counted}}^+ + \text{fi}(\text{Hf}_{\text{counted}}^+)^2$. For small $\text{Hf}_{\text{counted}}^+$, the image intensity should scale linearly with $\text{Hf}_{\text{counted}}^+$. Therefore, we can apply a saturation correction to our data: $\text{Hf}_{\text{real}}^+ = \text{Hf}_{\text{counted}}^+ (1 + \text{fi}/\text{ffHf}_{\text{counted}}^+)$. We then re-analyzed all our data with several saturation corrections applied: none/zero, two incorrect ones, and the “true” value of $\beta/\alpha = 8.52 \times 10^{-4}$. We see no observable shift in our science channel and so we do not include this effect in our systematic error budget.

3.8.2.4 Time-varying offsets

Our asymmetry fringes can be offset from zero, characterised by the parameters $O_{u/l}$ from Eq. 3.9. It is possible that these offsets might decay as a function of Ramsey time, for example if collisions slowly redistribute population between the two states in a doublet. Since we fit our fringes with a single offset for all Ramsey times, this could potentially lead to a systematic shift in the relative phases. To investigate this possibility, we refit our entire dataset with a fixed differential offset between the early and late time asymmetry data points and looked for changes in f^{BD} . We observe no shifts at the nHz level and so include no contribution in our systematic uncertainty budget.




Figure 3.17

Table 3.7: Systematic error budget

Effect	Section	Correction (μHz)	Uncertainty (μHz)
Magnetic			
Non-reversing \vec{B}^0	3.7.1.1	0.2	< 0.1
Second harmonic of \mathcal{E}_{rot} and transverse magnetic field	3.7.1.3		2
Third harmonic of \mathcal{E}_{rot} and magnetic field gradients	3.7.1.4		2
Higher harmonics of \mathcal{E}_{rot} and higher order magnetic field gradients	3.7.1.4		< 0.1
Ellipticity of \mathcal{E}_{rot} and magnetic field gradients	3.7.1.5		1
Berry's phase			
Phase modulation due to axial secular motion	3.7.2.3		3
Axial 2nd harmonic with ellipticity of \mathcal{E}_{rot}	3.7.2.4		2
Higher harmonics of \mathcal{E}_{rot}	3.7.2.5		3
Rotation induced mixing			
Leaking of f^{BR}	3.7.3.1		< 0.2
Axial magnetic fields	3.7.3.2		0.4
Other frequency			
Imperfect overlap of spatial distributions	3.7.4.2		3.5
Phase			
Improperly characterized imaging contrast	3.8.2.1		0.7
Swatch position	3.8.2.2		1.2
Total		0.2	7

3.9 Interpretation

Table 4.1 presents our error budget. We have demonstrated a significant advance in the characterization of our experiment resulting in a lower systematic uncertainty and a measurement which is statistics dominated. The use of rotating bias fields and trapped molecules is a powerful technique for suppressing systematics; the rotation means that most stationary field-induced systematics average away over a rotation cycle, and the trapped species allows us to measure both early- and late-time phase which eliminates a large class of state-preparation systematics one might otherwise have to worry about. As eEDM sensitivity is pushed into the next decade of accuracy, these advantages may prove essential.

3.10 Supplemental Information / Appendices to the paper

3.10.1 Overview of data

The data blocks can be divided into three main categories, based on the approximate value of f^0 , corresponding to roughly 77, 105, and 151 Hz. $f^D \simeq \delta g/gf^0$, thus came in respectively at -171, -230 and -327 mHz, and by choosing late-time Ramsey fringes centered at $1/2f^D$ or respectively roughly 2.9, 2.2 and 1.5 s, the late-time Ramsey fringes were taken at the re-phasing of the fringes for the upper and lower doublets. In the tables below, we present averages taken separately over the three ranges of f^0 , plus an average taken over all the blocks. In each case the quoted one- σ errors are based on combining the estimated error on all the relevant blocks, and then relaxing that value by multiplying by $\sqrt{\chi^2}$ to correct for overscatter. For reference we include the relevant value of χ^2 as calculated before the error bars were relaxed. The χ^2 associated with the relaxed error bars displayed here are all by construction equal to 1.

The tables under discussion in this section are presented in E.

3.10.2 Overall comments

The parameters which are even in all switch states (superscript o) are the noisiest, as they are unprotected from all the drifts in the experimental conditions. These drifts arise from many sources, for instance (i) changes in the temperature of the lab, or of the water cooling the power op-amps that drive the ion-trap electrodes. (ii) changes in the ambient magnetic field and its gradients, arising e.g. from a distant freight elevator, a less-distant optical table, or a still-closer weld in an only nominally nonmagnetic UHV chamber (iii) variation in the number, temperature, and density of the ion cloud (iv) drifts in the current supply that drives the \vec{B}^0 coils. Cloud size and thus ion density and trap temperature were not well-characterized parameters. Ion number in particular drifts from block to block, and due to mean-field coulomb repulsion within the cloud the ion number is coupled to cloud size, and from there to f^0 which is an average over spatial field inhomogeneities. The strength of radial and axial trap confinement were treated as parameters to be tweaked to maximize the precision per block of data. Ion production was frequently re-optimized as well, with resulting jumps in f^0 between blocks of tens of millihertz or more.

All \tilde{P} -odd values are in a category by themselves, in that they were not collected as part of a rigorously implemented intra-block chop. Instead, we reverse the direction of the relevant polarizer every other block. On some days, we took an odd number of blocks. In other instances we vetoed an entire block of data which left the next block of data unpaired with a block of opposite P switch. For these reasons, from our entire run of ~ 1300 blocks we do not try to generate ~ 650 sets of \tilde{P} -odd fit values based on pairs of matched, sequentially collected blocks. Instead, for each of the f^0 superblocs, we divide the blocks into the $\tilde{P} = 1$ and $\tilde{P} = -1$ piles, calculate weighted mean values for each fitting parameter and each of the $\{\tilde{D}, \tilde{B}, \tilde{R}, \tilde{I}\}$ parity channels, with error bars corrected for scatter. Then, element-wise across the large table of values, we take either a sum or difference between the $\tilde{P} = 1$ and $\tilde{P} = -1$ to create the overall \tilde{P} -even or \tilde{P} -odd parity channels. In this method, there is no particular meaning to χ^2 for the

\tilde{P} -odd parity channels, so no such value presented in the tables below.

After each block we attempt to servo the value of f^B back to zero for the subsequent block. Thus the mean value of f^B is very low even though f^B channel is susceptible to drifts in ambient conditions. These drifts can be tracked because we record the value of the current in the coils used to servo f^B .

Looking across the 480 parity-channel values associated with the three frequency superblocks, and the 160 fully aggregated values, we see the overwhelming majority of the values are quite close to zero. There are a few numbers which are dramatically different from zero, for good reason:

- (i) f^0 itself of course, the absolute value of the energy difference between $m_F = \pm 3/2$, which we apply a bias magnetic field to set to a value between 75 and 155 Hz, depending on the superblock.
- (ii) f^D comes in a -225 mHz, smaller than the mean value of f^0 by a factor of about $-1/450$, which is the fractional difference in g -factor between the upper and lower states, plus a few smaller correction factors.
- (iii) f^{BR} , at 212 mHz, comes from a magnetic field, rotating at 375 kHz, generated by oscillating currents that arise from charging the fins to the oscillating voltages that create the rotating electric bias field.
- (iv) f^{BDR} , at roughly 410 μ Hz, is f^{BR} echoed in corresponding \tilde{D} -odd channel by the same factor of $\delta g/g \sim -1/450$. After a correction for finite f^{BR} is applied, f^{BDR} is within 1.7σ of zero.
- (v) f^{BD} is proportional to the electron's electric dipole moment d_e and its value in principle could be quite large. If we take the ACME 90% confidence limit, $d_e < 1.1 \times 10^{-29}$ e cm e-cm at face value, we would have similar confidence that our measured value will lie between

$-120 \mu\text{Hz} < f^{BD} < 120 \mu\text{Hz}$. For all but the last few weeks of the experimental effort described here, the blinding procedure built into our analysis software made f^{BD} appear to have the arbitrary value near 16.1 mHz.

- (vi) γ^0 The mean decay rate of coherence is about 0.1 s^{-1} . Arising from ion-ion collisions and from spatial inhomogeneities, gamma is a noisy number that depends sensitively on number of trapped ions and fine details of the trap shimming.
- (vii) ϕ^0 , fairly large at tens of milliradians, is measured to be proportional to f^0 and is consistent with $\sim 150 \mu\text{s}$ (CK) offset in where we define $t_R = 0$. This is not unexpected given that the during the $\pi/2$ pulses the rabi frequency varies during the 1 ms pulse duration.
- (viii) O^0, O^D, C^D, ϕ^D . These nonzero terms all arise predominantly from the same basic cause. Due to the presence of the $F = 1/2$ manifold intermingled with the $F = 3/2$ levels, the coupling procedure we use to drive nominal $\pi/2$ pulses is characterized by a slightly different effective Rabi frequency for the upper and lower levels. The duration of our coupling pulses is chosen as a compromise, and results in a state mixing of slightly larger than $\pi/2$ for one doublet and slightly smaller than $\pi/2$ for the other. This is the origin of the fringe offset O^0 and this compromise also contributes to the deviation of fringe contrast from unity. Our compromise $\pi/2$ pulse duration was in retrospect chosen imperfectly, such that the deviation from perfect $\pi/2$ duration was not equal and opposite for the upper and lower levels. This led to finite values of O^D and C^D . With the $\pi/2$ pulses not applied at exactly zero detuning, a ϕ^D term also results.

Beyond these large (and largely understood) nonzero fit values, in a perfect version of the experiment we would like to see the remainder of the 480 values be very small, and ideally within measurement error of zero.

3.10.3 Frequencies

Of the remaining 27 frequency channels, there are five that differ from zero by more than 4σ , i.e. by more than twice their scatter-adjusted estimated error.

The four largest are f^R, f^{RI}, f^I , and f^{BI} with frequencies of -4.2, -2.2, -0.8 and -0.42 mHz respectively. What f^R, f^{RI} , and f^I have in common is that they correspond to switches which involve changing the inputs to our direct digital synthesis (DDS) boards. Each fin is driven by a distinct DDS board. To change the sign of rotation of $\vec{\mathcal{E}}_{rot}$, the \tilde{R} -switch, we change the relative digital phases we load into the different DDS boards. To effect the \tilde{I} switch, which has to do with the direction of the the $\vec{\mathcal{E}}_{rot}$ electric field at the instant we do the photodissociation pulse, we switch the sign of the amplitude input into all 8 boards. The digital math performed within the DDS boards must always include at least implicitly a truncation to the least significant bit (LSB) of the digital-to-analog converter, and this process is repeatable but subject to tiny discontinuities from the round-off error. Our fringe frequency is directly proportional to the rotating electric field amplitude. The commanded voltage on any fin is sinusoidal in time. A shift of 4 mHz could result in a change in the magnitude of $\vec{\mathcal{E}}_{rot}$ corresponding to much less than 1 LSB of commanded voltage. In this context, our nonzero values of f^R, f^{RI} and f^I seem unsurprising. f^R is additionally affected by small drifts in the equilibrium location of the center of the ion cloud, and by drifts in the magnitude of residual second harmonic contamination in V_{rot} .

After the discussion above, we are left with 10 \tilde{D} -even and 14 \tilde{D} -odd channels which we believe should be “quiet”, near zero, after a “ $\delta g/g$ correction” for the nonzero value in the corresponding \tilde{D} -even channel, (the correction leaves f^{DR} at 3.0σ from zero, and f^{BDR} at 1.6σ from zero). We have 24 “quiet” channels, and an additional channel f^{BD} , which we believe should be quiet except for any nonzero value of d_e , the electron’s EDM. We don’t have an independent measure of the correct value of d_e , but the 24 channels offer a chance for test of our accuracy and precision claims. We take each quiet frequency channel, divide it by its corresponding sigma, and

then we can ask what is the rms amount by which the items in the ensemble differ from zero? If we omit the three-sigma value, f^{DR} , and look at only the remaining 23 values, we get a very pleasing answer - the frequencies differ from zero by an rms amount of 0.98. If instead we include f^{DR} , the rms normalized difference from zero is 1.3. This is not a catastrophically large number for an ensemble of 24 independent points, but is a larger than one might hope. Put another way, if you have 24 normally distributed points, there is about an 8% chance that at least one will be as far away from its "correct" value as 3σ . It is a thing that happens, sometimes. We have no explanation for why f^{DR} should be nonzero, and we can't rule out that an uncharacterized systematic error contributed to its nonzero value. We put a lot more effort into thinking about and characterizing systematic errors that will move f^{BD} around compared to errors in f^{DR} , so even if one in 24 measured frequencies does have an uncharacterized systematic error, this is not a very worrisome impeachment of f^{BD} . In the end we have chosen to comment on this slightly aberrant observation but to make no corresponding relaxation in our error budget for our main measurement.

3.10.4 Contrasts, and contrast decays

We've already mention C^0 and C^D . The fit value C , the initial contrast in the Ramsey fringes, is affected by parameters of the ion cloud exactly at the moment of photodissociation. The cloud mean position, its mean velocity, and the direction in which Hafnium ions is ejected after photodissociation, all impact where the atomic Hafnium ions eventually impact on the imaging ion detector. Different regions of the multi-channel ion detector have different sensitivity and different propensity to saturate. We have a procedure which region of the detector receives ions ejecting from upper doublet and which from lower doublet, but this procedure can be subject to biases. All this means that changes in contrast can appear in multiple channels. In particular C^{DI} corresponds to a different contrast detected on the right or on the left side of the detector. This is a particularly large term. C^{RI} corresponds to the direction of the HfF^+ velocity, along the direction towards the detector, at the moment of photodissociation. This in turn affects the

mean time-of-flight for the Hf^+ ions and the HfF^+ ions. We impose a hardware gate to cause the image of ions to display only the Hf^+ ions and not the heavier, slower, later-arriving HfF^+ ions. The timing of this gate is designed to maximize count rate from Hf^+ ions but suppress the contrast-destroying HfF^+ ions. Changing the molecule velocity at the moment of dissociation changes the optimum time for the gate, but we do not make any adjustment to the detection circuitry. Therefore a substantial value of C^{RI} comes as no surprise. C^R, C^{DR} and C^{DRI} likely arise from similar but smaller effects. The entire point of including the contrast C as a fitting term is to prevent the nonlinear fitting routine from interpreting changes in contrasts as changes in frequencies. Frequencies are further isolated from misinterpreted changes in contrast (as for instance, if the detector loses more sensitivity due to large-count rate saturation) by ensuring we choose Ramsey times separately for data each switch state so that we take points as close as possible to the exactly halfway up the sides of fringes.

What we call γ is really just a measure of the ratio of the contrast in the early-time fringe to that in the late-time fringe. We don't routinely collect over the intermediate times between the first, "early-time" sinusoidal cycle of the Ramsey fringe and the last "last-time" fringe. When we have filled in some of the intermediate times, the resulting full fringe does not fit well to a pure exponential decay in contrast. The various effects that limit contrast can't be combined in a purely multiplicative way, and thus our approach to fitting (Eq. 3.13) means that any channel with a distinctly nonzero contrast, say C^{RI} , will have a nonzero decay rate, ie γ^{RI} .

3.10.5 Phases

The largest observed nonzero channels in phase and offset (O, O^D, ϕ, ϕ^D) are well understood. Phase errors typically arise from small imperfections in the $\pi/2$ pulses. The observed value of ϕ^{BR} is for instance consistent with the presence of an otherwise imperceptibly small axial gradient in the magnitude of $\vec{\mathcal{E}}_{\text{rot}}$. Other channels of ϕ differ from zero by statistically significant amounts, but all less than 3 mrad. With exception of ϕ^D , no \tilde{D} -odd channel differs from zero by

more than 1 mrad, and the all important ϕ^{DB} is measured to be 0.03(14) mrad. The frequencies we measure are isolated from phase errors by the early fringe/late fringe chop. Our best effort at creating a nonzero fringe phase resulted in a 170 mrad phase, which is 1000 times larger than the 1σ limit on our measured ϕ^{DB} , but even in that case we could resolve no frequency shift in the corresponding frequency channel.

3.10.6 Offsets

Offsets in our Ramsey fringes can occur due to imperfections in $\pi/2$ pulses, as occurs in O^0, O^D and presumably O^{BR} . There are scenarios where offsets in some channels can occur due to minor spatial variations in the polarization of laser beams, such as could happen from a dusty waveplate. This may account for the anomalous value of O^{DIP} . Other than O^0 and O^D , the magnitudes of offset channels are less than one thousandth of the average fringe amplitude. Even we did not allow the fitting routine to relax around a nonzero offset, the spacing of our Ramsey time points is such that it would strongly suppress coupling of offsets to frequencies in our fits.

This marks the end of the manuscript draft

3.11 Not Systematics

In this section I'm going to do a quick review of some effects we observed which are not systematics, but which were interesting/edifying/time-wasting nonetheless.

3.11.1 Shifts vs z position

We have recorded f , f^B , & f^{BR} all varying with the z position of the ions in the trap. Plotting script for the plots in this section are in the X drive in 211129, titled `plotting_z_shim`. We were able to explain f and f^B fairly quickly, but it took us some time to figure out f^{BR} .




Figure 3.18: Observed frequency shifts in z

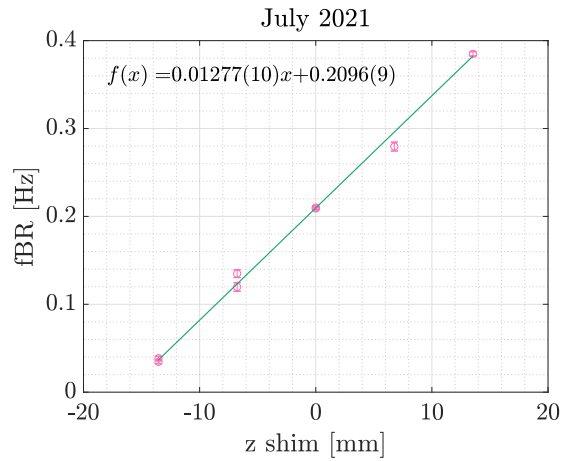


Figure 3.19: Observed frequency shifts in z

3.11.1.1 f vs z

The parabolic shift in $f(z)$ can be explained by constructing a simple model of the axial gradient (\mathcal{B}_{rot}) coils, and expanding the resulting magnetic field out to second order. The field from a single coil is (up to a factor of current and $4\pi\epsilon_0$):

$$k = \frac{4R\rho}{(R + \rho)^2 + (z - z_0)^2} \tag{3.94}$$

where R is the radius of the coil, ρ is the radial distance from the center, and z is the distance along the axis. We can take this expression and add a second term at the correct distance (and

sign for the current) to model an anti-Helmholtz configuration, then expand that expression, with our actual coil parameters substituted, to second order in ρ and z . This gives a term like $(r - Krz^2)B_{\text{axgrad}}$, with $K = 6 \times 10^{-5}$. The fit to the data gives $K = 9 \times 10^{-5}$.




Figure 3.20: Fitting $f(z)$

Note the model is Luke's mathematica notebook titled Baxgrad model which is in 201112 as is the analysis code which generated the $f(z)$ plot, named analyzing_frequencies_vs_zshim.

3.11.1.2 f^B vs z

$f^B(z)$ shows a shift of 0.007 Hz/mm or 0.07 Hz/cm. We can approximate $f^B \sim (3g_F\mu_B/h)B_{\text{rot}}^{nr} = (3g_F\mu_B/h)B_{\text{axgrad}}^{nr}r_{\text{rot}}$, where B^{nr} is the non-reversing magnetic gradient. We can differentiate both sides to get $\frac{\partial f^B}{\partial z} = (3g_F\mu_B/h)\frac{\partial B_{\text{axgrad}}^{nr}}{\partial z}r_{\text{rot}}$. Rearranging, $(\frac{\partial f^B}{\partial z})/(3g_F\mu_B r_{\text{rot}}/h) = \frac{\partial B_{\text{axgrad}}^{nr}}{\partial z}$. Plugging in the known values ($\frac{\partial f^B}{\partial z} = 0.07$ Hz/cm, $r_{\text{rot}} = 0.05$ cm, etc), we get $\frac{\partial B_{\text{axgrad}}^{nr}}{\partial z} = 1 \times 10^{-6}$ G/mm², or $\frac{\partial B_{\text{axgrad}}^{nr}}{\partial z} = 1 \times 10^{-4}$ G/cm². For comparison, $B_{\text{rot}} = 11.9$ mG, which for $r_{\text{rot}} = 0.5$ mm means we have 0.238G/cm B_{axgrad} . So to get our observed change in f^B we need a change in B_{axgrad}^{nr} over 1 cm which is about 1/1000 of B_{axgrad} , which seems quite believable.

It remains to determine the source of $f^{BR}(z)$. The next sections detail the our investigations.

3.11.1.3 f^{BR} vs z : is it from the magnetic field shimming Z coils?

In the Gen 1 trap $f^B(z)$ turned out to be caused by the z coils - turning them off removed the dependence (see notes in lab notebook Nov 17 2016). Simulations of the Gen 1 coils showed that the z coils were causing a quadratic field near $z = 0$. We repeated this experiment in the Gen 2 trap and found that the Z coils being either on or off had no effect on the shifts observed.

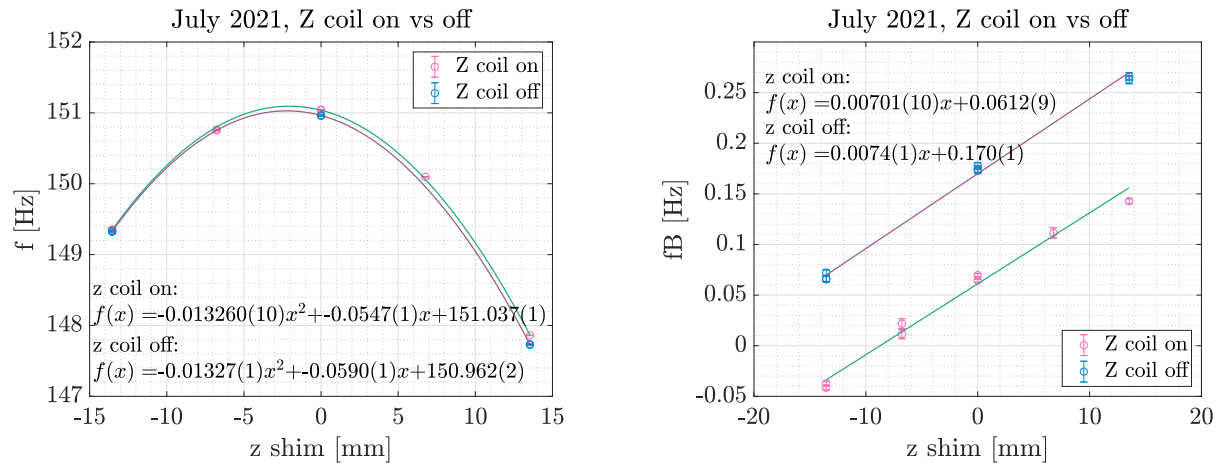


Figure 3.21: Comparing frequency shifts in z with the magnetic field shimming Z coil on or off




Figure 3.22: Comparing frequency shifts in z with the magnetic field shimming Z coil on or off

3.11.1.4 f^{BR} vs z : is it from axial confinement?

Changing the axial confinement changes the spatial extent of the ion cloud in the z direction. We checked to see what kind of effect this might have on the observed frequency shifts. We did not observe any notable changes.




Figure 3.23: Comparing observed frequency shifts in z with 2 different values for the axial confinement




Figure 3.24: Comparing observed frequency shifts in z with 2 different values for the axial confinement

3.11.1.5 f^{BR} vs z : is it from anharmonicity in z ?

We know that our z trapping potential is anharmonic. We found this by giving the cloud varying magnitudes of kicks and looking at the resultant slosh frequency. See the script from Oct 22 2021 titled `trap_anharmonicity_analysis` (lab notebook day Oct 22 2021) (the original data is from 2020). See also the mathematica notebook attached in the 'trap anharmonicity analysis continued' lab notebook, it should be titled 'modelling endcap anharmonicity 3'.

Using Mathematica, I modelled the axial confinement as quadratic with higher order even perturbations (whose strengths are parameterized by α and β , while the perturbation is parameterized by ϵ) plus a z shift (whose strength is parameterized by k_2):

$$U(x) = \frac{1}{2}m\omega_0^2x^2 + \frac{1}{4}m\alpha\epsilon x^4 + \frac{1}{6}m\beta\epsilon x^6 + k_2mx. \quad (3.95)$$

We solve for the equilibrium position (which is a rather complicated expression), then expand our potential about the equilibrium position for small ϵ to get an expression for the potential that we will use which eliminates the shift. Then we take the negative gradient of this expression to get the force:

$$F(x) = -m\epsilon \left(\frac{\beta k_2^5 + \alpha k_2^3 \omega_0^4}{\omega_0^{10}} + \alpha \left(x - \frac{k_2}{\omega_0^2} \right)^3 + \beta \left(x - \frac{k_2}{\omega_0^2} \right)^5 \right) - mx\omega_0^2. \quad (3.96)$$

We are interested in periodic functions of $s \equiv \omega t$, so we write $x(t) = z(s)$, which leads to $\ddot{x}(t) = \omega^2 z''(s)$. We define ω and z as series:

$$\omega = \omega_0 + \epsilon\omega_1 + \dots \quad (3.97)$$

and

$$z = z_0 + \epsilon z_1 + \dots \quad (3.98)$$

then we substitute the series and $x(t) = z(s)$ into our equation of motion:

$$m\omega^2 (z_0'' + \epsilon z_1'') = -m\epsilon \left(\frac{\beta k_2^5 + \alpha k_2^3 \omega_0^4}{\omega_0^{10}} + \alpha \left(z - \frac{k_2}{\omega_0^2} \right)^3 + \beta \left(z - \frac{k_2}{\omega_0^2} \right)^5 \right) - m\omega_0^2 z. \quad (3.99)$$

Then we expand the result for small ϵ to get:

$$m\omega_0^2 z_0'' + \epsilon (2m\omega_0 \omega_1 z_0'' + m\omega_0^2 z_1'') + O(\epsilon^2) = \quad (3.100)$$

$$-m\omega_0^2 z_0 + \epsilon \left(-\frac{5\beta k_2^4 m z_0}{\omega_0^8} + \frac{10\beta k_2^3 m z_0^2}{\omega_0^6} - \frac{10\beta k_2^2 m z_0^3}{\omega_0^4} - \frac{3\alpha k_2^2 m z_0}{\omega_0^4} + \right. \quad (3.101)$$

$$\left. \frac{5\beta k_2 m z_0^4}{\omega_0^2} + \frac{3\alpha k_2 m z_0^2}{\omega_0^2} + \beta(-m)z_0^5 - \alpha m z_0^3 - m\omega_0^2 z_1 \right) + O(\epsilon^2) \quad (3.102)$$

Then we group terms by order of ϵ , and solve each resulting equation using the solution from the previous order. The zeroth order term, $z_0''(t) = z_0(t)$ has the obvious solution:

$$z_0(t) = A \sin(s). \quad (3.103)$$

We substitute this into the equation which is first order in ϵ . We set the integration constant to zero and set the “secular” terms (terms with the form $s \cos(s)$) to zero to get

$$\omega_1 = \frac{5A^4 \beta \omega_0^8 + 6\alpha A^2 \omega_0^8 + 60A^2 \beta k_2^2 \omega_0^4 + 40\beta k_2^4 + 24\alpha k_2^2 \omega_0^4}{16\omega_0^9}. \quad (3.104)$$

Plugging this into the first order equation you get a pretty gnarly expression for $z_1(t)$. Then since we have $z_0(t)$ we can write down $z(t)$ to first order: $z(t) = z_0(t) + \epsilon z_1(t)$. We plug our expression for $z(t)$ into our expression for the force, which becomes pretty gross to look at,

then we expand the expression for the force to first order in ϵ . It's still pretty gross but for the record this is what we get:

$$\begin{aligned}
F(t) = & -Am\omega_0^2 \sin(\omega_0 t) + \\
& \frac{Am\epsilon}{384\omega_0^8} [128\omega_0^2 \cos(t\omega_0) (8A^3\beta k_2\omega_0^4 + 20A\beta k_2^3 + 6\alpha Ak_2\omega_0^4 - 3t\omega_0^8\omega_1) \\
& - 256Ak_2\omega_0^2 \cos(2t\omega_0) (\omega_0^4 (3\alpha + 5A^2\beta) + 10\beta k_2^2) + 256A^3\beta k_2\omega_0^6 \cos(4t\omega_0) - \\
& 6 \sin(t\omega_0) (A^2\omega_0^8 (46\alpha + 37A^2\beta) + 4k_2^2\omega_0^4 (48\alpha + 115A^2\beta) + 320\beta k_2^4) \\
& + 27A^2\omega_0^4 \sin(3t\omega_0) (\omega_0^4 (4\alpha + 5A^2\beta) + 40\beta k_2^2) - 25A^4\beta\omega_0^8 \sin(5t\omega_0)] + O(\epsilon^2)
\end{aligned} \tag{3.105}$$

Taking the electric field $E_z(t)$ as the force over the charge, and integrating

$$\frac{-1}{6} \frac{\omega}{2\pi} \int_0^{2\pi/\omega} \frac{E_z^3(t)}{E_{rot}^3} \tag{3.106}$$

then expanding the result as a series in ϵ , we get an expression for the Berry's phase we expect as a function of displacement (in terms of k_2) and the strength of the quartic (parameterized by α (we also let ϵ go to 1 as we finish with the perturbation theory.) This gives

$$\alpha_{berry} = - \frac{m^3\epsilon (5A^6\beta k_2\omega_0^4 + 10A^4\beta k_2^3 + 3\alpha A^4 k_2\omega_0^4)}{12E_{rot}^3 q^3 \omega_0^2} \tag{3.107}$$

If we plug in the known values for everything in the above equation plus $x_{eq}\omega_0^2$ for k_2 , where X_{eq} is the equilibrium position, we get that the berrys phase alpha is

$$\alpha_{berry} = x_{eq} \left(\beta (-3.722534205999362 x_{eq}^2 - 1.861267102999681 x_{eq}^{-19}) - 1.1167602617998085 x_{eq}^{-17} \alpha \right) \tag{3.108}$$

fBR is then $375000 \times 3\alpha_{berry}$ so the fBR due to this quartic perturbation should scale like

$$f^{BR} = 375000 \times 3 \times \text{xeq}(-1.12 \times 10^{-17}\alpha + \beta(-1.86 \times 10^{-19} - 3.72 \times 10^{-17}\text{xeq}^2)) \quad (3.109)$$

Since we have expressions for ω_0 and ω_1 we can say that the frequency is

$$f = \frac{1}{2\pi}(\omega_0 + \omega_1) \quad (3.110)$$

and for the data where we gave the cloud a kick and measured the slosh amplitude A and frequency then the equilibrium position is zero, so $k_2 = 0$ and our expression for the frequency is

$$f = \frac{5A^4\beta + 6\alpha A^2 + 16\omega_0^2}{32\pi\omega_0} \quad (3.111)$$

We then fit the slosh vs frequency data we have to determine ω_0 , α and β .




Figure 3.25: Fitting the measured frequency vs slosh amplitude in the non-displaced trap

Now we can plug the values for ω_0 , α , and β into equation 3.109:

$$f^{BR} = -0.000154x_{eq} - 2.677x_{eq}^3 \quad (3.112)$$

We can plot this prediction on a plot with our experimental results, and we see that this completely fails to account for the observed shift; the slopes differ by a factor of ~ -10 .




Figure 3.26: Comparing the prediction for f^{BR} from our model of the anharmonicity to the empirical results. We see that this model does not account for the observed shifts

3.11.1.6 f^{BR} vs z: we're pretty sure it's due to currents in the electrodes

Trevor Wright made a simple model of the currents in the electrodes which explained what we observed quite well. Look forward to his thesis for the write-up :)

In any case, we have no mechanism by which this could become a systematic.

3.11.2 Ion number vs f^0

At one point we saw what appeared to be a correlation between total trapped HfF^+ number and f^0 , but this has not held true through the dataset. What seems to happen reliably is that f^0 tends to drop over a day whether or not HfF^+ number is dropping, and some days HfF^+ number would also drop throughout the day so we saw false correlations.

3.11.3 f^R

In general we measure a non-zero f^R , the value depends on what primary fringe frequency we are operating at (see Figure 3.27). We suspect this is just due to the DAC resolution on our \mathcal{E}_{rot} generating channel; it's only a 10-bit DAC so a difference in a single DAC step would be enough to provide the f^R we observe.


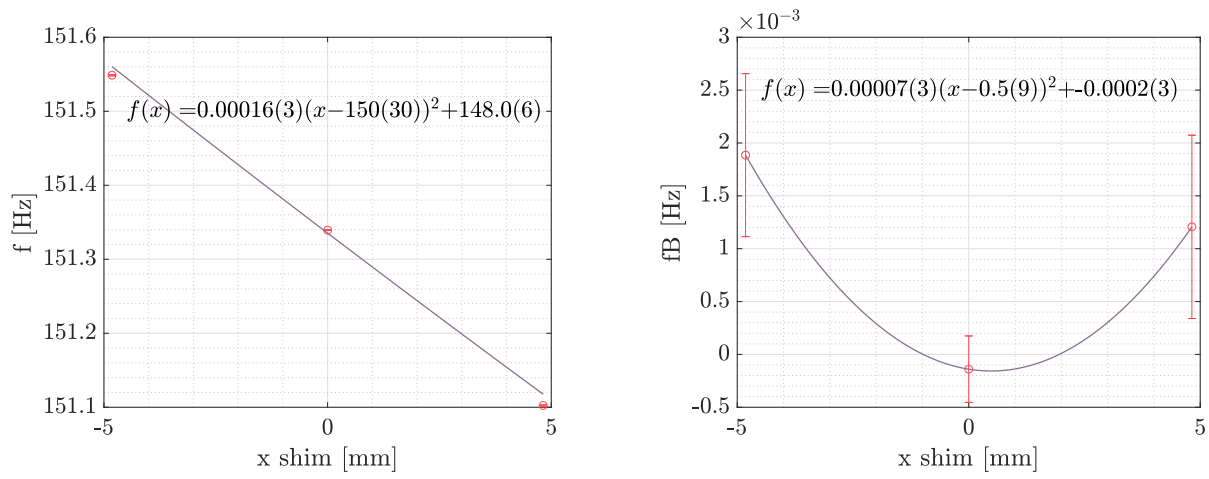



Figure 3.27: Averaging f^R over our entire dataset, per primary fringe frequency f^0 . f^R scales with f^0 with a slope of 0.00018, which is very close to $1/1024 = 0.00098$ or a single step in a 10-bit DAC.

3.11.4 Shifts vs x-y position

We observe frequency shifts if we move the ions in the x-y direction (see Figures 3.28 and 3.29). We have no clue why! Baxgrad gradients would explain f , but f^R is a little harder to understand. This one remains a mystery, but we have no mechanism by which this could be a systematic⁴.

⁴plotting script is in 220125\analyze_radial_shifts.

Figure 3.28: Comparing observed frequency shifts with x translationFigure 3.29: Comparing observed frequency shifts with x translation

Chapter 4

A new limit on the electron's electric dipole moment

This chapter is a reproduction of a manuscript in preparation co-written by myself and Luke Caldwell

4.1 Summary of the eEDM result

Physicists have been able to describe our entire universe using just two mathematical models: General Relativity, for gravity, and the Standard Model, for everything else. While the Standard Model is generally successful it has some serious inconsistencies: it fails, for example, to explain why matter dominates over antimatter in our universe. New physics models aim to eliminate these inconsistencies and often introduce additional symmetry violation, which in turn results in predicted symmetry violation in fundamental particles like electrons. In particular, many new physics models predict that the electron will have a measurable non-zero electric dipole moment aligned with the spin axis, which violates time-reversal symmetry. Here we present a new experimental limit on the electron's electric dipole moment using trapped HfF^+ in rotating bias fields. We find $d_e = (-9.7 \pm 20.1_{\text{stat}} \pm 6.2_{\text{syst}}) \times 10^{-31}$ e·cm, which corresponds to an upper bound of $|d_e| < 4.4 \times 10^{-30}$ e·cm (90% confidence). Our measurement is a factor of 2.5 better than the previous best limit, and confirms previous results in a radically different system. Our limit directly constrains new physics theories at energies beyond the reach of the Large Hadron Collider [32].

The Standard Model of physics predicts that the electron has a very small displacement of its center of charge from its center of mass, giving it an electric dipole moment (EDM) which violates time-reversal (T) symmetry. So far, nobody has detected an electron electric dipole moment (eEDM, d_e), and the Standard Model prediction lies ~ 8 orders of magnitude below

current experimental sensitivity. Extensions to the Standard Model aimed at solving inconsistencies like matter-antimatter asymmetry often do so by adding new particles at energies higher than any so far discovered, accompanied by new T-violating phases. These new particles result in a much larger prediction for the eEDM - often within reach of near-term experiments. Measurements of the eEDM serve as a direct constraint on extensions to the Standard Model, and current eEDM measurements probe for particles with energies above ~ 30 TeV [32], well beyond the reach of any current or planned particle collider. In addition, a non-zero measurement at current experimental sensitivities would unambiguously signal new physics. Here we present the most precise measurement of the eEDM so far. We achieved this with a unique experimental setup, trapping molecular ions and polarizing them in rotating bias fields. Our result, $d_e = (-9.7 \pm 20.1_{\text{syst}} \pm 6.2_{\text{stat}}) \times 10^{-31} e \cdot \text{cm}$ ('syst', systematic uncertainty; 'stat', statistical uncertainty), is consistent with zero and constrains d_e to an upper limit of $|d_e| < 4.4 \times 10^{-31} e \cdot \text{cm}$ at 90% confidence.

In general, eEDM searches aim to measure the energy difference ΔU when \mathbf{d}_e is aligned vs anti-aligned with an electric field \mathcal{E} ; $\Delta U = -2\mathbf{d}_e \cdot \mathcal{E}$. Assuming there are no new degrees of freedom the eEDM must be aligned with the electron spin, and an unpaired electron spin inside a polarized paramagnetic molecule is subject to a huge intramolecular electric field \mathcal{E}_{eff} [41]. This is why our measurement uses trapped HfF^+ molecules; with an applied field of ~ 50 V/cm the science states of our molecule are oriented and fully polarised ($\mathcal{E}_{\text{eff}} = 23$ GV/cm). By using trapped ions we are able to realize spin coherence times 3 orders of magnitude larger than any other molecular eEDM experiment. As shown in Figure 4.1, the intramolecular axis, which defines the direction of \mathcal{E}_{eff} , points either parallel or antiparallel to the applied electric field. A differential measurement performed between these two orientation states is at the heart of our experiment.

4.2 Experimental Overview

We begin our experiment with the production and radiofrequency trapping of roughly 20,000 HfF^+ ions in the $^1\Sigma_+$ ground state [55] [50]. In order to both polarize and trap our ions we rotate our external bias field \mathcal{E}_{rot} at angular frequency $\omega_{\text{rot}} = 2\pi \times 375$ kHz and do our spectroscopy in the rotating frame. The rotating electric field polarizes the ions, which have a molecule-frame dipole moment, and states of good parity are mixed into states of good orientation (either aligned or anti-aligned with \mathcal{E}_{rot}), known as Stark doublets (see Figure 4.1, top). We also apply a quadrupole magnetic field gradient \mathcal{B} , which drives the primary component of the spin precession (see Figure 4.2).

Once we have trapped our ions and applied our bias fields, we prepare them for spectroscopy. We perform Ramsey spectroscopy on the stretched ($m_F = \pm 3/2$) states in the $^3\Delta_1$, $J = 1$, $F = 3/2$ manifold of HfF^+ . As described in Methods, we transfer our ensemble into $^3\Delta_1$ then prepare a pure spin state before applying a $\pi/2$ pulse to create a coherent superposition of the $m_F = \pm 3/2$ states. We allow the superposition to evolve for a variable amount of time, then we apply a second $\pi/2$ pulse to map the relative phase onto a population difference between the $m_F = \pm 3/2$ states. We clean out one of the stretched states, then we detect and count the number of ions in the remaining stretched state by state-selectively photodissociating the molecules [19], using the opposing orientations in the trap to send the dissociation products to opposite sides of our imaging microchannel plate and phosphor screen assembly [48, 47]. We count the ions in each doublet then repeat the measurement reading out the other stretched state. From these two measurements we can construct a spin asymmetry, $\mathcal{A}_{u/l} = \frac{N_{\uparrow} - N_{\downarrow}}{N_{\uparrow} + N_{\downarrow}}$, where the subscripts u/l indicate the upper or lower Stark doublet, and the subscripts \uparrow / \downarrow indicate the spin state / stretched state we read out. We repeat our measurement at different free-evolution times, which generates a Ramsey fringe whose frequency is proportional to the energy splitting between stretched states (see Figure 4.1).

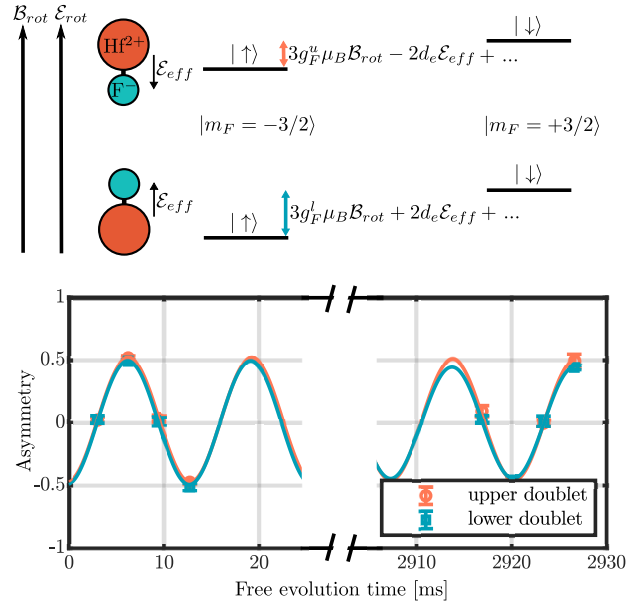


Figure 4.1: Ramsey spectroscopy in HfF^+ . Top, level structure in the eEDM-sensitive $^3\Delta_1$, $J = 1$, $F = 3/2$ manifold in external electric \mathcal{E}_{rot} and magnetic \mathcal{B}_{rot} bias fields (not to scale). The oriented molecules are either aligned or anti-aligned with \mathcal{E}_{rot} , which we call the lower and upper Stark doublets respectively. The component of the energy splitting due to the eEDM, $2d_e \mathcal{E}_{\text{eff}}$, flips sign between doublets. Bottom, example Ramsey fringes from our dataset. The fringes from the two doublets are collected simultaneously.




Figure 4.2: Top down schematic of our experiment's main chamber. The neutral ions fly into the center of the trap where they are ionized and trapped. We apply the rotating electric bias field $\vec{\mathcal{E}}_{\text{rot}}$ and the quadrupole magnetic field gradient $\vec{\mathcal{B}}$. The ions undergo uniform circular motion, either aligned or anti-aligned with $\vec{\mathcal{E}}_{\text{rot}}$, and their unpaired spins precess about $\vec{\mathcal{E}}_{\text{rot}} \cdot \vec{\mathcal{B}}$.

In practice, our experiment is an electron spin precession measurement simultaneously executed on two orientation states (Stark doublets, \tilde{D}) of trapped HfF^+ . The primary component of the spin precession frequency that we measure, f_B , is due to the spin precessing about a laboratory magnetic field \vec{B} : $hf_B = 3g_F^{u/l}\mu_B\hat{\mathcal{E}}_{\text{rot}} \cdot \vec{B} = 3g_F^{u/l}\mu_B\mathcal{B}_{\text{rot}}$, where h is Planck's constant, g_F is the magnetic g-factor, the superscripts u/l indicate the upper or lower Stark doublet, μ_B is the Bohr magneton, and \mathcal{B}_{rot} is the magnetic field in the ion's rotating frame. As indicated in Figure 4.1, the relativistically enhanced eEDM-induced energy shift, $hf_{\text{EDM}} \sim d_e\mathcal{E}_{\text{eff}}$, adds or subtracts from the primary frequency hf_B depending on the relative orientation of \mathcal{E}_{eff} and \mathcal{B}_{rot} (aligned or anti-aligned). We take measurements in both orientation states **simultaneously** but **separately** (on either side of the screen), which allows us to perform a differential measurement between the \tilde{D} states (cancelling most common-mode noise [47]).

We collect our Ramsey fringes in each one of 8 possible experimental states, where we switch the sign of the magnetic field \vec{B} (in or out), the rotation direction $\tilde{\mathcal{R}}$ (up or down), and the side of the screen/image we read a given Stark doublet out on $\tilde{\mathcal{I}}$ (left or right). A set of Ramsey fringes in all possible switch states is called a *block*, and we fit those fringes to a set of parameters (frequency, phase, contrast, offset and decay rate). Using all the fringes in a block we form linear combinations of the fitted fringe parameters to create what we call *parity channels*. The \tilde{B}, \tilde{D} -odd frequency channel, f^{BD} , contains the eEDM signal. The other channels allow us to diagnose systematics and monitor experimental performance.

We collected our blocks in “superblock” chunks where we changed an experimental parameter not listed above for a large set of blocks to search for unknown systematic effects. Our superblocks consist of: different magnitudes of \mathcal{B} to adjust the average fringe frequency f^0 (we took data at $f^0 = 77, 105$ and 151 Hz), inverse orientations of the fixed waveplates used for state preparation and readout (to change the handedness of the light in the lab frame), and changing which stretched state we prepare (we routinely read out both stretched states, and we interleaved blocks preparing one or the other stretched state). During data collection and analysis, we

‘blinded’ our measurement to prevent experimenter bias from affecting our results. We generated a random number from a uniform distribution, and when frequency channels are computed the random number is added to the f^{BD} channel before any results are displayed. We did not remove the blind until our systematics search and analysis was complete.

4.3 Accuracy evaluation

To evaluate the accuracy of our measurement we searched extensively for systematic shifts prior to data collection; full details are in 3 and our systematic error budget is delineated in Table 4.1. In general, we tuned a variety of experimental parameters over a large range and observed the response in our data channels. The only shift we were able to observe directly in the eEDM channel is that which is due to a non-reversing quadrupole magnetic field and the difference in magnetic g-factors g_F between Stark doublets δg_F . The f^B channel, which is acquired concurrently with the f^{BD} channel, provides a direct measurement of the non-reversing magnetic field and allows us to apply a correction to our science channel: $\delta f_{\text{corr}}^{BD} = f^B \frac{\delta g_F}{g_F}$, see Figure 4.3. Before applying any corrections to the science channel, we suppress this systematic by actively shimming the currents through the \mathcal{B} coils in each direction to minimize f^B . This shimming was so effective that the mean correction we applied was well below our statistical sensitivity: $\delta f_{\text{corr}}^{BD} = 181$ nHz. Unfortunately, this type of shimming and correction leaves us susceptible to other possible effects which can cause shifts in f^B but not f^{BD} . One of the most important sources of such a shift was the combination of a transverse magnetic field and an \mathcal{E}_{tot} second harmonic caused by harmonic distortion in our op-amps. This combination can cause shifts in both f^B and f^{BR} . We shim out the second harmonic on each electrode by feeding forward a second harmonic signal with the opposite phase. With this technique, we are able to suppress the second harmonic by ~ 80 dB. We also used our magnet coils to null the ambient uniform magnetic field at the RF trap center below 10 mG. All remaining known systematic shifts are enumerated in Table 4.1.

Table 4.1: Systematic error budget

Effect	Correction (μHz)	Uncertainty (μHz)
Magnetic		
Non-reversing \vec{B}^0	0.2	< 0.1
Second harmonic of \mathcal{E}_{rot} and transverse magnetic field		2
Third harmonic of \mathcal{E}_{rot} and magnetic field gradients		2
Higher harmonics of \mathcal{E}_{rot} and higher order magnetic field gradients		< 0.1
Ellipticity of \mathcal{E}_{rot} and magnetic field gradients		1
Berry's phase		
Phase modulation due to axial secular motion		3
Axial 2nd harmonic with ellipticity of \mathcal{E}_{rot}		2
Higher harmonics of \mathcal{E}_{rot}		3
Rotation induced mixing		
Leaking of f^{BR}		< 0.2
Axial magnetic fields		0.4
Other frequency		
Imperfect overlap of spatial distributions		3.5
Phase		
Improperly characterized imaging contrast		0.7
Swatch position		1.2
Total	0.2	7




Figure 4.3: Systematic shift in the eEDM channel f^{BD} due to non-reversing axial magnetic field, which shows up in the f^B channel.

4.4 Results

We collected 1370 blocks over roughly two months. We applied cuts to the blinded data based on non-EDM channels indicating signal quality: blocks with late-time contrast below 0.2 were cut due to low signal to noise, as were blocks containing fitted fringe frequencies which were over 3.5σ different from the mean fringe frequency for that switch state. The cuts applied did not shift our result, see Figure 4.4. After applying cuts we were left with 1328 blocks, or ~ 620 hours of data.

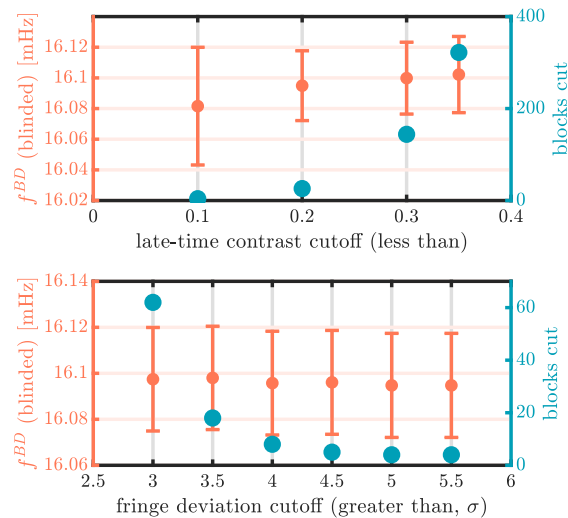


Figure 4.4: Results as we change the threshold for cutting blocks from our dataset. In each plot, only one cut is applied (the other cut is not applied). Applying cuts does not systematically shift our result.

Each block results in one value of f^{BD} and thus one eEDM measurement. In general our fringe measurement uncertainty came in at the quantum projection noise limit [47], and the reduced chi-squared statistic for taking a weighted mean of all the blocks is $\chi_r^2 = 1.07$. To correct for the over scatter we scale our final statistical error bar by $\sqrt{\chi^2} = 1.03$. The unblinded dataset is shown in Figure 4.5. Inspection of a normal probability plot indicates that the distribution of normalized and centered eEDM measurements $(f^{BD} - \langle f^{BD} \rangle) / \delta f$ is consistent with a normal distribution.

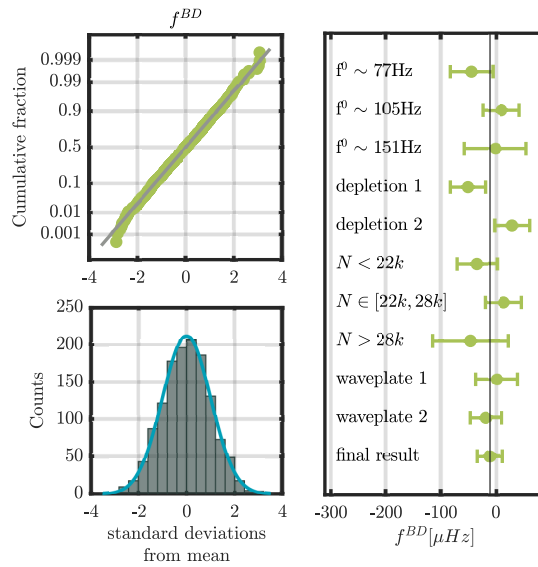


Figure 4.5: Summary of our dataset after cuts have been applied and scaling δf by $\sqrt{\chi^2}$ to account for over-scatter. Top left, normal probability plot of f^{BD} showing the data are consistent with a normal distribution. Bottom left, histogram of normalized and centered f^{BD} data. Overlaid in blue is a normal distribution. Right, superblocks taken under different experimental parameters compared to the overall average value of f^{BD} . Here N is the average number of trapped HfF^+ ions per experimental trail.

We removed our blind on November 1, 2022 and obtained a final value for the eEDM-sensitive frequency channel

$$f^{BD} = -10.9 \pm 22.8_{\text{stat}} \pm 7_{\text{syst}} \mu\text{Hz}. \quad (4.1)$$

Dividing by $-2|\mathcal{E}_{\text{eff}}|\text{sgn}(g_F)/h \simeq 1.13 \times 10^{31} \mu\text{Hz}/e\cdot\text{cm}$, we obtain a value for the eEDM

$$d_e = (-9.7 \pm 20.1_{\text{stat}} \pm 6.2_{\text{syst}}) \times 10^{-31} e \cdot \text{cm}, \quad (4.2)$$

which is consistent with zero within one standard error.

The resulting upper bound on d_e is

$$|d_e| < 4.4 \times 10^{-30} e \cdot \text{cm} \text{ (90\% confidence)}, \quad (4.3)$$

Our result is consistent with the ACME collaboration's 2018 findings [32], and confirms their result with a radically different experimental approach. So far we have assumed that parity and time-reversal violation arise purely from d_e . Diatomic molecules are sensitive to other symmetry violating effects [56], and we can interpret our measurement as a linear combination of them all. Assuming that d_e is zero, we can instead attribute our measurement of f^{BD} to $2W_S C_S$, which arises from a pseudoscalar-scalar electron-nucleon coupling C_S [57] [58] (here the molecule-specific structure constant $W_S = 20.4$ kHz [59]). Dividing f^{BD} by $2W_S$, we find

$$C_S = (-2.7 \pm 5.6_{\text{stat}} \pm 1.7_{\text{syst}}) \times 10^{-10}. \quad (4.4)$$

4.5 Conclusions

Our measurement imposes constraints on beyond standard model theories, in a regime largely inaccessible to particle accelerators [32]. The standard model predictions for both d_e and C_S are well below current experimental sensitivity, making this measurement a background-free probe for new physics. In addition, our measurement confirms previous results [32] in a radically different system, improving confidence in both experiments. Measurements like ours can also be used to constrain possible couplings to dark matter [60].

4.6 Methods

4.6.1 Bias fields

We apply a rotating electric field $\mathcal{E}_{\text{rot}}(\mathbf{R}, t) = \mathcal{E}_{\text{rot}}[\hat{X} \cos(\omega_{\text{rot}}t) - \tilde{R}\hat{Y} \sin(\omega_{\text{rot}}t)]$, where $\omega_{\text{rot}} \equiv 2\pi f_{\text{rot}}$, $f_{\text{rot}} = 375$ kHz, and \tilde{R} indicates the rotation direction (up or down). The rotating field orients and polarizes the ions, which are driven in uniform circular micromotion with a radius $r_{\text{rot}} \sim 0.5$ mm.

Our experiment has 4 pairs of magnetic field coils and no magnetic shielding. All coils are outside of the main vacuum chamber. Three of the coil pairs are set up along our lab-frame X, Y, Z axes in Helmholtz configuration for tuning the magnetic field at the position of the molecules. The magnetic field at the center of the trap is measured by an array of fluxgate magnetometers bolted to the main experimental chamber. The fourth coil pair generates our so-called ‘rotating’ magnetic field \mathcal{B}_{rot} . The coils are aligned with the lab-frame Z axis and produce the quadrupole magnetic gradient $\mathcal{B} = \mathcal{B}'_{\text{axgrad}}(2Z - X - Y)$. In the rotating frame of the ions this appears as a magnetic bias field $\mathcal{B}_{\text{rot}} \equiv |\langle \mathcal{B} \cdot \mathcal{E}_{\text{rot}} / \mathcal{E}_{\text{rot}} \rangle| \simeq |\mathcal{B}'_{\text{axgrad}} r_{\text{rot}}|$.

4.6.2 $\text{HfF}^{+3}\Delta_1$ details

Once we have trapped our ions and applied our bias and shimming fields, we prepare them for spectroscopy. We perform Ramsey spectroscopy on the stretched ($m_F = \pm 3/2$) states in HfF^{+} 's $^3\Delta_1$, $J = 1$, $F = 3/2$ manifold. Here, J indicates the electronic plus rotational angular momentum, F indicates the total angular momentum, and $m_F = \mathbf{F} \cdot \hat{z}$ is the projection of \mathbf{F} onto the rotating quantization axis. The $^3\Delta_1$ state is well-described by Hund's case a) basis states with coupled nuclear spin from the Fluorine nucleus.

4.6.3 Quantum state preparation

Our state-selective two-photon ionization procedure produces HfF^+ in the ground electronic and vibrational state ($^1\Sigma^+, v = 0$), primarily distributed over the lowest 4 rotational levels. We connect $J = 0 \rightarrow 3$ using microwaves and perform incoherent transfer to the eEDM-sensitive $^3\Delta_1, v = 0, J = 1$ state via the $^3\Pi_{0+}, J = 0 \leftarrow ^1\Sigma^+$ transition at ~ 961 nm, which decays preferentially to $^3\Delta_1$. The decay from $^3\Pi_{0+}$ puts population in several vibrational levels in $^3\Delta_1$, which will decay into our ‘science’ state $^3\Delta_1, v = 0$ if left untreated. We clean the population in $^3\Delta_1, v = 1$ out by illuminating the cloud with radiation at ~ 818 nm which connects $^3\Sigma_{0+}^-, v = 1, J = 0 \leftarrow ^3\Delta_1, v = 1, J = 1$ and preferentially decays back to $^1\Sigma^+$. This laser is on for the duration of the experiment, and will clean out any population which leaks down from higher vibrational levels over the course of the measurement. The lasers used in our experimental sequence are summarized in Figure 4.6.

After transfer to $^3\Delta_1$ we turn on a strobed, circularly polarized, CW laser at ~ 1082 nm which is brought in from the side of the trap and connects $^3\Pi_{0-}, v = 0, J = 0 \leftarrow ^3\Delta_1, v = 0, J = 1$. The strobing is synchronous with the rotation of \mathcal{E}_{rot} such that the wavevector is either parallel or anti-parallel to \mathcal{E}_{rot} , which drives a $\sigma^{+/-}$ transition to an $F' = 3/2$ manifold in $^3\Pi_{0-}$, eventually leaving only one stretched state populated in $^3\Delta_1$ (in both Stark doublets).

4.6.4 Ramsey spectroscopy

Once we have a pure spin state in both doublets prepared, we apply a $\pi/2$ pulse to the ensemble of ions by temporarily ramping down the magnitude of \mathcal{E}_{rot} , which increases a rotation-induced coupling $\Delta^{u/l}$ between $m_F = \pm 3/2$ states [61]. The increased coupling causes the pure spin states to evolve into a coherent superposition of the $m_F = +3/2$ and $m_F = -3/2$ states in ~ 1 ms [43] [50]. We ramp \mathcal{E}_{rot} back up and allow the superposition to evolve for a variable amount of time. After the desired precession time has elapsed, we apply a second $\pi/2$ pulse to map the relative phase onto a population difference between the $m_F = \pm 3/2$ states. We project

the ions into their final state by applying the ‘depletion’ laser, which connects ${}^3\Delta_1, v = 0, J = 1 \rightarrow {}^3\Sigma_{0+}^-, v = 0, J = 0$ at ~ 814 nm, and preferentially decays to ${}^1\Sigma^+$ by a 10:1 ratio having weaker coupling to the ${}^3\Delta_1$ state. This laser is circularly polarized and strobed at f_{rot} such that it is parallel/antiparallel to \mathcal{E}_{rot} , which drives a $\sigma^{+/-}$ transition to ${}^3\Sigma_{0+}^-$, leaving only one stretched state populated.

4.6.5 Final state readout

Finally, we detect and count the number of ions in the stretched states via state-selective resonance-enhanced multiphoton dissociation [19] using two pulsed lasers, both of which are circularly polarized and timed such that their k-vectors are parallel/antiparallel to \mathcal{E}_{rot} to drive $\sigma^{+/-}$ transitions. The ions are oriented in the trap thanks to \mathcal{E}_{rot} , and we perform dissociation when they are parallel to our imaging microchannel plate (MCP) and phosphor screen assembly. Because our dissociation transitions obey the angular momentum requirements detailed in [13], the dissociated Hafnium atoms from the two Stark doublets will fly to opposite sides of the screen. We gate the phosphor screen such that we only image the Hf^+ atoms, the un-dissociated HfF^+ ‘spectator’ ions would otherwise ruin the image. We detect both the Hf^+ and the HfF^+ ions in time of flight. The technical details of our imaging and counting system are detailed in [48] and [47].




Figure 4.6: Cartoon depicting the transitions used during our state preparation.

We repeat our measurement at different free-evolution times, which generates a Ramsey fringe whose frequency is proportional to the energy splitting between stretched states. Because we simultaneously measure both doublets, we get two Ramsey fringes. We can prepare one of 2 stretched states and read out one of 2 stretched states, and we *chop* between those possible preparation and readout choices by strobing the circularly polarized pumping and depletion lasers either in phase or out of phase with \mathcal{E}_{rot} to drive σ^+ or σ^- transitions. There are 4 possible combinations: in/in, in/out, out/in, and out/out.

4.6.6 d_e signal extraction

When we take measurements we do them in one of 8 possible experimental *switch states*, where we switch the sign of the magnetic field ($\vec{\mathcal{B}}$, *in* or *out*), the rotation direction ($\vec{\mathcal{R}}$, again *in* or *out*), and the side of the screen/image we read a given Stark doublet out on ($\vec{\mathcal{I}}$, *left* or *right*). Note, we measure both doublets ($\vec{\mathcal{D}}$) **simultaneously**, but **separately** (on either side of the screen).

We take a set of Ramsey fringes in all possible switch states (we call this a *block*). We fit each fringe in the block to extract the fringe parameters (frequency f , phase ϕ , contrast C , decay rate γ , and offset O), and we form linear combinations of those parameters to create what we call *parity channels*. Taking linear combinations of the frequencies in a block gives us *frequency channels*, which are of primary concern. For example, the mean of all the measured frequencies in a block is $f^0 \sim 3g_F\mu_B\mathcal{B}_{\text{rot}}/h$. If instead we take the difference over the $\vec{\mathcal{B}}$ switch we get $f^B \sim 3g_F\mu_B\delta\mathcal{B}_{\text{nr}}/h \sim f^0(\frac{\delta\mathcal{B}_{\text{nr}}}{\mathcal{B}_{\text{rot}}})$, which measures how accurately the magnetic field is reversed. f^D is the difference in frequency between the two Stark doublets: $f^D \sim -3\delta g_F\mu_B\mathcal{B}_{\text{rot}}/h$. If the g-factor was identical for each doublet and the rotation-induced coupling terms were identical, the leading order f^D term would vanish. In general, we indicate the odd parity components with a superscript, so f^{BR} is the antisymmetric combination of switching both the magnetic field and the rotation direction. Forming frequency channels allows us to diagnose experimental issues and get a handle on many systematics.

The frequency channel of primary interest contains the eEDM signal: $f^{BD} \sim 2d_e \mathcal{E}_{\text{eff}}/h - 3\delta g_F \mu_B \delta \mathcal{B}_{nr}/h + \dots$. Here, the subscript on $\delta \mathcal{B}$ indicates a quadrupole magnetic field gradient $\delta \mathcal{B}_{nr}$ that does not switch with our applied switching quadrupole field gradient.

Because we collect data in both doublets simultaneously, we are able to take differential measurements which cancel out most common-mode noise. After we define the asymmetry $\mathcal{A}_{u/l} = \frac{N_{\uparrow} - N_{\downarrow}}{N_{\uparrow} + N_{\downarrow}}$ in the upper and lower doublets, we can take the difference D and sum S over the doublets:

$$\begin{aligned} \mathcal{A}_D &= \mathcal{A}_u - \mathcal{A}_l \\ \mathcal{A}_S &= \mathcal{A}_u + \mathcal{A}_l, \end{aligned} \tag{4.5}$$

which allows us to fit the difference and sum fringes to reduce uncertainty in our \tilde{D} -odd channels. Full details about our fitting and analysis routine are in 3, and a discussion of the advantages of our simultaneous measurement scheme is in [47].

Chapter 5

Dark Matter Exclusion

5.1 Introduction

Sometime in 2017 Victor Flambaum approached us with a curious idea: to use our first-generation EDM measurement as the basis for a dark matter search. Dark matter? Indeed. As it turns out, a certain class of dark matter could leave an oscillatory echo in our spectroscopic signal - in short, even though the DC value of the transition we measured was zero it could still have a non-zero AC amplitude.

One might reasonably ask - why use AMO/precision measurements to try and detect dark matter at all? Isn't new particle detection an accelerator's job? And what about dedicated




Figure 5.1: Generation 1 EDM data vs time. Could there be an oscillatory signal buried in here?

dark matter detection experiments? To answer the first question, I think it's fair to predict that accelerator energies are likely to remain on a plateau for quite a while to come. So if the particles we are searching for are higher energy than the reach of the LHC, the dark matter detective is out of luck. So we really need to look outside the accelerator paradigm to solve the issue of dark matter. Which leads us to the next question: what about dedicated direct detection searches? Well, of course there are plenty of those, and they're all interesting and worth doing. But one big advantage that AMO/precision measurements have over their direct detection counterparts is their ability to exclude huge portions of the parameter space (in a relative sense) at once. For example, our exclusion covered seven orders of magnitude: from $10^{-22} - 10^{-15}$ eV in mass. Competing direct detection searches are often limited to tiny fractions of a single decade (see, for example, this paper [62] from the ADMX collaboration, which excluded axion dark matter in the 3.3-4.3 eV mass range). There are, of course, some direct detection experiments that have large mass ranges, but while we still have no idea what dark matter is and we are still searching over 40 mass decades, it seems prudent to expand our toolbox to include a few tools that allow for broadband searches.

Separate from the question of the width of the search in mass range, there is also the depth of the search in coupling strength, which we can call our "sensitivity". Because AMO experiments are not designed specifically for dark matter detection they can be (and in our case, was) much less sensitive than one might hope. But again, I think pursuing these searches is still worth it - simply because of how quickly precision measurements are becoming more sensitive. The electron EDM limit has moved a factor of 100 in less than 10 years, and one can see similarly impressive progress with things like atomic clocks. There are tons of precision measurements like this out there, and a lot of them are so sensitive now that they are really well-positioned to make a big dent in the search for dark matter. This is exciting, and something that I think we as a field should take advantage of. Assuming we do, this leads us to the most important question of all: at the threshold of this era, we should not be asking 'how is the analysis of hypothesis

exclusion typically done?’ but rather ‘how can we get the analysis of hypothesis exclusion *right*?’

You see, after taking on this project there was about a year of confusion and suffering trying to understand dark matter (in general) and statistical exclusion methods for dark matter searches (in particular). Part of the difficulty was that it became clear that nobody really agreed on how to do dark matter exclusions, and based on my preliminary and fairly shaky analysis it seemed like some groups were doing stuff that was patently *wrong*. Most of the stuff I read about was fairly standard threshold-based frequentist methods. One issue I noticed right away was that none of those methods had a particularly appealing way of dealing with the look-elsewhere effect, and after more reading and struggling it seemed like none of the methods I looked into had a reasonable way of dealing with the stochastic nature of the dark matter field itself. There were more issues: the scale of the stochasticity changing as your measurement time approaches the dark matter coherence time, the differential sensitivity of different measurement quadratures, and even the finite sample time for fast hypothesized oscillations. All of these had us (Eric and I) spinning our wheels somewhat, trying to make it all work using threshold-based frequentist methods.

Ultimately, and thankfully, we teamed up with Dan Palken and Ben Brubaker from the Lehnert group for their much-needed expertise on statistical methods (Dan) and dark matter (Ben). Dan ultimately confirmed my suspicion about the general confusion and disagreement and was able to demonstrate exactly what some groups were doing wrong. I’d recommend reading his thesis for all the gory details [63]. Dan introduced us to a Bayesian exclusion method which could neatly deal with the issues we had identified, and we set to work extending the analysis framework he had devised for to suit our needs. The details of that are, obviously, below.

Before we get into the details I’d like to make two last points. The first is that the issues listed above, that we were the first group to properly deal with, really are the difference between

getting it wrong and getting it at least closer to right. In fact, Gary Centers and co [64] wrote a lovely paper about how failing to address the stochasticity issue (the field amplitude itself is temporally stochastic) can mean that seven currently published exclusions are too stringent, some by about a factor of 10! So one really must try to address these things if they want to claim they gave it the ol' college try. The second point is that some folks can be hesitant about Bayesian analysis methods, because they feel that things like the 'prior belief' are too subjective. I completely understand this sentiment, I actually began my dark matter experimentalist journey as a staunch frequentist for this exact reason. It wasn't until I was doing this analysis that it became clear how superior the Bayesian approach is. But to ease your mind before you dig into the details, dear reader, I will point out three things. The first is that our analysis is largely independent of the priors. The prior update, introduced to you soon, is actually the key parameter in this search and this is really just the ratio of two probabilities: that of getting your measurement if dark matter really exists vs that of getting your measurement if it's just noise. We also made our matrix of prior updates freely available to anyone who wants it so that they can explore the results with whatever priors they want. Our exclusion does not even depend on the absolute *scale* of the priors, but only on their profile. But the profile really matters, which brings me to the second point. We chose a log-uniform prior profile and we believe that it is a natural choice given that there isn't really any reason for the new dark matter particle to appear in any specific mass decade. The masses of known standard model particles are distributed over many decades, and there is no particular reason to expect new physics to be any different. In addition, the body of literature surrounding this type of dark matter has only constrained it to exist within 25 (!) mass decades, which means that we have already *implicitly agreed*, as a community, on logarithmically uniform priors. The third point is that the choice of priors is isomorphic to the choice of power threshold set for detection used in a frequentist analysis. You'll often read about a 95% confidence interval for each frequency examined. All of the published frequentist analyses of this type that I've seen use a constant threshold for detection over a large frequency range which is, frankly, silly - because it is equivalent to choosing uniform priors over several mass

decades, which amounts to only actually doing your search in the highest decade. So if you're going to do that, you might as well throw out the lower decades entirely. A frequentist analysis would solve this problem by choosing a frequency-dependent threshold, but then you have done something which is isomorphic to choosing a profile for the priors, and then you are being just as "subjective" as Bayesians. I really think that it is a meaningful practical advantage, both as the scientist doing the analysis, and for the critical readers examining the results, to have clearly states priors and justifications for them - precisely because it makes your assumptions explicit. One can compare this situation to a frequentist analysis where they can be hidden, unexamined, and sometimes, indefensible.

Well I'm sure you've heard just about enough of me by now. Let's get on to the main event: a copy-paste of the paper and supplement :)

5.2 Experimental Constraint on Axionlike Particles over Seven Orders of Magnitude in Mass

This section is a reproduction of the main text of ref. [60]

A measurement of the permanent electric dipole moment of the electron (eEDM, d_e) well above the standard model prediction of $|d_e| \leq 10^{-38} e \text{ cm}$ would be evidence of new physics [65, 66, 67, 68, 69]. The search for permanent EDMs of fundamental particles was launched by Purcell and Ramsey in the 1950s [21, 70], but it was not until recently that physicists began to consider the possibility of **time-varying** EDMs inspired in large part by hypothetical axionlike particles (ALPs) [71, 72, 6, 73, 74, 75]. Coupling to ALPs can cause EDMs of paramagnetic atoms and molecules to vary, due to variations in d_e or the scalar-pseudoscalar nucleon-electron coupling C_S . In this paper, we extend our analysis of the EDM data we collected in 2016 and 2017 [2] to place a limit on the possibility that C_S oscillates in time. This allows us to put a constraint on the hypothetical ALP-gluon coupling over seven orders of magnitude in mass.

Dark matter is one of the largest unsolved mysteries in modern physics: the microscopic nature of 84% of our universe’s cosmic matter density remains inscrutable [10]. A family of candidate dark matter particles called axion-like particles are pseudoscalar fields favored for their potential role in solving problems such as the strong CP problem, baryogenesis, the cosmological constant, and small-scale cosmic structure formation [76]. ALPs can be anything from the QCD axion [77, 78, 79, 80] to ultralight axions from string theory [81], with different models favouring different mass ranges [82, 83, 84, 76, 85, 5]. In this paper, we will consider ultralight spin-0 bosonic fields whose mass m_a can range from $10^{-24} - 10^1$ eV (or $10^{-10} - 10^{15}$ Hz) [5]. ALPs in the low end of this mass range would solve some issues pertaining to small scale structure raised by astrophysical results [85, 86, 76, 87].

Ultralight dark matter particles such as these must be bosonic: the observed matter density can’t be reproduced with light fermions because their Fermi velocity would exceed the galactic escape velocity. The resulting field has high mode occupation, is well described classically as a wave, and oscillates primarily at the Compton frequency $\nu = m_a c^2/h$. Astrophysical observations tell us that a quasi-Maxwellian dark matter velocity distribution would have, near earth, a dispersion of $\delta v \sim 10^{-3}c$, which in turn leads to a coherence for the dark matter (DM) wave of about 10^6 cycles [88]. This finite linewidth ultimately limits the sensitivity of any experiment whose total observation time approaches or exceeds the coherence time.

Several distinct couplings to ALP fields can result in an oscillating CP-odd quantity such as d_e or C_S , which amounts to an oscillating molecular EDM signal in our data. The most straightforward effect is the derivative coupling of the ALP field to the pseudovector electron current operator — in other words, a direct coupling between ALP and electron. Unfortunately, in an atomic or molecular system the observable effect scales linearly with m_a , leading to a large suppression for the mass range we explore [75]. A second potential coupling is from the ALP-induced modification of the electron-nucleon Coulomb interaction [74]. Regrettably, the nature of this effect has not been elucidated in detail, but it is also expected to scale linearly with m_a

[89]. Finally, there are two distinct effects due to the coupling of ALPs to gluons. The first leads to a partially screened nuclear EDM which would be enhanced in polar molecules but scales linearly with m_a , which again is problematic in the mass range we probe [90]. The second effect will result in an oscillating C_S and this is independent of m_a [91, 92, 93]. In this paper, we use the hypothesized oscillation in C_S to constrain the ALP-gluon coupling.

To understand the analysis discussed herein, the reader should be familiar with a few key details from the original measurement. Complete details can be found in [2] and references therein. The heart of our experiment is an electron spin precession measurement of $^{180}\text{Hf}^{19}\text{F}^+$ molecular ions confined in a radio-frequency Paul trap and polarized in a rotating bias field to extract the relativistically-enhanced CP-violating energy shift $2hW_S C_S$ between the stretched and oriented $m_F \Omega$ Stark sublevels [94]. The molecule-specific structure constant $W_S = 20.4$ kHz [59], h is Planck's constant, and we assume for the remainder of this paper that d_e is zero.

In an individual run (or **shot**) of the experimental procedure, many HfF^+ ions are loaded into an ion trap, and the internal quantum population is concentrated into one of four internal sublevels by means of pulses of polarized light. The energy difference between two sublevels is interrogated with Ramsey spectroscopy [2]. The population is put into a superposition of two levels by means of a coherent mixing pulse, and then allowed to coherently evolve for a variable time t_R , before a second coherent pulse remixes the two levels. The resulting population in a sublevel is destructively read out by means of final series of laser pulses, which induce photodissociation of HfF^+ ions into Hf^+ ions with a probability proportional to the population in the interrogated sublevel. The procedure requires about one second and the ion dissociation signal has a component which oscillates with t_R as $\cos(2\pi f_R t_R)$, where f_R is proportional to the energy difference (averaged over the ~ 0.5 second coherent evolution time t_R) between the $m_F = \pm 3/2$ magnetic sublevels of a given Stark level. There are multiple contributions to f_R , including Zeeman shifts and Berry's phase shifts. Over the course of a **block** of data, comprising 1536 experimental shots, we chop the direction of the ambient magnetic field, the rotation direction of the bias field, and

the internal Stark level probed, and for each of these eight different experimental configurations we sample multiple Ramsey times, in order to extract f_R under eight different laboratory conditions. A certain linear combination of all the various f_R measurements isolates, to first order, the CP-violating contribution to the energy difference, $f^{C_S} = 2W_S C_S$, where in this case we are sensitive to the average value of C_S during the ~ 22 minutes required to collect a block of data.

To probe for a possible time-variation in C_S , our analysis is broken into two parts: ‘low frequency’ (27 nHz – 126 μ Hz), corresponding to hypothesized variations in W_S with periods shorter than the overall 14 month span of data collection but still long compared to the duration of a block, and ‘high frequency’ (126 μ Hz – 400 mHz), corresponding to periods shorter than our most sensitive period of data collection (~ 11 days) but still long compared to the duration of a shot (~ 1 second)[95, 96].

The low-frequency Least Squares Spectral Analysis (LSSA) is straightforward: we take the bottom-line results for each of our 842 blocks of data ($f^{C_{Si}}, \sigma_i, T_i$), where T_i is the time in seconds at which the i^{th} block was collected, relative to the center of our data set (at 00:17 UTC, Aug 16, 2016), and σ_i is the standard error on $f^{C_{Si}}$. Then for each hypothesized ALP frequency $\nu_i = \omega/(2\pi)$ in a set of equally spaced trial frequencies in the low-frequency range, we assume

$$f^{C_S}(T) = A \cos(\omega T) + B \sin(\omega T), \quad (5.1)$$

and do a weighted least-squares fit [96] to extract our best values $A_0(\nu)$ and $B_0(\nu)$.

For the high-frequency LSSA analysis we want to resolve hypothesized changes in f^{C_S} that occur over the time scales ≤ 22 minutes (the time required to collect a block), so we disaggregate each block into its 1536 component measurements: the number of dissociated ions N measured for each of the shots that make up a block. The expected value for N for each shot is an oscillatory function of the Ramsey time t_R . The frequency of the oscillation in t_R is modulated because a component of f_R is proportional to C_S , hypothesized to be an oscillatory function of T_i , and the sign of the modulation amplitude depends on the sign of the bias magnetic field applied for




Figure 5.2: Least-squares spectral analysis of our EDM data.

that shot, and on the Stark doublet selected for probing for that shot. The least-squares fitting routine is correspondingly more complicated [96] but at the end of the day we again extract best-fit values $A_0(\nu)$ and $B_0(\nu)$ for each hypothesized value of the ALP oscillation frequency ν in the equally spaced grid of trial frequencies across the high-frequency range.

Combining the two analyses, we obtain a measure of the best fit oscillation amplitude over seven orders of magnitude (Figure 5.2). Now we can ask: to what extent do our fit values convince us that C_5 is **actually** oscillating, given experimental noise and the look-elsewhere effect [97]? To answer this question we used the Bayesian Power Measured (BPM) analysis framework developed by Palken et al. [3].

The essence of the BPM framework is a comparison between the probability of measuring a given oscillation amplitude assuming the existence of an ALP with a specified mass and coupling and the probability of measuring the same amplitude assuming it doesn't exist. This means we need probability distributions for both cases at each frequency ν , which we can evaluate at the actual measured values A_0, B_0 . For the case where ALPs do not exist, we get the distributions by generating simulated data which by construction has no coherent C_5 oscillation in it. For the low

frequency data we generate a single point of ‘noise’ at each acquisition time in the real dataset. Each point of noise is drawn from a Gaussian distribution with mean of zero and a variance which matches that of the real dataset. For the high frequency analysis we randomly shuffle the timestamps of the individual Ramsey measurements in such a manner as to effectively erase any coherent oscillation which may be present in the C_5 channel while preserving both the basic structure of the data (including the general coherent oscillation present in each Ramsey dataset) and the technical noise (see [96] for more details).

Once we have this simulated data we perform LSSA on it to get best fit values for A, B at each frequency ν of interest. We repeat this process 1000 times. The resultant distribution of A_0, B_0 at each frequency ν represents what we would expect to find in a universe where ALPs do not exist. Each distribution is bivariate normal with some rotation angle. We rotate each distribution into the primary axes A', B' where the variance of the distribution is maximized along A' and minimized along B' , to define the no-ALP distribution:

$$\mathcal{N} = \frac{1}{2\pi\sigma_{A'}\sigma_{B'}} e^{-\frac{1}{2}\left(\frac{A'^2}{\sigma_{A'}^2} + \frac{B'^2}{\sigma_{B'}^2}\right)}. \quad (5.2)$$

We still need to determine the probability distributions in the case that ALPs do exist. Note that A' and B' are still random variables in this case due to the stochastic nature of the ALP field: different spatio-temporal modes of the field interfere with each other (the ALP field at any point in space is a sum over many contributions with random phases), so the field amplitude is stochastic over timescales longer than the coherence time [98, 99]. Except at the very high end of our analysis range, the ALP coherence time is much larger than the duration of our measurement, so we resolve only a single mode of the ALP field. In this limit the local ALP field may be written in the form

$$a(T) = \alpha\sqrt{\rho_{\text{DM}}}/m_a \cos(\omega T + \phi), \quad (5.3)$$

where we take the mean local dark matter density ρ_{DM} to be $0.4 \text{ GeV}/\text{cm}^3$ [4], ω includes a small contribution from the ALP kinetic energy, $\phi \in [0, 2\pi)$ is a uniform-distributed random variable,

and $\alpha \geq 0$ is a Rayleigh-distributed random variable: $P(\alpha) = \alpha e^{-\alpha^2/2}$. The modifications to our analysis due to the decoherence of the ALP field are discussed in our Supplemental Material [96], otherwise we assume α and ϕ are constant over the entire data collection time. The oscillating C_S term in the Ramsey fringe frequency is then written as

$$f^{C_S}(T) = 2W_S\eta C_G / f_a a(T), \quad (5.4)$$

where C_G / f_a is the ALP-gluon coupling (see [96] for alternative parameterizations) and $\eta \approx 0.03$ is a coefficient relating C_S induced in a generic heavy nucleus to the QCD theta angle [91, 100]. The quadrature amplitudes of the ALP signal are then

$$\begin{aligned} A' &= 2W_S\eta C_G / f_a \alpha \sqrt{\rho_{\text{DM}}} / m_a \cos(\phi) \\ B' &= -2W_S\eta C_G / f_a \alpha \sqrt{\rho_{\text{DM}}} / m_a \sin(\phi). \end{aligned} \quad (5.5)$$

As mentioned earlier, A' and B' must still be treated as random variables even in the absence of noise; the fact that α is Rayleigh-distributed implies that A' and B' are each Gaussian distributed with mean zero and variance

$$\sigma_X^2 = (2W_S\eta C_G / f_a \sqrt{\rho_{\text{DM}}} / m_a)^2. \quad (5.6)$$

The equal quadrature variances reflect our ignorance of the local phase of the ALP field oscillations. The fact that A' and B' can simultaneously be small reflects the possibility that the earth may have been near a null in ALP density during our data collection time. We can now construct the expected 'ALP distribution' by adding the variance σ_X^2 to our no-ALP variances (the no-ALP variances characterize the experimental noise). This defines a new bivariate normal distribution

$$\mathcal{X} = \frac{1}{2\pi\sigma_{AX}\sigma_{BX}} e^{-\frac{1}{2}\left(\frac{A'^2}{\sigma_{AX}^2} + \frac{B'^2}{\sigma_{BX}^2}\right)}, \quad (5.7)$$

which we call our 'ALP distribution'. Here, $\sigma_{AX}^2 = \sigma_{A'}^2 + \zeta\sigma_X^2$ and $\sigma_{BX}^2 = \sigma_{B'}^2 + \zeta\sigma_X^2$. We encapsulate sources of attenuation, including decoherence of the ALP field due to the finite linewidth, in the frequency-dependent factor ζ which we discuss in detail in the Supplemental Material [96].

Now we can compare probabilities for the ALP vs no-ALP case. For each hypothesized coupling strength $C_G/(f_a m_a)$ and frequency ν of interest we take the ratio of the probability distributions \mathcal{N} and \mathcal{X} evaluated at our actual measured amplitudes A'_0 and B'_0 :

$$U_i(\nu_i, C_G/(f_a m_a)) = \frac{\mathcal{X}(A'_0(\nu_i), B'_0(\nu_i), \nu_i, C_G/(f_a m_a))}{\mathcal{N}(A'_0(\nu_i), B'_0(\nu_i), \nu_i)}, \quad (5.8)$$

which we call the **prior update**: the number U_i is a multiplicative factor which, when multiplied with our logarithmically-uniform priors, updates our prior belief to our post-experimental (or posterior) belief (see Figure 5.3). Note the prior update is equal to the **Bayes factor** in the limit of very low priors, which is the the limit we are operating in. A 95% exclusion corresponds to the prior update dropping to 0.05 [3].

So far, our analysis has not taken into account the look-elsewhere effect even though we are looking at more than 10^6 frequencies. The colormap of U in Figure 5.3 simply shows how our belief in the existence of ALPs has changed as a function of frequency and coupling. To account for the look-elsewhere effect one must move from a notion of how **locally** unlikely an event is to how **globally** unlikely it is. For example, what might seem unlikely in a single frequency bin (say, you only expect it to happen 1/1000 times) becomes rather unsurprising when you look in 1000 bins - in this case you should expect to see that event approximately once. Our local measure, the prior update, compares the probability distributions in each bin for the ALP vs the no-ALP case. To move to a global measure we generate the **aggregate** prior update:

$$\mathcal{U}(C_G/(f_a m_a)) = \frac{\text{aggregate posterior}}{\text{aggregate prior}} = \frac{\sum_i U_i \epsilon / \nu_i}{\sum_i \epsilon / \nu_i}, \quad (5.9)$$

where ϵ / ν_i is our prior belief (for an ALP of frequency ν_i), using logarithmically-uniform priors (ϵ being constant). The aggregate prior update correctly accounts for our logarithmic priors, which is necessary in broad searches like ours where the number of hypotheses tested varies strongly as a function of frequency but we expect the ALP is equally likely to be found in any frequency decade. Using equation 5.9 we can sub-aggregate over any subset of the analysis range and at any frequency we like, and any interested reader can do the same - the full un-aggregated




Figure 5.3: Bayesian power measured analysis of the EDM data. The grayscale matrix corresponds to a sub-aggregation of the prior updates U over 100 logarithmically-spaced frequency bins. Darker shading on the logarithmic color scale corresponds to increased prior update, which can be interpreted as an increased belief in an ALP of given frequency ν at ALP-gluon coupling $C_G/(f_a m_a)$. We use a blue line to indicate the 95% exclusion (which corresponds to the sub-aggregated update dropping to 0.05 [3]) for sub-aggregation over each decade. The corresponding greyscale matrix for this sub-aggregation is not shown. Finally, a dotted line corresponds to the aggregated exclusion over the full analysis range. The full un-aggregated matrix of U values is available upon request.

matrix of U values is available upon request. The aggregate prior update taken over the **entire** set of frequencies analyzed represents the fractional change in our belief that an ALP of a given coupling strength exists anywhere in the full analysis range, which appropriately accounts for the look-elsewhere effect. The aggregate prior update as a function of coupling strength is plotted in the Supplemental Material, where we also discuss our selection of the priors [96].

Our results yield no strong indication of a dark matter signal over the range $10^{-22} - 10^{-15}$ eV, so we use the above procedure to exclude $C_G/f_a m_a > 1.72 \times 10^{16} \text{ GeV}^{-2}$ at 95% confidence. If we assume the density of the ALP field is held constant at its mean value, we can convert our exclusion back to frequency modulation amplitude in Hz via eq (5.3) and (5.4) which gives us an exclusion of 37 mHz. This is considerably less impressive-seeming than the limit of 1.2 mHz which we extracted from our original DC analysis [2]. This relaxation of the exclusion is an unavoidable consequence of correctly accounting for the stochastic nature of the ALP field and the look-elsewhere effect associated with setting a limit across seven orders of magnitude in frequency.

Our constraint on the ALP-gluon coupling C_G/f_a is plotted in Figure 5.4 along with existing direct and indirect limits. Our results are the first laboratory constraints on the ALP-gluon coupling in the mass range $10^{-17} - 10^{-15}$ eV.

The analysis techniques applied herein are general enough to apply to most time-series datastreams, and in particular will be applied to our next-generation EDM measurement which we expect to be about an order of magnitude more sensitive. More generally, the field of precision measurement is in a remarkable period of progress reminiscent of Moore's Law, and we expect that analyses of this type will become more and more common. At the threshold of this era, we believe that it is critical to ask: how we can get the analysis of hypothesis exclusion right?

This is especially pertinent given that the existing frameworks and analyses already published have so far failed to account for important effects, such as the stochastic nature of the ALP




Figure 5.4: Astrophysical and laboratory limits on the ALP-gluon interaction. All limits assume that ALPs saturate the local dark matter content and we have taken the mean local dark matter density to be $0.4 \text{ GeV}/\text{cm}^3$ [4]. All constraints correspond to a 95% confidence level. The purple region embodies the constraint that the virial deBroglie wavelength is smaller than the size of a dwarf galaxy [5]. The blue region is the constraint derived from looking at neutron EDM oscillation [6]. The green region represents constraints from big bang nucleosynthesis [7], and the yellow region represents constraints from supernova energy loss bounds [8, 9]. The pink region is the constraint described in this work.

field or the look-elsewhere effect. At least seven other dark matter exclusion publications have overestimated their exclusion by factors ranging from 3 to 10 in their failure to account for the stochasticity of the field [98]. To the best of our knowledge, the Bayesian analysis framework we have chosen here is optimal for this kind of exclusion for several reasons: first, it easily incorporates the stochastic nature of the dark matter field, as well as the fact that the scale of the stochastic fluctuations changes as the measurement time approaches the coherence time of the dark matter field. In addition, this framework seamlessly accounts for the differential sensitivity of the two quadratures of our measurement to the dark matter field, and can be easily scaled up to even higher dimensional parameter spaces. Finally, this framework correctly accounts for the look-elsewhere effect with respect to exclusion [63]. Note that [6] in Figure 5.4 fails to account for the stochastic fluctuations in the ALP field, which could mean that their exclusion is overestimated by nearly an order of magnitude [98]. More generally, we appreciate how this framework preserves all the information in the signal compared to a threshold-based framework, which reduces the information in the signal to only ‘above’ or ‘below’ threshold.

5.3 Experimental Constraint on Axionlike Particles over Seven Orders of Magnitude in Mass - Supplemental Material

this section is a reproduction of the supplemental material of ref. [60]

5.3.1 Frequency Ranges

Our analysis was broken into two regimes: ‘low frequency’ (27 nHz – 126 μ Hz) and ‘high frequency’ (126 μ Hz – 400 mHz). The fundamental limits of the low frequency analysis were set by the inverse timespan of the entire dataset (27 nHz, see Figure 5.5a) and the ‘Nyquist’ limit for the ‘block’ collection (each block took 22 minutes to collect, so this is 378 μ Hz). The fundamental lower limit of the high frequency analysis was set by the inverse timespan of the most sensitive data, which was taken over an 11-day period in 2017 (1.12 μ Hz, see Figure 5.5b). The upper limit is slightly below the ‘Nyquist’ limit – each point took about 1 second to collect and we chose

400 mHz. Our data is not evenly spaced in time, so we are using the term ‘Nyquist limit’ in a loose sense: with unevenly spaced data, the Nyquist frequency is often **higher** than one might expect [101]. The frequency ranges above have considerable overlap, and we used this overlap




Figure 5.5: (Top) Individual C_5 measurements plotted as a function of their acquisition time since August 16 2016 (the temporal ‘center’ of our dataset). (Bottom) Our most sensitive C_5 measurements were collected over 11 days in 2017. For the low frequency analysis we used the entire dataset (top), and for the high frequency analysis we used only the most sensitive 11 days (bottom).

to validate our high frequency analysis. We analyzed the overlapping portion of the dataset using both the ‘low frequency’ and ‘high frequency’ methods and found virtually the same LSSA result in both cases. Taken together, the entire dataset was enormous, so we chose to truncate both datasets to eliminate overlap and reduce the size of the dataset, which made computing slightly less resource-intensive.

In each regime, the frequencies we interrogated were evenly spaced between maxima and minima; the spacing was chosen as $1/6$ of our spectral resolution – intentionally oversampling so that we wouldn’t miss any potential peaks. Our spectral resolution is given by the inverse of the timespan to collect each dataset, in the low frequency case this is ~ 450 days and in the high

frequency case this is ~ 11 days.

5.3.2 Generating Least-Squares Spectra

Our experimental signal is a count of Hf^+ ions generated by state-selective dissociation of the HfF^+ ions. From this signal we generate an asymmetry $\mathcal{A}'_i = \pm \frac{N_i - \langle \text{Hf}^+ \rangle}{\langle \text{Hf}^+ \rangle}$. Here, N_i is the Hf^+ count in a single **shot** (run) of the experiment and $\langle \text{Hf}^+ \rangle$ is the mean Hf^+ count for the **data point** (set of shots, averaged) to which the shot belongs. The sign of \mathcal{A}'_i is given by the spin preparation/readout conditions [102, 103]. A set of asymmetries, taken by sandwiching a variable free-evolution time t_R between 2 coherent m_F mixing pulses, generates a Ramsey fringe. The frequency of the Ramsey fringe is directly proportional to the energy difference between the $m_F = \pm 3/2$ states in a given Stark doublet.

The Ramsey fringe frequency has multiple contributions:

$$f_R = 3g_F\mu_B B_{\text{rot}}/h + 3\alpha' f_{\text{rot}} \tilde{R} \tilde{B} - 2W_S C_S \tilde{D} \tilde{B} + \dots, \quad (5.10)$$

where g_F is the $F = 3/2$ state g-factor, μ_B is the Bohr magneton, B_{rot} is the magnitude of the rotating bias magnetic field, f_{rot} is the frequency of rotation, \tilde{B} is the sign of the magnetic field, \tilde{D} is the populated Stark doublet, \tilde{R} is the sense of the electric bias field rotation, and α' is a parameter describing Berry's phase [45]. For any given measurement, the experiment will be in a given **switch state**, characterized by the values of $\{\tilde{B}, \tilde{D}, \tilde{R}\} = \{\pm 1, \pm 1 \pm 1\}$. To isolate the CP-violating energy $2hW_S C_S$, we repeat our spin precession measurement in each of the 2^3 unique states at varying free-evolution times to form eight different Ramsey fringes and respective **switch frequency** measurements $f^{\{\tilde{B}\tilde{D}\tilde{R}\}}$. From each set of eight fringes (called a **block**) we form linear combinations of our measured switch frequencies to form **parity frequencies**: $\mathbf{f}^{\text{parity}} = M^{-1} \mathbf{f}^{\{\tilde{B}\tilde{D}\tilde{R}\}}$, where M^{-1} is a transformation matrix and \mathbf{f} indicates the set of parity or switch frequencies. One of these linear combinations isolates, to first order, the CP-violating C_S term $f^{C_S} = 2W_S C_S$.

As mentioned above, our analysis is broken into two parts, low and high frequency. In

general terms, the low frequency analysis consists of taking the f^{C_s} measurement from each block as a function of block acquisition time (Figure 1) and reanalyzing the data to check for oscillations in our signal. In contrast, for the the high frequency analysis we break apart each block and fringe to extract a (potentially) time-varying C_s signal from each shot of the experiment, allowing us to extend the bandwidth of our search over an additional three orders of magnitude.

In detail, the low frequency analysis consists of least-squares spectral analysis (LSSA) on our set of f^{C_s} measurements as a function of block acquisition time. We minimize

$$\chi^2 = \sum_{i=1}^N \frac{1}{\sigma_i^2} [A \cos(\omega T_i) + B \sin(\omega T_i) - f^{C_s}(T_i)]^2 \quad (5.11)$$

over A and B for every potential ALP oscillation frequency $\nu = \omega/(2\pi)$ of interest, where σ_i is the standard error on a given f^{C_s} measurement and T_i is the block acquisition time in seconds with respect to the temporal center of the dataset. This equation is analytically solvable, so we directly compute the minima of the resulting quadratic equation to obtain the “best fit” values of A and B : $A_0(\nu)$ and $B_0(\nu)$.

For the high frequency analysis, in contrast, we have to perform LSSA directly on the set of **asymmetry** measurements as a function of **shot** acquisition time T . It may be helpful at this point to emphasize that we have two different time streams: the **free-evolution time** t_R (which ranges between 0 – 0.7 seconds) which is an entirely distinct parameter from the **shot/data acquisition time** T (which ranges between $\pm 2 \cdot 10^7$ seconds). The asymmetry measurements already oscillate as a function of free-evolution time (t_R), and for a static (DC) EDM measurement it is assumed this oscillation frequency is constant (for a given switch state, up to experimental noise). In the presence of an oscillating $f^{C_s}(T)$, however, we expect the already-present oscillation in t_R to be frequency modulated in T , and we need to determine the form of this frequency modulation. We begin with the fits generated from our DC analysis in [2]: the asymmetry points $\mathcal{A}'(t_R)$ from each fringe in a given switch state were fit to $\mathcal{A}^{\text{fit}}(t_R) = -C e^{-\gamma t_R} \cos(2\pi f_R t_R + \phi) + O$, where C is the fringe contrast, f_R is the fringe frequency, ϕ is the initial phase, O is the offset, and γ

is the coherence decay rate. As mentioned earlier, the frequency f_R measured in a given fringe does not provide a direct measurement of C_S ; it contains multiple contributions. The frequency due to non- C_S terms is $f_0 = f_R - M' f^{C_S}$, where M' are the matrix elements used to transform between switch and parity bases and f^{C_S} is the frequency component due to C_S . Since we are only interested in the potentially time-varying C_S component (and we assume it has a mean of zero [2]) we write

$$\begin{aligned} f_R(T) &= f_0 + M' f^{C_S}(T), \\ f^{C_S}(T) &= A \cos(\omega T) + B \sin(\omega T). \end{aligned} \tag{5.12}$$

If we include this frequency modulation (which we assume is very small), we can expand our fit asymmetry \mathcal{A}^{fit} to second order in $f^{C_S}(T)$ to obtain

$$\begin{aligned} \mathcal{A}^{\text{osc}}(T) &= -C e^{-\gamma t_R} \cos(2\pi(f_0 + M' f^{C_S}(T))t_R + \phi) + O \\ &\sim -C e^{-\gamma t_R} \cos(2\pi f_0 t_R + \phi) + O \\ &\quad + C e^{-\gamma t_R} (2\pi t_R M' f^{C_S}(T)) \sin(2\pi f_0 t_R + \phi) \\ &\quad + C e^{-\gamma t_R} 2(\pi t_R M' f^{C_S}(T))^2 \cos(2\pi f_0 t_R + \phi), \end{aligned} \tag{5.13}$$

which is our desired expression for the $f^{C_S}(T)$ -induced frequency modulation. With this, we perform LSSA on the difference between $\mathcal{A}^{\text{osc}}(T)$ and $\mathcal{A}'(T)$,

$$\chi^2 = \sum_{i=1}^N \frac{1}{\sigma_i^2} [\mathcal{A}_i^{\text{osc}} - \mathcal{A}'_i]^2, \tag{5.14}$$

where the subscript indicates the i^{th} acquisition time point in our dataset.

5.3.3 Timestamp Shuffling

As described in the main text, our statistical analysis requires simulated data which, by construction, contains no coherent oscillations in the f^{C_S} channel. For the low frequency analysis this is a rather simple procedure which relies on our knowledge that each block-based f^{C_S} measurement is drawn from a Gaussian distribution with mean zero and well-characterized variance [2]. For the high frequency analysis our simulation technique warrants further discussion.

The challenge is to simulate shot-by-shot data which preserves the basic structure of the fringes and blocks as well as the inherent noise of the measurements. The fringes themselves have a coherent oscillation in them, and each block contained eight fringes – each of which may oscillate at a slightly different frequency depending on the switch state. The fringes also had a limited coherence time, which varied throughout the dataset as we changed experimental parameters. While generating simulated data that faithfully preserves these features is in principle possible, the most difficult part for our purposes was preserving the noise. There were a number of technical noise sources which were not easily characterized, i.e. they didn't conform to some known distribution, and were transient, i.e. they were present or absent to different degrees through the dataset.

We found that the most effective way to preserve the qualitative features of the noise and the specific details of the fringes and blocks while erasing any real-time coherent oscillation in the f^{Cs} channel was to randomly shuffle the timestamps of the data by block. As described in the previous section, a block of data formed one single f^{Cs} measurement and was composed of eight Ramsey fringes, which together comprise 96 data points or 1536 shots. When we collected a block of data we would **not** collect each Ramsey fringe one at a time, slowly increasing the free-evolution time until completion then changing the switch state to collect the next fringe. Instead, we would increase the Ramsey time while scrambling the various switch states, effectively collecting points from each fringe in random order. In other words, we randomized the order in which we collected the 96 data points corresponding to each block, such that there is no fixed relationship between either the switch state or Ramsey time of adjacent data points. When we shuffle the timestamps by block, that means we take all timestamps of the data points belonging to a given block and swap them with timestamps for data points belonging to some other block. One can imagine this as swapping whole blocks around in time. Since a block corresponds to a f^{Cs} measurement, this will eliminate any coherent oscillation in the f^{Cs} channel which may have been present, so that it is broken instead into many lower-amplitude features that neither

appear as candidates for discovery, nor, when combined with many reshuffles, as significant perturbation to the noise distributions we use for determining exclusions.

We checked to make sure the reshuffled data quantitatively reproduced the coarse-grained frequency structure of the noise in the actual data. We were also able to compare the spectra generated by shuffling the f^{C_s} data to spectra of the non- f^{C_s} parity frequencies, which by construction should be free of coherent oscillation. Each such non- f^{C_s} channel yields a single representation of a noise spectrum which we should expect to be similar to that of f^{C_s} in the absence of an ALP signal – indeed, these shuffled spectra qualitatively reproduced our non- f^{C_s} noise spectra quite well. As an additional check, we can take an individual data set from our collection of reshuffled data sets and treat it to the full analysis as if it were our actual data set, including creating a subaggregated exclusion heat-map plot as in Figure 3 of our main text, and an aggregate prior update ‘discovery plot’ as in Figure 4 below. The fake exclusion heat maps reproduce the qualitative features of our actual exclusion heat maps. For the twenty fake ‘discovery plots’ we generated, the largest peak aggregate update value was < 20 , and all but the three largest were < 10 ; i.e., none were indicative of discovery. The peak aggregate update value occasionally being slightly larger than in our real data set is to be expected as a consequence of a sort of higher-order look-elsewhere effect.

To confirm that our simulation method really did sufficiently attenuate coherent oscillations while preserving the spectral structure of the noise in our dataset, we performed least-squares spectral analysis (LSSA) on simpler versions of simulated datasets which by construction contained a coherent oscillation as well as noise. The simulated data was sampled at the same times as our real dataset, and we performed simulations at frequencies throughout the bandwidth analyzed. In each case we found that the detected amplitude of the oscillation was significantly attenuated while the noise at other frequencies was not affected. Conversely, when we shuffled by data point or shot we found that while any coherent oscillation would be erased, the noise at other frequencies would also be attenuated.

It is clear why shuffling the timestamps by block should erase any oscillations on the timescale of the block itself or longer, but it may not be clear why the technique works for oscillations with a higher frequency. The reason this works is due to the manner in which the data was collected – the random scrambling of when we collected the Ramsey times and switch states is unique to every block, so the shuffle introduces phase jumps into the data on the timescale of single data points – in other words, it effectively shuffles the data on the **data point** (16 second) timescale.

One might ask: why not just shuffle the timestamps by data point? The reason is that shuffling by data point, while very effective at erasing any coherent oscillation present, also appears to drastically change the spectral structure of the noise in our dataset. One might then ask why we don't simply simulate the data directly, and again we find that assuming a specific noise model (such as shot noise in the HfF^+ count) is insufficient for simulating the real noise we had present. It should go without saying that arbitrarily increasing the noise in the simulated data led to the risk of hiding any real oscillation which would otherwise be detected, has no real scientific basis, and should be avoided.

5.3.4 Estimating the no-ALP distribution

Our analysis hinges on having a good estimate of the no-ALP distribution, which itself relies on knowing the variances $\sigma_{A'}^2$ and $\sigma_{B'}^2$. Because we are searching over such a large frequency range, and examining so many different frequencies, there are many upward fluctuations in the oscillation amplitudes obtained from LSSA, in both the real and simulated datasets. Small fluctuations in the estimated variances can result in large changes in the apparent significance of these amplitudes. If our estimates of $\sigma_{A'}^2$ and $\sigma_{B'}^2$ are poor, our exclusion will reflect that.

One might think that the best way to ensure we have good estimates of the variances would be to do many more simulations. Unfortunately, especially for the high frequency data, the simulations were very computationally expensive so running more simulations was not an

option. This high computational cost is not a feature specific to our choice of Bayesian analysis framework, but rather a generic consequence of the large number of independent frequencies at which we must construct the no-ALP distribution, taken together with the size of the full set of asymmetry measurements on which the high frequency LSSA is performed. Because we only ran 1000 simulations our error in estimating the variances is dominated by statistical noise. To account for this, we chose to adjust the raw variances obtained from the simulated data to better match what we know to be true.

In the low frequency range, we made no adjustments to the variances. We know that there will be some phases and frequencies which are favoured due to the temporal structure of the data (see Figure 5.5, top), so we wanted to preserve any structure in the variances which we found via simulation (including significant differences between the two quadratures A', B'), even if it was a bit more noisy than desired due to the limited number of simulations.

In contrast, for the high frequency range we have no reason to believe there would be any meaningful difference between the two quadratures thanks to the near-continuous temporal spacing of the data we analyze (see Figure 5.5, bottom). In addition, we know that statistical fluctuations are independent, bin-to-bin. So to attenuate the effect of statistical fluctuations we smoothed the data using a 36-bin running average (this is 6 times our spectral resolution) and then averaged the the two quadratures A', B' .

5.3.5 Sources of Attenuation

To reduce the risk of artificially over-constraining the axion-like particle (ALP) coupling it was important to account for all the known sources of attenuation present in our experimental signal. To date we are aware of three possible sources of attenuation:

- (1) When the total observation time T is smaller than the coherence time τ_c , i.e. $T < \tau_c$, the ALP field is built from a sum over many phases of a freely oscillating bosonic field

in three dimensions, resulting in a random distribution of the field amplitude. We must account for the probability that the local ALP field amplitude is significantly smaller than its RMS value throughout the galaxy. For example it is possible that when we took the data we could have been, by sheer bad luck, be sitting near a null of the field [98].

- (2) When the total observation time T is comparable to the coherence time τ_c , i.e. $T \sim \tau_c$, decoherence due to the finite linewidth of the ALP field will decrease the measured signal amplitude [99].
- (3) When the observation time for a given **sample** (i.e. a block in the low frequency analysis or a shot in the high frequency analysis) is not small compared to the period of the ALP-induced oscillation, we are effectively averaging over some portion of the wave, leading to attenuation of the signal.

5.3.5.1 Effect 1: Random distribution of the ALP field amplitude

This effect is dealt with automatically in the Bayesian Power Measured analysis framework. When we generate the ALP distribution \mathcal{X} from the no-ALP distribution \mathcal{N} we increase the variance of both quadratures without increasing the mean. See the main text between equations (3)-(7) for more details.

5.3.5.2 Effect 2: Decoherence

To quantify the effect of decoherence we compared a realistic (finite linewidth) ALP field to an infinitely narrow linewidth field at several possible mass values. We built the realistic field from its constituent parts as described in Ref. [99], implementing eqs (3) – (10) numerically. The result is an expression for the local ALP field, which we can ‘fit’ using LSSA as described in the main text. To generate the ‘infinitely narrow’ ALP distribution, we follow the same steps but instead sample the velocity distribution only once before building up the full field. We ‘fit’ this field with LSSA as well, then repeat the process hundreds of times to get distributions for the

quadrature amplitudes A_0, B_0 in both cases. We call the amplitudes A, B and $A_{\text{narrow}}, B_{\text{narrow}}$. We rotate these distributions into the primary axes where the variance is maximized along one axis, then define the attenuation due to decoherence in each quadrature $\zeta_{d,i}$ as the ratio of the variances: $\zeta_{d,A} = \sigma_A^2 / \sigma_{A,\text{narrow}}^2$, $\zeta_{d,B} = \sigma_B^2 / \sigma_{B,\text{narrow}}^2$. In practice we find $\zeta_{d,A} \sim \zeta_{d,B}$, so we define $\zeta_d = (\zeta_{d,A} + \zeta_{d,B})/2$. Unsurprisingly we see no decoherence effect in our low frequency band, but we do at the extreme high end of the high frequency band. We fit the results using a second-order polynomial to obtain a functional form for the attenuation due to decoherence over the entire analysis band: $\zeta_d(\nu)$ (Figure 5.6).




Figure 5.6: Attenuation factor due to ALP field decoherence, calculated over the high frequency analysis range. Grey-filled circles are the calculated values, and the blue curve is the fit.

5.3.5.3 Effect 3: Finite sample time

At the high end of each analysis regime, we are analyzing frequencies fast enough that there is a non-negligible change in the signal during the sample time. This results in an ‘averaging’ of the oscillation over the sample time and an effective attenuation in perceived signal amplitude. For instance, at the upper end of our high frequency range we analyze 0.4 Hz but each sample takes about one second to acquire. We have a similar issue at the high frequency

end of the low frequency analysis. To quantify this attenuation, we generated artificial sinusoidal data sampled more finely than the timestamps in our real dataset, then averaged the signal over the actual intervals of data collection. We did this for a number of frequencies spaced through our analysis band, then averaged these results over phases spanning $[0, 2\pi)$. We then perform LSSA on both the original simulated data and the averaged simulated data and define the attenuation due to finite sample time $\zeta_{s,i}$ in each quadrature as $\zeta_{s,A} = \sigma_A^2 / \sigma_{A'}^2$, $\zeta_{s,B} = \sigma_B^2 / \sigma_{B'}^2$, where A, B are from the averaged signal and A', B' are from the original signal. We average the two quadrature attenuation factors to determine the attenuation due to finite sample time ζ_s . We did this for both the high frequency analysis and the low frequency analysis, and used the fits to the data to generate a piecewise function describing the attenuation due to this effect as a function of frequency (Figure 5.7).

Finally, our overall attenuation factor over the entire analysis band is given by the product of the attenuation factors for the two effects: $\zeta(\nu) = \zeta_d(\nu)\zeta_s(\nu)$. This is the quantity introduced below Eq. (7) in the main text.

5.3.6 Parameterizing the ALP-gluon coupling

Several different conventions have been used in the literature to parameterize the ALP-gluon coupling we constrain in this letter. Here, we use the parameterization of Ref. [6]. The most fundamental parameter in any ALP model is the energy scale f_a at which a new $U(1)$ symmetry is spontaneously broken. The ALP field then arises as the pseudo-Goldstone boson of this broken symmetry; its mass and all its couplings will generically scale as f_a^{-1} , with different dimensionless coefficients C_G, C_N, C_γ, C_e for couplings to gluons, nucleons, photons, and electrons, respectively. For the canonical QCD axion, the gluon coupling coefficient is fixed at $C_G = 1$ by the requirement that the axion solve the Strong CP problem, but for more general ALPs C_G can be regarded as a free parameter, which we would typically expect to be $O(1)$. Because of this,




Figure 5.7: Attenuation due to finite sample time over the entire analysis range. The blue curve shows the piecewise function generated by the fit to the grey circles.

many papers (such as Ref. [9]) express bounds on the gluon coupling C_G/f_a as bounds on f_a . Other papers (e.g., Refs. [7, 8]) parameterize the ALP-gluon coupling as the coefficient g_d of the oscillating nucleon EDM that the ALP field would generate. Because an EDM has dimensions of $ecm \sim eV^{-1}$ in natural units and the ALP field itself has dimensions of eV , this implies g_d has dimensions of eV^{-2} , and is related to our parameterization through Eq. (7) of Ref. [8], with $f_a \rightarrow f_a/C_G$.

5.3.7 Aggregate Prior Update

As described in the main text, the aggregate prior update (Figure 5.8) describes the fractional change in our belief that an ALP of a given coupling strength exists **anywhere** in the full analysis range, which accounts for the look-elsewhere effect. Computing the aggregate prior update requires making a choice about the frequency-dependence of the priors (though not their absolute scale). While this may seem to introduce an element of subjectivity into the analysis, it in fact

merely makes explicit a choice of assumptions which is also present, though implicit, in other analysis frameworks [3]. In equation (9) of the main text and Fig 4 below, we adopt logarithmically uniform priors, which are simple, well-suited to broadband exclusion, and well-motivated by how particle masses seem to distribute themselves in nature.

If our analysis were composed of evenly-spaced frequencies between 27 nHz — 400 mHz, computing the aggregate update would have been a direct implementation of equation (9) of the main text, where the logarithmically uniform priors account for the fact that there are far fewer frequencies in the lowest bin vs the highest bin. But since our analysis is split into two regimes, so-called ‘low frequency’ and ‘high frequency’, each with different frequency spacing, we needed to ensure our computation of the aggregate prior update accounted for the fact that the number of frequency bins in each decade wasn’t a smooth function of the frequency (or, in other words, that it **is** a smooth function of the frequency, **except for the discontinuity** at the point where the two analysis regimes are stitched together). To account for this we computed the aggregate prior update in each regime separately, then took a weighted average of the two results, where the weights were the logarithm of the number of frequency bins in each analysis regime.




Figure 5.8: Aggregate prior update \mathcal{U} (equivalently, exclusion) as a function of ALP-gluon coupling $C_G/(f_a m_a)$. The horizontal dotted line indicates 95% exclusion (where the aggregate prior update drops to 0.05), while the vertical dotted line indicates the lower limit of $C_G/(f_a m_a)$ which is excluded at 95% confidence over the entire analysis band. Had our data presented us with strong affirmative evidence for the existence of an ALP in our frequency range, the curve of \mathcal{U} vs $C_G/(f_a m_a)$ might have peaked not at $\mathcal{U} = 2.5$ but for instance at $\mathcal{U} = 1,000$ or higher.

Chapter 6

Conclusion

Conclusion chapters are often corny, pedantic, and/or unnecessary. “Tell me what you told me” is the general approach to chapters with this title. Why bother with that? Even I am tired of listening to myself talk at this point, and you, dear reader, have already slogged through 5 chapters with me. But I made the executive decision (it is my thesis after all), to include this chapter for one specific reason: this thesis marks an end to the original conception of the JILA trapped ion eEDM experiment. Moving forwards, JILA’s eEDM experiment will take on a totally new, and rather revolutionary¹, form.

Let’s try to re-adjust our perspective for a moment: this thesis was about performing a precision measurement of the electron’s electric dipole moment using trapped molecular ions. The measurement detailed in this thesis was the second, and last, of its kind: no other group in the world is using trapped ions to perform any measurement that requires electrically polarizing the species - the idea of a rotating electric field is still so new no-one else has picked it up yet. Since 2014 I have been working on a physics experiment which, in form at least, had no peers. Honestly, I found that delightful.

With such a unique experimental setup it almost seems a shame to put the whole thing out

¹In my opinion.

to pasture², but we are left with little choice (plus, and it's hard to understate how important this is to the whole operation, Eric was obviously already kind of bored with it). Why is our hand forced? Well, let's imagine what we need to do to squeeze another result out: if we aim for a 10x improvement in sensitivity with each subsequent measurement (which is admittedly ambitious but also in keeping with our current pace) then we would need to either read out 100x more ions with each shot or find a way to get our fringe contrast up to 1, then read out 25x more ions each shot, or use a new molecule to get higher $|E_{\text{eff}}|$, then get more contrast, then, etc etc. Choose any way you like to manipulate our old friend, equation 2.1 (reproduced below as Eq. 6.1 for your convenience), to get a factor of 10 out, and then think about *what's actually possible*.

$$\delta d_e \sim \frac{1}{|E_{\text{eff}}| \tau \sqrt{N}} \quad (6.1)$$

Let's consider τ first: we could get a straight factor of 10 in sensitivity if we can make τ 10x larger. Imagine we have a molecule which has the science state as a ground state, so that we are not limited by excited state lifetimes. Then what happens? Our experimental rep rate goes *way down*. We would prepare our ions, start the Ramsey sequence, then wait for something like 15 seconds before reading out the result and starting all over again. We would still need lots of data to get the statistical sensitivity we need, and now what previously took about 4 weeks of dedicated data-taking time now takes a year! Not to mention the $1/f$ noise issues that would inevitably crop up with such a low rep rate. So in general, if we are looking to extend τ we are looking immediately at having a rep rate, and a pink noise, problem. Not acceptable.

What about \sqrt{N} ? Could we leave τ alone and just read out 100x more ions per shot? Well I don't want to declare anything like that *impossible*³, but I can confidently say (quite sarcastically): good luck. First of all, we can't really trap more ions: space charge effects will always put a

²Presumably, some kind of UHV safe pasture with no shortage of copper gaskets and pure quantum states

³I am far too cowardly for that

fairly hard limit on the total number of ions one can trap in one place. But! We were routinely trapping something on the order of 20,000 ions while we only read out ~ 500 of them. This massive discrepancy is due to state preparation issues: each step in our state prep was much less efficient than 1 and so the ultimately *useful* fraction of ions we could read out is the result of a bunch of values < 1 being multiplied together. Ok fine, then why can't we just 'fix' each step in the state preparation to massively improve our efficiency? I suppose we could, in theory. But in practice? Nah. Our species was selected for its sensitivity to the eEDM, not its amenability to quantum state control. It is possible that there is a magic species out there, with highly diagonal frank-condon factors as well as a large enough E_{eff} etc, but if it exists nobody has found it yet. Even if we did know of this magical species, and we could, through sheer force of will, prepare all 20,000 ions in the trap in the science state and read them out with perfect efficiency, this would still only offer us $\sqrt{(40)} \sim 6\times$ improvement in sensitivity, so we would also need to fight some other battles to get the factor of 10 we require. But more realistically: this magic species is unknown and the time it would take to find it, do all the required spectroscopy, etc etc, would put us so far behind our competitors temporally that we would suddenly be needing a factor of 100 in improvement to stay competitive. So again, we are stuck.

That leave us with E_{eff} . This is (obviously) fixed for our current species, but of course one is free to change species! Could we get a factor of 10 that way? Maybe. But changing species is no small task: since we work with bespoke molecules (they simply *must* be sensitive to the eEDM, darling) we have to do all the spectroscopy and figure out all the state preparation steps ourselves before we can even think about performing an EDM measurement. And given a quick survey of the molecules currently in use for eEDM experiments, it's unlikely that one would enjoy a full $10\times$ improvement in E_{eff} . It's more likely to be something like $3\times$. So again, as expected, we are stuck.

Now you may be thinking: 'well, just figure out the right combination of those three things, dummy!' And sure, that *is* a possibility. But how far will that get us? Another $10\times$, then after

that? At some point the potential of this approach is limited: you can't, even if you *reeeeally* want to, trap arbitrarily large amounts of charged species in one location, so you will always end up severely statistics limited. And so we end back up where we started: trapped ion experiments have a rep rate problem.

What one needs is a way to combine the 'unlimited' rep rate of a *beam* experiment with the long coherence times available with *trapped* ions. Eric realized that instead of trapping the ions in one place one could instead 'trap' the ions in a clumpy sort of beam : one can shuffle the ion trap itself along in space, then load another one behind it and another one behind that and so on, until what one has is a *beam* (more like a train) of N ion *clumps*, with state preparation being done at one end, free evolution happening along the travel distance, and state readout happening at the other end. With this method, it's like running N traps in parallel, without multiplying the surrounding infrastructure by the same factor. We have been calling this approach the 'bucket brigade'. Thankfully⁴, the details of the bucket brigade are well beyond the scope of this thesis.

So while our once-revolutionary trapped ion EDM experiment is being burnt to the ground out behind the institute by a group of disgruntled graduate students (proverbially, of course), a new phoenix rises from the ashes, crying out 'good luck, suckers'.

Oh! And before I go, one last thing: unionize your workplace. United we bargain, divided we beg.

⁴At this point we are *all* getting pretty tired of my cutesy use of italics and breaking the 4th wall

Bibliography

- [1] Alan Carrington John M. Brown. Rotational Spectroscopy of Diatomic Molecules. Cambridge University Press, 2003.
- [2] 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):1–5, 2017.
- [3] D. A. Palken et al. Improved analysis framework for axion dark matter searches. Phys. Rev. D, 101:123011, Jun 2020.
- [4] Riccardo Catena and Piero Ullio. A novel determination of the local dark matter density. Journal of Cosmology and Astroparticle Physics, 2010(8), 2010.
- [5] Andrei Derevianko. Detecting dark-matter waves with a network of precision-measurement tools. Physical Review A, 97(4), 2018.
- [6] C. Abel et al. Search for axionlike dark matter through nuclear spin precession in electric and magnetic fields. Physical Review X, 7(4):1–9, 2017.
- [7] Kfir Blum, Raffaele Tito D'Agnolo, Mariangela Lisanti, and Benjamin R. Safdi. Constraining axion dark matter with Big Bang Nucleosynthesis. Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics, 737:30–33, 2014.
- [8] Peter W. Graham and Surjeet Rajendran. New observables for direct detection of axion dark matter. Physical Review D - Particles, Fields, Gravitation and Cosmology, 88(3):1–13, 2013.
- [9] Georg G. Raffelt. Astrophysical methods to constrain axions and other novel particle phenomena. Physics Reports, 198(1-2):1–113, 1990.
- [10] P. A. R. Ade et al. (Planck Collaboration). Planck 2015 results. XIII. Cosmological parameters. Astronomy & Astrophysics, 594(A13):A13, oct 2016.
- [11] Jonathan L. Feng. Naturalness and the status of supersymmetry. Annual Review of Nuclear and Particle Science, 63:351–382, 2013.

- [12] Andrei D Sakharov. Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe. Soviet Physics Uspekhi, 34(5):392–393, may 1991.
- [13] William B Cairncross. Searching for time-reversal symmetry violation with molecular ions: Quantum state control and photofragment imaging. PhD thesis, 2019.
- [14] Markus Greiner, Olaf Mandel, Tilman Esslinger, Theodor W. Hänsch, and Immanuel Bloch. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. Nature, 415(6867):39–44, jan 2002.
- [15] R. Stutz and E. Cornell. Search for the electron edm using trapped molecular ions. Bull. Am. Soc. Phys, 89(76), 2004.
- [16] 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, 2002.
- [17] P. G H Sandars. Measurability of the proton electric dipole moment. Physical Review Letters, 19(24):1396–1398, 1967.
- [18] 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.
- [19] Kang Kuen Ni, Huanqian Loh, Matt Grau, Kevin C. Cossel, Jun Ye, and Eric A. Cornell. State-specific detection of trapped HfF^+ by photodissociation. Journal of Molecular Spectroscopy, 300:12–15, 2014.
- [20] 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, jan 2014.
- [21] E. M. Purcell and N. F. Ramsey. On the Possibility of Electric Dipole Moments for Elementary Particles and Nuclei. Physical Review, 78(6):807–807, jun 1950.
- [22] J. Goldemberg and Y. Torizuka. Upper Limit of the Electric Dipole Moment of the Electron. Physical Review, 129(6):2580–2581, mar 1963.
- [23] E. E. Salpeter. Some Atomic Effects of an Electronic Electric Dipole Moment. Physical Review, 112(5):1642–1648, dec 1958.
- [24] D F Nelson, A A Schupp, R W Pidd, and H R Crane. Search for an Electric Dipole Moment of the Electron. Physical Review Letters, 2(12):492–495, jun 1959.
- [25] E. S. Ensberg. Erratum: Experimental upper limit for the permanent electric dipole moment of Rb85 by optical pumping techniques (Physical Review (1967) 164, 1). Physical Review, 164(1):270, 1967.
- [26] 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, dec 1964.
- [27] Harvey Gould. Search for an Electric Dipole Moment in Thallium. Physical Review Letters, 24(20):1091–1093, may 1970.
- [28] M. A. Player and P G H Sandars. An experiment to search for an electric dipole moment in the $3P_2$ metastable state of xenon. Journal of Physics B: Atomic and Molecular Physics, 3(12):1620–1635, dec 1970.
- [29] S. K. Lamoreaux, J. P. Jacobs, B. R. Heckel, F. J. Raab, and N. Fortson. New constraints on time-reversal asymmetry from a search for a permanent electric dipole moment of Hg_{199} . Physical Review Letters, 59(20):2275–2278, 1987.
- [30] 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, nov 1990.
- [31] 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, oct 1994.
- [32] ACME Collaboration. Improved limit on the electric dipole moment of the electron. Nature, 562(7727):355–360, 2018.
- [33] Ivan Kozyryev and Nicholas R. Hutzler. Precision Measurement of Time-Reversal Symmetry Violation with Laser-Cooled Polyatomic Molecules. Physical Review Letters, 119(13):1–6, 2017.
- [34] Yuhe Zhang, Zixuan Zeng, Qian Liang, Wenhao Bu, and Bo Yan. Doppler cooling of buffer-gas-cooled barium monofluoride molecules. Physical Review A, 105(3):1–5, 2022.
- [35] L. Caldwell, J. A. Devlin, H. J. Williams, N. J. Fitch, E. A. Hinds, B. E. Sauer, and M. R. Tarbutt. Deep Laser Cooling and Efficient Magnetic Compression of Molecules. Physical Review Letters, 123(3):33202, 2019.
- [36] N. J. Fitch, J. Lim, E. A. Hinds, B. E. Sauer, and M. R. Tarbutt. Methods for measuring the electron’s electric dipole moment using ultracold YbF molecules. Quantum Science and Technology, 6(1), 2021.
- [37] R. F. Garcia Ruiz, R. Berger, J. Billowes, C. L. Binnorsley, M. L. Bissell, A. A. Breier, A. J. Brinson, K. Chrysalidis, T. E. Cocolios, B. S. Cooper, K. T. Flanagan, T. F. Giesen, R. P. de Groote, S. Franchoo, F. P. Gustafsson, T. A. Isaev, Koszorús, G. Neyens, H. A. Perrett, C. M. Ricketts, S. Rothe, L. Schweikhard, A. R. Vernon, K. D.A. Wendt, F. Wienholtz, S. G. Wilkins, and X. F. Yang. Spectroscopy of short-lived radioactive molecules. Nature, 581(7809):396–400, 2020.
- [38] Mohit Verma, Andrew M. Jayich, and Amar C. Vutha. Electron Electric Dipole Mo-

- ment Searches Using Clock Transitions in Ultracold Molecules. Physical Review Letters, 125(15):153201, 2020.
- [39] A. C. Vutha, M. Horbatsch, and E. A. Hessels. Orientation-dependent hyperfine structure of polar molecules in a rare-gas matrix: A scheme for measuring the electron electric dipole moment. Physical Review A, 98(3):1–8, 2018.
- [40] Confidential correspondence.
- [41] Op Sushkov and Vv Flambaum. Parity breaking effects in diatomic molecules. Sov. Phys. JETP, 48(4):608–611, 1978.
- [42] S. Eckel, P. Hamilton, E. Kirilov, H. W. Smith, and D. Demille. Search for the electron electric dipole moment using Ω -doublet levels in PbO. Physical Review A - Atomic, Molecular, and Optical Physics, 87(5):1–19, 2013.
- [43] A. E. Leanhardt, J. L. Bohn, H. Loh, P. Maletinsky, E. R. Meyer, L. C. Sinclair, R. P. Stutz, and E. A. Cornell. High-resolution spectroscopy on trapped molecular ions in rotating electric fields: A new approach for measuring the electron electric dipole moment. Journal of Molecular Spectroscopy, 270(1):1–25, 2011.
- [44] Stuart J. Freedman Dmitry Budker, Philip H. Bucksbaum, editor. Art and symmetry in experimental physics : festschrift for Eugene D. Commins, Berkeley, California, 20-21 May 2001, volume 596. American Institute of Physics, AIP conference proceedings, 2001.
- [45] Edmund R. Meyer, Aaron E. Leanhardt, Eric A. Cornell, and John L. Bohn. Berry-like phases in structured atoms and molecules. Physical Review A - Atomic, Molecular, and Optical Physics, 80(6), 2009.
- [46] Edmund R. Meyer, John L. Bohn, and Michael P. Deskevich. Candidate molecular ions for an electron electric dipole moment experiment. Physical Review A - Atomic, Molecular, and Optical Physics, 73(6):1–10, 2006.
- [47] Yan Zhou, Yuval Shagam, William B. Cairncross, Kia Boon Ng, Tanya S. Roussy, Tanner Grogan, Kevin Boyce, Antonio Vigil, Madeline Pettine, Tanya Zelevinsky, Jun Ye, and Eric A. Cornell. Second-Scale Coherence Measured at the Quantum Projection Noise Limit with Hundreds of Molecular Ions. Physical Review Letters, 124(5):53201, 2020.
- [48] Yuval Shagam, William B. Cairncross, Tanya S. Roussy, Yan Zhou, Kia Boon Ng, Daniel N. Gresh, Tanner Grogan, Jun Ye, and Eric A. Cornell. Continuous temporal ion detection combined with time-gated imaging: Normalization over a large dynamic range. Journal of Molecular Spectroscopy, 368:111257, 2020.
- [49] Kevin C Cossel. Techniques in molecular spectroscopy : from broad bandwidth to high resolution. PhD thesis, University of Colorado Boulder and JILA, 2014.
- [50] Huanqian Loh. Search for an electron electric dipole moment with trapped molecular ions. PhD thesis, University of Colorado Boulder and JILA, 2013.

- [51] Matt Grau. Measuring the Electron Electric Dipole Moment with Trapped Molecular Ions. PhD thesis, University of Colorado Boulder, 2016.
- [52] Kevin C. Cossel, Daniel N. Gresh, Laura C. Sinclair, Tyler Coffey, Leonid V. Skripnikov, Alexander N. Petrov, Nikolai S. Mosyagin, Anatoly V. Titov, Robert W. Field, Edmund R. Meyer, Eric A. Cornell, and Jun Ye. Broadband velocity modulation spectroscopy of HfF^+ : Towards a measurement of the electron electric dipole moment. Chemical Physics Letters, 546:1–11, 2012.
- [53] N. J. Stone. Table of nuclear magnetic dipole and electric quadrupole moments. Atomic Data and Nuclear Data Tables, 90(1):75–176, 2005.
- [54] Timo Fleig and Malaya K. Nayak. Electron electric-dipole-moment interaction constant for HfF^+ from relativistic correlated all-electron theory. Physical Review A - Atomic, Molecular, and Optical Physics, 88(3):1–6, 2013.
- [55] Huanqian Loh, Jia Wang, Matt Grau, Tyler S. Yahn, Robert W. Field, Chris H. Greene, and Eric A. Cornell. Laser-induced fluorescence studies of HfF produced by autoionization. Journal of Chemical Physics, 135(15), 2011.
- [56] M. G. Kozlov and L. N. Labzowsky. Parity violation effects in diatomics. Journal of Physics B: Atomic, Molecular and Optical Physics, 28(10):1933–1961, 1995.
- [57] Timothy Chupp and Michael Ramsey-Musolf. Electric dipole moments: A global analysis. Physical Review C - Nuclear Physics, 91(3):1–11, 2015.
- [58] Malika Denis, Morten S. Norby, Hans Jorgen Aa Jensen, André Severo Pereira Gomes, Malaya K. Nayak, Stefan Knecht, and Timo Fleig. Theoretical study on ThF^+ , a prospective system in search of time-reversal violation. New Journal of Physics, 17, 2015.
- [59] Timo Fleig and Martin Jung. Model-independent determinations of the electron EDM and the role of diamagnetic atoms. Journal of High Energy Physics, 2018(7):12, jul 2018.
- [60] Tanya S. Roussy, Daniel A. Palken, William B. Cairncross, Benjamin M. Brubaker, Daniel N. Gresh, Matt Grau, Kevin C. Cossel, Kia Boon Ng, Yuval Shagam, Yan Zhou, Victor V. Flambaum, Konrad W. Lehnert, Jun Ye, and Eric A. Cornell. Experimental Constraint on Axionlike Particles over Seven Orders of Magnitude in Mass. Physical Review Letters, 126(17):171301, 2021.
- [61] William B. Cairncross, Daniel N. Gresh, Matt Grau, Kevin C. Cossel, Tanya S. Roussy, Yiqi Ni, Yan Zhou, Jun Ye, and Eric A. Cornell. A precision measurement of the electron's electric dipole moment using trapped molecular ions. 2017.
- [62] C. Bartram, T. Braine, E. Burns, R. Cervantes, N. Crisosto, N. Du, H. Korandla, G. Leum, P. Mohapatra, T. Nitta, L. J. Rosenberg, G. Rybka, J. Yang, John Clarke, I. Siddiqi, A. Agrawal, A. V. Dixit, M. H. Awida, A. S. Chou, M. Hollister, S. Knirck, A. Sonnenschein, W. Wester, J. R. Gleason, A. T. Hipp, S. Jois, P. Sikivie, N. S. Sullivan, D. B. Tanner, E. Lentz, R. Khatiwada, G. Carosi, N. Robertson, N. Woollett, L. D. Duffy, C. Boutan, M. Jones, B. H.

- Laroque, N. S. Oblath, M. S. Taubman, E. J. Daw, M. G. Perry, J. H. Buckley, C. Gaikwad, J. Hoffman, K. W. Murch, M. Goryachev, B. T. McAllister, A. Quiskamp, C. Thomson, and M. E. Tobar. Search for Invisible Axion Dark Matter in the 3.3-4.2 μeV Mass Range. Physical Review Letters, 127(26):261803, 2021.
- [63] Daniel A Palken. Enhancing the scan rate for axion dark matter: Quantum noise evasion and maximally informative analysis. PhD thesis, pages 167-168, University of Colorado, 2020.
- [64] Gary P. Centers, John W. Blanchard, Jan Conrad, Nataniel L. Figueroa, Antoine Garcon, Alexander V. Gramolin, Derek F. Jackson Kimball, Matthew Lawson, Bart Pelssers, Joseph A. Smiga, Yevgeny Stadnik, Alexander O. Sushkov, Arne Wickenbrock, Dmitry Budker, and Andrei Derevianko. Stochastic amplitude fluctuations of bosonic dark matter and revised constraints on linear couplings. arxiv:1905.13650v4, 2019.
- [65] M. E. Pospelov and I. B. Khriplovich. Electric dipole moment of the W boson and the electron in the Kobayashi-Maskawa model. Sov. J. Nucl. Phys., 53:638–640, 1991.
- [66] M. S. Safronova, D. Budker, D. Demille, Derek F. Jackson Kimball, A. Derevianko, and Charles W. Clark. Search for new physics with atoms and molecules. Reviews of Modern Physics, 90(2):25008, 2018.
- [67] Maxim Pospelov and Adam Ritz. Electric dipole moments as probes of new physics. Annals of Physics, 318(1):119–169, jul 2005.
- [68] Yuichiro Nakai and Matthew Reece. Electric dipole moments in natural supersymmetry. Journal of High Energy Physics, 2017(8):31, aug 2017.
- [69] 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, 2013.
- [70] J. H. Smith, E. M. Purcell, and N. F. Ramsey. Experimental limit to the electric dipole moment of the neutron. Physical Review, 108(1):120–122, oct 1957.
- [71] Dmitry Budker, Peter W. Graham, Micah Ledbetter, Surjeet Rajendran, and Alexander O. Sushkov. Proposal for a cosmic axion spin precession experiment (CASPER). Physical Review X, 4(2):1–10, 2014.
- [72] Peter W. Graham and Surjeet Rajendran. Axion dark matter detection with cold molecules. Physical Review D, 84(5):055013, sep 2011.
- [73] Christopher T. Hill. Axion induced oscillating electric dipole moments. Physical Review D, 91(11):111702, jun 2015.
- [74] V. V. Flambaum, B. M. Roberts, and Y. V. Stadnik. Comment on "axion induced oscillating electric dipole moments". Physical Review D, 95(5):1–2, 2017.

- [75] Y. V. Stadnik and V. V. Flambaum. Axion-induced effects in atoms, molecules, and nuclei: Parity nonconservation, anapole moments, electric dipole moments, and spin-gravity and spin-axion momentum couplings. Physical Review D, 89(4):1–14, 2014.
- [76] David J.E. Marsh. Axion cosmology. Physics Reports, 643:1–79, jul 2016.
- [77] R. D. Peccei and Helen R. Quinn. CP conservation in the presence of pseudoparticles. Physical Review Letters, 38(25):1440–1443, 1977.
- [78] R. D. Peccei and Helen R. Quinn. Constraints imposed by CP conservation in the presence of pseudoparticles. Physical Review D, 16(6):1791–1797, 1977.
- [79] F. Wilczek. Problem of strong P and T invariance in the presence of instantons. Physical Review Letters, 40(5):279–282, 1978.
- [80] Steven Weinberg. A New Light Boson? Physical Review Letters, 40(4):223–226, jan 1978.
- [81] Asimina Arvanitaki, Savas Dimopoulos, Sergei Dubovsky, Nemanja Kaloper, and John March-Russell. String axiverse. Physical Review D, 81(12):123530, jun 2010.
- [82] Michael Dine and Willy Fischler. The not-so-harmless axion. Physics Letters B, 120(1-3):137–141, jan 1983.
- [83] L.F. Abbott and P. Sikivie. A cosmological bound on the invisible axion. Physics Letters B, 120(1-3):133–136, jan 1983.
- [84] John Preskill, Mark B. Wise, and Frank Wilczek. Cosmology of the invisible axion. Physics Letters B, 120(1-3):127–132, jan 1983.
- [85] Wayne Hu, Rennan Barkana, and Andrei Gruzinov. Fuzzy Cold Dark Matter: The Wave Properties of Ultralight Particles. Physical Review Letters, 85(6):1158–1161, aug 2000.
- [86] Hsi-Yu Schive, Tzihong Chiueh, and Tom Broadhurst. Cosmic structure as the quantum interference of a coherent dark wave. Nature Physics, 10(7):496–499, jul 2014.
- [87] Lam Hui, Jeremiah P Ostriker, Scott Tremaine, and Edward Witten. Ultralight scalars as cosmological dark matter. Physical Review D, 95(4):043541, feb 2017.
- [88] Andrzej K. Drukier, Katherine Freese, and David N Spergel. Detecting cold dark-matter candidates. Physical Review D, 33(12):3495–3508, jun 1986.
- [89] Y. V. Stadnik (unpublished calculations).
- [90] V. V. Flambaum and H. B. Tran Tan. Oscillating nuclear electric dipole moment induced by axion dark matter produces atomic and molecular EDM. Physical Review D, 100(11):111301, 2019.
- [91] V. V. Flambaum, M. Pospelov, A. Ritz, and Y. V. Stadnik. Sensitivity of EDM experiments in paramagnetic atoms and molecules to hadronic CP violation. arxiv:1912.13129, 2019.

- [92] V. V. Flambaum, I. B. Samsonov, and H. B. Tran Tan. Limits on CP-violating hadronic interactions and proton EDM from paramagnetic molecules. [arxiv:2004.10359](https://arxiv.org/abs/2004.10359), 2020.
- [93] The authors of [91, 92] calculate the magnitude of the scalar-pseudoscalar electron-nucleon coupling C_S that would be generated by a heavy nucleus like Hf by a nonzero QCD θ parameter. An ALP coupled to gluons behaves like a dynamical version of θ_{QCD} , and thus generates an oscillating C_S .
- [94] Here, m_F refers to the Zeeman sublevel and Ω refers to the projection of the sum of electronic and rotational angular momentum onto the molecular axis.
- [95] The reader would be justified to scoff at calling μHz to mHz ‘high frequency’, but we beg their patience.
- [96] See Supplemental Material, which includes Refs. [101, 45].
- [97] If one searches a large area in parameter space for some effect (in our case, over 10^6 possible ALP frequencies for an oscillation signal), one is bound to see a statistical fluctuation. The look-elsewhere effect is the name for this phenomenon: if one continues to ‘look elsewhere’, they may find what they are looking for.
- [98] G. P. Centers et al. Stochastic amplitude fluctuations of bosonic dark matter and revised constraints on linear couplings, 2019.
- [99] Joshua W Foster, Nicholas L Rodd, and Benjamin R Safdi. Revealing the dark matter halo with axion direct detection. *Physical Review D*, 97(12):123006, jun 2018.
- [100] We expect forthcoming nucleus-specific calculations to improve the inferred sensitivity [92].
- [101] L. Eyer and P. Bartholdi. Variable stars: Which Nyquist frequency? *Astronomy and Astrophysics Supplement Series*, 135(1):1–3, 1999.
- [102] \mathcal{A}'_i is positive when we prepare the same state as we read out, and negative when we prepare the opposite state to the one we read out.
- [103] In our previous constraint of the static eEDM, we used a slightly different asymmetry, which required 2 shots to collect: $\mathcal{A} = \frac{N_{\uparrow} - N_{\downarrow}}{N_{\uparrow} + N_{\downarrow}}$, where $N_{\uparrow/\downarrow}$ indicates the number of Hf^+ molecules we counted from the $m_F = \pm 3/2$ stretched states, respectively. For this analysis, we use a ‘single-shot’ definition instead, as it allows us to extend our analysis band out to 400 mHz.
- [104] W. M. Itano, J C Bergquist, J J Bollinger, J. M. Gilligan, D. J. Heinzen, F. L. Moore, M. G. Raizen, and D. J. Wineland. Quantum projection noise: Population fluctuations in two-level systems. *Physical Review A*, 47(5):3554–3570, may 1993.

Appendix A

Lasers

Our experiment used a lot of lasers. Here is a master list, much too late to be of any use for anyone actually doing the experiment.

A.1 Neutral HfF production / ablation: "Abby" 532 nm

Continuum 'Minilite' pulsed Nd:YAG Laser

Abby is a Class 4 Pulsed Q-switched Nd:YAG laser, with a 2nd harmonic crystal for 532 nm output. The output pulses are 10 ns long and deliver ~ 10 mJ per pulse.

We use a dichroic mirror on the output to dump 1064, while 532 goes to the chamber. Abby is a variable rep rate laser, and is triggered by the LabView program's DIO.

The table below is from the manual, which is stored in the bottom drawer of the filing cabinet.

Wavelength	1064 nm	532 nm
Energy	50 mJ	25 mJ
Peak Power	8.3 MW	6.3 MW
Average Power	750 mW	375 mW
Pulsewidth	5-7 nsec	3-4 nsec
Stability	2%	3%
Polarization	Horz	Vert
Jitter	± 0.5 nsec	± 0.5 nsec
Beam Size	< 3 mm	< 3 mm
Divergence	< 3 mrad	< 3 mrad
Rep Rate	1-15 Hz	1-15 Hz

Table A.1: Minilite II specs from manual

A.2 Photionization 1+2 pump: "Mia" 532 nm

Spectra Physics 'Quanta Ray Pro Series' pulsed Nd:YAG laser

Part specification: Quanta-Ray-Pro-250-10. This 10 Hz YAG pumps Harambe and Bertha. The undoubled 1064 nm fundamental gets dumped internally. 532 nm is the only output. We use about 25 mJ to pump Harambe and 250 mJ to pump Bertha. The serial number is 2344PR.

The table below is from the manual, which is stored in the bottom drawer of the filing cabinet.

Wavelength	1064 nm	532 nm	266 nm
Energy	1250 mJ/p	650 mJ/p	130 mJ/p
Pulsewidth	8-12 ns	1-2 ns	3-4 ns
Short Term Energy Stability	$\pm 2\%$	$\pm 3\%$	$\pm 8\%$

Table A.2: Quanta-Ray-Pro-250-10 specs from manual

A.3 Photoionization 1 seed: "Tuna" 1237.55598 nm

Toptica 'TA Pro' external cavity diode laser

Tuna is a commercial ECDL with a TA, whose fundamental wavelength is set to 1237.55598 nm. We send about 350 mW into a fibre, which goes to a waveguide doubler (NTT Electronics Wavelength Conversion Module WH-0619-000-A-B-C) with a maximum power input of 200 mW. Out of the waveguide doubler we get about 70 mW of doubled output (618.77799) which is used to seed Harambe, our pulsed dye amplifier. The wavelength of the doubled light is controlled by the temperature of waveguide doubler. The serial number is 025112, The TA Article number is TA-1250-1500-1.

Tuna's datasheets are stored in the tall black cabinet as well as digitally in the 'Lasers' lab notebook. The table below summarizes a few key parameters. The 'actual' parameters drift around a bit, but they give a good starting point if you're looking to get it locked to the nominal wavelength.

Wavelength (nominal)	1237.5 nm
Output Power (Actual)	370 mW into fiber
Output Power (Max)	2000 mW
Diode Current (Actual)	267.9 mA
Diode Current (Max)	272 mA
Piezo Voltage (Actual)	94.643 V
Lasing Threshold	43 mA
TA Current (Actual)	3840 mA
TA Current (Max)	7000 mA
Linewidth (from OSA spectra I took)	0.1 nm FWHM

Table A.3: A few Tuna (Toptica TA Pro) specs from the data sheet

A.4 Photoionization 1 (first photon): "Harambe" 309.388 nm

Radiant 'NarrowAmp' pulsed dye amplifier

Harambe is a 3 stage pulsed dye amplifier¹. Note there isn't much of a manual for this (apparently it was shipped without one and when Yan requested one they sent him a hastily written 10 page Word document).

Harambe is pumped by Mia (532 nm) and seeded by doubled Tuna (618.778 nm). The photon used for ionization is the doubled output of Harambe (309.388 nm). We use around 30 μ J per pulse, and when the amplifier is well-aligned with 50 mJ of pump power and 60 mW of seed power we are able to get \sim 150 μ J out.

Harambe's dyes are Rhodamine 610 chloride and Rhodamine 640 perchlorate. We pre-make 1 g/L concentrate of both dyes which we then dilute when we fill the circulators. Harambe's dye volumes are (all assuming 1g/L concentrate):

- first stage: 165 mL Rhodamine 610, 185 mL Rhodamine 640, 725 mL ethanol
- second stage: 90 mL Rhodamine 610, 100 mL Rhodamine 640, 390 mL ethanol
- third stage: 90 mL Rhodamine 610, 100 mL Rhodamine 640, 880 mL ethanol

Photoionization is a two-photon process, which promotes a valence electron to a Rydberg state 54 cm^{-1} above the ionization potential. Once in this state, the molecule quickly autoionizes.

Here are some specs I found at the following website: <https://www.radiant-dyes.com/PDF/NarrowAmp.pdf>, plus specs for our specific usage - generating 618nm light then doubling it to 309 nm.

¹Technically not a laser, just an amplifier, but nonetheless relevant to put here. Sorry to be pedantic, but this is a PHD thesis after all.

Tuning Range	540-800 nm
Linewidth (dependent on pump laser)	approx. < MHz 100 mS
Maximum Output Power, doubled (309 nm) (50 mJ pump, 60 mW seed)	150 μ J
Maximum Output Power, doubled (309 nm) (21 mJ pump, 60 mW seed)	70 μ J
Output pulse energy (from url) (at 730 nm, pumped 45 mJ with 532 nm)	> 9mJ
Power Stability	< 6%
Pointing Stability	< 50 μ rad
Divergence	< 0.5mrad

Table A.4: Select Harambe (Radiant Narrow Amp) specs

A.5 Photoionization 2 (second photon): "Bertha" 735.464 nm

Sirah 'Precision Scan' dye laser

Bertha is a tunable pulsed dye laser pumped by Mia (we typically use about 250 mJ pump power). Currently, we don't use the tuneability. The output beam is S polarized. The fundamental is doubled, and it is the doubled output which is used for ionization. We use around 2 mJ of doubled power for ionization.

Bertha's dye is Pyridine 2 (LDS 722). We pre-make 0.5 g/L concentrate we then dilute when we fill the circulators. Bertha's dye volumes are (all assuming 0.5 g/L concentrate):

- Resonator: 252 mL dye concentrate, 448 mL ethanol
- Pre-Amp: 173 mL dye concentrate, 1027 mL ethanol
- Cap Amp: 90 mL dye concentrate, 1910 mL ethanol

Here are some specs I found at the following website: <https://www.sirah.com/wp-content/uploads/documents/PrecisionScan.pdf>, plus specs for our specific usage - generating 735 nm light then doubling it to 367.7 nm.

Linewidth (dependent on grating, our specific are unknown)	1 – 4 pm
Absolute wavelength accuracy	< 15 pm
Wavelength Stability	1.5pm/°C
Divergence (typical)	0.4mrad
Polarization	> 98% (vertical)
Pump Power (532 nm, our use)	250 mJ
Maximum Output Power (doubled, with 250 mJ pump)	5 mJ

Table A.5: Select Harambe (Radiant Narrow Amp) specs

A.6 State transfer: "Toptica" 961.43495 nm

Toptica 'TA Pro' external cavity diode laser

Toptica is a commercial ECDL with a TA, it's named 'Toptica' because it was the very first Toptica laser we got in the lab and apparently it never occurred to anyone that we would get another one. It's probably easier to call this laser 'Transfer' but it's too late for that! We couple about 1W into a PM fiber which is sent over to the chamber then directed through the top of the trap (with linear polarization). We send about 600 mW into the trap.

Toptica's datasheets are stored in the tall black cabinet as well as digitally in the 'Lasers' lab notebook. The table below summarizes a few key parameters.

Wavelength (nominal)	986.45 nm
Output Power (Actual)	1.3 W before fiber, 620 mW after fiber
Output Power (Max)	2060 mW
Diode Current (Actual)	94 mA
Diode Current (Max)	110 mA
Lasing Threshold	27 mA
TA Current (Actual)	2877 mA
TA Current (Max)	2930 mA

Table A.6: A few Toptica/Transfer (Toptica TA Pro) specs

A.7 Cleanup: "Vibational Cleanup" 818.372 nm

Homebuilt external cavity diode laser

Vibrational Cleanup is a homebuilt ECDL, using a ThorLabs diode centered at 818 nm. We get about 68 mW out which we send into a fiber then into the top of the trap (linearly polarized). We get about 17.5 mW into the chamber.

There are no datasheets or manuals for this laser. Here's our current operating parameters:

Wavelength (nominal)	818 nm
Output Power (Actual)	68.5 mW before the fiber, 17.5 mW into the chamber
Diode Current (Actual)	291.4 mA
Piezo Voltage (Actual)	73.18 mV

Table A.7: A few Vibrational Cleanup (Homebuilt ECDL) specs

A.8 mF pumping: "DBR/1083" 1082.4137 nm

I couldn't find much documentation for this one

Operating Temperature (Actual)	7.5 °C
Diode Current (Actual)	318.36 mA
Power into isolator (Actual)	170 mW
Power into AOM (Actual)	120 mW
Power into fiber at 49% duty cycle (Actual)	46 mW
Power into trap at 49% duty cycle (Actual)	21 mW

Table A.8: A few Vibrational Cleanup (Homebuilt ECDL) specs

A.9 Depletion: "814/Depletion" 814.508 nm

Toptica 'TA Pro' external cavity diode laser

Depletion is a commercial ECDL with TA. The laser puts out about 1W. We couple the light into an AOM which is modulated at a 50% duty cycle at f_{rot} . The AOM's zeroth order output is coupled into a fiber and sent into the top of the trap for single doublet depletion. The AOM's first order diffracted light is sent into a polarization setup (known as a Berek Polarization Compensator) to obtain circularly polarized light, which is then sent into the side of the trap. We send about 500 mW into the side of the trap for mF level depletion and about 50 mW into the top of the trap for single doublet depletion.

Depletion's datasheets are stored in the tall black cabinet as well as digitally in the 'Lasers' lab notebook. The table below summarizes a few key parameters.

Wavelength (nominal)	815.0 nm
Output Power (Actual)	540 mW into trap at 50% duty cycle
Output Power (Max)	2100 mW
TA Current (Set)	2610 mA
TA Current (Max)	2800 mA
Diode Current (Set)	111.665 mA
Diode Current (Max)	128 mA
Piezo Voltage (Actual)	90.83 V
Lasing Threshold	56 mA

Table A.9: A few Depletion (Toptica TA Pro) specs

A.10 Photodissociation 1 (first photon): "Cora" 736.989/2 nm or 736.986/2 nm

Sirah 'Cobra Stretch' dye laser

Cora is a tunable pulsed dye laser pumped by Olga. We use the tunability to dissociate out of different states. All values here are for dissociating out of ${}^3\Delta_1, v = 0, J = 1$. The output beam is S polarized. The fundamental is doubled, and it is the doubled output which is used for dissociation. We use around 1.5 mJ of doubled power for dissociation.

Cora's dye is Pyridine 2 (LDS 722). We pre-make 0.5 g/L concentrate we then dilute when we fill the circulators. Cora's dye volumes are (all assuming 0.5 g/L concentrate):

- Resonator: 280 mL dye concentrate, 420 mL ethanol
- Amplifier: 144 mL dye concentrate, 936 mL ethanol

Here are some specs I found at the following website: <https://www.sirah.com/wp-content/uploads/documents/CobraStretch.pdf>, plus specs for our specific usage - generating 736 nm light then doubling it to 368.5 nm.

Linewidth (dependent on grating, our specific are unknown)	1 – 4 pm
Absolute wavelength accuracy	< 15 pm
Wavelength Stability	1.5pm/°C
Divergence (typical)	1.0mrad
Polarization	> 98% (vertical)
Pump Power (532 nm, our use)	170 mJ
Fundamental measured before doubling (at 736.989 nm)	20 mJ
Power measured right after doubling (at 736.989/2 nm)	3 mJ
Power measured after the chamber, (at 736.989/2 nm)	1.6 mJ

Table A.10: Select Cora (Syrah Cobra Stretch) specs

A.11 Photodissociation 2 (second photon) + Cora pump: "Olga" 266 nm

Spectra Physics 'Quanta Ray Lab Series' pulsed Nd:YAG laser

Part specification: Quanta-Ray-Lab-170-30H - note we bought this as a 30 Hz laser so the last part of the part number is 30 H. But we had it converted to 10 Hz, so the last part of the part number should be changed to 10H. This 10 Hz YAG pumps Cora and provides 266 nm for dissociation. The undoubled 1064 nm fundamental gets dumped internally. 532 nm output pumps Cora (170 mJ, remainder is dumped). 266 nm output is sent straight to the trap for dissociation. The serial number is 4364L.

The table below is from the manual, which is stored in the bottom drawer of the filing cabinet.

Wavelength	1064 nm	532 nm	266 nm
Energy	850 mJ/p	450 mJ/p	90 mJ/p
Pulsewidth	8-12 ns	1-2 ns	3-4 ns
Short Term Energy Stability	±2%	±3%	±8%

Table A.11: Quanta-Ray-Lab-170-10H specs from manual. Note the manual specified the pulse width in nm but I'm pretty sure that's a typo, especially given footnote 1 on table 3-2.

Appendix B

One shot asymmetry

We normally define our asymmetry as $\mathcal{A} = (N_{\downarrow} - N_{\uparrow}) / (N_{\downarrow} + N_{\uparrow})$, which takes two ‘shots’ of the experiment, but one can instead define it with only one shot. We did this for our dark matter analysis [60]. We define the single-shot asymmetry \mathcal{A} as

$$\mathcal{A} = \pm \frac{N_i - \langle N \rangle}{\langle N \rangle}, \quad (\text{B.1})$$

where we choose $+$ when the preparation and the depletion have the same phase (known colloquially as ‘in-in’ or ‘anti-anti’) and $-$ when preparation and depletion have the opposite phase. N_i is the Hf⁺ count in a single shot, and $\langle N \rangle$ is the mean Hf⁺ count for the set of shots (datapoint) to which the single shot belongs.

Below is a little thing I worked out on how background affects the measured fringe contrast, in case it’s useful to someone one day.

B.0.1 One-shot asymmetry and measured fringe contrast

As mentioned above, we define the single-shot asymmetry \mathcal{A} as

$$\mathcal{A} = \pm \frac{N_i - \langle N \rangle}{\langle N \rangle} \quad (\text{B.2})$$

where we choose $+$ when the preparation and the depletion have the same phase (known colloquially as ‘in-in’ or ‘anti-anti’) and $-$ when preparation and depletion have the opposite phase.

N_i is the Hf^+ count in a single shot, and $\langle N \rangle$ is the mean Hf^+ count for the set of shots (datapoint) to which the single shot belongs. Let's see how this definition works. Assume, in some perfect experiment, that we prepare the same number of ions each time, and that our state preparation and dissociation are perfect. If you consider zero free-evolution time, then the in-in and anti-anti shots will result in N Hf^+ ions being read out, while the in-anti and anti-in shots will result in 0 Hf^+ ions being read out. For a given set of shots, called a **data point** we will alternate between in-in and anti-in **or** anti-anti and in-anti. For any set of these chops, the **mean** number of ions we read out will be $N/2$. For the in-in and anti-anti shots we have the asymmetry

$$\mathcal{A} = +\frac{N - N/2}{N/2} = 1 \quad (\text{B.3})$$

while for the in-anti or anti-in shots we will have the asymmetry

$$\mathcal{A} = -\frac{0 - N/2}{N/2} = 1 \quad (\text{B.4})$$

so you see that we can combine the shots with opposite depletion appropriately. In this perfect experiment, as we increase the free-evolution time, the asymmetry \mathcal{A} traces out a sinusoid which oscillates between ± 1 at a frequency proportional to the energy splitting.

Now what happens if you add a constant background? This is an important question because we know that there is a background which is independent of the free-evolution time. It seems to be due to imperfect ionization and dissociation. For simplicity let's continue to assume the ion number is constant and so on, and that the background is constant. Call the background Hf number b . Now the mean Hf (for a data point) is $N/2 + b$, the minimum number is b , and

the maximum number is $N + b$. Our in-in or anti-anti shots will have the asymmetry

$$\begin{aligned}
 \mathcal{A} &= + \frac{(N + b) - (N/2 + b)}{(N/2 + b)} \\
 &= \frac{N + b}{N/2 + b} - 1 \\
 &= \frac{2N + 2b}{N + 2b} - 1 \\
 &= \frac{2N}{N + 2b} + \frac{2b}{N + 2b} - 1
 \end{aligned}$$

while the anti-in or in-anti shots will give

$$\begin{aligned}
 \mathcal{A} &= -1 - \frac{b}{N/2 + b} \\
 &= 1 - \frac{2b}{N + 2b}
 \end{aligned}$$

Averaging these two will give

$$\begin{aligned}
 \mathcal{A} &= \left(\frac{2N}{N + 2b} + \frac{2b}{N + 2b} - 1 + 1 - \frac{2b}{N + 2b} \right) / 2 \\
 &= \frac{N}{N + 2b} \\
 &< 1
 \end{aligned}$$

Which we see leads to a loss in contrast.

Our measured fringe contrast varies between 0.6-0.85. It is possible that the imperfection is due to more than just a constant background, such as imperfect depletion, pumping, or $\pi/2$ pulses, but it is at least in part due to the background that we always observe.

Appendix C

A list of useful transitions

There's a lot to keep track of in our experiment, and at one point I made the diagram below and put it as the desktop background on the main experiment computer. Until then, there wasn't any central place where all the useful transitions we use were catalogued. Turns out that kind of thing is quite useful so I am saving that figure here, for posterity¹. Enjoy.

¹I don't actually know what the word 'posterity' means and I will not look it up

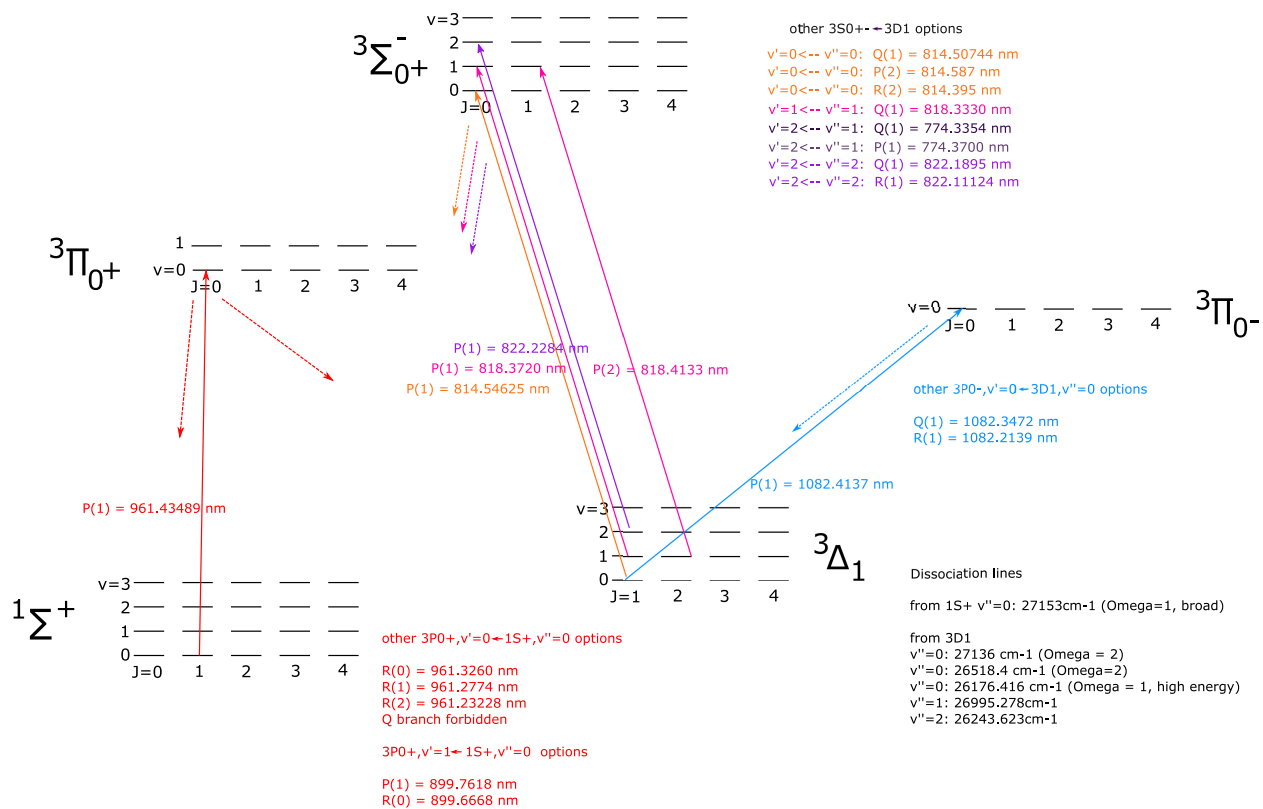


Figure C.1: The infamous 'laser diagram'

Appendix D

An approximation to our shot noise limit

Here's my derivation of the shot noise limit if we knew the total number of ions we were starting with.

Ultimately, the core of our experiment is spin precession experiment (Ramsey spectroscopy) between stretched states. We prepare an ensemble of ions (for now, let's say that we prepare the same number each time N_{tot}) whose valence electrons are in a pure spin state $|\uparrow\rangle$ or $|\downarrow\rangle$, we take that pure spin state and rotate it by $\pi/2$ into an equal superposition $(|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$, then we wait a varying amount of time before we project the superposition back onto one of the basis vectors ($|\uparrow\rangle$ or $|\downarrow\rangle$), and we count N_{\uparrow} or N_{\downarrow} (for now let's assume that $N_{\uparrow} + N_{\downarrow} = N_{\text{tot}}$). For any quantum two-level system prepared in a superposition $|\psi\rangle = c_{\uparrow}|\uparrow\rangle + c_{\downarrow}|\downarrow\rangle$, the probability of measuring (projecting the state into) a given basis vector is given by the modulus squared of the coefficient¹: $p_{\uparrow} = |c_{\uparrow}|^2$, $p_{\downarrow} = |c_{\downarrow}|^2$. This probabilistic/non-deterministic essence is inescapable (it's one of the core *features* of quantum mechanics, after all²), so unless either c_{\downarrow} or c_{\uparrow} is 0 or 1, there will always be some uncertainty in the outcome. When an experiment has a binary outcome with some probability (eg p_{\downarrow} , like we have here), and it is repeated over and over, one will find that the measured outcomes obey a binomial distribution: if we have N_{tot} ions, and N_{\uparrow} is the number of ions in state $|\uparrow\rangle$ (and N_{\downarrow} is the number of ions in state $|\downarrow\rangle$, $N_{\uparrow} + N_{\downarrow} = N_{\text{tot}}$), then the probability

¹Assuming that the state vector is properly normalized, ie $|c_{\uparrow}|^2 + |c_{\downarrow}|^2 = 1$.

²As far as we know

of measuring a given value of N_{\uparrow} , ie $P(N_{\uparrow})$, is³

$$P(N_{\uparrow}, N_{\text{tot}}, p_{\uparrow}) = \frac{N_{\text{tot}}!}{N_{\uparrow}!(N_{\text{tot}} - N_{\uparrow})!} p_{\uparrow}^{N_{\uparrow}} (1 - p_{\uparrow})^{N_{\text{tot}} - N_{\uparrow}}. \quad (\text{D.1})$$

The variance of the binomial distribution is

$$\sigma^2 = N_{\text{tot}} p_{\uparrow} (1 - p_{\uparrow}), \quad (\text{D.2})$$

which goes to zero when $p_{\uparrow} = 0$ or $p_{\uparrow} = 1$, and has its maximum value of $N_{\text{tot}}/4$ when $p_{\uparrow} = 1/2$.

It is this variance in the measured outcome that is the ultimate limit to our errorbars ⁴.

Now there's one final detail: during our Ramsey experiment the spin *precesses*, so the initial state $(|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$ evolves such that at a later time t_R after a second $\pi/2$ pulse the state vector can be written as $\sim \cos(\omega t_R) |\uparrow\rangle + \sin(\omega t_R) |\downarrow\rangle$. So now the probabilities of measuring a given basis vector become sinusoidally time-varying: $p_{\uparrow} = |c_{\uparrow}|^2 = |\cos(\omega t_R)|^2$, which goes right into the binomial distribution of equation D.1 and the corresponding variance (eq D.2):

$$\sigma^2(t_R) = N_{\text{tot}} |\cos(\omega t_R)|^2 (1 - |\cos(\omega t_R)|^2). \quad (\text{D.3})$$

This is the quantum projection noise limited uncertainty for any measurement we will make of an ensemble of independent ions. In other words, this is the minimum uncertainty we can expect in any measurement of the number of ions in a given stretched state. How will this uncertainty affect our *total* uncertainty?

To compute the quantum projection noise limited total uncertainty, we need to go from an uncertainty in a single spin asymmetry measurement to an uncertainty in the precession frequency, then from that to a measurement of d_e . I'll sketch out a heuristic argument using only the sides of the fringe, I expect the rigorous argument is more complex than we need. Let's step

³If you haven't studied probability distributions, the way to read $P(N_{\uparrow}, N_{\text{tot}}, p_{\uparrow})$ is "the probability of measuring N_{\uparrow} given N_{tot} and p_{\uparrow} is...

⁴If I've explained this poorly, or if you're just looking for more details, I would strongly recommend you read [104]

through a perfect, quantum-projection-noise-limited version of our experiment and keep track of the uncertainties:

First, we prepare an ensemble of ions N_{tot} in a superposition state $(|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$, let it evolve, apply a second $\pi/2$ pulse, then measure the number of ions in the up state, N_{\uparrow} . The probability of measuring a given value of N_{\uparrow} (on the side of a fringe, we expect $N_{\uparrow} = N_{\text{tot}}/2$) is given by the binomial distribution with variance $\sigma_{\uparrow/\downarrow}^2 = N_{\text{tot}}p_{\uparrow}(1 - p_{\uparrow}) = N_{\text{tot}}/4$, since $p_{\uparrow} = 1/2$ on the side of a fringe.

Then we repeat the same process to measure N_{\downarrow} and again $\sigma_{\uparrow/\downarrow}^2 = N_{\text{tot}}/2$ since we are on the side of a fringe ($p_{\uparrow} = p_{\downarrow} = 1/2$).

Now we take these two measurements and form the fringe asymmetry: $\mathcal{A} = (N_{\downarrow} - N_{\uparrow})/(N_{\downarrow} + N_{\uparrow})$. Combining the errors we have a new uncertainty: the numerator and denominator have the same variance $\delta\mathcal{A}_{n/d}^2 = N_{\text{tot}}/2$, and the quotient of that has variance

$$\begin{aligned}\delta\mathcal{A}^2 &= \left(\frac{N_{\downarrow} - N_{\uparrow}}{N_{\downarrow} + N_{\uparrow}}\right)^2 \left(\frac{N_{\text{tot}}/2}{(N_{\downarrow} - N_{\uparrow})^2} + \frac{N_{\text{tot}}/2}{(N_{\downarrow} + N_{\uparrow})^2}\right) \\ &= \frac{N_{\text{tot}}/2}{(N_{\downarrow} + N_{\uparrow})^2} + \frac{N_{\text{tot}}/2(N_{\downarrow} - N_{\uparrow})^2}{(N_{\downarrow} + N_{\uparrow})^4} \\ &= \frac{1}{2N_{\text{tot}}}\end{aligned}\tag{D.4}$$

where in the last line we used $N_{\downarrow} + N_{\uparrow} = N_{\text{tot}}$

Now we repeat this procedure r times, and we get $r/2$ asymmetry measurements. We average these measurements to get a better estimate of the spin asymmetry, and the sampling distribution of the mean⁵ is usually approximated by a normal distribution, so the quantum-projection-noise-limited uncertainty on an asymmetry measurement is $\delta\mathcal{A}_{\text{mean}}^2 = \delta\mathcal{A}^2/r$:

$$\delta\mathcal{A}_{\text{QPN}}^2 = \frac{1}{N_{\text{tot}}f_{\text{rep}}T_{\text{tot}}},\tag{D.5}$$

⁵For $N_{\text{tot}} \gg 1$

f_{rep} is the rep rate of the experiment, and T_{tot} is the total time the experiment was running, and $r = f_{\text{rep}} T_{\text{tot}}/2$.

On the side (zero-crossing) of a sinusoidally-varying asymmetry measurement, the slope (in units 1/s) is equal to ω ⁶, which we approximate as a line with zero offset: $y = \omega t$. If we are fitting a line to a point with errorbar σ' (and $\sigma' \ll 1$), then the line can be moved up or down σ' , so the fitted line can range from $y = \omega t$ to $y = \omega t \pm \sigma'$. The horizontal distance (in this case, the time shift) between the "true line" $y = \omega t$ and the "line inferred due to the errorbar" $y = \omega t \pm \sigma'$ is given by finding the point on the x -axis where $y = \omega t = \sigma'$, so $t = \sigma'/\omega$. In a cycle, the time for a single degree is $t_{\text{deg}} = 1/360f = T/360$, where f is the frequency and T is the period. The time uncertainty due to our errorbar divided by the time for a single degree is the uncertainty in degrees: $\sigma'/\omega/1/360f = \sigma'/2\pi f/1/360f = \sigma'$ (remembering that the 2π is radians and 360 is degrees). So our phase error, in radians on the side of a fringe, is the same as the errorbar: σ' . Now if we say that our frequency measurement is given by $(\phi_L - \phi_E)/2\pi t_R$, where ϕ_L is the late-time phase, ϕ_E is the early-time phase and t_R is the ramsey free evolution time, then the frequency uncertainty is $\delta f = \frac{\sqrt{(\sigma')^2 + (\sigma')^2}}{2\pi t_R} = \frac{\sqrt{2}\sigma'}{2\pi t_R}$. This is all for a fringe with perfect contrast, for a non-unity contrast C we can scale it $\delta f = \frac{\sqrt{2}\sigma'}{C2\pi t_R}$. Finally, to get from a frequency uncertainty to a d_e uncertainty, we divide by $-2|\mathcal{E}_{\text{eff}}|\text{sgn}(g_F)/h$ [2], where $|\mathcal{E}_{\text{eff}}|$ is the effective field experienced by the valence electron and $\text{sgn}(g_F)$ is the sign on the g -factor. Thus,

$$\delta d_e = \frac{\sqrt{2}\sigma' h}{2|\mathcal{E}_{\text{eff}}|C2\pi t_R} \quad (\text{D.6})$$

is the uncertainty in d_e for a single frequency measurement.

Finally, we will take the mean of all our frequency measurements, and the distribution of frequency measurements is again normal (we go to some effort to quantify this, see Chapter 4.) So our new uncertainty is just equation D.6 divided by the square root of the number of

⁶The side of the fringe $\cos(\omega t)$ is at $\omega t = \pi/2$, the derivative is $d/dt \cos(\omega t) = -\omega \sin(\omega t)$, plug in $\omega t = \pi/2$

measurements:

$$\delta d_e = \frac{\sqrt{2}\sigma' \hbar}{2|\mathcal{E}_{\text{eff}}|C2\pi t_R \sqrt{f_{\text{rep}} T_{\text{tot}}/2}}. \quad (\text{D.7})$$

Let's clean this up by taking $\hbar = h/2\pi$ and cleaning up the factors of 2:

$$\delta d_e = \frac{\sigma' \hbar}{|\mathcal{E}_{\text{eff}}|Ct_R \sqrt{f_{\text{rep}} T_{\text{tot}}}}. \quad (\text{D.8})$$

Now all that remains is plug in the quantum-projection noise limited errorbar on a single frequency measurement, $\sigma' = 1/\sqrt{N_{\text{tot}} f_{\text{rep}} T_{\text{tot}}}$. Finally, we have the quantum-projection-noise-limited expression for our uncertainty in d_e :

$$\delta d_e = \frac{\hbar}{|\mathcal{E}_{\text{eff}}|Ct_R f_{\text{rep}} T_{\text{tot}} \sqrt{N_{\text{tot}}}}. \quad (\text{D.9})$$

The expression above assumes that in each experimental cycle we read out an identical number of ions, N_{tot} . In reality, N varies in every cycle due to fluctuations in the spatial mode and power of the lasers that ablate, ionize, and dissociate, as well as other experimental noise sources. We avoid most of that noise by performing a differential measurement between the upper and lower doublets: in a single shot of the experiment we measure both the lower and upper Stark doublets, and we perform our measurements when the fringes for both are in phase, so the proportional experimental noise is effectively cancelled out. Details of this scheme are in [47] and [48].

Appendix E

All parity channel results

Below are all parity channel results for our 2022 dataset. Because we took data at three distinct values of f^0 , the results are disaggregated over those values then re-aggregated afterwards to give an overall average and error. All errorbars are adjusted for χ^2 . Before taking weighted means and adjusting for χ^2 , the value and error for each parity channel for each block was multiplied by 1000, so the units for contrast are milli “contrast units”, for decay rate the units are per millisecond, for frequency the units milli hertz, and for offset the units are milli offset units.

parameter	$f^0 < 80$			$80 < f^0 < 140$			$f^0 > 140$			aggregate over f^0					
	mean	σ	χ^2	mean/ σ	mean	σ	χ^2	mean/ σ	mean	σ	χ^2	mean/ σ	mean	σ	mean/ σ
Co	552.43	1.86	14.23	296.29	541.93	1.49	16.08	364.09	545.97	1.05	5.53	519.81	546.00	0.78	700.61
CoP	-2.01	1.86	14.23	1.08	-0.98	1.49	16.08	0.66	-1.74	1.05	5.53	1.66	-1.58	0.78	-2.03
CoB	-0.67	0.54	1.18	1.24	-0.22	0.41	1.24	0.54	-0.25	0.48	1.17	0.52	-0.34	0.27	-1.27
CoBP	-0.33	0.54	1.18	0.62	-0.02	0.41	1.24	0.06	0.07	0.48	1.17	0.15	-0.07	0.27	-0.26
CoR	-2.07	0.57	1.33	3.63	-1.62	0.41	1.24	3.92	-3.82	0.49	1.18	7.85	-2.42	0.28	-8.79
CoRP	-0.55	0.57	1.33	0.96	-0.07	0.41	1.24	0.16	0.35	0.49	1.18	0.71	-0.05	0.28	-0.18
CoI	-1.14	0.57	1.31	2.02	1.19	0.42	1.31	2.80	-1.57	0.53	1.39	2.98	-0.19	0.29	-0.67
CoIP	-0.92	0.57	1.31	1.63	0.05	0.42	1.31	0.12	-0.56	0.53	1.39	1.05	-0.37	0.29	-1.29
CoBR	-0.67	0.53	1.15	1.27	1.01	0.42	1.26	2.43	0.56	0.53	1.39	1.06	0.42	0.28	1.49
CoBRP	-0.12	0.53	1.15	0.22	0.11	0.42	1.26	0.26	-0.88	0.53	1.39	1.68	-0.23	0.28	-0.82
CoBI	-0.54	0.57	1.34	0.94	0.85	0.39	1.13	2.15	0.29	0.50	1.26	0.58	0.37	0.27	1.38
CoBIP	0.13	0.57	1.34	0.23	0.08	0.39	1.13	0.21	-0.32	0.50	1.26	0.63	-0.03	0.27	-0.10
CoRI	3.79	0.56	1.30	6.72	5.26	0.44	1.40	11.99	6.42	0.51	1.29	12.67	5.24	0.29	18.31
CoRIP	0.96	0.56	1.30	1.71	0.79	0.44	1.40	1.79	-0.58	0.51	1.29	1.14	0.40	0.29	1.41
CoBRI	0.15	0.55	1.25	0.27	-0.69	0.42	1.29	1.63	-1.11	0.54	1.43	2.07	-0.58	0.28	-2.05
CoBRIP	-0.97	0.55	1.25	1.75	-0.81	0.42	1.29	1.92	0.56	0.54	1.43	1.05	-0.47	0.28	-1.67
CoD	11.85	0.43	4.86	27.51	16.70	0.24	3.46	70.35	24.34	0.39	5.48	62.62	17.52	0.18	94.89
CoDP	0.27	0.43	4.86	0.63	-0.25	0.24	3.46	1.07	0.18	0.39	5.48	0.45	-0.06	0.18	-0.31
CoDB	0.03	0.22	1.24	0.13	-0.14	0.14	1.14	1.06	-0.08	0.19	1.25	0.41	-0.09	0.10	-0.88
CoDBP	-0.46	0.22	1.24	2.11	-0.03	0.14	1.14	0.19	0.16	0.19	1.25	0.87	-0.07	0.10	-0.66
CoDR	-1.21	0.21	1.19	5.65	-0.31	0.14	1.16	2.28	-0.91	0.18	1.23	4.94	-0.68	0.10	-6.98
CoDRP	-0.35	0.21	1.19	1.64	0.08	0.14	1.16	0.55	-0.02	0.18	1.23	0.10	-0.04	0.10	-0.44
CoDI	6.28	0.24	1.50	26.19	6.93	0.14	1.28	48.01	8.71	0.18	1.18	48.14	7.37	0.10	73.42
CoDIP	0.00	0.24	1.50	0.00	0.15	0.14	1.28	1.06	0.16	0.18	1.18	0.91	0.13	0.10	1.26
CoDBR	0.12	0.22	1.32	0.53	-0.32	0.13	1.09	2.36	-0.61	0.17	1.04	3.63	-0.33	0.09	-3.51
CoDBRP	-0.02	0.22	1.32	0.07	-0.01	0.13	1.09	0.07	-0.28	0.17	1.04	1.63	-0.09	0.09	-1.00
CoDBI	-0.01	0.20	1.09	0.05	-0.03	0.14	1.12	0.21	-0.20	0.19	1.27	1.08	-0.07	0.10	-0.72
CoDBIP	0.09	0.20	1.09	0.43	-0.09	0.14	1.12	0.64	-0.20	0.19	1.27	1.05	-0.08	0.10	-0.77
CoDRI	0.33	0.22	1.32	1.49	0.43	0.15	1.32	2.93	0.55	0.19	1.30	2.93	0.44	0.10	4.27
CoDRIP	-0.05	0.22	1.32	0.21	0.15	0.15	1.32	1.01	0.25	0.19	1.30	1.33	0.14	0.10	1.30
CoDBRI	-0.17	0.24	1.46	0.70	0.00	0.14	1.18	0.02	-0.13	0.19	1.28	0.69	-0.07	0.10	-0.67
CoDBRIP	-0.09	0.24	1.46	0.37	-0.01	0.14	1.18	0.09	0.16	0.19	1.28	0.83	0.02	0.10	0.24

Table E.1: All contrast parity channels, disaggregated over distinct f^0 values and aggregated over all 3 f^0 values.

parameter	$f^0 < 80$			$80 < f^0 < 140$			$f^0 > 140$			aggregate over f^0					
	mean	σ	χ^2	mean/ σ	mean	σ	χ^2	mean/ σ	mean	σ	mean/ σ	mean/ σ			
g	75.87	1.70	9.64	44.68	87.81	1.57	8.30	56.01	121.69	2.41	7.52	50.50	89.65	1.04	86.17
gP	-2.29	1.70	9.64	1.35	-0.97	1.57	8.30	0.62	0.26	2.41	7.52	0.11	-1.24	1.04	-1.19
gB	-1.20	0.59	1.16	2.03	0.57	0.59	1.18	0.96	0.37	1.00	1.29	0.37	-0.21	0.39	-0.55
gBP	-0.80	0.59	1.16	1.36	0.15	0.59	1.18	0.25	-0.38	1.00	1.29	0.38	-0.33	0.39	-0.87
gR	2.67	0.60	1.22	4.42	0.93	0.59	1.18	1.58	-11.29	1.00	1.27	11.34	-0.18	0.39	-0.47
gRP	-0.05	0.60	1.22	0.08	-0.59	0.59	1.18	0.99	0.13	1.00	1.27	0.13	-0.26	0.39	-0.66
gI	-0.29	0.58	1.12	0.49	1.74	0.61	1.24	2.87	2.83	0.99	1.27	2.85	1.00	0.39	2.59
gIP	-1.20	0.58	1.12	2.07	-0.12	0.61	1.24	0.19	-0.61	0.99	1.27	0.62	-0.68	0.39	-1.75
gBR	-1.42	0.55	1.02	2.58	-0.33	0.61	1.25	0.54	0.46	1.00	1.28	0.46	-0.73	0.38	-1.94
gBRP	0.48	0.55	1.02	0.86	0.44	0.61	1.25	0.72	-0.71	1.00	1.28	0.71	0.29	0.38	0.78
gBI	-0.44	0.62	1.30	0.70	0.36	0.58	1.14	0.61	-0.95	1.00	1.30	0.94	-0.16	0.39	-0.40
gBIP	-0.24	0.62	1.30	0.39	-0.05	0.58	1.14	0.09	-0.46	1.00	1.30	0.45	-0.19	0.39	-0.48
gRI	5.38	0.58	1.11	9.33	6.41	0.64	1.37	10.07	8.36	0.96	1.19	8.72	6.26	0.39	15.97
gRIP	0.60	0.58	1.11	1.04	0.12	0.64	1.37	0.19	-1.38	0.96	1.19	1.44	0.09	0.39	0.23
gBRI	0.69	0.57	1.10	1.20	-0.60	0.59	1.15	1.02	-0.72	0.98	1.25	0.73	-0.05	0.38	-0.13
gBRIP	-1.05	0.57	1.10	1.84	-1.41	0.59	1.15	2.41	0.91	0.98	1.25	0.93	-0.91	0.38	-2.40
gD	-2.12	0.43	1.78	4.95	-6.16	0.38	2.12	16.00	-6.25	0.59	1.71	10.68	-4.74	0.26	-18.48
gDP	-0.30	0.43	1.78	0.71	0.05	0.38	2.12	0.13	0.20	0.59	1.71	0.35	-0.05	0.26	-0.18
gDB	0.04	0.34	1.15	0.12	-0.02	0.29	1.18	0.06	0.03	0.48	1.15	0.07	0.01	0.20	0.05
gDBP	-0.07	0.34	1.15	0.21	-0.30	0.29	1.18	1.04	0.11	0.48	1.15	0.24	-0.15	0.20	-0.74
gDR	-1.13	0.37	1.31	3.07	-0.44	0.30	1.32	1.45	-0.39	0.48	1.15	0.81	-0.65	0.21	-3.11
gDRP	-0.20	0.37	1.31	0.54	-0.11	0.30	1.32	0.35	-0.05	0.48	1.15	0.09	-0.13	0.21	-0.61
gDI	2.59	0.37	1.34	6.99	4.14	0.30	1.32	13.64	4.97	0.49	1.21	10.06	3.79	0.21	18.02
gDIP	0.02	0.37	1.34	0.05	-0.42	0.30	1.32	1.40	0.12	0.49	1.21	0.24	-0.18	0.21	-0.85
gDBR	-0.04	0.36	1.29	0.12	-0.09	0.29	1.20	0.31	-0.36	0.47	1.10	0.77	-0.12	0.20	-0.61
gDBRP	-0.23	0.36	1.29	0.62	-0.08	0.29	1.20	0.28	-0.53	0.47	1.10	1.14	-0.21	0.20	-1.04
gDBI	0.37	0.35	1.20	1.06	0.06	0.29	1.19	0.20	-0.41	0.50	1.23	0.83	0.09	0.20	0.43
gDBIP	0.28	0.35	1.20	0.79	0.28	0.29	1.19	0.98	0.19	0.50	1.23	0.37	0.27	0.20	1.30
gDRI	-0.32	0.35	1.22	0.89	-0.58	0.30	1.33	1.90	-0.48	0.50	1.22	0.96	-0.47	0.21	-2.28
gDRIP	-0.29	0.35	1.22	0.81	0.72	0.30	1.33	2.35	0.57	0.50	1.22	1.15	0.34	0.21	1.64
gDBRI	-0.63	0.39	1.45	1.62	-0.43	0.28	1.15	1.52	-0.60	0.48	1.16	1.23	-0.52	0.21	-2.51
gDBRIP	-0.48	0.39	1.45	1.25	0.58	0.28	1.15	2.04	-0.05	0.48	1.16	0.11	0.17	0.21	0.83

Table E.2: All decay rate parity channels, disaggregated over distinct f^0 values and aggregated over all 3 f^0 values.

parameter	$f^0 < 80$			$80 < f^0 < 140$			$f^0 > 140$			aggregate over f^0					
	mean	σ	χ^2	mean/ σ	mean	σ	χ^2	mean/ σ	mean	σ	χ^2	mean/ σ	mean	σ	mean/ σ
f	76991.99	1.06	167.84	72780.41	105043.72	1.35	263.99	77654.72	151225.70	4.66	944.34	32485.90	89661.08	0.82	109252.06
fP	-0.17	1.06	167.84	0.16	-0.04	1.35	263.99	0.03	-3.18	4.66	944.34	0.68	-0.22	0.82	-0.26
fB	-0.11	0.14	3.05	0.79	0.00	0.15	3.07	0.01	0.04	0.25	2.76	0.14	-0.04	0.09	-0.47
fBP	0.05	0.14	3.05	0.34	0.06	0.15	3.07	0.39	0.30	0.25	2.76	1.21	0.09	0.09	0.95
fR	-1.62	0.13	2.44	12.68	-4.75	0.14	2.73	34.50	-15.01	0.28	3.51	52.97	-4.31	0.09	-47.77
fRP	0.22	0.13	2.44	1.71	0.02	0.14	2.73	0.17	0.06	0.28	3.51	0.21	0.12	0.09	1.34
fI	-0.66	0.10	1.50	6.61	-0.95	0.12	2.02	8.04	-1.17	0.23	2.21	5.21	-0.82	0.07	-11.23
fIP	-0.07	0.10	1.50	0.66	0.00	0.12	2.02	0.01	-0.24	0.23	2.21	1.07	-0.06	0.07	-0.84
fBR	213.28	0.09	1.30	2294.25	210.86	0.10	1.56	2029.86	209.14	0.18	1.46	1141.35	211.83	0.06	3378.08
fBRP	0.11	0.09	1.30	1.14	-0.15	0.10	1.56	1.48	-0.41	0.18	1.46	2.23	-0.06	0.06	-0.88
fBI	-0.27	0.09	1.14	3.09	-0.54	0.09	1.28	5.68	-0.62	0.17	1.25	3.66	-0.43	0.06	-7.24
fBIP	-0.06	0.09	1.14	0.71	0.04	0.09	1.28	0.47	-0.05	0.17	1.25	0.27	-0.01	0.06	-0.25
fRI	-1.72	0.09	1.23	19.05	-2.51	0.10	1.45	25.08	-3.35	0.18	1.39	18.77	-2.23	0.06	-35.54
fRIP	0.07	0.09	1.23	0.77	0.00	0.10	1.45	0.02	-0.09	0.18	1.39	0.49	0.02	0.06	0.37
fBRI	0.10	0.09	1.31	1.10	-0.05	0.09	1.25	0.51	-0.10	0.16	1.12	0.63	0.01	0.06	0.13
fBRIP	-0.10	0.09	1.31	1.10	-0.03	0.09	1.25	0.31	-0.19	0.16	1.12	1.16	-0.08	0.06	-1.39
fD	-171.27	0.04	1.09	4394.16	-230.12	0.03	1.06	7089.88	-327.98	0.06	1.18	5507.94	-225.35	0.02	-10113.08
fDP	0.06	0.04	1.09	1.66	-0.04	0.03	1.06	1.11	-0.07	0.06	1.18	1.25	-0.01	0.02	-0.59
fDB	16.07	0.04	1.03	423.35	16.12	0.03	1.08	491.54	16.11	0.06	1.06	286.34	16.10	0.02	722.65
fDBP	-0.13	0.04	1.03	3.39	0.01	0.03	1.08	0.19	0.01	0.06	1.06	0.20	-0.03	0.02	-1.50
fDR	0.05	0.04	1.18	1.15	0.08	0.03	1.03	2.49	0.12	0.05	0.87	2.25	0.08	0.02	3.64
fDRP	-0.01	0.04	1.18	0.34	0.00	0.03	1.03	0.03	-0.11	0.05	0.87	2.08	-0.02	0.02	-1.09
fDI	-0.01	0.04	1.03	0.18	-0.09	0.03	1.08	2.72	0.01	0.06	1.07	0.23	-0.05	0.02	-2.31
fDIP	0.02	0.04	1.03	0.57	-0.02	0.03	1.08	0.63	-0.04	0.06	1.07	0.70	-0.01	0.02	-0.46
fDBR	-0.37	0.04	1.16	9.13	-0.45	0.03	0.99	14.41	-0.35	0.06	1.03	6.21	-0.41	0.02	-18.46
fDBRP	-0.05	0.04	1.16	1.12	0.05	0.03	0.99	1.47	0.06	0.06	1.03	1.14	0.02	0.02	0.91
fDBI	-0.01	0.04	1.04	0.19	0.05	0.03	1.06	1.58	0.05	0.05	0.96	0.85	0.03	0.02	1.50
fDBIP	-0.05	0.04	1.04	1.22	-0.03	0.03	1.06	0.98	0.05	0.05	0.96	0.91	-0.02	0.02	-0.96
fDRI	0.04	0.04	1.11	0.93	0.00	0.03	0.99	0.15	0.06	0.06	1.11	1.05	0.02	0.02	0.93
fDRIP	-0.03	0.04	1.11	0.77	0.03	0.03	0.99	0.94	-0.01	0.06	1.11	0.18	0.01	0.02	0.26
fDBRI	-0.01	0.04	1.05	0.16	0.02	0.03	1.11	0.52	-0.01	0.06	1.02	0.27	0.01	0.02	0.29
fDBRIP	0.04	0.04	1.05	1.08	0.01	0.03	1.11	0.25	-0.03	0.06	1.02	0.47	0.01	0.02	0.62

Table E.3: All frequency parity channels, disaggregated over distinct f^0 values and aggregated over all 3 f^0 values.

parameter	$f^0 < 80$				$80 < f^0 < 140$				$f^0 > 140$				aggregate over f^0			
	mean	σ	χ^2	mean/ σ	mean	σ	χ^2	mean/ σ	mean	σ	χ^2	mean/ σ	mean	σ	mean/ σ	
p	-804.15	167.09	39670.19	4.81	-3045.98	0.64	1.13	4790.86	-3017.21	0.80	1.18	3749.48	-3034.73	0.50	-6072.46	
pP	-195.53	167.09	39670.19	1.17	-0.85	0.64	1.13	1.33	0.93	0.80	1.18	1.15	-0.16	0.50	-0.31	
pB	0.85	0.88	1.09	0.97	-0.95	0.66	1.22	1.45	-3.11	0.83	1.25	3.75	-1.11	0.45	-2.49	
pBP	0.59	0.88	1.09	0.67	-0.39	0.66	1.22	0.59	0.31	0.83	1.25	0.38	0.06	0.45	0.14	
pR	-1.07	0.93	1.22	1.16	0.36	0.61	1.05	0.58	0.04	0.77	1.09	0.05	-0.04	0.43	-0.09	
pRP	-1.29	0.93	1.22	1.39	0.25	0.61	1.05	0.41	-0.15	0.77	1.09	0.19	-0.19	0.43	-0.46	
pI	0.26	0.81	0.92	0.32	-0.22	0.61	1.04	0.37	-1.00	0.78	1.12	1.27	-0.31	0.41	-0.76	
pIP	1.12	0.81	0.92	1.39	-0.52	0.61	1.04	0.85	-0.27	0.78	1.12	0.34	-0.02	0.41	-0.06	
pBR	5.23	0.87	1.07	6.02	5.97	0.67	1.26	8.90	7.27	0.80	1.15	9.13	6.18	0.44	13.96	
pBRP	0.94	0.87	1.07	1.09	1.15	0.67	1.26	1.72	1.74	0.80	1.15	2.19	1.28	0.44	2.88	
pBI	-2.00	0.91	1.17	2.20	-1.50	0.65	1.18	2.30	-4.55	0.78	1.12	5.80	-2.58	0.44	-5.89	
pBIP	-0.20	0.91	1.17	0.22	0.77	0.65	1.18	1.19	-0.39	0.78	1.12	0.50	0.18	0.44	0.41	
pRI	-0.76	0.88	1.09	0.87	-0.85	0.61	1.03	1.40	0.21	0.74	1.00	0.28	-0.50	0.42	-1.20	
pRIP	-1.26	0.88	1.09	1.44	-0.12	0.61	1.03	0.20	0.54	0.74	1.00	0.73	-0.17	0.42	-0.40	
pBRI	0.18	0.88	1.08	0.21	2.08	0.62	1.08	3.35	3.24	0.77	1.08	4.21	1.99	0.42	4.70	
pBRIP	1.06	0.88	1.08	1.21	0.74	0.62	1.08	1.20	1.48	0.77	1.08	1.92	1.04	0.42	2.45	
pD	-8.80	0.26	1.06	33.85	-15.66	0.20	1.12	79.00	-11.45	0.26	1.04	43.70	-12.66	0.14	-93.53	
pDP	-0.27	0.26	1.06	1.02	-0.07	0.20	1.12	0.37	-0.28	0.26	1.04	1.06	-0.18	0.14	-1.34	
pDB	-0.26	0.27	1.14	0.97	0.18	0.20	1.14	0.90	-0.20	0.27	1.11	0.73	-0.03	0.14	-0.25	
pDBP	0.24	0.27	1.14	0.88	0.27	0.20	1.14	1.38	-0.40	0.27	1.11	1.48	0.09	0.14	0.63	
pDR	-0.01	0.27	1.15	0.03	-0.28	0.19	1.02	1.51	0.10	0.27	1.14	0.36	-0.12	0.13	-0.88	
pDRP	0.04	0.27	1.15	0.15	0.05	0.19	1.02	0.27	0.55	0.27	1.14	2.02	0.17	0.13	1.28	
pDI	-0.19	0.26	1.03	0.74	0.88	0.20	1.18	4.31	0.17	0.28	1.19	0.61	0.41	0.14	2.95	
pDIP	0.16	0.26	1.03	0.63	-0.24	0.20	1.18	1.17	0.51	0.28	1.19	1.81	0.05	0.14	0.40	
pDBR	-0.59	0.27	1.14	2.17	-0.57	0.19	1.02	3.04	-1.23	0.27	1.11	4.52	-0.74	0.13	-5.49	
pDBRP	0.44	0.27	1.14	1.62	-0.34	0.19	1.02	1.77	-0.16	0.27	1.11	0.59	-0.10	0.13	-0.75	
pDBI	0.12	0.25	0.96	0.47	0.16	0.21	1.21	0.80	0.48	0.26	1.03	1.83	0.24	0.14	1.73	
pDBIP	-1.05	0.25	0.96	4.21	0.13	0.21	1.21	0.62	1.07	0.26	1.03	4.11	0.04	0.14	0.27	
pDRI	-0.21	0.27	1.12	0.77	0.12	0.20	1.19	0.58	-0.41	0.27	1.09	1.52	-0.10	0.14	-0.76	
pDRIP	0.03	0.27	1.12	0.10	-0.03	0.20	1.19	0.16	0.24	0.27	1.09	0.90	0.06	0.14	0.41	
pDBRI	0.03	0.26	1.09	0.12	-0.75	0.20	1.14	3.75	-0.61	0.26	1.04	2.31	-0.50	0.14	-3.70	
pDBRIP	-0.30	0.26	1.09	1.12	0.03	0.20	1.14	0.16	0.55	0.26	1.04	2.10	0.08	0.14	0.60	

Table E.4: All phase parity channels, disaggregated over distinct f^0 values and aggregated over all 3 f^0 values.

parameter	$f^0 < 80$			$80 < f^0 < 140$			$f^0 > 140$			aggregate over f^0					
	mean	σ	χ^2	mean/ σ	mean	σ	χ^2	mean/ σ	mean	σ	χ^2	mean/ σ	mean	σ	mean/ σ
O	26.07	0.34	1.86	76.65	26.11	0.23	1.61	112.75	24.39	0.30	1.76	81.85	25.61	0.16	159.23
OP	0.68	0.34	1.86	2.00	-0.33	0.23	1.61	1.45	0.00	0.30	1.76	0.00	-0.01	0.16	-0.06
OB	-0.23	0.31	1.54	0.75	-0.10	0.23	1.52	0.45	-0.10	0.26	1.39	0.38	-0.13	0.15	-0.87
OBP	-0.14	0.31	1.54	0.45	-0.27	0.23	1.52	1.19	-0.35	0.26	1.39	1.32	-0.27	0.15	-1.77
OR	0.23	0.31	1.55	0.73	0.04	0.22	1.39	0.19	-0.03	0.24	1.15	0.14	0.06	0.14	0.39
ORP	-0.16	0.31	1.55	0.51	-0.19	0.22	1.39	0.90	0.22	0.24	1.15	0.92	-0.04	0.14	-0.25
OI	0.15	0.29	1.36	0.52	-0.11	0.21	1.28	0.52	-0.23	0.28	1.52	0.83	-0.08	0.15	-0.53
OIP	-0.54	0.29	1.36	1.86	0.05	0.21	1.28	0.25	0.17	0.28	1.52	0.62	-0.07	0.15	-0.45
OBR	0.18	0.28	1.22	0.66	0.19	0.22	1.50	0.85	0.67	0.25	1.24	2.67	0.34	0.14	2.41
OBRP	-0.27	0.28	1.22	0.99	0.20	0.22	1.50	0.91	-0.05	0.25	1.24	0.21	-0.00	0.14	-0.02
OBI	-0.17	0.28	1.25	0.60	0.30	0.22	1.48	1.36	-0.50	0.27	1.42	1.87	-0.06	0.15	-0.41
OBIP	0.46	0.28	1.25	1.66	-0.12	0.22	1.48	0.52	-0.12	0.27	1.42	0.44	0.04	0.15	0.25
ORI	-0.20	0.32	1.67	0.61	-0.22	0.23	1.53	0.96	-0.28	0.27	1.43	1.06	-0.23	0.15	-1.53
ORIP	-0.01	0.32	1.67	0.03	0.05	0.23	1.53	0.22	-0.06	0.27	1.43	0.22	0.00	0.15	0.00
OBRI	0.02	0.28	1.27	0.09	-0.61	0.22	1.45	2.78	0.03	0.26	1.36	0.10	-0.25	0.14	-1.71
OBRIP	-0.27	0.28	1.27	0.95	-0.50	0.22	1.45	2.28	-0.34	0.26	1.36	1.28	-0.39	0.14	-2.71
OD	10.38	0.15	2.34	67.88	10.86	0.11	2.69	98.81	1.39	0.34	15.01	4.10	10.10	0.09	117.66
ODP	0.12	0.15	2.34	0.82	-0.09	0.11	2.69	0.80	-0.17	0.34	15.01	0.51	-0.03	0.09	-0.31
ODB	0.13	0.11	1.12	1.27	-0.10	0.07	1.17	1.37	-0.15	0.10	1.28	1.56	-0.06	0.05	-1.25
ODBP	0.04	0.11	1.12	0.42	0.03	0.07	1.17	0.39	-0.05	0.10	1.28	0.51	0.01	0.05	0.23
ODR	0.03	0.11	1.13	0.26	-0.03	0.08	1.34	0.40	0.00	0.09	1.17	0.03	-0.01	0.05	-0.12
ODRP	0.01	0.11	1.13	0.13	0.04	0.08	1.34	0.55	-0.04	0.09	1.17	0.43	0.01	0.05	0.11
ODI	0.05	0.11	1.20	0.45	-0.01	0.07	1.14	0.09	0.15	0.10	1.24	1.50	0.04	0.05	0.87
ODIP	0.66	0.11	1.20	6.08	0.50	0.07	1.14	7.01	0.56	0.10	1.24	5.75	0.55	0.05	10.81
ODBR	0.28	0.12	1.41	2.35	0.28	0.07	1.17	3.90	0.36	0.10	1.20	3.76	0.30	0.05	5.83
ODBRP	0.06	0.12	1.41	0.48	-0.01	0.07	1.17	0.19	0.03	0.10	1.20	0.36	0.01	0.05	0.27
ODBI	-0.04	0.11	1.13	0.36	-0.07	0.07	1.15	0.94	-0.02	0.09	1.16	0.21	-0.05	0.05	-0.99
ODBIP	-0.05	0.11	1.13	0.43	0.10	0.07	1.15	1.39	-0.03	0.09	1.16	0.30	0.03	0.05	0.62
ODRI	-0.18	0.11	1.19	1.65	-0.58	0.08	1.26	7.72	-0.46	0.10	1.24	4.70	-0.45	0.05	-8.23
ODRIP	-0.18	0.11	1.19	1.63	-0.13	0.08	1.26	1.76	0.08	0.10	1.24	0.79	-0.08	0.05	-1.48
ODBRI	-0.08	0.11	1.26	0.68	-0.01	0.07	1.21	0.15	0.05	0.10	1.20	0.51	-0.01	0.05	-0.19
ODBRIP	-0.07	0.11	1.26	0.63	0.10	0.07	1.21	1.36	0.07	0.10	1.20	0.72	0.06	0.05	1.10

Table E.5: All offset parity channels, disaggregated over distinct f^0 values and aggregated over all 3 f^0 values.

parameter	$f^0 < 80$			$80 < f^0 < 140$			$f^0 > 140$			aggregate over f^0			
	mean	σ	χ^2	mean/ σ	χ^2	σ	mean	χ^2	σ	mean/ σ	mean	σ	mean/ σ
f	76991.994	1.058	167.844	72780.413	1.353	263.993	77654.719	151225.696	4.655	944.337	32485.897	89611.303	0.820
fP	-0.167	1.058	167.844	0.158	1.353	263.993	0.031	-3.183	4.655	944.337	0.684	-0.215	0.820
fB	-0.113	0.143	3.054	0.795	0.146	3.067	0.008	0.036	0.252	2.759	0.143	-0.044	0.095
fBP	0.048	0.143	3.054	0.339	0.056	3.067	0.387	1.303	0.252	2.759	1.205	0.087	0.095
fR	-1.617	0.127	2.436	12.685	0.138	2.734	34.497	-15.005	0.283	3.505	52.970	-4.229	0.089
fRP	0.218	0.127	2.436	1.707	0.024	2.734	0.174	0.059	0.283	3.505	0.210	0.122	0.089
fI	-0.660	0.100	1.496	6.606	0.118	2.022	8.035	-1.173	0.225	2.206	5.210	-0.822	0.072
fIP	-0.066	0.100	1.496	0.663	0.118	2.022	0.012	-0.241	0.225	2.206	1.073	-0.060	0.072
fBR	213.275	0.093	1.296	2294.245	0.104	1.555	2029.864	209.142	0.183	1.463	1141.354	211.819	0.065
fBRP	0.106	0.093	1.296	1.142	0.154	1.555	1.479	-0.409	0.183	1.463	2.230	-0.060	0.065
fBI	-0.269	0.087	1.142	3.086	0.094	1.281	5.678	-0.618	0.169	1.249	3.661	-0.420	0.060
fBIP	-0.062	0.087	1.142	0.708	0.094	1.281	0.475	-0.046	0.169	1.249	0.270	-0.017	0.060
fRI	-1.723	0.090	1.226	19.054	0.100	1.446	25.081	-3.346	0.178	1.386	18.774	-2.232	0.063
fRIP	0.070	0.090	1.226	0.769	0.100	1.446	0.016	-0.088	0.178	1.386	0.495	0.022	0.063
fBRI	0.103	0.093	1.305	1.099	0.093	1.249	0.505	-0.101	0.161	1.122	0.628	0.010	0.061
fBRIP	-0.102	0.093	1.305	1.098	0.093	1.249	0.307	-0.186	0.161	1.122	1.160	-0.083	0.061
fD	-171.268	0.039	1.089	4394.159	0.032	1.057	7089.878	-327.977	0.060	1.182	5507.936	-224.100	0.023
fDP	0.065	0.039	1.089	1.656	0.036	1.057	1.110	-0.074	0.060	1.182	1.250	-0.007	0.023
fDB	16.072	0.038	1.029	423.351	0.033	1.079	491.535	16.111	0.056	1.056	286.337	16.102	0.023
fDBP	-0.129	0.038	1.029	3.387	0.006	1.079	0.191	0.011	0.056	1.056	0.201	-0.042	0.023
fDR	0.047	0.041	1.179	1.153	0.080	1.033	2.487	0.115	0.051	0.872	2.252	0.077	0.023
fDRP	-0.014	0.041	1.179	0.335	0.001	1.033	0.033	-0.107	0.051	0.872	2.083	-0.026	0.023
fDI	-0.007	0.038	1.029	0.175	0.089	1.083	2.719	0.013	0.057	1.067	0.229	-0.043	0.023
fDIP	0.021	0.038	1.029	0.565	0.021	1.083	0.627	-0.039	0.057	1.067	0.696	-0.009	0.023
fDBR	-0.369	0.040	1.163	9.129	0.453	0.990	14.405	-0.346	0.056	1.033	6.214	-0.409	0.022
fDBRP	-0.045	0.040	1.163	1.119	0.046	0.990	1.468	0.063	0.056	1.033	1.138	0.020	0.022
fDBI	-0.007	0.038	1.041	0.189	0.052	1.062	1.585	0.046	0.054	0.963	0.854	0.030	0.023
fDBIP	-0.047	0.039	1.114	1.221	0.032	1.062	0.984	0.049	0.054	0.963	0.909	-0.023	0.023
fDRI	0.037	0.039	1.114	0.930	0.005	0.985	0.153	0.061	0.058	1.111	1.050	0.019	0.022
fDRIP	-0.030	0.039	1.114	0.774	0.029	0.985	0.940	-0.011	0.058	1.111	0.184	0.004	0.022
fDBRI	-0.006	0.038	1.052	0.164	0.017	1.111	0.518	-0.015	0.055	1.021	0.267	0.003	0.023
fDBRIP	0.041	0.038	1.052	1.084	0.008	1.111	0.251	-0.026	0.055	1.021	0.472	0.014	0.023

Table E.6: All frequency parity channels (but with 3 digits after the decimal), disaggregated over distinct f^0 values and aggregated over all 3 f^0 values.