

The fading of Cassiopeia A, and improved models for the absolute spectrum of primary radio calibration sources

A. S. Trotter,¹ D. E. Reichart,^{1★} R. E. Egger,¹ J. Stýblová,¹ M. L. Paggen,¹ J. R. Martin,¹ D. A. Dutton,¹ J. E. Reichart,² N. D. Kumar,¹ M. P. Maples,¹ B. N. Barlow,^{1,3} T. A. Berger,^{1,4} A. C. Foster,¹ N. R. Frank,¹ F. D. Ghigo,⁵ J. B. Haislip,¹ S. A. Heatherly,⁵ V. V. Kouprianov,¹ A. P. LaCluyzé,^{1,6} D. A. Moffett,⁷ J. P. Moore,¹ J. L. Stanley⁸ and S. White⁵

¹Department of Physics and Astronomy, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA

²Jordan High School, Durham, NC 27707, USA

³Department of Physics, High Point University, High Point, NC 27268, USA

⁴Institute for Astronomy, University of Hawaii, Honolulu, HI 96822, USA

⁵Green Bank Observatory, Green Bank, WV 24944, USA

⁶Department of Physics, Central Michigan University, Mount Pleasant, MI 48859, USA

⁷Department of Physics, Furman University, Greenville, SC 29613, USA

⁸Michigan State University, Lansing, MI 48824, USA

Accepted 2017 March 29. Received 2017 March 29; in original form 2017 March 20

ABSTRACT

Based on 5 yr of observations with the 40-foot telescope at Green Bank Observatory (GBO), Reichart & Stephens found that the radio source Cassiopeia A had either faded more slowly between the mid-1970s and late 1990s than Baars et al. had found it to be fading between the late 1940s and mid-1970s, or that it had rebrightened and then resumed fading sometime between the mid-1970s and mid-1990s, in the L band (1.4 GHz). Here, we present 15 additional years of observations of Cas A and Cyg A with the 40-foot in the L band, and three and a half additional years of observations of Cas A, Cyg A, Tau A and Vir A with GBO's recently refurbished 20-m telescope in the L and X (9 GHz) bands. We also present a more sophisticated analysis of the 40-foot data, and a reanalysis of the Baars et al. data, which reveals small, but non-negligible differences. We find that overall, between the late 1950s and late 2010s, Cas A faded at an average rate of 0.670 ± 0.019 per cent yr^{-1} in the L band, consistent with Reichart & Stephens. However, we also find, at the 6.3σ credible level, that it did not fade at a constant rate. Rather, Cas A faded at a faster rate through at least the late 1960s, rebrightened (or at least faded at a much slower rate), and then resumed fading at a similarly fast rate by, at most, the late 1990s. Given these differences from the original Baars et al. analysis, and given the importance of their fitted spectral and temporal models for flux-density calibration in radio astronomy, we update and improve on these models for all four of these radio sources. In doing so, we additionally find that Tau A is fading at a rate of $0.102^{+0.042}_{-0.043}$ per cent yr^{-1} in the L band.

Key words: ISM: individual objects: Cassiopeia A – ISM: individual objects: Taurus A – galaxies: individual: Cygnus A – galaxies: individual: Virgo A – radio continuum: general.

1 INTRODUCTION

The bright, but fading, supernova remnant Cassiopeia A serves as a primary calibration source for radio astronomy (e.g. Baars

et al. 1977). As such, it is important that its absolute spectrum and frequency-dependent fading rate be known as accurately as possible. To this end, Baars et al. (1977) compiled (1) absolute measurements of the flux density of Cas A, made between 1959 and 1973, and at frequencies ranging from 10 MHz to 22 GHz and (2) measurements of the fading rate of Cas A, made over 8–20 yr spans between 1949 and 1976, and at frequencies ranging from 82 MHz to 9.4 GHz. To

* E-mail: dan.reichart@gmail.com

Table 1. Summary of GBO telescopes and bands.

Telescope	Band	Central frequency	Bandwidth	Beamwidth
40-foot	<i>L</i>	1405 MHz	110 MHz	1°2
20-m	<i>L</i>	1395 MHz	80 MHz	0°75
20-m	<i>X</i>	9000 MHz	2000 MHz	0°12

the former data set, they fit their spectral model (see Section 4.2). To the latter data set, they fit their temporal model, and found that Cas A was fading at a frequency-dependent rate of (also see Dent, Aller & Olsen 1974)

$$-\frac{100}{F_\nu} \frac{dF_\nu}{dt} = 0.97(\pm 0.04) - 0.30(\pm 0.04) \log \nu_{\text{GHz}} \text{ per cent yr}^{-1}. \quad (1)$$

This corresponds to a fading rate of 0.93 ± 0.04 per cent yr^{-1} at the *L*-band frequency $\nu = 1405$ MHz.

However, Reichart & Stephens (2000) made additional 1405-MHz measurements of Cas A, relative to Cyg A, between 1995 and 1999 with the 40-foot telescope at Green Bank Observatory (GBO) in West Virginia, and found a different result. In combination with the fitted Cas A and Cyg A spectral models of Baars et al. (1977), corresponding to an epoch of 1965 and evaluated at the same *L*-band frequency, Reichart & Stephens (2000) found a fading rate of only 0.62 ± 0.12 per cent yr^{-1} .

In this paper, we present 18 more years of GBO observations, for a total of (1) 20 yr of 40-foot observations of Cas A and Cyg A in the *L* band and (2) three and a half years of observations with GBO's recently refurbished 20-m telescope (Martin et al., in preparation), of Cas A, Cyg A, Tau A and Vir A in the *L* and *X* bands (see Table 1 for telescope and band information). We present these GBO observations, and review the Baars et al. (1977) data, in Section 2. In Section 3, we (1) present a more sophisticated analysis of the 40-foot data, accounting for small pointing errors and focus differences; (2) measure the fading rate of Cas A between 1994 and 2017 in the *L* band and (3) confirm the discrepancy with Baars et al. (1977) identified by Reichart & Stephens (2000). In Section 4, we reanalyze the Baars et al. (1977) data, and find somewhat different results, but still the discrepancy. In Section 5, we combine the Baars et al. (1977) data, which also include absolute measurements of Cyg A and Tau A, and both absolute and relative measurements of Vir A, with our GBO data, and present new, and simultaneously fitted, spectral and temporal models for all four of these sources, so they may better serve as calibration sources in radio astronomy. We summarize our conclusions in Section 6.

2 DATA

2.1 Baars et al. (1977)

In addition to measurements of the fading rate of Cas A (Section 1), Baars et al. (1977) compiled absolute measurements of the flux density of Cas A, Cyg A and Tau A, and both absolute and relative measurements of the flux density of Vir A, at frequencies ranging from 10 MHz to 35 GHz, from multiple papers published between 1960 and 1975. These data were collected using a variety of instruments, including dipole, horn and dish antennas. They assigned each measurement a weight that was inversely proportional to the square of its published uncertainty. However, for dish antennas, they decreased this weight by a factor of 4, given concerns over how accurately dish antennas could be absolutely calibrated. In their spectral fits, they included only (1) measurements with less

than 6 per cent uncertainty and (2) measurements between 22 MHz and 22 GHz.

In our own fits, we include the same data that Baars et al. (1977) included in their fits (with a few, relatively minor exceptions¹), and with the same weights. However, before carrying out their Cas A spectral fits, Baars et al. (1977) first corrected, or referenced, their data to epoch 1965, using their fitted temporal model (equation 1). We on the other hand do not reference their, or our, data to a common epoch; rather, we fit our spectral and temporal models simultaneously (see Sections 4.2 and 5).

2.2 GBO 40-foot

From 1994 through 2015, we observed Cas A and Cyg A in the *L* band (centred at 1405 MHz) using GBO's 40-foot telescope. The 40-foot is a transit telescope, so we observed each source no more than once per day, as it transited the meridian. Each year, observations were carried out over a single week, resulting in 1–5 successful Cas A observations and 0–2 successful Cyg A observations. The telescope is controlled manually, and observations were carried out in drift-scan mode, in the following way. Well before transit, we would move the telescope to the approximate transit elevation of the source, and then begin moving the telescope back and forth in elevation, (1) searching for the source, and (2) collecting background information at different elevations.² Output powers were saved digitally, but were also displayed in real time using an analog strip-chart recorder. Once the source was detected, we narrowed our elevation range, until, using the strip-chart recorder, we found the elevation for which the signal was greatest. Experienced operators had enough time to check this multiple times before transit, and were likely able to point the telescope to within a tenth of a beam (i.e. $\approx 0.1 \times 1^\circ 2$) of the source's true transit elevation.

Once this was achieved, and prior to the transit of the source, the telescope's elevation was locked, to prevent the telescope from slowly drifting off of the source. The observation would then continue until the background level was again achieved, and maintained for a few minutes. Both before and after the observation, a noise diode in the receiver was activated for one minute, and then deactivated for one minute, permitting the observation to be gain calibrated, which we did by linearly interpolating between the two, albeit nearly identical, calibrations. In post-processing, we excised data that were not collected near the final, locked-in elevation.

We then modelled each drift scan, $F(\theta)$, as the sum of two functions: a source function, $F_s(\theta)$, and a background function, $F_b(\theta)$ (see Fig. 1). We took the source function to be a non-negative, even function of the form:

$$F_s(\theta) = \left\{ \sum_{n=1}^N a_n \cos \left[2\pi n \frac{(\theta - \theta_0)}{w_\theta} \right] \right\}^2. \quad (2)$$

¹ In their Cyg A and Tau A spectral fits, Baars et al. (1977) included a few additional, relative measurements from Baars & Hartsuijker (1972) and Baars, Mezger & Wendker (1965), but these are not clearly specified so we did not include them in our fits. Additionally, in their table 4, Baars et al. (1977) averaged three relative measurements of Vir A from Baars & Hartsuijker (1972) and, separately, three absolute measurements of Vir A from Janssen, Golden & Welch (1974). We have replaced these with their original, individual measurements.

² Due to terrestrial spillover, the background level is sensitive to the angle that the telescope makes with respect to the ground. Since the final elevation is not known at the beginning of the observation, we were careful to collect background information at all possible final elevations.

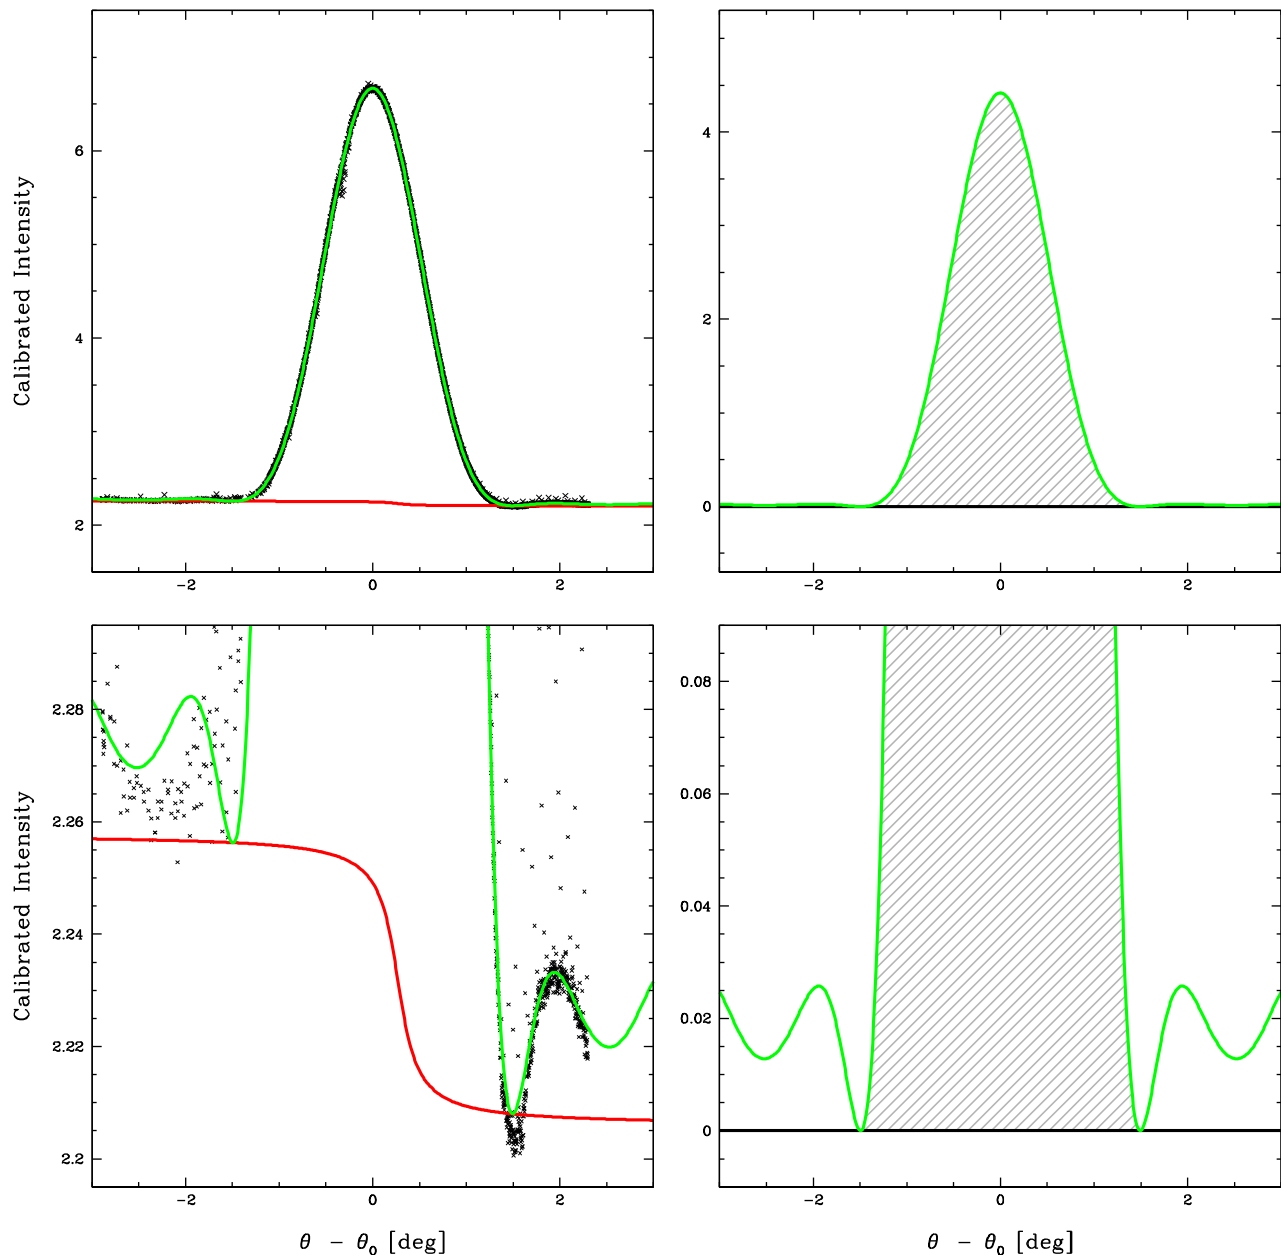


Figure 1. Example drift scan of Cas A, acquired with the 40-foot in the L band. Top left: measured intensity, in gain-calibration units, versus angular position (black crosses), and best-fitting source (green curve) plus background (red curve) functions. Data not collected near the final, locked-in elevation have already been excised. Bottom left: an expanded view, to better see the base of the best-fitting source function and the best-fitting background function. Right: best-fitting source function only (i.e. the background-subtracted best-fitting model). The shaded area, divided by the source function's peak value, F_m , defines I_m (equation 5).

The angle at which this function peaks, θ_0 , the parameter that determines its width, w_θ , and the N coefficients, a_n , were fitted to each drift scan. We found that $N = 6$ was sufficient in each case.

The background function, $F_b(\theta)$, models emission from the Milky Way and terrestrial spillover, as well as any signal drift over the course of the observation, characteristic of the instrument's stability (generally, a very small effect). We took the background function to be odd, to prevent degeneracies with the source-function parameters:

$$F_b(\theta) = b_0 + b_1 \tan^{-1} \left(\frac{\theta - b_2}{b_3} \right), \quad (3)$$

where b_0 , b_1 , b_2 and b_3 were also fitted to each drift scan, simultaneously with the source-function parameters.³ This background function is sufficiently flexible to model most likely behaviours, including everything from constant, linear, and some non-linear functions to step functions (corresponding to a sudden jump or drop in the signal level).

³ For all of the fits in this paper, we used *Galapagos*, which is in-house, parallelized, genetic algorithm-based model-fitting software. It is particularly good at fitting large-dimensional models (in the above case, 12-dimensional) without being waylaid by local solutions.

Table 2. 20 yr of GBO 40-foot *L*-band measurements of (1) the Cas A-to-Cyg A flux-density ratio, $R_m \equiv F_{m,Cas}/F_{m,Cyg}$ (see Section 3.2); (2) our measurements of the square of the integral of the corresponding peak-normalized source functions, $I_{m,Cas}^2$ and $I_{m,Cyg}^2$ and (3) the weight that we have assigned to each ratio (Section 2.2).

Epoch	R_m	$I_{m,Cas}^2$	$I_{m,Cyg}^2$	Weight
1995.62	1.316	1.474	1.435	1
1996.62	1.335	1.407	1.360	1
1996.62	1.333	1.394	1.360	1
1997.62	1.335	1.394	1.340	1
1999.62	1.244	1.395	1.398	1
2000.62	1.256	1.394	1.394	1
2000.62	1.257	1.395	1.394	1
2002.62	1.246	1.472	1.455	1
2002.62	1.260	1.492	1.455	1
2003.59	1.181	1.484	1.450	1
2004.61	1.213	1.485	1.454	2/3
2004.61	1.200	1.460	1.454	2/3
2004.61	1.196	1.476	1.454	2/3
2004.61	1.267	1.485	1.439	2/3
2004.61	1.253	1.460	1.439	2/3
2004.61	1.249	1.476	1.439	2/3
2006.53	1.248	1.462	1.425	1
2006.53	1.245	1.468	1.425	1
2007.45	1.167	1.418	1.404	2/3
2007.45	1.170	1.442	1.404	2/3
2007.45	1.183	1.413	1.404	2/3
2007.45	1.169	1.418	1.415	2/3
2007.45	1.172	1.442	1.415	2/3
2007.45	1.185	1.413	1.415	2/3
2009.57	1.243	1.450	1.391	1
2009.57	1.219	1.423	1.391	1
2010.49	1.133	1.435	1.436	1
2010.49	1.220	1.490	1.436	1
2010.49	1.200	1.510	1.436	1
2011.47	1.156	1.497	1.456	1
2011.47	1.106	1.461	1.456	1
2011.47	1.108	1.472	1.456	1
2011.47	1.164	1.496	1.456	1
2011.47	1.112	1.453	1.456	1
2012.54	1.106	1.454	1.466	2/3
2012.54	1.018	1.422	1.466	2/3
2012.54	1.108	1.455	1.466	2/3
2012.54	1.122	1.454	1.451	2/3
2012.54	1.032	1.422	1.451	2/3
2012.54	1.123	1.455	1.451	2/3
2013.51	1.032	1.427	1.503	1
2013.51	1.050	1.516	1.503	1
2013.51	0.993	1.396	1.503	1
2013.51	1.039	1.473	1.503	1
2014.43	1.153	1.446	1.428	1
2014.43	1.137	1.464	1.428	1
2014.43	1.152	1.449	1.428	1

From each best fit, we calculated (1) the measured peak intensity of the source:

$$F_m \equiv F_s(\theta_0) = \left(\sum_{n=1}^N a_n \right)^2 \quad (4)$$

and (2) the integral of the measured source function, normalized to its peak value:

$$I_m \equiv \frac{1}{F_m} \int_{\theta_0 - \Delta\theta}^{\theta_0 + \Delta\theta} F_s(\theta) d\theta, \quad (5)$$

Table 3. Summary of GBO 20-m *L*- and *X*-band observations.

Dates	Band	Source	Number of observations
2014 Jul 10–14	<i>L</i>	Cas A	24
2014 Jul 10–14	<i>L</i>	Cyg A	24
2014 July 17–Aug 5	<i>X</i>	Cas A	31
2014 July 17–Aug 5	<i>X</i>	Cyg A	30
2016 Sep 24–29	<i>L</i>	Cas A	20
2016 Sep 24–29	<i>L</i>	Cyg A	23
2017 Jan 6–8	<i>L</i>	Cas A	23
2017 Jan 5–8	<i>L</i>	Cyg A	25
2017 Jan 5–8	<i>L</i>	Tau A	23
2017 Jan 5–8	<i>L</i>	Vir A	24

where the integration limits are the symmetric, first minima on either side of the source function (Fig. 1, bottom right). In Section 3, we show that (1) the Cas A-to-Cyg A flux-density ratio, R_m , depends not only on $F_{m,Cas}$ and $F_{m,Cyg}$, but also on $I_{m,Cas}^2$ and $I_{m,Cyg}^2$ and (2) $I_{m,Cas}^2$ and $I_{m,Cyg}^2$ can also be used to account for small errors in pointing and differences in focus, at least in part.

We list these ratios, R_m , and their corresponding values of $I_{m,Cas}^2$ and $I_{m,Cyg}^2$ in Table 2, for all same-week combinations of our Cas A and Cyg A measurements. For most years, we have only one Cyg A measurement, having made other observations during Cyg A's transit time the rest of the week. For years where we have two Cyg A measurements, and three Cas A measurements, we assigned each ratio a weight of 2/3, given redundancy of information in the six possible combinations. In total, we obtained 47 (41 non-redundant) *L*-band measurements of R_m with the 40-foot over a 20-yr period.

2.3 GBO 20-m

In 2014, 2016 and 2017, we observed different combinations of Cas A, Cyg A, Tau A and Vir A in different combinations of *L* (centred at 1395 MHz) and *X* (centred at 9000 MHz) bands using GBO's 20-m telescope (see Table 3). The 20-m had recently been refurbished and integrated into the Skynet Robotic Telescope Network (Martin et al., in preparation). We submitted each observation through Skynet's web-based user interface,⁴ after which they were automatically scheduled and completed. Each *L*-band observation consisted of a 13.5-beam \times 13.5-beam on-the-fly raster, with 0.2-beam spacing between measurements, both along and between scans, and at a telescope slew speed of 0.5 s⁻¹ took 23.0 min to complete. Each *X*-band observation consisted of an 18-beam \times 18-beam on-the-fly raster, also with 0.2-beam spacing between measurements, and at a telescope slew speed of 0.3 s⁻¹ took 10.5 min to complete. In total, 247 observations were completed.

Each observation was gain calibrated, background subtracted, time-delay corrected, RFI cleaned, and surface modelled using Skynet's single-dish radio telescope image-reduction pipeline (Martin et al., in preparation; see Fig. 2). We then photometered each image using 3- and 1.5-beam radius apertures for the 2014 *L*- and *X*-band images, respectively, and using 1.25-beam radius apertures for the 2016 and 2017 *L*-band images (focus had been improved for the latter *L*-band images; see Fig. 3).

We determined flux-density ratios by pairing each Cas A, Tau A or Vir A measurement with its closest-in-time Cyg A measurement, and by pairing each Cyg A measurement with its closest-in-time Cas A, Tau A or Vir A measurement. We then measured the mean

⁴ <https://skynet.unc.edu>




Figure 2. Example on-the-fly raster of Cas A, acquired with the 20-m in the X band. For all of our 20-m observations, signal was integrated and read out every 0.2 beams (corresponding to every 7 arcmin in this X-band case) along scans in declination (plotted here vertically), and scans were separated also by 0.2 beams, in right ascension (plotted here horizontally). In this case, we have zoomed into the central 5.7 beams \times 5.7 beams, corresponding to 40 arcmin \times 40 arcmin. Top left: gain-calibrated measurements. Top right: after background subtraction on the recommended, 6-beam scale. Bottom left: after time-delay correction, by which we mean that occasionally the clocks of the telescope's separate coordinate and signal computers lose synchronization, which results in a hysteretic splitting of the source along the scan direction, prior to correction. Bottom right: after RFI cleaning, on the recommended 0.8-beam scale. In all four panels, scale-weighted surface models have been fitted to the data about each pixel, for visualization. In the first three panels, a minimal scale is used, for maximal surface flexibility. In the fourth panel, the recommended, 1/3-beam (near-Nyquist) scale is used, and this is the final image. A non-linear colour spectrum is used to emphasize fainter structures (units are dimensionless, with one corresponding to the noise diode). Details regarding the Skynet Robotic Telescope Network's single-dish image-reduction pipeline will be available in Martin et al. (in preparation).




Figure 3. Example *L*-band images of Cas A, made with the 20-m in 2014 (left) and 2016 (right). Each is $13.5 \text{ beams} \times 13.5 \text{ beams}$, corresponding to $10''.1 \times 10''.1$, with declination plotted vertically and right ascension plotted horizontally. A non-linear colour spectrum is used to emphasize fainter structures (units are dimensionless, with one corresponding to the noise diode). The 2014 images are somewhat out of focus, requiring a larger aperture (green circle) for photometry. This had been corrected by 2016.

Table 4. GBO 20-m *L*- and *X*-band flux-density ratio measurements, with respect to Cyg A.

Mean epoch	Band	Source	Number of ratios	Number rejected	R_m
2014.53	<i>L</i>	Cas A	48	3	1.123 ± 0.019
2014.58	<i>X</i>	Cas A	61	0	2.46 ± 0.19
2016.74	<i>L</i>	Cas A	43	1	1.115 ± 0.021
2017.02	<i>L</i>	Cas A	48	16	1.103 ± 0.021
2017.02	<i>L</i>	Tau A	48	11	0.529 ± 0.012
2017.02	<i>L</i>	Vir A	49	17	0.1370 ± 0.0015

and the uncertainty in the mean for each of these (six; see Table 4) collections of ratios, after rejecting outliers using Robust Chauvenet Rejection (RCR; Maples et al. 2017). More outliers were found with the Tau A and Vir A ratios than with the Cas A ratios, because, being on the other side of the sky than Cyg A, these pairs were generally observed with a greater separation in time. The 2017 Cas A ratios, however, were an exception, in that observations of Cyg A began a full day before observations of Cas A.

3 GBO MODELS AND FITS

3.1 GBO 40-foot pointing-error and non-constant focus models

From our first 5 yr of GBO 40-foot observations, Reichart & Stephens (2000) found that the ratio of Cas A's peak intensity to that of Cyg A, $F_{m,Cas}/F_{m,Cyg}$, was significantly greater than expected, given the fitted spectral and temporal models of Baars et al. (1977).

In this section, we investigate if our 40-foot measurements could be biased high, due to (1) imperfect pointing of the telescope and/or (2) non-constant focus of the telescope.

Imperfect pointing. If during a drift scan the operator fails to zero-in on the true transit elevation of the source (Section 2.2), the peak intensity that we would then measure for that source would be underestimated. Consequently, our higher-than-expected Cas A-to-Cyg A peak-intensity ratios could, at least in theory, be due to systematically greater pointing errors when we observed Cyg A. However, we regard this as unlikely: each year, our scans of Cyg A were taken by the same, most-experienced member of our team, while most (but not all) of our scans of Cas A were taken by first-time, student operators. Hence, if anything, we would expect our Cas A-to-Cyg A peak-intensity ratios to be biased low due to imperfect pointing, not high.

Non-constant focus. Changes in temperature between observations of Cas A and Cyg A, and/or elevation-dependent differences in how gravity pulls on the dish and the receiver and its support structure could, at least in theory, change the focus of the telescope. Non-ideal focus would result in a broader point spread function, and consequently, a lower peak intensity. Hence, if the conditions under which we took our Cyg A observations were such that focus was systematically poorer than for our Cas A observations, this could also explain our higher-than-expected Cas A-to-Cyg A peak-intensity ratios. But again, we regard this as unlikely: at the time of year that we took our observations, between June and August, both Cyg A and Cas A transit the meridian at night, or early in the morning, within 3.4 h of each other, when temperatures at GBO are fairly stable. Furthermore, both sources transit the meridian at similar high elevations (87° for Cyg A and 69° for Cas A), so we do not

expect gravitational distortion of the telescope to differ significantly either.

Although we anticipate that neither of these effects, imperfect pointing and non-constant focus, can account for our higher-than-expected Cas A-to-Cyg A peak-intensity ratios, we still model both effects (1) to be sure and (2) to ensure that our fits to these data (see Sections 3.2 and 5) are as accurate as possible. As we will see in these sections, these effects are indeed small, generally offsetting, and insufficient to explain the discrepancy with Baars et al. (1977).

Pointing-error model. Imperfect pointing has two effects: (1) the measured peak intensity is less than the true peak intensity ($F_m < F_t$) and (2) the measured peak-normalized source function is likely different from the true, or perfect-pointing, peak-normalized source function, resulting in $I_m^2 \neq I_t^2$.⁵ We model the consequent correlation, between measured peak intensity and the square of the integral of the measured peak-normalized source function, to first order, since pointing errors, if not negligible, were at least small (Section 2.2):

$$\frac{F_m}{F_t} = \left(\frac{I_m^2}{I_t^2} \right)^\alpha, \quad (6)$$

where α is determined empirically in Sections 3.2 and 5.⁶

Non-constant focus model. Non-constant focus would imply that the 40-foot's true, or perfect-pointing, peak-normalized source function differs between our observations of Cas A and Cyg A.⁷ If there is a systematic component to this, we take it to be small, and, as above, model it to first order:⁸

$$I_{t,Cas}^2 = \beta I_{t,Cyg}^2, \quad (7)$$

where β is also determined empirically in Sections 3.2 and 5. That said, we can place a prior on β , by fitting equation (7) to our measured values, $I_{m,Cas}^2$ and $I_{m,Cyg}^2$.⁹ As can be seen in Fig. 4, these values appear to have changed in concert from year to year. We find




Figure 4. Measured values of the square of the integral of the peak-normalized source function from Table 1, plotted $\log I_{m,Cas}^2$ versus $\log I_{m,Cyg}^2$. The line is the best fit of equation (7) to these, weighted data. Open circles have been rejected using Robust Chauvenet Rejection (RCR; Maples et al. 2017), and correspond to pairs of measurements that likely suffer from significantly different pointing errors; they do not contribute to the fit. The shaded regions show the 1σ , 2σ and 3σ statistical uncertainty in the best-fitting value of $\log \beta$. The uncapped error bars show the best-fitting 1σ sample scatter, which in this case is two-dimensional (see Trotter 2011, or Reichart 2001).

that $\log \beta = 0.0070 \pm 0.0010$,¹⁰ with a best-fitting sample scatter of 0.0059 dex. Consequently, in our upcoming fits to 40-foot data (see Sections 3.2 and 5), we adopt a Gaussian prior in $\log \beta$, of mean 0.0070 and width 0.0010 dex.

3.2 GBO 40-foot and 20-m model and fit

Of course, it is not the ratio of Cas A's and Cyg A's *peak* intensities that we must model, but the ratio of their *integrated* intensities, which do not depend on the telescope's precise pointing or focus, nor on the intrinsic size of the source. With the 20-m, this is straightforward. We densely sample each source's 2D distribution on the sky, and can reconstruct and integrate these distributions (Section 2.3). But with the 40-foot, we have only single, 1D slices through these 2D distributions (Section 2.2). However, if one assumes radial symmetry, or at least the same shape, if not the same size, for these sources' 2D distributions, integrated intensity is then simply

$$S_t \propto F_t I_t^2, \quad (8)$$

where subscript t again refers to true, or perfect-pointing, values. Given this, we define the measured Cas A-to-Cyg A flux-density ratio analogously:

$$R_m = \frac{S_{m,Cas}}{S_{m,Cyg}} \equiv \frac{F_{m,Cas} I_{m,Cas}^2}{F_{m,Cyg} I_{m,Cyg}^2}. \quad (9)$$

¹⁰ Note, if due only to the difference in these sources' intrinsic angular sizes, given the telescope's large beamwidth, $\log \beta$ would be no more than ≈ 0.003 , and probably half of that, both of which are ruled out.

⁵ In the idealized case of a Gaussian point spread function, and a point, or Gaussian, source, $I_m^2 = I_t^2$, regardless of pointing.

⁶ If the point spread function has narrower wings than a Gaussian, which is usually the case, $I_m^2 < I_t^2$ and hence α must be positive; else, $\alpha < 0$.

⁷ Note, for the 40-foot at the L band, this function is not appreciably different from the telescope's beam function, for either of these sources. Cas A is ≈ 6 arcmin across and Cyg A is ≈ 2 arcmin across. Convolution of the telescope's beam function with such intrinsic distributions can be, at least roughly, approximated by adding their FWHMs in quadrature, which yields functions that are, at most, ≈ 0.3 per cent and ≈ 0.04 per cent wider, respectively.

⁸ We also tried two-parameter models, such as $I_{t,Cas}^2 = \beta I_{t,Cyg}^2 + \gamma$ and $I_{t,Cas}^2 = \beta I_{t,Cyg}^2$, but the additional parameter was never requested by the data.

⁹ Specifically, we fit $\log I_{t,Cas}^2 = \log \beta + \log I_{t,Cyg}^2$ to the logarithm of our measured values, while simultaneously rejecting outliers using Robust Chauvenet Rejection (RCR; Maples et al. 2017), to avoid the fit being biased by pairs of measurements with significantly different pointing errors. Indeed, $I_{m,Cas}^2 < I_{m,Cyg}^2$ for five of the six rejected points (see Fig. 4), corresponding to some of our student-led observations of Cas A having significantly greater pointing errors than their corresponding professional-led observations of Cyg A, assuming $\alpha > 0$ (see Sections 3.2 and 5). We measure the best-fitting value of $\log \beta$, the statistical uncertainty in this best-fitting value, and the sample scatter. By the latter, we mean the scatter of the data about the model that is not accounted for by the data's error bars (which in this case, we take to be zero). Like $\log \beta$, sample scatter is a model parameter, which we fit to the data (see Trotter 2011, or Reichart 2001, for a description of the maximum-likelihood technique).

These are the values that we list in Table 2. Note, equation (9) reduces to $R_m = F_{m,Cas}/F_{m,Cyg}$ in the limit of perfect pointing and identical 2D distributions.

Substituting our pointing-error (equation 6) and non-constant focus (equation 7) models into equation (9) yields

$$R_m = \frac{F_{t,Cas} I_{t,Cas}^2}{F_{t,Cyg} I_{t,Cyg}^2} \left(\frac{I_{m,Cas}^2}{\beta I_{m,Cyg}^2} \right)^{\alpha+1} \equiv R_t \left(\frac{I_{m,Cas}^2}{\beta I_{m,Cyg}^2} \right)^{\alpha+1}, \quad (10)$$

where $R_t \equiv F_{t,Cas} I_{t,Cas}^2 / F_{t,Cyg} I_{t,Cyg}^2$ is the true, or perfect-pointing, Cas A-to-Cyg A flux-density ratio at the time of the observations.

Given the size of Cyg A's emitting region,¹¹ Cyg A's brightness should not change over the time-scale of these observations. Consequently, the observed decrease in R_t with time must be due to Cas A fading. In keeping with Baars et al. (1977), we model this fading exponentially:¹²

$$R_t(t) = R_0 e^{m_{v_0}(t-t_0)}, \quad (11)$$

where $R_0 \equiv R_t(t_0)$; t_0 is a reference time, which in this paper we set equal to the weighted mean of whichever observations we are fitting to;¹³ and m_{v_0} is the fading rate, in this case at $v_0 = 1405$ MHz.¹⁴ In units of percent per year, this is $-100 \times \ln 10 \times m_{v_0}$.

Substituting equation (11) into equation (10) and taking the logarithm of both sides yields

$$\log R_m(t) = \log R_0 + m_{v_0}(t - t_0) + (\alpha + 1) \log \left(\frac{I_{m,Cas}^2}{I_{m,Cyg}^2} \right) - (\alpha + 1) \log \beta, \quad (12)$$

where the first two terms model the normalization and fading of the flux-density ratio, respectively; the third term accounts for imperfect pointing of the telescope, for 40-foot data; and the fourth term accounts for systematic differences in focus between observations of Cas A and Cyg A, again for 40-foot data.¹⁵ Equation (12) can also, and simultaneously, be fitted to our 20-m *L*-band Cas A-to-Cyg A flux-density ratios (Table 4), by dropping the last two terms for these data.¹⁶

¹¹ Cyg A's lobes dominate the emission, with its nucleus contributing only ≈ 1 Jy at GHz frequencies (e.g. Kellermann et al. 1981).

¹² Vinyaikin (2006) (also, see Shklovskii 1960) argues that a power-law model would be more appropriate, physically, but also points out that this matters little, given the time-scale over which these observations were made, compared to the age of Cas A.

¹³ Choice of reference value, be it a time or a frequency, affects the best-fitting value of the normalization parameter, but not the best-fitting model itself. Given this freedom, one often selects simple, easy-to-use or easy-to-remember reference values; e.g. Baars et al. (1977) used 1965 and 1 MHz, respectively. However, choice of reference value also affects the degree to which the uncertainties in the best-fitting parameter values are correlated (e.g. the uncertainties that Baars et al. 1977 measured are highly correlated). By calculating reference values as we do here, we minimize these correlations (e.g. Maples et al. 2017), the consequence being that we can easily, and analytically, approximate the uncertainty in our best-fitting models, as a function of time in this case, and as a function of time and frequency in latter cases (e.g. see equation 15).

¹⁴ In Sections 4 and 5, we instead use a frequency-dependent model.

¹⁵ The fourth term can also be used to account for systematic differences in the shape of each source's 2D distribution, but only in combination with data for which this does not need to be accounted for, such as our 20-m data; see below.

¹⁶ Note, since the 40-foot and these 20-m measurements were made at slightly different frequencies, 1405 MHz versus 1395 MHz, there should be a slight offset between their best-fitting models. However, given the




Figure 5. 23 yr of GBO 40-foot (blue) and 20-m (red) *L*-band Cas A-to-Cyg A flux-density ratios versus time, and best-fitting model (solid line). Open circles are measured values. Filled circles are after subtracting off the best-fitting third and fourth terms of equation (12), which account for imperfect pointing and non-constant focus of the 40-foot telescope. The shaded regions show the 1σ , 2σ and 3σ statistical uncertainty in the best-fitting values of $\log F_0$ and m_{v_0} . The uncapped error bars show the best-fitting 1σ sample scatter of the 40-foot measurements (that of the 20-m measurements is almost negligible). The significantly lower dashed line and shaded regions are from the fitted spectral and temporal models of Baars et al. (1977).

We plot our best fit of equation (12) to all 23 yr of GBO 40-foot and 20-m *L*-band measurements in Fig. 5. We find that $\log R_0 = 0.0557^{+0.0011}_{-0.0019}$ (at $t_0 \equiv 2013.65$), $m_{v_0} = -0.00348^{+0.00022}_{-0.00023}$ (corresponding to $0.802^{+0.053}_{-0.050}$ per cent yr^{-1}), $\alpha = 0.64 \pm 0.15$, and $\log \beta = 0.00754^{+0.00090}_{-0.00089}$, with best-fitting sample scatters of 0.0099 dex for the 40-foot data and 0.0014 dex for the 20-m data.

The fitted pointing-error and non-constant focus ‘correction’ terms in equation (12) are, as expected, small, but not negligible. The mean pointing-error correction is 0.00774 ± 0.00073 dex, or 1.80 ± 0.17 per cent. This suggests that, on average, our pointing was a bit better for Cyg A than for Cas A, as expected (Section 3.1). The non-constant focus correction is -0.0123 ± 0.0019 dex, or -2.80 ± 0.42 per cent (this uncertainty is 44 per cent anticorrelated with the uncertainty in the mean pointing-error correction). This suggests that the 40-foot’s focus was a bit worse when observing Cas A than when observing Cyg A, for either of the reasons postulated in Section 3.1. The former correction, visibly, accounts for some of the scatter of the data about the best-fitting model (Fig. 5), and raises the best-fitting model. The latter correction lowers the best-fitting model, to be in accordance with the 20-m data, although $\log \beta$ is also constrained by the prior that we were able to place on it in Section 3.1. The sum of these corrections is -0.0046 ± 0.0015 dex or -1.05 ± 0.35 per cent, which falls significantly short of the roughly -0.06 ± 0.01 dex or -15 ± 2 per cent correction that would be required to explain the discrepancy with the fitted spectral and temporal models of Baars et al. (1977), first identified by Reichart

spectral indices of Cas A and Cyg A at the *L* band (see Baars et al. 1977, Sections 4.2 and 5), this offset should be only ≈ 7 per cent of the size of these 20-m measurements’ error bars (Table 4), and consequently can be safely ignored.




Figure 6. Fading-rate measurements from table 1 of Baars et al. (1977), and best-fitting models. The original result of Baars et al. (1977) is on the left; ours is on the right. The shaded regions show the 1σ , 2σ and 3σ statistical uncertainty in the best-fitting values of m_{ν_0} and $m_{\Delta \log \nu}$. The best-fitting sample scatter (at least in the case of our fit) is negligible. We were unable to reproduce the original result specifically, but do so generally.

& Stephens (2000) (Fig. 5). Furthermore, the 20-m measurements, which are not subject to these corrections, confirm this discrepancy.

4 MODELS AND FITS TO THE BAARS ET AL. (1977) DATA

Now that we have established that Cas A is, and has been since at least the mid-1990s, significantly brighter than expected, given the fitted spectral and temporal models of Baars et al. (1977), at least in the L band, two likely explanations remain: (1) Cas A rebrightened sometime between the end of the Baars et al. (1977) observations and the beginning of our GBO observations, so between the mid-1970s and mid-1990s, after which it resumed fading or (2) the fitted spectral and/or temporal models of Baars et al. (1977) are erroneous. Given this, and given the importance of these fitted models for flux-density calibration in radio astronomy, we now reanalyse the original Baars et al. (1977) data.

4.1 Temporal model and fit

We begin by fitting the same, frequency-dependent fading-rate model to the same fading-rate measurements that Baars et al. (1977) did. The fading-rate model is given by

$$-\frac{100}{F_\nu} \frac{dF_\nu}{dt} = m_{\nu_0} + m_{\Delta \log \nu} \log \left(\frac{\nu}{\nu_0} \right), \quad (13)$$

where we have decided to normalize the function at the weighted mean of the data's $\log \nu$ values (corresponding to $\nu_0 \equiv 1315$ MHz), instead of at 1 GHz. The fading-rate measurements are from table 1 of Baars et al. (1977).

We plot our best fit of equation (13) to these measurements in the right-hand panel of Fig. 6. We find that $m_{\nu_0} = 0.894 \pm 0.021$ per cent yr^{-1} and $m_{\Delta \log \nu} = -0.284 \pm 0.053$ per cent yr^{-1} , with a best-fitting sample scatter of zero. This differs from the original result of Baars et al. (1977) (left-hand panel), which has

marginally different best-fitting values, but significantly different uncertainties in these best-fitting values. Baars et al. (1977) appears to have overstated the uncertainty in the best-fitting model around 1 GHz, and understated it below ≈ 100 MHz.¹⁷ Despite this, our fitted temporal models are sufficiently similar that this is unlikely to be the source of the discrepancy in Fig. 5.

4.2 Spectral model and fits

Baars et al. (1977) used their best-fitting temporal model to reference all of their measurements to epoch 1965, after which they fitted their spectral model to these ‘corrected’ measurements. Generally, this is a poor modelling practice, in that it is seldom clear how all of the various uncertainties propagate.¹⁸ Instead, we fit our spectral and temporal models simultaneously (Section 2.1), adopting the results of the previous section as priors. Otherwise, our only difference is that we use a slightly different spectral model; where Baars

¹⁷ In an attempt to determine what Baars et al. (1977) might have done incorrectly, or at least differently, we also tried an unweighted fit, with $\nu_0 \equiv 1$ GHz (as in equation (1)). Although this was much closer to the original result, it still did not reproduce it exactly.

¹⁸ If the measurement uncertainties are small, and the uncertainties in the best-fitting values of m_{ν_0} and $m_{\Delta \log \nu}$ are also small, and uncorrelated, this can be done with some reliability – for a single measurement. But for multiple measurements, this only introduces correlations between the ‘corrected’ measurements, which is problematic when one attempts to fit a model to them, because most statistics assume fully independent data. Alternatively, if the uncertainties in the best-fitting values of m_{ν_0} and $m_{\Delta \log \nu}$ are simply ignored, as was likely the case here, each ‘corrected’ measurement will be in error by some, different amount, and these errors are not guaranteed to be offsetting. In short, it is always better to ‘bring the model to the data’ (e.g. as we did in Section 3), than to attempt the opposite. This said, note that one often references data to a common frequency or epoch, etc., when plotting. But this is for visualization only. Model fits to these data should always be ‘forward folded’.




Figure 7. Absolute (black) and relative (green) measurements of the flux density of Cas A from Baars et al. (1977) and best-fitting temporal models. Relative measurements have been multiplied by the best-fitting spectral model of the reference source, and all measurements have been referenced to 1405 MHz using the best-fitting spectral model of Cas A, for better visualization. Open circles are the GBO 40-foot (blue) and 20-m (red) measurements from Fig. 9, and are not included in these fits. The fit on the left is to just the plotted, absolute Cas A data. The fit on the right is a simultaneous fit to all Baars et al. (1977) flux-density data. Both fits are additionally constrained by the fading-rate priors of Section 4.1, from the right-hand panel of Fig. 6. The shaded regions show the 1σ , 2σ and 3σ statistical uncertainty in the best-fitting values of $\log F_0$, a_1 , a_2 , a_3 , m_{v_0} and $m_{\Delta \log v}$. The best-fitting sample scatter is negligible. This confirms that the Baars et al. (1977) data do indeed imply a trend that undercuts our GBO data at later times (Fig. 5).

et al. (1977) used piecewise linear and quadratic models, we use a third-order polynomial model over the entire, 22 MHz to 22 GHz, spectral range. Our combined spectral and temporal model is then given by

$$\log F_m(\nu, t) = \log F_0 + a_1 \log \left(\frac{\nu}{\nu_{\text{Cas}}} \right) + a_2 \left[\log \left(\frac{\nu}{\nu_{\text{Cas}}} \right) \right]^2 + a_3 \left[\log \left(\frac{\nu}{\nu_{\text{Cas}}} \right) \right]^3 + \left[m_{v_0} + m_{\Delta \log v} \log \left(\frac{\nu}{\nu_0} \right) \right] (t - t_0), \quad (14)$$

where the first term is the normalization; the next three terms are our spectral model, normalized at the weighted mean of the data's $\log \nu$ values; and the last term is our frequency-dependent temporal model, where, a priori, we take m_{v_0} to be distributed as a Gaussian of mean 0.894 per cent yr^{-1} and width 0.021 per cent yr^{-1} , $m_{\Delta \log v}$ to be distributed as a Gaussian of mean -0.284 per cent yr^{-1} and width 0.053 per cent yr^{-1} , and $\nu_0 \equiv 1315$ MHz (Section 4.1).

In the left-hand panel of Fig. 7, we fit equation (14) to the exact same measurements that Baars et al. (1977) did, the absolute measurements of Cas A listed in their table 2. In the right-hand panel, we expand upon this by additionally fitting to all of their relative measurements of Cas A, which are with respect to Vir A. This requires simultaneously fitting to Vir A, and as some of these measurements are relative to Cyg A and Tau A, this requires simultaneously fitting to all four sources. For this, we adopt the same, third-order polynomial spectral model for each source, but with independent parameters and normalization reference

frequencies.¹⁹ In either case, we confirm that the Baars et al. (1977) data do indeed imply a trend that undercuts our GBO data at later times, albeit by a marginally smaller amount, but also at a greater level of significance, compared to the original result (Fig. 5).

5 COMBINED MODELS AND FITS

We have now established that neither our more recent, GBO measurements (Section 3) nor the original, fitted spectral and temporal models of Baars et al. (1977) (Section 4) have misdirected us (although we were not able to reproduce the latter exactly, and with more powerful statistical and computing techniques, have been able to improve on these models, especially within certain frequency ranges; see below). Consequently, we conclude that Cas A rebrightened, or at least faded at a much slower rate, between the end of the Baars et al. (1977) observations and the beginning of our GBO observations (Section 4).

In this section, we begin by modelling the average fading rate of Cas A between the late 1950s and late 2010s, by ignoring the fading-rate priors of Section 4.1, which clearly apply only to the epoch of the Baars et al. (1977) data. We then reintroduce these priors, and model the fading of Cas A in two ways: (1) assuming a rebrightening event between the last Baars et al. (1977) observation and our first GBO observation and (2) assuming a period of much slower fading between these epochs.

For both of these models, we fit to all data (Cas A, Cyg A, Tau A and Vir A) simultaneously, as we did in Section 4.2. Given the longer time-scale over which these observations were made, we also

¹⁹ Altogether, this is a 22-parameter model, but not a problem for our *Galapagos* software (Section 2.2).




Figure 8. Absolute (black) and relative (green) measurements of the flux density of Cas A from Baars et al. (1977), our (also relative) GBO 40-foot (blue) and 20-m (red) measurements, and our best fit, one-segment temporal model of both epochs. Relative measurements have been multiplied by the best-fitting spectral model of the reference source, and all measurements have been referenced to 1405 MHz using the best-fitting spectral model of Cas A, for better visualization. 40-foot measurements have additionally been corrected for imperfect pointing and non-constant focus using the best-fitting values of α and β , as in Fig. 5, also for better visualization. The fading-rate priors of Section 4.1 apply only to the epoch of the Baars et al. (1977) data and consequently have not been included in this fit. The shaded regions show the 1σ , 2σ and 3σ statistical uncertainties in the best-fitting values of $\log F_0$, a_1 , a_2 , a_3 , m_{v_0} and $m_{\Delta \log v}$. The uncapped error bars show the best-fitting 1σ sample scatter of the 40-foot measurements (that of the other measurements is negligible).

permit Tau A to fade (neither Cyg A nor Vir A should). However, we do not have enough late-time data to measure if this fading also has a frequency dependence, so for Tau A, we simply measure its fading rate in the L band. Furthermore, we do not have enough late-time data to measure if the frequency dependence of Cas A's fading rate is also changing, so we take it to be constant across all epochs.

5.1 Average fading model and fit

Our fading-rate priors (Section 4.1) derive from eight fading-rate measurements. On average, these measurements were made over 12.9 ± 4.4 yr spans of time, beginning in 1959.1 ± 4.8 and ending in 1972 ± 2.1 . As such, these priors are germane really only to this relatively brief period, which matches the epoch of the rest of the Baars et al. (1977) data almost perfectly. So, if we wish to characterize Cas A with a single fading rate, representative of all of the epochs for which we have data, we should not necessarily bind ourselves with these priors. Dropping them for this fit only yields: $\log(F_0/1 \text{ Jy}) = 3.2494^{+0.0041}_{-0.0042}$ (at $t_0 = 2005.64$), $m_{v_0} = 0.663 \pm 0.019$ per cent yr^{-1} , $m_{\Delta \log v} = -0.317 \pm 0.067$ per cent yr^{-1} , and a best-fitting sample scatter for all but the 40-foot data (i.e. for all data with measured uncertainties) of zero (see Fig. 8). At 1405 MHz, this corresponds to a fading rate of 0.670 ± 0.019 per cent yr^{-1} , which is consistent with our initial finding of 0.62 ± 0.12 per cent yr^{-1} (Reichart & Stephens 2000). This is also consistent with the finding of Vinyaikin (2007a), who measured a fading rate of 0.67 ± 0.04 per cent yr^{-1} between 1977 and 2004 at the nearby frequency of 927 MHz using the 10-m diameter radio telescope at Staraya Pustyn' Radio Astron-

omy Observatory in Russia. However, all of these measurements are inconsistent with the fading rate measured during the epoch of the Baars et al. (1977) data: 0.93 ± 0.04 per cent yr^{-1} at 1405 MHz as measured by Baars et al. (1977), or 0.886 ± 0.021 per cent yr^{-1} at 1405 MHz as we have remeasured it, in Section 4.1.

5.2 Rebrightening models and fits

Next, we fit two- and three-segment models to the same data. In the case of the two-segment model, we fit one segment to the Baars et al. (1977) data, one segment to our GBO data, and we assume that an, unmodelled, rebrightening event took place between these segments. In the case of the three-segment model, we assume a less-dramatic rebrightening, consisting merely of a period of slower fading, between two faster-fading segments. In both cases, we include the fading-rate priors of Section 4.1, applying them to the first segment only. We plot our best-fitting two- and three-segment models in the left- and right-hand panels of Fig. 9, respectively.

Compared to a one-segment model with the fading-rate priors also applied, these two- and three-segment, rebrightening models are favoured at the 6.8σ ($\Delta\chi^2 = 49.8$, $\Delta\nu = 2$) and 6.3σ ($\Delta\chi^2 = 50.4$, $\Delta\nu = 4$) credible levels, respectively.

Vinyaikin (2006, 2007a, 2014) (and references therein) observed Cas A and Cyg A between the two primary epochs considered in this paper, and at a variety of frequencies. We did not include these data in our analysis, because the scatter of these data (compared to their measured uncertainties) suggests that systematics might be an issue, and we do not know enough about these telescopes to model possible systematics (e.g. as we did in Section 3). However, these data are still useful to us in that they show no evidence of a sudden rebrightening, but are generally consistent with monotonic fading, which favours our three-segment model over our two-segment model. We present best-fitting parameter values and uncertainties for this, 27-parameter, model in Table 5.

Both of our rebrightening models, as well as the one-segment model from Section 5.1, result in nearly identical fitted spectral models. We plot these for our favoured, three-segment model, and for all four sources, in Fig. 10 (also see Perley & Butler 2016 for an independent analysis).

To examine the quality of the original, fitted spectral models of Baars et al. (1977), we plot their difference from our updated and improved versions in Fig. 11. As we have already established, Baars et al. (1977) overpredicted the fading of Cas A, and consequently their Cas A spectrum is now too low. Conversely, they assumed no fading for Tau A, and consequently their Tau A spectrum now appears to be too high (see below).

However, even if one restricts oneself to an apples-to-apples comparison at epoch 1965, there are still significant differences. In particular, the Baars et al. (1977) Cyg A spectrum is (1) discontinuous, (2) significantly high (up to ≈ 7 per cent) in the few-GHz range and (3) significantly low (up to ≈ 6 per cent) in the many-GHz range. And the Baars et al. (1977) Tau A spectrum is significantly high (up to ≈ 11 per cent) in the couple-GHz range. These differences are much larger than the ≈ 2 per cent accuracy quoted by Baars et al. (1977).

Finally, for Tau A, we measure a fading rate of $m_{v_0} = 0.102^{+0.042}_{-0.043}$ per cent yr^{-1} in the L band (see Fig. 12). Vinyaikin (2007b) found a similar result: 0.18 ± 0.10 per cent yr^{-1} , from observations made between 1977 and 2000 at 927 MHz.

For all of our fitted models, we have made every effort to select reasonable reference frequencies and times, in an attempt to minimize correlations between the fitted parameters. Given this, one may approximate the positive and negative 1σ uncertainties in any




Figure 9. The same as Fig. 8, but for our best-fitting two-segment (left) and three-segment (right) temporal models. The fading-rate priors of Section 4.1 have been applied to the first segment of each model.

Table 5. Best-fitting parameter values and 1σ statistical uncertainties in these values for our favoured, three-segment model.

Parameter	Cas A segment specific		
	Segment 1	Segment 2	Segment 3
$\log F_0$	3.3731 ± 0.0037	3.3173 ± 0.0047	$3.2530^{+0.0051}_{-0.0050}$
m_{ν_0}	$-0.003884^{+0.000092}_{-0.000093}$ (0.894 ± 0.021 per cent yr^{-1})	-0.00227 ± 0.00024 ($0.522^{+0.056}_{-0.054}$ per cent yr^{-1})	$-0.00350^{+0.00022}_{-0.00025}$ ($0.806^{+0.058}_{-0.050}$ per cent yr^{-1})
Reference time	1963.9	1983.8	2006.9

Parameter	Source specific			
	Cas A	Cyg A	Tau A	Vir A
$\log F_0$	Segment specific	$3.1861^{+0.0046}_{-0.0047}$	$2.9083^{+0.0044}_{-0.0045}$	$2.3070^{+0.0045}_{-0.0046}$
a_1	$-0.732^{+0.011}_{-0.010}$	-1.038 ± 0.011	$-0.226^{+0.014}_{-0.013}$	-0.876 ± 0.017
a_2	$-0.0094^{+0.0058}_{-0.0059}$	$-0.1457^{+0.0075}_{-0.0077}$	$-0.0113^{+0.0081}_{-0.0083}$	$-0.047^{+0.031}_{-0.033}$
a_3	0.0053 ± 0.0058	$0.0170^{+0.0075}_{-0.0076}$	$-0.0275^{+0.0077}_{-0.0079}$	$0.073^{+0.030}_{-0.029}$
Reference frequency	1477 MHz	1416 MHz	1569 MHz	1482 MHz
m_{ν_0}	Segment specific	–	$-0.00044^{+0.00019}_{-0.00018}$ ($0.102^{+0.042}_{-0.043}$ per cent yr^{-1})	–
Reference time	Segment specific	–	2009.05	–
$m_{\Delta \log \nu}$	0.00124 ± 0.00018 (-0.286 ± 0.042 per cent yr^{-1})	–	–	–
Reference frequency	1315 MHz	–	–	–
Reference time	2005.64	–	–	–

Parameter	Telescope specific	
	40-foot	All Others
α	$0.64^{+0.16}_{-0.15}$	–
β	$0.00757^{+0.00091}_{-0.00089}$	–
Scatter	0.0099	0




Figure 10. Our best-fitting spectral models (specifically, corresponding to the best-fitting three-segment temporal model of Fig. 9). The data are colour-coded as in Figs 8 and 9. Note, however, that we have not plotted the GBO 40-foot data here, for greater clarity at this frequency. Relative measurements (green and red) have again been multiplied by the best-fitting spectral model of the reference source, and (1) the Baars et al. (1977) Cas A measurements have been referenced to 1965 (dotted curve), (2) our GBO Cas A measurements have been referenced to 2015 (dashed curve) and (3) all Tau A measurements have been referenced to 2015 (dashed curve) using the best-fitting temporal models of Cas A and Tau A, respectively, for better visualization. The shaded regions show the 1σ , 2σ and 3σ statistical uncertainty in the best-fitting values of $\log F_0$, a_1 , a_2 , a_3 and, when applicable, m_{ν_0} and $m_{\Delta \log \nu}$. The best-fitting sample scatter of all but the 40-foot measurements is negligible.

of our best-fitting spectral and temporal models, at any frequency and time, by taking the square root of the following function:

$$\begin{aligned} \sigma_{\log F_m}^2(\nu, t) \approx & \sigma_{\log F_0}^2 + \sigma_{a_1}^2 \left[\log \left(\frac{\nu}{\nu_{\text{source}}} \right) \right]^2 \\ & + \sigma_{a_2}^2 \left[\log \left(\frac{\nu}{\nu_{\text{source}}} \right) \right]^4 + \sigma_{a_3}^2 \left[\log \left(\frac{\nu}{\nu_{\text{source}}} \right) \right]^6 \\ & + \sigma_{m_{\nu_0}}^2 (t - t_{\text{segment}})^2 \\ & + \sigma_{m_{\Delta \log \nu}}^2 \left[\log \left(\frac{\nu}{\nu_0} \right) \right]^2 (t - t_{\text{source}})^2 + \sigma_{\text{scatter}}^2, \end{aligned} \quad (15)$$

where σ_X is the positive or negative 1σ statistical uncertainty, respectively, in the best-fitting value of parameter X ; the two temporal terms are only for Cas A and Tau A, and in the case of Tau A, the reference time $t_{\text{segment}} = t_{\text{source}}$; and for all four sources, the best-fitting scatter $\sigma_{\text{scatter}} = 0$ in all of these fits (e.g. Table 5).

6 CONCLUSIONS

The fading of Cas A appears to be more complicated than originally measured, with the remnant fading quickly between the late 1940s and late 1960s/early 1970s (0.886 ± 0.021 per cent yr^{-1} in the L band), then more slowly until the mid- to late-1990s




Figure 11. Difference between the original, best-fitting spectral models of Baars et al. (1977) and our updated and improved best-fitting spectral models (Fig. 10). The Baars et al. (1977) Cas A and Tau A models have been referenced to 1965 (dotted) and 2015 (dashed), respectively. Our Cas A and Tau A models have been referenced to 2015. Clearly, Baars et al. (1977) overpredicted the fading of Cas A, and appear to have also underpredicted (specifically, set to zero) the fading of Tau A. But even the 1965 versions of these curves differ by up to ≈ 6 per cent to ≈ 11 per cent in the couple- to many-GHz range of some of these spectra (Cyg A and Tau A in particular).

($0.516^{+0.056}_{-0.054}$ per cent yr^{-1} in the L band), and then quickly again until at least the late 2010s ($0.798^{+0.058}_{-0.050}$ per cent yr^{-1} in the L band; Section 5.2).²⁰ Whether Cas A will continue to fade at this rate, or over longer time-scales will more closely adhere to the average that we measured between the late 1950s and late 2010s (0.654 ± 0.019 per cent yr^{-1} in the L band; Section 5.1) is unknown. Consequently, care should be taken if calibrating against Cas A, or Tau

A for that matter, although Tau A is fading much more slowly, and consequently should remain reasonable as a calibrator longer.

Furthermore, we have identified potential weaknesses in the absolute spectral calibrations provided by Baars et al. (1977), particularly for Cyg A and Tau A in the couple- to many-GHz range, not due to measurement errors, but due to, perhaps, statistical and computing limitations at the time (Section 5.2). We have updated and improved these calibrations, for Cas A, Cyg A, Tau A and Vir A, and have provided a better way to characterize the uncertainty in these calibrations, as a function of frequency, and, in the case of Cas A and Tau A, as a function of time.

This work can be improved upon in the future with additional measurements of Cas A at other frequencies, to see if the frequency

²⁰ It should be noted that Cas A also appears to have flared, at the $\approx 4\sigma$ confidence level, for a brief time in the 1970s, but only at low frequencies (38 MHz; Erickson & Perley 1975; Read 1977a,b).




Figure 12. The same as Fig. 8, but for Tau A. The fading of Tau A appears to be measurable over this time-scale.

dependence of its fading rate is also changing, and with additional measurements of Tau A, both at the L band and at other frequencies, to measure the frequency dependence of its fading rate.

ACKNOWLEDGEMENTS

We gratefully acknowledge the support of the National Science Foundation, through the following programmes and awards: ESP 0943305, MRI-R² 0959447, AAG 1009052, 1211782 and 1517030, ISE 1223235, HBCU-UP 1238809, TUES 1245383 and STEM+C 1640131. We are also appreciative to have been supported by the Mt Cuba Astronomical Foundation, the Robert Martin Ayers Sciences

Fund, and the North Carolina Space Grant Consortium. We also thank the referee, and the editor, for their helpful comments and suggestions. Finally, we thank the approximately 400 students and educators who have participated in *Educational Research in Radio Astronomy (ERIRA)* at Green Bank Observatory since 1992, many of whom helped to collect the 40-foot data that we present in this paper. We also thank Green Bank Observatory for hosting our group for all of these years.

REFERENCES

- Baars J. W. M., Hartsuijker A. P., 1972, *A&A*, 17, 172
 Baars J. W. M., Mezger P. G., Wendker H., 1965, *ApJ*, 142, 122
 Baars J. W. M., Genzel R., Pauliny-Toth I. I. K., Witzel A., 1977, *A&A*, 61, 99
 Dent W. A., Aller H. D., Olsen E. T., 1974, *ApJ*, 188, L11
 Erickson W. C., Perley R. A., 1975, *ApJ*, 200, L83
 Janssen M. A., Golden L. M., Welch W. J., 1974, *A&A*, 33, 373
 Kellermann K. I., Downes A. J. B., Pauliny-Toth I. I. K., Preuss E., Shaffer D. B., Witzel A., 1981, *A&A*, 97, L1
 Maples M. P. et al., 2017, *ApJS*, in press
 Perley R. A., Butler B. J., 2016, *ApJS*, preprint ([arXiv:1609.05940](https://arxiv.org/abs/1609.05940))
 Read P. L., 1977a, *MNRAS*, 178, 259
 Read P. L., 1977b, *MNRAS*, 181, 63P
 Reichart D. E., 2001, *ApJ*, 553, 235
 Reichart D. E., Stephens A. W., 2000, *ApJ*, 537, 904
 Shklovskii I. S., 1960, *SvA*, 4, 243
 Trotter A. S., 2011, PhD dissertation, University of North Carolina at Chapel Hill
 Vinyaikin E. N., 2006, *Astron. Rep.*, 50, 143
 Vinyaikin E. N., 2007a, *Astron. Rep.*, 51, 87
 Vinyaikin E. N., 2007b, *Astron. Rep.*, 51, 570
 Vinyaikin E. N., 2014, *Astron. Rep.*, 58, 626

This paper has been typeset from a \LaTeX file prepared by the author.