To appear in the Astrophysical Journal
Preprint typeset using LaTeX style emulateapj v. 5/2/11

DISCOVERY OF A ~ 5 DAY CHARACTERISTIC TIMESCALE IN THE KEPLER POWER SPECTRUM OF ZW 229-15

R. Edelson
Department of Astronomy, University of Maryland, College Park, MD 20742-2421, USA

S. Vaughan
University of Leicester, X-ray and Observational Astronomy Group, Department of Physics and Astronomy, University Road, Leicester LE1 7RH, UK

M. Malkan
Department of Physics & Astronomy, University of California, Los Angeles, CA 90095-1547, USA

B. C. Kelly
Department of Physics, Broida Hall, University of California, Santa Barbara, CA 93106-9530, USA

K. L. Smith
Department of Astronomy, University of Maryland, College Park, MD 20742-2421, USA

P. T. Boyd
Astrophysics Science Division, NASA/GSFC, Code 660, Greenbelt, MD 20771, USA

R. Mushotzky
Department of Astronomy, University of Maryland, College Park, MD 20742-2421, USA
To appear in the Astrophysical Journal

ABSTRACT

We present time series analyses of the full Kepler dataset of Zw 229–15. This Kepler light curve — with a baseline greater than three years, composed of virtually continuous, evenly sampled 30-minute measurements — is unprecedented in its quality and precision. We utilize two methods of power spectral analysis to investigate the optical variability and search for evidence of a bend frequency associated with a characteristic optical variability timescale. Each method yields similar results. The first interpolates across data gaps to use the standard Fourier periodogram. The second, using the CARMA-based time-domain modeling technique of Kelly et al. (2014), does not need evenly-sampled data. Both methods find excess power at high frequencies that may be due to Kepler instrumental effects. More importantly both also show strong bends ($\Delta\alpha \sim 2$) at timescales of ~ 5 days, a feature similar to those seen in the X-ray PSDs of AGN but never before in the optical. This observed ~ 5 day timescale may be associated with one of several physical processes potentially responsible for the variability. A plausible association could be made with light-crossing, dynamical or thermal timescales, depending on the assumed value of the accretion disk size and on unobserved disk parameters such as α and H/R. This timescale is not consistent with the viscous timescale, which would be years in a ~ $10^7 M_\odot$ AGN such as Zw 229–15. However there must be a second bend on long (& 1 year) timescales, and that feature could be associated with the viscous timescale.

Subject headings: accretion disks, black hole physics, galaxies: active, galaxies: individual (Zw 229–15), galaxies: Seyfert

1. INTRODUCTION

The accretion disks of Active Galactic Nuclei (AGN) are much too distant to image directly. Simple physical arguments place an upper limit on the accretion disk size at ~ 10^{-2} parsecs (Hawkins 2007), corresponding to . 1 milliarcsec for even the closest AGN. Thus indirect methods must be used to study their centers. AGN's strong and rapid aperiodic optical variability provides a powerful tool for constraining physical conditions and processes. For example reverberation mapping allows estimation of Seyfert 1 emission-line region sizes and central black hole masses (Peterson et al. 2004). However recent progress has been slowed by limitations inherent in ground-based optical monitoring where it is nearly impossible to obtain continuous light curves longer than ~ 12 hr and errors better than ~ 1%.

The Kepler mission breaks these barriers with fast (30 min) sampling, high (> 90%) duty cycle, and ex-

redelson@astro.umd.edu

cellent (~ 0.1% for a 16th magnitude source) precision. Zw 229–15 (z = 0.0025; Falco et al. 1999) was observed during Quarters 4-17 (Q4-17), totaling 3.4 years. As the longest-monitored, brightest, and one of the most strongly variable Kepler AGN, it is ideal for time-series analyses that estimate the Power Spectral Density (PSD) function. Previous studies used only a small fraction of these data, finding a power law-like PSD with a steep index ($P_f \propto f^{\alpha}$; Mushotzky et al. 2011; Kelly et al. 2014) with indications of a break on ~ 90 day timescales (Carini & Ryle 2012).

This paper reports light curve extraction and PSD analysis of the full Kepler dataset. The paper is organized as follows: Section 2 reports the data reduction, Section 3 presents the PSD analyses, Section 4 discusses the theoretical implications of these analyses, and Section 5 concludes with a brief summary of this work and future plans.

2. DATA REDUCTION

Since Kepler was designed to detect exoplanets, the standard pipeline processing removes long-term trends from light curves to optimize detection of short, shallow dips. This renders data unusable for AGN, which show broad intrinsic variability power over timescales of hours to years. Thus previous Kepler Zw 229–15 studies (e.g., Mushotzky et al. 2011) used "Simple Aperture Photometry" (SAP) data from earlier in the pipeline, but this also contained uncorrected systematics due to differently-sized "optimal" extraction apertures.

2.1. Extraction from pixel data

In order to avoid these systematics, our approach starts earlier with the 2-dimensional calibrated pixel data. We used the PyKE[1] programs `kepmask` and `kepextract` to build large 32-pixel masks (see Figure 1). These larger masks mitigate excursions due to thermally-induced focus changes and to differential velocity aberration (Kinemuchi et al. (2012); see also `http://keplergo.arc.nasa.gov/PyKEprimerTPFs.shtml`). These extraction masks are identical within each season (e.g. Season 2 = Q4/8/12/16) even though the pixel downloads are not.

The resulting 3.4 year (55,653 cadence) light curve is shown in the top panel of Figure 2. The interquarter jumps arise because quarterly spacecraft rolls move the source to a different chip with a different aperture. Thanks to Kepler's stable pointing and these identical seasonal masks, the jumps between the same seasons are highly repeatable.

2.2. Interquarter scaling

Previous Kepler AGN studies dealt with interquarter jumps by restricting the analyses to single quarters (Mushotzky et al. 2011; Kelly et al. 2014), performing simple scaling to match fluxes across the gaps (Wehrle et al. 2013) or using ground-based data to normalize the offsets (Carini & Ryle 2012). Our approach starts by averaging the five good cadences immediately before and after each gap and taking the difference, then sorting these by season (Table 1). We then measured and applied mean seasonal corrections (e.g., the mean difference between end/start of quarters 4/5, 8/9, 12/13, 16/17). A good measure of the systematics remaining after these corrections is the standard deviation of the residuals: 63 ct s^{-1}. Since the mean flux was 15,933 ct s^{-1}, this indicates ~ 0.4% residual errors after quarterly offset correction.

2.3. Filtering

Next, bad cadences are filtered on three criteria: 1) "manual exclude" (`SAP_QUALITY` bit 9 set true), indicating Solar coronal mass ejections, 2) four cadence ranges (54,941-54,960, 64,084-64,129, 71,054-71,063, 72,324-72,332) deemed problematic on the basis of other sources' light curves, or 3) cadences that deviated by more than 5 times the reported error from both cadences before and after. These 444 bad cadences (in red in Figure 2) were eliminated, leaving a total of 55,209 good cadences.

2.4. Moiré pattern drift noise

Kepler data can also suffer from Moiré pattern drift (MPD) noise, which arises from crosstalk between the four fine guidance sensors and the 84 science channel readouts (Kolodziejczak, J. et al. 2010). Both Wehrle et al. (2013) and Revalski et al. (2014) have noted that MPD noise can be a serious problem for Kepler AGN variability studies, adding to the apparent source variablity or even mimicking variability in a non-variable source.

There is no known procedure to flag or mitigate MPD errors, although the Kepler project is investigating the phenomenon (Clarke et al. 2014). Table 13 of the Kepler Instrument Handbook[2] lists as problematic only one module.output used to observe Zw 229–15: 14.4, during Season 2. The possibility that MPD noise is affecting these data is discussed in Section 3.3.

3. PSD ESTIMATION

We performed two power spectral analyses of these data. The first uses standard Fourier methods to directly estimate the PSD, the second uses a continuous random process model fitted in the time domain to infer the PSD.

3.1. Periodogram analysis

The first approach used the standard periodogram of the full dataset to directly estimate the PSD. Gaps within each segment were filled by interpolation (using the LOWESS method; Cleveland (1981)) to give one evenly sampled light curve with sampling of ≈ 29.4 min. The light curve was end-matched (Fougere 1985) to suppress the effects of spectral "leakage" (Uttley et al. 2002). This involves subtracting a linear term such that the mean fluxes for the first and last 20 data points are the same The resulting periodogram is shown in Figure 3.

At high frequencies the power spectrum flattens, as expected from independent (white) flux measurement errors, but the observed level is higher than expected given the pipeline errors by a factor ≈ 1.57, suggesting the flux

[1] Software available at `http://keplergo.arc.nasa.gov/PyKE.shtml`

[2] See `https://archive.stsci.edu/kepler/manuals/KSCI-19033-001.pdf`

measurement errors are 25% larger than the pipeline errors. Further, the PSD rises slowly from 10 day^{-1} down to 1 day^{-1}, which could be explained in terms of some degree of correlation in the measurement errors, perhaps resulting from MPD. At lower frequencies the periodogram rises steeply with decreasing frequency and shows a bend around a timescale of ~ 5 days.

We fitted simple models to the periodogram by maximizing likelihood to estimate of the model parameters (see e.g. Vaughan 2010) using XSPEC 12.8.1 (Arnaud 1996) with the Whittle statistic. A model comprising a simple power law plus a constant gave a poor match to the data, with power law slope -3.15 and fit statistic $D = -2\log(\text{likelihood}) = -270489.0$ using 3 free parameters. Including an additional power law to model the excess power at ~ 1 day^{-1} improved the fit by $\Delta D = 157.9$ using five free parameters. Including a bend, in the steep power law, to a flatter slope at low frequencies (using a simple bending power law as in Edelson et al. 2013; González-Martín & Vaughan 2012) improved the fit by a further $\Delta D = 231.3$, and has a total of seven free parameters. The best-fitting parameters of this model are as follows: power law slopes of -2.00 ± 0.12 at low frequencies and -4.51 ± 0.20 at high frequencies, with a bend at $f_b = 0.18 \pm 0.03$ day^{-1} (~ 5.6 day timescale). The additional (unbending) power law had a slope -1.28 ± 0.13 and contributes significantly only around frequencies ~ 1 day^{-1}. Replacing the simple bending power law with a "Nuker" law (equation 1 of Lauer et al. 2005), which includes an extra parameter to adjust the sharpness of the bend between power law slopes, did not significantly improve the fit ($\Delta D = 0.3$ improvement for one additional free parameter).

3.2. CARMA-based PSD analysis

The second analysis utilized the continuous-time autoregressive moving average (CARMA) modeling technique of Kelly et al. (2014). This method naturally handles data gaps, measuring the power spectrum down to the lowest frequencies available. The CARMA modeling technique assumes that the light curve is a Gaussian process and that the power spectrum can be approximated as a mixture of Lorentzian functions. For computational purposes, we reduced the sampling by binning on 2.5 hour intervals. We considered CARMA($p, p-1$) models and used the Deviance Information Criterion (DIC, Spiegelhalter et al. 2002) to choose the value of $p = 5$; higher values of p produced worse DIC and did not lead to significantly different power spectra with the exception of higher uncertainty at the low and high frequency ends.

The inferred power spectrum (Figure 4a) shows evidence for a bending power-law shape, and a flattening toward the highest frequencies. The apparent kink in the PSD estimate is an artifact of the use of Lorentzians in the CARMA-PSD process. By analogy with a bending power-law model, we can quantify the effective power-law slopes of the power spectrum above and below the bend frequency as $d\log P(f)/d\log f$, and the effective bend frequency as the argumentative maximum of $|d^2 \log P(f)/(d\log f)^2|$. Based on the CARMA model we infer the effective bend frequency to be located 0.25 ± 0.01 day^{-1} (~ 4 day timescale), and effective power-law slopes of -1.99 ± 0.01 and -3.65 ± 0.07 at frequencies of 0.01 day^{-1} and 0.6 day^{-1}.

Figure 4b shows the same data with the y-axis multiplied by frequency so equal power per logarithmic frequency interval would be a flat line (e.g. Psaltis et al. 1999, Uttley et al. 2002). Note the continued rise to the longest timescales sampled by Kepler, indicating that the PSD must undergo second flattening because otherwise the total variability power would diverge.

3.3. Excess high frequency power

One unusual aspect of these PSD estimates is the apparent excess power at high temporal frequencies. In particular the upward bend around ~ 1 day timescales is unlikely to be intrinsic to the source, since both X-ray binary and AGN power $*$ frequency plots are generally downward-bending or flat at both high and low frequencies, consistent with finite total power (Psaltis et al. 1999, McHardy et al. 2004). No such feature has ever been seen in the X-rays for AGN (e.g., Markowitz et al. 2003), or, to our knowledge, at any wavelength in any astrophysical source. In this section we investigate the possibility that this is due to a poorly-studied Kepler instrumental effect: MPD noise.

As mentioned earlier, the Kepler Instrument Handbook listed only one of the modules used to observe Zw 229–15 as problematic: the one used for Season 2 data collection. To investigate this we eliminated Q17 (which is much shorter than all others), then segregated the four quarters in Season 2 (Q4/8/12/16) from the other three seasons (nine quarters in total). We truncated each quarter's data to the first 67.0 days, the length of the next shortest quarter (Q8). Then we interpolated across missing cadences in each dataset as before, measured the periodogram, and averaged the Season 2 quarters to produce one PSD estimate (the "MPD-flagged" estimate, shown in red in the top panel of Figure 5) and the non-Season 2 quarters to produce another ("MPD-unflagged," in blue). This assures both PSDs cover exactly the same range of temporal frequencies.

Note that the fits to the MPD-flagged PSD show that both the unbroken $\alpha \sim -1$ power-law and the Poisson noise term are higher for the MPD-flagged PSD, although both PSDs show an excess above the Poisson level expected solely on the basis of the quoted Kepler errors (see Vaughan et al. 2003 for details). This is clear evidence that there is additional variance in the data on short timescales that cannot be explained solely with the quoted Kepler errors.

The ratio of powers (MPD-flagged divided by MPD-unflagged) is shown in the bottom panel of Figure 5. Note that the MPD-flagged PSD estimate significantly exceeds that of the MPD-unflagged PSD at all temporal frequencies above 1.3 day^{-1} (timescales below ~0.75 day). This suggests, but by no means proves, that MPD noise could be responsible for the excess high frequency noise, as both the PSD levels and the fits to the two highest-frequency components are larger for the MPD-flagged data than for the MPD-unflagged data. It also suggests that even the MPD-unflagged quarters suffer from some effects of MPD, because the high-frequency fits and PSD data lie significantly above the levels expected from pure Poisson noise. However because MPD noise is a detector-wide phenomenon that is poorly suited to study with small "postage-stamp" downloads, this in-

dication must be considered tentative until the Kepler project completes a systematic detector-wide study of MPD noise.

3.4. Emission line variability

Although Barth et al. (2011) find that the optical Hβ emission lags the optical continuum by ~4 days, it is not likely that the ~ 5 day PSD break is due to emission line variability. Examination of spectra taken in that campaign indicates roughly 2-4% of the total flux in this small (~ 2x4 arcsecond) slit due to emission lines. Further, optical imaging suggests the underlying galaxy contributes no more than ~60% of the light in this aperture (Q. Wang, priv. comm.), so no more than 40% is nonstellar. This means that the emission lines contribute no more than 10% of the nonstellar Kepler-band flux from the AGN. (Galactic starlight does not contribute to the PSD because it is not variable.)

In AGN, the fractional continuum variability amplitude is always greater or similar to that of the nearby emission lines. Power is the square of amplitude, so the < 10% contribution of the lines to the total variability amplitude means that a negligible fraction (<1%) of the variability power can be due to the lines. Thus we conclude that the PSD features observed in this source are not due to the emission lines, but instead must be due to the nuclear continuum variability.

4. DISCUSSION

The main finding of this paper is that the Kepler PSD of Zw 229–15 shows smooth, power law-like shape with a bend from an index of ~ −2 at low frequencies to ~ −4 at high frequencies at a temporal frequency corresponding to a ~ 5 day timescale. Previous PSD analyses of Zw 229–15 either reported no evidence of a break (Mushotzky et al. 2011; Kelly et al. 2014) or a possible break on a ~ 90 day timescale (Carini & Ryle 2012). For the ~ −2 slope measured in the current paper, the (Carini & Ryle 2012) analysis does not require a break at ~ 90 day timescales (<60% likelihood). Further the Carini & Ryle (2012) binning limits the high frequency sampling such that it would be difficult to find a high frequency break similar to that reported in this paper. We thus conclude that we do not have a fundamental disagreement with any previous work.

4.1. Accretion disk size and time scales

In this section we calculate disk sizes implied by the measured ~ 5 day timescale under different assumptions about the physical processes that may be responsible for the observed variability. The optical emission from AGN is thought to arise in an accretion disk so we utilize the standard α-disk scaling formulae from King (2008).

Ground-based emission line monitoring of Zw 229–15 (Barth et al. 2011) yielded a ~ 4 day continuum-Hβ lag, which implies a black hole mass of $M_{BH} \sim 10^7\ M_\odot$ and Schwarzschild radius of $r_S = 2GM/c^2 \approx 3 \times 10^{12}$ cm ≈ 100 lt-sec. Barth et al. (2011) also estimated the bolometric luminosity $L_{bol} = 9 \times \lambda L_\lambda(5100A) = 10^{43.8}$ erg/s, corresponding to $L/L_{Edd} \sim 0.05$.

We consider a range of assumed accretion disk sizes because AGN accretion disk sizes are not currently well-constrained. First, Horne et al. (2014) used contemporaneous Kepler, Swift and Suzaku monitoring of Zw 229–15 to estimate the X-ray/optical delay map and thus infer the size of the optically-emitting accretion disk. That analysis yielded a mean X-ray/optical lag of ~ 1.7 days, corresponding to a disk radius of R ~ $1500 r_S$, albeit with large uncertainties. Second, we note that standard α-disk models (e.g., Shakura & Sunyaev 1973) yield significantly smaller sizes, R ~ $100 r_S$. In order to capture the large uncertainty in this important but poorly-constrained parameter, we consider the implications of the observed ~ 5 day break timescale in light of two assumed order-of-magnitude disk size scales: 100 r_S and 1000 r_S. Recent observations of microlensing in AGN (e.g., Jimenez-Vicente et al. 2014) typically yield disk size estimates between these values.

4.1.1. Reprocessing

"Reprocessing" models posit that the observed optical variations are due to irradiation of the disk by the central X-ray source (a corona or jet). The disk would act as a low-pass spatial filter so that the long timescale X-ray variations would be reproduced in the optical while timescales shorter than the light-crossing time would be smoothed out. This model has been successful in reproducing interband optical correlations in AGN (e.g., Cackett et al. 2007) but there is often insufficient X-ray luminosity to power the total disk luminosity.

The light-crossing time for a $100 - 1000 r_S$ emitting region would be $10^4 - 10^5$ sec or 0.12 – 1.2 day. While the smaller estimate is too small to be consistent with the observed bend timescale, the larger size is marginally consistent, especially given that the light-crossing time is more likely to be associated with the source diameter, not its radius. Further even in the case of the smaller size estimate, reprocessing would remain viable if the (currently undetermined) PSD of the driving X-ray light curve already had a ~ 5 day bend similar to that observed in the optical. Alternatively it could also be that a small fraction of the optical variability was due to reprocessing, while the bulk due to "intrinsic" processes discussed below, in which case the signature would not be visible in the optical PSD.

4.1.2. Dynamical processes

For a $10^7 M_\odot$ black hole, the dynamical (orbital) timescale is $t_{dyn} \sim 140 (R/r_S)^{3/2}$ sec (equation 11, King 2008). Optical emission distances of $R = 100 - 1000 r_S$ yield effective timescales of $t_{dyn} \sim 1.6 - 50$ days. This range encompasses the observed ~ 5 day beak timescale. It is thus tempting to associate the bend with the dynamical timescale at the radius at which the optical photons are produced. If this is indeed the proper scaling then we predict that for the other objects observed by Kepler that the bend frequency will scale weaker than linearly with the mass and weakly with the Eddington ratio of the sources.

4.1.3. Thermal processes

The thermal and viscous timescales are considerably longer (King 2008): $t_{dyn} \sim \alpha t_{th} \sim \alpha (H/R)^2 t_{visc}$, where α is the Shakura & Sunyaev (1973) viscosity parameter and (H/R) is the ratio of disk height to radius. This assures that $t_{lc} < t_{dyn} < t_{th} < t_{visc}$, since both α and H/R must be less than one. For α ~ 0.1, an emission distances

$R = 100 - 1000 r_S$ yields $t_{th} \sim 16 - 500$ days. The lower end of this range could be considered a marginally acceptable match to the observed ~ 5 day timescale given the uncertainties.

4.1.4. Viscous processes

Emission distances of $R = 100 - 1000 r_S$ yield $t_{visc} \sim 4 - 140$ yr for H $R \sim \alpha \sim 0$. These are not consistent with the observed ~ 5 day timescale. This might be considered to be evidence against the propagating fluctuation model (Lyubarskii 1997; King et al. 2004; Arevalo & Uttley 2006), which posits that the variations are generated internally by random variations in viscosity (and therefore local accretion rate) over a wide range of spatial scales in the accretion flow. Since there must be a second, currently unobserved bend in the PSD at longer timescales than probed by the Kepler data in order for the total variability power to be finite (see Figure 4b) this feature may turn out to be associated with viscous or thermal timescales. The true origin of the observed bend frequency awaits the development of detailed accretion disk models (Schnittman et al. 2013) capable of self consistently calculating the emitted radiation.

If the reprocessing model explains the origin of the rapid optical variations, these should be correlated with the X-rays, but delayed and smoothed. By contrast if the disk is varying due to intrinsic dynamical/thermal fluctuations, the optical variations may be largely independent of the X-ray variations, or may even generate delayed modulations in the X-rays. These predictions can be tested with coordinated X-ray and optical campaigns that provide sufficient signal/noise and temporal resolution.

4.2. Caveats

While this is the first optical AGN PSD to span such a large range of temporal frequencies, it is important to be cautious interpreting these new, relatively untested Kepler data. As discussed in Section 3.3, this PSD seems to show excess power on timescales < 1 day. It is tempting to associate this feature with MPD noise (which also seems to show correlated variability on timescales of order ~ 1 day) but this has not been definitively established (Kolodziejczak, J. et al. 2010). Further, other possible systematic errors may be present in ways not yet considered.

Finally a more fundamental limitation may be that the underlying PSDs of AGN are not fully described by the simple multiple power-law models used here and in previous analyses. The highly variable blazar W2R1926+42 exhibits an even more complex Kepler PSD than Zw 229–15 (Edelson et al. 2013). Because Kepler PSDs cover an unprecedented range of temporal frequency (~ 4 decades), we should remain skeptical of assigning physical significance to parameters derived from simple model fits.

5. CONCLUSIONS

This paper presents the best AGN optical light curve ever measured, showing strong variability on long timescales (Figure 2b) yet appearing remarkably "smooth" on timescales shorter than a few days (Figure 2c). The resulting PSD (Figures 3 and 4) is the first to cover nearly ~ 4 decades of temporal frequency. The PSD shows no evidence of QPO, instead it has a smooth, power law-like shape with a bend from an index of ~ –2 at low frequencies to ~ –4 at high frequencies at a frequency corresponding to a ~ 5 day timescale, and excess power near ~ 1 day timescales that may be due to systematic MPD noise generated by Kepler. This measured timescale can be plausibly associated with the light-crossing, dynamical or thermal timescales, depending on the assumed size of the optically-emitting accretion disk as well as other unobserved disk parameters. It cannot be associated with the viscous timescale, although the fact that the PSD continues to rise to the longest timescales measured means that there must be a second low-frequency break, and that feature could be associated with the viscous timescale.

These results should also motivate future work in this area. First, while Zw 229–15 is the best-observed Kepler AGN, there are others with data that are nearly as good, and we will be analyzing them in a future paper to see if they share the same characteristics. Second, the steep PSD at the longest timescales probed by Kepler means that there must be a second, currently unobserved break. It is important to continue to monitor Kepler AGN at lower cadence so we can locate that break. Third, MPD noise is the result of spurious "waves" across large regions of Kepler detectors, so it cannot be properly addressed by studies limited to tiny "postage stamp" downloads such as this. The Kepler project's close-out plans include a final update to the data processing pipeline to identify and flag data impacted by MPD in all Kepler light curve files. Fourth, Kepler data are of such high precision that they are pushing the limits of PSD analysis followed by simple broadband model fits. It would be useful for theorists develop the stationary conditional probability distribution that describes the evolution of the light curve, instead of only the first and second moments of the distribution, which is what the PSD supplies.

The authors appreciate the assistance of Martin Still and Tom Barclay of the Kepler GO office in helping us understand the Kepler data, and Keith Horne and Qian Wang for sharing their unpublished results on Zw 229–15. We also thank the anonymous referee and the editors of the Astrophysical Journal for a prompt and helpful refereeing process. This research utilized the HEASARC, IRSA, NED and MAST data archives and the NASA Astrophysics Data System Bibliographic Service. RE acknowledges support from the Kepler GO and ADAP programs through NASA grants NNX13AC26G and NNX13AE99G.

Facilities: Kepler.

REFERENCES

Arevalo, P., Uttley, P. 2006, MNRAS, 367, 801

Arnaud, K.A., 1996, Astronomical Data Analysis Software and Systems V, eds. Jacoby G. and Barnes J., p17, ASP Conf. Series volume 101

TABLE 1
Interquarter offset correction

(1) Gap	(2) Quarters	(3) Seasons	(4) Offset (c/s)	(5) Correction (c/s)	(6) Residual (c/s)
1	4/5	2/3	1471	-1547	-76
5	8/9	2/3	1550	-1547	4
9	12/13	2/3	1600	-1547	54
13	16/17	2/3	1566	-1547	19
2	5/6	3/0	-870	750	-120
6	9/10	3/0	-702	750	48
10	13/14	3/0	-677	750	72
3	6/7	0/1	-234	252	18
7	10/11	0/1	-291	252	-39
11	14/15	0/1	-230	252	22
4	7/8	1/2	-397	481	84
8	11/12	1/2	-485	481	-3
12	15/16	1/2	-562	481	-81
				Standard deviation of residual	63

Note. — Column 2 lists the Kepler quarter numbers bounding each gap. Column 3 is the same referenced to the four Kepler seasons. Column 4 gives the measured interquarter offset. Column 5 gives resulting seasonal correction. Column 6 gives the residual after correction. The standard deviation of these residuals is given in the bottom line.

Barth, A. et al. 2011, ApJ, 732, 121
Borucki, W. et al. 2010, Science, 327, 977
Breedt, E., et al., 2009, MNRAS, 403, 605
Carini, M., Ryle, W. 2012, ApJ, 749, 70
Cackett, E. M., Horne, K., Winkler, H., 2007, MNRAS, 380, 669
Cleveland, W. S. 1981, LOWESS: A program for smoothing scatterplots by robust locally weighted regression. The American Statistician 35, 54
Clarke, B., Kolodziejczak, J, Caldwell, D. 2014, BAAS, 224, 120.07
Edelson et al. 2013, ApJ, 766, 16
Falco, E. et al. 1999, PASP, 111, 438
Fougere, P. F., 1985, J. Geophys. Res., 90, 4355
González-Martín, O., Vaughan, S., 2012, A&A, 544, A80
Hawkins, M. 2007, A&A, 462, 581
Horne, K., Edelson, R., Mushotzky, R., Olden, P., Simpson, J., Vaughan, S. 2014, in prep.
Jenkins, J. et al. 2010, ApJ, 713, L87
Jiminez-Vicente et al. ApJ, in press, arXiv:1401.2785
Kelly, B. et al. 2014, ApJ, 788, 33
King, A. et al. 2004, MNRAS, 348, 111
King, A. 2008; 52, 253
Kinemuchi, K., Fanelli, M., Pepper, J., Still, M., Howell, S. 2012, PASP, 124, 963
Kolodziejczak, J. et al. 2010, SPIE, 7742, 38
Lauer, T. et al. 2005, AJ, 129, 2138
Lyubarskii, Y. 1997, MNRAS, 292, 679
Markowitz, A. et al. 2003, ApJ, 593, 96
McHardy, I. et al. 2004, MNRAS, 348, 783
Mushotzky, R. et al. 2011, ApJL, 743, L12
Peterson, B. et al. 2004, ApJ, 613, 682
Psaltis, D. 1999, ApJ, 520, 262
Revalski, M. et al. 2014, ApJ, 785, 60
Schnittman, J. et al 2013, ApJ, 769, 156
Shakura, N., Sunyaev, R. 1973, A&A, 24, 337
Spiegelhalter D. J., Best N. G., Carlin B. P., Van Der Linde A., 2002, J. Roy. Stat. Soc. Series B, 64, 583
Uttley, P., McHardy, I. M., Papadakis, I., 2002, MNRAS, 332, 231
Vaughan, S., Edelson, R., Warwick, R., Uttley, P. 2003, MNRAS, 345, 1271
Vaughan S. 2010, MNRAS, 402, 307
Wehrle, A. et al. 2013, ApJ, 773, 89

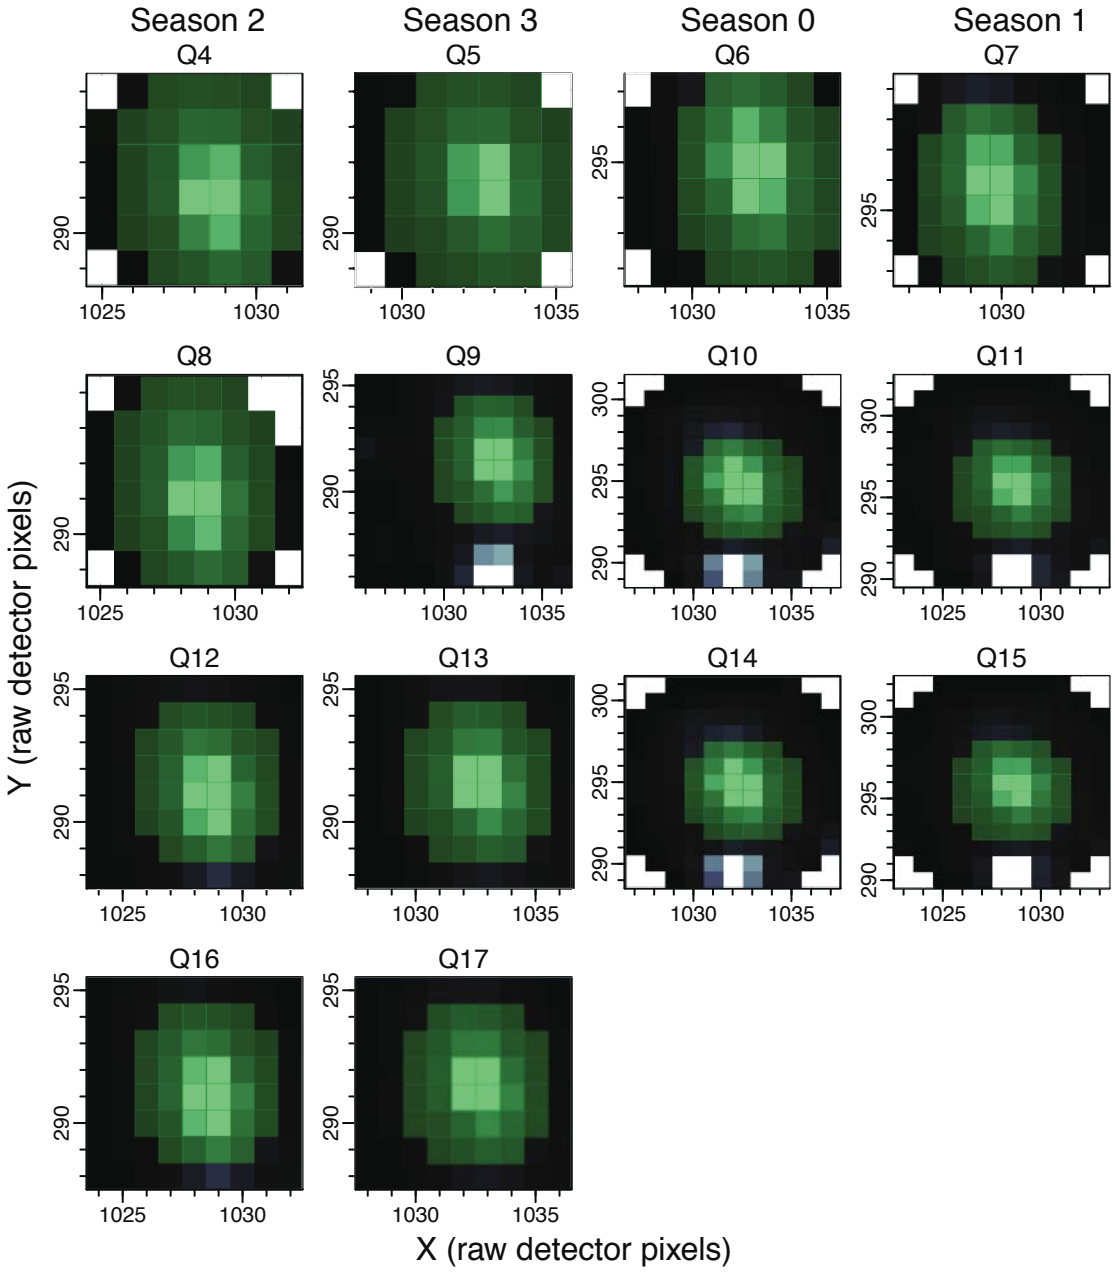

Fig. 1.— Kepmask output showing downloaded data "postage stamp" in greyscale, except occasional missing corner pixels (white). Extraction masks (green) are identical within each season and matched as closely as possible between seasons. A nearby star is visible e.g. in Q9-11, however its light is not included in the extracted light curve because it lies outside the green extraction mask. Only data from pixels within the mask are included in the light curve. The 32-pixel masks are the largest possible symmetrical masks given the small Q4/Q5 downloads. We find that this makes interquarter repeatability better than the SAP pipeline "optimal" masks, which vary in size and location. The larger sizes mitigate excursions due to thermally-induced focus changes and to differential velocity aberration Kinemuchi et al. (2012).

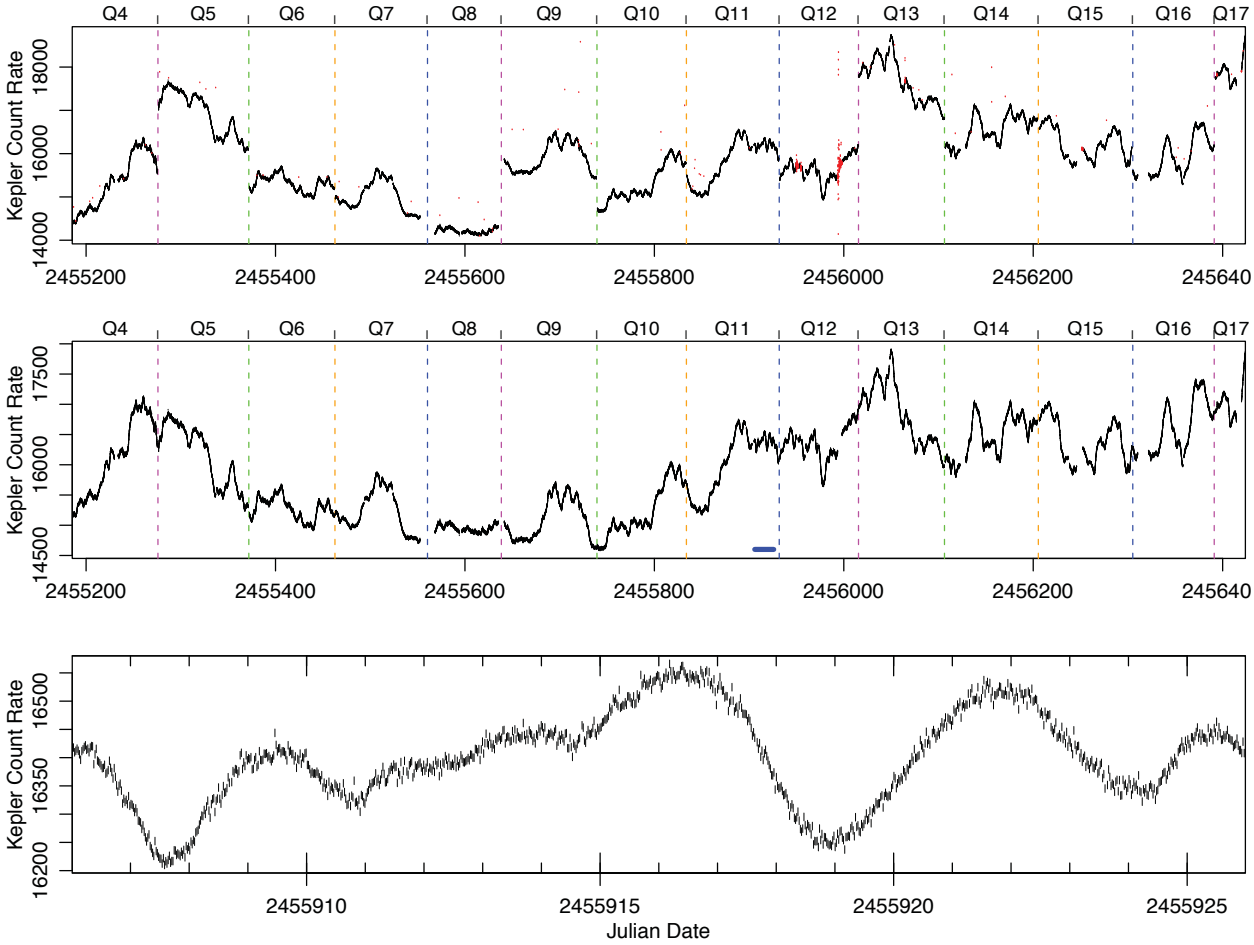

Fig. 2.— 2a: Uncorrected light curve. Bad data (in red) were eliminated and seasonal jumps were corrected as discussed in text. 2b: Corrected light curve. 2c: 20 day (960 cadence) snippet from Q11, showing the quality of the Kepler data. This time range is shown with a horizontal blue line in Figure 2b.

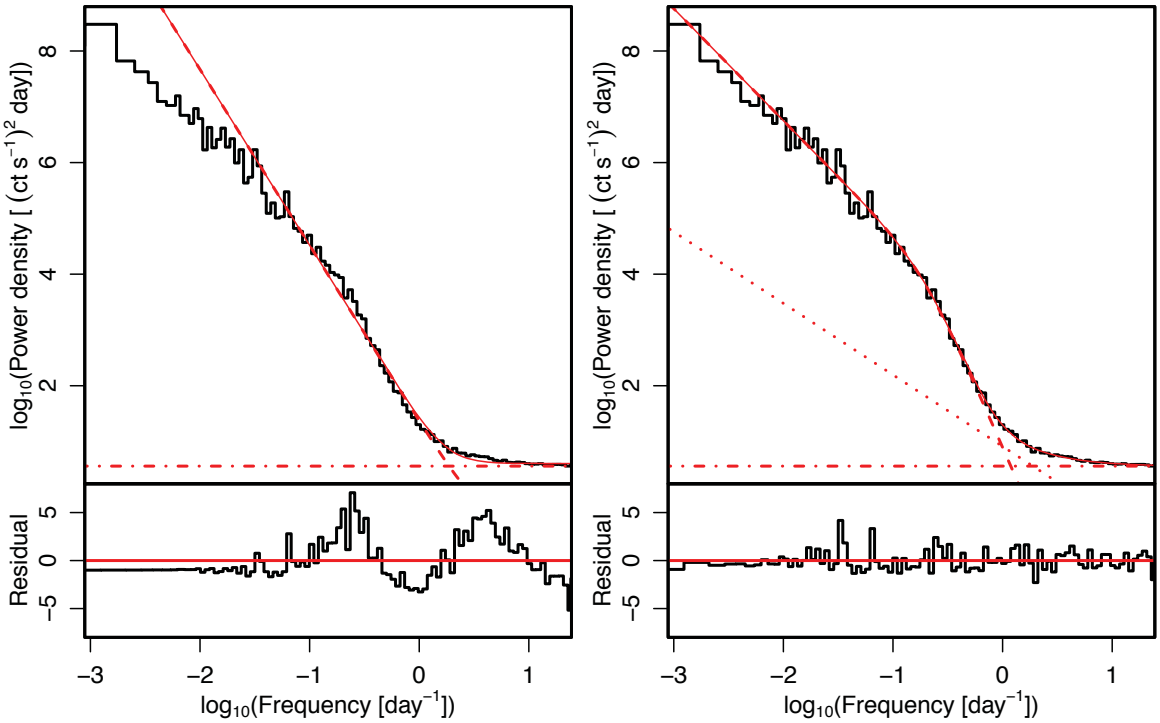

Fig. 3.— Standard periodogram analysis. The data (black) have been rebinned for display purposes only, such that the lowest frequency data are not binned, and at higher frequencies data are averaged over bins spanning a factor 1.1 in frequency. The fits are shown in solid red at the top and residuals, computed as (data − model)/($\sqrt{\text{model}}$) at the bottom. 3a: A single power-law (dashed line) plus Poisson noise (dot-dash line) model yields a poor fit with large coherent features in the residual plot. 3b: A bending power law (dashed) plus a second power-law (dotted line) plus noise model yields an acceptable fit with smaller and better-distributed residuals.

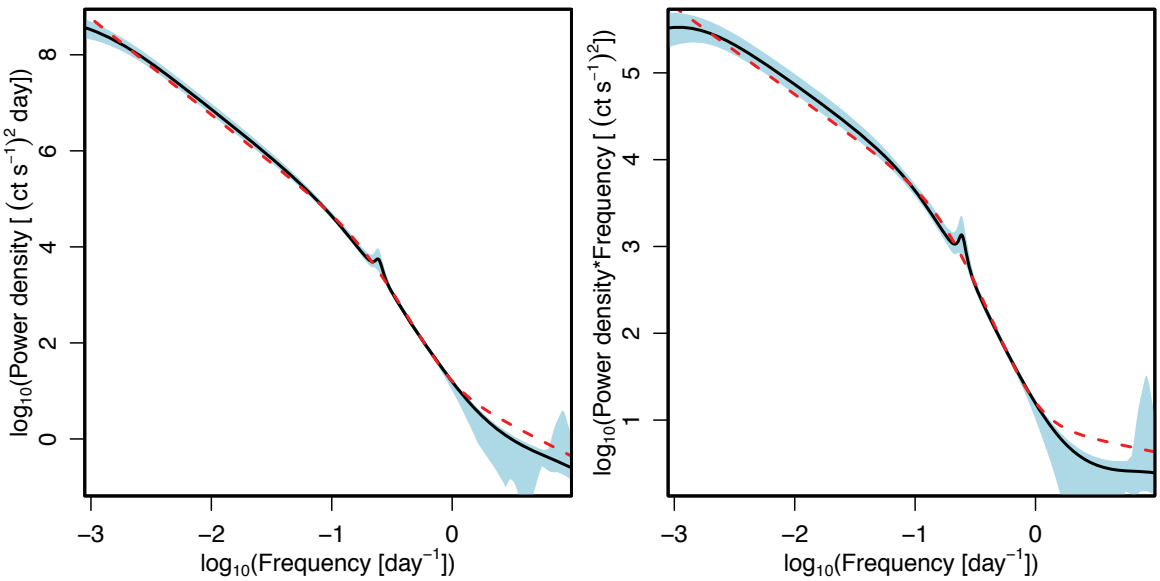

Fig. 4.— 4a: CARMA-based periodogram using the same normalization as Figure 3. The best estimate of the PSD is shown as a black line and the shaded blue region shows the 95% confidence interval. Note the good agreement with the Figure 3b fit (dashed red line). The feature near ~ 0.3 day^{-1} is an artifact of the modeling and does not indicate QPO (Kelly et al. 2014). 4b: Same data but with the y-axis multiplied by frequency.

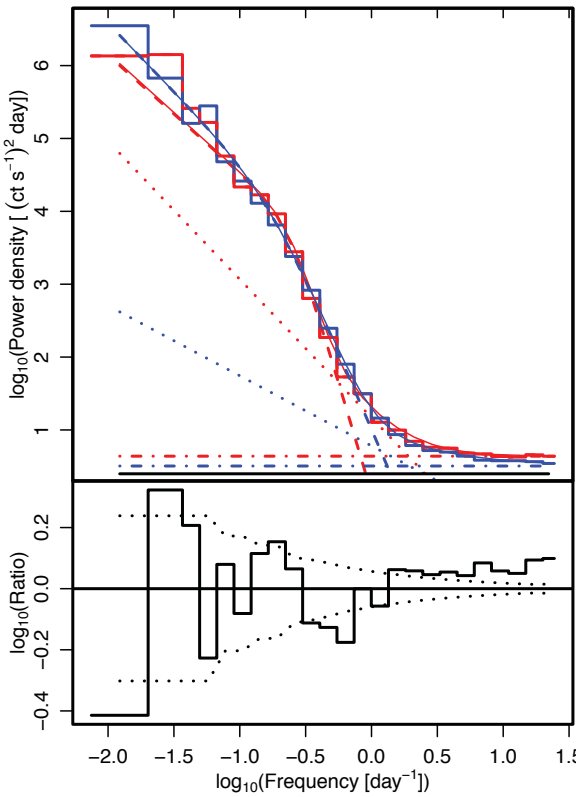

Fig. 5.— Top: Averaged PSDs for the MPD-flagged quarters (in red) and MPD-unflagged quarters (blue). As in Figure 3, the data are shown as a step plot, full fit as a thin solid line, broken power-law component as a dashed line, second power-law as a dotted line, and Poisson noise as a dot-dash line. The expected Poisson noise, based on the quoted Kepler errors, is shown as a horizontal solid black line. Note that the MPD-flagged PSD fit at high frequencies shows higher Poisson noise and second power-law than the unflagged PSD fit, and that both stronger Poisson noise than expected from the quoted Kepler errors. Bottom: Logarithm of the ratio of MPD-flagged and unflagged PSDs. This would be a flat line consistent with zero if MPD was not a source of noise. Dotted lines show the 1-σ confidence interval for a result consistent with zero. Note that the flagged PSD shows significantly more power than the unflagged PSD at all high frequencies above 1.3 day^{-1}.