We also used the NH22 test field (NHTF01) as part of the survey since, by definition, it meets all the above requirements. This brings the total number of COB survey fields to 16. With the exception of NHTF01, we denote the COB science fields with the prefix NCOB.

We selected an additional set of eight fields to perform an improved self-calibration of the relation of FIR intensity to optical DGL. The calibration fields for the FIR–DGL relation are denoted with the prefix DCAL. The DCAL fields were explicitly selected to cover fields with progressively higher 100 μm surface brightness, up to a limit of $\square\, 3\,\mathrm{MJy\,sr^{-1}}$. This limit was selected to avoid dust optical depths large enough that nonlinear behavior between the FIR intensity and scattered light amplitude might start to come into play. All DCAL fields were selected to have SSL levels in a narrow range ($6 < \mathrm{SSL} < 9\,\mathrm{nW\,m^{-2}\,sr^{-1}}$) that was similar to the NCOB science fields.

We also selected eight fields to help verify the SSL estimates. These fields are denoted with the prefix SCAL. The SCAL fields were chosen to be closer to brighter stars than we would otherwise permit for the science observations to test the reliability of the scattered light estimates. All SCAL fields were selected to have DGL levels in a narrow range ($3 < \mathrm{DGL} < 6\,\mathrm{nW\,m^{-2}\,sr^{-1}}$) that is similar to that in the NCOB science fields.

Lastly, we selected four fields at low ecliptic latitude ($|\beta| \leqslant 6^\circ$) solely to verify the lack of significant brightness from interplanetary dust (IPD) at the large heliocentric distance of the NTf.39cest
the

The coordinates, observation dates, and spacecraft helio-

LORRI is an unfiltered

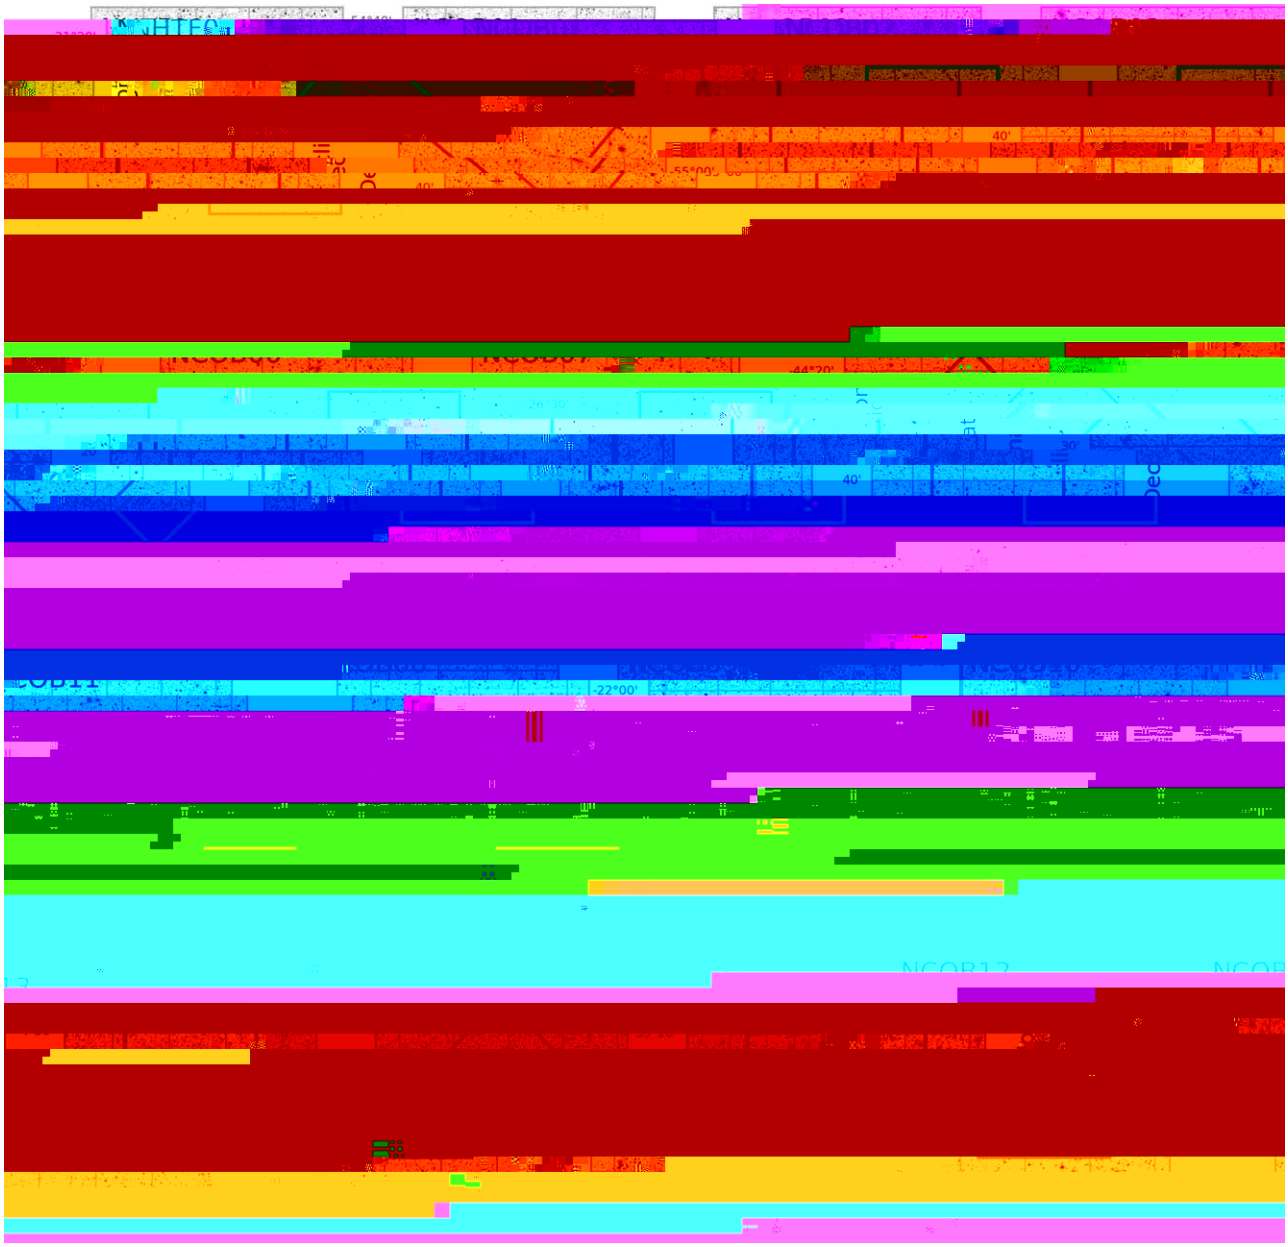

Figure 2. The positions of the NCOB fields are shown with respect to images from the Deep Legacy Survey (Dey et al. 2019)

is the the average of the individual sky levels measured for the 16 or eight images obtained of the field, as opposed to the single sky value measured from a stack of all the images. As noted in Weaver et al. (2020) and NH21, LORRI exhibits a slowly varying pattern of row-wise low-amplitude (<1 DN) streaks in its bias level. This pattern is treated as a random noise source, which is captured in the dispersion of sky values measured for any field.

2.5. B[illegible] D[illegible]

In NH21 we discovered that images taken shortly after LORRI was powered on had elevated background levels, which appeared to decay away during the initial four minutes of operation. The cause of this background effect is unknown. The NH22 test COB images were therefore obtained only after this "cool-down" interval had elapsed following activation of the instrument. As noted in NH22, the sky levels in the 16 images obtained of the test field appeared to be constant over the sequence, validating this solution. We thus used the same procedure to obtain the present data.

Examining the sky levels measured for each image in our present richer data set, however, we have found that the background decay still continues even after the four minute delay, albeit at a low level. Comparison of the amplitude of the decaying background between the NCOB and DCAL exposures, which covers roughly a factor of two in total sky level, shows that the background excess was not tied to the exposure level. It is thus modeled as an additive effect.

All NCOB exposures comprise the same sequence of eight images of 65 s taken in rapid succession, followed by an 80 s pause to adjust the spacecraft pointing, followed by the final

eight images, again taken in rapid succession. Subtracting the mean total sky level from the complete set of 16 images for any field showed that the first images had generally positive residuals compared to the average level over the sequence, with the final images having slightly negative residuals. Figure 3 shows the average residual trend for all the NCOB fields as a function of time.

An exponential decay model appears to be an excellent description of the behavior of the residuals with time. We fitted the trace as

$$\Delta_{\rm sky}(t) = ae^{-t/} + b, \quad (1)$$

where $\Delta_{\rm sky}(\ )$ is the average sky residual at any position on the NCOB exposure sequence, is the time since the start of the sequence, and is a constant background, which accounts for the fact that the initial mean sky for any field will include the background excess. For the NCOB sequence, a least-squares fit recovers $= 0.315$ DN, $\tau = 295$ s, and $= -0.075$ DN, which corresponds to $1.72\,\mathrm{nW\,m^{-2}\,sr^{-1}}$ in intensity units. The background term in essence is the correction needed for the initial total sky levels to account for the presence of the decaying background. In practice, we use the model to apply a correction to each image in the sequence, which removes a source of variance in determination of the fi

DCAL images in the figure. These fields have only eight images, and the shorter duration of the exposure sequence following the start of the sequence means that they are affected more strongly by the decaying background. The exponential model also provides the correction for the DCAL sky residuals, which are in excellent agreement with the residuals from the NCOB images taken at the same time lag. Use of the decay correction does add a systematic uncertainty of 0.16 $\mathrm{nW\,m^{-2}\,sr^{-1}}$

± 0.01 nW m^{-2} sr^{-1} and is derived from the photometric errors given in the DLS catalogs. The systematic error for the bright IGL, typically ± 0.07 nW m^{-2} sr^{-1}, is derived from the

intensity reflects the 10% scatter in the LORRI scattering function, and thus is systematic over all fields. We used the Gaia ESA Archive to retrieve all of the above star catalogs for each field. The SSL values for each field are given in Table 4.

The scattered galaxy light (SGL) term is the analogous scattered light contributed by bright galaxies outside the LORRI field. As with the SSL calculation, the contribution to the SGL is calculated out to an off-axis angle of 45°. Because no uniform all-sky galaxy catalog yet exists to perform this calculation using the actual positions and fluxes of known galaxies, we estimate the SGL as described in NH21. Briefly, we use the galaxy number counts from well-calibrated surveys to compute the mean surface brightness of galaxies with < 20 and then compute the contribution to each annular bin extending out to a radius of 45°. The flux contributions in each bin are convolved with the LORRI scattering function and are then summed up to provide the final SGL estimate. The same SGL value is adopted for all fields. The surface density of bright galaxies is so low that this intensity, 0.10 ± 0.01 nW m^{-2} sr^{-1}, is almost negligible. Hence, even using actual galaxy positions and brightnesses would not make a substantial difference in the final results. As with the SSL, the uncertainty in SGL is taken to be 10%.

As noted in Section 2.1, we observed eight "SCAL" calibration fields to test the SSL corrections. Figure 4 shows the predicted SSL corrections for the NCOB and SCAL fields as compared to the inferred SSL corrections estimated by subtracting all other intensity components from the total observed sky intensity. In detail, this means subtracting the fi

obtained by observing fields over a range of 100 μm intensity in order to derive an empirical relation between the DGL level and the FIR indicator intensity. This approach was later attempted by Symons et al. (2023) in their independent measurement of the COB intensity from NH archival LORRI observations.

As discussed in Section 2.1, we used 100 μm intensity to select both the NCOB science fields and the DCAL fields, and had planned to derive an empirical DGL estimator based on the average 100 μm intensity in any given field. In the initial phases of this work, however, we discovered an error in our previous DGL estimates. As detailed in NH21, we concluded that the IRIS 100 μm map included an amount of residual ZL, and thus derived a correction to the input 100 μm intensity as a function of ecliptic latitude of the fields (see Figure 16 of NH21). We now understand that the rise in intensity with decreasing (absolute) ecliptic latitude is due an extended zone of dust emission at high Galactic latitude that overlaps the ecliptic plane, which we simply missed in our earlier analysis. We thus no longer apply a residual ZL correction to our FIR input intensities, although we do still retain a $|\beta| < 15^\circ$ exclusion zone for COB fields, given concern with potential systematic effects in the ZL corrections. We also note that this error caused us to underestimate the amplitude of the DGL foreground in NHTF01 analyzed in NH22; we rework the decomposition of that field as part of the present analysis.

In the course of redeveloping a 100 μm-based DGL estimator,

to as the generalized needlet internal linear combination (GNILC) method, to separate the structure of the anisotropies from that of the thermal dust emission. These are treated as a field-dependent intensity correction to the overall CIB monopole. Specifically, the CIB-subtracted FIR intensity for a given field is

$$I_c(\ ) = I(\ ) - \mathrm{CIB}_{\mathrm{GNILC}}(\ ) - \mathrm{CIB\_MONOPOLE}, \quad (6)$$

where $I(\lambda)$ and $\mathrm{CIB}_{\mathrm{GNILC}}(\lambda)$ are the averages over the LORRI field from the HFI and GNLIC CIB maps hosted at IRSA, and

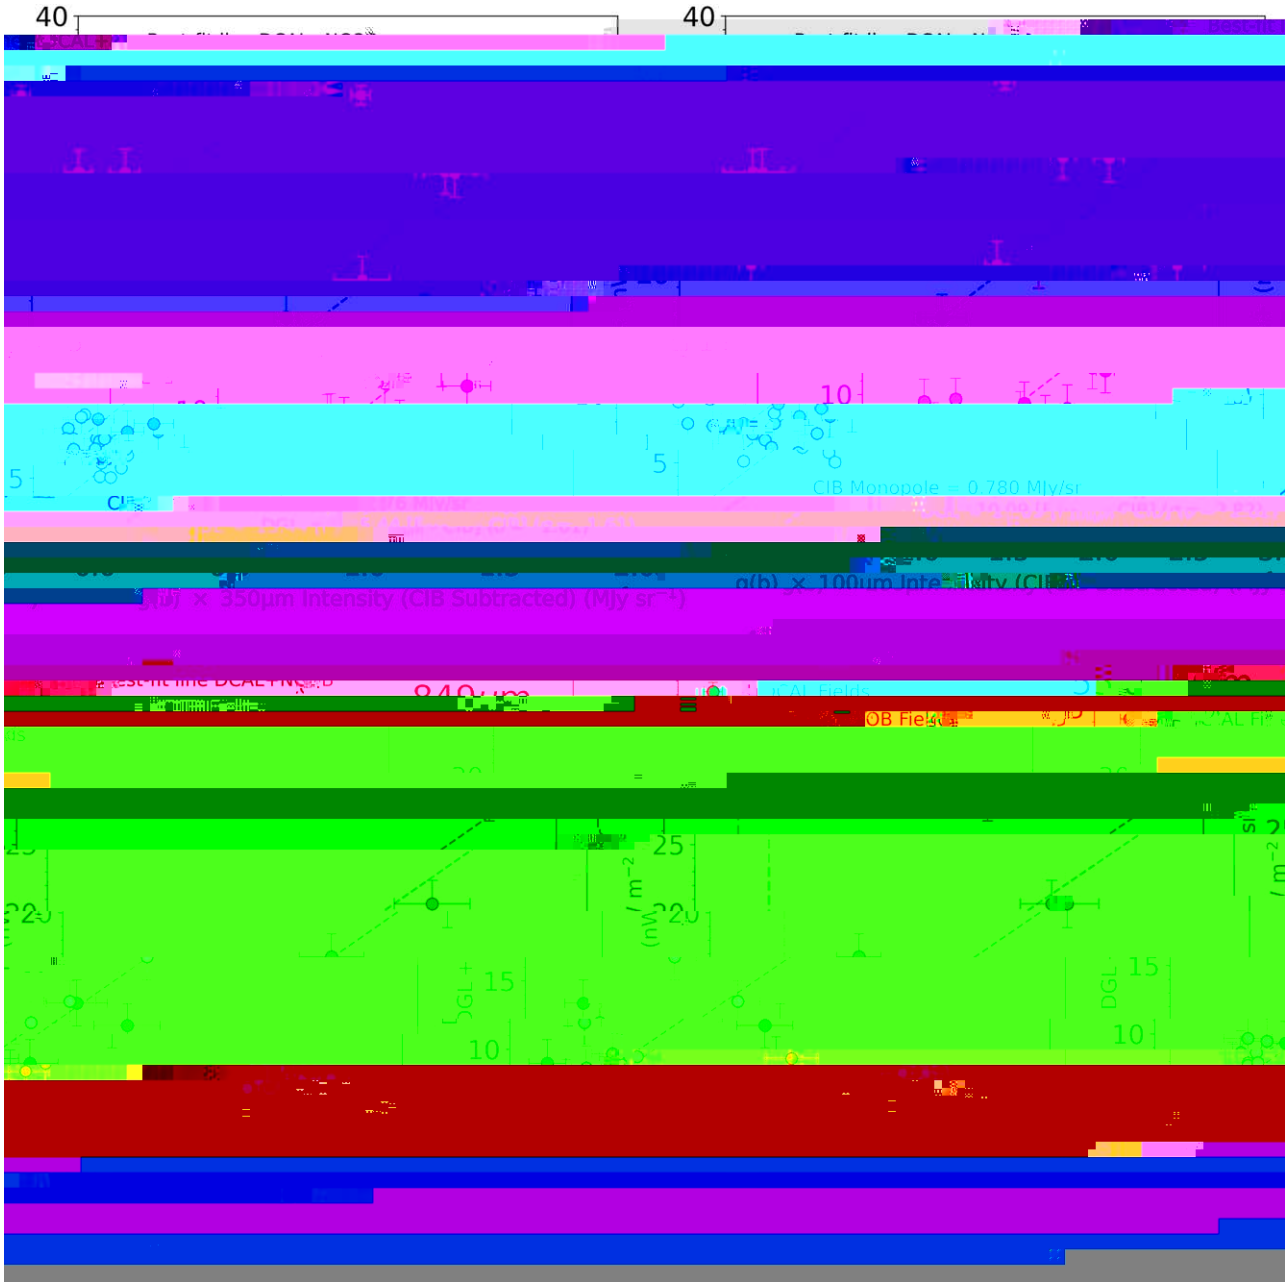

Figure 7. The relationships between DGL+ intensity in the NCOB and DCAL fields as a function of IRIS 100 µm, Planck HFI 350 µm, 550 µm, or 849 µm intensity averaged over the LORRI field. The CIB background intensity has been subtracted from the input FIR intensities to isolate the FIR emission from dust alone. ( ) is a function that accounts for the phase-angle d8.2(VarD(One)4447TJ-3.53351.65560ne)446007P339.JF71Tf7.4(or)3: strengr9ntensity

4.3. • DGL E , ▼ .

Figure 7 shows the four single-band DGL estimators based on 100, 350, 550, and 849 µm intensities, derived by fitting Equation (4) to the NCOB and DCAL DGL+ values (Equation (3)). The linear fit parameters are given in Table 5. The slope of the 100 µm estimator is nearly the same as that in the Zemcov et al. (2017) theoretical estimator but clearly has much smaller errors. The 350 µm and 550 µm estimators have even tighter trends, with the scatter in the 550 µm estimator nearly a factor of two smaller than that in the 100 µm estimator. In contrast, the 849 µm estimator offers the poorest performance of the four bands tested.

One possible reason for larger scatter in the estimators at 100 µm versus 550 µm is field-to-field variation in the dust temperature. The average dust temperature for the NCOB and DCAL fields is 20.1 K, with a dispersion of 1.9 K. Variation in dust temperature can cause variations in the strength of the FIR intensity emitted for the same surface density of dust and the same optical light scattered. At 20 K, 100 µm falls on the short-wavelength side of the peak of the blackbody spectrum, and thus 100 µm emission is more sensitive to small temperature changes than is the intensity at 550 µm, which falls on the long-wavelength side of the peak. We thus investigated two-band DGL estimators, finding that using the 350 µm and 550 µm

intensities in combination gives the best performance, returning smaller scatter than the 550 μm single-band estimator.

In detail, the two-band estimator fits the DGL+ value for any field as

$$\mathrm{DGL}+ = c_1 + g(b)\left[c_2 I_c(550\ \ \mathrm{m}) + c_3\left(\frac{I_c(350\ \ \mathrm{m})}{I_c(550\ \ \mathrm{m})} - \left\langle \frac{I_c(350\ \ \mathrm{m})}{I_c(550\ \ \mathrm{m})} \right\rangle\right)\right], \quad (7)$$

where the subscripts on the intensities indicate that the CIB intensity has been subtracted from them. The mean 350–550 μm intensity ratio is subtracted from ratio for each field to strongly reduce covariance of the intensity-ratio term with the overall intercept term. The fit of this estimator is shown in Figure 8, with coefficient values $,'_1 = 2.60$, $,'_2 = 48.01$, and $,'_3 = 0.96$, with rms residuals of 1.39 nW m^{-2} sr^{-1}. The mean 350–550 μm intensity ratio for our sample is 3.66.

To use the two-band estimator to predict just the DGL value, we subtract the $,'_1$ coefficient, as the predicted DGL must go to zero when the FIR intensity goes to zero. Specifically, our DGL predictor is

$$\mathrm{DGL(nW\ m^{\ 2}\ sr^{\ 1})} \quad g(b)\left[48.01\ I_c(550\ \ \mathrm{m}) \quad 0.96\left(\frac{I_c(350\ \ \mathrm{m})}{I_c(550\ \ \mathrm{m})} \quad 3.66\right)\right],$$

random and systematic errors. Those individual random and systematic errors are given in Table 6. The errors for DGL are the rms values obtained from the Monte Carlo analysis described in Section 5.

To visualize the relative importance of all sky components we represent the results as a stacked bar chart for each field in Figure 9. While the summed intensity from every field is less than the total sky level, the dominance of systematic errors in most of the intensity components means the uncertainty in the combined data set will not be diminished by the $1/\sqrt{N}$ factor. We present our approach for determining the proper propagation of errors for the combined suite of data in the next section.

5. The Cosmic Optical Background

5.1. A M C A E COB I

Having identified all the sources of foreground optical emission known to us in the two previous sections, the task is then to recover an optimal estimate of the COB intensity with accurate errors that reflect the appropriate random and systematic uncertainties (see Table 6), as well as any covariances among the parameters used to estimate the COB intensity. We will use a Monte Carlo approach that randomly generates complete realizations of the COB observations as based on our error model. As noted at the start of Section 3, we do this in two steps. With the exception of the DGL component, our knowledge of all other foregrounds, as well as the total sky level itself, comes from mutually independent information. As such, we start with Equation (3) to estimate DGL+, the DGL intensity plus any anomalous intensity component, which is not affected by covariance among the terms that are subtracted off to isolate it. The second step is to then use the FIR background intensity to isolate the DGL component itself. As this uses the observations to develop a self-calibrated DGL estimator, this step does account for covariance between the observational parameters.

For the first step then, the Monte Carlo routine generates 10,000 realizations of the observations of total sky intensity and the non-DGL foreground components (T, IGL, SSL, SGL, ISL, BIGL, and $I(H\alpha)$) for each field. We also generate random realizations of the independent variables (FIR intensities and cosmic IR backgrounds) that will be used to estimate the DGL. In detail, for any observed component or observational input, $D_{\rm obs}(\,)$, in field $\,$ ($1 \leqslant \,\leqslant 24$; 16 NCOB + 8 DCAL fields), the Monte Carlo routine generates a set of simulated values for $1 \leqslant \,\leqslant 10^4$:

$$D_{\rm sim}(i,j) = D_{\rm obs}(j) \; + \; {}_{\rm RAN}(j)G_{\rm D}(i,j) \; + \; {}_{\rm SYS}(j)F_{\rm D}(i), \tag{9}$$

where $G_{\rm D}(\,,\,)$ and $F_{\rm D}(\,)$ are sequences of Gaussian random variables with zero mean and unit variance. Note that $F_{\rm D}(\,)$ scales the systematic error term $\sigma_{\rm SYS}$()()()((((()()()()()()))(())()))()(((()))()(

we note that while Equation (9) is fully general in allowing for both systematic and random errors for any component, in practice the systematic errors strongly dominate the total error budget, and for many components the random error term is negligible.

5.2. E , ▼, - • DGL I ,,

A two-band DGL estimator of the form given by Equation (7) is generated for every Monte Carlo realization from its set of 24 DGL+ values produced, which is used to predict the DGL backgrounds for that particular realization. As is described by Equation (6), the independent variables used by the estimator are the FIR background intensities at 350 μm and 550 μ , with the estimated CIB backgrounds subtracted. Fitting the DGL+ values to the FIR background intensities generates the

We compute $_U$ for each field using the 10,000 realizations of each parameter on the right-hand side of Equation (10). We then generate the cumulative distribution function of $_U$ for each field and find its 68.3% confidence limits. We adopt the 1σ value on $_U$ to be half the difference between high and low 68.3% confidence limits for that specific field. We then generate a summed distribution of all the $_U$ values for all fields, weighting the $_U$ values for each field by the inverse square of that field's corresponding 1σ value. The value of $_U$ for the full sample is then taken to be the mean value of that weighted sum distribution, and the error in the full sample $_U$ is taken as the half width of the 68.3% confidence limit range of that weighted sum distribution. The resulting summed distribution of $_U$ values for all 16 COB survey fields is shown in Figure 10.

To get the optimal full-survey COB intensity, we repeat the above f.316.1(he)3.y34718006.641288.6235249.5622TmF11161

5.5. R [illegible] I [illegible] D F [illegible]

As noted in Section 2.1, we obtained observations of four fields at low ecliptic latitude to verify the assumption that there is no significant ZL emission at heliocentric distances beyond 50 au. We perform this test by computing the COB and $_{U}$ intensity values for the four IPD fields and look for any significant deviations from the main survey results listed in Equation (12). The results are given in Table

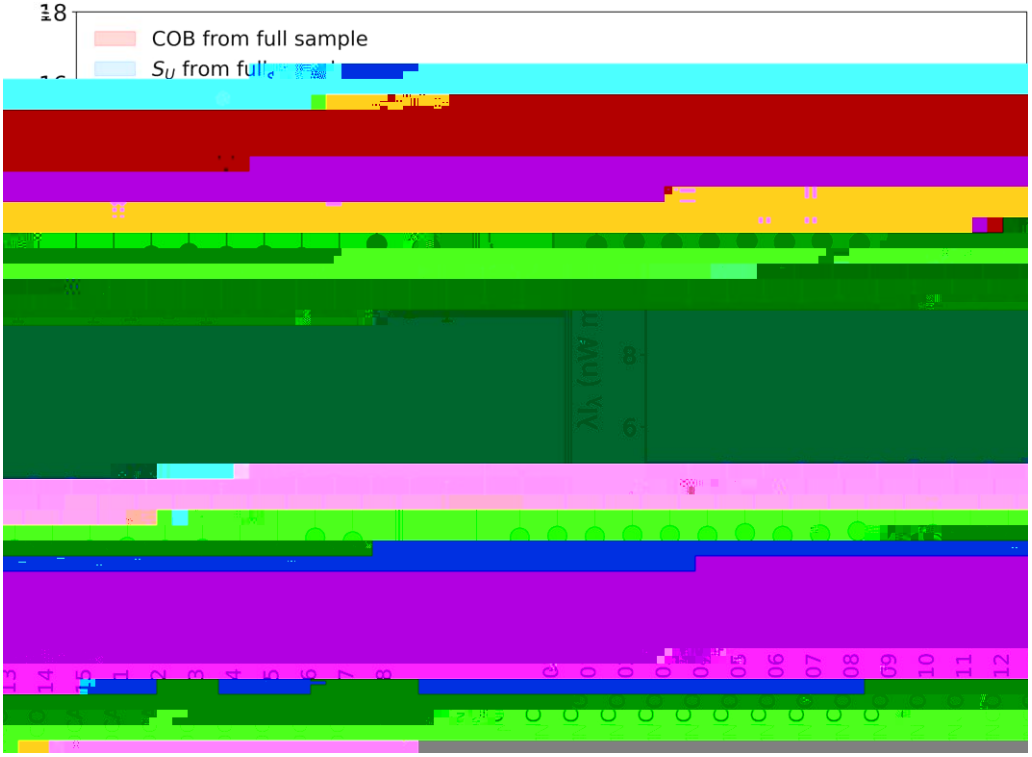

Figure 12. Jackknife test results are shown for the COB and $_{U}$ intensity measurements. The data points show the intensities calculated when the indicated NCOB or DCAL field (shown along the -axis) is excluded from the analysis. The central -values and heights of the red and blue horizontal bands are the derived values and their 1σ ranges for the full-sample COB and $_{U}$, respectively. No field is seen to have an outsized influence on the results using the total sample.

Table 7
COB and $_{U}$ Results for the IPD Fields

Field ID	COB Value	ΔCOB (NCOB – IPDF)	$_{U}$ Value	Δ $_{U}$ (NCOB – IPDF)
IPDF01	8.78 ± 2.59	− 2.38 ± 3.07	0.61 ± 2.88	− 2.38 ± 3.52
IPDF02	11.24 ± 2.79	−0.08 ± 3.24	3.18 ± 3.05	−0.19 ± 3.66
IPDF03	12.10 ± 2.14	−0.94 ± 2.70	2.69 ± 2.47	− 0.30 ± 3.20
IPDF04	11.15 ± 2.02	− 0.01 ± 2.61	3.44 ± 2.36	−0.45 ± 3.11

Note. All values are in units of nW m^{-2} sr^{-1}.

Table 8
COB and $_{U}$ Results for the North and South Galactic Hemispheres

Galactic Hemisphere	COB Value	Stat. Signif.	$_{U}$ Value	Stat. Signif.	Number of Fields
North	11.91 ± 1.34	8.92	4.16 ± 1.96	2.12	3
South	10.93 ± 1.64	6.65	2.80 ± 2.00	1.40	13

Note. COB and $_{U}$ values are in units of nW m^{-2} sr^{-1}.

intensity. A second approach is inferential. The existence of the COB implies that very-high-energy (VHE) γ-rays cannot freely traverse the Universe. Their observed extinction as a function of cosmological distance to their source active galactic nuclei provides an estimate of the COB intensity. The third and final approach is that attempted here: direct observation of the COB intensity. This requires care to isolate and correct for irrelevant foreground intensity sources, but also allows for the discovery of previously unknown intensity sources.

At the outset of this work we posed the question: Is the COB intensity as expected from our census of faint galaxies, or does the Universe contain additional sources of light not yet recognized? With our present result, it appears that these diverse approaches are converging to a common answer. Galaxies are the greatly dominant and perhaps even complete source of the COB. There does remain some room for interesting qualifications and adjustments to this picture, but in broad outline it is the simplest explanation for what we see.

Figure 14 shows our present result in the context of COB measurements from all three methods. We presented a previous version of this figure in NH22, but we revisit it here in light of our revised estimate of the COB intensity and the greatly

reduced allowance for an anomalous COB component. As we noted in NH22, there is excellent agreement on the IGL level over the ensemble of estimates. Driver et al. (2016), Saldana-Lopez et al. (2021), and our own estimate (see NH21) all imply a contribution to the COB intensity of ▫ 8 nW m^{-2} sr^{-1} over the passband sampled by LORRI. At the same time, the galaxy counts feeding into the IGL are obtained from similar, if not the same, observational sources, and thus may have common systematic errors. For example, Conselice et al. (2016) argued that the galaxy counts are seriously incomplete, while Cooray et al. (2012), Zemcov et al. (2014), and Matsumoto & Tsumura (2019) argued that the COB includes a substantial component of light from stars tidally removed from galaxies, or from a population of faint sources in extended halos. Our present COB intensity would indeed allow for a modest enhancement in the implied starlight contribution to the COB, but not a wholesale revision of it. To explain the $_{U}$ value of 2.91 ± 2.03 nW m^{-2} sr^{-1} as extragalactic in origin would require a ▫ 37(±7)% increase in light from galaxies or intergalactic space, corresponding to the ratio of our COB value (11.16 nW m^{-2} sr^{-1}) to that predicted from deep galaxy counts (= IGL + BIGL = 8.17 nW m^{-2} sr^{-1}). Driver et al. (2016) suggest a diffuse component to extragalactic background light could be present at the 20% level, possibly due to low-surface-brightness galaxies and/or intrahalo light in the specific case of the COB spectrum accessible to LORRI. We note that if we extend the integration limit on galaxy counts at the faint end to $= 34$ mag instead of $= 30$ mag and we assume that the faint-end slope of the relation between galaxy number counts and magnitude remains unchanged, our IGL estimate would increase by ▫ 8% with a corresponding reduction in $_{U}$.

We show the COB constraints from five recent VHE (0.1 − 30 TeV) γ-ray studies: Ahnen et al. (2016), H.E.S.S. Collaboration et al. (2017), Fermi-LAT Collaboration et al. (2018), Desai et al. (2019), and Acciari et al. (2019) in Figure 14