Optimal reconstruction of the Hellings and Downs correlation

Bruce Allen
Max Planck Institute for Gravitational Physics
(Albert Einstein Institute, AEI)
Hannover, Germany

- Work done with Joe Romano
- Details arXiv:2407.10968v2
- Related work, mostly with Joe:
  - arXiv:2308.05847, *Answers to frequently asked questions about the pulsar timing array Hellings and Downs curve*, **CQG**
  - arXiv:2208.07230**,** *Hellings and Downs correlation of an arbitrary set of pulsars*, **PRD**
  - arXiv:2205.05637, *Variance of the Hellings-Downs Correlation***, PRD**

Pulsar Timing Arrays (PTA)

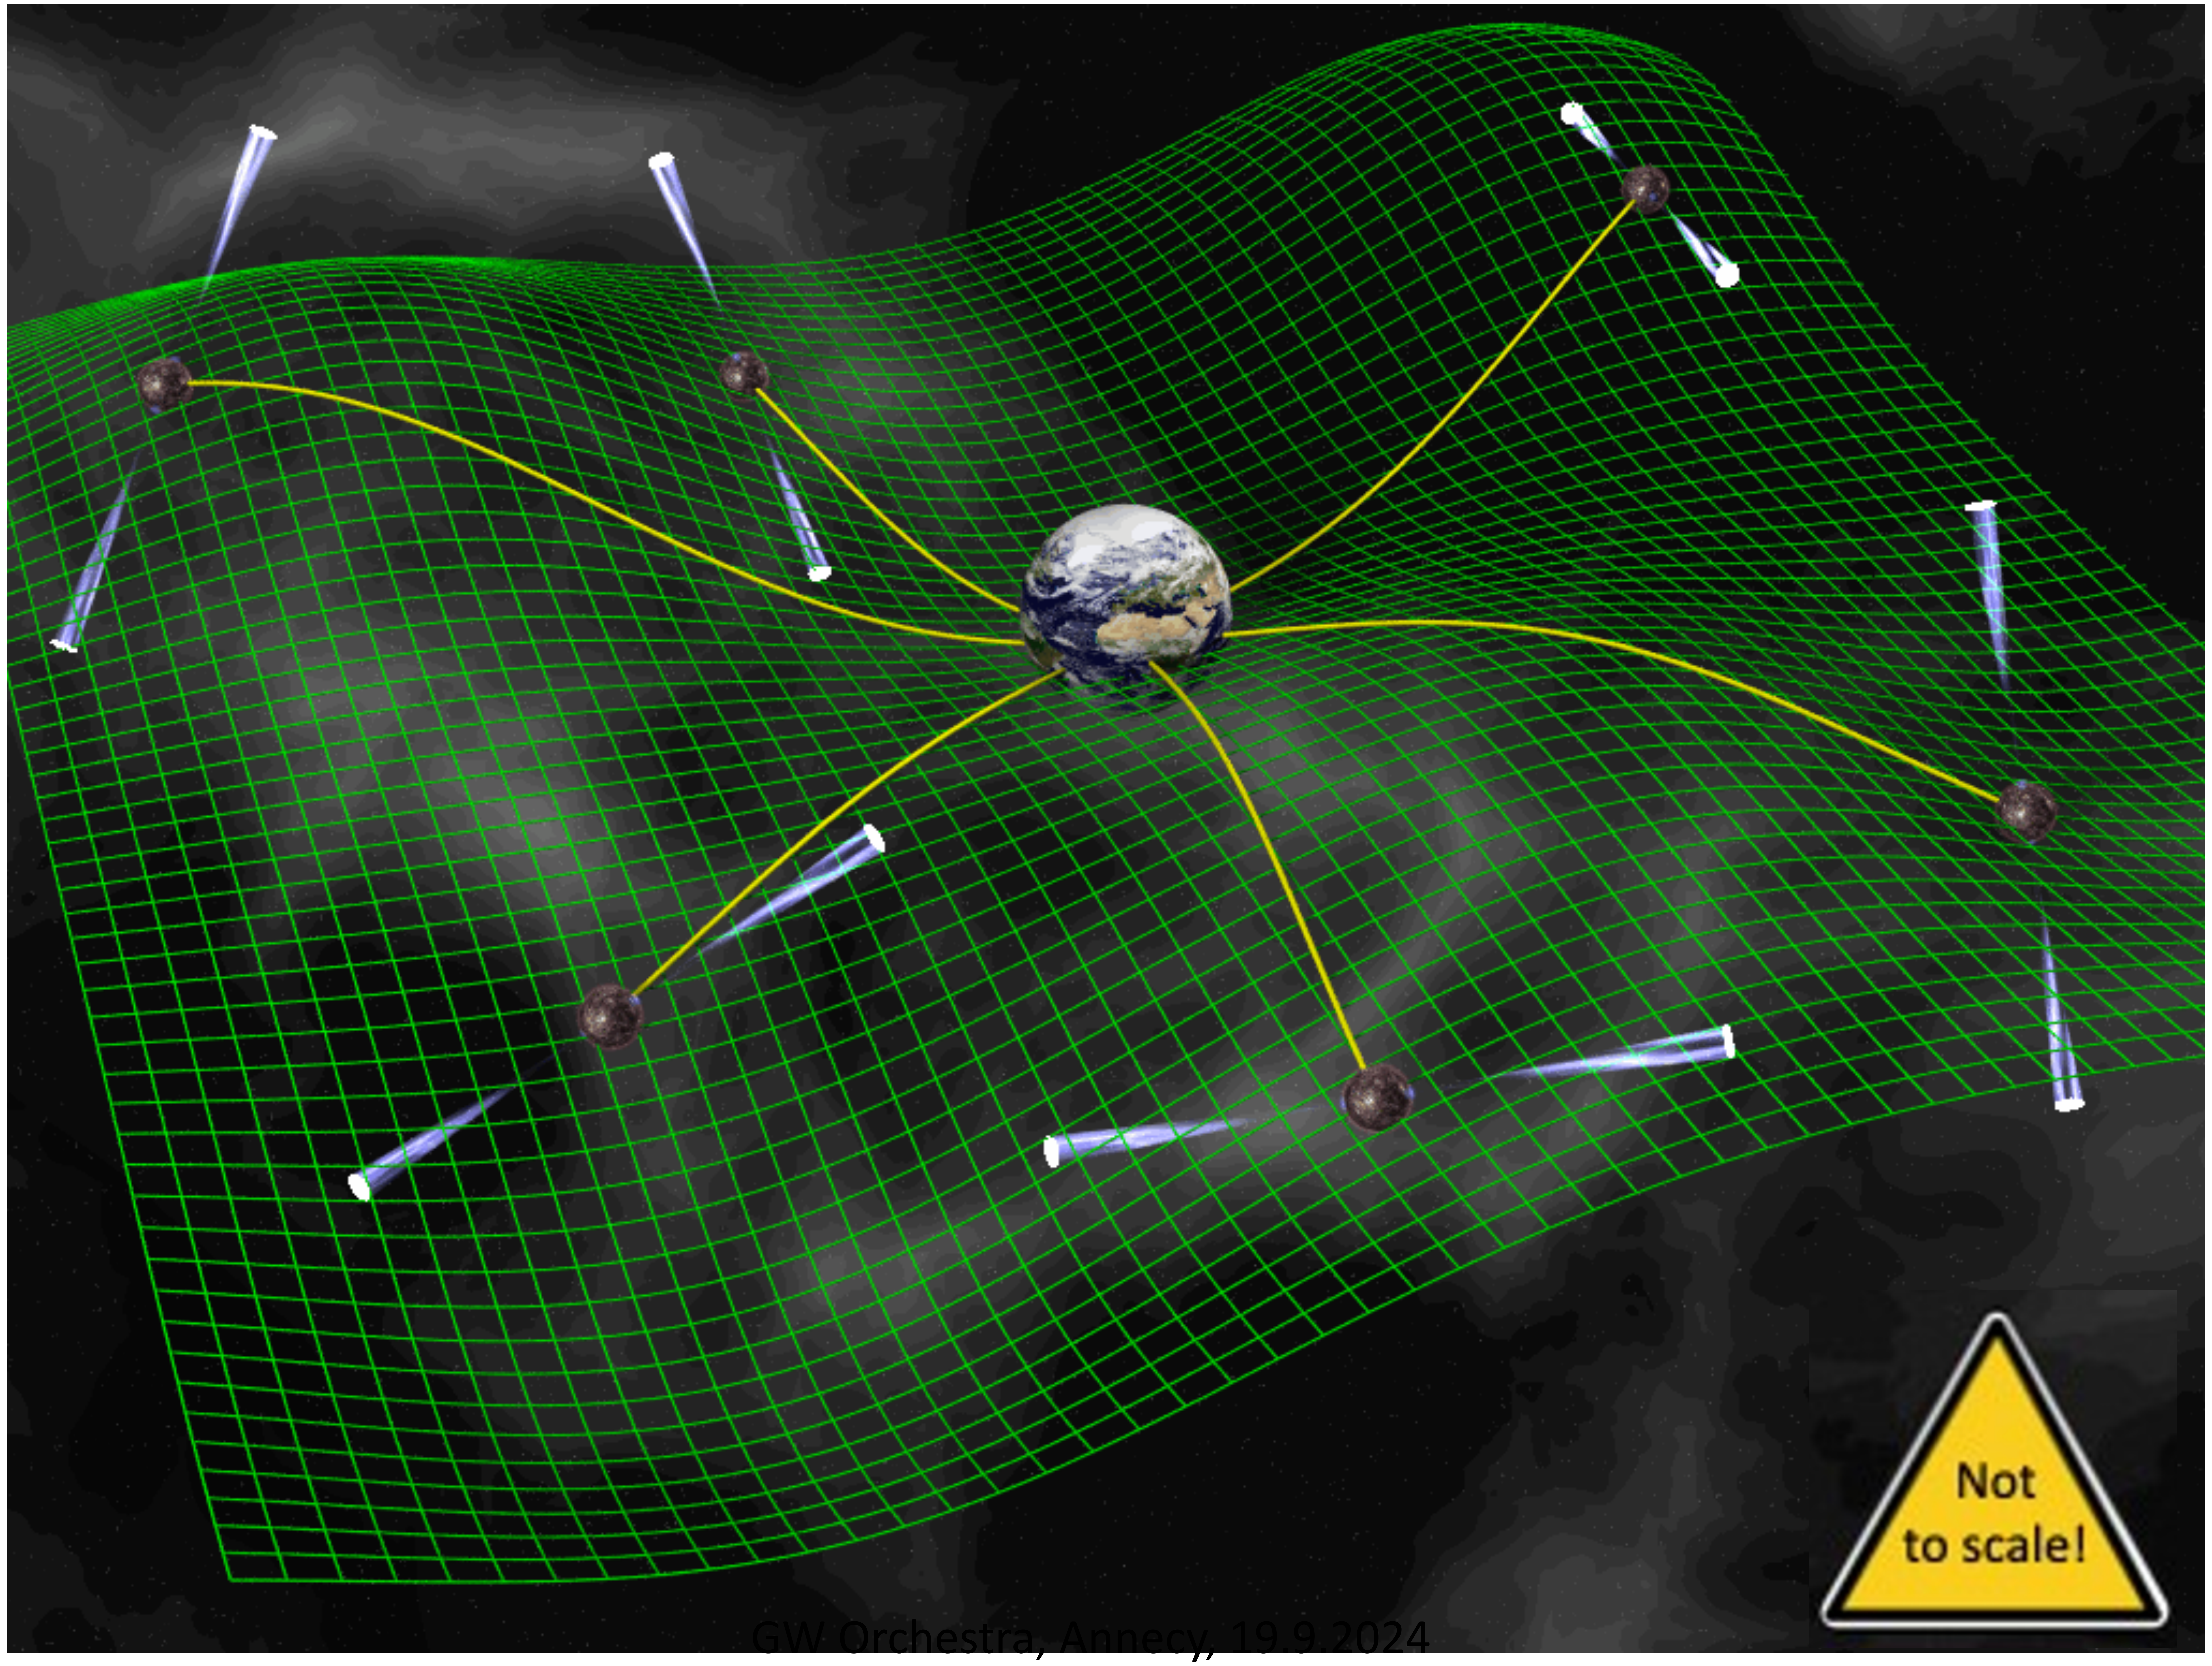

$\gamma \in [0,180°]$ is angle between directions to two pulsars

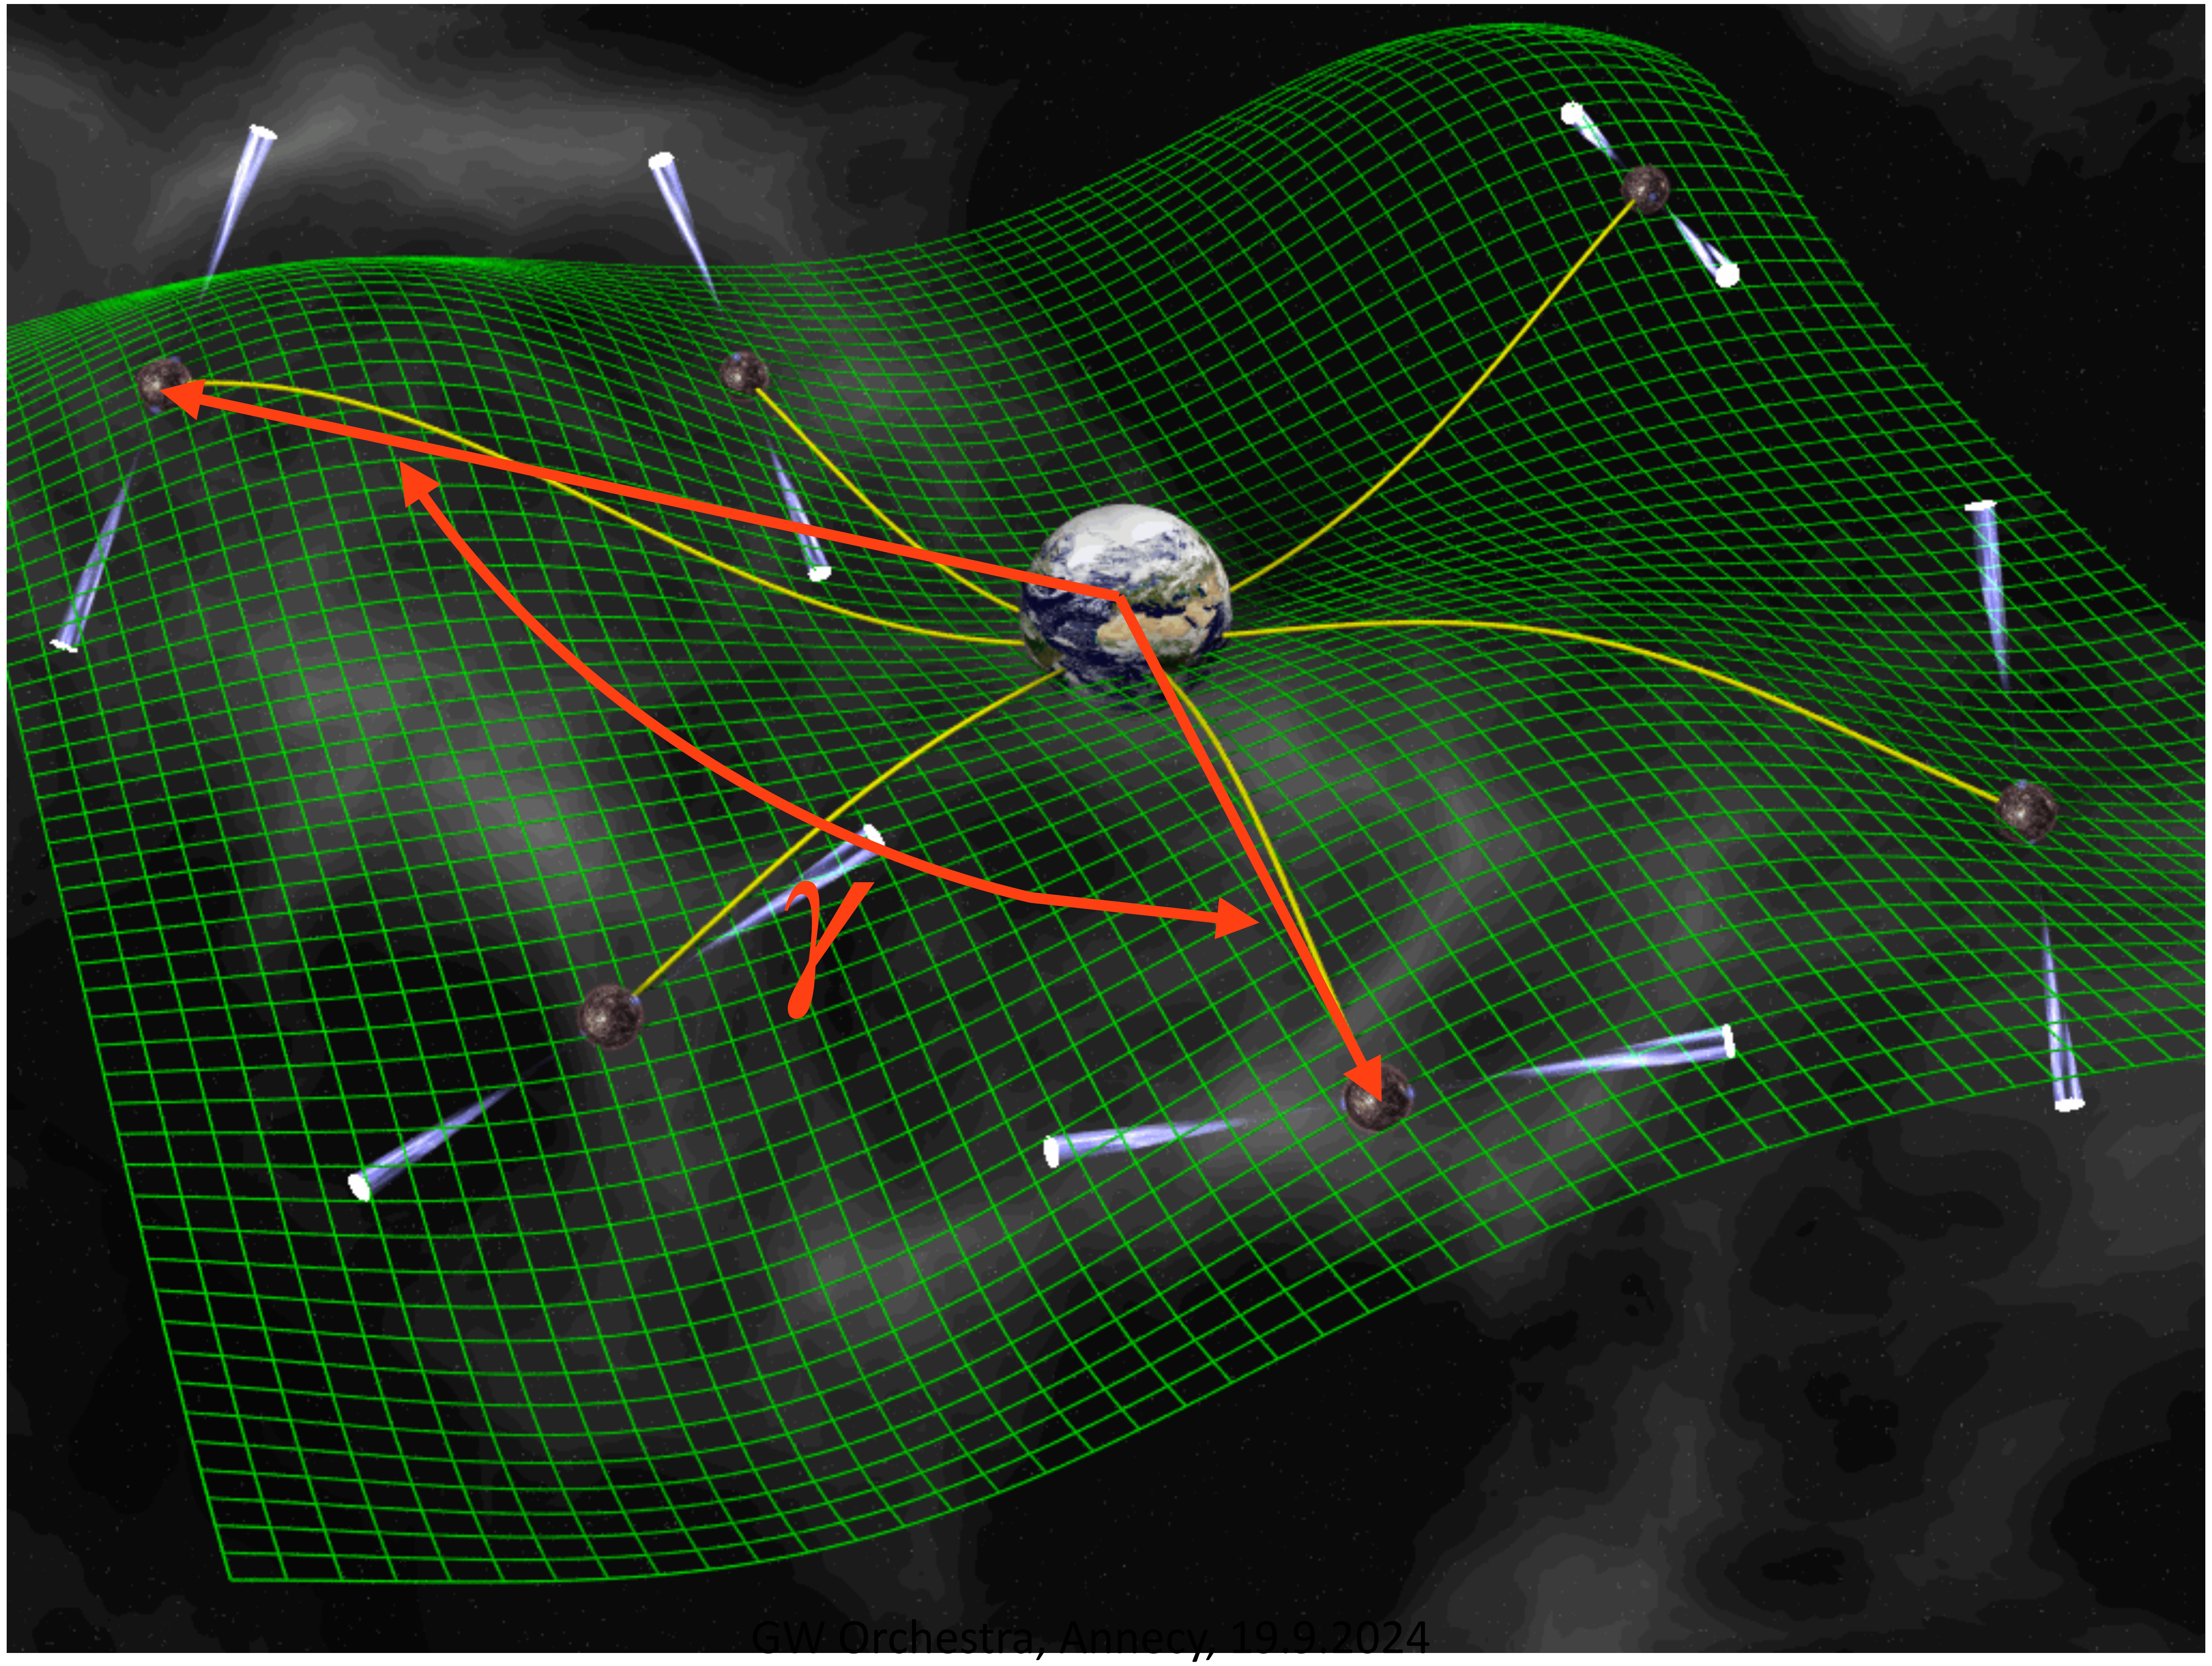

$\gamma \in [0,180°]$ is angle between directions to two pulsars

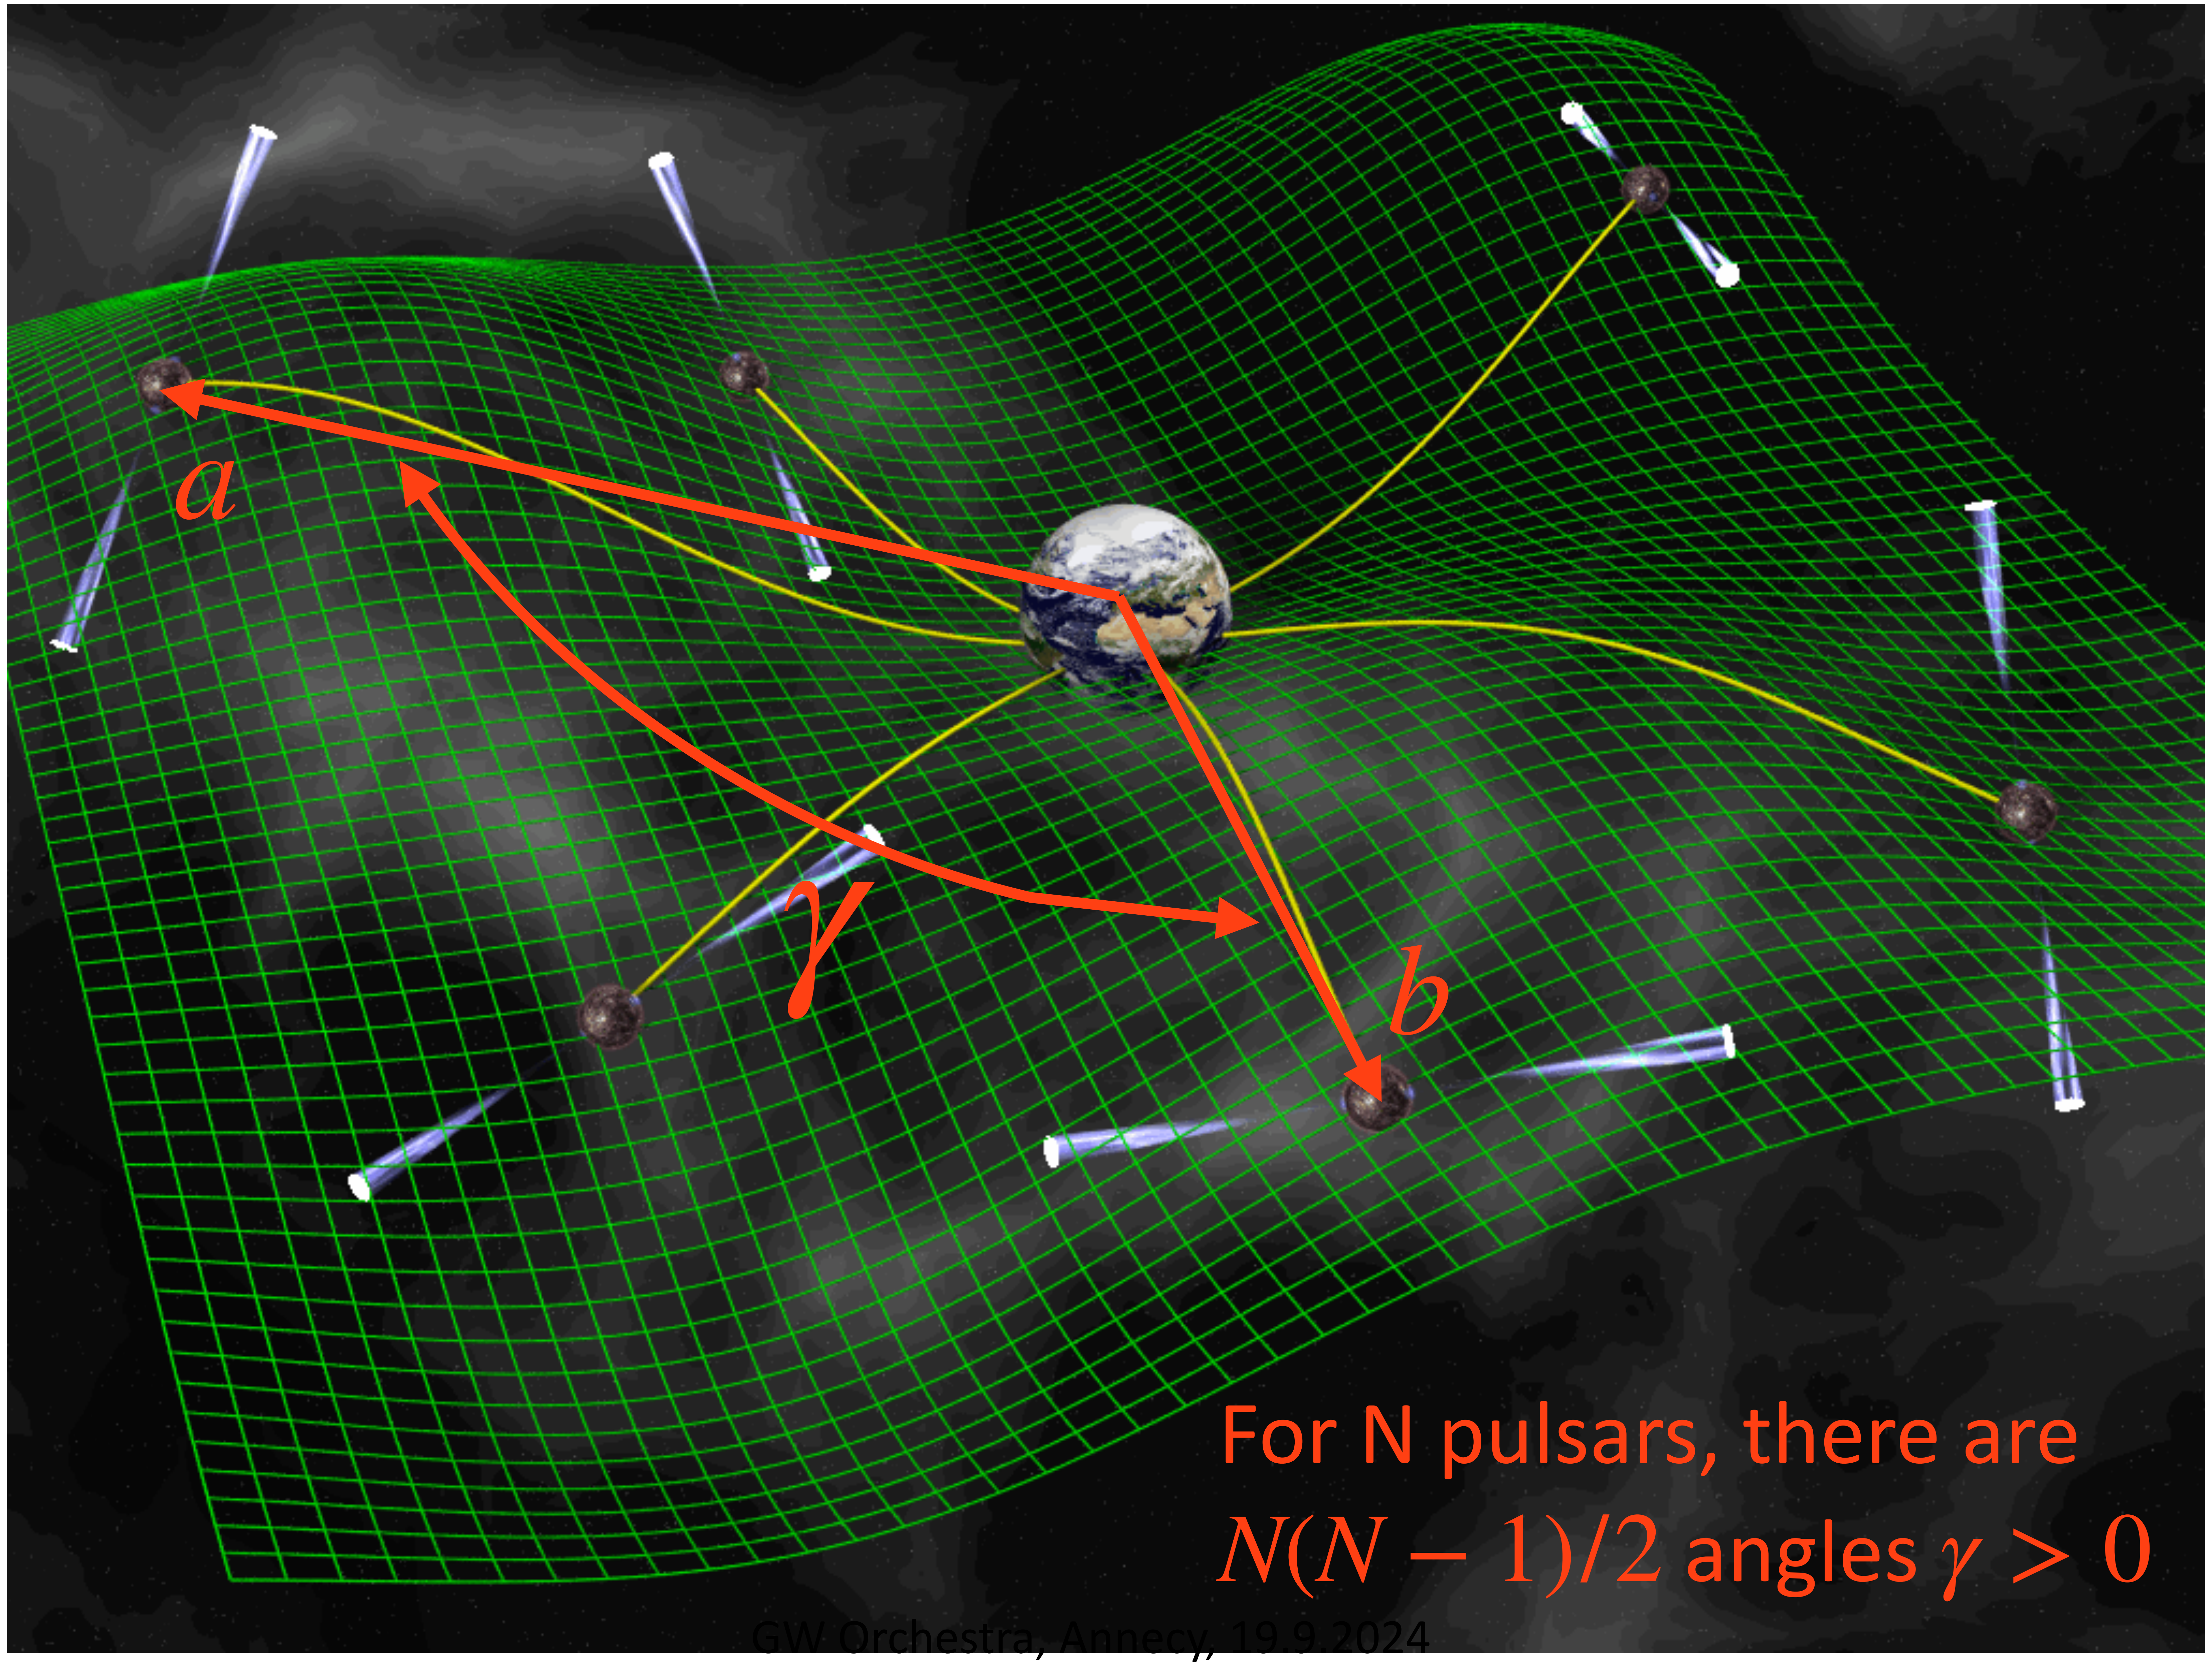

Hellings and Downs curve
For PTAs, like LIGO/Virgo binary "chirp" waveform

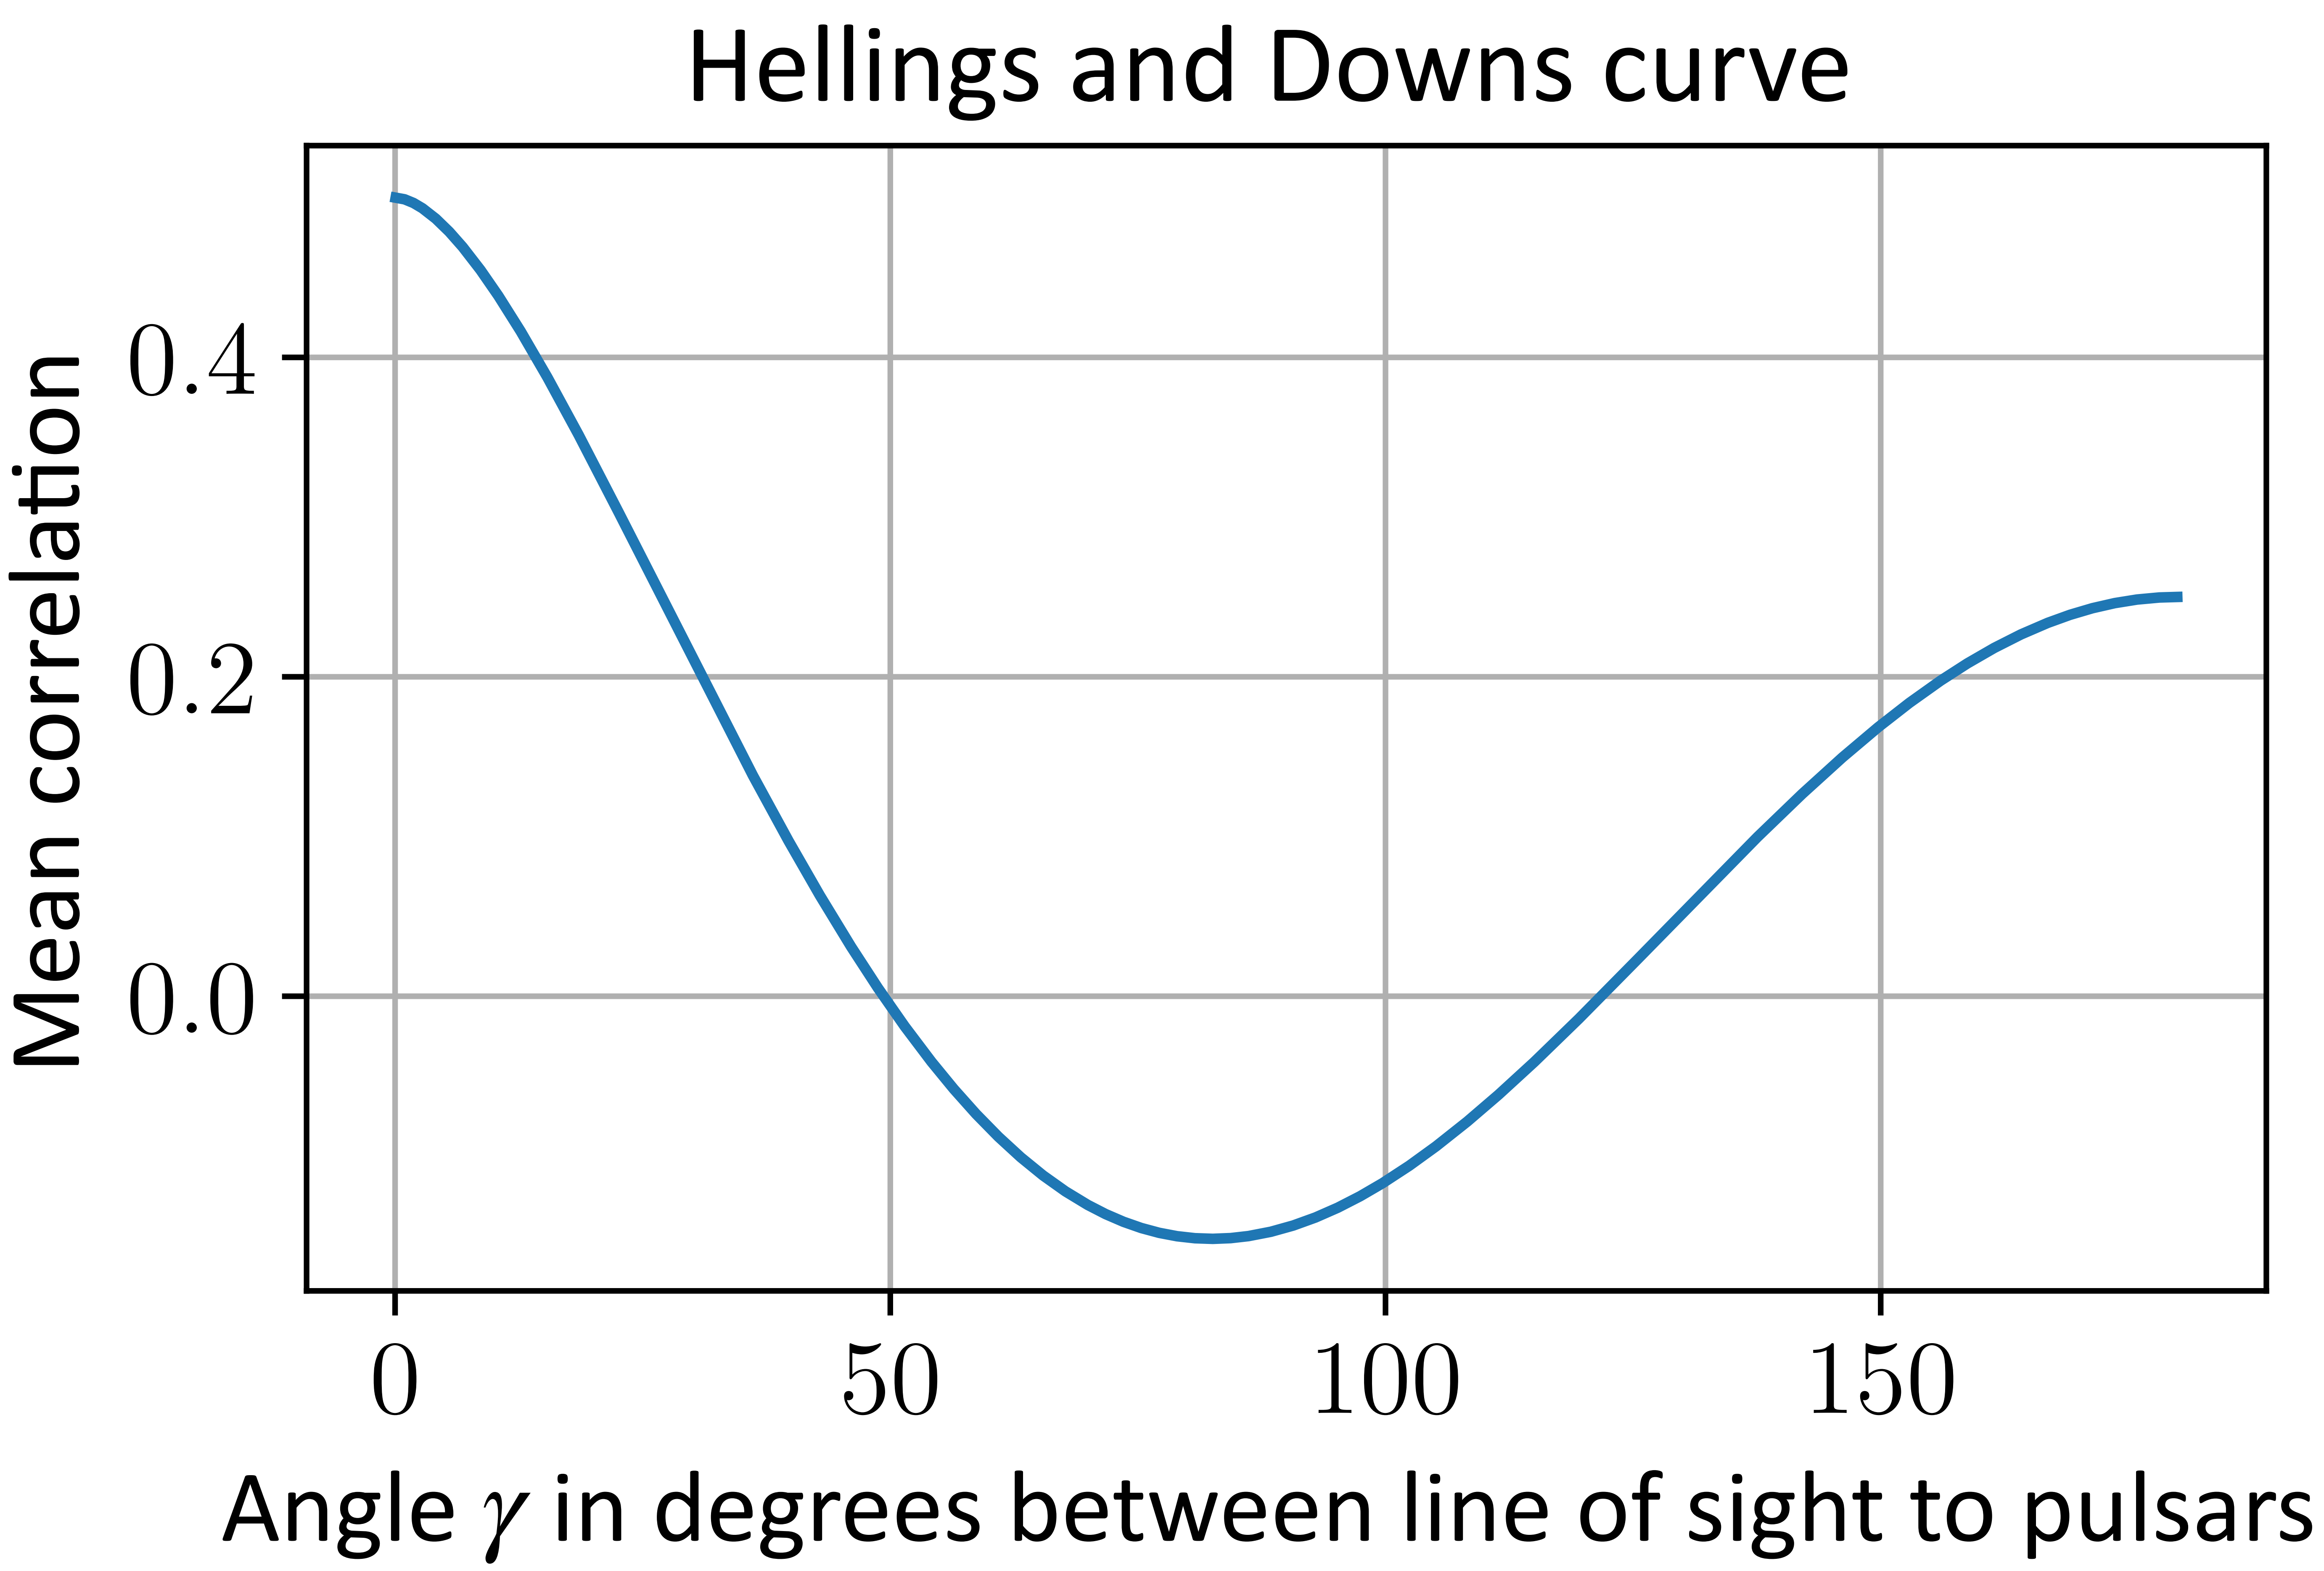

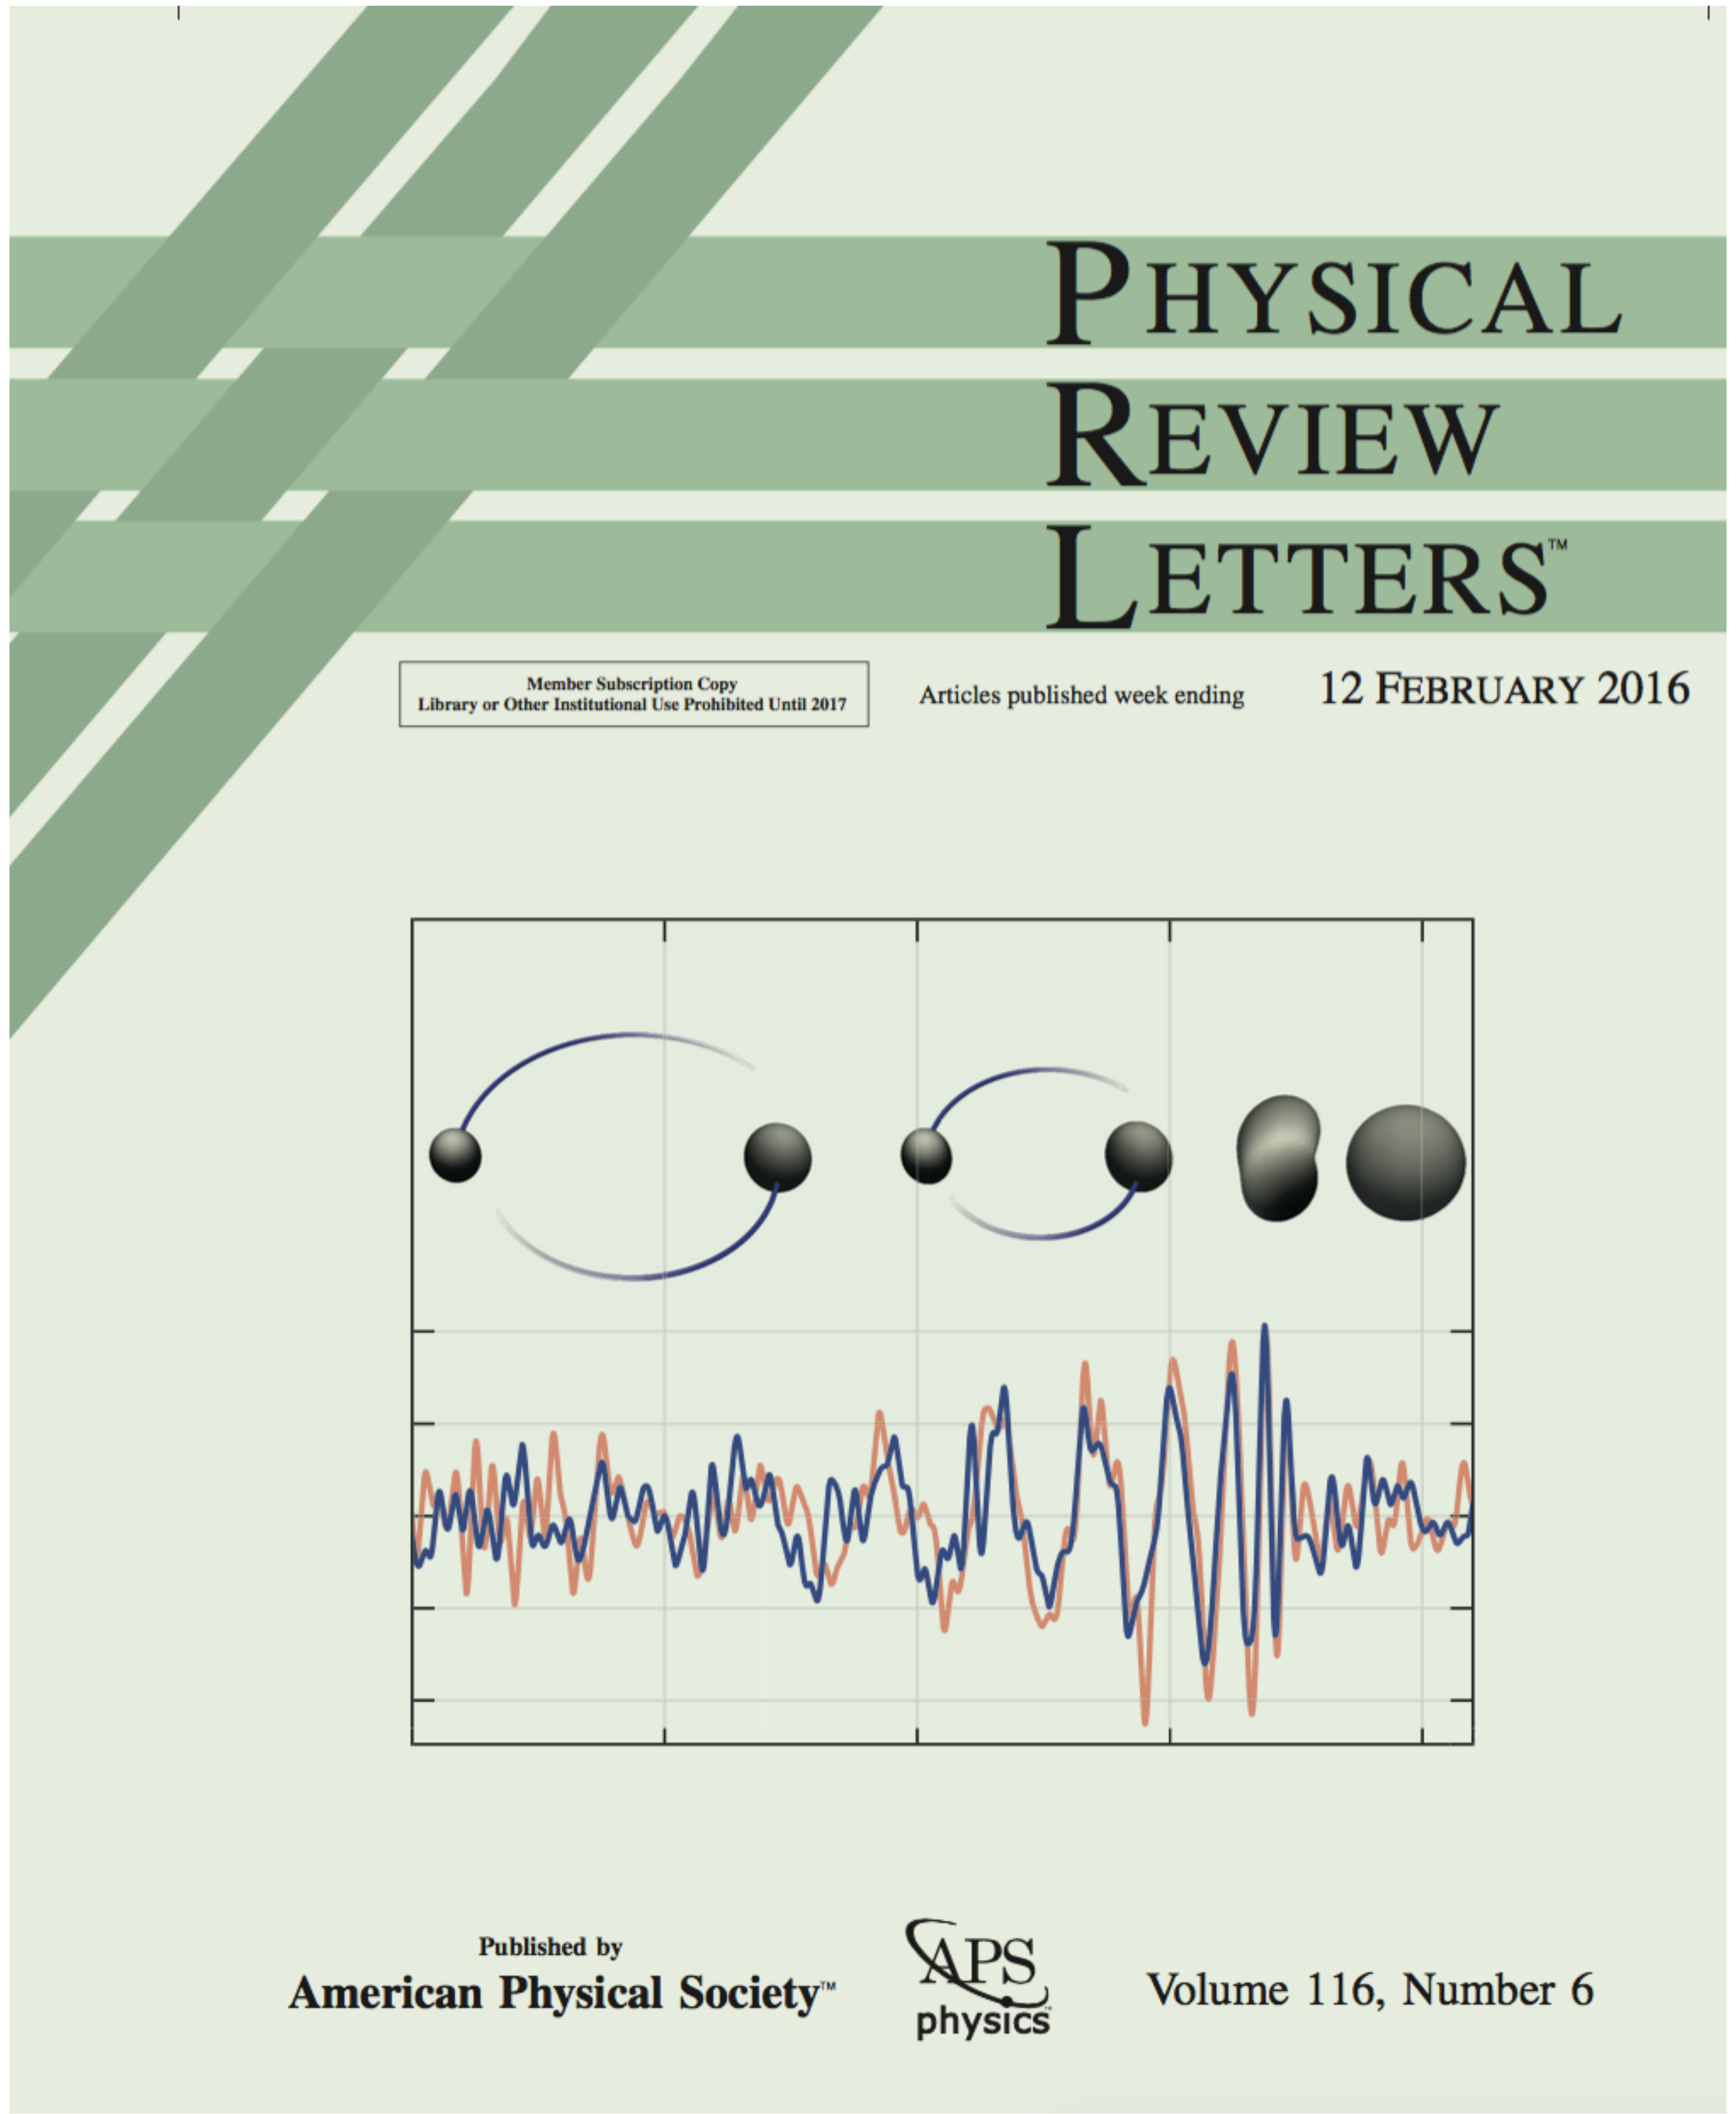

Hellings and Downs curve
For PTAs, like LIGO/Virgo binary “chirp” waveform

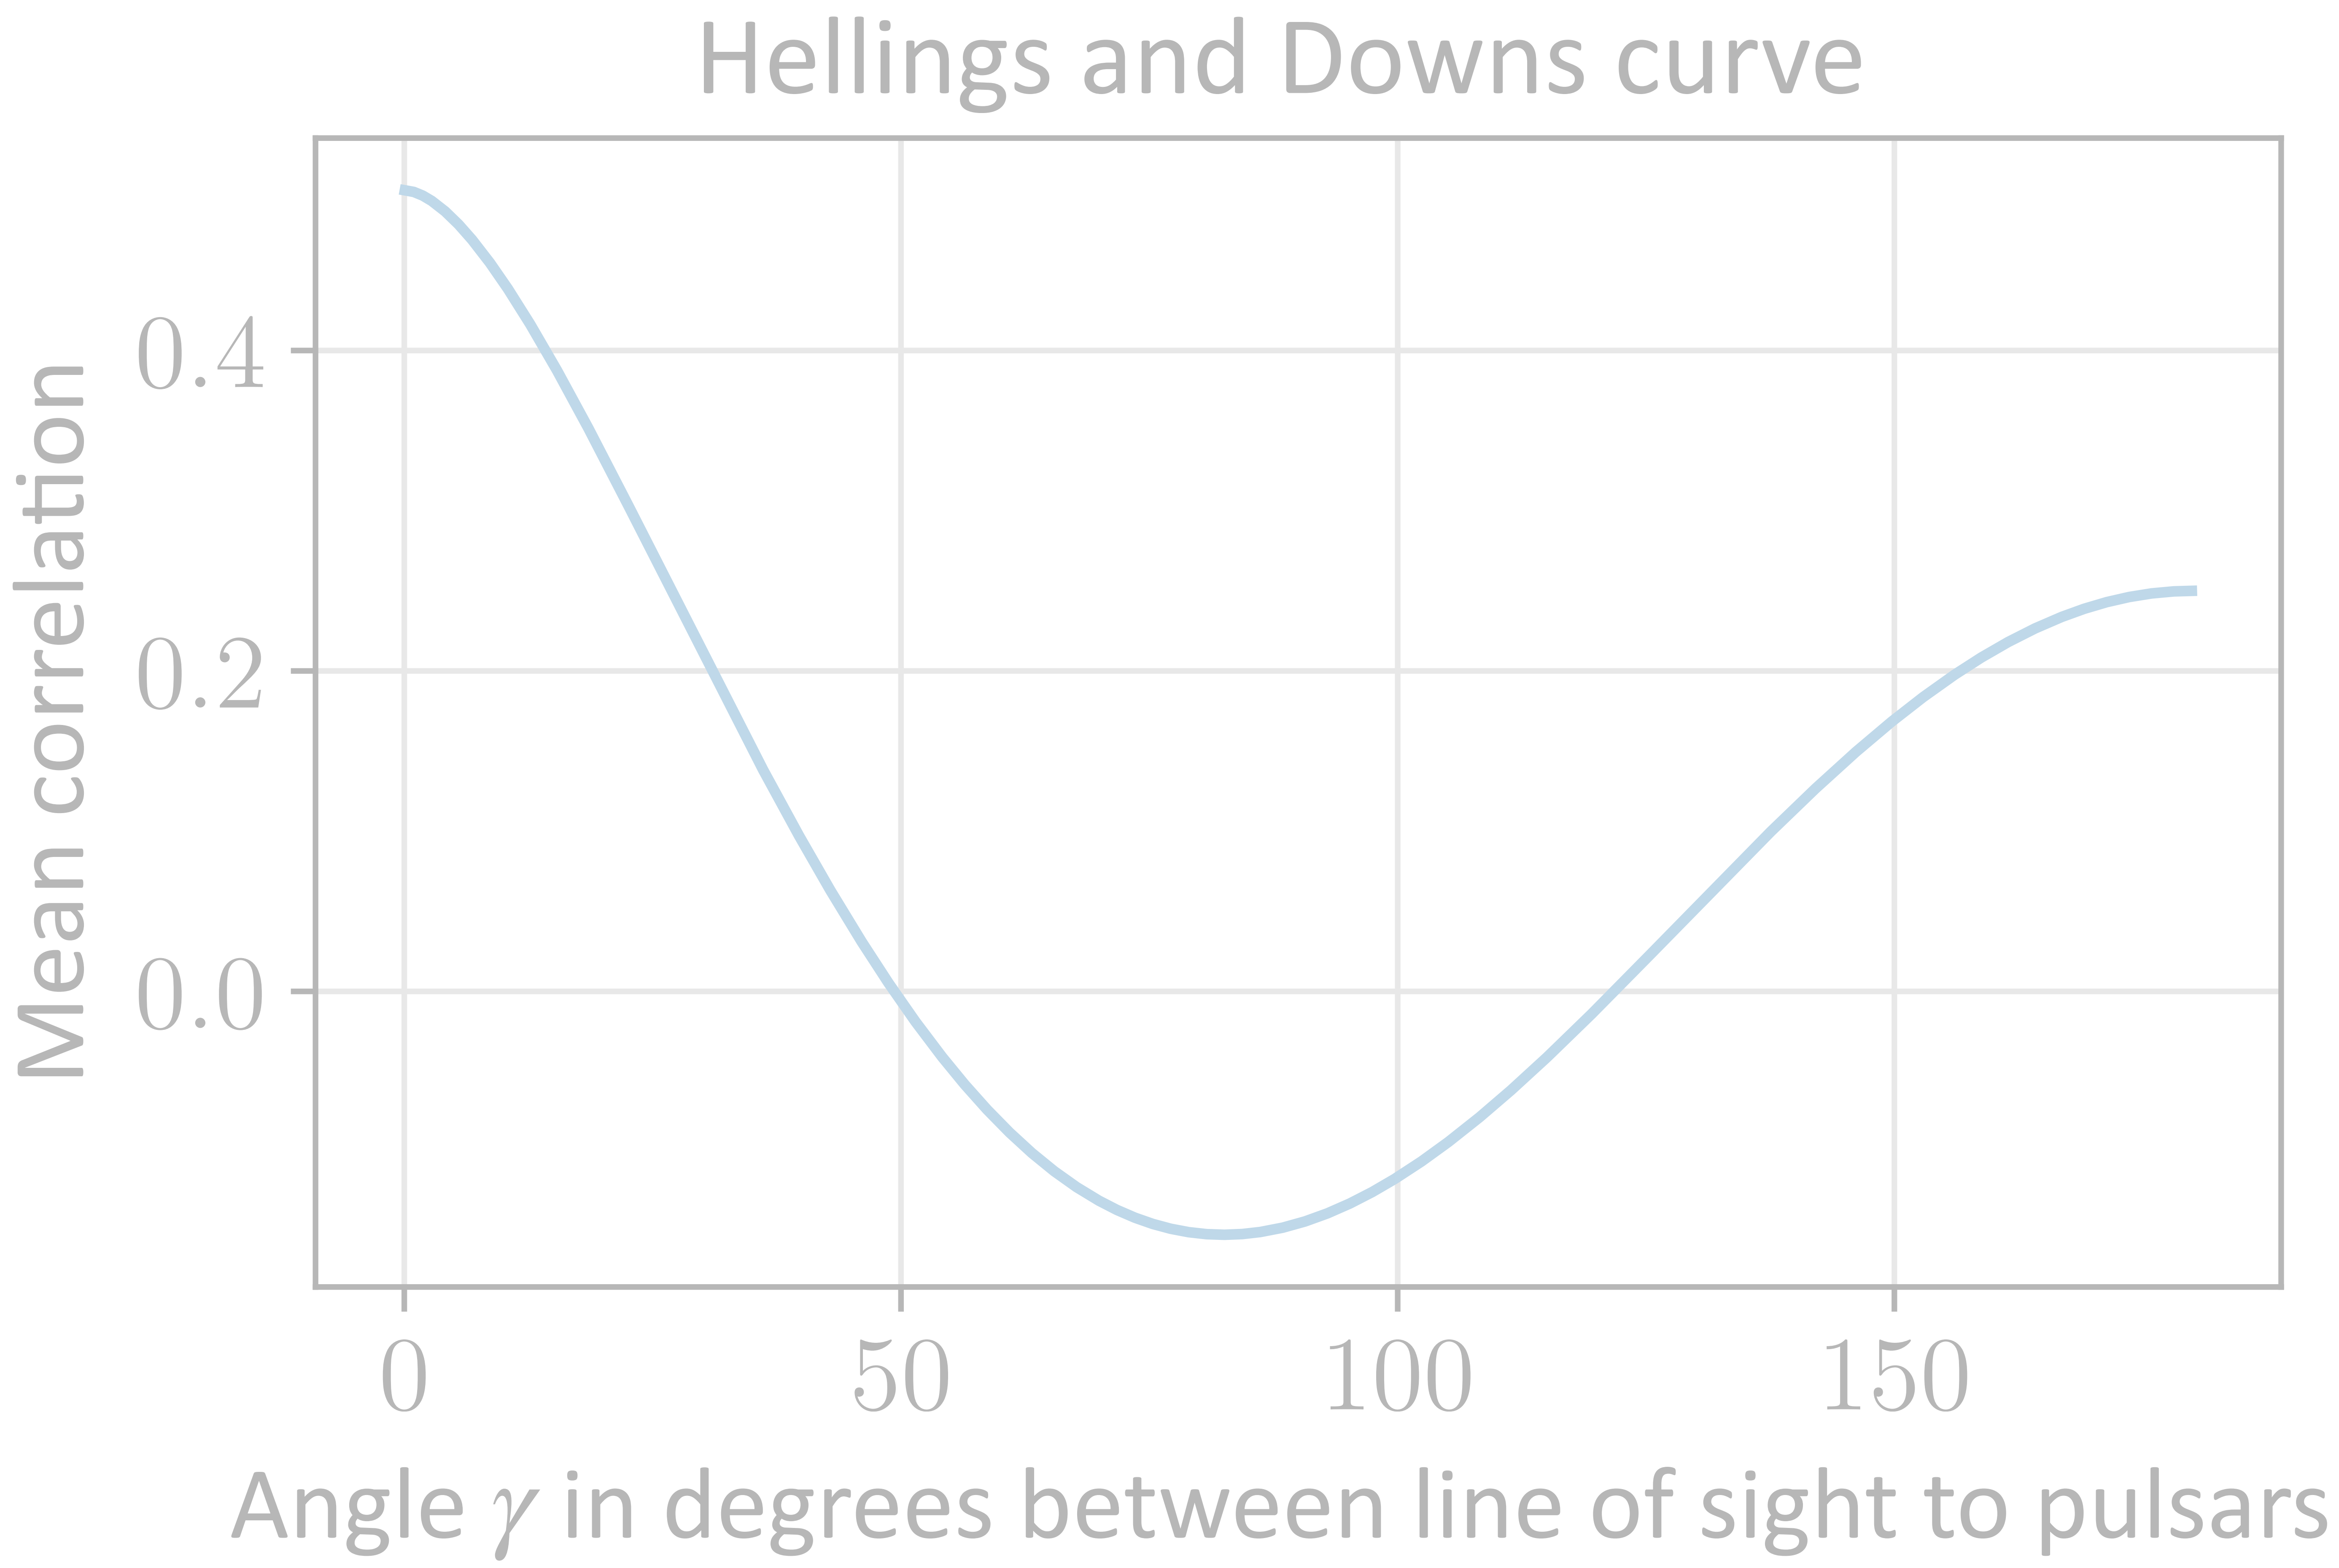

“To demonstrate that gravitational waves are creating some of the noise in the pulsar-timing data sets, observations must exhibit the Hellings and Downs curve.”

The International Pulsar Timing Array checklist for the detection of nanohertz gravitational waves, arXiv:2304.04767

Example: NANOGrav 15-year data

- $N = 67$ pulsars $\Longrightarrow$ 2211 pulsar pairs
- 15 angular bins, so $2211/15 \approx 147$ pulsar pairs per angular bin

NANOGrav reconstruction with N = 67 pulsars

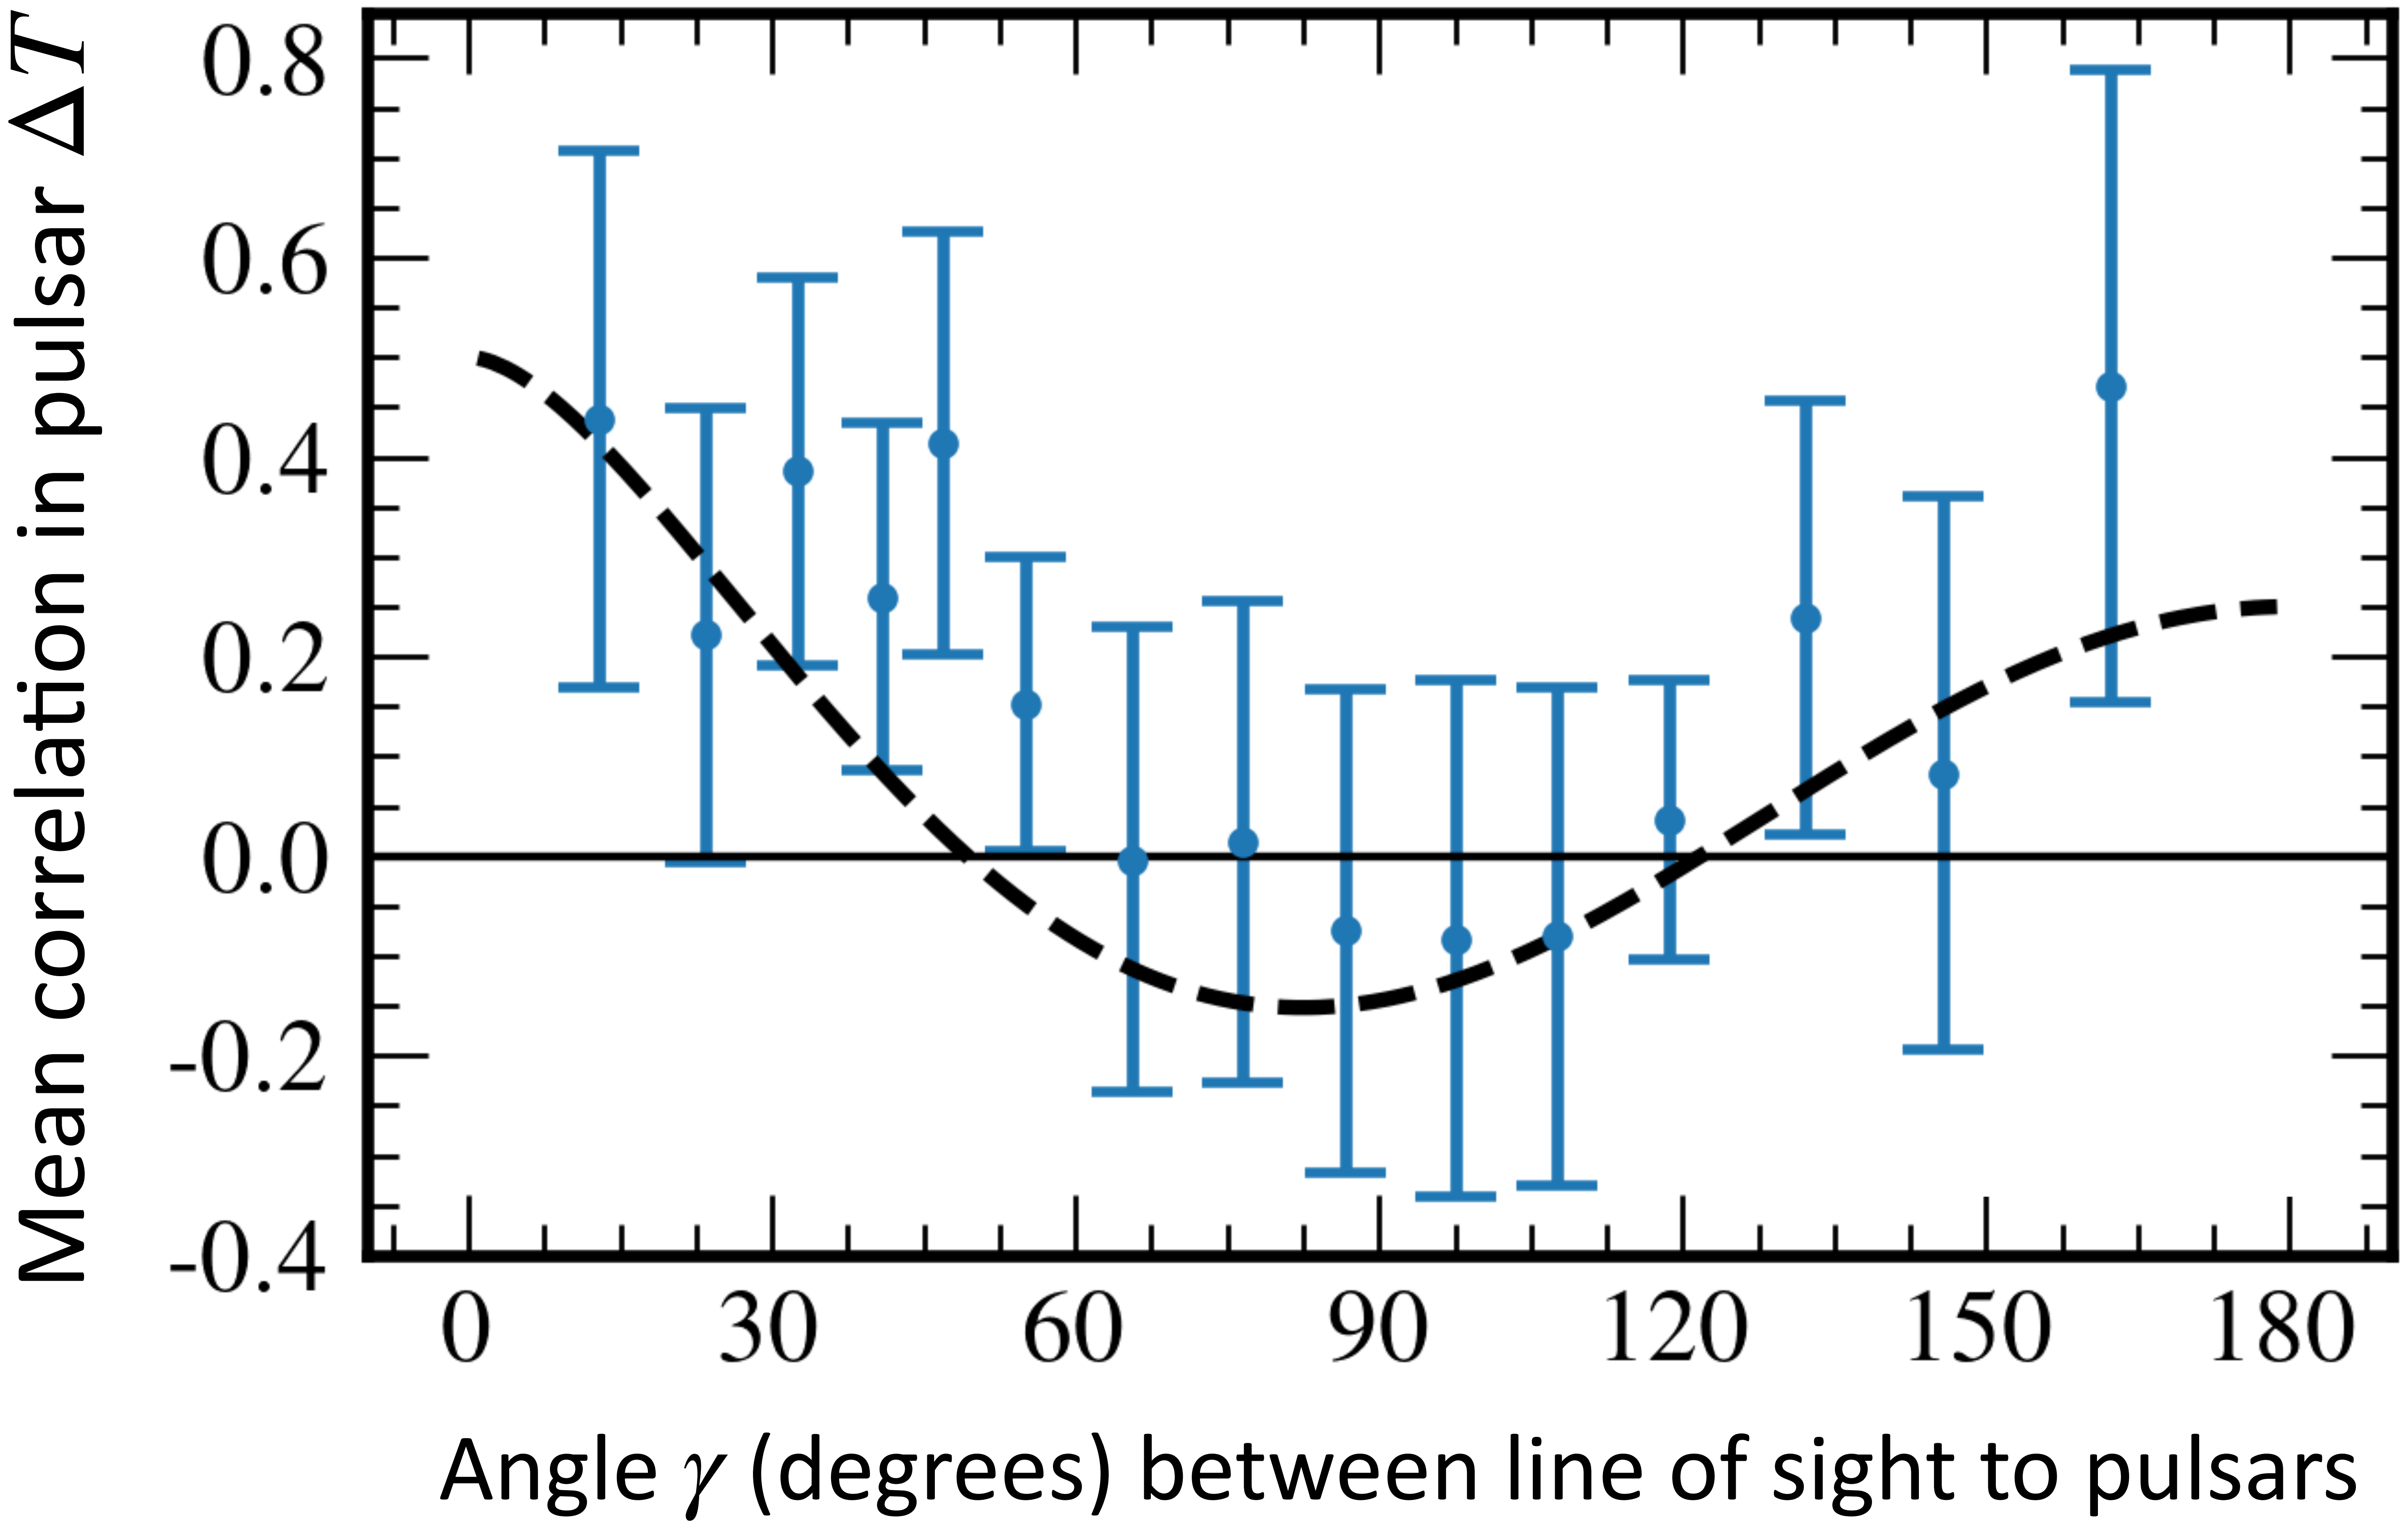

Hellings and Downs curve

- One distant GW point source, sky direction $-\hat{\Omega}$
- Earth at $\vec{0}$, pulsar a at position $L_a \hat{p}_a$
- GWs produce time-dependent redshift/blueshift:
  $$Z_a(t) = \left[h_{\mu\nu}\left(t, \vec{0}\right) - h_{\mu\nu}\left(t - L_a, L_a \hat{p}_a\right)\right] F_a^{\mu\nu}(\hat{\Omega})$$
- Pulsar a antenna pattern:
  $$F_a^{\mu\nu}(\hat{\Omega}) = \frac{1}{2} \frac{\hat{p}_a^\mu \hat{p}_a^\nu}{(1 + \hat{\Omega} \cdot \hat{p}_a)}$$
- Pulsar a response to circular polarization:
  $$F_a(\hat{\Omega}) = F_a^{\mu\nu}(\hat{\Omega}) \left[e_{\mu\nu}^{+}(\hat{\Omega}) + i\, e_{\mu\nu}^{\times}(\hat{\Omega})\right]$$
- Hellings and Downs curve:
  $$\mu_{\mathrm{u}}(\gamma) = \left\langle F_a(\hat{\Omega}) F_b^*(\hat{\Omega}) \right\rangle_{ab \in \gamma}$$
- Key idea for detection in noise: *the pulsar redshifts are correlated among different pulsars*

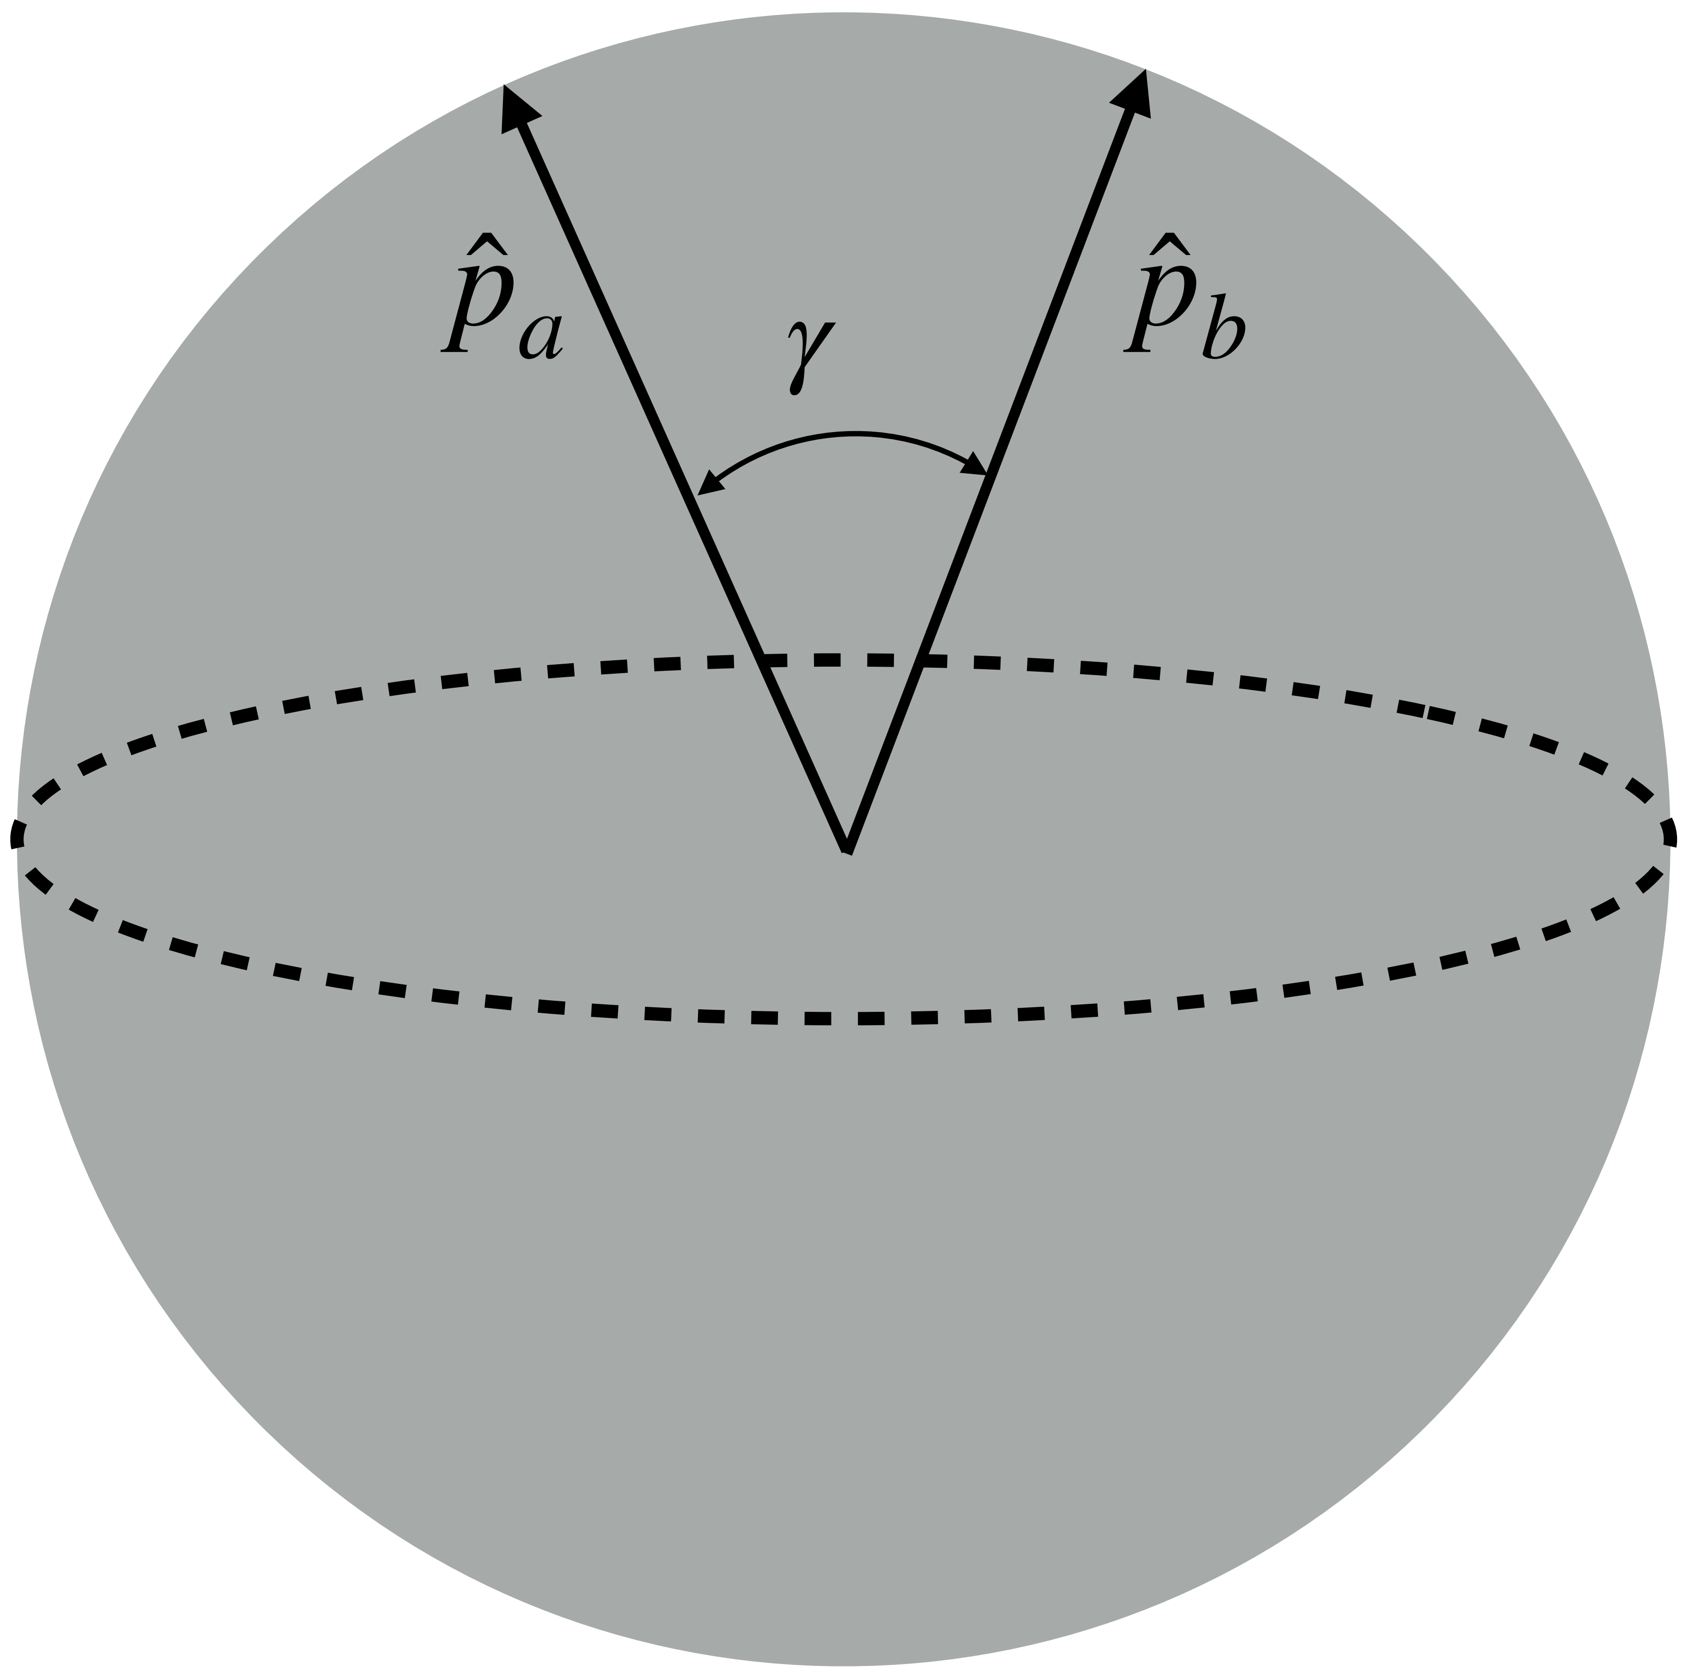

Average over all pulsar directions $\hat{p}_a$ and $\hat{p}_b$ which are uniformly distributed on the sphere and separated by angle γ.

Example: NANOGrav 15-year data

- $N = 67$ pulsars $\Longrightarrow$ 2211 pulsar pairs
- 15 angular bins, so $2211/15 \approx 147$ pulsar pairs per angular bin
- Deviations from the Hellings and Downs curve arise from:
  - *Pulsar/measurement noise*
  - *Finite set of pulsars at particular sky locations*
  - *Cosmic variance*

NANOGrav reconstruction with N = 67 pulsars

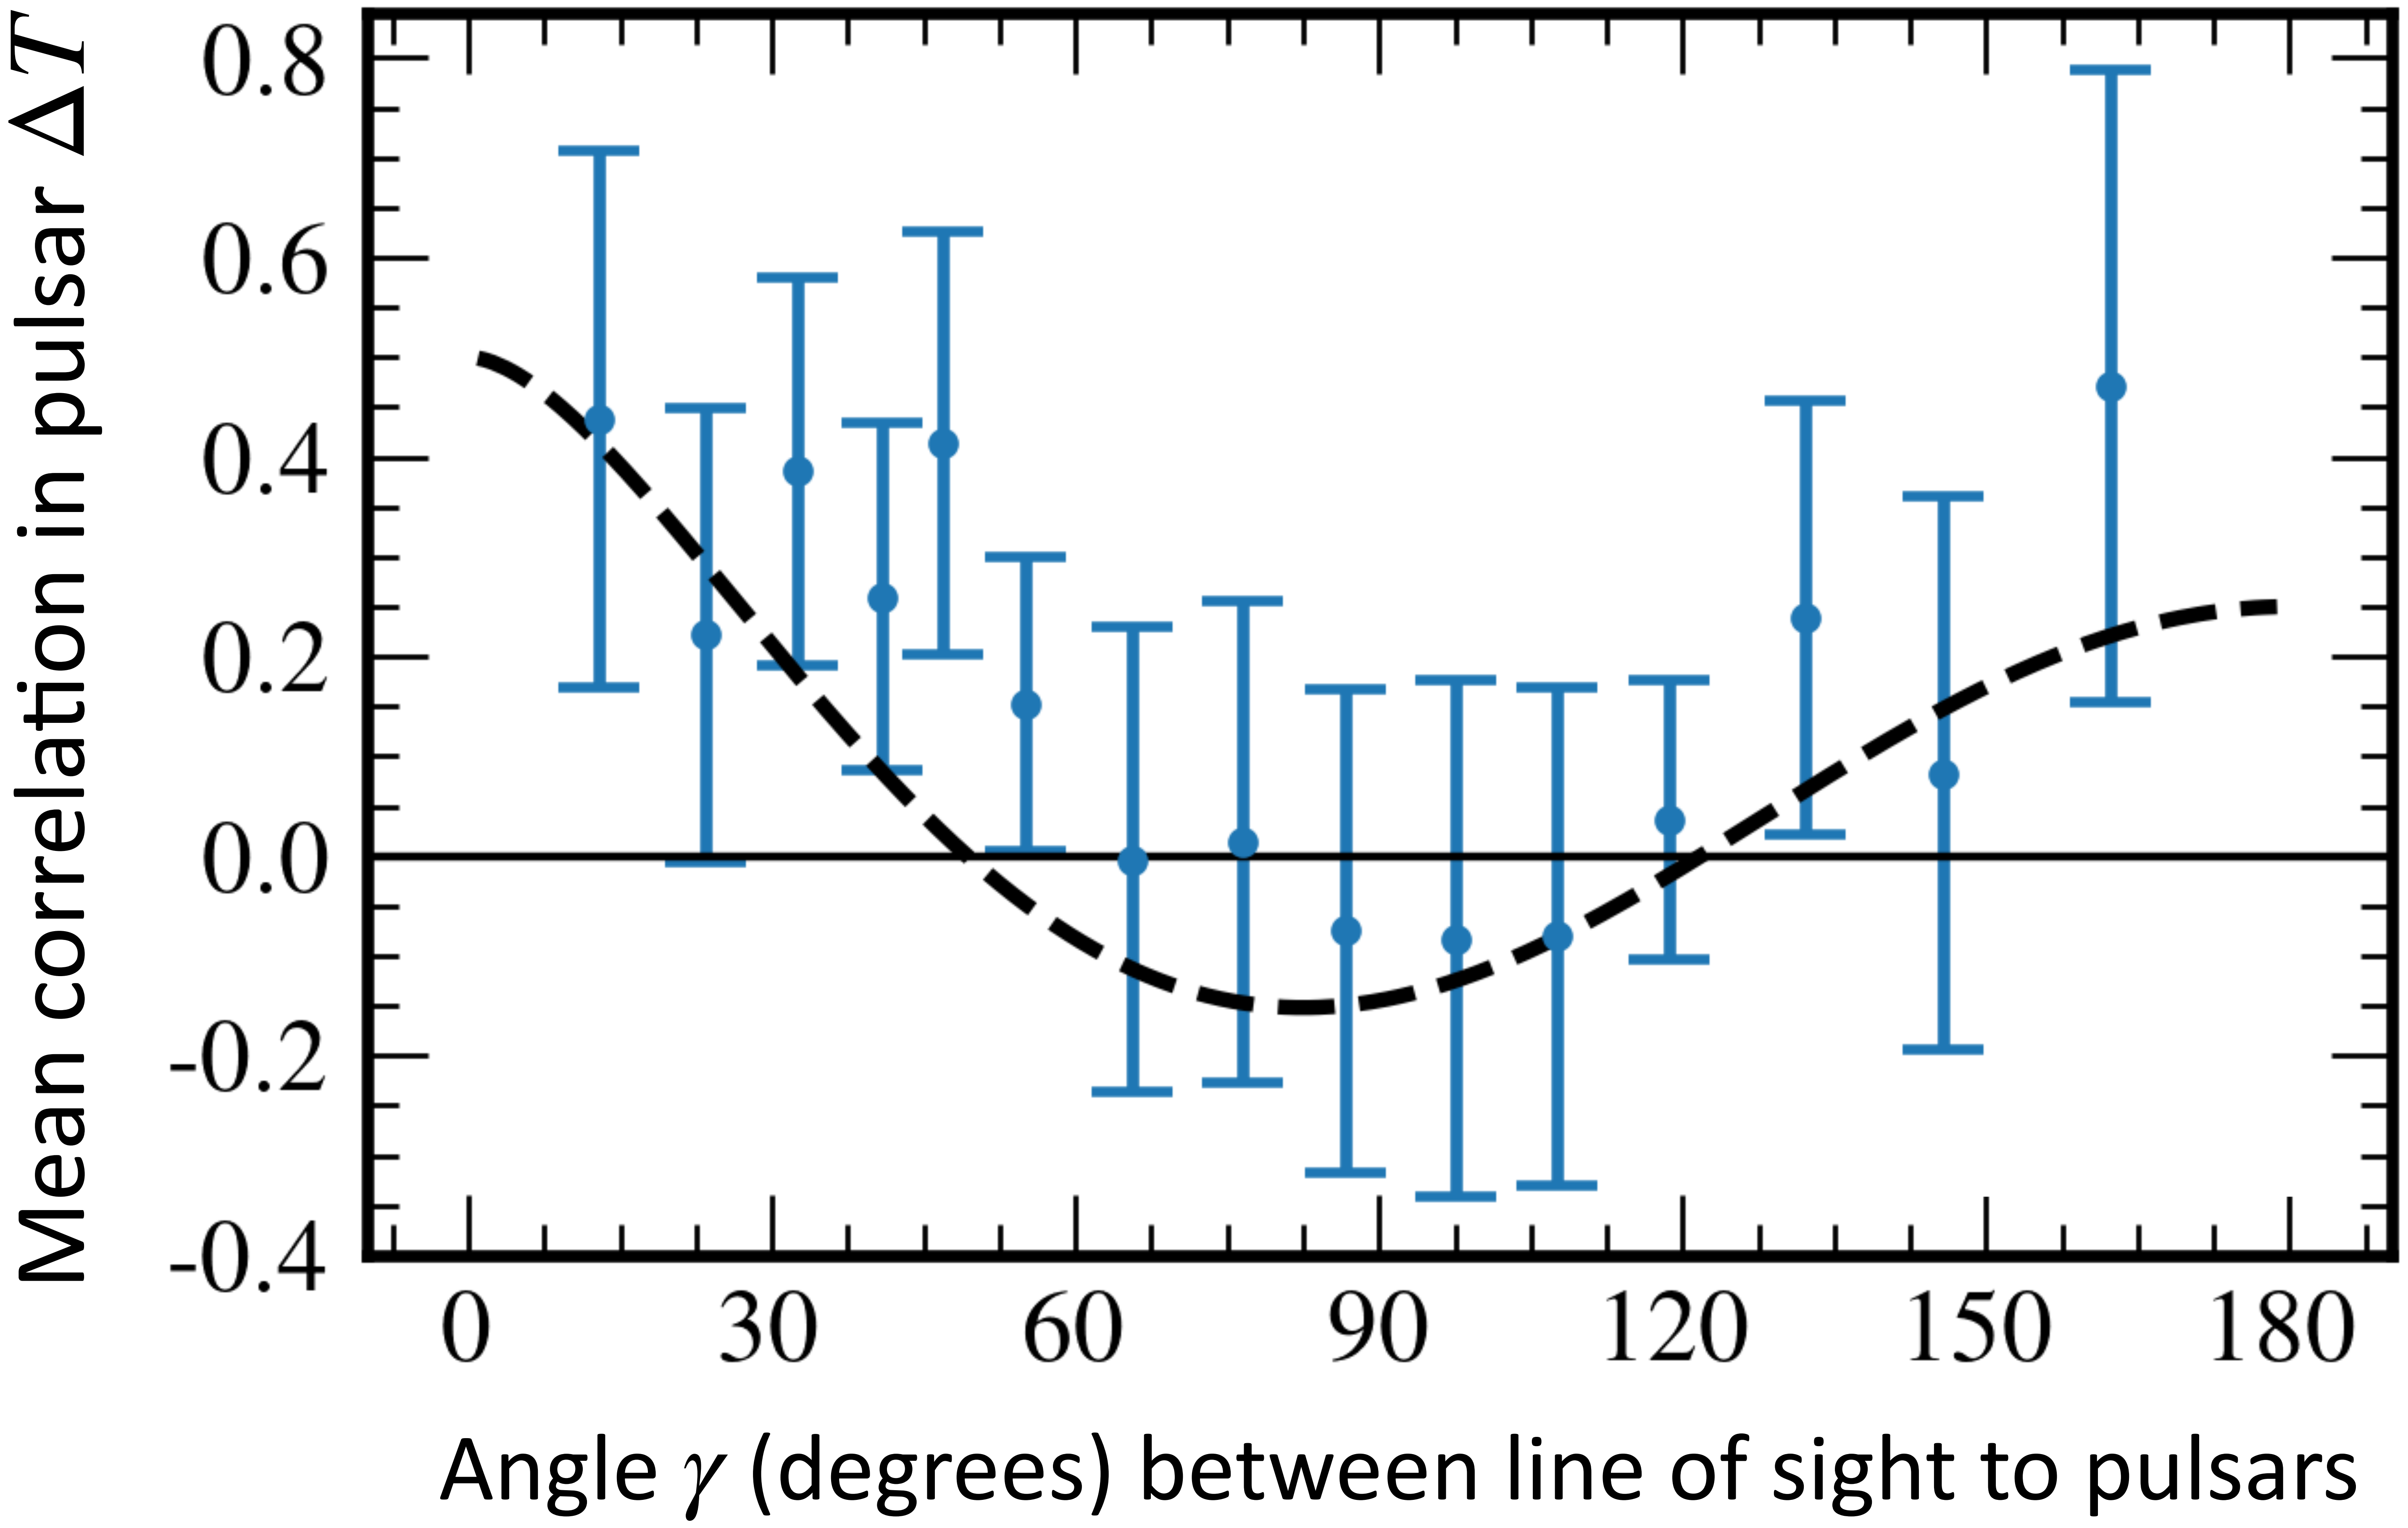

Example: NANOGrav 15-year data

- $N = 67$ pulsars $\Longrightarrow$ 2211 pulsar pairs
- 15 angular bins, so $2211/15 \approx 147$ pulsar pairs per angular bin
- Deviations from the Hellings and Downs curve arise from:
  - *Pulsar/measurement noise*
  - *Finite set of pulsars at particular sky locations*
  - *Cosmic variance*

THE QUESTION:
With more/better data, how close will the reconstructions come to the Hellings and Downs curve?

NANOGrav reconstruction with N = 67 pulsars

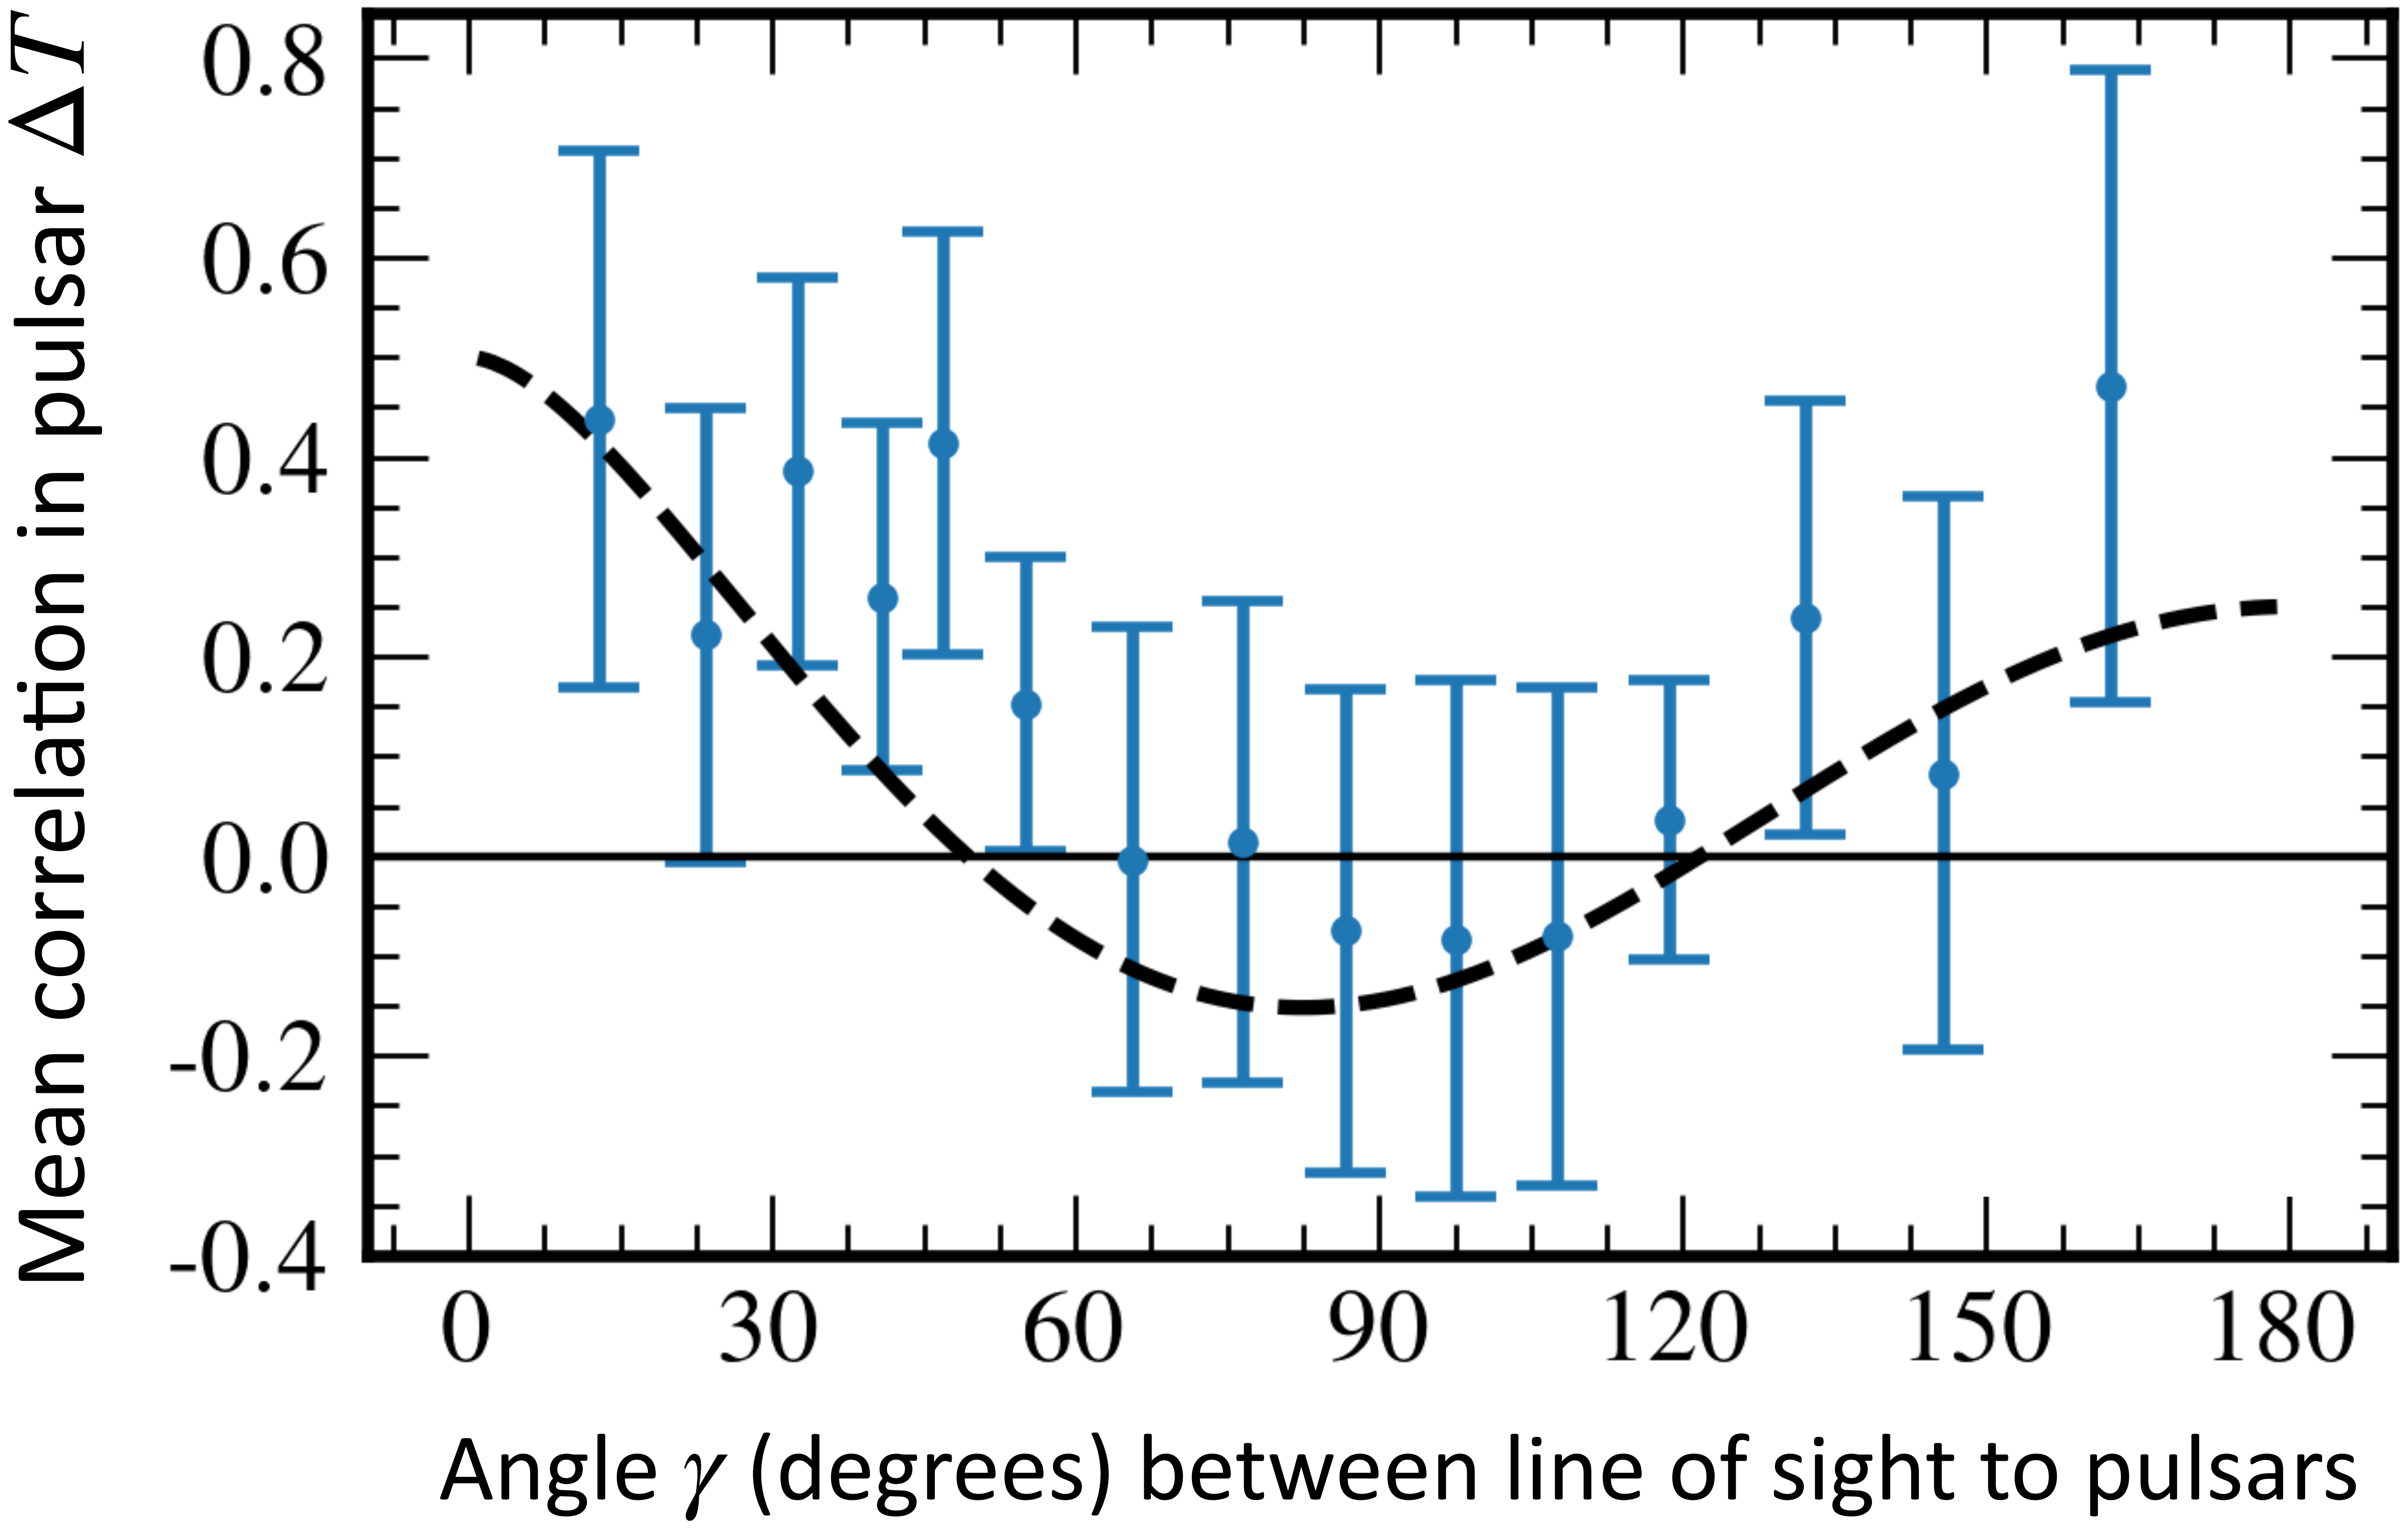

Example: NANOGrav 15-year data

- $N = 67$ pulsars $\Longrightarrow$ 2211 pulsar pairs
- 15 angular bins, so $2211/15 \approx 147$ pulsar pairs per angular bin
- Deviations from the Hellings and Downs curve arise from:
  - *Pulsar/measurement noise*
  - *Finite set of pulsars at particular sky locations*
  - *Cosmic variance*

THE QUESTION:
With more/better data, how close will the reconstructions come to the Hellings and Downs curve?

QUANTITATIVE VERSION:
After averaging over pulsar pairs at angle γ, write

$$\mu(\gamma) = \sum_l c_l P_l(\cos\gamma) \, ,$$

for some set of constants c_l. How much do these differ from the Hellings and Downs curve? That has

$$\langle c_l \rangle = \frac{2l+1}{(l+2)(l+1)l(l-1)} \, .$$

NANOGrav reconstruction with N = 67 pulsars

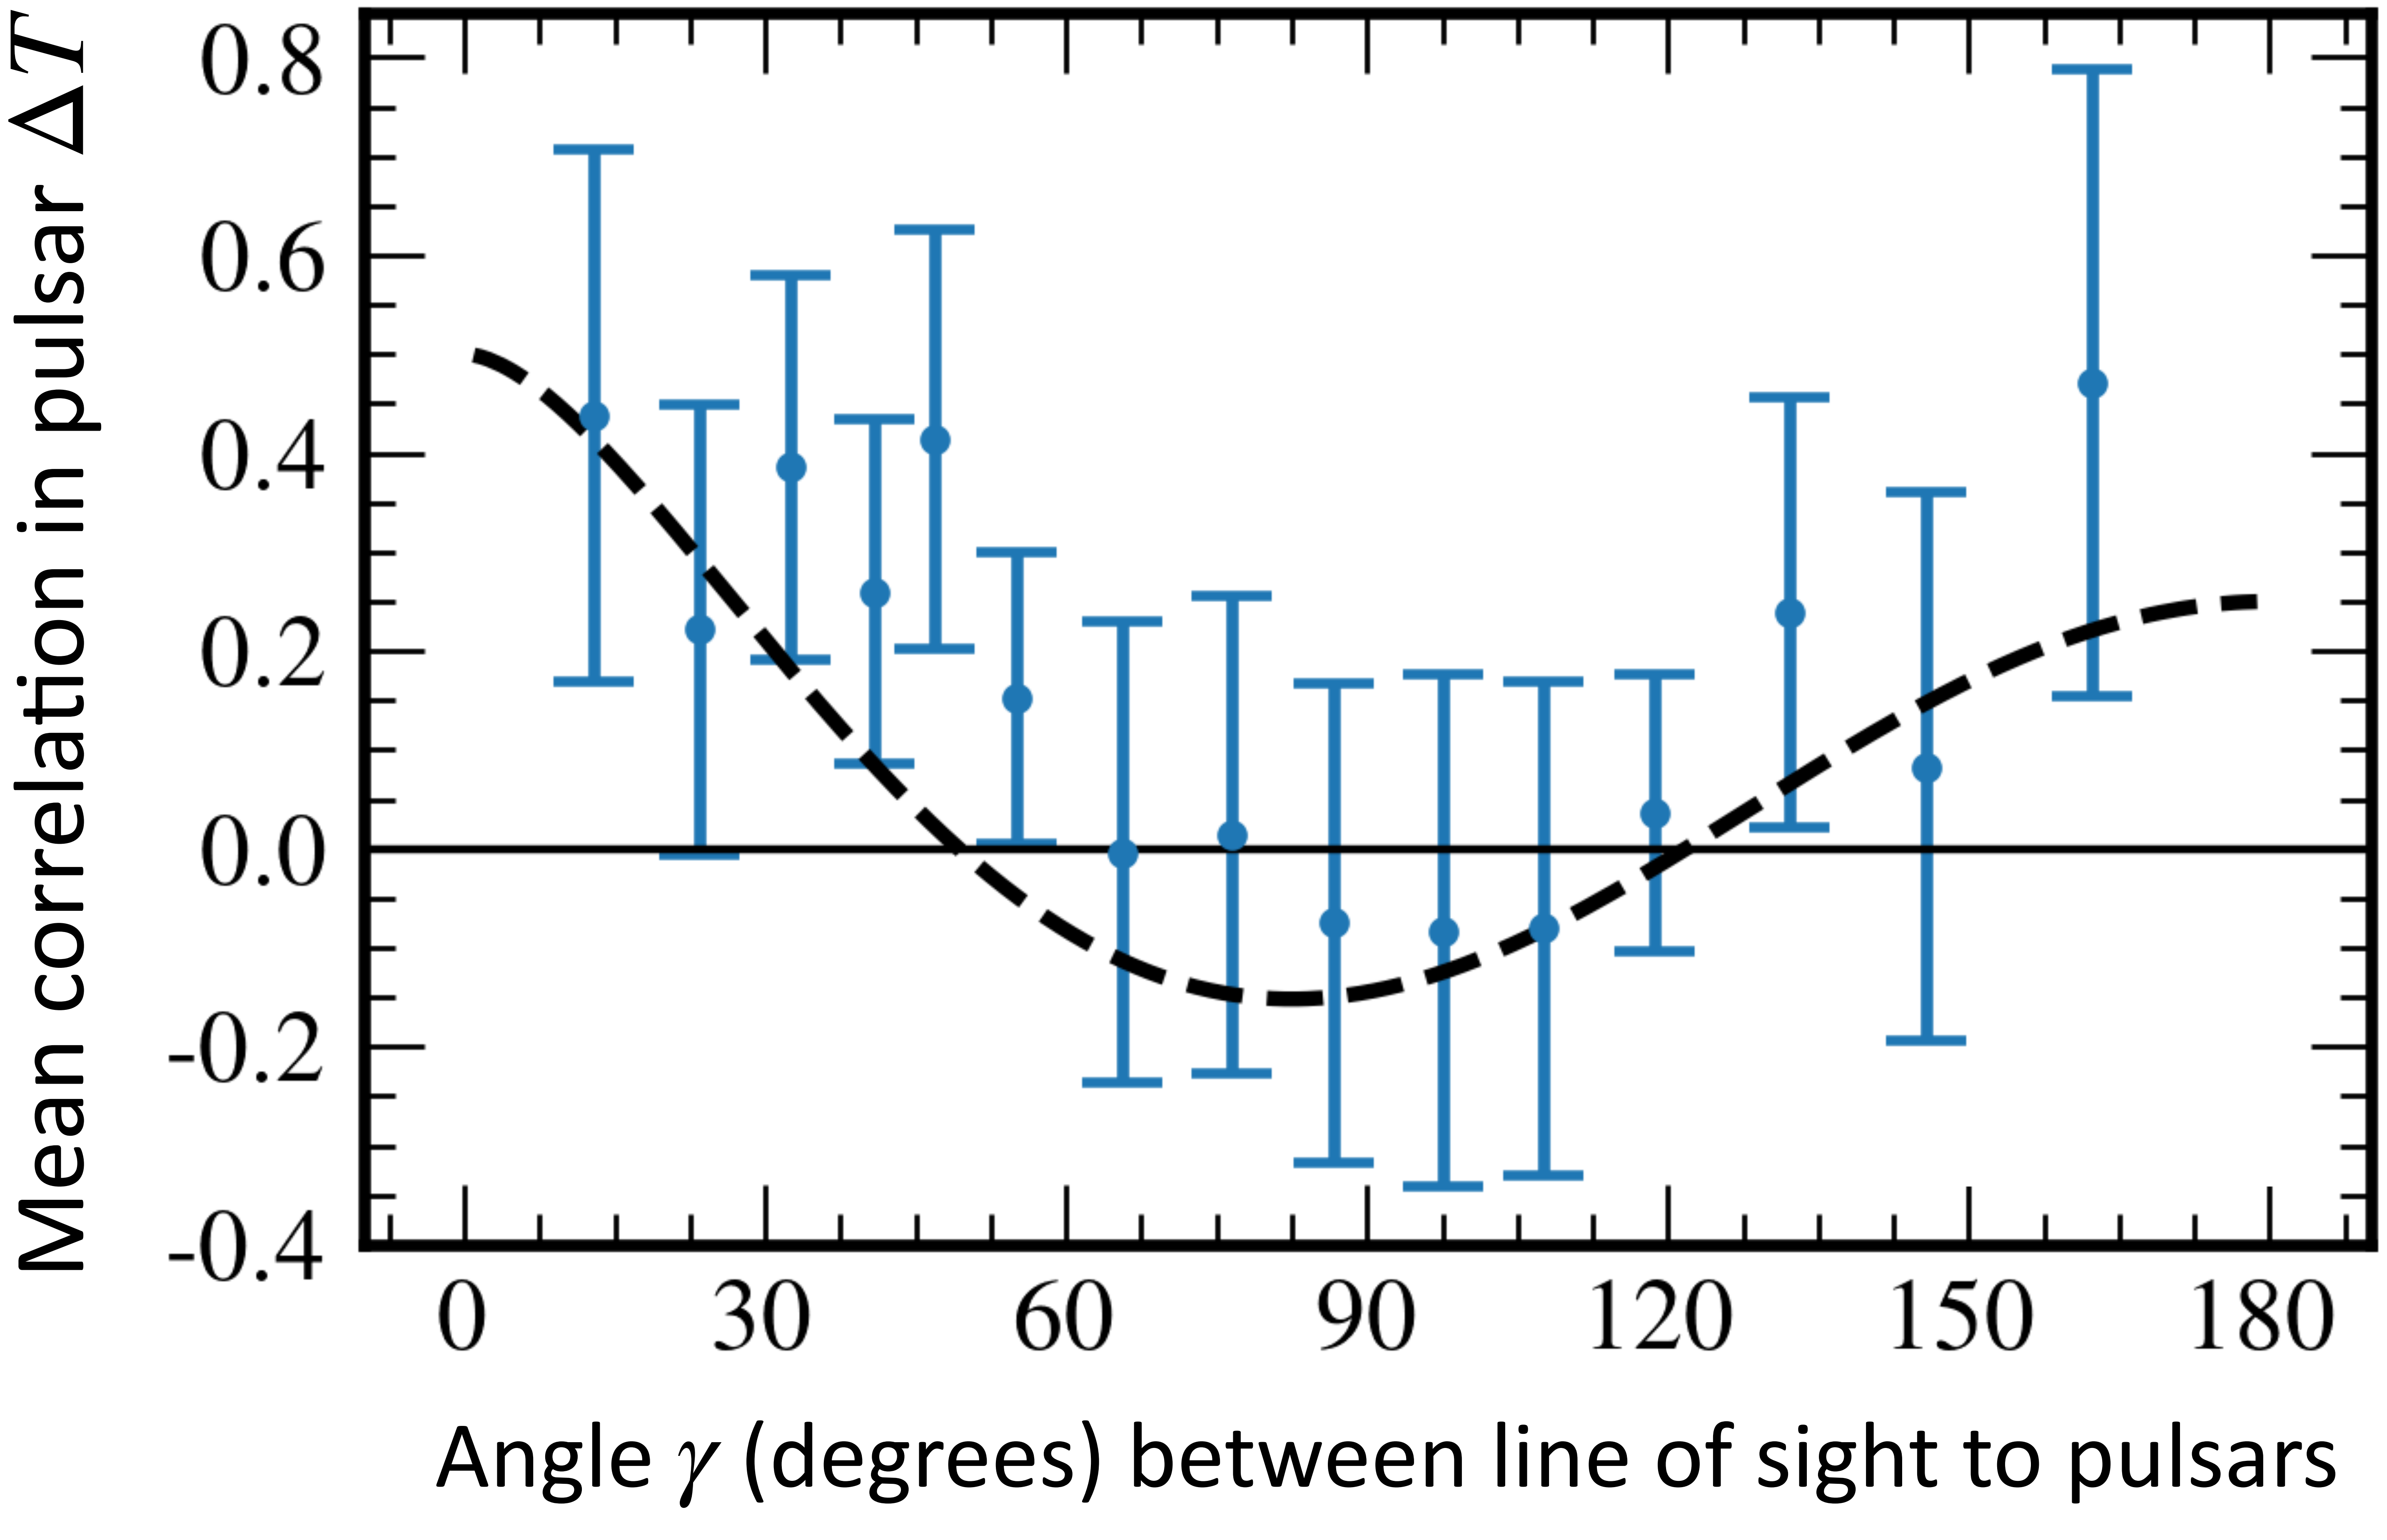

The answer: arXiv:2407.10968 (for Gaussian ensemble)

- Take angular bin at angle γ containing a set of pulsar pairs denoted $ab \in \gamma$.
- Let Z_a^j denote pulsar a's redshift Fourier amplitude in frequency bin j.
- Compute weights W_{ab}^{jk} according to "recipe".
- Best unbiased estimator of correlation is:

$$\hat{\mu} = \sum_{ab\in\gamma} \sum_{j,k} W_{ab}^{jk} Z_a^j Z_b^k \, .$$

- Variance of this estimator:

$$\sigma_{\hat{\mu}}^2 = \frac{\sigma_G^2}{N_f} \, .$$

- N_f = effective number of frequency bins in which the GW signal dominates the noise.
- σ_G^2 = geometrical quantity that depends upon where the pulsars are located on the sky

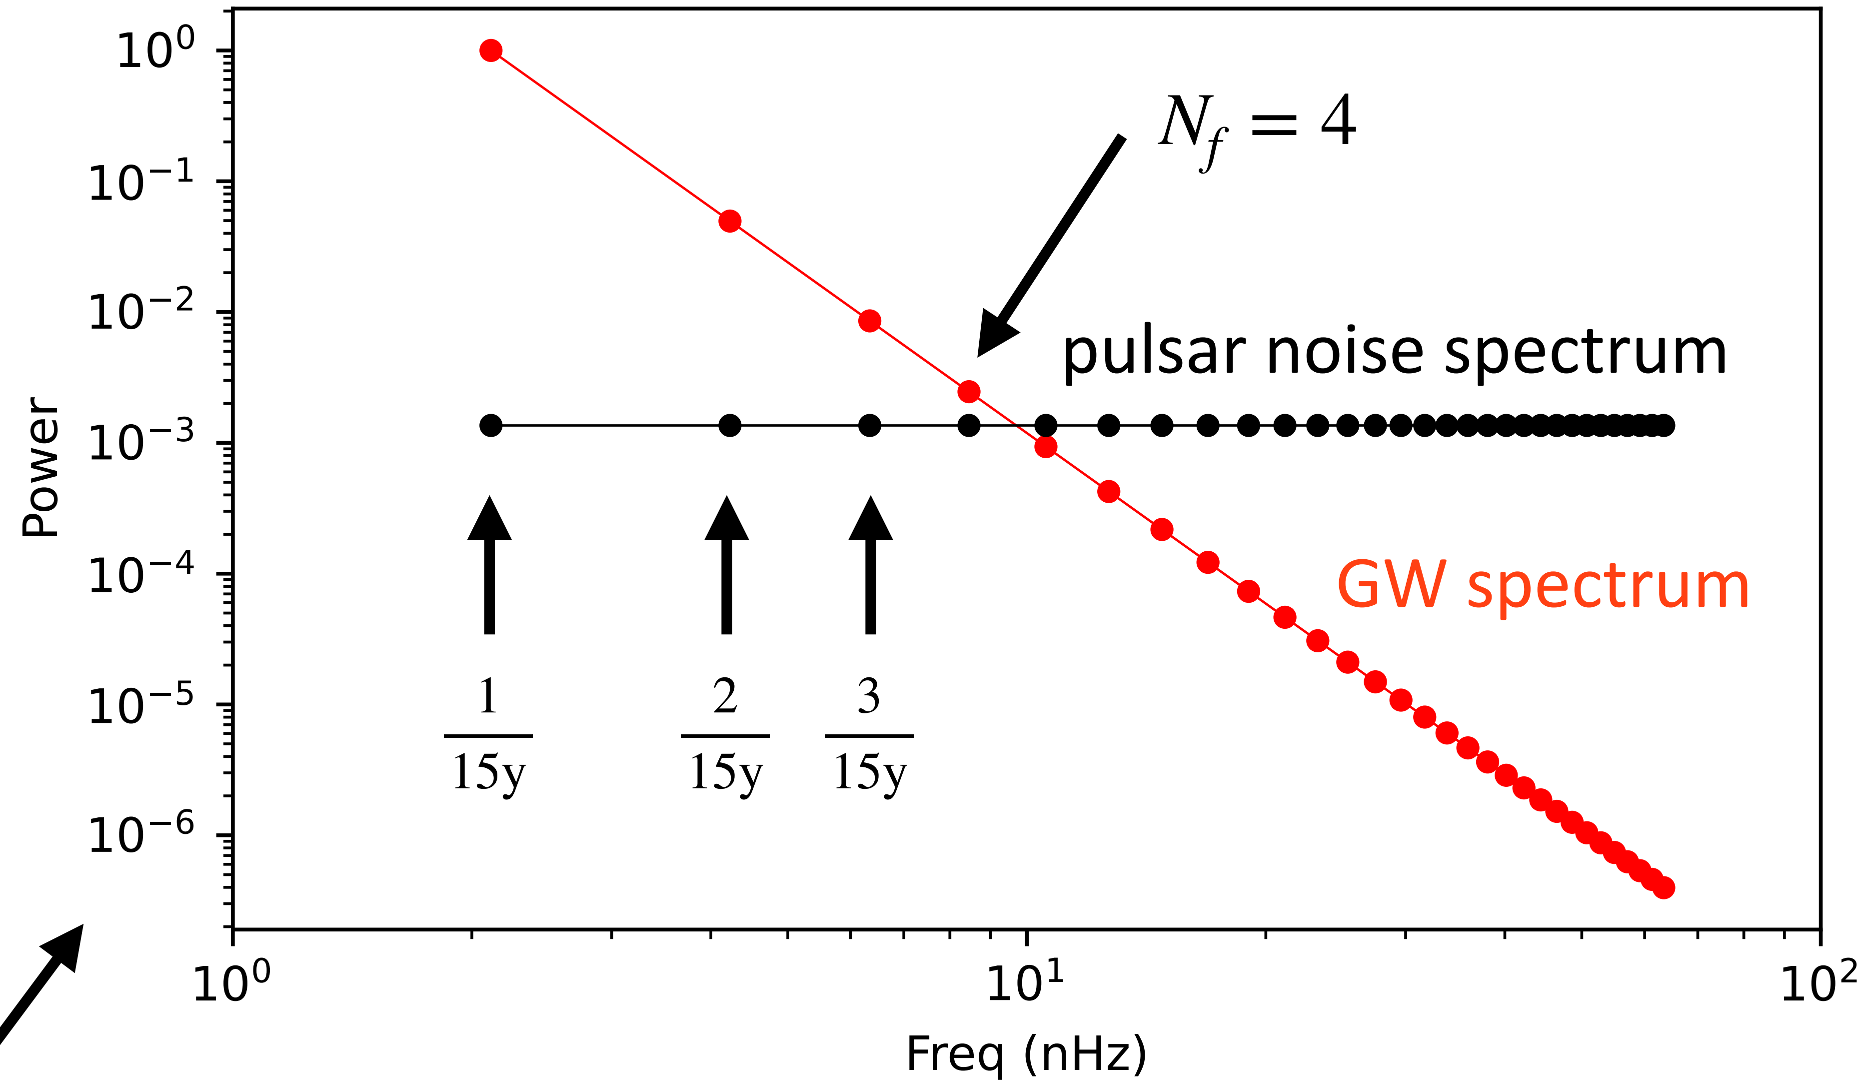

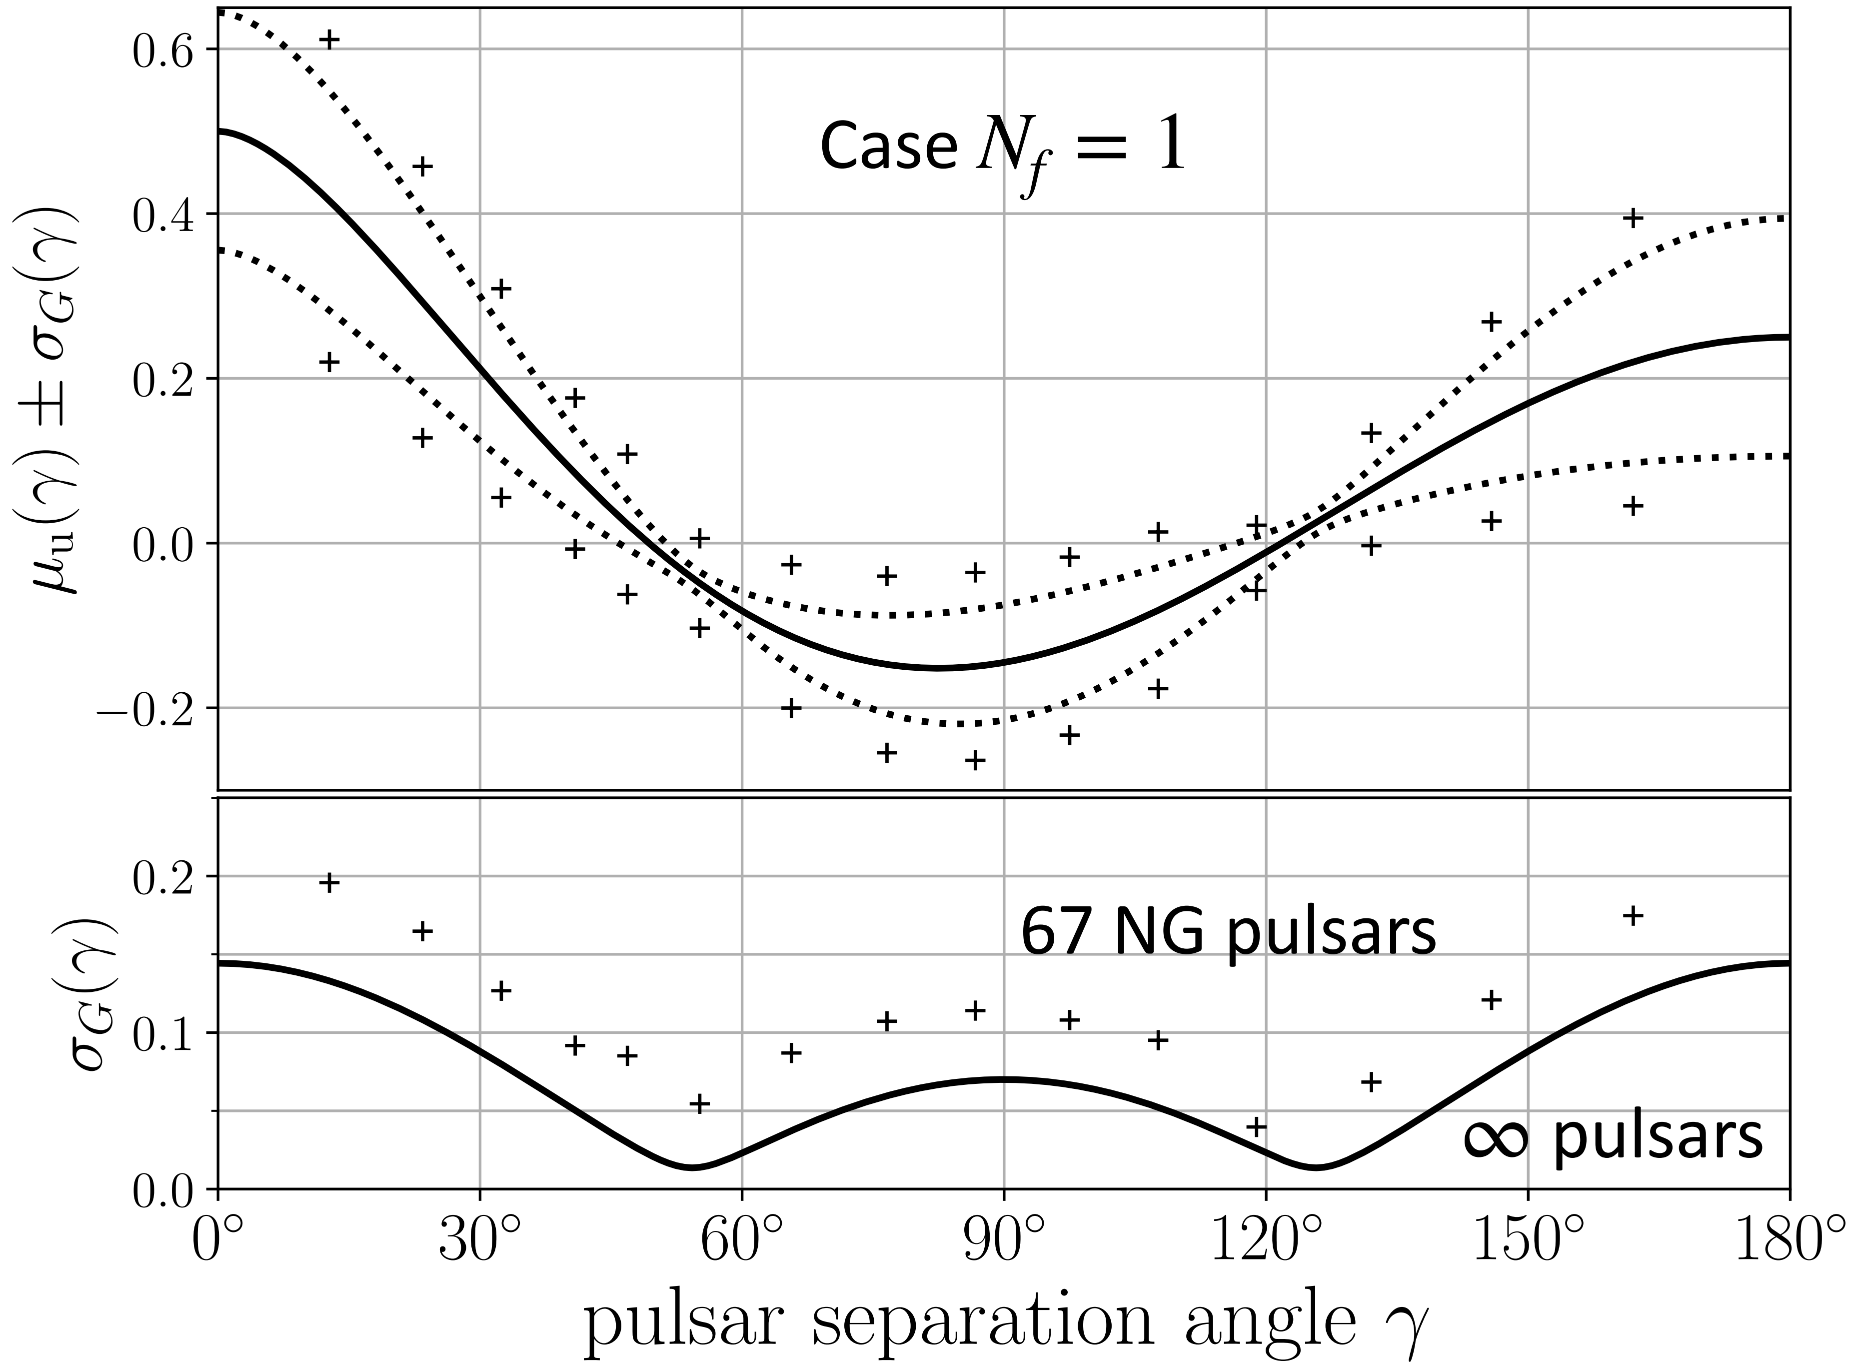

Generalizes Roebber & Holder 2017

- Considers a special case:
  - Infinitely many pulsars, all over the sky.
  - No pulsar or measurement noise.
  - Gaussian ensemble of unpolarized GW sources, all radiating at the same frequency $f = 1/T$, where $T = \text{observation time}$
  - Case $N_f = 1$: pure cosmic variance

The Astrophysical Journal, 835:21 (10pp), 2017 January 20
doi:10.3847/1538-4357/835/1/21

CrossMark

HARMONIC SPACE ANALYSIS OF PULSAR TIMING ARRAY REDSHIFT MAPS

Elinore Roebber[1] and Gilbert Holder[1,2]
[1] Department of Physics and McGill Space Institute, McGill University, 3600 rue University, Montréal, QC, H3A 2T8, Canada; roebbere@physics.mcgill.ca
[2] Department of Physics, University of Illinois at Urbana-Champaign, 1110 W Green Street, Urbana, IL 61801, USA
Received 2016 September 27; accepted 2016 November 11; published 2017 January 16

ABSTRACT

In this paper, we propose a new framework for treating the angular information in the pulsar timing array (PTA) response to a gravitational wave (GW) background based on standard cosmic microwave background techniques. We calculate the angular power spectrum of the all-sky gravitational redshift pattern induced at the Earth for both a single bright source of gravitational radiation and a statistically isotropic, unpolarized Gaussian random GW background. The angular power spectrum is the harmonic transform of the Hellings & Downs curve. We use the power spectrum to examine the expected variance in the Hellings & Downs curve in both cases. Finally, we discuss the extent to which PTAs are sensitive to the angular power spectrum and find that the power spectrum sensitivity is dominated by the quadrupole anisotropy of the gravitational redshift map.

Key words: gravitational waves – large-scale structure of universe – methods: analytical – pulsars: general

1. INTRODUCTION

Pulsar timing arrays (hereafter PTAs) are galactic-scale gravitational wave (GW) detectors based on the precise timing of millisecond pulsars across the sky (Foster & Backer 1990). The nanohertz frequency band of GWs accessible to PTAs has several potential production mechanisms, the most prominent of which is due to the inspiral of subparsec supermassive binary black holes (SMBBHs; see Lommen 2015, and references therein).

SMBBHs with chirp mass $\mathcal{M} > 10^8\,M_\odot$ at redshifts $z \lesssim 2$ are expected to produce most of the signal (e.g., Sesana et al. 2008). Since there should be many such sources evolving over times much longer than human timescales, the GW signal is expected to form a stochastic background with considerable source confusion. However, individual strong sources may stand out (Sesana et al. 2008; Ravi et al. 2012).

A passing GW induces compression and rarefaction of spacetime along its polarization axes. Periodic signals such as rays of light or pulse trains propagating through this region will be blue- or redshifted according to the strain of the GW. For periodic signals with frequency much higher than that of the GW, the shift will build up, producing a potentially measurable effect. This is the principle on which several models of GW detection are founded, including interferometers such as LIGO (Abbott et al. 2016) and LISA (eLISA Consortium 2013) as well as for PTAs (Lommen 2015). There are three PTA consortia: EPTA (Lentati et al. 2015), NANOGrav (Arzoumanian et al. 2016), and PPTA (Shannon et al. 2015). They combine together to form the IPTA (Verbiest et al. 2016).

PTAs search for integrated red- and blueshifts produced by GWs passing the Earth through the careful timing of a network of millisecond pulsars across the sky. Each millisecond pulsar produces an extraordinarily regular train of high-frequency pulses. If this pulse train is redshifted by a GW with typical strain $\lesssim 10^{-14}$ (e.g., Lommen 2015), no effect will be immediately visible, but after the passage of many pulses, a difference between the expected and actual time of arrival of pulses will become apparent. This timing residual is the basic measurable quantity for a PTA.

A GW of a given polarization will induce red- and blueshifts according to the geometry set by the direction of propagation of the GW and the projection of its polarization axes onto the sky. In order to sample this effect as fully as possible, PTAs time many millisecond pulsars across the sky and search for a correlation in their timing residuals which reflects the redshift pattern induced by GWs.

The expected form of this correlation is the Hellings & Downs curve (Hellings & Downs 1983), which was originally derived for a statistically isotropic unpolarized Gaussian random field of GWs. It also represents the expected correlation pattern for a single SMBBH source of GWs (Cornish & Sesana 2013).

However, the gravitational wave background (GWB) expected to be produced by a population of inspiraling SMBBHs will be neither completely dominated by a single source nor a completely stochastic Gaussian field. In general, it should be somewhere in between (e.g., Sesana et al. 2008).

Although much work has made use of the assumption that a stochastic background would have Gaussian statistics, single sources should not be neglected in the PTA search for GWs (Rosado et al. 2015). This is because the distribution of SMBBH sources is such that the rarest brightest sources dominate the signal in the GWB (Sesana et al. 2008; Kocsis & Sesana 2011; Ravi et al. 2012; Cornish & Sesana 2013; Roebber et al. 2016).

In light of this, it is of interest to search for angular information in the GWB. PTAs can be likened to a collection of GW antennas: their angular resolution is limited but not nonexistent. This has been taken advantage of in the attempt to search for individual sources and hotspots (e.g., Corbin & Cornish 2010; Sesana & Vecchio 2010; Babak & Sesana 2012; Simon et al. 2014). Additionally, recent works have characterized the correlation patterns expected for statistically anisotropic backgrounds made up of a large number of sources (Mingarelli et al. 2013; Taylor & Gair 2013) as well as attempting to map general GWBs (Cornish & van Haasteren 2014; Gair et al. 2014).

Many of these recent works have focused on estimating the distribution of GW signals produced by the source population, either in terms of power or components of the GW tensor. However, the GW strain is not directly measured by PTAs. The large effective beam patterns smear power out across the sky,

1

Generalizes Roebber & Holder 2017

- Considers a special case:
  - Infinitely many pulsars, all over the sky.
  - No pulsar or measurement noise.
  - Gaussian ensemble of unpolarized GW sources, all radiating at the same frequency $f = 1/T$, where $T = \text{observation time}$
  - Case $N_f = 1$: pure cosmic variance
- Averaging over pulsar pairs at angle γ

$$\mu(\gamma) = \sum_l c_l P_l(\cos\gamma)$$

- Gaussian ensemble $\Longrightarrow$

$$\langle c_l \rangle = \frac{2l+1}{(l+2)(l+1)l(l-1)}$$

- While R & H do not write it like this, variance is

$$\sigma_G^2(\gamma) = \sum_{l=2}^{\infty} \frac{\langle c_l \rangle^2}{2l+1} P_l^2(\cos\gamma)$$

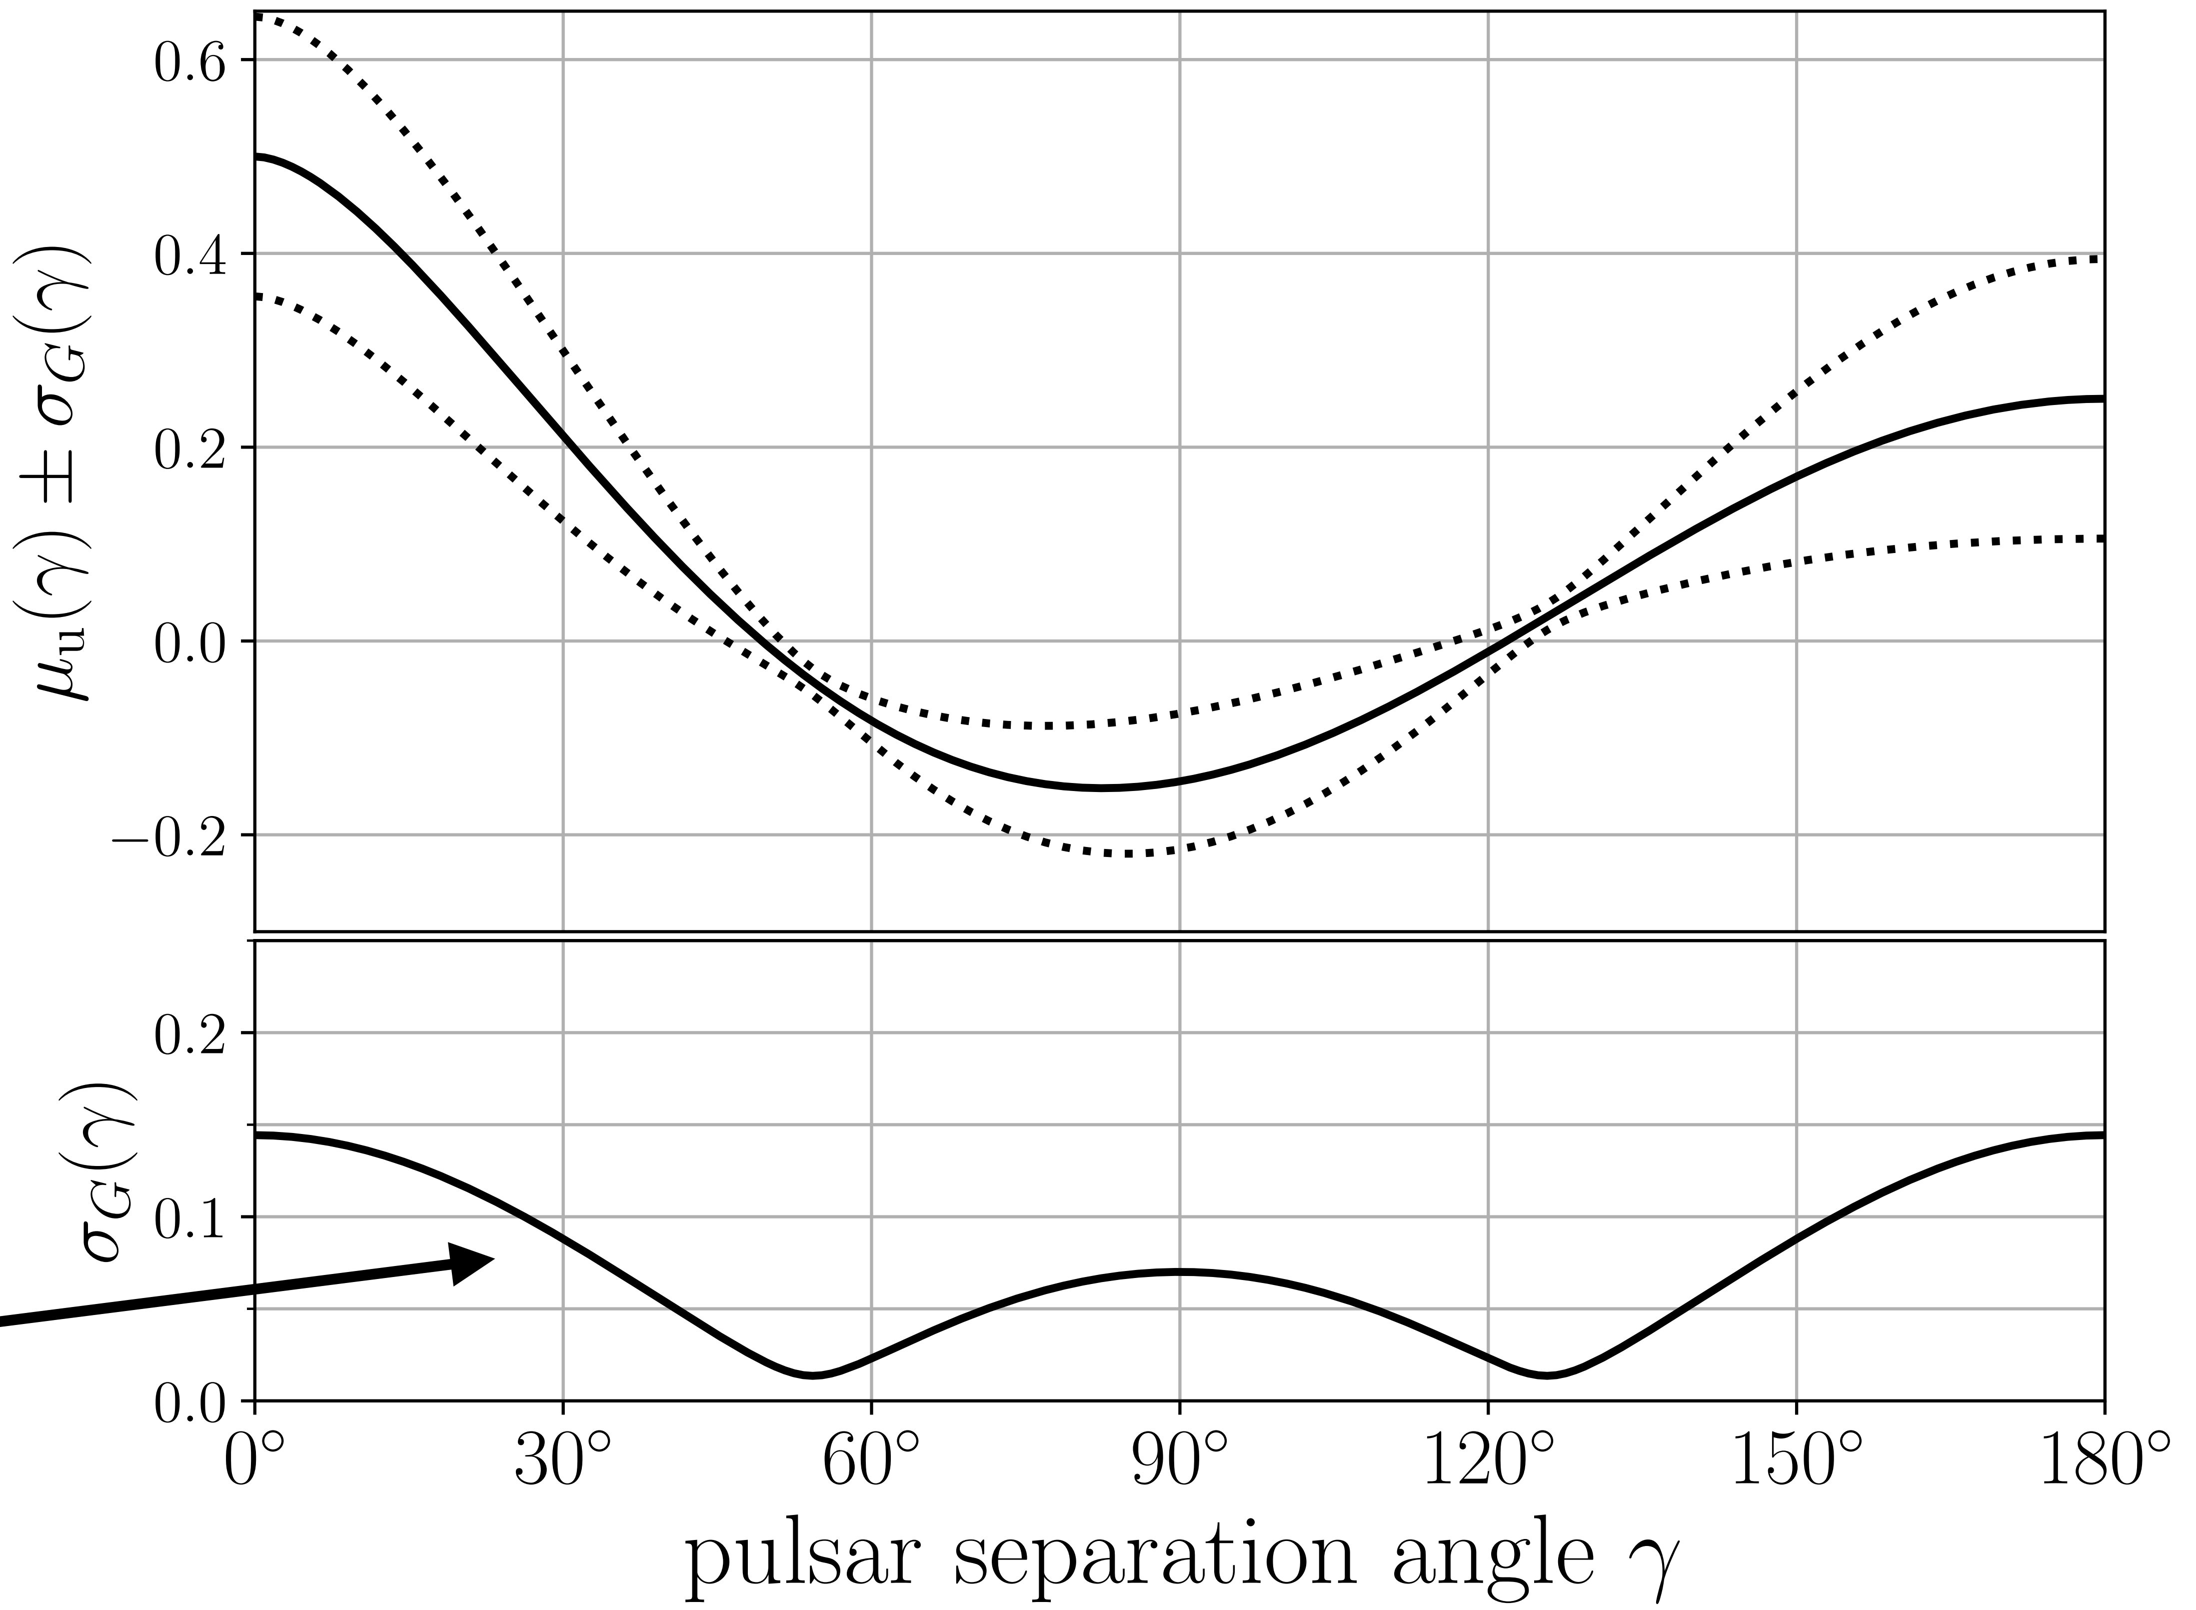

Noise-dominated or cosmic variance dominated

- Plots show effect of varying noise for fixed GW signal in a toy model
- Number of bins below crossing: N_c
- Expected SNR in angular bin: ρ
- Variance in Hellings & Downs correlation estimator:
  $$\sigma_{\hat{\mu}}^2 = \frac{\sigma_G^2}{N_f}$$
- Noise dominated case:
  ρ is small
  $N_f \ll 1$
  $\sigma_{\hat{\mu}}^2 \approx \mu_{\mathrm{u}}^2(\gamma)/\rho^2$
  Deviations from Hellings-Downs curve are dominated by noise
- Signal dominated case:
  ρ is big
  $N_f \gg 1$
  $\sigma_{\hat{\mu}}^2 \approx \sigma_G^2/N_c$
  Deviations from Hellings-Downs curve are dominated by cosmic variance

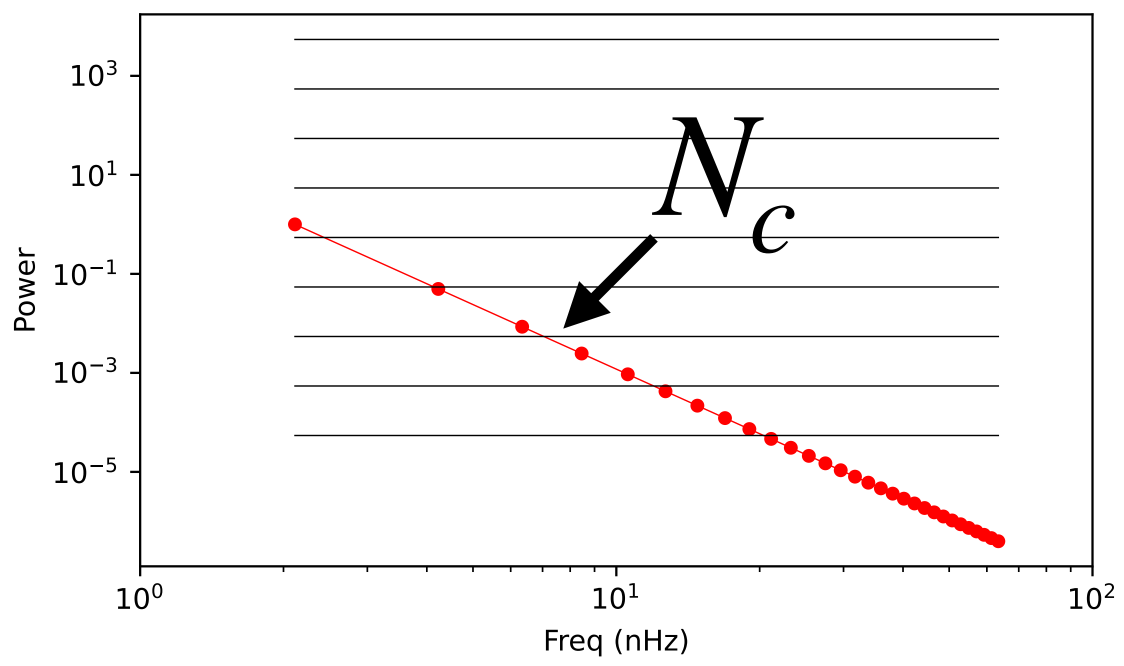

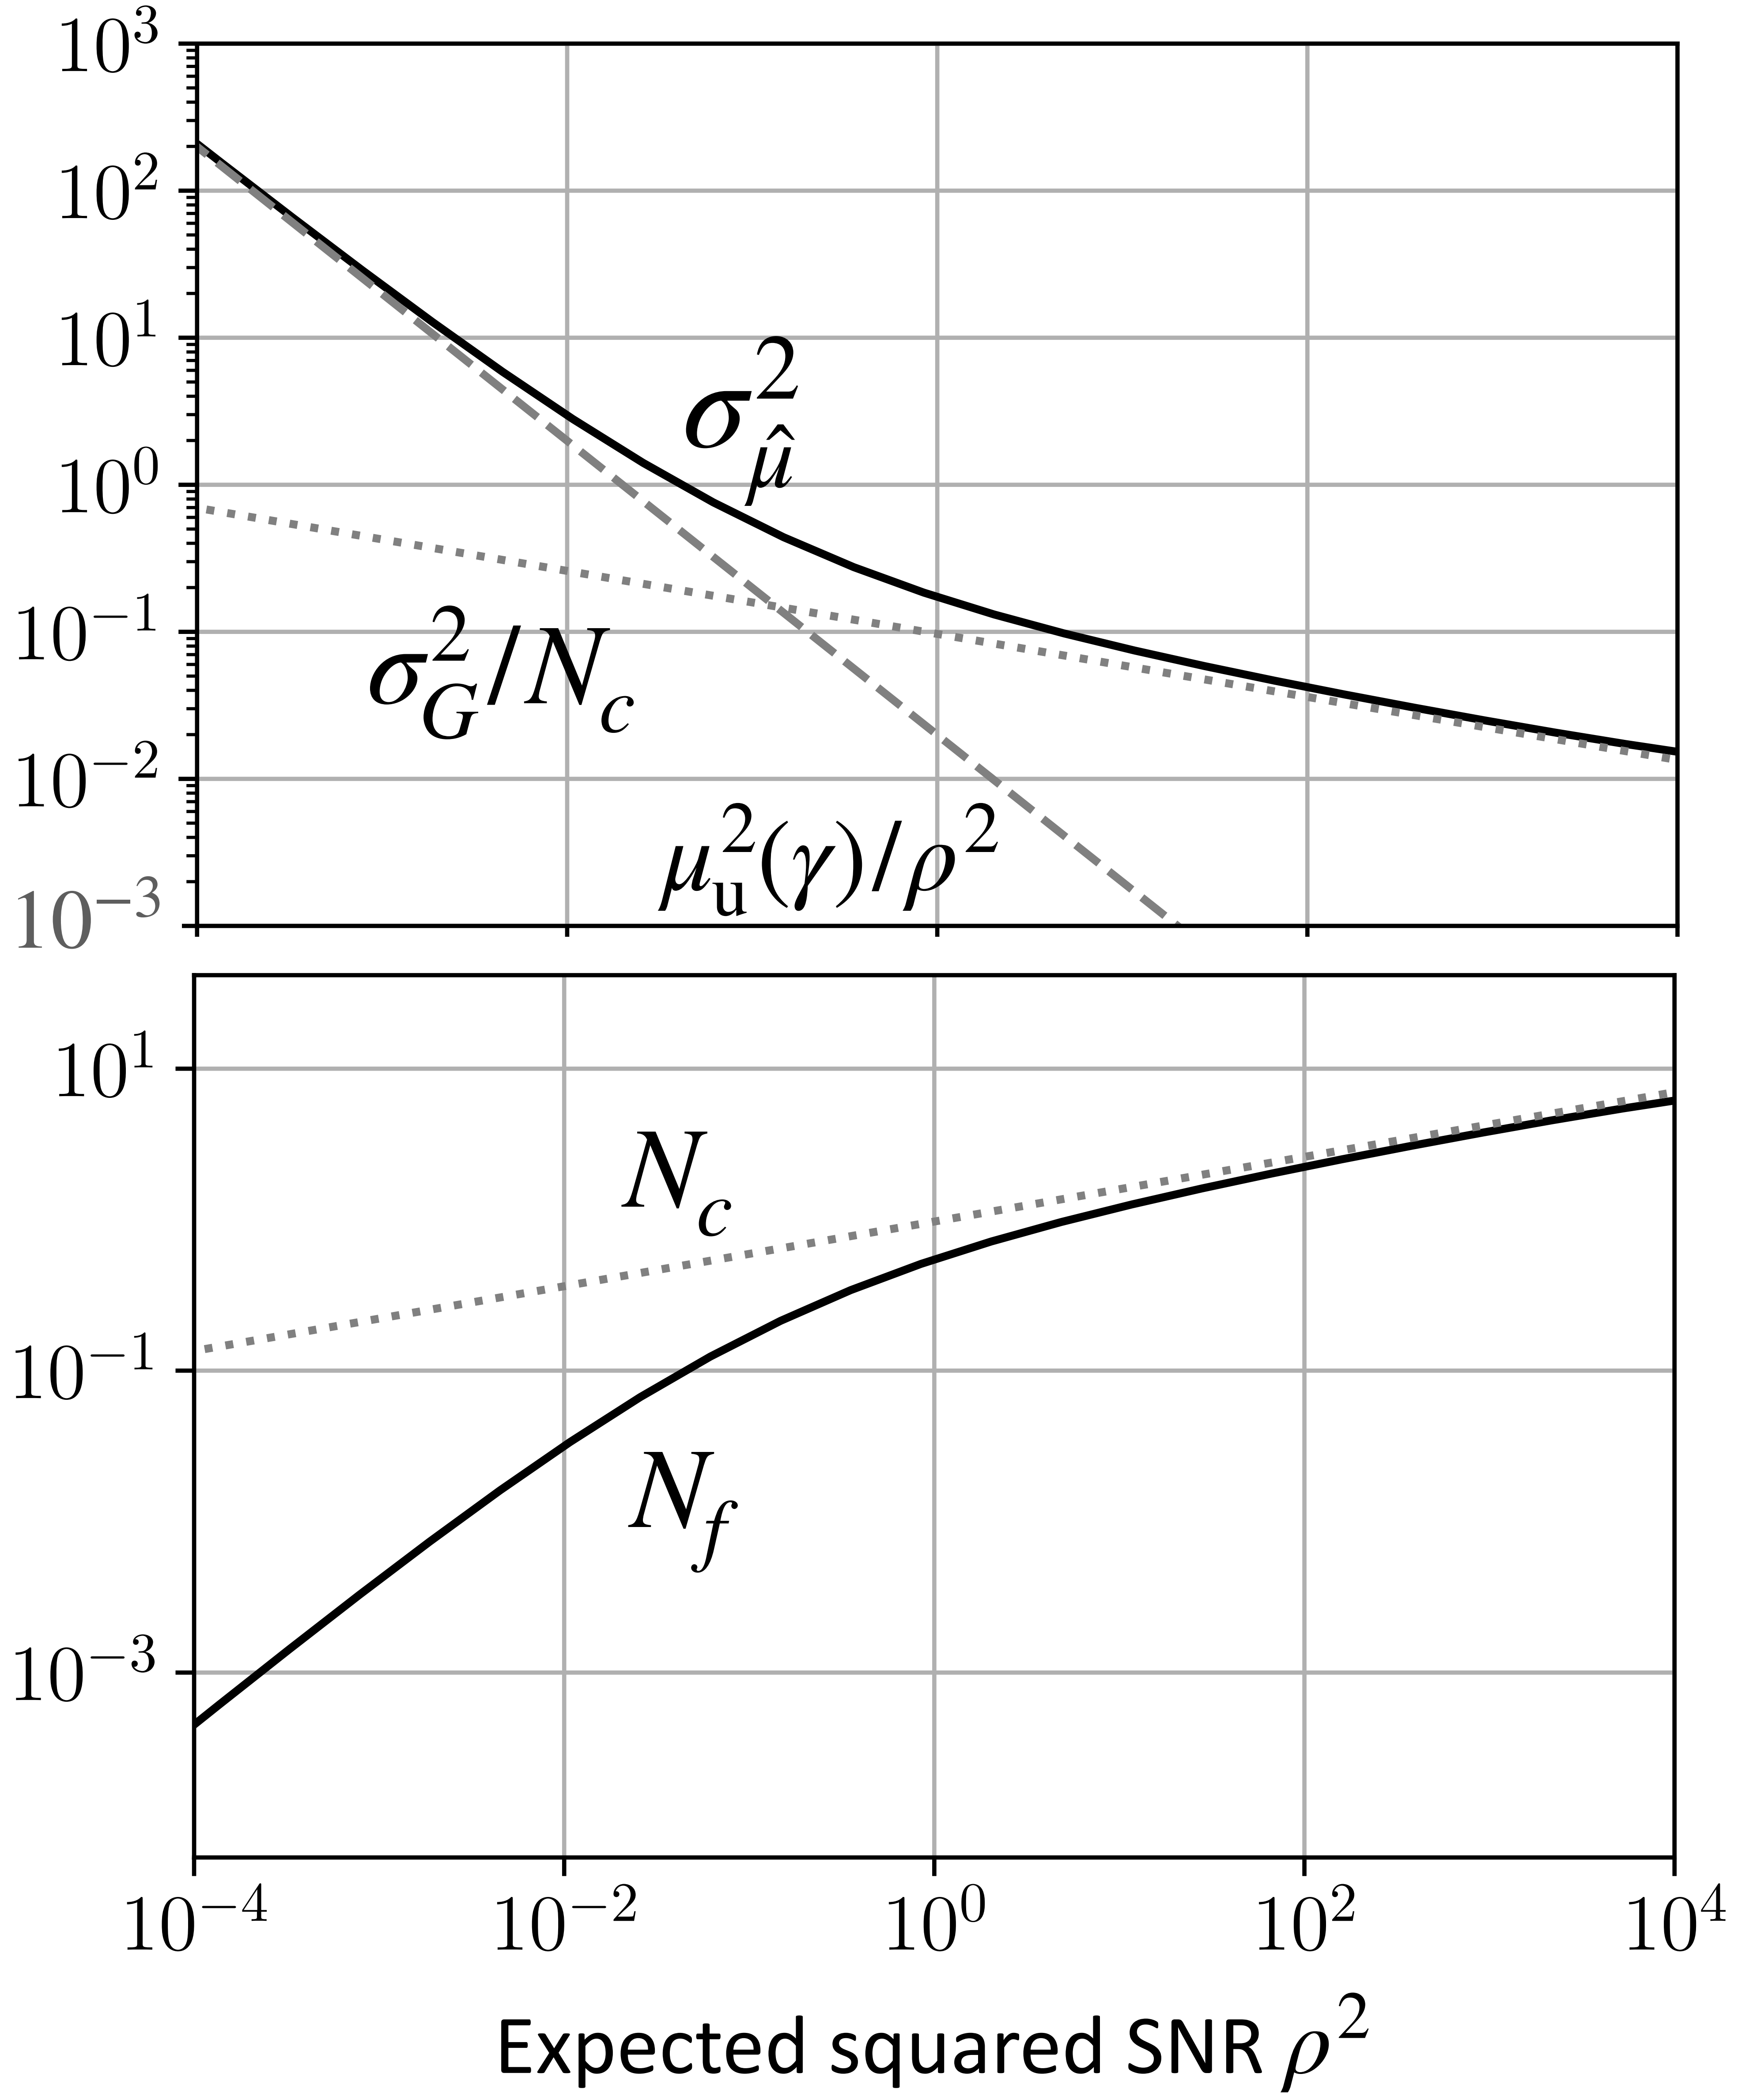

PTA EXPERTS ONLY

Work with pulsar pairs ab in angular bin at angle γ

Pulsar labels: a, b, c, d

HD correlation:

$$\boldsymbol{\mu} \equiv \mu_{ab} \equiv (1 + \delta_{ab})\mu_{\rm u}(\gamma_{ab})$$

Geometric covariance matrix:

$$\boldsymbol{G} \equiv G_{ab,cd} \equiv \mu_{ac}\mu_{bd} + \mu_{ad}\mu_{bc}$$

Geometric factor in the variance:

$$\sigma_G^2 \equiv \frac{1}{2}\frac{\mu_{\rm u}^2(\gamma)}{\boldsymbol{\mu}^t\,\boldsymbol{G}^{-1}\,\boldsymbol{\mu}}$$

Frequency bin labels: j, k, ℓ, m

Frequency of bin j is $f_j = j/T$, for observation time T

Covariance matrix for GW power:

$$\boldsymbol{H} \equiv H_{jk} \equiv 4\pi \int df\, H(f)\, \mathrm{sinc}\big(\pi(f - f_j)T\big)\, \mathrm{sinc}\big(\pi(f - f_k)T\big)$$

Covariance matrix for power of pulsar a:

$$N_a^{jk} \equiv 4\pi \int df\, N_a(f)\, \mathrm{sinc}\big(\pi(f - f_j)T\big)\, \mathrm{sinc}\big(\pi(f - f_k)T\big)$$

Reflect covariance matrix across anti-diagonal:

$$\overline{\boldsymbol{H}} \equiv \overline{H}_{jk} \equiv H_{j,-k}$$

Covariance matrix of pulsar redshifts:

$$\begin{aligned}\mathsf{C}_{ab,cd}^{jk,\ell m} = {} & \mu_{ac}\mu_{bd}H_{j\ell}H_{km} + \mu_{ad}\mu_{bc}H_{jm}H_{k\ell} + \\ & \delta_{ac}\mu_{bd}N_a^{j\ell}H_{km} + \delta_{ad}\mu_{bc}N_a^{jm}H_{k\ell} + \\ & \mu_{ac}\delta_{bd}H_{j\ell}N_b^{km} + \mu_{ad}\delta_{bc}H_{jm}N_b^{k\ell} + \\ & \delta_{ac}\delta_{bd}N_a^{j\ell}N_b^{km} + \delta_{ad}\delta_{bc}N_a^{jm}N_b^{k\ell}\end{aligned}$$

Symmetric part:

$$\boldsymbol{C} \equiv \mathsf{C}_{ab,cd}^{(jk),(\ell m)}$$

Weights for optimal estimator:

$$\boldsymbol{W} = \mu_{\rm u}(\gamma)\frac{(\boldsymbol{\mu}\overline{\boldsymbol{H}})^t\boldsymbol{C}^{-1}}{(\boldsymbol{\mu}\overline{\boldsymbol{H}})^t\boldsymbol{C}^{-1}(\boldsymbol{\mu}\overline{\boldsymbol{H}})}$$

Number of signal-dominated frequency bins:

$$N_f \equiv \frac{(\boldsymbol{\mu}\overline{\boldsymbol{H}})^t\boldsymbol{C}^{-1}(\boldsymbol{\mu}\overline{\boldsymbol{H}})}{2\boldsymbol{\mu}^t\boldsymbol{G}^{-1}\boldsymbol{\mu}}$$

Summary

- When the Hellings and Downs correlation is reconstructed, the values do not agree with the famous Hellings and Downs curve.
- If observations are noise dominated, then the deviations are due to noise.
- If the observations are signal dominated, then the deviations are due to cosmic variance.
- Cosmic variance arises because our universe is one instance. Interference between GW sources causes the pulsar-averaged correlation to differ from the mean.
- This cosmic variance can be predicted. It is a geometric quantity that depends upon the pulsar sky locations, divided by the (effective) number of signal-dominated frequency bins.
- In effect, each signal-dominated frequency bin provides an independent sky map.