
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

O'Neill, S.; Kiehlmann, S.; Readhead, A. C.S.; Aller, M. F.; Blandford, R. D.; Lioudakis, I.; Lister, M. L.; Mróz, P.; O'Dea, C. P.; Pearson, T. J.; Ravi, V.; Vallisneri, M.; Cleary, K. A.; Graham, M. J.; Grainge, K. J.B.; Hodges, M. W.; Hovatta, T.; Lähteenmäki, A.; Lamb, J. W.; Lazio, T. J.W.; Max-Moerbeck, W.; Pavlidou, V.; Prince, T. A.; Reeves, R. A.; Tornikoski, M.; Vergara De La Parra, P.; Zensus, J. A.

The Unanticipated Phenomenology of the Blazar PKS 2131-021: A Unique Supermassive Black Hole Binary Candidate

Published in:
Astrophysical Journal Letters

DOI:
[10.3847/2041-8213/ac504b](https://doi.org/10.3847/2041-8213/ac504b)

Published: 01/02/2022

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
O'Neill, S., Kiehlmann, S., Readhead, A. C. S., Aller, M. F., Blandford, R. D., Lioudakis, I., Lister, M. L., Mróz, P., O'Dea, C. P., Pearson, T. J., Ravi, V., Vallisneri, M., Cleary, K. A., Graham, M. J., Grainge, K. J. B., Hodges, M. W., Hovatta, T., Lähteenmäki, A., Lamb, J. W., ... Zensus, J. A. (2022). The Unanticipated Phenomenology of the Blazar PKS 2131-021: A Unique Supermassive Black Hole Binary Candidate. *Astrophysical Journal Letters*, 926(2), Article 35. <https://doi.org/10.3847/2041-8213/ac504b>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.



The Unanticipated Phenomenology of the Blazar PKS 2131–021: A Unique Supermassive Black Hole Binary Candidate

S. O’Neill¹, S. Kiehlmann^{2,3}, A. C. S. Readhead^{1,3}, M. F. Aller⁴, R. D. Blandford⁵, I. Liodakis⁶, M. L. Lister⁷, P. Mróz⁸, C. P. O’Dea⁹, T. J. Pearson¹, V. Ravi¹, M. Vallisneri¹⁰, K. A. Cleary¹, M. J. Graham¹¹, K. J. B. Grainge¹², M. W. Hodges¹, T. Hovatta^{6,13}, A. Lähteenmäki^{13,14}, J. W. Lamb¹, T. J. W. Lazio¹⁰, W. Max-Moerbeck¹⁵, V. Pavlidou^{2,3}, T. A. Prince¹¹, R. A. Reeves¹⁶, M. Tornikoski¹³, P. Vergara de la Parra¹⁶, and J. A. Zensus¹⁷

¹Owens Valley Radio Observatory, California Institute of Technology, Pasadena, CA 91125, USA; acr@caltech.edu

²Department of Physics and Institute of Theoretical and Computational Physics, University of Crete, 71003 Heraklion, Greece

³Institute of Astrophysics, Foundation for Research and Technology-Hellas, GR-71110 Heraklion, Greece

⁴2 Department of Astronomy, University of Michigan, 323 West Hall, 1085 S. University Avenue, Ann Arbor, MI 48109, USA

⁵Kavli Institute for Particle Astrophysics and Cosmology, Department of Physics, 20 Stanford University, Stanford, CA 94305, USA

⁶Finnish Centre for Astronomy with ESO (FINCA), University of Turku, FI-20014 University of Turku, Finland

⁷Department of Physics and Astronomy, Purdue University, 525 Northwestern Avenue, West Lafayette, IN 47907, USA

⁸Astronomical Observatory, University of Warsaw, Al. Ujazdowskie 4, 00-478 Warszawa, Poland

⁹Department of Physics and Astronomy, University of Manitoba, Winnipeg, MB R3T 2N2, Canada

¹⁰Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

¹¹Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA

¹²Jodrell Bank Centre for Astrophysics, University of Manchester, Oxford Road, Manchester M13 9PL, UK

¹³Aalto University Metsähovi Radio Observatory, Metsähovintie 114, FI-02540 Kylmäla, Finland

¹⁴Aalto University Department of Electronics and Nanoengineering, P.O. Box 15500, FI-00076 Aalto, Finland

¹⁵Departamento de Astronomía, Universidad de Chile, Camino El Observatorio 1515, Las Condes, Santiago, Chile

¹⁶CePIA, Astronomy Department, Universidad de Concepción, Casilla 160-C, Concepción, Chile

¹⁷Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, D-53121 Bonn, Germany

Received 2021 October 30; revised 2022 January 20; accepted 2022 January 21; published 2022 February 23

Abstract

Most large galaxies host supermassive black holes in their nuclei and are subject to mergers, which can produce a supermassive black hole binary (SMBHB), and hence periodic signatures due to orbital motion. We report unique periodic radio flux density variations in the blazar PKS 2131–021, which strongly suggest an SMBHB with an orbital separation of ~ 0.001 – 0.01 pc. Our 45.1 yr radio light curve shows two epochs of strong sinusoidal variation with the same period and phase to within $\lesssim 2\%$ and $\sim 10\%$, respectively, straddling a 20 yr period when this variation was absent. Our simulated light curves accurately reproduce the “red noise” of this object, and Lomb–Scargle, weighted wavelet Z-transform and least-squares sine-wave analyses demonstrate conclusively, at the 4.6σ significance level, that the periodicity in this object is not due to random fluctuations in flux density. The observed period translates to 2.082 ± 0.003 yr in the rest frame at the $z = 1.285$ redshift of PKS 2131–021. The periodic variation in PKS 2131–021 is remarkably sinusoidal. We present a model in which orbital motion, combined with the strong Doppler boosting of the approaching relativistic jet, produces a sine-wave modulation in the flux density that easily fits the observations. Given the rapidly developing field of gravitational-wave experiments with pulsar timing arrays, closer counterparts to PKS 2131–021 and searches using the techniques we have developed are strongly motivated. These results constitute a compelling demonstration that the phenomenology, not the theory, must provide the lead in this field.

Unified Astronomy Thesaurus concepts: Relativistic jets (1390); Blazars (164); Supermassive black holes (1663); Active galactic nuclei (16)

1. Introduction

The identification of supermassive black hole binaries (SMBHBs) will open the field to multimessenger astronomy through the gravitational radiation they produce. Pulsar timing arrays provide a powerful technique for searching for nanohertz signals from gravitational waves from SMBHBs through the timing of millisecond pulsars (Holgado et al. 2018; Burke-Spolaor et al. 2019). However, in spite of the fact that galaxy mergers are not uncommon, there are relatively few instances of two galaxies with supermassive black holes (SMBHs) in

their nuclei being seen in the actual process of the merging of the spheres of influence of their SMBHs, or of the following stages, when an SMBHB forms by ejecting stars from the merging central clusters, spirals in more closely owing to gravitational radiation, and finally coalesces (Begelman et al. 1980). A particularly fine example of the early stage of possible evolution toward an SMBHB is that of 3C 75 (Owen et al. 1985), where both SMBHs are producing radio jets, and their projected separation is 7.2 kpc. On parsec scales the best SMBHB candidate is B3 0402+379 (Rodriguez et al. 2006; Bansal et al. 2017), with a projected separation of 7.3 pc, a deduced period of 3×10^4 yr, and a deduced SMBHB mass of $\approx 1.5 \times 10^{10} M_{\odot}$. The strongest SMBHB candidate with a separation of $\ll 1$ pc is OJ 287 (Sillanpaa et al. 1988; Valtonen et al. 2016; Dey et al. 2021), for which the separation is

~ 0.1 pc, the deduced primary mass is $\approx 1.8 \times 10^{10} M_{\odot}$, and the deduced secondary mass is $\approx 1.5 \times 10^8 M_{\odot}$. At separations $\ll 1$ pc, even with high-frequency very long baseline interferometry (VLBI), for all but the closest active galactic nuclei (AGNs), we lack the angular resolution required to demonstrate the existence of an SMBHB through imaging, and we have to look for other signatures.

In principle, light curves offer a way forward (Haiman et al. 2009; Burke-Spolaor et al. 2019) because SMBHBs may reasonably be expected to exhibit periodicities. However, it has been pointed out (Vaughan et al. 2016; Covino et al. 2019) that, notwithstanding the rich literature on periodicities and quasi-periodic oscillations (QPOs) in blazars going back over five decades, there are very few statistically solid results. In their detailed analysis of 10 blazars in which QPOs have been reported, Covino et al. (2019) show that no strong cases for \sim -year-long periodicities can be confirmed. They are all consistent with the power spectra of the variations in these objects. It requires careful modeling of the red noise in the power spectrum of a blazar to evaluate the significance of any claimed periodicity. Sandrinelli et al. (2017, 2018) had estimated that $\sim 10\%$ of bright γ -ray blazars are QPOs, but after their detailed analysis they withdrew this estimate (Covino et al. 2019).

In a search for strong periodic signals showing at least 1.5 cycles in the optical light curves of 243,500 quasars, Graham et al. (2015a) found 111 candidates, of which the strongest is PG 1302–102 (Graham et al. 2015b). In this object approximately sinusoidal variations have been seen over a span of ~ 20 yr. However, Vaughan et al. (2016) have challenged the SMBHB interpretation of the PG 1302–102 optical light curve, attributing the periodicity to red noise.

Since blazar light curves at radio wavelengths have a red-noise spectrum with a non-Gaussian probability density function (pdf; Liodakis et al. 2017), except where stated otherwise, we assume a red-noise pdf throughout this paper.

Given the history and the persisting problems, it is clear that great caution is needed in the identification of periodicities and quasi-periodicities in active galactic nuclei. For this reason we regarded the $P_{\oplus} = 4.69 \text{ yr} \pm 0.14 \text{ yr}$ Earth-frame observed periodicity reported in the blazar PKS 2131–021 by Ren et al. (2021) as an interesting result to follow, but by no means yet shown to be a strong QPO or SMBHB candidate. The Ren et al. (2021) paper is based entirely on 11 yr (2008–2019) of our own 15 GHz monitoring observations of PKS 2131–021 with the Owens Valley Radio Observatory (OVRO) 40 m Telescope.

We recently came across observations of PKS 2131–021 (O’Dea et al. 1986) made at the Haystack Observatory between 1975 and 1983, which show the same periodicity to $\lesssim 2\%$ and phase to within 10% of the period. As we show with extensive tests in this paper, the level of significance of this periodicity, when 45.1 yr of radio monitoring data are combined, is 4.6σ , and it is certainly not a red-noise phenomenon. This makes PKS 2131–021 a strong QPO or SMBHB candidate. In addition to the Haystack data, we have also added the 14.5 GHz light curve of the University of Michigan Radio Astronomy Observatory (UMRAO), which covers the period 1980–2012 and is in excellent agreement with the Haystack and OVRO light curves in the regions of overlap, thereby giving us an uninterrupted, well-sampled, 45.1 yr 14.5–15.5 GHz light curve of PKS 2131–021.

All of the data presented in this paper are from targeted observations of PKS 2131–021, i.e., they are not serendipitous observations in which PKS 2131–031 was observed in the fields of other objects. Given the inhomogeneous processing procedures, the diverse sets of flux calibrators, and the observed matched flux densities in the overlapping regions, it is clear that the unusual light curve of PKS 2131–021 is not a result of faulty processing. These demonstrate the basic light-curve integrity.

At the start of this project we confirmed four predictions related to the periodicity seen in the OVRO data. This convinced us that the periodicity is telling us something important about the physics of this object and is not simply a random variation due to red noise. We began with a least-squares sine-wave fit to the OVRO data, on the basis of which we predicted that the sine-wave oscillations had begun before the OVRO observations, so we extrapolated the sine wave backward and began a search for earlier data. Our prediction was that we would find earlier sinusoidal variations in phase with the OVRO observations. The first confirmed prediction came when we looked at MOJAVE (Lister et al. 2018) and UMRAO data going back to 1995 and found an in-phase peak in 2005, immediately preceding the first OVRO peak. We subsequently obtained UMRAO data going back to 1980 and found a second in-phase peak in 1982. This was the second confirmed prediction. The peak in 1982 was much larger than the peaks from 2005 to 2020, and we predicted that had we been observing prior to 1982 we would have seen a larger sinusoidal oscillation prior to 1982. At the time we thought there were no data on PKS 2131–021 earlier than the UMRAO data. However, through a literature search, we discovered the Haystack data, which began in 1975. This did indeed show an earlier in-phase peak, in 1976. This was the third confirmed prediction. Furthermore, as predicted, like the peak in 1982, the peak in 1976 was much larger than the sinusoidal variations from 2005 to 2020. This was the fourth confirmed prediction. The chain of events occurred exactly as described here and convinced us of the significance of the OVRO periodicity, which, as we show in this paper, has been confirmed by rigorous statistical analysis.

There are clearly several possible explanations for the periodicities observed in PKS 2131–021. Precession of the relativistic jet due to misalignment of the spin axis of the SMBH and accretion disk (Caproni et al. 2004) is one possibility. Alternatively, the periodicity could be the result of precession due to misalignment of the orbital plane of an SMBHB with the accretion disk of the more massive SMBH (Caproni et al. 2004). Another possibility is precession due to warping of the accretion disk (Abraham 2018; Britzen et al. 2018). While these may all be viable explanations of the periodicity we see in PKS 2131–021, there is a more straightforward explanation—namely, that the periodicity is simply due to the orbital motion of the SMBHB. We show here that all the observations can be explained by this simple model and that no precession is needed to explain the light curve of PKS 2131–021, although, as we show, it might well explain the large-scale morphology. It should be noted that, given the characteristic timescales for variability in blazars, which range from months to years, and which are likely dominated by fueling of the central engine, and given the multiple sites of radio emission along the jets, it is not difficult to invent models in which the sinusoidal signal switches on and off. Thus, the

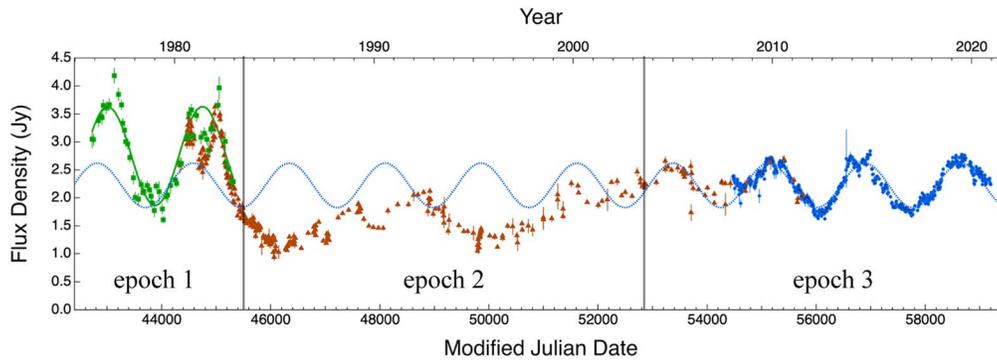


Figure 1. The 14.5–15.5 GHz light curve of PKS 2131–021. The Haystack data are shown by the green squares, the UMR AO data by the brown triangles, and the OVRO data by the blue circles. Note the excellent agreement between Haystack and UMR AO and between UMR AO and OVRO in the regions of overlap. The light curve shows three distinct epochs of activity: In epoch 1 (MJD < 45,500) there is a strong periodic signal, $P_1 = 1729.1 \pm 32.4$ days; in epoch 2 (45,500 < MJD < 52,850) this periodic signal is absent; in epoch 3 (MJD > 52,850) the periodic signal of epoch 1 returns, with $P_3 = 1760.4 \pm 5.3$ days, in phase but with lower amplitude than that of epoch 1. The solid green line shows the least-squares sine-wave fit to the Haystack data, and the dotted blue line shows the least-squares sine-wave fit to the OVRO data extrapolated backward to provide a comparison with the Haystack data. The periods of the Haystack and OVRO sine-wave fits match to $\lesssim 2\%$. The phase of the OVRO periodic signal extrapolated back to epoch 1 matches to $\sim 10\%$ of the period.

appearance and disappearance of the sinusoidal variability is easily accommodated in any model, and for this reason we do not discuss it further in this paper. Since the SMBHB orbital motion explanation is the simplest, we apply Occam’s razor. We deliberately do not consider other possible explanations in this paper, since it seems to us that nature is pointing the way. Adopting the simplest explanation is the best way to proceed at this early stage in our understanding of the phenomenology of SMBHBs with relativistic jets. Simple orbital motion was suggested as an explanation of blazar periodicities by Sobacchi et al. (2017), but their model is a complex one, which does not produce sinusoidal variations in all circumstances, and not at all unless the jet itself is assumed to consist of a fast-moving “spine” surrounded by a slower-moving “sheath.” While this might well be the case in PKS 2131–021, it is not required by our model, which is the simplest possible SMBHB–relativistic jet model.

We wish, therefore, to make clear at the outset that in this paper we deliberately focus on a particular model, to the exclusion of other viable models, because the phenomenology of the sinusoidal flux density variations—which has not been anticipated in previous studies—is an inevitable and inescapable consequence of the SMBHB orbital motion of the relativistic jet, and we are strongly of the opinion that we should pursue this approach until the phenomenology *requires* more complex models.

In this paper we consider all periodicities with significance levels below 3σ to be red noise, unless other, uncorrelated and independent observations raise the significance to the $\geq 3\sigma$ level. We analyze the PKS 2131–021 light curve and show that it is unique among AGNs that have been considered as possible SMBHB candidates and is indeed a prime SMBHB candidate. In Section 2 we describe the observations. In Section 3 we analyze the light curve using three different approaches and taking great care to model the red noise correctly and hence to derive robust measures of the significance of our results, and we derive an upper limit to the chirp mass of the putative SMBHB based on the radio observations alone. In Section 4 we present a model of PKS 2131–021, in which the observed periodicity is the orbital period of the putative black hole binary, which can account for the sinusoidal shape and amplitude of the periodic variability that we see. In Section 5

we discuss the expected gravitational-wave (GW) strain and derive an upper limit on the chirp mass of the SMBHB based on the upper limits derived from pulsar NANOGrav observations (Aggarwal et al. 2019).

The redshift of PKS 2131–021 is $z = 1.285$ (Drinkwater et al. 1997; Rector & Stocke 2001; Sbarufatti et al. 2006), which we have recently confirmed (see Section 2.6). For consistency with our other papers, we assume the following cosmological parameters: $H_0 = 71 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.27$, $\Omega_\Lambda = 0.73$ (Komatsu et al. 2009). On this model the comoving coordinate distance of PKS 2131–021 is 3.97 Gpc, the angular diameter distance is 1.74 Gpc, and the luminosity distance is 9.08 Gpc. None of the conclusions would be changed were we to adopt the best model of the Planck Collaboration (Planck Collaboration et al. 2020).

2. The Observations

In this section we describe our radio monitoring and VLBI observations and, in addition, infrared and optical observations of PKS 2131–021. Note that for all the single-dish observations the angular extent of PKS 2131–021 is \ll the telescope beamwidth, so that the total flux density is always measured.

2.1. The Radio Light-curve Observations

The 45.1 yr radio light curve of PKS 2131–021 is shown in Figure 1. It is immediately apparent to the eye that there are three distinct epochs, which are demarcated by the vertical black lines in Figure 1: In epoch 1 a strong 1.5-cycle periodic oscillation is seen in the Haystack data (green squares), of which the last half-cycle is also seen in the UMR AO data (brown triangles). There follows a 20 yr period (epoch 2) in which the oscillation of epoch 1 is not seen. In 2003 we enter the third epoch, in which a strong oscillation is again seen and which has very nearly the same period (within 2%) and also the same phase (within 10% of a cycle) as the periodic oscillation of epoch 1. This analysis by inspection is well corroborated by both the Lomb–Scargle (LS; Lomb 1976; Scargle 1982) and weighted wavelet Z-transform (WWZ; Foster 1996) analyses we carried out (see Section 3).

Haystack Observations: The 15.5 GHz ($\lambda 1.9 \text{ cm}$) data (O’Dea et al. 1986) are shown by the green square data points

in Figure 1. The bandwidth was 1200 MHz, and the primary beamwidth (FWHM) was $2\prime.2$. These data were obtained at roughly 1-month intervals with the 120 ft diameter telescope of the Haystack Radio Observatory located in Westford, MA. Details of the observing procedures (Dent et al. 1974b; Balonek 1982) are briefly summarized here. The observations were taken with the dual-feed beam-switched system, which reduces the effect of atmospheric fluctuations. Pointing corrections were applied, and the source antenna temperature was measured by an “on–off” procedure. Visual inspections of a chart record were used to identify observations affected by interference, which were then deleted. Corrections for atmospheric extinction and elevation-dependent gain were applied. The primary flux density calibrators were the compact H II region DR 21, 3C 274, and 3C 123. A correction of a few percent was applied to account for partial resolution of DR 21 and 3C 274 by the Haystack observing beam.

UMRAO Observations: The UMRAO light curve is shown by the brown triangular data points in Figure 1. The long-term 14.5 GHz total flux density data for PKS 2131–021 were obtained with the 26 m equatorially mounted, prime-focus University of Michigan paraboloid as part of the UMRAO monitoring program that operated from the mid-1960s until 2012.5 and included observations of total flux density and linear polarization at three centimeter-band wavelengths for hundreds of blazars in a dedicated AGN program. A description of the general observation and calibration procedures is given in Aller et al. (1985). In brief, at 14.5 GHz dual, rotating, linearly polarized feed horns placed symmetrically about the prime focus were employed and alternated beams on the source; these fed a broadband uncooled high electron mobility transistor (HEMT) amplifier with a bandwidth of 1.68 GHz. Each observation shows the average from 1 day of a series of on–on measurements obtained by alternating between rotating feed horns during an interval of 30–40 minutes. The cadence was one of these averages per week. The adopted flux density scale (Baars et al. 1977) uses Cas A as the primary calibrator, accounting for its measured decay rate (Dent et al. 1974a). Observations of secondary calibrators were interleaved approximately every 2 hr throughout each 14.5 GHz observing run to monitor pointing and gain changes. The standard deviation associated with each daily flux density observation is computed from the system noise temperature and the number of individual on–on measurements made on the particular day. The standard error estimates shown include the effects of measurement noise, the errors introduced by uncertainties in the pointing corrections applied to the observations, and the uncertainties in determining the antenna gain as a function of time. The goal of the UMRAO program was to detect and follow blazar flares, so in general the cadence of the observations was set by the variability state (high cadence during flaring). However, as a member of the UMRAO BL Lac sample (Aller et al. 1999), PKS 2131–021 was also routinely observed approximately every 3 months. Additionally, during the last decade of operation, quasi-simultaneous observations (within a week) were matched to MOJAVE epochs for calibration of the Very Long Baseline Array (VLBA) data.

OVRO Observations: The OVRO light curve is shown by the blue dot data points in Figure 1. The OVRO data were part of the OVRO 40 m Telescope Monitoring Program (Richards et al. 2011). The telescope uses off-axis dual-beam optics in which the beamwidth (FWHM) is $157\prime$. The cryogenic receiver

uses an HEMT amplifier and is centered at 15 GHz with 2 GHz equivalent noise bandwidth. Gain fluctuations and atmospheric and ground contributions are removed with the double switching technique (Readhead et al. 1989), where one of the beams is always pointed on the source. Until 2014 May the two beams were rapidly alternated using a Dicke switch. In 2014 May a new pseudo-correlation receiver replaced the old receiver. Since then a 180° phase switch has been used to alternate the beams. A temperature-stable noise diode is used for relative calibration to compensate for gain drifts. The primary flux density calibrator is 3C 286 with an assumed value of 3.44 Jy (Baars et al. 1977), and DR 21 is used as a secondary calibrator source. Details of the observation and data reduction (Richards et al. 2011) cover the absolute calibration and the uncertainties, which include both the thermal fluctuations in the receiver and systematic errors that have been added in accordance with a rigorous procedure (Richards et al. 2011).

MRO Observations: The 22 and 37 GHz observations were made with the 13.7 m diameter Aalto University Metsähovi radio telescope, which is a radome-enclosed antenna with Cassegrain optics in Finland ($60^\circ 13' 04''$ N, $24^\circ 23' 35''$ E). The receivers have HEMT front ends operating at room temperature. The bandwidth at both frequencies is 1 GHz, and the beamwidths (FWHM) are $4\prime 0$ and $2\prime 4$ at 22 and 37 GHz, respectively. The observations are Dicke-switched ON–ON observations, alternating the source and the sky in each feed horn. A typical integration time to obtain one flux density data point is between 1200 and 1800 s. The detection limit of the system at 22/37 GHz is on the order of 0.2 Jy under optimal conditions, but it is heavily weather dependent. The flux density scale is set by observations of the H II region DR 21. Sources NGC 7027, 3C 274, and 3C 84 are used as secondary calibrators. A detailed description of the data reduction and analysis is given in Teraesranta et al. (1998). The error estimate in the flux density includes the contribution from the measurement rms and the uncertainty of the absolute calibration.

2.2. 15 GHz VLBI Observations

PKS 2131–021 was observed with the VLBA at 15 GHz on numerous occasions between 1995 and 2012 as part of the 2 cm Survey (Kellermann et al. 1998) and MOJAVE (Lister et al. 2018) programs. The last eight maps made on this program, six of which overlap the OVRO monitoring observations, are shown in Figure 2. These show the basic core–jet structure of PKS 2131–021 on parsec scales. The results have been analyzed for 24 VLBA epochs (Lister et al. 2019; including 6 from other projects obtained from the NRAO data archive) by fitting Gaussian model features to the interferometric visibilities and tracking their evolution in flux density (Figure 3) and separation (Figure 4). The VLBA images have a restoring beam with typical FWHM dimensions $\sim 1 \text{ mas} \times 0.5 \text{ mas}$, but the positional determination accuracy of the individual model features is approximately 10 times better. The parsec-scale radio morphology of the source consists of a strong core feature with brightness temperature varying between $10^{10.6}$ and $10^{11.9}$ K and a jet structure that extends 2.5 mas (21 pc projected) to the east. The core accounts for roughly half of the total 15 GHz flux density of the source (Figure 3). The jet and core are linearly polarized, with fractional polarization levels increasing from 5% to 15% downstream. These values are typical of other core-dominated blazars in the MOJAVE

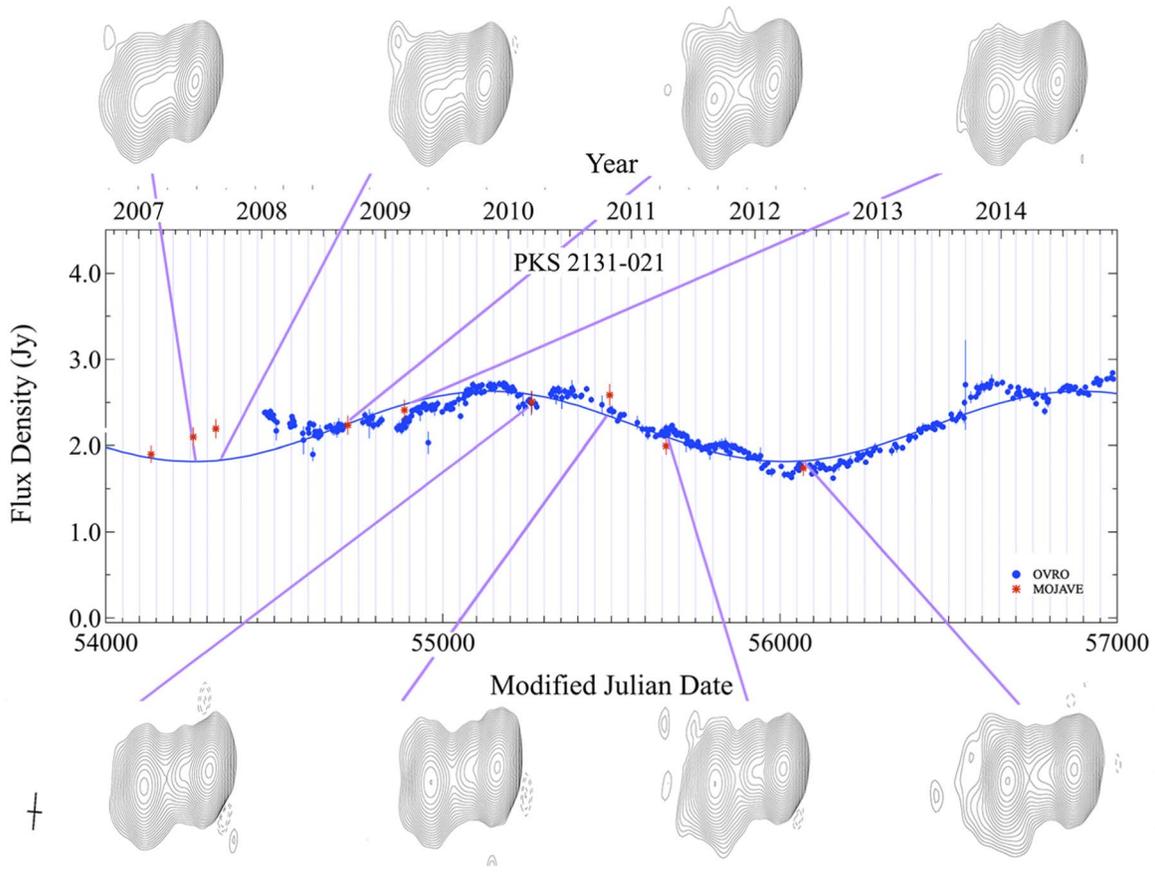


Figure 2. MOJAVE 15 GHz VLBI maps (Lister et al. 2018) matched to the OVRO 15 GHz light curve (blue circles) and least-squares sine-wave fit to the OVRO data (blue line). The total flux densities measured by MOJAVE are shown by the red asterisks. The maps show the basic core–jet structure of PKS 2131–021 on parsec scales. The bar at the lower left indicates the scale of 5 mas (~ 42 pc) of all the maps, and the small cross indicates the typical FWHM beam size and orientation. Detailed comparison of the maps, models, and light curve shows that it is the core (primarily) and innermost component (partly), at ~ 0.3 mas from the core, that are responsible for the periodic flux density variations.

sample (Lister & Homan 2005). Observations in 2006 (Hovatta et al. 2012) showed low Faraday rotation measure, with values ranging up to 131 rad m^{-2} , and the electric vectors are generally aligned with the jet direction, indicative of a magnetic field oriented perpendicular to the flow. The one exception was a jet feature (component 3) that had electric vectors oriented in the transverse jet direction.

In Figure 3 we show the light curves of the individual milliarcsecond jet components measured by MOJAVE (Lister et al. 2019). These overlap the 2005 peak in the sine-wave signal seen by UMRAO and the 2010 peak seen by both UMRAO and OVRO. Note that the core component and the component nearest to the core (component 5) are the components that produce the periodic signal. The sum of these two components is shown by the solid line in Figure 3. In Figure 4 we show the separations of the individual components from the core. Component 5 is stationary ($\mu = 3.0 \pm 1.9 \mu\text{as yr}^{-1}$; Lister et al. 2019) at a separation of 0.3 mas (0.2 pc) from the core. We note that, while we might expect all of the periodic variability to arise in the core, in view of the small separation of component 5 from the core, it is not unlikely that the core and component 5 will vary in phase on the timescale of the observed periodicity.

Based on near-simultaneous OVRO and VLBI observations, the VLBI flux density is $\gtrsim 98\%$ of the single-dish flux density, so that, at most, ~ 25 mJy of flux density is missing on small

angular scales and is distributed on large scales. PKS 2131–021 has the typical blazar jetted-AGN morphology of a one-sided jet with a flat-spectrum core at one end of a steep-spectrum jet. The components in the jet have been modeled (Lister et al. 2019) with the results shown in Figures 3 and 4. The parsec-scale jet of PKS 2131–021 showed significant activity during the period between mid-1993 and 2000, with several features emerging from the core and moving downstream at apparent superluminal speeds. No new moving features appeared in the jet between 2000 and 2013. As can be seen in Figure 3, at all epochs except the first, the core dominates the flux density. From 2001 March 15 a stationary component (#5) is seen at a distance of 0.3 mas from the core, and together the core plus component 5 dominate the light curve from then on. The sinusoidal flux density fluctuations seen in epoch 3 are clearly due to these two components and dominated by the core. Thus, the periodicity we observe in PKS 2131–021 is coming primarily from unresolved components closer to the central engine and the base of the jet than the structures observed with 15 GHz VLBI.

Kinematic analysis (Lister et al. 2019) identified five distinct features in the jet, each having a different apparent speed. The innermost feature (id = 5) had no detectable motion over a 12 yr period, while the speeds of the other features ranged from $2.6c \pm 0.2c$ (id = 4) to $20c \pm 2.6c$ (id = 2). Most of the moving features display accelerated and/or nonradial motions on the sky

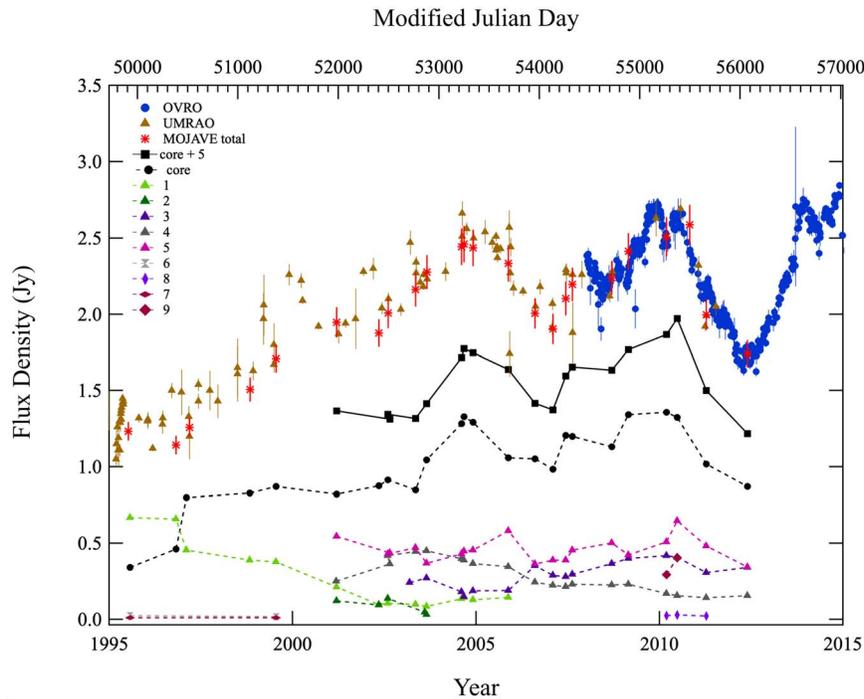


Figure 3. Radio light curves of the entire source from UMRAO at 14.5 GHz (brown triangles), OVRO at 15.0 GHz (filled blue circles), and MOJAVE at 15.0 GHz (red asterisks), as well as individual parsec-scale jet features as measured by the MOJAVE program.

that are exaggerated by projection effects. Given the maximum apparent speed, the jet viewing angle, θ , is constrained to lie within $5^\circ.7$ of the line of sight, and an analysis based on the median core brightness temperature yields $\theta = 3^\circ.8$ and a Doppler factor of $\delta = 14$ (Homan et al. 2021).

The exact dates of emergence of the individual moving features are difficult to determine since that would require extrapolations of unknown accelerations close to the core. Also, in the case of components 3 and 4, there is blending with other jet emission that made it impossible to measure their positions when they were located near the core. Simple constant-velocity radial fits to components 1 and 2 give ejection dates of 1993.7 ± 0.4 and 1998.6 ± 0.5 , respectively. For component 3, the first five epochs (when the feature was moving approximately radially away from the core) imply an ejection date of 1999.2 ± 0.3 . In the case of the strongly accelerating component 4, its time of ejection can only be narrowed down to the interval 1997–2000.

2.3. The Large-scale Radio Structure of PKS 2131–021

The radio structure of PKS 2131–021 on different angular scales is shown in Figure 5. On arcsecond scales PKS 2131–021 has an unusual structure, which is most clearly seen in the 4.86 GHz VLA map of Rector & Stocke (2003)—reproduced here in Figure 5(b), which shows two linear features extending to about $8'$ from one side of the nucleus, and nothing on the opposite side of the nucleus. An EVN-MERLIN map by Cassaro et al. (2002) shows this bifurcated structure extending inward to within 100 mas of the core. Both gravitational lensing and precession have been suggested as possible causes for the strange structure. Another possibility, which has not previously been suggested to the best of our knowledge, are twin jets, such as are seen in 3C 75 (Owen et al. 1985). It is interesting that two of the possible explanations for

the large-scale structure of PKS 2131–021 (precession and twin jets) could well require an SMBHB.

2.4. Higher-frequency Radio Observations

The MRO 22 GHz and 37 GHz observations are shown in Figure 6(a). We have cross-correlated the MRO 37 GHz observations with the OVRO 15 GHz light curve, with the results shown in Figure 6(b). There is clearly a correlation, as can also be seen by visual examination of the light curves, and it also appears that the 15 GHz variations lag those at 37 GHz, but it is hard to determine the lag—it could be anywhere between ~ 0 and ~ 200 days.

2.5. Infrared Observations

We have extracted the Wide Field Infrared Explorer (WISE) data in the $2.8\text{--}3.8 \mu\text{m}$ and $4.1\text{--}5.2 \mu\text{m}$ bands. These are shown in Figures 6(c) and (e). The cross-correlation between the flux densities of these two bands with the radio light curve can be seen in Figures 6(d) and (f). The cross-correlation function is rather noisy, but it likely shows a real correlation when it is at, or just above, the 2σ level. If this is to be believed, then the radio variations lag the infrared variation by $\sim 250\text{--}350$ days.

2.6. Optical Observations

We observed PKS 2131–021 on 2021 October 6 (HST) with the Low Resolution Imaging Spectrometer (LRIS) mounted on the Keck I telescope. The observations were carried out in dark time, under photometric conditions with a median seeing disk FWHM of $0''.5$. A 900 s exposure was obtained at an airmass of 1.08, with the D560 dichroic, the 400/3400 blue-side grism, and the 400/8500 red-side grating (central wavelength 7830 \AA). Spectral flux and telluric-absorption calibration was carried out using an observation of the standard star Feige 34. Data reduction followed standard procedures using the lpipe

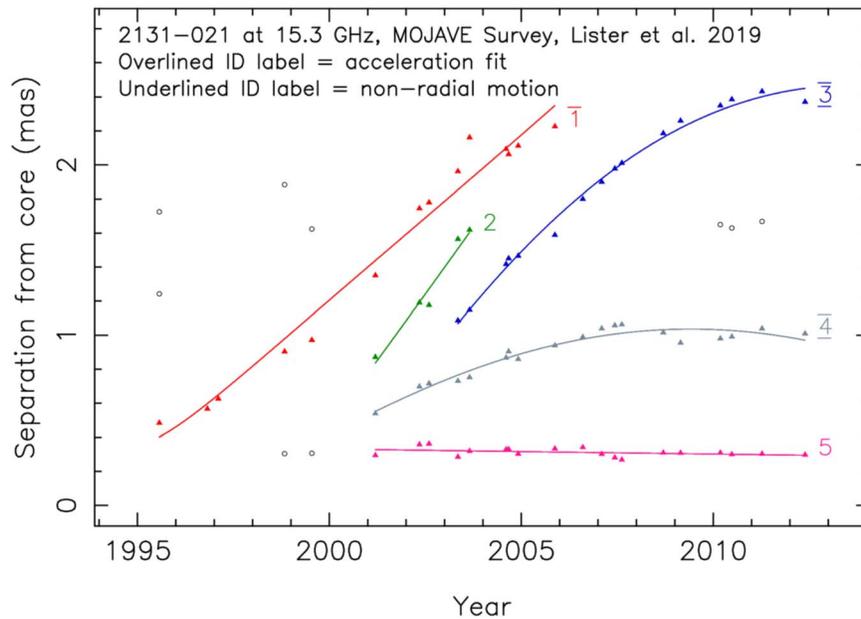


Figure 4. The angular separation of jet components from the core vs. time for individual Gaussian model fitted features in PKS 2131–021 on the MOJAVE program. Colored symbols indicate robust features for which kinematic fits were obtained. The identification number is overlined if an acceleration model was fitted and indicated a $\geq 3\sigma$ acceleration. An underlined identification number indicates a feature with nonradial motion. The black circles represent Gaussian components fitted to the jet emission that could not be clearly cross-identified over a minimum of five epochs.

software (Perley 2019). The software was used to perform bias subtraction using the overscan levels, flat-fielding using dome-flat exposures, automatic cosmic-ray rejection, wavelength calibration using internal comparison-lamp exposures and sky emission lines, optimal sky-line subtraction, and optimal extraction of the spectral trace. No significant difference was observed between our spectrum of PKS 2131–021 and a previous spectrum obtained on 1995 November 20 (Rector & Stocke 2001). We confirm the redshift of 1.285, together with the BL Lac nature of the object.

The archival optical light curve of the blazar is presented in Figure 7. It consists of 119 V-band epochs from the Catalina Real-Time Transient Survey (CRTS; Drake et al. 2009) and 203 g-band epochs from the Zwicky Transient Facility (ZTF; Graham et al. 2019; Masci et al. 2019). The CRTS data were collected between 2005 May and 2015 December using the 0.7 m Catalina Schmidt Telescope located north of Tucson, Arizona. The ZTF data were taken between 2018 May and 2021 September with a wide-field camera mounted on the Palomar 48-inch Oschin (Schmidt) telescope.

We searched for periodic variations in the optical data. We ran generalized Lomb–Scargle (GLS; Lomb 1976; Scargle 1982) and analysis of variance period-searching algorithms (Schwarzenberg-Czerny 1989) on the CRTS and ZTF data, and we did not find any significant periodicity in the range 1–1850 days. The optical light curve is not correlated with the radio data.

3. Analysis of the 14.5–15.5 GHz Radio Light Curve

We have long been aware of the dangers of overinterpretation of AGN light curves. Before analyzing our radio light curve of PKS 2131–021, we outline the approach we have adopted to avoid the red-noise pitfall and other systematic issues that could undermine our key results.

In order to make statistically robust estimates of the chance probability of occurrence of any feature in a radio light curve,

we need a process of simulating the light curve of any AGN such that all of the statistical and variational characteristics of the AGN are preserved in the simulations, while at the same time taking into account that the data are not equi-spaced owing to weather and hardware problems. We can then simulate a large number of light curves for any AGN and estimate the probability of chance occurrence of any feature we observe in the light curve of any particular AGN. More details are given in Appendix A.

The problem of generating simulated light curves has been addressed in a number of papers (Welsh 1999; Uttley et al. 2002; Emmanoulopoulos et al. 2010; Max-Moerbeck et al. 2014), which simulate light curves with PSD of the same variability power-law slope as observed in the AGN. However, the underlying pdf is Gaussian and does not produce realistic light curves. Radio AGN light curves exhibit “burst”-like events, which can often yield long-tailed pdf’s.

This problem was elegantly solved by Emmanoulopoulos et al. (2013) using a simple method that precisely reproduces light curves that match both the PSD and the pdf of the observed light curves. As desired, the final artificial light curves have all the statistical and variability properties of the observed light curves and are statistically (and visually) indistinguishable from the true light curves. We have implemented this approach¹⁸ and applied it to the light curve of PKS 2131–021 in our analysis. For each epoch we simulated light curves that have the same PSD, pdf, sampling, and “observational” noise as the real data of that specific epoch. We are therefore confident that the significances we calculate are robust, having correctly taken into account the red noise in the observed pdf.

In our analysis we have carried out a detailed WWZ analysis (Foster 1996), as well as a detailed GLS periodogram analysis (Lomb 1976; Scargle 1982; Zechmeister & Kürster 2009), and

¹⁸ <https://github.com/skiehl/lcsim>

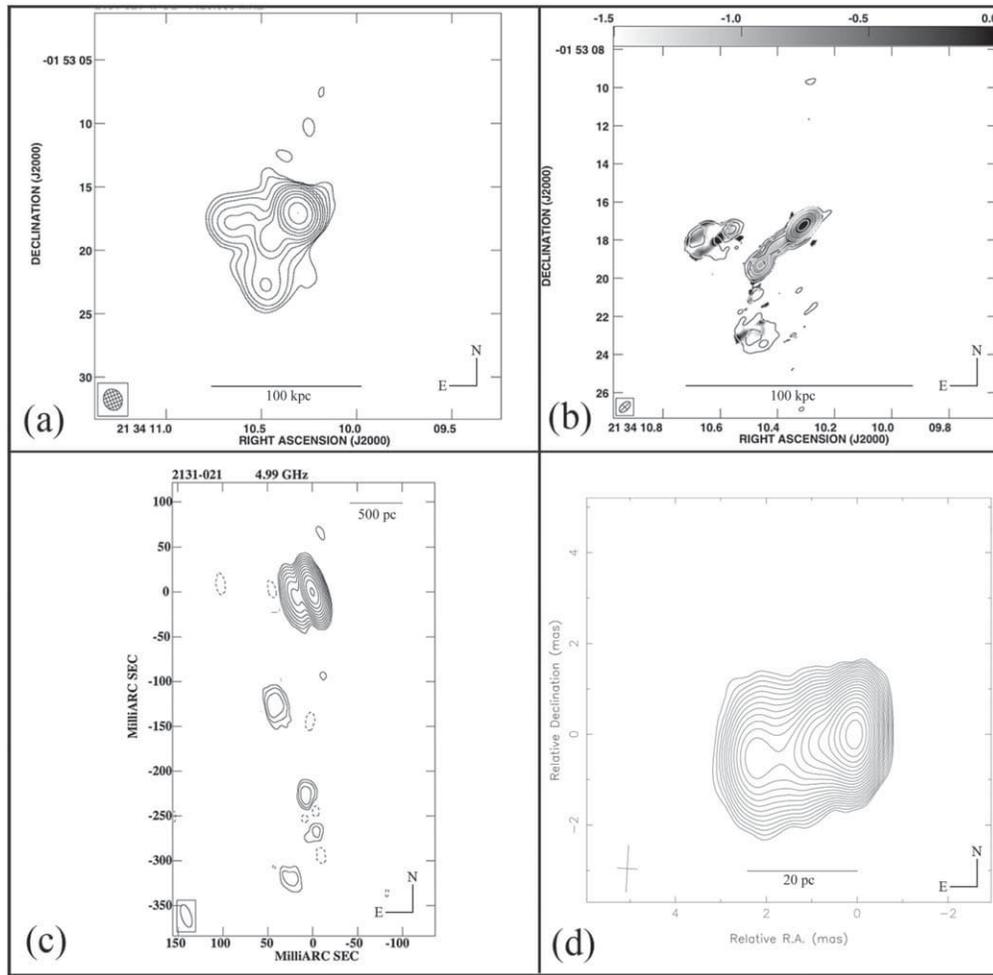


Figure 5. Maps of PKS 2131–021 on different angular scales. (a) VLA 1.4 GHz map from Rector & Stocke (2001). (b) VLA 4.86 GHz map from Rector & Stocke (2003); the gray scale shows the 4.86–8.46 GHz spectral index, which shows that the core has a flat spectrum and the jet has a steep spectrum. (c) EVN+MERLIN 4.99 GHz map from Cassaro et al. (2002). The rms noise level is $0.2 \text{ mJy beam}^{-1}$, and the peak flux density is $1349 \text{ mJy beam}^{-1}$. (d) Stack of 24 MOJAVE images (Lister et al. 2019) from 1995 July 28 to 2012 May 24. FWHM beamwidth: $1.04 \times 0.45 \text{ mas}$ in position angle -3° . Contours: $\sqrt{2} \times 1.714 \text{ mJy beam}^{-1}$.

we find the results to be fully consistent with our results derived from least-squares sine-wave fitting of the light curve.

3.1. WWZ Analysis of the Light Curve

The WWZ is a Z-transform (Ragazzini & Zadeh 1952) that uses the technique of wavelet analysis to detect time-dependent periodicity in data (Foster 1996). We have performed a thorough WWZ analysis¹⁹ of the light curves of the three epochs and all combinations thereof. These showed very clearly that there are three distinct epochs. This is most easily seen in the WWZ plot covering the whole 45.1 yr light curve shown in Figure 8. The power across this period varies strongly, and the raw WWZ plot (top panel) does not have the dynamic range to enable us to see the key features very clearly—more details are given in Appendix B. For this reason, we calculated time-normalized WWZ estimates by dividing the WWZ power in each 100-day time bin by the maximum power in the time bin in order to bring out the details of any changes in period (bottom panel). This WWZ plot supports the conclusion we arrived at by visual inspection of the light curve that there are three distinct epochs of activity in

PKS 2131–021 over the 45.1 yr span of our observations, with a periodic oscillation of approximately the same period dominating in epochs 1 and 3 and a periodic oscillation of about twice that period dominating in epoch 2 (see Table 2). It is interesting to note that the boundary between epochs 1 and 2 is at the point where the WWZ power shifts from one period to another, whereas the boundary between epochs 2 and 3 is later than the point where the WWZ power shifts from one period to another. We ascribe the dominant variability in epoch 2 to red noise. It is possible that the transition from epoch 2 to epoch 3 occurred earlier, but we decided to retain the boundary we picked out by eye since the periodicity of epoch 3 is very clear from this point on and not nearly as clear before that, so that we suspect that all of the interval we demarcate as epoch 2 is dominated by red noise. Furthermore, we tried shifting the boundary between epoch 2 and epoch 3 between MJD 51,200 and MJD 53,800, and we found that it had little effect on the results, as shown in the notes to Table 1. Clearly, the precise choice of the boundary between epoch 2 and epoch 3 is immaterial for the purposes of this study.

Since the results of the WWZ analysis were entirely consistent with those from our GLS analysis and least-squares sine-wave analysis, we do not reproduce further details of our WWZ analysis here other than to give the periods associated

¹⁹ <https://github.com/skiehl/wwz>

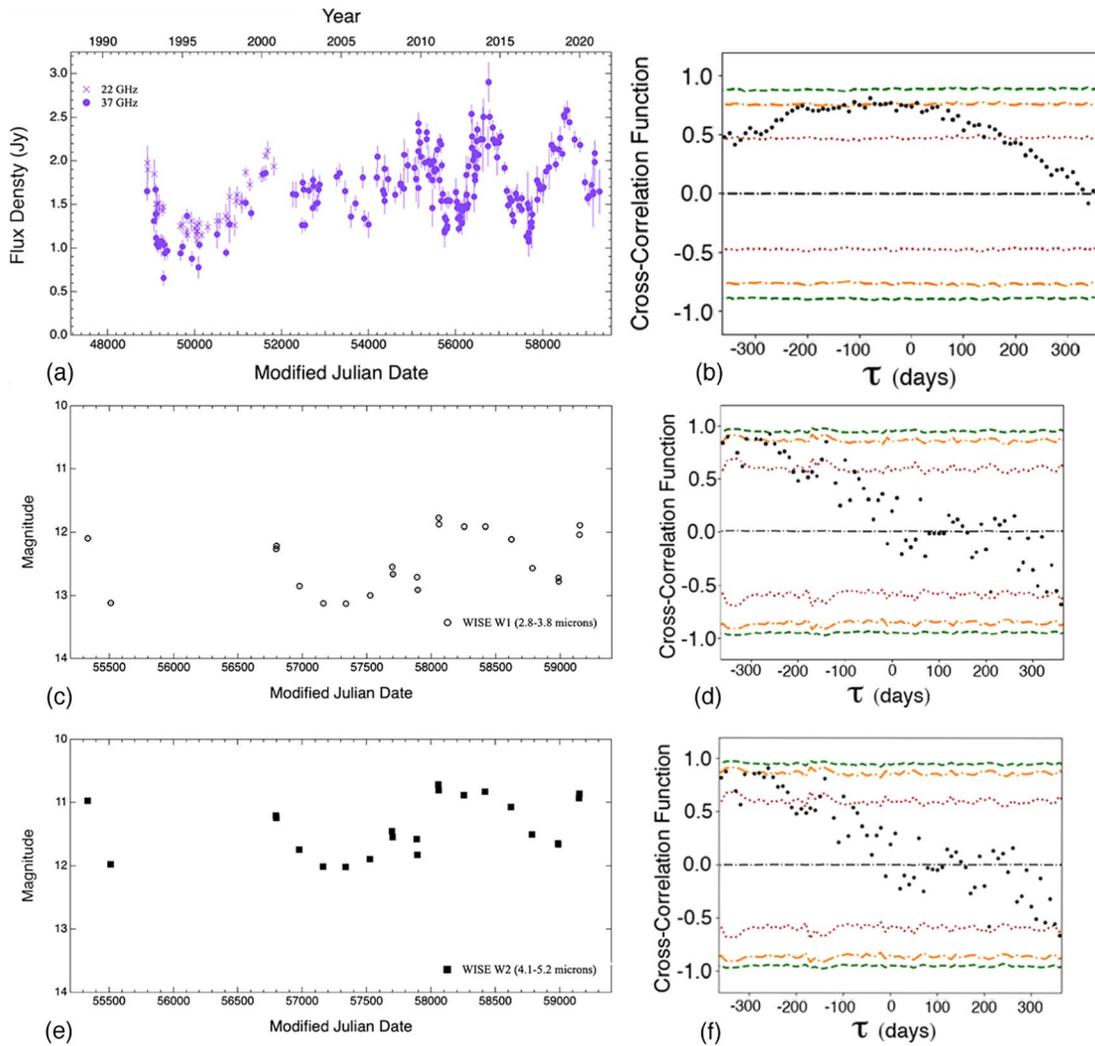


Figure 6. PKS 2131–021 multifrequency light curves and cross-correlations. (a) Light curves at 22 GHz (purple crosses) and 37 GHz (purple circles) from MRO. (b) Cross-correlation of the MRO 37 GHz and OVRO 15 GHz light curves. (c) WISE infrared light curve in the W1 (2.8–3.8 μm) band. (d) Cross-correlation of the WISE W1 (using flux density) and OVRO 15 GHz light curves. (e) WISE infrared light curve in the W2 (4.1–5.2 μm) band. (f) Cross-correlation of the WISE W2 (using flux density) and OVRO 15 GHz light curves. In both panels (c) and (e) the error bars are smaller than the symbols. The pink, orange, and green dashed curves indicate the 1σ , 2σ , and 3σ significance levels of the cross-correlation, respectively, using our simulated light curves, which assume a simple power-law form with a slope of 2, typical for blazars in these bands.

with the maximum power level in our WWZ analysis for the three epochs (see the notes to Table 1).

3.2. Lomb–Scargle Analysis of the Light Curve

We carried out a GLS periodogram analysis (Lomb 1976; Scargle 1982) using the GLS () periodogram²⁰ (Zechmeister & Kürster 2009) on each epoch of the light curve of PKS 2131–021, as well as all combinations thereof. In order to distinguish it from the period of the periodicity, P , we use the symbol \mathcal{P} to denote the GLS power. Following Scargle (1982), we evaluate the GLS at frequencies between $f_{\text{low}} = 1/T$, where T is the total time of the corresponding epoch, and $f_{\text{high}} = N/(2T)$, corresponding to one over twice the average time interval between measurements, with N the number of data points. The selection of a pseudo-Nyquist frequency as the highest frequency does not have any effect on the results, as all relevant peaks are found at much lower

frequencies. We sample this frequency range uniformly in steps of $\Delta f = 1/(\zeta T)$. For our analysis we conservatively chose $\zeta = 10$, since with this value we can be confident that we have adequately sampled periods up to the duration of the observations. For comparison, Scargle (1982) used $\zeta = 5$.

We also tested the effect of the frequency grid spacing with coarser and finer grids, $\zeta = 5$ and $\zeta = 20$, and found only slight differences compared to the values for $\zeta = 10$. The variations with ζ were far too small to affect our conclusions in the analysis presented here.

The results of the GLS analyses of epochs 1, 2, and 3, both independently and combined, are shown in Figure 9 and Table 2. There are about 1.5 cycles of the periodic variation in the Haystack+UMRAO data and about 3 in the UMRAO+OVRO data. While this is a small number of cycles, our simulation procedure takes long-term fluctuations into account via the red tail in the spectrum, and so we are confident that our significance levels are robust. Clearly, epoch 3 has the statistically strongest periodicity, and the similarity of the period with that of epoch 1, as well as the phase

²⁰ <https://github.com/sczesla/PyAstronomy>

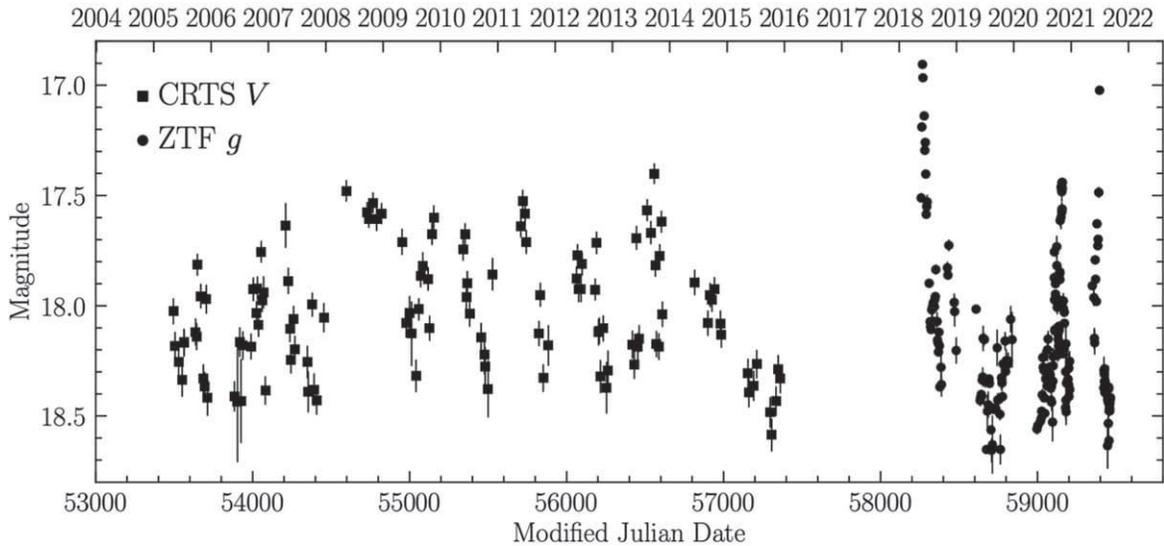


Figure 7. Archival optical light curve of PKS 2131–021.

correspondence, which are what immediately caught our attention when we first saw the Haystack observations, when taken together, are highly significant, as we will show.

In the combined GLS analysis in the bottom panel of Figure 9 the largest peak corresponds to the epoch 3 peak, the second-largest peak corresponds to the epoch 1 peak, and the third-largest peak corresponds to the epoch 2 peak. All three of these peaks in the combined analysis of the bottom panel are slightly shifted in period relative to their period values in the individual epoch GLS analyses. Given that the epoch 3 peak is the only peak detected at greater than 3σ significance, and that the spectrum has a powerful red-noise component, the peaks other than the epoch 3 peak should not be regarded as physically significant unless supported by other data, as is true for the epoch 1 peak, but not for any of the other peaks seen in the combined analysis. This, plus the WWZ analysis presented in Figure 8, illustrates the potential danger of carrying out GLS, or WWZ, analyses on astronomical light curves without adequately simulating the variability in the object and, in addition, applying rigorous statistical criteria.

3.3. Sine-wave Least-squares Fitting of the Light Curve

We fitted the light curves in the three epochs with a sine-wave function:

$$S_{\text{model},i} = A \sin(\phi_i - \phi_0) + S_0, \quad (1)$$

where $\phi_i = 2\pi(t_i - t_0)/P$ is the phase of the i th data point. Our model has five free parameters: P —period, A —amplitude, ϕ_0 —phase of the sine wave, S_0 —mean flux, and σ_0 —characteristic amplitude of an intrinsic AGN variability. We keep $t_0 = 51,000$ fixed (which is in the midpoint between the start of epoch 1 and end of epoch 3 observations). We find the best-fitting parameters by maximizing the following likelihood function:

$$\ln \mathcal{L} = -\frac{1}{2} \sum_{i=1}^N \left(\frac{(S_i - S_{\text{model},i})^2}{\sigma_i^2 + \sigma_0^2} + \ln(\sigma_i^2 + \sigma_0^2) \right). \quad (2)$$

The uncertainties are estimated using the EMCEE sampler (Foreman-Mackey et al. 2013) and represent the 68%

confidence range of the marginalized posterior distribution. The σ_i are taken from the observations, assuming independent Gaussian errors, and the intrinsic variability is represented by Gaussian white noise of variance σ_0^2 . We have not used the red-noise spectrum because to include it in the fit would mean including the full covariance matrix rather than just the diagonal term. We know that including a red-noise term of the type seen in epoch 2 can change the period at the $\sim 10\%$ level over periods of a few years (see Section 3.5).

Fit results, separately for epochs 1, 2, and 3, are presented in Table 1. Note that the uncertainties of ϕ_0 are relatively large because there are strong degeneracies between period and phase, especially as the adopted t_0 is relatively far from the actual observations. If we chose t_0 to be closer to epoch 1 (epoch 3) observations, the respective uncertainties of ϕ_0 would be much smaller. The degeneracy between the period and phase reflects the fact that, over time, information about the phase is lost.

We also simultaneously fitted both epoch 1 and epoch 3 data by fitting a single wave, with the same period and phase extending across both epochs, but permitting the amplitude and offset between epoch 1 and epoch 3 to vary. Results of this joint fit are presented in the last column of Table 1.

Prior to MJD 45,500 (epoch 1), the strong sinusoidal variation seen by Haystack+UMRAO has period $P_1 = 1729.1 \pm 32.4$ days. In epoch 2 ($45,500 < \text{MJD} < 52,850$), the UMRAO data show sinusoidal variability, but no periodicity at the frequency or in phase with that seen in the first and third epochs. We therefore think that this is a red-noise phenomenon, and we will show that it fails our 3σ criterion by a substantial margin. In epoch 3 ($\text{MJD} > 52,850$) strong sinusoidal variation returns with period $P_3 = 1760.4 \pm 5.3$ days. In addition, the phase of the sinusoidal variation from epoch 3, when extrapolated back to epoch 1, matches to within 10% of P , or 20% of $P/2$, which is the relevant number here since we associate the extrapolated curve with the nearest peak. Combining the Haystack, UMRAO, and OVRO data from epoch 1 and epoch 3 yields a period $P_{13} = 1737.9 \pm 2.6$ days, i.e., the fractional uncertainty in period $\delta P/P \sim 1.5 \times 10^{-3}$.

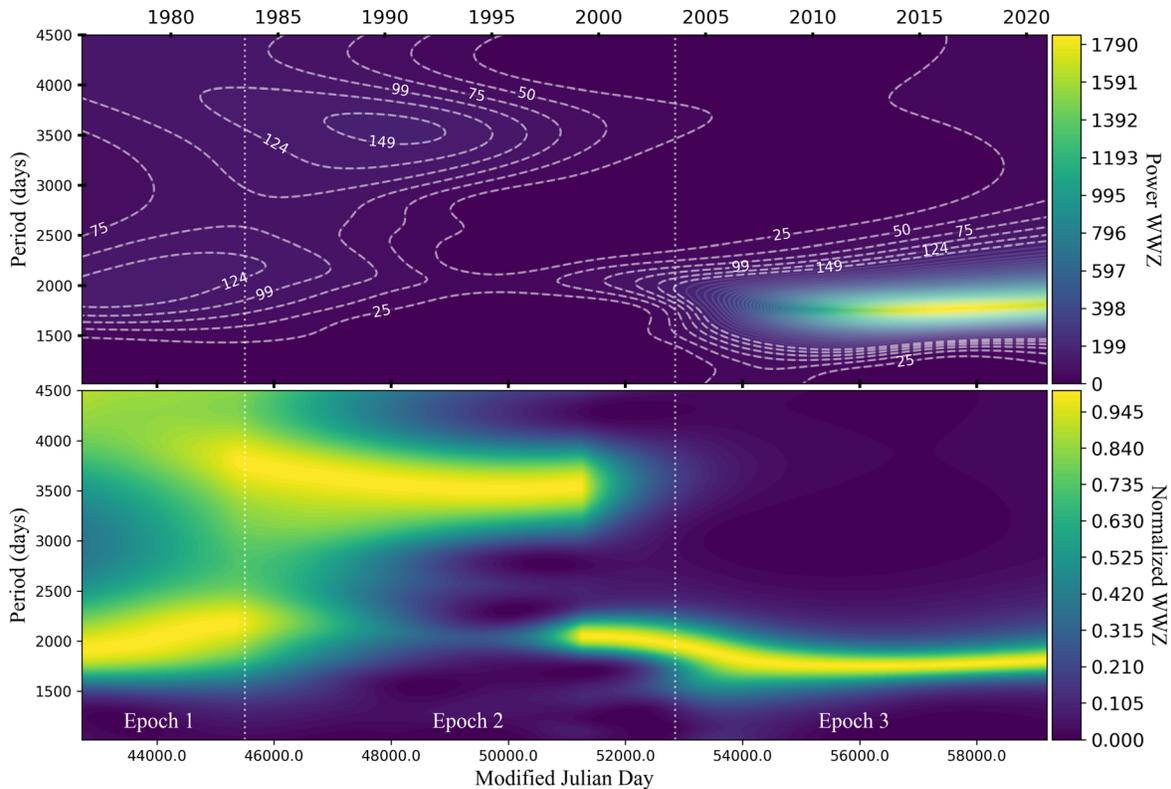


Figure 8. Plots of the WWZ statistic computed from the PKS 2131–021 45.1 yr light curve. Top panel: WWZ statistic plot, in which the color scale at the right and the contour levels on the plot represent the value of the WWZ statistic (see Appendix B). In this panel all the features apart from the feature in epoch 3 are almost invisible because they are more than an order of magnitude weaker than the epoch 3 feature. Bottom panel: time-normalized WWZ plot of the 45.1 yr light curve of PKS 2131–021. Here we normalized each of the 164 100-day time bins by dividing the power across all the periods in each 100-day time bin by the maximum power in the bin. Thus, every time bin has a maximum power of unity. This enables us to see clearly how the periods change over the whole 45.1 yr span of the observations.

3.4. The Significances of the Observed Periodicities

To calculate the significance of the power levels observed in epochs 1, 2, and 3, we count the number of simulations in which the p -value, p_{sim} , of the strongest peak in the simulation is less than the observed peak p -value, p_{peak} , for that epoch. The procedure is explained in more detail in Appendix A. The results are shown in Table 2. We see that the peaks in epoch 1 and epoch 2 are significant only at the 1.06σ and 1.53σ levels, respectively (Tests 1.1 and 1.2). However, the peak in epoch 3 is significant at the 3.35σ level (Test 1.3). This makes it clear that PKS 2131–021 has a periodicity in epoch 3 that is not due to red noise, which must therefore be associated with some physical mechanism in the object.

We next turn to the good agreement in the periods, straddling the 20 yr gap in periodic fluctuations, seen in epochs 1 and 3 (see Table 1), which, to our knowledge, is unprecedented in observations of AGNs. The question we address is this: having observed the periodicity in epoch 3, which is what first drew our attention to this source, what is the probability of observing the same periodicity in epoch 1, at an acceptable LS power level and within the 3σ period window of our epoch 3 period?

We first determine what we deem to be an acceptable GLS power level, \mathcal{P}_{lim} . From Table 2, we see that, of the three periodicities we have been discussing, the lowest power level is that of the peak in epoch 1, at $\mathcal{P}_{\text{peak}} = 0.71$. From Figure 9 we see that, apart from the largest peak in each epoch, all the other peaks in power are below the $\mathcal{P} = 0.4$ power level. We therefore select a threshold level of power of $\mathcal{P}_{\text{lim}} = 0.5$ as a

power level suitably above the level of sidelobes in the GLS spectrum and suitably low that we do not miss any strong sinusoidal features in the light curves. We will therefore count all simulations that give an GLS power level $\mathcal{P}_{\text{lim}} \geq 0.5$ in epoch 1, in our test of the significance of the epoch 1 periodicity, given the epoch 3 periodicity.

We now turn to the range of periods to be considered. From Table 1 we see that the uncertainty in the difference between P_1 and P_3 is $\sqrt{32.4^2 + 5.3^2} = 32.8$ days. Thus, the 3σ uncertainty in the difference is 98.5 days. Since we wish to calculate the probability of observing the periodicity in epoch 1, given that it had been observed in epoch 3, we will count simulations in epoch 1 where the strongest peak lies in the period range $P_3 \pm 98.5$ days, i.e., from 1661.9 to 1858.9 days. So our two criteria are GLS power $\mathcal{P} \geq \mathcal{P}_{\text{lim}} = 0.5$ and $1661.9 \text{ days} < P < 1858.9 \text{ days}$.

The probability of observing in epoch 1 a peak within the $\pm 3\sigma$ window of the periodicity observed in epoch 3 and with $\mathcal{P} \geq 0.5$ is shown in Test 2 of Table 2 to be 1.97×10^{-2} , which is significant at the 2.06σ level. We may now combine Test 1.3 and Test 2 to determine the probability of observing both the observed periodicity in epoch 3 and a periodicity in epoch 1 with period within the $\pm 3\sigma$ window of the periodicity found in epoch 3. This is determined in Test 3 by multiplying $p(\text{Test 1.3}) \times p(\text{Test 2}) = 7.88 \times 10^{-6}$, which is significant at the 4.3σ level. When we add the correspondence in phase to within one-fifth of a half-cycle, the probability drops to $0.2 \times p(\text{Test 3}) = 1.58 \times 10^{-6}$, which is significant at the 4.6σ level.

Table 1
Sine-fitting Results for Epoch 1, Epoch 2, Epoch 3, and the Joint Epoch 1 + Epoch 3 Data Sets

| | epoch 1 | epoch 2 | epoch 3 | epoch 1 + epoch 3 |
|----------------------|-------------------|-------------------|-------------------|-------------------|
| P (days) | 1729.1 ± 32.4 | 3779.1 ± 46.0 | 1760.4 ± 5.3 | 1737.9 ± 2.6 |
| ϕ_0 | 0.89 ± 0.46 | 0.35 ± 0.08 | 0.60 ± 0.07 | 0.88 ± 0.03 |
| A (epoch 1) | 0.709 ± 0.047 | ... | ... | 0.679 ± 0.045 |
| S_0 (epoch 1) | 2.553 ± 0.036 | ... | ... | 2.584 ± 0.034 |
| σ_0 (epoch 1) | 0.333 ± 0.022 | ... | ... | 0.337 ± 0.023 |
| A (epoch 2) | ... | 0.392 ± 0.020 | ... | ... |
| S_0 (epoch 2) | ... | 1.724 ± 0.021 | ... | ... |
| σ_0 (epoch 2) | ... | 0.140 ± 0.009 | ... | ... |
| A (epoch 3) | ... | ... | 0.400 ± 0.007 | 0.400 ± 0.007 |
| S_0 (epoch 3) | ... | ... | 2.225 ± 0.005 | 2.229 ± 0.005 |
| σ_0 (epoch 3) | ... | ... | 0.118 ± 0.004 | 0.120 ± 0.004 |

Note. We also determined the least-squares sine fit to epoch 3 for shifted boundaries between epoch 2 and epoch 3 at MJD 51,200 and MJD 53,800, with the results $P = 1762.9 \pm 6.1$ days and $P = 1756.0 \pm 5.5$ days, respectively. These may be compared with the result above of $P = 1760.4 \pm 5.3$ days for the boundary set at MJD 52,850. The peak periods identified in our WWZ analyses of epochs 1, 2, and 3 were 1740, 3919, and 1779 days, respectively. Those in our GLS analyses were 1730, 3937, and 1788 days, respectively.

We believe that these significance estimates, based as they are on simulations that have the same PSD and the same pdf as the observed light curves in epochs 1, 2, and 3, and therefore with the same red-noise tail as PKS 2131–021, are robust, and therefore that this analysis demonstrates conclusively that the periodicities observed in epochs 1 and 3 are connected and that there is a physical process that maintains this period over our 45.1 yr observing period even when it is not manifested in the light curve.

From Table 1 we see that the observed difference in period between the epoch 3 data and the epoch 1 + epoch 3 data is significant at the 3.8σ level, so this might indicate that the period is slowly changing. Separate sine-wave fits to the first half and the second half of the OVRO data show a difference of $\sim 10\%$ in period. PKS 2131–021 is a blazar and, as such, highly variable even in the absence of the periodicity we have been discussing, as is easily seen from the deviations from sinusoidal variations observed in epochs 1 and 3 and throughout epoch 2. Thus, we ascribe the apparent changes in period to random red-noise variability in this blazar, such as that seen in epoch 2. Indeed, because of the strong red-noise component in their variability, such variations in the observed period in radio light curves of an SMBHB in a blazar like PKS 2131–021 are to be expected, but these should average out in the long term.

The reader may have noticed the apparent harmonic relationship between the epoch 2 and epoch 1+3 periodicities. In this paper we do not wish to draw undue attention to this for two reasons:

1. The epoch 2 periodicity has significance of only 1.53σ , as shown in Table 2. We do not discuss this periodicity since we think any discussion runs counter to the approach we are adopting in this paper. We encourage workers in this field to carefully assess red noise before entering discussions of apparent periodicities.
2. There is no obvious phase relationship in between the real, highly significant, periodic signals we see in epochs 1 and 3 and the low-significance apparent periodicity in epoch 2.

We reemphasize here our deliberate decision to consider only signals that are demonstrably not due to red noise, i.e.,

only those having significance greater than 3σ , unless supported by other evidence.

3.4.1. The “Look Elsewhere” Effect

In the GLS analysis the periods are quantized because they are based on discrete frequencies (see Appendix D of Scargle 1982). It is common, in studies of QPOs, to see significance contours, derived from simulations, plotted on GLS and WWZ plots. To the best of our knowledge, such contours are determined by counting the number of simulations at each quantized period and dividing by the total number of simulations, as we have done in Figure 9 (top panel) for the sole purpose of illustrating this common pitfall. This is a legitimate approach *only* if that periodicity has been pre-selected, for example, a known period in a double star, with known uncertainty in the period, in which case the value of ζ must be chosen so as to include the corresponding range of periods at the peak period sampling interval (see Section 3.2). However, there are two problems with this approach in general. The first is that since the sampling in periodicity is discrete and depends on ζ , the single-period p -value $\propto 1/\zeta$; the second is that if there is no a priori reason for selecting a particular period, then the significance of any peak must take into account all of the simulations in which the largest peaks have p -values that are less than or equal to the p -value of the peak under consideration, as explained in Appendix A. Our three epochs on PKS 2131–021 provide a striking example of the necessity for taking the p -values of all the simulated light curves into account when assessing the significance of an observed feature. The results are shown in Table 3. The true significances address the question, “What is the probability of finding a peak of *any* periodicity that has a p -value less than or equal to that of the observed peak?” In other words, one has to look elsewhere than just at the period corresponding to the peak in the GLS plot. The problem of using single-period probabilities, and of propagating these by plotting misleading significance curves on GLS plots, further exacerbates, in addition to the red-noise problem, the problem of random noise peaks being considered as if they are physically significant.

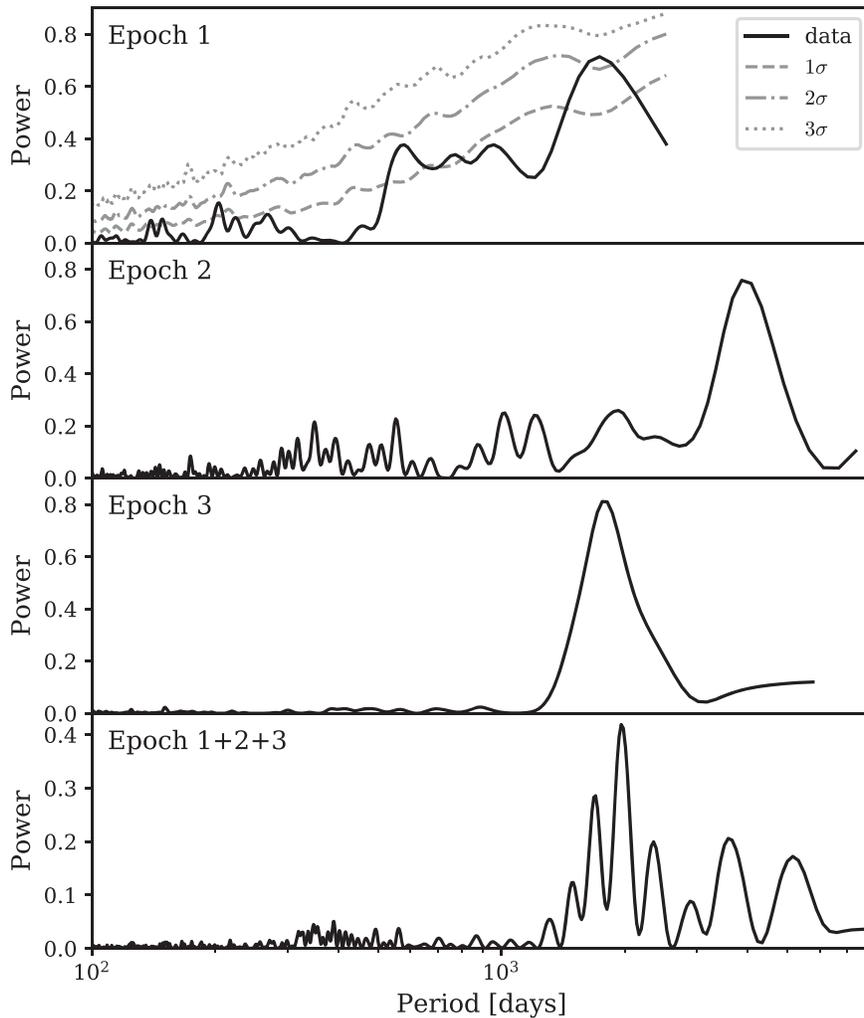


Figure 9. GLS analyses of the light curve of PKS 2131–021. In the upper three panels we show the GLS analyses for the three distinct epochs of activity. As is common practice, in the top panel the contours purport to show the significance levels of the features. The apparent significance of the peak in epoch 1 (top panel) is 2.3σ , whereas the true significance is 1.06σ (see Table 3). We do not plot contours in the middle and bottom panels because such contours are misleading and further exacerbate the problem of the misidentification of random features as significant peaks (see Section 3.4.1). The bottom panel shows the GLS power plot for the full 45.1 yr duration of the observations, i.e., for epochs 1, 2, and 3 combined. (see text).

3.5. Variations in the Period due to Noise

The fits of the sine waves shown in Figure 1, while good, are not perfect. There is a clear epoch from 2015 to 2017 when the sine wave is systematically above the data, and similarly in 2020 the data are systematically above the sine wave. If the two halves of the OVRO data are analyzed separately, the periods differ by $\sim 10\%$. Does this mean that PKS 2131–021 is simply another QPO? The unique properties we have drawn attention to above suggest that it is not.

We now provide an illustration, based on the variability during epoch 2, that refines the period derived for epoch 3 to $P = 1737.6 \pm 3.6$ days, shown by the gray curve in Figures 10(a) and (b), in good agreement with the period and phase observed over the whole 45.1 yr. This is merely an example of what must be happening, under our hypothesis of an SMBHB. We go through the following argument to illustrate the effect that we think the nonperiodic varying components must be having on the observations. Of course, this does not prove that PKS 2131–021 is an SMBHB. But the exercise is illuminating, and in our view it does therefore strengthen the case for an SMBHB on the grounds of plausibility.

We have fitted the period when the sinusoidal variations we are investigating were absent using the epoch 2 data between MJD 45,500 (1983 June 15) and MJD 52,850 (2003 July 30) and found that a 6° polynomial, shown by the brown curve in Figure 10(a), follows the slowly varying component well, whereas lower-degree polynomials do not. We adopted the least-squares sine-wave fit to epoch 1 + epoch 3 and subtracted this, adjusted for the amplitudes of the Haystack and OVRO periodic signals, from the OVRO data to give the residual light curve shown by the black points in Figure 10(a). This shows the flux density variations in PKS 2131–021, under the SMBHB hypothesis and in the absence of the sinusoidal signal. We then adopted the same approach with the OVRO data as with the UMRAO data and fitted a 6° polynomial to the residual light curve to give the black curves shown in Figures 10(b) and (c). These show what the slowly varying components of the source were doing, under the SMBHB hypothesis and apart from the periodic signal. Correcting for the slowly varying components, we derive the gray points shown in Figure 10(c). The sine-wave fit to the corrected OVRO data is shown by the gray curve in Figures 10(a) and (b), which has period $P = 1737.6 \pm 3.6$ days. Thus, as

Table 2
Probabilities and Significance Levels of GLS Tests Computed from Simulations with Matched Red-noise Tail

| Test Number | Test | GLS $\mathcal{P}_{\text{peak}}$ (max = 1) | Period Range (ΔP) (days) | Total Simulations | Number of Simulations That Pass Test | p -value | Significance (σ) |
|-------------|--|--|---------------------------------------|-------------------|---|-----------------------|------------------------------|
| 1.1 | epoch 1 $\mathcal{P}_{\text{peak}}, p_{\text{sim}} \leq p_{\text{peak}}$ | 0.71 | All | 10,000 | 1446 | 1.45×10^{-1} | 1.06 |
| 1.2 | epoch 2 $\mathcal{P}_{\text{peak}}, p_{\text{sim}} \leq p_{\text{peak}}$ | 0.76 | All | 10,000 | 632 | 6.32×10^{-2} | 1.53 |
| 1.3 | epoch 3 $\mathcal{P}_{\text{peak}}, p_{\text{sim}} \leq p_{\text{peak}}$ | 0.81 | All | 100,000 | 40 | 4.0×10^{-4} | 3.35 |
| 2 | epoch 1 $\mathcal{P}_{\text{lim}}, \Delta P_{\text{epoch 3}}$ | ≥ 0.50 | 1661.9–1858.9 | 10,000 | 197 | 1.97×10^{-2} | 2.06 |
| 3 | 1.3+2 | ... | ... | ... | ... | 7.88×10^{-6} | 4.32 |
| 4 | 3+phase | ... | ... | ... | ... | 1.58×10^{-6} | 4.66 |

Note. \mathcal{P} is the GLS power, and ΔP is the range of periods included in the test. In Tests 1.1, 1.2, and 1.3 we count simulations at all periods with p -values less than the p -value of the peak in that epoch. Since the significance of the peak in epoch 3 was so high, we needed 100,000 simulations to determine the significance. The low individual significance levels for epochs 1 and 2, shown in Tests 1.1 and 1.2, indicate that they are easily produced by red noise. Test 2 is the number of simulations of epoch 1 that were seen in the $\pm 3\sigma$ wide period window centered on the epoch 3 period ($P = 1760.4$ days).

Table 3
Single-period (Spurious) and All-period (True) Probabilities

| Epoch and Test | N_{tot} | n_{pass} | p -value | σ |
|-----------------------|------------------|-------------------|-----------------------|-------------------|
| epoch 1 single period | 10000 | 108 | 1.08×10^{-2} | 2.3 |
| epoch 1 all periods | 10000 | 1446 | 1.45×10^{-1} | 1.06 ^a |
| epoch 2 single period | 10000 | 122 | 2.20×10^{-3} | 2.85 |
| epoch 2 all periods | 10000 | 632 | 6.32×10^{-2} | 1.53 ^a |
| epoch 3 single period | 100000 | 0 | $< 10^{-5}$ | > 4.26 |
| epoch 3 all periods | 100000 | 40 | 4.00×10^{-4} | 3.35 ^a |

Note. The tests using all periods are the “Look Elsewhere” tests.

^a These are the true significances; the others are totally spurious unless the periods have been selected “a priori.”

expected, it has a period closer to the period of the epoch 1 + epoch 3 sine-wave fit ($P = 1737.9 \pm 2.6$ days) than the original sine-wave fit to the epoch 3 data ($P = 1760.4 \pm 5.3$ days). In addition, and *as expected*, the phase during epoch 1, as shown by the gray curve in Figure 10(a), is much closer to that of the joint epoch 1 + epoch 3 fit (black curve) than the original epoch 3 fit (blue dotted curve). This example simply shows how small, slow variations in the nonperiodic signal can have a significant effect on the period and phase of the sinusoidal variations in the periodic signal in PKS 2131–021, a point that, while obvious, is worth investigating to illustrate the levels of correction that are needed to achieve much better coherence. The fractional changes in the flux density required for perfect coherence, which would be grossly overfitting the data, are shown by the black points in Figure 10(c). Note that they are nearly all $< 10\%$.

Itemization of the above procedure:²¹

1. We fitted separate least-squares sine waves to the epoch 1 and epoch 3 data.
2. We subtracted these sine waves from the Haystack data of epoch 1 and the OVRO data of epoch 3.
3. We assumed that the sinusoid represented all of the SMBHB emission, and therefore that the SMBHB signal was now removed from the Haystack data and the OVRO data.
4. We assume that epoch 2 is already devoid of SMBHB signal. Thus, in Figure 10 the black data in epoch 1, the

brown data in epoch 2, and the black data in epoch 3 represent the data with the SMBHB sinusoid subtracted.

5. We fitted a six-order polynomial to the black data in epoch 3 in Figure 10. The six-order polynomial is thus assumed to represent the red-noise variation of the SMBHB.
6. We then subtracted the polynomial fit given by the black line from the original data, which gives the gray points.
7. Finally, we did a least-squares sine-wave fit to obtain the period for the gray points.

Summary: Our motivation in carrying out the above procedure is as follows. We have, quite deliberately, “made the data fit the SMBHB model” in order to provide an illustration of what we think the nonsinusoidal random flux density variations very likely were during the OVRO observations, and then to compare these to the nonsinusoidal random variation during epoch 2. This does not prove that PKS 2131–021 is an SMBHB, but, given the significance of the similarities in period and phase in epochs 1 and 3, as shown in Table 2, this is almost certainly the case. Thus, understanding and anticipating the effects of the nonsinusoidal variability in a jetted SMBHB is bound to be important as this field opens up.

It is interesting to note that the periodic fluctuations occur at the times of highest flux density in the overall light curve, so that it appears that the periodic signal is in addition to the usual flux density level in this source. Thus, on our SMBHB model we are proposing that PKS 2131–021 has an extremely stable period and that the small changes in period that we see are due to other sources of variation that distort the pure sine-wave variation. We also propose that the emission associated with the periodic variability can turn off and on in a time that is short compared to the period, and that the amplitude of the periodic signal can vary significantly between epochs of periodic emission.

Our observations thus lead to three strong predictions based on the SMBHB hypothesis: (1) we expect the future long-term averaged period to agree with our measured period spanning 45.1 yr of observations ($P = 1737.9 \pm 2.6$ days); (2) in the short term we expect to see small ($\sim 10\%$) variations in the period, due to the corrupting effects of the other varying emission components in this object; and (3) we expect the SMBHB sinusoidal signal to continue to appear and disappear at random times.

²¹ This helpful summary was provided by our anonymous referee.

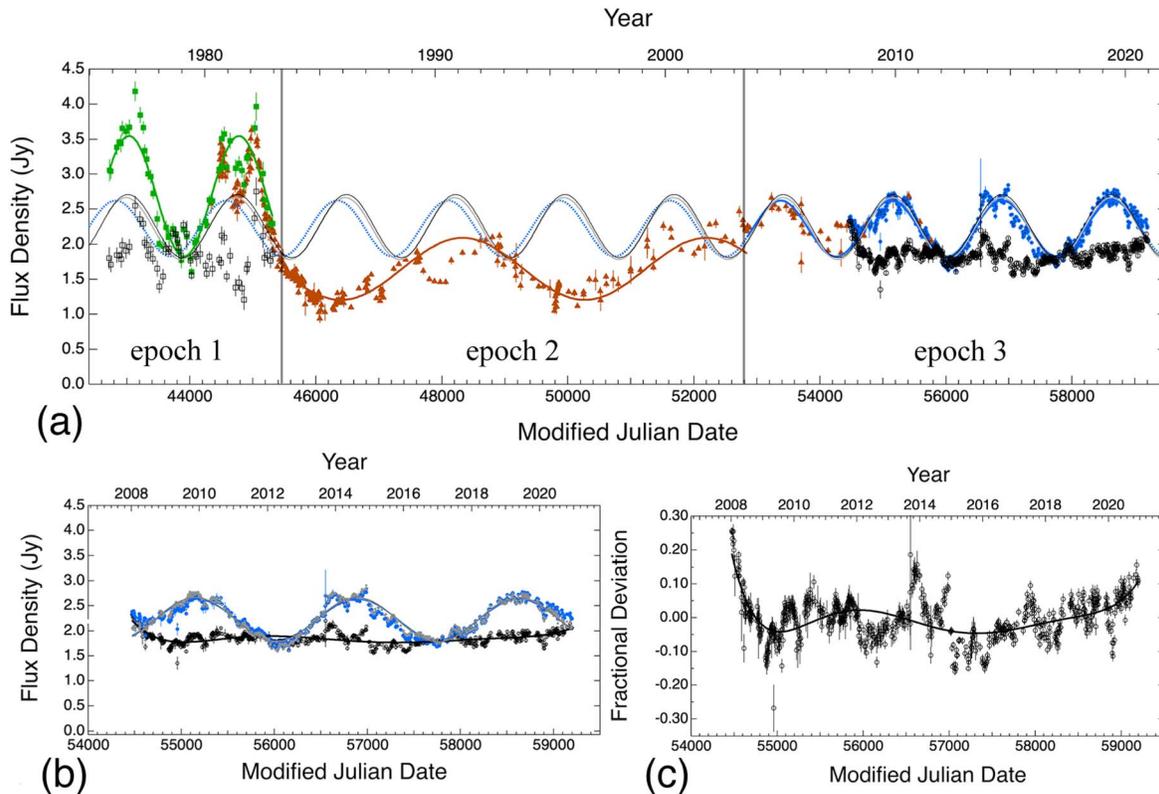


Figure 10. Period variations in the 14.5–15.5 GHz light curve of PKS 2131–021. (a) The Haystack data are shown by the green squares, the UMRAO data by the brown triangles, and the OVRO data by the blue circles. The solid green line shows the least-squares sine-wave fit to the Haystack data, and the dotted blue line shows the least-squares sine-wave fit to the OVRO data extrapolated backward to provide a phase comparison with the Haystack data. In epoch 2 we show a 6th-degree polynomial fit to the light curve, and it is seen to be sinusoidal in character. This oscillation has about twice the period of that in epoch 1, but it bears no obvious phase relationship with the oscillation of epoch 1, and it also bears no obvious phase relation with the epoch 3 periodicity. In epoch 3 we have fitted the combined Haystack and OVRO data (black sine wave). This fit, adjusted for the different amplitudes of the Haystack and OVRO periodic signals, has been subtracted from the OVRO data to provide the residual light curve in absence of the periodic signal (black data points). The black curve in panels (b) and (c) shows a 6th polynomial fit to the residual OVRO data. The gray curve is the sine-wave fit to the OVRO data after correcting for the slowly varying component in the residual signal of epoch 3 (black point), which brings the phase into alignment with that of epoch 1, as expected. (b) The blue points and blue curve show the raw OVRO data and sine-wave fit, and the gray points show the corrected OVRO light curve after applying the slowly varying component given by the black 6th polynomial fit to the residuals. The gray sine wave is the least-squares fit to the adjusted (gray) OVRO data. (c) Fractional flux density deviation from the mean value with the OVRO sine wave subtracted. This shows the fractional flux density variations that would be required to bring the Haystack and OVRO data into phase coherence.

3.6. The Long-term Stability of the Periodic Variations in the Light Curve

It is clearly important to investigate the stability of the periodic variations seen in PKS 2131–021, since this has potential implications for the SMBHB hypothesis, under which the period should decrease with time. For the purposes of comparison we will denote the periods determined in our least-squares sine-fitting analysis in the different epochs (see Table 1) by $P_1 = 1729.1 \pm 32.4$ days, $P_2 = 3779.1 \pm 46.0$ days, $P_3 = 1760.4 \pm 5.3$ days, and $P_{13} = 1737.9 \pm 2.6$ days.

The most interesting of these is P_{13} since it represents a single least-squares sine-wave fit across the 45.1 yr span. On the hypothesis that PKS 2131–021 is an SMBHB, this represents our best estimate of the observed orbital period. On the SMBHB hypothesis we therefore predict that, in future long-term monitoring, this period will be maintained until the effects of gravitational radiation produce a noticeable reduction in the period.

We therefore predict on the SMBHB hypothesis that the next minimum is occurring now (2021 October), and we will be able to confirm this in the next 6 months. The next maximum is predicted in 2024 February, and the following minimum in 2026 July. It should be possible to confirm or disprove the

SMBHB hypothesis within the next 5–10 yr, provided that we do not go back into an epoch 2–like phase, in which this periodicity vanishes.

We now consider the stability of the period as can best be determined from the observations. We have already seen that $\sim 10\%$ variations occur over time spans of a few years, which we ascribe to the corrupting effects of the variability in PKS 2131–021 that are unrelated to the underlying periodicity we seek to study. Taking their associated uncertainties into account, we see that $\delta P_{13} = P_3 - P_1 = 31.3 \pm 32.8$ days. So the 3σ window is $\delta P_{13} = -67.2 \rightarrow +129.8$ days.

On the SMBHB hypothesis we are interested in negative δP_{13} , and in this case our fractional 3σ limit is $\delta P_{13}/P_{13} = -67.2/1737.9 = -3.9 \times 10^{-2}$. The midpoint of epoch 1 is MJD = 44,109.8, and that of epoch 3 is MJD = 56,023.6, so the span between the midpoints at which P_1 and P_3 were measured is 11,913.8 days, or $\delta t = 32.6$ yr. Thus, the 3σ upper limit on the fractional period decrease per year is $\delta P_{13}/P_{13} \delta t = -1.19 \times 10^{-3} \text{ yr}^{-1}$. This is the maximum rate of fractional period decrease per year consistent with our observations. From Maggiore (2008), using our rest-frame quantities, we derive a chirp-mass upper limit of $5.4 \times 10^9 M_\odot$. This may be compared with the slightly lower upper limit on

the chirp mass derived from the GW limit discussed in Section 5.

3.7. Summary of Our Approach and Statistical Tests

In this section, before presenting our model for PKS 2131–021 and its implications, we summarize briefly our approach and the logical thread of our statistical tests. The key steps in our approach are as follows:

1. The recognition of the role of red noise in the γ -ray, optical, and radio light curves of blazars, as exemplified in the work of Vaughan et al. (2016) and Covino et al. (2019). As a consequence, we require strong statistical evidence that any apparent periodicity is not simply the result of red noise in the light curve, and we apply a 3σ threshold that must be met by any apparent periodicity before it should be discussed. Without such a threshold, workers in this field will expend a great deal of time and energy seeking physical explanations for apparent periodicities that are the result of random fluctuations and have no other physical significance. Consideration of such phenomena will obscure the physics and hinder progress in the field.
2. The recognition that there are three distinct epochs in the 45.1 yr radio light curve of PKS 2131–021, as indicated by the periodicities seen in Figure 1. The WWZ analysis of Section 3.1, including Figure 8, and the analysis of different MJD dates for the transition from epoch 2 to epoch 3 given in Section 3.1, and in the notes to Table 1, in addition to our original visual inspection, are our justification for this.
3. Starting with epoch 3, the epoch that initially drew our attention to the periodicity in PKS 2131–021, we carried out a GLS analysis making no assumptions about periodicities, and we considered all periodicities via the “look elsewhere” effect. We found that the probability at, or above, the observed power level in the GLS analysis is 4×10^{-4} , i.e., it is significant at the 3.35σ level, and thus it is above the threshold of point #1 above.
4. We then investigated the sinusoidal periodicity of epoch 1 as follows: we selected a $\pm 3\sigma$ -wide periodicity window, based on the uncertainties in the least-squares periodicities in epoch 1 and 3, and determined the probability of a chance periodicity at or above the GLS power observed in epoch 1 in the window of this width and centered on the periodicity of epoch 3. This probability is 1.97×10^{-2} , i.e., it is significant at the 2.06σ level.
5. Since the data from epochs 3 and 1 are independent, the probability of the agreement in periods observed in epoch 1 and epoch 3 is given by the product of the two probabilities and is 7.88×10^{-6} , i.e., it is significant at the 4.32σ level.
6. Finally, we took into account the agreement, to within 20% of half a period, of the phase of the epoch 3 periodicity extrapolated to epoch 1. This has a random probability of 0.2. Thus, the probability of the phase agreement in addition to the periodicity agreement between epoch 1 and epoch 3 is 1.58×10^{-6} , i.e., it is significant at the 4.66σ level. This leaves us in no doubt that the periodicity observed in epochs 1 and 3 is a significant physical property of PKS 2131–021 and is not simply a random manifestation of red noise.

4. A Model of PKS 2131–021

We propose the simple model shown in Figure 11, of a black hole binary. We express the masses of the binary components in units of $10^8 M_\odot$, so the primary mass is given by $M_1 = 10^8 M_{1,8} M_\odot$, and similarly for M_2 . The secondary, with mass M_2 , orbits the primary with period in the rest frame of the binary $P = 760.6$ days and angular momentum that makes an angle i with the line-of-sight unit vector \mathbf{n} . We assume that the motion is circular, but elliptical orbits work as well. We assume that the jet is launched along the spin axis of M_1 with fixed velocity $c\beta$ relative to the black hole. The jet could originate in M_2 , but this makes no difference to the discussion. The source is a BL Lac object, and we know from the above VLBI observations that the line of sight is inclined at a small angle, θ , to the jet axis. Let the orbital velocity of M_1 be $\beta_1 = 0.036M_{2,8}/(M_{1,8} + M_{2,8})^{2/3}$, where, and henceforth in this section, we set $c = 1$.

On our model the orbital motion changes the velocity of the emitting material in the jet relative to the observer, and hence the Doppler factor and beaming. In the case of PKS 2131–021 the Doppler factor of the jet is high, since this is a superluminal source with the jet axis closely aligned with the line of sight, as discussed in Section 2.2. In such a case, the orbital motion can have a significant effect on the Doppler factor, as we show below.

Suppose that we have a source at rest emitting isotropically and observed distantly with a flux density S' . The observed flux density will be given by $S = \mathcal{D}^{2-\alpha} S'$ (Scheuer & Readhead 1979), where $\alpha = d \ln S / d \ln \nu$ is the spectral index, \mathcal{D} is the Doppler factor

$$\mathcal{D} = \frac{1}{\gamma(1 - \beta \cdot \mathbf{n})}, \quad (3)$$

and $\gamma = (1 - \beta^2)^{-1/2}$ is the Lorentz gamma factor. The orbital motion causes both γ and β to change.

Applying the Lorentz transformation from the rest frame of M_1 to the rest frame of the binary barycenter, we find

$$\delta\beta = \beta_1 - (\beta \cdot \beta_1)\beta + O(\beta_1^2), \quad (4)$$

and we also have

$$\frac{\partial \ln \mathcal{D}}{\partial \beta} = \frac{\mathbf{n}}{1 - \mathbf{n} \cdot \beta} - \frac{\beta}{1 - \beta^2} \quad (5)$$

assuming that \mathbf{n} is fixed. Hence, the fractional change in S is given by

$$\delta \ln S = (2 - \alpha) \delta \ln \mathcal{D} = (2 - \alpha) \delta\beta \cdot \frac{\partial \ln \mathcal{D}}{\partial \beta} \quad (6)$$

$$= \frac{(2 - \alpha)(\mathbf{n} - \beta) \cdot \beta_1}{(1 - \mathbf{n} \cdot \beta)} + O(\beta_1^2). \quad (7)$$

A continuous jet comprises many such sources, starting and finishing at a supposed fixed rate. Their combined flux density should then satisfy the same relation, to $O(\beta_1^2)$.

In PKS 2131–021, as we have seen in Section 2.2, $\theta \sim \gamma^{-1} \ll 1$. Expanding, we find that $\delta \ln S$ varies sinusoidally,

$$\delta \ln S = \frac{2(2 - \alpha)\gamma^2\theta\beta_1 \cos i}{(1 + \gamma^2\theta^2)} \cos(2\pi t/P). \quad (8)$$

The observed amplitude, $\delta \ln S \sim 0.2$, is compatible with masses $M_{1,8} \sim M_{2,8} \sim 1$ and typical blazar values $\gamma \sim \theta^{-1} \sim 10$.

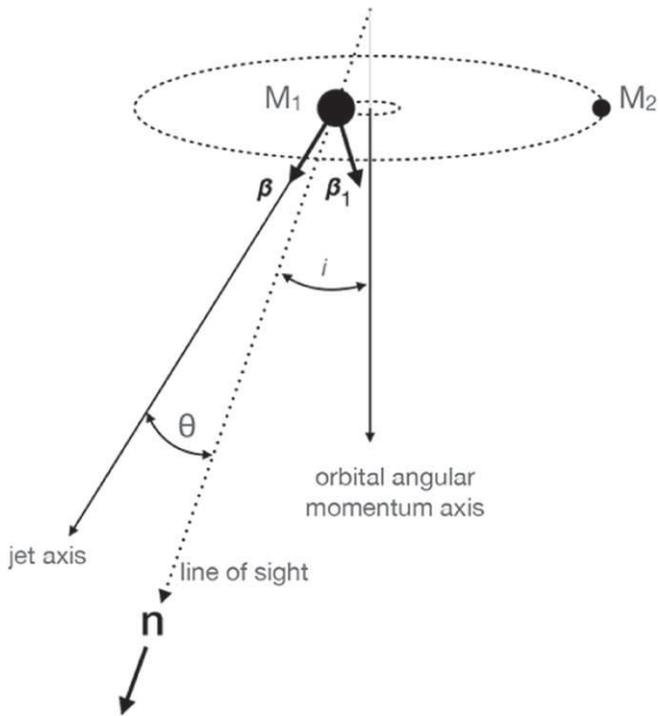


Figure 11. Simple model that produces sinusoidal flux density variations in an SMBHB blazar. In the proposed SMBHB, the larger mass, M_1 , and smaller mass, M_2 , orbit the center of mass with period in the rest frame of the binary, P . We assume that the jet originates in M_1 , although this is not necessary for the model. The angle between the jet axis and the line of sight is θ . We suppose that the relativistic jet is launched along the black hole spin axis with constant velocity $c\beta$ relative to the black hole. The orbital velocity can significantly change the relativistic jet Doppler factor if $\gamma \gg 1$ in the jet, as is the case in PKS 2131–021.

Intuitively, one would expect the greatest effect on the observed Doppler factor to occur if the orbital motion is parallel or antiparallel to the jet, and the least effect to occur if the orbital motion is perpendicular to the jet. However, as can be seen in Equation (7), this is not the case owing to the $(\mathbf{n} - \beta)$ term. Since $|\mathbf{n}|$ is unity and β is very close to unity, $\mathbf{n} - \beta \sim \theta$. Now θ is in the plane of the sky, and hence orthogonal to the line of sight, and it is θ that operates on β_1 , as we see in Equation (7). Hence, the effect is proportional to $\cos i$, rather than $\sin i$, contrary to what one would expect intuitively, i.e., a binary whose orbital plane is normal to the line of sight will give the largest effect, whereas a binary whose orbital plane lies along the line of sight will give zero effect to order $O(\beta_1)^2$.

This result is applicable to the observed variation and is sinusoidal, as measured. However, the quadratic and higher-order terms in an expansion will contribute harmonics. This model is highly simplistic. Real jets are unlikely to be ballistic, as they interact with their surroundings. They can accelerate and decelerate. Also the emission is quite likely to originate from radii that are not much smaller than cP , introducing retarded time effects that will also lead to higher harmonics.

However, this simple model demonstrates that with plausible values of the black hole masses and jet speed it is possible to account for the sinusoidal flux density variations observed in PKS 2131–021. Numerical simulations of jets can be used to explore the observed behavior of orbiting jets more carefully.

As discussed in Section 2.3, precession has been suggested to explain the unusual large-scale radio structure of PKS 2131–021.

Begelman et al. (1980) show that in an SMBHB the more massive component will undergo geodetic precession with period

$$P_{\text{prec}} \sim 400 r_{16} \left(\frac{M_1}{M_2} \right) \frac{P}{M_{18}}, \quad (9)$$

in which the orbit is assumed to be circular, and where r_{16} is the separation of the binary in units of 10^{16} cm. In PKS 2131–021 $r_{16} \sim 1$ for $M_{18} = 1$ (see Equation (13)).

It is therefore entirely possible, and even likely, that the unusual large-scale radio structure in PKS 2131–021 is due to geodetic precession, while the radio and infrared light-curve periodicity is due to orbital motion.

From Figures 5(a), (b), and (c), we measure the opening angle of the jet to be $73^\circ \pm 18^\circ$. As discussed in Section 2.2, the viewing angle in PKS 2131–021 is $\theta = 3^\circ.8$ (Homan et al. 2021). Thus, the precession cone opening angle would be $4^\circ.8 \pm 1^\circ.2$, which would seem entirely reasonable.

A more detailed model for the nature of the activity observed in PKS 2131–021 goes beyond the scope of this paper but should be able to reproduce the following features of the light curve:

- (a) Periodic variations with period observed on Earth $P_{\oplus} = 1737.9 \pm 2.6$ days are episodic and dominate the light curve at times and are at other times invisible. As mentioned in the Introduction, it is not difficult to invent models in which the sinusoidal signal switches on and off. However, the appearance and disappearance of the sinusoidal signal is an added complexity that would have to be accommodated in any complete model of the PKS 2131–021 system.
- (b) When the periodic behavior is observed, it manifests itself with a remarkable stability of period and phase.
- (c) The flux density is lower when the periodic behavior is absent.
- (d) The amplitude of the periodic behavior, during different epochs of its manifestation, can be very different.

All these features can be naturally explained assuming that we are observing the superposition of the output from two distinct emission processes:

- (I) Periodic, originating in a part of the jet directly affected by the SMBHB.
- (II) Nonperiodic, originating over a wide range of locations along the jet.

Process (II) is the usual mode of activity seen in VLBI-resolved emission in typical blazar jets: strongly variable, largely stochastic, as shown by the sinusoid-subtracted data in Figure 10. Features (c) and (d) of the light curve, pointed out above, strongly suggest that process (I) is also strongly variable in radiative output. If this were not the case, it would be hard to explain how it is possible for the periodicity to disappear when the overall flux density decreases. The qualitative features of the summed light curve during different epochs clearly depend on the relative level of activity between the two processes:

1. If both processes (I) and (II) are in a high state, we see a noisy periodic signal, as in epoch 1.
2. If process (I) is in a low state, we see only aperiodic variations produced by process (II). The flux density level depends on the level of activity of (II), as seen in epoch 2

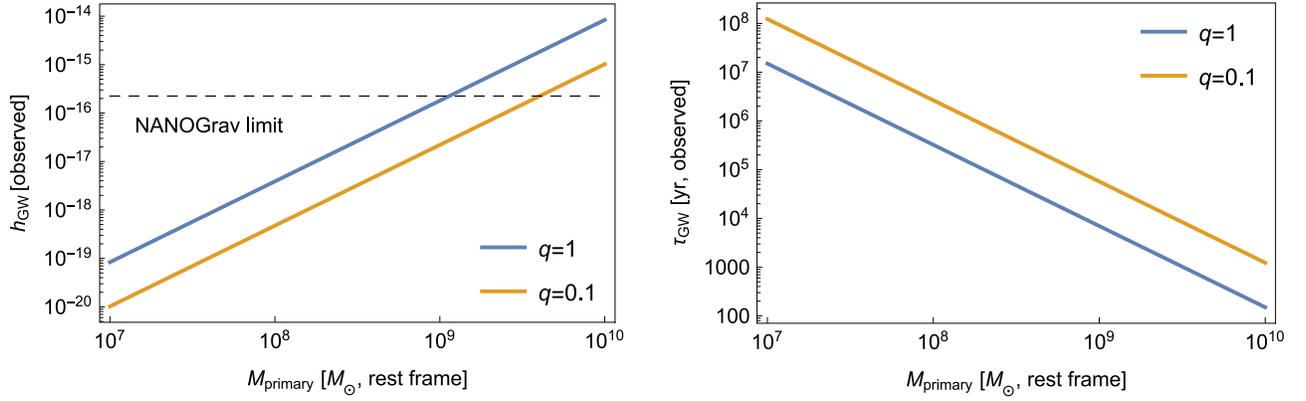


Figure 12. Mass limits from GW observations. Both the period and the location of PKS 2131–021 are in the windows of maximum sensitivity for the NANOGrav pulsar timing array (Aggarwal et al. 2019). Here we show the observed GW strain and time to coalescence for primary–SMBH rest masses ranging from 10^7 to $10^{10} M_\odot$ and for binary mass ratios of $q = 1.0$ and $q = 0.1$.

3. If emission due to process (II) is low, or stable, while process (I) is in its high state, then we observe a strongly periodic signal with little noise, as seen in epoch 3.

5. Implications for Gravitational-wave Emission

The characteristic GW strain produced by a circular SMBHB, averaged over binary inclination, is (Aggarwal et al. 2019; Arzoumanian et al. 2021)

$$h_{\text{GW}} = \frac{2(G\mathcal{M}_\oplus)^{5/3}(\pi f_{\text{GW}\oplus})^{2/3}}{c^4 D_L}, \quad (10)$$

where \mathcal{M}_\oplus is the “chirp mass” in the observer frame, D_L is the luminosity distance, $f_{\text{GW}\oplus}$ is the GW frequency as observed on Earth, and G and c are the standard physical constants. The observed chirp mass \mathcal{M}_\oplus is related to the rest-frame chirp mass \mathcal{M} and individual masses M_1, M_2 by Maggiore (2008),

$$\mathcal{M} = \frac{(M_1 M_2)^{3/5}}{(M_1 + M_2)^{1/5}} = \frac{\mathcal{M}_\oplus}{1+z}. \quad (11)$$

The observed GW frequency $f_{\text{GW}\oplus}$ and the observed orbital period of the binary P_\oplus are related as $f_{\text{GW}\oplus} = 2/P_\oplus$. Adopting values appropriate for PKS 2131–021, namely, $D_L = 9.08$ Gpc and $f_{\text{GW}\oplus} = 1.3 \times 10^{-8}$ Hz, and assuming that the binary is “face on,” we find

$$h_{\text{GW}} = 1.5 \times 10^{-16} (\mathcal{M}/10^9 M_\odot)^{5/3}, \quad (12)$$

while the rest-frame separation of the binary is

$$r = \left(\frac{Gm}{4\pi^2}\right)^{1/3} \left(\frac{P_\oplus}{1+z}\right)^{2/3} \simeq 9.6 \left(\frac{m}{10^9 M_\odot}\right)^{1/3} \text{ lt-day}, \quad (13)$$

where m is the total rest-frame mass. Thus, for $m \sim 3 \times 10^6 M_\odot - 3 \times 10^9 M_\odot$, we have $r \sim 0.001\text{--}0.01$ pc.

The PKS 2131–021 period is well matched to the sensitivity window of current GW searches with pulsar timing arrays, and its sky location is almost optimal for the NANOGrav array, which reported a 95% Bayesian upper limit on characteristic strain $h_{\text{GW}} \lesssim 2.25 \times 10^{-16}$ (Aggarwal et al. 2019). This translates to a limit $\mathcal{M} \lesssim 4 \times 10^9 M_\odot$ on rest-frame chirp mass, or $\lesssim 9 \times 10^9 M_\odot$ on rest-frame total mass for equal-mass components (see also Figure 12).

If PKS 2131–021 is indeed an SMBHB with such a high chirp mass, it will be within detection range for the joint data set of the International Pulsar Timing Array, enhanced with pulsars found in planned surveys in the first few years of the Square Kilometer Array (Xin et al. 2021). However, it is interesting to consider the likelihood that we would happen to observe an SMBHB in this particular stage of its evolution. GW strain and (observed) time to coalescence t_\oplus are related by

$$h_{\text{GW}} = \frac{5 c P_\oplus^2}{512 \pi^2 t_\oplus D_L}, \quad (14)$$

independent of chirp mass, as can be derived, for instance, using the equations in Maggiore (2008). Evaluating Equation (14) for PKS 2131–021, we obtain

$$h_{\text{GW}} \simeq 8 \times 10^{-19} (t_\oplus/\text{Myr})^{-1}; \quad (15)$$

equating expressions (12) and (15) yields $t_\oplus \sim 5000$ yr for $\mathcal{M} = 10^9 M_\odot$, which implies that the higher chirp masses allowed by GW limits would require implausible observational serendipity. However, this is mitigated by the fact that PKS 2131–021 is by far the best SMBHB candidate we have found in the sample of ~ 1830 blazars that we have been monitoring for 13 yr at the OVRO. In this regard we note that Holgado et al. (2018) estimate, from mock population studies based on the luminosity functions for BL Lacertae objects and flat-spectrum radio quasars with redshifts $z \leq 2$, that a fraction $\leq 10^{-3}$ of blazars host a binary with an orbital period $P < 5$ yr, which is not inconsistent with our statistic here of one strong SMBHB candidate out of a sample of ~ 1800 blazars. It should be clear that adding more years of OVRO data is essential for identifying other examples that may have longer periods or that may have been in a phase of not showing sinusoidal variations during the 13 yr OVRO time window.

6. Discussion

We have shown that PKS 2131–021 exhibits unique periodicity behavior over a 45.1 yr observing span: two epochs of periodic emission, separated by 20 yr, agree in both period and phase. This periodic feature in the PKS 2131–021 light curve is significant at the 4.6σ level. This is strongly suggestive of an SMBHB. We have also shown that, due to red noise, small $\sim 10\%$ variations in the periodicity are easily explained

and are, in fact, to be expected in the light curves of blazars. The form of the periodic variability in epochs 1 and 3 is unexpectedly sinusoidal in character, and this must be telling us something important about the source. We have demonstrated in a simple model that this can be accounted for by the Doppler effect of orbital motion in an SMBHB blazar due to the strong effect of the orbital motion on the Doppler factor of the relativistic jet.

Occam's Razor: While we have not yet proven definitively that PKS 2131–021 is an SMBHB, we believe that this is by far the most likely scenario. We know that PKS 2131–021 has a highly relativistic jet oriented close to the line of sight, which exhibits random flux density variations on timescales of months to years. If it is an SMBHB, then we also know that the SMBH generating the jet is in orbital motion. We have shown that this will produce a sinusoidal modulation of the observed flux density of approximately the magnitude observed. It is also highly likely that features appear and disappear in the light curves or blazars as a result of fueling of the central engine. Therefore, the economy of requiring no additional assumptions, other than that of being an SMBHB, to explain all the observations in PKS 2131–021 satisfies Occam's razor.

If PKS 2131–021 is indeed an SMBHB, it is of interest for the constraints it could provide on models of SMBHB merger progenitors (Shen et al. 2021). If sufficiently massive, which appears unlikely, this SMBHB would be a future candidate for GW detection with pulsar timing arrays, which can already provide an upper limit on its chirp mass. Its confirmation as a binary would therefore be consequential, and it can be expected within the next 5–10 yr from continuing light-curve observations with the 40 m Telescope of the OVRO, provided that the periodic variations continue as they have over the past 16 yr.

By far the most important conclusions of this paper are as follows: (i) regardless of whether or not PKS 2131–021 is an SMBHB, sinusoidal flux density fluctuations, due to orbital motion, and the inevitable apparent variations in period caused by red noise are observational properties of SMBHBs with relativistic jets that clearly should be anticipated; (ii) we should expect to see gaps in the sinusoidal variations, possibly lasting decades; and (iii) in this field the observed phenomenology must lead, rather than the theory, because the detailed models that are required cannot be predicted by theory. Thus, this object provides an extremely useful blueprint for the analysis of SMBHB candidate light curves, including accounting for the effects of random variability that is not associated with the periodicity of interest. We hope that it provides a useful demonstration of the complementary advantages of the GLS, WWZ, and least-squares sine-wave fitting analyses for the investigation of light curves in the pursuit of SMBHBs. We trust that these legacy results from Haystack, UMRAO, and OVRO have convinced the reader of the importance of long-term monitoring in astronomy and its potential for making discoveries.

We thank the anonymous referee for carrying out a very detailed and thorough review of the original manuscript and for many insightful suggestions and questions, which have greatly improved this paper. We thank the California Institute of Technology and the Max Planck Institute for Radio Astronomy for supporting the OVRO 40 m program under extremely difficult circumstances over the last 5 yr in the absence of agency funding. Without this private support this paper would

not have been written. We also thank all the volunteers who have enabled this work to be carried out. Prior to 2016, the OVRO program was supported by NASA grants NNG06GG1G, NNX08AW31G, NNX11A043G, and NNX13AQ89G from 2006 to 2016 and NSF grants AST-0808050 and AST-1109911 from 2008 to 2014. The UMRAO program received support from NSF grants AST-8021250, AST-8301234, AST-8501093, AST-8815678, AST-9120224, AST-9421979, AST-9617032, AST-9900723, AST-0307629, AST-0607523, and earlier NSF awards, and from NASA grants NNX09AU16G, NNX10AP16G, NNX11AO13G, and NNX13AP18G. Additional funding for the operation of UMRAO was provided by the University of Michigan. The NANOGrav project receives support from National Science Foundation (NSF) Physics Frontier Center award No. 1430284. T.H. was supported by the Academy of Finland projects 317383, 320085, and 322535. S.K. acknowledges support from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program under grant agreement No. 771282. W.M. acknowledges support from ANID projects Basal AFB-170002 and FONDECYT 11190853. : R.R. gratefully acknowledges support by the ANID BASAL projects ACE210002 and FB210003, and ANID Fondecyt 1181620. C.O. acknowledges support from the Natural Sciences and Engineering Research Council (NSERC) of Canada. M.V. and T.J.W.L. performed part of this work at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). V.P. acknowledges support from the Foundation of Research and Technology—Hellas Synergy Grants Program through project MagMASim, jointly implemented by the Institute of Astrophysics and the Institute of Applied and Computational Mathematics and by the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the “First Call for H.F.R.I. Research Projects to support Faculty members and Researchers and the procurement of high-cost research equipment grant” (Project 1552 CIRCE). S.O. gratefully acknowledges the support of the Caltech Summer Undergraduate Research Fellowship program. P.VdIP. acknowledges support from Núcleo Milenio TITANs (NCN19-058).

Appendix A

Estimating the Significance of Periodicities in Blazar Light Curves

The flaring behavior of blazars can lead to spurious periodicity “detections.” It is therefore important to make reliable estimates of the significance of any likely detections. Here we suggest an approach for avoiding the known pitfalls of red noise and single-epoch significance tests.

For each epoch we simulate light curves that match the power spectral density (PSD), the pdf, sampling, and observational noise of the specific epoch. Assuming a pure power-law $\text{PSD} \sim \nu^{-\beta}$, where ν is the frequency, we use our own implementation of the method introduced by Uttley et al. (2002) to estimate the following power-law slopes for epochs 1–3: $\beta_{\text{epoch 1}} = 1.71 \pm 0.19$, $\beta_{\text{epoch 2}} = 1.75 \pm 0.26$, $\beta_{\text{epoch 3}} = 1.82 \pm 0.14$, showing no significant difference between the three epochs. An Anderson–Darling test shows that the pdf of each epoch is not Gaussian distributed at significance level $< 10^{-3}$. Therefore, we use the algorithm of Emmanoulopoulos et al. (2013) to create artificial light curves that match the

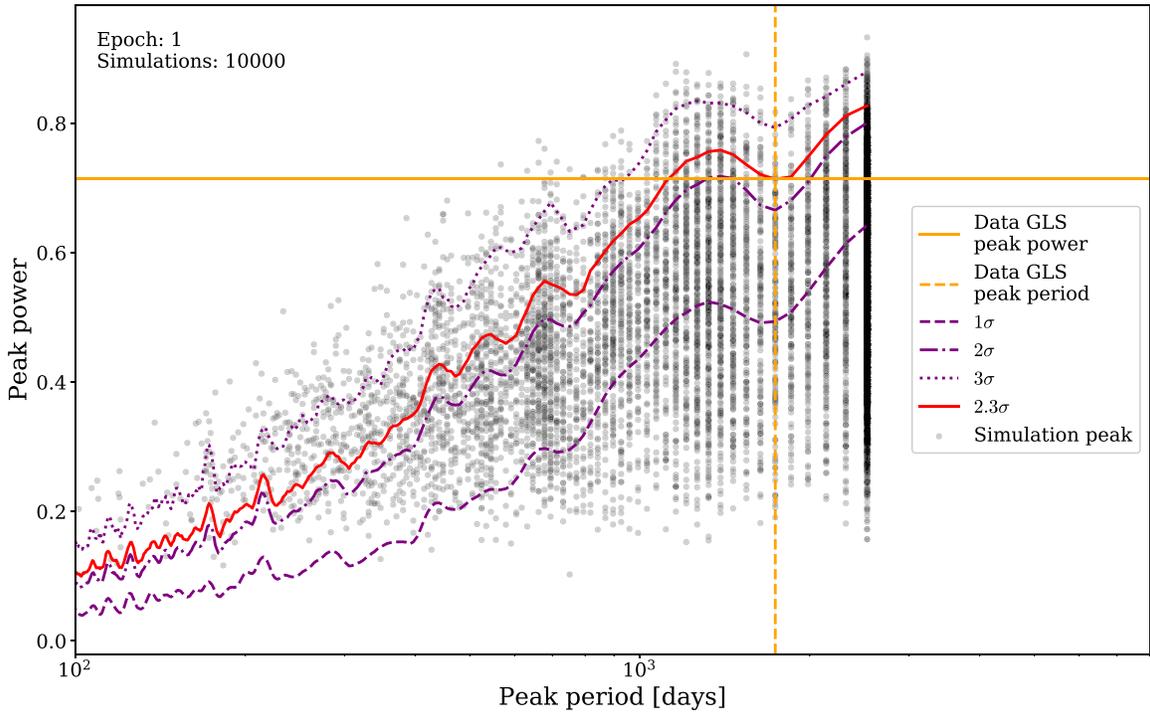


Figure 13. Calculation of the global p -value using epoch 1 as an example. The quantization of the periods is due to the quantization in frequency, $\Delta f = 1/(\zeta T)$, with $\zeta = 10$, which we impose in running the GLS. So the sampled periods are separated by $1/(10\Delta f)$. Step 1: Orange vertical and horizontal lines mark the period and power, respectively, of the strongest peak identified in the GLS of the epoch 1 data. The curves plotted here are merely a visualization tool; they do not represent levels of significance *unless* there are a priori reasons for selecting particular periods. Hence, in the next step (Step 2), the detection of the strongest peak in the GLS of the data at the period P_{peak} is a 2.3σ event only in the case that there is an a priori reason for selecting this particular period. The significance is obtained by summing the simulations with p -values less than that observed at this particular period. The red line connects the 2.3σ significance level at each period. The purple curves indicate the 1σ , 2σ , and 3σ significance levels at each specific period, as also shown in the top panel of Figure 9, which are misleading unless there is an a priori reason for selecting a specific period. Step 3: For each simulation the strongest peak is identified in the GLS. The peak periods and powers are shown as gray circles. Step 4: To estimate the global p -value, simulations at all periods are counted in which the power of the strongest peak lies above the 2.3σ significance level (red line).

estimated PSDs and pdf's. The initial simulations are sampled such that low- and high-frequency power is included beyond the data ranges. The simulations are then resampled to the time stamps of the data and Gaussian noise is added based on the estimated flux density uncertainties.

The *single-period p -value* is the probability of finding a periodic signal at the observed frequency and with the same or greater power under the null hypothesis that the signal is the result of a red-noise process with the same PSD, pdf, sampling, and observational noise as the real data. We use the following procedure—illustrated for epoch 1 in Figure 13—to estimate the *global p -value*, i.e., the probability of finding a periodic signal under the above null hypothesis as significant as the one observed at any frequency:

1. We calculate the GLS periodogram of the data, select the strongest peak, and measure its peak period, P_{peak} , and power, $\mathcal{P}_{\text{peak}}$.
2. At period P_{peak} we count all simulations with power $\mathcal{P}_{\text{sim}} \geq \mathcal{P}_{\text{peak}}$ to estimate the single-period p -value, p_{peak} .
3. For each simulation,
 - (a) we calculate the GLS periodogram at the same discrete frequencies as for the data, and we select the strongest peak and measure its peak period, P_{sim} , and power, \mathcal{P}_{sim} .
 - (b) At P_{sim} we calculate the single-period p -value, p_{sim} .
 - (c) We count all simulations for which $p_{\text{sim}} \leq p_{\text{peak}}$ to estimate the global p -value.

As explained in Section 3.4.1, the single-period p -value is commonly used to express the significance of a detected periodicity. However, it is an incorrect and misleading estimate of the significance, unless this period has been selected a priori, since it does not take into account the fact that spurious periodicities from a red-noise process could arise at any period, and not just the one detected.

We report the global p -values estimated for epochs 1–3 in Table 2 as Tests 1.1, 1.2, and 1.3. Test 1.3 is used in the additional significance estimates, starting from this highly significant detection, described in Section 3.4.

Appendix B The WWZ Transform

Following Foster (1996), the WWZ transform is defined here as

$$\text{WWZ} = \frac{(N_{\text{eff}} - 3)V_y}{2(V_x - V_y)}, \quad (\text{B1})$$

where N_{eff} is the effective number of data points and V_x and V_y are the weighted variances of the data and the model, respectively, as defined by Foster (1996). For PKS 2131–021, we found that the c -value of 0.125 proposed by Foster was a good choice, since other values degraded either the resolution of the time or the period. The density of data points, and hence N_{eff} , varies strongly between the three epochs

(see Figure 1). Also, as can be seen in Equation (B1), when the data and model variances approach each other, which happens when the data approach a sinusoidal waveform, the denominator approaches zero and hence the WWZ power increases. This is why the WWZ magnitude in epoch 3 seen in the top panel of Figure 8 reaches a high value (1790). This is $\sim 20\text{--}30\times$ higher than the peak WWZ power in epoch 1. The signal in epoch 1 is noisier, with higher variance, and therefore the WWZ power is considerably smaller than in epoch 3. However, although the difference in power for the two epochs is large, the peak in epoch 1 is highly relevant since its period agrees with that of epoch 3 to within $\sim 2\%$, and in addition, as we have learned from the least-squares sine-wave fitting, the phase of the epoch 1 periodicity agrees, to within 10% of the periodicity period, with the extrapolated phase from epoch 3.

ORCID iDs

S. O'Neill  <https://orcid.org/0000-0003-2454-6828>
 S. Kiehlmann  <https://orcid.org/0000-0001-6314-9177>
 A. C. S. Readhead  <https://orcid.org/0000-0001-9152-961X>
 M. F. Aller  <https://orcid.org/0000-0003-2483-2103>
 R. D. Blandford  <https://orcid.org/0000-0002-1854-5506>
 I. Liodakis  <https://orcid.org/0000-0001-9200-4006>
 M. L. Lister  <https://orcid.org/0000-0003-1315-3412>
 P. Mróz  <https://orcid.org/0000-0001-7016-1692>
 C. P. O'Dea  <https://orcid.org/0000-0001-6421-054X>
 T. J. Pearson  <https://orcid.org/0000-0001-5213-6231>
 V. Ravi  <https://orcid.org/0000-0002-7252-5485>
 M. Vallisneri  <https://orcid.org/0000-0002-4162-0033>
 K. A. Cleary  <https://orcid.org/0000-0002-8214-8265>
 M. J. Graham  <https://orcid.org/0000-0002-3168-0139>
 K. J. B. Grainge  <https://orcid.org/0000-0002-6780-1406>
 T. Hovatta  <https://orcid.org/0000-0002-2024-8199>
 A. Lähteenmäki  <https://orcid.org/0000-0002-0393-0647>
 W. Max-Moerbeck  <https://orcid.org/0000-0002-5491-5244>
 V. Pavlidou  <https://orcid.org/0000-0002-0870-1368>
 T. A. Prince  <https://orcid.org/0000-0002-8850-3627>
 R. A. Reeves  <https://orcid.org/0000-0001-5704-271X>
 M. Tornikoski  <https://orcid.org/0000-0003-1249-6026>
 P. Vergara de la Parra  <https://orcid.org/0000-0001-5957-1412>
 J. A. Zensus  <https://orcid.org/0000-0001-7470-3321>

References

- Abraham, Z. 2018, *NatAs*, 2, 443
 Aggarwal, K., Arzoumanian, Z., Baker, P. T., et al. 2019, *ApJ*, 880, 116
 Aller, M. F., Aller, H. D., Hughes, P. A., & Latimer, G. E. 1999, *ApJ*, 512, 601
 Aller, H. D., Aller, M. F., Latimer, G. E., & Hodge, P. E. 1985, *ApJS*, 59, 513
 Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2021, *ApJ*, 914, 121
 Baars, J. W. M., Genzel, R., Pauliny-Toth, I. I. K., & Witzel, A. 1977, *A&A*, 500, 135
 Balonek, T. J. 1982, PhD thesis, University of Massachusetts
 Bansal, K., Taylor, G. B., Peck, A. B., Zavala, R. T., & Romani, R. W. 2017, *ApJ*, 843, 14
 Begelman, M. C., Blandford, R. D., & Rees, M. J. 1980, *Natur*, 287, 307
 Britzen, S., Fendt, C., Witzel, G., et al. 2018, *MNRAS*, 478, 3199
 Burke-Spolaor, S., Taylor, S. R., Charisi, M., et al. 2019, *A&ARv*, 27, 5
 Caproni, A., Abraham, Z., Motter, J. C., & Monteiro, H. 2004, *ApJL*, 851, L39
 Caproni, A., Mosquera Cuesta, H. J., & Abraham, Z. 2004, *ApJL*, 616, L99
 Cassaro, P., Stanghellini, C., Dallacasa, D., Bondi, M., & Zappalà, R. A. 2002, *A&A*, 381, 378
 Covino, S., Sandrinelli, A., & Treves, A. 2019, *MNRAS*, 482, 1270
 Dent, W. A., Aller, H. D., & Olsen, E. T. 1974a, *ApJL*, 188, L11
 Dent, W. A., Kapitzky, J. E., & Kojoian, G. 1974b, *AJ*, 79, 1232
 Dey, L., Valtonen, M. J., Gopakumar, A., et al. 2021, *MNRAS*, 503, 4400
 Drake, A. J., Djorgovski, S. G., Mahabal, A., et al. 2009, *ApJ*, 696, 870
 Drinkwater, M. J., Webster, R. L., Francis, P. J., et al. 1997, *MNRAS*, 284, 85
 Emmanoulopoulos, D., McHardy, I. M., & Papadakis, I. E. 2013, *MNRAS*, 433, 907
 Emmanoulopoulos, D., McHardy, I. M., & Uttley, P. 2010, *MNRAS*, 404, 931
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
 Foster, G. 1996, *AJ*, 112, 1709
 Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015a, *MNRAS*, 453, 1562
 Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015b, *Natur*, 518, 74
 Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, *PASP*, 131, 078001
 Haiman, Z., Kocsis, B., & Menou, K. 2009, *ApJ*, 700, 1952
 Holgado, A. M., Sesana, A., Sandrinelli, A., et al. 2018, *MNRAS*, 481, L74
 Homan, D. C., Cohen, M. H., Hovatta, T., et al. 2021, *ApJ*, 923, 67
 Hovatta, T., Lister, M. L., Aller, M. F., et al. 2012, *AJ*, 144, 105
 Kellermann, K. I., Vermeulen, R. C., Zensus, J. A., & Cohen, M. H. 1998, *AJ*, 115, 1295
 Komatsu, E., Dunkley, J., Nolta, M. R., et al. 2009, *ApJS*, 180, 330
 Liodakis, I., Pavlidou, V., Hovatta, T., et al. 2017, *MNRAS*, 467, 4565
 Lister, M. L., Aller, M. F., Aller, H. D., et al. 2018, *ApJS*, 234, 12
 Lister, M. L., & Homan, D. C. 2005, *AJ*, 130, 1389
 Lister, M. L., Homan, D. C., Hovatta, T., et al. 2019, *ApJ*, 874, 43
 Lomb, N. R. 1976, *Ap&SS*, 39, 447
 Maggiore, M. 2008, *Gravitational Waves: Vol. 1: Theory and Experiments*, Gravitational Waves (Oxford: Oxford Univ. Press)
 Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, *PASP*, 131, 018003
 Max-Moerbeck, W., Richards, J. L., Hovatta, T., et al. 2014, *MNRAS*, 445, 437
 O'Dea, C. P., Dent, W. A., Kinzel, W. M., & Balonek, T. J. 1986, *AJ*, 92, 1262
 Owen, F. N., O'Dea, C. P., Inoue, M., & Eilek, J. A. 1985, *ApJL*, 294, L85
 Perley, D. A. 2019, *PASP*, 131, 084503
 Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, *A&A*, 641, A6
 Ragazzini, J. R., & Zadeh, L. A. 1952, *AIEE*, 71, 225
 Readhead, A. C. S., Lawrence, C. R., Myers, S. T., et al. 1989, *ApJ*, 346, 566
 Rector, T. A., & Stocke, J. T. 2001, *AJ*, 122, 565
 Rector, T. A., & Stocke, J. T. 2003, *AJ*, 125, 2447
 Ren, G.-W., Ding, N., Zhang, X., et al. 2021, *MNRAS*, 506, 3791
 Richards, J. L., Max-Moerbeck, W., Pavlidou, V., et al. 2011, *ApJS*, 194, 29
 Rodriguez, C., Taylor, G. B., Zavala, R. T., et al. 2006, *ApJ*, 646, 49
 Sandrinelli, A., Covino, S., Treves, A., et al. 2017, *A&A*, 600, A132
 Sandrinelli, A., Covino, S., Treves, A., et al. 2018, *A&A*, 615, A118
 Sbarufatti, B., Treves, A., Falomo, R., et al. 2006, *AJ*, 132, 1
 Scargle, J. D. 1982, *ApJ*, 263, 835
 Scheuer, P. A. G., & Readhead, A. C. S. 1979, *Natur*, 277, 182
 Schwarzenberg-Czerny, A. 1989, *MNRAS*, 241, 153
 Shen, Y., Chen, Y.-C., Hwang, H.-C., et al. 2021, *NatAs*, 5, 569
 Sillanpää, A., Haarala, S., Valtonen, M. J., Sundelius, B., & Byrd, G. G. 1988, *ApJ*, 325, 628
 Sobacchi, E., Sormani, M. C., & Stameria, A. 2017, *MNRAS*, 465, 161
 Teraesranta, H., Tornikoski, M., Mujunen, A., et al. 1998, *A&AS*, 132, 305
 Uttley, P., McHardy, I. M., & Papadakis, I. E. 2002, *MNRAS*, 332, 231
 Valtonen, M. J., Zola, S., Ciprini, S., et al. 2016, *ApJL*, 819, L37
 Vaughan, S., Uttley, P., Markowitz, A. G., et al. 2016, *MNRAS*, 461, 3145
 Welsh, W. F. 1999, *PASP*, 111, 1347
 Xin, C., Mingarelli, C. M. F., & Hazboun, J. S. 2021, *ApJ*, 915, 97
 Zechmeister, M., & Kürster, M. 2009, *A&A*, 496, 577