

Coherent and Incoherent Quantum Feedback in an Atom–Cavity System

by

Julian Choate Wolf

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Dan M. Stamper-Kurn, Chair

Professor Holger Müller

Professor Boubacar Kanté

Summer 2023

Coherent and Incoherent Quantum Feedback in an Atom–Cavity System

Copyright 2023
by
Julian Choate Wolf

Abstract

Coherent and Incoherent Quantum Feedback in an Atom–Cavity System

by

Julian Choate Wolf

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Dan M. Stamper-Kurn, Chair

Feedback control allows a wide range of systems to be stabilized to out-of-equilibrium states. In quantum systems, feedback control takes two forms: incoherent control, in which projective measurements are made of the system state and the measurement results are used to inform changes to the system Hamiltonian, bringing the system towards its desired state; and coherent control, in which the system is allowed to interact coherently with an auxiliary quantum system, such that Hamiltonian dynamics coherently drive the system of interest into its desired state.

An ultracold atomic gas coupled to a high-finesse optical cavity offers a convenient testbed for both of these forms of quantum feedback. Light escaping from the cavity mode carries information about the state of the atomic sample; this information can be processed externally and used to inform changes to external magnetic fields, to trapping parameters, and to the amplitude or frequency of light used to pump the cavity, effecting changes to the system’s Hamiltonian and driving the atoms towards a desired state. Meanwhile, the atomic sample continually exchanges information with the cavity field through coherent interactions; by treating the atomic sample as the system of interest and the light field as an auxiliary controller, a Hamiltonian can be engineered such that these coherent interactions drive the atomic sample towards a desired state without the need for any external control.

In this dissertation, I will re-introduce the elder of the two atom–cavity experiments currently active in the Stamper-Kurn group at Berkeley—dubbed the E3 apparatus—and will describe how it can be used to examine quantum feedback. I will discuss two instances of this, in particular: a coherent quantum feedback system, in which the energy of the collective atomic spin is autonomously stabilized to a set point conditioned on the detuning of the pump light from cavity resonance; and an incoherent feedback system, in which light escaping from the cavity offers a real-time measurement of the number of atoms present in the cavity during evaporative cooling, which is then used to stabilize to a desired atom number. Along the

way, I will discuss other interesting findings that have been made along the way, as well as techniques that we've found particularly useful.

Contents

Contents	i
List of Figures	iv
List of Tables	vi
Acronyms and initialisms	1
1 Background	3
1.1 Ultracold atomic gasses	4
1.2 High-finesse optical cavities and cavity QED	5
1.3 Simple atom–cavity interactions	7
1.4 Feedback in quantum systems	9
1.4.1 Incoherent, or measurement-based, feedback	9
1.4.2 Coherent, or Hamiltonian, feedback	10
1.5 Definitions and conventions	11
1.6 A note on the anthropomorphization of scientific instruments	12
2 The E3 apparatus at Berkeley	14
2.1 Trapping, cooling, and preparation of the atomic sample	15
2.1.1 Magnetic transport using the atom chip conveyor	16
2.1.2 Final RF evaporation and loading into the cavity ODT	17
2.2 Intricacies of the science cavity	18
2.2.1 Heating of the cavity by the atom chip wires	19
2.2.2 Chip quakes: mechanical bistability of the atom chip	20
2.2.3 Controlling the cavity pump polarizations with LCVRs	22
2.3 Collecting information from the system	22
2.3.1 Heterodyne detection	23
2.3.2 Single photon counting	25
2.3.3 Optical imaging of the atomic cloud	25
2.4 Locking and stabilization of the E3 lasers	27
2.4.1 Locking to pump–probe spectroscopy of rubidium	27
2.4.2 Actively stabilized reference cavity	30

2.4.3	Using a wavemeter as a broadband relative reference	31
2.5	Significant experimental setbacks	31
2.5.1	The flood	32
2.5.2	The global pandemic, and replacing the rubidium source	34
2.5.3	The labor strike	37
3	Spin optodynamics and coherent quantum feedback	38
3.1	Dispersive interactions between a collective spin and a cavity	38
3.2	Nondestructive readout of atomic spin states	40
3.2.1	Simultaneous readout of both circular cavity modes using linearly polarized pump light	42
3.3	Autonomous stabilization of a cavity–spin system	43
3.3.1	Building the spin stabilization Hamiltonian	44
3.3.2	Analytic model for the autonomous stabilization system	46
3.3.3	Preparation of the collective spin ensemble	49
3.3.4	Experiments measuring the autonomous stabilization of the collective spin	51
3.3.5	Conclusions	54
3.4	Decoherence of the collective spin	55
3.4.1	Measuring and characterizing decoherence	56
3.4.2	Methods for improving spin coherence	57
3.5	High-field effects: the quadratic Zeeman shift	62
4	Real-time tracking and control of atom loss	66
4.1	Dispersive readout of the instantaneous atom number	67
4.2	Examining the statistics of evaporative cooling	69
4.2.1	Nonlinearity of the evaporative cooling process	74
4.2.2	Stochastic character of the evaporative cooling process	77
4.2.3	Conclusions	78
4.3	Preparing ensembles with deterministic atom numbers	78
4.3.1	Minimizing shot-to-shot atom number variation	79
4.3.2	Limitations on precise atom number control	81
5	Closing remarks	86
5.1	Summary	86
5.2	The future of E3	87
5.2.1	Simultaneous feedback stabilization of atom number and temperature	87
5.2.2	Cavity-assisted molecular photoassociation	87
	Bibliography	89
A	Calibration of magnetic bias fields	101

B Measuring the cavity linewidth	106
C Machine learning optimization of the experimental sequence	110
D Data management for high-throughput experiments	113

List of Figures

1.1	Energy level diagram of ^{87}Rb D_2 transition	6
1.2	Schematic of measurement-based quantum feedback	9
1.3	Schematic of coherent quantum feedback	10
2.1	Schematic of the E3 apparatus	16
2.2	Final atom numbers and temperatures after different RF evaporation	18
2.3	Controllable position-dependent atom-cavity coupling	19
2.4	Observation of chip quakes using cavity transmission	21
2.5	Typical heterodyne spectrum	24
2.6	Typical absorption image of the atomic sample	26
2.7	Measuring the temperature of the atomic sample using time-of-flight imaging	27
2.8	Generic pump-probe spectroscopy setup	28
2.9	Saturation Absorption Spectroscopy (SAS) locking setup	29
2.10	Modulation Transfer Spectroscopy (MTS) locking setup	29
2.11	Racks housing analog and digital control breakouts during the flood	33
2.12	Analog isolator board corroded by flood water	34
2.13	Temperature readings from the successful bakeout	36
3.1	Scalar, vector, and tensor coupling coefficients	39
3.2	Nondestructively measuring the vector shift to cavity resonance due to the collective atomic spin	41
3.3	Measuring the vector shift using single sweeps and linearly polarized light	43
3.4	Experimental setup for autonomous spin stabilization	44
3.5	Schematic of the frequency landscape for spin stabilization	45
3.6	Fitting to the spin stabilizations response curves	48
3.7	Preparation of the collective spin using RF drives	50
3.8	DC Response of the autonomous stabilization system	52
3.9	Sensitivity of the autonomous stabilization system to input noise	53
3.10	Response of the autonomous stabilization system to impulses	54
3.11	Measurement of the damping rate of the autonomous stabilization system	55
3.12	Spin response to modulations outside of the system bandwidth	56
3.13	Spin echo measurement of the collective spin coherence	58
3.14	Decoherence of the collective spin as observed in the Larmor sidebands	59

3.15	Comparison of spin decoherence with and without pump light	60
3.16	Measuring spin decoherence at different ODT polarizations	61
3.17	Tuning the ODT polarization to minimize spin decoherence	62
3.18	Reviving the collective spin using RF drives	63
3.19	Measured spin coherence beat frequencies	64
3.20	Qualitative differences between beating in the $f = 1$ and $f = 2$ manifolds	65
4.1	Dispersive readout of the instantaneous atom number	69
4.2	Sensitivity of cavity output quadratures	70
4.3	Watching evaporative atom losses in real time	72
4.4	Measuring correlated losses during evaporative cooling	74
4.5	Investigating the linearity of evaporative atom loss	75
4.6	Unexplained variance in atom atom number after evaporation	77
4.7	Narrowing the shot-to-shot distribution of atom numbers in the final sample	80
4.8	Dependence of atom number variance on atom-cavity coupling	83
4.9	Reaching sub-single-atom shot-to-shot number variance	84
A.1	Spectrograms showing spin precession	102
A.2	Example field calibration using cavity modulation	103
A.3	Reliable horizontal field calibrations using RF spectroscopy	104
A.4	Reliable vertical field calibration using RF spectroscopy	105
B.1	Failed attempt at measuring the cavity lifetime via ringdown	107
B.2	Measuring the cavity lifetime via population and extinction	109
C.1	Example Cicero control sequence	110
C.2	Example machine learning optimization of the experimental sequence	112

List of Tables

1.1	Parameter definitions and values.	12
2.1	Laser models, wavelengths, and locking schemes	31

Acknowledgments

I first became aware of physics as a concept in high school at Berkeley High, when I was required to take a class on the subject during 12th grade. The class was taught by Stephen Salser, who not only introduced me to the ideas of physics, but made me aware that it was something one could spend their life studying, and began to lay the way for my active interest in doing so. It was as a direct result of this class, and, in particular, of Mr. Salser's exceptional instruction and contagious interest in the subject, that I entered university with a declared physics major.

The first year of university at McGill was a slog, working through a heavy course load of prerequisites without any concept of whether research might be something I was interested in, let alone what it might look like in practice. My interest was eventually sparked, during my second year of studies, by Prof. Andreas Warburton, who was leading a class on experimental methods. Through conversations with Prof. Warburton, I became inspired to pursue research in experimental high-energy physics. I spend the summers after my second and third years of university as a part of the ATLAS group at LBNL, working under Dr. Sasha Pranko and Dr. Maurice Garcia-Sciveres, where I got my first feel for what experimental physics might feel like. I learned quickly that high-energy physics wasn't for me; all the same, this introduction to research science was important to my development, and I owe my early interest in pursuing experimental research to Prof. Warburton's enthusiastic mentorship.

During my final year at McGill and the summer thereafter, I was involved in biophysics and polymer physics research in the group of Prof. Sabrina Leslie. It was at Leslie Lab that I got my first real physics experience. My honors thesis surrounded a small side project of the group: myself and a graduate student mentor, Daniel Berard, along with two graduate student collaborators from other universities. The scale of the project meant that I got to play a significant role in all of its phases, gaining experience in design, planning, data collection, and analysis, as well as helping to put together the resulting paper. Looking back, this level of involvement would have been unlikely had I found myself working in an AMO lab, and would have been impossible had I remained in high-energy physics. Throughout the year that I spend working in her lab, Sabrina took on a very active mentorship role, walking me through many of the early steps of research, scientific communication, teaching and mentoring others, and productive scientific thinking in general. It is also through Sabrina that I eventually connected with Dan, under whom she had completed her Ph.D.

My interest in AMO physics in general (and, as it happens, in working with Dan in particular) was spurred by conversations with Prof. Jack Sankey. Jack taught my final undergraduate class in quantum physics. Assuming fairly basic background knowledge, the class explored black holes, atomic interactions, the Aharonov–Bohm effect, and other Really Neat Things. Hearing about Jack's research in quantum optics and optomechanics introduced me to the prospect of science that both could be built up from first principles (like my work with ATLAS) and existed from top to bottom within a single room (like my work in Leslie Lab). Jack's guidance, both in helping me come to the realization that I was inter-

ested in pursuing tabletop quantum science and in helping figure out how best to make this a reality, is in large part what led me to UC Berkeley.

By Jack's suggestion, and with Sabrina's introduction, I eventually made contact with Prof. Dan M. Stamper-Kurn and expressed my interest in working in his Ultracold Atomic Physics group as a post-baccalaureate researcher. This transitioned, the following year, to research as a graduate student; six years later, we find ourselves transitioning from the past to the past participle. Dan has taught me how to be a physicist. His mentorship and guidance have often been frustrating, but have led me to consistently push myself and to think in ways that would not have come naturally to me. Grad school in Dan's group has been extremely difficult, but also incredibly fun.

Alongside Dan, I've received important mentorship over the past six years from a number of close colleagues. Dr. Jonathan Kohler was the senior grad student on E3 before I joined, and he was responsible for teaching me how to run an AMO lab—and how to deal with many of the frustrations that come along with doing so. Dr. Johannes Zeiher joined the experiment during my second year; he was responsible for showing me how to bring difficult projects to completion. Dr. Josh Isaacs worked with me on E3 for nearly three years, spanning the global COVID-19 pandemic and the ensuing rebuild of the apparatus. It was in large part through Josh's mentorship that I learned how to make it through grad school in one piece, and how to be a real-life scientist.

Conversations with other students in the Ultracold group have been indispensable in growing my understanding of quantum science: Justin Gerber, Scott Eustice, Tsz-Him (Zephy) Leung, and Olive Eilbott have each taught me a great deal. I've had the benefit of working closely with Olive in E3 over the past two years; they have now taken on the role of senior grad student, and are pushing the experiment in exciting new directions.

Grad school is made up of much more than just physics. I owe my success over the past years in large part to the expansive backcountry of the American west, where I've often escaped, explored, cleared my mind; to the Berkeley farmers' markets, where I've found deep-rooted community and life-long friends; to the local trails and climbing gyms, where I've kept myself active enough during the evenings to be able to consistently find joy in long days of windowless knob-turning.

Likewise, a dissertation is made up of more than words typed into a computer. I owe my success over the past months to comfort while writing found in simple spaces: my house in the hills, and the balconies, kitchens, and basements of good friends; local public libraries, particularly the Claremont branch; many cafes; the Eastern-most table at Schmidt's, by the window.

The work reported in this dissertation was supported by the National Science Foundation Quantum Leap Challenge Institutes program, the National Science Foundation, the Air Force Office of Scientific Research, the US Army Research Office through the Multidisciplinary University Research Initiative program, and others.

Acronyms and initialisms

AMO	atomic, molecular, and optical 28, 31, 106, 111
AOM	acousto-optical modulator 19, 23, 28, 29, 41, 43, 108
APD	avalanche photodiode 18
CCD	charge-coupled device 25, 26
ECDL	external-cavity diode laser 31, 88
EIT	electromagnetically-induced transparency 88
EOM	electro-optical modulator 30
GSR	graduate student researcher 37
HWP	half-wave plate 59, 61, 62
IC	integrated circuit 32, 34
LCVR	liquid crystal variable retarder 22, 41, 42
LO	local oscillator 23, 24, 25
LPF	low-pass filter 23, 81
ML	machine learning 111
MOT	magneto-optical trap 15, 16, 17, 20, 35
MTS	modulation transfer spectroscopy 28, 29, 30, 31
NIR	near infrared 4, 23
ODT	optical dipole trap 15, 17, 18, 19, 20, 21, 22, 25, 26, 27, 30, 31, 38, 49, 57, 59, 61, 62, 68, 69, 82, 113
PBS	polarizing beamsplitter 24
PCIe	Peripheral Component Interconnect Express 32, 113
PDH	Pound–Dreaver–Hall [1–3] 18, 30, 31
PG	polarization gradient 16
PSD	power spectral density 24
PZT	piezoelectric transducer 19
QED	quantum electrodynamics 12
QWP	quarter-wave plate 22, 59
RAM	residual amplitude modulation 28
rf	radio-frequency 5, 12, 16, 17, 18, 19, 23, 24, 28, 30, 42, 49, 50, 57, 58, 63, 72, 73, 79, 101, 103, 104, 108

RGA	residual gas analyzer 35, 37
SAS	saturation absorption spectroscopy 28, 29, 30
SNR	signal-to-noise ratio 69, 70
SOF	side-of-fringe 68, 69, 70, 72, 80, 81, 82
SPAD	single-photon avalanche detector 25, 41, 43, 107, 109
SRS	Stanford Research Systems 30, 35
TEC	thermo-electric cooler 20, 30
TOF	time-of-flight 26
UPS	uninterrupted power supply 32

Chapter 1

Background

Cold atomic gasses offer a convenient building block for complex quantum systems. At sufficiently low temperatures, many degrees of freedom can be frozen out, allowing for internal energy levels to be isolated and treated as simple anharmonic multi-level systems. Individual atoms are identical and indistinguishable, which allows for many measurements to be carried out in a calibration-free manner, and which makes scaling systems to larger sizes more straightforward.

Optical cavities allow for the population, control, and readout of well-defined optical modes. Modern engineering offers cavities with extremely high finesse and low mode volumes, which allow strong interactions to be engineered between cavity modes and intracavity oscillators (such as atomic samples), even at very low photon occupations.

Feedback control can be used to stabilize quantum systems to out-of-equilibrium states, and can broadly be broken into two categories. Incoherent (or measurement-based) feedback allows for the stabilization of quantum systems to particular states, conditioned on the results of (destructive) measurements of the systems' states. Coherent (or Hamiltonian) feedback allows for the stabilization of one quantum subsystem by another quantum subsystem, without the need for projective measurement; this allows for the autonomous generation of statistically interesting states, but is limited by the variety of coherent interactions that can be engineered between subsystems. Cold atomic gasses coupled to high-finesse optical cavities offer the necessary components to realize and explore quantum feedback in a variety of contexts.

In this chapter, I will give a simple introduction to ultracold atomic gasses, high-finesse optical cavities, and common interactions between the two. I will also introduce some common varieties of feedback stabilization in quantum systems, as well as other background information useful for understanding the material in this dissertation.

1.1 Ultracold atomic gasses

Atomic gasses have long been a staple of experiments investigating quantum processes. Unlike other candidate systems such as superconducting qubits [4–10], quantum dots [11–13], or crystalline impurities [14–20], all atoms of a particular species are identical, allowing samples to be used interchangeably and allowing calibrations to be trusted across generations of science. Cooling atoms to ultracold temperatures allows for motional degrees of freedom to be studied in detail or frozen out completely, and allows interactions between atoms to be controlled or ignored [21–23].

Reducing the temperature of an atomic sample sufficiently far that the relative Doppler shifts to incident light as seen by different individual atoms are small compared to the linewidths of the transitions between internal atomic states allows the internal states of the atoms to be isolated and studied in a controlled environment. The work described in this dissertation was carried out using ^{87}Rb , an alkali atom. With 37 bound electrons, the structure of ^{87}Rb is complex, comprising many bound states. In the context of cold atomic physics experiments, and particularly of this work, it is beneficial to focus only a small number of quantum levels of the atom, reducing each atom to an effective two- or three-level system.

There are many approaches to achieving this simplification. In the work described in this dissertation, we have relied on two basic methods. First, atoms can be prepared in a specific quantum state of interest, either by deterministically driving them to the state (*e.g.*, optically pumping) or by applying state-selective filters to remove any atoms in undesired states from the ensemble (*e.g.*, magnetic trapping). Second, atoms can be driven with electromagnetic radiation that is either narrowly close to a specific transition between two levels, or is deliberately very far from such a transition. In general, an atom in a selected initial state can be driven radiatively to a number of different final states. For example, as discussed below, light at a wavelength near 780 nm can drive a ground-state ^{87}Rb atom to any of several electronically excited hyperfine manifolds, separated by frequencies of 100s of MHz, and also any of several magnetic sublevels of each of those hyperfine manifolds. Coherent state transfer to just one of these excited states can be achieved by driving the atom with light that is resonant with a particular one of these transitions. Resolving a particular transition resonance generally also requires that Doppler shifts to the light, as seen by the atom, be suppressed, as is indeed achieved for ultracold atoms, and also that the transitions not be overly power-broadened. Alternately, one can in fact drive the atom with light *very far* from *all* of the accessible transitions. In this situation, the complexity of the many separate excited atomic states can be simplified greatly through adiabatic elimination of the states; the total result of the drive, then, is that of a Stark shift to the energy of the ground state. This latter approach is used extensively throughout this dissertation.

In the case of alkali atoms, the relevant optical transitions fall in the visible and near infrared (NIR) frequency ranges (wavelengths 400–1000 nm). These are generally between manifolds of hyperfine states (defined by their total spins f), the individual magnetic levels of which are often not resolved by the optical drives discussed above. A small magnetic bias field

can be used to effect Zeeman splitting between states within a hyperfine manifold, resolving the individual levels. For atoms with ground (or metastable) state hyperfine number f , this allows for the treatment of each atom as a spin- f system with $2f+1$ evenly spaced levels. The level splitting can be tuned with the magnetic field, and is generally in the radio-frequency (rf) or microwave frequency ranges (10–1000 kHz).

For our experiment, we use atomic ^{87}Rb . The ground-state $f = 2$ hyperfine manifold is made to interact dispersively with light (or with a vacuum cavity mode) off-resonant from the $5S_{1/2} \rightarrow 5P_{3/2}$ (D_2) transition at frequency ω_a (wavelength 780 nm); see Fig. 1.1a. The excited state hyperfine detunings are given by $\Delta_{f'=\{0,1,2,3\}} = \{-302, -230, -73, +194\}$ MHz, such that the dispersive ($|\Delta_{\text{ca}}| \gg |\Delta_{f'}|$) regime is reached at roughly a 2 GHz atom–cavity detuning (Fig. 3.1 shows the relevant calculations). Scattering events are uncommon, meaning that atoms remain in the ground-state manifold, which is treated as a spin-2 system and addressed using rf drives (Fig. 1.1b). Commonly in our experiment, the spins can be made to couple symmetrically to the environment, allowing the ensemble of spin-2 atoms to be treated as a single macroscopic collective spin; this is discussed in detail in Section 1.3 and in Chapter 3.

1.2 High-finesse optical cavities and cavity QED

Optical Fabry–Pérot cavities are linear resonators that can operate in the quantum regime. They are made up of two high-reflectivity mirrors (transmission $T = 10^{-5} - 10^{-3}$) mounted to face each other in such a way that they support standing-wave modes of light. Sufficiently well-aligned cavities can reach finesse (\mathcal{F}) in the hundreds of thousands, meaning that light will, on average, make hundreds of thousands of round trips between the mirrors before escaping. Meanwhile, small cavity mode volumes allow relatively large intracavity fields to be generated using very small mean photon occupations on the order of $\bar{n} = 0.1 - 10$. In this regime, the Hamiltonian dynamics of an optical cavity are simple:

$$\hat{H} = \hbar\omega_c \hat{n}. \quad (1.1)$$

The circulating energy in the cavity is proportional to the frequency ω_c of the cavity mode and to the number of photons $\hat{n} \equiv \hat{c}^\dagger \hat{c}$ present in the cavity, where \hat{c} is the annihilation operator on the cavity mode.

Eqn. 1.1 describes a perfect cavity with no loss, which does not exist. In practice, photons will escape from any real cavity at a finite rate which is dependent on the cavity geometry as well as the reflectivity of the cavity mirrors; this rate is characterized by the “cavity linewidth” 2κ . This can be accounted for by including a simple non-Hermitian term:

$$\hat{H} = \hbar\omega_c \hat{n} - \hbar\kappa \hat{c} e^{-i\omega_c t}. \quad (1.2)$$

In order to compensate for this loss, it is common to supply a steady-state “pump” to the cavity mode by shining coherent laser light through one of the cavity mirrors. When

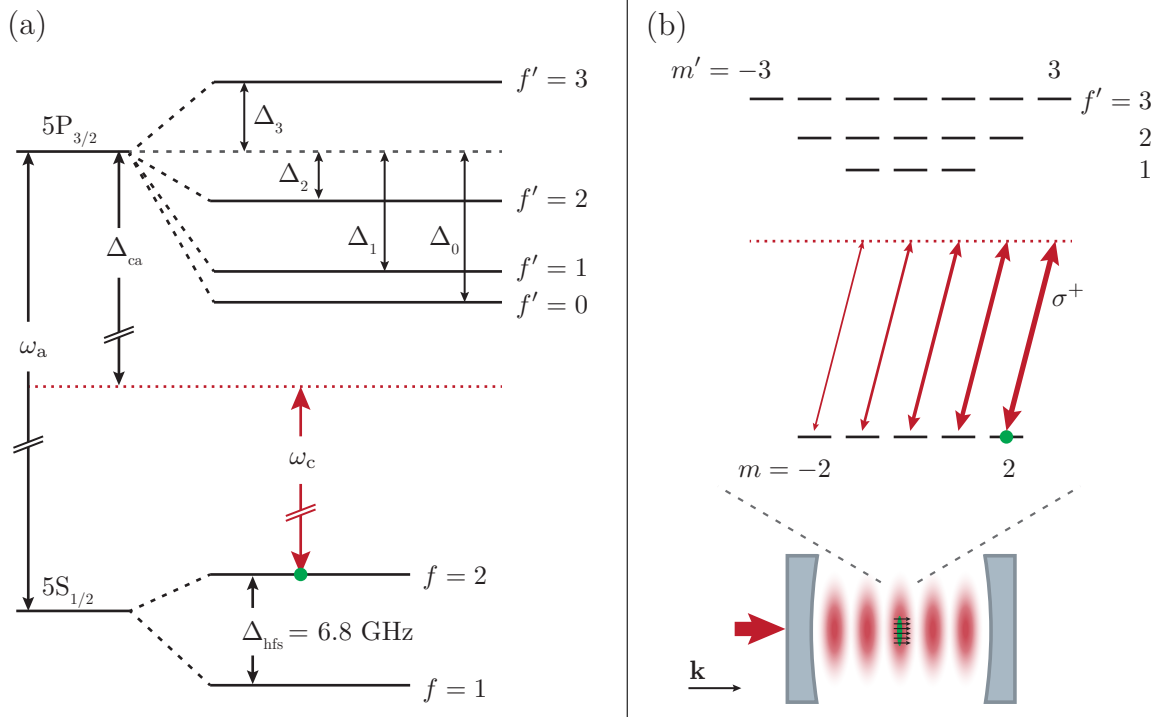


Figure 1.1: **a.** Energy levels relevant to the ^{87}Rb D_2 transition. The cavity mode is detuned by $|\Delta_{\text{ca}}| \gg |\Delta_{f'}|$ from the $5S_{1/2} \rightarrow 5P_{3/2}$ (D_2) line. **b.** The large atom–cavity detuning prevents the atoms from being excited out of the $f = 2$ ground-state hyperfine manifold, allowing this to be treated as a spin-2 subspace. Spin-dependent dispersive coupling to the cavity mode results in scalar and vector Stark shifts to the hyperfine levels. The relative coupling strength for each spin state, $\alpha_0 + \alpha_1 m$, is determined by the sum of Clebsch–Gordan coefficients for all accessible transitions to states in the $f' = 3$ manifold, summarized by the scalar and vector coupling strengths, $\alpha_0 = 2/3$ and $\alpha_1 = 1/6$, respectively (see Fig. 3.1). *Graphics adapted from Ref. [24].*

the pump light is resonant with the cavity mode, interference effects allow it to enter the cavity efficiently despite the high reflectivity of the mirrors. For incident light of coherent mode amplitude η , this can be accounted for by including a coherent pumping term in the Hamiltonian:

$$\hat{H} = \hbar\omega_c \hat{n} - \hbar\kappa \hat{c} e^{-i\omega_c t} + \hbar\kappa\eta (\hat{c}^\dagger e^{i\omega_c t} + \hat{c} e^{-i\omega_c t}). \quad (1.3)$$

The resulting Heisenberg equation of motion for \hat{c} can be solved directly in a frame rotating at ω_c , giving a photon occupation of $\langle \hat{n} \rangle = \eta^2$.

Indeed, it is not necessary that the pump light be at exactly the cavity resonance frequency. Pumping the cavity at arbitrary frequency ω_p will still result in some cavity occupation, though at a lower rate. This can be seen by noting that the cavity mode will now

evolve at ω_p rather than ω_c :

$$\hat{H} = \hbar\omega_c\hat{n} - \hbar\kappa\hat{c}e^{-i\omega_p t} + \hbar\kappa\eta(\hat{c}^\dagger e^{i\omega_p t} + \hat{c}e^{-i\omega_p t}). \quad (1.4)$$

Redefining $\hat{c} \rightarrow \hat{c}e^{-i\omega_p t}$ in order to enter a rotating frame at the pump frequency, this Hamiltonian can be rewritten as

$$\hat{H} = -\hbar\Delta_{pc}\hat{n} - \hbar\kappa\hat{c} + \hbar\kappa\eta(\hat{c}^\dagger + \hat{c}), \quad (1.5)$$

where $\Delta \equiv \omega_p - \omega_c$ is the detuning of the pump light from cavity resonance. Solving the Heisenberg equation of motion for \hat{c} in this frame gives the photon occupation for a cavity pumped at arbitrary frequency, which takes a simple Lorentzian line shape:

$$\langle \hat{n} \rangle = \eta^2 \frac{\kappa^2}{\Delta^2 + \kappa^2}. \quad (1.6)$$

It should be noted that, here, Δ refers to the detuning from the complete system's resonance condition, which in some cases may differ from the bare cavity resonance condition (if, *e.g.*, the cavity resonance is shifted by dispersive effects due to the presence of an atomic sample, as discussed below). In general, throughout this dissertation, Δ_{pc} will refer specifically to the pump's detuning from the bare cavity resonance frequency ω_c , while Δ will be used to mean something more specific, dependent on context.

One troubling step that I've in the discussion above is allowing the Hamiltonian to be a non-Hermitian operator. This offends a basic tenet of quantum mechanics, which is that measurable quantities (here, the energy of the system) should correspond to Hermitian quantum operators (here, the Hamiltonian). Consistency with this tenet is achieved in the theory of open quantum systems. In summary, the optical cavity can be considered to be part of a larger quantum system that includes also the continuum of electromagnetic modes incident on the cavity boundaries (the outside faces of the Fabry–Pérot mirrors). Deriving expressions for the dynamics of just the Fabry–Pérot cavity itself and treating the outside modes as a zero-temperature Markovian bath leads to (1) an effective Hamiltonian for the cavity that is non-Hermitian, and (2) a jump operator (not included here) that describes sudden changes to the cavity system when photons are emitted, and also the injection of noise from the outside bath. A full description of this procedure can be found in any of several quantum optics texts [25–27].

1.3 Simple atom–cavity interactions

In total, the interaction between an ensemble of two-level atoms and a cavity field is described by the Tavis–Cummings Hamiltonian [28–30]:

$$\hat{H} = \hbar\omega_c\hat{c}^\dagger\hat{c} + \sum_{i=1}^{N_a} \hbar\omega_a\hat{a}_i^\dagger\hat{a}_i + \hbar g(\vec{r}_i) \left[\hat{a}_i\hat{c}^\dagger e^{i\omega_c t} + \hat{a}_i^\dagger\hat{c} e^{-i\omega_c t} \right]. \quad (1.7)$$

Here, the local atom–cavity coupling can be separated as $g(\vec{r}) = g_0 U(\vec{r})$, where g_0 is the maximum atom–cavity coupling and U is the (normalized) field intensity profile of the cavity mode,

$$\vec{U}(r, z) = e^{-2r^2/w_0^2} \sin k_c z, \quad (1.8)$$

with $w_0 = 25 \mu\text{m}$ the waist of the cavity mode [31], $k_c = 2\pi/780 \text{ nm}$ its wavenumber, and r and z the radial and axial coordinates, respectively, relative to the cavity axis (with phase chosen arbitrarily).

In this two-level atom approximation, the complexity of the atom has been reduced to just one ground state $|g\rangle$ and one excited state $|e\rangle$; experimental methods for achieving this isolation of states are briefly described in Section 1.1. Using this approximation, the lowering operators used above are simply defined as $\hat{a} = |g\rangle\langle e|$.

For atoms localized at the maximum of the cavity field, $U(\vec{r}) = 1$, and the coupling is given by [32]

$$g(\vec{r}_i) = g_0 = \sqrt{\frac{\omega_a}{2\hbar\epsilon_0 V_m}} d = 2\pi \times 13 \text{ MHz}, \quad (1.9)$$

where $d = 2.989 ea_0$ is the dipole moment of the atomic transition [33] and $V_m = \int d^3 \vec{r} |U(\vec{r})|^2$ is cavity mode volume. For a near-planar cavity such as that used in this work, the mode volume simplifies to $V_m = \pi w_0 L_c / 4$. Likewise, for a near-planar cavity with length much less than the Rayleigh range of the cavity modes, the Gouy phase and mode divergence can both be ignored when calculating the mode profile as in Eqn. 1.8.

In practice, the interactions between the atom and the cavity can be tuned—both in strength and in character—by varying the detuning $\Delta_{ca} \equiv \omega_c - \omega_a$ between the atomic transition (here, the ^{87}Rb D₂ line at 780 nm) and the cavity mode. For sufficiently large detunings (as compared to the collective atom–cavity coupling $g_0\sqrt{N_a}$), the probability of absorbing cavity photons and transitioning to the excited atomic state becomes vanishingly small, and the system can be described dispersively [34]. In this limit, adiabatically eliminating the excited atomic states leads to an effective atom–cavity interaction of the form

$$\hat{H}_{\text{int}} = \hbar g_c \hat{c}^\dagger \hat{c} \sum_{i=1}^{N_a} U(r_i, z_i), \quad (1.10)$$

where the dispersive coupling strength has been defined as $g_c \equiv g_0^2 / \Delta_{ca}$. This amounts to a frequency shift to the cavity resonance, dependent on the total sum of the individual atom–cavity couplings; the effect parallels the corresponding AC stark shift to the atomic energy levels.

When all atoms couple symmetrically to the cavity mode, $U(r_i, z_i) = 1$, and the interaction simplifies to

$$\hat{H}_{\text{int}} = \hbar g_c N_a \hat{c}^\dagger \hat{c}, \quad (1.11)$$

a shift to cavity resonance directly proportional to the number of atoms coupling to the cavity mode. This effect, in its simplest form, is at the core of the work described in Chapter 4.




Figure 1.2: Schematic of measurement-based quantum feedback.

Further, accounting for the separate couplings of the ground state to the multiple hyperfine manifolds of the excited state (Fig. 1.1a) leads to a cavity shift which is dependent also on the internal magnetic states of the atoms (Fig. 1.1b); this interaction is derived and explored in detail in Chapter 3.

1.4 Feedback in quantum systems

In classical systems, it is common to want to control, or stabilize, the state of a particular subsystem. In the language of classical feedback, this subsystem is generically referred to as the “plant,” and the characteristic of the plant which is to be stabilized is referred to as the “control variable”. The “sensor” offers a readout of the control variable, which is fed to a comparator, along with an externally supplied “set point”: the value to which the control variable is to be stabilized. The comparator generates an “error signal” as the difference between the control variable and the set point. Finally, a “controller” conditions the error signal and acts back on the state of the plant, closing the feedback loop. Under proper negative feedback conditions, this system stabilizes the control variable to the set point unconditionally [35].

1.4.1 Incoherent, or measurement-based, feedback

The evolution of a quantum system is controlled by its Hamiltonian \hat{H} . In practice, this is set by variables such as magnetic field orientation and strength, light frequency and amplitude, and inter-atom interaction strengths. These control variables don’t need to be constant in time, and by varying the Hamiltonian the system’s trajectory can be altered.

Measuring the system’s state classically allows a feedback loop to be closed that can stabilize some quadrature \hat{X} of the system to a desired set point. The quantum system plays the role of the plant, with \hat{X} acting as the control variable. Dissipative measurement of $X_{\text{meas}} = \langle \hat{X} \rangle$ offers a sensor, the output of which can be compared to an external set point X_{set} to provide an error signal $\Delta X = X_{\text{meas}} - X_{\text{set}}$. Finally, altering \hat{H} by way of the Hamiltonian’s control variables, as informed by ΔX , closes the feedback loop. This process is outlined in Fig. 1.2, and is examined experimentally in Section 4.3.




Figure 1.3: Schematic of coherent quantum feedback. Some sketches from Section 3.3 are used for concreteness. In particular, the Bloch sphere within the Plant represents something (a collective spin) which might be stabilized; the cavity transmission profile within the Comparator represents a dissipative interaction (a red Stokes sideband scattering preferentially into a cavity over a blue anti-Stokes sideband) which might act as a controller. The part of the set point is played, in this sketch, by the location of the green carrier light, while the part of the error signal is played by the difference in amplitude between the two sidebands.

Measurement-based feedback [36–40] allows \hat{X} to be steered to a desired state; however, the dissipative measurement also reduces the coherence of the system, and, if \hat{X} does not commute with \hat{H} , it adds projection noise to other quadratures.

1.4.2 Coherent, or Hamiltonian, feedback

A conceptually simple way around the issue of projection noise inherent to measurement-based feedback is to avoid projective measurement altogether and rely, in stead, on coherent interactions between quantum subsystems. This can be achieved by dividing the total system Hamiltonian into a plant and a comparator:

$$\hat{H} = \hat{H}_{\text{plant}} + \hat{H}_{\text{comp}} + \hat{H}_{\text{int}}. \quad (1.12)$$

Here, \hat{H}_{int} comprises interactions between the plant and the comparator, and plays the role of both the sensor and the controller.

Generally, it is necessary that \hat{H}_{comp} either have some time dependence or incorporate some dissipation. For the work described in this dissertation, I will make use of the latter; in particular, I will use the frequency-dependent dissipation provided by the cavity to condition the system and act as a controller.

If the Hamiltonian \hat{H}_{plant} depends on some quadrature \hat{X}_p that is to be stabilized, \hat{X}_p must be made to interact with some other quadrature \hat{X}_c , through \hat{H}_{int} . An error signal is then formed by \hat{H}_{comp} , which uses dissipation to determine where the system lies relative to its set point and conditions \hat{X}_c accordingly. Finally, through \hat{H}_{int} , interactions with \hat{X}_c drive \hat{X}_p towards its final set point. This process is outlined in Fig. 1.3, and is examined experimentally in Section 3.3.

Coherent feedback [41–51] allows \hat{X}_p to be steered to a desired state without incurring added projection noise or loss of coherence, but is limited by the necessity of engineering the required interactions between quantum subsystems.

1.5 Definitions and conventions

Some symbols commonly used throughout this dissertation are defined in Tab. 1.1, and standard values are listed. The values of some parameters (such as Δ_{ca} and Δ_{pc}) are made to vary widely between experiments, and sometimes within an experiment. These values will be noted explicitly in the text, where appropriate.

It is often useful to speak of the physical apparatus in terms of its cardinal axes. The optical cavity used in this work is aligned along the vertical axis, which will sometimes be referred to as \hat{k} . One horizontal axis is defined by the atom chip waveguide wires used for magnetic transport (Section 2.1.1); this has been dubbed the “waveguide axis”. The remaining horizontal axis is defined by the optical imaging system (Section 2.3.3), which runs transverse to the waveguide axis; this has been dubbed the “imaging axis”. See Fig. 2.1. An alternate set of cardinal axes— \hat{x} , \hat{y} , and \hat{z} —is used in Chapter 3. These are determined by the quantization axis set by an externally applied magnetic field $\vec{B} = B_0 \hat{z}$. See Fig. 3.4.

Throughout this dissertation, I will attempt to be consistent with some notational conventions.

- Variable names such as g_0 (atom–cavity coupling), ω_c (bare cavity resonance frequency), Δ_{ca} (atom–cavity detuning) will be italicized. Operator names such as d (the infinitesimal unit) and \sin (the trigonometric function), as well as constants such as i (the imaginary unit) and π will be upright. One notable exception is \hbar (the reduced Planck constant), which will be italicized in keeping with standard convention.
- When referring to atomic hyperfine spins, mean single-atom spin vectors will be described by lowercase \vec{f} and symmetrized collective spins will be described by uppercase \vec{F} . The vector notation $\langle \hat{f}_z \rangle$ (likewise $\langle \hat{F}_z \rangle$) will often be used as shorthand for the expectation of the spin quantum number $\langle m_f \rangle$. Spin vectors will be defined without dimensions, meaning that factors of \hbar will appear explicitly in cases where dimensions of angular momentum are required.
- Hats will be used to indicate both quantum operators (q-numbers) as well as unit vectors; context should remove any ambiguity between the two uses. Occasionally, expectation values (c-numbers) corresponding to operators will be indicated by missing hats: $X \equiv \langle \hat{X} \rangle$.

Symbol	Definition	Value
L_c	Cavity mirror separation	250 μm
R_c	Cavity mirror radius of curvature	5 cm
w_0	Cavity mode waist (radius)	25 μm
κ	Cavity half-linewidth	$2\pi \times 1.81 \text{ MHz}$
m_a	Atomic mass of ^{87}Rb	$1.44 \times 10^{-25} \text{ kg}$
ω_a	^{87}Rb D ₂ $f = 2 \rightarrow f' = 3$ transition frequency	$2\pi \times 384.228 \text{ THz}$
Γ_e	Excited state decay rate	$2\pi \times 6 \text{ MHz}$
g_0	quantum electrodynamics (QED) coupling strength	$2\pi \times 13 \text{ MHz}$
C_{QED}	Resonant cavity QED coöperativity	16
Δ_{ca}	Cavity-atom detuning	$2\pi \times (-42 - - 2 \text{ GHz})$
g_c	Dispersive coupling strength	$2\pi \times -4 \text{ kHz}$
α_0	Relative D ₂ scalar coupling strength	2/3
α_1	Relative D ₂ vector coupling strength	1/6
α_2	Relative D ₂ tensor coupling strength	0
N_a	Atom number	1000–4000
\bar{n}	Intra-cavity photon number	0–10
Δ_{pc}	Cavity-probe detuning	$2\pi \times (-2 - 2 \text{ MHz})$
ω_m	Axial trap frequency	$2\pi \times (90 - 150 \text{ kHz})$
ω_s	Larmor frequency	$2\pi \times (100 - 700 \text{ kHz})$
ϵ	Total cavity photon detection efficiency	2–5 %

Table 1.1: Parameter definitions and values.

1.6 A note on the anthropomorphization of scientific instruments

Early in one’s career as a research scientist, there is a strong inclination to try to understand everything. There are many moving pieces in an atomic physics lab, and understanding how they work with each other, and under what conditions they do and don’t operate as intended, is indispensably helpful for running a successful experiment.

In publications and other formal communications, it is not appropriate to refer to a sample of trapped atoms being happy or sad, nor to refer to an rf amplifier as complaining or as wanting to be treated differently. In day-to-day lab interactions, however, these characterizations form a useful shorthand. When scientific explanations aren’t available for how and when things do and don’t work, a more heuristic, or “vibes-based,” vocabulary can allow for easy discussion of the problems at hand while avoiding inaccuracy and ambiguity. Noting that a microwave supply seems to become unhappy when a certain room light is turned on, without needing to speculate as to how the two might interact, can allow for science to move forward. Throughout this dissertation, I will attempt to refrain from using this sort of language, but I would encourage any young students who find themselves reading this to

acknowledge that the instruments we work with are often too complicated to understand in their completeness, and to embrace their (mild) anthropomorphization as a natural way of working through this complexity.

Chapter 2

The E3 apparatus at Berkeley

The apparatus on which I've worked throughout my time at Berkeley is E3, the third experiment to have been brought online by the Stamper-Kurn group. E3 was constructed from 2009–2011, primarily by Tom Purdy and Daniel Brooks, in the same room as its predecessor E2. In the decade since it was brought online, the apparatus has produced many significant results [52–58]. The lifetime of a modern atomic physics apparatus, however, rarely stretches beyond 3–4 generations of graduate students. New technologies are born, and what was once on the cutting edge of research becomes commonplace. More significantly, perhaps, old machines become increasingly difficult to maintain, and become increasingly cobbled together after years of repairs carried out by rushed, semi-qualified scientists.

By the time I joined the Ultracold group at Berkeley, E3 was already an old experiment. The plan was to shut down the apparatus after Jonathan Kohler, my predecessor, graduated, and for me to join a different experiment in the group for the remainder of my Ph.D. Every time we finished a project, though, another goal seemed well within reach. Things broke more often than they might have on a younger experiment, but not more quickly than we could fix them. New postdocs and undergraduate researchers—and eventually another new graduate student, Olive Eilbott—joined the project, and the apparatus kept producing interesting results. Not only will I have completed my Ph.D. on the E3 apparatus, but I won't be the last one to do so.

A big part of what has allowed E3 to continue producing interesting results is its unique combination of technologies. Historically, experiments interfacing cold atomic ensembles with optical cavities have done so by trapping large, dilute clouds of gas such that they overlap with the (macroscopic) cavity mode [59, 60]. More recently, modern experiments have evolved to trapping small ensembles of atoms using optical tweezers, which can be placed arbitrarily within cavity modes [61–63]. The E3 apparatus is unique in its ability to offer precise, arbitrary control of the placement of the atomic sample within the cavity mode while still making use of mesoscopic ($\sim 10^3$) numbers of atoms. This is achieved by interfacing the optical cavity directly with an atom chip [31, 52, 64] (Fig. 2.1).

The E3 cavity mirrors lie on either side of the atom chip, and the cavity mode passes through a small hole in the chip. During science experiments, the atomic sample is trapped

in a standing-wave optical dipole trap (ODT) which takes the form of a far off-resonant cavity mode. The proximity of the magnetic wires on the atom chip to the cavity allows the atoms to be magnetically trapped in a sufficiently small volume, before being transferred to the ODT, that they primarily occupy a single well of the standing-wave lattice. Further, the location of the magnetic trap can be moved with sufficient precision to control into which well of the ODT they are loaded. This, coupled with the finite beat frequency between the trapping mode and the probing mode of the cavity, allows for precise control of the coupling between the atomic ensemble and the probing mode of the cavity (Fig. 2.3), enabling a wide variety of unique studies to be carried out using this apparatus.

In this chapter, I will introduce the core elements of the E3 apparatus and of the experimental sequence, highlighting some technologies which we have found to be of particular interest. I will also describe the tools that we've found useful for troubleshooting the sequence and the apparatus itself. Finally, I will give a brief account of some significant experimental setbacks which have occurred during my time in graduate school, and of some ways in which they've each conditioned me to be a better experimentalist.

2.1 Trapping, cooling, and preparation of the atomic sample

As is standard for cold atomic physics experiments, the E3 experimental sequence begins with a magneto-optical trap (MOT) [65–67]. Conventionally, this requires 3 pairs of counterpropagating laser cooling beams, a single current coil to produce a magnetic quadrupole field, and one or more pairs of Helmholtz coils to produce a uniform magnetic bias field. In order to minimize the distance over which the atoms must be transported, it is desirable to place the MOT near to where the atoms will next be used: in our case, the atom chip. This places some considerable limitations on the construction of the MOT. First, the quadrupole coil is placed in-vacuum. This, along with the mounts for the cavity mirrors, limits optical access to the center of the science chamber, preventing 2 pairs of horizontal counter-propagating cooling beams from being used. In stead, a single pair of horizontal counterpropagating beams is used, which is formed by a single beam that is retroreflected by an external mirror after passing through the science chamber. For the remaining two pairs of cooling beams, light is sent into the science chamber from below at 45° angles and reflected off of the coated surface of the atom chip. (Incidentally, due to spatial constraints, one of these beams takes a path directly through the center of the in-vacuum quadrupole coil.) The initial MOT captures $\sim 10^8$ atoms at ~ 1 mK.

In order to transport the atomic sample into the volume subtended by the cavity mode, it is necessary to trap the sample using magnetic fields generated by the atom chip. The fields generated by the chip are not large enough to produce a MOT with sufficient capture velocity to load from a background vapor; however, once the initial MOT has been generated using the external quadrupole coil, it can be transferred to a smaller MOT which uses the




Figure 2.1: Above: Schematic of the E3 apparatus. Atoms are initially captured in a standard MOT formed using an external quadrupole coil, then transferred to a chip mot (“U-MOT”) formed using wires on the atom chip. From there, they are transferred to a magnetic trap and magnetically transported to their final position within the cavity. Below: Image of the E3 cavity and chip. The chip is etched to a $\sim 100 \mu\text{m}$ thickness near the cavity, and the cavity mode shines through a hole in the thin section of the chip. *Graphics adapted from Ref. [24].*

same cooling light but relies only on a chip-generated quadrupole field. This acts as an intermediate step between the initial large mot and the eventual chip-generated magnetic trap. Before loading into the chip MOT, compression of the initial MOT by increasing the detuning of the cooling beams from the atomic transition increases the phase-space density of the sample, compensating for the heating caused by its release and recapture and allowing for similar atom numbers and temperatures to be achieved in the chip MOT as in the initial MOT. See Fig. 2.1.

From the chip MOT, the atomic sample is briefly released, cooled using polarization gradient (PG) cooling [68, 69], optically pumped into the most trappable state, and transferred directly to the chip-based magnetic trap. This quick transfer is made possible by the close proximity of the chip wires to the trap, and by their low inductance. The magnetic trap is formed by the first link of the conveyor used for magnetic transport (Section 2.1.1), and initially holds $\sim 10^7$ atoms at $\sim 15 \mu\text{K}$. While held by this initial magnetic trap, the atomic sample is cooled using forced rf evaporative cooling, resulting in $\sim 5 \times 10^5$ atoms at $\sim 2 \mu\text{K}$.

For a more complete reference, see Tom Purdy’s dissertation [31].

2.1.1 Magnetic transport using the atom chip conveyor

In order to form a MOT, sufficient optical access is required to illuminate the trap with three pairs of counterpropagating beams, each with a large radius and clean wavefront. Designing an optical cavity that met the requirements of the E3 experiment, meanwhile, required

placing the cavity mirrors very close together, resulting in poor optical access to the final location of the atomic sample. To reconcile these design requirements, the initial MOT is formed away from the cavity and then transported, via magnetic conveyor wires, to its final location. The magnetic trap in which the first stage of rf evaporative cooling takes place is formed in the first link of the magnetic conveyor, and transport takes place immediately after evaporative cooling (outlined above).

This process is described in detail in Ref. [31], and will not be outlined here. Rather, I will briefly discuss methods used to troubleshoot the transport process. Some further complications related to the heat generated by the magnetic wires during transport are detailed in Section 2.2.1 and Section 2.2.2.

For a variety of reasons, it may be the case that no atoms are detected in their final location within the cavity. Commonly, this is due to issues with the trapping and cooling stages of the early sequence; absorption images of the cloud at various points during preparation (Section 2.3.3) makes it straightforward to determine when this is the case. Less commonly, a lack of atoms is due to issues with magnetic transport or with loading into the cavity ODT; troubleshooting this is more difficult, as it is less straightforward to optically image the atoms at any point between the start of transport and their final position.

The most consistent way that we’ve found of checking whether magnetic transport is working as intended has been to fold it back on itself, first transporting the atoms forward to the cavity, then backward to the initial magnetic trap. This allows the ensemble to be optically imaged under the well-understood conditions of the initial magnetic trap, and directly tells us whether transport is working as intended. Isolating transport from the final stage of evaporative cooling and loading into the cavity ODT makes it possible to troubleshoot with increased confidence.

2.1.2 Final RF evaporation and loading into the cavity ODT

The magnetic trap effected by the atom chip is sufficiently tight to confine the atomic sample to a volume comparable to that of a single well of the standing-wave cavity ODT [31]. When the sample arrives at the cavity after magnetic transport, however, it is hot enough that it explores a much larger region of the trap. In order to combat the heating caused by magnetic transport, a second round of rf evaporation is employed before leading the atoms into the ODT. Evaporating “farther”—ramping down the rf evaporation tone to a lower final frequency—results in a colder, and thus more tightly confined, sample, but also in a smaller final atom number (Fig. 2.2).

Because the cavity pump mode (wavelength 780 nm) and ODT mode (842 nm) are of incommensurate frequencies, the local intensity of the pump mode varies between neighboring wells of the trap, with maxima and minima occurring at a $5.3\ \mu\text{m}$ period corresponding to the beating between the two modes (Fig. 2.3). The extent of the localization of the sample into a single well of the ODT is measured using the contrast C of the mean atom–cavity coupling g with respect to loading position: if the size of the sample along the cavity axis is large compared to the $5.3\ \mu\text{m}$ beating period between the pump and ODT standing waves,




Figure 2.2: Final atom numbers and temperatures of the atomic sample after evaporatively cooling to different final rf tones. Atom numbers are measured dispersively (Section 3.2); error bars represent statistical uncertainty. Temperatures are measured using time-of-flight imaging (Section 2.3.3); error bars represent fit uncertainties. Atom numbers and temperatures also depend on a wide variety of other experimental parameters.

g will not depend on loading position ($C = 0$); whereas, if the sample is perfectly localized, g will vary from 0 to g_0 ($C = 1$).

In the past, collective mechanical motion of the atomic sample has been studied by placing the sample in a well (or wells) where the probe intensity is strongly spatially dependent [56, 57, 70, 71]. For all of the work described in this dissertation, it has been desirable to place the sample primarily in a single well which overlaps with an antinode of the pump mode, where the coupling to cavity is maximized and is only quadratically sensitive to mechanical motion along the axial direction, and where all atoms in the sample can be treated as coupling symmetrically to the pump mode.

2.2 Intricacies of the science cavity

The core of the E3 apparatus is the so-called “science cavity,” generally referred to, throughout this dissertation, simply as the cavity. Some specifications of this cavity are given in Tab. 1.1. The experiments described in this dissertation have often required the cavity resonance condition to lie at a precise frequency relative to the frequency of the pump laser and to atomic resonance. Meanwhile, the atomic ensemble is trapped, during science, in a standing-wave ODT which must itself be kept on-resonance with the cavity. Simultaneously meeting these conditions requires some care.

The ODT and pump lasers are each locked to the actively stabilized reference cavity, as discussed in Section 2.4.2. For most experiments, the science cavity is then locked to the ODT using a Pound–Dreaver–Hall [1–3] (PDH) scheme based on the cavity transmission, as measured by an avalanche photodiode (APD) after the ODT and pump are split at




Figure 2.3: The atom–cavity coupling g depends on the loading position of the atomic sample along the cavity axis, as well as on the spatial extent of the sample. The coupling is determined by measuring the dispersive shift to the cavity resonance condition due to the presence of the atomic sample. **a.** With less rf evaporation, we see a larger atom number, but atoms are spread out over a larger number of ODT wells resulting in a lower contrast between the maximal and minimal coupling strengths. **b.** Evaporating farther results in fewer atoms in the final sample, but allows them to be better localized in a single well of the ODT, resulting in high contrast between the maximal and minimal coupling strengths.

the cavity output. This lock feeds back to piezoelectric transducers (PZTs) controlling the cavity’s length. This references the pump to the cavity; the detuning between the two is then controlled by an offset acousto-optical modulator (AOM) in the pump path, located after the pickoff for locking to the reference cavity.

Further details of the locking and stabilization of the various lasers are given in Section 2.4.

2.2.1 Heating of the cavity by the atom chip wires

A primary context in which magnetic transport has been relevant to the day-to-day operation of the experiment is the heat generated by the chip wires. In order to maintain relative mechanical stability, the atom chip and the optical cavity share a mount. This allows for consistently loading into the same position in the cavity, but also means that the thermalization of the chip wires directly affects the thermalization of the cavity mounts, and thus the resonance condition of the cavity. This effect is mitigated by actively stabilizing the

temperature of the chip. The stabilization is achieved using thermo-electric coolers (TECs) that couple the chip mount to a large copper heat sink, which is externally water-cooled; this configuration is detailed in Ref. [31].

The maximum load of the TECs, the thermal conduction of the heat sink, and the capacity of the water-cooling each limit the rate at which the chip wires can be allowed to generate heat. This rate is exceeded significantly during magnetic transport, meaning that the system heats up beyond its ideal equilibrium temperature. During periods of the sequence when the chip wires are not active, the system cools to equilibrium. This is the limiting factor in the E3 cycle length: The $1 - 1/e$ MOT loading time is around 6 s, but in practice we have generally been required to load the MOT for at least 20 s, not to increase the number of atoms present, but to allow sufficient dead time in the sequence for the system to cool; otherwise, the copper heat sink heats up enough over the course of many cycles that the TECs controller reaches its rail.

More realistically, the heat generated by transport might be reduced by reducing the time spent transporting the sample. Historically, a transit time of 100 ms has been allotted to each of the 32 links in the transport chain, resulting in a total magnetic transport time of 3.2 s during which the system is being significantly heated. This transit time was chosen based on the tradeoff between going sufficiently slowly to meet the relevant adiabaticity conditions and sufficiently quickly to minimize losses due to collisions with background gas (corresponding to a $1/e$ loss rate of around 6 s) [31].

Transporting more quickly would result in increased losses due to diabatic heating, and a lower final atom number, but would allow for the experiment cycle time to be sped up significantly by also reducing the amount of dead time required for cooling the apparatus. In 2021, we reduced the per-link transit time from 100 ms to 66.6 ms, which allowed the MOT loading time to be reduced from 20 s to 12 s without increasing the overall load to the TECs. This reduced the overall cycle time from around 32 s to around 23 s, significantly increasing productivity.

2.2.2 Chip quakes: mechanical bistability of the atom chip

A surprising result of the thermal cycling of the apparatus during the experimental sequence is that the chip itself appears to be brought through a mechanical bistability. Like a deformed baking pan in a heating oven, the chip shifts between two stable mechanical configurations as it is heated and cooled during each cycle of the experiment. This has two significant effects:

1. The two mechanical chip states correspond to different relative positions between the chip and the cavity. This means that, if the chip is in the wrong state when the atoms are loaded from the magnetic trap into the ODT, the loading position will be incorrect. This is easy to avoid by ensuring that the equilibrium temperature of the apparatus remains consistent between experimental runs.
2. When the chip shifts between states, the sudden impulse excites mechanical vibrations in the cavity mounts. These vibrations are not sufficient to unlock the cavity, but they




Figure 2.4: Effect of chip quakes on the cavity. The transmitted intensity of the ODT light is modulated by the cavity length, giving a direct readout of the mechanical vibrations. Below: Ringing is observed at the 6.8 kHz fundamental mechanical mode, as well as the first few harmonics. Above: The chip quake, beginning at 46 ms, initially excites a mode common to the two cavity mirrors, which only has higher-order effects on the cavity transmission; it takes some time for coupling between the common mode and the breathing mode between the two mirrors to make the effects of the quake maximally apparent (here, the range 46–48 ms). After the initial ring-up, beating between the common and breathing modes is clearly visible (50–64 ms).

do result in atom loss from the ODT if they occur when atoms are present. We refer to these bistability-induced vibrations as “chip quakes,” and they have been the subject of considerable frustration during my time in E3.

Chip quakes can be observed directly, as the relative motion of the cavity mirrors is imprinted on the intensity of transmitted light (Fig. 2.4). If a chip quake happens between the time when atoms are loaded into the ODT and the end of the sequence, it results in significant atom loss. The times at which chip quakes occur can be roughly controlled by tuning the set point temperature of the apparatus, and the temperature can be tuned such that, when the system is running stably, the problematic chip quake occurs after atoms have been released from the trap at the end of the sequence.

2.2.3 Controlling the cavity pump polarizations with LCVRs

The cavity portion of the natural eigenmodes of the atom–cavity system is approximated well by coherent states of circularly polarized (σ_+ and σ_-) light. Pumping the cavity with different circular polarizations, as defined relative to the projection of the background magnetic field onto the cavity axis, results in materially different interaction Hamiltonians for the system (Section 3.1). As such, it is often desirable to have computer control over the pump polarization in order to change the form of interactions throughout the experimental sequence. This requires controlling two waveplates: one at the input to the cavity, before the pump and the trap light are combined, and one at the output of the cavity, before the light is detected.

Rather than mount quarter-wave plates (QWPs) to motorized rotational stages, which would be both prohibitively slow and mechanically noisy, we choose to use a pair of liquid crystal variable retarders (LCVRs), which can be tuned smoothly from 0-waveplates (no-ops) to 1/2-waveplates. Tuning the retardance over its full range can be accomplished in ~ 30 ms in the slow direction and in ~ 300 μ s in the fast direction. This fast tuning between σ_+ and σ_- pump light enables the nondestructive measurement of both the total atom number and the collective atomic spin at the end of each experiment (Section 3.2).

2.3 Collecting information from the system

For a large majority of atomic physics experiments, optical imaging is an indispensable tool. Atomic samples are prepared and manipulated, experiments are run, and images are recorded—either absorption [72] or by fluorescence [73, 74]—which are used to back out the full story of what has occurred. Imaging atoms immediately after they are released from their traps gives information about their initial position, while waiting some time before imaging gives information about their initial momentum by mapping it onto their final position [75]. Modern control sequences allow for optically imaging atoms in a state-sensitive manner, enabling local readout of which atoms end up in which states. These myriad use cases have led to a modern era in which quantum gas microscopes [74, 76–80] are a staple of many atomic physics labs.

In E3, however, optical access to the atomic sample is very limited. Designing for the ability to precisely load an atomic sample primarily into a single well of the standing-wave intra-cavity ODT dictated that the atom chip and the cavity almost completely surround the final position at which science takes place. Because of this, optical imaging of the atomic sample is only possible at a very low resolution, and does not offer the ability to collect sufficient information about what has occurred. In stead, the light leaking from the cavity is detected directly. Any interactions between the atomic sample and the cavity are imprinted on the light populating the cavity mode as amplitude and phase modulations, and, by collecting information about these modulations, a full story can be reconstructed.

2.3.1 Heterodyne detection

All of the information necessary to recreate the physics happening in the cavity is imprinted on the cavity output, but directly detecting and analyzing frequency and phase information in the visible and NIR bands is technologically infeasible. Optical heterodyne detection offers a way of mixing the relevant signals down to the rf band, allowing for straightforward detection and analysis.

Consider light pumping the cavity with amplitude A and frequency ω_p :

$$\xi_{\text{in}}(t) = A \cos \omega_p t. \quad (2.1)$$

The light escaping from the cavity will generically have had its amplitude, frequency, and phase modulated by its interaction with the atom-cavity system:

$$\xi_{\text{out}}(t) = b(t) A \cos ([\omega_p + \omega'(t)]t + \phi(t)). \quad (2.2)$$

Here, in practice, ω' and $d\phi/dt$ lie within or below the rf band or below. The goal is to read out the information contained in $b(t)$, $\omega'(t)$, and $\phi(t)$ without requiring the resolution necessary to resolve ω_p , which lies within the visible or NIR bands. To achieve this, the output light can be overlapped, on the face of a photodiode, with light of a constant frequency that is offset by some amount ω_{het} from the pump:

$$\xi_{\text{LO}}(t) = C \cos ([\omega_p + \omega_{\text{het}}]t + \phi_{\text{LO}}). \quad (2.3)$$

This second beam of light is referred to as the local oscillator (LO). Mixing the cavity output with the LO on the face of the photodiode results in spectral components near ω_{het} and near $2\omega_p + \omega_{\text{het}}$. Putting the signal through a low-pass filter (LPF) results in the so-called heterodyne output:

$$v_{\text{het}}(t) = b(t) D \cos ([\omega_{\text{het}} + \omega'(t)]t + \phi_{\text{LO}} + \phi(t)). \quad (2.4)$$

Here, the signal amplitude D depends on the pump and LO amplitudes as well as the overlap efficiency of the beams, the quantum efficiency of the photodiode, and the transimpedance gain of the detector. The heterodyne output contains all of the desired information, and has been mixed down to a band that is much easier to work with (Fig. 2.5).

Taking the case $\omega_{\text{het}} = 0$, this reduces to a homodyne detection scheme (the basis for lock-in measurement), which is technologically simpler but removes the ability to differentiate between positive and negative components of ω' and ϕ . This results in a factor of 2 increase in detection efficiency, but it requires a stable LO phase, which is limiting in practice.

In E3, the LO for the optical heterodyne detector split from the pump immediately after the fiber launcher that brings the light to the science table. Whereas the pump is shifted by $\omega_{\text{het}} = 2\pi \times 10$ MHz by a pair of single-pass AOMs (one AOM that shifts the frequency up, followed by another that shifts the frequency down, with ω_{het} being the net shift), the LO is sent directly to the region of the table where detection takes place. Upon leaving the cavity




Figure 2.5: Typical power spectral density (PSD) of the light leaking from the cavity, as measured by the heterodyne detector. Power is referenced to the carrier at 10 MHz. Peaks corresponding to modulation of the cavity field by the Larmor precession of the atomic sample are visible, in this example, at ± 1.4 MHz.

and the vacuum chamber, the pump light is collimated at a different size than the cavity input (and thus the LO); this requires the pump to be put through a 3 : 2 reducer telescope in order to match the beam waists w_0 of the pump and LO. In order to further optimize the mode matching between the pump and the LO at the heterodyne detector, it is desirable to make sure that their waists occur at similar locations z_0 ; to this end, the LO is sent through a 1 : 1 telescope with a final lens at the same location, and of the same focal length, as the final lens of the pump reducer telescope. Once beam shaping is complete, the pump and LO are combined on a polarizing beamsplitter (PBS). The combined beam is sent through a 1/4-waveplate mounted to a precision micrometer rotational stage before being split on a final PBS and sent to the two detectors of a balanced photodiode (Newport 1807-FS). This allows the beam balance between the two detectors to be fine-tuned in order to zero the DC component of the heterodyne signal.

The output of the balanced photodiode is sent through an rf bias tee. The DC component is monitored for alignment purposes, and the AC component is amplified and recorded.

Some care must be taken when designing (and re-designing) an optical heterodyne detector. First, in order to reduce detection shot noise, it is important for us to ensure that none of the 842 nm trap light makes it to the balanced photodiode. The pump and trap are separated from each other on a 45° long-pass dichroic mirror immediately after leaving the vacuum chamber; however, this still reflects $\sim 1\%$ of trap light. In order to further minimize stray light, a second dichroic mirror is placed in the path immediately before the pump and LO are combined. Second, because of the very small amount of pump light used for many of our experiments, it is important to ensure that none of the light from the LO reflects back through the detection system and illuminates the cavity. To this end, an optical isolator is placed in the pump path between the reducer telescope and the final dichroic. In the past, when this isolator has not been present, light from the LO has interacted with the atomic ensemble in ways which have been particularly difficult to detect due to the spectral proximity of the interaction modes to ω_{het} .

Details of how the heterodyne detection efficiency is measured, and of how the cavity mode occupation and other useful quantities are derived from the heterodyne spectrum, are described in Ref. [24].

2.3.2 Single photon counting

Heterodyne detection offers complete information about the light escaping the cavity, but has decreased efficiency due to the need to overlap the cavity output and LO beams, and has a bandwidth limited by the heterodyne frequency ω_{het} . If information is required describing very small amounts of light or very precise timing, direct detection of the cavity output by a single-photon avalanche detector (SPAD) offers significant benefits.

A SPAD offers a single digital output which goes high for a brief period of time when a single photon is detected at its input. The quantum efficiency of the photon collection is $\epsilon \gtrsim 50\%$, meaning that the overall detection efficiency is limited primarily by losses in the cavity substrate, which are unavoidable. Meanwhile, the detection bandwidth is limited only by the ~ 45 ns dead time after a photon is detected.

The SPAD used in E3 (Laser Components COUNT) is coupled directly to a fiber at the cavity output, after the pump light is split from the ODT by a dichroic mirror. It has been used, in particular, for measuring the cavity linewidth (Appendix B), as well as for working towards detecting superradiance on the $^{87}\text{Rb} + ^{87}\text{Rb} \rightarrow ^{87}\text{Rb}_2$ molecular transition (Section 5.2.2).

2.3.3 Optical imaging of the atomic cloud

Despite limited optical access to the final location of the atomic sample, absorption imaging can still be a useful diagnostic tool. There is roughly a $100\ \mu\text{m}$ gap between the lower cavity mirror and the bottom of the atom chip (Fig. 2.1), which can be accessed from the side with a numerical aperture of roughly 0.2. After each cycle of the experiment, the ODT is diabatically turned off in order to release the atomic sample, and imaging light is sent in from the side to illuminate the sample. At the opposite side of the vacuum chamber, the light is magnified through a telescope and collected by a charge-coupled device (CCD) camera (Andor iXon, model DU-885K-CS0); this is referred to as the “signal” image. The atomic sample is then illuminated with high-intensity near-resonant light (picked off from the light used for trapping and cooling), causing it to disperse quickly, and another image is collected with the imaging light turned on but with no atoms present; this is referred to as the “reference” image. Finally, the imaging light is turned off and a third “background” image is collected. The spatially varying power attenuation of the imaging light by the atomic sample can then be reconstructed by combining the three images:

$$t^2 = \frac{\text{reference} - \text{signal}}{\text{background}}, \quad (2.5)$$




Figure 2.6: Typical absorption image of an ensemble of ~ 2000 atoms within the cavity. The image is recorded in time-of-flight (TOF), 0.4 ms after the sample is diabatically released from the ODT. Lower local transmission corresponds to a higher column-density of atoms.

where each pixel is treated independently. Conversely, the transmission of on-resonant light through the atomic sample is given by

$$t = \exp\left(-\frac{\rho_c \sigma_0}{2}\right), \quad (2.6)$$

where ρ_c is the local column density of the atomic cloud along the imaging axis (with units of inverse length squared), and where $\sigma_0 = 3\lambda^2/2\pi = 2.9 \times 10^{-9} \text{ cm}^2$ is the interaction cross-section on the transition being probed. This assumes that the transition is saturated by the imaging light, $I \gg I_{\text{sat}}$. Reconstructing ρ_c at each pixel allows the spatially dependent density of the atomic sample to be extracted, averaged over the imaging axis (Fig. 2.6).

The pixel size of the CCD and the magnification of the imaging telescope result in a resolution of $2.8 \mu\text{m}/\text{pixel}$, which corresponds to a 35-pixel region, vertically, over which the sample can be imaged. In practice, diffractive effects distort the top and bottom of the image significantly, making only the central ~ 25 -pixel area reliable. The size of the atomic sample, after 0.4 ms time of flight, is roughly $10 \mu\text{m}$ in the radial (horizontal, “waveguide-axis”) direction and $25 \mu\text{m}$ in the axial (vertical) direction; the difference is due to the ODT offering much tighter confinement in the vertical direction. These scales do not offer sufficient resolution for measuring the specifics of the motional state of the atomic ensemble, nor of which well or wells of the ODT it has been loaded; however, it is possible to fit a simple 2D Gaussian distribution to the optical density of the cloud, giving an approximate atom number and cloud size. This relatively simple test of whether an atomic sample of the expected size is present in the expected location has often proven useful when trying to debug whether the cavity probing light is acting as intended.

Further, this imaging can be used to measure the temperature of the atoms. The rate of expansion of the atomic sample, after it is diabatically released from the ODT, is proportional to the mean speed \bar{v} of the atoms before their release:

$$\sigma(t) = \sqrt{\sigma_0^2 + (\bar{v}t)^2}, \quad (2.7)$$

where $\sigma(t)$ is the (linear) size of the cloud at time t after release and $\sigma_0 \equiv \sigma(t=0)$ is the size of the sample in the ODT. The mean speed is, in turn, related to the temperature of




Figure 2.7: Measuring the temperature of the atomic sample using absorption imaging. After the atomic sample is released diabatically from the ODT, it expands at a rate related to its temperature by Eqn. 2.8. Fitting to the size of the cloud at a variety of times after release allows the temperature to be extracted reliably.

the gas as dictated by equipartition:

$$\frac{3}{2}k_{\text{B}}T = \frac{1}{2}m\bar{v}^2, \quad (2.8)$$

where $m = 1.443 \times 10^{-25}$ kg is the atomic mass of ^{87}Rb . Measuring the size of the atomic cloud at a series of different times after its release from the ODT, then, allows for \bar{v} and σ_0 to be extracted directly, and for the temperature to be calculated (Fig. 2.7).

2.4 Locking and stabilization of the E3 lasers

As alluded to in Section 2.2, keeping all of the E3 lasers at the appropriate frequencies (both relative and absolute) requires some care.

2.4.1 Locking to pump–probe spectroscopy of rubidium

Optically trapping and cooling atomic vapors requires light that is at precise frequencies relative to the relevant atomic transitions. The most direct way of achieving this is to reference (“lock”) the light directly to an atomic sample. At a basic level, this can be achieved by sending light directly through a vapor cell containing a room-temperature gas of atoms; light near the atomic resonances will be absorbed by the sample, while light far away from resonances will see a transparent gas and pass through the cell freely. This approach, known as absorption spectroscopy, is limited in practice: atoms moving along the axis of light propagation will see Doppler-shifted frequencies of light, meaning that light will be absorbed at a range of frequencies according to the thermal distribution of atomic velocities.

The technique of pump–probe spectroscopy (Fig. 2.8) gets around the issue of Doppler broadening by requiring that two counterpropagating beams be simultaneously resonant with




Figure 2.8: Generic pump–probe spectroscopy setup. Some light is picked off from a laser’s output before being sent to do science. The picked-off light is split on a PBS, and the two beams are sent to counterpropagate through a room-temperature ^{87}Rb vapor cell. Before being sent through the vapor cell, one of the beams is offset and modulated by a double-pass AOM. Either beam can be detected after passing through the vapor cell, giving information about the ^{87}Rb spectrum which can be fed back to the laser directly.

an atomic transition [81–84]. This technology is not new, and I won’t detail the operating principles here. I will note, though, that in August, 2021, the decision was made to switch from using saturation absorption spectroscopy (SAS) to using modulation transfer spectroscopy (MTS) for locking the repump laser. The benefits were significant enough that, a few months later, when the spectroscopy setup for the main cooling laser next needed tuning up, it was deemed worthwhile to convert this setup from SAS to MTS as well.

SAS relies on modulating the frequency of a low-power probe beam which is sent through an atomic vapor cell and then detected by a fast photodiode; a counterpropagating high-power pump beam of constant frequency, meanwhile, saturates the transition (Fig. 2.9) [85, 86]. This is a standard technique in atomic, molecular, and optical (AMO) physics labs, and works well over a wide range of rf modulation frequencies, which allows for tunable locks. A significant downside of the SAS scheme, though, is residual amplitude modulation (RAM): The background transmission of the probe laser through the vapor cell, in the vicinity of the Doppler-broadened transition saturated by the pump laser, depends on the frequency of the probe. Accordingly, modulating the frequency of the probe laser in order to lock to a narrow Doppler-free transition line results in a background modulation of the amplitude of the probe. The result is in an offset to the lock frequency which, notably, is dependent on laser power. This makes long-term upkeep difficult, as each tune-up of the setup potentially necessitates small changes to the offset frequencies used for locking.

MTS, meanwhile, relies in stead on modulating the frequency of the high-power pump beam that is used to saturate the transition; the low-power probe beam is picked off directly from the laser output before being sent through the atomic vapor cell and detected by a fast photodiode (Fig. 2.10) [87]. This removes the issue of RAM, resulting in a lock frequency that




Figure 2.9: SAS locking setup. The high-power pump beam (solid line) is picked off directly from the laser output and is used to saturate the atomic transition of the sample in the vapor cell. The low-power probe beam (dashed line) is modulated by a double-pass AOM before being sent through the vapor cell, and is detected, after interacting with the sample, by a fast photodiode.

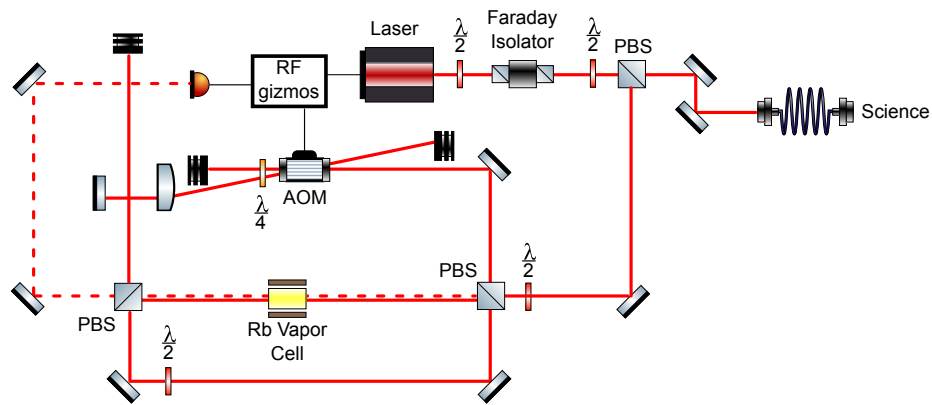


Figure 2.10: MTS locking setup. The high-power pump beam (solid line) is modulated by a double-pass AOM before being used to saturate the atomic transition of the sample in the vapor cell. The low-power probe beam (dashed line) is picked off directly from the laser output and sent through the vapor cell to interact with the sample, then detected by a fast photodiode. Box labeled “RF gizmos” mimics the demodulation and locking chain depicted in Fig. 2.9.

is independent of laser power. The tradeoff is that MTS relies on a higher-order four-wave mixing process, as compared to the simple two-photon process of SAS. Accordingly, the lock features may be smaller, and may depend more strongly on the rf modulation frequencies used. In our case, the benefits of switching to MTS have well outweighed the inconveniences, and I would strongly recommend that any new experiment, or any old experiment considering upgrading their pump–probe locking setups, consider this approach.

2.4.2 Actively stabilized reference cavity

Most of the science that has been done using the E3 apparatus has been in the far-detuned regime $|\Delta_{\text{ca}}| \gg |\Delta_{f'}|$ (see Fig. 1.1), where the precise frequencies of lasers relative to the ^{87}Rb D₂ transition are not critically important, and can be allowed to drift by some tens or hundreds of MHz. The optical cavity that is used to do science (the “science cavity,” for disambiguation), however, is not so forgiving. It is common to need to hold 2–4 lasers at precise frequencies relative to each other and to the bare cavity resonance ω_c . Because the lasers might not all see the bare cavity resonance (due, *e.g.*, to interactions with the atomic sample, see Section 1.3), and because of detection limitations that preclude locking to the science cavity directly, this requires mutually locking the cavity lasers before sending them to the science cavity. To this end, we use an external optical Fabry–Pérot cavity in a separate vacuum chamber (the “reference cavity”). Locking the science cavity to a single laser (the ODT) then guarantees that it will be stable relative to all of the lasers locked to the reference cavity.

Rather than sending the light outputted by the lasers directly to the reference cavity, the lasers are modulated using fiber electro-optical modulators (EOMs), and the resulting sidebands are modulated and locked to the reference cavity using a PDH scheme. This allows the effective frequencies of the lasers to be varied, relative to the reference cavity resonances and to each other, by varying the frequencies of the EOM sidebands.

The reference cavity is passively stable at the 100 kHz scale on timescales shorter than tens of seconds. On longer timescales, the cavity drifts thermally. This is accounted for by measuring the cavity’s temperature using a thermocouple, and feeding back to it using a TEC. The feedback is carried out using a Toptica DTC 110 module. Active temperature feedback accounts for the slow thermal drift of the reference cavity, and maintains the base stability at the 100 kHz scale over long times. For science carried out in the dispersive ($|\Delta_{\text{ca}}| \gg |\Delta_{f'}|$) regime, this is sufficient.

For some recent projects (Section 5.2.2), we have not operated in a relevant dispersive regime, and it has been necessary to lock the science cavity—and all of the light connected to it—to an absolute reference. For this, we use a circular scheme by which we send light from the repumper laser both to a MTS setup using a ^{87}Rb vapor cell and to the reference cavity. The reference cavity is locked to the spectroscopy signal, and the repumper laser is locked to the reference cavity using PDH. The MTS lock is carried out using a Stanford Research Systems (SRS) SR510 lock-in amplifier along with a Newport LB1005-S servo controller and the PDH lock is carried out using a homebuilt lock box. When both locks are simultaneously

Name	Model	Wavelength (nm)	Locking scheme
Cooling	Toptica DL 100	780.248	MTS to ^{87}Rb vapor cell
Repumper	Newport Vortex II	780.235	combined lock (see text)
Pump	Toptica DL Pro	780.3	PDH to reference cavity
ODT	Toptica DL Pro	842	PDH to reference cavity
Ti:Sapph	Coherent 899	780	PDH to reference cavity
Ti:Sapph pump	Coherent Verdi V18	532	free-running
Beatnote	Toptica DL Pro	781.887	beatnote lock to Ti:Sapph

Table 2.1: Laser models, wavelengths, and locking schemes. Given wavelength precisions are representative of the ranges used during regular operation. Specific uses and locking schemes are described in the main text.

active, the reference cavity and repumper laser are both kept stable at the kHz level at a known frequency relative to the chosen transition on the ^{87}Rb D₂ line.

2.4.3 Using a wavemeter as a broadband relative reference

When all of the E3 lasers are well-behaved, the reference cavity (along with the individual spectroscopy setups) maintains a good relative measurement of each of their frequencies. In practice, this is often not the case. External-cavity diode lasers (ECDLs) hop between modes, and occasionally need to be rebuilt all together, either of which requires re-tuning to the correct frequency. In order to carry out this tuning, it is helpful to have a direct (numerical) reference for the absolute laser frequency. This is offered by an industrial wavemeter (High Finesse Wavelength Meter WS-7), which takes fiber-coupled light as an input and measures its primary frequency component. In order to measure the frequencies of multiple lasers simultaneously, homebuild fiber switch (constructed by Sydney Schreppler [88]) takes multiple inputs on different fiber ports and cycles through them in sequence, sending one input at a time to the wavemeter to be measured.

The wavemeter is used to track the frequencies outputted by each of the E3 lasers, as listed in Tab. 2.1.

2.5 Significant experimental setbacks

A common expectation is that each graduate student in experimental AMO physics will spend some significant amount of their time devising and running experiments, some time analyzing experimental results and working through corresponding theoretical problems, and some time designing and building new experimental apparatus. As a student who began graduate school and joined an experiment the construction of which had already been long completed, it was not immediately obvious how I would fulfill the third of these expectations.

Over time, it became clear that the E3 apparatus would require enough significant repair and reconstruction to render this concern moot.

2.5.1 The flood

On the night of Friday, April 26, 2018, the lab directly above E3 sprung a leak. No one was present in either lab at the time, and the leak only became apparent the next morning when someone passing through the hallway outside of E3 noticed a large volume of water pouring out from under the door. By this point, the lab above had flooded completely, and water had worked its way through the gap in the structural concrete below the sink which was meant to allow plumbing to pass between floors—which, in a sense, it had done well. This gap happened to be above the eastern wall of E3, where the main work stations are located, resulting in a fast flow of water directly onto two of the most-used computers, an uninterrupted power supply (UPS), and the three most contemporary lab notebooks, all by way of the inkjet printer, which first imparted a rainbow of muted colors to the flood water.

The process of recovering the lab to working order took around six months. Luckily, the optics tables were largely unaffected. A majority of the damage was to the digital and analog output breakout boards used for computer control of the apparatus (Fig. 2.11). The digital breakouts comprise two stages: direct breakouts from the control computer’s Peripheral Component Interconnect Express (PCIe) board, and switch boards which allow external overriding and control of the digital channels by physical toggle switches. The analog breakouts also comprise two stages: direct breakouts from the control PCIe, and one-way optical isolators which protect from any possibility of downstream faults overvoluting the control computer. Neither the digital nor the analog direct breakouts, both housed in a computer box below the desk, were damaged.

Many channels of the switch boards for the digital outputs failed as a result of corrosion. Replacing these was straightforward, though time-consuming, as the digital channels do not require any particular calibration.

The analog isolator boards were damaged significantly (Fig. 2.12). Rebuilding the boards themselves would have been possible in principle; however, the experimental sequence depends sensitively on their exact transfer functions, and it would have been difficult to replicate these given that calibrations of the original boards from before the flood were not available. This motivated us to salvage the original boards. To do so, we planned to sonicate them in hydrogen peroxide in order to remove the corrosion, and then to replace individual components on a case-by-case basis, only as necessary. Central to each board is the IC containing the one-way optical isolator itself, which we worried might not handle sonication well. These were removed, and were tested while the boards were sonicated in order to ensure that they were still functional. After the boards were sonicated, the optical isolator ICs were replaced and the completed boards were re-installed. Finally, calibrations were recorded for posterity.




Figure 2.11: Racks housing analog and digital control breakouts, during the flood. When the first people arrived at the scene, the racks were quickly covered by a tarp to prevent further damage. The flood water was entering the room from the ceiling, just above and to the left of the frame of the photo.




Figure 2.12: An analog isolator board that was corroded significantly after getting wet while live during the flood. The one-way optical isolator itself is the central integrated circuit (IC) (labeled BB).

2.5.2 The global pandemic, and replacing the rubidium source

In early March, 2020, it became clear that lab work at Berkeley would be affected by the increasingly prominent global COVID-19 pandemic. The effects were significant. From late March through late June, there was no lab access; in this time, without regular tune-ups, enough components of the E3 apparatus became misaligned that no atoms could be detected in the cavity.

From late June, 2020, through March, 2021, lab occupancy was limited to one person at a time. This meant that the recovery process was slow; the E3 apparatus requires significant daily and weekly tune-ups, and, even when spirits are high, it is difficult to make positive net process with only one pair of hands. During this period, spirits were low, and many months passed without seeing evidence of the successful preparation of an atomic sample.

A core issue was that the background pressure of ^{87}Rb in the science chamber (as indicated by measurements of the background-limited lifetime of the initial magnetic trap) was too low. The background vapor used to load the E3 mot has historically been supplied via ^{87}Rb getter [31]. The chamber design includes 4 getters attached semi-rigidly to an electrical feedthrough blank on the lower portion of the main chamber; this offers some redundancy in case one or more getters malfunction or are used up over time. One of the getters appears never to have been functional. Over the course of E3's operation, the other three had been used in sequence, each until they ceased to produce sufficient ^{87}Rb . It became clear in late 2020 that it would be necessary to introduce more ^{87}Rb to the system in order to continue operation.

Replacing the original getters would not be possible. The feedthrough to which the getters are attached is also used to supply current to the in-vacuum quadrupole coil used for the initial MOT. During the initial construction of E3, this feedthrough was attached to the lower chamber while the chamber was open on top, before the installation of the chip assembly and the upper chamber. This meant that the MOT coil could be semi-rigidly connected to the feedthrough in situ. As a result, detaching the MOT coil from the feedthrough and removing the feedthrough, as would be required for replacing the original ^{87}Rb getters, would necessitate disassembling the entire main chamber. It was determined that building a new arm onto the chamber would be preferable. The extra complexity involved in using new getters, combined with unfavorable lead times for vacuum feedthroughs and for the necessary custom mounts, led us to consider switching to using an oven containing a metal ampule as a ^{87}Rb supply.

A ^{87}Rb oven was constructed from a Conflat nipple and connected to the pump arm, outside of an angle valve that was already in place. The new pump + oven arm was baked using standard procedures, with the ^{87}Rb ampule left unbroken and with the angle valve to the main chamber left closed. After the arm cooled to room temperature, the ampule was broken and the angle valve opened to the main chamber. Time passed. A residual gas analyzer (RGA) connected to the pump arm made it clear that ^{87}Rb was present in the upper chamber, but no spectroscopic signal could be measured in the lower (main) chamber.

The explanation, finally, was that the conductivity from the oven to the ion pump (also connected to the upper chamber) was sufficiently greater than the conductivity from the oven to the lower chamber (which required transiting the narrow gap between the copper block and the flange between chamber sections) that very little ^{87}Rb was reaching the main chamber before being removed from the system. It would be necessary to install the oven arm directly to the lower chamber. This would necessitate breaking vacuum to the main chamber and re-baking the entire system.

An oven arm was constructed, comprising a ^{87}Rb ampule inside a Conflat nipple connected via angle valve to one end of a tee, which would be connected directly to the main chamber. On the other end of the tee, an SRS RGA was attached, with the plan to leave it permanently connected to the main chamber.

Space was a limiting factor: of the 8 horizontal viewports on the lower main chamber, 2 are used for the MOT beams, 2 for transverse pumping light (optical imaging, optical pumping, photoassociation), 1 for input to the cavity, and 1 for diagnostic white-light imaging of the system using a security camera (EverFocus EQ150A). Considering the 1 viewport used for the permanent electrical feedthroughs, this left only 1 spot potentially available for mounting an oven: the so-called northwest viewport, currently occupied by an unused window.

The bolts connecting this window to the chamber appeared to have been liberally coated with cold-weld. There was no record of this having been applied, and there was some question as to whether the hard gray substance may actually have been standard vacuum grease that had been baked with the chamber and then allowed to sit for a decade. With no other options available, we moved forward, with the strong hope that it would not be necessary to




Figure 2.13: Temperature readings from across the science chamber during the final, successful bake. Readings are recorded using thermocouples paired with a Pico Technologies PC-08 reader.

chisel epoxy off of the viewport in situ in order to break vacuum.

In preparation for opening the main chamber, we connected a nitrogen tank to the pump arm, between the turbo pump and the roughing pump. The nitrogen was connected through a blowoff valve which would maintain pressure at just over 1 atm. The roughing and turbo pumps were then shut off, allowing nitrogen to flow into the chamber. A positive pressure bag was prepared, which would be held over the flange and the new oven arm while vacuum was broken in order to maintain a nitrogen-rich environment and minimize backflow of air into the chamber. The fully assembled arm was situated inside the bag, as near as possible to its eventual location. The vacuum was broken.

Luckily, the epoxy covering the bolts on the chamber window turned out not, in fact, to be epoxy; removing the window was straightforward, although awkward due to its location. The oven arm was connected to the main chamber, using a copper gasket that was kept warm using a hot plate inside the inert environment of the positive pressure bag. The flange was sealed, the turbo and roughing pumps engaged, and the nitrogen valve closed.

Once the system was evacuated, the chamber was outfitted with thermocouples, wrapped in heater tape and foil, and baked. This process has been described in numerous theses, and the details won't be repeated here. The temperatures of some significant points on the chamber during the bake are recorded in Fig. 2.13.

After the bake, a Titanium sublimator was fired. Before removing the mechanical pump arm, the ion pump was turned on and off so that any buildup released by this and the

sublimation pump could be removed. Some time later, the angle valve connected to the pump arm was closed, and the ion pump was fully powered on. We had vacuum.

Moving forward, the rebuild process was long and difficult; it wasn't until 4 months later that we observed atoms in the cavity. The baking process was also much more difficult than suggested by this short description. In fact, there were two bakes, the first one failing; an earlier attempt at a ^{87}Rb oven failed due to oxidation of the source by a small amount of water in the vacuum system; at one point, a power outage led to the failure of the turbo pump and to backflowing of roughing pump oil into the chamber, the heavy hydrocarbons of which could be detected on the RGA for quite some time. I've avoided detailing many of these issues only in order to save myself the frustration of recalling their details. In the end, 4 months after the final bake, the apparatus was operational, and in better shape than during any other period of my tenure in E3.

2.5.3 The labor strike

On November 14, 2022, graduate student researchers (GSRs) across the UC system went on strike, along with graduate student instructors, postdoctoral researchers, and other academic workers, in response to unfair labor practices by the UC. The strike was successful, and all three of the union locales that were bargaining collectively left with significantly improved contracts. The time lost, however, was also significant.

The final agreement between GSRs and the UC was reached on December 23, six weeks after the start of the strike. Without regular touch-ups during this time, many pieces of the experiment became significantly misaligned. As of the writing of this dissertation, five months after the resolution of the strike, atoms have not yet been measured in the cavity. While the timing of this downtime has been such that it has not directly affected the work presented in this dissertation, the rebuild is proving to be a major setback toward the intermediate-term goals of the experiment, at a great cost of both money and graduate student time.

This is a small price to pay, given the relative success of the strike; still, significant time and money could have been saved had an agreement been reached earlier. Attendance on the picket lines was very high from students, postdocs, and other academic workers in the physics department, and the pro-labor sentiment within the department was clear. Despite this, and despite the high standing of the UC Berkeley physics department—and the corresponding weight that any action by the physics faculty would have held—neither the physics department itself nor any individual faculty members spoke out in favor of the UC ending the strike. By doing so, the faculty and the department may have been able to expedite the process appreciably.

Chapter 3

Spin optodynamics and coherent quantum feedback

Interactions between spins and electromagnetic cavities give rise to a rich variety of physical systems [25, 58, 89–95]. Much of my time in E3 has been spent investigating the family of interactions that can be engineered, in particular, between an optical Fabry–Pérot cavity and the collective (mesoscopic) spin mode of an atomic ensemble.

By loading an ensemble of atoms into a well of the E3 cavity ODT that overlaps with an antinode of the pump mode, the system can be made to depend only quadratically on the motion of the atoms (Section 2.1.2). This leaves only interactions with the internal states of the atoms, which are described, in the on-resonant case, by the symmetrized Tavis–Cummings Hamiltonian, Eqn. 1.7, and in the off-resonant, spin-independent case by Eqn. 1.11. It is of further interest to explore the physics that arises in the remaining category of off-resonant, spin-dependent interactions.

In this chapter, I will describe how the E3 optical cavity can be made to interact with, as well as to nondestructively probe, the collective spin of an atomic ensemble. I will then investigate how the coherent cavity–spin interaction, along with the lossy nature of the cavity mode, can be used to construct a coherent quantum feedback system by which the collective atomic spin is robustly stabilized to a definite energy, conditioned by the externally applied pump field.

3.1 Dispersive interactions between a collective spin and a cavity

The Hamiltonian describing the coupled spin–cavity system can be written, generically, as a sum of cavity, spin, and interaction terms: $\hat{H} = \hat{H}_c + \hat{H}_s + \hat{H}_{\text{int}}$.

In the lab frame, the cavity Hamiltonian is simply given by $\hat{H}_c^{\text{lab}} = \hbar\omega_c \hat{c}^\dagger \hat{c}$. We find it




Figure 3.1: Coupling coefficients describing the interaction given by Eqn. 3.5. At $\Delta_{ca}/2\pi = -33.5$ GHz (experimental value, gray dashed line), the tensor coupling α_2 is negligible, and the scalar and vector coupling coefficients have reached their asymptotic values.

helpful to move to a frame rotating at the frequency ω_p of the cavity pump laser, such that

$$\hat{H}_c = -\hbar\Delta_{pc}\hat{c}^\dagger\hat{c} = -\hbar\Delta_{pc}\hat{n}, \quad (3.1)$$

where $\Delta_{pc} \equiv \omega_p - \omega_c$ is the pump-cavity detuning. Here, the cavity annihilation operator \hat{c} and the photon occupation operator \hat{n} include light of both right- and left-handed circular polarizations. Although the left- and right-handed cavity modes interact differently with the atomic ensemble, their bare energies are approximately degenerate, and here they can be considered together as $\hat{n} \equiv \hat{n}_R + \hat{n}_L$.

For an ensemble of noninteracting atoms indexed i at positions \vec{r}_i and spin projection $f_z^{(i)}$ along the direction of the magnetic field, the spin Hamiltonian is given by

$$\hat{H}_s = \sum_i \hbar\omega_s(\vec{r}_i) \hat{f}_z^{(i)}. \quad (3.2)$$

Here, the local spin precession frequency is given by $\hbar\omega_s(\vec{r}) = g_F\mu_B|\vec{B}(\vec{r})|$, where g_F is the Landé g -factor and μ_B is the Bohr magneton. For a localized ensemble of atoms, the magnetic field is approximately constant, such that this can be rewritten in terms of an average spin precession frequency ω_s and a total spin projection $\hat{F}_z = \sum_i \hat{f}_z^{(i)}$:

$$\hat{H}_s = \hbar\omega_s\hat{F}_z. \quad (3.3)$$

Generically, the interaction between the cavity and atom i is described by

$$\hat{H}_{\text{int}}^{(i)} = \hbar \sum_{g,e} g_{g;e}^+(\vec{r}_i) \hat{c}_+^\dagger \hat{\sigma}_{e;g}^{(i)} \delta_{m+1,m'} + g_{g;e}^-(\vec{r}_i) \hat{c}_-^\dagger \hat{\sigma}_{e;g}^{(i)} \delta_{m-1,m'} + \text{h.c.} \quad (3.4)$$

Here, the summation runs over all possible transitions from the ground-state manifold $g \equiv |f = 2, m\rangle$ to the excited states $e \equiv |f' = 3, m'\rangle$, with polarization-dependent coupling strengths $g_{g,e}^\pm$. When the cavity–atom detuning $|\Delta_{\text{ca}}|$ is large compared to the hyperfine splittings $|\Delta_{f'}|$ in the excited ($f' = 3$) manifold that is being addressed (Fig. 1.1), the excited states can be eliminated. This approximation results in a spin-dependent dispersive interaction Hamiltonian, describing dynamics within the ground-state manifold:

$$\hat{H}_{\text{int}}^{(i)} = \hbar g_c |U(\vec{r}_i)|^2 \left\{ \alpha_0 (\hat{n}_+ + \hat{n}_-) + \alpha_1 (\hat{n}_+ - \hat{n}_-) \hat{f}_k^{(i)} + \alpha_2 \left[(\hat{n}_+ + \hat{n}_-) \left(\hat{f}_k^{(i)} \right)^2 - \hat{c}_- \hat{c}_+ \left(\hat{f}_+^{(i)} \right)^2 - \hat{c}_+ \hat{c}_- \left(\hat{f}_-^{(i)} \right)^2 \right] \right\}, \quad (3.5)$$

where $|U(\vec{r}_i)|^2$ is the local relative intensity of the cavity pump mode, where \hat{c}_\pm are the annihilation operators for left- and right-handed cavity modes, which are approximately degenerate in our system, and where \hat{f}_k and \hat{f}_\pm are the spin operators relative to a quantization axis along the cavity axis \hat{k} . Here, the scalar, vector, and tensor interactions between the spin and the cavity field are described by coupling coefficients $(\alpha_0, \alpha_1, \alpha_2) \rightarrow (2/3, 1/6, 0)$ in the limit of large Δ_{ca} (Fig. 3.1). The calculation of these coefficients is carried out in detail in Ref. [24]. In our system, the atomic ensemble is primarily localized within a single antinode of the cavity pump field, which allows the local cavity field $U(\vec{r}_i)$ to be treated as approximately constant. This leaves

$$\hat{H}_{\text{int}} = \hbar g_c \left\{ \alpha_0 N_a \hat{n} + \alpha_1 (\hat{n}_+ - \hat{n}_-) \hat{F}_k \right\}, \quad (3.6)$$

such that the total system Hamiltonian, in the limit $|\Delta_{\text{ca}}| \gg |\Delta_{f'}|$, is given by

$$\hat{H} = -\hbar \Delta_{\text{pc}} \hat{n} + \hbar \omega_s \hat{F}_z + \hbar g_c \left\{ \alpha_0 N_a \hat{n} + \alpha_1 (\hat{n}_+ - \hat{n}_-) \hat{F}_k \right\}. \quad (3.7)$$

When the cavity is pumped with only right-handed light, this reduces to the stabilization Hamiltonian Eqn. 3.14 used in Section 3.3. When the cavity is pumped with only left-handed light, the sign of the cavity–spin interaction flips.

3.2 Nondestructive readout of atomic spin states

When the externally applied magnetic field is parallel to the cavity axis ($\hat{k} = \hat{z}$), the system Hamiltonian (Eqn. 3.7) for right-handed (σ_-) pump light commutes with the total spin energy \hat{F}_z since this becomes equivalent to the projection \hat{F}_k of the spins along the cavity axis:

$$\hat{H}^- = -\hbar \left(\Delta_{\text{pc}} - \frac{2}{3} g_c N_a + \frac{1}{6} g_c \hat{F}_z \right) \hat{c}^\dagger \hat{c} + \hbar \omega_s \hat{F}_z, \quad (3.8)$$

where the superscript on \hat{H}^- indicates that this Hamiltonian only considers the right-handed cavity mode, and where the coupling coefficients have been replaced with their numerical




Figure 3.2: An offset AOM is used to sweep the frequency of the pump light back and forth across cavity resonance. This is repeated four times: twice with the atomic ensemble present in the cavity (blue), and twice after the atoms have been expelled (red). For the first of each pair of sweeps, right-handed (σ_+) pump light is used; between sweeps, a pair of LCVRs is used to switch the polarization of the pump light to left-handed (σ_-) (gray dashed line). The power of the light transmitted through the cavity is measured either by a SPAD (not shown) or by an optical heterodyne detector. Comparing the resonance conditions with and without atoms present in the cavity offers a direct measure of Δ_N^\pm , which can be used to calculate the final collective spin \hat{F}_z and atom number N_a of the ensemble.

values. If the dispersive shift Δ_N to the cavity resonance condition is measured by comparing the resonance frequencies with and without the presence of atoms, it will be given by

$$\Delta_N^- = -\frac{2}{3}g_c N_a + \frac{1}{6}g_c \langle \hat{F}_z \rangle, \quad (3.9)$$

where the superscript on Δ_N^- indicates that this is the dispersive shift to the right-handed cavity mode. If the atom number N_a were known exactly, this measurement would be sufficient to determine the collective spin energy \hat{F}_z of the ensemble; however, variable atom loss between state preparation and readout makes this impractical.

By pumping the cavity with left-handed light, different information can be acquired. Examining Eqn. 3.7 and, again, considering the case of an external field parallel to the cavity axis, the Hamiltonian can be derived which describes the left-handed (σ_+) cavity mode:

$$\hat{H}^+ = -\hbar \left(\Delta_{pc} - \frac{2}{3}g_c N_a - \frac{1}{6}g_c \hat{F}_z \right) \hat{c}^\dagger \hat{c} + \hbar \omega_s \hat{F}_z, \quad (3.10)$$

which corresponds to a dispersive shift

$$\Delta_N^+ = -\frac{2}{3}g_c N_a - \frac{1}{6}g_c \langle \hat{F}_z \rangle. \quad (3.11)$$

Using a pair of LCVRs at the input and the output of the cavity, the polarization of the light pumping the cavity can be switched rapidly between left- and right-handed without otherwise affecting the detection chain. By measuring the resonance frequencies of each of the polarizations with the atomic ensemble present in the cavity, and then repeating both measurements with the atoms absent, the total atom number and collective spin can be recovered (Fig. 3.2):

$$N_a = -\frac{3}{4} \frac{\Delta_N^+ + \Delta_N^-}{g_c}; \quad (3.12)$$

$$\langle \hat{F}_z \rangle = -3 \frac{\Delta_N^+ - \Delta_N^-}{g_c}. \quad (3.13)$$

The same effect can be achieved by changing the orientation of the magnetic field to $\theta_B = 180^\circ$ between the first and second measurements of Δ_N , such that $\hat{F}_k = -\hat{F}_z$, but this takes too long to be practical due to the self-inductance of the coils used to generate the field. The effect can also be achieved by using an external rf field to drive a π -pulse (or, nearly equivalently, a Landau–Zener sweep) on the collective spin, taking $\hat{F}_z \rightarrow -\hat{F}_z$ between measurements; this approach has been successfully used in the past ([24]), but its dependence on the calibration of the rf drive makes it less appealing than switching the polarization of the pump light.

3.2.1 Simultaneous readout of both circular cavity modes using linearly polarized pump light

The eigenmodes of the atom–cavity system in the presence of an external magnetic field aligned with the cavity axis (Eqn. 3.7) correspond to circular light, and the system experiences being pumped with linear light as being pumped with a sum of its circular components. As a result, the large and small dispersive shifts Δ_N^\pm , which separately affect the modes corresponding to σ_\pm light, can be measured simultaneously.

To achieve this, we pump the cavity with linear light and measure the two resonances corresponding to the two component polarizations. Rather than spatially separating σ_+ and σ_- light at the output of the cavity and detecting them simultaneously, we temporally separate the detection of the two polarizations by sweeping the pump frequency across cavity resonance, first detecting one resonance and then the other. Comparing the offsets Δ_N^\pm of these resonance conditions from the empty-cavity resonance condition (which is, to the relevant level of precision, polarization-independent) allows all of the information present in Fig. 3.2 to be recorded in a single set of sweeps (Fig. 3.3).

3.3 Autonomous stabilization of a cavity–spin system

This section centers around the work which we completed in 2023 and published in Ref. [96], and draws heavily from the text thereof.

As in the case of classical systems, the state and evolution of quantum systems can be tailored by feedback control [97–99]. At a scientific level, the development of a quantum control theory, one that integrates entanglement and non-classical effects of dissipation and measurement, opens a new line of inquiry into non-equilibrium and open quantum systems. At an applied level, feedback control allows quantum devices to operate robustly, mitigating errors in system preparation and calibration as well as decoherence. Feedback control underpins important tasks such as error correction in quantum computation [49, 100–102] and sensing [103–105], entanglement purification [106], quantum state preparation [43, 44, 48], and adaptive measurement [107].

In this section, I discuss a coherent feedback scheme to stabilize the energy of an ensemble of quantum spins. The scheme employs optical backaction in a driven cavity to realize closed-loop autonomous feedback. Under negative-feedback conditions, we observe that cavity spin optodynamics [58, 95] deterministically steer the collective spin toward a steady-state energy that is set by the frequency of the driving optical field. By examining both the light that drives the system and the atomic spins that respond to this drive, we quantify the tuning




Figure 3.3: The cavity is pumped with linearly polarized ($\sigma_+ + \sigma_-$) light, and an offset AOM is used to sweep the frequency of the light back and forth across cavity resonance (only one direction is depicted here). This is repeated two times: once with the atomic ensemble present in the cavity (blue), and once after the atoms have been expelled (red). The power of the light transmitted through the cavity is measured either by a SPAD (not shown) or by an optical heterodyne detector. The eigenmodes of the cavity–atom system corresponding to different polarizations of light experience different dispersive shifts, allowing for simultaneous measurement of Δ_N^+ and Δ_N^- ; together, as in Fig. 3.2, these measurements can be used to calculate the final collective spin \hat{F}_z and atom number N_a of the ensemble.




Figure 3.4: Experimental setup for autonomous spin stabilization. Atoms are localized at an antinode of the intracavity pump field, such that they are coupled symmetrically to the cavity mode. A magnetic field $\vec{B} = B \hat{z}$ is applied at an angle θ_B relative to the cavity axis \hat{k} . The cavity is pumped with circularly polarized (σ_R) light. After interacting with the atomic ensemble, light leaks out of the cavity, and the σ_R portion of the light is detected using a balanced heterodyne scheme.

sensitivity as well as the damping rate of the autonomous feedback system and find close agreement with a theoretical model.

3.3.1 Building the spin stabilization Hamiltonian

When the cavity is pumped with only right-handed (σ_- , relative to the cavity axis) light, the generic Hamiltonian describing the coupled cavity–spin system (Eqn. 3.7), written in a frame rotating at the pump frequency ω_p , reduces to [95]

$$\hat{H} = -\hbar\Delta_{pc}\hat{c}^\dagger\hat{c} + \hbar\omega_s\hat{F}_z + \hbar g_c[\alpha_0 N_a - \alpha_1 \hat{F}_k] \hat{c}^\dagger\hat{c}. \quad (3.14)$$

The bare cavity energy depends only on its resonance frequency, and the bare atomic energy depends only on the magnetic field strength and collective spin state. The effect of the interaction is to shift the cavity energy by an amount that depends on the number of atoms present and on the projection \hat{F}_k of the collective atomic spin onto the cavity axis. Because the atom number N_a is kept constant during these experiments, it is useful to define the effective detuning $\Delta_{\text{set}} \equiv \Delta_{pc} - g_c\alpha_0 N_a$.

The Hamiltonian can be re-written in terms of the more natural coordinates defined by the magnetic field (Fig. 3.4). In particular, for an applied magnetic field at some angle θ_B to the cavity axis \hat{k} (Fig. 3.4a), $\hat{F}_k = \hat{F}_x \sin \theta_B + \hat{F}_z \cos \theta_B$. The total energy \hat{F}_z of the spins varies slowly, and can be considered as a shift to the bare cavity resonance (Fig. 3.4b). We define $\hat{\Delta} \equiv \Delta_{\text{set}} + g_c\alpha_1 \hat{F}_z \cos \theta_B$ to be the mean detuning from cavity resonance, averaged over the fast effects of the Larmor precession. This allows Eqn. 3.14 to be simplified to

$$\hat{H} = -\hbar\hat{\Delta}\hat{c}^\dagger\hat{c} + \hbar\omega_s\hat{F}_z - \hbar g_s\hat{c}^\dagger\hat{c}\hat{F}_x \sin \theta_B, \quad (3.15)$$




Figure 3.5: Schematic of the landscape of frequencies relevant to autonomous spin stabilization. The projection of the collective spin onto the cavity axis, $\langle \hat{F}_k \rangle$, results in a dispersive shift to the cavity resonance. For $\theta_B < 90^\circ$, positive-energy (negative-energy) spin states shift the cavity resonance, on average, by $\Delta > 0$ ($\Delta < 0$). For pump light detuned from the cavity resonance (vertical green line), the Stokes and anti-Stokes sidebands (red and blue lines) are filtered by the cavity, and have different amplitudes, resulting in net energy transfer between the light and the collective spin.

where $g_s \equiv \alpha_1 g_c$ is the cavity–spin coupling. Here, the third term $\hat{F}_x = \frac{1}{2}(\hat{F}_+ + \hat{F}_-)$ is reminiscent of the interaction term in the standard optomechanical Hamiltonian, where it gives rise to Stokes and anti-Stokes processes, exchanging energy between mechanical (here, spin) and optical degrees of freedom [71, 95, 108]. For pump light injected into the cavity at frequency ω_p , this term results in output light detected at $\omega_p \pm \omega_s$.

In free space, the Stokes and anti-Stokes processes would occur at similar rates. When the collective spin is placed inside an optical cavity, however, the light is filtered by the cavity spectrum, and the system scatters preferentially into whichever process produces light nearer to cavity resonance (Fig. 3.4b). This, in turn, shifts \hat{F}_z , which, for a cavity pumped with σ_R light and for $\theta_B < 90^\circ$, brings Δ nearer to resonance (Fig. 3.5). This dependence closes a feedback loop whereby the equilibrium state of the collective spin is determined solely by the pump detuning Δ_{set} . Eventually, the system will reach a resonance condition at $\Delta = 0$, at which point Stokes and anti-Stokes processes will occur at equal rates and \hat{F}_z will remain constant. This stable condition amounts to a final spin state

$$F_z^{\text{set}} = -\frac{\Delta_{\text{set}}}{g_s \cos \theta_B} \quad (3.16)$$

that is conditioned by the pump frequency ω_p (through Δ_{set}), and that, notably, is independent of the initial state of the collective spin and insensitive to many perturbations.

Here, The longitudinal spin \hat{F}_z plays the role of the control variable, and the cavity shift proportional to \hat{F}_z acts as a coherent sensor. The net detuning $\hat{\Delta}$ represents the control system error signal, proportional to the difference between the instantaneous value of \hat{F}_z and F_z^{set} . The final term in Eqn. 3.15 completes the autonomous control system, serving as

the feedback actuator. As described in Refs. [71, 95, 108], the Larmor precessing transverse spin modulates the cavity field intensity through the spin-dependent dispersive interaction. In turn, this modulation, conditioned by cavity dynamics, acts resonantly on the precessing spin and alters its energy.

3.3.2 Analytic model for the autonomous stabilization system

Considering the Hamiltonian Eqn. 3.15, and including terms accounting for pumping into and (non-Hermitian) leakage out of the cavity mode, the cavity field evolves according to

$$\frac{d}{dt}\hat{c} = i(\Delta_{\text{set}} + g_s\hat{F}_k)\hat{c} - \kappa\hat{c} + \kappa\eta, \quad (3.17)$$

where η is the coherent-state amplitude of the field pumping the cavity. Here, the field operator \hat{c} corresponds to the cavity field at frequency ω_p .

Without any coupling to the cavity, $g_s = 0$, and Eqn. 3.17 can be solved directly, giving

$$\hat{c}_0 = \eta \frac{\kappa}{\kappa - i\Delta_{\text{set}}}. \quad (3.18)$$

If the effects of the coupling between the cavity and the collective spin are small, the perturbation to the field can be approximated by

$$\hat{c}(t) = \hat{c}_0 + \hat{c}'(t). \quad (3.19)$$

The collective spin, meanwhile, evolves according to

$$\begin{aligned} \frac{d}{dt}\hat{F}_x &= -\omega_s\hat{F}_y + g_s\hat{c}^\dagger\hat{c}\hat{F}_y \cos\theta_B; \\ \frac{d}{dt}\hat{F}_y &= \omega_s\hat{F}_x + g_s\hat{c}^\dagger\hat{c}(\hat{F}_z \sin\theta_B - \hat{F}_x \cos\theta_B); \\ \frac{d}{dt}\hat{F}_z &= -g_s\hat{c}^\dagger\hat{c}\hat{F}_y \sin\theta_B. \end{aligned} \quad (3.20)$$

For $\omega_s \gg g_s$, as in our experiment, the transverse components admit solutions $\hat{F}_y \propto F_\perp \sin\omega_s t$. For spins precessing at frequency ω_s at polar angle χ and azimuthal angle $\phi = \omega_s t$, the projection of the spin along the cavity axis looks like $\hat{F}_k = F_\perp \sin\theta_B \cos\phi + \hat{F}_z \cos\theta_B$, where $F_\perp \equiv F \sin\chi$ and $F_z \equiv F \cos\chi$. Substituting this into Eqn. 3.17 gives

$$\frac{d}{dt}\hat{c}' = i(\Delta_{\text{set}} + g_s F_\perp \sin\theta_B \cos\omega_s t + g_s \hat{F}_z \cos\theta_B)(\hat{c}_0 + \hat{c}') - \kappa(\hat{c}_0 + \hat{c}') + \kappa\eta. \quad (3.21)$$

To lowest order, it seems reasonable to expect a solution that looks like effective cavity drives at frequencies $\pm\omega_s$ due to the modulation of the bare pump field \hat{c}_0 by the precessing spins. This leads to the ansatz

$$\hat{c}'(t) = \hat{c}_+ e^{i\omega_s t} + \hat{c}_- e^{-i\omega_s t}. \quad (3.22)$$

Plugging this into Eqn. 3.21 and ignoring quickly rotating terms $\sim e^{\pm 2i\omega_s t}$ as well as terms of order $[(g_s/2\kappa)F_\perp \sin \theta_B]^2$ gives

$$\hat{c}_0 = \eta \mathcal{L}(\Delta_{\text{set}} + g_s F_z \cos \theta_B); \quad (3.23)$$

$$\hat{c}_+ = \frac{i}{2} \frac{g_s}{\kappa} F_\perp \sin \theta_B \mathcal{L}(\Delta_{\text{set}} + g_s F_z \cos \theta_B + \omega_s) \hat{c}_0; \quad (3.24)$$

$$\hat{c}_- = \frac{i}{2} \frac{g_s}{\kappa} F_\perp \sin \theta_B \mathcal{L}(\Delta_{\text{set}} + g_s F_z \cos \theta_B - \omega_s) \hat{c}_0. \quad (3.25)$$

Here, $\mathcal{L}(\Delta) \equiv \kappa/(\kappa - i\Delta)$ refers to a Lorentzian line at center frequency Δ with width κ .

It is desirable to find the effect of the cavity field on the spin energy F_z . This is given by Eqn. 3.20, and depends on the instantaneous occupation number $\hat{n} \equiv \hat{c}^\dagger \hat{c}$ of the cavity mode:

$$\begin{aligned} \hat{n} &= \hat{n}_0 + (\hat{c}_+^\dagger \hat{c}_0 + \hat{c}_0^\dagger \hat{c}_-) e^{-i\omega_s t} + (\hat{c}_-^\dagger \hat{c}_0 + \hat{c}_0^\dagger \hat{c}_+) e^{i\omega_s t} \\ &= \hat{n}_0 - \frac{i}{2} \frac{g_s}{\kappa} F_\perp \sin \theta_B \hat{n}_0 \\ &\quad \times \left[\mathcal{L}(-\Delta_{\text{set}} - g_s F_z \cos \theta_B + \omega_s) e^{-i\omega_s t} \right. \\ &\quad - \mathcal{L}(\Delta_{\text{set}} + g_s F_z \cos \theta_B + \omega_s) e^{-i\omega_s t} \\ &\quad + \mathcal{L}(-\Delta_{\text{set}} - g_s F_z \cos \theta_B - \omega_s) e^{i\omega_s t} \\ &\quad \left. - \mathcal{L}(\Delta_{\text{set}} + g_s F_z \cos \theta_B - \omega_s) e^{i\omega_s t} \right] \end{aligned} \quad (3.26)$$

Again, terms of order $(\omega_s/\kappa)^2$, corresponding to second-order sidebands, have been ignored. Using the cycle-averages $\overline{e^{\pm i\phi} \sin \phi} = \pm i/2$, this gives the mean change in energy of the collective spin to be

$$\begin{aligned} \overline{\frac{d}{dt} \langle \hat{F}_z \rangle} &= -\frac{1}{2} F_\perp^2 \sin^2 \theta_B \frac{g_s^2}{\kappa} \langle \hat{c}_0^\dagger \hat{c}_0 \rangle \\ &\quad \times \left[\text{Re } \mathcal{L}(\Delta_{\text{set}} + g_s F_z \cos \theta_B + \omega_s) \right. \\ &\quad \left. - \text{Re } \mathcal{L}(\Delta_{\text{set}} + g_s F_z \cos \theta_B - \omega_s) \right], \end{aligned} \quad (3.27)$$

where the overline indicates time averaging over a Larmor precession cycle. As expected, for $\Delta_{\text{set}} < -g_s F_z \cos \theta_B$, the first Lorentzian term is larger, and F_z decreases; conversely, F_z increases for $\Delta_{\text{set}} > -g_s F_z \cos \theta_B$.

Experimentally, the system can be characterized in a variety of ways. One accessible quality is the damping rate, which measures the speed at which the system exponentially approaches its equilibrium condition when perturbed by a small amount. It is thus interesting to calculate how quickly the system approaches resonance in the limit $g_s F_z \cos \theta_B \rightarrow -\Delta_{\text{set}}$. In the unresolved sideband regime $\omega_s \ll \kappa$, the asymmetry between the sidebands reduces




Figure 3.6: The coupled differential equations Eqn. 3.31 are propagated forward in time, and the results are used as a model function to fit to the measured values of F_z (black triangles) for varying set points Δ_{set} . The output of the fit (gray line) can be used to calculate the sensitivity of the autonomous feedback system to input noise.

to

$$\begin{aligned} \text{Re } \mathcal{L}(\Delta_{\text{set}} + g_s F_z \cos \theta_B + \omega_s) - \text{Re } \mathcal{L}(\Delta_{\text{set}} + g_s F_z \cos \theta_B - \omega_s) \\ \approx -\frac{4\omega_s(\Delta_{\text{set}} + g_s F_z \cos \theta_B)}{\kappa^2}. \end{aligned} \quad (3.28)$$

Using this approximation along with Eqn. 3.27 and Eqn. 3.23, and noting that $\eta^2 = \bar{n}_{\text{max}}$ corresponds to the mean on-resonance photon occupation of the cavity, the damping rate of the system looks like

$$\begin{aligned} \beta &\equiv -\frac{1}{\Delta} \frac{d\Delta}{dt} \\ &= -\frac{g_s \cos \theta_B}{\Delta_{\text{set}} + g_s F_z \cos \theta_B} \overline{\frac{d}{dt} \langle \hat{F}_z \rangle} \\ &= -2F_{\perp}^2 \sin^2 \theta_B \cos \theta_B \frac{g_s^3}{\kappa^3} \omega_s \bar{n}, \end{aligned} \quad (3.29)$$

where $\bar{n} = \bar{n}_{\text{max}} \kappa^2 / (\kappa^2 + [\Delta_{\text{set}} + g_s F_z \cos \theta_B]^2)$ is the true cavity-filtered photon occupation of the cavity. For the parameters used in our experiment (Section 3.3.4), this amounts to a damping rate of $\beta = 1600 \text{ s}^{-1}$. Notably, however, these parameters do not fall well within the low-modulation regime used to approximate Eqn. 3.23.

The inclusion of higher-order terms $[(g_s/2\kappa)F_{\perp} \sin \theta_B]^2$ has the effect of reducing the carrier amplitude found in Eqn. 3.23. In particular, the full expression for the amplitude looks like

$$\hat{c}_0 = \eta \tilde{\mathcal{L}}(0) \times \left[1 + \left(\frac{g_s F_{\perp} \sin \theta_B}{2\kappa} \right)^2 \tilde{\mathcal{L}}(0) [\tilde{\mathcal{L}}(\omega_s) + \tilde{\mathcal{L}}(-\omega_s)] \right]^{-1}, \quad (3.30)$$

where $\tilde{\mathcal{L}}(\nu) \equiv \mathcal{L}(\Delta_{\text{set}} + g_s F_z \cos \theta_B + \nu)$ has been written for brevity. For the parameters used in our experiment (Section 3.3.4), this amounts to a correction factor of 0.7, resulting in a correction factor of 0.5 to \bar{n} and to the final damping rate: $\beta = 790 \text{ s}^{-1}$. The solutions given by Eqn. 3.24 and Eqn. 3.25 still ignore quickly rotating terms corresponding to higher-order sidebands; however, these effects are confirmed experimentally to be small.

In order to simulate the dynamics of the system, $F_z \equiv \langle \hat{F}_z \rangle$ can be treated as a c-number and Eqn. 3.29 can be used to propagate F_z forward in time. In the absence of any dephasing, this treatment can be made complete by requiring that the total spin is conserved, $F_z^2 + F_\perp^2 = 4N_a^2$. Dephasing can be included heuristically by decreasing F_\perp exponentially over time, with a timescale Γ_2 . In practice, this decrease can take many functional forms, but a simple exponential decay captures much of the system dynamics. Simulating the feedback process, then, amounts to propagating two coupled differential equations:

$$\begin{aligned} \frac{d}{dt} F_z &= \beta(F_\perp) F_z; \\ \frac{d}{dt} F_\perp &= -\frac{\beta(F_\perp) F_z^2}{F_\perp} - \Gamma_2 F_\perp. \end{aligned} \tag{3.31}$$

The resulting values of F_z can be used as a model function for least-squares fitting, where Γ_2 , as well as an overall offset to F_z which accounts for systematic measurement errors, are allowed to vary (Fig. 3.6). These fits are used to extract the sensitivities reported in Fig. 3.9.

3.3.3 Preparation of the collective spin ensemble

When the atomic ensemble is transferred from the magnetic trap to the ODT, each atom is in the maximally trappable ($|f, m_f\rangle = |2, 2\rangle$) state. Commonly, as in Section 3.3, it's been desirable to initiate the collective spin in a different state at the start of an experiment. In order to achieve this, we address the collective spin using rf drives, with the chip wires acting as antennas. The drives are calibrated spectroscopically; this section describes the calibration process.

Initially, a drive is applied near the expected resonance frequency, and then the final spin is measured. A new ensemble is prepared and the process is repeated, with the frequency ω_d of the drive varying between shots. A sufficiently low amplitude and drive length are chosen to ensure that the spin will not approach the low-energy pole. This allows the drive to be treated as a small perturbation, and the observed resonance feature to be described by a sinc function. Fitting to the feature provides a measure of the exact spin frequency ω_s (Fig. 3.7a).

Once ω_s is known, the drive frequency is fixed on resonance and the amplitude of the drive is varied from shot to shot. The resulting Rabi flops are observed; fitting with a sin function provides a measure of the Rabi frequency Ω , and, in particular, of the exact drive strength necessary to effect a $\pi/2$ - or π -pulse (Fig. 3.7b).

The calibration of Ω could just as easily be achieved by holding the drive amplitude constant and varying the length of the drive. In practice, this is less robust to small changes




Figure 3.7: The rf drives used to prepare the collective spin are calibrated spectroscopically by initiating the system at the high-energy pole of the collective Bloch sphere, effecting a variable rf drive, and then measuring the final collective spin state (black circles; error bars represent standard errors on the mean). **a.** First, the amplitude of the drive is set to a constant low value and the frequency ω_d is varied from shot to shot. Near resonance, the drive perturbs the spin slightly from the high-energy pole. Fitting a sinc function to the resonance feature reveals the frequency ω_s of the resonance (gray line). **b.** Next, the frequency of the drive is fixed at $\omega_d = \omega_s$, and the amplitude of the drive is varied from shot to shot. This drives Rabi flops between the poles of the collective Bloch sphere. Fitting a sinusoid to the flops reveals the Rabi frequency Ω of the drive (gray line). The length of the drive is held constant at 15 cycles. Dashed gray lines show the maximum and minimum possible spin projections.

in the ω_s (which is set by the external magnetic fields, and can vary slightly from shot to shot). Varying the length of the drive varies the sharpness of the resonance feature (Fig. 3.7a), meaning that moving away from resonance will affect Ω differently. Holding the drive length constant fixes the shape of the resonance feature. Varying the drive amplitude, meanwhile, does not have any similar nonlinear effects, and has generally been our preferred approach.

3.3.4 Experiments measuring the autonomous stabilization of the collective spin

We first confirm experimentally that the spin ensemble is autonomously stabilized to a state determined by the external set point. To this end, the collective spin is initiated to $\langle \hat{F}_z(t=0) \rangle = 0$ using a coherent rf $\pi/2$ -pulse at drive frequency ω_s , such that $\Delta(t=0) = \Delta_{\text{set}}$. The cavity is then pumped with light at a constant Δ_{set} and allowed to evolve. The light emitted by the cavity is measured on a balanced heterodyne detector [53, 54, 109] (Fig. 3.4a), allowing the power in the Stokes and anti-Stokes sidebands to be detected as independent time traces (Fig. 3.8a). The difference in the power of the two sidebands directly measures the instantaneous energy transfer from the pump light to the collective spin. The cumulative sum of this difference measures the total energy $\delta E(t)$ added to the collective spin, leading up to time t . As shown in Fig. 3.8a and b, an initial $\Delta_{\text{set}} < 0$ leads to an enhancement of the anti-Stokes sideband and a net energy transfer $\delta E < 0$, driving the spin to a low energy state, while $\Delta_{\text{set}} > 0$ has the opposite effect.

The spin state achieved after long evolution times under autonomous feedback, at a given Δ_{set} , is shown in Fig. 3.8c. Here, we measure the longitudinal spin by terminating the feedback, reorienting the magnetic field along the cavity axis, and measuring the spin-dependent cavity shift (Section 3.2). For each θ_B , the measured response shows a tuning range, centered about $\Delta_{\text{set}} = 0$, within which the steady-state spin energy varies linearly with Δ_{set} . The sideband-based energy transfer measurements show the same trend as the spin measurements, through the relation $\hbar\omega_s\langle \hat{F}_z \rangle(t) = \delta E(t)$, but are found to be less accurate. By fitting the response curves to an analytical model (discussed below and in Section 3.3.2), we determine the linear sensitivity of the steady-state spin to Δ_{set} near $\Delta_{\text{set}} = 0$. This linear sensitivity (Fig. 3.9), which can equivalently be thought of as an actuator gain of the feedback system, matches well to the prediction of the set-point equation (Eqn. 3.16) for a range of field angles.

Outside the linear tuning range $|\Delta_{\text{set}}| > |2N_a g_s \cos \theta_B|$, one would expect the feedback system to saturate, driving the spin ensemble to one of its extremal energy states. Such a saturated response is observed for $\theta_B \gtrsim 55^\circ$. Here, the cavity field modulation amplitude, proportional to $\sin \theta_B$, is large, driving the spin quickly to its steady state. In contrast, for shallower angles ($\theta_B \lesssim 55^\circ$), the collective spin undergoes dephasing during feedback, reducing its total magnitude to $\|\vec{F}\| < 2N_a$ before the system can reach its steady state.

Having confirmed that our feedback system qualitatively acts as one would expect, we probe its dynamics in two experiments. First, we characterize the system's closed-loop

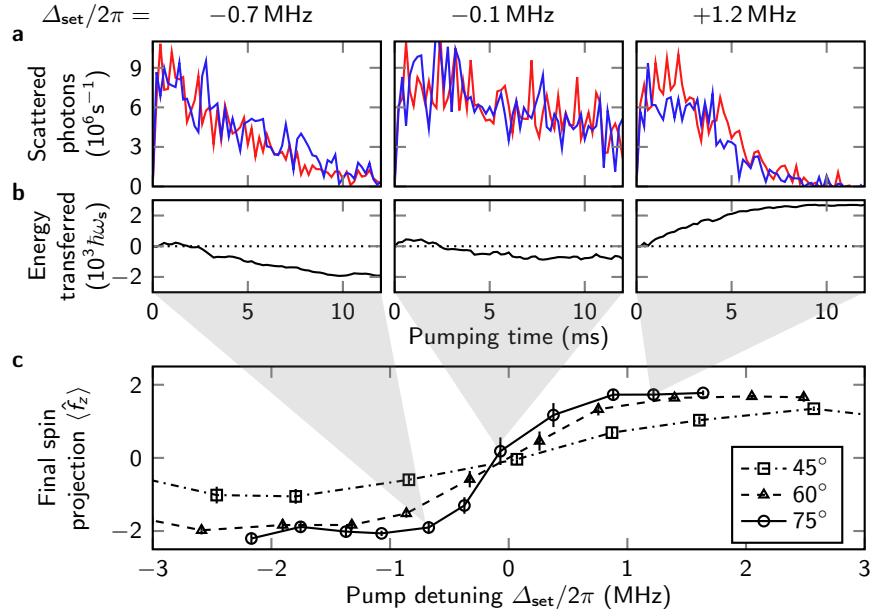


Figure 3.8: Response of the cavity–spin system to a constant pump tone. **a.** Power asymmetry between the Stokes (red) and anti-Stokes (blue) sidebands of the optical cavity emission records the exchange of energy between the cavity field and the spin ensemble. **b.** Integrating this asymmetry yields the net energy transfer $\delta E(t)$ to the spin ensemble vs. time. For $\Delta_{\text{set}} < 0$ (left panel), light is preferentially scattered into the Stokes sideband, pumping energy into the spin mode; for $\Delta_{\text{set}} > 0$ (right panel), energy is pumped out of the spin mode. For $\Delta_{\text{set}} \approx 0$, light is scattered equally into the two sidebands, yielding no net energy transfer. **c.** The final longitudinal spin after 2 ms of feedback is controlled by Δ_{set} . For large $\sin \theta_B$ (e.g., $\theta_B = 75^\circ$, circles; and $\theta_B = 60^\circ$, triangles), the final spin tunes over its entire range, up to $\langle \hat{f}_z \rangle = \pm 2$ for sufficiently large Δ_{set} . For smaller $\sin \theta_B$ (e.g., $\theta_B = 45^\circ$, squares), the collective spin dephases before it can be pulled all the way to either extreme. Error bars represent standard errors on the mean, averaged over 7 repetitions of the experiment.

transfer function by pumping the cavity with a time-varying tone $\Delta_{\text{set}}(t) = \Delta_0 \sin \omega_{\text{mod}} t$ and measuring the response $\delta E(t)$ (Fig. 3.10a). At each modulation frequency, the closed-loop gain is calculated as the ratio between the response and the perturbation:

$$G_{\text{cl}}[\omega_{\text{mod}}] = \frac{2i}{\hbar \Delta_0 T} \int_0^T dt \delta E(t) \exp(-i\omega_{\text{mod}} t), \quad (3.32)$$

where $T = 2\pi s/\omega_{\text{mod}}$, for integer s (Fig. 3.10b, black circles). For a pure integrator system such as ours with damping rate β , we expect a closed-loop gain of

$$G_{\text{cl}}[\omega] = \frac{i(\beta/\omega)}{1 + i(\beta/\omega)}, \quad (3.33)$$




Figure 3.9: The system’s sensitivity to pump detuning, $2\pi d\langle\hat{f}_z\rangle/d\Delta_{\text{set}}|_{\Delta_{\text{set}}=0}$ (black squares: measurement; gray line: Eqn. 3.16 prediction), depends on magnetic field orientation θ_B . Error bars represent fit uncertainties.

which should describe the system well for $\Delta_0 \ll \kappa, |2N_a g_s \cos \theta_B|$ (Fig. 3.10b, gray line). Our data match this expectation qualitatively, but the data quality is limited by a low signal-to-noise ratio and also by saturation at the large set-point modulation depth Δ_0 used for this measurement.

Second, and more quantitatively, we characterize the impulse response function of the feedback system. Here, we initialize the collective spin near $\langle\hat{f}_z\rangle = +1$, and then suddenly impose feedback with a set point of $F_z^{\text{set}}/N_a = -1$ (Fig. 3.11). Time-resolved direct spin measurements track the system evolution toward the set point (Section 3.2).

For regions over which β is approximately constant (namely, $|\langle\hat{f}_z\rangle| \leq 1$), Eqn. 3.29 states that $\langle\hat{F}_z\rangle$ should approach F_z^{set} exponentially. This allows the damping rate of the system to be directly measured, giving a value of $\beta = 450(60)\text{ s}^{-1}$ (Fig. 3.11). For the same experimental parameters ($\theta_B = 60^\circ$, $\omega_s = 2\pi \times 300\text{ kHz}$, $\bar{n}_{\text{max}} = 2.4$, $N_a = 1100$), Eqn. 3.29 predicts $\beta = 1600\text{ s}^{-1}$. This disagreement is due, in part, to the large modulation depth used for this experiment: here, $(g_s/2\kappa)F_\perp \sin \theta_B = 0.4$, which warrants the inclusion of higher-order terms. Accounting for these corrections reduces the expected gain to $\beta = 790\text{ s}^{-1}$ (Section 3.3.2). The analytical model still does not account for the dephasing of the spin ensemble. Constructing an accurate model for dephasing in our system is not straightforward, but any form of dephasing will have the effect of decreasing F_\perp , and thus β , which may explain the remaining discrepancy.

An interesting effect is observed when the pump is modulated with an amplitude $\Delta_0 \gtrsim 2\kappa$. If the collective spin cannot follow the varying set point quickly enough, this makes it possible for the modulation to bring the system sufficiently far from resonance that the cavity population decreases significantly, stopping the stabilization process (Fig. 3.12). This imposes an extra bandwidth restraint on the system, dependent on the amplitude of the modulation. It also introduces a highly nonlinear response: whereas, for $\Delta_0 < \kappa$, the spectrum of the spin response mirrors that of the drive (Fig. 3.10a), the spectrum for larger modulation contains significant second- and third-order components.




Figure 3.10: Response of the autonomous cavity–spin feedback system to impulses. **a.** The energy $\delta E(t)$ transferred to the collective spin reveals that the system responds directly to a time-dependent effective detuning $\Delta_{\text{set}}(t)$. Dotted lines show the response of the cavity with no atoms present, as a baseline for measurement noise. **b.** The measured closed-loop gain of the system is measured using Eqn. 3.32 (black circles). Error bars represent standard errors on the mean, taken over 10 repetitions of the experiment. The expected gain for a perfect pure-integrator system (Eqn. 3.33) is also shown (gray line). All data are recorded at $\theta_B = 60^\circ$.

3.3.5 Conclusions

In this section, I have shown that autonomous feedback generated by optical backaction of a driven cavity onto a spin ensemble stabilizes the ensemble energy at an energy determined by the cavity pump frequency. The optical cavity emission provides a real-time record of the feedback dynamics. In future work, information from this real-time optical signal may also be used to enhance the feedback stabilization through additional measurement-based feedback [110, 111].

This system can equivalently be described as autonomous feedback stabilization of the optical cavity’s resonance frequency. From this viewpoint, the control variable is Δ . The spin ensemble now plays the part of the controller by which the cavity is autonomously tuned to be in resonance with the light with which it is driven.

The feedback setup stabilizes the spin ensemble to a specific value of the longitudinal spin, but does not control the phase at which this spin undergoes Larmor precession because of the time translation symmetry of this scheme. In future work, it will be interesting to investigate methods for fuller control of the quantum spin state, *e.g.*, applying phase coherent modulation at the Larmor frequency, either to the optical pump field or to an

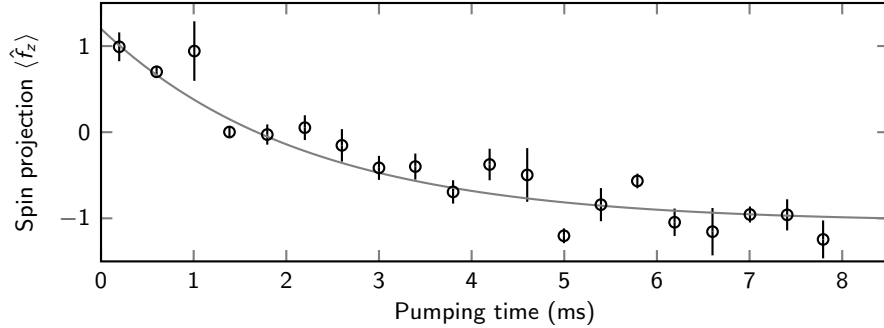


Figure 3.11: Response of the autonomous cavity–spin feedback system to impulses. The collective spin is prepared near $\langle \hat{f}_z \rangle = +1$ and the pump is left on at a constant detuning Δ_{pc} for a varying amount of time, after which the spin state is measured nondestructively (black circles). This allows the spin’s trajectory as the system approaches its equilibrium condition near the $\langle \hat{f}_z \rangle = -1$ to be reconstructed. An exponential fit is used to extract the gain $\beta = 450(60)$ Hz (gray line). Error bars represent standard errors on the mean, taken over 20 repetitions of the experiment. Data are recorded at $\theta_B = 60^\circ$.

applied magnetic field, so as to stabilize the Larmor precession phase.

Another target for future investigation is the fluctuation of the spin ensemble under steady-state feedback. In steady state, the ensemble should respond to the quantum noise of the cavity field, generating fluctuations in the longitudinal spin as well as the Larmor precession phase. At the same time, coherent feedback suppresses longitudinal spin fluctuations. The balance between quantum-optical fluctuations and coherent dissipation, achieved in the steady state and away from thermal equilibrium, may be revealed in the spectrum of the cavity output. However, in the current setup, technical noise on ω_c , the pump light spectrum, and optical detectors obscures this quantum noise signature.

3.4 Decoherence of the collective spin

Modeling the sum of the interactions between the cavity and individual atomic spins as a single interaction term between the cavity and one collective atomic spin greatly reduces the complexity of the system; however, it requires the assumption that each atomic spin sees an identical environment. In practice, this is only approximately the case. The atomic ensemble has some finite spatial extent, and neither external DC magnetic fields nor light modes are perfectly spatially uniform. Further, individual atoms are not stationary, and will each explore different regions of space. This leads to different atomic spins picking up different phases (equivalently, precessing at different rates), reducing the magnitude of the total collective spin over time. This is known as dephasing.

This is observed in Fig. 3.8a, particularly the center column: in the ideal case, the collective spin would continue to precess for arbitrarily long. In practice, the precession




Figure 3.12: Response of the autonomous cavity–spin feedback system to modulations outside of the system bandwidth **a**. The cavity–pump detuning Δ_{set} (solid line) is modulated over a large range $\Delta_0 > 2\kappa$ (dotted lines). **b**. The modulation of Δ_{set} only allows for population of the cavity mode during the brief periods of time when $|\Delta_{\text{set}}| \lesssim |2\kappa|$ (highlighted regions). **c**. When the cavity mode is populated and $\Delta_{\text{set}} > 0$ (red highlighted regions), energy is transferred into the collective spin; when $\Delta_{\text{set}} < 0$ (blue highlighted regions), energy is transferred out of the collective spin. The cavity is populated briefly enough, in total, that the time spent in these two conditions is roughly equal, and there is no net energy transfer, with the system returning approximately to its initial condition at the end of each cycle. All data are recorded at $\theta_B = 60^\circ$.

decays within a short time. From the center column of Fig. 3.8b, it is clear that no energy is lost by the spins during this period, and, accordingly, that the reduction in precession results not from the longitudinal spin decaying to the low-energy pole, but from the transverse spin decaying towards the axis. This is known, in the language of NMR, as T_2 decay [112]. In the example of Fig. 3.8, we observe T_2 decay with $T_2 \approx 3$ ms, with no measurable T_1 decay. In the remainder of this section, I discuss the effects of T_2 decay on our system, how we measure it, and some ways in which it can be avoided.

3.4.1 Measuring and characterizing decoherence

First, an important note on the following figures: The spin dephasing measurements presented in this section (Fig. 3.13, Fig. 3.14, Fig. 3.15, Fig. 3.16) represent a wide variety of different experimental conditions. My intention in presenting them is not to report the specific spin decoherence time observed in E3; rather, it is to present the variety of methods used to make these measurements. The true decoherence time has varied widely over the experimenter’s lifetime, dependent on experimental configurations as well as experimenters’ levels of

motivation. At all times, the methods described below have been useful to characterize the system.

Generally dephasing is measured using a Hahn–Ramsey sequence (Fig. 3.13). This gives a good picture of the spin decoherence time; however, it has often been observed, in E3, that some recorded Ramsey fringes have unexpectedly low contrasts. This leads to some difficulty in interpreting the results, and has led us to seek other methods of measuring dephasing in our system.

While Hahn–Ramsey sequences offer a reliable measure of the spin decoherence time, they rely on accurate calibration of rf drives (even if the method of varying $\delta\tau$ is employed to remove the phase dependence), and they require many experimental realizations in order to gather sufficient statistics while varying T as well as $\delta\tau$. Continuous readout of the collective spin by the optical cavity in E3 affords a convenient alternative. Fig. 3.14 shows an exponential fit to the decay in sideband power after the collective spin is initially prepared on the equator of the collective Bloch sphere using a coherent $\pi/2$ -pulse, corresponding directly to the decay in spin coherence F_{\perp} (the beating is unrelated to the T_2 decoherence, and is described above in Section 3.5).

In E3, the primary cause of dephasing is the spatial dependence of the effective magnetic field due to circularly polarized cavity light (both pump and ODT). Direct comparisons have been made between “bright” and “dark” dephasing (*i.e.*, dephasing with and without pump light present in the cavity) by turning on pump light during the free evolution periods of the Hahn–Ramsey sequence. This, like the standard Hahn–Ramsey sequence, requires many experimental realizations, making it impractical for daily checkups. A less involved comparison can be made by directly viewing the modulation of the cavity by the precessing spins, as in Fig. 3.14, but waiting a variable amount of time T between preparing the collective spin on the equator and beginning to pump the cavity (Fig. 3.15). The decay in the sideband power after the pump light is turned on measures the bright decoherence, and the initial sideband power when the pump light is turned on corresponds to the remaining coherence after dark decoherence for time T . The bright T_2 time is found to be independent of T .

3.4.2 Methods for improving spin coherence

In order to study cavity–spin interactions, it is essential for the timescales of any relevant dynamics to be faster than the T_2 time of the collective spin. For the experiments described above in Section 3.3, this requirement is only met marginally. While increasing the rates at which interesting physics happens is sometimes possible, it is often more straightforward to decrease the dephasing rate. A large portion of my time on E3 has been spent pursuing this.

Tuning the ODT polarization to reduce effective field inhomogeneities

As noted above, a primary source of decoherence in E3 is the spatial dependence of the effective fields due to the cavity light. The effective field due to the pump light cannot




Figure 3.13: Typical spin echo measurement of the collective spin coherence. The collective spin is prepared on the equator using a coherent rf $\pi/2$ -pulse, allowed to evolve for a time $T/2$, reversed with a π -pulse, allowed to evolve for another time $T/2 + \delta\tau$, then brought back to near the pole with a final $\pi/2$ -pulse before the z -component is measured. Varying $\delta\tau$ controls the axis about which the final $\pi/2$ -pulse rotates the collective spin. **a.** Measured Ramsey–Hahn fringes for varying evolution time T (corresponding to varying vertical offset). At each evolution time, the projection of the final collective spin along the z -axis is measured for a variety of $\pi/2$ -pulse phases $\omega_s \delta\tau$ (black circles). Sinusoidal fits are used to extract the contrast of each fringe (gray lines). **b.** The fringe contrast at each evolution time T (gray circles) offers a measurement of the remaining spin coherence. An exponential fit to the decoherence reveals a decay time $T_2 = 2.58 \pm 0.26$ ms (gray line). Error bars represent fit uncertainties. At some early evolution times, fringes have unexpectedly low contrasts (red exes); this is not well-understood.




Figure 3.14: The total power scattered into the Larmor precession sidebands at $\pm\omega_s$ (black line) decreases approximately exponentially over time, indicating the dephasing of the collective atomic spin. Also visible is an overall collapse and revival of the sideband power (which are proportional to the transverse spin F_{\perp}^2) due to the quadratic Zeeman shift (Section 3.5). A fit to the combination of these two terms (red line) reveals a T_2 decay time of 1.85 ± 0.04 ms. This dephasing is observed in the presence of cavity pump light, and is referred to as “bright”.

be avoided, as the parameters of the pump field must be chosen carefully to construct the desired physical systems. The trap, however, will only lead to an energy shift to the atoms if they see circularly polarized light, which can be avoided by carefully tuning the polarization of the trap before it is injected into the cavity.

For a perfectly azimuthally symmetric cavity, it would be sufficient to send the ODT light through a linear polarizer (at any angle) before sending it to the cavity. In practice, linear birefringence in the cavity mirror substrates, as well as in the optics immediately preceding the cavity, results in some change of the ellipticity of the input light which, itself, depends on the axis of the light’s polarization. In order to ensure that the light populating the cavity is linear, then, it is necessary that the light sent towards the cavity is polarized along a particular (unknown) axis with a particular (unknown) ellipticity. Accordingly, an HWP and a QWP are placed in sequence in the ODT path before the pump and the ODT are combined at the cavity input.

The input polarization of the ODT is tuned heuristically by scanning the angle of the HWP at the cavity input and measuring the spin decoherence time at each point (Fig. 3.16, Fig. 3.17). The same process is carried out scanning the QWP, with the HWP at the optimum angle. These steps are repeated until the optima for the two waveplates converge.




Figure 3.15: After initiating the collective spin on the equator, a variable amount of dead time is allowed to pass and then the cavity is pumped and the output observed (colored lines; different colors represent different dead times). During the dead time, the collective spin dephases some amount. The remaining coherence is measured by fitting an exponential decay to the sideband power, as in Fig. 3.14, and extracting the initial amplitude (black circles; error bars represent fit uncertainties). Fitting an exponential decay (gray line) to these data reveals a T_2 decay time of 5.6 ± 0.1 ms. This dephasing is observed in the absence of cavity pump light, and is referred to as “dark”. As in Fig. 3.14, an overall collapse and revival of the sideband power is visible due to the quadratic Zeeman shift (Section 3.5).

Harnessing spin locking to reduce decoherence

Even in the presence of inhomogeneous light fields, coherence of the collective spin can be maintained through a process known as spin locking. This occurs when a system of spins in the presence of a static magnetic (or light) field is strongly driven, resonant with the field, along an axis transverse to the field. In a frame rotating with the drive, inhomogeneities in the static field add in quadrature to the field effected by the drive, reducing the effect of the inhomogeneity [113].

In our system, this occurs when the cavity is driven off-resonance and when the external magnetic field is not oriented along the cavity axis. Here, the modulation of the cavity field by the precessing spin (Fig. 3.8) acts as a (partially transverse) resonant drive, significantly reducing the effects of field inhomogeneity. This has prevented significant dephasing from being observed when the system is kept off-resonance in previous works [58], and allows the system described in Section 3.3 to maintain coherence until the stable condition is reached (Fig. 3.8a and b, left and right columns). In the context of Section 3.3, this spin locking effect can be actively taken advantage of by intermittently reflecting the collective spin about the equator before it is pulled onto resonance; see Fig. 3.18.




Figure 3.16: Spectrograms of the cavity output reveal the varying lifetime of the collective spin as the polarization of the ODT is tuned; each spectrogram corresponds to a different angle of the half-wave plate (HWP) at the cavity input. The Heterodyne frequency is 10 MHz, and the external magnetic field is expected to result in a precession frequency of $\omega_s = 450$ kHz in the absence of any light shift. Each spectrogram is the average of 2–3 experimental realizations.




Figure 3.17: Spin lifetimes are extracted from the spectrograms shown in Fig. 3.16 for varying ODT polarization, as controlled by an HWP at the cavity input. **a.** The lifetime is found to be sharply peaked at an HWP angle of 290° , indicating that the presence of circularly polarized trap light is a significant source of decoherence. **b.** The spin precession frequency is found to depend linearly on the HWP angle near the optimum. This is explained by noting that different circular polarizations of light will result in different signs of light shift to the atomic transitions.

3.5 High-field effects: the quadratic Zeeman shift

Throughout the early iterations of the experiments described in Section 3.3, a collapse and revival was observed in the modulation of the cavity by the precession of the collective spin (Fig. 3.14). This was eventually understood to be due to the quadratic Zeeman shift, and, in particular, to beating between vector and tensor modes of the collective spin.

For alkali atoms such as ^{87}Rb with $j = 1/2$ ground states, the effects of the quadratic Zeeman shift can be calculated directly by observing the Zeeman energy splitting, as predicted by the Breit–Rabi formula [114]

$$\Delta E_z/\hbar = -\omega_{\text{hfs}} \frac{1}{2I(I+1)} + \omega_s m_f \frac{g_I}{g_f} \pm \frac{1}{2} \omega_{\text{hfs}} \sqrt{1 + \frac{2xm}{I+1/2} + x^2}. \quad (3.34)$$

Here, ω_{hfs} is the ground-state hyperfine splitting, $\omega_s = \mu_B g_f |\vec{B}|/\hbar$ is the small-field Larmor precession frequency (ignoring quadratic effects), and the so-called field strength parameter has been defined as

$$x = \frac{g_j - g_I}{f} \frac{\omega_s}{\omega_{\text{hfs}}}. \quad (3.35)$$




Figure 3.18: In the autonomous stabilization configuration of Section 3.3, keeping the atom–cavity system away from its resonance condition with the pump laser by periodically driving the spins greatly decreases decoherence (top row), allowing much more energy to be transferred between the collective spin and the cavity field (bottom row). Left: The collective spin sees T_2 decoherence on a ~ 2 ms timescale in the presence of cavity pump light, greatly limiting the amount of energy that can be transferred between the collective spin and the cavity field. Right: Effecting a π -pulse every 2 ms using an external rf drive increases the coherence time of the system sufficiently by preventing it from reaching resonance, allowing 1–2 orders of magnitude more energy to be transferred.

For intermediate fields $g_f \ll \omega_s \ll \omega_{\text{hfs}}$, it is useful to expand Eqn. 3.34 in x :

$$\Delta E_z/\hbar = \frac{-1 \pm (2I + 1)}{2(2I + 1)}\omega_{\text{hfs}} + \frac{(4 \mp 1)g_I \pm g_j}{4g_f}\omega_s m_f \mp \left(\frac{g_j - g_I}{4g_f} \frac{\omega_s}{\omega_{\text{hfs}}} m_f \right)^2 \omega_{\text{hfs}}. \quad (3.36)$$

For ^{87}Rb , $I = 3/2$, $g_I \approx 0$, $g_j \approx 2$; within the $f = 2$ ground-state hyperfine manifold, $g_f = 1/2$. Putting this together,

$$\Delta E_z/\hbar = \frac{3}{8}\omega_{\text{hfs}} + \omega_s m_f - \left(\frac{\omega_s}{\omega_{\text{hfs}}} m_f \right)^2 \omega_{\text{hfs}}. \quad (3.37)$$

For the $f = 1$ manifold, meanwhile, $g_f = -1/2$, giving

$$\Delta E_z/\hbar = -\frac{5}{8}\omega_{\text{hfs}} + \omega_s m_f + \left(\frac{\omega_s}{\omega_{\text{hfs}}} m_f \right)^2 \omega_{\text{hfs}}. \quad (3.38)$$




Figure 3.19: Measured revival frequencies ω_q for different known spin precession frequencies ω_s (black circles). Revival frequencies are measured by fitting to the total power scattered into the Larmor precession sidebands at $\pm\omega_s$, as in Fig. 3.14; error bars represent fit uncertainties. The expected revival frequency is also shown, with no free parameters (gray line).

Comparing Eqn. 3.37 and Eqn. 3.38, the bare hyperfine splitting and linear Zeeman shift fall out cleanly, leaving the quadratic Zeeman shift

$$\omega_q = \mp \frac{\omega_s^2}{\omega_{\text{hfs}}}, \quad (3.39)$$

where the shift is seen to be negative for states in the $f = 2$ manifold and positive for states in the $f = 1$ manifold.

Here, the spin Hamiltonian Eqn. 3.3 can be rewritten as

$$\hat{H}_s = \hbar\omega_s\hat{F}_z + \hbar\omega_q\hat{F}_z^2 \quad (3.40)$$

to include the quadratic term. The effect of this addition is to couple the vector and tensor components of the collective spin. This results in a collapse and revival of the observed transverse spin, as seen in Fig. 3.14 and Fig. 3.15, with a characteristic beating time given by $\tau_q = \pi/\omega_q$. The beat frequencies can be measured directly by propagating Eqn. 3.40 forward in time and fitting the result, along with a heuristic decoherence term, to the observed transverse spin amplitude (Fig. 3.14). Fig. 3.19 shows clear agreement between the measured beat frequencies and the prediction of Eqn. 3.39.

Interestingly, this measurement also confirms, qualitatively, that the spins being measured are in the $f = 2$ hyperfine manifold. (It would be very concerning if they were not.) Propagating Eqn. 3.40 forward in time and measuring \hat{F}_x at each point shows sharper revivals for larger f , which is understood intuitively as the beating that results in the collapse and revival occurring between more spin components. This is shown for $f = 1$ and $f = 2$ in Fig. 3.20.




Figure 3.20: The beating of the spin coherence $\langle |f_{\perp}|/f \rangle$ varies qualitatively for different f . Plots are generated by propagating Eqn. 3.40 forward in time and measuring \hat{F}_x at each point.

Chapter 4

Real-time tracking and control of atom loss

Evaporative cooling is ubiquitous in the preparation of ultracold atomic gasses. It allows for decreasing temperatures and increasing phase-space densities beyond the limits imposed by Doppler and polarization gradient cooling; rather, it is limited only by the number of atoms available to act as a thermal sink [115, 116].

A gas of interacting atoms in a magnetic or optical trap continuously sees the exchange of energy between its constituent atoms via collisions. As the gas approaches equilibrium, this results in an energy distribution which is dependent on the trap potential as well as the specifics of the intra-atomic interactions, but which is, in general, qualitatively similar to a Maxwell–Boltzmann distribution: a soft peak near the mean energy with a long tail at higher energies. The mean energy depends on the depth of the trap. Reducing the depth of the trap causes some atoms at the high-energy tail of the distribution to escape, reducing the mean energy; over time, as the gas re-equilibrates, this reshapes the energy distribution, bringing its peak to a lower energy as well. This allows the temperature of the gas to be decreased smoothly, at the expense of lost atoms.

For appropriately shaped trapping potentials, as the trap depth is reduced, the decreasing temperature of the gas leads to increased density, which in turn leads to higher collision rates and allows the trap depth to be reduced more rapidly while maintaining thermalization. Once this “runaway” regime is reached, the evaporative cooling process becomes somewhat robust to the specific trajectory with which the trap depth is lowered: Ensembles which become dense too rapidly will thermalize quickly but will lose atoms more slowly, decreasing their evaporation rate and bringing the evaporation process back in tune with the lowering of the trap.

At its core, evaporative cooling is a tool, and it often does its job well without the need for further understanding. Some systems lend themselves more naturally to evaporative cooling than others, though, and detailed theoretical models are helpful when treating (and troubleshooting) systems that require more care. These models have been tested with only moderate thoroughness. By performing highly sensitive, minimally invasive real-time atom

number measurements on samples undergoing evaporative cooling, we have endeavored to test these models more completely.

Previous studies of evaporative cooling have largely focused on average quantities, such as the average rates of atom loss, temperature change, or phase space density increase. Real-time dispersive measurement by our optical cavity of the atom number during evaporative cooling allows for studies of the fluctuations in these quantities, as well as higher-order correlations. A thorough description of evaporative cooling, or of any equilibrium or non-equilibrium thermodynamic process, should be able to describe such fluctuations. Our experimental approach provides a first-time view into these fluctuations, specifically within a mesoscopic (100s of atoms) sample where they are relatively pronounced.

Access to real-time information about the evaporative cooling process also provides an opportunity to apply feedback control, tailoring the final state of the atomic ensemble to experimental needs. A simple use-case for this is offered by the common desire for atomic ensembles to contain similar numbers N_a of atoms between experimental realizations. Evaporative atom loss is a stochastic and uncorrelated (or not highly correlated) process, meaning that, without intervention, shot-to-shot variations in N_a , after the process is complete, will be roughly Poissonian. By feeding back to the trajectory of the evaporative cooling process, however, we find that this distribution can be narrowed significantly.

In this chapter, I will describe how our system can be used to probe the process of evaporative atom loss in real-time, and how the information that we recover can be used to learn about the statistics of evaporative cooling. I will then outline how this real-time readout can be used to close a measurement-based quantum feedback loop, stabilizing the number of atoms present after forced evaporation, as well as their temperature, to predetermined values.

4.1 Dispersive readout of the instantaneous atom number

When an ensemble of atoms is coupled symmetrically to an optical cavity mode which is detuned from atomic resonance, the combined system can be viewed dispersively: the effect of the cavity field on the atoms is an AC Stark shift to their resonance, and the effect of the atoms on the cavity field is dispersive. The phase velocity of light traveling through the atoms is reduced relative to that in free space, which makes the cavity appear longer than its true length, and results in a so-called dispersive shift to its resonance condition.

This can be already seen by examining the Jaynes–Cummings model of a two-level atom (transition frequency ω_a) interacting with a cavity (bare resonance frequency ω_c) with coupling strength g_0 [117]:

$$\hat{H} = \frac{1}{2}\hbar\omega_a\hat{\sigma}_z + \hbar\omega_c\hat{c}^\dagger\hat{c} + \frac{1}{2}\hbar g_0(\hat{\sigma}_+\hat{c} + \hat{\sigma}_-\hat{c}^\dagger). \quad (4.1)$$

Here, $\hat{\sigma}_z \equiv |2\rangle\langle 2| - |1\rangle\langle 1|$ and the raising and lowering operators $\hat{\sigma}_+ \equiv |2\rangle\langle 1|$ and $\hat{\sigma}_- \equiv |1\rangle\langle 2|$ act on the atomic state and \hat{c} is the photon annihilation operator on the cavity field. This

Hamiltonian has eigenenergies

$$\omega_{\pm}(n) = \omega_c \left(n + \frac{1}{2} \right) \pm \frac{1}{2} \sqrt{4g_0^2(n+1) + \Delta_{ca}^2}, \quad (4.2)$$

where n corresponds to the photon occupation number of the cavity mode. At $\Delta_{ca} = 0$, and with no light present in the cavity, this produces the vacuum Rabi splitting $\Delta = 2g_0$ [118–121]. In the far-detuned ($|\Delta_{ca}| \gg g_0\sqrt{N_a}$) regime, meanwhile, the eigenenergies can be rewritten as

$$\omega_{\pm}(n) = \omega_c \left(n + \frac{1}{2} \right) \pm \frac{1}{2} \Delta_{ca} \left[1 + 2 \frac{g_0^2}{\Delta_{ca}^2} (n+1) \right]. \quad (4.3)$$

This amounts to a dispersive shift in the energy per photon of

$$\Delta_1 = \frac{g_0^2}{\Delta_{ca}}. \quad (4.4)$$

When multiple atoms couple with equal strength to the same cavity field, the collective coupling between the cavity and the atomic ensemble is modified as $g = g_0\sqrt{N_a}$ [28, 29], resulting in an N_a -atom dispersive shift of [122–124]

$$\Delta_N = \frac{g_0^2}{\Delta_{ca}} N_a = N_a \Delta_1. \quad (4.5)$$

This can be seen, equivalently, from Eqn. 1.11, and allows the number of atoms present in the cavity to be probed directly by measuring the shift that they effect to the cavity's resonance condition [124].

The cavity resonance frequency $\omega_c + \Delta_N$ can be measured by sweeping the frequency of the pump light across cavity resonance and fitting to the cavity profile (Section 3.2); however, this requires the resonant frequency to be slowly varying, and, in any case, doesn't offer a real-time readout. The cavity frequency can be measured in real time by locking the frequency ω_p of the pump light to the cavity and reading out the output of the feedback controller, which will be linearly related to ω_p , and thus to Δ_N .

The simplest way of locking the pump light to the cavity, and that which we have chosen to implement, is a side-of-fringe (SOF) lock. Light is pumped into the cavity at a constant intensity, and the intensity I of the light escaping from the cavity is monitored. The pump frequency is then conditioned to maintain an output intensity $I = I_0/2$, where I_0 is the peak intensity observed to escape from the cavity on resonance. This corresponds to a detuning of $\kappa/2$ from resonance:

$$\omega_c + \Delta_N = \omega_p \pm \frac{\kappa}{2}, \quad (4.6)$$

where the side of resonance is determined by the sign of the feedback.

An interesting use case for this is to examine evaporative cooling in real time. Fig. 4.1 shows the trajectory of atom loss during forced evaporative cooling from the ODT, as measured using a SOF lock to the cavity. As compared to the standard technique of observing




Figure 4.1: Dispersive readout of the instantaneous atom number, recorded during forced evaporative cooling from the ODT. Noise on individual traces (black) results primarily from detection shot noise, and from the limited bandwidth of the SOF lock of the pump to the cavity. The mean trace (red) shows a smooth decrease in atom number over time.

evaporative cooling by stopping the process after some (variable) time and destructively imaging the ensemble, this has the benefit of allowing the same realization of the ensemble to be observed consistently, allowing for multi-time correlations to be calculated; this is discussed in detail in Section 4.2.

In principle, a higher signal-to-noise ratio (SNR) could be achieved by locking on resonance using the phase quadrature of the heterodyne signal (Fig. 4.2). This would use all of the available information about the pump light’s perturbations from resonance, since the amplitude quadrature of the light does not carry any information about frequency deviations in the immediate vicinity of resonance. Meanwhile, locking to the SOF based only on the amplitude quadrature, as we do in this work, uses only 1/3 of the available information. At the time of this work, the technology for an on-resonance phase lock was not yet available to us. Such a lock has since been designed, implemented, and tested, and details of the lock will be present in a future E3 dissertation.

4.2 Examining the statistics of evaporative cooling

This section centers around the work which we completed in 2020 and published in Ref. [109], and draws heavily from the text thereof.

Tracking out-of-equilibrium dynamical processes and their fluctuations in mesoscopic systems is central to thermodynamics at intermediate scales [125, 126] and transport in solid state systems [127]. For example, current fluctuations in mesoscopic electronic devices reveal the charge quantization of elementary or emergent particles, shedding light on the underlying




Figure 4.2: Sensitivities of the amplitude T and phase ϕ quadratures of the cavity output. Top: Transmission amplitude through (blue) and phase shift due to (orange) the cavity. Bottom: Sensitivities of these quadratures to changes in pump-cavity detuning. At $\Delta_{\text{ca}} = \pm\kappa/2$ (dashed line), information is contained in both quadratures; a SOF lock to the amplitude at this point ignores the information contained in the phase quadrature, leading to a sub-optimal SNR. At $\Delta_{\text{ca}} = 0$, all of the sensitivity to changes in Δ_{ca} is contained in the phase quadrature, and a phase lock to the cavity at this point optimizes SNR.

microscopic physics [128, 129]. Advanced experimental control and precise measurements make ultracold atomic gasses an ideal testbed for studying transport phenomena with solid-state analogs and beyond [130–132]. Furthermore, the achievable system sizes, ranging from single to millions of atoms in different setups, naturally provide access to explore the mesoscopic domain with ensembles of cold atoms.

However, solid-state and ultracold atom mesoscopic systems differ in their fragility against measurement. Solid-state devices are coupled to large thermal reservoirs, which rapidly dissipate the backaction of measurement, and are refreshed with large particle reservoirs. In contrast, ultracold atom systems are well isolated from thermal environments. Technical and backaction disturbance from measurement, such as optical force fluctuations caused by light scattering, is absorbed within the mesoscopic system itself and can change the properties of the system significantly. Thus, the measured fluctuations within a mesoscopic cold-atom system can be strongly altered by continuous or stroboscopic measurements performed on the system. Accessing real-time information in such systems therefore requires strategies to maximize the extracted information for a given heating rate associated with the measurement.

The enhanced atom-light interaction in high-finesse optical cavities [133] provides a means for performing minimally invasive, extremely sensitive measurements on atomic gasses.

Demonstrations span from recording transient signals of single or few atoms passing through an optical cavity [134, 135] to measurements on static and dynamically evolving mesoscopic trapped atomic ensembles [58, 122–124, 136, 137], or the probing of dynamical evolution of novel states of matter realized in the cavity [138–140]. Dispersive atom-number readout, with precision below atomic shot noise and sensitivity reaching down to the single-atom level in mesoscopic ensembles [123], makes cavity-aided measurements particularly interesting for non-invasive dynamical transport measurements in cold gasses [141–143].

In this work, we employ cavity-aided measurements to observe the non-equilibrium process of evaporative cooling [116, 144–147]. Evaporative cooling occurs in a gas of temperature T and atom number N when collisions drive atoms to energies above the finite trap depth U , whereupon they escape the trap, reducing the number of atoms remaining as well as their temperature.

The ensuing dynamics depend on dimensionality, atom number, and temperature of the gas, all features also at the heart of transport phenomena studied with destructive measurements [141]. A simple model captures the interplay between temperature and atom number in an evaporatively cooled atomic ensemble [116]: To evaporate from the trap, atoms have to be collisionally transferred to the high-energy tail of the Maxwell–Boltzmann distribution, such that the average evaporation rate $\dot{N} \propto -\eta e^{-\eta}$, depends exponentially on $\eta = U/k_{\text{B}}T$. This implies that samples with initially higher temperature evaporate atoms more quickly than samples with lower temperature. Moreover, due to the density-dependent thermalization rate of an evaporatively cooled gas [116, 145] and the presence of three-body collisions [148], the evaporation dynamics can be expected to be nonlinear in atom number. The stochastic character of collisions and atom loss enriches this setting through the addition of fluctuations and calls for an in-depth study of the characteristic mesoscopic features of evaporative cooling.

Here, we observe the non-equilibrium dynamics of an ultracold quantum gas during forced evaporation in a tilted trap potential by collecting real-time traces of the atom number dynamics (see Fig. 4.3a). We reveal the interplay of atom number and temperature and their fluctuations during evaporative cooling via two-time correlations of the continuous atom number record. Such two-time correlations also enable us to shed light on the stochastic fluctuations inherent in the evaporative cooling process itself, and discriminate them from fluctuations from one atomic ensemble to another.

The instantaneous atom number is measured using the dispersive shift to cavity resonance (Eqn. 4.5). During the measurement, the probe power in the cavity is held constant at a level characterized by the intracavity photon number $\bar{n} = P/(2\hbar\omega_c\kappa)$, where P is the power transmitted through the cavity. The frequency shift of the cavity resonance is extracted through the feedback signal (Fig. 4.3b). We use knowledge of g [58, 71, 122] and Δ_{ca} to estimate the instantaneous “equivalent atom number” $N(t) = \Delta_N(t)/\Delta_1$ (Fig. 4.3d) where “equivalent” reflects a slight reduction of the vacuum Rabi coupling in our experiment, and is implied if not stated explicitly otherwise.

We monitor the evolution of the intracavity gas during evaporation dynamics initiated by an applied magnetic field gradient (Fig. 4.3a). We quantify the imprecision in the real-time




Figure 4.3: **a** Schematic of the tilted evaporation of atoms (red) with temperature T in the trap potential (blue) with dynamically lowered depth U . Atoms with energies exceeding the trap depth are spilled. **b** The atoms dispersively couple to the cavity, resulting in an atom-number-dependent shift Δ_N of the transmission line shape (blue solid traces) compared to an empty cavity (blue dashed line). Tracking the SOF (marked by red dot and line) provides a dynamical measurement of atom number. **c** The cavity resonance is tracked using a feedback loop involving the cavity-coupled atomic cloud (dark red), located at the intensity maximum of a probe beam in the cavity (light red). The transmitted intensity of probe light is detected using a heterodyne receiver with local oscillator (LO) and a rf power detector (Det). This power is kept constant using a feedback loop (PID) adjusting the frequency of probe and LO through an acousto-optic modulator (AO). The atom number is derived from an in-loop measurement of the voltage used to adjust a voltage controlled oscillator driving the AO. **d** A single unfiltered trace of equivalent atom number N vs. time with the filter procedure indicated (gray shaded areas). **e** The Allan deviation is dominated by photonic shot noise for small integration times τ and by dynamics of the atom number for large τ . A low pass filter of the traces with optimal integration time $\tilde{\tau}$ minimizes the noise associated with both effects. The solid lines represent guides to the eye. Larger photon number \bar{n} (orange: $\bar{n} = 1.9(1)$, blue: $\bar{n} = 3.2(1)$, red: $\bar{n} = 6.0(1)$ and green: $\bar{n} = 9.7(1)$) results in reduced shot noise, but also faster loss of the atoms and therefore increased imprecision at longer integration times. The colored ticks on the right mark the noise level set by Poissonian statistics for our lowest measured mean atom number for each \bar{n} . Graphics adapted from Ref. [109].

measurement for different integration times by the Allan deviation

$$\Delta N(\tau) = \sqrt{\langle (N_\tau(t_{i+1}) - N_\tau(t_i))^2 \rangle}. \quad (4.7)$$

Here, the trace $N_\tau(t_i)$ is obtained by low-pass filtering the full trace $N(t)$ with an integration time τ and then resampling at discrete times t_i separated by time intervals of length τ . The angle brackets denote the average over all i [122, 123]. Photon shot noise limits the measurement precision of the dispersive cavity shift for short integration times, (Fig. 4.3e). For integration times above approximately 1 ms, the dynamics of the evaporation process start dominating the Allan deviation. Choosing such a long integration time leads to a loss of information about the dynamical system under observation. As a consequence, the dynamics set an upper bound on the achievable integration times and therefore suppression of photonic shot noise in the measurement. Despite this, the minimal imprecision and therefore the measurement noise of our cavity-assisted detection is well below the level set by Poissonian fluctuations of size \sqrt{N} for N atoms for all measured traces and all times presented in the following (Fig. 4.3e).

We started our experiment with an atomic cloud of about 2200 ^{87}Rb atoms with a mean temperature of approximately $T_0 = 2.6 \mu\text{K}$. The gas was prepared in its hyperfine state $|F, m_F\rangle = |2, 2\rangle$ and was trapped predominantly in a single well of a far-detuned optical lattice potential with an initial depth $U_0/k_B = 31(1) \mu\text{K}$ in an optical cavity [52, 58, 71]. The lattice provided strong confinement in the z -direction with a trapping frequency of $\omega_z/2\pi = 91(2) \text{kHz}$, putting the gas in the quasi two-dimensional regime $\hbar\omega_z > k_B T_0$. The magnetic field control provided by an atom chip allowed us to compress the atomic cloud in the axial z -direction before loading into the optical cavity. This enabled the accurate positioning of the atomic cloud along the cavity axis at the peak of the probe standing wave, where the atom-cavity coupling was maximized, nearly identical for all atoms, and optomechanical backaction heating was minimal [52, 54]. The cavity length was stabilized such that its resonance frequency was kept at a near-constant red detuning $\Delta_{\text{ca}}/2\pi \approx -42 \text{GHz}$ with respect to the D_2 line of ^{87}Rb , for which the atomic resonance linewidth is $\Gamma/2\pi \approx 6 \text{MHz}$. The detuning of the cavity probe from atomic resonance together with the vacuum Rabi coupling $g/2\pi = 13.1 \text{MHz}$ led to a maximal cavity shift $|\Delta_1/2\pi| \approx 4 \text{kHz}$ per atom. The cavity probe was maintained at a constant detuning of $\Delta_{\text{pc}}/2\pi = \kappa/2\pi = 1.8 \text{MHz}$ from the atom-shifted cavity resonance frequency. The frequency lock of the probe to the cavity was realized by stabilizing the rf power output from a heterodyne receiver monitoring the probe transmission through the cavity (Fig. 4.3c). The large single-atom cooperativity $C = g^2/\kappa\Gamma = 15.9$ and consequently low cavity probe powers of few picowatts in our experiment allowed us to keep the probe-induced off-resonant scattering, and associated heating rate, at a minimal level.

In a first set of experiments, we tracked the equivalent atom number for different intracavity photon numbers \bar{n} , while slowly lowering the trap potential U/k_B from $33 \mu\text{K}$ to approximately $8 \mu\text{K}$ by ramping up a magnetic field gradient within 330 ms. The resulting ensemble of traces is shown in Fig. 4.4, with every individual trace representing a new run of the experiment. The traces taken together form a statistical ensemble that encompasses both




Figure 4.4: Traces of equivalent atom number N vs. observation time t , calculated from the number-dependent cavity shift, for varying intracavity photon number \bar{n} (indicated in the upper right corner). The bright curves represent individual runs of the experiment (approximately 150 per panel). The dark curves are the mean over all such traces. The traces were filtered using their corresponding optimal integration time (Fig. 4.3e). The vertical lines mark the two times $t_2 = 275$ ms ($t_2 = 275$ ms) used to produce the scatter plots shown above the time traces. They illustrate the standardized atom number $N_{2,s}$ measured at the earlier (later) t_2 vs. $N_{1,s}$ measured at $t_1 = 5$ ms. The diagonal (dashed line) and the slope extracted from a linear ordinary least squares fit to the data (dash-dotted line) are indicated in all scatter plots. The corresponding measurement noise is indicated by the ellipse in the center. *Graphics adapted from Ref. [109].*

the variation in evaporation trajectories with different initial atom number and temperature, and also the fluctuations in atom number generated by evaporation dynamics in individual trajectories. In order to minimize the effect of noise, each trace was filtered with a bandwidth corresponding to the optimal integration time $\tilde{\tau}$ extracted from the Allan deviation shown in Fig. 4.3c. For clarity we suppress the subscript and write $N(t) \equiv N_{\tilde{\tau}}(t)$ in the following. Our measurements of $N(t)$ show a clear trend to lower final atom numbers as the intracavity photon number is increased. This trend reflects the larger measurement-induced heating at increasing probe power, which leads to an increase in the number of atoms ejected from the trap during evaporation.

4.2.1 Nonlinearity of the evaporative cooling process

In general, decorrelation of two measured atom numbers $N(t_1) \equiv N_1$ and $N(t_2) \equiv N_2$ at two points in time, indicated by t_1 and t_2 , arises from three sources: technical noise, measurement noise and stochastic noise due to the evaporation process itself. In our experiment, we observe dominant linear correlations of around 97% at all intracavity photon numbers for two measurements closely spaced in time (Fig. 4.4, upper panel). We extract the correlations

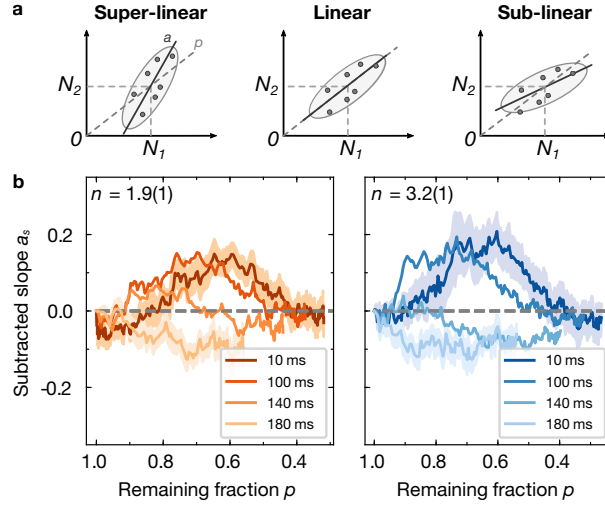


Figure 4.5: **a** Schematic illustrating the nonlinearity measure. Super- (sub-) linearity is equivalent to a larger (smaller) slope a of the scatter of N_2 vs. N_1 than the remaining fraction of atoms p . For a linear process, $a = p$. **b** Subtracted slope $a_s = a - p$ for different initial times t_1 (dark to bright color, indicated in legend) used to quantify the nonlinearity of the evaporation process vs. remaining fraction of atoms p . The shaded region shows the standard deviation calculated by bootstrapping for the smallest and largest initial time. *Graphics adapted from Ref. [109].*

from the slope of the standardized atom numbers $N_{2,s}$ vs. $N_{1,s}$, which equals the Pearson correlation coefficient ρ_{12} . The standardized atom numbers are defined as

$$N_{1(2),s} = \frac{N_{1(2)} - \langle N_{1(2)} \rangle}{\sigma_{1(2)}} \quad (4.8)$$

, where $\langle N_1 \rangle$ ($\langle N_2 \rangle$) and σ_1 (σ_2) are the mean and the standard deviation of the non-standardized atom number distribution measured at time t_1 (t_2). The strong observed correlation indicates a small influence of all noise sources at these early times. In particular, it confirms that our measurement noise is small, consistent with our previous analysis of the imprecision. For a time $t_2 = 275$ ms, later in the evaporation trajectory, the fluctuations are much more prominent. However, there is still some degree of linear correlation present. The reduction of the correlation from 75.9(1)% at the smallest intracavity photon number $\bar{n} = 1.9(1)$ to 35.0(1)% at the largest intracavity photon number $\bar{n} = 9.7(1)$ indicates a larger impact of stochastic noise coupled into the system at larger lost fraction of atoms.

The nonlinear dynamics of a system are captured already by the evolution of statistical averages. In the following, we outline how a continuous measurement of the atom number of the same cloud enables us to shed light on the nonlinear character of evaporative cooling by analyzing two-time correlations. To this end, we compare the remaining fraction of atoms,

$$p = \frac{\langle N_2 \rangle}{\langle N_1 \rangle}, \quad (4.9)$$

to the slope extracted from scatter plots such as those shown in Fig. 4.4, calculated as the least squares estimate

$$a = \rho_{12} \frac{\sigma_2}{\sigma_1}. \quad (4.10)$$

Here, the (two-time) correlation coefficient is given as

$$\rho_{12} = \frac{\text{cov}(N_1, N_2)}{\sigma_1 \sigma_2} \quad (4.11)$$

For a linear process, defined by a constant evaporation rate and $\dot{N} \propto -N$, the two predictions should coincide and hence $a = p$, whereas nonlinear effects lead to a deviation from this expectation (Fig. 4.5a). As an example, the case $\dot{N} \propto -N^\alpha$ with $\alpha > 1$ leads to $a < p$, which we term “*sublinear*” to indicate a scatter slope smaller than in the linear case.

For our evaporation sequence, we find systematic deviations from a simple linear relationship (Fig. 4.5b). Recording the difference $a_s = a - p$ versus the remaining fraction of atoms p , we observe a pronounced superlinear behavior with $a > p$ when referencing to a small initial time t_1 . We interpret this as a manifestation of initial temperature variations of our cloud arising from our sample preparation, which is backed by simulations of the evaporation process. The dynamics of evaporative cooling are governed by an interplay between a reduction of atom number and temperature of a gas, and the latter can have a moderating or accelerating effect on the former. Our observation can be understood as an effective nonlinearity in atom number as a result of the temperature of the gas, which is hidden in an atom number measurement at a single time, but reveals itself in our two-time correlation measurements. Initially hotter clouds will start with slightly fewer atoms due to atom loss before we start our real-time measurement, resulting in an initial anti-correlation between atom number and temperature. The further evaporation process is less efficient in hotter clouds, which therefore spill more atoms than initially colder clouds, leading to $a > p$. The rapid evaporation of hotter gasses also quickly reduces their temperature, such that over time different realizations of the gas arrive at nearly the same final temperature irrespective of their initial temperature. This interesting transient behavior is often implicitly assumed in literature quoting that the temperature locks to a fixed fraction $1/\eta$ of the trap depth [116, 145, 147] and ultimately originates from the competition between exponential truncation and temperature-dependent thermalization of atoms in the evaporation process.

At later times in the evaporation process, after the effect of initial temperature fluctuations is suppressed, we find slightly sublinear behavior, where an initial excess of atoms at time t_1 is reduced during the evolution. This behavior can be explained by increased three-body losses towards the end of the evaporation ramp, which reduce ensemble fluctuations due to their nonlinear character [148]. Interestingly, our observations imply that evaporation first moderates temperature fluctuations and then, thereafter, atom number fluctuations in different realizations, ultimately resulting in a stabilizing effect for both atom number and temperature.




Figure 4.6: The unexplained variance vs. remaining fraction of atoms p for different intracavity photon numbers (color code as in Fig. 4.4 and Fig. 4.5) and two different times t_1 indicated in the top left corner, with the standard deviation calculated by bootstrapping indicated by the shaded region. The solid lines show the model prediction including Poissonian stochastic noise, measurement noise and residual initial ensemble fluctuations, the dashed lines show the prediction without the contribution of Poissonian stochastic noise. *Graphics adapted from Ref. [109].*

4.2.2 Stochastic character of the evaporative cooling process

The mesoscopic nature of our samples together with the high measurement precision enable us to characterize also the fluctuations inherent in evaporative cooling. In addition to measurement and technical noise, fluctuations in the initially prepared ensemble can easily mask such stochastic fluctuations. Initial atom number fluctuations can become significant especially in the mesoscopic regime, reaching up to a point where they can also limit the achievable measurement sensitivity for the study of transport phenomena [141].

Minimally invasive atom number measurements allow for suppressing the effect of fluctuations in the prepared ensemble, and therefore provide a means for an in-depth study of the stochastic properties of evaporative cooling. Concretely, we calculate the unexplained variance

$$\sigma_u^2 = \sigma_2^2 (1 - \rho_{12}^2) \quad (4.12)$$

as the amount of variance in atom number measured at time t_2 that is not explained by correlations with the atom number measured at time t_1 .

For our evaporation sequence, the unexplained variance exhibits a clear peak when referencing to an early measurement at time $t_1 = 10$ ms, see Fig. 4.6, left column. Comparing with a simple parameter-free theoretical model, we find that the fluctuations are up to a factor of three above the fluctuations expected for a purely uncorrelated atom loss, which is described by a Poissonian stochastic process and a spilling rate constant in time. We attribute these large atom number fluctuations to additional initial temperature fluctuations beyond those anti-correlated with the initial atom number, which are then converted into large final atom number fluctuations. The large temperature-induced fluctuations are suppressed strongly during the further evaporation process with decreasing remaining fraction

p (later t_2). Evaluating the unexplained variance starting at a later initial time $t_1 = 180$ ms, we find that for large remaining fraction of atoms (t_2 close to t_1), the unexplained variance is consistent with that expected from the known measurement noise in N_1 and N_2 , see Fig. 4.6, right column. For smaller remaining fraction, other noise sources become dominant, with the increase in σ_u above the measurement noise limit, *i.e.*, the sum of process noise and technical noise, being consistent with the noise expected for a linear, Poissonian stochastic loss process.

4.2.3 Conclusions

The minimally invasive measurement of atom number dynamics in our cavity-coupled atomic gas provides an ideal starting point for further studies of low-dimensional mesoscopic quantum gasses. The densities reached in our two-dimensional system are close to the regime where corrections due to Bose statistics and interactions in the gas become relevant. This calls for further studies focusing on how evaporation dynamics and stochastic fluctuations as well as transport are modified by quantum statistics and atomic interactions, complementing recent studies in three-dimensional gasses [149, 150]. Furthermore, our experiments pave the way for future non-invasive two-terminal transport measurements of strongly correlated quantum gasses in optical cavities [142], for dynamical probing of fluctuation dissipation relations [151, 152], or for realizing novel non-destructive local probes of cold gasses [143]. Finally, our results indicate that feedback on the atom number [153, 154] based on the non-destructive real-time record of an evaporating gas can prepare a samples at a fixed temperature with controllable atom numbers fluctuating less than the amount set by Poissonian statistics. This provides an ideal starting condition to study atom-number-dependent collective phenomena in optical cavities such as dynamical instabilities [71].

4.3 Preparing ensembles with deterministic atom numbers

Forced evaporative cooling [116, 144, 145] is a stochastic process by which the depth of an optical or magnetic trap is steadily reduced, with atoms leaving the trap in a manner that is both uncorrelated in time and uncorrelated between individual atoms. The result of this stochasticity is a probabilistic distribution of how many atoms are lost in a given amount of time from shot to shot. Dynamics upstream of evaporative cooling, technical noise, and the dynamics of evaporative cooling itself all combine to determine the width and shape of this distribution. In ignorance of the details of all of these effects, one might start with a guess that the distribution is Poissonian, with the variance in atom number from shot to shot being roughly equal to the atom number. In our experimental work on monitoring evaporatively gases (Section 4.2), we saw different results: super-Poissonian fluctuations in initial preparation (before evaporative cooling), leading to Poissonian fluctuations after some evaporation. These results are somewhat in contradiction to the presumed role of nonlinear

dynamics and atom loss on the gas; however, the empirical results are clear: the fluctuations are roughly Poissonian.

When studying stochastic processes such as this, the Fano factor offers a useful measure of the deviation from the expected Poissonian nature [155]:

$$F \equiv \frac{\sigma^2}{\mu}, \quad (4.13)$$

where σ^2 is the measured variance in the quantity of interest (here, atom number) between experimental realizations and μ is the measured mean. For a perfect Poissonian process, $\sigma^2 = \mu$ and $F = 1$. Positively correlated losses (corresponding roughly to the “superlinear” case treated in Section 4.2.1) will result in a narrower distribution, $F < 1$. Shot-to-shot systematic noise (*e.g.*, variation in the trap depth, corresponding roughly to the “sublinear” case) will result in a wider distribution, $F > 1$.

Commonly, it is desirable to prepare an atomic sample which comprises a consistent number of atoms between realizations. Optimizing this has traditionally corresponded to minimize systematic noise, aiming to approach the ideal case of $F = 1$. In this section, I examine the possibility of feeding back to the evaporative cooling process in real time in order to reach a fixed atom number set point, and show that it is possible to achieve $F < 1$ in real experimental conditions.

4.3.1 Minimizing shot-to-shot atom number variation

As described in Section 4.1, the dispersive coupling of the atomic ensemble to the pump mode of the optical cavity offers a real-time nondestructive readout of the number of atoms present in the cavity. In Section 4.2, this was used to examine the statistics of atom loss during evaporative cooling, including information about two-time correlations which would not be possible to obtain through destructive measurements. This approach takes advantage of the nondestructive aspect of the measurement, but all of the processing is done after the fact, and does not rely on the real-time nature of the data acquisition.

In practice, real-time knowledge of the number of atoms instantaneously present in the cavity can be used to inform the trajectory of the evaporative cooling. Experimental realizations which are observed to have fewer atoms present, after a certain amount of time of stochastic losses, can be made to lose atoms more slowly by decreasing the rate at which the trap depth is decreased. Conversely, realizations which are observed to have more atoms present can be made to lose atoms more quickly (Section 5.2.1). This closes a measurement-based (*i.e.*, incoherent) quantum feedback loop [36–40], whereby the information leaving the cavity is used by an external controller to stabilize the number of atoms interacting with the cavity pump mode. A similar process of atom number stabilization has been realized previously by interleaving periods of forced rf evaporation with nondestructive phase contrast imaging [154]. Under appropriate conditions, it is possible to reduce shot-to-shot variations in atom number sufficiently to reach $F \ll 1$.




Figure 4.7: Distribution of atom numbers before (red) and after (blue) controlled evaporation with a target atom number of 1400. The atom number is measured nondestructively using the dispersive shift to the cavity resonance at the start and end of each sequence. The narrow width of the distribution of final atom numbers relative to that of initial atom numbers exhibits the ability of controlled evaporation to reduce shot-to-shot variations in the final atomic ensemble.

The simplest form of this feedback loop—and the form that we first considered, *ca.* December, 2018—is to simply ramp the trap depth down at a constant rate to effect evaporation, wait until the measured atom number reaches a certain value, and then abruptly increase the trap depth to prevent any further atom loss. This can be achieved by pumping the cavity with light at a constant frequency ω_p and observing the intensity of light that is transmitted through the cavity. Because the atom–cavity interaction results in a dispersive shift to the resonance condition of the cavity pump mode (Eqn. 4.5), the amount of light that populates the cavity will depend on the number of atoms present. In particular, for $\Delta_N(t=0) < \Delta_{pc} < 0$ (for positive Δ_{ca} , the ordering could be flipped), the loss of atoms will result in an increase in the amount of light transmitted through the cavity. This can be observed in real time, and when a certain power threshold is reached, the trap depth can be increased.

In practice, the approach of triggering based on transmitted intensity works well, but is difficult to standardize and characterize. A more consistent approach is to lock the frequency of the pump light to the SOF of the cavity transmission signal, as in Section 4.2. This is still dependent on the pump light at the cavity input being well intensity-stabilized, but results in an output that is linearly related to the number of atoms present in the cavity, rather than convolved with the cavity susceptibility.

Repeating this procedure over many experimental realizations allows the distribution of final prepared atom numbers to be measured (Fig. 4.7). In particular, by first measuring the initial atom number nondestructively, then using the feedback control scheme to prepare an ensemble with a desired atom number, before finally measuring the final achieved atom number, the shot-to-shot distribution of atom numbers can be compared with and without feedback. This shows clear narrowing of the distribution by the feedback scheme.

4.3.2 Limitations on precise atom number control

The limits on how precisely the atom number of an ensemble can be measured have been explored in a variety of systems [95, 123, 153, 154, 156]. Here, we also face the limitation of how reliably the atom number can be kept constant between the time at which evaporative losses are nominally turned off and the time at which the final number measurement is made. Broadly, we expect these limitations to fit into two categories:

1. Heating of the atomic sample. This will result in atom losses during measurement, broadening the distribution of final atom numbers from shot to shot. Causes for heating include on-resonant scattering and backaction heating.
2. Uncertainty in the readout of the atom number. This will result in the trigger to stop evaporation being sent at inconsistent times, broadening the true distribution of final atom numbers, and will also add errors to our final measurement of the atom number, further broadening the apparent (measured) distribution. Causes for measurement uncertainty include detection shot noise and uncertainty in the bare cavity resonance frequency.

In order to predict how these limitations will affect our system, we examine the signal which is measured by the heterodyne detector:

$$S(\Delta_{\text{pc}}) = 2\kappa\tau\epsilon\bar{n} = 2\kappa\tau\epsilon\frac{\bar{n}_{\text{max}}\kappa^2}{\kappa^2 + (\Delta_N - \Delta_{\text{pc}})^2}, \quad (4.14)$$

where τ is the interrogation time (or, equivalently, the inverse bandwidth of the LPF at the output of the detector). The SOF lock holds this signal at half of its maximum. Inverting this and substituting in Eqn. 4.5 explicitly, the instantaneous atom number can be extracted as

$$N_{\text{a}}(S) = \frac{\Delta_{\text{ca}}}{g_0^2} \left(\kappa \sqrt{\frac{2\epsilon\bar{n}_{\text{max}}\kappa\tau}{S}} - 1 + \Delta_{\text{pc}} \right). \quad (4.15)$$

Detection shot noise

Taking $\bar{n} = \bar{n}_{\text{max}}/2$ and differentiating Eqn. 4.15 with respect to \bar{n} gives an approximation to the sensitivity of changes ΔN in measured atom number to fluctuations $\Delta\bar{n}$ in photon number:

$$\Delta N_{\text{sn}} = \kappa \frac{\Delta_{\text{ca}}}{g_0^2} \frac{\epsilon\bar{n}\kappa\tau}{\sqrt{\epsilon\bar{n}\kappa\tau S} - S^2} \Delta\bar{n} \quad (4.16)$$

The shot noise variations in photon number should look like $\Delta\bar{n}/\sqrt{\bar{n}} = 1/\sqrt{\epsilon\kappa\tau}$. Meanwhile, for our SOF lock of the pump light to the cavity profile, $S = \epsilon\bar{n}_{\text{max}}\kappa\tau$. Together, this allows the sensitivity to shot noise to be simplified to

$$\Delta N_{\text{sn}} = \kappa \frac{\Delta_{\text{ca}}}{g_0^2} \sqrt{\frac{\bar{n}_{\text{max}}}{\epsilon\kappa\tau}}. \quad (4.17)$$

Uncertainty in the bare cavity resonance frequency

In order to measure the number of atoms present before and after carrying out feedback stabilization, we compare the resonant frequency of the atom–cavity system to that of the empty cavity in order to back out the dispersive shift effected by the atomic sample (Section 3.2). Due to the limited bandwidth with which the cavity is locked to the frequency of the ODT, there is a constant noise floor $\Delta\omega_c \approx 2\pi \times 100$ kHz on the measured frequency of the empty cavity resonance, which results in a noise floor on the measured atom number:

$$\Delta N_{\text{ec}} = \Delta\omega_c \frac{\Delta_{\text{ca}}}{g_0^2}. \quad (4.18)$$

On-resonant scattering from the pump

When the cavity is sufficiently far-detuned from the atomic transition being probed, the effective on-resonant scattering rate from the excited state of a single atom (with natural decay rate Γ) is given by

$$\Gamma_{\text{eff}} = \bar{n} \frac{g_0^2}{\Delta_{\text{ca}}} \Gamma, \quad (4.19)$$

where, for a SOF lock, $\bar{n} = \frac{1}{2}\bar{n}_{\text{max}}$. Each scattering event imparts $\hbar\omega_r$ of energy to the atomic sample, where the recoil frequency is given by $\omega_r = \hbar k_{\text{Rb}}^2 / 2m_{\text{Rb}} = 2\pi \times 3.77$ kHz [33]. For an ensemble of N_a atoms being probed for a time τ , this results in an energy $E_{\text{sc}} = \Gamma_{\text{eff}}\hbar\omega_r\tau$ per atom being imparted on the sample. For low Δ_{ca} , where scattering is most relevant, the temperature of the sample appears to be dominated by heating, allowing us to approximate it as $k_{\text{B}}T = E_{\text{sc}}$. Further, approximating the sample as being described by a Boltzmann distribution $N(E) = N_a E_{\text{sc}}^{-1} \exp[-E/E_{\text{sc}}]$, the number of atoms lost from a trap of depth V_0 will then be given by

$$N_{\text{sc}} = \int_{V_0}^{\infty} dE N(E) = N_a e^{-V_0/E_{\text{sc}}}. \quad (4.20)$$

Since this variety of atom loss should be uncorrelated, the fluctuations in atom number that result should have Poissonian statistics:

$$\Delta N_{\text{sc}} = \sqrt{N_{\text{sc}}} = \sqrt{N_a} e^{-V_0/2E_{\text{sc}}}. \quad (4.21)$$

Backaction heating

If the shot noise fluctuations in the amplitude of the cavity field have spectral components near harmonics of the trapping frequency, they can lead to parametric heating. The power spectral density of the photon shot noise is described by

$$S_{\text{sn}}(\omega) = \frac{2\bar{n}_{\text{max}}}{\kappa} \frac{\kappa^2}{\kappa^2 + (\Delta_{\text{pc}} + \omega)^2}. \quad (4.22)$$




Figure 4.8: Shot-to-shot variance in samples with feedback-controlled atom numbers. The achievable Fano factor F has a complex dependence on atom-cavity detuning Δ_{ca} . At large Δ_{ca} , readout noise is expected to limit the precision with which samples can be prepared; whereas, at small Δ_{ca} , the limit will come from heating effects (dashed and solid lines). The atom number variances achieved in a series of experiments matches well with this expectation (black circles; error bars are derived statistically). Prediction lines are calculated using the same experimental parameters as the recorded data.

Treating the mechanical motion of the atoms to act as a simple harmonic oscillator, the expected transition rates Γ_{sn} due to shot noise are calculated perturbatively by [157], and these are converted to heating rates by [31]. Considering only the axial motion in the 1D lattice:

$$\Gamma_{\text{sn}}(\omega) = N_{\text{a}} \left(\frac{g_0^2}{\Delta_{\text{ca}}} \right)^2 \left(\frac{2\omega_{\text{r}}}{\omega} \right)^2 S_{\text{sn}}(2\omega), \quad (4.23)$$

where ω is the axial trapping frequency. When examining the system for a time τ , this results in $E_{\text{ba}} = \Gamma_{\text{sn}} \hbar \omega_{\text{r}} \tau$ per atom of energy being imparted on the sample. As in the case of heating induced by on-resonant scattering (Eqn. 4.20), this results in atom number fluctuations that look like

$$\Delta N_{\text{ba}} = \sqrt{N_{\text{a}}} e^{-V_0/2E_{\text{ba}}}. \quad (4.24)$$

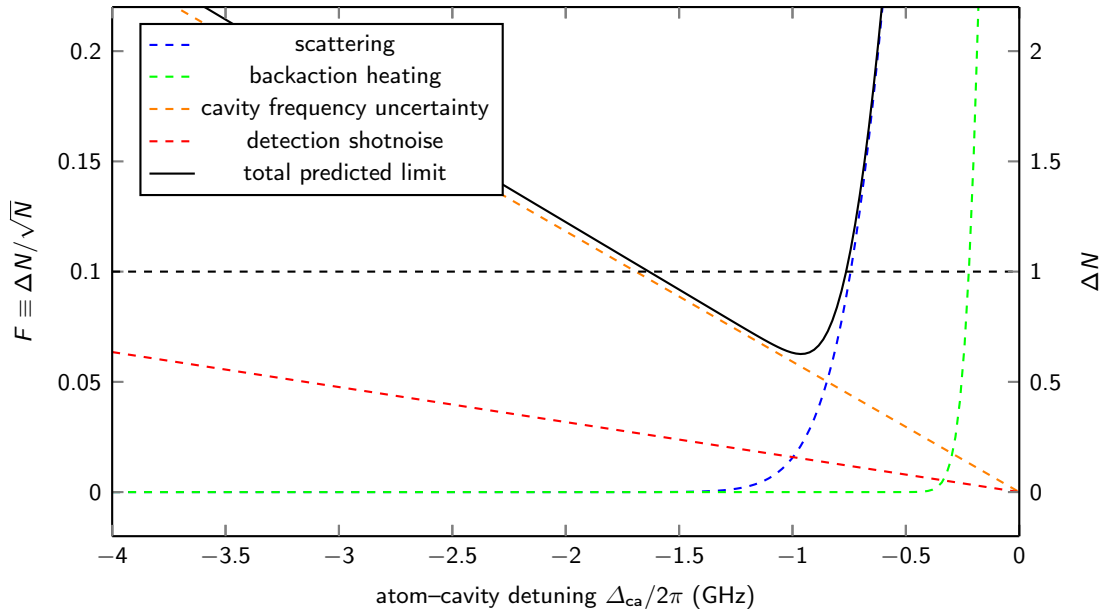


Figure 4.9: Predicted shot-to-shot atom number variance achievable by the proposed feedback system with more favorable experimental conditions. By reducing \bar{n} from 3 to 0.1, decreasing the final desired atom number from 1500 to 100, increasing the trap depth from 1 MHz to 2 MHz, and decreasing the interrogation time τ from 3 ms to 1 ms, we expect to find that in the region $-1.5 \text{ GHz} < \Delta_{\text{ca}}/2\pi < -1 \text{ GHz}$ it is possible to achieve sub-unity shot-to-shot atom number variations. This would correspond to the deterministic generation of samples with exactly $N_{\text{a}} = 100$ atoms.

Comparing measured variances to expectations based on E3 numbers

The total expected shot-to-shot variance in the measured atom number after feedback stabilization can be found by combining Eqn. 4.17, Eqn. 4.18, Eqn. 4.21, and Eqn. 4.24:

$$\Delta N = \sqrt{\left(\kappa \frac{\Delta_{\text{ca}}}{g_0^2}\right)^2 \frac{\bar{n}_{\text{max}}}{\epsilon \kappa \tau} + \left(\Delta \omega_c \frac{\Delta_{\text{ca}}}{g_0^2}\right)^2 + N_{\text{a}} e^{-V_0/E_{\text{sc}}} + N_{\text{a}} e^{-V_0/E_{\text{ba}}}}. \quad (4.25)$$

Fig. 4.8 compares the expected limitations of our system to the measured Fano factors for different coupling strengths (tuned by varying Δ_{ca}). As predicted by Eqn. 4.25, we are able to achieve $F < 1$ for a wide range of detunings $-50 \text{ GHz} < \Delta_{\text{ca}} < -20 \text{ GHz}$. This provides a very positive proof of concept for feedback stabilization of the preparation of atomic ensembles.

Indeed, for appropriately chosen conditions—namely, reduced atom number and cavity occupation, increased trap depth, and reduced interrogation time—Eqn. 4.25 suggests that the shot-to-shot variance in atom number can be pushed significantly farther. Fig. 4.9 shows an example of such conditions. In particular, this suggests that it may be possible to achieve

stabilization to the single-atom level, which provides an enticing goal for future work. An improved atom number readout using the phase locking scheme outlined near the end of Section 4.1 could further broaden the region of parameter space in which these very low shot-to-shot variations could be achieved.

Chapter 5

Closing remarks

In this chapter, I will summarize the techniques and results presented throughout this dissertation. I will also speculate as to what the E3 apparatus might achieve in the remaining months of its operation.

5.1 Summary

My time in the Ultracold Atomic Physics group at Berkeley, and, in particular, in E3, has seen both a wide range of scientific focuses and a wide range of engineering problems.

Using the spin-dependent dispersive interaction between an atomic ensemble and a high-finesse optical cavity, I have spent a considerable amount of time studying coherent quantum feedback (Chapter 3). I built a system to autonomously stabilize the collective spin of the atomic ensemble to a set energy, conditioned on the frequency of light used to pump the cavity; this feedback system was found to be well-described by a simple theoretical model, and to be robust to a variety of perturbations (Section 3.3; Ref. [96]). Along the way, I learned about many sources of decoherence of the collective spin, and about how to combat them (Section 3.4). I also encountered some interesting limiting cases resulting from working in a regime in which the high-field effects of the quadratic Zeeman shift became relevant to the dynamics being observed (Section 3.5).

Using the optical cavity as a tool to make real-time minimally destructive measurements of the instantaneous number of atoms present in the ensemble, I have examined the process of evaporative cooling in detail (Chapter 4). Studying the two-time correlations in atom number, as well as the sources of shot-to-shot variance in atom number under different conditions, offered new insight to the physics of evaporative atom loss, providing a description of the ways in which temperature and number fluctuations map onto each other over time (Section 4.2; Ref. [109]). Using control over the trap depth during evaporation to close a measurement-based feedback loop, I realized a scheme to control the number of atoms present after the evaporative cooling process, reaching shot-to-shot number variations appreciably lower than those achievable with standard forced evaporation (Section 4.3).

5.2 The future of E3

The E3 apparatus is nearing the end of its life span. Before we allow it to stop running, though, there are a few more areas of science that we hope to explore.

5.2.1 Simultaneous feedback stabilization of atom number and temperature

Section 4.3 explores the ability to use real-time readout of the number of atoms present in an ensemble to feed back to the evaporative cooling process, allowing for the preparation of ensembles with minimal shot-to-shot variations in atom number. This work has not yet been completed, and doing so is an immediate priority. In particular, the ability to prepare ensembles with single-atom precision, which may be possible using the current E3 apparatus, would be of fundamental interest to the atomic physics community.

Even taken to this greatest extent, though, the form of feedback presented in Section 4.3 is still very primitive: we allow atoms to be lost until we reach the desired number, and then we close the loss channels. The findings of Section 4.2 suggest that the temperature of an ensemble, as well as its atom number, follows a particular trajectory during the evaporation process. We expect that it should be possible, by feeding back in real time to the specific trajectory with which the trap depth is decreased during evaporation, to control both the atom number and the temperature of an ensemble. This would amount, in practice, to feedback control of the atom number along with simultaneous feedforward control, as informed by Section 4.2, of the temperature. The ability to narrow the shot-to-shot distributions of both the atom number and the temperature of the atomic ensemble between experimental realizations would allow for greatly reduced postprocessing in many experiments.

5.2.2 Cavity-assisted molecular photoassociation

The photoassociation of cold gasses of diatomic molecules—particularly dialkalis—has become commonplace in modern atomic physics. Ultracold molecules are interesting to study in the context of quantum simulation and computation, as their permanent dipole moments allow for controllable long-range interactions. Associating bound molecules from unbound atoms is a two-photon process, with two atoms being stimulated first to an excited (unbound) molecular state and then to their final (bound) molecular state. Using a vacuum cavity mode rather than a free-space laser to excite one of these transitions may prove interesting.

In the context of this cavity-stimulated molecular photoassociation, the cavity can be considered a catalyst, leading to the preferential formation of molecules in a particular ground state. If the cavity is tuned to be near single-photon resonance with the second leg of the two-photon transition, then being in the high-cooperativity regime will mean that the presence of a single molecule in the desired ground state may shift the cavity frequency significantly. This could be used to herald the association of a single molecule. Alternatively, the shift in

the cavity frequency may be used to prevent the stimulated formation of further molecules, allowing for the deterministic formation of single molecules.

The atom–molecule system can also be treated as an effective spin, and can be used to study cavity electromagnetically-induced transparency (EIT) or cavity superradiance. It's not immediately clear what benefits the molecular system would offer, for this use case, over other three-level Raman systems, but exploring the possibilities could be interesting.

So far, we have simulated ^{87}Rb – ^{87}Rb interactions to find the Franck–Condon factors and expected transition rates between different potential excited and bound molecular states; this has allowed us to pick which vibrational levels are most likely to prove useful, and at what frequencies we should expect to find them. Using a tunable refurbished Ti:Sapphire laser, we have successfully used loss spectroscopy to locate the desired up-leg (unbound to excited) transition. By beatnote locking an ECDL to the Ti:Sapphire output to achieve sub-Hz relative phase noise, we have successfully used EIT spectroscopy to locate the desired down-leg (excited to bound) transition. This lays the groundwork for future studies, which will use the E3 apparatus for cavity-assisted photoassociation of dirubidium, and which will be the topic of a future dissertation.

Bibliography

- [1] R. V. Pound, “Electronic frequency stabilization of microwave oscillators”, *Review of Scientific Instruments* **17**, 490 (1946).
- [2] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, “Laser phase and frequency stabilization using an optical resonator”, *Applied Physics B* **31**, 97 (1983).
- [3] E. D. Black, “An introduction to Pound–Drever–Hall laser frequency stabilization”, *American Journal of Physics* **69**, 79 (2001).
- [4] M. H. Devoret, J. M. Martinis, and J. Clarke, “Measurements of macroscopic quantum tunneling out of the zero-voltage state of a current-biased josephson junction”, *Phys. Rev. Lett.* **55**, 1908 (1985).
- [5] Y. Nakamura, C. D. Chen, and J. S. Tsai, “Spectroscopy of energy-level splitting between two macroscopic quantum states of charge coherently superposed by josephson coupling”, *Phys. Rev. Lett.* **79**, 2328 (1997).
- [6] J. E. Mooij, T. P. Orlando, L. Levitov, L. Tian, C. H. van der Wal, and S. Lloyd, “Josephson persistent-current qubit”, *Science* **285**, 1036 (1999).
- [7] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf, “Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics”, *Nature* **431**, 162 (2004).
- [8] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, “Charge-insensitive qubit design derived from the cooper pair box”, *Phys. Rev. A* **76**, 042319 (2007).
- [9] J. Q. You and F. Nori, “Atomic physics and quantum optics using superconducting circuits”, *Nature* **474**, 589 (2011).
- [10] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, “Superconducting qubits: current state of play”, *Annual Review of Condensed Matter Physics* **11**, 369 (2020).
- [11] M. A. Kastner, “Artificial atoms”, *Physics Today* **46**, 24 (1993).
- [12] R. C. Ashoori, “Electrons in artificial atoms”, *Nature* **379**, 413 (1996).

- [13] U. Banin, Y. Cao, D. Katz, and O. Millo, “Identification of atomic-like electronic states in indium arsenide nanocrystal quantum dots”, *Nature* **400**, 542 (1999).
- [14] F. Jelezko, T. Gaebel, I. Popa, A. Gruber, and J. Wrachtrup, “Observation of coherent oscillations in a single electron spin”, *Phys. Rev. Lett.* **92**, 076401 (2004).
- [15] R. J. Epstein, F. M. Mendoza, Y. K. Kato, and D. D. Awschalom, “Anisotropic interactions of a single spin and dark-spin spectroscopy in diamond”, *Nature Physics* **1**, 94 (2005).
- [16] J. Wrachtrup and F. Jelezko, “Processing quantum information in diamond”, *Journal of Physics: Condensed Matter* **18**, S807 (2006).
- [17] R. Hanson, F. M. Mendoza, R. J. Epstein, and D. D. Awschalom, “Polarization and readout of coupled single spins in diamond”, *Phys. Rev. Lett.* **97**, 087601 (2006).
- [18] T. Gaebel, M. Domhan, I. Popa, C. Wittmann, P. Neumann, F. Jelezko, J. R. Rabeau, N. Stavrias, A. D. Greentree, S. Prawer, J. Meijer, J. Twamley, P. R. Hemmer, and J. Wrachtrup, “Room-temperature coherent coupling of single spins in diamond”, *Nature Physics* **2**, 408 (2006).
- [19] M. V. Gurudev Dutt, L. Childress, L. Jiang, E. Togan, J. Maze, F. Jelezko, A. S. Zibrov, P. R. Hemmer, and M. D. Lukin, “Quantum register based on individual electronic and nuclear spin qubits in diamond”, *Science* **316**, 1312 (2007).
- [20] J. R. Maze, P. L. Stanwix, J. S. Hodges, S. Hong, J. M. Taylor, P. Cappellaro, L. Jiang, M. V. G. Dutt, E. Togan, A. S. Zibrov, A. Yacoby, R. L. Walsworth, and M. D. Lukin, “Nanoscale magnetic sensing with an individual electronic spin in diamond”, *Nature* **455**, 644 (2008).
- [21] J. Stenger, S. Inouye, D. M. Stamper-Kurn, H.-J. Miesner, A. P. Chikkatur, and W. Ketterle, “Spin domains in ground-state bose–einstein condensates”, *Nature* **396**, 345 (1998).
- [22] H.-J. Miesner, D. M. Stamper-Kurn, J. Stenger, S. Inouye, A. P. Chikkatur, and W. Ketterle, “Observation of metastable states in spinor bose-einstein condensates”, *Phys. Rev. Lett.* **82**, 2228 (1999).
- [23] D. M. Stamper-Kurn, H.-J. Miesner, A. P. Chikkatur, S. Inouye, J. Stenger, and W. Ketterle, “Quantum tunneling across spin domains in a bose-einstein condensate”, *Phys. Rev. Lett.* **83**, 661 (1999).
- [24] J. L. Kohler, “Optodynamical measurement and coupling of atomic motion and spin”, PhD thesis (UC Berkeley, 2018).
- [25] S. Haroche and J.-M. Raimond, *Exploring the quantum: atoms, cavities, and photons* (Oxford University Press, Oxford, UK, 2006).
- [26] G. Grynberg, A. Aspect, and C. Fabre, *Introduction to quantum optics* (Cambridge University Press, Cambridge, UK, 2010).

- [27] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Photons and atoms: introduction to quantum electrodynamics* (Wiley, Hoboken, NJ, USA, 1997).
- [28] M. Tavis and F. W. Cummings, “Exact solution for an N -molecule—radiation-field Hamiltonian”, *Phys. Rev.* **170**, 379 (1968).
- [29] M. Tavis and F. W. Cummings, “Approximate solutions for an N -molecule-radiation-field Hamiltonian”, *Phys. Rev.* **188**, 692 (1969).
- [30] K. Hepp and E. H. Lieb, “On the superradiant phase transition for molecules in a quantized radiation field: the Dicke maser model”, *Annals of Physics* **76**, 360 (1973).
- [31] T. P. Purdy, “Cavity QED with ultracold atoms on an atom chip”, PhD thesis (UC Berkeley, 2009).
- [32] H. J. Kimble, “Strong interactions of single atoms and photons in cavity QED”, *Physica Scripta* **1998**, 127 (1998).
- [33] D. A. Steck, “Rubidium 87 D line data”, 2010.
- [34] S. Leslie, N. Shenvi, K. R. Brown, D. M. Stamper-Kurn, and K. B. Whaley, “Transmission spectrum of an optical cavity containing N atoms”, *Phys. Rev. A* **69**, 043805 (2004).
- [35] G. F. Franklin, D. Powell, and A. F. Emami-Naeini, *Feedback control of dynamic systems* (Pearson, London, UK, 2019).
- [36] A. C. Doherty, S. Habib, K. Jacobs, H. Mabuchi, and S. M. Tan, “Quantum feedback control and classical control theory”, *Phys. Rev. A* **62**, 012105 (2000).
- [37] C. Sayrin, I. Dotsenko, X. Zhou, B. Peaudecerf, T. Rybarczyk, S. Gleyzes, P. Rouchon, M. Mirrahimi, H. Amini, M. Brune, J.-M. Raimond, and S. Haroche, “Real-time quantum feedback prepares and stabilizes photon number states”, *Nature* **477**, 73 (2011).
- [38] R. Vijay, C. Macklin, D. H. Slichter, S. J. Weber, K. W. Murch, R. Naik, A. N. Korotkov, and I. Siddiqi, “Stabilizing rabi oscillations in a superconducting qubit using quantum feedback”, *Nature* **490**, 77 (2012).
- [39] D. Ristè, C. C. Bultink, K. W. Lehnert, and L. DiCarlo, “Feedback control of a solid-state qubit using high-fidelity projective measurement”, *Phys. Rev. Lett.* **109**, 240502 (2012).
- [40] P. Campagne-Ibarcq, E. Flurin, N. Roch, D. Darson, P. Morfin, M. Mirrahimi, M. H. Devoret, F. Mallet, and B. Huard, “Persistent control of a superconducting qubit by stroboscopic measurement feedback”, *Phys. Rev. X* **3**, 021008 (2013).
- [41] H. Mabuchi, “Coherent-feedback quantum control with a dynamic compensator”, *Phys. Rev. A* **78**, 032323 (2008).
- [42] M. H. Schleier-Smith, I. D. Leroux, and V. Vuletić, “Squeezing the collective spin of a dilute atomic ensemble by cavity feedback”, *Phys. Rev. A* **81**, 021804 (2010).

- [43] K. W. Murch, U. Vool, D. Zhou, S. J. Weber, S. M. Girvin, and I. Siddiqi, “Cavity-assisted quantum bath engineering”, *Phys. Rev. Lett.* **109**, 183602 (2012).
- [44] Z. Leghtas, U. Vool, S. Shankar, M. Hatridge, S. M. Girvin, M. H. Devoret, and M. Mirrahimi, “Stabilizing a Bell state of two superconducting qubits by dissipation engineering”, *Phys. Rev. A* **88**, 023849 (2013).
- [45] D. Ristè, M. Dukalski, C. Watson, G. de Lange, M. Tiggelman, Y. Blanter, K. Lehnert, R. Schouten, and L. DiCarlo, “Deterministic entanglement of superconducting qubits by parity measurement and feedback”, *Nature* **502**, 350 (2013).
- [46] S. Shankar, M. Hatridge, Z. Leghtas, K. M. Sliwa, A. Narla, U. Vool, S. M. Girvin, L. Frunzio, M. Mirrahimi, and M. H. Devoret, “Autonomously stabilized entanglement between two superconducting quantum bits”, *Nature* **504**, 419 (2013).
- [47] Y. Lin, J. P. Gaebler, F. Reiter, T. R. Tan, R. Bowler, A. S. Sørensen, D. Leibfried, and D. J. Wineland, “Dissipative production of a maximally entangled steady state of two quantum bits”, *Nature* **504**, 415 (2013).
- [48] C. K. Andersen, J. Kerckhoff, K. W. Lehnert, B. J. Chapman, and K. Mølmer, “Closing a quantum feedback loop inside a cryostat: autonomous state preparation and long-time memory of a superconducting qubit”, *Phys. Rev. A* **93**, 012346 (2016).
- [49] J. M. Gertler, B. Baker, J. Li, S. Shirol, J. Koch, and C. Wang, “Protecting a bosonic qubit with autonomous quantum error correction”, *Nature* **590**, 243 (2021).
- [50] S. Puri, L. St-Jean, J. A. Gross, A. Grimm, N. E. Frattini, P. S. Iyer, A. Krishna, S. Touzard, L. Jiang, A. Blais, S. T. Flammia, and S. M. Girvin, “Bias-preserving gates with stabilized cat qubits”, *Science Advances* **6**, eaay5901 (2020).
- [51] W. Ma, S. Puri, R. J. Schoelkopf, M. H. Devoret, S. M. Girvin, and L. Jiang, “Quantum control of bosonic modes with superconducting circuits”, *Science Bulletin* **66**, 1789 (2021).
- [52] T. P. Purdy, D. W. C. Brooks, T. Botter, N. Brahms, Z.-Y. Ma, and D. M. Stamper-Kurn, “Tunable cavity optomechanics with ultracold atoms”, *Phys. Rev. Lett.* **105**, 133602 (2010).
- [53] D. W. C. Brooks, T. Botter, S. Schreppler, T. P. Purdy, N. Brahms, and D. M. Stamper-Kurn, “Non-classical light generated by quantum-noise-driven cavity optomechanics”, *Nature* **488**, 476 (2012).
- [54] N. Brahms, T. Botter, S. Schreppler, D. W. C. Brooks, and D. M. Stamper-Kurn, “Optical detection of the quantization of collective atomic motion”, *Phys. Rev. Lett.* **108**, 133601 (2012).
- [55] T. Botter, D. W. C. Brooks, S. Schreppler, N. Brahms, and D. M. Stamper-Kurn, “Optical readout of the quantum collective motion of an array of atomic ensembles”, *Phys. Rev. Lett.* **110**, 153001 (2013).

- [56] S. Schreppler, N. Spethmann, N. Brahms, T. Botter, M. Barrios, and D. M. Stamper-Kurn, “Optically measuring force near the standard quantum limit”, *Science* **344**, 1486 (2014).
- [57] N. Spethmann, J. Kohler, S. Schreppler, L. Buchmann, and D. M. Stamper-Kurn, “Cavity-mediated coupling of mechanical oscillators limited by quantum back-action”, *Nature Physics* **12**, 27 (2016).
- [58] J. Kohler, N. Spethmann, S. Schreppler, and D. M. Stamper-Kurn, “Cavity-assisted measurement and coherent control of collective atomic spin oscillators”, *Phys. Rev. Lett.* **118**, 063604 (2017).
- [59] S. Gupta, K. L. Moore, K. W. Murch, and D. M. Stamper-Kurn, “Cavity nonlinear optics at low photon numbers from collective atomic motion”, *Phys. Rev. Lett.* **99**, 213601 (2007).
- [60] K. W. Murch, K. L. Moore, S. Gupta, and D. M. Stamper-Kurn, “Observation of quantum-measurement backaction with an ultracold atomic gas”, *Nature Physics* **4**, 561 (2008).
- [61] J. A. Gerber, “Cavity quantum electrodynamics with a locally addressable quantum gas”, PhD thesis (UC Berkeley, 2020).
- [62] E. Deist, J. A. Gerber, Y.-H. Lu, J. Zeiher, and D. M. Stamper-Kurn, “Superresolution microscopy of optical fields using tweezer-trapped single atoms”, *Phys. Rev. Lett.* **128**, 083201 (2022).
- [63] E. Deist, Y.-H. Lu, J. Ho, M. K. Pasha, J. Zeiher, Z. Yan, and D. M. Stamper-Kurn, “Mid-circuit cavity measurement in a neutral atom array”, *Phys. Rev. Lett.* **129**, 203602 (2022).
- [64] T. P. Purdy and D. M. Stamper-Kurn, “Integrating cavity quantum electrodynamics and ultracold-atom chips with on-chip dielectric mirrors and temperature stabilization”, *Applied Physics B* **90**, 401 (2008).
- [65] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard, “Trapping of neutral sodium atoms with radiation pressure”, *Phys. Rev. Lett.* **59**, 2631 (1987).
- [66] H. J. Metcalf and P. van der Straten, *Laser cooling and trapping* (Springer, New York, NY, USA, 1999).
- [67] C. J. Foot, *Atomic physics* (Oxford University Press, Oxford, UK, 2005).
- [68] A. Aspect, E. Arimondo, R. Kaiser, N. Vansteenkiste, and C. Cohen-Tannoudji, “Laser cooling below the one-photon recoil energy by velocity-selective coherent population trapping”, *Phys. Rev. Lett.* **61**, 826 (1988).
- [69] J. Dalibard and C. Cohen-Tannoudji, “Laser cooling below the doppler limit by polarization gradients: simple theoretical models”, *J. Opt. Soc. Am. B* **6**, 2023 (1989).

- [70] L. F. Buchmann, S. Schreppler, J. Kohler, N. Spethmann, and D. M. Stamper-Kurn, “Complex squeezing and force measurement beyond the standard quantum limit”, *Phys. Rev. Lett.* **117**, 030801 (2016).
- [71] J. Kohler, J. A. Gerber, E. Dowd, and D. M. Stamper-Kurn, “Negative-mass instability of the spin and motion of an atomic gas driven by optical cavity backaction”, *Phys. Rev. Lett.* **120**, 013601 (2018).
- [72] M. Inguscio, S. Stringari, and C. E. Wieman, eds., *Making, probing and understanding Bose–Einstein condensates*, Vol. CXL, Proceedings of the International School of Physics Enrico Fermi (IOS Press, Amsterdam, 1999), pp. 67–176.
- [73] M. Karski, L. Förster, J. M. Choi, W. Alt, A. Widera, and D. Meschede, “Nearest-neighbor detection of atoms in a 1D optical lattice by fluorescence imaging”, *Phys. Rev. Lett.* **102**, 053001 (2009).
- [74] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr, “Single-atom-resolved fluorescence imaging of an atomic mott insulator”, *Nature* **467**, 68 (2010).
- [75] P. D. Lett, R. N. Watts, C. I. Westbrook, W. D. Phillips, P. L. Gould, and H. J. Metcalf, “Observation of atoms laser cooled below the doppler limit”, *Phys. Rev. Lett.* **61**, 169 (1988).
- [76] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, “A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice”, *Nature* **462**, 74 (2009).
- [77] E. Haller, J. Hudson, A. Kelly, D. A. Cotta, B. Peaudecerf, G. D. Bruce, and S. Kuhr, “Single-atom imaging of fermions in a quantum-gas microscope”, *Nature Physics* **11**, 738 (2015).
- [78] A. Omran, M. Boll, T. A. Hilker, K. Kleinlein, G. Salomon, I. Bloch, and C. Gross, “Microscopic observation of pauli blocking in degenerate fermionic lattice gases”, *Phys. Rev. Lett.* **115**, 263001 (2015).
- [79] M. F. Parsons, F. Huber, A. Mazurenko, C. S. Chiu, W. Setiawan, K. Wooley-Brown, S. Blatt, and M. Greiner, “Site-resolved imaging of fermionic ${}^6\text{Li}$ in an optical lattice”, *Phys. Rev. Lett.* **114**, 213002 (2015).
- [80] G. J. A. Edge, R. Anderson, D. Jervis, D. C. McKay, R. Day, S. Trotzky, and J. H. Thywissen, “Imaging and addressing of individual fermionic atoms in an optical lattice”, *Phys. Rev. A* **92**, 063406 (2015).
- [81] C. Wieman and T. W. Hänsch, “Doppler-free laser polarization spectroscopy”, *Phys. Rev. Lett.* **36**, 1170 (1976).
- [82] G. C. Bjorklund, “Frequency-modulation spectroscopy: a new method for measuring weak absorptions and dispersions”, *Opt. Lett.* **5**, 15 (1980).

- [83] K. L. Corwin, Z.-T. Lu, C. F. Hand, R. J. Epstein, and C. E. Wieman, “Frequency-stabilized diode laser with the zeeman shift in an atomic vapor”, *Appl. Opt.* **37**, 3295 (1998).
- [84] T. Petelski, M. Fattori, G. Lamporesi, J. Stuhler, and G. M. Tino, “Doppler-free spectroscopy using magnetically induced dichroism of atomic vapor: a new scheme for laser frequency locking”, *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics* **22**, 279 (2003).
- [85] P. W. Smith and R. Hänsch, “Cross-relaxation effects in the saturation of the 6328-Å neon-laser line”, *Phys. Rev. Lett.* **26**, 740 (1971).
- [86] K. B. MacAdam, A. Steinbach, and C. Wieman, “A narrow-band tunable diode laser system with grating feedback, and a saturated absorption spectrometer for Cs and Rb”, *American Journal of Physics* **60**, 1098 (1992).
- [87] R. K. Raj, D. Bloch, J. J. Snyder, G. Camy, and M. Ducloy, “High-frequency optically heterodyned saturation spectroscopy via resonant degenerate four-wave mixing”, *Phys. Rev. Lett.* **44**, 1251 (1980).
- [88] S. F. Schreppler, “Quantum measurement with atomic cavity optomechanics”, PhD thesis (UC Berkeley, 2016).
- [89] Q. A. Turchette, C. J. Hood, W. Lange, H. Mabuchi, and H. J. Kimble, “Measurement of conditional phase shifts for quantum logic”, *Phys. Rev. Lett.* **75**, 4710 (1995).
- [90] J. M. Geremia, J. K. Stockton, A. C. Doherty, and H. Mabuchi, “Quantum Kalman filtering and the Heisenberg limit in atomic magnetometry”, *Phys. Rev. Lett.* **91**, 250801 (2003).
- [91] P. E. Barclay, K.-M. C. Fu, C. Santori, and R. G. Beausoleil, “Chip-based microcavities coupled to nitrogen-vacancy centers in single crystal diamond”, *Applied Physics Letters* **95**, 191115, 10.1063/1.3262948 (2009).
- [92] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, “Demonstration of two-qubit algorithms with a superconducting quantum processor”, *Nature* **460**, 240 (2009).
- [93] P. F. Herskind, A. Dantan, J. P. Marler, M. Albert, and M. Drewsen, “Realization of collective strong coupling with ion Coulomb crystals in an optical cavity”, *Nature Physics* **5**, 494 (2009).
- [94] I. D. Leroux, M. H. Schleier-Smith, and V. Vuletić, “Implementation of cavity squeezing of a collective atomic spin”, *Phys. Rev. Lett.* **104**, 073602 (2010).
- [95] N. Brahm and D. M. Stamper-Kurn, “Spin optodynamics analog of cavity optomechanics”, *Phys. Rev. A* **82**, 041804 (2010).
- [96] J. Wolf, O. H. Eibl, J. A. Isaacs, K. Mours, and D. M. Stamper-Kurn, “Autonomous feedback stabilization of a cavity-coupled spin oscillator”, arXiv:2307.04808.

- [97] S. Lloyd, “Coherent quantum feedback”, *Phys. Rev. A* **62**, 022108 (2000).
- [98] H. Mabuchi and N. Khaneja, “Principles and applications of control in quantum systems”, *International Journal of Robust and Nonlinear Control* **15**, 647 (2005).
- [99] J. Zhang, Y. Liu, R. Wu, K. Jacobs, and F. Nori, “Quantum feedback: theory, experiments, and applications”, *Physics Reports* **679**, Quantum feedback: theory, experiments, and applications, 1 (2017).
- [100] P. W. Shor, “Scheme for reducing decoherence in quantum computer memory”, *Phys. Rev. A* **52**, R2493 (1995).
- [101] A. M. Steane, “Error correcting codes in quantum theory”, *Phys. Rev. Lett.* **77**, 793 (1996).
- [102] K. Geerlings, Z. Leghtas, I. M. Pop, S. Shankar, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret, “Demonstrating a driven reset protocol for a superconducting qubit”, *Phys. Rev. Lett.* **110**, 120501 (2013).
- [103] H. Keßler, J. Klinder, M. Wolke, and A. Hemmerich, “Steering matter wave superradiance with an ultranarrow-band optical cavity”, *Phys. Rev. Lett.* **113**, 070404 (2014).
- [104] W. Dür, M. Skotiniotis, F. Fröwis, and B. Kraus, “Improved quantum metrology using quantum error correction”, *Phys. Rev. Lett.* **112**, 080801 (2014).
- [105] S. Zhou and L. Jiang, “Optimal approximate quantum error correction for quantum metrology”, *Phys. Rev. Res.* **2**, 013235 (2020).
- [106] J.-W. Pan, C. Simon, Č. Brukner, and A. Zeilinger, “Entanglement purification for quantum communication”, *Nature* **410**, 1067 (2001).
- [107] M. A. Armen, J. K. Au, J. K. Stockton, A. C. Doherty, and H. Mabuchi, “Adaptive homodyne measurement of optical phase”, *Phys. Rev. Lett.* **89**, 133602 (2002).
- [108] T. J. Kippenberg and K. J. Vahala, “Cavity optomechanics: back-action at the meso-scale”, *Science* **321**, 1172 (2008).
- [109] J. Zeiher, J. Wolf, J. A. Isaacs, J. Kohler, and D. M. Stamper-Kurn, “Tracking evaporative cooling of a mesoscopic atomic quantum gas in real time”, *Phys. Rev. X* **11**, 041017 (2021).
- [110] H. F. Hofmann, G. Mahler, and O. Hess, “Quantum control of atomic systems by homodyne detection and feedback”, *Phys. Rev. A* **57**, 4877 (1998).
- [111] J. Wang and H. M. Wiseman, “Feedback-stabilization of an arbitrary pure state of a two-level atom”, *Phys. Rev. A* **64**, 063810 (2001).
- [112] C. Slichter, *Principles of magnetic resonance* (Springer, New York, NY, USA, 1990).
- [113] M. Lee and W. I. Goldburg, “Nuclear-magnetic-resonance line narrowing by a rotating rf field”, *Phys. Rev.* **140**, A1261 (1965).
- [114] G. Breit and I. I. Rabi, “Measurement of nuclear spin”, *Phys. Rev.* **38**, 2082 (1931).

- [115] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, “Observation of bose-einstein condensation in a dilute atomic vapor”, *Science* **269**, 198 (1995).
- [116] W. Ketterle and N. J. V. Druten, “Evaporative cooling of trapped atoms”, *Advances In Atomic, Molecular, and Optical Physics* **37**, 181 (1996).
- [117] E. T. Jaynes and F. W. Cummings, “Comparison of quantum and semiclassical radiation theories with application to the beam maser”, *Proceedings of the IEEE* **51**, 89 (1963).
- [118] J. P. Reithmaier, G. Sek, A. Löffler, C. Hofmann, S. Kuhn, S. Reitzenstein, L. V. Keldysh, V. D. Kulakovskii, T. L. Reinecke, and A. Forchel, “Strong coupling in a single quantum dot–semiconductor microcavity system”, *Nature* **432**, 197 (2004).
- [119] T. Yoshie, A. Scherer, J. Hendrickson, G. Khitrova, H. M. Gibbs, G. Rupper, C. Ell, O. B. Shchekin, and D. G. Deppe, “Vacuum rabi splitting with a single quantum dot in a photonic crystal nanocavity”, *Nature* **432**, 200 (2004).
- [120] D. Englund, A. Faraon, I. Fushman, N. Stoltz, P. Petroff, and J. Vučković, “Controlling cavity reflectivity with a single quantum dot”, *Nature* **450**, 857 (2007).
- [121] R. Gehr, J. Volz, G. Dubois, T. Steinmetz, Y. Colombe, B. L. Lev, R. Long, J. Estève, and J. Reichel, “Cavity-based single atom preparation and high-fidelity hyperfine state readout”, *Phys. Rev. Lett.* **104**, 203602 (2010).
- [122] N. Brahms, T. P. Purdy, D. W. C. Brooks, T. Botter, and D. M. Stamper-Kurn, “Cavity-aided magnetic resonance microscopy of atomic transport in optical lattices”, *Nat. Phys.* **7**, 604 (2011).
- [123] H. Zhang, R. McConnell, S. Ćuk, Q. Lin, M. H. Schleier-Smith, I. D. Leroux, and V. Vuletić, “Collective state measurement of mesoscopic ensembles with single-atom resolution”, *Phys. Rev. Lett.* **109**, 133603 (2012).
- [124] Z. Chen, J. G. Bohnet, J. M. Weiner, K. C. Cox, and J. K. Thompson, “Cavity-aided nondemolition measurements for atom counting and spin squeezing”, *Phys. Rev. A* **89**, 043837 (2014).
- [125] M. Campisi, P. Hänggi, and P. Talkner, “*Colloquium*: quantum fluctuation relations: foundations and applications”, *Rev. Mod. Phys.* **83**, 771 (2011).
- [126] M. Campisi, P. Talkner, and P. Hänggi, “Fluctuation theorems for continuously monitored quantum fluxes”, *Phys. Rev. Lett.* **105**, 140601 (2010).
- [127] Y. M. Blanter and M. Büttiker, “Shot noise in mesoscopic conductors”, *Physics Reports* **336**, 1 (2000).
- [128] W. Schottky, “Über spontane Stromschwankungen in verschiedenen Elektrizitätsleitern”, *Ann. Phys.* **362**, 541 (1918).

- [129] S. Gustavsson, R. Leturcq, B. Simovič, R. Schleser, T. Ihn, P. Studerus, K. Ensslin, D. C. Driscoll, and A. C. Gossard, “Counting statistics of single electron transport in a quantum dot”, *Phys. Rev. Lett.* **96**, 076605 (2006).
- [130] J.-P. Brantut, J. Meineke, D. Stadler, S. Krinner, and T. Esslinger, “Conduction of ultracold fermions through a mesoscopic channel”, *Science* **337**, 1069 (2012).
- [131] J.-P. Brantut, C. Grenier, J. Meineke, D. Stadler, S. Krinner, C. Kollath, T. Esslinger, and A. Georges, “A thermoelectric heat engine with ultracold atoms”, *Science* **342**, 713 (2013).
- [132] C.-C. Chien, S. Peotta, and M. Di Ventra, “Quantum transport in ultracold atoms”, *Nat. Phys.* **11**, 998 (2015).
- [133] H. Tanji-Suzuki, I. D. Leroux, M. H. Schleier-Smith, M. Cetina, A. T. Grier, J. Simon, and V. Vuletić, “Chapter 4 - interaction between atomic ensembles and optical resonators: classical description”, *Advances In Atomic, Molecular, and Optical Physics* **60**, 201 (2011).
- [134] C. J. Hood, T. W. Lynn, A. C. Doherty, A. S. Parkins, and H. J. Kimble, “The atom-cavity microscope: single atoms bound in orbit by single photons”, *Science* **287**, 1447 (2000).
- [135] P. W. H. Pinkse, T. Fischer, P. Maunz, and G. Rempe, “Trapping an atom with single photons”, *Nature* **404**, 365 (2000).
- [136] M. A. Norcia and J. K. Thompson, “Strong coupling on a forbidden transition in strontium and nondestructive atom counting”, *Phys. Rev. A* **93**, 023804 (2016).
- [137] K. Roux, V. Helson, H. Konishi, and J.-P. Brantut, “Cavity-assisted preparation and detection of a unitary Fermi gas”, *New Journal of Physics* **23**, 043029 (2021).
- [138] J. Léonard, A. Morales, P. Zupancic, T. Esslinger, and T. Donner, “Supersolid formation in a quantum gas breaking a continuous translational symmetry”, *Nature* **543**, 87 (2017).
- [139] J. Léonard, A. Morales, P. Zupancic, T. Donner, and T. Esslinger, “Monitoring and manipulating Higgs and Goldstone modes in a supersolid quantum gas”, *Science* **358**, 1415 (2017).
- [140] R. M. Kroeze, Y. Guo, and B. L. Lev, “Dynamical spin-orbit coupling of a quantum gas”, *Phys. Rev. Lett.* **123**, 160404 (2019).
- [141] S. Krinner, T. Esslinger, and J.-P. Brantut, “Two-terminal transport measurements with cold atoms”, *J. Phys.: Condens. Matter* **29**, 343003 (2017).
- [142] S. Uchino, M. Ueda, and J.-P. Brantut, “Universal noise in continuous transport measurements of interacting fermions”, *Phys. Rev. A* **98**, 063619 (2018).
- [143] D. Yang, C. Laflamme, D. V. Vasilyev, M. A. Baranov, and P. Zoller, “Theory of a quantum scanning microscope for cold atoms”, *Phys. Rev. Lett.* **120**, 133601 (2018).

- [144] K. B. Davis, M. Mewes, and W. Ketterle, “An analytical model for evaporative cooling of atoms”, *Appl. Phys. B* **60**, 155 (1995).
- [145] O. J. Luiten, M. W. Reynolds, and J. T. M. Walraven, “Kinetic theory of the evaporative cooling of a trapped gas”, *Phys. Rev. A* **53**, 381 (1996).
- [146] H. Wu and C. J. Foot, “Direct simulation of evaporative cooling”, *J. Phys. B: At. Mol. Opt. Phys.* **29**, L321 (1996).
- [147] K. M. O’Hara, M. E. Gehm, S. R. Granade, and J. E. Thomas, “Scaling laws for evaporative cooling in time-dependent optical traps”, *Phys. Rev. A* **64**, 051403 (2001).
- [148] S. Whitlock, C. F. Ockeloen, and R. J. C. Spreeuw, “Sub-poissonian atom-number fluctuations by three-body loss in mesoscopic ensembles”, *Phys. Rev. Lett.* **104**, 120402 (2010).
- [149] M. A. Kristensen, M. B. Christensen, M. Gajdacz, M. Iglicki, K. Pawłowski, C. Klempt, J. F. Sherson, K. Rzażewski, A. J. Hilliard, and J. J. Arlt, “Observation of atom number fluctuations in a Bose–Einstein condensate”, *Phys. Rev. Lett.* **122**, 163601 (2019).
- [150] M. B. Christensen, T. Vibel, A. J. Hilliard, M. B. Kruk, K. Pawłowski, D. Hryniuk, K. Rzażewski, M. A. Kristensen, and J. J. Arlt, “Observation of microcanonical atom number fluctuations in a bose-einstein condensate”, *Phys. Rev. Lett.* **126**, 153601 (2021).
- [151] D. Yang, A. Grankin, L. M. Sieberer, D. V. Vasilyev, and P. Zoller, “Quantum non-demolition measurement of a many-body Hamiltonian”, *Nat. Commun.* **11**, 775 (2020).
- [152] A. Schuckert and M. Knap, “Probing eigenstate thermalization with the emergence of fluctuation-dissipation relations in quantum simulators”, arXiv:2007.10347 (2020).
- [153] M. Gajdacz, P. L. Pedersen, T. Mørch, A. J. Hilliard, J. Arlt, and J. F. Sherson, “Non-destructive Faraday imaging of dynamically controlled ultracold atoms”, *Review of Scientific Instruments* **84**, 083105 (2013).
- [154] M. Gajdacz, A. J. Hilliard, M. A. Kristensen, P. L. Pedersen, C. Klempt, J. J. Arlt, and J. F. Sherson, “Preparation of ultracold atom clouds at the shot noise level”, *Phys. Rev. Lett.* **117**, 073604 (2016).
- [155] U. Fano, “Ionization yield of radiations. ii. the fluctuations of the number of ions”, *Phys. Rev.* **72**, 26 (1947).
- [156] D. B. Hume, I. Stroescu, M. Joos, W. Muessel, H. Strobel, and M. K. Oberthaler, “Accurate atom counting in mesoscopic ensembles”, *Phys. Rev. Lett.* **111**, 253001 (2013).
- [157] T. A. Savard, K. M. O’Hara, and J. E. Thomas, “Laser-noise-induced heating in far-off resonance optical traps”, *Phys. Rev. A* **56**, R1095 (1997).
- [158] A. Keshet and W. Ketterle, “A distributed, graphical user interface based, computer control system for atomic physics experiments”, *Rev. Sci. Instrum.* **84**, 015105 (2013).

- [159] J. A. Nelder and R. Mead, “A Simplex Method for Function Minimization”, *The Computer Journal* **7**, 308 (1965).
- [160] P. B. Wigley, P. J. Everitt, A. van den Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. McDonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. N. Kuhn, I. R. Petersen, A. N. Luiten, J. J. Hope, N. P. Robins, and M. R. Hush, “Fast machine-learning online optimization of ultra-cold-atom experiments”, *Scientific Reports* **6**, 25890 (2016).

Appendix A

Calibration of magnetic bias fields

Precise characterizations of the background magnetic fields present at the location of the cavity, as well as precise calibrations of the bias fields effected by the external bias coils, are necessary for reliable preparation and measurement of the atomic ensemble. Constructing Hamiltonians that commute with the spin energy, such as those used in Chapter 4 to measure atom number dynamics and in Section 3.2 to nondestructively measure the collective spin state, requires generating fields that point directly along the cavity (“vertical”) axis—or, equivalently, precisely zeroing fields along the two other (“waveguide” and “imaging”) axes. Meanwhile, rf drives are required for preparing the collective spin ensemble in a desired initial state, as in Chapter 3, and the fidelity of these drives depends on the precise characterization of the magnetic field strength.

When a spin ensemble is placed in a magnetic field and perturbed away from one of the poles of its collective Bloch sphere, it will precess at a frequency ω_s that is proportional to the magnitude B of the magnetic field:

$$\hbar\omega_s = g_F\mu_B B. \quad (\text{A.1})$$

The spin ensemble can be made to interact with an optical cavity, as in Chapter 3. For a magnetic field that is not aligned with the cavity axis, the spin precession will modulate the resonance condition of the cavity (Eqn. 3.15), giving rise to Stokes and anti-Stokes sidebands on the pump mode at $\mp\omega_s$, which can be measured on the light escaping from the cavity.

The cavity is pumped with light at a constant frequency, and the spectrum of the light escaping from the cavity is measured using a balanced heterodyne detector (Section 2.3.1). Fig. A.1 shows the cavity spectra for a variety of magnetic field conditions with constant (small) B_z and varying B_x . The spin precession frequency, proportional to the total magnetic field, is extracted from each spectrum. The fields measured at each field condition can be used to extract the background magnetic field along the axis, as well as the (linear) field generated per current supplied (Fig. A.2). It is important that some background field B_z be kept on as B_x is varied, as this prevents the total field from being set to 0, which would depolarize the collective spin. Accordingly, the fits to the total field as B_x is varied are hyperbolic rather than linear.

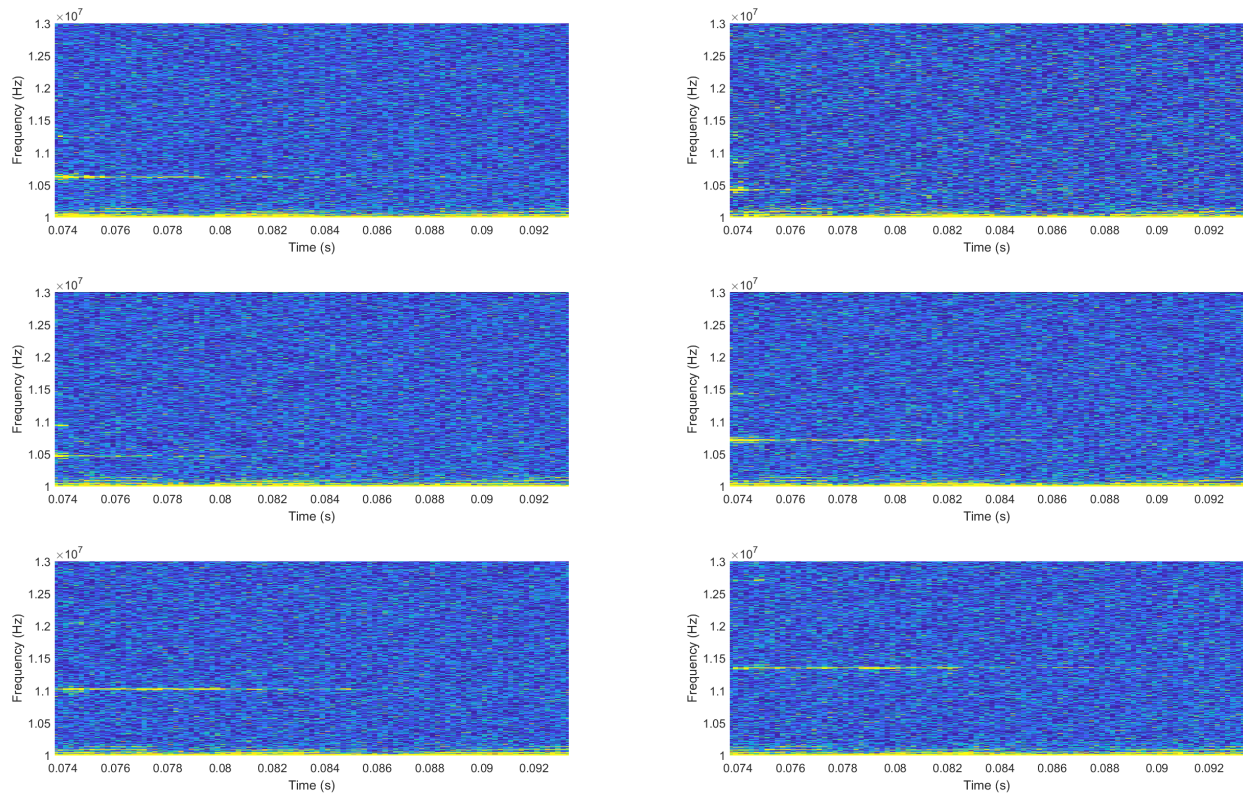


Figure A.1: Spectrograms showing the precession of the collective atomic spin in varying magnetic fields. The current sent to the imaging-axis bias coils is varied between panels. The collective spin is placed near the equator and allowed to precess freely, modulating the cavity field at a frequency ω_s that depends on the magnitude of the total magnetic field while the cavity is pumped near resonance with a constant tone. The heterodyne carrier frequency is $f_{\text{car}} = 10$ MHz; features are seen at $f_{\text{car}} \pm \omega_s/2\pi$. A small bias field is effected in the vertical direction in order to avoid zero-field conditions. Each spectrogram shows an average of ~ 10 experimental runs.




Figure A.2: Calibration of the imaging-axis bias field. The spectrograms shown in Fig. A.2 are averaged over time, and the resulting spectra are used to find the precession frequency at each setpoint. This is used to find the offset current required to compensate for the background field, as well as the slope of the field generated per current supplied.

The above calibrations, using spectroscopy of the cavity mode to measure the modulation imprinted by the precessing collective spin, rely on the presence of pump light in the cavity. The light used to pump the cavity, however, results in an extra effective magnetic field, and biases the calibrations. To avoid this, we can use rf spectroscopy to measure the magnetic field seen by the collective spin without needing to pump the cavity. The initial and final states are measured as in Section 3.2, and, in between, an rf drive is applied. This is repeated for varying drive frequencies, and the response of the collective spin is observed (Fig. A.3 and Fig. A.4, left panels) and used to back out the magnitude of the field. Repeating this for a variety of bias fields, as in the “bright” procedure outlined above, allows the magnetic field to be calibrated in the absence of any pump light (Fig. A.3 and Fig. A.4, right panels).

Comparing the “bright” (Fig. A.2) and “dark” (Fig. A.4, top right) calibrations of the imaging-axis bias field, the effect of the cavity occupation is clear: The field generated per current supplied to the coils is measured to be the same using either method of calibration; however, there is an apparent offset between the measured currents required to zero the total fields. This offset corresponds to the effective field due to the pump light.

In order for these calibrations to be effective, it must be possible to reach a field of either sign using each of the bias coils (*e.g.*, upwards or downwards, in the case of the vertical coils). This is not generically possible, since the coils are driven by unipolar supplies. To account for this, small “anti-bias” coils are installed along each axis and are driven with constant current at all times, opposing the fields generated by the main bias coils. This means that, when the main bias coil along a given axis is not driven, there is a small negative field present. When a true zero-field condition is desired, a small current must be run through the main bias coils; this corresponds to the (positive) horizontal offset on each of the fits in the right

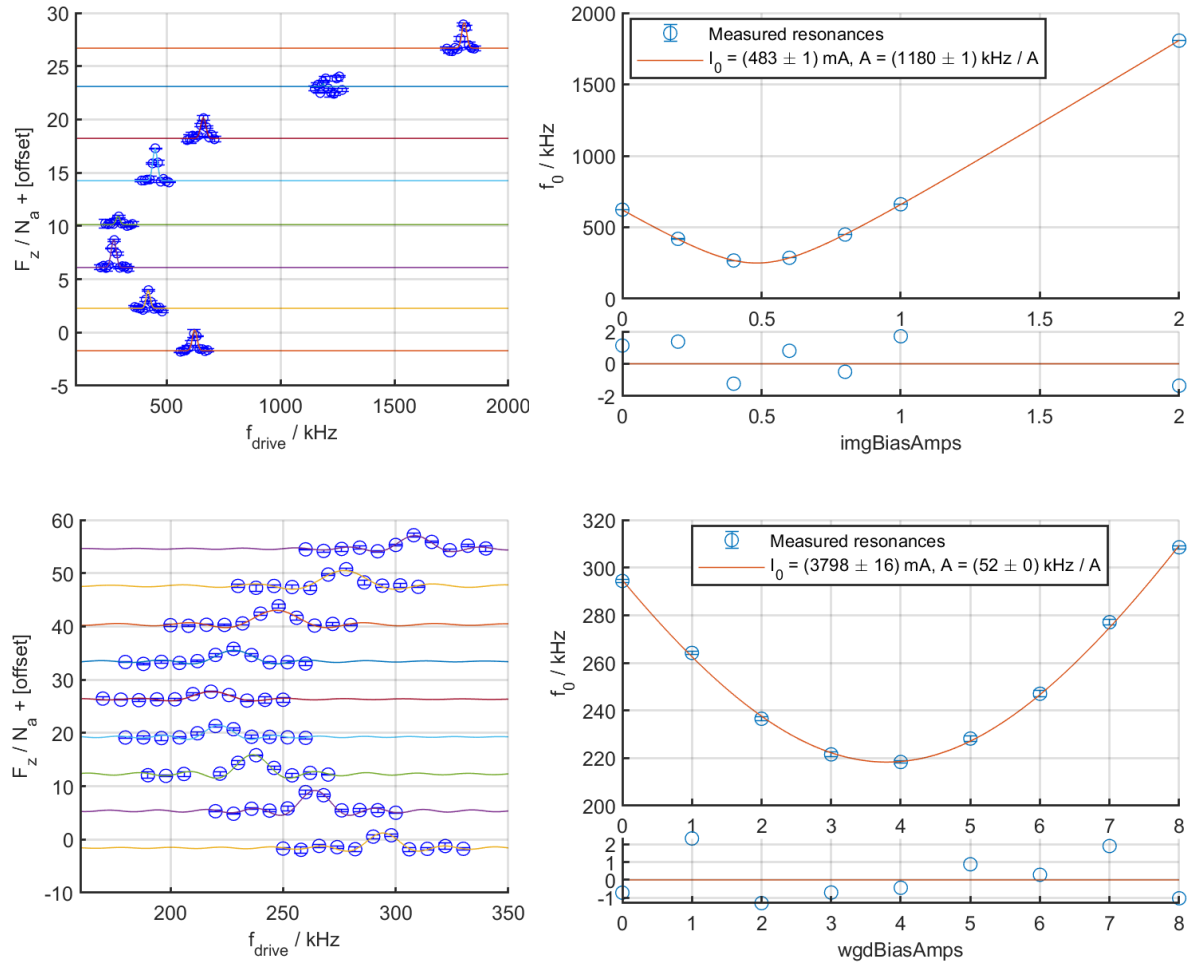


Figure A.3: Calibration of the magnetic bias field in the imaging-axis (top) and waveguide-axis (bottom) directions. **left:** At each magnetic coil current setpoint, rf spectroscopy is used to find the magnitude of the magnetic field in the absence of any pump light. Error bars represent standard errors on the mean. **right:** The total field magnitudes measured at each magnetic coil current setpoint are used to find the offset current required to compensate for the background field, as well as the slope of the field generated per current supplied. Error bars represent fit uncertainties. A small bias field is effected in the vertical direction in order to avoid zero-field conditions. Each point is an average of ~ 5 experimental runs.

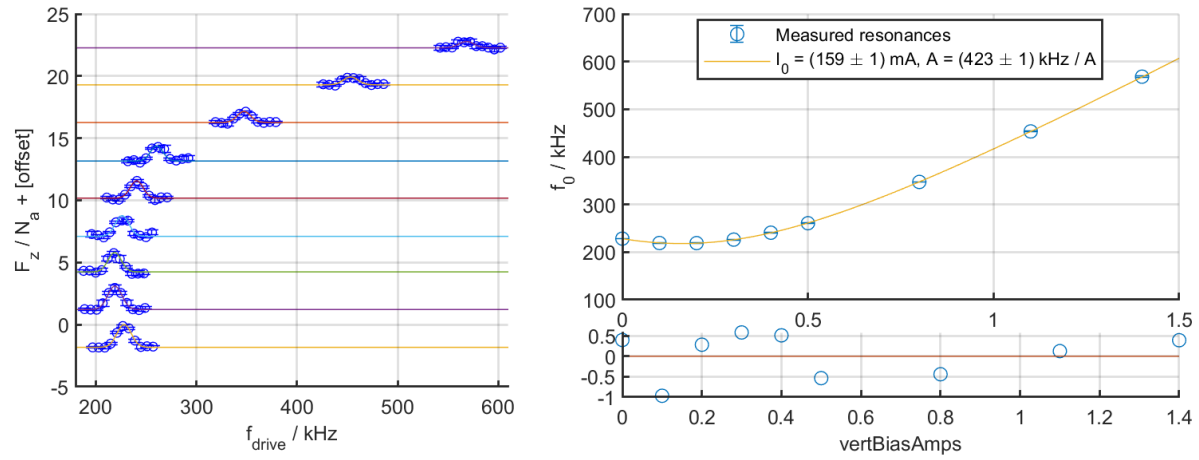


Figure A.4: Calibration of the magnetic bias field in the vertical direction. Same procedure is used as in Fig. A.3; see caption for details. A small bias field is effected in the imaging-axis direction, rather than the vertical direction, in order to avoid zero-field conditions.

columns of Fig. A.3 and Fig. A.4.

Appendix B

Measuring the cavity linewidth

Many of the measurements carried out in a cavity experiment such as E3 depend sensitively on the cavity’s half-linewidth κ . This is a property of the cavity itself, and is not expected to change over time, so in most cases—barring any major changes to the apparatus—it is measured only once in the active lifetime of an experiment. The linewidth of the E3 cavity was measured in 2010 by Tom Purdy to be $\kappa = 2\pi \times 1.82$ MHz [31]. In late 2022, there was reason to believe that the quality of the cavity may have decreased, leading to an increase in κ . There was no obvious explanation for why properties of the cavity would have changed over time; however, over the course of the previous year, the science chamber had been opened and re-baked (twice) in order to replace the rubidium supply (Section 2.5), and, although measures were taken to ensure minimal air entered the chamber, it seemed plausible that the surfaces of the mirrors may have oxidized during the process. As such, we endeavored to re-measure the cavity linewidth.

Cavity ringdown measurement (unsuccessful)

A common way of measuring κ is a so-called cavity ringdown; Justin Gerber’s dissertation offers a concise description of a ringdown measurement in the context of AMO experiments [61]. When the detuning between the cavity resonance and the light pumping the cavity is rapidly swept across resonance, the population of the cavity is not able to respond immediately. The finite response time of the cavity results in different detunings of light being present in the cavity at the same time; these interfere with each other, leading to beating in the cavity output:

$$P_{\text{out}}^{\text{rd}}[\kappa, \eta](t) = 2\kappa\eta^2 \frac{2\pi}{|\tilde{\nu}|} \left| e^{-t'} \operatorname{erfc} \left(\frac{i - \tilde{\nu}t'}{\sqrt{2i\tilde{\nu}}} \right) \right|^2. \quad (\text{B.1})$$

Here, $t' \equiv \kappa t/2$ and

$$\tilde{\nu} \equiv \frac{4}{\kappa^2} \frac{d}{dt} \Delta_{\text{pc}} \quad (\text{B.2})$$




Figure B.1: A failed attempt at measuring κ by scanning the cavity across resonance and measuring the ringdown. The cavity output is measured by a SPAD as the pump-cavity detuning is scanned across resonance; data are averaged over ~ 10000 repetitions (black circles). A fit to Eqn. B.1 (orange line) shows that the expected beating between frequency components is not present, likely due to improper averaging of the data. Adding a temporal blur to the fitting function, as in Eqn. B.3 (blue line), appears to fit the data well, but claims a significantly smaller value of κ than that which was measured previously; this is not treated as trustworthy, although the specific failure mode this method remains poorly understood.

measures the rate at which the detuning is swept. For the desired effect, it is necessary that $\tilde{\nu} \gg 1$. The beating (or “ringing”) between different frequencies of light offers a very precise measurement of κ , as long as $d\Delta_{pc}/dt$ is known precisely.

When attempting to carry out the ringdown measurement, we were unable to send sufficient power to the cavity to record $P_{\text{out}}^{\text{rd}}$ in a single shot. In order to achieve a strong signal, it was necessary to average over 1000s of repetitions, measuring the (very small) output of each repetition on a SPAD at the cavity output. This introduced the problem of how to appropriately average these data together in the presence of the slow drift of the cavity resonance frequency. Even when some significant care was taken, clear ringing was not observable (Fig. B.1, orange lines), and the quality of the fit of $P_{\text{out}}^{\text{rd}}$ to the data (with κ and η as free parameters) was poor.

In order to account for the apparent timing jitter between shots, a fit function was developed which included some temporal blur:

$$P_{\text{out}}^{\text{rd}}[\kappa, \eta, \sigma](t) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} d\tau P_{\text{out}}^{\text{rd}}[\kappa, \eta](\tau) \exp\left(-\frac{1}{2} \frac{\tau^2}{\sigma^2}\right). \quad (\text{B.3})$$

Here, σ (dimensions of time), which is allowed to vary as a free parameter when fitting,

quantifies the amount of jitter from shot to shot. This assumes Gaussian jitter, an assumption which is not justified *a priori* but which seems to describe the data well (Fig. B.1, blue lines).

The results of this fit were surprising: it appeared that κ had decreased, corresponding to an increased cavity lifetime since the time of its original characterization. This did not seem plausible. Any number of things (damaged mirrors, alkali deposits, misalignments) could result in a decreased cavity lifetime; however, increasing the cavity lifetime should have been very difficult. Accordingly, although it was not clear what had gone wrong, this measurement was taken to be erroneous.

Cavity extinction measurement (successful)

A more simple, although less precise, measurement of κ can be achieved by populating the cavity mode with an external pump and then suddenly extinguishing the pump. If the pump is extinguished on a timescale that is fast compared to the cavity lifetime, then the cavity occupation (and, thus, the power of the light leaving the cavity) will decrease exponentially with a time constant $1/2\kappa$.

One way to quickly extinguish light is with AOM. The speed at which this can be done is limited by the time that it takes the acoustic wave traveling through the AOM crystal to traverse the spot of light being diffracted. The speed of sound in most AOM crystals is around 4000 m s^{-1} , meaning that, in order to ensure a shuttering time of 10 ns or less (roughly what we expect to need for this characterization), a beam diameter of $40 \mu\text{m}$ is required. This is easily achievable by focusing the pump beam through the AOM. All that remains is to turn off the rf signal driving the AOM sufficiently quickly, which can be achieved easily by a modern rf switch.

The turn-off time of the pump light by the AOM was measured experimentally (Fig. B.2, right panel) and found to be much faster than the expected extinction time of the cavity mode. The extinction measurement, then, allowed the cavity linewidth to be measured reliably (Fig. B.2, left panel). The measured value of $\kappa = 2\pi \times 1.81(8) \text{ MHz}$ was in agreement with the initial characterization of the cavity at the time of the experiment's construction.




Figure B.2: The cavity half-linewidth κ is measured by populating the cavity and then quickly extinguishing the pump light. Left: The cavity population decays exponentially with a time-constant $1/2\kappa = 44(2)$ ns, corresponding to a cavity half-linewidth of $\kappa = 2\pi \times 1.81(8)$ MHz. This is measured using a fiber-coupled SPAD at the output of the cavity. Right: The pump light is extinguished in $7(1)$ ns, which is sufficiently fast to render this measurement trustworthy. This is measured by redirecting the pump light directly to the SPAD, rather than coupling it to the cavity.

Appendix C

Machine learning optimization of the experimental sequence

The sequence required to prepare a sample of atoms before doing science comprises many steps and many moving pieces. An overview is given in Chapter 2, with greater detail provided by previous E3 theses [24, 31]. Throughout the sequence, optimizations over high-dimensional parameter spaces may often offer significant improvements to the number of atoms available in the final sample; however, because of (often poorly understood or poorly characterized) interdependencies between parameters, individually optimizing each parameter may not lead to good global optima. Further, the long cycle time (30 second for the full sequence, or a minimum of 4 second to reach the first opportunities for quantitative readout) makes it impractical to simply scan over all possible combinations of parameters. This offers an ideal use-case for purpose-built optimization algorithms which may be allowed to feed back directly to the experimental parameters based on previous outputs.




Figure C.1: Example Cicero Word Generator sequence used to control the experimental apparatus.

Cicero Word Generator sequence control frontend

The experimental sequence is controlled by the Cicero Word Generator software suite [158]. This offers a graphical interface for sequencing analog and digital channels (Fig. C.1). The sequence depends on 2–400 control variables which may be stepped, from shot to shot, in order to optimize parts of the sequence.

At the beginning of each shot, Cicero offers the option to read values of control variables from an external file. This, along with the ability to read out diagnostic information about the system during the sequence (using, *e.g.*, absorption imaging or dispersive atom number readout), provides the requisite components of a feedback loop: all that remains is to build a controller. The controller must maintain a list referencing the atom number at each previous shot to the corresponding control variables, and must use this information to determine how the control variables should be changed going forward. The difficulty lies in this final step: the parameter space is complex, and building an algorithm to map it out is not straightforward.

Using M-LOOP for high-dimensional optimization

The above is not a novel problem; optimization algorithms specialized for high-dimensional parameter spaces have been explored thoroughly. One such algorithm was described by Nelder and Meade [159]. This offers a strong starting point for the optimization process, and does an efficient job of initially mapping out the parameter space. Once the parameter space is roughly understood, however, modern machine learning (ML) techniques can be much more efficient. The technique which we have chosen to use involves first “learning” the system using the Nelder–Meade algorithm for $2q$ steps (where q is the size of the parameter space being optimized, generally 10–30 in our case) and then handing over this information to an ML optimizer which “iterates” on the system until it reaches an apparent optimum; see Fig. C.2.

This process is carried out with the help of M-LOOP, a software suite designed to reduce the complexity of general ML problems, specializing on the specific use case of AMO experiments [160]. With some extra application-specific scaffolding, M-LOOP makes it straightforward to read in initial parameter values; learn the parameter space by using Nelder–Meade or a similar simplex algorithm and feeding back to the list of Cicero control variables, aiming to maximize the measured atom number; and efficiently optimize over the parameter space using ML techniques (Fig. C.2). This has allowed us to optimize significant portions of the sequence which would have otherwise been inaccessible to simple iteration techniques due to a lack of scalability. In doing so, it has offered increased overhead, making it possible to go longer periods of time without touching up physical systems which are prone to drift away from their optima; this has increased productivity significantly.




Figure C.2: Machine learning optimization of the experimental sequence. The ensemble is imaged by absorption after the initial evaporative cooling stage, and the atom number is derived from the image. In this example, magnetic field strengths during the early trapping and cooling stages are varied. The negative of the atom number is used as a cost function for the optimization (vertical axis). For the first $2q = 24$ shots, the Nelder–Meade learner maps the parameter space (orange points). After the initial mapping stage, the machine learning algorithm (blue points) is able to quickly increase the atom number from $\sim 3.5 \times 10^5$ to $\sim 6 \times 10^5$ while the experimenters are out at lunch. Periodically, the system returns to the Nelder–Meade learner to avoid leaving a favorable region of the parameter space entirely.

Appendix D

Data management for high-throughput experiments

The light outputted from the E3 cavity is measured, in most contexts, on a heterodyne detector at $\omega_{\text{het}} = 10\text{--}20$ MHz (Section 2.3.1). Processing the output of the heterodyne measurement digitally requires a minimum bandwidth of $2\omega_{\text{het}}$. This is achieved using a CompuScope GaGeScope board, externally clocked at 80 MHz to ensure sufficient bandwidth overhead, which connects to our data acquisition computer as a PCIe extension. GaGeScope records data from its analog inputs at a 14b depth. Along with the output of the heterodyne detector, we also use the GageScope board to measure the amplitude of the ODT light outputted from the cavity, as well as 1–2 other signals dependent on the science we’re pursuing.

Recording on 3–4 channels at 80 MHz produces a significant amount of data, particularly when experiments (such as those discussed in Chapter 4) often require recording data for 100s of ms: this amounts to ~ 200 MB per experiment, or roughly 30 GB/h. Processing and storing this amount of data—as well as making it available for easy referencing—requires significant care.

Preprocessing raw data

The first step taken to streamline data storage and processing is to preprocess the data as they are recorded. This serves the purpose of allowing only the necessary density of information to be saved to drives.

The heterodyne signal must be stored at 80 MHz in order to be demodulated with sufficient overhead, but the same is not true for most other channels; for example, the ODT intensity is generally downsampled by a factor of 10 before being saved. The result is that, in practice, a majority of space is taken up only by the heterodyne signal itself.

Similarly, while data are recorded over the course of 100s of ms, it is common to only need information from a few small segments of time: an initial (12 ms) dispersive atom number measurement, an intermediate (10–400 ms) period of time during which experiments are

carried out, and a final (27 ms) dispersive measurement of both the final atom number and final total spin energy, as described in Section 3.2. These segments are separated by 10s or 100s of ms of dead time, the length of which is limited by the need to change the currents in large magnetic field coils. Rather than save the data recorded throughout this entire time span, it is beneficial to segment the data during preprocessing and only save what is recorded during the periods of interest.

These preprocessing steps are carried out by a homebuilt software suite, `gage_acquire`, written initially by Jonathan Kohler and maintained collectively by E3. As well as resampling and segmenting the raw data, `gage_acquire` offers convenient shot-by-shot visualization of resampled and (where appropriate) demodulated data, which is essential for debugging new experimental sequences.

Redundancy and backups

Data storage is cheap, and data loss can be catastrophic. As such, redundant backups of all data should be viewed as essential to any high-throughput experiment. In E3, this takes a few forms.

Data read in from GageScope, as well as from the Andor camera used for absorption imaging, is saved directly to a RAID0 housed in the same box as the GageScope board and mapped to the local E: drive. This computer is used for data analysis as well as acquisition, making it a convenient access point. Each evening, the relevant base directories of this internal drive array are backed up to a Synology DiskStation configured as a RAID10 and conventionally mapped to S:; this is automated using `relone`. This provides an immediate back up, protecting against damage to the local E: drive. It also offers a greater storage capacity than the local E: drive: generally, data are removed from E: after ~ 2 years, while they remain on S: indefinitely, with extra capacity being added as necessary.

Nightly backups to S: provide some security against data loss due to computer issues. Due to being housed in separate (neighboring) buildings, some protection is even provided against slightly less local issues such as floods. Larger-scale calamities such as fires or earthquakes, however, could still easily damage both drive arrays simultaneously. Indeed, even in the absence of any such calamities, the chances of two RAIDs failing simultaneously is very small, but not vanishing. As such, it is prudent to also keep redundant off-site backups of all data.

For E3, off-site backups are available twofold, in the forms of Box and Google Drive. Both are available through UC Berkeley, and each has specific upsides. Google Drive offers relatively quick recovery of complete datasets (ones of days, for our tens of TB of data); however, Google's implementation of indexing makes it much less straightforward to retrieve specific subsets of data. Box, meanwhile, indexes files in a much more straightforward manner, but recovery of a complete dataset is much more time-consuming (tens of days, roughly). Accordingly, once per month the entirety of the S: server is backed up to both Google Drive and Box. This doubly redundant offsite backup scheme, along with the singly redundant (although limited in scope) onsite backup provided by S:, guarantees long-term

data security, and is completely automated by a series of simple scripts built around rclone. These scripts can be made available upon request.