Version June 17, 2026
Preprint typeset using LaTeX style openjournal v. 09/06/15

BLACK HOLE STARS ACROSS THE UNIVERSE: IDENTIFYING CENTRAL ENGINE DOMINATED LITTLE RED DOTS AT $z \sim 1.5 - 9.5$

Andrea Weibel[1,*], Rohan P. Naidu[2,3,*,†], Pascal A. Oesch[1,4,5], Anna de Graaff[6,7,‡], Raphael E. Hviding[7], Zhaoran Liu[2], Jorryt Matthee[8], Christina C. Williams[9,10], Gabriel Brammer[4,5], Alba Covelo Paz[1], Jenny E. Greene[11], Christian Kragh Jespersen[11], Zhiyuan Ji[10], Michael V. Maseda[12], David J. Setton[11], Wendy Q. Sun[2], Alberto Torralba[8], Callum Witten[1], and Mengyuan Xiao[1]

[1] Department of Astronomy, University of Geneva, Chemin Pegasi 51, 1290 Versoix, Switzerland
[2] MIT Kavli Institute for Astrophysics and Space Research, 70 Vassar Street, Cambridge, MA 02139, USA
[3] Institute for Astronomy, University of Hawai'i, 2680 Woodlawn Drive, Honolulu, HI 96822, USA
[4] Cosmic Dawn Center (DAWN), Copenhagen, Denmark
[5] Niels Bohr Institute, University of Copenhagen, Jagtvej 128, København N, DK-2200, Denmark
[6] Center for Astrophysics, Harvard & Smithsonian, 60 Garden St, Cambridge, MA 02138, USA
[7] Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany
[8] Institute of Science and Technology Austria (ISTA), Am Campus 1, A-3400 Klosterneuburg, Austria
[9] NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, USA
[10] Steward Observatory, University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721, USA
[11] Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ 08544, USA and
[12] Department of Astronomy, University of Wisconsin-Madison, Madison, WI 53706, USA

Version June 17, 2026

ABSTRACT

Photometric selections of Little Red Dots (LRDs) largely rely on identifying their "V-shaped" spectral energy distribution (SED). Recent work suggests this V-shape stems from a combination of a central engine – also referred to as a Black Hole Star (BH*) – and a star-forming host galaxy. We present a new and highly complementary photometric selection that is based on incorporating BH* templates in the `eazy` redshift fitting code. Selecting compact sources where a BH* template contributes $> 80\%$ to the best fitting SED in the rest-optical, we compile a sample of 241 BH*-dominated candidates from $\sim 1000\,\mathrm{arcmin}^2$ of legacy and pure parallel JWST imaging. Our selection does not require a blue UV-component, and it successfully identifies objects that resemble the paradigmatic sources "MoM-BH*-1" and "The Cliff". We find that BH*-dominated sources exist across a wide range of redshifts ($z \sim 1.7 - 9.3$) and optical luminosities ($\log(L_{5100}/\mathrm{erg\,s^{-1}}) \sim 42 - 44.5$), and we measure a median Balmer break strength of ~ 3, with some breaks reaching values > 10. We estimate bolometric luminosities in the range $\log(L_{\mathrm{bol}}/\mathrm{erg\,s^{-1}}) \sim 42 - 45$, which, assuming accretion at the Eddington-limit, would translate to black hole masses of $M_{\mathrm{BH}} \sim 10^4 - 10^7 \mathrm{M}_\odot$, spanning the intermediate mass black hole to the quasar regime. The number density of BH*-dominated candidates peaks at $z \sim 5 - 6$ ($\sim 10^{-5}\,\mathrm{Mpc}^{-3}$) and it declines by an order of magnitude down to $z \sim 2$. Tentatively, comparing to V-shaped LRD samples suggests that the fraction of BH*-dominated sources among the broader LRD population does not decrease towards lower redshift. Crucially, our work demonstrates that BH*-dominated sources are not merely an early-Universe phenomenon but rather persist at least until cosmic noon.[a]

1. INTRODUCTION

Every massive galaxy in the local Universe hosts a supermassive black hole (SMBH; Kormendy & Ho 2013). It is well established that feedback from SMBHs has a significant impact on galaxy evolution throughout the history of the Universe (e.g., Weinberger et al. 2018), in particular through regulating star formation and quenching (e.g., De Lucia et al. 2025). However, the formation and growth of SMBHs in the early Universe is still poorly understood (e.g., Volonteri et al. 2021). Thanks to its powerful near-infrared (NIR) capabilities, the James Webb Space Telescope (JWST) is revealing the early phases of SMBH evolution. Using grism data from the FRESCO (Oesch et al. 2023) and EIGER (Kashino et al. 2023) surveys, Matthee et al. (2024) identified a surprisingly large number of broad-line Hα emitters at redshifts of $z \sim 4 - 6$, which they interpreted as being powered by active galactic nuclei (AGN). Motivated by the finding that the broad-line emitters showed red NIRCam colors and point-like morphologies they named them Little Red Dots (LRDs). LRDs have since become one of the most hotly debated discoveries of JWST. Apart from their broadened Balmer lines, they show various features that are not typically associated with AGN: weak X-ray (Ananna et al. 2024; Yue et al. 2024) and radio emission (Latif et al. 2025); a flattening or turnover in the NIR spectral energy distribution (SED), as probed by MIRI (e.g., Pérez-González et al. 2024; Williams et al. 2024) as well as weak far-infrared (FIR) emission (e.g., Xiao et al. 2025; Setton et al. 2025a), inconsistent with expectations from an AGN dust torus; a lack of variability (e.g., Kokubo & Harikane 2025; Liu et al. 2026b); Balmer absorption on top of the line emission (e.g., Matthee et al. 2026); high Balmer decrements (e.g., Nikopoulos et al. 2025; Torralba et al. 2026b; Sun et al. 2026); strong Balmer breaks (e.g., Wang et al. 2024; Naidu et al. 2025; de Graaff et al. 2025a); and "V-shaped" SEDs (e.g., Pérez-González et al. 2024; Kocevski et al. 2025) that inflect around the Balmer limit (H_∞, Setton et al. 2025b).

The difficulty of modeling the full SED of LRDs with standard AGN models (e.g., Ma et al. 2025b) has led to

*E-mail: andrea.weibel@unige.ch, rnaidu@hawaii.edu
† NASA Hubble Fellow, Pappalardo Fellow
‡ Clay Fellow
[a] A machine-readable catalog of our sample is available at: https://doi.org/10.5281/zenodo.20611334

the proposal of a wide variety of alternative scenarios. A broader class of models interprets LRDs as a central ionizing source (typically an accreting black hole) that is enshrouded by a dense gas envelope or cocoon (e.g., Naidu et al. 2025; Kido et al. 2025; Rusakov et al. 2026; Umeda et al. 2025; Liu et al. 2026a). Re-processing of the light in this gas envelope may account for both the Balmer break (Inayoshi & Maiolino 2025; Naidu et al. 2025; Ji et al. 2025), as well as the broadening of the Balmer lines (e.g., Chang et al. 2026; Naidu et al. 2025; Sneppen et al. 2026). This scenario bears similarities with so-called quasi-stars (Begelman et al. 2008), as explored in e.g., Begelman & Dexter (2026); Santarelli et al. (2026). Other models and scenarios include super-Eddington accretion (e.g., Lupi et al. 2024; Liu et al. 2025; Secunda et al. 2026; Madau 2026; Madau & Maiolino 2026a), self-interacting dark matter (Grant Roberts et al. 2025; Jiang et al. 2026), primordial black holes (Huang et al. 2024; Zhang et al. 2025, 2026), direct collapse black holes (Pacucci et al. 2026), supermassive stars (Nandal & Loeb 2026), and globular clusters in formation (Chisholm et al. 2026). Alternatively, Madau & Maiolino (2026b) propose a unification scheme in which LRDs are the counterparts of blue AGN (see also Brazzini et al. 2026), viewed from a high-inclination angle, and with a specific gas and dust geometry. Most of these scenarios share the characteristic that LRDs are a combination of a compact, non-stellar engine and a surrounding host galaxy.

Photometric selections of LRDs largely focus on identifying point sources with V-shaped SEDs, either based on NIRCam-colors (e.g., Labbe et al. 2025; Greene et al. 2024; Kokorev et al. 2024; Pérez-González et al. 2024; Rinaldi et al. 2026), or by fitting power-laws to the rest-optical and the rest-UV (e.g., Kocevski et al. 2025). Using a spectroscopic sample of LRDs from RUBIES (de Graaff et al. 2025c), Hviding et al. (2025) showed that the vast majority of rest-optical point sources with V-shaped SEDs do show broad Balmer lines, suggesting that these three properties are linked, and that photometric LRD-selections successfully identify broad-line AGN. However, these selections crucially rely on significant detections in the rest-frame UV to measure a blue color or slope. Furthermore, there is nothing inherently physical about a V-shape. Billand et al. (2026) recently defined an "LRDness" that quantifies how compact and how strongly V-shaped the SED of a given source is. They found a continuum of "LRDness" across the galaxy population, with no distinct locus for LRDs.

Two recently discovered objects may contribute a key piece to the puzzle of understanding LRDs: "The Cliff" at $z_{\rm spec} = 3.55$ (de Graaff et al. 2025a), and the "MoM-BH*-1" at $z_{\rm spec} = 7.76$ (Naidu et al. 2025). They show red rest-optical SEDs that can be well-approximated by a blackbody with a temperature of $\sim$ 5000 K, broad Balmer lines, extremely strong Balmer breaks ($\sim 7-8$), and very little UV-emission. Naidu et al. (2025) combined the spectrum of the MoM-BH*-1 with that of a star-forming galaxy at the same redshift to show that this reproduces the typical V-shaped SED of LRDs, inspiring a picture where sources similar to The Cliff or the MoM-BH*-1 lie at the heart of every LRD. This scenario has been further explored by, e.g., de Graaff et al. (2025b), Barro et al. (2026), and Sun et al. (2026), showing that the diversity of LRDs can be explained as a diversity of host galaxies and central engines, and their relative contribution to the total SED (see also Pan et al. 2026; Mérida et al. 2026; Cloonan et al. 2026). However, this insight has yet to be incorporated in photometric LRD selections. Specifically, objects like The Cliff or the MoM-BH*-1 do not show a strong V-shape and are missing from current photometric samples because of their faintness in the rest-frame UV (Hviding et al. 2025). Identifying such objects is crucial because they let us probe the physics of the central engine without being impeded by the host galaxy. Following the terminology established in Naidu et al. (2025), de Graaff et al. (2025b), and Sun et al. (2026), we refer to the LRD engine as a black hole star (BH*) noting that the analysis in this paper is independent of the physics of the engine, as long as the decomposition of LRDs into an engine that resembles sources like The Cliff and the MoM-BH*-1, and a host galaxy is valid.

In this work, we present a new way of photometrically selecting LRDs as a composite of a BH* and a host galaxy. This method is naturally more sensitive to BH*-dominated sources, and does not require detections in the rest-UV. The goal of this selection method is to enlarge our sample of BH*-dominated sources akin to objects like The Cliff and the MoM-BH*-1. As such, it is complementary to existing photometric LRD samples, and does not itself deliver a complete sample of LRDs. It is essential to photometrically identify BH*-dominated sources as prime candidates for spectroscopic follow-up. Furthermore, understanding the redshift evolution in the number density of both BH*-dominated sources and V-shaped LRDs may inform us about how the formation and evolution of the LRD engines and their hosts are connected, and whether the engines are increasingly outshone by their host galaxies at later cosmic times.

The paper is structured as follows: In Section 2, we present the imaging data and photometric catalogs that form the basis of the subsequent analysis. We outline our sample selection in detail in Section 3, followed by a spectroscopic validation of the sample, and a comparison to V-shape selections in the literature. Section 4 provides an overview of the sample, highlighting the variety of identified sources, as well as outstanding candidates. We turn to sample properties in Section 5 where we investigate optical and bolometric luminosities, Balmer break strengths, as well as the number density evolution of BH*-dominated candidates. We discuss our findings, along with limitations and caveats of the sample selection in Section 6, and end with a summary and conclusions in Section 7.

Whenever relevant, we assume a ΛCDM cosmology with parameters from the Planck Collaboration et al. (2020) and we specify magnitudes in the AB-system (Oke & Gunn 1983).

2. IMAGING AND CATALOGS

This paper is based on the imaging data and photometric catalog presented in Weibel et al. (2026). Below, we briefly describe the data and catalog production, and refer the reader to Weibel et al. (2026) and references therein for further details.

2.1. *Imaging*

We retrieve imaging mosaics across legacy fields from the DAWN JWST Archive (DJA),[1] using version 7.0 or higher. The reduction of these mosaics starts from level 2 calibrated data products from the Mikulski Archive for Space Telescopes (MAST) using the software package `grizli` (Brammer 2018) as described in e.g., Valentino et al. (2023). Specifically, these legacy fields are the EGS, UDS, and COSMOS fields, the GOODS fields (North and South), and the Abell-2744 cluster field. A complete list of all JWST programs that contribute imaging data to these mosaics is provided in Section 2.1 of Weibel et al. (2026). JWST imaging is complemented with HST/ACS and HST/WFC3 imaging where available, most importantly from CANDELS (Grogin et al. 2011; Koekemoer et al. 2011). Beyond the legacy data from the DJA, we include pure parallel imaging from PANORAMIC (GO-2514, PIs Williams & Oesch, Williams et al. 2025), using the first data release.[2] PANORAMIC adds $\sim 250\,\mathrm{arcmin}^2$ of NIRCam imaging along 28 independent lines of sight, as well as additional depth and area to various legacy fields, so that we obtain a total area of $\sim 1000\,\mathrm{arcmin}^2$ in six or more NIRCam filters. To remove the brightest cluster galaxies from the Abell-2744 science mosaics before catalog construction, we subtract a running median filter with a box-size of 101×101 pixels ($4.04''\times 4.04''$; see Naidu et al. 2024).

2.2. *Photometry*

We create a photometric master catalog that contains sources from all the fields mentioned above following the methods outlined in Weibel et al. (2024) for the legacy fields, and Williams et al. (2025) for PANORAMIC. In short, we run `SourceExtractor` in dual-mode using an inverse-variance weighted stack of the F277W, F356W, and F444W images as the detection image. We then measure fluxes through circular apertures ($r = 0.16''$) on convolved versions of each available JWST and HST band, so as to match the PSF resolution in F444W. We scale these fluxes to total based on the Kron correction in the detection image, and an additional correction to account for flux beyond the Kron aperture based on the F444W PSF. The full master catalog contains 1,380,654 sources and forms the basis of the subsequent sample selection.

3. SAMPLE SELECTION AND VALIDATION

In the following, we outline the selection of a robust sample of BH*-dominated sources from our photometric catalogs, and validate the sample with publicly available spectra.

3.1. *Pre-Selection*

We start by pre-selecting compact sources from our photometric catalogs as c(F444W) < 1.7, where c(F444W) $= f_{\mathrm{F444W}}(r = 0.2'')/f_{\mathrm{F444W}}(r = 0.1'')$ is the ratio of the fluxes measured through circular apertures with radii of $0.2''$ and $0.1''$ in the F444W image. This broadly removes sources that are extended, following the definition of compactness for LRDs originally proposed in Labbe et al. (2025). Since compactness can only be robustly measured for sufficiently bright sources, we complement this with a magnitude cut as mag(F444W) < 27 (see, e.g., Hviding et al. 2025). We further require (1) SNR(F444W) > 10, (2) at least two additional filters with SNR > 3, and (3) data in at least the six NIRCam filters F115W, F150W, F200W, F277W, F356W, and F444W, concurrent with the minimal PANORAMIC setup (see Williams et al. 2025). This ensures a robust detection at the red end of the NIRCam wavelength range, and continuous photometric coverage from $1-5\,\mu$m. These cuts reduce the initial master catalog to 11,529 sources. We cross-match this reduced catalog with objects in the GAIA DR3 (Gaia Collaboration et al. 2016, 2023) that have a proper motion measured at SNR > 3, as well as a SNR > 3 photometric detection in the G-band, and that are neither flagged as a galaxy, nor as a QSO in the GAIA catalog. This yields 113 matches with secure bright stars which we remove from our catalog, leaving us with 11,416 compact objects.

3.2. *BH* Templates*

The key idea of our sample selection is to incorporate BH* templates in the `eazy` redshift fitting code (Brammer et al. 2008), and to identify objects whose best-fitting SED is dominated by those templates.

We start from the `blue_sfhz` template set (`sfhz` hereafter)[3] which contains 13 templates generated with the Flexible Stellar Population Synthesis (`FSPS`) code (Conroy et al. 2009; Conroy & Gunn 2010) with a broad range of redshift-dependent properties, and an additional template with extremely strong emission lines, based on a JWST/NIRSpec spectrum of a $z = 8.5$ galaxy (Carnall et al. 2023). We then complement this template set with one empirical BH* template at a time, re-running `eazy` for six different BH* templates that are described in the following.

Sun et al. (2026) decomposed LRD spectra into a host galaxy and a BH* component by subtracting the spectra of star-forming galaxies at the same redshift from LRD spectra, after matching them in [O III] luminosity. This assumes that the [O III] doublet originates from HII regions associated with star formation, and is backed up by the fact that the [O III] line widths are usually much narrower than the Balmer lines (see Sun et al. 2026 for details). They then stacked both the host galaxies and the BH*s to investigate their average properties. Here, we use three BH*-stacks that were created by binning the individual BH* components in luminosity as $\log(L_{5500}/\mathrm{erg\,s^{-1}}) < 43.2$, $43.2 < \log(L_{5500}/\mathrm{erg\,s^{-1}}) < 44.2$ and $\log(L_{5500}/\mathrm{erg\,s^{-1}}) > 44.2$. Interestingly, these stacks also show different effective temperatures, as derived by parameterizing their continuum shape with a blackbody, such that there is a positive correlation between luminosity and temperature ($T_{\mathrm{eff}} = 3505^{+182}_{-149}$, 4409^{+124}_{-170}, and $5122^{+32}_{-416}\,\mathrm{K}$ for the three luminosity bins respectively; see also de Graaff et al. 2025b). We extend the three stacks that cover $\lambda_{\mathrm{rest}} \lesssim 1\,\mu$m to longer wavelengths assuming a single blackbody of the respec-

[1] https://dawn-cph.github.io/dja/imaging/v7/
[2] https://panoramic-jwst.github.io/
[3] https://github.com/gbrammer/eazy-photoz/tree/master/templates/sfhz

tive temperature to obtain our first three templates that represent average BH* properties across a range of temperatures and luminosities.

Next, we include templates based on two paradigmatic BH*s, the MoM-BH*-1 (Naidu et al. 2025) and The Cliff (de Graaff et al. 2025a). Due to the high redshift of the MoM-BH*-1 ($z_{\rm spec} = 7.76$), the NIRSpec PRISM spectrum only covers $\lambda_{\rm rest} \lesssim 0.6\,\mu$m. We therefore use the best-fitting `CLOUDY` (Ferland et al. 2017) model from Naidu et al. (2025) for this source. Briefly, this model is selected to match the steep Balmer break, Balmer emission lines, and the overall SED shape from the UV to the NIR, providing a self-consistent template across the rest-wavelengths of interest for our search.

For The Cliff, we retrieve the public spectrum from the DJA which covers $\lambda_{\rm rest} \lesssim 1.15\,\mu$m. To extend the template to longer wavelengths, we fit a blackbody to the spectrum at $\lambda_{\rm rest} > 4000$Å. Masking the H$\alpha$ as well as the [O III] and Hβ emission lines, we find a best-fitting $T_{\rm eff} = 4635\,$K and we stitch the spectrum and the best-fitting blackbody together at $\lambda_{\rm rest} = 1.15\,\mu$m to create our fifth template.

Finally, we wish to include a template that represents the most extreme BH*s at the hot end of the temperature distribution. To this end, we add another empirical template that is based on a source tagged GN-9771, initially published as part of the NIRCam/grism sample in Matthee et al. (2024), and followed up with NIRSpec/IFU spectroscopy as presented in Torralba et al. (2026b). It is one of the most luminous LRDs known to date at $z_{\rm spec} = 5.5$ and its spectrum closely resembles that of A2744-45924 (Labbe et al. 2024) at $z_{\rm spec} = 4.47$, another extremely luminous LRD ($L_V \approx 10^{45}{\rm erg\,s^{-1}}$). Specifically, GN-9771 is characterized by strong Balmer emission lines, a forest of optical [Fe II] lines, and a sharp Balmer break. We use the host-subtracted spectrum of GN-9771 from Sun et al. (2026) which - compared to the other templates constructed so far - displays the brightest UV-emission with strong Fe II lines as well as He IIλ3203 Å, Mg IIλ2799 Å, and C III]λ1909 Å, and a downturn at $\lambda_{\rm rest} \lesssim 2000$ Å. This template thus accounts for the possibility that some UV-emission originates from the BH* as recently discussed in e.g., Ando et al. (2026). Indeed, the Sun et al. (2026) BH* templates allow for $\approx 0{-}50\%$ contribution in the UV from the central engine, with the brightest sources showing the highest UV fraction (see their Fig. 12). Similar to the procedure applied to the spectrum of The Cliff, we fit a single blackbody to the continuum at $0.4\,\mu{\rm m} < \lambda_{\rm rest} < 0.82\,\mu$m, masking prominent emission lines. We obtain a temperature of $T_{\rm eff} = 5410\,$K, and use this fit to extend the template to $\lambda_{\rm rest} > 1\,\mu$m. We discuss the inclusion of an additional template representing the cold end of the temperature distribution in Section 6.1.1.

Extending our template spectra using single blackbodies may not accurately represent the true continuum shape of BH*-dominated sources. For example, the source A2744-45924, with a similar spectrum to GN-9771 that we use as a template, shows a significant excess in MIRI relative to the expectation of a pure blackbody continuum (Setton et al. 2025a). However, this only affects our sample selection at the low redshift end because at $z \gtrsim 4$, the NIRCam data used in this work do not probe the extrapolated part of the templates. Further, the rest-optical continuum shapes of LRDs have been shown to be well described by blackbody continuum at least out to $\sim 1\,\mu$m (de Graaff et al. 2025b; Umeda et al. 2025). To avoid any impact of the noisy far UV part of our empirical templates on the selection, we further set all fluxes to 0 at $\lambda_{\rm rest} < 0.16\,\mu$m. We show our BH* templates in Figure 1.

3.3. *Template Fitting with* `eazy`

We run `eazy` six times on the pre-selected catalog, complementing the `sfhz` template set with a different BH* template each time. We do not include multiple BH* templates simultaneously to avoid potentially unrealistic combinations of different BH* templates and to limit the number of degrees of freedom in each fit. We allow the redshift to vary freely in the range $z \in (0.01, 20)$ and apply a noise floor at 5% of the flux in each filter to allow for additional flexibility. One concern in selecting BH*s from photometry is contamination by Milky Way stars, and in particular brown dwarfs (BDs), which can show strong breaks in their spectra and are known to contaminate photometric LRD samples (e.g., Greene et al. 2024). We use the internal `eazy` function `fit_phoenix_stars` to fit our candidates against a large grid of stellar templates that is composed of (1) the grid of PHOENIX templates available directly through `eazy`, (2) a grid of low-temperature, low-metallicity BDs from Meisner et al. (2021), (3) a set of 22 empirical BD templates based on NIRSpec fixed-slit PRISM spectra from GO-2302 (PI Cushing, Beiler et al. 2024), and (4) the Sonora Elf Owl grid (Mukherjee et al. 2024) consisting of 43,251 templates spanning a wide range of BD and giant exoplanet properties.

3.4. *Selection Cuts*

To obtain our sample of BH*-dominated candidates, we apply the following selection cuts to each of the six `eazy` runs corresponding to the six different BH* templates:

- The integrated BH* template contribution to the best-fitting (lowest χ^2) `eazy`-SED is $> 80\,\%$ between $\lambda_{\rm rest,\,min} = 0.4\,\mu$m and $\lambda_{\rm rest,\,max} = \min(1\mu{\rm m}\,, 5\mu{\rm m}/(1+z))$, i.e., red-wards of the Balmer break, and out to either $1\,\mu$m rest-frame (the range covered by all our empirical templates), or to the maximum rest-frame wavelength that is covered by NIRCam at the best-fitting redshift, if this is $< 1\mu$m.
- $1.5 < z_{\rm phot,\,eazy} < 9.5$
- The source is nominally better fit with the `sfhz`+BH* template set, than with the large grid of stellar templates, $\Delta\chi^2({\rm stars}) > 0$, as well as than with the `sfhz` template set only: $\Delta\chi^2(\texttt{sfhz}) > 0$.
- $\chi^2_{\rm eazy}/{\rm N_{filters}} < 5$ where $\rm N_{filters}$ is the number of available NIRCam and HST filters for each object.

For sources that satisfy all these criteria in more than one `eazy` run (i.e., with more than one of our BH* templates), we choose the fit with the lowest χ^2. Of the final

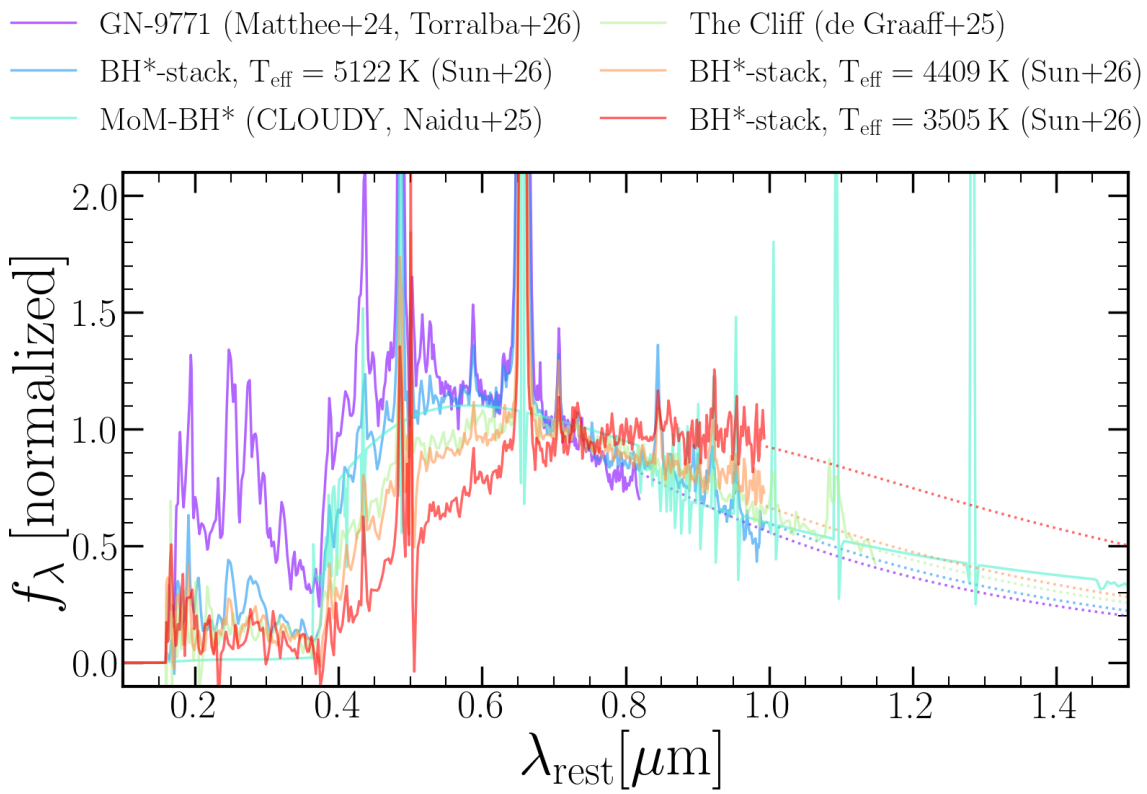

FIG. 1.— BH* templates used to select BH*-dominated candidates with `eazy`, normalized to the median in the range $0.7\,\mu{\rm m} < \lambda_{\rm rest} < 0.8\,\mu{\rm m}$. The extrapolated part of the templates is shown as a dotted line respectively. This illustrates that our templates cover a broad range of continuum shapes, with the SED peaking at longer wavelengths for lower blackbody temperatures. As a consequence, the Balmer break is stronger and sharper for the hotter sources where there is more continuum to absorb at $\sim 0.4\,\mu{\rm m}$.

sample, 44.8, 22.2, 12.3, 5.6, and 0.4% of all sources are selected with 2, 3, 4, 5, and all six templates.

The first and most important cut requires that the best-fitting SED from `eazy` is dominated by a BH* template. We discuss the implications of choosing a somewhat arbitrary cut of a template contribution of $> 80\%$ in Section 6.2. This cut alone yields a sample of 1,434 objects. We then remove sources fit at $z < 1.5$ and at $z > 9.5$. At the low redshift end, the distinction between robust BH*-dominated candidates and Milky Way stars becomes increasingly difficult, and we lack deep and uniform photometric coverage blue-wards of the Balmer break. On the other hand, at $z \sim 9.5$, the NIRCam-coverage is limited to $\lambda_{\rm rest} \lesssim 0.48\,\mu{\rm m}$, meaning that we lose sufficient access to the rest-optical continuum at even higher redshifts. The redshift cuts reduce the sample size to 297 objects. The third cut makes sure that the BH* candidates are better fit with our `sfhz`+BH* template set than with the large stellar grid described above, as well as than with the `sfhz` template set only. For five sources, `eazy` does not find any fit with the `sfhz` templates and returns an arbitrary value of $\chi^2(\texttt{sfhz}) = -1$. However, it does find a decent fit with a BH* template, so while these would nominally fail our $\Delta\chi^2(\texttt{sfhz})$ cut, we consider them to be promising candidates and include them in our sample. The $\Delta\chi^2$-based cuts are complemented by the fourth cut which requires that the best fit with a BH* template has a reasonably low $\chi^2/{\rm N_{filters}}$ to remove cases where either fit is bad. The third and fourth cuts combined remove 45 objects, so that we are left with 252. We visually inspect imaging cutouts and SED-fits, and remove eight sources that we identify as spurious detections or hot pixels, or whose photometry is affected by data quality issues, diffraction spikes, or a bright foreground source. We further remove a source in the Abell-2744 field that corresponds to a secondary lensed image of the triply imaged LRD A2744-QSO1 (Furtak et al. 2023a), and only keep the brightest to avoid double-counting the source, leaving us with 243 candidates. We identify two of those as contaminants in Section 4.3, yielding our final sample size of 241 BH*-dominated candidates shown from Section 5.

Finally, we correct all fluxes in the Abell-2744 cluster field for lensing magnification using the UNCOVER DR4 magnification catalog which is based on the lens modeling from Furtak et al. (2023b), building on photometric redshifts from Weaver et al. (2024) and Suess et al. (2024) and updated with spectroscopic redshifts in Price et al. (2025).

Figure 2 shows an example of a galaxy in the EGS field with imaging from CEERS that is selected in the `eazy` run with The Cliff template which produces a good fit at $z_{\rm phot} = 3.53$. The best-fitting SED without The Cliff template (`sfhz` only) is that of a quiescent galaxy (QG) at $z_{\rm phot} = 4.29$, and it does not match the flux in the F277W filter which is likely boosted by Hα emission. This illustrates that Hα emission is an important distinguishing feature between BH*-dominated sources and contaminants such as stars or QGs, since LRDs show typical Hα equivalent widths of $\sim 300 - 1500$ Å (see de

Graaff et al. 2025b). However, depending on the redshift of the source, and the photometric filter coverage, different solutions still remain degenerate. Further, the QG solution for EGS-73885 does not reproduce the strength of the Balmer break as probed by deep imaging in F150W, and the detected rest-UV flux at $\lambda_{\rm obs} < 1\mu$m. To quantify how significantly the fit with the BH* template is favored, we compute the Bayesian information criterion (BIC) as BIC $= k\log(\mathrm{N_{filters}}) - 2\log(L_{\rm max})$ where k is the number of degrees of freedom of the fit, $\mathrm{N_{filters}}$ is the number of available photometric filters, and $L_{\rm max}$ is the maximum likelihood, or equivalently the minimum χ^2 of the fit, $L_{\rm max} \propto e^{-\chi^2_{\rm min}/2}$. We can then approximate the Bayes factor as $B(\texttt{sfhz}) \approx e^{-\Delta\mathrm{BIC}/2} = e^{(\Delta\chi^2(\texttt{sfhz})+\log(\mathrm{N_{filters}}))/2}$. A high Bayes factor suggests strong evidence in favor of the fit with the BH* template (e.g., $B > 100$ is a common threshold for decisive evidence following Kass & Raftery 1995). In principle, the BIC is a good approximation only for well-constrained models (i.e., for $\mathrm{N_{filters}} \gg k$), which is not true in our case. Further, a high $\Delta\chi^2$ (and thus a low BIC and high Bayes factor) does not guarantee that the underlying model – in our case a combination of galaxy templates and a BH* template – is correct. Because we do not claim to have access to the "true" model for each spectrum, the exact values of the resulting Bayes factors should be interpreted cautiously. Nevertheless, this formalism takes into account that different sources are covered by different numbers of photometric filters, ranging from $\mathrm{N_{filters}} = 6$ for most PANORAMIC pointings to $\mathrm{N_{filters}} = 23$ in some parts of the Abell-2744 cluster field, and is therefore better suited than only comparing $\chi^2_{\rm min}$ values. For EGS-73885 shown in Figure 2, we find $B(\texttt{sfhz}) = 1.38 \times 10^6$. This example illustrates how our sample selection successfully identifies promising BH*-dominated candidates. We note that EGS-73885 has not been observed by any (public or proprietary) spectroscopic program to date (see Section 3.5). The example also shows a potential degeneracy between QG and BH* solutions in cases where the Hα-excess is less prominent or not well probed by the available photometry. We discuss this further in Section 4.3.

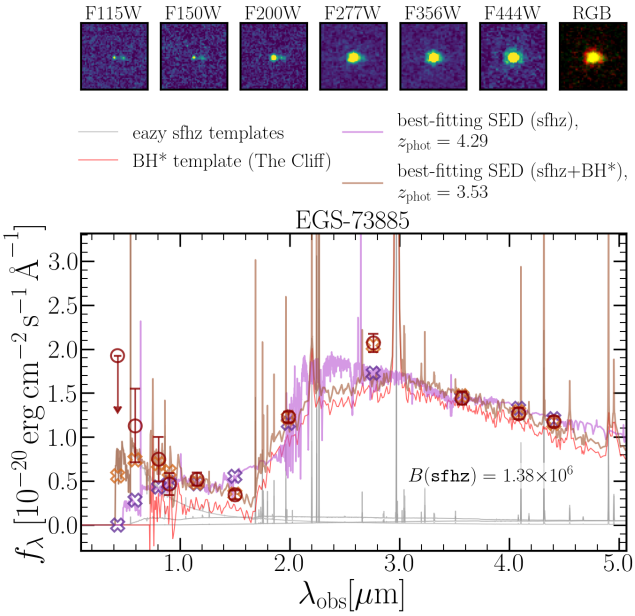

FIG. 2.— Example of a BH*-dominated candidate selected with the template based on The Cliff (de Graaff et al. 2025a). The stamps show $2.04'' \times 2.04''$imaging cutouts in JWST filters, as well as an RGB image which is composed of the F444W (red), F200W (green) and F115W (blue) filters. With the **sfhz**+Cliff template set, the photometry of EGS-73885 is fit at $z_{\rm phot} = 3.53$, with the fit being dominated by The Cliff template whose contribution to the total SED is shown in red. Contributions of **sfhz** templates to the best-fitting SED are displayed in gray. The best-fitting SED with the **sfhz** template set only (purple line) corresponds to a QG at $z_{\rm phot} = 4.29$ but does not reproduce the flux in F277W, the strong Balmer break, and the detected rest-UV flux at $\lambda_{\rm obs} < 1\,\mu$m, as can be seen from the model photometry (crosses). This results in a Bayes factor $B(\texttt{sfhz}) = 1.38 \times 10^6$ suggesting strong evidence in favor of the BH*-dominated solution.

3.5. *Spectroscopic Validation*

To validate our sample, we leverage the large and growing archive of spectroscopic data from JWST. We start by cross-matching our sample with all public spectra available on the DJA as of April 29, 2026, finding 55 spectra with grade 3 redshifts (i.e., robust redshifts as determined by visual inspection). Of those, 51 have at least one robust redshift from a PRISM spectrum, the other four sources have robust redshifts from G235M and G395M medium resolution grating spectra (two each). A full list of programs contributing public spectra used in this work is provided in Appendix A. We additionally include six PRISM spectra with grade 3 redshifts from two recently observed programs, GO-7511 (PI Covelo-Paz, four spectra), and GO-8915 (PI Weibel, two spectra), increasing the total number of sources with robust spectroscopic redshifts to 61, i.e., $\sim 25\%$ of the sample.

This spectroscopic sub-sample includes various well-known LRDs. Besides The Cliff and the MoM-BH*-1 which are selected by design (i.e., because they are used as templates), we select A2744-45924 (Labbe et al. 2024), A2744-QSO1 (Furtak et al. 2023a), UNCOVER-20466 (Kokorev et al. 2023), CAPERS-LRD-z9 (Taylor et al. 2025), and RUBIES-BLAGN-1 (Wang et al. 2025).

To estimate the possible contamination rate of our photometric sample, we focus on the sub-set of 57 sources with grade 3 PRISM spectra, as a more in-depth analysis of the medium resolution spectra is beyond the scope of this work. First, we identify one source as a BD and return to this contaminant in Section 4.3. Then, we cross-match with the spectroscopic LRD-sample from de Graaff et al. (2025b), and find 34 matches. Their selection method builds on Hviding et al. (2025) where LRDs are defined through a combination of a (spectroscopic) V-shape, compactness, and the detection of broad lines. The latter requires a medium grating spectrum with a sufficiently high SNR. However, de Graaff et al. (2025b) show that sources that satisfy the V-shape and compactness criteria turn out to have broad lines in 98% of the cases when such data are available. Thus, we consider candidates that are part of their sample to be robust LRDs (additionally selected here to have a dominant BH*-component). To further investigate the remaining 22 sources with PRISM spectra, we apply the V-shape classification from Hviding et al. (2025). The details of this analysis are described in Appendix B. Here, we directly move to estimating contamination rates. In doing so, we exclude the MoM-BH*-1 and The Cliff, be-

cause they were used as templates in the sample selection. Only considering sources with confirmed V-shapes as per Hviding et al. (2025) to be robust LRDs, we find a very conservative upper limit on the contamination rate of 14/55 (25.5%). However, except for the confirmed BD contaminant, we cannot rule out a V-shaped SED for any of these candidates due to insufficient SNR in the measured slopes (see Appendix B). If we additionally consider sources that show breaks or turnovers around H_∞, along with strong Hα-emission and/or tentative broadened lines to be confirmed LRDs, the potential contamination rate decreases to 7/55 (12.7%). In addition to the BD contaminant, we only identify one likely contaminant (UDS-27591, showing a blue UV-slope and a tentatively blue optical slope, along with strong narrow emission lines, see Appendix B), resulting in a lower limit on the contamination rate of 2/55 (3.6%).

In Figure 3, we show a gallery of 16 PRISM spectra of sources selected as BH*-dominated candidates in this work. This includes known sources from the literature which are labeled accordingly, as well as spectra presented here for the first time, which are highlighted with red frames. GDN-59854 and GDN-59983 from GO-7511 (PI Covelo-Paz) are new BH*-dominated sources at $z_{\rm spec} > 6.5$ (see Witten et al. in preparation). The spectra of The Cliff and the MoM-BH*-1 are plotted in purple to mark them as templates used for the sample selection (Section 3.2). The gallery highlights that the selection presented here successfully identifies sources that resemble these paradigmatic BH*s across a wide range of redshifts. Importantly, many of the objects shown in Figure 3 are not picked up by prevalent photometric LRD searches. We cross-match the 16 sources with the public catalogs from Kokorev et al. (2024) and Kocevski et al. (2025), noting that both works do not include the GOODS-N field, and Kokorev et al. (2024) also do not include the Abell-2744 cluster field. We find that Kokorev et al. (2024) only select three of the 12 sources in overlapping fields, and Kocevski et al. (2025) select five out of 13; the only object selected by both is GDS-16213. Most notably, both selections miss The Cliff as already discussed in Hviding et al. (2025). Kokorev et al. (2024) also miss the MoM-BH*-1 (Naidu et al. 2025) as well as the RUBIES-BLAGN-1 (Wang et al. 2025), while Kocevski et al. (2025) do not include A2744-45924 (Labbe et al. 2024) and our GDS-57503. We further discuss how our sample compares to photometric LRD samples in the literature in Section 3.7.

Next, we cross-match our sample in the Abell-2744 cluster field with the DR1 catalog from the All the Little Things (ALT) survey (GO-3516, PIs Naidu & Matthee, Naidu et al. 2024). We find three matches, two of which already have a NIRSpec-based spectroscopic redshift. The ALT catalog thus adds one grism redshift to our spectroscopic sample. Interestingly, the corresponding source at $z_{\rm grism} = 2.49$ (ID 49263, see Section 4.1) was highlighted in Figure 19 of the ALT survey paper (Naidu et al. 2024) to show a broadened Pa-γ line as well as HeI in absorption close to the systemic redshift. The authors therefore hypothesized that it may be a lower redshift version of an LRD.

We further cross-match our sample in the GOODS-N and GOODS-S fields with the public Hα-emitter catalog from Covelo-Paz et al. (2025) that is based on the FRESCO survey (Oesch et al. 2023), finding two matches with additional spectroscopic redshifts that are consistent with our $z_{\rm phot}$ estimates. One of these sources is a confirmed broad line Hα emitter from Matthee et al. (2023, their GOODS-N-13733), the other is too low SNR to determine whether the Hα line in the grism spectrum is broadened.

Another of our BH*-dominated candidates has been observed as a backup target of program 116.294D (PI Matthee) with VLT/X-Shooter, and it was then followed-up by the DDT program ID 116.2AQ0 (PI Matthee). These observations confirmed its redshift of $z_{\rm spec} = 1.73$ through the detection of broad Hα and Hβ lines. The source with ID 1115 is highlighted again in Section 4.1; it has been named PAN-BH*-1 and is discussed in detail in Torralba et al. (2026a).

Overall, 65 of our 243 initial candidates (including the contaminants discussed in Section 4.3) have robust spectroscopic redshifts, corresponding to 26.7% of the sample. We plot them against our photometric redshifts in Figure 4. The comparison shows remarkably good agreement between $z_{\rm phot}$ and $z_{\rm spec}$ with no significant systematic offset. Defining the photometric redshift deviation as $\Delta = |z_{\rm phot} - z_{\rm spec}|/(1+z_{\rm spec})$, we find a median deviation of $\Delta_{\rm med} = 0.017$ and a scatter of $\sigma_{\rm NMAD} = 0.021$. We then measure the outlier fraction, $\eta_{\rm outliers}$, as the number of sources satisfying $\Delta > 0.1$, and find $\eta_{\rm outliers} = 1.5\,\%$. We compare these numbers to Hviding et al. (2025) who derived photometric redshifts for their sample of spectroscopic LRDs using `eazy` with the `blue_sfhz` template set and found much larger values of $\Delta_{\rm med} = 0.045$, $\sigma_{\rm NMAD} = 0.127$ and $\eta_{\rm outliers} = 0.44$. They identified photometric redshift failures, represented by the high outlier fraction, as a major reason for incompleteness of photometric LRD samples. This comparison shows that the inclusion of BH* templates vastly improves photometric redshift estimates for (BH*-dominated) LRDs, confirming the suitability of our templates and the corresponding composite BH$*$ + host galaxy fits. We note that the four sources at $z_{\rm phot} \sim 5-6$ and $z_{\rm spec} \sim 7-8$ correspond to confusion between Hα and [OIII]/Hβ boosting the F444W filter, as is reflected by the error bars on $z_{\rm phot}$ in three cases. The remaining outlier at $z_{\rm spec} = 0$ is the BD which we discuss in Section 4.3.

We finally use the DJA to identify sources for which NIRSpec spectra have been taken but where extractions are not available, usually because the data are still proprietary, and we find 13 such sources. This means that 169 (69.5%) of our BH*-dominated candidates do not yet have any NIRSpec spectroscopy.

3.6. *Selection Completeness*

As is illustrated in Figure 3, the selection presented here is sensitive to BH*-dominated LRDs, i.e., LRDs with relatively weak UV-emission, due to our requirement of a sub-dominant (< 20%) host galaxy contribution in the rest-optical. This causes us to miss a substantial fraction of the total LRD population, in particular strongly V-shaped, UV-blue LRDs, meaning that the sample compiled here is an incomplete subset of the total LRD population. We explore this by comparing to various photometric LRD selection methods in the next Section. Here, we wish to roughly quantify the incompleteness of our sample due to photometric noise, taking

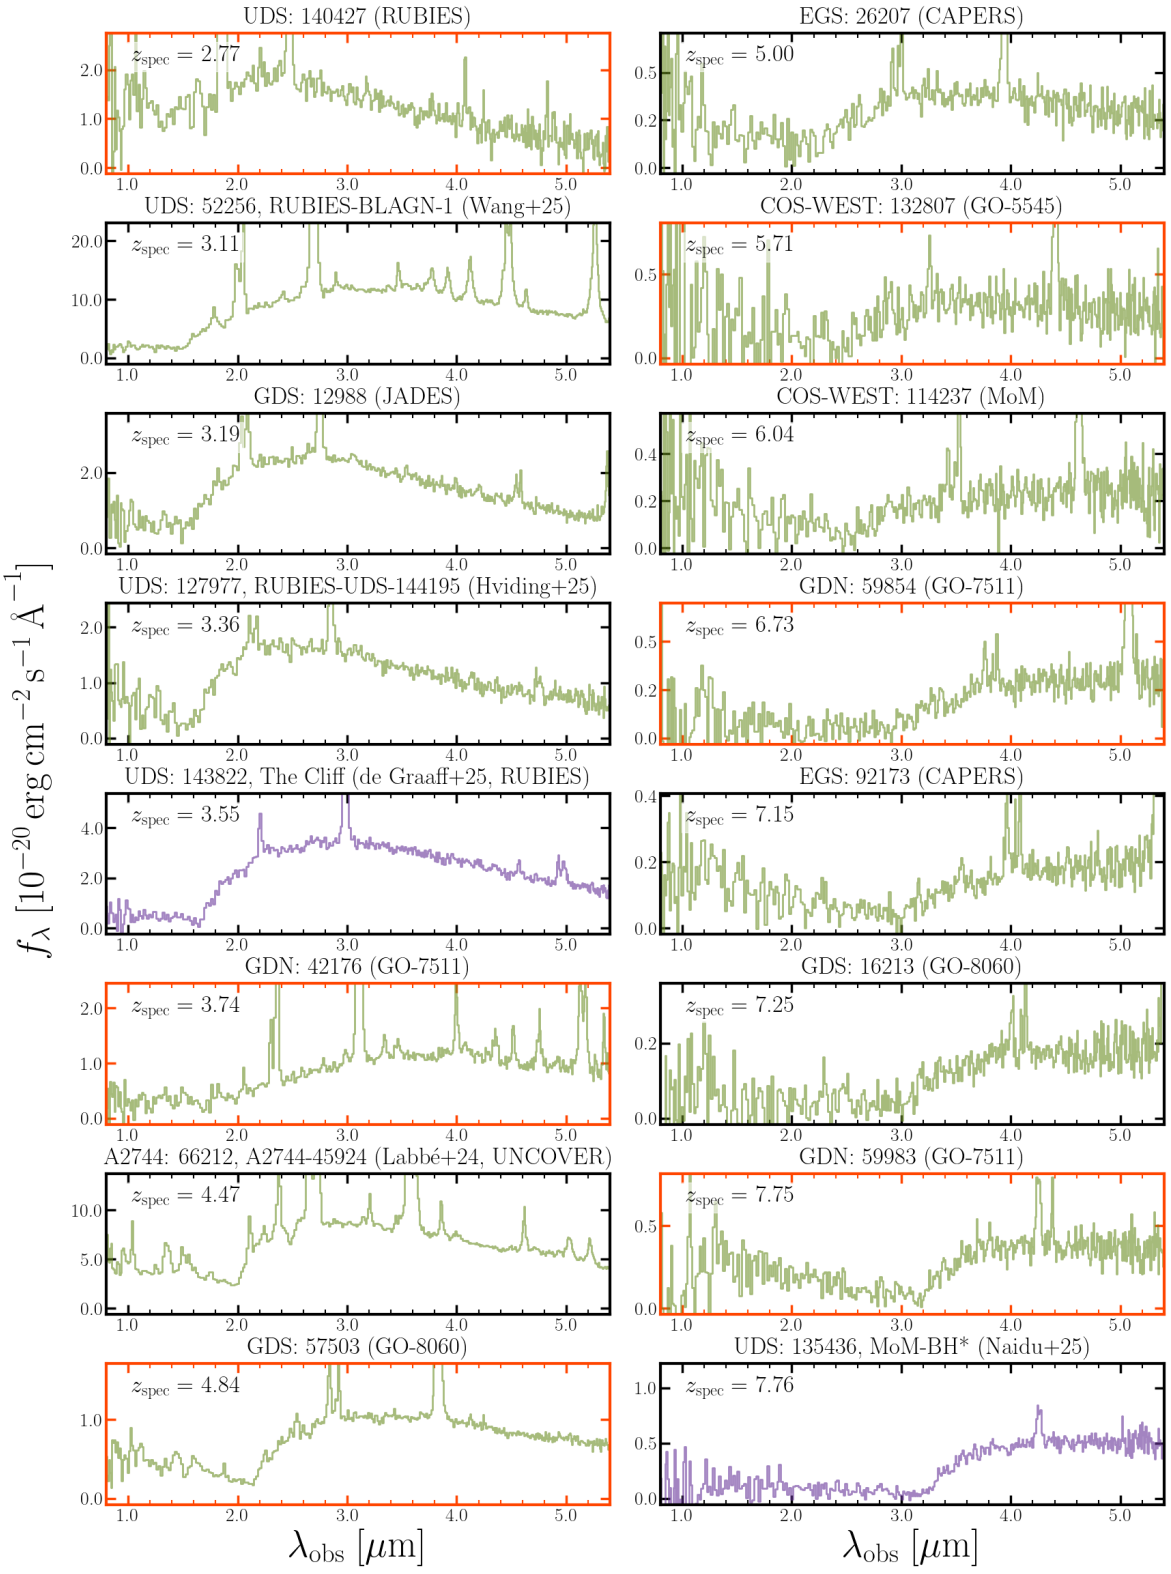

FIG. 3.— Example PRISM spectra of sources selected in this work, ordered by redshift. Spectra of well-known LRDs/BH*s (single object papers) are labeled accordingly, and spectra presented for the first time here are highlighted with red frames. This illustrates that our selection successfully identifies objects that resemble the paradigmatic sources The Cliff (de Graaff et al. 2025a) and the MoM-BH*-1 (Naidu et al. 2025) over a wide redshift range. These two sources are plotted in purple to indicate that they were used as templates in the sample selection. Most of the objects shown here are not picked up by prevalent photometric LRD selection methods.

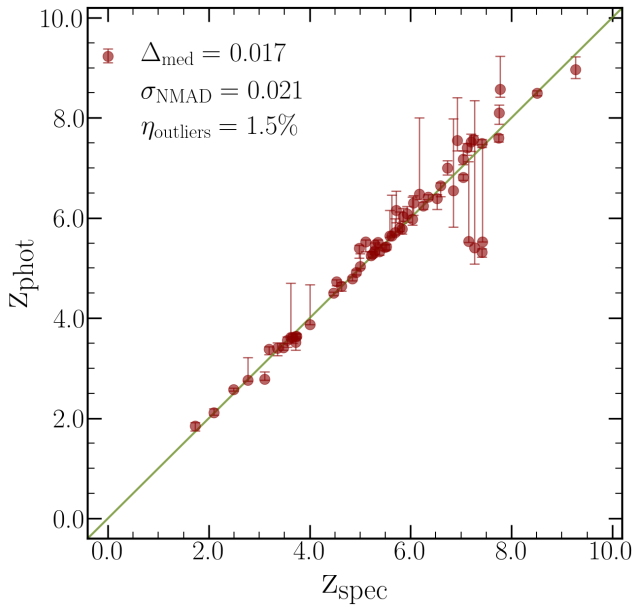

FIG. 4.— Photometric redshifts from `eazy` against spectroscopic redshifts for 65 of our BH*-dominated candidates that have robust redshifts from NIRSpec spectra (61), public grism catalogs (3), and the source named PAN-BH*-1 (Torralba et al. 2026a) with a $z_{\rm spec} = 1.73$ from VLT/X-Shooter.

into account the varying imaging depth across the different fields used in this work.

For simplicity, we only calculate the incompleteness caused by our requirement of SNR(F444W) > 10, and mag(F444W) < 27. This part of the selection is crucial because the F444W filter probes different rest-frame wavelengths as a function of redshift, introducing redshift-dependent completeness at, e.g., fixed optical luminosity. While the `eazy`-based requirement of a $> 80\%$ BH* template contribution introduces substantial incompleteness relative to the full LRD population, it is less sensitive to photometric noise, in particular given the stringent SNR-cut in F444W. We therefore start from the six BH* templates presented in Section 3.2 and scale each template to different values of redshift and monochromatic luminosity at 5100Å (L_{5100}). We compute synthetic photometric fluxes in F444W and then loop through the different fields in the master catalog to add noise to the photometric fluxes according to the respective median depth. We create 1000 realizations of the photometry in F444W for each BH* template, redshift, luminosity, and field, and then apply the F444W-based pre-selection cuts. The fraction of the 1000 realizations that passes these cuts provides a rough estimate of the photometric completeness as a function of redshift, L_{5100}, imaging depth, and for the different BH* templates. We return to these estimates in Section 5.1.

3.7. *Comparison to "Classic" V-Shape Selections*

In the following, we compare our sample selection to photometric LRD selections from the literature. Building on the color selection cuts from Labbe et al. (2025), Greene et al. (2024) defined two complementary selections named *red1* and *red2*, and combined them with a compactness cut to select LRDs at lower and higher redshifts, corresponding to $z \lesssim 6$ and $z \gtrsim 6$ (Kokorev et al. 2024). Specifically, their color cuts are

$$\begin{aligned} \textit{red1} = {} & \text{F115W} - \text{F150W} < 0.8 \wedge \\ & \text{F200W} - \text{F277W} > 0.7 \wedge \\ & \text{F200W} - \text{F356W} > 1.0 \\ & \text{and} \\ \textit{red2} = {} & \text{F150W} - \text{F200W} < 0.8 \wedge \\ & \text{F277W} - \text{F356W} > 0.7 \wedge \\ & \text{F277W} - \text{F444W} > 1.0 \end{aligned} \tag{1}$$

with all quantities in AB-magnitudes. This selects for a V-shaped SED, requiring a red color at the red end of the NIRCam-coverage (rest-optical), and a moderately blue color at the blue end (rest-UV). Kokorev et al. (2024) modified the *red2* cut to be somewhat less stringent, requiring F277W − F356W > 0.6 and F277W − F444W > 0.7. Simplifying the selection, Pérez-González et al. (2024) only required F277W − F444W > 1, combined with a slightly more stringent cut at the blue end, F150W − F200W < 0.5. Barro et al. (2024) and Akins et al. (2025) instead opted for a selection more targeted at extremely red sources, without imposing cuts on the rest-UV, as F277W − F444W > 1.5. Slightly changing the approach, Kocevski et al. (2025) defined a V-shape based on $\beta_{\rm opt}$ and $\beta_{\rm UV}$, the slopes obtained from fitting a power law to the photometry red-wards and blue-wards of H_∞. More recently, Rinaldi et al. (2026) defined more relaxed color cuts as F277W − F444W > 0.5 and F150W − F200W < 1 with the goal of obtaining a more complete sample of LRDs. Some color-based selections add further cuts to remove BDs by requiring, e.g., F115W − F200W > -0.5 (Greene et al. 2024).

To directly compare our template-based selection to a color-based V-shape selection, we apply the selection from Kokorev et al. (2024) to our photometric catalog. For consistency, we adopt the same SNR and magnitude cuts as used in this work (SNR(F444W) > 10 and mag(F444W) < 27), similar but not identical to the cuts applied in Kokorev et al. (2024, SNR(F444W) > 14 and mag(F444W) < 27.7). Conveniently, their compactness cut is identical to the cut used in this work. We then apply their color selection including the BD removal cut from Greene et al. (2024). We require 3σ detections in all bands needed to measure colors, or replace non-detections ($< 3\sigma$) by 2σ upper limits where appropriate to constrain the respective colors, consistent with Kokorev et al. (2024). We visually inspect all selected sources and remove spurious detections to obtain a sample of 545 V-shaped LRDs. In Figure 5, we show an Euler diagram illustrating the overlap between this color-based V-shape selection and our template-based search for BH*-dominated sources, showing that the two selections are highly complementary. Only 20.4% of the V-shaped LRDs are also selected as BH*-dominated candidates. Conversely, only 45.9% of our BH*-dominated candidates are picked up by the color-based V-shape selection. The main reason for this is that many BH*-dominated candidates are faint in the rest-UV, making

it impossible to constrain the blue part of the V-shape. We illustrate this in the bottom of Figure 5 where we plot the SNR in F150W against redshift for our sample of BH*-dominated candidates, and the V-shaped LRD sample. The vast majority of our BH*-dominated candidates that are missed by the V-shape selection lie at SNR(F150W) < 3 (dashed line). The importance of rest-UV SNR-cuts for the selection of UV-faint LRDs has already been discussed in Hviding et al. (2025).

Irrespective of the SNR in F150W, most of our sample galaxies do show red colors in F277W $-$ F444W. Specifically, 71.1% of our sample galaxies satisfy F277W $-$ F444W > 1. For sources with $z_{\rm phot} > 3$, 4, and 5 this fraction increases to 77.1, 89.1, and 100%. Many of the lower redshift candidates that are missed by a F277W $-$ F444W-based selection (*red2* in the Kokorev et al. 2024 selection, see Equation 1) instead show a red color in F200W $-$ F356W (*red1*). E.g., 85.5% of all our sample galaxies satisfy F200W $-$ F356W > 1, though even this color selection fails at $z \lesssim 2.5$ (see the examples in Section 4.1).

Towards the low redshift end of our sample ($z \lesssim 3$), the SNR in the F150W filter generally increases as it starts probing rest-optical wavelengths. While the *red1* selection relies on a detection in F115W instead of F150W, we note that $\sim 90\%$ of the sources undetected in F150W are also undetected in F115W. We note that there is no overlap between BH*-dominated candidates and V-shaped LRDs at $z \lesssim 3$, and there are no V-shaped LRDs at $z < 2$, reflecting that color-based selections are inherently redshift-dependent, while our template-based search is, in principle, redshift-independent.

For further comparison, we apply various other color cuts from the literature to our sample. The selections from Greene et al. (2024), Pérez-González et al. (2024), and Rinaldi et al. (2026) recover 47.1%, 34.7%, and 42.6% of our candidates, similar to the 45.9% inferred for the Kokorev et al. (2024) selection above. Interestingly, the stringent single color cut from Akins et al. (2025) selects 58.7 % of our sample because it does not rely on rest-UV fluxes. By design, it is however biased towards the reddest, most extreme, and typically high redshift sources ($z \gtrsim 5$). We compare these findings to Hviding et al. (2025) who inferred the completeness of the Kocevski et al. (2025), Kokorev et al. (2024) and the single-color F277W $-$ F444W > 1.5 (e.g., Akins et al. 2025) samples based on their spectroscopic sample of RUBIES LRDs, and found values of 61.8%, 50.0%, and 35.29%. This means that the selection based on $\beta_{\rm UV}$ and $\beta_{\rm opt}$ achieves the highest completeness for LRDs in general, and also picks up close to half of our BH*-dominated candidates. A single color cut such as used in Akins et al. (2025) is relatively successful at selecting the most extreme LRDs, including BH*-dominated sources, but is highly incomplete for the general population of LRDs. Typical color-based selections pick up roughly half of the population of LRDs, and an even smaller fraction of BH*-dominated objects. We emphasize that none of the photometric selection methods discussed above picks up any of our BH*-dominated candidates at $z_{\rm phot} < 2.5$ where the NIRCam wide filters at $\lambda_{\rm obs} \gtrsim 2\mu$m probe the declining part of the rest-frame optical to NIR SED (see Section 4.1).

The findings of this Section illustrate that the sample presented here is a subset of the LRD population that consists of objects with a dominant BH* contribution and that are relatively faint in the rest-UV, due to the faint, sub-dominant host galaxy. As such, our sample is highly complementary to previous photometric selections, and sensitive to a different part of the LRD parameter space.

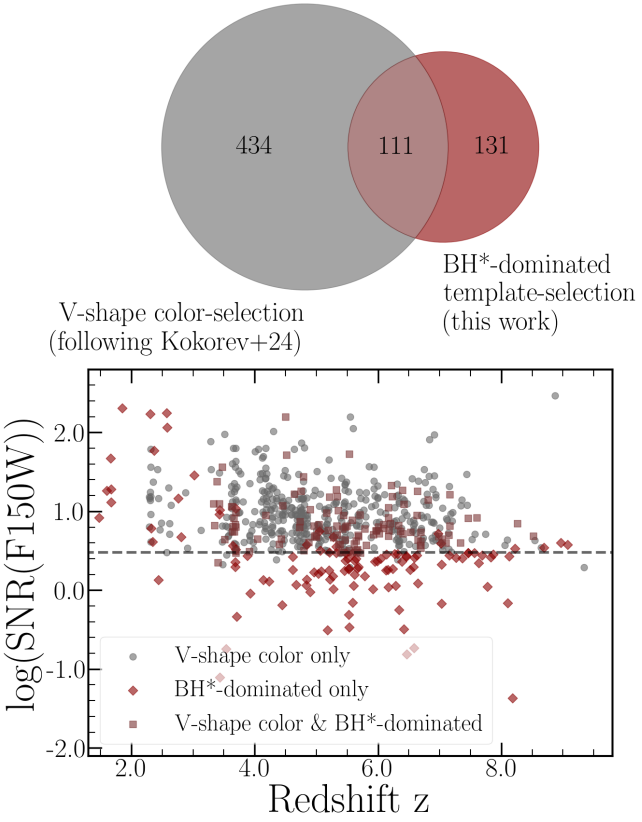

FIG. 5.— Top: Euler diagram showing the overlap between a color-based V-shape LRD selection following Kokorev et al. (2024), and the templated-based selection of BH*-dominated sources applied here. This highlights how these selections are highly complementary. Bottom: SNR in the F150W filter against redshift for our sample galaxies, as well as the color-selected V-shaped LRDs, illustrating the new parameter space probed in this work: low SNR in the rest-UV (as probed by F150W for $z \gtrsim 3$) and the lowest redshifts ($z \lesssim 2.5$). Sources picked up by both selections are plotted as squares.

4. SAMPLE OVERVIEW

To provide a first overview of our photometric sample, we show six BH*-dominated candidates that do not have public spectra on the DJA in Figure 6. We choose one source selected by each of the six BH* templates, and a wide range of photometric redshifts ($z \sim 3-9$). Lower redshift candidates are discussed separately in Section 4.1. The displayed candidates illustrate the varying filter coverage across fields ranging from six NIRCam wide filters in PANORAMIC to a wealth of HST and NIRCam wide and medium bands in legacy fields such as the GDS. We emphasize that BH*-dominated candidates at different redshifts can be successfully selected from pure parallel imaging in six NIRCam filters. The candidates in Figure 6 further show different levels of rest-frame UV flux that can be reconciled with a $> 80\%$ contribution of

the BH* template in the rest-optical. Subsequently, we discuss the lowest and highest redshift candidates in our sample in more detail.

4.1. *Low Redshift BH*s ($z < 2.5$)*

BH*-dominated sources at low redshift are particularly valuable for various reasons: First, their apparent brightness is relatively high at a fixed intrinsic luminosity, enabling high-resolution and high-SNR JWST spectroscopy in relatively short integration times. For the brightest objects at $z \sim 1.5-2.5$, even spectroscopy from the ground becomes viable (see Torralba et al. 2026a). Second, at lower redshift, JWST/NIRSpec has access to different spectral features such as Paschen lines, HeI at 1.08μm, and even potential water absorption features at $\sim 1.4\mu$m (e.g., Wang et al. 2026) which may provide crucial insights into the physical conditions that prevail in BH*-dominated sources. Third, the rapid decline in the number density of LRDs at $z \lesssim 4$ is poorly understood (e.g., Ma et al. 2026) and it may in part be related to selection effects inherent to photometric color selections that are mitigated by the selection strategy applied here.

In Figure 7, we present our four most promising $z < 2.5$ BH*-dominated candidates. Sorted by increasing redshift, COS-WEST-154491 at $z_{\rm phot} = 1.67$ in the top left shows a continuum shape that is remarkably similar to the BH* stack at $T_{\rm eff} = 4409$ K (Sun et al. 2026), including its extension to longer wavelengths as a pure blackbody. PANO-J024000M0142-1115 is the source tagged PAN-BH*-1 in Torralba et al. (2026a) with a spectroscopic redshift of $z_{\rm spec} = 1.73$ from VLT/X-Shooter. UDS-134217 in the bottom left has a G235M spectrum from GO-3567 (PI Valentino), confirming its photometric redshift as $z_{\rm spec} = 2.1$ through various prominent emission lines, in particular broadened Paβ and Paγ as well as HeI at $1.08\,\mu$m in both emission and absorption. Finally, A2744-49263 in the bottom right has a NIRCam/grism spectrum from the ALT survey ($z_{\rm grism} = 2.49$), as discussed in Section 3.5. In our photometric sample, the source is fit with The Cliff template which yields $z_{\rm phot} = 2.55$, close to the spectroscopic redshift. We refer to this source as the ALT-BH* hereafter.

4.2. *High Redshift BH*s ($z > 7$)*

Next, we turn to the other end of the redshift distribution, and highlight four of our highest redshift ($z > 7$) BH*-dominated candidates in Figure 8. None of these sources has a public spectrum on the DJA. GDS-52764 at $z_{\rm phot} = 7.10$ shown in the top left has deep HST and JWST photometry, including medium bands that constrain the Balmer break/turnover as well as the rest-frame UV. UDS-39621 in the top right has sparser wavelength coverage and shallower imaging from PRIMER, but is significantly brighter than GDS-52764, and fit at a slightly higher redshift ($z_{\rm phot} = 7.46$). UDS-169425 in the bottom left is again fainter at $z_{\rm phot} = 8.27$ and thus the least robust of the displayed candidates. Finally, PANO-J033224M2756-5088 in the bottom right is selected from PANORAMIC pure parallel imaging adjacent to legacy imaging in the GOODS-S field. At $z_{\rm phot} = 9.16$ it is nominally the source with the highest bolometric luminosity in the entire sample as we discuss in Section 5.3.

4.3. *Contaminants*

In Section 3.5, we estimate a contamination rate of $3.6 - 25.5\%$ for our photometric sample based on the subset of our candidates with grade 3 PRISM spectra. The only unambiguously identified contaminant is shown in Figure 9. It is a candidate BH*-dominated source at $z_{\rm phot} = 9.24$, selected with the template based on a CLOUDY model of the MoM-BH*-1 (Naidu et al. 2025). It shows a huge break in the photometry of a factor > 5 between F444W and F356W, and even of > 20 between F444W and F277W, the strongest break measured in our entire sample. The source was observed by MoM (GO-5224, PIs Oesch & Naidu), and the PRISM spectrum reveals that it is a cold BD, with a steeply rising continuum that peaks at $\sim 4.1\,\mu$m (see Hainline et al. in preparation).

Importantly, this type of contaminant only affects the high redshift end of the sample, where the supposed Balmer break of the BH* shifts to the red end of the NIRCam wavelength coverage, so that only one or two filters (F444W and F410M) are substantially boosted with respect to the other filters. As opposed to the BH* SED, the BD spectrum drops again rapidly at $\lambda_{\rm obs} \gtrsim 4.1\,\mu$m, meaning that it can be easily distinguished with sufficiently deep MIRI data.

The photometry of COS-WEST-39450, MoM-BD-1 hereafter, resembles that of the object tagged *Capotauro* in Gandolfi et al. (2026). The latter is roughly 3× fainter in F444W which may explain why it is not detected in any filters blue-wards of F410M. Gandolfi et al. (2026) discuss various physical origins of the source, including the possibility that it is a $z \sim 30$ galaxy, a BH*, as well as a BD. In the latter case, a temperature of $\lesssim 400$ K is required to produce a sufficiently strong break between F444W and F356W, as well as to not produce emission at shorter wavelengths that would be detected by JWST. MoM-BD-1 illustrates that such sources indeed exist and produce NIRCam photometry that is largely degenerate with BH* solutions at $z \gtrsim 7$. This is reflected in the measured small $\Delta\chi^2$(stars) $= 1.3$ between the `sfhz`+BH* templates and the large grid of stellar templates (Section 3.3), resulting in a Bayes factor B(stars) $= 0.6$, suggesting no evidence in favor of the BH*-dominated solution (Kass & Raftery 1995). We note that our sample selection cut based on $\Delta\chi^2$(stars) (Section 3.4) likely removes most stellar contaminants from our sample, so that we are left with this extreme and potentially rare type of BD that is not entirely captured by our stellar templates. Future work may include this source as another stellar template to further refine searches for BH*-dominated sources at high redshift.

We further identify a source, COS-WEST-29852, that is fit at $z_{\rm phot} = 5.61$ with the BH* template based on GN-9771. If true, its magnitude of mag(F444W) $= 19.81$ would make it more than an order of magnitude brighter than any other BH*-dominated candidate at similar redshifts. We show imaging cutouts, the measured photometry, as well as the best-fitting SED of this source in Appendix C. The photometry does not show distinctive features such as a break or filters boosted by emission lines, but suggests a smoothly declining red SED over the wavelength range probed by NIRCam. A more detailed investigation of COS-WEST-29852 is beyond the

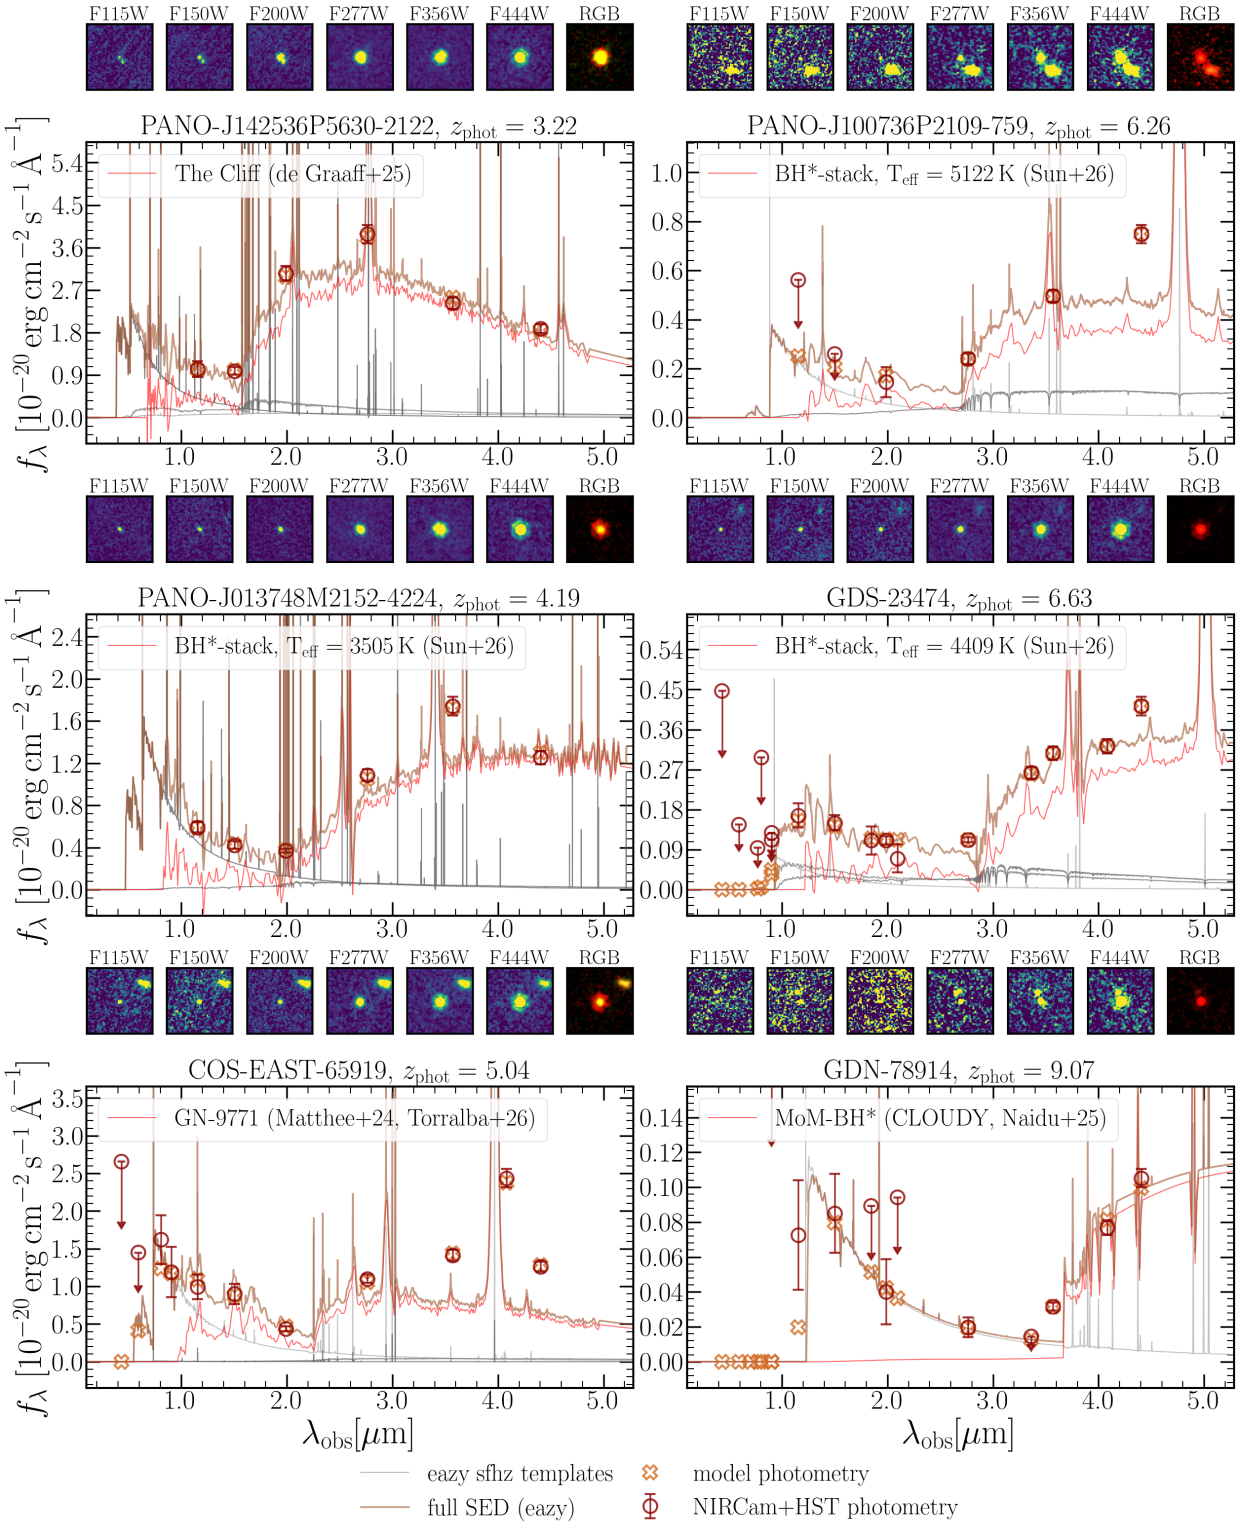

FIG. 6.— Overview of the sample selection. We show one BH*-dominated candidate selected by each of the six BH* templates used in this work (Section 3.2). The displayed candidates span a wide range in redshift ($z \sim 3-9$), and they highlight various levels of rest-frame UV emission, as well as different photometric constraints depending on redshift and wavelength coverage. The RGB stamp shown for each source is constructed from F444W (red), F277W (green) and F115W (blue). Gray lines show the contribution of different `eazy` templates to the best-fitting SED (brown line). The contribution of the respective BH* template that dominates the fit is shown in red.

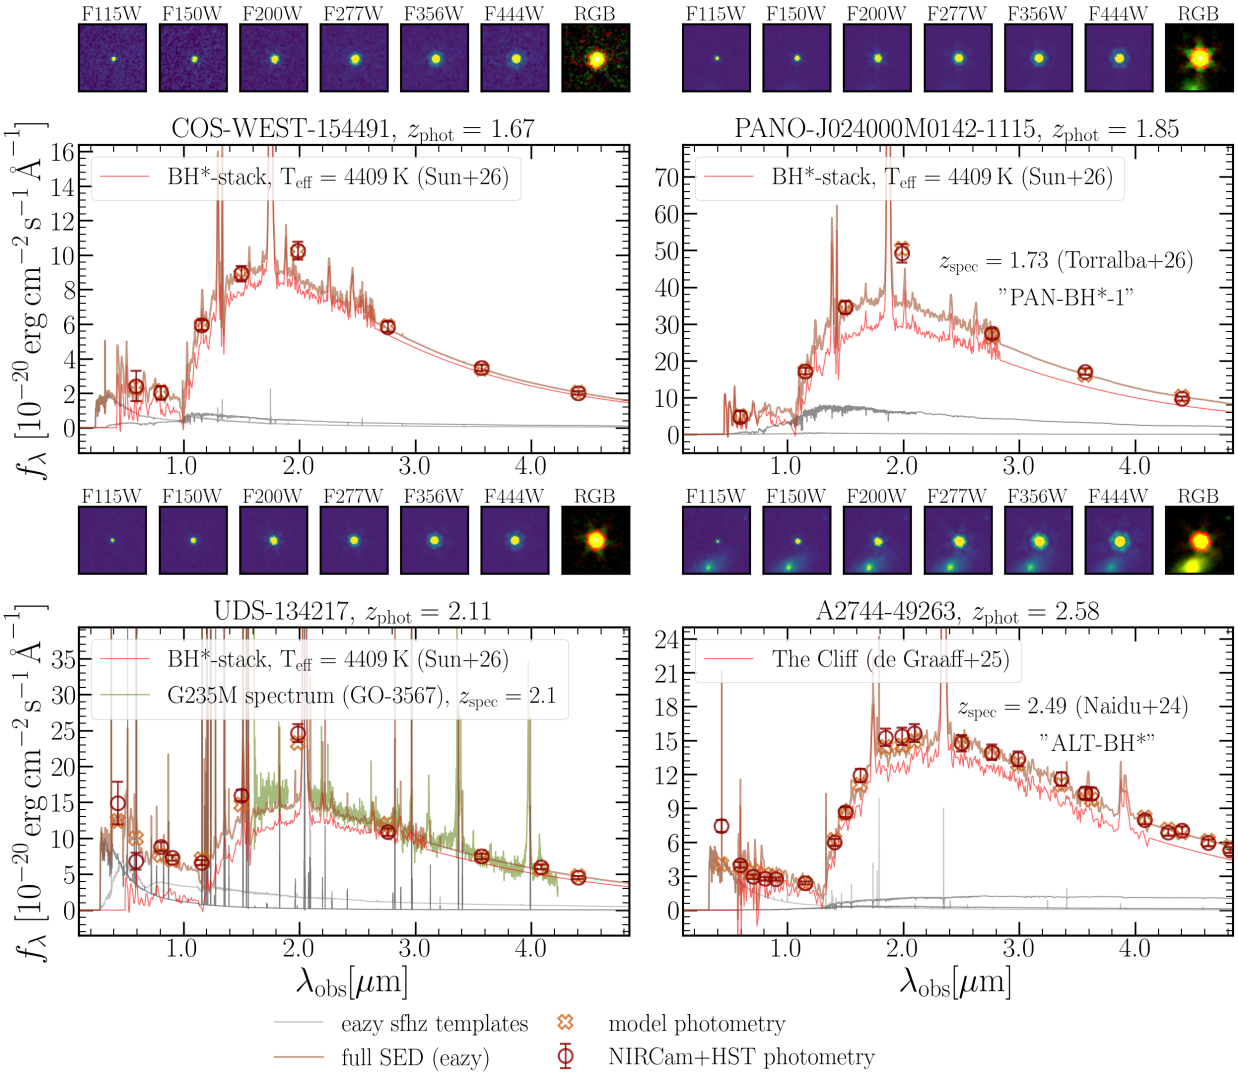

FIG. 7.— Most promising BH*-dominated candidates at low redshift ($1.5 < z < 2.5$), in analogy to Figure 6. Top left: Pure photometric candidate at $z_{\rm phot} = 1.67$ where the photometry closely follows the BH* stack at $T_{\rm eff} = 4409$ K, including the Balmer break, and the extension to $\lambda_{\rm rest} > 1.15\,\mu$m assuming a pure blackbody. Top right: PAN-BH*-1 spectroscopically confirmed in Torralba et al. (2026a) with a spectrum from VLT/X-Shooter showing remarkably strong absorption on top of a broad Hα line at $z_{\rm spec} = 1.73$. Bottom left: A source with a NIRSpec/G235M spectrum from GO-3567 (PI Valentino), confirming its redshift as $z_{\rm spec} = 2.1$ through broadened Paschen-lines and HeI/1.08 μm in emission and absorption. Bottom right: BH*-dominated candidate in the Abell-2744 cluster field that was highlighted in the ALT survey paper (Naidu et al. 2024, thus we name it "ALT-BH*") with a broad Paγ line as well as HeI in absorption at $z_{\rm spec} = 2.49$.

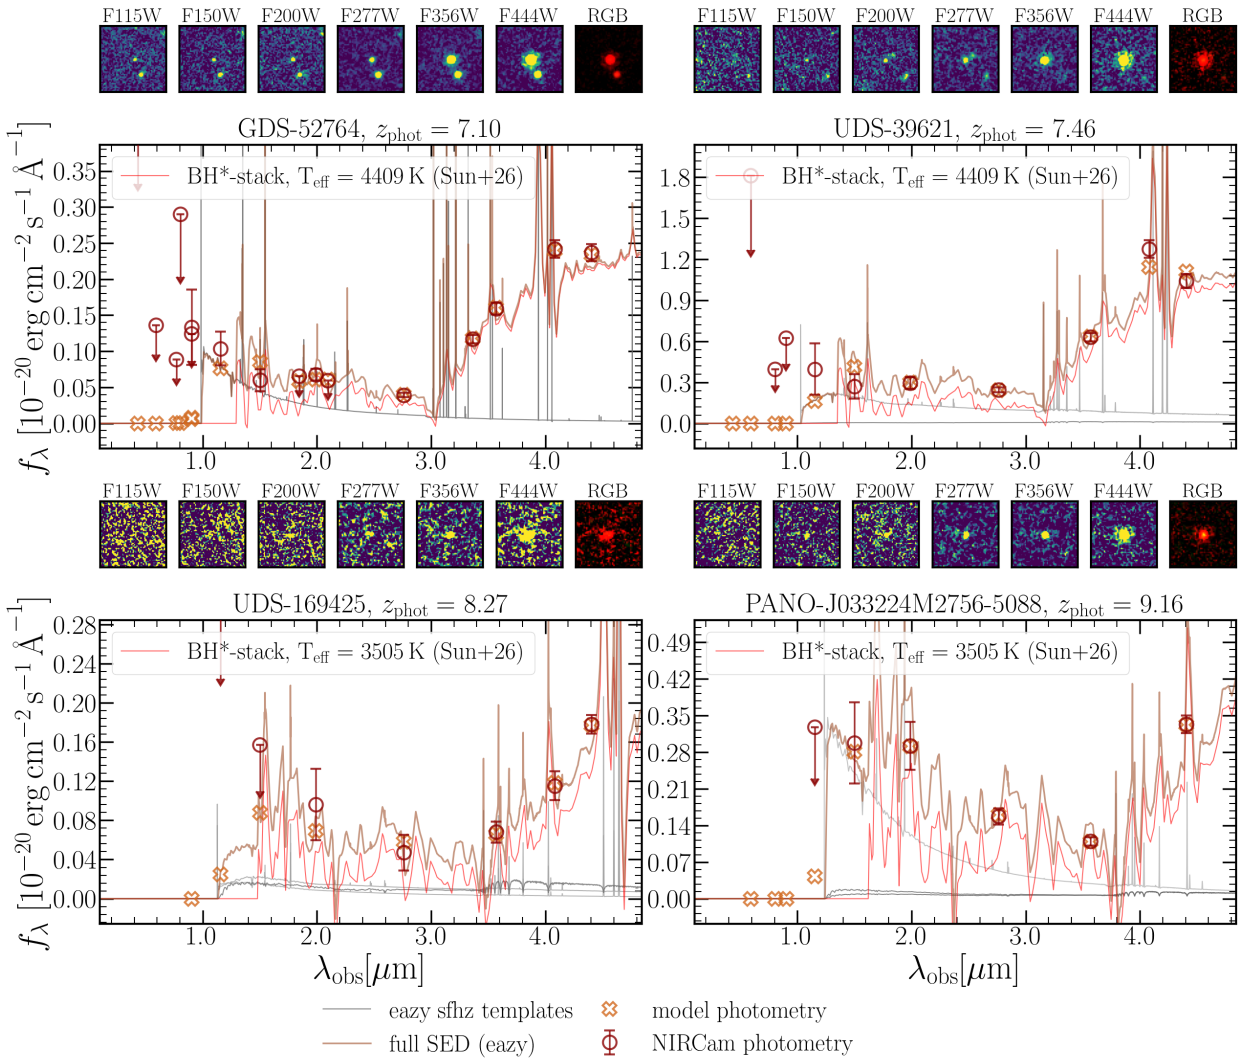

FIG. 8.— Most promising BH*-dominated candidates at high redshift ($z > 7$), in analogy to Figure 6. Top left: Candidate at $z_{\rm phot} = 7.10$ with deep photometry from JWST and HST in the GOODS-S field. Top right: Slightly higher redshift ($z_{\rm phot} = 7.46$), brighter candidate in the UDS field. Bottom left: Fainter candidate at $z_{\rm phot} = 8.27$ with a smoothly rising continuum from $\lambda_{\rm obs} \sim 3.5 - 4.5\mu$m that is fit with the coldest BH*-stack ($T_{\rm eff} = 3505$ K). Bottom right: A candidate at $z_{\rm phot} = 9.16$ identified in pure parallel imaging from PANORAMIC adjacent to legacy imaging in the GOODS-S field.

scope of this work, although we argue that it is unlikely to be a BH*-dominated source at $z_{\rm phot} = 5.61$. As noted in Section 3.4, we therefore remove this source, as well as the MoM-BD-1 from our sample, reducing our total sample size to 241.

For many of our BH*-dominated candidates, in particular at the low redshift end ($z \lesssim 3.5$), the best fitting SED with the `sfhz` galaxy templates is that of a QG. We show and discuss a corresponding example in Figure 2 where the fit with the BH* template is clearly preferred due to a strong Balmer break and Hα emission boosting a photometric filter. Some objects in our sample indeed have Balmer breaks that are too strong to be reproduced by a stellar population (see the discussion regarding The Cliff in de Graaff et al. (2025a), and Figure 10), but not all BH*-dominated candidates have strong Balmer breaks. The SEDs of the colder sources peak at longer wavelengths and their SEDs decline towards the rest-UV. At the Balmer limit there is thus not much continuum to break (see Figure 1). Further, in some cases, the limited depth of the photometric bands bluewards of the Balmer break does not allow us to probe its full strength. On the other hand, the detection of an Hα-excess relative to the continuum depends on the redshift of the source, and the available photometric filters, causing the photometry of some candidates to be more degenerate with QG solutions, even if many photometric bands are available. For example, our source GDS-57503, spectroscopically confirmed to be a BH*-dominated LRD through GO-8060 (PI Egami) as shown in Figure 3, has been confidently modeled to be a QG based on 22 photometric bands including 7 NIRcam and 8 MIRI filters in Williams et al. (2024). Crucially, any degeneracy between BH*-dominated and QG solutions is captured by the Bayes factor B(`sfhz`) which we use to define sample tiers in the following.

4.4. *Gold and Silver Sample*

We use the Bayes factors B(`sfhz`) and B(stars) to define a gold sample of BH*-dominated candidates that

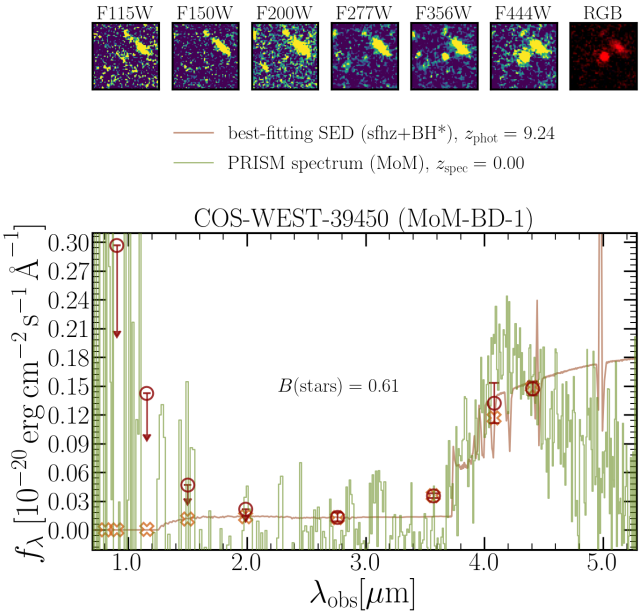

FIG. 9.— BH*-dominated candidate at $z_{\rm phot} = 9.24$, spectroscopically identified as a cold BD through the MoM program. We refer to this source as the MoM-BD-1 hereafter. The best-fitting `eazy` SED and the photometry are displayed in analogy to Figure 6. The PRISM spectrum from MoM is shown in green. The Bayes factor B(stars) = 0.61 indicates no evidencde in favor of the fit with the BH* template, reflecting the degeneracy between solutions when only NIRCam photometry is considered.

we consider the most robust because the fit with the BH* template is clearly favored over fits with standard galaxy templates as well as over fits with stellar templates. Specifically, we require B(`sfhz`) > 100 and B(stars) > 100, following Bayes factor thresholds for decisive evidence as provided in Kass & Raftery (1995). We consider the five sources for which `eazy` does not find any fit with the `sfhz` templates (see Section 3.4) to have high Bayes factors. This yields a sample of 127 BH*-dominated sources that we refer to as the gold sample hereafter; the remaining 114 sources are referred to as the silver sample. We show the photometry, best-fitting SEDs from `eazy`, as well as RGB cutouts of all our gold sample candidates in Appendix D.

5. SAMPLE PROPERTIES

We now turn to exploring sample properties of both our gold and silver samples of BH*-dominated candidates. We start with purely empirical quantities, specifically L_{5100} and the Balmer break strength, before moving on to estimating bolometric luminosities.

5.1. *Optical Luminosity and Balmer break strength*

We use the best-fitting `eazy` SED to infer the monochromatic luminosity L_{5100}, as well as the Balmer break strength, where the latter is given by the ratio between the median fluxes in the ranges $\lambda_{\rm rest} \in [4060, 4140]$ Å and $\lambda_{\rm rest} \in [3560, 3640]$ Å, measured in f_ν. We plot L_{5100} against redshift (photometric or spectroscopic where available) and the derived Balmer break strength in Figure 10. Uncertainties in $z_{\rm phot}$ are computed by excluding any probability density at $z > 10$ because solutions at $z > 10$ largely correspond to unphysical combinations of galaxy and BH* templates. We use uncertainties in rest-frame fluxes from `eazy` to compute uncertainties for L_{5100} and the Balmer break strength.

The Figure highlights that we identify BH*-dominated sources across the entire redshift range we select for $(1.5 < z < 9.5)$, and over almost three orders of magnitude in L_{5100}. Our sample is mostly concentrated around intermediate redshifts of $z \sim 3-7$, and luminosities of $L_{5100} \sim 10^{43} - 10^{44}\,{\rm erg\,s^{-1}}$ with the fraction of objects belonging to the silver sample slightly increasing towards the upper right (bright, high redshift) as well as the lower left (faint, low redshift) parts of the z-L_{5100} diagram. The lower right edge of the distribution reflects a selection effect where we only detect intrinsically brighter objects at higher redshifts. This is illustrated by the 50% completeness limit at the median imaging depth in F444W across all fields, computed according to Section 3.6 for the BH* template based on The Cliff. There are only a handful of candidates at $L_{5100} \lesssim 10^{42.5}$, only two of which are part of the gold sample, and we return to these candidates in Section 5.3. At the other end of the distribution, we identify relatively few objects at $L_{5100} \gtrsim 10^{44}$, with the source A2744-45924 (Labbe et al. 2024) standing out as the most optically luminous source. Our BH*-dominated candidates also show a broad range of Balmer break strengths, with most sources showing strong breaks. Specifically, 95% of all candidates show Balmer breaks > 1.85, and the median break strength is 3.05, close to the strongest breaks observed in high redshift QGs (~ 3.1; de Graaff et al. 2025a). A few objects show some of the strongest Balmer breaks observed in the Universe to date (e.g., 21 objects have break strengths > 6). We highlight various spectroscopically confirmed objects that are part of our sample: PAN-BH*-1 (Torralba et al. 2026a), the ALT-BH* (Naidu et al. 2024), RUBIES-BLAGN-1 (Wang et al. 2025), The Cliff (de Graaff et al. 2025a), A2744-45924 (Labbe et al. 2024), A2744-QSO1 (Furtak et al. 2023a), the MoM-BH*-1 (Naidu et al. 2025), UNCOVER-20466 (Kokorev et al. 2023), and CAPERS-LRD-z9 (Taylor et al. 2025). All these sources lie at high optical luminosities relative to the sample at their respective redshift, except for A2744-QSO1 which is strongly lensed (Furtak et al. 2023a), illustrating that well-known literature sources tend to probe an extreme part of the parameter space. We finally note a weak correlation between L_{5100} and Balmer break strength, consistent with the findings of de Graaff et al. (2025b).

5.2. *Bolometric Luminosities*

We calculate bolometric luminosities ($L_{\rm bol}$) of all our BH*-candidates by integrating over the BH* template contribution to the best-fitting `eazy` SED. For sources with spectroscopic redshifts, we instead use the best-fitting SED at $z = z_{\rm spec}$. The resulting $L_{\rm bol}$ is meant to quantify the total energy output of the BH* itself, and to not include any host galaxy contribution. We note that, by construction, the `eazy` fits are dominated by the BH* template in the rest-frame optical. The templates then fall off rapidly towards longer wavelengths due to their blackbody extensions (Section 3.2). Based on rest-frame FIR constraints on two of the most luminous LRDs, Setton et al. (2025a) showed that the IR-luminosity of LRDs is very limited. Greene et al. (2026) specifically studied

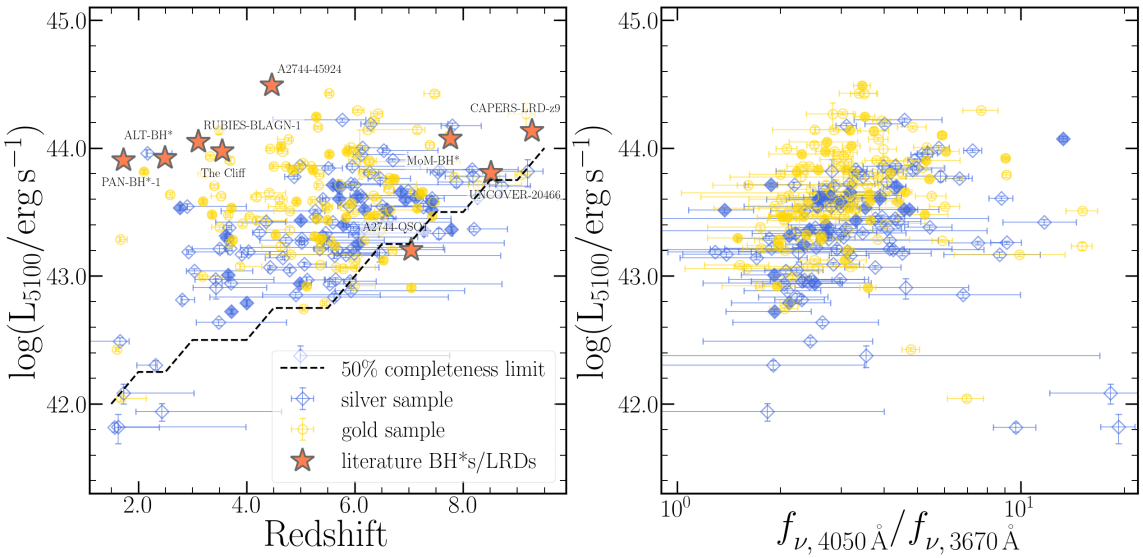

FIG. 10.— Monochromatic optical luminosity L_{5100} against redshift ($z_{\rm phot}$, replaced by $z_{\rm spec}$ where available) on the left and against Balmer break strength ($f_{\nu,\,4050\text{Å}}/f_{\nu,\,3670\text{Å}}$) on the right. The gold and silver sample are displayed with different markers and colors, and sources with spectroscopic redshifts are plotted as filled markers. In the left panel, we further plot the rough 50% completeness limit of our sample selection based on The Cliff template at the median depth across all fields, illustrating that the lower right edge of the distribution in the z-L_{5100} diagram reflects a selection effect. These plots show that we identify BH*-dominated candidates across a wide range of redshifts ($z \sim 1.5-9.5$) and luminosities $L_{5100} \sim 10^{42}-10^{44.5}\,{\rm erg\,s^{-1}}$, with the bulk of the sample being concentrated around $z \sim 3-7$ and $L_{5100} \sim 10^{43}-10^{44}\,{\rm erg\,s^{-1}}$. The sample displays a broad range of Balmer break strengths of $\sim 1.9-8.7$ (5th and 95th percentiles) with a high median of 3.05, and some of the strongest Balmer breaks measured in the Universe to date. There is a weak apparent correlation between Balmer break strength and L_{5100}, consistent with the trend seen for LRDs in de Graaff et al. (2025b).

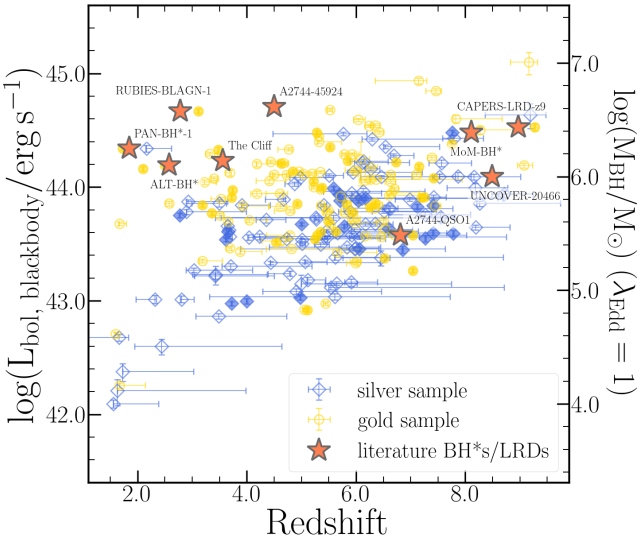

FIG. 11.— Bolometric luminosities derived by integrating the BH* template contribution to the best-fitting SED, plotted against redshifts ($z_{\rm phot}$, replaced by $z_{\rm spec}$ where available). Filled markers indicate sources with spectroscopic redshifts and various objects known from the literature are highlighted as stars, and labeled accordingly (Torralba et al. 2026a; Naidu et al. 2024; Wang et al. 2025; de Graaff et al. 2025a; Labbe et al. 2024; Naidu et al. 2025; Furtak et al. 2023a; Taylor et al. 2025; Kokorev et al. 2023). The secondary y-axis specifies black hole masses derived from $L_{\rm bol}$ using Equation 2, and assuming $\lambda_{\rm Edd} = 1$. This again highlights the broad range of redshifts and luminosities spanned by our sample, as well as the rapid decline in the number of candidates towards the highest ($L_{\rm bol} \gtrsim 10^{44.5}\,{\rm erg\,s^{-1}}$) and lowest ($L_{\rm bol} \lesssim 10^{43.5}\,{\rm erg\,s^{-1}}$) luminosities.

implications for $L_{\rm bol}$ estimates and found that more than half of $L_{\rm bol}$ is emitted in the rest-frame optical, justifying the approach followed here. Figure 11 shows the derived bolometric luminosities plotted against our photometric redshifts from `eazy`.

$L_{\rm bol}$ is directly related to the black hole mass as

$$M_{\rm BH} \approx 0.81 \times 10^5\,{\rm M_\odot}\,\lambda_{\rm Edd}^{-1}\left(\frac{L_{\rm bol}}{10^{43}\,{\rm erg\,s^{-1}}}\right) \quad (2)$$

where $\lambda_{\rm Edd}$ is the Eddington ratio (see, e.g., Umeda et al. 2025). Assuming $\lambda_{\rm Edd} = 1$, i.e., accretion at the Eddington limit, then yields a black hole mass that, in principle, constitutes a lower limit. However, recent studies proposed that LRDs may accrete at or even above the Eddington limit (e.g., Lambrides et al. 2026; Liu et al. 2025; Madau & Maiolino 2026a), suggesting $\lambda_{\rm Edd} = 1$ may be a reasonable assumption, or even over-estimate the true black hole masses. We assume $\lambda_{\rm Edd} = 1$, noting that the resulting $M_{\rm BH}$ represents a rough reference and has to be interpreted cautiously which is why we only show $M_{\rm BH}$ as a secondary y-axis on the right of Figure 11.

In rough analogy to the results for L_{5100}, we detect BH*-dominated candidates over three orders of magnitude in $L_{\rm bol}$. The full sample, and especially the gold sample, is mostly concentrated around intermediate to high luminosities ($L_{\rm bol} \sim 10^{44}\,{\rm erg\,s^{-1}}$), while the number of candidates declines rapidly towards higher ($L_{\rm bol} \gtrsim 10^{44.5}\,{\rm erg\,s^{-1}}$) and lower ($L_{\rm bol} \lesssim 10^{43.5}\,{\rm erg\,s^{-1}}$) luminosities. The implied black hole masses (for $\lambda_{\rm Edd} = 1$) span $M_{\rm BH} \sim 10^4 - 10^7\,{\rm M_\odot}$.

5.3. *Intrinsically Brightest and Faintest BH*s Candidates*

The objects at the extremes of the luminosity distribution are of particular interest because they may inform us about the physics driving the engines of the LRDs. On

the one hand, the intrinsically brightest sources may host the most massive black holes and/or have the highest accretion rates and they are the easiest and most obvious targets for spectroscopic follow-up. On the other hand, intrinsically faint LRDs remain relatively unexplored, as the best-studied objects so far are also among the brightest (see Figure 10). Further, if confirmed, our faintest candidates would imply that BH*-dominated sources exist over three orders of magnitude in L_{5100} (and $L_{\rm bol}$), and may push us into the intermediate mass black hole regime (e.g., Greene et al. 2020).

In Figure 10, the source A2744-45924 (Labbe et al. 2024) stands out as the most optically luminous source. In terms of $L_{\rm bol}$, it is slightly less extreme because it is fit with the BH* template corresponding to the hottest stack ($T = 5122\,$K; Sun et al. 2026) which peaks around $\lambda_{\rm rest} \approx 5000\,$Å and falls off rapidly towards longer wavelengths. Interestingly, we measure even higher $L_{\rm bol}$ for three candidates at $z_{\rm phot} > 7$, with the highest luminosity ($\log(L_{\rm bol}/{\rm erg\,s^{-1}}) = 45.1$) measured for the source PANO-J033224M2756-5088 which is shown in the bottom right of Figure 8. However, its high redshift of $z_{\rm phot} = 9.16$ implies that only the F444W filter probes wavelengths red-wards of the Balmer break, and it is limited to $\lambda_{\rm rest} \lesssim 4850\,$Å at the red end. Nominally, the best-fitting BH* template is the coldest stack ($T = 3505\,$K; Sun et al. 2026), which peaks at $\lambda_{\rm rest} \sim 0.8\mu$m, meaning that most of the bolometric luminosity is emitted in a wavelength range that is not directly constrained by NIRCam. Further, a cold temperature for the most luminous source is in contrast to the trend that more luminous LRDs are associated with higher temperatures (i.e., shorter peak wavelengths, de Graaff et al. 2025b). Indeed, the fits with other BH* templates are almost equally good, pointing to a degeneracy between templates at high redshift, and suggesting that the true $L_{\rm bol}$ may be lower. To test this, we infer $L_{\rm bol}$ again from the best-fitting `eazy` SED with the hottest stack ($T = 5122\,$K) and find $\log(L_{\rm bol}/{\rm erg\,s^{-1}}) = 44.6$ which still places the source among the most luminous in the sample, but below, e.g., A2744-45924. We further elaborate on the degeneracy between BH* templates at high redshift in Section 6.1.2 and subsequently focus on L_{5100} as a more directly constrained quantity.

Despite the larger area ($\sim 0.3\,{\rm deg}^2$) searched here compared to previous JWST-based searches, A2744-45924 remains the most optically luminous source in our sample. This is in line with the sharp cutoff found at the bright end of the optical luminosity function (LF) in Ma et al. (2025a), although we lack sufficient area to put tighter constraints on the LF (see Section 5.4). In Figure 12, we show the second and fourth most optically luminous BH*-dominated candidates (apart from A2744-49524 and the third most luminous source, UDS-39621 shown in Figure 8), as well as the two faintest sources that are part of the gold sample. The most luminous sources are outstanding BH*-dominated candidates that warrant spectroscopic follow-up, in particular because they appear to be rare.

Moving to the intrinsically faintest sources, COS-WEST-11734 in the bottom left is well fit with the hottest BH*-stack ($T = 5122\,$K), though it lacks sufficiently deep rest-UV constraints to fully probe the Balmer break. On the other hand, EGS-7039 in the bottom right shows a stronger break of ~ 3 based on the photometry (using a 2σ upper limit on the flux) and a good fit with the GN-9771 template, including its extension to longer wavelengths as a pure blackbody. We note that these are the only two sources in the gold sample with $\log(L_{5100}/{\rm erg\,s^{-1}}) < 42.5$. All other faint candidates show $B(\texttt{sfhz}) < 100$, i.e., the photometry is more degenerate with galaxy template fits. They are further preferentially found at low redshift where they appear brighter, but where we often lack strong constraints blue-wards of the Balmer break as is illustrated by COS-WEST-11734. We also note that even if we consider the full sample, we only find nine candidates at $\log(L_{5100}/{\rm erg\,s^{-1}}) < 42.5$, and only three silver candidates at $\log(L_{5100}/{\rm erg\,s^{-1}}) < 42$, in line with the findings of de Graaff et al. (2025b) who discussed a possible physical origin of the lack of LRDs at low luminosities but could not rule out a spectroscopic selection effect. The 50% completeness limit shown in Figure 10 suggests that the faint cut-off observed here is indeed a selection effect, and that current photometric data sets are insufficient to robustly detect fainter BH*-dominated sources, if they exist.

5.4. *Number Densities*

The mere fact that BH*-dominated sources exist over such a wide range of redshifts, and at least down to $z = 1.73$ (Torralba et al. 2026a) implies that they are not merely an early Universe phenomenon requiring primordial conditions or primordial chemistry (e.g., Dayal & Maiolino 2026). Note that lower redshift LRDs such as J1025+1402 ("The Egg"; Lin et al. 2026) would likely not qualify as BH*-dominated due to a stronger host galaxy contribution.

To provide a rough estimate of the number density evolution of BH*-dominated sources as a function of redshift, we build on the completeness calculation described in Section 3.6. Calculating 80% completeness limits in terms of L_{5100} at the median F444W imaging depth in analogy to the 50% completeness limits shown in Figure 10, we find that our sample is $> 80\%$ complete for $\log(L_{5100}/{\rm erg\,s^{-1}}) > 43.5$ at $z < 7.5$ for all six BH* templates. We thus define six redshift bins centered at $z = 2$, 3, 4, 5, 6, and 7 with a width of $\Delta z = 1$. In each bin, we count the number of BH*-dominated candidates with optical luminosity $\log(L_{5100}/{\rm erg\,s^{-1}}) > 43.5$ in fields whose F444W depth is sufficient to detect a BH*-dominated source at the minimum optical luminosity and at the central redshift of the bin. We calculate the survey area across the sufficiently deep fields from the number of un-flagged pixels in all six NIRCam wide filters (F115W, F150W, F200W, F277W, F356W, and F444W) that are required in the sample selection (Section 3), and convert this to a co-moving volume in each redshift bin by which we divide the respective number count. The resulting number densities are shown for both the full sample and the gold sample in Figure 13. We compare to number densities inferred through a ground-based search at $z < 4$ from Ma et al. (2026). These are derived for sources below a treshold in the absolute magnitude at 5500Å $M_{5500} < -20.5$, which corresponds to $\log(L_{5500}/{\rm erg\,s^{-1}}) = 43.57$, close to our threshold of

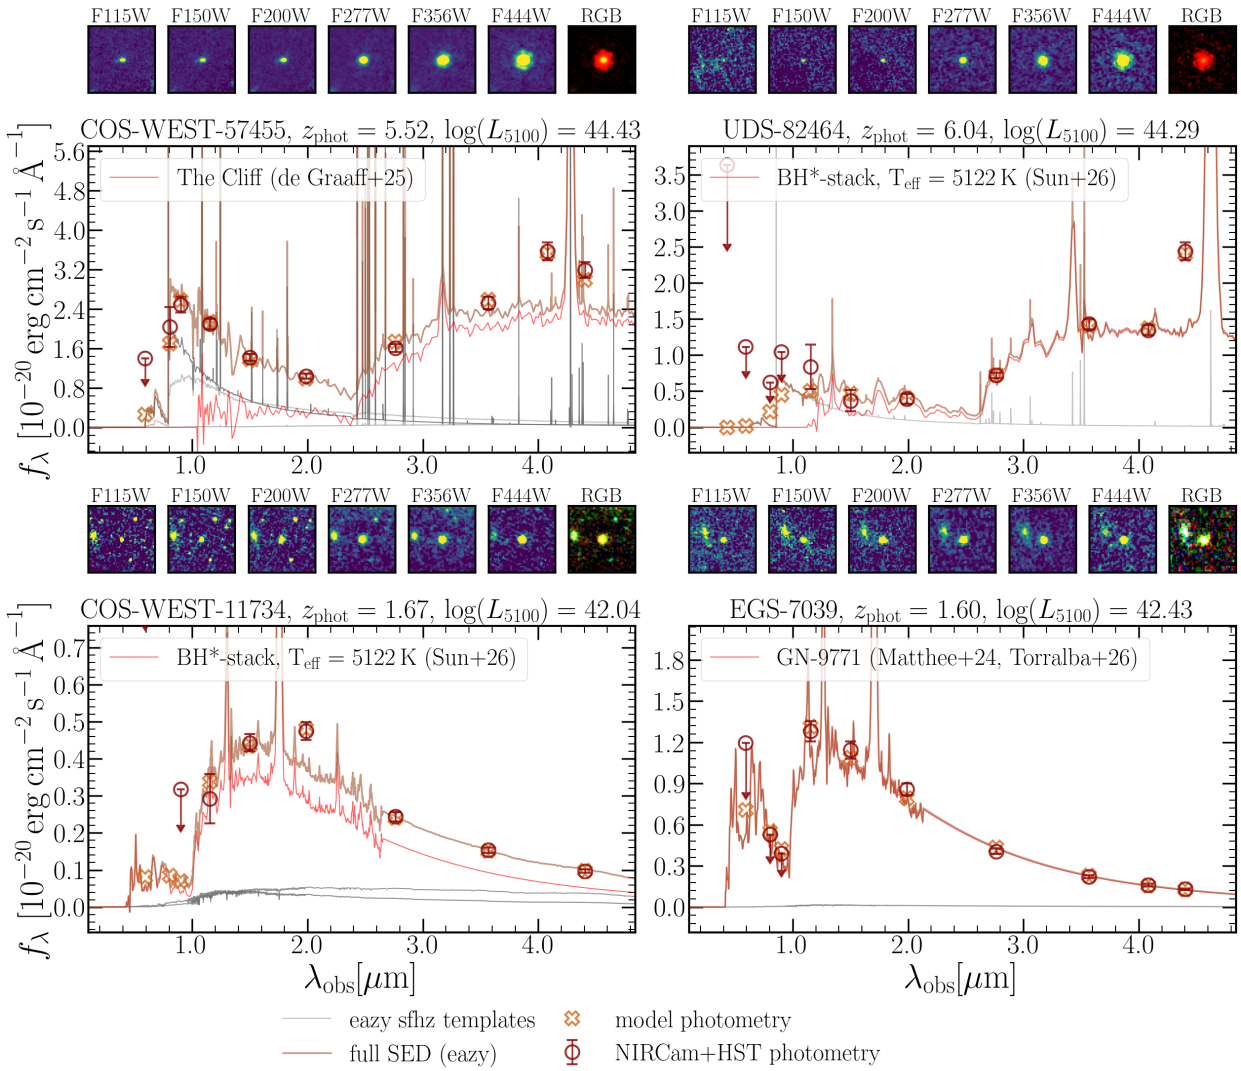

FIG. 12.— BH*-dominated candidates at the highest (top panels) and lowest (bottom panels) optical luminosities. The former constitute some of the most outstanding candidates and are ideal targets for spectroscopic follow-up. The latter probe a relatively unexplored part of the LRD parameter space given most well studied objects are among the intrinsically brightest. However, we can only detect the faintest sources at the lowest redshifts where we often lack strong constraints blue-wards of the Balmer break. Top left: $z_{\rm phot} = 5.52$ candidate in the COSMOS field and the second most luminous source in the sample. Top right: Fourth most luminous candidate at $z_{\rm phot} = 6.04$, identified in the UDS field. Bottom left: Intrinsically faintest candidate in the gold sample at $z_{\rm phot} = 1.67$ showing a good fit with the BH*-stack at $T_{\rm eff} = 5122$ K, but lacking tight constraints blue-wards of the Balmer break. Bottom right: Slightly brighter candidate at a similar redshift ($z_{\rm phot} = 1.60$) in the EGS field with tighter constraints on the Balmer break from NIRCam/F090W and HST imaging.

$\log(L_{5100}/{\rm erg\,s^{-1}}) > 43.5$. Ma et al. (2026) also provide equivalent number density estimates at higher redshifts based on the luminosity functions from Kokorev et al. (2024) which we include in the comparison. Finally, we plot the log-normal relation from Inayoshi (2025) which is inferred from the LRD-sample of Kocevski et al. (2025).

Our measurements suggest that the number density of BH*-dominated sources peaks at $z \sim 5-6$ and declines by about an order of magnitude from $z \sim 6$ to $z \sim 2$. We also identify a weak decline from $z \sim 6$ to $z \sim 7$ which is, however, not significant within uncertainties. At even higher redshifts, the number of candidates in our sample drops rapidly, though this is affected by a selection effect where we only detect the most luminous sources at the highest redshifts. Those appear to be rare at all redshifts as can be seen in Figure 10. Moreover, the limited coverage of the rest-frame optical by NIRCam at those high redshifts, makes it harder to robustly identify BH*-dominated sources using the method outlined here, which may further decrease our number densities. We argue that larger samples, and/or the inclusion of longer wavelength data, in particular from MIRI is required to shed further light on the number density evolution at $z \gtrsim 7$.

Intriguingly, our number density estimates for the full (gold) sample are a factor of 2.5 to 4.5 (5 to 8) lower than those based on the LRD-samples from Kokorev et al. (2024) and Kocevski et al. (2025) at $z \sim 5-8$. Moving to lower redshifts, our number densities are only factors of 2.3 (3.5) and 1.7 (2.2) below those from Ma et al. (2026) at $z \sim 2$ and $z \sim 3$, and lie slightly above the log-normal curve from Inayoshi (2025), albeit with large uncertainties. Taken at face value, and assuming that BH*-dominated LRDs are a subset of the full LRD pop-

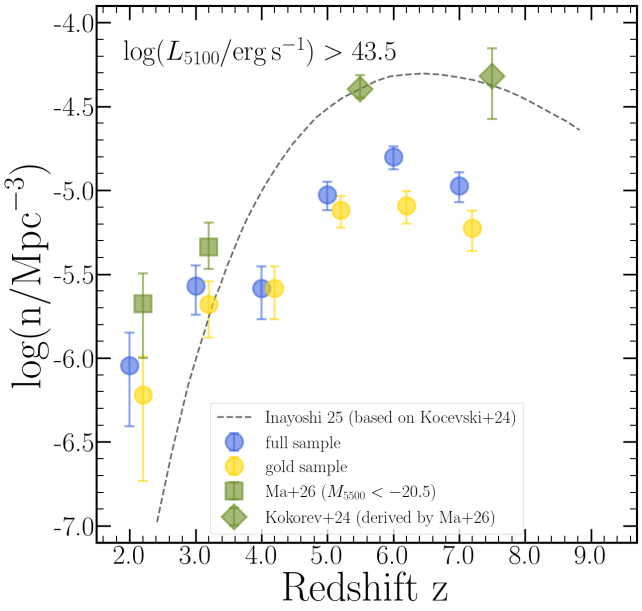

FIG. 13.— Rough number density of our BH*-dominated candidates as a function of redshift for both the full and the gold sample. These are derived from dividing the number of sources in redshift bins centered at $z = 2$, 3, 4, 5, 6, and 7 by the respective co-moving volume. The error bars represent Poisson uncertainties, and the markers for the gold sample are displaced by 0.2 on the x-axis for better visual separation. For comparison, we show number densities at $z < 4$ from Ma et al. (2026), as well as their numbers provided for the sample from Kokorev et al. (2024), and the log-normal relation from Inayoshi (2025) based on the sample from Kocevski et al. (2025).

ulation, this indicates that while the total LRD number density declines towards lower redshift, the fraction of BH*-dominated sources among them increases. However, there are a few important caveats to consider.

As we have discussed in Section 3.7, photometric V-shape selections tend to miss objects with a dominant BH*-component because they are too faint in the rest-UV to measure a blue UV-color. Specifically, applying the selection criteria from Kokorev et al. (2024) to our photometric catalog in Section 3.7 yields a sample of 545 LRDs, compared to our 241 BH*-dominated candidates. Moreover, the V-shaped LRD sample only contains about half of our sources (see Figure 5). Similarly, Ma et al. (2026) selected LRDs at $z \sim 2-4$ defining a V-shape criterion in the r, z, J, and K_s bands which likely also misses BH*-dominated sources. Different number densities in Figure 13 therefore trace partially disjoint samples, complicating their interpretation. Nevertheless, our detection of robust BH*-dominated sources around $z \sim 2$, including some with spectroscopic confirmation (see Figure 7) suggests that the number density of such objects remains surprisingly high around cosmic noon, relative to expectations from LRD number densities. While this may point to a significant incompleteness in existing photometric samples of (V-shaped) LRDs at $z \sim 2-3$, Ma et al. (2026) also point out that their sample likely contains a significant fraction of contaminants, and they argue their number densities should be considered as upper limits. We further note that our extension of the BH* templates to longer wavelengths assuming single blackbodies more strongly affects the selection at low redshift where NIRCam probes rest-frame wavelengths in the extended part of the templates. While we do find candidates whose SED shape closely follows the extended template (e.g., COS-WEST-154491 in the top left of Figure 7), this may not be true for all BH*-dominated sources, and could thus be a source of incompleteness of our low redshift sample. In summary, this points to a surprising abundance of BH*-dominated sources around cosmic noon which remains to be confirmed through spectroscopic follow-up and larger samples in the future.

6. DISCUSSION

In the preceding sections, we have developed a novel method of selecting BH*s from photometric data based on template fitting with `eazy`. Using six different empirical BH* templates covering a wide range of intrinsic luminosities and blackbody temperatures, we compiled a sample of 241 BH*-dominated candidates. We showed that our robust candidates range from $z \sim 1.7$ to $z \sim 9.3$ and span more than two orders of magnitude in optical luminosity ($\log(L_{5100}/\mathrm{erg\,s^{-1}}) \sim 42-44.5$). Subsequently, we discuss some caveats and limitations of our sample selection, along with possible future avenues for improvement. We end the discussion by exploring the potential of this new selection method to compile samples of LRDs more broadly.

6.1. *Caveats and Limitations*

6.1.1. *Intrinsic Variation in BH* Spectra*

The completeness of our sample selection hinges upon the extent to which our BH* templates cover the intrinsic variation in BH* spectra. To this day, only few spectroscopically confirmed BH*-dominated objects are known in the literature (e.g., de Graaff et al. 2025a; Naidu et al. 2025; Torralba et al. 2026a). In principle, future spectra may reveal new features and add diversity to the library of observed BH*s that is not yet captured by our templates. On the other hand, if the larger population of LRDs can be reliably decomposed into a host galaxy and a BH* (e.g., Barro et al. 2026), then the host-subtracted spectra and stacks from Sun et al. (2026) used in this work provide a solid baseline for the identification of BH*-dominated objects over a wide range of intrinsic properties.

Nevertheless, current samples of LRDs, both spectroscopic and photometric, may still be missing extreme and rare sources. Focusing on the extremes in blackbody temperature, the hottest sources are represented by our template based on GN-9771 (Torralba et al. 2026b). While hot sources are characterized by strong Balmer lines and Balmer breaks (de Graaff et al. 2025b) that make them easy to identify in principle, they may also show significant UV-emission that is intrinsic to the BH* component (Labbe et al. 2024; Torralba et al. 2026b; Ando et al. 2026). If true, this complicates their selection from photometry because they are more degenerate with LRDs that show a significant host galaxy contribution.

We explore the cold end of the temperature distribution by adding a host-subtracted version of the source UNCOVER-20698 (de Graaff et al. 2025b; Wang et al. 2026) to our template set. It is highlighted as one of the coldest sources ($T_{\rm eff} \sim 2300\,\mathrm{K}$) in the sample of de Graaff et al. (2025b, their Fig. 5), and was shown to display a

potential water absorption feature at $\lambda_{\rm rest} \sim 1.4\mu$m, indicative of a cold atmosphere ($\lesssim 3000\,$K) in Wang et al. (2026). First, this template only adds a small number of candidates to the sample which may indicate that extremely cold sources ($T \lesssim 3000\,$K) are rare. However, the template picks up a confirmed QG at $z_{\rm spec} = 3.25$ as a BH*-dominated candidate, illustrating that the degeneracy between BH* and QG solutions is particularly pronounced for cold sources (see also Section 4.3). At the same time, we do not select the source UNCOVER-20698 itself from our photometric catalog due to degeneracies with galaxy templates, also related to its host galaxy. Further, with this extremely cold template, some spectroscopically confirmed sources are fit at the wrong redshift, e.g., the ALT-BH* at $z_{\rm spec} = 2.59$ (bottom right panel in Figure 7) is fit at $z_{\rm phot} = 1.68$, indicating further degeneracy with other BH* templates and making it extremely difficult to identify (potentially rare) cold BH*s from photometric data alone. Additional medium band and MIRI imaging may help to put tighter constraints on the emission line strength (colder sources tend to show weaker Hα emission, de Graaff et al. 2025b), as well as the continuum shape of such objects. Nevertheless, spectroscopy will be indispensable to unambiguously identify and characterize cold BH*-dominated sources, as well as potentially new and unexplored types of LRDs or BH*s.

6.1.2. *Limitations at High Redshift*

At the highest redshifts, our sample selection is increasingly limited by the rest-frame optical coverage of the available NIRCam filters. In particular, at $z \gtrsim 9$, only F444W (and the medium bands probing similar wavelengths) trace the light red-wards of the Balmer break. This leads to a degeneracy with at least a certain type of BDs as exemplified by the MoM-BD-1 (Section 4.3). Such BDs may, however, be rare and most stellar contaminants are likely captured by our large grid of stellar templates. In general, the rest-optical continuum shape of the highest redshift candidates is poorly constrained by NIRCam imaging causing degeneracies between different BH* templates. Thus, it is nearly impossible to constrain the peak wavelength of the optical continuum and the blackbody temperature. As discussed in Section 5.3, this introduces substantial uncertainty in the bolometric luminosity of the highest redshift candidates, and may specifically cause us to overestimate $L_{\rm bol}$ for some of them. We also note that we select the highest redshift spectroscopically confirmed LRD, CAPERS-LRD-z9 (Taylor et al. 2025) at $z_{\rm spec} = 9.27$, showing that our selection successfully identifies LRDs at such high redshift.

In the high redshift regime, MIRI data will be extremely useful. First, Figure 9 shows how the spectrum of the MoM-BD-1 declines much more sharply at $\lambda_{\rm obs} > 4\,\mu$m than the BH* template, so that the two can be easily distinguished based on MIRI fluxes. Besides removing this remaining type of contaminant, MIRI can probe Hα at $z \gtrsim 7$, and it provides valuable constraints on the rest-optical continuum shape of BH*-dominated candidates. A major limitation of MIRI is the depth that can be achieved, though many of our BH*-dominated candidates are sufficiently bright to be detected in existing or upcoming MIRI surveys.

6.1.3. *Limitations at Low Redshift*

At the lower end of the redshift distribution, our sample selection is mainly limited by the availability of sufficiently deep imaging at rest-frame wavelengths blue-wards of the Balmer break. At $z \lesssim 2.5$, this regime is mainly probed by HST/ACS and the NIRCam SW filters F070W and F090W which are not uniformly available across the fields studied here, and are often significantly shallower than the NIRCam/LW imaging. Without sufficient constraints below the Balmer break, i.e., without actually constraining the full Balmer break strength, the photometry becomes degenerate with QGs or even sources showing relatively weak stellar Balmer breaks, such as "mini"-quenched or "napping" galaxies (e.g., Looser et al. 2024; Covelo-Paz et al. 2026). However, the few robust sources we identify at $z < 2.5$ (specifically the four sources highlighted in Figure 7) are particularly valuable as they allow us to probe redder rest-frame wavelengths with NIRSpec, and to explore lines such as Paβ, Paγ and HeI at 1.08μm (see the spectrum of UDS-134217 in the top left panel of Figure 7), or even molecular absorption features (Wang et al. 2026).

6.2. *A New Way of Selecting LRDs*

In Section 3.7, we discussed how the selection method applied here is highly complementary to photometric LRD selections in the literature that are based on colors between NIRCam filters and/or a V-shaped SED. Specifically, we do not require any detection in the rest-frame UV which is often very faint for the sources presented here. Our requirement of a BH* template contribution of $> 80\%$ at $\lambda_{\rm rest} > 4000\,$Å leaves some room for UV flux from a host galaxy which is why we recover some "classic" V-shaped LRDs. However, with increasing host emission in the rest-UV, the contribution of the host galaxy also increases in the rest-optical (depending on the age of the stellar population and its dust content), so that the BH* contribution quickly drops to $< 80\%$.

In principle, we could use the search method presented here to select LRDs more broadly by lowering our threshold below 80%. As a simple test, we compare to the spectroscopic LRD sample of de Graaff et al. (2025b). Considering only overlapping fields, and sources that pass our pre-selection cuts (in particular SNR(F444W) > 10 and mag(F444W) < 27), we find that our sample contains 34/84 of their sources, corresponding to 40.5%. If we decrease our template contribution threshold to 75% and 50%, these numbers increase to 44/84 (52.4%) and 70/84 (83.3%), while the total sample size increases by a factor of 1.4 and 3.1. This indicates that while it is possible to recover most LRDs with the method presented here, this likely comes at the cost of a significantly higher contamination fraction. The extra BH* template adds a degree of freedom to the `eazy` fits, so that the fits to some galaxies can be nominally improved with a limited contribution of a BH* template which is not physically motivated but simply reflects degeneracies with other templates. Our choice of a relatively high threshold of 80% is motivated by our objective of finding the most extreme, BH*-dominated sources, resembling the paradigmatic sources The Cliff and the MoM-BH*-1. In other words, we wish to identify BH*-dominated sources with a small host galaxy contribution. With a threshold of 80%,

we achieve a high completeness for such sources with a limited contamination fraction. At the same time, we are increasingly incomplete for sources with a stronger host galaxy contribution, thus missing a substantial fraction of V-shaped LRDs as explored in detail in Section 3.7.

7. SUMMARY AND CONCLUSIONS

In this paper, we present a new selection method for LRDs that focuses on sources whose SED is dominated by the central engine, the BH*. This is motivated by the two first spectroscopically confirmed BH*s, The Cliff and the MoM-BH*-1, which are characterized by an optical continuum that resembles a blackbody at a temperature of $T_{\rm eff} \sim 5000\,$K, a strong Balmer break, and very weak UV emission. Our method is based on incorporating six different BH* templates in the `eazy` redshift-fitting code to then identify objects whose best-fitting SED is dominated by these templates. We construct empirical templates that cover the range of spectroscopically observed BH* properties, in particular a wide range of intrinsic luminosities and temperatures, based on individual sources and BH* stacks from Sun et al. (2026). Applying the method to a photometric catalog across 1000 arcmin^2 of legacy and pure parallel NIRCam-imaging in six or more filters, we identify 241 BH*-dominated candidates. Cross-matching with publicly available spectroscopic data, we find that $\sim 27\%$ of our candidates have spectroscopic redshifts, largely confirming our SED-modeling and photometric redshifts. The vast majority of the available spectra with sufficiently high SNR show V-shaped SEDs, blackbody shaped continua, strong Balmer breaks and relatively weak rest-UV emission, indeed resembling objects like The Cliff and the MoM-BH*-1. Among all the spectra, we only identify one secure contaminant, a BD, and one likely contaminant with strong emission lines boosting the NIRCam/LW photometry. To mitigate the possible impact of such contaminants, we calculate Bayes factors between our fits with galaxy and a BH* template, and pure galaxy fits, as well as fits with stellar templates. We use these Bayes factors to define a gold sample of 127 particularly robust BH*-candidates showing Bayes factors > 100.

Our results can be summarized as follows:

- BH*-dominated sources exist across the Universe, and at least over the range $z = 1.7 - 9.3$, indicating they are not merely an early-Universe phenomenon but persist to cosmic noon.

- Comparing our sample to photometric V-shape selections in the literature, we show that our selection is highly complementary, and identifies objects that are missing from existing photometric samples, in particular because they are too faint in the rest-UV to measure a blue UV-color. Conversely, and by design, our sample misses strongly V-shaped LRDs with significant host galaxy contributions.

- We identify four promising candidates at $z < 2.5$, three of which have been confirmed spectroscopically through NIRCam/grism, NIRSpec/G235M and ground based VLT/X-Shooter spectroscopy, but lack continuous spectroscopic coverage in the rest-frame optical. These sources provide unique opportunities to study the engines of the LRDs in great detail owing to their apparent brightness ($21.9 - 23.6\,$mag at $\sim 2\,\mu$m), and by enabling NIRSpec spectroscopy out to $\lambda_{\rm rest} \sim 1.8\,\mu$m.

- We measure optical luminosities spanning at least two orders of magnitude, $\log(L_{5100}/{\rm erg\,s^{-1}}) \sim 42.5 - 44.5$, with candidates as faint as $L_{5100} \sim 10^{42}\,{\rm erg\,s^{-1}}$ at $z < 2$.

- The Balmer break strengths of our candidates are centered around a median of ~ 3, close to the maximum possible stellar Balmer breaks, and with a long tail out to break strengths > 10, the strongest Balmer breaks measured in the Universe to date.

- We estimate bolometric luminosities by integrating over the BH* template contribution to the best-fitting SED, and find that our sample covers $\log(L_{\rm bol}/{\rm erg\,s^{-1}}) \sim 42 - 45$. Converting the bolometric luminosities to black hole masses assuming accretion at the Eddington limit, we find $M_{\rm BH} \sim 10^5 - 10^7\,{\rm M_\odot}$, with the faintest candidates approaching $\sim 10^4 {\rm M_\odot}$, suggesting that BH*-dominated sources range from the intermediate mass black hole to the quasar regime.

- We compute approximate number densities of BH*-dominated sources at $L_{5100} > 10^{43.5}\,{\rm erg\,s^{-1}}$ and find that they peak at $z \sim 5 - 6$ ($\sim 10^{-5}\,{\rm Mpc^{-3}dex^{-1}}$), and decline towards higher and lower redshifts. While the weak decline towards higher redshift is not significant, not least due to possible selection effects, we measure a drop in the number density by an order of magnitude from $z \sim 6$ to $z \sim 2$.

- Our number densities at $z \sim 5 - 7$ are a factor of $2.5 - 8$ below LRD number densities based on photometric V-shape selections in the literature, but they are only a factor of $1.7 - 3.5$ below estimates based on V-shapes identified from ground-based imaging at $z \sim 2 - 3$. Despite some caveats regarding the comparison of these partly complementary samples, this suggests that the fraction of BH*-dominated sources among the total LRD population does not decline towards lower redshifts, or that lower redshift samples of LRDs are still highly incomplete.

The sample presented here provides ideal targets for spectroscopic follow-up, enabling unique insights into the physics driving the SEDs of LRDs through a direct view into their engines that outshines the host galaxy in the rest-frame optical. Further, it highlights the power of identifying such sources through template-based searches. The robustness of our sample selection is limited by the photometric wavelength-coverage and depth of the available NIRCam-imaging, by the extent to which our templates capture intrinsic variations in the spectra of LRD engines, and by how accurately they trace the continuum shapes, especially at $\lambda_{\rm rest} \gtrsim 1\,\mu$m. Future studies including MIRI data as well as additional medium-band imaging as is for example obtained through the cycle 4 program MINERVA (GO-7814, PI

Muzzin, Muzzin et al. 2025) will likely help improve the sample selection, and shed further light on the rest-frame optical to NIR properties of BH*-dominated objects and LRDs more generally, as well as their number density evolution. NIRSpec spectra of a larger sample of BH*-dominated sources will be crucial to better understand the diversity in their continuum shapes, and the physical mechanisms behind their peculiar and still enigmatic appearance.

ACKNOWLEDGMENTS

This work has received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number MB22.00072, as well as from the Swiss National Science Foundation (SNSF) through project grant 200020_207349. RPN, WQS, and ZL acknowledge funding from *JWST* programs GO-3516, GO-5224, and the MIT Undergraduate Research Opportunities Program (UROP). Support for this work was provided by NASA through the NASA Hubble Fellowship grant HST-HF2-51515.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555. RPN thanks Neil Pappalardo and Jane Pappalardo for their generous support of the MIT Pappalardo Fellowships in Physics, and for their enthusiasm and encouragement for pursuing the earliest galaxies and black holes. This research was supported by the International Space Science Institute (ISSI) in Bern, through ISSI International Team 25-659, "Little Red Dots, Big Open Questions: Unraveling the Mystery of the James Webb Space Telescope's Most Debated Discovery" led by MX and RPN. AdG acknowledges support from a Clay Fellowship awarded by the Smithsonian Astrophysical Observatory. REH acknowledges support by the German Aerospace Center (DLR) and the Federal Ministry for Economic Affairs and Energy (BMWi) through program 50OR2403 'RUBIES'. JM and AT acknowledge funding from the European Union (ERC, AGENTS, 101076224). Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. The data products presented herein were retrieved from the Dawn JWST Archive (DJA). DJA is an initiative of the Cosmic Dawn Center (DAWN), which is funded by the Danish National Research Foundation under grant DNRF140. This work is based 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. Support for programs #2514, was provided by NASA through grants from the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-03127.

Software used in developing this work includes: `matplotlib` (Hunter 2007), `jupyter` (Kluyver et al. 2016), `IPython` (Pérez & Granger 2007), `numpy` (Oliphant 2015), `scipy` (Virtanen et al. 2020), and `Astropy` (Astropy Collaboration et al. 2013).

REFERENCES

Akins, H. B., Casey, C. M., Lambrides, E., et al. 2025, ApJ, 991, 37
Ananna, T. T., Bogdán, Á., Kovács, O. E., Natarajan, P., & Hickox, R. C. 2024, ApJ, 969, L18
Ando, M., Harikane, Y., Katz, H., Inayoshi, K., & Tanaka, T. S. 2026, arXiv e-prints, arXiv:2606.03522
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33
Barro, G., Pérez-González, P. G., Kocevski, D. D., et al. 2024, ApJ, 963, 128
Barro, G., Pérez-González, P. G., Kocevski, D., et al. 2026, ApJ, 1003, 96
Begelman, M. C., & Dexter, J. 2026, ApJ, 996, 48
Begelman, M. C., Rossi, E. M., & Armitage, P. J. 2008, MNRAS, 387, 1649
Beiler, S. A., Cushing, M. C., Kirkpatrick, J. D., et al. 2024, ApJ, 973, 107
Bezanson, R., Labbe, I., Whitaker, K. E., et al. 2024, ApJ, 974, 92
Billand, J.-B., Elbaz, D., Franco, M., et al. 2026, arXiv e-prints, arXiv:2604.11677
Brammer, G. 2018, Gbrammer/Grizli: Preliminary Release, Zenodo, doi:10.5281/zenodo.1146905
Brammer, G. B., van Dokkum, P. G., & Coppi, P. 2008, ApJ, 686, 1503
Brazzini, M., D'Eugenio, F., Maiolino, R., et al. 2026, arXiv e-prints, arXiv:2601.22214
Carnall, A. C., Begley, R., McLeod, D. J., et al. 2023, MNRAS, 518, L45
Chang, S.-J., Gronke, M., Matthee, J., & Mason, C. 2026, MNRAS, 545, staf2131
Chisholm, J., Berg, D. A., Boylan-Kolchin, M., et al. 2026, arXiv e-prints, arXiv:2602.15935
Cloonan, A. P., Whitaker, K. E., Manning, S. M., et al. 2026, arXiv e-prints, arXiv:2603.24700
Conroy, C., & Gunn, J. E. 2010, ApJ, 712, 833
Conroy, C., Gunn, J. E., & White, M. 2009, ApJ, 699, 486
Covelo-Paz, A., Giovinazzo, E., Oesch, P. A., et al. 2025, A&A, 694, A178
Covelo-Paz, A., Meuwly, C., Oesch, P. A., et al. 2026, A&A, 705, A155
Dayal, P., & Maiolino, R. 2026, A&A, 706, A72
de Graaff, A., Rix, H.-W., Naidu, R. P., et al. 2025a, A&A, 701, A168
de Graaff, A., Hviding, R. E., Naidu, R. P., et al. 2025b, arXiv e-prints, arXiv:2511.21820
de Graaff, A., Brammer, G., Weibel, A., et al. 2025c, A&A, 697, A189
De Lucia, G., Fontanot, F., Hirschmann, M., & Xie, L. 2025, arXiv e-prints, arXiv:2502.01724
Donnan, C. T., Dickinson, M., Taylor, A. J., et al. 2025, ApJ, 993, 224
Eisenstein, D. J., Willott, C., Alberts, S., et al. 2026, ApJS, 283, 6
Fei, Q., Fujimoto, S., Brammer, G., et al. 2026, arXiv e-prints, arXiv:2602.12325
Ferland, G. J., Chatzikos, M., Guzmán, F., et al. 2017, Rev. Mexicana Astron. Astrofis., 53, 385
Furtak, L. J., Zitrin, A., Plat, A., et al. 2023a, ApJ, 952, 142
Furtak, L. J., Zitrin, A., Weaver, J. R., et al. 2023b, MNRAS, 523, 4568
Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, A&A, 595, A1
Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, A&A, 674, A1
Gandolfi, G., Rodighiero, G., Castellano, M., et al. 2026, A&A, 706, A364
Grant Roberts, M., Braff, L., Garg, A., et al. 2025, J. Cosmology Astropart. Phys., 2025, 060
Greene, J. E., Strader, J., & Ho, L. C. 2020, ARA&A, 58, 257
Greene, J. E., Labbe, I., Goulding, A. D., et al. 2024, ApJ, 964, 39
Greene, J. E., Setton, D. J., Furtak, L. J., et al. 2026, ApJ, 996, 129
Grogin, N. A., Kocevski, D. D., Faber, S. M., et al. 2011, ApJS, 197, 35
Huang, H.-L., Jiang, J.-Q., He, J., Wang, Y.-T., & Piao, Y.-S. 2024, arXiv e-prints, arXiv:2410.20663
Hunter, J. D. 2007, Computing In Science & Engineering, 9, 90

Hviding, R. E., de Graaff, A., Miller, T. B., et al. 2025, A&A, 702, A57
Inayoshi, K. 2025, ApJ, 988, L22
Inayoshi, K., & Maiolino, R. 2025, ApJ, 980, L27
Ito, K., Valentino, F., Brammer, G., et al. 2025, arXiv e-prints, arXiv:2506.22642
Ji, X., Maiolino, R., Übler, H., et al. 2025, MNRAS, 544, 3900
Jiang, F., Jia, Z., Zheng, H., et al. 2026, ApJ, 996, L19
Kashino, D., Lilly, S. J., Matthee, J., et al. 2023, ApJ, 950, 66
Kass, R. E., & Raftery, A. E. 1995, J. Am. Statist. Assoc., 90, 773
Kido, D., Ioka, K., Hotokezaka, K., Inayoshi, K., & Irwin, C. M. 2025, MNRAS, arXiv:2505.06965
Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in Positioning and Power in Academic Publishing: Players, Agents and Agendas, ed. F. Loizides & B. Schmidt, IOS Press, 87 – 90
Kocevski, D. D., Finkelstein, S. L., Barro, G., et al. 2025, ApJ, 986, 126
Koekemoer, A. M., Faber, S. M., Ferguson, H. C., et al. 2011, ApJS, 197, 36
Kokorev, V., Fujimoto, S., Labbe, I., et al. 2023, ApJ, 957, L7
Kokorev, V., Caputi, K. I., Greene, J. E., et al. 2024, ApJ, 968, 38
Kokubo, M., & Harikane, Y. 2025, ApJ, 995, 24
Kormendy, J., & Ho, L. C. 2013, ARA&A, 51, 511
Labbe, I., Greene, J. E., Matthee, J., et al. 2024, arXiv e-prints, arXiv:2412.04557
Labbe, I., Greene, J. E., Bezanson, R., et al. 2025, ApJ, 978, 92
Lambrides, E., Larson, R. L., Garofali, K., et al. 2026, Nature Astronomy, arXiv:2409.13047
Latif, M. A., Aftab, A., Whalen, D. J., & Mezcua, M. 2025, A&A, 694, L14
Lin, X., Fan, X., Cai, Z., et al. 2026, ApJ, 997, 364
Liu, H., Jiang, Y.-F., Quataert, E., Greene, J. E., & Ma, Y. 2025, ApJ, 994, 113
Liu, H., Jiang, Y.-F., Quataert, E., et al. 2026a, arXiv e-prints, arXiv:2603.02317
Liu, Z., Naidu, R. P., Secunda, A., et al. 2026b, arXiv e-prints, arXiv:2604.13000
Looser, T. J., D'Eugenio, F., Maiolino, R., et al. 2024, Nature, 629, 53
Lupi, A., Trinca, A., Volonteri, M., Dotti, M., & Mazzucchelli, C. 2024, A&A, 689, A128
Ma, Y., Greene, J. E., Volonteri, M., et al. 2025a, arXiv e-prints, arXiv:2509.02662
Ma, Y., Greene, J. E., Setton, D. J., et al. 2025b, ApJ, 981, 191
—. 2026, ApJ, 1000, 59
Madau, P. 2026, A&A, 708, A116
Madau, P., & Maiolino, R. 2026a, arXiv e-prints, arXiv:2602.22386
—. 2026b, arXiv e-prints, arXiv:2605.05074
Matthee, J., Mackenzie, R., Simcoe, R. A., et al. 2023, ApJ, 950, 67
Matthee, J., Naidu, R. P., Brammer, G., et al. 2024, ApJ, 963, 129
Matthee, J., Torralba, A., Pezzulli, G., et al. 2026, arXiv e-prints, arXiv:2603.17667
Meisner, A. M., Schneider, A. C., Burgasser, A. J., et al. 2021, ApJ, 915, 120
Mérida, R. M., Sawicki, M., Gaspar, G., Willott, C. J., & Iyer, K. G. 2026, arXiv e-prints, arXiv:2605.07976
Mukherjee, S., Fortney, J. J., Morley, C. V., et al. 2024, ApJ, 963, 73
Muzzin, A., Suess, K. A., Marchesini, D., et al. 2025, arXiv e-prints, arXiv:2507.19706
Naidu, R. P., Matthee, J., Kramarenko, I., et al. 2024, arXiv e-prints, arXiv:2410.01874
Naidu, R. P., Matthee, J., Katz, H., et al. 2025, arXiv e-prints, arXiv:2503.16596
Nanayakkara, T., Glazebrook, K., Schreiber, C., et al. 2025, ApJ, 981, 78
Nandal, D., & Loeb, A. 2026, ApJ, 998, 124
Nikopoulos, G. P., Watson, D., Sneppen, A., et al. 2025, arXiv e-prints, arXiv:2510.06362
Oesch, P. A., Brammer, G., Naidu, R. P., et al. 2023, MNRAS, 525, 2864
Oke, J. B., & Gunn, J. E. 1983, ApJ, 266, 713
Oliphant, T. E. 2015, Guide to NumPy (Continuum Press)
Pacucci, F., Ferrara, A., & Kocevski, D. D. 2026, arXiv e-prints, arXiv:2601.14368
Pan, Z., Zhuang, M.-Y., Shen, Y., et al. 2026, arXiv e-prints, arXiv:2606.09721
Pérez, F., & Granger, B. E. 2007, Computing in Science and Engineering, 9, 21
Pérez-González, P. G., Barro, G., Rieke, G. H., et al. 2024, ApJ, 968, 4
Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, A&A, 641, A6
Price, S. H., Bezanson, R., Labbe, I., et al. 2025, ApJ, 982, 51
Rinaldi, P., Hainline, K., D'Eugenio, F., et al. 2026, arXiv e-prints, arXiv:2604.07138
Rusakov, V., Watson, D., Nikopoulos, G. P., et al. 2026, Nature, 649, 574
Santarelli, A. D., Farag, E., Bellinger, E. P., et al. 2026, ApJ, 998, L4
Secunda, A., Somerville, R. S., Jiang, Y.-F., et al. 2026, ApJ, 996, 6
Setton, D. J., Greene, J. E., Spilker, J. S., et al. 2025a, ApJ, 991, L10
Setton, D. J., Greene, J. E., de Graaff, A., et al. 2025b, ApJ, 995, 118
Sneppen, A., Watson, D., Matthews, J. H., et al. 2026, arXiv e-prints, arXiv:2601.18864
Suess, K. A., Weaver, J. R., Price, S. H., et al. 2024, ApJ, 976, 101
Sun, W. Q., Naidu, R. P., Matthee, J., et al. 2026, arXiv e-prints, arXiv:2601.20929
Taylor, A. J., Kokorev, V., Kocevski, D. D., et al. 2025, ApJ, 989, L7
Torralba, A., Matthee, J., Weibel, A., et al. 2026a, arXiv e-prints, arXiv:2603.28335
Torralba, A., Matthee, J., Pezzulli, G., et al. 2026b, A&A, 707, A75
Umeda, H., Inayoshi, K., Harikane, Y., & Murase, K. 2025, arXiv e-prints, arXiv:2512.04208
Valentino, F., Brammer, G., Gould, K. M. L., et al. 2023, ApJ, 947, 20
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261
Volonteri, M., Habouzit, M., & Colpi, M. 2021, Nature Reviews Physics, 3, 732
Wang, B., Leja, J., de Graaff, A., et al. 2024, ApJ, 969, L13
Wang, B., de Graaff, A., Davies, R. L., et al. 2025, ApJ, 984, 121
Wang, B., Leja, J., Labbe, I., et al. 2026, arXiv e-prints, arXiv:2602.06024
Weaver, J. R., Cutler, S. E., Pan, R., et al. 2024, ApJS, 270, 7
Weibel, A., Oesch, P. A., Barrufet, L., et al. 2024, MNRAS, 533, 1808
Weibel, A., Oesch, P. A., Williams, C. C., et al. 2026, ApJ, 1002, 136
Weinberger, R., Springel, V., Pakmor, R., et al. 2018, MNRAS, 479, 4056
Williams, C. C., Alberts, S., Ji, Z., et al. 2024, ApJ, 968, 34
Williams, C. C., Oesch, P. A., Weibel, A., et al. 2025, ApJ, 979, 140
Xiao, M., Oesch, P. A., Bing, L., et al. 2025, A&A, 700, A231
Yue, M., Eilers, A.-C., Ananna, T. T., et al. 2024, ApJ, 974, L26
Zhang, B., Feng, W.-X., & An, H. 2025, arXiv e-prints, arXiv:2507.07171
Zhang, S., Liu, B., Bromm, V., & Kühnel, F. 2026, ApJ, 1000, L19

APPENDIX

A. LIST OF JWST PROGRAMS CONTRIBUTING NIRSPEC SPECTRA

As described in Section 3.5, we cross-match our catalog of BH*-dominated candidates with all public spectra available on the DJA as of April 29, 2026. Considering only the highest SNR spectrum of each source (PRISM if available, medium resolution grating otherwise), these come from the following programs: 19 PRISM spectra and one G395M spectrum are from RUBIES (GO-4233, PIs de Graaff & Brammer, de Graaff et al. 2025c); CAPERS (GO-6368, PI Dickinson, e.g., Donnan et al. 2025) and Mirage or Miracle (MoM, GO-5224, PIs Oesch & Naidu) each contribute seven PRISM spectra; another six are from UNCOVER (GO-2561, PI Labbé, Bezanson et al. 2024); JADES (e.g., Eisenstein et al. 2026) through GTO-1180, GTO-1181 (PI Eisenstein), and GTO-1287 (PI Isaak), and GO-5545 (PI Barrufet) each contribute three; another two PRISM spectra are from GO-8060 (PI Egami); finally, one spectrum is associated with each of the programs GO-2565 (PI Glazebrook, Nanayakkara et al. 2025), GO-4106 (PI Nel-

son), GO-5997 (PI Looser), and DIVER (GO-8018, PI Lin). The remaining three medium grating spectra (one in G395M and two in G235M) are from GO-3567 (PI Valentino, Ito et al. 2025).

B. SPECTROSCOPIC V-SHAPE CLASSIFICATION

To provide an estimate of the contamination rate of our photometric selection, we apply the spectroscopic V-shape classification from Hviding et al. (2025) to the subset of 57 sources with grade 3 PRISM spectra. One spectrum is immediately identified as a BD (Section 4.3), and removed from the sample, another 34 sources are part of the spectroscopic LRD sample of de Graaff et al. (2025b) and considered robust LRDs. Therefore, we only apply the V-shape classification to the remaining 22 spectra. Specifically, we fit a power-law, $f_\lambda = a \cdot \lambda_{\rm rest}^{\beta}$, to both the photometry and the spectrum on either side of H_∞. The V-shape is then defined following Section 3.2.2 of Hviding et al. (2025) as (1) $\beta_{\rm UV} < -0.2$ at the 2σ level, and $a_{\rm UV} > 0$ from either spectroscopy or photometry, (2) $\beta_{\rm opt}$(Spec.) > 0 at the 2σ level and $a_{\rm opt}$(Spec.) > 0, and (3) $\beta_{\rm opt} - \beta_{\rm UV} > 0.5$ with $\beta_{\rm UV}$ from spectroscopy if it satisfies (1) or from photometry otherwise. According to this definition, 9/22 sources display V-shaped continua. However, for none of the other 13 sources can a V-shaped SED be confidently ruled out. Most sources do show $\beta_{\rm UV} < -0.2$ and $\beta_{\rm opt} > 0$ but lack the SNR to confirm these slopes at the $> 2\sigma$ level. For example, the source GDN-59854 which is shown in Figure 3 shows a clear red optical slope ($\beta_{\rm opt} = 1.39 \pm 0.09$), a tentatively broadened Hα line and a clear break/turnover around the Balmer limit. However, the SNR in both the photometry and the spectrum is insufficient to confirm a blue UV-slope for this source (e.g., $\beta_{\rm UV} = -0.21 \pm 0.54$ from the spectrum). Further, the spectrum of the source COS-EAST-85239 is affected by the chip gap and only covers $\lambda_{\rm rest} \gtrsim 0.7\mu$m, which means that neither slope can be measured. However, this spectrum shows tentatively broadened Paschen lines, and the photometry clearly favors a BH*-dominated interpretation. We only find one likely contaminant, UDS-27591, whose spectrum shows blue UV- and optical slopes ($\beta_{\rm UV} = -3.72 \pm 0.60$ and $\beta_{\rm opt} = -2.99 \pm 2.04$), albeit with large uncertainties. The spectrum shows strong lines that boost the NIRCam/LW photometry which has likely caused `eazy` to wrongly prefer a fit with a BH* template. Resulting estimates of the contamination rate of our sample are provided in Section 3.5.

C. LIKELY CONTAMINANT IN THE COSMOS FIELD

Upon visual inspection of all our BH*-dominated candidates, we identify COS-WEST-29852 as an extremely bright source that is fit with the BH* template based on GN-9771 (Torralba et al. 2026b) at $z_{\rm phot} = 5.61$. We show imaging cutouts, the measured photometry, and the best-fitting `eazy` SED of the source in Figure 14. If at $z_{\rm phot} = 5.61$, the implied optical luminosity of COS-WEST-29852 would be more than an order of magnitude higher than that of any other candidate in our sample ($\log(L_{5100}/{\rm erg\,s^{-1}}) = 45.68$). While the fit with the BH* template is nominally robust (Bayes factors $\gg 100$), the photometry shows no distinctive features such as a (Lyman or Balmer) break, or individual filters boosted by emission lines. Instead, it suggests a smoothly declining red SED which happens to be consistent with the GN-9771 template at the best-fitting redshift. The imaging cutouts further shows a dominant point source component in the NIRCam/LW filters, and a clearly extended morphology in the SW bands, suggesting it could be an obscured AGN or quasar (see, e.g., Fei et al. 2026) whose host galaxy dominates the light at $\lesssim 2\mu$m. Further investigation of COS-WEST-29852 at (R.A., Dec.) = (150.0646281, 2.1909212) is beyond the scope of this work, but we argue that it is likely a peculiar contaminant, and not a BH*-dominated source at $z_{\rm phot} = 5.61$, and therefore remove it from our sample.

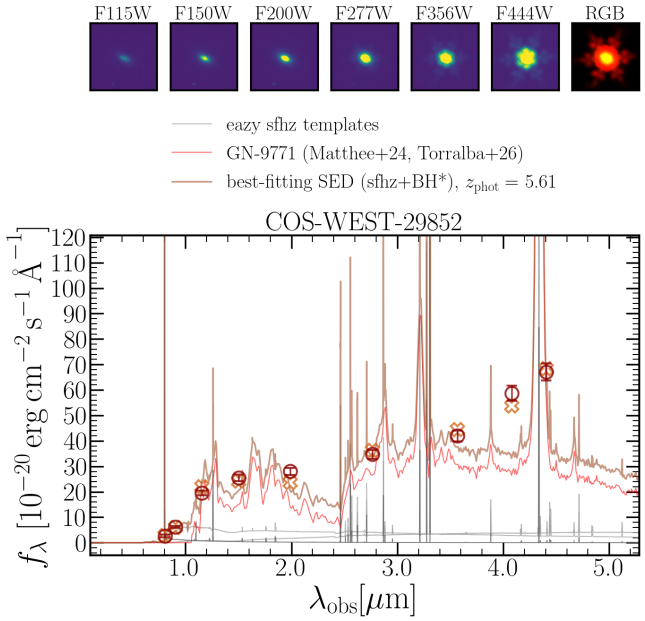

FIG. 14.— BH*-dominated candidate identified from a fit with the template based on GN-9771 at $z_{\rm phot} = 5.61$. Its brightness would imply an optical luminosity that is more than an order of magnitude higher than that of any other source in our sample. We argue that it is likely a contaminant and do not further investigate it here.

D. GALLERY OF ALL CANDIDATES IN THE GOLD SAMPLE

In Figures 15 to 27, we show the photometry, best-fitting `eazy` SEDs, and an RGB imaging stamp for each BH*-dominated candidate that is part of our gold sample (see Section 4.4). The gold candidates are sorted in redshift, and we specify their photometric and - if available - spectroscopic redshift in the title of each panel respectively. The contribution of the respective BH* template to the best-fitting SED, as well as the contributions of `sfhz` templates are shown in each panel.

This paper was built using the Open Journal of Astrophysics LaTeX template. The OJA is a journal which provides fast and easy peer review for new papers in the `astro-ph` section of the arXiv, making the reviewing process simpler for authors and referees alike. Learn more at `http://astro.theoj.org`.

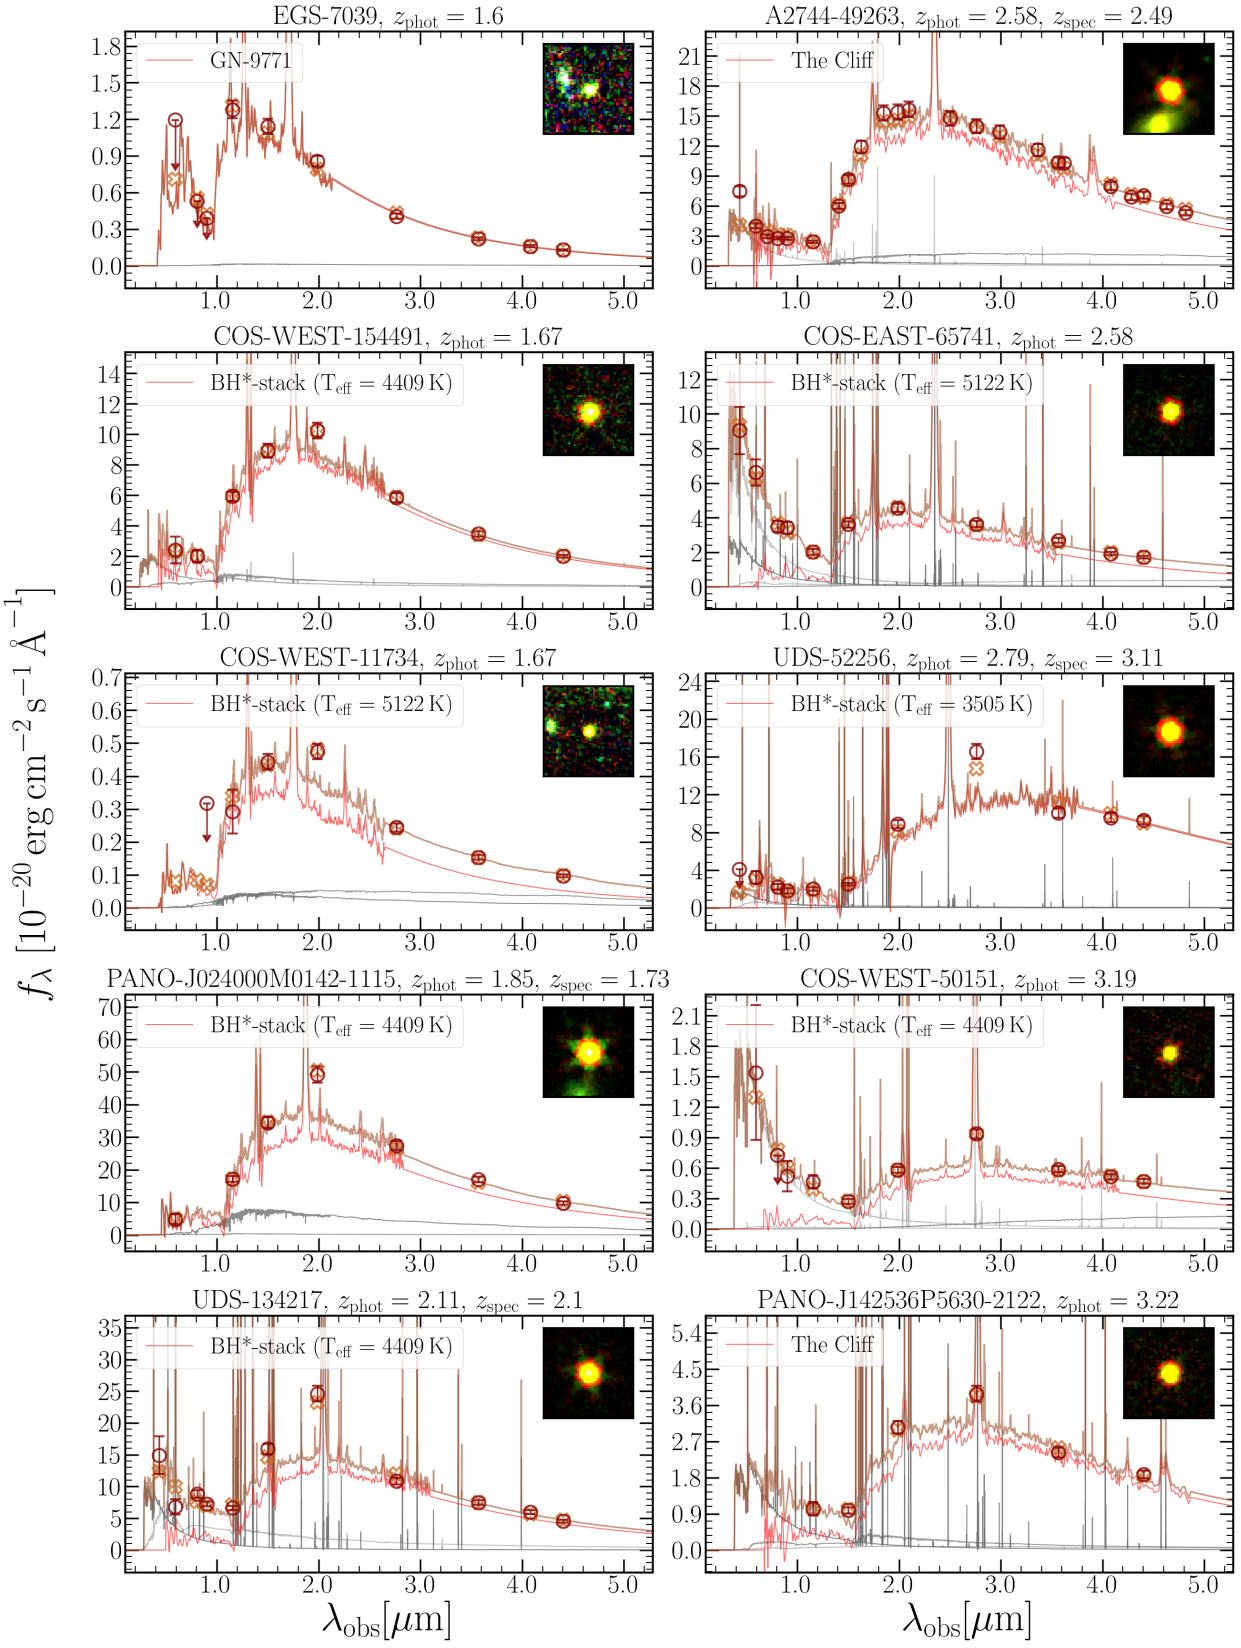

FIG. 15.— Gallery of BH*-dominated candidates that are part of our gold sample (B(`sfhz`) > 100 and B(stars) > 100), in analogy to Figure 6. The RGB stamp shown for each source is constructed from F444W (red), F277W (green) and F115W (blue). Gray lines show the contribution of different `eazy` templates to the best-fitting SED (brown line). The contribution of the respective BH* template that dominates the fit is shown in red.

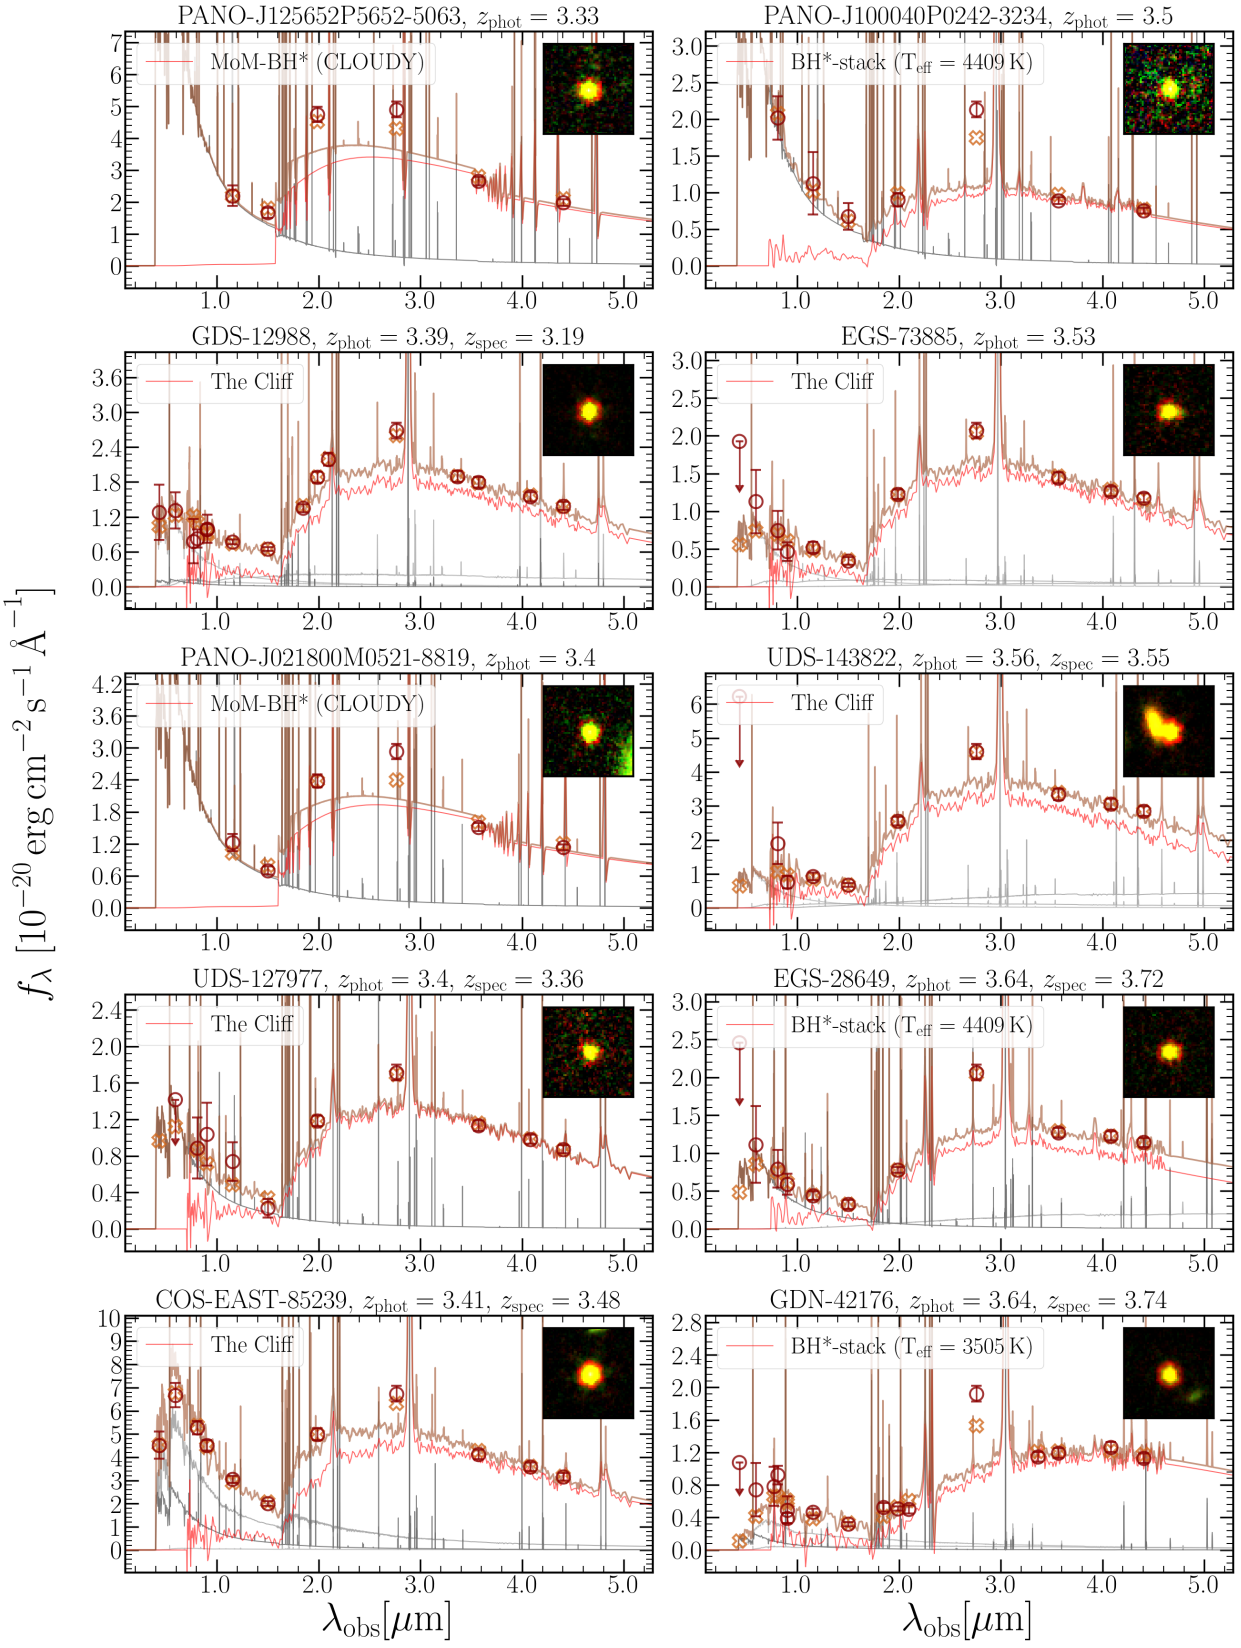

FIG. 16.— Same as Figure 15 (continued).

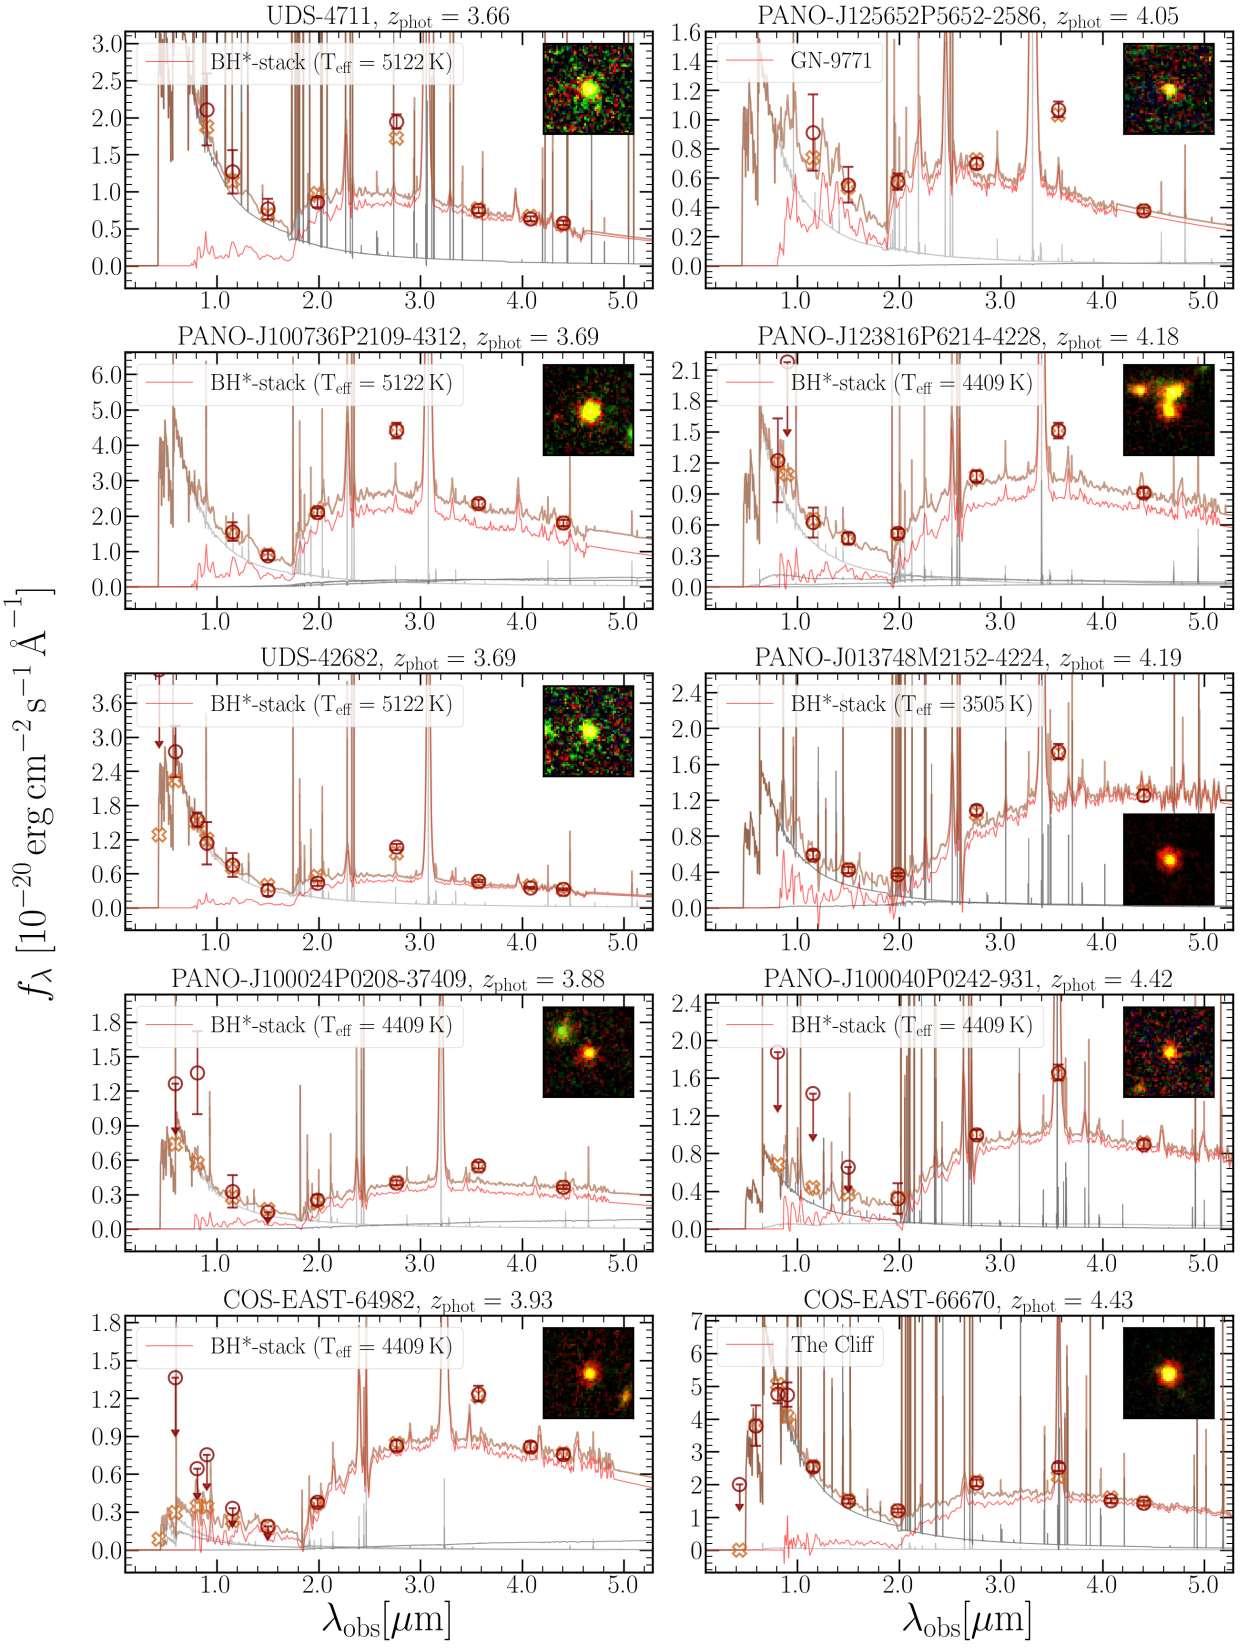

FIG. 17.— Same as Figure 15 (continued).

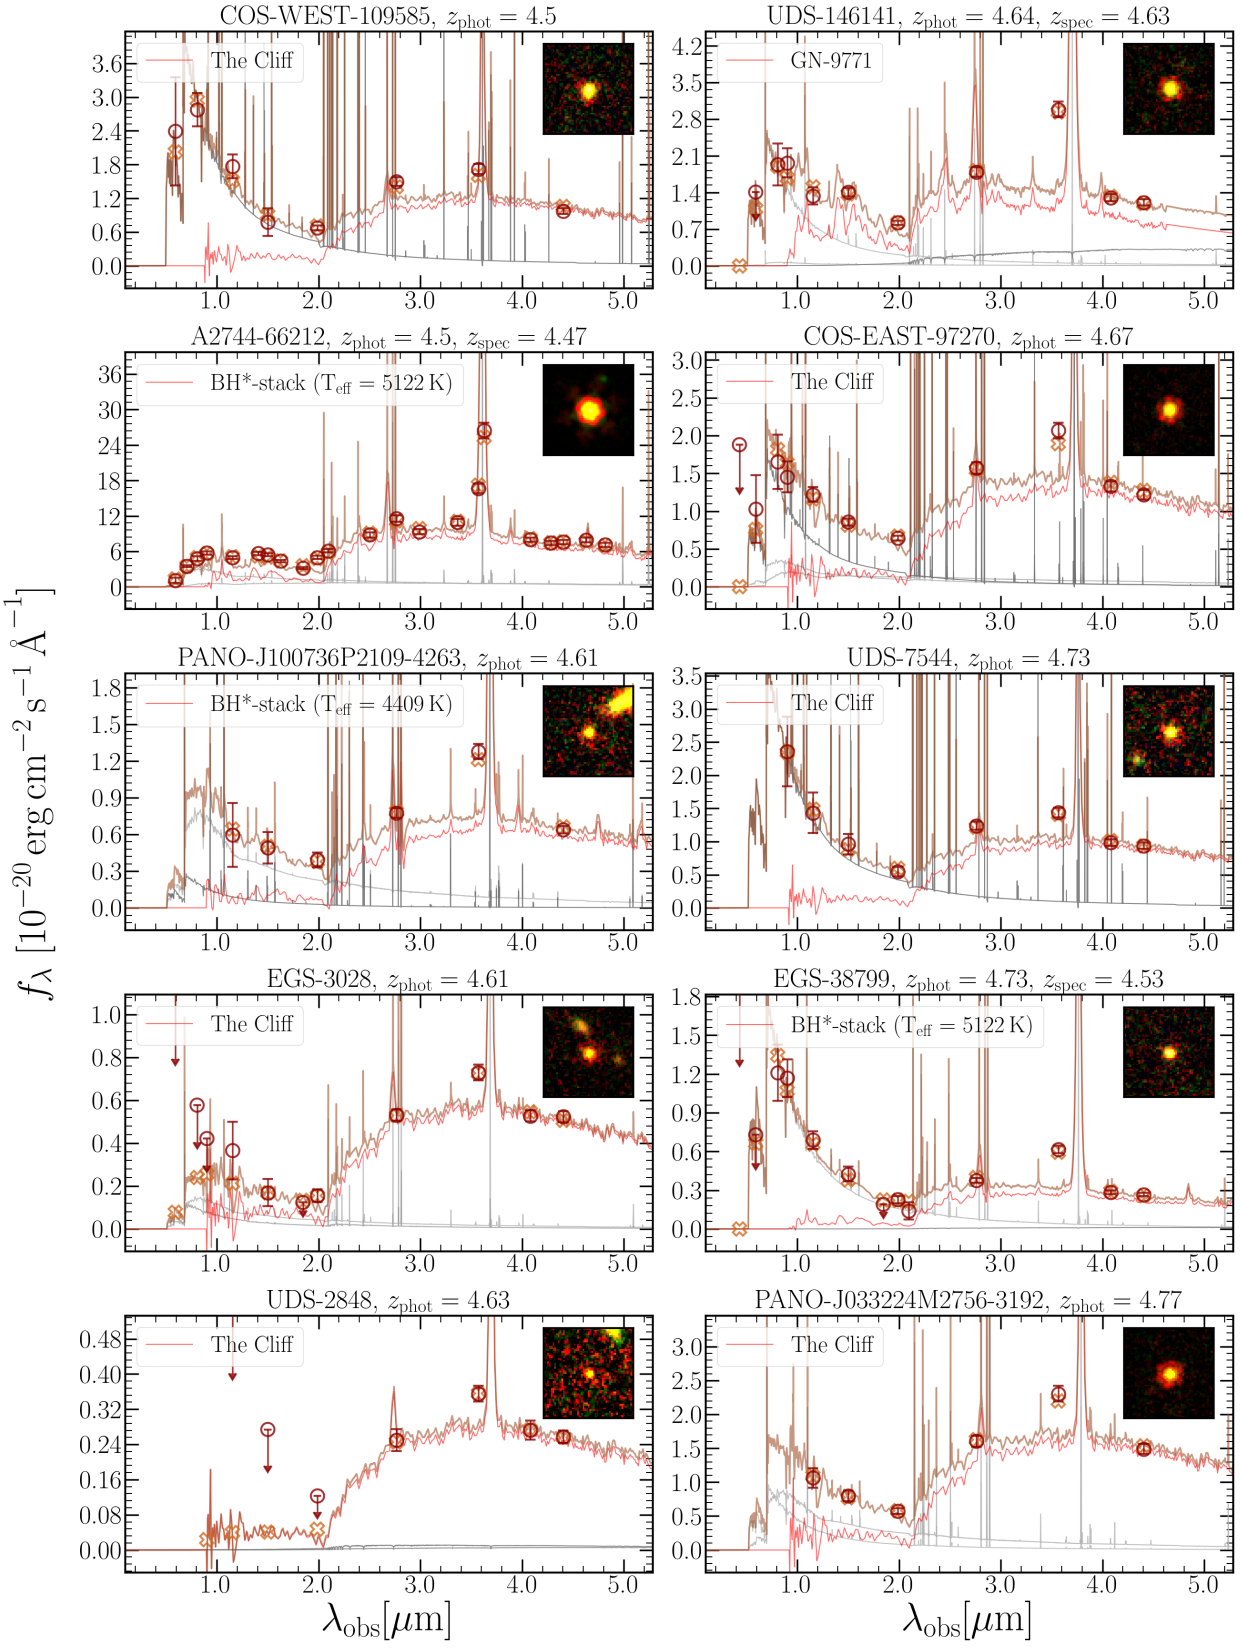

FIG. 18.— Same as Figure 15 (continued).

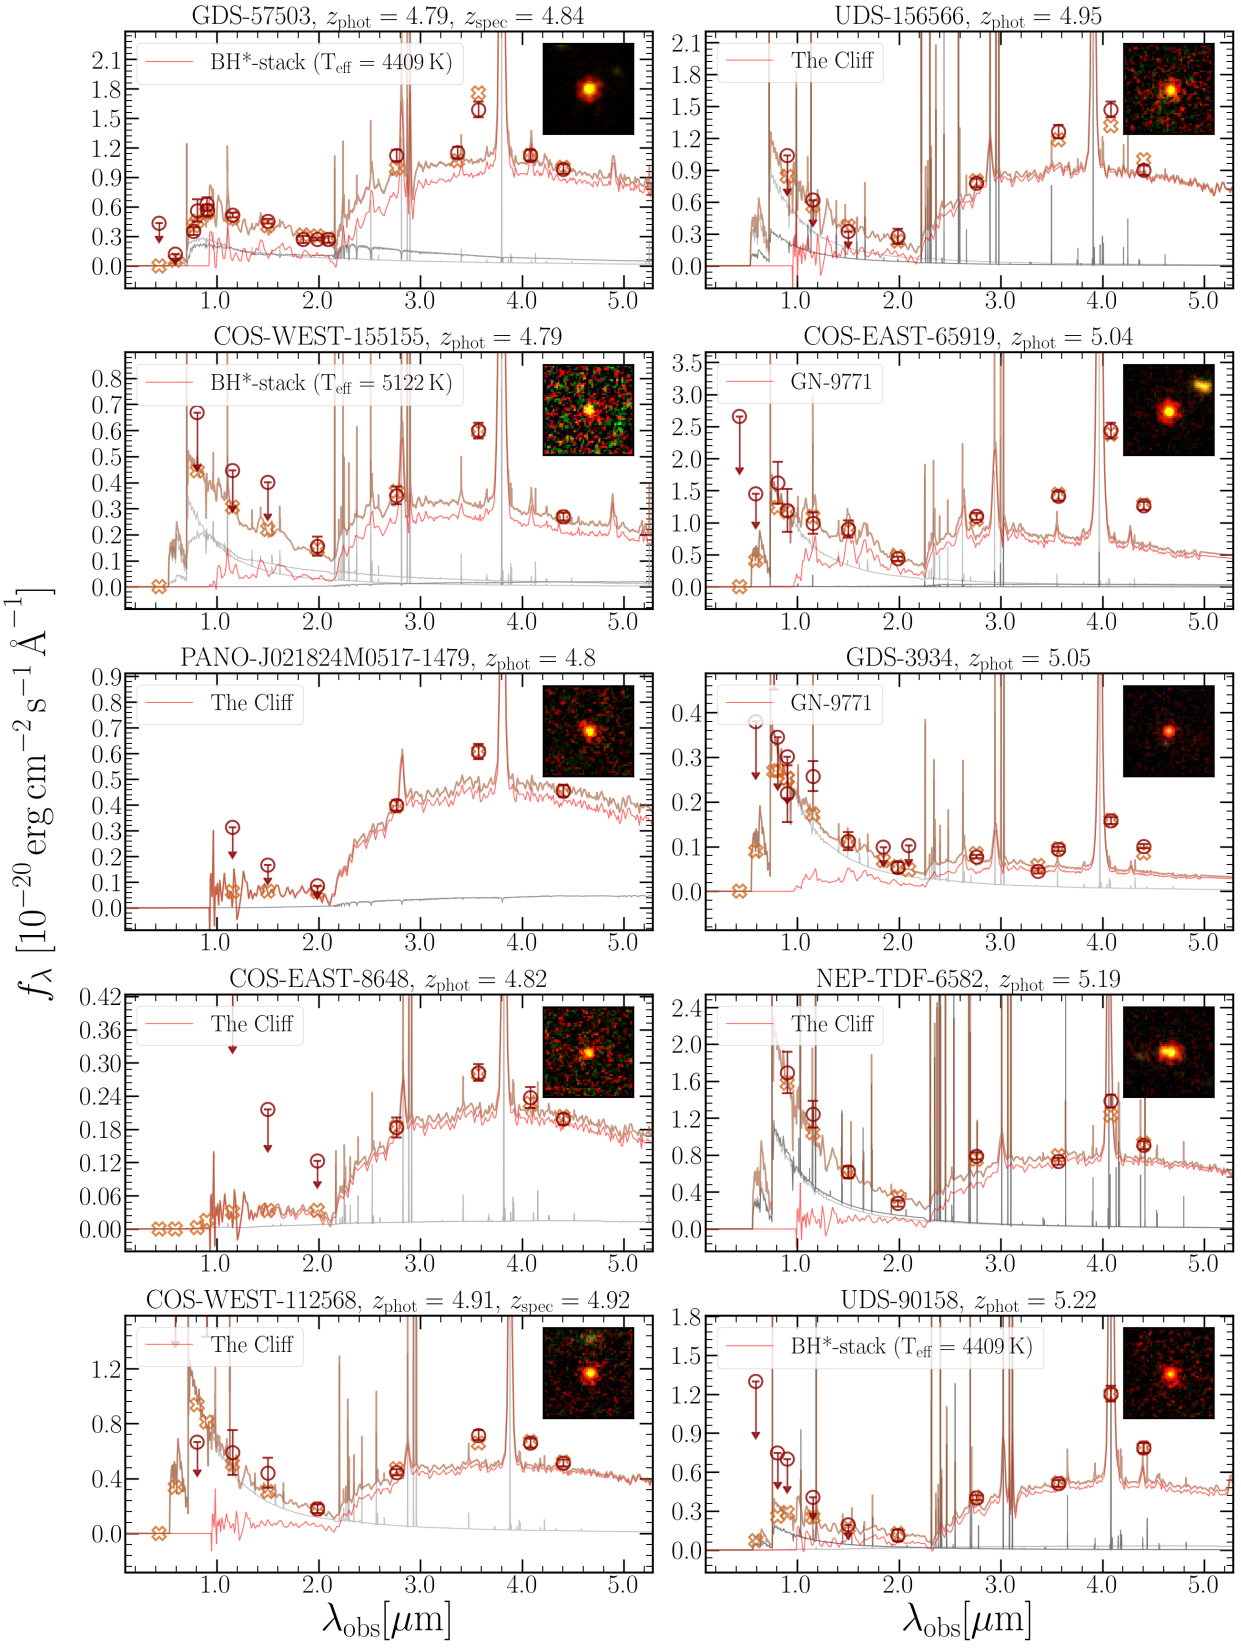

FIG. 19.— Same as Figure 15 (continued).

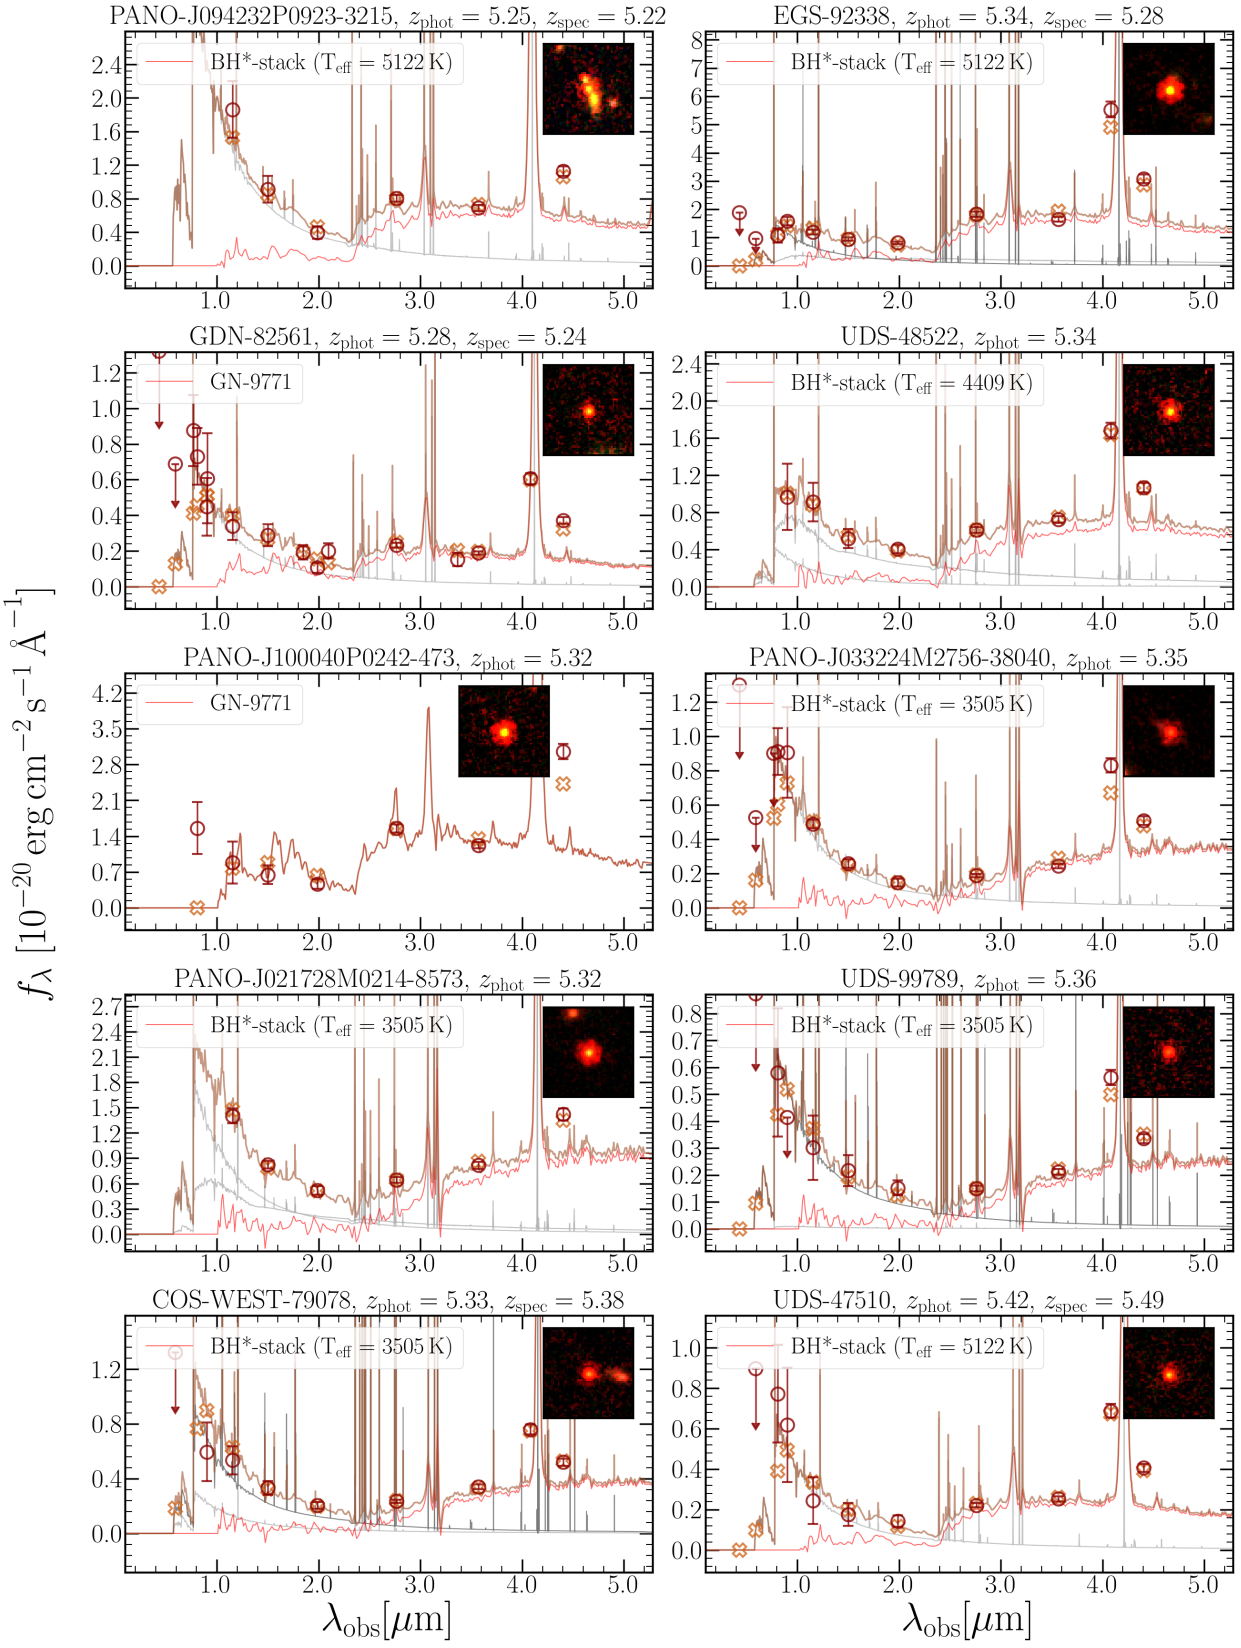

FIG. 20.— Same as Figure 15 (continued).

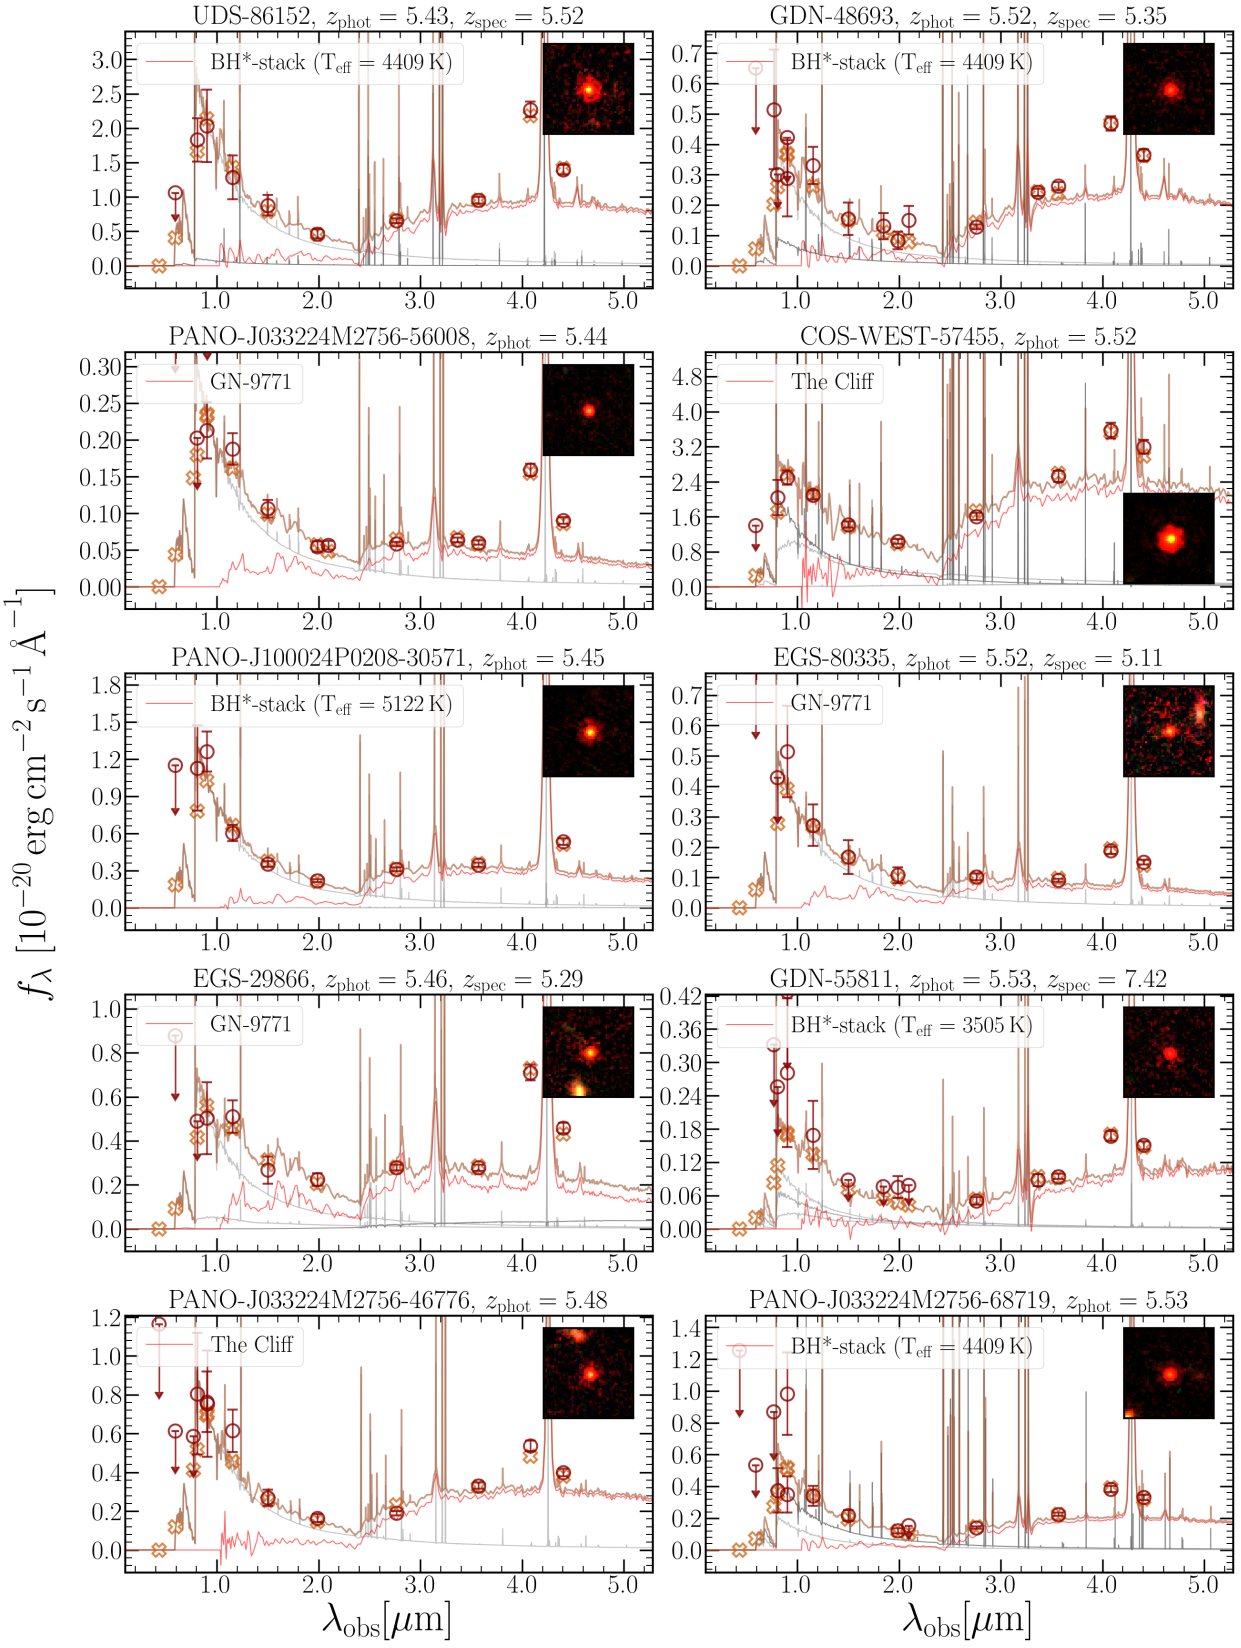

FIG. 21.— Same as Figure 15 (continued).

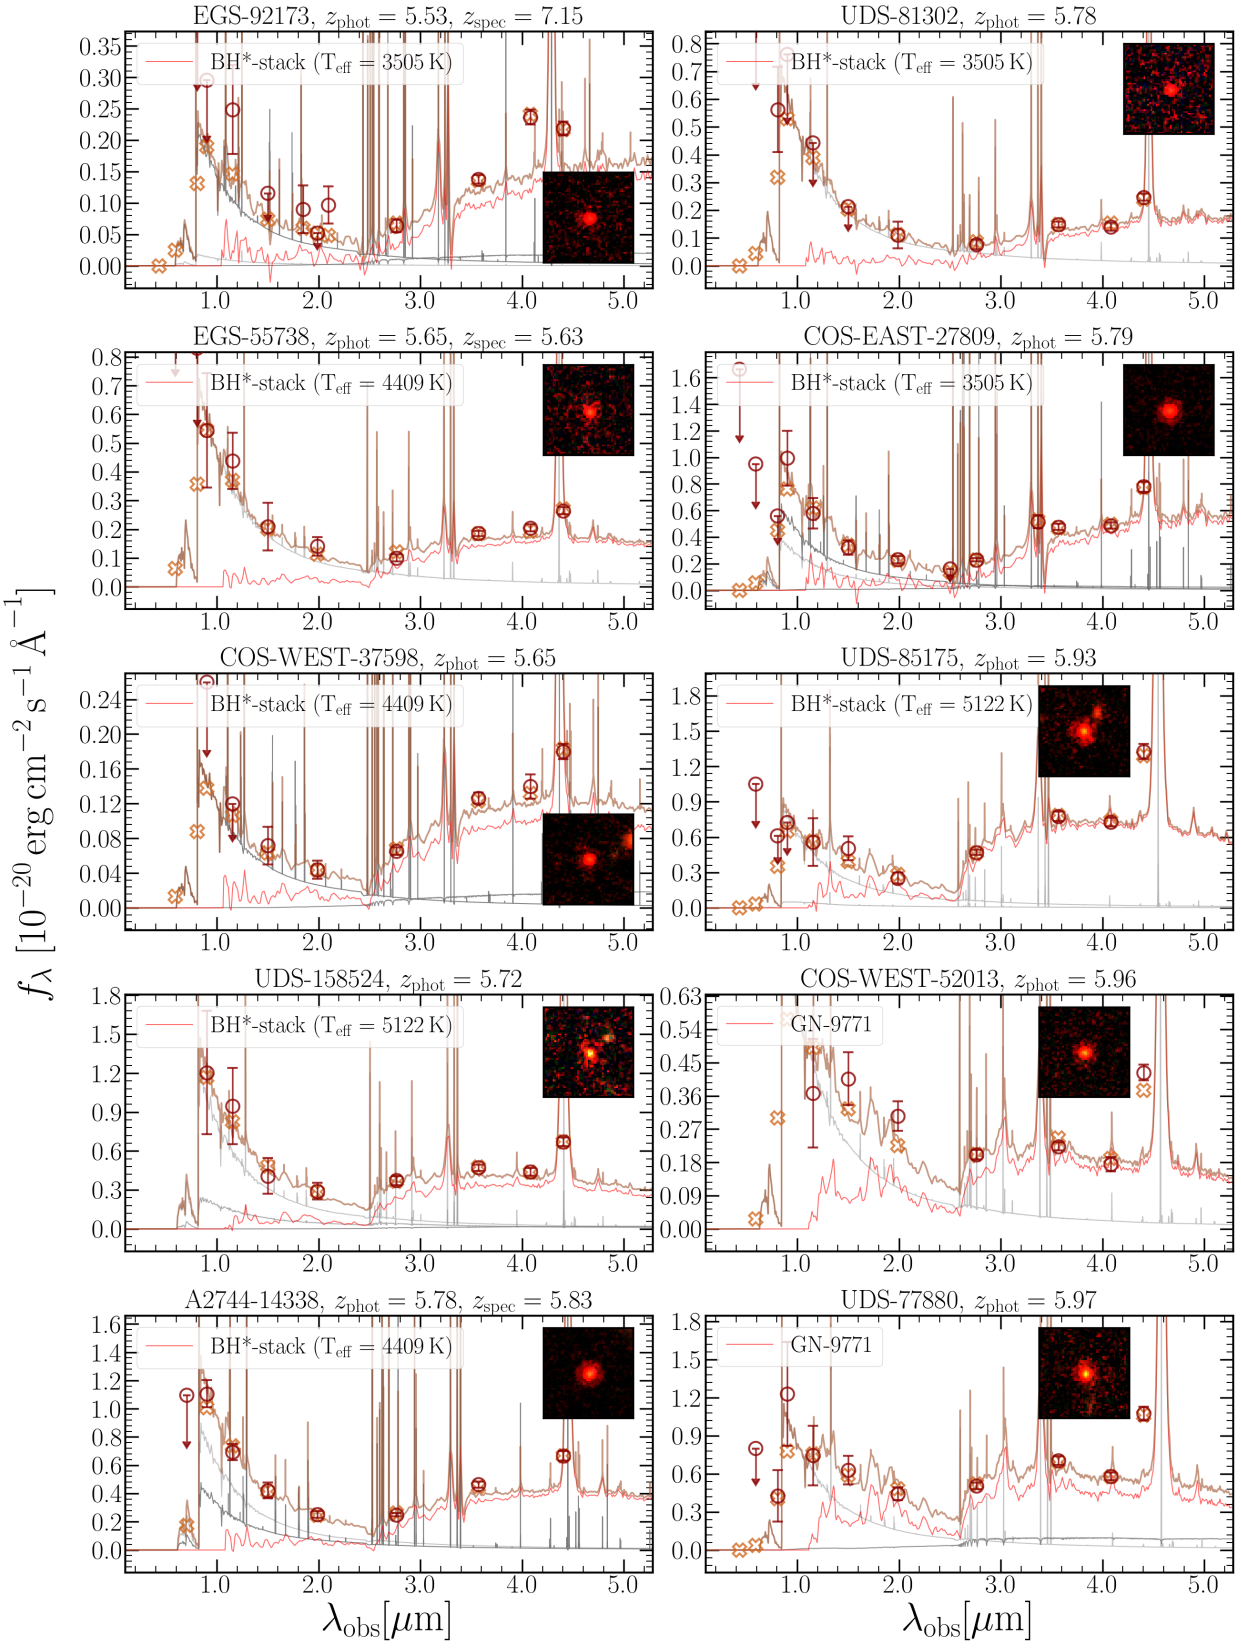

FIG. 22.— Same as Figure 15 (continued).

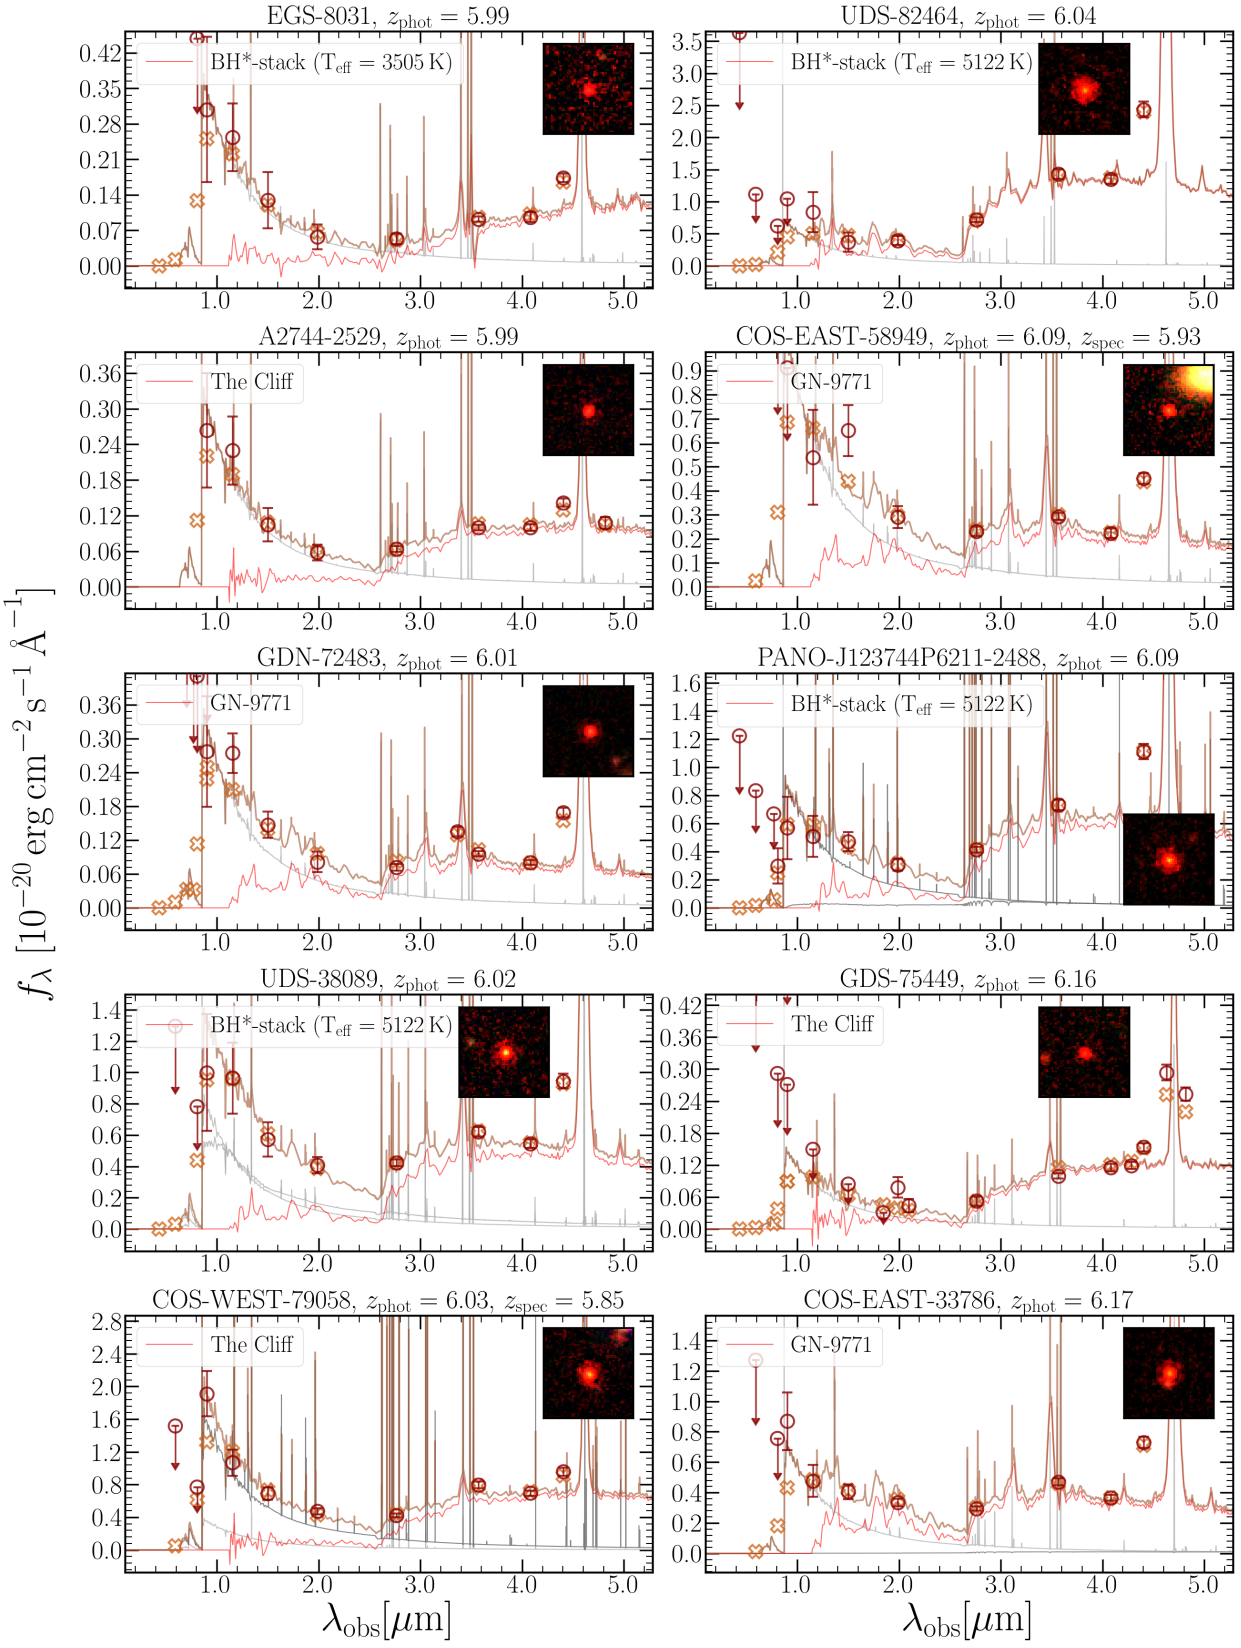

FIG. 23.— Same as Figure 15 (continued).

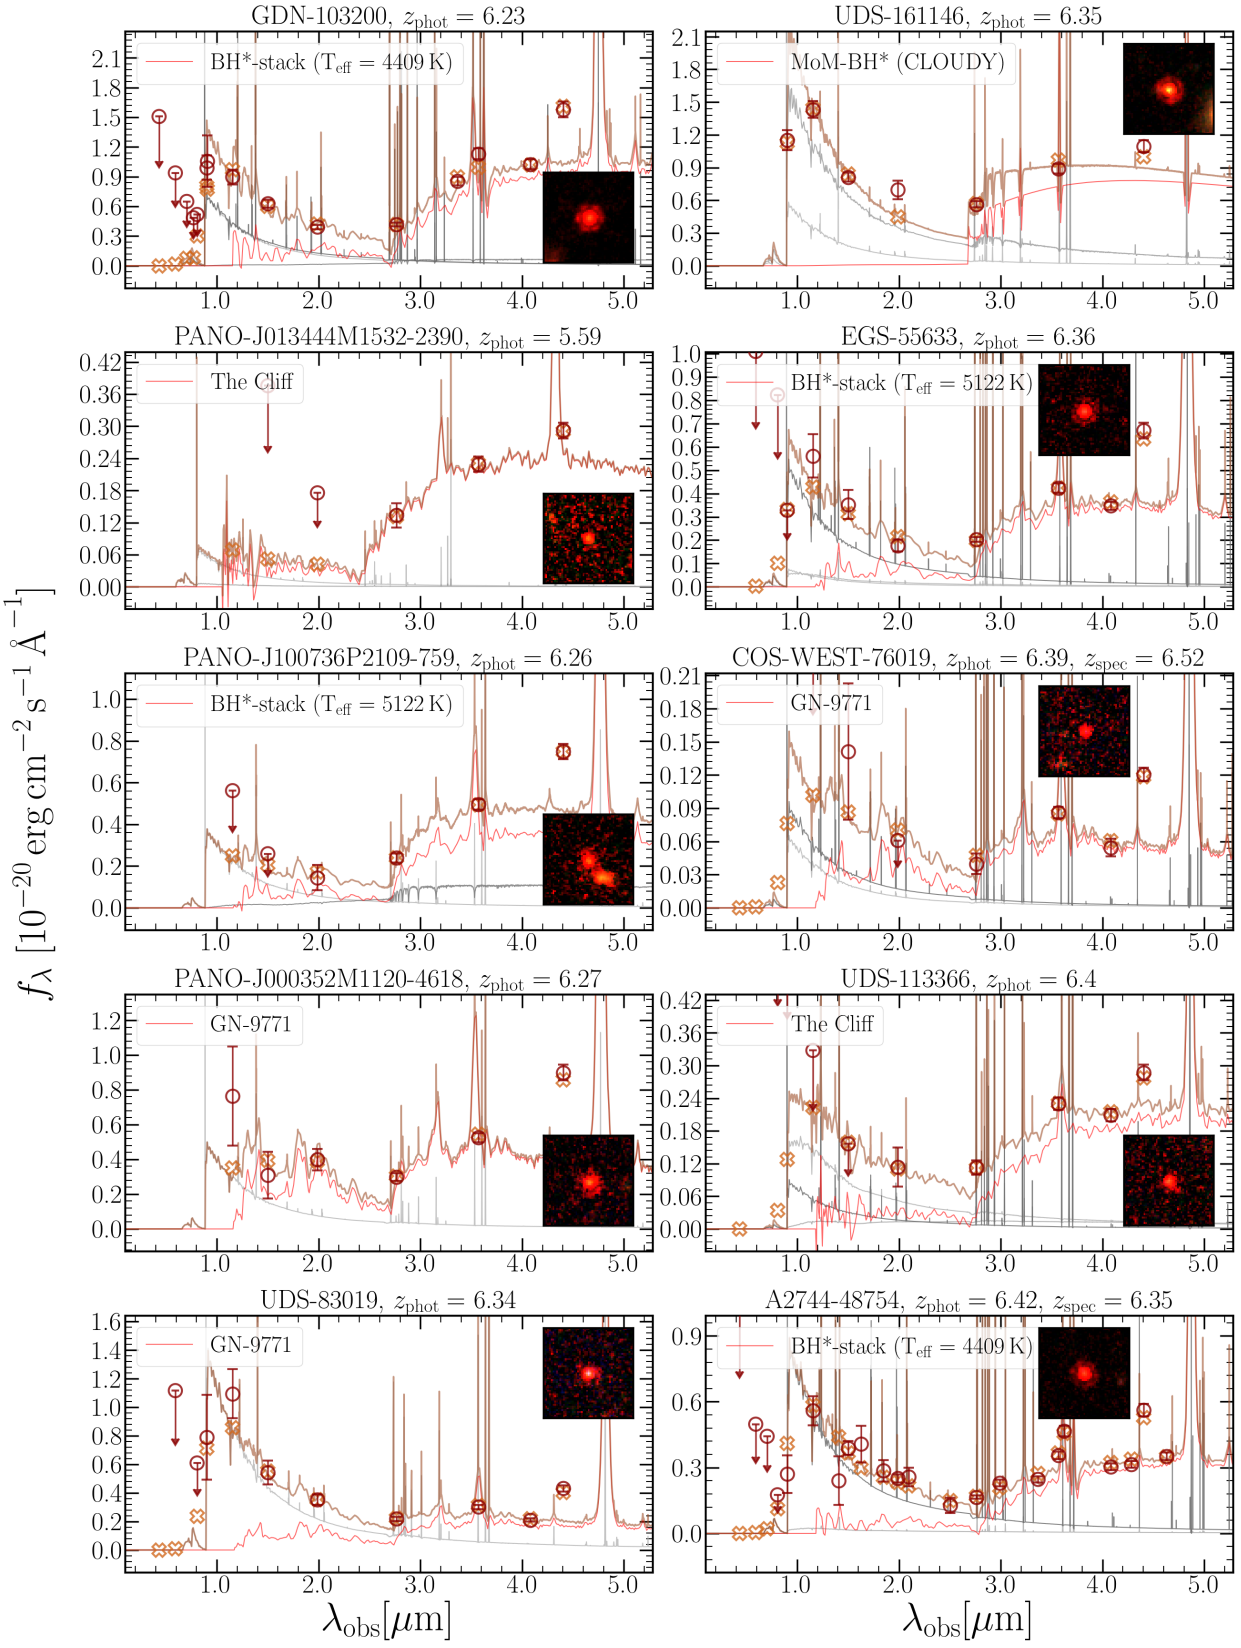

FIG. 24.— Same as Figure 15 (continued).

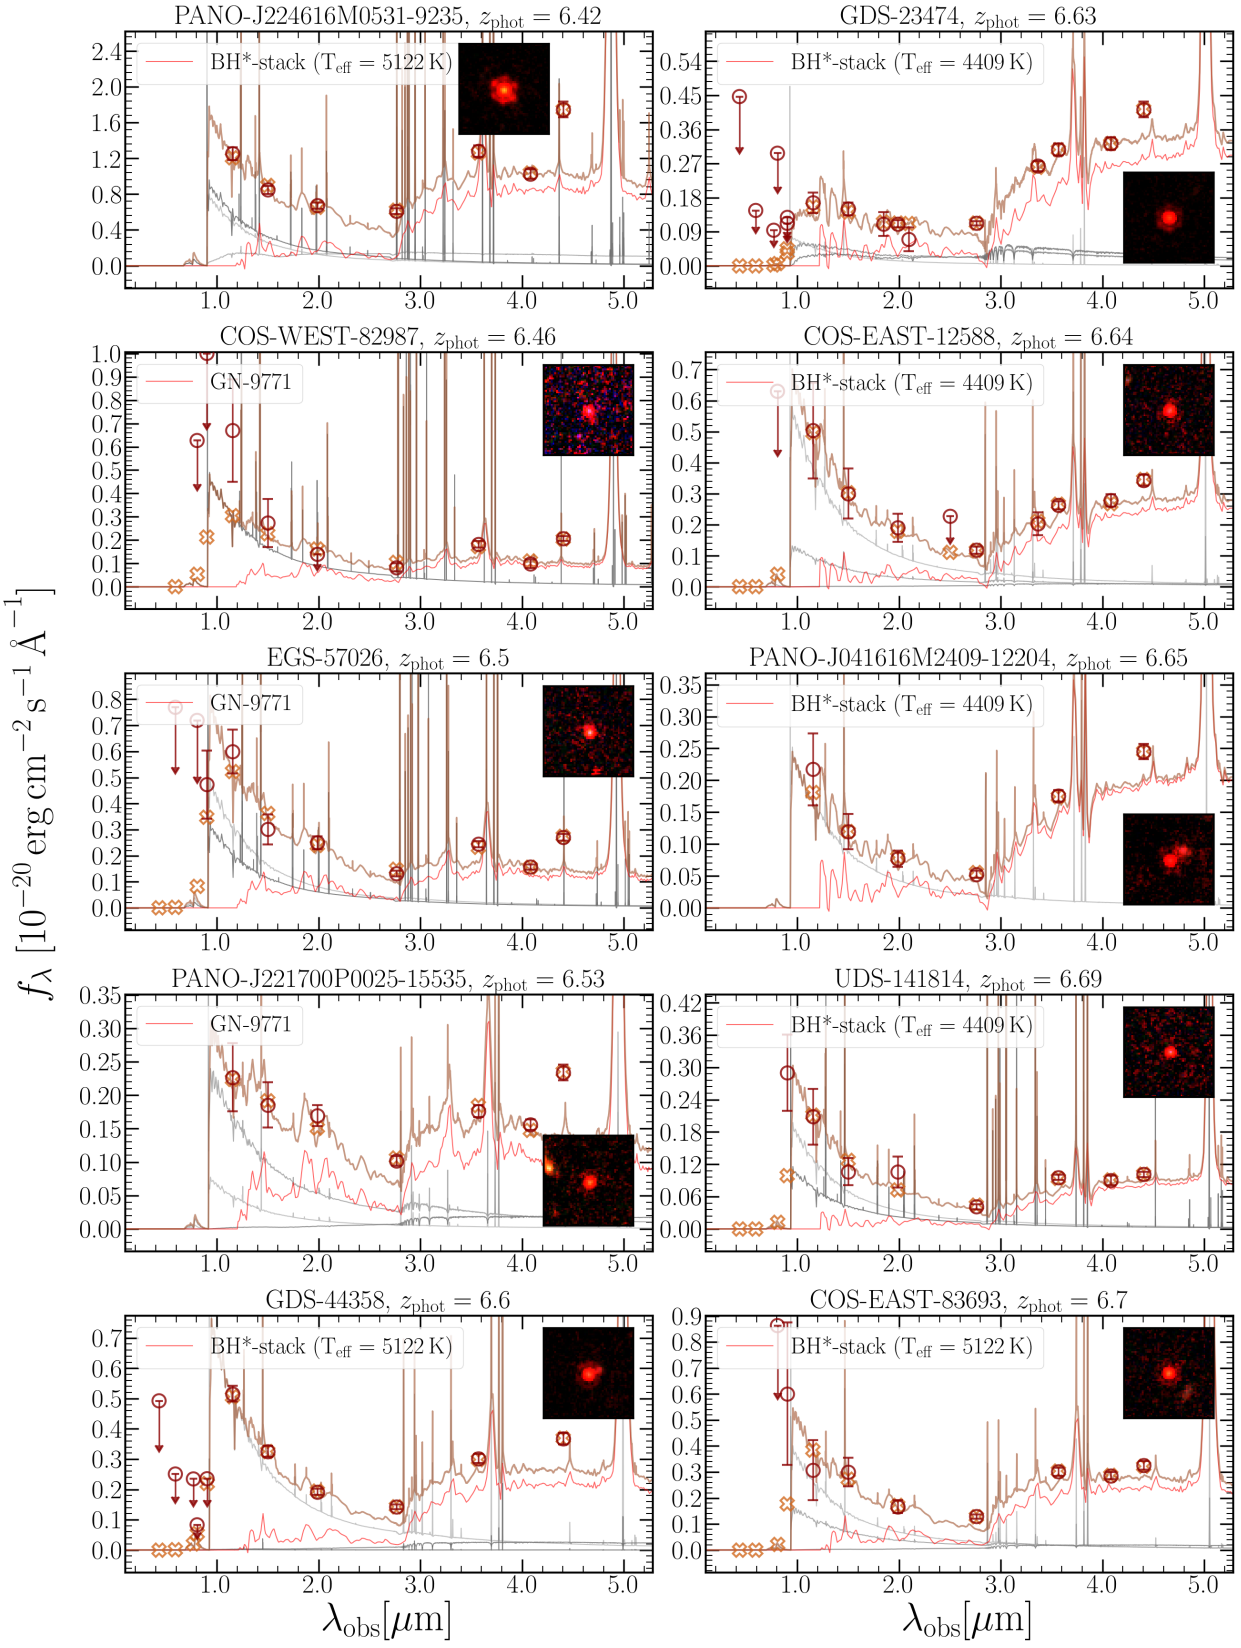

FIG. 25.— Same as Figure 15 (continued).

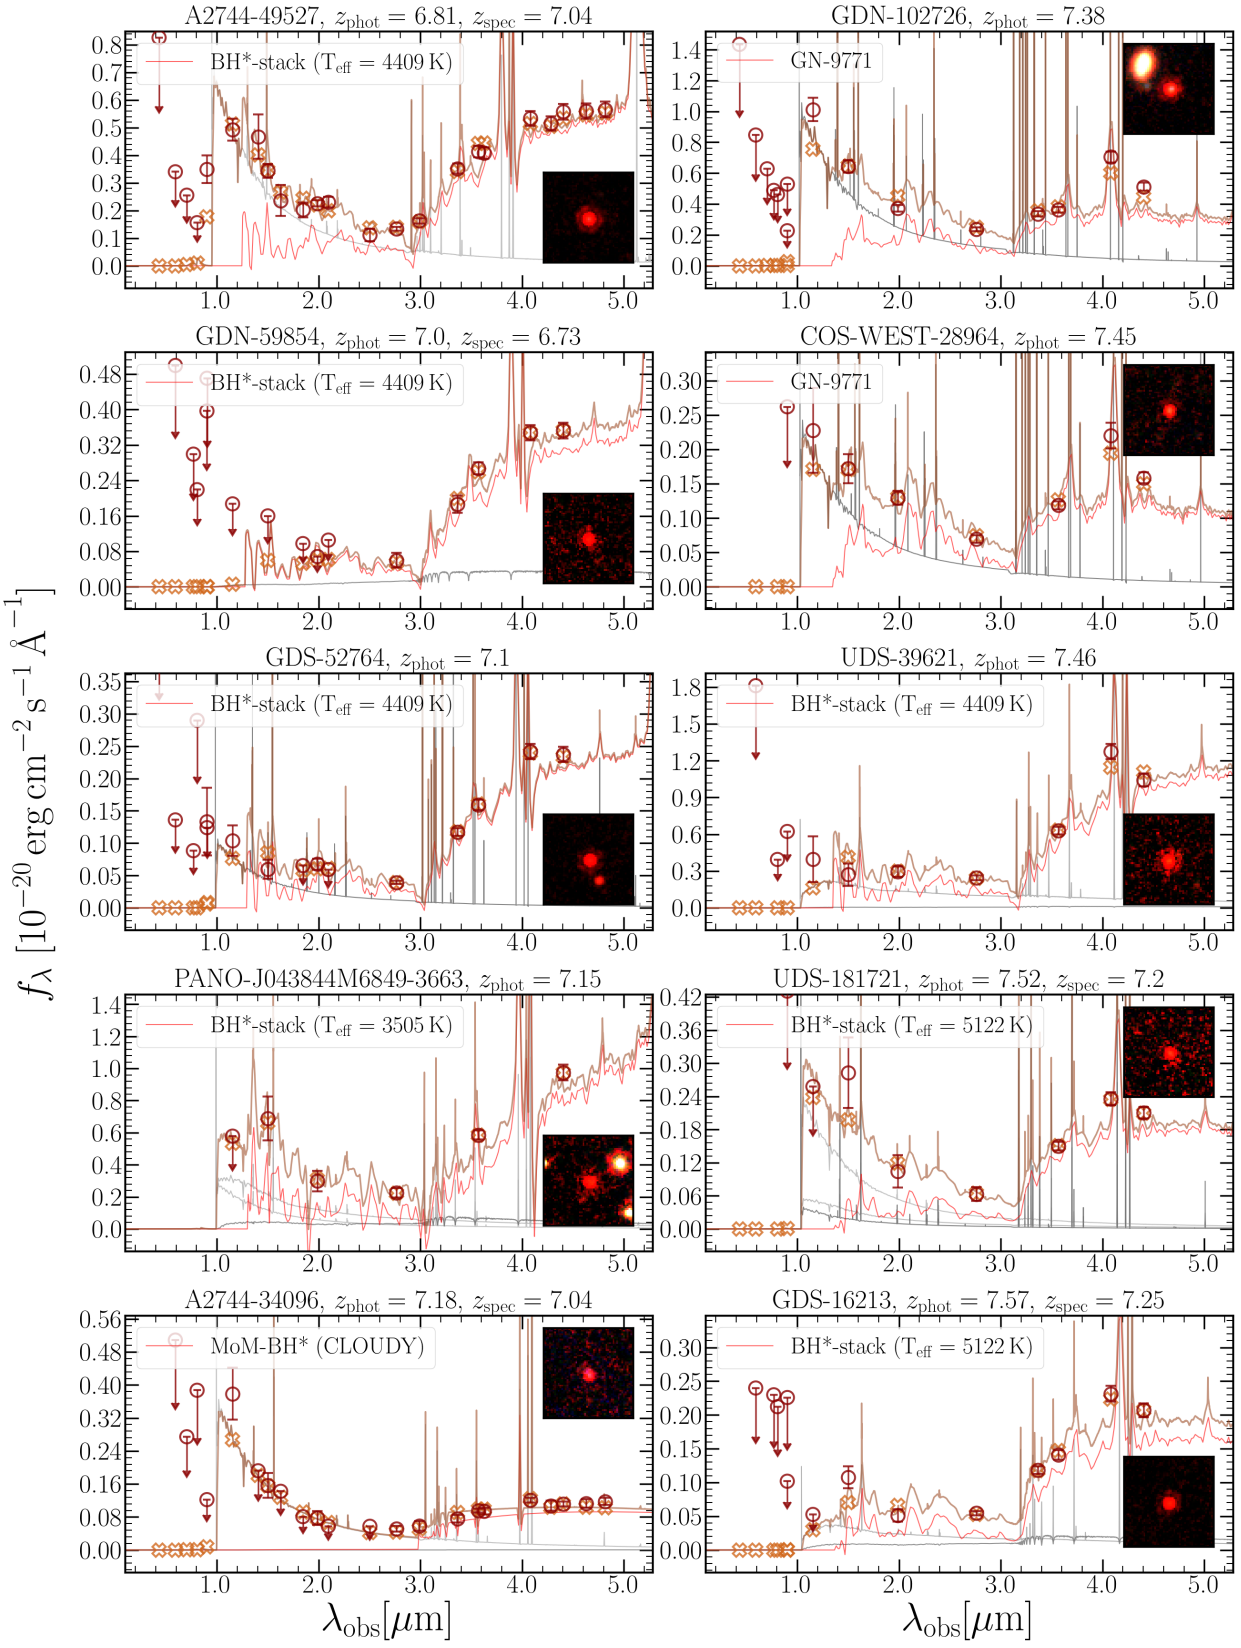

FIG. 26.— Same as Figure 15 (continued).

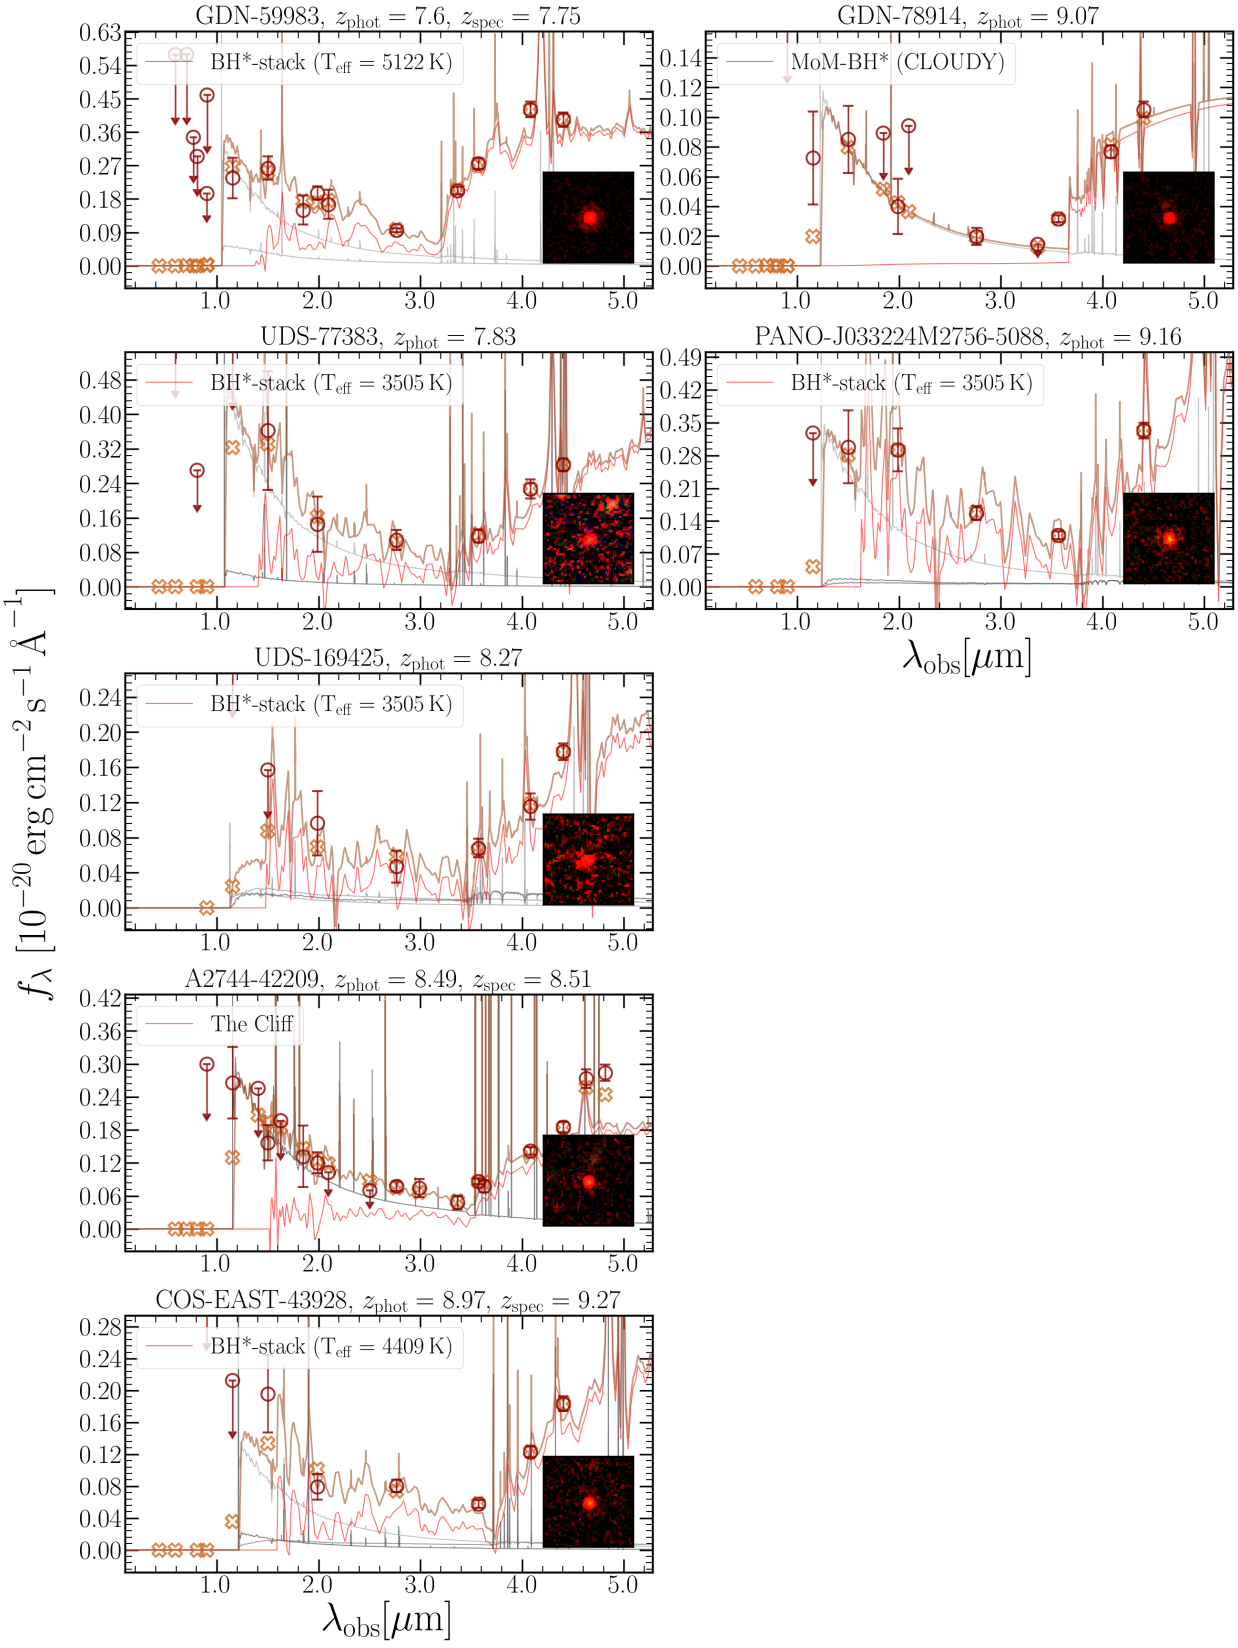

FIG. 27.— Same as Figure 15 (continued).