Probing the Low-Frequency Gravitational Wave Background with Fermi-LAT Pulsar Timing

Matthew Kerr (NRL)
& Aditya Parthasarathy (MPIfR)
on behalf of the Fermi-LAT Collaboration

Image credit: Danielle Futselaar (artsource.nl)

nHz Gravitational Waves from Binary Supermassive Black Holes

- Basic physics: GR (Einstein, 1918) $\frac{\mathrm{d}E_{\mathrm{GW}}}{\mathrm{d}t} = \frac{c^3}{16\pi G}\iint |\dot{h}|^2 \mathrm{d}S = \frac{1}{5}\frac{G}{c^5}\sum_{i,j=1}^{3}\frac{\mathrm{d}^3 Q_{ij}}{\mathrm{d}t^3}\frac{\mathrm{d}^3 Q_{ij}}{\mathrm{d}t^3}$

  + Mechanics ($Q \propto r^2$) + Kepler's Laws ($r^3 = \frac{GM}{(2\pi f)^2}$) lead to a simple relation:

  A circular binary with frequency $\frac{f}{2}$ emits GW with frequency f and amplitude $h \propto f^{\frac{2}{3}}$.

- For typical SMBH masses, GW emission dominates at binary separations of <0.1pc, i.e. orbital periods $\gg$ 10 years.

- Binaries spend more time at wide separations. Adding them up over the cosmological population and adding in the $\frac{1}{d}$ amplitude scaling (Phinney (2003), Sesana et al. (2004)) leads to the prediction of a power-law "stochastic" cosmological background of nanohertz GW with spectrum ($\alpha = -\frac{2}{3}$):

$$\boxed{h_c(f) = A_{\mathrm{gwb}}\left(\frac{f}{\mathrm{yr}^{-1}}\right)^{\alpha}}$$

Properties of the GWB
(see Burke-Spolaor et al. 2019 for a review)

- GW amplitude A_{gwb} scales with masses of merging black holes and the efficiency of forming BH binaries.
- BHs must reach center of merged galaxy and must shed angular momentum to close from 1pc to <0.1pc, the "Last Parsec Problem".
  - If BHBs don't solve the Last Parsec Problem efficiently, fewer binaries feed in at the low-frequency end. This produces a turnover in the power law and an overall reduction in amplitude.

- **Thus, detecting the amplitude A_{gwb} and ultimately, the shape of the GWB spectrum provides a direct constraint on a process whose large dynamic range makes it difficult to observe/simulate, as provides a check on black hole mass functions.**

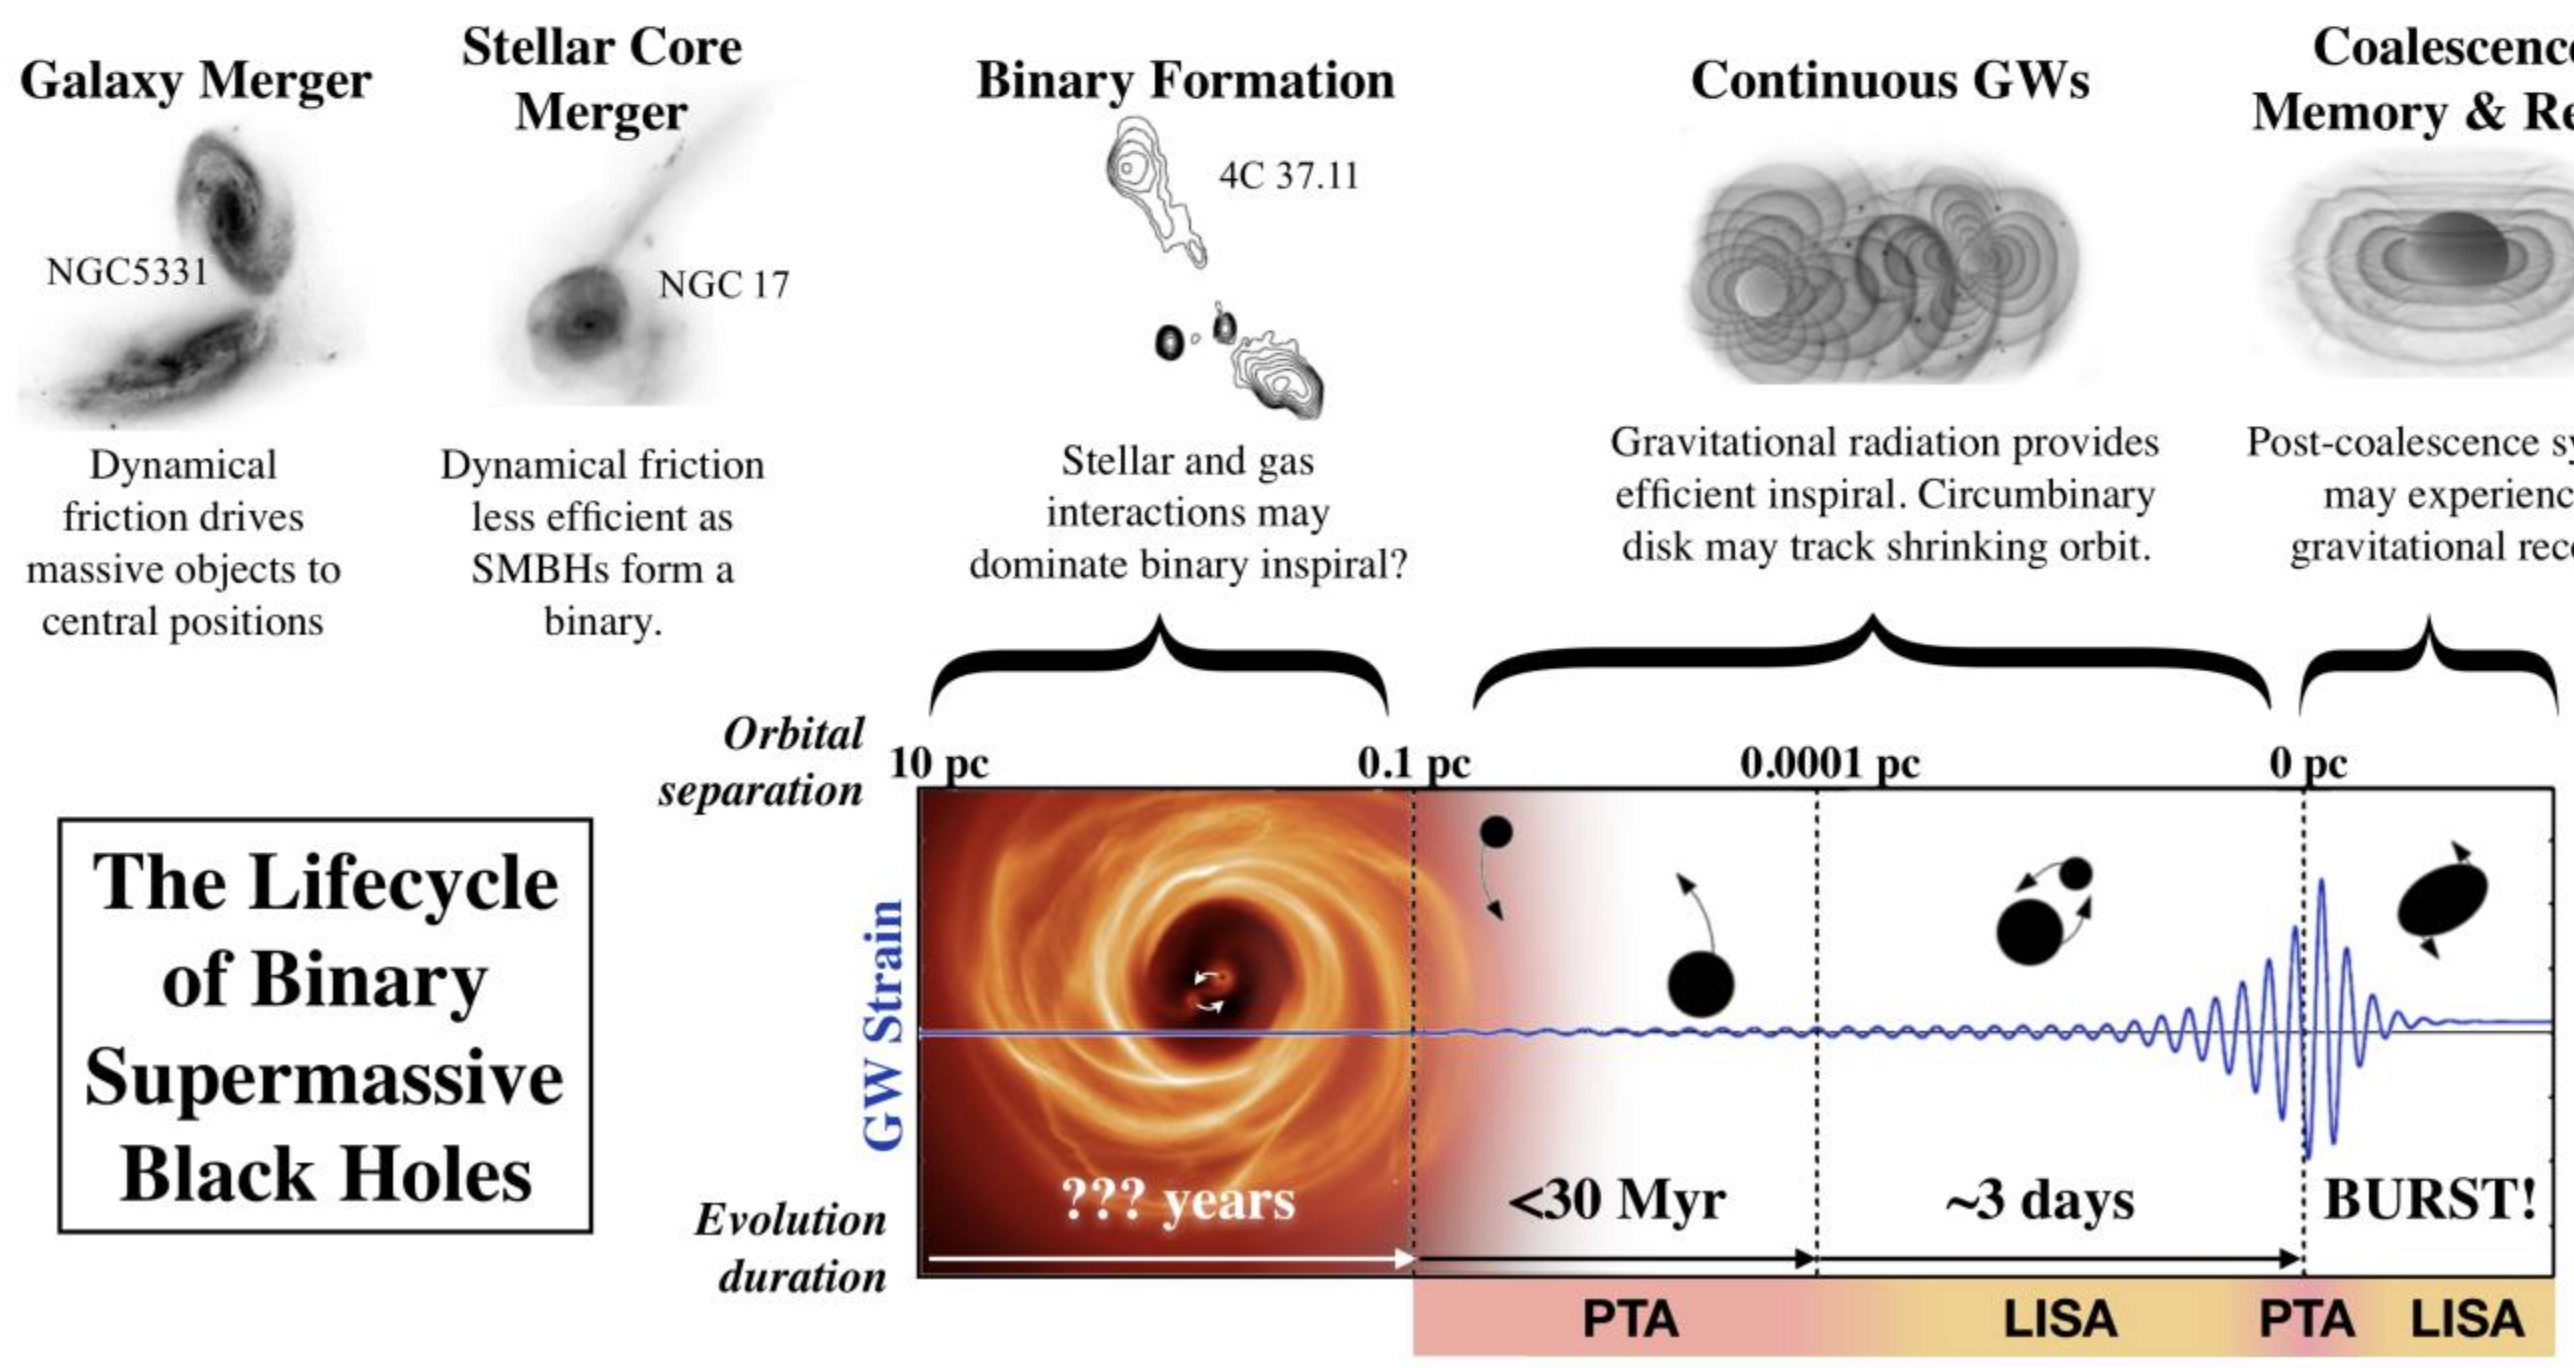

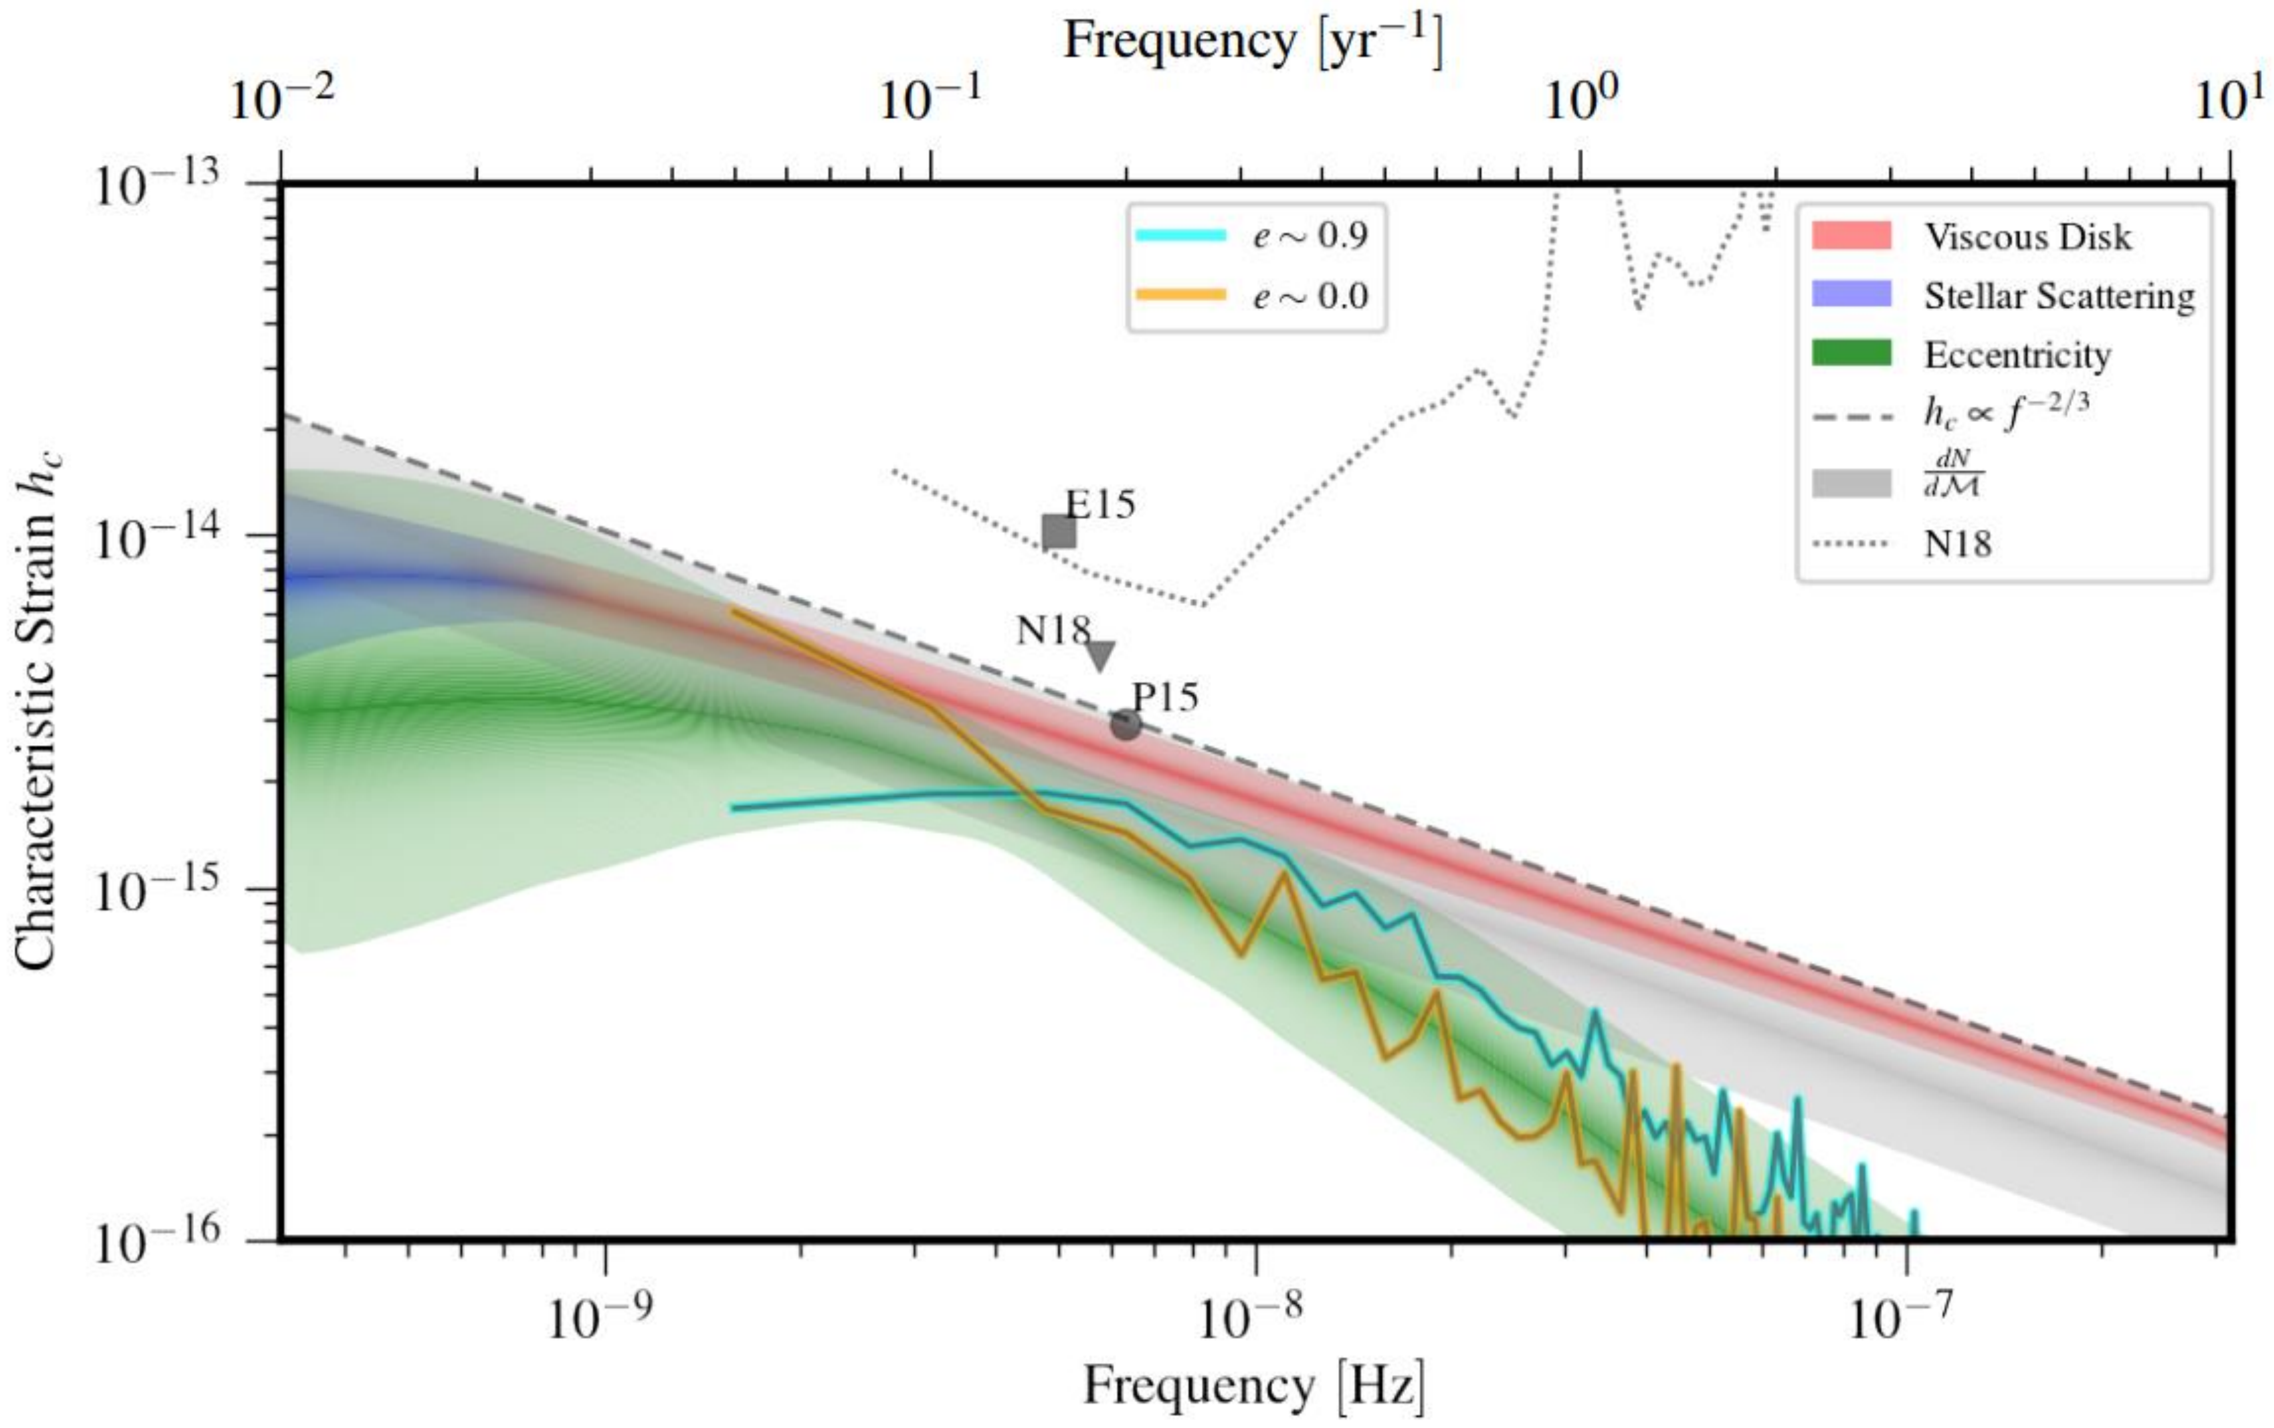

How to detect nHz GWs?

- Wavelengths > c*yr too long for even for LISA.
- Pulsars are celestial clocks, and many wavelengths of nHz GWs fit along one “detector arm”.
  - The longer one monitors a pulsar, the lower frequencies one can access.
- How do GW affect signals pulsar timing signals?
  - Intuitively, the bulk effect washes out, and the result depends on the GW strain at the “detector endpoints”.

RESPONSE OF DOPPLER SPACECRAFT TRACKING
TO GRAVITATIONAL RADIATION†

FRANK B. ESTABROOK and HUGO D. WAHLQUIST

Jet Propulsion Laboratory,
California Institute of Technology,
Pasadena, California 91103

Revised version received 16 *January* 1975

ABSTRACT

A calculation is made of the effect of gravity waves on the observed Doppler shift of a sinusoidal electromagnetic signal transmitted to, and transponded from, a distant spacecraft. We find that the effect of plane gravity waves on such observations is not intuitively immediate and in fact can have surprisingly different spectral signatures for different spacecraft directions and distances. We suggest the possibility of detecting such plane waves by simultaneous coherent Doppler tracking of several spacecraft.

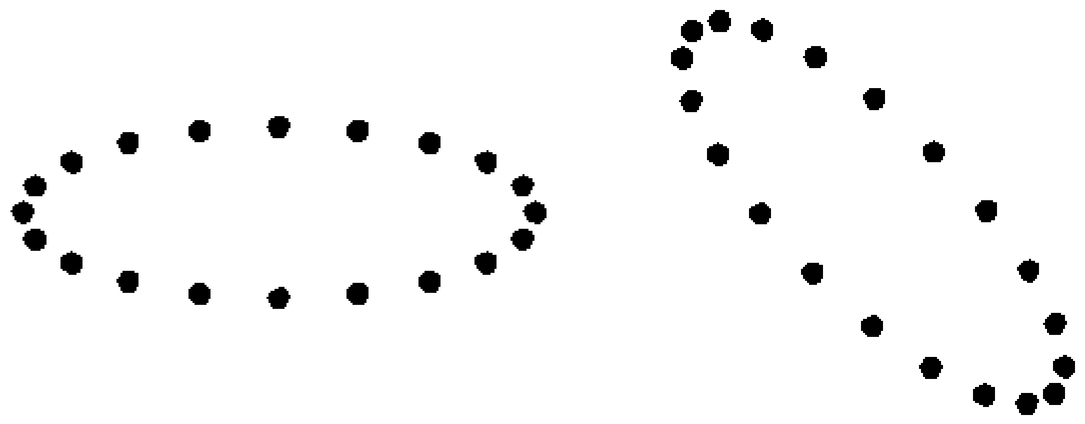

Wave polarization averages out for stochastic background.

$$\frac{\Delta\nu}{\nu} = \frac{1}{2}\cos 2\phi[1 - \cos\theta]$$

$$\times[h(t) - h(t - l - l\cos\theta)],$$

Angle between line-of-sight and GW propagation vector also averages, but leaves hallmark correlations (next slide).

The “earth” term: The GW amplitude at Earth.

The “pulsar” term: The GW amplitude at the pulsar with distance l. Unknown, but could be measured with ~1ly accurate distances!

Pulsar Timing Arrays

(PTAs are arrays of *pulsars* not of telescopes!)

- PTAs are monitored collections of high-precision millisecond pulsars (MSPs).
- The GWB induces time-dependent residuals in pulse arrival times with a power spectrum (with $\Gamma = {}^{13}/_{3}$) which is <u>common to every pulsar</u>:

$$P(f)\,[\mathrm{yr}^{-3}] = \frac{A_{\mathrm{gwb}}^2}{12\pi^2}\left(\frac{f}{\mathrm{yr}^{-1}}\right)^{-\Gamma}$$

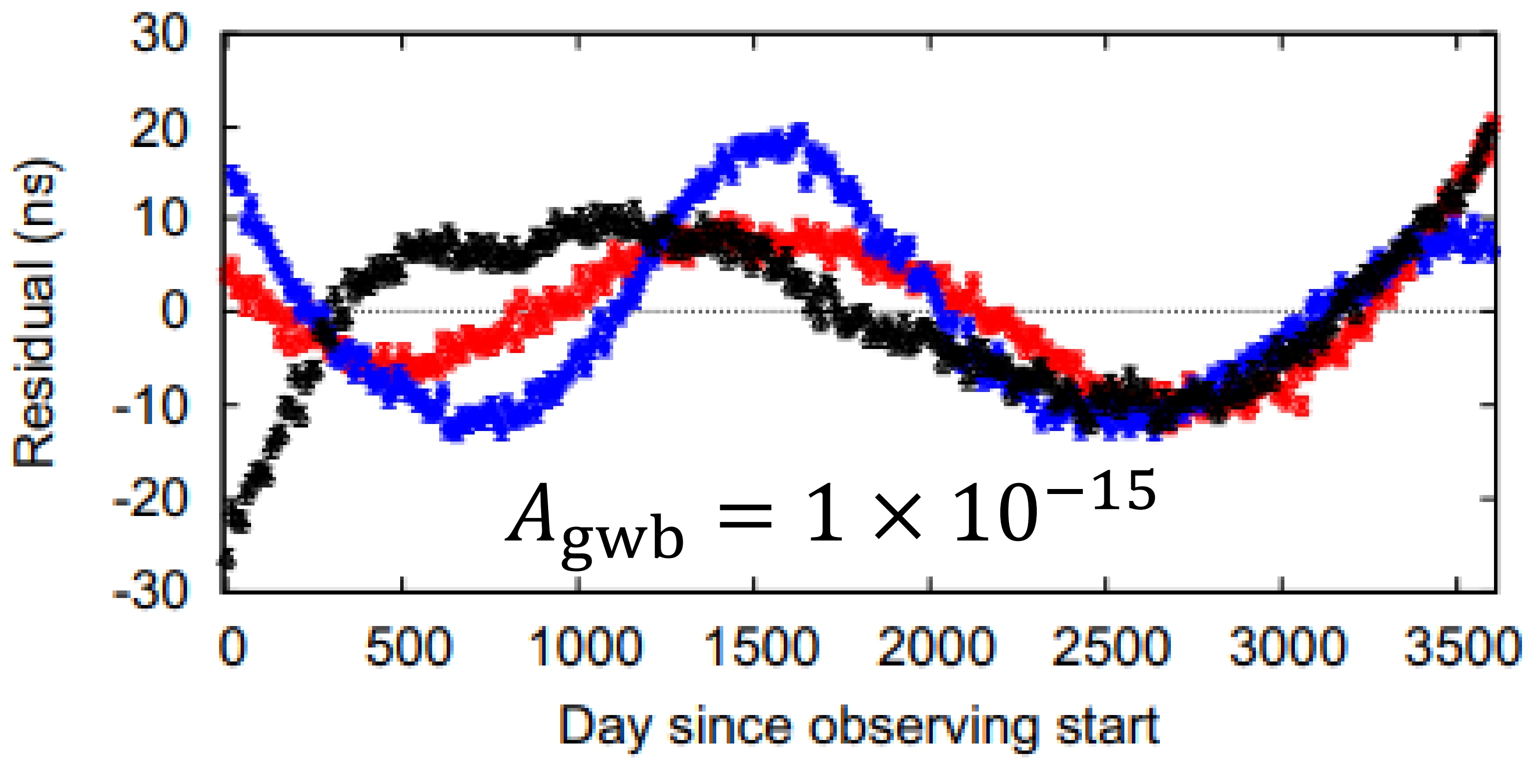

- Because pulsars share the “earth” term, the noise is correlated between pulsars depending on their angular separation. This is the famous “Hellings-Downs” curve.

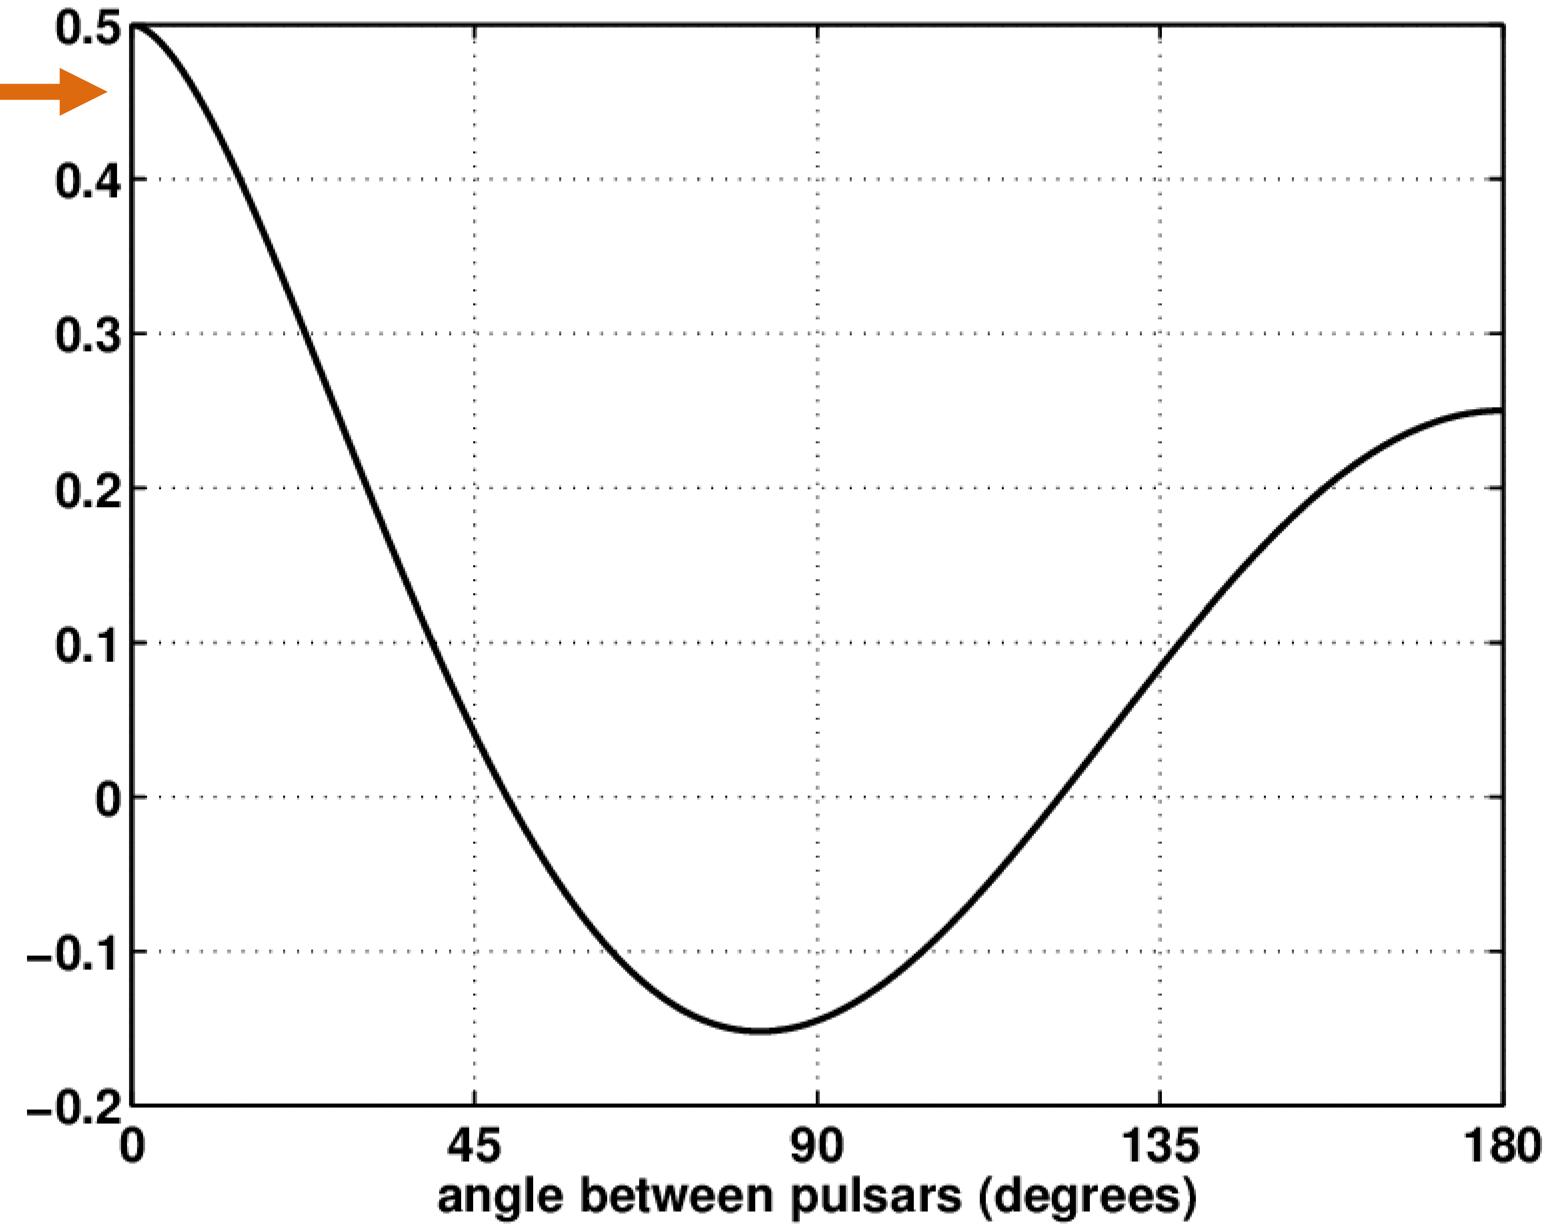

- So searching for the GWB has two prongs:
  1. Identification of noise processes with the right spectral shape, present at the same amplitude in *every* pulsar..
  2. Detection of the HD curve. However, because the typical correlation coefficient is small (absolute value < 0.2), <u>it is likely that the first method will yield the first detection</u>.

- Ancillary: the power spectrum is very steep, so the first detection will also come from the lowest frequencies. Long data sets are good!

Confounding Effects

The (ionized) interstellar medium (IISM) disperses, diffracts, and refracts radio waves. Main effect is dispersion, which introduces a frequency-dependent delay (relative, and absolute.)

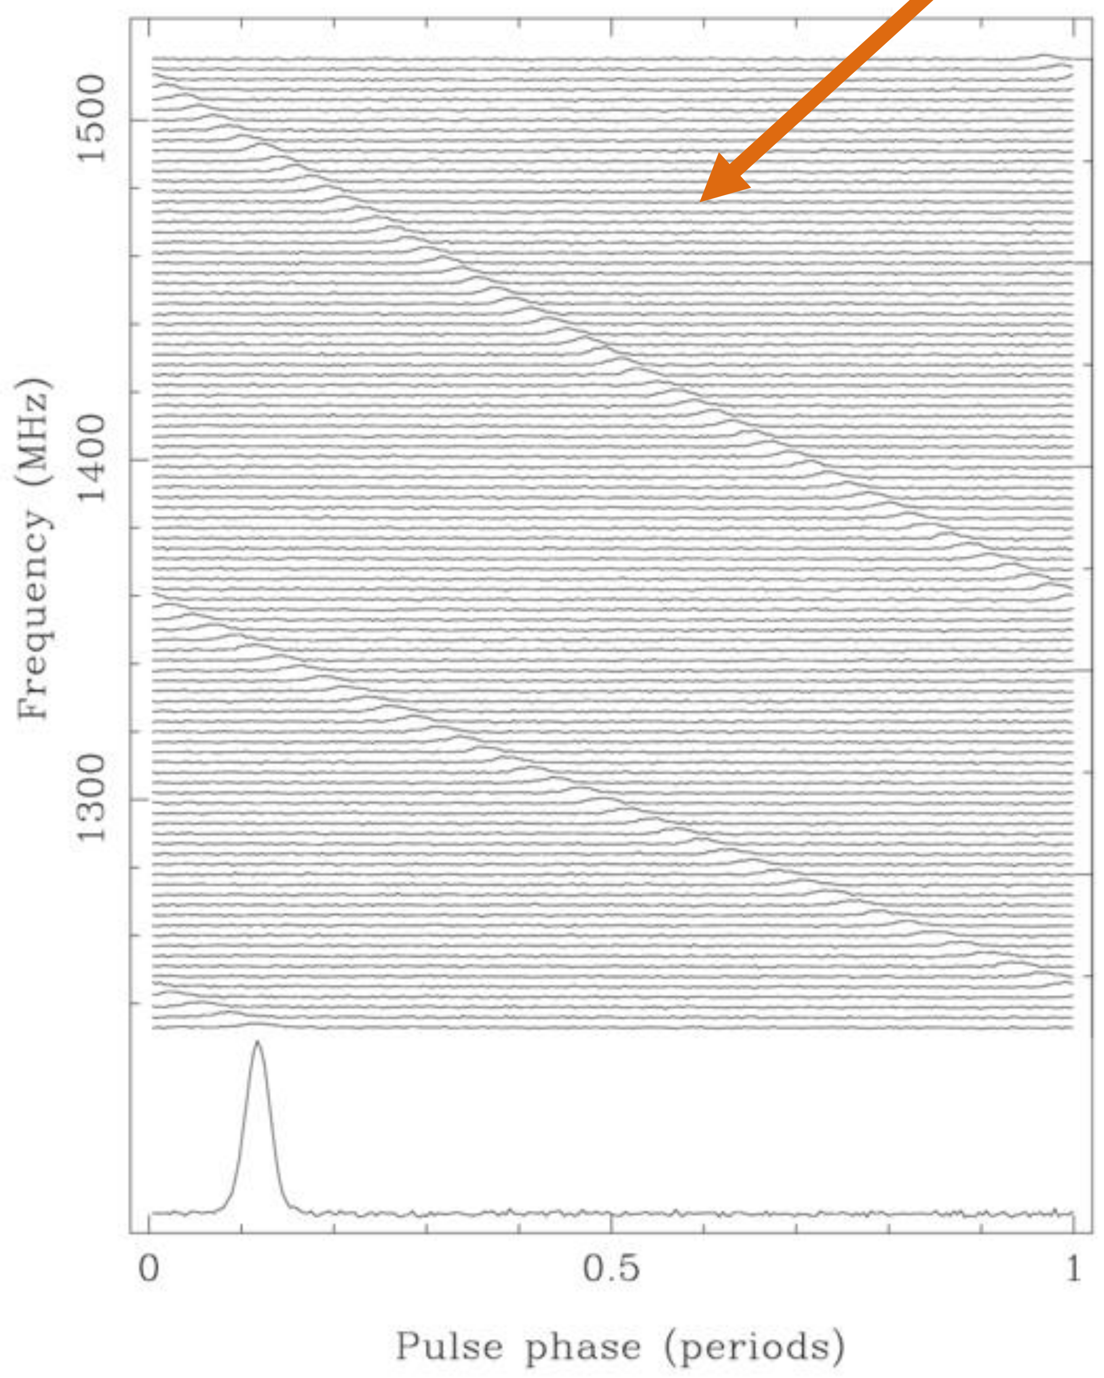

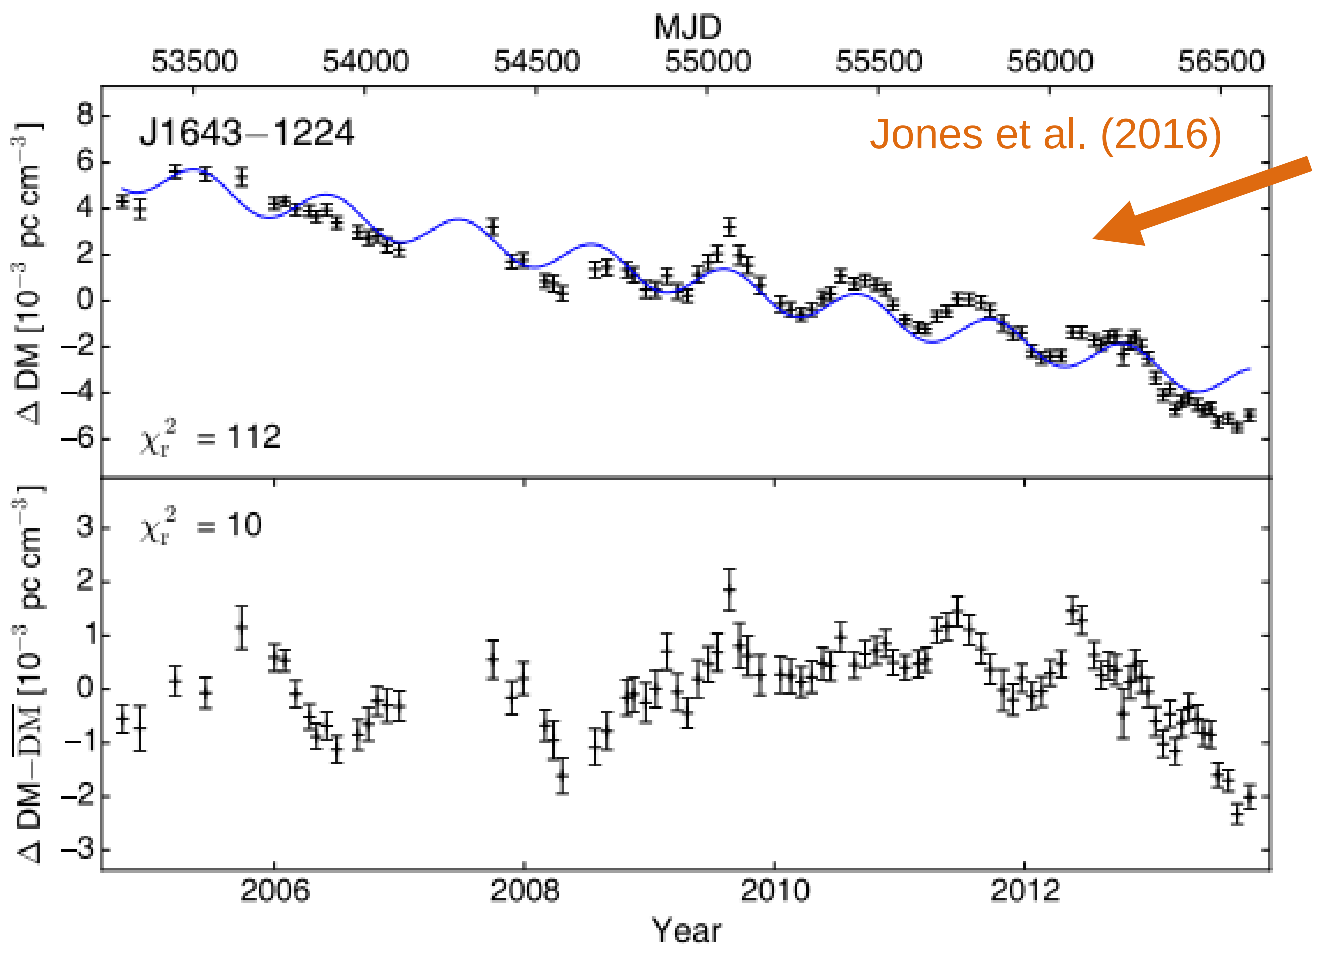

The dispersion measure (electron column density) varies with time because of relative motion. Here, note parallactic terms and a gradient from proper motion.

DM is usually estimated by fitting for the $\propto \nu^{-2}$ delay, but the pulse profile itself varies as a function of frequency.

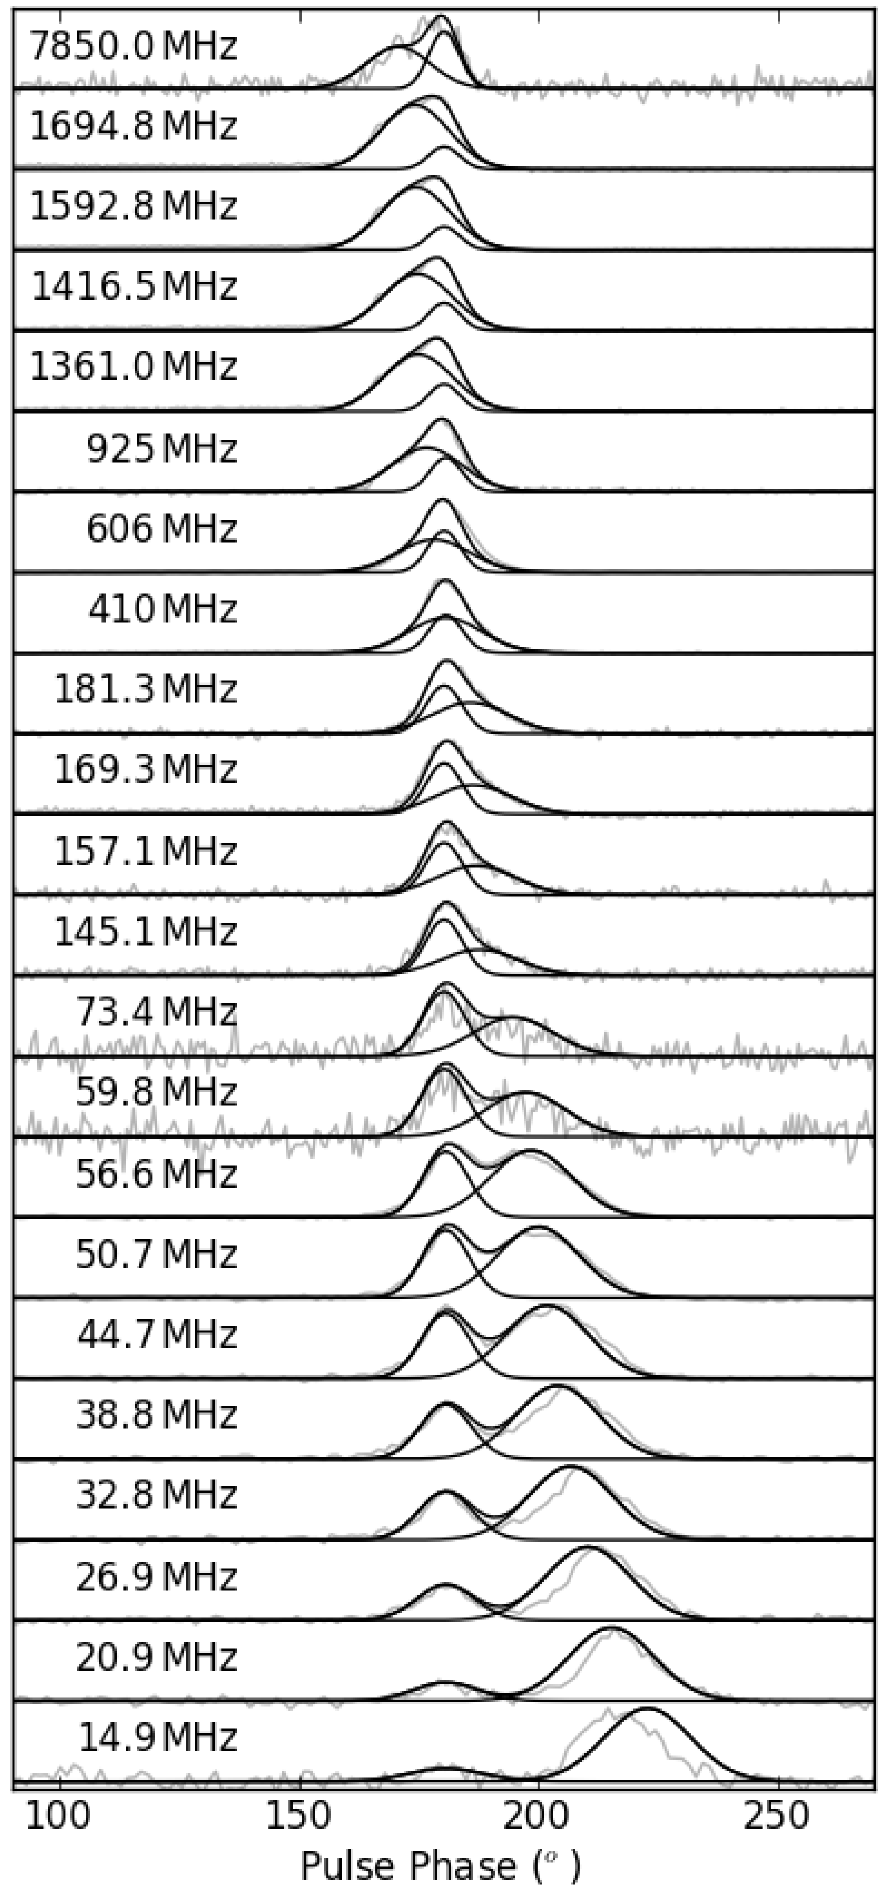

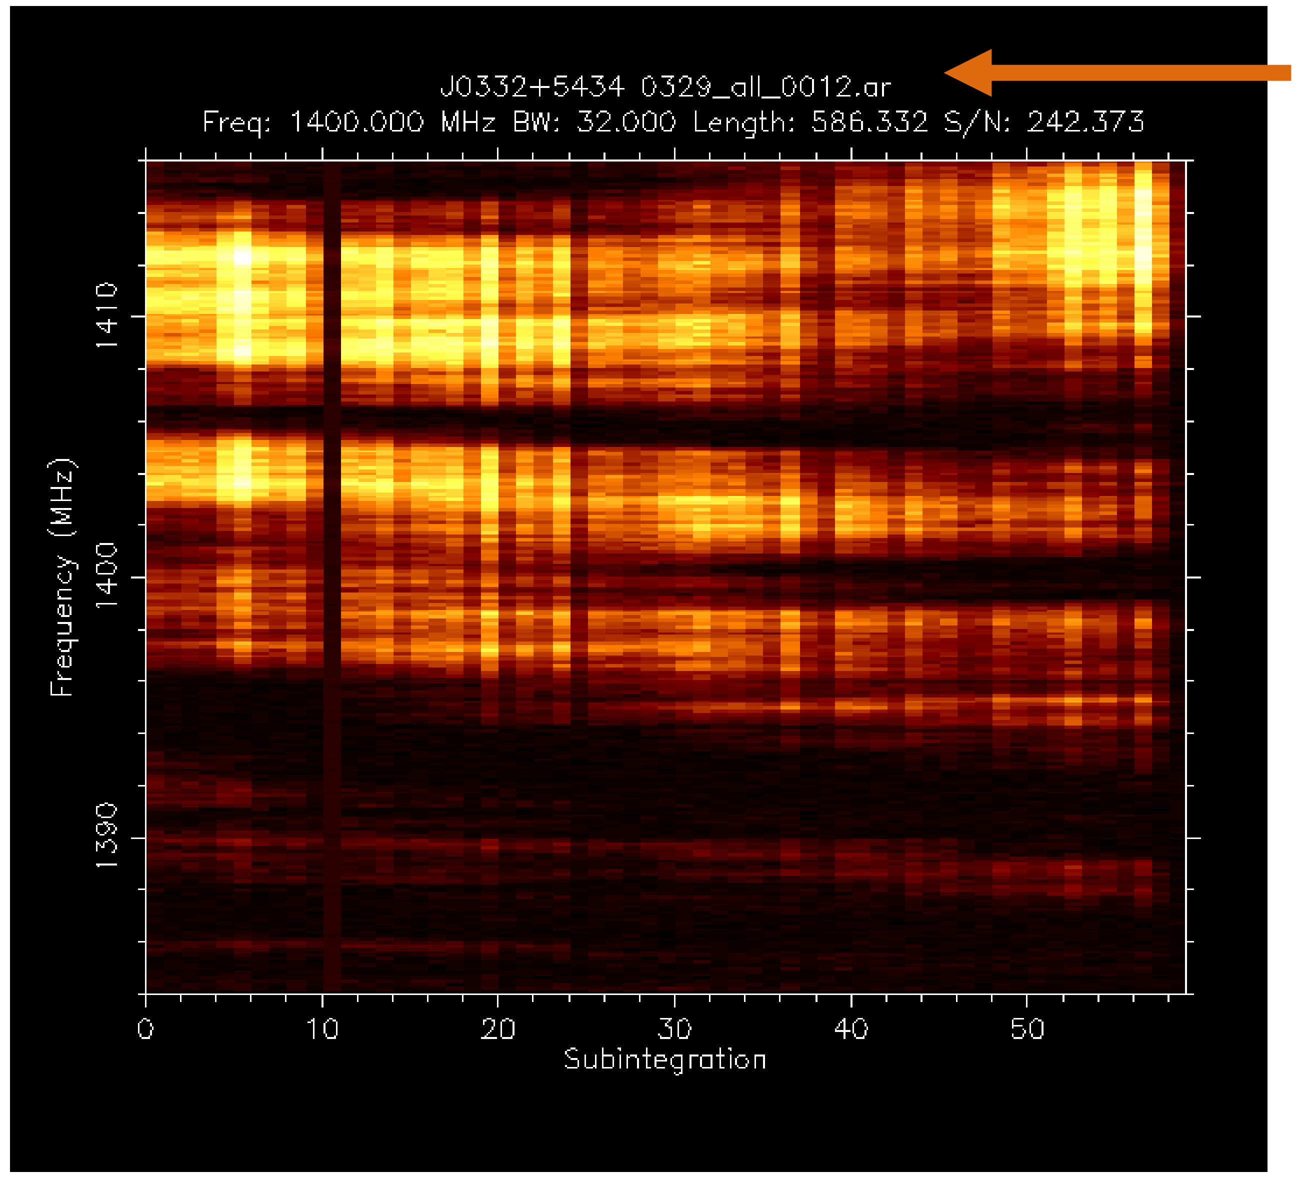

And diffraction/refraction change the received intensity as a function of time and frequency.

Other IISM effects don't have simple analytic models. All told, removing IISM effects is a major theoretical and practical challenge!

Further effects:

- Radio frequency interference (RFI)
- Instrumental "jumps"
- Jitter and other "white noise"
- Pulsar red/timing/spin noise

Recent Results

- PTAs are good! Have mitigated confounding effects to <1μs.
- Due to favorable time scaling ($\propto t_{\mathrm{span}}^{-13/6}$), with longer data sets, published limits reached as low as $A_{\mathrm{gwb}} < 1.0 \times 10^{-15}$.
- Recently, PTAs have accumulated ~3σ evidence for a "common mode" process in agreement with the predictions for "Prong 1" of the GWB with $A_{\mathrm{gwb}} \sim 2 - 3 \times 10^{-15}$.

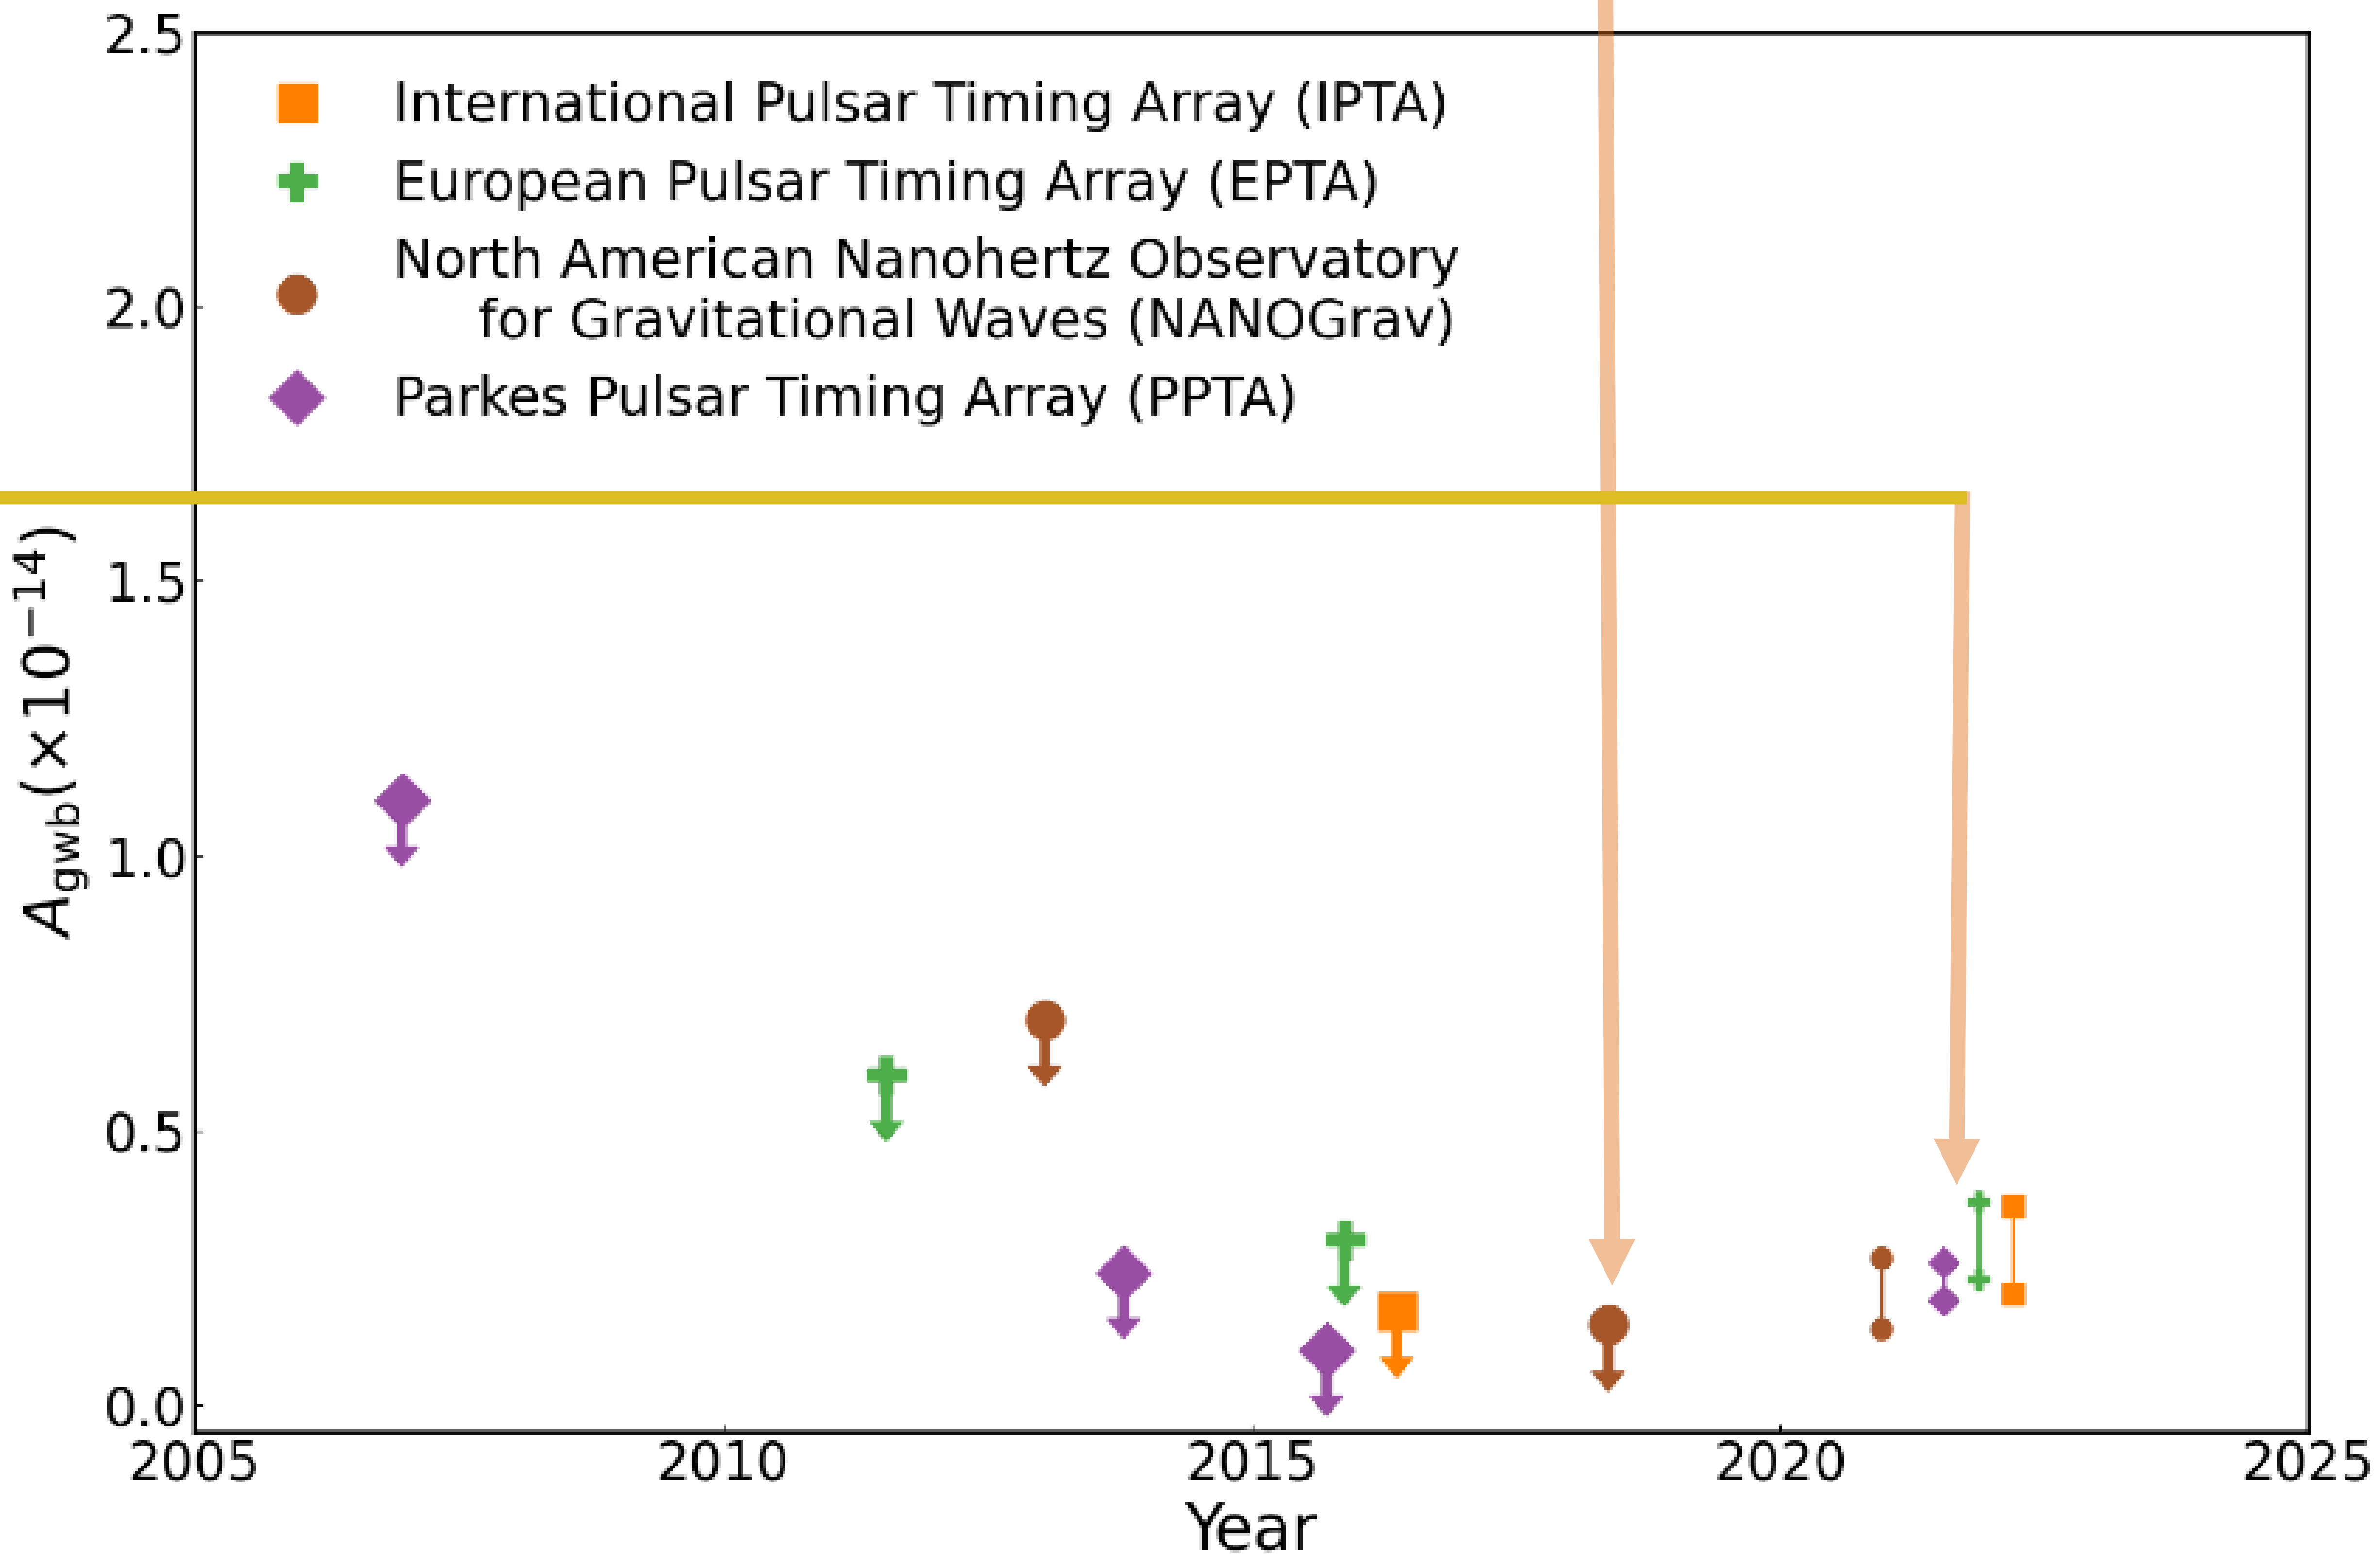

- Based on current PTA data quality, **if** this is a GWB signal, will require 5-8 years to detect Hellings-Downs correlations (Pol et al., 2021)
  - Detections are slower than limits because PTAs are self-noise limited by the GWB itself.

- Other possibilities: spin noise, residual IISM or other correlated systematics.

Fermi Large Area Telescope

- Fermi-LAT is a widefield pair-conversion telescope operating between ~50 MeV and ~1 TeV, most sensitive at ~1 GeV.
- Major sub-subsystems: anticoincidence detector, silicon strip tracker, and CsI hodoscopic calorimeter.
- Operating since 2008: long uninterrupted dataset.
- Timestamping accurate to <300 ns and onboard GPS provides accurate absolute time reference.

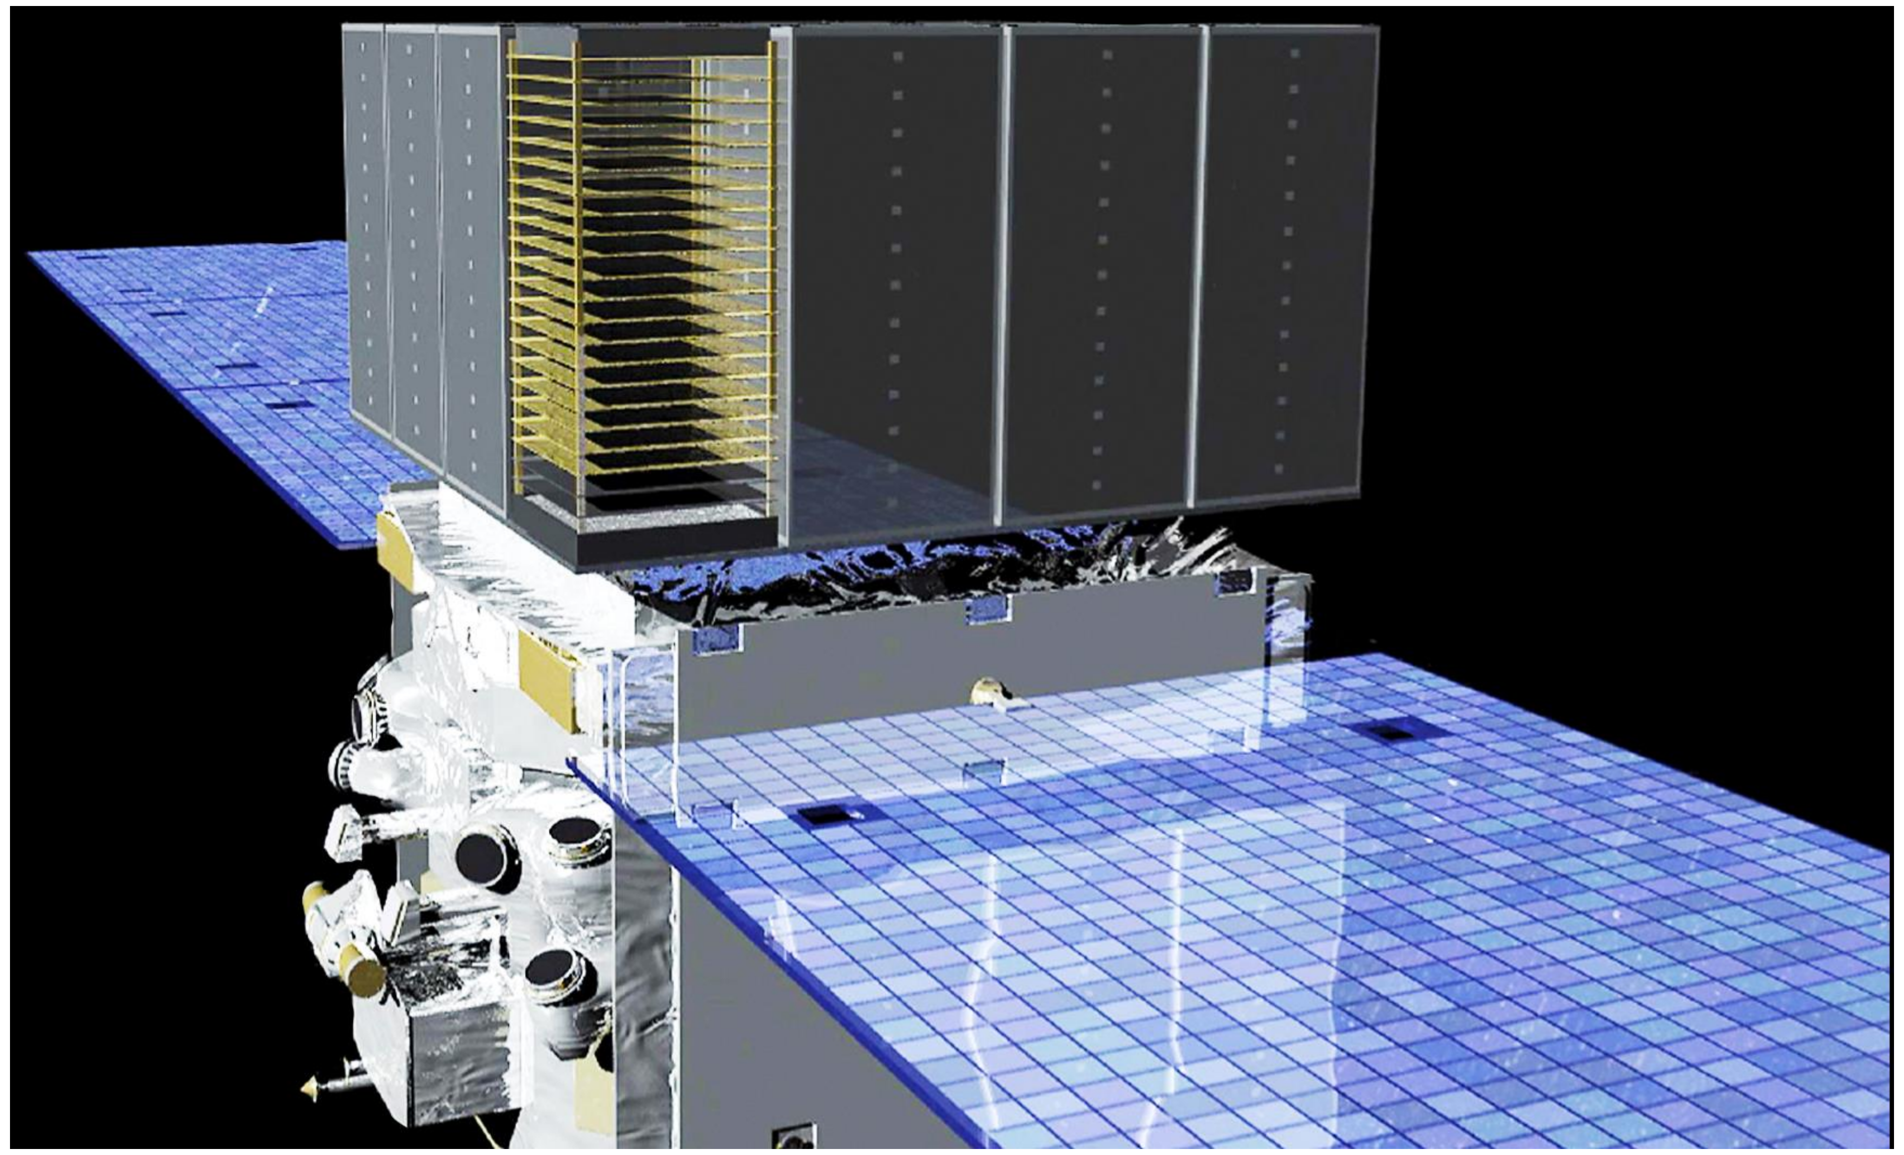

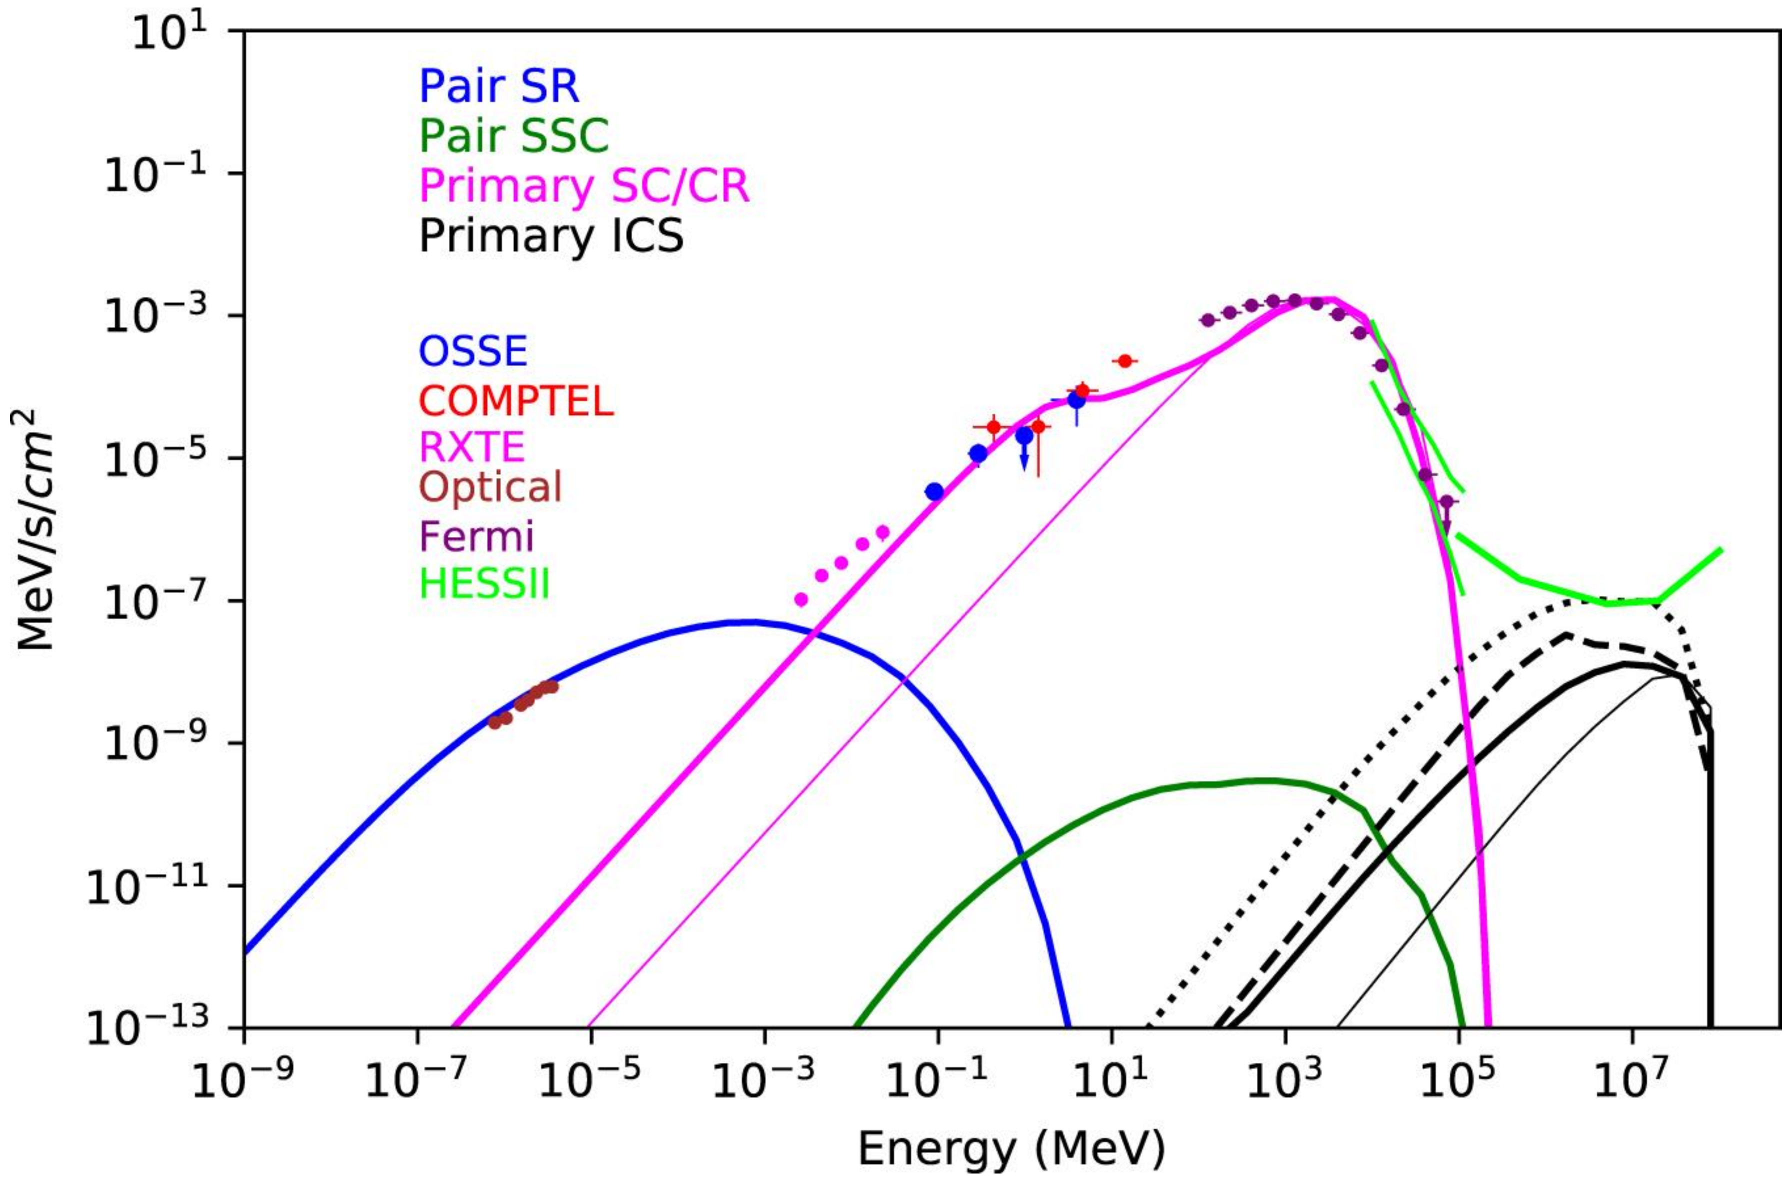

Most importantly, many pulsars are very strong GeV emitters, converting up to or more than 10% of their spin-down power into gamma rays.

This includes millisecond pulsars!

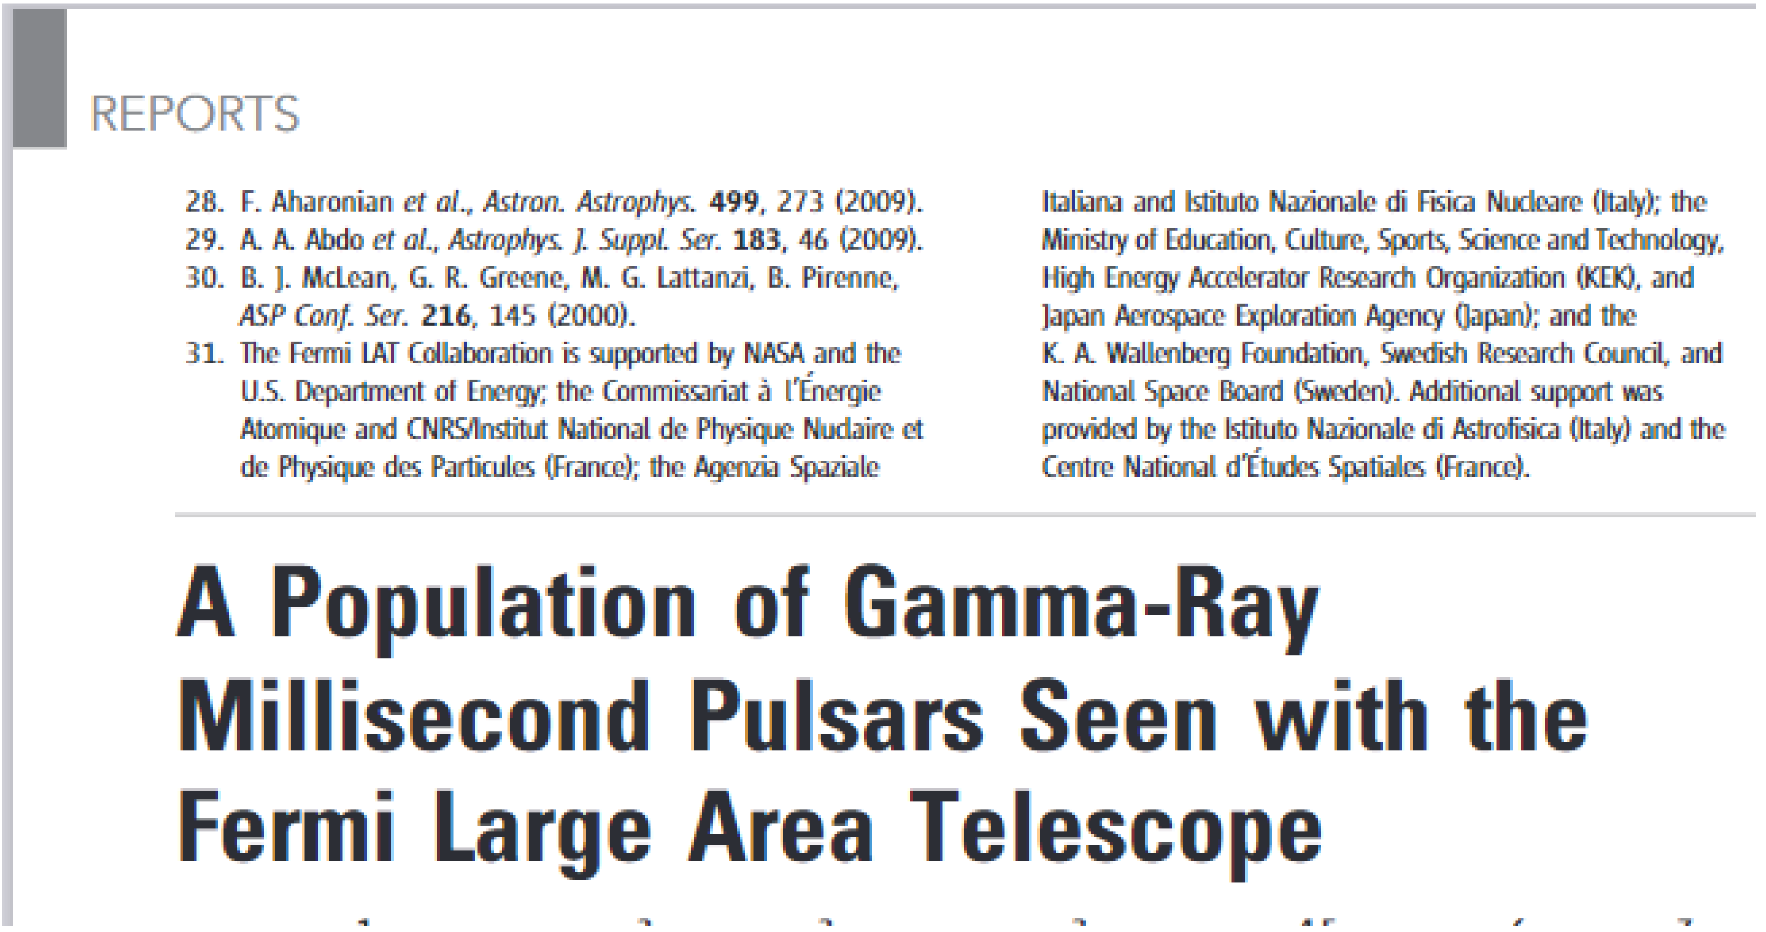
REPORTS

28. F. Aharonian *et al.*, *Astron. Astrophys.* **499**, 273 (2009).
29. A. A. Abdo *et al.*, *Astrophys. J. Suppl. Ser.* **183**, 46 (2009).
30. B. J. McLean, G. R. Greene, M. G. Lattanzi, B. Pirenne, *ASP Conf. Ser.* **216**, 145 (2000).

31. The Fermi LAT Collaboration is supported by NASA and the U.S. Department of Energy; the Commissariat à l'Énergie Atomique and CNRS/Institut National de Physique Nucléaire et de Physique des Particules (France); the Agenzia Spaziale Italiana and Istituto Nazionale di Fisica Nucleare (Italy); the Ministry of Education, Culture, Sports, Science and Technology, High Energy Accelerator Research Organization (KEK), and Japan Aerospace Exploration Agency (Japan); and the K. A. Wallenberg Foundation, Swedish Research Council, and National Space Board (Sweden). Additional support was provided by the Istituto Nazionale di Astrofisica (Italy) and the Centre National d'Études Spatiales (France).

A Population of Gamma-Ray Millisecond Pulsars Seen with the Fermi Large Area Telescope

Fermi Large Area Telescope

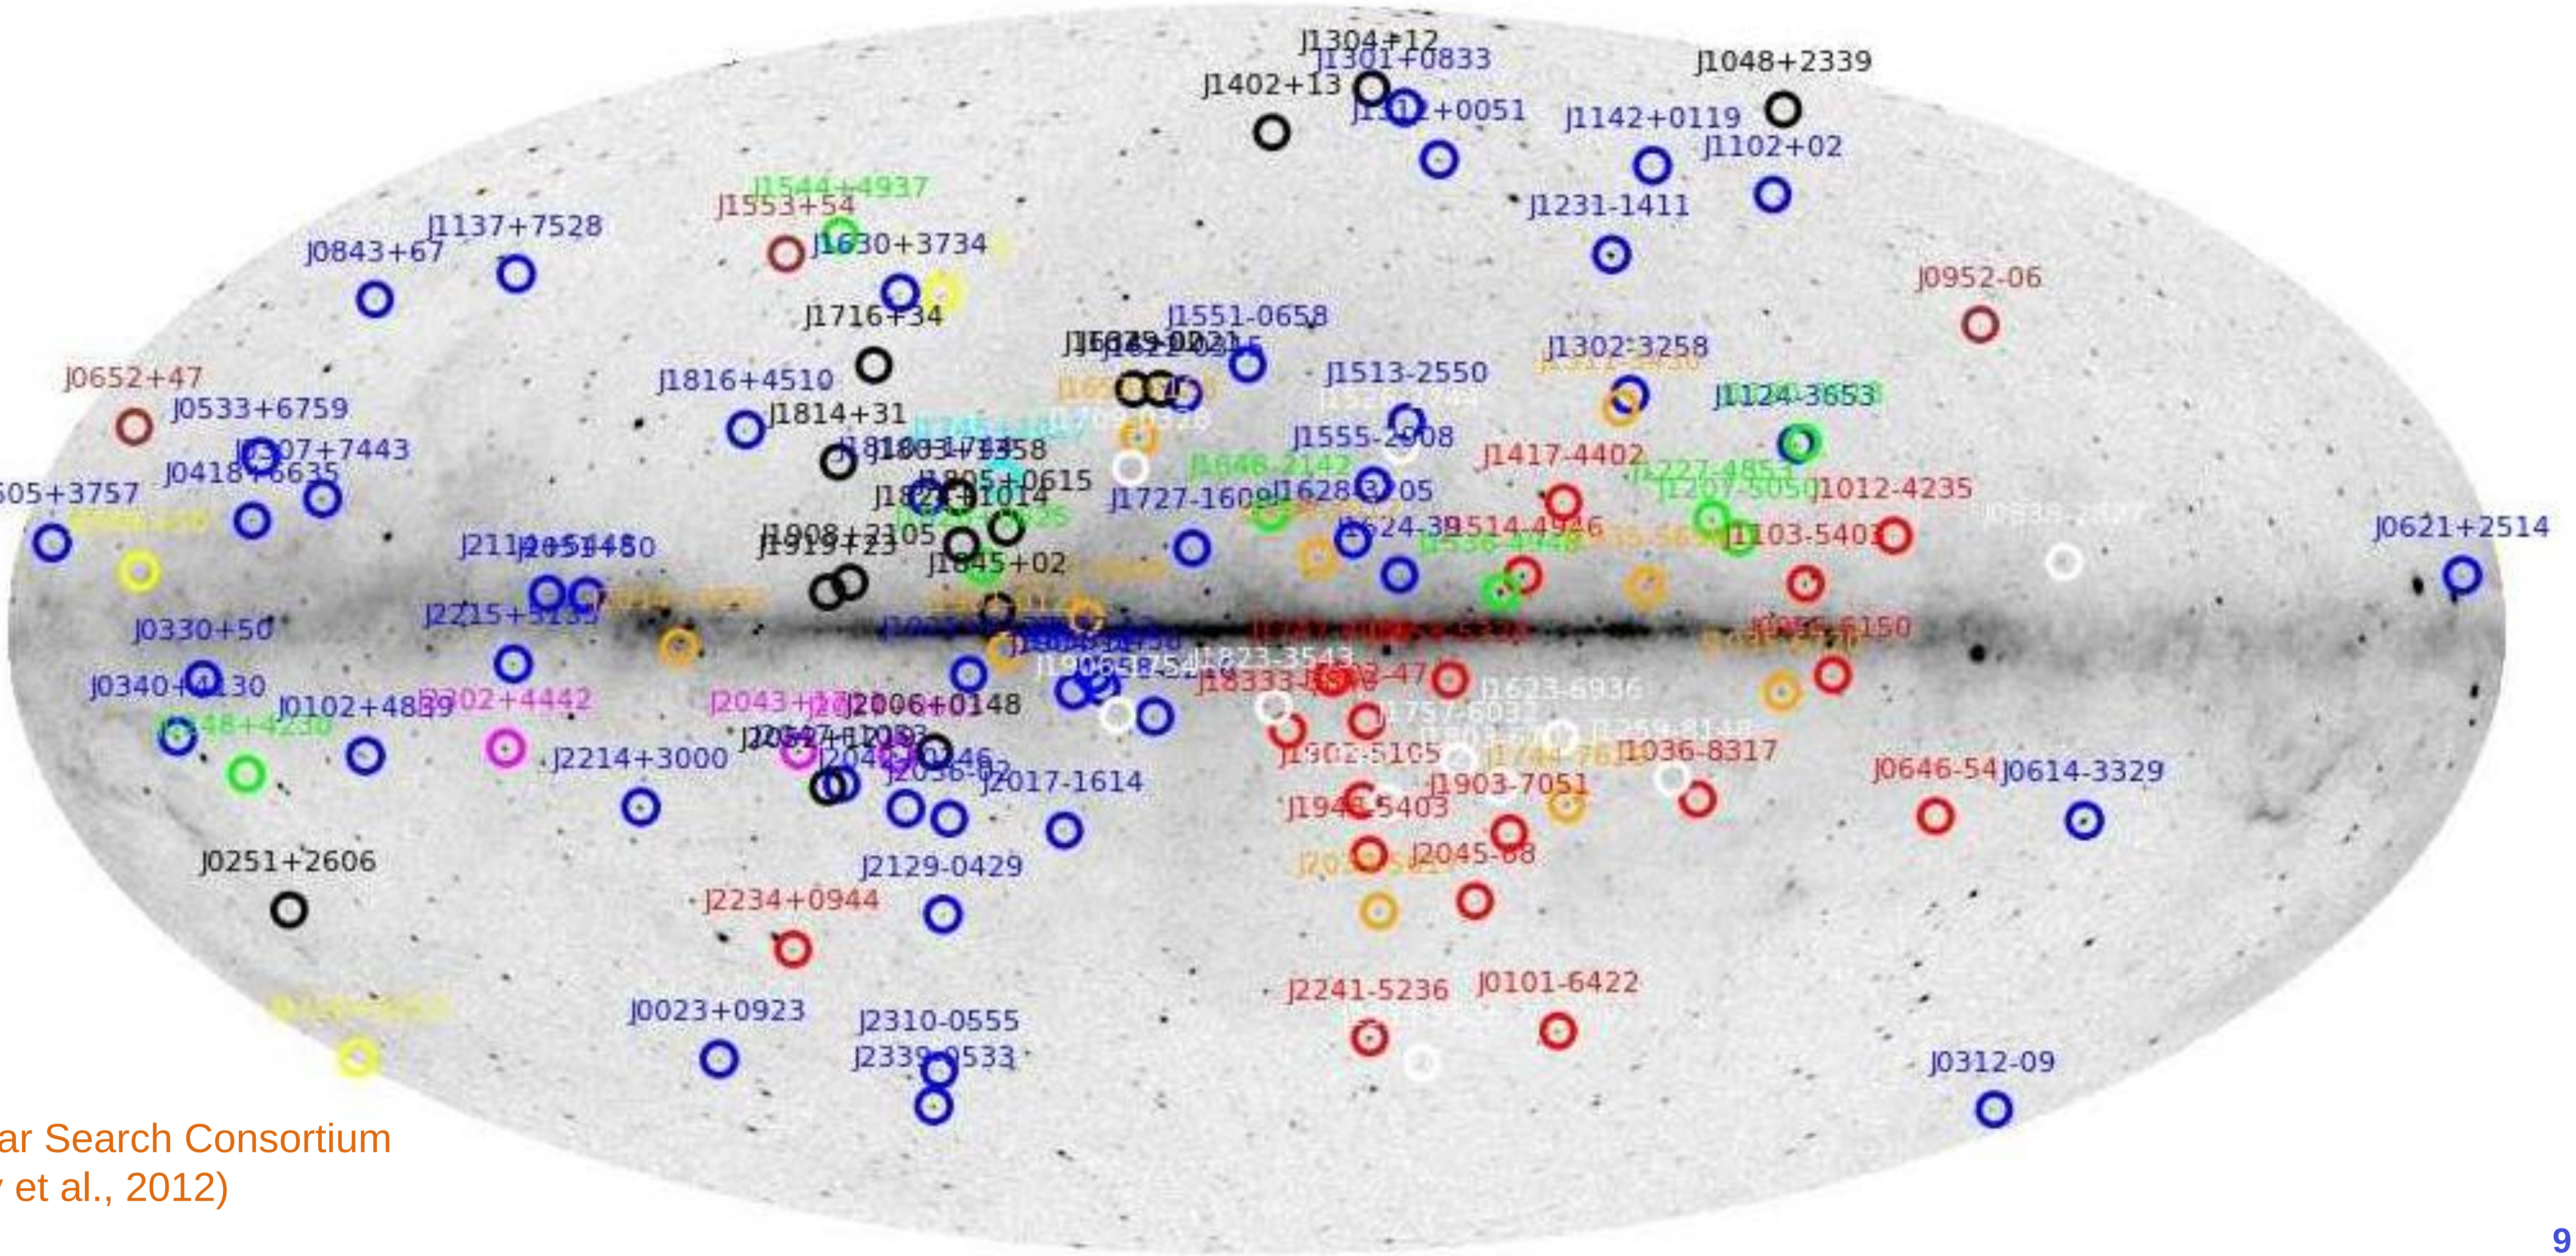

Pulsar Search Consortium
(Ray et al., 2012)

LAT Data vs. Radio Data

- In some cases pulsar timing is similar to radio:
  - Observe a pulsar “long enough” to detect its pulse profile and reference it to a good clock (“TOA”).
  - Use Poisson likelihood instead of gaussian for LAT.
- Integration times vary wildly: ~10 minutes for Vela, up to 1 year for faintest pulsars.
  - Averaging so much data together smears out signals, e.g. from the 1-year annual sinusoid from position fitting.
- Best to use an “unbinned” approach – compute the spin phase of each photon and maximize the likelihood.
- We have developed pipelines for both TOA-based and unbinned gamma-ray pulsar timing.

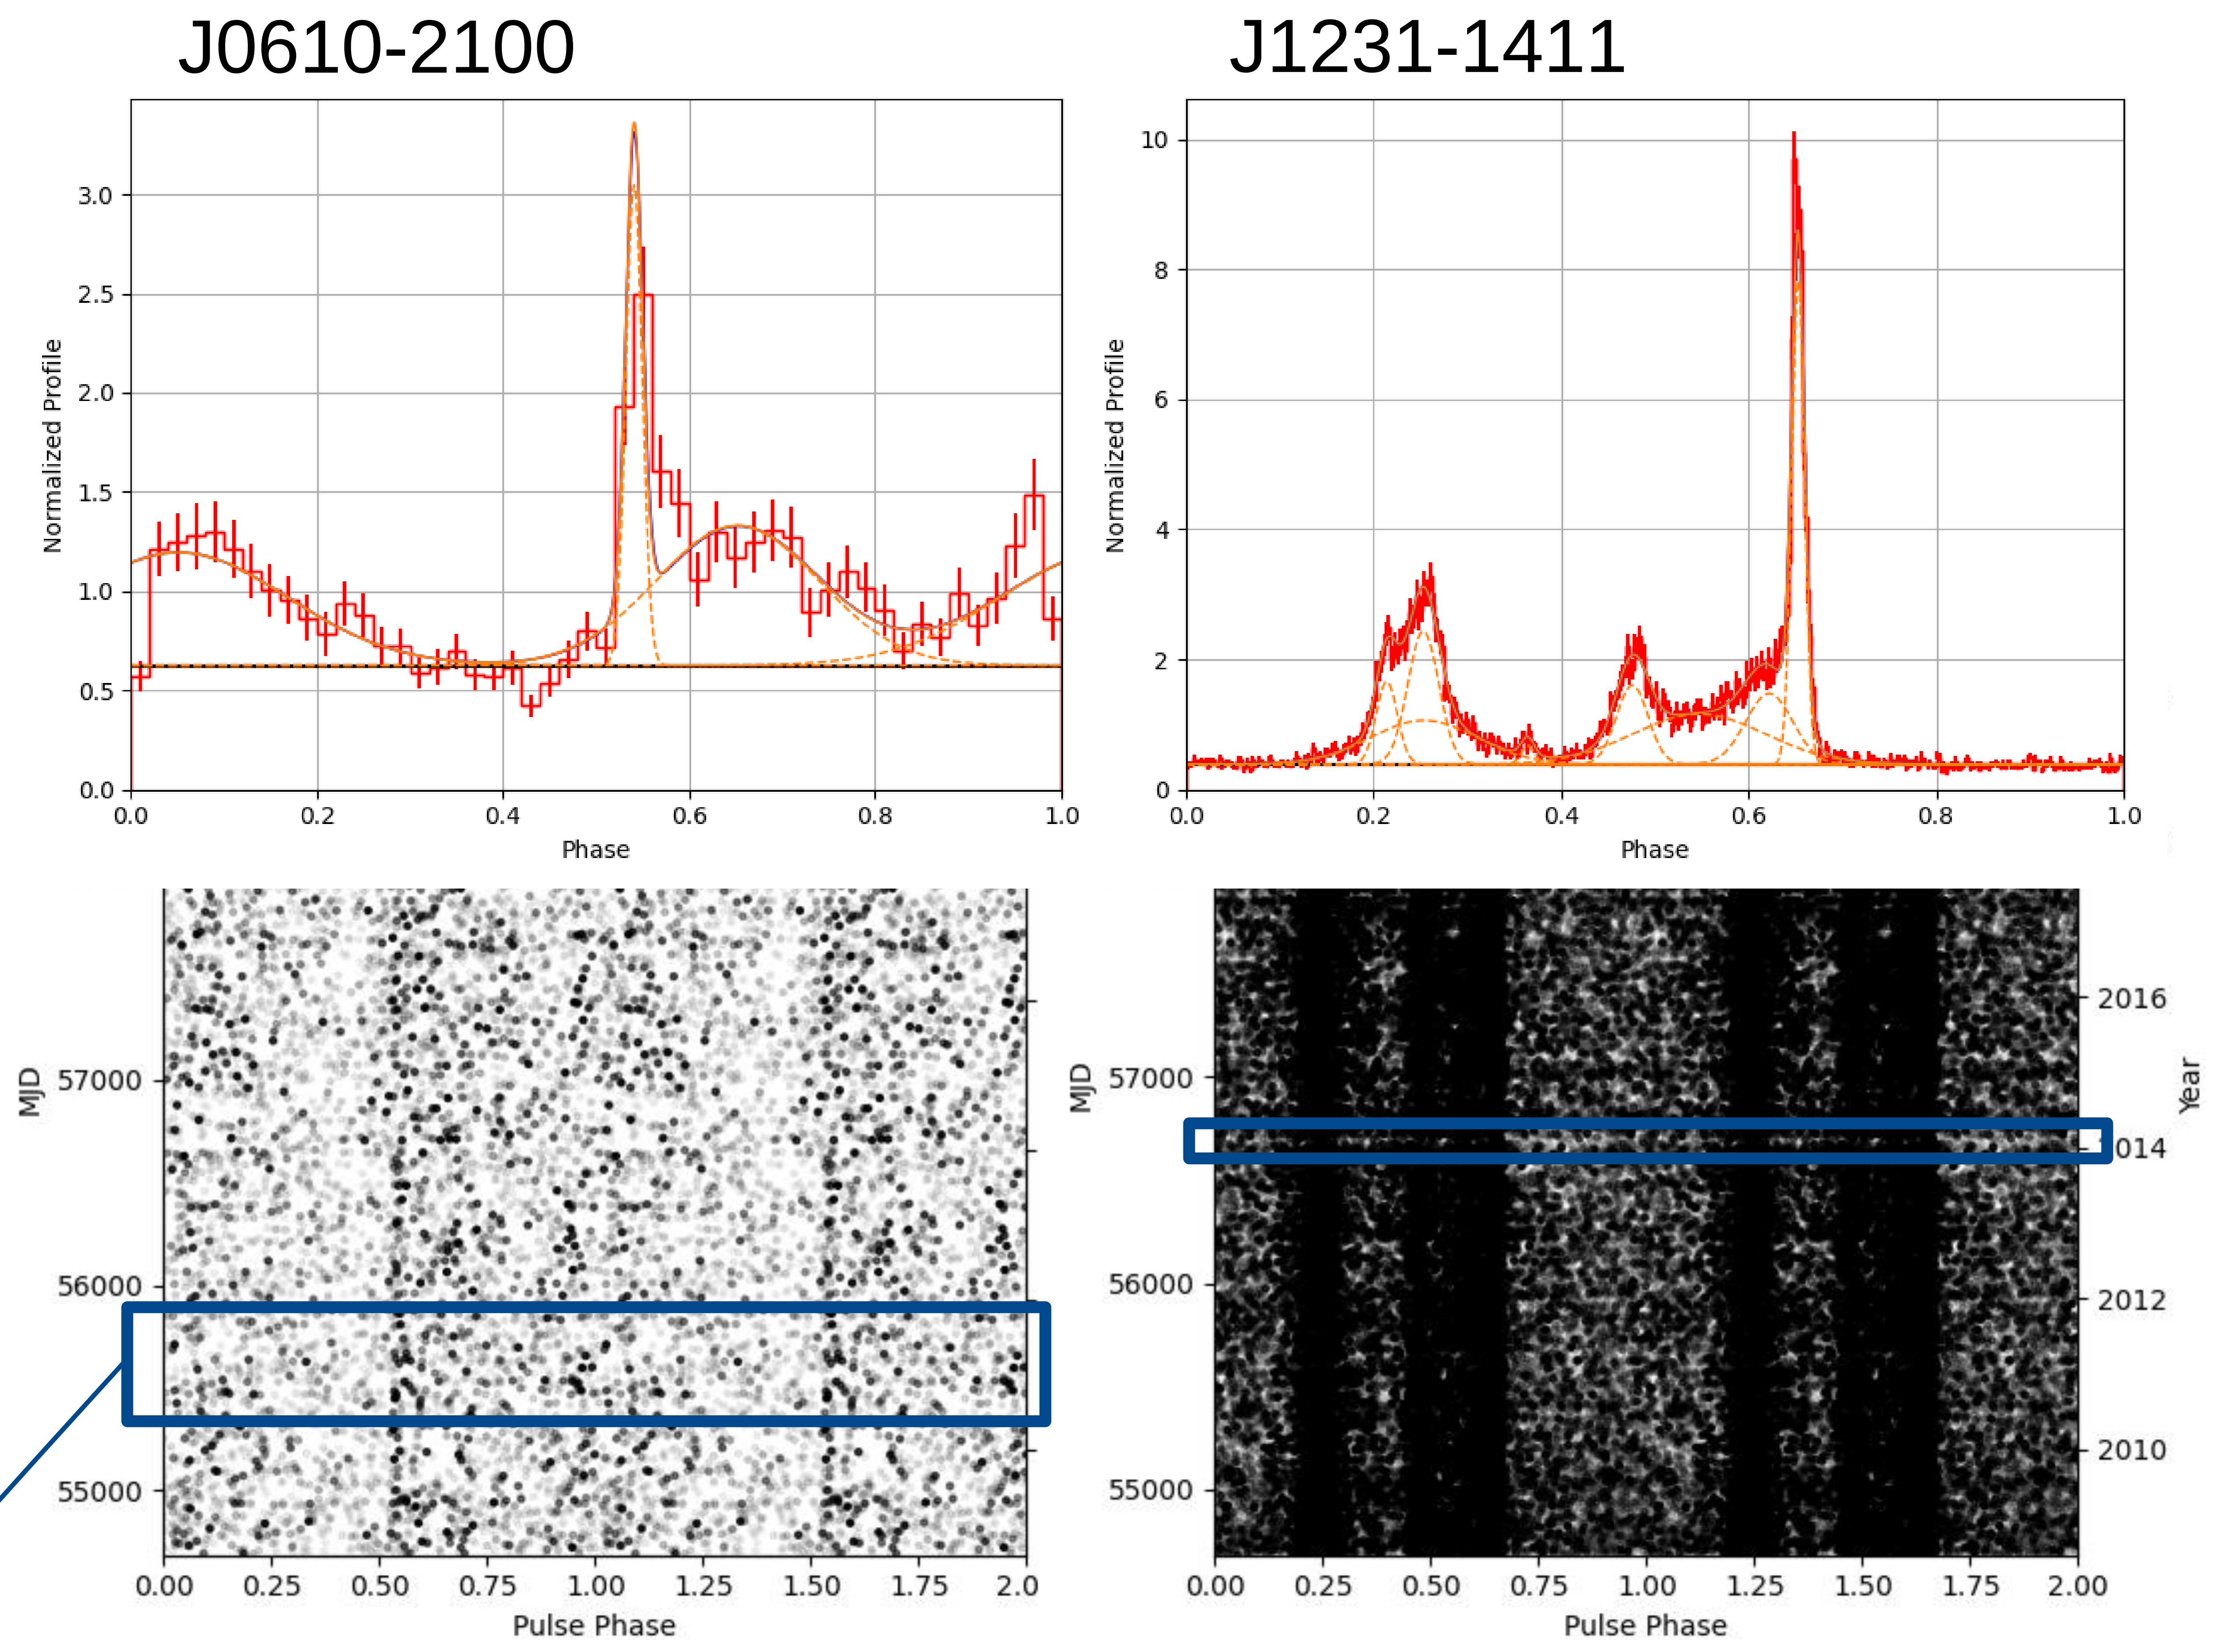

Increase integration window until it encircles enough photons to significantly see the pulse profile.

For brighter pulsars, the window is narrower.

Thus, for some applications, e.g. determining a position, only bright pulsars (many windows per year) are suitable.

Placing Constraints on the GWB

- Of the ~130 MSPs, ~30 are useful for PTA work.
  - Balance of timing precision with overall brightness.

- Used two techniques to search for a GWB:
  - Measured pulse arrival times (TOAs) like radio PTAs and adopted widely used codes (TempoNest and enterprise) to characterize noise and GWB signals, including Hellings-Downs correlations.
  - Developed an unbinned method which retains full time resolution and avoids systematic uncertainty from TOA measurement.
  - **Excellent agreement between methods!**

- Primary results are single and ensemble limits on the "Prong 1" common mode manifestation of A_{gwb}.

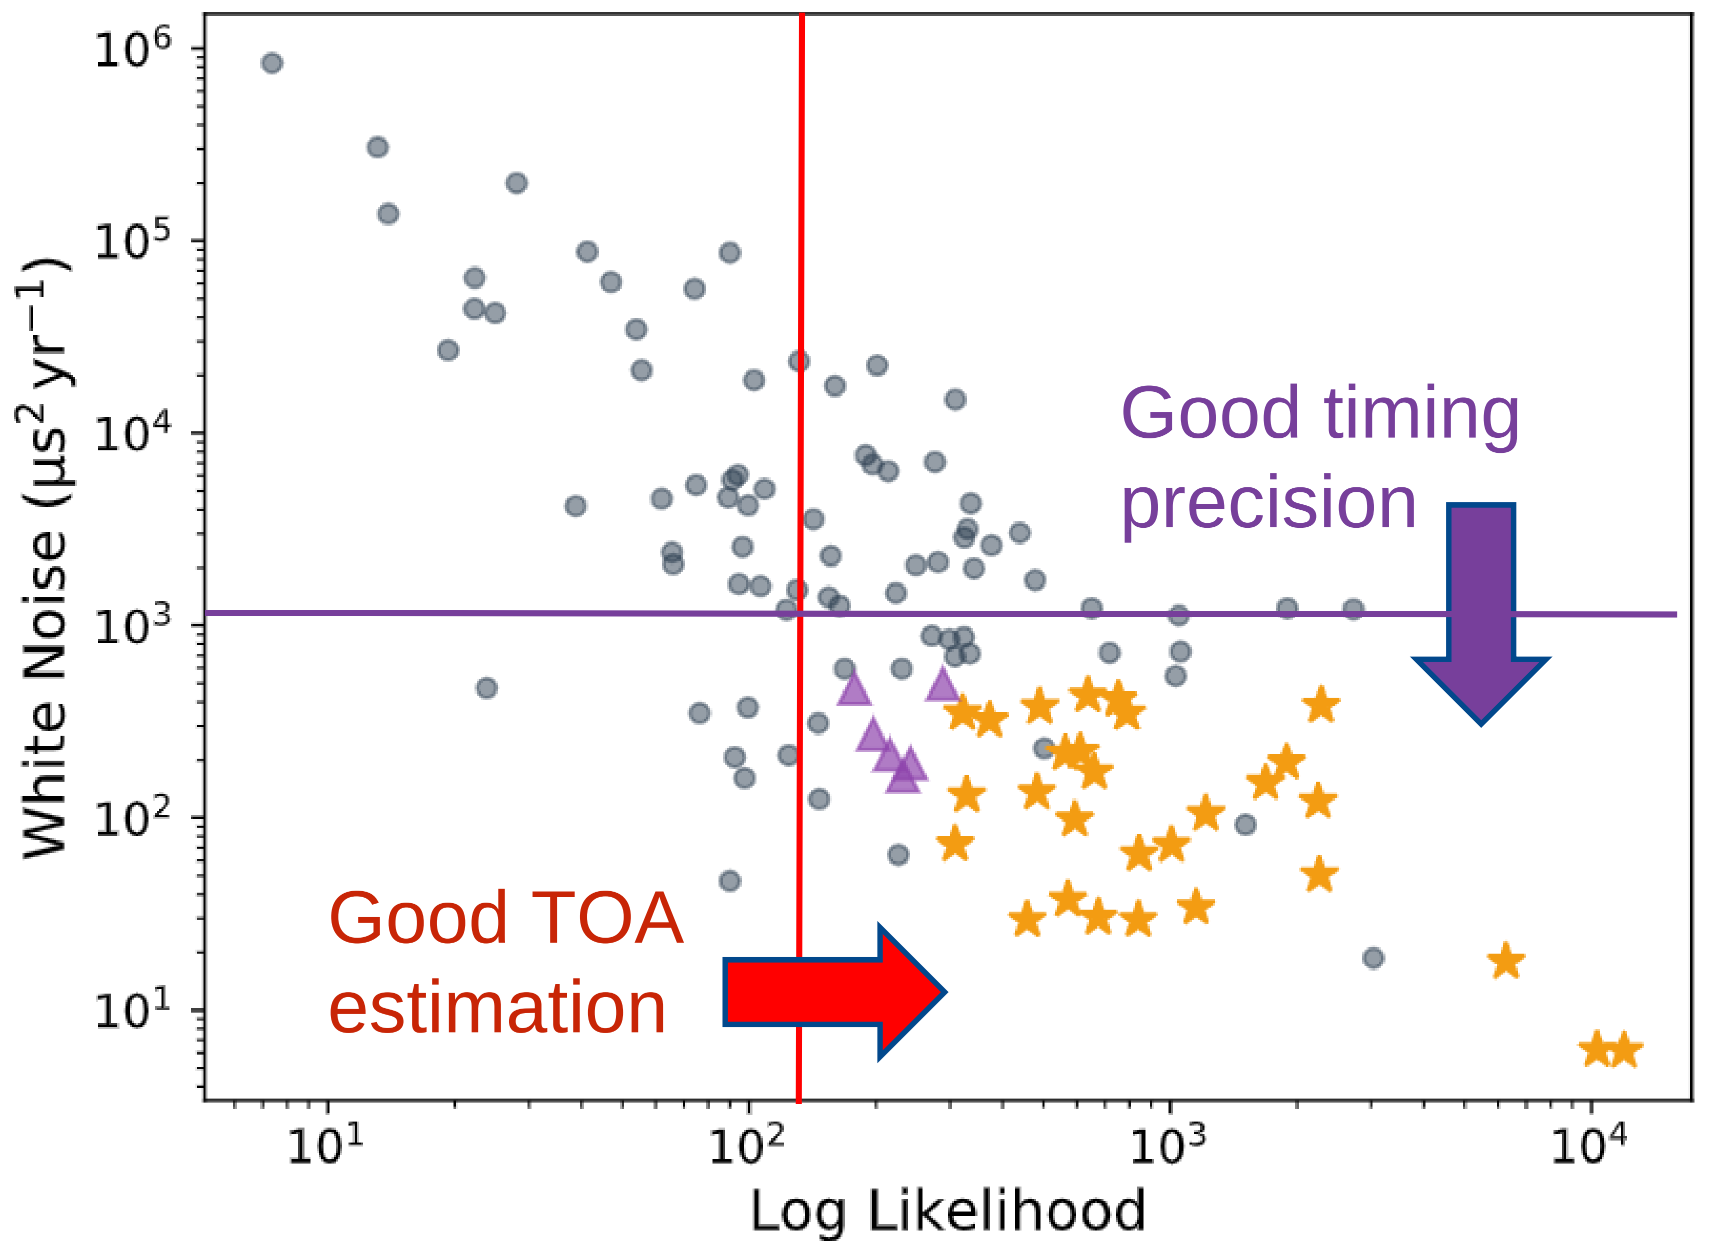

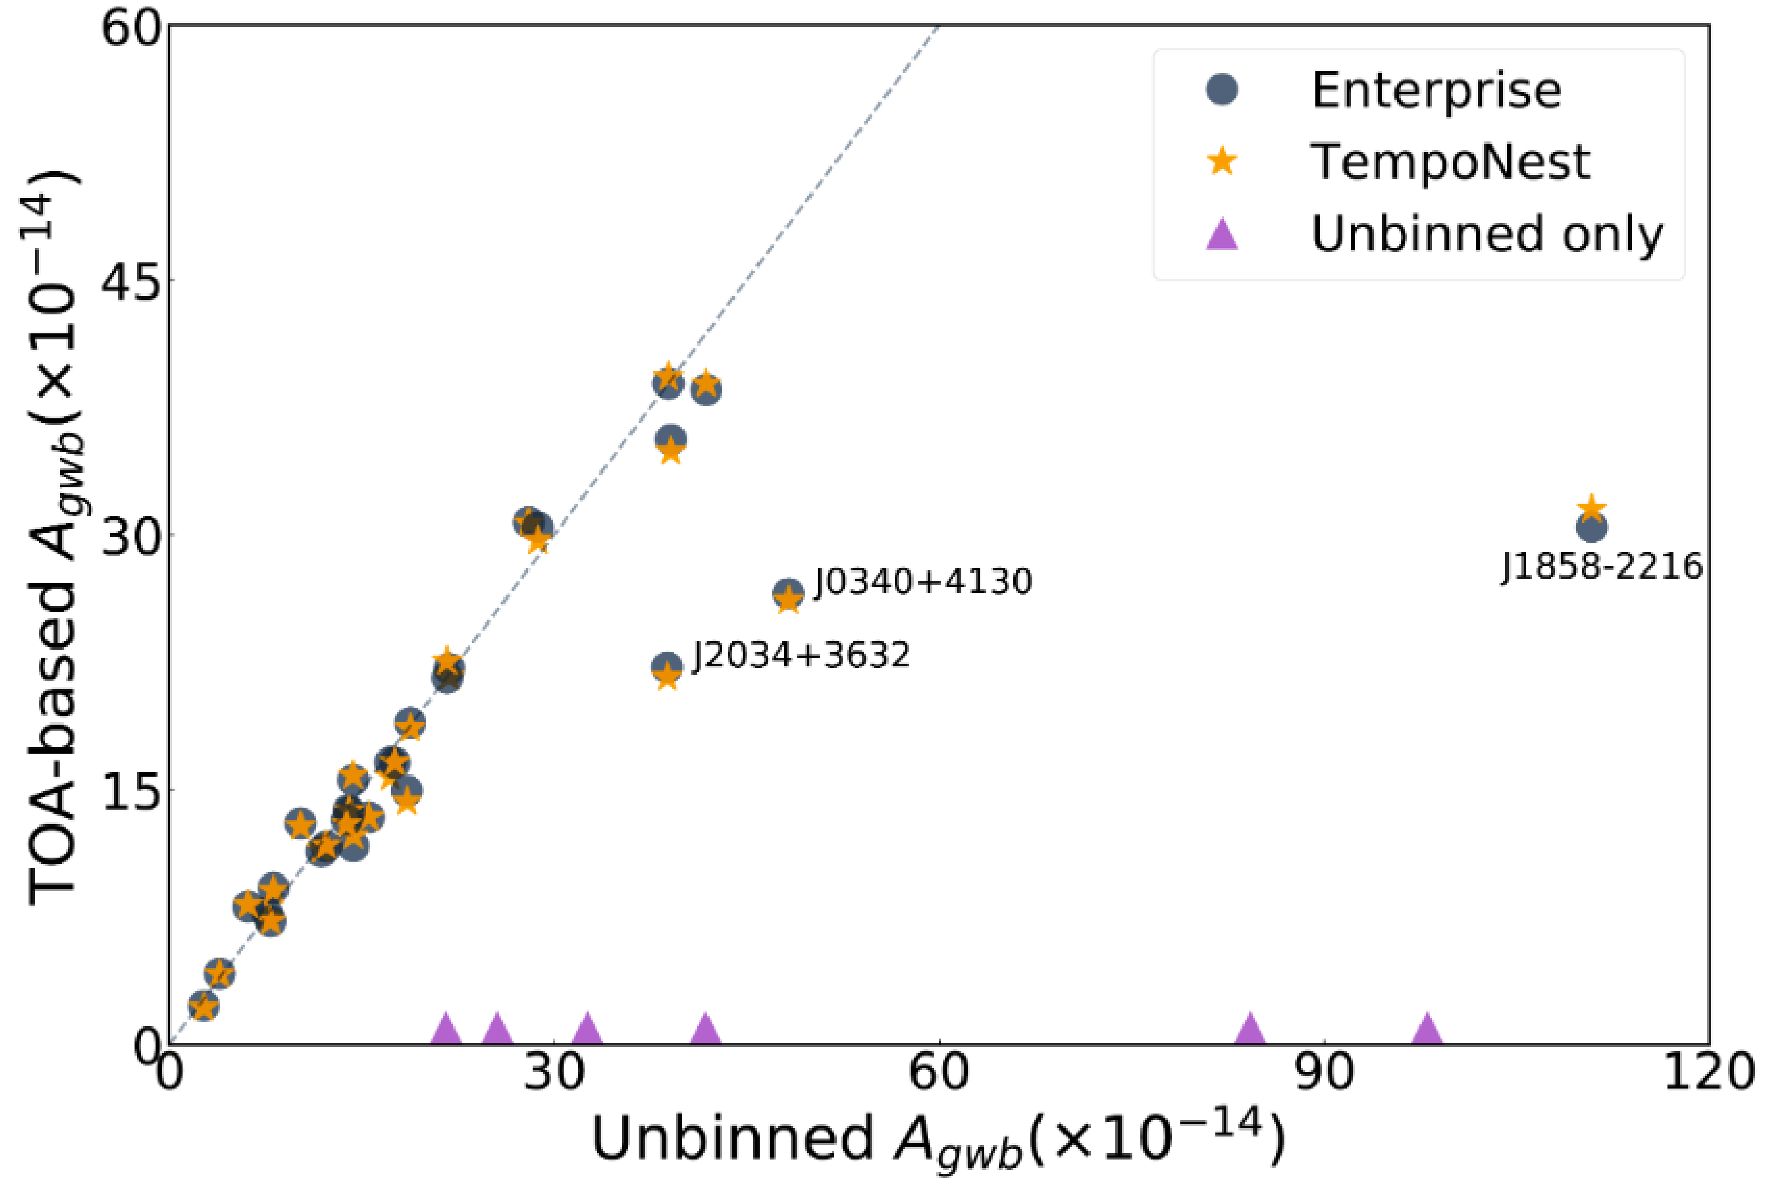

Common Mode limit and time dependence

Result: $A_{\mathrm{gwb}} < 1.0 \times 10^{-14}$

- Assuming no GWB low-frequency cutoff and applying time scaling, will reach $A_{\mathrm{gwb}} = 2 \times 10^{-15}$ with another 12 years of data (double).
  - Fermi has no consumables, orbit is good for decades.

- Improvements on the method can give ~20% better limits:
  - Energy-resolved pulse profiles, better code to allow compact binaries, data selection optimization…
- Additional MSP discoveries can give another 5-20%.

- <u>**With improvements, reach upper end of candidate range by 2025, probe full range within 10 years!**</u>

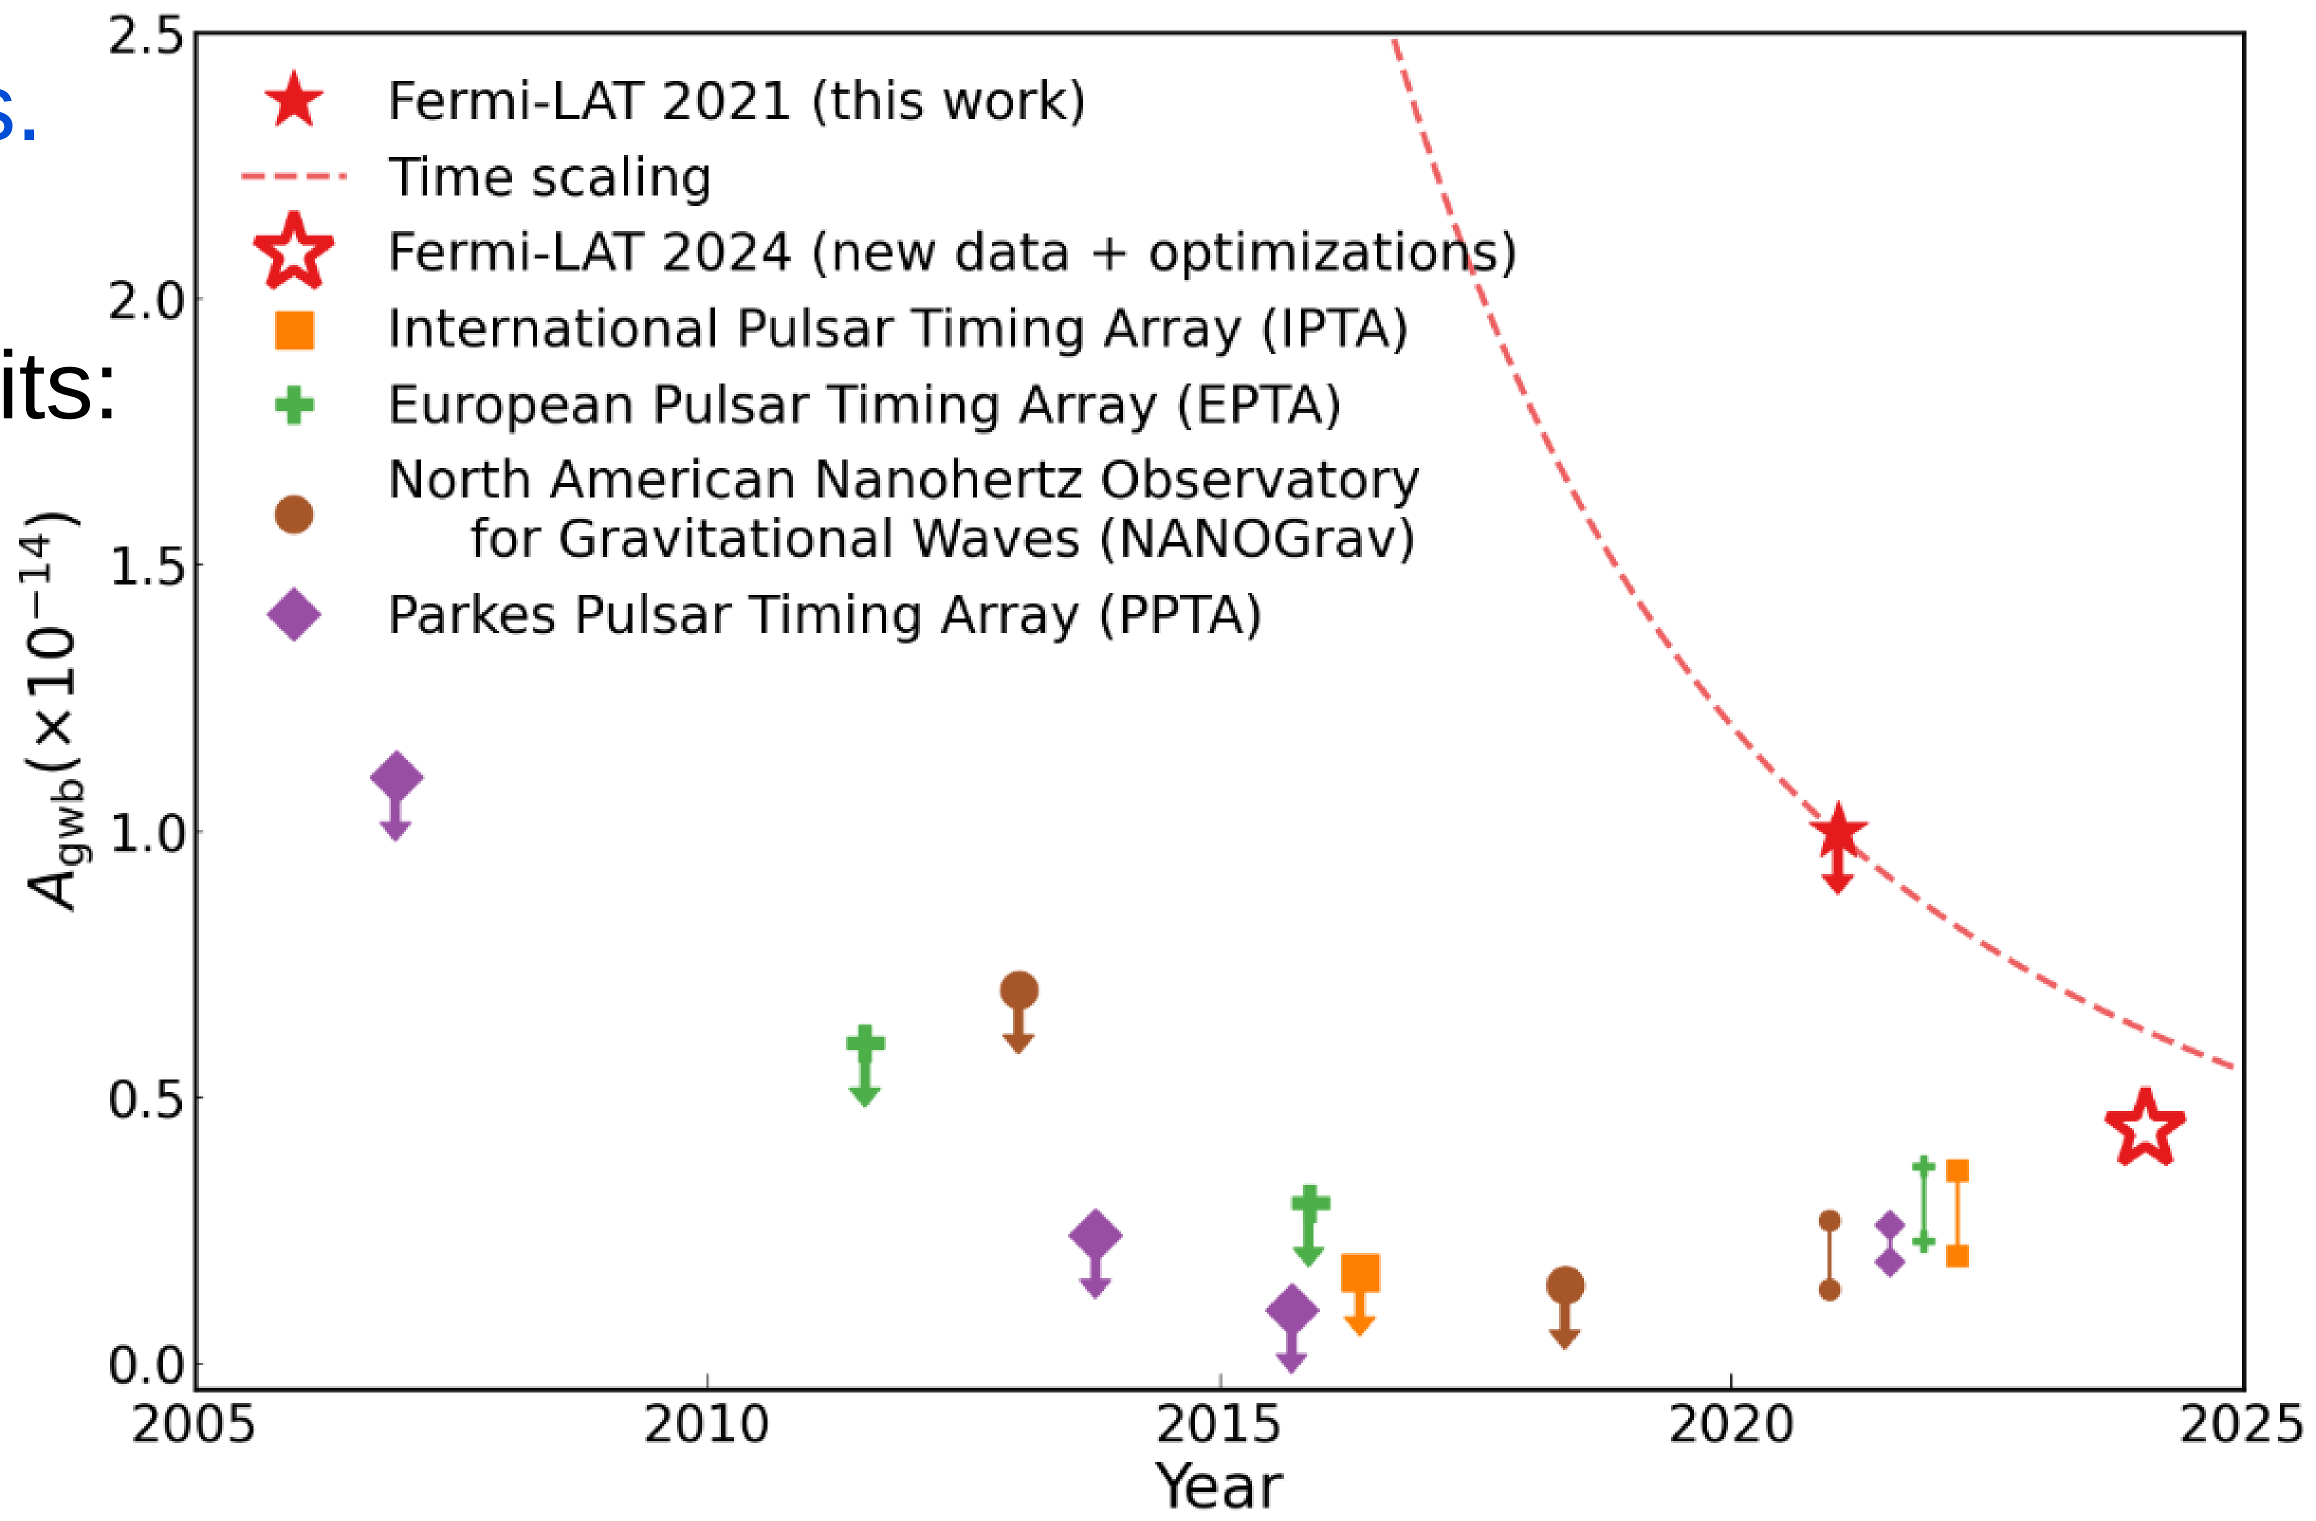

- **The overlap of "well-characterized" MSPs common to radio and gamma-ray PTAs is small (3), but 2/3 show evidence of uncorrected IISM noise.**

Where Are We Now?

- Update Fermi data: 12.5 to 14.1 years. 13% increase in data → Expect limit = 75%.
  - Use PSF types for weights calculation (c.f. Front/Back).
  - MSPs have the highest cutoff energies → raise cut from 10 to 30 GeV. (But see next slide.)

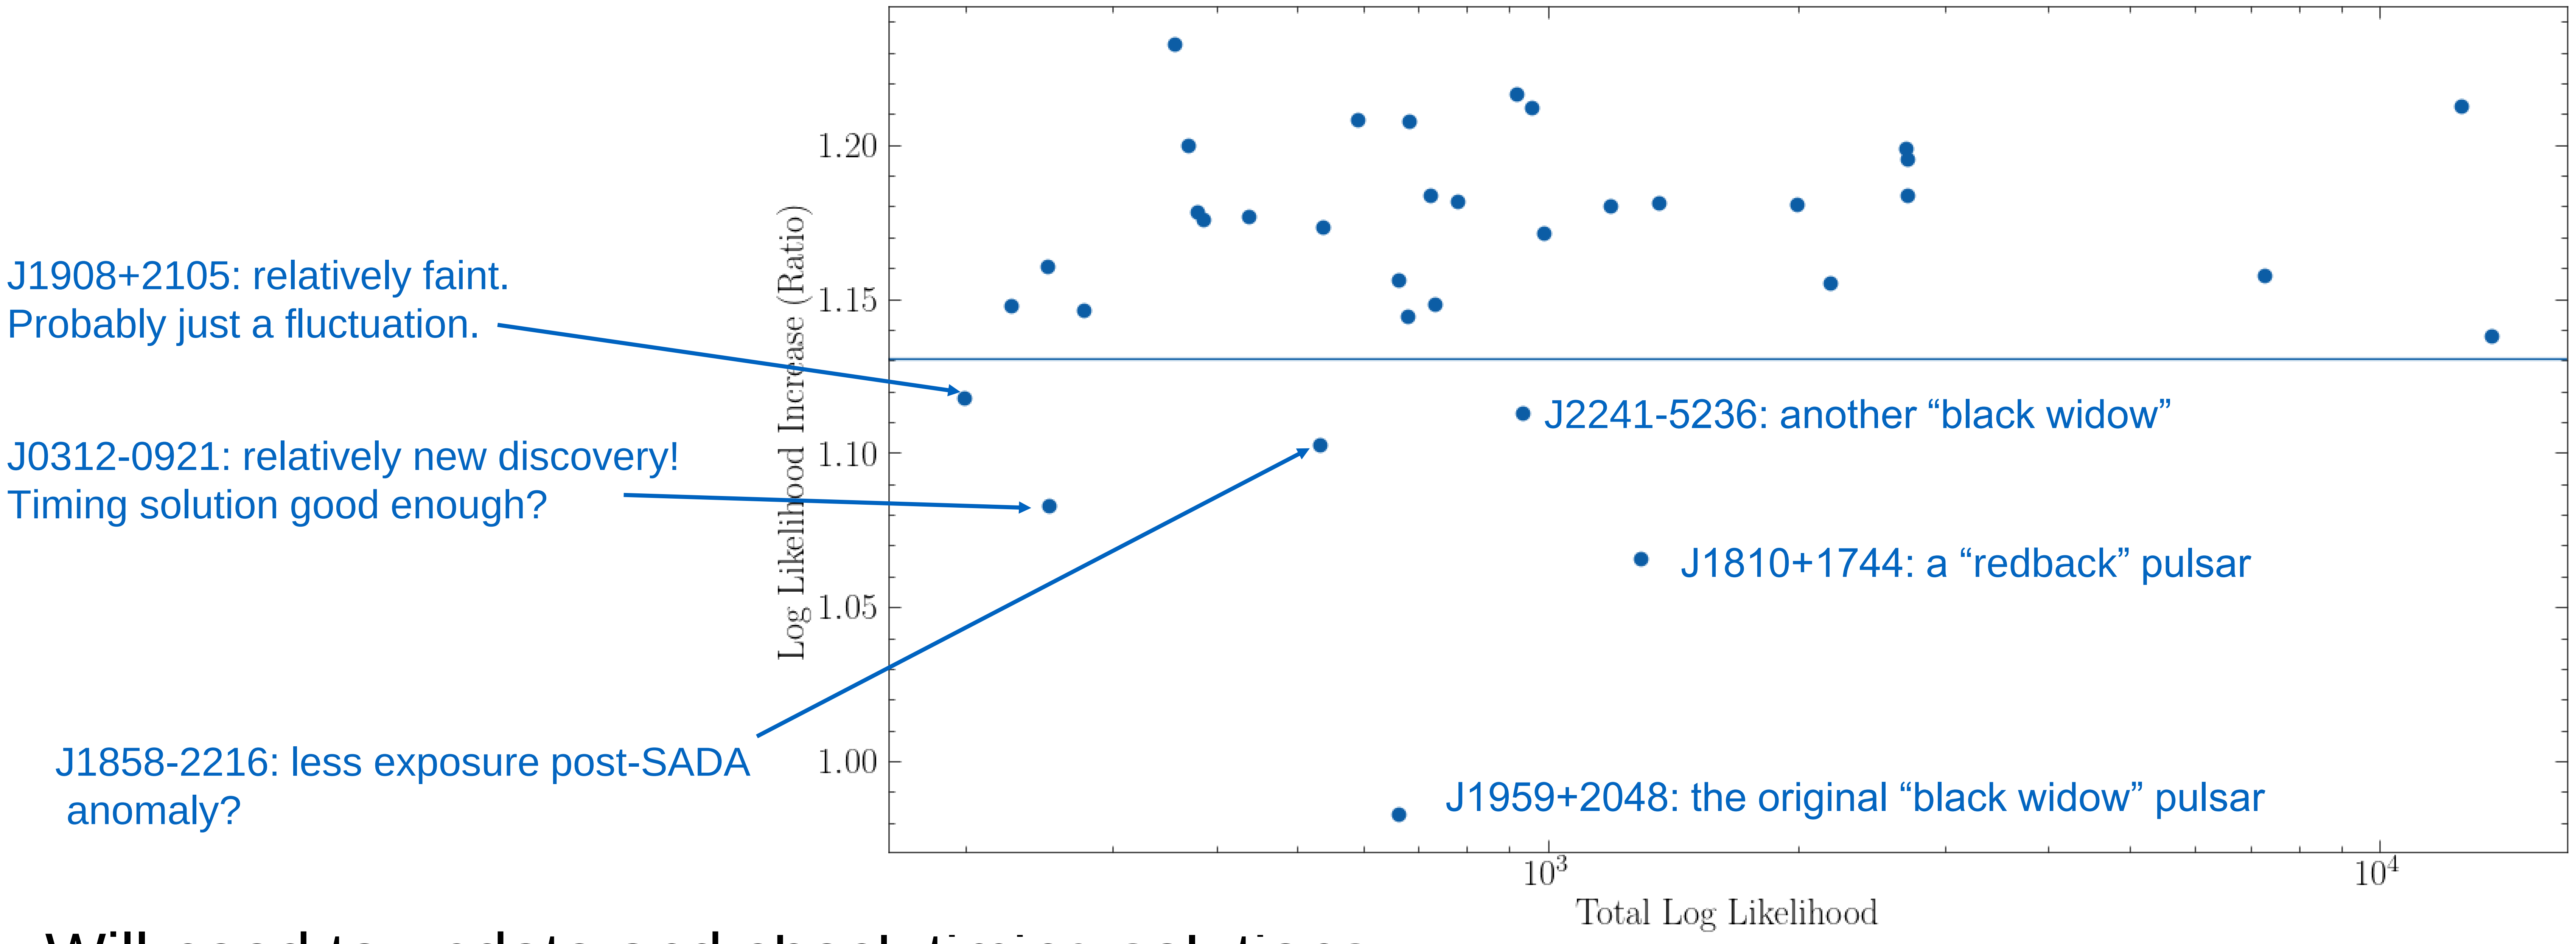

- Will need to update and check timing solutions.
  - Will want to coordinate with radio astronomers to collect constraints on timing model parameters.
  - If possible, introduce constrained orbital-period variation modeling to improve extrapolation / reduce degreee of freedom.

Energy Dependence

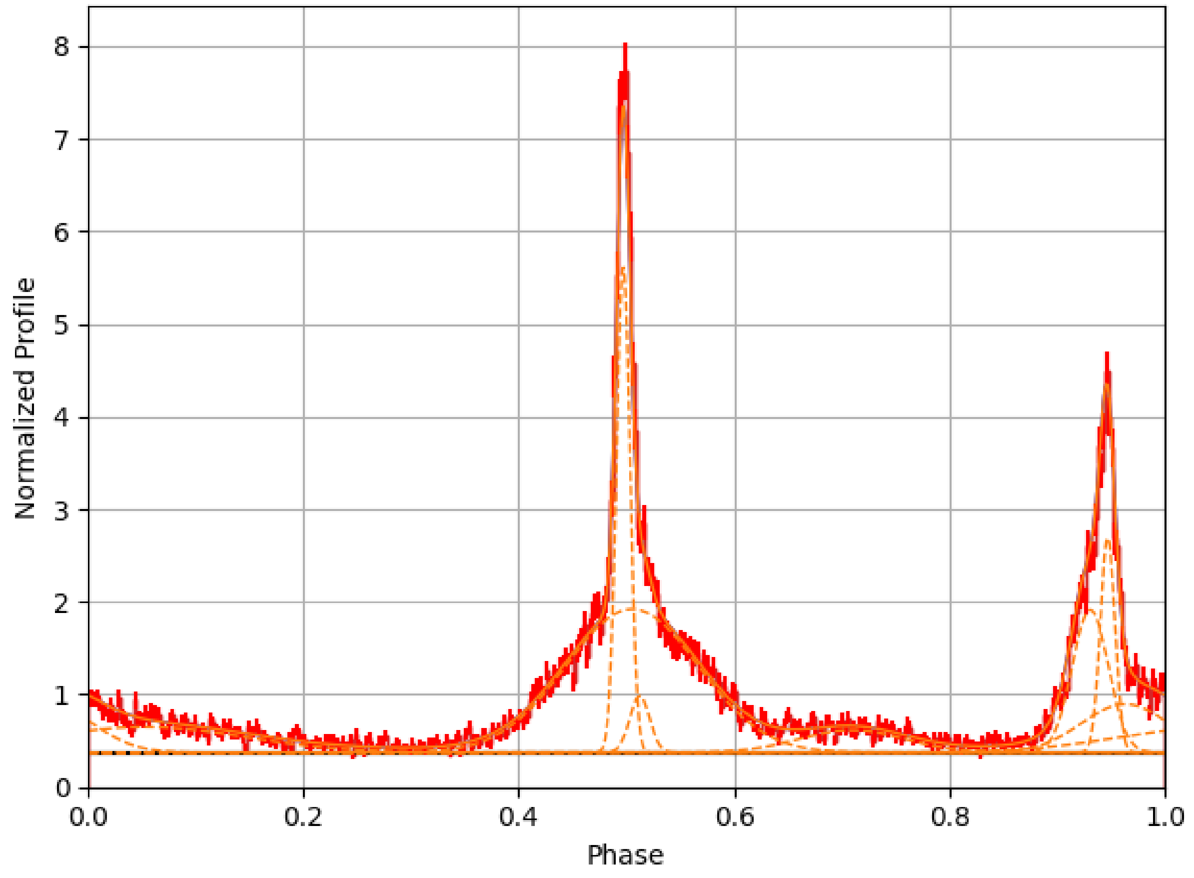

0.1-10 GeV profiles **lovingly** crafted for *Science* analysis, all generally good descriptions of the data.

This is J0614-3329, the 2nd-brightest MSP (almost the brightest).

But what about energy dependence? We know pulse components evolve with energy.

- In young pulsars, it's the famous P1/P2 ratio (decreases) with energy.
- In MSPs it's a little more complicated, but there.

How strong is the effect and:

- What is the "statistical" relevance?
- What is the systematic relevant?

Energy Dependence

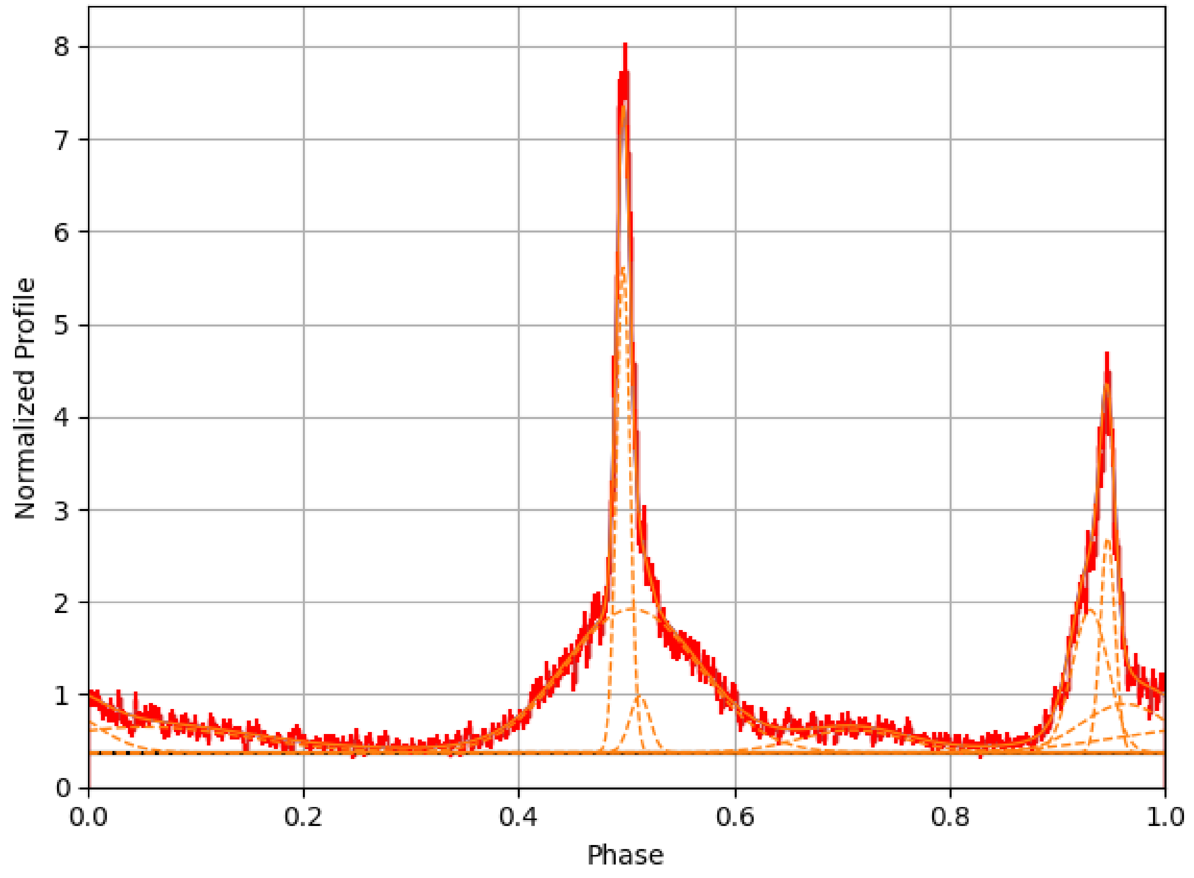

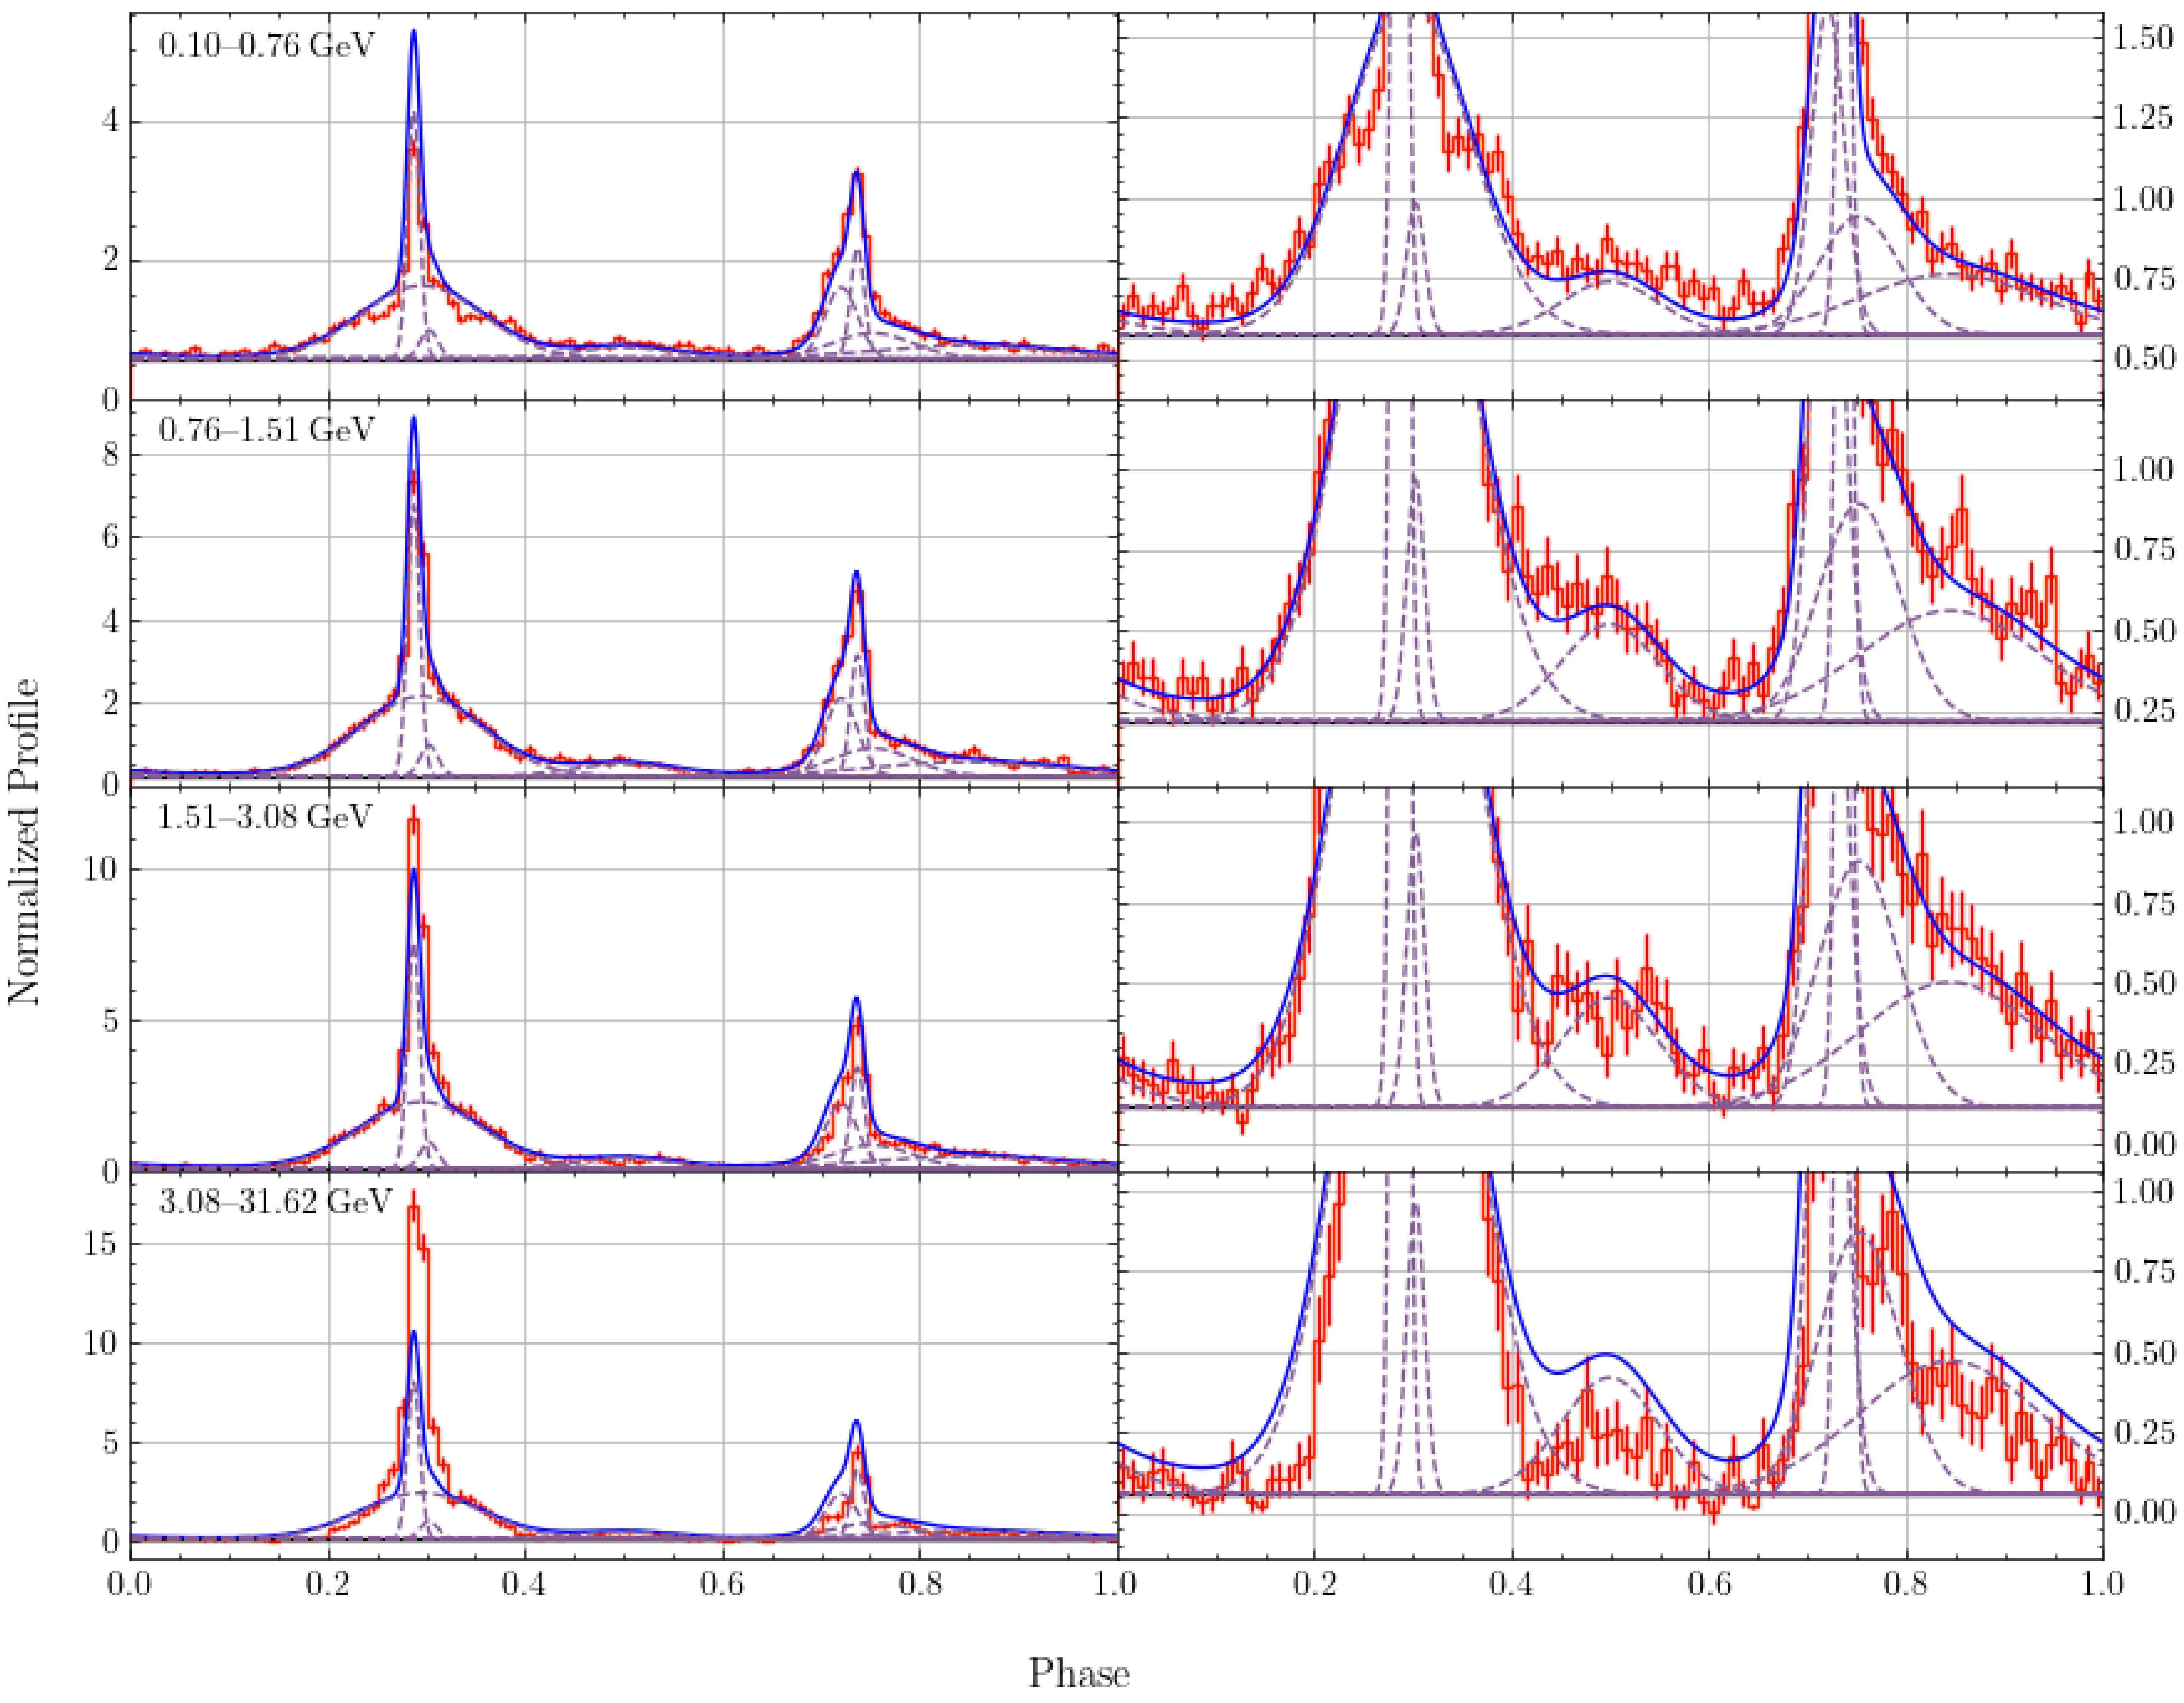

The plot shows 4 energy bands selected such that the significance of each band is equal. (Note the sweet spot is clearly 1-3 GeV!)

It is clear that the energy-independent template really only matches the ~1 GeV data.

For J0614-3329, there is clear peak evolution – leading peak gets (much) stronger at high energy, and the overall emission complex narrows.

The *effective centroid* shifts right on this plot as you go up in energy. This potentially couples energy variations with (apparent) phase variations.

It also diminishes the “worth” of our cleanest data, >3 GeV while overweighting the less-good low-energy data.

Energy Dependence

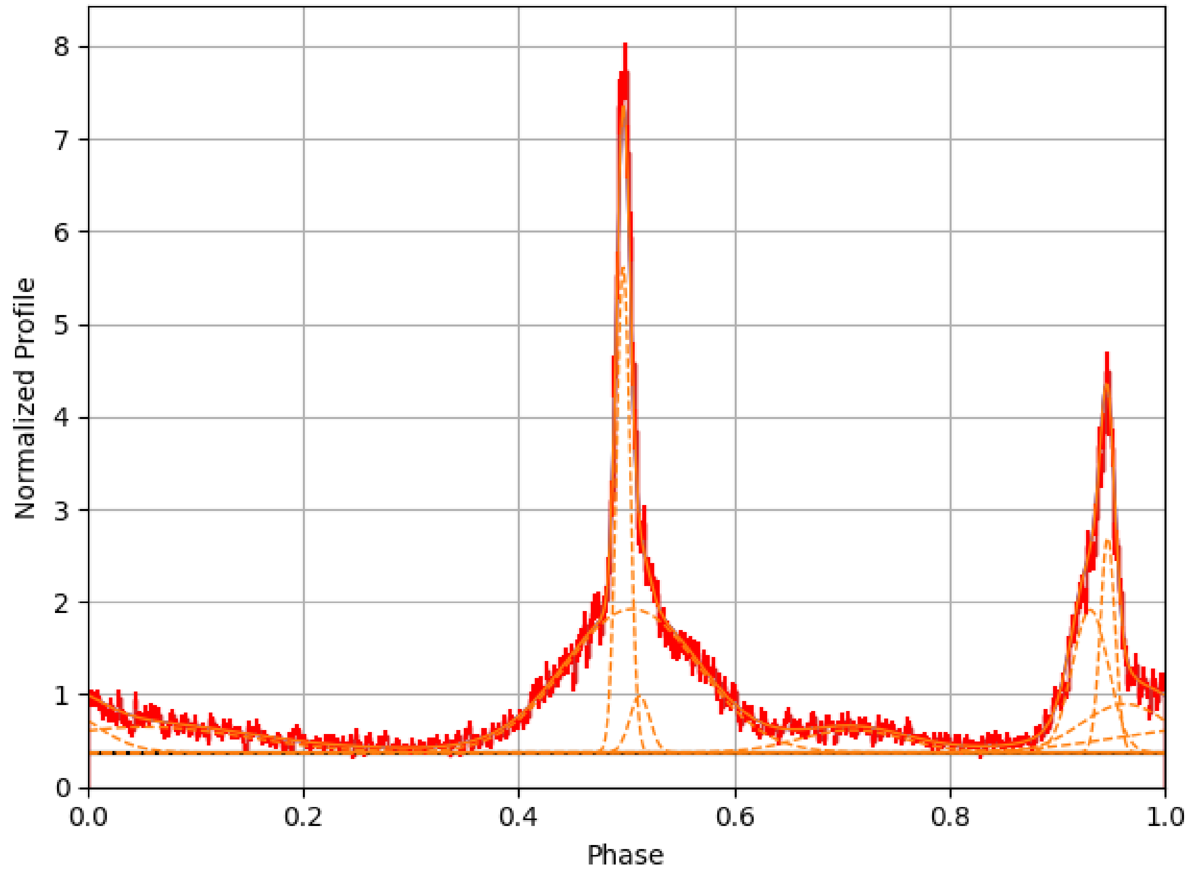

Modifying code to account for energy-dependence with a simple model: component parameters evolve linearly with log_10(energy).

WORK IN PROGRESS!

But, a simple fit with the amplitudes free improves the log likelihood by 788 (!!, about 6%).

- A similar analysis for other pulsars reveals typical improvements of a few percent. Better models with some love will help.
- This is a **BIG** potential systematic. If the GWB signal can conspire to soak up even a tiny fraction of this "spare" log likelihood, it will impact the limit/detection.
- (Except, it can't be too big, or else our limit wouldn't have been good…)

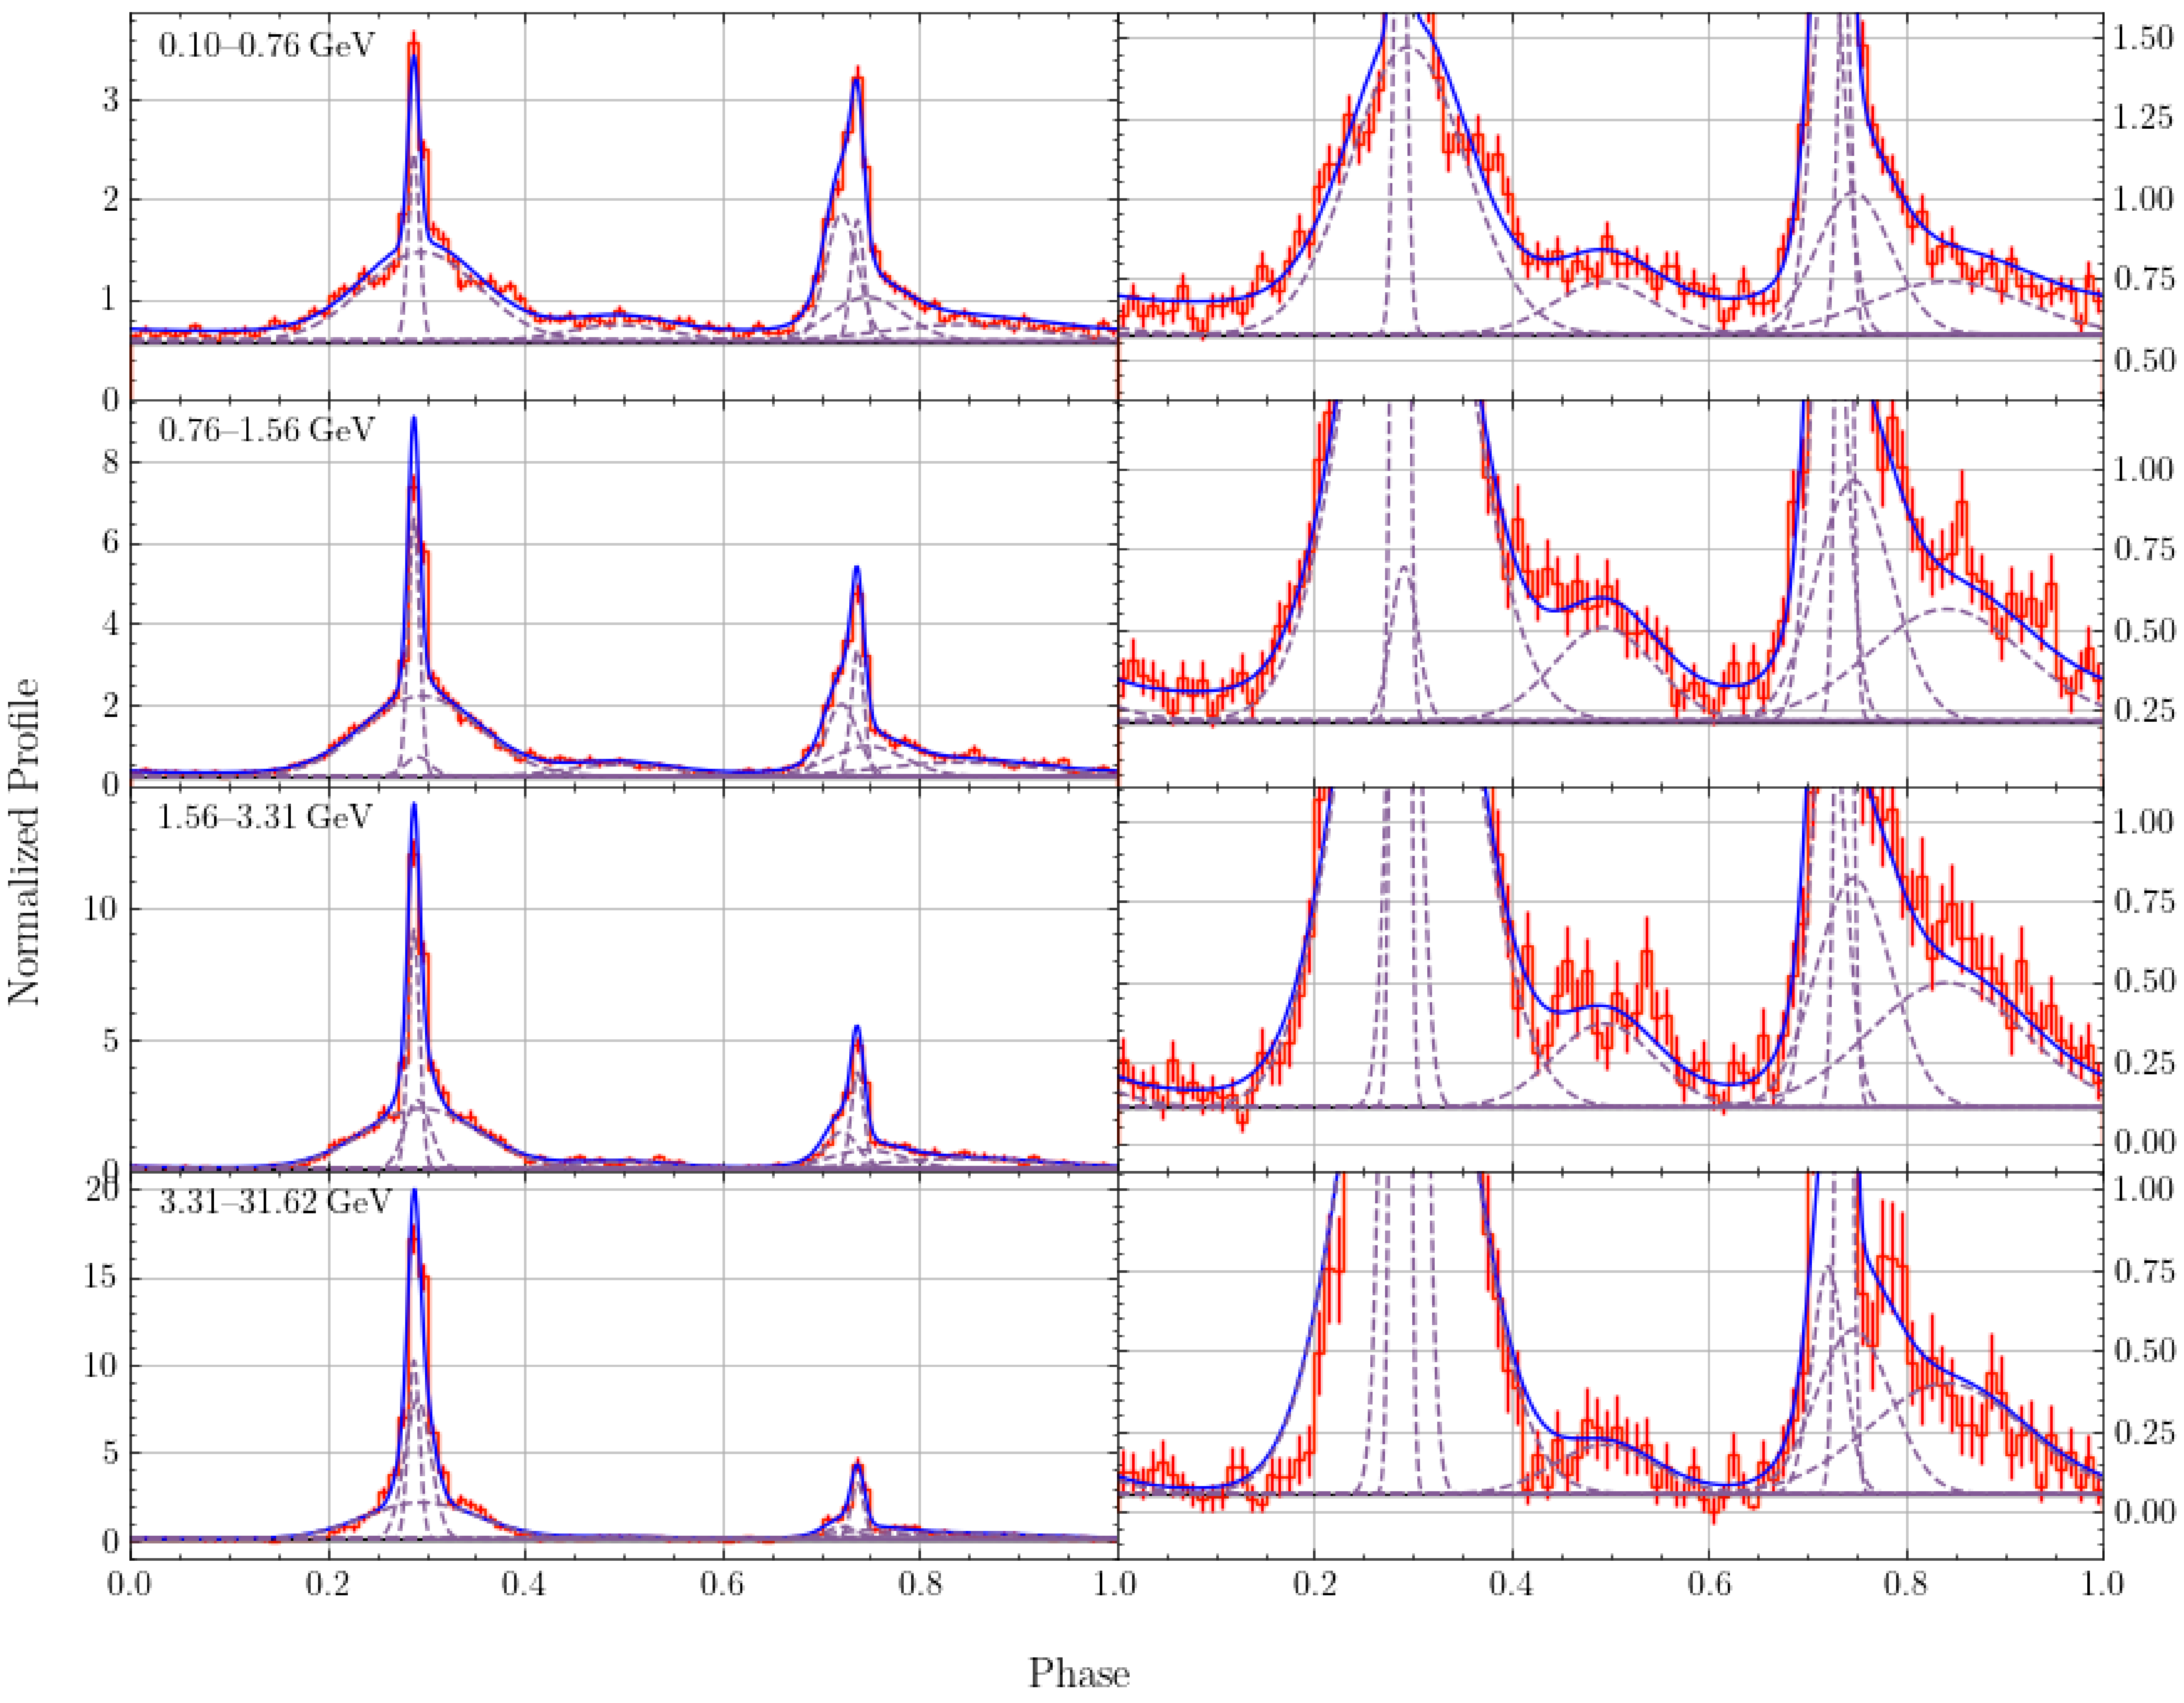

Energy Dependence

J0613-2200: substantial shift in centroid!

J0030+0451: modest centroid shift, more of a loss of precision

Systematics and Simulations

- Big selling point of Fermi for PTA work is unchanged experimental setup.
  - But, there have been some (modest) changes! Rocking profile after Year 1, Galactic center stare, and post-Solar Array Drive Assembly (SADA) anomaly rocking profile.
- For a PTA, the “sweet spot” timescale is ~T/2, and the post-SADA data is getting ever closer to T/2.
  - We **don’t** see evidence that this is a factor, but…
  - Now is the time to see if we truly are systematic free.
    - With energy-dependent templates plus exposure, can simulate a truly faithful pulsar dataset for the first time.
    - Many simulations will give true sensitivity and allow quantification of whether energy-dependent profiles can “leak” through exposure variations into the GWB signal.
- If there are systematic effects from exposure variations, we can build it into the analysis…
  - … and that would be a huge pain!

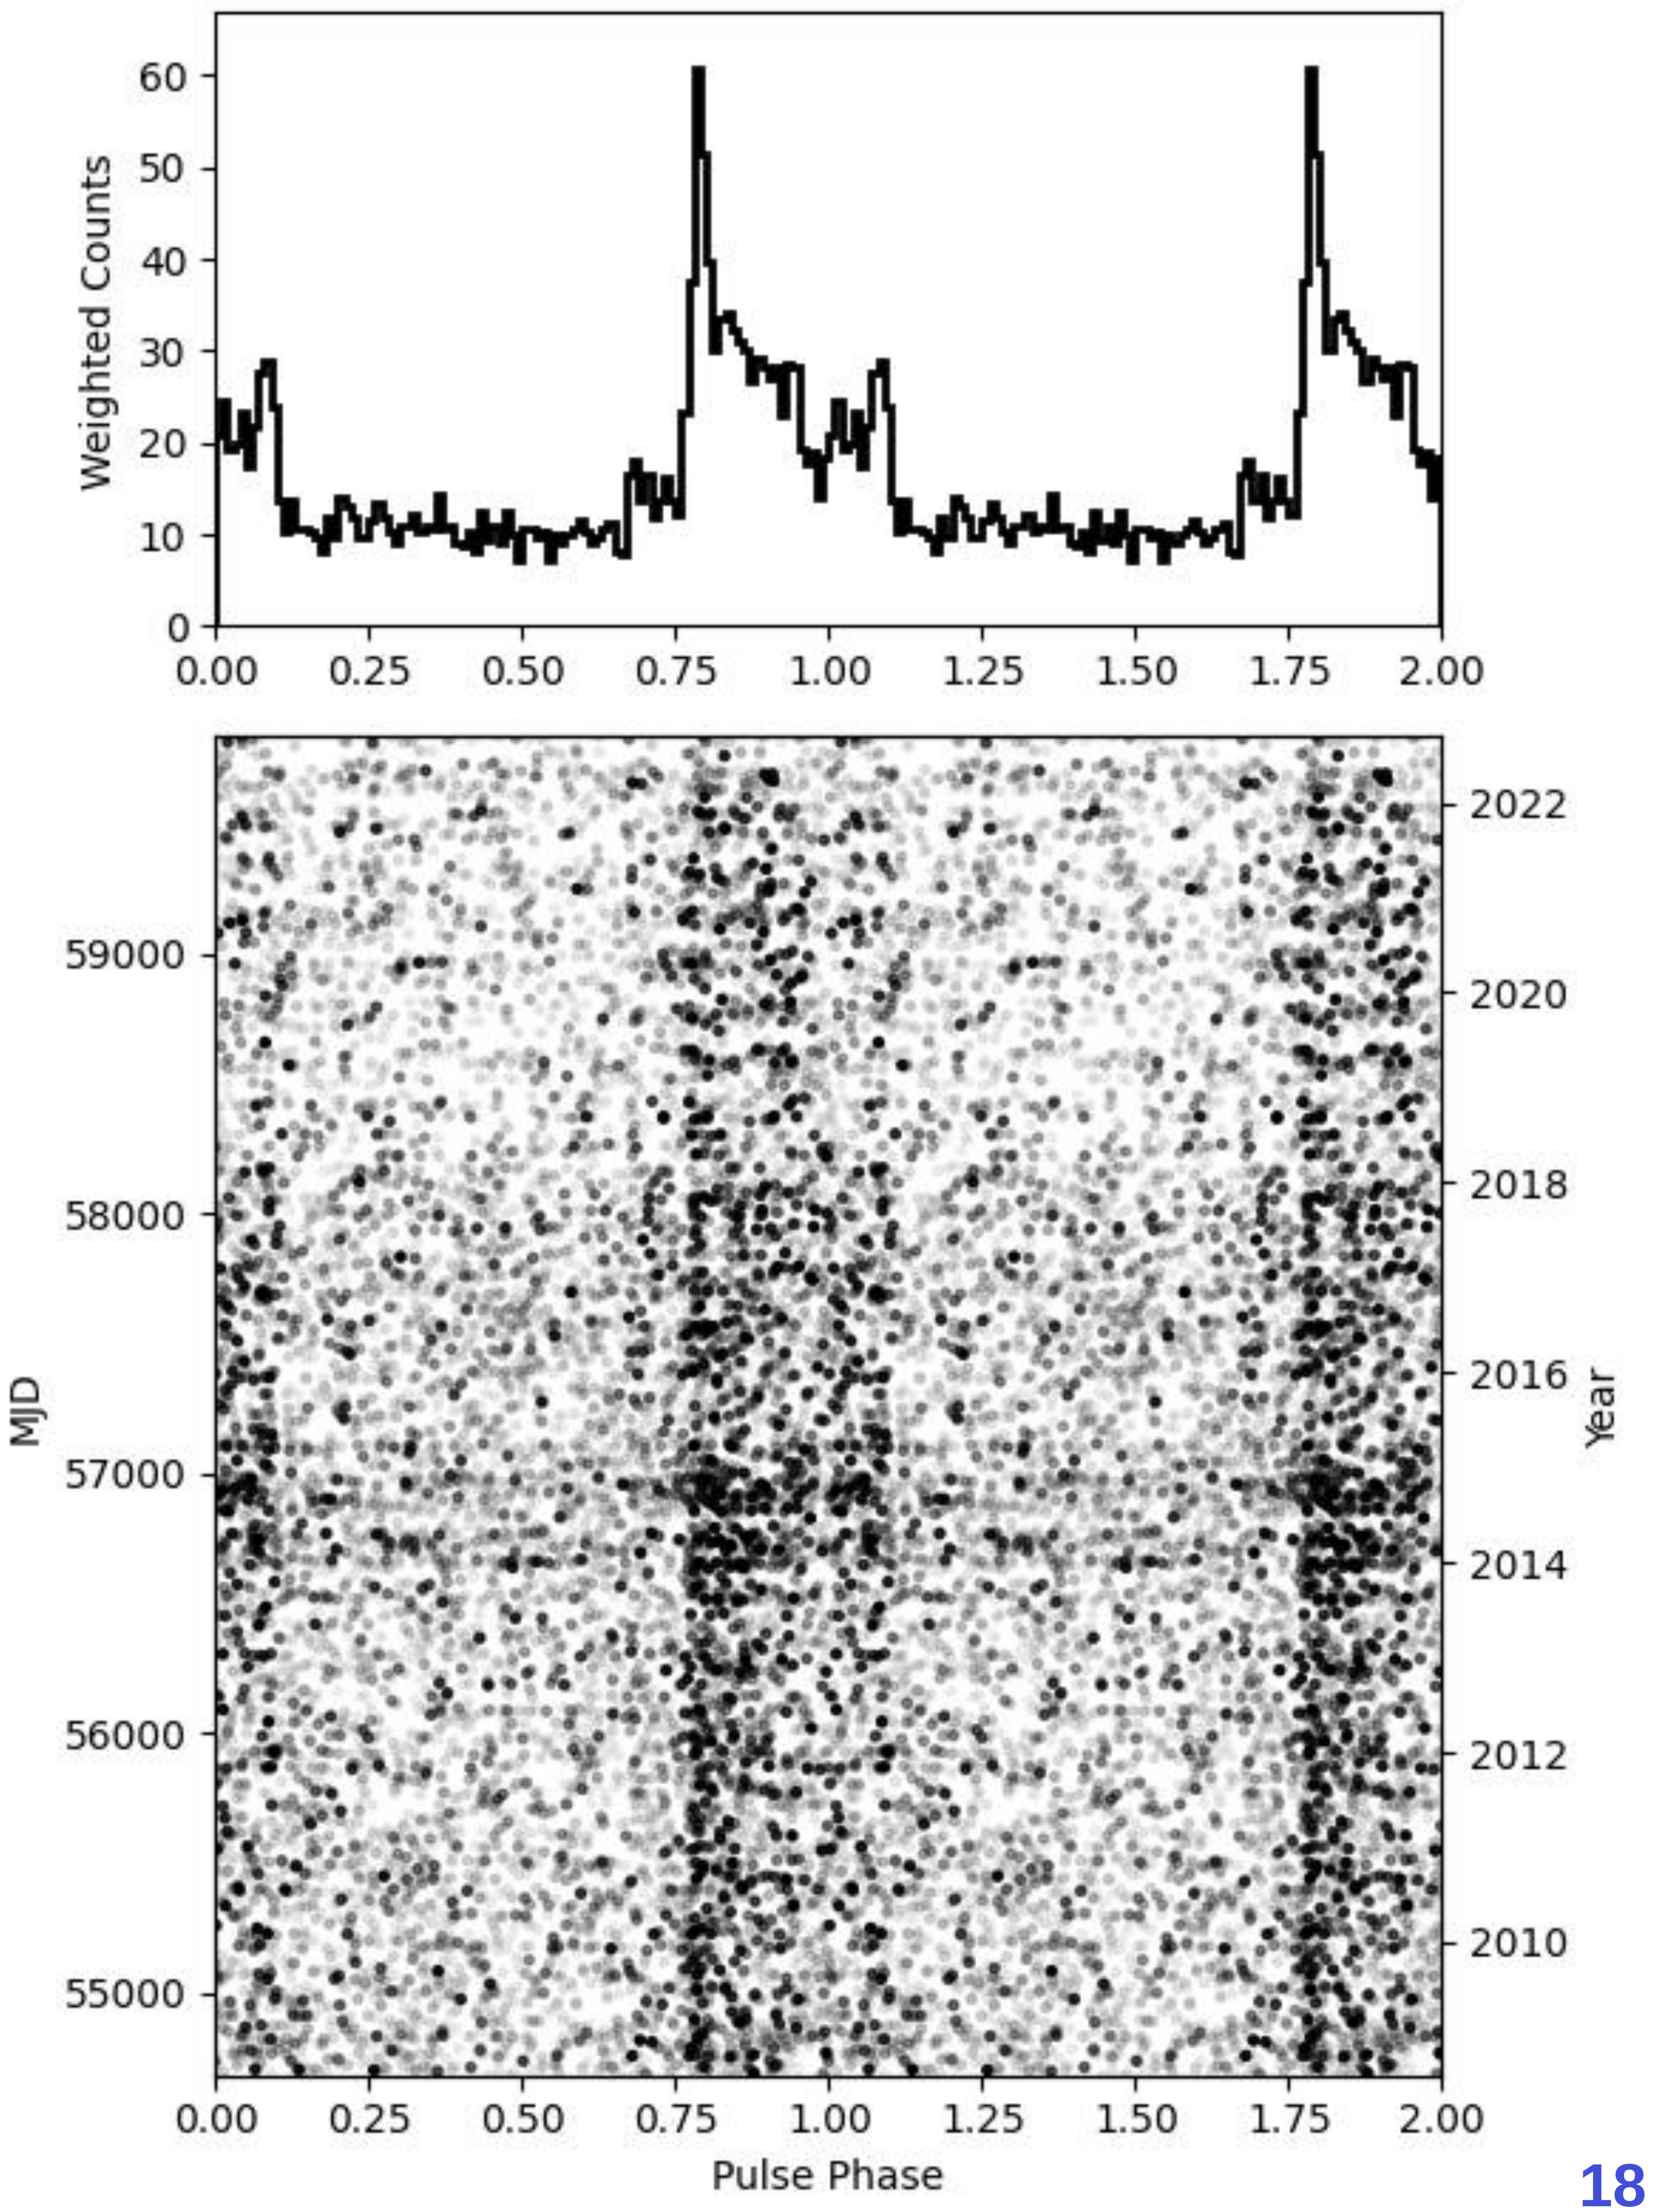

Conclusions

- Fermi LAT can deliver a surprisingly good constraint on the nHz GWB, and it is subject to a reduced—or at least, independent—set of systematic uncertainties.

- Comparison of Fermi noise models to radio suggest some level of contamination by IISM.
  - Need more radio timing of gamma-ray MSPs.

- Time scaling very favorable, will probe candidate signal in 5—10 years.
  - Need to keep photons coming.
  - Technique improvements key to reaching GWB candidate signals quickly.

- Updating data set – collaborate with radio astronomer colleagues, updating timing solutions, study systematics, perform simulations, …
  - A new limit by early 2023? (May bump data up to 15 yr for a nice round number.)

- Support for project with regular data release and analysis through Fermi GI.

Pulsar Timing with Radio Telescopes

- Obtain “folded” radio profiles using big dishes.
- Cross-correlate with “standard” to estimate offset (relative to observatory clock).
- This is a “pulse time of arrival” or TOA.
- Compare to a predictive model to estimate parameters. (Position, proper motion, parallax, binary period, post-Keplerian parameters…)

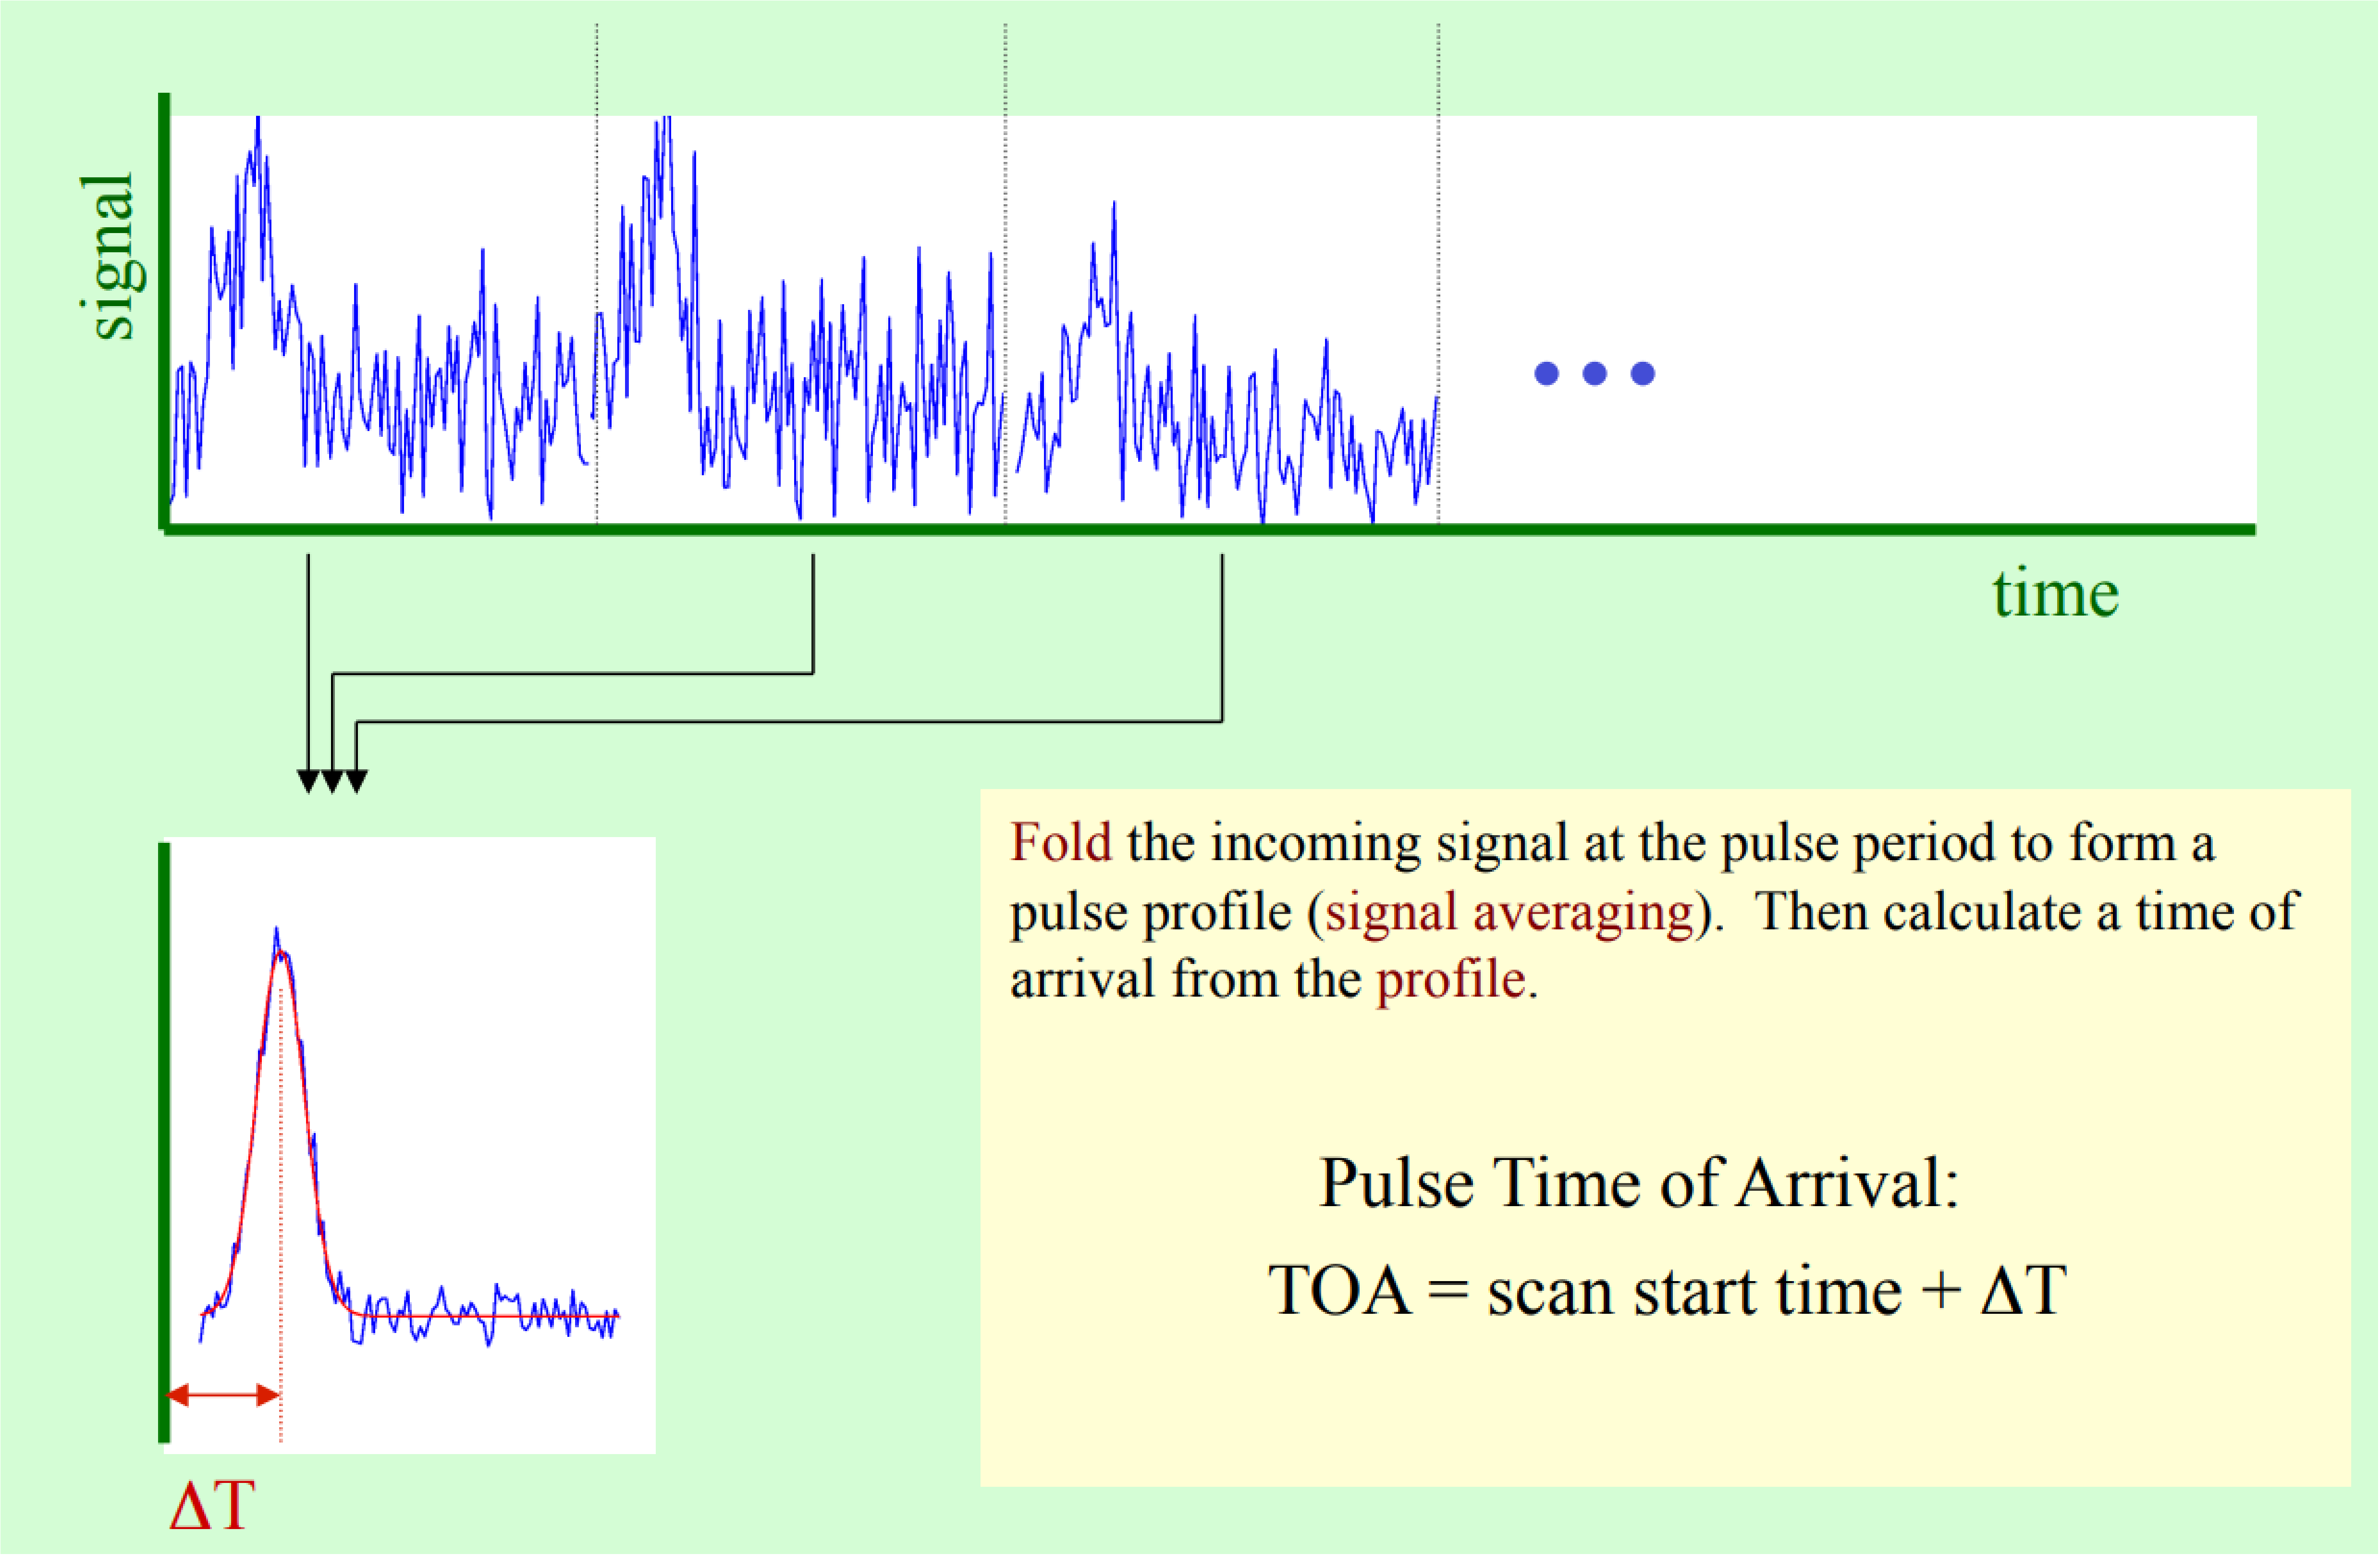

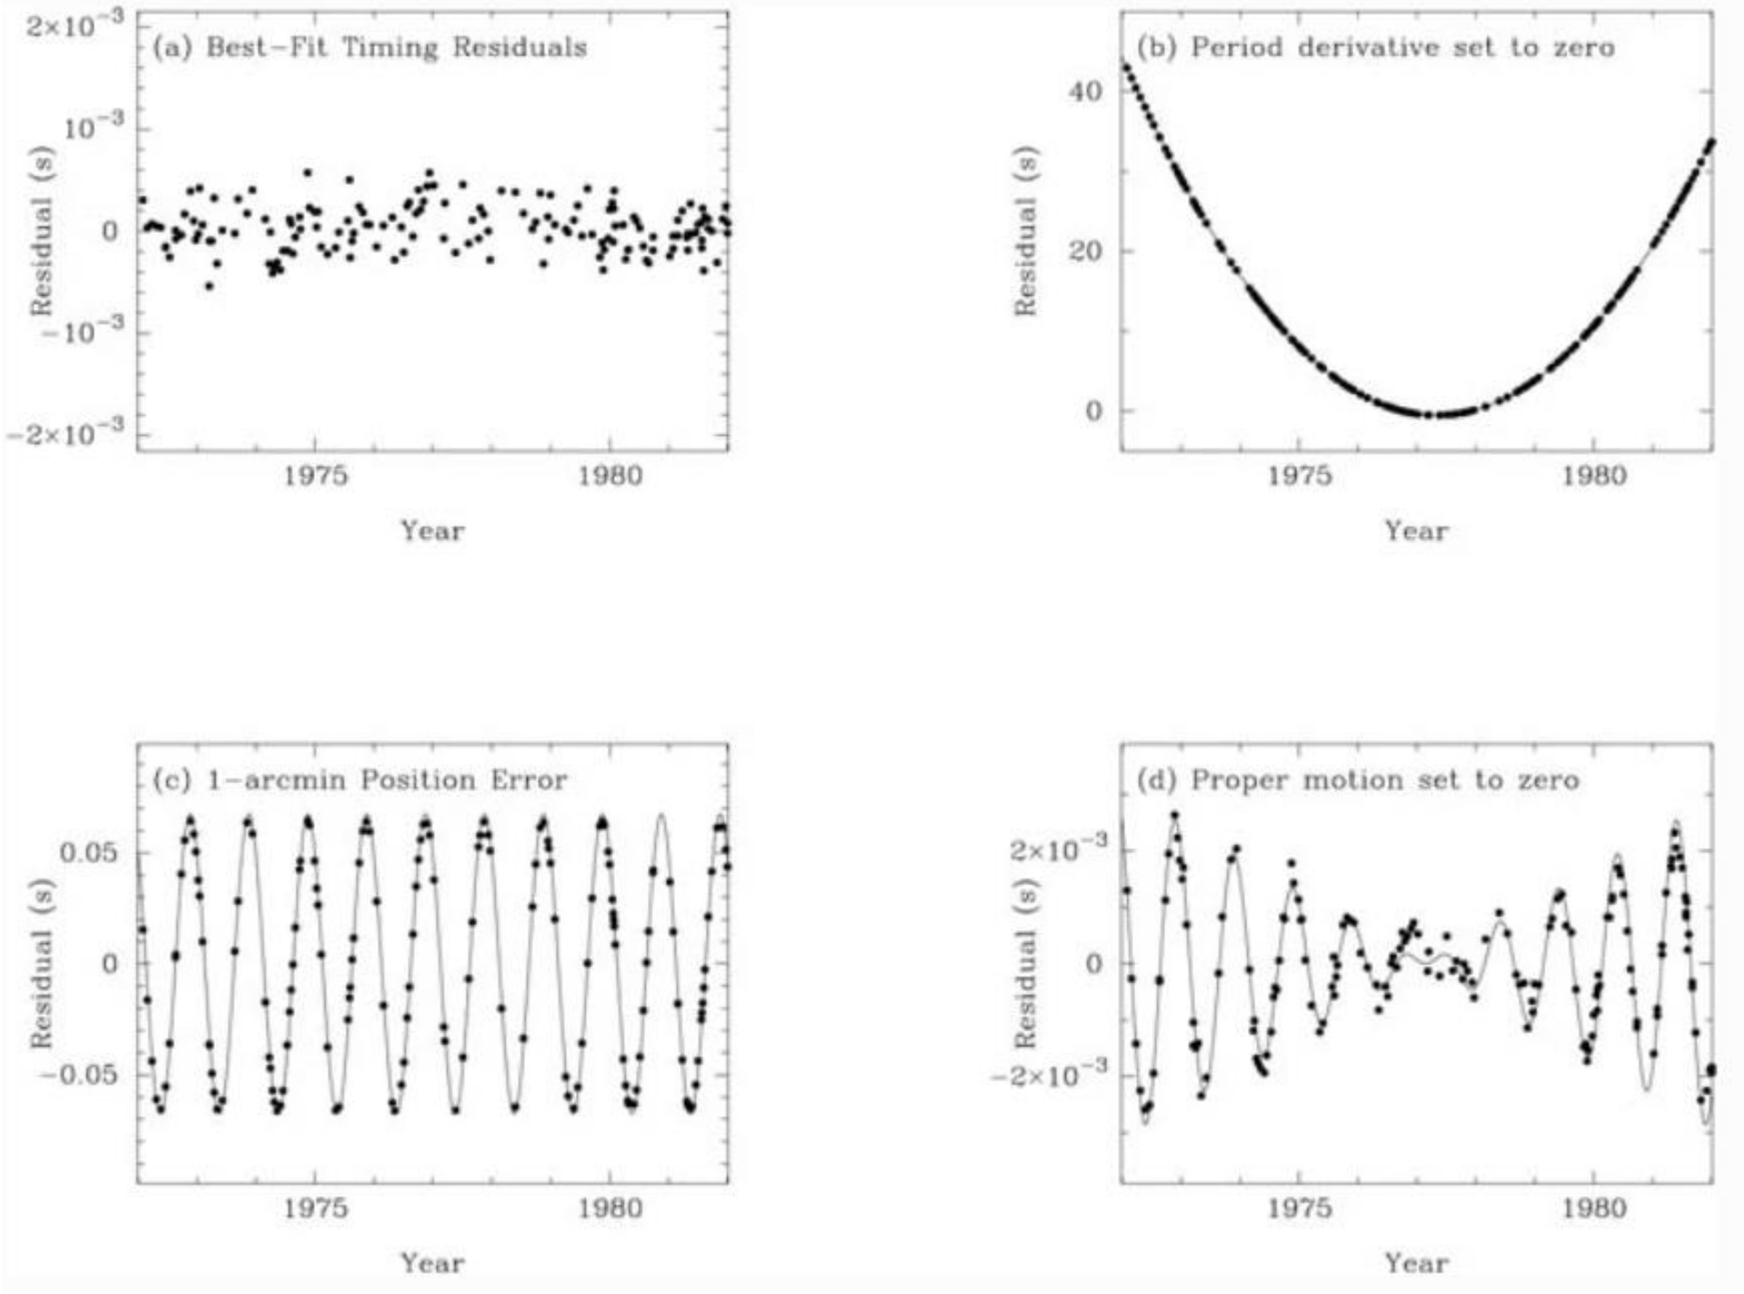

Major PTAs

- NANOGrav: Green Bank Observatory, Arecibo Observatory, limited VLA
- EPTA: Effelsberg, Nancay, Sardinia Radio Telescope, Jodrell Bank, …
- PPTA: Parkes Observatory
- IPTA: “consortium of consortia” – combines data from member PTAs

- These PTAs have >10 years of data now.

- New PTAs / data sets: FAST, In(dia)PTA, MeerTIME (MeerKAT)

- Ultimately, SKA: pulsars underpin many key science projects.

Fermi and MSPs

Fermi has detected known MSPs (with “timing solutions” from radio colleagues) and has discovered many MSPs by guiding radio telescope searches (Pulsar Search Consortium, Ray et al., 2012).

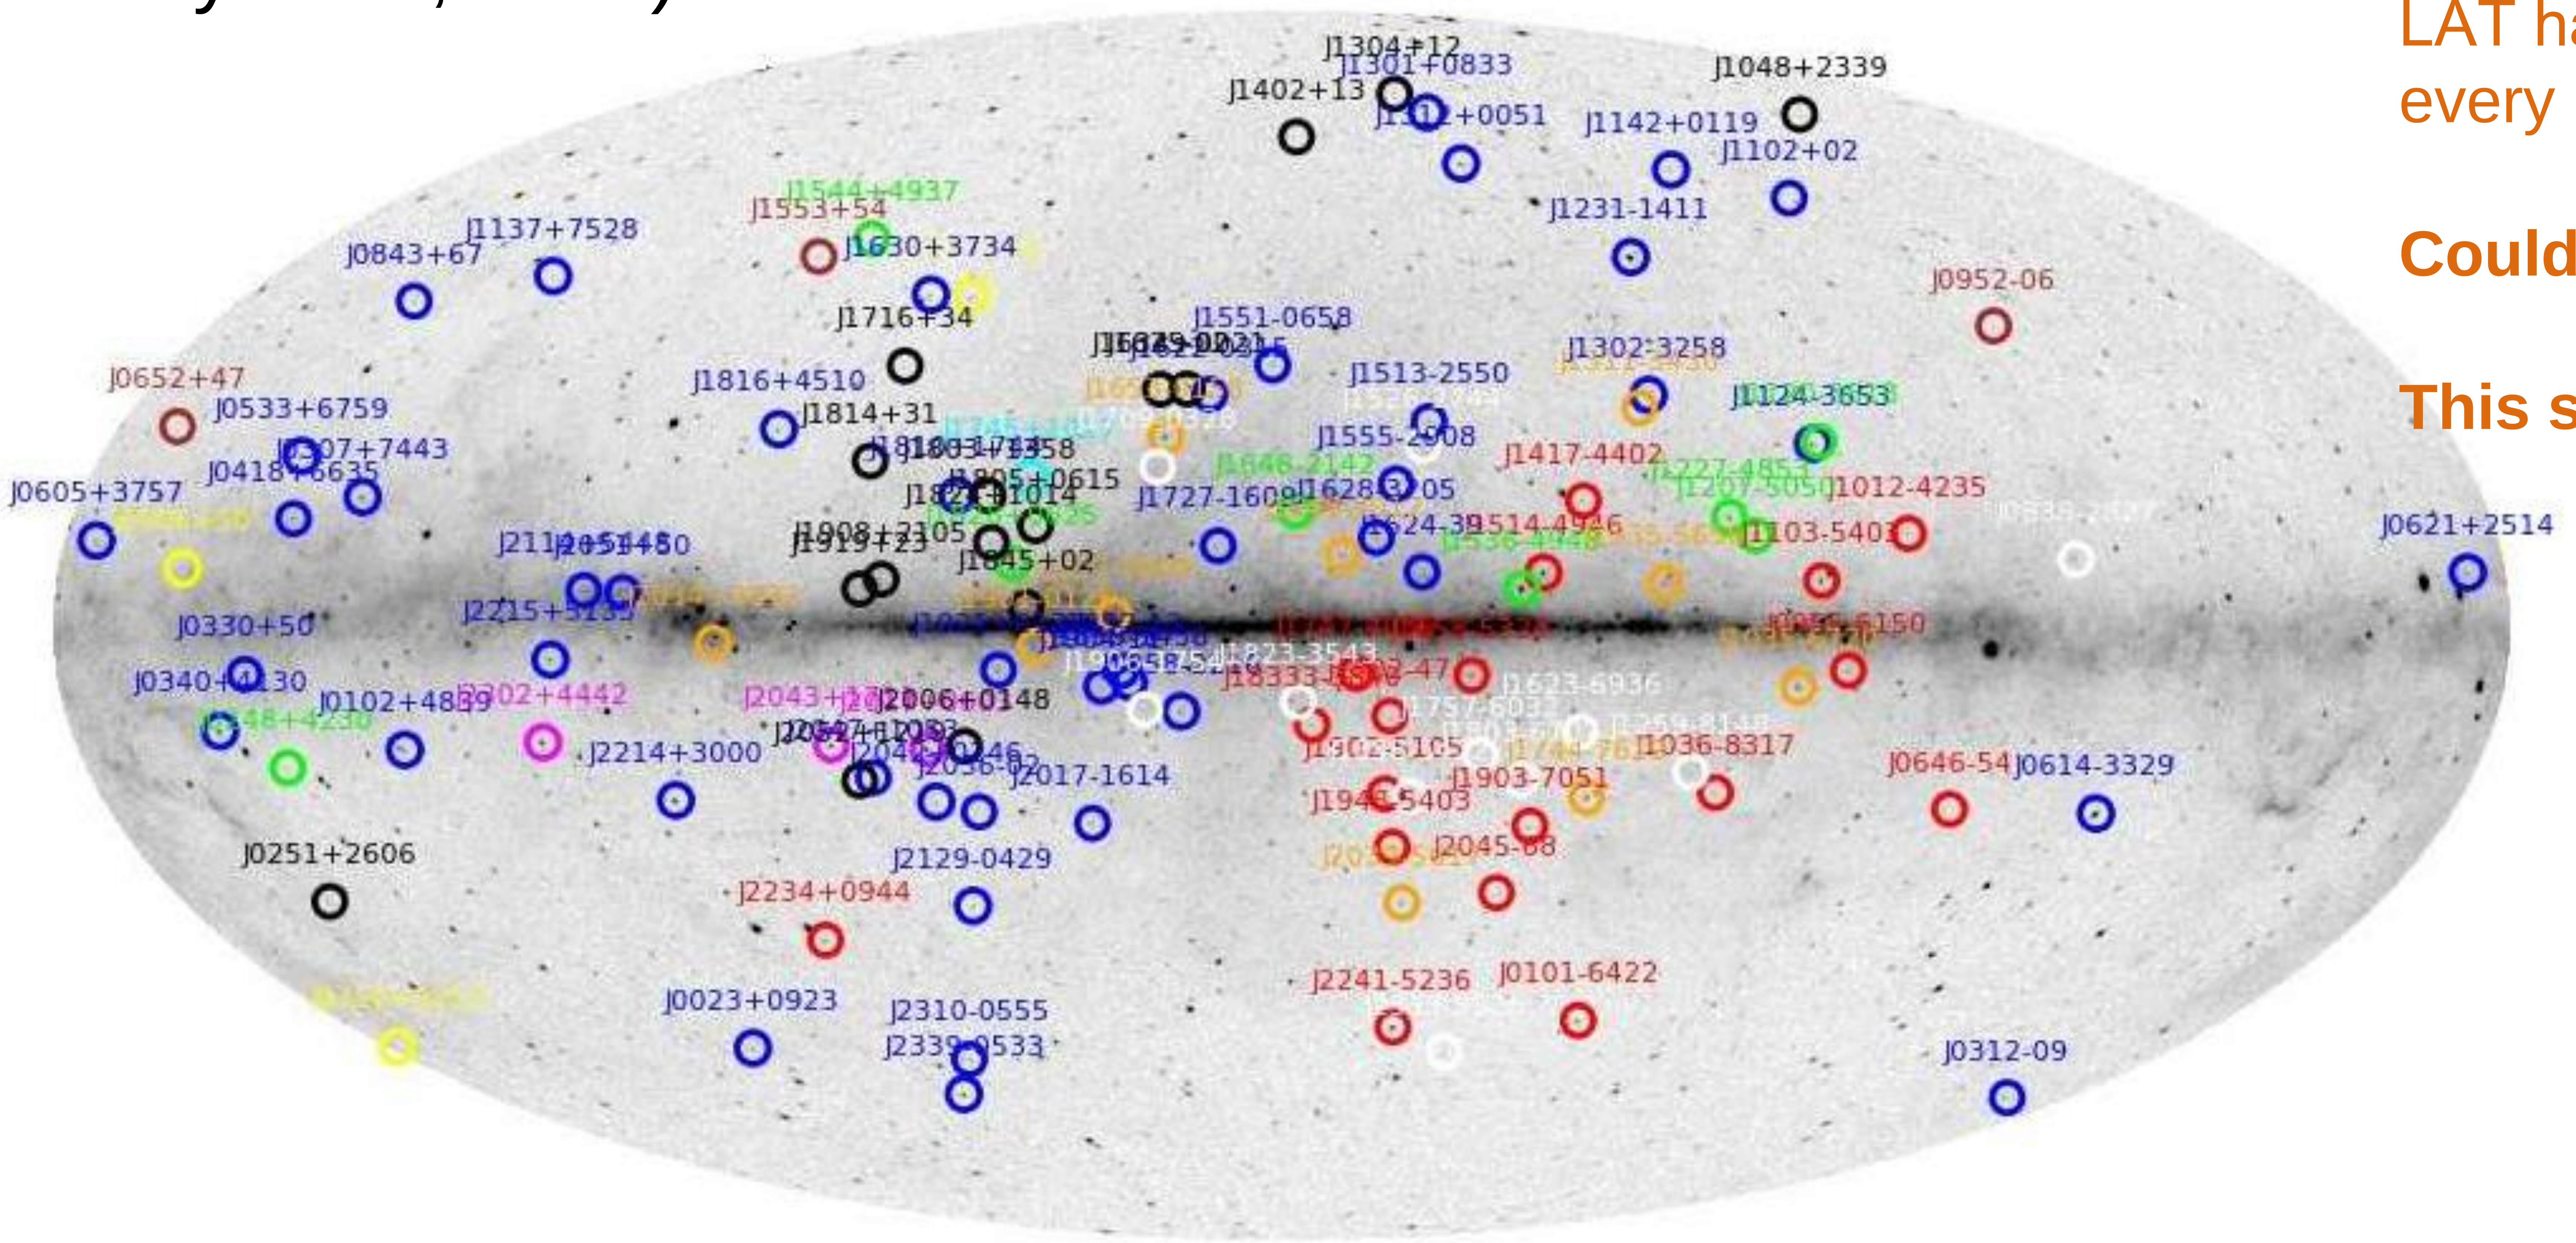

Over 130 MSPs now detected with Fermi-LAT.

LAT has a good clock… and it “observes” every MSP in the sky every day.

Could it be a good PTA?

This slide could go or be consolidated.

What do Fermi data look like?

- **Every photon is time tagged and archived.**
  - **We are continually recording data from pulsars we don't even know about!**

- **Photons can be selected and converted to pulse phase with an ephemeris and a tool like PINT, the "Fermi" plugin for tempo2, etc.**

- **Photon weights help separate photons from the pulsar from all other gamma-ray sources.**
  - **Sidebar: the LAT has a relatively large PSF, ~1 deg at 1 GeV. So essentially all sources are confused.**
  - **A photon weight is based on a model of every gamma-ray source, folded through the instrument response, to estimate the probability the photon comes from a specific source.**

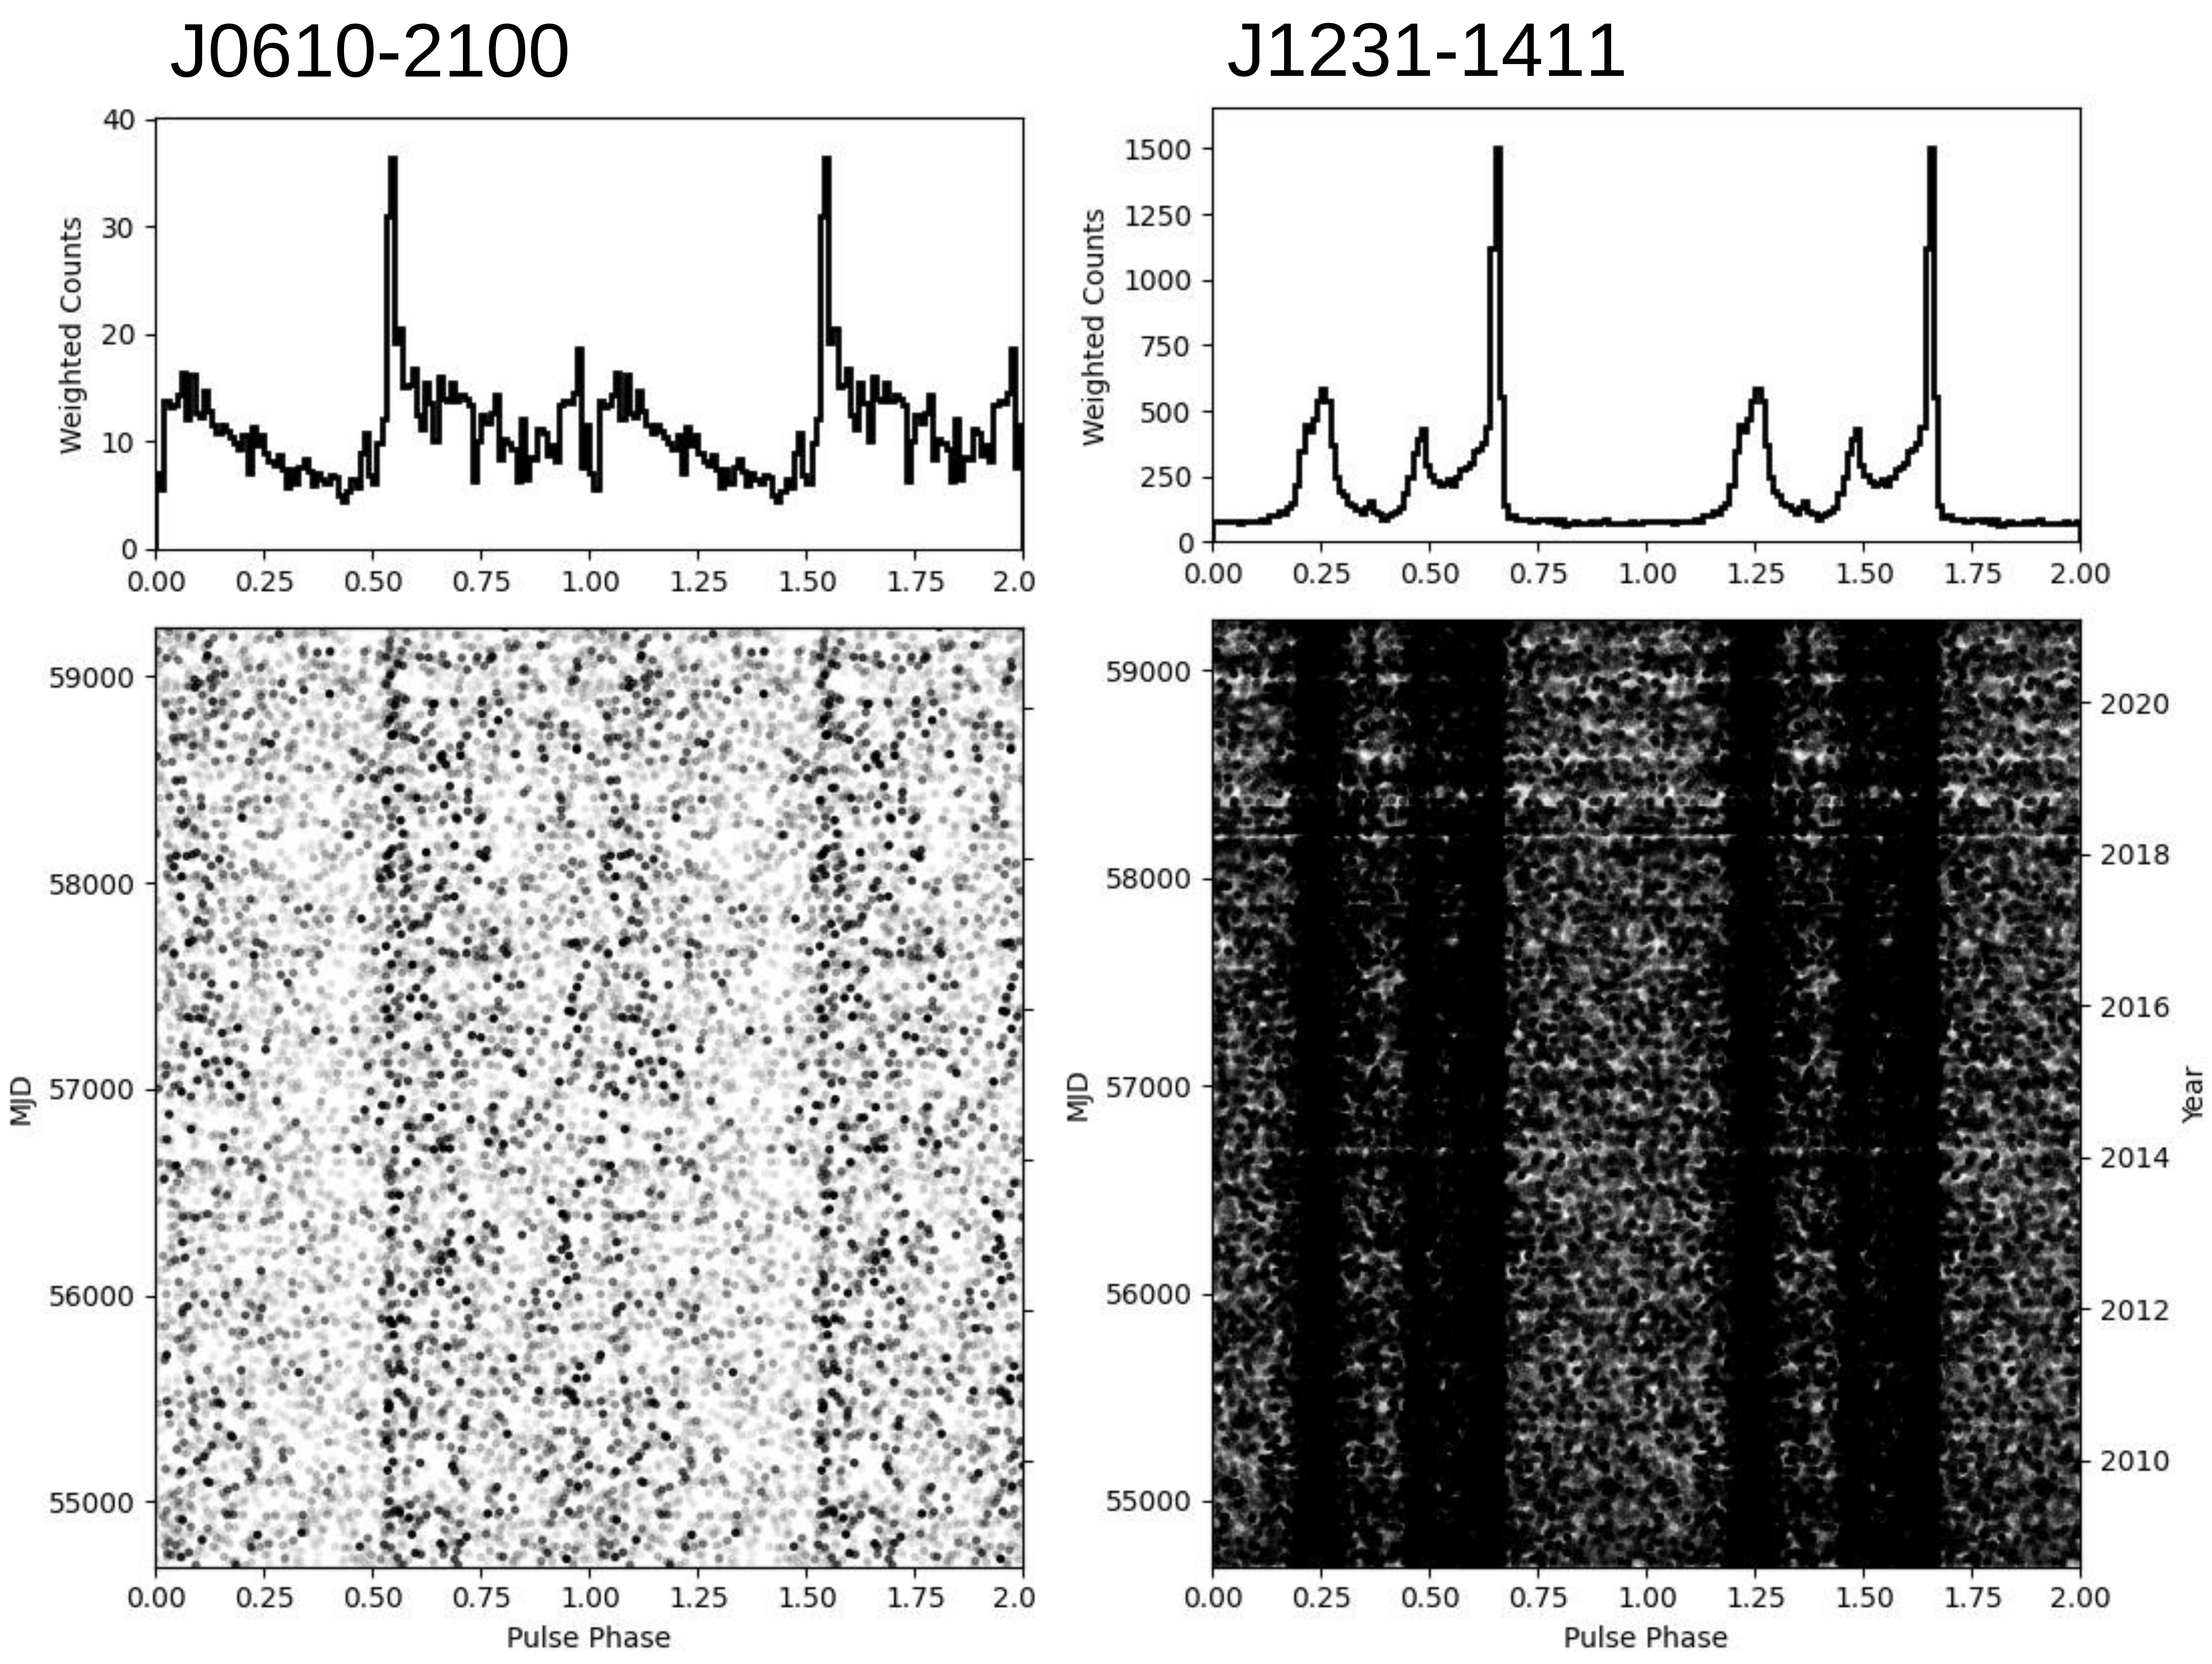

High-energy Pulsar timing Pros and Cons

- No ISM effects
  - No dispersion, no scattering, no ESEs, no phase wander, no nothin'.
- Simple noise model
  - No jitter, no EFAC, no EQUAD, no ECORR.
- Simple (and stable?) pulse profiles.
  - Some slight energy dependence, but not correlated with ISM!
  - Pulse profile changes – open question! But so far not observed in MSPs.
- No calibration / polarization issues.
  - LAT is not sensitive to polarization; response essentially unchanged.
- No JUMPS.
  - LAT configuration very stable, clock continuously running. No JUMPs at all.
- Wide field of view gives continual monitoring.
- Archival data give a >10 year pulsar timing dataset any time a new MSP is discovered.

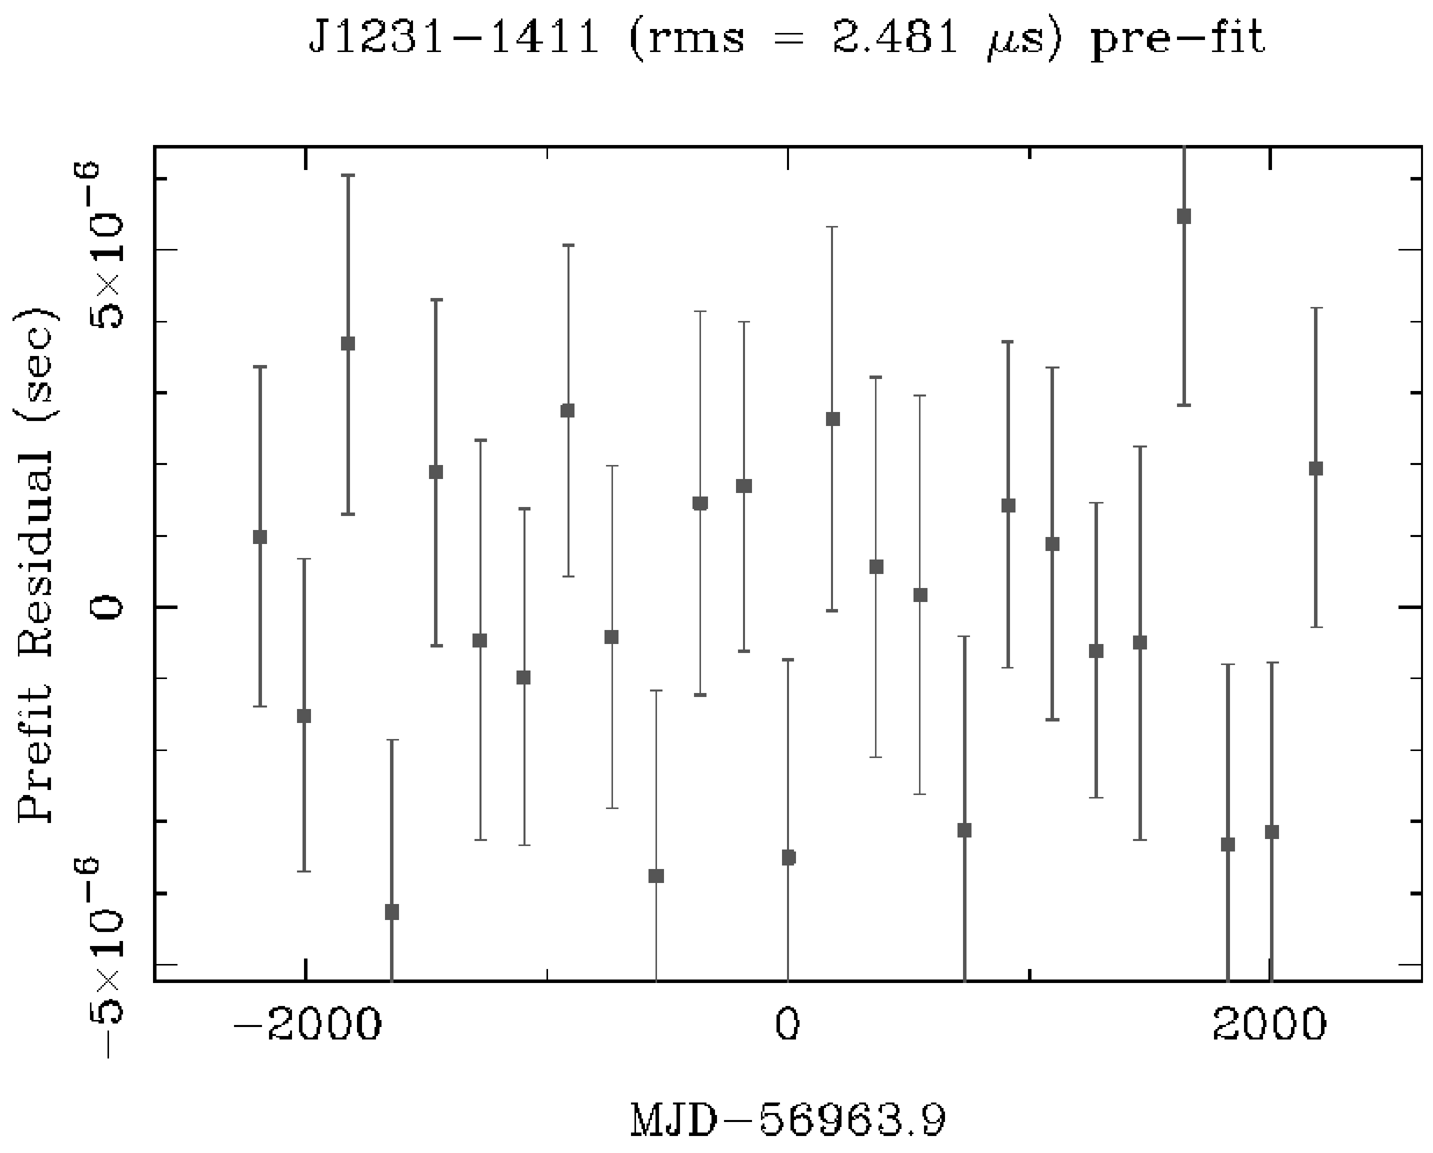

- The LAT effective area is only ~1m^2.
  - The *very best* MSP produces the equivalent of a 1.7 mus TOA each year.
- Unbinned analysis makes some procedures complicated.

MSP Sample and Method

- Require white noise level equivalent to TOA precision (2/yr cadence) <= 22 mus.
  - Test multiple cadences, 2/yr, 0.667/yr, 1/yr, to determine good log likelihood threshold.
    - Results surprisingly robust; consistent results with only 12 TOAs!
  - Require logL > 400 for ~12.5 year dataset.
- Run single pulsar limits with Temponest, enterprise, and unbinned method.
  - Also test intrinsic timing noise models.
- Run common mode with unbinned and enterprise.

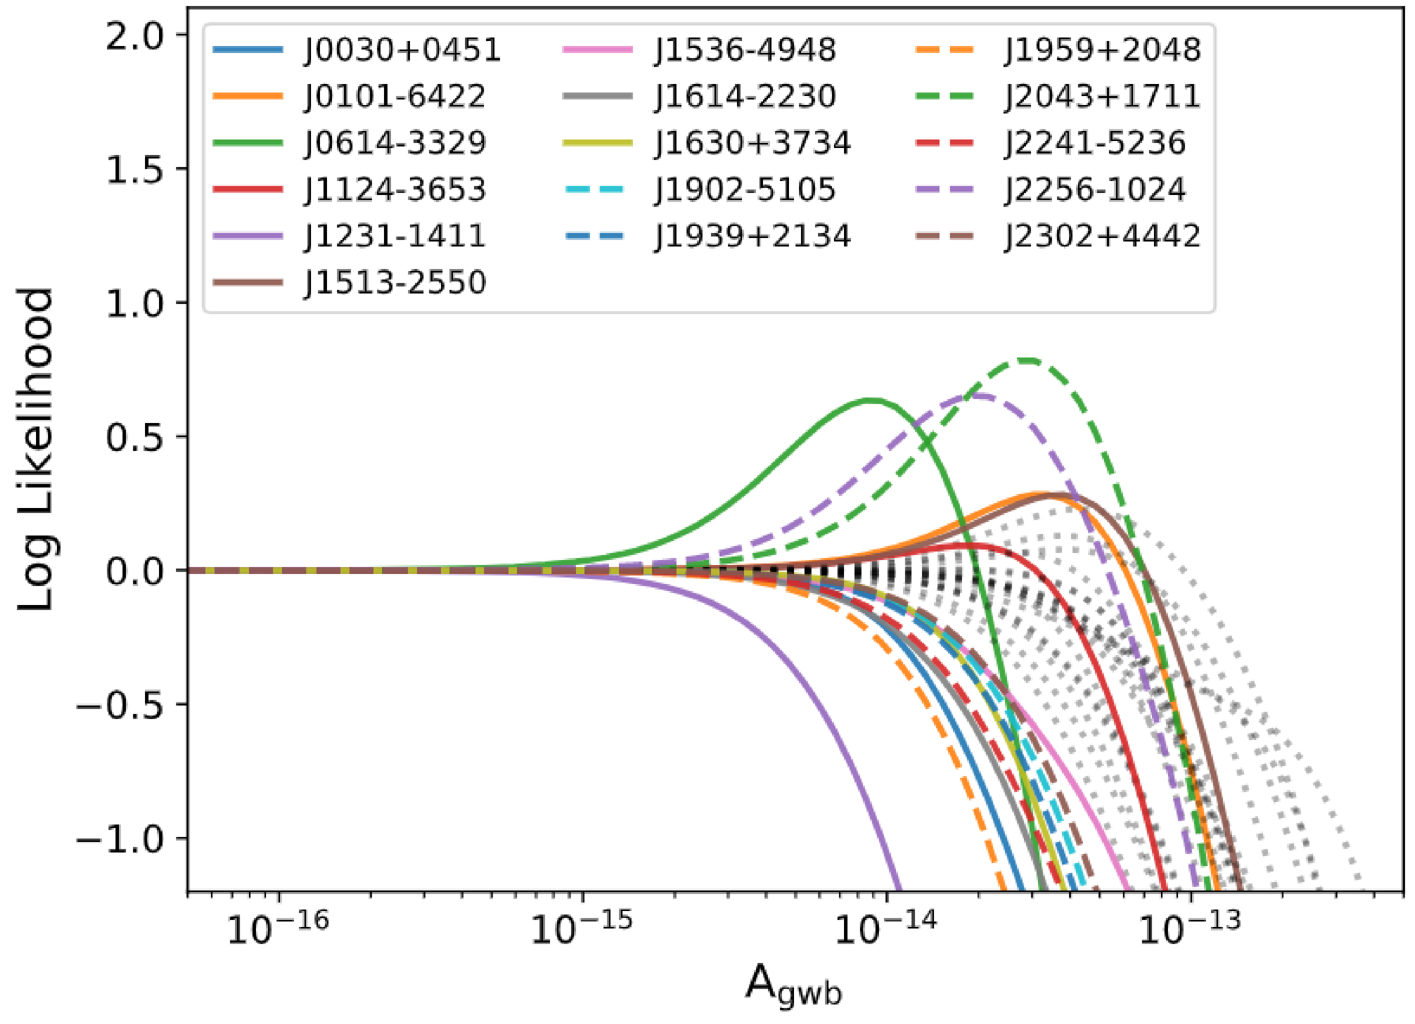

Preliminary results: all methods broadly similar and deliver comparable limits!

- Both single pulsar and common mode.
- This is comforting because unbinned method retains sensitivity to "fast" parameters like position, proper motion, orbital period…
- Outliers on this plot are understood.

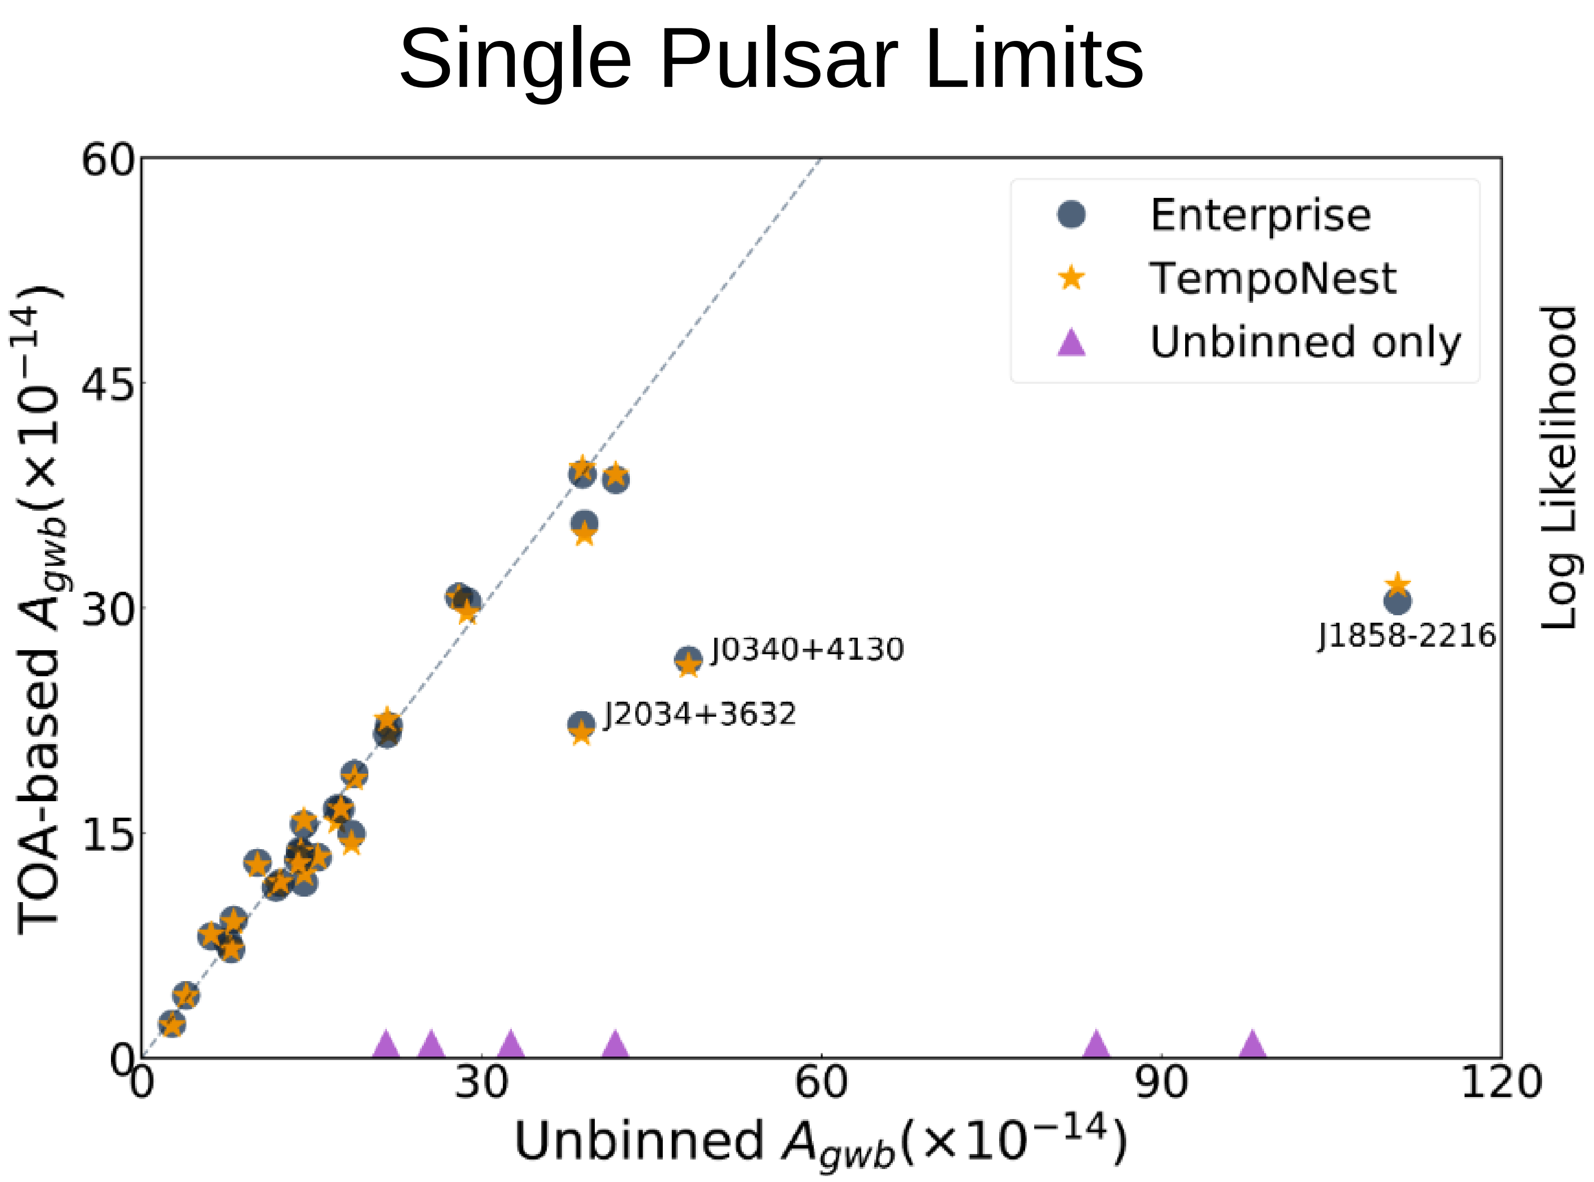

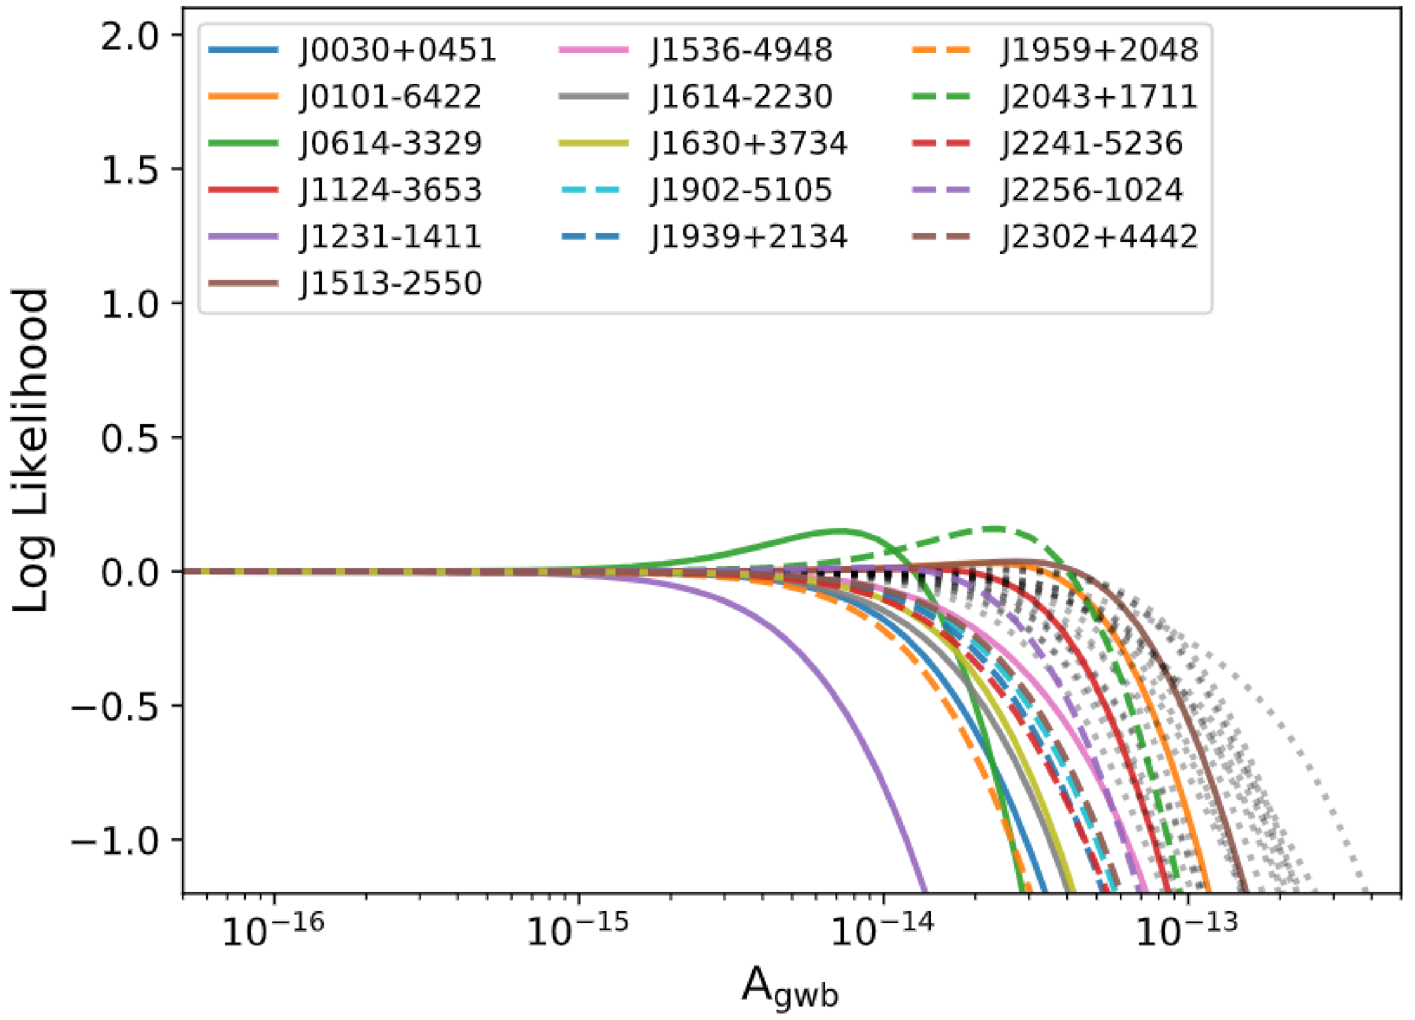

Log likelihood from unbinned method: (top) only GWB (bottom) GWB + intrinsic per-pulsar timing noise (marginalized)

Other GWB models

- There are many other models for GWs detectable by PTAs than merging SMBHs!
- Calculate the limit for a range of (strain) spectral indices.
  - Recall SMBH → -2/3
- Some relic GWs from inflation predict an even steeper spectrum.
- Cosmic strings are very model dependent, but in some models would have a yet steeper spectrum.
- Some phase transitions could produce GWs with frequencies ~1/30 yrs.
  - Fermi's long dataset and stability really help with these! But we have a long way to go, only at ~1/6 yrs now…