STScI | SPACE TELESCOPE SCIENCE INSTITUTE

Instrument Science Report ACS 2023-02

Systematic Effects in ACS/WFC Absolute Gain Measurements

G.S. Anand, N.A. Grogin, J. Anderson, J. Ryon

April 4, 2023

Abstract

The ACS team takes calibration data yearly to measure the absolute gain for the four ACS/WFC amplifiers. Initiated in Cycle 25, this yearly monitoring of the gain allows us to monitor the health of the CCD. Here, we present a reanalysis of the gain with a revised photon transfer curve method, as well as with a new "full-frame" methodology that exploits the use of Poisson statistics on pairs of Tungsten-lamp exposures. Testing for the presence of systematic effects reveals: 1) an increase in the measured gain as a function of increasing row number away from the serial register, and 2) an increase in the measured gain as a function of a increasing flux level of the underlying data. We show evidence that the cause of the former systematic is likely CTE-related effects, and that the cause of the latter may be related to charge-diffusion processes. As the latter has also been tied to the brighter-fatter effect in the literature, we encourage future investigations along this matter for ACS/WFC images. After accounting for the two systematic effects, we find that our measured gains are ∼5% lower than what is currently in the CCDTAB reference file. However, we see that the relative amp-to-amp gains are consistent with the values in CCDTAB to better than ∼0.5%. Hence, while we do not recommend any changes to the relative amp-to-amp gains, the ACS team is considering future adjustments to the absolute gain levels (along with appropriate changes to the photometric zeropoints to compensate).

Copyright ©2023 The Association of Universities for Research in Astronomy, Inc. All Rights Reserved.

1 Introduction

The absolute gain of a CCD provides the conversion between the electrons that accumulate within an individual CCD pixel and the number of digital counts that are read out after the conversion by the analog-to-digital converter (ADC). The only currently supported gain setting for the four ACS/WFC CCD amplifiers is the setting of GAIN = 2.0 e^-/DN. This is because it allows for complete full-well sampling before analog-to-digital saturation, while also aiding in the mitigation of artifacts such as cross-talk between amplifiers (Giavalisco, 2004).

Since any deviation from the "true" gain can be compensated for in the absolute-flux calibration, it is the relative gain values *across amplifiers* that are relevant for obtaining accurate amp-to-amp photometry. These relative gains have been measured down to uncertainties of ∼0.2% using flat-field observations (Bohlin et al., 2009). However, knowledge of the absolute gain is also important, as it is needed to accurately populate the image's error arrays. Additionally, monitoring the absolute gain over time can give indications as to the present health of the ACS/WFC CCD, a key point given the increasing age of the instrument (launched in 2002).

A popular technique for measuring the gain from CCD imaging has been the Photon Transfer Curve (PTC) technique (Janesick et al., 1987). In this method, sub-regions from two identically timed exposures are used to generate a plot of the variance as a linear function of the signal level (i.e., the photon transfer curve). The inverse of the slope of this line is the gain value. This PTC methodology has long been used to determine the gains of the four ACS/WFC amplifiers (Martel et al., 2001). Beginning in Cycle 25, a yearly ACS/WFC gain monitor program was established using F435W flat-field exposures with the onboard Tungsten lamp.

One notable issue with the application of the PTC method to the ACS/WFC images is that there is a noticeable non-linearity present in the photon transfer curve. The origin of this systematic effect has been uncertain, but deserves to be explored in further detail. In this ISR, we perform gain measurements with both the traditional PTC technique, as well as two new methods, one of which exploits the statistics of difference-squared images obtained from pairs of flat-field exposures, and the second which uses the post-flashed short (∼0.5 second) dark exposures taken as part of the CCD daily monitoring program to measure the per-pixel signal and variance to derive the gain. We discuss our findings, which include a systematic effect that inflates the measured gain as a function of increasing row number away from the serial register, as well as the earlier finding of non-linearity in the photon transfer curve method (and an analogous effect for the other methods). We then propose a CTE-related origin as the cause for the new systematic effect, and a charge-diffusion process (the same which underlies the brighter-fatter effect) for the previously known one.

2 Methods

The data we use in this analysis has been taken once per cycle, beginning with Cycle 25. It consists of two consecutive images taken with the Tungsten lamp through F435W filter with identical exposure times, with five different exposure times (0.8, 4.4, 8.8, 16.2, and

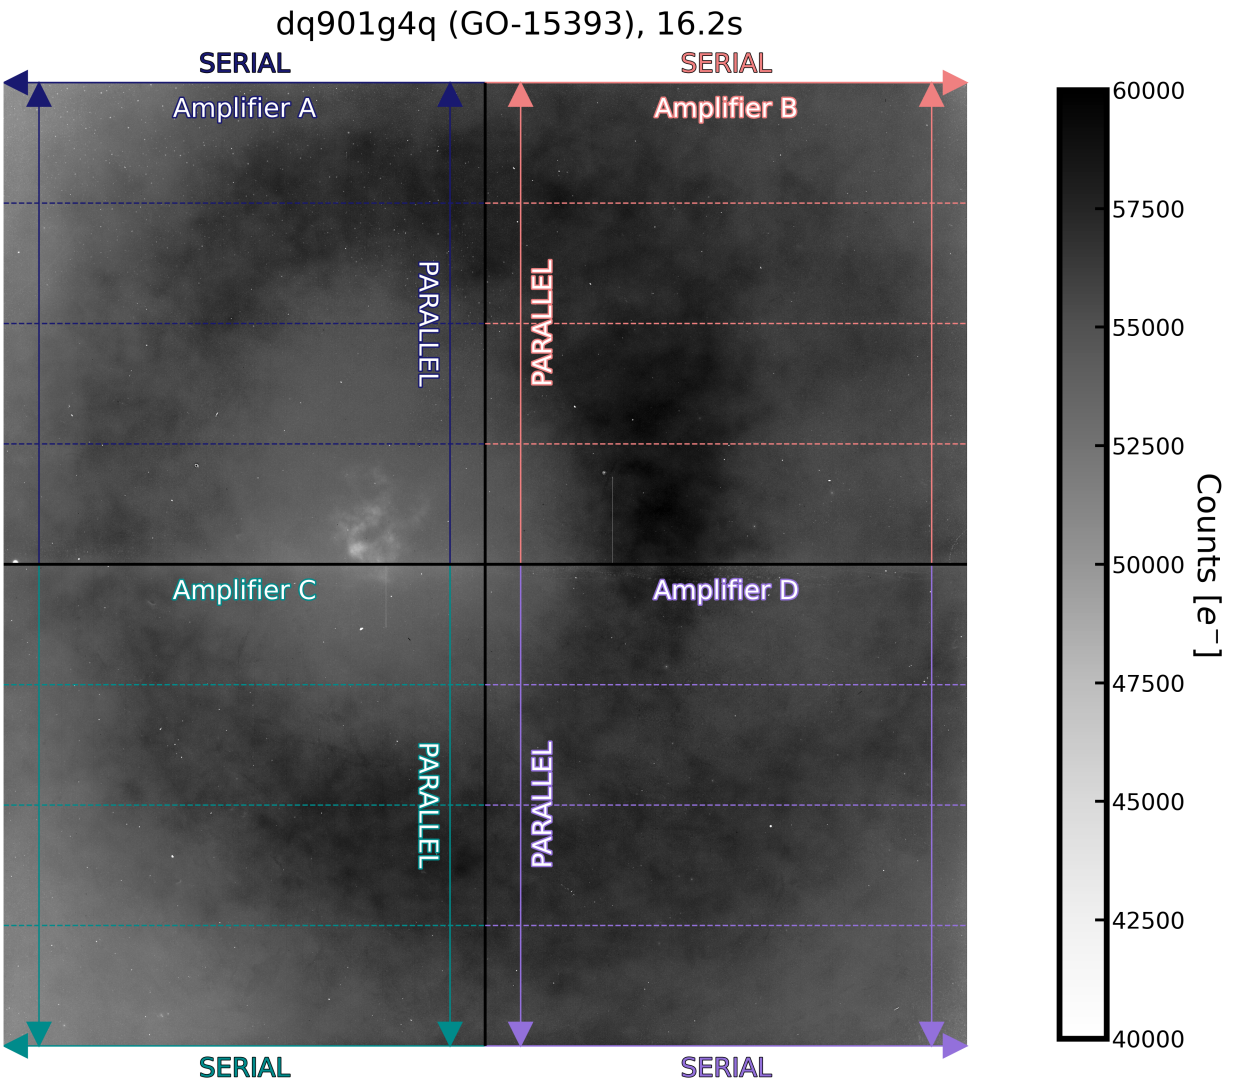

(a)

Figure 1: Example of a 16.2 second F435W Tungsten lamp image from the Cycle 25 ACS/WFC gain monitoring program. The ACS/WFC CCD is shown split into its four distinct amplifier regions, and arrows highlight the directions of parallel readout of the columns, and the serial transfer of charge to the amplifier. Additionally, shown in dashed lines, are the four sub-regions each amplifier is split into in order to test for systematic effects in our gain measurements (see Section 3).

21.4 seconds). To prepare the raw data for analysis, we follow the same steps outlined by Desjardins & Grogin (2018). To briefly summarize, the data is bias-corrected with `CALACS` version 10.3.2 with the `BLEVCORR` and `BIASCORR` steps from `acsccd`. The resulting files (∗blv_tmp.fits) are output in units of electrons.

A careful look at the longest exposure time (21.4s) shows that pixel values are close to saturation over large regions of the detector, and so we exclude images of this exposure time from our analysis (and will be replaced with more intermediate exposures in the calibration program in future cycles). The median flux values of exposures from 0.8 – 16.2 seconds is ∼2,800 – 56,0000 e^-, where the reported values are taken from the peak of the distributions

for each exposure time. An example 16.2 second exposure from Cycle 25 (GO–15393) is shown in Figure 1, with the greyscale highlighting the structure present in this F435W flat-field image. The ACS/WFC is shown separated into four main regions, which read out to its four distinct amplifiers. The parallel readout (along the y-direction) occurs simultaneously for all columns on the CCD. Once the accumulated charge in a given row reaches the serial register, the row is transferred pixel-by-pixel in the serial direction towards each amplifier.

For the purposes of testing systematic effects present in our gain measurements, we further split each amplifier into four distinct regions (shown by the dashed lines in Figure 1), such that each sub-region is a different average readout distance from the amplifier. The reason for this selection will be made clear later, in the results (Section 3).

2.1 Method 1: Photon Transfer Curve Methodology

As mentioned in the introduction, the PTC method has been used previously to measure the absolute gain from ACS/WFC images. We use a methodology similar to the previous ACS/WFC gain pipeline as described in Desjardins & Grogin (2018), except that we increase the box sizes from which the individual data points are calculated from 40×40 pixels to 100×100 pixels to reduce uncertainties, although we note that this change was already in place in an updated version of the pipeline from Desjardins & Grogin (2018). We also note that while Desjardins & Grogin (2018) reports a 5σ clipping threshold for excluding hot pixels and cosmic rays, their pipeline actually uses a 3σ threshold. Again, to be comparable with the results from their pipeline, we retain their 3σ threshold. Lastly, instead of only 25 boxes placed sparingly throughout the ACS/WFC detector, we use contiguous boxes placed over the entirety of ACS/WFC, with the exception of the outer edges of the amplifiers. This change allows us to fit the PTC relation in each of the four sub-regions for each amplifier while still retaining adequate statistics. We also note that the PTC method as described in Desjardins & Grogin (2018) does not include a formal treatment for the read-noise contribution to the variance. Although the effect of this non-treatment is small (as the read-noise is only $\sim 4e^-$), it can play a noticeable role at lower flux levels (i.e. shorter exposure times), where the variance from the read-noise contributes a proportionally higher amount than it does at higher flux levels.

2.2 Method 2: Full-Frame Methodology

Our new method involves the use of Poisson statistics and difference-squared imaging to determine the absolute gain for each amplifier. With two Tungsten lamp exposures of equal exposure times, A and B (in units of DN):

$$(A - B) = 0 \pm \frac{K + [(A + B)g]^{\frac{1}{2}}}{g} \tag{1}$$

where 0 is the baseline difference between the two images, $K = 2RN^2$, RN is the read noise for each amplifier (in units of DN), and g is the gain (in units of e^-/DN). Rearranging further to solve for the gain, we find:

$$g = \left[\frac{(A-B)^2 - 2RN^2}{(A+B)}\right]^{-1} \tag{2}$$

where all quantities on the right hand side are in units of DN. However, we note earlier that we are using blv_tmp files for our analysis as we require data that has been bias-corrected. These files have already been transformed into units of e^- by `CALACS` using the hard-coded gains found in the CCDTAB reference file (g_{CCDTAB}), which have remained static since 2009[1]. We can then multiply the right-hand side of the above equation by $(g_{CCDTAB}/g_{CCDTAB})^2$ and rearrange to arrive at a final equation for the gain:

$$g = \left[\frac{(A-B)^2_{e^-} - 2RN^2_{e^-}}{(A+B)_{e^-}}\right]^{-1} g_{CCDTAB} \tag{3}$$

where the inputs are in units of electrons, to match the form of our data, and Equation 3, not including the factor of g_{CCDTAB}, is a multiplicative offset from the CCDTAB gain.

There are two data issues that need to be dealt with before proceeding with the procedure outlined above. The first is an issue with the fractional decimal values of the pixels in the flat-field images. Due to digitization errors, these fractional values are not randomly distributed but are quantized at discrete values. To correct this issue, we simply drop the fractional value for each pixel and replace it with a random value pulled between .00 and .99. As we are dealing with thousands of counts in even the lowest signal levels in our data, this procedure has a negligible effect on the measurements. The second issue involves the fact that the second exposure in a pair of back-to-back lamp exposures has been noticed to be slightly brighter than the first, although the magnitude of this effect is filter-dependent (typically $\sim$0.3%). The cause of this effect is likely due to the lamp becoming warmer after the first exposure (Y. Cohen, private communication). We deal with this by first finding the peaks of the distributions of each image using Gaussian kernel density estimation, and then scaling the second image so that its peak matches that from the first image. Due to inherent uncertainties in determining the modes of the distributions, we correct for the offset only if we find that it is greater than 0.5%. We note that this only occurs in two cases (both of which are pairs of 0.8 second exposures), and that typically the measured offset is $\sim$0.1%.

With the above issues handled, we then proceed to calculate the right-hand side of Equation 3 and solve for the gain separately for each amplifier. Although the procedure results in a map of multiplicative gain offsets relative to their corresponding CCDTAB gains (the right-hand side of Equation 3 before multiplying by g_{CCDTAB}), we are not able to make use of this map due to the presence of hot pixels and cosmic rays, which artificially inflate the value of the individual pixels they affect. To circumvent this issue, we collapse the 2D images into single 1D arrays, and then sort them by their values. We then remove the top 10% of values and take the mean of the remaining 1D array. This threshold is chosen given that warm and hot pixels effect >2% of pixels on the images, and cosmic rays affecting *either* of the two individual images would skew our distributions. Ordinarily, the mean of this array would provide us with our multiplicative gain offsets (from g_{CCDTAB}), however we must first

[1]See history at `https://hst-crds.stsci.edu/browse/72m18219j_ccd.fits`

divide the mean by 0.623. This value is obtained from the 90^{th} percentile of a chi-squared distribution, which we observe the data to follow. Now, with the resultant value, we simply multiply by the hard-coded gain (g_{CCDTAB}) to retrieve the final gain value for each amplifier, as shown in Equation 3.

3 Results

With our two distinct methodologies in place, we run our procedures on all the ACS/WFC gain monitor data from Cycles 25–29, analyzing each exposure time separately. Our initial results show gains that are consistent between the two methods. In order to test potential systematic issues with the new full-frame method, we split each amplifier into four distinct sub–regions and re-run our analysis (see dashed regions in Figure 1). We find intriguing results, in that we see the measured gain to vary as 1) a function of distance from the serial register, and 2) the flux level of the exposures at which the gain is measured. While we show only figures for amplifier C in the main text (chosen because of a later comparison in Section 3.2.2 only performed for amplifier C), similar figures for all amplifiers are available in Appendix A.

3.1 Dependence on Row Number– CTE-Related Effects?

As mentioned above, when examining results from the full-frame methodology, we find that the measured gain appears to vary as a function of distance from the serial register. We perform the same sub-region splitting as was done with the PTC method, and find very similar results. An example of this effect on the 0.8 second exposures from Cycle 25 can be seen in Figure 2, where the gain measured from each sub-region is plotted against its average readout distance. We see that the effect is persistent with both methods, and appears to have a similar slope with each amplifier.

Figure 3 shows the same procedure for the 16.2 second exposures. Here we see a similar trend, but the slope, which gives the magnitude of the effect, is diminished by a factor of $\sim$3. At this point we also note a discrepancy in the retrieved gain values from these two different flux levels, an effect that we will discuss in the next subsection. We also note that there is a $\sim$2% disagreement between the two methods for at the fainter flux level (as seen in the 0.8 second exposures; Figure 2), whereas the disagreement is $<$1% at the brighter flux levels (as seen in the 16.2 second exposure; Figure 3). As mentioned earlier, this may be due to the fact that the PTC method used does not include a formal consideration for the small, but non-negligible read-noise, which plays a proportionally higher role in the measurements at lower flux levels.

What is the cause of this systematic effect with row number? What we are seeing may be the effect of charge-transfer efficiency (CTE) losses and related effects on the gain measurements. As ACS/WFC spends more time in the harsh radiation environment of low-earth orbit, the efficiency of individual transfers during readout continues to degrade. In general, the degradation of CTE generates a variety of trapping and release effects that result in CTE "blurring", and these effects conspire to result in a net lowering of the measured variance, thus resulting in artificially higher estimates of the gain.

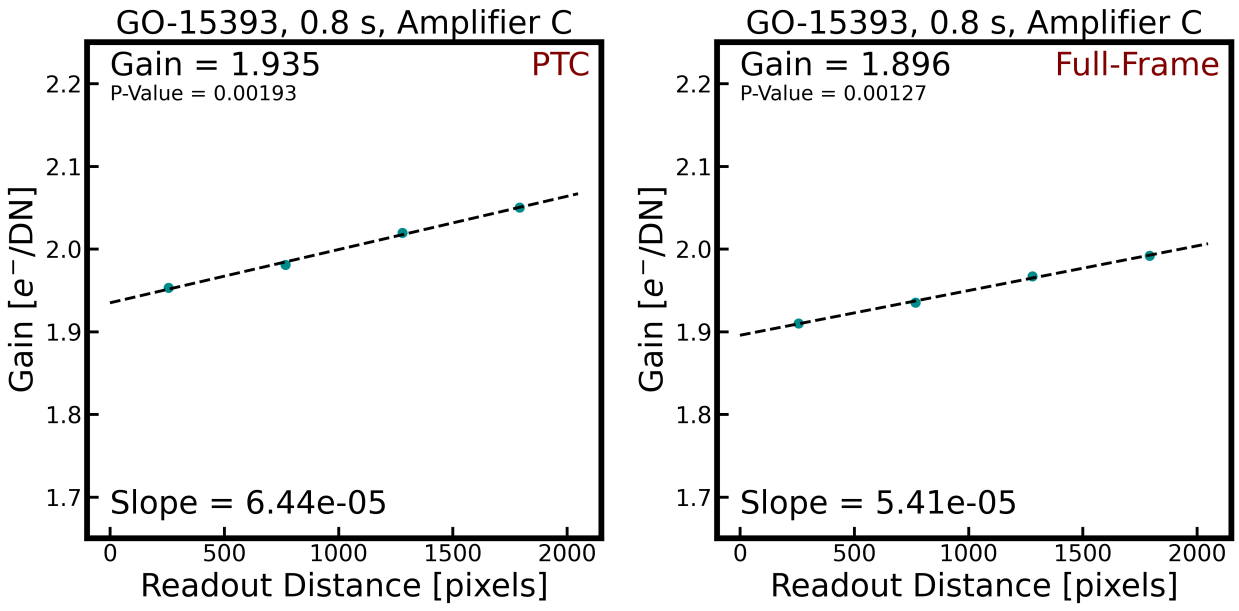

(a) Method #1: Photon-Transfer Curve

(b) Method #2: Full-Frame

Figure 2: Gain results for amplifier C obtained from the sub-regions of each amplifier for the 0.8 second exposures from Cycle 25, from both the a) PTC and b) full-frame methodologies. The typical statistical uncertainties on the measured intercept (gain at the zeroth row) are of order $\sim$0.01 e^-/DN. Results from both methodologies are similar, although the results from the PTC method are likely skewed as it does not include a formal treatment for read noise. There is a notable increase in measured gain with the distance from readout, which is likely due to CTE-related effects. A linear fit to each amplifier allows us to infer the gain at a readout distance of zero pixels, where there are no issues due to CTE-related effects.

There is more evidence to suggest that CTE-related effects are the root cause. We have tested splitting the sub-regions such that they are not horizontal as shown in Figure 1, but instead vertical. In this case, we do not see any variation in the measured gain between the distinct sub-regions. As each tall sub-region is the same average readout distance from its amplifier (1024 pixels), this result is also consistent with the cause of the issue being due to CTE-related effects preferentially acting in the parallel direction. Additionally, we can take a look at the evolution of the derived slope for the gain as a function of readout distance. Figure 4 shows the change in slope as a function of time for the 0.8 second exposures, where the magnitude of the effect is the strongest. With both methods, we see that the slope continues to increase with time. This evolution is consistent with the idea that CTE worsens as the CCD continues to degrade due to radiation damage.

Lastly, in the next subsection, we examine synthetic ACS/WFC images produced with and without the inclusion of the effects of CTE forward model, and find that the inclusion of CTE-related effects produces a very similar effect to what we observe here. It is also worthwhile to note that very recently, the WFC3 team has also found that changes in CTE over time inflate their measured gain values on WFC3/UVIS (Huynh & Kuhn, 2022), showing that this effect, perhaps unsurprisingly, is not limited to just the ACS/WFC CCD.

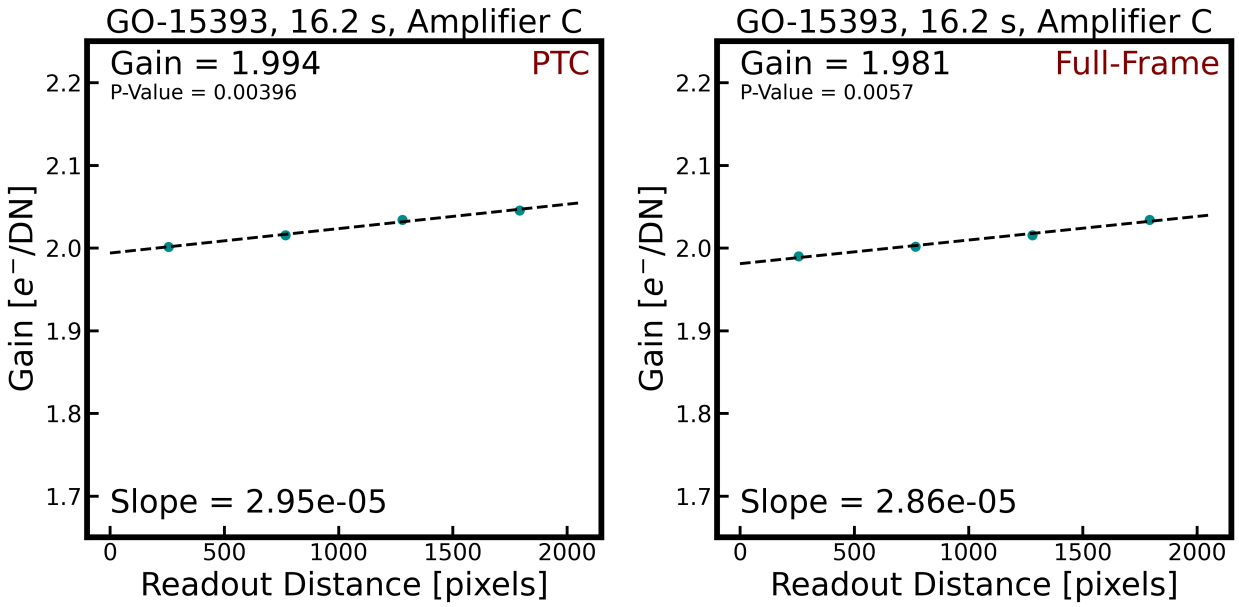

(a) Method #1: Photon-Transfer Curve

(b) Method #2: Full-Frame

Figure 3: Same as Figure 2, but for the 16.2 second exposures. The results from both methodologies are very similar– the difference between the two seen in Figure 2 has diminished, likely due to the reduced effects of read noise on the PTC measurements. The slope of the readout distance dependency is decreased but still measurable. Notably, the gain as measured from this set of data is ∼3-5% higher than that from the 0.8 second exposures.

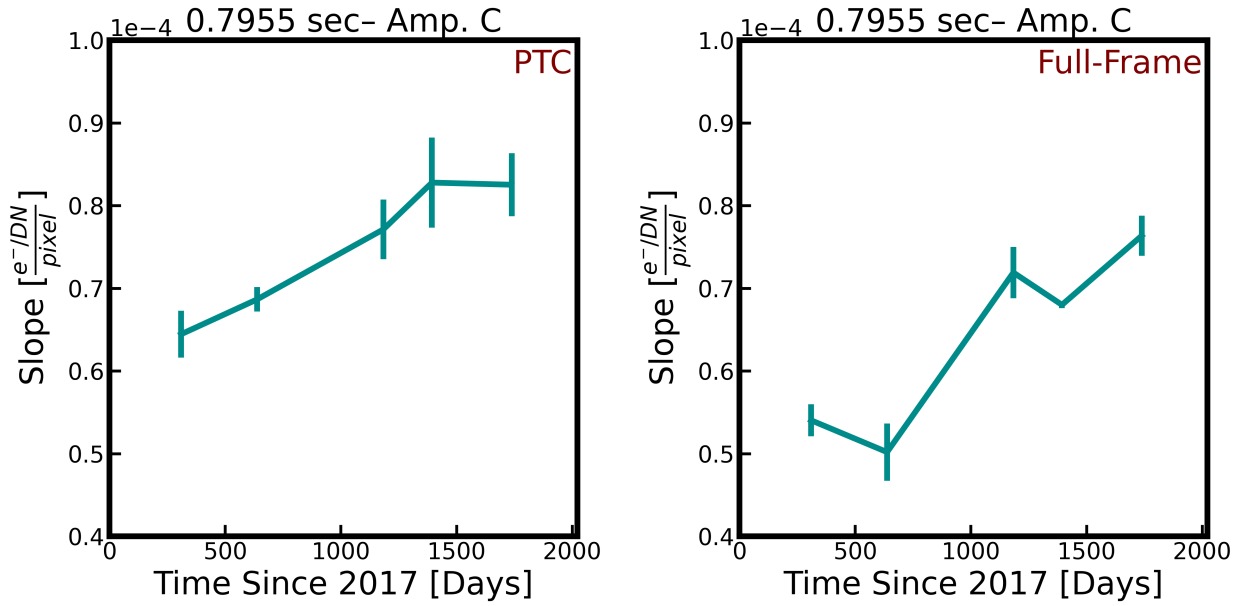

(a) Method #1: Photon-Transfer Curve

(b) Method #2: Full-Frame

Figure 4: The evolution of the slope of the effect of CTE-related effects on the gain as measured from the 0.8 second exposures, using both the a) PTC and b) full-frame methodologies. The increase in slope as a function of time is likely due to further degradation of the ACS/WFC detector resulting from the radiation damage it endures in low earth orbit. Results from both methodologies are similar.

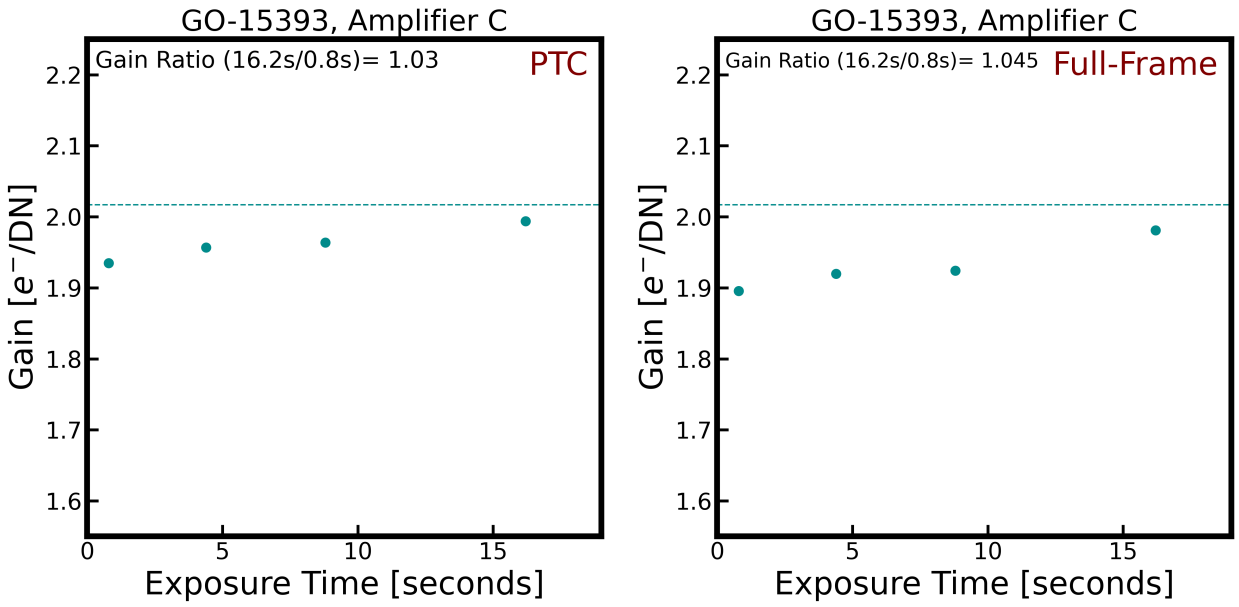

(a) Method #1: Photon-Transfer Curve

(b) Method #2: Full-Frame

Figure 5: Gain values for amplifier C obtained from different exposure times from the Cycle 25 data, analyzed with both the a) PTC and b) full-frame methodologies. Results from both methodologies are similar. The dashed lines indicate the hard-coded gain values, found in the CCDTAB reference file. We find that the gain values measured from the longest exposures are ∼5% higher than as measured from the shortest exposures. We see similar results for the rest of the Cycles in our dataset.

3.2 Dependence on Flux Level: Charge Diffusion Effects?

Given the circumstances, how can we make accurate estimates of the gain? Shown in Figures 2 and 3 are simple linear fits to the gains derived from within each sub-region, as a function of their average distance from readout. In all of these cases, we see that a linear fit is an accurate representation of the data. To retrieve the gain unaffected by CTE, we can now simply read off the gain corresponding to the zeroth row. However, it is at this point that we find that the y-intercepts of the lines also vary as a function of exposure time (and thus flux level), in that the shorter exposures provide a lower estimate of the gain. Figure 5 shows the gain values obtained by extrapolating to a readout distance of zero for the Cycle 25 data, shown with both the PTC and full-frame methodologies. Over the dynamic range of our exposures (∼2,800 – 56,000 e^-), we find that there is a difference in the measured gain of ∼3–5%. We also find that the magnitude of this effect remains consistent over the time range of data that we present in this report (∼5 years).

To investigate the origin of this dependence of the gain measurement on the flux level, we employ a variety of tests. First, we restrict ourselves to columns close to and far from the readout amplifier in each quadrant, to see whether CTE losses in the x-direction might be the cause. We find negligible differences for all the data. Next, we try adjusting the number of sub-regions with both the PTC and full-frame methods. We are able to increase the number of sub-regions for the PTC method to 8 (from the original 4) while still retaining adequate statistics, and find no measurable differences. With the full-frame methodology, we can increase the number of sub-regions to 256 and again find no differences. We can also try

to increase the percentage of pixels that we removed on the bright end with the full-frame method, which also results in no measurable differences. Investigating the linearity of the different exposure times also reveals no hints of signal non-linearity from our shortest to longest exposures.

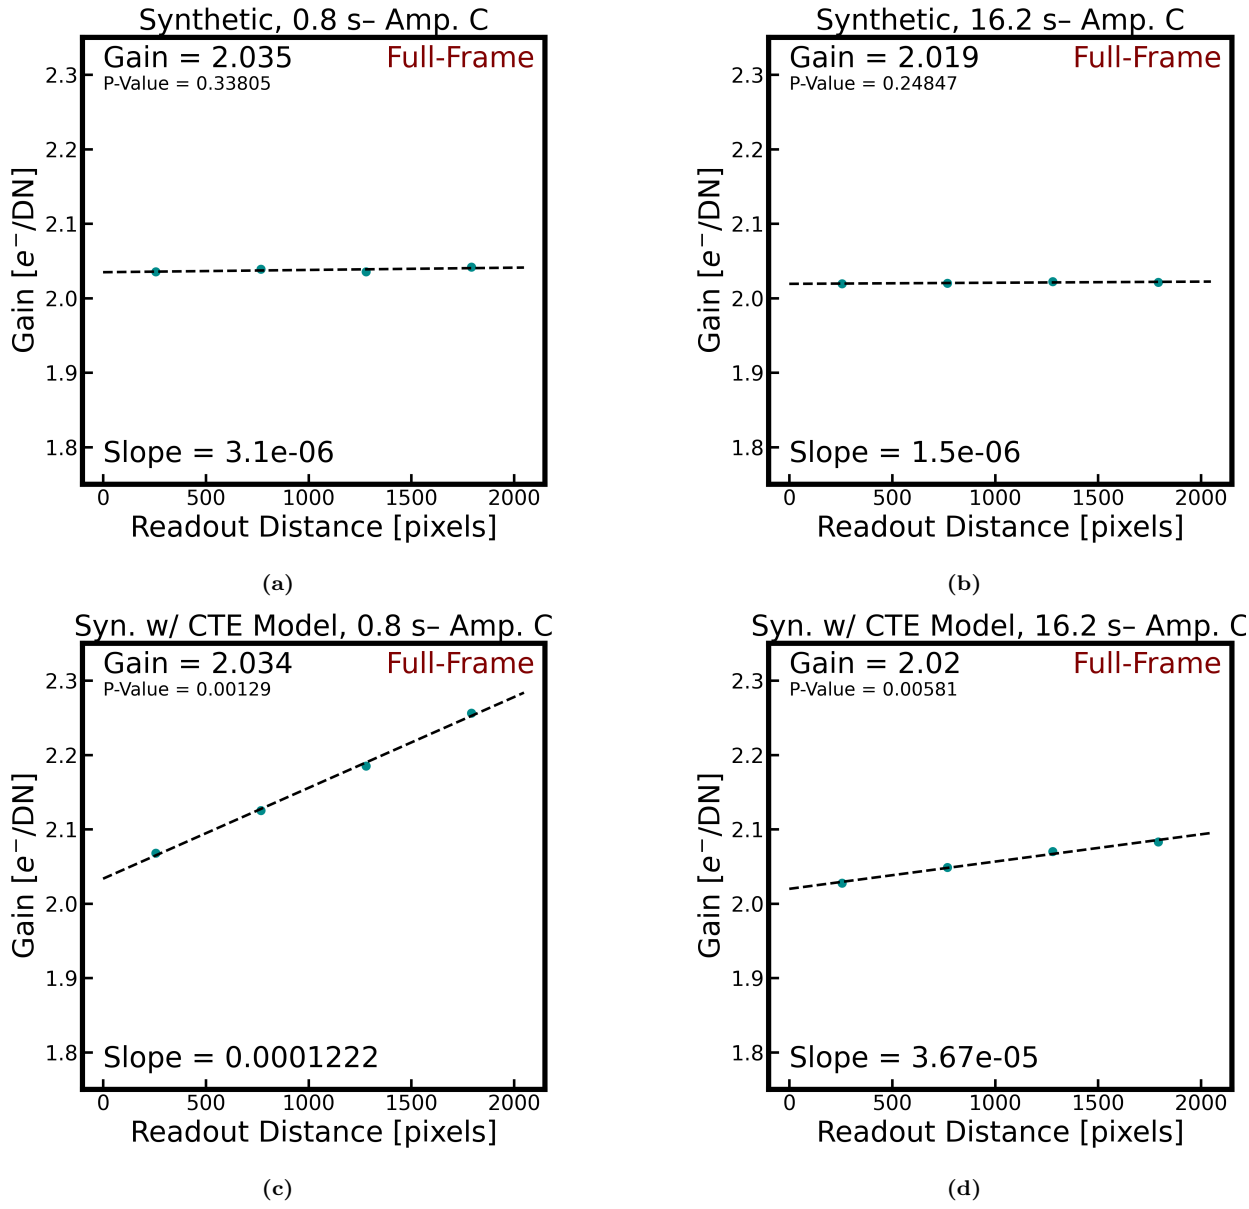

Figure 6: The results of full-frame gain measurements on synthetic ACS/WFC imaging for amplifier C, generated to match the median counts from our gain calibration program exposures. The top half, panels (a) and (b), show the results when no CTE related modelling is included, and the bottom half, panels (c) and (d), show results from the inclusion of CTE-related effects via forward modelling. We see that when no CTE-relayed effects are included, there is no dependence of the measured gain on the distance from readout. When CTE-related effects are included, we see a strong slope, similar to the one observed in the real data. Additionally, this effect is more severe for lower flux levels (modelled to match the median counts from the 0.8s observations) However, we do not recover the ∼4% offset in the final gains measured in the real data over the range of sampled flux levels. While not shown here, the results from the PTC method are similar.

Lastly, we created synthetic images that have measured counts equivalent to the ones in our gain calibration program. We made one set without modelling CTE-related effects, and another set that was run through the latest CTE forward model[2]. If we perform our gain measurement procedures on the set of images that do not include modelled CTE effects, we see no evidence for any slope at all in the gain measurements as a function of readout distance (see the top panel of Figure 6). However, running the same procedures on the set of images that include modelled CTE effects, we find evidence for a strong row-dependence in our gain measurements. Just as in the observed data, the slope of this effect is stronger for lower flux levels (e.g. those modelled to match the median counts from the 0.8s observations). In fact, the magnitude of the row-dependent effect in the forward-modelled CTE images is even greater than in the observed data (see the bottom panel of Figure 6), by a factor of $\sim$1.5–2[3]. While Figure 6 shows results only from the full-frame method, results from the PTC method look very similar.

It is reassuring to see that we are able to reproduce the observed effect of row-dependence on the gain measurement by including CTE-related effects within our synthetic images, but what about the mismatching gain values found at different flux levels in the real data? In the synthetic images with modelled CTE-related effects, we do not see a significant mismatch in the measured gains, unlike the $\sim$3–5% offset seen in the real data. As none of our tests hinted at a possible origin for this offset, we turned to the broader literature to search for similar effects.

3.2.1 Charge Diffusion as the Culprit?

The first reference to differences in gain values measured from varying flux levels can be found in Downing et al. (2006), who noticed an offset in the linearity of the photon transfer curve at higher flux levels. Importantly, while their *signal* was increasing linearly with exposure time, the corresponding *variance* was increasing more slowly. The net effect of this is an artificial increase in the measured gain at higher flux levels, where the variance is no longer increasing to the degree that is predicted by Poisson statistics. The authors propose that the underlying cause is some process that results in the sharing of charge between adjacent pixels, though a precise mechanism was unable to be identified. We can see the manifestation of a similar effect in the ACS/WFC photon transfer curve by examining the residuals of a simple linear fit to the data. Figure 7 shows the residuals from a linear PTC fit to the full range of flux levels considered in this ISR, using only the $\sim$15% of amplifier C that is closest to the serial register (to minimize CTE-related effects). We can clearly see that the residuals are clearly non-linear, as also seen by Downing et al. (2006).

Later, (Guyonnet et al., 2015) explore this non-linearity in the photon transfer curve as well as the "brighter-fatter" effect, which is a broadening of the PSF of stars at higher flux levels. They argue that both of these effects are a result of the same underlying physical mechanism. They propose that the build-up of charge in an individual pixel generates an electric field which begins to repel some incoming flux towards other pixels, in effect decreasing the effective pixel area for pixels with higher fluxes relative to their neighbors.

[2] `https://acstools.readthedocs.io/en/latest/calacs.html#cte-forward-model`

[3] Given the difficulties with modelling CTE losses and other CTE-related effects in a relatively flat background, we would not expect these synthetic results to fully replicate the observed values.

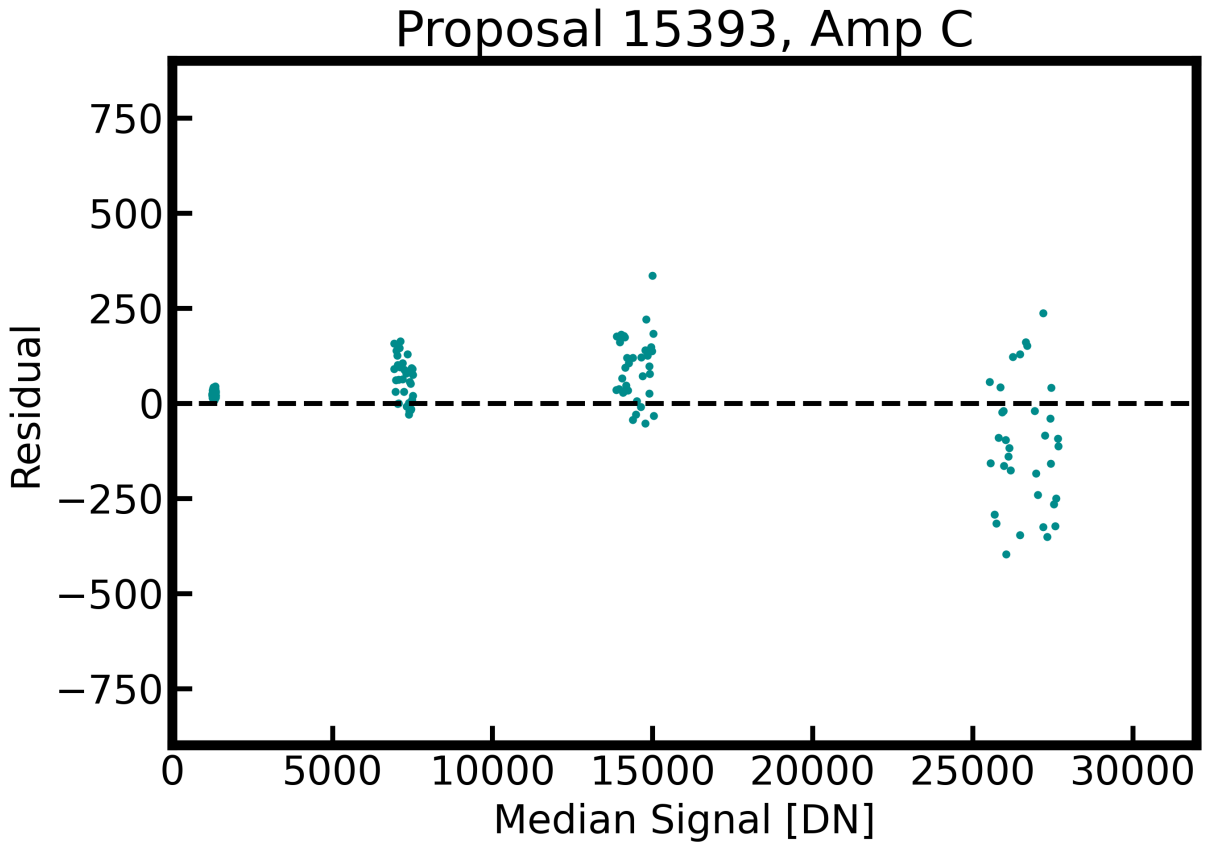

Figure 7: Residuals from a linear fit to the photon transfer curve for the data from Cycle 25. Each of the four clusters of points represents a pair of exposure times, and contain the same number of data points. We see a strong non-linearity in the residuals, similar to what is seen in several results in the literature. These analyses suggest that a charge diffusion process that preferentially bleeds brighter pixels into their slightly fainter neighbors may be the root cause.

This charge diffusion process is similar to the mechanism proposed by Downing et al. (2006). It is also consistent with the work of Coulton et al. (2018), who use data from Hyper Suprime-Cam on the Subaru Telescope to develop a method to correct for this charge diffusion process, resulting in both the return of linearity to the photon transfer curve, as well as a much reduced brighter-fatter effect.

Given the explanations in the literature for this effect, the question arises as to whether there may be a measurable bright-fatter effect in ACS/WFC in addition to the non-linearity of its photon-transfer curve. While such an investigation is outside the scope of this study, we encourage future work on this matter.

3.2.2 Gain Measurements from Short Darks

Given that something like charge diffusion is affecting gain measurements at higher flux levels, how can we be assured that even the lowest flux levels in our gain calibration dataset ($\sim$2,800 e^-) are probing the regime where the photon-transfer curve is not affected by these non-linearities? To provide another test point, we employ the use of the "short darks", i.e.

the 0.5 second dark exposures taken as part of the CCD Daily Monitor calibration program. As the actual dark current in these exposures is entirely negligible ($\sim$0.005 e^-/pixel, see Figure 3 in Anand et al. (2022)), we use only the subset of these darks after the time when post-flashing was introduced in early 2015 (Ogaz et al., 2015). The flash duration in use for these short darks is 4.6 seconds, corresponding to a nominal flash level of $65e^-$, although there is significant spatial variation throughout the image (Ogaz et al., 2014).

We process these short darks as we would images for our other gain measurement procedures, resulting in `blv_tmp` files that are in units of electrons. We translate them back to DN by using the CCDTAB gain conversion values. With $\sim$2300 short darks, the idea now is to measure the average signal and variance *per pixel* over the nearly 8 years of this dataset. Performing this procedure on a per pixel basis allows us to bypass systematic effects that would be introduced by the spatial variations present via post-flashing. Before proceeding further, we must also detrend the images for the temporal evolution of the post-flash LED (Miles & Grogin, 2019). To do this, we calculate a mean illumination level based on the lowest 95 percentile of pixels (to avoid cosmic rays) after masking bad pixels using the DQ arrays. We correct each image by the difference between each image's mean and the overall mean of the images in the timeline. This procedure is performed on a per-amplifier basis. Afterwards, for each detrended image, we first mask the image using its corresponding DQ array, and then again mask pixels above the 95 percentile threshold. From now on, we consider only the 10 rows of pixels closest to the serial register in order to minimize the effects of CTE on the gain measurements (which is likely to be significant at these very low flux levels).

To tabulate the signal and the variance of each pixel, we take the median and variance of all unmasked values for each pixel, separately, throughout the timeline. Following the basic premise that the total variance is the sum of the variances from the read-noise and the signal level [4] (assuming Poisson statistics), we have:

$$\sigma^2_{e^-} = RN^2_{e^-} + S_{e^-} \tag{4}$$

where $\sigma^2_{e^-}$ is the total variance, $RN^2_{e^-}$ is the read-noise contribution to the variance, and S_{e^-} is the contribution to the variance from the measured signal. Using the definition of the gain, we can see that:

$$\sigma^2_{DN} = \frac{\sigma^2_{e^-}}{g^2} \tag{5}$$

which allows us to convert our equation to:

$$\frac{RN^2_{e^-}}{g^2} + \frac{S_{DN}}{g} - \sigma^2_{DN} = 0 \tag{6}$$

where our quantities are now in units which match the values we have in hand (signal in DN, and read noise in electrons). We can now numerically solve this equation for each pixel, while

[4] `https://cxc.cfa.harvard.edu/mta/ASPECT/aca_read_noise/`

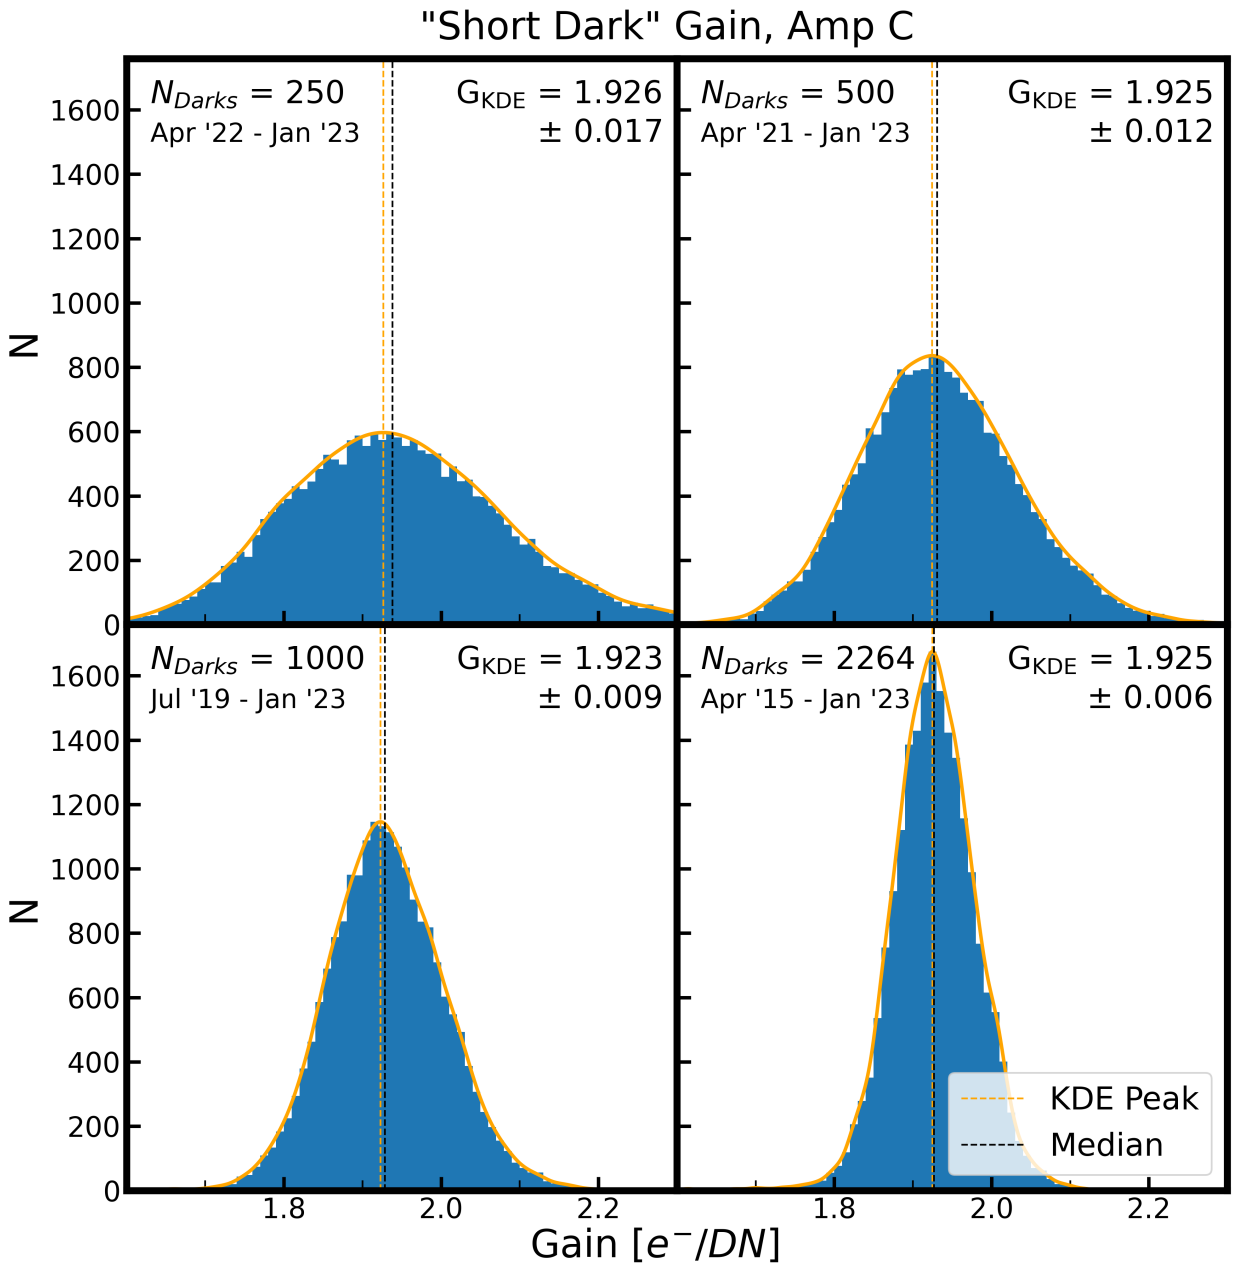

Figure 8: Gain results from our analysis of the short darks for amplifier C. Each data point in the underlying histogram represents the gain derived from the evolution of a single pixel. The final gain value is then derived from the peak of a kernel density estimation (KDE) fit to the underlying, unbinned data. The Silverman bandwidth (Silverman, 1986) used for calculating the KDE is adopted as the nominal statistical error. We find that there is no measurable difference when using the latest 250, 500, 1000, or 2264 (all) post-flashed short darks, although the distribution becomes significantly sharper with the inclusion of more data.

removing the negative (non-physical) solution from the solved pair. In this way, we now have a distribution of gain values, one for each pixel in the 10 rows closest to the serial register. Additionally, as the read noise is known to fluctuate, for now we consider only amplifier C which has had a very constant read noise for the length of time we are considering[5] ($\sim$4.05 $\pm$0.01 e^-). We tabulate an overall gain for the data by finding the peak of the distribution using kernel density estimation. To provide a nominal error from this procedure, we simply adopt the chosen bandwidth for calculating the kernel density estimation (KDE), tabulated using Silverman's rule-of-thumb (Silverman, 1986).

Figure 8 shows the results of this analysis, for four continuously increasing timelines. The first panel has only the most recent 250 short darks. Subsequent panels then show the results from the most recent 500, 1000, and 2264 (all) of the short darks in our data set. With the increasing amount of data, we see a tightening of the distribution, and the KDE derived gain is consistent through each of these partitions of the data set (although the median is skewed towards higher values with the use of less short darks). The gain derived from the overall timeline for amplifier C (1.93 ± 0.01 e^-/DN) is appreciably lower than the current CCDTAB value (2.017 e^-/DN for amplifier C), similar to what is seen with the PTC and full-frame methods after accounting for the gain inflation due to both 1) CTE-related effects and 2) charge diffusion. In detail, it lies right near the value from the PTC method using only the 0.8 second exposures, averaged over the five epochs of data (median = 1.93 e^-/DN), and slightly above the full-frame results from the same set of data (median = 1.90 e^-/DN). We recall that the PTC method as implemented does not include a treatment for the read noise variance (small, but non-negligible compared to the $\sim$3,000 e^- signal present in that data). Additionally, it is possible that there are still CTE-related effects at play at the very low signal levels probed in the post-flashed dark, even when looking only at the 10 rows closest to the serial register. While CTE effects are not well quantified in this regime (very low signals with no astrophysical sources), it is noteworthy that the results from all three of our methods fall below the present CCDTAB values (2.017 e^-/DN for amplifier C) by $\sim$4-6%.

3.3 Time Evolution of the Gain

Overall we find that the gain values appear to be consistent over time when looking at any individual set of exposure times, or separately with the short dark method. Figure 9 shows the gain values obtained from the 0.8 second exposures, with both the PTC and full-frame methods. These plots also show the range (highest–lowest) of gain values obtained for each amplifier over the five cycles of data. Given that the ranges are comparable to the typical nominal error from the individual gain measurements ($\sim$0.01-0.02 e^-/DN), there is no evidence for temporal evolution of the gain in any of the amplifiers. It is worthwhile to note that the PTC method shows less scatter over time. This may be due to the increased sensitivity that the full-frame method has to the small increase in brightness that is exhibited by the Tungsten lamp for the second of two consecutive exposures, despite our attempt at corrections. While showing less statistical scatter, we also note that the PTC method as

[5]See Figure 1 here: `https://hst-docs.stsci.edu/acsdhb/chapter-4-acs-data-processing-considerations/4-1-read-noise-and-a-to-d-conversion`

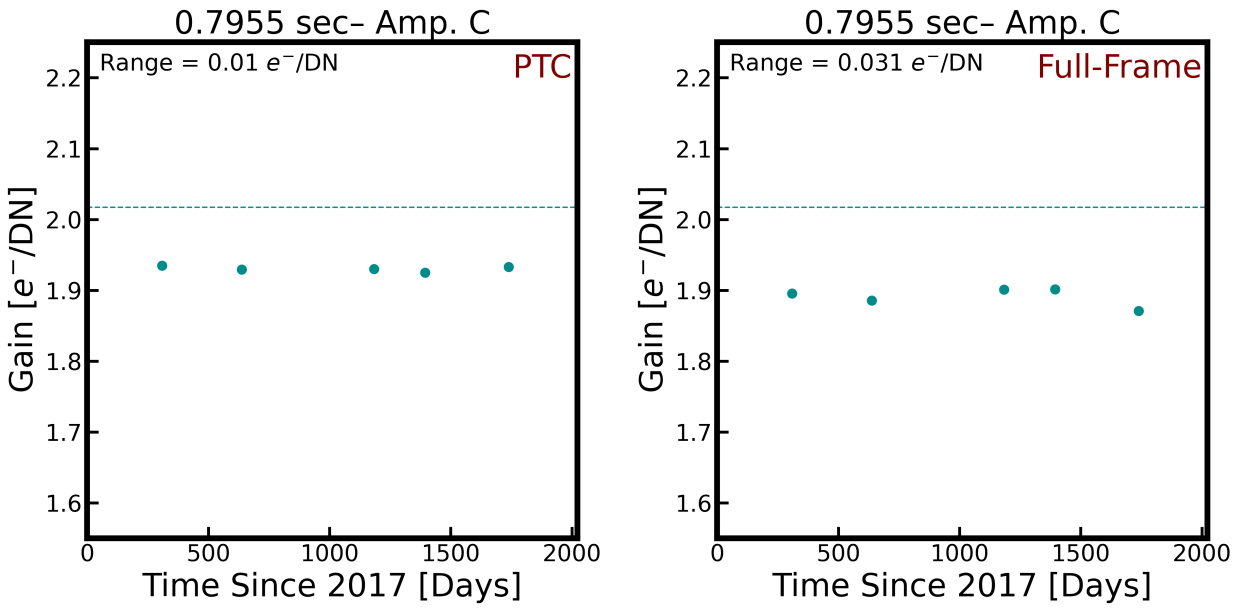

(a) Method #1: Photon-Transfer Curve

(b) Method #2: Full-Frame

Figure 9: The time-evolution of the gain of amplifier C, as determined from the 0.8 second exposures and extrapolated to the zeroth row, with both the PTC and full-frame methods. As the PTC method does not take into account the effects of read noise, the results from that method may be skewed at these lower flux levels. The nominal statistical errors from fitting for the data points are not plotted due to the small size of the error bars, which are typically $\sim$0.01 e^-/DN. The dashed lines indicate the hard-coded gain values, found in the CCDTAB reference file. The range of gain values (highest$-$lowest) for each amplifier are shown in the top left. With both methods, there is no strong evidence for temporal evolution of the gain.

Amplifier	g_{PTC}/g_{CCDTAB}	g_{FF}/g_{CCDTAB}
B	0.9636	0.9413

Table 1: The ratio of measured gain values (as determined by the PTC and Full-Frame methods) to their CCDTAB counterparts for Amplifier B, where the ACS/WFC absolute flux calibrations are determined (Bohlin et al., 2020).

described in Desjardins & Grogin (2018) does not formally include the read-noise, which likely skews the results at the lowest flux levels. Similarly, Figures 10 and 11 show the results from the short dark gain method over nine consecutive timelines. There is no evidence for evolution in the gain when analyzed with this method either.

3.4 Gain Offsets

Given our results, what can be said about the absolute values of the gain, as well as the relative gains between amplifiers? Regarding the former, we can compare our measured gain values to the current CCDTAB values for Amplifier B, where the absolute flux calibration for ACS/WFC takes place (Bohlin et al., 2020). To do this comparison, we consider only the data from the 0.8 second exposures, to minimize the effects of charge diffusion inflating the measured gain. We take the median of the five epochs shown in Figure 9, where we again

Amplifier	g_{PTC}/g_{CCDTAB}	g_{FF}/g_{CCDTAB}
A	1.0034	1.0084
B	1.0*	1.0*
C	0.9930	0.9984
D	0.9945	1.0022

Table 2: The ratio of measured gain values (as determined by the PTC and Full-Frame methods) to their CCDTAB counterparts, after normalization by the same ratio for amplifier B (from Table 1). We can see that the relative gains between amplifiers as measured in this study are consistent with the current CCDTAB values, on average, to better than ∼0.5%. *By definition, the values in the row for amplifier B are identically one.

note that these are extrapolated to the zeroth row, in an efffort to minimize CTE-related effects (which also inflate the measured gain). The results are shown in Table 1. We find that with both the PTC and full-frame methods, our measured gains are ∼4% (PTC) – ∼6% (full-frame) *lower* than what is currently in the CCDTAB reference file.

While an offset in the overall, absolute gain values can be corrected for without changing results for amp-to-amp photometry (i.e. with an appropriate change to the overall zeropoint), the same is not true if we were to find a large discrepancy between our relative gain values between amplifiers, versus those provided within the CCDTAB reference file. Table 2 shows the results for the ratios of measured gains, after normalizing/correcting for the offset seen in Amplifier B. Here, we see that the relative amp-to-amp gains are consistent with what is in CCDTAB to better than ∼0.5%, indicating that we *do not* see strong evidence for offsets in these relative amp-to-amp gains.

4 Conclusions

Using a revised photon-transfer curve and a new methodology based on a "full-frame" analysis, we have performed an analysis on five contiguous cycles of ACS/WFC gain calibration data. We find that results from the two methodologies are consistent, and that both show no evidence for evolution of the gain over this time period. Results from a third method which uses the short darks taken as part of the CCD Daily Monitoring program show similar results.

Our measured gain values are lower than those previously reported by Desjardins & Grogin (2018). The cause of this decrease lies mainly within two systematic effects. The first is that we have found that CTE-related issues artificially inflate the measured gain, and the magnitude of this effect increases with distance from the serial register. A second effect is found that appears to inflate the gain measurements taken from data with higher flux levels (longer exposure times). The difference between the gains that we measure from the lowest and highest flux levels (∼2,800–56,000 e^-) is ∼5%. Based on numerous results from the literature that also show a similar non-linearity in the photon transfer curve at higher signal levels, we believe the root cause of this issue is due to a charge-diffusion process that preferentially affects pixels with higher levels of accumulated flux. Numerous authors attribute this process to be the cause for both a non-linearity in the photon transfer curve *and* the brighter-fatter effect, highlighting the need for future investigations of the presence

and/or magnitude of the latter in ACS/WFC data.

We find that with both our PTC and full-frame methods, our measured gains are ∼4% (PTC) – ∼6% (full-frame) *lower* than what is currently in the CCDTAB reference file. However, we see that the relative amp-to-amp gains are consistent with what is in CCDTAB to better than ∼0.5%. Hence, while we do not recommend any changes to the relative amp-to-amp gains, the ACS team is considering future adjustments to the absolute gain levels (along with appropriate changes to the photometric zeropoints to compensate).

Lastly, from Cycle 30 onwards, we have initiated several adjustments to our gain monitoring program. As we find that images from the highest exposure time (21.4s) are running near saturation, we have removed this pair of exposures from the program going forward, as well as the second highest exposure-time pair (as it is redundant with the flat-field monitoring program). We have added another intermediate exposure time, such that our exposures will be in time steps of 1, 2, 4, and 8 seconds, as well as the 16-second exposures from the flat-field program. Finally, in an effort to reduce the issue of increasing lamp brightness with consecutive exposures, we have divided the exposure sets into two identically-ordered orbits, with each orbit consisting of a single Tungsten lamp exposure in each of 1, 2, 4, and 8 seconds.

Acknowledgements

We thank Meaghan McDonald for assistance with the use of the ACS short darks, and Yotam Cohen for details about the ramp-up in brightness in consecutive flat-field images. We also thank Ralph Bohlin, Nimish Hathi, and David Stark for providing comments on an earlier draft of this report.

References

Anand, G., Grogin, N., & Anderson, J. 2022, Revisiting ACS/WFC Sky Backgrounds, Instrument Science Report ACS 2022-1, 30 pages

Bohlin, R. C., Maybhate, A., & Mack, J. 2009, ACS after SM4: Relative Gain Values Among the Four WFC Amplifiers, Instrument Science Report ACS 2009-03, 7 pages

Bohlin, R. C., Ryon, J. E., & Anderson, J. 2020, Update of the Photometric Calibration of the ACS CCD Cameras, Instrument Science Report ACS 2020-8

Coulton, W. R., Armstrong, R., Smith, K. M., Lupton, R. H., & Spergel, D. N. 2018, AJ, 155, 258

Desjardins, T. D., & Grogin, N. A. 2018, Remeasuring the ACS/WFC Absolute Gains, Instrument Science Report ACS 2018-6, 8 pages

Downing, M., Baade, D., Sinclaire, P., Deiries, S., & Christen, F. 2006, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 6276, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. D. A. Dorn & A. D. Holland, 627609

Giavalisco, M. 2004, Cross-Talk in the ACS WFC Detectors. II: Using GAIN=2 to Minimize the Effect, Instrument Science Report ACS 2004-13, 13 pages

Guyonnet, A., Astier, P., Antilogus, P., Regnault, N., & Doherty, P. 2015, AAP, 575, A41

Huynh, K., & Kuhn, B. 2022, WFC3/UVIS Gain Stability Results for Cycles 26 - 29, Instrument Science Report WFC3 2022-8, 12 pages

Janesick, J. R., Klaasen, K. P., & Elliott, T. 1987, Optical Engineering, 26, 972

Martel, A., Hartig, G., & Sirianni, M. 2001

Miles, N. D., & Grogin, N. A. 2019, Temporal Stability of the ACS/WFC OD-800W LED, Instrument Science Report ACS 2019-8

Ogaz, S., Anderson, J., & Golimowski, D. 2015, Post-Flash Calibration Darks for the Advanced Camera for Surveys Wide Field Channel (ACS/WFC), Instrument Science Report ACS 2015-03, 10 pages

Ogaz, S., Chiaberge, M., & Grogin, N. A. 2014, Post-Flash Capabilities of the Advanced Camera for Surveys Wide Field Channel (ACS/WFC), Instrument Science Report ACS 2014-01, 11 pages

Silverman, B. W. 1986, Density estimation for statistics and data analysis

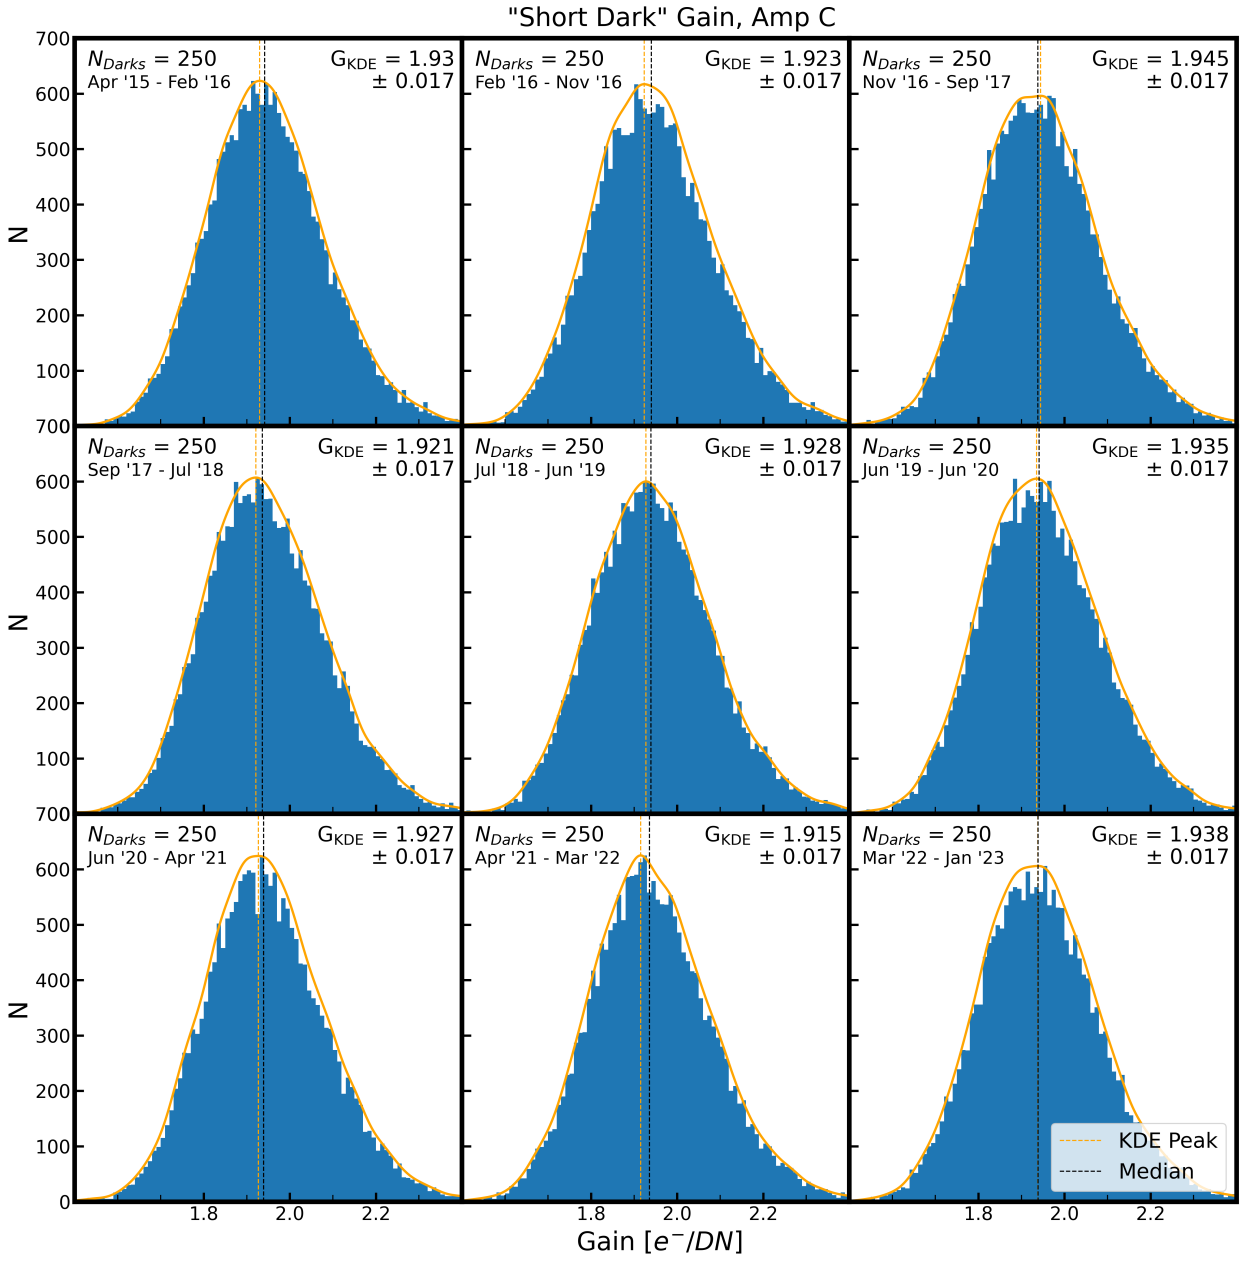

Figure 10: Similar to Figure 8, but now for nine consecutive bins of time. We find that the resulting gains are consistent over time (see Figure 11 as well).

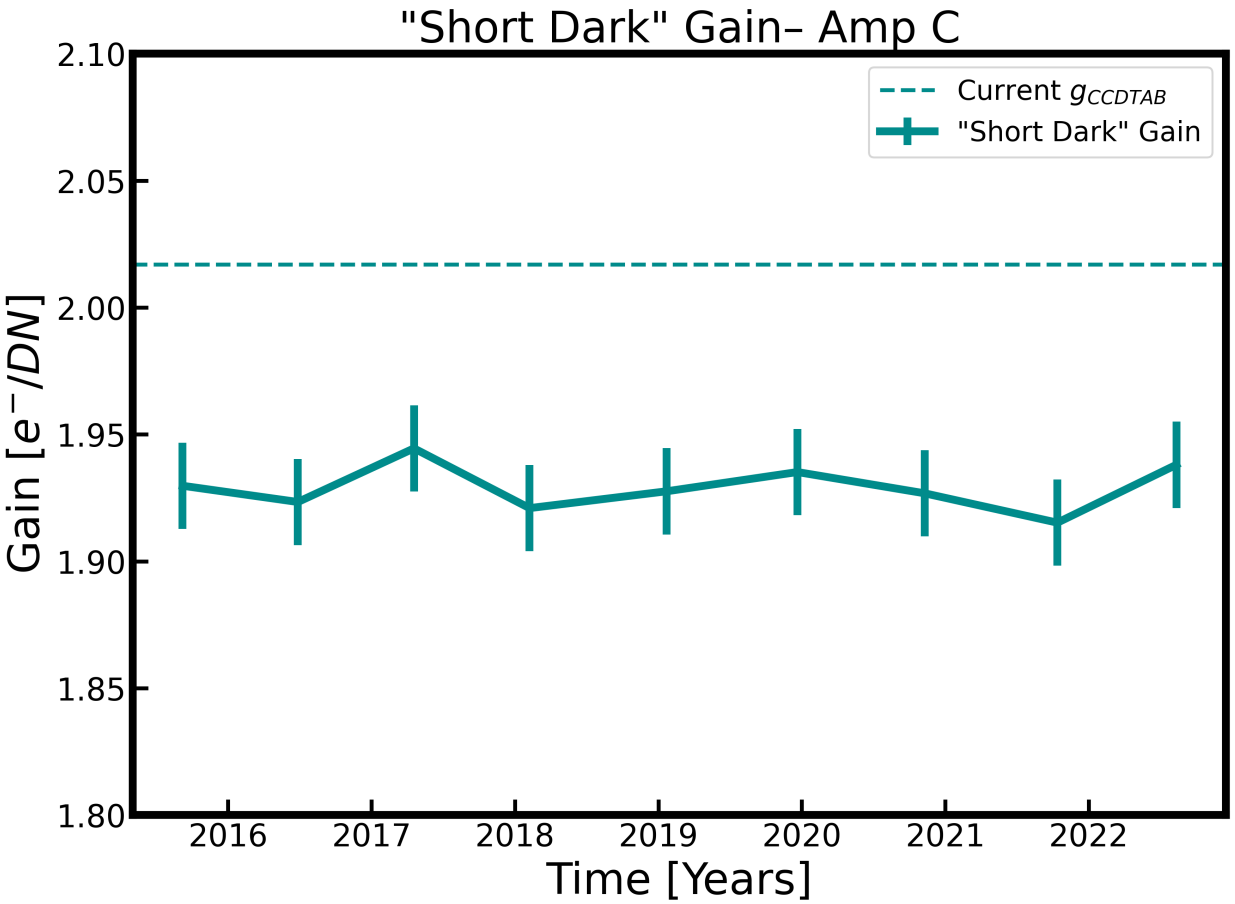

Figure 11: The resulting gains and associated error bars from Figure 10. We see no evidence for the evolution of the gain over time.

Appendix A

Here we show versions of the PTC and full-frame figures, but with data for all four of the ACS/WFC amplifiers.

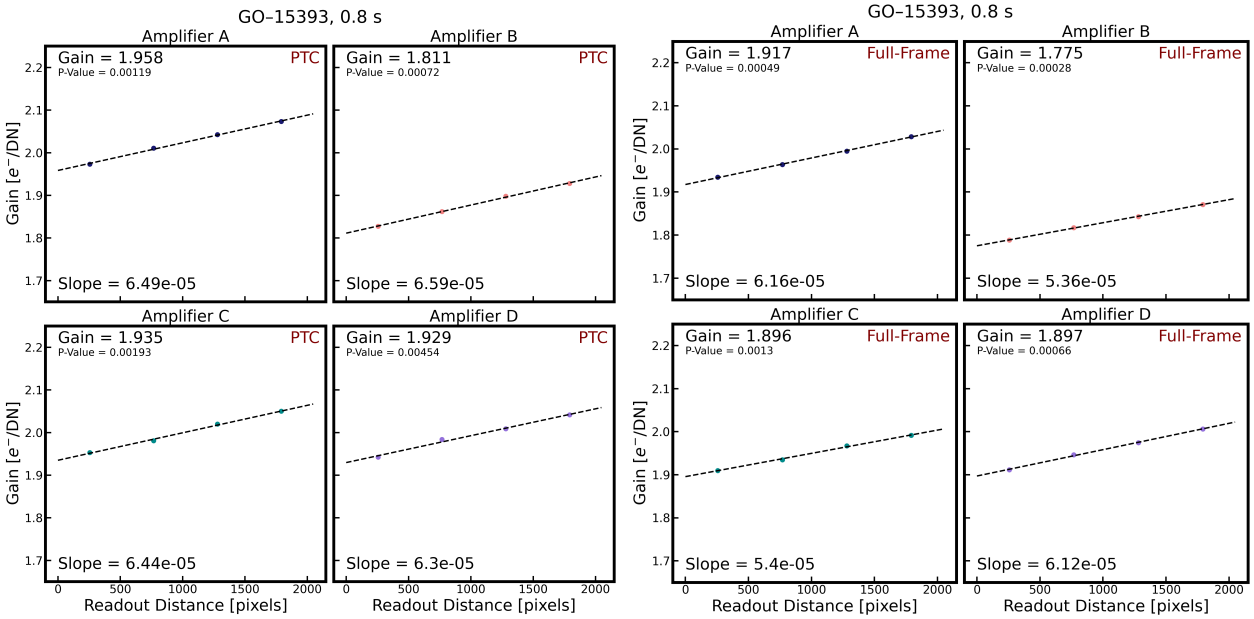

(a) Method #1: Photon-Transfer Curve

(b) Method #2: Full-Frame

Figure 12: Same as Figure 2, but for all four amplifiers.

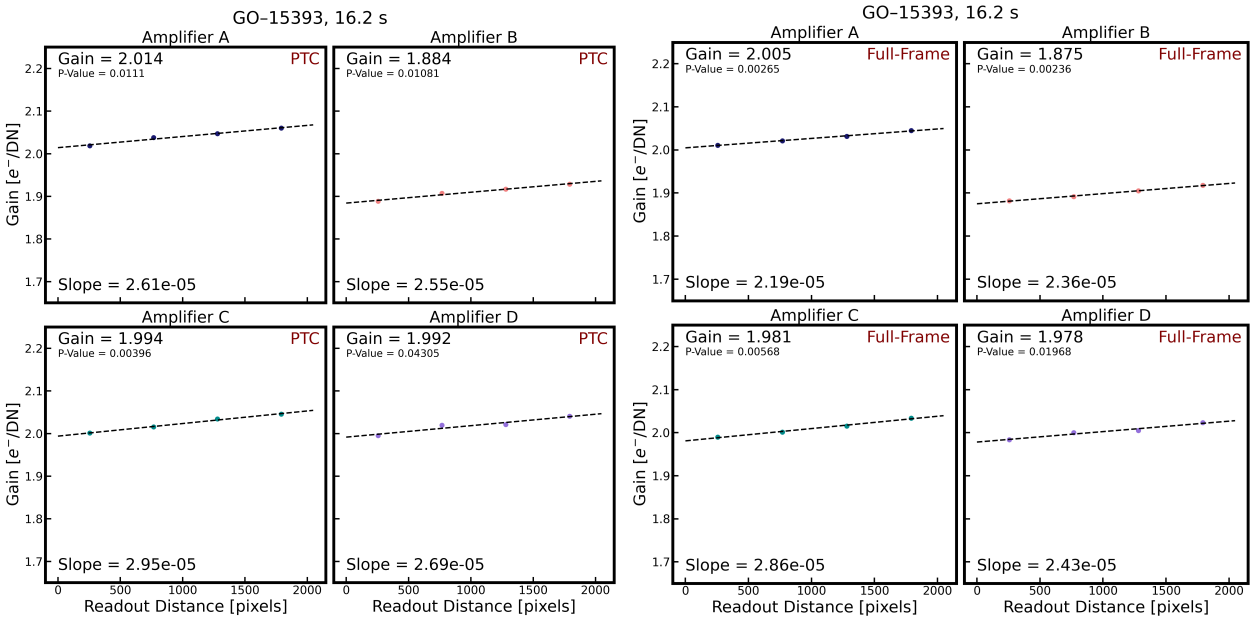

(a) Method #1: Photon-Transfer Curve **(b)** Method #2: Full-Frame

Figure 13: Same as Figure 3, but for all four amplifiers.

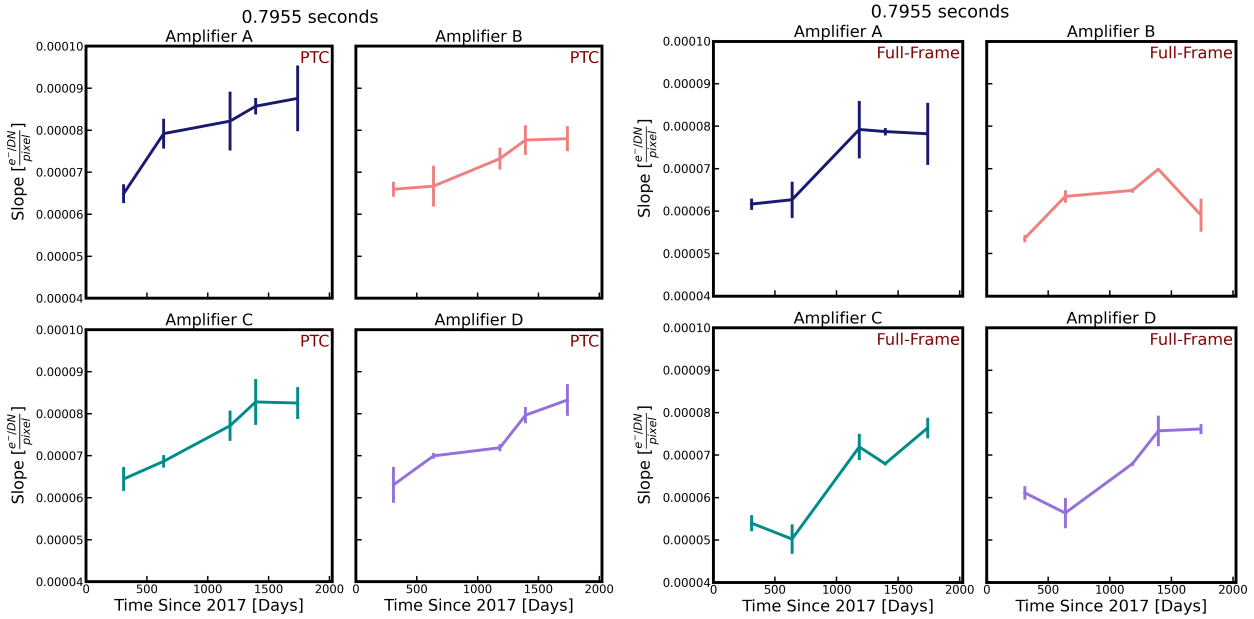

(a) Method #1: Photon-Transfer Curve **(b)** Method #2: Full-Frame

Figure 14: Same as Figure 4, but for all four amplifiers.

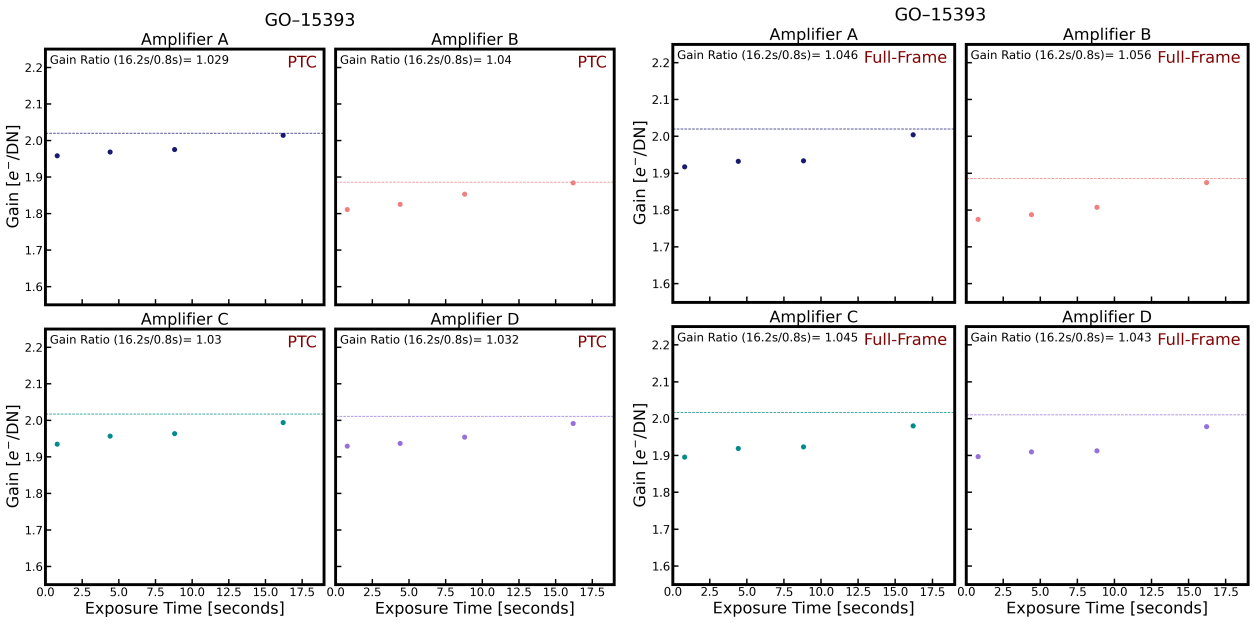

(a) Method #1: Photon-Transfer Curve

(b) Method #2: Full-Frame

Figure 15: Same as Figure 5, but for all four amplifiers.

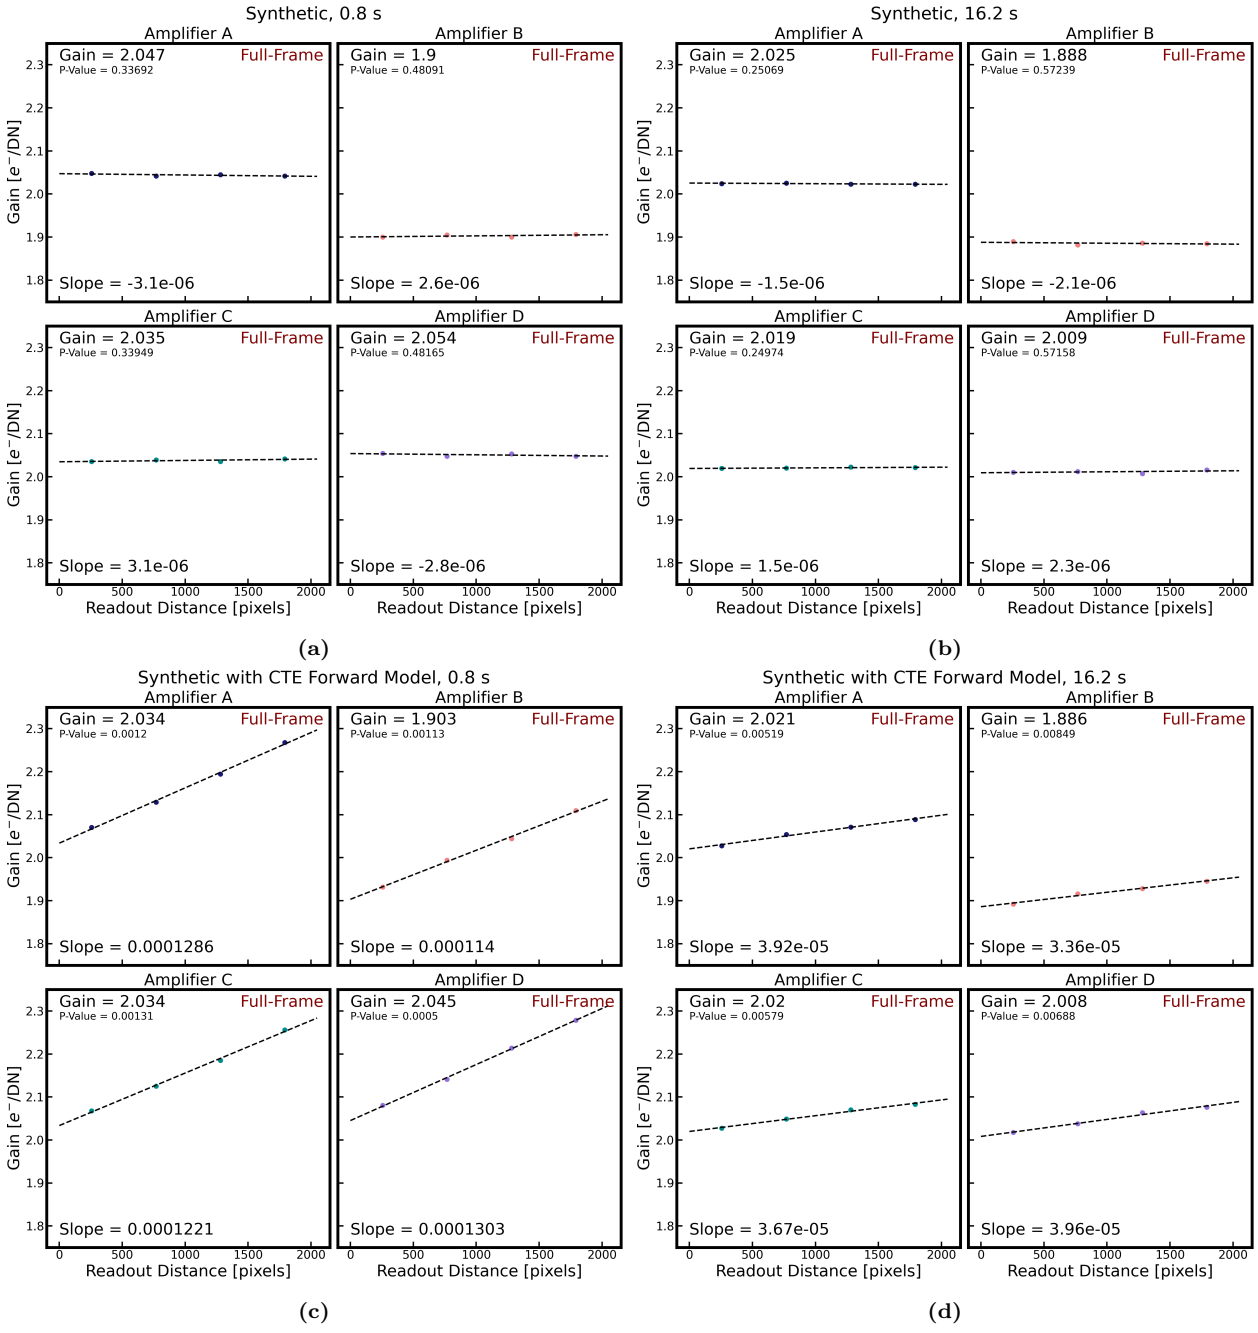

Figure 16: Same as Figure 6, but for all four amplifiers.

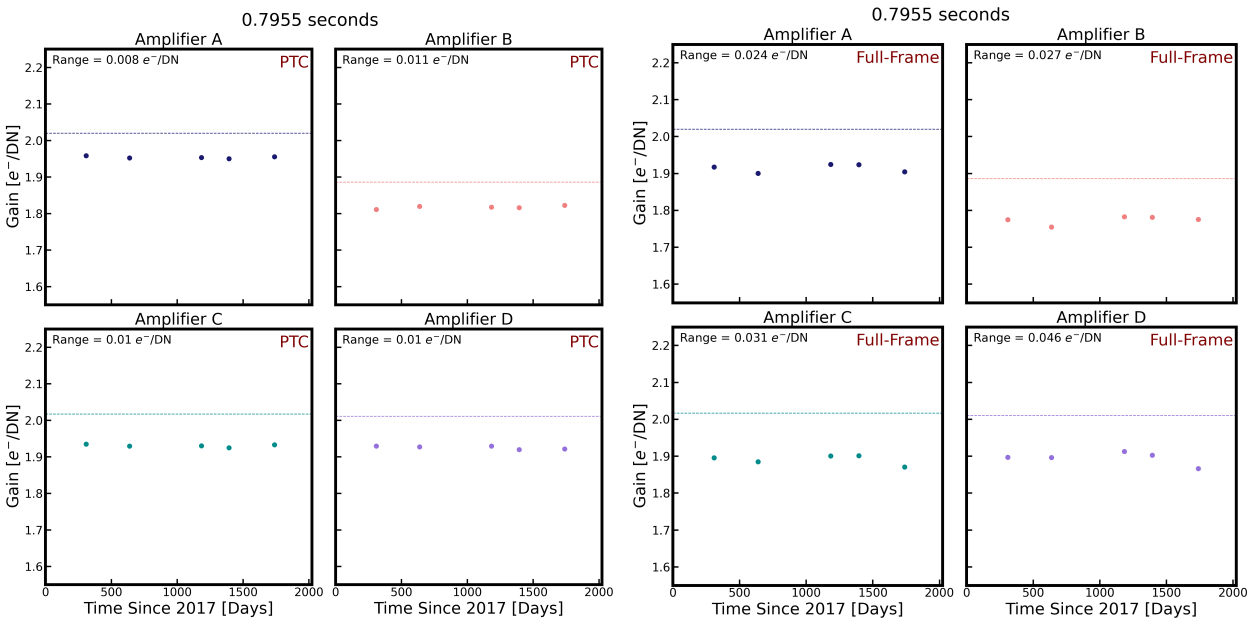

(a) Method #1: Photon-Transfer Curve

(b) Method #2: Full-Frame

Figure 17: Same as Figure 9, but for all four amplifiers.