Searches for gravitational waves from known pulsars: constraining the pulsar ellipticity distribution

Matthew Pitkin

University of Glasgow

Background

Known pulsars are a great target in searches for continuous gravitational wave signals.

Known frequency evolution and sky location allow coherent searches over long data stretches (~years) using (*almost*) a single phase template

Pulsars with periods accessible to the LIGO/Virgo gravitational wave detectors

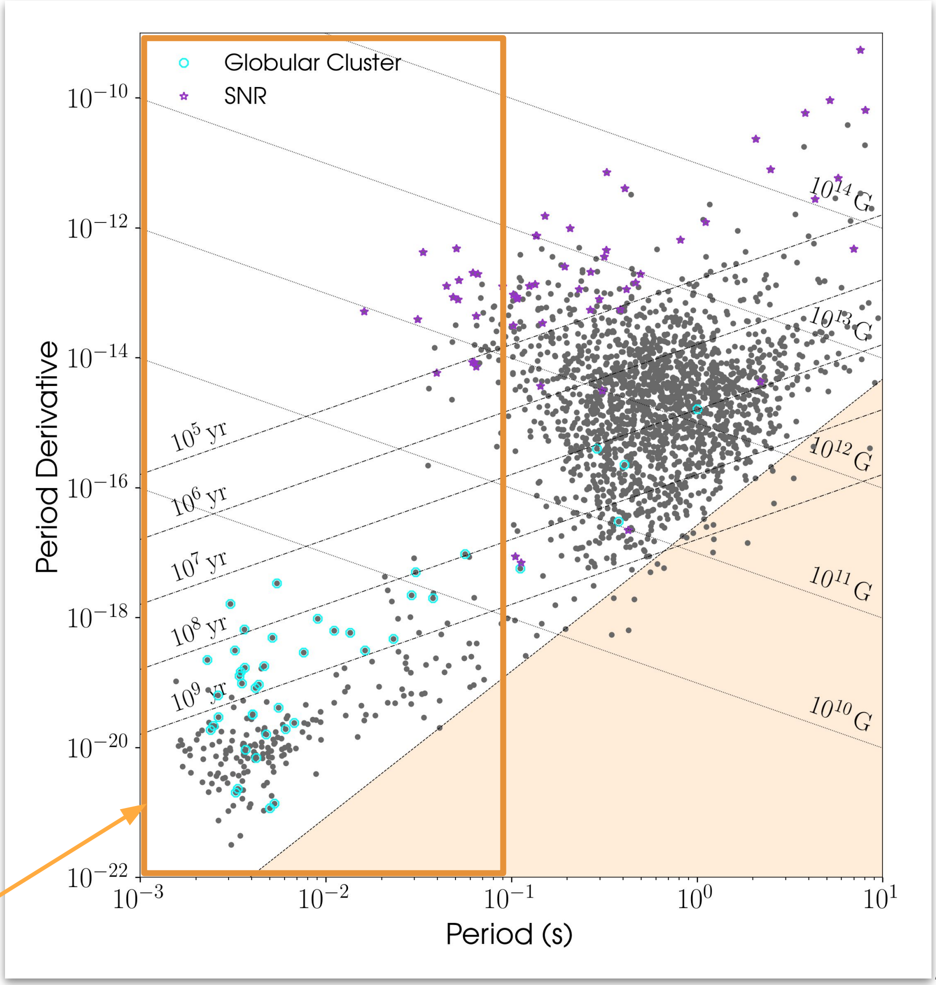

Image produced with psrqpy, Pitkin, JOSS (2018)

Searches

Previous known pulsar searches carried out by the LIGO Scientific Collaboration & Virgo Collaboration (LVC) have made the following assumption:

- signals are emitted from a triaxial star ($l=m=2$ mass quadrupole mode) rotating about its principal moment of inertia I_{zz} (no precession)
- GW signals are phase locked with the electromagnetic emission (which is itself locked to the star's rotation)† giving emission at twice the rotation rate f_{rot}
- each star's amplitude is independent (i.e. each target is independent of all others)

†Some targeted searches have been performed relaxing the very strong assumption about GW emission being phase locked to the rotation, e.g., Abbott *et al*, *ApJL*, 683 (2008) & Abbott *et al*, *PRD* 96, 122006 (2017)

Searches

For each pulsar, searches attempt to evaluate the probability distribution of:

- h_0: the gravitational wave strain detected at the Earth
- $\cos\iota$: the cosine of the inclination of the rotation axis to the line-of-site
- ϕ_0: the phase of the signal at some epoch
- ψ: the polarisation angle

When no signal is found an upper limit on h_0 can be set (often at 95% credible level). This can be compared to the *spin-down limit* set by assuming all rotational kinetic energy is dissipated through l=m=2 mass quadrupole GW emission:

$$h_0^{\rm sd} = 2.55\times10^{-25}\left(\frac{I_{zz}}{10^{38}\,{\rm kg\,m^2}}\right)^{1/2}\left(\frac{1\,{\rm kpc}}{d}\right)\left(\frac{100\,{\rm Hz}}{f_{\rm rot}}\right)^{1/2}\left(\frac{|\dot{f}_{\rm rot}|}{10^{-11}\,{\rm Hz}\,s^{-1}}\right)^{1/2}$$

LVC Searches

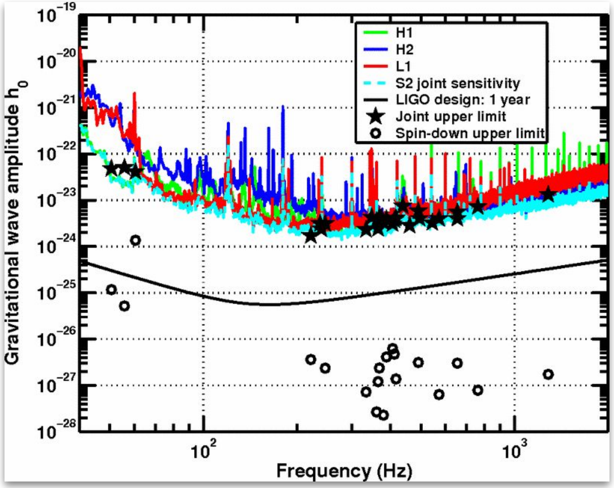

LIGO S2: 28 pulsars. Abbott *et al*, *PRL* 94, 181103 (2005)

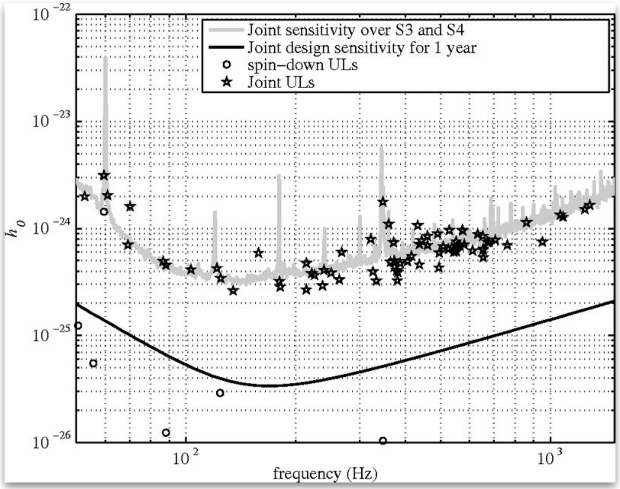

LIGO S3+S4: 78 pulsars. Abbott *et al*, *PRD* 76, 042001 (2007)

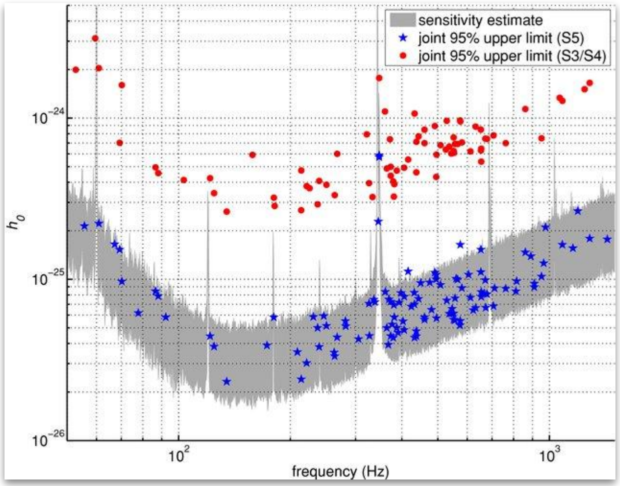

LIGO S5: 116 pulsars. Abbott *et al*, *ApJ*, 713 (2010)

Rely on up-to-date ephemerides from EM pulsar observations (radio, X-ray, γ-ray) preferably overlapping GW observing runs.

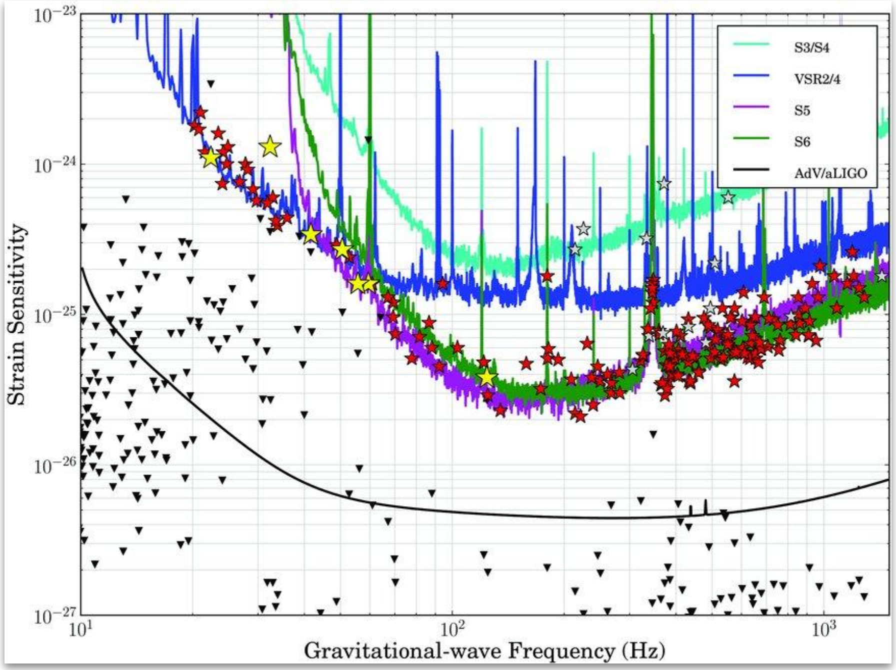

LIGO S6, Virgo VSR2/4 (+previous): 195 pulsars. Aasi *et al*, *ApJ*, 785 (2014)

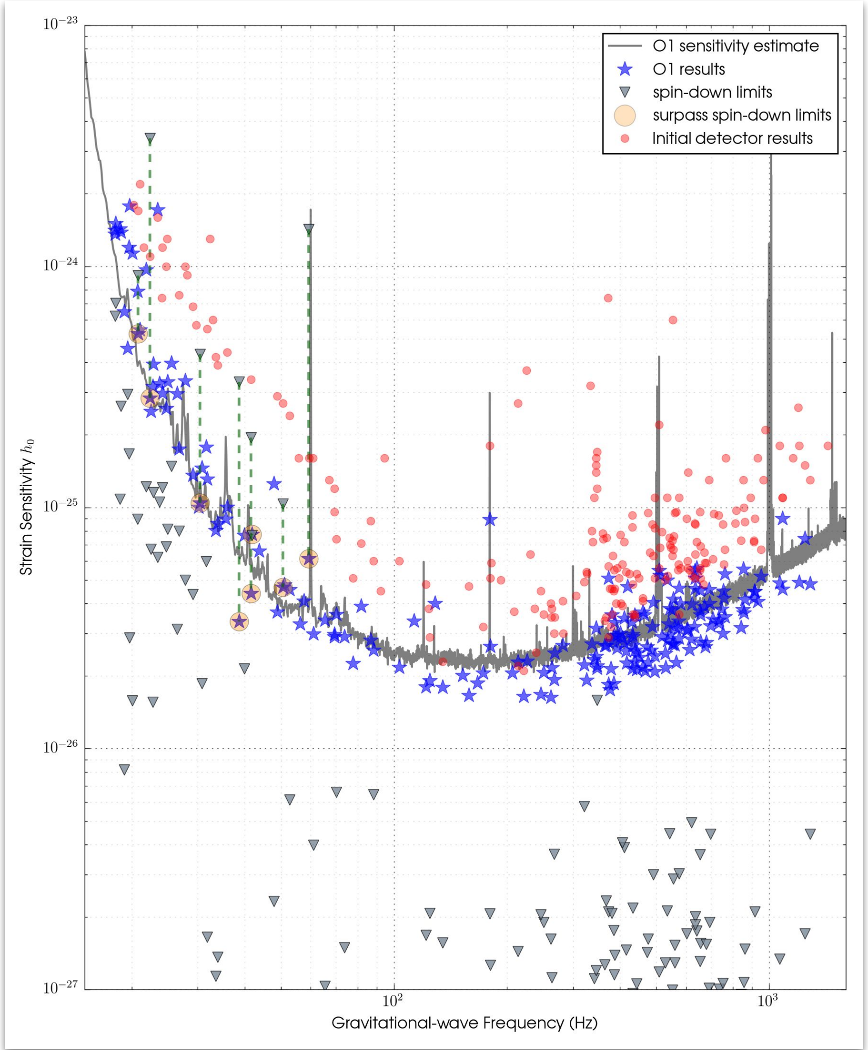

LIGO O1 200 pulsars. Abbott *et al*, *ApJ*, 839 (2017)

Searches

†see, e.g., Johnson-McDaniel & Owen, *PRD* 88, 044004 (2013)

The probability distribution of h_0 can be converted into a distribution on the mass quadrupole moment Q_{22}, or fiducial ellipticity† ε assuming a known distance (often known to ~20%):

$$\varepsilon = 10^{-6}\left(\frac{h_0}{4.2\times10^{-26}}\right)\left(\frac{100\,\mathrm{Hz}}{f_\mathrm{rot}}\right)^2\left(\frac{10^{38}\,\mathrm{kg\,m^2}}{I_{zz}}\right)\left(\frac{d}{1\,\mathrm{kpc}}\right)$$

and (following Ushomirsky, Cutler & Bildsten, *MNRAS* 319 (2000))

$$Q_{22} = \sqrt{\frac{15}{8\pi}}I_{zz}\varepsilon$$

This can in-turn be converted to a limit on the model-dependent internal B-field strength (e.g. Cutler, *PRD* 66, 084205 (2002) for toroidal field with $B < 10^{15}$ G):

$$\varepsilon_B \approx -1.6\times10^{-6}\left(\frac{\langle B_t\rangle}{10^{15}\,\mathrm{G}}\right)$$

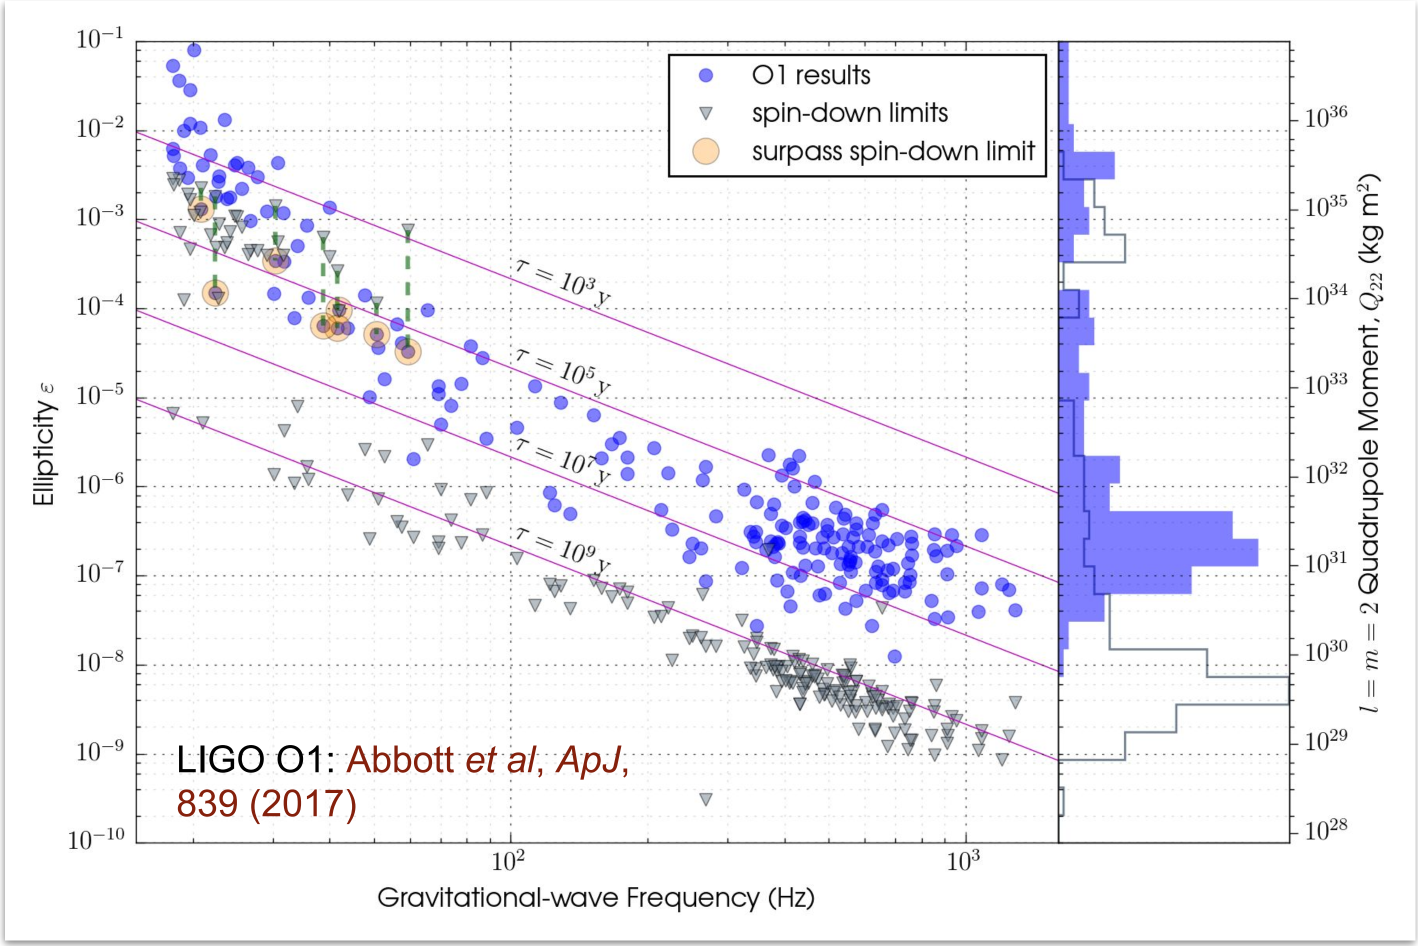

We can convert to surface deformation, maximised over EoS, using†:

$$R\varepsilon_{\mathrm{surf},22} \approx 25 \left(\frac{\varepsilon}{10^{-4}}\right) \mathrm{cm}$$

† Johnson-McDaniel, *PRD* 88, 044016 (2013)

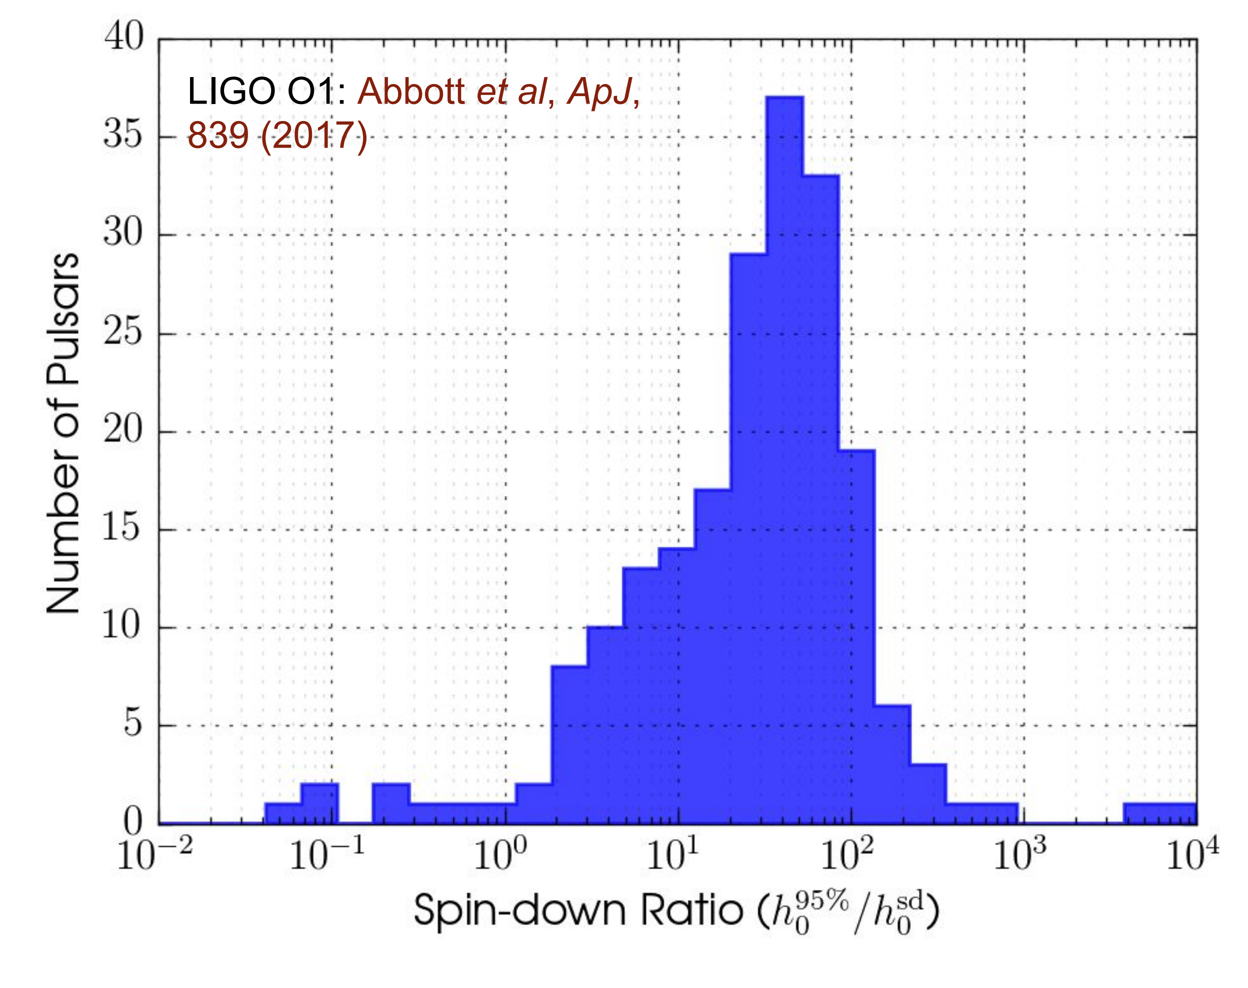

Smallest spin-down ratio: Crab pulsar, $\varepsilon < 3.6 \times 10^{-5}$, which is 20 times below the spin-down limit

MSP closest to spin-down limit: J0437-4715 (GW frequency 347 Hz, at 0.16 kpc) $\varepsilon < 2.8 \times 10^{-8}$, which is only 1.4 times spin-down limit. Limits toroidal B field to $\lesssim 10^{13}$ G

Narrow-band search

LIGO O1 data: 11 pulsars. Abbott *et al*, *PRD*, 96 122006 (2017)

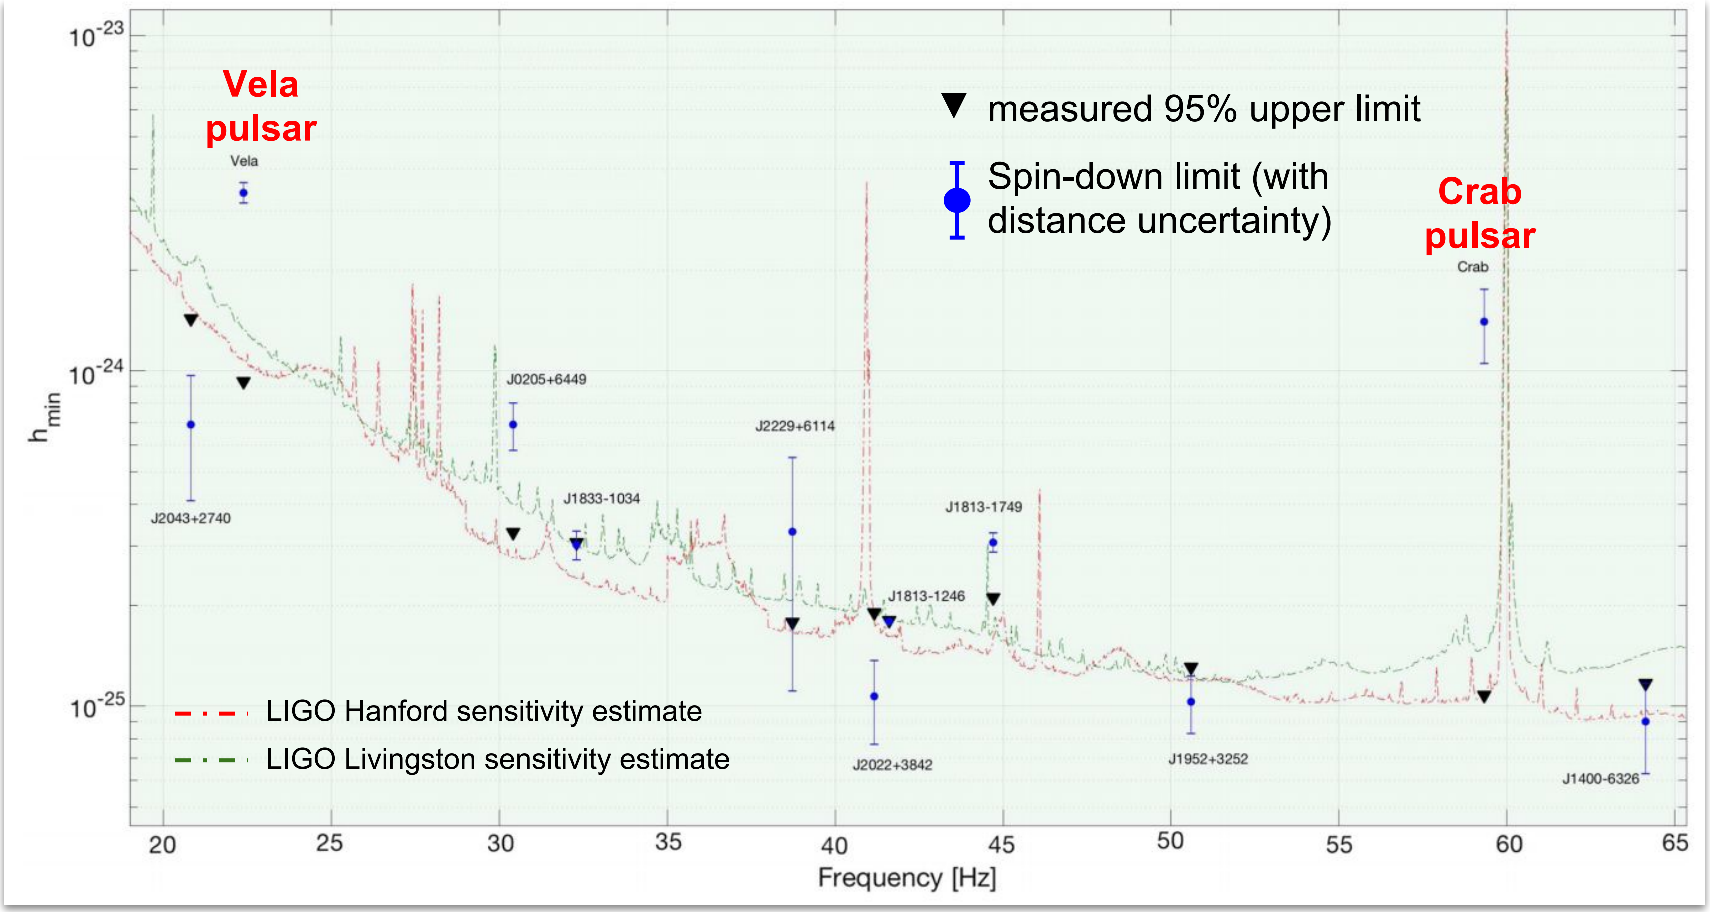

Ellipticity distribution (with C. Messenger & X. Fan)

Known pulsar searches assume each pulsar's amplitude is independent. However, we could instead assume that there is an underlying distribution from which the ellipticities are drawn, e.g., an **exponential** distribution

$$p(\varepsilon|\mu_\varepsilon, I) = \frac{1}{\mu_\varepsilon} e^{\varepsilon/\mu_\varepsilon}$$

defined by the *hyperparameter* μ_ε, the mean of the distribution.

We can combined data from all pulsars to estimate the probability distribution of μ_ε: ***hierarchical Bayesian inference*** (already used in GW field for black holes mass and spin distributions, e.g., Abbott *et al.*, *PRX* 6, 041015 (2016) & Stevenson, Berry & Mandel, *MNRAS*, 471 (2017)).

Ellipticity distribution

We form a **joint likelihood** of the data from all pulsars marginalised over independent $\cos\iota$, ψ, and ϕ_0 values for each pulsar, and also marginalising over the uncertainty on distance (assuming a 20% Gaussian error):

$$p(\mu_\varepsilon|\{\mathbf{D}\}, I) \propto \int_{\varepsilon_{\min}}^{\varepsilon_{\max}} \left(\prod_{i=1}^{N_{\text{pulsars}}} \int^{\theta_i} p(\mathbf{D}_i|\varepsilon, \theta_i, d_i, I) p(\theta_i|I) \boxed{p(d_i|I)} \mathrm{d}\theta_i \mathrm{d}d_i \right) \boxed{p(\varepsilon|\mu_\varepsilon, I) p(\mu_\varepsilon|I)} \mathrm{d}\varepsilon$$

The same kind of analysis could be used on results from a blind all-sky search by changing the prior on the distance

The probability distribution for the ellipticity distribution (and prior on hyperparameters) can be changed to your favourite function

where $\mathbf{D}_i$ is the data for each pulsar, $\theta_i = \{\phi_{0i}, \psi_i, \cos\iota_i\}$, and I assumes knowledge such as pulsar frequencies.

Calculating this for all pulsars at once would be a difficult problem (100s of parameters), but instead we can make use of posteriors distributions on ε (or equivalently Q_{22}) individually marginalised over other parameters.

Simulations

We generated independent realisations of signals from a population of 200 known pulsars (those searched for in O1 data) with ε drawn from exponential distributions with a range of μ_ε values, injected into Gaussian noise based on the aLIGO design sensitivity.

We also generated a set of “background” realisations, in which all pulsars had zero ellipticity (i.e. detector noise only).

For each population realisation we calculate the “evidence” that the data contains pulsars with ε drawn from an exponential distribution, and also that the population is consistent with noise.

Exponential mean, μ_ε, prior (Jeffreys)

$$p(\mu_\varepsilon|I) = \frac{(\ln \mu_{\varepsilon\mathrm{max}} - \ln \mu_{\varepsilon\mathrm{min}})^{-1}}{\mu_\varepsilon} \text{ if } \mu_{\varepsilon\mathrm{min}} < \mu_\varepsilon < \mu_{\varepsilon\mathrm{max}}$$

Bayesian evidence for population with ellipticities drawn from an exponential distribution

$$p(\{\mathbf{D}\}|H_{\mathrm{exp}}, I) = \int^{\mu_\varepsilon} p(\{\mathbf{D}\}|\mu_\varepsilon, H_{\mathrm{exp}}, I) p(\mu_\varepsilon|H_{\mathrm{exp}}, I) \mathrm{d}\mu_\varepsilon$$

Bayesian odds comparing evidence for exponential distribution to data consistent with noise

$$\mathcal{O} = \frac{p(\{\mathbf{D}\}|H_{\mathrm{exp}}, I)}{p(\{\mathbf{D}\}|H_{\mathrm{noise}}, I)}$$

Results

90% credible intervals from the posterior distributions on μ_ε for the simulated population distributions (orange/green intervals are for the ensemble with the largest/smallest odds)

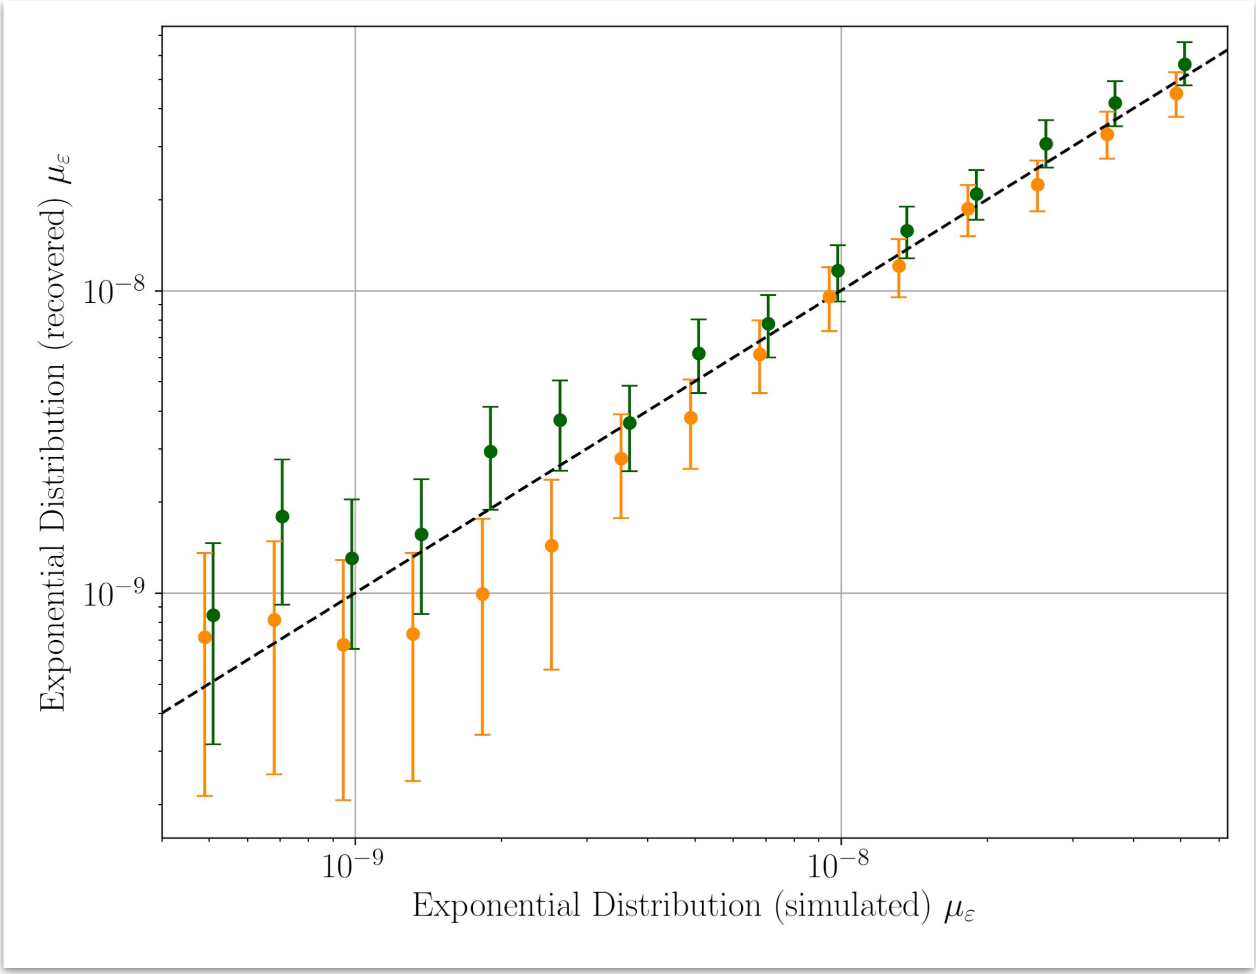

Results

Odds comparing a hypothesis that the data is consistent with signals having ellipticities drawn from an exponential distribution against all data being consistent with noise

Odds comparing a hypothesis that the data is consistent with *any* combination of pulsars containing a signal (signal amplitudes are independently marginalised over) against all data being consistent with noise

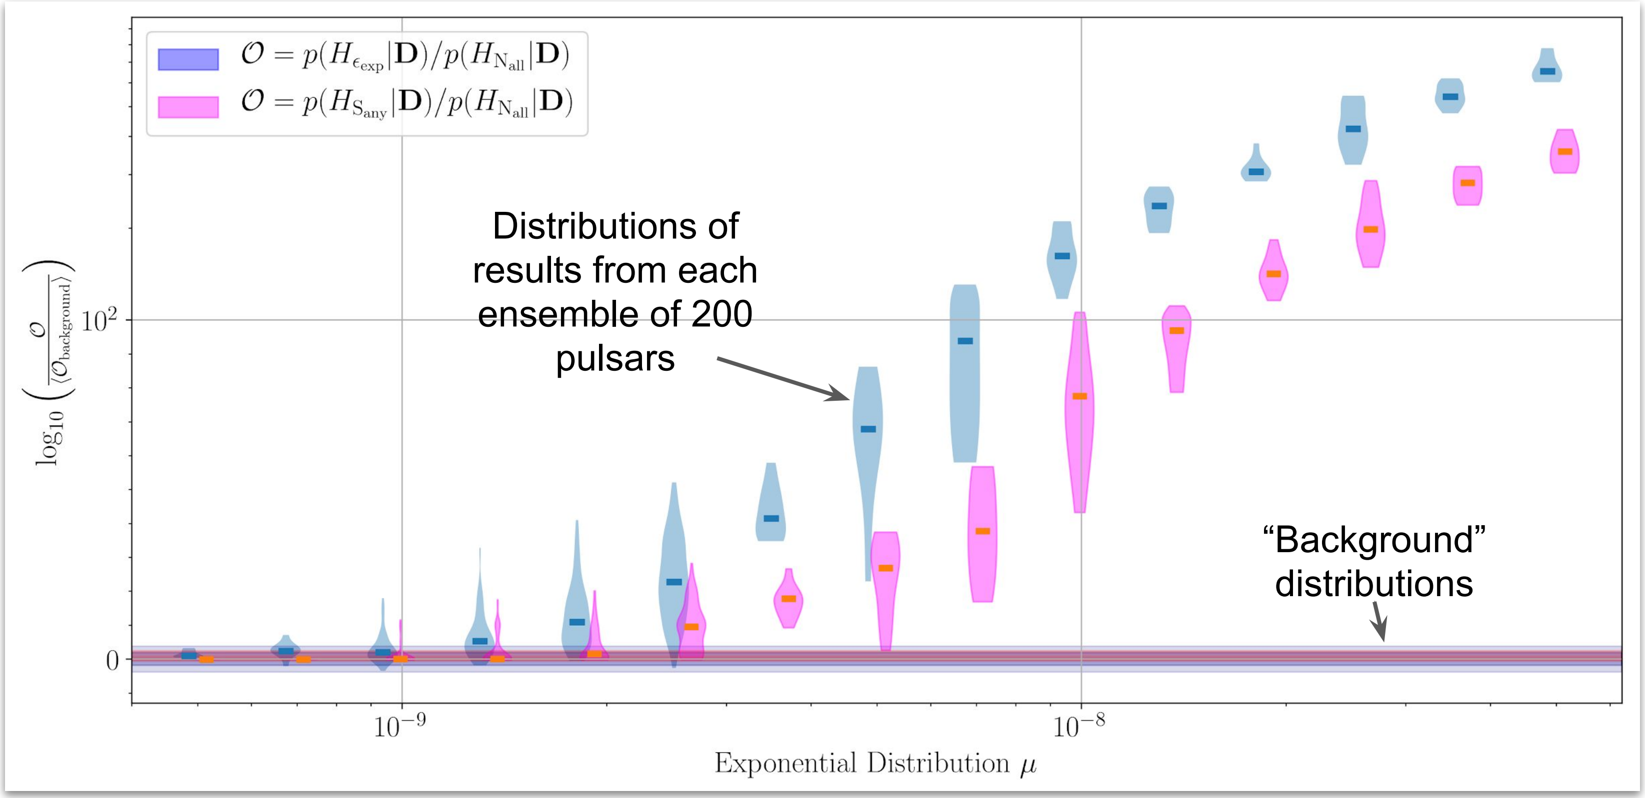

Results

SNR histograms for the pulsar ensemble producing the minimum/maximum odds for each simulated distribution

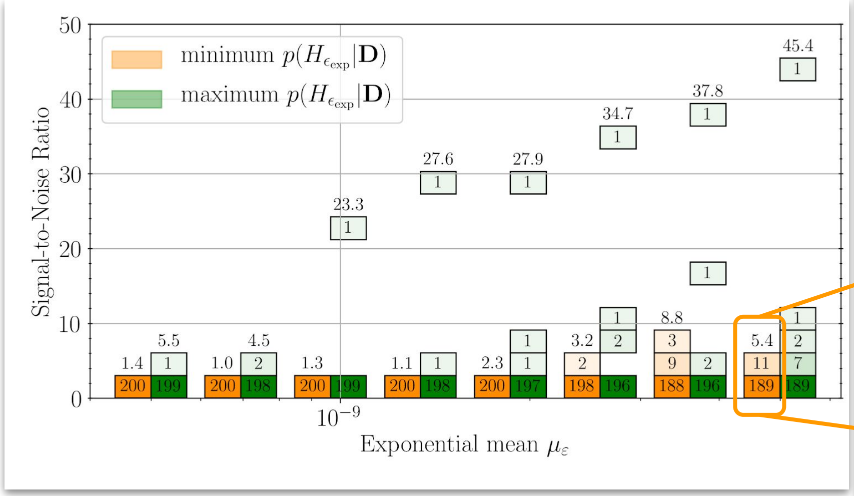

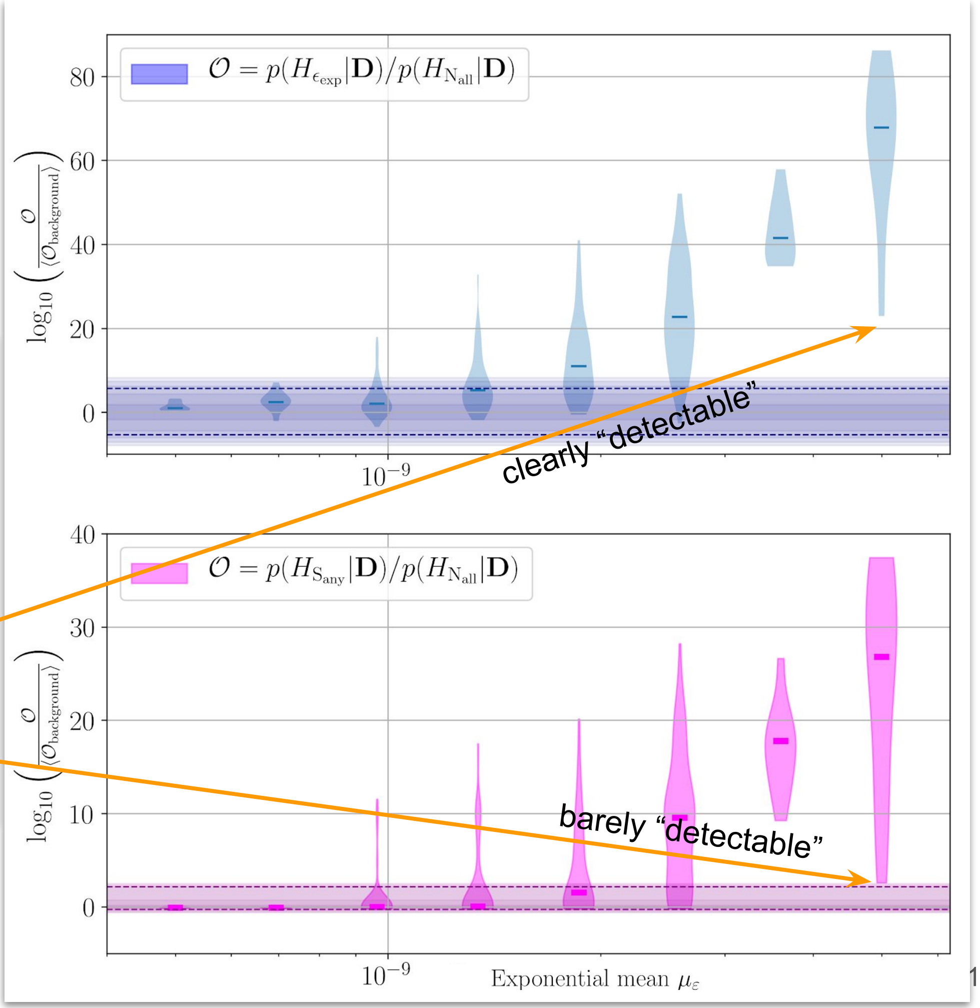

Results

Efficiency curves for detection of ensemble of pulsars.
Left: using a “false alarm probability” set by the number of background realisations.
Right: extrapolating the background to an equivalent “5σ” level (using a KDE of the background).

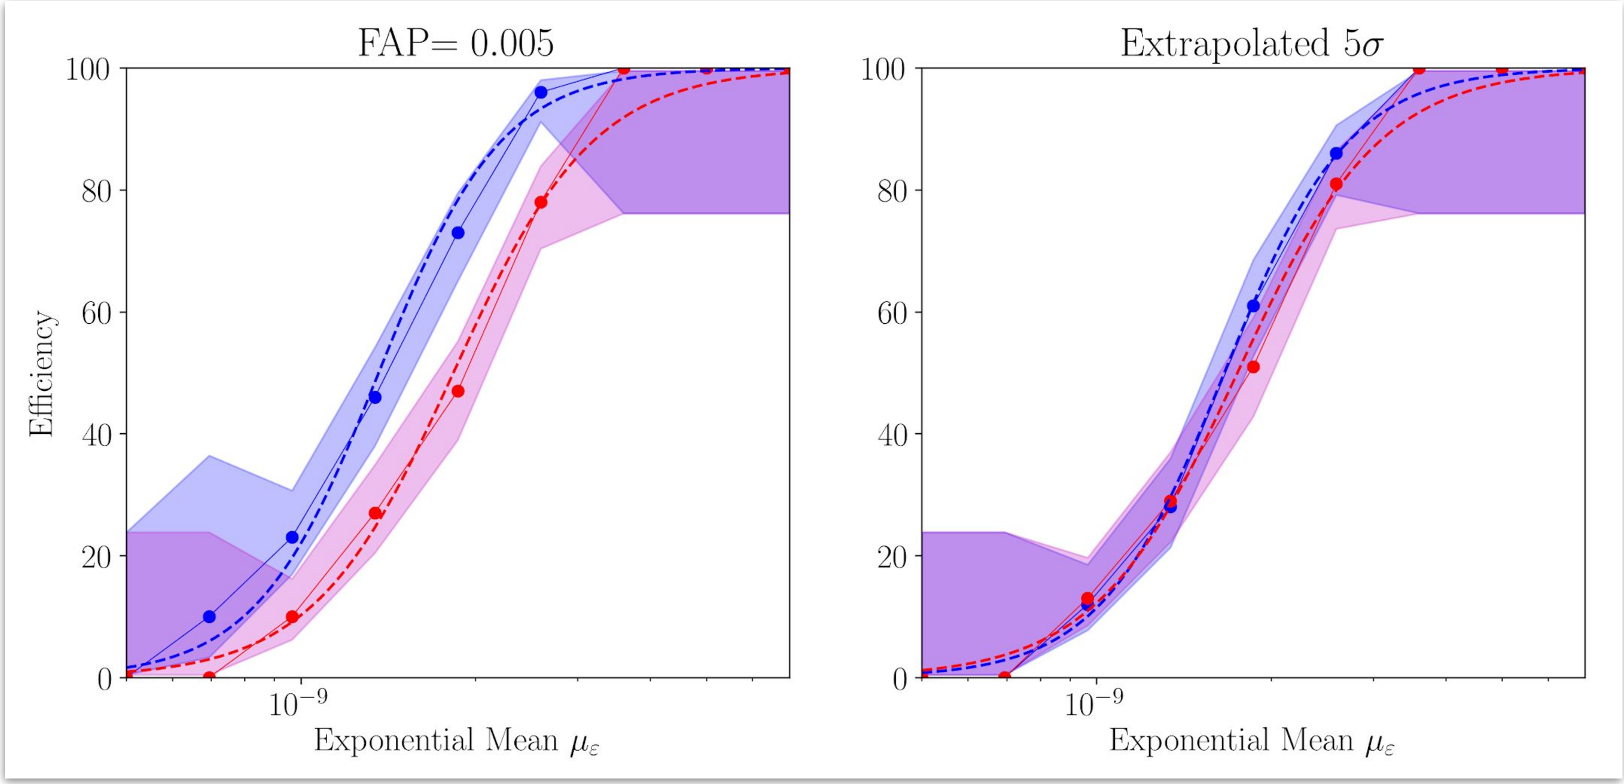

Results

What about a different distribution? We can assume the distribution is a half-Gaussian (using the same prior on σ_ε as we had for μ_ε)and compare models

$$p(\varepsilon|\sigma_\varepsilon, I) = \frac{2}{\sqrt{2\pi}\sigma_\varepsilon} e^{-\frac{1}{2}(\varepsilon/\sigma_\varepsilon)^2}$$

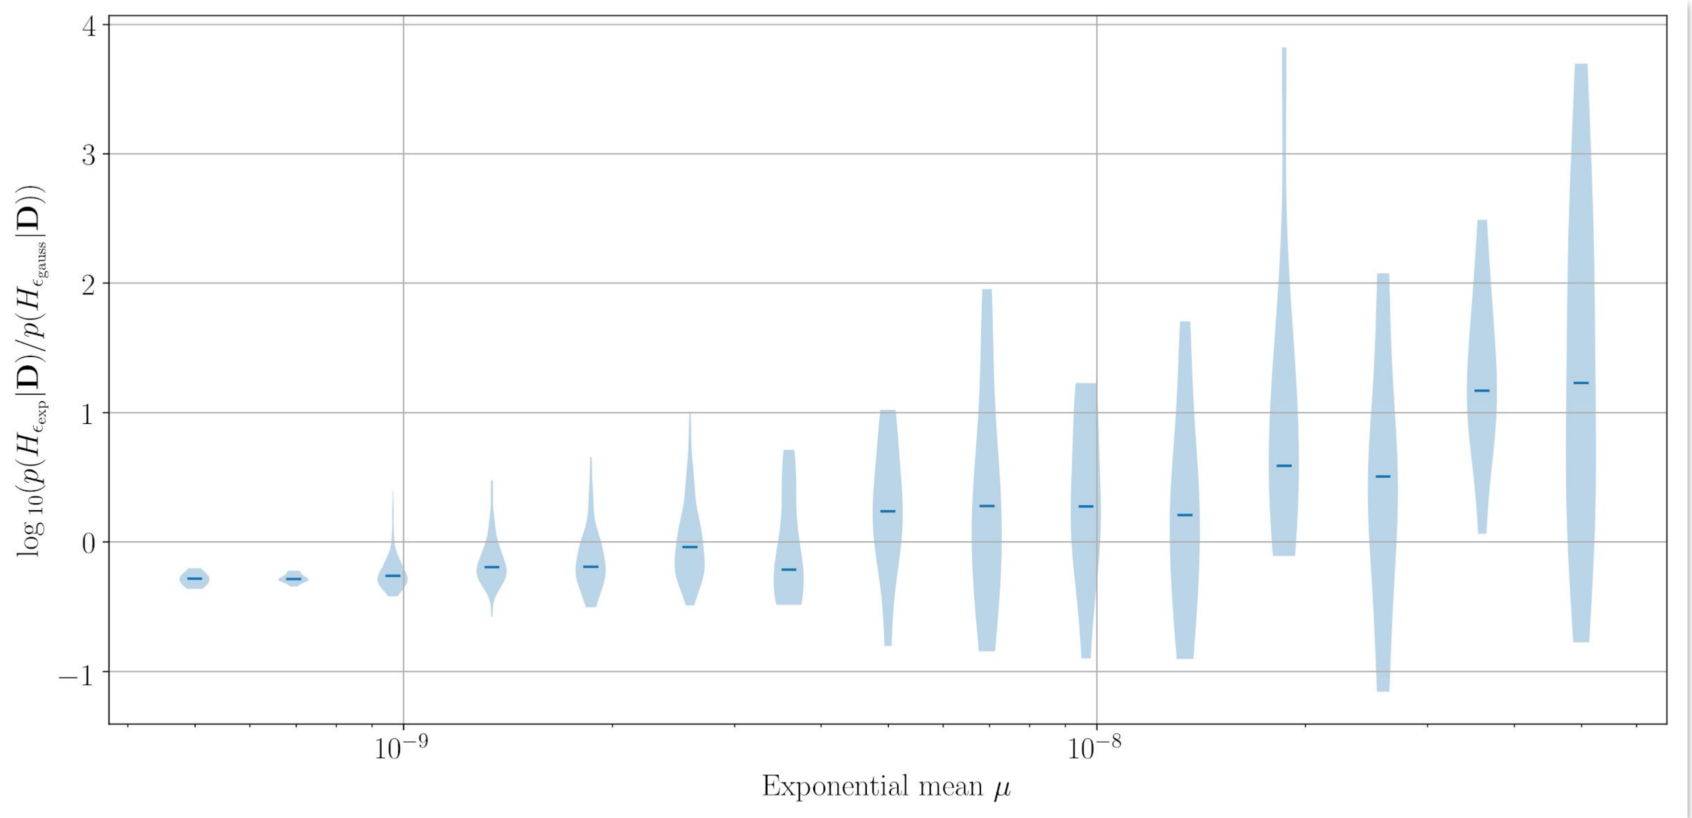

S6 results

Using 92 pulsars from the S6 known pulsar search (Aasi *et al*, *ApJ*, 785 (2014)) we applied the analysis assuming a 20% Gaussian error on pulsar distances (and upper bound on ε of ~10^{-4}). **No detection** of an ensemble of sources.

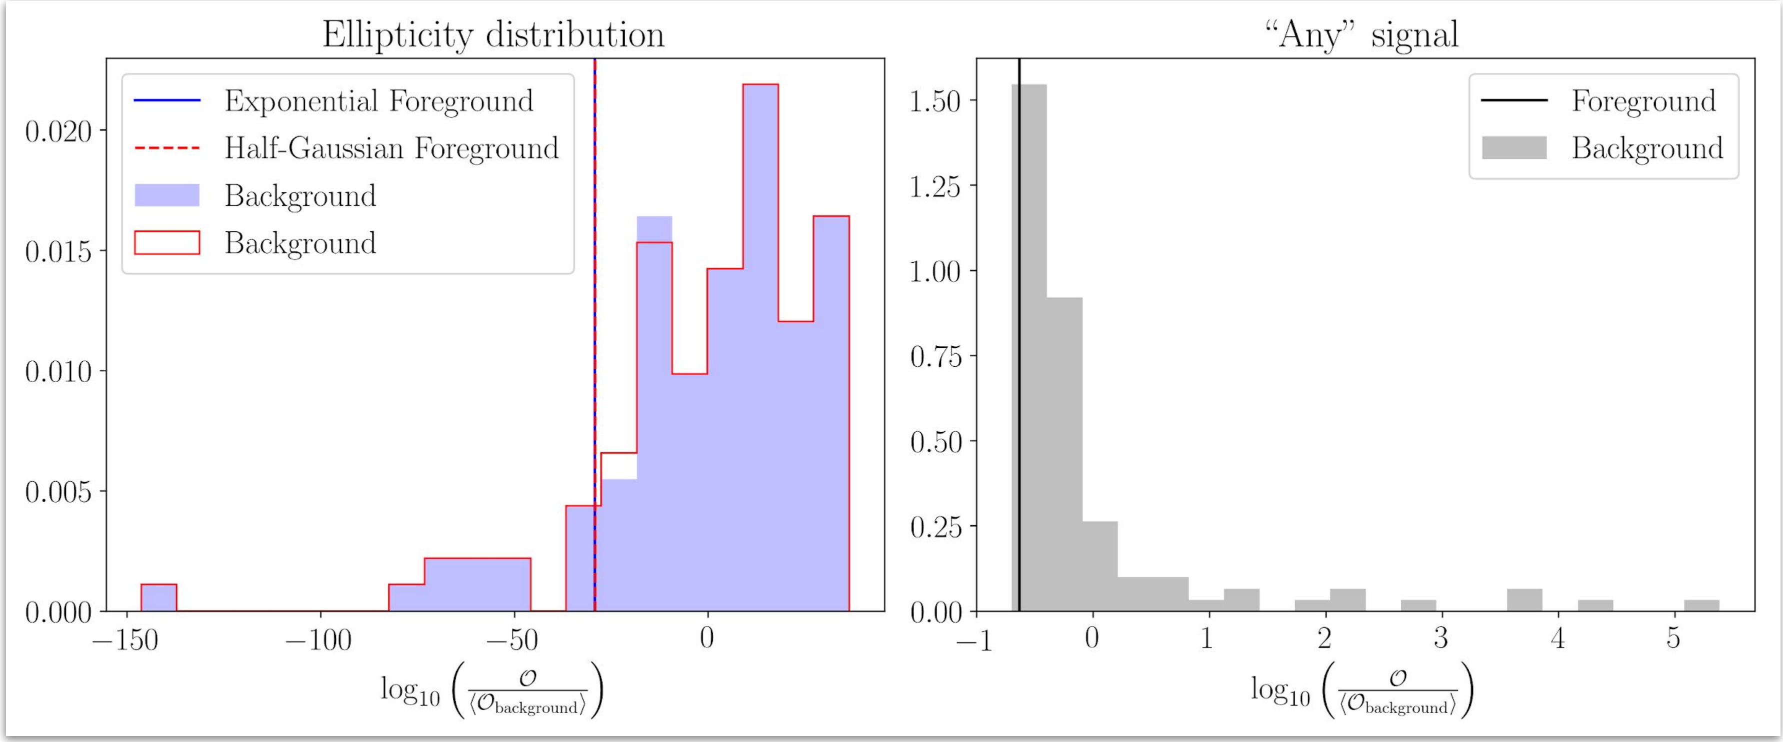

S6 results

Set 95% credible upper limits of:

- $\mu_\varepsilon < 5.4 \times 10^{-8}$ for an exponential distribution
- $\sigma_\varepsilon < 6.5 \times 10^{-8}$ for a half-Gaussian distribution

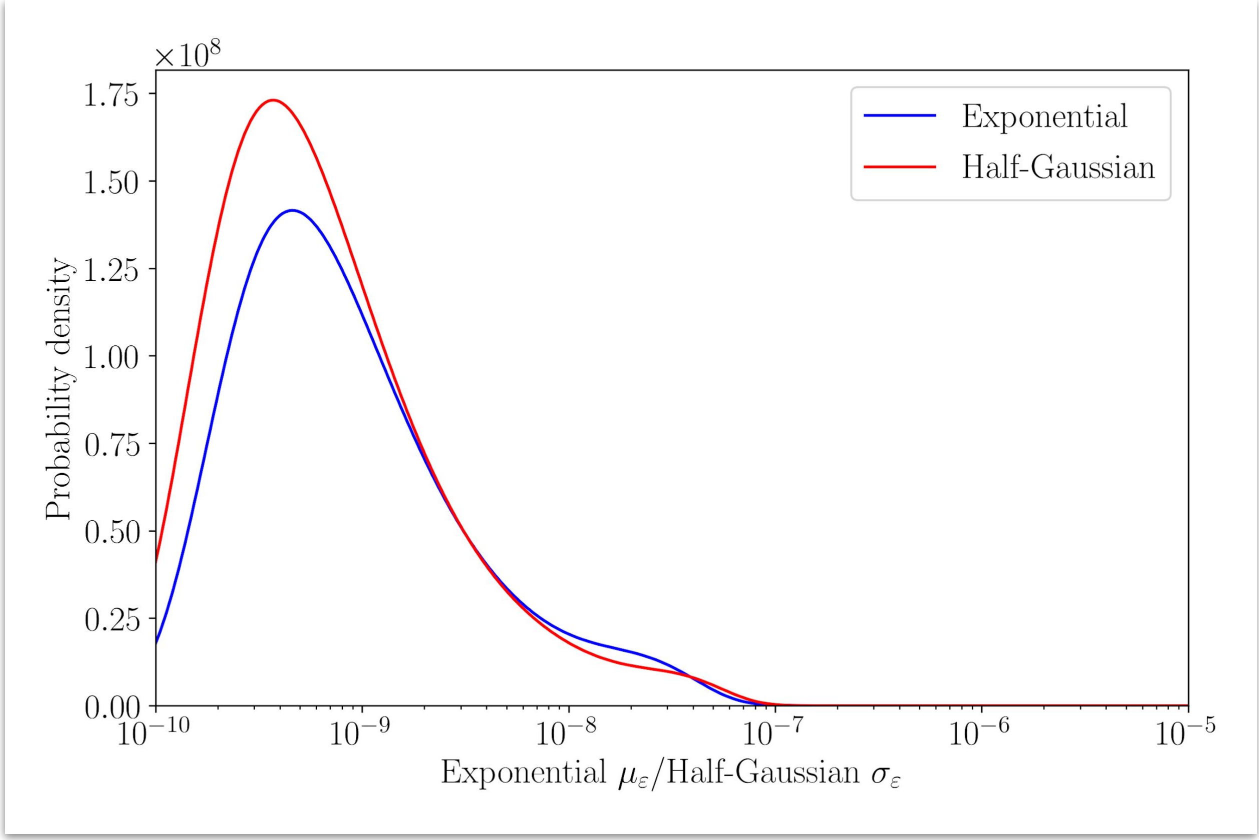

Conclusions

Searches for known pulsars provide tight constraints on neutron star ellipticity that are starting to compete with *spin-down limits*. Further searches are underway - we're currently looking at O2 data.

We can apply the method to constrain (a parameterised model-dependent) ellipticity distribution using results from real searches.

This could give us a new detection statistic for the ensemble of pulsars rather than relying on detecting individual sources (see also, e.g., Cutler & Schutz, *PRD* 72, 063006 (2005), Fan, Chen & Messenger, PRD 94, 084029 (2016) & Smith & Thrane, arXiv:1712.00688).

Questions

What might be a good model for the underlying distribution?

- Exponential?
- Power law?
- Gaussian mixture?
- Non-parametric (i.e. histogram, IGMM)?

Should we fit separate distributions for MSPs & young pulsars / globular cluster & field pulsars / binary pulsars and isolated pulsars?

Should/could we also constrain model-dependent internal B-field distributions?

What distance priors are reasonable if using sources from blind searches?