The Astrophysical Journal Supplement Series, 285:61 (24pp), 2026 August https://doi.org/10.3847/1538-4365/ae8444

QhX Catalog of Coherent and Quasi-coherent Active Galactic Nucleus Variability from Vera C. Rubin LSST DP1

Andjelka B. Kovačević[1,36], Dragana Ilić[1,2], Uroš Meštrić[1,3], Milena Stanković Cerović[4], Natale De Bonis[1], Luka Č. Popović[1,5], Saša Simić[6,7], Marta Fatović[8,9,10], Lorena Hernandez-García[11,12], Donald P. Schneider[13,14], Roberto J. Assef[11], Swayamtrupta Panda[15], Claudia M. Raiteri[16], Christian Wolf[17], Christopher A. Onken[17], Neven Caplar[18], Maria Charisi[19,20], Claudio Ricci[11,21,22], Chengcheng Xin[23], Ilsang Yoon[24], Mara Salvato[25], Matthew J. Temple[26], Sarath Satheesh-Sheeba[27], Demetra De Cicco[10,28], Maurizio Paolillo[8,10,28], Paula Sánchez-Sáez[29], Angela Bongiorno[10], Weixiang Yu[30], Akke Viitanen[31], Bindu Rani[32], Megan C. Davis[19], Marcin Marculewicz[33], Tomislav Jurkić[34], and Đorđe Savić[5,35]

[1] Department of Astronomy, Faculty of Mathematics, University of Belgrade, Studentski trg 16, 11000 Belgrade, Serbia; andjelka@matf.bg.ac.rs
[2] Hamburger Sternwarte, Universität Hamburg, Gojenbergsweg 112, D-21029 Hamburg, Germany
[3] ICRAR, The University of Western Australia, 35 Stirling Highway, Crawley, WA 6009, Australia
[4] Independent Researcher
[5] Astronomical Observatory, Volgina 7, 11060 Belgrade, Serbia
[6] Faculty of Science, University of Kragujevac, Radoja Domanovića 12, 34000 Kragujevac, Serbia
[7] Center for Astrophysics and Cosmology, University of Nova Gorica, Vipavska 13, 5000 Nova Gorica, Slovenia
[8] Dipartimento di Fisica "Ettore Pancini," Università di Napoli Federico II, Via Cintia 80126, Naples, Italy
[9] Ruđer Bošković Institute, Bijenička cesta 54, 10000 Zagreb, Croatia
[10] INAF—Osservatorio Astronomico di Capodimonte, Via Moiariello 16, 80131 Naples, Italy
[11] Instituto de Estudios Astrofísicos, Facultad de Ingeniería y Ciencias, Universidad Diego Portales, Av. Ejército Libertador 441, Santiago, Chile
[12] Centro Interdisciplinario de Data Science, Facultad de Ingeniería y Ciencias, Universidad Diego Portales, Av. Ejército Libertador 441, Santiago, Chile
[13] Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA
[14] Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA
[15] International Gemini Observatory, NSF/NOIRLab: La Serena, Coquimbo, Chile
[16] INAF—Osservatorio Astrofisico di Torino, Via Osservatorio 20, 10025, Pino Torinese, Italy
[17] Research School of Astronomy and Astrophysics, Australian National University, Canberra, ACT 2611, Australia
[18] Department of Astronomy, University of Washington, Box 351580, Seattle, WA 98195, USA
[19] Institute of Astrophysics, Foundation for Research and Technology–Hellas (FORTH), Vassilika Vouton, GR-70013 Heraklion, Crete, Greece
[20] Department of Physics and Astronomy, Washington State University, Pullman, WA 99164, USA
[21] Department of Astronomy, University of Geneva, ch. d'Ecogia 16, 1290, Versoix, Switzerland
[22] Kavli Institute for Astronomy and Astrophysics, Peking University, Beijing 100871, People's Republic of China
[23] Black Hole Initiative, Harvard & Smithsonian, 20 Garden Street, Cambridge, MA 02138, USA
[24] National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903, USA
[25] Max Planck Institute for Extraterrestrial Physics (MPE) & Excellence Cluster Universe. Giessenbachstr. 1, D-85748 Garching, Germany
[26] Centre for Extragalactic Astronomy, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK
[27] Instituto de Astrofísica, Facultad de Ciencias Exactas, Universidad Andrés Bello, Fernández Concha 700, 7591538 Las Condes, Santiago, Chile
[28] INFN—Sez. di Napoli, Via Cintia 21, I-80126 Napoli, Italy
[29] European Southern Observatory, Karl-Schwarzschild-Strasse 2, 85748 Garching bei München, Germany
[30] Department of Physics & Astronomy, Bishop's University, 2600 rue College, Sherbrooke, QC J1M 1Z7, Canada
[31] Department of Astronomy, University of Geneva, Chemin Pegasi 51, 1290 Versoix, Switzerland
[32] NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA
[33] Department of Physics and Astronomy, Wayne State University, 666 W. Hancock Street, Detroit, MI 48201, USA
[34] University of Rijeka, Faculty of Physics, Radmile Matejcic 2, 51000 Rijeka, Croatia
[35] Institut d'Astrophysique et de Géophysique, Université de Liège, Allée du 6 Août 19c, 4000 Liège, Belgium
Received 2026 March 18; revised 2026 June 22; accepted 2026 June 28; published 2026 August 3

Abstract

Coherent and quasi-coherent variability in active galactic nuclei (AGN) can serve as a useful diagnostic of dynamical processes operating in the immediate environments of supermassive black holes. The time-domain capabilities of the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) enable systematic searches for such signals across large AGN samples. We identify and classify coherent oscillatory processes in AGN cores using LSST Data Preview 1 (DP1), and compile the results in the QhX Hub, an interactive catalog providing statistical variability features and visualization tools. AGN are selected by cross-matching LSST DP1 sources with external catalogs, and light curves are analyzed with the multiband wavelet-based Quasar harmonic eXplorer (`QhX`), designed to detect coherent signals in irregularly sampled, red-noise-dominated data. The catalog contains 230 objects stratified into variability classes. A single strong coherent candidate is identified: 3D-HST 2351, $z = 0.198$, LSST DP1 `object_id` = 611255072642321324, with a period of $P = 2.76^{+2.65}_{-0.06}$ days, where the large upper bound reflects the limited observational baseline. Its coherent nuclear variability is the

[36] Corresponding author.

primary motivation for follow-up. In subsequent LSST data releases, all catalog sources will be recalculated on progressively longer baselines, enabling more robust population-level characterization of coherent AGN variability.

Unified Astronomy Thesaurus concepts: Active galactic nuclei (16); Astroinformatics (78); Sky surveys (1464); Catalogs (205); Astrostatistics (1882); Light curves (918); Period search (1955); AB photometry (2168); Time series analysis (1916)

1. Introduction

The photometric variability of active galactic nuclei (AGN)[37] is commonly modeled as a stochastic process dominated by correlated (red) noise (S. Vaughan et al. 2003). On timescales from several months to a few years, this variability is often described using damped random walk (DRW) models (e.g., P. Sánchez-Sáez et al. 2018; K. L. Suberlak et al. 2021) or more flexible continuous-time autoregressive processes (B. C. Kelly et al. 2009; C. L. MacLeod et al. 2010; S. Kozlowski 2016; J. Moreno et al. 2019; W. Yu et al. 2022; W. Yu et al. 2025). However, deviations from simple DRW behavior have been reported on timescales shorter than a few months (V. P. Kasliwal et al. 2015; Z. Stone et al. 2022). More generally, quasar variability is often better described by higher-order autoregressive or continuous-time autoregressive moving-average models, with the preferred stochastic description likely differing from source to source (B. C. Kelly et al. 2014; E. D. Feigelson et al. 2018; Z. Stone et al. 2022; J. Xu et al. 2025). In such stochastic systems, oscillation-like or quasiperiodic-looking behavior may arise without an underlying deterministic clock (T. J. Ulrych & T. N. Bishop 1975; A. V. Vecchia 1985). Accordingly, throughout this paper, the term "period" is used operationally to denote a recovered characteristic oscillation timescale, and not necessarily a deterministic period in the classical sense of stationary sinusoidal modulation. Quasar light curves may also exhibit fractal-like scaling and stochastic self-similarity (R. Vio et al. 1991; A. B. Belete et al. 2018). Recently, noise-driven damped harmonic oscillator (DHO) models have been shown to provide an improved statistical description of quasar variability, while also revealing correlations between variability parameters and physical properties such as black hole mass (M_{BH}) and Eddington ratio λ_{Edd} (see, e.g., M. Paolillo & I. Papadakis 2025; W. Yu et al. 2025).

Observationally, temporal brightness correlations have been reported in a subset of quasars, including phase-stable modulations, quasiperiodic oscillations (QPOs), and, in rare cases, candidate periodic signals (M. Gierlinski et al. 2008; M. J. Graham et al. 2015; M. Charisi et al. 2016; T. Liu et al. 2018; A. De Rosa et al. 2019; Y.-C. Chen et al. 2020, 2023; C. A. Witt et al. 2022; D. J. D'Orazio & M. Charisi 2024; L. Hernández-García et al. 2024). At the same time, red-noise fluctuations, particularly when nonstationary, can generate transient oscillatory features that persist for several cycles, leading classical peak-based methods to misidentify such patterns as genuine periodicity (B. C. Kelly et al. 2009; S. Vaughan et al. 2016; B. J. Kelly et al. 2017; J. Robnik et al. 2024; K. El-Badry et al. 2026). This ambiguity is especially important because autoregressive stochastic processes can themselves produce oscillation-like behavior over several cycles without any uniquely defined deterministic driver. In this regime, the central question is not only whether a light curve shows excess power at a characteristic timescale, but also whether the signal exhibits persistent multiband structure and temporal coherence beyond what is expected from correlated noise alone. This has motivated detailed investigations of the limitations of linear time–frequency techniques and the development of new approaches aimed at isolating physically meaningful, coherence-rich oscillatory behavior in quasar light curves (A. B. Kovačević et al. 2018, 2019, 2020b; Y.-J. Chen et al. 2023; M. Fatović et al. 2023; M. C. Davis et al. 2024; K. Park et al. 2025; F. Rigamonti et al. 2025; X. Zhang 2025; L. Bertassi et al. 2026; A. Lin et al. 2026).

At the same time, there are growing indications that at least part of quasar variability may arise from dynamical processes operating within the accretion flow. In particular, eccentric accretion disks can introduce modulation on characteristic orbital and precessional timescales (e.g., M. Eracleous et al. 1995; A. Pancoast et al. 2014; T. Storchi-Bergmann et al. 2017; C.-H. Chan et al. 2024). Within this framework, variability can originate from spatially and dynamically distinct emitting regions along eccentric streams or precessing inner flows, such that the resulting light curves need not be sinusoidal or strictly periodic (e.g., H. Deng 2026). Simulations of eccentric and torn accretion disks show that radial oscillations in the inner regions can generate quasiperiodic luminosity variations at a few-percent level, with amplitudes and characteristic timescales that evolve due to periapsis shifts, shocks, and relativistic precessional dynamics (G. Musoke et al. 2022; N. Kaaz et al. 2026). In particular, N. Kaaz et al. (2026) showed that even intraday quasiperiodic oscillatory signatures in optical and ultraviolet light curves may arise from torn disk configurations. More broadly, coherence-rich variability may also originate from disk instabilities, warped or precessing accretion flows, Lense–Thirring precession, or interactions in circumbinary black hole systems (S. Kato 2001; A. Ingram et al. 2009; L. Z. Kelley et al. 2019; L. Č. Popović et al. 2021; Y.-Y. Songsheng & J.-M. Wang 2023; J.-M. Wang et al. 2024; M. Fatović et al. 2025; C. Xin et al. 2026). These processes do not usually produce perfect sinusoids (e.g., K. El-Badry et al. 2026), but instead generate dynamically evolving oscillatory signatures with varying phase and amplitude structures. Among the possible interpretations of such signals are subparsec binary SMBH systems, which are of particular interest as potential nHz gravitational-wave sources (e.g., C. Xin & Z. Haiman 2024).

A key distinction in this context is the concept of *temporal coherence*, which refers to the persistence of phase structure and harmonic organization across time, even when frequency or amplitude evolve.[38] Unlike classical periodogram-based approaches that assume stationarity and detect excess power at fixed frequencies (P. Abry et al. 1995; L. Cohen 1995), coherence-based analyses seek oscillatory behavior that retains

[37] Hereafter, we use the terms AGN and quasars interchangeably.

[38] Coherent variability is usually called "autocorrelated" or "autoregressive" variability in the time-series literature.

phase stability and structural persistence under nonstationary conditions (W. H. Press et al. 1992; S. Vaughan et al. 2016). Our aim in this work is therefore not to assign a unique stochastic generative model or to claim a definitive physical interpretation for each light curve. Instead, we demonstrate that a coherence-based catalog can identify and characterize structured oscillatory behavior with sufficient multiband consistency and statistical robustness in data from modern large-scale observing programs such as the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST; Ž. Ivezić et al. 2019).

The LSST is uniquely suited to investigate coherence-rich variability owing to its wide-area, cadence, and multiband monitoring capabilities (Ž. Ivezić et al. 2019; F. B. Bianco et al. 2022). The LSST Data Preview 1 (DP1, NSF-DOE Vera C. Rubin Observatory 2025) provides the first commissioning-quality light curves, capturing observational complexities such as irregular sampling, photometric noise, seasonal gaps, and filter heterogeneity that cannot be fully reproduced in simulations (J. L. Carlin et al. 2025; J. Freeburn et al. 2025; G. Li et al. 2026; W. Yu et al. 2026). Despite its limited temporal and spatial baseline, with only seven pointings observed over a few months, DP1 serves as a crucial testbed for assessing whether coherent variability signatures can be extracted under realistic survey conditions.

Motivated by these challenges, we apply the `Quasar Harmonic eXplorer` (`QhX`; A. B. Kovačević et al. 2018, 2020a, 2026; M. Fatović et al. 2023, 2025; A. Kovacevic 2024), a framework designed to identify coherence-rich oscillatory behavior in quasar light curves through time–frequency and cross-correlation analyses, to the DP1 data. Rather than relying on isolated spectral peaks, `QhX` compares time–frequency representations across filters and temporal segments to distinguish physically driven oscillations from stochastic red-noise variability.

This paper is organized as follows. Section 2 describes the LSST DP1 dataset, AGN sample selection via cross-matching with external catalogs, and light-curve preparation. Section 3 presents the catalog construction methodology. Section 4 overviews the resulting prototype catalog and illustrates how it can be used for population-level studies of accretion flows in AGN. Section 5 outlines future improvements and prospects enabled by upcoming LSST data releases. For all cosmological calculations, we adopt a flat ΛCDM cosmology consistent with that of the Planck Collaboration et al. (2020), with $H_0 = 67.66\ \mathrm{km\ s^{-1}\ Mpc^{-1}}$, $\Omega_{\mathrm{m}} = 0.3111$, and $\Omega_\Lambda = 0.6889$.

2. Data

DP1 represents the first public dataset processed with the Rubin Science Pipelines and provides commissioning-era observations reduced in a manner consistent with survey operations (NSF-DOE Vera C. Rubin Observatory 2025). The release contains data from seven distinct observing fields, covering a total sky area of ∼15 deg^2. The observations were obtained with the LSST Commissioning Camera between 2024 October 24 and December 11. In total, 1791 individual exposures are available, each acquired in one of the six LSST filters (*u*, *g*, *r*, *i*, *z*, and *y*). The Rubin/LSST photometric system is not identical to the Sloan Digital Sky Survey (SDSS) filter set. Although the band nomenclature is similar, Rubin employs six broad bands, *u, g, r, i, z,* and *y*, rather than the five SDSS bands, *u, g, r, i,* and *z*, and the corresponding throughput curves differ significantly, particularly at the blue and red ends of the spectral range. For reference, the approximate wavelength coverage of the Rubin bands is *u* (338–395 nm), *g* (405–552 nm), *r* (553–690 nm), *i* (690–817 nm), *z* (818–920 nm), and *y* (920–1010 nm). Each pointing covers a $40' \times 40'$ field of view, enabling multiband sampling of astrophysical sources under LSST observing conditions (see Vera C. Rubin Observatory Team et al. 2026).

The cadence of DP1 differs substantially from that of the main LSST survey (F. B. Bianco et al. 2022). The Wide–Fast–Deep (WFD) survey is the core 10 yr LSST program, designed to cover ∼18,000–20,000 deg^2 of the southern sky in six filters (*u, g, r, i, z,* and *y*) with ∼30 s visits and a typical revisit cadence of ∼3–4 days, optimized for uniform wide-area mapping and population-level science. By contrast, the Deep Drilling Fields (DDFs) target five predefined regions of roughly ∼10 deg^2 each with much higher cadence, often daily or subdaily with multiple visits per night, and deeper cumulative exposure, reaching at least ∼1 mag deeper than WFD[39] (Ž. Ivezić et al. 2019).

The DP1 observing strategy included multiple observations of the same field within a single night, often in up to three filters. This enhanced temporal sampling permits exploration of short-timescale variability (see also J. Freeburn et al. 2025; K. Malanchev et al. 2025) and enables multiband quasar periodicity searches on timescales as short as days. Although the total DP1 baseline is limited to lasting a total of 7 weeks, the high intranight cadence provides a useful testbed for evaluating algorithms designed to detect short-period signals. We therefore use data from all seven fields (see Table 1).

2.1. *Compilation of Known AGN Detected in LSST DP1*

To construct a sample of quasars detected in the LSST DP1 fields, we cross-matched DP1[40] object catalogs accessed through `LSDB` (N. Caplar et al. 2025; K. Malanchev et al. 2025) with several major publicly available AGN catalogs from complementary surveys: the Gaia mission quasar catalog Quaia (K. Storey-Fisher et al. 2024), the Dark Energy Spectroscopic Instrument (DESI DR1; DESI Collaboration et al. 2024; DESI Collaboration et al. 2026), the Zwicky Transient Facility 5 Million Quasars catalog (QZO; F. J. Masci et al. 2019; S. J. Nakoneczny et al. 2025), and the Million Quasars catalog (MilliQ; E. W. Flesch 2023). We selected sources classified as AGN in these external cross-matched catalogs.

To associate LSST DP1 object IDs[41] with sources from the external catalogs, we adopted a matching radius of $0\overset{\prime\prime}{.}1$, consistent with the centroiding accuracy expected from the LSST DP1 pipelines.[42] No additional selection criteria were imposed, except for DESI, for which only sources flagged as `SPECTYPE=QSO` were retained prior to matching.

We identified 1456 unique quasars from the combined literature sample within the LSST DP1 fields. The same result

39 https://survey-strategy.lsst.io/baseline/index.html

40 https://dp1.lsst.io/overview/iqsummary.html

41 For readability, we quote only the numerical portion of the official LSST DP1 `objectId`. This identifier is a 64-bit integer constructed by the Rubin Science Pipelines `IdGenerator` class by encoding sky map quantities such as tract and patch together with a per-table counter, https://pipelines.lsst.io/index.html; it is not a coordinate-based designation. Explicit J2000 coordinates for all candidates are provided in the text.

42 https://dp1.lsst.io/tutorials/notebook/306/notebook-306-3.html

Table 1
Summary of LSST DP1 Fields and Number of Quasars Analyzed in This Work, Extracted by Cross-matching with Other External Catalogs

Field Name	R.A., Decl. (deg)	Epochs	Mean Visits (per Epoch)	Number of Quasars
47 Tuc Globular Cluster	6.02, −72.08	4	16.5	14
Low Ecliptic Latitude Field	37.86, +6.98	5	31.8	683
Fornax Dwarf Spheroidal Galaxy	40.00, −34.45	2	21.0	44
Extended Chandra Deep Field South (ECDFS)	53.13, −28.10	21	40.7	602
Euclid Deep Field South (EDFS)	59.10, −48.73	9	30.2	53
Low Galactic Latitude Field	95.00, −25.00	10	29.2	48
Seagull Nebula	106.23, −10.51	4	25.0	12

Note. For each field the columns are: coordinates for epoch J2000, number of observing nights (epochs), number of mean visits per epoch, and number of detected quasars.

was independently recovered using the Tool for Operations on Catalogues and Tables (TOPCAT; M. B. Taylor 2005). The spatial distribution of AGN uniquely identified in the different catalogs is shown in Figure 1. The top panel shows the selection geometry of the cross-matched LSST DP1 AGN relative to the larger external catalogs, while the lower panels highlight field-to-field differences in the distribution of cross-matched AGN and their redshift content. The main properties of the seven DP1 fields are summarized in Table 1. The fields also differ in temporal sampling: ECDFS received the most uniform near-nightly cadence, Rubin SV 38 7 was revisited with a structured dither pattern, and fields such as Seagull and Fornax were sampled more sparsely across filters and nights.

Table 2 summarizes, for each external catalog, the number of matched quasars, the number uniquely contributed by that catalog, and the overlap with the others. The source overlap is illustrated further by the Venn diagram in Figure 2. Multi-epoch photometry was then extracted from the DP1 object catalog in all available filters for each of the 1456 quasars.

The resulting sample spans a redshift range of $0.01 < z < 5.33$ (Figure 3). Most redshifts are spectroscopically confirmed (1417 sources), while 21 sources have photometric redshifts. For sources present in multiple catalogs, the adopted redshift priority was DESI, MilliQ, Quaia, and QZO. Of the 1456 quasars, 18 objects, all from the QZO catalog, have no reported redshift.

Figure 4 illustrates the redshift–luminosity distribution of the LSST DP1 quasar sample analyzed in this work. The bolometric luminosities, derived from the continuum luminosity L_{3000} (see Section 4.4), span $\sim 10^{42}$–10^{48} erg s^{-1} and fall within the expected range of known quasar populations (e.g., I. Pâris et al. 2018). At $z \lesssim 2$, the sample covers several orders of magnitude in luminosity, whereas at higher redshifts it becomes increasingly dominated by intrinsically luminous quasars, reflecting the flux-limited nature of the parent catalogs. Although sampling was evaluated across all six LSST bands, u, g, r, i, z, and y, only the g, r, i, and z bands met the minimum threshold of 50 epochs per source. The working sample was therefore defined by this four-band subset. In the luminosity–redshift plane, the 50+ epoch subset overlaps the parent sample but is preferentially concentrated toward lower redshift and, correspondingly, lower luminosity.

2.2. Selection of Light Curves

From the cross-matched sample of 1456 objects (Section 2.1), we searched for sources with at least 50 epochs per band. We found that 428 objects satisfied this requirement in the g, r, i, and z bands, and we refer to this subset as the four-band DP1 sample. This subset forms the basis of our analysis. The threshold of $\geqslant$50 points (see Table 2) is more conservative than the minimum adopted in previous variability-based studies (e.g., C. J. Burke et al. 2021; D. De Cicco et al. 2025, 2026) and is intended to ensure adequate sampling of short- to intermediate-timescale periodicities. It also ensures sufficient per-band sampling to evaluate the vetting metrics described in Section 3.1.

The final working sample consists of 428 quasar light curves satisfying the criteria above. While future regular LSST Data Releases are expected to provide a suite of precomputed time-series features (M. Jurić et al. 2023), these are not yet available for DP1. We therefore compute a baseline set of 54 model-independent light-curve features using the `Feature Extractor for Time Series (FEETS)`[43] package (J. Cabral et al. 2018). Such features have been widely used in variability studies of astronomical light curves (e.g., P. Sánchez-Sáez et al. 2021; M. Pavez-Herrera et al. 2025).

For illustration, Figure 5 shows a subset of features calculated with the `FEETS` package that summarize key aspects of quasar variability: (i) variability amplitude (J. W. Richards et al. 2011) and mean-variance ratios (I. Nun et al. 2015), (ii) smoothness and short-timescale structure, as quantified by the von Neumann ratio (D.-W. Kim et al. 2014), and (iii) temporal sampling properties such as the effective baseline and interepoch separations across filters. These compact, survey-agnostic descriptors provide a useful benchmark for contextualizing the coherence-rich variability later detected by `QhX` within the broader variability behavior of the DP1 sample. In all panels, the vertical dashed line marks the value of the corresponding feature for the strong coherent candidate in that band (see Section 3 for the candidate classification).

The top panel shows the distribution of the percentile-based amplitude computed by `FEETS`, defined as the difference between the 95th and 5th percentile magnitudes normalized by the median magnitude. This robust estimator indicates that typical month-scale AGN variability is at the few-percent level.

The second panel of Figure 5 shows the mean variance. Most objects lie between 10^{-3} and 10^{-2} mag^2, with a tail extending down to $\sim 10^{-5}$ mag^2 that likely reflects noise-dominated light curves. The modal variance decreases gradually from g to z, consistent with the known wavelength

[43] The full list of features is available at https://feets.readthedocs.io/en/latest/features.html.

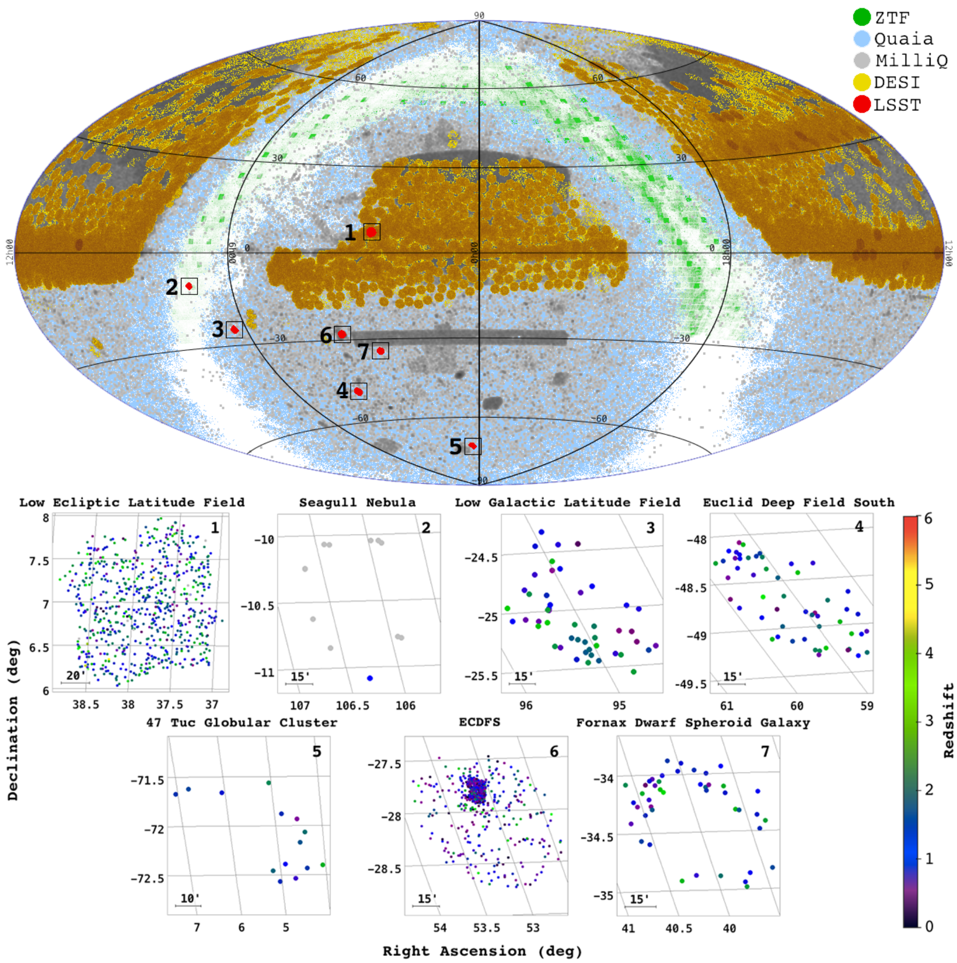

Figure 1. LSST DP1 sky map overplotted with the distribution of AGN from external catalogs that are cross-matched with LSST DP1 sources in seven fields (red circles with labels 1 to 7), shown in equatorial coordinates (J2000.0). The top panel illustrates the selection geometry of the analyzed sample with respect to the external AGN catalogs used in the crossmatch, highlighting the restricted DP1 footprint. The lower panels are insets of these fields, including the field-level distribution of the cross-matched AGN, color coded by redshift; sources without available redshift measurements are shown in gray.

Table 2
Number of Quasars Per Catalog, Extracted After Cross-matching LSST DP1 with Other External Catalogs

Catalog	Number of Quasars	Number of Unique Quasars
DESI	622	475
Quaia	333	140
MilliQ	646	546
ZTF	149	72
Overlapping	⋯	223
Total	1750	1456

dependence of variability strength (D. E. Vanden Berk et al. 2004).

The third panel shows the von Neumann ratio η, which quantifies smoothness and short-timescale randomness in the light curves. All filters show broad peaks near $\eta \sim 1$–2, consistent with stochastic but temporally correlated variability. The distributions extend from $\eta \approx 0.2$ to $\eta \approx 3$, with larger values corresponding to noisier or more sparsely sampled curves. Given the short DP1 baseline of only ~32 days, most sources do not sample variability over a sufficiently long

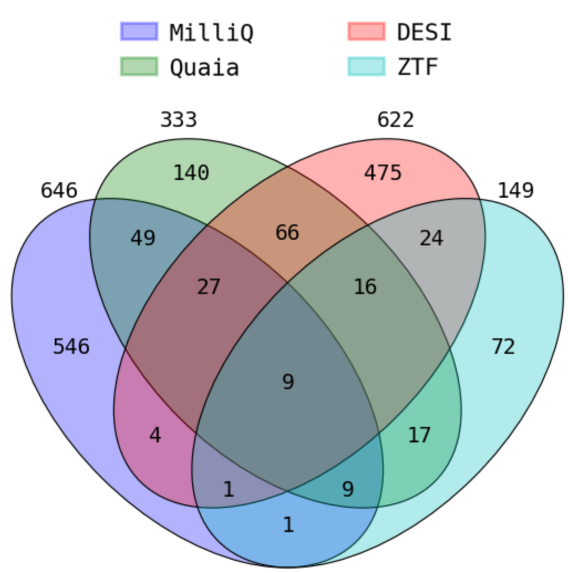

Figure 2. Venn diagram of the combined LSST DP1 sample of known AGN obtained by cross-matching LSST DP1 with the MilliQ, Quaia, DESI, and ZTF catalogs. Values indicate unique and intersected identifications across the four catalogs.

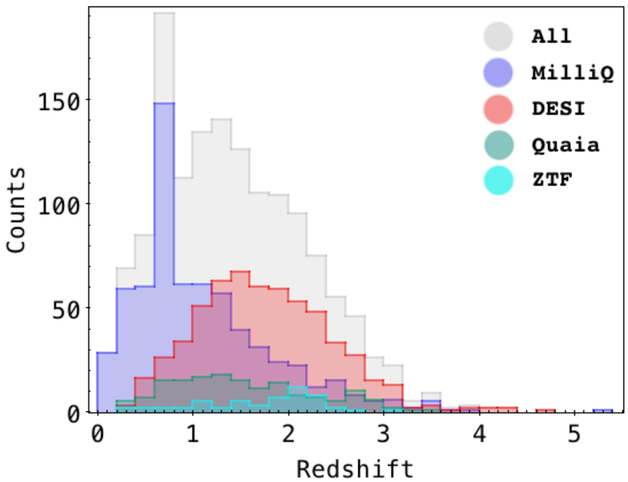

Figure 3. The redshift distribution of the extracted AGN sample includes 1438 out of 1456 sources contributing to the presented histogram. This sample resulted from cross-matching LSST DP1 data with publicly available external AGN catalogs. The gray histogram represents the entire sample, while the colored histograms show the redshift distributions of subsamples based on the external catalogs from which the AGN were extracted. The LSST DP1 AGN cover $0.01 < z < 5.33$. The vast majority of selected objects (1417) have spectroscopic and only 21 have photometric redshifts.

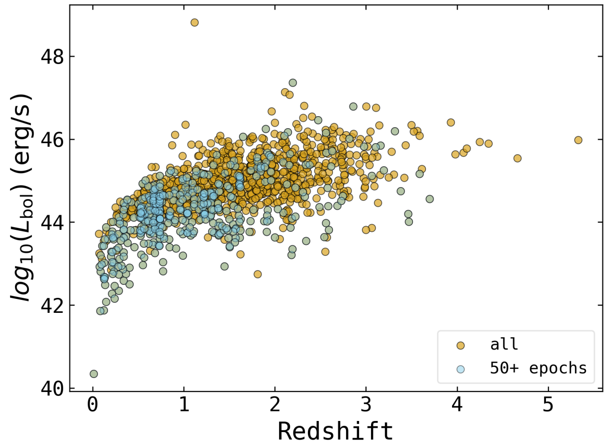

Figure 4. Bolometric luminosity–redshift distribution of the LSST DP1 quasar sample. Golden points show the full sample, while blue points indicate the subsample with at least 50 epochs in the LSST g, r, i, and z bands used in the time-domain analysis. The overall trend reflects the flux-limited nature of the parent catalogs.

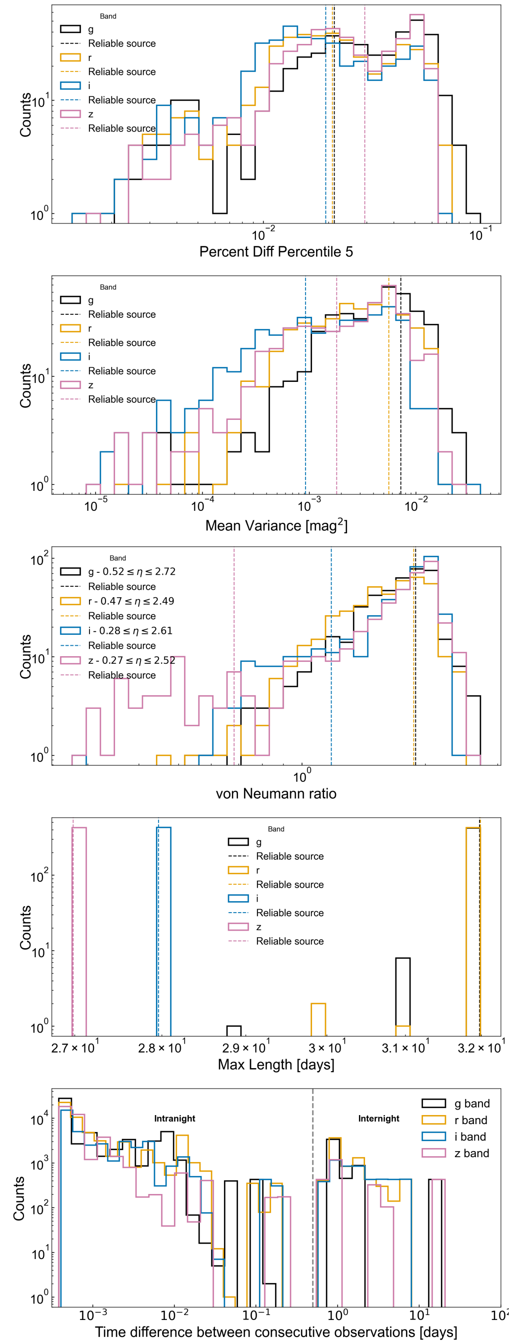

Figure 5. Distributions of selected light-curve properties in different photometric bands. From top to bottom: amplitude, mean variance, von Neumann ratio, continuous observing-window length in days, and time differences between consecutive observations. In the first four panels, dashed lines indicate the value for the detected candidate (Section 4). In the bottom panel, the gray dashed line marks the separation between intranight and internight cadences.

interval to exhibit strong serial correlation, so η naturally clusters near the uncorrelated expectation, $\eta \simeq 2$ (M.-S. Shin et al. 2009).

The fourth panel shows the maximum continuous observing-window length, computed using the standard FEETS implementation as the longest uninterrupted temporal interval without gaps exceeding the package-defined threshold. This feature captures the effective baseline over which variability can be probed without strong fragmentation by the window function, and is therefore relevant for assessing sensitivity to longer-timescale variability. All bands cluster between 27 and 32 days, reflecting the DP1 cadence. The strong periodic candidate source 3D-HST 2351 ($\alpha = 03^{\rm h}32^{\rm m}32.78^{\rm s}$, $\delta = -27°54'23.0''$, J2000; $z = 0.198$; LSST DP1 object_id = 611255072642321324) lies within the dominant cluster, indicating that its temporal coverage is representative of the sample. Hereafter, for brevity, we refer to this object simply as "the candidate."

The bottom panel shows the separations between consecutive observations. All filters display a dense intranight component at 10^{-3}–10^{-1} days, followed by a clear gap before internight cadences of 1–20 days.

3. Methods

3.1. Catalog Construction and Pipeline Overview

We construct the first catalog of AGN light-curve periodicity signatures in LSST DP1 by cross-matching LSST sources with external AGN catalogs and processing the resulting multiband light curves with the `QhX` pipeline. `QhX` identifies characteristic timescales in irregularly sampled multiband light curves through wavelet-based cross-correlation analysis followed by statistical vetting (A. B. Kovačević et al. 2026). The package was developed in `Python` as an in-kind contribution to the LSST AGN and TVS Science Collaborations. A compact summary of the full workflow is provided in Appendix A. The theoretical background and broader methodological context of `QhX` are described by A. B. Kovačević et al. (2018, 2019, 2020a), M. Fatović et al. (2023), A. Kovacevic (2024), and A. B. Kovačević et al. (2026); here, we retain only the elements needed to understand the catalog construction and vetting logic adopted in this work.

`QhX` represents each light curve by means of the weighted wavelet Z-transform (WWZ; G. Foster 1996). Denoting the time–period representation by $W(t, P)$, the variability signal is embedded as a square-integrable time–frequency surface

$$W(t, P) \in L^2(\mathbb{R}^2), \tag{1}$$

where $L^2(\mathbb{R}^2)$ denotes the space of square-integrable functions over the (t, P) domain.

To identify oscillatory coherence, `QhX` compares pairs of WWZ surfaces, either from different photometric bands or, in the autocorrelation case, from the same light curve. The comparison is performed through a nonlinear cross-correlation operator,

$$\mathrm{Corr}(P_1, P_2) = \iint \widetilde{W}_A(t, P_1)\, \widetilde{W}_B(t, P_2)\, dt\, dP, \tag{2}$$

where $\widetilde{W}_A$ and $\widetilde{W}_B$ denote normalized WWZ surfaces. A phase-coherent oscillatory process produces structured features in period–period space, whereas stochastic red-noise variability tends to generate diffuse correlation patterns with no stable geometric organization.

3.2. Period Estimation and Candidate-level Significance

For each light curve, `QhX` computes the WWZ map on a uniform angular-frequency grid $\omega \in [\omega_{\min}, \omega_{\max}]$, with $\omega = 2\pi/P$ (G. Foster 1996; S. Mallat 2008). To assess candidate oscillatory timescales, the two-dimensional cross-band coherence information is reduced to a one-dimensional coherence spectrum

$$S(\omega) \equiv \frac{1}{\max_\omega \sum_\tau |\mathcal{C}(\tau, \omega)|} \sum_\tau |\mathcal{C}(\tau, \omega)|, \tag{3}$$

where $\mathcal{C}(\tau, \omega)$ is the coherence statistic derived from the cross-correlation of WWZ surfaces evaluated at matched periods.

Candidate frequencies are identified from $S(\omega)$ using `scipy.signal.find_peaks`. We adopt conservative peak-height and prominence thresholds in order to suppress low-contrast extrema induced by stochastic variability and cadence structure. To improve frequency localization, $S(\omega)$ is linearly interpolated onto a refined grid with spacing $\Delta\omega/2$.

For a time series of total duration T_{obs}, the characteristic frequency resolution scales as $\Delta f \sim 1/T_{\mathrm{obs}}$. Since $f = 1/P$, first-order propagation implies

$$\frac{\Delta P}{P} \sim \frac{P}{T_{\mathrm{obs}}}. \tag{4}$$

This motivates our cross-band consistency tolerance: restricting attention to candidate oscillations satisfying $T_{\mathrm{obs}}/P \gtrsim 10$ implies $\Delta P/P \lesssim 0.1$. Accordingly, taking band i as a reference, candidate periods detected in band j are considered mutually consistent when

$$\frac{|P_j - P_i|}{P_i} \leqslant 0.1. \tag{5}$$

For each retained candidate peak, we estimate a frequentist per-peak significance using a permutation procedure. We generate N_{perm} surrogate light curves by randomly permuting the observed magnitudes while preserving the observational cadence. For each surrogate, the full WWZ → correlation → $S(\omega)$ procedure is repeated, and the surrogate coherence is evaluated at the same frequency index as the observed peak. The local p-value at candidate frequency ω_k is

$$\hat{p}(\omega_k) = \frac{1}{N_{\mathrm{perm}}} \sum_{r=1}^{N_{\mathrm{perm}}} \mathbb{I}\{S^{(r)}(\omega_k) \geqslant S(\omega_k)\}, \tag{6}$$

where $\mathbb{I}[\cdot]$ denotes the indicator function. We report the corresponding per-peak significance as $1 - \hat{p}(\omega_k)$.

Let ω_* denote the frequency of the most significant retained peak. The corresponding period estimate is

$$P_* = \frac{1}{\omega_*}. \tag{7}$$

Following standard spectral-estimation practice, candidate-level period uncertainties are obtained from the width of the dominant coherence peak (J. H. Horne & S. L. Baliunas 1986; A. Schwarzenberg-Czerny 1991; W. H. Press et al. 1992; P. Stoica et al. 2005). Specifically, we measure the full width at half-maximum relative to the local background of $S(\omega)$ and determine the left and right half-maximum crossings, ω_{L} and ω_{R}. This yields asymmetric period uncertainties

$$\Delta P_{\mathrm{lo}} = \frac{1}{\omega_*} - \frac{1}{\omega_{\mathrm{R}}}, \qquad \Delta P_{\mathrm{up}} = \frac{1}{\omega_{\mathrm{L}}} - \frac{1}{\omega_*}. \tag{8}$$

When these quantities are unavailable or ill-defined, conservative fallback proxy uncertainties are assigned later during object-level catalog construction; the detailed prescription is given in Appendix B.

3.3. Internal Multiband Consistency Checks

Before astrophysical interpretation, `QhX` applies an internal vetting stage designed to retain only candidates that are statistically significant, spectrally localized, and mutually consistent across bands. This stage is not intended as formal model selection; rather, it suppresses contamination from stochastic red-noise fluctuations, aliases, and poorly constrained peaks.

For each detected peak, we define the feature vector

$$\boldsymbol{x} = (\mathrm{significance}, \epsilon_{\max}, \mathrm{IoU}), \tag{9}$$

where significance $= 1 - \hat{p}(\omega_k)$ is the permutation-based significance, $\epsilon_{\max} = \max(\Delta P_{\mathrm{lo}}, \Delta P_{\mathrm{up}})/P$ is the fractional localization width, and IoU is the intersection over union of the period intervals inferred in two bands.

For band pair (1, 2), we define the corresponding uncertainty intervals as

$$I_b = [P_b - \Delta P_{\mathrm{lo},b},\ P_b + \Delta P_{\mathrm{up},b}], \qquad b = 1, 2, \tag{10}$$

and compute

$$\mathrm{IoU}(I_1, I_2) = \frac{|I_1 \cap I_2|}{|\, I_1 \cup I_2 \,|}, \tag{11}$$

where $|I|$ denotes the length of the interval I (i.e., the measure of the interval on the real line). Large IoU values indicate that two bands yield mutually compatible period estimates within their uncertainties; this is equivalent to a one-dimensional Jaccard-type interval-overlap measure (P. Jaccard 1912; E. B. Fowlkes & C. L. Mallows 1983; M. Everingham et al. 2010).

Using these quantities, each candidate is assigned an internal QhX class:

$$\begin{aligned}
\text{high:}\quad & \begin{cases} \text{significance} \geqslant 0.99, \\ \epsilon_{\max} \leqslant 0.1, \\ \mathrm{IoU} \geqslant 0.99; \end{cases} \\
\text{moderate:}\quad & \begin{cases} 0.5 \leqslant \text{significance} < 0.99, \\ 0.1 < \epsilon_{\max} \leqslant 0.3, \\ 0.8 \leqslant \mathrm{IoU} < 0.99; \end{cases} \\
\text{poor:}\quad & \text{otherwise.}
\end{aligned} \tag{12}$$

For the highest-quality internally supported candidates, we additionally test a physically motivated red-noise null hypothesis using correlation-preserving bootstrap surrogates. The technical details of this auxiliary red-noise assessment are given in Appendix C. Detection tests on simulated data are presented in Appendix D.

3.4. Postdetection Multiband Vetting

The postdetection vetting step tests whether a candidate period identified by QhX corresponds to a coherent oscillatory signal that can be detected independently in multiple photometric bands. At this stage the candidate frequency is fixed; we do not re-search for a period, but instead project each band onto a sinusoidal basis at the detected angular frequency ($\omega = \frac{2\pi}{P_*}$)

$$\hat{m}_b(t) = A_b \cos(\omega(t - t_0)) + B_b \sin(\omega(t - t_0)) + C_{0,b}. \tag{13}$$

From the fitted coefficients we derive the oscillation amplitude and phase

$$\mathrm{Amp}_b = \sqrt{A_b^2 + B_b^2}, \qquad \phi_b = \arctan 2(B_b, A_b). \tag{14}$$

Here, ϕ_b is measured in radians.

Residuals are defined as $\epsilon_{b,i} = m_{b,i} - \hat{m}_b(t_i)$ and normalized using the median absolute deviation

$$r_{b,i} = \frac{\epsilon_{b,i}}{s_b}, \qquad s_b = 1.4826\ \mathrm{median}_j|\epsilon_{b,j} - \mathrm{median}_k(\epsilon_{b,k})|, \tag{15}$$

where s_b is a robust estimate of the residual scatter in band b. The sinusoidal fit is then obtained by minimizing the Huber loss applied to the normalized residuals

$$\rho_c(r) = \begin{cases} \frac{1}{2}r^2, & |r| \leqslant c, \\ c|r| - \frac{1}{2}c^2, & |r| > c, \end{cases} \tag{16}$$

where c is the tuning threshold at which the loss changes from quadratic least-squares behavior to linear, outlier-robust behavior. Thus, the procedure combines the median absolute deviation (MAD)-based residual scaling with robust Huber-loss fitting, reducing the influence of outliers and heteroscedastic photometric errors (P. J. Huber 1964; T. C. Beers et al. 1990; P. J. Rousseeuw & C. Croux 1993). We set $c = 1$, so that residuals larger than approximately one MAD-scaled scatter unit are down-weighted (P. Amaro Seoane 2026).

For each band we compute the diagnostics

$$\mathrm{amp\_snr}_b = \frac{\mathrm{Amp}_b}{\sigma_{A_b}}, \qquad R_b^2 = 1 - \frac{\sum_i (m_{b,i} - \hat{m}_{b,i})^2}{\sum_i (m_{b,i} - \bar{m}_b)^2}, \tag{17}$$

where σ_{A_b} is the formal amplitude uncertainty and $\bar{m}_b$ is the mean magnitude in band b.

A reference band b_{ref} defines the phase zero-point, and interband phase lags are measured in cycles as

$$\Delta\phi_{b|\mathrm{ref}} = \frac{\phi_b - \phi_{b_{\mathrm{ref}}}}{2\pi}. \tag{18}$$

A band is marked as supporting the candidate when it satisfies conservative trigger requirements on both oscillation strength and fit quality

$$\mathrm{good\_band}_b \coloneqq (\mathrm{amp\_snr}_b \geqslant 1.5) \wedge (R_b^2 \geqslant 0.10). \tag{19}$$

The threshold $\mathrm{amp\_snr}_b \geqslant 1.5$ ensures a nonnegligible fitted oscillatory amplitude, while $R_b^2 \geqslant 0.10$ requires the fixed-frequency projection to explain at least 10% of the band variability. A supporting band is further classified as coherent if its phase agrees with that of the reference band within the adopted tolerance:

$$\mathrm{coherent\_band}_b \coloneqq \mathrm{good\_band}_b \wedge (|\Delta\phi_{b|\mathrm{ref}}| \leqslant 0.20). \tag{20}$$

At the object level we then count

$$n_{\mathrm{good}} = \sum_b \mathbb{I}(\mathrm{good\_band}_b), \tag{21}$$

$$n_{\mathrm{coherent}} = \sum_b \mathbb{I}(\mathrm{coherent\_band}_b). \tag{22}$$

where $\mathbb{I}(\cdot)$ denotes the indicator function: it equals 1 if the condition in brackets is satisfied and 0 otherwise, so that n_{good} and n_{coherent} simply count the number of bands satisfying the respective trigger conditions.

Requiring coherent detection in more than one photometric band provides an additional coincidence test across separate observational channels. This mitigates look-elsewhere effects and suppresses cadence-induced aliases, since noise- or sampling-generated peaks are less likely to recur with compatible phase across bands (J. H. Horne & S. L. Baliunas 1986; S. Vaughan 2005).

The final catalog verdict is assigned according to

$$\text{verdict} = \begin{cases} \text{strong}, & n_{\text{coherent}} \geqslant 2, \\ \text{medium}, & n_{\text{good}} \geqslant 1, \\ \text{discarded}, & \text{otherwise}. \end{cases} \tag{23}$$

Finally, each detected object is assigned both a main catalog verdict and, where useful, a diagnostic bucket: *coherent*, *pseudoperiodic*, *transiently coherent*, or *red-noise dominated*. These buckets should not be interpreted as unique physical classes. Rather, they summarize why a candidate did or did not pass the strongest coherence requirements: a strong candidate is coherent according to its definition, pseudoperiodic objects show local oscillatory support but weak phase agreement across bands, transiently coherent objects show coherence only in a limited band context, and red-noise-dominated objects lack significant single-band oscillatory support. The formal implementation-level definitions of these buckets are given in Appendix C. For the strong selected candidate, we additionally test whether the observed coherence could plausibly arise from correlated red-noise variability. This test is performed with moving-block-bootstrap surrogate light curves that preserve local temporal correlations. Because the DP1 baseline is short, this nonparametric red-noise test is used as a baseline-appropriate check rather than as a full stochastic process model; the detailed formulation is given in Appendix E.

3.5. Catalog-level Multiple-testing Control

Because the periodicity search is performed over an ensemble of objects, we additionally control the catalog-level false discovery rate (FDR) using the Benjamini–Hochberg procedure (Y. Benjamini & Y. Hochberg 1995; Y. Benjamini 2010). For each promoted object, we retain a single object-level p-value associated with the retained candidate period P_*, namely the permutation-based p-value of the peak that enters the postdetection vetting stage. We then apply the Benjamini–Hochberg procedure across this family of promoted object-level candidates. This step is postdetection: it does not alter the per-object identification or multiband vetting, but quantifies the expected fraction of false positives among the reported candidates. We adopt a reference FDR threshold of $\alpha = 0.05$, with the usual caution that p-values and FDR thresholds should be interpreted as catalog-level error-control summaries rather than absolute probabilities of astrophysical truth (R. L. Wasserstein & N. A. Lazar 2016; N. Galwey 2025).

4. Results and Discussion

4.1. Overview of the Catalog

We present the first catalog of coherent and quasi-coherent oscillatory signatures identified in AGN light curves from Rubin LSST DP1, constructed with the `QhX` v0.2.1 pipeline (A. B. Kovačević et al. 2026). The catalog (M. Stanković Cerović et al. 2025) is open source and publicly available through a dedicated data portal,[44] which provides an interactive interface for inspection of selected candidates and their associated diagnostics (see Appendix F). The portal is accessible on both desktop and mobile devices and is designed to support detailed candidate assessment, with future potential for broader community use, including citizen-science-oriented exploration. In addition to tabulated numerical outputs, the portal includes interactive visual diagnostics that allow direct inspection of signal consistency and provide greater interpretability than static plots alone.

The catalog contains both rare highest-confidence cases (reliable) and the more numerous low (poor)-confidence detections, which are further stratified into pseudoperiodic, transiently coherent, and red-noise-dominated behaviors. Two complementary data products are provided: (1) an *object-level table*, with one row per AGN summarizing the most significant detected period, its uncertainty, and cross-band diagnostics, and (2) a *per-band table*, with one row per AGN×band containing the underlying fit quality and phase-coherence metrics. A detailed schema of all catalog columns, together with access to the full machine-readable products, is provided through the portal. Starting from an initial sample of 428 DP1 quasars, 230 objects remain after applying the internal `QhX` selection criteria (Section 3.3). These 230 objects were then evaluated using the postdetection vetting framework introduced in Section 3.4. Of these, only one source, hereafter referred to as *the candidate*, is retained as a high-confidence periodic detection in the LSST DP1 sample and classified as *strong* (i.e., reliable), as it satisfies the strictest significance and cross-band consistency requirements. In comparison, 15 sources satisfy intermediate robustness conditions and are labeled *medium*, while the remaining 214 fail one or more robustness criteria and are flagged as *discard*. A schematic representation of the full selection workflow is shown in Appendix A.

Of these, only one source is retained as a high-confidence periodic detection in the LSST DP1 sample: the candidate that satisfies the strictest significance and cross-band consistency requirements and is therefore classified as *strong (i.e., reliable)*, 15 satisfy intermediate robustness conditions and are labeled *medium*, and the remaining 214 fail one or more robustness criteria and are flagged as *discard*. A schematic representation of the full selection workflow is shown in Appendix A.

Applying FDR control at $\alpha = 0.05$ (Section 3.5), no vetted candidates are rejected. This indicates that the combination of internal `QhX` classification and cross-band postdetection evaluation yields a conservative catalog with respect to false positives.

4.2. Distribution of Postdetection Vetting Metrics

Figure 6 summarizes the per-band distributions of the postdetection vetting diagnostics for the medium-ranked objects, namely the detected period, amplitude signal-to-noise ratio (S/N), coherence, and R^2. The single high-confidence object is shown separately as a star symbol.

Across all four analyzed filters (g, r, i, and z), the recovered periods (Figure 6(a)) cluster around $P \simeq 2.5$–3 days, with interquartile ranges typically spanning $\Delta P \lesssim 2$ days. No strong systematic dependence of the detected period on wavelength is apparent, consistent with a common underlying driver of variability across the optical continuum. The amplitude S/N (Figure 6(b)) shows broad but positively skewed distributions, with median values of ∼3–5 in the bluer bands and a gradual increase toward the z band, where S/N ≈ 6. This trend may reflect a combination of band-dependent flux contrast, sampling, and photometric precision,

[44] https://ser-sag.pmf.kg.ac.rs:8081/agn_catalogue.php

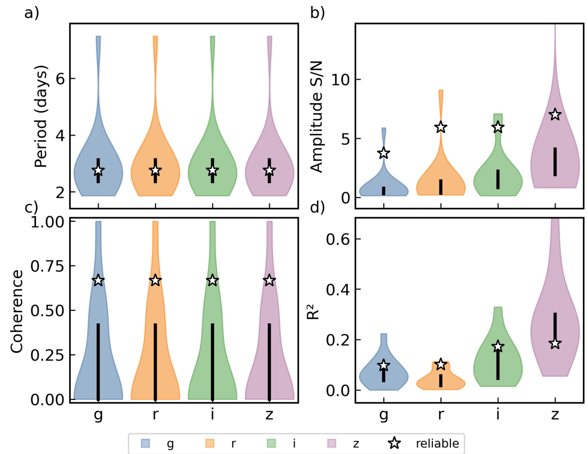

Figure 6. Violin diagrams of vetting metrics for AGN flagged as medium (distributions with interquartile bars) and the single strong candidate (star). Panels show (a) detected period, (b) amplitude S/N, (c) object-level coherence fraction $C_{\rm obj}$ (Equation (C2)), and (d) R^2 across the bands. The strong candidate, with $P = 2.76$ days, lies within the dominant 2 and 4 day cadence-imposed modes but is distinguished by higher amplitude S/N and coherence.

although the present sample remains too small to support a firmer physical interpretation.

The coherence parameter (Figure 6(c)) has median values of $C \sim 0.4$–0.6 among the medium-ranked candidates, while the corresponding R^2 statistic (Figure 6(d)) is typically $\sim$0.2–0.3. This indicates that in most medium-ranked cases the periodic component accounts for only a modest fraction of the total variance.

Taken together, these results show that even within the limited temporal baseline of DP1, Rubin LSST photometry can already isolate candidates whose joint amplitude, multiband coherence, and goodness of fit are difficult to explain as purely stochastic variability. At the same time, the injection tests show that the detailed periodogram structure remains strongly shaped by cadence-induced aliasing and must therefore be interpreted with caution.

4.2.1. Dawid–Skene and Alias-based Estimates of Catalog Purity

To quantify the reliability of the identified periodic candidates, we estimated two complementary purity metrics: the Dawid–Skene (DS; Appendix G) ensemble purity and the alias-floor purity.

The DS purity, $P_{\rm DS}$, provides a probabilistic measure of the internal consistency of the catalog. In this framework, each LSST photometric band is treated as an independent detector that votes for or against periodicity using strict criteria based on phase coherence, amplitude S/N (S/N > 1.5), coefficient of determination ($R^2 > 0.10$), and interband phase-lag tolerance ($|\Delta\phi| < 0.1$ cycles). The DS model then infers the true-positive and false-positive rates of these band-level detectors and assigns to each object a posterior probability p_i of being genuinely periodic.

For the present catalog, we obtain $P_{\rm DS} \simeq 0.023$, implying that $\sim$2% of the sample might correspond to genuinely periodic sources under this conservative multiband-consensus model. This low global purity does not imply poor performance. Rather, it reflects the intentionally stringent design of the pipeline, which prioritizes robustness and suppression of marginal detections at the expense of completeness.

In parallel, we evaluate the alias-floor purity, $P_{\rm alias}$, which measures the fraction of candidates whose best periods do not lie near well-known sampling aliases such as 1 day, 0.5 days, 29.53 days, and 365.25 days, following the general logic of alias avoidance discussed by D. Kramer et al. (2023). It is defined as

$$P_{\rm alias} = 1 - \frac{N_{\rm alias}}{N}, \tag{24}$$

where $N_{\rm alias}$ is the number of alias-flagged sources. For our sample, we obtain $P_{\rm alias} \simeq 0.975$, indicating that nearly all *medium*-ranked periods are not located at the dominant known sampling aliases.[45]

The contrast between $P_{\rm DS}$ and $P_{\rm alias}$ highlights the distinct roles of these two metrics. The alias-floor purity measures the cleanliness of the sampling pattern with respect to a predefined set of major cadence aliases, whereas the DS purity measures the strength of statistically significant cross-band consensus. In this sense, the catalog is clean with respect to the dominant known alias frequencies, but deliberately conservative in retaining only candidates supported by multiband coherence rather than marginal single-band detections.

Applying a strict reliability threshold of $p_i > 0.801$,[46] only one source, `object_id` = 611255072642321324, is retained as a high-confidence periodic detection in the LSST DP1 sample.

4.3. Observational Properties of the High-confidence Candidate

We now examine the highest-confidence object in the catalog in order to characterize its multiband photometric behavior, coherence, and variability properties. The source is located in the ECDFS at $z \sim 0.198$ and exhibits a quasi-sinusoidal modulation with a period of $P = 2.76^{+2.65}_{-0.06}$ days, consistently recovered across all four LSST filters, while the coherence criterion is satisfied in the r and i bands. The variability diagnostics shown in Figure 5 further distinguish this source from the bulk of the AGN population. In the percentile-amplitude distribution (top panel), the candidate lies above the bulk of objects in all bands, indicating consistently enhanced variability rather than isolated outlier behavior. In the mean variance (second panel), the source lies toward the upper part of the distribution in each band, indicating variability amplitudes that are elevated relative to much of the sample. The von Neumann ratio (third panel) places the source within the main body of the distribution in each band, indicating that its point-to-point variability is not unusually erratic relative to the sample. Finally, its maximum baseline (fourth panel) and sampling properties (fifth panel) lie within the transition regime between limited and more favorable temporal coverage.

To quantify the effect of the sampling pattern, we evaluated the normalized spectral-window power $W(P) \in [0, 1]$ (Figure 7). The detected period of the high-confidence object, $P_{\rm det} = 2.76^{+2.65}_{-0.06}$ days, lies away from the strongest diurnal

[45] Alias amplitudes are governed primarily by the spectral window function and by phase clustering of visits, rather than by cadence density alone (K. J. Bell et al. 2018). This suggests that the full LSST survey cadence should yield improved alias suppression and period recoverability relative to DP1.

[46] The cutoff $p_i > 0.801$ corresponds to the high-consensus tail of the DS posterior distribution, retaining approximately 95% of strict periodic candidates while suppressing marginal detections.

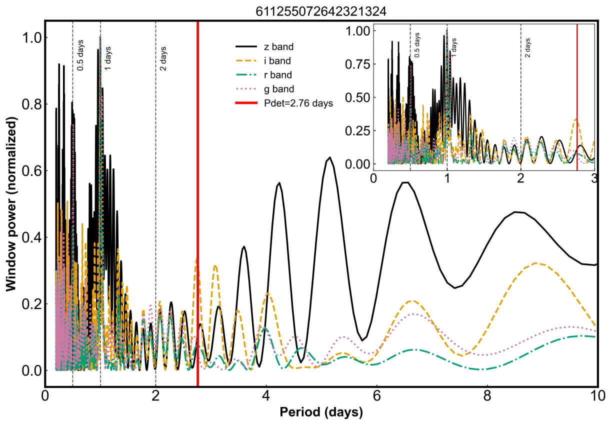

Figure 7. Spectral window function of source 611255072642321324 across the g, r, i, and z bands. The curves show the normalized power spectrum of the discrete sampling pattern in each band, illustrating how the observational cadence redistributes signal power across frequencies. The gray vertical line marks the `QhX`-detected period, $P_{\rm det} = 2.76$ days, while the black dashed vertical lines indicate the principal aliases associated with observing patterns at 0.5, 1, and 2 days. Inset: zoom of the sub–3 day region, showing the dense alias structure common to all bands and dominated by diurnal sampling and its harmonics.

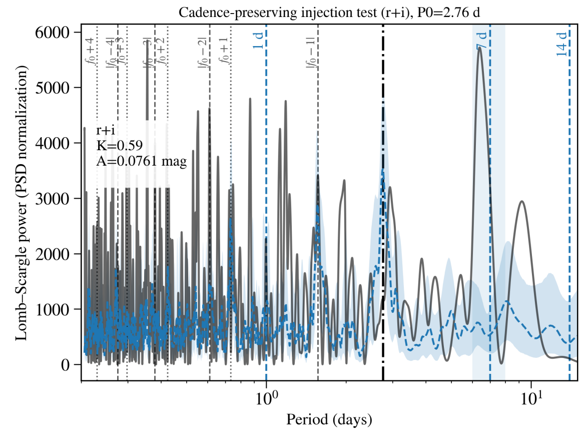

Figure 8. Lomb–Scargle alias injection test for the combined $r + i$ light curve of the high-confidence `QhX` candidate with $P_0 = 2.76$ days, marked by the black dashed–dotted vertical line. The dark gray curve shows the observed combined $r + i$ Lomb–Scargle periodogram (power spectral density normalization), while the dashed blue curve and shaded region indicate the median and 16th–84th percentile range from $N_{\rm sim} = 5000$ OU+sinusoid simulations sampled at the exact observation times. Vertical gray lines mark the expected diurnal alias locations, $|f_0 \pm n\ \mathrm{day}^{-1}|$ (J. T. VanderPlas 2018), illustrating how the survey cadence redistributes power into characteristic alias peaks around the injected period.

resonances at 0.5, 1, and 2 days, and the window response at $P_{\rm det}$ is modest in all bands ($W \lesssim 0.4$), remaining well below the dominant spectral-window peaks. This indicates that the detected period is not coincident with the dominant cadence-induced aliases, although the survey window can still broaden or distort the recovered peak structure. The asymmetric period uncertainty reflects the shape of the corresponding `QhX` correlagram peak, which is broadened preferentially toward longer periods, with the asymmetry further influenced by the finite DP1 baseline and the nonuniform sampling pattern. As new LSST data become available, this source will be reanalyzed on progressively longer baselines to assess the persistence of the detected period, the evolution of its uncertainty, the stability of its phase, and the long-term coherence of the modulation. We also tested this source against the red-noise hypothesis using a fixed-frequency multiband bootstrap analysis of the coherence statistic (see Appendix E). Across 9745 valid surrogate realizations, the empirical joint probability is 2.1×10^{-4}, indicating that stochastic red-noise variability only rarely reproduces the observed level of multiband coherence at the candidate frequency.

However, as emphasized by J. T. VanderPlas (2018), the interaction between the survey window function and the intrinsic variability spectrum can still redistribute power such that the dominant Lomb–Scargle peak appears at a diurnal alias rather than at the true driving frequency. To examine this failure mode more directly, beyond the scalar diagnostic $W(P)$, we performed cadence-preserving injection tests on the combined $r + i$ light curve of the high-confidence `QhX` candidate.

Each band was demeaned independently using a weighted mean, and the resulting measurements were concatenated prior to period analysis. The observed Lomb–Scargle periodogram was computed on a fixed frequency grid spanning $f_0 \pm N_{\rm alias}\ \mathrm{day}^{-1}$, where $f_0 = 1/P_0$.

The stochastic variability was modeled as red noise using an Ornstein–Uhlenbeck (OU) process, equivalent to a DRW with exponential covariance. We characterize this process by a damping timescale τ and a long-term variability amplitude σ_∞ (the asymptotic rms variability of the OU/DRW process). We fixed $\tau = 150$ days[47] and estimated σ_∞ from the observed combined $r + i$ light curve by maximum likelihood, explicitly accounting for the reported photometric uncertainties.

Each simulated light curve was constructed as

$$m_{\rm sim}(t) = m_{\rm OU}(t) + A_{\rm inj} \sin\left(\frac{2\pi t}{P_0} + \phi\right), \tag{25}$$

where $m_{\rm OU}(t)$ is a realization of the OU process and $\phi \sim \mathrm{Uniform}(0, 2\pi)$ is a random phase. The injected sinusoidal amplitude was parameterized as $A_{\rm inj} = K\sigma_\infty$. We adopted $K = 0.59$, selected empirically as the value that produced the closest match between the simulated and observed Lomb–Scargle periodograms near the candidate period and its principal daily aliases. We generated an ensemble of $N_{\rm sim} = 5000$ realizations.

For each realization, we computed the Lomb–Scargle periodogram using the same normalization, frequency grid, and uncertainty weights as for the observed data. The ensemble was then summarized by the median periodogram and the 16th–84th percentile interval at each frequency. Figure 8 compares the observed combined $r + i$ Lomb–Scargle periodogram of the high-confidence candidate at $P_0 = 2.76$ days with the cadence-preserving OU+sinusoid simulations. The simulations reproduce not only the power near P_0 but also the prominent secondary peaks at the expected diurnal alias locations, $|f_0 \pm n\ \mathrm{day}^{-1}|$ (gray vertical lines), demonstrating how the survey window redistributes power away from the true frequency. Small offsets between the

[47] The smallest characteristic damping timescale reported by C. L. MacLeod et al. (2010) is of order ~100 days in the observer frame, while M. J. Graham et al. (2014) found a minimum timescale of ~54 days in the source rest frame, corresponding to ~150 days in the observer frame. These values remain sufficiently long compared with P_0 so the precise choice of τ does not materially affect the inferred alias pattern.

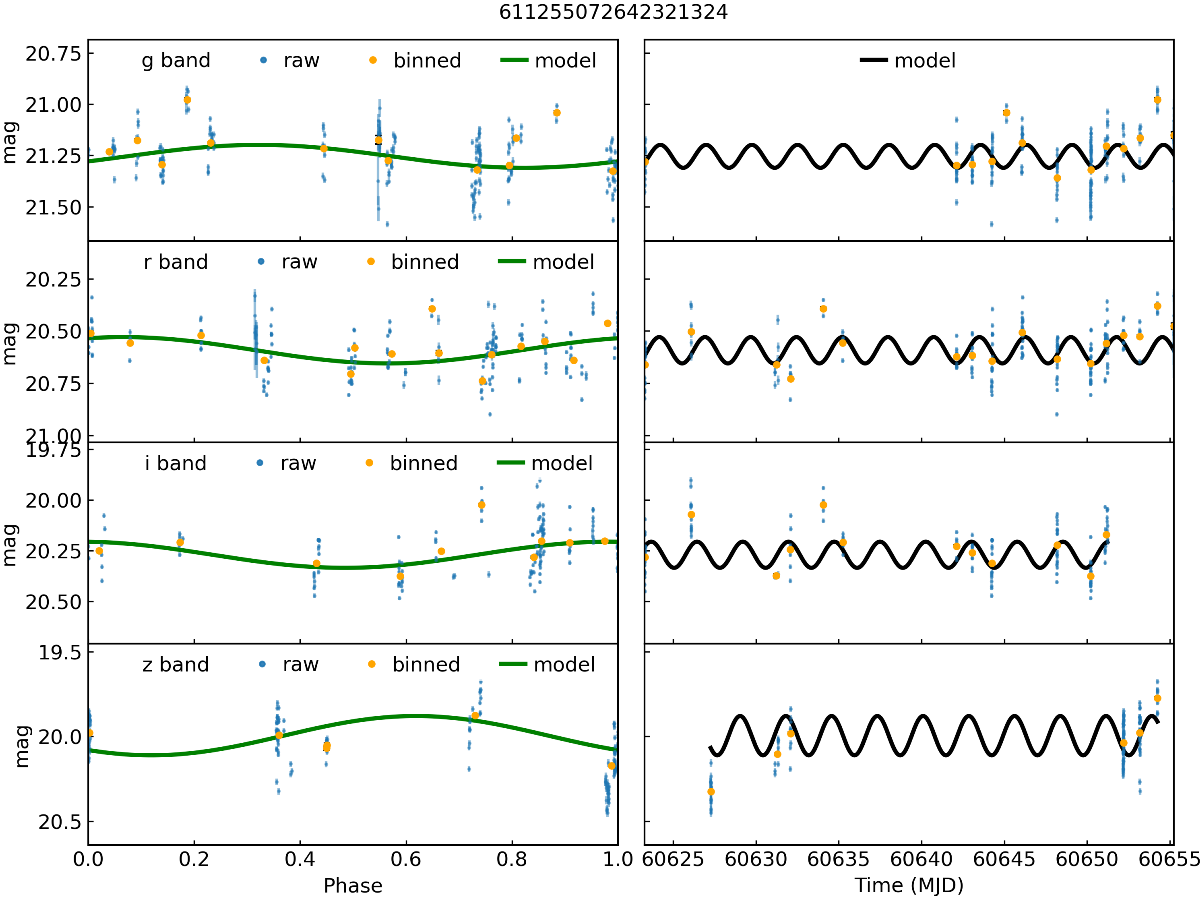

Figure 9. Phase-folded (left) and time-domain (right) light curves of the candidate (ID = 611255072642321324) across the four LSST filters (g, r, i, and z). Blue points show the raw photometric measurements, orange points indicate phase-binned means, and the solid curves show the best-fit sinusoidal models from the verification step. The recovered period is $P = 2.76^{+2.65}_{-0.06}$ days, with broadly consistent behavior across the four bands, amplitude S/N of order S/N $\approx$ 8–10, and coherence $C \approx 0.8$. The modulation is quasi-sinusoidal and nearly achromatic, consistent with a coherent oscillatory component under the adopted vetting criteria.

observed peak and the median simulated response are expected because the plotted median is a bin-wise ensemble statistic over many realizations, whereas the observed Lomb–Scargle peak is evaluated on a discrete frequency grid and is further modulated by the specific cadence window. The close agreement between the observed and simulated alias structure indicates that the detailed morphology of the periodogram is strongly shaped by sampling-induced aliasing, consistent with the Lomb–Scargle failure modes discussed by J. T. VanderPlas (2018).

Its phase-folded light curves (Figure 9; Table 3) show smooth oscillatory behavior, with residual scatter comparable to the observed point-to-point dispersion and nearly constant amplitude ratios among the g, r, i, and z bands, suggesting largely achromatic variability. The source also exhibits moderate-to-high amplitude-to-noise ratios ($A/\tilde{\sigma}_{\rm phot} \approx 4$–7) across multiple bands, together with nonnegligible sinusoidal fit quality ($R^2 \approx 0.10$–0.18), indicating that the detected modulation is not dominated by photometric noise (Table 3).

The phase lags relative to the reference band remain below $\sim$0.45 cycles, with two bands satisfying the adopted multiband coherence criterion, supporting the presence of a temporally correlated signal across filters. Given the short $\sim$32 day baseline, these phase offsets should be interpreted primarily as relative alignment diagnostics rather than as precise interband time delays. The high coherence ($C \approx 0.8$) and amplitude S/N ($\approx$9) further argue against stochastic red-noise fluctuations as the dominant explanation (X. Zhang 2022).

The optical colors calculated from mean magnitudes (see P. Sánchez-Sáez et al. 2019, 2021) place the high-confidence source slightly redward ($g - r > 0.6$ and $r - i < 0.5$) of the median quasar locus in $g - r$, while remaining consistent with the population trend in $r - i$ (Figure 10). This behavior is consistent with modest reddening or mild host-galaxy contribution, which may attenuate variability amplitudes but cannot generate spurious periodicity. The detected modulation is therefore interpreted as intrinsic to the AGN. The distributions of variability amplitude, mean variance, and sampling-corrected time differences (Figure 5) do not reveal a separate high-variability regime within the population.

4.4. Illustration for Future LSST AGN Population Studies

The `QhX` coherent variability analysis presented here is an early prototype of how future Rubin LSST AGN population studies may be performed on progressively longer time baselines.

To place the detected periods in broader physical context, we compare them with approximate black-hole-mass proxies in the $M_{\rm BH}$ period plane. This is motivated by population-level studies showing that characteristic quasar variability time-scales correlate statistically with source properties, including

Table 3
Band-level Statistical and Phase Diagnostics for the Periodic Candidate with Identifier 611255072642321324

Band	A (mag)	$\tilde{\sigma}_{\mathrm{phot}}$ (mag)	$A/\tilde{\sigma}_{\mathrm{phot}}$	S/N	R^2	Median R^2	Phase	ΔPhase (versus Ref)	Phase Lag (Cycle)	N_{used}	N_{coh}	Baseline (day)	Good	Coherent
g	0.0559	0.0149	3.74	5.94	0.098	0.137	0.823	0.256	0.256	192	2	31.99	F	F
r	0.0629	0.0106	5.94	5.94	0.101	0.137	0.566	0.000	0.000	193	2	31.99	T	T
i	0.0643	0.0108	5.94	5.94	0.172	0.137	0.484	0.082	0.082	137	2	31.99	T	T
z	0.1157	0.0165	7.01	5.94	0.185	0.137	0.119	0.447	0.447	127	2	31.99	F	F

Notes. Rows correspond to individual photometric bands; only an excerpt of the full diagnostic table is shown.
A denotes the fitted oscillation amplitude, $\tilde{\sigma}_{\mathrm{phot}}$ the median photometric uncertainty, and $A/\tilde{\sigma}_{\mathrm{phot}}$ the corresponding amplitude S/N. The coefficient R^2 quantifies the quality of the sinusoidal fit. Phase is the fitted oscillation phase expressed in cycles, while phase lag is measured relative to the reference band. N_{used} gives the number of photometric points used in the fit, and N_{coh} the number of bands supporting the coherent detection.

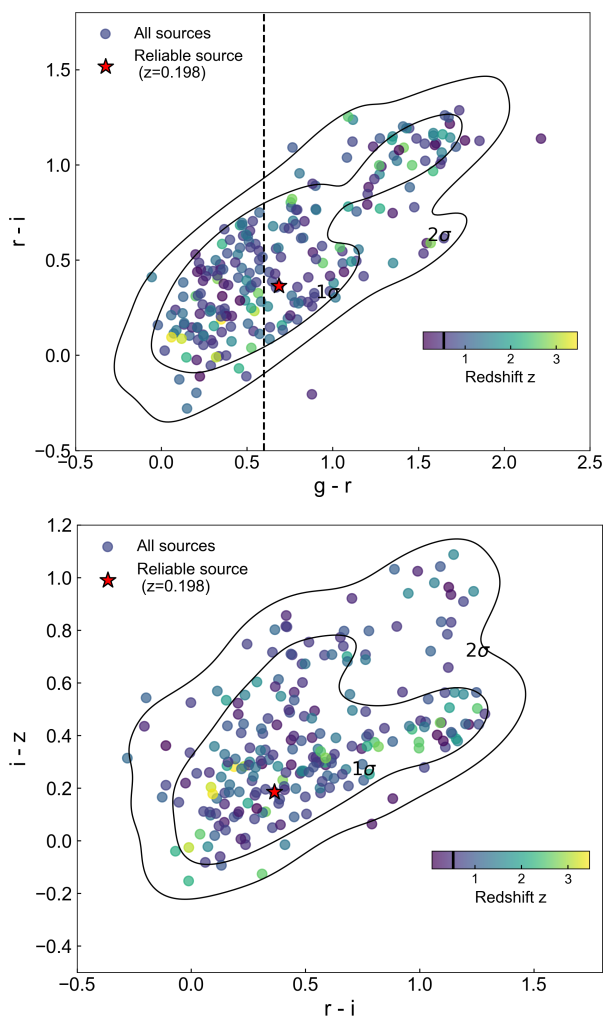

Figure 10. Distributions of the studied sample in the $g - r$ versus $r - i$ (top) and $r - i$ versus $i - z$ (bottom) spaces. The color bar shows the redshift of the sources. Contour lines show the 68% and 95% highest-density regions derived from the two-dimensional kernel density estimate. The red star represents the only source classified as reliable and is within 1σ from the mean locus in both cases. The dashed vertical line (top plot) marks the empirical color division $g - r = 0.6$ adopted by P. Sánchez-Sáez et al. (2019), separating blue ($g - r \leqslant 0.6$) from redder AGN ($g - r > 0.6$). The black stripe on the color bar indicates the redshift of the reliable source.

M_{BH}, particularly in large samples with long or densely sampled light curves (A. Bauer et al. 2009; N. Caplar et al. 2017; S. Kozłowski 2017; P. Sánchez-Sáez et al. 2018; Y. Tachibana et al. 2020; K. L. Suberlak et al. 2021; P. Arévalo et al. 2024; W. Yu et al. 2025; V. Petrecca et al. 2026); in particular, several works report a strong positive correlation between the DRW damping timescale and M_{BH} (see Table 1 in the review by M. Paolillo & I. Papadakis 2025). For this purpose, we estimate L_{3000} by fitting an empirically derived quasar spectral energy distribution template to the Rubin *u, g, r, i, z,* and *y* photometry (M. J. Temple et al. 2021), excluding bands sampling rest-frame wavelengths shorter than 1215 Å, and adopt $L_{\mathrm{bol}} = 5.15\,L_{3000}$. Defining the Eddington ratio as $\lambda_{\mathrm{Edd}} \equiv L_{\mathrm{bol}}/L_{\mathrm{Edd}}$, with $L_{\mathrm{Edd}} = 1.26 \times 10^{38}(M_{\mathrm{BH}}/M_{\odot})\,\mathrm{erg\,s^{-1}}$, gives

$$\frac{M_{\mathrm{BH}}}{M_{\odot}} = \frac{L_{\mathrm{bol}}}{\lambda_{\mathrm{Edd}}\,1.26 \times 10^{38}\,\mathrm{erg\,s^{-1}}}. \tag{26}$$

Because quasar Eddington ratios span several orders of magnitude and are typically sub-Eddington (e.g., Y. Shen et al. 2011; B. C. Kelly et al. 2013; A. Wójtowicz et al. 2026), we adopt, for illustrative purposes, a heuristic two-component Gaussian-mixture prior in $\log_{10}\lambda_{\mathrm{Edd}}$ to represent plausible moderate- and low-accretion populations. Here, "prior" denotes an assumed probability distribution used to represent plausible values of λ_{Edd} in the absence of source-specific constraints. The two-component form is adopted only as a simple phenomenological description of broad population scatter, rather than as a fit to the present sample. We use this prior only to propagate uncertainty through 500 Monte Carlo realizations of M_{BH} at fixed L_{bol}, so the resulting values should be interpreted as order-of-magnitude population-level proxies rather than precise per-object black-hole-mass measurements. With these statistical caveats in mind, we briefly illustrate how the detected candidate timescales relate to standard accretion-disk scalings, purely as a reference for future LSST population studies. The modulation amplitude increases mildly toward longer wavelengths. In a standard thin disk, longer wavelengths are emitted predominantly at larger radii, with $R_{\lambda} \propto \lambda^{4/3}$ (E. M. Cackett et al. 2007). For a representative proxy mass of $M_{\mathrm{BH}} \approx 10^{7}\,M_{\odot}$, the observed period of $P = 2.76$ days corresponds to a Keplerian orbital radius of approximately $150\,r_{\mathrm{g}}$ (M. A. Abramowicz & P. C. Fragile 2013), placing the variability scale within the inner accretion flow, outside the innermost stable circular orbit and well interior to the broad-line region.

Figure 11 (left panel) places the detected rest-frame periods in approximate relation to the thermal timescale of a standard

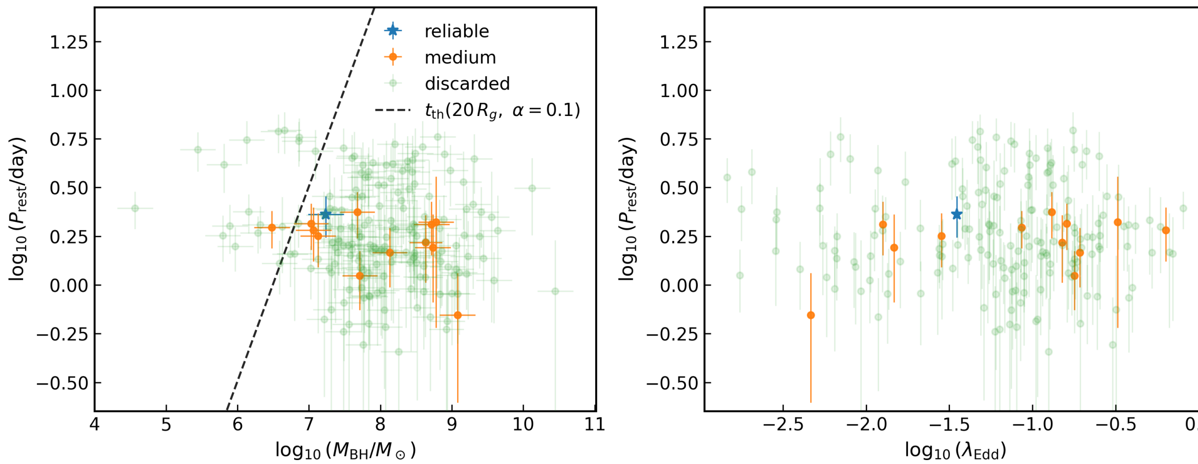

Figure 11. Period–mass (left) and period–Eddington ratio (right) distributions of AGN with coherent oscillation signatures identified in LSST DP1. The candidate is shown in blue, medium-ranked candidates in orange, and rejected candidates in green. Black hole masses and Eddington ratios are estimated using the proxy framework. Error bars reflect propagated internal uncertainties within the adopted framework and do not include the full systematic uncertainty of the proxy assumptions. The dashed line in the left panel marks the thermal timescale of a standard thin accretion disk at 20 R_g for $\alpha = 0.1$ (A. R. King et al. 2007), shown for order-of-magnitude reference only. The dominant feature of both panels is that the overwhelming majority of the sample—including all rejected candidates—shows no coherent variability, consistent with the low theoretical incidence of periodically detectable AGN (M. Volonteri et al. 2009).

thin accretion disk

$$t_{\rm th}(R) = \alpha^{-1}\, t_{\rm orb}(R) = \frac{2\pi}{\alpha}\left(\frac{R}{R_g}\right)^{3/2}\frac{GM_{\rm BH}}{c^3}, \tag{27}$$

evaluated at $R = 20\,R_g$ and adopting $\alpha = 0.1$ (A. R. King et al. 2007). Most sources lie to the right of the $t_{\rm th}(20\,R_g)$ reference line (i.e., at $P_{\rm rest} > t_{\rm th}$ for fixed proxy mass), and the candidate occupies a similar region. Because the masses shown here are proxy values derived from the heuristic framework described in Section 4.3, rather than direct per-object measurements, this comparison is intended only as an order-of-magnitude illustration and is not used to infer a specific physical mechanism in the single candidate. The right panel shows no obvious trend between rest-frame period and proxy Eddington ratio, and is similarly treated as illustrative only. The figure might be compatible with scenarios in which variability on timescales longer than the local thermal time may be associated with global disk structure or geometric distortions such as eccentricity or tilt (H. Deng 2026; N. Kaaz et al. 2026), but we do not draw strong conclusions about the physical origin of any individual detection given the commissioning-stage nature of the present dataset. More broadly, coherent variability on these spatial and temporal scales is compatible with mechanisms such as orbiting inhomogeneities (J. D. Schnittman & E. Bertschinger 2004), spiral density waves (M. Tagger & R. Pellat 1999), or quasiperiodic modulation of the local accretion rate (P. Arévalo et al. 2024). For completeness, Figure 12 indicates where the present candidate sample falls in the proxy mass–period plane relative to theoretical regions discussed in connection with ultralight-boson superradiance models. In such scenarios, ultralight bosons may form long-lived bound clouds around spinning black holes, and characteristic timescales could indirectly influence the surrounding accretion flow (R. Brito et al. 2015; C. Ünal et al. 2021). Mapping the observed periods to an

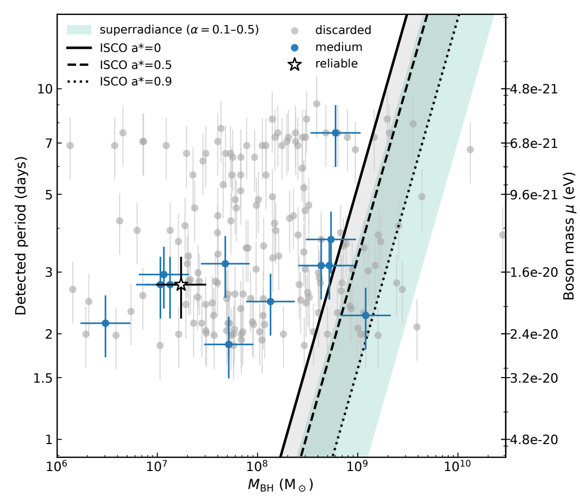

Figure 12. Mass–period diagram illustrating possible avenues for AGN studies in the LSST era. Each point is a candidate object with oscillatory signal identified with the QhX pipeline: the candidate (star), medium (blue), and discarded candidates (gray). Error bars are propagated internal uncertainties within the adopted mass framework. Black curves show the orbital period at the innermost stable circular orbit (ISCO) for black hole spins $a_* = 0, 0.5, 0.9$. The light-blue shaded region marks the range $\alpha \in [0.1, 0.5]$, often discussed in connection with efficient ultralight-boson superradiance, and the right-hand axis shows the corresponding illustrative boson-mass scale μ obtained by mapping the observed period to the Compton relation. The figure is intended only as a comparison with selected theoretical timescale loci and broader applications beyond standard AGN studies, and should not be interpreted as evidence for any specific nonstandard scenario (see R. Brito et al. 2015; C. Ünal et al. 2021; P. Du et al. 2022).

approximate boson-mass scale using the Compton relation,

$$\mu = \frac{h}{T\,c^2}, \tag{28}$$

yields $\mu \sim 10^{-21}$–10^{-20} eV, while the corresponding dimensionless gravitational coupling,

$$\alpha_g = \frac{G\,M_{\rm BH}\,\mu}{\hbar\,c^3}, \tag{29}$$

falls in the range $\alpha_g \sim 0.1$–0.5 for proxy masses $M_{\rm BH} \sim 10^7$–$10^9\,M_\odot$. We stress that this mapping is purely illustrative: it combines proxy masses with observed characteristic periods and is not presented as evidence for any exotic-field interpretation of the present candidate sample. At present, no statistically significant clustering is apparent in either diagram, consistent with the predominantly near-null character of the current DP1 result. The sole purpose of Figures 11 and 12 is to illustrate how future LSST population studies may eventually place coherent variability candidates within broader physical parameter spaces once longer baselines and much larger samples become available. Such analyses could serve as a complementary phenomenological probe alongside more direct searches for ultralight bosonic fields through tidal disruption events (TDEs; see P. Du et al. 2022) and black hole spin demographics (R. Brito et al. 2015; C. Ünal et al. 2021). We conclude by emphasizing that the near-null detection rate reported here—one strong coherent candidate among 1456 cross-matched and 428 quality-selected DP1 AGN—is not a null result in the pejorative sense. Rather, it provides an early empirical estimate of the detectable coherent variability incidence in commissioning-era Rubin LSST AGN data. Ruling out coherent variability across the broad AGN population is as constraining as finding it: a substantially higher prevalence of detectable coherent variability would likely have produced a larger number of robust candidates than observed in the present DP1 baseline. We also note that A. D. Hincks et al. (2026) demonstrated for the confirmed SMBHB candidates PKS 2131−021 and PKS J0805−0111 that sinusoidal variability dominating their radio and millimeter light curves yields only marginal optical periodicity detections when optical data are analyzed independently (GLS power = 0.22, global p-value = 0.937), because stronger short-term red-noise variability at optical wavelengths suppresses the coherent signal. This demonstrates that the near-null coherent variability detection rate in our LSST DP1 optical sample is physically consistent with expectations from multiwavelength studies of known periodic AGN, and further motivates combining future LSST optical data with radio monitoring to search for and confirm SMBHB candidates—including objects currently classified as `pseudoperiodic`, `transiently coherent` in the QhX catalog, which should not be dismissed a priori from multiwavelength follow-up, as their marginal optical detections may correspond to genuine periodicities more clearly revealed at radio wavelengths.

As multiyear LSST data releases accumulate, the `QhX` catalog will be recomputed on progressively longer baselines, transforming this early prototype into a statistically robust, population-level census of coherent, pseudocoherent, and stochastic AGN variability across orders of magnitude in sample size.

4.5. External Crossmatch with Blazar and Radio-loud AGN Catalogs

We performed an external catalog crossmatch to identify QhX-catalog sources that are already known blazars or radio-loud AGN, since blazars can exhibit strong quasiperiodic variability signatures associated with jet-related processes (C. M. Raiteri 2025). The `QhX` methodology is general and can be applied to variable AGN regardless of whether the dominant variability arises from the accretion flow, the jet, or a combination of both.

We cross-matched the `QhX` catalog against several widely used blazar and radio-selected reference catalogs, including BZCAT5 (E. Massaro et al. 2015), 4LAC (M. Ajello et al. 2020), CRATES (S. E. Healey et al. 2007), CGRaBS (S. E. Healey et al. 2008), BROS (R. Itoh et al. 2020), 3HSP (Y.-L. Chang et al. 2019), WIBRaLS2 (R. D'Abrusco et al. 2019), KDEBLLACS (R. D'Abrusco et al. 2019), the Atacama Large Millimeter/submillimeter Array (ALMA) calibrator catalog (A. Paggi et al. 2020), the SPT-SZ survey (W. B. Everett et al. 2020), and the Radio Fundamental Catalog (RFC; L. Y. Petrov & Y. Y. Kovalev 2025). Positional associations were performed using a cone-search radius of 0.″1, with additional inspection out to 1.″0 to account for possible small astrometric offsets; no additional counterparts or discrepancies were found at larger separations.

Only a small number of blazar matches were identified. The source J062032-251521 (CRATES designation) was consistently cross-identified in the radio-selected catalogs BZCAT5, CRATES, CGRaBS, BROS, and RFC. Additional isolated matches were found in 3HSP (3HSP J035532.7−485134), WIBRaLS2 (ATLAS S336 2), and KDEBLLACS (CDFS J03324-2741B). No counterparts were found in 4LAC, the ALMA calibrator catalog, or the SPT-SZ survey within either matching radius.

Overall, the low overlap fraction indicates that most `QhX` catalog sources are not previously identified blazar-dominated or radio-selected AGN, confirming that the detected quasi-coherent variability population is not dominated by jet-selected sources. Nevertheless, the small number of blazar matches identified here demonstrates that the `QhX` framework can detect coherent oscillatory signatures across the full AGN population, including jet-dominated sources. Characterizing quasiperiodic variability in blazars with future longer-baseline LSST data may prove equally valuable as in more numerous AGN varieties, since blazars potentially offer an unambiguous face-on view of the inner accretion disk, whereas orientation effects are more difficult to disentangle in other classes of AGN.

4.6. Role of the QhX Catalog in Supporting Rubin LSST AGN Community Follow-up Efforts

Tools and derived catalogs such as `QhX`, together with complementary efforts within the Rubin LSST AGN Science Collaboration, are expected to provide an important discovery and triage layer for rare time-domain AGN sources through targeted spectroscopic, spectropolarimetric, high-cadence photometric, and multiwavelength follow-up campaigns.

One particularly important application concerns searches for supermassive black hole binaries (SMBHBs). Quasars exhibiting optical periodicity may be associated with emission from accretion disks around such systems. The most compact SMBHBs can exhibit ultrashort periods of days to weeks (C. Xin & Z. Haiman 2021). Such systems may enter the Laser Interferometer Space Antenna (LISA) band within the next decade (C. Xin & Z. Haiman 2021), underscoring the synergy between time-domain and gravitational-wave searches. These ultracompact systems are expected to be rare, typically at the low-mass end of the SMBHB population ($M \sim 10^4$–$10^7\, M_\odot$), with at most ∼150 predicted across the full LSST AGN catalog of ∼100 million quasars. They are therefore unlikely to be present in our DP1 sample. Nevertheless, this work lays the foundation for systematic searches for such ultrarare compact binaries in future LSST data releases, prior to the launch of LISA. Early identification of short-period candidates with our framework will also enable more targeted follow-up analyses, including more computationally intensive and shape-specific Bayesian periodicity methods (M. C. Davis et al. 2024; L. Bertassi et al. 2026; C. Xin et al. 2026), to further test the periodicity, search for electromagnetic chirp signatures, and help distinguish such signals from intrinsic red noise.

Nuclear transient phenomena may provide an alternative origin for some genuine day-scale periodicity in the present sample. Quasiperiodic eruptions (QPEs) are nuclear transients characterized by X-ray bursts on hour-to-day timescales (R. Arcodia et al. 2021), potentially triggered by the interaction of a stellar-mass object with an AGN disk (H. Tagawa & Z. Haiman 2023). To date, QPEs are predominantly observed in X-rays, while clear detections of strictly periodic UV/optical counterparts remain elusive (e.g., T. Wevers et al. 2025). In some systems, particularly those associated with TDEs, UV and optical emission is observed, but it appears to trace the broader transient evolution rather than the periodic QPE bursts themselves (M. Nicholl et al. 2024). Approximately $\sim 10^4$ galaxies are expected to host QPE-like systems exhibiting ultrashort periodic activity (R. Arcodia et al. 2024), derived by dividing the eROSITA volumetric QPE abundance ($\mathcal{R}_{\rm vol} \approx 0.6 \times 10^{-6}\,{\rm Mpc}^{-3}$, from four secure detections corrected for efficiency) by the local galaxy number density ($\sim 0.01\,{\rm Mpc}^{-3}$); this is a survey-depth-limited lower bound, as fainter or more distant hosts below the eROSITA completeness threshold ($\log L^{\rm peak}_{0.5-2.0\,{\rm keV}} > 41.7$) remain uncounted.

Recent theoretical studies further suggest that stellar-mass binaries embedded within AGN disks may produce short-duration bursts. Gas torques and dynamical interactions in the disk can trigger tidal disruptions of stars by stellar-mass compact objects, including black holes (K. Kremer et al. 2022; C. Xin et al. 2024). If detectable, such "micro-TDE" events may exhibit ultrashort periodic signatures on hour-to-day timescales, particularly in low-eccentricity encounters. Although only two candidates have been proposed so far, the predicted volumetric rate could reach $\sim 100\,{\rm Gpc}^{-3}\,{\rm yr}^{-1}$ (Y. Yang et al. 2022). A comparable ∼3 day X-ray modulation has recently been reported in NICER observations of a TDE and interpreted as arising from an inclined stellar impact onto an AGN disk, potentially leading to disk precession or transient geometric modulation (A. Malyali et al. 2026). While this detection concerns a different source class, it demonstrates that day-scale periodicities can naturally emerge in TDE–disk interaction scenarios.

5. Outlook: Future Improvements and Synergies

In this work, we constructed the `QhX` catalog of coherent and quasi-coherent oscillatory behavior in selected LSST DP1 AGN using the `QhX` pipeline. It should be regarded as a prototype, i.e., an initial demonstration of `QhX`-based searches for coherent oscillations in Rubin LSST data rather than a statistically complete census of periodic AGN. The temporal baseline and cadence of LSST DP1, spanning less than a full observing season for most fields, limit sensitivity to long-period or weakly coherent signals and introduce window-function aliases near integer multiples and fractions of the nightly sampling frequency. Future LSST data releases will enable repeated recalculation of the full catalog on progressively longer baselines, incorporating both the sources analyzed here and newly available AGN entering subsequent releases. This will allow systematic tracking of candidate behavior over time, including tests of signal persistence, phase stability, secular drift, and the appearance of new periodic candidates.

Early photometric calibration (Y. Choi et al. 2025) and instrumental systematics may also introduce low-amplitude or phase-dependent modulations in multiband light curves. Our use of postdetection per-band robust fitting and cross-band coherence criteria reduces sensitivity to such effects, while future Rubin data releases will benefit from refined global calibration and improved instrument characterization, further suppressing residual systematics.

The present postdetection vetting step adopts a single-harmonic (sinusoidal) model, a pragmatic choice for sparse commissioning data but not necessarily representative of the full complexity of AGN variability. Importantly, AGN and blazar light curves may be described by stochastic autoregressive processes, and even low-order linear models such as AR(2) can generate short-lived QPOs in the absence of an underlying periodic physical driver (T. J. Ulrych & T. N. Bishop 1975; A. V. Vecchia 1985; R. J. Adams & G. C. Goodwin 1995). Robust discrimination between stochastic quasi-periodicity and genuine periodic components therefore requires likelihood-based comparison between aperiodic autoregressive moving average (ARMA)/autoregressive fractionally integrated moving average (ARFIMA) models and their periodic extensions (periodic autoregressive moving average; PARMA), following both classical time-series methodology and modern software implementations (e.g., G. E. P. Box et al. 2015; R. Hyndman & G. Athanasopoulos 2018; E. D. Feigelson et al. 2018; C. Chatfield & H. Xing 2019), together with developments in periodic ARMA modeling (K. J. Bell et al. 2018). The limited baseline and cadence of LSST DP1 currently preclude statistically reliable fitting and model selection among such parametric time-domain models, which motivates our wavelet-based framework as an initial detection layer, analogous in spirit to earlier ARMA-based variability studies in astronomy (A. Stanislavsky et al. 2009; G. A. Caceres et al. 2019; M. Tarnopolski et al. 2020). As multiyear LSST light curves become available, the method of Z. Hu & H. Tak (2020), specifically designed for LSST-like irregular time series, can be incorporated into future releases of the `QhX` catalog as a complementary statistical validation layer.

An important future direction is the integration of the catalog with complementary AGN datasets. Cross-matching LSST candidates with spectroscopic surveys, reverberation-mapping samples, radio and X-ray monitoring programs, and future LSST-derived variability and classification catalogs will enable population studies linking coherent optical variability to $M_{\rm BH}$, accretion state, and multiwavelength behavior. Such synergies will be essential for distinguishing between competing physical scenarios and for placing LSST-detected oscillatory behavior into a broader astrophysical context.

More broadly, LSST-era surveys may open a new avenue for exploring possible indirect connections between galaxy interactions, angular momentum redistribution, and inner-disk variability. Future LSST data releases will make it possible to test such connections at population scale.

Another important direction for future development is to augment the current rule-based framework with supervised machine learning classifiers trained on labeled variability samples. In the present work, the selection criteria and class boundaries are intentionally kept transparent and physically interpretable, guided by injection–recovery tests and multi-band consistency diagnostics. Future LSST data releases, together with expanding samples of confirmed AGN and control sources, will enable a multistage framework in which learned classifiers complement the present coherence-based catalog construction. This would provide a machine learning classification layer while preserving the interpretability of the underlying `QhX` diagnostics.

Finally, as LSST sample sizes increase by orders of magnitude, statistically robust population studies will become possible, enabling quantitative discrimination between stochastic variability, accretion-disk dynamics, binary modulation, and other sources of coherent AGN variability.

Acknowledgments

A.B.K., D.I., L.Č.P., U.M., Đ.S., and S.S. acknowledge funding provided by the University of Belgrade - Faculty of Mathematics (the contract 451-03-33/2026-03/200104), Astronomical Observatory Belgrade (the contract 451-03-33/2026-03/200002) and the Faculty of Science of the University of Kragujevac (451-03-33/2026-03/200122) through the grants by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia. This material is based upon work supported in part by the International Contributions and Data Rights Agreement (IDRA) for Vera C. Rubin Observatory NO. R0013. The research leading to these results has received funding from the EU HORIZON-MSCA-2023-DN Project 101168906 "TALES: Time-domain Analysis to study the Life-cycle and Evolution of Supermassive black holes." M.F. acknowledges the financial support by Italian Ministry for Education University and Research (MIUR) grant PRIN 2022 2022383WFT "SUNRISE" and from the Time-domes grant within the "INAF 2023 Finanziamento della Ricerca Fondamentale." C.R. acknowledges support from SNSF Consolidator grant F01−13252, Fondecyt Regular grant 1230345, ANID BASAL project FB210003 and the China-Chile joint research fund. R.J.A. and S.S.-S were supported by FONDECYT grant No. 1231718 and by the ANID BASAL project FB210003. L.H.G. acknowledges financial support from ANID program FONDECYT Iniciación 11241477. M.C. acknowledges support by the European Union (ERC, MMMonsters, 101117624). S.P. is supported by the International Gemini Observatory, a program of NSF NOIRLab, which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the U.S. National Science Foundation, on behalf of the Gemini partnership of Argentina, Brazil, Canada, Chile, the Republic of Korea, and the United States of America. T.J. acknowledges the support of the European Union through NextGenerationEU grants RubinMining-uniri-iz-25-282 and AstroVar-uniri-iz-25-107. W.Y. acknowledges support from the Dunlap Institute for Astronomy & Astrophysics at the University of Toronto.

The authors acknowledge support of the HPC AI Platform of Office for Information Technologies and e-Government of Republic Serbia. The authors gratefully acknowledge the support of the Center for Advanced Computing and Modeling, University of Rijeka (Croatia), for providing supercomputing resources at HPC (High Performance Computing) Bura, and the Slovenian Research and Innovation Agency, grant P1-0031, the HPC RIVR consortium (http://www.hpc-rivr.si) and EuroHPC JU (http://eurohpc-ju.europa.eu) for funding this research by providing computing resources of the HPC system Vega at the Institute of Information Science (http://www.izum.si).

This work has made use of data from the European Space Agency (ESA) mission Gaia (https://www.cosmos.esa.int/web/gaia), processed by the Gaia Data Processing and Analysis Consortium (DPAC, https://www.cosmos.esa.int/web/gaia/dpac/consortium). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the Gaia Multilateral Agreement. This work is based in part on observations made with the NASA/ESA/CSA James Webb Space Telescope. The data were obtained from the Mikulski Archive for Space Telescopes at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-03127 for JWST. JADES DR5 includes NIRCam data from JWST programs 1176, 1180, 1181, 1210, 1264, 1283, 1286, 1287, 1895, 1963, 2079, 2198, 2514, 2516, 2674, 3215, 3577, 3990, 4540, 4762, 5398, 5997, 6434, 6511, and 6541, and MIRI data from JWST programs 1180, 1181, and 1207. We acknowledge the teams of programs 1895, 1963, 2079, 2514, 3215, 3577, 3990, 6434, and 6541 for developing their observing programs with zero exclusive-access periods. This work also makes use of observations taken by the 3D-HST Treasury Program (HST-GO-12177 and HST-GO-12328) with the NASA/ESA Hubble Space Telescope, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555.

Facilities: Rubin:Simonyi (LSSTComCam), Gaia, Sloan, Mayall (DESI), HST, JWST

Software: `LSST Science Pipelines` (NSF-DOE Vera C. Rubin Observatory 2025), `astropy` (Astropy Collaboration et al. 2013, 2018; Astropy Collaboration et al. 2022), `LSDB` (N. Caplar et al. 2025), `QhX` (A. B. Kovačević et al. 2026), `feets` (J. Cabral et al. 2018), `QSOGEN` (M. J. Temple et al. 2021), `TOPCAT` (M. B. Taylor 2005).

Appendix A
Workflow of the `QhX` Pipeline

The `QhX` pipeline is not intrinsically limited to LSST data. It requires irregularly or regularly sampled AGN light curves, preferably in multiple photometric bands, with sufficient temporal sampling to construct time–frequency representations and evaluate cross-band coherence. In this work, the pipeline is applied to LSST DP1 AGN light curves, but the same workflow can be applied to other time-domain survey data after adapting the input formatting, photometric uncertainties, and cadence-dependent thresholds. This workflow combines WWZ-based time–frequency analysis (G. Foster 1996), wavelet-domain coherence ideas (S. Mallat 2008; A. B. Kovačević

Table 4
General Workflow of the QhX Periodicity Detection Pipeline

Step	Description
1	Compute WWZ representations for each band or band pair
2	Compute coherence maps and derive $S(\omega)$
3	Identify candidate peaks and estimate significance
4	Estimate periods and uncertainties from peak widths
5	Apply internal consistency checks
6	Perform sinusoidal vetting per band
7	Compute amplitude, S/N, R^2, and phase lag
8	Assign object-level verdict
9	Apply FDR control across objects

Note. The method is applicable to irregularly or regularly sampled AGN light curves from LSST or other time-domain surveys, provided that the data have sufficient sampling for time–frequency and coherence-based analysis.

et al. 2018, 2020a), peak-localization and spectral-estimation diagnostics (W. H. Press et al. 1992; P. Stoica et al. 2005), red-noise-aware validation (S. N. Lahiri 2003; S. Vaughan 2005; S. Vaughan et al. 2016), and catalog-level FDR control (Y. Benjamini & Y. Hochberg 1995; Y. Benjamini 2010).

Table 4 provides a compact summary of the principal stages of the QhX pipeline, from period detection to object-level vetting and catalog-level false-discovery-rate control.

Appendix B
Fallback Proxy Period Uncertainties

When candidate-level lower and upper period errors are unavailable, malformed, or numerically unstable, we assign conservative fallback proxy uncertainties during object-level catalog construction. These quantities are used only for homogeneous reporting and do not replace the preferred peak-width-based uncertainties described in Section 3.2.

We first define the representative object-level S/N as

$$\mathrm{med\_snr} \equiv \mathrm{median}_b(\mathrm{amp\_snr}_b), \tag{B1}$$

where $\mathrm{amp\_snr}_b$ is the band-level amplitude S/N from Equation (17).

When med_snr is finite, we assign the fallback period uncertainty

$$\Delta P_{\mathrm{fb}} \equiv \frac{P_*^{\,2}}{T_{\mathrm{obs}}\ \mathrm{med\_snr}}, \tag{B2}$$

where P_* is the adopted period and T_{obs} is the effective observational baseline. This scaling is motivated by resolution-limited frequency estimation. For a finite baseline T_{obs}, the characteristic frequency resolution is $\delta f \sim 1/T_{\mathrm{obs}}$, and since $P = 1/f$, one obtains by first-order propagation $\Delta P \simeq |dP/df|\ \delta f = P^2\ \delta f$ (e.g., W. H. Press et al. 1992). Finite measurement noise degrades attainable precision approximately in inverse proportion to the S/N, consistent with classical bounds for frequency and period estimation (e.g., P. Stoica et al. 2005; J. T. VanderPlas 2018).

If med_snr is unavailable or if Equation (B2) yields an unphysical value, we impose the conservative floor

$$\sigma_P = \max\left[0.1\ P_*,\ 0.2\ \mathrm{day},\ 0.3\ \frac{T_{\mathrm{obs}}}{\max(1,\ T_{\mathrm{obs}}/P_*)}\right]. \tag{B3}$$

The first term enforces a minimum relative uncertainty of 10%, the second a minimum absolute uncertainty of 0.2 days, and the third a "few-cycles" floor that scales with the effective number of observed cycles, $N_{\mathrm{cyc}} \approx T_{\mathrm{obs}}/P_*$, inflating the uncertainty when N_{cyc} is small. Reported periodic candidates are generally quasiperiodic rather than strictly phase-coherent (M. C. Davis et al. 2024). In such systems, formal error estimates can become artificially small even when long-term phase stability is not physically supported by the data. A fractional period error $\Delta P/P$ accumulates into an approximate phase drift $\Delta\phi \sim N_{\mathrm{cyc}}(\Delta P/P)$ over N_{cyc} cycles. Thus $\Delta P/P \sim 0.1$ implies order-unity phase uncertainty after $\sim$10 cycles, equivalent to an effective quality factor $Q \equiv P/\sigma_P \sim 10$. This conservative floor therefore prevents reporting unrealistically high coherence in red-noise-dominated quasar light curves, while remaining subdominant to the few-cycle resolution term in short-baseline cases. In our sample, the distribution of consecutive time separations, $p(\Delta t)$, is bimodal, with a dense intranight component and an internight component near $\sim$1 day; the local minimum between these modes occurs at $\delta t_{\mathrm{gap}} \approx 0.1$–0.2 day. Period uncertainties below this cadence-transition scale are not constrained by independent cycle-to-cycle sampling, so we impose a cadence-based absolute uncertainty floor of

$$\sigma_P \geqslant c_t\ \delta t_{\mathrm{gap}}, \tag{B4}$$

and adopt $c_t = 1$, giving $\sigma_P \gtrsim 0.2$ day. In addition, the finite baseline gives a frequency resolution $\Delta f \sim 1/T_{\mathrm{obs}}$ and therefore

$$\Delta P \sim \frac{P_*^2}{T_{\mathrm{obs}}} = \frac{P_*}{N_{\mathrm{cyc}}}, \qquad N_{\mathrm{cyc}} \equiv \frac{T_{\mathrm{obs}}}{P_*}, \tag{B5}$$

so for few-cycle detections we include the conservative baseline-dependent term used in Equation (B3).

When fallback errors are used, we report symmetric lower and upper uncertainties

$$\Delta P_{\mathrm{lo}} = \Delta P_{\mathrm{up}} = \sigma_P. \tag{B6}$$

Appendix C
Detailed Object-level Statistics and Bucket Construction

For implementation and diagnostic purposes, we also compute an extended object-level summary vector

$$\begin{aligned}\boldsymbol{T}_{\mathrm{obj}} = (&n_{\mathrm{good}},\ n_{\mathrm{coherent}},\ \mathrm{med\_snr},\ \mathrm{med\_R2},\\ &\mathrm{max\_lag},\ \mathrm{support\_count},\ T_{\mathrm{obs}})\end{aligned} \tag{C1}$$

where n_{good} and n_{coherent} are defined in Equation (22), med_snr is given by Equation (B1), $\mathrm{med\_R2} = \mathrm{median}_b(R_b^2)$, $\mathrm{max\_lag} = \max_b |\Delta\phi_{b|\mathrm{ref}}|$, `support_count` is the number of statistically independent contexts in which the candidate is detected above threshold, and T_{obs} is the observational baseline. In the present implementation, an independent context corresponds to either a distinct photometric band or, when applicable, a separately analyzed nonoverlapping temporal segment.

We also define the coherence fraction

$$C_{\mathrm{obj}} \equiv \frac{n_{\mathrm{coherent}}}{\max(1,\ n_{\mathrm{good}})}, \tag{C2}$$

which is used only as a compact descriptive diagnostic.

To provide a statistical interpretation of these quantities, we consider a stochastic null hypothesis H_0 in which band-level triggers arise from noise. Under a stochastic null hypothesis H_0, let $p_{\rm good}$ denote the probability that a single photometric band produces a spurious single-band trigger, and let $p_{\rm coh}$ denote the probability that it also satisfies the stricter phase-coherence requirement, with $p_{\rm coh} \leqslant p_{\rm good}$. Assuming approximate independence between bands, the probability that at least one band produces a spurious trigger scales as $\sim B\, p_{\rm good}$, where B is the number of photometric bands. By contrast, the probability of obtaining two or more coherent detections by chance scales as $\sim\binom{B}{2} p_{\rm coh}^2$ for $p_{\rm coh} \ll 1$, and therefore decreases much more rapidly. This motivates the use of $n_{\rm good} \geqslant 1$ as a lower-confidence trigger and $n_{\rm coherent} \geqslant 2$ as the criterion for the strongest object-level class.

The medians are preferred over means because they are more robust to outlying or poorly sampled bands. Likewise, the statistic

$$\text{max\_lag} = \max_b |\Delta\phi_{b|\rm ref}| \tag{C3}$$

records the largest absolute phase disagreement among bands and therefore provides a conservative summary of cross-band coherence. This parameter records the largest absolute phase disagreement among bands and therefore provides a conservative summary of cross-band coherence. Since a coherent band is required to satisfy $|\Delta\phi_{b|\rm ref}| \leqslant 0.20$, values of max_lag > 0.20 indicate that at least one band lies outside the adopted phase-coherence window. This criterion is therefore used to identify candidates with local oscillatory support but unstable or inconsistent multiband phase alignment.

In the catalog-production code, these quantities are used to organize candidates into implementation-oriented diagnostic buckets. These buckets are auxiliary bookkeeping classes used to separate clearly supported candidates from marginal or noise-dominated cases. The bucket thresholds follow directly from the band-level trigger definitions and the multiband coincidence logic used in the main verdict rule. In particular, $n_{\rm coherent} \geqslant 2$ identifies candidates with repeated phase-consistent support across more than one band, $n_{\rm coherent} = 1$ marks cases with only partial coherent support, and $n_{\rm coherent} = 0$ with $n_{\rm good} \geqslant 1$ identifies candidates with single-band oscillatory evidence but no confirmed multiband coherence. The red-noise-dominated condition combines $n_{\rm good} = 0$ with median amplitude S/N and R^2 below the adopted single-band thresholds, indicating that the typical band does not show significant oscillatory support. A compact implementation-level bucket scheme is

$$\text{bucket} = \begin{cases} \text{reliable/coherent}, & n_{\rm coherent} \geqslant 2, \\ \text{transiently coherent}, & n_{\rm coherent} = 1, \\ \text{pseudoperiodic}, & \begin{array}{l} n_{\rm coherent} = 0, \\ n_{\rm good} \geqslant 1 \end{array}, \\ \text{red-noise dominated}, & \begin{array}{l} n_{\rm good} = 0, \\ \text{med\_snr} < 1.5, , \\ \text{med\_R2} < 0.10 \end{array} \\ \text{undefined}, & \text{otherwise}. \end{cases} \tag{C4}$$

Here, *reliable/coherent* denotes candidates with multiband coherence, *transiently coherent* denotes candidates with coherence in only one band, *pseudoperiodic* denotes candidates with at least one internally significant single-band detection but no coherent multiband support, and *red-noise dominated* denotes objects without significant detections and with weak statistical support. All remaining cases are classified as *undefined*.

Appendix D
Additional Tests on Simulated Data

We adopt the LSST AGN Data Challenge dataset as a conservative benchmark for algorithm testing. The cadence properties of these light curves, including seasonal gaps and irregular sampling, provide a stringent test of period detection performance (D. V. Savić et al. 2023). Since Rubin Observatory operations are expected to achieve improved cadence coverage, our results should be interpreted as lower-bound performance estimates rather than optimistic projections. The injected amplitude grid was designed to span physically motivated Doppler-boost variability regimes reported in the literature (see Table 5). M. Charisi et al. (2018) demonstrate that Doppler modulation becomes reliably detectable when peak-to-peak fractional flux variations exceed approximately 5%–10%. M. C. Davis et al. (2024) further parameterize simulated binary AGN signals using a magnitude semiamplitude A, defining "weak" binaries as $0.1 \leqslant A \leqslant 0.5$ mag and "strong" binaries as $0.5 < A < 1.0$ mag.

Following these definitions, we adopted a magnitude semiamplitude ladder spanning 0.1–0.9 mag. This range probes the transition from marginal detectability to high signal-to-noise Doppler modulation while avoiding unrealistically small amplitudes that are systematically undetectable under LSST-like cadence and noise conditions.

The upper bound of the long-period injection grid was determined empirically from the sampling properties of the real light curves. The cadence properties of the light curves are characterized using the distribution of time separations between consecutive observations, $\Delta t = t_{i+1} - t_i$, computed

Table 5
Injection Configuration and Cadence Constraints Used in the Simulated Detection Tests

Parameter	Value
Injected period ranges	Short: 5, 10, 15, 20, 30, 50 days; long: 120, 180, 220 days
Global cadence statistics	$\Delta t_{\rm med} \approx 1$ day, $\Delta t_{95} \approx 35$ days, $\Delta t_{\rm max} \approx 350$ days
Injected amplitude range ($A_{\rm mag}$)	0.12, 0.18, 0.25, 0.35, 0.45, 0.60, 0.90 mag
Flux modulation parameter	$\epsilon = \frac{10^{0.4A_{\rm mag}} - 1}{10^{0.4A_{\rm mag}} + 1}$
Injection model	Doppler-like sinusoidal modulation
Photometric bands	g, r, i
Phases per (P, A) cell	10
Injection assignment	One case per parent object
Minimum points per band	80
Early/late baseline fraction	20%/20%
Temporal bin width	200 days
Minimum nonempty bins	Three per band
Total injections (short-period grid)	540
Total injections (long-period grid)	180

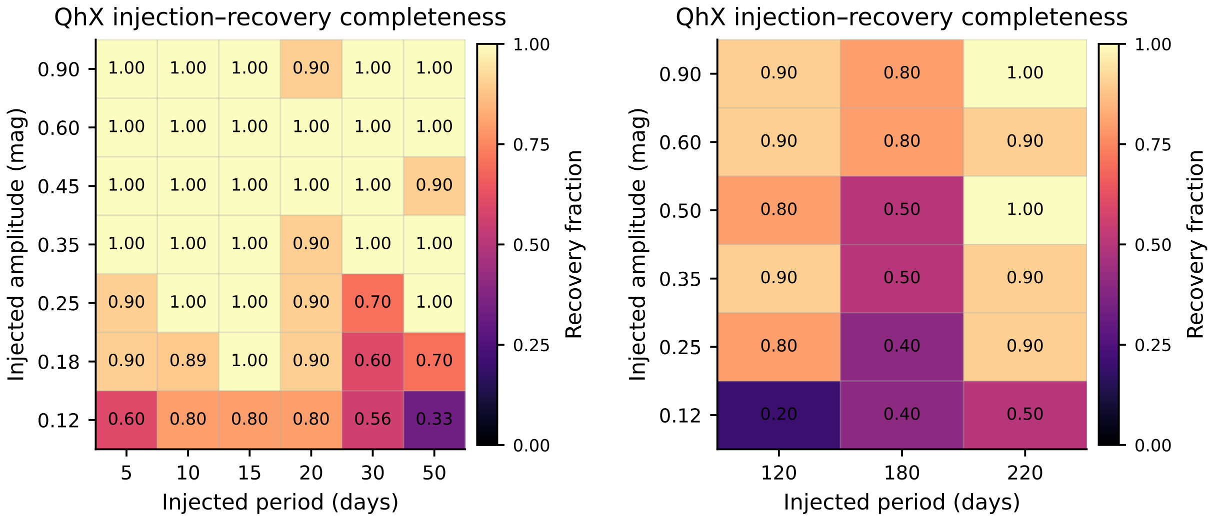

Figure 13. Injection–recovery completeness as a function of injected period (*x*-axis) and modulation amplitude (*y*-axis). Numbers indicate bin-wise recovery fractions.

across all objects and bands. We summarize this distribution using three robust statistics: the median sampling interval $\Delta t_{\rm med}$, the 95th percentile spacing Δt_{95}, and the maximum observed gap $\Delta t_{\rm max}$. For the AGN Data Challenge dataset we find $\Delta t_{\rm med} \approx 1.0$ day, $\Delta t_{95} \approx 35$ days, $\Delta t_{\rm max} \approx 350$ days.

The median cadence reflects the typical intraseason sampling, while the 95th percentile captures the onset of seasonal gaps. The maximum gap corresponds to full seasonal observing interruptions and sets a practical upper bound on coherent phase detection. Periods approaching this scale become increasingly degenerate with window-function aliases and produce unstable period estimates.

To remain within a cadence-stable regime, we conservatively limited the long-period injection grid to $P \leqslant 220$ days, corresponding to approximately $0.6\,\Delta t_{\rm max}$. Injection tests beyond this threshold exhibited rapidly increasing uncertainty widths and systematic period bias dominated by seasonal sampling rather than intrinsic signal properties.

Two complementary period regimes were therefore explored: a short-period grid (5–50 days), sensitive to compact systems and dense cadence, and a long-period grid consisting of $P = \{120, 180, 220\}$ days, probing progressively more challenging but still reliably sampled timescales.

Injected signals follow a Doppler-like sinusoidal modulation applied coherently across all *gri* observations of each object: $F'(t) = F(t)\,[1 + \epsilon\,\sin(2\pi(t - t_0)/P + \phi)]$, where P is the injected period, ϵ the modulation strength which enforces symmetry between flux maxima and minima in logarithmic magnitude space, ϕ a random phase drawn uniformly from $[0, 2\pi)$, and t_0 the first observation epoch. Each injection case is applied to a distinct parent light curve, preserving realistic LSST sampling patterns, photometric uncertainties, and seasonal gaps.

To ensure robust detection statistics and suppress cadence pathologies, parent light curves were filtered using three quality-control criteria:

1. At least 80 valid observations in each of the g, r, and i bands;
2. A minimum of 10 observations per band in both the early and late 20% segments of the survey baseline;
3. At least three nonempty temporal bins per band using a bin width of 200 days.

These criteria enforce broad baseline coverage and mitigate biases introduced by seasonal clustering or truncated observing windows. For each $(P, A_{\rm mag})$ grid point, ten independent phase realizations were injected. Injection cases were randomly permuted and assigned to distinct parent objects. This strategy statistically averages phase-dependent cadence effects while preserving the irregular temporal sampling of the survey data.

Figure 13 summarizes the recovery fraction as a function of injected period and amplitude. The left panel shows the short-period regime (5–50 days), where recovery completeness exceeds 90% for moderate-to-high amplitudes ($\gtrsim$0.25 mag) and remains above 70%–80% even at the lowest tested amplitudes. This demonstrates robust sensitivity to rapidly varying periodic signals under LSST-like cadence and noise conditions.

The right panel presents the long-period regime (120–220 days). High recovery fractions are maintained at large amplitudes, with completeness approaching unity at 220 day periods for amplitudes $\gtrsim$0.35 mag. A moderate decrease in recovery efficiency is observed near 180 days, corresponding to approximately half-year timescales. This behavior is expected given seasonal visibility gaps and cadence windowing effects, which reduce phase coverage and introduce aliasing near annual harmonics.

Importantly, the overall performance remains stable across both regimes, indicating that `QhX` preserves sensitivity from short-timescale oscillations through multimonth periodicities without requiring regime-specific tuning.

These tests confirm that the injection–recovery completeness trends of the detection pipeline are driven primarily by survey cadence and signal amplitude, rather than algorithmic failure modes, and provide a quantitative basis for interpreting detection efficiencies in the released periodicity catalog.

Appendix E
Red Noise Null Hypothesis Test

Because quasar variability is intrinsically correlated and often well described by red-noise or damped-random-walk-like processes (S. Vaughan 2005; C. L. MacLeod et al. 2010; S. Vaughan et al. 2016), we further evaluate the internally strong candidates against a red-noise null hypothesis by means of correlation-preserving surrogate light curves generated with a moving-block bootstrap (S. N. Lahiri 2003). At candidate frequency ω_k, the empirical band-level false-alarm probability is

$$\hat{p}_b(\omega_k) = \frac{1 + k_{\geqslant}}{1 + B}, \tag{E1}$$

where $k_{\geqslant}$ is the number of surrogates with a statistic at least as large as the observed one and B is the total number of surrogate realizations.

For multiband coherence we additionally compute the joint surrogate probability

$$p_{\mathrm{joint}} = \frac{1}{B}\sum_{i=1}^{B} 1(S_1^{(i)} \geqslant S_1^* \wedge \cdots \wedge S_{N_b}^{(i)} \geqslant S_{N_b}^*), \tag{E2}$$

where $1(\cdot)$ is the indicator function. These quantities are used as auxiliary diagnostics and not as replacements for the primary object-level verdict defined in Equation (23).

Appendix F
Catalog Organization and Access

The detailed technical documentation on the catalog organization and access is published by M. Stanković Cerović et al. (2025). Figure 14 summarizes the most important technical elements as follows. User interaction is performed exclusively through a web-based portal, which provides authentication, query construction, and visualization services. End users do not access the underlying database directly; all read operations are mediated through controlled interfaces to ensure consistency, provenance tracking, and long-term maintainability.

Catalog ingestion and updates are handled by an offline validation pipeline, which applies schema verification, value-range checks, and internal consistency tests prior to promotion into the core catalog tables. This separation between ingestion and access ensures that scientifically validated data products are exposed to users, while preventing accidental or unverified modifications to published records.

The catalog database contains both core tables, representing the primary per-object and per-band results, and derived tables that store secondary quantities such as variability metrics, coherence measures, and physical inferences. All tables exposed to users are versioned and read only.

Programmatic access is provided through Virtual Observatory (VO) services, including TAP and ConeSearch endpoints, compliant with International Virtual Observatory Alliance (IVOA) standards. These services enable interoperable queries using external tools such as `TOPCAT` and `pyVO`, facilitating integration with existing astronomical workflows and cross-survey analyses. The VO interfaces mirror the content of the public catalog tables and do not bypass the validation or access-control layers shown in Figure 14. The publication and hosting of the `QhX` catalog infrastructure were supported by the Regional Storage Support, through the LSST Science Preparation Kickstarter grants initiative.[48]

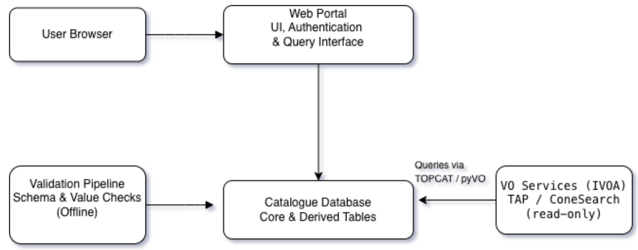

Figure 14. High-level organization and access model of the `QhX` catalog. User interaction is mediated through a web interface providing authentication and query controls. Catalog ingestion proceeds through an offline validation pipeline enforcing schema and value checks before promotion to core and derived tables. Read-only access is provided via IVOA-compliant TAP and ConeSearch services, enabling external queries through tools such as `TOPCAT` and `pyVO`.

Appendix G
Dawid–Skene Catalog Purity Estimation

We estimated the purity based solely on the expected fraction of genuinely periodic sources, without distinguishing between subclasses (e.g., periodic AGN, transiently coherent, or pseudoperiodic). We applied DS ensemble purity (A. P. Dawid & A. M. Skene 1979), a probabilistic expectation maximization (EM) model originally developed to estimate the accuracy of multiple imperfect classifiers in the absence of ground truth.

In our application, each LSST photometric band was treated as an independent *detector* that could vote either 'periodic" or nonperiodic" for a given source. A detector cast a positive vote only when all strict quality criteria were simultaneously satisfied: coherence flag set to `True`, $S/N \geqslant 1.5$, goodness of fit $R^2 \geqslant 0.10$, interband phase lag $\leqslant 0.1$ cycles (for nonreference bands), and relative period deviation $|\Delta P|/P \leqslant 10\%$.

The DS model assumes that each detector j is characterized by a true-positive rate $t_j = \Pr(v_{ij} = 1 \mid z_i = 1)$ and a false-positive rate $f_j = \Pr(v_{ij} = 1 \mid z_i = 0)$, where v_{ij} is the vote of detector j for object i, and $z_i \in \{0, 1\}$ represents the (unknown) true label of that object (periodic or not). It also estimates the prevalence of true positives in the sample, $\pi = \Pr(z_i = 1)$. The joint likelihood of the votes $\boldsymbol{v}_i$ across all detectors is written as

$$\begin{aligned}\Pr(\boldsymbol{v}_i|\pi,\, \{t_j, f_j\}) = \pi \prod_j\; t_j^{v_{ij}}(1 - t_j)^{1-v_{ij}} \\ + (1 - \pi)\prod_j f_j^{v_{ij}}(1 - f_j)^{1-v_{ij}}.\end{aligned} \tag{G1}$$

The model iteratively maximizes this likelihood by alternating two steps:

[48] Additional details are seen at https://lsst-sci-prep.github.io/kickstarter_grants.html.

E-step. Given the current detector parameters (π, t_j, f_j), the algorithm computes for every object a posterior probability

$$p_i = \Pr(z_i = 1|\boldsymbol{v}_i)$$
$$= \frac{\pi \prod_j t_j^{v_{ij}}(1 - t_j)^{1-v_{ij}}}{\pi \prod_j t_j^{v_{ij}}(1 - t_j)^{1-v_{ij}} + (1 - \pi) \prod_j f_j^{v_{ij}}(1 - f_j)^{1-v_{ij}}}. \quad \text{(G2)}$$

which represents the likelihood that the detected periodicity is genuine given the votes from all bands.

M-step. Using these posteriors, the algorithm updates the detector parameters:

$$\pi \leftarrow \frac{1}{N}\sum_i p_i, \quad \text{(G3)}$$

$$t_j \leftarrow \frac{\sum_i p_i\ \mathbb{I}[v_{ij} = 1]}{\sum_i p_i}, \quad \text{(G4)}$$

$$f_j \leftarrow \frac{\sum_i (1 - p_i)\ \mathbb{I}[v_{ij} = 1]}{\sum_i (1 - p_i)}. \quad \text{(G5)}$$

Here, $\mathbb{I}[\cdot]$ denotes the indicator function, which equals 1 if the condition in brackets is satisfied and 0 otherwise.

These updated values are then fed back into the E-step, and the process is repeated until the parameters converge (changes below a numerical tolerance of 10^{-6}).

In our implementation, all calculations are performed in log-space to ensure numerical stability, and missing or undefined votes are conservatively treated as negatives ($v_{ij} = 0$). The result is a per-object posterior probability p_i that quantifies the probability that the detected periodicity is real, as well as per-detector reliability parameters (t_j, f_j) that describe the consistency of each band.

The overall catalog purity is then defined as the mean posterior

$$P_{\rm DS} = \frac{1}{N}\sum_i p_i, \quad \text{(G6)}$$

which represents the expected fraction of true periodic sources in the strictly selected sample.

Appendix H Imaging of Detected Candidate

Figure 15 shows LSST DP1 (u, g, r, i, z, and y) cutouts of the `QhX` candidate. The figure is included as an image-plane validation of the source associated with the time-domain detection. The candidate is compact and consistently detected in the g, r, i, and z bands, with no obvious severe blending or imaging artifact at the catalog position. The u-band cutout is substantially shallower and the y-band cutout is noisier, consistent with the band-quality considerations discussed in the main text.

Archival high-resolution imaging is shown in Figure 16; it was obtained from the Hubble Space Telescope (3D-HST; G. B. Brammer et al. 2012; I. G. Momcheva et al. 2016) and the James Webb Space Telescope (JADES; M. A. C. Johnson et al. 2018). These data provide morphological context for the candidate host environment, which is unresolved in ground-based imaging.

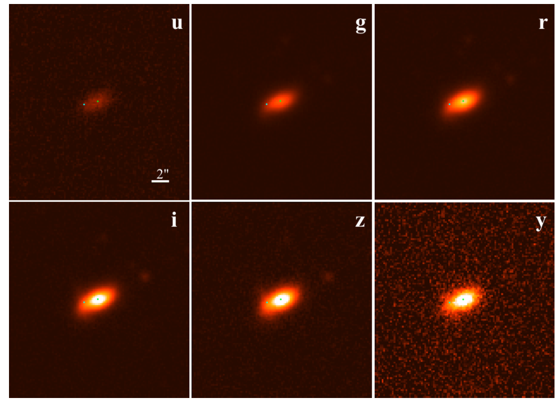

Figure 15. LSST DP1 u, g, r, i, z, and y cutouts of the discovered QhX candidate. The source is more prominent in the redder g, r, i, and z bands, where its morphology is compact and stable. The detected candidate is barely seen in the u band.

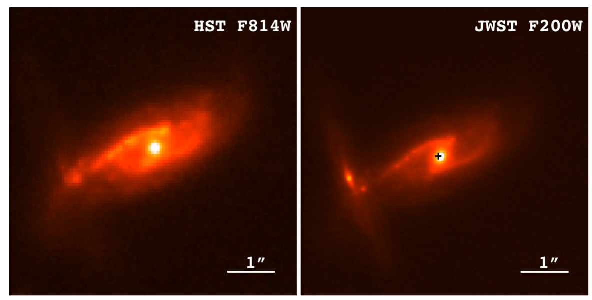

Figure 16. High-resolution imaging centered on the discovered candidate in the optical (HST F814W, left panel) and near-infrared (JWST F200W, right panel). Each image covers 6″ × 6″ in size.

ORCID iDs

Andjelka B. Kovačević https://orcid.org/0000-0001-5139-1978
Dragana Ilić https://orcid.org/0000-0002-1134-4015
Uroš Meštrić https://orcid.org/0000-0002-0441-8629
Milena Stanković Cerović https://orcid.org/0009-0002-7891-1755
Natale De Bonis https://orcid.org/0009-0001-1547-6648
Luka Č. Popović https://orcid.org/0000-0003-2398-7664
Saša Simić https://orcid.org/0000-0001-7453-2016
Marta Fatović https://orcid.org/0000-0003-1911-4326
Lorena Hernandez-García https://orcid.org/0000-0002-8606-6961
Donald P. Schneider https://orcid.org/0000-0001-7240-7449
Roberto J. Assef https://orcid.org/0000-0002-9508-3667
Swayamtrupta Panda https://orcid.org/0000-0002-5854-7426
Claudia M. Raiteri https://orcid.org/0000-0003-1784-2784
Christian Wolf https://orcid.org/0000-0002-4569-016X
Christopher A. Onken https://orcid.org/0000-0003-0017-349X
Neven Caplar https://orcid.org/0000-0003-3287-5250
Maria Charisi https://orcid.org/0000-0003-3579-2522
Claudio Ricci https://orcid.org/0000-0001-5231-2645
Chengcheng Xin https://orcid.org/0000-0003-3106-8182
Ilsang Yoon https://orcid.org/0000-0001-9163-0064
Mara Salvato https://orcid.org/0000-0001-7116-9303

Matthew J. Temple https://orcid.org/0000-0001-8433-550X
Sarath Satheesh-Sheeba https://orcid.org/0009-0003-0654-6805
Demetra De Cicco https://orcid.org/0000-0001-7208-5101
Maurizio Paolillo https://orcid.org/0000-0003-4210-7693
Paula Sánchez-Sáez https://orcid.org/0000-0003-0820-4692
Angela Bongiorno https://orcid.org/0000-0002-0101-6624
Weixiang Yu https://orcid.org/0000-0003-1262-2897
Akke Viitanen https://orcid.org/0000-0001-9383-786X
Bindu Rani https://orcid.org/0000-0001-5711-084X
Megan C. Davis https://orcid.org/0000-0001-9776-9227
Marcin Marculewicz https://orcid.org/0000-0002-1380-1785
Tomislav Jurkić https://orcid.org/0000-0002-4993-2939
Đorđe Savić https://orcid.org/0000-0003-0880-8963

References

Abramowicz, M. A., & Fragile, P. C. 2013, LRR, 16, 1
Abry, P., Gonçalvès, P., & Flandrin, P. 1995, in Wavelets and Statistics, ed. A. Antoniadis & G. Oppenheim (Springer), 15
Adams, R. J., & Goodwin, G. C. 1995, J. Time Ser. Anal., 16, 449
Ajello, M., Angioni, R., Axelsson, M., et al. 2020, ApJ, 892, 105
Amaro Seoane, P. 2026, A&A, 706, A315
Arcodia, R., Merloni, A., Buchner, J., et al. 2024, A&A, 684, L14
Arcodia, R., Merloni, A., Nandra, K., et al. 2021, Natur, 592, 704
Arévalo, P., Churazov, E., Lira, P., et al. 2024, A&A, 684, A133
Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, ApJ, 935, 167
Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33
Bauer, A., Baltay, C., Coppi, P., et al. 2009, ApJ, 696, 1241
Beers, T. C., Flynn, K., & Gebhardt, K. 1990, AJ, 100, 32
Belete, A. B., Bravo, J. P., Canto Martins, B. L., et al. 2018, MNRAS, 478, 3976
Bell, K. J., Hambleton, K. M., Lund, M. B., & Szabó, R. 2018, arXiv:1812.03142
Benjamini, Y. 2010, RSPSB, 72, 405
Benjamini, Y., & Hochberg, Y. 1995, RSPSB , 57, 289
Bertassi, L., Charisi, M., Buscicchio, R., et al. 2026, A&A, 711, A126
Bianco, F. B., Ivezić, Ž., Jones, R. L., et al. 2022, ApJS, 258, 1
Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. 2015, Time Series Analysis: Forecasting and Control (5th ed.; Wiley), 720
Brammer, G. B., van Dokkum, P. G., Franx, M., et al. 2012, ApJS, 200, 13
Brito, R., Cardoso, V., & Pani, P. 2015, Superradiance: New Frontiers in Black Hole Physics (Springer), 1
Burke, C. J., Shen, Y., Blaes, O., et al. 2021, Sci, 373, 789
Cabral, J., Sanchez, B., Ramos, F., et al. 2018, feets: feATURE eXTRACTOR for tIME sERIES, Astrophysics Source Code Library, ascl:1806.001
Caceres, G. A., Feigelson, E. D., Jogesh Babu, G., et al. 2019, AJ, 158, 58
Cackett, E. M., Horne, K., & Winkler, H. 2007, MNRAS, 380, 669
Caplar, N., Beebe, W., Branton, D., et al. 2025, arXiv:2501.02103
Caplar, N., Lilly, S. J., & Trakhtenbrot, B. 2017, ApJ, 834, 111
Carlin, J. L., Ferguson, P. S., Vivas, A. K., Caplar, N., & Malanchev, K. 2025, RNAAS, 9, 161
Chan, C.-H., Piran, T., & Krolik, J. H. 2024, ApJ, 973, 103
Chang, Y.-L., Arsioli, B., Giommi, P., Padovani, P., & Brandt, C. H. 2019, A&A, 632, A77
Charisi, M., Bartos, I., Haiman, Z., et al. 2016, MNRAS, 463, 2145
Charisi, M., Haiman, Z., Schiminovich, D., & D'Orazio, D. J. 2018, MNRAS, 476, 4617
Chatfield, C., & Xing, H. 2019, The Analysis of Time Series: An Introduction with R (7th ed.; CRC Press/Taylor and Francis)
Chen, Y.-C., Liu, X., Liao, W.-T., et al. 2020, MNRAS, 499, 2245
Chen, Y.-J., Zhai, S., Liu, J.-R., et al. 2023, MNRAS, 527, 12154
Choi, Y., Olsen, K. A. G., Carlin, J. L., et al. 2025, ApJ, 992, 47
Cohen, L. 1995, Time-Frequency Analysis (Prentice Hall)
D'Abrusco, R., Álvarez Crespo, N., Massaro, F., et al. 2019, ApJS, 242, 4
Davis, M. C., Grace, K. E., Trump, J. R., et al. 2024, ApJ, 965, 34
Dawid, A. P., & Skene, A. M. 1979, AppSt, 28, 20
De Cicco, D., Cavuoti, S., Paolillo, M., et al. 2026, A&A, 707, A155
De Cicco, D., Zazzaro, G., Cavuoti, S., et al. 2025, A&A, 697, A204
De Rosa, A., Vignali, C., Bogdanović, T., et al. 2019, NewAR, 86, 101525
Deng, H. 2026, PhRvD, 113, 083021
DESI Collaboration, Abdul-Karim, M., Adame, A. G., et al. 2026, AJ, 171, 285
DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2024, AJ, 168, 58
D'Orazio, D. J., & Charisi, M. 2024, in Black Holes in the Era of Gravitational-Wave Astronomy, ed. M. Arca Sedda, E. Bortolas, & M. Spera (Elsevier), 379
Du, P., Egaña-Ugrinovic, D., Essig, R., Fragione, G., & Perna, R. 2022, NatCo, 13, 4626
El-Badry, K., Hogg, D. W., & Rix, H.-W. 2026, PASP, 138, 024102
Eracleous, M., Livio, M., Halpern, J. P., & Storchi-Bergmann, T. 1995, ApJ, 438, 610
Everett, W. B., Zhang, L., Crawford, T. M., et al. 2020, ApJ, 900, 55
Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J., & Zisserman, A. 2010, IJCV, 88, 303
Fatović, M., Ilić, D., KovaÄević, A. B., et al. 2025, A&A, 695, A208
Fatović, M., Palaversa, L., Tisanić, K., et al. 2023, AJ, 165, 138
Feigelson, E. D., Babu, G. J., & Caceres, G. A. 2018, FrP, 6, 80
Flesch, E. W. 2023, OJAp, 6, 49
Foster, G. 1996, AJ, 112, 1709
Fowlkes, E. B., & Mallows, C. L. 1983, JASA, 78, 553
Freeburn, J., Andreoni, I., de Soto, K. M., et al. 2025, ApJL, 994, L24
Galwey, N. 2025, The False Discovery Rate: Its Meaning, Interpretation and Application in Data Science (Wiley)
Gierlinski, M., Middleton, M., Ward, M., & Done, C. 2008, Natur, 455, 369
Graham, M. J., Djorgovski, S. G., Drake, A. J., et al. 2014, MNRAS, 439, 703
Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015, Natur, 518, 74
Healey, S. E., Romani, R. W., Cotter, G., et al. 2008, ApJS, 175, 97
Healey, S. E., Romani, R. W., Taylor, G. B., et al. 2007, ApJS, 171, 61
Hernández-García, L., Muñoz-Arancibia, A. M., Lira, P., et al. 2024, A&A, 692, A84
Hincks, A. D., Ma, X., Mróz, P., et al. 2026, A&A, 706, A206
Horne, J. H., & Baliunas, S. L. 1986, ApJ, 302, 757
Hu, Z., & Tak, H. 2020, AJ, 160, 265
Huber, P. J. 1964, AnnMS, 35, 73
Hyndman, R., & Athanasopoulos, G. 2018, Forecasting: Principles and Practice (2nd ed.; OTexts)
Ingram, A., Done, C., & Fragile, P. C. 2009, MNRAS, 397, L101
Itoh, R., Utsumi, Y., Inoue, Y., et al. 2020, ApJ, 901, 3
Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, ApJ, 873, 111
Jaccard, P. 1912, NewPh, 11, 37
Johnson, M. A. C., Gandhi, P., Chapman, A. P., et al. 2018, MNRAS, 484, 19
Jurić, M., Axelrod, T., Becker, A., et al. 2023, LSST Data Products Definition Document LSE-163, Vera C. Rubin Observatory
Kaaz, N., Liska, M., Ward, C., & Davelaar, J. 2026, ApJ, 1003, 161
Kasliwal, V. P., Vogeley, M. S., & Richards, G. T. 2015, MNRAS, 451, 4328
Kato, S. 2001, PASJ, 53, 1
Kelly, B. C., Bechtold, J., & Siemiginowska, A. 2009, ApJ, 698, 895
Kelly, B. C., Becker, A. C., Sobolewska, M., Siemiginowska, A., & Uttley, P. 2014, ApJ, 788, 33
Kelly, B. C., Treu, T., Malkan, M., Pancoast, A., & Woo, J.-H. 2013, ApJ, 779, 187
Kelly, B. J., Baker, J. G., Etienne, Z. B., Giacomazzo, B., & Schnittman, J. 2017, PhRvD, 96, 123003
Kelley, L. Z., Haiman, Z., Sesana, A., & Hernquist, L. 2019, MNRAS, 485, 1579
Kim, D.-W., Protopapas, P., Bailer-Jones, C. A., et al. 2014, A&A, 566, A43
King, A. R., Pringle, J. E., & Livio, M. 2007, MNRAS, 376, 1740
Kovacevic, A. 2024, Quasar harmonic eXplorer (QhX): Supporting LSST Data Classification of Periodic Phenomena, v1, Zenodo, doi:10.5281/zenodo.10561219
Kovačević, A. B., Ilić, D., Tošić, M., et al. 2026, JOSS, 11, 8163
Kovačević, A. B., Pérez-Hernández, E., Popović, L. Č, et al. 2018, MNRAS, 475, 2051
Kovačević, A. B., Popović, L. Č, & Ilić, D. 2020a, OAst, 29, 51
Kovačević, A. B., Popović, L. Č., Simić, S., & Ilić, D. 2019, ApJ, 871, 32
Kovačević, A. B., Yi, T., Dai, X., et al. 2020b, MNRAS, 494, 4069
Kozlowski, S. 2016, MNRAS, 459, 2787

Kozłowski, S. 2017, ApJS, 228, 9
Kramer, D., Gowanlock, M., Trilling, D., McNeill, A., & Erasmus, N. 2023, A&C, 44, 100711
Kremer, K., Lombardi, J. C., Lu, W., Piro, A. L., & Rasio, F. A. 2022, ApJ, 933, 203
Lahiri, S. N. 2003, Resampling Methods for Dependent Data (Springer)
Li, G., Assef, R. J., Brandt, W. N., et al. 2026, ApJ, 1000, 165
Lin, A., Charisi, M., & Haiman, Z. 2026, ApJ, 997, 316
Liu, T., Gezari, S., & Miller, M. C. 2018, ApJL, 859, L12
MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, ApJ, 721, 1014
Malanchev, K., DeLucchi, M., Caplar, N., et al. 2025, arXiv:2506.23955
Mallat, S. 2008, A Wavelet Tour of Signal Processing: The Sparse Way (3rd ed.; Academic)
Malyali, A., Rau, A., Baldini, P., et al. 2026, A&A, 708, A163
Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, PASP, 131, 018003
Massaro, E., Maselli, A., Leto, C., et al. 2015, Ap&SS, 357, 75
Momcheva, I. G., Brammer, G. B., van Dokkum, P. G., et al. 2016, ApJS, 225, 27
Moreno, J., Vogeley, M. S., & Richards, G. 2019, AAS Meeting , 233, 431.02
Musoke, G., Liska, M., Porth, O., van der Klis, M., & Ingram, A. 2022, MNRAS, 518, 1656
Nakoneczny, S. J., Graham, M. J., Stern, D., et al. 2025, ApJ, 992, 153
Nicholl, M., Pasham, D. R., Mummery, A., et al. 2024, Natur, 634, 804
NSF-DOE Vera C. Rubin Observatory 2025, Legacy Survey of Space and Time Data Preview 1, doi:10.71929/RUBIN/2570308
Nun, I., Protopapas, P., Sim, B., et al. 2015, arXiv:1506.00010
Paggi, A., Bonato, M., Raiteri, C. M., et al. 2020, A&A, 641, A62
Pancoast, A., Brewer, B. J., Treu, T., et al. 2014, MNRAS, 445, 3073
Paolillo, M., & Papadakis, I. 2025, NCimR, 48, 537
Pâris, I., Petitjean, P., Aubourg, É., et al. 2018, A&A, 613, A51
Park, K., Xin, C., Davelaar, J., & Haiman, Z. 2025, PhRvD, 111, 063011
Pavez-Herrera, M., Sánchez-Sáez, P., Hernández-García, L., et al. 2025, A&A, 696, A153
Petrecca, V., Papadakis, I. E., Paolillo, M., et al. 2026, A&A, 706, A10
Petrov, L. Y., & Kovalev, Y. Y. 2025, ApJS, 276, 38
Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, A&A, 641, A6
Popović, L. Č., Simić, S., KovaÄević, A., & Ilić, D. 2021, MNRAS, 505, 5192
Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical Recipes in C. The Art of Scientific Computing (Cambridge Univ. Press)
Raiteri, C. M. 2025, A&ARv, 33, 8
Richards, J. W., Starr, D. L., Butler, N. R., et al. 2011, ApJ, 733, 10
Rigamonti, F., Bertassi, L., Buscicchio, R., et al. 2025, A&A, 702, A242
Robnik, J., Bayer, A. E., Charisi, M., et al. 2024, MNRAS, 534, 1609
Rousseeuw, P. J., & Croux, C. 1993, JASA, 88, 1273
Rubin Observatory Science Pipelines Developers 2025, The LSST Science Pipelines Software: Optical Survey Pipeline Reduction and Analysis Environment PSTN-019, Vera C. Rubin Observatory
Sánchez-Sáez, P., Lira, P., Cartier, R., et al. 2019, ApJS, 242, 10
Sánchez-Sáez, P., Lira, P., Mejía-Restrepo, J., et al. 2018, ApJ, 864, 87
Sánchez-Sáez, P., Reyes, I., Valenzuela, C., et al. 2021, AJ, 161, 141
Savić, D. V., Jankov, I., Yu, W., et al. 2023, ApJ, 953, 138
Schnittman, J. D., & Bertschinger, E. 2004, ApJ, 606, 1098
Schwarzenberg-Czerny, A. 1991, MNRAS, 253, 198
Shen, Y., Richards, G. T., Strauss, M. A., et al. 2011, ApJS, 194, 45
Shin, M.-S., Sekora, M., & Byun, Y.-I. 2009, MNRAS, 400, 1897
Songsheng, Y.-Y., & Wang, J.-M. 2023, ApJ, 945, 89
Stanislavsky, A., et al. 2009, ApJ, 693, 1877
Stanković Cerović, M., Kovačević, A. B., Ilić, D., et al. 2025, QhX/AGN Catalogue Portal—Technical Documentation, v2, Zenodo, doi:10.5281/zenodo.17914817
Stoica, P., Moses, R. L., et al. 2005, Spectral Analysis of Signals (Pearson Prentice Hall)
Stone, Z., Shen, Y., Burke, C. J., et al. 2022, MNRAS, 514, 164
Storchi-Bergmann, T., Schimoia, J. S., Peterson, B. M., et al. 2017, ApJ, 835, 236
Storey-Fisher, K., Hogg, D. W., Rix, H.-W., et al. 2024, ApJ, 964, 69
Suberlak, K. L., Ivezić, Z., & MacLeod, C. 2021, ApJ, 907, 96
Tachibana, Y., Graham, M. J., Kawai, N., et al. 2020, ApJ, 903, 54
Tagawa, H., & Haiman, Z. 2023, MNRAS, 526, 69
Tagger, M., & Pellat, R. 1999, A&A, 349, 1003
Tarnopolski, M., Żywucka, N., Marchenko, V., & Pascual-Granado, J. 2020, ApJS, 250, 1
Taylor, M. B. 2005, ASPC, 347, 29
Temple, M. J., Hewett, P. C., & Banerji, M. 2021, MNRAS, 508, 737
Ulrych, T. J., & Bishop, T. N. 1975, RvGSP, 13, 183
Ünal, C., Pacucci, F., & Loeb, A. 2021, JCAP, 2021, 007
Vanden Berk, D. E., Wilhite, B. C., Kron, R. G., et al. 2004, ApJ, 601, 692
VanderPlas, J. T. 2018, ApJS, 236, 16
Vaughan, S. 2005, A&A, 431, 391
Vaughan, S., Edelson, R., Warwick, R. S., & Uttley, P. 2003, MNRAS, 345, 1271
Vaughan, S., Uttley, P., Markowitz, A. G., et al. 2016, MNRAS, 461, 3145
Vecchia, A. V. 1985, Technometrics, 27, 375
Vera C. Rubin Observatory Team, Acero-Cuellar, T., Acosta, E., et al. 2026, AJ, 171, 360
Vio, R., Cristiani, S., Lessi, O., & Salvadori, L. 1991, ApJ, 380, 351
Volonteri, M., Miller, J. M., & Dotti, M. 2009, ApJ, 703, L86
Wang, J.-M., Xu, D.-W., Cao, X., et al. 2024, ApJ, 970, 85
Wasserstein, R. L., & Lazar, N. A. 2016, Am. Stat., 70, 129
Wevers, T., Guolo, M., Lockwood, S., et al. 2025, ApJL, 980, L1
Witt, C. A., Charisi, M., Taylor, S. R., & Burke-Spolaor, S. 2022, ApJ, 936, 89
Wójtowicz, A., Werner, N., Stawarz, Ł., & Cheung, C. C. 2026, A&A, 708, A2
Xin, C., & Haiman, Z. 2021, MNRAS, 506, 2408
Xin, C., & Haiman, Z. 2024, MNRAS, 533, 3164
Xin, C., Haiman, Z., Perna, R., Wang, Y., & Ryu, T. 2024, ApJ, 961, 149
Xin, C., Isi, M., Farr, W. M., & Haiman, Z. 2026, ApJ, 999, 149
Xu, J., Hu, S., Chen, X., & Huang, S. 2025, ApJ, 984, 45
Yang, Y., Bartos, I., Fragione, G., et al. 2022, ApJL, 933, L28
Yu, W., Richards, G. T., Ruan, J. J., et al. 2025, ApJ, 992, 130
Yu, W., Richards, G. T., Vogeley, M. S., Moreno, J., & Graham, M. J. 2022, ApJ, 936, 132
Yu, W., Ruan, J. J., Burke, C. J., et al. 2026, ApJ, 998, 144
Zhang, X. 2022, MNRAS, 516, 3650
Zhang, X. 2025, A&A, 703, L4