PECULIAR VARIATIONS OF WHITE DWARF PULSATION FREQUENCIES AND MAESTRO

by

James Ruland Dalessio

A dissertation submitted to the Faculty of the University of Delaware in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Physics

Spring 2013

PECULIAR VARIATIONS OF WHITE DWARF PULSATION FREQUENCIES AND MAESTRO

by

James Ruland Dalessio

Approved: _____
Edmund R. Novak, Ph.D.
Chair of the Department of Physics and Astronomy

Approved: _____
George H. Watson, Ph.D.
Dean of the College of Arts and Sciences

Approved: _____
James G. Richards, Ph.D.
Vice Provost for Graduate and Professional Education

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: _____

Henry L. Shipman, Ph.D.
Professor in charge of dissertation

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: _____

Judith L. Provencal, Ph.D.
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: _____

James MacDonald, Ph.D.
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: _____

Donald E. Winget, Ph.D.
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: _____

Stephen M. Barr, Ph.D.
Member of dissertation committee

ACKNOWLEDGEMENTS

Firstly, I thank Judi Provencal, who has put up with my antics for over 6 years. For all practical purposes, Judi served as my research advisor throughout my graduate career. Judi found the perfect balance of guidance and freedom for me to grow throughout graduate school and we've had a lot fun along the way. I'd also like to thank Harry Shipman, who served as my official advisor. Harry was always there to mentor me in both my teaching and my research.

The data used for this research would not have been obtained if it weren't for the diligent work of Denis Sullivan, Fergal Mullally, and JJ Hermes. To them, I am forever indebted for the many hours they spent observing and reducing data. I'd also like to thank Roy Østensen for his work searching for a DB in the Kepler field and spearheading Kepler observations of KIC 8626021.

Whenever I had a tough theoretical question, Mike Montgomery was always happy to help whether it was through email or Skype. Fergal Mullally offered extremely helpful suggestions and discussion as referee of my first major publication. I learned a lot about computers from Staszek Zola, and we had a lot of laughs. I also enjoyed time spent with members of the WET team, including Susan Thompson, Kepler, Agnes Kim, Denis Sullivan, Don Winget, and the many many others.

In the early years of graduate school, I spent many hours studying with fellow graduate students including Josh Wickman, Dana Saxson, and Mary Oksala. Thanks to you folks for keeping me sane in the first few years. The *Final Countdown* will forever remind me of the qualifiers. I also met many graduate students in my field at other institutions. I always look forward to seeing JJ Hermes, Ross Falcon, Brad Barlow, Paul Chote, Bart Dunlap, and others. I have greatly enjoyed our conversations over the last few years.

I've had the opportunity to TA, instruct several courses, and develop lab curriculum. I enjoyed TA'ing for Peter Georgopolus and working with Almas Khan, Abby Pillitteri, and Mary Oksala along the way.

I thank Dermott Mullan and those at the Delaware Space Grant Consortium for two full years of funding. I also thank Mount Cuba/Crystal Trust for funding me for the first few semesters.

Without some physical activity I would have gone (completely) crazy. Special thanks to Mark Poindexter and the folks at Wilbur for organizing our weekly basketball games and to Joe Wiseman, Lori Marinucci, and the New Castle county softball crew for organizing the many years of softball. I've recently also enjoyed playing softball (and post-game activities) with the Honey Badgers, despite our incredibly terrible performance on the field.

In last few years, John Meyer and I have become good friends and have had many epic physics and philosophy conversations over lunch. I've also enjoyed our mini video game marathons. Its been a blast and I look forward to starting a software company with John in the very near future.

Thanks to my mom, my in-laws Art and Diane, and my sister in-law Amanda, who all took turns taking care of Aria while I worked on the final stretch of my Ph.D. Extra thanks to my mom, who has always supported me, even in those years when I was especially "difficult". A special thanks to my step-father for his helpful advice and diligent proof-reading, and for keeping me in touch with mainstream astronomy. My father did not get to see me finish graduate school, but no doubt he would be very proud.

Lastly and most of all, I thank my wife, who is as patient as she is beautiful. During the course of graduate school Lindsay and I were engaged, married, and welcomed our first child. Lindsay has supported me during every step of this long process and I wouldn't have made it without her.

To my wife

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

In Part I we report on variations of the normal mode frequencies of the pulsating DB white dwarfs EC 20058-5234 and KIC 8626021 and the pulsating DA white dwarf GD 66. The observations of EC 20058-5234 and KIC 8626021 were motivated by the possibility of measuring the plasmon neutrino production rate of a white dwarf, while the observations of GD 66 were part of a white dwarf pulsation timing based planet search. We announce the discovery of periodic and quasi-periodic variations of multiple normal mode frequencies that cannot be due to the presence of planetary companions. We note the possible signature of a planetary companion to EC 20058-5234 and show that GD 66 cannot have a planet in a several AU orbit down to half a Jupiter mass. We also announce the discovery of secular variations of the normal mode frequencies of all three stars that are inconsistent with cooling alone. Importantly, the rates of period change of several modes of KIC 8626021 are consistent with evolutionary cooling, but are not yet statistically significant. These modes offer the best possibility of measuring the neutrino production rate in a white dwarf. We also observe periodic and secular variations in the frequency of a combination mode that exactly matches the variations predicted by the parent modes, strong observational evidence that combination modes are created by the convection zone and are not *normal* modes. Periodic variations in the amplitudes of many of these modes is also noted. We hypothesize that these frequency variations are caused by complex variations of the magnetic field strength and geometry, analogous to behavior observed in the Sun.

In Part II we describe the *MAESTRO* software framework and the *MAESTRO REDUCE* algorithm. *MAESTRO* is a collection of astronomy specific *MatLab* software developed by the Whole Earth Telescope. *REDUCE* is an an algorithm that can extract the brightness of stars on a set of CCD images with minimal configuration and

human interaction. The key to this algorithm is automatic identification of stars and a sophisticated implementation of geometric hashing.

Part I

Peculiar Variations of White Dwarf Pulsation Frequencies

Chapter 1

BACKGROUND, FORMALISM, AND METHODOLOGY

1.1 White Dwarf Asteroseismology

1.1.1 Introduction

Just a teaspoon of white dwarf matter has about the mass of a school bus. However exotic this may seem, the vast majority of stars, including the Sun, will live out their post-main sequence days as white dwarfs. In a sense, the "death" of most stars is the birth of a white dwarf. Like pathologists, we can study these "dead" stars to learn about how stars live and die. We can also study white dwarfs to learn about the behavior of matter in some of the most extreme and unique conditions in the universe.

When a typical main sequence star exhausts its hydrogen fuel source, the core collapses and hydrogen shell burning begins around the inert helium core. This causes the star to swell into a red giant. Eventually, helium burning will begin in the core and the star will contract. More massive stars begin helium burning immediately after the initial core collapse. After the core helium has been fused into carbon and oxygen, helium shell burning begins and the star again swells into a red giant. Variability in the helium shell burning rate causes substantial mass loss and ejection of the envelope. What is left behind is essentially the inert core of a star, now an extremely hot pre-white dwarf. When the pre-white dwarf cools to below about $75,000\,\mathrm{K}$, gravitational contraction ceases due to the rigid resistance of electron degeneracy pressure. At this point the star is considered a young white dwarf. Without the ability to further collapse and with no other energy sources, the white dwarf will spend its years cooling into electromagnetic oblivion.

White dwarfs are classified by their dominant spectroscopic features. The most common are the hydrogen white dwarfs (DA) followed by the helium I (DB), helium

II (DO), metallic (DZ), featureless (DC), and carbon (DQ) white dwarfs. Some white dwarfs are given multiple classifications, for example the DAZs show prominent hydrogen and metallic lines. Note that this classification scheme is purely observational and white dwarfs may transition between flavors as they evolve. Of these many flavors of white dwarfs, the DAs and DBs make up the vast majority. Although the DAs and DBs are spectroscopically very different, their interiors are structurally similar. They are composed of mostly carbon and oxygen surrounded by a relatively thin layer of helium, and for the DAs, an additional layer of hydrogen.

One of the primary ways we study individual white dwarfs is by observing the way they *pulsate*. The study of pulsating stars is called *asteroseismology*. Consider striking a bell with a hammer; the bell rings. The quality and tone of the sound depends on many factors, such as the shape and composition of the bell. Just like the ringing of a bell, nature has its own way of ringing stars. The ringing, or *pulsating*, often causes stars to change in brightness. Just as the ear can distinguish the properties of a church bell from that of a gong without visual inspection, the properties of a white dwarf can be inferred from the nature of the pulsations. We can also study changes in the physical properties of white dwarfs as the nature of the pulsations change. Many other types of stars are pulsating, including our own Sun. The pulsations of the Sun have allowed us to map out the Sun's internal structure in incredible detail (Thompson et al., 2003).

Pulsations are the result of *normal mode oscillations*. A normal mode oscillation is a periodic variation of physical quantities with fixed phase. The frequencies (or periods) of the possible normal mode oscillations are determined by the physical nature of the system. The *amplitude* of a normal mode oscillation is a measurement of the size of the variations and an *excited* normal mode has an amplitude greater than zero. *Driving* processes, like striking a bell with a hammer or pushing a child on a swing, increase the amplitudes of normal mode oscillations while *damping* processes decrease the amplitudes of normal mode oscillations. A system may have multiple normal modes excited in superposition, or in other words, the brightness of a star may vary as the

sum of many sinusoids, one for each excited normal mode oscillation.

Each normal mode of a white dwarf can be identified by a set of numbers that represent the way the oscillations propagate throughout the star. These numbers are called the radial overtone, spherical degree, and spherical order and are represented by the integers k, l, and m respectively. They are analogous to the quantum numbers n, l, and m of electrons in an atom. The radial overtone, or k, is the number of nodes between the center and the surface of the star and the spherical degree and order (l and m) describe the angular geometry of the mode in terms of Laplace's spherical harmonics Y_l^m.

White dwarfs are observed to pulsate with periods of minutes to tens of minutes. These oscillations correspond to normal *gravity* modes with spherical degree $l = 1$ (and possibly $l = 2$) and k values of less than 50. White dwarfs that pulsate have a "V", for "variable", appended to their spectral classification. For example, a pulsating DA is known as a DAV. DA white dwarfs are observed to pulsate when their effective temperatures are between about $11,000\,\mathrm{K}$ and $12,500\,\mathrm{K}$ and DB white dwarfs are observed to pulsate when their effective temperature is between about $22,000\,\mathrm{K}$ and $28,000\,\mathrm{K}$ (Gianninas et al., 2005; Nitta et al., 2009). Its important to note that all typical white dwarfs pulsate in these temperature ranges; the bulk properties of pulsating white dwarfs are the bulk properties of non-pulsating white dwarfs (Fontaine et al., 1985). These two temperature ranges are known as the classical white dwarf *instability strips*. A DO instability strip exists at very high temperatures, but by most standards the DOVs are pre-white dwarfs. There are also two new types of white dwarfs that were recently discovered to vary in brightness: the DQVs (Montgomery et al., 2008; Dufour et al., 2008) and the ELMVs (Hermes et al., 2012; Van Grootel et al., 2013). The ELMVs, or "extremely low mass variables", are a special class of DAs and are technically part of a low mass extension of the DA instability strip. The DQVs are carbon atmosphere pulsators whose exact nature is currently being debated (Williams et al., 2012).

The two classical instability strips correspond to a temperature range where

there is a thin convective region at the surface of the white dwarf. It is thought that the convection, or at the very least the change in opacity that causes the convection, is responsible for "ringing" the star, i.e. driving the pulsations. Remarkably, calculations of the temperature of partial ionization for DB atmospheres predicted the existence and temperature of the DBVs before a DB was actually discovered to pulsate (Winget et al., 1982).

1.1.2 Cooling Processes

The periods of the normal modes are sensitive to the thermal properties of the white dwarf; as the star cools, the periods of the normal modes should increase. Due to the separation of thermal and mechanical structure in a white dwarf, the rate at which the period of a mode increases, $\dot{P}$, is very directly related to the total luminosity. This simplicity is highlighted by the highly successful model of white dwarf cooling presented by Mestel (1952). Essentially, a white dwarf cools as an isothermal conducting sphere surrounded by an insulating ideal gas. No substantial improvements have been made to this model except for superimposing other physical effects, such as surface convection, neutrino production, and crystallization.

For the DAVs, the energy loss is primarily due to thermal radiation and $\dot{P}$ is expected to be of order 10^{-15} (Bradley et al., 1992). However, if the axion (a dark matter candidate) mass is large, it could increase $\dot{P}$ by a measurable amount (Isern et al., 1992). The periods of the DAs G117-B15A and R548 have been monitored since the 1970s and $O-C$ analysis (see Section 1.4) yields statistically significant $\dot{P}s$ that are aligned with theoretical predictions (Kepler et al., 2005; Mukadam et al., 2009). The $\dot{P}$ of G117-B15A was used to put an upper limit on the DFSZ axion mass of around 30 meV (Bischoff-Kim et al., 2008).

In free space, the massless photon cannot decay into a neutrino/anti-neutrino pair. This is not the case in the core of a white dwarf, where a photon couples to the plasma and picks up a "mass" proportional to the plasma frequency. This allows the decay of the photon into a neutrino/anti-neutrino pair to conserve energy and

momentum. The neutrinos freely escape the core of the white dwarf. Production of plasmon neutrinos is the dominant cooling mechanism for hot white dwarfs, including the hottest DBVs. A measurement of a hot DBVs $\dot{P}s$ would be a unique low-energy test of the unified electroweak theory of lepton interactions (Winget et al., 2004).

For the DBVs, the expected $\dot{P}s$ are more than an order of magnitude greater than the DAVs (Córsico and Althaus, 2004) and a measurement of the neutrino production rate could be obtained with less than a decade of observations. For some time, the leading candidate for the measurement of the neutrino production rate in a DB was EC 20058-5234 but the periods of this star were actually found to vary much faster than expected from neutrino cooling alone (Dalessio et al., 2013). These results are shown in Chapter 2. Very recently, a hot DBV was discovered in the Kepler field of view (Østensen et al., 2011; Bischoff-Kim and Østensen, 2011; Córsico et al., 2012) and is now the leading prospect for measuring the neutrino production rate in a white dwarf. Currently, the frequency of only a few of the normal modes of KIC 8626021 are behaving in a manner consistent with evolutionary cooling. These results are shown in Chapter 3.

1.1.3 Rotation and Magnetic Fields[1]

Normal modes with the same k and l have identical frequencies in a non-rotating non-magnetic star. For a rotating star, the frequencies of the normal modes are altered by the Coriolis force. For slow rotation and assuming the normal modes are aligned with the rotation axis, the observed frequency change is of order the rotation frequency and is proportional to m. In other words, the frequencies of the normal modes are *split* into $2l - 1$ frequencies. For example, if all three $k = 1$, $l = 1$ modes are excited, the difference in frequency between the $m = 1$ and $m = 0$ modes would be equal to the frequency difference between the $m = 0$ and $m = -1$ modes. These three frequencies would be referred to as a *triplet* and the difference between the frequencies is referred to as a *splitting*. If the star is differentially rotating radially or latitudinally,

[1] For a detailed treatment of these effects see Unno et al. (1989).

the splittings are altered based on a complicated integral involving the rotation rate and the eigenfunctions of the normal modes. However, the splittings would still be equal for modes with the same k and l.

A weak magnetic field affects the normal mode frequencies in a similar manner as rotation, except that the frequency changes depend on the absolute value of m, not the sign of m. For a rotating star, this essentially breaks the symmetry in the splitting of modes with the same k and l. For stronger fields, the situation gets extremely complicated and the normal modes can be misaligned with the rotation axis.

For white dwarfs, normal modes with large k have more nodes near the surface of the star, and would generally be more affected by physical processes near to the surface. If the surface of a star was rotating faster than the core, an $l = 1$ triplet of modes with high k would generally have a larger splitting than an $l = 1$ triplet with lower k. This should also hold true for magnetic fields that are confined to the surface; the asymmetry in splittings should generally increase with k.

1.1.4 Pulsation Timing Based Planet Detection

A planetary companion will cause a small variation in the position of a host star as the two masses orbit their barycenter. The motion of the star will cause variations in the time it takes for the star's light to reach an observer. If the star is providing regular signals, like pulsations, an observer would perceive periodic variations between the expected and actual arrival times of signals. This is the basis of the *pulsation timing* planet detection method.

While the Doppler shift method is more sensitive to planets very close to the star, the pulsation timing method is just the opposite. Take, for example, Jupiter and the Sun. Jupiter causes the Sun to orbit roughly about its radius with a period of nearly 12 years. The Doppler shift caused by this slow motion is extremely small. However, the Sun's orbit is several lightseconds large, causing the delay or acceleration of the arrival time of the light from the Sun of up to several seconds (for some external observer). If Jupiter were twice as far away, the delay would be twice as large.

The Sun may not provide reliable regular signals, but a white dwarf certainly can. Take as an example the DAV G117-B15A, the most stable known optical clock (Kepler et al., 2005). For G117-B15A, the analysis of Mullally et al. (2008) ruled out a *remarkably large* amount of planet parameter space.

In the early 2000s, astronomers at the University of Texas began a search for the signature of a planetary companion to a white dwarf. The most notable result was the discovery of periodic variations between the expected and actual pulsation arrival times of the 302 s normal mode of the DA GD 66 (Mullally et al., 2008). One explanation offered was a $> 2\,\mathrm{M_J}$ planet in a 4.5 yr orbit. Upper limits on the size of the possible planet were placed with Spitzer spectroscopy (Mullally et al., 2009), and Hubble Space Telescope (HST)/United States Naval Observatory (USNO) precision astrometry (Farihi et al., 2012). Then, with the discovery of similar variations between the expected and actual pulsation arrival times for the DBV EC 20058-5234 that were not due to a planetary companion, a large amount of skepticism was raised about the planet hypothesis for GD 66 (Dalessio et al., 2013). Further analysis of GD66's pulsation arrival times reveal that the variations are not due to a planetary companion. These results are presented in Chapter 4.

1.1.5 Combination and Harmonic Modes

White dwarfs are observed to pulsate at frequencies that are equal to the sum of two or more normal mode frequencies (*combinations*) or an integer times the normal mode frequencies (*harmonics*). In some cases, *difference* frequencies are also observed. These variations are not thought to be *normal mode* oscillations, but rather artifacts created by the distortion of normal mode oscillations as they pass through the outer convective layers of the star (Brickhill, 1992; Wu, 2001). The relative amplitudes of these artifacts are dependent on observational effects like inclination angle, along with physical effects like normal mode geometry and the properties of the outermost layers of the star. Montgomery (2005) shows that by fitting the observed variations of a white dwarfs brightness, the properties of the convection zone can be determined.

This method is effectively fitting the relative amplitudes of these combination and harmonic "modes". If the combination and harmonic modes are not entirely artificial and oscillations are actually occurring at these frequencies, the results obtained using the methods of Montgomery (2005) could be unreliable. While this notion is against conventional wisdom, observational confirmation is still important. In Chapter 2 we will present the most convincing observational evidence to date that combination modes are indeed artificial.

The ratio between the amplitude of a combination or harmonic mode to the product of its parents amplitudes, R_C, is a quantity partially dependent on the properties of the convection zone and the geometry of the pulsations (Wu, 2001) and can be used to investigate relative mode geometry (see e.g. Yeates et al. 2005; Handler et al. 2002). We investigate R_C as a function of time for GD 66 in Chapter 4.

1.2 Miscellaneous Considerations

1.2.1 Creating a Lightcurve

Modern measurements of the brightness of pulsating white dwarfs are made using digital photography. Here we briefly describe the general procedure used to *reduce* a series of raw images into a series of brightness measurements, or a *lightcurve*. For a detailed description of the methods involved, see Howell (2006). The images are first calibrated by removing electronic bias, thermal variations, and pixel to pixel sensitivity differences. The flux of each star on images is then calculated using synthetic apertures or PSF (point spread function) techniques. The background flux is also measured and removed from the flux of the stars. The measurements of flux versus time of the white dwarf is assembled into a raw lightcurve. Flux variations introduced by weather conditions are removed by dividing the white dwarf's lightcurve by the lightcurve of a comparison star. The lightcurve is then divided by a low order polynomial to remove variations caused by color differences between the variable and comparison stars. The lightcurve is then subtracted by unity to put it in units of fractional flux deviation

(ma). The times of the measurements are converted to the corresponding times at the barycenter of the solar system(see e.g. Eastman et al. 2010; Stumpff 1980).

1.2.2 Linearizing a Non-Linear Model

Many of the models in Chapters 2, 3, and 4 will consist of a polynomial plus a sinusoid. This model is is only linear if the period of the sinusoid is fixed. If the period is allowed to vary, initial guesses must be made for the coefficients of the polynomial and the amplitude, period, and phase of the sinusoid. The final fit parameters are only guaranteed to correspond to a *local* minimum in χ^2. A way to avoid this problem is to obtain linear fits over a grid of fixed range of periods. This procedure is commonplace in this dissertation.

1.3 Fourier Analysis

1.3.1 The Fourier Transform

1.3.1.1 Definition

The Fourier transform is a mathematical operation that converts a function in time to complex sinusoidal amplitudes as a function of frequency. The Fourier transform of a function $a(t)$ is

$$\tilde{A}(f) = \int_{-\infty}^{\infty} a(t) e^{-2\pi i f t} dt \tag{1.1}$$

Consider a lightcurve consisting of a set of N brightness measurements a_n at times t_n. The normalized *discrete Fourier transform* (DFT) of the lightcurve is

$$\tilde{A}(f_k) = \frac{2}{N} \sum_{n=0}^{N} a_n e^{-2\pi i f_k t_n} \tag{1.2}$$

$\tilde{A}^2$ is referred to as the *spectral power* and $|\tilde{A}|$ is often referred to as the *spectral amplitude*. Spectral power or amplitude as a function of frequency is referred to as the *Fourier spectrum*.

If the lightcurve is sampled at regular intervals, the effective frequency limits of the DFT are 0 and half the sampling rate of the lightcurve (the Nyquist frequency) and

the maximum frequency resolution goes as the inverse length of the lightcurve. For a lightcurve where the data is not sampled in uniform intervals, the effective frequency limits and resolution are difficult to quantify. In this case, it is a common and prudent practice to calculate the DFT with a frequency sampling rate an order of magnitude larger than the inverse of the length of the lightcurve. A common choice for the upper frequency limit of non-uniformly sampled data is half the median sampling rate of the lightcurve.

Computationally, the DFT is an expensive calculation; the number of operations goes as $O(N^2)$. In situations where the lightcurve is long or the DFT must be calculated many times, speed can become an issue. A solution to this problem is to interpolate the measurements into a uniformly spaced lightcurve and use the $O(N\log(N))$ fast Fourier transform (Press and Rybicki, 1989), i.e. the "fast periodogram". Here all calculations of the DFT are made using *MAESTRO*'s implementation of the fast periodiogram.

1.3.1.2 Properties

Each amplitude of the DFT, $\tilde{A}(f_k)$, represents a sinusoid defined as

$$A_k \sin\left(2\pi f_k t - \phi_k\right) \tag{1.3}$$

where

$$A_k = \mid \tilde{A}_k \mid, \;\; \phi_k = \pi + \arctan(\frac{\Re\tilde{A}_k}{\Im\tilde{A}_k}) \tag{1.4}$$

$\tilde{A}(f_k)$ minimizes the sum of the square difference between the lightcurve and the sinusoid at frequency f_k. In other words, the DFT is just a non-simultaneous linear least squares fit to a sinusoid at each frequency. Consider a lightcurve that is sampling some continuous variations described by the function $h(t)$. The lightcurve can be written as

$$a(t) = h(t)\sum_{0}^{N}\delta(t - t_n) \tag{1.5}$$

and the Fourier transform of the time series, $\tilde{A}(f)$, can be written in terms of the Fourier operator, $\mathcal{F}$,

$$\tilde{A}(f) = \mathcal{F}\Big(a(t)\Big) = \mathcal{F}\Big(h(t)\sum_{0}^{N}\delta(t-t_n)\Big) \tag{1.6}$$

Applying the convolution theorem, i.e. $\mathcal{F}(A \cdot B) = \mathcal{F}(A) * \mathcal{F}(B)$, Equation 1.6 becomes

$$\tilde{A}(f) = \mathcal{F}\Big(h(t)\Big) * \mathcal{F}\Big(\sum_{0}^{N}\delta(t-t_n)\Big) \tag{1.7}$$

where the $*$ symbol indicates convolution. Equation 1.7 shows that the Fourier transform of a discrete time series is the convolution of the Fourier transform of the actual function and the Fourier transform of the sampling of the function. The Fourier transform of the sampling of the function, the right most term in Equation 1.7, is known as the *spectral window*.

The spectral window couples the spectral power at all frequencies. The spectral window is a common problem in time series astronomy. Minimization of the amplitude of secondary spikes of power in the spectral window, or *aliases*, is the primary motivation for the Whole Earth Telescope (Nather, 1989).

1.3.2 Identifying Sinusoidal Variations in a Lightcurve

The lightcurves of pulsating white dwarfs can typically be reproduced by a sum of a small number of sinusoids. In the ideal case of a infinitely long continuous lightcurve, there will be a spike in the Fourier spectrum at the frequency of each sinusoid and the remaining Fourier spectrum will be zero. For real world observations, it cannot be assumed that power in the Fourier spectrum directly corresponds to the amplitudes of the sinusoids needed to reproduce the lightcurve; the situation is complicated by the spectral window. However, the spectral window introduced by a sinusoid with known amplitude and frequency can be removed from the Fourier spectrum by first subtracting the signal from the lightcurve. This is known as *pre-whitening*. Note that it is impossible to truly remove the effects of the spectral window; pre-whitening merely clarifies information that already exists in the DFT. A common procedure, here

referred to as *recursive pre-whitening*, is commonly used to identify the frequencies of sinusoidal variations that reproduce a lightcurve.

Given an initially empty table of frequencies and a continuous lightcurve, each iteration of the recursive pre-whitening procedure begins by *linearly* fitting a sinusoid to the lightcurve at each of the frequencies simultaneously (during the first iteration, there are no frequencies to fit). The derived amplitudes and phases are then used as initial guesses for *non-linearly* fitting a sinusoid to the lightcurve at each frequency simultaneously, now allowing the frequencies to vary. The frequencies in the table are then updated. The DFT of the residuals is calculated and the frequency with the largest spectral power is identified. If the spectral power is statistically significant, the frequency is added to the table and the procedure is repeated. If the spectral power is not statistically significant the procedure is terminated. A prudent criteria is that the spectral amplitude be at least 6 times greater than the median of the DFT amplitudes.

This process has a major caveat: variations in the frequency, amplitude, or phase of a single sinusoidal signal are indistinguishable from a sum of sinusoids with fixed frequencies, amplitudes, and phases. Consider a single sinusoidal variation at frequency f with an amplitude that varies sinusoidally at frequency f_A. The Fourier transform of this signal will have power at 3 frequencies f, $f + f_A$, and $f - f_A$. It is impossible to mathematically distinguish this signal from one composed of three sinusoidal signals with fixed amplitudes. Care must be taken to note the multiple interpretations and to only chose one if it is supported by a physical model.

1.3.3 Frequency Bootstrapping

Frequency Bootstrapping is a procedure used to combine lightcurves in order to increase the precision of the measurement of the frequency. Here, each lightcurve should be nearly continuous with no large observational gaps. Given a table of frequencies identified in a lightcurve, the procedure begins by performing a linear fit of a sinusoid to the lightcurve at each frequency simultaneously. The resultant amplitudes and phases are used as initial guesses for a non-linear fit of a sinusoid to the lightcurve at

each frequency simultaneously, now allowing the frequencies to vary. The next closest lightcurve (in time) is identified. It is ensured that the standard error in all frequencies is small enough so that there is little uncertainty to the number of cycles in the gap between the lightcurves. Mathematically, the measurement of each frequency must satisfy

$$\sigma_f^{-1} \ll \Delta t \tag{1.8}$$

where σ_f is the standard error in the frequency and Δt is the time between the lightcurves. If this relationship is satisfied, the lightcurves are combined and the procedure is repeated. The procedure continues until either all of the lightcurves are combined, or when Equation 1.8 cannot be satisfied. If one or some subset of frequencies does not satisfy Equation 1.8 during an iteration it can be retained in the list of frequencies for future iterations but should be discarded at the end of the procedure.

Choosing the right lightcurve to begin this procedure can be critical, but if the procedure is sucessfull, the frequencies should converge to the same values no matter the starting point. Generally, the first lightcurve should be one that has good signa to noise, spans a large amount of time, and is close (in time) to other lightcurves.

1.4 The $O - C$ Method

1.4.1 Introduction

$O - C$, or "observed minus calculated", is the difference between when a event occurs and when it is predicted to occur. If a lightning strike is estimated to occur 700 m away and the sound speed is 350 $\frac{\mathrm{m}}{\mathrm{s}}$, the sound of the thunder would be predicted to arrive 2 s later. If the sound of the thunder actually arrives 3 s later, $O - C = 1\,\mathrm{s}$. In this simple case, $O - C$ reveals an error in the estimate of the distance to the lightning.

When events occur in regular time intervals, the period or frequency of the events can be used to predict the time of future and past events. Take, for example, local ocean tides on the fictional planet Kerbin. If after many high tides, the average time between high tides is measured to be 12 hr and a high tide is measured at precisely noon, a high tide would then be predicted to occur at noon and midnight daily.

Given a sufficient baseline, $O-C$ is extremely sensitive to variations in frequency and can reveal changes that are not detectable by direct measurement. To illustrate, consider a a poorly manufactured clock. The second hand of this clock ticks one microsecond before a full second has elapsed. To the naked eye, the difference between the ticks on this and a perfect clock are indistinguishable. However, if the clock is synchronized to local time at noon, when the clock strikes noon a week later the actual time will be six tenths of a second past noon, i.e. $O-C = .6\,\mathrm{s}$. A remarkable example of the ability of $O-C$ to measure small changes in frequency is the case of the white dwarf G117-B15A. Observations of the 215 s pulsation mode reveal a statistically significant rate of period change of less than 10^{-14} (Kepler et al., 2005).

The power of $O-C$ has put it at the forefront of some of the greatest discoveries in the history of astronomy. The $O-C$ of pulsars revealed the first exoplanets (Wolszczan and Frail, 1992; Thorsett et al., 1993) and the slow orbital decay that earned Hulse and Taylor the Nobel Prize for their detection of gravitational wave emission (Taylor et al., 1979). $O-C$ is used in many other sub-fields of time domain astronomy, see e.g. Oksala et al. (2012); Barlow et al. (2011); Mullally et al. (2008); Bischoff-Kim et al. (2008).

1.4.2 Mathematical Description

The examples in Section 1.4.1, i.e. a flash of lightning, high tide, or a tick of the clock, are all "events"; they all take place at discrete times. If an event is observed at time O_0 and the typical time between events is measured to be P, each event would be calculated (predicted) to occur at times

$$C = O_0 + P \cdot N \tag{1.9}$$

where N is an integer. If an event is then observed at some time O, then

$$O - C = O - O_0 - P \cdot N \tag{1.10}$$

Note that $N-1$ is the number of events that occur between O and O_0.

A event can be defined to occur when an oscillation reaches a special instantaneous phase. For example, each high tide is an event that takes place at a special instantaneous phase of the water depth oscillation. An oscillation can be modeled as a sinusoidal variation with semi-amplitude A, fixed frequency f, and time varying absolute phase $\tau = \tau(t)$, i.e.

$$H(t) = A \sin\big(2\pi f\,(t-\tau)\big) \tag{1.11}$$

The instantaneous phase, T, is the argument of the sinusoid, i.e.

$$T = 2\pi f(t-\tau) \mod 2\pi \tag{1.12}$$

If the oscillation is observed to be at instantaneous phase T at time $O_0 = \tau_0 + \frac{T}{2\pi f}$, the same instantaneous phase would be predicted at for all times

$$C = \tau_0 + \frac{T}{2\pi f} + P \cdot M \tag{1.13}$$

where M is and integer and $P \equiv \frac{1}{f}$. If the instantaneous phase is again observed to be T at some time O

$$O = \tau + \frac{T}{2\pi f} \tag{1.14}$$

than $O - C$ can be written as

$$O - C = \tau - \tau_0 - P \cdot M \tag{1.15}$$

The absolute phase is usually measured to be between $-\frac{1}{2}P$ and $\frac{1}{2}P$. M is a correction to the absolute phase, and is not equivalent to N as defined above. M is only non-zero if $O - C + \tau_0 < -\frac{1}{2}P$ or $O - C + \tau_0 > \frac{1}{2}P$. For example, if $O - C = P$ and $\tau - \tau_0 = 0$, $M = -1$.

Equation 1.15 shows the observational relationship between $O - C$ and the absolute phase at fixed frequency. A series of measurements of the absolute phase are directly related to $O - C$.

The examples shown so far are all of behaviors that occur at discrete times. However, the ocean tides could be modeled as a oscillation in the depth of the water.

Measurements of the absolute phase of the water depth oscillation are related to $O-C$ by Equation 1.15.

Invariant of physical cause, all variations of $O-C$ can be treated as variations of the frequency of the behavior. In this section, a derivation of the relationship between $O-C$ and the frequency of the behavior is presented.

The instantaneous amplitude, $H(t)$, of an oscillation with semi-amplitude A and a time dependent frequency $f' = f'(t)$ is

$$H(t) = A \sin\left(2\pi \int^t f'(t')dt'\right) \tag{1.16}$$

f' is then rewritten as a constant frequency f plus some time dependent perturbation $\gamma = \gamma(t)$

$$f' = f + \gamma \tag{1.17}$$

Integration yields

$$\int^t f'(t')dt' = ft + \int^t \gamma(t')dt' + const \tag{1.18}$$

The time varying absolute phase, $\tau = \tau(t)$, is then defined, and the constant is absorbed into the new variable τ_0

$$\tau \equiv \tau_0 - \frac{1}{f}\int^t \gamma(t')dt' \tag{1.19}$$

so that

$$\int^t f'(t')dt' = f\,(t-\tau) \tag{1.20}$$

and

$$H(t) = A \sin\big(2\pi f\,(t-\tau)\big) \tag{1.21}$$

A variation in frequency at fixed absolute phase has been rewritten as a variation in absolute phase at fixed frequency. The relationship between $O-C$ and the frequency is thus

$$\boxed{O - C = \tau_0 - \frac{1}{f}\int^t \gamma(t')dt'} \tag{1.22}$$

In the next few sections, the effect of various perturbations to the frequency will be investigated.

1.4.3 Causes of Perturbations to $O - C$

1.4.3.1 Error in the Frequency

In practice, it is impossible to measure the frequency to infinite precision. The difference between the actual and measured frequency, δf, can be written as a perturbation to the measured frequency

$$\gamma(t) = \delta f \tag{1.23}$$

Equation 1.22 yields

$$O - C = \tau_0 - \frac{\delta f}{f} t \tag{1.24}$$

The difference between the actual frequency and f causes a linear trend, or "drift", in $O - C$.

1.4.3.2 Movement of the Source

Changes in the distance from the source to the observer will cause variations in the time it takes for the signal to reach the observer. Note that variations in $O - C$ as caused by the changes in signal travel time are equivalent to the changes caused by the Doppler shift of the frequency.

The effects of the movement of the source and observer are greatly simplified if the observations are made in an inertial reference frame, i.e. the observer's position is fixed and only the position of the source is allowed to vary. For most time series astronomical observations, corrections are made so that the time of the observations are recorded in the (nearly) inertial reference frame of the barycenter of the solar system. Otherwise the orbital motion of the Earth can cause ~ 10 minute variations in $O - C$ over the course of a year. Tools exist to automatically perform the conversion from UTC/Julian Date to BJD, such as WQED (Thompson and Mullally, 2009), which contains a c port of the algorithm of Stumpff (1980), or an IDL implementation by Jason Eastman (Eastman et al., 2010). The latter is available via the web. In general, it will be assumed that the observer is in an inertial reference frame and that the position of the source is measured is the frame of reference of the observer.

If the distance between the source and observer is r, the Doppler effect gives the relationship between the radial velocity, $v_r \equiv \frac{dr}{dt}$, and the perturbation to the frequency

$$\gamma = -f\left(\frac{v_r}{v_r + c}\right) \tag{1.25}$$

where c is the speed of the signal. If the speed of the source is much less than the radial velocity, the perturbation to the frequency is

$$\gamma = -f\left(\frac{v_r}{c}\right) \tag{1.26}$$

and

$$O - C = \tau_0 + \frac{1}{c}\int^t v_r(t')dt' \tag{1.27}$$

which simplifies to

$$O - C = \tau_0 + \frac{\delta r}{c} \tag{1.28}$$

where δr is the time dependent perturbation to the distance between the observer and source.

If the source has a constant radial velocity, $\delta r = v_R t$, Equation 1.28 becomes

$$O - C = \tau_0 + \frac{v_R}{c}t \tag{1.29}$$

A constant radial velocity causes a linear trend in $O - C$. Note, however, that if the frequency of the behavior is measured when the source already has radial velocity v_r, the difference between the the actual and Doppler shifted frequency will exactly cancel this effect. Constant radial velocities are often no concern to $O - C$ as frequency of a behavior is *always* Doppler shifted and the measurement of the frequency inadvertently accounts for the Doppler shift.

Proper motion will also affect $O - C$. If the proper velocity of the source is v and closest approach occurs at a distance d at time $t = 0$, the radial velocity is

$$v_r = \frac{v^2 t}{\sqrt{d^2 + v^2 t^2}} \tag{1.30}$$

and if the movement of the source is small compared to the distance to the source, i.e. $vt \ll d$,

$$v_r = \frac{v^2 t}{d} \tag{1.31}$$

Equation 1.28 becomes

$$O - C = \tau_0 + \frac{v^2}{2cd} t^2 \tag{1.32}$$

Proper motion causes a parabolic variation in $O - C$.

The orbital motion of a star and a companion causes a perturbation to the distance of the star from an observer. If the star is pulsating or providing a regular time variable signal, the changes in distance will cause variations in $O - C$. If the star is in a circular orbit, the perturbation to to the distance can be approximated as

$$\delta r = a \sin(i) \sin(\frac{2\pi}{\Pi} t) \tag{1.33}$$

where a is the radius of the orbit, i is the inclination of the orbit relative to the observer, Π is the period of the orbit, and the time system has been defined so that the absolute phase of the sinusoidal distance perturbations 0. Plugging into Equation 1.28 yields

$$O - C = \tau_0 + \frac{a}{c} \sin(i) \sin(\frac{2\pi}{\Pi} t) \tag{1.34}$$

Orbital motion of the source will cause a sinusoidal variation in the $O - C$ diagram. Given a measurement of the semi-amplitude $\delta\tau \equiv \frac{a}{c} \sin(i)$ and Π from successive measurements of $O - C$, Kepler's 3rd law gives the relationship between the masses of the star (m_s) and companion (m_c) as

$$\frac{{m_2}^3 \sin^3(i)}{(m_1 + m_2)^2} = \frac{4\pi^2 c^3}{G} \frac{\delta\tau^3}{\Pi^2} \tag{1.35}$$

where G is the gravitational constant. In the case where $m_2 \ll m_1$

$$m_2 \sin(i) = (\frac{4\pi^2 c^3}{G})^{\frac{1}{3}} (\frac{m_1}{\Pi})^{\frac{2}{3}} \delta\tau \tag{1.36}$$

This effect of orbital motion on $O - C$ is used extensively as a tool for searching for exoplanets around pulsating stars, see e.g. Mullally et al. (2008); Silvotti et al. (2007); Wolszczan and Frail (1992).

1.4.3.3 Intrinsic Frequency Perturbations

The variations in $O - C$ caused by movement of the source and the difference between the actual and measured frequency are observational effects; they are not caused by actual changes in the behavior. Actual changes to the frequency of the behavior will also cause variations in $O - C$. Consider an extremely small perturbation to the frequency, $\gamma \approx \dot{f}t \ll f$. Equation 1.22 becomes

$$O - C = \tau_0 - \frac{1}{2}\frac{\dot{f}}{f}t^2 \tag{1.37}$$

It is often convenient to rewrite this in terms of the period, $P \equiv f^{-1}$

$$O - C = \tau_0 + \frac{1}{2}\frac{\dot{P}}{P}t^2 \tag{1.38}$$

A slow change in frequency (or period) will lead to a parabolic change in $O - C$.

Now consider a sinusoidal perturbation in frequency with semi-amplitude α, period Π, and phase ϕ, i.e.

$$\gamma(t) = \alpha \sin\left(\frac{2\pi}{\Pi}t - \phi\right) \tag{1.39}$$

Equation 1.22 becomes

$$O - C = \tau_0 + \frac{\alpha\Pi}{2\pi f}\cos\left(\frac{2\pi}{\Pi}t - \phi\right) \tag{1.40}$$

A sinusoidal variation in frequency will result in a sinusoidal variation in $O - C$.

1.4.3.4 Indistinguishable Processes

Some very different phenomena cause identical variations in $O - C$. Care must be taken to examine all possible interpretations. If the source provides multiple repetitive behaviors, e.g. a star pulsating at several frequencies, the effects of motion will cause similar variations in the $O - C$ diagram of each behavior.

A transverse velocity (Section 1.4.3.2) has the same effect as a slowly changing frequency. If the effects of transverse velocity is mistaken as a slow change in period, or vice-versa

$$\dot{P} = \frac{Pv^2}{dc}, \qquad v = \sqrt{\frac{dc\dot{P}}{P}} \tag{1.41}$$

This effect was first noted by Shklovskii (1970) as a possible explanation for the observed period changes in pulsars and must be taken into account when measuring small frequency changes of the pulsations of white dwarfs (see e.g. Kepler et al. 2005). Equation 1.41 can also be expressed in terms of more convenient units for astronomers (Pajdosz, 1995)

$$\dot{P} = 2.430 \times 10^{-18} \;\; P[s] \;\; (\mu[\frac{''}{yr}])^2 \;\; d[pc] \tag{1.42}$$

A long period sinusoidal variation in $O-C$ could be interpreted as a slow change in frequency. A sinusoid with semi-amplitude A and period Π varies near minimum as $\frac{2\pi^2 A}{\Pi^2}t^2$. If long period orbital motion were interpreted as a slow change in the period of the behavior

$$\frac{{m_2}^3 \sin^3(i)}{(m_1+m_2)^2} = \frac{c^3 \Pi^4}{16\pi^4 G}\left(\frac{\dot{P}}{P}\right)^3 \tag{1.43}$$

and if $m_2 \ll m_1$

$$m_2 \sin(i) = (\frac{c^3}{16\pi^4 G})^{\frac{1}{3}} \Pi^{\frac{4}{3}} {m_1}^{\frac{2}{3}} \left(\frac{\dot{P}}{P}\right) \tag{1.44}$$

If long period sinusoidal variation in frequency (Section 1.4.3.3) is interpreted as a slow change in frequency

$$\frac{\alpha}{\Pi} = \frac{\dot{P}}{2\pi P^2} \tag{1.45}$$

1.4.4 Calculating $O-C$

Given a lightcurve and a frequency, $O-C$ is calculated by performing a linear least squares fit of a sinusoid at this frequency. If $O-C$ is to be calculated at multiple frequencies the fit should be performed simultaneously. The fit model can be described as

$$M(t) = \sum_i \alpha_i \sin\left(2\pi f_i \left(t - \tau_i\right)\right) \tag{1.46}$$

and can be written in terms of linear parameters a_i and b_i

$$M(t) = \sum a_i \sin\left(2\pi f_i\right) - b_i \cos\left(2\pi f_i\right) \tag{1.47}$$

where

$$a_i = \alpha_i \cos\left(2\pi f_i \tau_i\right), \quad b_i = \alpha_i \sin\left(2\pi f_i \tau_i\right) \tag{1.48}$$

$O - C$, or τ, and the amplitude, α, can be determined from the linear parameters

$$\tau_i = \left(2\pi f_i\right)^{-1} \arctan\left(\frac{b_i}{a_i}\right), \quad \alpha_i = \left(a_i^2 + b_i^2\right)^{\frac{1}{2}} \tag{1.49}$$

A plot of a series of $O - C$ measurements is often called an $O - C$ diagram. The ordinate and abscissa of the $O - C$ diagram are both expressed in units of time, but are sometimes normalized by the period of the behavior into units of cycles (in astronomy often labeled "epochs" on the abscissa), radians, or degrees. We will refer to a series of $O - C$ measurements as an "$O - C$".

Equation 1.24 shows that a difference between the measured and actual time averaged mode frequency will cause a linear trend in $O - C$. If this difference is large enough, each value of $O - C$ can differ by more than a cycle. This would make it practically impossible to constrain $O - C$ to a particular value. The difference between the measured and actual time averaged mode frequency should be small enough that the trend in $O - C$ is small, or in other words the standard error in the number of cycles predicted to elapse between successive measurements of $O - C$ is less than one. This leads to the same relationship as Equation1.8; the *bootstrapping* process should be used to obtain a precise measurement of the frequency when calculating $O - C$.

1.4.5 Phase Ambiguity

If the dataset is discontinuous there is always phase ambiguity, i.e. all points in $O - C$ can be shifted by an integer N times the period, but this ambiguity might be considered resolved by Occam's Razor if $O - C$ (aka τ) is smooth i.e.

$$\frac{d\tau}{dt}\Delta t \ll P \tag{1.50}$$

where Δt is the gap between consecutive measurements of $O - C$ and P is the period of the mode. $\frac{d\tau}{dt}$ could be carefully calculated for parts of $O - C$ from spline fitting or other methods, but Equation 1.50 should be very clearly satisfied, enough that if

a numerical calculation is actually needed it ought be assumed that Equation 1.50 is violated. We refer to a sequence of $O - C$ values as *unambiguous* if it clearly satisfies Equation 1.50.

If a subset of $O - C$ measurements are well described by some model or a model is known a priori, this model can be used to constrain the other values of $O - C$ with phase ambiguity. If the model predicts some value M and the ambiguous value is $O + N \times P$ then N is a statistically good fit only if χ^2 is of order one, i.e.

$$(M - O - N \times P)^2 \sim \sigma_M^2 + \sigma_O^2 \tag{1.51}$$

and if the observed and modeled values of $O - C$ are actually constrained in phase, i.e.

$$\sigma_M, \sigma_O \ll P \tag{1.52}$$

A newly constrained point in $O - C$ can then be included in the model fit, refining any other measurements of M, which may allow constraint of N for additional points in $O - C$. Note, this process could be further improved by introducing a Bayesian approach, for example by including an estimation of the probability the model is valid given the existing data.

Chapter 2

EC 20058-5234

2.1 Observations and Data Preparation

EC 20058-5234 is the second most well observed DB pulsator; since its discovery in 1994, its brightness has been measured over 250,000 times. These observations were primarily motivated by the possibility of measuring the plasmon neutrino production rate of a white dwarf, one of the holy grails of asteroseismology (see Winget et al. 2004).

The observations began with the discovery of EC 20058-5234 in 1994 (Koen et al., 1995) at South African Astronomical Observatory (SAAO) as part of the Edinburgh-Cape (EC) survey (Stobie et al., 1992). In 1997, EC 20058-5234 was chosen as the primary southern target of the Whole Earth Telescope (WET) campaign XCOV15 (Sullivan et al., 2008). Seasonal photometric observations followed from the 1 m McLellan telescope at Mount John University Observatory (MJUO) from 1998 though the present, highlighted by over 40,000 photometric measurements over the course of 5 months in 2004 (Sullivan, 2005, 2009). In 2003, spectral and photometric charge-coupled device (CCD) observations were obtained on the 6.5 m Magellan telescopes at Las Campanas Observatory (LCO) (Sullivan et al., 2007). Seasonal CCD photometry was also obtained between 2004 and 2011 with the Small and Moderate Aperture Research Telescope System (SMARTS) 36" at Cerro Tololo Inter-American Observatory (CTIO). In the Spring of 2009, EC 20058-5234 was selected as a tertiary target for the WET campaign XCOV27. The observations were assembled into 36 individual lightcurves as shown in Table 2.1. A portion of the lightcurve obtained with one of the 6.5 m Magellan telescopes in 2003 is shown in Figure 2.1.

Table 2.1: Journal of Observations

Start	End	Observatory	CCD/PMT	Measurements
1994 May 14	1994 Jun 02	SAAO	PMT	5535
1994 Jun 12	1994 Jun 12	SAAO	PMT	1378
1994 Jul 06	1994 Jul 12	SAAO	PMT	6881
1994 Oct 03	1994 Oct 03	SAAO	PMT	1304
1997 Jul 02	1997 Jul 11	WET[1]	PMT	46289
1997 Oct 05	1997 Oct 06	MJUO	PMT	3162
1998 Jul 24	1998 Jul 25	MJUO	PMT	7057
1998 Aug 14	1998 Aug 14	MJUO	CCD	1153
1999 Sep 09	1999 Sep 13	MJUO	PMT	1517
2000 Jul 06	2000 Jul 08	MJUO	PMT	8025
2000 Sep 05	2000 Sep 05	MJUO	PMT	770
2001 Mar 30	2001 Apr 01	MJUO	PMT	2101
2001 Sep 21	2001 Sep 24	MJUO	PMT	7710
2002 Apr 12	2002 Apr 15	MJUO	PMT	8872
2002 Aug 01	2002 Aug 06	MJUO	PMT	12672
2002 Sep 05	2002 Sep 09	MJUO	PMT	3814
2003 Jul 11	2003 Jul 13	LCO	CCD	2317
2003 Jul 25	2003 Jul 31	LCO	CCD	8865
2003 Aug 27	2003 Sep 02	MJUO	PMT	12216
2003 Sep 22	2003 Sep 23	MJUO	PMT	4694
2003 Oct 31	2003 Oct 31	MJUO	PMT	1229
2004 Apr 23	2004 Apr 25	MJUO	PMT	2474
2004 May 16	2004 May 16	MJUO	PMT	1599

Continued on Next Page...

[1] This WET campaign, XCOV15, included observations from SAAO, CTIO, MJUO, and Observatorio Pico dos Dias (OPD).

Table 2.1 – Continued

Start	End	Observatory	CCD/PMT	Measurements
2004 Jun 09	2004 Jun 16	MJUO	PMT	14554
2004 Jul 09	2004 Jul 14	MJUO	PMT	17240
2004 Aug 07	2004 Aug 10	MJUO	PMT	5072
2004 Aug 24	2004 Aug 28	CTIO	CCD	2320
2005 Sep 21	2005 Sep 25	CTIO	CCD	979
2006 Aug 31	2006 Sep 13	CTIO	CCD	5019
2007 Jul 16	2007 Jul 17	MJUO	PMT	6281
2007 Aug 17	2007 Sep 06	CTIO	CCD	14117
2008 Aug 29	2008 Sep 01	CTIO	CCD	1736
2008 Aug 30	2008 Aug 31	MJUO	CCD	3092
2009 May 18	2009 May 26	WET[2]	CCD	4609
2010 Jul 05	2010 Jul 11	MJUO	CCD	20201
2011 Sep 08	2011 Sep 15	CTIO	CCD	3584

The photomultiplier tube (PMT) photometry from SAAO was reduced as described in Koen et al. (1995), while the XCOV15 and MJUO PMT photometry was reduced using the methods outlined in Sullivan et al. (2008). All CCD calibration and reduction was performed as described in Section 1.2.1. The barycentric corrections applied to the data were all derived from the same algorithm of Stumpff (1980).

2.2 Fourier Analysis

The 1997 WET, July 2004 MJUO, and 2007 CTIO lightcurves were selected as three of the best lightcurves for Fourier analysis. The Fourier spectrum of the 2007

[2] This WET campaign, XCOV27, included observations from the Southern Astrophysical Research (SOAR) telescope, SAAO, and CTIO.

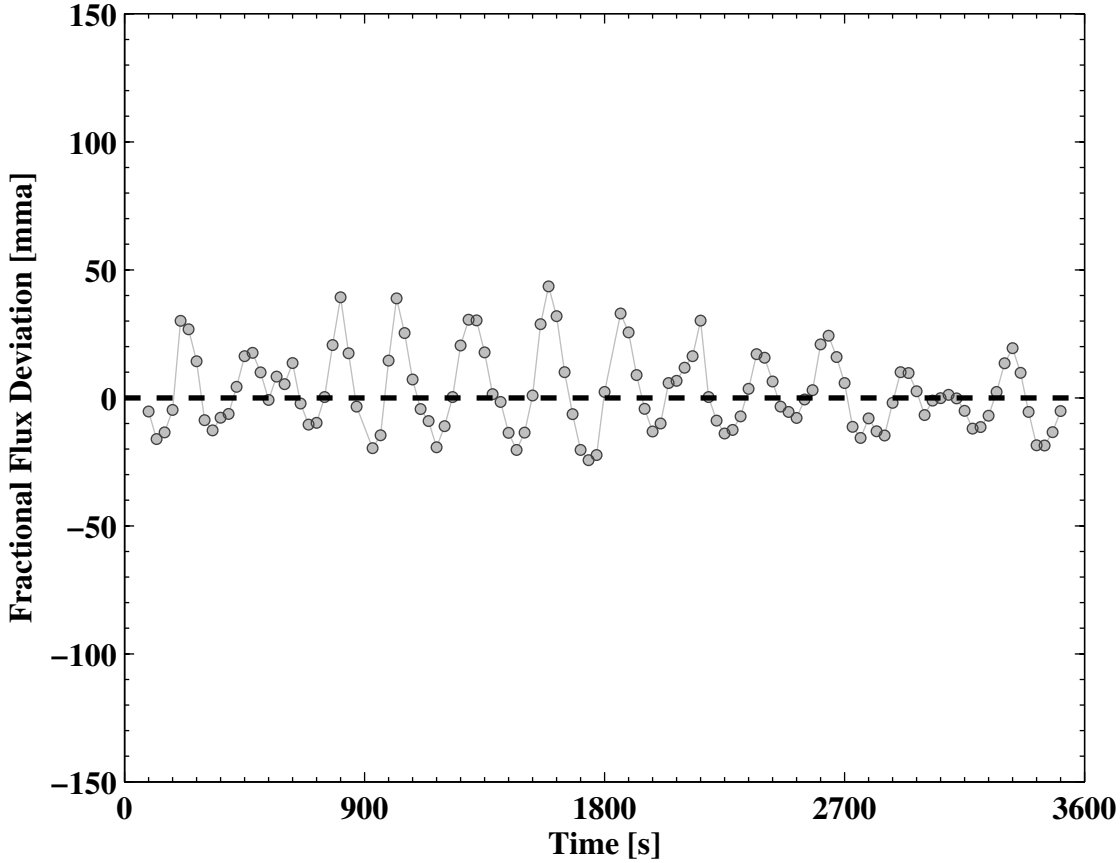

Figure 2.1: A portion of the lightcurve obtained with the Magellan telescopes in 2003.

CTIO lightcurve is shown in Figure 2.2. The Fourier spectra of the three lightcurves are extremely similar. There are two large ~ 10 mma modes at periods of approximately 257 s and 281 s and a handful of other modes with amplitudes less than 5 mma. Modes with a signal to noise above 6 were identified in the 1997 WET lightcurve and assigned a label. This process was also repeated for the July 2004 MJUO and 2007 CTIO lightcurves. The frequencies of the modes in the 2007 CTIO and July 2004 MJUO lightcurves very closely correspond to the frequencies of modes in 1997 WET lightcurve. Modes with frequencies within 1 μHz of each other were assigned the same label. Several combination modes were also found. The frequencies and amplitudes of modes identified in the three lightcurves are shown in Tables 2.2 and 2.3.

The most striking difference from lightcurve to lightcurve is the systematic increase of the amplitudes from 2004 to 2007. At least some of this difference can be attributed to an observational artifact caused by two visual companions within $2''$ and $4''$ of EC 20058-5234. A typical CCD image of EC 20058-5234 is shown in Figure 2.3. The PMT observations included a large amount of flux from the companions while the CCD observations allow at least some removal of the companions' flux. The result is that the same variation in flux appears fractionally larger for CCD observations. Variations in mode amplitude are investigated in greater detail in Section 2.4.

The July 2004 MJUO lightcurve was identified as the best starting point for frequency bootstrapping. The frequencies of the modes were bootstrapped using all lightcurves from 2002 through 2004. Modes $E + J$ and $A + C$ were thought to have too low of an amplitude to be considered and were removed from analysis. Additional bootstrapping beyond 2002 through 2004 resulted in a small increase in the standard error of several of the mode frequencies and was avoided. This is a sign that the frequencies of the modes are changing by a substantial amount during this time period. After bootstrapping, the frequencies of the combination modes were set to the sum of the final frequencies of the parent modes. The results are shown in Table 2.4.

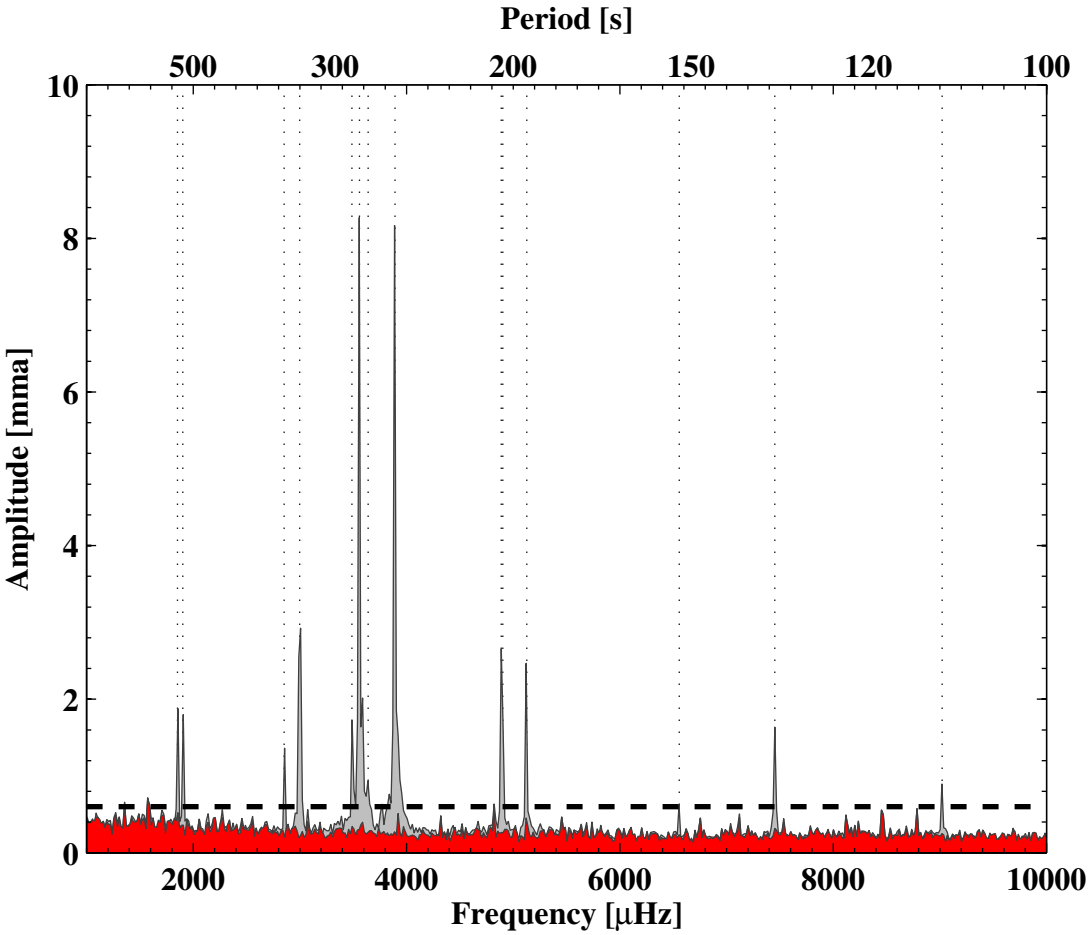

Figure 2.2: Fourier spectrum of the 2007 CTIO lightcurve. The Fourier spectra of the unmodified and prewhitened lightcurve are shown in grey and red respectively. The dotted lines indicate the frequencies listed in Table 2.2 and the dashed line indicates 6 times the median amplitude of the Fourier spectrum of the unmodified lightcurve.

Table 2.2: Frequencies of modes identified in the 1997 WET, July 2004 MJUO, and 2007 CTIO lightcurves.

Label	**Period** $[s]$	**Frequency** $[\mu Hz]$		
		1997	2004	2007
$A+G+I$	100	10030.7(0.1)	-	-
$A+C$	111	9021.68(0.08)	-	9021.94(0.04)
$C+E$	134	7452.21(0.05)	7452.6(0.1)	7452.30(0.02)
$E+G$	152	6557.6(0.1)	-	-
A	195	5128.56(0.03)	5128.37(0.07)	5128.70(0.02)
$G+I$	204	4902.11(0.05)	4902.37(0.08)	4902.12(0.02)
B	205	4887.86(0.03)	4887.92(0.07)	4887.89(0.02)
C	257	3893.205(0.009)	3893.31(0.02)	3893.245(0.005)
D	275	3640.38(0.07)	-	-
E	281	3559.003(0.009)	3559.01(0.03)	3559.032(0.005)
F	287	3489.04(0.05)	3489.0(0.1)	3489.07(0.02)
G	333	2998.70(0.03)	2998.83(0.05)	2998.71(0.01)
H	351	2852.38(0.06)	-	2852.46(0.04)
I	525	1903.53(0.04)	1903.5(0.1)	1903.47(0.02)
J	540	1852.62(0.04)	-	-

Table 2.3: Amplitudes of modes identified in the 1997 WET, July 2004 MJUO, and 2007 CTIO lightcurves.

Label	**Period** $[s]$	**Amplitude** $[mma]$		
		1997	2004	2007
$A+G+I$	100	0.67(0.09)	< 1.2	< 1.2
$A+C$	111	0.91(0.09)	< 1.2	1.4(0.2)
$C+E$	134	1.67(0.09)	1.6(0.2)	2.3(0.2)
$E+G$	152	0.67(0.09)	< 1.2	< 1.2
A	195	2.52(0.09)	2.8(0.2)	3.3(0.2)
$G+I$	204	1.48(0.09)	2.5(0.2)	2.6(0.2)
B	205	2.63(0.09)	2.7(0.2)	3.2(0.2)
C	257	8.39(0.09)	9.2(0.2)	11.3(0.2)
D	275	1.13(0.09)	< 1.2	< 1.2
E	281	8.45(0.09)	7.5(0.2)	10.5(0.2)
F	287	1.59(0.09)	1.7(0.2)	2.2(0.2)
G	333	3.01(0.09)	3.5(0.2)	4.3(0.2)
H	351	1.40(0.09)	< 1.2	1.4(0.2)
I	525	1.85(0.09)	1.7(0.2)	2.4(0.2)
J	540	1.87(0.09)	< 1.2	< 1.2

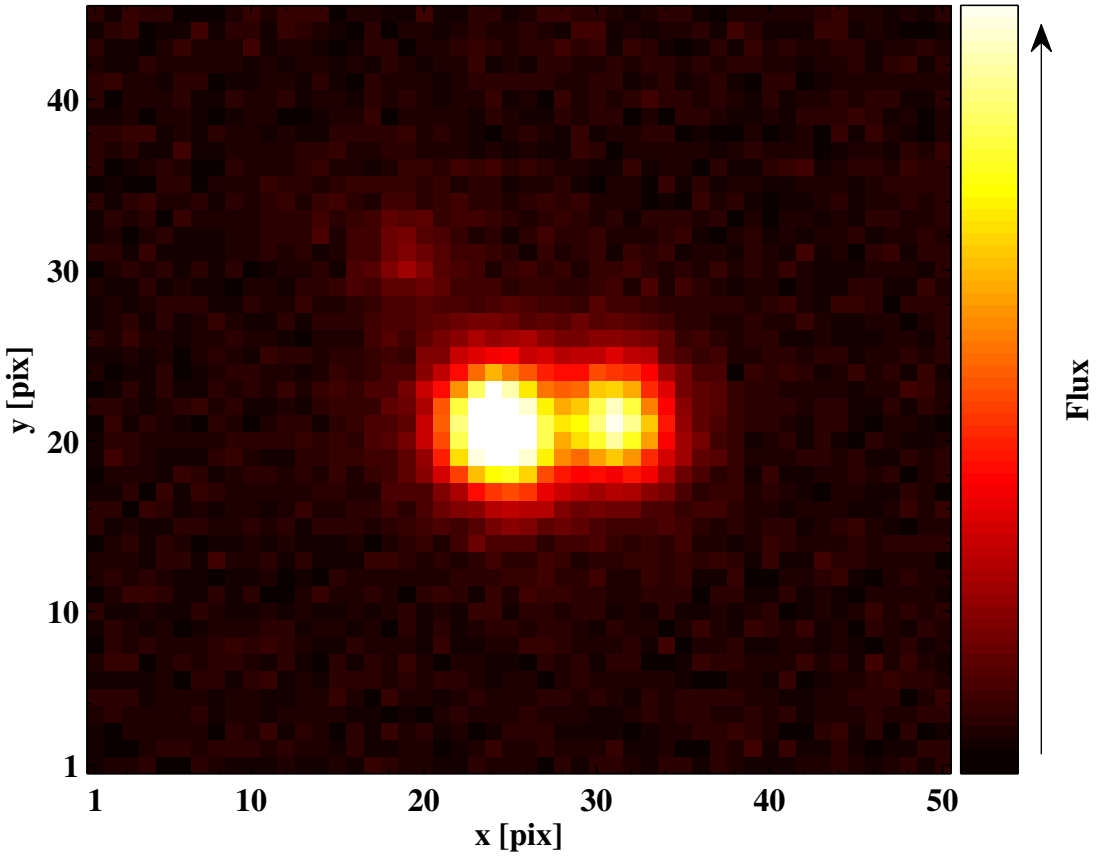

Figure 2.3: Portion of a CCD image taken with a BG 40 filter on the 36" CTIO SMARTS telescope in 2007. Relative flux is indicated by color as demonstrated in the colormap shown to the right of the image. The central bright source is EC 20058-5234. There is a bright visual companion just to the right and a faint visual companion to the upper left.

Table 2.4: Results of bootstrapping the frequencies. The parameter σ_f is the error in the frequency measurement and P_B is the probability that bootstrapping was successful if the frequencies remained constant.

Label	**Period** $[s]$	**Frequency** $[\mu Hz]$	σ_f^{-1} $[yr]$	P_B
$C+E$	134	7452.2476(0.0001)	235	1.00
A	195	5128.1290(0.0005)	58	1.00
$G+I$	204	4902.1820(0.0007)	44	0.98
B	205	4887.8462(0.0004)	88	1.00
C	257	3893.24790(0.00009)	354	1.00
E	281	3558.9997(0.0001)	315	1.00
F	287	3489.0516(0.0004)	78	0.98
G	333	2998.7163(0.0003)	120	1.00
I	525	1903.4657(0.0007)	47	0.98

2.3 $O-C$ Analysis

The $O-C$s of each mode are shown in Figures 2.12-2.11. No unambiguous sequence of $O-C$ values is identifiable for modes A, $G+I$, and I. These modes will not be considered any further in this section. An unambiguous sequence of $O-C$ values is identifiable for modes $C+E$, B, C, E, F, and G between 1997 and 2011. There is, however, ambiguity as to how these sequences of $O-C$ values bridge the $1995-1996$ observational gap.

The method described in Section 1.4.5 was used to constrain which 1994 $O-C$ values belong to the selected 1997 to 2011 sequences. Due to the apparent oscillatory behavior of $O-C$, a sinusoid plus a parabola was selected as a suitable model. The one sigma boundaries of the model predictions are shown in grey in Figures 2.12, 2.6, 2.7, 2.8, 2.9, and 2.10. For each $O-C$ in consideration, the model resolved all ambiguity as to which 1994 values connect to the 1997 to 2011 sequence. In other words, in each case of ambiguity there was only one statistically sensible solution given the model.

The now complete sequences of $O-C$ values were fit with a sinusoid plus

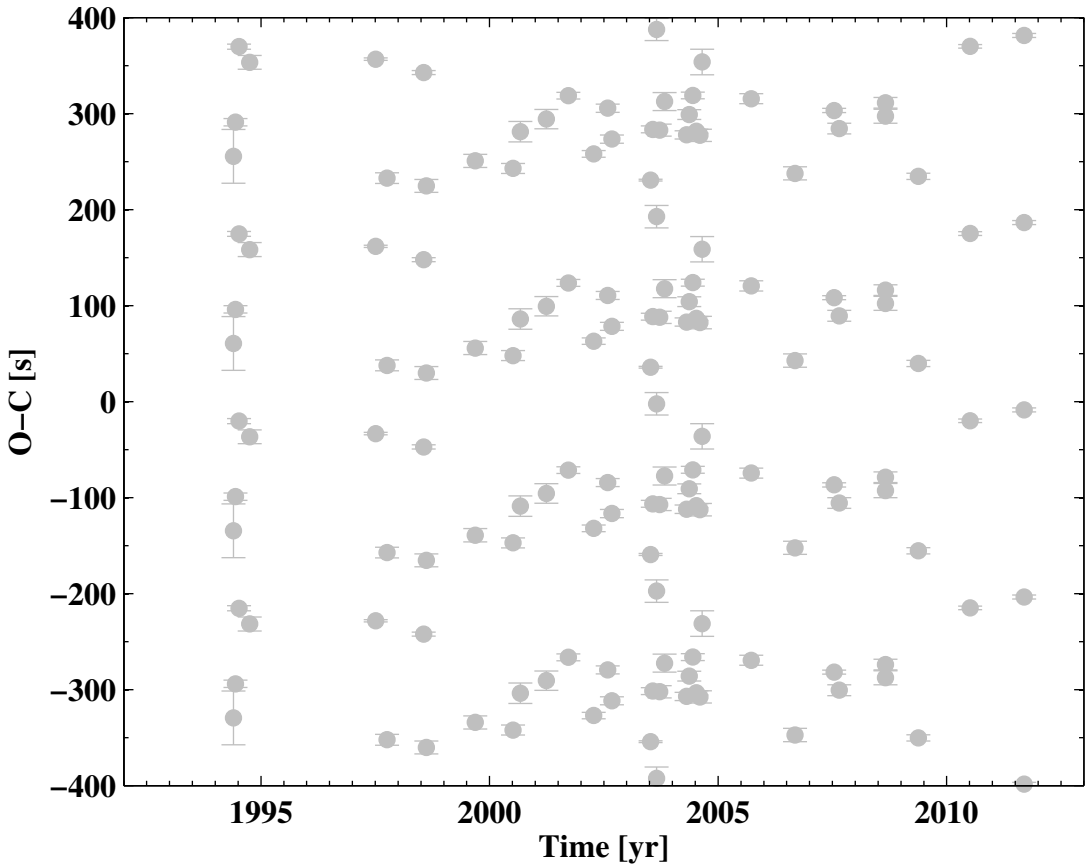

Figure 2.4: $O - C$ of mode A.No unambiguous sequence was identified.

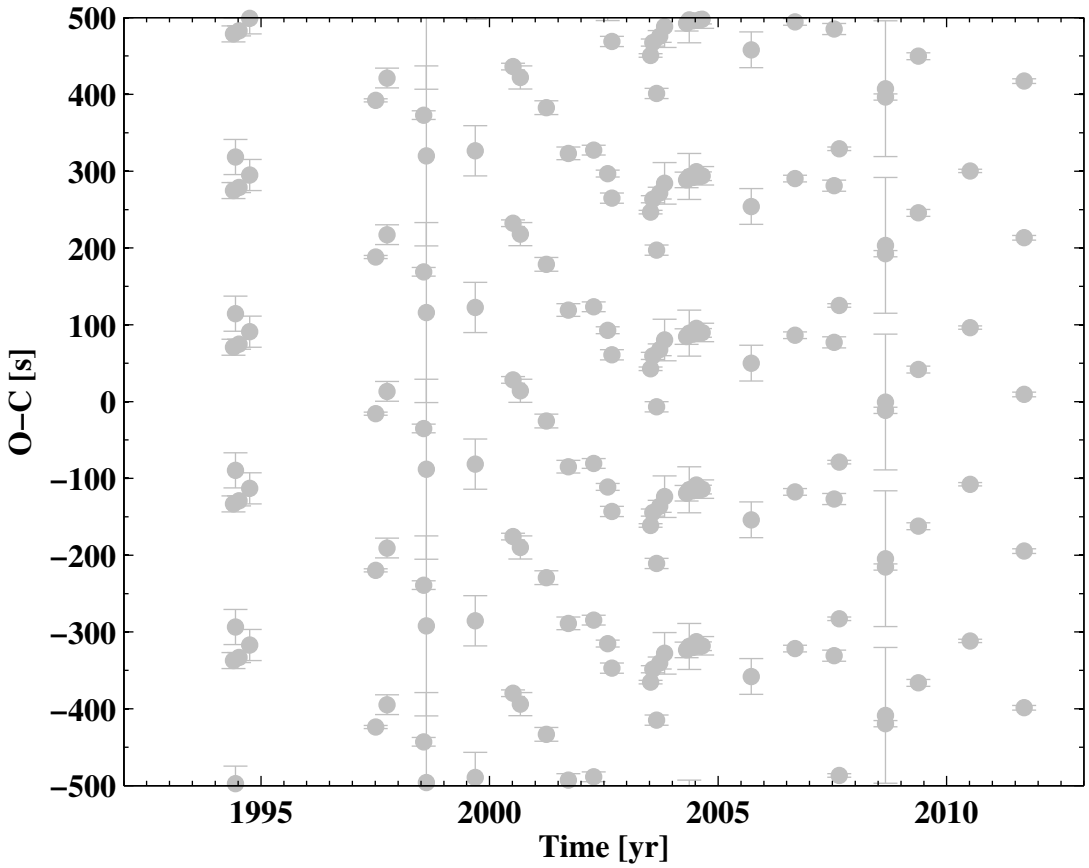

Figure 2.5: $O - C$ of mode $G + I$. No unambiguous sequence was identified.

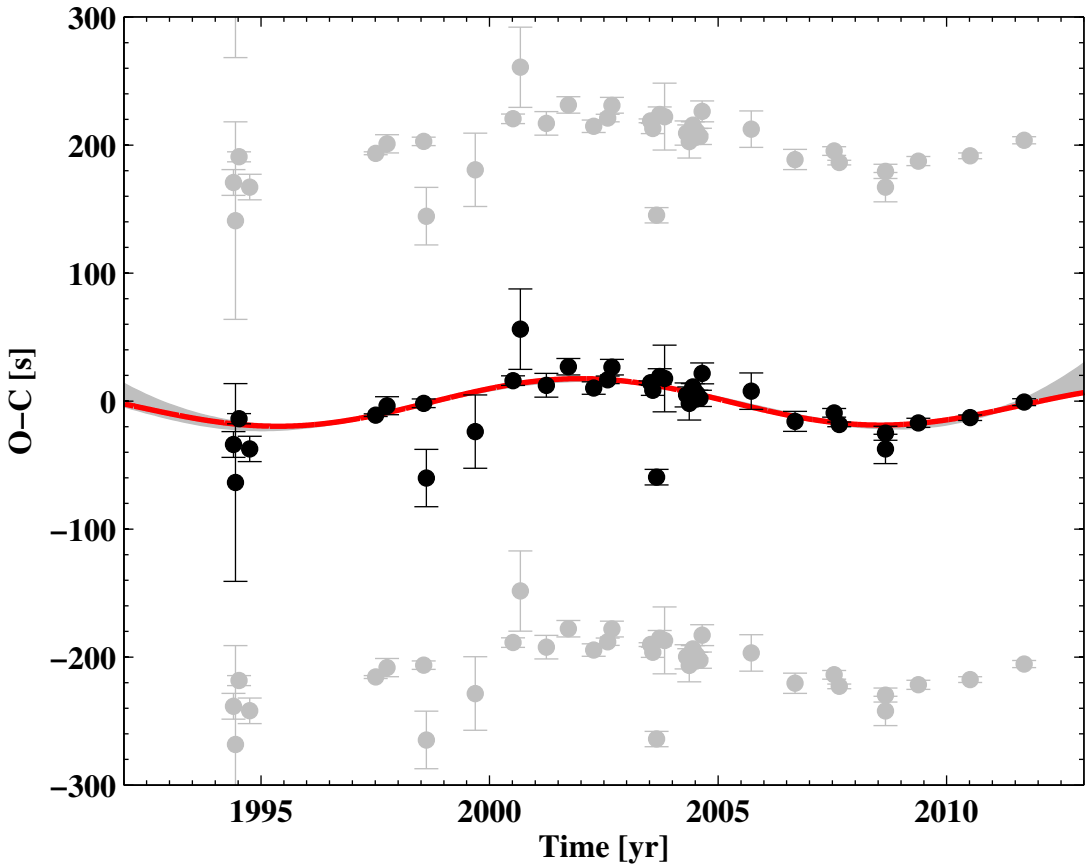

Figure 2.6: $O - C$ of mode B with a linear trend removed. The selected sequence is shown in black and other possible values are shown in grey. The red line is the best fit of a 12.93 year sinusoid plus a parabola. The grey region is the 1σ prediction of the $1997 - 2011$ values for a sinusoid plus parabola.

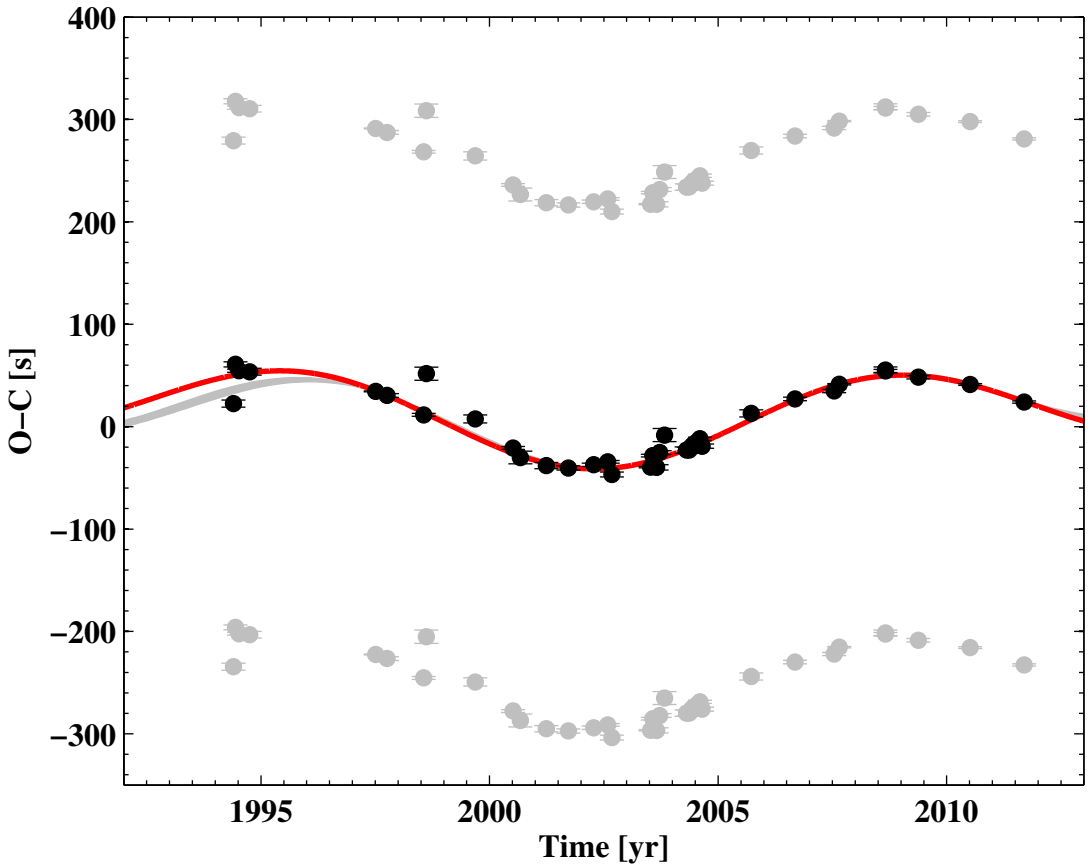

Figure 2.7: $O - C$ of mode C with a linear trend removed. The selected sequence is shown in black and other possible values are shown in grey. The red line is the best fit of a 12.93 year sinusoid plus a parabola. The grey region is the 1σ prediction of the $1997 - 2011$ values for a sinusoid plus parabola.

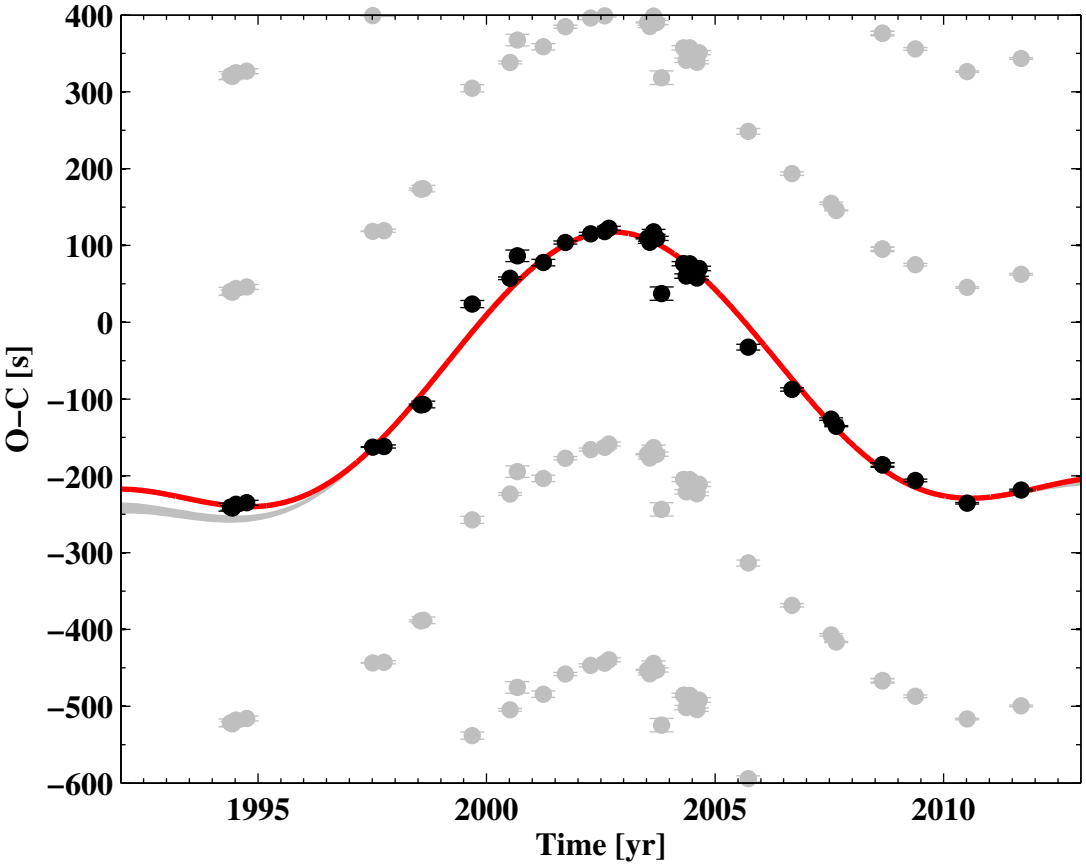

Figure 2.8: $O - C$ of mode E with a linear trend removed. The selected sequence is shown in black and other possible values are shown in grey. The red line is the best fit of a 12.93 year sinusoid plus a parabola. The grey region is the 1σ prediction of the $1997 - 2011$ values for a sinusoid plus parabola.

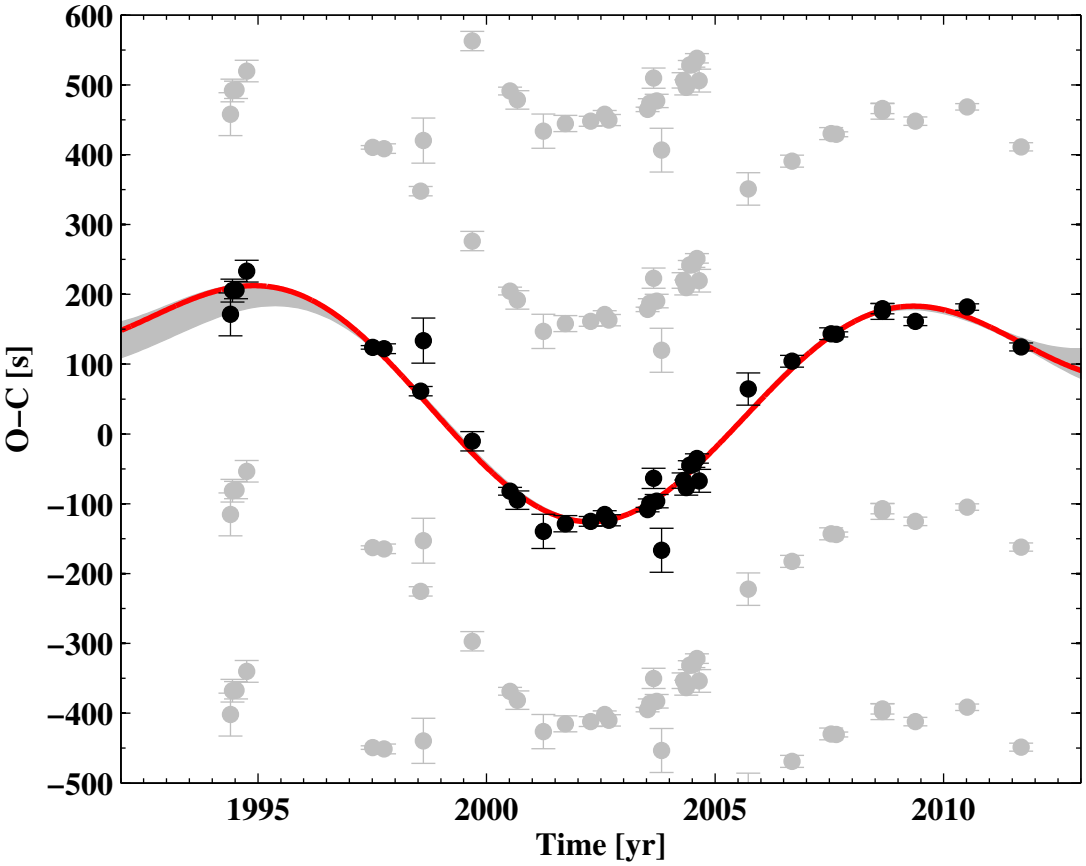

Figure 2.9: $O - C$ of mode F with a linear trend removed. The selected sequence is shown in black and other possible values are shown in grey. The red line is the best fit of a 12.93 yr sinusoid plus a parabola. The grey region is the 1σ prediction of the $1997 - 2011$ values for a sinusoid plus parabola.

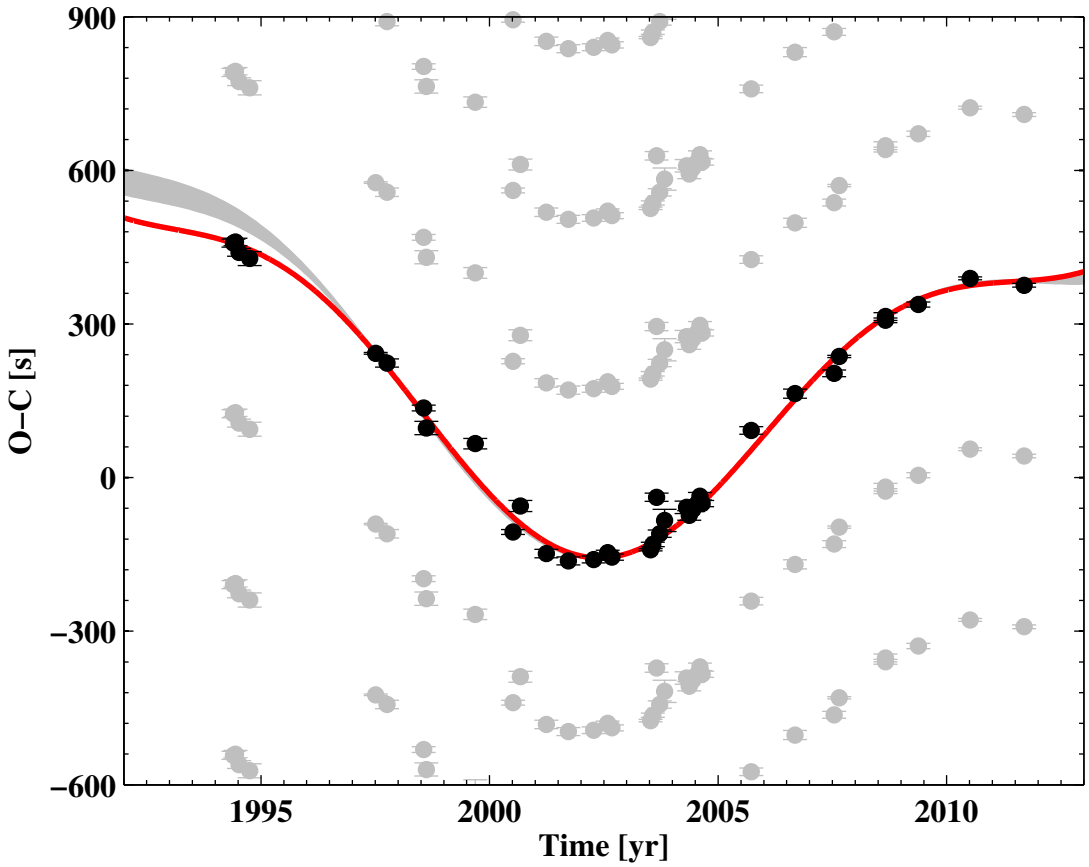

Figure 2.10: $O - C$ of mode G with a linear trend removed. The selected sequence is shown in black and other possible values are shown in grey. The red line is the best fit of a 12.93 year sinusoid plus a parabola. The grey region is the 1σ prediction of the $1997 - 2011$ values for a sinusoid plus parabola.

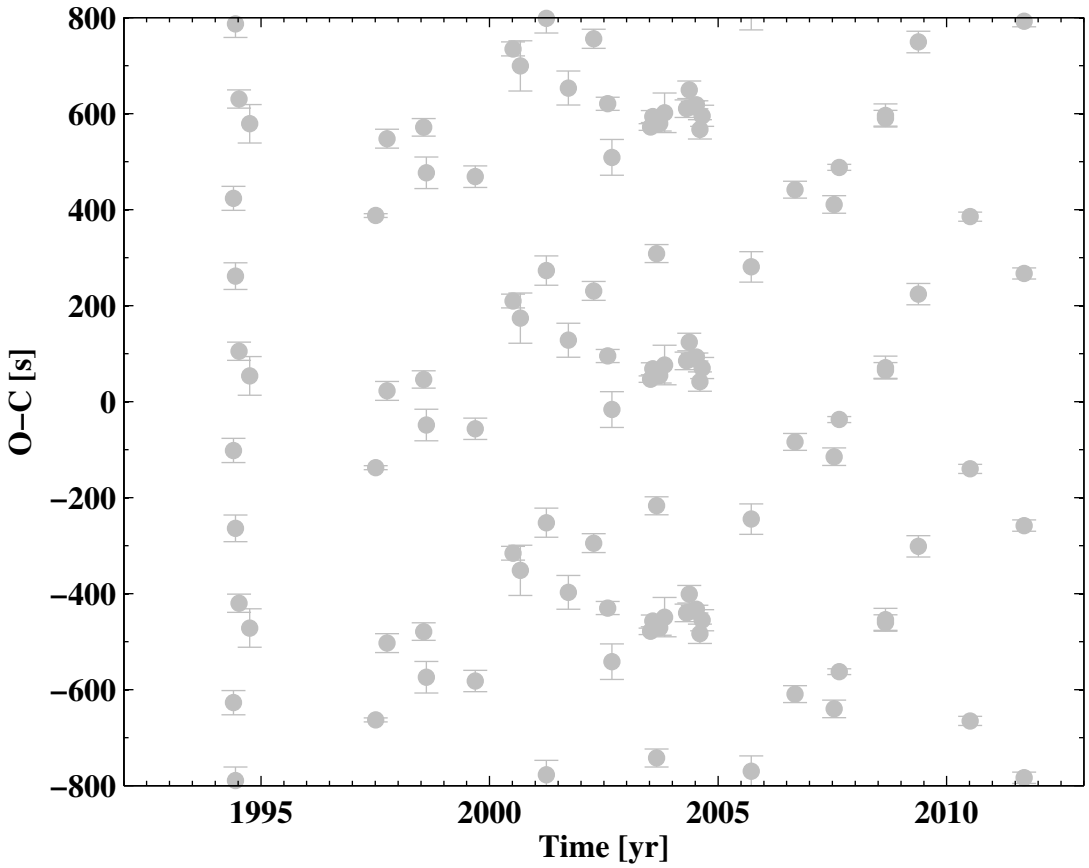

Figure 2.11: $O - C$ of mode I. No unambiguous sequence was identified.

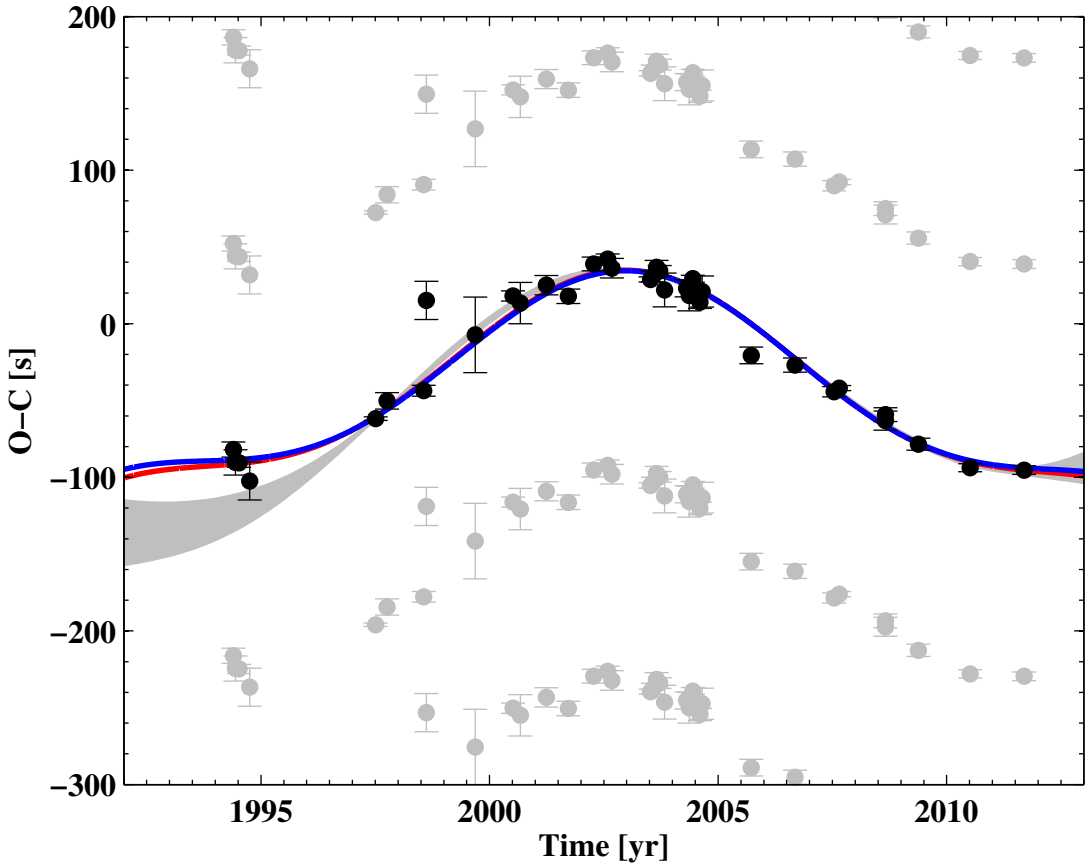

Figure 2.12: $O-C$ of mode $C+E$ with a linear trend removed. The selected sequence is shown in black and other possible values are shown in grey. The red line is the best fit of a 12.93 year sinusoid plus a parabola. The blue line is the best fit of the $O-C$ variations predicted by the parent modes C and E. The grey region is the 1σ prediction of the $1997-2011$ values for a sinusoid plus parabola.

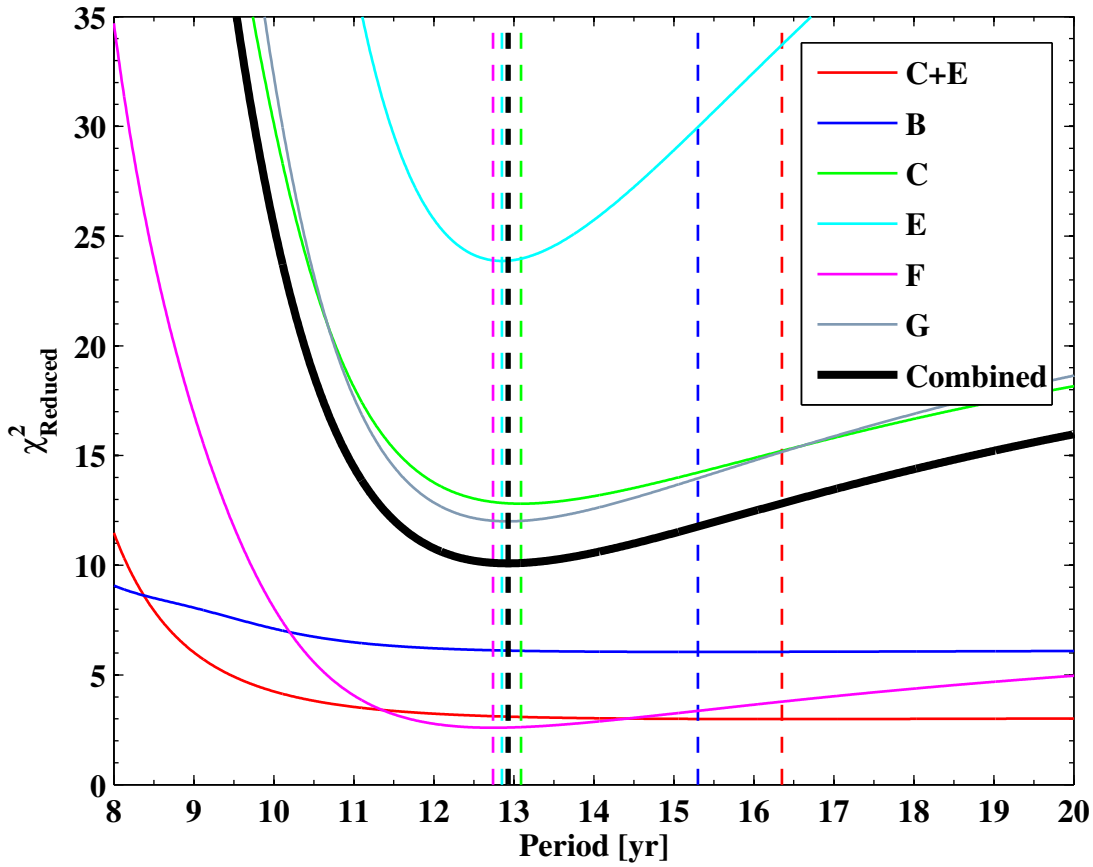

Figure 2.13: Reduced χ^2 as a function of period for fits of a sinusoid plus parabola to $O - C$. The dashed lines indicate minima in reduced χ^2.

parabola. Reduced χ^2 as a function of the period of the sinusoid is plotted in Figure 2.13. For modes C, E, F, and G, the minimum of reduced χ^2 occurs for a period around 13 yr. Reduced χ^2 is not particularly sensitive to the choice of period for modes $C + E$ and B, but is consistent with period around 13 yr. The total minimum of reduced χ^2 occurs for a period of 12.93 yr. Fixing the period to 12.93 yr results in a smaller total reduced χ^2 due to the reduction of free parameters, a sign that the oscillatory behavior of the six $O - C$ sequences can be effectively modeled with the same period. The best fits to each $O - C$ are shown in red in Figures 2.12, 2.6, 2.7, 2.8, 2.9, and 2.10. The fit parameters and are shown in Table 2.5 and plotted in Figure 2.14.

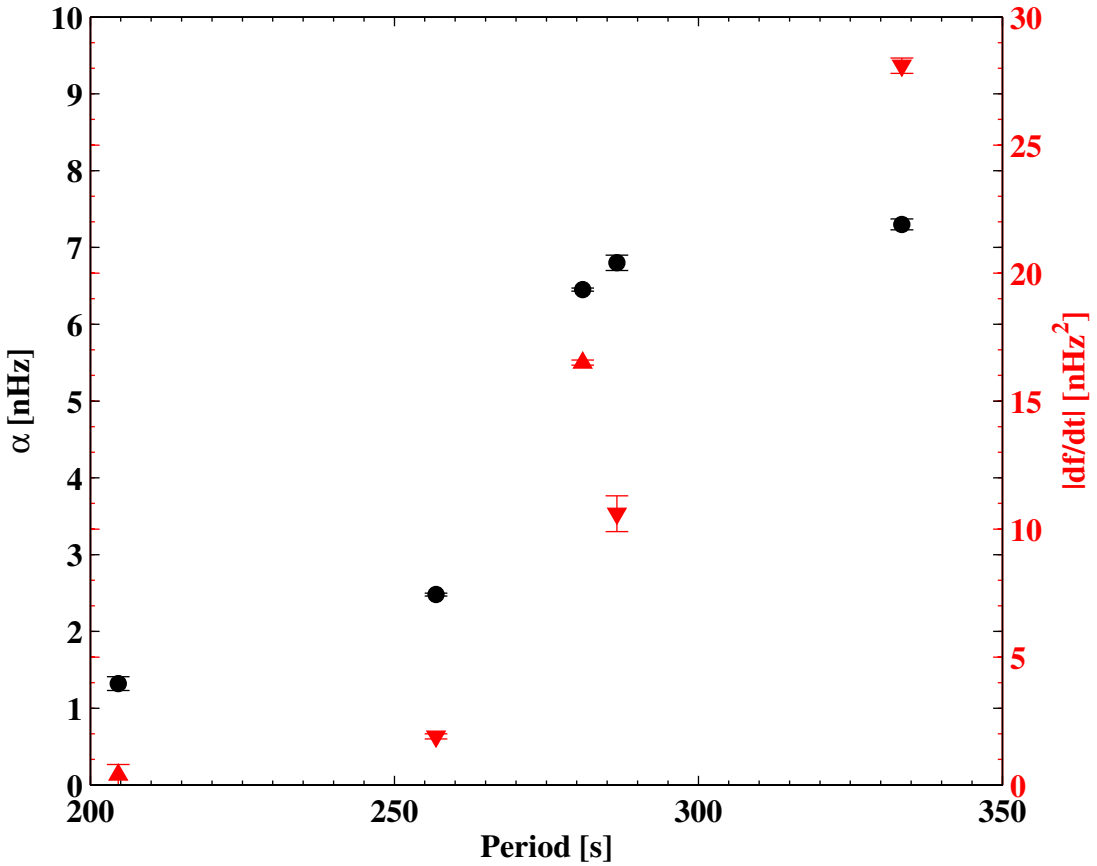

Figure 2.14: Frequency variations of the normal modes. The black points indicate the amplitude of the sinusoidal variations and the red triangles indicate the magnitude of the secular changes. The sign of the secular changes is indicated by the direction of the triangle.

Table 2.5: Results of fitting a 12.93 yr sinusoid plus a parabola to $O - C$. The parameter $\dot{f}$ is the rate of frequency change and the parameters α and ϕ are the amplitude and phase of the periodic variation.

Label	**Frequency** $[\mu Hz]$	$\dot{f}$ $[nHz^2]$	α $[nHz]$	ϕ $[deg]$	$\chi^2_{Reduced}$
$C + E$	7452.24564(0.00003)	15.3(0.5)	4.0(0.1)	200(2)	3.1
B	4887.84568(0.00002)	0.4(0.4)	1.32(0.09)	172(4)	6.1
C	3893.249385(0.000006)	$-1.9(0.1)$	2.48(0.02)	0.0(0.6)	12.9
E	3558.996306(0.000006)	16.5(0.1)	6.45(0.02)	193.1(0.2)	24.0
F	3489.05583(0.00003)	$-10.6(0.7)$	6.8(0.1)	-1(1)	2.6
G	2998.72105(0.00002)	$-28.1(0.3)$	7.30(0.07)	0.7(0.5)	12.1

The average residuals and the Fourier spectrum of the average residuals were calculated and are shown in Figures 2.15 and 2.16. There appears to be a statistically significant peak in the Fourier transform of the residuals at 2.88 yr. This could be the signature of a $> 2M_J$ planet in a circular orbit with a semi major axis of several astronomical unit (AU), but given the poor statistical nature of the fits, we approach this interpretation with a great amount of skepticism. Additionally, removing this signal from the lightcurve creates power at other frequencies in the DFT.

The modeled frequency variations of the parent modes C and E are compared to the modeled frequency variations of the combination mode $C + E$ in Table 2.6. Additionally, the $O - C$ variations predicted for mode $C + E$ by the modeled frequency variations of the parents are shown in blue in Figure 2.12. In both cases, the agreement is outstanding. This is extremely strong evidence that a sinusoid plus a parabola is effectively modeling the behavior of $O - C$ and that the fit parameters are statistically reliable.

2.4 Amplitude Modulation

EC 20058-5234 has two visual companions within several arcseconds (see Figure 2.3). PMT observations include flux from the companions, which will result in smaller

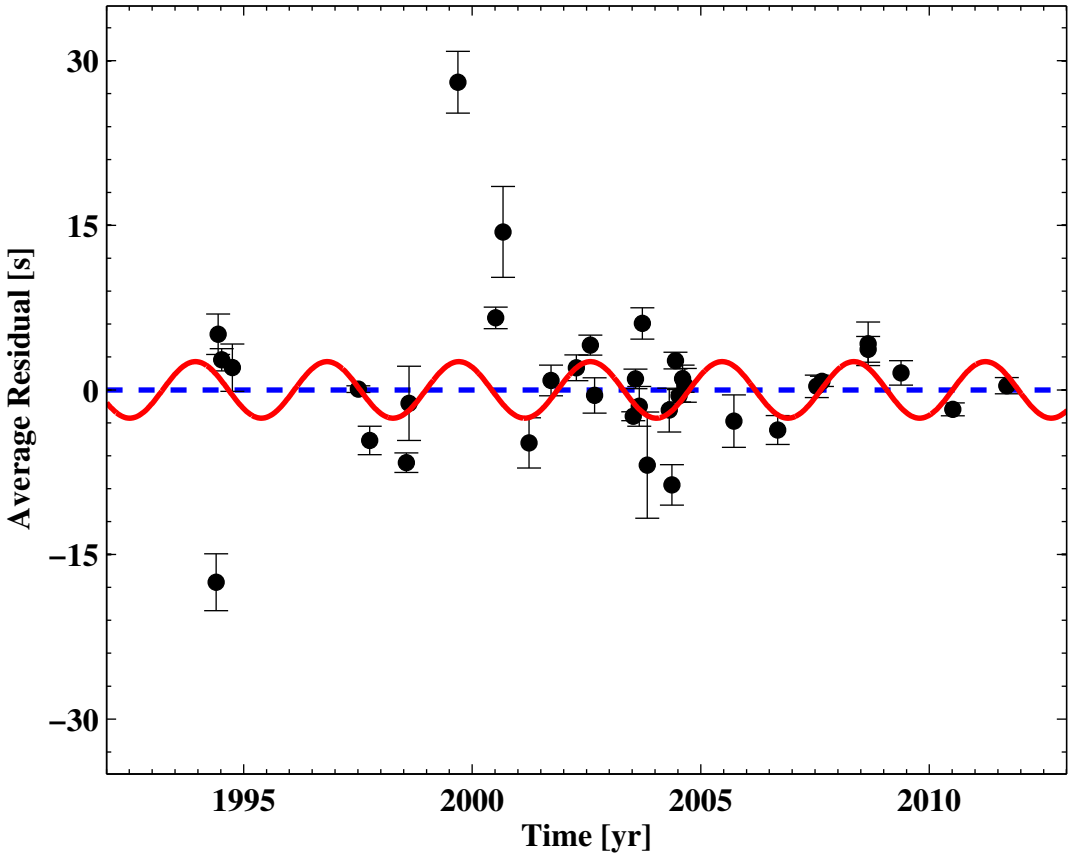

Figure 2.15: Average residuals of the fits to $O - C$. The red curve is the best fit of a sinusoid with a period of 2.88 yr.

Table 2.6: Comparison of the frequency of combination mode $C + E$ to the frequency of the parent modes C and E.

Parameter	Combination	Parents
f $[\mu Hz]$	7452.24564(0.00003)	7452.245691(0.000009)
$\dot{f}$ $[nHz^2]$	15.3(0.5)	14.6(0.2)
α $[nHz]$	4.0(0.1)	4.07(0.03)
ϕ $[deg]$	200(2)	201.0(0.5)

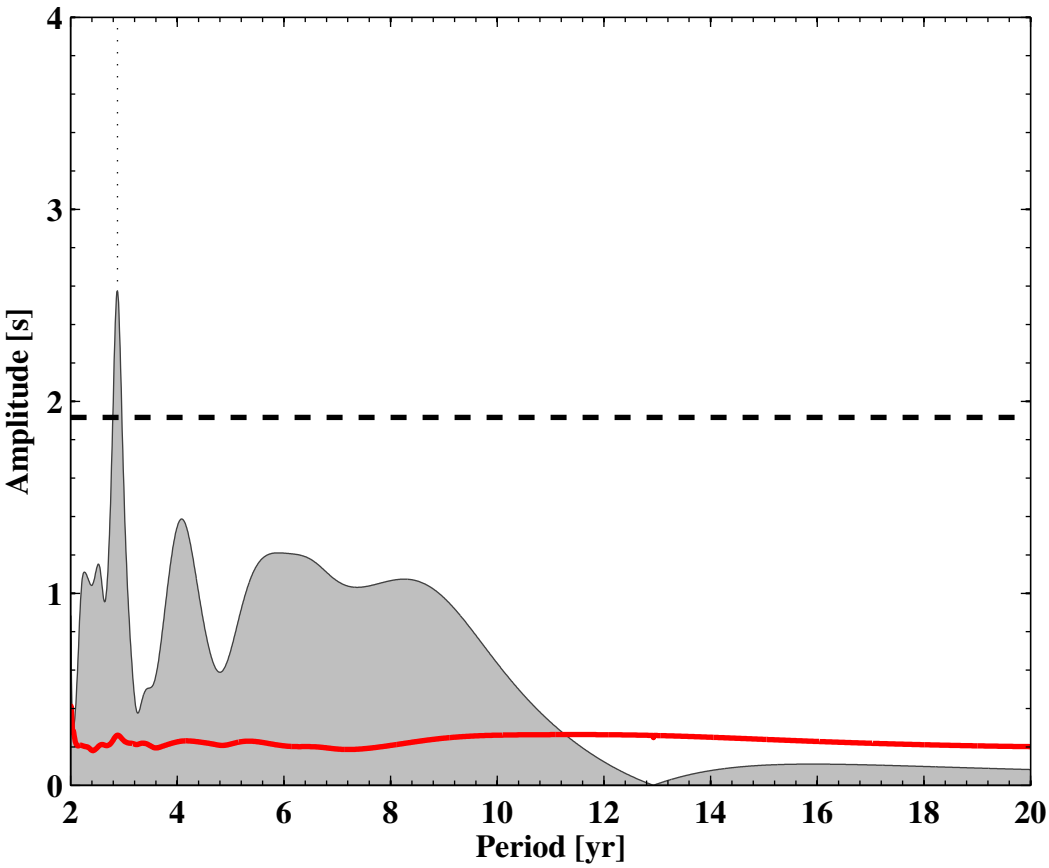

Figure 2.16: Fourier spectrum of the average residuals of the fits to $O - C$. The red line indicates the formal error and the black dashed line is 6 times the median amplitude. The dotted line indicates a peak at 2.88 yr.

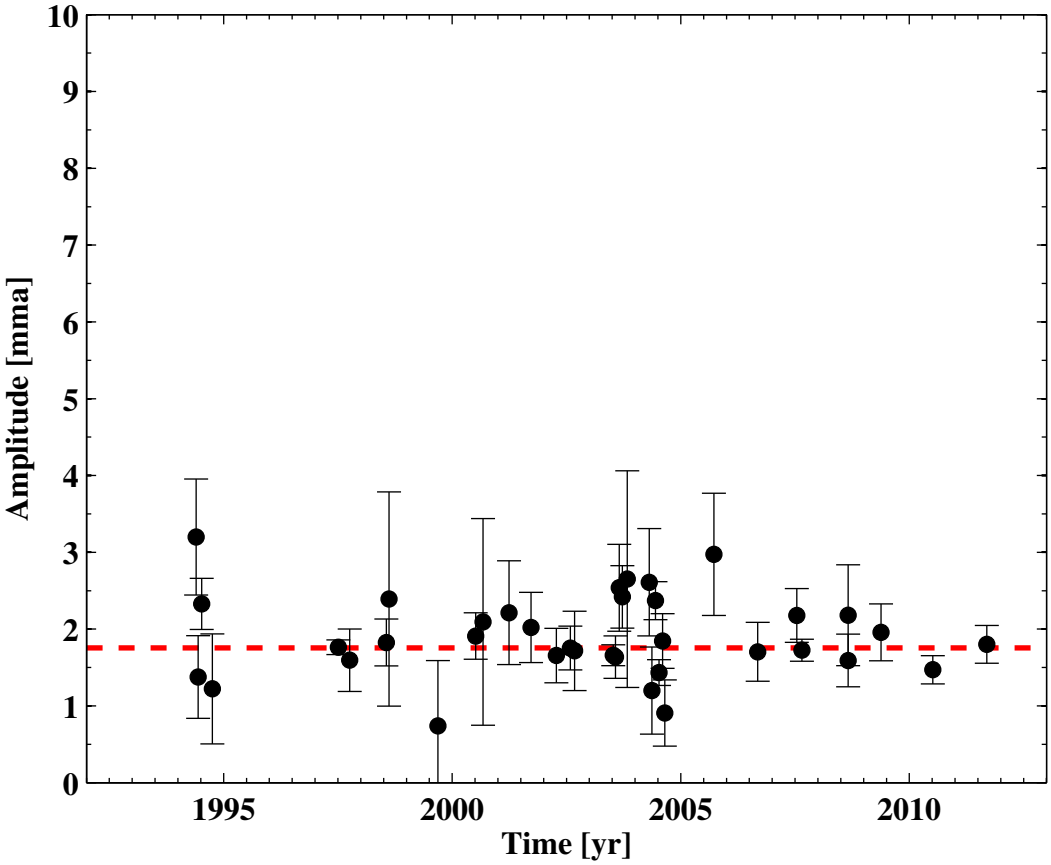

Figure 2.17: Amplitude of mode $C + E$. The dashed line is the weighted mean.

observed mode amplitudes. Because the amount of flux included is seeing dependent, this effect is not trivial to remove. Additionally, the CCD observations were not reduced with point spread function (PSF) photometry methods, so even the amplitudes extracted from CCD data are seeing dependent. To circumvent this problem, the set of measured amplitudes of mode C were normalized to have an average value of unity. All other amplitude measurements were divided by the corresponding normalized amplitude of mode C. The amplitudes are shown as function of time is shown in Figures 2.17-2.24.

There do not appear to be any notable variations in amplitude but for some of

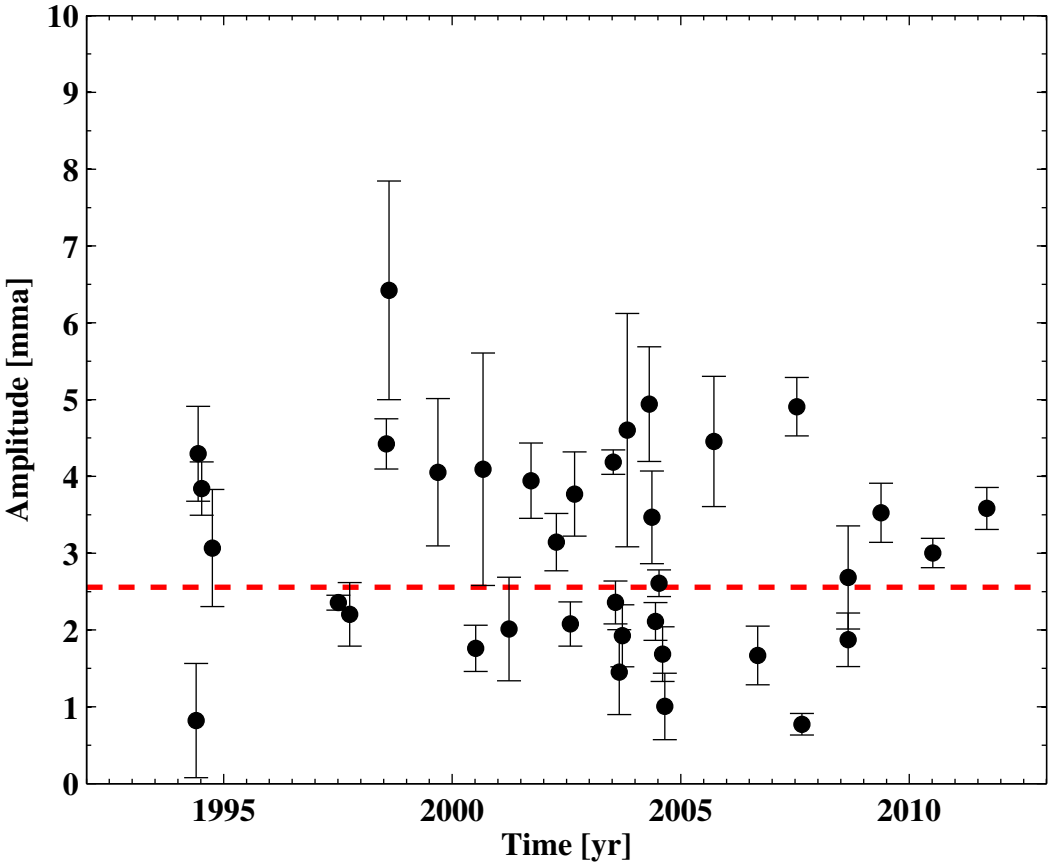

Figure 2.18: Amplitude of mode A. The dashed line is the weighted mean.

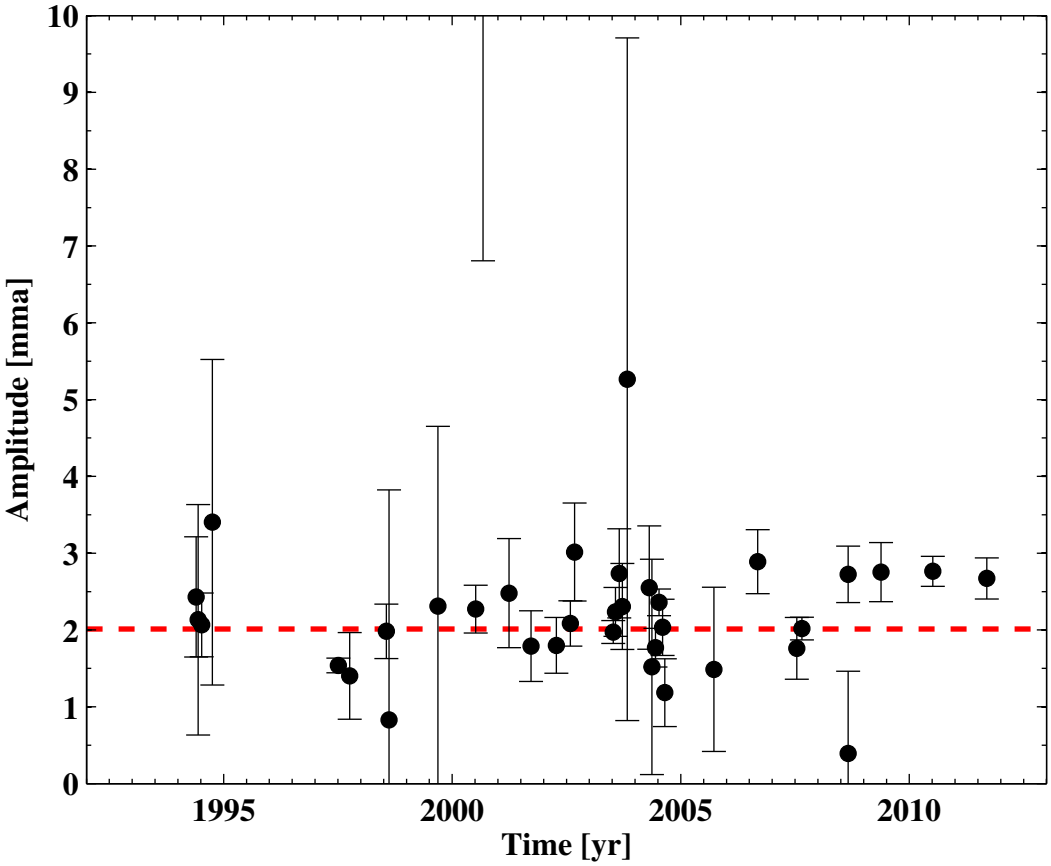

Figure 2.19: Amplitude of mode $G + I$. The dashed line is the weighted mean.

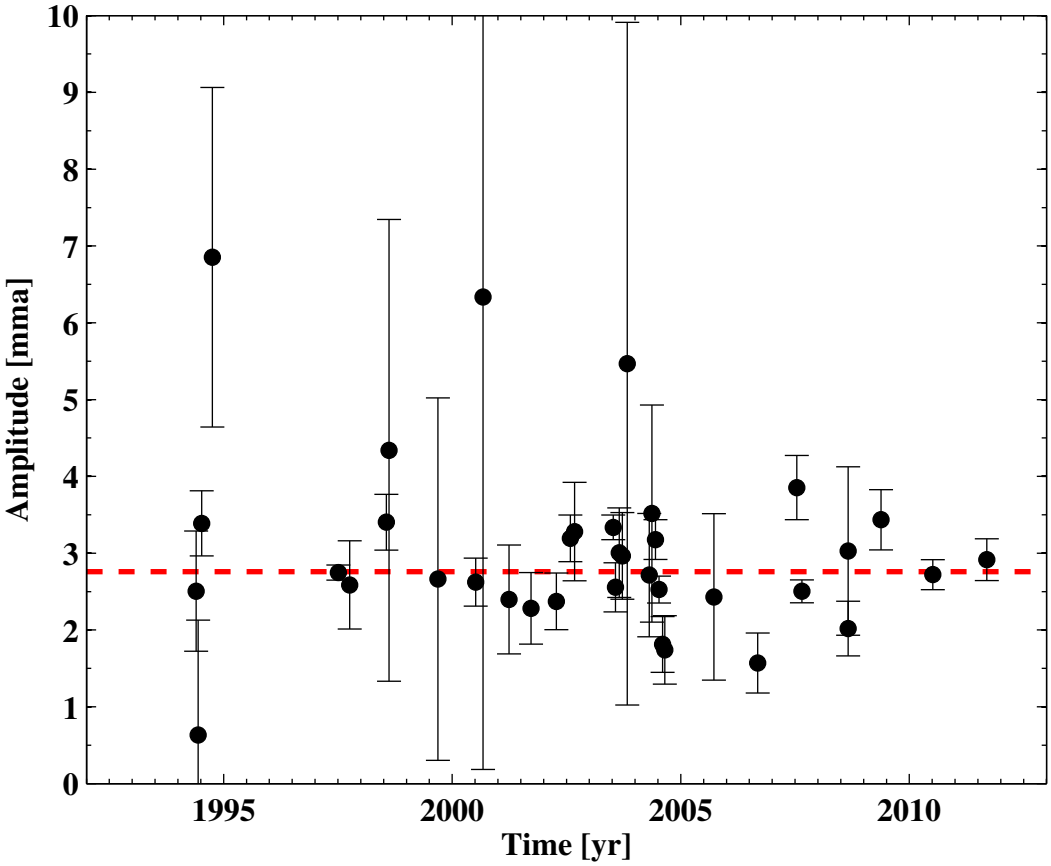

Figure 2.20: Amplitude of mode B. The dashed line is the weighted mean.

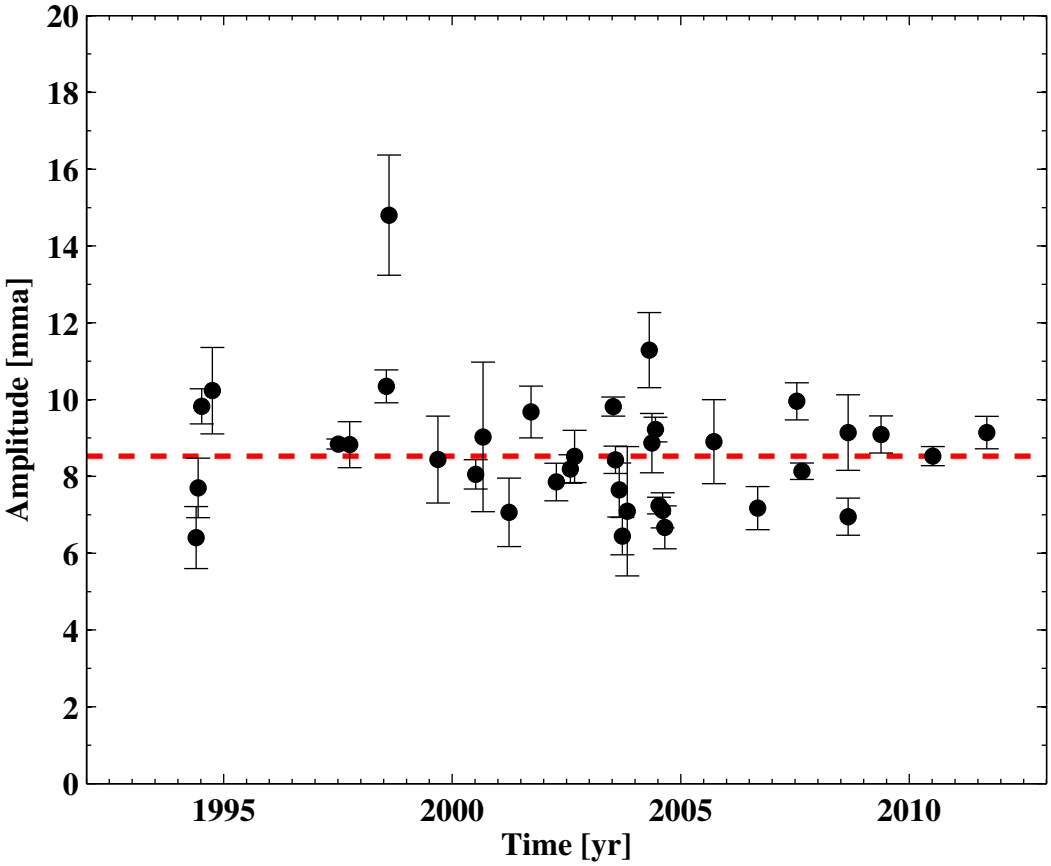

Figure 2.21: Amplitude of mode E. The dashed line is the weighted mean.

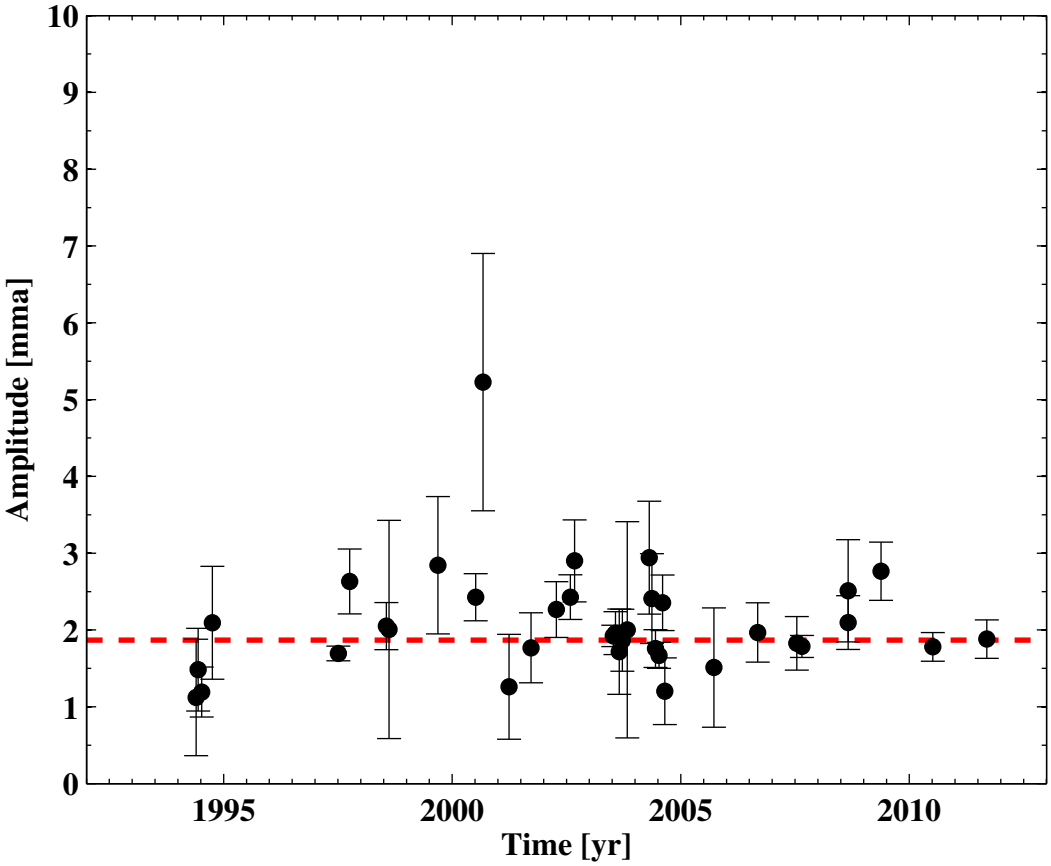

Figure 2.22: Amplitude of mode F. The dashed line is the weighted mean.

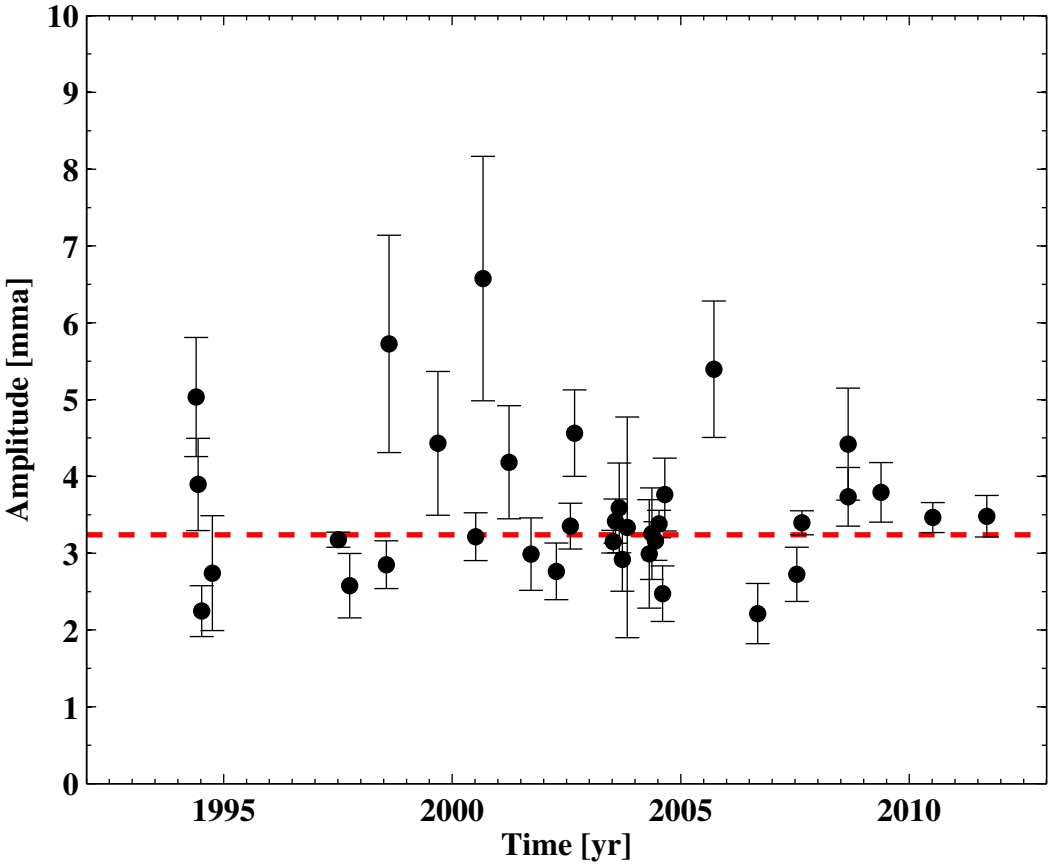

Figure 2.23: Amplitude of mode G. The dashed line is the weighted mean.

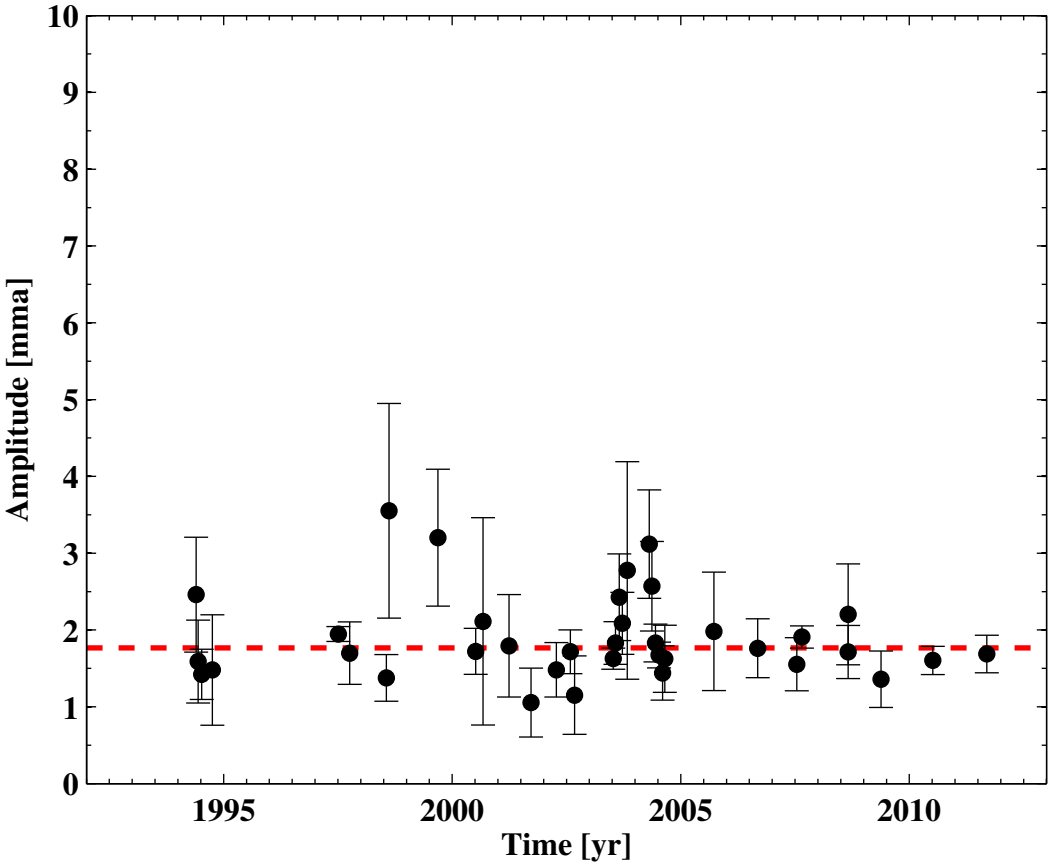

Figure 2.24: Amplitude of mode I. The dashed line is the weighted mean.

Table 2.7: Upper limits on the amplitude of sinusoidal variations in mode amplitude (α) and reduced χ^2 for the model of constant amplitude.

Label	**Period** $[s]$	α $[mma]$	$\chi^2_{Reduced}$
$C+E$	134	$< 0.07(0.07)$	1.1
A	195	$< 0.79(0.08)$	14.0
$G+I$	204	$< 0.40(0.08)$	2.3
B	205	$< 0.3(0.1)$	2.1
E	281	$< 0.5(0.1)$	6.2
F	287	$< 0.23(0.08)$	1.4
G	333	$< 0.19(0.08)$	2.0
I	525	$< 0.18(0.08)$	0.9

the modes the relative amplitudes are certainly not statistically consistent with being constant in time. A constant model was fit to the amplitudes and upper limits on the amplitude of a sinusoidal variation in mode amplitude were calculated from the residuals. The results are shown in Table 2.7.

Chapter 3

KIC 8626021

3.1 Observations and Data Preparation

KIC 8626021 was discovered to pulsate in late 2010 by Østensen et al. (2011) and is currently the only known Helium atmosphere white dwarf (DB) in the *Kepler* field of view (FOV). Bischoff-Kim and Østensen (2011) shows despite initial spectroscopic indications that KIC 8626021 is a cool V777 Her star (DBV), KIC 8626021 is actually a hot DBV and should cool primarily due to plasmon neutrino production. This hotter temperature was subsequently confirmed by the independent analysis of Córsico et al. (2012). Primarily due to the possibility of measuring the neutrino production rate in the core of a white dwarf, but also because of its uniqueness as the only DBV in the *Kepler* FOV, KIC 8626021 has occupied one of the competitive short cadence *Kepler* slots from June 2011 to the present.

Kepler observations are divided into "*Kepler* operational quarters" because the spacecraft is rotated 4 times per year. However, the data is actually downloaded from the spacecraft once per month so lightcurves are naturally divided into month long segments. Table 3.1 lists the currently available month long *Kepler* short cadence lightcurves of KIC 8626021. Each lightcurve was initially processed by the *Kepler* pipeline as described in Jenkins et al. (2010). A median cutoff filter was then applied to the lightcurves, which were subsequently divided by a low order polynomial. An hour of the 2010 observations are shown in Figure 3.1. Note that in Figure 3.1 the pulsations are not obvious due to the comparably long duty cycle and the relatively low signal to noise of the individual points.

Table 3.1: Journal of Observations

Start	End	Quarter	Observations
2010 Oct 24	2010 Nov 22	7.2	42701
2011 Jun 27	2011 Jul 27	10.1	44077
2011 Jul 28	2011 Aug 28	10.2	46028
2011 Aug 29	2011 Sep 28	10.3	44191
2011 Sep 29	2011 Oct 30	11.1	45451
2011 Oct 31	2011 Nov 30	11.2	44170
2011 Dec 01	2012 Jan 04	11.3	46669
2012 Jan 05	2012 Feb 01	12.1	38674
2012 Feb 02	2012 Feb 29	12.2	40067
2012 Mar 01	2012 Mar 28	12.3	38190
2012 Mar 29	2012 Apr 30	13.1	47220
2012 May 01	2012 May 30	13.2	42793
2012 May 31	2012 Jun 27	13.3	39876
2012 Jun 28	2012 Jul 29	14.1	36369
2012 Jul 30	2012 Aug 29	14.2	44157
2012 Aug 30	2012 Oct 03	14.3	50016

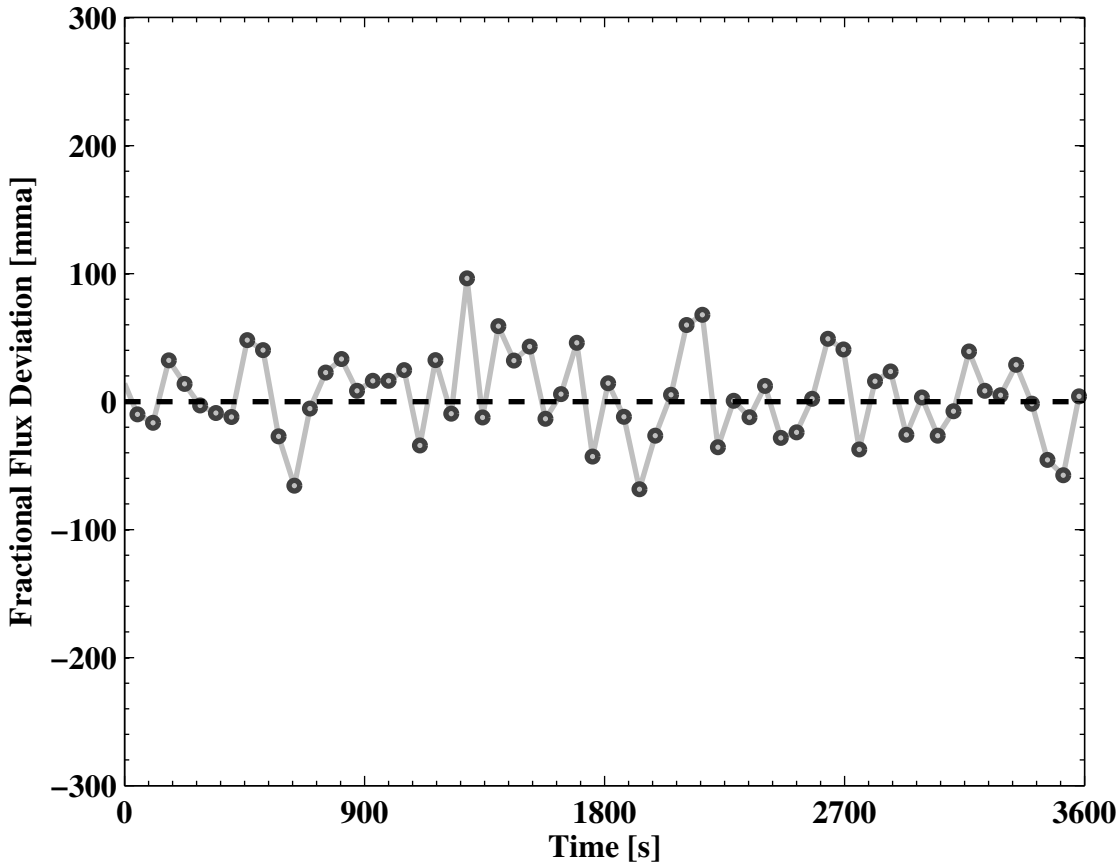

Figure 3.1: A portion of the Q7.2 lightcurve.

3.2 Fourier Analysis

All observations except the Q7.2 lightcurve were combined into a single nearly continuous lightcurve that spans more than a year. The Fourier spectrum of the combined lightcurve is shown in Figure 3.2. Compared to even the best WET observations of pulsating white dwarfs, the spectral window (see Figure 3.3) and frequency resolution are phenomenal. There is significant power at many frequencies although in terms of other known white dwarf pulsators, the amplitudes are notably small (< 5 mma). There are three groups of three closely spaced modes, i.e. "triplets", and several other modes in the range of $100 - 400$ s. There are also peaks very close to the $7^{\rm th}$, $8^{\rm th}$, $9^{\rm th}$, $10^{\rm th}$, and $14^{\rm th}$ harmonics of the *Kepler* 29.4244 min long-cadence cycle time. Figures 3.4, 3.5, and 3.6 show a small range of the Fourier spectrum around the three closely spaced triplets.

The remarkable length of the combined lightcurve results in an extremely narrow spectral window. This allows identification of modes without the need for iterative pre-whitening. The frequencies of each maximum in the Fourier spectrum that are at least 6 times the median amplitude and also the highest local maximum within $1\,\mu$Hz were identified. A sinusoid at each frequency was fit to the combined lightcurve simultaneously. The results of the fit and the labeling scheme are shown in Table 3.2. No harmonic modes and one combination mode was found. Note that the frequency of the modes labeled 7K, 8K, 9K, 10K, and 14K are within $1\,\mu$Hz of a harmonic of the *Kepler* long cadence cycle time. These modes are attributed to instrumental effects and do not merit further consideration.

3.3 $O - C$ Analysis

The $O - C$ diagrams of each mode are shown in Figures 3.7-3.19. For all modes, there is an unambiguous sequence in $O - C$ for the effectively continuous Q10-14 lightcurves, but for several modes there is ambiguity as to the appropriate Q7.2 $O - C$ value.

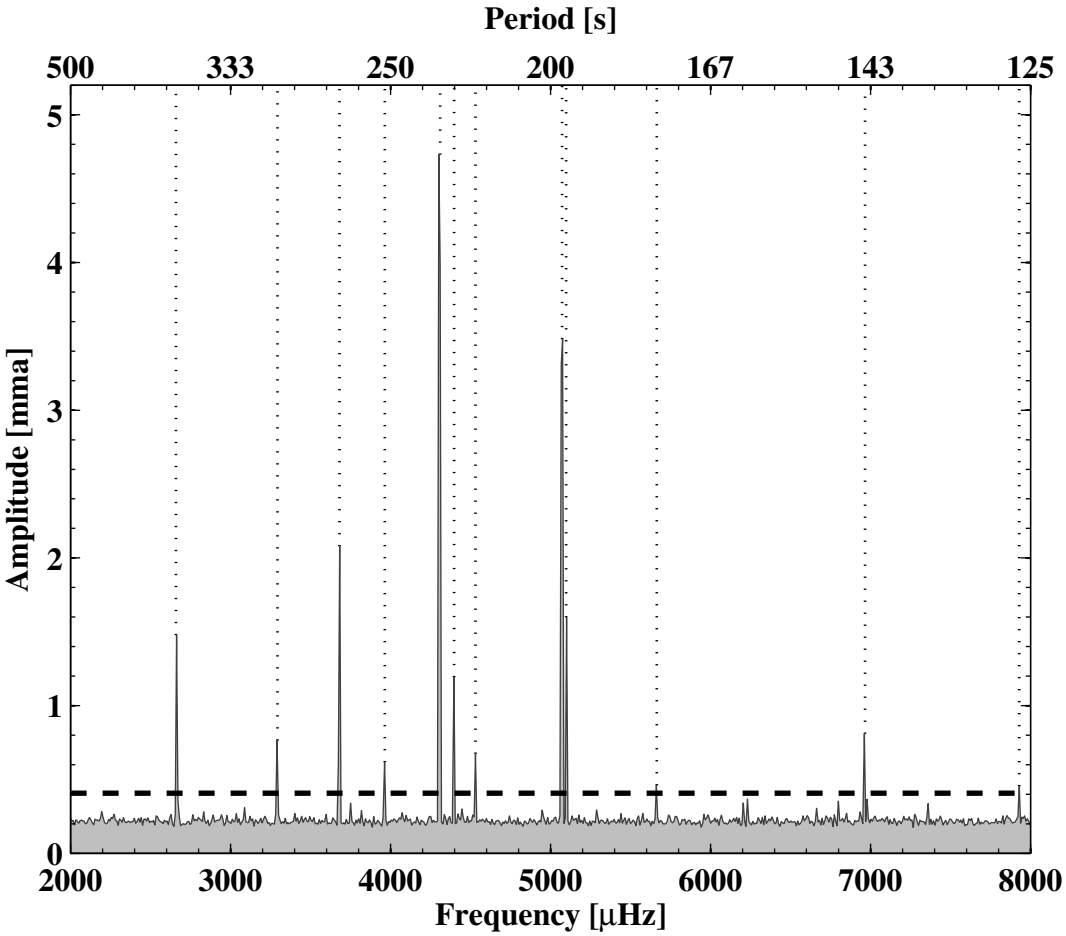

Figure 3.2: Fourier spectrum of the combined Q10-14 lightcurve. The dotted lines indicate frequencies listed in Table 3.2 and the dashed line indicates 6 times the median amplitude.

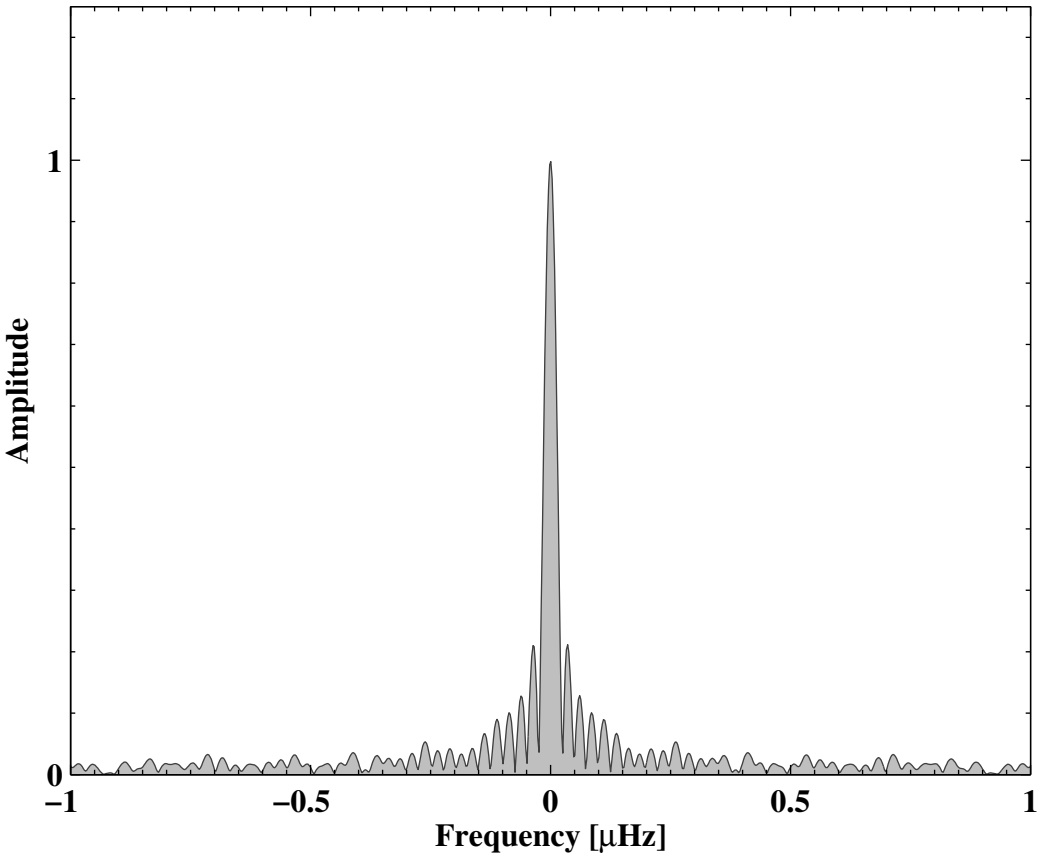

Figure 3.3: Spectral window of the combined Q10-14 lightcurve.

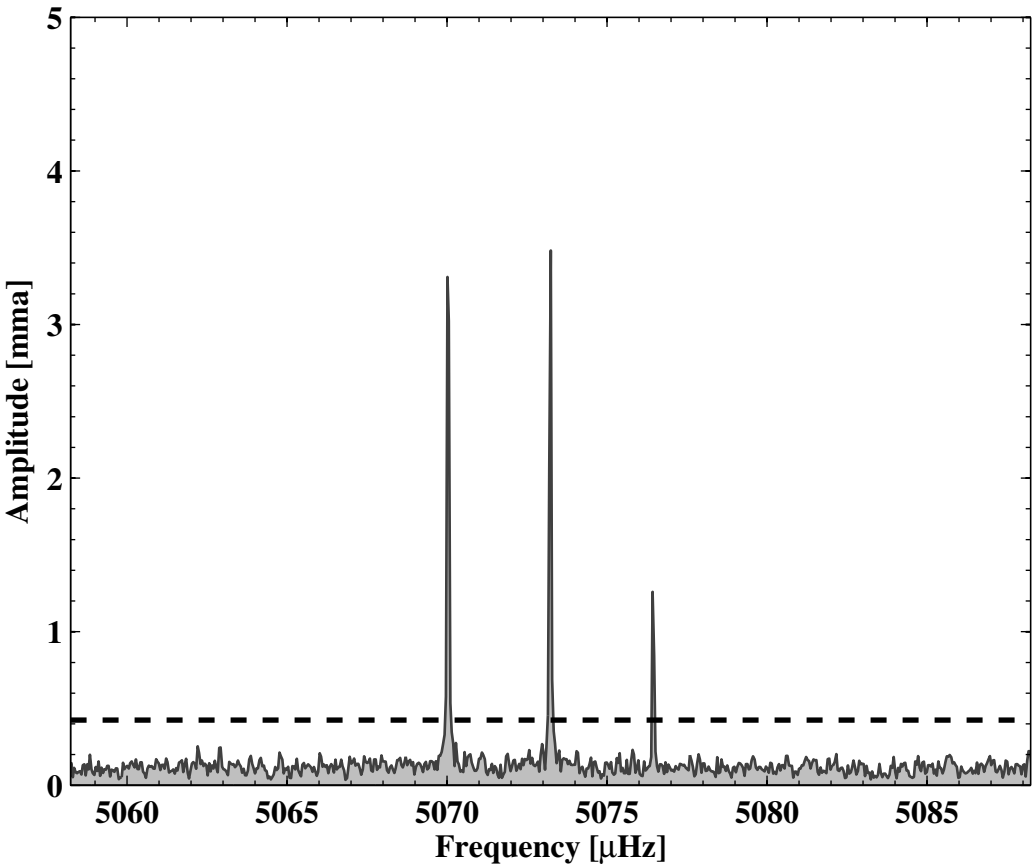

Figure 3.4: Region of the Fourier spectrum of the combined Q10-14 lightcurve that includes modes A^-, A, and A^+. The dashed line is 6 times the median amplitude.

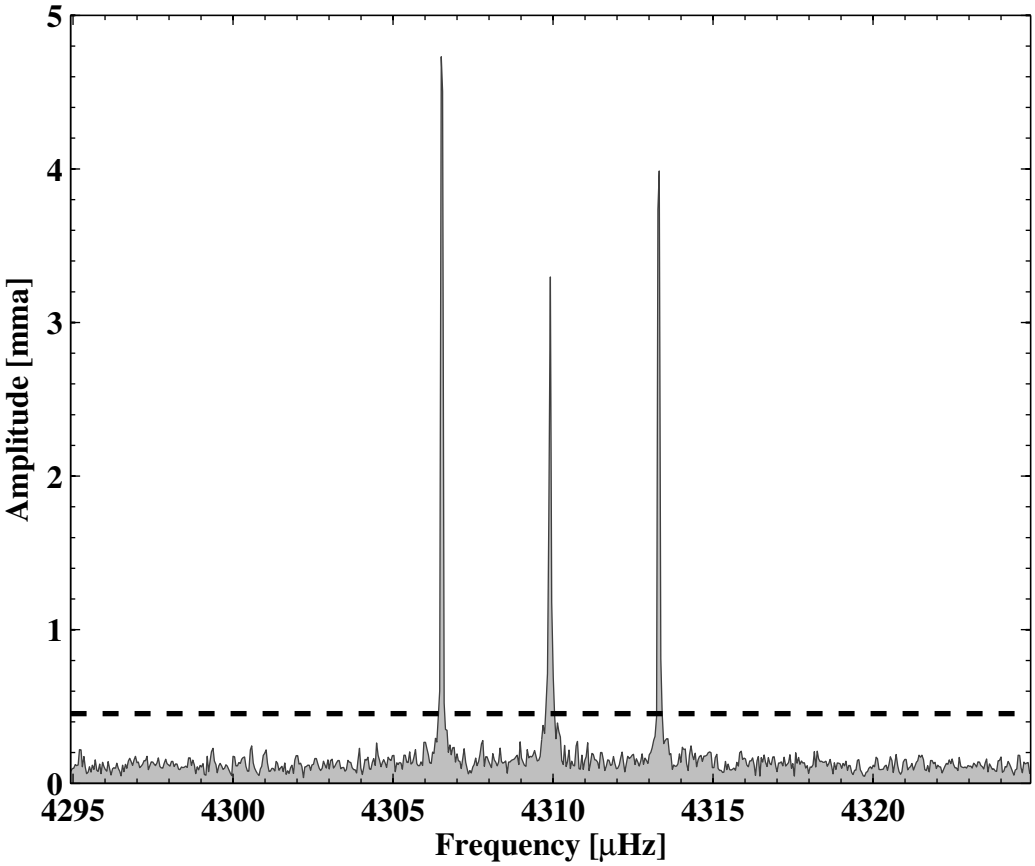

Figure 3.5: Region of the Fourier spectrum of the combined Q10-14 lightcurve that includes modes C^-, C, and C^+. The dashed line is 6 times the median amplitude.

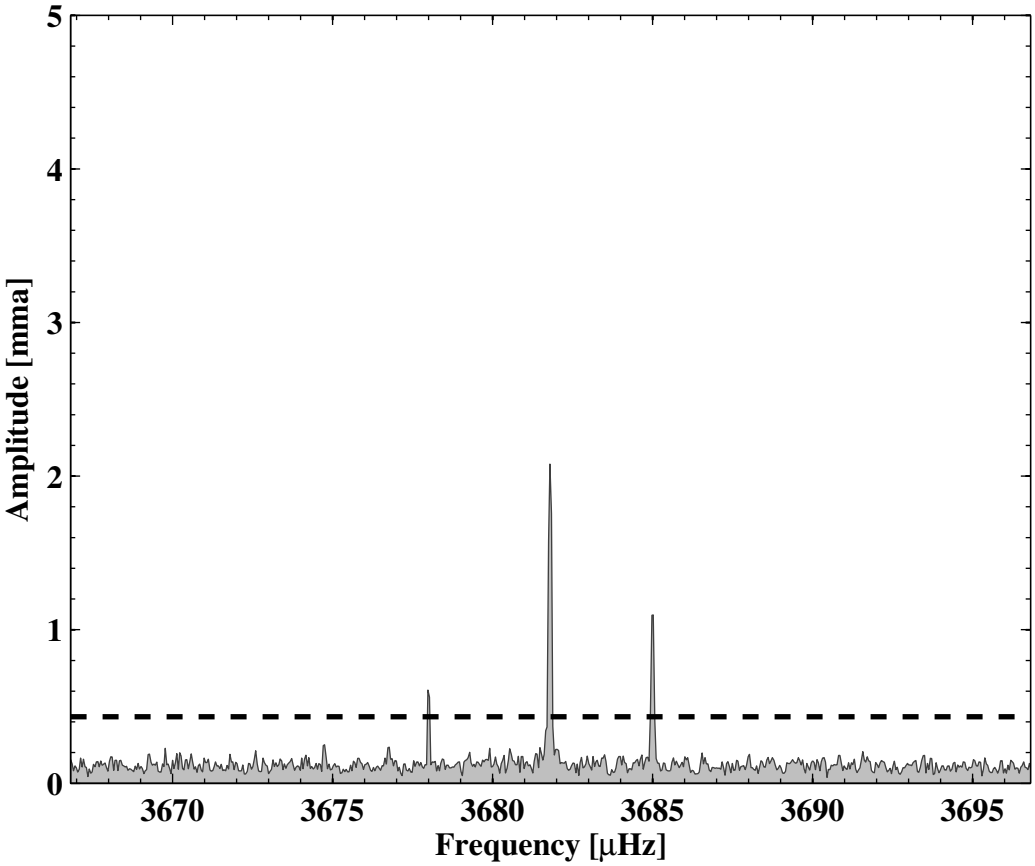

Figure 3.6: Region of the Fourier spectrum of the combined Q10-14 lightcurve that includes modes D^-, D, and D^+. The dashed line is 6 times the median amplitude.

Table 3.2: Frequencies and amplitudes of modes identified in the combined Q10-14 lightcurve.

Label	**Period** $[s]$	**Frequency** $[\mu Hz]$	**Amplitude** $[mma]$
$14K$	126.1	7930.055(0.002)	0.46(0.06)
$C^+ + F$	143.6	6965.292(0.001)	0.81(0.06)
$10K$	176.6	5664.066(0.002)	0.46(0.06)
$9K$	196.2	5097.7949(0.0005)	1.60(0.06)
A^-	197.0	5076.4360(0.0006)	1.25(0.06)
A	197.1	5073.2350(0.0002)	3.46(0.06)
A^+	197.2	5070.0305(0.0002)	3.31(0.06)
$8K$	220.7	4531.438(0.001)	0.68(0.06)
B	227.4	4398.3747(0.0007)	1.19(0.06)
C^-	231.8	4313.3027(0.0002)	3.99(0.06)
C	232.0	4309.9175(0.0002)	3.30(0.06)
C^+	232.2	4306.5207(0.0002)	4.74(0.06)
$7K$	252.2	3964.981(0.001)	0.62(0.06)
D^-	271.4	3685.0081(0.0007)	1.10(0.06)
D	271.6	3681.8066(0.0004)	2.08(0.06)
D^+	271.9	3677.996(0.001)	0.63(0.06)
E	303.5	3294.372(0.001)	0.77(0.06)
F	376.1	2658.7705(0.0005)	1.48(0.06)

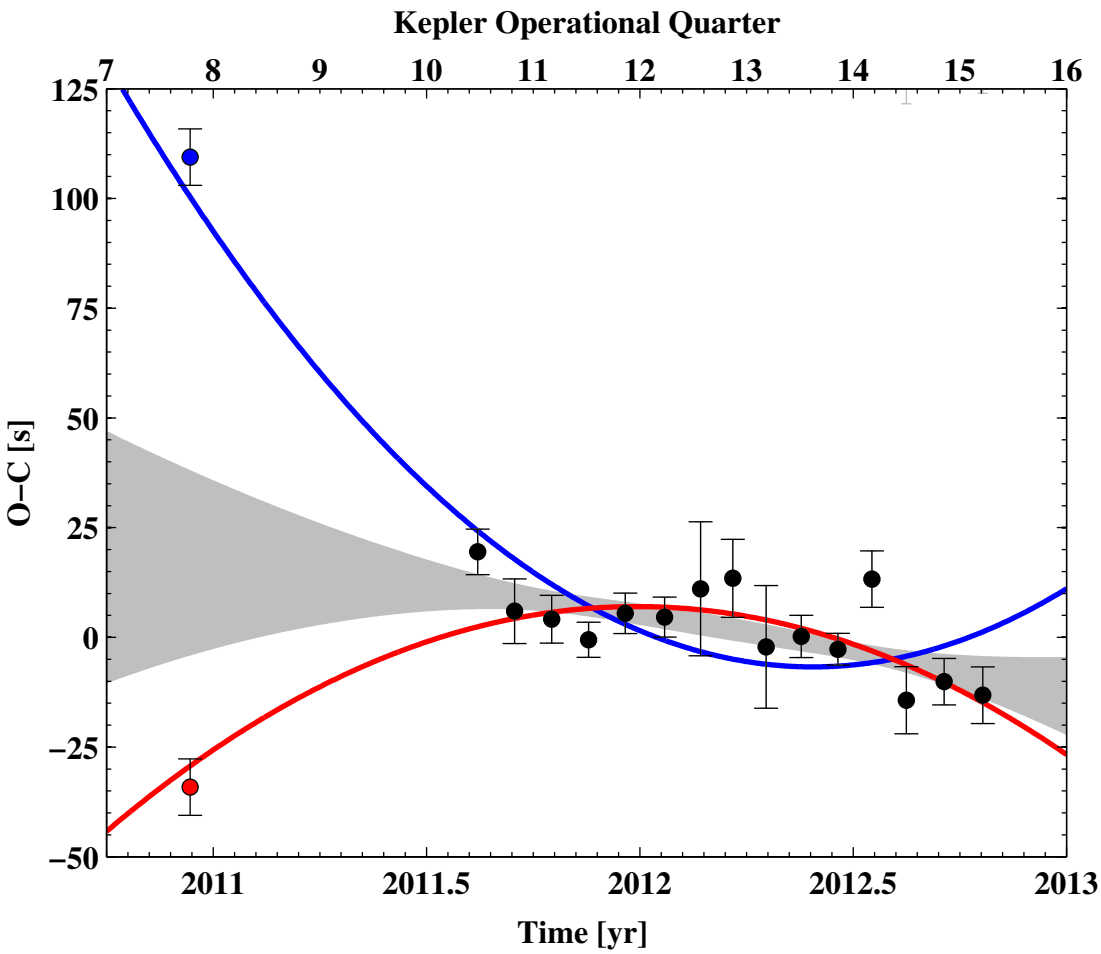

Figure 3.7: $O - C$ of mode $C^+ + F$ with a linear trend removed. The red and blue colored points are possible Q7.2 $O - C$ values. The red and blue lines are the best fit of a parabola given the selection of the correspondingly colored Q7.2 $O - C$ value. The grey region is the $1\,\sigma$prediction made by the Q10-14 values and a parabola.

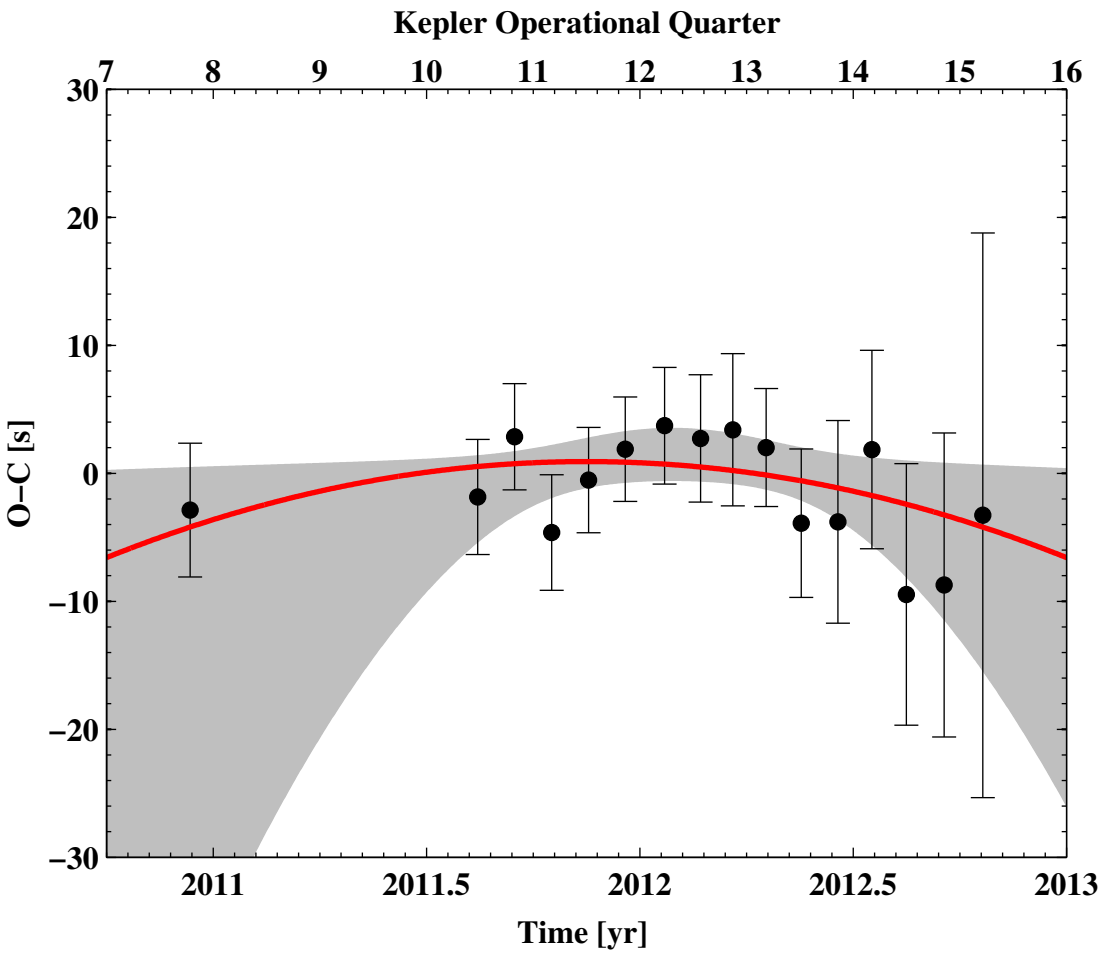

Figure 3.8: $O - C$ of mode A^- with a linear trend removed. The red line is the best fit of a parabola and the grey region is the $1\,\sigma$ prediction made by the Q10-14 values and a parabola.

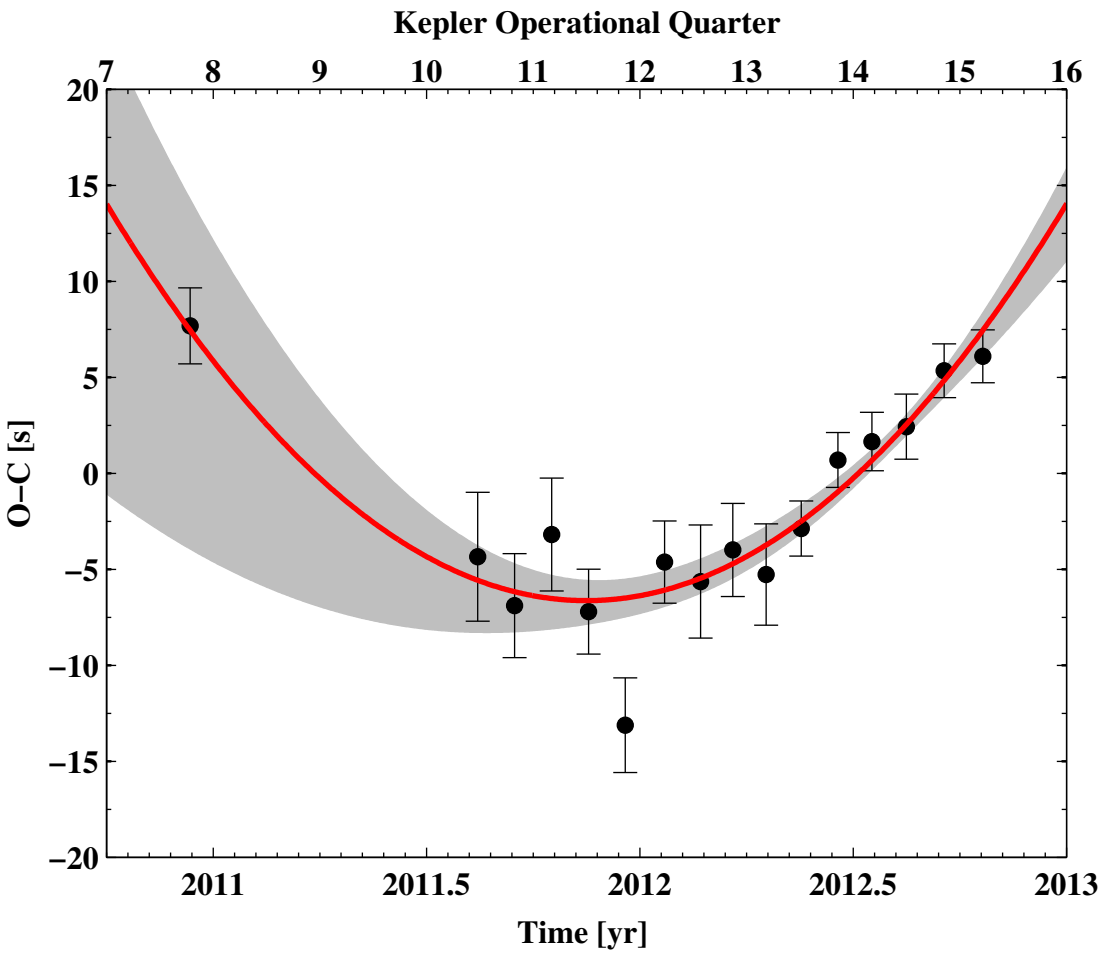

Figure 3.9: $O - C$ of mode A with a linear trend removed. The red line is the best fit of a parabola and the grey region is the $1\,\sigma$ prediction made by the Q10-14 values and a parabola.

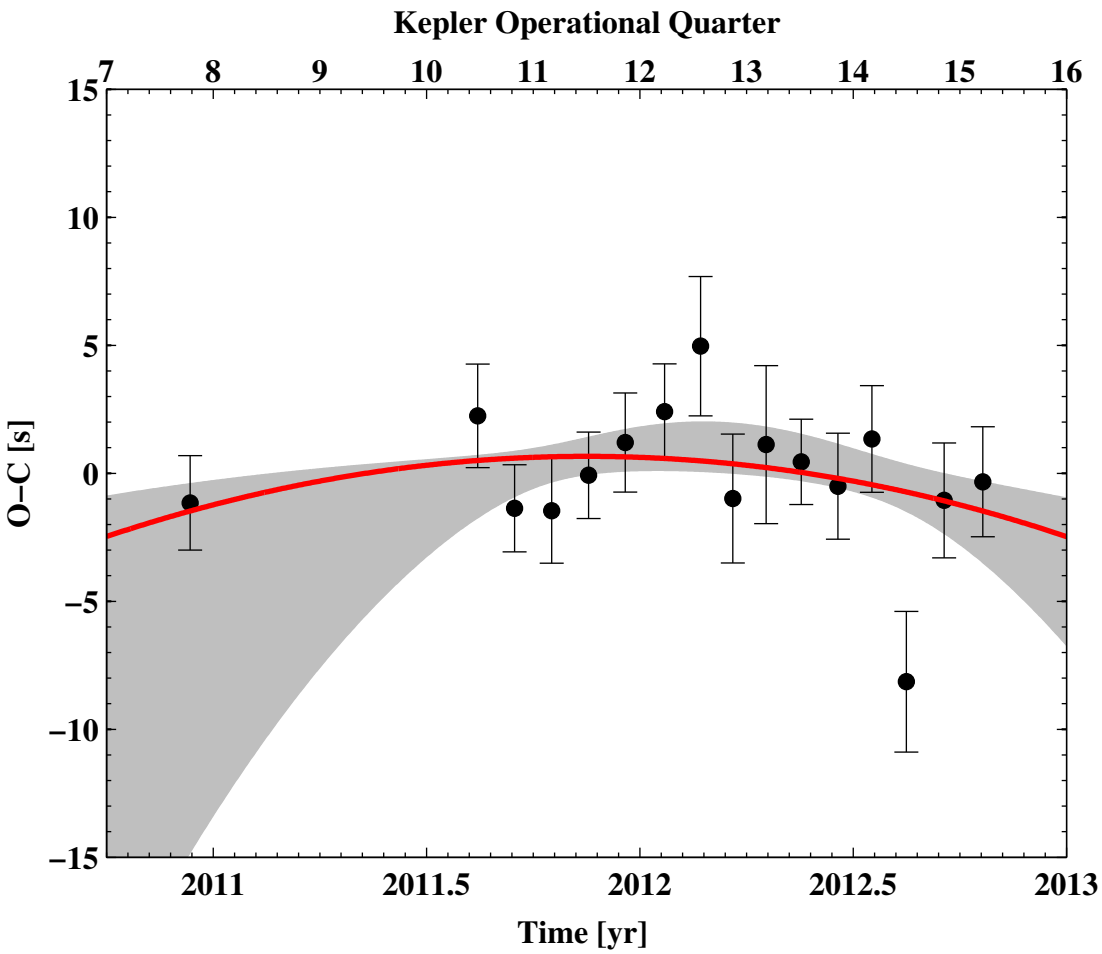

Figure 3.10: $O-C$ of mode A^+ with a linear trend removed. The red line is the best fit of a parabola and the grey shaded region is the $1\,\sigma$ prediction made by the Q10-14 values and a parabola.

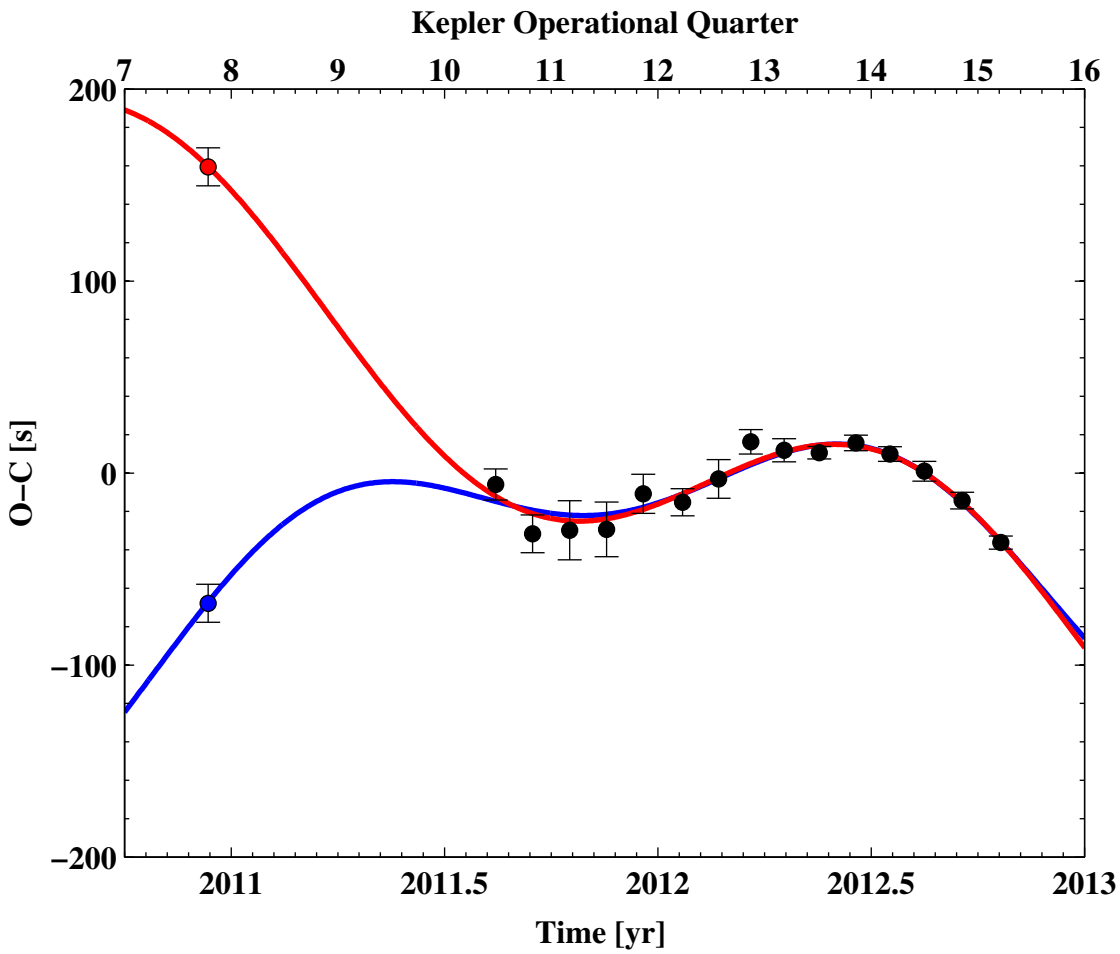

Figure 3.11: $O-C$ of mode B with a linear trend removed. The red and blue points are possible Q7.2 $O-C$ values. The red and blue lines are the best fit of a sinusoid plus a parabola given the selection of the correspondingly colored Q7.2 $O-C$ value.

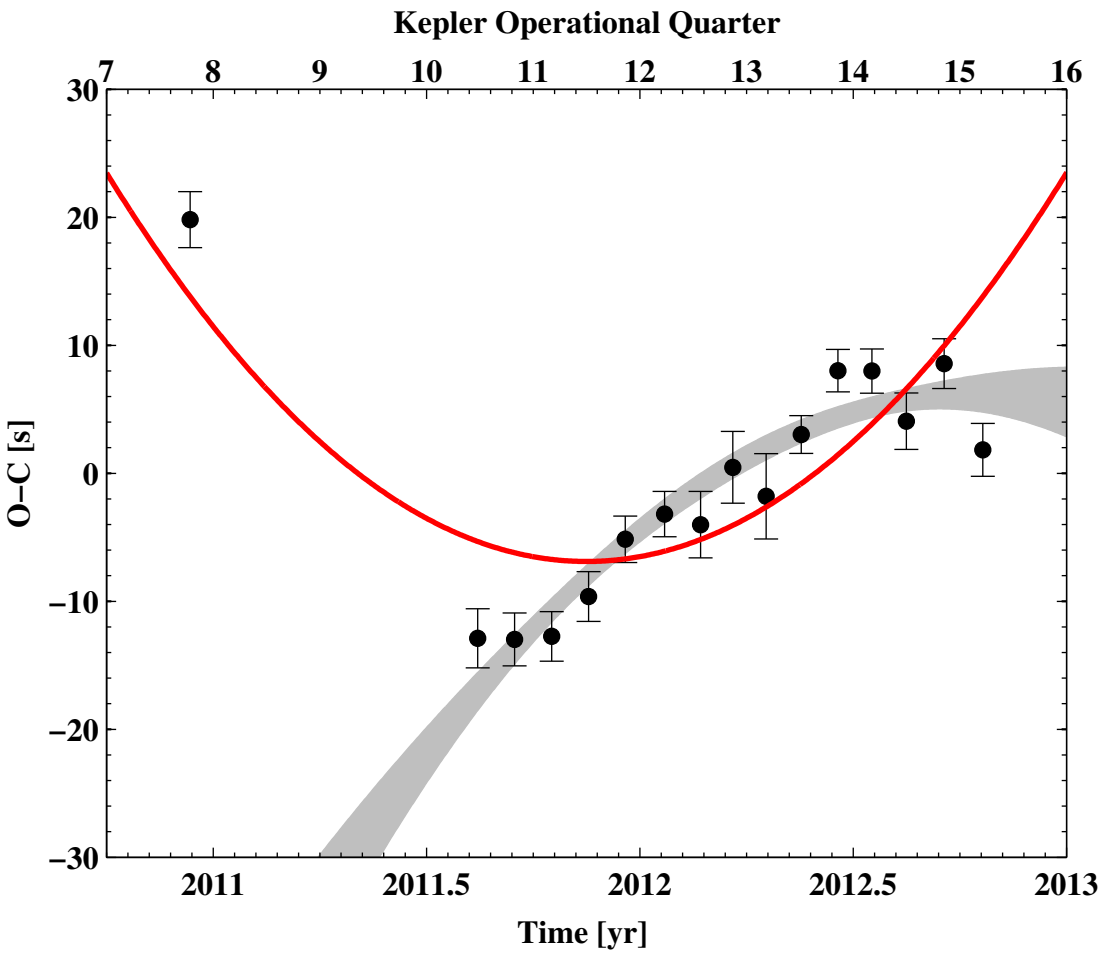

Figure 3.12: $O - C$ of mode C^- with a linear trend removed. The red line is the best fit of a parabola and the grey region is the $1\,\sigma$prediction made by the Q10-14 values and a parabola.

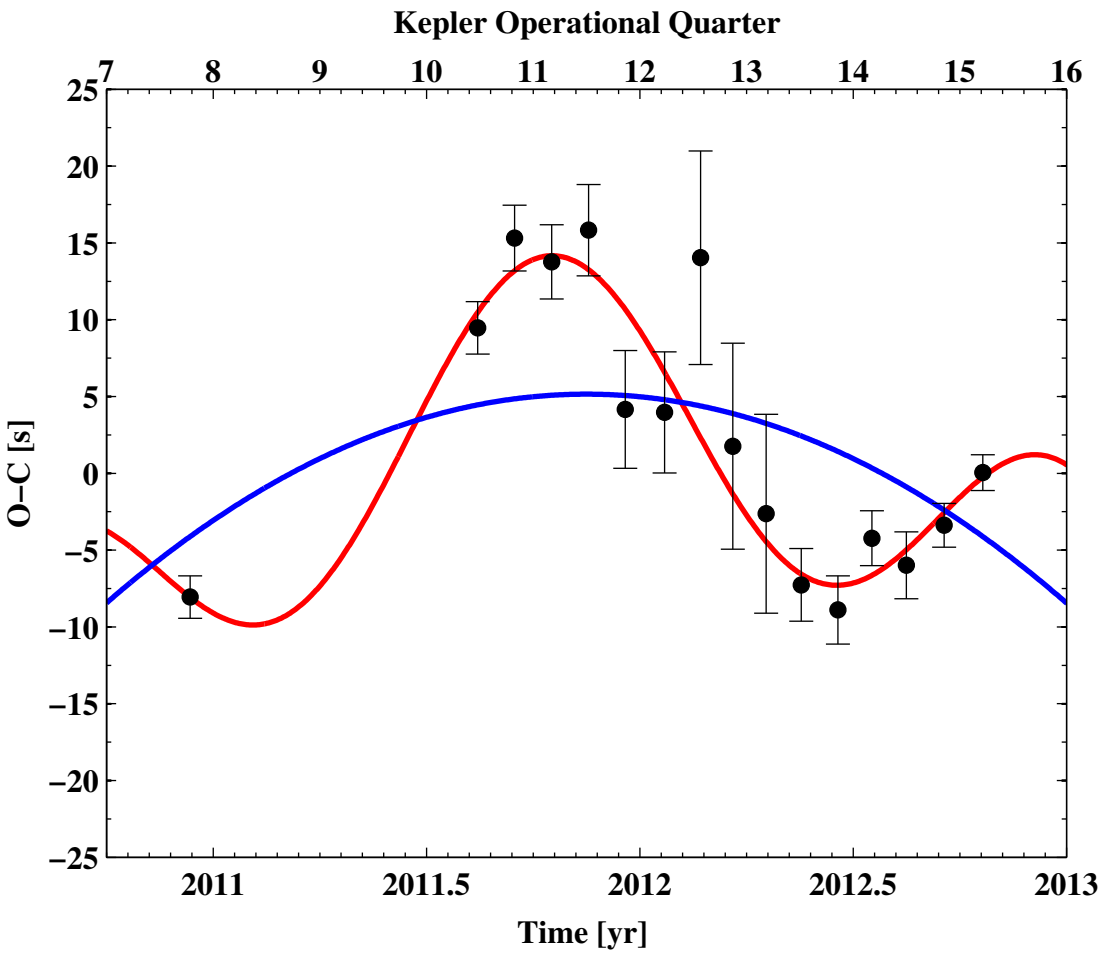

Figure 3.13: $O - C$ of mode C with a linear trend removed. The red line is the best fit of a sinusoid plus a parabola and the blue line is the best fit with the sinusoidal component removed.

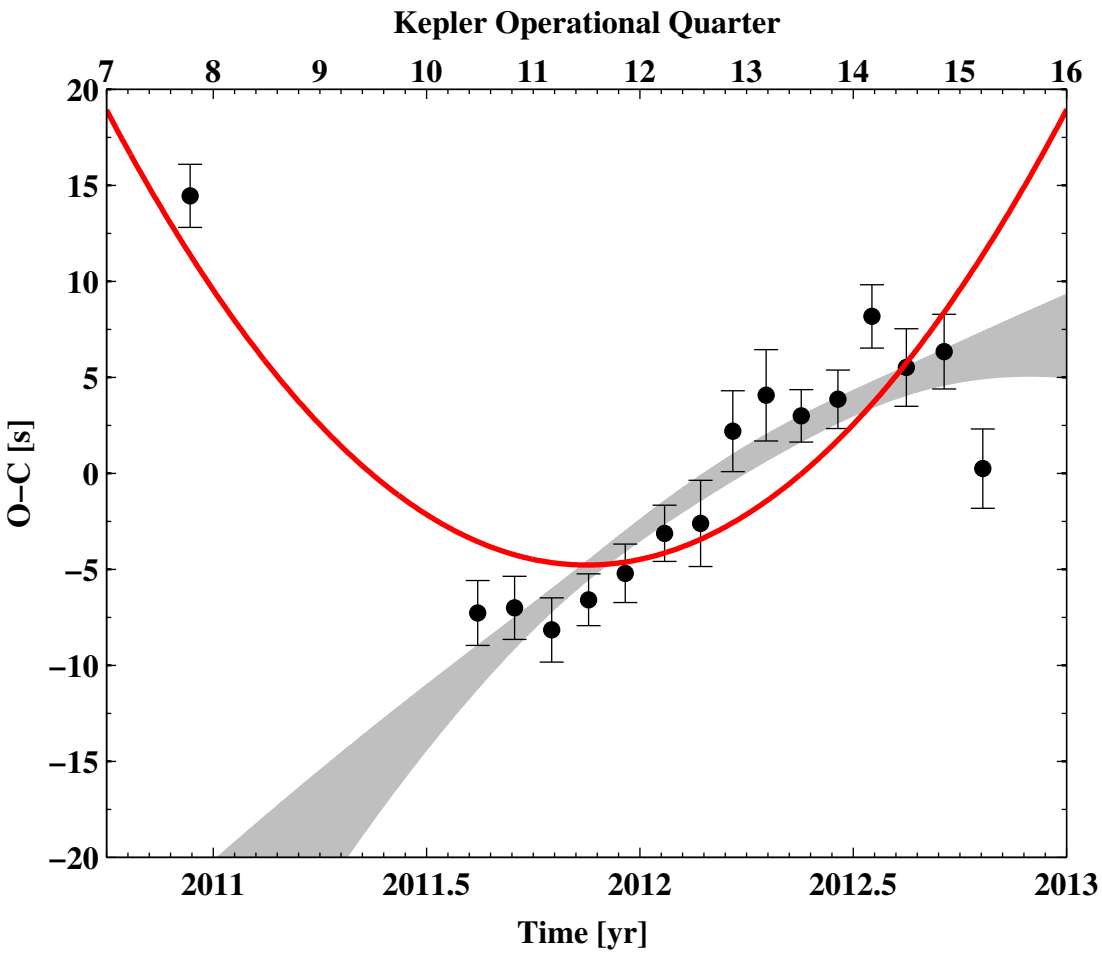

Figure 3.14: $O-C$ of mode C^+ with a linear trend removed. The red line is the best fit of a parabola and the grey region is the $1\,\sigma$ prediction made by the Q10-14 values and a parabola.

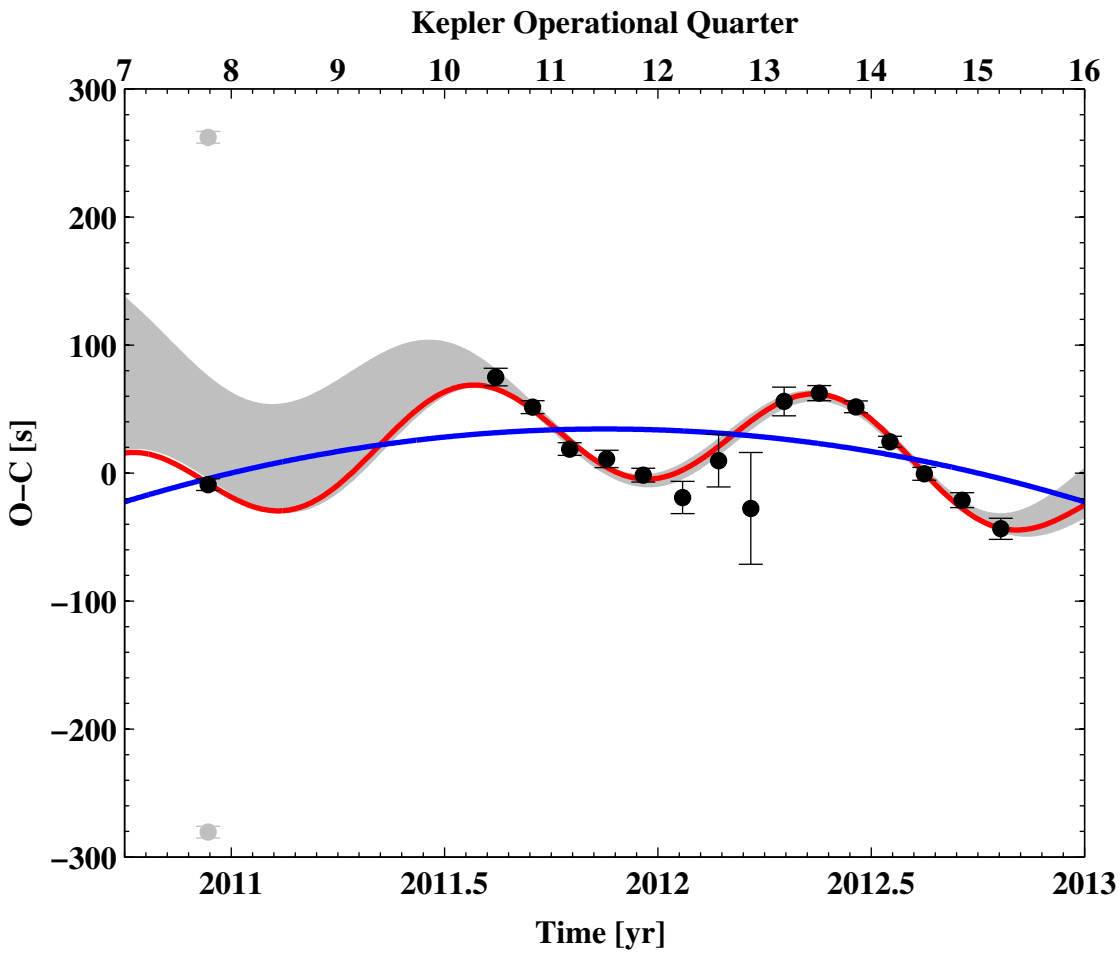

Figure 3.15: $O-C$ of mode D^- with a linear trend removed. The red line is the best fit of a sinusoid plus a parabola and the blue line is the best fit with the sinusoidal component removed. The grey region is the $1\,\sigma$ prediction made by the Q10-14 values and a sinusoid plus a parabola.

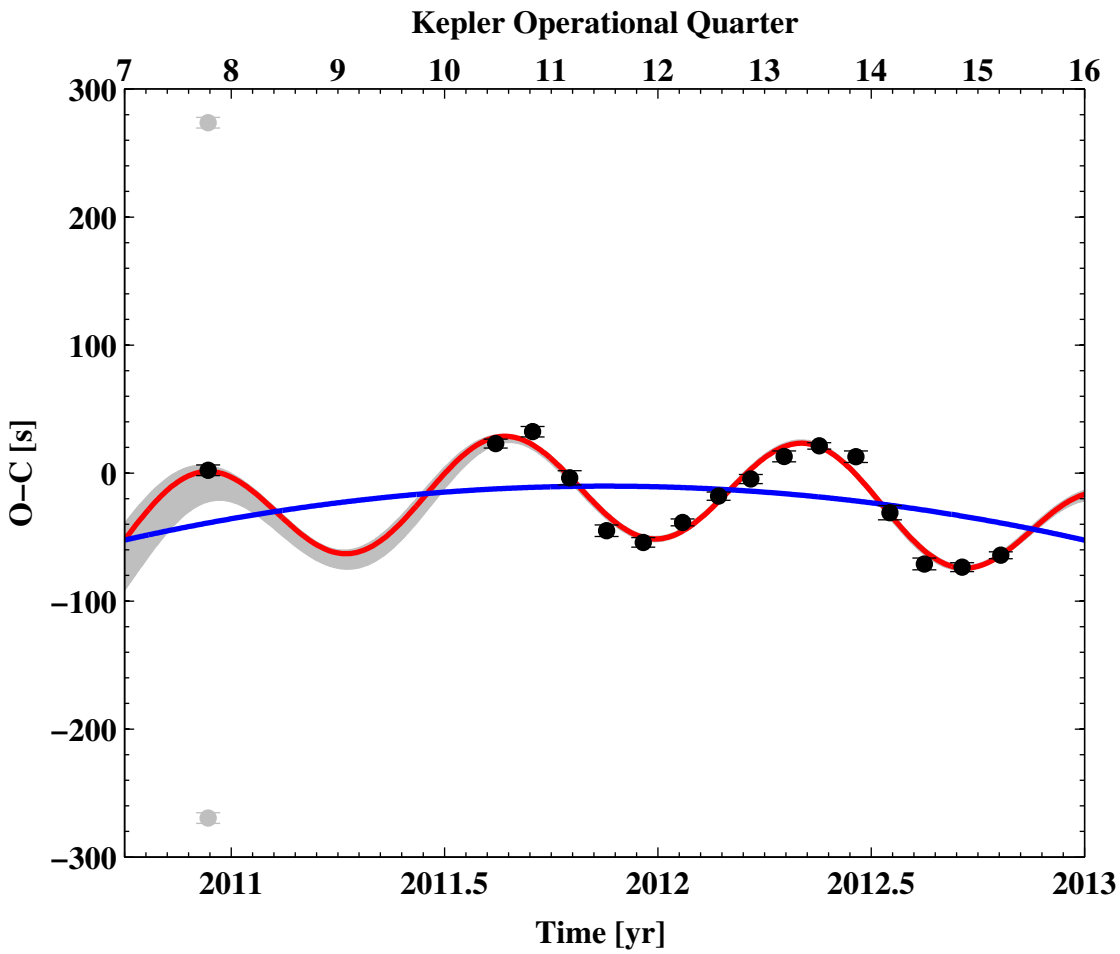

Figure 3.16: $O - C$ of mode D with a linear trend removed. The red line is the best fit of a sinusoid plus a parabola and the blue line is the best fit with the sinusoidal component removed. The grey region is the $1\,\sigma$ prediction made by the Q10-14 values and a sinusoid plus a parabola.

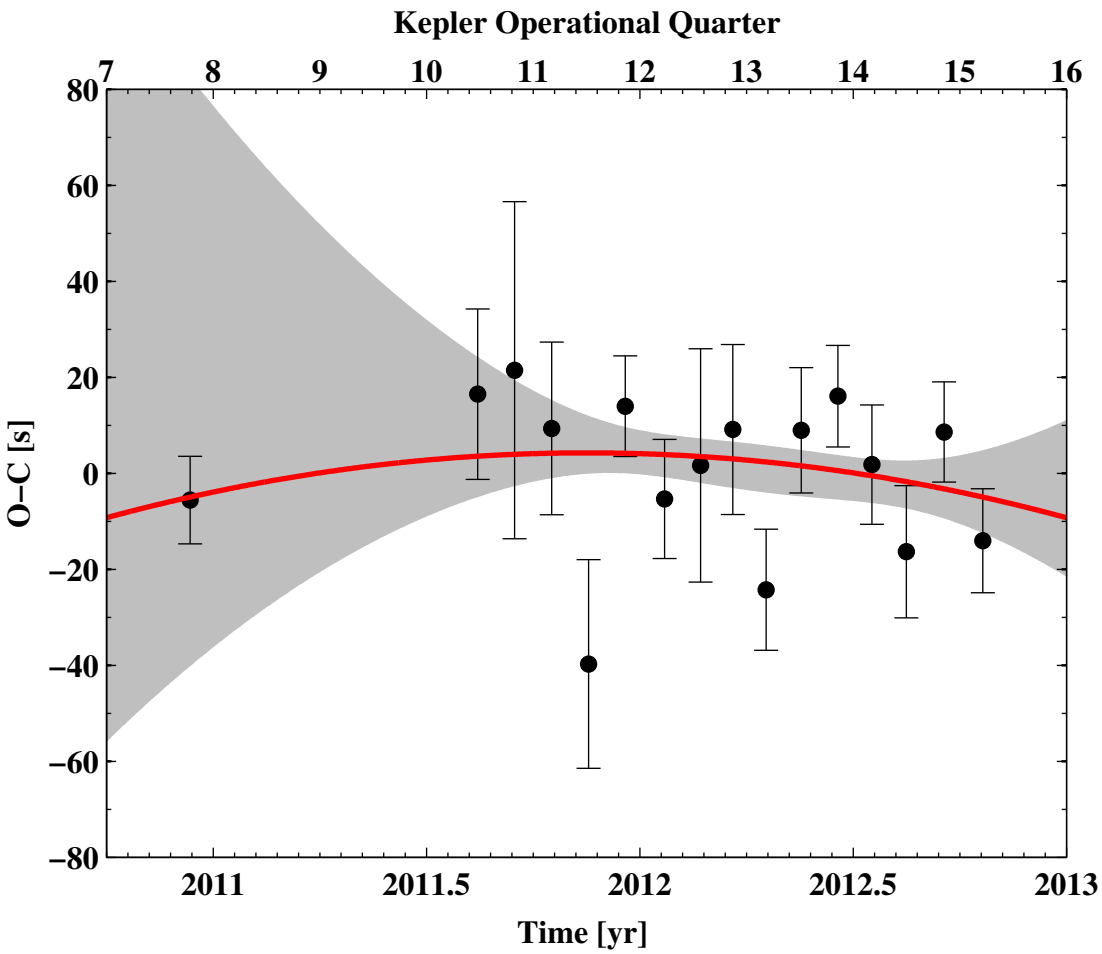

Figure 3.17: $O - C$ of mode D^+ with a linear trend removed. The red line is the best fit of a parabola and the grey region is the $1\,\sigma$ prediction made by the Q10-14 values and a parabola.

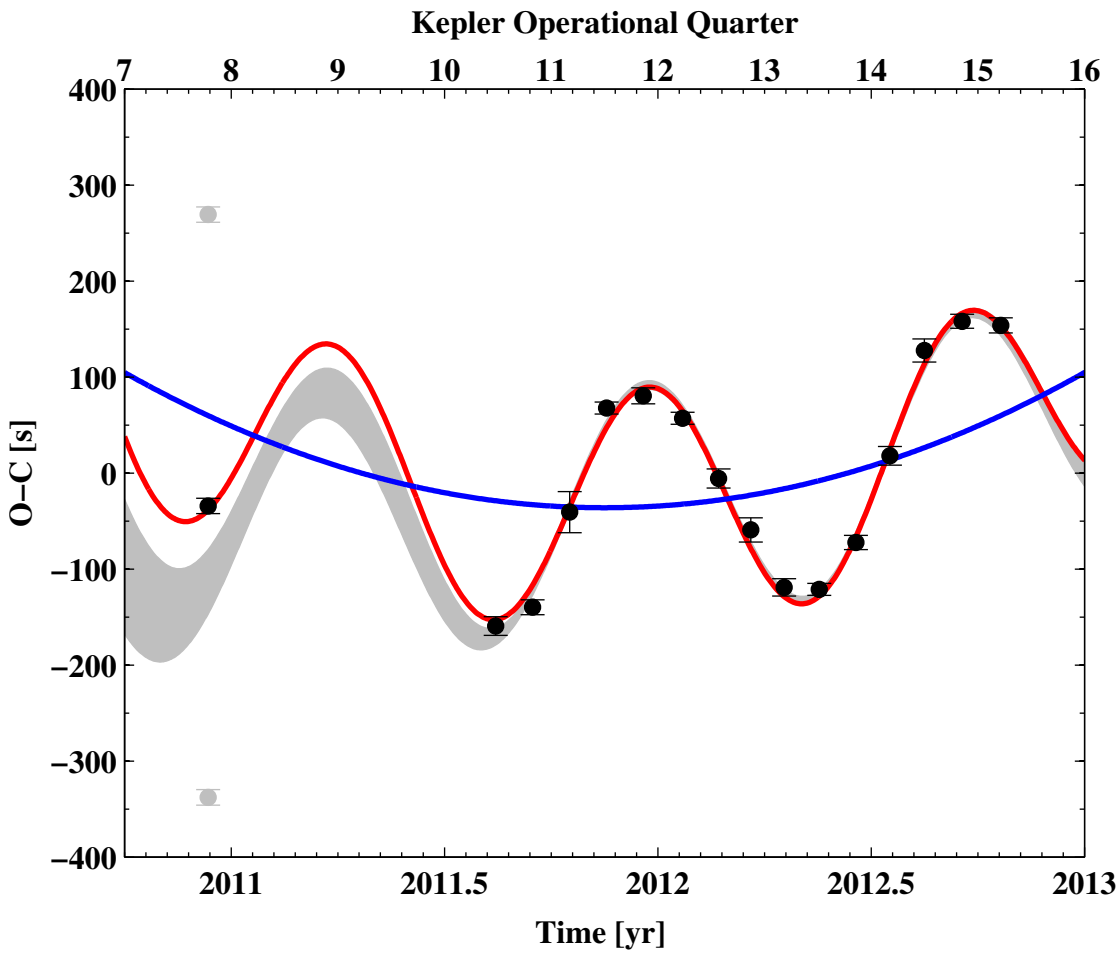

Figure 3.18: $O - C$ of mode E with a linear trend removed. The red line is the best fit of a sinusoid plus a parabola and the blue line is the best fit with the sinusoidal component removed. The grey region is the $1\,\sigma$ prediction made by the Q10-14 values and a sinusoid plus a parabola.

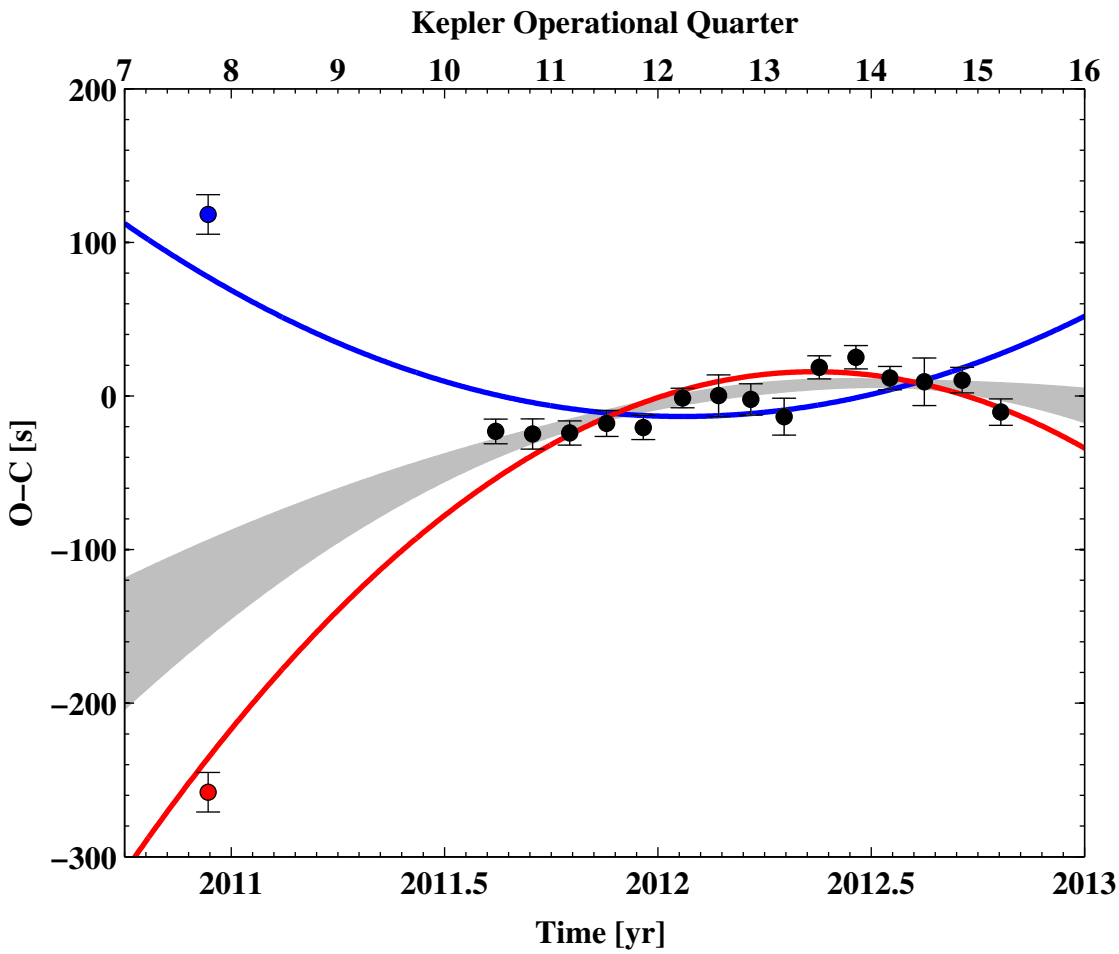

Figure 3.19: $O-C$ of mode F with a linear trend removed. The red and blue colored points are possible Q7.2 $O-C$ values. The red and blue lines are the best fit of a parabola given the selection of the correspondingly colored Q7.2 $O-C$ value. The grey region is the $1\,\sigma$ prediction made by the Q10-14 values and a parabola.

The $O-C$ of modes $C^+ + F$, A^-, A, A^+, C^-, C^+, D^+, and F does not appear to vary periodically. For each mode except $C^+ + F$ and F, there is little ambiguity as to the appropriate Q7.2 $O-C$ value. In an attempt to select the appropriate Q7.2 $O-C$ values for modes $C^+ + F$ and F and as a sanity check for the assumed Q7.2 $O-C$ values for the other modes, the method described in Section 1.4.5 was used with parabola and the Q10-14 values to calculate the statistical likelihood of each possible Q7.2 $O-C$ value. For modes A^-, A, A^+, and D^+, the model prediction confirmed the chosen Q7.2 $O-C$ values. For modes $C^+ + F$, C^-, C^+, and F, the model prediction did not agree with any of the possible Q7.2 $O-C$ values. Despite this, for modes C^- and C^+ only the selected Q7.2 $O-C$ values seem reasonable. The 1σ model predictions are shown as a shaded region in Figures 3.7, 3.8, 3.9, 3.10, 3.12, 3.14, 3.17, and 3.19. To further investigate, a sinusoid plus parabola was used to attempt to constrain the Q7.2 $O-C$ values for modes $C^+ + F$, C^-, C^+, and F. However, no minimum in χ^2 was found for any of these modes for any reasonable period. It is concluded that a sinusoid plus parabola is under-constrained by the Q10-14 $O-C$ values for these modes.

A parabola was fit to $O-C$ for modes $C^+ + F$, A^{-1}, A, A^+, C^-, C^+, D^+, and F. A fit was performed for selection of each of the possible Q7.2 $O-C$ values for modes $C^+ + F$ and F. The residuals of each fit were used to calculate the upper limit on the amplitude of any unmodelled sinusoidal variations in $O-C$. The results are shown in Table 3.3 and the fits are shown in red in Figures 3.7, 3.8, 3.9, 3.10, 3.12, 3.14, 3.17, and 3.19. The fit parameters are plotted in Figure 3.20.

The $O-C$s of the remaining modes, namely B, C, D^-, D, and E, appear to vary periodically. There is ambiguity as to the appropriate Q7.2 $O-C$ values for each mode except C. A sinusoid plus a parabola was used with the Q10-14 values to calculate the statistical likelihood of each possible Q7.2 $O-C$ value. The model prediction selected Q7.2 $O-C$ values for modes D^-, D, and E, but no minimum in χ^2 was found for modes B and C for any period. The previously selected Q7.2 $O-C$ value for mode C was still considered the only reasonable value, making the Q7.2 $O-C$ value for mode B the only case of ambiguity the periodic $O-C$s. For each of these modes, and for

Table 3.3: Results of fitting a parabola or a sinusoid plus a parabola to $O-C$. For modes $C^+ + F$, B, and F, the color of the label corresponds to the color of the selected Q7.2 $O-C$ values shown in Figures 3.7, 3.11, and 3.19 respectively.

Label	**Frequency** $[\mu Hz]$	$\dot{f}$ $[nHz^2]$	α $[nHz]$	Π $[yr]$	$\chi^2_{Reduced}$
$C^+ + F$	6965.2864(0.0008)	460(70)	$< 18(7)$	-	1.9
$C^+ + F$	6965.2999(0.0008)	$-700(70)$	$< 21(6)$	-	3.0
A^-	5076.4355(0.0006)	60(60)	$< 4(3)$	-	0.4
A	5073.2370(0.0002)	$-170(20)$	$< 3(2)$	-	1.0
A^+	5070.0303(0.0002)	30(20)	$< 4(2)$	-	1.2
B	4398.3887(0.0008)	0(200)	28(3)	1.79	0.8
B	4398.3722(0.0008)	680(70)	18(2)	1.34	1.0
C^-	4313.3052(0.0002)	$-210(20)$	$< 9(1)$	-	8.5
C	4309.9153(0.0002)	90(10)	6.3(0.8)	1.24	1.2
C^+	4306.5224(0.0001)	$-160(10)$	$< 4.6(0.8)$	-	5.6
D^-	3684.9669(0.0003)	320(40)	34(3)	0.84	1.3
D	3681.8001(0.0003)	240(20)	42(1)	0.72	4.8
D^+	3677.9946(0.0007)	80(70)	$< 14(9)$	-	1.3
E	3294.3410(0.0005)	$-710(50)$	109(3)	0.75	3.1
F	2658.7627(0.0005)	660(50)	$< 17(2)$	-	2.8
F	2658.7753(0.0005)	$-390(50)$	$< 17(3)$	-	6.0

selection of each reasonable Q7.2 $O-C$ value, a sinusoid plus a parabola was fit to $O-C$. Reduced χ^2 as a function of period is shown in Figure 3.21. For modes D^-, D, and E the minimum in χ^2 occur at very similar periods. The total minimum for those three modes occurs at a period of .74 yr. Forcing the period to .74 yr resulted in an increase in reduced χ^2 and a notably poor fit to mode D^-. It was decided to leave the period as a free parameter. The results are shown in Table 3.3 and the fit parameters are plotted in Figure 3.20. The fits and the model predictions are shown in Figures 3.11, 3.13, 3.15, 3.16, and 3.18.

Due to the ambiguity as to the appropriate Q7.2 $O-C$ value for combination mode $C^+ + F$ and the generally poor fits to $O-C$ for modes C^+ and F, no comparison of the parent and combination mode frequencies was performed.

The residuals of the fits that were deemed reliable, namely those for modes A^-, A, A^+, C, D^-, D, D^+, and E, were averaged and are shown in Figure 3.22. The Fourier spectrum of the average residuals is shown in Figure 3.23. There are no notable variations in the residuals and no significant power in the DFT.

3.4 Amplitude Modulation

The amplitude of each mode as a function of time is shown in Figures 3.25-3.37. There is quite a wide variety of trends in mode amplitude. For modes B, C, D^-, and D, a low order polynomial plus a sinusoid was fit to the amplitudes. For the remaining modes a low order polynomial was fit to the amplitudes and the upper limit on the amplitude of any sinusoidal variations in mode amplitude were calculated. The fits are shown in Figures 3.25-3.37 and the results are summarized in Table 3.4. The RC value of the combination mode $C^+ + F$ was also investigated and found to be statistically consistent with constant (see Figure 3.38).

The $O-C$ and amplitude variations of modes D^- and D are compared in Figure 3.24. It appears that the periodic components of the $O-C$ variations are in phase but the amplitude variations are 180 degrees out of phase. There is also a quarter cycle phase difference between the $O-C$ and amplitude variations.

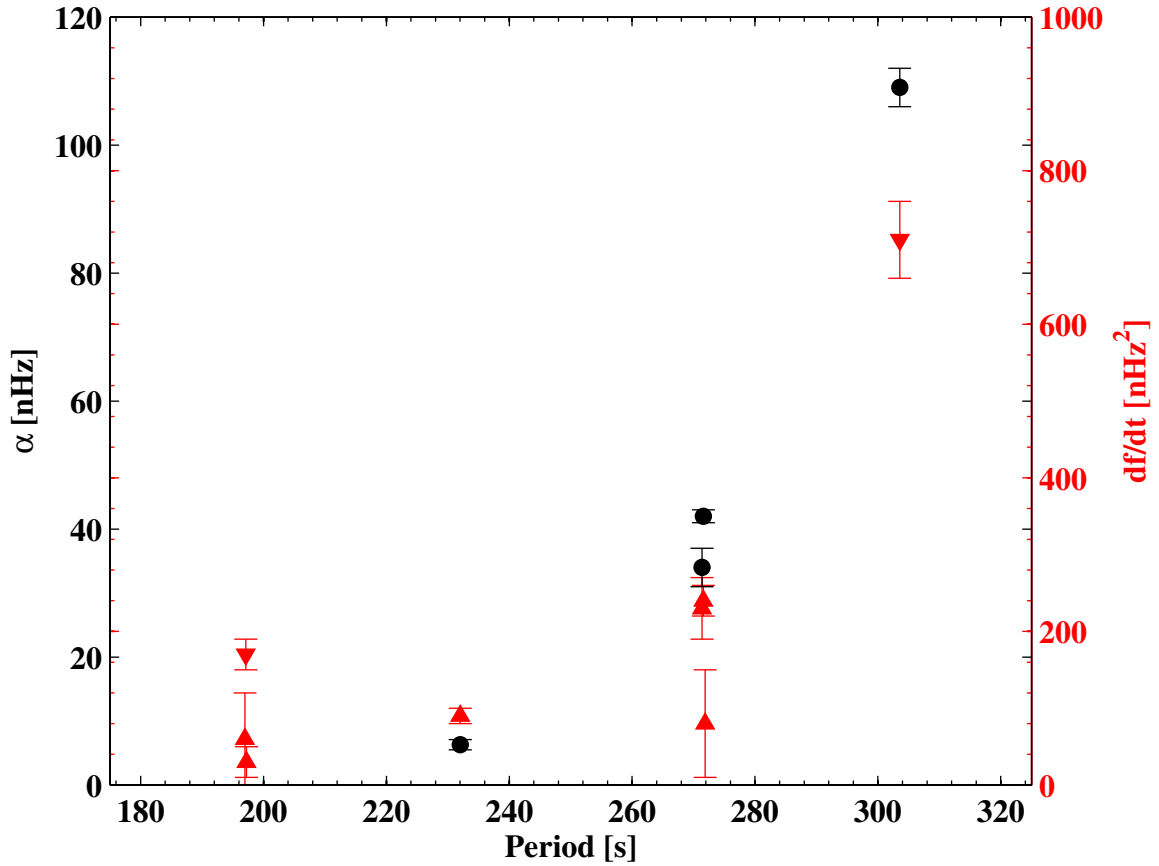

Figure 3.20: Frequency variations of the normal modes. The black points indicate the amplitude of the sinusoidal variations and the red triangles indicate the magnitude of the secular changes. The sign of the secular changes is indicated by the direction of the triangle.

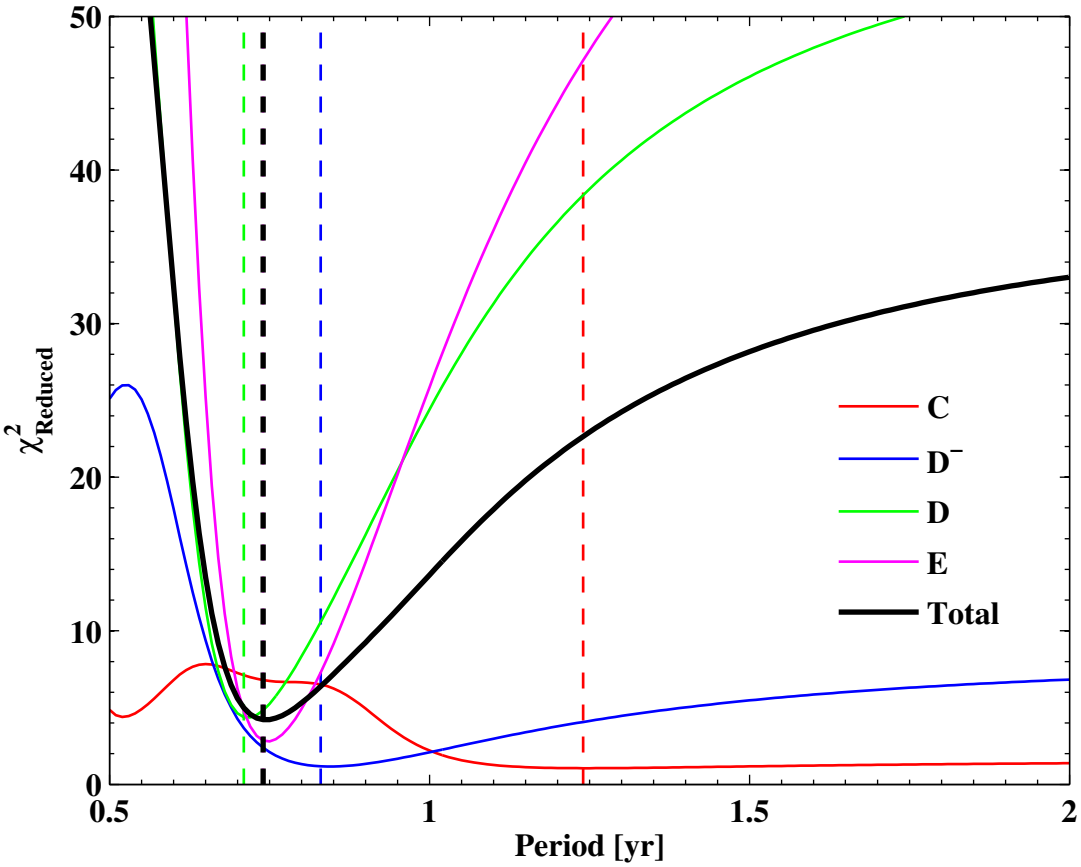

Figure 3.21: Reduced χ^2 as a function of the period for fits of a sinusoid plus parabola to $O - C$ for selected modes. The dotted lines indicate the global minima.

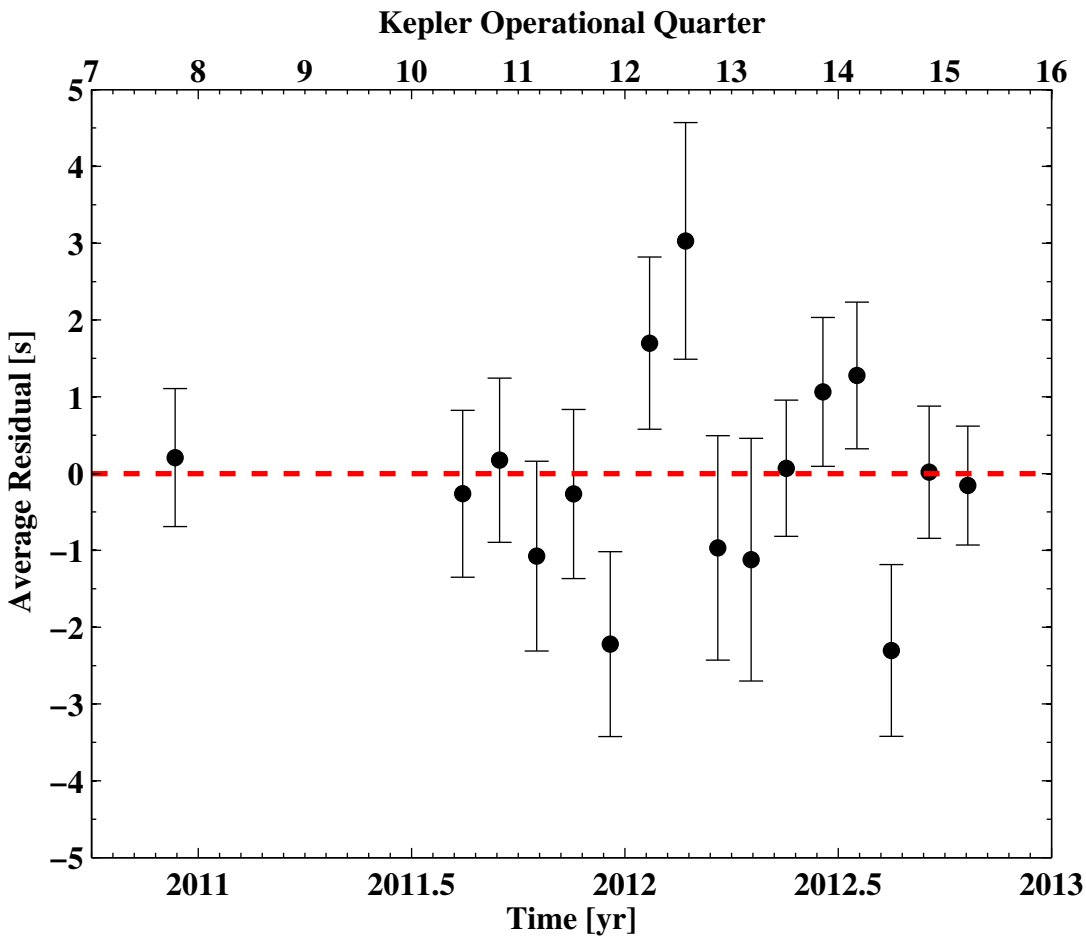

Figure 3.22: Average residuals of the fits to the $O - C$ of modes A^-, A, A^+, C, D^-, D, D^+, and E.

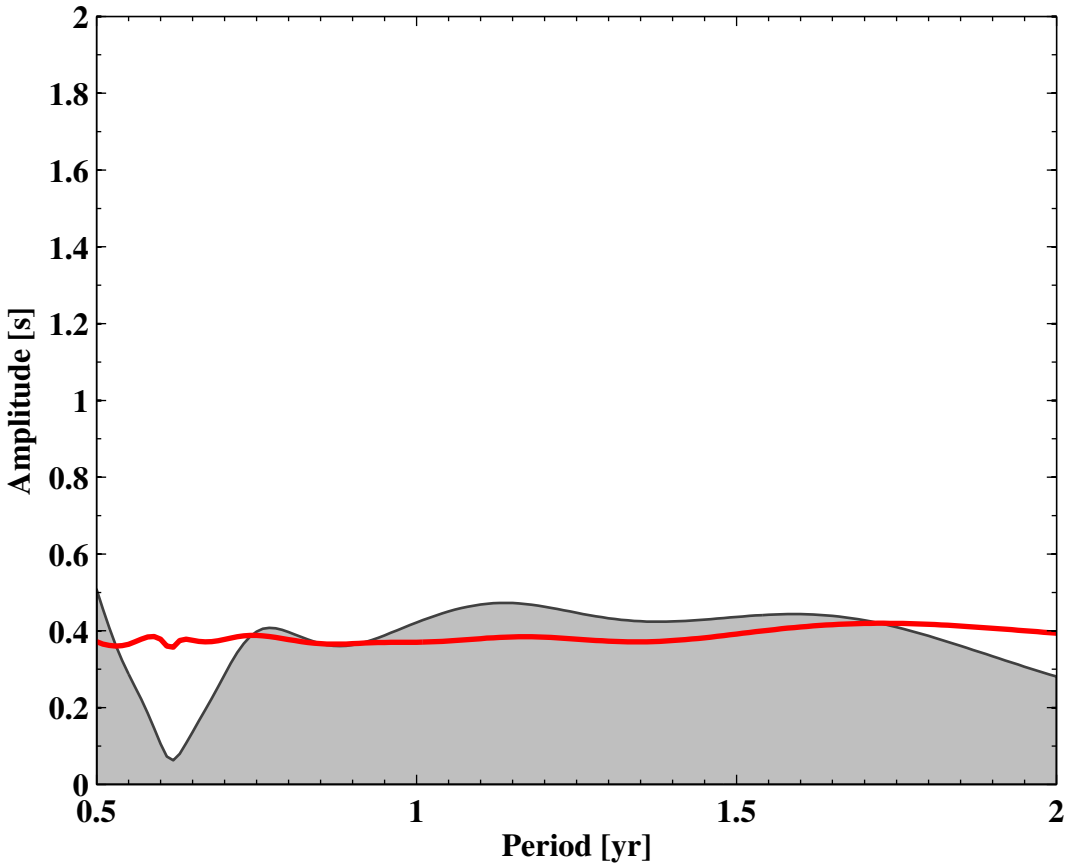

Figure 3.23: Fourier spectrum of the average residuals of the fits to the $O - C$ of modes A^-, A, A^+, C, D^-, D, D^+, and E. The red line is 6 times the formal error.

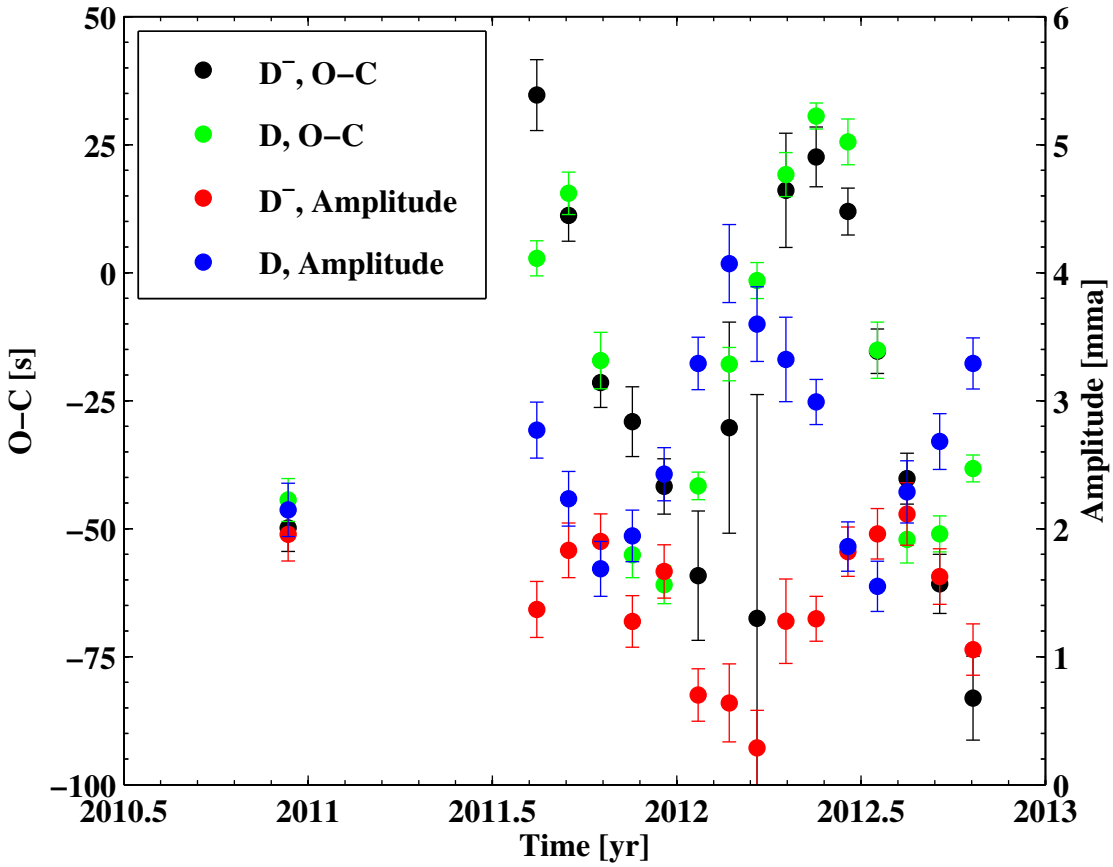

Figure 3.24: Comparison of the amplitude and $O - C$ variations of modes D^- and D.

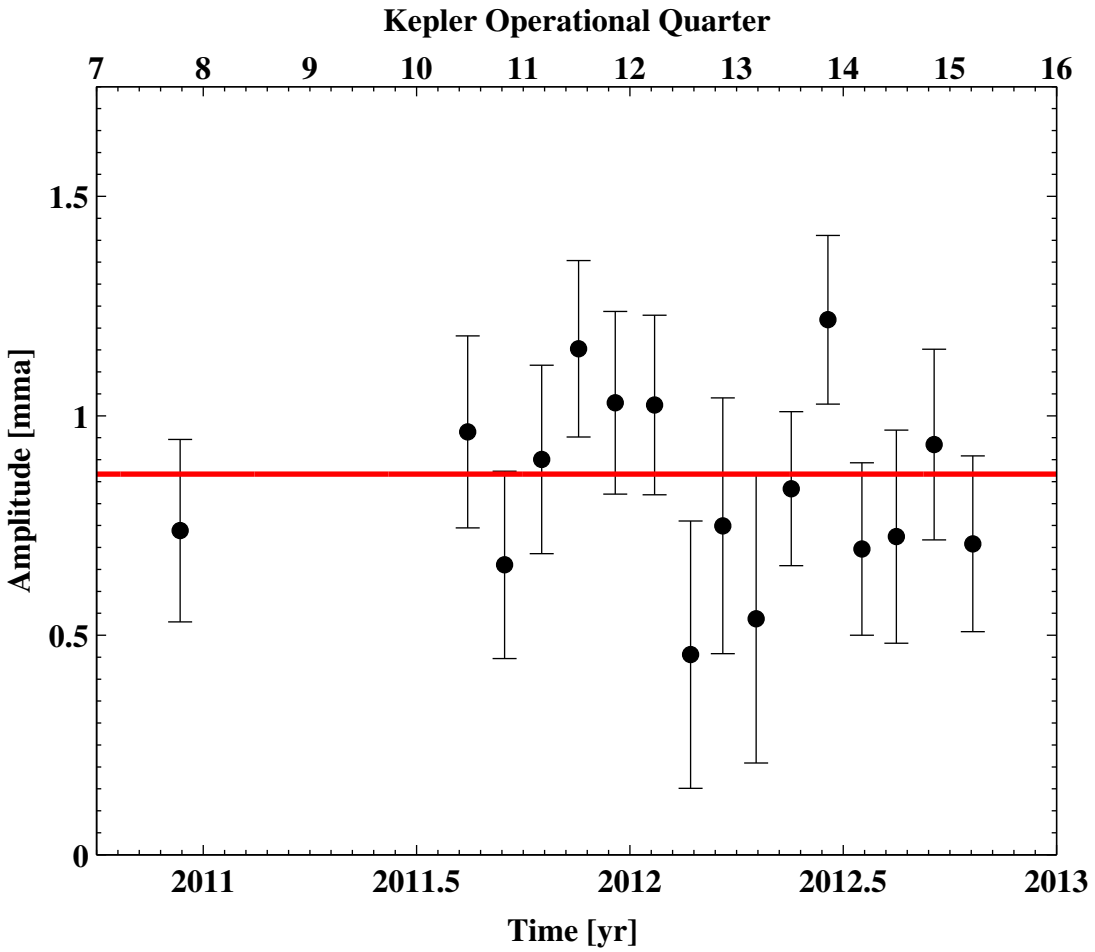

Figure 3.25: Amplitude of mode $C^+ + F$. The best fit to a constant amplitude is shown in red.

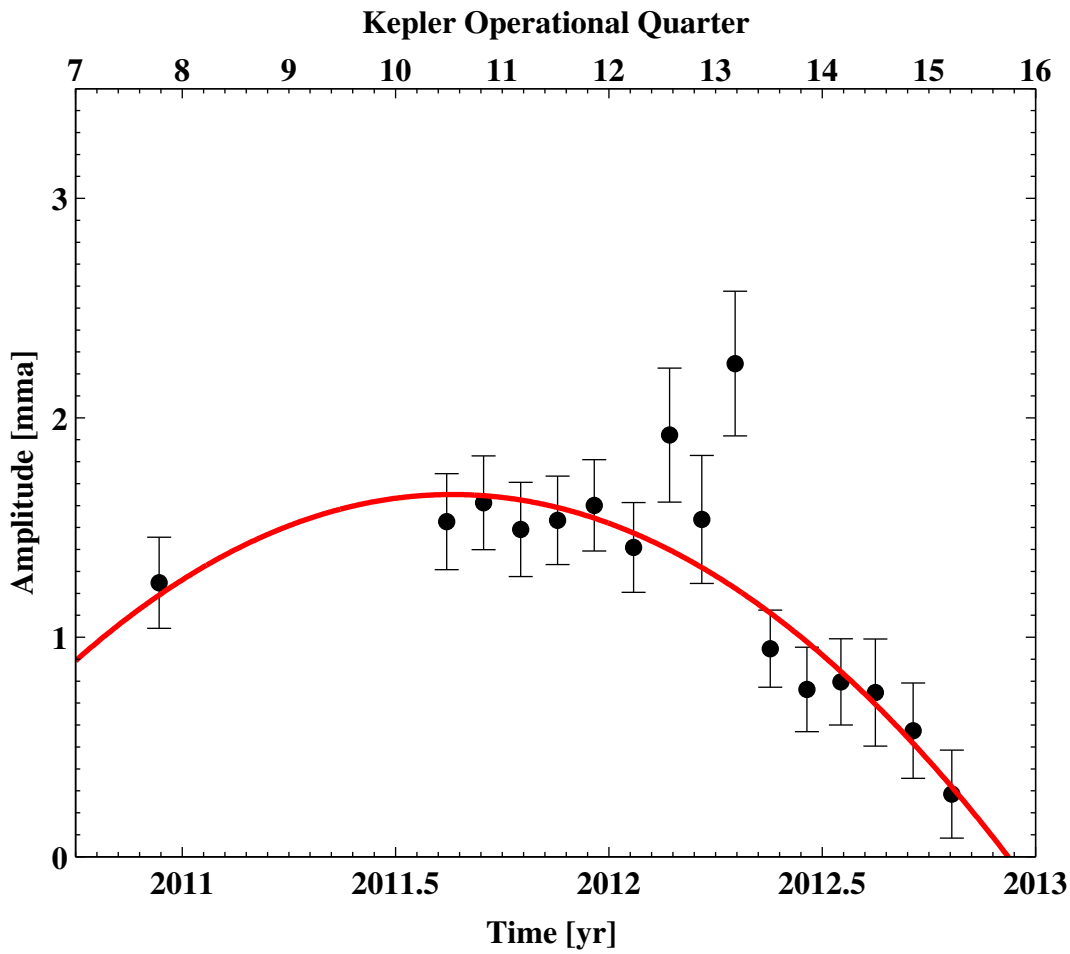

Figure 3.26: Amplitude of mode A^-. The best fit of a second order polynomial is shown in red.

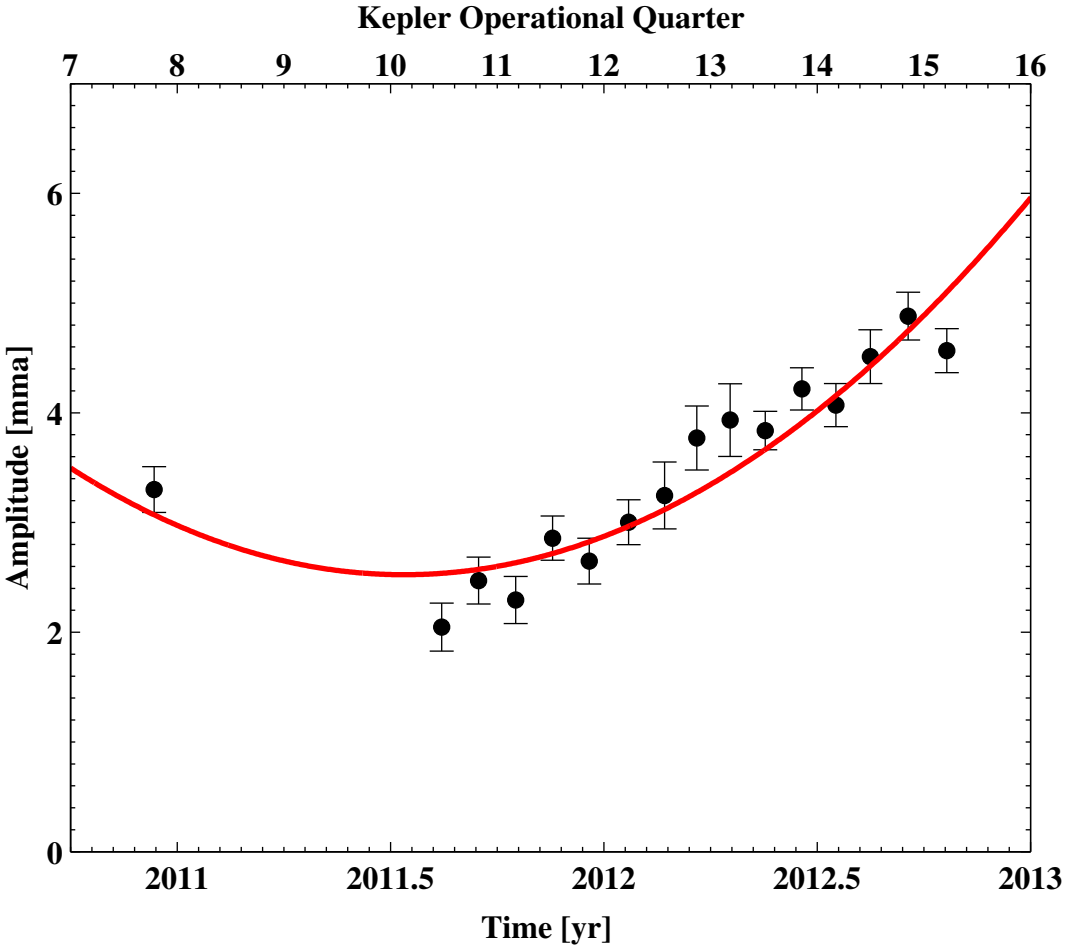

Figure 3.27: Amplitude of mode A. The best fit of a second order polynomial is shown in red.

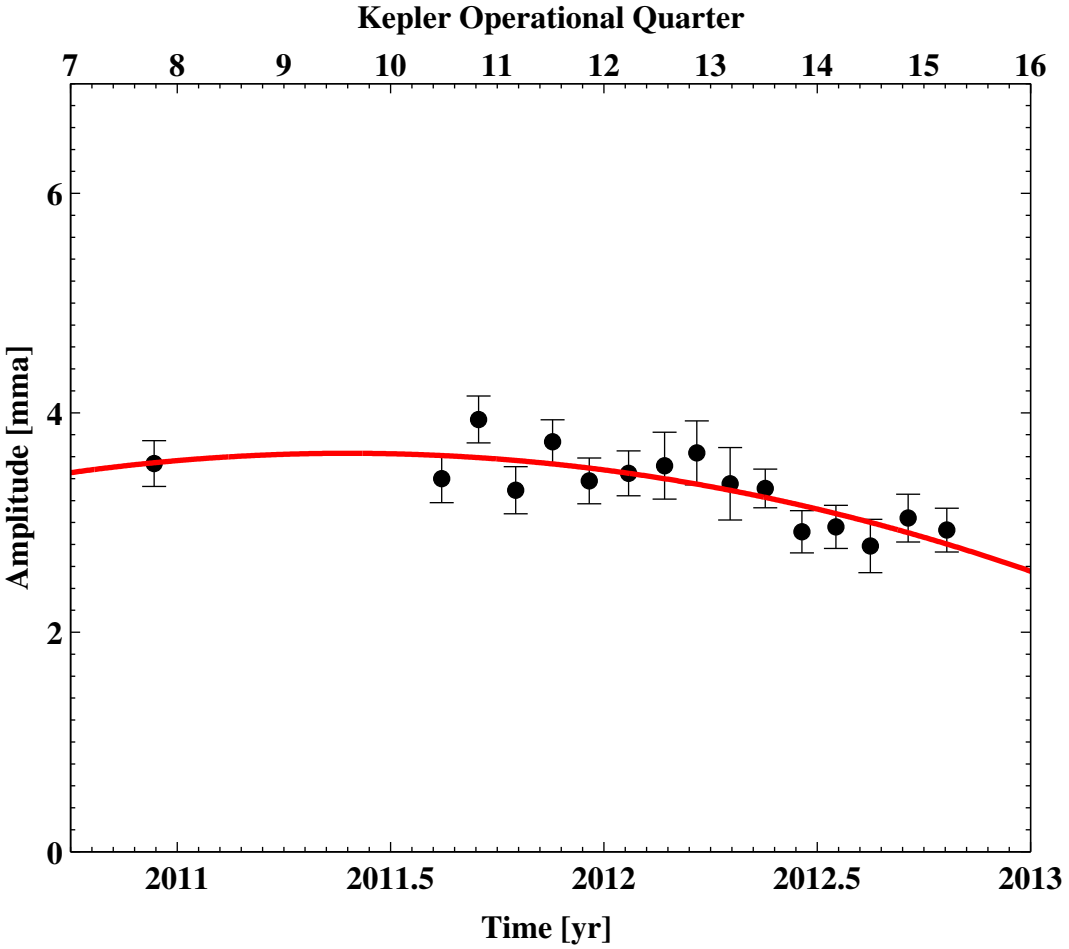

Figure 3.28: Amplitude of mode A^+. The best fit of second order polynomial is shown in red.

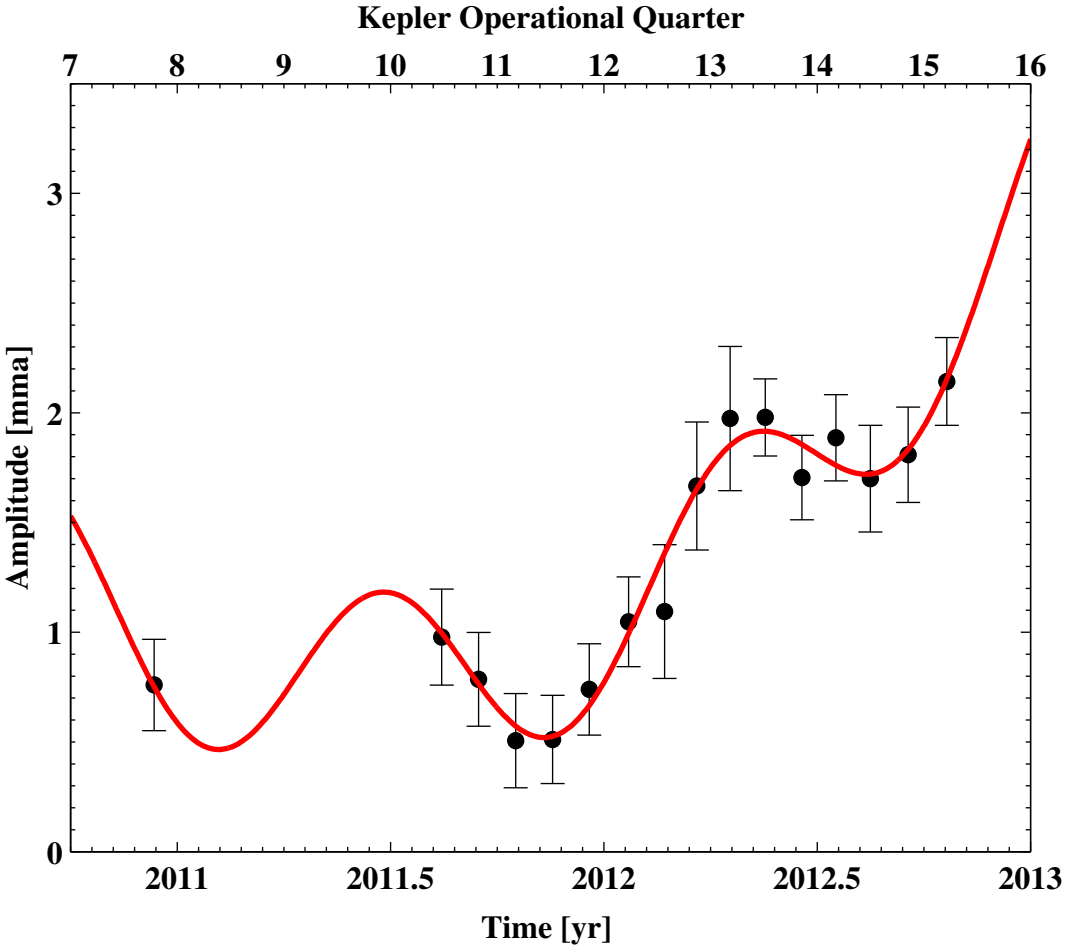

Figure 3.29: Amplitude of mode B. The best fit of a second order polynomial plus a sinusoid is shown in red.

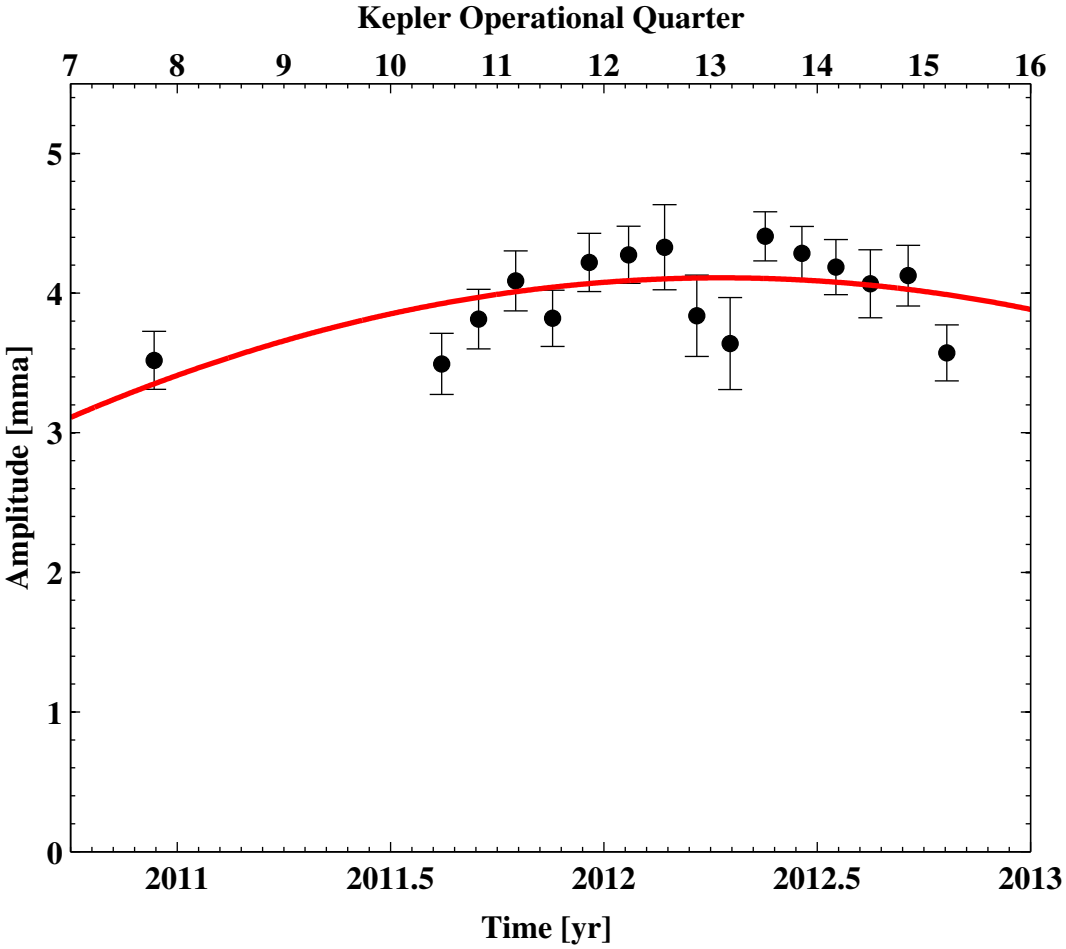

Figure 3.30: Amplitude of mode C^-. The best fit of a second order polynomial is shown in red.

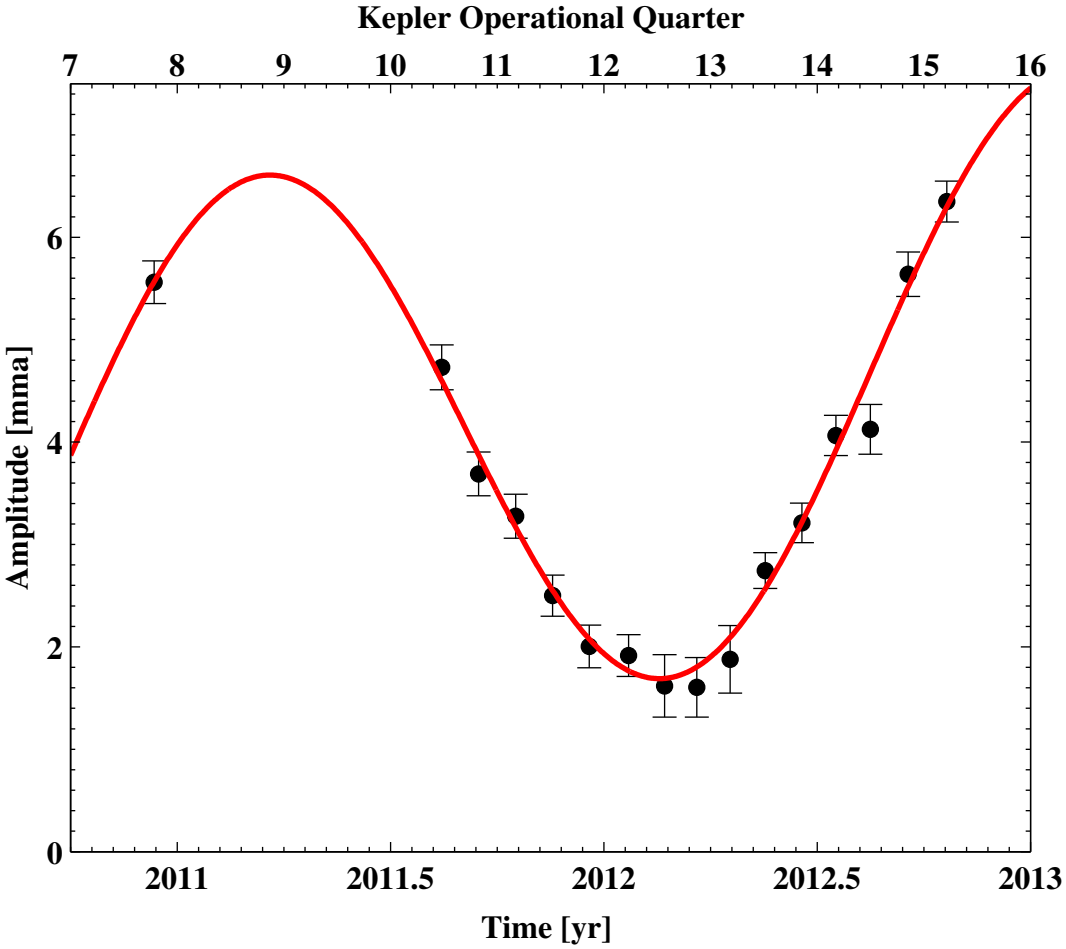

Figure 3.31: Amplitude of mode C. The best fit of a second order polynomial plus a sinusoid is shown in red.

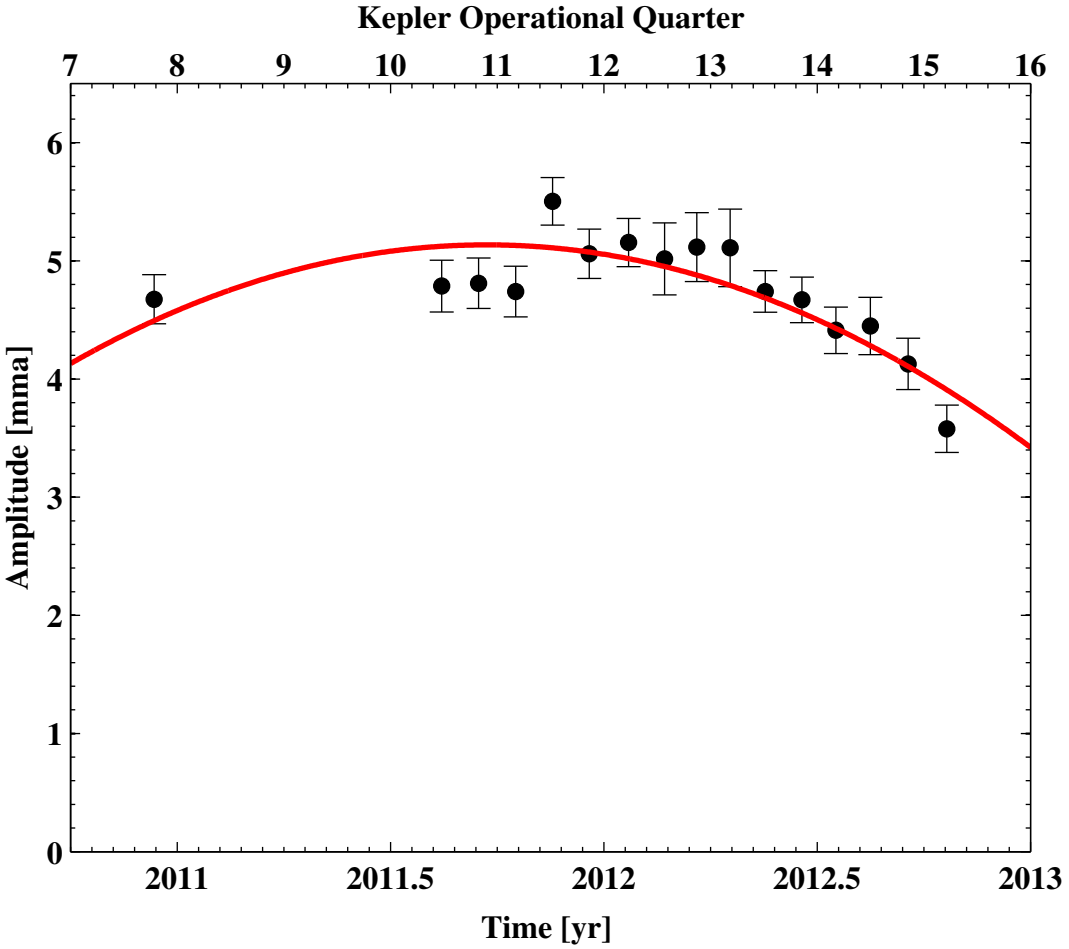

Figure 3.32: Amplitude of mode C^+. The best fit of a second order polynomial is shown in red.

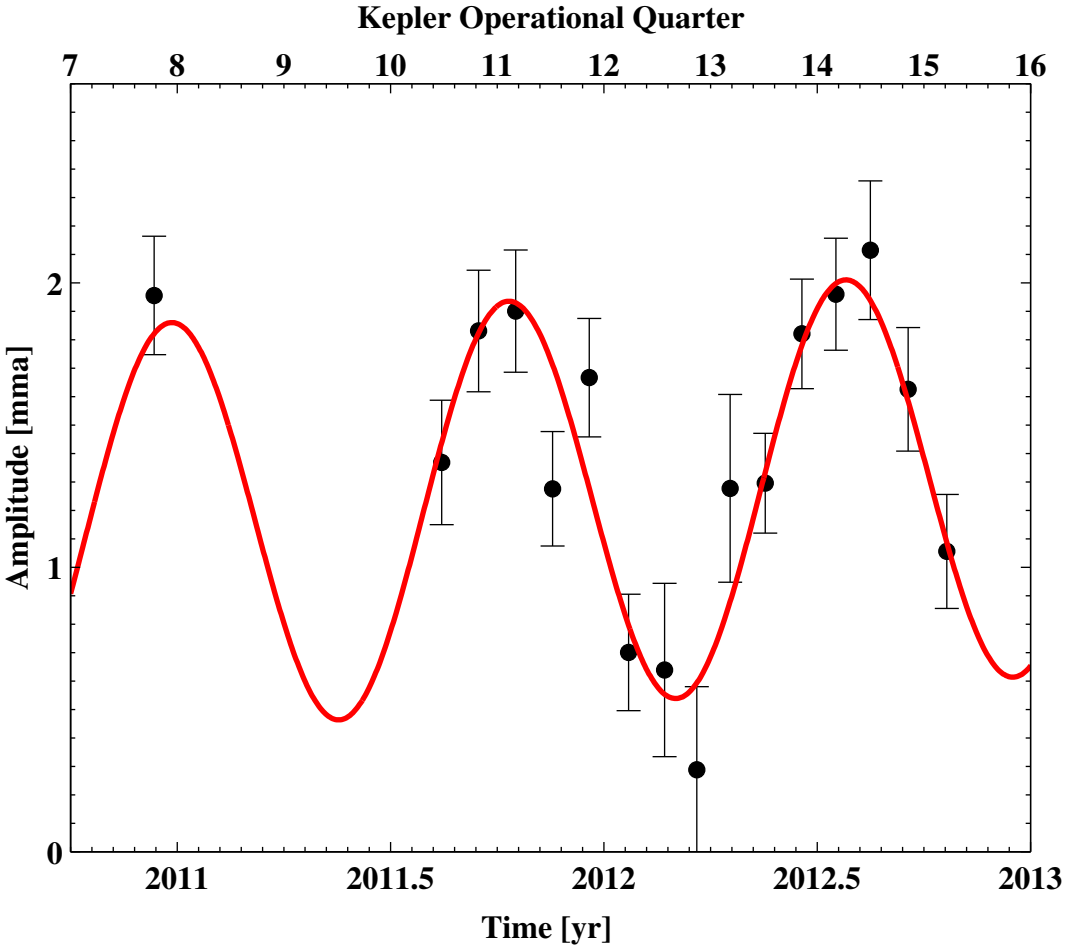

Figure 3.33: Amplitude of mode D^-. The best fit of a first order polynomial plus a sinusoid is shown in red.

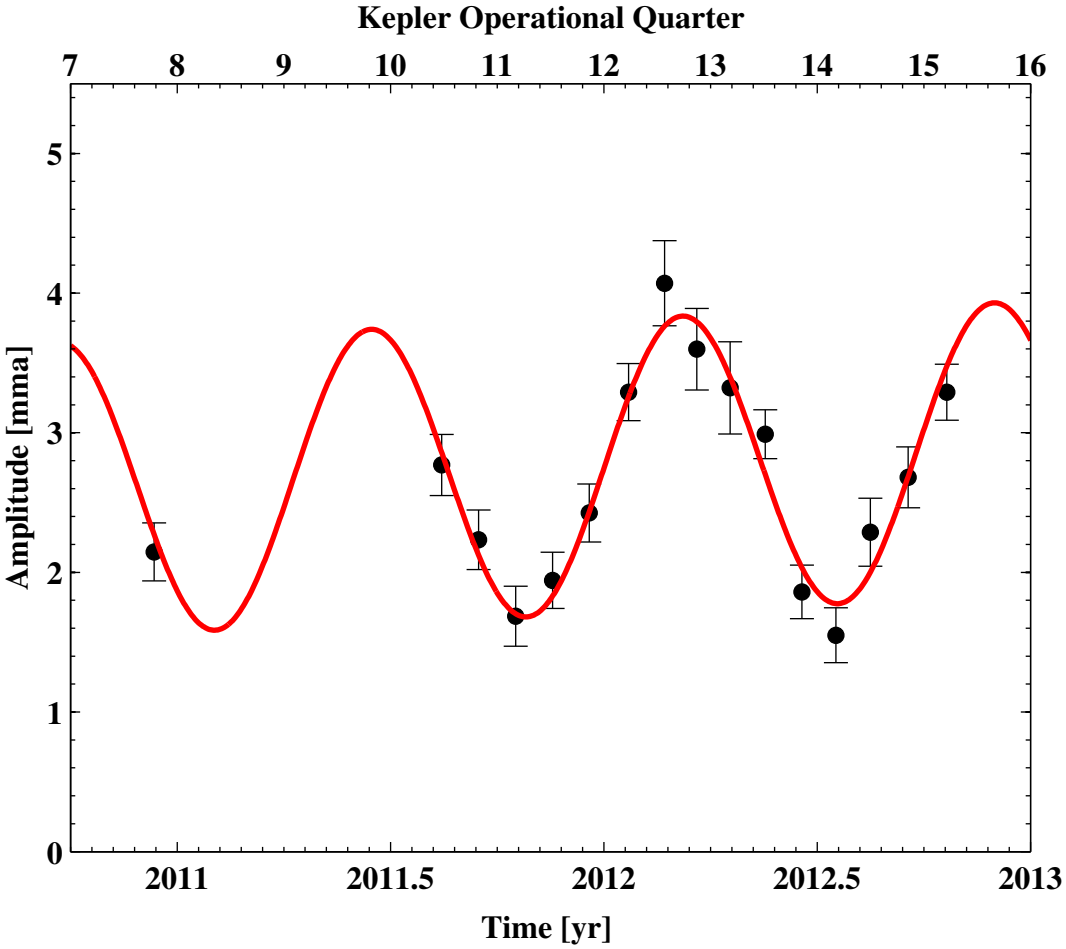

Figure 3.34: Amplitude of mode D. The best fit of a second order polynomial plus a sinusoid is shown in red.

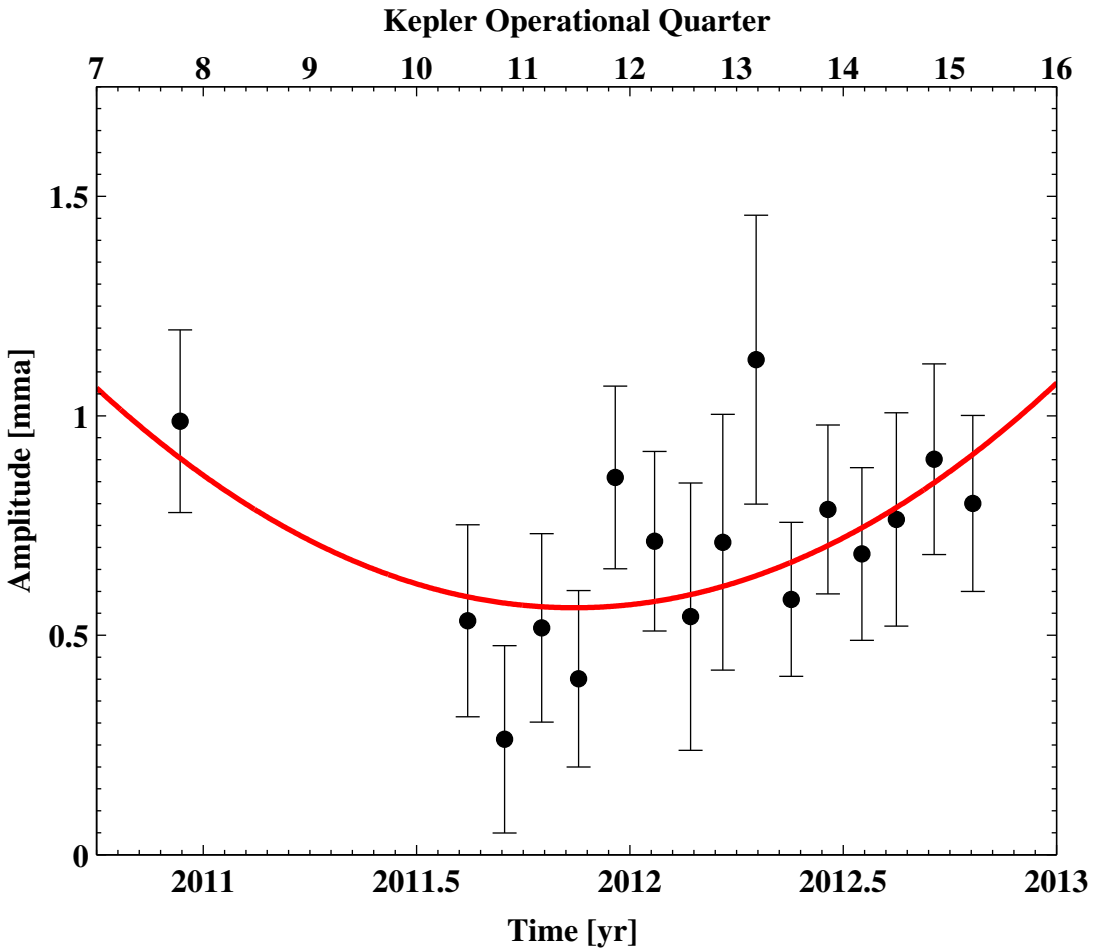

Figure 3.35: Amplitude of mode D^+. The best fit of a second order polynomial is shown in red.

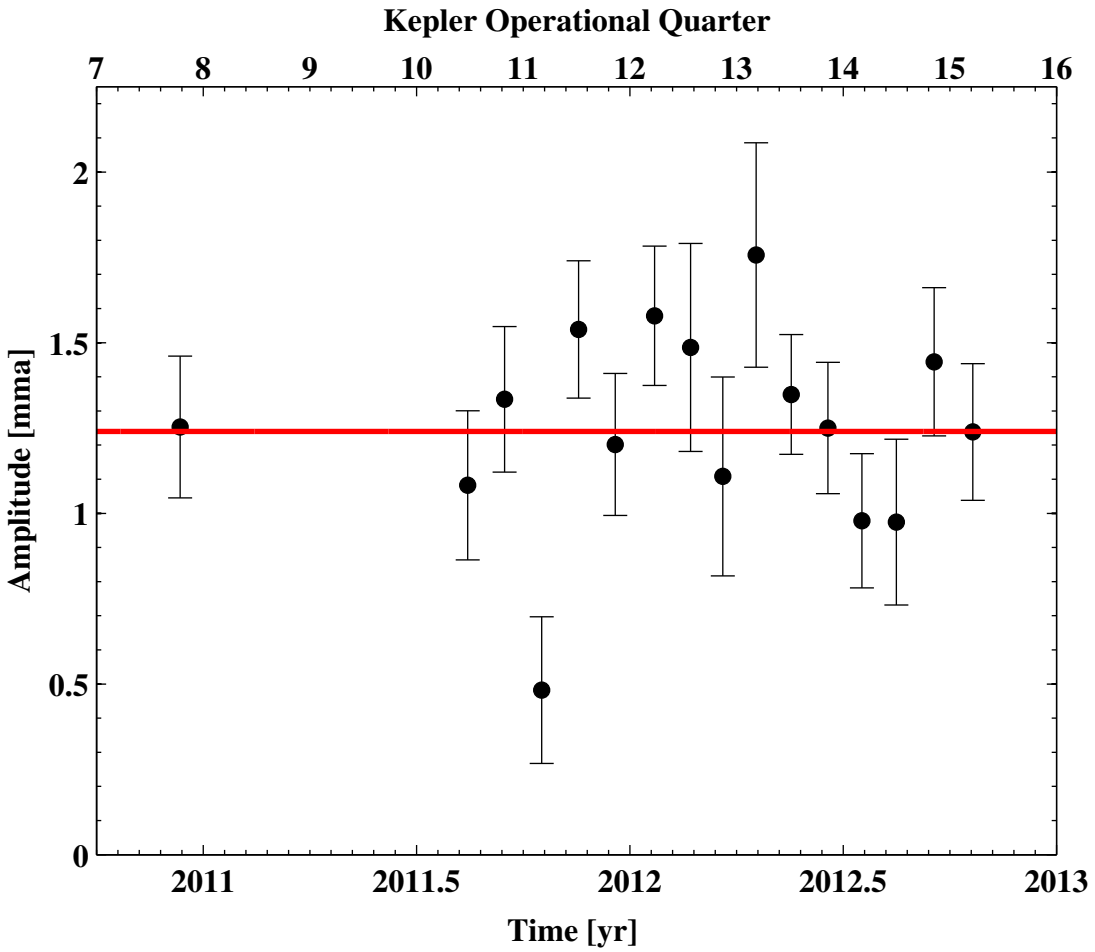

Figure 3.36: Amplitude of mode E. The best fit to a constant amplitude is shown in red.

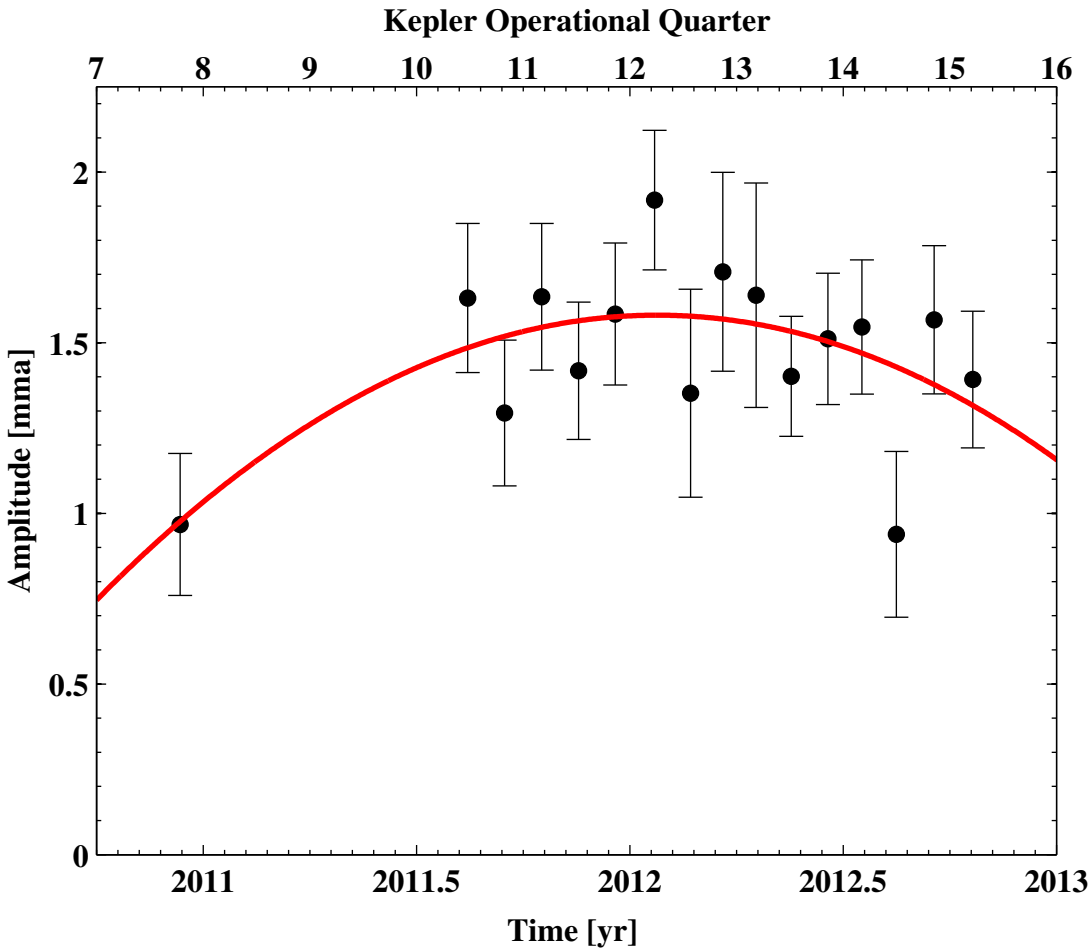

Figure 3.37: Amplitude of mode F. The best fit of a second order polynomial is shown in red.

Table 3.4: Values and upper limits of the amplitude of sinusoidal variations in mode amplitude (α). Also shown is the period of the sinusoidal variations in mode amplitude (Π) and reduced χ^2 for the fits to the amplitude shown in Figures 3.25-3.37.

Label	α $[mma]$	Π $[yr]$	$\chi^2_{Reduced}$
$C^+ + F$	$< 0.67(0.08)$	-	0.9
A^-	$< 0.19(0.08)$	-	1.1
A	$< 0.28(0.09)$	-	1.8
A^+	$< 0.14(0.08)$	-	0.8
B	$0.42(0.09)$	0.8	0.2
C^-	$< 0.21(0.08)$	-	1.3
C	$2.7(0.1)$	1.9	1.0
C^+	$< 0.19(0.08)$	-	1.2
D^-	$0.72(0.08)$	0.8	1.2
D	$1.07(0.08)$	0.7	0.9
D^+	$< 0.15(0.08)$	-	0.6
E	$< 0.73(0.08)$	-	1.7
F	$< 0.12(0.08)$	-	0.8

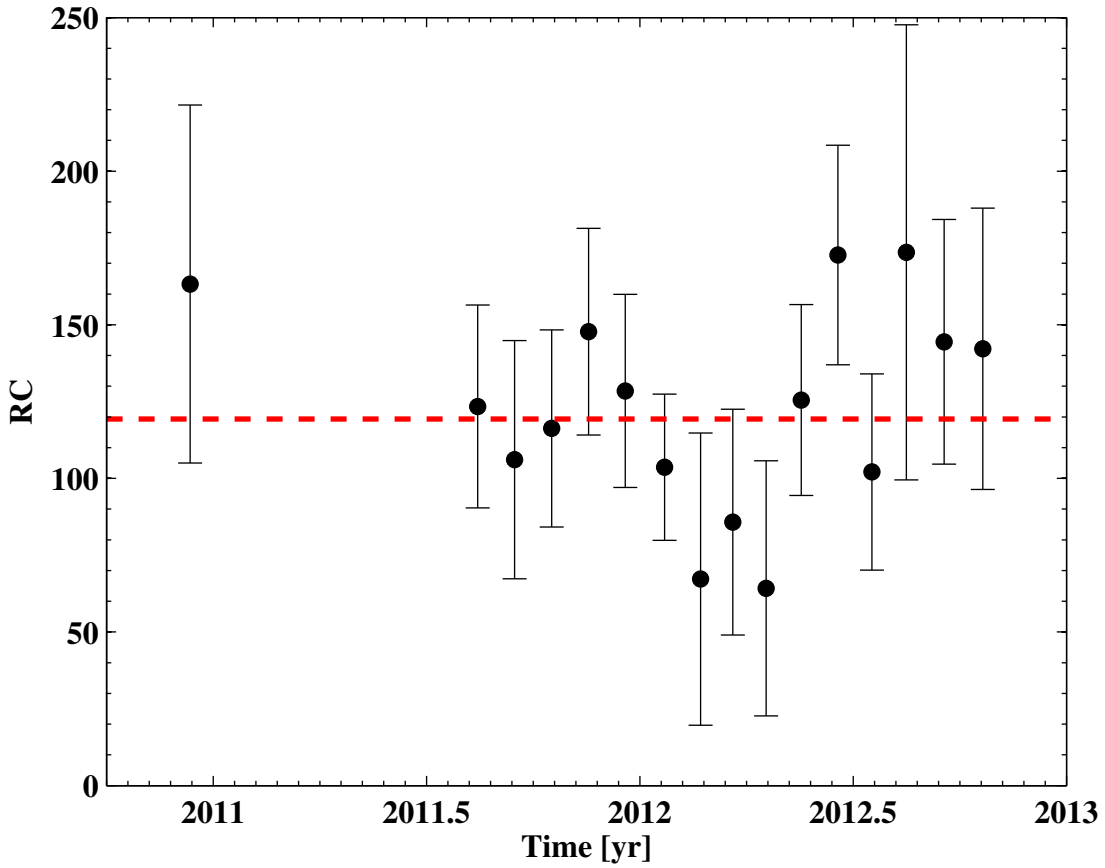

Figure 3.38: RC value of the $C^+ + F$ combination mode as a function of time.

Chapter 4

GD 66

4.1 Observations and Data Preparation

Over the last 10 years, over 450,000 images of GD 66 were taken from McDonald Observatory (MCDO) with the Argos frame transfer CCD on the 2.1 meter Otto Struve telescope. These observations were motivated by a pulsation timing based search for white dwarf planets (Mullally et al., 2008). Each series of images was reduced into a series of photometric measurements as described in Section 1.2.1 and all BJED corrections were calculated using the algorithm of Stumpff (1980). The photometric measurements were binned into 34 lightcurves by asserting that a single lightcurve cannot contain an observational gap greater than 10 days. It was also asserted that the total length of each lightcurve must be longer than 1 day. A small number of measurements could not meet these criteria and were removed from analysis. The lightcurves are listed in Table 4.1 and a portion of the January 2008 lightcurve is shown in Figure 4.1.

Table 4.1: Journal of Observations

Start	End	Observations
Oct 25 2003	Oct 31 2003	4641
Nov 19 2003	Nov 30 2003	9867
Jan 20 2004	Jan 21 2004	4432
Oct 15 2004	Oct 19 2004	9066
Nov 09 2004	Nov 10 2004	7189

Continued on Next Page...

Table 4.1 – Continued

Start	End	Observations
Feb 07 2005	Feb 13 2005	6529
Oct 30 2005	Nov 14 2005	17798
Nov 29 2005	Dec 06 2005	11039
Dec 28 2005	Jan 03 2006	8651
Mar 25 2006	Mar 30 2006	3284
Sep 18 2006	Sep 22 2006	6577
Oct 17 2006	Oct 23 2006	16393
Nov 22 2006	Nov 26 2006	16458
Dec 22 2006	Dec 26 2006	4251
Feb 16 2007	Feb 22 2007	20129
Mar 14 2007	Mar 16 2007	9036
Oct 06 2007	Oct 11 2007	7291
Nov 03 2007	Nov 08 2007	28467
Jan 08 2008	Jan 14 2008	30039
Feb 08 2008	Feb 11 2008	6112
Sep 30 2008	Oct 04 2008	17131
Oct 24 2008	Oct 31 2008	9402
Dec 23 2008	Jan 03 2009	13052
Jan 30 2009	Feb 03 2009	11798
Oct 15 2009	Oct 25 2009	15808
Nov 11 2009	Nov 23 2009	22015
Jan 12 2010	Jan 17 2010	7388
Oct 03 2010	Oct 12 2010	8272
Nov 09 2010	Nov 12 2010	13930
Dec 02 2010	Dec 13 2010	18318
Jan 01 2011	Jan 06 2011	11695

Continued on Next Page...

Table 4.1 – Continued

Start	End	Observations
Nov 22 2011	Nov 29 2011	10714
Oct 17 2012	Oct 19 2012	3261
Nov 11 2012	Nov 19 2012	4052

4.2 Fourier Analysis

The January 2008 lightcurve was selected as the best lightcurve for Fourier analysis; it contained the most observations. The Fourier spectrum is shown in Figure 4.2. There appear to be two large, > 10 mma modes near 272 s and 303 s along with smaller amplitude modes near 197 s, 256 s, and 522 s. There are also a few moderate amplitude modes where combination and harmonic modes would be expected. Further inspection reveals that there are likely multiple modes near 197 s, 256 s, 272 s, 303 s, and 522 s. Small regions of the Fourier spectrum near these periods are shown in Figures 4.3, 4.4, 4.5, and 4.6. All modes in the three lightcurves with an estimated signal to noise greater than 6 were identified and assigned a label. Several combination, harmonic, and difference modes were found. The frequencies and amplitudes of the identified modes are shown in Table 4.2.

Analysis of other lightcurves in the dataset indicate that some of the modes' amplitudes and frequencies may be changing from year to year. Some or all of the differences encountered can be attributed to to complications introduced by the poor spectral window (see Figure 4.7). The difference between many of the modes' frequencies is very near the two day alias, which can alter the interpreted amplitudes, phases, and frequencies. This should be taken into consideration as these results are interpreted.

The frequencies of the modes were bootstrapped beginning with the January 2008 lightcurve. The number of modes and observations made merging the entire

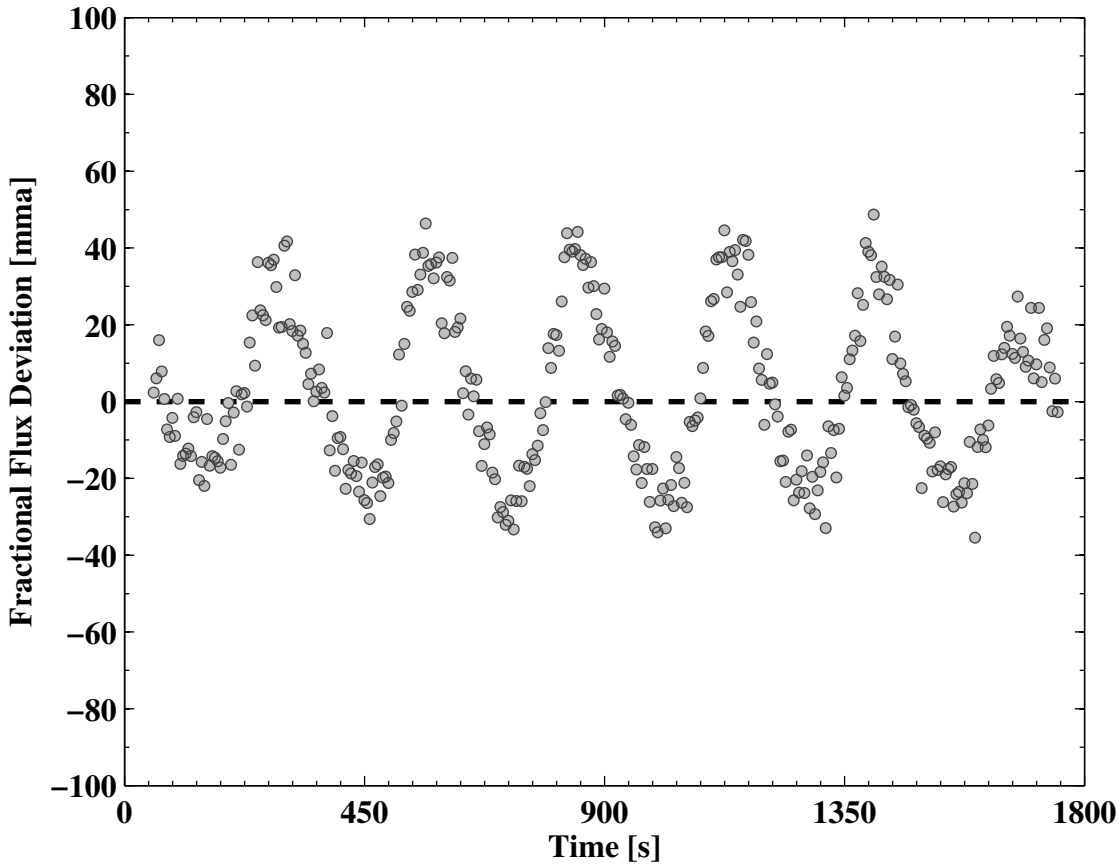

Figure 4.1: A portion of the January 2008 lightcurve.

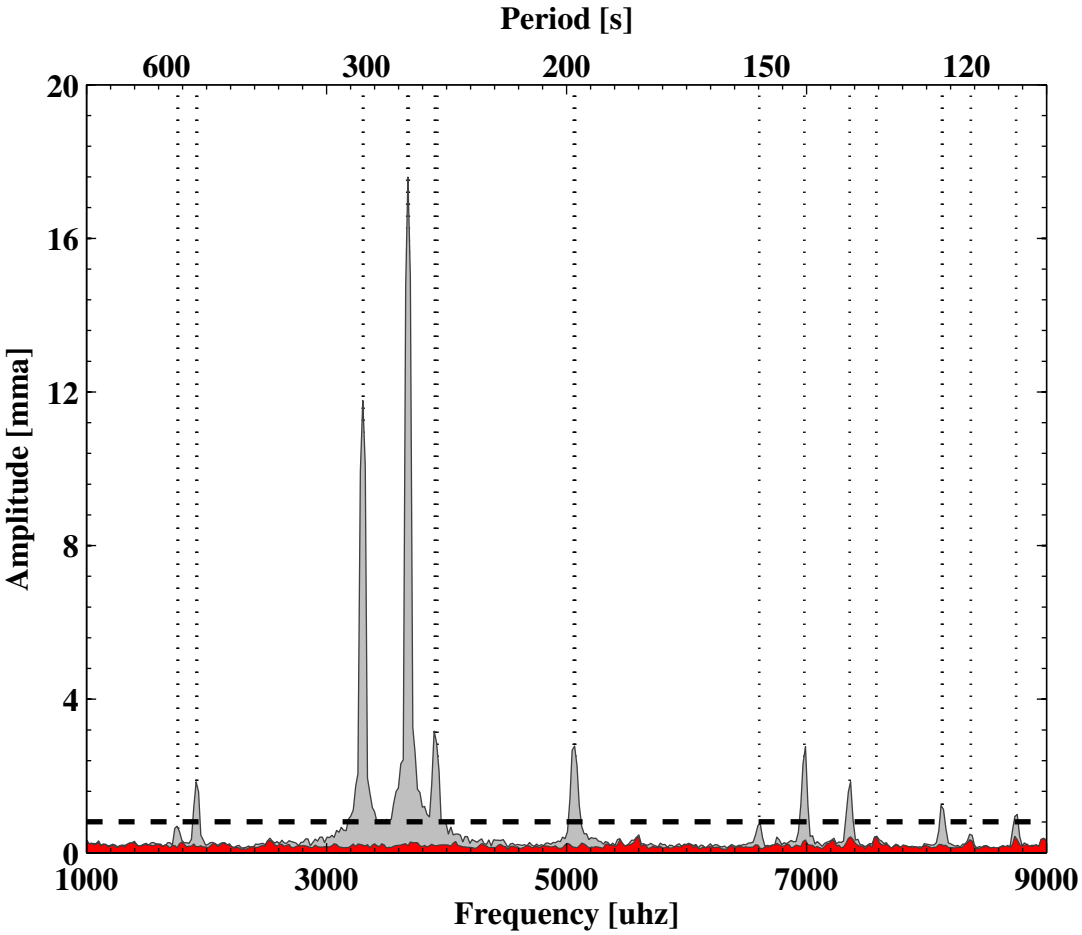

Figure 4.2: Fourier spectrum of the January 2008 lightcurve. The Fourier spectrum of the unmodified lightcurve is shown in grey and the Fourier spectrum of the prewhitened lightcurve is shown in red. The dotted lines indicate frequencies listed in Table 4.2 and the dashed line indicates 6 times the median amplitude of the Fourier spectrum of the unmodified lightcurve.

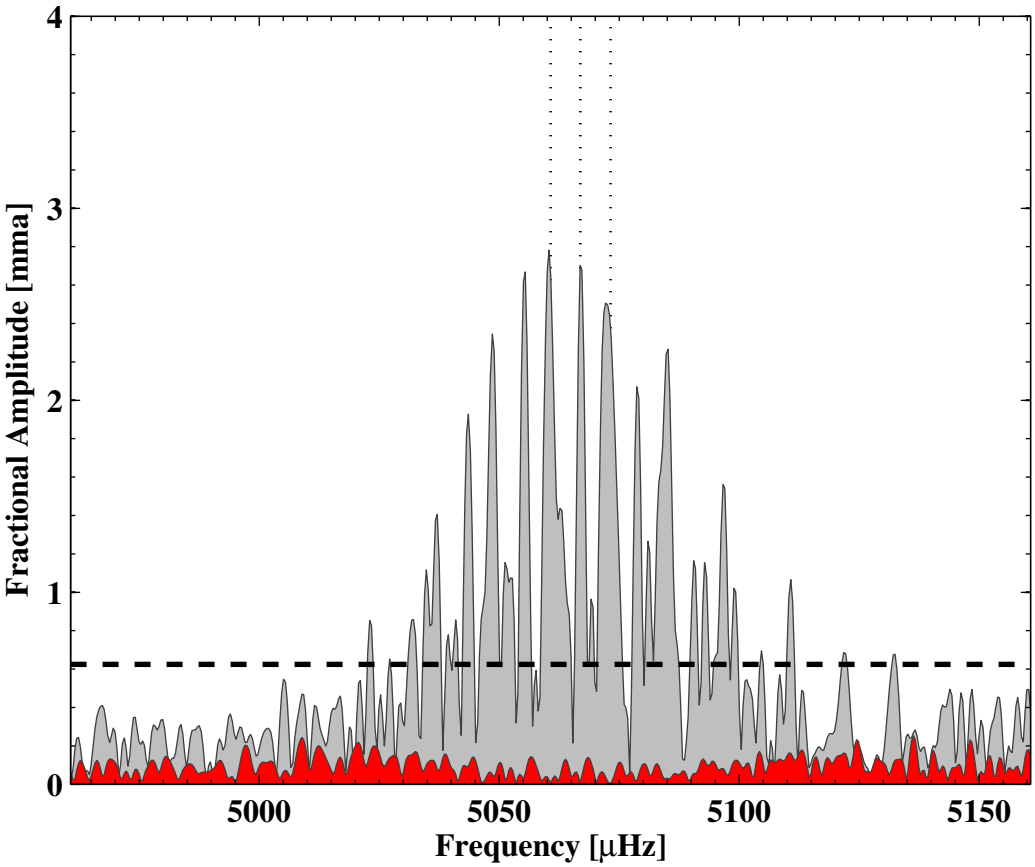

Figure 4.3: Fourier spectrum of the January 2008 lightcurve near the frequency of mode B. The Fourier spectrum of the unmodified lightcurve is shown in grey and the Fourier spectrum of the prewhitened lightcurve is shown in red. The dotted lines indicate frequencies listed in Table 4.2 and the dashed line indicates 6 times the median amplitude of the Fourier spectrum of the unmodified lightcurve. The spectral window is shown in Figure 4.7.

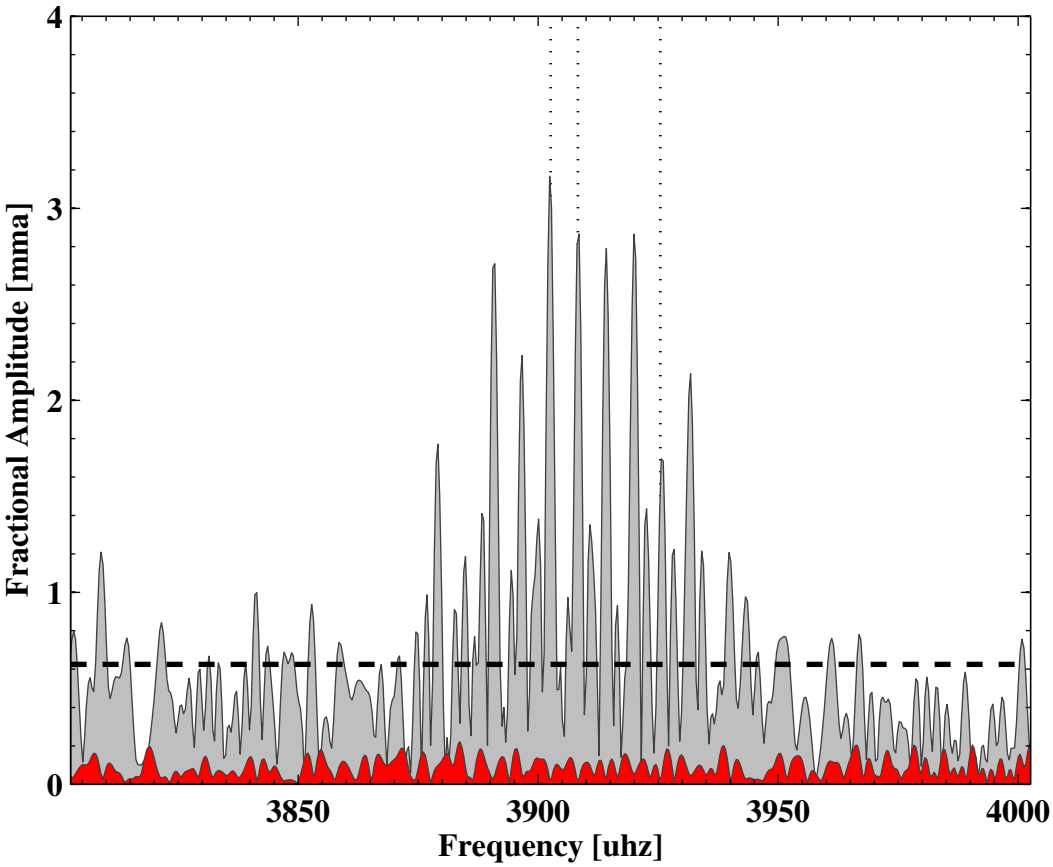

Figure 4.4: Fourier spectrum of the January 2008 lightcurve near the frequency of mode C. The Fourier spectrum of the unmodified lightcurve is shown in grey and the Fourier spectrum of the prewhitened lightcurve is shown in red. The dotted lines indicate frequencies listed in Table 4.2 and the dashed line indicates 6 times the median amplitude of the Fourier spectrum of the unmodified lightcurve. The spectral window is shown in Figure 4.7.

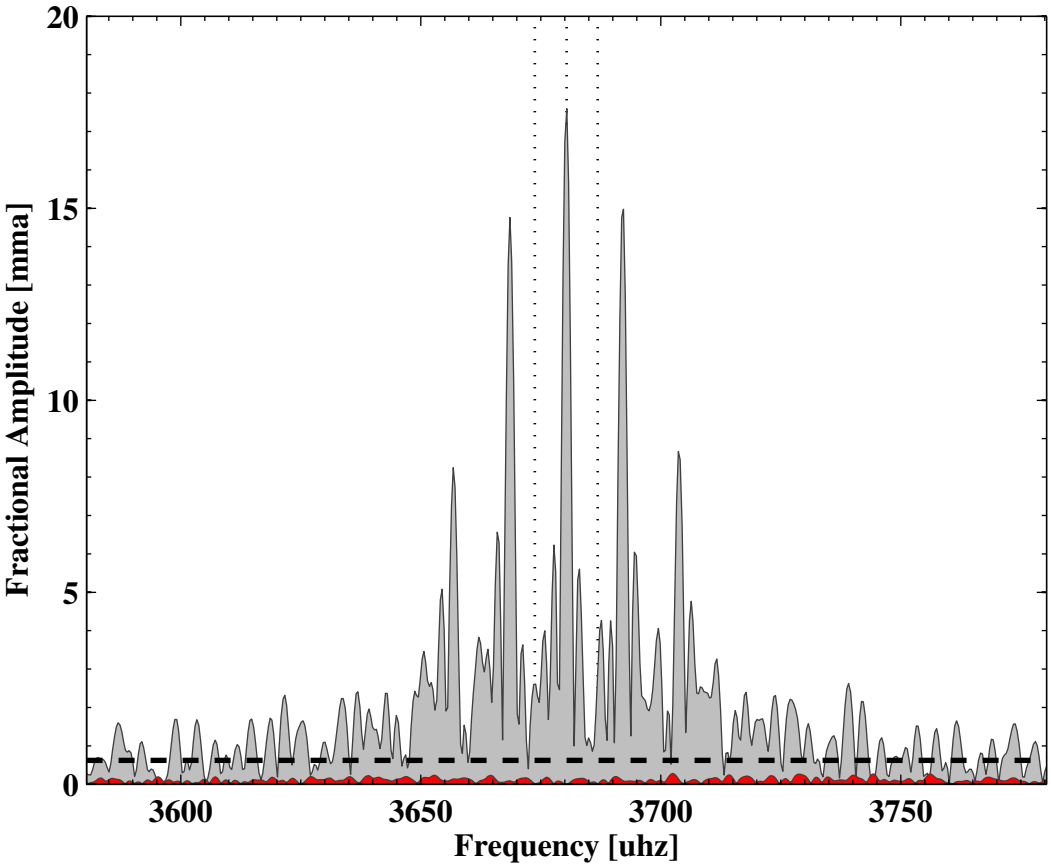

Figure 4.5: Fourier spectrum of the January 2008 lightcurve near the frequency of mode D. The Fourier spectrum of the unmodified lightcurve is shown in grey and the Fourier spectrum of the prewhitened lightcurve is shown in red. The dotted lines indicate frequencies listed in Table 4.2 and the dashed line indicates 6 times the median amplitude of the Fourier spectrum of the unmodified lightcurve. The spectral window is shown in Figure 4.7.

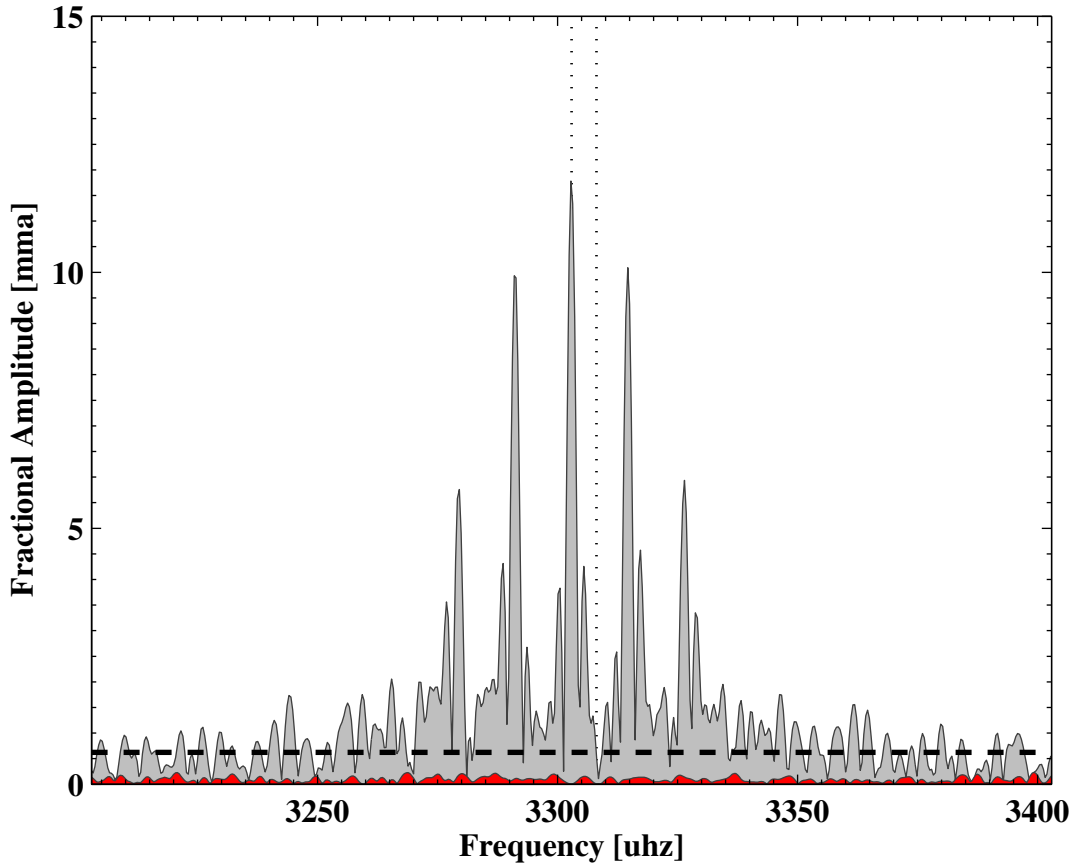

Figure 4.6: Fourier spectrum of the January 2008 lightcurve near the frequency of mode E. The Fourier spectrum of the unmodified lightcurve is shown in grey and the Fourier spectrum of the prewhitened lightcurve is shown in red. The dotted lines indicate frequencies shown in Table 4.2 and the dashed line indicates 6 times the median amplitude of the Fourier spectrum of the unmodified lightcurve. The spectral window is shown in Figure 4.7.

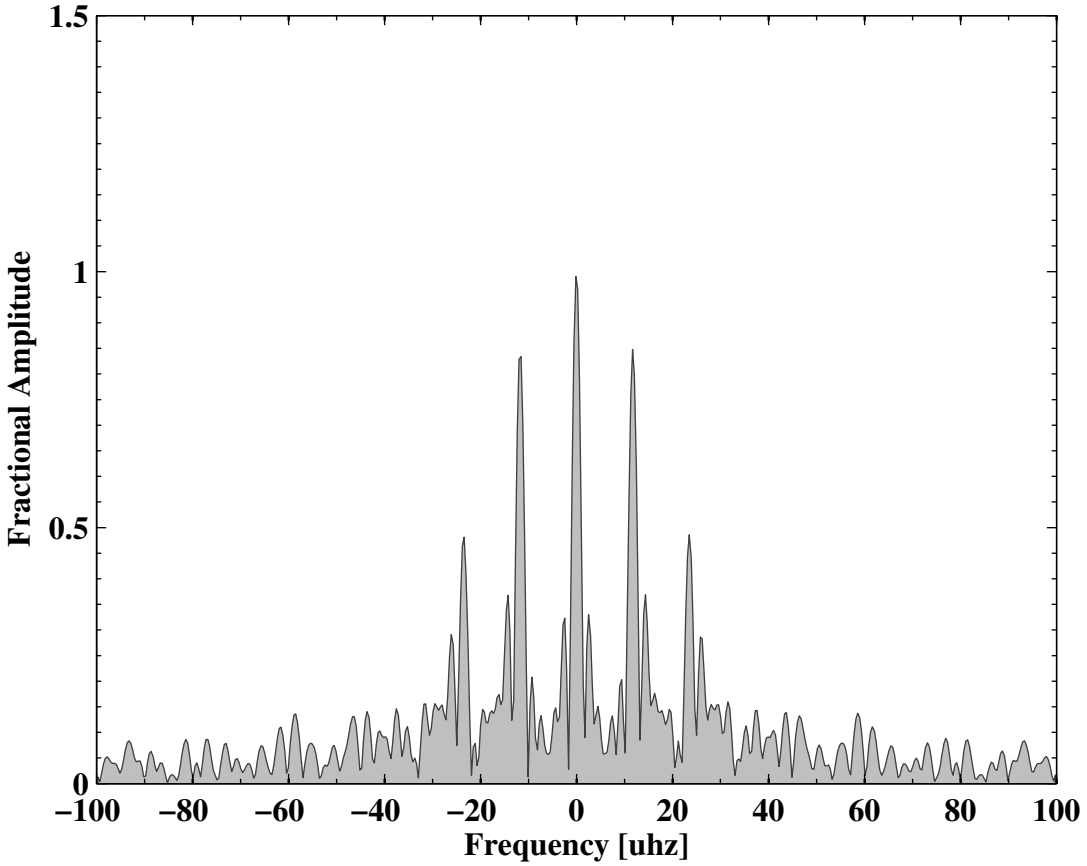

Figure 4.7: Spectral window of the January 2008 lightcurve.

Table 4.2: Frequencies and amplitudes of modes identified in the January 2008 lightcurve.

Label	**Period** $[s]$	**Frequency** $[\mu Hz]$	**Amplitude** $[mma]$
$2D + E_2$	93.8	10663.50(0.08)	0.64(0.06)
$B + D$	114.3	8747.21(0.05)	1.01(0.06)
$B + E_2$	119.5	8369.7(0.1)	0.47(0.06)
A_1	122.9	8133.54(0.08)	0.66(0.06)
A_2	123.0	8127.82(0.04)	1.24(0.06)
$C_3 + D$	131.9	7582.8(0.1)	0.48(0.06)
$2D$	135.9	7360.64(0.03)	1.88(0.06)
$D + E_2$	143.2	6983.17(0.02)	2.77(0.06)
$2E_2$	151.4	6605.81(0.06)	0.84(0.06)
B_-	197.1	5073.22(0.03)	3.0(0.1)
B	197.4	5066.91(0.02)	2.94(0.06)
B_+	197.6	5060.69(0.02)	3.6(0.1)
C_1	254.7	3925.43(0.08)	0.78(0.07)
C_2	255.9	3908.27(0.02)	2.80(0.06)
C_3	256.2	3902.55(0.02)	2.60(0.07)
D_-	271.2	3686.84(0.03)	2.41(0.07)
D	271.7	3680.338(0.003)	17.38(0.06)
D_+	272.2	3673.74(0.05)	1.82(0.07)
E_1	302.3	3308.07(0.08)	0.71(0.06)
E_2	302.8	3302.884(0.005)	11.92(0.06)
$D + E_2 - B_+$	520.1	1922.74(0.07)	0.78(0.06)
$D + E_2 - B$	521.8	1916.58(0.03)	1.86(0.06)
$B - E_2$	566.8	1764.2(0.1)	0.57(0.06)
$B_+ - E_2$	569.0	1757.52(0.08)	0.66(0.06)

dataset computationally cumbersome. Bootstrapping was halted with the merger of all data from the 2006/2007 and 2007/2008 observing seasons. After bootstrapping was completed, the final frequencies of the combination, harmonic, and difference modes were set to the sum of the final frequencies of the parent modes. The results are shown in Table 4.3.

4.3 $O - C$ Analysis

The $O - C$s are shown in Figures 4.8-4.31. The lack of any substantial gaps in the observations eliminates any ambiguity as to the appropriate sequence of $O - C$ values. However, the $O - C$s of modes A_1, $D + E_2 - B_+$, and $D + E_2 - B$ appear mostly consistent with scatter.

The $O - C$s of modes $2D + E_2$, $B + D$, $B + E_2$, $C_3 + D$, $D + E_2$, $2E_2$, B_-, B, B_+, C_2, $B - E_2$, $B_+ - E_2$ appear consistent with a line or parabola. A parabola was fit to each $O - C$ and the residuals were used to calculate upper limits on the amplitudes of any unmodelled periodic variations. The results are shown in Table 4.4 and the fits are shown in Figures 4.8, 4.9, 4.10, 4.13, 4.15, 4.16, 4.17, 4.18, 4.19, 4.21, 4.30, and 4.31. The fit parameters are plotted in Figure 4.32.

The $O - C$s of modes A_2, $2D$, C_1, C_3, D_-, D, D_+, E_1, and E_2 appear to vary periodically or are inconsistent with a parabolic model. These $O - C$s were further split into two groups. The first group, consisting of A_2, $2D$, D, and E_2 appear to vary sinusoidally with similar periods. A sinusoid plus a parabola was fit to each $O - C$. Reduced χ^2 as a function of model period is shown in Figure 4.33. A period of around 4 years is strongly preferred and the total minimum of χ^2 occurs at 4.15 years. For each of these modes, the period was fixed to 4.15 years for the final fit. The results are shown in Table 4.4 and the fits are shown in Figures 4.12, 4.14, 4.24, and 4.27. The fit parameters are plotted in Figure 4.32. The second group consists of modes C_1, C_3, D_-, D^+, and E_1. For these modes $O - C$ does not appear consistent with a parabolic or sinusoidal plus parabolic variation. The $O - C$s were fit with a sinusoid plus a parabola and reduced χ^2 as a function of model period is shown in Figure 4.34. Clearly, there

Table 4.3: Results of bootstrapping the frequencies. The parameter σ_f is the error is the frequency measurement and P_B is the probability that bootstrapping was successful if the frequencies remained constant.

Label	**Period** $[s]$	**Frequency** $[\mu Hz]$	σ_f^{-1} $[yr]$	P_B
$2D+E$	93.8	10663.55893(0.00005)	651	1.00
$B+D$	114.3	8747.2652(0.0001)	268	1.00
$B+E$	119.5	8369.8177(0.0001)	264	1.00
$A1$	122.9	8133.5584(0.0009)	36	0.99
A	123.0	8127.8431(0.0003)	117	1.00
$C+D$	131.8	7588.6203(0.0001)	262	1.00
$2D$	135.9	7360.67090(0.00004)	800	1.00
$D+E$	143.2	6983.22348(0.00003)	917	1.00
$2E$	151.4	6605.77607(0.00006)	560	1.00
B^-	197.1	5073.1685(0.0001)	278	1.00
B	197.4	5066.9297(0.0001)	272	1.00
B^+	197.6	5060.68860(0.00009)	354	1.00
$C1$	254.7	3925.5302(0.0004)	90	1.00
C	255.9	3908.2849(0.0001)	265	1.00
$C2$	256.2	3902.6431(0.0002)	195	1.00
D^-	271.2	3686.9314(0.0001)	238	1.00
D	271.7	3680.33545(0.00002)	1600	1.00
D^+	272.2	3673.7335(0.0001)	214	1.00
$E1$	302.3	3308.2501(0.0005)	64	1.00
E	302.8	3302.88803(0.00003)	1120	1.00
$D+E-B^+$	520.1	1922.5349(0.0001)	330	1.00
$D+E-B$	521.8	1916.2938(0.0001)	261	1.00
$B-E$	566.9	1764.0417(0.0001)	264	1.00
B^+-E	568.9	1757.80056(0.00009)	337	1.00

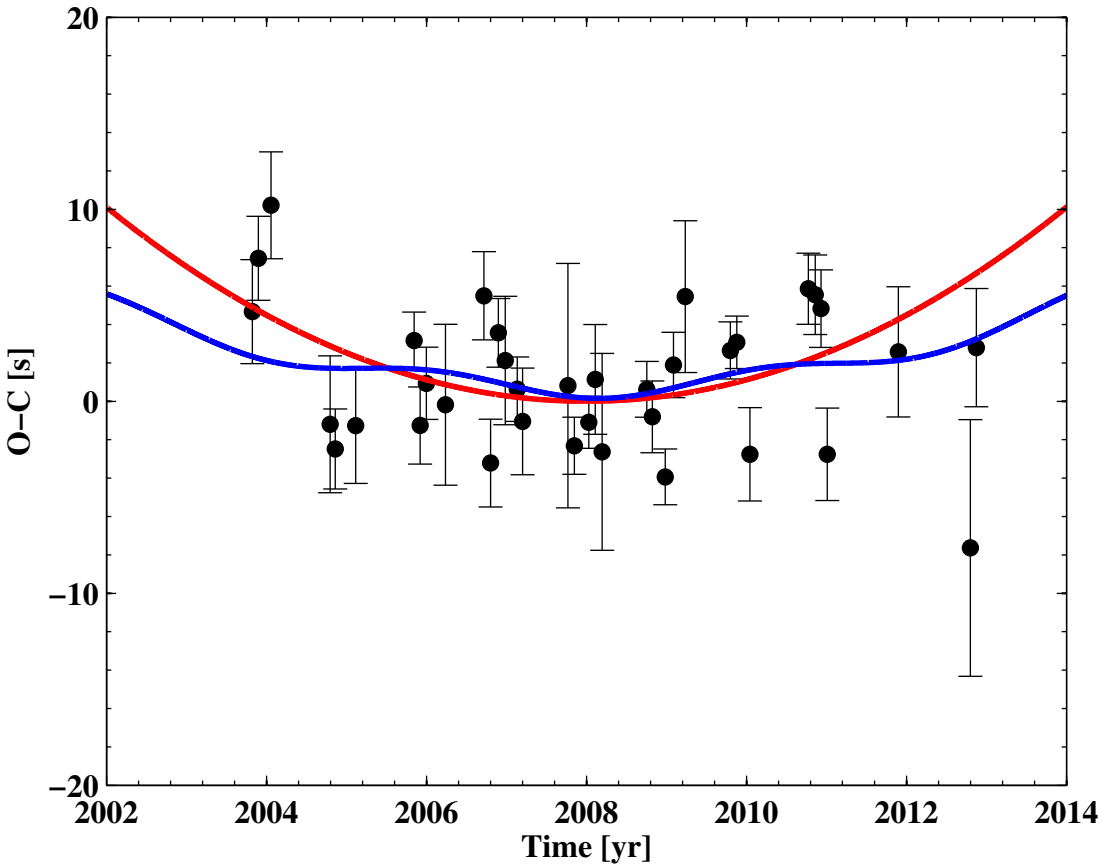

Figure 4.8: $O-C$ of mode $2D+E_2$ with a linear trend removed. The red line is the best fit of a parabola. The blue line is the $O-C$ predicted by fits to the $O-C$s of the parent modes D and E_2.

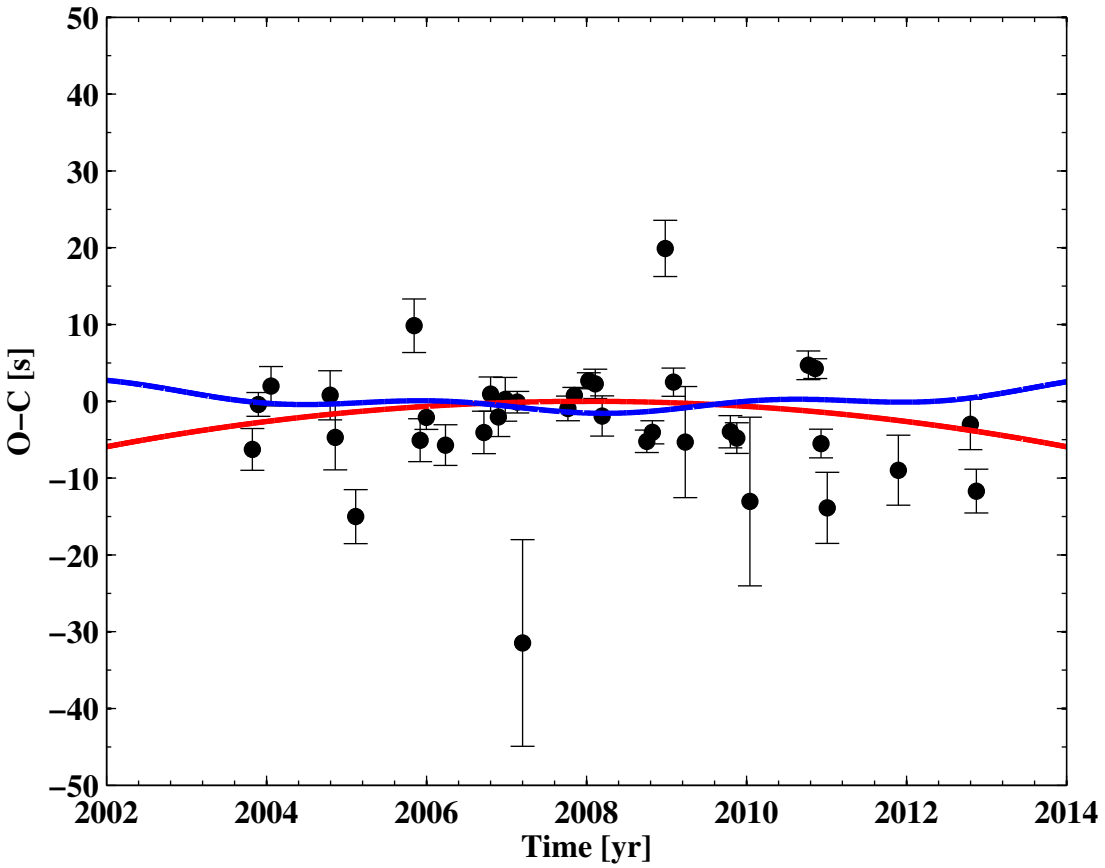

Figure 4.9: $O - C$ of mode $B + D$ with a linear trend removed. The red line is the best fit of a parabola. The blue line is the $O - C$ predicted by fits to the $O - C$s of the parent modes B and D.

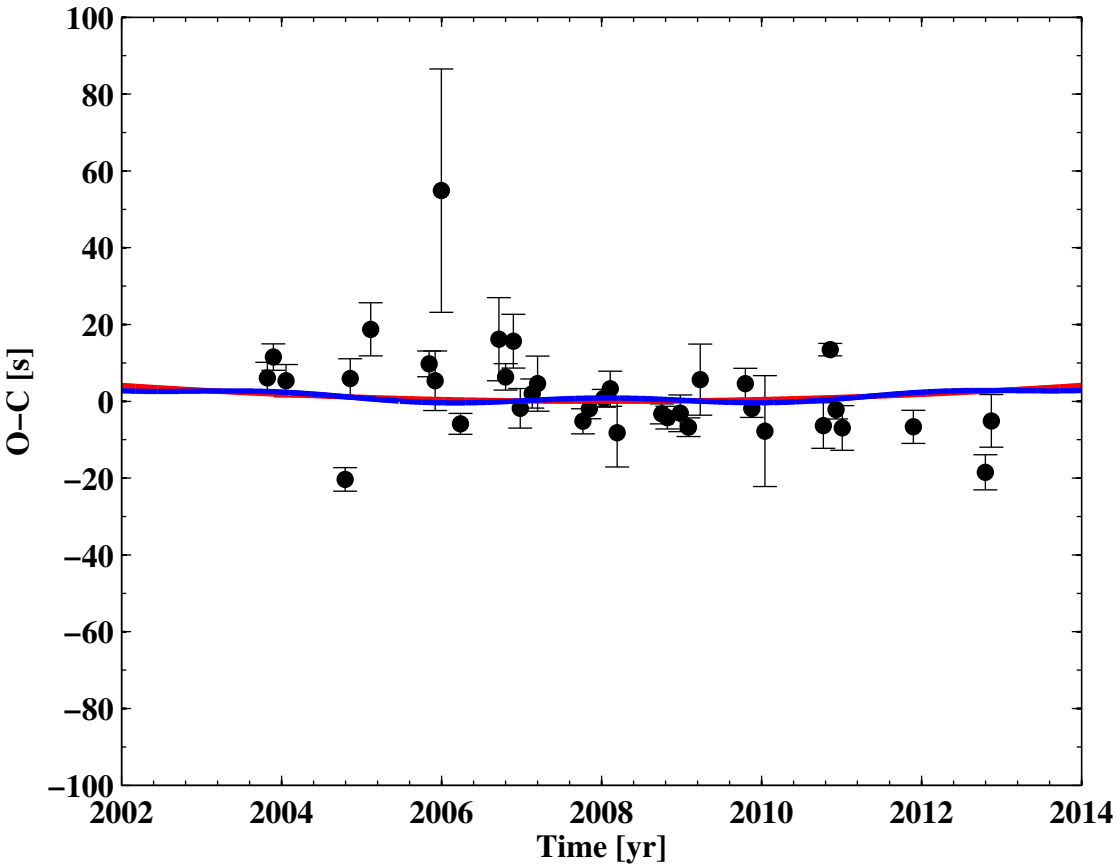

Figure 4.10: $O-C$ of mode $B+E_2$ with a linear trend removed. The red line is the best fit of a parabola. The blue line is the $O-C$ predicted by fits to the $O-C$s of the parent modes B and E_2.

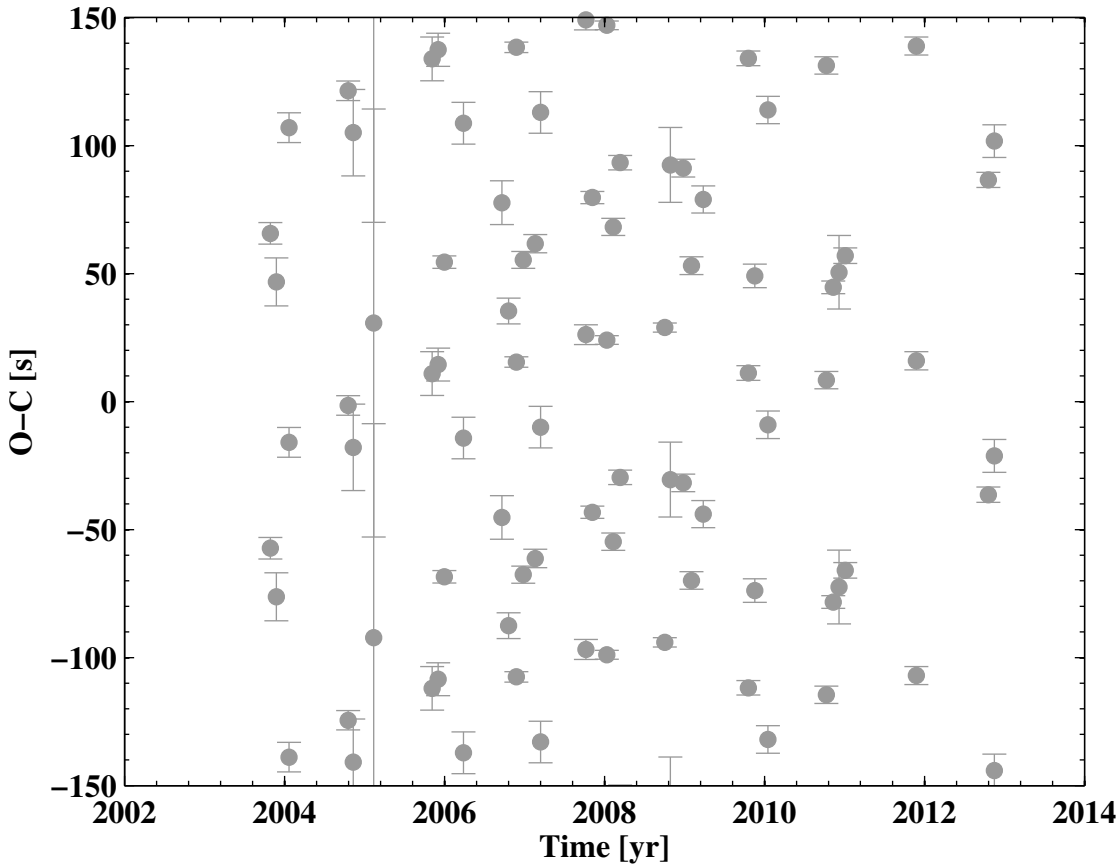

Figure 4.11: $O - C$ of mode A_1.

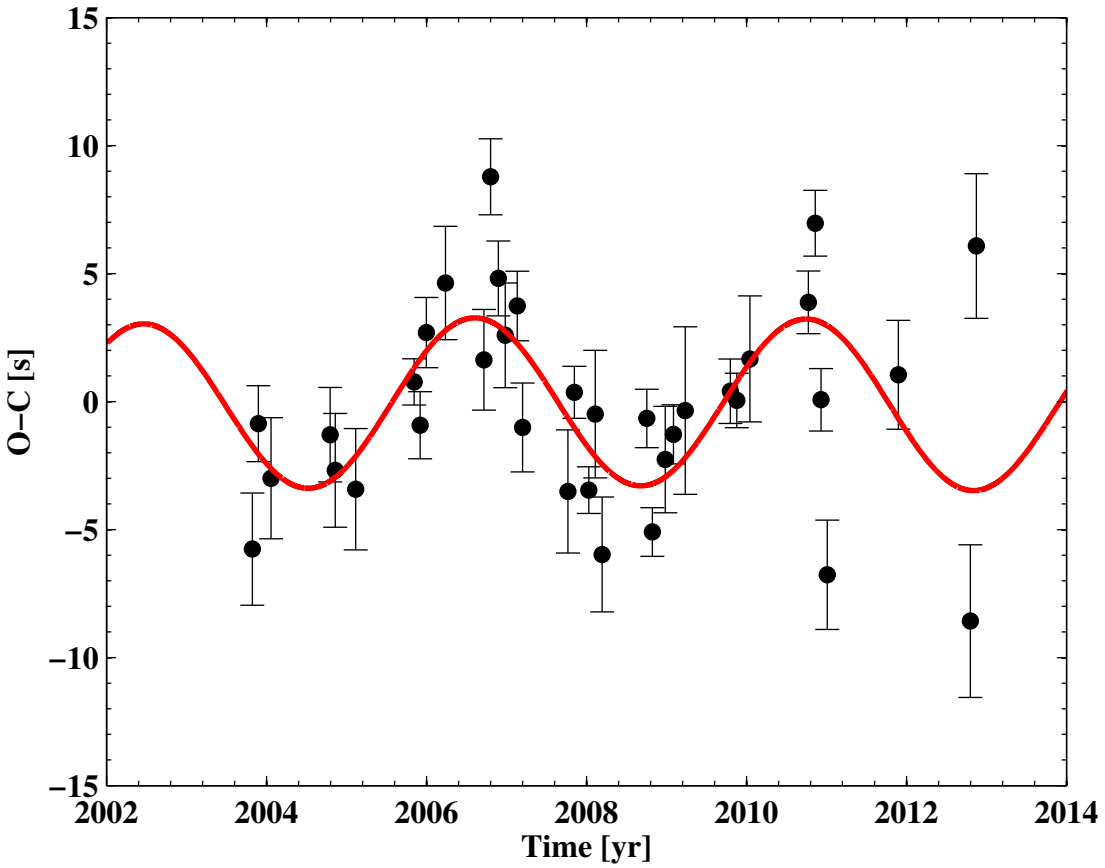

Figure 4.12: $O - C$ of mode A_2 with a linear trend removed. The red line is the best fit of a sinusoid plus parabola.

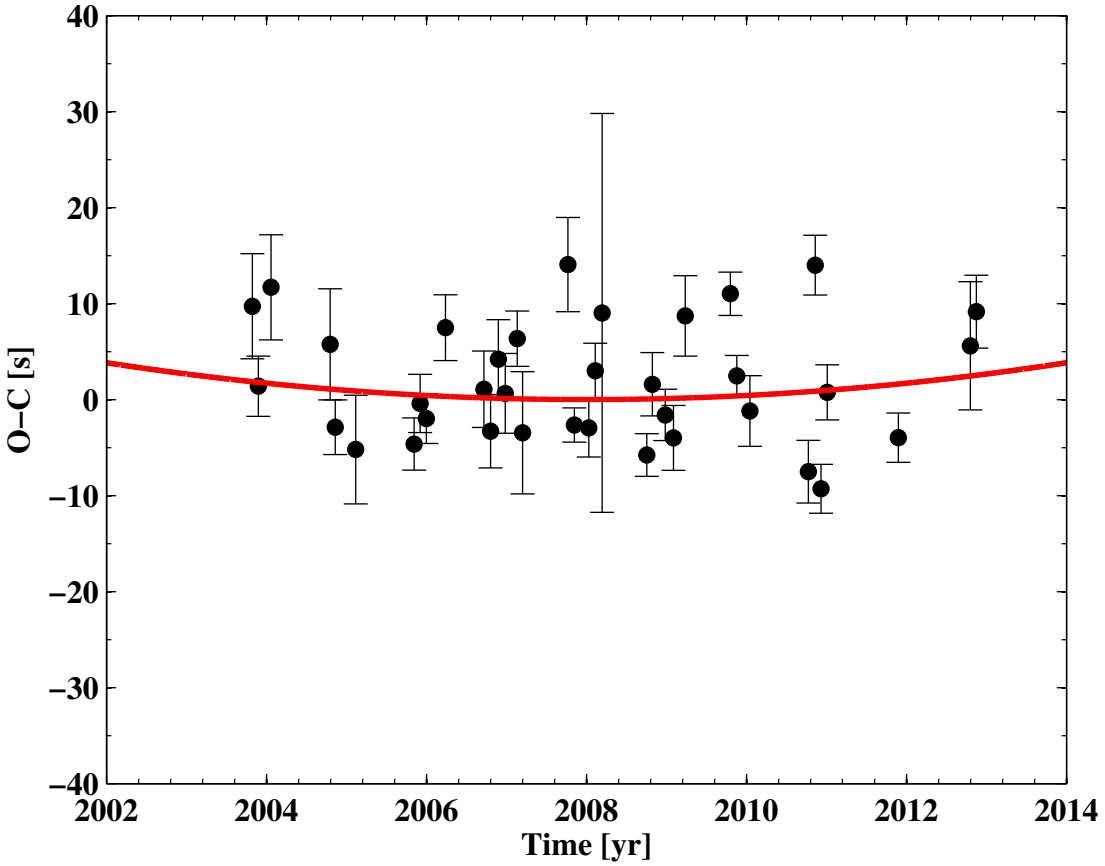

Figure 4.13: $O - C$ of mode $C_3 + D$ with a linear trend removed. The red line is the best fit of a parabola.

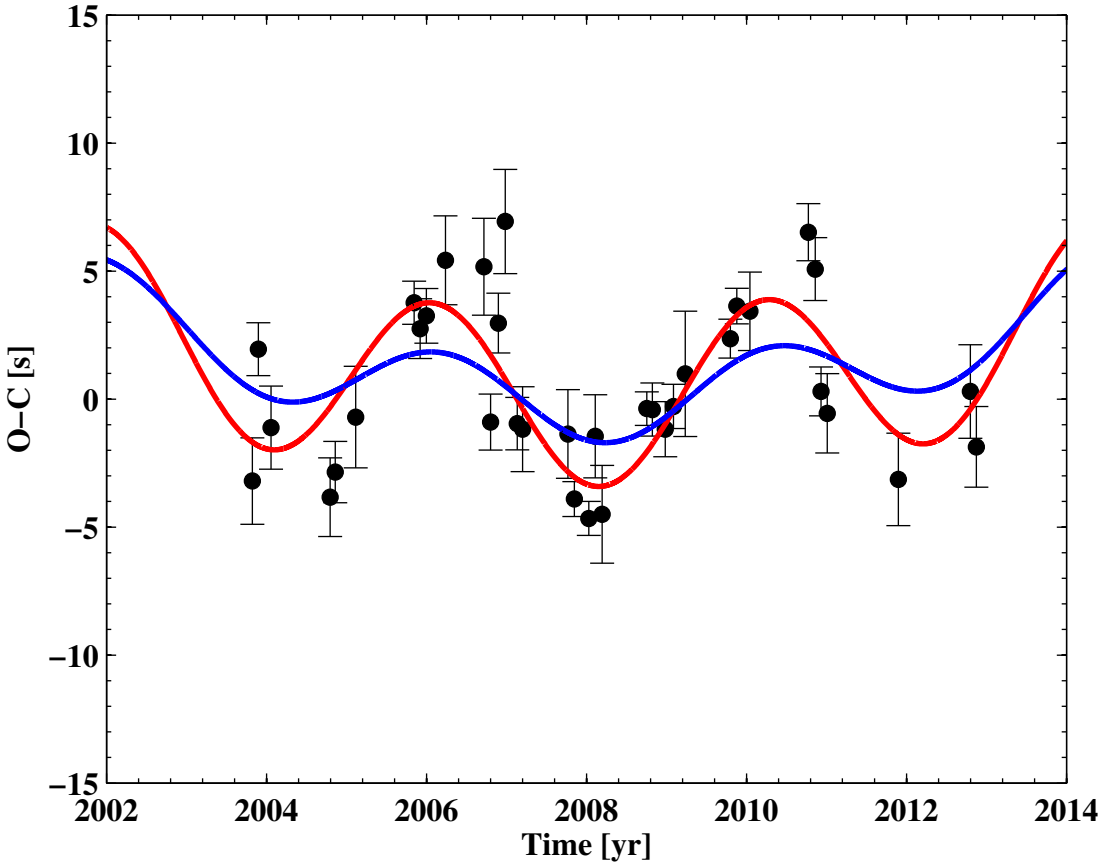

Figure 4.14: $O-C$ of mode $2D$ with a linear trend removed. The red line is the best fit of a parabola plus sinusoid. The blue line is the $O-C$ predicted by fits to the $O-C$s of mode D.

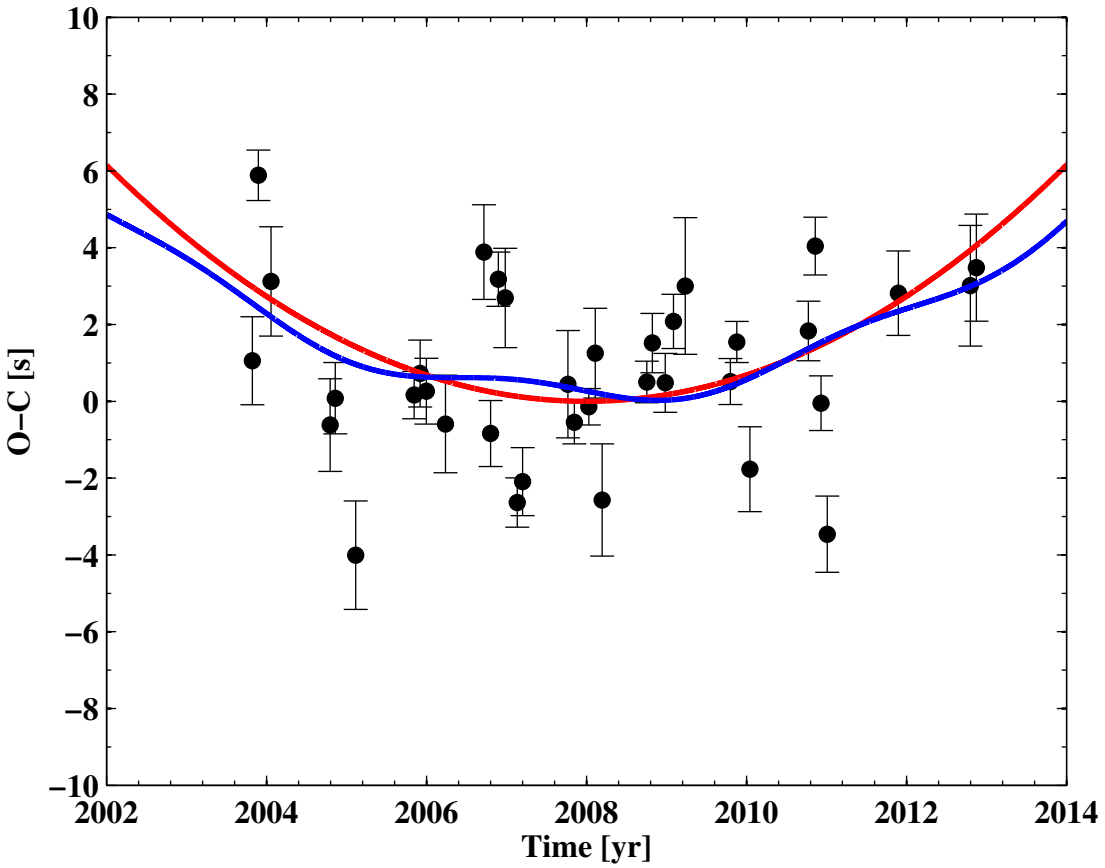

Figure 4.15: $O-C$ of mode $D+E_2$ with a linear trend removed. The red line is the best fit of a parabola. The blue line is the $O-C$ predicted by fits to the $O-C$s of the parent modes D and E_2.

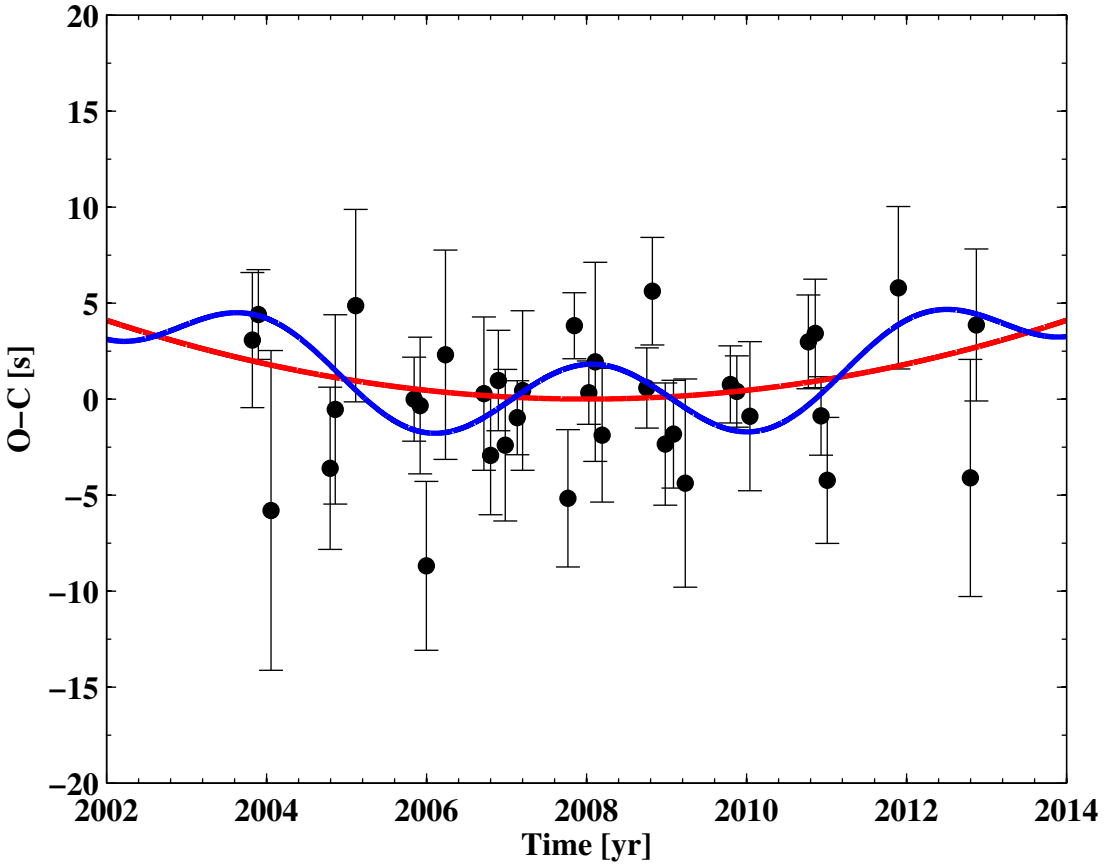

Figure 4.16: $O - C$ of mode $2E_2$ with a linear trend removed. The red line is the best fit of a parabola. The blue line is the $O - C$ predicted by fits to the $O - C$s of mode E.

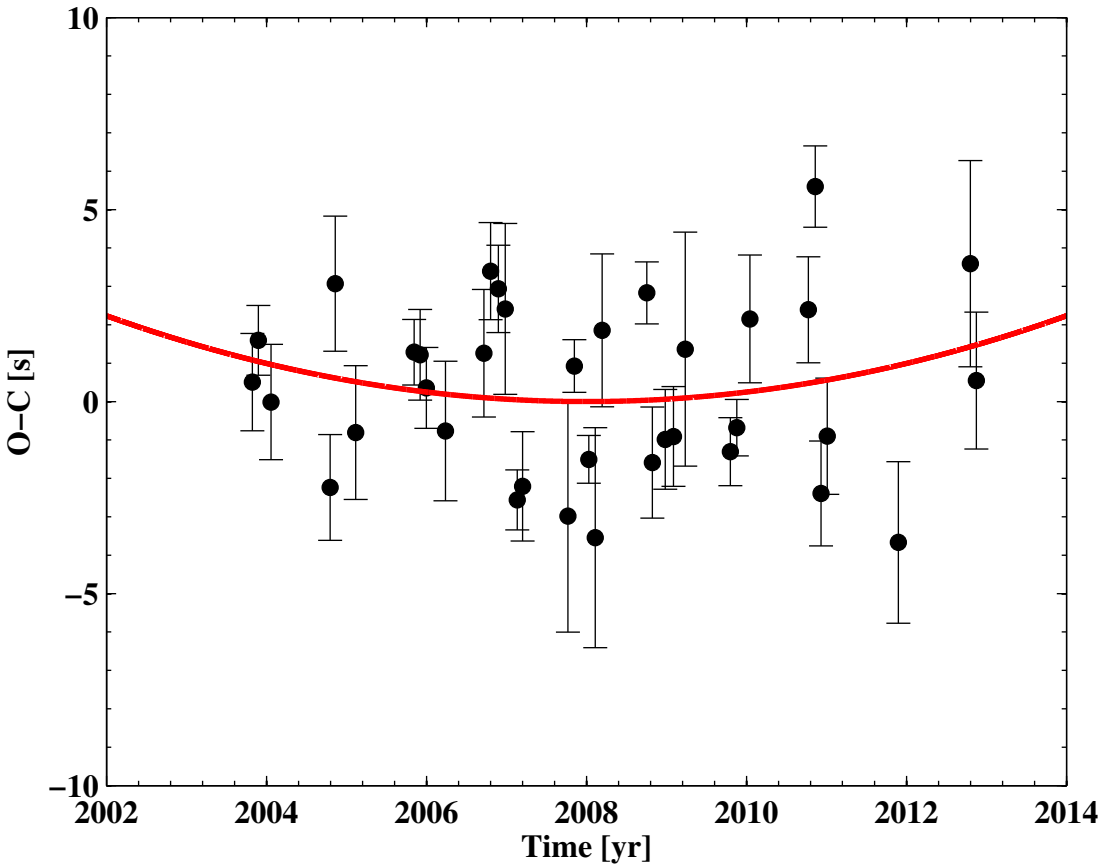

Figure 4.17: $O-C$ of mode B_- with a linear trend removed. The red line is the best fit of a parabola.

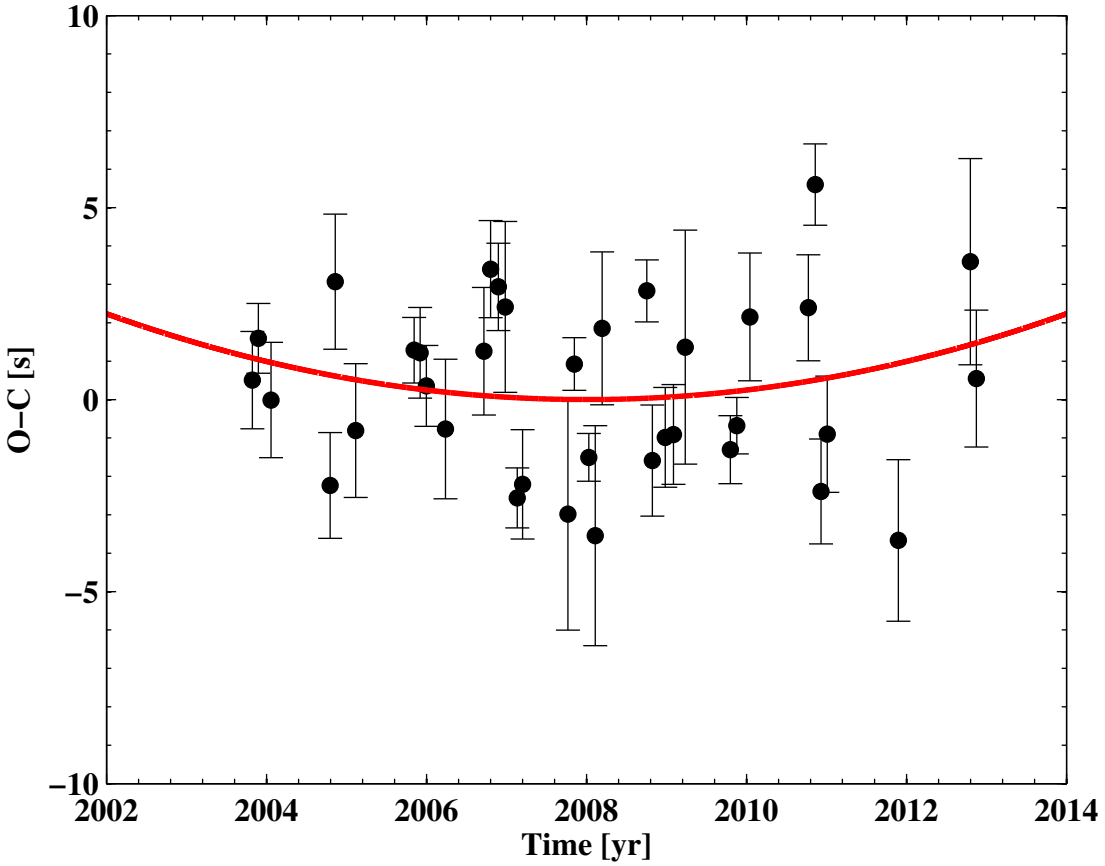

Figure 4.18: $O - C$ of mode B with a linear trend removed. The red line is the best fit of a parabola.

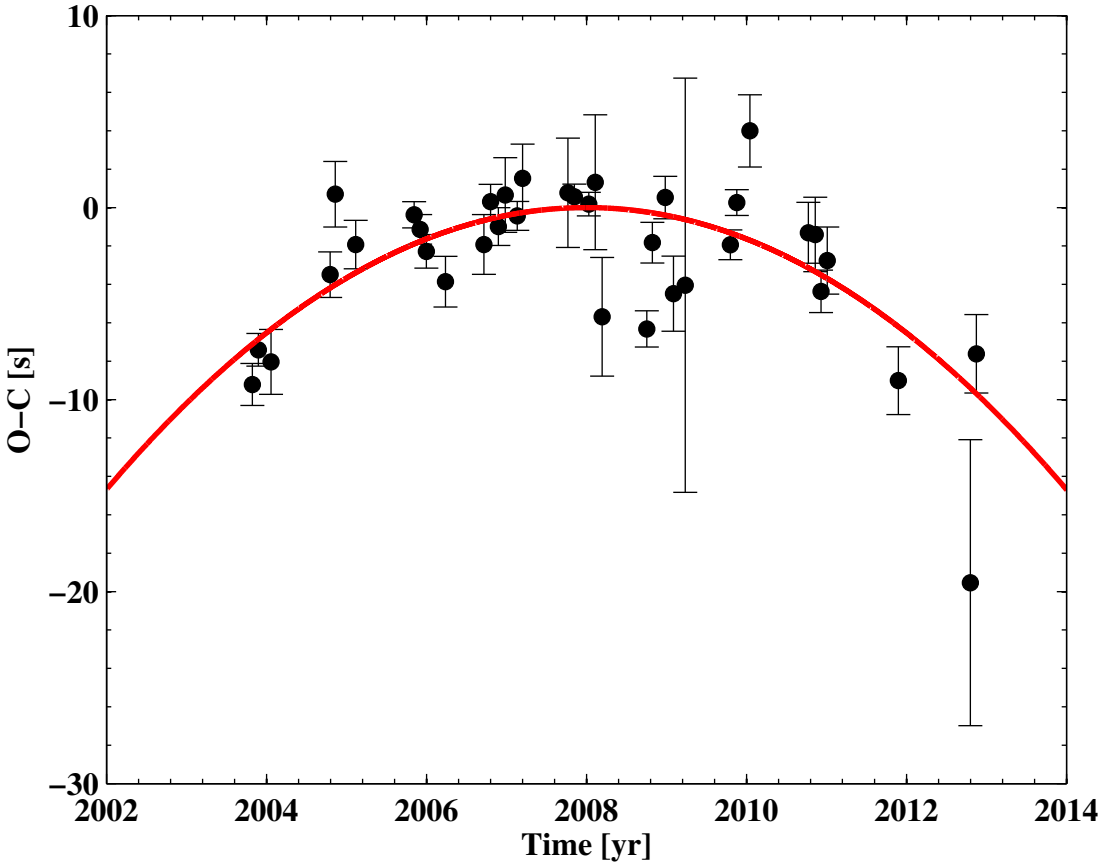

Figure 4.19: $O-C$ of mode B_+ with a linear trend removed. The red line is the best fit of a parabola.

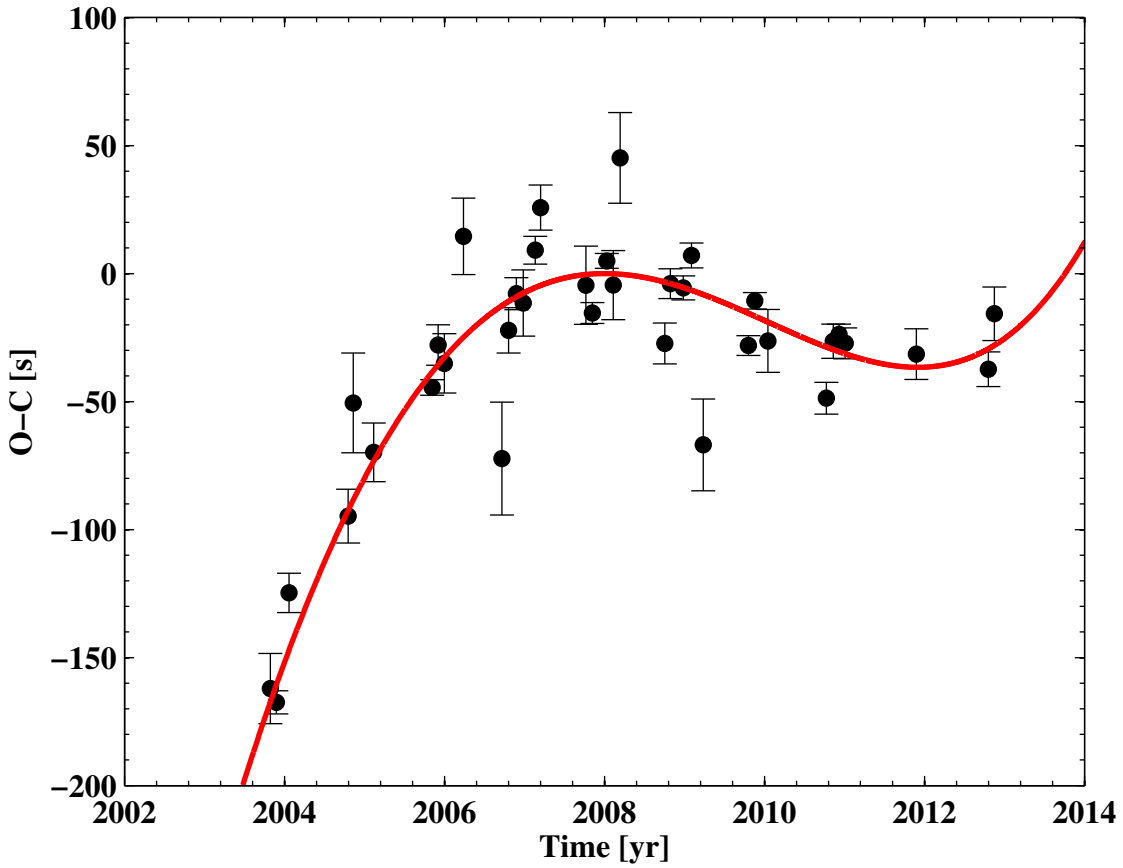

Figure 4.20: $O - C$ of mode C_1 with a linear trend removed. The red line is the best fit of a 5th order polynomial.

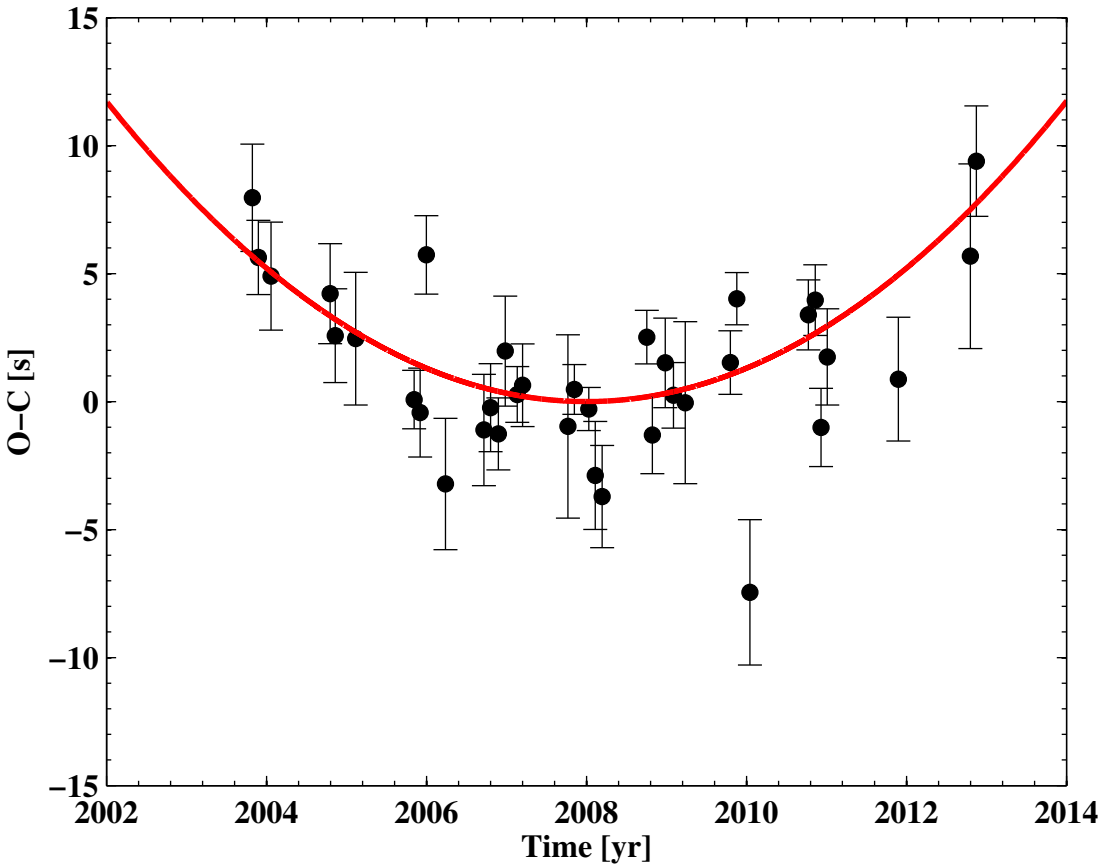

Figure 4.21: $O - C$ of mode C_2 with a linear trend removed. The red line is the best fit of a parabola.

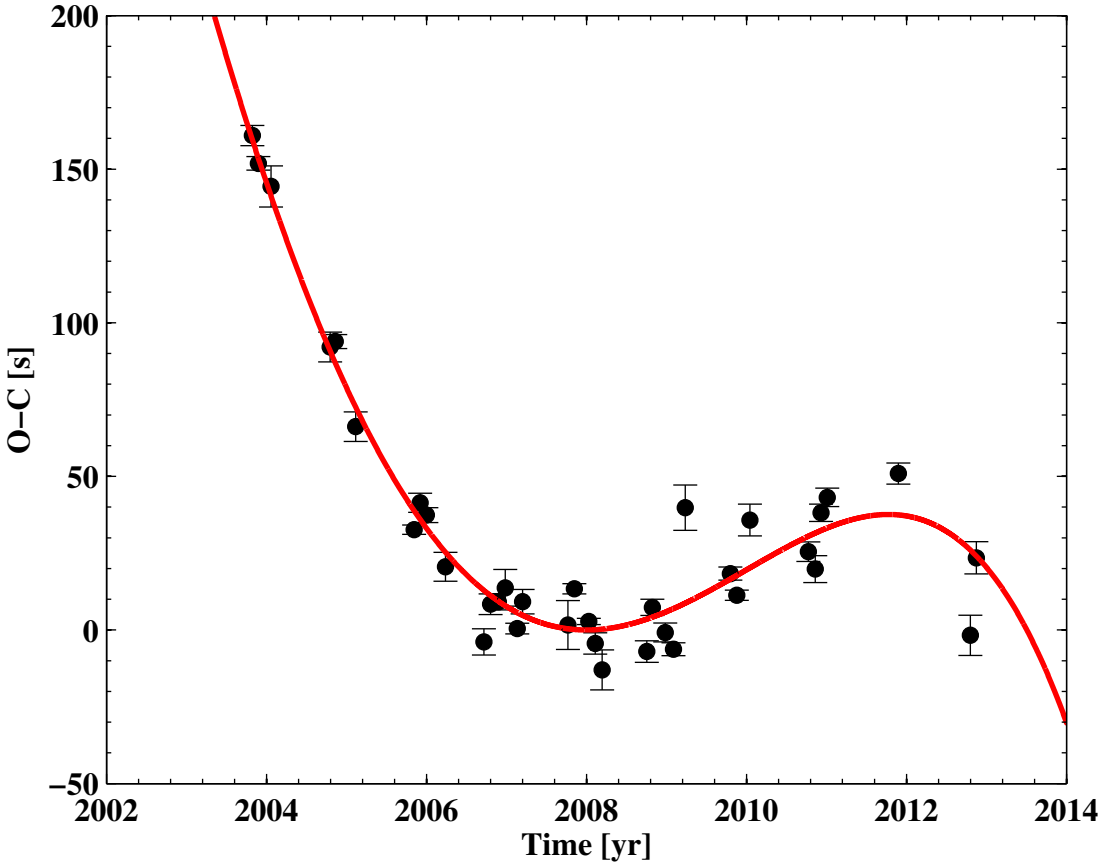

Figure 4.22: $O - C$ of mode C_3 with a linear trend removed. The red line is the best fit of a 5th order polynomial.

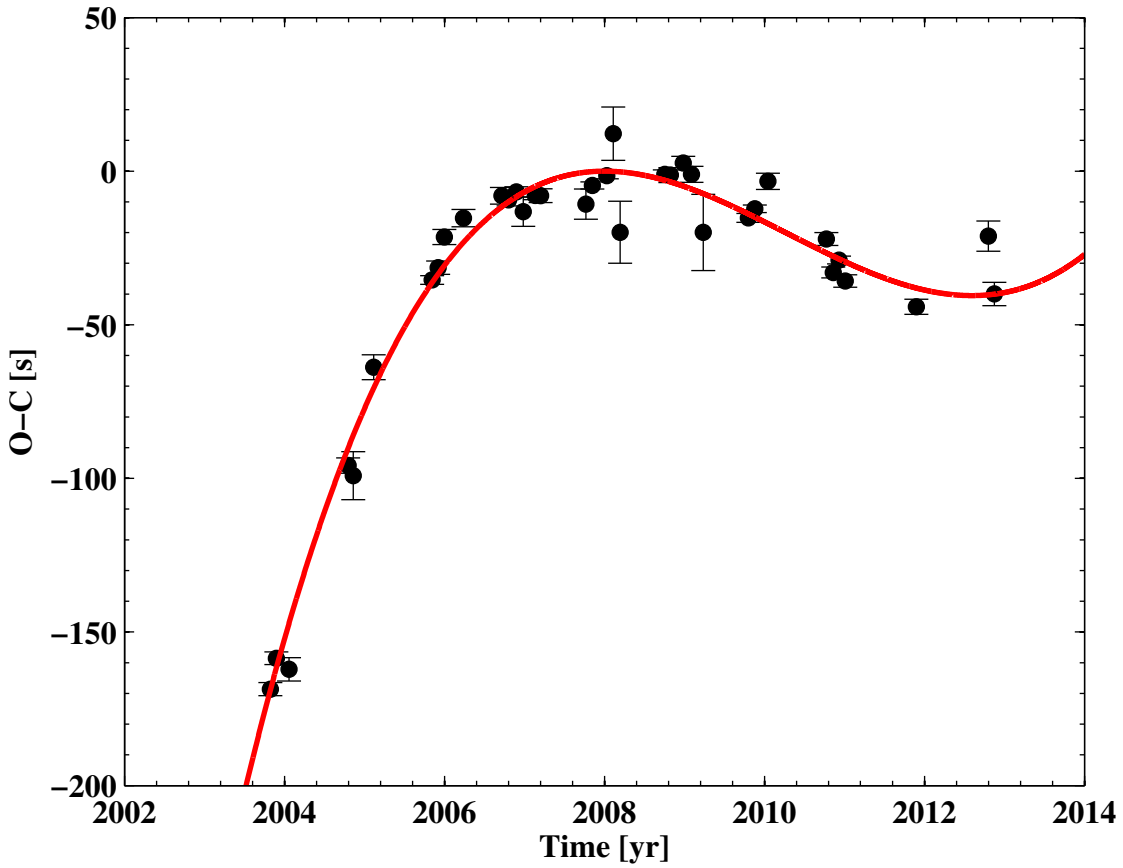

Figure 4.23: $O - C$ of mode D_- with a linear trend removed. The red line is the best fit of a 5th order polynomial.

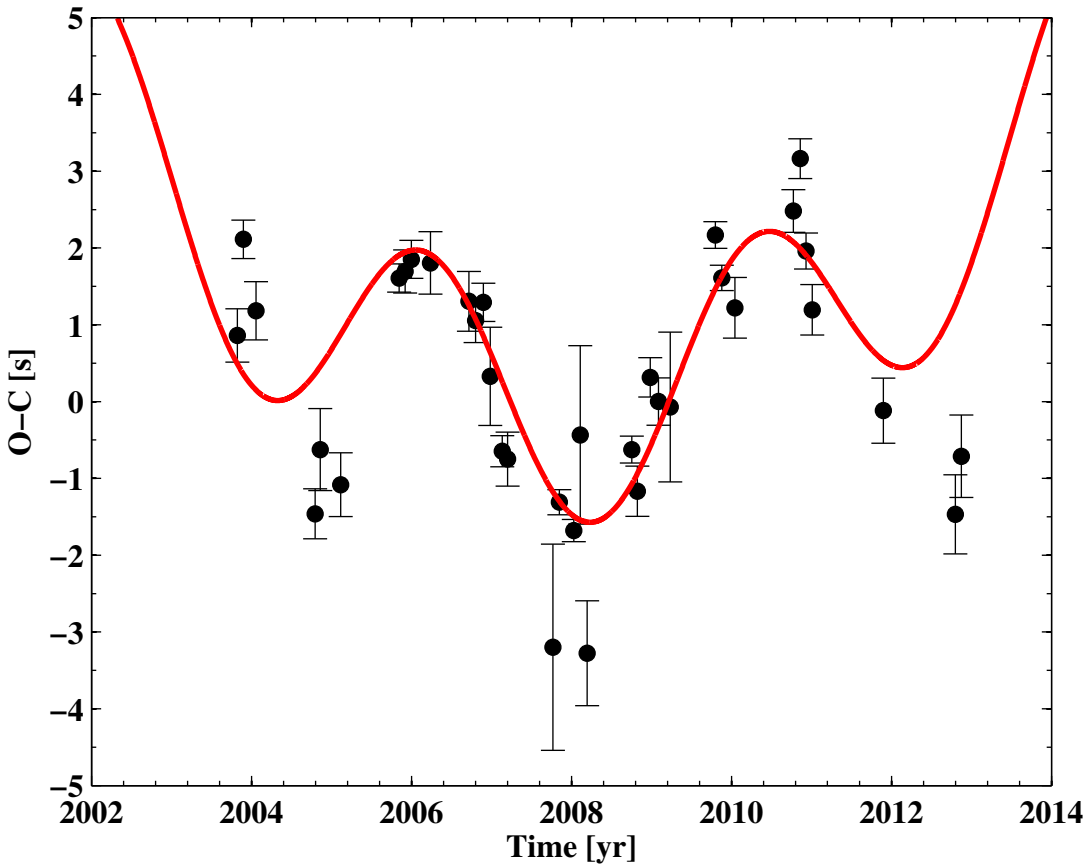

Figure 4.24: $O - C$ of mode D with a linear trend removed. The red line is the best fit of a parabola plus sinusoid.

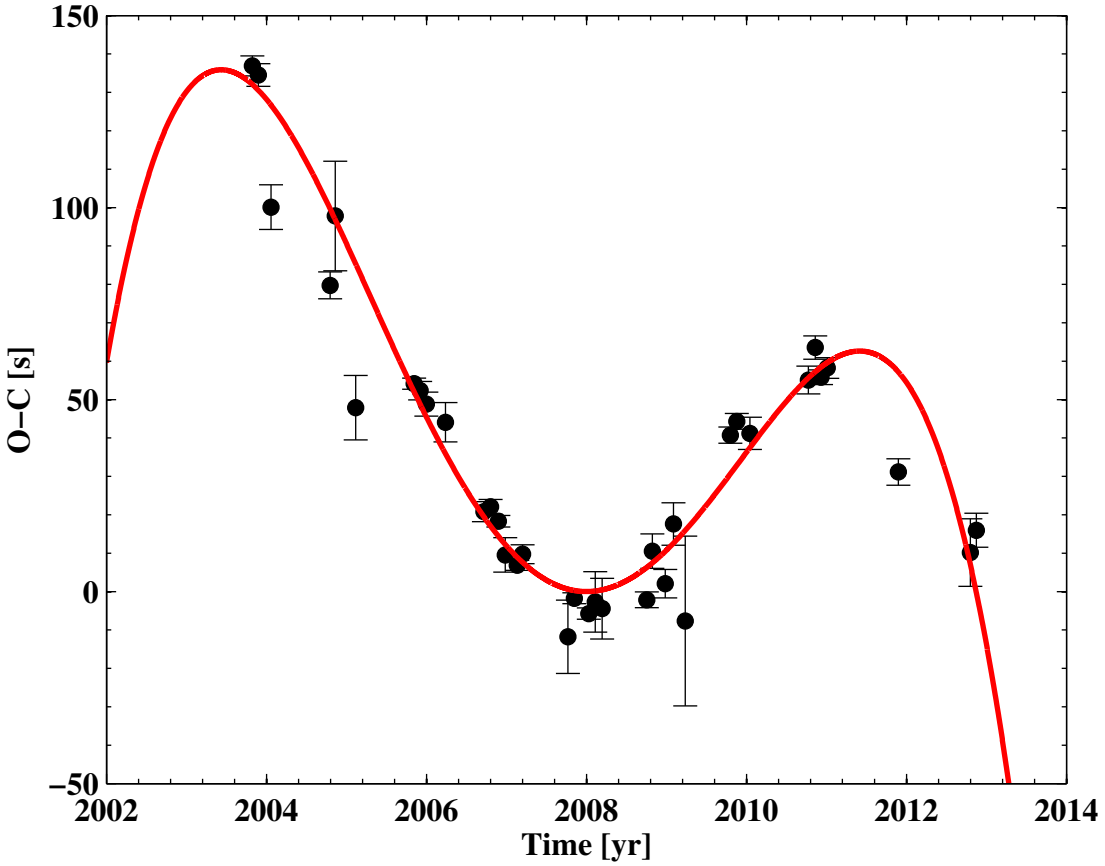

Figure 4.25: $O - C$ of mode D_+ with a linear trend removed. The red line is the best fit of a 5th order polynomial.

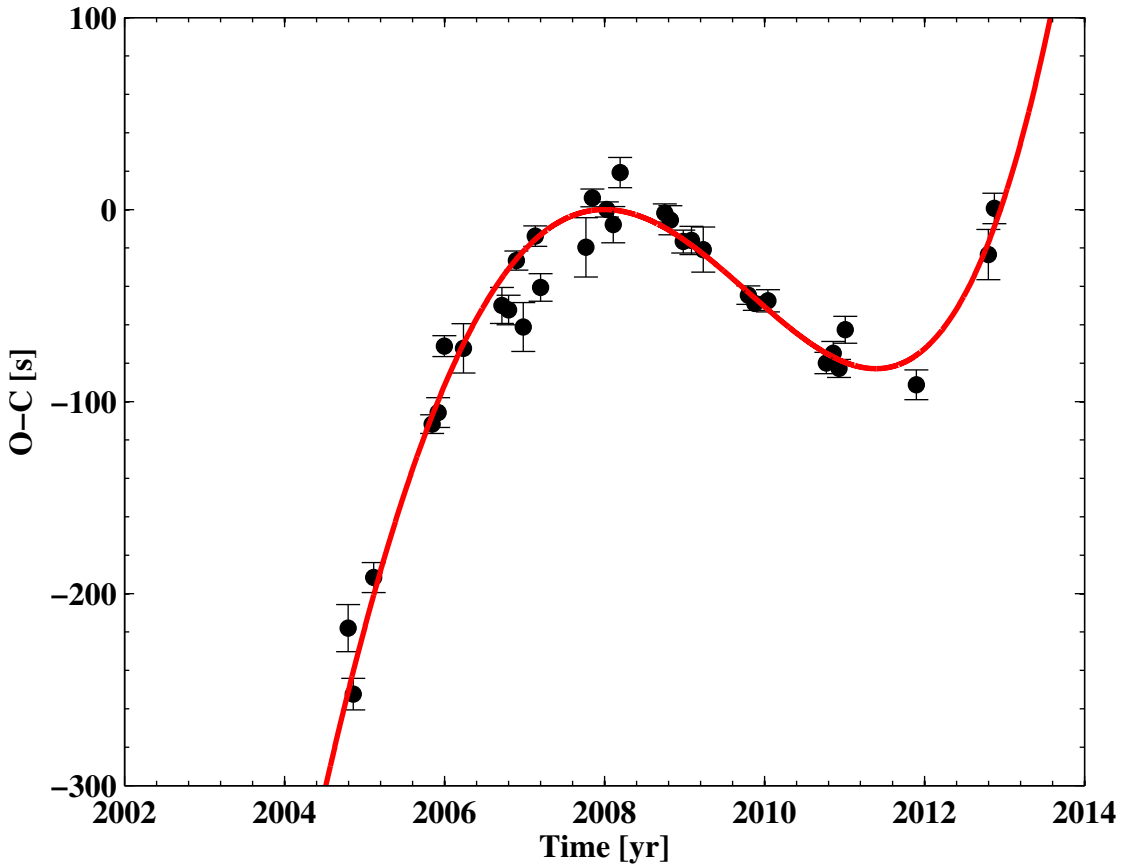

Figure 4.26: $O - C$ of mode E_1 with a linear trend removed. The red line is the best fit of a 5th order polynomial.

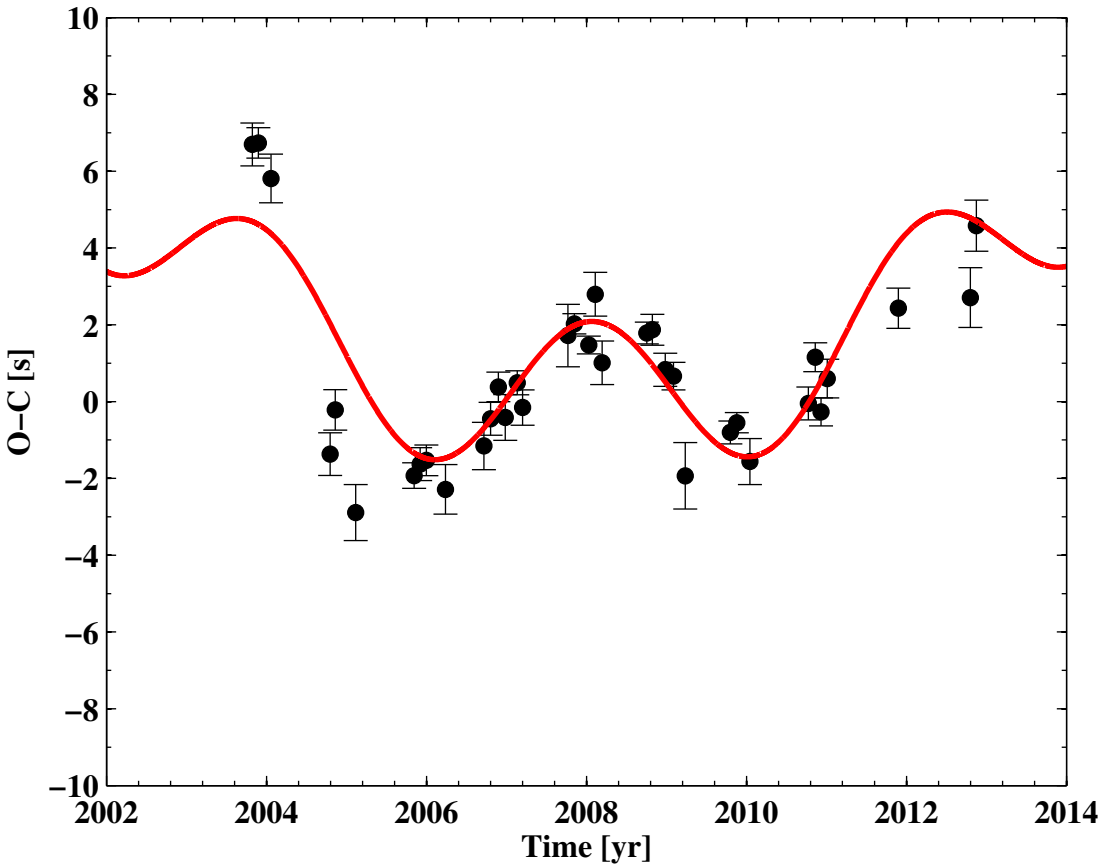

Figure 4.27: $O - C$ of mode E_2 with a linear trend removed. The red line is the best fit of a parabola plus sinusoid.

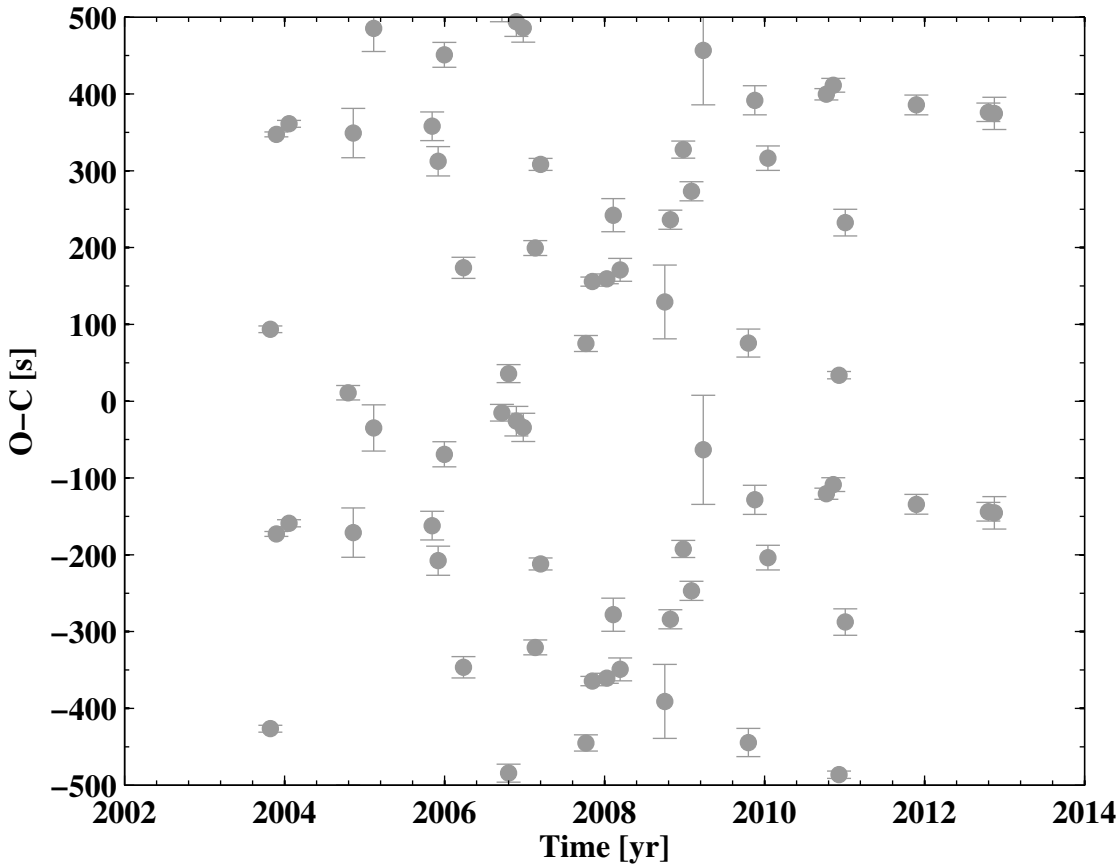

Figure 4.28: $O - C$ of mode $D + E_2 - B_+$.

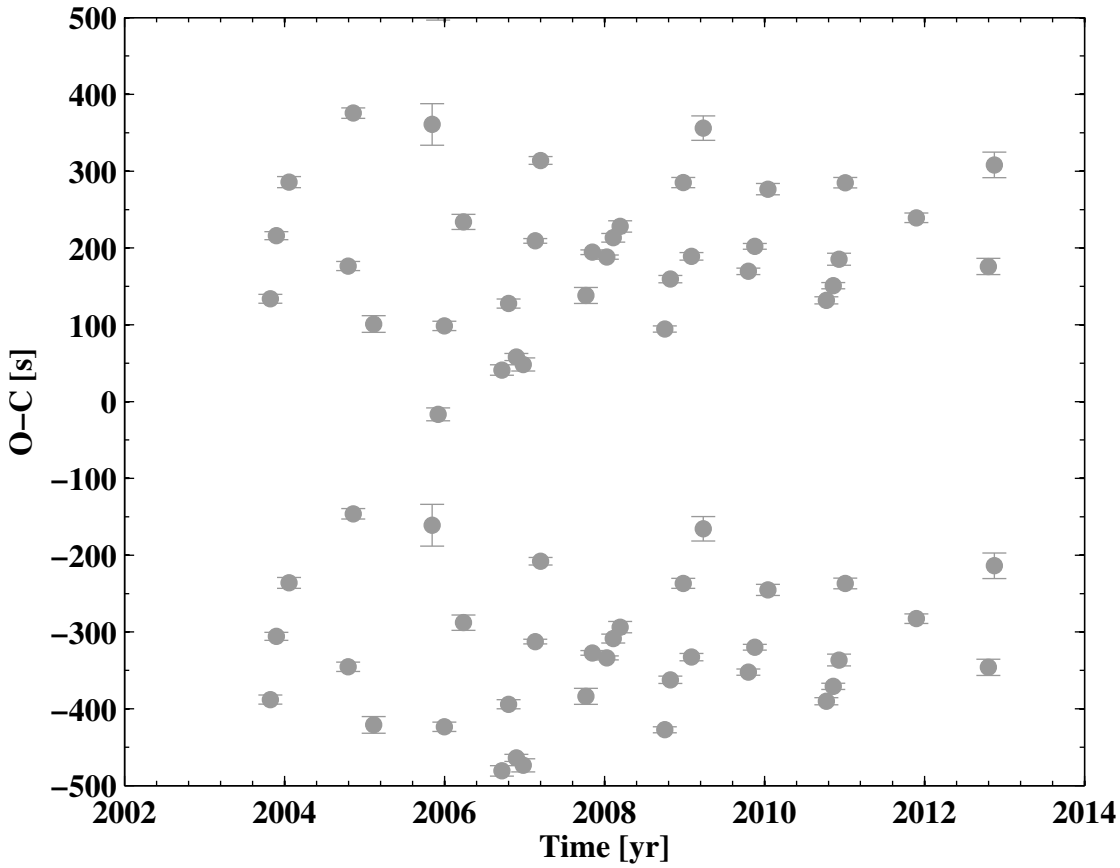

Figure 4.29: $O - C$ of mode $D + E_2 - B$.

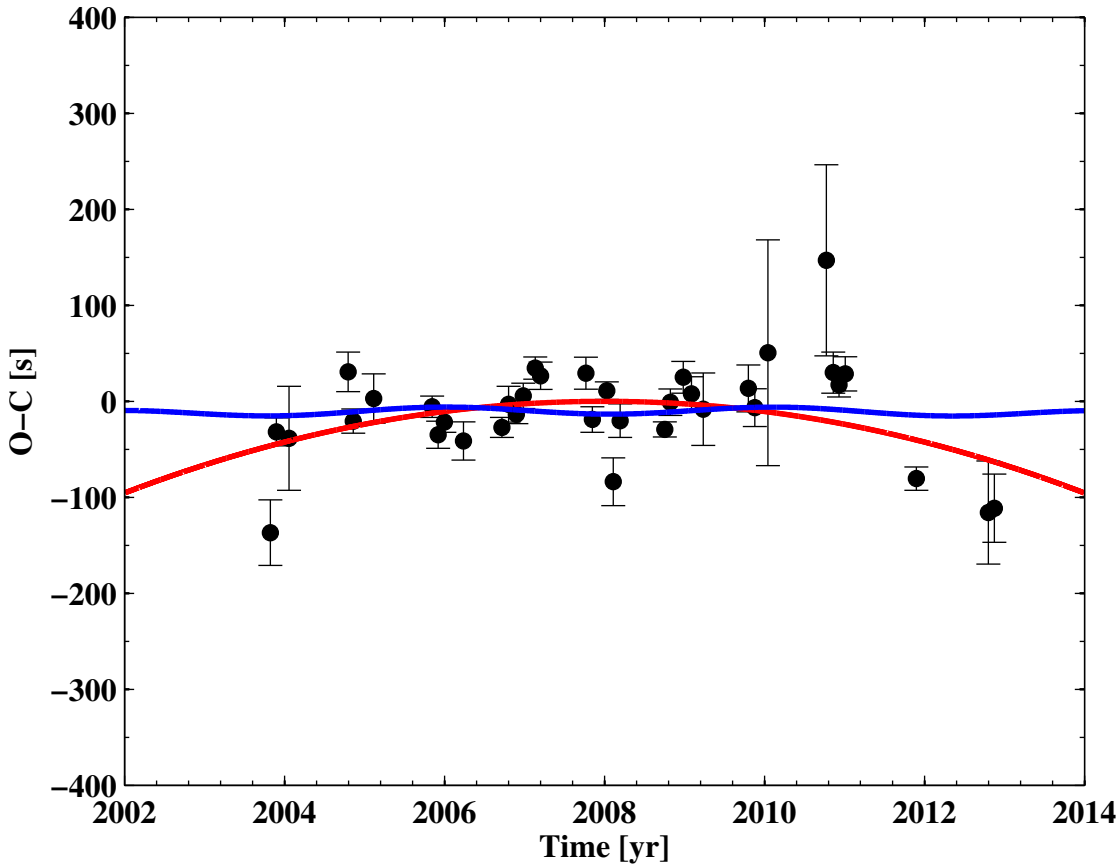

Figure 4.30: $O-C$ of mode $B-E_2$ with a linear trend removed. The red line is the best fit of a parabola. The blue line is the $O-C$ predicted by fits to the $O-C$s of the parent modes B and E_2.

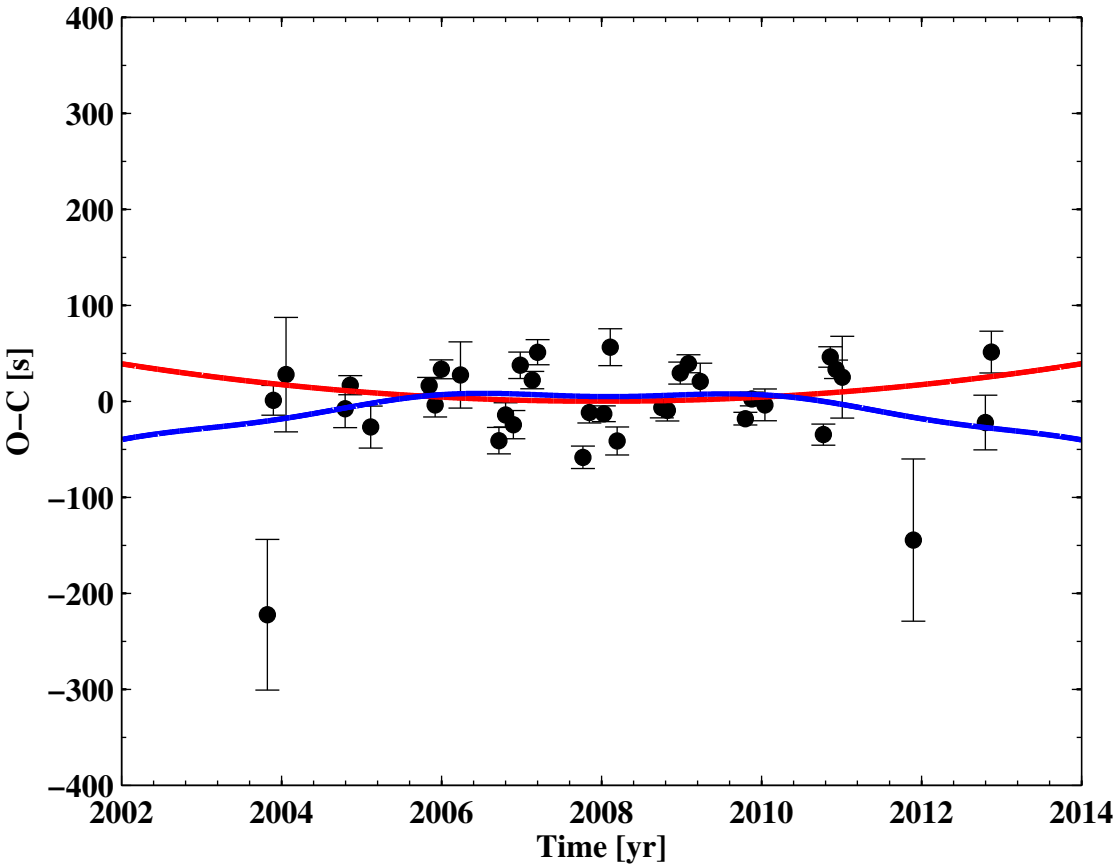

Figure 4.31: $O - C$ of mode $B_+ - E_2$ with a linear trend removed. The red line is the best fit of a parabola. The blue line is the $O - C$ predicted by fits to the $O - C$s of the parent modes B^+ and E_2.

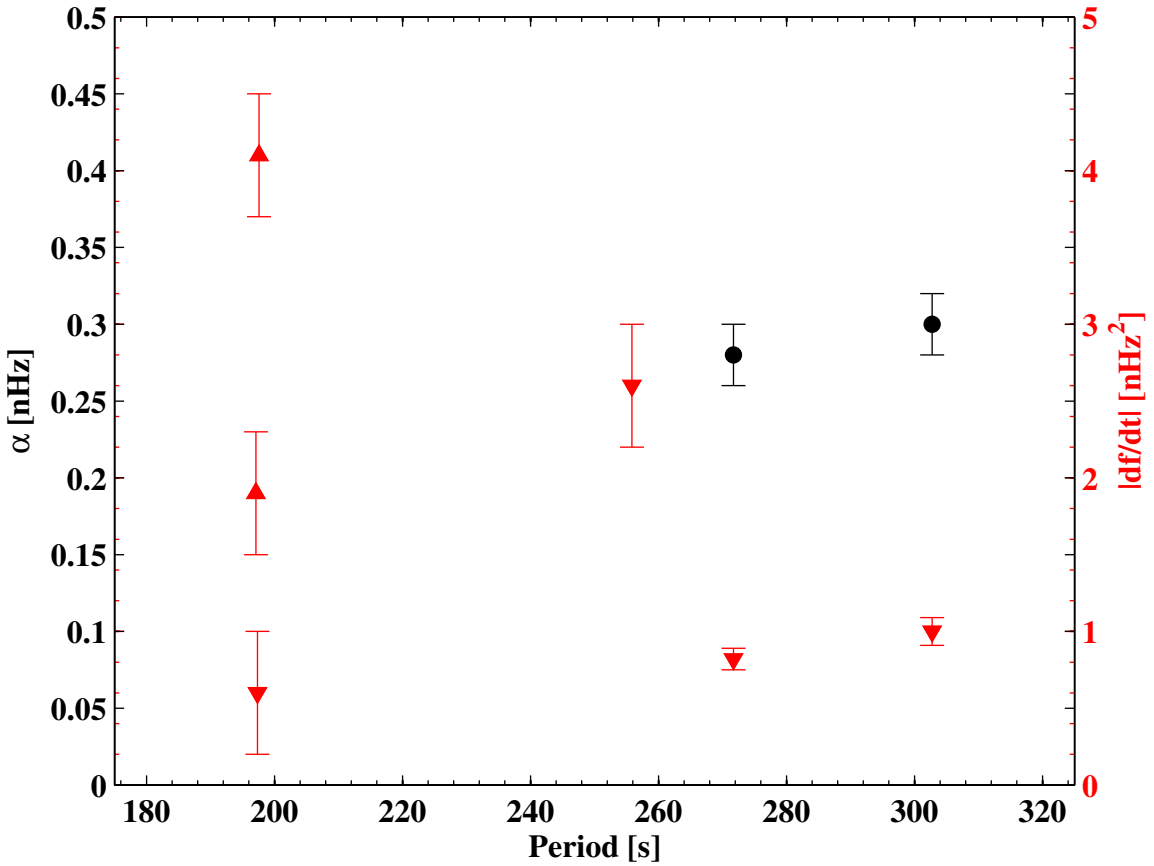

Figure 4.32: Frequency variations of the normal modes. The black points indicate the amplitude of the sinusoidal variations and the red triangles indicate the magnitude of the secular changes. The sign of the secular changes is indicated by the direction of the triangle.

Table 4.4: Results of fitting models to $O-C$. The parameter $\dot{f}$ is the rate of frequency change and the parameters α, and ϕ are the amplitude and phase of a sinusoidal model with a period of 4.15 yr.

Label	**Frequency** $[\mu Hz]$	$\dot{f}$ $[nHz^2]$	α $[nHz]$	ϕ $[deg]$	χ^2_{Red}
$2D + E_2$	10663.55866(0.00006)	$-6(1)$	$< 1.2(0.3)$	-	2.1
$B + D$	8747.26442(0.00005)	3(1)	$< 1.0(0.3)$	-	5.1
$B + E_2$	8369.81724(0.00007)	$-2(2)$	$< 2.9(0.4)$	-	6.2
A_2	8127.84092(0.00003)	0.1(0.8)	1.3(0.1)	31(6)	3.4
$C_3 + D$	7588.61983(0.00006)	$-2(1)$	$< 2.1(0.4)$	-	3.7
$2D$	7360.67041(0.00002)	$-1.4(0.5)$	1.2(0.1)	$-14(5)$	3.0
$D + E_2$	6983.22345(0.00001)	$-2.4(0.4)$	$< 0.7(0.2)$	-	5.3
$2E_2$	6605.77654(0.00005)	$-2(1)$	$< 0.6(0.6)$	-	0.9
B_-	5073.16893(0.00002)	1.9(0.4)	$< 0.7(0.3)$	-	3.2
B	5066.92924(0.00002)	$-0.6(0.4)$	$< 1.1(0.2)$	-	3.2
B_+	5060.68889(0.00002)	4.1(0.4)	$< 0.8(0.2)$	-	3.2
C_2	3908.28461(0.00002)	$-2.6(0.4)$	$< 0.2(0.2)$	-	1.8
D	3680.335184(0.000003)	$-0.82(0.07)$	0.28(0.01)	$-7(2)$	9.1
E_2	3302.888210(0.000004)	$-1.00(0.09)$	0.33(0.02)	157(3)	6.9
$B - E_2$	1764.04089(0.00007)	9(2)	$< 6(1)$	-	3.7
$B_+ - E_2$	1757.80044(0.00006)	$-4(2)$	$< 4.1(0.8)$	-	5.8

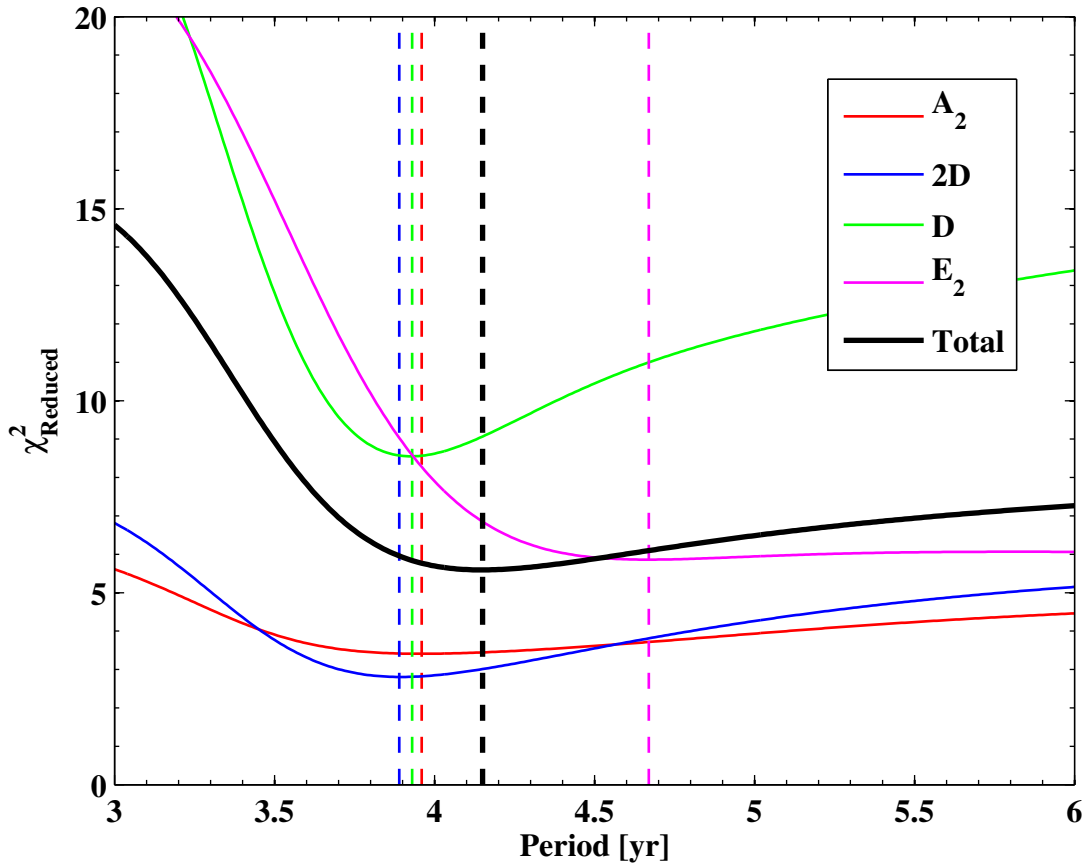

Figure 4.33: Reduced χ^2 as a function of period for fits of a sinusoid plus parabola to the $O-C$ of selected modes. The dashed lines indicate a minimum in reduced χ^2.

is no favored period and its difficult to say if a sinusoidal model is appropriate. Mode D^+ does seem to have a clearly defined minimum in χ^2 but at almost twice the period of modes A_2, $2D$, D, and E_2. These $O-C$s were fit with a "high" order polynomial. The fits are shown in Figures 4.20, 4.22, 4.23, 4.25, and 4.26.

A simulation was performed to test whether or not $O-C$ could be reliably recovered from artificial data with similar properties. First, each lightcurve was pre-whitened at the frequencies listed in Table 4.3. The values of the lightcurve were shuffled to remove any additional signals from the DFT. A sinusoid for each amplitude

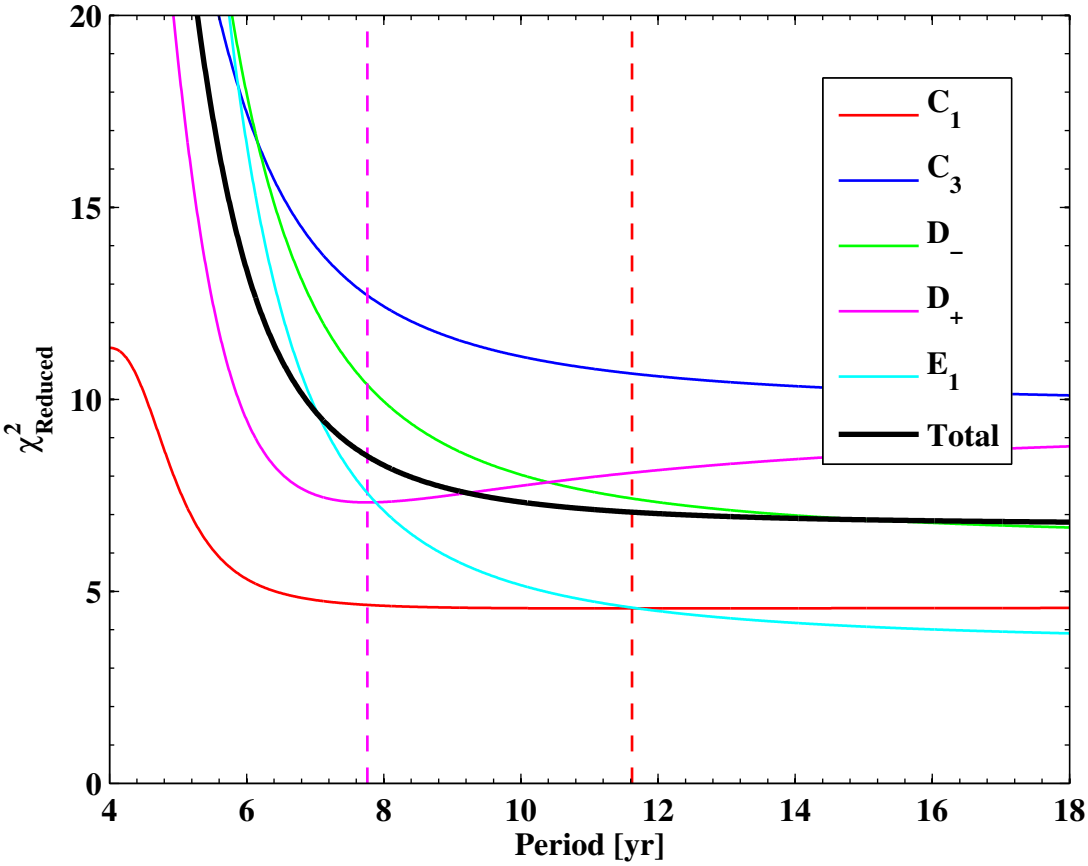

Figure 4.34: Reduced χ^2 as a function of period for fits of a sinusoid plus parabola to the $O - C$ of selected modes.

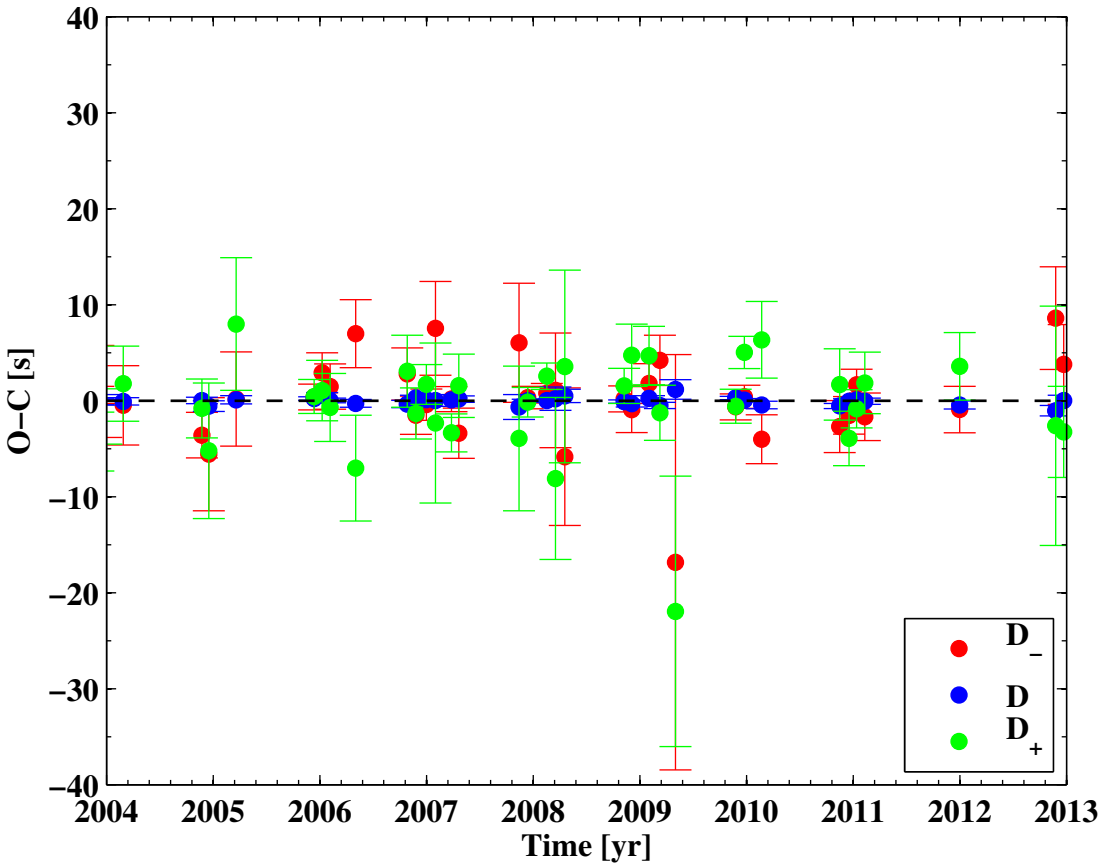

Figure 4.35: $O-C$ diagrams of simulated data for the D triplet. The frequencies were held constant in the simulations.

and frequency listed in Table 4.3 was then added to the signal. The absolute phase of the sinusoids were held constant. Finally, $O-C$ was calculated at each frequency for each lightcurve. For every mode, the $O-C$s were consistent with constant. The artificial $O-C$s of the D triplet are shown in Figure 4.35.

The fits to the parent modes' $O-C$s were used to predict the $O-C$s of the combination and harmonic modes. The predictions are shown in blue in Figures 4.8, 4.9, 4.10, 4.14, 4.15, 4.16, 4.30, and 4.31. The fit parameters shown in Table 4.4 can also be compared. The agreement is generally excellent.

The residuals of the parabolic fits to the $O-C$s of modes B^-, B, B^+, C_1,

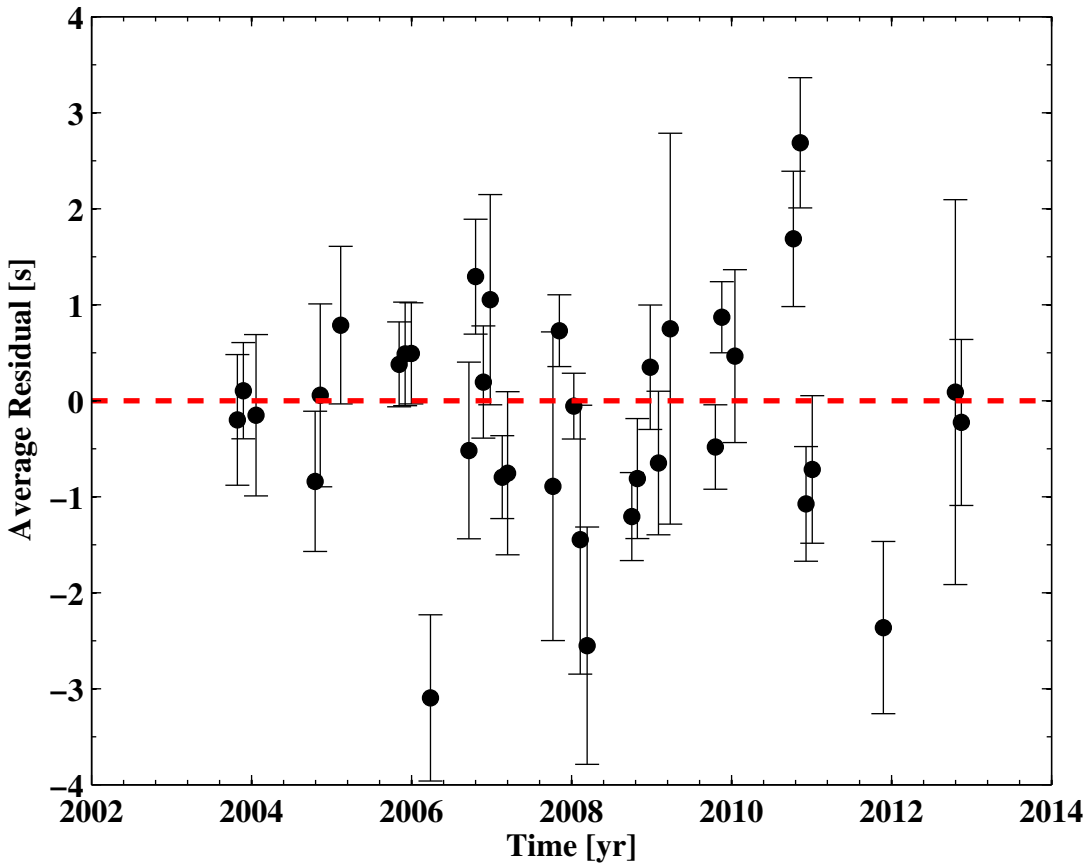

Figure 4.36: Average residuals of the fits to the $O-C$s of modes B^-, B, B^+, C_1, and C_2.

and C_2 were averaged and are shown in Figure 4.36. The Fourier transform of the averaged residuals is shown in Figure 4.37. There are no statistically significant periodic variations in the residuals down to about a half a second. This puts an upper limit on the mass of a planetary companion in a several AU orbit at less than half the mass of Jupiter. This, combined with the out of phase difference between the periodic variations between mode E_2 and D, puts an end to the possibility that any of these $O-C$ variations are being caused by a planetary companion.

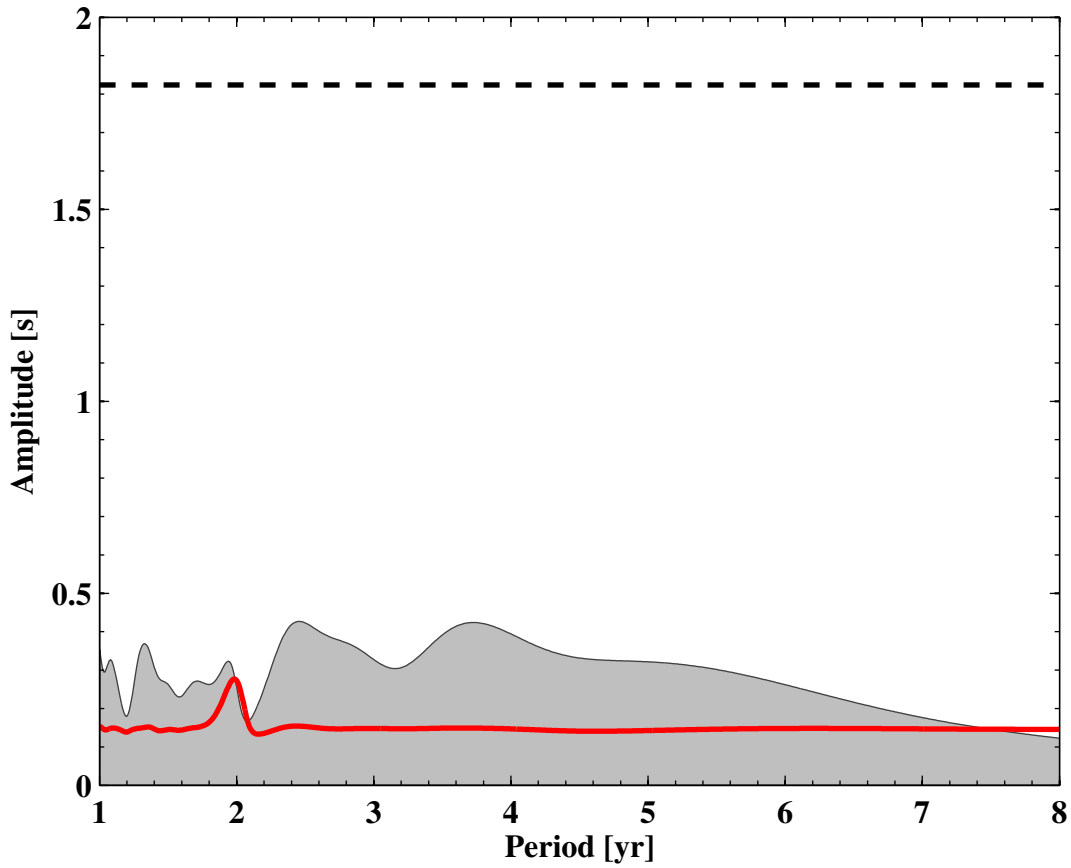

Figure 4.37: Fourier spectrum of the average residuals of the fits to the $O-C$s of modes B^-, B, B^+, C_1, and C_2. The red line is the formal error and the dashed black line is 6 times the median amplitude.

4.4 Amplitude Modulation

The amplitudes are shown as a function of time in Figures 4.38-4.61. The amplitudes of many of the modes are clearly periodic and were fit with a first order polynomial plus a sinusoid. For the remaining modes, the amplitudes were fit with a constant amplitude. Upper limits on the amplitude of a periodic variation were calculated using the residuals. The results are summarized in Table 4.5.

The RC values of the combination and harmonic modes were also calculated. The RC values of modes $2D$ and $2E$ were subtracted by their mean are shown in Figure 4.62. The RC values don't appear consistent with constant. A fit of a constant to either set of RC values has a reduced χ^2 of over 4. The correlated behavior of the two trends could be caused by variations in the reduction process.

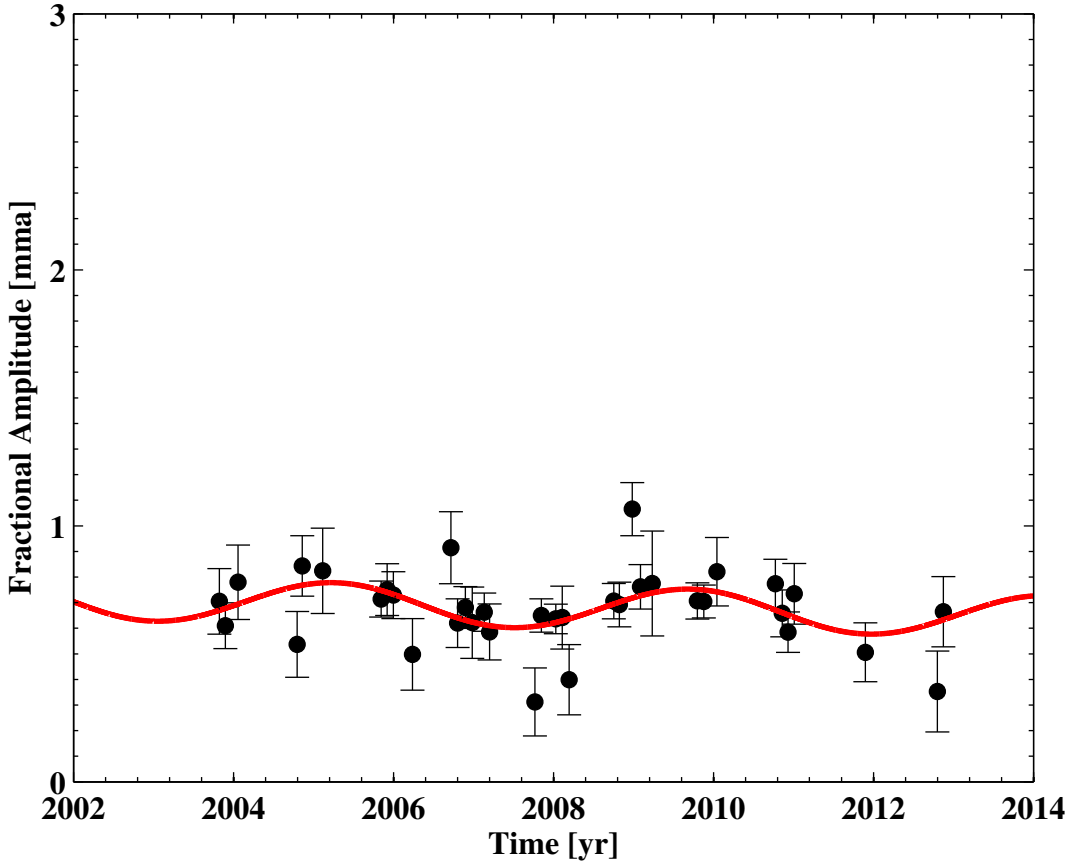

Figure 4.38: Amplitude of mode $2D + E_2$. The red dashed line is the weighted mean.

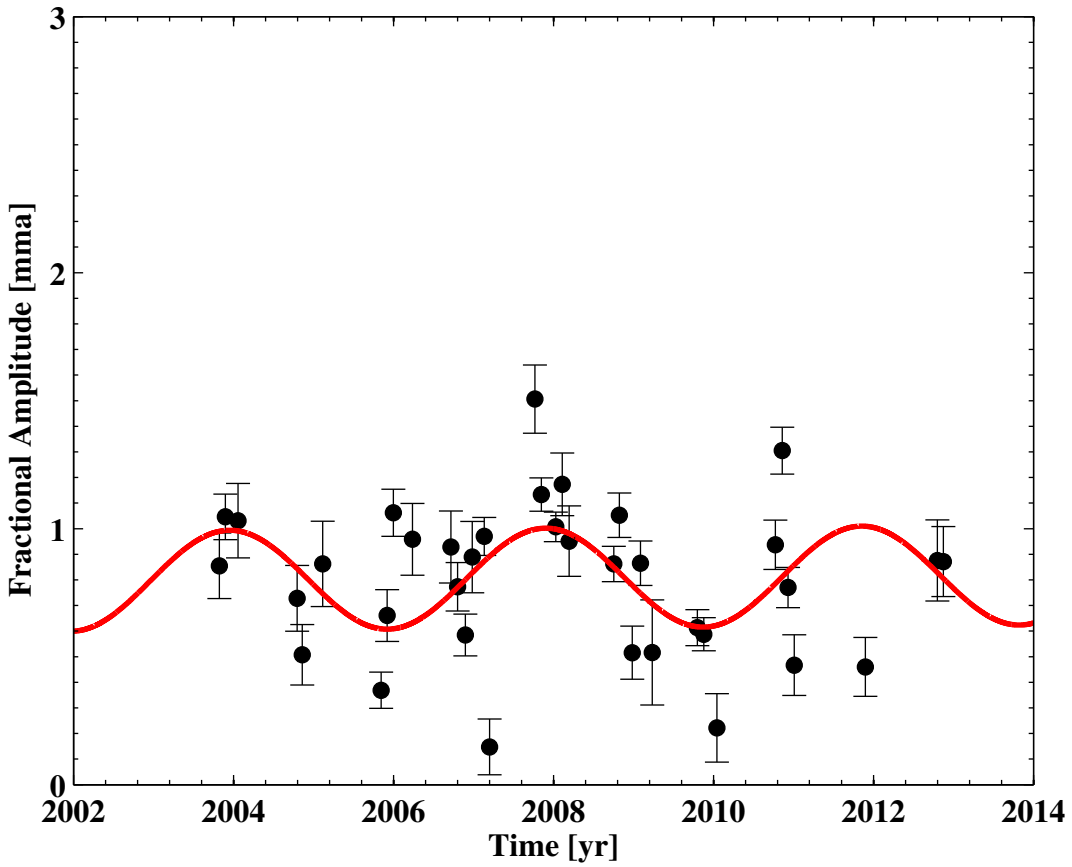

Figure 4.39: Amplitude of mode $B + D$. The red curve is the best fit of a line plus a sinusoid.

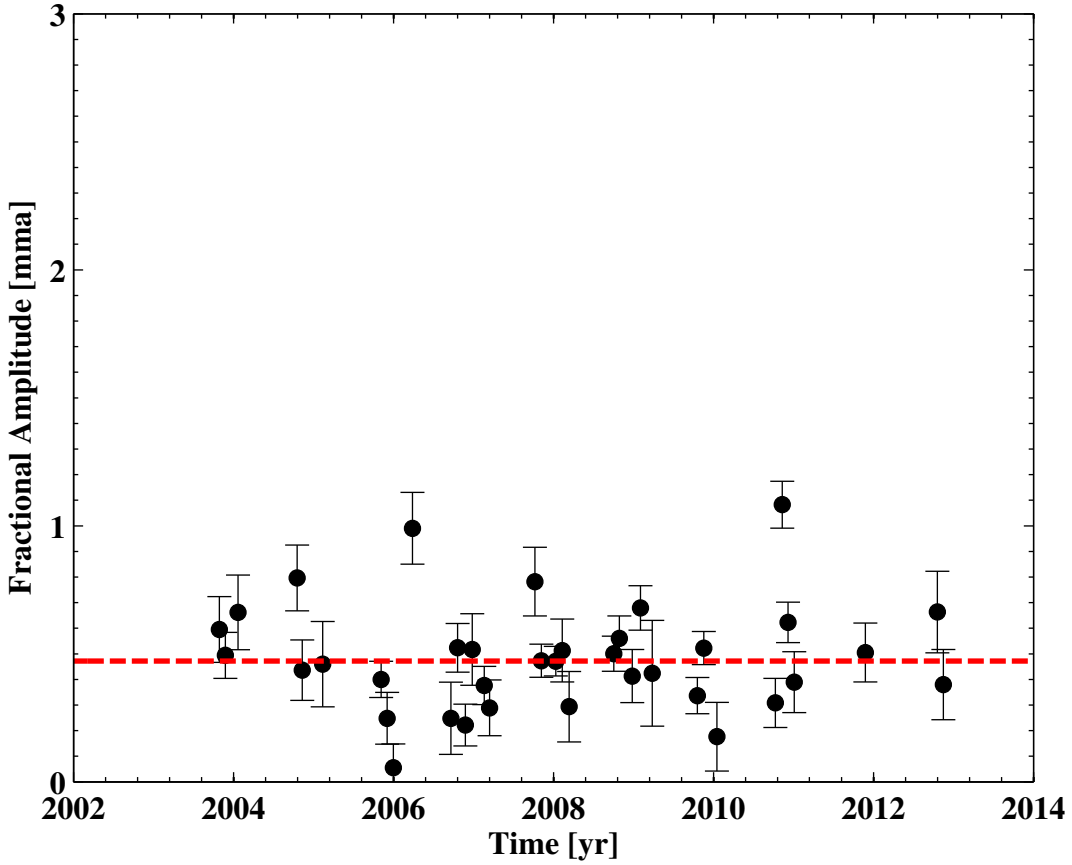

Figure 4.40: Amplitude of mode $B + E_2$. The red dashed line is the weighted mean.

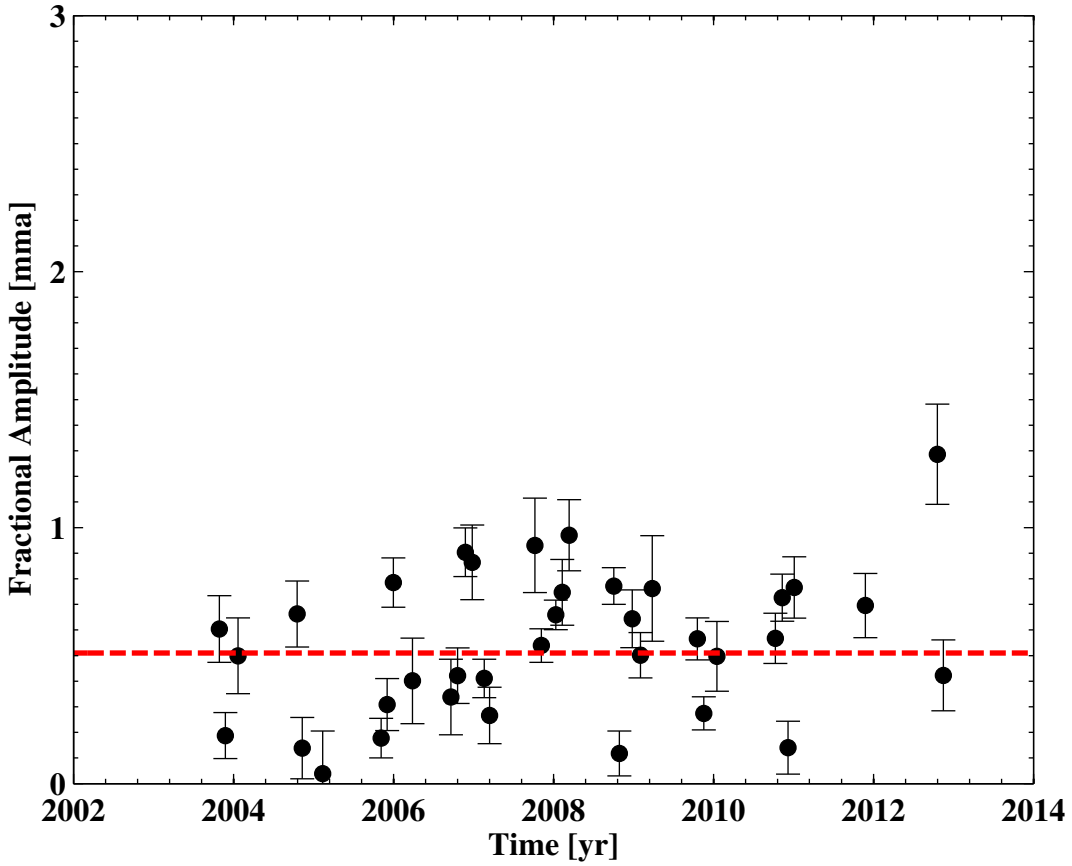

Figure 4.41: Amplitude of mode A_1. The red dashed line is the weighted mean.

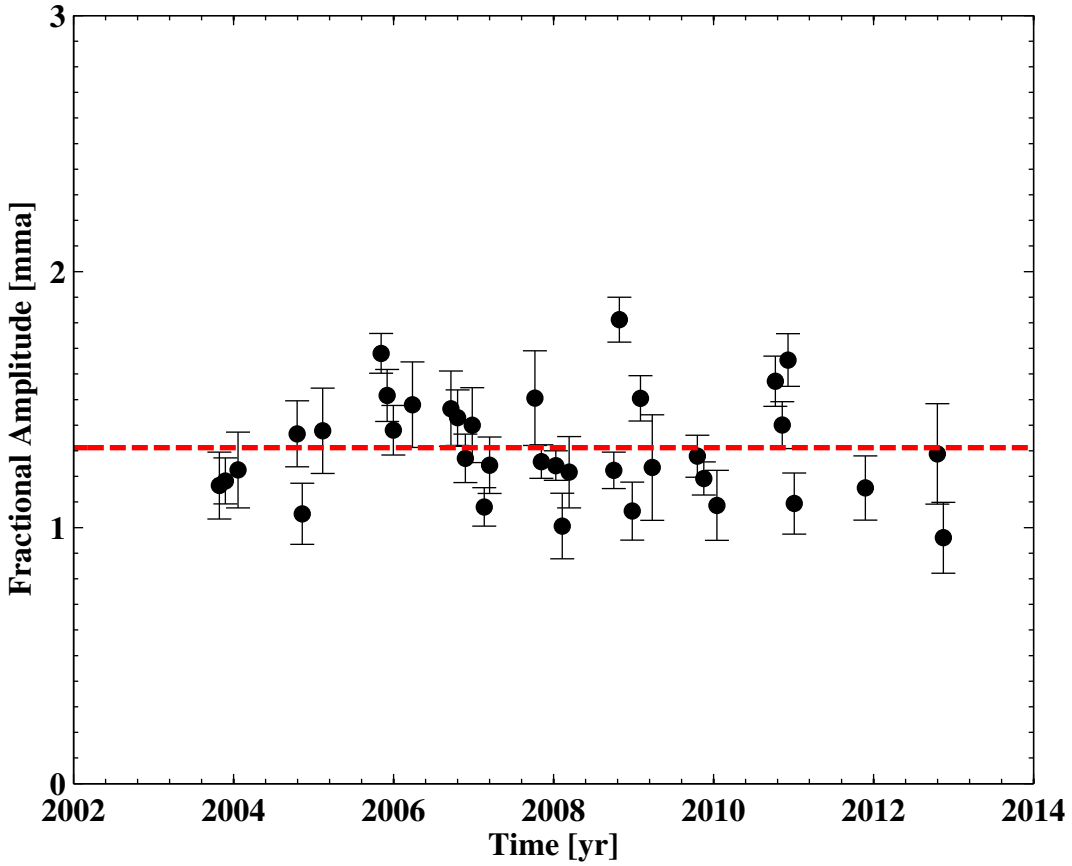

Figure 4.42: Amplitude of mode A_2. The red dashed line is the weighted mean.

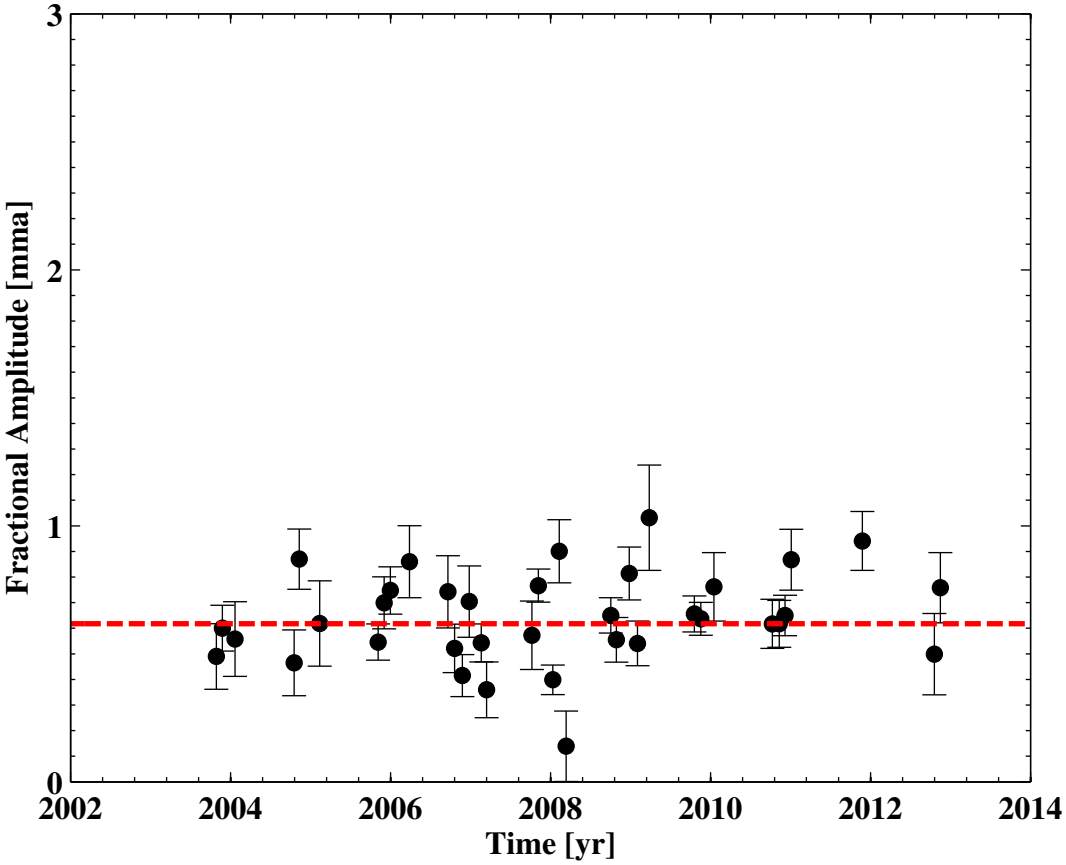

Figure 4.43: Amplitude of mode $C_3 + D$. The red dashed line is the weighted mean.

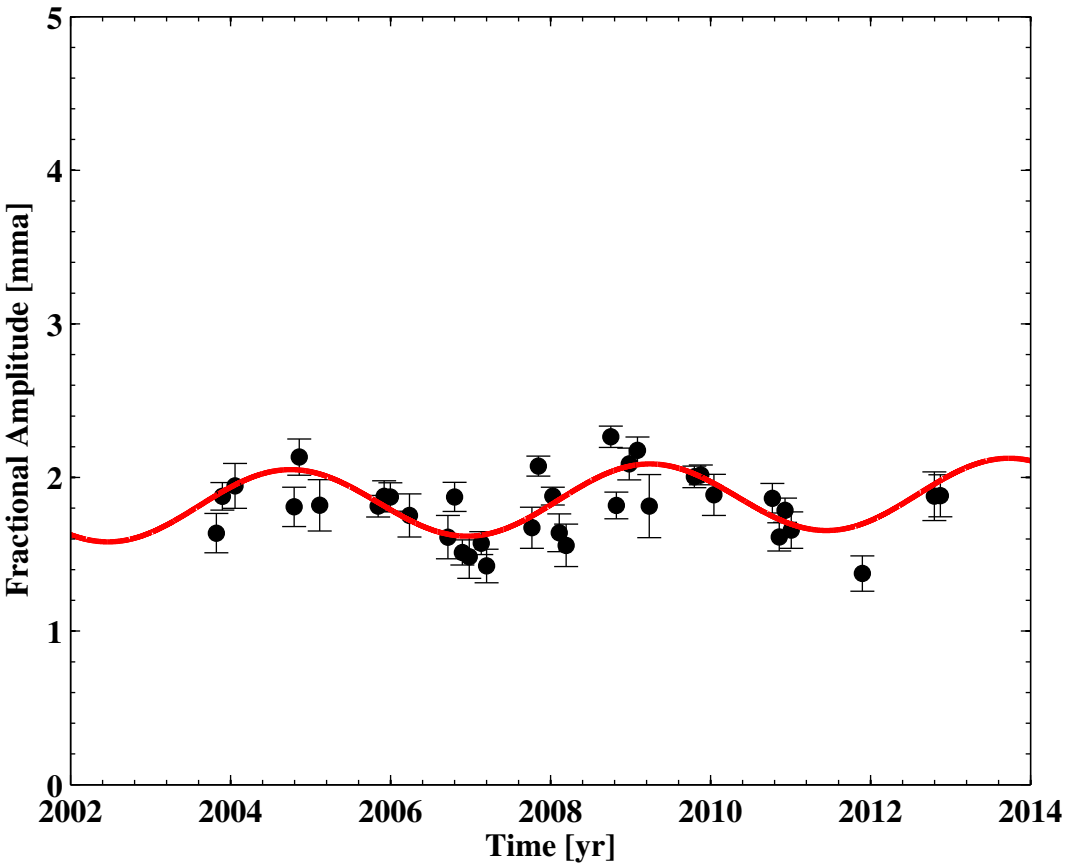

Figure 4.44: Amplitude of mode $2D$. The red curve is the best fit of a line plus a sinusoid.

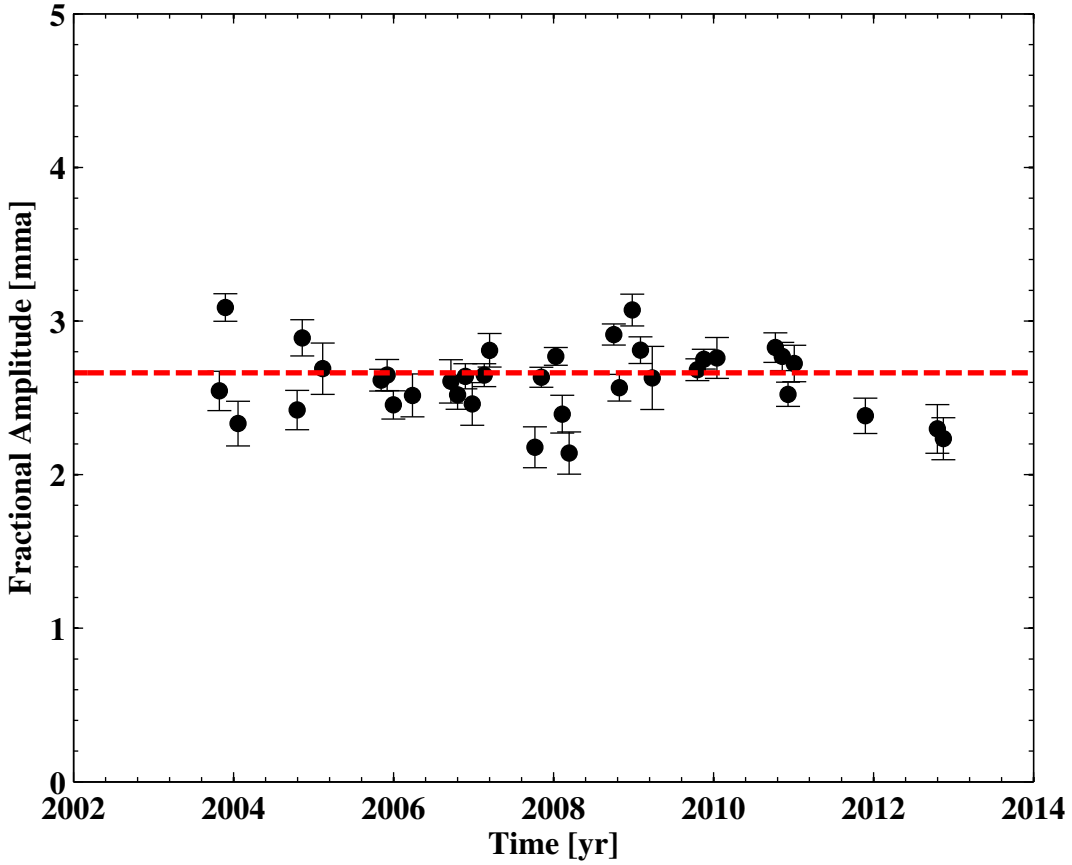

Figure 4.45: Amplitude of mode $D + E_2$. The red dashed line is the weighted mean.

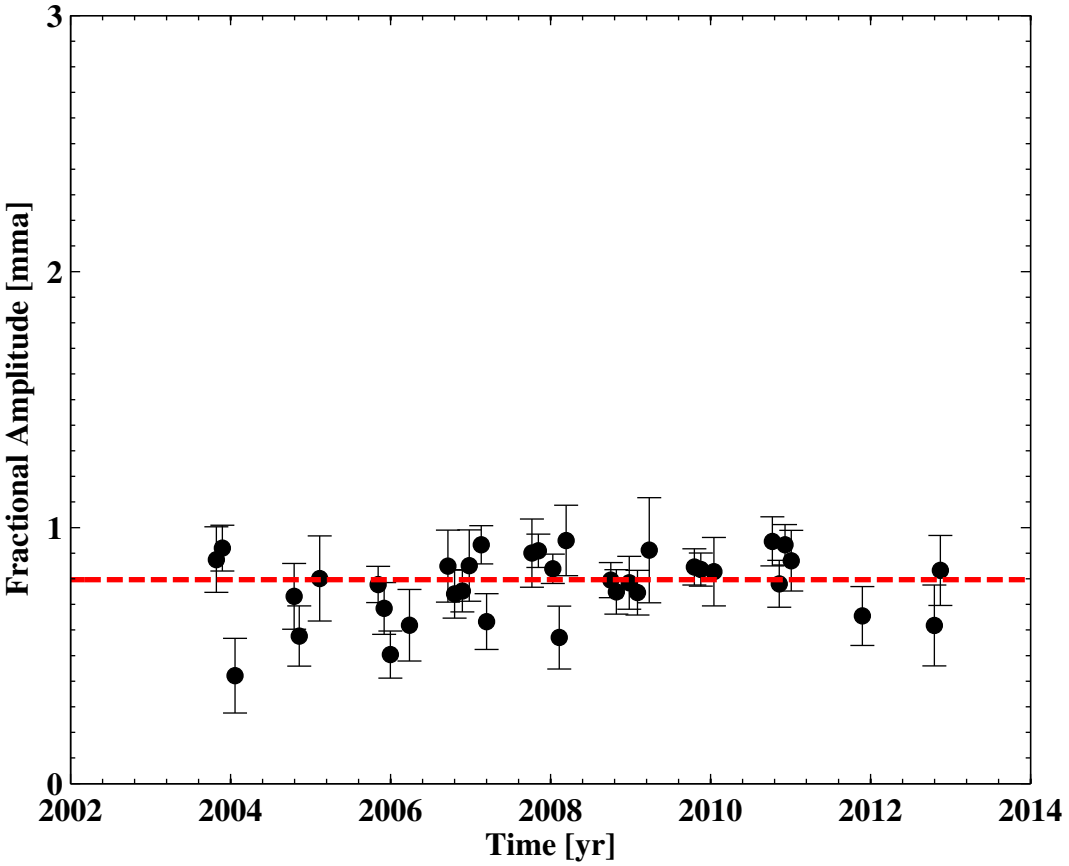

Figure 4.46: Amplitude of mode $2E_2$. The red dashed line is the weighted mean.

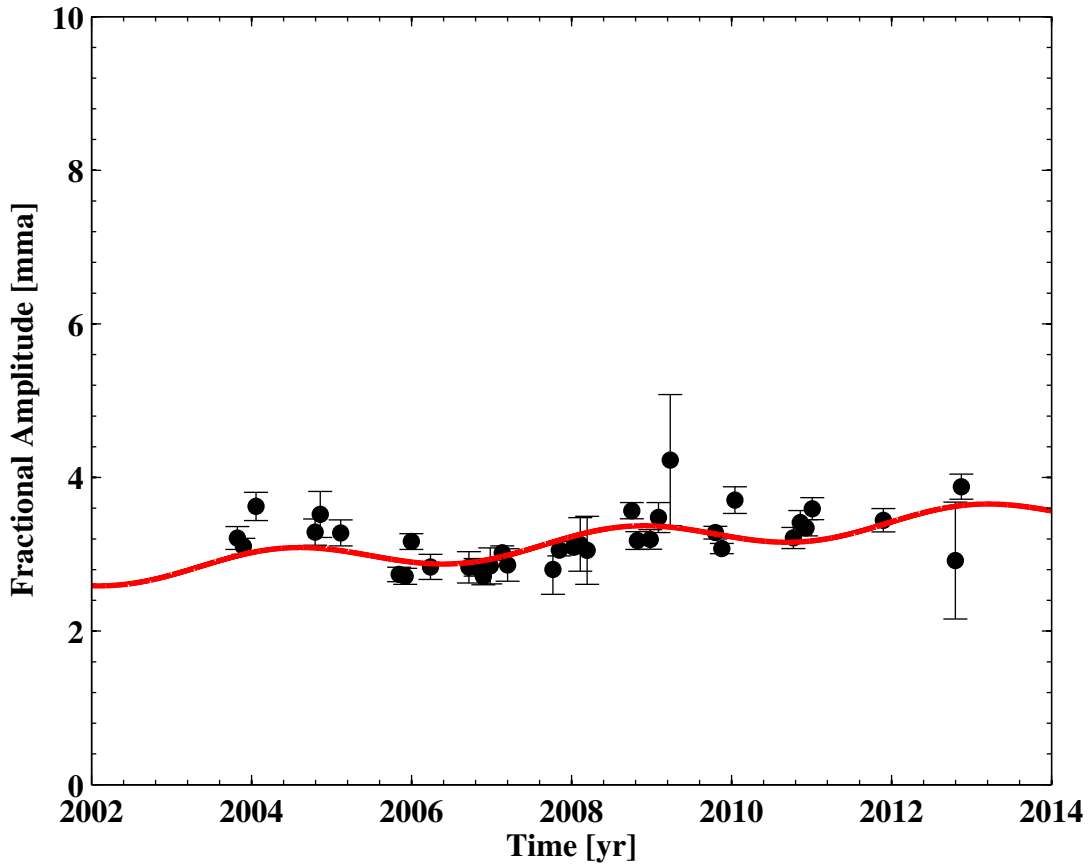

Figure 4.47: Amplitude of mode B_{-}. The red curve is the best fit of a line plus a sinusoid.

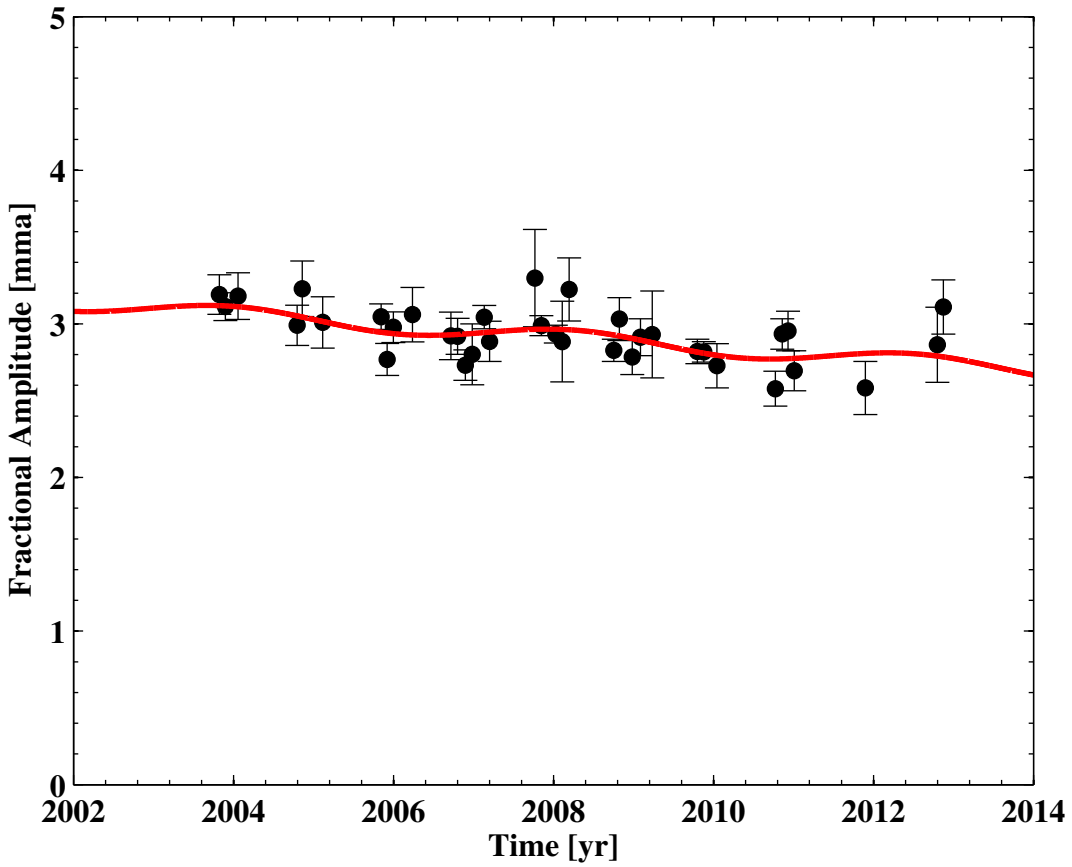

Figure 4.48: Amplitude of mode B. The red curve is the best fit of a line plus a sinusoid.

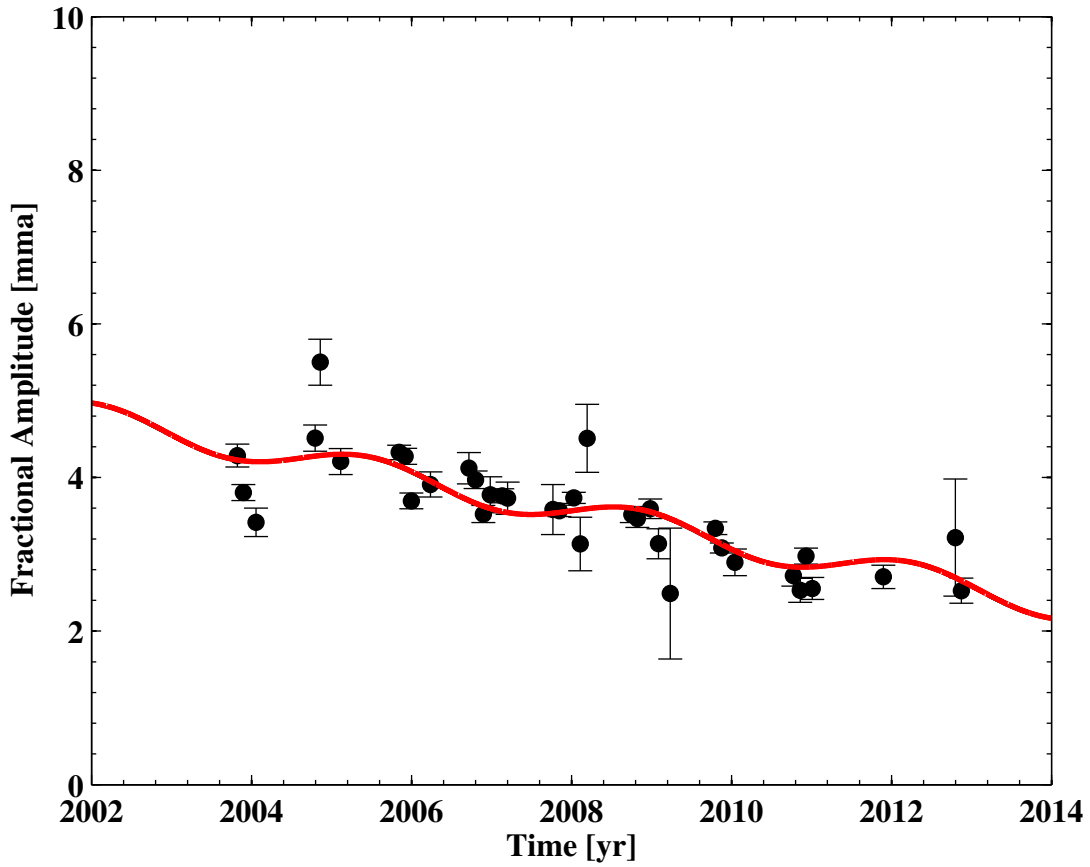

Figure 4.49: Amplitude of mode B_+. The red curve is the best fit of a line plus a sinusoid.

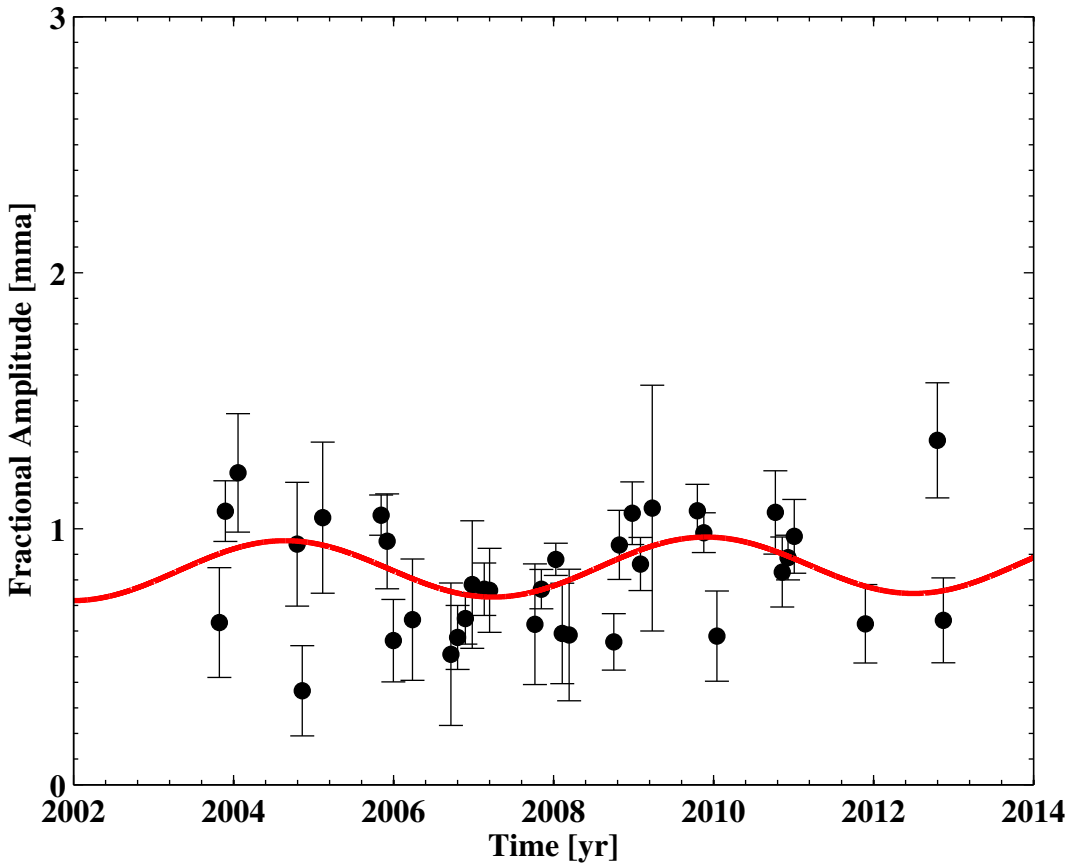

Figure 4.50: Amplitude of mode C_1. The red curve is the best fit of a line plus a sinusoid.

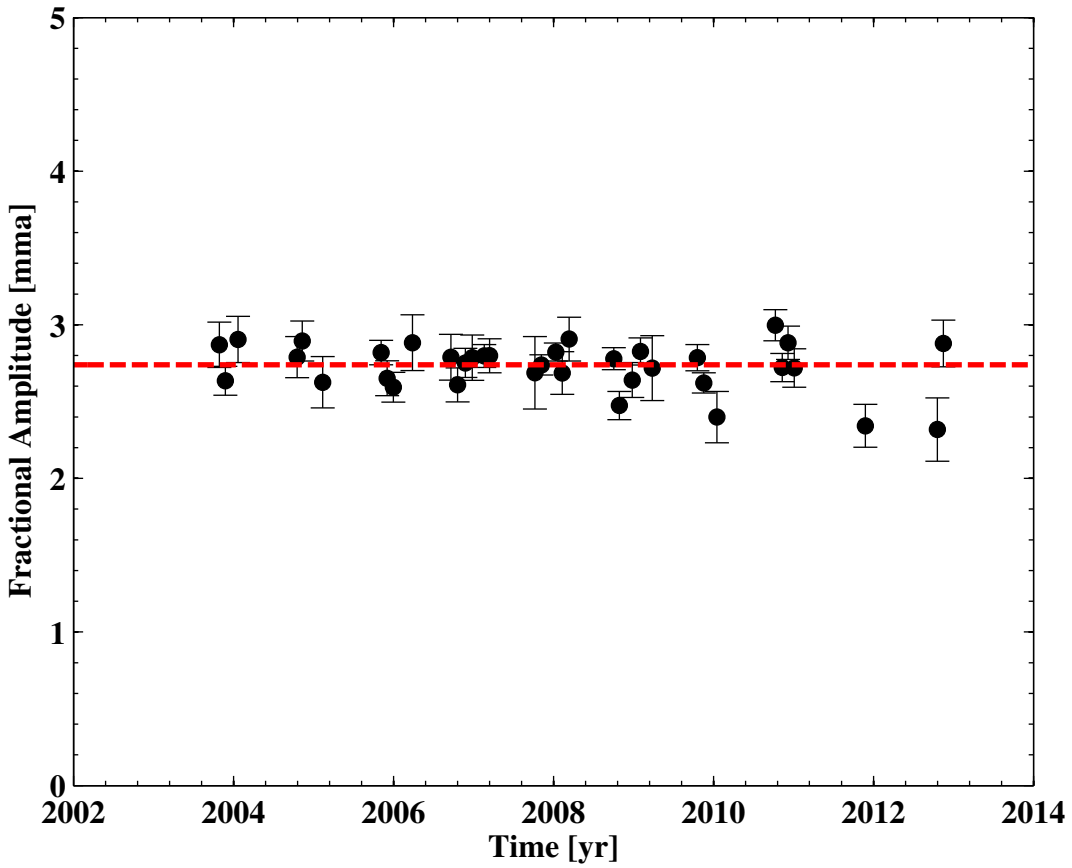

Figure 4.51: Amplitude of mode C_2. The red dashed line is the weighted mean.

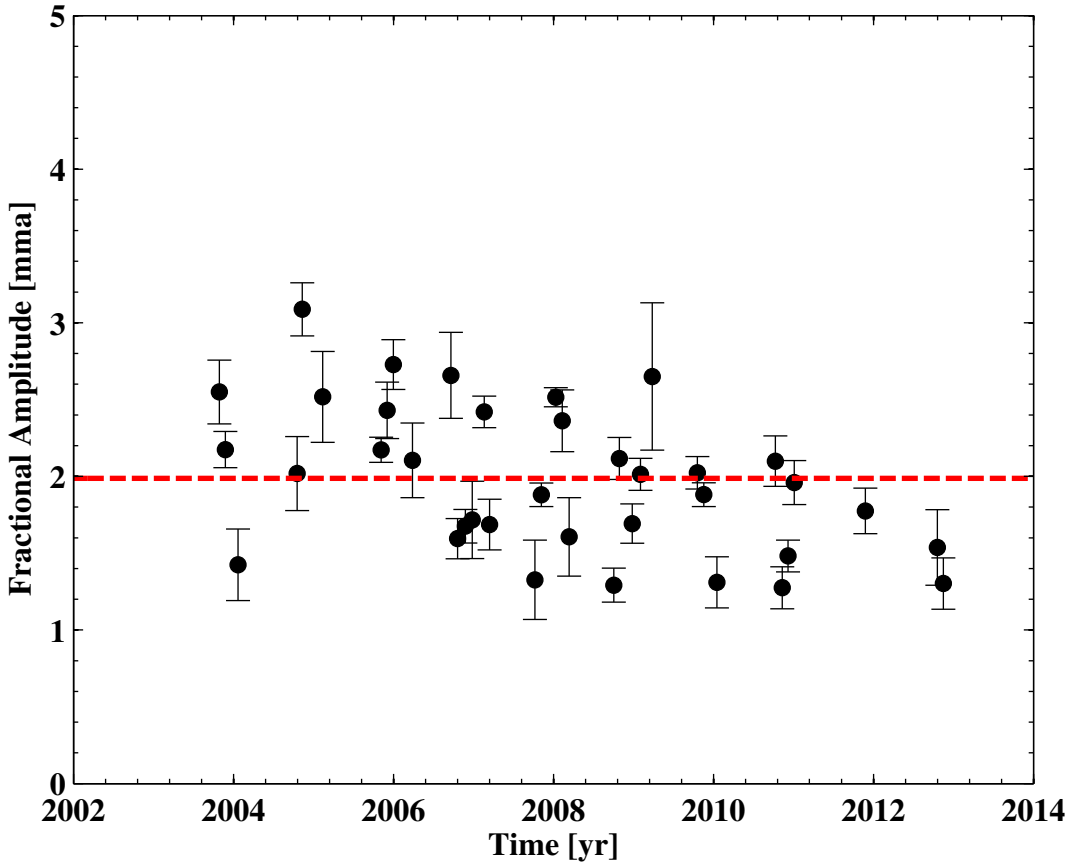

Figure 4.52: Amplitude of mode C_3. The red dashed line is the weighted mean.

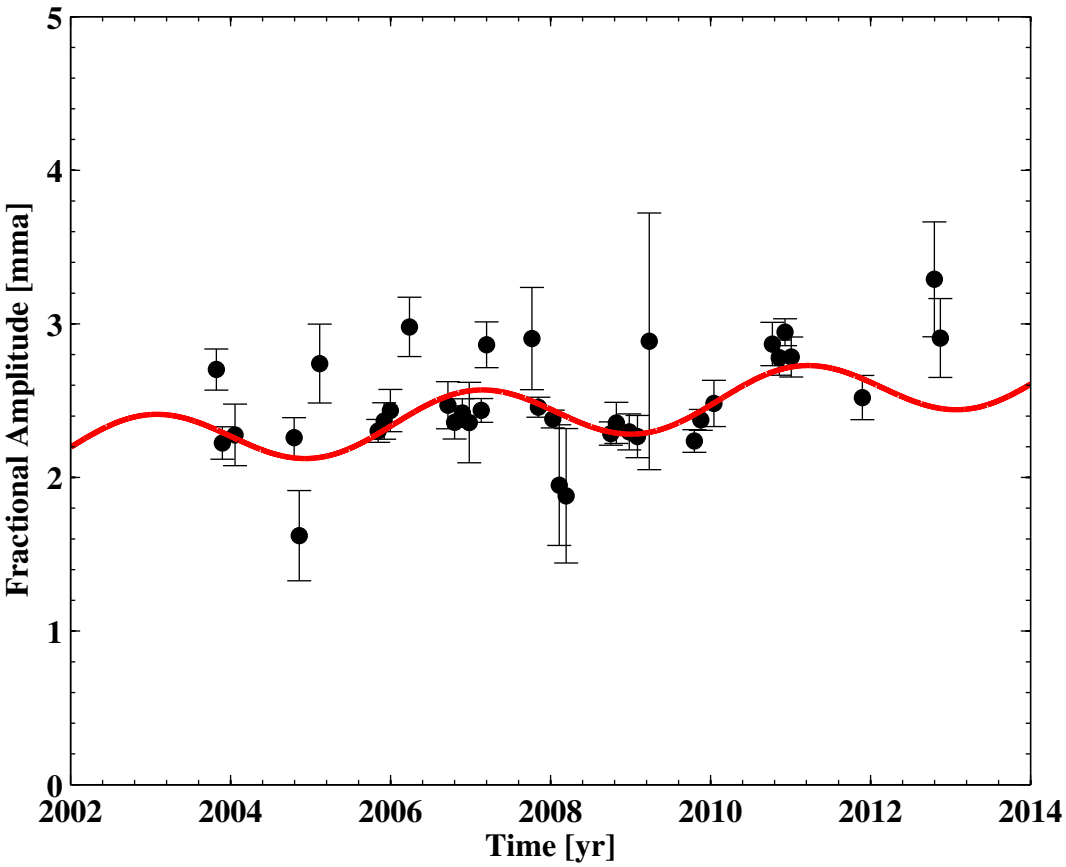

Figure 4.53: Amplitude of mode D_-. The red curve is the best fit of a line plus a sinusoid.

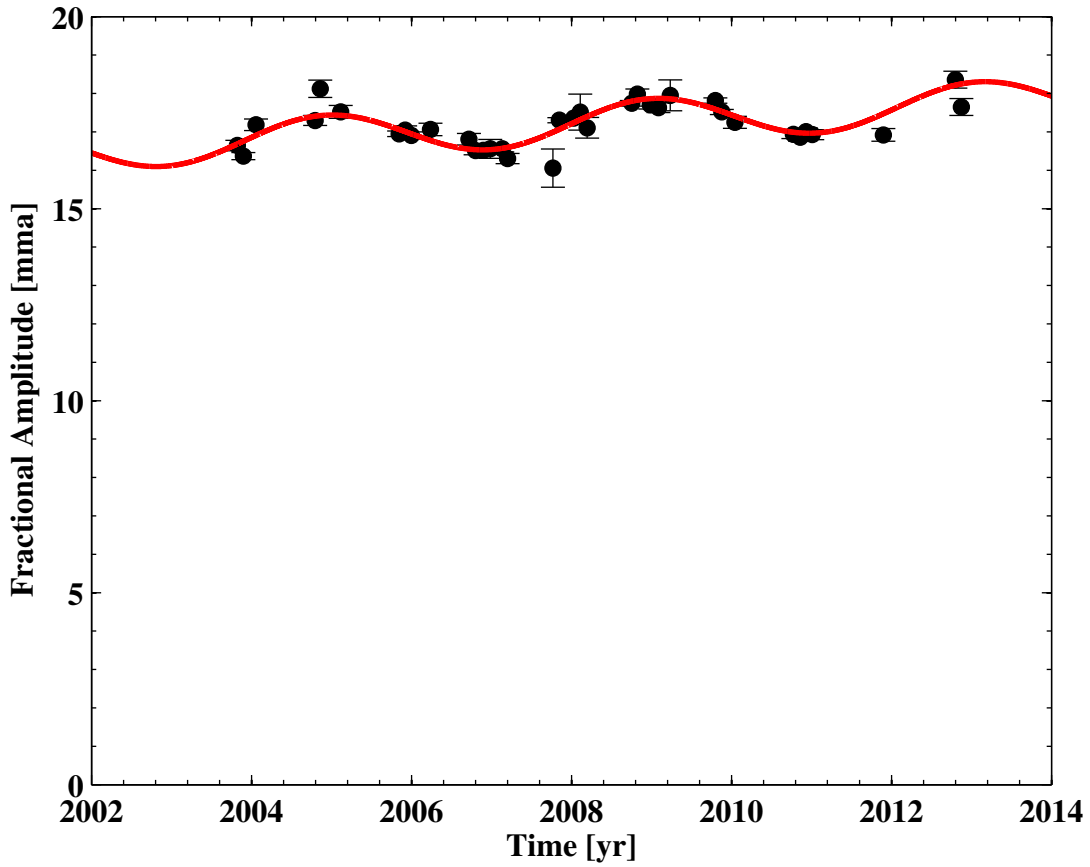

Figure 4.54: Amplitude of mode D. The red curve is the best fit of a line plus a sinusoid.

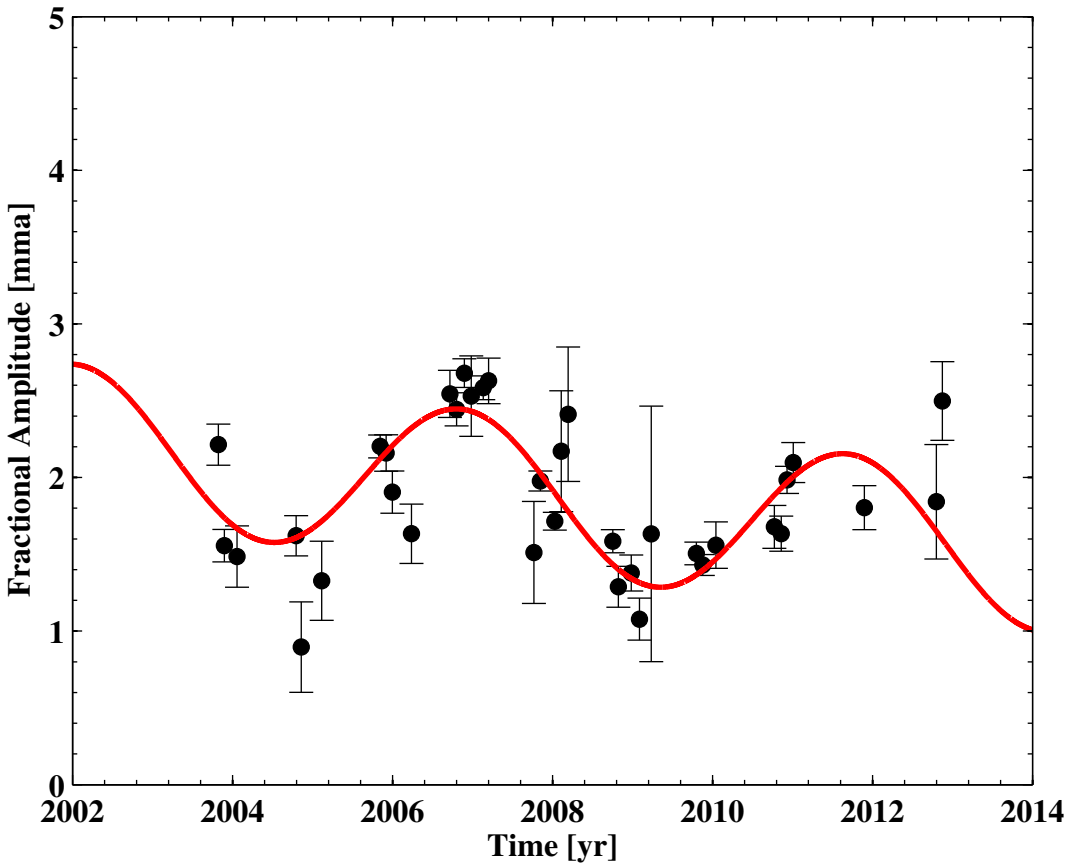

Figure 4.55: Amplitude of mode D_+. The red curve is the best fit of a line plus a sinusoid.

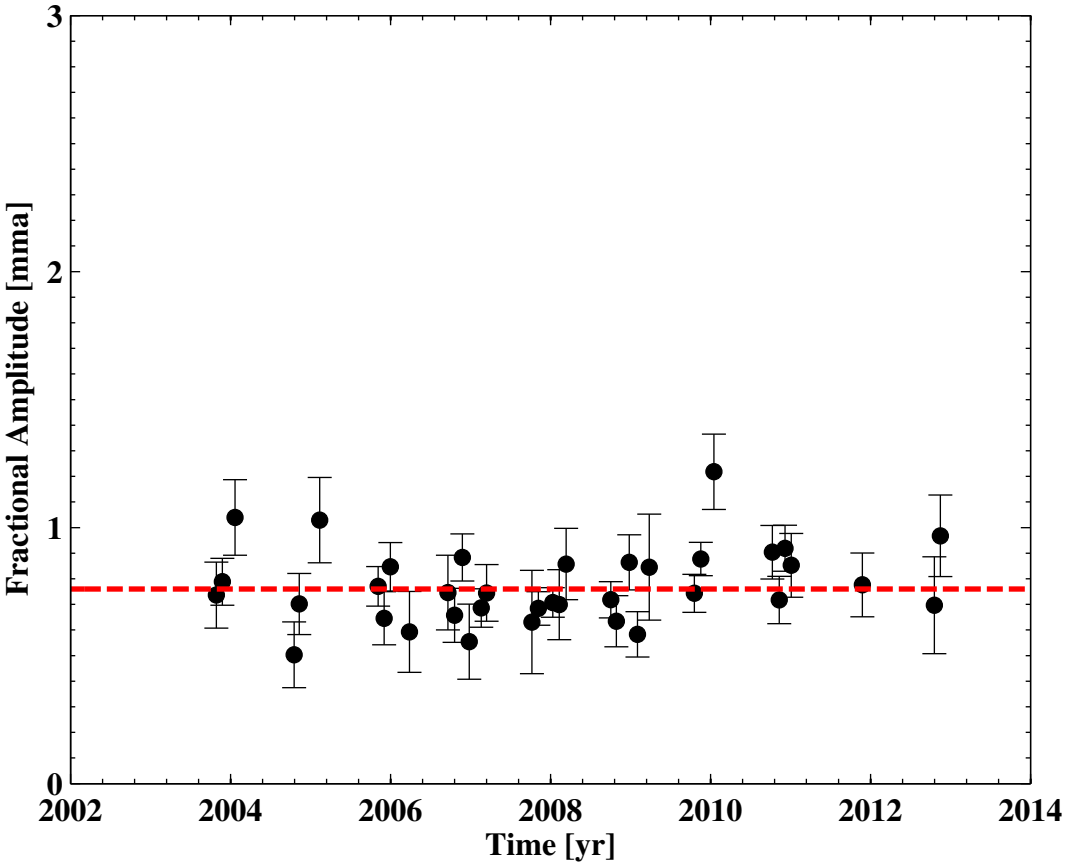

Figure 4.56: Amplitude of mode E_1. The red dashed line is the weighted mean.

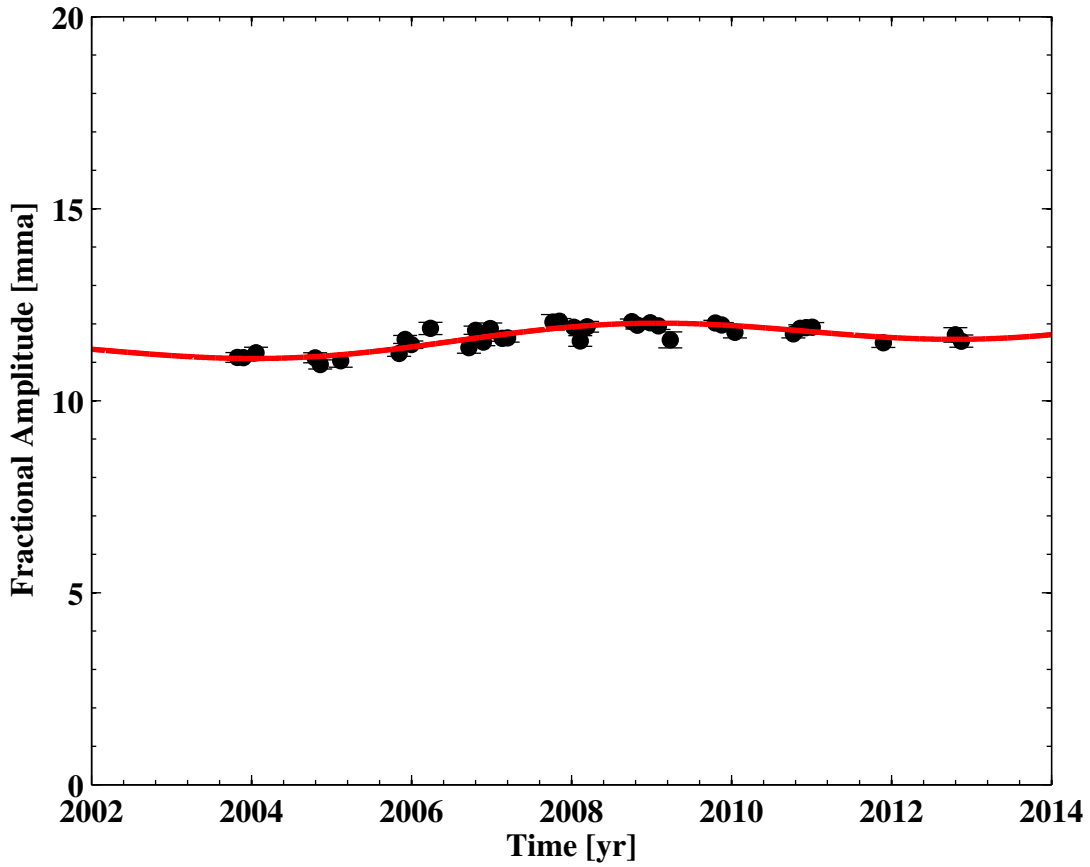

Figure 4.57: Amplitude of mode E_2. The red curve is the best fit of a line plus a sinusoid.

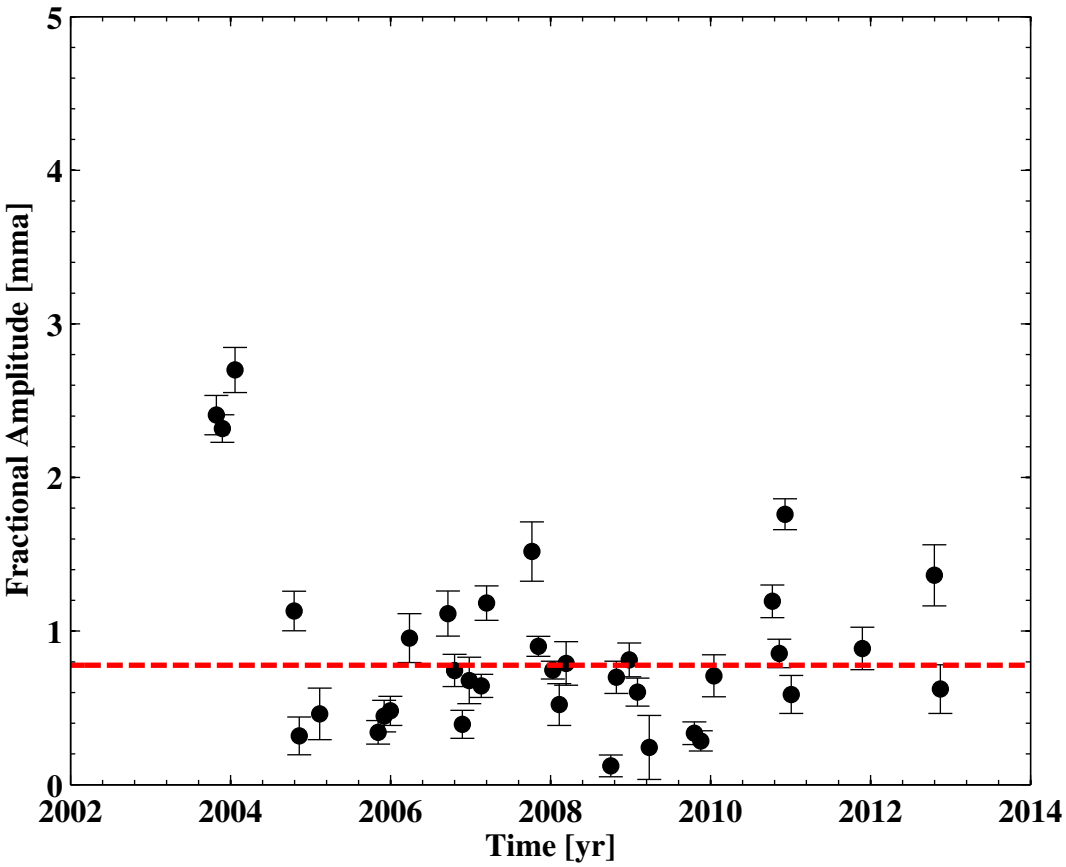

Figure 4.58: Amplitude of mode $D + E_2 - B_+$. The red dashed line is the weighted mean.

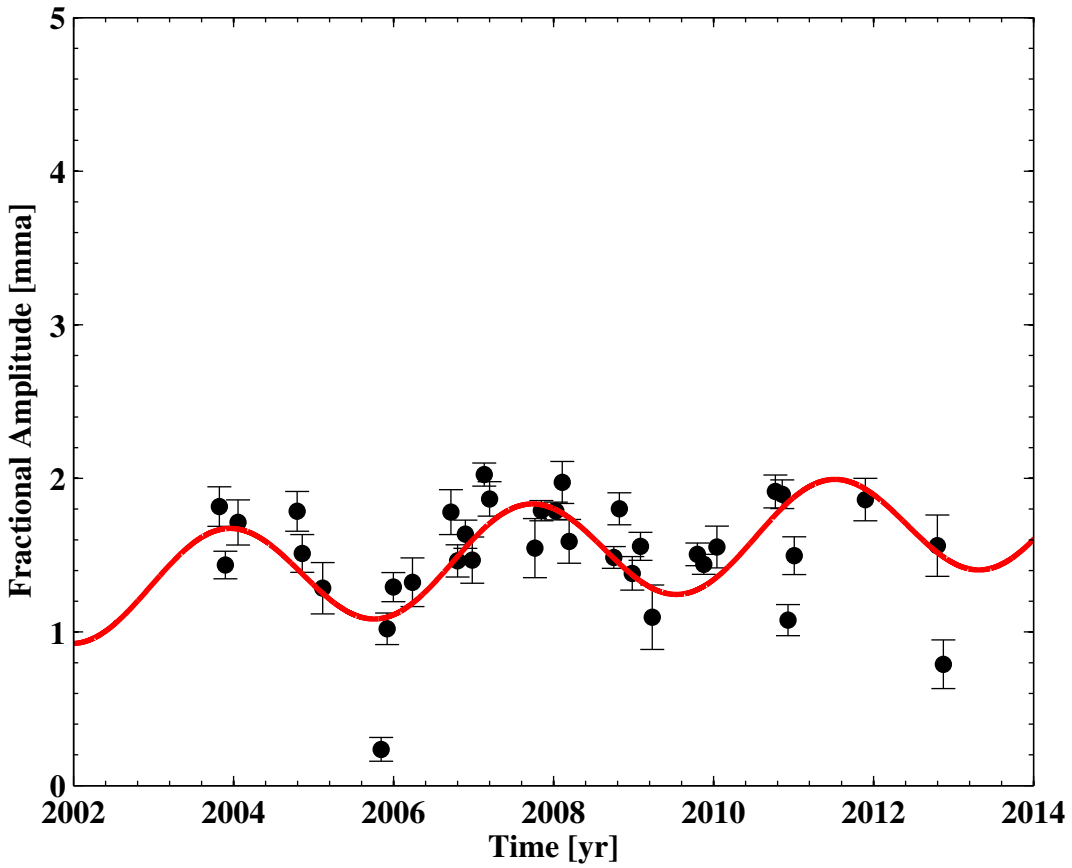

Figure 4.59: Amplitude of mode $D + E_2 - B$. The red curve is the best fit of a line plus a sinusoid.

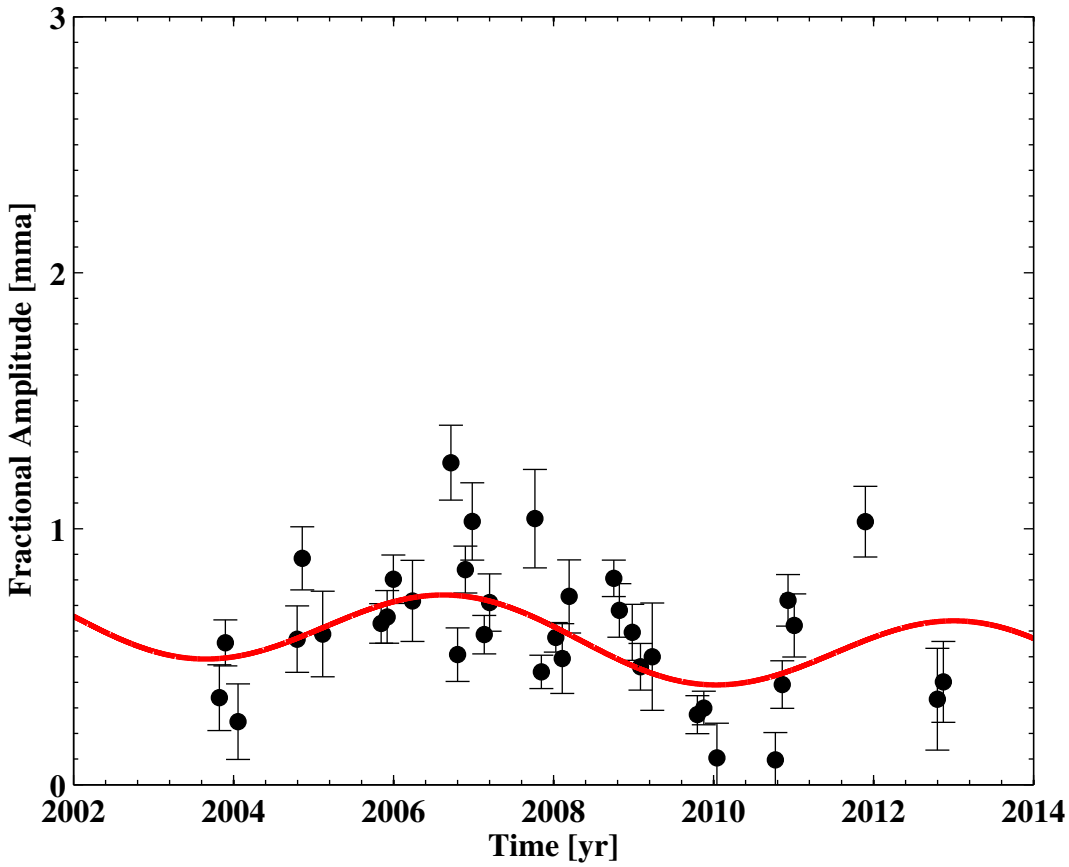

Figure 4.60: Amplitude of mode $B - E_2$. The red curve is the best fit of a line plus a sinusoid.

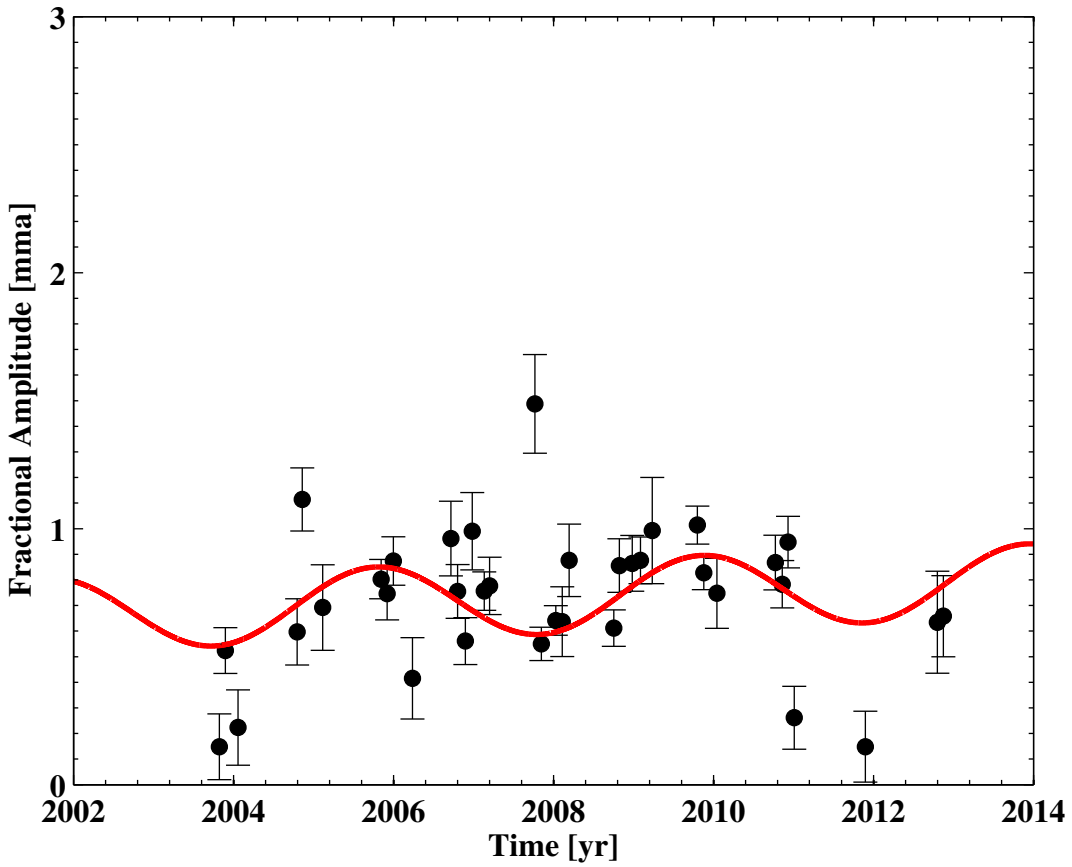

Figure 4.61: Amplitude of mode $B_+ - E_2$. The red curve is the best fit of a line plus a sinusoid.

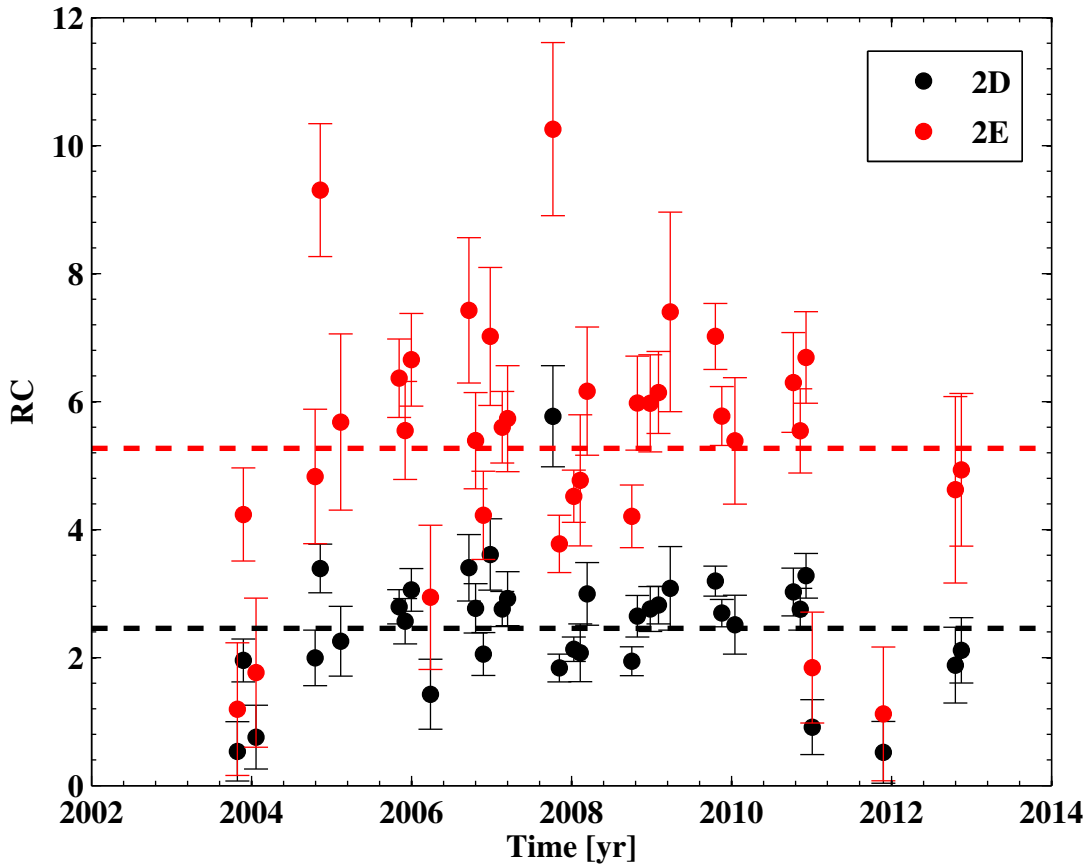

Figure 4.62: RC values of the harmonic modes $2D$ and $2E$. The dashed lines are the weighted means.

Table 4.5: Results of fitting models to the amplitudes. The parameters α and Π are the amplitude and period of a periodic model.

Label	**Period** $[s]$	α $[mma]$	Π $[yr]$	$\chi^2_{Reduced}$
$2D + E_2$	93.8	0.08(0.02)	4.46	1.2
$B + D$	114.3	0.20(0.02)	3.95	7.1
$B + E_2$	119.5	$< 0.09(0.02)$	-	4.1
A_1	122.9	$< 0.19(0.04)$	-	6.1
A_2	123.0	$< 0.14(0.02)$	-	4.0
$C_3 + D$	131.8	$< 0.06(0.02)$	-	2.6
$2D$	135.9	0.23(0.02)	4.49	2.9
$D + E_2$	143.2	$< 0.13(0.02)$	-	4.3
$2E_2$	151.4	$< 0.08(0.04)$	-	1.5
B_-	197.1	0.17(0.03)	4.30	2.9
B	197.4	0.05(0.03)	4.26	1.1
B_+	197.6	0.19(0.03)	3.39	4.2
C_1	254.7	0.11(0.03)	5.26	1.9
C_2	255.9	$< 0.07(0.02)$	-	1.6
C_3	256.2	$< 0.27(0.03)$	-	10.3
D_-	271.2	0.18(0.03)	4.07	2.3
D	271.7	0.56(0.02)	4.08	3.5
D_+	272.2	0.51(0.03)	4.83	3.6
E_1	302.3	$< 0.07(0.02)$	-	1.4
E_2	302.8	0.33(0.03)	8.76	2.0
$D + E_2 - B_+$	520.1	$< 0.45(0.02)$	-	30.9
$D + E_2 - B$	521.8	0.33(0.02)	3.78	9.9
$B - E_2$	566.9	0.15(0.02)	6.38	3.9
$B_+ - E_2$	568.9	0.14(0.02)	4.07	3.7

Chapter 5

DISCUSSION

5.1 Pulsation Timing Based White Dwarf Planet Detection

We have studied the pulsation arrival times of three white dwarf pulsators. For EC 20058-5234, we have detected a signal that may be produced by a several Jupiter mass planet in a ~ 3 year orbit, but we approach this interpretation with reservation. The high temperature of this star would provide great contrast for direct imaging of a planet with the next generation infrared space telescope. We were also able to show that GD 66 does not have a planet at several AU up to about half a Jupiter mass.

This makes a handful of white dwarfs that are now known to not have Jupiter sized planets at distances greater than several AU. As planets are commonplace around main sequence stars, the lack of planets detected around white dwarfs is growing in statistical significance and importance. Our knowledge about white dwarf planetary systems can place strong constrains on post main sequence planetary and stellar evolution, and help us learn the fate of our own solar system. White dwarf pulsation-timing remains an important technique for the search of planetary systems around white dwarfs.

We also discover periodic variations in the $O - C$s of EC 20058-5234, KIC 8626021, and GD 66 that cannot be caused by planetary companions or movement of the star; the $O - C$ variations in every mode in each star do not have the same amplitude and phase. We suggest significant changes to the standards for white dwarf pulsation timing based planet detection. It would be prudent to now require detection of periodic $O-C$ variations of *multiple* modes with the same amplitude and phase. We emphasize that pulsation timing is still one of the best techniques for planet detection and that the signature of a planet would still have been detected in our analysis.

5.2 Cooling and Secular Processes

We announce the discovery of a new timescale for period change in white dwarf evolution. A similar result was simultaneously obtained for a DA by Hermes et al. (2013). The secular variations found in the majority of these $O-C$s correspond to $\dot{P}s$ orders of magnitude larger than expected from cooling. In many cases the sign of the secular variations is also in the direction one would expect if the star were actually warming up.

The secular frequency variations observed for modes A^-, A^+, and D^+ of KIC 8626021 and modes A_2 and B of GD 66 are still consistent with conventional cooling models. Mode B of EC 20058-5234 has a $\dot{P}$ of around $-2(2) \times 10^{-14}$ which is of order the expected cooling rate. For KIC 8626021 and EC 20058-5234, these modes offer the best possibility of measuring the neutrino production rate in a white dwarf. We note that the $\dot{P}s$ measured for the DAs G117-B15A and R548 are generally consistent with theoretical cooling predictions. Comparing the properties of G117-B15A and GD 66 may give insight as to what physical processes are causing these large secular variations.

5.3 Combination and Harmonic Modes

Table 2.6 compares the frequency of the combination mode $C+E$ to the additive sum of the frequencies of modes C and E for EC 20058-5234. The parent modes are varying periodically with distinct amplitudes and phases and the sum of these variations exactly matches that of the combination to high precision. This is the most conclusive evidence obtained to date that the frequency of combination modes are indeed exactly the sum of the frequencies of the parent modes. Our analysis of GD 66 also shows that the frequencies of the combination and harmonic modes are directly related to the frequencies of the parent modes. These results, in addition to the similar results recently obtained by Hermes et al. (2013), strongly support the assertion that combination and harmonic modes are not *normal* modes.

The RC values of several combination modes of GD 66 and the combination mode in KIC 8626021 were also investigated. The RC values of GD 66 do not appear

constant in time, but the deviation from constant is rather small. We suspect this is an effect of data reduction. It could also be a sign that the properties of the convection zone or the geometry of the modes is changing or failure of the theoretical model. The variations in RC did not obviously correspond to any variations seen in $O-C$. The RC values for mode $C^+ + F$ of KIC 8626021 were consistent with being constant in time which supports the theoretical predictions of Wu (2001).

5.4 Model Validation and the Reliability of the Results

Inspection of many of the fits to $O-C$ reveals that the *overall* trends in $O-C$ are well modeled although many of the fits have large reduced χ^2 statistics. For KIC 8626021, only the fits of a parabola to the $O-C$ of modes $C^+ + F$, C^-, C^+, and F are suspect. For these $O-C$s, we suspect future observations will reveal that a sinusoid plus a parabola fits very well. For GD 66, we were unable to find a suitable model for a handful of modes and we do not consider the sinusoidal fits to be reliable. These variations clearly require more complicated mathematical models.

To test the ability of other models to fit the $O-C$s, polynomials of order 5 through 10 were fit to each $O-C$. Generally the fits resulted in poorer reduced χ^2 statistics. We believe this is strong evidence that the chosen models were suitable.

For EC 20058-5234, comparison of the fits to the $O-C$s of the parent modes C and E to fits of the $O-C$ of the combination mode $C+E$ are the most compelling evidence that the chosen models are effectively fitting $O-C$ for any of the stars (see Table 2.6). The match can only be this superb if the models are effectively fitting the overall trends in the EC 20058-5234's $O-C$s *and* if the fit parameters and confidence intervals are reliable.

The generally high values of reduced χ^2 could be indicative of severe underestimation of the standard error in the $O-C$ measurements, but Monte Carlo simulations and conventional wisdom say that the quoted errors are actually slightly overestimated and cannot be underestimated by a large amount. The increase in the total squared residual caused by variations in the lightcurves that are not included in the fit generally

causes an overestimation of the standard error. The most plausible explanation is that there are high frequency or chaotic variations in $O-C$ that are not included in the models.

We again note that aliasing could have major effects on GD 66's $O-C$s, specifically for modes closely spaced in frequency. Although simulations show that we should be recovering $O-C$ effectively, these $O-C$s should be interpreted with caution.

5.5 Characterization of the Frequency and Amplitude Variations

For 2 DBs and a DA, we observe 13 normal modes with periodically varying frequencies, 18 with secular frequency variations orders of magnitude larger than expected from cooling, and 5 with quasi-periodically varying frequencies (all in GD 66). Both the secular and periodic variations are generally larger for normal modes with longer periods (perhaps with the exception of GD 66, see Figures 2.14, 3.20, and 4.32). This hints to the possibility that the physical processes behind these frequency perturbations are near the surface of the white dwarfs.

One of the most interesting ways to investigate the cause of these frequency variations is by studying the triplets. There are two cases of symmetric secular variations of triplets: 1) The $m \neq 0$ modes of the B triplet of GD 66 have large positive $\dot{f}$s and the $m = 0$ mode has an $\dot{f}$ consistent with zero and 2) the $m = 0$ mode of the A triplet of KIC 8626021 has a large negative $\dot{f}$ and the $m \neq 0$ modes have $\dot{f}$s consistent with zero. The only other triplet with reliable fits to $O-C$ is the D triplet of KIC 8626021. The $m = -1$ and $m = 0$ modes have large and similar positive $\dot{f}$s and the $m = 1$ mode has an $\dot{f}$ that is consistent with 0. The D triplet of KIC 8626021 was the only triplet measured where the frequencies of more than one component were observed to vary periodically. The frequencies of the $m = 1$ and $m = 0$ modes appear to vary periodically with similar amplitudes, periods, and phase, and the frequency of the $m = -1$ mode does not appear to vary periodically.

Some of the most interesting results are the strange behavior of several of GD 66's modes, namely C_1, C_3, D_-, D_+, and E_1. These quasi-periodic frequency variations

cannot be well modeled by sinusoidal or parabolic variations and appear to be varying on a different timescale than the observed periodic variations.

Periodic variations in amplitude were also observed. For KIC 8626021, the frequencies of 5 modes were observed to vary periodically and the amplitudes of 4 of the 5 of those modes were also observed to vary periodically (B, C, D_-, and D). The period of the amplitude variations agreed with the period of the frequency variations for modes D_- and D, but disagrees by a significant amount for modes B and C. For D and D^-, the frequency variations had a similar period and phase, while the variations in amplitude of the two modes were 180 degrees out of phase with each other. There also appears to be a quarter/three-quarter cycle difference between the phase of the amplitude and frequency variations. For GD 66, 8 of 13 of the amplitudes were observed to vary periodically. The period of the amplitude variations were generally consistent with 4-5 years, similar to the period of the observed $O - C$ variations of several of the modes. We suspect the amplitude variations could be hinting at some sort of geometric effect, such as a change in the relative inclination of the pulsations.

While there are no known physical processes that operate on $\sim$ year long timescales in white dwarfs, variations like this are not entirely unexpected. The frequencies of other (cooler) white dwarfs are known to change by much larger amounts on similar timescales (see e.g. Provencal et al. 2012, 2009). This is simply the first time the modes of a white dwarf have been monitored in such detail. Of great importance is that the modes are not disappearing and reappearing at new frequencies, but changing frequency slowly and smoothly.

Its reasonable to hypothesize that these variations in frequency are caused by magnetic activity. Pulsating white dwarfs are convective at the surface, which could lead to a dynamo driven magnetic field. Given the complexity of magnetic activity in the Sun, it would not be surprising for such complex magnetic activity to be occurring at the surface of a white dwarf. The solar magnetic cycle has been directly connected to changes in the p-mode frequencies (Libbrecht and Woodard, 1990) and those changes have been used to infer the geometry of the magnetic field variations (see e.g. Goldreich

et al. 1991). Its possible that these observations could be used to create a detailed model of magnetic field variations at the surface of these white dwarfs.

Part II

Maestro

Chapter 6

THE MAESTRO FRAMEWORK

6.1 Introduction

MAESTRO is a software framework and library written in *MatLab*. The primary purpose of *MAESTRO* is to act as a centralized software repository for the typical data reduction and analysis needs of the Whole Earth Telescope (WET). As a library, users with a *MatLab* license can exploit *MAESTRO*'s code base in their own program. Users without a *MatLab* license can still use *MAESTRO* as a stand alone application with full access to all pre-configured *directives*. In Chapter 7, the *MAESTRO*'s *REDUCE* directive will be discussed. In this chapter we will describe how to install and use *MAESTRO*, how *MAESTRO*'s files are organized, and how to create a new *MAESTRO* directive.

6.2 Installation

The recommended installation procedure is to download and execute the *MAESTRO* bash installation script. The installation script is known to work on most flavors of *GNU/Linux* and *OSX*. The script installs both *MAESTRO* and the Matlab Component Runtime (MCR). The MCR is required for users without a *MatLab* license. If performing an install as root, *MAESTRO* will be installed in **/opt/maestro** and the MCR will be installed in **/opt/MCR**. The *MAESTRO* start up script, **maestro**, and an update script, **update-maestro**, will be placed in **/usr/local/bin/**. If the install script is run as a regular user, the MCR and *MAESTRO* will be installed in **~/.Maestro/program** and **~/.MCR** (*GNU/Linux* only) respectively and the executables will be placed in **~/.Maestro/program/binaries/**. The MCR cannot be

installed as a non-root user on *OSX* and the install path is chosen by the operating system.

6.3 File Organization

MAESTRO relies on information from two distinct locations. The *root path* is where the executables, program configuration, default configuration, and source code are stored, while the *user path* is where user created configuration files and other user created files are stored. The root path is typically either **/opt/maestro** or **~/.Maestro/program/** depending on whether or not *MAESTRO* was installed as root. The user path is typically **~/.Maestro**. Note that when *MAESTRO* is executed the location of the root and user path are extracted from the environment variations *MAESTRO_ROOT_PATH* and *MAESTRO_USER_PATH* respectively.

The root path contains a file named **manifest** containing a list of available directives, a file named **version** containing the version and revision of *MAESTRO* and required version of the MCR, and a file named **default** containing the default *MAESTRO* configuration. The directory **commands** (soon to be renamed **directives**) contains the flag, argument, and help files for each *MAESTRO* directive, and the **source** directory contains the source code for *MAESTRO*. The **scripts** directory contains the start-up, compilation, installation, and update scripts, and the **user** directory contains a copy of the files that are installed to the user path the first time *MAESTRO* is executed.

The user path currently contains two directories: **config** and **fields**. The **config** directory contains all user created overrides to the MAESTRO configuration (see 6.5) and the *fields* directory contains all fields generated by the **buildfield** directive (see Chapter 7).

6.4 Using *MAESTRO*

MAESTRO is called by executing a bash start up script, named **maestro**, that automatically sets the environment variables *MAESTRO_ROOT_PATH* and *MAESTRO_USER_PATH* (among others) and invokes the actual *MAESTRO MatLab* executable. The start up script would typically be placed in the user's binary path so that from a shell the user could call *MAESTRO* by executing:

$ **maestro**

MAESTRO generally needs to be called with an instruction, or "directive". A example call of *MAESTRO* may be:

$ **maestro do_my_laundry**

where **do_my_laundry** is the directive. Every directive has its own set of possible arguments and flags. For example, **do_my_laundry** may allow the following syntax:

$ **maestro do_my_laundry hot cold --use-dryer-sheet --detergent Tide**

where **hot** and **cold** are arguments to the directive and **--use-dryer-sheet** and **--detergent** are flags. The flag **--detergent** has a flag argument **Tide**. All directives follow the same general format of arguments followed by naked flags and flags with arguments. A directive may or may not require arguments or flags. In addition to directive specific flags, *MAESTRO* has universal flags that can be specified for all directives. A list of universal flags is shown in Table 6.1. The *MAESTRO* source code can be used as a library by simply adding the source code to the *MatLab* path. Some functionality will require setting the environment variables *MAESTRO_ROOT_PATH* and *MAESTRO_USER_PATH* to their appropriate values.

Table 6.1: *MAESTRO*'s universal flags as of version 00.91.35. These flags can be specified when calling any directive.

Flag	**Effect**
-c [name], --config [name]	Load a configuration file with the given name.
-h, --help	Display general help for MAESTRO.
-q, --quiet	Suppress all output.
-v, --verbose	Display extra output.
-l, --loud	Display as much output as possible. Typically used for debugging.
--version	Write out the running version of MAESTRO.
--about	Write out a short block of text about MAESTRO.
--copyright	Display licensing and copyright information.
-d, --debug	Turns on debug mode (for development only).

6.5 Configuration

Many *MAESTRO* directives rely upon global configuration options. Each configuration option has a *name* and *value*. The values of any configuration object are simple *MatLab* built-in types. For example, the configuration option *SPAM* might be set to the string "EGGS". The default value for all configuration options are defined in the file **default** in the root path. This file should not be modified. To modify a configuration option, one should create a file in the user path in the **config** sub-directory. A configuration file that specifies a configuration option *NUMBER_LLAMAS* to have the value 4.2 and the configuration option *RETICULATING* to have the value *'splines'* would have the contents:

```
NUMBER_LLAMAS 4.2
RETICULATING 'splines'
```

If this configuration file is specified when *MAESTRO* is called, the value of the the default configuration options will be overwritten to the values given in this file. If more than one configuration file is specified, the configuration options are overwritten in the order of specification.

6.6 Creating a New Directive

6.6.1 The Manifest

The first step to adding a new directive is adding it to the manifest file. The manifest provides *MAESTRO* with a list of enabled directives. The first column of the manifest is the name of the *MatLab* function to be executed when the directive is called. The second column is a quotation enclosed space separated list of aliases for the directive, e.g. the directive **cal** can be the same directive as **calibrate**. The third column is a quotation enclosed description. An example configuration line for a new directive could be:

f_do_make_spam_and_eggs "cook" "Make breakfast."

6.6.2 Argument Configuration

A file with the same name as each directive and extension "args" located in the **commands** sub-directory configures the arguments that can be passed to the directive. For example, the argument configuration for directive **visualize_whirled_peas** would be located in the **commands** sub-directory and named **visualize_whirled_peas.args**. A template for a argument configuration file named **ARGS_TEMPLATE** is located in the **commands** sub-directory. An argument configuration does not need to be supplied for a directive.

Each configured argument will be captured in the order they are specified. In other words, if two arguments are configured for a directive, the first argument supplied to the directive and the second argument supplied to the directive will be assumed to correspond to the first and second configured argument. The exception is that all arguments supplied beyond the number of configured arguments will be treated as belonging to the last configured argument.

An argument configuration file contains a line for each configured argument. The first column is the internal name for the argument, used when referencing the value in the code. The second column is the class of the argument, which can be set to **string**, **boolean**, or **double**. *MAESTRO* will parse the argument into one of these classes. The third column should be set to **True** or **False** to indicate if the argument is required. If set to **True** and no corresponding argument is supplied, an error will be thrown. The last two columns are a double quote enclosed human name and description for the argument that are used to display syntax help for the directive.

6.6.3 Flag Configuration

Like the configuration for an argument, a file with the same name as each directive and extension "flags" located in the **commands** sub-directory configures the

flags that can be passed to the directive. For example, the flag configuration to the directive **rinse_and_repeat** would be named "rinse_and_repeat.flags". A template for a flag configuration named FLAGS_TEMPLATE is located in the **commands** sub-directory. A flag configuration does not need to be supplied for a directive.

Each line in the flag configuration file corresponds to a flag that can be set for a directive. The first column is an internal name for the flag, used for accessing the data internally. The second column is the class of the argument, which can be set to **string**, **boolean**, or **double**. *MAESTRO* will parse the flag data into one of these classes. The next two columns are the short and long flag names, e.g. "-d" and "–digest-lean-beef". The next column is the number of args that can be given to the flag. There are a few special values for the number of allowed flags. A -1 requires that one or more arguments to the flag be specified (perhaps a list of files) and -2 forces there to either be 0 or 1 arguments to the flag. If either of these conditions are not met, an error occurs and the syntax of the directive is shown to the user. The next column is the flag group. If a flag is assigned to a flag group other than zero, only one flag from this group can be specified otherwise an error is displayed to the user. The last two columns are a double quote enclosed brief description of the flag and help text displayed with syntax help for the directive.

6.6.4 Help Text Configuration

In addition to the brief description provided in the manifest, a more complete description or help text can be configured by providing a directive help file. The help text is provided in a file in the "commands" sub-directory of the root path with a base filename identical to the directive name with the extension "help". The file should contain a double quote enclosed block of text.

6.6.5 Creating the *MatLab* Code for a Directive

The code for a directive should be a *MatLab* function with the same name as supplied by in the manifest. *MAESTRO* will parse the arguments and flags to the

directive and pass an structure as a single argument to the function. This structure will contain fields containing the value of every configured flag and argument. The name of the fields are the internal names specified in the flag and argument configuration. All flags and arguments configured as class *string* will be passed as a cell array of strings or if the flag or argument was not specified, an empty cell array. Those configured as class *boolean* will be a logical array. If a *boolean* class flag or argument was not specified, its value will be false. If a *boolean* class flag with 0 arguments was specified its value will be *true*. Flags and arguments configured as class *double* will be passed as arrays containing the specified values, or by default the empty array *[]*. The structure will also contain fields with information such as the raw call to the directive. The directive should return no values, and throw an exception if an error occurs.

Chapter 7

THE MAESTRO REDUCTION ALGORITHM

7.1 Introduction

Twenty years ago the majority of photometric observations were made with photomultiplier tubes (PMTs). A raw PMT observation consists of the integrated flux recorded from a solid angle of the sky. For PMT observations, *reducing* the observations into a record of flux vs time, i.e. a *lightcurve*, is as trivial as a few simple mathematical operations. In contrast, a modern digital image recorded with a charge-coupled device (CCD) consists of a measurement of the flux of hundreds of thousands to millions of overlapping solid angles of the sky. This dramatic increase in the sheer quantity of information recorded, combined with the superior efficiency of CCDs has revolutionized astronomy. It has, however, also dramatically increased the complexity of data reduction. For a comprehensive review of the general procedures and techniques involved in CCD data reduction see Howell (2006).

In this chapter, we describe a nearly robotic reduction algorithm implemented as a *MAESTRO* directive named *REDUCE*. The goal of *REDUCE* is for the following set of instructions to produce tables of flux versus Julian day for each star in a set of images:

Step 1: Go to the directory where the images are stored.

Step 2: Type "maestro reduce".

Step 3: Press return.

This goal may seem unreasonable, but *REDUCE* is now close to being a fully robotic reduction algorithm. Only in situations where a complicated setup is required,

such as multiple amplifiers, overscan and trim regions, non-standard Flexible Image Transport System (FITS) headers, or extremely strange observing conditions (e.g. seeing $\ll 1$ pixel) does *REDUCE* require the user to specify a few configuration options.

We would like to make a special note that we do not purport to have developed a new method for increasing the *quality* of extracted lightcurves, although in some situations *REDUCE* does outperform existing professional scientific software for various reasons. The actual methods used to extract the flux, e.g. weighted vs unweighted apertures, growth curves (Stetson, 1987), other forms of point spread function (PSF) photometry, etc, are not discussed here in any detail. Some discussion of these techniques can be found in Mighell (1999).

REDUCE was developed by the Whole Earth Telescope (WET) (Nather et al., 1990), a collaboration of currently ≈ 30 observatories around the world. One to two times per year the telescopes of the WET observe a set of pre-selected targets for a period of a few weeks to months. Given the number of factors to consider (weather conditions, instrument setup, telescope size, the sometimes unpredictable behavior of the targets, etc.) target allocation is critical. For this reason, the observatories send their images to WET headquarters (HQ) for immediate analysis. During recent WET campaigns hundreds of thousands of photometric observations are received at HQ from observatories all over the world. During the peak of a campaign, as many as $15,000$ images from 20 observatories are received in a single day. Though the 2009 campaign, "XCOV27", calibration and reduction were performed with a series of Image Reduction and Analysis Facility (IRAF) scripts. At the height of XCOV27, HQ was staffed 24 hours a day for close to a week. During the daytime at the height of the run, 8-10 astronomers worked feverishly to process the stream of incoming data. Even with this much manpower, keeping up was a challenge. There was typically a list of data sets that even the best reduction experts would spend hours re-reducing, whether it was because of poor signal to noise (S/N), poor tracking, or other quirks in the data. At the time of XCOV27, the WET was developing *REDUCE* to replace the existing IRAF scripts. In 2010, a graphical alpha version of *REDUCE* (Dalessio, 2010) was tested in

parallel with the existing IRAF scripts. What became apparent was that in the same amount of time, one person could reduce with *REDUCE* what many people could with IRAF. When situations arose where the IRAF scripts became unstable, particularly in cases of low S/N, *REDUCE* had few problems. In 2011, a command line version of *REDUCE* was deployed as the primary photometry reduction algorithm for the WET. *REDUCE* currently performs the calibration and reduction of the images, while WQED (Thompson and Mullally, 2009) performs all post reduction processing. In late 2011, the largest WET campaign in history was handled quite easily by four HQ staffers working a few hours a day. The number of required HQ staffers is now at the minimum for 24 hour communication with the participating observatories, which is more than a 5 fold reduction in manpower during the peak of a run. Thanks to *REDUCE* and web teleconferencing, the next planned campaign, XCOV29, will be the first major WET run with no centralized HQ. *REDUCE* is now being used by astronomers all over the world.

For a robust algorithm like *REDUCE*, proper differentiation between reliable and unreliable quantities is paramount. The most critical pieces of information *REDUCE* relies upon are the readnoise and gain of each amplifier of the CCD. The readnoise is the noise associated with the readout process and the gain is the number of analog-to-digital units (ADU) recorded per electron registered by the CCD. The readnoise and gain of the CCD allow the statistics to be propagated through any operation, which is critical for creating normalized configuration parameters that work under all conceivable conditions. In cases where a large sample of non-robust measurements are present, quantities like the median and corrected median absolute deviation (CMAD) are used in lieu of the mean and standard deviation.

The two most difficult challenges of automated CCD photometry reduction are finding stars and reliably tracking stars from image to image. In this Chapter, we will describe how *REDUCE* solves these problems and provide an outline of the algorithm in general.

7.2 Preparation

Before the reduction can begin, *REDUCE* obtains a list of science and calibration FITS files and extracts information about the files from the FITS headers. There are four types of files that can be specified: biases, darks, flats, and science images. Only science images are required with one caveat: the gain and readnoise must be specified if flats and biases are not provided. *REDUCE* allows files to be specified manually, specified through contextually named filelists, or discovered automatically. Manual specification is accomplished providing a flag to the command, such as "-b" or "–bias-files" followed by a list of files or a wildcard selector, i.e. **.fits*. An apetail symbol (@) can also be prefixed to a filename to indicate that it is an ASCII list of files. Files can be specified prefixed with their relative or absolute path. *REDUCE* will automatically look for contextually named ASCII lists of files named *biaslist*, *darklist*, *flatlist*, and *objlist* (and other variations) if no files are specified manually. This is the preferred method of the WET. On a UNIX-like system generating these files is simple, i.e. "ls bias*.fits > biaslist". The third method, automatic detection, is still experimental. This method finds FITS files in the working directory and checks the FITS header keywords like *IMAGETYPE* for specific values such as "BIAS" or "FLAT".

There are several critical pieces of information that must be extracted from the FITS header. The first piece of information is the images size and format. This is specified with the BITPIX, NAXIS, NAXISX, and related keywords. These keywords are standard, and its extremely uncommon to find a FITS image that doesn't conform to at least this structure. The real challenge is extracting the date and time of the midpoint of the exposure. The problem is that there is no standardization followed for the date and time. Many observatories use DATE-OBS to store the date and time, while some use DATE-OBS for the date and TIME-OBS for the time. Other commonly used keywords are UTSHUT and GPS-TIME. The date can come in two flavors, *yyyy/mm/dd* and *yyyy-mm-dd* and the time can sometimes include fractional seconds. In almost every case, the time is recorded at the open of the CCD shutter, so half of the exposure time can be added to the time to get the average, or midpoint

time of the of the exposure.

REDUCE uses a clever scheme to handle all of these permutations of keywords and formats. We will illustrate this scheme with an example. The configuration option *FITS_HEADER_TIME_KEYWORD* is a *MatLab* cell array of keywords specified in order of priority. For each image, *REDUCE* cycles through these keywords and attempts to extract the corresponding value from the FITS header. If the keyword is found, *REDUCE* then cycles through the regular expressions in *FITS_HEADER_TIME_REGEXP* looking for matching regular expressions in the keyword's value. If a match is found, *REDUCE* cycles through the time formats in *FITS_HEADER_TIME_FORMAT* trying to match the format of the time. This process is also repeated for the date, exposure time, and darktime. If a unique keyword, regexp, or format is encountered, the user can overwrite these configuration options accordingly. There is also a configuration option *FITS_HEADER_SHUTTER_STATE* that can be set to *mid* or *closed* if the time and date are not specified at shutter open.

7.3 Calibration

There are three types of calibration images: biases, darks, and flats. A bias is a measurement of the counts caused by powering on the CCD and a dark is a measurement of thermal noise recorded with the shutter closed. The flat contains information about individual pixel response to incident photons. In this section we will describe the general process *REDUCE* follows to obtain calibration images and how *REDUCE* applies calibrations to science frames. We will also discuss how the readnoise and gain of the CCD are calculated, how overscan is applied, and how images are trimmed.

Overscan is a region of an image that does not physically exist on the CCD. Overscan is measured when the amplifier is instructed to read without any charge being transferred from the CCD. The difference between pixel values in the overscan region are due to the noise of the readout process, i.e. the *readnoise*, and the mean value of the overscan region is representative of a bulk offset in pixel values due to CCD electronics.

In cases where multiple amplifiers are used to produce a single image, multiple overscan regions may be present. Many modern CCDs do not record an overscan region because the electronic offset can be held stable throughout a night's observations. If the *MAESTRO* configuration option *APPLY_OVERSCAN* is set to *true*, *REDUCE* will use the configuration options *OS_SOURCE_REGION* and *OS_APPLY_REGION* to apply overscan to all FITS images. *OS_SOURCE_REGION* and *OS_APPLY_REGION* are be *MatLab* arrays of numbers of length $4N$, where N is the number of overscan regions. Each set of 4 numbers specifies the boundaries of a rectangular region with the form $X_{min}, X_{max}, Y_{min}, Y_{max}$. The mean of each source region is calculated and subtracted from each application region. Note that this does not remove the overscan regions from the images. Removal of the overscan regions must be accomplishing with trimming.

Regions of CCD images may need to be removed or geometrically rearranged, i.e. *trimmed*. The most common uses for trimming is to remove parts of the CCD image that contain spurious or non-physical values but the overscan region(s) should also be removed from the images. For *REDUCE*, trimming is generalized as a series of copy and paste operations of rectangular sections from a raw image into an initially empty image. Instead of removing unwanted sections from a raw image, the wanted sections are copied into a new image. If the *MAESTRO* configuration option *TRIM_IMAGE* is set to *true*, the configuration options *TRIM_CUT_REGION* and *TRIM_PASTE_-REGION* specify sets of copy and paste sections respectively. *TRIM_CUT_REGION* and *TRIM_PASTE_REGION* are *MatLab* arrays of length $4N$, where N is the number of sections to be copied to the trimmed image. Each set of 4 numbers specifies a rectangular region of the form $X_{min}, X_{max}, Y_{min}, Y_{max}$. The coordinates for the paste sections are be relative to the origin of the new trimmed image. Each trim section is assumed to have its own readnoise and gain.

A zero second CCD image, or a *bias*, can be used to characterize variations around the mean electronic background. Multiple bias images can be combined to form a master bias with minimal effective readnoise. If no overscan has been applied, the master bias will contain both the bulk offset and the pixel to pixel variations caused

by powering the CCD. The master bias can be subtracted from all subsequent images to remove these effects. Note that by default, *REDUCE* calculates the master bias by mean combining the bias images.

REDUCE uses the pixel to pixel statistics throughout the reduction process; quantifying the readnoise is absolutely critical. The readnoise is estimated for each pixel from the sample standard deviation of the set of bias images. The median estimated readnoise of each trim section is recorded as the readnoise (in ADU) of each trim section. If no bias files are specified, the readnoise for each trim section must be specified in the *MAESTRO* configuration option *READNOISE* as an array of *MatLab* double precision numbers.

As a CCD image is exposed, random thermal variations cause charge accumulation that is not due to incident photons. There may be pixel to pixel variations in the rate these thermal counts accumulate. A *dark* image is an exposure where the CCD shutter remains closed so that no photons are striking the CCD. These images are combined into a master dark that is used to remove the effects of thermal charge accumulation from other images. *REDUCE* first subtracts the master bias from all dark images and divides the resultant images by their exposure time. These images are median combined to form a one second master dark.

Each pixel on the CCD has a slightly different response to incident photons. The optical system of the telescope may also contain dust or have other problems that cause non-uniform illumination of the CCD. To correct for these effects, images can be corrected by a master *flat*. To take a flat, the optical system is subjected to a uniform photon source and a CCD image is exposed. *REDUCE* subtracts the master bias and a master dark of the same effective exposure time from each flat image. Every pixel of each flat is divided by the median pixel value of the flat to create a normal response of unity. The flats are then median combined to form the master flat.

The gain of the CCD is the ratio between how many electrons were stored on each pixel of the CCD and the number of counts that are actually registered by the amplifier. The gain must be known to calculate statistical counting errors. Gain images

can be created by subtracting the master flat from the individual normalized flats. If each normalized flat had some median pixel value F and the readnoise (in ADU) is known to be σ_{RN}, the typical sample standard error of the gain image, σ_{Tot} is predicted to be

$$\sigma_{Tot}^2 G^2 F^2 \frac{N}{N-1} = G^2 \sigma_{RN}^2 + GF \tag{7.1}$$

where G is the gain in $\frac{e^-}{ADU}$ and N is the number of flat images. This can be solved for the gain, resulting in the following relationship:

$$G = \frac{F}{\sigma_{Total}^2 \frac{N}{N-1} - \sigma_{RN}^2} \tag{7.2}$$

REDUCE calculates the gain of each trim section from a series of flat images. If no flats are specified, the gain of each trim section must be specified in the *MAESTRO* configuration option *GAIN* as an array of *MatLab* double precision numbers.

Once the master bias, master dark, and master flat have been created and the gain and readnoise have been calculated, the calibrations are ready to be applied to science frames. For each science frame, overscan and trimming are applied first. Then the master bias and an appropriately scaled master dark are subtracted. Finally, the image is divided by the master flat. The resultant image is referred to as a *calibrated image*. For every calibrated image, a *noise image* is also created by calculating the formal standard error for each value of the calibrated image.

7.4 Star Finding

REDUCE is capable of identifying stars on any individual image with no user interaction. Given a calibrated image, the first step is to estimate and remove the background flux. *REDUCE* divides the image into a number of rectangular sections (the default is 4). The median, median absolute deviation, and the average coordinate of each section are calculated. A two dimensional first order polynomial is fit to these measurements. The fit is used to generate a background flux image which is then subtracted from the calibrated image. The resultant image is the gain corrected to form the *signal image*.

The next step is to calculate the standard error of each value in the signal image. The standard errors are assumed to be dominated by a combination of counting statistics and readnoise. The error in the master bias level or overscan mean is typically much smaller and is ignored, as is the noise introduced in mathematical operations such as application of the master flat, subtraction of the dark, and the removal of background flux. The square root of the gain corrected, bias subtracted raw image is added in quadrature to the gain corrected readnoise to obtain the *noise image*. The signal image is divided by the noise image to produce the *signal to noise image*. All pixels above a critical signal to noise threshold (the default is 8), that contain the largest signal within several pixels, are recorded as possible star positions. A 2-D circular Gaussian is then fit to the pixels surrounding every possible star. The 2-D circular Gaussian is mathematically described as

$$f(x, y) = Ae^{\frac{(x-x_0)^2+(y-y_0)^2}{2\sigma^2}} \tag{7.3}$$

Before the fitting process begins, the pixels adjacent to each possible star position are used to estimate the standard deviation by performing a linear fit of Equation 7.3 with the star position fixed. If the estimated value of σ lies outside of tolerable levels, the possible star is removed from analysis. The tolerable levels are set to exclude cosmic rays and extended objects like galaxies and planetary nebulae. Only pixels that lie within a few times the estimated σ (the default is 2) are included in the non-linear fit. If the fit does not converge, or if x_0, y_0, A, or σ lie outside of some reasonable tolerance limits, the possible star location is removed from analysis.

The resultant fit parameters and their standard errors are assembled into a table of data called a *field*. A field is essentially a list of possible star locations. The entries in an individual field are not considered robust, but rather are considered *likely* locations and estimated fluxes of stars on a particular image. In other words, it is not assumed that every star on an image is in the field or that every star in the field is actually a star on a particular image.

7.5 Tracking Stars

REDUCE does not attempt to use the previous known locations of stars to predict or guess the location of stars on any other images. *REDUCE* also does not use cross correlation techniques to identify similar patterns of stars between images. Instead, *REDUCE* aligns two images by searching for matching geometric patterns in their fields. This method is extremely robust because it makes no assumptions about the absolute position of stars on an image, just that the relative geometry does not change. This method is only unsuitable for objects with extremely high proper motion (e.g. asteroids).

We refer to this method as *geometric hashing*. A geometric hash is formed by first selecting some permutation of stars in a field. The coordinates of these stars are shifted so that the average coordinate is the origin. The order of appearance of stars in the hash is rearranged so that the stars increase in distance from the origin. The hash can be made rotationally invariant by rotating so that the most distant star from the origin lies along the x axis. The hash can also be made scale invariant by normalizing the coordinates so that the average distance of the stars from the origin is unity.

The final hash consists of an ordered set of x and y coordinates and corresponding standard errors. Hashes from two fields can be compared to identify similar patterns of stars in the two images. The chance of hash collision, or a statistically good match of hashes that do not correspond to the same group of stars, depends strongly on the number of stars in the hash minus the degrees of freedom of the hash, and the number of stars on image. *REDUCE* starts with permutations of many stars, and as long as the number of stars in each hash is greater than 4, collisions are very rarely experienced.

A flaw in this method is that it becomes ineffective when the number of stars on image is small. If only 2 stars are on the images, all rotationally invariant hashes are degenerate. In these cases, *REDUCE* uses the relative flux or known absolute positions of the stars to break the degeneracy. However, the relative flux is not very robust, especially considering photometric observations are typically of objects with

variable flux. If the only two objects on image have similar fluxes *REDUCE* will not reliably align the images if the telescope undergoes 180 degree rotation.

7.6 First Pass: Building a Master Field

REDUCE makes two passes through the set of science images. During the first pass, *REDUCE* extracts a field from each image and aggregates this information into a *master field*. The master field is a *robust* table of relative star positions and brightnesses found throughout the entire series of images.

To build the master field, every image is visited once in a random order. The field of stars on the first image encountered is recorded as the initial master field. Each star in the master field is given some set of weights representative of the statistical likelihood that the star actually exists. A match is searched for between the field of stars from each subsequent image and the master field. Matches are searched for in order of statistical uniqueness.

When a match is found and confirmed, master field stars that were part of the matching hash are given a large statistical boost. Master field stars that were not part of the matching hash, but given the geometric transformation applied to each hash, are found to align with a star in the image's field, are given a moderate statistical boost. Master field stars that should have been found but were not found on the image are statistically punished. Each time a match is found, the position of the stars and the standard error in the position is updated (along with other quantities, such as the average FWHM and flux). During each iteration, stars whose statistical likelihoods dip below some nominal threshold are removed from the master field, and any new stars encountered are added to the master field with nominal statistical likelihood. After all of the images have been visited once, all stars below some critical statistical likelihood are removed from the master field. At this point the "real" stars have distinguished themselves from a few statistical aberrations. The final master field is then considered a robust representation of the stars that exist on the series of science images.

While the master field was being built, the rotation angles required to align all selected matching hashes were recorded. After the master field is built, its determined if the images are non-rotating, rotating, or undergo a 180 degree flip at some point during the observations.

7.7 Second Pass: Final Alignment and Photometry

During the second pass, a hash match is searched for between the now robust master field, and the field from each image. The allowed rotation between the master field and individual fields is forced to obey the general pattern observed while building the master field. When a match is found, the position of stars in the master field are projected onto the image. For each projected star, a 2-D circular Gaussian is non-linearly fit to the image to obtain a precise measurement of the location of the star.

The background flux is calculated for each star using an annulus centered on the star's position. By default the boundaries of the annulus are $8\sigma - 12\sigma$. All pixels contained within this annulus that are greater than 2 times the CMAD are removed, and the mean of the remaining values is considered the background sky flux per pixel.

With the background flux and the center of each star on the image known, the fluxes can now be extracted. *REDUCE* currently calculates the flux using synthetic apertures with uniform weights per unit area, i.e. "top hat" or unweighted apertures. A series of apertures are used ranging from 2 to 20 pixels in size. For each aperture, each pixel value is multiplied by the fractional area of each pixel that is enclosed by the aperture and summed. The background flux is multiplied by the total area enclosed by the aperture and this value is then subtracted. This value is recorded for each aperture and each star. For pixels that are partially enclosed by the aperture, the fractional area is approximated with a 3 or 4 sided polygon.

After all images have been processed, files are created for each sized aperture. The first column of these files contains the Julian day of the image. The remaining columns contain the measured flux of every star in the master field. A similar file is

recorded for every aperture containing the standard error in each flux measurement. *REDUCE* also calculates which aperture has the highest formal signal to noise (SN) and indicates this with a special filename.

7.8 Typical Workflow Using *BUILDFIELD*

By default, the order of stars in the output of *REDUCE* is not pre-determined. This can make it difficult to identify which columns of the output files correspond to which star. A practical solution is to feed *REDUCE* a list of sorted star positions, i.e. a field file. The field file controls the order of stars in the output files.

The *BUILDFIELD* directive is intended to assist users in creating a field file. *BUILDFIELD* is essentially a truncated version of *REDUCE*. Instead of revisiting the images after building the master field, *BUILDFIELD* simply writes the master field to file. The file will be written to the *MAESTRO* user directory in the sub-directory **fields**. This file contains a single hash commented line containing the name of the FITS file the star coordinates are relative to, followed by a list of labels, coordinates, and fluxes of stars found in the set of images. The order of appearance of stars in this file can be rearranged to the users liking.

When *REDUCE* is called, the **-s** flag can be used to specify a field file. When *REDUCE* builds a master field during the first pass, the master field will be resorted to match the order of the stars in the field file. Note that a single field can be used to reduce multiple nights of data as long as the plate scale remains fixed. Testing with scale invariant hash matching of fields led to too many hash collisions.

A star that is not detected by *REDUCE* can also be added to a reduction using a field file. By simply adding a line to the file with the stars coordinates, *REDUCE* will now perform photometry at the specified position. This can be particular useful when a star has an extremely low signal to noise.

Bibliography

Barlow, B. N., Dunlap, B. H., Clemens, J. C., Reichart, D. E., Ivarsen, K. M., Lacluyze, A. P., Haislip, J. B., and Nysewander, M. C. (2011). Fortnightly fluctuations in the O-C diagram of CS 1246. MNRAS, 414:3434–3443.

Bischoff-Kim, A., Montgomery, M. H., and Winget, D. E. (2008). Strong Limits on the DFSZ Axion Mass with G117-B15A. ApJ, 675:1512–1517.

Bischoff-Kim, A. and Østensen, R. H. (2011). Asteroseismology of the Kepler Field DBV White Dwarf. It is a Hot One. ApJ, 742:L16.

Bradley, P. A., Winget, D. E., and Wood, M. A. (1992). Maximum rates of period change for DA white dwarf models with carbon and oxygen cores. ApJ, 391:L33–L36.

Brickhill, A. J. (1992). The Pulsations of Zz-Ceti Stars - Part Six - the Amplitude Spectra. MNRAS, 259:529.

Córsico, A. H. and Althaus, L. G. (2004). The rate of period change in pulsating DB white dwarf stars. A&A, 428:159–170.

Córsico, A. H., Althaus, L. G., Miller Bertolami, M. M., and Bischoff-Kim, A. (2012). Asteroseismology of the Kepler V777 Herculis variable white dwarf with fully evolutionary models. A&A, 541:A42.

Dalessio, J. (2010). MAESTRO: A Software Package Specializing in the Analysis of Pulsating White Dwarfs. In *American Astronomical Society Meeting Abstracts 215*, volume 42 of *Bulletin of the American Astronomical Society*, pages 452.09–+.

Dalessio, J., Sullivan, D. J., Provencal, J. L., Shipman, H. L., Sullivan, T., Kilkenny, D., Fraga, L., and Sefako, R. (2013). Periodic Variations in the O - C Diagrams of Five Pulsation Frequencies of the DB White Dwarf EC 20058-5234. ApJ, 765:5.

Dufour, P., Fontaine, G., Liebert, J., Schmidt, G. D., and Behara, N. (2008). Hot DQ White Dwarfs: Something Different. ApJ, 683:978–989.

Eastman, J., Siverd, R., and Gaudi, B. S. (2010). Achieving Better Than 1 Minute Accuracy in the Heliocentric and Barycentric Julian Dates. PASP, 122:935–946.

Farihi, J., Subasavage, J. P., Nelan, E. P., Harris, H. C., Dahn, C. C., Nordhaus, J., and Spiegel, D. S. (2012). Precision astrometry of the exoplanet host candidate GD 66. MNRAS, 424:519–523.

Fontaine, G., Bergeron, P., Lacombe, P., Lamontagne, R., and Talon, A. (1985). Stromgren photometry of ZZ Ceti and other DA white dwarfs. AJ, 90:1094–1103.

Gianninas, A., Bergeron, P., and Fontaine, G. (2005). Toward an Empirical Determination of the ZZ Ceti Instability Strip. ApJ, 631:1100–1112.

Goldreich, P., Murray, N., Willette, G., and Kumar, P. (1991). Implications of solar p-mode frequency shifts. ApJ, 370:752–762.

Handler, G., Romero-Colmenero, E., and Montgomery, M. H. (2002). The unusual pulsation spectrum of the cool ZZ Ceti star HS0507+0434B. MNRAS, 335:399–409.

Hermes, J. J., Montgomery, M. H., Mullally, F., Winget, D. E., and Bischoff-Kim, A. (2013). A New Timescale for Period Change in the Pulsating DA White Dwarf WD 0111+0018. ApJ, 766:42.

Hermes, J. J., Montgomery, M. H., Winget, D. E., Brown, W. R., Kilic, M., and Kenyon, S. J. (2012). SDSS J184037.78+642312.3: The First Pulsating Extremely Low Mass White Dwarf. ApJ, 750:L28.

Howell, S. B. (2006). *Handbook of CCD Astronomy.*

Isern, J., Hernanz, M., and Garcia-Berro, E. (1992). Axion cooling of white dwarfs. ApJ, 392:L23–L25.

Jenkins, J. M., Caldwell, D. A., Chandrasekaran, H., Twicken, J. D., Bryson, S. T., Quintana, E. V., Clarke, B. D., Li, J., Allen, C., Tenenbaum, P., Wu, H., Klaus, T. C., Middour, C. K., Cote, M. T., McCauliff, S., Girouard, F. R., Gunter, J. P., Wohler, B., Sommers, J., Hall, J. R., Uddin, A. K., Wu, M. S., Bhavsar, P. A., Van Cleve, J., Pletcher, D. L., Dotson, J. A., Haas, M. R., Gilliland, R. L., Koch, D. G., and Borucki, W. J. (2010). Overview of the Kepler Science Processing Pipeline. ApJ, 713:L87–L91.

Kepler, S. O., Costa, J. E. S., Castanheira, B. G., Winget, D. E., Mullally, F., Nather, R. E., Kilic, M., von Hippel, T., Mukadam, A. S., and Sullivan, D. J. (2005). Measuring the Evolution of the Most Stable Optical Clock G 117-B15A. ApJ, 634:1311–1318.

Koen, C., O'Donoghue, D., Stobie, R. S., Kilkenny, D., and Ashley, R. (1995). EC 20058-5234, a low-amplitude pulsating DB white dwarf. MNRAS, 277:913–919.

Libbrecht, K. G. and Woodard, M. F. (1990). Solar-cycle effects on solar oscillation frequencies. Nature, 345:779–782.

Mestel, L. (1952). On the theory of white dwarf stars. I. The energy sources of white dwarfs. MNRAS, 112:583.

Mighell, K. J. (1999). Algorithms for CCD Stellar Photometry. In Mehringer, D. M., Plante, R. L., and Roberts, D. A., editors, *Astronomical Data Analysis Software and Systems VIII*, volume 172 of *Astronomical Society of the Pacific Conference Series*, page 317.

Montgomery, M. H. (2005). A New Technique for Probing Convection in Pulsating White Dwarf Stars. ApJ, 633:1142–1149.

Montgomery, M. H., Williams, K. A., Winget, D. E., Dufour, P., De Gennaro, S., and Liebert, J. (2008). SDSS J142625.71+575218.3: A Prototype for a New Class of Variable White Dwarf. ApJ, 678:L51–L54.

Mukadam, A. S., Kim, A., Fraser, O., Winget, D. E., Kepler, S. O., Sullivan, D. J., Reaves, D., Robinson, E. L., von Hippel, T., Mullally, F., Shipman, H., Thompson, S. E., Silvestri, N. M., and Hynes, R. I. (2009). Watching ZZ Ceti evolve. *Journal of Physics Conference Series*, 172(1):012074.

Mullally, F., Reach, W. T., De Gennaro, S., and Burrows, A. (2009). Spitzer Planet Limits Around the Pulsating White Dwarf GD66. ApJ, 694:327–331.

Mullally, F., Winget, D. E., De Gennaro, S., Jeffery, E., Thompson, S. E., Chandler, D., and Kepler, S. O. (2008). Limits on Planets around Pulsating White Dwarf Stars. ApJ, 676:573–583.

Nather, R. E. (1989). The whole earth telescope. In Wegner, G., editor, *IAU Colloq. 114: White Dwarfs*, volume 328 of *Lecture Notes in Physics, Berlin Springer Verlag*, pages 109–114.

Nather, R. E., Winget, D. E., Clemens, J. C., Hansen, C. J., and Hine, B. P. (1990). The whole earth telescope - A new astronomical instrument. ApJ, 361:309–317.

Nitta, A., Kleinman, S. J., Krzesinski, J., Kepler, S. O., Metcalfe, T. S., Mukadam, A. S., Mullally, F., Nather, R. E., Sullivan, D. J., Thompson, S. E., and Winget, D. E. (2009). New Pulsating DB White Dwarf Stars from the Sloan Digital Sky Survey. ApJ, 690:560–565.

Oksala, M. E., Wade, G. A., Townsend, R. H. D., Owocki, S. P., Kochukhov, O., Neiner, C., Alecian, E., and Grunhut, J. (2012). Revisiting the Rigidly Rotating Magnetosphere model for σ Ori E - I. Observations and data analysis. MNRAS, 419:959–970.

Østensen, R. H., Bloemen, S., Vučković, M., Aerts, C., Oreiro, R., Kinemuchi, K., Still, M., and Koester, D. (2011). At Last — A V777 Her Pulsator in the Kepler Field. ApJ, 736:L39.

Pajdosz, G. (1995). Non-evolutionary secular period increase in pulsating DA white dwarfs. A&A, 295:L17–L19.

Press, W. H. and Rybicki, G. B. (1989). Fast algorithm for spectral analysis of unevenly sampled data. ApJ, 338:277–280.

Provencal, J. L., Montgomery, M. H., Kanaan, A., Shipman, H. L., Childers, D., Baran, A., Kepler, S. O., Reed, M., Zhou, A., Eggen, J., Watson, T. K., Winget, D. E., Thompson, S. E., Riaz, B., Nitta, A., Kleinman, S. J., Crowe, R., Slivkoff, J., Sherard, P., Purves, N., Binder, P., Knight, R., Kim, S. ., Chen, W.-P., Yang, M., Lin, H. C., Lin, C. C., Chen, C. W., Jiang, X. J., Sergeev, A. V., Mkrtichian, D., Andreev, M., Janulis, R., Siwak, M., Zola, S., Koziel, D., Stachowski, G., Paparo, M., Bognar, Z., Handler, G., Lorenz, D., Steininger, B., Beck, P., Nagel, T., Kusterer, D., Hoffman, A., Reiff, E., Kowalski, R., Vauclair, G., Charpinet, S., Chevreton, M., Solheim, J. E., Pakstiene, E., Fraga, L., and Dalessio, J. (2009). 2006 Whole Earth Telescope Observations of GD358: A New Look at the Prototype DBV. ApJ, 693:564–585.

Provencal, J. L., Montgomery, M. H., Kanaan, A., Thompson, S. E., Dalessio, J., Shipman, H. L., Childers, D., Clemens, J. C., Rosen, R., Henrique, P., Bischoff-Kim, A., Strickland, W., Chandler, D., Walter, B., Watson, T. K., Castanheira, B., Wang, S., Handler, G., Wood, M., Vennes, S., Nemeth, P., Kepler, S. O., Reed, M., Nitta, A., Kleinman, S. J., Brown, T., Kim, S.-L., Sullivan, D., Chen, W. P., Yang, M., Shih, C. Y., Jiang, X. J., Sergeev, A. V., Maksim, A., Janulis, R., Baliyan, K. S., Vats, H. O., Zola, S., Baran, A., Winiarski, M., Ogloza, W., Paparo, M., Bognar, Z., Papics, P., Kilkenny, D., Sefako, R., Buckley, D., Loaring, N., Kniazev, A., Silvotti, R., Galleti, S., Nagel, T., Vauclair, G., Dolez, N., Fremy, J. R., Perez, J., Almenara,

J. M., and Fraga, L. (2012). Empirical Determination of Convection Parameters in White Dwarfs. I. Whole Earth Telescope Observations of EC14012-1446. ApJ, 751:91.

Shklovskii, I. S. (1970). Possible Causes of the Secular Increase in Pulsar Periods. Soviet Ast., 13:562.

Silvotti, R., Schuh, S., Janulis, R., Solheim, J.-E., Bernabei, S., Østensen, R., Oswalt, T. D., Bruni, I., Gualandi, R., Bonanno, A., Vauclair, G., Reed, M., Chen, C.-W., Leibowitz, E., Paparo, M., Baran, A., Charpinet, S., Dolez, N., Kawaler, S., Kurtz, D., Moskalik, P., Riddle, R., and Zola, S. (2007). A giant planet orbiting the 'extreme horizontal branch' star V391 Pegasi. Nature, 449:189–191.

Stetson, P. B. (1987). DAOPHOT - A computer program for crowded-field stellar photometry. PASP, 99:191–222.

Stobie, R. S., Chen, A., O'Donoghue, D., and Kilkenny, D. (1992). The Edinburgh-Cape Blue Object Survey. In Warner, B., editor, *Variable Stars and Galaxies, in honor of M. W. Feast on his retirement*, volume 30 of *Astronomical Society of the Pacific Conference Series*, page 87.

Stumpff, P. (1980). Two Self-Consistent FORTRAN Subroutines for the Computation of the Earth's Motion. A&AS, 41:1.

Sullivan, D. J. (2005). EC 20058-5234: A DBV White Dwarf and a Possible Plasmon Neutrino Detector. In D. Koester & S. Moehler, editor, *14th European Workshop on White Dwarfs*, volume 334 of *Astronomical Society of the Pacific Conference Series*, pages 495–+.

Sullivan, D. J. (2009). The long term period stability of the hot DBV white dwarf EC20058-5234. *Journal of Physics Conference Series*, 172(1):012070.

Sullivan, D. J., Metcalfe, T. S., O'Donoghue, D., Winget, D. E., Kilkenny, D., van Wyk, F., Kanaan, A., Kepler, S. O., Nitta, A., Kawaler, S. D., Montgomery, M. H.,

Nather, R. E., O'Brien, M. S., Bischoff-Kim, A., Wood, M., Jiang, X. J., Leibowitz, E. M., Ibbetson, P., Zola, S., Krzesinski, J., Pajdosz, G., Vauclair, G., Dolez, N., and Chevreton, M. (2008). Whole Earth Telescope observations of the hot helium atmosphere pulsating white dwarf EC20058-5234. MNRAS, 387:137–152.

Sullivan, D. J., Metcalfe, T. S., O'Donoghue, D., Winget, D. E., Kilkenny, D., van Wyk, F., Kanaan, A., Kepler, S. O., Nitta, A., Kawaler, S. D., Montgomery, M. H., Nather, R. E., Steeghs, D., Koester, D., Bergeron, P., O'Brien, M. S., Wood, M., Jiang, X. J., Leibowitz, E. M., Ibbetson, P., Zola, S., Krzesinski, J., Pajdosz, G., Vauclair, G., Dolez, N., and Chevreton, M. (2007). The Hottest Known DBV White Dwarf. In Napiwotzki, R. and Burleigh, M. R., editors, *15th European Workshop on White Dwarfs*, volume 372 of *Astronomical Society of the Pacific Conference Series*, page 629.

Taylor, J. H., Fowler, L. A., and McCulloch, P. M. (1979). Measurements of general relativistic effects in the binary pulsar PSR 1913+16. Nature, 277:437–440.

Thompson, M. J., Christensen-Dalsgaard, J., Miesch, M. S., and Toomre, J. (2003). The Internal Rotation of the Sun. ARA&A, 41:599–643.

Thompson, S. E. and Mullally, F. (2009). Wqed: A lightcurve analysis suite. *Journal of Physics Conference Series*, 172(1):012081–+.

Thorsett, S. E., Arzoumanian, Z., and Taylor, J. H. (1993). PSR B1620-26 - A binary radio pulsar with a planetary companion? ApJ, 412:L33–L36.

Unno, W., Osaki, Y., Ando, H., Saio, H., and Shibahashi, H. (1989). *Nonradial oscillations of stars*. University of Tokyo Press.

Van Grootel, V., Fontaine, G., Brassard, P., and Dupret, M.-A. (2013). The Newly Discovered Pulsating Low-mass White Dwarfs: An Extension of the ZZ Ceti Instability Strip. ApJ, 762:57.

Williams, K. A., Montgomery, M. H., and Winget, D. E. (2012). Time-series UV Photometry Of Two Variable Carbon-atmosphere (DQV) White Dwarfs. In *American Astronomical Society Meeting Abstracts #219*, volume 219 of *American Astronomical Society Meeting Abstracts*, page #250.03.

Winget, D. E., Robinson, E. L., Nather, R. D., and Fontaine, G. (1982). Photometric observations of GD 358 - DB white dwarfs do pulsate. ApJ, 262:L11–L15.

Winget, D. E., Sullivan, D. J., Metcalfe, T. S., Kawaler, S. D., and Montgomery, M. H. (2004). A Strong Test of Electroweak Theory Using Pulsating DB White Dwarf Stars as Plasmon Neutrino Detectors. ApJ, 602:L109–L112.

Wolszczan, A. and Frail, D. A. (1992). A planetary system around the millisecond pulsar PSR1257 + 12. Nature, 355:145–147.

Wu, Y. (2001). Combination frequencies in the Fourier spectra of white dwarfs. MNRAS, 323:248–256.

Yeates, C. M., Clemens, J. C., Thompson, S. E., and Mullally, F. (2005). Mode Identification from Combination Frequency Amplitudes in ZZ Ceti Stars. ApJ, 635:1239–1262.