CAMBRIDGE
UNIVERSITY PRESS

Research Article

Modelling repetition in zDM: A single population of repeating fast radio bursts can explain CHIME data

C.W. James

International Centre for Radio Astronomy Research, Curtin University, Bentley, WA, Australia

Abstract

Regardless of whether or not all fast radio bursts (FRBs) repeat, those that do form a population with a distribution of rates. This work considers a power-law model of this population, with rate distribution $\Phi_r \sim R^{\gamma_r}$ between $R_{\min}$ and $R_{\max}$. The zDM code is used to model the probability of detecting this population as either apparently once-off or repeat events as a function of redshift, z, and dispersion measure, DM. I demonstrate that in the nearby Universe, repeating sources can contribute significantly to the total burst rate. This causes an apparent deficit in the total number of observed sources (once-off and repeaters) relative to the distant Universe that will cause a bias in FRB population models. Thus instruments with long exposure times should explicitly take repetition into account when fitting the FRB population. I then fit data from The Canadian Hydrogen Intensity Mapping Experiment (CHIME). The relative number of repeat and apparently once-off FRBs, and their DM, declination, and burst rate distributions, can be well explained by 50–100% of CHIME single FRBs being due to repeaters, with $R_{\max} > 0.75$ d^{-1} above 10^{39} erg, and $\gamma_r = -2.2^{+0.6}_{-0.8}$. This result is surprisingly consistent with follow-up studies of FRBs detected by the Australian Square Kilometre Array Pathfinder (ASKAP). Thus the evidence suggests that CHIME and ASKAP view the same repeating FRB population, which is responsible not just for repeating FRBs, but the majority of apparently once-off bursts. For greater quantitative accuracy, non-Poissonian arrival times, second-order effects in the CHIME response, and a simultaneous fit to the total FRB population parameters, should be treated in more detail in future studies.

Keywords: Radio transient sources (2008); astronomy data modelling (1859)

1. Introduction

Of the many mysteries surrounding fast radio bursts (FRBs; millisecond duration radio signals arising from cosmological distances Lorimer et al. 2007; Thornton et al. 2013) the question of whether or not they all repeat remains one of the greatest (e.g. Caleb et al. 2019). Since the discovery of the first repeating FRB, FRB 20121102A (Spitler et al. 2016), its localisation to a dwarf galaxy and association with a persistent radio source (PRS; Chatterjee et al. 2017), combined with its complex time–frequency burst structure (Hessels et al. 2019), have distinguished it from populations of apparently once-off FRBs, which arise from a plethora of galaxy types (Bhandari et al. 2020; Heintz et al. 2020; Bhandari et al. 2022; Gordon et al. 2023), and are more likely to exhibit broadband, single-component morphologies (Pleunis et al. 2021). Further discoveries of repeating FRBs have somewhat muddled this simple dichotomy however—while FRB 20190520B is similar in both its bursts and its host galaxy to FRB 20121102A (Niu et al. 2022), FRB 20180916B is located in a large spiral galaxy, close to—but offset from—a star-forming region (Marcote et al. 2020), while FRB 20200120E arises from a globular cluster (Bhardwaj et al. 2021b; Kirsten et al. 2022), which as is typical for such clusters, has no apparent star-forming activity.

Models of the FRB population paint a similarly ambiguous picture. Very early FRB data was shown to be consistent with all FRBs being similar to FRB 20121102A (Lu & Kumar 2016). As more data became available, Caleb et al. (2019) was able to rule out that all FRBs could repeat as rapidly as FRB 20121102A, while James (2019) showed that the number of strong repeaters similar to FRB 20121102A must be at most 27 Gpc^{-3} with 90% confidence. Strong limits on the repetition rate of individual FRBs now demonstrate that at least some repeat at most very rarely (James et al. 2020b; Lin et al. 2023). More recently, Gardenier et al. (2021) used FRBPOPPY (Gardenier et al. 2019) to match a population of repeating FRBs to the dispersion measure (DM) distribution from the Canadian Hydrogen Intensity Mapping Experiment (CHIME; CHIME/FRB Collaboration et al. 2018). They confirm that for a given population of repeating FRBs, those in the nearby Universe will preferentially be detected as repeaters, while those in the distant Universe will more likely be viewed as once-off bursts. This effect is qualitatively present in the CHIME data: repeaters at low DM have a higher rate (as noted by D. Good at the FRB 2021 online conference, using data from CHIME/FRB Collaboration et al. 2021), and the mean DM of repeaters is lower than that of once-off bursts (CHIME/FRB Collaboration et al. 2023). However, no quantitative analysis has attempted to fit this effect, which must be present in the data. Thus the degree to which this is evidence for repeating FRBs and intrinsically once-off bursts being distinct populations remains unknown.

Another result from FRB population analysis is the all-burst luminosity function, which is usually modelled as either a

power-law or a Schechter function (e.g. Luo et al. 2020a; James et al. 2022b; Shin et al. 2023). This is generally fitted to have a comparable slope to original measurements of the FRB 20121102A burst energy distribution (Law et al. 2017), indicating that the all-burst spectrum could be built up as the sum of repeating FRBs. More detailed measurements of burst energy spectra of repeaters has shown behaviour much more complex than a single power-law however (Li et al. 2021; Hewitt et al. 2022; Zhang et al. 2023, 2022), making comparisons difficult, though there are indications that the very high-energy tail might remain consistent with results from population models (Li et al. 2021; Hewitt et al. 2022; Zhang et al. 2023, 2022), Again, the evidence for or against an FRB population dominated by intrinsically repeating FRBs or once-off events is ambiguous.

Results from both host galaxy and population modelling therefore indicate that FRB progenitors either come from multiple populations, or if they do all repeat, must come from a broad distribution of repetition rates. The different morphologies observed for repeating and once-off bursts is consistent with both pictures (Pleunis et al. 2021); as is the observation of PRSs associated with some of the brightest and most rapid repeating objects (Chatterjee et al. 2017; Niu et al. 2022), since it is not implausible PRSs dissipate and pulse morphologies change as the progenitor ages. Even if some FRBs arise from intrinsically cataclysmic events, such as black hole formation following a binary neutron star merger (Zhang 2014; Falcke & Rezzolla 2014; Moroianu et al. 2023), it should be expected that the portion of the population that does intrinsically repeat has a broad distribution of properties. And from an observational or modelling perspective, there is no practical difference between FRBs which repeat very rarely, and those that are intrinsically once-off. This motivates studies which attempt to fit the properties of an intrinsically repeating FRB population to observational data.

FRB observations however give a biased picture of the true underlying FRB population. This is particularly the case with repeaters, which make better targets for host galaxy identification—of the 492 distinct FRBs published by CHIME/FRB Collaboration et al. (2021) four repeaters have been localised to their host galaxies (Marcote et al. 2020; Bhardwaj et al. 2021a; Michilli et al. 2022), with a further single burst in that catalogue since identified as a repeater and localised to its host (Ibik et al. 2023). However, only a single once-off CHIME FRB has a tentative host galaxy identification (Panther et al. 2022). Furthermore, once trends in FRB behaviour become identified, these result in preferential targeting with follow-up observations, for example, the identification of FRB 20180301A as a repeater from its time–frequency structure (Price et al. 2019; Luo et al. 2020b), or the targeting of FRB 20121102A during its identified activity phase (Rajwade et al. 2020; Li et al. 2021). Thus once a trend is identified, follow-up observations will naturally reinforce them. It is important to note that CHIME detects repeat bursts from FRBs in an almost unbiased manner (CHIME/FRB Collaboration et al. 2019; Fonseca et al. 2020)—the only other repeating FRB to be identified as such is FRB 20190520B (Niu et al. 2022), while all other identifications have been through targeted follow-up observations.

Observational biases thus limit the use of repeating FRBs in FRB population analyses, which model their cosmological source evolution, burst energy distribution, spectral properties, and local/cosmological/host DM contributions (e.g. Luo et al. 2020a; James et al. 2022b). When they are included, only the first burst of a repeating FRB tends to be used (e.g. Shin et al. 2023), in order to make the analysis as insensitive to their repeating nature as possible.

We are only aware of one work which fits the intrinsic properties of the repeating FRB population. James et al. (2020a) use results from follow-up observations of CRAFT FRBs with the Murriyang (Parkes) and Robert C. Byrd Green Bank Telescope (GBT) (James et al. 2020b), and assuming a power-law distribution of FRB repetition rates R, $dN(R)/dR \propto R^{\gamma_r}$, the authors placed limits on the power-law index $\gamma_r < -2$, and the maximum repetition rate $R_{\max}$. The data fit however incorporated only one observed repeater, FRB 20171019A (Shannon et al. 2018; Kumar et al. 2019), and 19 non-repeaters, in particular FRB 20171020A (Mahony et al. 2018), the proximity of which rules out repetition rates greater than 0.011 d^{-1} above 10^{39} erg (James et al. 2020b; Lee-Waddell et al. 2023). Another work, Law, Connor, & Aggarwal (2022), uses CHIME data to model the mean apparent (i.e. not intrinsic) repetition rate, finding of 25–440 yr^{-1}. The lower limit was calculated by assuming that only sources observed to repeat were true repeaters, while the upper limit assumed singly observed FRBs were also due to repeaters. Interestingly, the authors note the potential of using the population of PRS to constrain the repeating FRB population, and vice versa.

The aim of this paper is to incorporate a model for repeating FRBs into the framework of the zDM code (James et al. 2021, 2022b), as described in §2. Example FRB populations are then used to estimate the biasing effects of FRB repetition on the z–DM distribution of the FRB population in Section 3. In Section 4, a model of CHIME is described, and in Section 5, it's shown how the model compares to CHIME FRBs from Catalogue 1 (CHIME/FRB Collaboration et al. 2021). In Section 6, a fit is performed of FRB repetition parameters to CHIME data, and use the results to make predictions for future experiments in Section 7. Potential systematic effects are discussed in Section 8. The results are compared to those from FRB follow-up observations with the Australian Square Kilometre Array Pathfinder (ASKAP), and estimates of the number density of PRS, in Section 9; findings are summarised in Section 10.

Throughout, I use the nomenclature of 'burst' to refer to a single FRB (which nonetheless may have multiple components, typically on ms or sub-ms scales); while 'FRB' or 'progenitor' refers to the progenitor objects, such that a single repeating FRB emits multiple bursts.

A standard Planck cosmology (Planck Collaboration et al. 2020) with $H_0 = 67.4 \, \mathrm{km \, s^{-1} \, Mpc^{-1}}$ is used throughout. The Kolmogorov–Smirnov ('KS'; Kolmogorov 1933; Smirnov 1948) test is used to assess consistency between observed and expected distributions. This is a somewhat arbitrary choice, with no strong preference compared to, for example, the Anderson-Darling test ('AD'; Anderson & Darling 1954), though the extra sensitivity of the AD test to the tails of the distributions analysed herein may not be desirable.

2. Modelling repeating FRBs

I begin by characterising repeating FRBs solely by their time-averaged burst rate R, which is defined as their intrinsic rest-frame rate of producing bursts above 10^{39} erg. All FRBs are treated as repeating according to a Poissonian distribution (see Section 8.1 for a discussion of this assumption), such that the probability of them producing N observed bursts given an expectation $\lambda = R_{\mathrm{obs}} T$ (for some time interval T) is

$$P(N) = \frac{\lambda^N \exp{(-\lambda)}}{N!}. \tag{1}$$

It is assumed that each repeating FRB has an identical energy distribution, with cumulative probability described by an upper incomplete Gamma (Schechter) function. Thus the observable rate R_{obs} above some energy E_{th} (itself a function of z, DM, and position in the telescope beam) scales as

$$R_{obs}(E_{th}) = \frac{R}{1+z} \frac{\int_{E_{th}}^{\infty} (E/E_{max})^{\gamma} e^{-E/E_{max}} dE}{\int_{10^{39}}^{\infty} (E/E_{max})^{\gamma} e^{-E/E_{max}} dE}, \tag{2}$$

where γ is the cumulative power-law index, and E_{max} some cut-off energy.

The population density of repeating FRBs, Φ_r (progenitors Mpc^{-3}), is modelled with repetition rates above some rate R via a power-law distribution of their intrinsic repetition rates R,

$$\Phi_r(R) = \begin{array}{ll} C_r & (R < R_{min}) \quad (3) \\ 0 & (R > R_{max}) \\ C_r \frac{\left(\frac{R}{R_{min}}\right)^{\gamma_r+1} - \left(\frac{R_{max}}{R_{min}}\right)^{\gamma_r+1}}{1 - \left(\frac{R_{max}}{R_{min}}\right)^{\gamma_r+1}} & \text{otherwise,} \end{array}$$

between minimum and maximum rates R_{min} and R_{max} respectively, with differential power-law index γ_r, and constant population density C_r. The differential rate is more commonly used, which is defined as

$$\frac{d\Phi_r(R)}{dR} = C_r' R^{\gamma_r} \tag{4}$$

$$C_r' = \frac{(\gamma_r + 1)C_r}{R_{min}^{\gamma_r+1} - R_{max}^{\gamma_r+1}}.$$

In this model, the total burst density above 10^{39} erg, C (bursts Mpc^{-3} yr^{-1}), is related to the population of repeaters via

$$C = \int_{R_{min}}^{R_{max}} R C_r' R^{\gamma_r} dR$$
$$= \frac{C_r'}{\gamma_r + 2} \left[R_{max}^{\gamma_r+2} - R_{min}^{\gamma_r+2} \right]. \tag{5}$$

As per previous work, both C and C_r are defined as the number of bursts and FRB progenitors respectively observed at 1.3 GHz at $z = 0$. Both are treated as evolving with the star-formation rate, such that C_r evolves as

$$C_r(z) = C_r \left(\frac{\mathrm{SFR}(z)}{\mathrm{SFR}(z=0)} \right)^{n_{sfr}}. \tag{6}$$

n_{sfr} is used to scale smoothly between a non-evolving population ($n_{sfr} = 0$), one evolving with the star-formation rate ($n_{sfr} = 1$), and a population with a stronger peak at high redshift, such as AGN activity ($n_{sfr} > 1$). In the case of C, an additional factor of $(1 + z)^{-1}$ is included on the right-hand side of (6), due to the time dilation effect. However, the number of FRB progenitors is unaffected by time dilation; rather, the time dilation effect is instead included when scaling between the rate seen by an observer and the intrinsic rate.

2.1. Implementation in the zDM code

The zDM code was originally developed for the analysis of James et al. (2022b), and has since been extended as per James et al. (2022c) and Baptista et al. (2023). It calculates the expected number of FRBs from an FRB survey as a function of their redshift

z, dispersion measure DM, and relative fluence F compared to threshold fluence F_{th}. The key parameters of the code, and their current best-fit values, are given in Table 1.

For this analysis, only the measured values of z and DM for an FRB are considered—unlike previous works, where signal-to-noise ratio (SNR) is also used. I add an additional observable: whether or not an FRB is observed as a repeater. This is treated as a boolean, that is, the number of observed repeats is not used. Furthermore, only repetition information obtained through initial blind surveys is considered, that is, this analysis is not suited to modelling FRBs determined to repeat or not through targeted follow-up observations. While CHIME remains the only FRB instrument to observe repeat bursts by this definition, the non-observation of repeaters in other blind surveys can still be included.

The total number of repeaters in any given z–DM bin is given by:

$$N(z, \mathrm{DM}) = C_r(z) \frac{dV}{dzd\Omega} \Delta z f_z(\mathrm{DM}) \Delta\mathrm{DM}\Delta\Omega, \tag{7}$$

where $\Delta\Omega$ is the solid angle observed by the beam, Δz and ΔDM are the bin sizes of the z–DM grid, $dV(dzd\Omega)^{-1}$ is the size of the cosmological volume element, and $f_z(\mathrm{DM})$ is the distribution of DMs of FRBs at that redshift (itself a function of cosmological and host galaxy contributions, as given in James et al. 2022b).

The distribution of intrinsic FRB rates within that volume element is given by (4), which for any given threshold energy E_{th}, produces a distribution in R_{obs} as per (2). Thus the expected number of once-off FRBs from that volume becomes

$$\langle N_1(z, \mathrm{DM}) \rangle = \frac{dV}{dzd\Omega} \Delta z f_z(\mathrm{DM}) \Delta\mathrm{DM}\Delta\Omega \tag{8}$$
$$\cdot \int_{R_{min}}^{R_{max}} R_{obs} T_{obs} e^{-R_{obs}T_{obs}} \frac{d\Phi_r(R)}{dR} \frac{dR}{dR_{obs}} dR_{obs}.$$

Here, I have left off the dependence of R on R_{obs}, expressed through (2). Similarly,

$$\langle N_0(z, \mathrm{DM}) \rangle = \frac{dV}{dzd\Omega} \Delta z f_z(\mathrm{DM}) \Delta\mathrm{DM}\Delta\Omega \tag{9}$$
$$\cdot \int_{R_{min}}^{R_{max}} e^{-R_{obs}T_{obs}} \frac{d\Phi_r(R)}{dR} \frac{dR}{dR_{obs}} dR_{obs}$$

calculates the number of progenitors for which no bursts are detected. Thus, the expected number of repeating FRBs can be deduced as

$$\langle N_{reps} \rangle = N - \langle N_1 \rangle - \langle N_0 \rangle. \tag{10}$$

This therefore allows a likelihood to be assigned to both once-off and repeating FRBs as a function of redshift z and dispersion measure DM.

3. The effects of repetition on FRB redshift distributions

To qualitatively illustrate how the presence of FRB repetition affects the measured FRB redshift distribution, I perform example calculations for the CRAFT incoherent sum (ICS; Bannister et al. 2019) survey parameters as described in James et al. (2022b). Since this observing mode is mostly commensal, total time per field T_f to the end of 2022 has been at most $\sim$37 d (Shannon et al. in preparation). For illustrative purposes however, T_f is varied between 10, 100, and 1000 d, and a central frequency of 1.3 GHz is used.

Table 1. FRB population parameter sets used in this work. Shown are the best-fit parameter sets from Shin et al. (2023) and James et al. (2022c), and a set of 12 parameter sets from James et al. (2022c) when each parameter takes the minimum/maximum value within its 90% confidence interval. The *p*-values from KS tests against the observed rate of CHIME single bursts is given as p_{KS} for different sets of FRB population parameters.

Scenario		E_{max}[a] $\log_{10}$ [erg]	α[b]	γ[c]	n_{sfr}[d]	$\log_{10} \mu_{host}$[e] pc cm^{-3}	σ_{host}[f]	$\log_{10} C$[g] Gpc^{-3}yr^{-1}	p_{DM}^{dist}[h]	p_{DM}^{strong}[i]
Shin et al. (2023)		41.38	-1.39	-1.3	0.96	1.93	0.41	4.99	0.0028	10^{-12}
James et al. (2022c)		41.26	-1.0	-0.95	1.13	2.27	0.55	4.47	10^{-22}	10^{-54}
E_{max}	min	41.0	-1.0	-0.7	1.0	2.3	0.6	4.63	10^{-21}	10^{-38}
	max	41.8	-1.0	-1.1	1.25	2.2	0.5	4.57	10^{-80}	10^{-100}
α	min	41.3	-1.91	-0.9	0.75	2.2	0.6	5.00	0.56	10^{-23}
	max	41.3	0.24	-0.9	0.75	2.2	0.6	4.94	10^{-7}	10^{-26}
γ	min	41.8	-1.5	-1.18	1.75	2.2	0.5	4.35	10^{-125}	10^{-138}
	max	41.2	-1.0	-0.66	1.0	2.2	0.6	4.40	10^{-40}	10^{-61}
n_{sfr}	min	41.3	0.0	-0.8	0.49	2.2	0.6	4.87	10^{-8}	10^{-27}
	max	41.4	-1.9	-1.0	1.91	2.2	0.6	4.13	10^{-126}	10^{-135}
μ_{host}	min	41.5	-1.0	-0.9	1.25	1.98	0.6	4.40	10^{-60}	10^{-78}
	max	41.3	-1.0	-0.9	1.0	2.45	0.6	4.69	10^{-53}	10^{-74}
σ_{host}	min	41.5	-1.0	-1.0	1.25	2.2	0.43	4.56	10^{-54}	10^{-74}
	max	41.4	-1.0	-0.9	1.25	2.2	0.82	4.50	10^{-65}	10^{-83}

[a] Downturn ($\sim$ 'maximum') FRB energy, assuming a 1 GHz rest-frame emission bandwidth.
[b] Frequency scaling of rate: $C \propto \nu^{\alpha}$.
[c] Cumulative power-law index of the FRB luminosity function.
[d] Scaling of FRB density with star-formation: $C \propto \mathrm{SFR}^{n_{sfr}}$.
[e] Log-mean of host galaxy DM contribution.
[f] Log-standard deviation of host galaxy DM contribution.
[g] Absolute total burst rate above 10^{39} erg at 1.3 GHz at $z = 0$.
[h] *p*-value from the KS test assuming the distributed repeaters scenario.
[i] *p*-value from the KS test assuming the strong repeaters scenario.

For repeating FRB parameters, I consider two scenarios. The first is strong repeaters only, setting $R_{max} = R_{min} = 4\,\mathrm{d}^{-1}$ (thus making γ_r irrelevant), which is approximately the time-averaged rate observed above 10^{39} erg for FRB 20121102A by Li et al. (2021), divided by four to account for the off part of the activity cycle (Rajwade et al. 2020); this is also consistent with the time-averaged rate observed by Law et al. (2017). The second is a more realistic scenario, with a distribution of repeaters with $R_{max} = 10\,\mathrm{d}^{-1}$ (i.e. slightly above FRB 20121102A), $R_{min} = 10^{-3}\,\mathrm{d}$ (slightly below FRB 20171020A), and an index of $\gamma_r = -2.2$, which is consistent with the results of James et al. (2020a).

For the remaining properties of the FRB population, note that most previous models have analysed only once-off FRBs, or alternatively, included only the first burst of repeating FRBs. Several authors have derived values constraining (2), typically obtaining a cumulative fluence index $\gamma \sim -1$, $E_{max} \sim 10^{41-42}$ erg, with $\mathcal{O} \sim 10^5$ FRBs Mpc^{-3} yr^{-1} in the local Universe, and finding no evidence for a minimum energy (Luo et al. 2020a; James et al. 2022a,c; Shin et al. 2023). For now, the best-fit values from James et al. (2022c)—given in Table 1—are used, and other possibilities are considered in Section 5.
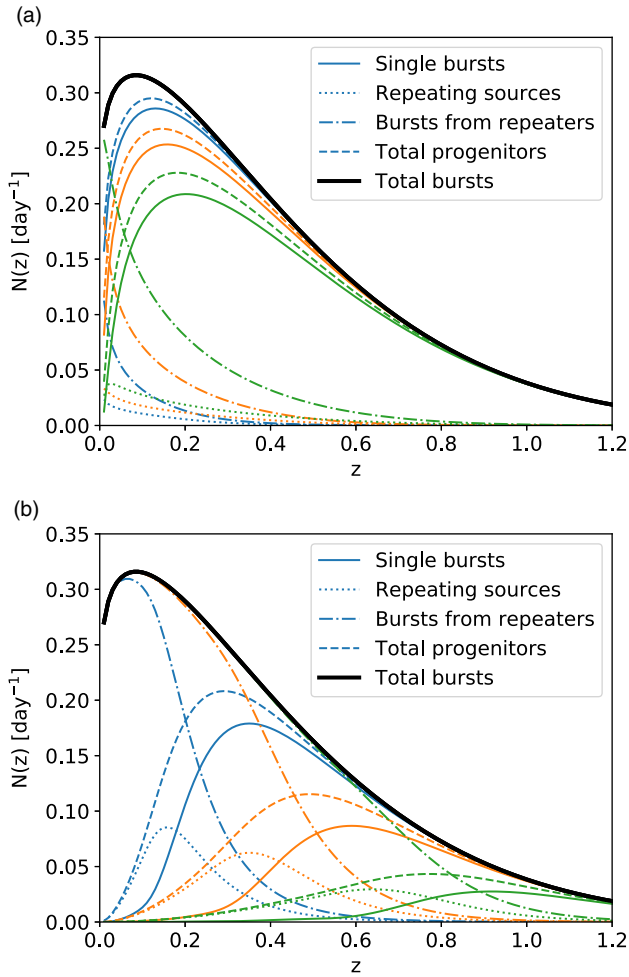
The results for strong and distributed repeaters are shown in Fig. 1, in units of expected bursts per day. The total number of expected detected bursts is identical in all scenarios. However, all other measurable quantities depend on T_f and the nature of the repeating FRB population. In all scenarios, intrinsically repeating FRBs (dotted lines) are more likely to be detected as such in the nearby Universe, which is the expected result. Necessarily, this decreases the expected number of once-off bursts in the local Universe (thin solid lines), since each FRB detected to repeat

removes its first detected burst from the once-off distribution. The standard method of including repeating FRBs in population models, that is, counting them only once using the first measured burst and adding apparently once-off bursts, is shown as 'Total progenitors' (dashed lines). This method still results in a deficit of the modelled population in the local Universe compared to the total burst population. Only when including all bursts from measured repeaters (dot-dashed lines) as well as single bursts will the bias against low-redshift FRBs disappear.

The effect of T_f is evident from Fig. 1. As T_f increases (shown via changing colour), more intrinsically repeating FRBs in the nearby Universe are detected with multiple bursts. Thus, all distributions are pushed to higher z, with repeating FRBs taking an increasingly large fraction of the measured population. For the distributed repeaters scenario, this effect is small, and even after 1000 d on a single field, by far the majority of FRBs are detected as once-off bursts, though the redshift peak for single bursts has shifted from 0.19 to 0.26 (with the unbiased peak for all bursts being at 0.15). For strong repeaters however, this effect is very important, with single bursts peaking at $z = 0.41$, 0.67, and 0.95 respectively—and even repeating sources peaking at higher redshifts than the true underlying total burst population. This latter effect is due to the small number of repeaters at low redshift being expected to produce a very large number of bursts, whereas for the distributed repeaters scenario, repetition remains dominated by rarely repeating objects at low redshift.

Repeating FRBs also introduce significant cosmic variance into FRB surveys. Since the volume of the nearby Universe is small, there is a small chance for a strongly repeating FRB to be located in that volume—however, if there is, many bursts will be detected
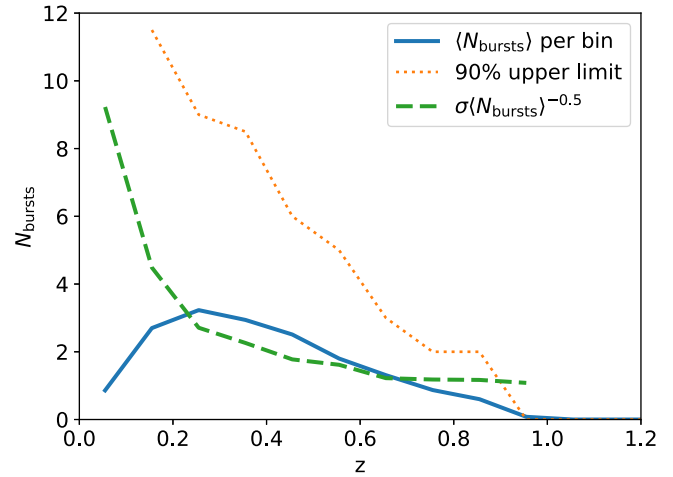
Figure 2. Burst statistics from 100 Monte Carlo simulations of an ASKAP/ICS 100 d pointing. Shown are the expected number of bursts $\langle N_{bursts} \rangle$, 90% upper limit, and standard deviation normalised by the square root of $\langle N_{bursts} \rangle$, histogrammed as a function of redshift.

Figure 1. Effects of the repeating FRB population on the redshift distributon of FRBs expected from the CRAFT/ICS survey. Top: repeating FRBs with a broad distribution of rates; bottom: repeating FRBs identical to FRB 20121102A; observation times T_f, with lines appearing from left to right, are 10 d (red), 100 d (green), 1000 d (blue).

and accuracy. I therefore proceed with the approach of modelling single bursts, and the number of repeaters, in this work, and do not directly model the number of bursts per repeater (though the distribution is fit via Monte Carlo in Section 6.3). This results in zero bias, and only a small reduction in accuracy.

4. Modelling CHIME Catalogue 1 data

The CHIME/FRB experiment is described by CHIME/FRB Collaboration et al. (2018). Not only does CHIME have the largest published sample of both repeating and non-repeating FRBs (CHIME/FRB Collaboration et al. 2021), but since CHIME is not re-pointed to target specific sources, CHIME FRBs are detected in an unbiased manner. The downside however is that CHIME's angular resolution is relatively poor (approximately 1'–10'; Michilli et al. 2021), so that the vast majority of FRBs are not localised with sufficient precision to identify their host galaxy. This makes CHIME data ideal for modelling the repeater versus non-repeater fraction, even if it is difficult to use it to model the FRB distribution in zDM space. The CHIME experiment is modelled as below.

4.1. Data sample

Data is taken from CHIME/FRB Collaboration et al. (2021) (hereafter, 'Cat1'), consisting of 536 FRBs. Nominally, this is divided into 474 once-off bursts, and 62 bursts from 18 repeaters. However, two repeaters—FRB 20190417A, and FRB 20181119D from FRB 20121102A—are only identified as such from observations external to Cat1. Thus it is more proper to say that Cat1 identifies 16 repeaters and 476 non-repeaters, which is the statistic used in this work.

CHIME have also recently published a search for repeating FRBs in an updated dataset, announcing the discovery of 25 new sources based on coincidences in DM–localisation space (CHIME/FRB Collaboration et al. 2023). I denote this the 'Gold25' sample, and discuss this in detail—including why it is not used for fitting—in Section 8.3.

from it due to its proximity. To illustrate this effect, a Monte Carlo sample of 100 instances of a 100-d ASKAP/ICS observation is generated, in a Universe consisting only of FRBs as strong as FRB 20121102A. The mean number of bursts detected, $\langle N_{bursts} \rangle$, is shown as a function of redshift in Fig. 2. As expected, it follows the shape of Fig. 1. For a Poisson distribution, the standard deviation σ would be expected to scale with $\langle N_{busts} \rangle^{0.5}$—therefore, I normalise σ by this value. At high redshift, it tends towards this expected value. However, this variance increases rapidly at low redshift. The number of bursts in the 90% upper limit over all 100 simulations increases even more rapidly. The most extreme example of this variance is that in 99 of the 100 simulations, no bursts were simulated in the redshift interval $0 \leq z \leq 0.1$—however, in a single instance, 86 bursts from a single repeating FRB were simulated.

This presents a problem for population modelling. The requirement to include all bursts from repeating FRBs to avoid a bias will result in large stochastic fluctuations in population statistics, dominated by the small population of strongly repeating FRBs. The choice to do so, or not, represents a trade-off between bias

The analysis in Cat1 used three data quality cuts which are not implemented here, for the following reasons:

1. Cut removing events with SNR < 12. This is implemented due to human inspection falsely rejecting low-SNR events. Since this work is primarily concerned with the ratio of repeaters to non-repeaters, rather than absolute FRB numbers, such a cut is not used. Repetitions were searched for at a slightly lower (but undocumented) threshold than initial bursts (CHIME/FRB Collaboration et al. 2021,; see Section 8.3 for further discussion of this effect) and using the SNR ≥ 12 cut would eliminate this effect. However, implementing it would leave only 7 of 16 repeating FRBs. I believe that the resulting loss of precision will be worse than the associated systematic error.

2. Cut removing events with DM < 1.5 max(DM$_{NE2001}$, DM$_{YMW16}$). Here, such events are retained—it is true that DM$_{MW}$ is poorly known, but this also means that the effect of such a cut is equally unknown.

3. Far-sidelobe events. These will have poor localisation, and are identified using the FRB spectrum. However, this work is not concerned with localisation accuracy, but rather the number of FRBs detected, for which precise sidelobe details are not as important. Thus such a cut is not implemented.

In the following analysis, the declination δ of each FRB, its DM, and whether or not it has been observed to repeat is used.
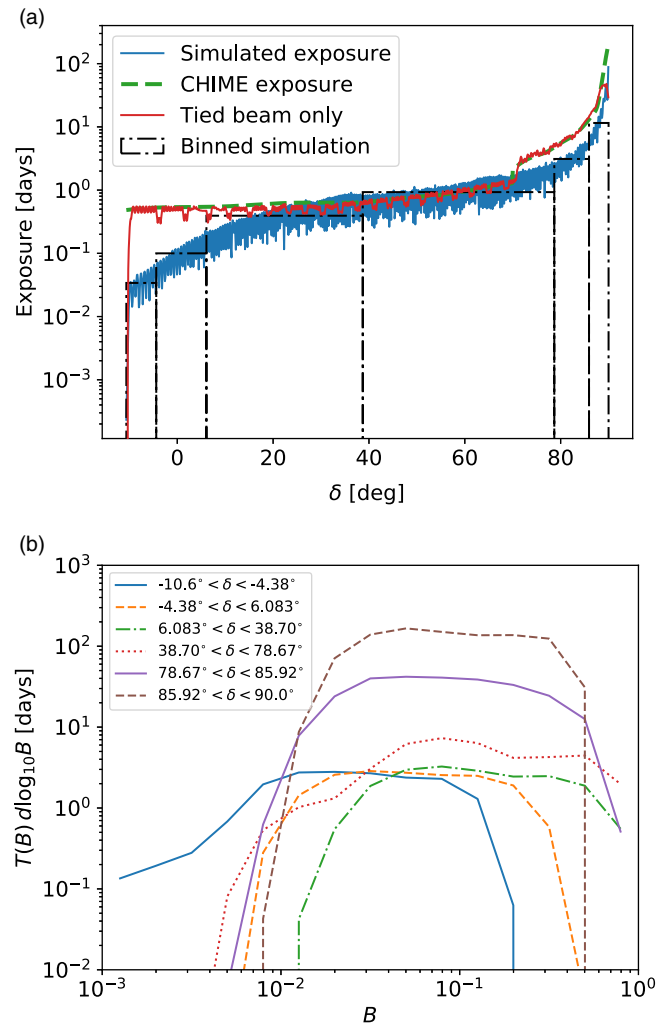
4.2. Beamshape

The CHIME beamshape is both unique and complex, with a broad primary beam with full-width at half-power of approximately 60°–120° in the N–S direction and 1.3°–2.5° in the E–W direction (CHIME/FRB Collaboration et al. 2018, 2021). Within this envelope are 1024 coherently formed beams used for FRB detection, with full-width, half-max (FWHM) of 20'–40'. I use the frequency-dependent beamshape given in CHIME/FRB Collaboration et al. (2021), averaged over XX and YY polarisations, and implemented in the GitHub library CHIME-FRB-BEAM-MODEL.[a] This is sampled at 16 frequencies in a 300 × 1000 grid in RA, DEC, taking points within 8° of the meridian. The final beamshape is taken as the envelope over all 1024 formed beams after averaging over frequency. I discuss an alternative method in Appendix A.

In James et al. (2022b), a telescope's beam pattern on sky is described via the 'inverse beamshape', $\Omega(B)$, being the solid angle of sky viewed at any given sensitivity. The total time observing that solid angle is then a simple scaling constant between rate and the number of observes FRBs. When considering the response of a transit instrument such as CHIME to repeating FRBs, the relevant metric is $T(B)$, being the time spent observing a given sky position at beam sensitivity B (relative to the nominal sensitivity at beam centre, where $B = 1$). For a source at a given declination, $T(B)$ can be calculated from the RA-dependence of the beam pattern—this is similar to the approach used by Gardenier et al. (2021).

4.2.1. Declination dependence

I calculate the declination dependence of CHIME's beamshape by calculating $T(B)$ for 1000 declinations between −11° and +90°.

Figure 3. Top: declination-dependent exposure of the CHIME experiment—simulation from this work based on the beamshape described in CHIME/FRB Collaboration et al. (2021) and scaled using a total of 220 d's observation time, 'CHIME' taken directly from CHIME/FRB Collaboration et al. (2021). Bottom: $T(\bar{B})$, calculated from the simulated beam pattern, and averaged over the indicated declination ranges.

The effective exposure, T_{eff}, is defined as the sum of times spent observing the source T_i, weighted by beam sensitivity b, assuming a Euclidean distribution of event rates:

$$T_{eff} = \sum_i T_i B^{1.5}. \qquad (11)$$

It is calculated for each declination as above, using discrete samples of source position (hence a sum, rather than an integral, in (11)), and is shown in Fig. 3. The oscillatory behaviour is due to the 256 rows of tied beams, with sources passing over beam centre having significantly more exposure than those passing between beams. In the region $\delta \gtrsim 70°$, the exposure is additionally increased by CHIME viewing sources transiting twice daily, both sides of the North Celestial Pole (NCP). Also plotted is the CHIME exposure, taken from CHIME/FRB Collaboration et al. (2021), defined as the time in which a source is within the full-width half-max (FWHM) of a tied beam at 600 MHz. This is used to normalise the total effective observation time to 311 d. by using CHIME's simulation of the 600 MHz tied beam, measuring the

time a source spends within the FWHM, and fitting to the published CHIME exposure. This total effective time is less than the 342 d duration of the catalogue, which is expected when allowing for equipment down-time. Since CHIME's total exposure is a function only of the tied beamshape, it is therefore significantly greater than T_{eff} at low declinations where the primary beam is less sensitive.

For calculation purposes, declination is divided into six regions, spanning the full $[-11°, 90°]$ range. The bounds of each region are chosen so that the change in exposure between regions is not too large (at most a factor of three), while preserving reasonable statistics in each region. The mean exposures in each region are also shown as histograms in Fig. 3.

4.3. Sensitivity

The zDM code by default uses a two-dimensional function of FRB width w and DM to calculate the effective threshold F_{th} to an FRB with those properties. An important input into this is the telescope's time- and frequency resolution used for FRB searches: in the case of CHIME, 0.983 ms, and 0.0244 MHz respectively (CHIME/FRB Collaboration et al. 2021). Integrating over a modelled distribution of intrinsic FRB widths w (James et al. 2022b) and intrinsic scattering measures τ (CHIME/FRB Collaboration et al. 2021; James et al. 2022c), this produces an efficiency function ϵ that acts to decrease measured SNR or equivalently, increase the effective detection threshold F_{th} above a nominal threshold F_0 in a DM-dependent manner. Including the effects of the beam B, F_{th} is given by

$$F_{th} = \frac{F_0}{B\epsilon(w, \tau, DM)}. \qquad (12)$$

CHIME/FRB Collaboration et al. (2021) extend this to a more complicated selection function, $P(SNR|F, DM, \tau, w, \gamma, r)$, giving SNR as a function of DM, τ, w, F, and also including spectral shape parameters γ and r. This is based on a sophisticated pulse injection system (Merryfield et al. 2022), and thus accounts for not only the dedispersion code, BONSAI, but also the full detection pipeline, including RFI rejection. The full selection function P is not published, although it can presumably be inferred from the published library of injected pulses. Rather, one-dimensional selection functions, integrated over all other variables and modelled population probability distributions in those variables, are given.

The DM selection function, $s(DM)$, gives the relative fraction of FRBs passing selection cuts as a function of DM—which is precisely what is required to model CHIME's DM distribution. To use CHIME's bias selection function, it is fit with a 4th-order polynomial as shown in Fig. 4. It is then corrected to a peak value of unity, such that $N_{obs} \leq N_{true}$, and converted to a modifier to the measured burst SNR by assuming a Euclidean relationship between event number and SNR, that is, $SNR_{bias} \sim s^{2/3}(DM)$. This is compared to the efficiency function ϵ produced by the zDM code, normalised such that the two efficiencies are equal at a DM of 1500 pc cm^{-3}.

There are two clear regimes present in Fig. 4. Above 1000 pc cm^{-3}, the CHIME SNR bias agrees well with that estimated by zDM, which is indicative of sensitivity being fundamentally limited by the time–frequency resolution of the instrument. Below 1000 pc cm^{-3}, zDM flattens to represent efficient detection, while the efficiency of CHIME decreases. This is likely due to the effects of CHIME's system for mitigating RFI (CHIME/FRB Collaboration et al. 2021; Merryfield et al. 2022).
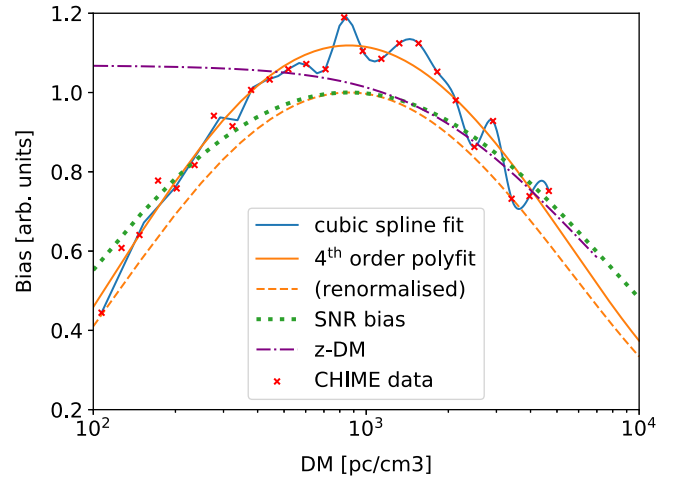
Figure 4. DM bias correction for CHIME data. Shown are values of $s(DM)$ from CHIME/FRB Collaboration et al. (2021) (red points), a cubic spline fit (blue solid line), the 4th order polynomial fit from this work (orange solid line), the renormalised fit (orange dashed line), and implied SNR bias (green dotted line).
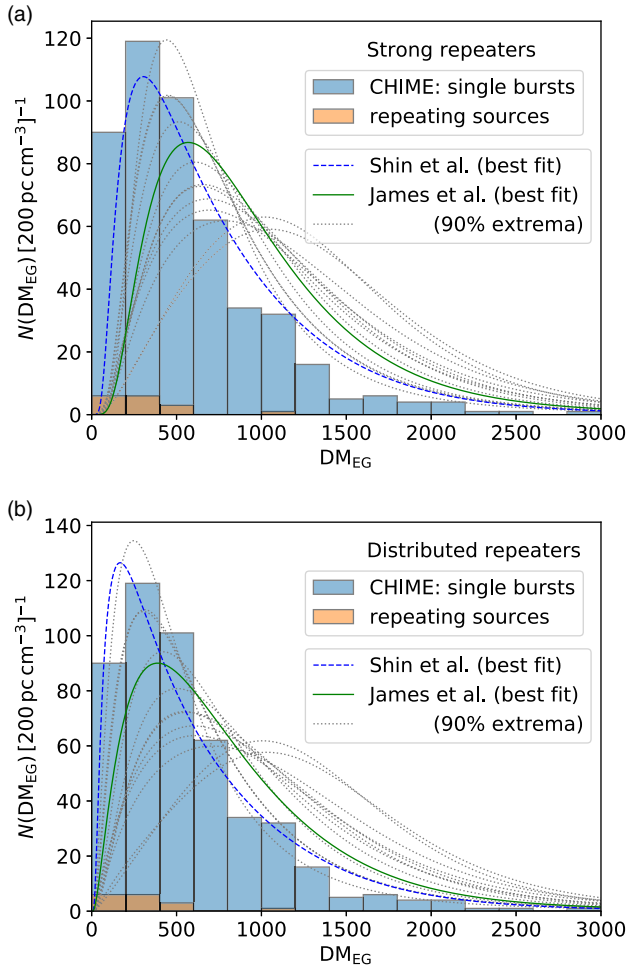
To estimate the fluence threshold at beam centre and peak DM efficiency, F_0, CHIME quotes a 95% completeness threshold of 5 Jy ms, compared to a theoretical minimum detection threshold of ~1 Jy ms (CHIME/FRB Collaboration et al. 2021). The factor of five between these thresholds is comparable to the factor of 6.6 by which the zDM efficiency had to be increased to match the CHIME efficiency in Fig. 4. Thus the fitted SNR bias of CHIME's selection function has been implemented in the zDM code, such that it acts in (12) as an efficiency factor; and $F_0 = 5$ Jy ms to a 1 ms burst is used. The value of the beam B is parameterised according to Sections 4.2 and 4.2.1, so that repeating FRBs in each declination bin are modelled as being exposed to sensitivity thresholds F_{th} for times T_{obs}.

5. Preliminary results—initial comparison

I begin by comparing the results of previous population modelling to CHIME single-burst data. While the aim of the present manuscript is to model the relative once-off and repeat burst rates by varying the repeater properties of the population, the properties of the total FRB population—luminosity function, source evolution etc.—will be correlated. Furthermore, if a good fit for the single-burst population cannot be obtained, it will be impossible to determine whether goodness of fit for repeat bursts is a function of repeating or total population parameters.

FRB population parameters from Shin et al. (2023) and James et al. (2022c), hereafter S22 and J22c respectively, are considered. In the former case, only the best-fit set is used, since the allowed ranges are very broad. In the latter case, both the best fit, and sets compatible with the 90% upper and lower limits of each parameter, are considered. To obtain these, the parameter in question is first set to its min/max value at 90% confidence, and then a search is performed over all evaluated parameter sets with a similar parameter value for the best-fitting set of other parameters. These parameter sets are listed in Table 1.

Fig. 5 shows fits to the single-burst population for both the strong repeaters (top) and distributed repeaters (bottom) scenarios, for all considered population parameters. The most striking comparison is that models generally have difficulty fitting the large

Figure 6. Critical value of repetition, R^*, as a function of the fraction F_{single} of apparently once-off bursts that are attributed to repeaters.

Figure 5. Observed rates of CHIME FRBs, showing sources observed as single and repeating, summed over declination. These are compared to estimates for the number of single bursts using the best-fit results from Shin et al. (2023) (solid, blue) and James et al. (2022c) (dashed, green), and 90% extreme (dotted, grey) values of population parameters from Table 1, and assuming a population of repeating FRBs with 'strong' (top) and 'distributed' (bottom) repetition rates. Predicted singles rates are normalised to observed singles rates.

number of low-DM FRBs observed by CHIME, especially given the strong bias against low-DM events from the CHIME selection function. Performing 1-sample KS tests (Kolmogorov 1933; Smirnov 1948) of the CHIME data against the predicted curves using SCIPY's 'stats' package, no parameter set provided a good fit to the strong repeater distribution. This is not surprising, since as discussed in Section 1, other analyses have already ruled out that all FRBs are strong repeaters. In the distributed repeaters scenario, only the minimum value of $\alpha = -1.91$ from James et al. (2022c) was compatible, with a p-value of 0.27.

The fact that the parameters of S22 provide a reasonable fit (p-value of 0.0028), but not the best fit, is a measure of the small but non-negligible systematic differences in the modelling. One possible cause is the use of F_0 set at the 95% completion threshold, which is likely high. Reducing it would make CHIME more sensitive, and push the expected DM distribution to higher values, making it compatible with measurements. The model using the S22 parameter set also predicts a low number of singles bursts (131), though reducing the threshold to 2.5 Jy ms to produce the correct number results in a poor fit to the DM distribution.
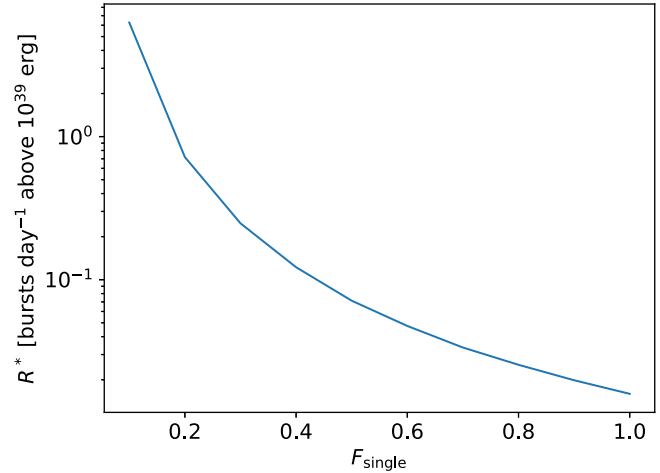
Another possibility is the complex interaction between burst shape and CHIME response, which is imperfectly captured here with the one-dimensional selection function $s(DM)$, but which S22 model explicitly using injected pulses. A third possibility is that S22 fit to all progenitors (i.e. single and repeat FRBs). However, performing the same comparison here, the difference between the S22 predictions and data becomes greater.

Lastly, I also check that the fitting is not strongly affected by errors in DM_{MW} at low Galactic latitudes. Such an error would smear the DM_{EG} distribution, creating excess low-DM events. Redoing the above analysis to include only FRBs with estimated Galactic contributions of less than 50 pc cm^{-3} changes the p-values by factors of order two, but this is minor compared to the different predictions between models. Thus this effect is ignored from hereon, and the entire sample is used to allow for greater precision and little cost of accuracy.

I therefore conclude that using the α 'min' parameter set from Table 1 in the zDM code is likely to provide a reasonable fit to the CHIME single burst rate, and hence it is used to fit to the repeating FRB population in the following Section.

6. Fitting results

To find a best-fit set of repeating FRB parameters ($R_{min}, R_{max}, \gamma_r$), I first find the critical value of repetition rate, R^*, such that when $R_{min} = R_{max} = R^*$ (and thus the value of γ_r is irrelevant), the correct number of repeating FRBs (in this case, 16) is reproduced. If $R_{min} > R^*$, then inevitably the model will produce too many repeating FRBs in the case that all bursts originate from repeaters; if $R_{max} < R^*$, then the model will not produce sufficiently many repeaters.

To illustrate, in Fig. 6, R^* is plotted as a function of F_{single}, being the fraction of all apparently once-off CHIME FRBs that are produced by true repeaters. Reducing the number of once-off bursts attributed to repeaters, while keeping the observed number of repeaters in Cat1 constant at 16, means that a higher fraction of true repeaters get detected as such, that is, they must be stronger. Indeed, for $F_{single} = 0.1$, the average repeater must repeat about 6 times per day above 10^{39} erg. For now, I continue with the case $F_{single} = 1$ (i.e. all FRBs are repeaters), and revisit $F_{single} < 1$ in Section 6.5.
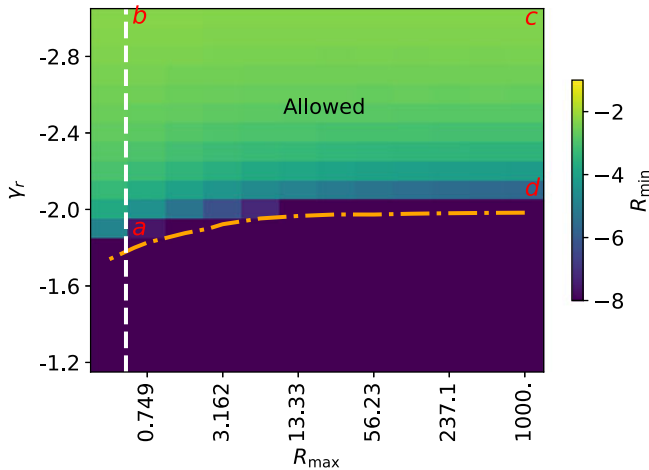
Figure 7. Value of R_{min} producing the observed number of 17 repeating FRBs in the CHIME catalogue (CHIME/FRB Collaboration et al. 2021) as a function of γ_r and R_{max}. Also shown are limits on R_{max} (white dashed) from FRB 20180916B, and the region excluded as producing too many repeaters (orange dot-dash curve). The total 'allowed' region is also indicated. Cases a–d used for §6.1 and onwards are indicated in red.

Figure 8. Predicted DM distribution of repeating FRBs compared to that from CHIME Catalogue 1 (CHIME/FRB Collaboration et al. 2021), calculated using cases a–d from Fig. 7, and the golden sample of repeaters from CHIME/FRB Collaboration et al. (2023) (renormalised to 16). Note that b and c overlap.

For any two values in the set $(R_{min} < R^*, R_{max} > R^*, \gamma_r)$, and fixed F_{single}, the third value can be found such that the number of repeating FRBs observed by CHIME is reproduced exactly. For reasons to do with code optimisation, calculations varying γ_r proceed very slowly, so γ_r is held fixed to a small number of values. Since reasonable estimates of R_{max} from observations of strong repeaters exist, it can be constrained to a sensible range. Therefore, I vary γ_r and $R_{max} > R^*$, and for each, calculate the value R_{min}.

The resulting values of R_{min} as a function of R_{max} and γ_r are shown in Fig. 7. For steep γ_r and for values of R_{max} not much above R^*, R_{min} must also be close to R^* (here, about 0.02 d^{-1}), while for flat γ_r and large R_{max} (the lower part of Fig. 7), R_{min} must be very small. For very flat γ_r, the contribution of low-R repeaters to the apparently once-off burst rate becomes sufficiently negligible that it cannot 'dilute' the number repeaters observed as such when R_{max} is large. This excludes the region in the lower right of the figure. For calculation purposes, R_{min} is set to 10^{-8} in this region—even though this over-produces the number of repeaters—it allows calculations to compare their DM, declination, and repeat rate distributions.

The lower limit on R_{max} given by FRB 20180916B is also shown—at a distance of approximately 150 Mpc (Marcote et al. 2020), its observed repetition rate above 5 Jy ms of $0.448^{+0.1}_{-0.086}$ h^{-1} (CHIME/FRB Collaboration et al. 2023) approximately translates to a rate above 10^{30} erg Hz^{-1} (i.e. 10^{39} erg assuming a 1 GHz bandwidth) of 0.5 d^{-1} when using a cumulative fluence index $\gamma = -1.5$.

An upper limit on R_{min} can be estimated from the lowest estimated rate for a low-DM FRB observed by CHIME. FRB 20190518D (CHIME/FRB Collaboration et al. 2021) has a DM of 202.2 pc cm^{-3}, with an estimated contribution by the Milky Way's interstellar medium (ISM) of 53.7 pc cm^{-3} according to the NE2001 model (Cordes & Lazio 2002). Conservatively assuming a low combined halo and host DM contribution of 35 pc cm^{-3}, similar to that found for FRB 20200120E (Bhardwaj et al. 2021b), and using a simplistic model of $z \sim 10^{-3}$ DM$_{EG}$, produces an approximate maximum redshift of $z \sim 0.11$. Again using $\gamma = -1.5$, this produces an upper limit on R_{min} of 0.05 d^{-1}.
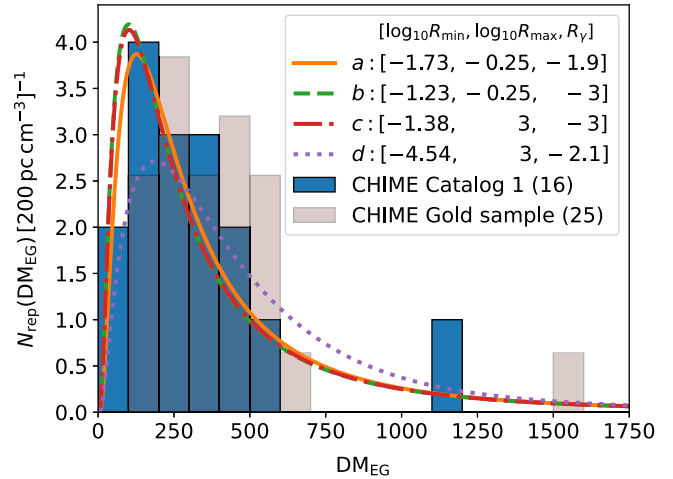
The combination of these three constraints produces the allowed region shown in Fig. 7. Note that in all cases, R_{min} is below (i.e. compatible with) the limit from FRB 20200120E.

6.1. Dispersion measure distribution

Within the range allowed by Fig. 7, the z–DM distribution predicted for each will in general be different. To illustrate, I take four scenarios, a–d, from the corners of the allowed region. In each case, the predicted DM distribution of repeating FRBs is plotted in Fig. 8.

From Fig. 8, it can immediately be seen that strong repeaters are more likely to be found at large distances (higher DM$_{EG}$). Case d has a significantly higher distribution of DMs compared to the other three, and it is the only case with a significant number of strong repeaters (cases a and b have low R_{max}, while case c has such a steep γ_r that the number of strong repeaters is negligible).

I also compare these DM distributions with those found from the Cat1 and Gold25 samples. While formally the DM distribution of the Gold25 sample is statistically consistent with that of Cat1 (CHIME/FRB Collaboration et al. 2023), the distribution is biased (see Section 8.3). That case d well reproduces the Gold25 sample DM distribution therefore should not be taken as evidence for it.

This comparison highlights a prediction of all repeating FRB models, which is the stronger upward skew of the DM distribution compared to single FRBs. This effect is seen in CHIME data, with two (one) high-DM repeaters in the Cat1 (Gold25) samples.

To quantify agreement in DM space, a KS test using the Cat1 FRBs and predicted DM distributions over R_{max}, γ_r space is performed. The results are shown in Fig. 9.

An implicit assumption of the above analysis is that the intrinsic distribution of DM$_{host}$ is identical regardless of the repetition rate of repeaters. The observation of PRSs at the locations of at least two bright repeaters (Marcote et al. 2017; Niu et al. 2022) suggests that these presumably young objects would be more likely to have a larger DM$_{host}$, bearing in mind that this term includes material in the vicinity of the progenitor, as well as the host galaxy's ISM and halo contributions. This could then be responsible for observing repeating FRBs (which are on-average intrinsically stronger repeaters) to have slightly more DM than expected, and would not
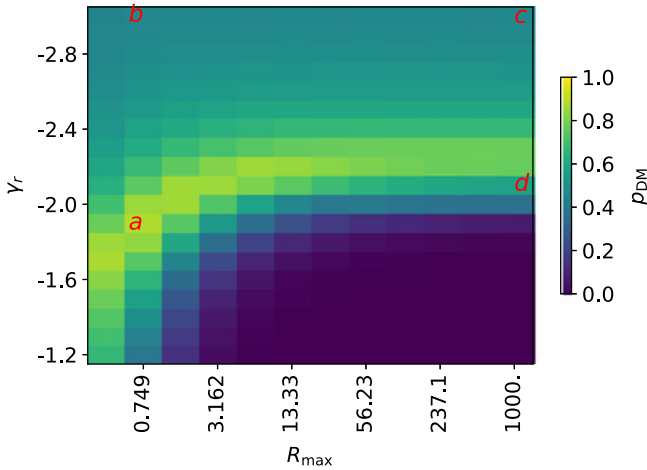
Figure 9. *P*-values from a KS test of the DM distribution of Cat1 repeating FRBs, $p_{ks}(DM_r)$, against predictions from models with different values of R_{max} and γ_r. Other features are identical to Fig. 7, including cases *b* and *c* overlapping.

Figure 10. Cumulative histogram of the CHIME repeating FRB declination (δ) distribution, for both the Cat1 and Gold25 samples, compared to Monte Carlo predictions from four example cases.

constitute hard evidence against the model. This might be the case for the $\gamma_r \lesssim -2.4$ region of Fig. 9, which is slightly disfavoured because it over-predicts the number of low-DM repeating FRBs. However, should observed repeaters have less DM than expected, this is clear evidence to reject the model. This is the case for the already ruled-out lower region of Fig. 9, which predicts more high-DM repeating FRBs than observed.
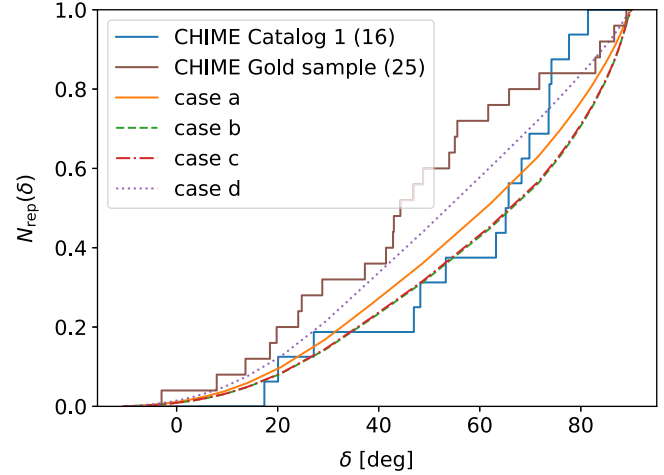
6.2. Declination distribution

The declination distribution of CHIME FRBs also holds information that allows us to discriminate between scenarios. For repeating FRBs, the difference between observing a small patch of the sky around the NCP almost continuously, and surveying several steradians near the equator for only a few minutes each day, is very important. Near the equator, only the strongest repeaters will be detected as such, while near the Pole, the small probed volume makes observations subject to cosmic variance.
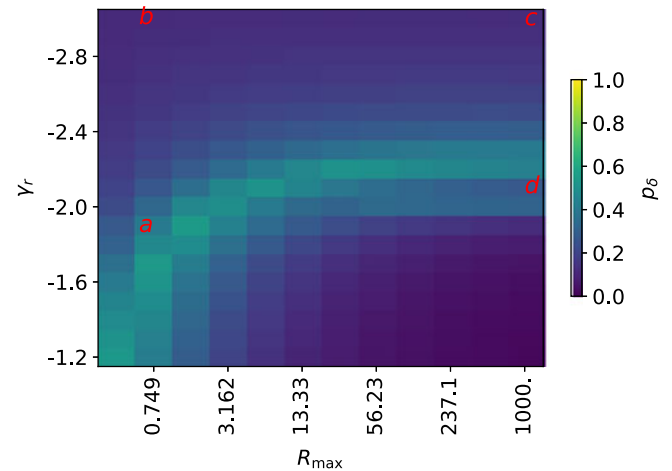
In Fig. 10, I plot the declination distribution of CHIME once-off and repeating FRBs, and compare this against model predictions. For this plot, the number of declination bins into which CHIME was divided was increased to 30, whereas six declination bins was found to be sufficient to model the total number of repeaters, and their DM distribution.

Since the *x*-axis of Fig. 10 is increasing linearly with δ, most of the solid angle is concentrated on the left-hand-side of the figure. Despite this, the number of repeaters—both observed and predicted—increases as fast as, or faster than linearly with δ. That the Gold25 sample shows the least steep rise with δ is likely because of the previously discussed bias against high declinations due to the increased background rate. This is evidence that the repeating population is dominated by progenitors with low apparent repetition rates that are best probed with deep observations (i.e. at high declinations), rather than sources with high apparent rates that are best detected in broad shallow surveys (i.e. at low declinations).

Of the four cases analysed, a–c show good agreement with Cat1 in the $\delta \lesssim 60°$ range, while not even d can match the rapid rise in repeater rates above this range. This suggests a simple fluctuation in the data, either a deficit at low declinations, or an excess at high

Figure 11. Results of the KS test against the declination distribution of identified repeating FRBs. Shown is the *p*-value as a function of R_{max} and γ_r.

declinatons—though an alternative explanation is the influence of non-Poissonian repetition (see Section 8.1).

I characterise the agreement in declination distributions via a KS test, with associated *p*-values given in Fig. 11. The greatest discrepancy with data is the aforementioned excess of high-δ repeaters, and the upper region of the figure is disfavoured because it reproduces this particularly poorly. The lower right region is disfavoured because this predicts mostly bright repeaters that should be found in the greater region of sky viewed at low declinations.

6.3. Repetition rate distribution

Most repeating CHIME FRBs are not localised, so that scaling between intrinsic and apparent repetition rates, which requires the luminosity distance to be known, is not possible. This precludes a direct fit to the rate distribution. Nonetheless, different combinations of R_{min}, R_{max}, and γ_r lead to more/less repeating FRBs being observed with different apparent repetition rates.

Directly computing the number of FRBs with any given repetition rate is highly inefficient however—the algorithm currently
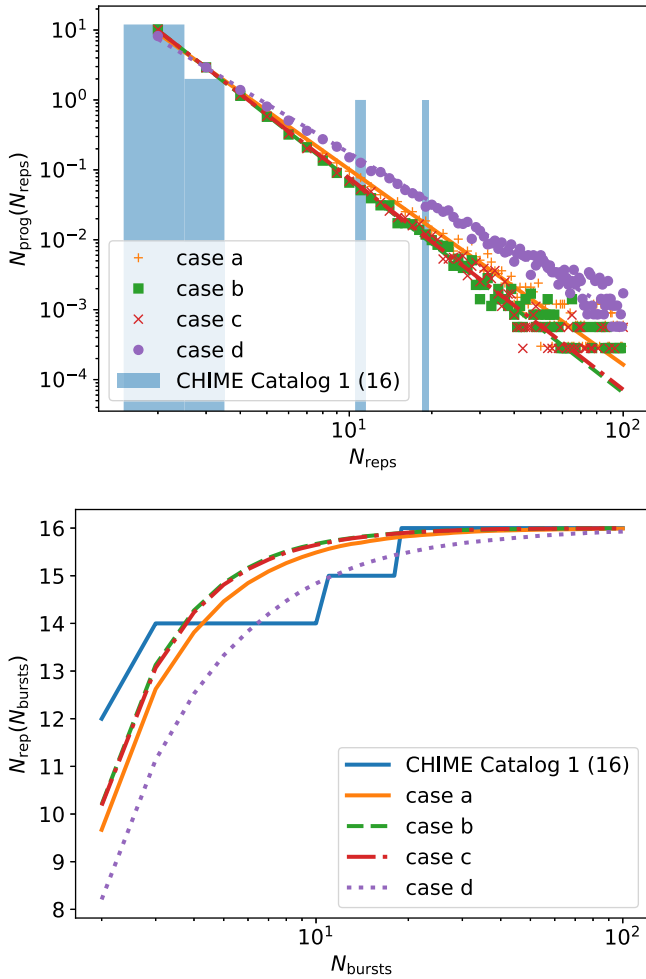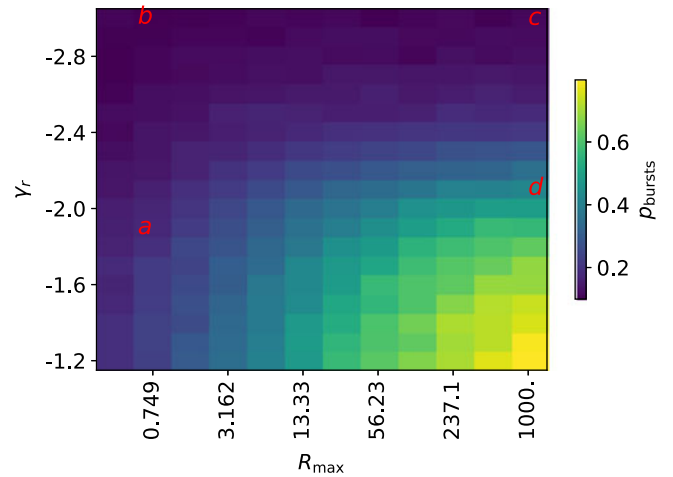
Figure 12. Top: histogram of observed number of repetitions in CHIME repeating FRBs from Cat1, compared to Monte Carlo predictions from four example cases, a–d (points). A power-law fit (lines) is given for each. Bottom: the same data, but shown as a cumulative distribution.
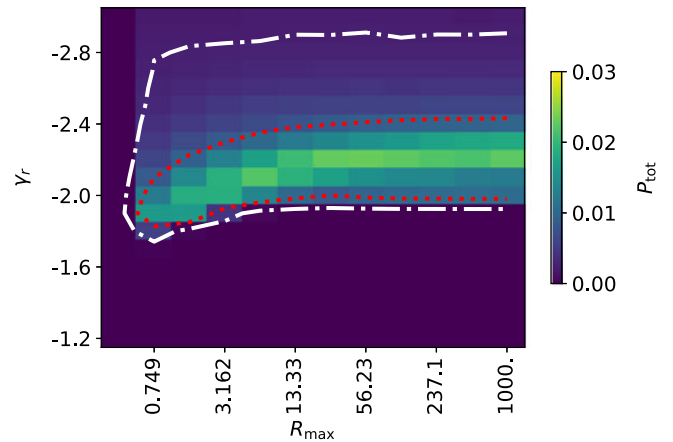

Figure 13. Maximum-likelihood estimates of FRB repeat parameters based on the distribution of the number of observed bursts from each repeater in the Cat1 sample.

Figure 14. Posterior probability of repeating FRB parameters assuming that all FRBs repeat. Shown are 68% (red dotted lines) and 95% (white dot-dash lines) confidence intervals.

estimates the number of repeaters by explicitly calculating N_0 and N_1 only. Extending this to a large number of N_{burst} values is computationally prohibitive. A Monte Carlo sampling algorithm was therefore implemented that generates repeating FRBs according to their underlying modelled distribution in z–DM–R space, and simulates the number of observed bursts assuming a Poissonian distribution.

To overcome Monte Carlo fluctuations, at least 1000 times as many repeating FRBs as expected are simulated, and a histogram produced in terms of the observed number of bursts by CHIME. This is then fit with a power-law distribution, and for histogram bins with less than 10 simulated repeaters, the observed number is replaced with the fitted number for purposes of evaluating likelihoods. An example of this procedure is shown in Fig. 12.

Since the data are discrete (integer numbers of bursts only), a KS test to assign a goodness of fit is inapplicable. Instead, the likelihood of the observed histogram of N_{burst} values for the 16 CHIME repeating FRBs from the Cat1 sample is calculated, given predictions from the Monte Carlo histogram. This is then repeated for at least 1000 sets of 16 Monte Carlo FRBs, and the fraction of likelihoods that are lower than that observed is determined.

This produces a p-value, p_{bursts}, under the null hypothesis that the Monte Carlo sample is the truth. Results are plotted in Fig. 13.

The repeat-rate distribution is best-reproduced by models with a large number of bright FRBs, since the two CHIME FRBs with high repetition rates in Cat1—FRB20180814A (11 bursts), and FRB20180916B (19 bursts)—are difficult to reproduce with models of low R_{max} and/or steep γ_r.

6.4. Combined likelihood

Combining the evidence from the DM, δ, and N_{burst} probabilities derived above, the combined probability p_{tot} is constructed as

$$p_{tot} = p_N \, p_\delta \, p_{DM} \, p_{bursts}, \qquad (13)$$

where p_N is a Poissonian probability of observing 16 repeaters, which suppresses the region of the parameter space that over-produces repeaters. The probabilities are renormalised to sum to unity over the investigated range, excluding $R_{max} < 0.5$ d^{-1}, and confidence intervals assuming flat priors in γ_r and $\log R_{max}$ are constructed. This results in the probability distribution, and confidence intervals (CIs), shown in Fig. 14.
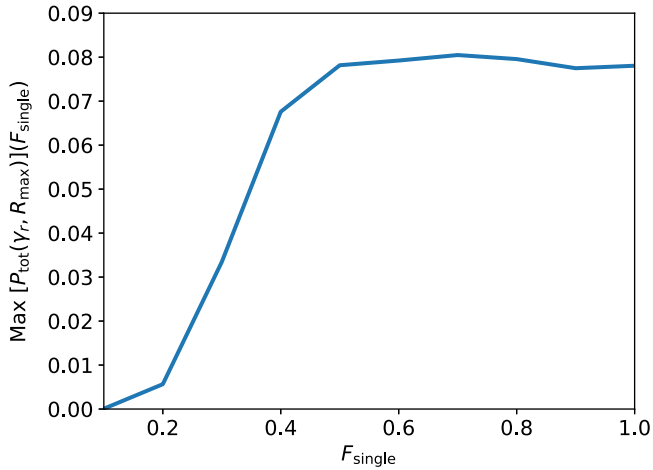
Figure 15. Maximum value of the joint probability P_{tot} over all analysed γ_r, R_{max}, as a function of the fraction of all CHIME single bursts explained by repeating FRBs, F_{single}.

Figure 16. Number of repeaters, N_{rep}, normalised by total observation time T, in units (and as a function) of the exposure from Cat1, for cases b and d.

The 95% CI encompasses almost the entire allowed region from Fig. 7, showing that DM, δ, and N_{burst} are not strong discriminators between different models of the repeating FRB population. However, a preference for $\gamma_r = -2.2^{+0.6}_{-0.8}$ (68% CI), and $R_{max} \geq 0.75$, is found.

The above formulation for p_{tot} ignores correlations between variables: closer (low DM) FRBs, and those viewed closer to the zenith, will be more likely to have more bursts detected. A better analysis would use the full 3D distribution of $p(DM, \delta, N_{bursts})$, similarly to the use of redshift, DM, and burst energy in standard z–DM analyses. However, generating the distribution of N_{bursts} is computationally intensive, and thus it is only performed in one dimension. This method should be revisited once other systematic effects, as discussed in Section 8, are treated.

6.5. What if not all FRBs are repeaters?

It is of course possible that repeating FRBs do not constitute the total FRB population. Evidence for this comes from the different spectro-temporal properties of repeaters compared to non-repeaters (Pleunis et al. 2021), and a tentative association of FRB 20190425A with binary neutron star merger GW190425 (Moroianu et al. 2023). If such a population exists, it is likely subdominant—most cataclysmic events, which would produce intrinsically once-off FRBs, have a rate which is much too low to explain the total FRB rate (Ravi 2019). Therefore, observations of the total FRB population still serve as good constraints on the total repeating population, and the predictions made here remain valid. Nonetheless, in this section, the case where repeating FRBs are responsible for a sub-dominant fraction of the total number of bursts observed by CHIME is investigated.

The above analysis is repeated by first optimising R_{min} to produce 16 CHIME repeaters and some fraction F_{single} of the total singles burst rate, and calculating the joint probability $P_{tot}(\gamma_r, R_{max}, F_{single})$. The peak likelihood over γ_r and R_{max} for each F_{single}, $Max[P_{tot}(\gamma_r, R_{max})](F_{single})$, is then plotted in Fig. 15.

In the range $0.5 \leq F_{single} \leq 1$, the peak probability is essentially identical, with fluctuations likely due to the coarse gridding in γ_r–R_{max} space. The likelihood decreases for lower values of F_{single}—this is driven almost entirely by p_{DM}, since decreasing F_{single} increases the fraction of true repeaters detected as such,
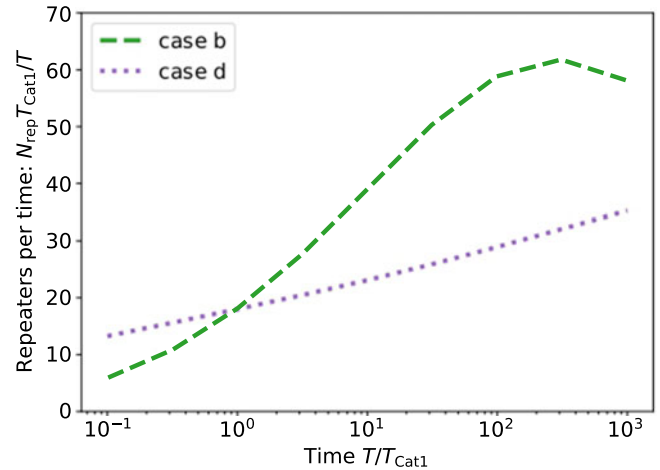
which requires on-average stronger repeaters that are invariably detectable at greater distances. This pushes the predicted DM distribution to higher values, inconsistent with CHIME data.

A note of caution is warranted however: if a small fraction of all single bursts are produced by repeaters, then the assumption that repeating FRB population parameters are the same as that of the total population is a bad one. Therefore, while it can be concluded that these results are consistent with a best-fit of all FRBs being from repeaters, it cannot be concluded that this excludes a large fraction of FRBs being from intrinsically once-off events.

7. Future prospects

Now that an estimate of the parameters of the repeating FRB population has been made, I make predictions for the effects of repetition on future observations. In the following, cases d (close to the best-fit values found in Section 6) and b (marginally excluded at the 90% level, albeit when considering random error only) are considered as two significantly different, but plausible, cases.

7.1. Rate of new repeater discoveries with CHIME

As time spent observing a particular field increases, the number of repeating FRBs should eventually saturate, as essentially all such objects in the field are detected. Seeing the rate of detected repeaters plateau at a level where a large number of once-off bursts have no associated repeater would be a clear indication of two populations. This raises the question: how long might CHIME have to wait until the rate of new repeating FRB detections decreases?

The answer is a very long time. Regardless of the scenario under consideration, the number of repeating FRB progenitors at high redshifts will vastly outnumber those at low redshifts due to the increased volume of the Universe. As observation time increases, the number of repeating FRBs in the nearby Universe will saturate, but the rate of repeater discoveries—both as single and repeat bursts—in the distant Universe increases. This effect is seen in Fig. 1—in Fig. 16, this is simulated for CHIME, by simply increasing the observation time in units of Cat1, T_{Cat1} (which is approximately a year's worth of exposure). In case b, there are relatively few strong repeaters, and saturation is expected to be seen
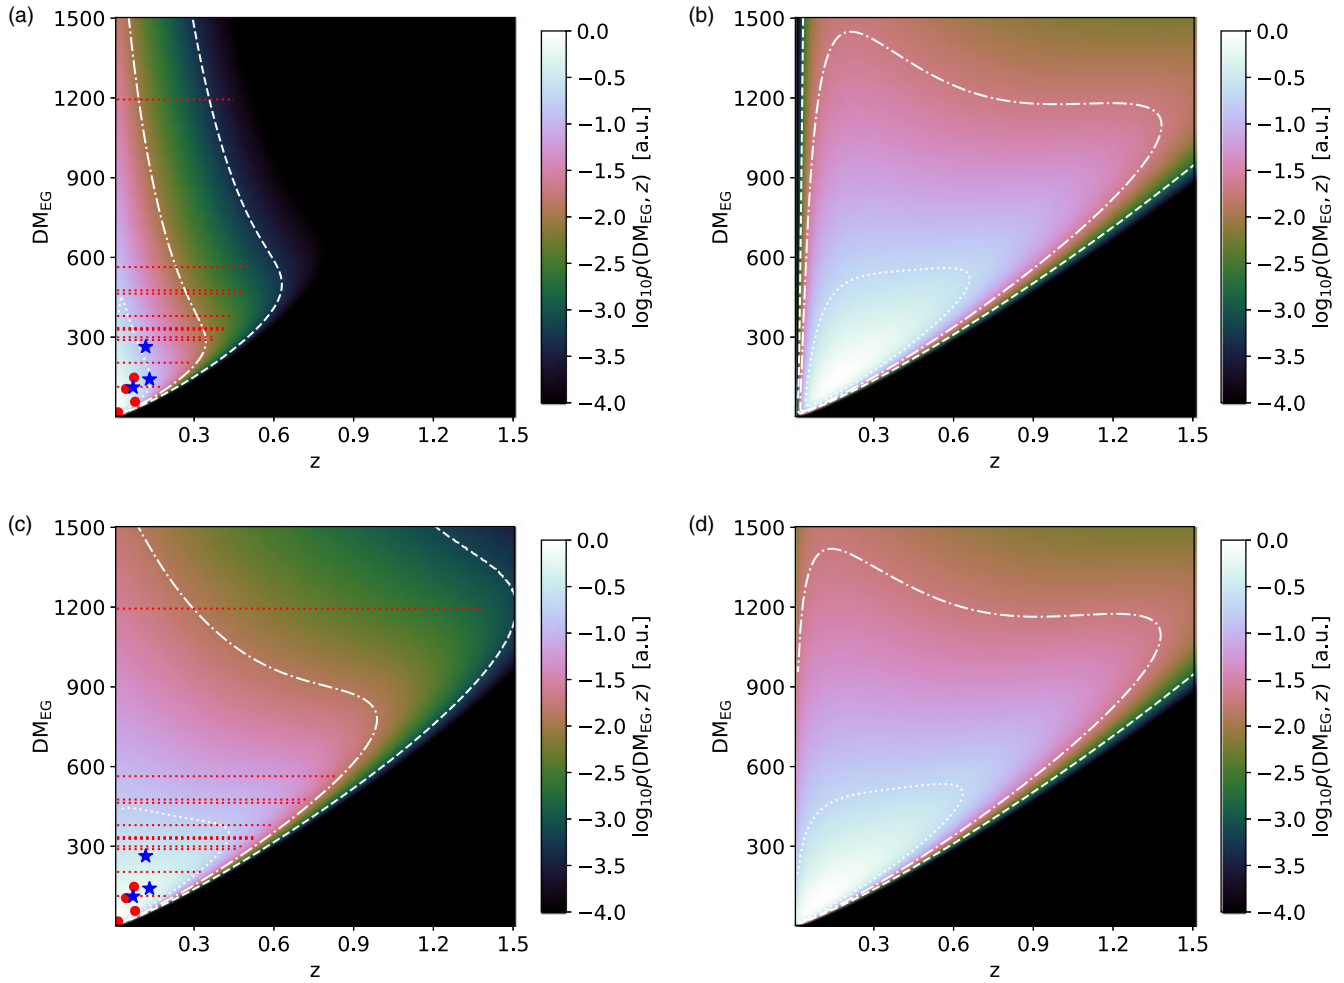
Figure 17. Predicted z–DM distribution of left: repeating and right: single FRBs for cases *b* (top) and *d* (bottom). Contours enclose 50% (dotted), 90% (dot-dash), and 99% (dashed) of the probability space. Repeating CHIME FRBs identified in Cat1 with host galaxies are shown as red circles; repeater hosts not from Cat1 are shown as blue stars; and repeating FRBs with no firm host association from Cat1 have their 0–99% probable redshift range indicated with red lines at their known DM_{EG}, calculated assuming the Cordes & McLaughlin (2003) model for DM_{ISM}, and a value of 50 pc cm^{-3} for DM_{halo}.

after ∼300 yr. In case *d*, with many strong repeaters, saturation will not occur in the next thousand years, and a steadily increasing repeat rate is expected.

7.2. z–DM distribution

Only six repeating CHIME FRBs have been localised to their host galaxies (Marcote et al. 2020; Bhardwaj et al. 2021b,a; Fong et al. 2021; Ibik et al. 2023), though three more associations are highly likely (Michilli et al. 2022; Ibik et al. 2023); and only one once-off CHIME FRB has a tentative host association (Panther et al. 2022). Of these repeaters, four were identified as such in Cat1. Furthermore, these have only been localised either because they are nearby, and hence CHIME's angular resolution—effectively enhanced when using multiple bursts (Michilli et al. 2022)—is sufficient to identify the host; or because they repeat rapidly, allowing follow-up observations with arrays with a better angular resolution to identify the host. Thus these represent a highly biased sample, and are unsuited to fitting to data. Nonetheless, the full z–DM distribution of repeating FRBs observed by CHIME can be predicted by the models. These distributions are given in Fig. 17

for cases *b* and *d*, and are compared to the distribution of singly detected FRBs.

Fig. 17 illustrates how the large tail of the DM distribution arises: it is almost entirely from objects lying well above the Macquart relation. Since repeating FRBs tend to only be detected as such in the nearby Universe, those repeating FRBs lying on the Macquart relation have a smaller range of DMs—using the 90% contours, up to ∼400 pc cm^{-3} in case *b*, and ∼800 pc cm^{-3} in case *d*. Thus the high-DM tail of low-z FRBs doesn't become over-ridden by the larger number of FRBs lying on the Macquart relation in the more-distant Universe, as it is for singly detected FRBs.

This closer proximity of repeating FRBs means that the reduced DM_{EG} of the CHIME repeater sample (436 ± 49 pc cm^{-3} for repeaters, 597 ± 24 pc cm^{-3} for apparently once-off bursts; CHIME/FRB Collaboration et al. 2023), which has been suggested to be evidence for two populations (Woods 2023), is entirely consistent with expectations from the models. These predict mean repeater DM_{EG} values in the range 460–540 pc cm^{-3}, and mean single DMs in the range 640–660 pc cm^{-3}. While the mean DMs of both samples are slightly over-predicted, the difference is a very good match with expectations.
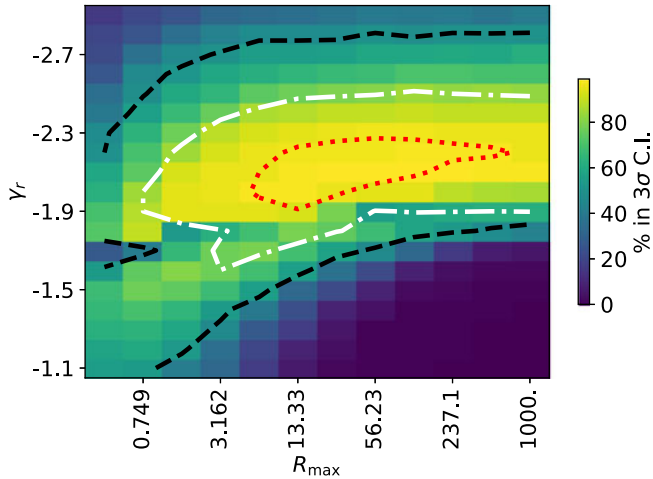
Figure 18. Fraction of Monte Carlo iterations in which trial values of γ_r, R_{max} fall within the 3σ confidence interval (CI), using a sample of simulated repeating FRBs with truth values $\gamma_r = -2$, $R_{max} = 31.62$. The contours correspond to regions that fall within the 1 (red, dotted) 2 (white, dash-dot) and 3 (black, dashed) confidence intervals 50% of the time.

Even more importantly for future observations, models b and d predict different z distributions. Case b has more low-rate repeaters, from which repeat bursts are only likely at low z; while case d has a significant population of high-rate repeaters, which can be detected as repeaters from the more-distant Universe. The z–DM values of repeaters with likely or confirmed host galaxies matches observations for both cases, but this has only been probed in the low-z, low-DM region. However, case d predicts redshifts will continue to increase with DM_{EG}, while case b predicts $z \lesssim 0.5$ for all repeaters. Thus, if a large fraction of repeating CHIME FRBs could be localised, this would enable much more powerful tests of the repeating FRB population.

I illustrate this using a toy example, using simulated true values $\gamma_r = -2.2$, $R_{max} = 30$, and 100 Monte Carlo instances of repeating FRBs from Cat1. All FRBs detected as repeaters are assumed to be localised to their host galaxies, yielding their correct z and DM values. For each Monte Carlo sample, the likelihood $p(z, DM)$ is calculated for all values on the γ_r, R_{max} grid. I do not calculate $p(N_{reps})$, that is, only the position in z, DM is accounted for, not the number of repeating FRBs or bursts per repeater. Bayesian 1-, 2-, and 3-σ confidence intervals are then constructed for each sample, and the number of MC iterations in which any given value lies in each interval is counted. The result is shown in Fig. 18, which shows the expected confidence intervals at each of the three levels. This shows the power of being able to localise repeating FRBs: if all Cat1 repeating FRBs could be localised, the expected 1σ accuracy on γ_r would be ± 0.2, and R_{max} would be determined to within a factor of ~ 10.

7.3. Predicted effects on other instruments

I now use cases b and d to estimate the relative rates of single and repeat observations for a sample of other FRB-hunting instruments. Four systems are considered: ASKAP, in Fly's Eye (FE), ICS, and coherent (CRACO) mode at 1.3 GHz; and the Five-hundred-metre Aperture Spherical Telescope (FAST). ASKAP/FE and ASKAP/ICS are modelled as per James et al. (2022b), while the model of the CRAFT Coherent Upgrade (CRACO) system is described in James et al. (2022c). The parameters for FAST

FRB searches are taken from Niu et al. (2021), namely a detection threshold of 0.0146 Jy ms for a 1 ms pulse width at a central frequency of 1.25 GHz, and time and frequency resolutions of 196.608 μs and 0.122 MHz respectively. The FAST receiver is a 19-beam multibeam (Li et al. 2018), similarly designed to the 13-beam Parkes multibeam (Staveley-Smith et al. 1996). I therefore take the inverse beamshape Ω_b used for Parkes, scale up by the ratio of the number of beams ($19/13 \approx 1.46$), and down by the ratio of effective collecting areas ($64^2/300^2 \approx 0.0456$).

All these instruments have searched for FRBs with different dwell times. Here, I consider the longest time spent on any given field for each instrument, which will prove most sensitive to the repeating FRB population: 1338.9 h for ASKAP/FE (James et al. 2020b), 879.1 h for ASKAP/ICS to the end of 2022 (Shannon et al. in preparation), and for ASKAP/CRACO predictions, the expected on-source time of 800 h for each of the Deep Investigation of Neutral Gas Origins (DINGO; Rhee et al. 2023, see also https://dingo-survey.org/) fields is used. For FAST, it is 59.5 h when performing follow-up observations on FRB 20121102A (Li et al. 2021). The normalised estimates are given in Fig. 19, as a function of DM for surveys with poor localisations, and as a function of z for those that typically identify host galaxies.

Qualitatively, all predictions are very similar. The total number of single bursts ranges from 27% of the total burst distribution (FAST, case b) to 70% (ASKAP/FE, case d). The difference between case b, which models a repeating FRB population spread over a narrow repetition rate, and case d, with a very broad distribution of rates, is marked, predicting 16% (case b) and 6% (case d) of FRB progenitors to repeat for all ASKAP models, and 31% and 10% for cases b and d for FAST.

The deficit between total burst number and total progenitors in the low-DM range for ASKAP/FE is not sufficient however to explain the observed deficit that has been previously noted by James et al. (2022b), especially when accounting for the average pointing time for that survey being less than that modelled here. Thus I conclude this effect—which originally motivated this work—is most likely a statistical fluctuation.

Jankowski et al. (2023) have noted that the FAST FRB rate is much lower than predicted. Here, the progenitor rate is predicted to be 40–60% of the burst rate, which certainly accounts for some, but not all, of the deficit. However, this would have no influence on the observations in drift-scan mode reported by Niu et al. (2021) due to the very short dwell times (~ 13 s). Thus this deficit must have some other explanation.

A single FRB (20220531A; Shannon et al. in preparation) has been discovered in the ASKAP field with 879.1 h of observations, against a mean ASKAP detection rate of 350 h FRB^{-1}. This could simply be a Poissonian under-fluctuation (p-value of 0.285 on a one-sided test), but repetition offers a partial explanation, which would reduce the expected number of progenitors from 2.5 to 1.75–1.86.

Overall, I expect that correct modelling of repeating FRBs will be important for these observations to account for repeater bias in the observed z–DM distribution.

8. Discussion of systematic effects

8.1. Non-Poissonian repetition

All FRBs with sufficiently many detected bursts to allow studies of their repetition rates show non-Poissonian behaviour. On
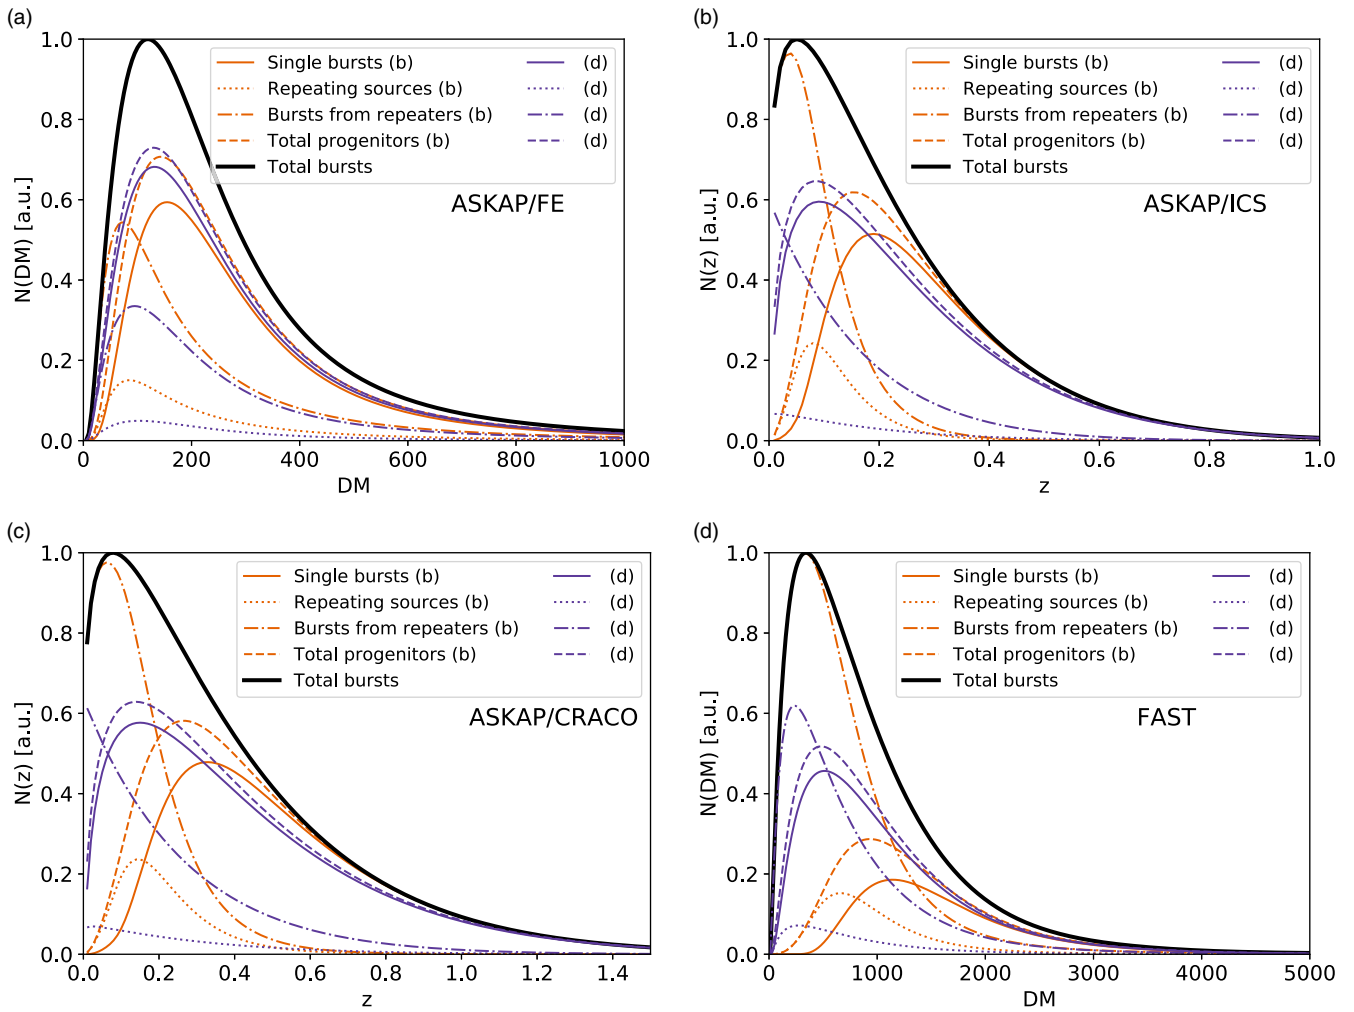
Figure 19. Predictions of the z or DM distributions of repeating FRBs for a selection of past and future FRB surveys for their longest pointing times (see Section 7.3), for cases (b) and (d). Shown are the distributions of those repeating FRBs detected as single bursts, as repeaters, the total progenitor distributions, and the total burst distributions, as per Fig. 1.

timescales of order seconds to hours, bursts from repeaters tend to be clustered (e.g. Gajjar et al. 2018; Zhang et al. 2021; Nimmo et al. 2023), in a process which is often modelled as a Weibull distribution (Oppermann, Yu, & Pen 2018). On longer timescales (~16–160 d), two repeating FRBs appear to have activity cycles (CHIME/FRB Collaboration et al. 2020; Rajwade et al. 2020), with evidence for frequency dependence in the timing of the windows (Pastor-Marazuela et al. 2021). Other behaviours include a rapidly increasing/decreasing event rate (Zhang et al. 2022), or 'turning on' despite several years of monitoring (see the time dependence of bursts in CHIME/FRB Collaboration et al. 2023). What the true underlying nature of the time distribution of repeat rates of FRBs is is still under debate—what is sure is that they are most certainly not Poissonian.

Thus it should be asked: what effect does this have on the modelling? On sufficiently long timescales, FRBs will become inactive, and new repeating FRBs will be born. Therefore, the repeating FRB population studied here can only refer to those FRBs which have been active during the approximate year (three years) corresponding to the CHIME Cat1 (Gold25) samples. However, FRB 20121102A has now been studied for over a decade since its first

detection (Spitler et al. 2014), and while its properties (DM, RM etc.) do vary (Michilli et al. 2018), no evidence of a systematically reducing rate has been published. Therefore, these considerations likely won't be relevant to current or near-future studies.

Of more relevance are FRBs with inactive windows comparable to, or longer than, the current survey. CHIME/FRB Collaboration et al. (2023) shows that at least three FRBs— 20201130A, 20200929C, and 20201124A—have numerous bursts in the latter six months studied, but none in the first two years. While some of this may be reflective of a changing search sensitivity and analysis methods (it would be useful for CHIME to release a time-dependent sensitivity to account for this), it is also suggestive that these objects have very long inactive phases. By generating either many bursts or none, a larger population of such objects would mimic a flatter value of γ_r than the true long-term rate, with repeaters being found at larger distances/DM values. Conversely, for a fixed observation, the fitted value of γ_r will be flatter than the true value. Since the observed DM distribution of repeaters already favours $F_{single} \geq 0.5$ and $\gamma_r \leq -1.4$, allowing for such behaviour would constrain F_{single} to higher, and γ_r to lower, values. That is, the limits from this work are sensitive to activity windows on
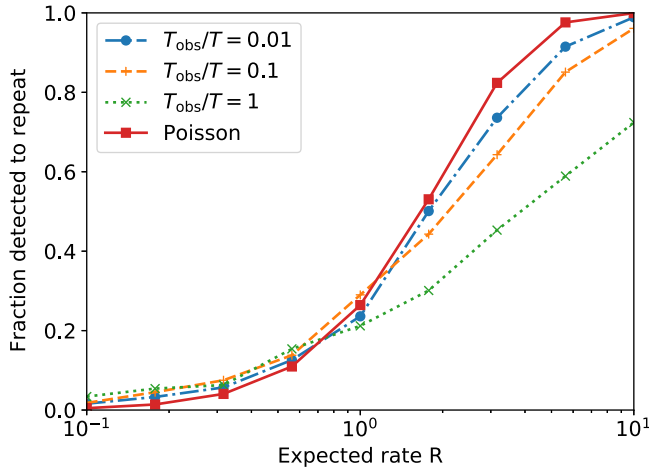
Figure 20. Probability of an FRB being detected as a repeater given its true expected rate R, as a function of the fractional observation time $f_{obs} = T_{obs}/T$, for a Weibull burst time distribution with shape index $k = 0.34$; and for a Poissonian distribution.

yearly timescales. Activity windows significantly shorter than a year however will have no consequence, since CHIME's coverage is spread uniformly in time, unless the period of these windows lies extremely close to a sidereal day.

Finally, the effect of bursty behaviour in general is to reduce the number of singly detected FRBs, and increase the likelihood of viewing zero or many bursts. The effect of time correlation in bursts will be most pronounced when observations occur all in one block—when observations are individually very short and spaced far apart in time, any intrinsically bursty distribution will exhibit a Poissonian distribution of burst numbers.

To gauge the impact of this effect, I simulate a Weibull distribution of arrival times, using a shape parameter $k = 0.34$, as found for FRB 20121102A by Oppermann et al. (2018). CHIME detections are simulated at three declinations: near the NCP, with sources observed continuously ($f_{obs} = 1$); approximately $7°$ away from the NCP, where sources will be observed a fraction $f_{obs} = 10\%$ of the time; and $\sim 30°$ from the NCP, which is modelled as a source being observed $f_{obs} = 1\%$ of the time. The expected repetition rates are modelled relative to the total time on-source in a calendar year, and this rate R_{obs} is varied from 0.1 to 10. I simulate 1000 Weibull sequences over 365 sidereal days, beginning each sequence well before the start of the year to ensure the sequence start time does not influence results. If a burst occurs in the first f_{obs} fraction of a day, it is counted as detected. The number of simulations resulting in none, one, or multiple detections are recorded.

The results are given in Fig. 20. As predicted, the fraction of FRBs detected to repeat twice or more is greater than Poissonian for low expected rates, and less for high expected rates. However, only very near the NCP is this effect large, where a deficit of $\sim 30\%$ of repeating FRBs relative to Poisson rates are found.

A more accurate estimate of the effect of burstiness on these results however is the ratio of single to repeat bursts. Weighting the results by R^{-2}, to represent there being less rapidly repeating FRBs than rarely repeating ones, $f_{obs} = 0.01, 0.1$, and 1 respectively produces 30%, 120%, and 310% more repeating FRBs relative to single bursts than expected for a Poissonian distribution. This effect would also alter the z–DM distribution of repeating FRBs, with bursty distributions favouring discoveries in the more-distant Universe; and flatten the distribution of observed burst rates,

undermining the fitting of Section 6.3. It is quite possible that the excess of repeating FRBs in Cat1 at high declinations, while not statistically significant, is due to this effect.

Nonetheless, I do not wish to over-emphasise this effect. For the majority of the sky seen by CHIME, the increase in observed repeaters due to burstiness with $k = 0.34$ is tens of percent. Furthermore, this choice of $k = 0.34$ is an extreme example— analysis of FRB 20121102A has shown that at timescales of seconds to hours, the wait time distribution is close the Poissonian, with a best-fit Weibull index of $k = 0.79$ (Nimmo et al. 2023).

Given the complexity of the issue, I defer an analysis of repetition that includes bursty behaviour to a future work—which should also include a fit of such behaviour to CHIME arrival time data.

8.2. Time-frequency structure of bursts

CHIME has shown that FRBs identified to repeat tend to have bursts which are broader in time by a factor of approximately two, with complex time-frequency structure (CHIME/FRB Collaboration et al. 2021; Pleunis et al. 2021). Whether this is due to two intrinsically different FRB populations, or a smooth transition in the properties of one population (e.g. an ageing effect whereby less active—and presumably older—objects have modified emission physics, perhaps due to a change in beaming angle, as suggested by Connor, Miller, & Gardenier 2020) is still up for debate. What is clearly true is that the increased time width of FRBs more likely to be detected as repeating will make them harder to detect for a given fluence. This effect is not included in the current work.

There are two methods of analysing this effect. The first is to modify the simulated burst width to increase with FRB repetition rate. This will act to suppress the number of repeating FRBs in the $z \lesssim 1$ range where intrinsic FRB width, rather than dispersion smearing, dominates the apparant FRB width, and hence sensitivity. The result will be that more active repeaters—which are in any case only observed as such at low z—become less detectable. Including this effect would require direct use of the CHIME pulse injection sample, since the published efficiency function analysed in Section 4.3 has already been averaged over the burst width distribution.

The second method to tackle this problem is to note that repetition rate R can be thought of as an effective repetition rate. After all, R can only ever be defined as the rate above a given energy threshold (here, 10^{39} erg), which must also be coupled to some assumption about the time–frequency properties of those bursts which affects their detectability. If more-strongly repeating FRBs tend to produce wider bursts, this reduces their detectability, and hence their apparent rate will decrease. In other words, this will cause γ_r to steepen slightly from its true value. In the regime where detectability is limited by intrinsic width w, SNR $\sim w^{0.5}$, and hence for a burst luminosity function with cumulative slope $\gamma = -1$, the detection rate will reduce as $w^{-0.5}$. If the width w scales linearly with the intrinsic rate R (and the real effect is unlikely to be this strong), this would then steepen the apparent value of γ_r by an extra factor of -0.5.

Since either method requires an accurate estimate of the relationship between intrinsic FRB rate and measured width w, and this is not currently possible, I consider the second method appropriate, which means that a little care must be taken in interpreting the estimate of γ_r found in this work.

8.3. Systematic effects in CHIME data

I conclude the discussion of systematics with an analysis of the data being used. As noted in Section 4.1, the threshold for identifying a repeat burst in Cat1 was lower than for an initial burst. The best way of removing this effect would be for CHIME to publish an FRB catalogue with those repeat bursts that passed threshold only because they were repeaters removed or otherwise identified as such. The dependence on SNR should be overcome by use of the CHIME pulse injection data (Merryfield et al. 2022), or a publication of the parameterised multi-dimensional selection function used in CHIME/FRB Collaboration et al. (2021). It has been checked that excluding FRBs with high Galactic DMs does not significantly change the DM fits, but future analyses with ZDM should in any case include an uncertainty term for this contribution. Overall, I do not think that this analysis is currently constrained by such systematic effects.

I have been reluctant to make quantitative comparisons in this work with CHIME's new Gold25 sample of 25 repeating FRBs however. There are several reasons why interpretation is difficult. Firstly, note that this work is in fact a discovery of 34 new repeating FRBs, given that the 'gold' sample of 25 sources has an estimated contamination of 0.5 from coincidences between two or more unrelated FRBs, and the authors also publish a 'silver' sample of 14 repeaters, with a total contamination rate in the combined gold and silver samples of 5. It is unknown whether or not CHIME have detected two or more bursts from a true repeater which has a higher contamination fraction and is thus not included in any sample. This measurement therefore has a $5^{0.5} = 2.2$ systematic uncertainty reflecting the uncertainty in the expected contamination, and a $34^{0.5} = 5.8$ statistical deviation from the expected mean number of repeater discoveries, for a total uncertainty of $(5 + 34)^{0.5} = 6.2$, or 18%.

The time period used for the search, from 2019 September 30 to 2021 May 1 (579 d), is approximately 70% longer than and does not overlap with, the period of 2018 July 25 to 2019 July 1 used for Cat1 (342 d) where 16 repeating FRBs were identified. Hence, the detection rate has increased by a factor of 1.26 ± 0.23, consistent both with the predictions of an increasing discovery rate in Section 7.1 and with a constant rate. Aside from Poisson error, this increase could also be due to a higher efficiency of CHIME data-taking, or a lower threshold for including bursts in the analysis. This question should be revisited once it becomes possible to normalise the two samples, for an accurate rate comparison.

Secondly, the identification of repeating FRBs in Gold25 placed a cut on the chance coincidence probability. This cut is strictest where the rate of FRBs in DM–δ space is highest: at high declinations and intermediate DMs. Therefore, the gold sample of 25 repeating FRBs is biased towards very low or high DMs at low declinations. Conversely, the silver sample may have the opposite bias, both because repeaters at high declinations and intermediate DMs will preferentially be placed there and because this region has more chance coincidences. CHIME/FRB Collaboration et al. (2023) do not publish estimates of the contamination probability in δ-DM space, which would be required to account for this effect.

9. Comparison with literature results

9.1. ASKAP follow-up observations

The only other result of which I am aware that has limited R_{min}, R_{max}, γ_r James et al. (2020a). Those authors used the observation
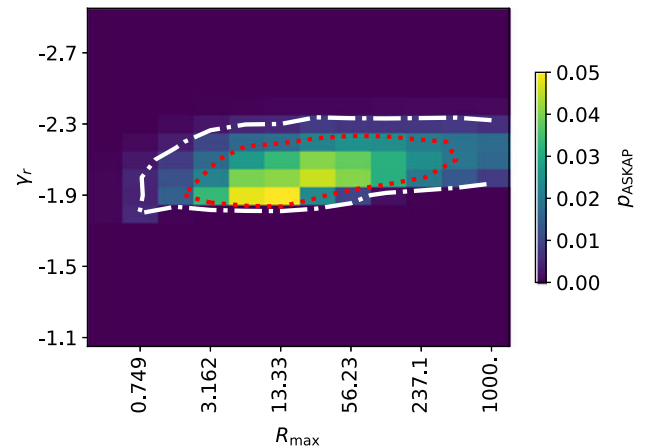
Figure 21. Bayesian posterior likelihoods, p_{ASKAP}, from follow-up observations of ASKAP repeaters (James et al. 2020a) for the parameter values investigated here.

of repetition in only one of 27 FRBs detected by ASKAP (Kumar et al. 2019; James et al. 2020b) to constrain $\gamma_r < -1.94$ (those authors use ζ for γ_r) and $R_{min} < 10^{-2.9}$ d^{-1}. Limits on R_{max} are not published, though values of $R_{max} \leq 100$ are investigated.

The values of R_{min} and R_{max} used in that work are applicable to 1.3 GHz observations and are measured relative to an energy threshold of 10^{38} erg. I therefore scale to a threshold of 10^{39} erg by reducing those rates by a factor of $10^{\gamma} \approx 0.13$, and to the mean CHIME frequency of 600 MHz by increasing the rates by a factor of $(600/1300)^{\alpha} \approx 4.4$, for a total adjustment factor of 0.55.

I extract the likelihoods from that work at the values of R_{min}, R_{max}, and γ_r used here and plot the Bayesian posterior in Fig. 21. The best-fit region agrees remarkably with results from CHIME FRB data, but has a much tighter constraint on γ_r. That two very different measurements with very different instruments converge to the same properties of the repeating FRB population is strong evidence that this model is a reasonable approximation to the underlying truth. This has implications for FRB progenitor models, as discussed in James et al. (2020a).

The one major discrepancy between the results which is hidden by Fig. 19 however is that the best-fit values of R_{min} found by James et al. (2020a) are a factor of $\sim$100 lower than that found for CHIME. Equivalently, the results of James et al. (2020a) would under-predict the number of repeating CHIME FRBs. There may be several causes for this.

Firstly: the results for James et al. (2020a) assumed a Weibull distribution with $k = 0.34$. As discussed in Section 8.1, a bursty distribution requires less-rapid repeaters to produce the same number of repeating FRBs and could therefore allow the values of R_{min} found here to be lower and more consistent with the ASKAP results. Secondly: the results here are only weakly constraining on R_{min}, and only the probabilities at the best-fit values have been used, rather than marginalising over R_{min}. Thirdly: repetition behaviour at 1.3 GHz and 600 MHz might be more different than the simple scaling above would suggest. A fuller investigation will require an improved model of FRB time–frequency structure. Fourthly: the limits from James et al. (2020a) used non-localised repeaters, with conservatively large distance estimates from the Macquart relation assuming no host contribution, thus introducing a bias in those results (albeit one which would push R_{min} and R_{max} to lower values). Fifthly and finally, not all FRBs may repeat. The upper limit on R_{min} found by James et al. (2020a) is driven
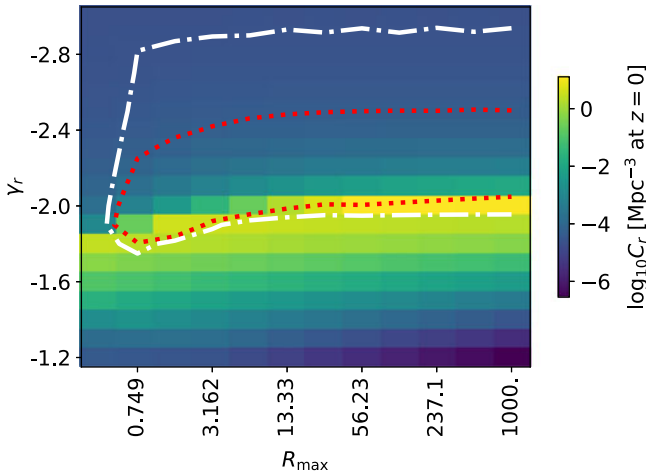
Figure 22. FRB progenitor population density C_r, as a function of R_{max} and γ_r with 68% and 95% contours from Fig. 14 overplotted.

primarily by the lack of observed repetition from FRB 20171020A. If this FRB is intrinsically a once-off event, then those results weaken significantly, allowing higher values of R_{min}.

The current implementation of repetition in the zDM code does not allow for easy estimation of FRB repetition parameters from follow-up observations. It would be useful to develop such a method to allow the results of follow-up observations to also be fit in a self-consistent manner.

9.2. Absolute number of repeaters and persistent radio sources

The total number of repeating FRBs, C_r, is a function of their rate of birth and active lifetime. Estimates of their birth rate vary greatly according to their progenitor model and range from 10^5 Gpc^{-3} yr^{-1} for core-collapse supernova (Taylor et al. 2014), to 0.02 Gpc^{-3} yr^{-1} for NS-NS binary mergers in globular clusters (Ye et al. 2020). Their lifetime is also unknown: while magnetar fields are expected to decay on timescales of order 10^4 yr (Colpi, Geppert, & Page 2000), precisely how this relates to their time as an active emitter of FRBs will depend on the microphysics of emission.

Fig. 22 plots the implied value of C_r as a function of R_{max} and γ_r, compared to limits on those parameters from Fig. 14. By far the largest number of FRB progenitors are found near $\gamma_r = -2$. At flatter (less negative) values of γ_r, the repeater distribution is dominated by so many rapidly repeating FRBs that very few are required, while at steep (more negative) values, R_{min} must be relatively high (and hence C_r not too large) to prevent singly detected objects overwhelming the total FRB rate. That this region is just allowed within the current limits means that the total uncertainty on C_r is huge, ranging approximately between 10^{-5} and 10 Mpc^{-3} at $z = 0$. Assuming a lifetime of 10^4 yr, this corresponds to birth rates of 1–10^6 Gpc^{-3} yr^{-1}, which is compatible with most FRB progenitor models.

The total number of strong repeating FRBs is also of interest. Aside from their obvious identification via detection of their bursts, these objects might be identified via their association with PRS, which can be identified in radio surveys. There is some evidence that PRS are more likely to be associated with strong repeaters (Chatterjee et al. 2017; Niu et al. 2022), though this evidence is not conclusive (Law et al. 2022).

In this context, Law et al. (2022) estimate the total number density, and differential power-law slope, of the observed repeating FRB population using CHIME data, as in Section 6.3. Assuming all CHIME FRBs come from intrinsic repeaters, they use the observed distribution of the number of bursts to have a power-law slope of $= -1.5$, calculated via the maximum-likelihood estimator of Crawford, Jauncey, & Murdoch (1970). This method thus determines the slope of the 'source counts' distribution of repeaters, which is largely insensitive to the intrinsic value of γ_r (see Fig. 12). They also estimate the mean number of observed repeat bursts per source, $\langle R_{obs} \rangle = 1.2$ d^{-1}, assuming a minimum repetition rate equal to the inverse of T_{Cat1}^{-1}. Again, this estimate is based on the observed repetition rates, rather than the intrinsic rates, which explains why the derived mean rate is somewhat higher than the values of R^* found here. Their derived number density of repeaters, being between 22 and $5.2 \cdot 10^3$ Gpc^{-3} in the no-beaming case, are both more constraining and on-average lower than derived here. Law et al. (2022) estimate that the fraction of FRB sources with associated PRS is between 0.06 and 0.36: the much higher uncertainty on the intrinsic source density found in this work suggests that surveys identifying PRS independently of FRBs might be a better tracer of the intrinsic number of FRB sources.

It should be noted that the estimates above assume isotropic emission. However, to first order, beaming should have no effect. Including beaming means that the intrinsic number of bursts per repeater is greater by $4\pi/\Omega_{frb}$, where Ω_{frb} is the solid angle subtended by each FRB. However, so is the intrinsic single-burst emission rate. Hence, the implied number of repeaters to reproduce the detected rate is identical, and beaming angle has no effect. Only in the case that Ω_{frb} varies with repetition rate R, as suggested by Connor et al. (2020), is this scaling broken.

10. Conclusions

I have implemented a model for repeating FRBs in the zDM code, allowing for a power-law distribution of intrinsic FRB repetition rates. A population of repeating FRBs will result in an apparent deficit of progenitors (single and repeat FRBs) in the nearby Universe, with the effect becoming increasingly strong with observation time per pointing. I show that this effect is significant for current observations with ASKAP, FAST, and CHIME, and hence that future FRB population modelling should include a simultaneous fit to FRB repetition parameters as well as the current set of cosmological parameters, host galaxy properties, population evolution, and the luminosity function. Such a fit is computationally infeasible with current methods implemented in the zDM code, and hence nested sampling techniques should be implemented in the future.

I have therefore fit a power-law model of repeating FRB rates, with differential slope γ_r between rates R_{min} and R_{max} (defined as bursts per day above 10^{39} erg, with a Poissonian distribution of arrival times), to CHIME Catalog 1 data. The model of the CHIME experiment includes beamshape, DM response, and declination-dependent exposure. This model can accurately reproduce the distribution of singly detected CHIME FRBs when using FRB population parameters with a steep dependence of FRB rate on frequency (rate$\propto \nu^{-1.91}$), consistent with previous fits to ASKAP and Parkes data at the 90% confidence level. Holding this parameter set fixed, I find that the distribution of repeating FRBs in DM,

δ, and N_{burst} space, as well as their number compared to apparently once-off bursts, is well reproduced when assuming the entire FRB population is explained by a single population of repeating FRBs with $\gamma_r = -2.2^{+0.6}_{-0.8}$ (68% CI), and $R_{\mathrm{max}} \geq 0.75$. Limits on R_{min} are less well constrained. In particular, results are consistent with the DM deficit of repeating FRBs found by CHIME/FRB Collaboration et al. (2023). This remains the case unless less than 50% of all CHIME single FRBs are due to intrinsic repeaters, at which point the predicted DMs of CHIME repeaters become too high to be consistent with data.

I also make predictions for the effects of repetition in current and future experiments. Localising CHIME repeating FRBs to obtain their redshifts would provide strong constraints on the repeating FRB population, and I urge that optical follow-up observations preferentially target these sources. Furthermore, I predict that the number of repeating FRBs identified by CHIME will continue to increase for at least the next hundred years' worth of observations and potentially for a much longer timespan. This is currently consistent with the new sample of repeaters released by CHIME/FRB Collaboration et al. (2023), but systematic effects in that dataset make more precise comparisons difficult.

An estimate is made of the systematic effects of correlations between repetition rates and FRB width, non-Poissonian arrival times, and differences between the full CHIME response to FRBs described by Merryfield et al. (2022) and used by Shin et al. (2023) and the method used here. Thus burstiness may be responsible for the apparent excess of repeating FRBs at high declinations, but that other effects are likely small, though may alter the best-fit value of γ_r away from its true value. Improved modelling of CHIME may be required in the future, though this current implementation already accounts for most effects described by CHIME/FRB Collaboration et al. (2023).

The most significant outcome is that limits on repetition parameters derived here agree with estimates of repeating FRB parameters produced from follow-up observations to ASKAP FRBs, despite the large differences in detection systems and method of estimation.

In conclusion, I emphasise that while this work is not conclusive evidence for the entire FRB population being explained by a single population of intrinsically repeating progenitors with a broad distribution of repetition rates, it certainly shows that current observations of repeating and single FRBs by CHIME are completely consistent in terms of DM, declination, number of bursts from repeaters, and the relative number of FRBs observed once and multiple times with this scenario.

References

Anderson, T. W., & Darling, D. A. 1954, JASA, 49, 765
Bannister, K. W., et al. 2019, Sci, 365, 565
Baptista, J., et al. 2023, arXiv e-prints, arXiv:2305.07022
Bhandari, S., et al. 2020, ApJ, 895, L37
Bhandari, S., et al. 2022, AJ, 163, 69
Bhardwaj, M., et al. 2021a, ApJ, 919, L24
Bhardwaj, M., et al. 2021b, ApJ, 910, L18
Caleb, M., Stappers, B. W., Rajwade, K., & Flynn, C. 2019, MNRAS, 484, 5500
Chatterjee, S., et al. 2017, Natur, 541, 58
CHIME/FRB Collaboration, et al. 2018, ApJ, 863, 48
CHIME/FRB Collaboration, et al. 2020, Natur, 582, 351
CHIME/FRB Collaboration, et al. 2021, ApJS, 257, 59
CHIME/FRB Collaboration, et al. 2019, ApJ, 885, L24
CHIME/FRB Collaboration, et al. 2023, ApJ, 947, 31
Colpi, M., Geppert, U., & Page, D. 2000, ApJ, 529, L29
Connor, L., Miller, M. C., & Gardenier, D. W. 2020, MNRAS, 497, 3076
Cordes, J. M., & Lazio, T. J. W. 2002, ArXiv Astrophysics e-prints, astro-ph/0207156
Cordes, J. M., & McLaughlin, M. A. 2003, ApJ, 596, 1142
Crawford, D. F., Jauncey, D. L., & Murdoch, H. S. 1970, ApJ, 162, 405
Falcke, H., & Rezzolla, L. 2014, A&A, 562, A137
Fong, W.-F., et al. 2021, ApJ, 919, L23
Fonseca, E., et al. 2020, ApJ, 891, L6
Gajjar, V., et al. 2018, ApJ, 863, 2
Gardenier, D. W., Connor, L., van Leeuwen, J., Oostrum, L. C., & Petroff, E. 2021, A&A, 647, A30
Gardenier, D. W., van Leeuwen, J., Connor, L., & Petroff, E. 2019, A&A, 632, A125
Gordon, A. C., et al. 2023, ApJ, 954, 40
Heintz, K. E., et al. 2020, ApJ, 903, 152
Hessels, J. W. T., et al. 2019, ApJ, 876, L23
Hewitt, D. M., et al. 2022, MNRAS, 515, 3577
Hunter, J. D. 2007, CSE, 9, 90
Ibik, A. L., et al. 2023, arXiv e-prints, arXiv:2304.02638
James, C. W. 2019, MNRAS, 486, 5934
James, C. W., Prochaska, J. X., & Ghosh, E. M. 2021, zdm, https://zenodo.org/record/5213780#.YRxh5BMzZKA
James, C. W., et al. 2022 a, MNRAS, 510, L18
James, C. W., et al. 2022 b, MNRAS, 509, 4775
James, C. W., et al. 2020 a, ApJ, 895, L22
James, C. W., et al. 2020 b, MNRAS, 495, 2416
James, C. W., et al. 2022 c, MNRAS, 516, 4862
Jankowski, F., et al. 2023, MNRAS, 524, 4275
Kirsten, F., et al. 2022, Natur, 602, 585
Kirsten, F., et al. 2023, arXiv e-prints, arXiv:2306.15505
Kolmogorov, A. 1933, G. Ist. Ital. Attuari. 4, 83
Kumar, P., et al. 2019, ApJ, 887, L30
Law, C. J., Connor, L., & Aggarwal, K. 2022, ApJ, 927, 55
Law, C. J., et al. 2017, ApJ, 850, 76
Lee-Waddell, K., et al. 2023, PASA, 40, e029
Li, D., et al. 2018, IEEE MM, 19, 112
Li, D., et al. 2021, Natur, 598, 267
Lin, H.-H., et al. 2023, arXiv e-prints, arXiv:2307.05261
Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, Sci, 318, 777
Lu, W., & Kumar, P. 2016, MNRAS, 461, L122

Luo, R., et al. 2020a, MNRAS, 494, 665
Luo, R., et al. 2020b, Natur, 586, 693
Mahony, E. K., et al. 2018, ApJ, 867, L10
Marcote, B., et al. 2017, ApJ, 834, L8
Marcote, B., et al. 2020, Natur, 577, 190
Merryfield, M., et al. 2023, AJ, 165, 11
Michilli, D., et al. 2018, Natur, 553, 182
Michilli, D., et al. 2021, ApJ, 910, 147
Michilli, D., et al. 2023, ApJ, 950, 12
Moroianu, A., et al. 2023, NatAs, 7, 579
Nimmo, K., et al. 2023, MNRAS, 520, 2281
Niu, C.-H., et al. 2021, ApJ, 909, L8
Niu, C. H., et al. 2022, Natur, 606, 873
Oppermann, N., Yu, H.-R., & Pen, U.-L. 2018, MNRAS, 475, 5109
Panther, F. H., et al. 2023, MNRAS, 519, 2235
Pastor-Marazuela, I., et al. 2021, Natur, 596, 505
PLANCK Collaboration, et al. 2020, A&A, 641, A6
Pleunis, Z., et al. 2021, ApJ, 923, 1
Price, D. C., et al. 2019, MNRAS, 486, 3636
Rajwade, K. M., et al. 2020, MNRAS, 495, 3551
Ravi, V. 2019, NatAs, 405
Rhee, J., et al. 2023, MNRAS, 518, 4646
Shannon, R. M., et al. 2018, Natur, 562, 386
Shin, K., et al. 2023, ApJ, 944, 105
Smirnov, N. 1948, AMS, 19, 279
Spitler, L. G., et al. 2014, ApJ, 790, 101
Spitler, L. G., et al. 2016, Natur, 531, 202
Staveley-Smith, L., et al. 1996, PASA, 13, 243
Taylor, M., et al. 2014, ApJ, 792, 135
Thornton, D., et al. 2013, Sci, 341, 53
van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, CSE, 13, 22
Virtanen, P., et al. 2020, NM, 17, 261
Woods, P. 2023, NatAs, 7, 374
Ye, C. S., et al. 2020, ApJ, 888, L10
Zhang, B. 2014, ApJL, 780, L21
Zhang, G. Q., et al. 2021, ApJ, 920, L23
Zhang, Y.-K., et al. 2022, RAA, 22, 124002
Zhang, Y.-K., et al. 2023, ApJ, 955, 12

Appendix A. Alternative beamshape modelling

To evaluate the effect of different beam approximations on CHIME's response to repeating FRBs, I use the 'strong repeaters' model of Section 3, and assume a source at an example declination of $\delta = 30°$. Three methods of parameterising the beam are considered. The first, and simplest, is to calculate a time-weighted effective beam sensitivity using

$$t'_{\text{eff}} = \sum_i t_i \overline{B}_i^{1.5} \qquad (A1)$$

which weights each frequency-averaged beam value, $\overline{B}$, by the cumulative source counts index of 1.5. This reduces the time spent observing a repeating FRB to a single effective time, t'_{eff}, at beam sensitivity $B = 1$. However, the characteristic sensitivity, B_{eff}, is better calculated by weighting these contributions by B, such that

$$B_{\text{eff}} = t'^{-1}_{\text{eff}} \sum_i \overline{B} t_i \overline{B}_i^{1.5}. \qquad (A2)$$

At that sensitivity, the effective time t_{eff} is greater than that at $B = 1$, that is,

$$t_{\text{eff}} = t'_{\text{eff}} B_{\text{eff}}^{-1.5}. \qquad (A3)$$
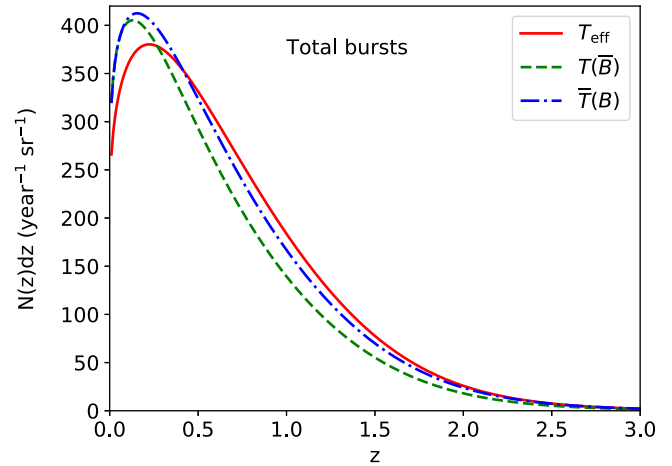
Figure A.1. Comparison of the total burst rate for three methods of calculating $T(B)$: a single effective value (T_{eff}), using the frequency-averaged beam $T(\overline{B})$, and averaging the time after calculations at each frequency $\overline{T}(B)$.

The second method parameterises $T(B)$ by histogramming the frequency-averaged beamshape $T(\overline{B})$, which is that standard method used in this work. The third method first creates a histogram for each frequency and then averages those to produce $\overline{T}(B)$. Here, 15 bins in B are used, equally log-spaced from 10^{-3} to unity.

The total expected number of bursts (per year per steradian) is given in Fig. A.1. Systematic errors due to the discretisation of $T(\overline{B})$ and $\overline{T}(B)$ are less than the line widths in the figure, whereas T_{eff} differs from $\overline{T}(B)$ purely because of discretisation to a single value. The difference between the latter two cases is clearly significant, with the use of a single characteristic value of sensitivity over-predicting the burst rate at high z, and underpredicting it at low z. As noted in James et al. (2022b), this is simply because using the T_{eff} method (method 1) is an over-simplification. However, the difference between $T(\overline{B})$ and $\overline{T}(B)$ is more subtle. When bursts are broadband, the instrumental response will be averaged over the total bandwidth, and hence the frequency-averaged sensitivity, $T(\overline{B})$, applies. However, when bursts exhibit limited band occupancy, the sensitivity at a specific frequency range is relevant. The total burst rate would then be better predicted by first calculating the time-average of the sensitivity in each fraction of the total bandwidth, that is, $\overline{T}(B)$. This latter method also leads to more bursts being detected when the slope of the cumulative source-counts ('logN-logS') distribution is steeper than -1: in a toy example of a beam with sensitivities 0.1 and 0.9 at two different frequencies and Euclidean source counts, the rate using the $T(\overline{B})$ method will be $[0.5(0.1 + 0.9)]^{1.5} = 0.35$, while using the $\overline{T}(B)$ method, it will be $0.5(0.1^{1.5} + 0.9^{1.5}) = 0.41$.

CHIME/FRB Collaboration et al. (2021); Pleunis et al. (2021) have found that the band occupancy varies significantly from burst-to-burst, with some bursts exhibiting a broadband morphology, and others—particularly those from observed repeaters—being band-limited. However, the majority are not significantly band-limited, and thus $T(\overline{B})$ has been adopted as the default beam parameterisation.

Appendix B. Declination binning

In order to account for the declination-dependent exposure of CHIME, six bins in declination have been chosen, and the CHIME exposure averaged over each bin. The choice of the number of bins is somewhat arbitrary—here, I characterise the error made through averaging by comparing predictions made with a much larger number of declination bins.

For this test, the 'min α' model with the CHIME exposure from Table 1 is used, and calculations are performed using $N_\delta = 6$ and 30. Predictions for the declination dependence of the CHIME FRB populations are compared to data in Fig. B.1. Other models produce qualitatively similar behaviour.

From Fig. B.1, it can be seen that while $N_\delta = 30$ produces quantitatively different predictions for the δ-dependence of the repeating FRB population, differences between predictions and observations are generally larger, such that using a small N_δ won't affect the conclusions here. Thus I have used $N_\delta = 30$ only for display purposes in Fig. 10.
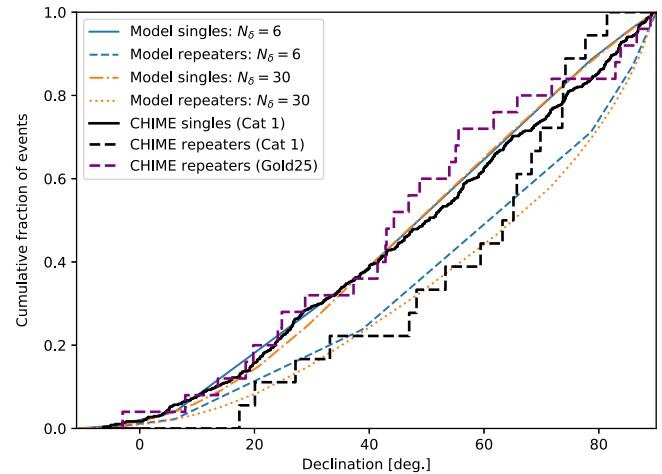
Figure B.1. Cumulative distribution of single (solid) and repeat (dashed) FRBs, as predicted from simulations with $N_\delta = 6$ (blue) and $N_\delta = 30$ (orange), with the 'min α' parameter set with CHIME's DM selection function. This is compared to CHIME Catalog 1 single and repeat bursts (black), and also repeat bursts including the 3-yr sample (purple).