STUDY OF CYTOP AND DEUTERATED POLYETHYLENE COATINGS FOR ULTRACOLD NEUTRON STORAGE

BY

THOMAS NEULINGER

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Physics
in the Graduate College of the
University of Illinois Urbana-Champaign, 2021

Urbana, Illinois

Doctoral Committee:

Professor Liang Yang, Chair
Professor S. Lance Cooper
Professor John Stack
Professor Douglas H. Beck

Abstract

In his 1920 Bakerian Lecture on the *Nuclear Constitution of Atoms*, Ernest Rutherford speculated on the properties of the then undiscovered neutron, remarking that “it may be impossible to contain it in a sealed vessel”. Nevertheless, trapping of very low energy—or “ultracold”—neutrons (UCN) is possible due to an effective “optical potential” experienced by neutrons in matter. At present, storage times in the best material traps are essentially limited only by the free neutron lifetime, $\tau_\beta \approx 880\,\mathrm{s}$, offering a precise method of measuring this quantity, whose value has important consequences in theoretical physics. Another important use of stored UCN is in experimental efforts to measure a nonzero neutron electric dipole moment (EDM). In contrast to neutron lifetime measurements, EDM experiments require storage materials that are capable of supporting large electric fields. In many cases, continued progress in these experiments will also be driven by the increased outputs expected from new UCN sources that will use novel storage materials suitable to cryogenic environments.

In this work, we report on UCN storage in vessels coated with the fluoropolymer CYTOP and with deuterated polyethylene (dPE), carried out at the SUN-2 prototype UCN source at the Institute Laue-Langevin (ILL) in Grenoble, France. The two materials were investigated for their potential use as coatings for the UCN production volume of the SuperSUN source, and for part of the storage cells of the PanEDM neutron EDM experiment, respectively, both of which are currently under construction at the ILL.

In Chapter 1, we motivate the development of improved coatings for UCN storage by discussing the use of neutron EDM measurements as probes of physics beyond the Standard Model. This is followed in Chapter 2 by a brief review of UCN physics, and a discussion in some detail of several models for nonideal UCN storage. The CYTOP and dPE storage results are then presented in Chapter 3, in addition to a measurement of the optical potential of CYTOP, which is found to be $115.2(2)\,\mathrm{neV}$. In Chapter 4, we present the results of simulations that can be used to extract an intrinsic storage property of the CYTOP and dPE coatings, the loss factor, η, from the experimental storage data. This analysis is discussed in Chapter 5, in which conservative estimates of η are made: the loss factor of CYTOP is found to be $\eta = 2.7(2) \times 10^{-5}$ at 11.7 K and $\eta = 1.1(1) \times 10^{-4}$ at 295.8 K, and the loss factor of dPE is found to be $\eta = 3.1(1) \times 10^{-5}$ at 11.9 K and $\eta = 2.6(1) \times 10^{-4}$ at 295.1 K.

Acknowledgments

First and foremost, I wish to thank my advisor, Professor Douglas Beck, who, despite having many responsibilities, has always been generous with his time. It would be difficult to overstate the degree to which I have benefited from his experience, insight, and sense of humor over the years. I am also indebted to Dr. Oliver Zimmer, without whom this research would not have been possible, and with whom it will be my great pleasure to continue working in the coming years.

In addition, I have been fortunate to know and to have worked with many wonderful people involved in the PanEDM collaboration: Euan Connolly, Professor Skyler Degenkolb, Dr. Hanno Filter, Professor Peter Fierlinger, Jürgen Hingerl, Chris Klau, Joachim Meichelböck, Robert Paddock, Magdalena Pieler, Florian Röhrer, Stefan Stuiber, and David Wurm. I would also like to express my gratitude for the help and support I have received from the staff in the Nuclear and Particle Physics group at the Institut Laue-Langevin, especially Thomas Brenner, Dr. Tobias Jenke, Dr. Michael Jentschel, Alex Quirk, and Dr. Stephanie Roccia.

I am especially grateful to Professor Lance Cooper, the Associate Head for Graduate Programs, whose encouragement and guidance throughout graduate school have been invaluable to me, and who is a credit to the Physics Department at the University of Illinois. Finally, I wish to thank my friends and family, whose love and support has seen me through it all.

This work was supported in part by the National Science Foundation under Grant No. PHY-2111046.

Table of Contents

List of Tables

List of Figures

List of Abbreviations and Jargon

CN	Cold neutron.
Cold	Of energy $E \sim 0.1$ to $10\,\mathrm{meV}$, wavelength $\lambda \sim 30$ to $3\,\mathrm{\AA}$ [5].
CYTOP	Polyfluoromer produced by A. G. C. Chemicals, monomer $C_6F_{10}O$.
dPE	Deuterated polyethylene, monomer C_2D_4.
EDM	Electric dipole moment.
ILL	Institut Laue-Langevin.
Thermal	Of energy $E \sim 5$ to $100\,\mathrm{meV}$, wavelength $\lambda \sim 1$ to $0.4\,\mathrm{\AA}$ [5].
Upscatter	Inelastic scattering of a neutron that carries it out of the UCN energy regime.
UCN	Ultracold neutron.
Ultracold	Of energy $E \lesssim 300\,\mathrm{neV}$, wavelength $\lambda \gtrsim 500\,\mathrm{\AA}$.

List of Symbols

E_b	The total UCN energy with respect to the storage bottle.
E_s	The total UCN energy with respect to the superfluid ^{4}He in the SUN-2 source.
η	The energy-independent neutron loss factor, W/V.
$\mu(E)$	The energy-dependent neutron loss probability per bounce.
$\bar{\mu}(E)$	The angle-averaged energy-dependent neutron loss probability per bounce.
$n(t,E)$	(Typically) an energy spectrum of UCN, units of inverse energy.
$n_b(t,E)$	The energy spectrum of UCN in the storage bottle, units of inverse energy.
$n_s(t,E)$	The energy spectrum of UCN in the SUN-2 source, units of inverse energy.
$\tau(E)$	The energy-dependent storage time constant.
τ_β	The free neutron lifetime, 879.4(16) s [6].
τ_0	The idealized, bottle- and spectrum-dependent storage time constant.
τ_d	An empirical time constant characterizing a decrease in storage data at low temperature.
τ_w	An approximate storage time constant due to the main coating in the weak patch model.
τ_p	An approximate storage time constant due to the weak patches in the weak patch model.
U_F	The effective neutron optical potential, $U_F = V - iW$.
V	The real part of the neutron optical potential, $\mathrm{Re}(U_F)$.
W	Minus the imaginary part of the neutron optical potential, $-\mathrm{Im}(U_F)$.

Chapter 1

Motivation

In this chapter, a motivation for the development of improved coatings for ultracold neutron storage is presented. First, the puzzle of baryon asymmetry within the Standard Model (SM) of particle physics is introduced, along with a discussion of its connection to the violation of the combined symmetries of charge conjugation and parity, that is, CP violation. Next, the example of measurements of a permanent electric dipole moment (EDM) is chosen as a class of experiments that can constrain CP violation, and ultimately act as a sensitive probe for new physics Beyond the Standard Model (BSM). The focus is then turned toward EDM measurements with neutrons, specifically using Ramsey spectroscopy, the most common experimental technique, to determine the EDM. Finally, the link to coatings for neutron storage is made in the context of the PanEDM collaboration experiment for the measurement of the neutron EDM and the SuperSUN source of ultracold neutrons, both currently under construction at the Institute Laue-Langevin neutron facility in Grenoble, France.

1.1 Baryon asymmetry and CP violation in the Standard Model

The Standard Model (SM) of particle physics describes how three fundamental forces—the strong force, the weak force, and the electromagnetic force—govern the interactions between all known elementary particles. It is not a theory of everything: if it were, it would include gravity, which is instead described by the separate theory of general relativity. Nevertheless, since the late 1970s [7], the SM has been extremely successful in explaining the results of many experiments in particle physics, and even in predicting the existence of a number of particles, such as the top quark [8], the tau neutrino [9], and the Higgs boson [10]. However, in its current formulation, there are clear problems with the SM.

One major issue is that the observable universe does not appear to contain very much baryonic antimatter [11]. (The same is true of leptonic antimatter.) This suggests that more matter than antimatter must have ultimately emerged from the Big Bang—the singular, violent event approximately 13.8 billion years ago that resulted in the universe as we observe it today. The alternative, that this baryon asymmetry was an initial condition, seems implausible within the framework of the SM, which treats particles and antiparticles on much the same mathematical basis. However, a line of reasoning may be constructed in which certain SM parameters are used to account for a spontaneous generation of some degree of baryon asymmetry. The problem is that they predict entirely too little [12].

The line of reasoning that connects SM parameters with baryon asymmetry is based on a set of assertions that are conventionally grouped as three criteria. These are essentially as follows [13]:

Criteria for the Production of Baryon Asymmetry during the Big Bang

1. Non-conservation of baryon number.
2. Departure from thermal equilibrium.
3. Violation of charge conjugation and time reversal symmetries.

Although the first of these is practically a tautology, no process violating baryon number conservation has ever been observed. However, without this *sine qua non*, no reaction could occur during the Big Bang to generate baryon asymmetry from an initial state of equal amounts of matter and antimatter. It may be shown that the second criterion is necessary because, otherwise, the number densities of every species X and its antiparticle pair $\overline{X}$ would be equal [12], and no net baryon asymmetry would be possible. The final statement in the above list is actually two conditions: violation of charge conjugation symmetry (C-symmetry), and violation of time reversal symmetry (T-symmetry). To motivate these last requirements, consider the baryon number violating reaction $A \rightarrow B$, in which an initial collections of species A with total baryon number b_A react to produce a final set of species B with total baryon number b_B, so that the net change in baryon number is $\Delta b = b_B - b_A \neq 0$. If C-symmetry were not violated, then the reaction rate of $A \rightarrow B$ would be equal to that of its charge conjugated pair $\overline{A} \rightarrow \overline{B}$, so that the overall change in baryon number from both reactions would be zero. Furthermore, if T-symmetry were not violated, then the rates of $A \rightarrow B$ and $B \rightarrow A$ would be equal, also resulting in no overall change in baryon number. Therefore, these two symmetries must have been violated during the Big Bang in order to generate baryon asymmetry.

It is the degree to which C-symmetry and T-symmetry are violated that can be used to calculate an expected amount of baryon asymmetry generated during the Big Bang. An estimate of the latter degree of violation can be made from the experimental results of a wide range of seemingly disparate measurements by means of the CPT theorem [14]. This theorem states that any Lorentz-invariant quantum field theory—like all those that make up the SM—must be symmetric under the combined action of C, P, and T, where C is the charge conjugation operator, T is the time reversal operator, and P is the parity operator. This implies that if a system is not symmetric under T, then it must also not be symmetric under the combined actions of C and P. For historical reasons, such systems are referred to as exhibiting CP violation [15].

There are two SM Lagrangians that have CP violation built in: that of the strong force and that of the weak force. In the case of the former, the CP violating term in the quantum chromodynamics (QCD) Lagrangian couples the gluon field to its dual. Its size is controlled by an angle $\overline{\theta}_{\mathrm{QCD}}$ [16], whose upper bound has been determined experimentally to be approximately 1.5×10^{-10} [17]. The diminutive size of this angle is referred to as the strong CP problem, as there is seemingly no presently-observed justification for its unnaturally small value. In the weak interaction, CP violation arises from two complex phase angles, the first in the parameterization of the quark-mixing Cabibbo-Kobayashi-Maskawa (CKM) matrix [18], and the second in the parameterization of the neutrino-mixing Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix [19]. Both matrices describe coupling of mass eigenstates to those of the weak interaction. In both cases, the experimentally-determined values of the complex phases, $\overline{\delta}_{\mathrm{CKM}}$ and $\overline{\delta}_{\mathrm{PMNS}}$, do not appear to be “fine-tuned” as in the strong interaction; see, for example, [20, 6]. Neither $\overline{\theta}_{\mathrm{QCD}}$ nor $\overline{\delta}_{\mathrm{CKM}}$ are large enough to explain the presently-observed baryon asymmetry without extensions to the SM; see also, for example, [21, 22]. (The PMNS matrix concerns leptons, not baryons; as with its baryon counterparts however, $\overline{\delta}_{\mathrm{PMNS}}$ is too small to account for the observed lepton asymmetry [23].)

A wide variety of theoretical schemes exist Beyond the Standard Model (BSM) that propose various mechanisms for the generation of additional CP violation. For example, one proposed resolution to the previously-mentioned fine-tuning problem of $\overline{\theta}_{\mathrm{QCD}}$ would be the existence of an “axion” field that couples to the gluon field [24]. Such a mechanism could potentially generate cancellations that suppress $\overline{\theta}_{\mathrm{QCD}}$ [16]. Another potential source of additional CP violation is proposed by Minimal Supersymmetric extensions to the SM (MSSM), which introduce a new cast of “superpartners” to SM fermions and bosons. Such squarks, sleptons, and gauginos would generate new couplings and CP violating phases [25]. In any case, whatever the proposed theoretical route to CP violation, experimental measurement has the final say.

1.2 CP violation and permanent electric dipole moments

An important class of experiments that place constraints on proposed BSM sources of CP violation are attempts to measure a permanent electric dipole moment (EDM). In contrast to an induced EDM, in which an applied electric field produces a physical separation of charges, a permanent EDM would exist in the absence of any external field. Any system possessing a non-zero permanent EDM must violate T-symmetry, and thus CP-symmetry by the CPT theorem.

To motivate this claim, consider a simplified example of a composite system possessing a permanent EDM. Further suppose that the constituents of the system are spinless particles, so that its non-zero angular momentum is completely orbital. In an electric field $\boldsymbol{E}$, the relevant part of the non-relativistic Hamiltonian is then [26]

$$H_{\boldsymbol{E}} = -\boldsymbol{d} \cdot \boldsymbol{E}\,, \tag{1.1}$$

where $\boldsymbol{d} = \sum q_i \boldsymbol{R}_i$ is the vector dipole moment operator, where the q_i and $\boldsymbol{R}_i$ are the charges and positions of the constituents. Due to the Wigner-Eckart theorem, the matrix elements of $\boldsymbol{d}$ and the total angular momentum operator $\boldsymbol{J}$ differ by only a constant factor (see e.g. Chapter 10, Complement E in [27]). Therefore, the above may equally-well be written

$$H_{\boldsymbol{E}} = -\frac{d}{J}\boldsymbol{J} \cdot \boldsymbol{E}\,, \tag{1.2}$$

where d/J is the constant of proportionality. The operator $\boldsymbol{J} = \sum \boldsymbol{R}_i \times \boldsymbol{P}_i$, where the $\boldsymbol{P}_i$ are the momenta of the constituents of the system. Transformation under T, the time reversal operator, gives (see e.g. Chapter 3, Section 3 in [28])

$$\begin{aligned} T\boldsymbol{R}_i T^{-1} &= \boldsymbol{R}_i \\ T\boldsymbol{P}_i T^{-1} &= -\boldsymbol{P}_i\,, \end{aligned} \tag{1.3}$$

which means that $\boldsymbol{J} \to -\boldsymbol{J}$, while $\boldsymbol{d}$ is unaffected. Therefore if the system possesses a non-zero permanent dipole moment, i.e. $d \neq 0$, then the Hamiltonian is not invariant under T, and thus violates T-symmetry. Coincidentally, this system also violates P-symmetry, where the parity operator P now sends $\boldsymbol{d} \to -\boldsymbol{d}$ while $\boldsymbol{J}$ is unaffected.

The results of P violation and CP violation for d hold in the relativistic limit [26], and are a general property of any system possessing a permanent EDM, be it a fundamental or composite particle, an atom, or a molecule [16, 29]. For example, if the electron, a spin-1/2 fermion, possesses a permanent EDM of magnitude d, the part of the Lagrangian density that results in an EDM is [26]

$$\mathcal{L}_{\mathrm{EDM}} = -i\frac{d}{2}\bar{\Psi}\sigma^{\mu\nu}\gamma^5 F_{\mu\nu}\Psi\,, \tag{1.4}$$

where Ψ is the fermion field, $\bar{\Psi} = \Psi^\dagger\gamma^0$ is the conjugate field, $\sigma^{\mu\nu} = (i/2)(\gamma^\mu\gamma^\nu - \gamma^\nu\gamma^\mu)$ where γ^μ are the Dirac matrices, $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3$, and $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the electromagnetic field tensor, where A_μ is the four-vector electromagnetic potential. In the above interaction, the presence of the γ^5 matrix indicates P violation, while the i is responsible for CP violation. In the non-relativistic limit, the Hamiltonian from the above Lagrangian density reduces to that of Equation 1.1 [26], and an experimental determination of the expectation value of $\boldsymbol{d}$ would thus entail measuring transition rates linear in the electric field $\boldsymbol{E}$. However, the main challenge in such an experiment involving a bare electron, or any charged particle for that matter, is that it is difficult to confine in an electric field. (Nevertheless, various storage ring experiments have been carried out; see Table X in [16] for a summary of those under development.) Therefore, EDM searches typically focus on neutral systems.

These systems, which are easier to confine in traps, face their own difficulties however. For example, in the case of neutral atoms, the nucleus, possibly possessing the EDM of interest, is largely shielded from external fields by its electrons [30]. However, if the nuclear EDM distribution is different from the charge distribution, as in large nuclei, an EDM can exist which is proportional to the "Schiff moment" [26]

$$\boldsymbol{S}_{\mathrm{Schiff}} = \left\langle I, I_z = I \left| \frac{1}{10}\sum_{p=1}^{Z}\left(r_p^2 - \frac{5}{3}\left\langle r^2\right\rangle\right)\boldsymbol{r}_p \right| I, I_z = I \right\rangle\,, \tag{1.5}$$

where the sum is carried out over nuclear protons, I is the nuclear spin, and $\left\langle r^2\right\rangle$ is evaluated with respect to the charge distribution. This effect can be most efficiently measured in systems with long nuclear spin relaxation times, such as in ^{199}Hg [31].

1.3 Measurement of the neutron electric dipole moment

Another neutral system in which measurement of a permanent EDM is theoretically possible is the neutron. In the SM, the neutron has contributions to its permanent EDM from two sources: the permanent EDMs of its constituent quarks, and "long-distance" quark interaction processes that can be characterized in terms of meson exchange. In the former case, the up and down quark EDMs arise only at the three loop level. Their combined contribution to the neutron EDM is approximately $-0.9 \times 10^{-34}\, e \cdot \mathrm{cm}$ [32]. On the other hand, long-distance meson-exchange processes are expected to give a much larger contribution, with an upper bound of $6 \times 10^{-32}\, e \cdot \mathrm{cm}$ [33]. The measurement of a neutron EDM substantially higher than this value would be a direct signal of BSM physics.

As with most neutral systems, the experimental determination of the neutron EDM is essentially a measurement of transition rates. In the non-relativistic limit, a neutron possessing a permanent EDM at rest in parallel magnetic and electric fields along the z-direction, say, is described by the Hamiltonian

$$H = -\sigma_z(\mu B_0 + dE_0)\,, \tag{1.6}$$

where μ is the magnitude of the magnetic dipole moment of the neutron, d is the magnitude of its permanent EDM, and σ_z is the familiar Pauli 2×2 matrix, $\mathrm{diag}(1, -1)$, reflecting the fact the the neutron is a spin-1/2 particle. This Hamiltonian has two energy eigenstates with $\varepsilon = \pm(\mu B_0 + dE_0)$. In an arbitrary superposition of these, the expectation values of the spin projections S_x and S_y will evolve in time with angular frequency $\omega_{\uparrow\uparrow} = 2(\mu B_0 + dE_0)/\hbar$, determined by the energy difference. If the fields are made antiparallel, the evolution frequency would instead be $\omega_{\uparrow\downarrow} = 2(\mu B_0 - dE_0)/\hbar$. The neutron EDM is the difference between these two frequencies, that is

$$d = \frac{\hbar}{4E_0}(\omega_{\uparrow\uparrow} - \omega_{\uparrow\downarrow})\,. \tag{1.7}$$

When d is on the order of the SM predictions, this frequency difference is rather small, even for large electric fields, and poses significant experimental challenges.

To appreciate these challenges, it helps to introduce the following simple picture. It may be shown that the expectation value of the angular momentum $\langle \boldsymbol{S}(t) \rangle$ obeys the classical equation of motion [34]

$$\frac{\mathrm{d}}{\mathrm{d}t}\langle \boldsymbol{S} \rangle = \langle \boldsymbol{S} \rangle \times \left(\frac{\mu}{S}\boldsymbol{B}(t) + \frac{d}{S}\boldsymbol{E}(t) \right)\,. \tag{1.8}$$

In the case of static, parallel fields $\boldsymbol{E} = E_0\hat{\boldsymbol{z}}$ and $\boldsymbol{B} = B_0\hat{\boldsymbol{z}}$, the vector $\langle \boldsymbol{S} \rangle$—often referred to as simply "the spin"—precesses on a cone about $\hat{\boldsymbol{z}}$ at angular frequency $\omega_0 = (\mu B_0 + dE_0)/S$, where $S = \hbar/2$. In a magnetic field only, this is called the Larmor frequency, $\omega_L = \gamma B_0$, where $\gamma \equiv \mu/S$ is the neutron gyromagnetic ratio.

This value has been determined experimentally to high precision, and is $-1.832\,471\,71(43) \times 10^{8}\,\mathrm{s^{-1}\,T^{-1}}$ [35], or approximately $-3 \times 10^{7}\,\mathrm{Hz\,T^{-1}}$. (The minus sign indicates that the spin and magnetic moment are antiparallel.) In contrast, the largest d predicted by the long-distance meson-exchange processes in the SM is $6 \times 10^{-32}\,e \cdot \mathrm{cm}$, or $10^{-53}\,\mathrm{C \cdot m}$ in SI units. Thus, $d/S \approx 3 \times 10^{-19}\,\mathrm{Hz\,m\,V^{-1}}$, where the units of the electric field are Volt per meter. Roughly speaking then, the signal to noise ratio is $(dE_0/S)/(\mu B_0/S) = 10^{-15} E_0/B_0$, where E_0 is now in kilovolt per centimeter and B_0 is in microtesla, units that are common in neutron EDM experiments. The best magnetically shielded rooms can reduce ambient fields to the level of 100 fT [36], but for technical reasons, B_0 is typically on the order of 1 µT. Electric fields achievable in neutron EDM experiments are typically on the order of $10\,\mathrm{kV\,cm^{-1}}$, with techniques currently in development to generate fields as high as $75\,\mathrm{kV\,cm^{-1}}$ [37]. Using this larger value for E_0 and rounding generously, the signal to noise is then only on the order of 10^{-13} at best in the case of d predicted by the SM.

The most common approach in EDM experiments to increase resolution of the precession frequency ω_0 is the technique of Ramsey spectroscopy [38]. A single measurement consists of the following steps:

A Single (Idealized) "Ramsey" Measurement

1. The neutron sits in colinear fields $\boldsymbol{E}_0$ and $\boldsymbol{B}_0$, and is prepared in the spin state $|S_z = \hbar/2\rangle$.
2. A magnetic field $\boldsymbol{B}_1$, orthogonal to $\boldsymbol{B}_0$ and rotating at frequency ω, is applied for a pulse of time t_1.
3. The neutron spin precesses about the static fields for a time T.
4. The rotating $\boldsymbol{B}_1$ field is applied again for time t_1, in phase with the original pulse.
5. The neutron spin projection S_z is measured.

The probability $P(\omega)$ to measure a transition to the state $S_z = -\hbar/2$ is [39]

$$P(\omega) = \frac{4\omega_1^2}{\Omega^2} \sin^2 \frac{\Omega t_1}{2} \left[\frac{\Delta}{\Omega} \sin \frac{\Omega t_1}{2} \sin \frac{\Delta T}{2} - \cos \frac{\Omega t_1}{2} \cos \frac{\Delta T}{2} \right]^2 , \tag{1.9}$$

where $\Delta = \omega - \omega_0$, $\omega_1 = \gamma B_1$, and $\Omega = \sqrt{\Delta^2 + \omega_1^2}$. This function is plotted in Figure 1.1 for the case when $\omega_1 t_1 = \pi/2$, so that the central peak occurs at $\omega = \omega_0$. The above steps are repeated at different values of ω, whereupon a fit of $P(\omega)$ to the experimental transition probabilities allows a determination of ω_0 to be made. Suppose that in this first determination of ω_0, the fields are parallel. When the entire procedure is repeated for $\boldsymbol{E}_0$ and $\boldsymbol{B}_0$ antiparallel, a fit to $P(\omega)$ should return a different value of ω_0 if d is non-zero. The difference in these frequencies will then give the neutron EDM according to Equation 1.7.

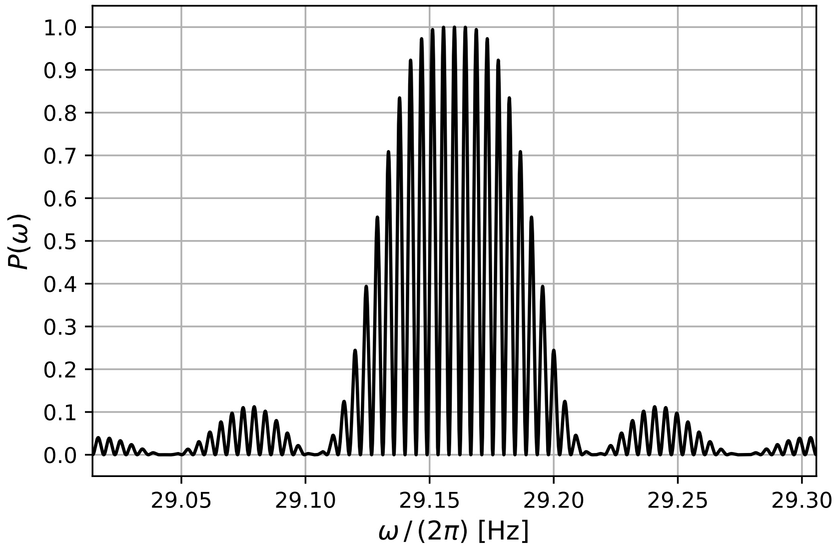

Figure 1.1: Example of ideal transition probability $P(\omega)$—"Ramsey fringes"—for $B_0 = 1.0\,\mu\mathrm{T}$, $B_1 = 0.5\,\mathrm{nT}$, $\omega_1 t_1 = \pi/2$, and $T = 12t_1 \approx 206\,\mathrm{s}$ so that the B_1 pulses are in phase.

An estimate of the frequency difference, and thus of d, from nonlinear fits of $P(\omega)$ to experimental data will have a statistical uncertainty that depends in some way on the number of data points, their sampling along ω, etc. In general, this uncertainty should shrink as more data are collected, but the specific way in which this occurs depends on the fitting routine used, and is best-suited to study by Monte Carlo methods. Nevertheless, an indicator of this uncertainty, referred to as the statistical sensitivity of an EDM experiment, is often given as [40]

$$\delta\hat{d} = \frac{\hbar}{E_0 T\sqrt{N}}\,. \tag{1.10}$$

This figure arises from the statistical uncertainties of the measured transition probabilities. At a given value of ω, the estimate of the transition probability is given by $\hat{P} = N_\downarrow/(N_\uparrow + N_\downarrow)$, where $N_\uparrow$ and $N_\downarrow$ are the numbers of neutrons measured with $S_z = +\hbar/2$ and $S_z = -\hbar/2$, respectively. $N_\downarrow$ is a binomial random variable with mean NP and standard deviation $\sqrt{NP\,(1-P)}$, where $N = N_\uparrow + N_\downarrow$. The estimate $\hat{P}$ will thus have mean P and uncertainty $\sqrt{P\,(1-P)/N}$. Now, Equation 1.9 can be Taylor expanded in $\delta\omega$ about $\omega = \omega_0$, in which case $P(\omega) \approx 1 - (\delta\omega\, T/2)^2$, where additional terms from ω_1 are neglected on the basis that $T \gg 1/\omega_1$ in typical experiments. Roughly speaking then, interpreting $\delta\omega$ as an uncertainty on $\hat{\omega}_0$ gives $\delta\hat{\omega}_0 \sim 1/T\sqrt{N}$. The EDM sensitivity is then determined using this result in combination with Equation 1.7. The details of this entire procedure, however, are murky in the literature, and different sources include various factors of 2 and so forth. (See, for example [41].)

Nevertheless, the qualitative behavior of the EDM statistical sensitivity in Equation 1.10 provides an intuitive notion of how to improve a measurement of the neutron EDM: generate as large an electric field E_0 as possible, allow the neutrons to precess for time T as long as possible, and measure as many neutrons N as possible. Of course, neutrons undergo beta decay at random, meaning that N is a decreasing function of T, and providing an immediate caveat regarding the latter two instructions of this three-part recipe. The motivation for, and focus of, this thesis will mainly be to reduce the EDM statistical sensitivity by increasing N subject to this and other considerations.

1.4 The PanEDM experiment at the SuperSUN source

In order to contextualize the goal of this thesis to increase N, the number of neutrons counted at the end of a single Ramsey measurement cycle, it is helpful to briefly review the history and current status of neutron EDM experiments. The earliest neutron EDM experiment was published in 1957 using the Ramsey spectroscopy technique on a beam of neutrons, and established $d \leq -0.1(24) \times 10^{-20}\, e \cdot \mathrm{cm}$ as the first upper bound to the permanent neutron EDM [42, 43]. Twenty years later, technological improvements had brought this limit down by four orders of magnitude to $d \leq 0.4(15) \times 10^{-24}\, e \cdot \mathrm{cm}$ [44], but the prevailing opinion at this time was that systematic effects arising from beam divergence and relativistic corrections—namely the additional magnetic field proportional to $\boldsymbol{v} \times \boldsymbol{E}_0$ in the neutron frame, where $\boldsymbol{v}$ is the neutron velocity in the lab frame—were preventing further progress from being made in beam experiments [16]. (Interestingly, a beam experiment is currently under development that proposes to take advantage of the high peak flux and time structure of next-generation neutron spallation sources [45, 46].)

Thereafter, the paradigm shifted away from beam experiments, and in 1984 a new limit $d \leq 0.3(48) \times 10^{-25}\, e \cdot \mathrm{cm}$ was established using the Ramsey spectroscopy technique on stored neutrons [47]. (The energy regime in which neutrons can be stored is referred to as ultracold; see the next chapter for a discussion of the physics of ultracold neutrons.) Since this measurement, incremental progress has been achieved by improvements in neutron storage technology and through a better understanding of systematic effects. As of 2020, the neutron EDM upper limit is now $d \leq (0.0 \pm 1.1_{\mathrm{stat}} \pm 0.2_{\mathrm{sys}}) \times 10^{-26}\, e \cdot \mathrm{cm}$ [48], where the $\pm 1.1_{\mathrm{stat}}$ refers to the EDM statistical sensitivity and $\pm 0.2_{\mathrm{sys}}$ refers to a systematic uncertainty.

The idealized discussion of the Ramsey spectroscopy technique in the previous section makes no mention of these systematic effects, nor of the experimental difficulties faced in real measurements. For one, neutrons can be depolarized in a variety of ways during storage. (See [49] for a particularly interesting mechanism for this.) Such depolarized neutrons will diminish the peak of the transition probability $P(\omega = \omega_0)$ to a level $\alpha = (N_\uparrow - N_\downarrow)/(N_\uparrow + N_\downarrow)$ referred to as the polarization visibility [41]. The probability of depolarization generally increases with increased precession time T, so that making T "as long as possible" is not desirable. Equally important are the spatial uniformity and temporal stability of the applied fields, critical to the management of systematic effects and to the the meaningful application of Equation 1.7 in determining d.

The PanEDM collaboration experiment for the measurement of the permanent neutron EDM is a stored neutron experiment currently under construction at the Institut Laue-Langevin neutron facility in Grenoble, France. Its goal is to decrease the current upper limit of d by up to two orders of magnitude [50], mainly through an increase in statistics—to be supplied in part by the SuperSUN ultracold neutron (UCN) source, also currently under construction [51]—and through a significant reduction in systematic effects, most notably by way of unprecedented magnetic shielding [36]. The two coating materials studied in this thesis—CYTOP and deuterated polyethylene (dPE), to be described in greater detail later on—were investigated with the goal of increasing N for the PanEDM measurement, and are envisioned for use in two specific situations.

The first of these concerns the UCN source SuperSUN. The number of neutrons delivered to the storage cells of the PanEDM experiment is dictated by the output of this source, which operates by building up a high density of UCN that are released to the EDM experiment before each Ramsey measurement cycle. This density is determined in large part by the coating materials used within its production volume. We will delay a more detailed discussion of the physics behind coatings for UCN storage until the next chapter, but, broadly speaking, a "good" coating is one that presents a large potential barrier and a low loss cross section to a trapped neutron. Thus, a good coating in the SuperSUN source produces a high density of trapped UCN, which in turn leads to a high output flux, and finally, enables a large N. There are also secondary properties that are desirable in coatings for specific applications. For example, cryogenic suitability is important for coatings inside the production volume of the SuperSUN source, which is filled with superfluid ^{4}He. One commonly-used coating material in low-temperature applications is Fomblin [52, 53], which is a fluorinated lubricant available in grease and oil form. One of the primary difficulties in implementing Fomblin coatings, however, is the lack of a scalable technique to uniformly apply the lubricant to large surfaces—it is typically prepared manually. It is also extremely delicate thereafter, and its durability throughout repeated thermal cycling of the source—upon maintenance, for example—is questionable. In this thesis, the material CYTOP is investigated as an alternative fluoropolymer coating to Fomblin lubricant. While completely new in UCN storage, it is widely-used as an insulating coating in industrial electronics applications [54]. Its formulation, designed to be easily applied to to surfaces in well-controlled thicknesses, makes it a particularly attractive potential substitute for Fomblin.

The second application envisioned for the other coating studied in this thesis, dPE, is for use in the storage cells of the PanEDM experiment, which is where the Ramsey measurements occur. There, the desired secondary properties of the coating are not as for in the UCN source, the most obvious difference being that a large electric field must be applied within the cells. This necessitates at least two coating materials: one conducting for the high voltage electrodes supplying the field, and the other insulating to separate the electrodes. It is expected that diamond-like carbon (DLC), which has been studied in great detail for use in UCN applications [55, 56, 57, 58], will be used as the coating for the conductive electrodes. dPE, on the other hand, is anticipated to be used as the insulating coating, and has been studied to some degree in previous works [59, 60]. However, owing to its rather niche use in experimental neutron physics, the purchase of quantities of dPE sufficient to produce relatively large coatings has proven challenging. Variation between suppliers necessitates the explicit testing of a given product with UCN storage measurements.

Chapter 2

The physics of ultracold neutron storage

In this chapter, an overview of the physics enabling ultracold neutron (UCN) storage is presented. All four fundamental interactions are relevant to the discussion: the strong force, the weak force, the electromagnetic force, and the gravitational force. The strong interaction, which gives rise to an effective neutron optical potential, V, and is responsible for setting the neutron energy scale of “ultracold”, is discussed first. A brief summary of the remaining forces is then given, as is relevant to the UCN energy regime. This is followed by a discussion of idealized UCN storage, in which the time evolution of a stored ensemble of neutrons in a closed vessel is characterized by a single, experimental value, τ_0, known as the storage time constant of the bottle. This constant is a proxy for the more fundamental, bottle-independent parameters, V, and η, the latter of which is known as the neutron loss factor. Several mechanisms for departure from a single storage time constant are then discussed in a manner relevant to the experimental results presented in the next chapter.

In principle, both V and η may be calculated from neutron scattering data; in practice, however, it is necessary to carry out bottle experiments in order to accurately measure the latter, which is arguably the most important parameter in UCN storage. In some sense, the goal of this thesis is the determination of η for the coating materials discussed in the previous section: CYTOP and dPE. Such an endeavor is, in general, complicated by an inability to directly measure the stored UCN energy spectrum, and the difficulty of experimentally determining the neutron scattering properties of possible defects in or contamination of *in situ* coating surfaces. The former challenge is the subject of Chapter 4, while the latter is considered here.

2.1 The neutron optical potential

Having zero net electric charge, the neutron’s external electric field is heavily suppressed. This property often allows neutrons to penetrate deeply into dense materials. Consider, for example, lead, which has number density $n = 3.3 \times 10^{22}\,\mathrm{cm}^{-3}$ and scattering cross section $\sigma = 11.1\,\mathrm{b}$ [2] (1b (barn) $= 10^{-24}\,\mathrm{cm}^2$) for "thermal" neutrons, i.e. those with kinetic energy $E = 25\,\mathrm{meV}$ and velocity $v = 2200\,\mathrm{m\,s}^{-1}$—the most probable velocity in a nuclear reactor for neutrons thermalized with a moderator at 290 K. The classical mean free path is therefore $(n\sigma)^{-1} = 2.7\,\mathrm{cm}$. (Inclusion of the absorption cross section decreases this by about 1.5 %.) An inch of lead shielding thus offers very little protection against thermal neutron radiation. Now, it is true that the classical mean free path is smaller for slower neutrons, because, ignoring nuclear resonances, $\sigma \propto 1/v$ [61]. (This relationship reflects the fact that the interaction probability is proportional to the time the neutron spends near the scatterer.) This trend predicts, for example, a mean free path of about 12 µm in lead for a neutron of velocity $v = 1\,\mathrm{m\,s}^{-1}$, or energy roughly $E = 5\,\mathrm{neV}$. However, at such low energies, the average potential presented by a collection of nuclei will actually cause an incident neutron to be completely reflected from, rather than scattered within, the material.

The starting point for studying this effective interaction is the potential between a nucleus and a nucleon that arises from the strong interaction. This may be approximated by the spherically symmetric Woods-Saxon potential [62]

$$V(r) = -\frac{V_0}{1 + e^{(r - r_0)/a_0}}\,, \tag{2.1}$$

where V_0 is the nucleon separation energy, r_0 is the mean nuclear radius, and a_0 is the "skin thickness". The depth of the potential well, V_0, varies by nuclide, but is generally on the order of tens of MeV. Scattering measurements give $r_0 \approx A^{1/3}\,1.25\,\mathrm{fm}$, for A nucleons, while $a_0 \approx 0.5\,\mathrm{fm}$ is roughly constant throughout the nuclide chart [63]. Figure 2.1 shows an example of the Woods-Saxon potential, plotted in black, for $r_0 = 5\,\mathrm{fm}$, corresponding to $A = 64$ nucleons. Unlike the Coulomb potential, for example, the nuclear potential drops off exponentially outside of the mean nuclear radius, and may therefore be said to be short range, or "localized" to within r_0.

A neutron with incident wavevector $\boldsymbol{k}$ scattering elastically from such a localized potential at the origin will have wavefunction

$$\psi(\boldsymbol{r}) = \psi_{\mathrm{inc}} + \psi_{\mathrm{sc}} = e^{i\boldsymbol{k}\cdot\boldsymbol{r}} + f(\theta)\frac{e^{ikr}}{r}\,, \tag{2.2}$$

far away, that is, when $r \gg r_0$. (See, for example, Chapter 11 in [64].) At low energies, when the wavelength, λ, is large compared to the size of the nucleus, r_0, the scattering amplitude, $f(\theta)$, is approximately constant. An intuitive, semiclassical argument for this result is as follows. For the neutron to have scattered from $V(r)$, its impact parameter must have been on the order of r_0. Its orbital angular momentum is then $L \sim$

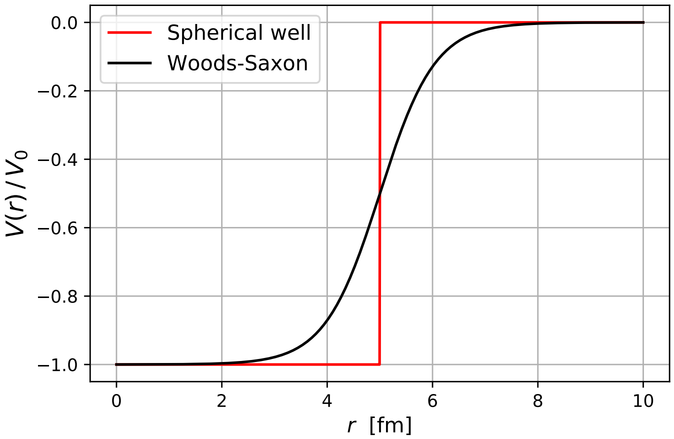

Figure 2.1: Example of the Woods-Saxon potential $V(r)$ of Equation 2.1 plotted in black, and the spherical "square" well plotted in red, for $r_0 = 5\,\mathrm{fm}$, corresponding to $A = 64$ nucleons.

$r_0\,h/\lambda \approx 0$. Therefore, $f(\theta)$ has no angular dependence, otherwise $L \neq 0$. If $V(r)$ is further approximated by a spherical "square" well—see red plot in Figure 2.1, the appropriate value of $f(\theta) = f$ may be determined by matching boundary conditions of the wavefunction at r_0, and taking the limit of large r. With this value in hand, Equation 2.2 says that $\psi \approx 1 + f/r$ for r small compared to λ: thus, it is approximately true that ψ vanishes at $r = -f$. If f is replaced by $-a$, that is

$$\psi(\boldsymbol{r}) = \psi_{\mathrm{inc}} + \psi_{\mathrm{sc}} = e^{i\boldsymbol{k}\cdot\boldsymbol{r}} - a\frac{e^{ikr}}{r}\,, \tag{2.3}$$

it may be seen that a repulsive, hard sphere potential of radius a gives the equivalent far field result for ψ [3], and a is thus referred to as the scattering length, and is on the order of several femtometer [3]. (For some values of V_0 and r_0, it happens that $a < 0$, in which case this intuitive picture does not work because the hard sphere cannot have a negative radius.)

We are interested in how collective scattering from many nuclei, rather than from a single nucleus, can give rise to an effective neutron potential. This is most readily examined by considering the integral form of the Schrödinger equation applied to the present case. (For this and many other results discussed in the present section, see again, for example, Chapter 11 in [64].) For an incident neutron with wavevector $\boldsymbol{k}$, this is

$$\psi(\boldsymbol{r}) = e^{i\boldsymbol{k}\cdot\boldsymbol{r}} - \frac{m}{2\pi\hbar^2}\int \mathrm{d}^3r'\,\frac{e^{ik|\boldsymbol{r}-\boldsymbol{r}'|}}{|\boldsymbol{r}-\boldsymbol{r}'|}V(r')\psi(\boldsymbol{r}')\,. \tag{2.4}$$

Now, if the potential $V(r)$ is small compared to the neutron energy, i.e. for long neutron wavelengths, this expression may be evaluated using the Born approximation, in which $\psi(r') \approx e^{i\boldsymbol{k}\cdot\boldsymbol{r}'}$ inside the integral [64]. In the long wavelength limit, Equation 2.4 simplifies greatly, and the scattering amplitude may be read off as

$$f = -\frac{m}{2\pi\hbar^2}\int \mathrm{d}^3r\, V(r)\,. \tag{2.5}$$

Of course, this result is not valid for the deep nuclear potential, where $V_0 \gg E = \hbar^2k^2/2m$. However, in light Equation 2.3, the "Fermi pseudopotential" [65]

$$V_{\mathrm{F}}(\boldsymbol{r}) = \frac{2\pi\hbar^2}{m} a\,\delta^{(3)}(\boldsymbol{r})\,, \tag{2.6}$$

gives the correct value of f due to a single scatterer, and thus the correct far-field behavior of ψ_{sc}, when inserted into Equation 2.5. The average potential, V, from many scatterers is thus

$$V(\boldsymbol{r}) \approx V = \frac{1}{\Omega}\int_{\Omega} \mathrm{d}^3r\, \frac{2\pi\hbar^2}{m}\sum_i a_i\,\delta^{(3)}(\boldsymbol{r}-\boldsymbol{r}_i) = \frac{2\pi\hbar^2}{m}\frac{\sum_i a_i}{\Omega} = \frac{2\pi\hbar^2}{m}\rho\,\langle a\rangle\ , \tag{2.7}$$

where the volume Ω that contains many nuclei, say N, is small compared to λ^3, $\rho = N/\Omega$ is the number density of the nuclei, and $\langle a\rangle$ is the average scattering length.

The average potential V is often referred to as the neutron optical potential because Equation 2.4 may be recast into the wave equation for a medium with index of refraction given by $\sqrt{1 - V/E}$ [3]. For many materials, the number density ρ is on the order of $10^{22}\,\mathrm{cm}^{-3}$, so that the "scattering length density" can become $\rho\,\langle a\rangle = \mathrm{SLD} \sim 10^{-6}\,\mathrm{\AA}^{-2}$. According to Equation 2.7—and with SLD in these peculiar units—the optical potential in neV is $V \approx 26 \times \mathrm{SLD}$. (See Table 2.1 later for some calculated values for several materials.) The corresponding velocity of a neutron at such energies is several meters per second.

If a neutron of energy $E < V$ is incident on the step potential V, it will be completely reflected, rather than scattered within the material. (See Appendix Section A.1 for a derivation of this result.) Therefore, if a box is made out of a material with optical potential V, and a neutron with energy $E < V$ is placed inside, it will not be able to leave. In this sense, the neutron can be said to be stored inside the box, and is thus an ultracold neutron, or UCN. With this definition of the UCN energy scale in hand, it is now appropriate to discuss the remaining neutron forces—the weak force, the electromagnetic force, and the gravitational force—as are relevant to UCN storage.

2.2 The other interactions

The weak force is responsible for the random beta decay of free neutrons via

$$n \to p + e + \bar{\nu}_e \,, \tag{2.8}$$

where the neutron decays into a proton, an electron, and an antielectron neutrino. The average decay time in the rest frame of the neutron—or "lifetime"—is $\tau_\beta = 879.4(16)\,\mathrm{s}$, according to the current estimate [6]. (At UCN energies, $E \sim 100\,\mathrm{neV}$, the velocity is $v \sim 4\,\mathrm{m\,s^{-1}}$; thus $v/c \ll 1$, where $c \approx 3 \times 10^8\,\mathrm{m\,s^{-1}}$ is the speed of light, and the effects of relativistic time dilation are much smaller than the uncertainty on τ_β.) The main source of uncertainty in the neutron lifetime is due to the discrepancy between the results of two categories of experiments in which τ_β is measured: beam measurements and storage measurements. (See, for example, [66].) In the former, a beam of known flux is monitored for decay products, while in the latter, UCN are stored for a time and then counted. The two experimental methods have produced values of τ_β which disagree by roughly 1 %, with the smaller values coming from storage measurements [67]. A number of potential explanations for this discrepancy have been proposed, ranging from systematic errors in the beam experiments [67], to exotic, dark matter decay channels in the storage experiments, although recent results appear to have ruled out this latter possibility [68]. Nevertheless, the random decay of neutrons places an upper limit of τ_β on any time scale for neutron storage inside a box.

Although not a focus of this thesis, neutrons may also be stored in magnetic traps. As mentioned in Section 1.3, the neutron spin is related to its magnetic moment by the gyromagnetic ratio, γ. The interaction of $\boldsymbol{\mu}$ with the magnetic field, $\boldsymbol{B}(\boldsymbol{r})$, may then be written as

$$V_m(\boldsymbol{r}) = -\boldsymbol{\mu} \cdot \boldsymbol{B}(\boldsymbol{r}) = +|\gamma|\, \boldsymbol{S} \cdot \boldsymbol{B}(\boldsymbol{r}) \,, \tag{2.9}$$

where the sign is appropriate for neutrons, whose spin and magnetic moment are opposite, i.e. $\gamma = \mu/S \approx -1.832 \times 10^8\,\mathrm{rad\,s^{-1}\,T^{-1}}$. The size of μ in the present units of interest is about $60.3\,\mathrm{neV\,T^{-1}}$, so that in a 1 Tesla field, a spin parallel to the magnetic field $\boldsymbol{B}$, will experience a potential roughly $120.6\,\mathrm{neV}$ greater than will a spin antiparallel to $\boldsymbol{B}$. The corresponding magnetic force may be expressed as

$$\boldsymbol{F}_m(\boldsymbol{r}) = -\boldsymbol{\nabla} V_m(\boldsymbol{r}) = -|\gamma|\, \boldsymbol{S} \cdot \boldsymbol{\nabla} \boldsymbol{B}(\boldsymbol{r}) \,. \tag{2.10}$$

Thus, a positive field gradient $\boldsymbol{\nabla}\boldsymbol{B}(\boldsymbol{r}) > 0$ will either repel or attract a neutron whose spin is parallel or antiparallel to $\boldsymbol{B}(\boldsymbol{r})$, respectively; in contrast to the neutron optical potential, the magnetic potential may only be used to store one spin state. The experimental hurdle to take advantage of such an effect in UCN storage would then seem to be the construction of a container in which the neutron always sees an increasing field gradient at the walls.

The difficulty of this problem is mitigated by an experimental property of UCN known as adiabatic spin transport. As discussed previously, the neutron spin—or expectation value of the spin—will precess about a static magnetic field at the Larmor frequency, $\omega_L = \gamma B$, or $\omega_L/2\pi \approx 29.2 \times 10^6\,\mathrm{Hz\,T^{-1}}$. (See Equation 1.8.) Now, if the field direction is changed "adiabatically"—roughly meaning that that Larmor frequency is very large compared to the rate of change of the field in the neutron frame of reference—it may be shown that a polarized neutron will continue to precess about the instantaneous field direction [34], and thereby maintain its orientation with respect to $\boldsymbol{B}(\boldsymbol{r})$. Therefore, as long as the total field does not vanish, a spin that is repelled by a positive field gradient in one corner of a trap, will likewise be repelled by a negative field gradient in another corner, because it will have followed the closed field lines present inside the container. Various "multipole" designs exist that present large field gradients at trap boundaries; for several examples, see Section 2.2.1 in [66]. In conjunction with high optical potential walls, such an approach is capable of greatly enhancing the UCN storage characteristics of a container [51], and, for example, is planned in the second development phase of the SuperSUN UCN source [50].

The remaining fundamental interaction, gravity, does not necessarily influence the amount of time for which a neutron may be stored in a box. However, there are several gravitational effects that play an important role in UCN storage. At extremely low energies, on the peV scale, neutrons have been shown to occupy bound states of the linear gravitational potential at the surface of the Earth [69]. For slightly higher energies, still within the UCN regime, the expectation value of the position of a stored UCN follows the classical trajectory of a bouncing ball [3]. The energy scale of the gravitational interaction is set by $mg \approx 103\,\mathrm{neV\,cm^{-1}}$, where m is the neutron mass and g is the acceleration due to gravity on the surface of the Earth. Therefore a UCN with energy $E = 100\,\mathrm{neV}$ can rise no higher than roughly one meter. This is an important design concept in the construction of some storage containers, which, if tall enough, may be left open at the top [55, 70].

Gravity is also responsible for a variation of UCN density with height in a storage container. The same is true for a classical gas of particles at temperature T, in which the density changes as a function of height y like $\exp(-mgy/k_BT)$. However, the thermodynamic notion of temperature for a gas of UCN is not applicable, because inelastic collisions with the walls of a container will tend to scatter UCN out of their low energy regime, and neutron-neutron collisions are so infrequent, due to the short range of the interaction and the low densities achievable, as to be entirely negligible. Nevertheless, an approximate spatial dependence of the density may be developed using Liouville's theorem, which states that the phase-space density of a collection of particles obeying Hamilton's equations is a constant in time [71]. Strictly speaking, this is not true for a stored ensemble of UCN, because, among other things, they may be lost via beta decay, absorption, etc. However, under the assumption that such effects are small—see Section 4.3.1 in [3]—it may be shown that a stored ensemble of UCN obeys a modified Liouville's theorem, in which the phase-space density decays uniformly as $\exp(-t/\tau_\beta)$. According to this modification, the number of UCN, $n(t, E, y)\,\mathrm{d}E$, at time t, height y, and with energy in the range E to $E + \mathrm{d}E$, varies as

$$n\big(t, E, h_2\big) = n\big(t, E, h_1\big)\sqrt{\frac{E - mg\,(h_2 - h_1)}{E}}\,, \tag{2.11}$$

where $h_2 > h_1$. This effect becomes important for tall storage containers, or those whose vertical dimensions are large compared to E/mg.

Finally, the gravitational force is often useful experimentally, in either creating velocity filtering chicanes—see, for example, Figure 2 in [72]—or in detecting UCN, say, after their storage. In its simplest form, a common neutron detector consists of a high voltage wire placed in a chamber filled with a mixture of ^{3}He and Ar gas. (See, for example, Chapter 14, Section III B in [73].) Upon neutron capture, the ^{3}He nucleus will split into a ^{3}H nucleus and a proton, and these fission fragments will ionize the Ar gas in the detector, liberating a cloud of electrons. These electrons are accelerated by the high voltage to the collection wire, whose current is monitored in order to detect UCN. The gas mixture is typically separated from the vacuum of the UCN storage experiment by a thin aluminum foil, whose optical potential is 54 neV [3]. If this window is positioned at the height of the storage bottle, UCN with $E < 54\,\mathrm{neV}$ emptied from the experiment are not counted, as they are not capable of entering the detector. Placing the entrance roughly 54 cm lower then ensures that all falling UCN emptied from the bottle are accelerated to velocities sufficient to overcome the potential barrier posed by the aluminum foil.

2.3 The storage time constant

The discussion so far has not carefully considered any mechanism for loss during UCN storage other than that of beta decay. This would suggest that the neutron population in all storage containers, or at least those without large holes, should decay exponentially like $\exp(-t/\tau_\beta)$. This is, however, not the case, as wall interactions play a large role in loss of UCN during storage. This effect can be considered analytically by the addition of an imaginary term $-iW$ to the real potential V. (The negative sign is required so that the effect of W is to decrease the probability density $|\psi|^2$.) The resulting complex optical potential is [3]

$$U = V - iW = \frac{2\pi\hbar^2}{m}\rho\,\langle a\rangle - i\frac{\hbar}{2}\rho\,\langle\sigma\rangle\,v\,, \tag{2.12}$$

where v is the neutron velocity, and $\langle\sigma\rangle$ is the sum of all loss cross sections—absorption, inelastic scattering, etc.—averaged over the different elements and isotopes of the material. In general, $\sigma \propto 1/v$ [61], so that the product $\langle\sigma\rangle\,v$ is velocity independent.

The reflection probability per wall interaction, or per "bounce", for a neutron of energy $E < V$ incident on the complex step potential U is no longer unity for nonzero W. Instead, if $W/V \ll 1$, which will be the case for reasonable storage materials, the loss probability for $E < V$ is

$$\mu(E) = 2\,\eta\sqrt{\frac{E}{V-E}}\,, \tag{2.13}$$

where the energy independent parameter $\eta = W/V$ is known as the loss factor. See Appendix Section A.1 for a derivation of this result, and Table 2.1 for a list of predicted values of η for several materials. Disregarding practical concerns such as ease of use, cost, and so forth, a good material for UCN storage is then one that possesses a large optical potential V and a small loss factor η: the former ensures that the material will store many neutrons over a large energy range, while that latter implies that the probability of loss with each wall interaction will be small.

Although Equation 2.12 may be used to predict good candidate storage materials, it is necessary to directly measure V and η in order to ensure that, for example, a given coating is free of contamination. The optical potential is frequently determined by neutron reflectometry. (See, for example, Chapter 4 in [74], or Figure 3.7 in the present work.) Typically, the technique consists of illuminating a sample with a beam of "cold" neutrons (CN)—$E \sim 0.1$ to $10\,\text{meV}$, wavelength $\lambda \sim 30$ to $3\,\text{Å}$ [5]—at a shallow angle to the surface and

	$\langle a \rangle$ [fm]	$\langle \sigma_{\text{abs}} \rangle$ [b]	V [neV]	$\eta \times 10^{-5}$
H_2O	-1.7	0.67	-15	-11^{*}
Superfluid ^{4}He	3.3	0	19	0
Al	3.5	0.23	54	1.9
Pb	9.4	0.17	81	0.5
Cu	7.7	3.8	171	13.6
Ni	10.3	4.5	242	12.1
BeO	13.6	0.007	261	$1.35^{\dagger}$
Diamond	6.6	0.004	305	0.01

Table 2.1: Average scattering length, $\langle a \rangle$, and average thermal neutron absorption cross section, $\langle \sigma_{\text{abs}} \rangle$, for several materials. Values for elements are given using isotopic averages for natural abundance, and are taken from references [1] and [2], respectively. For the case of compounds, these values are given per molecule. Values of the neutron optical potential, V, and loss factor, η, are calculated. (*): the negative value of η for water, H_2O, does not have the same meaning in the context of the loss probability $\mu(E)$, but nevertheless gives the ratio of W/V. (†): the value of η for BeO is taken from measurement at room temperature, as listed in Table 2.1 of [3], and does not correspond to the value of $\langle \sigma_{\text{abs}} \rangle$ listed here.

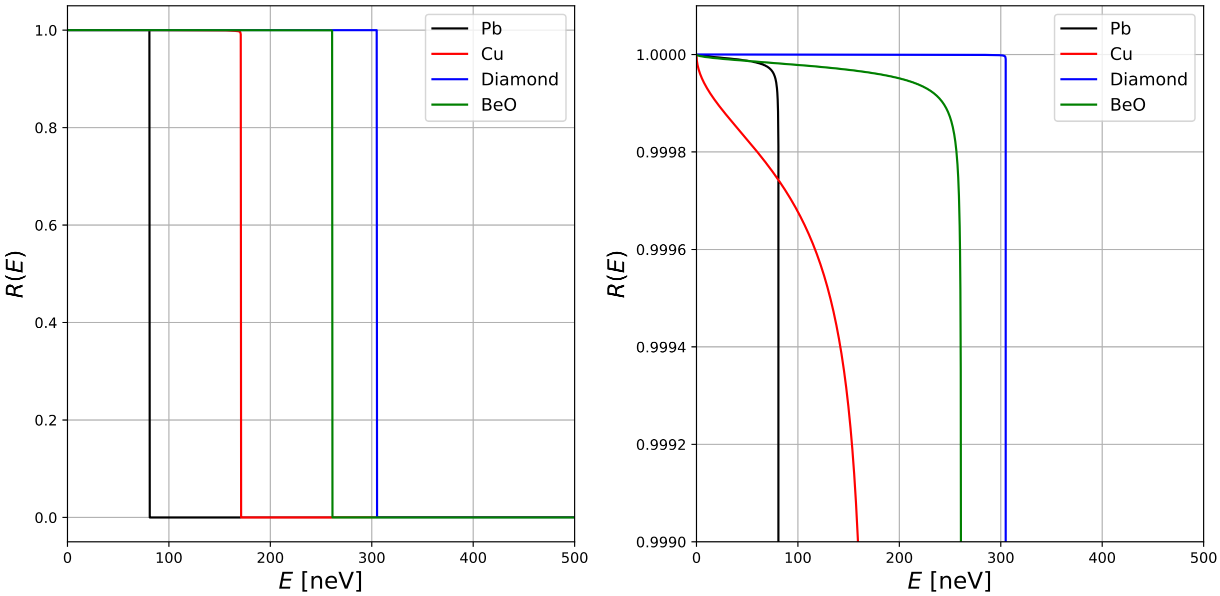

Figure 2.2: Idealized reflectivity, $R(E) = 1 - \mu(E)$, for several materials, see Table 2.1. E is the energy "normal" to the sample surface, i.e. $E = E_{\perp} \equiv \hbar^2 k_{\perp}^2 / 2m$. The plot at right shows an enlarged view of the "flat" portion of $R(E)$, that is, for energy below the optical potential.

measuring the reflection probability, $R(E)$. The independent variable is $k_\perp$, the momentum normal to the sample surface, or equivalently, the "normal" energy $E_\perp \equiv \hbar^2 k_\perp^2/2m$. An idealized form of $R(E)$ is displayed in two plots in Figure 2.2 for several materials, see Table 2.1. The energy at which the reflection probability rapidly drops from nearly unity—see plot on the left—is the optical potential, and can be determined by measuring $R(E)$ in the vicinity of this "critical edge".

While neutron reflectometry is well-suited to the precise determination of V, it is not generally capable of measuring the small loss factors of good materials for UCN storage. To motivate this claim in a quantitative manner, let "good" mean that the size of η will be such that UCN storage times t are on the order of the neutron lifetime τ_β. If the loss probability per wall interaction is μ, then after N collisions, the probability that the UCN has not been lost is $P = 1 - \mu - (1-\mu)\mu - ... - (1-\mu)^{N-1}\mu \approx 1 - N\mu$ for small μ. For a UCN of velocity $v = 1\,\mathrm{m\,s^{-1}}$ in a box of side length $L = 1\,\mathrm{m}$, $N \approx vt/L \sim v\tau_\beta/L = 10^4$. Thus, $P \approx 0$ when $\mu \approx 1/N \sim 10^{-4}$, which implies that a good storage material will have $\eta \sim \mu \sim 10^{-4}$. Now, recall that the reflection probability is $R(E) = 1 - \mu(E)$, and that $\mu(E)$ abruptly changes from $\sim \eta$ to ~ 1 when $E > V$. Thus, as can be seen in the plot on the right of Figure 2.2, η controls the slope of the "flat" region, $E < V$, of $R(E)$. The basic problem in reflectometry regarding the extraction of η from measurement is that the neutron interacts only once with the sample, so that the approximate signal to noise ratio is $\sim \eta$. Furthermore, the idealized plots of $R(E)$ in Figure 2.2 ignore the more complicated physics of real reflectometry measurements and various sources of experimental uncertainty, both of which conspire to make the precise determination of η impractical by such a technique. (See, for example Figure 3.7.)

Determination of η is typically achieved by means of a storage bottle experiment, in which there are a large number of UCN interactions with the sample surface, i.e. the bottle walls. The number of UCN in the bottle as a function of time, $N(t)$, is the measurable quantity. To connect this to η, it is first necessary to extend the one-dimensional formula for $\mu(E)$, Equation 2.13, to three dimensions. This is accomplished by replacing $E \to E\cos^2\theta$, where θ is the angle between the incident neutron and the normal to the surface of the step potential U. One is then interested in the angle-averaged value of $\mu(E,\theta)$ for the ensemble of UCN stored in the bottle. This may be calculated for an isotropic velocity distribution, in which case the flux of UCN on a surface is proportional to $\cos\theta$. (In time Δt, a UCN with velocity v impinging on a surface of

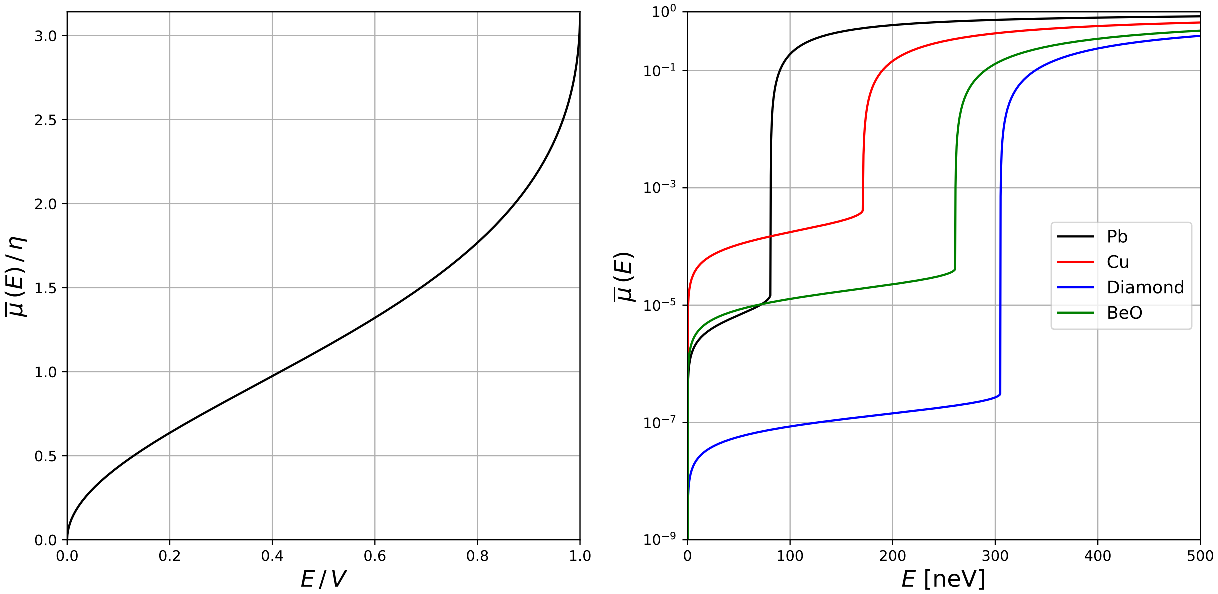

Figure 2.3: Left: $\bar{\mu}(E)/\eta$ plotted against E/V, according to the second line of Equation 2.14. The value rises from 0 at $E = 0$, to π at $E = V$. Right: $\bar{\mu}(E)$ plotted for various materials from Table 2.1. The curves are calculated by numerically evaluating the integral in the first line of Equation 2.14.

area A will sweep out a volume $Av\cos\theta\Delta t$.) Thus, $\mu(E,\theta)$ averaged over 2π steradians is

$$\begin{aligned}\bar{\mu}(E) &= \frac{1}{\pi}\int \mathrm{d}\Omega \cos\theta \left(2\,\eta\sqrt{\frac{E\cos^2\theta}{V - E\cos^2\theta}}\right) \\ &= 2\eta\left[\frac{V}{E}\arcsin\left(\sqrt{\frac{E}{V}}\right) - \sqrt{\frac{V}{E} - 1}\right], \qquad E \leq V\,.\end{aligned} \tag{2.14}$$

The closed form expression on the second line, valid for $E \leq V$, is plotted at left in Figure 2.3. For $E > V$, $\bar{\mu}(E)$ rapidly increases toward unity, and the above integral must be evaluated numerically over the range of θ such that $V > E\cos^2\theta$. This is shown at right in Figure 2.3.

An isotropic velocity distribution has been found to be a reasonable approximation for a stored ensemble of UCN in containers with rough inner surface on the basis of Monte Carlo simulations. (See Section 4.3.2 in [3].) By rough, it is meant that the microscopic structure of the container surface will produce some degree of non-specular reflection of UCN. A more precise notion can be developed by a consideration of the “micro-roughness model” [75], in which perturbation theory is used to calculate the probability distribution for final angles $(\varphi_\mathrm{f}, \theta_\mathrm{f})$ given an incident angle θ_i when a surface has “Gaussian” roughness of RMS height b and correlation length w. (The incident azimuthal angle is taken as $\varphi_\mathrm{i} = 0$.) This model is appropriate for highly-polished steel guides, which are frequently used to transport UCN. For such surfaces, the “settling time" for the UCN ensemble to achieve an isotropic velocity distribution is on the order of $3\,\mathrm{s}$ [3], which—it will be seen—is much smaller than the characteristic UCN storage time. Furthermore, typical storage bottle surfaces are significantly less smooth than polished guides, and so the settling time will in general be even smaller for situations of interest.

The choice of a UCN velocity distribution enables the rate of wall collisions to be determined. From kinetic theory, this rate per unit area is $\frac{1}{4}v\rho(E)\,\mathrm{d}E$ [76] for an isotropic distribution, where v and $\rho(E)\,\mathrm{d}E$ are the velocity and number density, respectively, for UCN with energy in the range E to $E+\mathrm{d}E$ near the wall. If wall interactions and beta decay are the only loss mechanisms considered, then for storage in a container of volume $\mathcal{V}$ and inner surface area A, the number of UCN in the container, that is $n(E)\,\mathrm{d}E = \rho(E)\,\mathrm{d}E\,\mathcal{V}$, will decay according to

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}t}\Big[n(E)\,\mathrm{d}E\Big] &= -\frac{1}{4}v\rho(E)\,\mathrm{d}E\,A\bar{\mu}(E) - \frac{1}{\tau_\beta}\Big[n(E)\,\mathrm{d}E\Big] \\
\frac{\mathrm{d}}{\mathrm{d}t}n(E) &= -\frac{A}{4\mathcal{V}}v\bar{\mu}(E)n(E) - \frac{1}{\tau_\beta}n(E) \\
\frac{\mathrm{d}}{\mathrm{d}t}n(E) &= -\frac{1}{\tau(E)}n(E)\,,
\end{aligned} \tag{2.15}$$

which serves to define the energy dependent, analytical characteristic storage time

$$\frac{1}{\tau(E)} = \frac{A}{4\mathcal{V}}v\bar{\mu}(E) + \frac{1}{\tau_\beta}\,. \tag{2.16}$$

This result neglects gravitational effects, and assumes that v and $\rho(E)$ are uniform over the volume of the bottle. (For the full result, see Equation 4.17 in [3]; Equation 2.16 is appropriate for short storage bottles, whose vertical dimensions are small compared to E/mg.) The absolute scale of the free neutron lifetime, $\tau_\beta \approx 880\,\mathrm{s}$, makes it necessary to choose the geometrical factor of $A/4\mathcal{V}$ in order to plot $\tau(E)$. The most relevant choice is for the bottle geometry used in theses studies, which is, briefly, a nearly cylindrical bottle of diameter $90\,\mathrm{mm}$ and length $589\,\mathrm{mm}$. (See Chapter 3.) When the bottle is closed, the inner surface area is $A = 1839\,\mathrm{cm}^2$ and the volume is $\mathcal{V} = 3814\,\mathrm{cm}^3$, so that $A/4\mathcal{V} = 12.1\,\mathrm{m}^{-1}$. The reciprocal of this quantity may be interpreted as a mean free path, and is approximately $8.3\,\mathrm{cm}$. Various examples of $\tau(E)$ are plotted in Figure 2.4.

The solution to Equation 2.15 is the decaying exponential

$$n(t,E) = n(0,E)e^{-t/\tau(E)}\,. \tag{2.17}$$

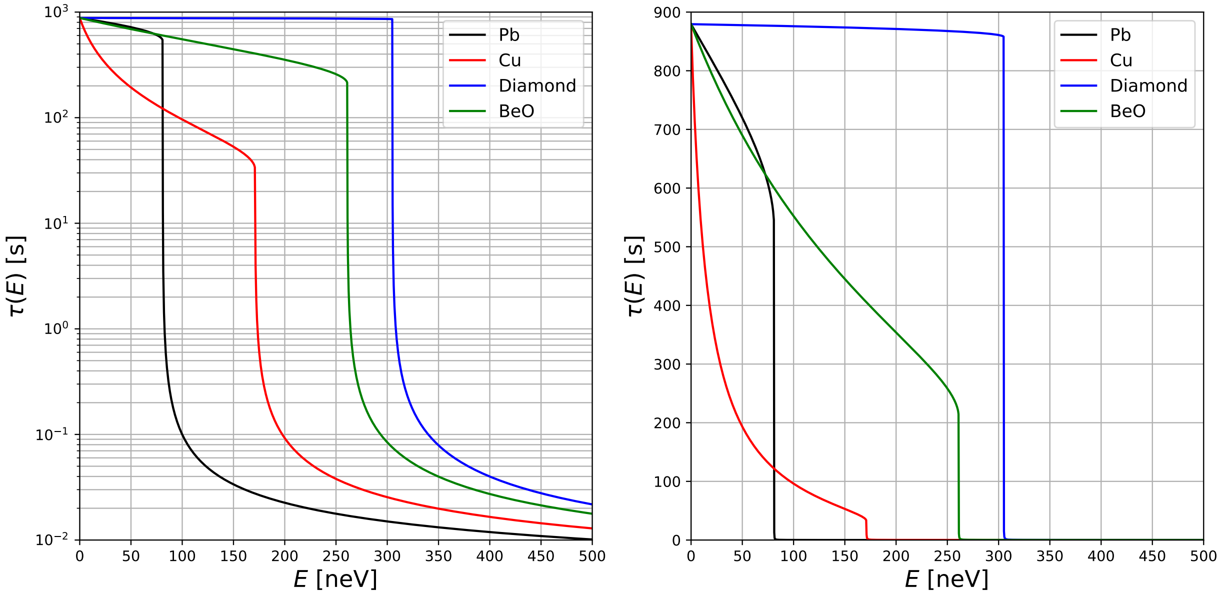

Figure 2.4: Examples of $\tau(E)$ of Equation 2.16 plotted on logarithmic (left) and linear (right) scales for various materials from Table 2.1 and using $A/4\mathcal{V} = 12.1\,\mathrm{m}^{-1}$—mean free path of roughly 8.3 cm—as for the nearly cylindrical bottle geometry used in these studies. (See Chapter 3.) Values were calculated by numerically integrating the first line of Equation 2.14 for $\bar{\mu}(E)$. The asymptotic value of all curves as $E \to 0$ is the free neutron lifetime $\tau_\beta \approx 880\,\mathrm{s}$.

Thus, integrating over all energies E present in the bottle, the total number of UCN is given by

$$N(t) = \int \mathrm{d}E\, n(0, E) e^{-t/\tau(E)} . \tag{2.18}$$

This expression is, finally, the connection between $N(t)$ and the loss factor η. The determination of the latter from the former, however, requires one to "invert" Equation 2.18, which may only be done with knowledge of the initial UCN energy spectrum, $n(0, E)$. We consider this problem in Chapter 4.

When $\tau(E)$ does not vary significantly over the energy range of $n(0, E)$, the function

$$N_0\, e^{-t/\tau_0} \tag{2.19}$$

often provides a reasonable fit to $N(t)$. The quantity τ_0 is then referred to as the experimental "storage time constant". It is important to note that it is characteristic not only of the bottle, with its unique geometry, but of an initial spectrum $n(0, E)$. This single, decaying exponential will typically be a reasonable description when the loss factor η is very small. (The notion of "small" will be developed in the next section.)

At some point, the question arises as to why such elaborate storage measurements to determine η are necessary. After all, neutron absorption and inelastic scattering cross sections for a wide range of energies are tabulated in the literature, and may be inserted into Equation 2.12 to calculate the loss factor, which one may believe should be reasonably close to measured values for well-prepared surfaces. However, it is almost invariably the case that such calculations predict values of η several orders of magnitude smaller than experimental values. (For histories of this "anomalous loss", see Chapter 5 of [3] and Chapter 6 of [77].) In many materials, this increase of η is caused by additional absorption by and upscattering from surface contaminants, such as hydrogen [78, 79], which is almost impossible to completely remove. (Upscattering is jargon for the inelastic scattering of a neutron that carries it out of the UCN energy regime.) The details of this effect are difficult to predict, making it necessary to measure η by UCN storage.

2.4 Multiple storage time constants

Sometimes, the single time constant function of Equation 2.19 does not adequately describe experimental storage data. In such cases, a better fit is typically given by

$$\sum_i N_i \, e^{-t/\tau_i} \, , \tag{2.20}$$

where the series is somewhat arbitrarily terminated depending on the goodness of fit achieved after a given number of terms. Of course, if the initial spectrum of UCN has more than one "peak", that is, if, for example, $n(0,E) \sim \sum_i A_i \delta(E - E_i)$, then the number of neutrons in the bottle, $N(t)$, will decay with multiple time constants, $\tau_i = \tau(E_i)$. Such energy distributions, however, are not relevant to the present work—see Chapter 4—where $n(0,E)$ is expected to be unimodal. Nevertheless, it is still possible to obtain $N(t)$ that appear to have multiple time constants with "singly-peaked" $n(0,E)$.

One of the ways in which this can occur is, roughly speaking, when η is "not small". The meaning of this phrase depends on the size of $A/4\mathcal{V}$, the shape of the initial UCN energy spectrum $n(0,E)$, and the size of the optical potential V. To provide a quantitative example, we choose specific values of these parameters. In keeping with previous example calculations of $\tau(E)$, the value of the geometrical factor $A/4\mathcal{V}$ most relevant to this work is taken as that of the nearly cylindrical bottle geometry used in these studies: that is, $12.1\,\mathrm{m}^{-1}$, corresponding to a mean free path of roughly $8.3\,\mathrm{cm}$. The initial UCN energy spectrum is modeled after the

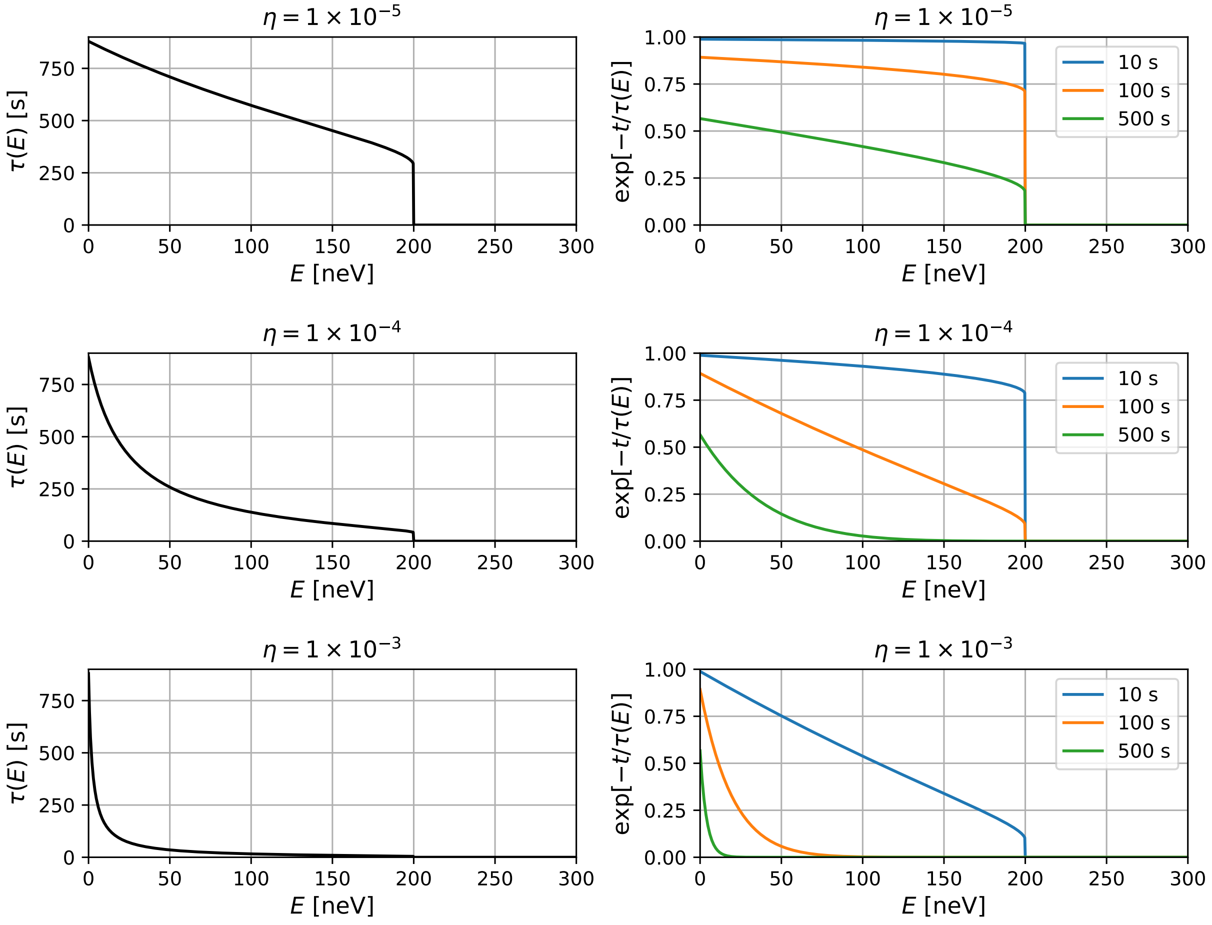

Figure 2.5: Example plots of $\tau(E)$ and $e^{-t/\tau(E)}$ for different values of the loss factor η. The angle-averaged loss probably $\bar{\mu}(E)$ is approximated as unity when $E \geq V$. The geometrical factor $A/4\mathcal{V} = 12.1\,\mathrm{m}^{-1}$, corresponding to a mean free path of roughly 8.3 cm. The optical potential is $V = 200\,\mathrm{neV}$.

Maxwell-Boltzmann distribution, or

$$n(0, E) = 2\sqrt{\frac{E}{\pi}}\frac{1}{\alpha^{3/2}}\exp\left(-\frac{E}{\alpha}\right), \tag{2.21}$$

with $\alpha = 300\,\mathrm{neV}$ playing the role of $k_B T$; such a distribution is roughly equivalent to that at the UCN-EDM beamline at the the ILL [80, 81]. Finally, the optical potential is chosen to be $V = 200\,\mathrm{neV}$, similar that of dPE as reported in the literature [59, 60].

Several plots of $\tau(E)$ are shown at left in Figure 2.5 for different values of η, in which the angle-averaged loss probability $\bar{\mu}(E)$ is approximated as unity when $E \geq V$. (See Equation 2.14.) The number of UCN in the bottle as a function of time, $N(t)$, is obtained by integrating the spectrum $n(0, E)$ against the factor $e^{-t/\tau(E)}$,

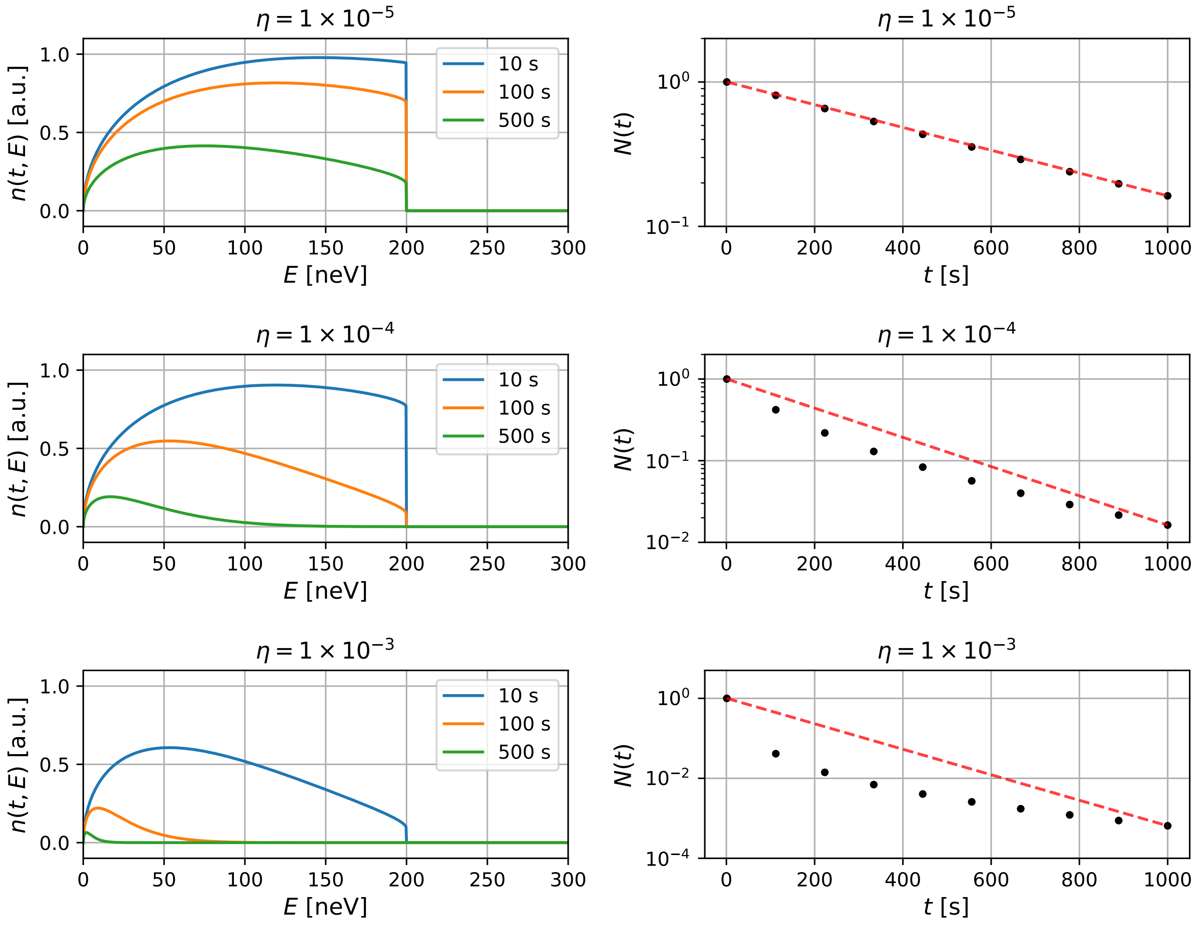

Figure 2.6: Example plots of $N(t)$ and its integrand, $n(t,E)$, for different values of the loss factor η. The angle-averaged loss probably $\bar{\mu}(E)$ is approximated as unity when $E \geq V$. The geometrical factor $A/4\mathcal{V} = 12.1\,\mathrm{m}^{-1}$, corresponding to a mean free path of roughly 8.3 cm. The optical potential is $V = 200\,\mathrm{neV}$. At right, the black dots show the value of $N(t)$ calculated by Equation 2.18, rescaled to unity at time zero. The red dashed line is a guide to the eye in order to see the degree of deviation of $N(t)$ from a single exponential.

and is plotted at right in Figure 2.5. (See Equation 2.18.) The results are shown in Figure 2.6, where $N(t)$ has been rescaled to unity at zero time. At right, it may be seen that already at $\eta = 1 \times 10^{-4}$, $N(t)$ is not described by a single, decaying exponential. This behavior becomes more pronounced as η increases, and may be understood by examining the plots at left. In all cases, the spectrum at long storage times, $n(t,E)$ of Equation 2.17, becomes biased toward low energy, where $\tau(E)$ is larger. For $\eta \gtrsim 10^{-4}$ in this example, the average energy of $n(t,E)$ decreases rapidly enough that its effect on $N(t)$ is clear.

However, in real UCN storage experiments, $N(t)$ may still depart from single, decaying exponential behavior even when η is small. One possible source of this deviation is the presence of some degree of defects in or contamination on virtually all real bottle surfaces. This may be considered analytically by supposing that a small fraction f of the bottle interior consists of "weak patches", with optical potential V_p and loss factor η_p, while the majority of the surface has optical potential $V_w > V_p$ and loss factor η_w. (To the author's knowledge, the name "weak patches" was coined by Dr. Kent Leung.) Equation 2.16 then becomes

$$\begin{aligned} \frac{1}{\tau(E)} &= f\frac{A}{4\mathcal{V}}v\bar{\mu}_p(E) + (1-f)\frac{A}{4\mathcal{V}}v\bar{\mu}_w(E) + \frac{1}{\tau_\beta} \\ &= f\xi\sqrt{E}\bar{\mu}_p(E) + (1-f)\xi\sqrt{E}\bar{\mu}_w(E) + \frac{1}{\tau_\beta} \end{aligned} \tag{2.22}$$

where $\xi = \sqrt{2/m}A/4\mathcal{V}$, and $\bar{\mu}_p(E)$ and $\bar{\mu}_w(E)$ are the angle-averaged loss probabilities of Equation 2.14 using V_p and η_p, and V_w and η_w, respectively. This can be further rearranged as

$$\begin{aligned} \frac{1}{\tau(E)} &= \xi\sqrt{E}\bar{\mu}_w(E) + f\xi\sqrt{E}\Big[\bar{\mu}_p(E) - \bar{\mu}_w(E)\Big] + \frac{1}{\tau_\beta} \\ &= w(E) + p(E) + \frac{1}{\tau_\beta}\,, \end{aligned} \tag{2.23}$$

which defines $w(E)$ and $p(E)$.

To provide a quantitative illustration of the "weak patch" model, we will use the same "main" optical potential, $V_w = 200\,\mathrm{neV}$, the same initial spectrum $n(0, E)$, and the same bottle geometry as in the previous example. ($\xi = \sqrt{2/m}A/4\mathcal{V} \approx 5\,\mathrm{s}^{-1}$ when E is in units of neV.) The main loss factor is chosen to be $\eta_w = 1 \times 10^{-5}$, so that it alone is not sufficient to produce large deviation of $N(t)$ from single, decaying exponential behavior. In this model, as will be demonstrated, UCN with energy above V_p are effectively stored with a loss rate that is $\sim f$, the patch area fraction. In this example then, we choose $f = 1 \times 10^{-4}$, corresponding to an area of roughly $20\,\mathrm{mm}^2$, to produce a loss rate comparable to $\sim 1/\tau_\beta$. The optical potential of the patches is set at $V_p = 100\,\mathrm{neV}$, and $\eta_p = 1 \times 10^{-5}$ is also chosen to be small. The resulting $w(E)$, $p(E)$, and their "exponentiated" values at different times are plotted in Figure 2.7.

The behavior of the plot at left of Figure 2.7, showing $w(E)$ and $p(E)$ may be roughly divided into three regions: for $0 < E < V_p$, the loss rate is dominated by $w(E)$; for $V_p < E < V_w$, the loss rate is dominated by $p(E)$; and for $E > V_w$, the loss rate is dominated by $\xi\sqrt{E} = vA/4\mathcal{V}$, and is relatively large. As may be seen at right of Figure 2.7, these same energy ranges delineate the different regions of suppression of $e^{-t(w+p)}$. The number of stored UCN in the bottle, $N(t)$ is then obtained from the integral in Equation 2.18, and is plotted as the black line at right in Figure 2.8. Its integrand, with the dependence on the free neutron

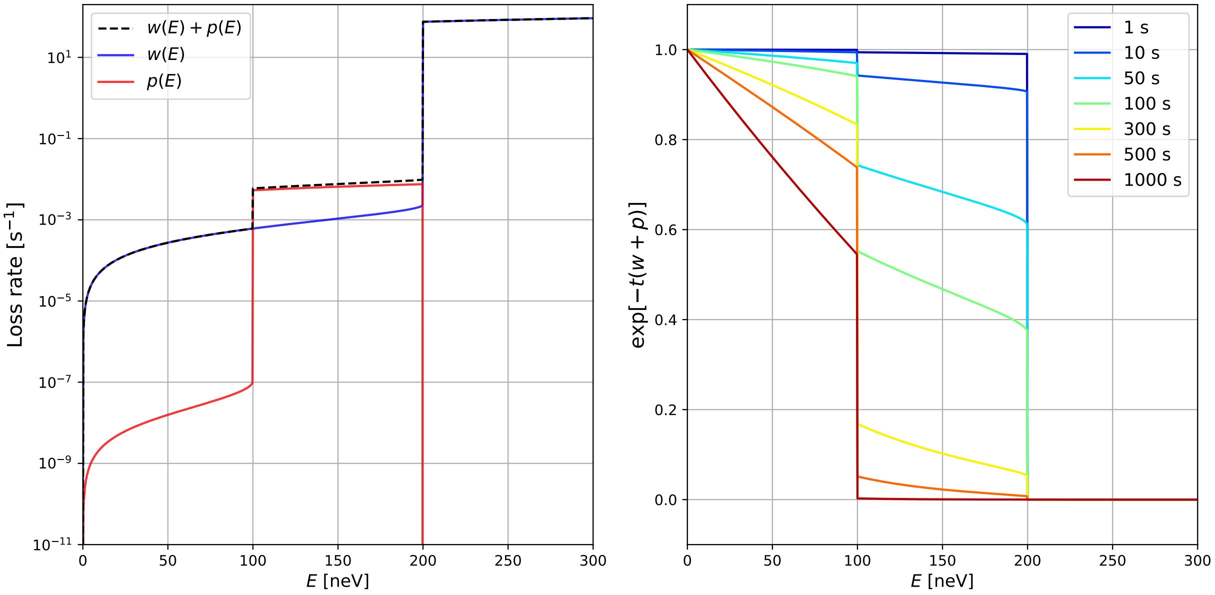

Figure 2.7: Left: Plots of $w(E)$ and $p(E)$ from Equation 2.23 and their sum. $V_w = 200\,\text{neV}$, $\eta_w = 1 \times 10^{-5}$, $V_p = 100\,\text{neV}$, $\eta_p = 1 \times 10^{-5}$, and $f = 1 \times 10^{-4}$. Right: plots of $e^{-t(w+p)}$ at different storage times t.

lifetime τ_β factored out, is shown at left in Figure 2.8.

In this example, the plot of $N(t)$ in Figure 2.8 appears to have roughly two time constants. This can be understood by breaking up the energy integral of Equation 2.18 into three parts:

$$\begin{aligned} N(t) &= e^{-t/\tau_\beta} \left(\int_0^{V_p} + \int_{V_p}^{V_w} + \int_{V_w}^{\infty} \right) \mathrm{d}E\, n(0,E) e^{-t[w(E)+p(E)]} \\ &\approx e^{-t/\tau_\beta} \left(\int_0^{V_p} \mathrm{d}E\, n(0,E) e^{-tw(E)} + \int_{V_p}^{V_w} \mathrm{d}E\, n(0,E) e^{-tp(E)} \right) \\ &\approx I_w(t) + I_p(t)\,, \end{aligned} \tag{2.24}$$

which defines $I_w(t)$ and $I_p(t)$. The contribution to $N(t)$ from the third integral is negligible for times greater than a few seconds, and $w(E)$ and $p(E)$ are approximated as being the dominant contributions to the loss in the energy regions $0 < E < V_p$ and $V_p < E < V_w$, respectively. For short storage times t, so that $t \ll \tau_\beta$

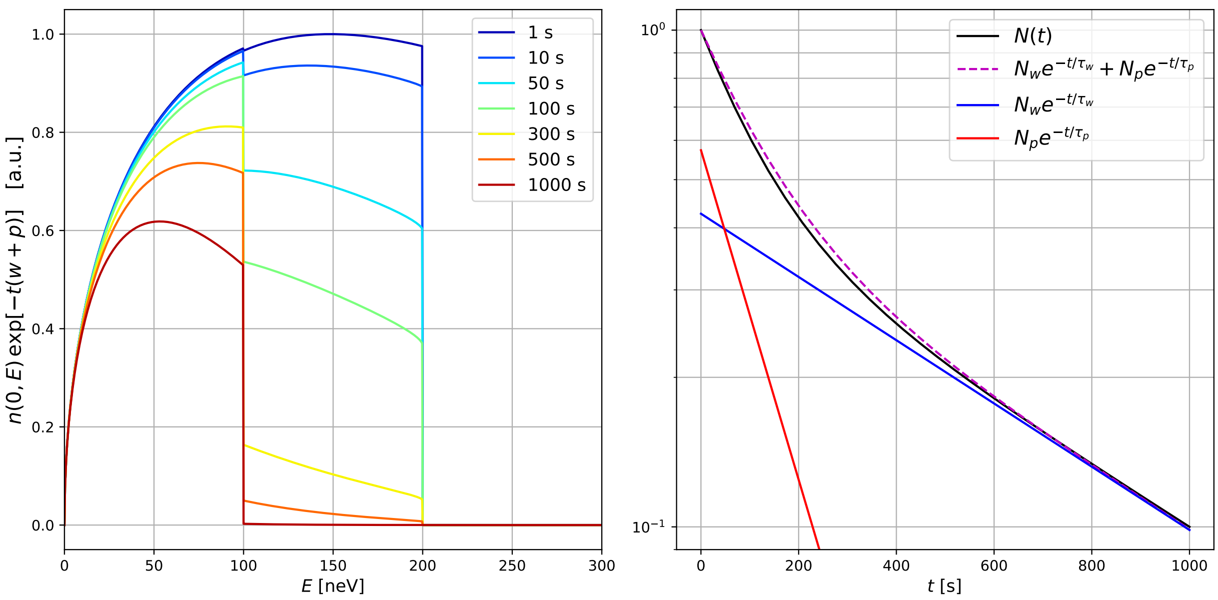

Figure 2.8: Left: plots of $n(0,E)e^{-t(w+p)}$ against energy E for different storage times t. $V_w = 200\,\mathrm{neV}$, $\eta_w = 1\times10^{-5}$, $V_p = 100\,\mathrm{neV}$, $\eta_p = 1\times10^{-5}$, and $f = 1\times10^{-4}$. Right: $N(t)$ of Equation 2.23 plotted against storage time t in black. The sum of the approximations to $I_w(t)$ and $I_p(t)$ slightly overestimates $N(t)$.

and $tw(E) \ll 1$ in the energy region $0 < E < V_p$, an approximation to the the first integral may made as

$$\begin{aligned}
I_w(t) &= e^{-t/\tau_\beta}\int_0^{V_p} \mathrm{d}E\, n(0,E)e^{-tw(E)} \\
&\approx e^{-t/\tau_\beta}\int_0^{V_p} \mathrm{d}E\, n(0,E)\left(1 - tw(E)\right) \\
&\approx \int_0^{V_p} \mathrm{d}E\, n(0,E)\left(1 - t\left[\frac{1}{\tau_\beta} + w(E)\right]\right) \\
&\approx N_w\left(1 - \frac{t}{\tau_w}\right) \\
&\approx N_w e^{-t/\tau_w}\,,
\end{aligned} \tag{2.25}$$

where the constants N_w and τ_w are given by

$$\begin{aligned}
N_w &= \int_0^{V_p} \mathrm{d}E\, n(0,E) \\
\frac{1}{\tau_w} &= \frac{1}{\tau_\beta} + \frac{1}{N_w}\int_0^{V_p} \mathrm{d}E\, n(0,E)w(E) \\
&= \frac{1}{\tau_\beta} + \langle w(E)\rangle_0^{V_p}\,.
\end{aligned} \tag{2.26}$$

N_w is thus the spectral weight between 0 and V_p, and τ_w is an effective time constant arising from the combination of beta decay and the average loss rate due to the "main" coating material $\langle w(E)\rangle_0^{V_p} = \xi \left\langle \sqrt{E}\bar{\mu}_w(E)\right\rangle_0^{V_p}$. In this example, $\tau_w \approx 680\,\mathrm{s}$, and N_w corresponds to roughly 43 % of the UCN population.

A similar operation may be carried out on the second integral, giving

$$\begin{aligned} I_p(t) &= e^{-t/\tau_\beta} \int_{V_p}^{V_w} \mathrm{d}E\, n(0,E) e^{-tp(E)} \\ &\approx N_p e^{-t/\tau_p}\,, \end{aligned} \tag{2.27}$$

where

$$\begin{aligned} N_p &= \int_{V_p}^{V_w} \mathrm{d}E\, n(0,E) \\ \frac{1}{\tau_p} &= \frac{1}{\tau_\beta} + \frac{1}{N_p}\int_{V_p}^{V_w} \mathrm{d}E\, n(0,E) p(E) \\ &= \frac{1}{\tau_\beta} + \langle p(E)\rangle_{V_p}^{V_w} \approx \frac{1}{\tau_\beta} + f\xi \left\langle \sqrt{E}\right\rangle_{V_p}^{V_w}\,. \end{aligned} \tag{2.28}$$

N_p is thus the spectral weight between V_p and V_w, and τ_p is an effective time constant arising from the combination of beta decay and the average loss rate due to the weak patches $\langle p(E)\rangle_{V_p}^{V_w} \approx f\xi \left\langle \sqrt{E}\right\rangle_{V_p}^{V_w}$. In this example, $\tau_p \approx 130\,\mathrm{s}$, and N_p corresponds to roughly 57 % of the UCN population. The plot at right in Figure 2.8 shows the approximations to $I_w(t)$ and $I_p(t)$. As can be seen, their sum lies slightly above the true $N(t)$, which implies that the storage time constants τ_w and τ_p are overestimates. In this example, this overestimation results mainly from both the absence of $w(E)$ in the approximation to $I_p(t)$.

When UCN storage data are not well described by a single time constant, the weak patch model discussed in this section can potentially offer insight into deviations of $N(t)$ from single, decaying exponential behavior. For example, it is typically the case that the loss factor η is smaller at low bottle temperatures, so that storage times increase with decreasing temperature. (There are a variety of possible physical explanations for this trend; for a theoretical review, see [82].) Consider then a hypothetical storage experiment whose data are described by a two-term decaying exponential—c.f. Equation 2.20, and in which the smaller time constant remains unchanged with varying bottle temperature. (This appears to be the case for the dPE storage experiments; see the next chapter for a presentation of the results.) Such behavior is consistent with the weak patch model, in which $I_p(t)$, having no dependence on η_w or η_p, should not change with temperature. A series of fits of experimental data to Equation 2.24 that produce a constant value of f would

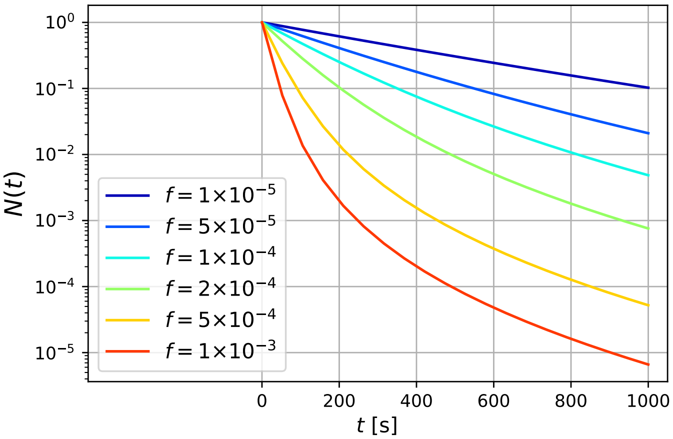

Figure 2.9: Plots of $N(t)$ with different patch area fractions f for $V_p = 0$, as would be the case for a completely absorbing material or a hole. $f = 1 \times 10^{-4}$ corresponds to a patch area of approximately $20\,\mathrm{mm}^2$ in the present example.

then seem to suggest the presence of some degree of contamination.

Unfortunately, the shape of $N(t)$—at least up to experimental resolution—is not uniquely determined by a single set of weak patch model parameters. It is therefore necessary to perform complementary characterization measurements of the bottle surface. As discussed in Section 2.3, neutron reflectometry may be used to determine the primary optical potential V_w. It then remains to determine V_p, η_w, and f. (The variable η_p is removed from the parameter list by the approximation of Equation 2.24.) For coating materials—like those studied in this work, and as opposed to those from which an entire bottle may be made, e.g. copper—there are two "natural" choices for V_p. The first choice is to set V_p equal to the optical potential of the material over which the coating has been placed, which would become visible to UCN as a result of holes or gaps in the coating. The second is to set $V_p = 0$, as would be true for gaps at the edges of any valves used for loading and emptying of UCN. This is an important limiting case of the weak patch model. However, as may be seen in Figure 2.9, a hole area fraction $f \sim 10^{-4}$ in this example is not capable of producing comparable levels of curvature as a similar patch area fraction does. (See Figure 2.8.) Depending on the experimental shape of $N(t)$, this may rule in favor of one choice of V_p. Finally then, with V_w and V_p fixed, and with the initial UCN energy spectrum $n(0, E)$ known—see Chapter 4—the optimal fitting values of η_w and f are uniquely determined for a given experimental data set.

Chapter 3

Storage measurements

In this chapter, the results of the UCN storage experiments are presented. The first section contains a description of the SUN-2 UCN source and the storage measurement apparatus. The following sections present neutron reflectometry data for CYTOP, and the UCN storage results for bottles coated in CYTOP and dPE. Storage measurements took place at the SUN-2 UCN source during two consecutive reactor cycles at the Institut Laue-Langevin: Cycle 188 and Cycle 189. Measurements during Cycle 188, which ran with a nominal reactor power of 56 MW, were performed between the 16th and 27th of September 2020. During this time, the SUN-2 source was operated in a “continuous” mode. (See the following section for an explanation of the source operating modes.) Measurements during the next cycle, Cycle 189, which ran with a lower nominal reactor power of 43 MW, were performed between the 1st and 27th of March 2021. During this cycle, the SUN-2 source was operated in both “continuous” and “accumulation” modes. (A smaller set of measurements were also performed in a bottle coated with “diamond-like carbon”, and are reported on in Appendix C.)

The CYTOP storage data appear to be well described by a single storage time constant that increases with decreasing temperature. (See Table 3.1 and Figure 3.10.) At a given temperature, its value remains roughly constant between different experimental configurations that are expected to influence the UCN energy spectrum. Its value was also observed to be insensitive to a systematic effect associated with prolonged measurements at low temperatures in the storage bottle. The dPE storage data are not well described by a single storage time constant. Fits to a double exponential function produce one time constant that increases with decreasing temperature, and one that remains approximately unchanged. (See Table 3.2 and Figure 3.15.) This behavior, in addition to other experimental observations, suggests that the weak patch model is an appropriate description of the dPE storage data. In addition, dPE storage data show a systematic effect that appears to be parameterized by the same quantity as that observed in the CYTOP measurements, implying that the effect is due to the storage apparatus, and is not inherent to either coating.

3.1 The setup

SUN-2 UCN source The UCN storage experiments described in this work were carried out at the SUN-2 prototype UCN source at the Institut Laue-Langevin (ILL) in Grenoble, France. The ILL is an international research facility for neutron science, whose nuclear reactor produces the most intense continuous neutron flux in the world: $1.5 \times 10^{15}\,\mathrm{s}^{-1}\,\mathrm{cm}^{-2}$ for a nominal reactor power of 57 MW [83]. The reactor consists of a highly-enriched ^{235}U fuel element moderated by heavy water, D_2O, which is maintained at a temperature slightly above 30 °C [84]. The heavy water also serves as the reflector material within the reactor pool, in which the entrances to a number of neutron guides are placed to deliver neutrons of various energies to more than forty experimental installations [85].

There are three guide sections leading from the heavy water reactor pool to the SUN-2 UCN source. The first of these is the "H1" guide, whose entrance sits in the immediate vicinity of an in-pile "cold source". The cold source consists of a 38 cm diameter sphere of liquid deuterium maintained at roughly 25 K [86]. Upon multiple scattering in the liquid deuterium, neutrons thermalized in the heavy water pool are further reduced in speed, and become "cold neutrons" (CN). These CN, having energy $E \sim 0.1$ to 10 meV [5], enter the H1 beam guide with a roughly Maxwellian energy distribution. The second guide section, "H172" is one of several that branch off of the larger H1 guide. At its end, a graphite crystal monochromator [87], which operates on the principle of Bragg reflection [88], may be rotated to direct a narrow spectrum beam of 8.9 Å wavelength CN [87] to one of two terminal beam guides: "H172a", which leads to the predecessor of SUN-2 [87, 89], and "H172b", which leads to SUN-2 itself [90, 91, 70].

The SUN-2 source produces UCN by converting a portion of the beam of 8.9 Å cold neutrons (CN) from the H172b guide into a stored collection of UCN via an inelastic "single-phonon" scattering process that is possible between neutrons and superfluid ^{4}He [92, 93, 94]. (The energy spectrum of the UCN so produced is considered in Chapter 4.) The source's "production volume" is essentially a storage bottle that traps the UCN converted from CN in this manner. It consists of a long, 4 liter aluminum box filled with a bath of superfluid ^{4}He—optical potential $V = 18.5\,\mathrm{neV}$, calculated from a density of $0.145\,\mathrm{g\,mL}^{-1}$ [95]—and coated with, first a layer of beryllium—$V = 252\,\mathrm{neV}$ [3]—and then an additional layer of Fomblin grease—$V = 106\,\mathrm{neV}$ [52]. To deliver UCN to the experiment, the source may be operated in either a "continuous" or "accumulation" mode. In the former, the CN beam is incident on the superfluid bath, and a continuous stream of UCN is allowed to leak out of the production volume via an uncovered 21 mm exit hole in its top. In the latter, a high density of UCN are first built up in the source by covering this hole with a UCN valve, and are then released in a burst to the experiment by opening the valve. This mode produces a softer—i.e. lower energy—spectrum compared to continuous mode operation. One practical drawback of operating the source in accumulation mode is the increased complexity to the user; the large number of cryogenic components needed to cool the source require frequent adjustment when the UCN valve is repeatedly opened and closed, placing a varying thermal load on the production volume.

Storage experiment The UCN storage experiment is connected to the SUN-2 source via a series of highly-polished stainless steel—$V = 189\,\mathrm{neV}$ [96]—UCN guides. The first guide is oriented vertically, and results in a height offset of approximately 356 mm between the horizontal source and bottle axes. Thereafter, all guides leading to the bottle are also oriented horizontally, and are of 50 mm inner diameter. A polished copper block sits approximately 250 mm from the axis of the vertical guide section. This block acts as a thermal screen to the source, allowing the production volume to be cooled to its base temperature of approximately 0.6 K. Approximately 270 mm thereafter, a polypropylene vacuum-separation foil—$V \approx 0\,\mathrm{neV}$, 4 µm thick [97]—reinforced by a thin, laser-cut steel mesh isolates the source and storage experiment vacuums during measurements. A vacuum gate valve sits roughly 30 mm downstream of this foil and is left permanently open during measurements, but can be closed to independently vent either the source or the storage experiment. The next component in the guide system is a UCN switch, which consists of a 120° stainless steel "elbow" positioned inside of a rotating drum whose vertical axis sits approximately 190 mm from the gate vale. This switch has three positions: in the "filling " position, the switch directs UCN from the source into the horizontal guides leading to the storage bottle; in the "bypass" position, the switch directs UCN from the source toward a UCN detector; and in the "emptying" position, the switch connects the bottle volume to the the detector volume, enabling stored UCN to be counted by falling into the detector.

The guides leading to the detector from the UCN switch are first a 90° elbow, and then a vertical guide section that widens to 80 mm inner diameter. (See (5) in Figure 3.1; the detector is normally shielded with several layers of boron loaded rubber.) The detector is a ^{3}He proportional counter—see, for example, Chapter 14, Section III B in [73]—belonging to the ILL. Its 50 mm wide, 80 mm diameter cylindrical detection volume was filled on September 10, 2018 with 37 mbar of ^{3}He gas, 15 mbar of CO_2 gas, and 1.1 bar of Ar gas. The entrance window of the detector, made of aluminum—$V = 54\,\mathrm{neV}$ [3]—and with a thickness of approximately 100 µm, lies 650 mm below the horizontal guide axes, far enough to ensure that all UCN incident on the detector have velocity sufficient to enter the detection volume. In this configuration, when the detector is "freshly filled", its efficiency is nearly 100 % [89]. (The absolute detection efficiency in these

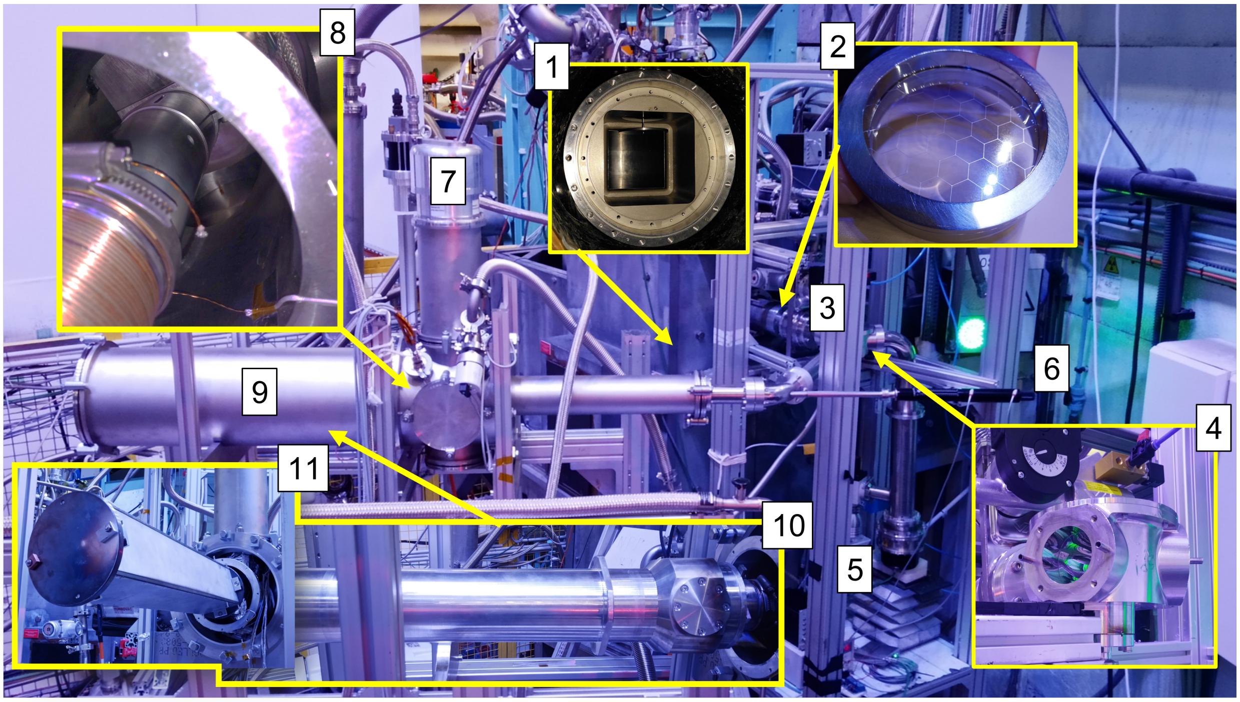

Figure 3.1: Picture of the storage experiment. (1) SUN-2 production volume with UCN valve open; (2) polypropylene vacuum-separation foil before the (3) gate valve between the source and the experiment; (4) rotating UCN switch; (5) UCN detector—normally shielded; (6) air-actuated bottle valve feedthrough; (7) pulse-tube refrigerator, or "cold head"; (8) thermal anchoring of cold head to copper sleeve; (9) vacuum vessel around modular storage bottles of either the (10) cylindrical-style or (11) box-style—not reported on in this work.

measurements, however, is not known as, over time, the gas mixture in the detector can leak out at an unknown rate.) The detector was operated at 1.4 kV via a high voltage supply passing through a model "MAD2016" amplifier, built and designed by the Instrument Control Group (SCI) at the ILL. Counts from storage measurements were recorded in the form of pulse-height spectra using the Amptek "MCA-8000D" digital multichannel analyzer (MCA) connected to the analog pulse signal from the ILL amplifier.

The first guide leading from the UCN switch to the storage bottle is approximately 500 mm in length. It is terminated by a 120° elbow that contains the linear motion vacuum feedthrough used to actuate the storage bottle valve. (See (6) in Figure 3.1.) An approximately 1000 mm long guide with a copper sleeve brazed to its exterior then runs directly to the storage bottle. This has a 3 mm hole in its side, through both the guide and sleeve, leading to the primary vacuum space of the storage experiment—see (9) in Figure 3.1—and enabling pumping of the guides and bottle. On the bottle end of this final guide, the sleeve has a large-bore, M60×1.5 male thread that mates to the aluminum front plate of the storage bottle assembly. This modular design allows bottles with different coatings and geometries to be easily installed. (See, for

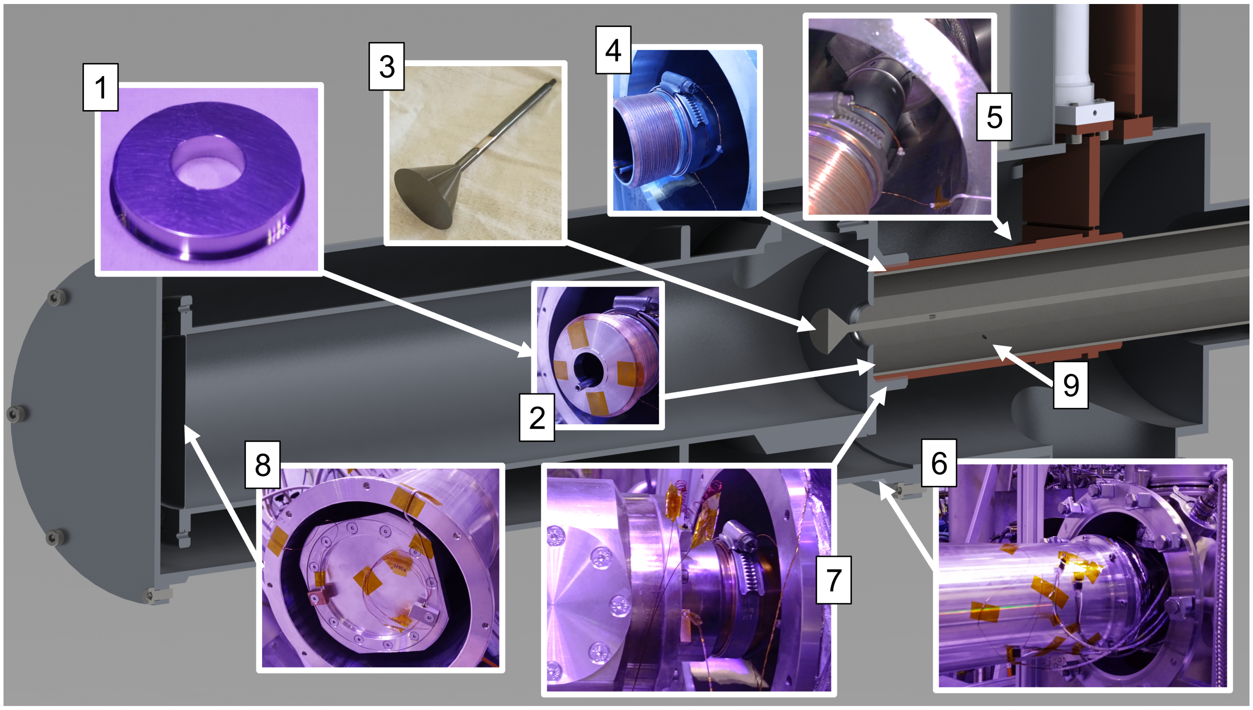

Figure 3.2: Cross section of storage bottle assembly installed in the system. The outermost shell is a heat screen thermally anchored to the first stage of a pulse-tube refrigerator. (1) polished stainless steel insert; (2) the insert in place; (3) example of exchangeable bottle plug; (4) threaded end of copper sleeve with hose-clamp around heater visible; (5) thermal anchoring to cold head; (6) routing of temperature sensors outside of thermal screen; (7) front temperature sensor placement and heater under hose-clamp; (8) back temperature sensor placement; (9) 3 mm diameter pumping hole through the guide and copper sleeve.

example, (11) in Figure 3.1.) Approximately 180 mm before the bottle, this sleeve is thermally anchored via a clamp to the second stage of a pulse-tube refrigerator, allowing the bottle to be cooled. (See Figures 3.1 and 3.2.) The first stage of the pulse-tube refrigerator is thermally anchored to a heat screen enclosing the storage bottle. This configuration enables an average base bottle temperature of roughly 10 K for the cylindrical-style bottles used in this work—to be described in the next paragraph—as registered by four sensors: one calibrated Cernox 1050 and one PT-100 on both the front and back of the bottle—with a temperature difference between ends at the lowest temperature of about 2 K. Pressures of around 2×10^{-6} mbar at room temperature and 7×10^{-7} mbar at the base temperature are typical, as measured in the primary vacuum space. (Immediately inside (9) in Figure 3.1.) A heater wire is clamped to the copper guide sleeve to achieve temperatures between room temperature and the base temperature, and is operated manually. (See (7) in Figure 3.2.)

As was mentioned in the previous paragraph and in Chapter 2, the aluminum storage bottle design used for CYTOP and dPE in these measurements is nearly cylindrical in shape, with a diameter of 90 mm and a

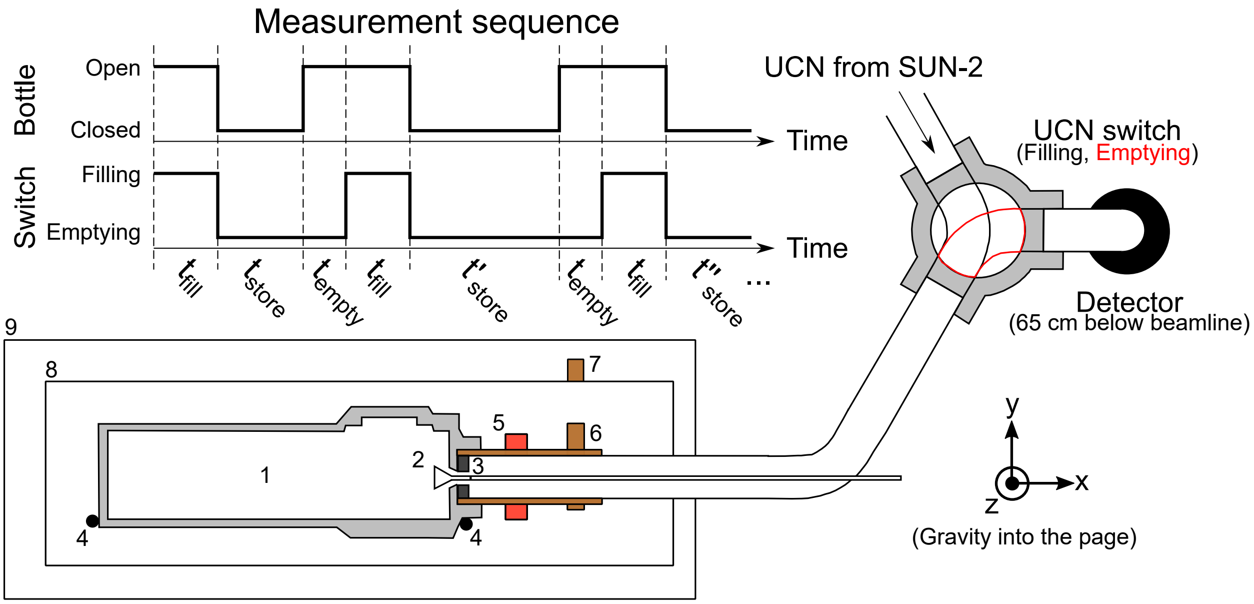

Figure 3.3: Overhead diagram of storage experiment, not to scale. Measurement timing sequence indicated at top with varying t_{store}. Gravity into the page. All elements are shown in the horizontal plane, except for the UCN detector, whose aluminum entrance window sits 650 mm below the axis of the bottle. (1) storage bottle; (2) coated steel bottle plug; (3) removable steel entrance cover; (4) temperature sensors—only two are shown, there are four in total; (5) heater; (6) thermal anchoring to first stage of pulse-tube refrigerator (not shown); (7) thermal anchoring to second stage; (8) thermal screen; (9) vacuum vessel.

length of 589 mm. (See bottle at bottom of Figure 4.3 for a cross section view.) The bottle is mounted with its axis horizontal. (See Figure 3.1.) The UCN entrance and exit to the bottle consists of a 20 mm diameter hole in the front plate of the bottle assembly, which can be opened or closed by pushing or pulling a stainless steel plug—coated with the same material as the interior of the bottle—into or away from the bottle. (See (3) in Figure 3.2.) As with the bottle, this plug can be exchanged, and is installed by screwing it into one end of a 4 mm diameter steel rod that runs down the center of the guide leading to the bottle. The other end of this rod passes through the linear motion vacuum feedthrough discussed in the last paragraph. The exterior of the front plate, which is seen by UCN approaching the bottle as an annulus, is uncoated; some measurements were performed with a stainless steel insert covering this exposed aluminum. (See (1) and (2) in Figure 3.2.) Additional details of the bottle will be discussed in the Sections 3.2 and 3.3.

Storage measurement procedure A simplified UCN storage measurement consists of three steps: filling, storage, and emptying. (See timing sequence in Figure 3.3). In the first step, the UCN switch is rotated to the filling position with the bottle valve open, allowing UCN from the source to enter the storage bottle for a time t_{fill}. In the second step, the bottle is closed, and the UCN switch is rotated to the emptying position, allowing any UCN marginally stored in the guides between the bottle and the switch to fall into the detector. (These guides were observed to be completely emptied in roughly 30 s.) The last step is to open the bottle after the storage time t_{store} has elapsed, and empty UCN into the detector for a time t_{empty}. The counts are recorded by the MCA, which is triggered to take data 1 s before the bottle opens. A new measurement begins by returning the UCN switch to the "filling" position, after which the three steps are repeated using different values of t_{store}.

The appropriate values for t_{fill} and t_{empty} are different for continuous and accumulation mode operation of the source. In the former, the number of UCN detected, N, for a fixed storage time is roughly $\propto 1 - \exp(-t_{\mathrm{fill}}/\tau_{\mathrm{fill}})$, where τ_{fill} is a time constant associated with filling. A value of $t_{\mathrm{fill}} = 200\,\mathrm{s}$ was chosen in this mode as a compromise between large N and efficient use of measurement "beamtime" at the SUN-2 source. In accumulation mode, before filling begins, UCN are built up in the production volume. This accumulation time was chosen to be 600 s based on similar considerations of large N and measurement efficiency. The filling time was then roughly optimized by finding the value of t_{fill} that maximizes N. This was found to be $t_{\mathrm{fill}} = 50\,\mathrm{s}$. (See Figure 3.4.) In both modes of operation of the source, the emptying time was chosen to be long enough for the detector count rate to return to the background level. (See later in this section for a description of background measurements.) In continuous mode, this was chosen to be $t_{\mathrm{empty}} = 120\,\mathrm{s}$, while in accumulation mode, this was chosen to be $t_{\mathrm{empty}} = 140\,\mathrm{s}$, on account of the number of UCN in the bottle being slightly higher in this mode of operation.

The timed coordination of the UCN switch, the bottle valve, and the MCA trigger are performed automatically using custom LabVIEW programs. One "main" program was written for repeated storage measurements made while operating the source in continuous mode. In this mode, the position of the CN beam, which is incident on the SUN-2 production volume, is unchanged. This is not the case in accumulation mode measurements, in which the CN beam is moved away when the UCN valve in the source is opened. Control of the UCN valve motion and the CN beam position is only possible through the ILL's "NOMAD" instrument control system. A separate main LabVIEW program was therefore written for accumulation mode measurements, in which timing control was slaved to the NOMAD software. An independent LabVIEW program was also written to continuously monitor the temperature and pressure in the experiment.

Immediately prior to and following a run of storage measurements at a given temperature, "source flux" and "background" rate measurements were performed. (See Figures 3.5 and 3.6.) The former were made by measuring with the source UCN valve open, the UCN switch in the bypass position, and the CN beam

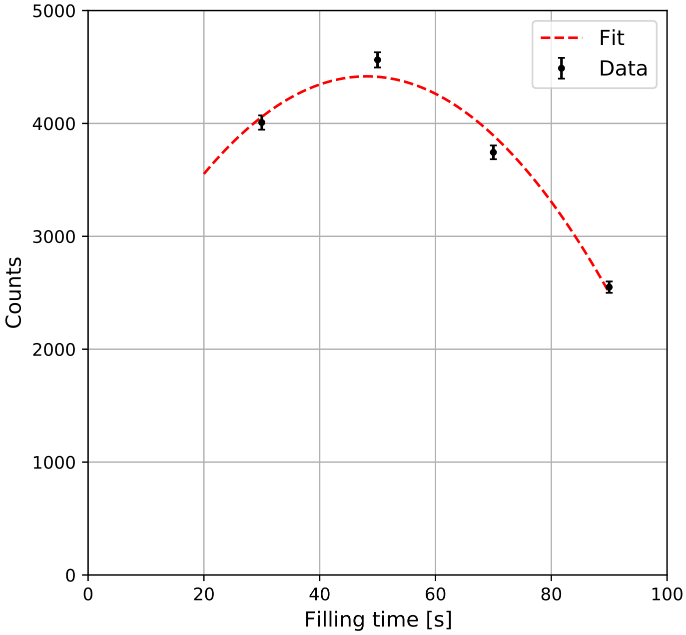

Figure 3.4: Data taken for a rough optimization of the filling time, t_{fill}, while operating the SUN-2 source in accumulation mode. UCN were counted after storage for 100 s in the dPE bottle. The red dashed line shows a second-order polynomial fit to the data.

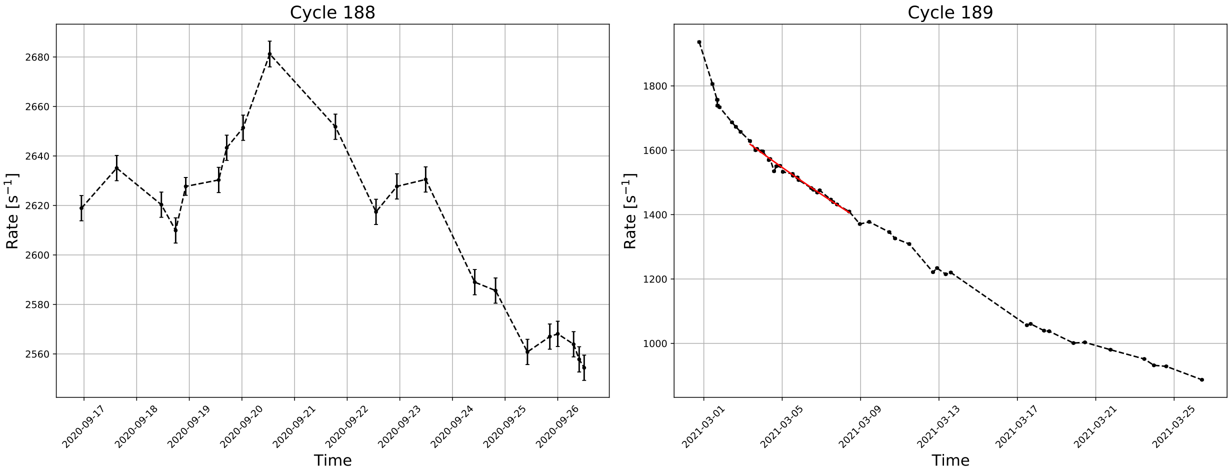

Figure 3.5: SUN-2 “source flux” rate measurements during Cycle 188 and Cycle 189. The rate is measured with the source UCN valve open, the UCN switch in the bypass position, and the CN beam incident on the production volume. See text for an explanation of the different trends. The red line in the plot at right is an exponential fit to the source flux data between the 3rd and 8th of March, corresponding to a dPE measurement series. (See Section 3.3.) In this region, the time constant is 863(151) h, corresponding to a decrease of about 2.7(4) % per day.

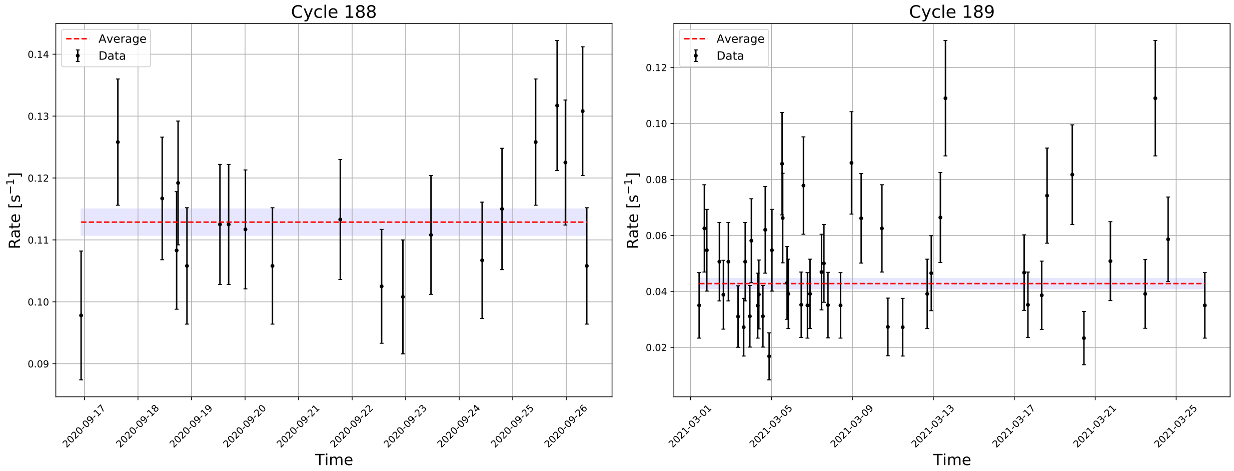

Figure 3.6: Background rates during Cycle 188 and Cycle 189, as measured for continuous mode operation of the SUN-2 source. The rate is measured with the source UCN valve open, the UCN switch in the emptying position, and the CN beam incident on the production volume. The red dashed line indicates the weighted average rate, and the blue band shows the 1-σ uncertainty.

incident on the production volume. During Cycle 188, an average rate of approximately $2600\,\mathrm{s}^{-1}$ was observed, varying by less than 3 % over the 11 day period of measurements. During cycle 189, a leak in the vacuum-separation foil between the source and the storage experiment resulted in a gradual decrease in the source flux rate from roughly $1900\,\mathrm{s}^{-1}$ to $900\,\mathrm{s}^{-1}$ over a 26 day period. It is suspected that this decrease was caused by residual gases, like water vapor, slowly leaking into the source and freezing on the walls of the extraction guides, a volume normally filled only with helium gas. When additional pumps were placed on a guide section near the source, the improved vacuum resulted in a change in the rate of decrease of the flux from about 6 % per day to about 3 % per day. (See the kink in this plot roughly around the 2nd of March.) Strangely, despite the gradual decrease in the source flux rate, no obvious corresponding decrease in the maximum number of stored UCN in the bottles was observed. (See Section 3.3.)

The background rate for continuous mode storage measurements was determined by measuring with the source UCN valve open, the UCN switch in the emptying position, and the CN beam incident on the production volume. (See Figure 3.6.) During Cycle 188, a rate of $0.113(2)\,\mathrm{s}^{-1}$ was observed. (About 13 counts average in 120 s.) During Cycle 189, this rate was measured to be $0.043(2)\,\mathrm{s}^{-1}$. (About 5 counts average in 120 s) The decrease is mainly the result of a smaller source flux; the UCN switch is not completely “tight” against UCN leakage. This fact is also evident in measurements of the background rate appropriate for accumulation mode. These were made with the source UCN valve open, the UCN switch in the emptying position, and the CN beam moved away from the production volume. A rate of $0.0156(7)\,\mathrm{s}^{-1}$ was observed.

(About 2 counts average in 140 s.)

3.2 CYTOP

CYTOP is an amorphous “fluoropolymer”—repeating units of a $C_6F_{10}O$ monomer—developed by AGC Chemicals, and is used in various industrial coating applications [54]. Its constituent elements have small absorption cross sections, which, in conjunction with its hydrophobic chemical properties [54], makes CYTOP a promising candidate for producing low-loss, hydrogen-free coating surfaces for UCN storage. Using tabulated neutron scattering lengths [1] and absorption cross sections [2], along with a manufacturer-specified density of $2.03\,\mathrm{g\,cm^{-3}}$, Equation 2.12 predicts a neutron optical potential of $V = 117\,\mathrm{neV}$ and a loss factor of $\eta \sim 3 \times 10^{-7}$ for CYTOP. This optical potential is slightly higher than that of Fomblin, 106 neV [52], a fluorinated compound commonly used for UCN storage coatings [98, 99]. The measured loss factor of Fomblin is roughly $\eta \sim 2 \times 10^{-5}$ [52], whereas its predicted value is also $\eta \sim 3 \times 10^{-7}$ [100]. This discrepancy between predicted and measured values of η is ubiquitous in UCN storage experiments, and it is unlikely that such a low estimate of the loss factor is correct for CYTOP. (See Chapter 2.) Coatings reported here were prepared mainly by Jürgen Hingerl [101] and Euan Connolly [102] using the liquid solution CTL-107MK, which consists of a ratio of 7 g of type-M CYTOP dissolved in 300 mL of CT-SOLV100K fluorosolvent. The manufacturer recommends curing of the product—intended for surface application by dip-coating—by first drying it in air to evaporate the solvent, and then baking it to harden and achieve cross-linking of the polymer chains.

Bottle preparation The UCN storage bottle coating was prepared using a custom oven originally developed to produce deuterated polyethylene coatings [103, 60, 104, 101]. The main body of the bottle assembly is an approximately cylindrical tube of aluminum: inner diameter 90 mm, length 589 mm. (See lower bottle in Figure 4.3 for a cross section view.) A side port and the back end of the main body are capped by aluminum blanks, while an aluminum front plate with a 20 mm hole in its center allows UCN to enter and exit the bottle. To coat, the fully-assembled bottle was filled with CYTOP solution, and was installed in the coating oven with its axis oriented horizontally. Coating of the entrance plate was achieved by inclining the entire oven at an angle chosen to avoid spillage of the solution from the bottle entrance hole. The bottle was rotated about its axis at 6 rpm for 15 min with the oven temperature set to 100 °C. The oven was then inclined in the opposite sense and rotated for an additional 15 min to coat the back plate. The temperature was then increased to 150 °C, and the oven was tilted back and forth in intervals of 10 min until no solvent remained. This was monitored by means of a solvent recovery system. Finally, the bottle was baked at 200 °C for 30 min to complete hardening of the coating. It is estimated that the average coating thickness is approximately 15 µm.

In the UCN storage experiment, the filleted bottle entrance hole is opened and closed by a conical stainless-steel plug. This piece was prepared in a separate dip-coating process. For a pull speed of approximately $1\,\mathrm{mm\,s^{-1}}$, the manufacturer gives a nominal coating thickness of 60 nm. The plug was coated multiple times by pulling it at this speed from an identical CYTOP solution. A mechanical device was used to pull and slowly rotate the plug to avoid forming large drops. The plug was dried between 1 and 5 min in an oven at 80 °C after each dip. The plug was finally placed in an oven at 200 °C and baked for 30 min to harden the coating. The coating thickness is estimated to be no less than 1.5 µm.

Reflectometry Two thin films of CYTOP on silicon substrate were prepared for neutron reflectometry measurements using a dip-coater: samples 1 and 2 were prepared using pull speeds of $0.5\,\mathrm{mm\,s^{-1}}$ and $1\,\mathrm{mm\,s^{-1}}$, respectively. Both samples were dried in air for 15 min, and baked at 130 °C for 45 min to harden. The neutron optical potential was determined via time-of-flight reflectometry with cold neutrons at the ILL instrument D17 [105] [106]. Raw data refinement, consisting of background subtraction, normalization, and correction for instrumental effects, was performed using the ILL's LAMP/COSMOS software [107] to obtain reflectivity as a function of momentum transfer Q, as shown in Figure 3.7. ($Q = 2k_\perp$, see Chapter 2.) The refined data were fit to a "target function" that is the convolution of an ideal reflectivity curve and a resolution function. (For a description of the method of calculation of the ideal reflectivity curve for a layered system, see Appendix Section A.2.) The resolution function is a set of distributions that describe the range of physically-different Q values that contribute to the reflectivity at each measured value of the momentum transfer on a real instrument. When the width of a given distribution at some nominal value of Q is small compared to the features of the reflectivity curve, it may be approximated by a Gaussian having the same variance [108]. The collection of these widths are returned by the refinement software—as are appropriate for a given configuration of the D17 instrument—in terms of the full-width half maximum (FWHM) of the distributions: these are displayed as the horizontal error bars in Figure 3.7. Several models of the ideal reflectivity were tested, incorporating different combinations of possible parameters: film thickness and optical potential, interfacial roughness, and the thickness and optical potential of a "native oxide" layer on the silicon substrate. (Appendix Section A.2 also provides a definition of the surface roughness.)

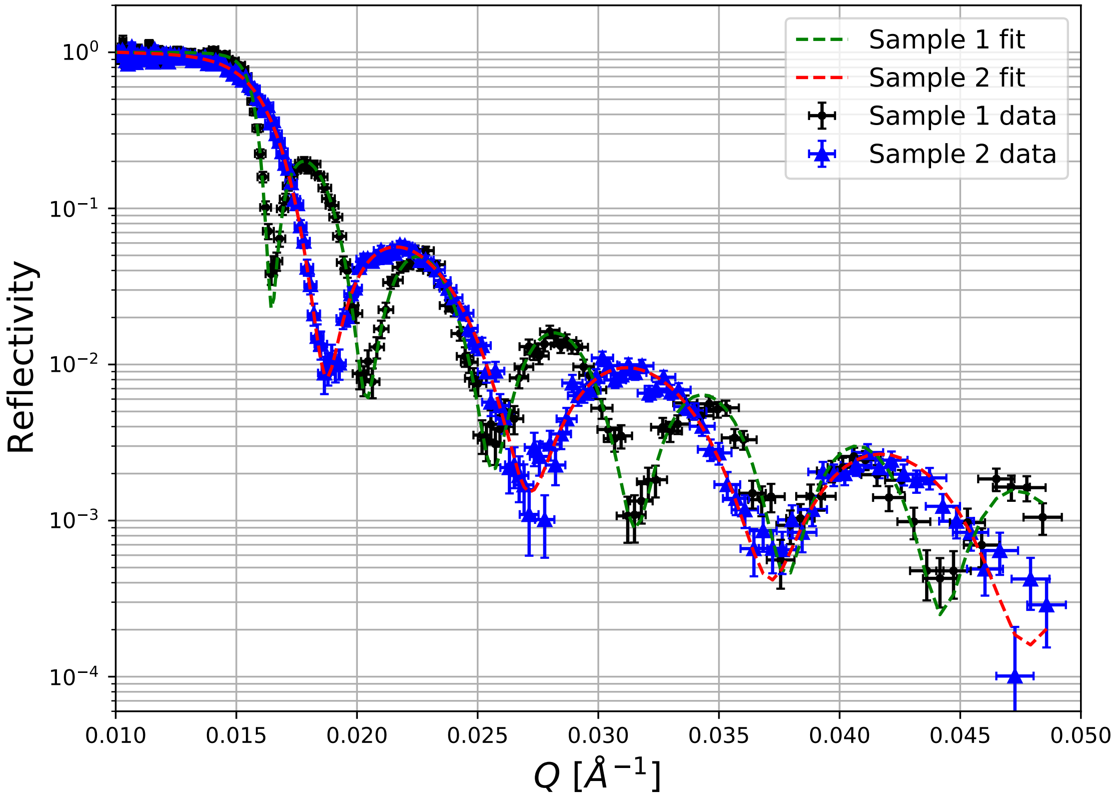

Figure 3.7: Neutron reflectivity of two CYTOP thin film samples on silicon substrate as a function of momentum transfer Q. Horizontal error bars are the FWHM of the "resolution function" of the D17 instrument. (See the text.) Example fits in this figure are for a minimal model, containing only the film thickness, t, and (real) optical potential, V, as parameters. Sample 1: $\chi^2/(185-2) = 1.58$, $V = 115.3(1)\,\text{neV}$, $t = 905.8(4)\,\text{Å}$. Sample 2: $\chi^2/(184-2) = 1.03$, $V = 114.9(1)\,\text{neV}$, $t = 554.9(4)\,\text{Å}$.

For each model, optimal fit parameters were determined by numerically minimizing χ^2, that is, the sum of squares of the weighted residuals formed by the difference between the refined data and the target function. This was performed with the `curve_fit` function in SciPy [109], using the Nelder-Mead algorithm [110]. The parameter uncertainties were determined using numerical estimates of the Hessian, H, of χ^2, evaluated at the optimal fit values. (The inverse of the covariance matrix $(\Sigma^{-1})_{ij}$ is estimated by $\frac{1}{2}H = \frac{1}{2}\frac{\partial^2\chi^2}{\partial\theta_i\partial\theta_j}$; see, for example, Equation 40.12 in [6].) The weighted averages of the optical potentials returned from fits to the different models were calculated for each sample separately, and were found to be 115.41(4) neV and 115.04(5) neV for samples 1 and 2, respectively. To report a single value for the optical potential of CYTOP with an uncertainty that reflects the disagreement between these two results, we take a weighted average in the following way. First of all, the procedure of taking the weighted averages of a group of fit results—that is, the different optical potentials returned by the different models—is equivalent to fitting a constant to the optical potentials: the reduced chi-squared values associated with these "fit of the fits" are rather large: 4.70 and 2.31, respectively. The statistical uncertainties of the optical potentials going into the weighted averages are therefore enlarged by rescaling by these respective values of $\sqrt{\chi_r^2}$. Taking the weighted average a second time then produces $\chi_r^2 = 1$, without changing the optimal values of the weighted average optical potentials. (Such an approach to expanding uncertainties is taken, for example, in the meta-analysis of the free neutron lifetime, see [6].) The weighted average optical potential of samples 1 and 2 likewise results in a large value of χ_r^2. This rescaling operation is then performed again, producing a final estimate of 115.2(2) neV for the optical potential of CYTOP.

Storage measurements For CYTOP, after background subtraction, the number of UCN counted by the detector after being emptied from the bottle are well described by a single-term exponential model

$$N_0 \exp(-t_s/\tau_0)\,, \tag{3.1}$$

where t_s is the storage time, and τ_0 is the storage time constant. (Compare to Equation 2.19.) See Figure 3.8 for an example. At low bottle temperatures, a systematic decrease in counts over time was observed in long measurement series. For such measurements, an additional term was added to the storage model:

$$\exp(-t_m/\tau_d)\, N_0 \exp(-t_s/\tau_0)\,, \tag{3.2}$$

where t_m is the absolute measurement time, taken to be zero at the beginning of a run of storage measurements, and τ_d is the time constant of the systematic decrease in the counts. (See Figure 3.9 for an example.)

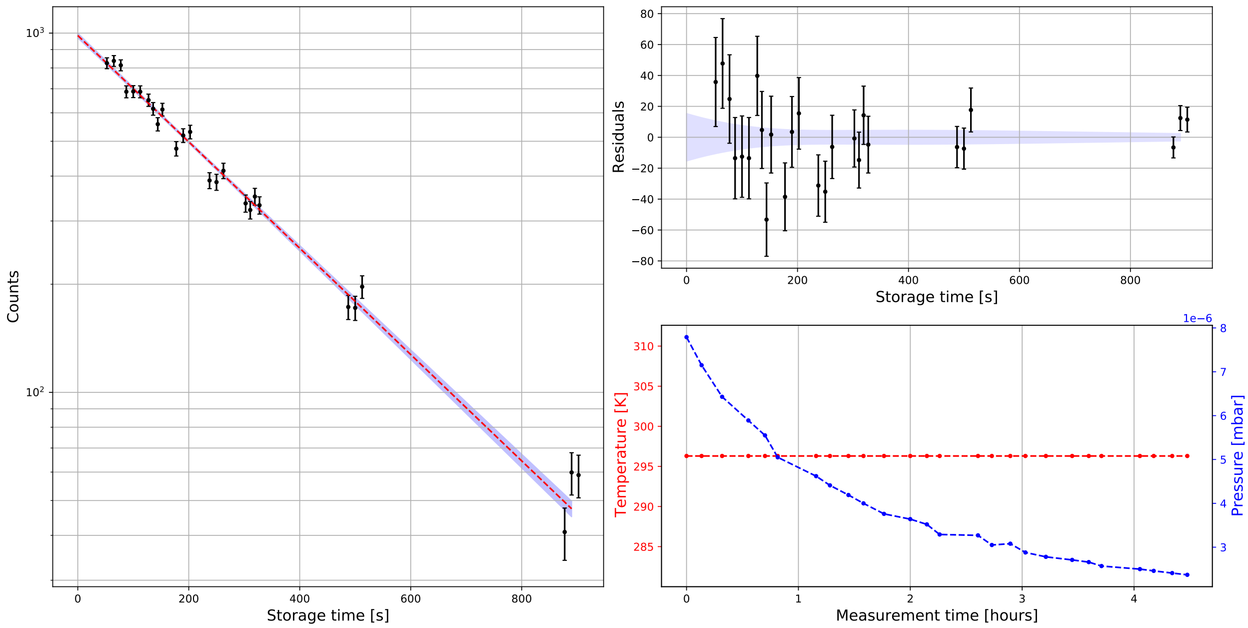

Figure 3.8: Example CYTOP storage data set taken in continuous mode during Cycle 189 at 296.3 K average and with the aluminum entrance to the bottle. $N_0 = 985(15)$, $\tau_0 = 293(6)$ s, $\chi^2/(26-2) = 1.30$. Left: Storage data are shown in black, and are displaced horizontally to show vertical spread. The fit is shown in red, with 1-σ error bands in light blue. (See Equation 3.1.) Top right: residuals from the plot at left. Bottom right: Average temperature and pressure during the measurement.

The motivation for the form of this empirical, decreasing term is that the least squares estimates of the fit parameters may be solved for exactly by linearizing Equation 3.2 using the natural logarithm. (See, for example, [111].) The same is true for the simpler expression in Equation 3.1. Statistical uncertainties are obtained by standard propagation of errors. The fit parameters are reported in Table 3.1, along with other information of interest. See Figure 3.10 for plots of N_0 and τ_0 against the average bottle temperature.

To determine whether the additional term in t_m was appropriate to a given measurement, a statistical test was used. The following quantity,

$$F = \frac{\chi^2_{\nu-2} - \chi^2_{\nu-3}}{\chi^2_{\nu-3}/(\nu-3)}\,, \tag{3.3}$$

was computed, where $\chi^2_{\nu-2}$ and $\chi^2_{\nu-3}$ are the chi-squared values from the fits of Equations 3.1 and 3.2, respectively. ν is the number of storage measurements in the data set, so that, for example, for the two-parameter fit to Equation 3.1, $\nu - 2$ is the number of degrees of freedom. Under the assumption that $\chi^2_{\nu-2}$ and $\chi^2_{\nu-3}$ are chi-squared random variables with $\nu-2$ and $\nu-3$ degrees of freedom, respectively, F follows the F-distribution with 1 “numerator” and $\nu-3$ “denominator” degrees of freedom. (See, for example, Chapter 11 in [112].) If the gain in goodness-of-fit from including the term in t_m is appreciable, F will be larger than

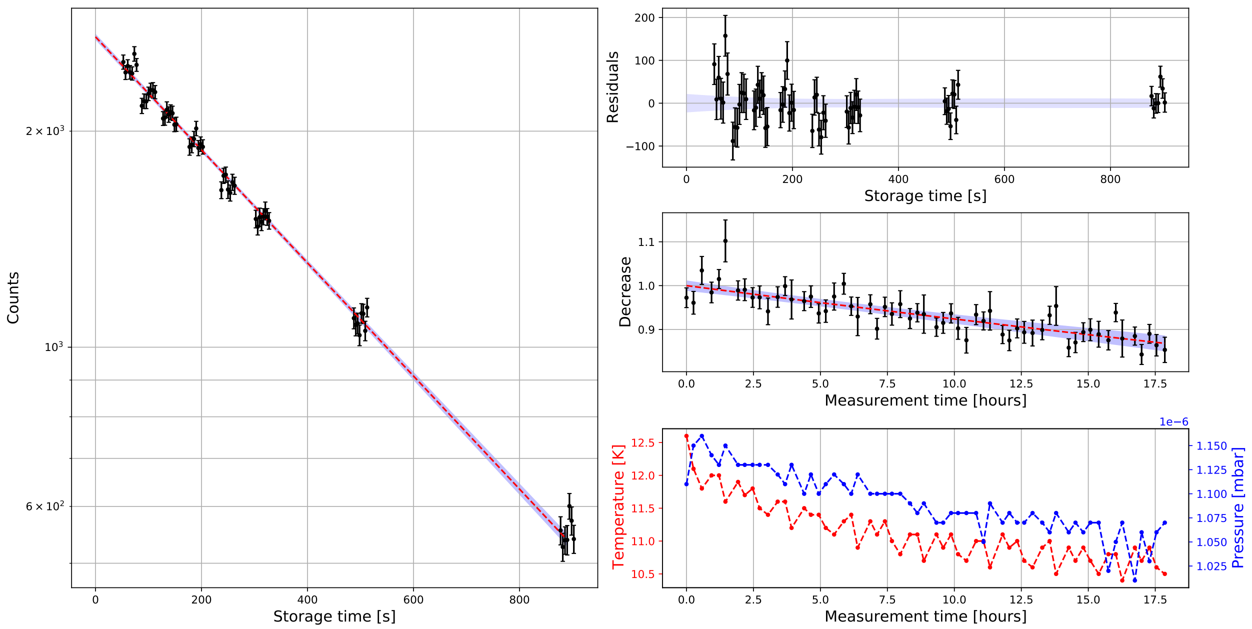

Figure 3.9: Example CYTOP storage data set taken in accumulation mode during Cycle 189 at 11.1 K average and with the aluminum entrance to the bottle. $N_0 = 2707(20)$, $\tau_0 = 551(5)\,\mathrm{s}$, $\tau_d = 126(10)\,\mathrm{h}$, $\chi^2/(59-3) = 1.08$. Left: Storage curve with the systematic decrease corrected for. (Equation 3.2 divided by the factor $\exp(-t_m/\tau_d)$.) Data are shown in black, and are displaced horizontally to show vertical spread. The fit is shown in red, with 1-σ error bands in light blue. Top right: residuals from the plot at left. Middle right: the systematic decrease. (Equation 3.2 divided by the factor $N_0\exp(-t_s/\tau_0)$.) Bottom right: Average temperature and pressure during the measurement.

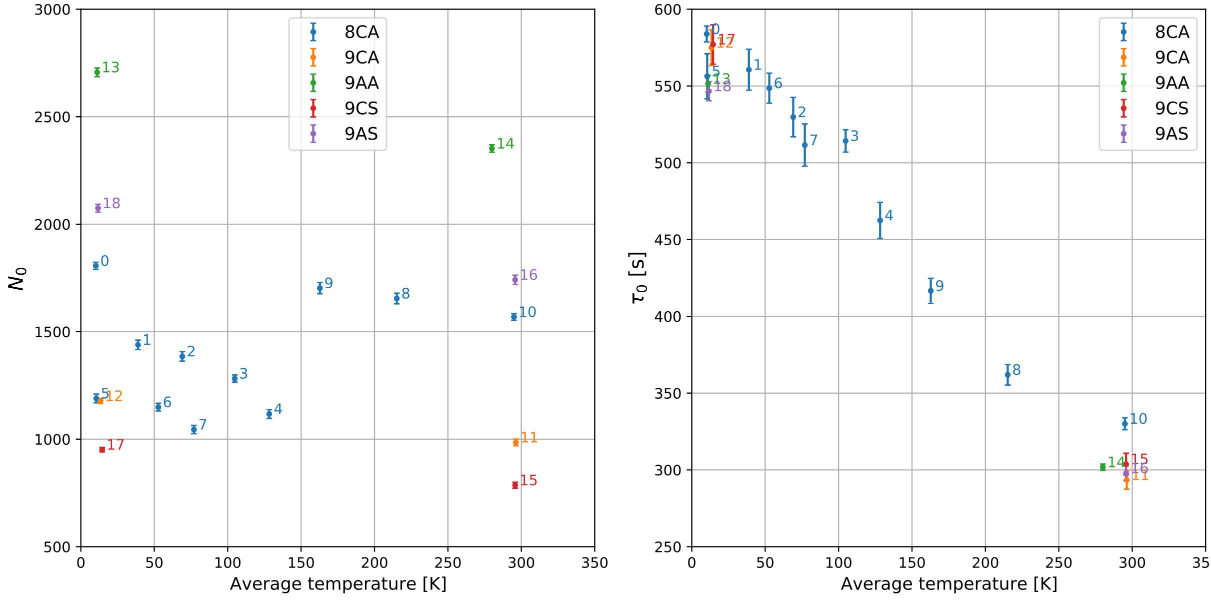

Figure 3.10: Fit parameters N_0 and τ_0 for CYTOP data. Points are labeled by a chronological index. (See Table 3.1.) Data are grouped by cycle, source mode, and bottle entrance. For example, “8CA” corresponds to Cycle 188, continuous mode, and aluminum bottle entrance, and “9AS” corresponds to Cycle 189, accumulation mode, and stainless steel bottle entrance. Left: N_0 plotted against average bottle temperature. Right: The storage time constant τ_0 plotted against the average bottle temperature.

Table 3.1: Summary table of CYTOP storage measurements. The first column assigns a chronological index to the measurement. The "Date" column lists the start date and time of a measurement. "Front" indicates the entrance to the storage bottle: either the uncovered aluminum, "Alu.", or with the polished stainless steel insert, "SS.", in place. "Size" is the number of points in the data set. "Mode" is either continuous, "C", or accumulation, "A". "T" is the average bottle temperature during the measurement. (Various problems with the temperature sensors in Cycle 188 necessitated their replacement after the cycle. The average bottle temperatures reported for these data sets were reconstructed using the same heater settings, but with new temperature sensors.) The parameter estimates for τ_0, N_0, and τ_d—when appropriate, see text—are least squares estimates. χ_r^2 gives the value of the reduced chi-squared statistic.

Index	Cycle	Date	Front	Size	Mode	T [K]	τ_0 [s]	N_0	τ_d [h]	χ_r^2
0	188	20/09/17 15h	Alu.	65	C	10.2	583(5)	1806(16)	99(7)	1.23
1	188	20/09/18 13h	Alu.	15	C	38.9	560(13)	1438(21)	–	1.21
2	188	20/09/18 18h	Alu.	15	C	69.1	529(12)	1385(22)	–	0.71
3	188	20/09/19 00h	Alu.	49	C	104.8	514(7)	1281(16)	104(20)	1.05
4	188	20/09/19 13h	Alu.	15	C	128.3	462(11)	1117(20)	–	0.92
5	188	20/09/19 20h	Alu.	15	C	10.5	556(14)	1189(20)	–	1.11
6	188	20/09/20 00h	Alu.	35	C	52.8	548(9)	1149(17)	107(38)	0.81
7	188	20/09/20 09h	Alu.	15	C	77.0	511(13)	1044(19)	–	1.42
8	188	20/09/21 19h	Alu.	16	C	215.2	361(6)	1654(24)	–	1.93
9	188	20/09/22 13h	Alu.	15	C	162.8	416(8)	1702(26)	–	1.57
10	188	20/09/22 23h	Alu.	38	C	295.0	330(3)	1568(15)	–	2.06
11	189	21/03/12 16h	Alu.	26	C	296.3	293(5)	984(15)	–	1.30
12	189	21/03/13 08h	Alu.	33	C	13.4	574(11)	1177(11)	–	1.12
13	189	21/03/13 14h	Alu.	59	A	11.1	551(5)	2706(20)	126(10)	1.08
14	189	21/03/14 17h	Alu.	100	A	280.0	301(2)	2352(17)	315(34)	1.97
15	189	21/03/17 11h	SS.	24	C	295.8	303(7)	785(13)	–	1.30
16	189	21/03/17 16h	SS.	22	A	295.8	297(4)	1741(21)	–	1.22
17	189	21/03/18 08h	SS.	34	C	14.5	577(12)	951(10)	–	0.97
18	189	21/03/18 15h	SS.	57	A	11.7	546(6)	2074(18)	149(16)	1.72

some critical value—chosen here to correspond to a 99 % confidence level—and the use of Equation 3.2 in fitting the storage data is valid. For example, for measurement 13—see Table 3.1—the value of the above test statistic is approximately 148. The probability that F, for an F-distribution with 1 numerator degrees of freedom and $(59-3) = 56$ denominator degrees of freedom, takes a value of 148 or larger is approximately 2×10^{-15} %. Such a low probability indicates that the gain in goodness of fit from the addition of the τ_d explanatory variable is significant. Interestingly, it was observed that only the chi-squared value of the fit changes appreciably with inclusion of the term in t_m; the estimates of τ_0 and N_0 remain largely uninfluenced by the presence or absence of the extra factor.

It is suspected that the physical reason for the systematic decrease in counts with t_m at low temperature is a "cyropumping" of residual gas by the cold walls of the experiment, similar in nature to that believed to have caused the gradual degradation in the source during Cycle 189. (See the discussion in the previous section.) Several observations support this hypothesis. For one, the pressure in the experiment drops from roughly 2×10^{-6} mbar to 7×10^{-7} mbar when the system is cold, i.e. it appears to be cryopumping. Secondly, when the cold system is returned to room temperature, the decrease in counts ceases. On recooling, it returns. This patterns persists with thermal cycling of the system. This behavior also indicates that CYTOP is a cyrogenically durable coating, and is resistant against cracking or other aging effects from repeated warming and cooling. An additional observation in favor of the cryopumping scenario is that similar values of τ_d were observed in tests involving different bottles and plugs. (See next section.) This suggests a mechanism for the systematic decrease that is related to the storage measurement apparatus, and is not inherent to CYTOP as a coating.

With respect to τ_d, there is one clear outlier in Table 3.1: an accumulation-mode measurement, near room temperature at 280 K, gives a value of 315(34) h for τ_d. This is the largest data set, and represents the longest uninterrupted storage measurement made. Fits to these data should therefore be more sensitive to systematic decreases with longer time constants. However, the mechanism for this effect is not understood. (See the next section, in which two sets of storage data show similar outliers.)

In the plot at right in Figure 3.10, the storage time constants τ_0 appear to cluster together, independently of the bottle entrance type, the mode of operation of the SUN-2 source, or the source flux rate. The gradual increasing trend of τ_0 with decreasing temperature is most clear in the data taken during Cycle 188, in which all measurements were performed with a bare aluminum bottle entrance, and while the source flux

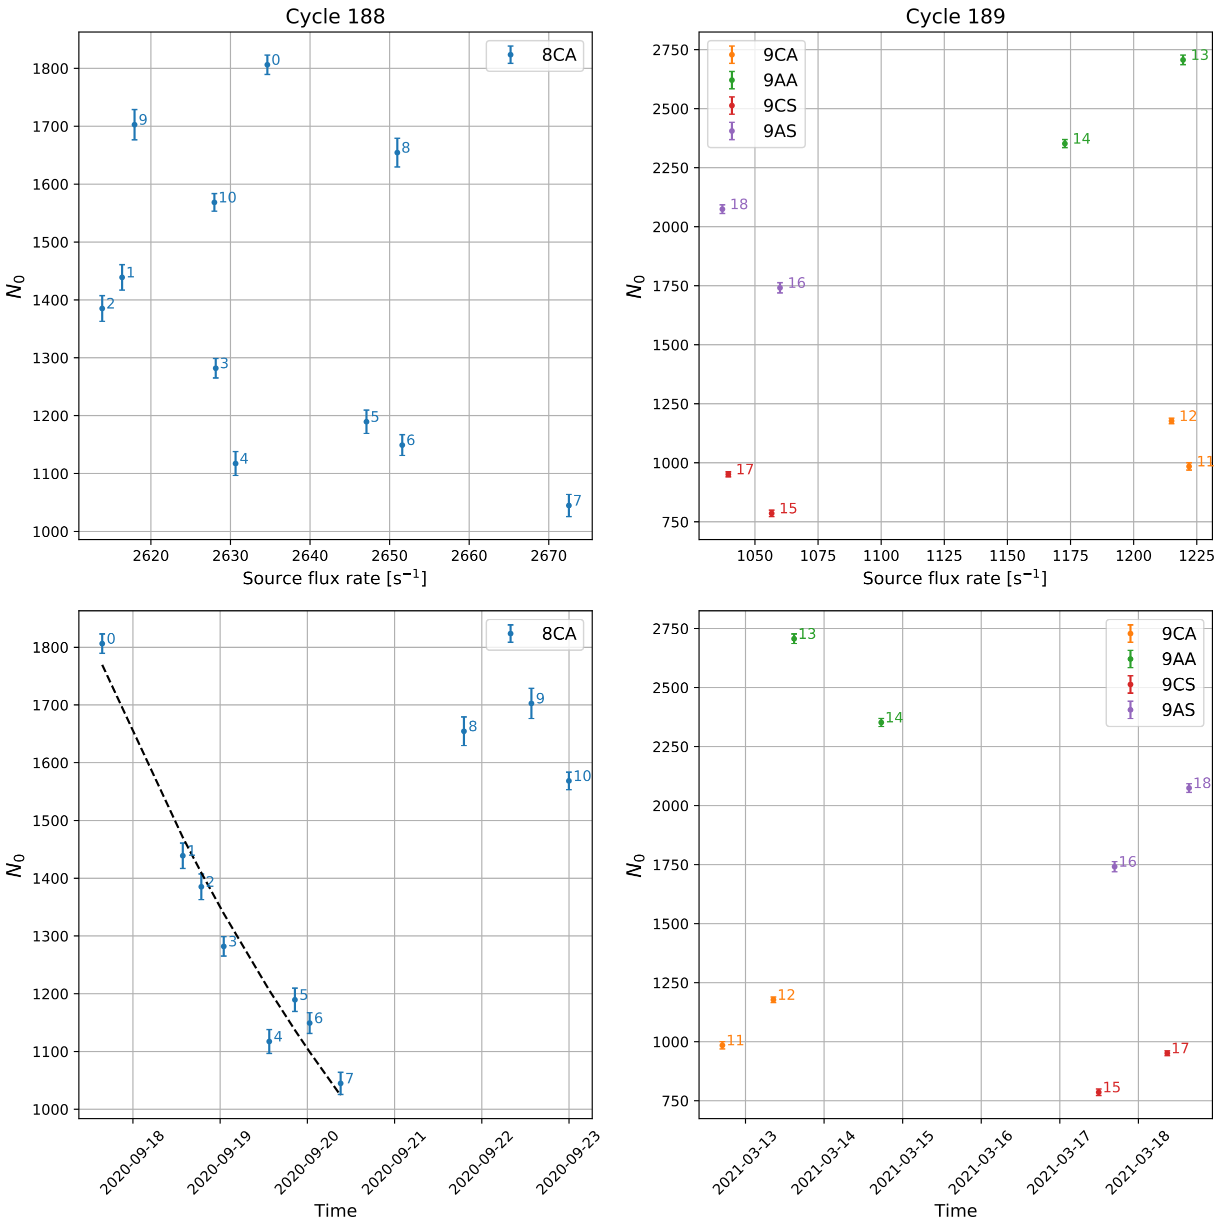

Figure 3.11: Trends in N_0. Data are labeled and indexed as in Figure 3.10. Top: N_0 fit values plotted against the source flux rate—interpolated from points in Figure 3.5—for Cycle 188 and Cycle 189. Bottom: N_0 plotted against measurement time. The black dashed line shows an exponential fit to values of N_0 for CYTOP measurements 0 through 7, all of which were performed while the system was below roughly 150 K. (See the text for more details.)

was relatively stable and operated in continuous mode. (Labeled as "8CA" in Figure 3.10.) In contrast, the plot at left in Figure 3.10 shows a wide variation of N_0 with bottle temperature. This is perhaps surprising at first. For one, the increase in τ_0 with decreasing temperature means that the bottle is better able to store UCN, and one therefore expects that the initial number of UCN in the bottle after filling should be larger at lower temperatures. However, the relatively large hole in the front of the bottle dominates the time constants associated with filling and emptying, so that even when $\tau_0 \sim \tau_\beta$, one does not expect a large increase in N_0. (This point is treated in greater detail in Section 4.3.)

Such a mechanism, however, does not explain the seemingly random variation of N_0 with bottle temperature. It is therefore instructive to plot N_0 against other experimental variables. The top two plots in Figure 3.11 show the variation of N_0 in Cycle 188 and Cycle 189 against the source flux rate, while the bottom two show its variation with absolute measurement time. The former do not seem to show a clear trend. The Cycle 189 N_0 data could be interpreted as increasing with increasing source flux rate, but it could also be the case that the bare aluminum entrance has lower loss than the steel. These two explanations could also be used to describe the trends in the bottom right plot of Figure 3.11, but there are insufficient data to rule in favor of one possibility or the other. The bottom left plot of Figure 3.11 contains the most clear pattern. There, N_0 values from storage measurements 0 through 7—see first column in Table 3.1—follow a clear decreasing trend. All of these measurements were also taken while the system was below 130 K. It would then seem possible that this trend is related to the previously discussed systematic decrease, characterized by τ_d. To compare the two, a fit to these N_0 was made using a function $\propto \exp(-t_m/\tau_d)$, which returned $\tau_d = 120(14)\,\mathrm{h}$. (See black dashed line in Figure 3.11.) The three values of τ_d from Table 3.1 to which this may be compared are for measurements 0, 3, and 6. These give a weighted average of $\tau_d = 100(7)\,\mathrm{h}$, in reasonable agreement.

3.3 Deuterated polyethylene (dPE)

Deuterated polyethylene (dPE) is a polymer made up of repeating units of C_2D_4. Its chemical structure is identical to that of "regular" polyethylene (PE)—repeating units of C_2H_4—but with hydrogen replaced by deuterium. This substitution increases the optical potential of dPE with respect of PE: deuterium has a positive scattering length, $a = 6.7\,\mathrm{fm}$, whereas hydrogen has a negative scattering length, $a = -3.7\,\mathrm{fm}$ [1]. It also decreases the loss probability per wall interaction: the absorption cross section of hydrogen is larger than that of deuterium by a factor of about 67 [2]. The precise values of the optical potentials and loss factors of these materials, however, depend on their densities, which in turn depend on their method of preparation.

To make PE or dPE coatings, a solution of the powder form of the polymer, melted into a hot solvent, is applied to a surface. This solution is then heated to remove the solvent, and finally baked to "cure" the remaining plastic. The density of the coating after this operation depends not only on the distribution of chain lengths and branching ratios of the polymer [113]—which vary between manufacturers, but on the temperature history of the curing process [114]. Owing to a lack of widespread use of deuterated plastics, the quantitative effects of these factors on the density of dPE is not readily available in the literature. Attempts were made to measure the optical potential of dPE coatings on wafers via neutron reflectometry, as was done with CYTOP, however, sufficiently thin and uniform coatings were not achieved. (Similar issues were faced in attempts by others to perform reflectometry on dPE coatings [59, 60], in which reflectivity curves show large sample geometry effects.) We therefore estimate the optical potential of dPE in this work based on the possible range of densities of PE. The density of PE depends on its degree of "crystallinity", with the theoretical densities of fully-amorphous and fully-crystalline PE being $0.88\,\mathrm{g\,cm^{-3}}$ and $1.00\,\mathrm{g\,cm^{-3}}$, respectively [115]. Replacing hydrogen by deuterium will increase the density by a factor of 1.14, that is, by the ratio of the molar mass of C_2D_4 to C_2H_4. This predicts a density range for dPE of $1.01\,\mathrm{g\,cm^{-3}}$ to $1.14\,\mathrm{g\,cm^{-3}}$, or, taking the average, roughly $1.07(7)\,\mathrm{g\,cm^{-3}}$. Such a density corresponds to an optical potential of 210(10) neV—see Equation 2.12—which is in rough agreement with values reported in the literature: 214(10) neV [59] and 214.8(52) neV [60].

Bottle preparation The dPE coatings of the storage bottle components were mainly prepared by Euan Connolly [102], using batch number P41793-dPE product from Polymer Source, Inc. [116]. The coating of the main body of the storage bottle assembly was prepared in a process similar to that of the CYTOP coating. This component was filled with a mixture of 0.81 g of solid dPE in 110 mL deuterated p-Xylene, and mounted horizontally in the custom coating oven, with its open ends sealed by two Teflon disks, and the side port facing up. The oven temperature was brought to 130 °C and held there for approximately 1.5 hours, at which point the dPE was assumed to be completely melted. The bottle was then rotated at 12 rpm for one hour. After this, the bottle was dried using a supply of hot nitrogen gas, until the level of reclaimed solvent stopped increasing: approximately 1.5 hours. The bottle was gradually cooled to 75 °C at a rate of approximately $0.5\,\mathrm{^\circ C\,min^{-1}}$, at which point the oven was turned off, and the bottle came to room temperature over approximately 5 hours, with an initial cooling rate of roughly $0.2\,\mathrm{^\circ C\,min^{-1}}$. If the dPE

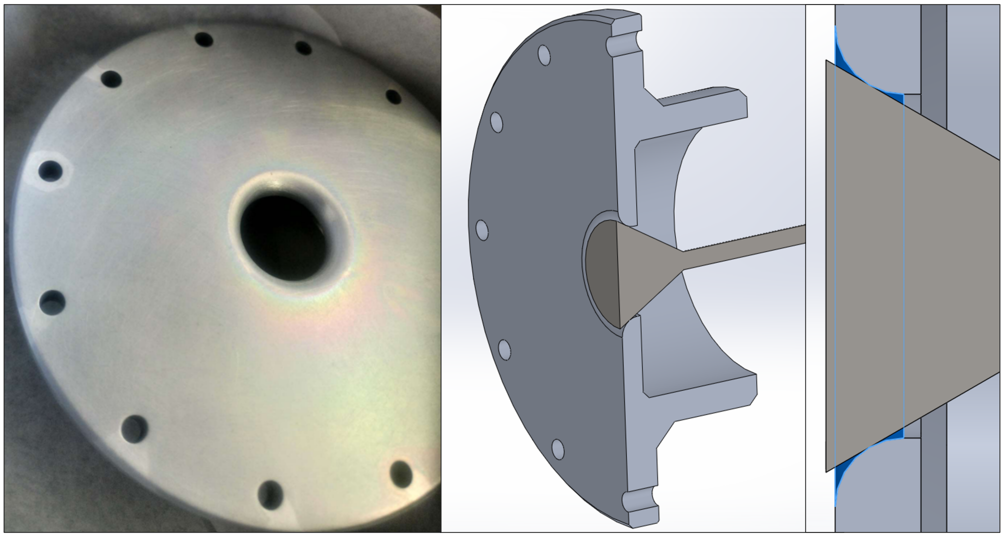

Figure 3.12: Left: photo of dPE coating of the inner surface of the storage bottle front flange. Note the rainbow coloring near the filleted entrance, evidence of possible thinning of the coating. Middle: cross section view in perspective of the front plate and plug assembly. Right: cross section view of plug and front flange interface. The entrance hole is 20 mm in diameter, and has a fillet radius of 4 mm. The plug widens from a 4 mm minimum diameter to a 24 mm maximum diameter at an angle of 30° to the plug axis. With these specifications, when the bottle is closed, the front face of the plug is offset from the inner face of the front plate by approximately 0.5 mm toward the inside of the bottle, and the surface area of the fillet exposed to the inside of the bottle—left portion of the surface highlighted in blue—is approximately 320 mm^2.

is assumed to be evenly distributed over the entire surface of the main body piece, approximately 1735 cm^2 in area, the estimated density of 1.07(7) $\mathrm{g\,cm^{-3}}$ for dPE gives an average layer thickness of approximately 4.3(2) µm.

The front and back bottle flanges were coated using a spin-coating process [102]. First, solid dPE was mixed with deuterated p-Xylene in a flask, and then fully melted by immersion of the flask in an oil bath at 120 °C. The hot solution was manually syringed onto the flanges while they were spun at 100 rpm. Prior to application of the solution to the flanges, they were heated for several minutes with a heat gun set to 160 °C and held approximately 20 cm away. (10 minutes for the back flange, 25 minutes for the front flange.) Both solutions were prepared with the same concentration: 80.3 mg of dPE in 6 mL of deuterated p-Xylene for the back flange, and 125.6 mg of dPE in 9.5 mL of deuterated p-Xylene for the front. In both cases, approximately 6.0(5) mL of solution was transferred to the flanges, where some splatter makes the exact amount uncertain. Under the assumption of a uniform coating of dPE over the entire surfaces of the back and front flanges, the average thickness should be approximately 7.9(9) µm in both cases. (It should be noted that the total surface areas of the back and front flanges which were coated, approximately 95 cm^2 and 93 cm^2, respectively, are not the same as those presented to the inside of the closed storage bottle, approximately

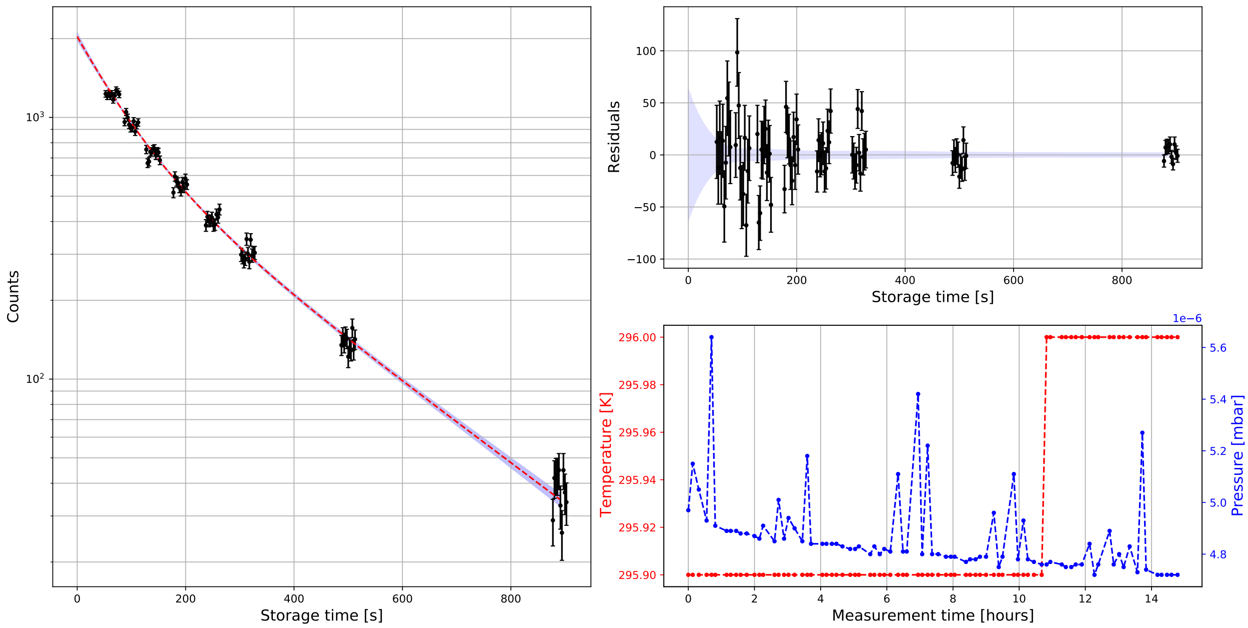

Figure 3.13: Example dPE storage data set taken in continuous mode during Cycle 189 at 295.9 K average and with the aluminum entrance to the bottle. $N_1 = 1207(56)$, $\tau_1 = 85(8)$ s, $N_2 = 829(77)$, $\tau_2 = 280(12)$ s, $\chi^2/(83-4) = 1.22$. Left: Storage data are shown in black, and are displaced horizontally to show vertical spread. The fit is shown in red, with 1-σ error bands in light blue. (See Equation 3.4.) Top right: residuals from the plot at left. Bottom right: Average temperature and pressure during the measurement.

64 cm^2 and 60 cm^2, respectively.) However, this assumption of a uniform thickness is clearly not applicable to front flange, whose coating shows signs of thinning around the filleted entrance hole, as can be seen in Figure 3.12.

The steel bottle plug coating was prepared by hand [102]. It was repeatedly dipped into and extracted from a hot solution of 106.2 mg of dPE dissolved in 27.5 mL of deuterated p-Xylene, drying in air for approximately 3 minutes between dips. The pull out speed is estimated to be roughly 2 mm s^{-1}. During each dip, the plug was submerged to about halfway up the rod section, with the conical section totally immersed in the hot solution. After 10 dips, 22.0 mL of solution remained. Assuming a uniform coating over the area of the cone and half of the rod section, approximately 17.4 cm^2, the coating thickness should be approximately 11.3(7) µm.

Storage measurements In contrast to storage measurements in CYTOP, the dPE data are not well described by a single exponential model. Instead, a double exponential model is used, as that of Equation 2.20:

$$N_1 \exp(-t_s/\tau_1) + N_2 \exp(-t_s/\tau_2)\,, \tag{3.4}$$

where t_s is the storage time, and where τ_1 is taken to be the shorter of the two time constants. (See Figure 3.13 for an example.) As with the CYTOP-coated bottle at low temperatures, a systematic decrease in counts was observed in long storage measurements. In such situations, equation 3.4 is modified as

$$\exp(-t_m/\tau_d)\left[N_1\,\exp(-t_s/\tau_1) + N_2\,\exp(-t_s/\tau_2)\right], \tag{3.5}$$

where t_m is the absolute measurement time, taken to be zero at the beginning of a measurement, and τ_d is a time constant associated with the systematic decrease in the counts. (See Figure 3.14 for an example.) Unfortunately, it is not the case that the exact least squares estimates of the fit parameters in Equations 3.4 and 3.5 may be solved for analytically. Therefore, a numerical search of parameter space was performed using the `curve_fit` function in SciPy [109], using the MINPACK implementation [117] of the "Levenberg-Marquardt" algorithm[118, 119]. This function also returns an estimate of the inverse of the Hessian, from which uncertainties are estimated. The results are summarized in Table 3.2, and plots are shown against bottle temperature in Figure 3.15. (In Figure 3.15, two sets of data have been excluded: indices 2 and 24. It is not definitively known why these produce such large uncertainties in their fit parameters, but these data sets are among those with the fewest points.)

In the plot at right in Figure 3.15, τ_1 does not appear to vary strongly with bottle temperature, while τ_2 shows a gradual increase with decreasing temperature. As with CYTOP, the latter trend implies a reduction in the loss factor η at low temperature. Both behaviors are consistent with the weak patch model, which predicts two approximate time constants: one that is related to the "main" wall coating, τ_w of Equation 2.26, and the other that is related to the weak patches, τ_p of Equation 2.28. The thinning of the dPE coating, visible in in Figure 3.12, further supports such a scenario.

Just as for τ_d extracted from CYTOP storage measurements, the outliers in τ_d from dPE storage data—that is, 191(20) h and 362(78) h in Table 3.2—come from accumulation mode measurements. As discussed previously, it is suspected that these longer time constants are related to the varying thermal load placed on the SUN-2 production volume by a prolonged activity of the source UCN valve. When these outliers are excluded, the weighted average of τ_d across cycles is 107(3) h, which is in rough agreement with the values found for CYTOP: 120(14) h and 100(7) h. (See previous section.) This agreement supports the hypothesis that the mechanism for the systematic decrease is inherent to the experiment being at low temperature, and not to either CYTOP or dPE as coatings.

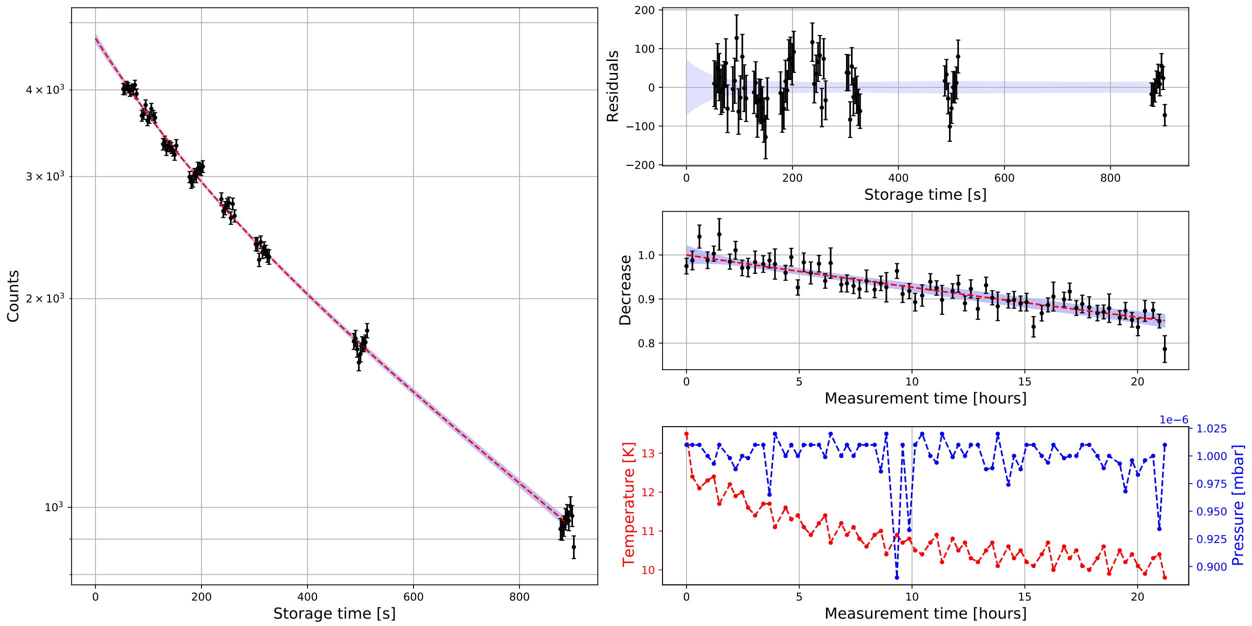

Figure 3.14: Example dPE storage data set taken in accumulation mode during Cycle 189 at 10.9 K average and with the aluminum entrance to the bottle. $N_1 = 1332(215)$, $\tau_1 = 163(33)\,\mathrm{s}$, $N_2 = 3416(259)$, $\tau_2 = 691(40)\,\mathrm{s}$, $\tau_d = 131(7)\,\mathrm{h}$, $\chi^2/(70-5) = 1.09$. Left: Storage curve with the systematic decrease corrected for. (Equation 3.5 divided by the factor $\exp(-t_m/\tau_d)$.) Data are shown in black, and are displaced horizontally to show vertical spread. The fit is shown in red, with 1-σ error bands in light blue. Top right: residuals from the plot at left. Middle right: the systematic decrease. (Equation 3.5 divided by Equation 3.4.) Bottom right: Average temperature and pressure during the measurement.

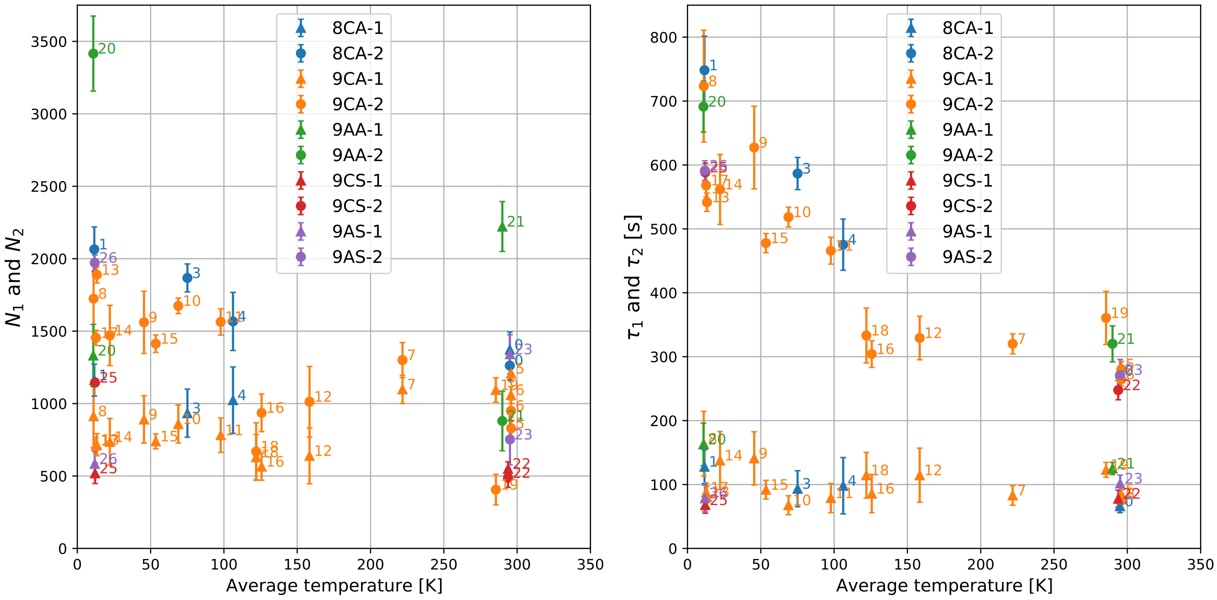

Figure 3.15: Fit parameters for dPE data. Points are labeled by a chronological index. Points with very large error bars are not plotted. (Indices 2 and 24 in Table 3.2.) Data are grouped by cycle, source mode, and bottle entrance, and i for N_i and τ_i, with the triangles indicating $i = 1$ and circles indicating $i = 2$. For example, "8CA-1" corresponds to Cycle 188, continuous mode, aluminum bottle entrance and both N_1 and τ_1, while "9AS-2" corresponds to Cycle 189, accumulation mode, stainless steel bottle entrance, and both N_2 and τ_2. Left: N_1 and N_2 plotted against average bottle temperature. Right: The storage time constants τ_1 and τ_2 plotted against the average bottle temperature. (These plots use vector graphics, making it possible to distinguish the sometimes-clustered points above on a computer by zooming in.)

Table 3.2: Summary table of dPE storage measurements. The first column assigns a chronological index to the measurement. The "Date" column lists the start date and time of a measurement. "Front" indicates the entrance to the storage bottle: either the uncovered aluminum, "Alu.", or with the polished stainless steel insert, "SS.", in place. "Size" is the number of points in the data set. "Mode" is either continuous, "C", or accumulation, "A". "T" is the average bottle temperature during the measurement. (Various problems with the temperature sensors in Cycle 188 necessitated their replacement after the cycle. The average bottle temperatures reported for these data sets were reconstructed using the same heater settings, but with new temperature sensors.) All parameter estimates are least squares estimates. χ_r^2 gives the value of the reduced chi-squared statistic.

Index	Cycle	Date	Front	Size	Mode	T [K]	τ_1 [s]	τ_2 [s]	N_1	N_2	τ_d [h]	χ_r^2
0	188	20/09/24 11h	Alu.	34	C	295.0	66(10)	270(11)	1372(122)	1262(99)	–	1.49
1	188	20/09/25 11h	Alu.	32	C	11.6	127(26)	748(53)	1161(110)	2064(155)	101(20)	1.23
2	188	20/09/25 20h	Alu.	15	C	40.6	43(31)	613(21)	2111(3063)	2215(71)	–	0.77
3	188	20/09/26 00h	Alu.	31	C	75.1	93(28)	586(25)	933(166)	1867(96)	–	1.87
4	188	20/09/26 09h	Alu.	10	C	106.2	98(44)	475(40)	1023(229)	1567(200)	–	0.46
5	189	21/02/28 19h	Alu.	83	C	295.9	84(7)	280(12)	1207(56)	829(77)	–	1.22
6	189	21/03/01 19h	Alu.	82	C	295.8	77(8)	264(9)	1059(57)	950(74)	–	1.13
7	189	21/03/02 15h	Alu.	30	C	221.8	82(15)	320(15)	1098(97)	1300(120)	85(32)	0.88
8	189	21/03/03 08h	Alu.	34	C	11.1	163(50)	723(87)	912(215)	1724(261)	86(18)	1.02
9	189	21/03/03 17h	Alu.	26	C	45.5	140(41)	627(64)	890(163)	1561(214)	–	0.74
10	189	21/03/04 00h	Alu.	38	C	68.9	67(14)	518(15)	858(131)	1674(53)	–	1.13
11	189	21/03/04 09h	Alu.	24	C	97.8	78(22)	465(21)	781(118)	1563(90)	76(30)	0.71
12	189	21/03/04 17h	Alu.	24	C	158.4	114(42)	329(34)	638(193)	1012(243)	–	0.87
13	189	21/03/05 00h	Alu.	65	C	13.3	80(16)	541(14)	715(76)	1888(56)	134(18)	1.30
14	189	21/03/05 13h	Alu.	28	C	22.2	137(45)	561(54)	735(161)	1470(208)	98(38)	0.65
15	189	21/03/05 20h	Alu.	83	C	53.5	91(14)	477(14)	739(51)	1413(60)	97(8)	1.47
16	189	21/03/06 14h	Alu.	24	C	125.7	85(29)	304(21)	564(93)	936(130)	–	1.26
17	189	21/03/06 22h	Alu.	73	C	12.7	87(15)	568(17)	709(57)	1453(52)	103(10)	0.90
18	189	21/03/07 14h	Alu.	24	C	122.0	114(35)	333(43)	629(156)	668(198)	–	0.86
19	189	21/03/08 10h	Alu.	70	C	285.5	123(11)	360(41)	1094(84)	406(105)	–	1.04
20	189	21/03/09 13h	Alu.	70	A	10.9	162(33)	691(39)	1331(214)	3415(259)	131(6)	1.09
21	189	21/03/11 12h	Alu.	73	A	289.9	125(9)	320(28)	2221(172)	879(205)	362(78)	0.98
22	189	21/03/19 20h	SS.	75	C	293.8	77(13)	247(15)	549(48)	487(66)	–	1.07
23	189	21/03/20 10h	SS.	42	A	295.1	101(13)	271(24)	1340(134)	752(175)	–	1.44
24	189	21/03/21 10h	SS.	24	A	14.6	339(155)	1500(2766)	2029(1346)	824(1392)	–	1.62
25	189	21/03/21 18h	SS.	103	C	12.1	67(12)	588(14)	518(70)	1144(27)	95(6)	1.32
26	189	21/03/22 14h	SS.	64	A	11.9	78(20)	591(15)	584(84)	1972(53)	191(20)	0.97

We now consider the UCN "population" parameters N_1 and N_2. (See plots at left in Figure 3.15.) The plots of N_1 and N_2 against the average bottle temperature do not seem to show any clear trend, other than that at higher temperatures, N_1 seems to be larger than N_2. To search for systematic effects, N_1 and N_2 are plotted in Figure 3.16 against the source flux rate and the absolute measurement time. Of all data presented there, the clearest trend is for that labeled "9CA-1" and "9CA-2". (This corresponds to N_1 and N_2, respectively, from measurements in Cycle 189, while operating the source in continuous mode, and with an aluminum entrance to the storage bottle.) Roughly speaking, the collection of N_1 and N_2 seem to increase with increasing source flux, and decrease with increasing time. (In fact, these two statements are equivalent, because the source flux decreases over time.) On closer inspection, it is seen that the N_1 and N_2 are "mixed" together. Consider, for example, the lower right plot of Figure 3.16. Initially, $N_1 > N_2$. Then, for intermediate measurements $N_1 < N_2$. Finally, for the last measurement, $N_1 > N_2$. Examining Table 3.2, it seems possible that this behavior is correlated with the temperature history of the bottle: early measurements occur at high temperature, after which the system is cooled and held there, until the very last measurement when the system is returned to near room temperature. In the previous section, for the CYTOP measurements, a decrease in N_0 data was found to occur while the system was cold: see lower left plot and the fit in Figure 3.11. It therefore seems possible that a similar mechanism may be responsible for the decrease in N_1 and N_2 over time.

To compare the decrease in the CYTOP data at low bottle temperatures with the decrease in the dPE data—possibly due to low bottle temperatures—the appropriate quantity to consider is the sum $N_1 + N_2$. As with N_0, this parameter should be related to the total UCN population at time zero, that is $\int \mathrm{d}E\, n(0, E)$. (See Equation 2.18.) The top left plot in Figure 3.17 shows $N_1 + N_2$ against absolute measurement time, and the decrease appears to occur roughly between the 3rd and 8th of March. The plot also shows two fits to these data, for measurement indices 8 through 19. (See Table 3.2.) The black dashed curve shows a fit to the exponential function $\alpha \exp(-t_m/\beta)$, and the red dashed curve shows a fit to the linear function $A + Bt_m$, where t_m is the measurement time in hours. Both fits yield time constants that are in rough agreement: $\beta = 265(43)\,\mathrm{h}$, and $-A/B = 302(51)\,\mathrm{h}$. However, these time constants are larger than the value of $\tau_d = 107(3)\,\mathrm{h}$ calculated above, in which two large outliers were excluded. Interestingly, these outliers appear to lie just outside the error bars of β and $-A/B$, suggesting a possible connection. These values could also be related to the $\tau_d = 315(34)\,\mathrm{h}$ outlier from the CYTOP measurements. (See Table 3.1.)

It would also seem to be the case that the decrease in the measured source flux is not related in an obvious way to the decrease in $N_1 + N_2$. In Figure 3.5 in Section 3.1, an exponential fit has been made to the source flux rate during the period of time in which the "9CA" data were taken. This returns a time constant of 863(151) h, substantially longer than those estimated by β and $-A/B$. (See the red fit and caption in the plot at right in Figure 3.5.) Now, it must be the case that the source flux and $N_1 + N_2$ are related, for if the source is "turned off" there will be no UCN in the bottle. However, as was discussed previously, the flux is measured while the source is operated in continuous mode, with the UCN switch in the bypass position. In this configuration, the UCN spectrum source should have a higher average energy than when the switch is in the filling position and the storage bottle is open, because in the latter situation, UCN are stored in the entire volume consisting of the source, the bottle, and the guides leading to the bottle. If the mechanism responsible for the slow decrease in the source flux preferentially attenuates faster UCN—as would be the case for a contaminant, with a lower optical potential than steel, freezing onto the walls of the guide—$N_1 + N_2$ may not change significantly. That is, depending on the specifics of the energy filtering process, the "storable" portion of the UCN spectrum may be largely unaffected. Such a scenario, however, is purely speculation, and further measurements are necessary to understand the effect.

The lack of a clear understanding of the systematic effects responsible for the variation in N_1 and N_2 makes it useful to examine the relative population fractions, $x_1 = N_1/(N_1+N_2)$ and $x_2 = N_2/(N_1+N_2)$. These are plotted at bottom in Figure 3.17, against absolute measurement time and bottle temperature. In the bottom left plot, the "mixing" of N_1 and N_2 discussed previously seems to be present. The plot immediately to its right, however, suggests that, when viewed in terms of x_1 and x_2, such mixing is most strongly controlled by the bottle temperature. There, x_1 clearly decreases at low temperature, while x_2 clearly increases. At first, this behavior might be interpreted as suggesting that at room temperature, there are more UCN of higher average energy—associated with N_1 and the shorter time constant τ_1—in the initial spectrum $n(0, E)$, while at low temperature, the spectrum is shifted toward lower energy. Such a conclusion would be inconsistent with the weak patch model, which offers one possible explanation for the lack of variation of τ_1 with bottle temperature. In the model, the relative fractions of the UCN population are fixed by the value of the weak patch potential, which does not change with temperature. (See Section 2.4.) It is important, however, to note that the multiple time constant approximation to the storage curve, Equation 2.20, does not include any physics *per se*, and that N_1 and N_2 do not necessarily correspond to, for example, N_p and N_w discussed in Section 2.4.

To develop this claim, recall that in the example of the weak patch model given in Section 2.4, which was for the present bottle geometry and a particular choice of $n(0, E)$, the loss factor was chosen to be $\eta_w = 10^{-5}$. This choice was made on the basis of the previous discussion in this section, in which it was shown that the storage curve deviates from single, exponential behavior for larger loss factors, $\eta_w \sim 10^{-4}$. (See Figure 2.6.) Consider, then, a hybrid situation, in which the loss factor at room temperature is $\eta_w \sim 10^{-4}$, and a portion of the bottle surface consists of weak patches. In such a case, a fit of Equation 3.4 to the number of UCN in the bottle, $N(t)$, would not be able to distinguish between a short time constant due to large η_w or due to the presence of weak patches, and N_1 would thus be related to the sum of contributions from the two effects. If the loss factor were to decrease to $\eta_w \sim 10^{-5}$ for low bottle temperatures, one would recover the example discussed in Section 2.4, and N_1 and N_2 would approximately correspond to N_p and N_w, respectively. Of course, it remains to be seen whether such an explanation is reasonable when the experimental bottle spectrum is considered. The simulation of this distribution is the subject of the next chapter.

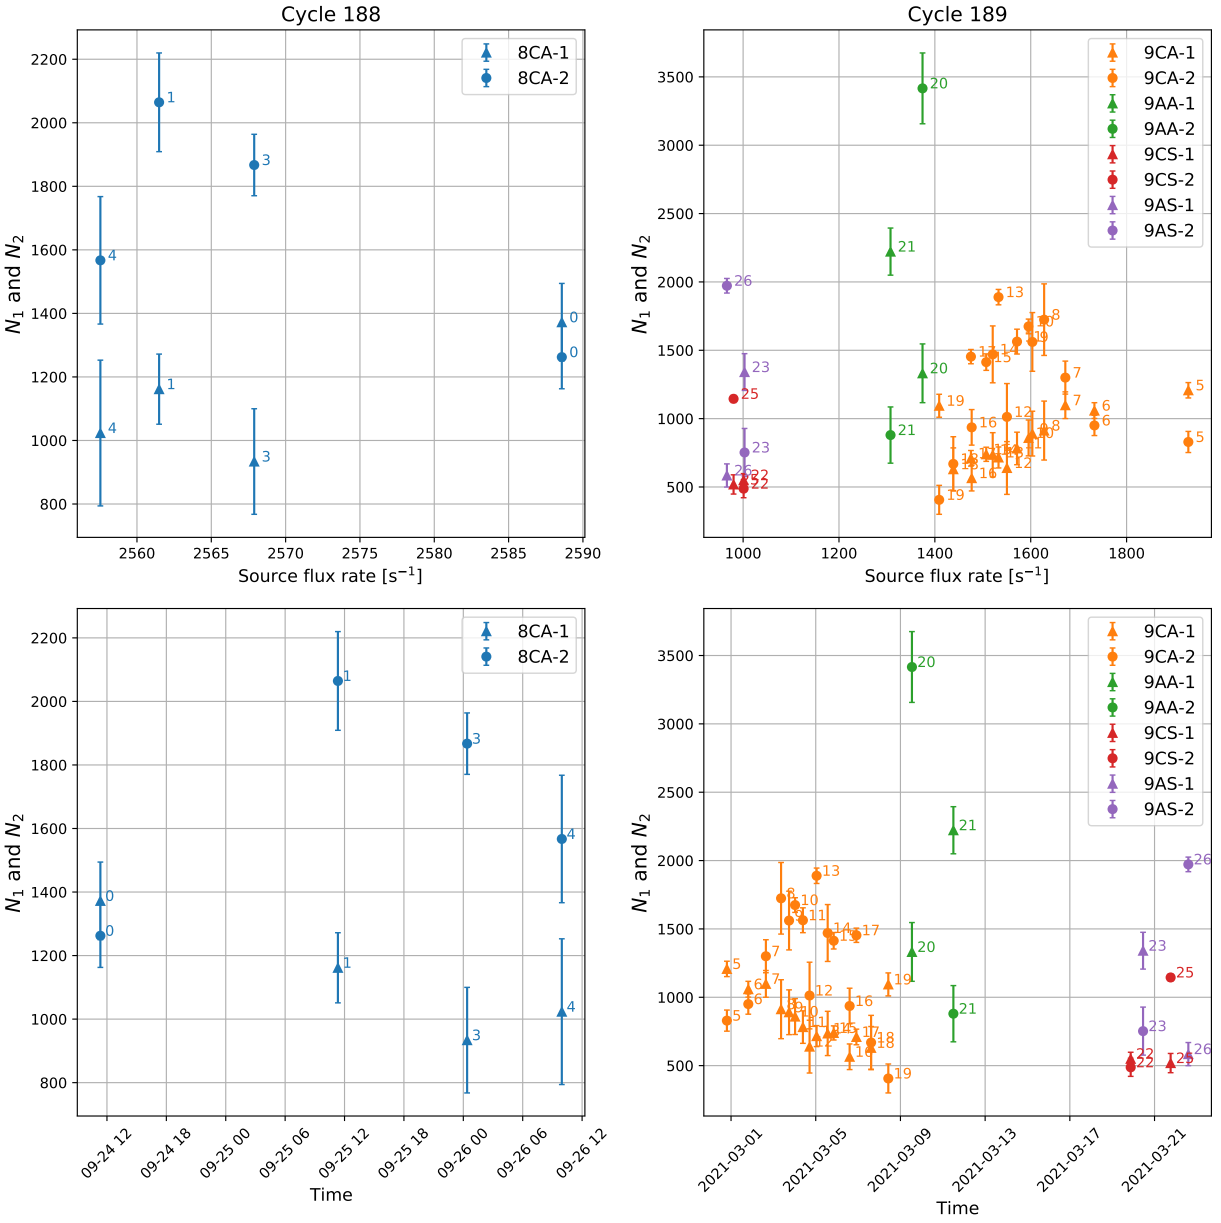

Cycle 188

Cycle 189

Figure 3.16: Trends in N_1 and N_2 for dPE. Data are labeled and indexed as in Figure 3.15. Top: N_1 and N_2 fit values plotted against the source flux rate—interpolated from points in Figure 3.5—for Cycle 188 and Cycle 189. Bottom: N_1 and N_2 plotted against measurement time.

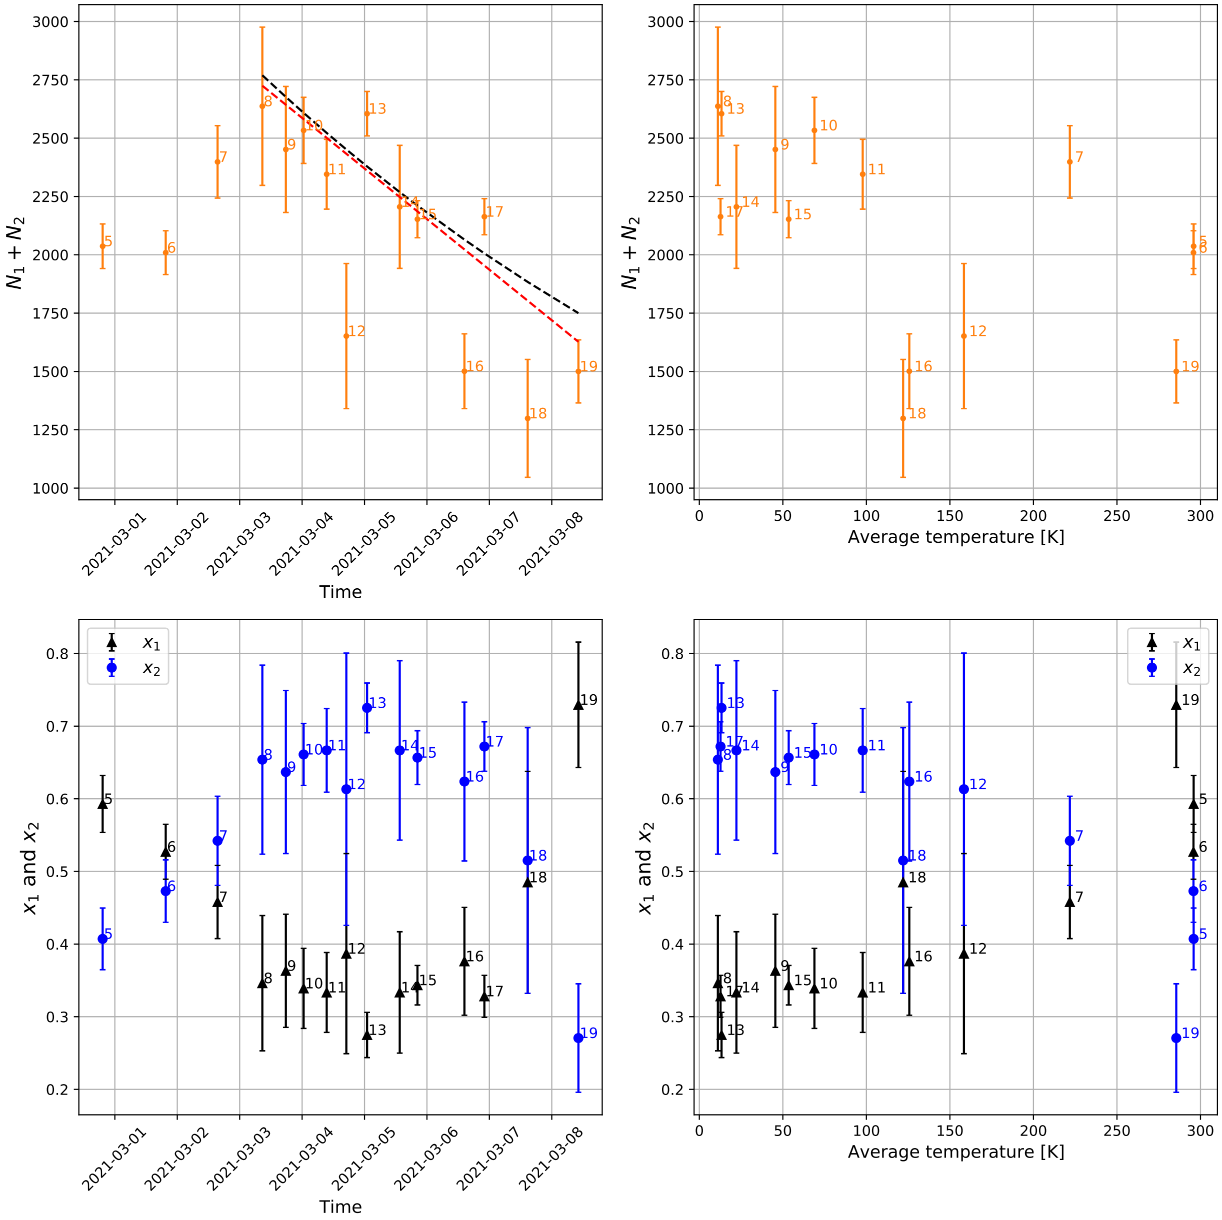

Figure 3.17: Trends in $N_1 + N_2$, and $x_1 = N_1/(N_1 + N_2)$ and $x_2 = N_2/(N_1 + N_2)$ for dPE, specifically from Cycle 189, operating the source in continuous mode, and with the aluminum bottle entrance. (“9CA” in Figures 3.15 and 3.16.) Data are indexed chronologically. (See Table 3.2.) Top left: $N_1 + N_2$ fit values plotted against measurement time. The black and red lines are exponential and linear fits, respectively. See the text. Top right: $N_1 + N_2$ fit values plotted against the average bottle temperature. Bottom left: x_1 and x_2 plotted against measurement time. Bottom right: x_1 and x_2 plotted against the average bottle temperature.

Chapter 4

Simulations

One of the primary goals in this thesis is to extract the loss factors, η, of CYTOP and dPE coatings from experimental storage data. In principle, this could be achieved by fitting these data to Equation 2.18 using the storage models presented in Chapter 2. However, the initial energy spectrum of UCN in the bottle, $n(t = 0, E)$, is unknown. Furthermore, this spectrum depends not only on the details of the UCN source and the guide system, but on the storage properties of the bottle itself. It would seem, then, that one must know either both $n(t = 0, E)$ and η, or neither.

Some insight into this problem may be gained from studying it with a computer simulation. For example, suppose the entire experiment—source, guides, and storage bottle—were simulated using one particular set of values for all the parameters characterizing the setup. (These parameters include, but are not limited to, the loss factors, optical potentials, and roughnesses of each component of the simulation.) The results of such a calculation could then be compared to experimental storage data. The process could then be repeated, iteratively refining the unknown input parameters until the simulation output best matches experimental data. However, a very large number of measurements would be required to unambiguously determine these parameters.

We therefore truncate the problem, and, in this chapter, present a modified version of such an approach. Several assumptions are made so that $n(t = 0, E)$ may be estimated using only a subset of the quantities that would be calculated in the "full" simulation described in the previous paragraph. In the first section of this chapter, this subset and its underlying assumptions are discussed. This is followed by a presentation of the simulation results for each subset. Finally, the individual results are synthesized in order to produce an estimate of the initial energy spectrum of UCN in the storage bottle.

Computer simulations were performed with Geant4, an extensive set of C++ libraries developed by the European Organization for Nuclear Research (CERN), primarily to simulate the passage of particles through matter in high-energy physics [120, 121]. Referred to as a toolkit, its core functionality enables the simulation of particle transport as a result of a variety of physical interactions, spanning a wide energy range, from eV to TeV. Since its initial development, the Geant4 toolkit has been extended by an external collaboration into the neV range in order to simulate UCN transport in gravitational and optical potentials [55, 122]. It is this functionality that is used in the present work.

4.1 Overview

In this section we introduce and motivate the proposed subset of simulations needed to estimate an initial energy spectrum of UCN in the storage bottle. First, a few comments about notation are in order. In the following discussion, we will need to consider the energy spectrum of UCN in the SUN-2 source. Therefore, to avoid confusion, we will write $n_b(t, E)$ for the time-dependent spectrum in the storage bottle, and $n_s(t, E)$ for the time-dependent spectrum in the SUN-2 source. Additionally, care must be taken not to confuse the UCN energy label E. UCN are produced in the "production volume" of the SUN-2 source, which is filled with ^{4}He, having an optical potential $V_m = 18.5\,\text{neV}$. (See Chapter 3.) When UCN exit the superfluid medium—as they must in order to reach the storage bottle—they will gain V_m kinetic energy. They must then travel a distance Y vertically, thereby losing mgY kinetic energy. Let E_b be the energy of a UCN with respect to vacuum in the storage bottle, and E_s be the energy of a UCN with respect to superfluid ^{4}He in the SUN-2 source. We define the zeros of these energies as being on the axes of the bottle and source, respectively, which are perpendicular to gravity. They are then related by

$$E_b = E_s + V_m - mgY\,, \tag{4.1}$$

where $Y = 356\,\text{mm}$ is the vertical distance between the bottle and source axes. One consequence of this is that only those UCN produced with $E_s > |V_m - mgY| \approx 18\,\text{neV}$ may reach the bottle. (Actually, this is not quite true, because the entrance hole to the bottle is on its axis, and is 20 mm in diameter; however, this is accounted for in the simulations.)

We begin by considering the UCN energy distribution in the SUN-2 source, which may be operated in either a continuous or accumulation mode. In the former mode, UCN are stored not only in the source, but in the section of guide leading to the UCN switch. This greatly increases the size of parameter space for exploration via simulation, and **we therefore only consider the initial spectrum of UCN in the storage bottle generated by operating the source in accumulation mode.** In accumulation mode then, the spectrum of UCN in the source is governed by the following differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t}\Big[n_s(t,E_s)\,\mathrm{d}E_s\Big] = R_p(t,E_s) - \frac{1}{\tau_s(E_s)}\Big[n_s(t,E_s)\,\mathrm{d}E_s\Big]\,, \tag{4.2}$$

where $\tau_s(E_s)$ is the storage time constant of the source volume when it is closed—see Equation 2.16, and $R_p(t,E_s)$ and $n_s(t,E_s)\,\mathrm{d}E_s$ are the rate of production of UCN and number of UCN in the source, respectively, at time t for energies between E_s and $E_s + \mathrm{d}E_s$. (For the sake of brevity in the following discussion, we will absorb the factor $\mathrm{d}E_s$ into $n_s(t,E_s)$. Also, the factor of $\mathrm{d}E_s$ is already contained in $R_p(t,E_s)$; see below.) For long "accumulation times", i.e. $t \gg \tau_s(E)$, the UCN density in the source saturates asymptotically. We assume that the cold neutron (CN) flux on the production volume does not vary with time after it is "switched on", so that the UCN production rate, $R_p(E_s)$, is also time-independent. In this limit, the solution to the above differential equation is then $n_s(t\to\infty,E_s) = R_p(E_s)\tau_s(E_s)$. We must therefore calculate $R_p(E_s)$ and $\tau_s(E_s)$. The latter quantity can be derived from simulations in which UCN of energy E_s are generated in the source, and the times at which they are lost are recorded. This procedure is discussed Section 4.2. It then remains to calculate $R_p(E_s)$, which may be done analytically. (Useful references in what follows are [3, 92, 93, 123, 124, 125].)

In general [3],

$$R_p(E) = \int_{\mathcal{V}} \mathrm{d}^3r \int \mathrm{d}E'\,\phi(E',\mathbf{r})\Sigma(E' \to E)\,, \tag{4.3}$$

where two integrals are performed: the first over the volume $\mathcal{V}$ of the source, and the second over the range of energies E' of the CN beam incident on the production volume. The quantity $\phi(E',\mathbf{r})\,\mathrm{d}E'$ is the magnitude of the flux of CN incident on the converter at position $\mathbf{r}$ with energies between E' and $E' + \mathrm{d}E'$, and has units of number per area per second. $\Sigma(E' \to E)$ is the macroscopic inelastic scattering cross section to convert a CN with energy E' to a UCN with energy E, and has units of inverse length. We approximate the beam profile as uniform over the volume of the source, so that the $\boldsymbol{r}$ dependence is removed from the above expression. It may then be shown (see e.g. Equation 3.37 in [3]) that the production rate of UCN is

$$R_p(E_s) \propto \sqrt{E_s}\,, \tag{4.4}$$

where the proportionality constant is related to terms that depend on the CN flux, the temperature of the production volume, and properties of the excitations of superfluid ^{4}He. Thus, $n_s(t\to\infty,E_s) \propto \sqrt{E_s}\tau_s(E_s)$. (We are only interested in the shape of $n_b(t,E_b)$, and it will be shown that knowledge of the proportionality constant is not necessary.)

It now remains to propagate this distribution of UCN from the source to the storage bottle. The differential equation governing the UCN energy distribution in the bottle is

$$\frac{\mathrm{d}}{\mathrm{d}t} n_b(t, E_b) = R_f(t, E_b) - \frac{1}{\tau_e(E_b)} n_b(t, E_b) \,, \tag{4.5}$$

where $\tau_e(E_b)$ is the storage time constant of the bottle when it is open, $R_f(t, E_b)$ is the rate of filling into the bottle, and $n_b(t, E_b)\,\mathrm{d}E_b$ is the number of UCN in the bottle. (As with $n_s(t, E_s)$, we have absorbed the factor of $\mathrm{d}E_b$ into $n_b(t, E_b)$ in the above expression and in what follows. Also, the factor of $\mathrm{d}E_b$ is already absorbed into $R_f(t, E_b)$; see below. Finally, for notational uniformity, we might have called the time constant $\tau_b(E_b)$ rather than $\tau_e(E_b)$; however, the meaning of $\tau_e(E_b)$ is closest to an "emptying" time constant, and we will therefore use the "e" subscript.) Just as with the determination of $n_s(t, E_s)$, there are two quantities to calculate: $R_f(t, E_b)$ and $\tau_e(E_b)$. We first consider how to simulate the latter.

We guide this process by considering analytical expressions for $\tau_e(E_b)$. Ignoring the effects of gravity, the storage time constant of a bottle with a hole in its side is

$$\begin{aligned} \frac{1}{\tau_e(E_b)} &= \frac{A_b}{4\mathcal{V}} v_b \bar{\mu}(E_b) + \frac{1}{\tau_\beta} + \frac{A_h}{4\mathcal{V}} v_b \\ &= \left(\frac{A_b}{4\mathcal{V}} v_b \bar{\mu}(E_b) + \frac{1}{\tau_\beta} \right) + \frac{1}{\tau_h(E_b)} \\ &\approx \frac{1}{\tau(E_b)} + \frac{1}{\tau_h(E_b)} \,. \end{aligned} \tag{4.6}$$

where $v_b = \sqrt{2E_b/m}$ is the UCN velocity, A_h is the area of the hole, and A_b is the area of the inner bottle surface. (Compare to, for example, Equation 2.16.) The second line serves to define the time constant $\tau_h(E_b)$ due to the hole, while the third line makes the approximation that, when $A_h \ll A_b$, the sum of the first and last terms on the right-hand-side of the above expression is very nearly equal to the storage time constant of the bottle when it is closed. In light of the relatively long experimental time constants measured for CYTOP and dPE—see Tables 3.1 and 3.2—we make the approximation that the dominant contribution to $\tau_e(E_b)$ is from the hole. For example, for the bottle geometry studied in this work, $\tau_h(E_b) \approx 34\,\mathrm{s}$ for $E_b = 10\,\mathrm{neV}$, and $\tau_h(E_b) \approx 16\,\mathrm{s}$ for $E_b = 50\,\mathrm{neV}$. However, $\tau_e(E_b)$ cannot be approximated as only, $\tau_h(E_b)$, because this quantity diverges as $v_b \to 0$. We therefore use

$$\frac{1}{\tau_e(E_b)} \approx \frac{1}{\tau_h(E_b)} + \frac{1}{\tau_\beta}\,, \tag{4.7}$$

including only contributions from the hole and beta decay. As with $\tau_s(E_s)$, $\tau_e(E_b)$ will be determined from simulations in which UCN of energy E_s are generated in the bottle, and the times at which they are lost are recorded. (See Section 4.3 for details and the simulation results.)

It now remains to simulate $R_f(t, E_b)$. This is complicated by the fact that UCN may leave the bottle during filling and return at later times. In principle, this filling and "refilling" can be monitored in the simulation by keeping track of the number of UCN in the bottle at a given time. At present, however, such a scheme is difficult to implement in the simulation tool. Nevertheless, $R_f(t, E_b)$ can be recast in terms of two quantities that are straight forward to simulate: the rate of "direct filling", $R_f^{(0)}(t, E_s)$, and the rate of "refilling", $R_f^{(1)}(t, E_s)$.

We discuss first the rate of direct filling. Consider the following simulation. First, place some number of UCN, say N, with energy E_s in the closed source. Then, at time $t = 0$, open the source and allow the UCN to travel to the storage bottle. Finally, record the times at which they arrive. Supposing that a UCN that enters the bottle never leaves, the function

$$F^{(0)}(t, E_s) = \frac{1}{N}\int_0^t \mathrm{d}t' \sum_{\{t_a\}} \delta(t' - t_a) \tag{4.8}$$

gives the fraction of UCN in the bottle at time t, where $\delta(t)$ is the Dirac delta function, and $\{t_a\}$ is the set of arrival times, which depend in some way on E_s. In general, $F^{(0)}(t, E_s) \leq 1$, because some UCN may never reach the bottle. The rate of increase of UCN in the bottle is then

$$R_f^{(0)}(t, E_s) = \frac{\mathrm{d}}{\mathrm{d}t} N F^{(0)}(t, E_s) = N \frac{\mathrm{d}}{\mathrm{d}t} F^{(0)}(t, E_s)\,, \tag{4.9}$$

which is the rate of direct filling.

To simulate refilling, N UCN of energy E_b are released from the bottle rather than from the source, and the times at which they return are recorded. The function

$$F^{(1)}(t, E_b) = \frac{1}{N} \int_0^t \mathrm{d}t' \sum_{\{t_r\}} \delta(t' - t_r) \tag{4.10}$$

gives the fraction of UCN that return at time t, where $\{t_r\}$ is the set of return times. Now, the number of UCN that leave the bottle between time t' and $t' + \mathrm{d}t'$, say, is $[n(t')/\tau_h(E_b)]\,\mathrm{d}t'$. Therefore, the rate at which these return to the bottle at time $t > t'$ is

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{n_b(t')}{\tau_h(E_b)}\,\mathrm{d}t'\,F^{(1)}(t-t', E_b)\right) = \frac{n_b(t')}{\tau_h(E_b)}\,\mathrm{d}t'\,\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t-t', E_b)\,, \tag{4.11}$$

where the derivative is evaluated at $t - t'$. To find the rate of refilling at time t, we must add up all the contributions from the earlier times t'. Thus

$$R_f^{(1)}(t, E_b) = \int_0^t \frac{n_b(t')}{\tau_h(E_b)}\,\mathrm{d}t'\,\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t-t', E_b)\,. \tag{4.12}$$

Now, as discussed previously, $\tau_h(E_b)$ should diverge as $E_b \to 0$, which means that the computation time required to simulate $\tau_h(E_b)$ should also diverge. However, the analytical prediction for $\tau_h(E_b)$—see Equation 4.6—only becomes comparable to τ_β when $E_b \sim 10^{-5}\,\mathrm{neV}$. We shall see that the relative fraction of UCN with such low energies that reach the bottle is very small. We thus make the approximation that $\tau_h(E_b) \to \tau_e(E_b)$, i.e. the right hand side of Equation 4.7.

In a complete simulation of $R_f^{(0)}(t, E_s)$ and $R_f^{(1)}(t, E_b)$, or rather of $F^{(0)}(t, E_s)$ and $F^{(1)}(t, E_b)$, the guide losses should be taken into account. However, as there are many different components in the guide section between the source and the bottle, with many different loss factors, optical potentials, roughnesses, etc. that are not known, we assume that UCN are only lost in the guides via beta decay. (See Sections 4.4 and 4.5 for a presentation of these results.)

The full differential equation governing $n_b(t)$ is then approximated as

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}t}n_b(t, E_b) &= R_f(t, E_b) - \frac{1}{\tau_e(E_b)}n_b(t, E_b) \\
&= R_f^{(0)}(t, E_s) + R_f^{(1)}(t, E_b) - \frac{1}{\tau_e(E_b)}n_b(t, E_b) \\
&= N\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t, E_s) + \int_0^t \mathrm{d}t'\,\frac{n_b(t', E_b)}{\tau_e(E_b)}\,\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t-t', E_b) - \frac{1}{\tau_e(E_b)}n_b(t, E_b)\,,
\end{aligned} \tag{4.13}$$

where the E_b and E_s are related through Equation 4.1. For known $F^{(0)}(t)$ and $F^{(1)}(t)$, this differential equation could potentially be solved analytically, for example, by Laplace transform methods. However, as $F^{(0)}(t)$ and $F^{(1)}(t)$ are simulated, numerical methods are used to solve for $n_b(t, E_b)$. (For a discussion of this procedure, refer to Appendix B.1.)

In summary then, there are four quantities that will be simulated:

Quantities to Simulate

1. $\tau_s(E_s)$: the storage time constant of the SUN-2 source.
2. $\tau_e(E_b)$: the "emptying time constant" of the storage bottle.
3. $F^{(0)}(t, E_s)$: the direct filling fraction from the source to the storage bottle.
4. $F^{(1)}(t, E_b)$: the refilling fraction from the bottle back to itself.

The last three quantities enable calculation of $n_b(t, E_b)$ for a given number of UCN, N, with energy E_s, released from the source. The result at each energy value is then rescaled by $\sqrt{E_s}\tau_s(E_s)$. (See Equation 4.4.) We discuss the results for each of these in the following sections.

4.2 The storage time constant of the UCN source

We first describe the geometry and properties of the SUN-2 source used in the simulation. The production volume in which the UCN are generated is a two-section aluminum box, which is coated on its interior with a layer of Fomblin, $V = 106\,\mathrm{neV}$ [52], over a beryllium coating, $V = 252\,\mathrm{neV}$ [3], and filled with superfluid ^{4}He, $V = 18.5\,\mathrm{neV}$ (calculated). The two sections are a long box, $500 \times 80 \times 80\,\mathrm{mm}$, and a wide box, $100 \times 100 \times 80\,\mathrm{mm}$. (See the image at right in Figure 4.1.) An exit hole, 21 mm in diameter, sits on the top surface of the wide box section. During accumulation mode operation of the source, this is covered by the UCN valve, which is a vertically-actuated, 3 mm thick, 63 mm diameter disk that is also coated with Fomblin and beryllium. During the installation of the production volume of the SUN-2 source, the Fomblin grease was manually applied as a top-layer coating to all inner surfaces. On account of its application "by-hand", its thickness and uniformity are not well-characterized. Another unknown quantity is the level of superfluid ^{4}He in the production volume, which is only estimated during preparation of the source by indirect methods. In the simulations, this level was chosen so that when the UCN valve is fully open—travel distance 15 mm

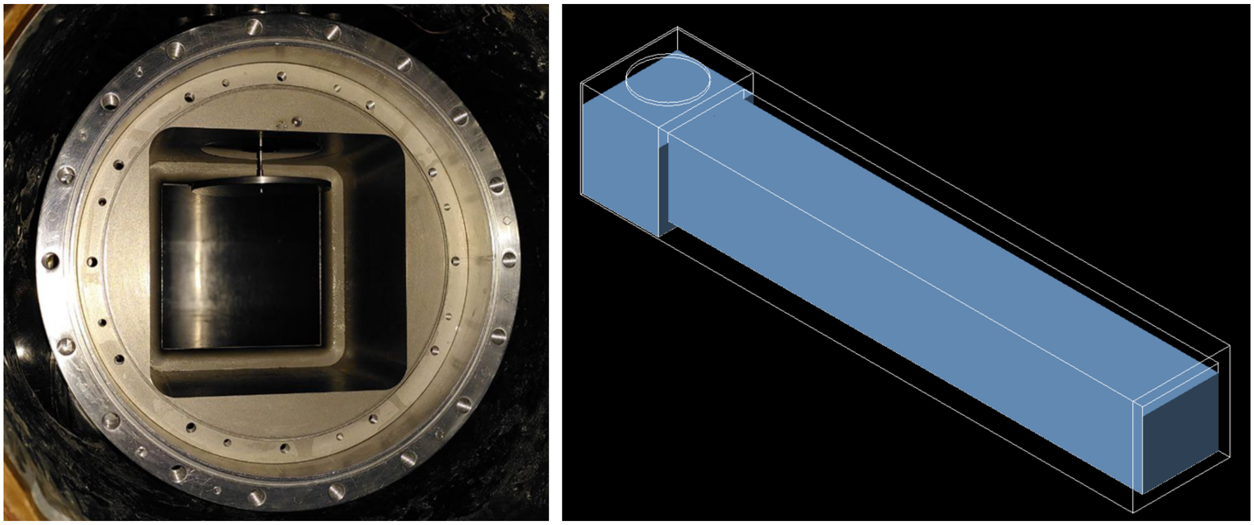

Figure 4.1: Left: Picture of the partially-disassembled inside of the converter volume, with the UCN valve disk open. Right: Image of the simulated converter volume geometry filled with superfluid helium-4, when closed, as in accumulation mode operation.

downward—the lower valve disk surface sits 1 mm above the surface of the lake. The level of the lake is likely this height or lower, because otherwise, the (warmer) disk of the UCN valve would make thermal contact with the lake when it was fully open, heating the liquid bath above its base temperature of about 0.6 K. The effects of capillary forces that will tend to generate a thin film of superfluid ^{4}He over surfaces higher than the level of the lake are neglected.

In the simulations, the surface roughness of the source walls is characterized by a single parameter p, which is the probability that a wall reflection is non-specular, or "diffuse". The probability of a specular reflection is thus $1 - p$. This approach differs from the "micro-roughness model" [75], in which perturbation theory is used to calculate the probability distribution for final angles $(\varphi_\mathrm{f}, \theta_\mathrm{f})$ given the incident angle θ_i when a surface has "Gaussian" roughness of RMS height b and correlation length w. (The incident azimuthal angle is taken as $\varphi_\mathrm{i} = 0$.) This theory, however, is only applicable in the long wavelength limit, i.e. when $b/\lambda \ll 1$. As was indicated in the previous paragraph, neither b nor w are known for the top Fomblin coating layer, so it is not possible to determine b/λ. Therefore, the single parameter p model is used in place of the micro-roughness model. The value of p was arbitrarily chosen to be 10 % to avoid long settling times and closed UCN orbits, in which the results of the UCN kinetic theory are not applicable. (See Section 4.3.2 in [3].) In the case of a diffuse reflection then, the probability distributions for the $(\varphi_\mathrm{f}, \theta_\mathrm{f})$ are uniform for the azimuthal angle, $\varphi_\mathrm{f} \sim \mathrm{Unif}\,[0, 2\pi]$, and cosine for the polar angle $\theta_\mathrm{f} \sim \cos(\theta)$. This collection of distributions is often referred to as a Lambert model because of its connection to optics, in which a perfectly diffuse, or

“Lambertian” surface, has a radiance that is independent of viewing angle. (See, for example, [126].)

We must also specify the loss factor and optical potential of the walls of the source. To guide our selection of these parameters, we recall the goal stated at the beginning of this chapter: we wish to determine the loss factors, η, of CYTOP and dPE. Briefly, this is accomplished by first estimating the initial energy spectrum of UCN in the storage bottle, $n_b(t = 0, E_b)$, and then inserting this distribution into Equation 2.18 along with the storage models developed in Chapter 2. Equation 2.18 is then fit against storage data in order to extract the loss factor, η, of the bottle coating. In carrying out this procedure, we wish to avoid producing estimates of η that are too small, that is, that promise storage properties that are “too good”. Now, roughly speaking, the larger the average energy of $n_b(t = 0, E_b)$, the smaller the implied value of η. This is due to the fact that spectra with higher average energy are associated with a greater number of wall interactions. Therefore, fits using such spectra will produce smaller values of η in order to describe a given set of storage data. Therefore, we should choose η and V of the source walls to generate a conservatively low estimate of the energy spectrum of $n_b(t, E_b)$. This is accomplished by choosing η in the source conservatively large, and V in the source conservatively small. Therefore, we model the source walls as being made of Fomblin. (The thin layer of Fomblin over the higher potential beryllium is intended to enhance the storage properties of the source by effectively presenting a larger optical potential to stored UCN than that of the top coating of Fomblin, and at the same time, lowering the effective loss factor from that of beryllium, which is presumably higher than that of Fomblin at low temperatures [100].) Since the storage properties of Fomblin at low temperature are not well-reported in the literature, we opt to use the room temperature value for its loss factor, which is in keeping with a conservative approach. For Fomblin then, $\eta = 1.8 \times 10^{-5}$ has been reported at room temperature [52].

In addition to contributions to $\tau_s(E_s)$ from beta decay and the factors discussed above, there are also contributions from the parameters chosen to describe the superfluid ^{4}He medium, and potentially from leakage through holes in the source walls. Contributions from the latter are disregarded on the grounds that the size of such leaks are not known. Regarding the former, the neutron absorption cross section of ^{4}He is essentially zero [2], and its optical potential in superfluid form, $V_m = 18.5\,\text{neV}$, may be calculated from a consideration of its density of $0.145\,\text{g}\,\text{mL}^{-1}$ [95]. In general, there is also a contribution to $\tau_s(E_s)$ from “upscattering” inside of the superfluid medium, that is, inelastic scattering of UCN and the excitations of superfluid ^{4}He that will carry the neutron out of the UCN energy regime. The problem of upscattering in superfluid ^{4}He has been studied theoretically in reference [124], in which it was found that the upscattering rate is approximately [127]

$$\frac{1}{\tau_{\rm up}} \approx A e^{-12/T} + B(T) T^7 + C(T) T^{3/2} e^{-8.6/T} \,, \tag{4.14}$$

where the temperature T is in Kelvin, and A, $B(T)$, and $C(T)$ are constants calculated using perturbation theory. In order, they represent the upscattering rates from single phonon, two-phonon, and phonon-roton processes. However, at temperatures below 0.7 K, $\tau_{\rm up}$ is already greater than 1600 s [90]. Since the superfluid ^{4}He in the production volume of SUN-2 is cooled to approximately 0.6 K, upscattering is neglected. At this point, all parameters characterizing the storage properties of the source have been specified.

The $\tau_s(E_s)$ simulation consists of generating $N = 2{,}000$ UCN per energy E_s within the source, storing them, and recording the loss times. The UCN are generated with random direction on a plane inside the superfluid ^{4}He which contains the axis of the production volume and is perpendicular to gravity. This allows comparison of E_s with E_b via Equation 4.1. A histogram of the loss times for different E_s is shown in Figure 4.2 at left. Note especially the abrupt change in slope when $E_s \gtrsim 85\,\text{neV}$, large enough for the UCN to escape the source.

The histograms plotted on a logarithmic scale at left in Figure 4.4 appear to lie on straight lines, indicating that the loss times are exponentially distributed. One way to extract the set of $\tau_s(E_s)$ from these data would then involve fitting a series of exponential functions to these histograms. However, such a procedure relies on an arbitrary choice of binning for the data, and further presupposes that the loss times are indeed exponentially distributed. We therefore first apply a modified Kolmogorov-Smirnov test [128], originally due to Lilliefors [4], to the data to see whether or not the observed loss times are exponentially distributed. In the event that they are, the maximum likelihood estimate for $\tau_s(E_s)$ is simply the mean of the loss times. Thereafter, a 1-σ confidence interval may be calculated. (For details of this process, refer to Appendix B.2.)

The values of $\tau_s(E_s)$ estimated in this way from the simulations are listed in Table B.1 in Appendix B.2, and are plotted as the black points in Figure 4.2 at right. These data will be used in the expressions developed in the previous section for calculations of the initial spectrum in the storage bottle. (See Section 4.6.) However, several analytical cases are considered and compared to the simulated data as a "sanity check". The first of these is plotted as the blue curve in Figure 4.2, corresponding to a calculation of $\tau_s(E_s)$ using Equation 2.16, and assuming that the source is filled completely with superfluid ^{4}He. The superfluid medium has the effect of altering the loss probability per wall interaction according to Equation A.9 in Appendix A.1.

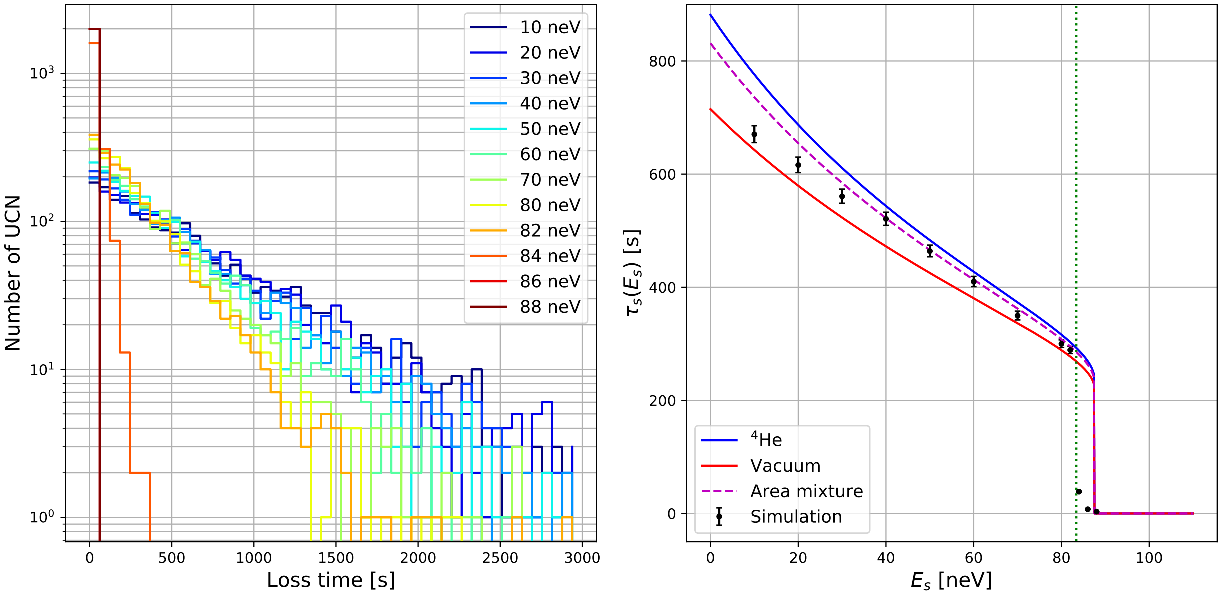

Figure 4.2: Left: A histogram of the loss times in the simulation of $\tau_s(E_s)$, shown on a logarithmic scale. $N = 2{,}000$ per energy. Note especially the abrupt change in slope when $E_s \gtrsim 85\,\mathrm{neV}$, large enough for the UCN to escape the source. Right: (black) simulated $\tau_s(E_s)$; (blue) an analytical calculation of $\tau_s(E_s)$ assuming the source is filled completely with superfluid ^{4}He; (red) another calculation assuming the source is filled with vacuum; (dashed purple) a mixture of the two previous curves based on the fractions of wall areas in contact with helium and vacuum. The dotted green line marks $E_s = V_w - V_m - mgH$, where V_w is the potential of the walls, V_m is the potential of the superfluid medium, and $H = 40\,\mathrm{mm}$ is the height from the plane of UCN production in the simulation to the bottom of the source.

Unfortunately, the associated modification of the loss factor η_w therein is not implemented in the simulation tool, and so a value of 1.8×10^{-5}—as in the simulation—was chosen for this analytical calculation. One immediately obvious feature of the blue curve is that its high energy "cutoff" is greater than that of the simulated data. This offset in E_s is a gravitational effect, resulting from UCN being generated on the axis of the source, at a height $H = 40\,\text{mm}$ from its bottom. When these UCN interact with the bottom wall of the source, they have energy $E_s + mgH$, and are therefore more likely to be lost. The horizontal offset between the blue curve and the simulated data is indicated by the dashed green line at right in Figure 4.2, whose location is at $E_s = V_w - V_m - mgH$, where V_w is the potential of the walls, and V_m is the potential of the superfluid medium. The red curve is the result of another analytical calculation, now assuming that the source is "filled" with vacuum. It has been shifted horizontally by $-V_m$ to account for the fact that a UCN of energy E_s with respect to the superfluid medium has energy $E_s + V_m$ with respect to vacuum. The red curve is lower than the blue because the energy of the UCN in the former situation is larger relative to the optical potential of the wall, and is therefore more likely to be lost per bounce. Finally, the dashed purple curve is a mixture of the red and blue curves. The mixing fraction is determined by the surface areas of the walls in contact with the vacuum and superfluid ^{4}He. If the level of the helium bath were on the axis of the production volume, the areas would be equal, and the mixing fraction would be 50 %. In the simulations, however, the level of the liquid helium is raised by 31 mm above this axis, and the bath therefore touches roughly 70 % of the available surface area of the walls.

4.3 The emptying time constant of the bottle

The next quantity to simulate is $\tau_e(E_b)$, the emptying time constant of the bottle. The UCN storage measurements performed with CYTOP and dPE coatings, the results of which are presented in the main body of the present work in Chapter 3, used bottles with a roughly cylindrical design. A small number of storage measurements were also carried out with a "diamond-like carbon" (DLC) coating using a bottle of a different design. (See (11) in Figure 3.1.) These measurements do not constitute a "complete" data set however, and are therefore described in Appendix C. Nevertheless, simulations of UCN emptying from this bottle geometry—a box of $500 \times 80 \times 80\,\text{mm}$—were also carried out, and are reported on in this section. Figure 4.3 shows cross sections of both bottle geometries when they are open. In both cases, the circular face of the plug on the inside of the storage vessel is 100 mm from the front inner surface of the bottle when the valve is open. Due to small variations in the machining of the large-bore thread on the entrance plate used to install

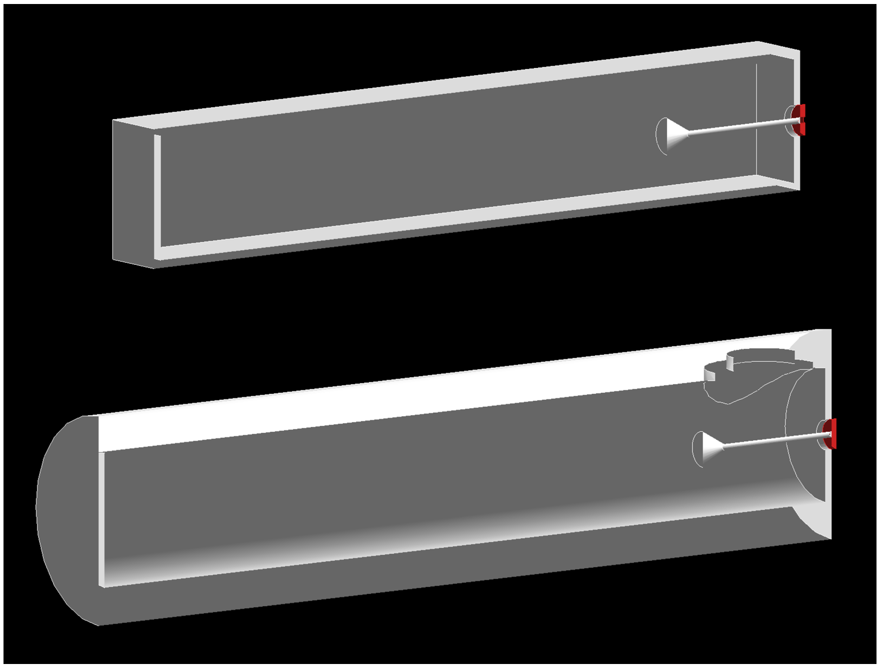

Figure 4.3: Top: Cross section view of the box-style bottle geometry in the simulation. The red annulus is a detector volume used to monitor outgoing UCN. Bottom: cross section view of the cylindrical-style bottle geometry with detector in red. Orientations of this bottle in the actual simulations is not as shown here; for the CYTOP measurements, the bottle is rotated by 180° about its long axis; for the dPE measurements, the bottle is rotated by 45° about its axis.

the bottle—see Section 3.1—the circular side port of the CYTOP and dPE bottles, whose face is parallel to the bottle axis, was oriented differently with respect to the plane of the floor for the two bottles. For the former, the outer normal to the side port face points nearly straight down, while for the latter, this normal makes an angle of roughly 45° to the vertical. In the case of the box-style DLC bottle, the sides of the bottle lie in roughly horizontal and vertical planes. Simulations are therefore performed for three different bottles in the orientations just described.

As discussed previously, the simulation of the emptying of the storage bottle does not take into account loss due to interactions with the bottle walls, because these are not known. The only two mechanisms for loss are beta decay and exit through the front of the bottle. As with simulations of the source, the probability for a non-specular reflection is arbitrarily chosen to be 10 %. The simulations consist of generating N UCN with kinetic energy E_b and random direction in a plane containing the bottle axis and perpendicular to the direction of gravity, and recording the detection times, that is, the times when the UCN exit the bottle. The histograms of the detection times for the various bottle orientations in these simulations are shown in Figure 4.4.

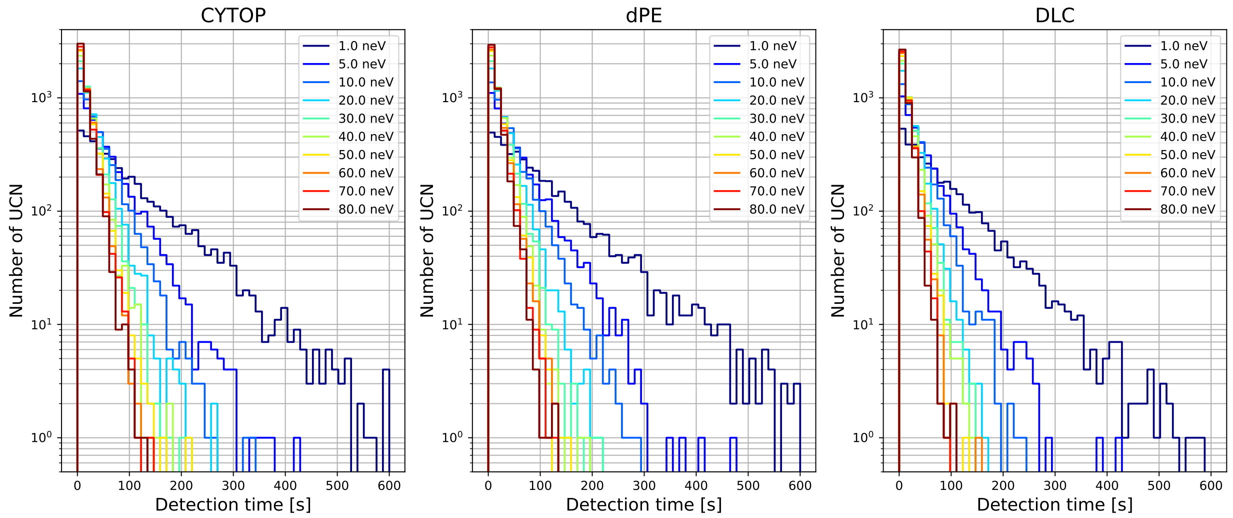

Figure 4.4: From left to right, histograms of the detection times for UCN of different energies E_b exiting the cylindrical-style bottle with the orientation of the CYTOP bottle, the same bottle with the orientation of the dPE bottle, and the box-style bottle with the orientation of the DLC bottle.

As in the previous section, the extraction of $\tau_e(E_b)$ from simulated data is determined via the method described in Appendix B.2. The results are given in Table B.2 therein, and are also shown in Figure 4.5. At left is a plot of $\tau_e(E_b)$ for the different bottles and orientations: CYTOP data are shown in blue, dPE in orange, and DLC in green. The CYTOP and dPE results are rather close, indicating that the orientation of the bottle in this approximation has a small effect on the emptying time constant. The dashed lines show the analytical predictions for $\tau_e(E_b)$ from Equation 4.7, where the emptying time constant of the hole, $\tau_h(E_b)$, is calculated using the surface to volume (SV) ratios $A_h/\mathcal{V}$, where A_h is the area of the annulus presented by the bottle exit, and $\mathcal{V}$ is the volume of the bottle interior taking into account the presence of the plug. These ratios are approximately $0.082\,\mathrm{m}^{-1}$ for the cylindrical-style bottle (black dashed line, CYTOP and dPE) and $0.097\,\mathrm{m}^{-1}$ for the box-style bottle (purple dashed line, DLC), respectively. Immediately to the right of this plot, at the top, are the residuals of $\tau_e(E_b)$ against the dashed lines in the former, using the calculated values of $A_h/\mathcal{V}$. (The red dashed line is a guide to the eye.) For both the cylindrical-style and box-style bottles, this tends to underestimate the results of the simulation. Such disagreement could be interpreted as $A_h/\mathcal{V}$ being too large. This may be due to the fact that the effective area A_h is slightly smaller in the simulations: the detector sits at the end of a short tube of length 5 mm, and the central rod of the plug extending into the bottle casts a "shadow" over the annulus of the detector. A fit to $A_h/\mathcal{V}$, shown in the middle residual plot at right of Figure 4.5, returns $A_h/\mathcal{V} \approx 0.075\,\mathrm{m}^{-1}$ for both the cylindrical-style bottles and $0.092\,\mathrm{m}^{-1}$ for the box-style bottle, all of which are smaller than the values calculated from geometry. However, such a fit very clearly does not describe the trend of $\tau_e(E_b)$ at low energy. (See the middle right plot in Figure 4.5.)

This is due to gravitational effects on the UCN trajectories at low velocity. For moderate velocities, so that UCN trajectories are not especially curved by gravity, the average time between wall collisions is $\sim D/v$, where D is some dimension of the bottle, like its diameter or width. For low velocity, however, the trajectories are curved enough so that the average time between bounces is $\sim v/g$, where g is the acceleration due to gravity. (The transition to such trajectories is for v smaller than $\sim \sqrt{Dg}$.) Therefore, the rate of exploration of the bottle by the UCN should be proportional to the reciprocal of the time between bounces, or $\sim 1/(v/g) = g/v$. The result of extending Equation 4.7 by adding such a term, $\alpha\, g/v$, and then fitting this expression to the simulated data, can be seen in the final residual plot of Figure 4.5 at bottom right. The fit gives $\alpha = 2.2(7) \times 10^{-5}$ for the cylindrical-style bottle in the CYTOP orientation, $4.0(7) \times 10^{-5}$ for the dPE orientation, and $1.8(9) \times 10^{-5}$ for the box-style bottle in the DLC orientation. The value of α presumably contains the combined information about the relative size of the hole and the bottle inner surface, and the average fraction of time along its trajectory that the UCN spends at the height of the hole. The fact that α is larger—albeit within relatively large error bars—for the bottle in the dPE orientation, in which the side port is above the bottle axis, would seem to make sense: in the CYTOP orientation, the side port sits at the bottom of the bottle, and thus enlarges the effective volume for low velocity UCN compared to in the dPE bottle. This additional term, however, does not capture all the physics—see lower right plot of residuals in Figure 4.5—which is why the simulation results, rather than these fitting functions, are used for $\tau_e(E_b)$.

4.4 The rate of direct filling

The next quantity to determine by simulation is $R_f^{(0)}(t, E_s)$, or the rate of direct filling from the SUN-2 source to the storage bottle. (See Equation 4.9.) The simulation geometry is shown in Figure 4.6. The source is modeled as in Section 4.2, except that UCN valve is open. As stated previously, the loss factors, optical potentials, and roughnesses of the guide elements between the source and the bottle are not known. In the simulation, these are treated as "ideal", meaning that they have zero loss probability per bounce, and the only loss mechanism in the guides is via beta decay. The Lambert model is again used in place of a micro-roughness model, and the probability of a non-specular reflection was chosen arbitrarily to be 1 %. This ensures that most reflections preserve the UCN direction of momentum, that is, toward the bottle during filling. However, careful study of the effect of changing this parameter has not been made.

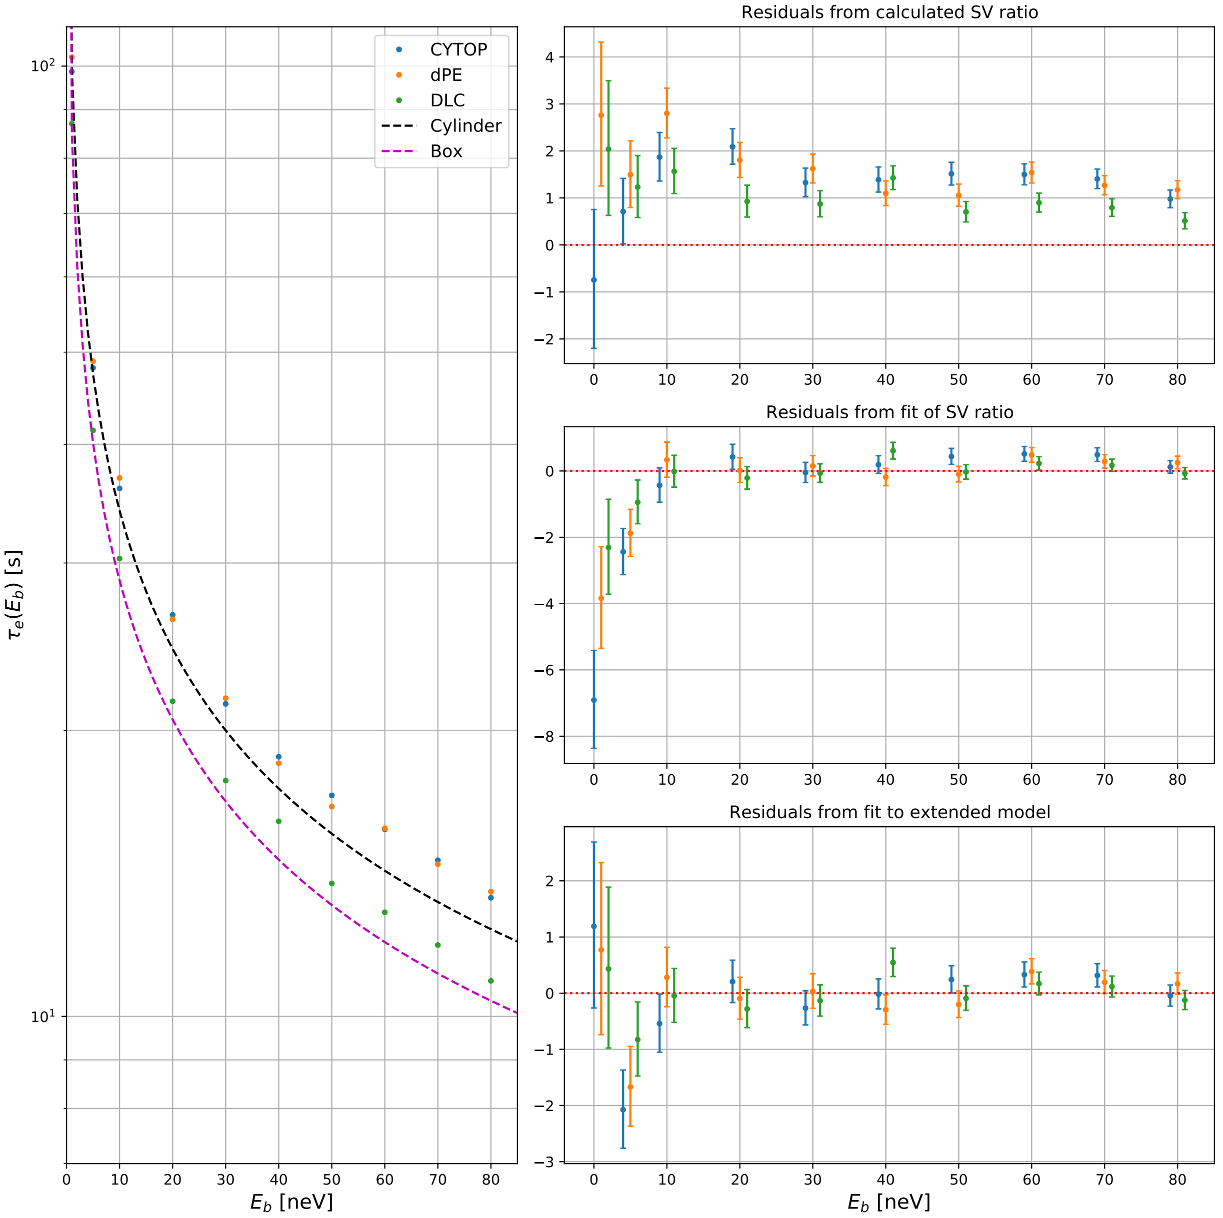

Figure 4.5: Left: plot of the $\tau_e(E_b)$ simulation results for the orientations of the CYTOP (blue points), dPE (orange points), and DLC (green points) bottles. (See plots on the right for the error bars.) The dashed lines indicate the analytical predictions for $\tau_e(E_b)$ via Equation 4.7, where the emptying time constant due to the hole, $\tau_h(E_b)$, is calculated using the surface to volume (SV) ratios $A/\mathcal{V} \approx 0.082\,\mathrm{m}^{-1}$ and $0.097\,\mathrm{m}^{-1}$ for the cylindrical-style bottle (black dashed line, CYTOP and dPE) and box-style bottle (purple dashed line, DLC), respectively. Top right: residuals from plot immediately to the left, for the calculated values of the SV ratios. Points are displaced horizontally to enable easier viewing of uncertainties; these correspond to 1-σ confidence intervals, and are very nearly symmetric. Middle right: residuals from a fit to the SV ratios, see text. Bottom right: residuals from a fit to the SV and the additional factor $\alpha g/v$, see text.

The UCN valve is displaced 15 mm vertically into the production volume, and UCN are able to escape through the 21 mm diameter hole in the front section of the source. (See bottom left inset of Figure 4.6.) The diameter of the vertical guide section after this hole widens to 50 mm, which is the diameter of all following guides. After exiting the production volume, UCN travel approximately 356 mm vertically—measured from the axis of the source to the axis of the storage bottle—toward a 90° "cross piece", which directs the UCN into the first section of horizontal guide. The vertical guide section also has a 2 mm stainless steel rod running down its center and holding up the UCN valve disk. A polished copper block sits approximately 250 mm from the axis of the vertical guide section. This block acts as a thermal screen for the source, allowing the production volume to be cooled to its base temperature of approximately 0.6 K. The block has had several modifications since its initial machining, and its approximate dimensions have been inferred from reference [70]. The remaining horizontal guides leading to the storage bottle are in the same configuration as described in Section 3.1, except that no vacuum separation foil is modeled. Due to the 4 mm diameter rod running down the center of the guide leading to the storage bottle, the bottle entrance is presented as an annulus of inner diameter 4 mm and outer diameter 20 mm. A detector in this shape, shown in red in Figure 4.6, has been placed just at the end of this guide. An additional detector is placed just outside of the 3 mm pumping hole, approximately 120 mm from the end of the last horizontal guide, visible in the bottom right inset of Figure 4.6.

In this simulation, UCN are generated as described in Section 4.2: with random direction in the superfluid ^{4}He, and on a plane containing the axis of the production volume and perpendicular to gravity. Their arrival times are then recorded by the detector at the end of the guide leading to the bottle, and used to estimate $F^{(0)}(t)$ via Equation 4.8. This function is plotted in Figure 4.7 on the left, for various energies with respect to the bottle, E_b. (For reference, the experimental filling time in accumulation mode operation of the source was 50 s.) This plot shows that UCN with high energy arrive at the bottle sooner than those with low energy. The plot at right in Figure 4.7 shows the estimates of the derivatives of $F^{(0)}(t, E_b)$, that is $\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t, E_b) = R_f^{(0)}(t, E_b)/N$. For the $N = 6{,}000$ UCN generated per simulation, the derivative of $F^{(0)}(t)$, taken numerically by finite differences, is noisy. The considerable computation time required for these simulations makes it infeasible to reduce this noise by simply increasing N. Therefore, an averaging process—described in Appendix Section B.3—was used to determine the derivatives and uncertainties of $F^{(0)}(t)$.

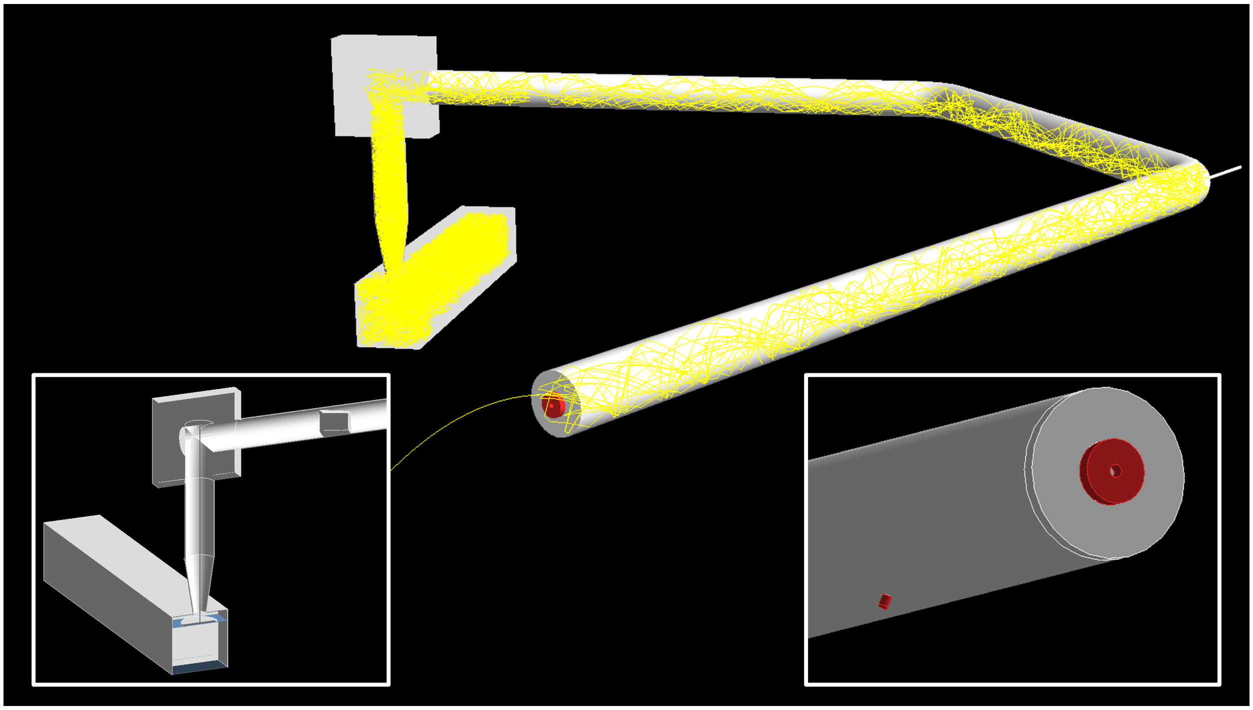

Figure 4.6: Main: simulation geometry for the calculation of $R_f^{(0)}(t, E_s)$. UCN are generated in the SUN-2 source and travel to the bottle. A single UCN trajectory is shown in yellow. An annular detector (red) sits at the bottle entrance and is used to record the times at which the UCN arrive. Bottom left inset: cut through the front of the source and guide sections, showing the UCN valve open and just above the level of the ^{4}He lake. This valve is supported by the thin steel rod leading up to the 90° cross. The copper block is shown in the horizontal guide, approximately 250 mm from the axis of the vertical guide section. Bottom right inset: view from below the horizontal guide leading to the bottle, showing an additional detector (red) placed just outside the 3 mm pumping hole in the guide.

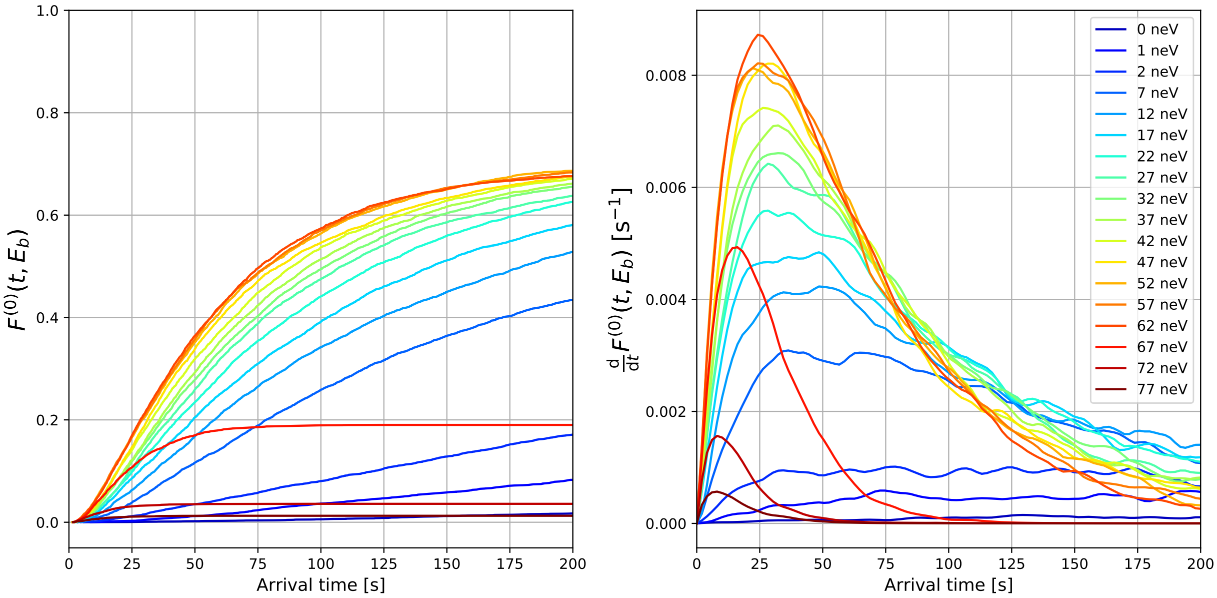

Figure 4.7: Left: simulated $F^{(0)}(t, E_b)$, where t is the arrival time, and where E_s has been translated to E_b via Equation 4.1. See legend at right for E_b values. Right: the derivative $\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t, E_b) = R_f^{(0)}(t, E_b)/N$.

4.5 The rate of refilling

The final quantity to determine via simulation is the rate of “refilling” $R_f^{(1)}(t, E_b)$, that is, the rate at which UCN return to the storage bottle after having left. (See Equation 4.12.) The simulation geometry is exactly as was described in the previous section, except that UCN are no longer generated in the source. Instead, they are produced at two points, both of which are 1 mm away from the bottle detector in the simulation. These two points are situated on either side of the 4 mm rod that runs down the guide leading to the bottle, and are centered between the inner and outer radii of the annular detector. The UCN generated there point in a random direction within 2π-steradians facing away from the bottle detector. UCN are produced from these two points, rather than from the annular shape of the bottle entrance, in order to group them by their initial values of E_b.

The $F^{(1)}(t, E_b)$ estimated—Equation 4.10—from the results of these simulations are shown at left in Figure 4.8. As with $F^{(0)}(t, E_b)$, the higher energy UCN have a greater contribution at early times. This can also be seen in the right plot of Figure 4.8, where the peak rate $\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t, E_b)$ is sharper for higher energy UCN. The derivatives $\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t, E_b)$ are estimated by the technique described in Appendix Section B.3.

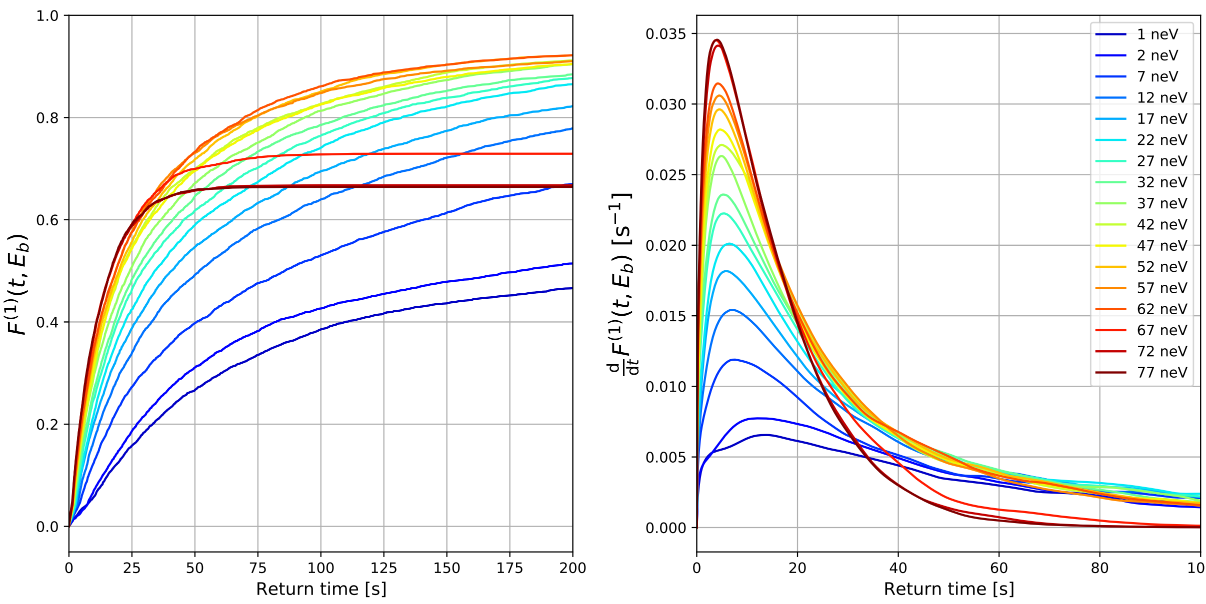

Figure 4.8: Left: simulated values of $F^{(1)}(t, E_b)$, where t is the return time. See legend at right for E_b values. Right: the derivative $\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t, E_b) = R_f^{(1)}(t, E_b)/N$.

4.6 The initial spectrum of UCN

We now apply the mathematics developed in Section 4.1 to the simulation results in order to calculate $n_b(t = 0, E_b)$, the initial spectrum of UCN in the storage bottle. This is accomplished by numerically solving Equation 4.13—see Appendix B.1—for different values of E_b using $N = 1$, where N is the number of UCN with energy E_b released from the source, and using the $\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t, E_s)$ and $\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t, E_b)$ derived from simulation. The solutions are then rescaled by $N \propto \sqrt{E_s}\tau_s(E_s)$, as discussed in Section 4.1, to give the set of $n_b(t, E_b)$ with the correct proportionality.

The energies at which UCN were simulated in the study of $\tau_s(E_s)$ and $\tau_e(E_b)$ are not the same as those used in simulations for the determination of $F^{(0)}(t, E_s)$ and $F^{(1)}(t, E_b)$. We therefore linearly interpolate between the known values of $\tau_s(E_s)$ and $\tau_e(E_b)$—and their errors bars—to determine appropriate values for use in Equation 4.13. (See Tables B.1 and B.2 in Appendix Section B.2.) Doing so allows one to generate, for example, Figure 4.9, which shows a plot of $\tau_s(E_s)$ at left, and a plot of N at right. Some care must be taken when considering energies outside the range of those simulated for $\tau_s(E_s)$ and $\tau_e(E_b)$. The lowest UCN energy in the simulations for $\tau_s(E_s)$ was $E_s = 10\,\mathrm{neV}$, but it is not necessary to know $\tau_s(E_s)$ for lower energy because such UCN do not reach the storage bottle. The highest UCN energy in the simulations for $\tau_s(E_s)$ was $E_s = 88\,\mathrm{neV}$, at which $\tau_s(E_s) = 3.4(1)\,\mathrm{s}$, and the storage time constant of the source is

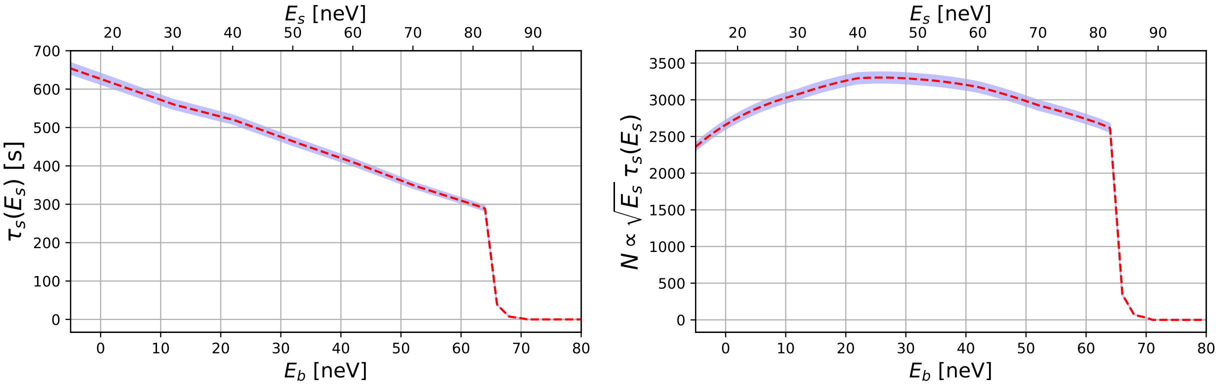

Figure 4.9: Left: Interpolated values of $\tau_s(E_s)$ shown with 1-σ error bands. Right: N from Equation 4.13, taken to be the numerical value of $\sqrt{E_s}\tau_s(E_s)$.

approximated as zero for UCN with greater energy. Thus $E_s = 88\,\text{neV}$, or $E_b \approx 70\,\text{neV}$, is the highest UCN energy. The highest UCN energy simulated for $\tau_e(E_b)$ was $E_b = 77\,\text{neV}$, which is already above this value. The lowest energy simulated for $\tau_e(E_b)$ was $E_b = 1\,\text{neV}$, at which $\tau_e(E_b) = 98.7(15)\,\text{s}$ for the CYTOP orientation, $102.2(16)\,\text{s}$ for the dPE orientation, and $87.0(15)\,\text{s}$ for the DLC orientation. Values of $\tau_e(E_b)$ at lower energy are approximated using the values of $\tau_e(E_b = 1\,\text{neV})$.

The uncertainty in the four quantities derived from simulations—$\tau_s(E_s)$, $\tau_e(E_b)$, $\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t, E_s)$, and $\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t, E_b)$—must be propagated to the uncertainty in n_b. For the sake of brevity in the present discussion, let us refer to these quantities in order as Θ_1, Θ_2, Θ_3, and Θ_4, respectively. To determine the effect that uncertainty in the Θ_i have on n_b, we treat these former quantities as independent random variables. Their mean values, θ_i, are then the simulated values of these quantities, and their variances, $\delta\theta_i$, are the estimates of their uncertainties.

The uncertainty in n_b due to Θ_1 is the easiest to treat, because this quantity only scales the solution n_b. It is calculated as follows. First, the mean value of $n_b(t, E_b)$ is found by numerical integration of Equation 4.13, using the mean values of the simulated quantities θ_2, θ_3, and θ_4. The solution is then rescaled as $n_b(t) \to \sqrt{E_s}\theta_1 n_b(t)$. The uncertainty due to Θ_1 is then simply $\delta n_{b,1}(t) = \sqrt{E_s}\delta\theta_1 n_b(t)$.

The uncertainty in n_b due to Θ_2 is more complicated to treat, because the solution n_b depends on the value of Θ_2. It is estimated here as follows. First, as in the last paragraph, the mean value of $n_b(t, E_b)$ is found by numerical integration of Equation 4.13. Then, this process is repeated multiple times, but using randomly-generated values of Θ_2, in which Θ_2 is normally distributed with mean θ_2 and standard deviation $\delta\theta_2$. Finally, the standard deviation of the difference between these "random" $n_b(t, E_b)$ and the mean

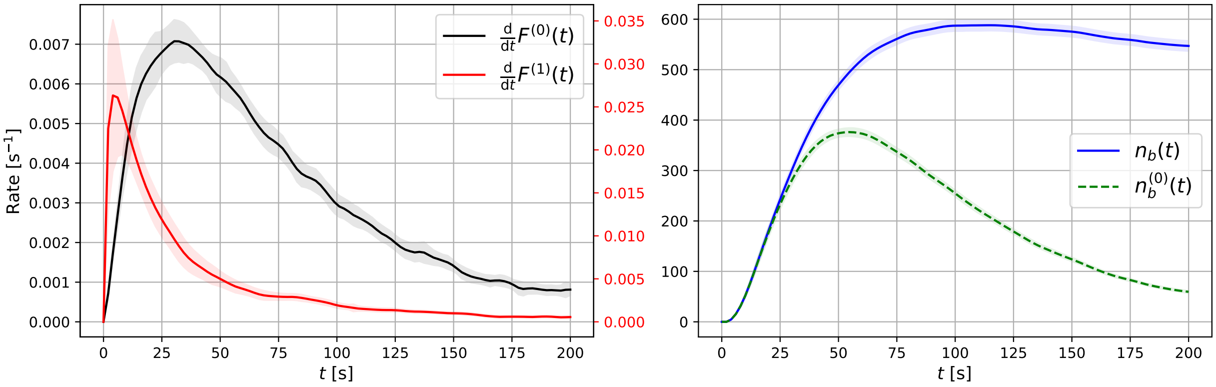

Figure 4.10: Left: Estimates of $\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t)$ and $\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t)$ and their uncertainties for the storage bottle in the CYTOP orientation, and $E_b = 37\,\mathrm{neV}$. Right: $n_b(t)$ and $n_b^{(0)}(t)$ for $E_b = 37\,\mathrm{neV}$. Uncertainties are calculated using $N = 100$ randomly generated $n_b(t)$ and $n_b^{(0)}(t)$ per Θ_i. (See text.) In this example, the rate of direct filling reaches its peak around 30 s, which is roughly where $n_b(t)$ and $n_b^{(0)}(t)$ begin to differ. After this time, the peak refilling rate occurs within roughly 5 s, so that, for no wall loss in the guides, the two solutions quickly separate.

$n_b(t, E_b)$ is taken as the estimate of the uncertainty $\delta n_{b,2}(t)$ for a given time t.

A similar procedure is used to calculate the uncertainties due to Θ_3 and Θ_4. However, now, for example, $\Theta_3 = \Theta_3(t)$ depends on time, so that rather than generating a single random value, each calculation of the "random" $n_b(t, E_b)$ uses the vector of $\Theta_3(t)$, whose points are normal random variables with mean $\theta_3(t)$ and standard deviation $\delta\theta_3(t)$. The uncertainty $\delta n_{b,3}(t)$ is then calculated in the same manner as $\delta n_{b,2}(t)$. (The uncertainty in $\Theta_4(t)$ is calculated just as for $\Theta_3(t)$.) Finally, the uncertainties are added in quadrature to arrive at an estimate of the total uncertainty for $n_b(t, E_b)$, that is

$$\delta n_b(t) \approx \sqrt{[\delta n_{b,1}(t)]^2 + [\delta n_{b,2}(t)]^2 + [\delta n_{b,3}(t)]^2 + [\delta n_{b,4}(t)]^2}\,. \tag{4.15}$$

An example of the results of this process for the storage bottle in the CYTOP orientation at a single energy, $E_b = 37\,\mathrm{neV}$, are shown in Figure 4.10. The left plot shows the estimates of $\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t)$ and $\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t)$ with their relatively large uncertainties. The right plot shows $n_b(t)$ and $n_b^{(0)}(t)$, where this latter quantity is the solution of Equation 4.13 for $\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t) = 0$, that is, for no "refilling". The results of calculating these quantities for all E_b, and for the CYTOP orientation of the storage bottle, is shown in Figure 4.11 as a function of the filling time. Both $n_b(t, E_b)$ and $n_b^{(0)}(t, E_b)$ are biased toward high energy at early filling times, and low energy at long filling times. This is because high energy UCN arrive at the bottle sooner than low energy UCN, and because the latter are stored more efficiently than the former. The lower half of Figure 4.11 shows 1-σ

uncertainty bands for a selection of filling times. An examination of these error bands for $n^{(0)}(t)$ reveals that the uncertainty in the spectrum decreases with filling time. Roughly speaking, this is because the relatively large uncertainty in $\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t)$ is "integrated out" with time. More precisely, when $\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t) = 0$ in Equation 4.13, $n^{(0)}(t)$ may be found using the Green's function $G(t, t') = u(t - t') \exp\left[-(t - t')/\tau_e\right]$, where $u(t - t')$ is the Heaviside step function. That is

$$\begin{aligned}
\frac{1}{N}\, n^{(0)}(t) &= \int_{-\infty}^{\infty} \mathrm{d}t'\, G(t, t') \frac{\mathrm{d}}{\mathrm{d}t'} F^{(0)}(t') \\
&= \int_{-\infty}^{t} \mathrm{d}t' \exp\left[-(t - t')/\tau_e\right] \frac{\mathrm{d}}{\mathrm{d}t'} F^{(0)}(t') \\
&= \Big(F^{(0)}(t) \exp\left[-(t - t')/\tau_e\right]\Big)_{-\infty}^{t} - \int_{-\infty}^{t} \mathrm{d}t' \exp\left[-(t - t')/\tau_e\right] \frac{F^{(0)}(t')}{\tau_e} \\
&= F^{(0)}(t) - \int_{-\infty}^{t} \mathrm{d}t' \exp\left[-(t - t')/\tau_e\right] \frac{F^{(0)}(t')}{\tau_e}\,,
\end{aligned} \tag{4.16}$$

where integration by parts is performed going to the third line, and the fact that $F^{(0)}(t \to -\infty) = 0$ is used going to the last line. This result is simply the statement that $n_b^{(0)}(t)$ is equal to everything that has entered the bottle, $NF^{(0)}(t)$, less everything that has left, which has instantaneous outgoing rate $NF^{(0)}(t)/\tau_e$.

The inclusion of a non-zero $\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t)$ in Equation 4.13 has two pronounced effects on the time-evolution of the spectrum of UCN in the storage bottle. Firstly, the total number of UCN in the storage bottle, which is proportional to the area under a given spectrum, decays much slower for $n_b(t)$ than for $n_b^{(0)}(t)$. This is because, in the case of the former, there is a non-zero probability for a UCN that exits the bottle to return at a later time. Secondly, the total number of UCN in the storage bottle reaches its maximum around 50 s for $n_b^{(0)}(t)$, and around 100 s for $n_b(t)$. As described in Section 3.1, the optimal experimental filling time was found to be approximately $t_{\text{fill}} = 50\,\text{s}$. (See Figure 3.4.) This agrees rather well with the behavior of $n_b^{(0)}(t)$, and indicates that, unsurprisingly, guide losses play an important role in the determination of the initial UCN energy spectrum in the storage bottle. Unfortunately, however, potential sources of loss in the guides have not been characterized experimentally.

Ultimately, we will fit experimental storage data to Equation 2.18, using the simulated UCN energy spectrum, in order to estimate the loss factors, η, of CYTOP and dPE. As discussed previously, we would like such estimates to be conservative. We will therefore use $n_b^{(0)}(E_b)$, rather than $n_b(E_b)$, as the estimate of the initial UCN energy spectrum in the storage bottle. This is due to the fact that spectra with higher average energy are associated with a greater number of wall interactions. Therefore, fits using such spectra

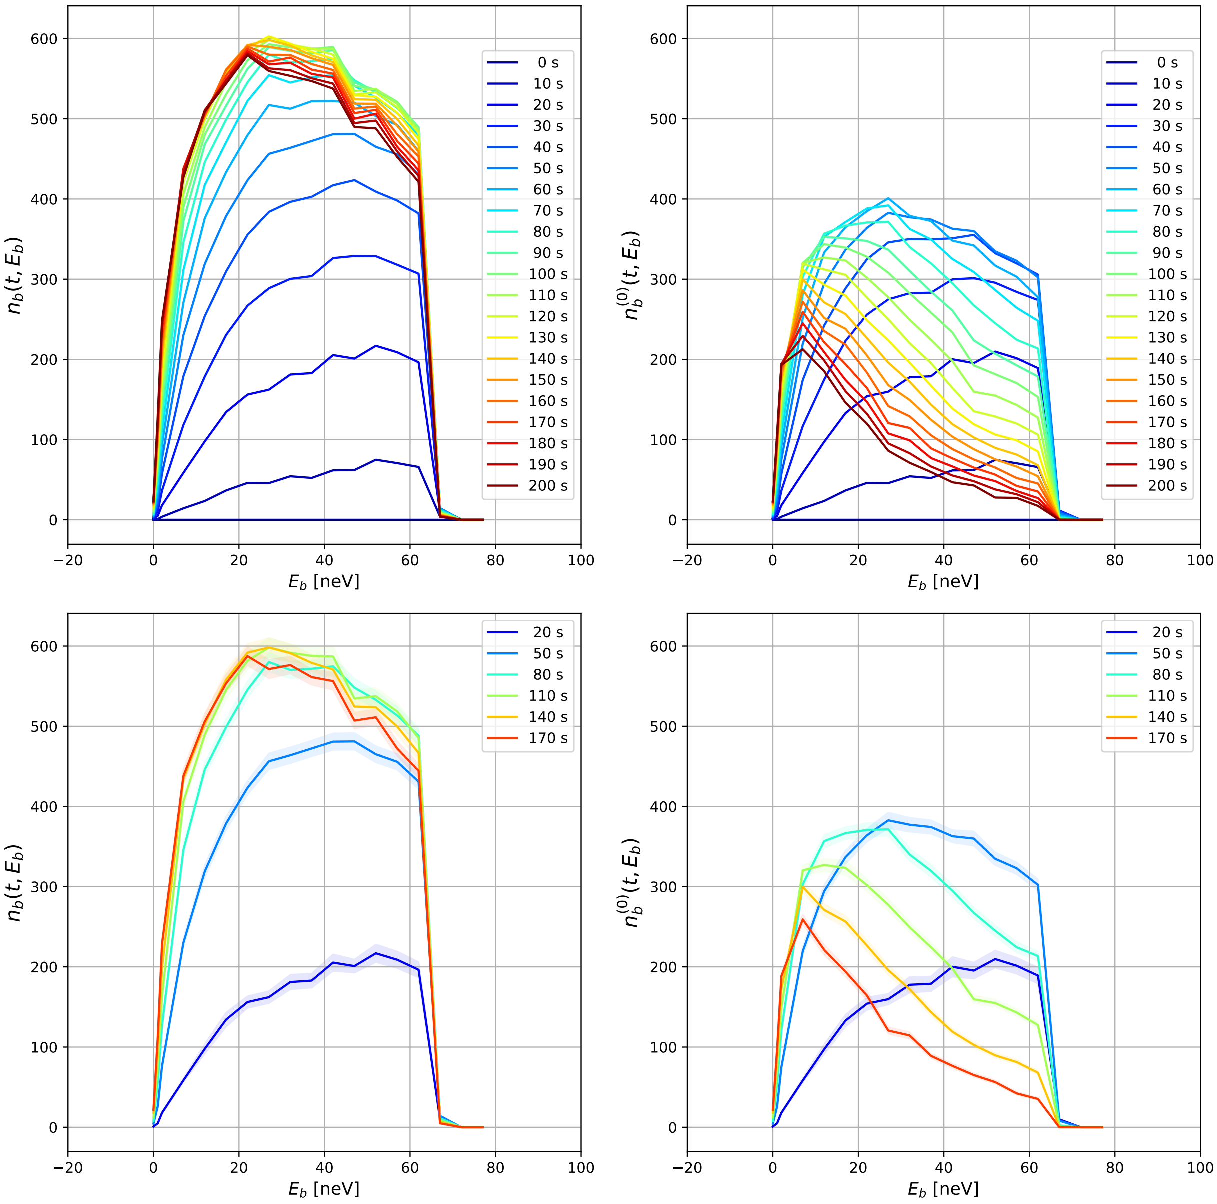

Figure 4.11: Upper: $n_b(t, E_b)$ and $n_b^{(0)}(t, E_b)$ for wide range of filling times t. $n_b^{(0)}(t, E_b)$ is the solution to Equation 4.13 in the case that $R^{(1)}(t) = 0$, i.e. no refilling occurs. Lower: Selected filling times showing the 1-σ uncertainty bands on $n_b(t, E_b)$ and $n_b^{(0)}(t, E_b)$.

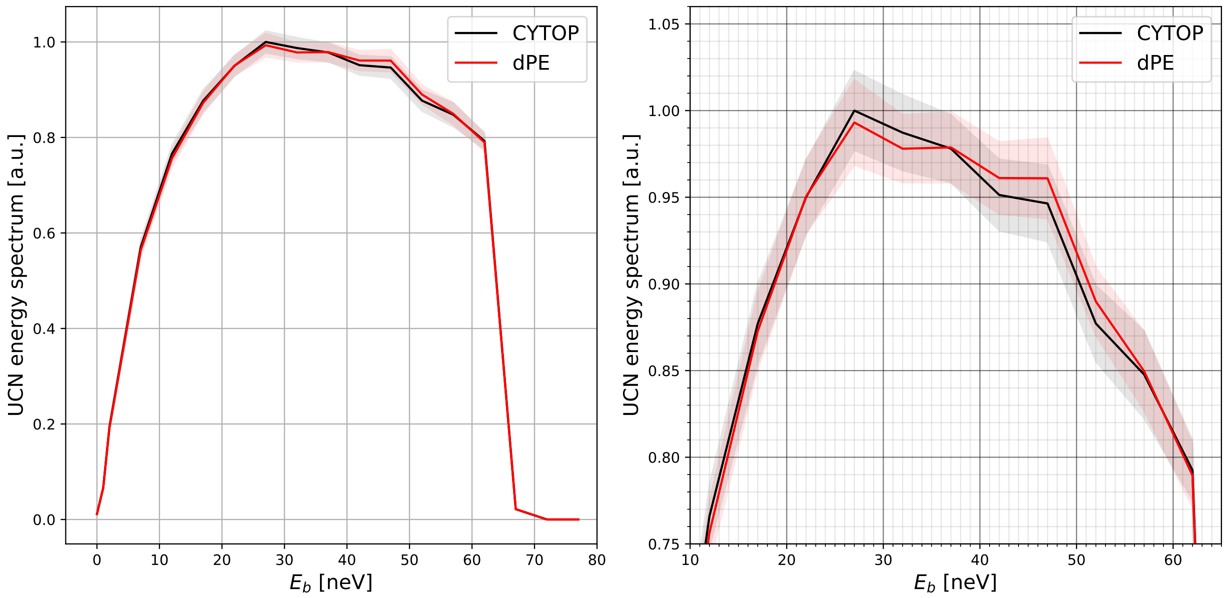

Figure 4.12: The estimates at $t_{\text{fill}} = 50\,\text{s}$ of the initial UCN energy spectra for the cylindrical-style bottles in the CYTOP and dPE orientations. Left: View of the entire energy range of the spectra. Right: Enlarged view where the uncertainty is greatest.

will produce smaller values of η in order to describe a given set of storage data.

Figure 4.12 shows plots of the estimates at $t_{\text{fill}} = 50\,\text{s}$, the experimental filling time, of the initial UCN energy spectra in the storage bottles in the CYTOP and dPE orientation. The plot at left covers the entire energy range of the spectra, while that at right provides an enlarged view in the region where the uncertainty is greatest. The spectral values at any given energy differ by at most about 3 %, while the average energies differ by less than 0.3 %: $\langle E_b \rangle \approx 34.8\,\text{neV}$ for the CYTOP orientation, and $\langle E_b \rangle \approx 34.9\,\text{neV}$ for the dPE orientation. This discrepancy could be due to the fact that is slightly more difficult for UCN to escape the bottle in the CYTOP orientation—see the discussion at the end of Section 4.3—but the difference is small.

Chapter 5

Loss factors

In this chapter, estimates of the loss factors, η, of CYTOP and dPE are made from the experimental storage data presented in Chapter 3 using the expressions developed in Chapter 2 and the results of the simulations described in Chapter 4. The storage data used for these estimates are only from measurements in which the SUN-2 source was operated in accumulation mode. (See Section 4.1 for the reasoning behind this choice.) This analysis results in estimates of $\eta = 2.7(2) \times 10^{-5}$ at 11.7 K and $\eta = 1.1(1) \times 10^{-4}$ at 295.8 K for the loss factor of CYTOP, and estimates of $\eta = 3.1(1) \times 10^{-5}$ at 11.9 K and $\eta = 2.6(1) \times 10^{-4}$ at 295.1 K for the loss factor of dPE.

The first section of this chapter gives the details of the models used to extract estimates of η from the storage data. (In all models, gravitational effects are not included; see Appendix D for a justification of this approximation.) The following sections then present the results of fits using these models to the storage data, in addition to the method used to make the final estimates of the loss factors.

5.1 Overview

The basic expression used to fit the experimental storage curves $N(t)$, that is, the number of UCN measured after storage for time t—c.f. Equation 2.18—is

$$N_0 \int \mathrm{d}E_b \, n_b(0, E_b) e^{-t/\tau(E_b)} \,, \tag{5.1}$$

where E_b is the UCN energy with respect to the storage bottle axis. The initial UCN energy spectrum in the bottle, $n_b(0, E_b)$, is derived from simulations. (See Chapter 4.) It is normalized to unity, so that the fit parameter N_0 corresponds to the total number of UCN in the storage bottle when it is closed. We will use several models for the energy-dependent storage time $\tau(E_b)$, containing different fit parameters, in the following analysis. (In all models, the effects of gravity are neglected; see Appendix D for a justification of this approximation.)

The simplest expression for $\tau(E_b)$ is that of the "single coating" (SC) model, in which the only sources of UCN loss are free neutron decay, with lifetime $\tau_\beta = 879.4(16)\,\mathrm{s}$ [6], and wall interactions characterized by the loss probability $\bar{\mu}(E_b)$. (See Equation 2.14.) For the present case, this is (c.f. Equation 2.16)

$$\frac{1}{\tau(E_b)} = \frac{A}{4\mathcal{V}} v_b \bar{\mu}(E_b) + \frac{1}{\tau_\beta}\,, \tag{5.2}$$

where $A = 1839\,\mathrm{cm}^2$ is the inner surface area of the closed bottle, $\mathcal{V} = 3814\,\mathrm{cm}^3$ is its volume, and $v_b = \sqrt{2E_b/m}$ is the UCN velocity. In order to calculate the loss probability, $\bar{\mu}(E_b)$, one must specify the loss factor η and the optical potential V. The former quantity is treated as a fit parameter, while the latter is treated as an input. Thus, the SC model results in there being two fit parameters in Equation 5.1: N_0 and η. (See Table 5.1.) We make one additional comment about the use of $\bar{\mu}(E_b)$ in this model. Firstly, for CYTOP, $V = 115.2(2)\,\mathrm{neV}$ was measured by neutron reflectometry—see Section 3.2—and for dPE, V is estimated to be $210(10)\,\mathrm{neV}$, which is consistent with values reported in the literature [59, 60]. (See Section 3.3.) Secondly, the simulation results described in the previous chapter estimate the high energy cutoff of the UCN energy spectrum, $n_b(0, E_b)$, to be approximately $70\,\mathrm{neV}$, well below V for both CYTOP and dPE. Therefore, it is only necessary to use the analytical expression for $\bar{\mu}(E_b)$ of Equation 2.14, valid for $E_b < V$, when using Equation 5.2 in Equation 5.1 to fit experimental storage data.

One important extension to the SC model that was not discussed in Chapter 2 includes UCN loss due to the presence of "residual" gases in the storage bottle. (Such gases are unwanted contaminants that remain in the experiment despite efforts to pump them away.) For a UCN of velocity v moving through a single gas species of number density ρ_i, the loss probability per unit time is $\rho_i \sigma_i v$, where σ_i is the loss cross section due to UCN absorption by and upscattering from the gas [129]. In real experiments, several gas species are present, so that the quantity of interest is $X = \Sigma_i \rho_i \sigma_i v$, where the summation is performed over the different gas species. Despite the presence of v, X is energy-independent, because $\sigma \propto 1/v$ in general [61]. The total UCN loss rate is then given by adding X to the SC model to obtain

$$\frac{1}{\tau(E_b)} = \frac{A}{4\mathcal{V}} v_b \bar{\mu}(E_b) + \frac{1}{\tau_\beta} + X\,, \tag{5.3}$$

to which we will refer as the "single coating X" (SCX) model. In comparison to the SC model, the SCX model has one additional unknown quantity, so that its use in Equation 5.1 results in three fit parameters:

N_0, η, and X. (See Table 5.1.)

Another extension to the SC model that was discussed at length in Section 2.4 is the "weak patch" model. Briefly, this model supposes that the majority of the bottle interior is coated with a material possessing optical potential V_w and loss factor η_w, and that a small fraction of the area, f, consists of "weak" patches, having optical potential $V_p < V_w$ and loss factor η_p, present due to possible contamination of, or defects in, the main coating. The UCN loss rate may then be expressed as (c.f. Equation 2.22)

$$\frac{1}{\tau(E_b)} = f\frac{A}{4\mathcal{V}}v_b\bar{\mu}_p(E_b) + (1-f)\frac{A}{4\mathcal{V}}v_b\bar{\mu}_w(E_b) + \frac{1}{\tau_\beta}\,. \tag{5.4}$$

In $\bar{\mu}_w(E_b)$, V_w is taken to be either the optical potential of CYTOP or dPE, and there are four remaining parameters: V_p, η_p, f and η_w. As discussed in Section 2.4, in the absence of direct measurement of, for example, a possible contaminant, there are two reasonable guesses for the first of these parameters, V_p, which correspond to two ways in which the storage bottle surface can depart from ideal. The first choice is the limiting case of the weak patch model in which $V_p = 0$. This value would be appropriate for possible holes in the bottle surface, as might be present, for example, near the bottle valve. This choice of V_p also has the effect of removing the parameter η_p from Equation 5.4, because $\bar{\mu}_p(E_b) = 1$ for a hole. That is

$$\frac{1}{\tau(E_b)} = f\frac{A}{4\mathcal{V}}v_b + (1-f)\frac{A}{4\mathcal{V}}v_b\bar{\mu}_w(E_b) + \frac{1}{\tau_\beta}\,. \tag{5.5}$$

We shall refer to this case as the "weak patch hole" (WPH) model, the use of which in Equation 5.1 results in three fit parameters: N_0, η_w, and f. Because there is only one loss factor, we will omit the w subscript, and refer to η_w as simply η. (See Table 5.1.)

The other possibility for the weak patch model is to set V_p equal to the optical potential of aluminum, 54 neV [3], from which the main body of the storage bottle is made, and which would be exposed by any gaps or cracks in the coating. We will refer to this case as the "weak patch aluminum" (WPA) model. Unlike the WPH model, the loss factor, η_p, is still present in Equation 5.4 for the WPA model. However, when fitting storage data to Equation 5.1 using Equation 5.4, the fit is not particularly sensitive to η_p. This can be understood for two reasons that we discuss presently.

First, when $E_b < V_p$, the loss probability $\bar{\mu}_p(E_b)$ is on the order of η_p. Likewise, when $E_b < V_w$, the loss probability $\bar{\mu}_w(E_b)$ is on the order of η_w. However, $\bar{\mu}_p(E_b)$ appears in Equation 5.4 multiplied by the factor

f, and one expects that $f\bar{\mu}_p(E_b) \ll \bar{\mu}_w(E_b)$ when $E_b < V_p$. This is expected to be true based on the observed behavior of the fits, and because of one assumption. It will be shown that values of $f \sim 10^{-4}$, which correspond to areas of roughly $\sim 10\,\mathrm{mm}^2$, and values of $\eta_w \sim 10^{-4}$ or smaller are obtained from fits using the WPA model. Therefore, the relationship $f\bar{\mu}_p(E_b) \ll \bar{\mu}_w(E_b)$ follows from the assumption that η_p is "reasonably" small. (The loss factor of aluminum is given as 2.25×10^{-5} in [3], so that even if η_p is roughly three orders of magnitude larger, the relationship should still be valid.) We thus make the approximation that $f\bar{\mu}_p(E_b) \approx 0$ when $E_b < V_p$, so that Equation 5.4 contains only the terms in $\bar{\mu}_w(E_b)$ and τ_β in this energy region.

The second reason that a fit using the WPA model is not sensitive to η_p is that when $E_b > V_p$, the loss probability $\bar{\mu}_p(E_b)$ abruptly increases by several orders of magnitude. (See, for example, at right in Figure 2.3.) We therefore make the approximation that $\bar{\mu}_p(E_b) = 1$ for $E_b > V_p$. Finally then, the explicit statement of Equation 5.4 for the WPA model is

$$\frac{1}{\tau(E_b)} = \begin{cases} (1-f)\frac{A}{4V}v_b\bar{\mu}_w(E_b) + \frac{1}{\tau_\beta}\,, & 0 \le E_b < V_p \\ f\frac{A}{4V}v_b + (1-f)\frac{A}{4V}v_b\bar{\mu}_w(E_b) + \frac{1}{\tau_\beta}\,, & V_p \le E_b \le V_w\,. \end{cases} \tag{5.6}$$

Thus, when the WPA model is used in Equation 5.1, there are three fit parameters: N_0, η_w, and f. When reporting the latter value for both the WPA and WPH models, we replace the area fraction f by the area of the patch itself, that is $a_p = fA$, where $A = 1839\,\mathrm{cm}^2$ is the area of the inner bottle surface. Also, as with the WPH model, because there is only one loss factor, we will omit the w subscript, and refer to η_w as simply η. (See Table 5.1.)

A summary of the parameters of these four models for $\tau(E_b)$ is given in Table 5.1. The "Fit parameters" column lists the parameters discussed above, in addition to the parameter τ_d in parenthesis. Recall that for the CYTOP and dPE measurements, a systematic decrease in $N(t)$ with absolute measurement time t_m was observed. This decrease was typically associated with the measurement being at low temperature, and was characterized by the time constant, τ_d, which was on the order of 100 hours. (See Chapter 3.) In the case that a nonzero value of τ_d is obtained from fits of $N(t)$ to the empirical models presented in Chapter 3—see Equations 3.2 and 3.5—the storage curve used to fit against Equation 5.1 is $N(t)$ divided by the factor $\exp(-t_m/\tau_d)$. In such case, τ_d is included in the list of fit parameters.

For each model, optimal fit parameters were determined by numerically minimizing χ^2, that is, the sum of

Table 5.1: Summary table of the models for $\tau(E_b)$. (See the text.)

Model	Abbreviation	Fit parameters	Equation
Single coating	SC	η, N_0, (τ_d)	5.2
Single coating X	SCX	η, N_0, X, (τ_d)	5.3
Weak patch holes	WPH	η, N_0, a_p, (τ_d)	5.5
Weak patch aluminum	WPA	η, N_0, a_p, (τ_d)	5.6, $V_p = 54\,\mathrm{neV}$

squares of the weighted residuals. This was performed with the `curve_fit` function in SciPy [109], using the MINPACK implementation [117] of the "Levenberg-Marquardt" algorithm[118, 119]. The parameter uncertainties from the statistics of the fitting procedure were determined using numerical estimates of the Hessian, H, of χ^2, evaluated at the optimal fit values. (The inverse of the covariance matrix $(\Sigma^{-1})_{ij}$ is estimated by $\frac{1}{2}H = \frac{1}{2}\frac{\partial^2\chi^2}{\partial\theta_i\partial\theta_j}$; see, for example, Equation 40.12 in [6].)

The uncertainties on the optimal fit parameters due to both the uncertainty in $n_b(0, E_b)$, and due to the uncertainties in the optical potentials of CYTOP and dPE, were estimated by Monte Carlo methods. For the former, 100 random spectra were generated, where the value of the spectrum at a given E_b is a normal random variable with mean and variance given as indicated in Figure 4.12. For each random spectrum, a fit was made as described in the previous paragraph. The standard deviations of the optimal fit parameters were then taken as a measure of the error due to the uncertainty in $n_b(0, E_b)$. A similar procedure was followed for the error due to the uncertainty in the optical potentials. The error bars on the optimal parameter values due to the statistics of the fitting procedure, due to the uncertainty in $n_b(0, E_b)$, and due to the uncertainty in the optical potential were finally added in quadrature. (These are the values given in Tables 5.3 and 5.5.) The dominant source of uncertainty for all parameters is due to statistics, with the combined contribution from the latter sources being no more than 50 % of the error bars. When giving the value of the reduced chi-squared statistic for these fits, χ^2_r, as in Tables 5.3 and 5.5, the number of degrees of freedom is $\nu - p$, where ν is the size of the data set, and p is the number of parameters listed in Table 5.1.

5.2 CYTOP

For the CYTOP data, there are four storage measurements in which the SUN-2 source was operated in accumulation mode, and whose data may therefore be fit against Equation 5.1 using $n_b(0, E_b)$ derived from the simulations described in the previous chapter. These are measurements 13, 14, 16, and 18. (See Table

Table 5.2: Summary table of CYTOP storage measurements in which the SUN-2 source was operated in accumulation mode. This table is a subset of Table 3.1. The values of χ_r^2 in this table are for fits to the empirical storage curves of Equations 3.1 and 3.2.

Index	Front	Size	T [K]	τ_0 [s]	N_0	τ_d [h]	χ_r^2
13	Alu.	59	11.1	551(5)	2706(20)	126(10)	1.08
14	Alu.	100	280.0	301(2)	2352(17)	315(34)	1.97
16	SS.	22	295.8	297(4)	1741(21)	–	1.22
18	SS.	57	11.7	546(6)	2074(18)	149(16)	1.72

5.2, which is a subset of Table 3.1.) All four models for $\tau(E_b)$ discussed in the previous section are applied to these data; the fits are plotted in Figure 5.1, and the corresponding optimal parameter values are summarized in Table 5.3. From an examination of the reduced chi-squared statistic, χ_r^2, it is not always clear which models are most appropriate to which data. In order to use these fits to produce estimates of the loss factor of the CYTOP coating, both statistical and physical arguments are considered.

For the two low temperature measurements, 13 and 18, at 11.1 K and 11.7 K, respectively, the SCX and WPH models do not produce better fits than the simpler SC model. Table 5.3 shows that these models choose the same values for η and N_0 as the SC model—but with larger uncertainties, and that the respective optimal values of the X and a_p parameters are zero, meaning that these models are essentially the same as the SC model. The WPA model, on the other hand, does give a better fit. As an objective measure of whether or not this extended model is appropriate, we use the "F-test". (This test was also used in Section 3.2.) That is, we compute the statistic

$$F = \frac{\chi^2_{\nu-3} - \chi^2_{\nu-4}}{\chi^2_{\nu-4}/(\nu-4)}\,, \tag{5.7}$$

where $\chi^2_{\nu-3}$ is the chi-squared value of the fit for the SC model, and $\chi^2_{\nu-4}$ is that for the WPA model. Under the assumption that the $\chi^2_{\nu-3}$ and $\chi^2_{\nu-4}$ are chi-squared random variables with $\nu-3$ and $\nu-4$ degrees of freedom, respectively, the random variable F is described by the F-distribution, with 1 "numerator" degrees of freedom and $\nu-4$ "denominator" degrees of freedom. (The number of degrees of freedom is the size of the data set, ν, less the number of parameters in the model; these are nominally 2 and 3, respectively, for the SC and WPA models, plus 1 from the use of τ_d.) A large value of F suggests that the model with an additional fitting parameter is appropriate. (See, for example, Chapter 11 in [112].) For example, for measurement 13, $\chi^2_{\nu-3} = 0.91 \times (59-3) = 51.0$, and $\chi^2_{\nu-4} = 0.86 \times (59-4) = 47.1$, so that $F = (51.0-47.1)/0.86 = 4.6$. The probability that $F \geq 4.6$ in this case is approximately 3.7 %. A similar calculation gives that, for measurement 18, $F = 3.6$, and that this probability, or the "p-value" as it is sometimes called, is approximately

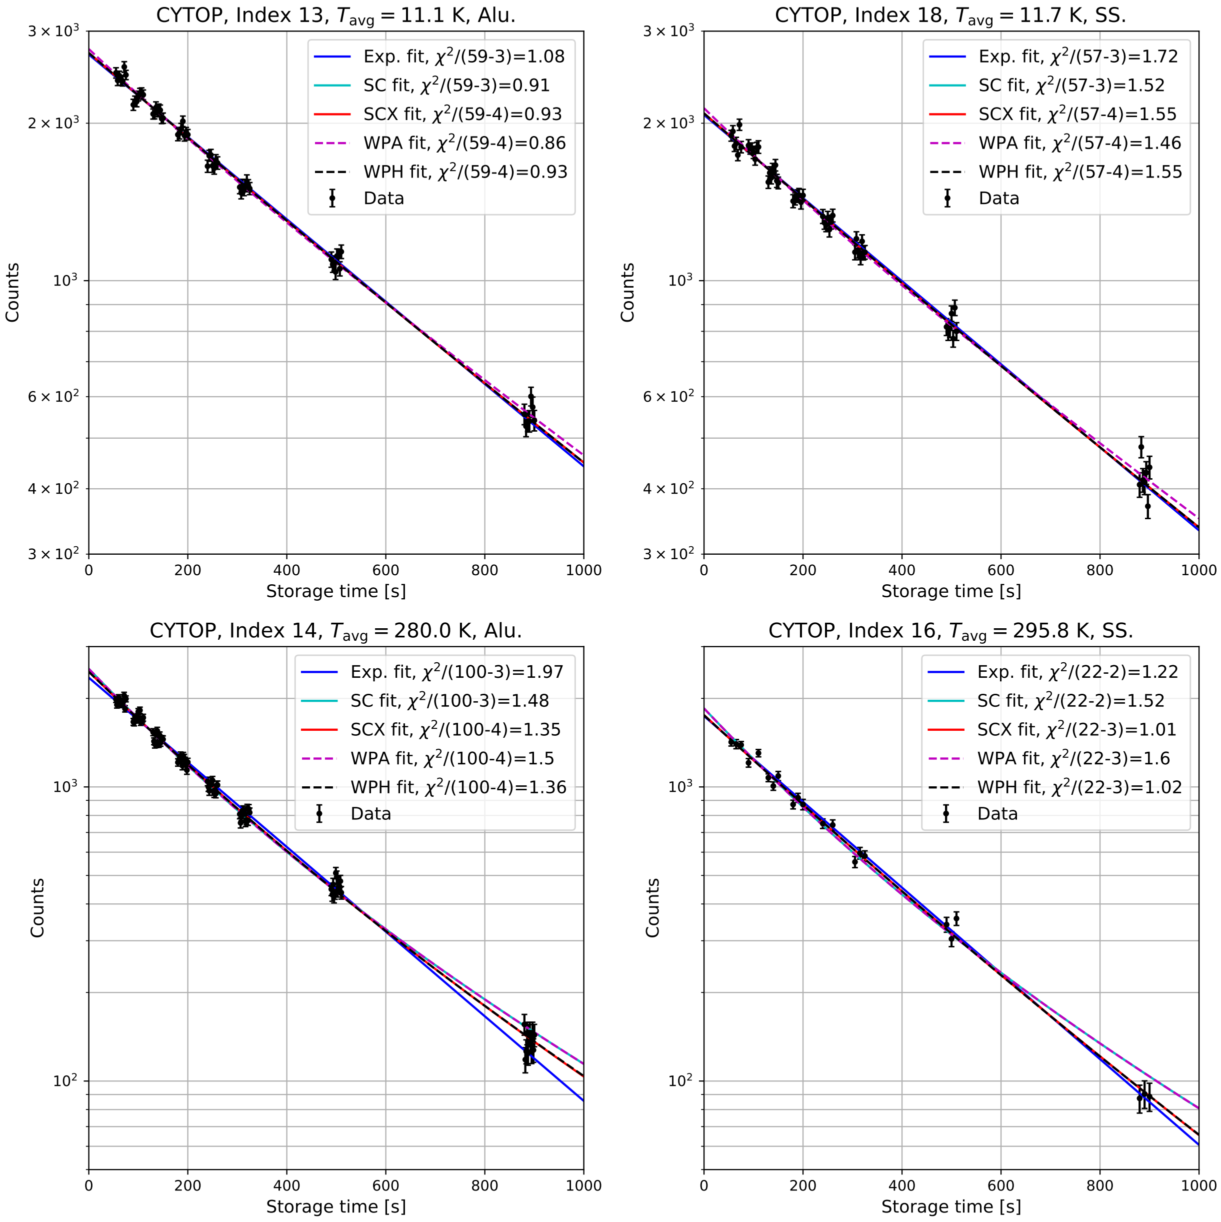

Figure 5.1: Fits using the models discussed in Section 5.1 to the storage data of measurements 13, 14, 16, and 18 for CYTOP. Fits to the empirical storage curves of Equations 3.1 and 3.2 are also plotted in dark blue, labeled as "Exp.". (See Table 5.2.) At low temperature, the SC, SCX, and WPH curves lie on top of one another, because for the latter two models, the optimal values of the additional parameters, X and a_p, respectively, are approximately zero. (Also, for low temperature, the storage data have been corrected for a systematic decrease by dividing $N(t)$ by the factor $\exp(-t_m/\tau_d)$, where t_m is the absolute measurement time, and τ_d is the time constant of the systematic effect; see the text.) At higher temperatures, the SC and WPA curves, and the SCX and WPH curves, respectively, are approximately the same. For the former two, this is similarly due to the fact that the optimal value of the additional a_p parameter in the WPA model is approximately zero.

Table 5.3: Fit parameters for CYTOP storage data using four different models. (See Table 5.1.) The p-value for the F-test—see the text—is given where appropriate.

Index	T [K]	–	SC	SCX	WPA	WPH
13	11.1					
		η	$2.82 \pm 0.09 \times 10^{-5}$	$2.8 \pm 1.3 \times 10^{-5}$	$2.38 \pm 0.13 \times 10^{-5}$	$2.8 \pm 2.6 \times 10^{-5}$
		N_0	2725 ± 24	2725 ± 29	2770 ± 40	2725 ± 31
		X	–	~ 0 s^{-1}	–	–
		a_p	–	–	7.8 ± 3.4 mm^2	~ 0 mm^2
		χ_r^2	0.91	0.93	0.86	0.93
		p-value	–	–	3.7 %	–
14	280.0					
		η	$10.97 \pm 0.16 \times 10^{-5}$	$8.7 \pm 0.7 \times 10^{-5}$	$10.97 \pm 0.21 \times 10^{-5}$	$6.7 \pm 1.2 \times 10^{-5}$
		N_0	2521 ± 24	2465 ± 27	2521 ± 33	2463 ± 27
		X	–	$(3.7 \pm 1.1) \times 10^{-4}$ s^{-1}	–	–
		a_p	–	–	~ 0 mm^2	5.4 ± 1.6 mm^2
		χ_r^2	1.48	1.35	1.50	1.36
		p-value	–	0.17 %	–	0.22 %
16	295.8					
		η	$11.3 \pm 0.4 \times 10^{-5}$	$5.7 \pm 2.2 \times 10^{-5}$	$11.3 \pm 0.9 \times 10^{-5}$	$1 \pm 4 \times 10^{-5}$
		N_0	1849 ± 33	1760 ± 40	1850 ± 150	1750 ± 40
		X	–	$(9 \pm 4) \times 10^{-4}$ s^{-1}	–	–
		a_p	–	–	~ 0 mm^2	14 ± 6 mm^2
		χ_r^2	1.52	1.01	1.60	1.02
		p-value	–	0.35 %	–	0.38 %
18	11.7					
		η	$2.91 \pm 0.10 \times 10^{-5}$	$2.9 \pm 1.5 \times 10^{-5}$	$2.43 \pm 0.13 \times 10^{-5}$	$2.9 \pm 3.0 \times 10^{-5}$
		N_0	2089 ± 19	2089 ± 26	2137 ± 33	2089 ± 26
		X	–	~ 0 s^{-1}	–	–
		a_p	–	–	10 ± 4 mm^2	~ 0 mm^2
		χ_r^2	1.52	1.55	1.46	1.55
		p-value	–	–	6.8 %	–

6.8 %. These relatively low p-values suggest that the introduction of the additional model parameter a_p in the WPA model is reasonable.

Although the SCX and WPH models are poorly suited to the low temperature storage data, Table 5.3 shows that fits to these models for the “room temperature” measurements, 14 and 16, at 280.0 K and 295.8 K, respectively, produce smaller values of χ_r^2 compared to fits to the SC model. Furthermore, the F-test suggests that both models are reasonable: for measurement 14, the associated p-values for the SCX and WPH models are 0.17 % and 0.22 %, respectively, and for measurement 16, 0.35 % and 0.38 % respectively. In contrast, for these data, the WPA model now gives a worse fit for both measurements: as was the case for the SCX and WPH models at low temperature, the WPA model gives essentially the same result as the SC model at room temperature. In the following, we consider several physical arguments for the differing success of the SCX, WPH, and WPA models when applied to room temperature and low temperature data.

For the SCX model, one possible explanation for a nonzero value of X at room temperature could be the increased pressure in the system compared to when it is at low temperature. As discussed in Chapter 3, this pressure difference is suspected to be caused by residual gases, like water vapor, freezing in the guides at low temperature. A rough estimate of the partial pressure of water vapor at room temperature required to produce $X \sim 10^{-4}\,\mathrm{s}^{-1}$ can be obtained from its measured UCN scattering cross section: approximately 20×10^3 b for UCN of velocity $6.6\,\mathrm{m\,s}^{-1}$ [130]. Using $X \sim \rho\sigma v$, one obtains a partial pressure of $\sim 10^{-4}$ mbar for water vapor, which seems too high considering the fact that typical pressures in the experiment at room temperature were in the low 10^{-6} mbar range. These pressures, however, were measured in a vacuum space that is connected to the storage bottle and guides through a 3 mm diameter hole—see (9) in Figure 3.2, which offers relatively inefficient pumping. (For reference, such a hole has a conductance of roughly $0.8\,\mathrm{l\,s}^{-1}$ for air at room temperature; see, for example, Equation 98 in Chapter 14 of [131].) Furthermore, there are other possible residual gases—mainly hydrocarbons—that have cross sections several times larger than that of water vapor [129, 130]. Without specific knowledge of the residual gas species present in the experiment, it is difficult to exclude the possibility that a variety of these gases are responsible for the nonzero value of X at room temperature. (A residual gas analyzer—essentially a mass spectrometer—will be installed on the experiment in the future to aid in the clarification of this point.)

At first glance, one might suspect that a nonzero value for the a_p fit parameter of the WPH model—i.e. the total area of the hole (or holes)—in going from low temperature to room temperature could arise from the different thermal expansion coefficients of the different metals used in the storage bottle assembly. As discussed in Chapter 3, the front flange of the storage bottle is made out of aluminum, while the bottle plug is made out of stainless steel. (See Figure 3.12 for a cross section view of these components and their dimensions.) Now, the relative linear thermal expansion coefficients of these materials are approximately $-0.42\,\%$ and $-0.29\,\%$ near 10 K, respectively. (See Figure 3.17 in [95]; as opposed to the linear thermal expansion coefficient, the *relative* linear thermal expansion coefficient gives the percent contraction in length of a material when it is cooled from room temperature.) This means that the diameter of the bottle entrance hole will shrink by $0.42\,\%$, while the maximum diameter of the plug, for example, will only shrink by $0.29\,\%$. It would seem possible, then, that this change in relative geometry could result in different size gaps between the plug and bottle entrance at different temperatures if these components are not perfectly round.

The different thermal expansion coefficients of the stainless steel plug and aluminum bottle entrance could also be responsible for a nonzero value of a_p in the WPA model. The change in relative geometry could expose some material underneath the CYTOP coating that is worn away along the line of contact of the bottle entrance fillet and the plug as a result of repeated opening and closing of the bottle. It is difficult, however, to estimate the size of such an effect. If the plug and bottle entrance axes are perfectly aligned, the only effect of the different thermal contractions is to make the plug protrude farther into the bottle; the line of contact is a circle, and is unchanged on the entrance hole fillet. Using the relative thermal expansion coefficients given above, the additional area on the plug that should become visible to UCN inside the bottle is a thin ring around the cone of the plug, approximately 2.5 mm from the front face of the plug, and roughly $1.2\,\mathrm{mm}^2$ in size. Of course, the optical potential of stainless steel, 189 neV [96], is greater than that of aluminum, so such an exposure would not have the same effect on $N(t)$ as would an aluminum patch. However, if the plug and entrance hole axes are not parallel, the line of contact would differ from a circle, and it is possible that an area of a similar order of magnitude could be exposed on the aluminum bottle entrance fillet.

It is also possible that the same area of aluminum patches is present at room temperature and at low temperature, but that, at room temperature, the nonzero value of a_p is “disguised” by a relatively large η. As discussed in Section 2.4, deviation of $N(t)$ from single exponential behavior can be caused by either large values of η, or by the presence of weak patches. At room temperature, fits to the WPA model suggest that the loss factor has increased appreciably, a result that will tend to mask the influence of a nonzero a_p on $N(t)$. (The variety of possible explanations for the increase of η at room temperature will not be discussed here; see, e.g. [82] for a review.)

It is difficult to rule out the possibility of any of these physical scenarios for the SCX, WPH, and WPA models. We will therefore estimate η by taking weighted averages of the loss factors returned by those fits that have reasonable values of χ_r^2. To this end, we will use the results of the SC, SCX, and WPH models to estimate η at room temperature, and the results of the SC and WPA models to estimate η at low temperature. Furthermore, the initial UCN energy spectrum in the bottle, $n_b(0, E_b)$, was simulated assuming ideal guide transmission: it should therefore be equally appropriate to storage data taken with and without the stainless steel insert at the bottle entrance. However, the value of the initial number of UCN in the storage bottle, N_0, is clearly different for the two cases, with the larger values coming from measurements in which the bottle entrance is uncovered. Because stainless steel has a higher optical potential than aluminum—$V = 189\,\mathrm{neV}$ for stainless steel [96], and $V = 54\,\mathrm{neV}$ for aluminum [3]—this is presumably as a result of the loss due to the stainless steel insert being greater than that of the bare aluminum at the bottle entrance. (Although the precise values of these losses are not known, iron, the primary ingredient in steel, has an absorption cross section that is larger than that of aluminum by a factor of approximately 11 [2].) Enhanced UCN loss at this location in the guide system should tend to shift the spectrum more toward a "direct filling only" scenario, rather than one that includes "refilling". (See Chapter 4.) Therefore, because $n_b(0, E_b)$ was constructed under the assumption of no refilling, we will make estimates of the loss factor of CYTOP at room temperature and low temperature using only the storage data taken with the stainless steel insert at the bottle entrance. (This has the effect of increasing the value of η, which is in keeping with our goal of making conservative estimates.)

The result of all of these considerations is that at an average bottle temperature of 11.7 K, the loss factor of CYTOP is $\eta = 2.7(2) \times 10^{-5}$, and that at an average bottle temperature of 295.8 K, the loss factor of CYTOP is $\eta = 1.1(1) \times 10^{-4}$. The error bars on these values of η are larger than those that would be obtained by taking the weighted averages of the quantities in Table 5.3 using the error bars reported there. This comes from rescaling the uncertainties by $\sqrt{\chi_r^2}$ of the original weighted average to produce $\chi_r^2 = 1$. (This rescaling approach was also used to estimate the uncertainty of the optical potential of CYTOP; see Section 3.2.)

Table 5.4: Summary table of dPE storage measurements in which the SUN-2 source was operated in accumulation mode. See Table 3.2. The values of χ_r^2 in this table are for fits to the empirical storage curves of Equations 3.4 and 3.5.

Index	Front	Size	T [K]	τ_1 [s]	τ_2 [s]	N_1	N_2	τ_d [h]	χ_r^2
20	Alu.	70	10.9	162(33)	691(39)	1331(214)	3415(259)	131(6)	1.09
21	Alu.	73	289.9	125(9)	320(28)	2221(172)	879(205)	362(78)	0.98
23	SS.	42	295.1	101(13)	271(24)	1340(134)	752(175)	–	1.44
26	SS.	64	11.9	78(20)	591(15)	584(84)	1972(53)	191(20)	0.97

5.3 Deuterated polyethylene (dPE)

As with the CYTOP measurements, there are four dPE storage data sets that may be analyzed as in the previous section: measurements 20, 21, 23, and 26. (See in Table 5.4, which is a subset of Table 3.2; measurement 24 is excluded from this section on the basis of the large spread in its data, as is evinced by the poor fits to the empirical storage model of Equation 3.4.) Again, all four models for $\tau(E_b)$, discussed in Section 5.1, are applied to these data; the fits are plotted in Figure 5.2, and their parameters are summarized in Table 5.5.

The trends of the fits to the dPE data bear a striking resemblance to those found for the CYTOP data. At low temperature, the SCX and WPH models are worse than the simpler SC model, while the WPA model is an improvement. This improvement is, in fact, rather significant: the F-test applied to the WPA model results gives p-values of roughly 1×10^{-12} % and 5×10^{-8} % for measurements 20 and 26, respectively. The clear success of the WPA model with respect to these statistical tests is in keeping with a number of experimental facts and observations. For one, unlike the CYTOP bottle, the components of the dPE bottle were coated individually. While it is possible that wear on the entrance hole fillet is responsible for some exposure of aluminum on the interior of the bottle (see Section 5.2), and there is even some possible evidence of thinning of the dPE coating on the entrance plate (see Figure 3.12), it is also true that the assembly of the individually-coated components will introduce “seams” along the inside of the bottle. Furthermore, the deviation of $N(t)$ away from a single, exponential decay is significantly more pronounced in the dPE data compared to the CYTOP data. This behavior is not only apparent at room temperature, but at low temperature. There, all models suggest that the loss factor is smaller, and it thus seems likely that another mechanism is responsible for the enhanced curvature of $N(t)$.

The trends of the fits to the dPE data and to the CYTOP data are also common for the room temperature measurements: the SCX and WPH models seem to be more appropriate descriptions of the room temper-

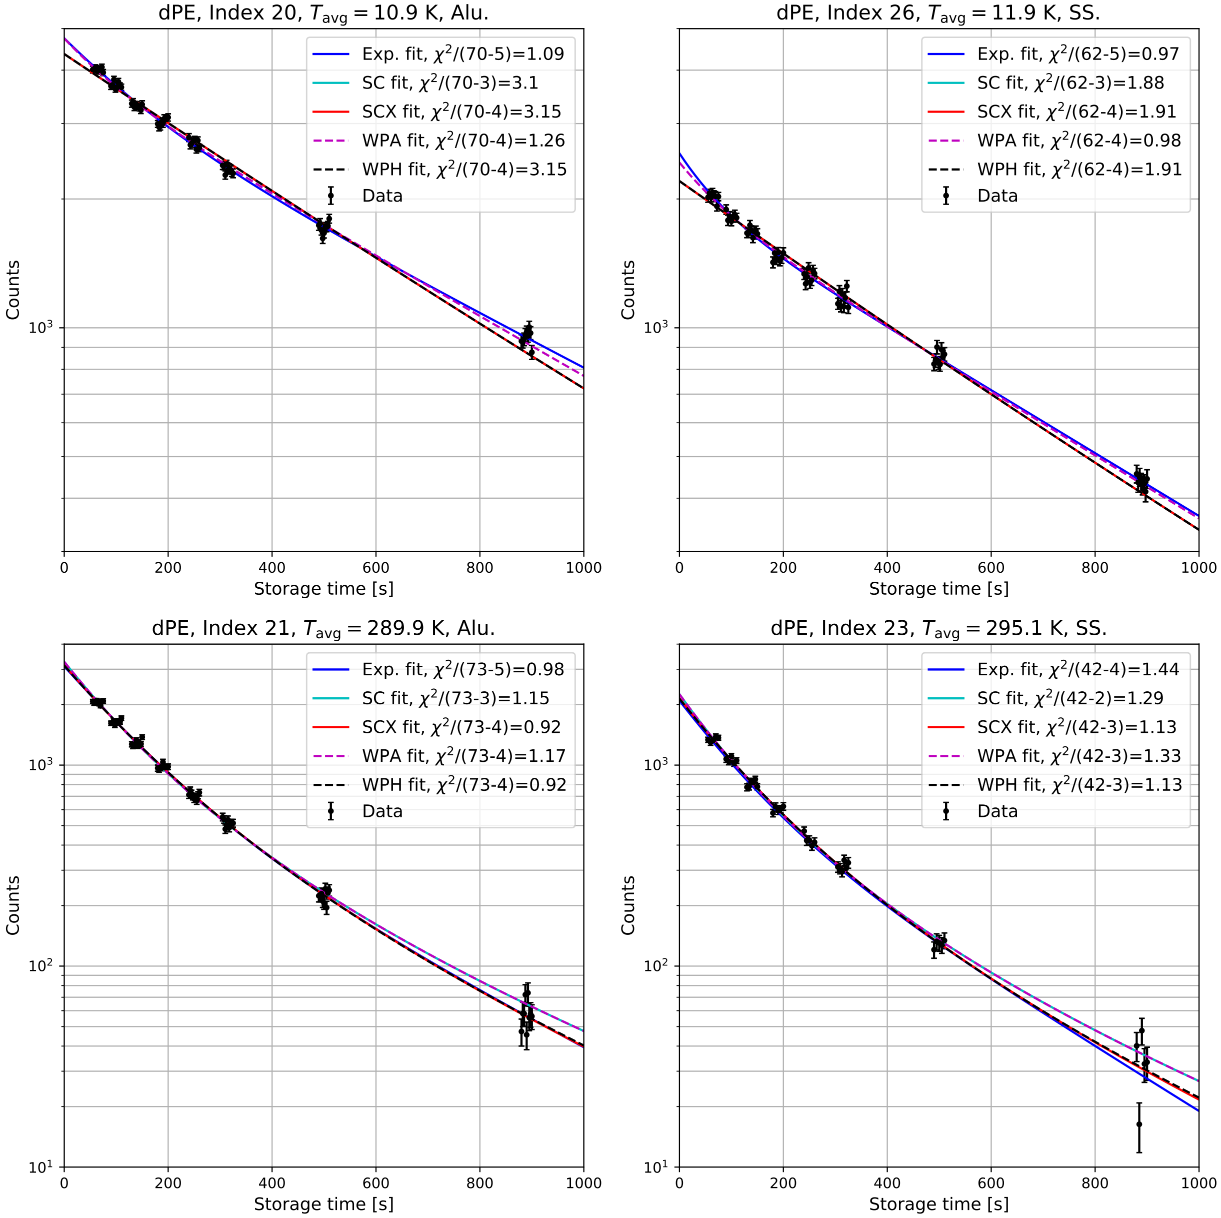

Figure 5.2: Fits using the models discussed in Section 5.1 to the storage data of measurements 20, 21, 23, and 26 for dPE. (See Table 5.4.) Fits to the empirical storage curves of Equations 3.4 and 3.5 are also plotted in dark blue, labeled as "Exp.". (See Table 5.4.) At low temperature, the SC, SCX, and WPH curves lie on top of one another, because for the latter two models, the optimal values of the additional parameters, X and a_p, respectively, are approximately zero. (Also, for low temperature, the storage data have been corrected for a systematic decrease by dividing $N(t)$ by the factor $\exp(-t_m/\tau_d)$, where t_m is the absolute measurement time, and τ_d is the time constant of the systematic effect; see the text.) At higher temperatures, the SC and WPA curves, and the SCX and WPH curves, respectively, are approximately the same. For the former two, this is similarly due to the fact that the optimal value of the additional a_p parameter in the WPA model is approximately zero.

Table 5.5: Fit parameters for dPE storage data using four different models. (See Table 5.1.) The p-value for the F-test—see the text—is given where appropriate.

Index	T [K]	–	SC	SCX	WPA	WPH
20	10.9					
		η	$2.79 \pm 0.06 \times 10^{-5}$	$2.8 \pm 1.0 \times 10^{-5}$	$2.41 \pm 0.15 \times 10^{-5}$	$2.8 \pm 1.9 \times 10^{-5}$
		N_0	4362 ± 35	4362 ± 40	4743 ± 110	4362 ± 41
		X	–	~ 0 s^{-1}	–	–
		a_p	–	–	33 ± 11 mm^2	~ 0 mm^2
		χ_r^2	3.10	3.15	1.26	3.15
		p-value	–	–	1×10^{-12} %	–
21	289.9					
		η	$23.9 \pm 0.4 \times 10^{-5}$	$20.5 \pm 0.9 \times 10^{-5}$	$23.9 \pm 0.4 \times 10^{-5}$	$17.6 \pm 1.6 \times 10^{-5}$
		N_0	3266 ± 41	3146 ± 48	3266 ± 84	3132 ± 48
		X	–	$(3.8 \pm 1.0) \times 10^{-4}$ s^{-1}	–	–
		a_p	–	–	~ 0 mm^2	6.6 ± 1.7 mm^2
		χ_r^2	1.15	0.92	1.17	0.92
		p-value	–	5×10^{-3} %	–	4×10^{-3} %
23	295.1					
		η	$27.1 \pm 0.7 \times 10^{-5}$	$23.0 \pm 1.6 \times 10^{-5}$	$27.1 \pm 0.8 \times 10^{-5}$	$19.6 \pm 2.8 \times 10^{-5}$
		N_0	2253 ± 45	2151 ± 53	2253 ± 111	2140 ± 55
		X	–	$(4.3 \pm 1.1) \times 10^{-4}$ s^{-1}	–	–
		a_p	–	–	~ 0 mm^2	7.4 ± 2.8 mm^2
		χ_r^2	1.29	1.13	1.33	1.13
		p-value	–	1.2 %	–	1.4 %
26	11.9					
		η	$3.16 \pm 0.10 \times 10^{-5}$	$3.2 \pm 1.3 \times 10^{-5}$	$2.93 \pm 0.27 \times 10^{-5}$	$3.2 \pm 2.6 \times 10^{-5}$
		N_0	2202 ± 21	2202 ± 26	2433 ± 84	2202 ± 26
		X	–	~ 0 s^{-1}	–	–
		a_p	–	–	42 ± 21 mm^2	~ 0 mm^2
		χ_r^2	1.88	1.91	0.98	1.91
		p-value	–	–	5×10^{-8} %	–

ature dPE storage data. (See Table 5.5.) In light of the similar behavior of the fits to the dPE storage data and the fits to the CYTOP storage data, estimates of the loss factor of dPE are made in the same way as for CYTOP. The data used to make the final estimates of the loss factor are from measurement 23 and 26, in which the stainless steel insert was installed at the bottle entrance. The weighted average of η from the SC, SCX, and WPH models is calculated for the storage measurement at room temperature, while the weighted average of η from the SC and WPA models are calculated at low temperature. This analysis results in estimates of $\eta = 3.1(1) \times 10^{-5}$ at an average bottle temperature of 11.9 K, and $\eta = 2.6(1) \times 10^{-4}$ at an average bottle temperature of 295.1 K.

Chapter 6

Conclusion

In this work, we have reported on UCN storage in two 3.8 liter “bottles”, coated with the fluoropolymer CYTOP and with deuterated polyethylene (dPE), carried out at the SUN-2 prototype UCN source at the Institute Laue-Langevin (ILL), during reactor cycles 188 and 189. The two materials were investigated for their potential use as coatings for the production volume of the SuperSUN UCN source [51], and for part of the storage cells of the PanEDM neutron electric dipole moment experiment [50], respectively, both of which are currently under construction at the ILL. (See Section 1.4.) The neutron optical potential of CYTOP was determined by complementary neutron reflectometry measurements, and was found to be 115.2(2) neV, slightly lower than the 117 neV predicted on the basis of the manufacturer-specified density. (See Section 3.2.)

For CYTOP and dPE, no evidence of coating deterioration was observed as a result of repeated thermal cycling of the storage bottles, suggesting that both of these coatings are, in fact, suitable for cryogenic applications. During storage measurements, a systematic effect was observed in both bottles at low temperatures, in which the initial number of UCN in the storage bottle gradually decreased with absolute measurement time, without appearing to affect the bottle storage time constants. It is suspected that this effect is the result of residual gases, such as water vapor, freezing inside the UCN guide section near the “cold head” refrigerator. (See Sections 3.2 and 3.3.) In future studies, a residual gas analyzer—essentially a mass spectrometer—will be installed on the experiment to clarify this point.

The CYTOP storage curves were found to be well described by a single, empirical storage time constant, τ_0, the values of which are reported in Table 3.1 for various temperatures and modes of operation of the SUN-2 source. To give a general sense of the size of τ_0, we report here weighted averages from different measurements: for bottle temperatures below 15 K, this gives an estimate $\tau_0 = 564(7)\,\mathrm{s}$, and for bottle temperatures above 290 K, this gives an estimate of $\tau_0 = 311(9)\,\mathrm{s}$. (See Section 3.2.) Conservative estimates of the loss factor of CYTOP are made from the synthesis of the experimental storage data presented in Chapter 3 and simulations of the SUN-2 UCN source presented in Chapter 4. These are $\eta = 2.7(2) \times 10^{-5}$ at an average bottle temperature of 11.7 K, and $\eta = 1.1(1) \times 10^{-4}$ at an average bottle temperature of 295.8 K. (See Section 5.2.)

For the dPE measurements, a single time constant was not sufficient to describe the storage data. Instead, a "double-exponential" model with two time constants was used: τ_1 and τ_2, with τ_1 being the shorter of the two. (See Equation 3.4.) The values of these time constants are given in Table 3.2 for various temperatures and modes of operation of the SUN-2 source. To give a general sense of the sizes of τ_1 and τ_2, we report here weighted averages from different measurements: for bottle temperatures below 15 K, this gives an estimate $\tau_1 = 87(11)$ s and $\tau_2 = 580(7)$ s, and for bottle temperatures above 290 K, this gives an estimate of $\tau_1 = 80(5)$ s and $\tau_2 = 267(7)$ s. (See Section 3.3.) Conservative estimates of the loss factor of dPE are made in the same manner as for CYTOP. These are $\eta = 3.1(1) \times 10^{-5}$ at an average bottle temperature of 11.9 K and $\eta = 2.6(1) \times 10^{-4}$ at an average bottle temperature of 295.1 K. (See Section 5.3.) The latter value is in rough agreement with estimates reported in the literature for composite bottles made partly out of dPE: $1.3(3) \times 10^{-4}$ [60] and $3(1) \times 10^{-4}$ [59].

There are two main limitations of the analysis used to produce estimates of the loss factors of CYTOP and dPE. These are related to the unknown UCN transmission characteristics of the guides between the source and bottle, and to the unknown storage properties of the SUN-2 source. As discussed in Appendix E, the expected effects of these factors on the initial UCN energy spectrum in the storage bottle, $n_b(0, E_b)$, would be to bias it in opposite directions, that is toward lower and higher energy, respectively. However, the degree to which this occurs—and the subsequent effect on estimates of η—are not known. We therefore report values of η that are calculated using the $n_b(0, E_b)$ presented in Chapter 4.

Appendix A

Reflectivity matters

A.1 Derivation of $\mu(E)$, the loss probability per "bounce"

Consider the one-dimensional problem of a plane wave of energy E incident on the real step potential of height V located at the origin. The wavefunctions to the right and left of the boundary are

$$\begin{aligned} \psi_L(x) &= \exp(ikx) + r\exp(-ikx) \\ \psi_R(x) &= t\exp(ik'x)\,, \end{aligned} \tag{A.1}$$

where $k = \sqrt{2mE/\hbar^2}$ and $k' = \sqrt{2m(E-V)/\hbar^2}$, the latter of which is imaginary when $E < V$. Matching boundary conditions gives

$$\begin{aligned} 1 + r &= t \\ ik(1-r) &= ik't\,, \end{aligned} \tag{A.2}$$

which may be solved for the reflection and transmission coefficients r and t to give

$$\begin{aligned} r &= \frac{k-k'}{k+k} \\ t &= \frac{2k}{k+k'}\,. \end{aligned} \tag{A.3}$$

The reflection probability is then $|r|^2$, and the transmission probability is $1-|r|^2$, which $\neq |t|^2$ in general. For the case $E < V$ we have that

$$r = \frac{\sqrt{E} - i\sqrt{V-E}}{\sqrt{E} + i\sqrt{V-E}}\,, \tag{A.4}$$

from which it is seen that $|r|^2 = 1$, and that there is no loss on reflection from a purely real step potential.

Now suppose that the potential is complex, that is $U = V - iW$ as for the neutron optical potential, and that $W \ll V$. For the case $E < V$ then, the reflection coefficient is

$$\begin{aligned} r &= \frac{\sqrt{E} - i\sqrt{V - iW - E}}{\sqrt{E} + i\sqrt{V - iW - E}} \\ &= \frac{\sqrt{E} - i\sqrt{V - E} - \frac{1}{2}\frac{W}{\sqrt{V-E}}}{\sqrt{E} + i\sqrt{V - E} + \frac{1}{2}\frac{W}{\sqrt{V-E}}} \,, \end{aligned} \tag{A.5}$$

where a Taylor expansion has been made in $W/(V - E) \ll 1$ if V and E are not too close. Computing the reflection probability, we have

$$\begin{aligned} |r|^2 &= \frac{V - W\sqrt{\frac{E}{V-E}} + \mathcal{O}(W^2)}{V + W\sqrt{\frac{E}{V-E}} + \mathcal{O}(W^2)} \\ &= 1 - 2\eta\sqrt{\frac{E}{V-E}} + \mathcal{O}(W^2) \,, \end{aligned} \tag{A.6}$$

where $\eta = W/V$ and higher order terms in $W/(V - E)$ and so on are neglected. We then obtain the result that, to first order, the loss probability $\mu(E)$ per wall interaction, or per “bounce”, is

$$\mu(E) = 1 - |r|^2 = 2\eta\sqrt{\frac{E}{V - E}} \,, \tag{A.7}$$

for $E < V$. If $E > V$, the primary loss is via transmission.

An important application of this result concerns neutron storage in a superfluid ^{4}He “medium”, real potential $U_m = V_m$, within the walls of a container with complex potential $U_w = V_w - iW_w$. In this case, the reflection coefficient is

$$\begin{aligned} r &= \frac{\sqrt{E - U_m} - \sqrt{E - U_w}}{\sqrt{E - U_m} + \sqrt{E - U_w}} \\ &= \frac{\sqrt{E - V_m} - \sqrt{E - V_w + iW_w}}{\sqrt{E - V_m} + \sqrt{E - V_w + iW_w}} \\ &= \frac{\sqrt{E'} - \sqrt{E' - (V_w - V_m) + iW_w}}{\sqrt{E'} + \sqrt{E' - (V_w - V_m) + iW_w}} \,, \end{aligned} \tag{A.8}$$

where a change of variables has been made to $E' = E - V_m$, the energy with respect to the superfluid medium. It follows that the loss probability for $E' < (V_w - V_m)$ is

$$\mu(E') = 2\frac{W_w}{V_w - V_m}\sqrt{\frac{E'}{(V_w - V_m) - E'}} = 2\eta_w \, \frac{V_w}{V_w - V_m}\sqrt{\frac{E'}{(V_w - V_m) - E'}} \,, \tag{A.9}$$

and the storage now depends on a decreased optical potential $V' = V_w - V_m$ and an increased loss factor $\eta' = \eta_w V_w/(V_w - V_m)$.

A.2 Reflectivity matrix method for a layered system

Consider a system of $n+1$ layers of thickness d_i, starting with the 0th layer (semi-infinite vacuum) on the left, and the nth layer on the right (semi-infinite substrate). There are n total interfaces, with e.g. the 0th interface between layer 0 and layer 1. We take the direction of the incident neutron to be $+x$, so that the wavefunction in the i-th layer is

$$\psi_i(x) = A_i \exp(ik_i x) + B_i \exp(-ik_i x)\,, \tag{A.10}$$

where A_i and B_i are the amplitudes of the right-going and left-going waves in the layer, respectively, and k_i is the wavevector in the layer of potential U_i. Taking the origin to be at the 0th interface, between vacuum and the top layer of the structure, and matching boundary conditions gives

$$\begin{aligned} A_{i-1}\exp(ik_{i-1}c_{i-1}) + B_{i-1}\exp(-ik_{i-1}c_{i-1}) &= A_i\exp(ik_i c_{i-1}) + B_i\exp(-ik_i c_{i-1}) \\ ik_{i-1}A_{i-1}\exp(ik_{i-1}c_{i-1}) - ik_{i-1}B_{i-1}\exp(-ik_{i-1}c_{i-1}) &= ik_iA_i\exp(ik_i c_{i-1}) - ik_iB_i\exp(-ik_i c_{i-1})\,, \end{aligned} \tag{A.11}$$

where $c_i = \sum_j^i d_j$ is the position of the ith interface. (Set $d_0 = 0$, so that for the 0th interface $c_0 = d_0 = 0$, and for the last interface $c_{n-1} = d_{n-1} + ... + d_1$, and so on.) Rearranging this expression to obtain the coefficients A_{i-1} and B_{i-1} in terms of the A_i and B_i gives

$$\begin{aligned} 2\frac{k_{i-1}}{k_{i-1}+k_i}A_{i-1} &= A_i + \frac{k_{i-1}-k_i}{k_{i-1}+k_i}B_i \\ 2\frac{k_{i-1}}{k_{i-1}+k_i}B_{i-1} &= \frac{k_{i-1}-k_i}{k_{i-1}+k_i}A_i + B_i\,, \end{aligned} \tag{A.12}$$

where we have temporarily absorbed the exponential terms into the A and B. Following Equation A.3 in the previous section, we make the following definitions

$$\begin{aligned} r_{i-1,i} &= \frac{k_{i-1}-k_i}{k_{i-1}+k_i} \\ t_{i-1,i} &= 2\frac{k_{i-1}}{k_{i-1}+k_i}\,. \end{aligned} \tag{A.13}$$

Inserting these into Equation A.12, in matrix form, yields

$$\begin{pmatrix} A_{i-1} \\ B_{i-1} \end{pmatrix} = \frac{1}{t_{i-1,i}} \begin{bmatrix} 1 & r_{i-1,i} \\ r_{i-1,i} & 1 \end{bmatrix} \begin{pmatrix} A_i \\ B_i \end{pmatrix} , \tag{A.14}$$

or, with the exponential terms returned,

$$\begin{bmatrix} e^{ik_{i-1}c_{i-1}} & 0 \\ 0 & e^{-ik_{i-1}c_{i-1}} \end{bmatrix} \begin{pmatrix} A_{i-1} \\ B_{i-1} \end{pmatrix} = \frac{1}{t_{i-1,i}} \begin{bmatrix} 1 & r_{i-1,i} \\ r_{i-1,i} & 1 \end{bmatrix} \begin{bmatrix} e^{ik_i c_{i-1}} & 0 \\ 0 & e^{-ik_i c_{i-1}} \end{bmatrix} \begin{pmatrix} A_i \\ B_i \end{pmatrix} . \tag{A.15}$$

Further simplification results in

$$\begin{aligned} \begin{pmatrix} A_{i-1} \\ B_{i-1} \end{pmatrix} &= \frac{1}{t_{i-1,i}} \begin{bmatrix} \exp\left[i(k_i + k_{i-1})c_{i-1}\right] & r_{i-1,i}\exp\left[-i(k_i - k_{i-1})c_{i-1}\right] \\ r_{i-1,i}\exp\left[i(k_i + k_{i-1})c_{i-1}\right] & \exp\left[-i(k_i - k_{i-1})c_{i-1}\right] \end{bmatrix} \begin{pmatrix} A_i \\ B_i \end{pmatrix} \\ \begin{pmatrix} A_{i-1} \\ B_{i-1} \end{pmatrix} &= \frac{1}{t_{i-1,i}} \begin{bmatrix} M_{i-1,i} \end{bmatrix} \begin{pmatrix} A_i \\ B_i \end{pmatrix} , \end{aligned} \tag{A.16}$$

which defines the matrix $M_{i-1,i}$, representing the interface between the ith and $(i-1)$th layers.

Taking the incident wavefunction to be $A_0 = 1$, $B_0 = R$, and the final wave to be $A_n = T$, $B_n = 0$, gives

$$\begin{pmatrix} 1 \\ R \end{pmatrix} = \frac{1}{t_{0,1} \ldots t_{n-1,n}} \begin{bmatrix} M_{0,1} \end{bmatrix} \ldots \begin{bmatrix} M_{n-1,n} \end{bmatrix} = \frac{1}{t_{0,1} \ldots t_{n-1,n}} \begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix} \begin{pmatrix} T \\ 0 \end{pmatrix} . \tag{A.17}$$

Finally then, the reflection probability, or reflectivity, is $|R|^2 = |\gamma/\alpha|^2$.

This calculation has been performed for perfectly flat surfaces. The result may be extended to include surface roughness for the case that, at each two-dimensional layer interface, the position the surface is a Gaussian random variable with mean zero and variance σ_i^2. The distorted-wave Born approximation can then be used to calculate the average effect of the surface roughness on the reflection coefficient $r_{i-1,i}$ using perturbation theory. The result is [132]

$$r'_{i-1,i} = r_{i-1,1} \exp\left(-2k_{i-1}k_i\sigma_i^2\right) , \tag{A.18}$$

where the prime denotes the modified reflection coefficient, and σ_i is roughness at the interface between layers $i-1$ and i.

Appendix B

Numerical matters

B.1 Numerical solution of the refilling equation

To simplify the notation of the "refilling equation"

$$\frac{\mathrm{d}}{\mathrm{d}t}n_b(t,E_b) = N\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t,E_s) + \int_0^t \mathrm{d}t'\,\frac{n_b(t',E_b)}{\tau_e(E_b)}\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t-t',E_b) - \frac{1}{\tau_e(E_b)}n_b(t,E_b)\,, \tag{B.1}$$

we make the following definitions:

$$\begin{aligned} y(t) &= n_b(t,E_b) \\ f(t) &= N\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t,E_s) \\ \lambda &= \frac{1}{\tau_e(E_b)} \\ g(t) &= \frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t-t',E_b)\,, \end{aligned} \tag{B.2}$$

where the label E_b has been dropped for the sake of brevity. (Recall that E_s is related to E_b by Equation 4.1.)

Stated in terms of these variables, the refilling equation is

$$\frac{\mathrm{d}}{\mathrm{d}t}y(t) = f(t) + \lambda\int_0^t y(t')g(t-t')\,\mathrm{d}t' - \lambda y(t)\,. \tag{B.3}$$

Now, if $f(t)$ and $g(t)$ are known, this equation can potentially be solved by Laplace transform methods. The Laplace transform is defined as

$$x(s) = \int_0^t x(t)e^{-st}\,\mathrm{d}t\ , \tag{B.4}$$

so that Equation B.3 becomes

$$\begin{aligned} sy(s) &= f(s) + \lambda y(s)g(s) - \lambda y(s) \\ y(s) &= \frac{f(s)}{s + \lambda[1 - g(s)]} \, . \end{aligned} \tag{B.5}$$

Depending on the form of $f(s)$ and $g(s)$, it may be possible to solve for $y(t)$ by computing the inverse Laplace transform of the right hand side of this expression. However, $f(t)$ and $g(t)$ are not known analytically, but are estimated by simulation. Therefore, Equation B.3 must be solved by numerical methods.

Equation B.3 is known as a "Volterra-type" integral equation, and we follow the techniques described in reference [133] in its solution. To that end, we seek an approximate solution to Equation B.3 over the uniform discrete domain

$$t_i = i\Delta t \, , \qquad i = 0, 1, 2, ..., N \, , \qquad \Delta t = \frac{t}{N} \, , \tag{B.6}$$

where the solutions at neighboring i are related by

$$\begin{aligned} y_{i+1} &= y_i + \Delta y_i \, , \qquad \Delta y_i = \frac{\mathrm{d}}{\mathrm{d}t} y(t_i) \Delta t \\ t_{i+1} &= t_i + \Delta t \, , \end{aligned} \tag{B.7}$$

and where $y_0 = 0$. We use the "trapezoid rule"—higher-order schemes are not straight forward, see Section 18.2 in [133]—to approximate the integral term, that is

$$\int_0^t y(t')g(t - t') \, \mathrm{d}t' \to \Delta t \left(\frac{1}{2} y_0 g_i + \sum_{j=1}^{i-1} y_j g_{i-j} + \frac{1}{2} y_i g_0 \right) \, . \tag{B.8}$$

Equation B.3 then becomes

$$\Delta y_i = \Delta t \left(f_i + \lambda \frac{1}{2} y_0 g_i + \sum_{j=1}^{i-1} y_j g_{i-j} + \frac{1}{2} y_i g_0 - \lambda y_i \right) \, . \tag{B.9}$$

The set of y_i are then built up by recursively evaluating this expression. As a check of this procedure, we study some simple examples.

We imagine an ideal case, in which no UCN are lost from the system, which consists of the source, the guides, and the bottle. Further suppose that the source and bottle sit at opposite ends of a tube, and that UCN released in a burst from the source reach the bottle with times clustered around T. Let it also be

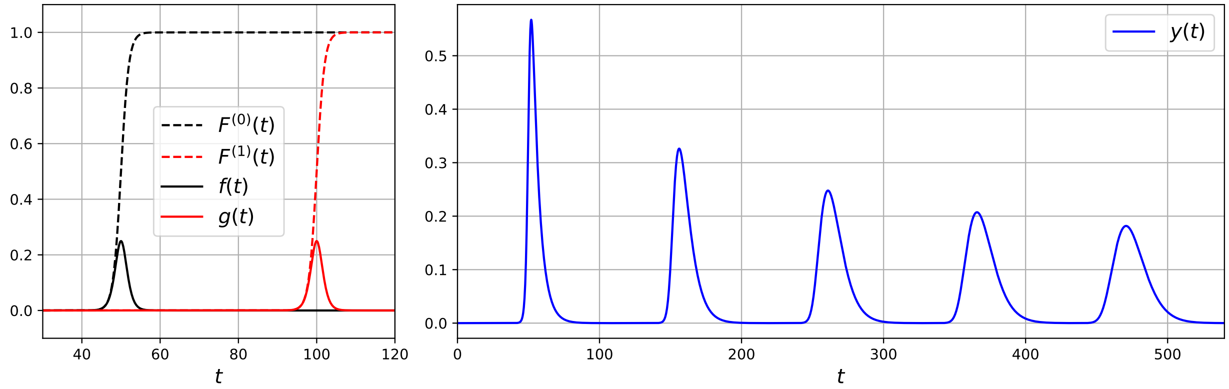

Figure B.1: Left: example $F^{(0)}(t)$ and $F^{(1)}(t)$, and corresponding $f(t)$ and $g(t)$ for $N = 1$ and $T = 50\,\mathrm{s}$. Right: numerical solution of Equation B.3 for $y(t)$ with $\lambda = 1/(5\,\mathrm{s})$. See the text.

the case that UCN exiting the bottle return with times clustered around $2T$. This qualitative behavior is described by the following functions:

$$\begin{aligned} F^{(0)}(t) &= \frac{1}{1+\exp[-(t-T)]} \\ F^{(1)}(t) &= \frac{1}{1+\exp[-(t-2T)]} \,. \end{aligned} \tag{B.10}$$

The plot at left in Figure B.1 shows these functions and their corresponding $f(t)$ and $g(t)$ for the case that $N = 1$ and $T = 50\,\mathrm{s}$. The peak direct filling rate occurs at T, and the peak refilling rate occurs at $2T$. The plot at right in this figure shows the numerical solution to $y(t)$ for $\lambda = 1/(5\,\mathrm{s})$. This solution is a train of UCN "pulses" of equal area spaced by approximately $2T$. Now, recall that the instantaneous rate of UCN leaving the bottle is $n(t)/\tau_e$, or $\lambda y(t)$. Thus, integrating over a single UCN pulse and multiplying by λ will give $\lambda \int \mathrm{d}t'\, y(t') = N$ if no UCN have left the system. Indeed, in this example, we find that the area under each UCN pulse is $5\,\mathrm{s}$, so that $\lambda \int \mathrm{d}t'\, y(t') = 1$.

The plot at left in Figure B.2 shows $F^{(0)}(t)$ and $F^{(1)}(t)$, and the corresponding $f(t)$ and $g(t)$, for the case that $N = 1$ and $T = 20\,\mathrm{s}$. Again, the peak direct filling rate occurs at T, and the peak refilling rate occurs at $2T$. The plot at right in this figure shows the numerical solution to $y(t)$ for $\lambda = 1/(10\,\mathrm{s})$. The solution is now a train of overlapping UCN pulses. Now, on average, a UCN spends time λ^{-1} in the bottle, and time $2T$ in the guides. Therefore, for long filling times, the "steady-state" fraction in the bottle should be $y(t \to \infty) = N\lambda^{-1}/(2T + \lambda^{-1})$. Indeed, in this example, the dashed purple line shows $(10)/(40 + 10) = 0.2$, the steady-state value of $y(t)$.

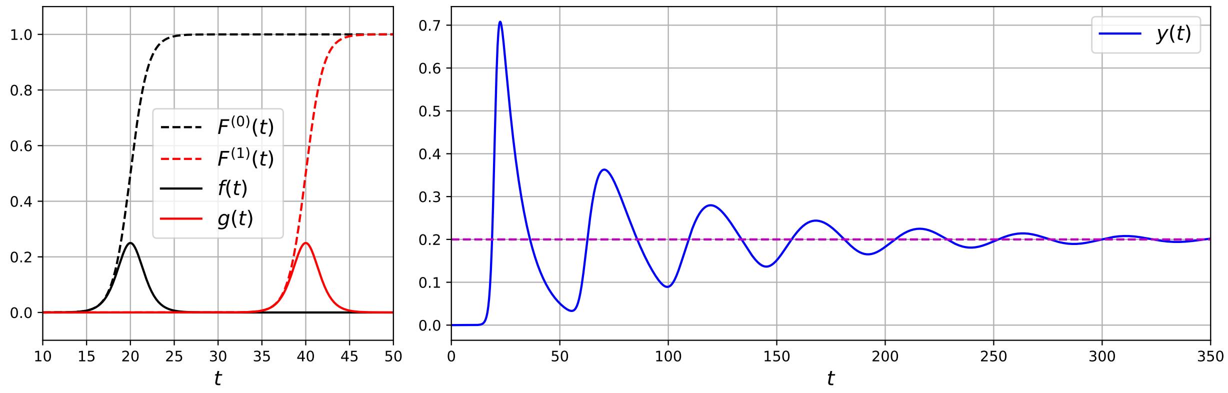

Figure B.2: Left: example $F^{(0)}(t)$ and $F^{(1)}(t)$, and corresponding $f(t)$ and $g(t)$ for $N = 1$ and $T = 20\,\mathrm{s}$. Right: numerical solution of Equation B.3 for $y(t)$ with $\lambda = 1/(10\,\mathrm{s})$. The dashed purple line shows the steady-state value of $y(t)$. See the text.

B.2 Extraction of time constants from simulation

In simulations of the storage time constant of the SUN-2 source, $\tau_s(E_s)$, and the emptying time constant of the storage bottle, $\tau_e(E_b)$, a set of times $\{t_i\}$ are generated, which are either the loss times of the UCN in the case of the former, or the detection times in the case of the latter. In both situations, the procedure to extract a time constant—and an uncertainty for that time constant—from $\{t_i\}$ is the same. We discuss the steps of this procedure in what follows.

One must first establish whether or not the $\{t_i\}$ data are exponentially distributed. This could potentially be accomplished by making a histogram of the data and fitting to it an exponential distribution. Then, for example, if the fit is very poor, one would conclude that the data are not exponentially distributed. Such an approach, however, would require an arbitrary choice of binning to be made for the histogram, which could potentially bias the results of the test in some unknown way. To avoid this problem, we use a modified version of the Kolmogorov-Smirnov (KS) test [128], originally due to Lilliefors [4], to test the distribution of the $\{t_i\}$. In a KS test, the maximum distance D_N between an empirical cumulative distribution function (CDF), $P_N(t)$, of the data of size N, and a hypothetical CDF $P(t)$, is calculated. The empirical CDF may be written as

$$P_N(t) = \frac{1}{N}\int_{-\infty}^{t} \mathrm{d}t' \sum_{\{t_i\}} \delta(t' - t_i)\,, \tag{B.11}$$

where $\delta(t)$ is the Dirac delta function. $P_N(t)$ is thus a "stepping function," increasing by 1 every time t passes one of the t_i. Intuitively, if $P(t)$ is the true distribution of the t_i, then as $N \to \infty$, one expects that $P_N(t) \to P(t)$, so that $D_N \to 0$. It may be shown that the quantity $K_N = \sqrt{N} D_N$ converges to the

Kolmogorov distribution, independent of the distribution $P(t)$ [134]. Therefore, one may either accept or reject the "null hypothesis"—that $P(t)$ is the true distribution of the t_i—by comparing K_N to critical values of the Kolmogorov distribution.

This is all well and good for a known distribution $P(t)$. However, the exponential CDF is

$$P(t) = \int_0^t \frac{1}{\tau}\exp(-t'/\tau)\,\mathrm{d}t' \;, \qquad t \geq 0\,, \tag{B.12}$$

and is parameterized by τ, which is not known *a priori*. Rather, an estimate of τ must be made, such as the maximum likelihood estimate (MLE), which is simply the mean of the $\{t_i\}$, that is $\hat{\tau} = \sum t_i/N$. The distribution of D_N when this particular $P(t)$ is estimated using the MLE $\hat{\tau}$ is considered via Monte Carlo methods in [4]. This reference also contains a set of tables of critical values to test whether or not a set of t_i are exponentially distributed for a given level of confidence. The results of this test applied to data of the SUN-2 storage simulations and the bottle emptying simulations are given in Tables B.1 and B.2, respectively. There is only one case in which the Lilliefors test indicates that the distribution is not exponential: for 1.0 neV UCN emptying from the box-style bottle in the DLC orientation. This case is marginal, and the critical value and D_N are roughly the same. We therefore disregard the results of this single test.

Tables B.1 and B.2 also give 1-σ confidence intervals for the $\hat{\tau}$, stated as "lower" and "upper" values (LB and UB), to be subtracted from or added to $\hat{\tau}$, respectively. These values are obtained as follows. Since the t_i are exponentially distributed, scaling them by the factor $2/\tau$ would give them a chi-squared distribution with $\nu = 2$ degrees of freedom. Therefore, the random variable $Z = 1/N \times \sum(2t_i/\tau) = 2\hat{\tau}/\tau$ has a chi-squared distribution with $\nu = 2N$ degrees of freedom. We may find symmetric bounds $[Z_{\mathrm{LB}}, Z_{\mathrm{UB}}]$ on Z at the α-level of significance (taken to be 1-σ) so that

$$P(Z_{\mathrm{LB}} \leq Z \leq Z_{\mathrm{UB}}) = 1 - \alpha\,, \tag{B.13}$$

where P is the CDF of a chi-squared distribution with $2N$ degrees of freedom. We may then solve for upper

Table B.1: Results of applying the Lilliefors test to the simulated SUN-2 storage time data: E_s is the energy of the UCN in the source with respect to the superfluid ^{4}He; N is the number of UCN simulated; $\hat{\tau}_s(E_s)$ is the MLE estimate of $\tau_s(E_s)$; LB gives the lower uncertainty to $\hat{\tau}_s(E_s)$ from the 1-σ confidence interval; UB gives the upper uncertainty; D_N is the KS statistic $K_N/\sqrt{N}$; the penultimate column lists the critical values at the $\alpha = 0.01$ significance level according to the tables in [4]; the last column gives the result of whether or not the data are exponentially distributed according to the Lilliefors test.

E_s [neV]	N	$\hat{\tau}_s(E_s)$ [s]	LB [s]	UB [s]	D_N	CV, $\alpha = 1\,\%$	Result
10	2000	670.3	14.6	15.2	0.018	0.028	True
20	2000	616.0	13.4	14.0	0.014	0.028	True
30	2000	560.7	12.2	12.8	0.015	0.028	True
40	2000	520.8	11.3	11.8	0.014	0.028	True
50	2000	463.9	10.1	10.6	0.015	0.028	True
60	2000	409.9	8.9	9.3	0.025	0.028	True
70	2000	349.8	7.6	8.0	0.011	0.028	True
80	2000	300.1	6.5	6.8	0.023	0.028	True
82	2000	289.0	6.3	6.6	0.017	0.028	True
84	2000	38.6	0.8	0.9	0.013	0.028	True
86	2000	7.5	0.2	0.2	0.009	0.028	True
88	2000	3.4	0.1	0.1	0.015	0.028	True

and lower bounds on τ as

$$\begin{gathered}
Z_{\mathrm{LB}} \leq Z \leq Z_{\mathrm{UB}} \\
\frac{Z_{\mathrm{LB}}}{2\hat{\tau}} \leq \frac{1}{\tau} \leq \frac{Z_{\mathrm{UB}}}{2\hat{\tau}} \\
\frac{2\hat{\tau}}{Z_{\mathrm{UB}}} \leq \tau \leq \frac{2\hat{\tau}}{Z_{\mathrm{LB}}} \\
\mathrm{LB} \leq \tau \leq \mathrm{UB}\,.
\end{gathered} \tag{B.14}$$

As can be seen in Tables B.1 and B.2, these bounds are very nearly symmetric, as should be the case for large N, in which the chi-squared distribution approaches a rescaled normal distribution.

B.3 Estimation of $\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t)$

Here we describe the procedure to estimate $\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t)$ and its uncertainty from $F^{(0)}(t)$. (The procedure for estimating $\frac{\mathrm{d}}{\mathrm{d}t}F^{(1)}(t)$ from $F^{(1)}(t)$ runs along similar lines.) As mentioned in Section 4.4, for the number of simulations carried out, computing the derivative of $F(t)$—the superscript is omitted for the sake of brevity—by finite differences produces an estimate of $\frac{\mathrm{d}}{\mathrm{d}t}F(t)$ that is noisy. (Roughly speaking, this occurs because the set of arrival times, $\{t_a\}$, from which $F(t)$ is estimated sometimes contains members that are

Table B.2: Results of applying the Lilliefors test to the simulated storage bottle emptying data: the first column gives the orientation of the storage bottle; E_b is the energy of the UCN in the storage bottle; N is the number of UCN simulated; $\hat{\tau}_e(E_b)$ is the MLE estimate of $\tau_e(E_b)$; LB gives the lower uncertainty to $\hat{\tau}_e(E_b)$ from the 1-σ confidence interval; UB gives the upper uncertainty; UB gives the upper bound; D_N is the KS statistic $K_N/\sqrt{N}$; the penultimate column lists the critical values at the $\alpha = 0.01$ significance level according to the tables in [4]; the last column gives the result of whether or not the data are exponentially distributed according to the Lilliefors test.

Orientation	E_b [neV]	N	$\hat{\tau}_e(E_b)$ [s]	LB [s]	UB [s]	D_N	CV, $\alpha = 1\%$	Result
CYTOP	1	4419	98.7	1.5	1.5	0.011	0.019	True
CYTOP	5	4726	48.1	0.7	0.7	0.008	0.018	True
CYTOP	10	4786	35.9	0.5	0.5	0.010	0.018	True
CYTOP	20	4861	26.5	0.4	0.4	0.008	0.018	True
CYTOP	30	4880	21.3	0.3	0.3	0.009	0.018	True
CYTOP	40	4884	18.8	0.3	0.3	0.008	0.018	True
CYTOP	50	4905	17.1	0.2	0.2	0.017	0.018	True
CYTOP	60	4921	15.7	0.2	0.2	0.008	0.018	True
CYTOP	70	4924	14.6	0.2	0.2	0.014	0.018	True
CYTOP	80	4938	13.3	0.2	0.2	0.015	0.018	True
dPE	1	4407	102.2	1.5	1.6	0.009	0.019	True
dPE	5	4690	48.9	0.7	0.7	0.016	0.018	True
dPE	10	4796	36.9	0.5	0.5	0.007	0.018	True
dPE	20	4858	26.2	0.4	0.4	0.009	0.018	True
dPE	30	4887	21.6	0.3	0.3	0.008	0.018	True
dPE	40	4891	18.5	0.3	0.3	0.011	0.018	True
dPE	50	4899	16.6	0.2	0.2	0.010	0.018	True
dPE	60	4907	15.8	0.2	0.2	0.007	0.018	True
dPE	70	4906	14.5	0.2	0.2	0.009	0.018	True
dPE	80	4928	13.5	0.2	0.2	0.009	0.018	True
DLC	1	3649	87.0	1.4	1.5	0.021	0.021	False
DLC	5	3896	41.4	0.6	0.7	0.019	0.020	True
DLC	10	3939	30.3	0.5	0.5	0.009	0.020	True
DLC	20	3963	21.5	0.3	0.3	0.013	0.020	True
DLC	30	4019	17.7	0.3	0.3	0.010	0.020	True
DLC	40	3980	16.0	0.2	0.3	0.013	0.020	True
DLC	50	3990	13.8	0.2	0.2	0.014	0.020	True
DLC	60	4006	12.9	0.2	0.2	0.008	0.020	True
DLC	70	4029	11.9	0.2	0.2	0.010	0.020	True
DLC	80	3988	10.9	0.2	0.2	0.011	0.020	True

very close.) The derivative calculated in this way can be smoothed by convolution with a "windowing" function, but the choice of such a function is somewhat arbitrary, and it also becomes necessary to correct for boundary effects "by hand".

One alternative method of estimating $\frac{\mathrm{d}}{\mathrm{d}t}F(t)$ involves making a histogram of the $\{t_a\}$, and dividing this by N_{sim}, the total number of UCN generated in the simulation at a given energy. The estimate between bins t_i and $t_i + \Delta t_i$ is then $\left(\frac{\mathrm{d}}{\mathrm{d}t}\hat{F}\right)_i = N_i/(\Delta t_i N_{\mathrm{sim}})$, where N_i is the number of counts in the bin. However, the choice of the number of bins, N_{bin}, and the bin widths, Δt_i, are arbitrary. We therefore compute multiple histograms of the same data set using different N_{bin} and random bin edges, and average the results.

To determine the appropriate weighting to be used in the averaging process, we consider the following simple case. The number of events, X, that occur in time T for a random process with rate R, is a Poisson random variable with mean and variance RT. Thus, the estimate of the rate from X, $\hat{R} = X/T$, has expectation value $\mathrm{E}(X/T) = RT/T = R$, and variance $\mathrm{Var}(X/T) = \mathrm{Var}(X)/T^2 = RT/T^2 = R/T$. An estimate of the uncertainty of $\hat{R}$ is thus $\sqrt{\hat{R}/T} = \sqrt{X}/T$. Applying this result to the present case gives the estimate of the uncertainty of $\left(\frac{\mathrm{d}}{\mathrm{d}t}\hat{F}\right)_i$ as $\sigma_i \equiv \sqrt{N_i}/\Delta t_i N_{\mathrm{sim}}$ for a single bin.

The process of computing multiple histograms with different N_{bin} and Δt_i will produce a collection of different estimates for $\frac{\mathrm{d}}{\mathrm{d}t}F(t)$ at a given time t. The weighted average of this collection is then calculated at the time t, using the weight $w_i = 1/\sigma_i$, that is

$$\frac{\mathrm{d}}{\mathrm{d}t}\hat{F}(t) = \frac{\sum_i w_i \left(\frac{\mathrm{d}}{\mathrm{d}t}\hat{F}\right)_i}{\sum_i w_i}. \tag{B.15}$$

The left hand side is the final estimate of $\frac{\mathrm{d}}{\mathrm{d}t}F(t)$, and the sums runs over all i for which $\left(\frac{\mathrm{d}}{\mathrm{d}t}\hat{F}\right)_i$ contributes at time t. A similar procedure is used to generate an estimate of the uncertainty of $\frac{\mathrm{d}}{\mathrm{d}t}\hat{F}(t)$:

$$\hat{\sigma}(t) = \sqrt{\frac{\sum_i w_i \left[\left(\frac{\mathrm{d}}{\mathrm{d}t}\hat{F}\right)_i - \frac{\mathrm{d}}{\mathrm{d}t}\hat{F}(t)\right]^2}{\sum_i w_i}}, \tag{B.16}$$

which is the weighted average of the square of the distances from $\frac{\mathrm{d}}{\mathrm{d}t}\hat{F}(t)$ to the various $\left(\frac{\mathrm{d}}{\mathrm{d}t}\hat{F}\right)_i$. Several examples of estimates of $\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t)$ calculated in this manner are displayed in Figure B.3.

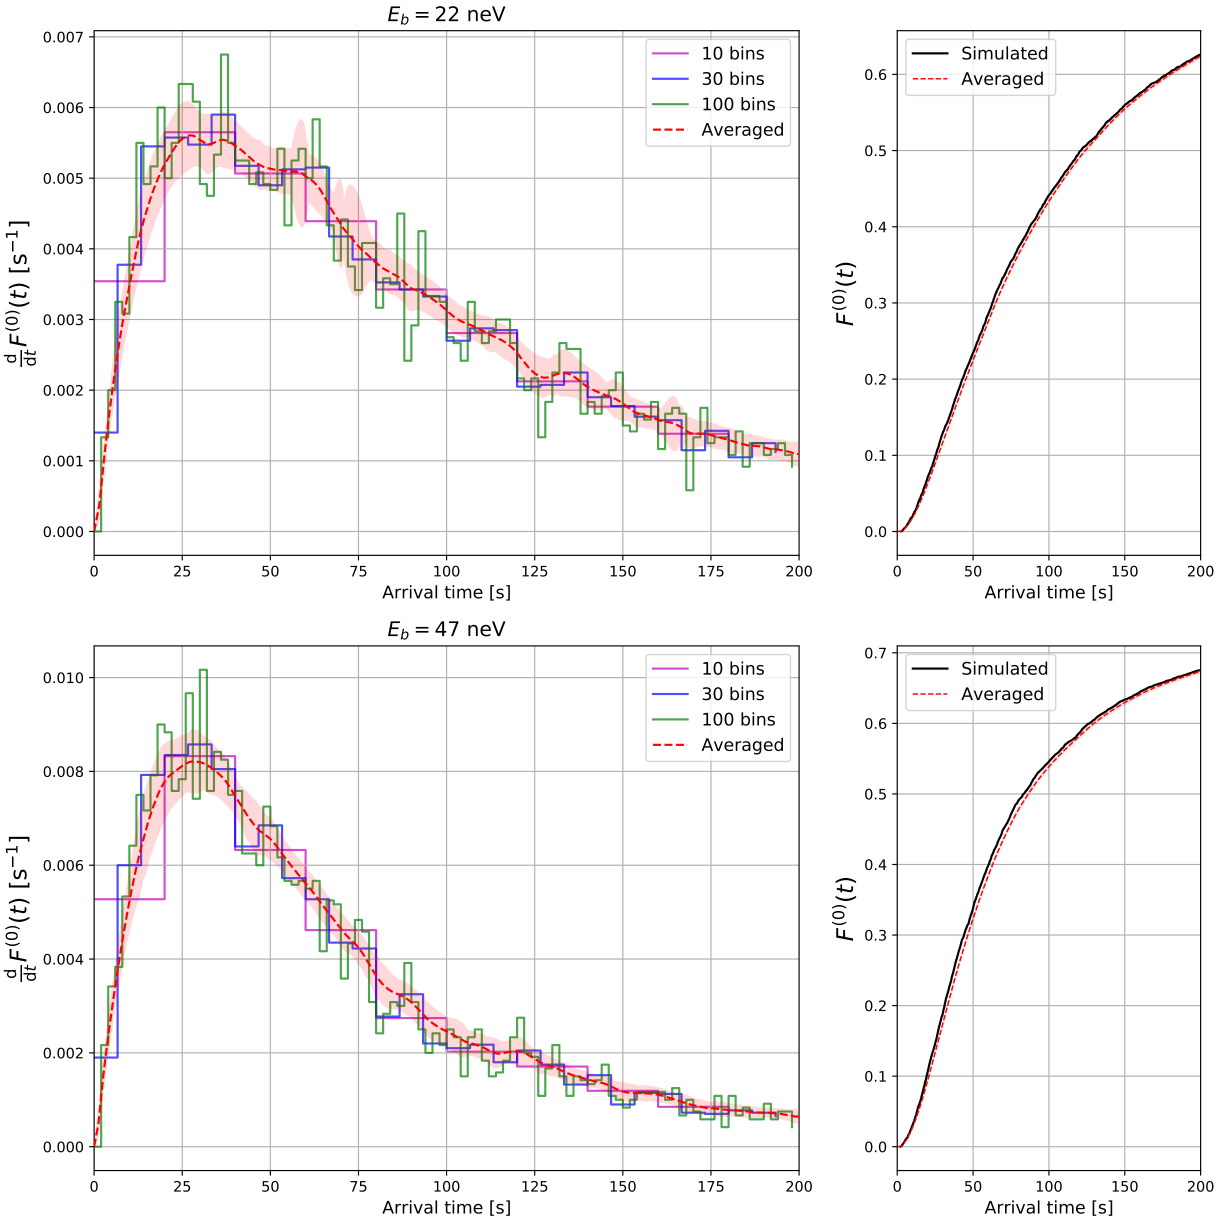

Figure B.3: Left: Two estimates, at $E_b = 22\,\mathrm{neV}$ and $47\,\mathrm{neV}$, of $\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t)$, shown in dashed red with 1-σ uncertainty bands. Histograms computed with equally spaced bin widths are plotted for comparison. Right: A comparison of $F^{(0)}(t)$ and the integral of the estimate of $\frac{\mathrm{d}}{\mathrm{d}t}F^{(0)}(t)$.

Appendix C

Diamond-like carbon measurements

Neutron reflectometry and UCN storage measurements were also performed on diamond-like carbon (DLC) coatings. However, these are not reported on in the main body of the present work because they do not constitute a complete data set.

The storage measurements were performed in a box-style bottle of inner dimensions 500 × 80 × 80 mm. (Labeled as (11) in Figure 3.1.) The bottle is assembled from six pieces: four 500 × 80 mm side plates, an 80 × 80 mm back plate, and an 80 × 80 mm front plate with a filleted 20 mm diameter hole in its center. At one point in time, the interior of the bottle was coated with a nickel-titanium "supermirror", and then coated with an approximately 1 µm thick layer of DLC. Measurements were then made in the bottle at cryogenic temperatures, which destroyed the DLC coating. The cracked coating was manually removed—the state of the supermirror coating underneath is not known—and the bottle was recoated with DLC by Argor-Aljba SA, based in Mendrisio, Switzerland. This coating is Argor-Aljba's "dialong", class "ta-C" DLC coating, which specifies a diamond content of greater than 80 %, and lists the density as $3.2\,\mathrm{g\,cm^{-3}}$. The bottle coating consists of a 0.9(3) µm (nominal) thick DLC layer over an adhesion layer of titanium, approximately 50 nm to 100 nm thick. The plug coating is of the same class, but does not have an adhesion layer underneath, and is of nominal thickness 1.0(2) µm.

As for the dPE storage data, the DLC storage data are not well described by a single time constant. Thus, the double exponential approximations, Equations 3.4 and 3.5, were used to fit the storage data. The use of the later was necessary for the single measurement at low temperature. (The construction of the bottle results in poor thermal contact between the edges of the plates, and the lowest bottle temperature is roughly 23 K with a 14 K temperature gradient between ends.) The results are listed in Table C.1, and two storage curves, taken while operating the source in accumulation mode are shown in Figures C.1 and C.2.

Only one reflectometry measurement exists for Argor-Aljba's DLC coating, performed during Cycle 188 on instrument D17 at the Institut Laue-Langevin. (See Section 3.2 for details of the instrument.) The sample belongs to Oliver Zimmer, and was prepared on a germanium wafer. Unfortunately, the sample quality was poor, and satisfactory reflectivity curves were not obtained. (See Figure C.3.)

Reflectometry measurements were also performed during Cycle 188 on DLC coatings from Diamond Hard Surfaces Ltd. (DHS), based in Towcester, United Kingdom. Unlike Argor-Aljba, this manufacturer possesses equipment capable of producing DLC coatings on the large storage cell electrodes to be used in the PanEDM experiment. (See the discussion at the end of Section 1.4.) Four samples were prepared on silicon wafers using DHS's "Adamant 010" preparation: two of nominal thickness 350(10) nm, and two of nominal thickness 75(10) nm. For each thickness, one sample is prepared using a hydrogen-based precursor, and the other using a deuterated precursor. The reflectivity curves are plotted in Figure C.4. The samples prepared with the deuterated precursor have a higher optical potential. Satisfactory fits could not be found for the thicker samples, but cursory fits to the thinner samples return optical potentials and thicknesses of $V \approx 208\,\mathrm{neV}$ and $t \approx 91\,\mathrm{nm}$ for the deuterated precursor (DP) sample, and $V \approx 134\,\mathrm{neV}$ and $t \approx 90\,\mathrm{nm}$ for the hydrogen-based precursor (HP) sample. A very rough comparison between the "critical edge" of the Argor-Aljba sample and those of the DHS samples would seem to indicate that $V \approx 200\,\mathrm{neV}$ for the former.

The density of DLC coatings can be determined from the measured optical potential via Equation 2.12. The 134 neV optical potential of the HP sample corresponds to a density of approximately $1.5\,\mathrm{g\,cm^{-3}}$, while the 208 neV optical potential of the DP sample corresponds to a density of approximately $2.4\,\mathrm{g\,cm^{-3}}$. (This assumes that both coatings contain only carbon.) The sp^3 fraction, k, of the DLC coating can be determined from the density via the relation [135].

$$\rho = 1.92 + 1.37k\,, \tag{C.1}$$

where ρ is in units of grams per cubic centimeter. Clearly, the HP sample must contain some degree of hydrogen, because even for fully-amorphous DLC ($k = 0$), the density should be at least $1.92\,\mathrm{g\,cm^{-3}}$. Considering the the DP sample was prepared in the same way, it seems likely that some degree of deuterium is present, so that a value of k estimated from Equation C.1 would be suspect.

Table C.1: Summary table of DLC storage measurements. The first column assigns a chronological index to the measurement. The "Date" column lists the start date and time of a measurement. "Front" indicates the entrance to the storage bottle: either the uncovered aluminum, "Alu.", or with the polished stainless steel insert, "SS.", in place. "Size" is the number of points in the data set. "Mode" is either continuous, "C", or accumulation, "A". "T" is the average bottle temperature during the measurement. All parameter estimates are least squares estimates. χ_r^2 gives the value of the reduced chi-squared statistic.

Index	Cycle	Date	Front	Size	Mode	T [K]	τ_1 [s]	τ_2 [s]	N_1	N_2	τ_d [h]	χ_r^2
0	189	21/03/23 23h	SS.	80	C	294.3	57(9)	201(14)	529(51)	304(45)	–	0.73
1	189	21/03/24 14h	SS.	29	A	295.1	84(15)	217(30)	1007(138)	492(186)	–	1.03
2	189	21/03/25 16h	SS.	55	A	22.6	130(23)	325(32)	1080(223)	951(258)	159(27)	1.13

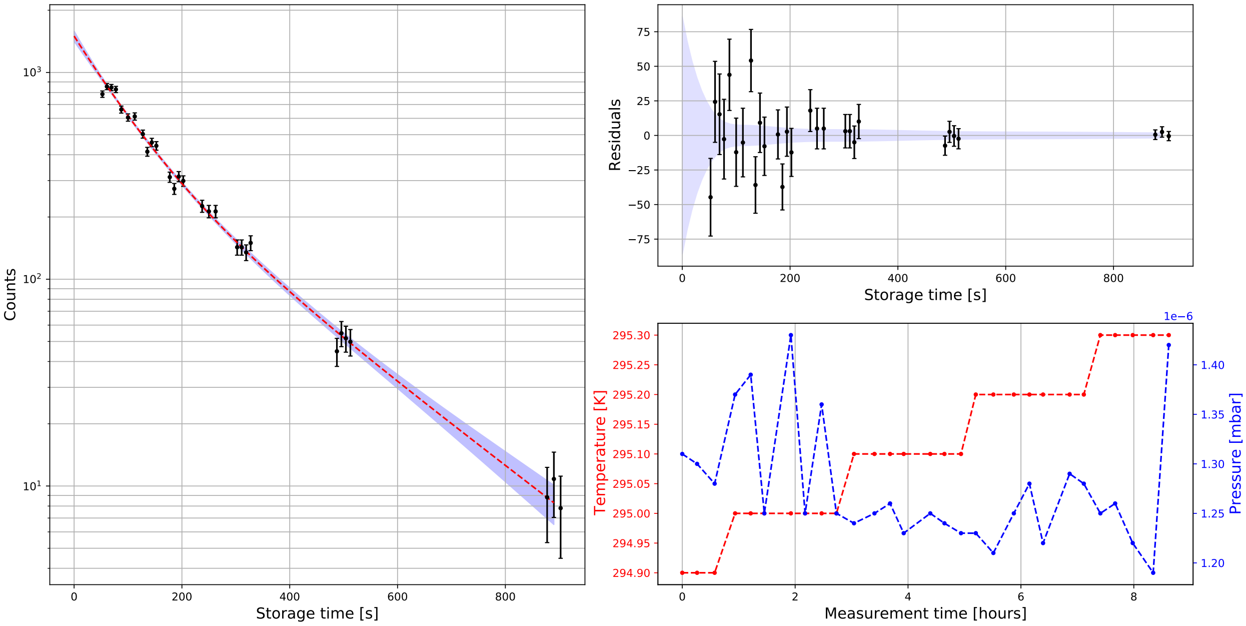

Figure C.1: Example DLC storage data set taken in accumulation mode during Cycle 189 at 295.1 K average and with the stainless steel entrance to the bottle. $N_1 = 1007(138)$, $\tau_1 = 84(15)$ s, $N_2 = 492(186)$, $\tau_2 = 217(30)$ s, $\chi^2/(29-4) = 1.03$. Left: Storage data are shown in black, and are displaced horizontally to show vertical spread. The fit is shown in red, with 1-σ error bands in light blue. (See Equation 3.4.) Top right: residuals from the plot at left. Bottom right: Average temperature and pressure during the measurement.

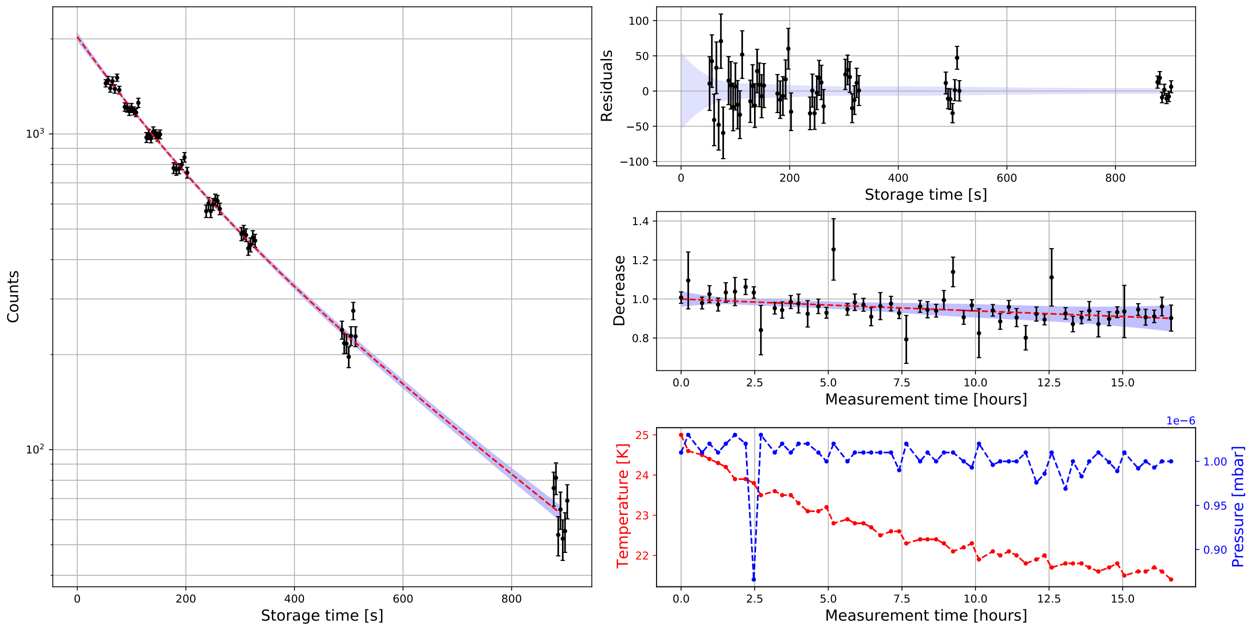

Figure C.2: Example DLC storage data set taken in accumulation mode during Cycle 189 at 22.6 K average and with the stainless steel entrance to the bottle. $N_1 = 1080(223)$, $\tau_1 = 130(23)\,\mathrm{s}$, $N_2 = 951(258)$, $\tau_2 = 325(32)\,\mathrm{s}$, $\tau_d = 159(27)\,\mathrm{h}$, $\chi^2/(55-5) = 1.13$. Left: Storage curve with the systematic decrease corrected for. (Equation 3.5 divided by the factor $\exp(-t_m/\tau_d)$.) Data are shown in black, and are displaced horizontally to show vertical spread. The fit is shown in red, with 1-σ error bands in light blue. Top right: residuals from the plot at left. Middle right: the systematic decrease. (Equation 3.5 divided by Equation 3.4.) Bottom right: Average temperature and pressure during the measurement.

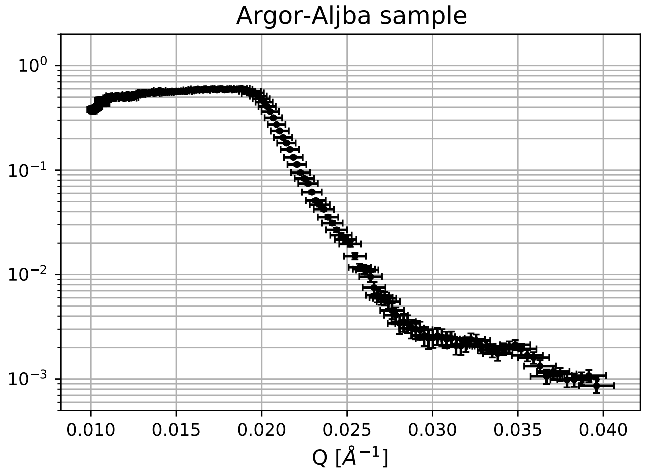

Figure C.3: Neutron reflectivity of DLC samples on germanium substrate, from Argor-Aljba SA, based in Mendrisio, Switzerland. Due to poor sample quality, a satisfactory reflectivity curve was not obtained. Horizontal error bars are the FWHM of the “resolution function” of the D17 instrument.

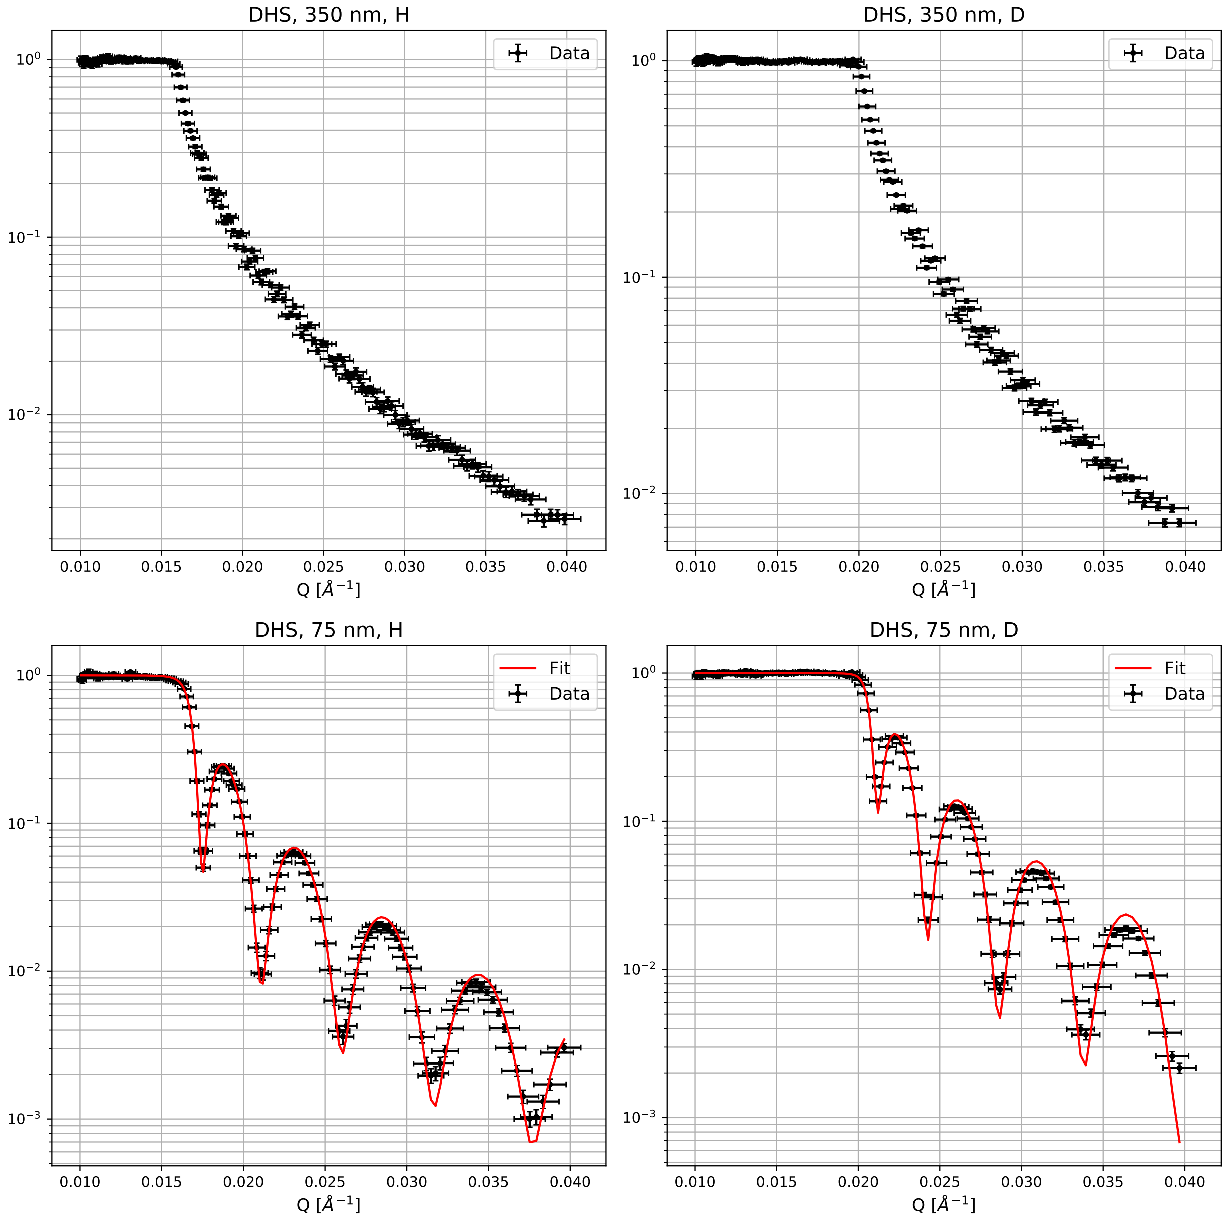

Figure C.4: Neutron reflectivity of four DLC thin film samples on silicon substrate. Samples were prepared by Diamond Hard Surfaces Ltd. (DHS), based in Towcester, United Kingdom. The titles list the nominal thicknesses specified by DHS, with the "H" and "D" indicating whether the sample was prepared using a hydrogen-based or deuterated precursor. See text. Satisfactory fits to the samples of nominal thickness 350 nm could not be found. Horizontal error bars are the FWHM of the "resolution function" of the D17 instrument.

Appendix D

The effect of gravity on the storage time constant

In light of the simulation results for the storage time constant of the SUN-2 source, in which the high energy "cutoff" is shifted by an amount that is related to the vertical dimensions of the production volume—see green dashed line in Figure 4.2 at right and the accompanying text in Section 4.2—it is necessary to investigate whether a model of $\tau(E)$ that accounts for gravitational effects, rather than that of Equation 2.16, should be used in fitting experimental storage data. The full expression for $\tau(E)$, taking into account gravity, is given by Equation 4.17 in reference [3], and is

$$\frac{1}{\tau(E_b)} = \frac{v_b}{4\gamma(E_b)} \int_{H_0}^{H_1^*} \mathrm{d}h \, \frac{E_b - mgh}{E_b} S(h) \bar{\mu}(E_b - mgh) + \frac{1}{\tau_\beta} \,, \tag{D.1}$$

written in terms of E_b, the UCN energy on the bottle axis. There are many quantities to define: h is a height coordinate inside the bottle; v_b is the UCN velocity on the bottle axis; $S(h)$ is the perimeter of the bottle interior at height h, obtained by calculating the length of the curve of intersection of the bottle walls and the plane normal to gravity at height h; $\bar{\mu}(E_b - mgh)$ is the wall loss probability, taking into account the variation of UCN kinetic energy with height, that is $E = E_b - mgh$; H_0 is the minimum height in the bottle; H_1 is the maximum height of the bottle, and H_1^* is the maximum height that the UCN can reach, which is H_1 if $E_b/mg \geq H_1$, and E_b/mg otherwise; $\gamma(E_b)$ is the effective volume accessible to the UCN, which is

$$\gamma(E_b) = \int_{H_0}^{H_1^*} \mathrm{d}h \, \mathcal{A}(h) \sqrt{\frac{E_b - mgh}{E_b}} \,, \tag{D.2}$$

where $\mathcal{A}(h)$ is obtained by calculating the area lying inside the curve of intersection of the bottle walls and the plane normal to gravity at height h. In the limit where gravity is ignored, Equation D.1 reduces to Equation 2.16: the integral over $S(h)$ yields A, the interior bottle surface area, and $\gamma(E_b) \to \mathcal{V}$, the bottle volume. The factors of $\sqrt{(E_b - mgh)/E_b}$ come from the variation of the UCN velocity and density with height. (See Equation 2.11.) Finally, written as it is, Equation D.1 appears to diverge as $E_b \to 0$. However, the factor of E_b^{-1} in the integral is canceled by the product of $v_b = \sqrt{2E_b/m}$ and the factor of $\sqrt{E_b}$ from $\gamma^{-1}(E_b)$.

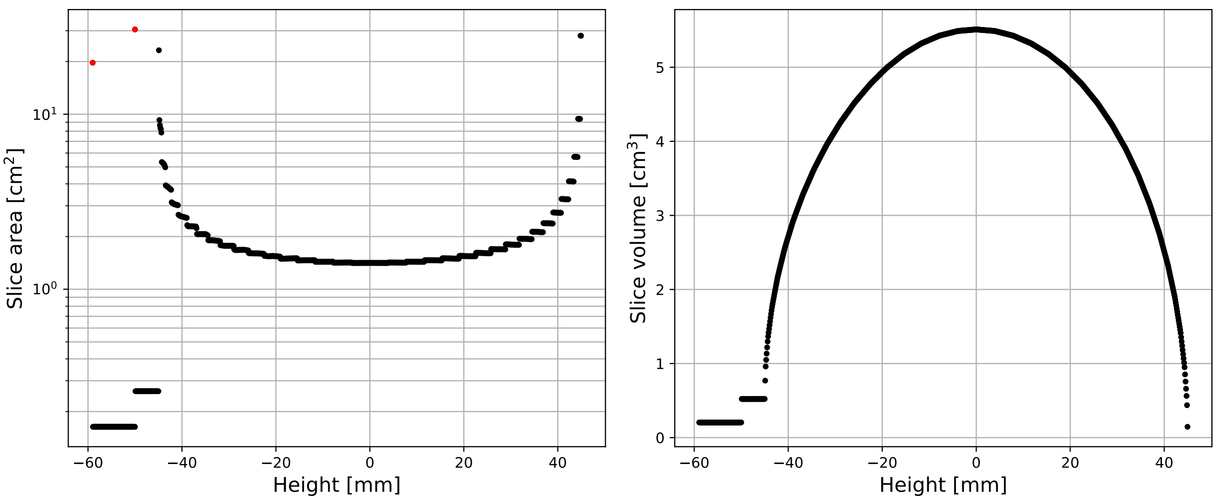

Figure D.1: Left: Plot of the "slice areas", $S(h)\Delta h$, for the bottle in the CYTOP orientation. $\Delta h = 0.1\,\mathrm{mm}$. The discontinuous points—red dots—correspond to the areas of the two downward facing surfaces; see blue shaded areas in Figure D.2. (N. B. the total surface area of the inside of the storage bottle, A, is not the integral of these data, which have units of area; rather, A is their sum.) Right: Plot of the "slice volumes", $\mathcal{A}(h)\Delta h$, for the bottle in the CYTOP orientation. $\Delta h = 0.1\,\mathrm{mm}$. (Also, $\mathcal{V} \approx \Sigma\mathcal{A}(h)\Delta h$.)

Examples of numerical approximations to $S(h)\,\mathrm{d}h$ and $\mathcal{A}(h)\,\mathrm{d}h$, for the inner surface of the storage bottle in the CYTOP orientation, are shown in Figure D.1. These were obtained by numerically computing the path lengths, $S(h)$, and enclosed areas, $\mathcal{A}(h)$, of curves formed by the intersection of the bottle inner surface and a set of planes normal to the direction of gravity and spaced by $\Delta h = 0.1\,\mathrm{mm}$. When calculated in this manner, the sums of the volume and area elements, $\Sigma S(h)\Delta h$ and $\Sigma\mathcal{A}(h)\Delta h$, differ by 0.1 % and 0.4 % from the exact values of A and $\mathcal{V}$, respectively.

Several examples of the storage time constant, $\tau(E_b)$, and $n_b(t, E_b)$, the time-dependent UCN energy spec-

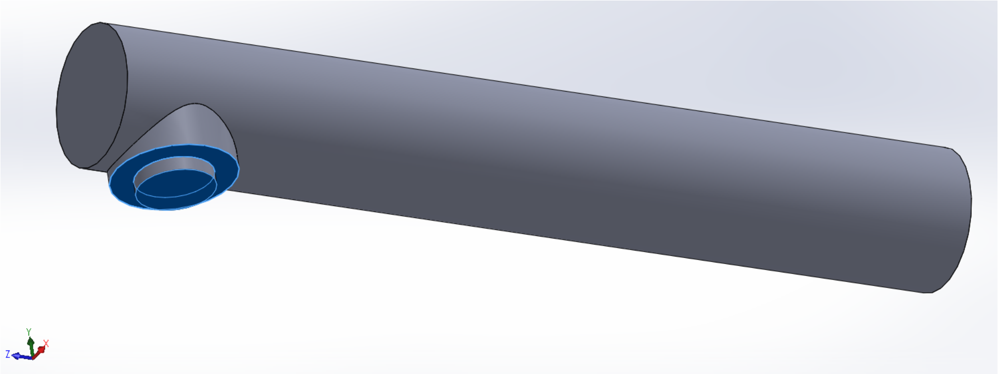

Figure D.2: Images of the inner bottle surface in the CYTOP orientation. Areas shaded in blue are downward facing, i.e. with surface normals colinear to the "Y" axis.

trum in the storage bottle, are plotted in Figure D.3 for different values of the loss factor, η, and using the optical potential $V = 115\,\mathrm{neV}$, as was measured for CYTOP via neutron reflectometry. (See Section 3.2.) In the plots at left, the storage time has been calculated in two ways: the dashed line shows $\tau(E_b)$ calculated using Equation D.1 and the numerical approximations to the bottle surface discussed in the previous paragraph, and the solid line shows $\tau(E_b)$ calculated from Equation 2.16. For the present discussion, let us refer to these as τ' and τ, respectively. The two values differ mainly at high and low energy. At high energy, for all values of η, the cutoff of τ' is roughly $6\,\mathrm{neV}$ below that of τ. (This latter quantity has a high energy cutoff of $V = 115\,\mathrm{neV}$, the optical potential of the walls.) This energy difference corresponds to mgH_0, where $H_0 = 59\,\mathrm{mm}$ is the distance between the bottle's axis and its lowest point. (See Figure D.2.) At low energy, τ approaches the neutron lifetime, τ_β, because in Equation 2.16, the wall loss rate vanishes at zero velocity. However, as is reflected in the plots of τ', some wall loss will always be present for UCN stored in a gravitational field, because even those with $E_b = 0$ will gain kinetic energy by falling.

The plots at right in Figure D.3 show the time-dependent UCN energy spectra, $n_b(t, E_b)$, calculated using the two models for $\tau(E_b)$, and using the initial UCN energy spectrum, $n_b(0, E_b)$, described in Chapter 4. The $n_b(t, E_b)$ plotted as dashed lines use τ', while the $n_b(t, E_b)$ plotted as solid line use τ. For the present discussion, let us refer to these as n' and n, respectively. For $\eta = 1 \times 10^{-5}$, there is no discernible difference between the two, while for $\eta = 1 \times 10^{-4}$, some deviation becomes visible by $500\,\mathrm{s}$. Now, the high energy cutoff of $n_b(0, E_b)$ is far below the high energy cutoffs of the τ' and τ, so that the disagreement between n' and n is mainly due to the difference between τ' and τ at low energy. However, because the spectral weight of $n_b(0, E_b)$ is relatively small at low energy range, the large $\sim 200\,\mathrm{sec}$ difference between τ' and τ at $E_b = 0$ only produces a relatively small difference in the n' and n.

In view of these observations, we will ignore the effects of gravity on UCN storage, and use the expressions developed in Chapter 2 in the fitting of experimental storage data. (The disagreement between n' and n is even milder for the dPE case: the optical potential, $V \approx 210\,\mathrm{neV}$—see Section 3.3, is even further away from the the high energy cutoff of $n_b(0, E_b)$; the bottle is also rotated by 135° on its axis compared to the CYTOP bottle orientation, resulting in a smaller effective surface to volume ratio for low energy UCN, and thus a smaller disagreement between τ' and τ at low energy.)

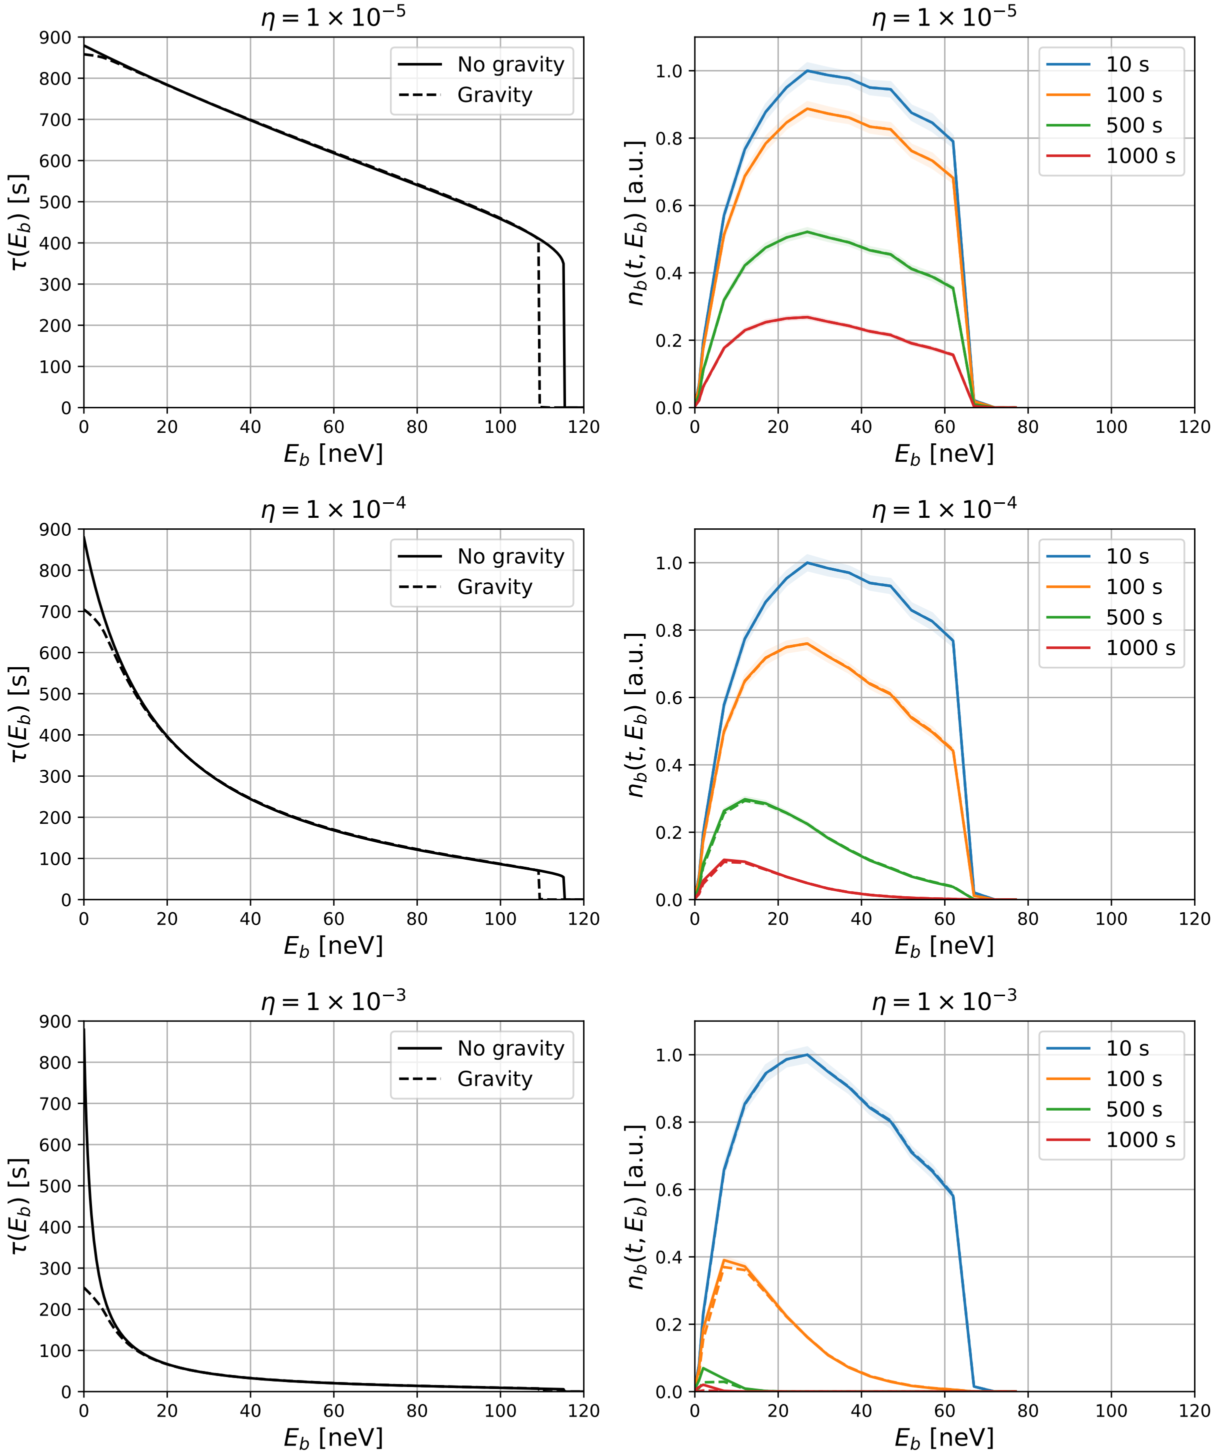

Figure D.3: Left: Plots of $\tau(E_b)$ for the storage bottle in the CYTOP orientation for different values of the loss factor, η. The solid lines and dashed lines correspond to $\tau(E_b)$ calculated via Equation 2.16 and Equation D.1, respectively. Right: Plots at various times of the UCN energy spectrum in the storage bottle, $n_b(t, E_b)$, for different values of the loss factor, η. Again, the solid and dashed lines corresponds to $\tau(E_b)$, calculated via Equation 2.16 and Equation D.1, respectively. The difference between the solid and dashed lines only becomes visible for large η and long times.

Appendix E

Limitations of the simulation method

As discussed in the beginning of Chapter 4, an unambiguous determination of the complete set of the parameters characterizing the experimental setup—including, for example, the loss factors, optical potentials, and roughnesses of each component—would require a very large number of additional measurements to be made. In the main body of this work, we therefore pursue an abbreviated version of such a "full" simulation, making several assumptions so that the initial UCN energy spectrum in the storage bottle, $n_b(0, E_b)$, may be estimated. As discussed in Chapter 5, this quantity is then used in Equation 5.1 in order to fit various storage models to experimental data, and thereby extract values for the loss factors, η, of CYTOP and dPE. There are, perhaps unsurprisingly, several limitations to this method of analysis. In this appendix, we discuss two unknown factors that are expected to have the largest influence on the shape of $n_b(0, E_b)$.

We will first consider the influence on $n_b(0, E_b)$ of the unknown characteristics of the guides between the source and the bottle. Simulations of direct filling and refilling in Chapter 4 assumed no wall losses in these guides, and that the probability of a non-specular reflection for a UCN in the guides is $p = 1\,\%$, independent of energy. Roughly speaking, UCN transported from the source to the bottle are briefly stored in the guides, so that a nonzero wall loss will tend to soften, or decrease the average energy of, the UCN spectrum that ends up in the storage bottle. Such softer spectra will, in turn, produce larger estimates of the loss factor of a bottle coating when used in fits to experimental storage data. The degree to which this will occur, however, is difficult to estimate.

In reference [96], a loss factor of $\eta = 3.5(1) \times 10^{-4}$ and a non-specular fraction of $p = 3.0(5)\,\%$ are reported for steel guides. These values were chosen in order to fit simulation outputs to measured UCN transmission data. The value of p found in this way, however, is specific to a certain distribution of UCN, and to the particular steel guides in which the transmission measurements were performed. The value of the loss factor is likewise specific to those particular steel guides. Without performing similar measurements, it is difficult to predict these values for the experimental setup described in this work.

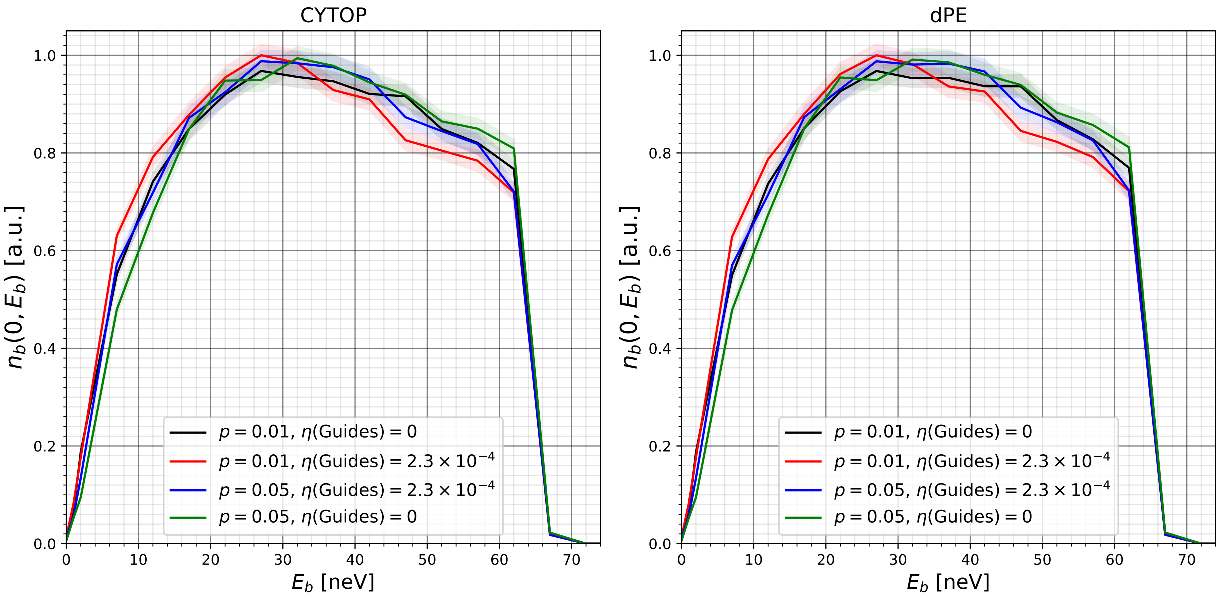

Figure E.1: Initial UCN energy spectra in the storage bottles, $n_b(0, E_b)$, calculated for the CYTOP and dPE orientations of the bottle, for the four possible combinations of $p = 1\,\%$ and $p = 5\,\%$, and $\eta(\text{Guides}) = 0$ and $\eta(\text{Guides}) = 2.3 \times 10^{-4}$. (See the text.) The difference between the CYTOP and dPE plots is very small.

Nevertheless, in this appendix, we report on simulations of direct filling—see Section 4.4—using the four possible combinations of $p = 1\,\%$ and $p = 5\,\%$, and $\eta(\text{Guides}) = 0$ and $\eta(\text{Guides}) = 2.3 \times 10^{-4}$. (The values of $p = 1\,\%$ and $\eta(\text{Guides}) = 0$ were used for the simulations reported on in Chapter 4.) The latter value of the loss factor is the average of that reported in reference [96] and the predicted loss factor—based on absorption cross sections [2]—of a hypothetical "304" stainless steel alloy: 70 % iron, 20 % chromium, and 10 % nickel. (Of course, there are other materials in the guides too, such as, for example, the copper block—see Section 4.2; however, in these calculations we use a single value for loss factor of all guide surfaces.) The same procedure described in Section 4.6 is then followed to generate $n_b(0, E_b)$ for the experimental filling time of 50 s. Plots of the four resulting spectra, for bottles in the CYTOP and dPE orientations, are shown in Figure E.1.

The list of loss factors for CYTOP and dPE, resulting from fits to experimental storage data with these four different spectra, and using the same weighted averaging procedure described in Chapter 5, are listed in Table E.1. As can be seen, for these values of p and η(Guides), the estimates of the loss factors for CYTOP and dPE are not changed by a significant amount. (In some cases, the estimates even decrease, although by an amount that is relatively small compared to the uncertainty.) However, without transmission

Table E.1: List of loss factors for CYTOP and dPE, resulting from fits to experimental storage data with the four different spectra shown in Figure E.1, and using the same weighted averaging procedure described in Chapter 5. (See Table 5.1.) .

η(Guides)	p	η(CYTOP), 11.7 K	η(CYTOP), 295.8 K	η(dPE), 11.9 K	η(dPE), 295.1 K
0	0.01	$(2.73 \pm 0.23) \times 10^{-5}$	$(11.05 \pm 0.93) \times 10^{-5}$	$(3.13 \pm 0.09) \times 10^{-5}$	$(26.14 \pm 1.45) \times 10^{-5}$
2.3×10^{-4}	0.01	$(2.83 \pm 0.22) \times 10^{-5}$	$(11.48 \pm 1.02) \times 10^{-5}$	$(3.24 \pm 0.09) \times 10^{-5}$	$(27.32 \pm 1.64) \times 10^{-5}$
2.3×10^{-4}	0.05	$(2.74 \pm 0.23) \times 10^{-5}$	$(11.04 \pm 0.89) \times 10^{-5}$	$(3.14 \pm 0.09) \times 10^{-5}$	$(26.06 \pm 1.17) \times 10^{-5}$
0	0.05	$(2.65 \pm 0.24) \times 10^{-5}$	$(10.60 \pm 0.78) \times 10^{-5}$	$(3.04 \pm 0.09) \times 10^{-5}$	$(24.85 \pm 0.71) \times 10^{-5}$

measurements of these guides, it is not clear whether such values of p and η(Guides) are appropriate.

Next, we consider the influence on $n_b(0, E_b)$ of the unknown storage properties of the UCN production volume in the SUN-2 source. In the simulations, the optical potential of the production volume walls was chosen to be that of Fomblin, 106 neV [52], while the loss factor was chosen to be slightly higher than that of Fomblin at room temperature. (These parameters were selected in order to produce a relatively soft UCN energy spectrum in the source, which would translate to a relatively soft initial UCN energy spectrum in the bottle, and therefore produce conservative estimates of η for the CYTOP and dPE coatings.) These choices result in a maximum UCN energy in the horizontal guides of approximately $E_{\max} \approx 106\,\mathrm{neV} - mg \times 356\,\mathrm{mm} \approx 70\,\mathrm{neV}$, where 356 mm is the vertical offset between the horizontal source and bottle axes. (See Section 4.1.) It is expected, however, that the optical potential of the walls of the production volume is effectively higher, resulting in a greater $E_{\max}$ in the horizontal guide section.

This expectation is based on two factors. Firstly, the Fomblin on the production volume walls is a coating that was applied manually over a beryllium layer, $V = 252\,\mathrm{neV}$ [3]. This means that, although a UCN with "perpendicular" energy, $E_\perp$, greater than 106 neV will be transmitted through the Fomblin layer, it may still be reflected from the beryllium underneath if $E_\perp \leq 252\,\mathrm{neV}$. ($E_\perp = E\cos^2\theta$, where θ is the angle between the surface normal to the wall and the UCN momentum at the point of interaction.) This will result in a smaller loss probability for higher energy UCN, and therefore, roughly speaking, an effectively larger optical potential than for the Fomblin coating alone. The details of this process, however, depend on the thickness of the Fomblin layer, which is not known due to its method of application.

The second indication that the optical potential of the production volume walls is effectively larger than that of Fomblin comes from previous measurements of the UCN "time of flight" (TOF) distributions, $f_t(t)$, for the source operated in continuous and accumulation modes. These are shown at top and bottom left, respectively, in Figure E.2. In these measurements, the beam chopper was separated from the detector by a 1 m long guide. Knowledge of this length, L, allows for a calculation of the "axial" velocity and energy distributions, $f_v(v)$ and $f_E(E)$, corresponding to the measured TOF data. (Here, axial means along the axis of the guide between the chopper and the detector; the variables are related by $vt = L$ and $E = \frac{1}{2}mv^2$.) However, these spectra have not yet been analyzed in depth, and precise information cannot be drawn from these plots. Notably, the zero of time may include an offset that is not corrected for in the data presented in Figure E.2. For example, there should be no data for UCN with axial energy below 54 neV, corresponding to the optical potential of the aluminum foil in front of the detector. (Or, equivalently, there should be no data for UCN with axial velocities below $3.2\,\mathrm{m\,s^{-1}}$, or with times of flight longer than 310 ms; the location of these cutoffs are indicated as the red dashed lines in Figure E.2.)

However, the gross features of the TOF data do seem to make sense: the average UCN energy is greater in continuous mode operation of the source. Although the precise energy values in the rightmost plots of Figure E.2 are suspect, both modes of operation seem to indicate the presence of a significant portion of UCN having axial energy within "plus or minus" 20 neV of 70 neV, higher than E_{max} in the simulations. (It should also be noted that the total UCN energy is, in general, larger than its "projection" along the axis of the guide between the chopper and the detector.) Therefore, it would seem that $E_{\mathrm{max}} = 70\,\mathrm{neV}$ is an underestimate, and that the true $n_b(0, E_b)$ extends to higher energy.

The two factors discussed in this Appendix—that is, the probable nonzero losses and higher UCN energies in the guides—are expected to bias the initial UCN energy spectrum in the bottle in opposite directions: lowering and raising the average energy of $n_b(0, E_b)$, respectively, and thus producing estimates of η for the CYTOP and dPE coatings that are larger and smaller, respectively. However, the degree of biasing for either effect is not known and is difficult to estimate. In the main body of this work, we therefore report values of η that are calculated using the $n_b(0, E_b)$ presented in Chapter 4.

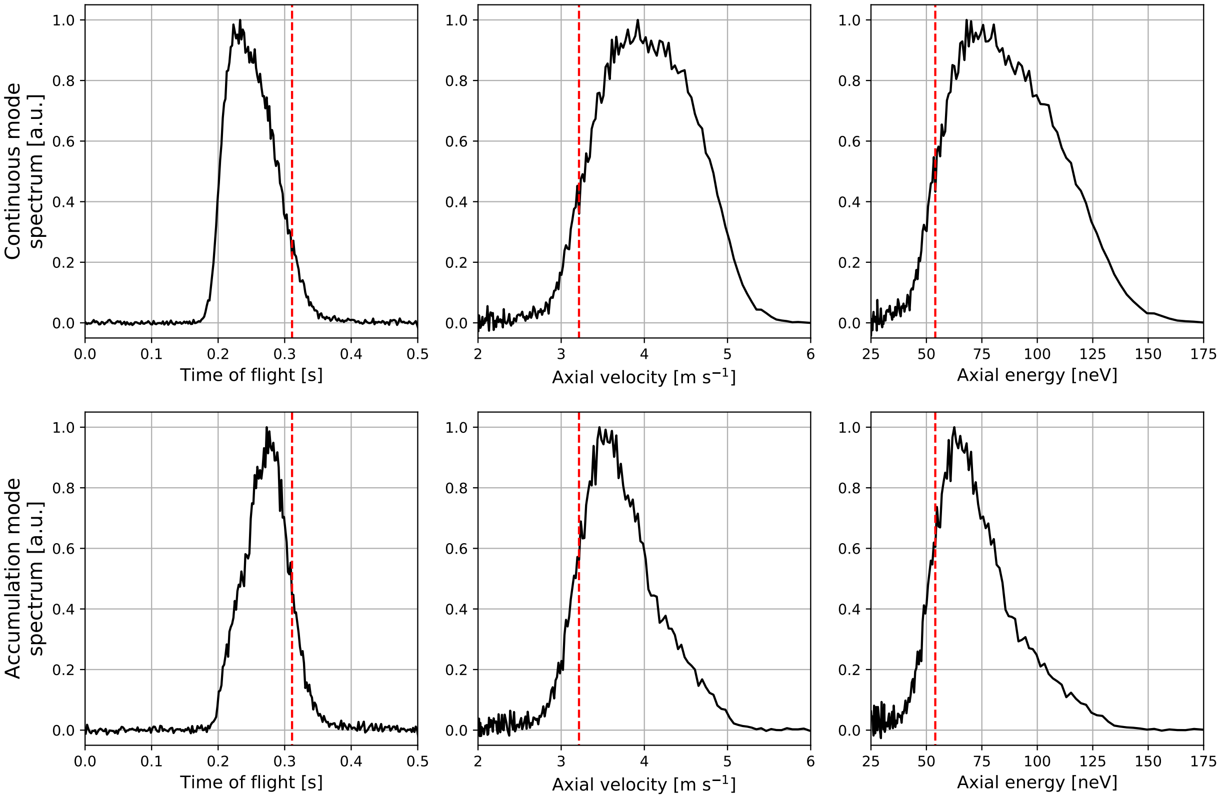

Figure E.2: Previously measured—and preliminary, meaning not to be quoted—"time of flight" (TOF) distributions, and derived axial velocity and energy distributions, at the output of SUN-2, for the source operated in continuous (top) and accumulation (bottom) modes. The beam chopper is separated from the detector by a 1 m long guide. The red dashed lines shows the location of the aluminum "cutoffs". (54 neV in the rightmost plots.) UCN should not be detected with "axial" energy below this value. (This also means that UCN with axial velocities greater than 3.2 m s^{-1}, or with times of flight less than 310 ms, should not be detected.)

References

[1] M Hainbuchner and E Jericha. Neutron Scattering Lengths, 2001. URL `http://www.ati.ac.at/~neutropt/scattering/table.html`. pages vii, 20, 42, 51

[2] V F Sears. Neutron scattering lengths and cross sections. *Neutron News*, 3(3):26–37, jan 1992. ISSN 1044-8632. doi: 10.1080/10448639208218770. URL `http://www.tandfonline.com/doi/abs/10.1080/10448639208218770`. pages vii, 13, 20, 42, 51, 72, 99, 129

[3] R Golub, D J Richardson, and S K Lamoreaux. *Ultra-Cold Neutrons.* CRC Press, jul 1991. ISBN 9780203734803. doi: 10.1201/9780203734803. URL `https://www.taylorfrancis.com/books/9781351406406`. pages vii, 14, 15, 17, 18, 19, 20, 22, 23, 25, 34, 35, 66, 70, 71, 91, 92, 99, 124, 130

[4] Hubert W. Lilliefors. On the Kolmogorov-Smirnov Test for the Exponential Distribution with Mean Unknown. *Journal of the American Statistical Association*, 64(325):387–389, mar 1969. ISSN 0162-1459. doi: 10.1080/01621459.1969.10500983. URL `http://www.tandfonline.com/doi/abs/10.1080/01621459.1969.10500983`. pages viii, 73, 113, 114, 115, 116

[5] G. L. Squires. *Introduction to the Theory of Thermal Neutron Scattering.* Cambridge University Press, Cambridge, 2012. ISBN 9781139107808. doi: 10.1017/CBO9781139107808. URL `http://ebooks.cambridge.org/ref/id/CBO9781139107808`. pages xv, 19, 34

[6] P A Zyla, R M Barnett, J Beringer, et al. Review of Particle Physics. *Progress of Theoretical and Experimental Physics*, 2020(8):1–2093, aug 2020. ISSN 2050-3911. doi: 10.1093/ptep/ptaa104. URL `https://academic.oup.com/ptep/article/doi/10.1093/ptep/ptaa104/5891211`. pages xvi, 3, 16, 45, 90, 93

[7] David Griffiths. *Introduction to Elementary Particles*, pages 11–53. Wiley, dec 1987. ISBN 9780471603863. doi: 10.1002/9783527618460. URL `https://onlinelibrary.wiley.com/doi/book/10.1002/9783527618460`. pages 1

[8] F. Abe, H. Akimoto, A. Akopian, et al. Observation of Top Quark Production in p p-bar Collisions with the Collider Detector at Fermilab. *Physical Review Letters*, 74(14):2626–2631, apr 1995. ISSN 0031-9007. doi: 10.1103/PhysRevLett.74.2626. URL `https://link.aps.org/doi/10.1103/PhysRevLett.74.2626`. pages 1

[9] K. Kodama, N. Ushida, C. Andreopoulos, et al. Observation of tau neutrino interactions. *Physics Letters B*, 504(3):218–224, apr 2001. ISSN 03702693. doi: 10.1016/S0370-2693(01)00307-0. URL `https://linkinghub.elsevier.com/retrieve/pii/S0370269301003070`. pages 1

[10] S. Chatrchyan, V. Khachatryan, A. M. Sirunyan, et al. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics*, 716(1):30–61, sep 2012. ISSN 03702693. doi: https://doi.org/10.1016/j.physletb.2012.08.021. pages 1

[11] Laurent Canetti, Marco Drewes, and Mikhail Shaposhnikov. Matter and antimatter in the universe. *New Journal of Physics*, 14(9), sep 2012. ISSN 1367-2630. doi: 10.1088/1367-2630/14/9/095012. pages 1

[12] Antonio Riotto. Theories of Baryogenesis, 1998. URL `https://arxiv.org/abs/hep-ph/9807454`. pages 2

[13] A. D. SAKHAROV. VIOLATION OF CP-INVARIANCE, C-ASYMMETRY, AND BARYON ASYMMETRY OF THE UNIVERSE. In *In the Intermissions...*, volume 34, pages 84–87. WORLD SCIENTIFIC, nov 1998. doi: 10.1142/9789812815941_0013. URL `http://www.worldscientific.com/doi/abs/10.1142/9789812815941_0013`. pages 2

[14] R F Streater and Wightman A S. *PCT, Spin and Statistics, and All That.* Princeton University Press, 1989. URL `https://www.jstor.org/stable/j.ctt1cx3vcq`. pages 3

[15] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay. Evidence for the 2π Decay of the K02 Meson. *Physical Review Letters*, 13(4):138–140, jul 1964. ISSN 0031-9007. doi: 10.1103/PhysRevLett.13.138. URL `https://link.aps.org/doi/10.1103/PhysRevLett.13.138`. pages 3

[16] T E Chupp, P Fierlinger, M J Ramsey-Musolf, and J T Singh. Electric dipole moments of atoms, molecules, nuclei, and particles. *Reviews of Modern Physics*, 91(1):015001, jan 2019. ISSN 0034-6861. doi: 10.1103/RevModPhys.91.015001. URL `https://link.aps.org/doi/10.1103/RevModPhys.91.015001`. pages 3, 5, 9

[17] M. V. Romalis, W. C. Griffith, J. P. Jacobs, and E. N. Fortson. New Limit on the Permanent Electric Dipole Moment of 199Hg. *Physical Review Letters*, 86(12):2505–2508, mar 2001. ISSN 0031-9007. doi: 10.1103/PhysRevLett.86.2505. URL `https://link.aps.org/doi/10.1103/PhysRevLett.86.2505`. pages 3

[18] M Kobayashi and T Maskawa. CP-Violation in the Renormalizable Theory of Weak Interaction. *Progress of Theoretical Physics*, 49(2):652–657, feb 1973. ISSN 0033-068X. doi: 10.1143/PTP.49.652. URL `https://academic.oup.com/ptp/article-lookup/doi/10.1143/PTP.49.652`. pages 3

[19] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the Unified Model of Elementary Particles. *Progress of Theoretical Physics*, 28(5):870–880, nov 1962. ISSN 0033-068X. doi: 10.1143/PTP.28.870. URL `https://academic.oup.com/ptp/article/28/5/870/1858382`. pages 3

[20] Ivan Esteban, M.C. Gonzalez-Garcia, Michele Maltoni, et al. The fate of hints: updated global analysis of three-flavor neutrino oscillations. *Journal of High Energy Physics*, 2020(9):178, sep 2020. ISSN 1029-8479. doi: 10.1007/JHEP09(2020)178. URL `http://link.springer.com/10.1007/JHEP09(2020)178`. pages 3

[21] Géraldine Servant. Baryogenesis from Strong CP Violation and the QCD Axion. *Physical Review Letters*, 113(17):171803, oct 2014. ISSN 0031-9007. doi: 10.1103/PhysRevLett.113.171803. URL `https://link.aps.org/doi/10.1103/PhysRevLett.113.171803`. pages 3

[22] Patrick Huet and Eric Sather. Electroweak baryogenesis and standard model CP violation. *Physical Review D*, 51(2):379–394, jan 1995. ISSN 0556-2821. doi: 10.1103/PhysRevD.51.379. URL `https://link.aps.org/doi/10.1103/PhysRevD.51.379`. pages 3

[23] M. Fukugita and T. Yanagida. Barygenesis without grand unification. *Physics Letters B*, 174(1): 45–47, jun 1986. ISSN 03702693. doi: 10.1016/0370-2693(86)91126-3. URL `https://linkinghub.elsevier.com/retrieve/pii/0370269386911263`. pages 3

[24] R. D. Peccei and Helen R. Quinn. CP Conservation in the Presence of Pseudoparticles. *Physical Review Letters*, 38(25):1440–1443, jun 1977. ISSN 0031-9007. doi: 10.1103/PhysRevLett.38.1440. URL `https://link.aps.org/doi/10.1103/PhysRevLett.38.1440`. pages 3

[25] John Ellis, Jae Sik Lee, and Apostolos Pilaftsis. Electric dipole moments in the MSSM reloaded. *Journal of High Energy Physics*, 2008(10):049–049, oct 2008. ISSN 1029-8479. doi: 10.1088/1126-6708/2008/10/049. URL `http://stacks.iop.org/1126-6708/2008/i=10/a=049?key=crossref.b9599a4600ef218f578b8856bdae1f25`. pages 3

[26] Eugene D. Commins. Electric Dipole Moments of Elementary Particles, Nuclei, Atoms, and Molecules. *Journal of the Physical Society of Japan*, 76(11):111010, nov 2007. ISSN 0031-9015. doi: 10.1143/JPSJ.76.111010. URL `http://journals.jps.jp/doi/10.1143/JPSJ.76.111010`. pages 4, 5

[27] Claude Cohen-Tannoudji, Bernard Diu, and Franck Laloë. *Quantum mechanics; 1st ed.* Wiley, New York, NY, 1977. URL `https://cds.cern.ch/record/101367`. pages 4

[28] R G Sachs. *The Physics of Time Reversal.* University of Chicago Press, 1987. ISBN 978-0-226-73331-9. pages 4

[29] Iosif B Khriplovich and Steve K Lamoreaux. *CP Violation Without Strangeness.* Springer Berlin Heidelberg, Berlin, Heidelberg, 1997. ISBN 978-3-642-64577-8. doi: 10.1007/978-3-642-60838-4. URL `http://link.springer.com/10.1007/978-3-642-60838-4`. pages 5

[30] L. I. Schiff. Measurability of Nuclear Electric Dipole Moments. *Physical Review*, 132(5):2194–2200, dec 1963. ISSN 0031-899X. doi: 10.1103/PhysRev.132.2194. URL `https://link.aps.org/doi/10.1103/PhysRev.132.2194`. pages 5

[31] J. P. Jacobs, W. M. Klipstein, S. K. Lamoreaux, et al. Limit on the electric-dipole moment of 199Hg using synchronous optical pumping. *Physical Review A*, 52(5):3521–3540, nov 1995. ISSN 1050-2947. doi: 10.1103/PhysRevA.52.3521. URL `https://link.aps.org/doi/10.1103/PhysRevA.52.3521`. pages 5

[32] Andrzej Czarnecki and Bernd Krause. Neutron Electric Dipole Moment in the Standard Model: Complete Three-Loop Calculation of the Valence Quark Contributions. *Physical Review Letters*, 78 (23):4339–4342, jun 1997. ISSN 0031-9007. doi: 10.1103/PhysRevLett.78.4339. URL `https://link.aps.org/doi/10.1103/PhysRevLett.78.4339`. pages 6

[33] Chien Yeah Seng. Reexamination of the standard model nucleon electric dipole moment. *Physical Review C - Nuclear Physics*, 91(2), feb 2015. ISSN 1089490X. doi: 10.1103/PhysRevC.91.025502. pages 6

[34] Charles P Slichter. *Principles of Magnetic Resonance*, volume 1 of *Springer Series in Solid-State Sciences*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1990. ISBN 978-3-642-08069-2. doi: 10.1007/978-3-662-09441-9. URL `http://link.springer.com/10.1007/978-3-662-09441-9`. pages 6, 17

[35] Eite Tiesinga, Peter J. Mohr, David B. Newell, and Barry N. Taylor. CODATA recommended values of the fundamental physical constants: 2018. *Reviews of Modern Physics*, 93(2):025010, jun 2021. ISSN 0034-6861. doi: 10.1103/RevModPhys.93.025010. URL `https://link.aps.org/doi/10.1103/RevModPhys.93.025010`. pages 7

[36] I. Altarev, M. Bales, D. H. Beck, et al. A large-scale magnetic shield with 10 6 damping at millihertz frequencies. *Journal of Applied Physics*, 117(18):183903, may 2015. ISSN 0021-8979. doi: 10.1063/1.4919366. URL `http://aip.scitation.org/doi/10.1063/1.4919366`. pages 7, 10

[37] S.M. Clayton, T.M. Ito, J.C. Ramsey, et al. Cavallo's multiplier for in situ generation of high voltage. *Journal of Instrumentation*, 13(05):P05017–P05017, may 2018. ISSN 1748-0221. doi: 10.1088/1748-0221/13/05/P05017. URL `https://iopscience.iop.org/article/10.1088/1748-0221/13/05/P05017`. pages 7

[38] Norman F. Ramsey. A Molecular Beam Resonance Method with Separated Oscillating Fields. *Physical Review*, 78(6):695–699, jun 1950. ISSN 0031-899X. doi: 10.1103/PhysRev.78.695. URL `https://link.aps.org/doi/10.1103/PhysRev.78.695`. pages 7

[39] F.M. Piegsa, B. van den Brandt, H. Glättli, et al. A Ramsey apparatus for the measurement of the incoherent neutron scattering length of the deuteron. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 589(2):318–329,

may 2008. ISSN 01689002. doi: 10.1016/j.nima.2008.02.020. URL `https://linkinghub.elsevier.com/retrieve/pii/S0168900208002544`. pages 7

[40] B. W. Filippone. Worldwide Search for the Neutron EDM. *arXiv:nucl-ex*, oct 2018. URL `http://arxiv.org/abs/1810.03718`. pages 8

[41] P. Schmidt-Wellenburg. The quest for an electric dipole moment of the neutron. In *AIP Conference Proceedings*, volume 1753. American Institute of Physics Inc., jul 2016. ISBN 9780735414112. doi: 10.1063/1.4955363. pages 8, 10

[42] J H Smith. *A Search for a Permanent Electric Dipole Moment of the Neutron.* PhD thesis, Harvard, 1951. pages 9

[43] J H Smith, E M Purcell, and N F Ramsey. Experimental Limit to the Electric Dipole Moment of the Neutron. *Physical Review*, 108(1):120–122, oct 1957. ISSN 0031-899X. doi: 10.1103/PhysRev.108.120. URL `https://link.aps.org/doi/10.1103/PhysRev.108.120`. pages 9

[44] W B Dress, P D Miller, J M Pendlebury, et al. Search for an electric dipole moment of the neutron. *Physical Review D*, 15(1):9–21, jan 1977. ISSN 0556-2821. doi: 10.1103/PhysRevD.15.9. URL `https://link.aps.org/doi/10.1103/PhysRevD.15.9`. pages 9

[45] Florian M. Piegsa. New concept for a neutron electric dipole moment search using a pulsed beam. *Physical Review C*, 88(4):045502, oct 2013. ISSN 0556-2813. doi: 10.1103/PhysRevC.88.045502. URL `https://link.aps.org/doi/10.1103/PhysRevC.88.045502`. pages 9

[46] Estelle Chanel, Zachary Hodge, Dieter Ries, et al. The pulsed neutron beam EDM experiment. *EPJ Web of Conferences*, 219:02004, dec 2019. ISSN 2100-014X. doi: 10.1051/epjconf/201921902004. URL `https://www.epj-conferences.org/10.1051/epjconf/201921902004`. pages 9

[47] J.M. Pendlebury, K.F. Smith, R. Golub, et al. Search for a neutron electric dipole moment. *Physics Letters B*, 136(5-6):327–330, mar 1984. ISSN 03702693. doi: 10.1016/0370-2693(84)92013-6. URL `https://linkinghub.elsevier.com/retrieve/pii/0370269384920136`. pages 9

[48] C. Abel, S. Afach, N. J. Ayres, et al. Measurement of the Permanent Electric Dipole Moment of the Neutron. *Physical Review Letters*, 124(8):81803, 2020. ISSN 10797114. doi: 10.1103/PhysRevLett.124.081803. URL `https://doi.org/10.1103/PhysRevLett.124.081803`. pages 9

[49] S. Afach, N.FIXMEJ. Ayres, C.FIXMEA. Baker, et al. Gravitational depolarization of ultracold neutrons: Comparison with data. *Physical Review D*, 92(5):052008, sep 2015. ISSN 1550-7998. doi: 10.1103/PhysRevD.92.052008. URL `https://link.aps.org/doi/10.1103/PhysRevD.92.052008`. pages 10

[50] D Wurm, D H Beck, T Chupp, et al. The PanEDM neutron electric dipole moment experiment at the ILL. *EPJ Web of Conferences*, 219:02006, dec 2019. ISSN 2100-014X. doi: 10.1051/epjconf/201921902006. URL `https://www.epj-conferences.org/10.1051/epjconf/201921902006`. pages 10, 17, 104

[51] O. Zimmer and R. Golub. Ultracold neutron accumulation in a superfluid-helium converter with magnetic multipole reflector. *Physical Review C*, 92(1):015501, jul 2015. ISSN 0556-2813. doi: 10.1103/PhysRevC.92.015501. URL `https://link.aps.org/doi/10.1103/PhysRevC.92.015501`. pages 10, 17, 104

[52] D J Richardson, J M Pendlebury, P Iaydjiev, et al. Measurement of the energy dependence of the neutron loss per bounce function on reflection from oil and grease surfaces using monochromatic ultracold neutrons. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 308(3):568–573, oct 1991. ISSN 01689002. doi: 10.1016/0168-9002(91)90069-3. URL `https://linkinghub.elsevier.com/retrieve/pii/0168900291900693`. pages 10, 34, 42, 70, 72, 130

[53] S. M. Cherniavsky, E. V. Lychagin, A. Yu. Muzychka, et al. Temperature dependence of the probability of "small heating" and total losses of UCNs on the surface of Fomblin oils of different molecular mass. *The European Physical Journal C*, 79(4):329, apr 2019. ISSN 1434-6044. doi: 10.1140/epjc/s10052-019-6835-z. URL `http://link.springer.com/10.1140/epjc/s10052-019-6835-z`. pages 10

[54] A. G. C. Chemicals Inc. Amorphous Fluoropolymer CYTOP, 2018. URL `https://www.agc-chemicals.com/file.jsp?id=jp/en/fluorine/products/cytop/download/pdf/CYTOP_EN_Brochure.pdf`. pages 11, 42

[55] Peter Fierlinger. *Losses and Depolarization of Stored Ultracold Neutrons on Diamond-like Carbon.* Phd, Universität Zürich, Zürich, 2005. pages 11, 17, 65

[56] F. Atchison, T. Bryś, M. Daum, et al. First storage of ultracold neutrons using foils coated with diamond-like carbon. *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics*, 2005. ISSN 03702693. doi: 10.1016/j.physletb.2005.08.066. pages 11

[57] F. Atchison, B. Blau, M. Daum, et al. Storage of ultracold neutrons in a volume coated with diamondlike carbon. *Physical Review C - Nuclear Physics*, 74(5):1–6, 2006. ISSN 05562813. doi: 10.1103/PhysRevC.74.055501. pages 11

[58] F. Atchison, B. Blau, M. Daum, et al. Diamondlike carbon can replace beryllium in physics with ultracold neutrons. *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics*, 642(1-2):24–27, nov 2006. ISSN 03702693. doi: 10.1016/j.physletb.2006.09.024. pages 11

[59] K. Bodek, M. Daum, R. Henneck, et al. Storage of ultracold neutrons in high resistivity, non-magnetic materials with high Fermi potential. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 597(2-3):222–226, dec 2008. ISSN 01689002. doi: 10.1016/j.nima.2008.09.018. URL `https://linkinghub.elsevier.com/retrieve/pii/S0168900208014381`. pages 11, 26, 52, 90, 105

[60] T Brenner, P Fierlinger, P Geltenbort, et al. Deuterated polyethylene coatings for ultra-cold neutron applications. *Applied Physics Letters*, 107(12):121604, sep 2015. ISSN 0003-6951. doi: 10.1063/1.4931388. URL `http://aip.scitation.org/doi/10.1063/1.4931388`. pages 11, 26, 42, 52, 90, 105

[61] Leonard Isaac Schiff. *Quantum mechanics.* International series in pure and applied physics. McGraw-Hill, New York, 3d ed. edition, 1968. pages 13, 19, 90

[62] Roger D. Woods and David S. Saxon. Diffuse Surface Optical Model for Nucleon-Nuclei Scattering. *Physical Review*, 95(2):577–578, jul 1954. ISSN 0031-899X. doi: 10.1103/PhysRev.95.577. URL `https://link.aps.org/doi/10.1103/PhysRev.95.577`. pages 13

[63] K Krane. *Introductory Nuclear Physics.* John Wiley & Sons, 3 edition, 1987. ISBN 978-0-471-80553-3. URL `https://www.wiley.com/en-au/Introductory+Nuclear+Physics{%}2C+3rd+Edition-p-9780471805533`. pages 13

[64] David J. Griffiths and Darrell F. Schroeter. *Introduction to Quantum Mechanics.* Cambridge University Press, aug 2018. ISBN 9781316995433. doi: 10.1017/9781316995433. URL `https://www.cambridge.org/core/product/identifier/9781316995433/type/book`. pages 13, 14, 15

[65] Enrico Fermi. Sul moto dei neutroni nelle sostanze idrogenate. *Ricerca scientifica*, 7(2):13–52, 1936. pages 15

[66] Stephan Paul. The puzzle of neutron lifetime. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 611(2-3):157–166, dec 2009. ISSN 01689002. doi: 10.1016/j.nima.2009.07.095. URL `https://linkinghub.elsevier.com/retrieve/pii/S0168900209015174`. pages 16, 17

[67] A P Serebrov, M E Chaikovskii, G N Klyushnikov, et al. Search for explanation of the neutron lifetime anomaly. *Physical Review D*, 103(7):074010, apr 2021. ISSN 2470-0010. doi: 10.1103/PhysRevD.103.074010. URL `https://link.aps.org/doi/10.1103/PhysRevD.103.074010`. pages 16

[68] D Dubbers, H Saul, B Märkisch, et al. Exotic decay channels are not the cause of the neutron lifetime anomaly. *Physics Letters B*, 791:6–10, apr 2019. ISSN 03702693. doi: 10.1016/j.physletb.2019.02.013. URL `https://linkinghub.elsevier.com/retrieve/pii/S0370269319301066`. pages 16

[69] V V Nesvizhevsky, H G Börner, A K Petukhov, et al. Quantum states of neutrons in the Earth's gravitational field. *Nature*, 415(6869):297–299, jan 2002. ISSN 0028-0836. doi: 10.1038/415297a. URL `http://www.nature.com/articles/415297a`. pages 17

[70] Loris Babin. *Characterisation of HOPE a magneto-gravitational trap for UCNs for neutron lifetime measurement on SUN2, a superthermal UCN source.* PhD thesis, Universite Grenoble Alpes, 2019. pages 17, 34, 80

[71] Herbert Goldstein. *Classical mechanics.* Addison Wesley, San Francisco, 3rd ed. edition, 2002. ISBN 0201316110. pages 18

[72] Yu.Yu. Kosvintsev, E.N. Kulagin, Yu.A. Kushnir, et al. Extraction of ultracold neutrons from high flux reactor SM-2. *Nuclear Instruments and Methods*, 143(1):133–137, may 1977. ISSN 0029554X. doi: 10.1016/0029-554X(77)90341-X. URL `https://linkinghub.elsevier.com/retrieve/pii/0029554X7790341X`. pages 18

[73] Glenn F Knoll. *Radiation detection and measurement* . John Wiley, Hoboken, N.J, 4th ed. edition, 2010. ISBN 9780470131480. pages 18, 35

[74] D.S. Sivia. *Elementary Scattering Theory.* Oxford University Press, jan 2011. ISBN 9780199228676. doi: 10.1093/acprof:oso/9780199228676.001.0001. URL `https://oxford.universitypressscholarship.com/view/10.1093/acprof:oso/9780199228676.001.0001/acprof-9780199228676`. pages 19

[75] A. Steyerl. Effect of surface roughness on the total reflexion and transmission of slow neutrons. *Zeitschrift für Physik A Hadrons and nuclei*, 254(2):169–188, apr 1972. ISSN 0939-7922. doi: 10.1007/BF01380066. URL `http://link.springer.com/10.1007/BF01380066`. pages 22, 71

[76] E H (Earle Hesse) Kennard. *Kinetic theory of gases, with an introduction to statistical mechanics.* International series in physics. McGraw-Hill Book Company, inc., New York, 1st ed. edition, 1938. pages 23

[77] V K (Vladimir Kazimirovich). Ignatovich. *The physics of ultracold neutrons* . Oxford series on neutron scattering in condensed matter. 5. Clarendon Press, Oxford, 1990. ISBN 0198510152. pages 25

[78] A. D. Stoika, A. V. Strelkov, and M. Hetzelt. Upscattering detected as the main reason for anomalous loss of ultra-cold neutrons in neutron storage experiments. *Zeitschrift fur Physik B Condensed Matter and Quanta*, 29(4):349–352, dec 1978. ISSN 0340-224X. doi: 10.1007/BF01324032. URL `http://link.springer.com/10.1007/BF01324032`. pages 25

[79] W. A. Lanford and R. Golub. Hydrogen Surface Contamination and the Storage of Ultracold Neutrons. *Physical Review Letters*, 39(24):1509–1512, dec 1977. ISSN 0031-9007. doi: 10.1103/PhysRevLett.39.1509. URL `https://link.aps.org/doi/10.1103/PhysRevLett.39.1509`. pages 25

[80] A. Steyerl, H. Nagel, F.-X. Schreiber, et al. A new source of cold and ultracold neutrons. *Physics Letters A*, 116(7):347–352, jun 1986. ISSN 03759601. doi: 10.1016/0375-9601(86)90587-6. URL `https://linkinghub.elsevier.com/retrieve/pii/0375960186905876`. pages 26

[81] P. Fierlinger, A. Pichlmaier, and H. Rauch. A time-of-flight chopper for ultracold neutrons. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 557(2):572–575, feb 2006. ISSN 01689002. doi: 10.1016/j.nima.2005.10.048. URL `https://linkinghub.elsevier.com/retrieve/pii/S0168900205018929`. pages 26

[82] V K Ignatovich. The Influence of Low-Frequency Oscillations on the Storage Time of Ultracold Neutrons. *physica status solidi (b)*, 71(2):477–486, oct 1975. ISSN 03701972. doi: 10.1002/pssb.2220710208. URL `https://onlinelibrary.wiley.com/doi/10.1002/pssb.2220710208`. pages 31, 98

[83] Bernard Jacrot. *Neutrons for Science.* Institut Laue-Langevin, Grenoble, 2018. URL `https://www.ill.eu/fileadmin/user_upload/ILL/1_About_ILL/Documentation/History/Jacrot-book-en/BookJacrot-Nov20.html`. pages 34

[84] M Chazalon and Y Lecuyer. Les circuits de refroidissement et l'installation de détritiation. *Bulletin d'Informations Scientifiques et Techniques*, 165:57–66, 1971. URL `https://wwwarchive.ill.fr/fileadmin/users_files/img/about_ill/history/founding_papers/BIST-165/BIST165-05-Chazalon.pdf`. pages 34

[85] Giovanna Cicognani and Mark Johnson. Institut Laue-Langevin Annual Report 2020. Technical report, Institut Laue-Langevin, Grenoble, 2020. URL `https://www.ill.eu/fileadmin/user_upload/ILL/1_About_ILL/Documentation/Annual_report/AR-20/Annual_report_2020.pdf`. pages 34

[86] P Ageron, J M Astruc, H Geipel, and Verdier J. La source .de neutrons froids du·réacteur à haut flux. *Bulletin d'Informations Scientifiques et Techniques*, 166:17–22, 1972. URL `https://wwwarchive.ill.fr/fileadmin/users_files/img/about_ill/history/founding_papers/BIST-166/BIST166-02-Ageron.pdf`. pages 34

[87] P. Schmidt-Wellenburg, K.H. Andersen, P. Courtois, et al. Ultracold-neutron infrastructure for the gravitational spectrometer GRANIT. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 611(2-3):267–271, dec 2009. ISSN 01689002. doi: 10.1016/j.nima.2009.07.096. URL `https://linkinghub.elsevier.com/retrieve/pii/S0168900209015411`. pages 34

[88] M. L. Goldberger and Frederick Seitz. Theory of the Refraction and the Diffraction of Neutrons by Crystals. *Physical Review*, 71(5):294–310, mar 1947. ISSN 0031-899X. doi: 10.1103/PhysRev.71.294. URL `https://link.aps.org/doi/10.1103/PhysRev.71.294`. pages 34

[89] F. M. Piegsa, M. Fertl, S. N. Ivanov, et al. New source for ultracold neutrons at the Institut Laue-Langevin. *Physical Review C*, 90(1):015501, jul 2014. ISSN 0556-2813. doi: 10.1103/PhysRevC.90.015501. URL `https://link.aps.org/doi/10.1103/PhysRevC.90.015501`. pages 34, 35

[90] Kent Leung. *Development of a new superfluid helium ultra-cold neutron source and a new magnetic trap for neutron lifetime measurements.* PhD thesis, Technical University of Munich, 2013. URL `https://www.researchgate.net/publication/308718316_Development_of_a_new_superfluid_helium_ultra-cold_neutron_source_and_a_new_magnetic_trap_for_neutron_lifetime_measurements/citation/download`. pages 34, 73

[91] K. K.H. Leung, S. Ivanov, F. M. Piegsa, et al. Ultracold-neutron production and up-scattering in superfluid helium between 1.1 K and 2.4 K. *Physical Review C*, 93(2):1–12, 2016. ISSN 24699993. doi: 10.1103/PhysRevC.93.025501. pages 34

[92] R Golub and J. M. Pendlebury. Super-thermal sources of ultra-cold neutrons. *Physics Letters A*, 53(2):133–135, jun 1975. ISSN 03759601. doi: 10.1016/0375-9601(75)90500-9. URL `https://linkinghub.elsevier.com/retrieve/pii/0375960175905009`. pages 34, 66

[93] R Golub and J M Pendlebury. The interaction of Ultra-Cold Neutrons (UCN) with liquid helium and a superthermal UCN source. *Physics Letters A*, 62(5):337–339, sep 1977. ISSN 03759601. doi: 10.1016/0375-9601(77)90434-0. URL `https://linkinghub.elsevier.com/retrieve/pii/0375960177904340`. pages 34, 66

[94] E. Korobkina, R. Golub, B.W. Wehring, and A.R. Young. Production of UCN by downscattering in superfluid He4. *Physics Letters A*, 301(5-6):462–469, sep 2002. ISSN 03759601. doi: 10.1016/S0375-9601(02)01052-6. URL `https://linkinghub.elsevier.com/retrieve/pii/S0375960102010526`. pages 34

[95] Frank Pobell. *Matter and Methods at Low Temperatures.* Springer Berlin Heidelberg, Berlin, Heidelberg, 2007. ISBN 978-3-540-46356-6. doi: 10.1007/978-3-540-46360-3. URL `http://link.springer.com/10.1007/978-3-540-46360-3`. pages 34, 72, 98

[96] A. Saunders, M. Makela, Y. Bagdasarova, et al. Performance of the Los Alamos National Laboratory spallation-driven solid-deuterium ultra-cold neutron source. *Review of Scientific Instruments*, 84(1): 013304, jan 2013. ISSN 0034-6748. doi: 10.1063/1.4770063. URL `http://aip.scitation.org/doi/10.1063/1.4770063`. pages 35, 98, 99, 128, 129

[97] O Zimmer. Unpublished, 2021. pages 35

[98] A Pichlmaier, V Varlamov, K Schreckenbach, and P Geltenbort. Neutron lifetime measurement with the UCN trap-in-trap MAMBO II. *Physics Letters B*, 693(3):221–226, oct 2010. ISSN 03702693. doi: 10.1016/j.physletb.2010.08.032. URL `https://linkinghub.elsevier.com/retrieve/pii/S0370269310009792`. pages 42

[99] O Zimmer, P Schmidt-Wellenburg, M Fertl, et al. Ultracold neutrons extracted from a superfluid-helium converter coated with fluorinated grease. *The European Physical Journal C*, 67(3-4):589–599, jun 2010. ISSN 1434-6044. doi: 10.1140/epjc/s10052-010-1327-1. URL `http://link.springer.com/10.1140/epjc/s10052-010-1327-1`. pages 42

[100] Yu.N. Pokotilovski. UCN anomaly and the possibility for further decreasing neutron losses in traps. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 554(1-3):356–362, dec 2005. ISSN 01689002. doi: 10.1016/j.nima.2005.07.026. URL `https://linkinghub.elsevier.com/retrieve/pii/S0168900205014634`. pages 42, 72

[101] J Hingerl. *A Versatile Device for Studying Ultracold Neutrons and Testing Cryogenic Storage Volumes.* Master's thesis, Technical University of Munich, jan 2019. pages 42

[102] E Connolly. ILL Internship Report: SUNino experiment preparation. Technical report, Institut Laue-Langevin, Grenoble, mar 2020. pages 42, 52, 53, 54

[103] D Ruhstorfer. *A setup for coating ultracold neutron guides with deuterated polyethylene.* Bachelor's thesis, Technical University of Munich, apr 2014. pages 42

[104] D Windmayer. *Using Deuterated Polyethylene Coatings for Manipulating Ultra-Cold-Neutrons in Different Applications.* Master's thesis, Technical University of Munich, jun 2016. URL `https://wiki.panedm.org/_media/nedm_theses/diplomarbeit_dpe_coatings_final_formatiert.pdf`. pages 42

[105] T Saerbeck, R Cubitt, A Wildes, et al. Recent upgrades of the neutron reflectometer D17 at ILL. *Journal of Applied Crystallography*, 51(2):249–256, apr 2018. ISSN 1600-5767. doi: 10.1107/S160057671800239X. URL `http://scripts.iucr.org/cgi-bin/paper?S160057671800239X`. pages 43

[106] S Degenkolb, P Nordin, and T Saerbeck. Measurements of SLD and sample characterization for thin films, 2021. URL `https://doi.ill.fr/10.5291/ILL-DATA.TEST-3178`. pages 43

[107] P Gutfreund, T Saerbeck, M Gonzalez, et al. Towards generalized data reduction on a chopper-based time-of-flight neutron reflectometer. *Journal of Applied Crystallography*, 51(3):606–615, jun 2018. ISSN 1600-5767. doi: 10.1107/S160057671800448X. URL `http://scripts.iucr.org/cgi-bin/paper?S160057671800448X`. pages 43

[108] Andrew Robert John Nelson and Charles David Dewhurst. Towards a detailed resolution smearing kernel for time-of-flight neutron reflectometers. Corrigendum. *Journal of Applied Crystallography*, 47 (3):1162–1162, jun 2014. ISSN 1600-5767. doi: 10.1107/S1600576714009595. URL `http://scripts.iucr.org/cgi-bin/paper?S1600576714009595`. pages 43

[109] SciPy: open-source software for mathematics, science, and engineering. URL https://scipy.org/. pages 45, 55, 93

[110] J. A. Nelder and R. Mead. A Simplex Method for Function Minimization. *The Computer Journal*, 7 (4):308–313, jan 1965. ISSN 0010-4620. doi: 10.1093/comjnl/7.4.308. URL https://academic.oup.com/comjnl/article-lookup/doi/10.1093/comjnl/7.4.308. pages 45

[111] A G Frodesen. *Probability and statistics in particle physics* . Universitetsforl., Bergen, 1979. ISBN 8200019063. pages 46

[112] Philip R Bevington. *Data reduction and error analysis for the physical sciences* . McGraw-Hill, New York, 2nd ed. edition, 1992. ISBN 0079112439. pages 46, 94

[113] Dusan Jeremic. Polyethylene. In *Ullmann's Encyclopedia of Industrial Chemistry*, pages 1–42. Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, Germany, aug 2014. ISBN 9780323358248. doi: 10.1002/14356007.a21_487.pub3. URL https://linkinghub.elsevier.com/retrieve/pii/B9780323358248000104http://doi.wiley.com/10.1002/14356007.a21_487.pub3. pages 52

[114] Allan P. Gray. Polymer crystallinity determinations by DSC. *Thermochimica Acta*, 1(6):563–579, dec 1970. ISSN 00406031. doi: 10.1016/0040-6031(70)80008-9. URL https://linkinghub.elsevier.com/retrieve/pii/0040603170800089. pages 52

[115] F G Kerry. *Industrial Gas Handbook: Gas Separation and Purification.* CRC Press, 2006. ISBN 0-8493-9005-2. URL http://www.pacificcrn.com/Upload/file/201612/12/20161212222434_36026.pdf. pages 52

[116] L Gataleta. dPE order 1076912/1. Technical report, Dorval, Quebec, may 2019. pages 52

[117] J. More, B. Garbow, and K. Hillstrom. User guide for MINPACK-1. [In FORTRAN]. Technical report, Argonne National Laboratory (ANL), Argonne, IL (United States), aug 1980. URL https://www.scipy.org/http://www.osti.gov/servlets/purl/6997568/. pages 55, 93

[118] Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168, jul 1944. ISSN 0033-569X. doi: 10.1090/qam/10666. URL http://www.ams.org/qam/1944-02-02/S0033-569X-1944-10666-0/. pages 55, 93

[119] Donald W. Marquardt. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, jun 1963. ISSN 0368-4245. doi: 10.1137/0111030. URL http://epubs.siam.org/doi/10.1137/0111030. pages 55, 93

[120] S Agostinelli, J Allison, K Amako, et al. GEANT4 - A simulation toolkit. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, jul 2003. ISSN 01689002. doi: 10.1016/S0168-9002(03)01368-8. URL https://www.sciencedirect.com/science/article/pii/S0168900203013688?casa_token=p-lFqT46g0sAAAAA:LtxHzTt3PUam_EP7_ATwLTr5S799dprfBit-n8rYoCBovCH6-tSMxQ-flT0L8haxeqFSIBUV4w. pages 65

[121] J Allison, K Amako, J Apostolakis, et al. Recent developments in GEANT4. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 835:186–225, nov 2016. ISSN 01689002. doi: 10.1016/j.nima.2016.06.125. URL https://www.sciencedirect.com/science/article/pii/S0168900216306957. pages 65

[122] F Atchison, T Bryś, M Daum, et al. The simulation of ultracold neutron experiments using GEANT4. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 552(3):513–521, nov 2005. ISSN 01689002. doi: 10.1016/j.nima.2005.06.065. URL https://linkinghub.elsevier.com/retrieve/pii/S0168900205013847. pages 65

[123] M Cohen and R Feynman. Theory of Inelastic Scattering of Cold Neutrons from Liquid Helium. *Physical Review*, 107(1):13–24, jul 1957. ISSN 0031-899X. doi: 10.1103/PhysRev.107.13. URL `https://link.aps.org/doi/10.1103/PhysRev.107.13`. pages 66

[124] R. Golub. On the storage of neutrons in superfluid 4He. *Physics Letters A*, 72(4-5):387–390, jul 1979. ISSN 03759601. doi: 10.1016/0375-9601(79)90505-X. URL `https://linkinghub.elsevier.com/retrieve/pii/037596017990505X`. pages 66, 72

[125] H. Yoshiki. The cross sections for one phonon emission and absorption by slow neutrons in superfluid liquid Helium. *Computer Physics Communications*, 151(2):141–148, mar 2003. ISSN 00104655. doi: 10.1016/S0010-4655(02)00819-6. URL `https://linkinghub.elsevier.com/retrieve/pii/S0010465502008196`. pages 66

[126] Frank L. Pedrotti, Leno M. Pedrotti, and Leno S. Pedrotti. *Introduction to Optics*. Cambridge University Press, dec 2017. ISBN 9781108428262. doi: 10.1017/9781108552493. pages 72

[127] H. Yoshiki, K. Sakai, M. Ogura, et al. Observation of ultracold-neutron production by 9-Å cold neutrons in superfluid helium. *Physical Review Letters*, 68(9):1323–1326, mar 1992. ISSN 0031-9007. doi: 10.1103/PhysRevLett.68.1323. URL `https://link.aps.org/doi/10.1103/PhysRevLett.68.1323`. pages 73

[128] Frank J. Massey. The Kolmogorov-Smirnov Test for Goodness of Fit. *Journal of the American Statistical Association*, 46(253):68–78, mar 1951. ISSN 0162-1459. doi: 10.1080/01621459.1951.10500769. URL `http://www.tandfonline.com/doi/abs/10.1080/01621459.1951.10500769`. pages 73, 113

[129] S. J. Seestrom, E. R. Adamek, D. Barlow, et al. Upscattering of ultracold neutrons from gases. *Physical Review C*, 92(6):065501, dec 2015. ISSN 0556-2813. doi: 10.1103/PhysRevC.92.065501. URL `https://link.aps.org/doi/10.1103/PhysRevC.92.065501`. pages 90, 97

[130] S. J. Seestrom, E. R. Adamek, D. Barlow, et al. Total cross sections for ultracold neutrons scattered from gases. *Physical Review C*, 95(1):015501, jan 2017. ISSN 2469-9985. doi: 10.1103/PhysRevC.95.015501. URL `https://link.aps.org/doi/10.1103/PhysRevC.95.015501`. pages 97

[131] C Kittel and H Kroemer. *Thermal physics* . W.H. Freeman, San Francisco, 2d ed. edition, 1980. ISBN 0716710889. pages 97

[132] R Pynn. Neutron scattering by rough surfaces at grazing incidence. *Physical Review B*, 45(2):602–612, jan 1992. ISSN 0163-1829. doi: 10.1103/PhysRevB.45.602. URL `https://link.aps.org/doi/10.1103/PhysRevB.45.602`. pages 109

[133] William H Press, Saul A Teukolsky, William T Vetterling, and Brian P Flannery. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, USA, 3 edition, 2007. ISBN 0521880688. pages 111

[134] Andrey Kolmogorov. Sulla determinazione empirica di una lgge di distribuzione. *Inst. Ital. Attuari, Giorn.*, 4:83–91, 1933. pages 114

[135] A. C. Ferrari, A. Libassi, B. K. Tanner, et al. Density, sp3 fraction, and cross-sectional structure of amorphous carbon films determined by x-ray reflectivity and electron energy-loss spectroscopy. *Physical Review B*, 62(16):11089–11103, oct 2000. ISSN 0163-1829. doi: 10.1103/PhysRevB.62.11089. URL `https://link.aps.org/doi/10.1103/PhysRevB.62.11089`. pages 120