COMPARISON OF METHODS FOR MODELING OCXO FREQUENCY AGING BEHAVIOR

by

SAMUEL J GRIFFIN

ANDREW LEMMON, COMMITTEE CHAIR
TODD FREEBORN
PAUL KOPPANG

A THESIS

Submitted in partial fulfillment of the requirements for the
degree of Master of Science in in the Department of
Electrical and Computer Engineering
in the Graduate School of
The University of Alabama

TUSCALOOSA, ALABAMA

2019

ABSTRACT

Frequency aging is an important factor in frequency standards. Aging is the portion of frequency drift that is intrinsic to the frequency standard. Improving the knowledge surrounding the mechanisms of aging provides means to minimize their impact. To study frequency aging further, several known models used to describe frequency aging were investigated. A set of oven-controlled oscillator datasets were obtained in order to evaluate two frequency aging models. The first model is a logarithmic model, which was selected due to its ubiquity. The second model is a linear Kalman Filter model, which was chosen for its dynamic capabilities. The results of each modeling procedure were analyzed using linear frequency aging baseline comparison, coefficients of determination, and residual analysis. It was determined that the linear Kalman Filter model provides better model fidelity than the logarithmic model. Additionally, the analysis found that the linear frequency aging baseline comparison and residual analysis provide a sufficient comparison suite, with the coefficient of determination adding little in terms of further insight.

DEDICATION

I dedicate this work to Emily Barbee, Bryan Owings, and Jamie McKelvey. I would also like to thank my committee for working with me on this thesis.

LIST OF ABBEVIATIONS AND SYMBOLS

XO	Crystal Oscillator
NIST	National Institute of Standards and Technology
UTC	Universal Coordinated Time
OCXO	Oven-Controlled Crystal Oscillator
TCXO	Temperature Compensated Crystal Oscillator
KF	Kalman Filter

ACKNOWLEDGEMENTS

I would like to acknowledge Bryan Owings and Microchip for providing the OCXO datasets for which this work is based upon. I also wish to acknowledge my committee for their contributions to this work.

CONTENTS

ABSTRACT ii

DEDICATION iii

LIST OF ABBEVIATIONS AND SYMBOLS iv

ACKNOWLEDGEMENTS v

LIST OF TABLES ix

LIST OF FIGURES x

CHAPTER 1: INTRODUCTION 1

CHAPTER 2: BACKGROUND AND LITERATURE REVIEW 13

2.1. Background 13

2.2. Literature Review 19

2.2.1. Simple Models 20

2.2.2. Compound Models 25

2.2.3. Adaptive Models 27

2.2.4. Comparison of Aging Models 38

2.3. Summary 39

CHAPTER 3: DATA COLLECTION AND METHODS 40

3.1. Microchip 1000B and 1000C OCXOs 40

3.2. OCXO Datasets: Short-term Measurement Process 42

3.3. OCXO Datasets: Long-term Measurement Process 43

3.4. Selection of Data for Analysis 46

3.5. Data Preprocessing 47
3.5.1. Dataset Gaps 47
3.5.2. Outlier Removal 47
3.5.3. Rebasing 49
3.6. Baseline Frequency Aging Estimation 50
3.7. Processed Datasets 51
3.7.1. OCXO A 52
3.7.2. OCXO B 54
3.7.3. OCXO H 56
3.8. Summary 58
CHAPTER 4: APPLICATION AND EVALUATION OF AGING MODELS 59
4.1. Logarithmic Model 59
4.1.1. Logarithmic Modeling Process 60
4.1.2. Logarithmic Modeling Results 63
4.1.2.1. Logarithmic Modeling: OCXO A 64
4.1.2.2. Logarithmic Modeling: OCXO B 66
4.1.2.3. Logarithmic Modeling: OCXO H 68
4.1.2.4. Logarithmic Modeling: Coefficient of Determination Summary 70
4.1.3. Logarithmic Modeling Discussion 71
4.2. Linear Kalman Filter Model 73
4.2.1. Linear Kalman Filter Modeling Process 73
4.2.1.1. Linear Kalman Filter: Covariance Calculation 73
4.2.1.2. Linear Kalman Filter: Initial Conditions 75

4.2.1.3. Linear Kalman Filter: Model Output 76
4.2.2. Linear Kalman Filter Modeling Results 77
4.2.2.1. Linear Kalman Filter Modeling: OCXO A 78
4.2.2.2. Linear Kalman Filter Modeling: OCXO B 80
4.2.2.3. Linear Kalman Filter Modeling: OCXO H 82
4.2.3. Linear Kalman Filter Modeling: Coefficient of Determination Summary 84
4.2.4. Linear Kalman Filter Discussion 85
4.3. Summary 87
CHAPTER 5: CONCLUSION 88
REFERENCES 91
APPENDIX 94
APPENDIX A. DATASETS 94
APPENDIX B. LOGARITHMIC MODEL RESULTS 104
APPENDIX C. LINEAR KALMAN FILTER MODEL RESULTS 112

LIST OF TABLES

1. Kalman Filter Components 30
2. Selected OCXOs 46
3. Baseline Frequency Aging Estimates 58
4. Logarithmic Model Coefficient of Determination Results 70
5. Power Law Noise Types 74
6. Power Law Noise Coefficients 75
7. Linear Kalman Filter Model Results 84

LIST OF FIGURES

1. Taxonomy of frequency standards [6][7] 3
2. LC Colpitts oscillator circuit [6] 4
3. Quartz resonator (top of canister removed) 5
4. Oscillator system diagram [9] 5
5. Microchip MHM-2020 Active Hydrogen Maser [11] 7
6. Hydrogen Maser conceptual diagram [7] 8
7. Relative clock timekeeping ability [1] 10
8. OCXO system diagram [9] 11
9. Aging characteristic of a precision frequency source [18] 16
10. Typical Aging Behaviors [16] 17
11. Retrace characteristic of a precision frequency source [18] 18
12. Logarithmic model fit on OCXO data 22
13. Weighted logarithmic model results, adapted from [24] 24
14. Kalman Filter algorithm flowchart [27] 31
15. Kalman Filter estimate [28] 36
16. Logarithmic Kalman Filter example from [29] 37
17. Microchip 1000X chassis 41
18. Expected fractional frequency behavior for the Microchip 1000C OCXO [18][30] 41
19. Squid setup 43
20. MMS setup 44

21. MMS setup – insulated OCXOs 44
22. MAD outlier detection and removal on OCXO A MMS dataset 49
23. OCXO B dataset baseline frequency aging estimates 51
24. OCXO A dataset 53
25. OCXO B datasets 55
26. OCXO H datasets 57
27. OCXO B MMS initial parameter search 60
28. OCXO B MMS initial logarithmic fit 61
29. OCXO B MMS logarithmic model after Least Squares Parameter Search 62
30. OCXO A logarithmic modeling results 64
31. OCXO A logarithmic modeling residuals 65
32. OCXO B logarithmic modeling results 66
33. OCXO B logarithmic modeling residuals 67
34. OCXO H logarithmic modeling results 68
35. OCXO H logarithmic modeling residuals 69
36. OCXO B Kalman Filter estimates 76
37. OCXO A linear Kalman Filter modeling results 78
38. OCXO A linear Kalman Filter modeling residuals 79
39. OCXO B linear Kalman Filter modeling results 80
40. OCXO B linear Kalman Filter modeling residuals 81
41. OCXO H linear Kalman Filter modeling results 82
42. OCXO H linear Kalman Filter modeling residuals 83
43. OCXO C datasets 95

44. OCXO D datasets 97
45. OCXO E datasets 99
46. OCXO F datasets 101
47. OCXO G datasets 103
48. OCXO A logarithmic modeling results 104
49. OCXO A logarithmic modeling residuals 104
50. OCXO B logarithmic modeling results 105
51. OCXO B logarithmic modeling residuals 105
52. OCXO C logarithmic modeling results 106
53. OCXO C logarithmic modeling residuals 106
54. OCXO D logarithmic modeling results 107
55. OCXO D logarithmic modeling residuals 107
56. OCXO E logarithmic modeling results 108
57. OCXO E logarithmic modeling residuals 108
58. OCXO F logarithmic modeling results 109
59. OCXO F logarithmic modeling residuals 109
60. OCXO G logarithmic modeling results 110
61. OCXO G logarithmic modeling residuals 110
62. OCXO H logarithmic modeling results 111
63. OCXO H logarithmic modeling residuals 111
64. OCXO A linear Kalman Filter modeling results 112
65. OCXO A linear Kalman Filter residuals 112
66. OCXO B linear Kalman Filter modeling results 113

67. OCXO B linear Kalman Filter residuals 113
68. OCXO C linear Kalman Filter modeling results 114
69. OCXO C linear Kalman Filter residuals 114
70. OCXO D linear Kalman Filter modeling results 115
71. OCXO D linear Kalman Filter residuals 115
72. OCXO E linear Kalman Filter modeling results 116
73. OCXO E linear Kalman Filter residuals 116
74. OCXO F linear Kalman Filter modeling results 117
75. OCXO F linear Kalman Filter residuals 117
76. OCXO G linear Kalman Filter modeling results 118
77. OCXO G linear Kalman Filter residuals 118
78. OCXO H linear Kalman Filter modeling results 119
79. OCXO H linear Kalman Filter residuals 119

CHAPTER 1:

INTRODUCTION

Timing is vital. People use time to coordinate their daily schedules. Financial institutions use time to definitively mark transactions with the time of arrival, thus preventing ambiguity on the order of transactions. Computers use timing internally to synchronize instructions within the processor and externally to manage peripheral interfaces. Scientists use time to make observations and quantify natural phenomena. The use of time and timing is far reaching.

Like the applications of time, the sources of time vary widely. Time can be marked by tracking observations of astronomical phenomenon, i.e. day-night cycles. Or, a wristwatch can be used to maintain the number of hours, minutes, and seconds that have elapsed. Regardless of the disparity in scale and duration, both day-night observations and wristwatch are clocks. Clocks are systems comprised of a frequency standard and counter, where the counter accumulates the cycles of a frequency standard to measure time [1].

A clock can be designed using infinite combinations of frequencies and counts. In order for time information to be meaningful between people and systems, a standard unit of time is required for synchronization between clocks. The second, which is the basic unit of time measurements, is defined by the unperturbed ground-state hyperfine transition frequency of the cesium-133 atom: 9,192,631,770 $s^{-}1$ [1][2]. This cesium isotope was selected since its inner orbitals are filled to capacity, with a single electron on the outermost orbital. This electron configuration reduces the sensitivity of the atom's quantum transitions to external influences [3].

This level of certainty of a frequency source's behavior is important for measuring small increments of time. Typically, when an application requires timing resolution on the order of fractions of a second, the time keeping is referred to as precise time. Applications which require precise timing can receive time and frequency synchronization from dedicated precise time organizations. Universal Coordinated Time (UTC) is a time scale used internationally for time synchronization, but there is not one true physical instance of this time scale. UTC is estimated by timing laboratories, and the collective estimates from all participating timing laboratories are monitored and reported by the International Bureau of Weights and Measures (BIPM). In the United States, the National Institute of Standards and Technology (NIST) maintains an estimate of UTC, denoted as UTC(NIST) [4]. Receiving time from dedicated organizations has the advantage of expert maintenance and continuous improvement of time systems. But the methods of transferring time and frequency information will impact synchronization performance. For instance, increasing transmission distance incurs cumulative risk for noise interference. Even for high-performance fiber-optic link systems, time transfers are subject to degradation [5]. Special attention is required to mitigate these effects. There are other situations where links to UTC or similar time scales are limited or impossible, as when satellites go out of range for time transfer or when telecommunications links are unavailable. Or simply, the system is meant to be standalone, without continuous outside synchronization. These situations require designers to develop or maintain time and frequency independently.

The core component of precise time is the frequency standard. When a frequency standard must be implemented for a specific application, there are certain classes of devices which are used. Figure 1 shows the broad categories of frequency standard devices, from which timing information can be derived. In order of increasing precision, the categories include

resonant circuit, piezoelectric, and atomic frequency standards. These frequency standard categories have significant differences in mechanism and behavior, which give rise to their performance ranges. Ergo, examining each category's implementation will aid in understanding these performance differences.

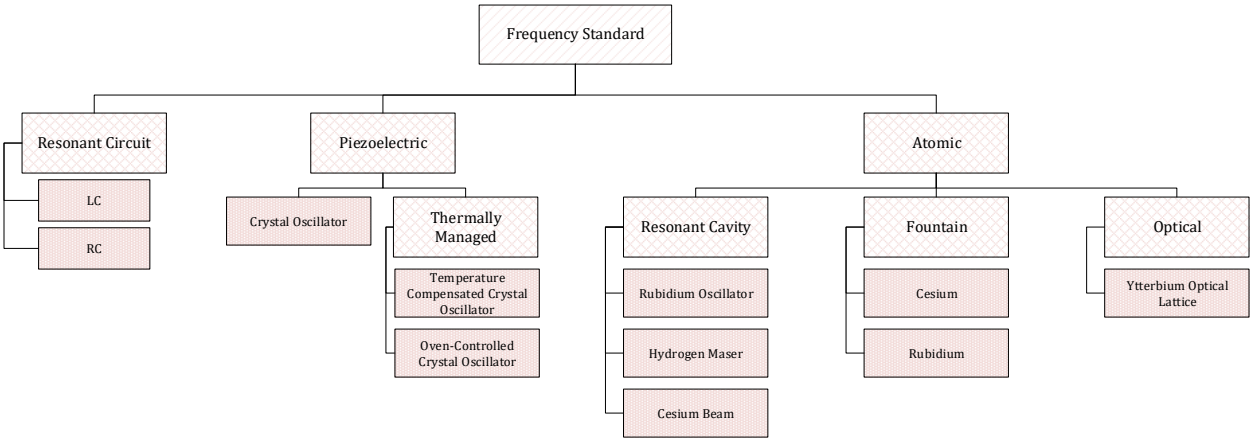

Figure 1: Taxonomy of frequency standards [6][7]

For a device to produce a harmonic oscillatory signal, it must contain an amplifier to amplify a signal; positive feedback to the amplifier to provide negative resistance; an element to select the frequency of oscillation; and some component of nonlinearity to limit the oscillatory signal's amplitude [6]. Where the categories of frequency standard differ most is in the frequency selecting element, known as a resonator.

Resonant circuit frequency standards implement the resonator with linear circuit components, such as resistors, inductors, and capacitors. These resonant circuit frequency standards are among the most abundant type of standard, since these components are easy to integrate into larger circuit assemblies. Figure 2 shows an idealized Colpitts oscillator, which is a standard oscillator design. This circuit implements an amplifier using a bipolar junction transistor, which produces a positive feedback path from the capacitor-inductor network resonator. The positive feedback produces negative resistance, which induces the circuit to

oscillate. The nonlinearity which limits these oscillations stems from the nonideal aspects of the transistor and other components. This type of oscillator construction results in mid to low performance [6]. Depending upon the application, this type of standard can be sufficiently stable. Consumer electronics generally use this type of oscillator.

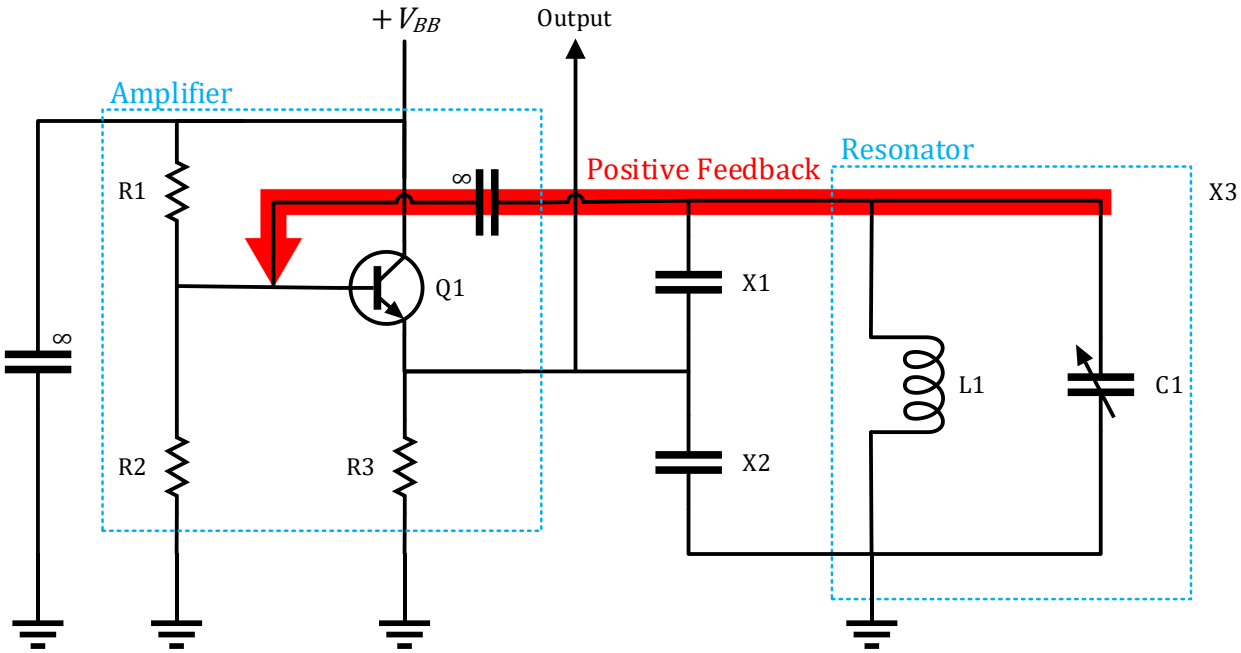

Figure 2: LC Colpitts oscillator circuit [6]

Piezoelectric frequency standards use materials with piezoelectric properties as a resonator. The direct piezoelectric effect occurs when a material develops an electric potential when mechanical stress is applied. The indirect piezoelectric effect occurs when a material mechanically deforms due to an applied electric potential [8]. The direct and indirect piezoelectric effects cause the material to resonate at specific frequencies when an AC signal is applied, which is ideal for a resonator component. Quartz crystals are a common piezoelectric resonator material used to make oscillators. Figure 3 shows the internal construction of one such quartz resonator. The quartz disk is attached to the assembly with two electrodes, which connect the crystal to the oscillator circuit.

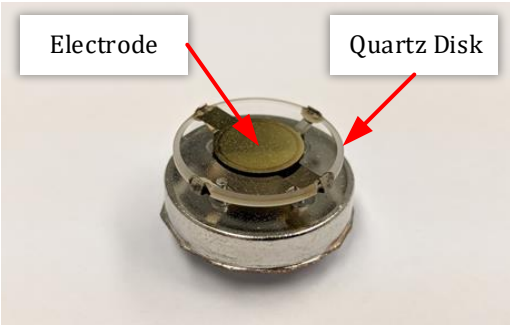

Figure 3: Quartz resonator (top of canister removed)

A generic quartz crystal oscillator (XO) diagram is shown in Figure 4, which follows the resonator-with-amplifier design paradigm. This can take a similar form to the LC oscillator presented in Figure 2, with the quartz resonator replacing the LC resonator. Due to the reduction of energy loss in the quartz resonator compared to resonators constructed from ordinary resonant circuits, these XOs can surpass resonant circuit performance [6]. Depending on an application's requirements, XOs can directly replace resonant circuit frequency standards due to the minor differences in form factor and circuit topology. For systems with high performance requirements, XOs are modules in standalone packages which contain additional features to attain superior frequency stability. Telecommunication and satellite systems are examples of applications which require this type of performance.

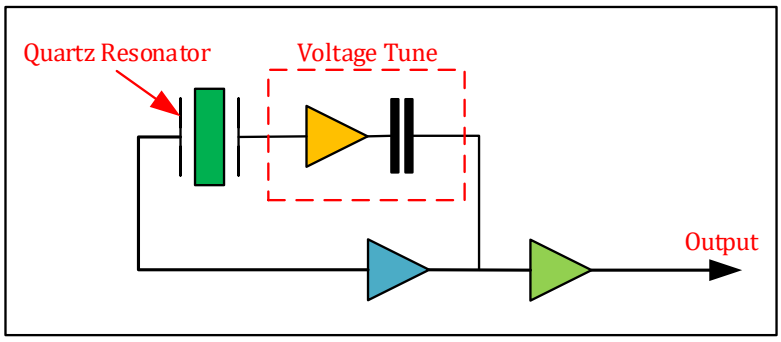

Figure 4: Oscillator system diagram [9]

Atomic frequency standards use the quantum properties of specific atoms to create a resonator. Atoms experience quantum transitions, meaning that they take in and release energy in the form of electromagnetic energy, or photons. For atomic resonators, the two important quantum events are absorption and stimulated emission. Absorption is when a photon is taken into an atom, and stimulated emission is when a photon is released [10]. From the electromagnetic wave properties of the released and absorbed photon, a specific frequency may be observed and used as an active frequency standard. Or, the atomic resonance can be used to steer the frequency of another frequency standard, creating a composite system known as a passive frequency standard [10]. This strategy is known as passive since the atoms are not continuously producing a signal.

The active hydrogen maser is one example of an atomic frequency standard. Maser stands for the microwave amplification by stimulated emission of radiation, meaning the quantum transitions of a particular atom (hydrogen, in this case) is facilitated by an electromagnetic field with a frequency in the microwave range. Figure 5 shows an example of one such standard.

Figure 5: Microchip MHM-2020 Active Hydrogen Maser [11]

Figure 6 demonstrates the general components necessary to produce the maser effect. Dihydrogen (H_2) molecules are released from a storage source and are disassociated into single hydrogen atoms. Next, the beam of hydrogen passes through a state selector, consisting of permanent magnets, which allows hydrogen atoms in the proper state to enter the cavity, while pulling atoms with improper states out of the hydrogen beam entering the storage bulb located inside the cavity. The storage bulb concentrates the population of hydrogen atoms, increasing the amount of time each atom can continue transitions. The microwave cavity provides the proper electromagnetic environment for the atoms to undergo quantum transitions, exchanging photons with each other, the cavity, and the antenna. The antenna couples a small portion of the energy from the photons via the current induced by the photons. The frequency of this induced current is

down converted from the high frequency of the hydrogen atomic resonance into a frequency reference suitable for utilization.

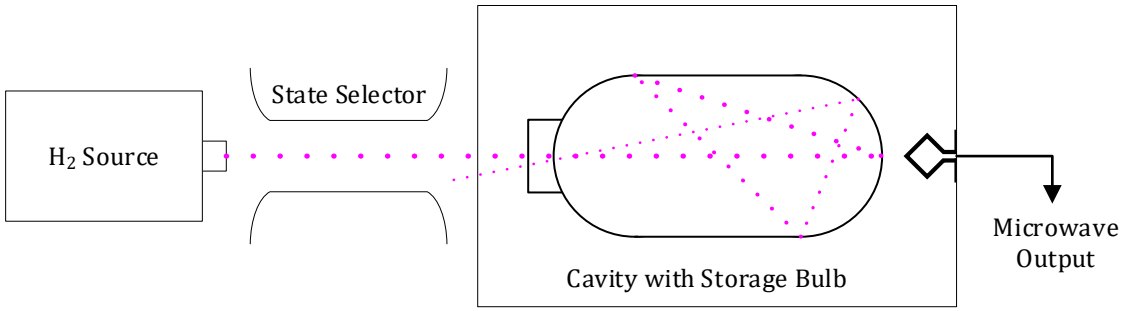

Figure 6: Hydrogen Maser conceptual diagram [7]

The resonant circuit, piezoelectric, and atomic frequency standards have diverse methods to produce a stable frequency signal. Although all of these devices are frequency standards, the stability performance of the devices are not identical. As of the time of this writing, and the foreseeable future, there is not a perfect frequency standard. Without a perfect frequency source, issues arise for systems dependent on accurate time keeping. For example, with high-frequency trading on financial markets the intervals between transactions are increasingly miniscule, and the financial losses due to time stamp errors can be high [12]. Or with time-based navigation using Global Navigation Satellite Systems (GNSS), small timing errors can result in significant positioning errors [13]. The discipline of precise time is dedicated to improving the reliability and accuracy of timing solutions. The different classes of frequency standard emerged from the requirements of precise time, via concentrated efforts to minimize the variation in timing systems such that time and frequency specifications of client systems are met.

The variation of frequency standards, and by extension clocks, is due to the nonidealities of real systems. These errors with respect to an idealized frequency standard interfere with the production of signals with known frequency and amplitude, manifesting as variation in frequency and amplitude with time. It is important to identify the causes of error, because once

the error source is known it can be addressed. Error can be segmented into systematic and stochastic error. Systematic error is the deterministic frequency drift due to internal and external influences of a frequency standard. Environmental effects, such as temperature, pressure, and magnetic fields, are a source of external systematic error. Internal systematic error arises from a frequency standard's design, as well as changes to internal components [14]. Stochastic error is the totality of all unpredicted events acting upon and within a system. These events include random mechanical vibrations, environmental transients, and other sporadic events that are of unknown origin [15].

Resonant circuit, piezoelectric, and atomic frequency standards are distinguished by how each addresses these random and deterministic sources of error. Given these issues with frequency error, the degree to which error influences a frequency standard determines its relative performance. This gives rise to differing capabilities on timing resolution. Figure 7 shows the relative accuracy of frequency standards. Stochastic error influences the smallest interval upon which a frequency standard is accurate, with clocks based on standards such as the hydrogen maser and quartz oscillator reaching timing resolution on the order of picoseconds. Systematic error generally limits how long a frequency standard maintains synchronization with other standards, due to frequency drift. In terms of synchronization, crystal oscillators can maintain synchronization for one day, whereas cesium beam atomic frequency standards can maintain synchronization out to a year.

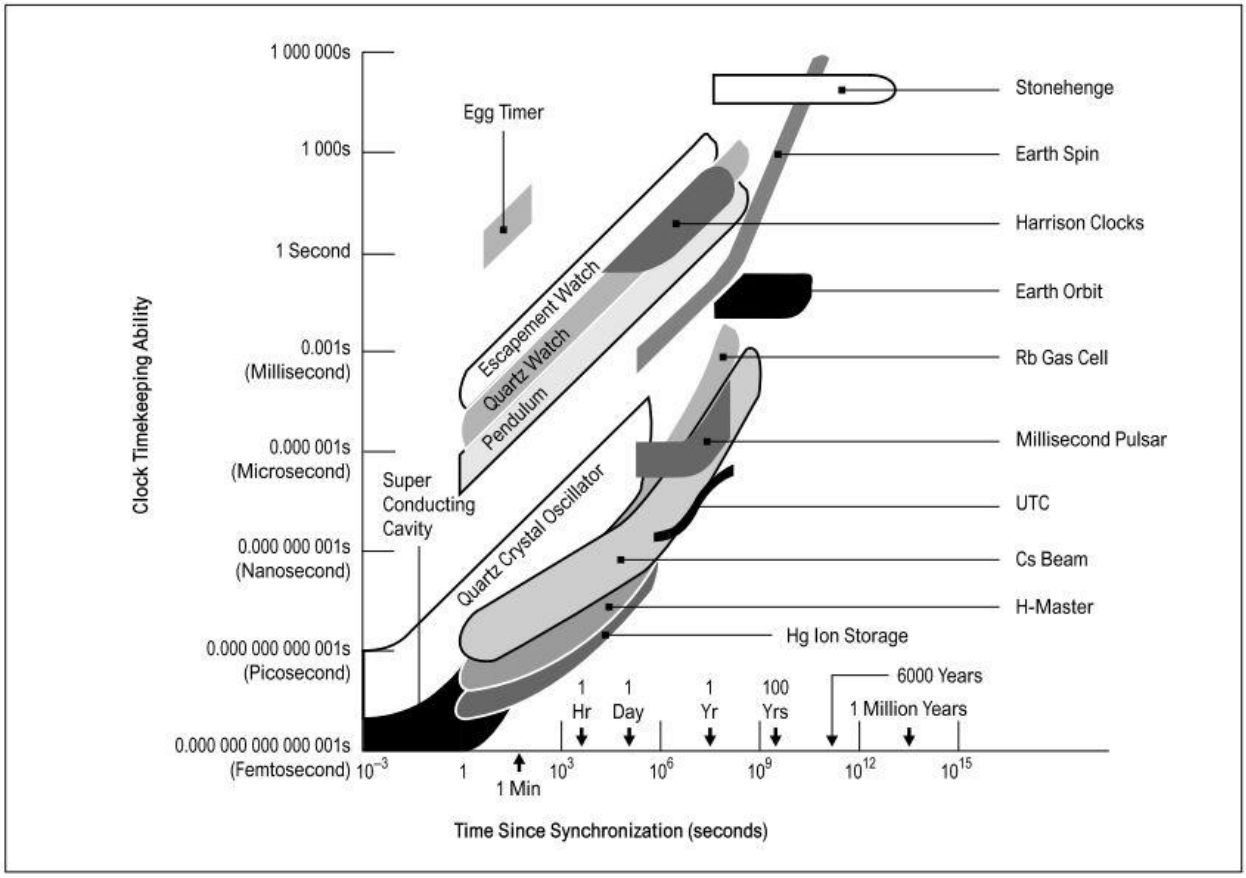

Figure 7: Relative clock timekeeping ability [1]

For practitioners of precise time, the mitigation of error in frequency standards is a primary goal to improve timekeeping overall. In order to address sources of error, the error must be characterized in some manner. Models are used to codify the predicted behavior of a system, which is used to identify root causes of error. Once these models are implemented, system designers can use various methods to reduce the impact of error. Generally, these methods involve changing a system's design to improve error resilience, eliminating external sources of error, or implementing a form of active control within a standard to compensate for system changes.

As an example of this characterization and alteration process, designers have modeled the thermal sensitivity of XOs [16], which is a systematic error. With a model of XO behavior

created, two XO design variants were created, the temperature-compensated crystal oscillator (TCXO) and the oven-controlled crystal oscillator (OCXO). These designs address the temperature sensitivity of the XO via mathematical compensation and temperature control, respectively. The TCXO measures the ambient temperature of the oscillator's environment, which is used to adjust the output of the standard. Versus the standard XO, the TCXO achieves an order of magnitude frequency stability improvement [9]. Figure 8 illustrates the OCXO design changes, which add an oven to apply heat to the oscillator components, as well as a temperature sensor and controller used to maintain the temperature. Versus the standard XO, the OCXO achieves three orders of magnitude improvement in frequency stability [9].

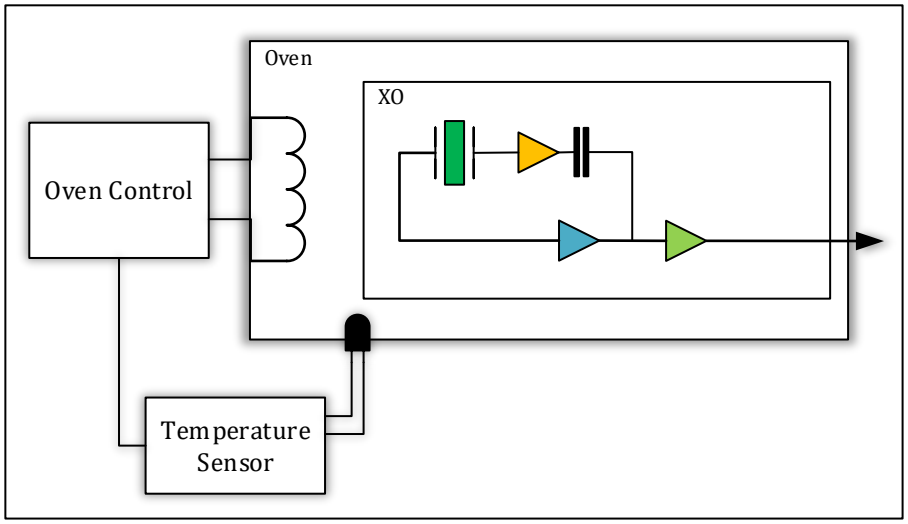

Figure 8: OCXO system diagram [9]

As exemplified by the handling of thermal systematic error, addressing sources of error leads to improved frequency standard performance. Addressing each of these error sources in turn is paramount to the maturation of frequency standard technology. Of error sources, frequency aging is of particular concern. Frequency aging is the combination of internal systematic errors, which result in long-term change in a source's frequency. The aging of frequency standards is the largest contributor to time error in clock systems [17]. In addition to

the large impact of frequency aging, characterizing its effects can be difficult. Whereas environmental effects such as temperature can be directly measured and accounted for with control systems, frequency aging can only be measured indirectly, has multiple competing sources, and occurs over long time periods. Trends occurring over long time periods are especially difficult to characterize by virtue of the length of time before an analysis can be completed. Due to the difficulty in directly observing frequency aging, the utilization of modeling techniques is paramount.

Although models have been developed and used to describe frequency aging in the literature, there is a lack of guidance on the relative performance of each model. In that, a procedure of comparison between model performance is not defined. Comparing modeling procedures by implementing figures of merit and baseline aging values can assist in determining a proper procedure for each frequency standard under evaluation.

As such, a method for comparing frequency aging model fidelity is proposed in this thesis. Two frequency aging models were selected for comparison: a logarithmic model and a linear Kalman Filter model. The relative performance of the logarithmic and linear Kalman Filter based frequency aging models is explored and quantified, using residual analysis and statistical metrics, with a focus on the benefits of each model for describing frequency aging.

CHAPTER 2:

BACKGROUND AND LITERATURE REVIEW

2.1. Background

In order to perform the procedures and analysis required to evaluate frequency aging models, some select concepts need to be understood. The signals produced by oscillators must be defined, in terms of voltage and extended to frequency. Utilizing proper units of frequency, frequency aging can then be described conceptually, providing a firm basis for modeling.

In the literature, the output signal of a frequency source is usually defined as a sinusoidal voltage signal, as expressed in Equation (1) [17]:

$$V(t) = [V_0 + \varepsilon(t)]sin[2\pi v_0 t + \phi(t)] \quad (1)$$

where V_0 is the nominal peak output voltage; $\varepsilon(t)$ is a term that represents variation of the signal amplitude; v_0 represents the nominal frequency of the output signal in hertz; $\phi(t)$ represents the phase deviation in radians; and t represents time in seconds.

Since the primary phenomenon under study in this thesis is frequency aging, the frequency component of Equation (1) is of greatest interest. However, the frequency aging phenomenon is captured by the phase deviation term $\phi(t)$ in Equation (1), so a relationship between phase and frequency must be established.

The concept of fractional frequency allows for such a relationship. Fractional frequency represents normalized frequency differences, as expressed in Equation (2) [17]:

$$y(t) = \frac{v(t) - v_0}{v_0} = \frac{\Delta f}{f} \tag{2}$$

where $v(t)$ is a frequency measurement series; v_0 is the nominal output frequency; and $\frac{\Delta f}{f}$ is the symbol for fractional frequency. In general terms, fractional frequency describes the difference of a signal's frequency to its designed frequency. Fractional frequency can also be expressed in terms of phase measurements, as in Equation (3) [17]:

$$\frac{\Delta f}{f} = \frac{1}{2\pi v_0}\frac{d\phi}{dt} = \frac{dx}{dt} \tag{3}$$

where v_0 is the nominal output frequency; $\frac{d\phi}{dt}$ is the derivative of phase measurements $\phi(t)$ in radians with respect to time; and $\frac{dx}{dt}$ is the derivative of phase measurements $x(t)$ in seconds with respect to time. The formulation in Equation (3) is useful for measurement systems which measure phase differences $x(t)$ of an oscillator in reference to a frequency standard.

Between fractional frequency and phase exists an important relationship: fractional frequency is the derivative with respect to time of phase, which can be computed as shown in Equation (4) [17]:

$$\frac{\Delta f}{f} = \frac{x_i - x_{i-1}}{\tau} \tag{4}$$

where $x(t)$ is a phase measurement series in seconds; and τ is the measurement interval between phase measurements. This transformation performs an important role when oscillator performance is recorded in phase, and the desired aging analysis must be in terms of fractional frequency.

The practice of analysis via fractional frequency is standard within the precise time community, since this allows for comparison of frequency standards with different nominal frequencies v_0. Thus, only the frequency error is compared between standards. Smaller fractional frequency measurements indicate better consistency in the production of a designed frequency. Due to the very small numbers involved, fractional frequency it is often reported in 'parts per' notation, where $1.5\ pp10^9$ means 1.5×10^{-9} [16].

Frequency stability describes the consistency of a frequency standard's output. The fractional frequency deviations in a standard's output come from the systematic and stochastic error present. The collective deviations from the designed frequency over a long timeframe is termed frequency drift. Frequency drift is the long-term change in frequency of a frequency standard due to systematic phenomena [16]. Meaning, drift is the deviation of frequency due to both environmental influences and internal changes in a device's components. Now, frequency aging is the subset of frequency drift that is attributed to internal changes to a frequency standard, such as physical and chemical changes in the resonator and changes to supporting electronics [16]. This differs from short-term fluctuations in frequency, which are stochastic in nature. Effects of frequency aging take the form of the long-term trends, which are typically measured in units of normalized fractional frequency per unit time [16]. In general, the aging rate is referenced to specific events, such as oscillator turn-on.

Frequency aging of most frequency standards is known to conform to a consistent form, as shown in Figure 9. This pattern has been shown to be demonstrated particularly by crystal oscillators [18]. As indicated in Figure 9, frequency aging has three distinct sections: a turn-on transient, initial aging (logarithmic region), and long-term aging (approximately linear region). Spanning turn-on and part of initial aging, oscillators experience warm-up, which is the period

during which the oscillator stabilizes to operating temperature. The length of warm-up is prescribed by the specifications of an oscillator's design, after which the oscillator will attain steady-state performance. The initial aging region contains the highest rate of frequency aging, which is then followed by a dramatic decrease in aging as an oscillator reaches the long-term aging region.

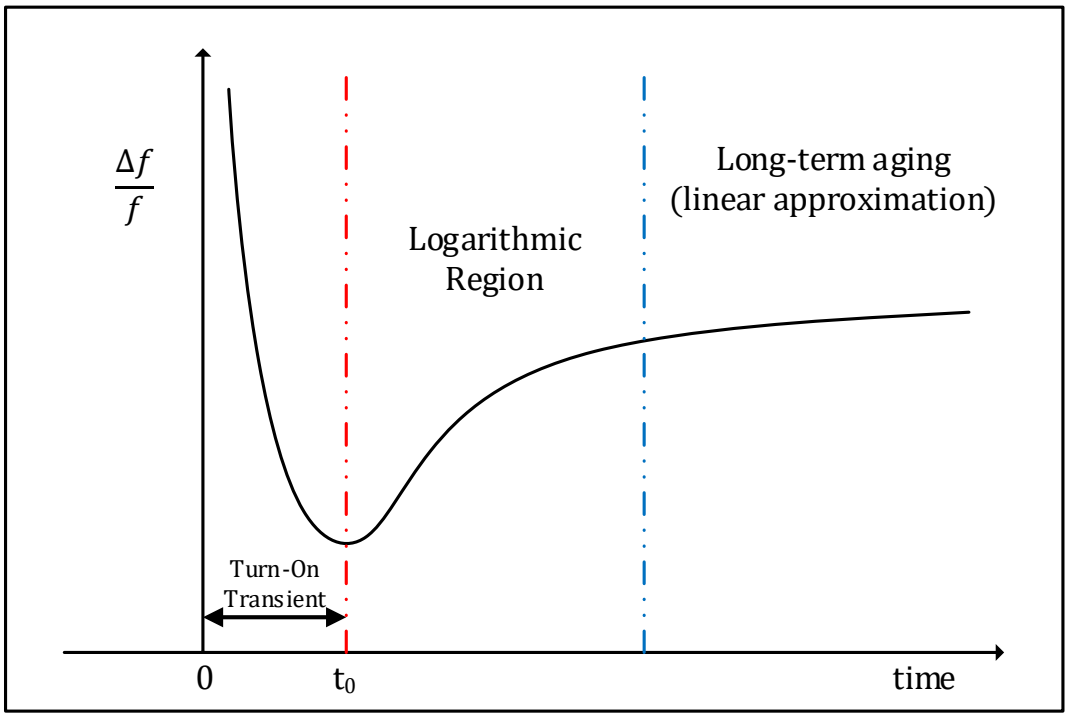

Figure 9: Aging characteristic of a precision frequency source [18]

Over all of the aging periods, the frequency aging trend can be positive or negative, and can reverse after arbitrary time intervals [16]. These changes in behavior are due to a combination of multiple mechanisms with different time constants and magnitudes. In Figure 10, there are two aging mechanisms, $A(t)$ and $B(t)$ acting simultaneously. These two phenomena interfere with each other, resulting in the true aging behavior of $A(t) + B(t)$; initially the aging trend is positive due to the high rate change in mechanism $A(t)$, until $A(t)$'s effects are dominated by mechanism $B(t)$.

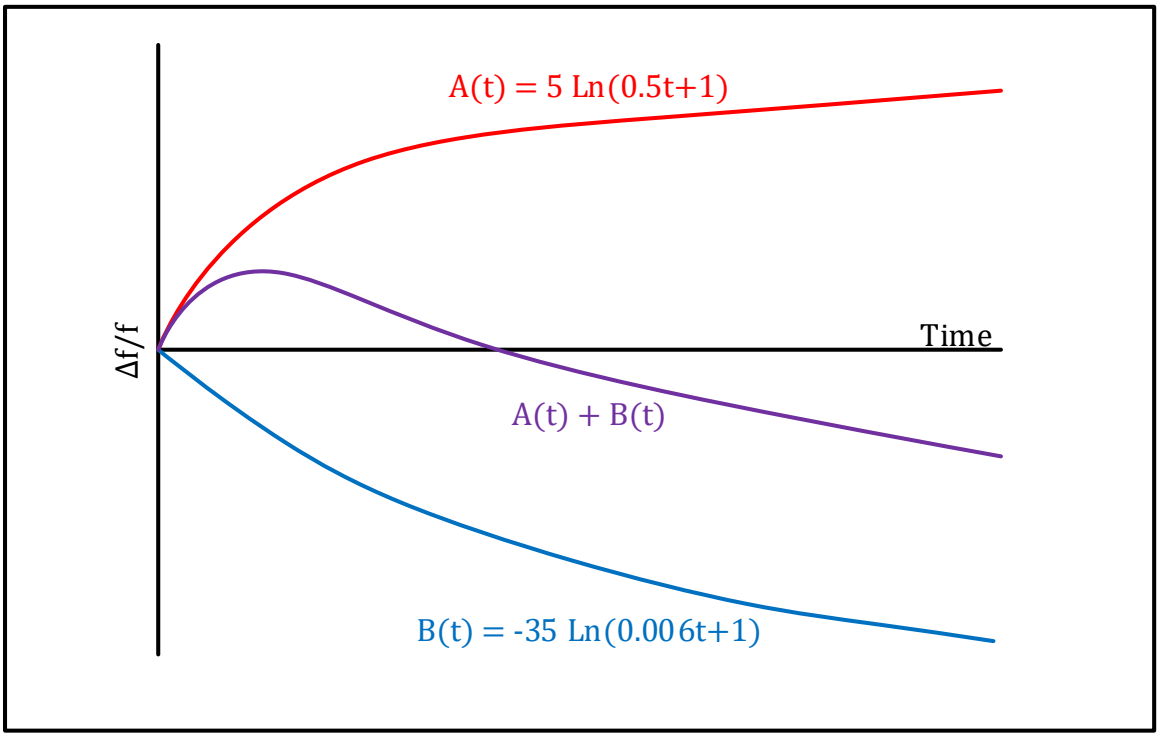

Figure 10: Typical Aging Behaviors [16]

Due to the combination of aging factors, it is difficult to predetermine the aging behavior of a frequency standard. Changes to an oscillator and its environment have wide-ranging impacts on its aging behavior. The most dramatic changes that oscillators undergo regularly are on/off events, which can cause the output signal to 'retrace' a logarithmic trend which has similar magnitude and period to initial aging but can settle at a different aging rate [18]. This phenomenon is shown in Figure 11.

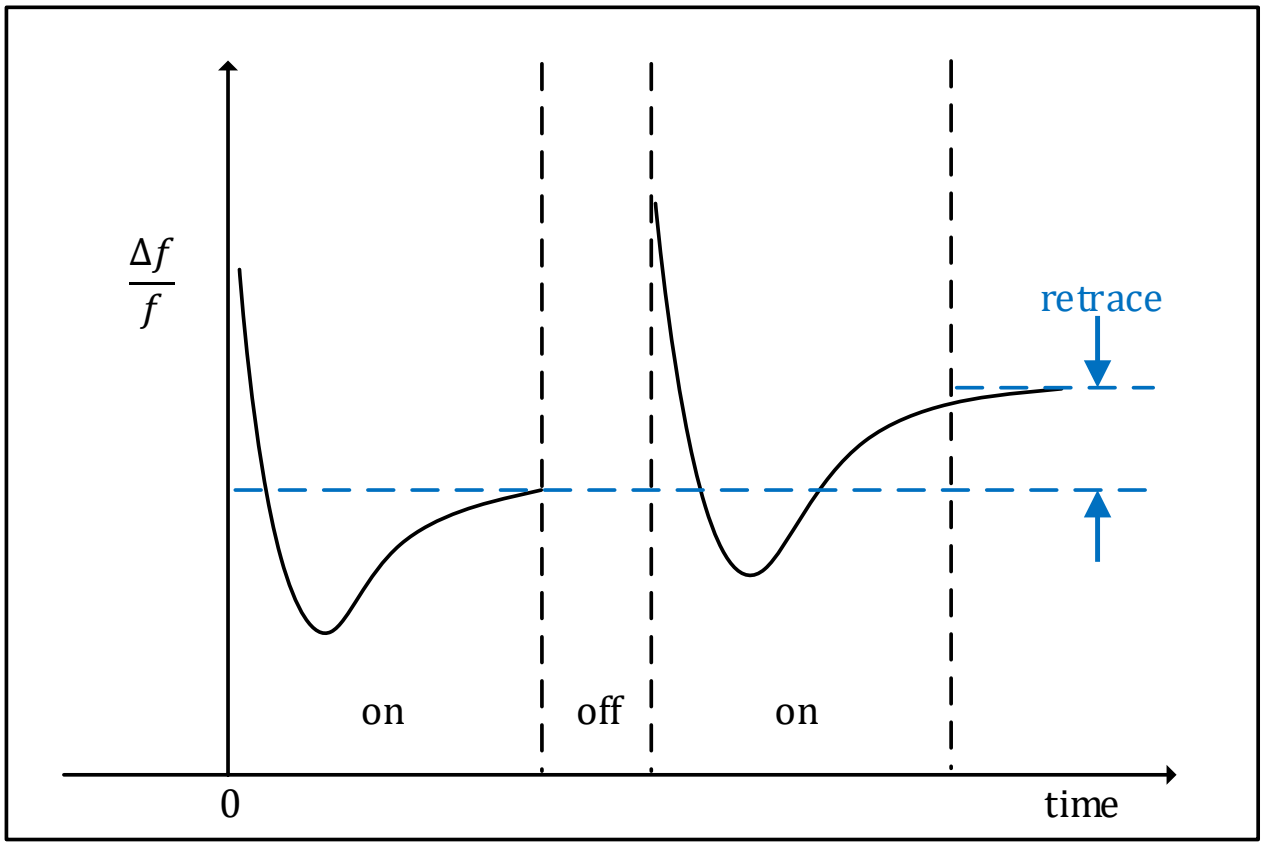

Figure 11: Retrace characteristic of a precision frequency source [18]

In summary, the important quantities of an oscillator's output signal are phase and frequency. The phase measurements of an oscillator signal can be transformed into fractional frequency, which is standard practice for analyzing a frequency standard's stability. With the signal data quantities and general aging behavior of oscillators defined, these concepts can be applied to the analysis and modeling of oscillator data.

2.2. Literature Review

In this literature review, the existing models of frequency aging are explored. One motivation for analyzing these models is that they can be used to predict future changes in a frequency standard's behavior. This allows these models to be used to correct for the effects of aging over time. For this work, the aging models will be presented in terms of crystal oscillator frequency aging.

Essentially, models describe relationships of systems to basic physical mechanisms [19]. Models employ functions with selectable numeric parameters, which are used to associate a model to a particular dataset. If a model is directly formed from theoretical concepts or first-principles, it is known as a mechanistic model. If a model is knowledge-based but does not directly derive from first-principles, it is referred to as an empirical model. Empirical models are derived from a combination of system behavior observation and theory of operation. In general, aging models are empirical, and attempt to describe aging behavior using the general trend of data or underlying physical chemical processes.

In this work, aging models will be categorized as simple, compound, or adaptive. Simple aging models assume that one aging mechanism is responsible for the totality of aging behavior. Compound models combine descriptions of multiple simultaneous aging mechanisms to represent an oscillator's frequency aging. Adaptive models can apply either simple or compound techniques but are distinct in that they update system parameter estimates iteratively - allowing parameters to vary in contribution over time.

In many of the studies that are described in this section, parametric regression analysis will be employed to find the parameters of each model.

The parameters are estimated using the least squares method shown in Equation (5) [20]:

$$LS(\beta) = \sum_{i=1}^{n} (y_i - x_i'\beta)^2 = \sum_{i=1}^{n} \varepsilon_i^2 = \varepsilon'\varepsilon \tag{5}$$

where n is the number of data points under analysis; y_i is the dependent variable matrix at instance i; x_i is the independent variable matrix at instance i; β is the parameter matrix which is varied to minimize error; and ε_i is the random deviation of the regression at instance i. The least squares method performs regression by varying the parameters in β, such that minimal solution of LS is found. When the least squares method is applied to find the optimal parameter values for each frequency aging model, the dependent variable will be the fractional frequency values for a dataset, and the independent variable will be time.

2.2.1. Simple Models

Filler introduces a fundamental model formulation for frequency aging in [21], which describes frequency aging as a logarithmic function, as shown in Figure 10. The logarithmic model presumes that the primary component of frequency aging is the mechanical stress on an oscillator's resonator [21]. The stress on a resonator's mass can be expressed as in Equation (6) [21]:

$$Y_s(t) = A \cdot ln(Bt + 1) + Y_s(0) \tag{6}$$

where Y_s is mechanical stress; A is a fitting parameter for the sign and magnitude of the stress function; B is a fitting parameter describing the decay of the natural logarithmic function; and $Y_s(0)$ is the first stress measurement.

Using the assumption that frequency aging is dominated by resonator stress, frequency aging is modeled as a logarithmic relationship in Equation (7) [21]:

$$y(t) = A \cdot ln(Bt + 1) + C \tag{7}$$

where A is the parameter responsible for the sign and magnitude of the aging process; B dictates the decay of the logarithmic expression; and C represents the starting frequency measurement of a dataset. Equations (6) and (7) form an analogy between mechanical stress on an oscillator's resonator and frequency aging. Subsequent work expands the explanation of this logarithmic behavior to include chemisorption (of a contaminant gas into the resonator), oxidation (of the resonator and electrodes), and mechanical stress (of the resonator)[16][22].

The frequency aging model in Equation (7) is typically applied to oscillator frequency data using the nonlinear least squares fitting process derived from Equation (5) to obtain the parameters A and B. The value of C may be obtained from the first frequency measurement. Estimating C via the least squares method is also valid and may have advantages for arriving at more accurate overall agreement than fitting only two parameters. Example results from this method can be seen in Figure 12, which displays the result of modeling one of the datasets collected for this thesis. Note that the fractional frequency y-axis is scaled by $pp10^9$.

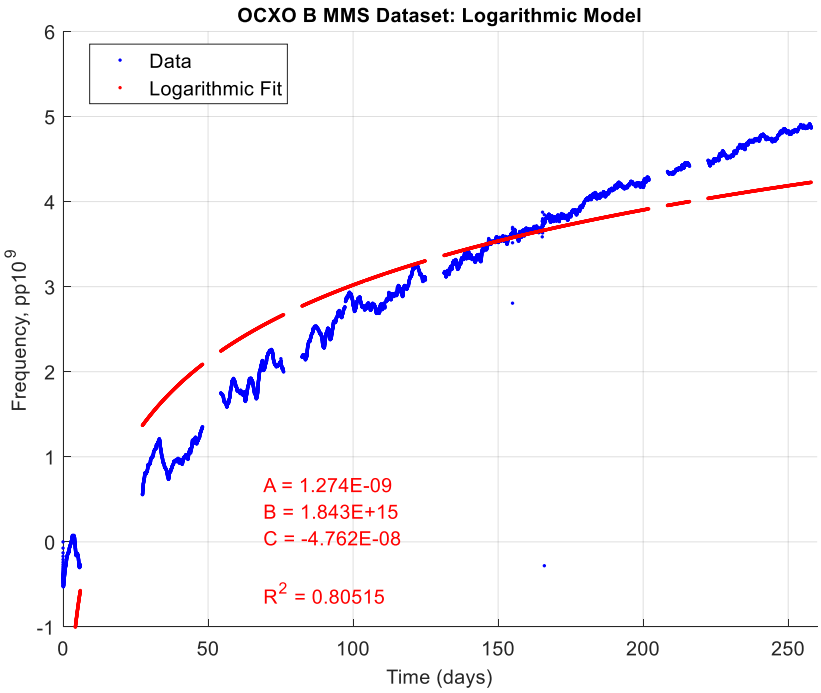

Figure 12: Logarithmic model fit on OCXO data

There are some important factors to keep in mind while using the fundamental logarithmic model. With this model, Filler assumes that frequency aging is monotonically increasing for the interval being modeled. If the aging is nonmonotonic, this model will perform poorly. To address the possibility of significant behavior changes in observed frequency data, Filler suggests the following method. First, the user should wait 12 days after a period of nonmonotonic aging to mark the beginning of a new analysis interval. Second, the user must collect a minimum of 28 days of data from this new starting point to produce accurate results from this logarithmic modeling process. Thus, this model is expected to have difficulty modeling datasets containing significant behavior changes. Another factor to attend to is the length of the analysis interval. Subsequent work from Filler and Vig in [23] recommends using the last 30 days of a 100 day period to properly capture long-term aging behavior with this logarithmic model.

Building upon the fundamental logarithmic model, Leibfried and Neubig took issue with the ability of the logarithmic model as utilized in [21] to predict long-term frequency aging [24]. Their solution was to modify the least squares regression with a weighting component to improve the logarithmic model's long-term accuracy. The least squares equation used to apply the logarithmic model with weighting is expressed in Equation (8) [24]:

$$LS(A,B,C) = \sum_{i=1}^{n}[(y(t) - A \cdot ln(Bt_i + 1) + C)^2 \cdot (1 - e^{-bt_i})] \quad (8)$$

where A represents the sign and magnitude of the logarithmic function; B is the decay term of the logarithmic function; C is the y-axis offset for the first frequency measurement; and b controls the decay of the exponential weighting term. The weighting increases the contribution of early frequency measurements in the least squares method, while decreasing the contribution of later frequency measurements. The reduction of weight for later values results in a flatter shape of the fitted curve, which attempts to reflect the diminished aging rate in the long-term linear aging region shown in Figure 9.

To determine the b weighting parameter, a preliminary least squares regression must be conducted. The difference between a dataset's frequency measurements and standard logarithmic model regression produces an error curve. This error curve is used to fit a first-order differential function as expressed in Equation (9) [24]:

$$y(t) = k(1 - e^{-bt}) \quad (9)$$

where k controls the sign and magnitude of the exponential function; and b controls the decay of the exponential function. After estimating b with the curve fit of Equation (9), the weighted least squares method described in Equation (8) can be directly used to determine parameters A, B, and C for a dataset. One example fit obtained by using this weighted logarithmic model approach is

shown in Figure 13. Note the unweighted curve predicts a greater fractional frequency magnitude than that exhibited by the measured data. In contrast, the weighted curve shows improved agreement during the long-term interval of this dataset.

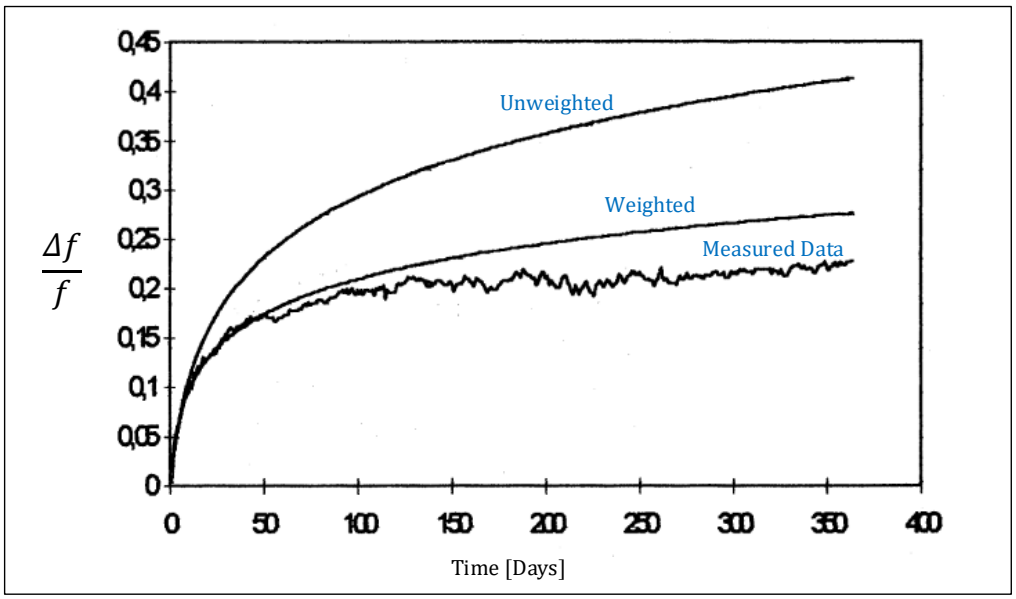

Figure 13: Weighted logarithmic model results, adapted from [24]

In addition to logarithmic models, other functions have been used with the least squares method to model frequency aging in oscillators. One such model uses an exponential function to describe chemical reactions within an oscillator's resonator, as well as the exchange of molecules between a resonator and its surrounding environment in adsorption and desorption processes - as shown in Equation (10) [22]:

$$y(t) = A\left[1 - e^{-\frac{t}{B}}\right] + C \tag{10}$$

where parameter A controls the magnitude and sign of the exponential expression; B determines the decay of the exponential; and C provides a y-axis offset for the first frequency measurement.

A power function is another expression which can be used to model frequency aging. It is based on the exchange of material between a resonator's quartz crystal and electrode attachments, with the relationship set in Equation (11) [22]:

$$y(t) = At^B + C, \qquad B \approx \frac{1}{2} \tag{11}$$

where parameter A controls the sign and magnitude of the expression; parameter B sets the rate of growth of the function, with a value expected to be approximately $\frac{1}{2}$; and parameter C which provides a y-axis offset in frequency based on the first frequency measurement of a dataset.

Logarithmic, exponential, and power functions are commonly used to model frequency aging. But, as with any model, there are behaviors in frequency aging which are not fully captured by specific modeling techniques, as evidenced by the addition of a weighting parameter to the logarithmic model. Therefore, modifications to the frequency aging modeling techniques are required to improve model fidelity on frequency aging behaviors.

2.2.2. Compound Models

A logical next step for improving modeling of frequency aging is to break the assumption that only one aging model dominates the total frequency aging behavior of a frequency standard. As illustrated in Figure 10, aging mechanisms can constructively and destructively interfere with each other, as these mechanisms are simultaneously active within an oscillator [22]. Therefore, models which can describe simultaneous frequency aging mechanisms are expected to have an advantage in tracking overall frequency aging behaviors.

The error in modeling performance of simple models can provide clues toward capturing multiple aging mechanisms in a compound model. In the weighted logarithmic model introduced in Section 2.2.1, the exponential model of Equation (10) has very similar form to the weight function of Equation (9), indicating that Leibfried and Neubig found that the logarithmic model

needed an exponential component to better capture long-term frequency aging behaviors. Effectively, the weighted logarithmic model implements a compound model.

Compound models can contain many configurations of expressions, such that the degrees of freedom in a model can be configured as needed. In one such attempt to consider many possible combinations of aging mechanisms, four copies of the logarithmic, exponential, and power functions were combined to provide 24 degrees of freedom in Equation (12) [25]:

$$y(t) = \sum_{L=1}^{4} A_L \, ln(B_L t) + \sum_{E=1}^{4} A_E \, e^{(B_E t)} + \sum_{P=1}^{4} A_P \, t^{B_P} \qquad (12)$$

where A_L determines the sign and magnitude of the L^{th} logarithmic component; B_L determines the decay of the L^{th} logarithmic component; A_E determines the sign and magnitude of the E^{th} exponential component; B_E determines the decay of the E^{th} exponential component; A_P determines the sign and magnitude of the P^{th} power component; and B_P determines the rate of growth of the P^{th} power component. Each of these parameters can adjust the contribution of its model segment. Thus, the compound model has the ability to describe multiple aging phenomena with different transient impacts.

The compound model presented in Equation (12) has substantially more flexibility to match frequency aging trends than the simple models previously presented; however, using the least squares method to determine this many parameters is error prone due to the multitude of possible parameter ratios. While compound models provide the flexibility needed to describe simultaneous aging mechanisms, the least squares regression used to find model parameters limits this utility.

2.2.3. Adaptive Models

To improve the usability of frequency aging models while maintaining behavioral flexibility, the method of parametric regression must be addressed. Thus far, the simple and compound models introduced used static parameters. Meaning, the parameters found via the least squares method are set once for a dataset. For well-behaved datasets which follow the traditional logarithmic aging trend, these simple aging models are sufficient. As an increasing number of dramatically changing behaviors are encountered in a dataset, a compound model must be used with unique components to describe each behavior instance to gain model fidelity improvements. This results in added complexity and reduced generality. Apart from model accuracy considerations, the least squares regression method has implications for how measurements are processed. To evaluate a model using the least squares method, the full measurement set is conventionally used to obtain model parameters. Given the slow progression of frequency aging in frequency standards, it can be quite a long time until a useful model output of frequency aging can be made utilizing this technique.

As an alternative to the previous regression method, adaptive models use dynamic parameters, such that the parameter values found using adaptive methods can change over time as required for a dataset. Since these parameters are iteratively updated, a new model output of a system's state can be produced on each iteration as well – so outputs of frequency aging can be made upon each measurement and not just on a complete dataset. Thus, adaptive models are well suited for dataset analysis or real-time control of systems.

Adaptive models are based on the concept of dynamic systems. Dynamic system equations are simultaneous differential (or difference) equations used to describe the behavior of

a system over time. The general form of discrete time-varying dynamic systems is as modeled in Equations (13) [26]:

$$\begin{aligned} x_k &= f_k(x_{k-1}, u_{k-1}) + w_{k-1} \\ z_k &= h_k(x_k, u_k) \quad\quad + v_k \end{aligned} \tag{13}$$

where x_k is a state vector that contains values which describe the current state of a system, u_k is a vector containing known control inputs to a system, f_k is the system model parameter matrix which updates the previous iteration's state vector x_{k-1} and control vector u_{k-1} to the current iteration k, and w_k is a vector representing the random error for an iteration. z_k is a vector containing the sensor measurements of a system's state, h_k is the sensor model parameter matrix which uses the state vector x_k and control vector u_k to calculate the measurement vector z_k, and v_k is the vector representing the error associated with sensor measurements.

The distinctive qualities in Equations (13) are the handling of time and separation of system behaviors. The inclusion of model time-dependence allows for model parameters to update on each analysis iteration. This system of equations allows for the separate description of a system and its measurement. Additionally, each equation separates deterministic and random phenomena, such that the full range of system behaviors can be modeled [26]. The random components and time-dependence are the primary attributes which separate adaptive models from the prior simple and complex models.

Due to the new attributes of adaptive models, the traditional least squares method from Equation (5) is insufficient to determine adaptive model parameters. As such, another algorithm is required which can update model parameters with time and maintain a concept of uncertainty: the Kalman Filter. The Kalman Filter algorithm is an optimal estimator, which executes a prediction and estimate cycle upon each iteration of the algorithm [27]. The prediction and estimate process applies the least squares method to determine model parameter values.

Any of the models implemented in the simple or compound aging model categories can be described as a dynamic system and integrated with the estimator.

To generate model predictions using the Kalman Filter, the dynamic system model must be translated into terms of the Kalman Filter's components, as summarized in Table 1 [26]. The first three components concern the state of the system. The state vector x_k holds the n state values of a system at iteration k. The measurement vector z_k contains the ℓ sensor measurements at iteration k. The error covariance matrix P_k contains the covariance of each state value n relative to the ℓ sensor measurements for each iteration k. The remaining Kalman Filter components define how the state estimates are updated for each iteration k. The state transition matrix A defines the relationships between state values and how they change on each iteration k. The state-to-measurement transition matrix H is used to convert the state vector x_k into terms of the measurement vector z_k. The state error covariance matrix Q relates the expected variance of each n^{th} state value to the n state values, with respect to the error associated with state transitions. Similarly, the measurement covariance error matrix R relates the expected variance of each ℓ^{th} state value to the ℓ measurement values, with respect to the error associated with sensor measurements. The Kalman gain K_k uses the prediction of the overall error covariance with the measurement error covariance to scale the influence of the k^{th} measurement vector on the k^{th} state estimate.

TABLE 1:
KALMAN FILTER COMPONENTS

Variable	Description
n	Scalar number of state variables
ℓ	Scalar number of senor measurements
x_k	$n\times1$ state vector
z_k	$\ell\times1$ measurement vector
P_k	$n\times n$ error covariance matrix
A	$n\times n$ state transition matrix
H	$\ell\times n$ state-to-measurement transition matrix
Q	$n\times n$ state error covariance matrix
R	$\ell\times\ell$ measurement error covariance matrix
K_k	$n\times\ell$ Kalman gain matrix

With the individual components of the Kalman Filter defined, the algorithm's estimation and prediction cycle can be traced as illustrated in Figure 14. Here, predicted and estimated quantities are distinguished by character accent. For example, $\hat{y}$ is the estimate of some variable y, and $\bar{y}$ is the prediction of the variable y. The algorithm generates a prediction of a system state using the previous system state. Then, this prediction is modified using measurements from the system to generate a state estimate. This Kalman Filter estimate is the model's prediction for that iteration.

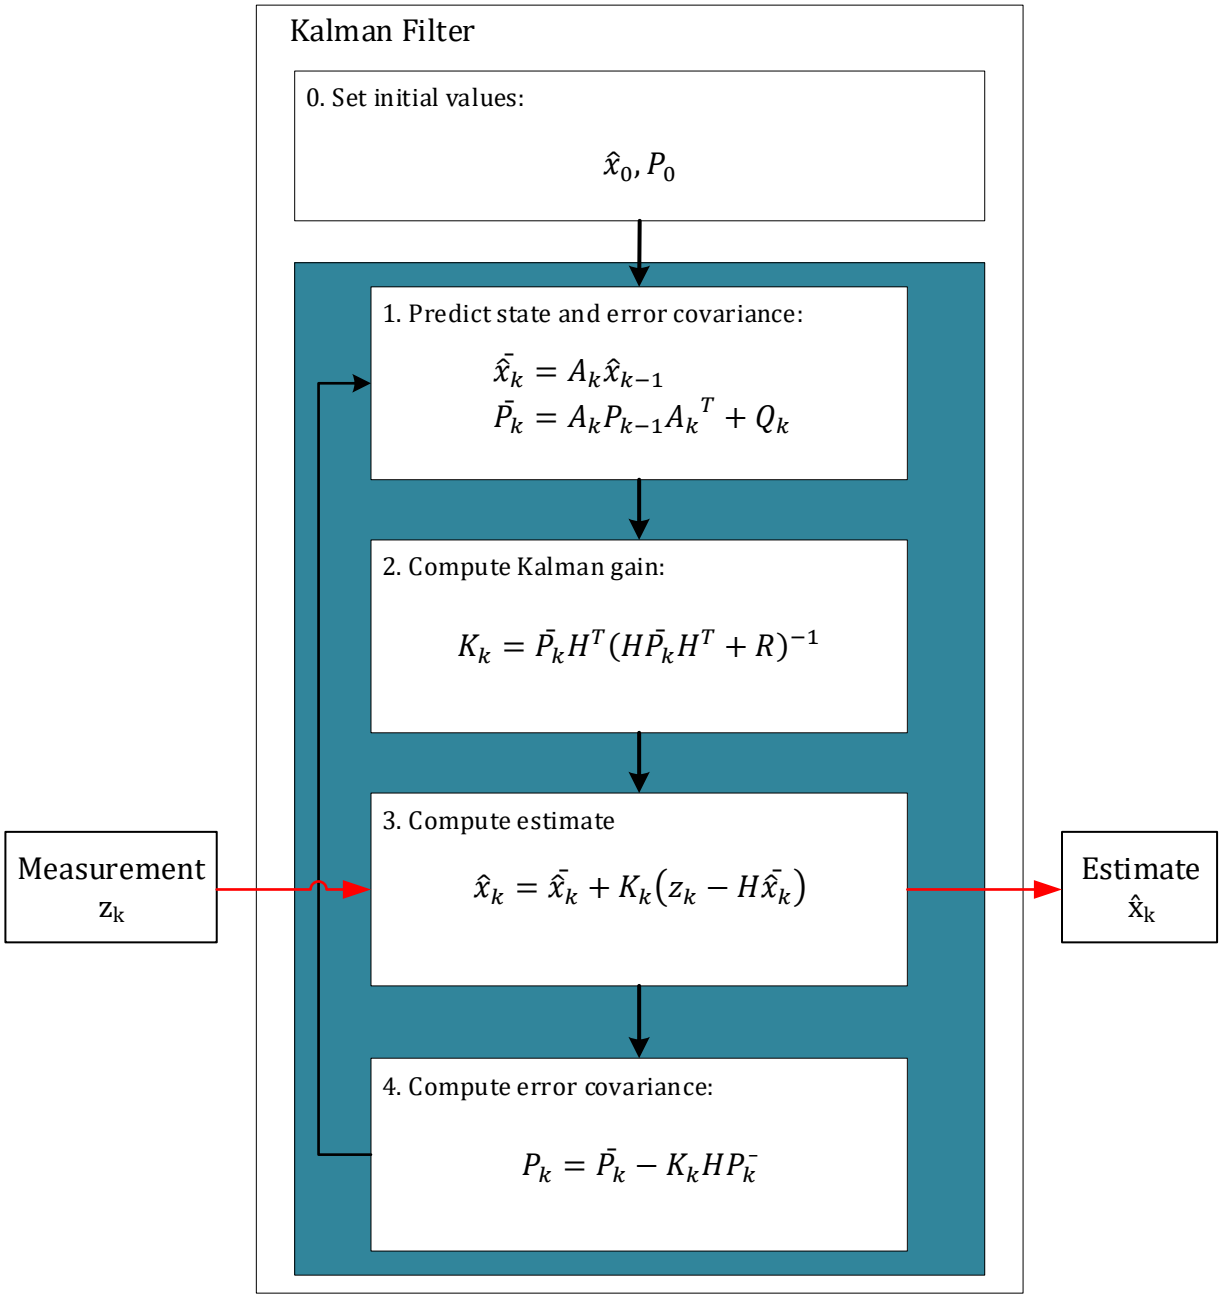

Figure 14: Kalman Filter algorithm flowchart [27]

First, the initial conditions for the algorithm must be defined in the state estimate $\hat{x}_0$ and covariance calculation P_0. These values are found by analyzing the modeled system, using any prior knowledge of the system and its behavior.

Next, the first prediction of the system state is made in Equation (14) [27]:

$$\overline{\hat{x}_k} = A\hat{x}_{k-1} \tag{14}$$

where the state transition matrix A transforms the initial state estimate $\hat{x}_0$ into the state prediction $\overline{\hat{x}_1}$.

Then, the error covariance prediction $\overline{P_1}$ is calculated in Equation (15) [27]:

$$\overline{P_k} = AP_{k-1}A^T + Q \tag{15}$$

where the previous error covariance calculation P_{k-1} is updated using the state transition matrix A and summed with the state error covariance matrix Q.

Next, the Kalman gain K_1 is computed in Equation (16) [27]:

$$K_k = \overline{P_k}H^T(H\overline{P_k}H^T + R)^{-1} \tag{16}$$

where the total error covariance prediction $\overline{P_k}$ is converted by the state-to-measurement transition matrix H such that it can be combined with the measurement error matrix R. This gain K_k determines to what extent the state estimate $\hat{x}_1$ will be influenced by the measurement z_1.

The state estimate $\hat{x}_1$ is then calculated using Equation (17) [27]:

$$\hat{x}_k = \overline{\hat{x}_k} + K_k\left(z_k - H\overline{\hat{x}_k}\right) \tag{17}$$

where the state prediction $\overline{\hat{x}_1}$ is summed with a scaled contribution from the measurement z_1. If the error associated with the measurement is less than the error associated with the system state, then the measurement will contribute more to the state estimate, and vice versa. This step generates the algorithm output for this iteration.

To setup the next iteration, the error covariance calculation P_1 must be created using Equation (18) [27]:

$$P_k = \overline{P_k} - K_k H \overline{P_k} \tag{18}$$

where the predicted error covariance $\bar{P}_1$ is differenced with the contributions of the measurements to the error covariance scaled by the Kalman gain K_1. This sequence of Equation (15) through Equation (18) is repeated until the k^{th} element of a dataset has been processed.

With the Kalman Filter algorithm set as the adaptive model regression method, various implementations of adaptive oscillator frequency aging models can be investigated. One such model describes the behavior of an oscillator in terms of linear relationships [28]. This model uses phase measurements $x(t)$ to directly compute fractional frequency $y(t)$ and frequency aging $w(t)$ by taking the first and second derivatives of $x(t)$, respectively. The dynamic system equation for the state vector is shown in Equation (19) [28]:

$$\frac{d}{dt}\begin{bmatrix} x(t) \\ y(t) \\ w(t) \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}\begin{bmatrix} x(t) \\ y(t) \\ w(t) \end{bmatrix} + \begin{bmatrix} \xi(t) \\ \mu(t) \\ \zeta(t) \end{bmatrix} \tag{19}$$

where the state vector x_k is the matrix on the left side, with $x(t)$ as the system's phase at time t, $y(t)$ as frequency, and $w(t)$ as frequency aging. The derivative of the left side represents the update of the system state. The next matrix is the system parameter matrix which multiplies the current state vector values, which is then summed with the system error vector with $\xi(t)$ as the phase error, $\mu(t)$ as the frequency error, and $\zeta(t)$ as the frequency aging error.

In order to use this adaptive linear aging model with the Kalman Filter, the dynamic system equation must be modified. The form of Equation (19) puts the state vector in terms of its first derivative. To deal with the state vector directly in the Kalman Filter, Equation (19) is

integrated and separated into the Kalman Filter components. The state vector is now defined as Equation (20) [28]:

$$x_k = \begin{bmatrix} x(t) \\ y(t) \\ w(t) \end{bmatrix} \tag{20}$$

where $x(t)$ is phase, $y(t)$ is frequency, and $w(t)$ is frequency aging.

The state transition matrix A is the integral of Equation (19)'s second matrix, and is defined as Equation (21) [28]:

$$A(\delta) = \begin{bmatrix} 1 & \delta & \delta^2/2 \\ 0 & 1 & \delta \\ 0 & 0 & 1 \end{bmatrix} \tag{21}$$

where δ is the elapsed time between the current and previous iterations. The error component of Equation (19) is represented in the state error matrix Q.

The measurement vector for this model implementation is defined in Equation (22) [28]:

$$z_k = \begin{bmatrix} 0 \\ y(k) \\ 0 \end{bmatrix} \tag{22}$$

where the only system value under measurement is the frequency measurements $y(t)$. If desired, this vector can be updated to include phase measurements $x(t)$, but the frequency aging component cannot be directly measured. To match this measurement vector, the state-to-measurement matrix is defined in Equation (23) [28]:

$$H = [0 \quad 1 \quad 0] \tag{23}$$

where the matrix is set to 1 for present state measurements, and 0 for unmeasured quantities. The error component of the measurement is contained in the measurement error matrix R.

Using the components defined in Equation (20) through Equation (23) and the error matrices Q and R, the Kalman Filter algorithm can be iterated as shown in Figure 14. Figure 15

contains example results of using the Kalman Filter to evaluate the linear frequency aging model [28]. Figure 15(a) contains the phase estimate that was obtained indirectly from frequency measurements. Figure 15(b) contains the frequency estimate based on frequency measurements, which were updated until about 50 weeks. Figure 15(c) contains the frequency aging estimate which was calculated based on the frequency estimate. After 50 weeks, the Kalman Filter iterated without new measurement input, predicting the behavior of the oscillator using the dynamic system equations described in Equation (19).

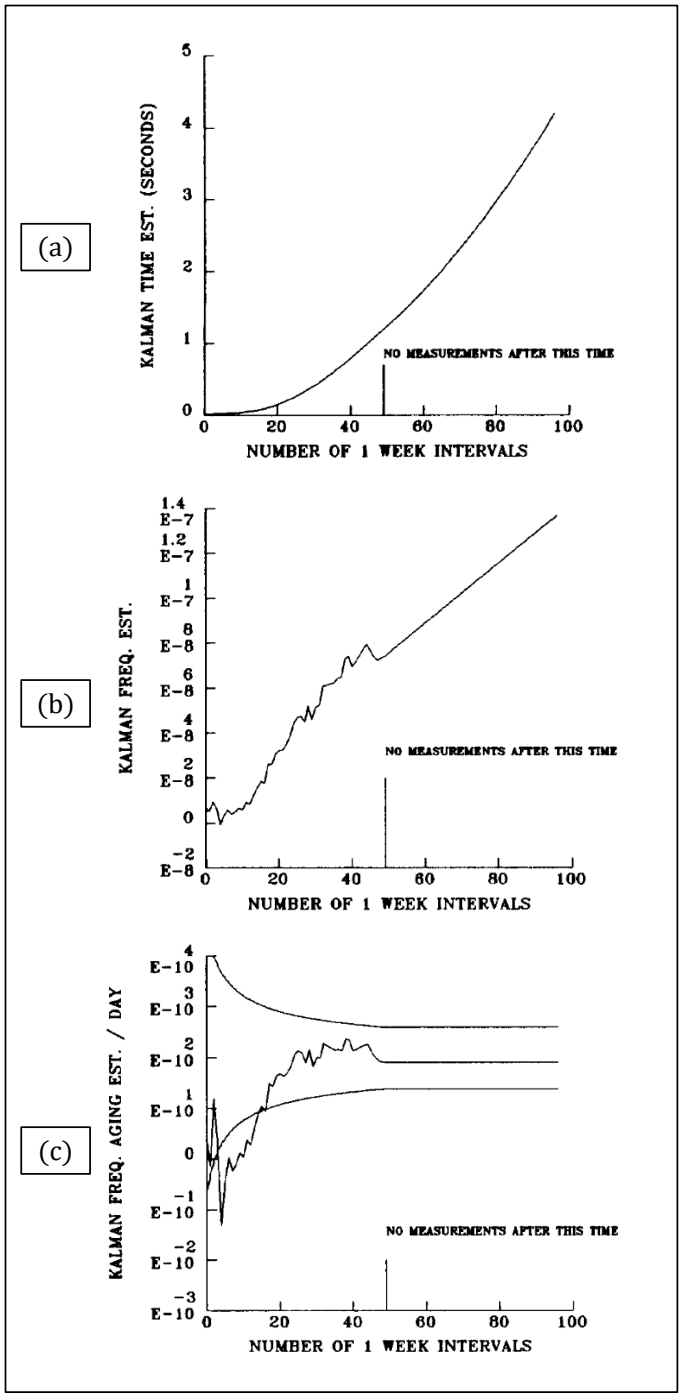

Figure 15: Kalman Filter estimate [28]

Another adaptive model uses a logarithmic description to consider non-constant frequency aging [29]. Instead of direct linear relationships present in Equation (19), the logarithmic equation introduced in Equation (7) was chosen as the starting model, specifically

for its ability to fit non-linear functions. The dynamic system equation for the adaptive logarithmic model is defined in Equation (32) [29]:

$$\begin{bmatrix}\eta_{k+1}\\ \alpha_{k+1}\end{bmatrix} = \begin{bmatrix}1 & log\frac{t_{k+1}}{t_k}\\ 0 & 1\end{bmatrix}\begin{bmatrix}\eta_k\\ \alpha_k\end{bmatrix} + \begin{bmatrix}0\\ \omega_k\end{bmatrix} \tag{24}$$

where η_k is the fractional frequency of the system's output, and α_k is the accumulation of prior η_k values. The total system error was ascribed to the gaussian white noise ω_k associated with the accumulation α_k. This dynamic system has the effect of implementing a summation of logarithmic functions, with a weighting factor applied to each. This allows the model to react to changes in an oscillator's aging behavior. Figure 16 shows the results of the logarithmic Kalman Filter. Note that this implementation can describe normal aging behaviors, as well as the sudden change in aging at 600 days. This type of dynamic prediction is not possible for static parameter models.

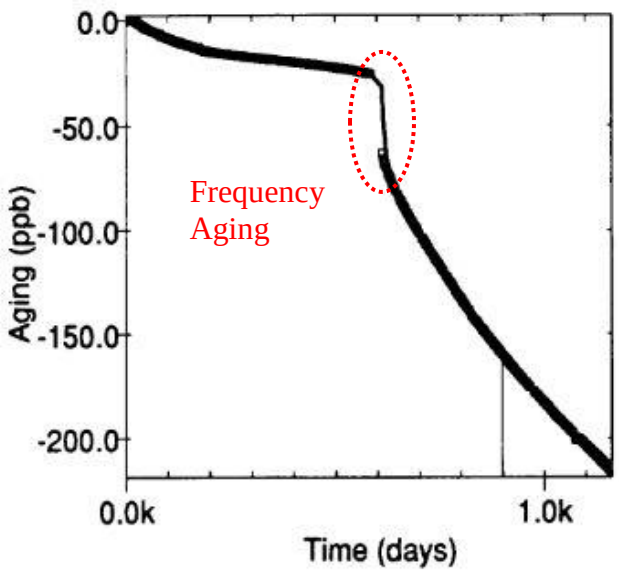

Figure 16: Logarithmic Kalman Filter example from [29]

2.2.4. Comparison of Aging Models

This section describes the process used in this thesis for model selection. Generally, comparison of models of any type can be divided into qualitative and quantitative categories. Qualitative model comparison evaluates the general trends of each model prediction against a representative data source. This involves visual examination of the two curves, noting where the model does and does not track details in the data. Qualitative comparisons introduce subjectivity into data analysis, which can be difficult to replicate between analysts. As such, quantitative comparisons which are calculated from each model output are preferred as a means of model performance evaluation.

It is often difficult to evaluate the relative performance of different models via direct curve comparison. As such, it is useful to reduce the overall model behavior to a single number for direct comparison. A common approach is the use of the coefficient of determination (R^2) as a metric of comparison, which is defined in Equation (25) [20]:

$$R^2 = \frac{\sum_{i=1}^{n}(\hat{y}_i - \bar{y})^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2} = 1 - \frac{\sum_{i=1}^{n}\hat{\varepsilon}_i^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2}, \qquad where\ 0 \leq R^2 \leq 1 \tag{25}$$

where y_i is the measured value at iteration i. $\bar{y}_i$ is the mean value of the measurements y_i. $\hat{y}_i$ is the value predicted by the model at iteration i. $\hat{\varepsilon}_i$ is the residual of the value predicted by the model at iteration i. R^2 is the complement of the squared sum of residuals divided by the squared sum of measurement to mean differences. Models with R^2 values approaching 1 have minimal output error; conversely, models with R^2 values approaching 0 have high output error.

The R^2 metric is useful for comparing model performances on a dataset, but there are some caveats to its applicability. The models being compared must consider the same estimated variable $\hat{y}_i$, in that no modification can be made to the value. Each model must have the same number of parameters, such as a model with the parameters a, b, and c cannot be compared to

another model with parameters a, b, c, and d. Lastly, the model must have a y-axis intercept. Fortuitously, these constraints are rarely a problem for models describing oscillator behavior.

2.3. Summary

In this literature search, there were three primary foci: oscillator frequency aging models, regression methods, and quantitative regression comparison metrics. The frequency aging models span a wide range of predictive abilities, from simple models with static parameters to adaptive models with dynamic parameters. Two regression techniques were introduced: the least squares method and the Kalman Filter algorithm. Lastly, the coefficient of determination was described as a model accuracy metric.

CHAPTER 3:

DATA COLLECTION AND METHODS

To study oscillator frequency aging models, representative oscillator frequency datasets are required. In these datasets, the frequency aging behavior must be well understood, such that dataset phenomena can be separated from model output phenomena. Additionally, an expectation of frequency aging behavior must be established to aid in frequency aging model comparative analysis.

3.1. Microchip 1000B and 1000C OCXOs

For the purposes of examining oscillator frequency aging models, phase and frequency datasets from Microchip 1000B and 1000C OCXOs were obtained for analysis. Eight of these OCXO datasets were provided by Microchip. These OCXOs are precise frequency sources, which can be stand-alone or used in systems with high performance requirements, such as the Microchip MHM-2010 Active Hydrogen Maser and other OEM systems. The design of the 1000X series OCXOs is the same between the 1000B and 1000C models, with minimal performance differences due to production changes. Figure 17 shows the chassis for the Microchip 1000X OCXO. Note that this chassis is from an earlier production run with older company branding.

Figure 17: Microchip 1000X chassis

To analyze OCXO data, a general notion of the expected frequency changes over the OCXOs lifespan must be understood. Figure 18 illustrates the general frequency behavior of these 1000X OCXOs. From the specifications for the Microchip 1000C [30], warm-up fractional frequency is specified to reach $\pm 2\ pp10^8$ per day after 15 minutes, and the expected duration for initial aging is 30 days. After this point the 1000C is expected to reach frequency stability $\pm 5\ pp10^9$ per day.

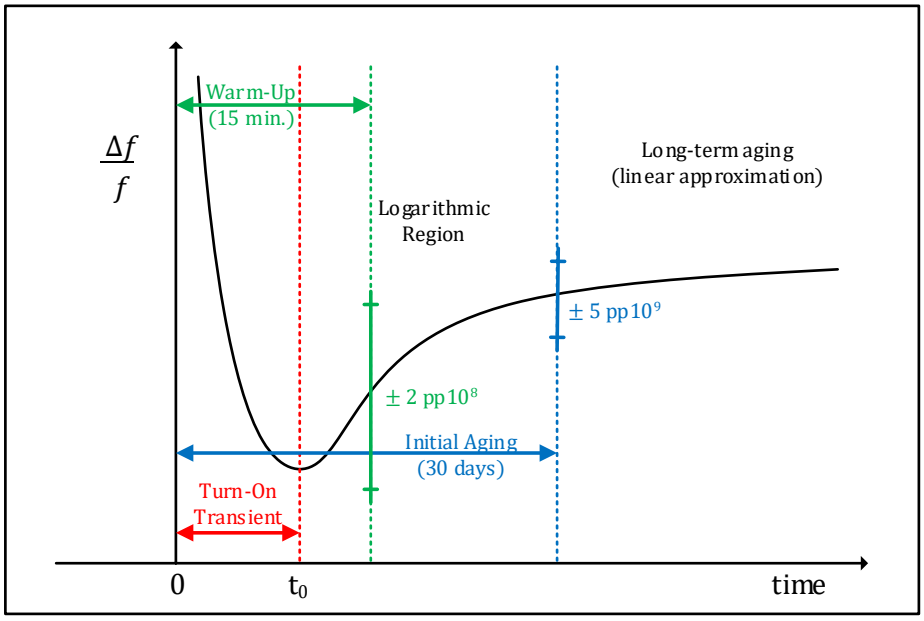

Figure 18: Expected fractional frequency behavior for the Microchip 1000C OCXO [18][30]

In the production process of the Microchip 1000X series, a variety of techniques are employed to minimize the aging rate of the complete OCXO. For example, the resonators undergo a lengthy pre-assembly burn-in process, during which they are kept powered and near operating temperature. After assembly, the completed oscillators are sequestered for a period of aging, during which phase and frequency measurements are taken in two distinct phases. The first phase consists of short-term measurements meant to help detect early issues with OCXO operation. This first evaluation phase generally lasts between 30 days and one year. The second phase is comprised of long-term measurements, during which each OCXO is qualified to its frequency aging specification. In addition, this period allows the OCXO to continue aging with the goal of reaching low aging rates. The long-term measurement period lasts from 90 days to multiple years. These production steps influence what types of aging behaviors may be present in each measurement set.

3.2. OCXO Datasets: Short-term Measurement Process

Data measurements during the first phase of analysis are recorded using the “Squid” Measurement System, which is described in [31]. The Squid is a clock measurement system which was developed specifically for exclusive use of Microchip and is not commercially available. This measurement setup is shown in Figure 19, with the Squid board mounted above the OCXOs. These measurements are taken in a laboratory environment without strict environmental controls on temperature, humidity, pressure, magnetic fields, etc.

Figure 19: Squid setup

The Squid Measurement System requires a precision frequency standard to use as a reference for making phase measurements of the OCXOs. In this case, an MHM-2010 Active Hydrogen Maser is used as the reference source. The squid measures the phase differences $x(t)$ in seconds between each oscillator under measurement and the MHM-2010. For Squid measurements reported in this thesis, the sample times (τ) between Squid phase measurements is configured at 8 seconds. The Squid also directly reports fractional frequency calculations as per Equation (4). For all Squid measurements reported in this thesis, the fractional frequency τ is selected as 16 seconds.

3.3. OCXO Datasets: Long-term Measurement Process

The second set of phase measurements are recorded using the Multi-Channel Measurement System (MMS) [32], which is a Microchip commercial clock measurement system. The setup used for the measurements described in this thesis is shown in Figure 20, with the MMS connected to the output of multiple OCXOs via coaxial cables. Just as in the first phase of analysis, measurements are taken in a laboratory environment without strict environmental controls on temperature, humidity, pressure, magnetic fields, etc. However, for this analysis,

insulation is added around the oscillators as shown in Figure 21. This insulation is added to improve the stability of the thermal environment during long-term aging.

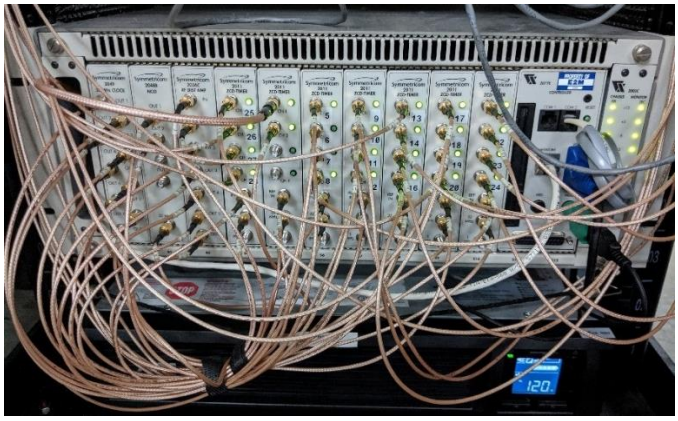

Figure 20: MMS setup

Figure 21: MMS setup – insulated OCXOs

The MMS requires a precision frequency standard to use as a reference for making phase measurements of the OCXOs. In this case, an MHM-2010 Active Hydrogen Maser is used as the reference source. The MMS the phase differences $x(t)$ in seconds between each oscillator under measurement and the MHM-2010. For MMS measurements reported in this thesis, the sample times (τ) between MMS phase measurements is configured at 1 second. These phase measurements are then averaged with a factor of 10 to minimize measurement noise. The MMS does not report fractional frequency measurements $y(t)$. As such, the $y(t)$ values must be

calculated from the phase differences $x(t)$ using Equation (4). For all MMS measurements reported in this thesis, the fractional frequency τ is selected as 10 seconds.

3.4. Selection of Data for Analysis

The eight OCXO Squid and MMS dataset pairs studied in this thesis were chosen from a significantly larger pool of data. To avoid over-fitting models to specific datasets, two requirements for each dataset were established. First, the datasets must have a start-to-end duration greater than 30 days. This requirement ensures that the end of each dataset will contain some amount of long-term aging. The second requirement is that at least 75% of the interval is sampled. This eliminates datasets with large gaps that can produce poor modeling results. These gaps arise from measurement interruptions, such as power outages. The OCXOs that were selected for this study are identified in Table 2, along with the duration of each measurement phase.

TABLE 2:
SELECTED OCXOS

OCXO Label	OCXO Type	Squid Duration (Days)	MMS Duration (Days)
A	1000C	71.7	195.6
B	1000C	63.6	258.6
C	1000C	55.0	69.5
D	1000C	253.0	112.6
E	1000C	33.9	34.6
F	1000B	36.1	35.7
G	1000B	62.8	132.6
H	1000B	43.8	105.6

3.5. Data Preprocessing

For frequency aging model analysis, the datasets must be preprocessed. First and foremost, the phase measurements from the Squid and MMS must be transformed to fractional frequency using Equation (4). This transformation exposes the frequency behavior of the OCXOs, but there are three issues remaining to consider: dataset gaps, outliers, and series starting point.

3.5.1. Dataset Gaps

Dataset gaps represent periods of missing measurements. A variety of situations can create data gaps. In the phase to frequency conversion process, repeated phase measurements result in a frequency gap. Other phase measurement gaps are from missing data, when an OCXO is removed from measurement temporarily or when there is a malfunction with the measurement system. For the datasets considered in this thesis, these gaps are represented with 1×10^{-99} instead of zero to avoid computation issues [17]. With a phase data gap of length n, the phase-frequency conversion results in a frequency gap of length $n+1$.

In this work, the 1×10^{-99} values are treated as a gap accounting measure, which are removed from the frequency datasets for analysis. This prevents the gaps from influencing model results.

3.5.2. Outlier Removal

This section details the process by which outliers are recognized and removed. Within the selected datasets, there exists spurious measurements that interfere with analysis. As such, it is common practice to remove outliers before data analysis [17]. One common method for outlier detection is the median absolute deviation (MAD), as shown in Equation (26) [17]:

$$MAD = Median\{|y(i) - m|/0.6745\} \tag{26}$$

where $y(i)$ is a measurement from a dataset y, m is the median value for a dataset y, and the factor 0.6745 scales the deviation to the standard deviation of the normal distribution. For each measurement in the dataset, the absolute difference from the median of a datapoint is found, then scaled based on the normal distribution. The median value of this resultant series is the MAD. A tolerance factor is used in conjunction with the MAD value offset from the median value m to set upper and lower bounds for outlier detection. The value typically used for tolerance factor is 5 [17]. This is shown in Equation (27) [17]:

$$Outlier\ Threshold = \pm[(5 \cdot \sigma) + m] \tag{27}$$

where σ is the MAD value, m is the median value for the dataset, and 5 is the tolerance factor.

Once outliers are detected, they are removed from the dataset by deleting both the measurement value and associated timestamp from the dataset. Figure 22 (a) shows an example of the outlier detection using MAD, with the outliers to be removed highlighted in red. In this OCXO dataset, the outliers indicated by the box are most readily attributed to a frequency jump of the oscillator caused by an internal malfunction. As such, in Figure 22 (b) these detected outliers are removed, providing a clearer view of the long-term frequency behavior of the dataset.

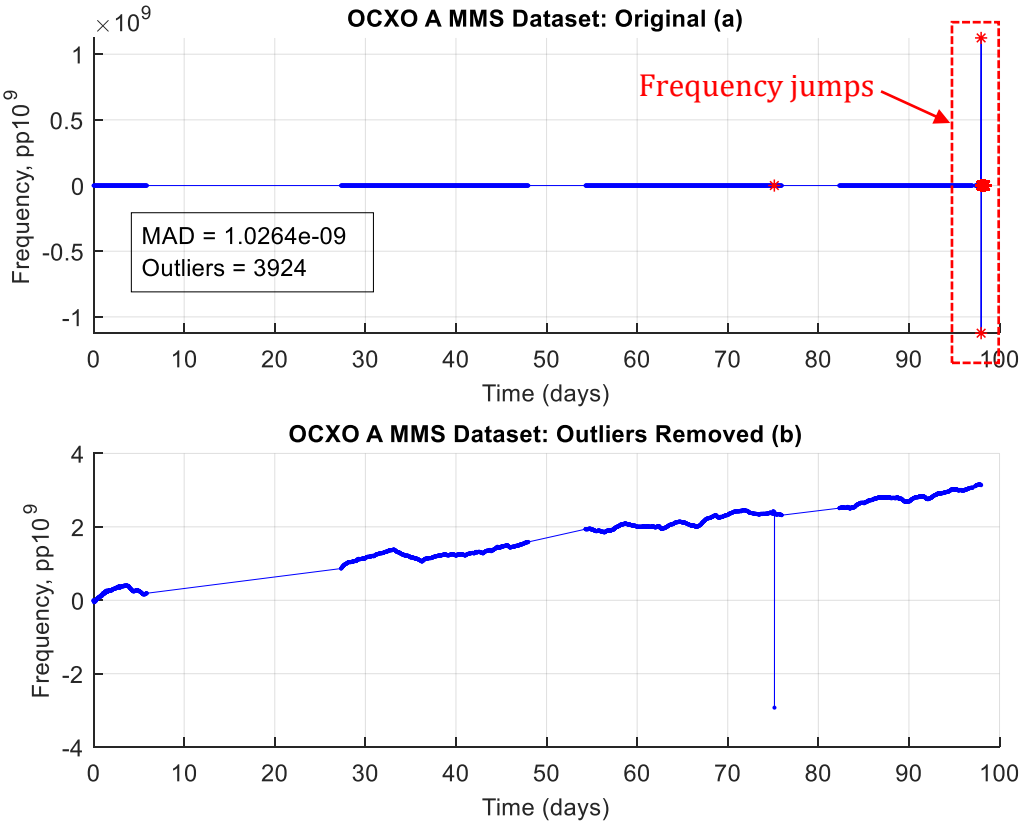

Figure 22: MAD outlier detection and removal on OCXO A MMS dataset

This outlier detection process also detects intervals containing retrace events as outliers. This is because the dramatic change in frequency value shifts the measurement mean. In this case, whichever measurement segment surrounding the retrace event contains the fewest number of samples will be removed. Typically, retrace events are removed from an analysis set before outlier detection, but in this thesis the outlier detection was used to identify these segments for removal.

3.5.3. Rebasing

As a final preprocessing step, each dataset is rebased with respect to origin. This involves taking the first data point from a data series and subtracting its value and timestamp from all

points in its series. This helps in presentation consistency, such that all data series plots start from the same value (zero).

3.6. Baseline Frequency Aging Estimation

With the datasets preprocessed, a method must be developed to find a baseline frequency aging estimate. This baseline aids in the comparison of processed datasets and will be highly useful in evaluating model accuracy results during the subsequent chapters of this thesis.

Due to the change of frequency aging over time, a baseline frequency aging value must be made with respect to the initial and long-term aging stages. This ensures that the aging rate of each stage does not subsume the other. As presented in Figure 18, the duration of initial aging is identified as 30 days for the 1000X series OCXO. Long-term aging is defined as the period after 30 days. For each dataset, a linear approximation is made on these intervals. The slope of the linear regression is taken as the baseline frequency aging for the interval. As the Squid data represents the early evaluation of 1000X performance, this data subset was selected to estimate initial aging. Conversely, the MMS dataset represents the long-term aging evaluation of 1000X performance, thus this subset was selected to estimate long-term aging. An example of this process is shown in Figure 23. Note the greater slope of the initial aging, which matches the expectation of oscillator frequency aging.

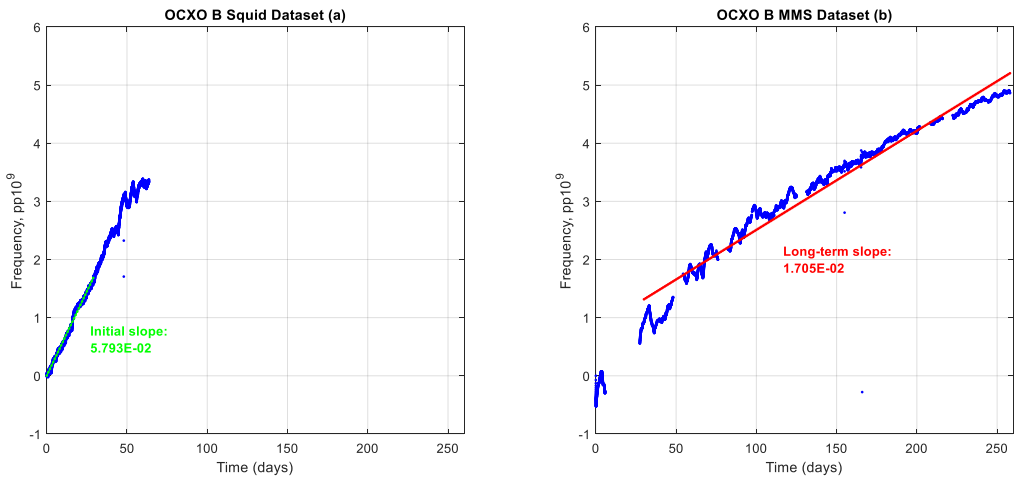

Figure 23: OCXO B dataset baseline frequency aging estimates

3.7. Processed Datasets

This section presents the 1000X OCXO frequency datasets A, B, and H that will be used for model fidelity analysis in this thesis. Special attention to unique frequency aging behavior will be noted here, which can be used as an important discriminator between models. Comparing the degree to which each model can predict these behaviors is a major component of model performance. The remaining datasets may be found in Appendix A.

Each dataset presented in Figure 24 through Figure 26 has undergone gap handling, outlier removal, and rebasing. The Squid data subsets are used to obtain an initial frequency aging baseline, and the MMS data subsets are used to obtain a long-term frequency aging baseline. Note that the MMS data subset does not necessarily directly follow the Squid subsets; there can be a significant time gap in-between.

3.7.1. OCXO A

Figure 24(a) presents a plot of the frequency behavior for OCXO A, which was measured over a period of 71.7 days utilizing the Squid system. A linear trend has been fit to the first 30 days of this dataset, which estimates an aging rate over this period of 6.446 $pp10^{11}$. The length of this dataset is average amongst the Squid-measured datasets. This dataset not contain any irregularities that suggest anomalous operation during the measurement window. In general, the first 30 days of Squid-measured frequency data for OCXO A appears to be well represented by a linear trend, which is consistent with the expectation of initial aging for a crystal oscillator. The logarithmic decrease in aging rate at the end of the dataset is typical of the transition from initial aging into long-term aging.

Figure 24(b) presents a plot of the frequency behavior for OCXO A, which was measured over a period of 195.6 days utilizing the MMS system. A linear trend has been fit to the last 165.6 days of this dataset, which estimates an aging rate over this period of $2.939\ pp10^{11}$. The length of this dataset is larger than most of the other MMS-measured datasets. This dataset contains anomalous behavior that manifests as short-term instability superimposed on the aging trend. This behavior is believed to be attributed to an internal oscillator failure, resulting in increased environmental sensitivity. The anomalous behavior potentially has minor impact on frequency aging analysis, since the long-term aging trend is not obscured. In general, the MMS-measured frequency data for OCXO A appears to be well represented by a linear trend, which is consistent with the expectation of long-term aging for a crystal oscillator.

Comparing the estimated aging rates for OCXO A detailed in Figure 24 corroborates the expected aging trends for crystal oscillators since the aging rate estimated from the Squid-measured data (short-term aging) is considerably higher than the aging rate estimated from the

MMS-measured data (long-term aging). The OCXO A dataset clearly displays typical frequency aging trends. Fitting each model to this dataset will provide an indication of each model's ability to describe standard frequency aging.

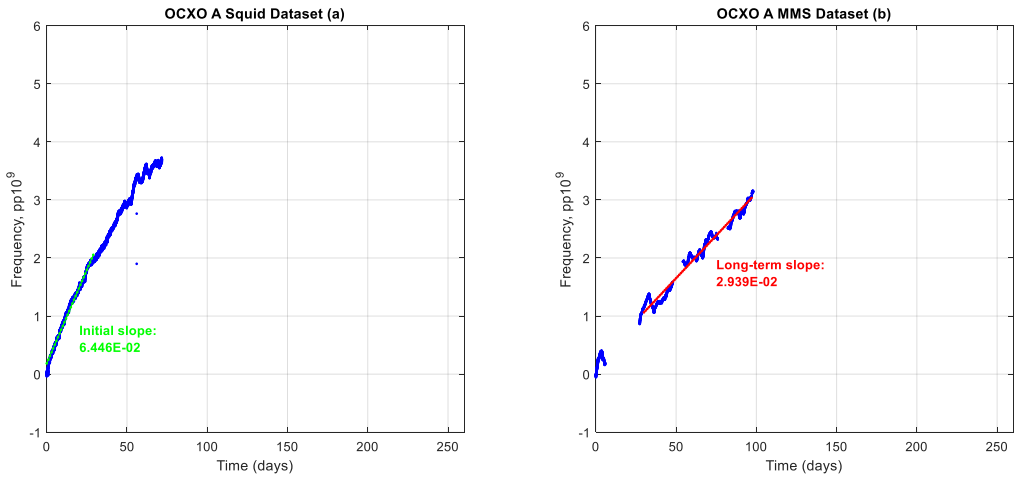

Figure 24: OCXO A dataset

3.7.2. OCXO B

Figure 25(a) presents a plot of the frequency behavior for OCXO B, which was measured over a period of 63.6 days utilizing the Squid system. A linear trend has been fit to the first 30 days of this dataset, which estimates an aging rate over this period of $5.793\ pp10^{11}$. The length of this dataset is average amongst the other Squid-measured datasets. This dataset does not contain any irregularities that suggest anomalous operation during the measurement window. In general, the first 30 days of Squid-measured frequency data for OCXO B appears to be well represented by a linear trend, which is consistent with the expectation of initial aging for a crystal oscillator. The logarithmic decrease in aging rate at the end of the dataset is typical of the transition from initial aging into long-term aging.

Figure 25(b) presents a plot of the frequency behavior for OCXO B, which was measured over a period of 258.6 days utilizing the MMS system. A linear trend has been fit to the last 228.6 days of this dataset, which estimates an aging rate over this period of $1.705\ pp10^{11}$. The length of this dataset is the largest of the MMS-measured datasets. This dataset does not contain any irregularities that suggest anomalous operation during the measurement window. In general, the MMS-measured frequency data for OCXO B appears to be well represented by a logarithmic trend, which is consistent with the expectation of the transition from initial to long-term aging for a crystal oscillator.

Comparing the estimated aging rates for OCXO B detailed in Figure 25 corroborates the expected aging trends for crystal oscillators since the aging rate estimated from the Squid-measured data (short-term aging) is considerably higher than the aging rate estimated from the MMS-measured data (long-term aging). The OCXO B dataset clearly displays typical frequency

aging trends. Fitting each model to this dataset will provide an indication of each model's ability to describe the transition from initial to long-term frequency aging.

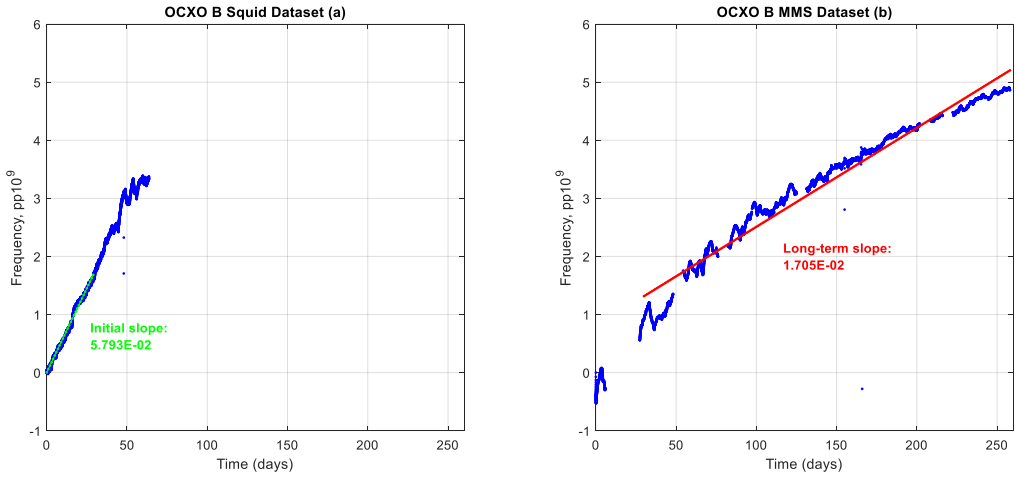

Figure 25: OCXO B datasets

3.7.3. OCXO H

Figure 26(a) presents a plot of the frequency behavior for OCXO H, which was measured over a period of 43.8 days utilizing the Squid system. A linear trend has been fit to the first 30 days of this dataset, which estimates an aging rate over this period of $-3.152\ pp10^{12}$. This is the only dataset to exhibit negative frequency aging. The length of this dataset is smaller than most of the other Squid-measured datasets. This dataset contains anomalous behavior that manifests as frequency jumps. This behavior is believed to be attributed to either an internal oscillator malfunction or counteracting aging mechanisms. Since the estimated aging rate is an order of magnitude smaller than the aging estimates of the other Squid and MMS-measured datasets, this anomaly may indicate that this oscillator has comparatively superior frequency aging behavior. The anomalous behavior potentially has major impact on frequency aging analysis. In general, the Squid-measured frequency data for OCXO H does not appear to be well represented by either a linear or logarithmic trend, due to the frequency jump anomalies.

Figure 26(b) presents a plot of the frequency behavior for OCXO H, which was measured over a period of 105.6 days utilizing the MMS system. A linear trend has been fit to the first 75.6 days of this dataset, which estimates an aging rate over this period of $7.886\ pp10^{12}$. The length of this dataset is average amongst the other MMS-measured datasets. This dataset does not contain any irregularities that suggest anomalous operation during the measurement window. In general, the MMS-measured frequency data for OCXO H appears to be well represented by a logarithmic trend, which is consistent with the expectation of initial to long-term aging for a crystal oscillator.

Comparing the estimated aging rates for OCXO H detailed in Figure 26 corroborates the expected aging trends for crystal oscillators since the aging rate estimated from the Squid-

measured data (short-term aging) is considerably higher than the aging rate estimated from the MMS-measured data (long-term aging). The OCXO H dataset does not exhibit typical frequency aging trends in the Squid-measured data but does clearly show initial to long-term frequency aging in the MMS-measured data. Fitting each model to this dataset will provide an indication of each model's ability to describe irregular frequency aging with small aging rates.

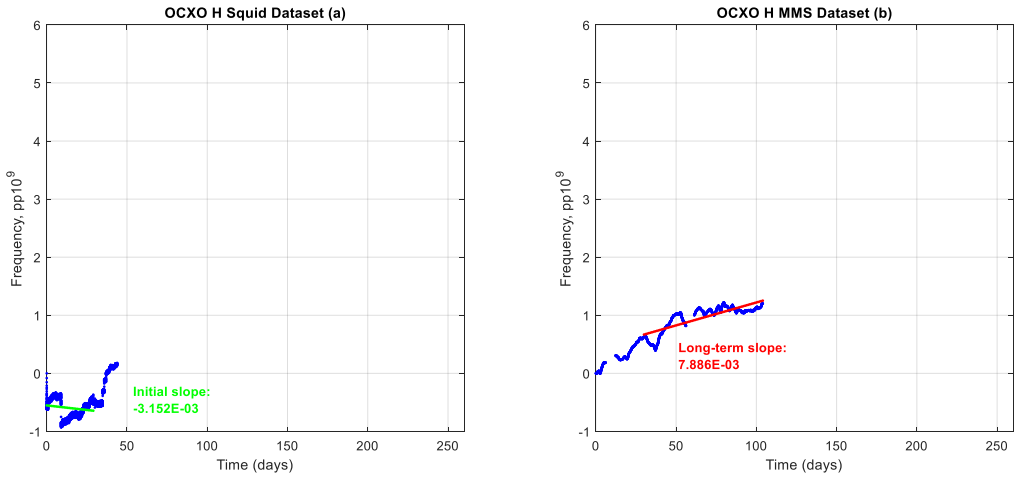

Figure 26: OCXO H datasets

Overall, each 1000X OCXO dataset is demonstrated to exhibit typical frequency aging, with exception of data shown in Figure 26(a) for OCXO H. Table 4 tabulates the baseline frequency aging estimates for each oscillator. Note that for each oscillator, the initial aging was estimated to be higher than the long-term aging, as expected for crystal oscillator frequency aging.

TABLE 3:
BASELINE FREQUENCY AGING ESTIMATES

OCXO Label	Initial Aging	Long-Term Aging
A	$6.446\ pp10^{11}$	$2.939\ pp10^{11}$
B	$5.793\ pp10^{11}$	$1.705\ pp10^{11}$
C	$8.312\ pp10^{11}$	$3.032\ pp10^{11}$
D	$2.234\ pp10^{11}$	$1.646\ pp10^{11}$
E	$5.568\ pp10^{11}$	$4.948\ pp10^{11}$
F	$3.905\ pp10^{11}$	$1.847\ pp10^{11}$
G	$2.844\ pp10^{11}$	$2.090\ pp10^{11}$
H	$-3.152\ pp10^{12}$	$7.886\ pp10^{12}$

3.8. Summary

In this work, analysis is conducted using multiple sets of 1000X OCXO phase data. This phase dataset is transformed into fractional frequency, such that frequency aging behaviors are exposed. The dataset is then preprocessed, such that analysis can focus on long-term trends instead of short disturbances. On the processed datasets, baseline frequency aging estimates are calculated. This preprocessing forms the groundwork for the comparative model analysis to follow in the subsequent chapters of this thesis.

CHAPTER 4:

APPLICATION AND EVALUATION OF AGING MODELS

The central component of this work is the analysis of two different modeling methods to be used on OCXO data in order to properly evaluate model performance. The Microchip 1000X datasets described in Chapter 3: are to serve as the analysis platform. A simple model and an adaptive model are chosen to illustrate the extrema of the modeling spectrum, with the goal of highlighting the differences between the approaches in high contrast. Subsequent to the evaluation of each model, the two techniques are compared.

4.1. Logarithmic Model

A logarithmic model is chosen as the first model for examination. The logarithmic aging model [21] is discussed in Section 2.2.1 of this thesis. One common method of using this model to describe the aging behavior of an oscillator is to fit the model parameters to fractional frequency data using the non-linear least squares method. The curve fitting process searches possible values for model parameters A, B, and C using Equation (5) to find the values that minimize error between the fitting equation and the measured data. The resulting system of Equation (7) and the identified values for A, B, and C can then model the OCXO's frequency aging behavior, in that it can estimate current aging behavior or predict future aging behavior.

A standard process by which the least squares method is applied for the logarithmic model is the "Modified Grid Search Method" [21]. In this method, linearization of the logarithmic model is employed to isolate model parameters, such that an initial parameter estimate is obtained via regression.

These initial parameter estimates are used to improve the likelihood that the least squares method will converge. The implementation of the parameter search algorithm in this work follows the Modified Grid Search Method.

4.1.1. Logarithmic Modeling Process

To illustrate the Modified Grid Search Method, this process will be applied to the OCXO B MMS dataset described in Chapter 3:. To make an initial estimation for parameters A, B, and C, the data is first fit to a linear model using the least squares method. The estimate of A is found from the slope of the linear model fit. The estimate of B is found via relating the slope and y-intercept of the linear model to the linearized logarithmic model, as shown in Equation (28) [21]:

$$B = 10^{\frac{b}{A}} \tag{28}$$

where b is the y-intercept of the linear model fit. The estimate of C is assumed to be 0. Figure 27 shows the results of this first step.

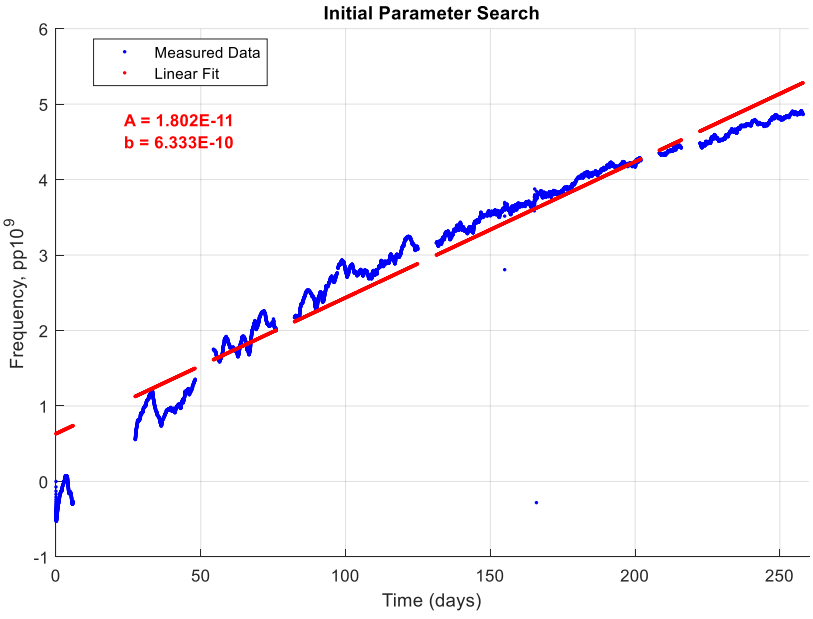

Figure 27: OCXO B MMS initial parameter search

Using the initial estimations for A, B, and C, the logarithmic model of Equation (7) is populated and compared to the measured frequency values. Figure 28 presents a comparison of this initial model output and the measured fractional frequency data for OCXO B. By visual inspection, the agreement between the model output and the data series demonstrated in this figure is poor. The initial model output exhibits a slope close to zero, whereas the measured data clearly shows increasing fractional frequency with increasing time. Notwithstanding this initial disagreement, the initial parameter estimates will assist in the search for the optimal parameter estimates.

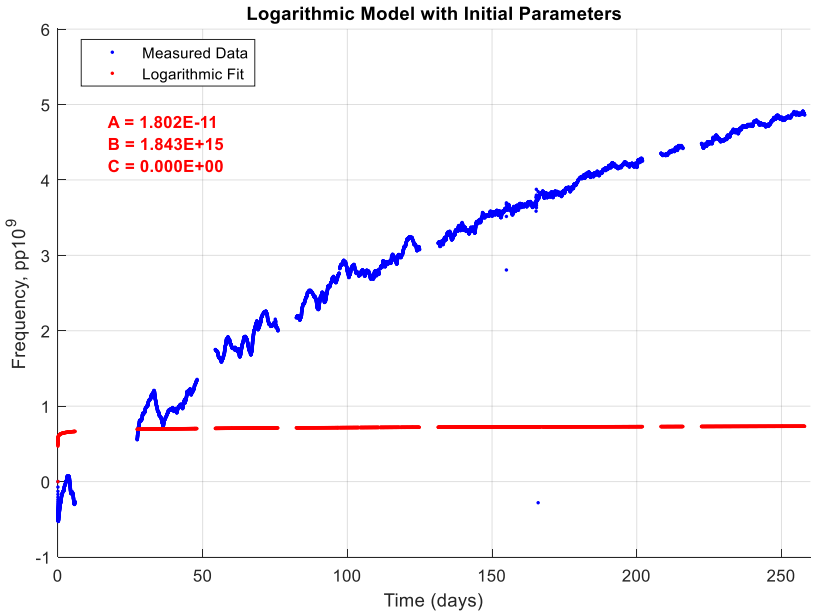

Figure 28:OCXO B MMS initial logarithmic fit

The last step of the "Modified Grid Search Method" involves updating the values of parameters A, B, and C by applying the least squares parameter estimation to Equation (7) and using the initial values for A, B, and C identified in the first step. Figure 29 displays the results of completing this process for the OCXO B dataset. It should be noted that Figure 29

demonstrates better model agreement than Figure 28, in that the amplitude at the end of the dataset is more representative of the input fractional frequency data.

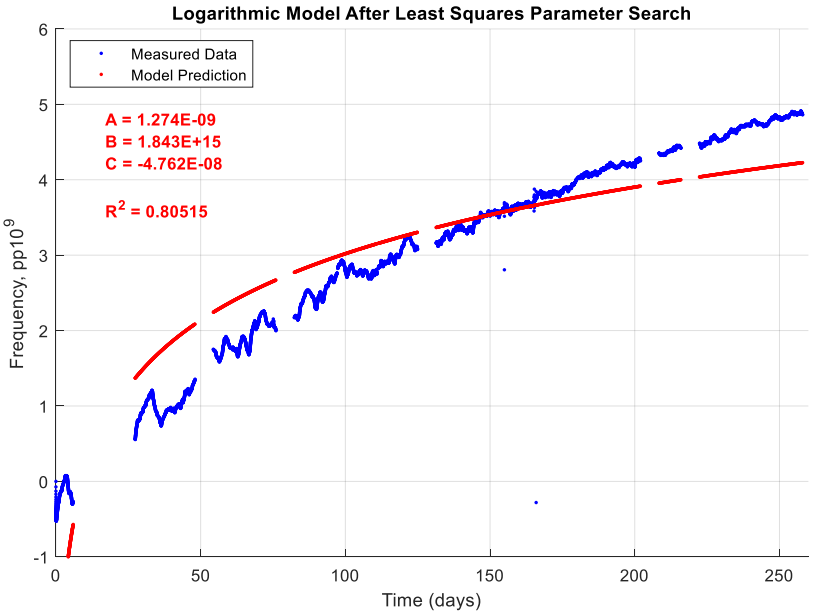

Figure 29: OCXO B MMS logarithmic model after Least Squares Parameter Search

4.1.2. Logarithmic Modeling Results

The Modified Grid Search method was utilized to fit all eight OCXO frequency datasets introduced in Chapter 3: to the logarithmic aging model. For brevity, only three of the eight datasets are detailed in this section. The three selected datasets are OCXO A, B, and H. The details for all eight datasets can be found in Appendix B. The parameter values used to generate the model output are listed in figure (a) of each OCXO data subset. For each Squid-measured and MMS-measured OCXO data subseries, two comparisons will be made in this section. The first directly compares the measured frequency data against a corresponding logarithmic model fit. The second compares the calculated frequency aging predicted by the model against the baseline frequency aging estimate for each OCXO data subset, which was obtained as described in Section 3.6. After these comparisons, a residual analysis is performed to ascertain how well the logarithmic model captures the frequency aging behavior of each data subset.

4.1.2.1. Logarithmic Modeling: OCXO A

Figure 30(a) displays the results of logarithmic modeling performed on the Squid-measured OCXO A dataset. On this interval, the measured frequency exhibits an approximately linear trend, whereas the fitted curve exhibits the purely logarithmic trend of the model. The R^2 metric for this measurement-model comparison is 0.81948. In general, this model fit tracks the long-term frequency aging behavior without notable influence from short-term effects. Figure 30(b) displays the calculated frequency aging behavior of the model output against the baseline initial aging estimate. The model fit for initial aging does not agree with the baseline estimate. Figure 30(c) displays the results of logarithmic modeling performed on the MMS-measured OCXO A dataset. On this interval, the measured frequency exhibits an approximately linear trend, whereas the fitted curve exhibits the purely logarithmic trend of the model. The R^2 metric for this measurement-model comparison is 0.70239. In general, this model fit tracks the long-term frequency aging behavior without notable influence from short-term effects. Figure 30(d) displays the calculated frequency aging behavior of the model output against the baseline long-term aging estimate. The model fit for long-term aging does not agree with the baseline estimate.

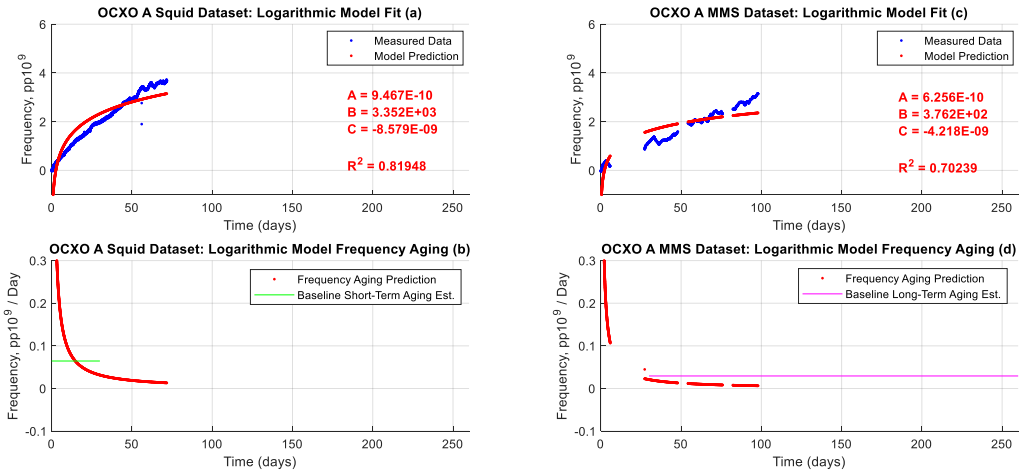

Figure 30: OCXO A logarithmic modeling results

Figure 31(a) contains the residuals from the logarithmic model fit to the Squid-measured dataset. On the measurement interval, there is a discernible logarithmic trend in the residuals. The remaining non-random structure of the residuals indicates that the logarithmic model does not fully capture the behaviors of the Squid-measured dataset. Figure 31(b) contains the residuals from the logarithmic model fit to the MMS-measured dataset. On the measurement interval, there is a discernible logarithmic trend in the residuals. The remaining structure of the residuals indicates that the logarithmic model does not fully capture the behaviors of the MMS-measured dataset.

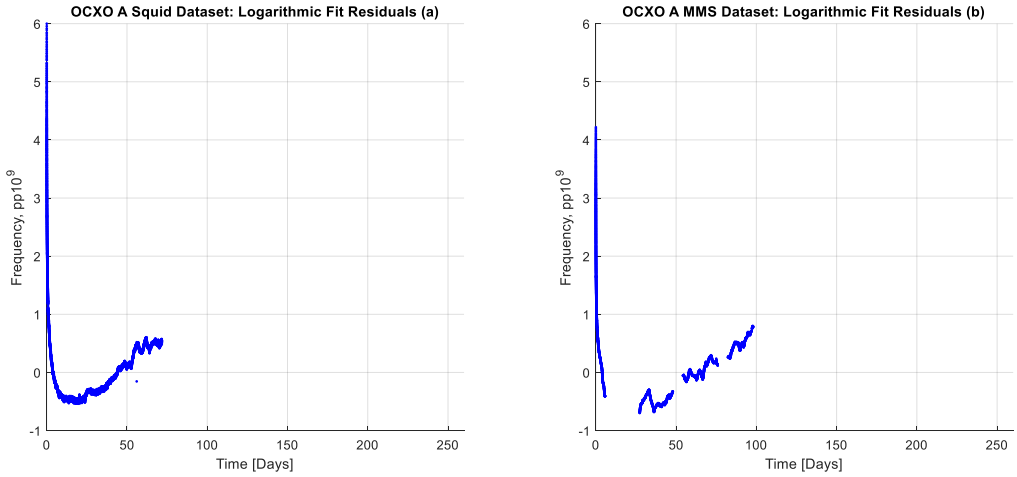

Figure 31: OCXO A logarithmic modeling residuals

For the initial aging behavior contained in the OCXO A Squid-measured dataset and the long-term aging in the OCXO A MMS-measured dataset, the logarithmic model demonstrates a poor fit to the measured data.

4.1.2.2. Logarithmic Modeling: OCXO B

Figure 32(a) displays the results of logarithmic modeling performed on the Squid-measured OCXO B dataset. On this interval, the measured frequency exhibits an approximately linear trend, whereas the fitted curve exhibits the purely logarithmic trend of the model. The R^2 metric for this measurement-model comparison is 0.82378. In general, this model fit tracks the long-term frequency aging behavior without notable influence from short-term effects. Figure 32(b) displays the calculated frequency aging behavior of the model output against the baseline initial aging estimate. The model fit for initial aging does not agree with the baseline estimate. Figure 32(c) displays the results of logarithmic modeling performed on the MMS-measured OCXO B dataset. On this interval, the measured frequency exhibits a quasi-logarithmic trend, whereas the fitted curve exhibits the purely logarithmic trend of the model. The R^2 metric for this measurement-model comparison is 0.80515. In general, this model fit tracks the long-term frequency aging behavior without notable influence from short-term effects. Figure 32(d) displays the calculated frequency aging behavior of the model output against the baseline long-term aging estimate. The model fit for long-term aging loosely approximates the baseline estimate.

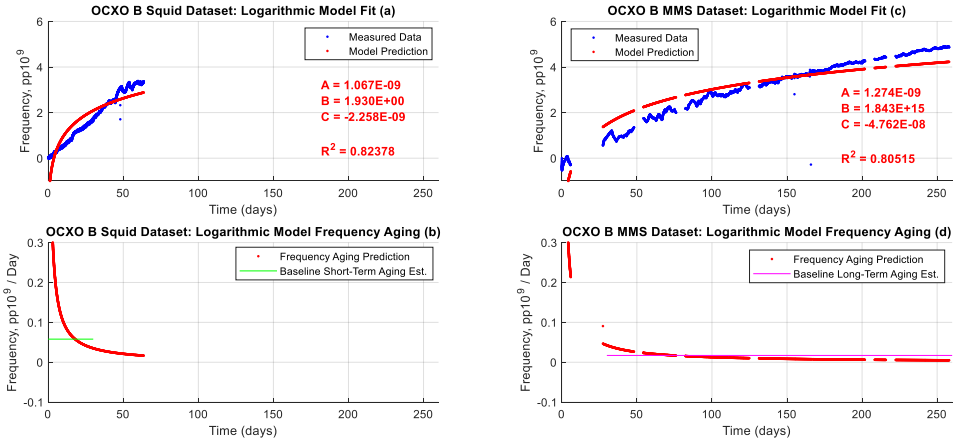

Figure 32: OCXO B logarithmic modeling results

Figure 33(a) contains the residuals from the logarithmic model fit to the Squid-measured dataset. On the measurement interval, there is a discernible logarithmic trend in the residuals. The remaining non-random structure of the residuals indicates that the logarithmic model does not fully capture the behaviors of the Squid-measured dataset. Figure 33(b) contains the residuals from the logarithmic model fit to the MMS-measured dataset. On the measurement interval, there is a discernible linear trend with a positive slope in the residuals. The remaining structure of the residuals indicates that the logarithmic model does not fully capture the behaviors of the MMS-measured dataset.

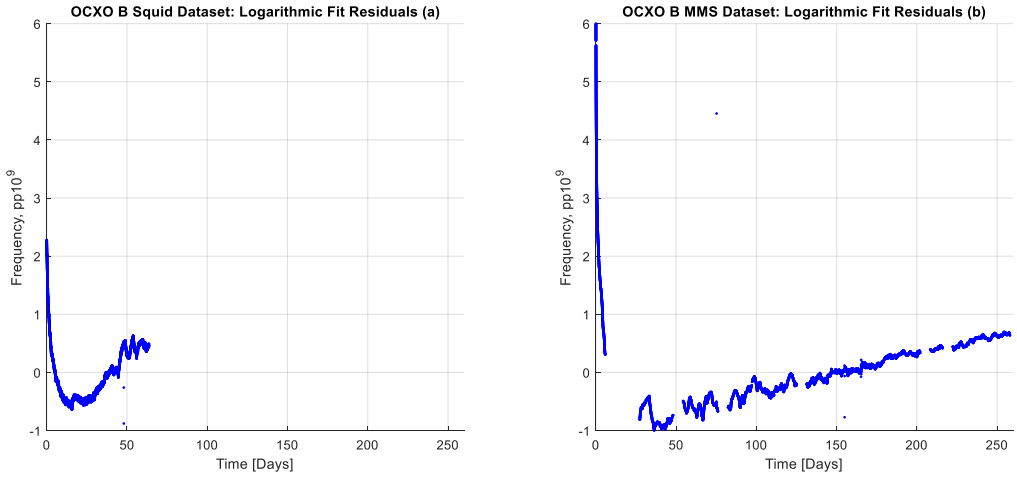

Figure 33: OCXO B logarithmic modeling residuals

For the initial aging behavior contained in the OCXO B Squid-measured dataset and the long-term aging in the OCXO B MMS-measured dataset, the logarithmic model demonstrates a poor fit to the measured data.

4.1.2.3. Logarithmic Modeling: OCXO H

Figure 34(a) displays the results of logarithmic modeling performed on the Squid-measured OCXO H dataset. On this interval, the measured frequency exhibits anomalous behavior, whereas the fitted curve exhibits the purely logarithmic trend of the model. The R^2 metric for this measurement-model comparison is 0.12878. In general, this model fit tracks the long-term frequency aging behavior without notable influence from short-term effects. Figure 34(b) displays the calculated frequency aging behavior of the model output against the baseline initial aging estimate. The model fit for initial aging does not agree with the baseline estimate. Figure 34(c) displays the results of logarithmic modeling performed on the MMS-measured OCXO H dataset. On this interval, the measured frequency exhibits a quasi-logarithmic trend, whereas the fitted curve exhibits the purely logarithmic trend of the model. The R^2 metric for this measurement-model comparison is 0.77798. In general, this model fit tracks the long-term frequency aging behavior without notable influence from short-term effects. Figure 34(d) displays the calculated frequency aging behavior of the model output against the baseline long-term aging estimate. The model fit for long-term aging agrees quite well with the baseline estimate.

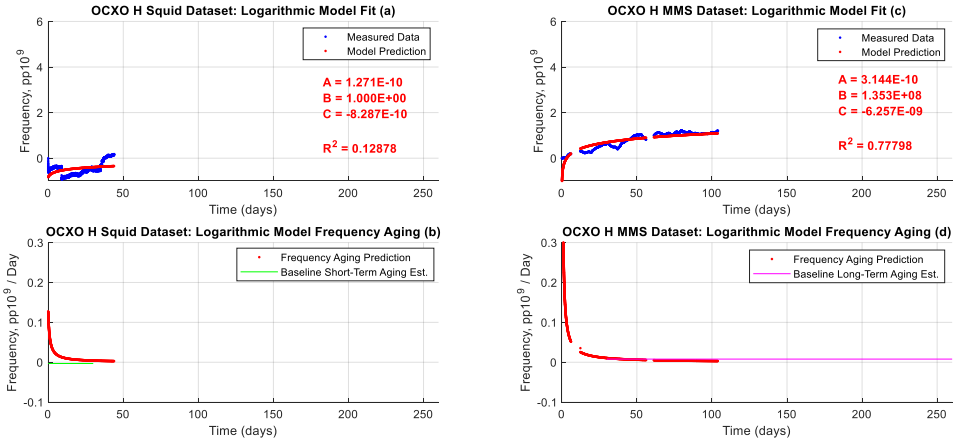

Figure 34: OCXO H logarithmic modeling results

Figure 35(a) contains the residuals from the logarithmic model fit to the Squid-measured dataset. On the measurement interval, the residuals exhibit the same shape as the measured dataset. The remaining non-random structure of the residuals indicates that the logarithmic model could not adequately describe the anomalous Squid-measured dataset. Figure 35(b) contains the residuals from the logarithmic model fit to the MMS-measured dataset. On the measurement interval, there is a slightly discernible linear trend with a positive slope in the residuals. The remaining structure of the residuals indicates that the logarithmic model does better in this case compared to the previous examples but still does not fully capture the behaviors of the MMS-measured dataset.

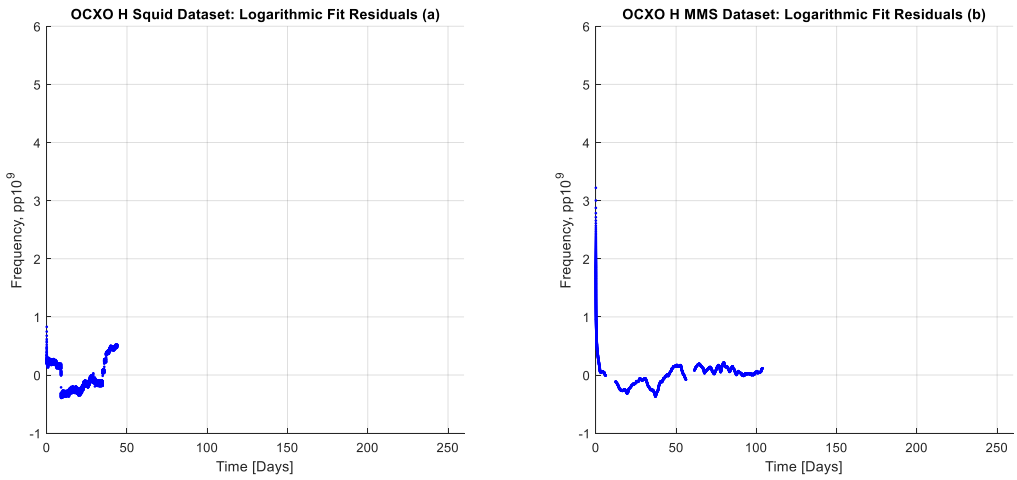

Figure 35: OCXO H logarithmic modeling residuals

For the initial aging behavior contained in the OCXO H Squid-measured dataset and the long-term aging in the OCXO B MMS-measured dataset, the logarithmic model demonstrates a poor fit to the measured data.

4.1.2.4. Logarithmic Modeling: Coefficient of Determination Summary

For all eight datasets, the coefficient of determination results are tabulated in Table 4. Note that the OCXO H Squid dataset demonstrates the lowest coefficient of determination value, where anomalous behavior dominated the measured fractional frequency data.

TABLE 4:
LOGARITHMIC MODEL COEFFICIENT OF DETERMINATION RESULTS

OCXO	Squid R^2	Fit Quality	MMS R^2	Fit Quality
A	0.81948	Poor	0.70239	Poor
B	0.82378	Poor	0.80515	Poor
C	0.82485	Poor	0.76085	Poor
D	0.73887	Good	0.77359	Good
E	0.83629	Good	0.83133	Poor
F	0.84006	Good	0.81445	Good
G	0.73467	Poor	0.77598	Poor
H	0.12878	Poor	0.77798	Good

4.1.3. Logarithmic Modeling Discussion

To evaluate the logarithmic model's overall performance, this section focuses on three views of performance. The first view evaluates the degree to which the computed frequency aging agrees with the baseline linear estimates of aging behavior. The second view centers on the R^2 error metric. The third view considers the results of residual analysis between the measured and fitted data.

Visually inspecting the overall agreement of each logarithmic fit to the short-term linear estimate, the aging estimates did not agree. There are two reasons for this. The first is that the length of the short-term aging is dwarfed by the length of long-term frequency aging data. The second is that the dramatically higher short-term aging rate is not well represented by a single-valued estimate. The logarithmic model can intersect the line representing the short-term aging estimate but is strongly influenced by the long-term aging plateau. The logarithmic model generally underpredicts the long-term aging of the OCXOs as estimated by the long-term linear aging estimates. In short, the linear frequency aging estimates indicate that the models approximate the baseline aging rates for the intervals under evaluation, but with non-negligible disagreement.

From the view of the R^2 values, the logarithmic regression applied to most of the OCXO datasets produces fits with relatively small error, as indicated by values above 0.7. The worst performance observed for the logarithmic modeling process is for the Squid-measured OCXO H data subset with a value of 0.12878. Using R^2 as the only figure of merit for logarithmic model performance, the most insight that can be obtained is the identification of datasets which contain aging anomalies.

The residual analysis of these logarithmic model fits indicates a lack of flexibility of the logarithmic function to describe the frequency aging datasets. The residual plots for all datasets contained a significant systematic component. These systematic components indicate that the logarithmic model is too simple for the observed behavior.

Combining the linear estimate comparison, R^2 metrics, and residual analysis provides a more complete picture of the logarithmic model performance. The linear estimate offsets demonstrate setpoint error within the model. The R^2 metrics highlight datasets with anomalous behavior but does not particularly help distinguish fit performance between datasets. The behavior discrepancy highlighted by the residual analysis is due to the logarithmic model's inability to model complex aging behaviors.

This performance analysis indicates that the logarithmic model generally poorly fits the frequency data. Since the logarithmic model is a simple model, only one type of aging mechanism is described. As such, the logarithmic model performs poorly when multiple aging mechanisms are present. To gain improvements in the model performance of the frequency aging model, a different strategy is required.

4.2. Linear Kalman Filter Model

To consider a vastly different strategy for modeling OCXO aging, the adaptive class of frequency aging models was selected. The dynamic parameter feature of adaptive models is a possible solution to address the lack of behavioral flexibility observed with the logarithmic model. Specifically, the linear Kalman Filter introduced in Section 2.2.3 was selected as the candidate adaptive model framework for consideration.

4.2.1. Linear Kalman Filter Modeling Process

The linear Kalman Filter uses a set of three linear difference equations to describe the output behavior of an oscillator, as shown in Equation (19). These equations form a dynamic system representation of the oscillator. The three state variables considered in this dynamic system are output phase, fractional frequency, and frequency aging. Before the Kalman Filter can estimate these quantities, covariance values and initial conditions are needed to complete the system model.

4.2.1.1. Linear Kalman Filter: Covariance Calculation

The covariance of the system model and associated measurement inputs is an important component of the Kalman Filter implementation. Selecting proper covariance values helps suppress noise content, as well as filter data outliers. For the calculation of system and measurement noise in precision timing systems, the concept of power law noise is typically employed. Power law noise types model the noise of a system using the relation in Equation (29) [17]:

$$S_y(f) = h_\alpha f^\alpha \tag{29}$$

where f is the sideband frequency in Hertz, h_α is the coefficient of the noise process, and α is the exponent of the power law noise. Integer values of α utilized to categorize noise contributors

into the specific types listed in Table 5 [17]. Coefficients associated with these noise types can be obtained by fitting a power law model to a plot of Allan Deviation (ADEV) for a given set of fractional frequency or phase data.

TABLE 5:
POWER LAW NOISE TYPES

α	Noise Type
2	White Phase Modulated (W PM)
1	Flicker PM
0	White Frequency Modulated (W FM)
-1	Flicker FM
-2	Random Walk FM
-3	Flicker Walk FM
-4	Random Run FM (RR FM)

For the measurement covariance matrix R, the Kalman Filter expects the frequency variance associated with frequency measurements to have a gaussian distribution. This frequency variance may be calculated using white noise frequency variance in Equation (30) [33]:

$$R \;=\; \sigma_y^2(\tau) = \frac{h_0}{2} \cdot |\tau|^{-1} \tag{30}$$

where h_0 is the power spectral density coefficient associated with White FM noise and τ is the sample period of the dataset.

A value for the Q system covariance matrix for OCXOs is also required for proper configuration of the Kalman Filter. For the linear Kalman Filter Model, the system covariance matrix is defined in Equation (31) [28]:

$$Q_k(\delta) = \begin{bmatrix} S_\xi\delta + S_\mu\delta^3/3 + S_\zeta\delta^5/20 & S_\mu\delta^2/2 + S_\zeta\delta^4/8 & S_\zeta\delta^3/6 \\ S_\mu\delta^2/2 + S_\zeta\delta^4/8 & S_\mu\delta + S_\zeta\delta^3/3 & S_\zeta\delta^2/2 \\ S_\zeta\delta^3/6 & S_\zeta\delta^2/2 & S_\zeta\delta \end{bmatrix} \tag{31}$$

where δ is the sample interval, S_ξ is the spectral density of phase noise, S_μ is the spectral density of frequency noise, and S_ζ is spectral density of frequency aging noise. S_ξ is found from the

white frequency modulated power spectral density coefficient h_0. S_μ is found by the relation in Equation (32) [28]:

$$S_\mu = (2\pi)^2 h_{-2} \tag{32}$$

where h_{-2} is the power spectral density coefficient associated with random walk FM noise. S_ζ is equal to the power spectral density coefficient associated with random run FM noise. It is important to note that it can take a long time to create an accurate estimate of the power spectral density coefficient for Random Run FM noise. As such, minor adjustments to this value may improve KF estimates.

For the model implementation presented in this section, the measurement covariance matrices are based on the power law noise coefficients found in Table 6. The White FM noise for the Squid and MMS measurement systems was calculated using their respective noise floor measurements. The 1000X OCXO power law noise coefficients were determined using the OCXO B MMS dataset, since it is the longest data subset and is considered representative of all analyzed oscillators.

TABLE 6:
POWER LAW NOISE COEFFICIENTS

Device	Coefficient	Value
Squid	h_0 (White FM)	$9.715E-24$
MMS	h_0 (White FM)	$2.963E-27$
1000X OCXO	h_0 (White FM)	$1.327E-18$
	h_{-2} (Random Walk FM)	$5.408E-25$
	h_{-4} (Random Run FM)	$2.992E-27$

4.2.1.2. Linear Kalman Filter: Initial Conditions

Proper initial conditions assist the Kalman Filter in producing meaningful and accurate estimates. For this work, the initial state vector x_0 values were selected as follows. The initial value for Phase is set to zero since phase is an accumulation from a particular instant. The initial

value for fractional frequency is set to the first input value. The initial value for frequency aging is set to zero since aging is not instantaneous.

The error covariance matrix P_k is dependent purely on the dynamic system model, not on data input. As such, the values in P_k can be iterated until steady state has been reached.

4.2.1.3. Linear Kalman Filter: Model Output

The Kalman Filter iterates on each time interval, using the dynamic system in Equation (19) with the system measurement from that interval to produce a state estimate. This process continues until the iterations are interrupted. The result of this Kalman Filtering process is shown in Figure 36, which uses the OCXO B dataset to illustrate the evolution of the three state variables over time. In this example, only fractional frequency samples are supplied to the Kalman Filter. The phase and frequency aging values are estimated from the frequency measurements.

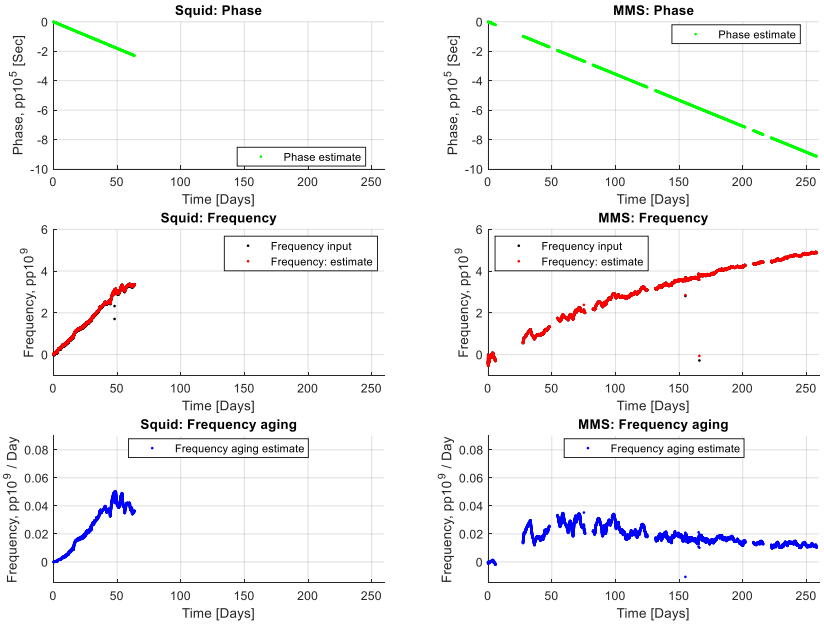

Figure 36: OCXO B Kalman Filter estimates

4.2.2. Linear Kalman Filter Modeling Results

The linear Kalman Filter was utilized to fit the OCXO frequency datasets to the linear dynamic system aging model. For brevity, three of the eight datasets are detailed in this section. The three selected datasets are OCXO A, B, and H. The details for all eight datasets can be found in Appendix C. For each Squid-measured and MMS-measured OCXO data subseries, two comparisons will be made in this section. The first directly compares the measured frequency data against its corresponding Kalman Filter estimate. The second compares the frequency aging predicted by the Kalman Filter against the baseline frequency aging estimate for each OCXO data subset, which was obtained as described in Section 3.6. After the comparisons, a residual analysis will be performed to ascertain how well the Kalman Filter estimates captured the frequency aging behavior of each data subset.

4.2.2.1. Linear Kalman Filter Modeling: OCXO A

Figure 37(a) displays the results of Kalman Filter modeling performed on the Squid-measured OCXO A dataset. On this interval, the measured frequency exhibits an approximately linear trend, with the Kalman Filter predicted values closely tracking the measurements. The R^2 metric for this measurement-model comparison is 0.99996. In general, this model fit tracks the long-term frequency aging behavior while filtering short-term effects. Figure 37(b) displays the frequency aging behavior of the model output against the baseline initial aging estimate. The model fit for initial aging demonstrates lower magnitude aging than the baseline estimate. Figure 37(c) displays the results of Kalman Filter modeling performed on the MMS-measured OCXO A dataset. On this interval, the measured frequency exhibits an approximately linear trend, with the Kalman Filter predicted values closely tracking the measurements. The R^2 metric for this measurement-model comparison is approximately 1. In general, this model fit tracks the long-term frequency aging behavior while filtering short-term effects. Figure 37(d) displays the frequency aging behavior of the model output against the baseline long-term aging estimate. The model fit for long-term aging agrees with the baseline estimate, with some modest variation.

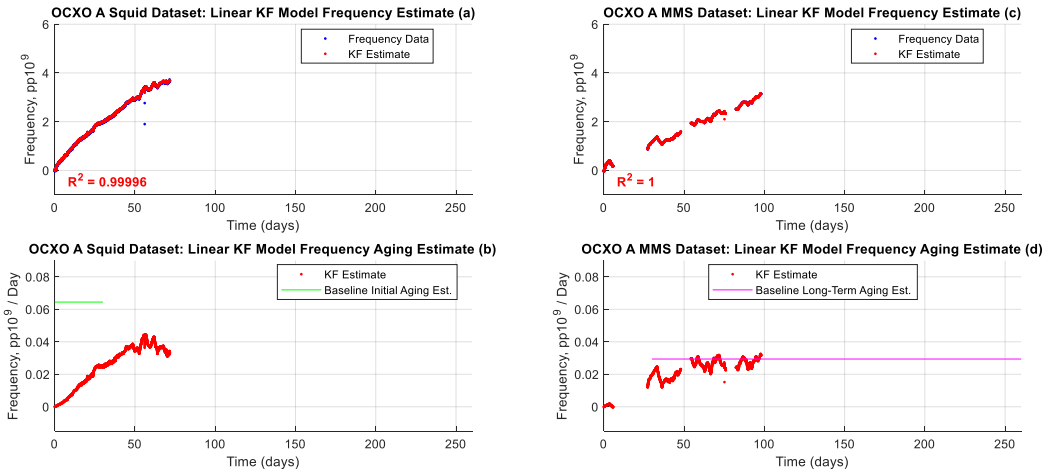

Figure 37: OCXO A linear Kalman Filter modeling results

Figure 38(a) contains the residuals from the Kalman Filter frequency estimate to the Squid-measured dataset. On the measurement interval, there is not a discernable residual pattern. The lack of pattern indicates that the model fully describes the fractional frequency behavior present. Figure 38(b) contains the residuals from the Kalman Filter frequency estimate to the MMS-measured dataset. On the measurement interval, there is not a discernable residual pattern. The lack of pattern indicates that the model fully describes the fractional frequency behavior present.

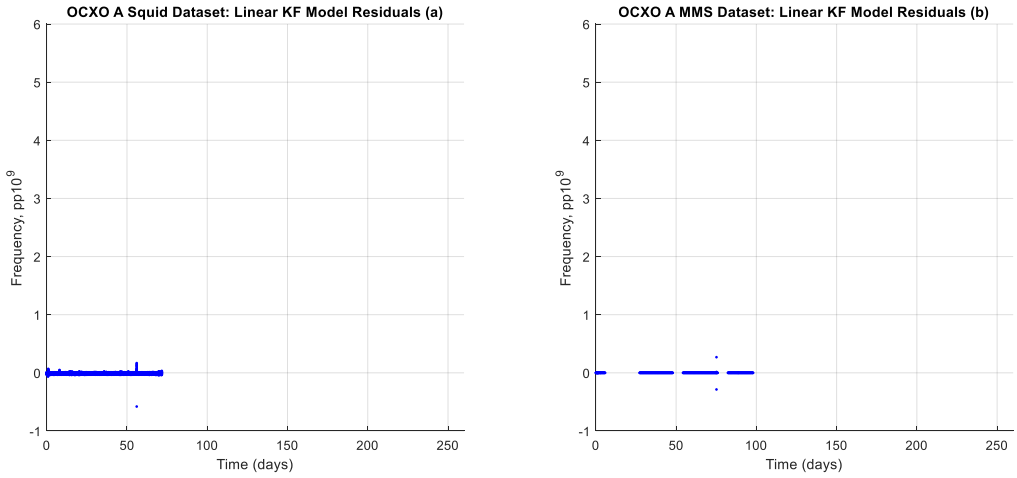

Figure 38: OCXO A linear Kalman Filter modeling residuals

For the initial aging behavior contained in the OCXO A Squid-measured dataset, the Kalman Filter estimate produces a good fit. For the long-term aging in the OCXO A MMS-measured dataset, the Kalman Filter estimate produces a good fit.

4.2.2.2. Linear Kalman Filter Modeling: OCXO B

Figure 39(a) displays the results of Kalman Filter modeling performed on the Squid-measured OCXO B dataset. On this interval, the measured frequency exhibits an approximately linear trend, with the Kalman Filter predicted values closely tracking the measurements. The R^2 metric for this measurement-model comparison is 0.99995. In general, this model fit tracks the long-term frequency aging behavior while filtering short-term effects. Figure 39(b) displays the frequency aging behavior of the model estimate against the baseline initial aging estimate. The model fit for initial aging demonstrates lower magnitude aging than the baseline estimate. Figure 39(c) displays the results of Kalman Filter modeling performed on the MMS-measured OCXO B dataset. On this interval, the measured frequency exhibits a quasi-logarithmic trend, with the Kalman Filter predicted values closely tracking the measurements. The R^2 metric for this measurement-model comparison is approximately 1. In general, this model fit tracks the long-term frequency aging behavior while filtering short-term effects. Figure 39(d) displays the frequency aging behavior of the model estimate against the baseline long-term aging estimate. The model fit for long-term aging agrees with the baseline estimate, with some modest variation.

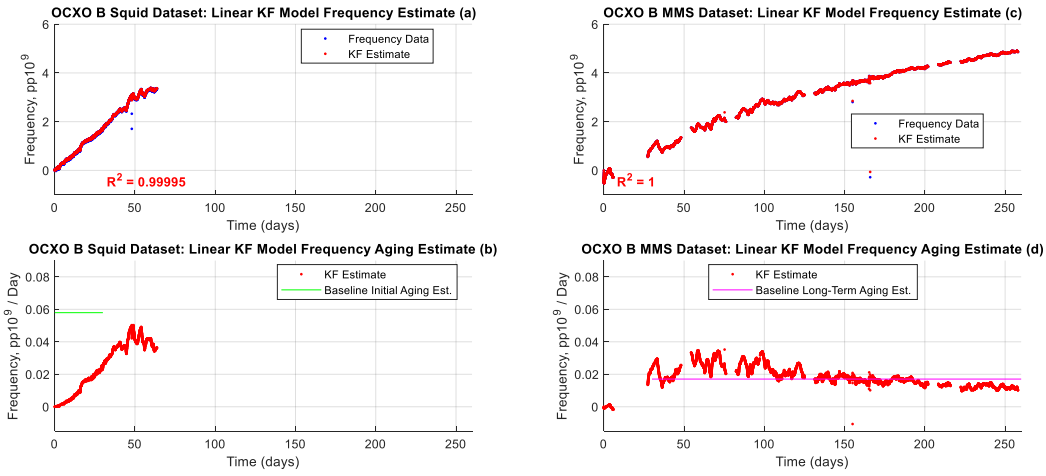

Figure 39: OCXO B linear Kalman Filter modeling results

Figure 40(a) contains the residuals from the Kalman Filter frequency estimate to the Squid-measured dataset. On the measurement interval, there is not a discernable residual pattern. The lack of pattern indicates that the model fully describes the fractional frequency behavior present. Figure 40(b) contains the residuals from the Kalman Filter frequency estimate to the MMS-measured dataset. On the measurement interval, there is not a discernable residual pattern. The lack of pattern indicates that the model fully describes the fractional frequency behavior present.

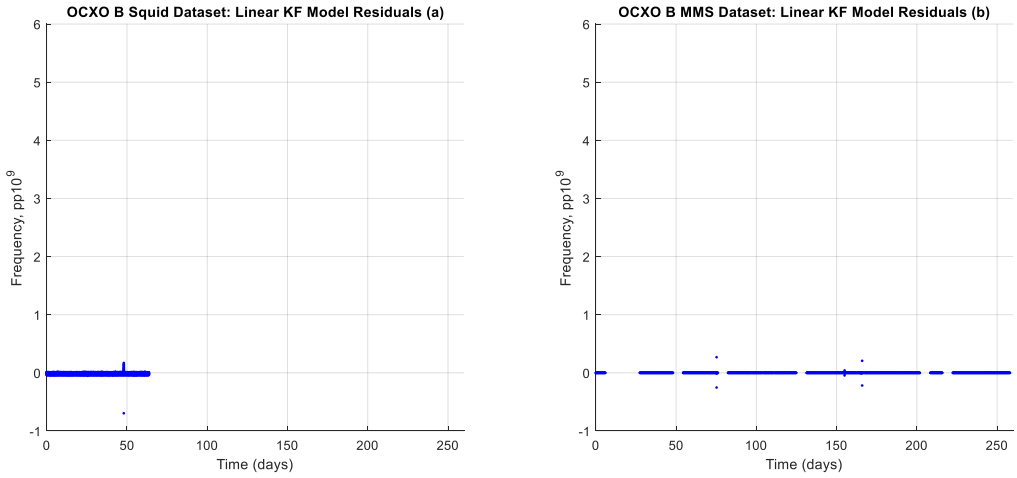

Figure 40: OCXO B linear Kalman Filter modeling residuals

For the initial aging behavior contained in the OCXO B Squid-measured dataset, the Kalman Filter estimate produces a good fit. For the long-term aging in the OCXO B MMS-measured dataset, the Kalman Filter estimate produces a good fit.

4.2.2.3. Linear Kalman Filter Modeling: OCXO H

Figure 41(a) displays the results of Kalman Filter modeling performed on the Squid-measured OCXO H dataset. On this interval, the measured frequency exhibits an anomalous trend, with the Kalman Filter predicted values closely tracking the measurements. The R^2 metric for this measurement-model comparison is 0.99672. In general, this model fit tracks the long-term frequency aging behavior while filtering short-term effects. Figure 41(b) displays the frequency aging behavior of the model estimate against the baseline initial aging estimate. The model fit for initial aging agrees with the baseline estimate, with some modest variation. Figure 41(c) displays the results of Kalman Filter modeling performed on the MMS-measured OCXO H dataset. On this interval, the measured frequency exhibits a quasi-logarithmic trend, with the Kalman Filter predicted values closely tracking the measurements. The R^2 metric for this measurement-model comparison is approximately 1. In general, this model fit tracks the long-term frequency aging behavior while filtering short-term effects. Figure 41(d) displays the frequency aging behavior of the model estimate against the baseline long-term aging estimate. The model fit for long-term aging agrees with the baseline estimate, with some modest variation.

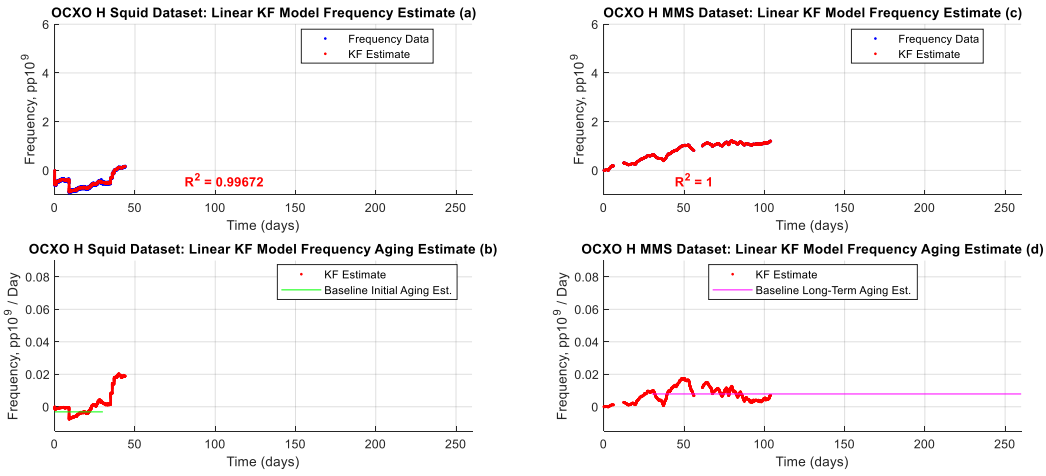

Figure 41: OCXO H linear Kalman Filter modeling results

Figure 42(a) contains the residuals from the Kalman Filter frequency estimate to the Squid-measured dataset. On the measurement interval, there is not a discernable residual pattern. The lack of pattern indicates that the model fully describes the fractional frequency behavior present. Figure 42(b) contains the residuals from the Kalman Filter frequency estimate to the MMS-measured dataset. On the measurement interval, there is not a discernable residual pattern. The lack of pattern indicates that the model fully describes the fractional frequency behavior present.

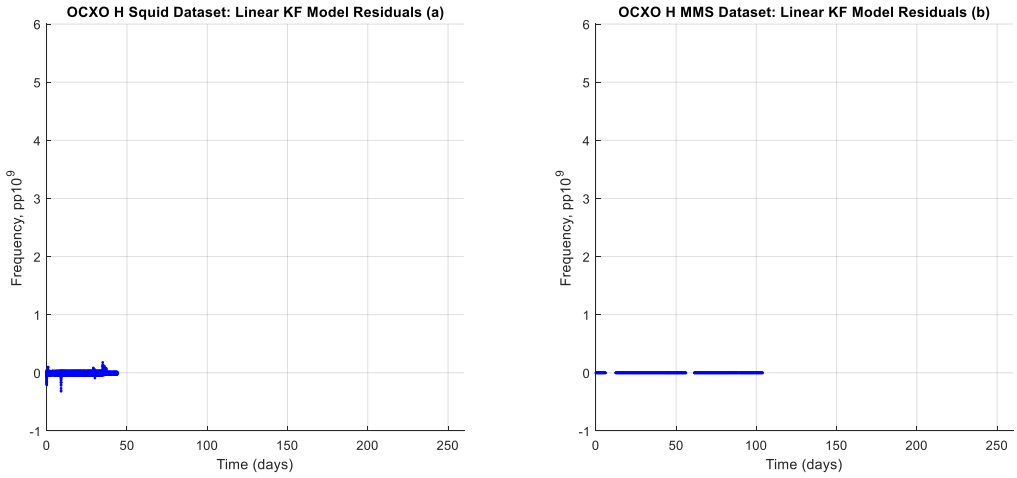

Figure 42: OCXO H linear Kalman Filter modeling residuals

For the initial aging behavior contained in the OCXO H Squid-measured dataset, the Kalman Filter estimate produces a good fit. For the long-term aging in the OCXO H MMS-measured dataset, the Kalman Filter estimate produces a good fit exhibited a good fit.

4.2.3. Linear Kalman Filter Modeling: Coefficient of Determination Summary

For all eight datasets, the R^2 results are tabulated in Table 7. All datasets exhibited R^2 in the 0.98 and above range, with fractional frequency fits exhibiting excellent agreement.

TABLE 7:
LINEAR KALMAN FILTER MODEL RESULTS

OCXO	Squid R^2	Fit Quality	MMS R^2	Fit Quality
A	0.99996	Good	1	Good
B	0.99995	Good	1	Good
C	0.99994	Good	1	Good
D	0.99678	Good	1	Good
E	0.99983	Good	1	Good
F	0.99997	Good	1	Good
G	0.98287	Good	1	Good
H	0.99672	Good	1	Good

4.2.4. Linear Kalman Filter Discussion

To evaluate the linear Kalman Filter model's overall estimate performance, this section focuses on three views of performance. The first view evaluates the degree to which the predicted frequency aging agrees with the baseline linear estimates of aging behavior. The second view centers on the R^2 error metric. The third view considers results of residual analysis between the measured and predicted data.

Visually inspecting the overall agreement of model output based on the Squid-measured datasets to the baseline initial aging estimate, the aging estimates are generally lower in magnitude than the baseline calculation. The long-term behavior of the Kalman Filter frequency aging estimate generally matches the baseline long-term estimates. With these results, it can be concluded that the Kalman Filter frequency aging estimates are representative of the actual state of the frequency aging present in the OCXOs analyzed.

From the view of the R^2 values, the Kalman Filter frequency estimate applied to the OCXO datasets produces fits with very small error, as indicated by values above 0.98. Using R^2 as the only figure of merit for logarithmic model performance, there is very little variation between the performance of the Kalman Filter model between datasets.

The residual analysis of these Kalman Filter frequency estimates indicates great flexibility of the Kalman Filter-based models to describe the frequency aging datasets. The residual plots for all datasets contain primarily random content. This random content is attributed to short-term effects on the OCXOs related to environmental transients, which are effectively removed by the Kalman Filter estimation process.

Combining the linear estimate comparison, R^2 metrics, and residual analysis provides a more complete picture of the linear Kalman Filter model performance. The linear estimate

agreement attests to the responsiveness of the Kalman Filter estimates. The R^2 metrics lend further confirmation of the small modeling error. The random content found via the residual analysis is indicates that the Kalman Filter performs well as far as isolating long-term trends related to frequency aging.

This performance analysis indicates that the linear Kalman Filter model fits the frequency data well. The dynamic system parameters allow the Kalman Filter to adjust estimate behavior to adapt input to changes, resulting in improved input tracking. The selection of frequency aging variance covariance values could be used to further adjust the performance of the Kalman filter, which is an exercise left for future work.

4.3. Summary

As part of this work, a comparison was performed between the model outputs of a logarithmic model and a linear Kalman Filter in terms of their ability to accurately describe the frequency aging of several example OCXO datasets. In general, the logarithmic model produces aging outputs that are highly immune to short-term noise, but also inflexible to changes in long-term aging behavior. The linear Kalman Filter estimate tracks the frequency data tightly while suppressing almost all short-term behavior. The flexibility of the dynamic parameters within the linear Kalman Filter facilitates estimates that agree well with baseline long-term aging estimates, but with some disagreement with the baseline initial aging estimates. The selection of frequency aging covariance controls the degree to which the frequency aging estimate adjusts to aging changes while suppressing noise. The inclusion of error covariance in the linear Kalman Filter model is a primary contributor to its performance improvement over the Logarithmic model.

In this chapter, several methods are utilized to evaluate the performance between these models. These methods include linear aging estimate comparison, R^2 values, and residual analysis. Through the analysis, it was demonstrated that a combination of linear aging estimate comparison and residual analysis are the best indicators of model performance, with R^2 presenting little added benefit.

CHAPTER 5:

CONCLUSION

The relationship of precision timing to critical system performance cannot be overstated. As such, understanding the underlying behaviors of frequency sources is paramount, particularly behavior related to long-term frequency drift. In this work, the contribution of frequency drift due to frequency aging is modeled using two different methods: a logarithmic model and a linear Kalman Filter model. Example datasets derived from a set of Microchip 1000X series oscillators are analyzed using these two modeling techniques.

The logarithmic model is established in the literature as the foundational model of frequency aging. In this thesis, the least squares method is used to fit the logarithmic model to measured fractional frequency data from each of the oscillators under consideration. The logarithmic modeling process involves adjusting the three available model parameters utilizing the least squares process. The resulting model can be used to estimate frequency aging or to extrapolate for missing data or predict future trends not included in the measured data. The static parameters used in this model provide robustness against the influence of short-term noise, but this comes at the cost of flexibility in describing changes in long-term frequency aging behaviors.

The linear Kalman Filter model assumes a direct linear relationship between phase, fractional frequency, and frequency aging. The dynamic nature of the parameters used in this model provides flexibility in describing changes in long-term frequency aging behaviors. The

degree to which the Kalman Filter tracks the fractional frequency input and accurately estimates frequency aging is controlled by the covariance assumptions of the model.

The method of obtaining numerical estimates of these covariance values is presented. This study finds that the behavior of the Kalman Filter frequency aging estimate may be adjusted with careful selection of the frequency aging state covariance parameters.

To aid in analysis of frequency aging model fidelity, it is suggested to include residual analysis and linear frequency aging estimate comparisons. One of the findings of this thesis is that a purely metric-driven approach, such as using the coefficient of determination, R^2, does not consistently discriminate between good model accuracy and poor model accuracy. However, R^2 can be used directly to detect anomalous oscillator behavior, due to the particularly low R^2 value (less than 0.5) associated with this phenomenon. Residual analysis helps to identify the degree to which the model describes present frequency aging behaviors. Baseline linear frequency aging estimate comparisons are also useful. These comparisons indicate the frequency aging fidelity of the model output.

Using these performance analysis tools, the linear Kalman Filter model was shown to be more useful than the logarithmic model for obtaining frequency aging values that are representative of realistic OCXO behavior. In the fractional frequency domain, the logarithmic model provides relatively poor agreement with the measured behavior of the OCXOs under consideration. In the frequency aging domain, the logarithmic model also does not accurately predict the long-term aging behavior of these oscillators. On the other hand, the linear Kalman Filter model is shown to accurately estimate both the fractional frequency values and the long-term aging trends present in such measurements. However, the linear Kalman filter has an additional benefit over traditional static models in terms of tunability. The covariance matrices

provide a powerful capability to adjust the properties of the linear Kalman filter to optimize its accuracy for frequency aging output.

This consideration of noise within the Kalman Filter model proves invaluable in describing the full behavior present in OCXO data. The flexibility to adjust the characteristics of this model provide a distinct advantage over traditional static approaches such as the logarithmic model for the purpose of describing OCXO behavior. For the purview of the wider precise time community, these results emphasize the need for error consideration in frequency standard modeling. This study focuses on a limited subset of oscillators from a single manufacturer. Future work can incorporate different types of oscillators and other frequency standards, as well as different model types.

REFERENCES

[1] D. W. Allan, N. Ashby, and C. C. Hodge, “The Science of Timekeeping.” Hewlett-Packard Company, 1997.

[2] “Base Units,” *BIPM*. [Online]. Available: https://www.bipm.org/en/measurement-units/base-units.html.

[3] “CESIUM ATOMIC CLOCKS,” *USNO*. [Online]. Available: https://tycho.usno.navy.mil/cesium.html. [Accessed: 17-Aug-2019].

[4] J. Levine and T. Parker, “The algorithm used to realize UTC(NIST),” in *Proceedings of the 2002 IEEE International Frequency Control Symposium and PDA Exhibition (Cat. No.02CH37234)*, 2002, pp. 537–542.

[5] Ł. Sliwczynski and J. Kolodziej, “Bidirectional Optical Amplification in Long-Distance Two-Way Fiber-Optic Time and Frequency Transfer Systems,” *IEEE Trans. Instrum. Meas.*, vol. 62, no. 1, pp. 253–262, Jan. 2013.

[6] B. Parzen and A. Ballato, *Design of Crystal and Other Harmonic Oscillators*. New York: Wiley, 1983.

[7] H. W. Hellwig, “Atomic frequency standards: A survey,” *Proc. IEEE*, vol. 63, no. 2, pp. 212–229, 1975.

[8] W. A. Marrison, “The Evolution of the Quartz Crystal Clock,” *Bell Syst. Tech. J.*, vol. 27, no. 3, pp. 510–588, Jul. 1948.

[9] J. Vig, “Quartz Crystal Resonators and Oscillators For Frequency Control and Timing Applications - A Tutorial.” 2008.

[10] A. O. McCoubrey, “A survey of atomic frequency standards,” *Proc. IEEE*, vol. 54, no. 2, pp. 116–135, 1966.

[11] “Microchip Gives Its Most Precise Atomic Clock a Performance Upgrade.” Microchip, 2019.

[12] M. A. Lombardi, A. N. Novick, G. Neville-Neil, and B. Cooke, “Accurate, Traceable, and Verifiable Time Synchronization for World Financial Markets,” *J. Res. Natl. Inst. Stand. Technol.*, vol. 121, p. 436, Oct. 2016.

[13] I. Skog and P. Handel, “Effects of time synchronization errors in GNSS-aided INS,” in *2008 IEEE/ION Position, Location and Navigation Symposium*, 2008, pp. 82–88.

[14] IEEE Standards Coordinating Committee 27 on Time and Frequency., American National Standards Institute., Institute of Electrical and Electronics Engineers., and IEEE-SA Standards Board., "1193-2003 - IEEE guide for measurement of environmental sensitivities of standard frequency generators." Institute of Electrical and Electronics Engineers, p. 74, 2004.

[15] IEEE-SA Standards Board, "IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology - Random Instabilities." IEEE, New York, 2009.

[16] J. R. Vig and T. R. Meeker, "The aging of bulk acoustic wave resonators, filters and oscillators," in *Proceedings of the 45th Annual Symposium on Frequency Control 1991*, 1991, pp. 77–101.

[17] W. J. Riley, *Handbook of Frequency Stability Analysis*, vol. 31, no. 1. 1994.

[18] J. Vanier, J. J. Gagnepain, W. J. Riley, F. L. Walls, and M. Granveaud, "Aging, warm-up time and retrace: important characteristics of standard frequency generators (Proposal for IEEE Standards Project P1193)," in *Proceedings of the 1992 IEEE Frequency Control Symposium*, pp. 807–815.

[19] D. C. Montgomery and G. C. Runger, *Applied Statistics and Probability for Engineers*, Third Edit. New York: Wiley, 2003.

[20] L. Fahrmeir, T. Kneib, S. Lang, and B. Marx, *Regression*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013.

[21] R. L. Filler, "Aging Specification, Measurement, and Analysis," in *7th Quartz Devices Conference*, 1985, pp. 93–104.

[22] M. R. Miljkovic, G. L. Trifunovic, and V. J. Brajovic, "Aging prediction of quartz crystal units," in *Proceedings of the 42nd Annual Frequency Control Symposium, 1988.*, pp. 404–411.

[23] R. L. Filler and J. R. Vig, "Long-term aging of oscillators," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 40, no. 4, pp. 387–394, Jul. 1993.

[24] O. Leibfried and B. Neubig, "Correlation of predicted and real aging behaviour of crystal oscillators using different fitting algorithms," in *11th European Frequency and Time Forum*, 1997, pp. 268–272.

[25] E. Seydel, "Relation between physical processes and aging," in *2009 IEEE International Frequency Control Symposium Joint with the 22nd European Frequency and Time forum*, 2009, pp. 927–930.

[26] M. S. Grewal and A. P. Andrews, *Kalman filtering: theory and practice using MATLAB*, Fourth. Hoboken, New Jersey: Wiley, 2015.

[27] P. Kim, *Kalman Filter for Beginners with MATLAB Examples*. A-JIN, 2010.

[28] S. R. Stein and R. L. Filler, “Kalman filter analysis for real time applications of clocks and oscillators,” in *Proceedings of the 42nd Annual Frequency Control Symposium, 1988.*, 1988, pp. 447–452.

[29] W. Su and R. L. Filler, “Application of Kalman filtering techniques to the precision clock with non-constant aging,” in *Proceedings of the 1992 IEEE Frequency Control Symposium*, 1992, pp. 231–237.

[30] “1000C | Microsemi.” [Online]. Available: https://www.microsemi.com/product-directory/embedded-clocks-frequency-references/3963-1000c. [Accessed: 22-Apr-2019].

[31] L. Cosart, B. Hamilton, and G. Zampetti, “Squid Measurement System.” Microchip.

[32] Microsemi, “Multi-Channel Measurement System (MMS).” .

[33] J. A. Barnes *et al.*, “Characterization of Frequency Stability,” *IEEE Trans. Instrum. Meas.*, vol. IM–20, no. 2, pp. 105–120, May 1971.

APPENDIX

This appendix contains the OCXO datasets which were omitted in Section 3.7 and the full modeling results for the logarithmic and linear Kalman Filter models.

APPENDIX A. DATASETS

Figure 43(a) presents a plot of the frequency behavior for OCXO C, which was measured over a period of 55 days utilizing the Squid system. A linear trend has been fit to the first 30 days of this dataset, which estimates an aging rate over this period of $8.312\ pp10^{11}$. The length of this dataset is smaller than most of the other Squid-measured datasets. This data does not contain any irregularities that suggest anomalous operation during the measurement window. In general, the Squid-measured frequency data for OCXO C appears to be well represented by a linear trend, which is consistent with the expectation of initial aging for a crystal oscillator. The logarithmic decrease in aging rate at the end of the dataset is typical of the transition from initial aging into long-term aging.

Figure 43(b) presents a plot of the frequency behavior for OCXO C, which was measured over a period of 69.5 days utilizing the MMS system. A linear trend has been fit to the last 39.5 days of this dataset, which estimates an aging rate over this period of $3.032\ pp10^{11}$. The length of this dataset is smaller than most of the other MMS-measured datasets. This dataset does not contain any irregularities that suggest anomalous operation during the measurement window. In general, the MMS-measured frequency data for OCXO C appears to be well represented by a linear trend, which is consistent with the expectation of long-term aging for a crystal oscillator.

Comparing the estimated aging rates for OCXO C detailed in Figure 43 corroborates the expected aging trends for crystal oscillators since the aging rate estimated from the Squid-measured data (short-term aging) is considerably higher than the aging rate estimated from the MMS-measured data (long-term aging). The OCXO C dataset clearly displays typical frequency aging trends. Fitting each model to this dataset will provide an indication of each model's ability to describe standard frequency aging.

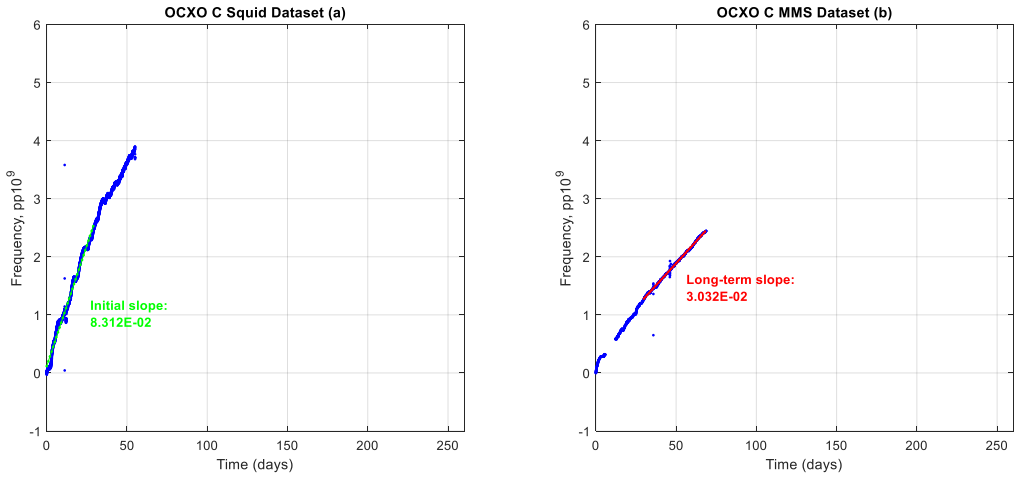

Figure 43: OCXO C datasets

Figure 44(a) presents a plot of the frequency behavior for OCXO D, which was measured over a period of 253 days utilizing the Squid system. A linear trend has been fit to the first 30 days of this dataset, which estimates an aging rate over this period of $2.234\ pp10^{11}$. The length of this dataset is larger than all other Squid-measured datasets. This dataset contains anomalous behavior that manifests as a logarithmic trend during the first 30 days of data, as well as high-frequency oscillations in the last 25 days of data. The initial logarithmic behavior is believed to be attributed to OCXO D having completed initial aging, such that the Squid measurements actually capture long-term aging behavior. The oscillatory measurements at the end of the dataset are believed to be attributed to an internal oscillator error. These anomalous behaviors potentially have minor impact on frequency aging analysis, since the long-term aging trends are unobscured. In general, the Squid-measured frequency data for OCXO D appears to be well represented by a logarithmic trend, which is consistent with the expectation of the transition from short-term to long-term aging for a crystal oscillator.

Figure 44(b) presents a plot of the frequency behavior for OCXO D, which was measured over a period of 112.6 days utilizing the MMS system. A linear trend has been fit to the last 82.6 days of this dataset, which estimates an aging rate over this period of $1.646\ pp10^{11}$. The length of this dataset is average amongst the other MMS-measured datasets. This dataset does not contain any irregularities that suggest anomalous operation during the measurement window. In general, the MMS-measured frequency data for OCXO D appears to be well represented by a linear trend, which is consistent with the expectation of long-term aging for a crystal oscillator.

Comparing the estimated aging rates for OCXO D detailed in Figure 44 corroborates the expected aging trends for crystal oscillators since the aging rate estimated from the Squid-measured data (short-term to long-term aging) is considerably higher than the aging rate

estimated from the MMS-measured data (long-term aging). The OCXO D dataset clearly displays typical frequency aging trends. Fitting each model to this dataset will provide an indication of each model's ability to describe long-term frequency aging.

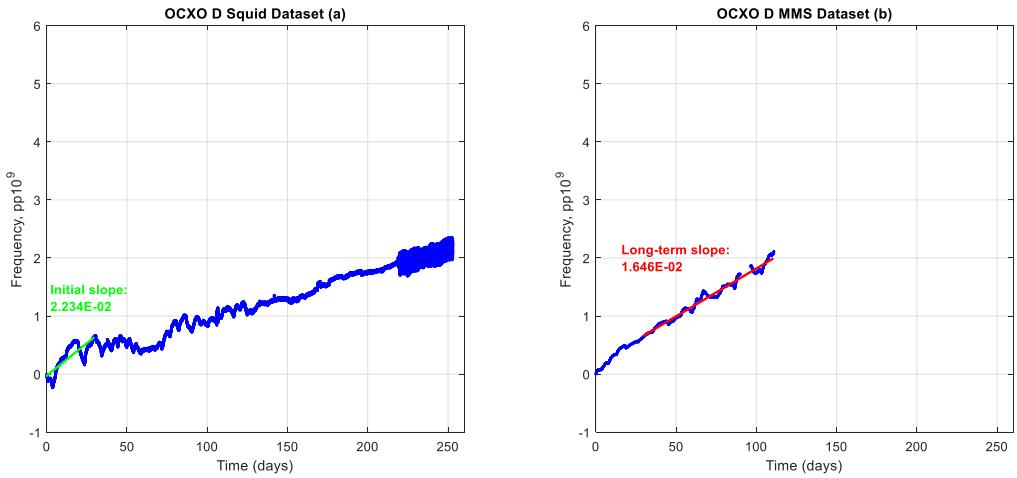

Figure 44: OCXO D datasets

Figure 45(a) presents a plot of the frequency behavior for OCXO E, which was measured over a period of 33.9 days utilizing the Squid system. A linear trend has been fit to the first 30 days of this dataset, which estimates an aging rate over this period of $5.568\ pp10^{11}$. The length of this dataset is the smallest of the Squid-measured datasets. This dataset does not contain any irregularities that suggest anomalous operation during the measurement window. In general, the Squid-measured frequency data for OCXO E appears to be well represented by a linear trend, which is consistent with the expectation of initial aging for a crystal oscillator.

Figure 45(b) presents a plot of the frequency behavior for OCXO E, which was measured over a period of 34.6 days utilizing the MMS system. A linear trend has been fit to the last 4.6 days of this dataset, which estimates an aging rate over this period of $4.948\ pp10^{11}$. The length of this dataset is the smallest of the MMS-measured datasets. Due to the short length of the fitted dataset, the estimated frequency aging may be heavily influenced by environmental transients. Otherwise, this dataset does not contain any irregularities that suggest anomalous operation during the measurement window. In general, the MMS-measured frequency data for OCXO E appears to be well represented by a linear trend, which is consistent with the expectation of long-term aging for a crystal oscillator.

Comparing the estimated aging rates for OCXO E detailed in Figure 45 corroborates the expected aging trends for crystal oscillators since the aging rate estimated from the Squid-measured data (short-term aging) is considerably higher than the aging rate estimated from the MMS-measured data (long-term aging). The OCXO E dataset clearly displays typical frequency aging trends. Fitting each model to this dataset will provide an indication of each model's ability to describe standard frequency aging over short datasets.

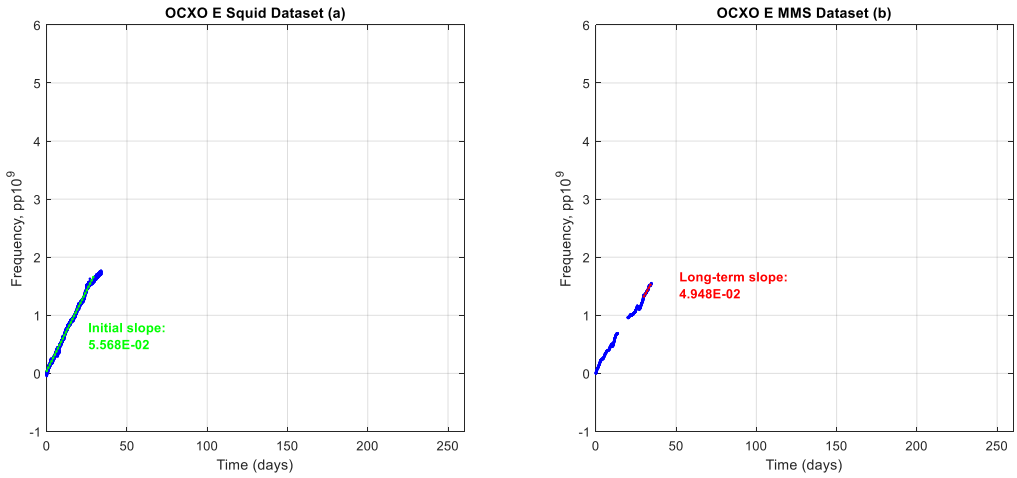

Figure 45: OCXO E datasets

Figure 46(a) presents a plot of the frequency behavior for OCXO F, which was measured over a period of 36.1 days utilizing the Squid system. A linear trend has been fit to the first 30 days of this dataset, which estimates an aging rate over this period of $3.905\ pp10^{11}$. The length of this dataset is smaller than most of the other Squid-measured datasets. This dataset does not contain any irregularities that suggest anomalous operation during the measurement window. In general, the Squid-measured frequency data for OCXO F appears to be well represented by a linear trend, which is consistent with the expectation of short-term aging for a crystal oscillator.

Figure 46(b) presents a plot of the frequency behavior for OCXO F, which was measured over a period of 35.7 days utilizing the MMS system. A linear trend has been fit to the last 5.7 days of this dataset, which estimates an aging rate over this period of $1.847\ pp10^{11}$. The length of this dataset is smaller than most of the other MMS-measured datasets. Due to the short length of the fitted dataset, the estimated frequency aging may be heavily influenced by environmental transients. Otherwise, this dataset does not contain any irregularities that suggest anomalous operation during the measurement window. In general, the MMS-measured frequency data for OCXO F appears to be well represented by a linear trend, which is consistent with the expectation of long-term aging for a crystal oscillator.

Comparing the estimated aging rates for OCXO F detailed in Figure 46 corroborates the expected aging trends for crystal oscillators since the aging rate estimated from the Squid-measured data (short-term aging) is considerably higher than the aging rate estimated from the MMS-measured data (long-term aging). The OCXO F dataset clearly displays typical frequency aging trends. Fitting each model to this dataset will provide an indication of each model's ability to describe standard frequency aging on short datasets.

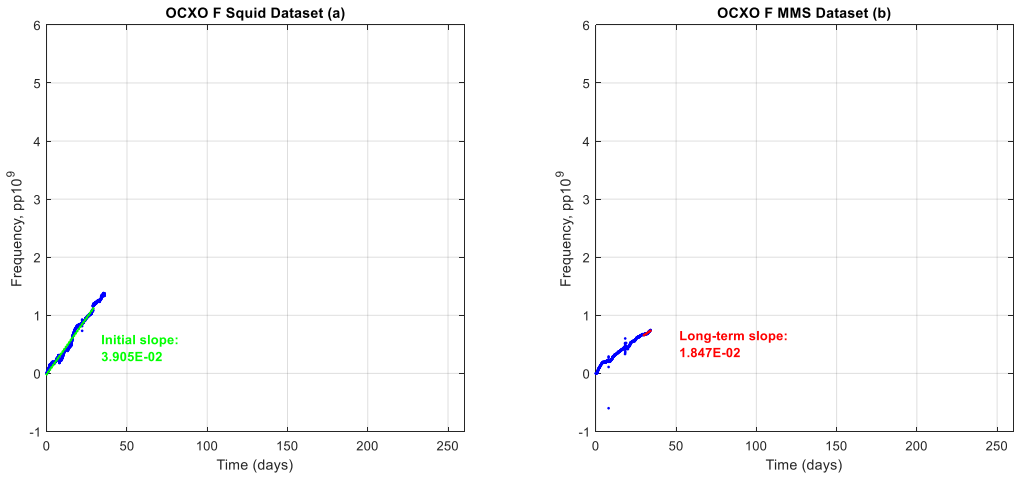

Figure 46: OCXO F datasets

Figure 47(a) presents a plot of the frequency behavior for OCXO G, which was measured over a period of 62.8 days utilizing the Squid system. A linear trend has been fit to the first 30 days of this dataset, which estimates an aging rate over this period of $2.844\ pp10^{11}$. The length of this dataset is average amongst the other Squid-measured datasets. This dataset contains anomalous behavior that manifests as high-frequency oscillations during the first 25 days of data. This behavior is believed to be attributed to an internal oscillator malfunction. The anomalous behavior potentially has minor impact on frequency aging analysis, since the long-term aging trend remains unobscured. In general, the Squid-measured frequency data for OCXO G appears to be well represented by a linear trend, which is consistent with the expectation of short-term aging for a crystal oscillator.

Figure 47(b) presents a plot of the frequency behavior for OCXO G, which was measured over a period of 132.6 days utilizing the MMS system. A linear trend has been fit to the last 102.6 days of this dataset, which estimates an aging rate over this period of $2.090\ pp10^{11}$. The length of this dataset is average amongst the other MMS-measured datasets. This dataset does not contain any irregularities that suggest anomalous operation during the measurement window. In general, the MMS-measured frequency data for OCXO G appears to be well represented by a logarithmic trend, which is consistent with the expectation of initial to long-term aging transition for a crystal oscillator.

Comparing the estimated aging rates for OCXO G detailed in Figure 47 corroborates the expected aging trends for crystal oscillators since the aging rate estimated from the Squid-measured data (short-term aging) is higher than the aging rate estimated from the MMS-measured data (long-term aging). The OCXO G dataset clearly displays typical frequency aging

trends. Fitting each model to this dataset will provide an indication of each model's ability to describe initial to long-term frequency aging.

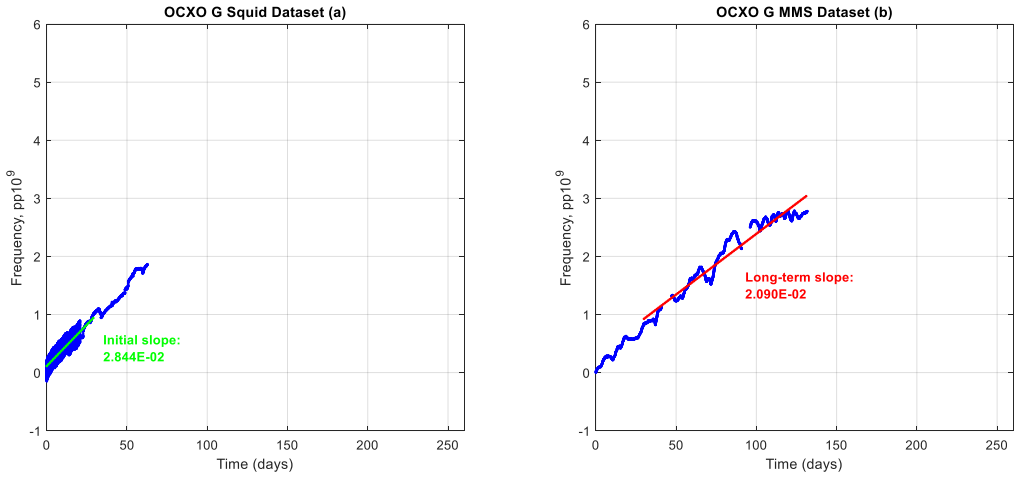

Figure 47: OCXO G datasets

APPENDIX B. LOGARITHMIC MODEL RESULTS

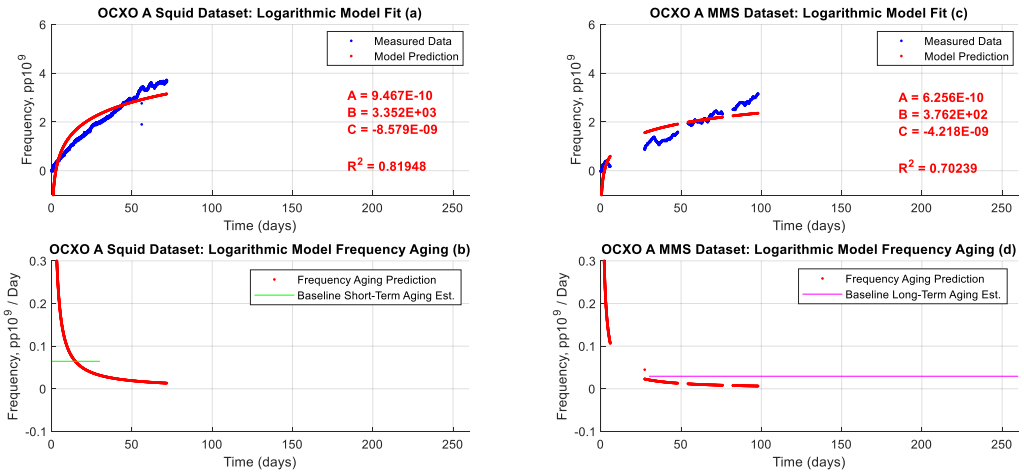

Figure 48: OCXO A logarithmic modeling results

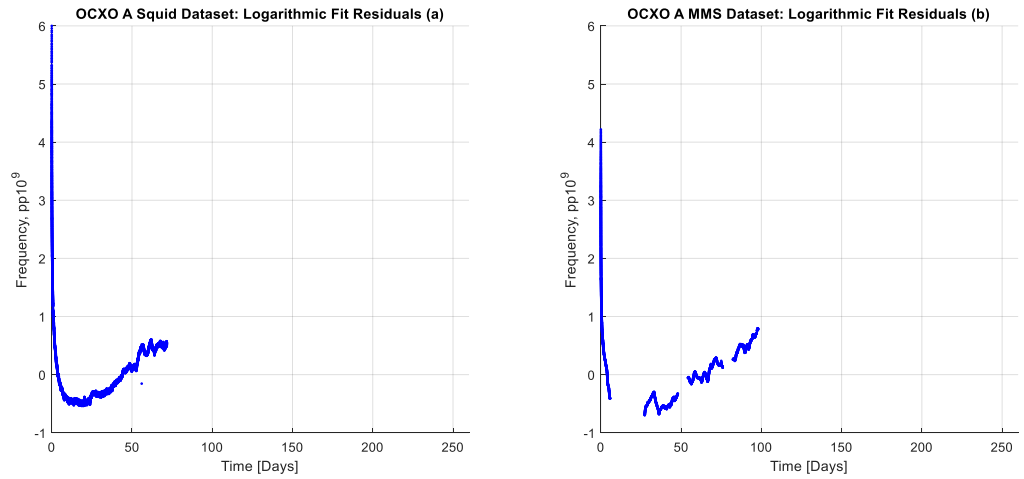

Figure 49: OCXO A logarithmic modeling residuals

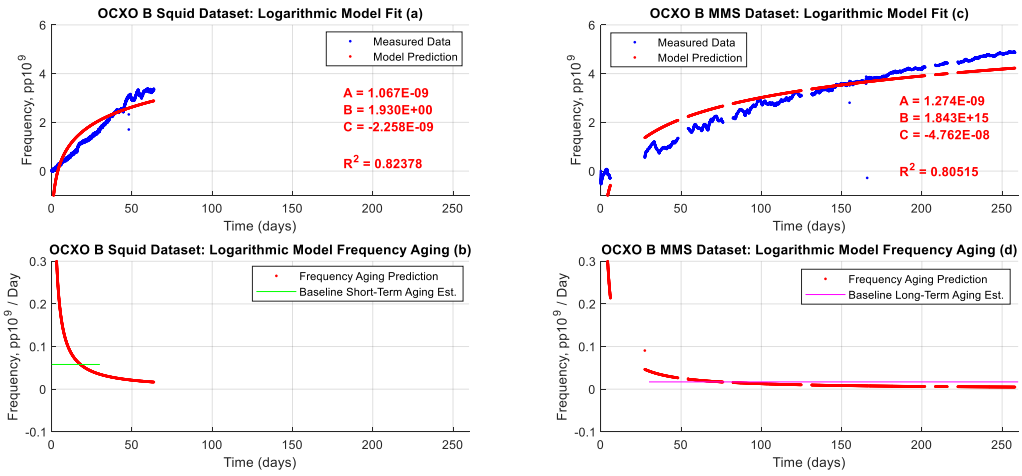

Figure 50: OCXO B logarithmic modeling results

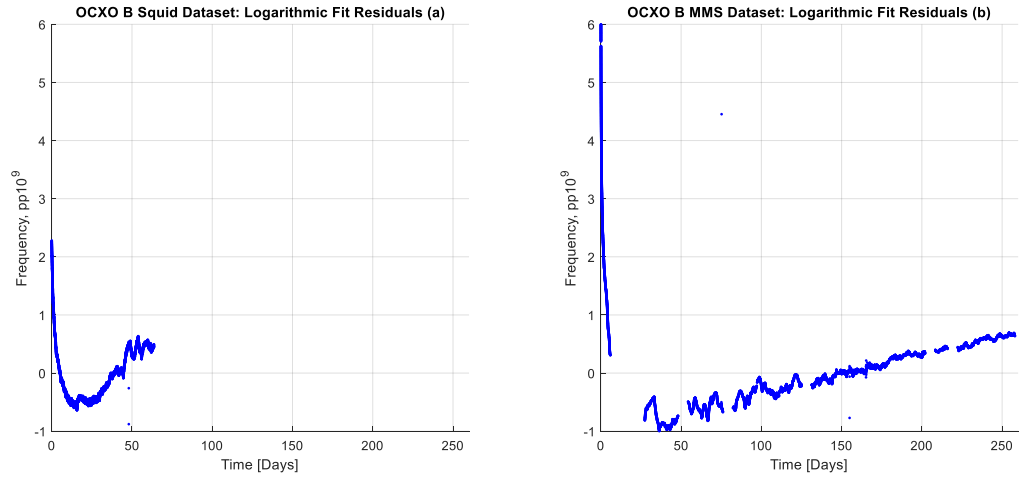

Figure 51: OCXO B logarithmic modeling residuals

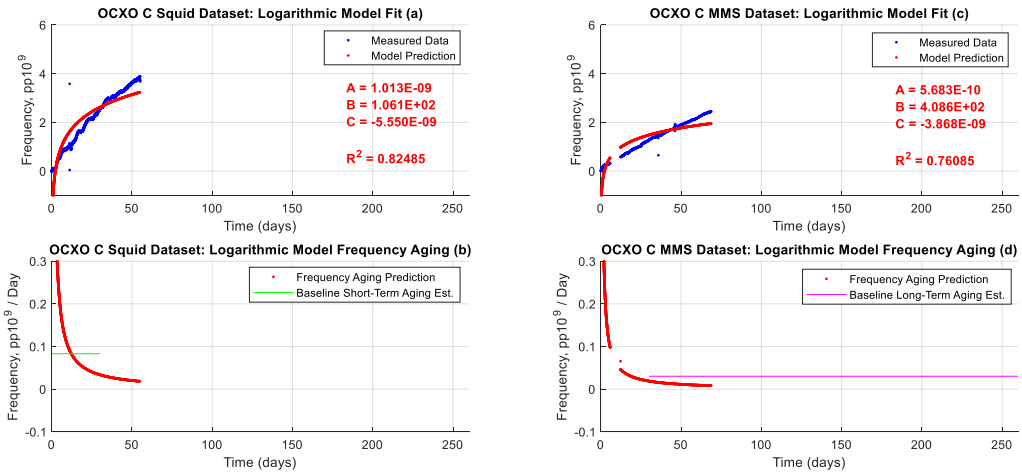

Figure 52: OCXO C logarithmic modeling results

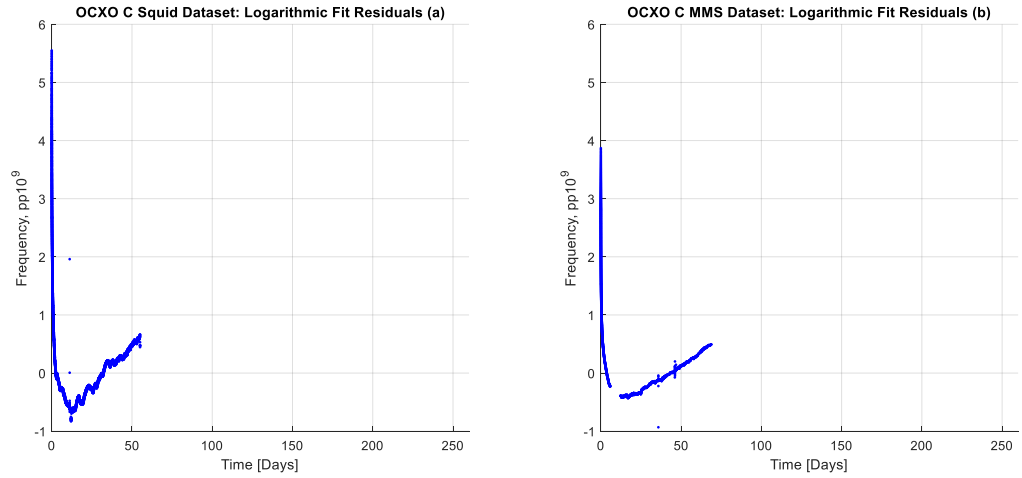

Figure 53: OCXO C logarithmic modeling residuals

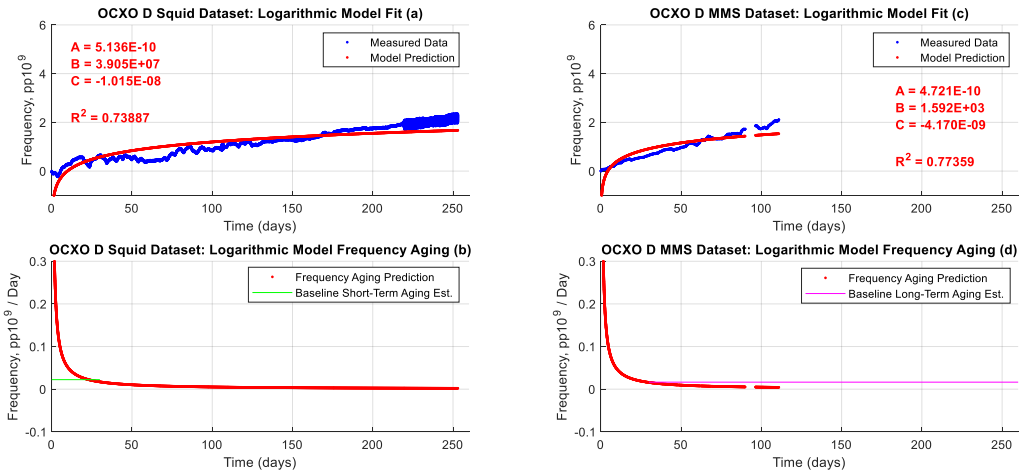

Figure 54: OCXO D logarithmic modeling results

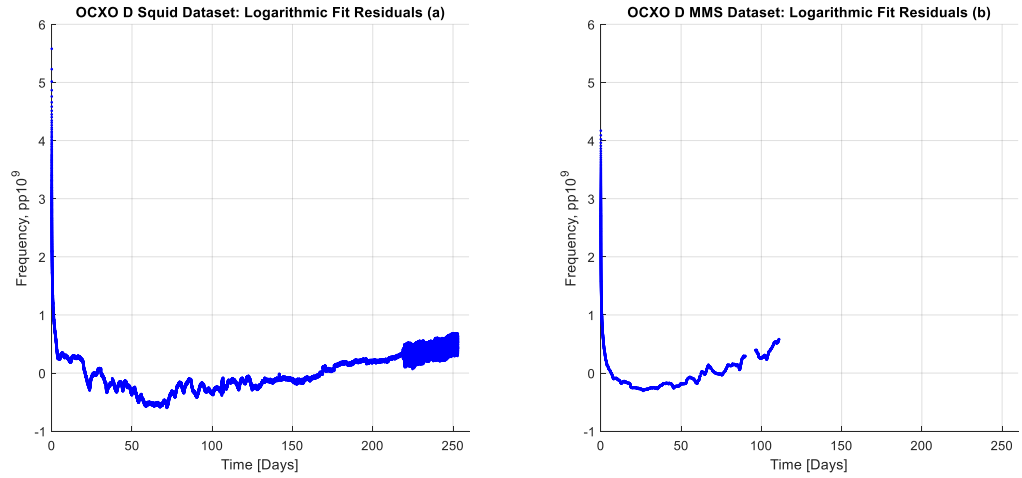

Figure 55: OCXO D logarithmic modeling residuals

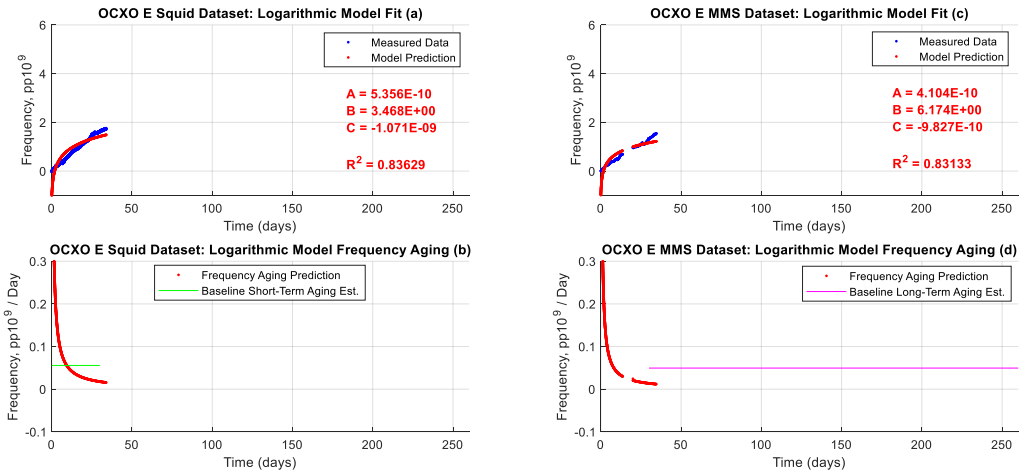

Figure 56: OCXO E logarithmic modeling results

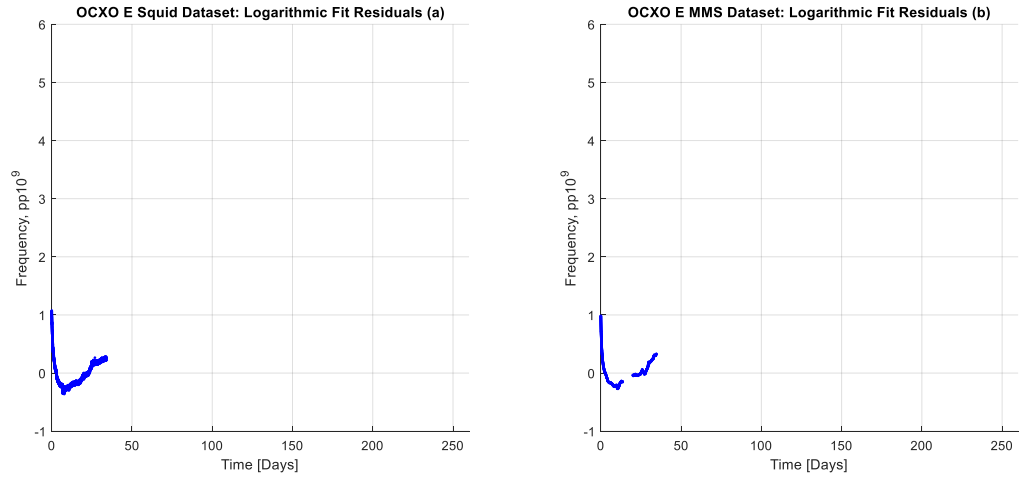

Figure 57: OCXO E logarithmic modeling residuals

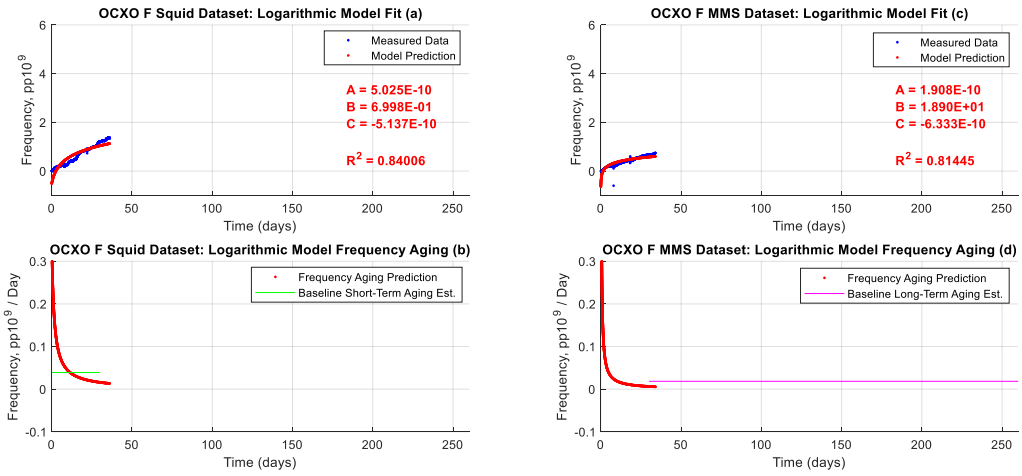

Figure 58: OCXO F logarithmic modeling results

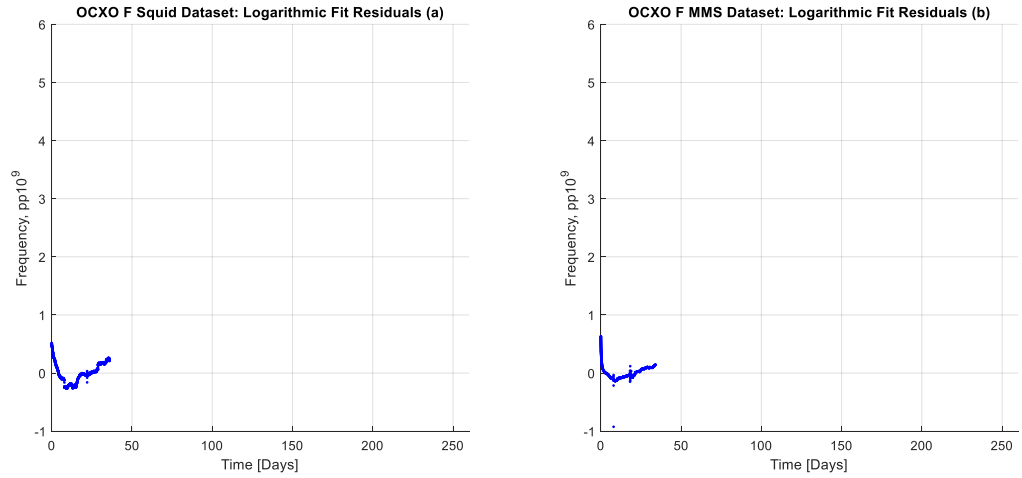

Figure 59: OCXO F logarithmic modeling residuals

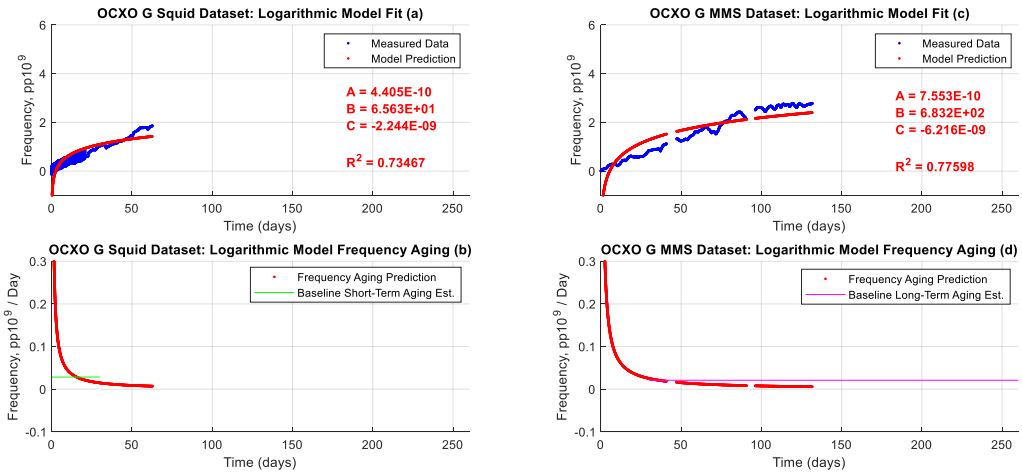

Figure 60: OCXO G logarithmic modeling results

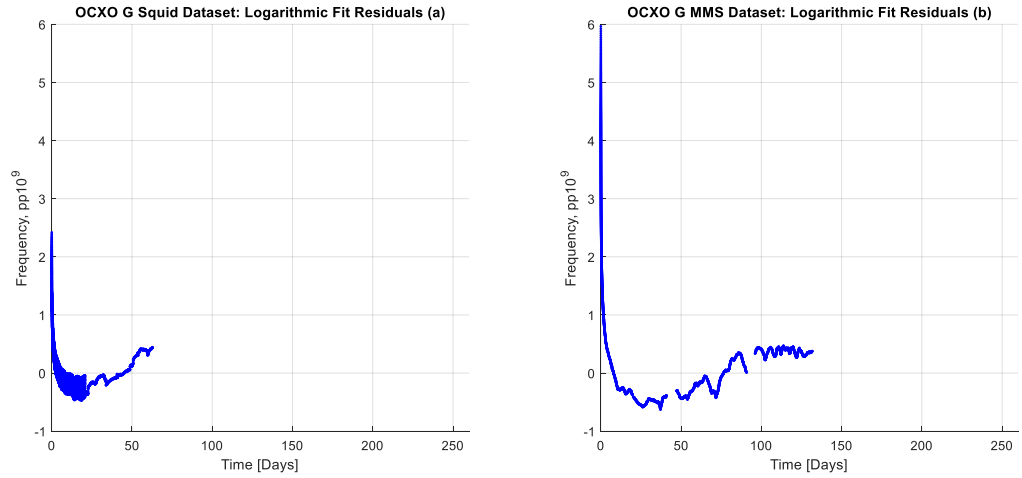

Figure 61: OCXO G logarithmic modeling residuals

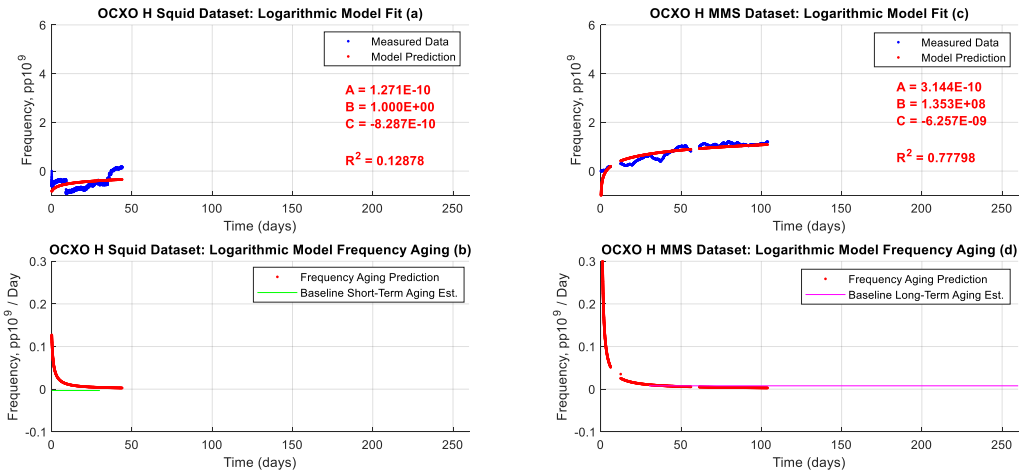

Figure 62: OCXO H logarithmic modeling results

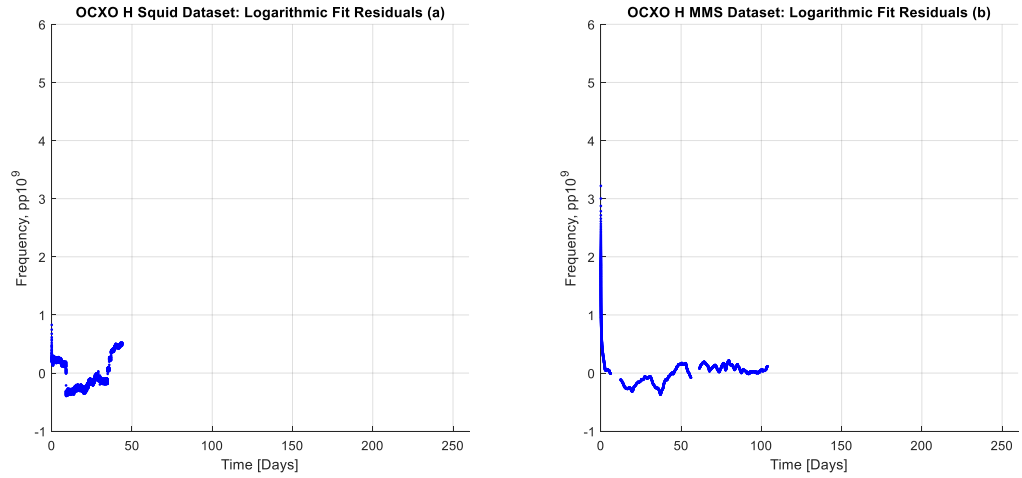

Figure 63: OCXO H logarithmic modeling residuals

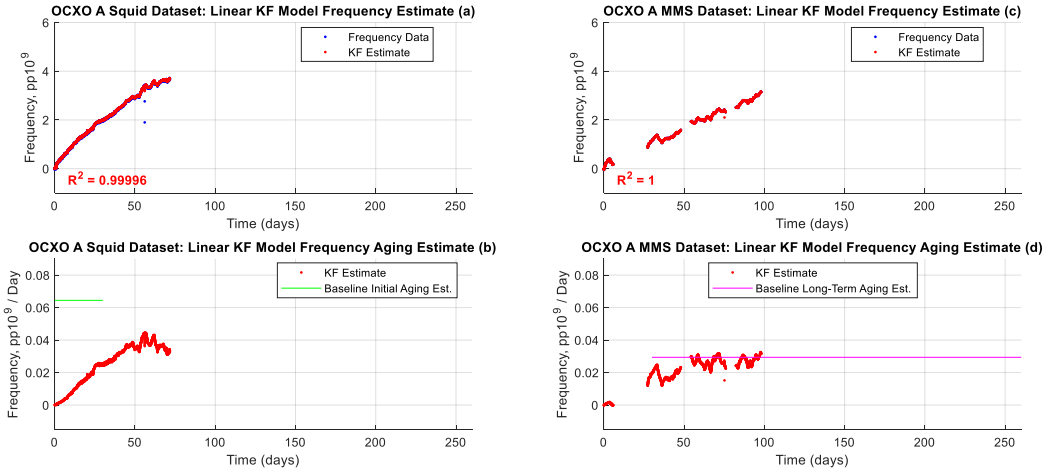

Figure 64: OCXO A linear Kalman Filter modeling results

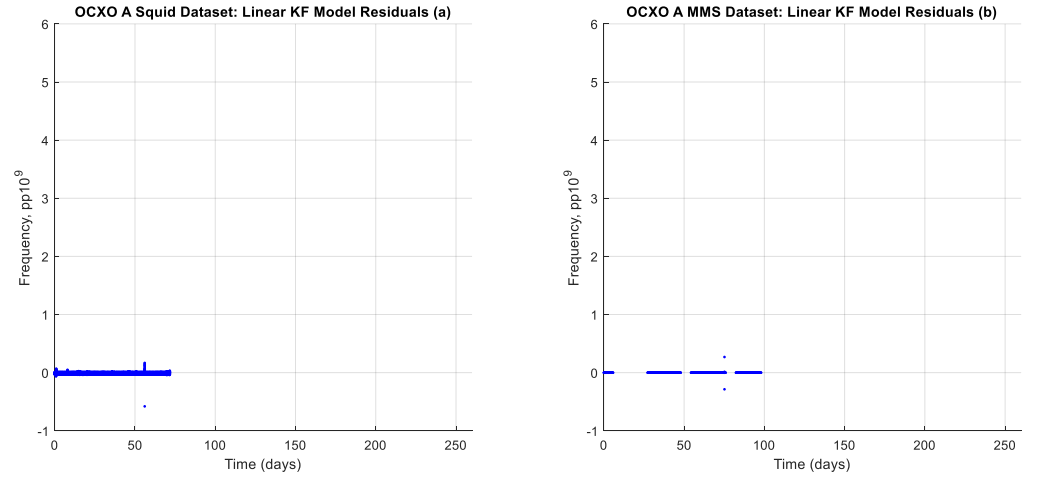

Figure 65: OCXO A linear Kalman Filter residuals

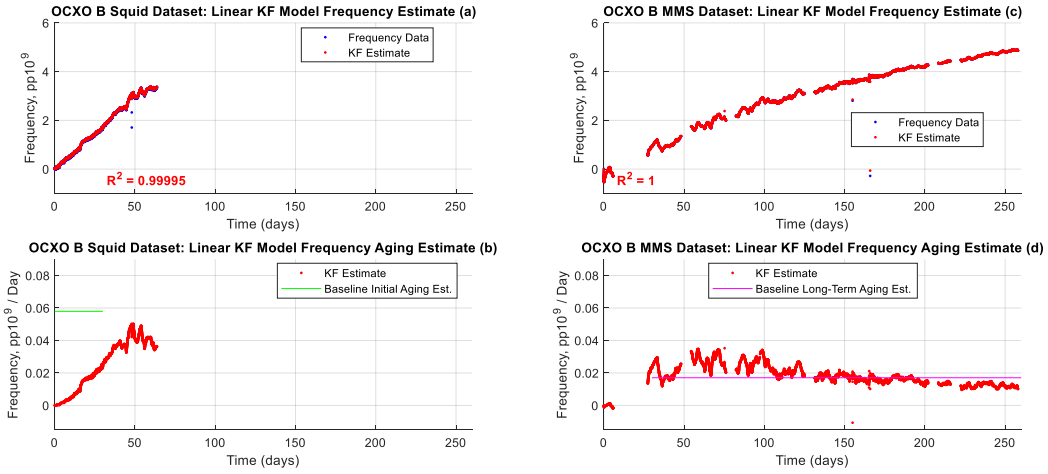

Figure 66: OCXO B linear Kalman Filter modeling results

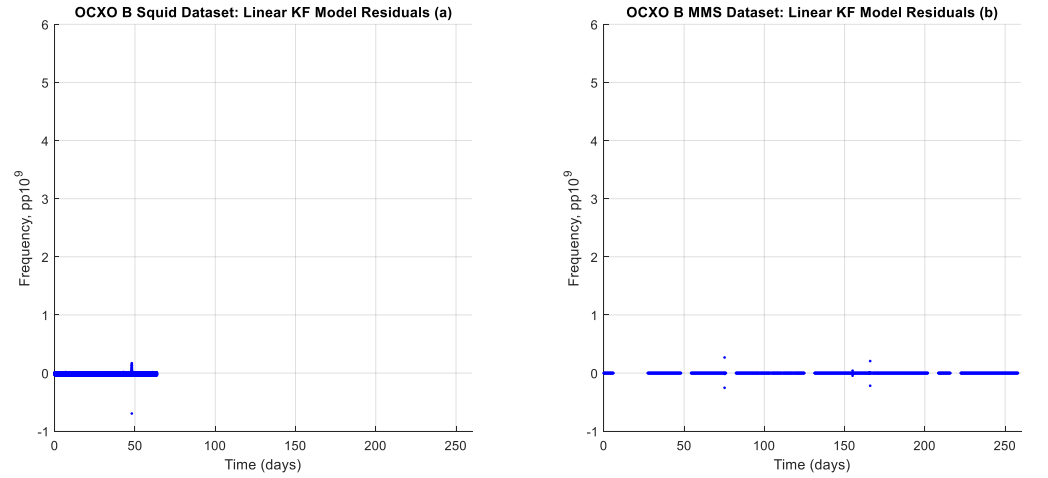

Figure 67: OCXO B linear Kalman Filter residuals

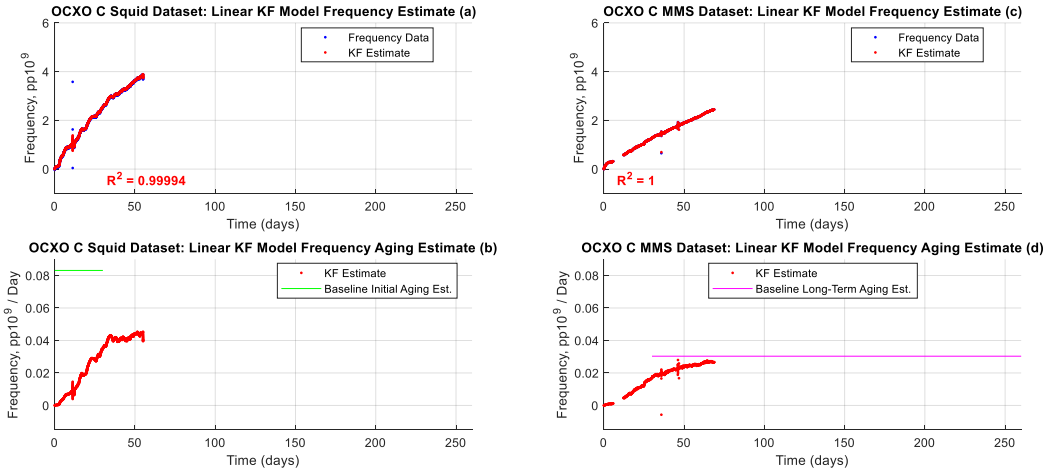

Figure 68: OCXO C linear Kalman Filter modeling results

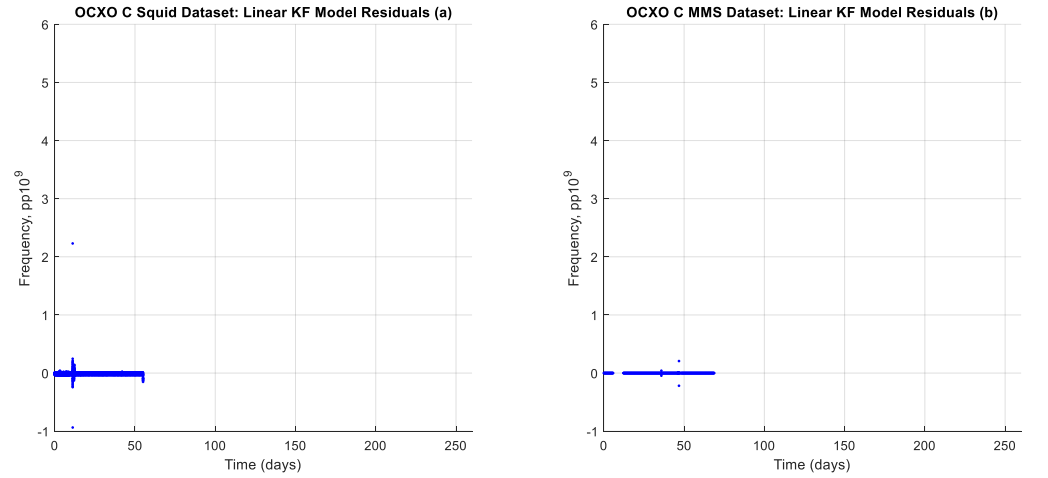

Figure 69: OCXO C linear Kalman Filter residuals

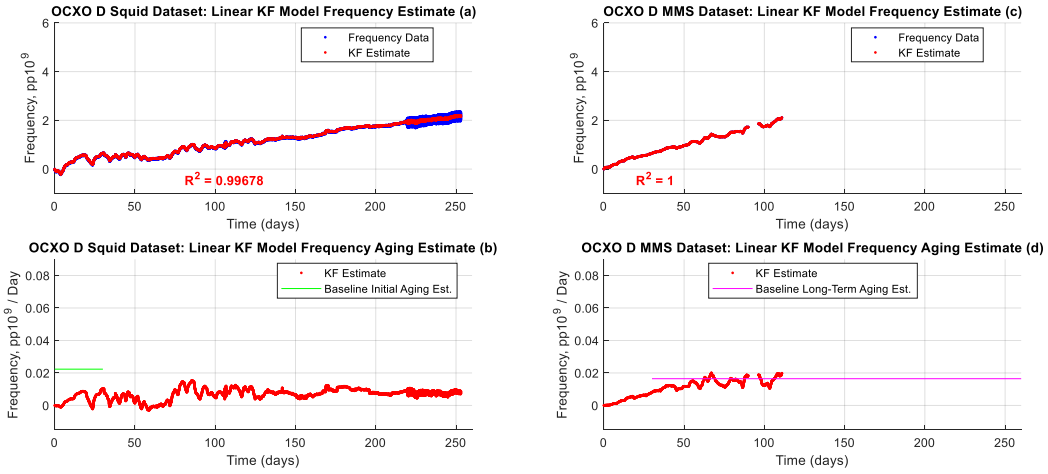

Figure 70: OCXO D linear Kalman Filter modeling results

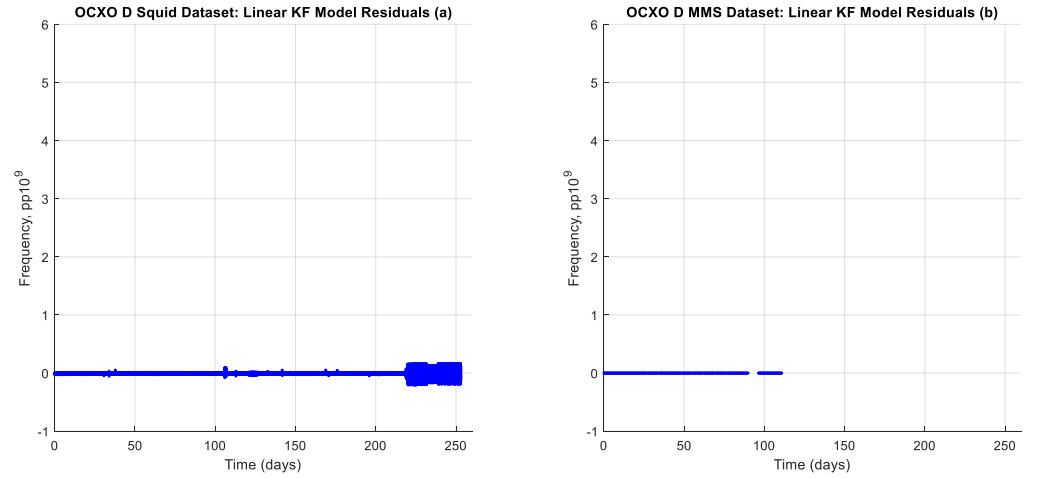

Figure 71: OCXO D linear Kalman Filter residuals

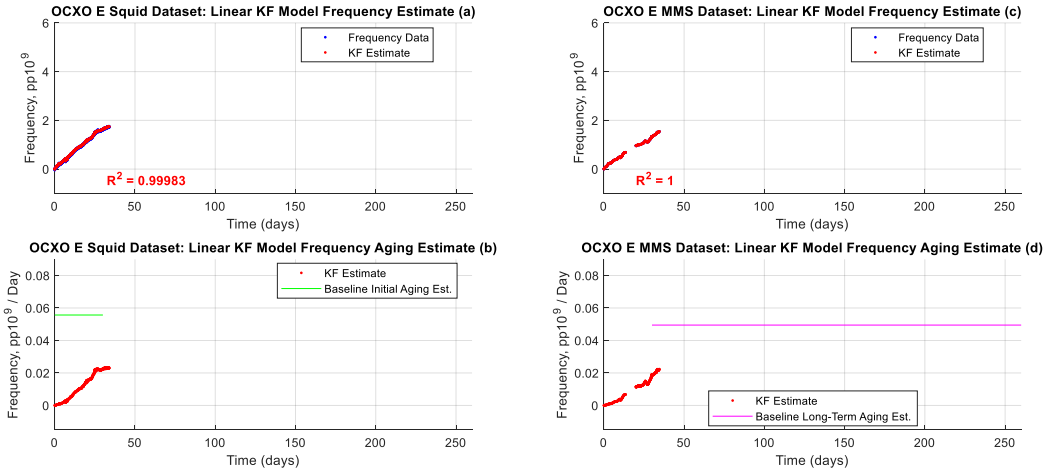

Figure 72: OCXO E linear Kalman Filter modeling results

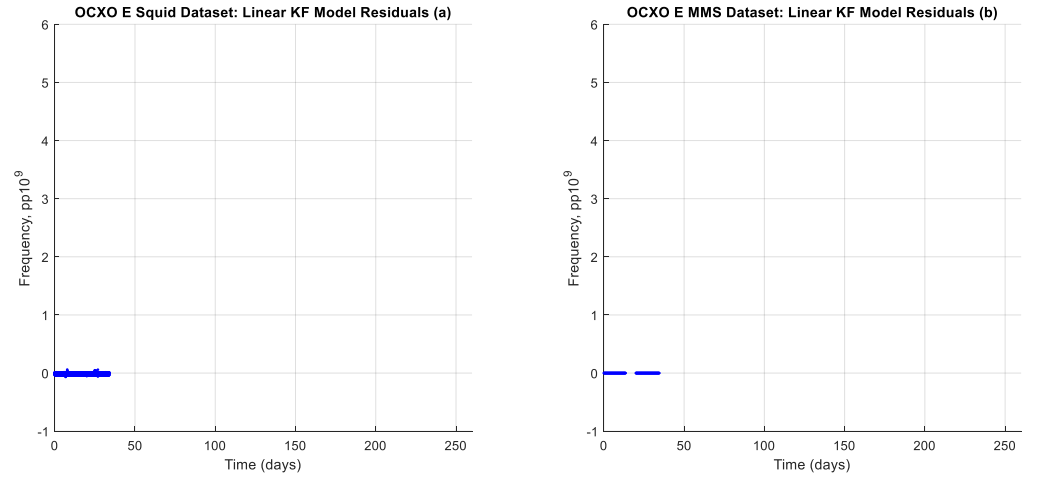

Figure 73: OCXO E linear Kalman Filter residuals

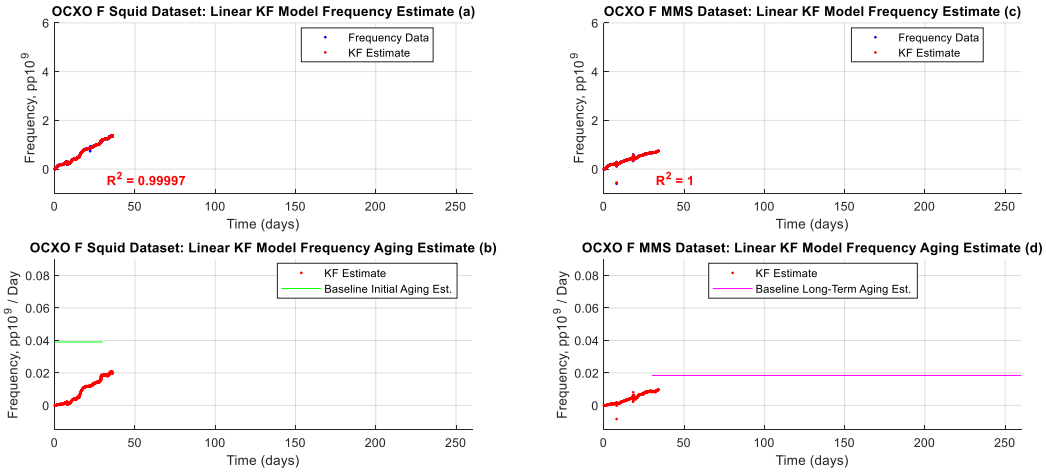

Figure 74: OCXO F linear Kalman Filter modeling results

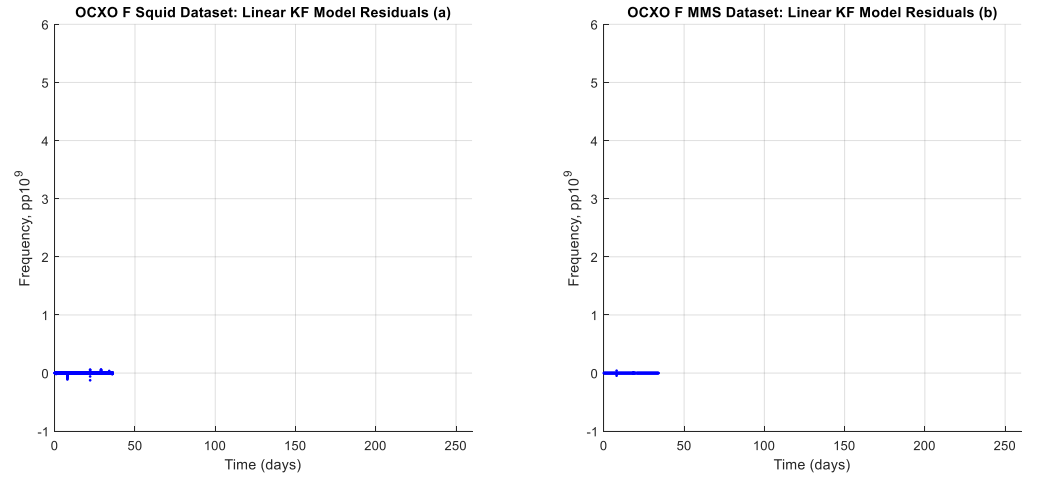

Figure 75: OCXO F linear Kalman Filter residuals

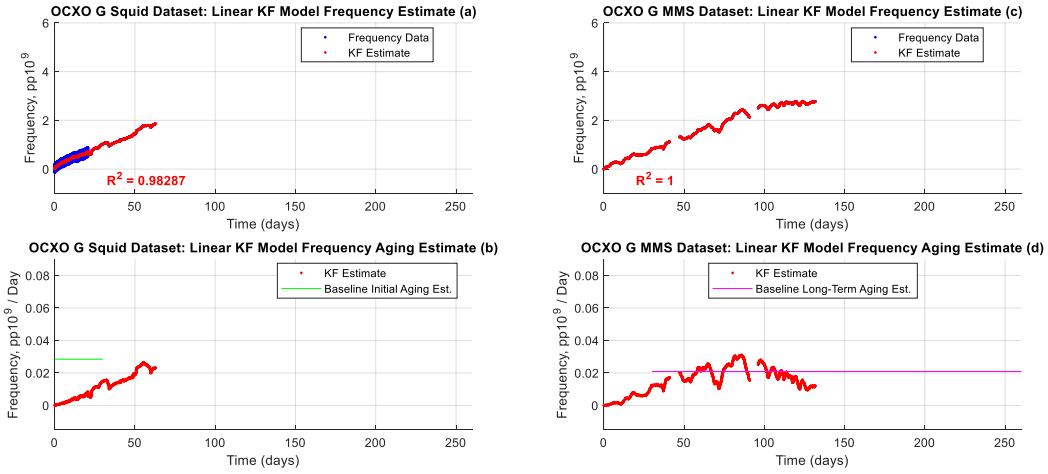

Figure 76: OCXO G linear Kalman Filter modeling results

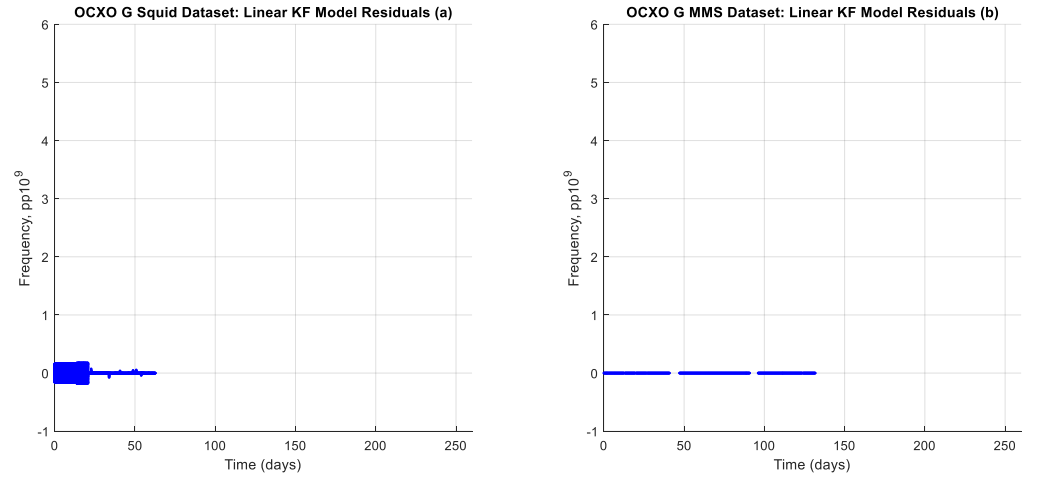

Figure 77: OCXO G linear Kalman Filter residuals

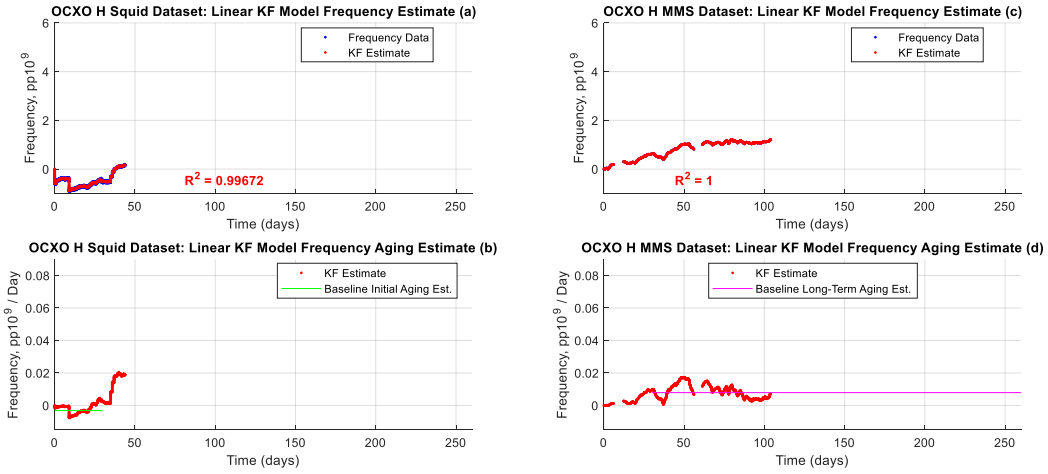

Figure 78: OCXO H linear Kalman Filter modeling results

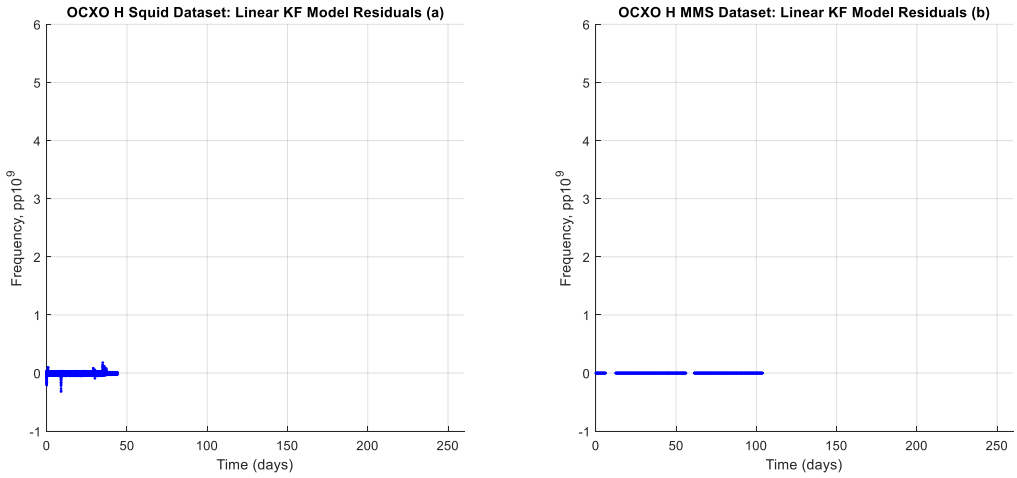

Figure 79: OCXO H linear Kalman Filter residuals