NR, EOB & DATA ANALYSIS
LECTURE 3: GW DATA ANALYSIS

AMPLITUDES, STRONG-GRAVITY GRAVITY AND RESUMMATION @ NORDITA, APRIL 2026

PATRICIA SCHMIDT, UNIVERSITY OF BIRMINGHAM

UNIVERSITY OF BIRMINGHAM | GRAVITATIONAL WAVE ASTRONOMY

GW DATA ANALYSIS

(Some) Recommended literature (mostly reviews):

- ★ Luc Blanchet, Living Reviews in Relativity: https://link.springer.com/article/10.12942/lrr-2014-2
- ★ Buonanno & Sathyaprakash: https://arxiv.org/abs/1410.7832
- ★ Talbot & Thrane: https://arxiv.org/abs/1809.02293
- ★ Sathyaprakash & Schutz: https://link.springer.com/article/10.12942/lrr-2009-2
- ★ Cutler & Flanagan: https://arxiv.org/abs/gr-qc/9402014
- ★ Veitch et al.: https://arxiv.org/abs/1409.7215
- ★ Finn & Chernoff: https://arxiv.org/abs/gr-qc/9301003
- ★ Moore et al.: https://arxiv.org/abs/1408.0740

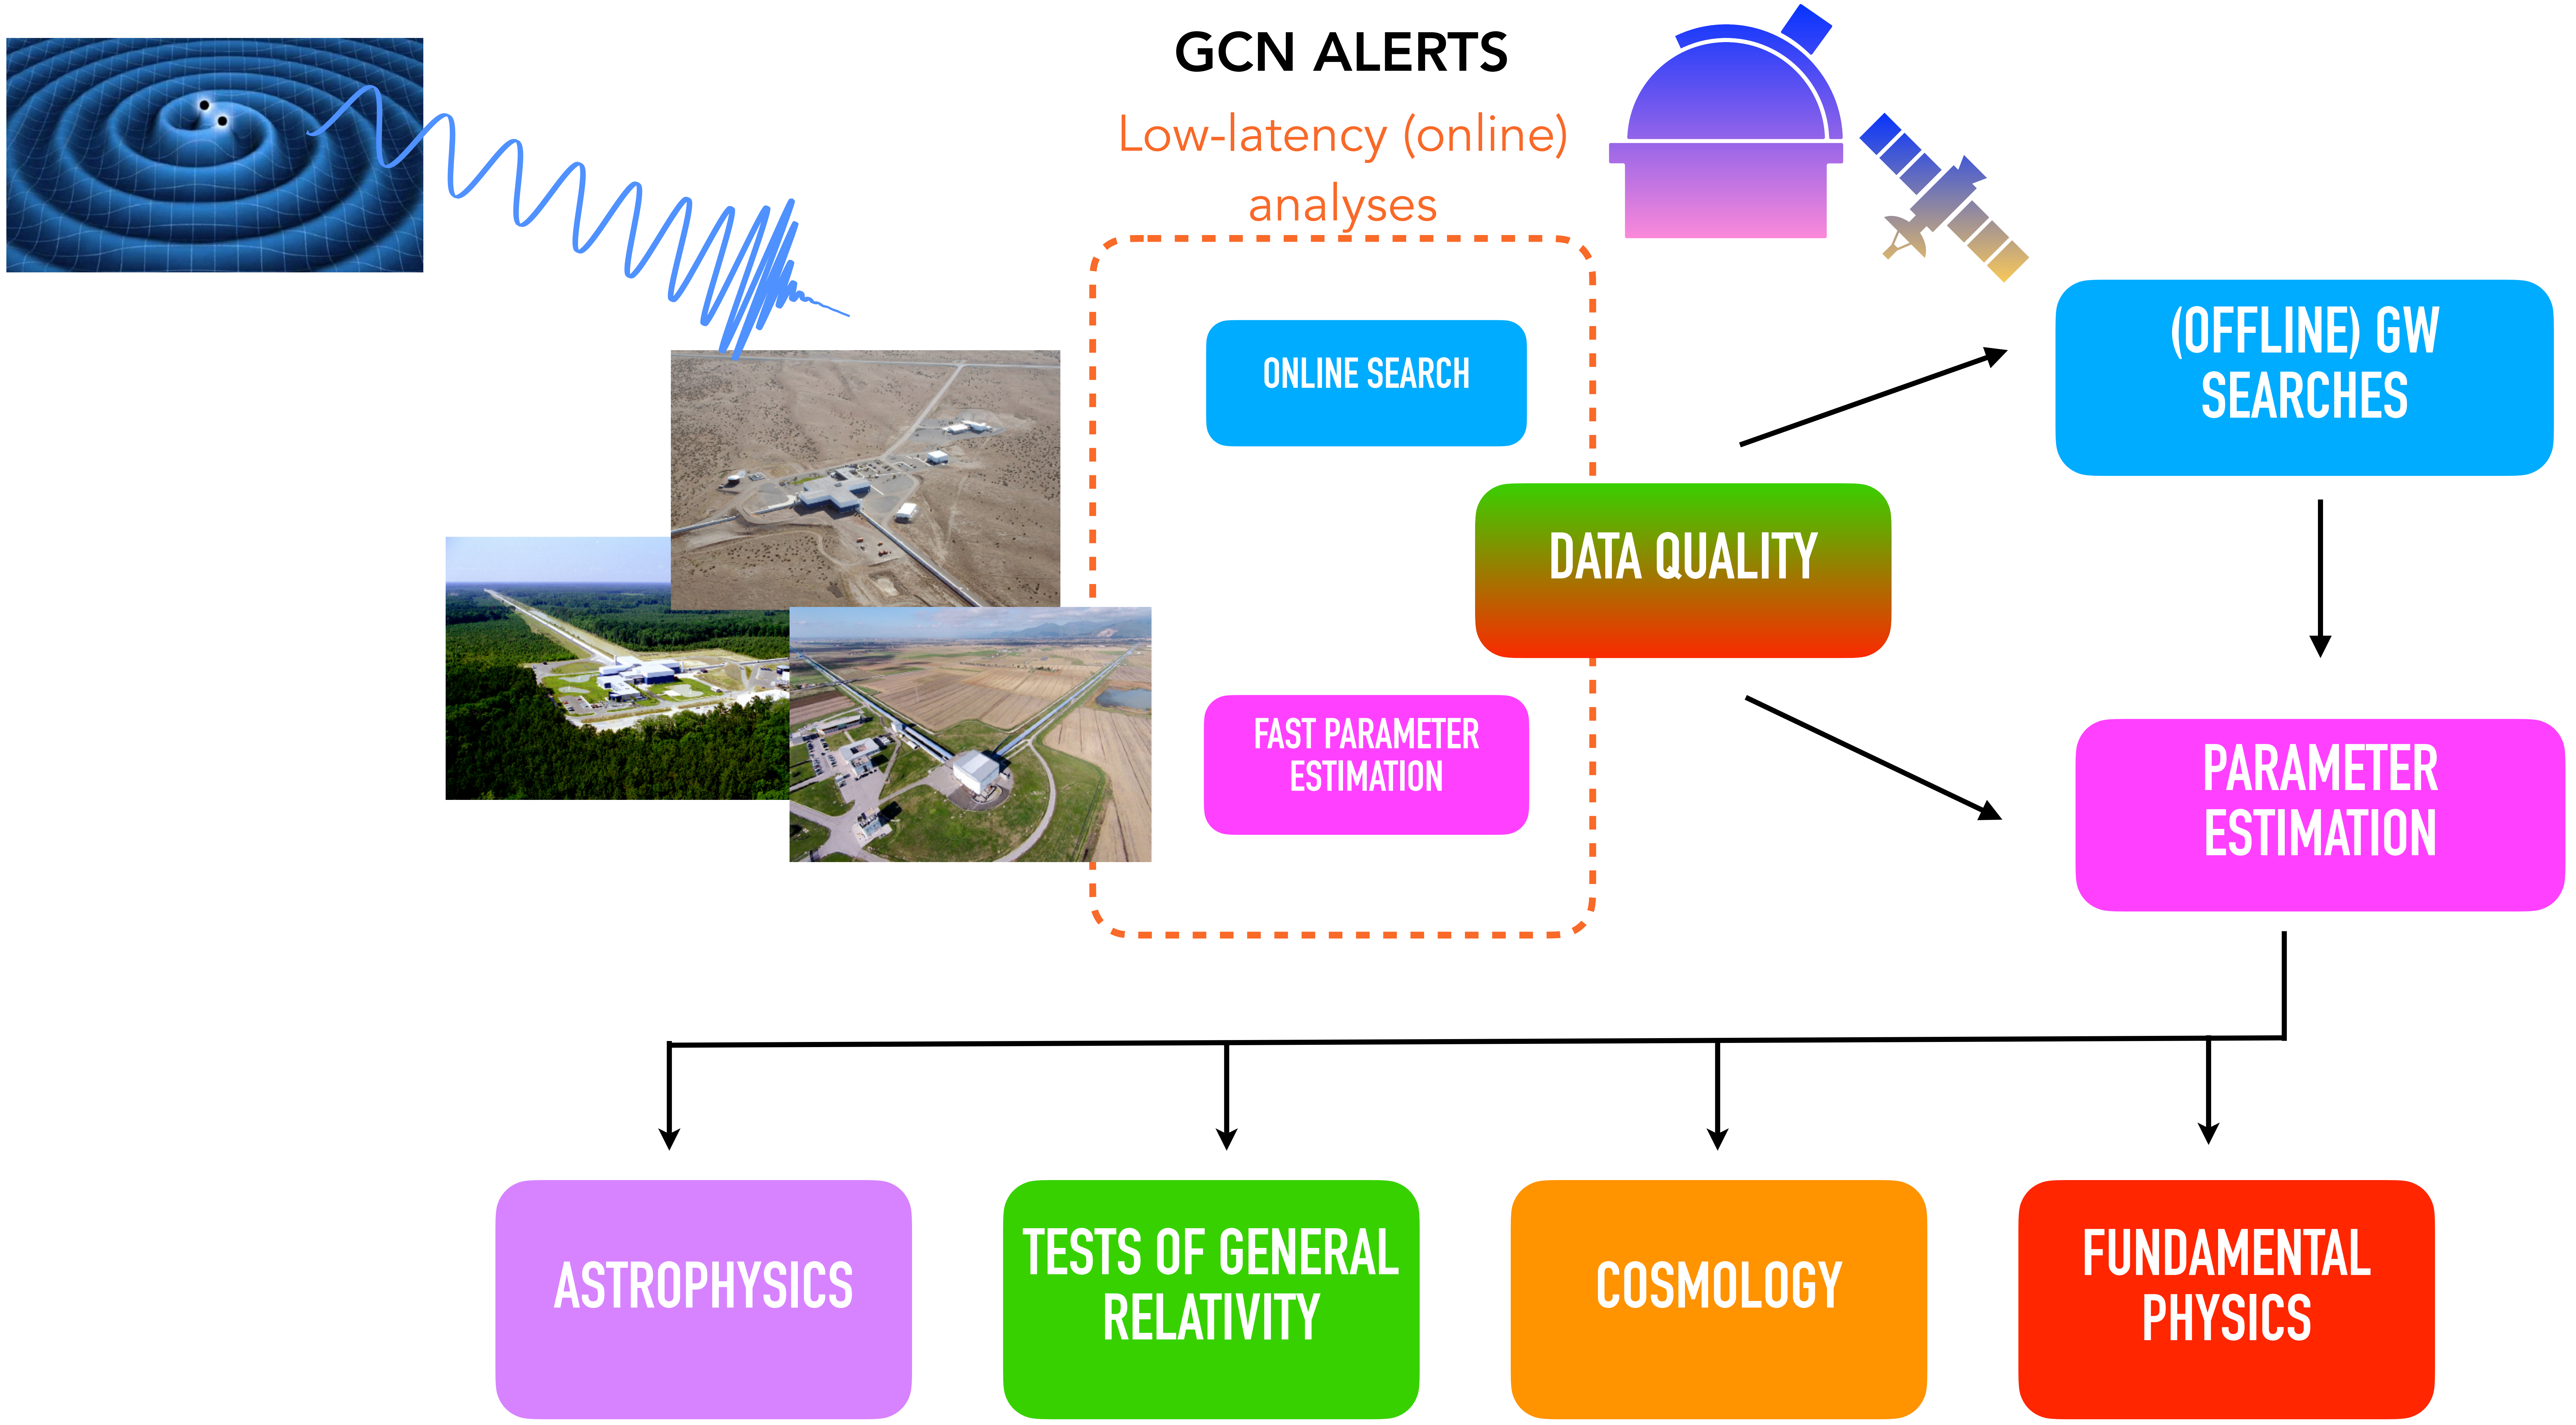

https://gracedb.ligo.org/superevents/public/O4/

DETECTION

CURRENT DETECTOR NETWORK

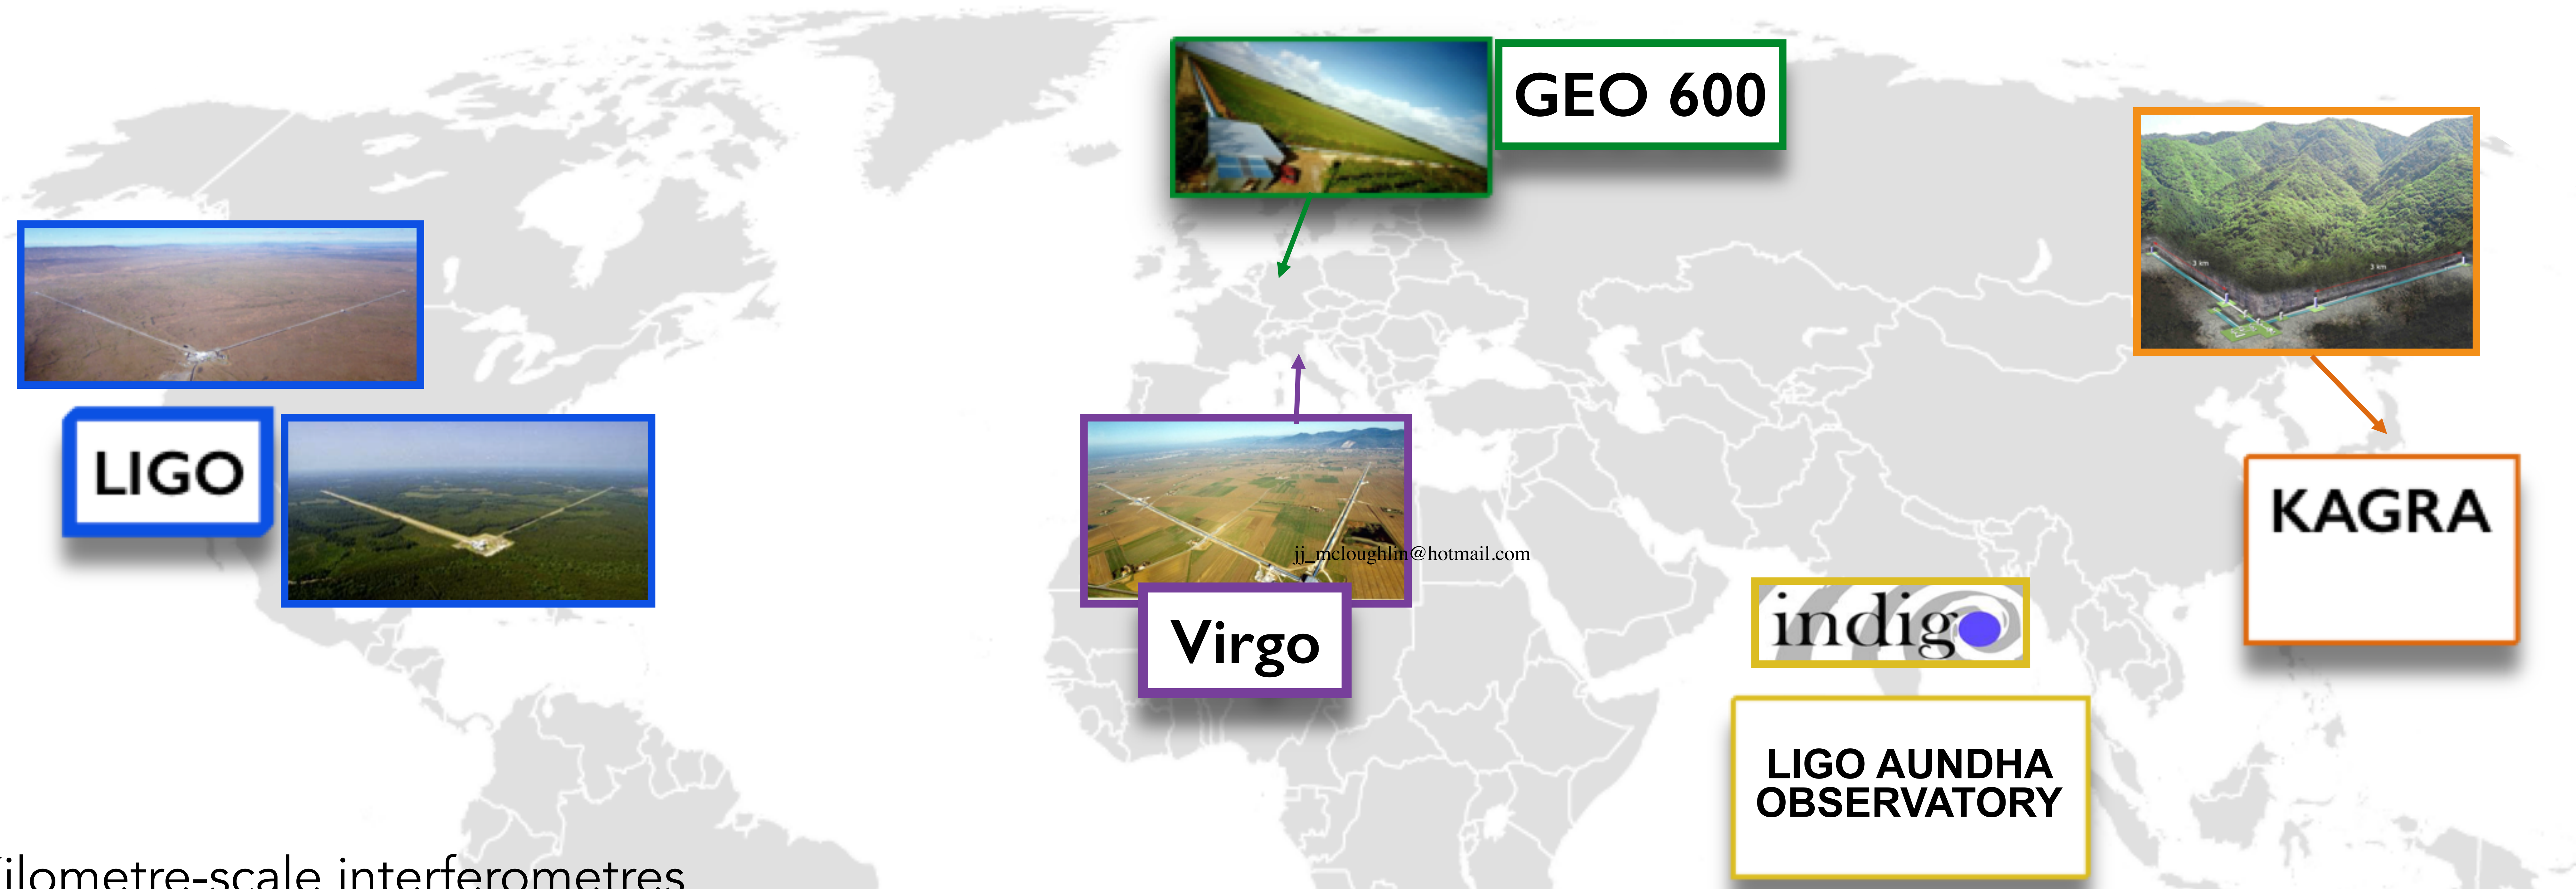

- Kilometre-scale interferometres
- Sensitive to GWs between a few Hz to a few kHz
- Simultaneous detection increases detection confidence
- Improved sky localisation & polarisation measurement
- Increased duty cycle

DETECTION PRINCIPLE

- Gravitational wave incident on a GW detector induces **a (dimensionless) strain**:

$$\frac{\delta L(t)}{L} = F_+(t;\alpha,\delta,\psi)h_+(t;\theta) + F_\times(t;\alpha,\delta,\psi)h_\times(t;\theta) \equiv h(t)$$

  - $(F_+, F_\times)$ are the **antenna patterns**.
  - The antenna response depends on the location and orientation of the source.
- **Detector output** is a real-valued time-series: $d(t) = n(t) + h(t)$

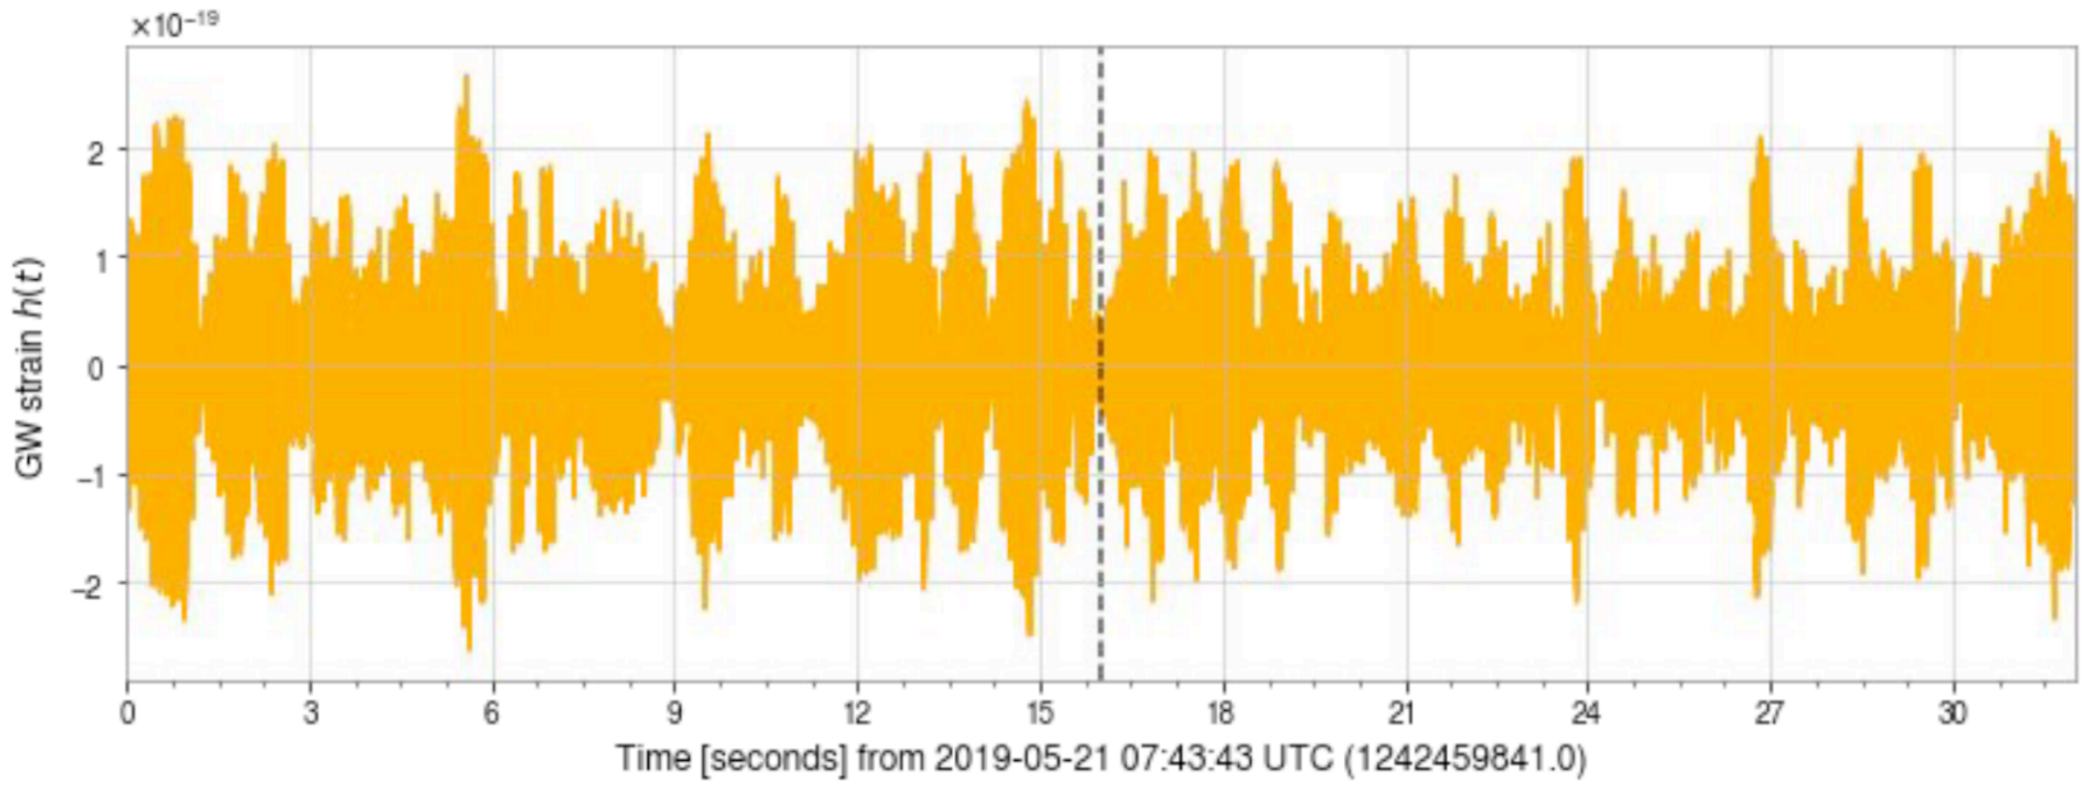

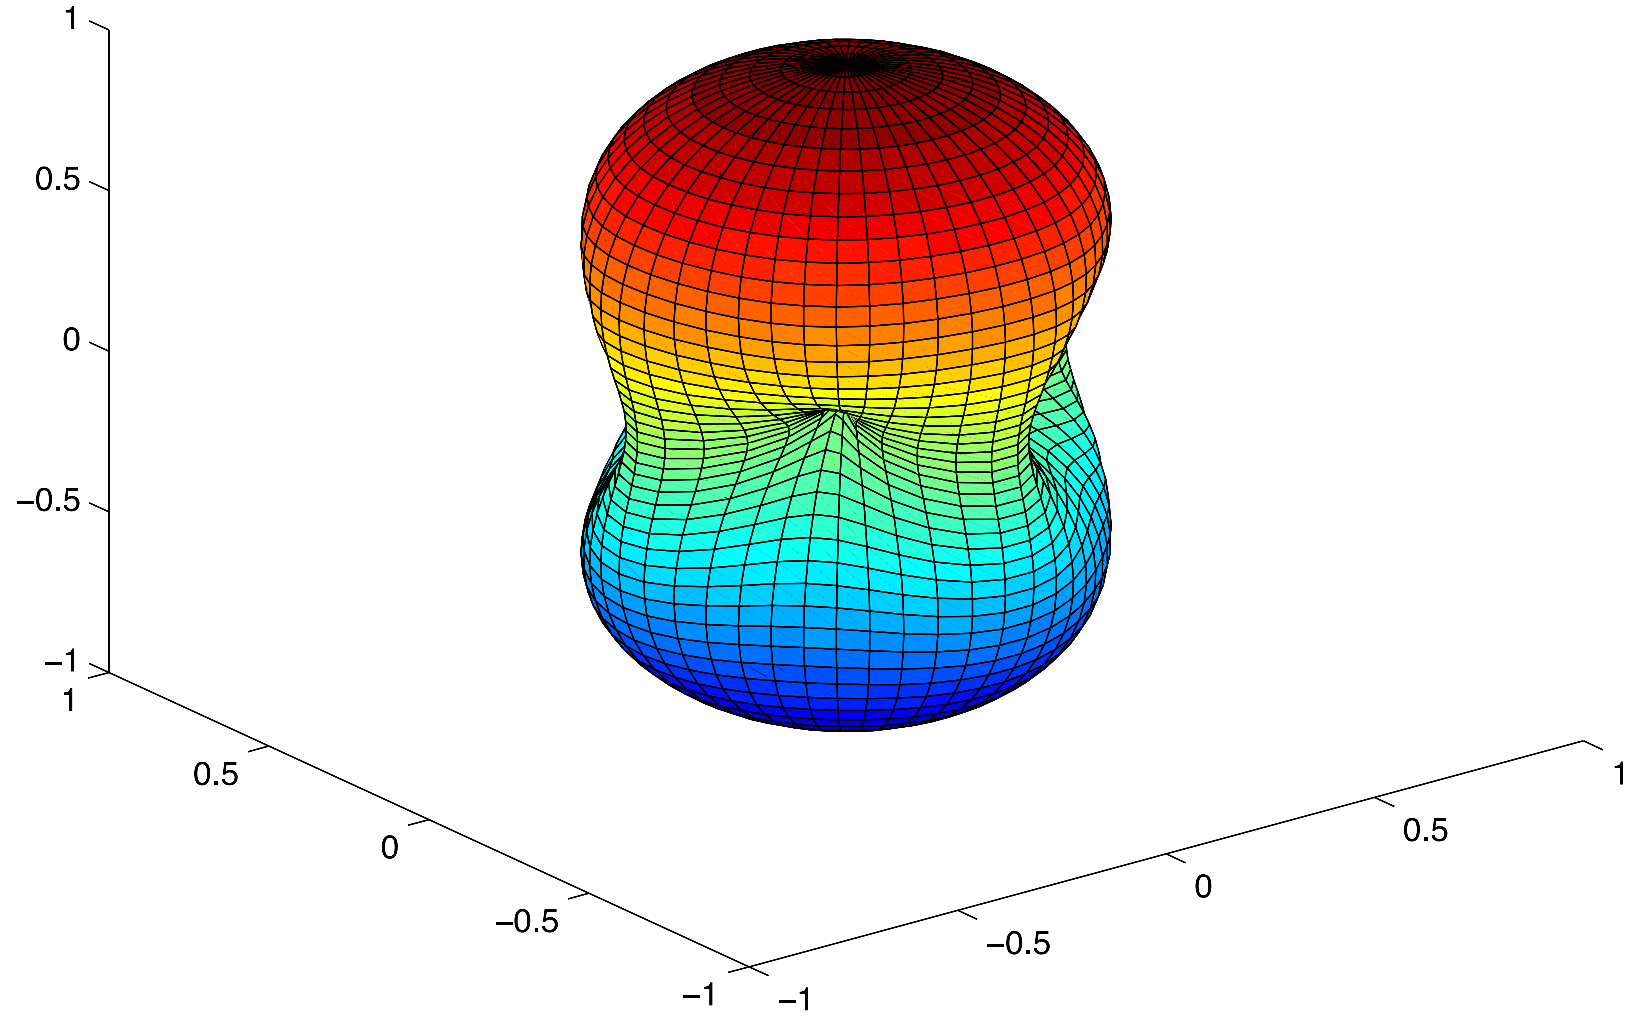

Antenna patterns for an L-shaped detector the xy-plane with arms along the x- and y-axis.

SOME FOURIER CONVENTIONS

- Fourier transform convention:

$$\tilde{x}(f) = \mathcal{F}[x(t)](f) = \int_{-\infty}^{\infty} dt\, x(t) e^{-2\pi i f t}$$

$$x(t) = \mathcal{F}^{-1}[\tilde{x}(f)](t) = \int_{-\infty}^{\infty} df\, \tilde{x}(f) e^{2\pi i f t}$$

- For a real function, e.g. a timeseries $x(t)$, we have:

$$\tilde{x}(-f) = \tilde{x}^*(f)$$

- Time domain translations are equivalent to a Fourier domain phase shift:

$$\text{If } \mathcal{F}[x(\tau)](f) = \tilde{x}(f) \quad \Rightarrow \quad \mathcal{F}[x(\tau - t)](f) = \tilde{x}(f) e^{2\pi i f t}$$

NOISE & SENSITIVITY

- Sensitivity of a GW detector is characterised by the **power spectral density (PSD)** of its background.
- The **noise autocorrelation** is defined as the average over an **ensemble of noise** realisations:

$$\kappa := \langle n(t_1) n(t_2) \rangle$$

ensemble average

- Assuming a **stationary stochastic process**, then κ depends only on $\tau := |t_1 - t_2|$.
  - Then, the **one-sided** (i.e. $f \geq 0$) noise PSD is the Fourier transform of $\kappa(\tau)$ (Wiener-Khinchin theorem):

$$S_n(f) = \frac{1}{2} \int_{-\infty}^{\infty} \kappa(\tau) e^{2\pi i f \tau} d\tau \in \mathbb{R} \qquad \text{with} \qquad S_n(f) = S_n(-f)$$

- **For stationary, Gaussian noise with zero mean:**

$$\langle \tilde{n}(f) \tilde{n}^*(f') \rangle = \frac{1}{2} S_n(f) \delta(f - f')$$

NOISE& SENSITIVITY

- In practice, we work with **one noise realisation.**
  - **Ergodic principle** allows us to replace the ensemble average by a **time average**:

$$\langle X \rangle \to \overline{X}$$

  - Measure $n(t)$ for a sufficiently long duration T.
  - Compute the Fourier transform with $\Delta f = T^{-1}$ frequency resolution.
  - Repeat many times and average.

- **Mean square amplitude of the noise** is then given by:

$$\overline{|n(t)|^2} = \lim_{T\to\infty} \frac{1}{2T} \int_{-T}^{T} n(t) n^*(t)\, dt = \ldots = \int_0^\infty S_n(f) df$$

- The dimension of the PSD is inverse frequency (time), i.e., [PSD] = Hz^{-1}

NOISE & SENSITIVITY

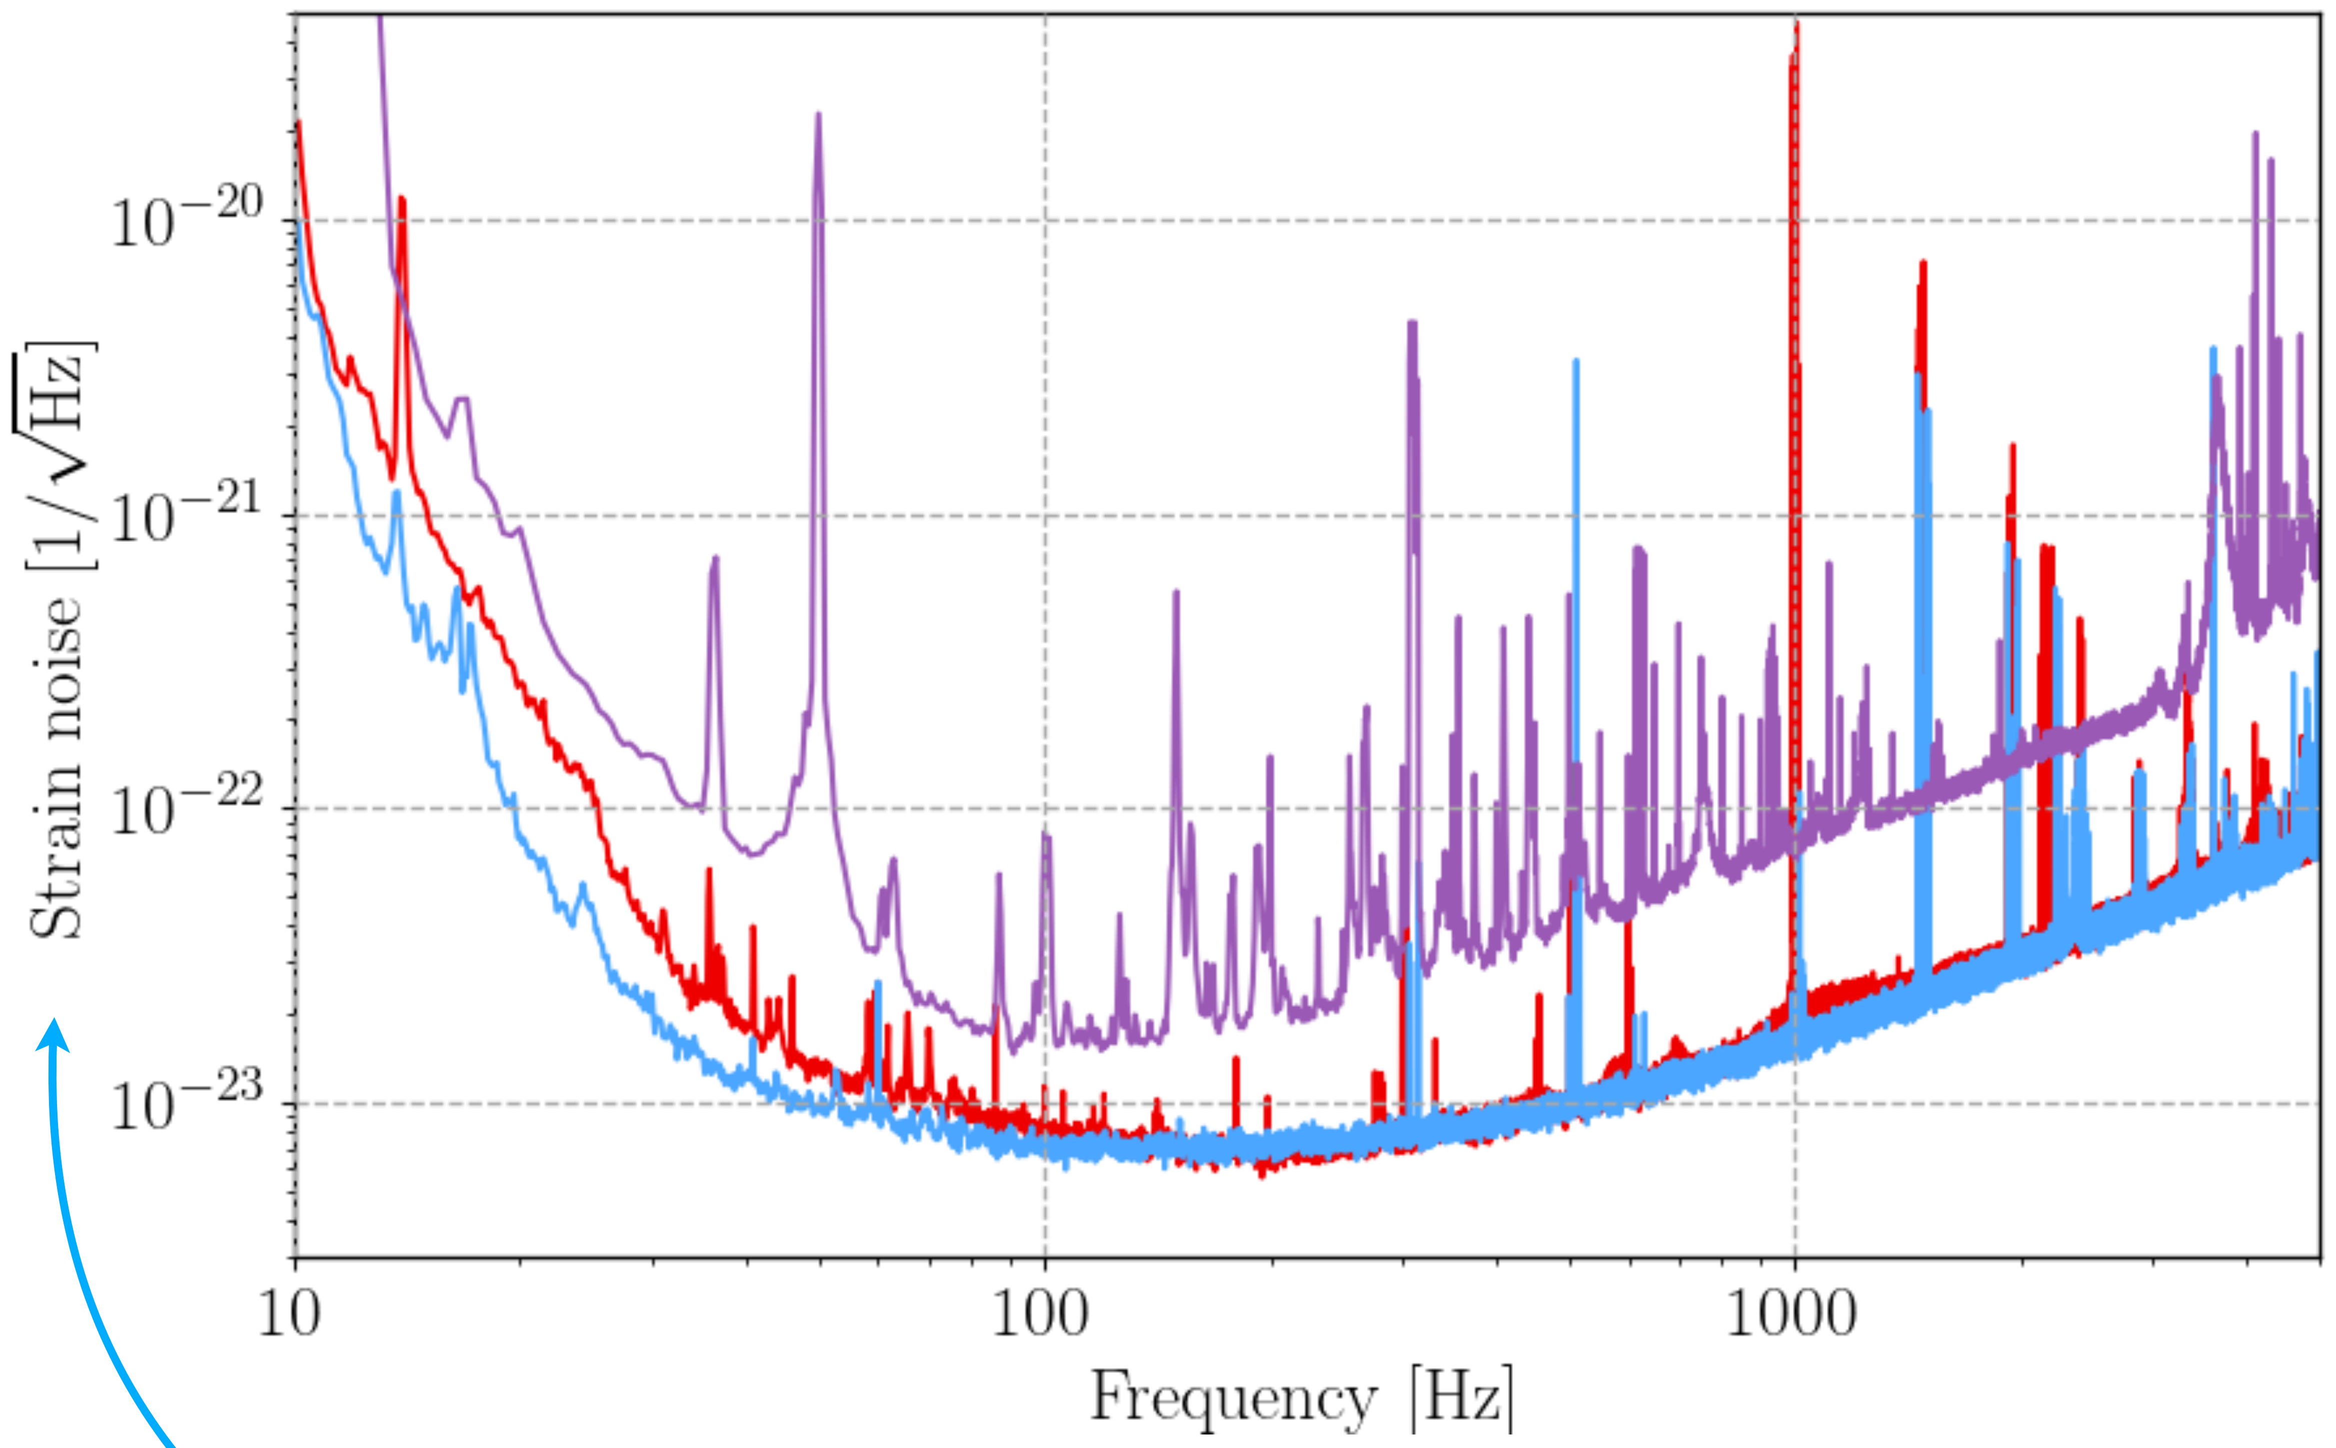

$$\mathrm{ASD} = \sqrt{\mathrm{PSD}}$$

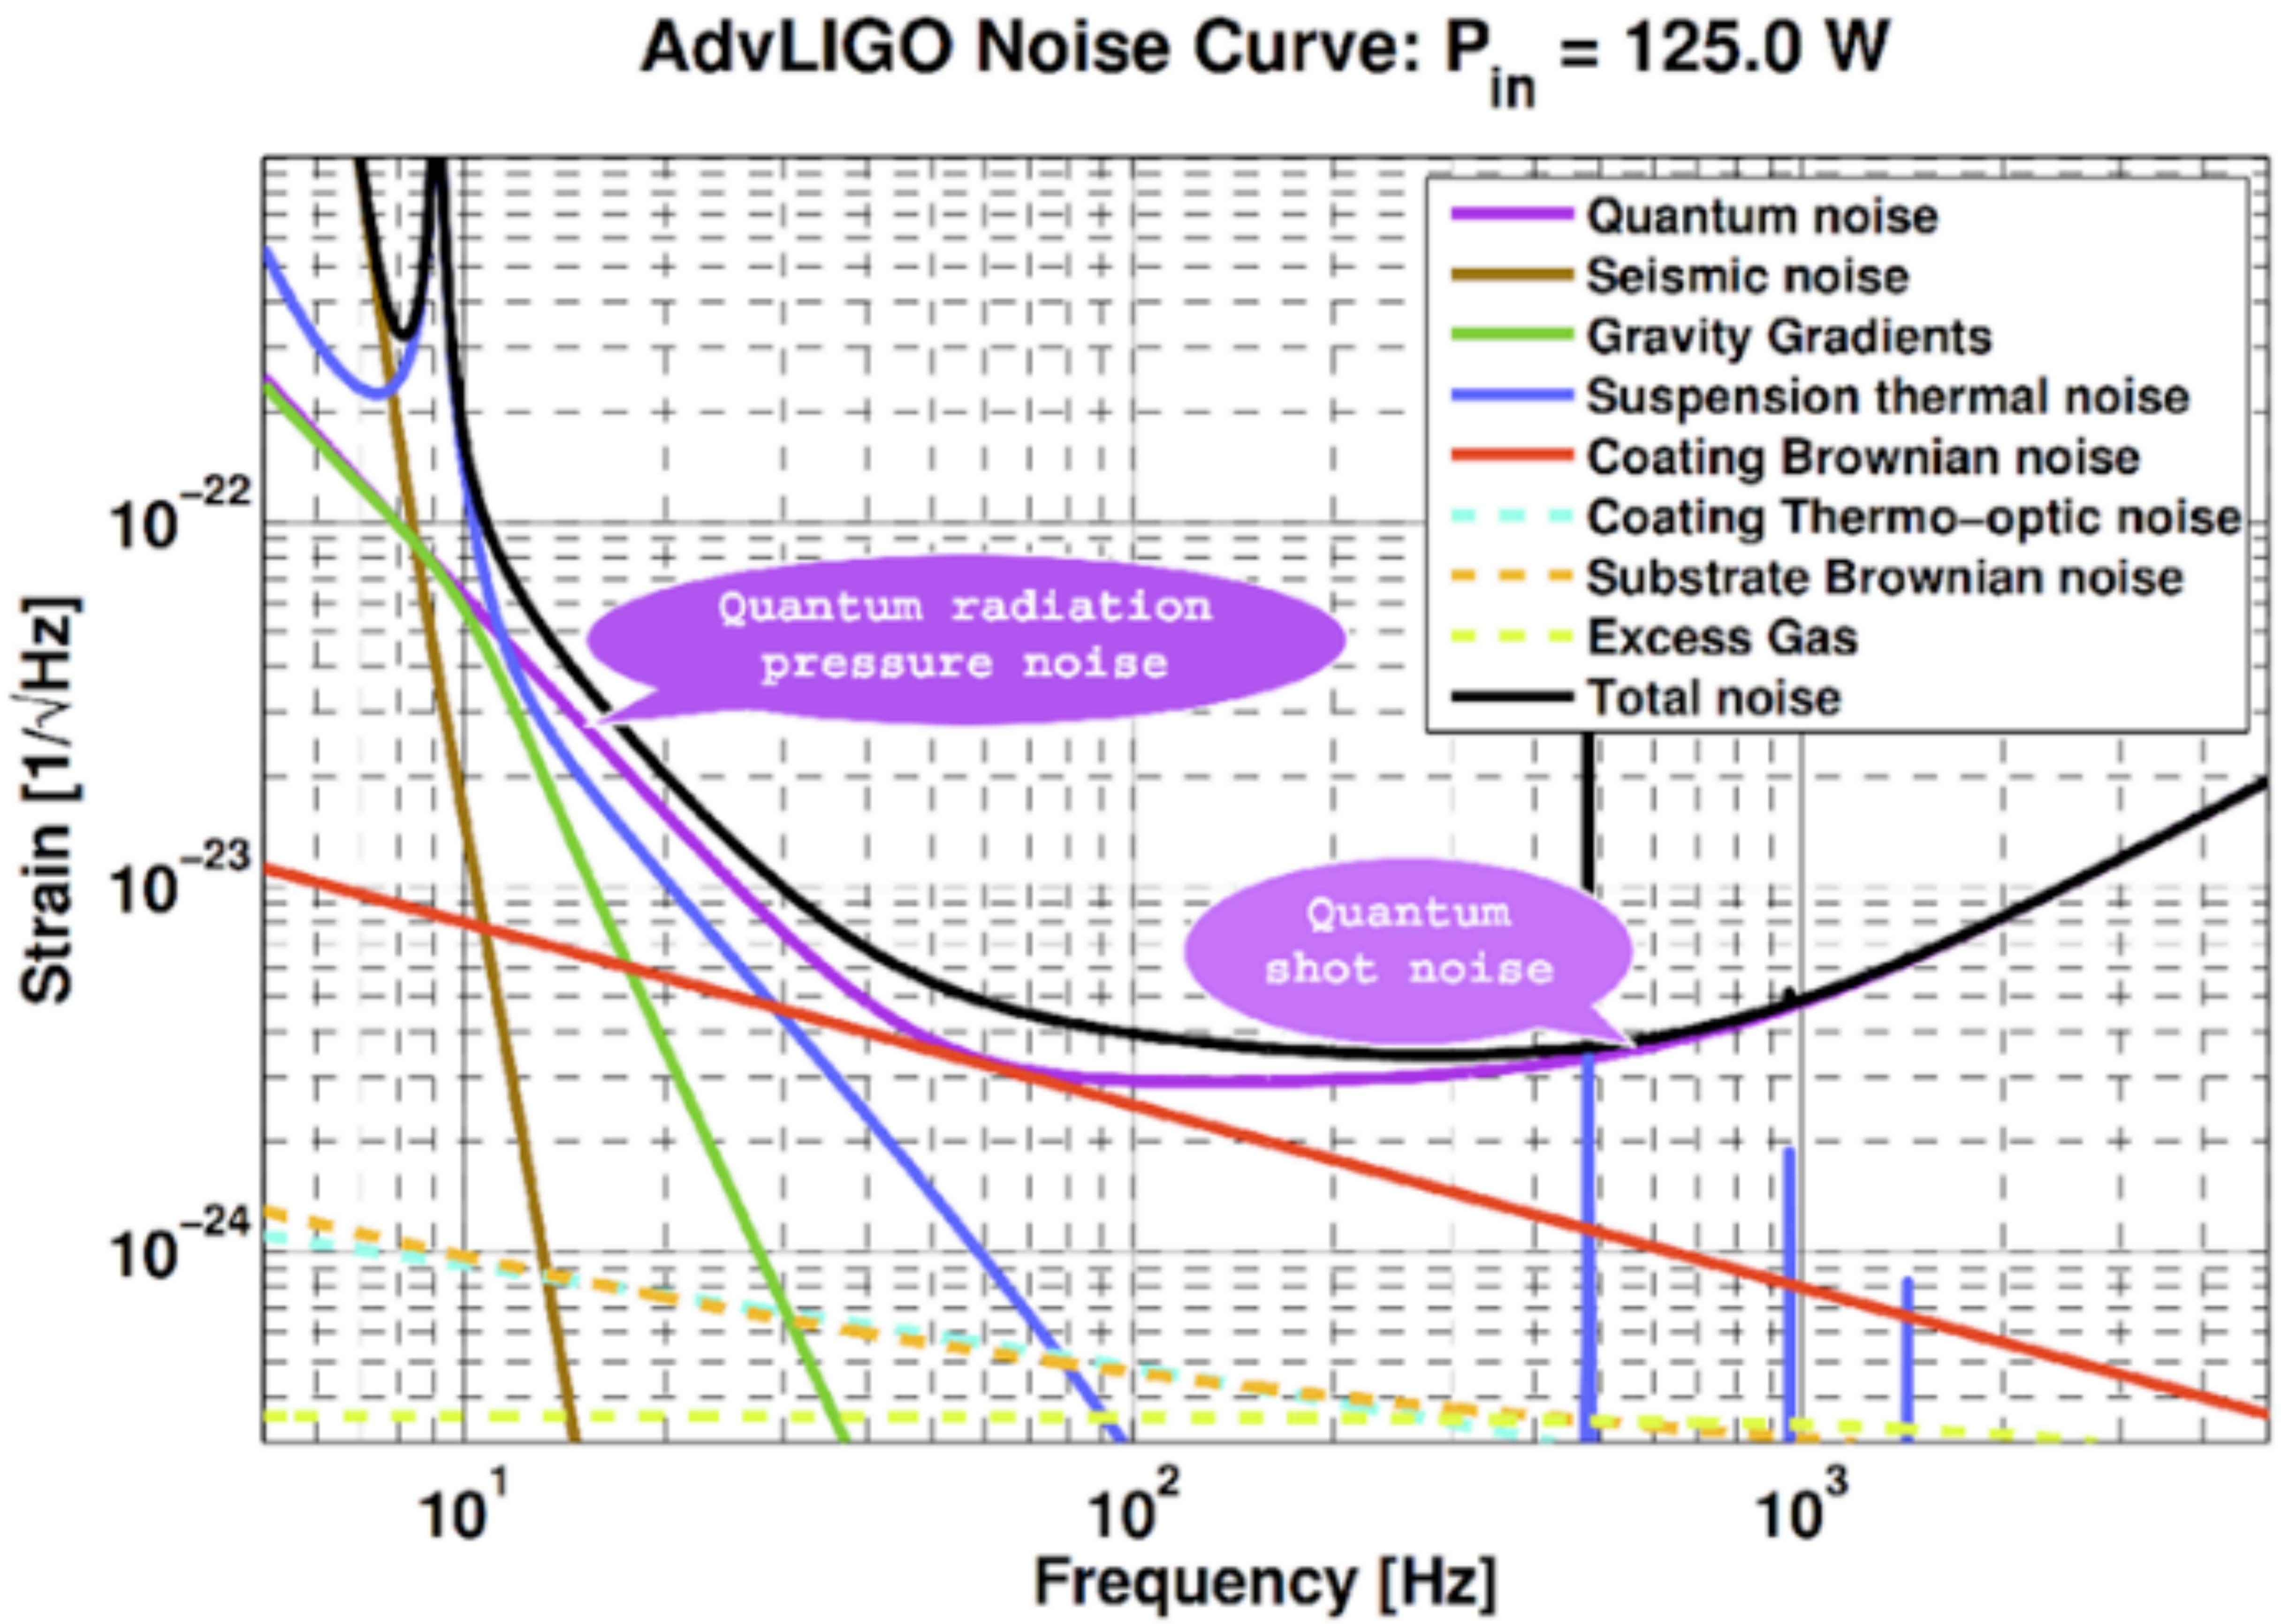

MEASURING THE NOISE IN PRACTICE

- **Welch's method:** obtains a point estimate of the PSD of a noisy time series
  - Split time series into N segments (chunks) of the same duration with a certain overlap (typically 50%)
  - Multiply each segment by a window function to reduce spectral leakage and avoid discontinuities
  - Compute the periodogram (= distribution of power per frequency bin) for each segment
  - Average over all segments to reduce the variance

- NB: Longer segments give a better frequency resolution but higher variance (lower N); higher overlap increases N.

- Commonly used in GW searches and online (fast) analyses.
- Drawbacks:
  - Assumes stationarity over the entire interval (typically 32 - 128s)
  - Computed ``off-source'', i.e. at times away from a signal
  - Windowing, averaging and finite-sample size can induce **biases** in the PSD.

EXTRACTING A KNOWN SIGNAL FROM THE NOISE

- For a known (modelled) signal, the optimal extraction strategy is the construction of a **Wiener filter $K(t)$.**
  - $K(t)$ is real and linear with Fourier transform $\tilde{K}(f)$.

- **Convolution** of the data with the filter gives:

$$(d * K)(\tau) = \int_{-\infty}^{\infty} (h(t) + n(t))K(t - \tau)\, dt \approx \mathcal{S} + \mathcal{N}$$

  - Signal contribution maximised over a time offset:

$$\mathcal{S} = \int_{-\infty}^{\infty} h(t)K(t)\, dt = \int_{-\infty}^{\infty} \tilde{h}(f)\tilde{K}^*(f)\, df$$

  - Noise contribution = mean square of the convolution in the absence of a signal:

$$\mathcal{N}^2 = \int_{-\infty}^{\infty} dt \int_{-\infty}^{\infty} K(t)K(t')\langle n(t)n(t')\rangle\, dt' = \ldots = \frac{1}{2}\int_{-\infty}^{\infty} S_n(f)\tilde{K}(f)\tilde{K}^*(f)\, df$$

EXTRACTING A KNOWN SIGNAL FROM THE NOISE

- The **signal-to-noise ratio (SNR)** is then given by:

$$\rho^2 \equiv \frac{\mathcal{S}^2}{\mathcal{N}^2} = \frac{\left\langle \frac{1}{2} S_n(f) \tilde{K}(f) \middle| \tilde{h}(f) \right\rangle^2}{\left\langle \frac{1}{2} S_n(f) \tilde{K}(f) \middle| \frac{1}{2} S_n(f) \tilde{K}(f) \right\rangle}$$

- With the following **noise-weighted inner product:**

$$\langle \tilde{A}(f) | \tilde{B}(f) \rangle := 4\Re \int_0^\infty \frac{\tilde{A}^*(f) \tilde{B}(f)}{S_n(f)}$$

NB: Defines a **metric** on the space of all possible detector outputs.

- The function $\tilde{K}(f)$ that maximises the SNR is the **optimal filter**. From the Cauchy-Schwarz inequality it follows that

$$\tilde{K}(f) = \frac{\tilde{h}(f)}{S_n(f)}$$

$$\rho^2_{\rm opt} = 4 \int_0^\infty \frac{|\tilde{h}(f)|^2}{S_n(f)} df = \langle \tilde{h}(f) | \tilde{h}(f) \rangle$$

EXTRACTING A KNOWN SIGNAL FROM THE NOISE

- Implementation = **matched filtering**

strain data $d(t)$

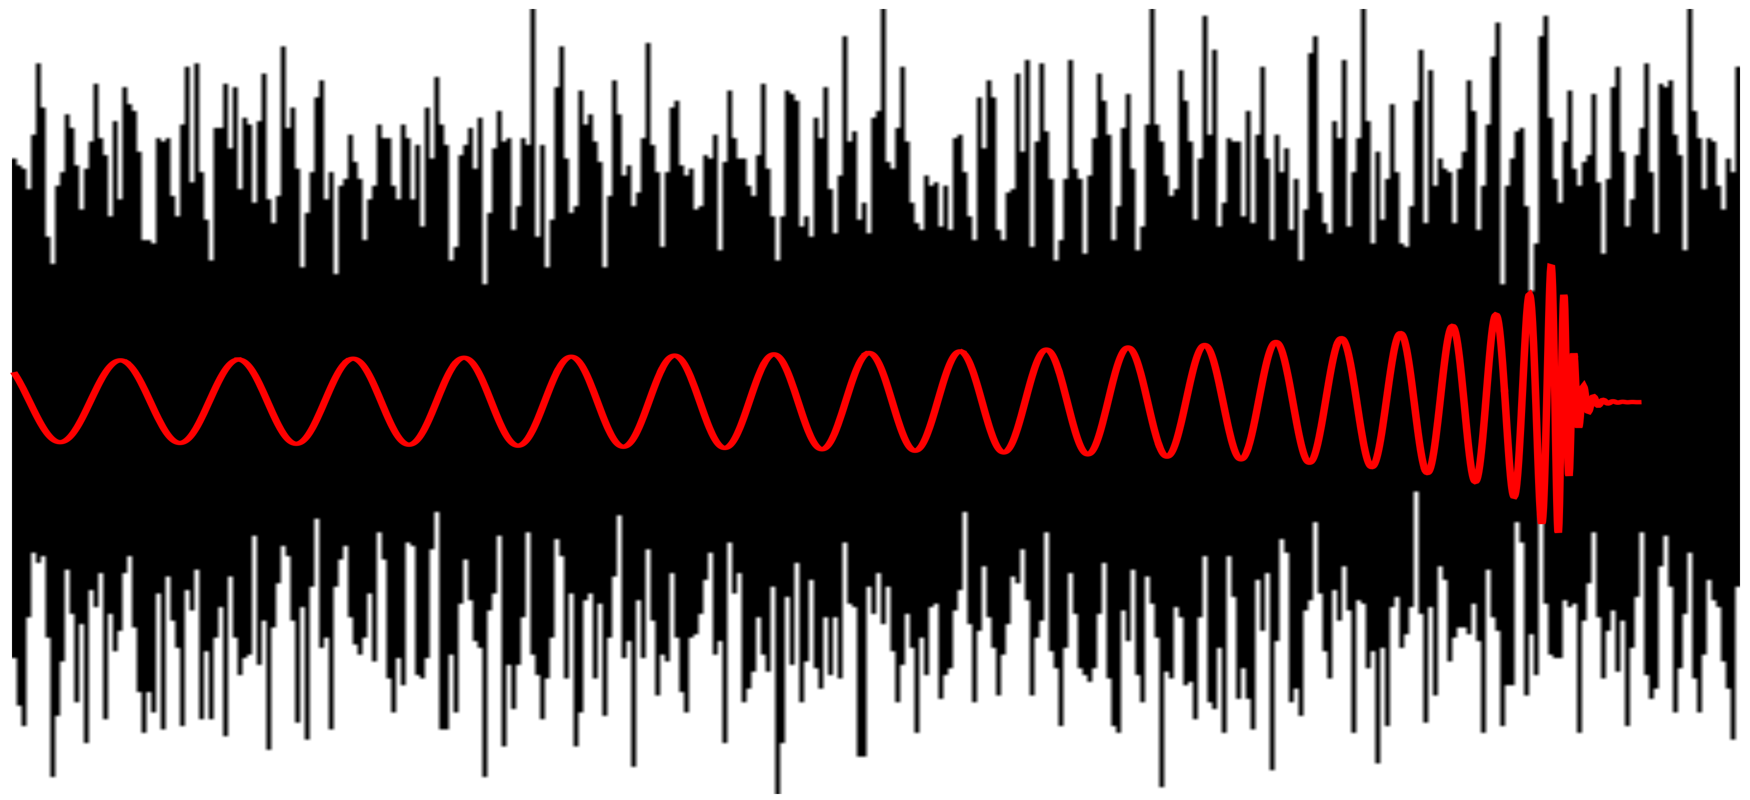

filters: template waveforms $h_i(t)$

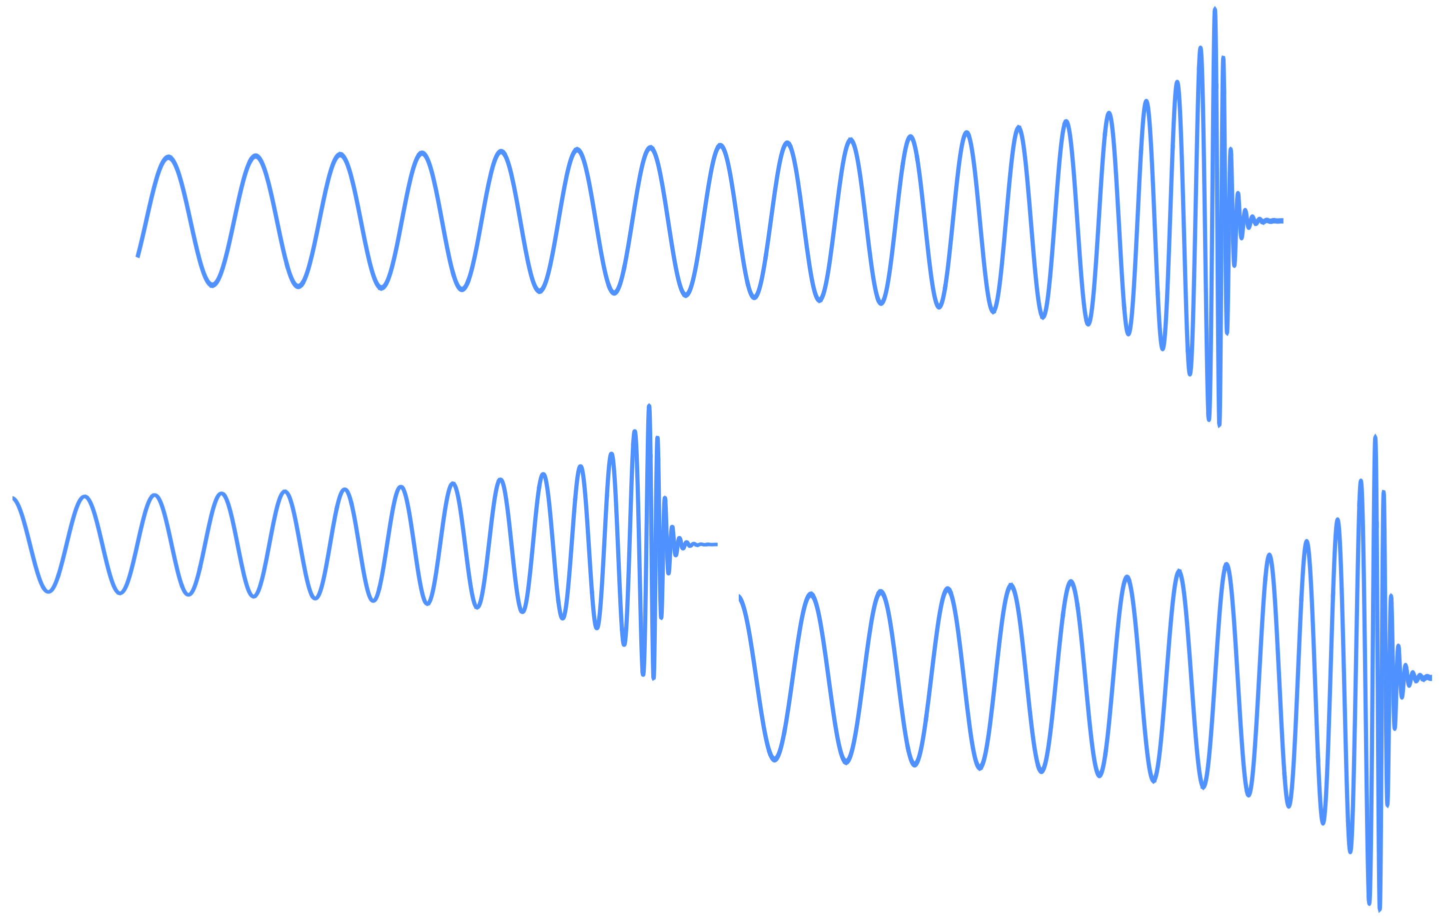

SNR time series

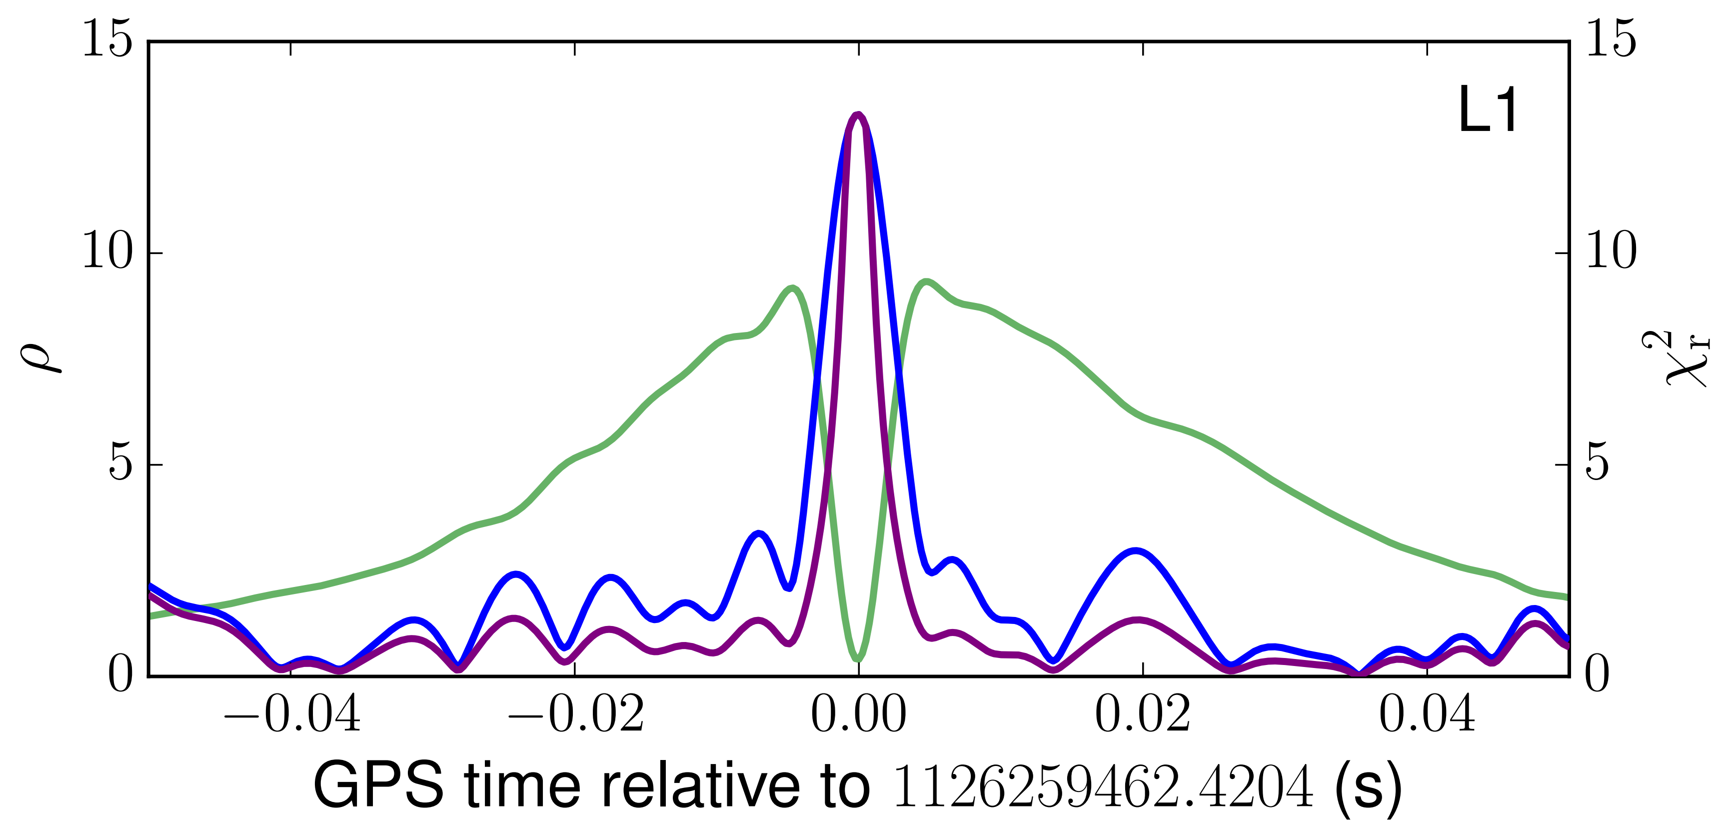

SNR for each template:

$$\rho_i = \frac{\langle d|h_i\rangle}{\sqrt{\langle h_i|h_i\rangle}} \leq \rho_{\rm opt}$$

ρ defines the statistic of the matched filter.

NB: Loss of SNR if your template does not match your signal.

MATCHED FILTERING IN PRACTICE

- **Template bank** = Large collection of precomputed theoretical waveforms
- Construction: hybrid method
  - Geometric lattice for known metric (green area)
    - Metric = mismatch between neighbouring templates
  - Stochastic placement
- What physics is currently included in LVK banks?
  - Component masses
  - Aligned-spins
  - Quadrupolar mode
- Ongoing work to include higher-order modes, eccentricity, precession, etc.

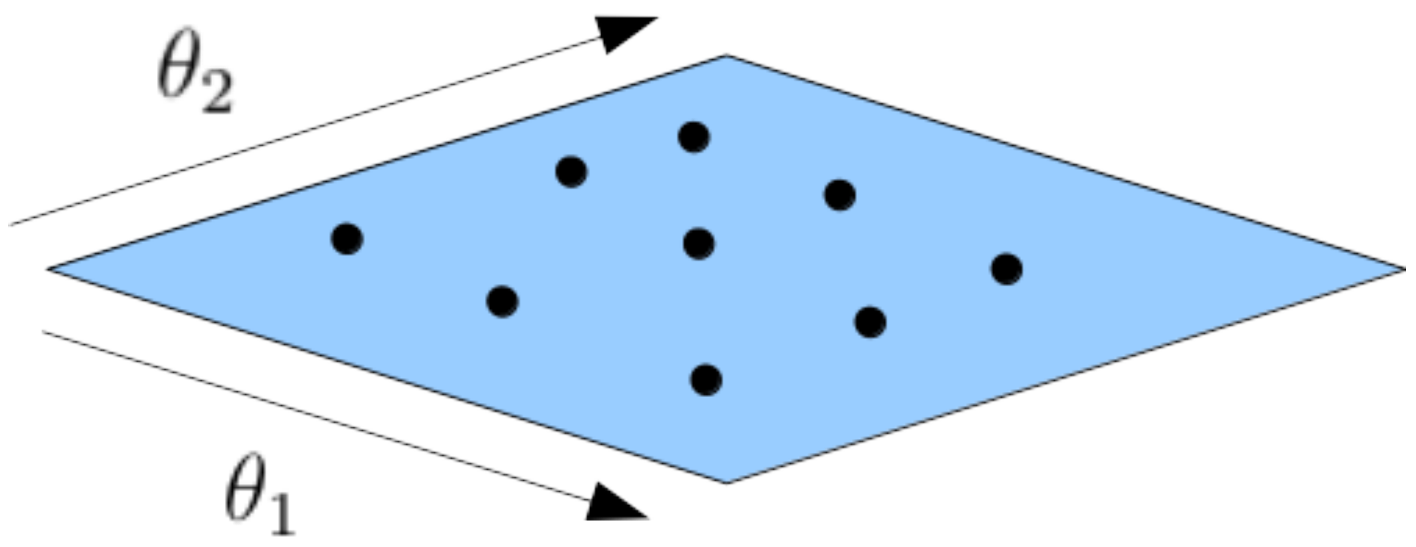

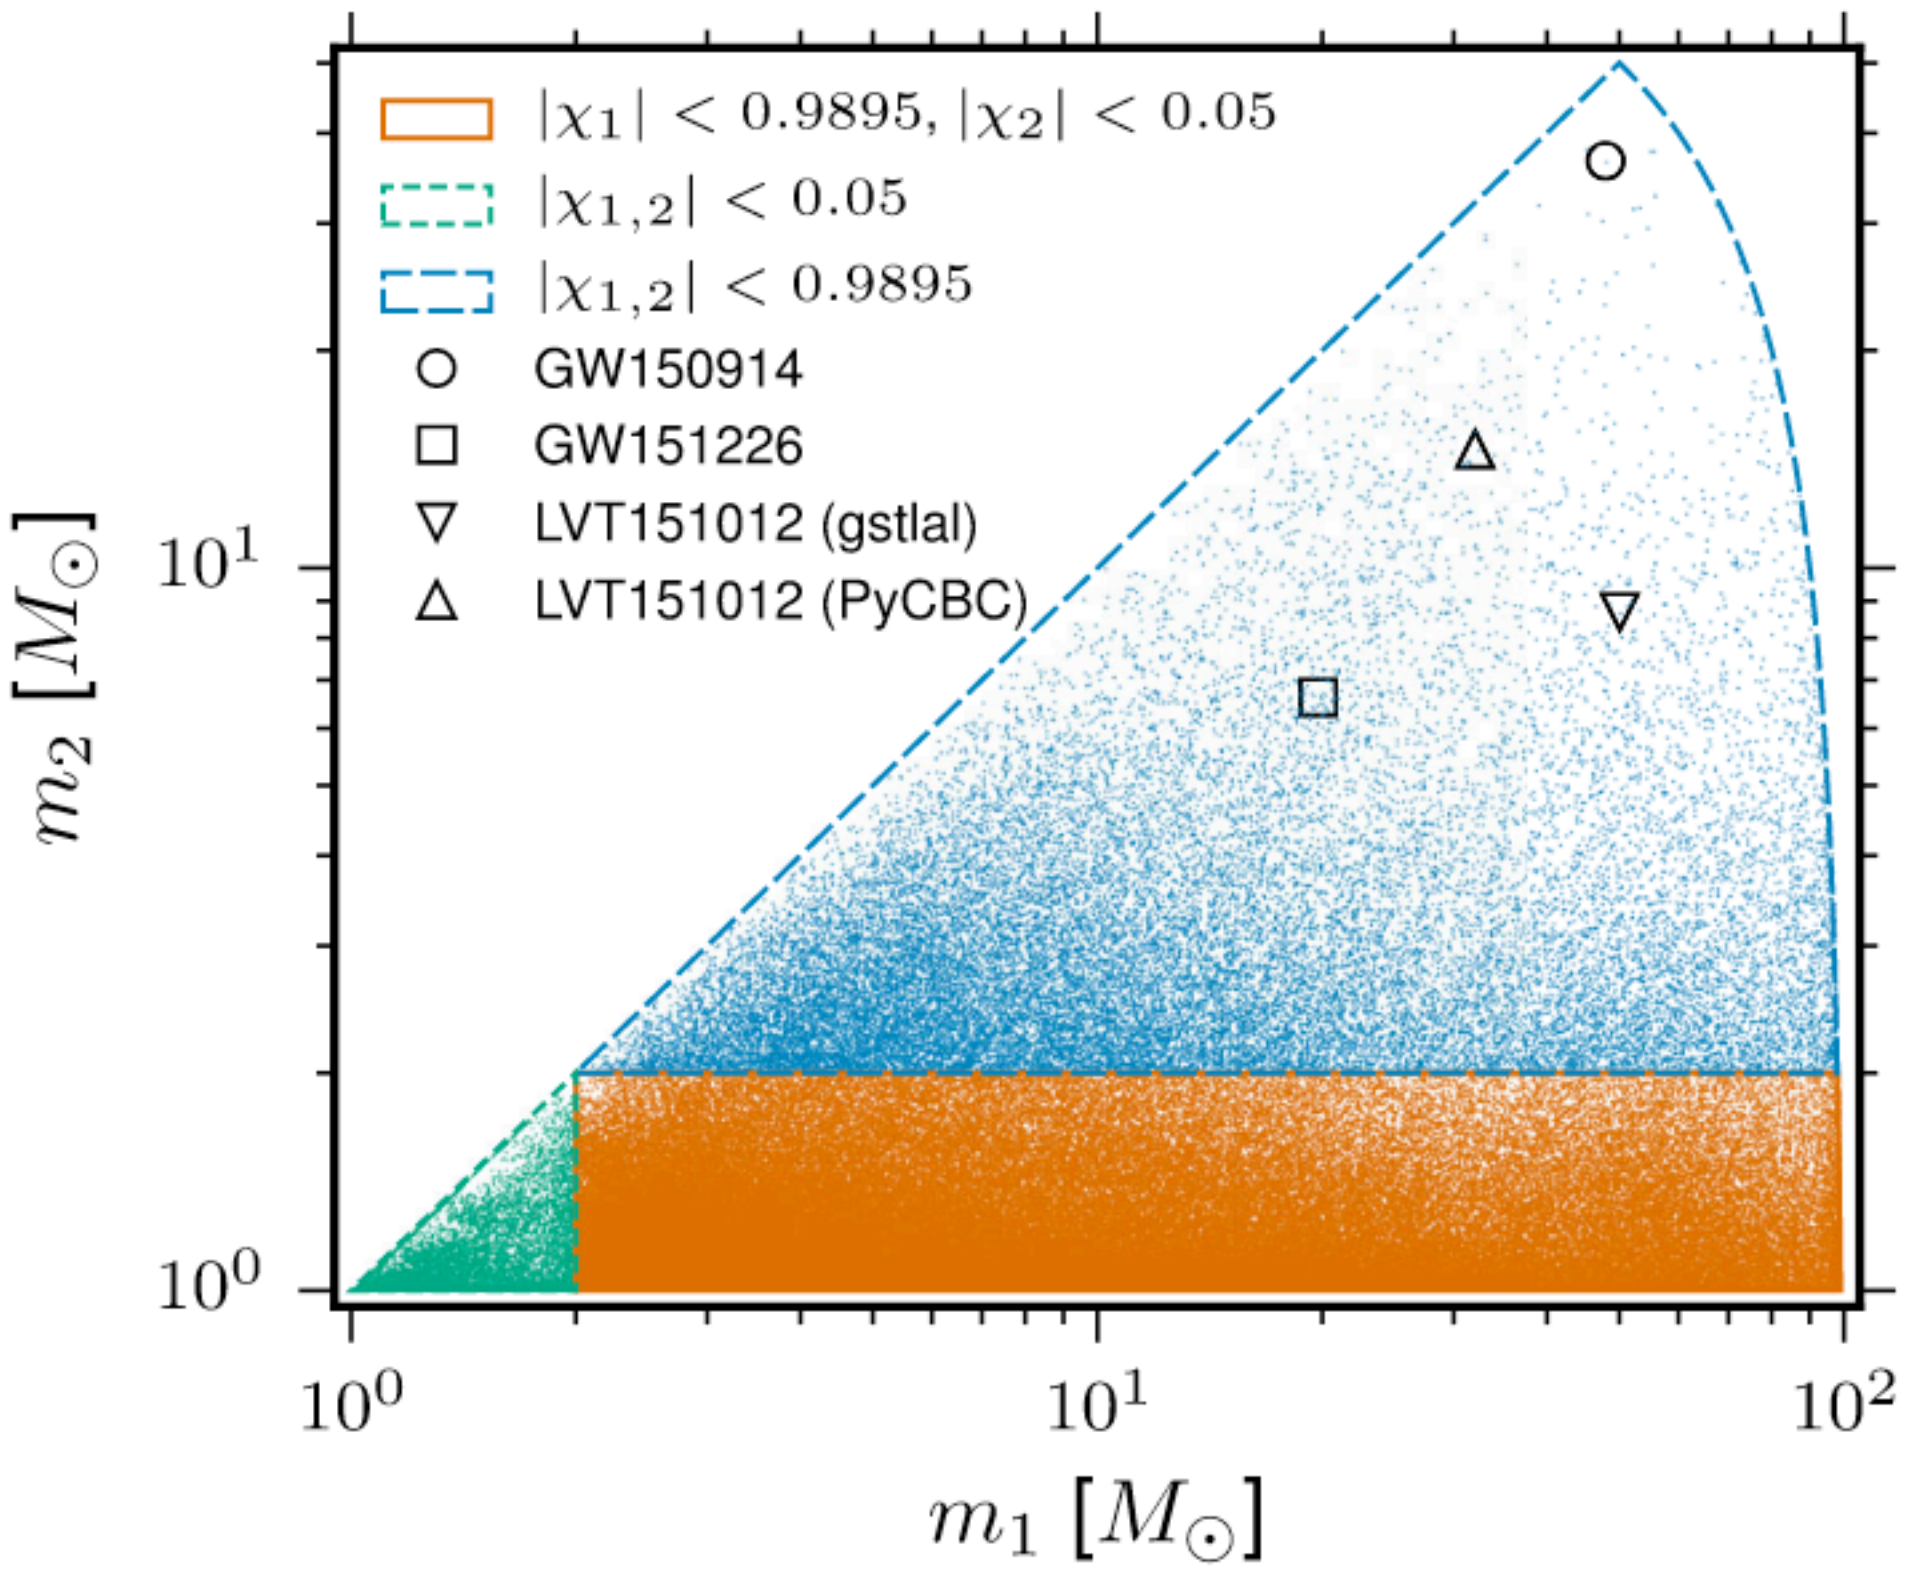

FALSE ALARMS

- If the **noise is Gaussian**, then the probability of observing a noise amplitude in the range $[n, n + dn]$ is

$$p(n)dn = \frac{1}{\sqrt{2\pi}\sigma} e^{-n^2/(2\sigma^2)} dn$$

- Therefore, the probability that the **noise amplitude is above some threshold** η is:

$$p(n|n \geq \eta) = \int_\eta^\infty p(n)dn = \frac{1}{\sqrt{2\pi}\sigma} \int_\eta^\infty e^{-n^2/(2\sigma^2)} dn$$

- Choosing an acceptable number of noise-generated false alarms, e.g. 1 per year, determines the threshold.
    - A high threshold guarantees a purer event list, i.e. the probability of contamination by false alarms is lower.
    - Low-SNR signals are discarded.
    - Signal enhancement techniques allow us to lower the threshold.

GAUSSIANITY

- Real data are neither Gaussian nor stationary!
    - Non-Gaussian transients occur all the time
        - O3: 24% of GW candidates occurred near "glitches".
- Glitches can mimic GW signals
- What to do?
    - Data cleaning
    - Glitch removal/mitigation
    - Use a more sophistic statistic than just the SNR
    - Modify the likelihood

Citizen science: Gravity Spy

https://www.zooniverse.org/projects/zooniverse/gravity-spy

Some example **"glitches"**

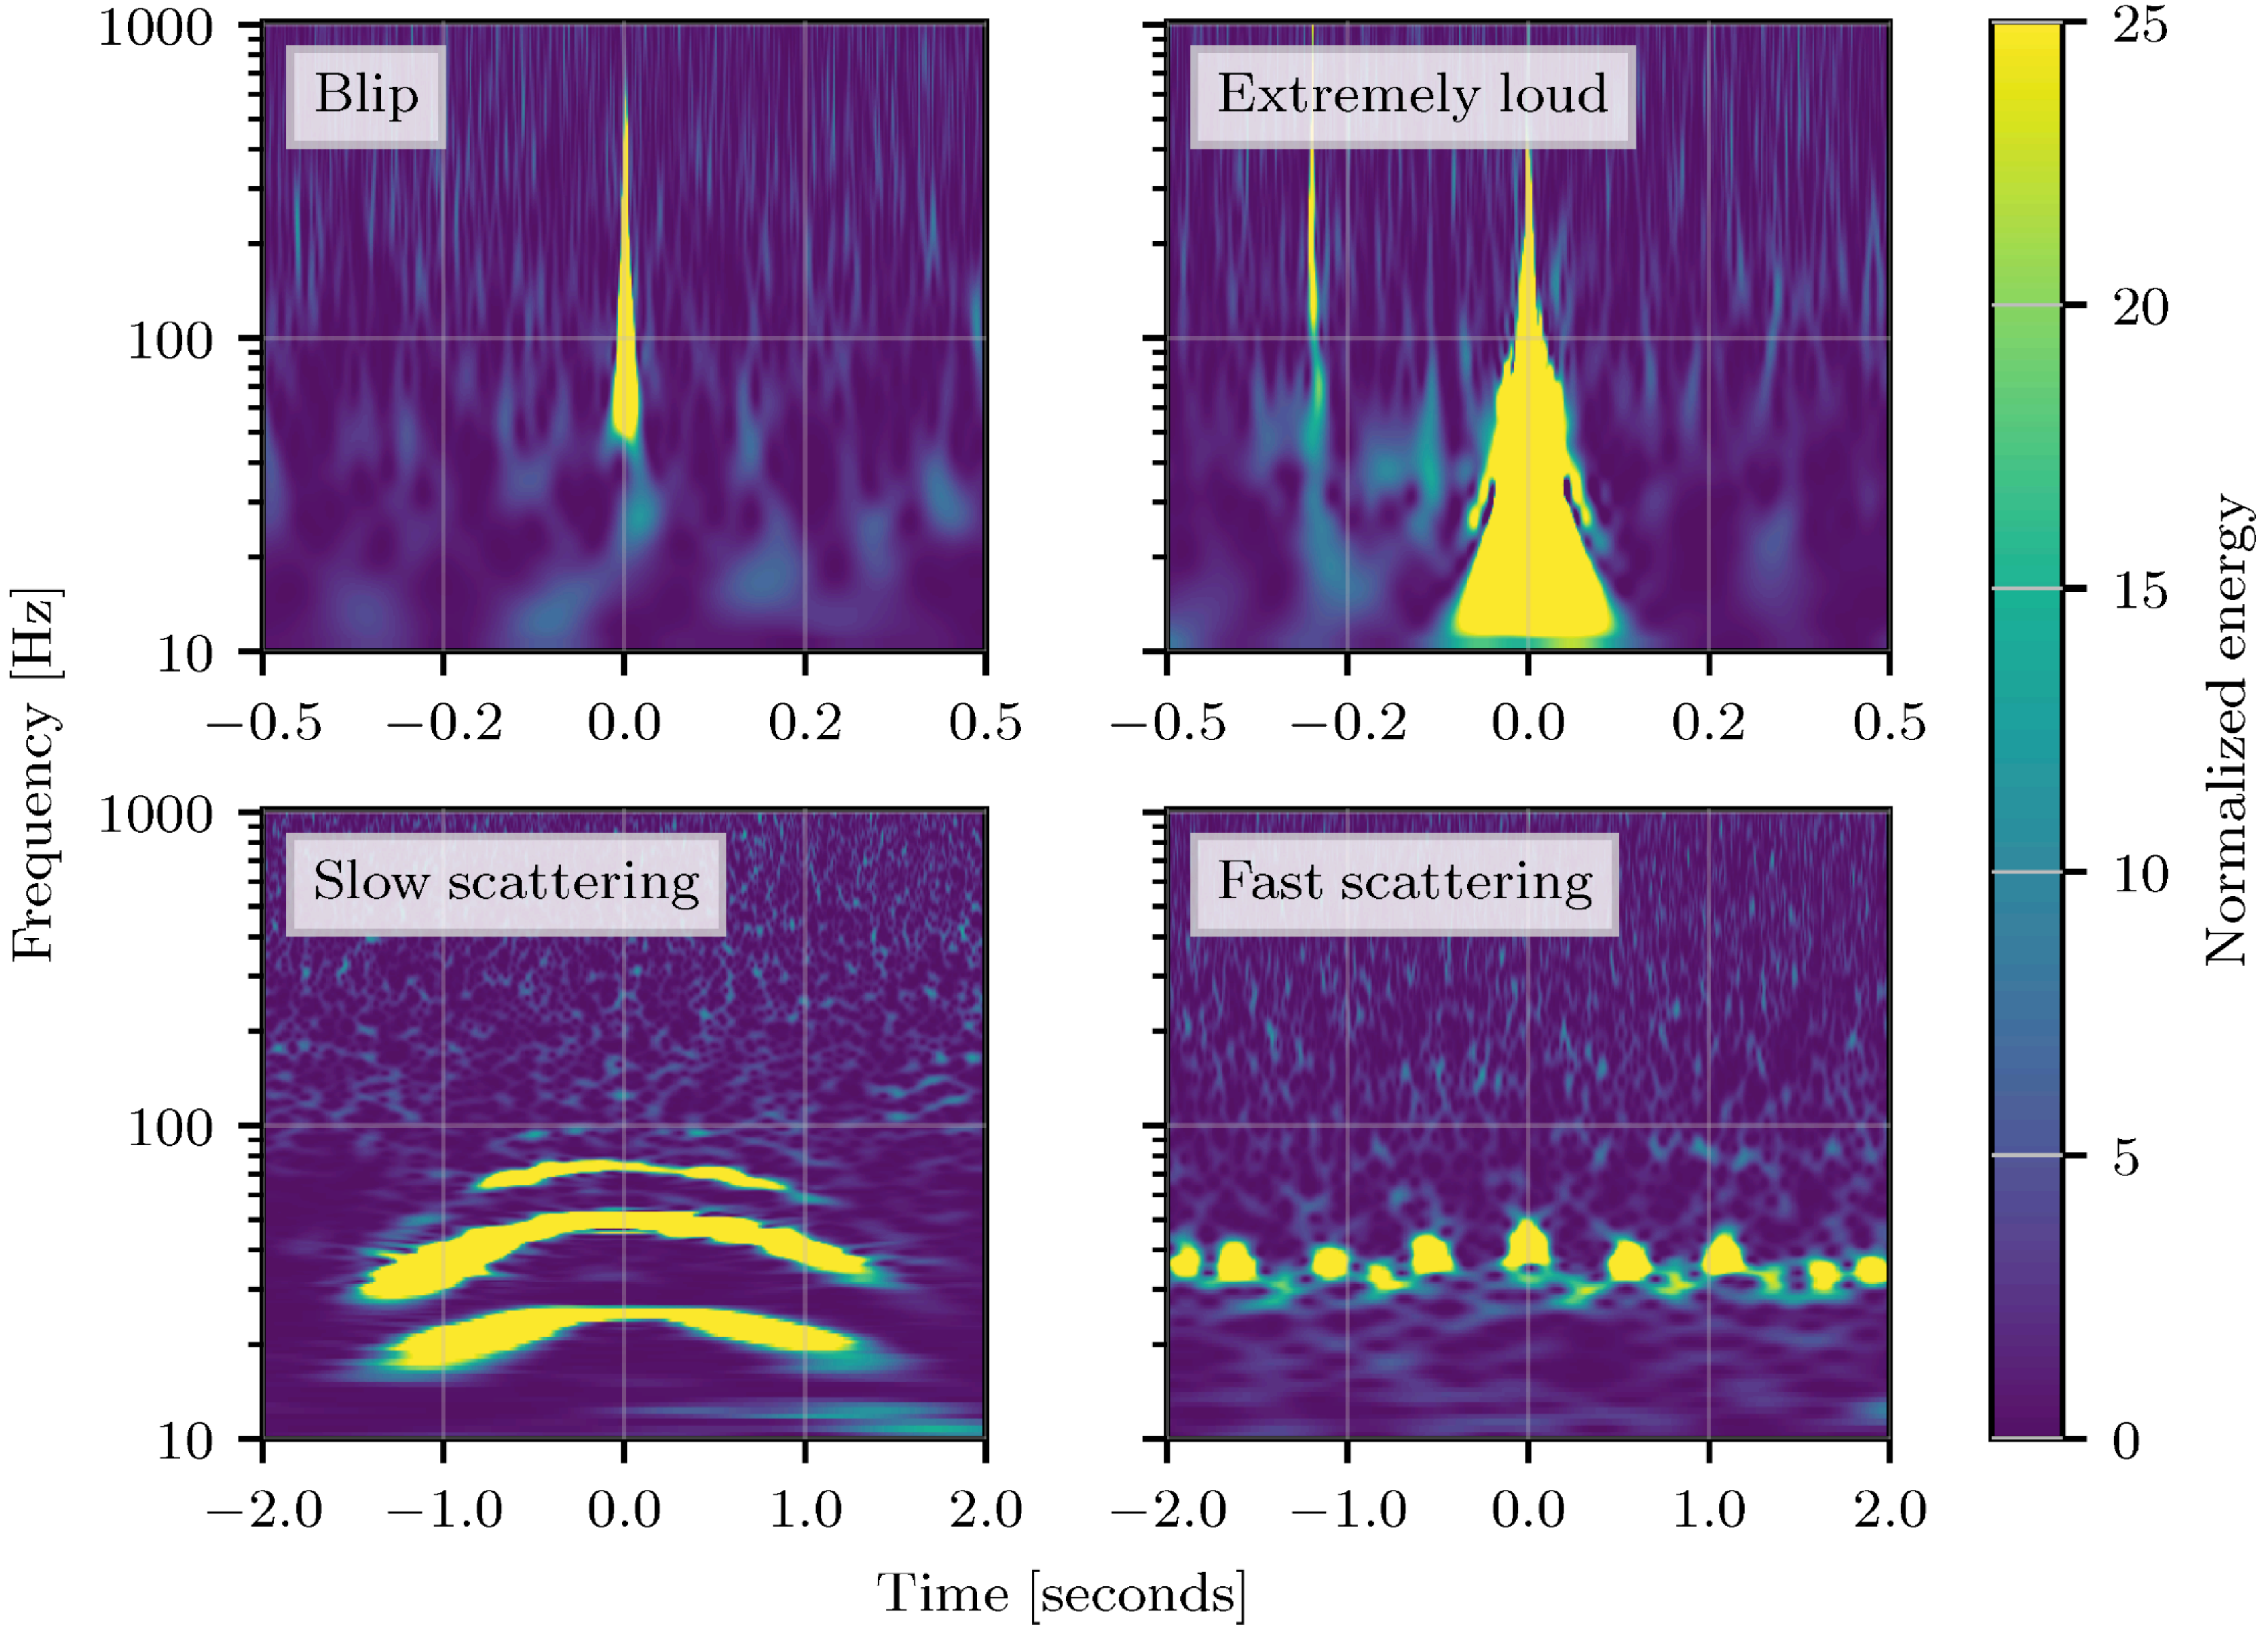

THE χ^2-VETO

- The matched filter statistic collapses the frequency evolution into a single number, the SNR ρ.
  - BUT: Signals from compact binaries are broadband and have a characteristic frequency evolution.

- To take this into account, split the signal into N frequency bins such that the SNR accumulated in each bin is the same:

$$\rho = \sum_{k=1}^{N} \rho_k$$

- We can now construct a new statistic based on the value in each bin compared to the expected value:

$$\chi^2 = N \sum_{k=1}^{N} \left(\rho_k - \frac{\rho}{N} \right)^2$$

- If the background noise is Gaussian, we will recover the standard χ^2-distribution with N-1 degrees of freedom.

SIGNIFICANCE

- Significance quantifies how unlikely it is that a signal arose from noise alone.
- Quantified via the **false alarm rate (FAR)** or equivalently the false alarm probability:

FAR = number of noise triggers above a threshold/total background time

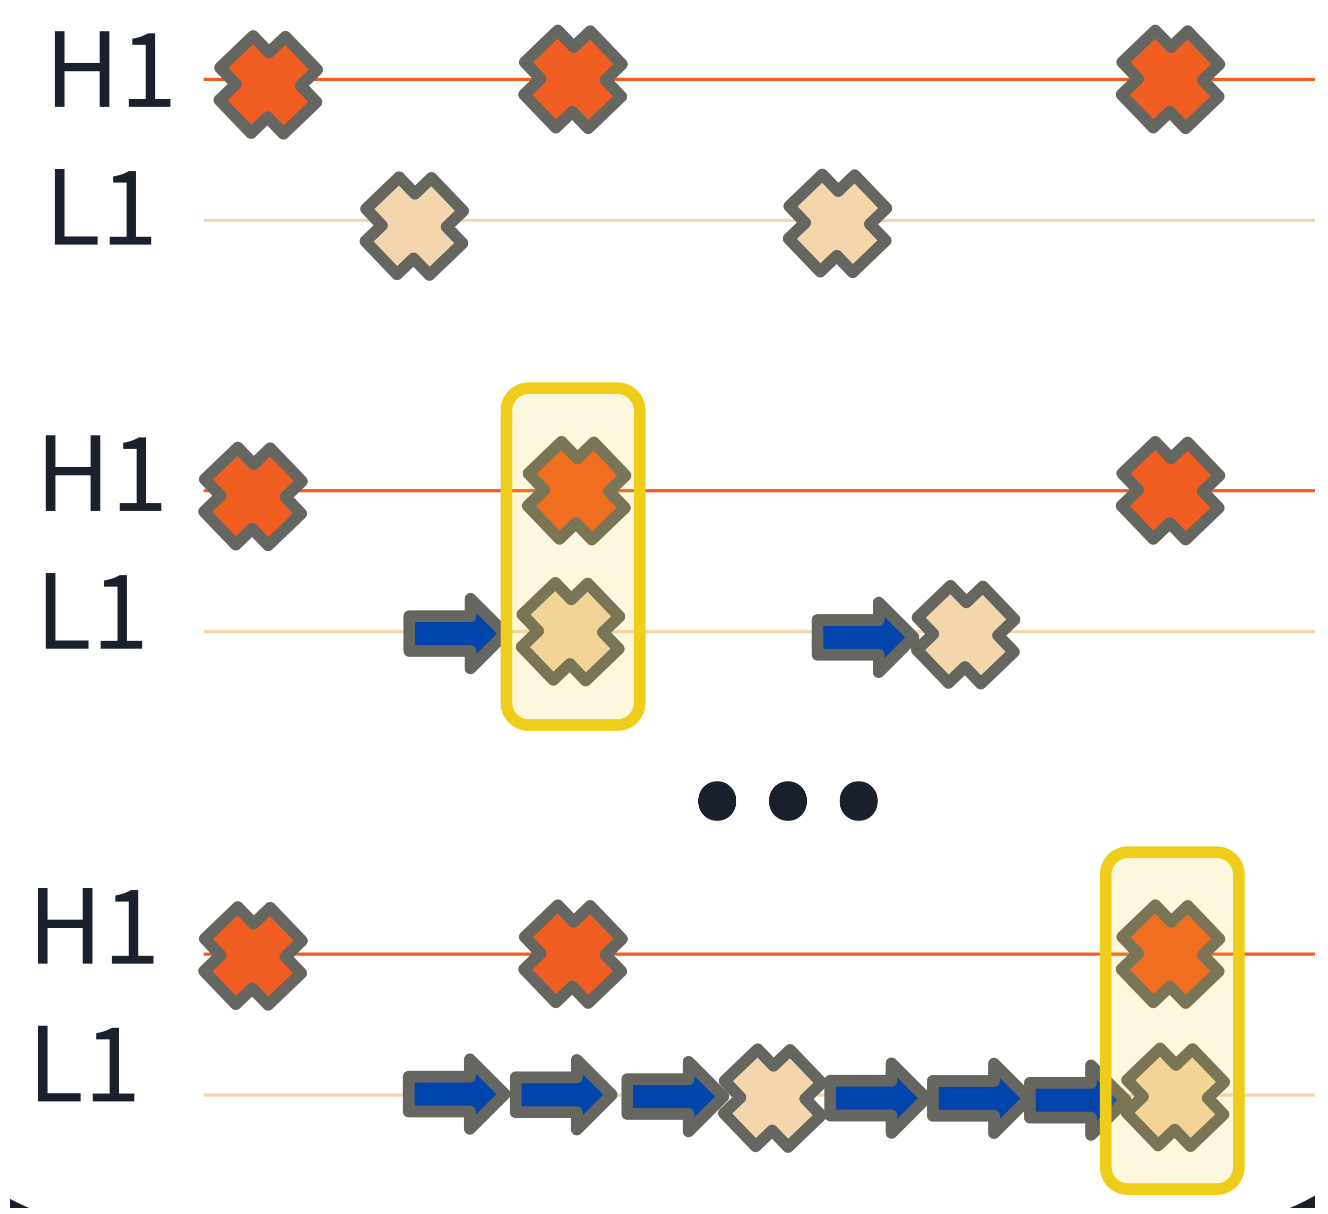

[Image: A. Nitz]

Use **time shifts** between data streams to determine the occurrence rate of coincident noise triggers = background.

Example: GWTC-1

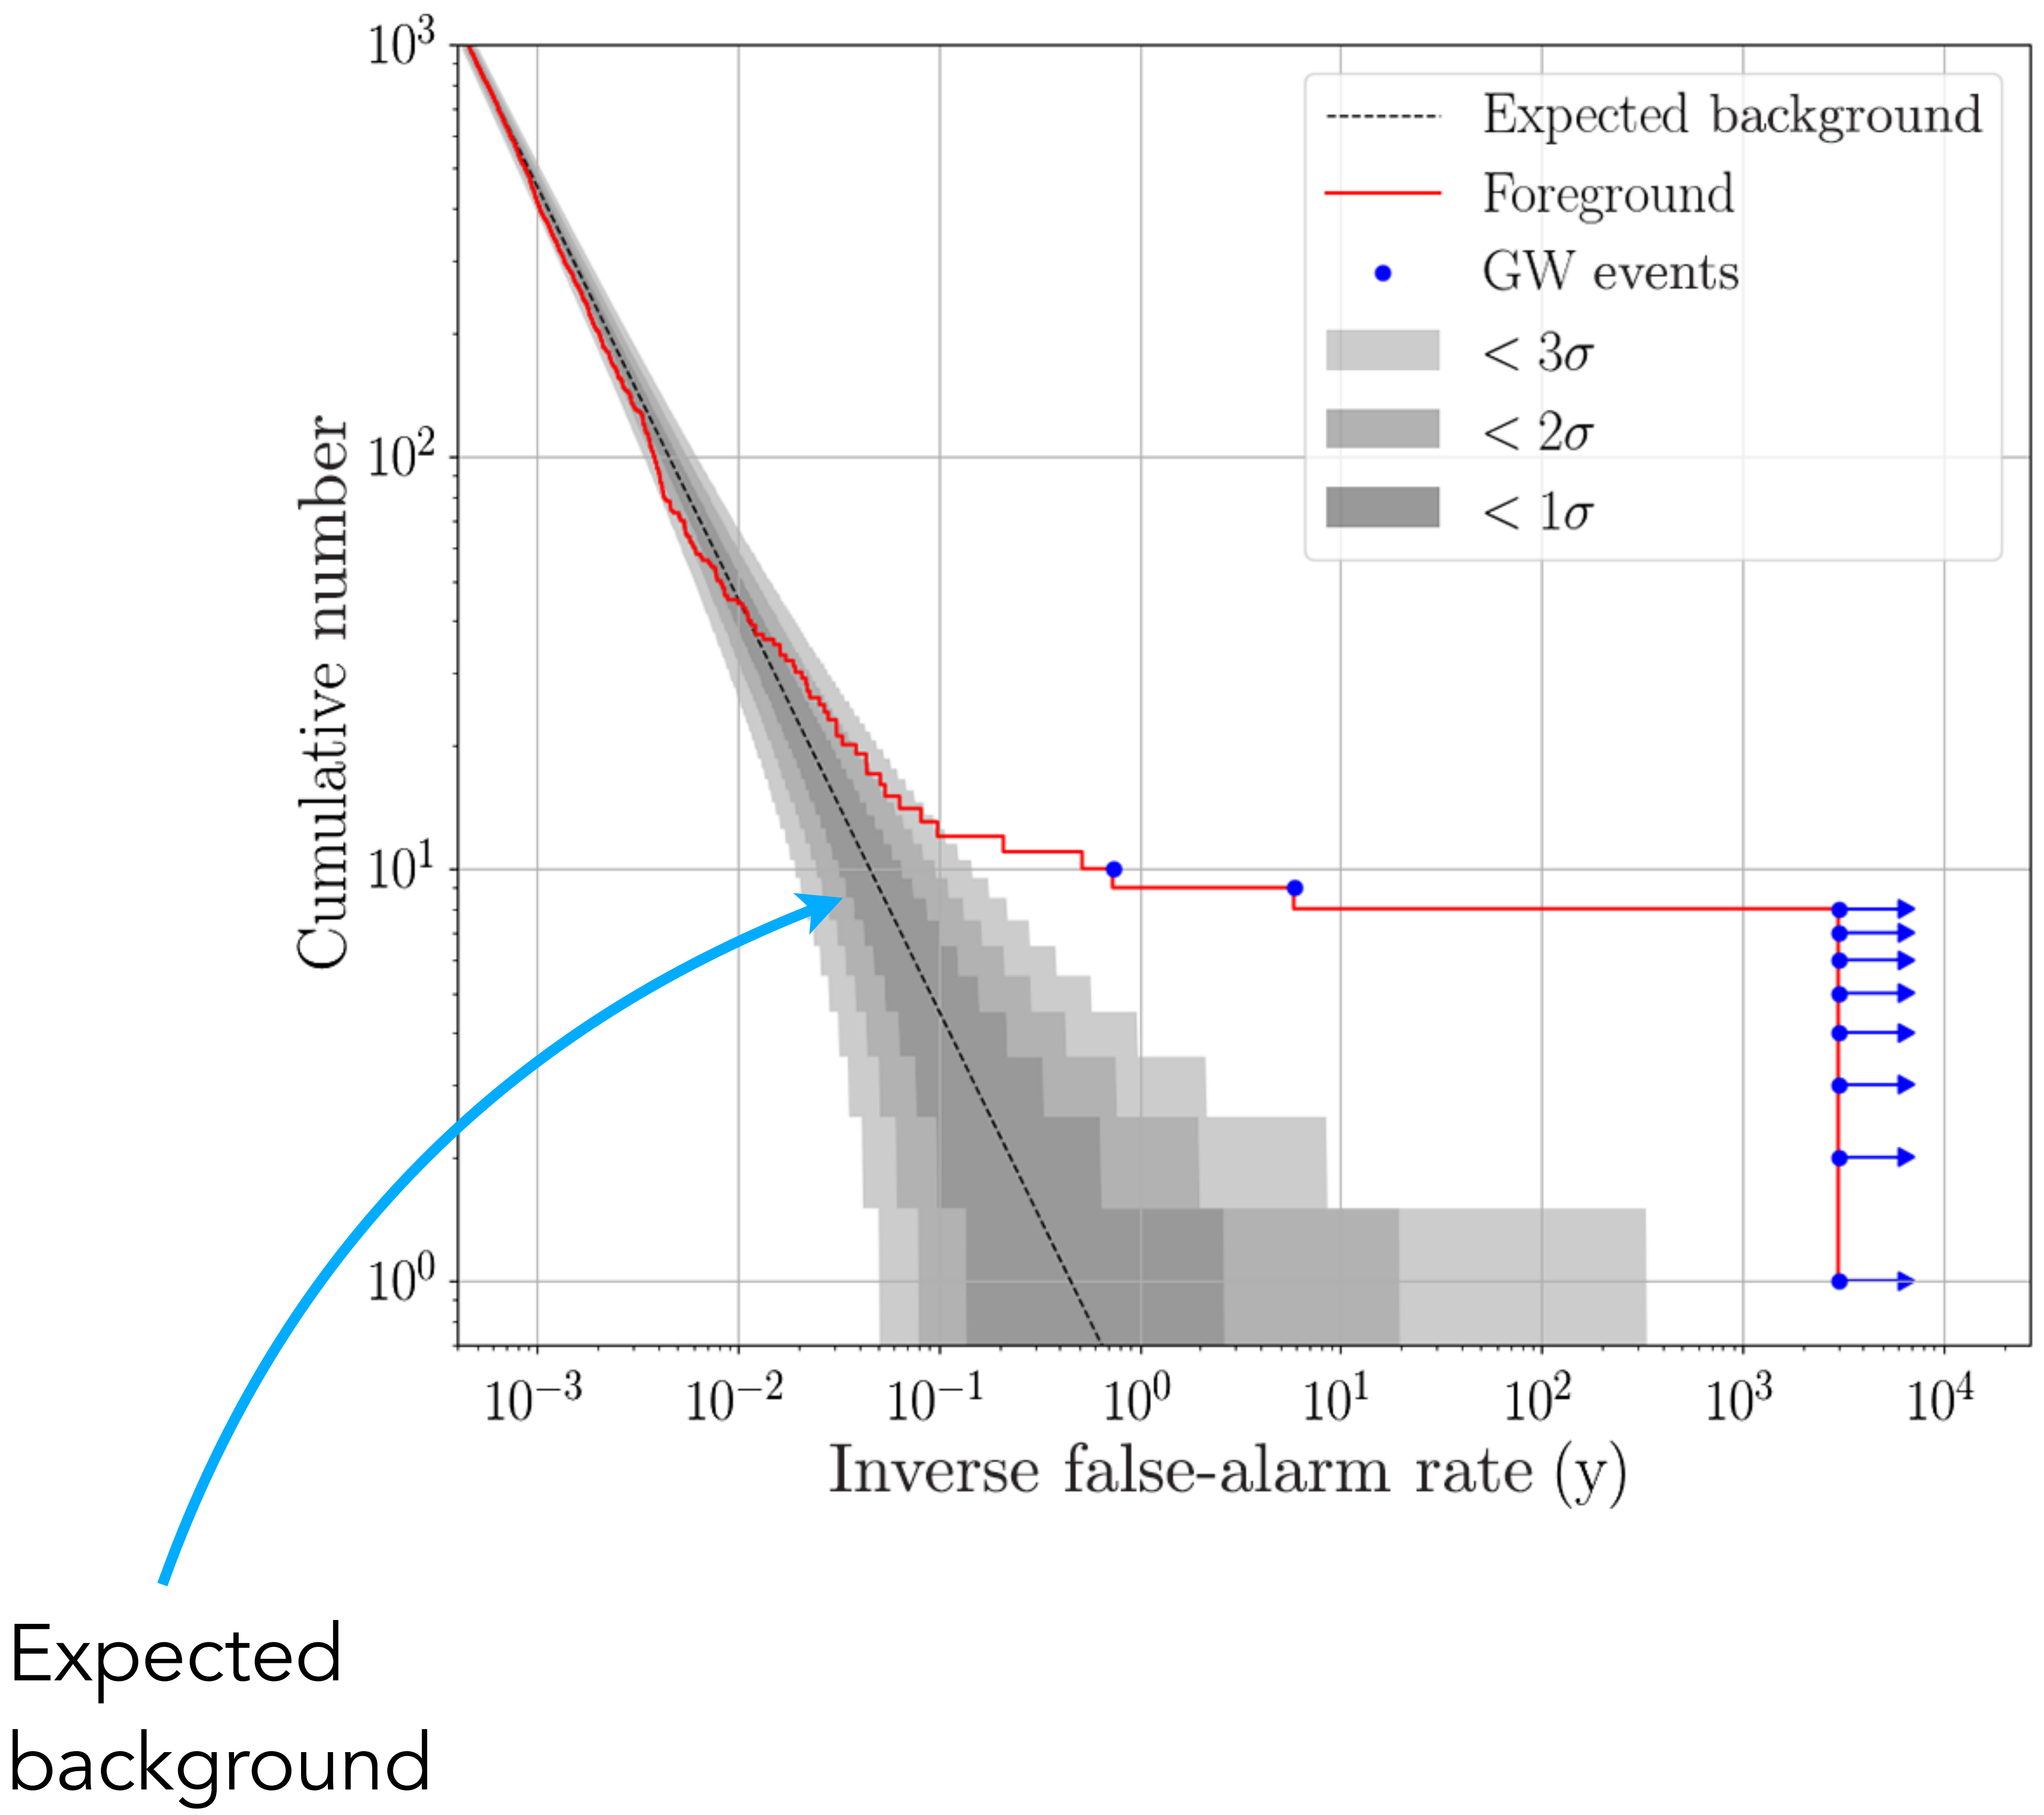

IF YOU DON'T KNOW WHAT YOU ARE LOOKING FOR

- **Unmodelled searches**
- Identification of **coincident excess power** in the time-frequency representation of the strain data
  - Identifies events that are coherent in multiple detectors and reconstructs the source sky location and signal waveforms by using the constrained maximum likelihood method
  - Does not rely on waveform models
  - Sensitive to a wide range of short-duration transient signals ("bursts")
  - Weak assumption of "chirpyness" of the signal
- Detection statistic: coherent energy constructed via cross-correlation $E_c \propto \rho_c$

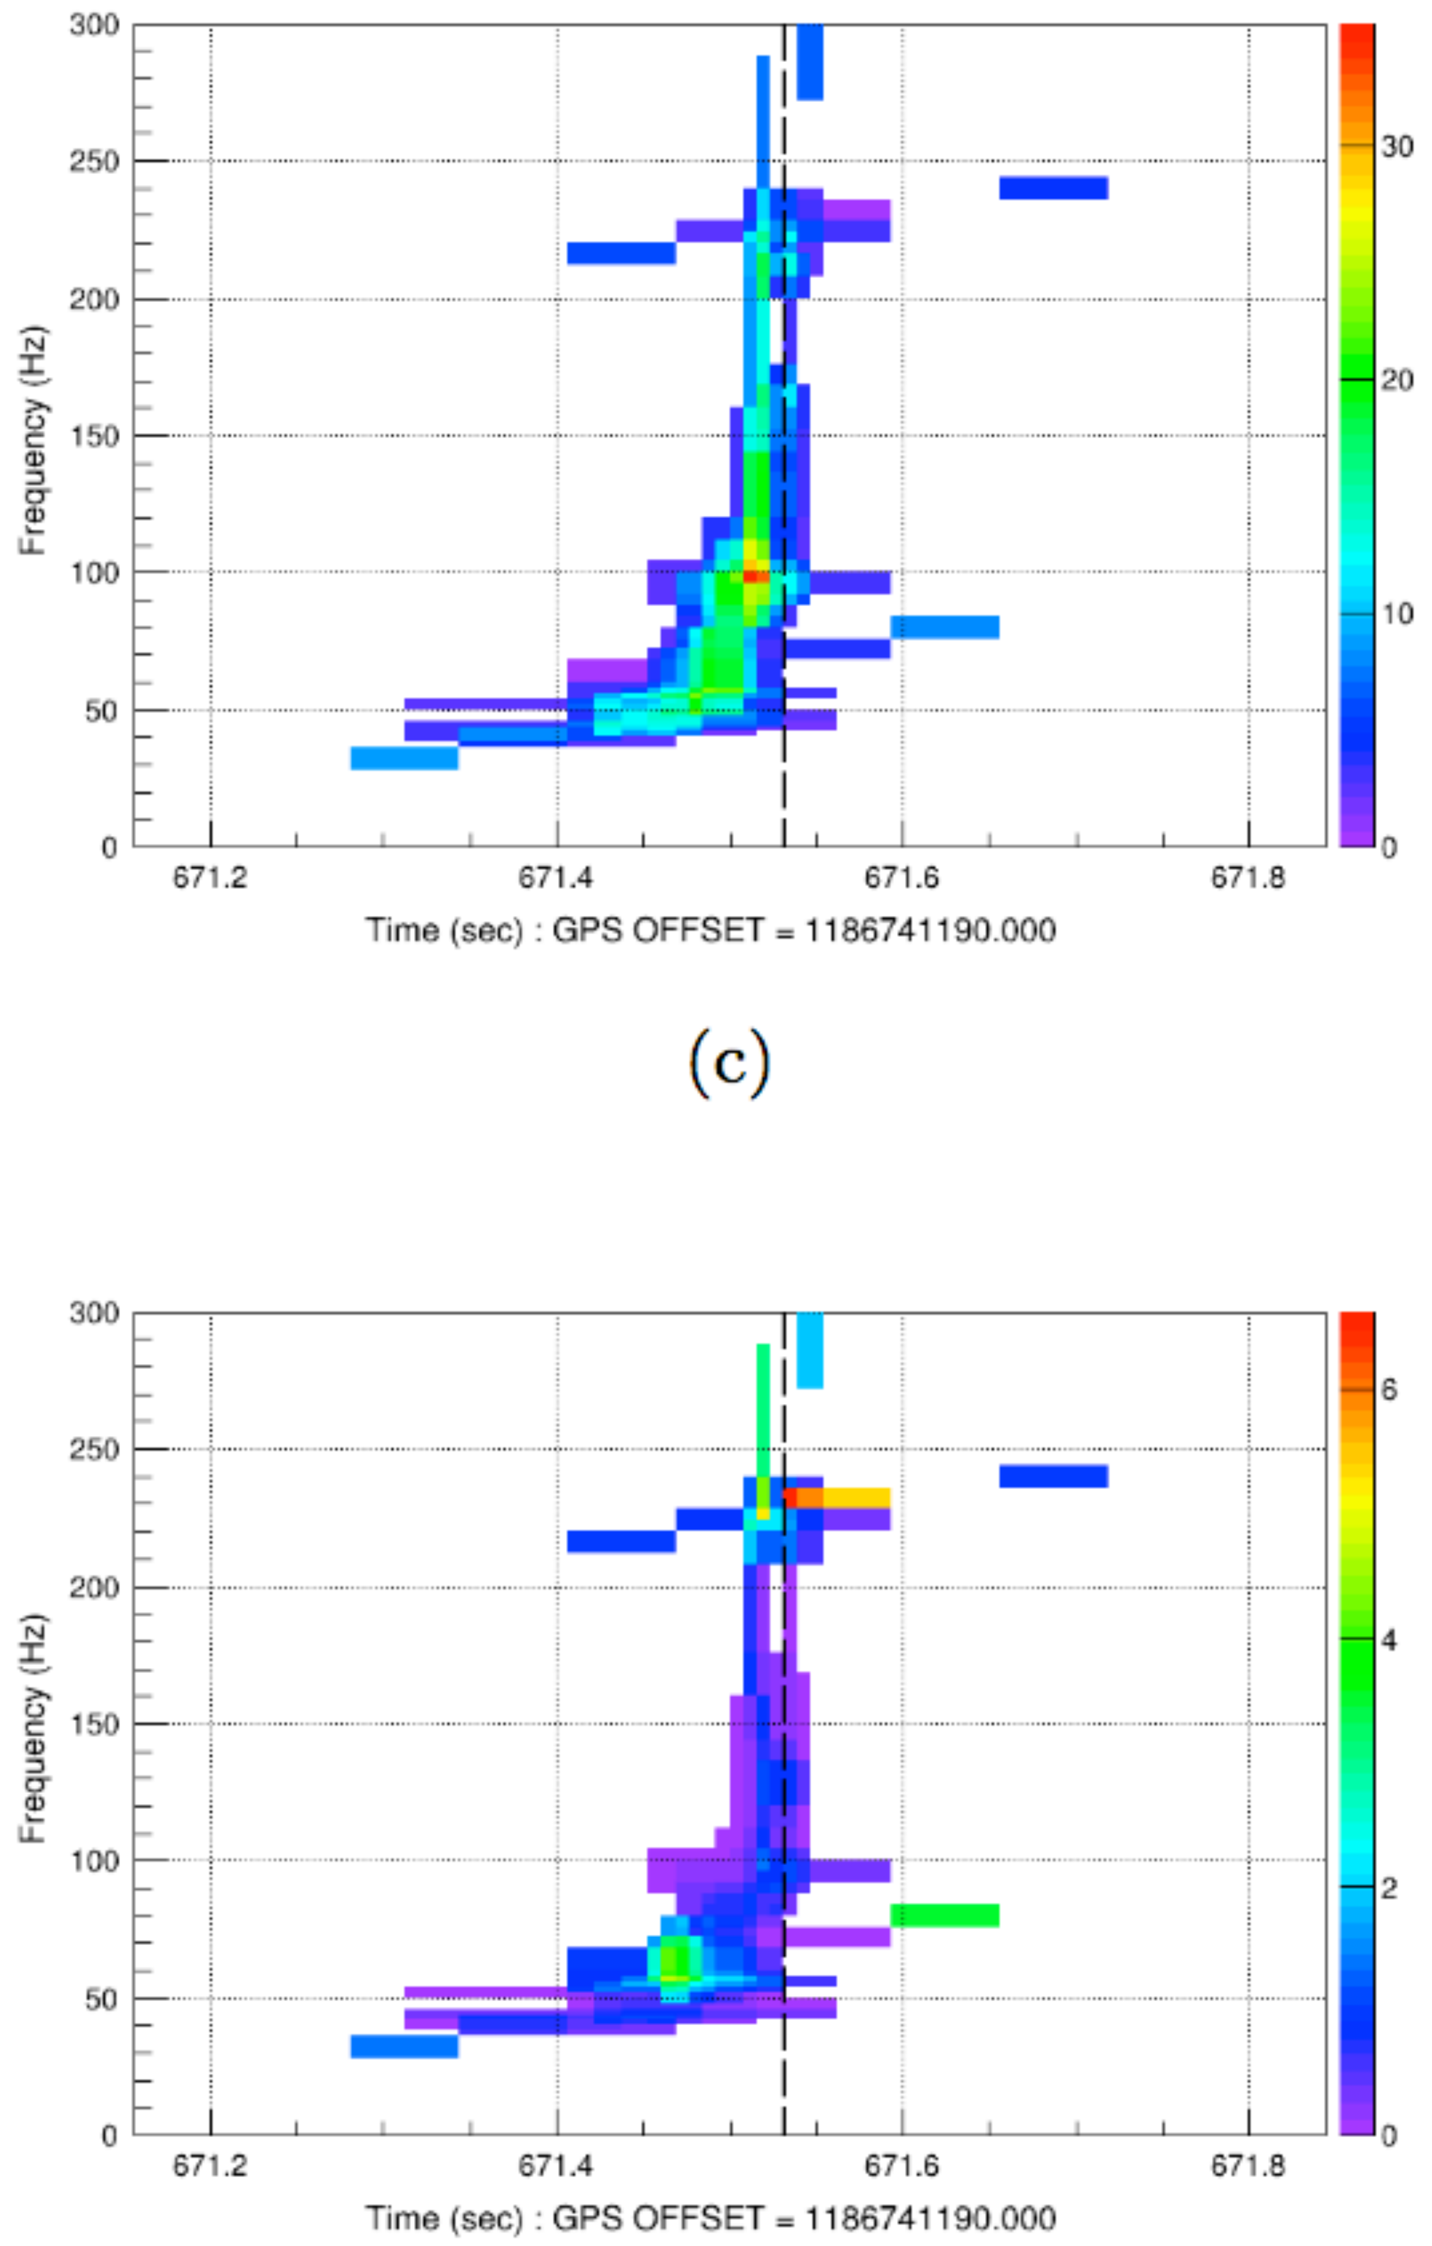

[Salemi+, 2019]

PARAMETER ESTIMATION

SOURCE CHARACTERISATION

Calibrated (raw-ish) strain data → **Astrophysical parameter distributions**

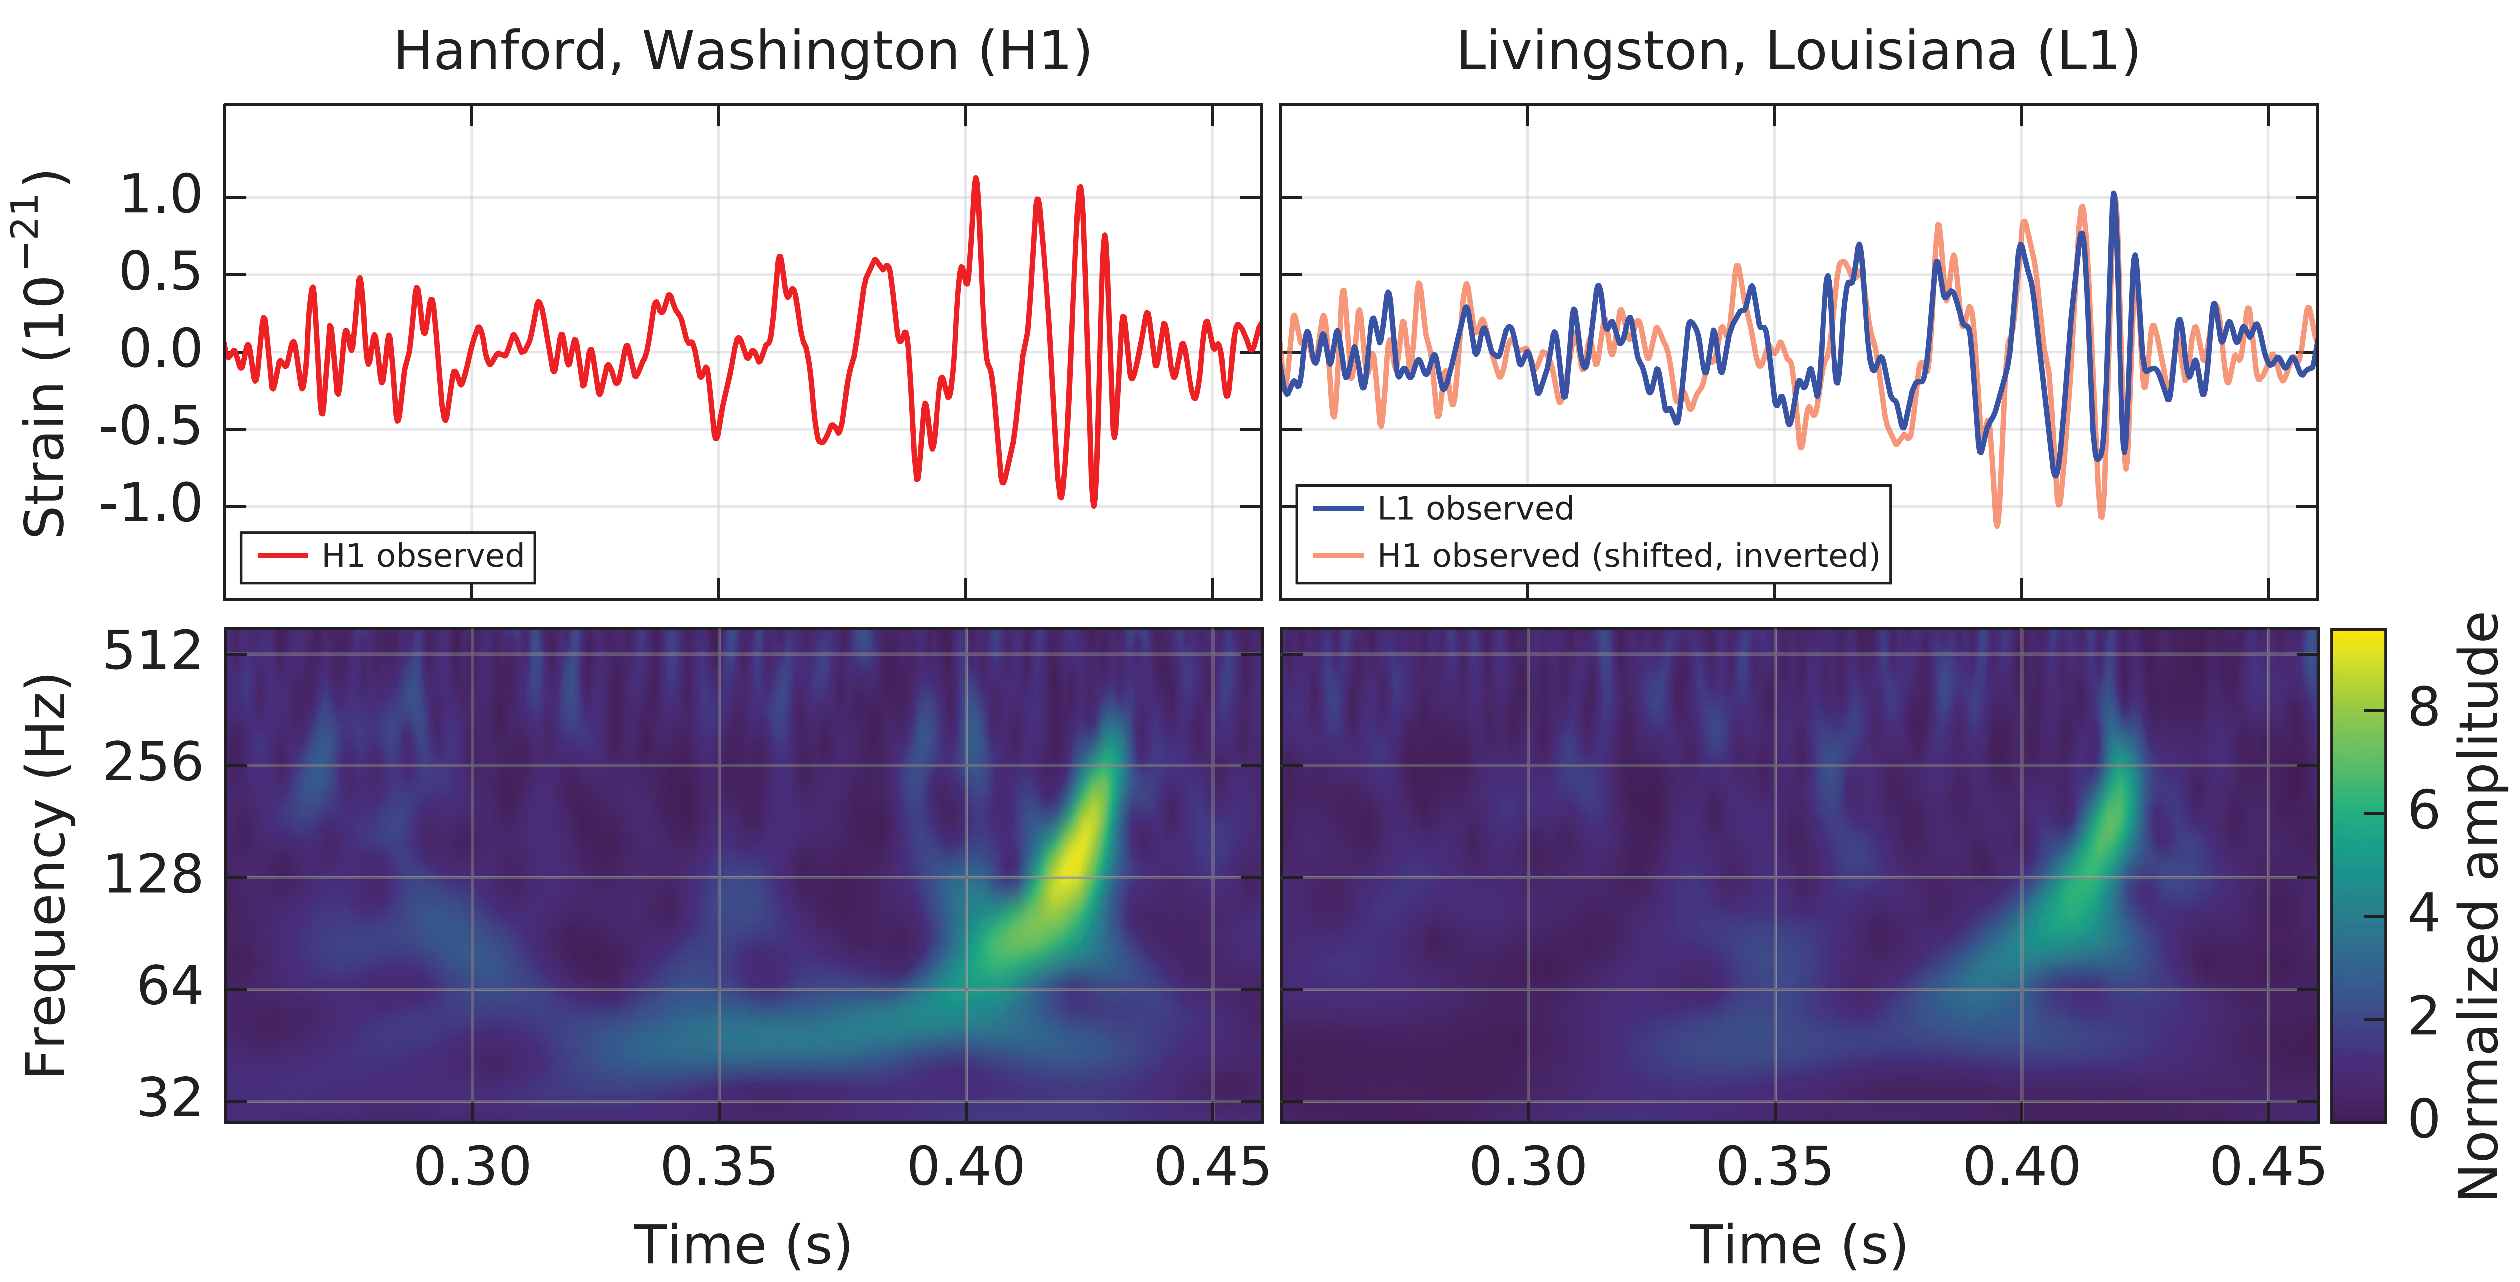

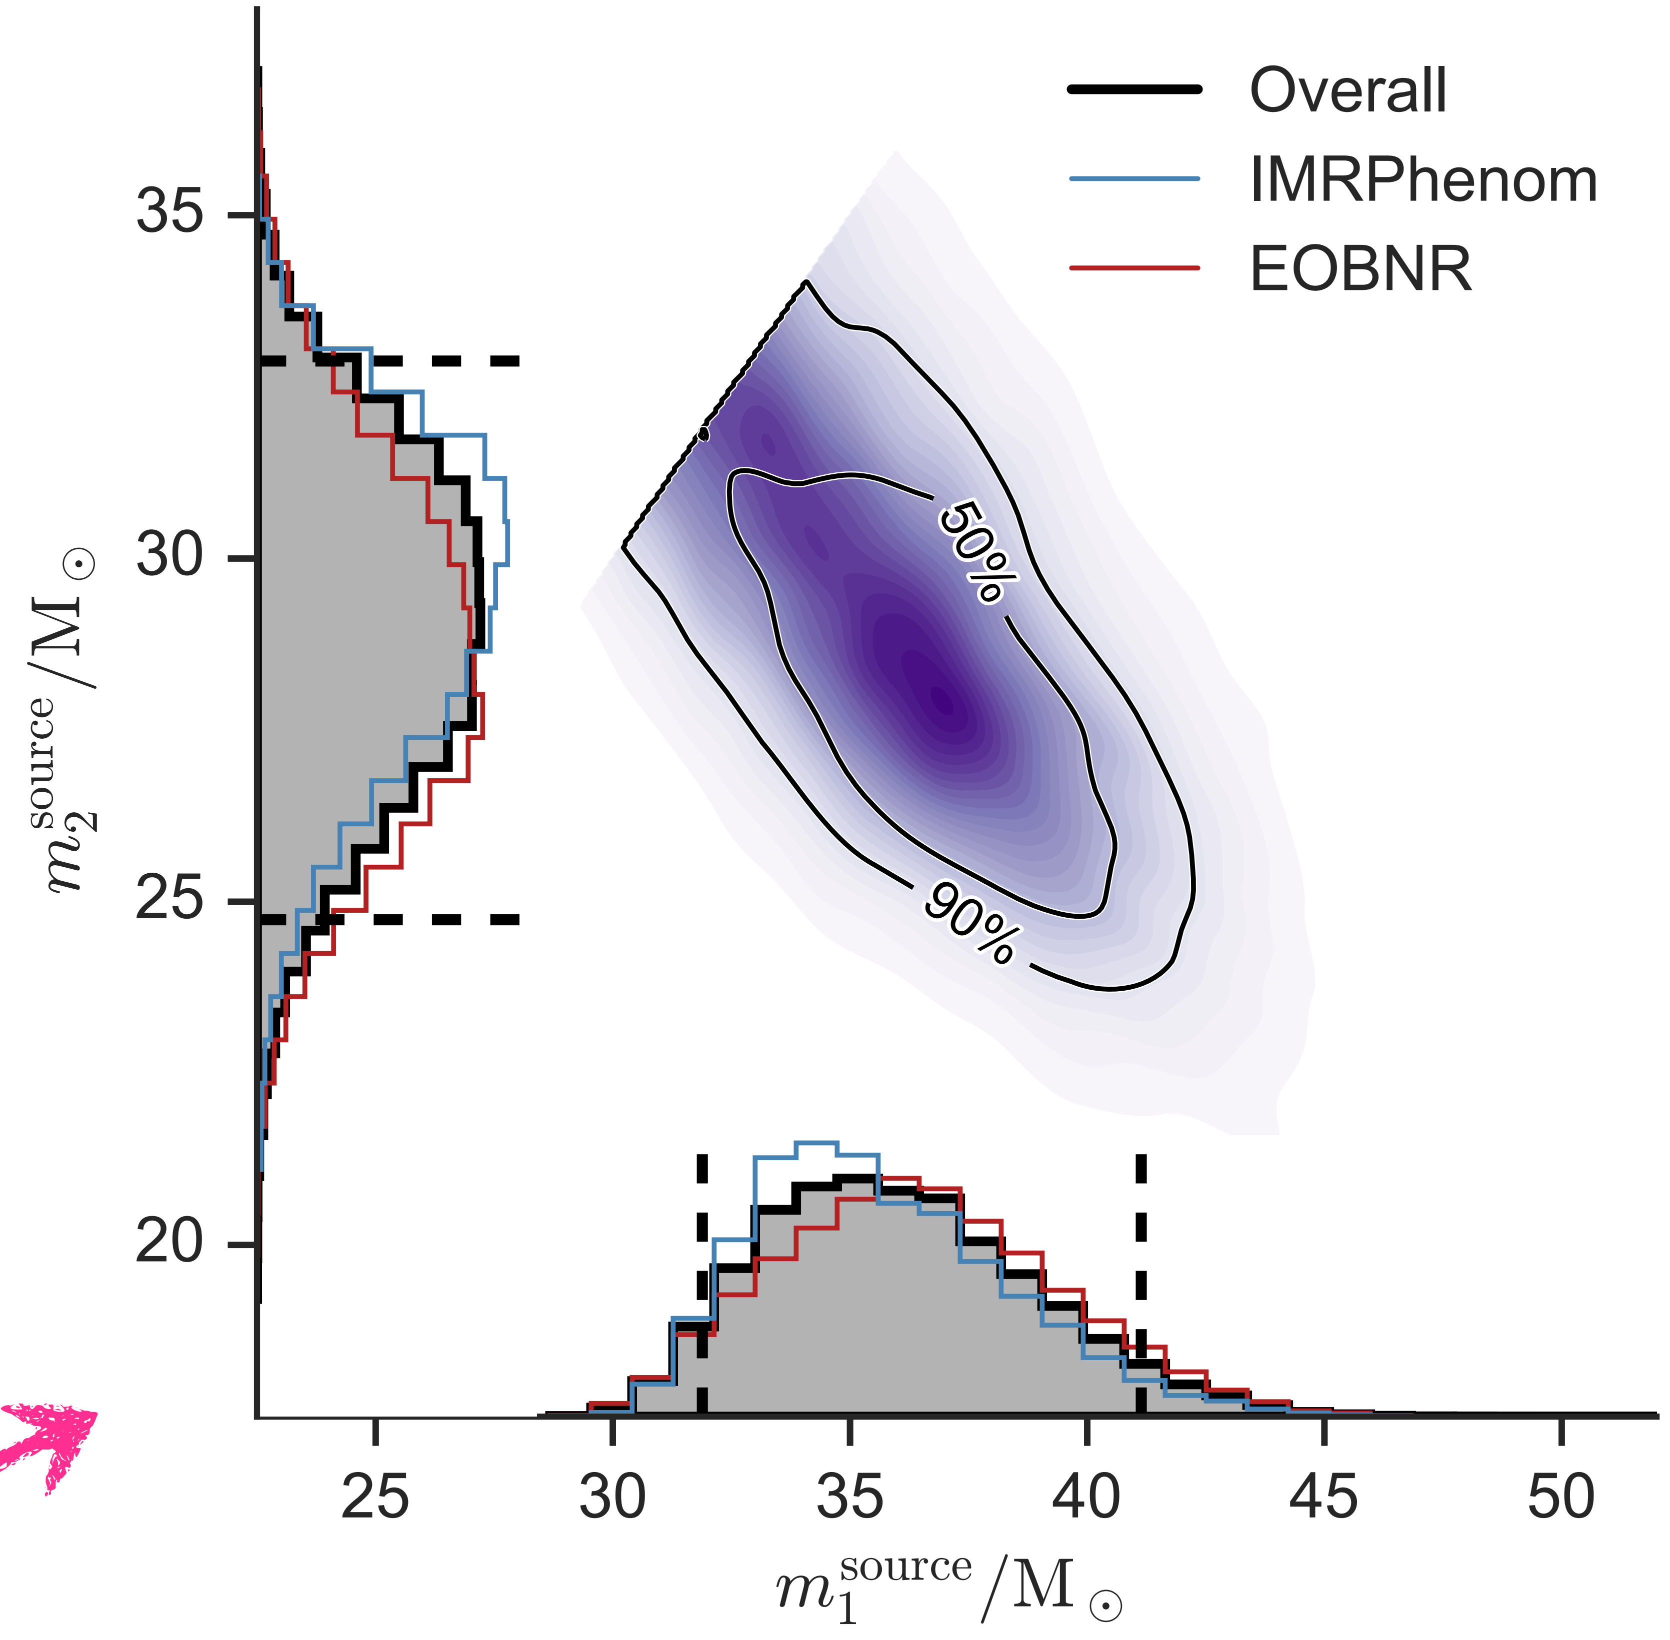

[LVC, GW150914 discovery papers]

THE SIGNAL MODEL

- Need a signal model $h(\theta)$ to compare against the data.

- Compact binary signals span a multi-dimensional parameter space:

$$\theta = \{\underbrace{m_1, m_2, \vec{\chi}_1, \vec{\chi}_2, e, \ell, \text{tidal params}}_{\text{intrinsic}}, \underbrace{D_L, \iota, \alpha, \delta, \psi, \phi_c, t_c}_{\text{extrinsic}}\}$$

- Eccentric orbits: Require two parameters to describe the ellipse and its orientation

- Binary neutron stars: Many parameters that characterise the tidal response of the star

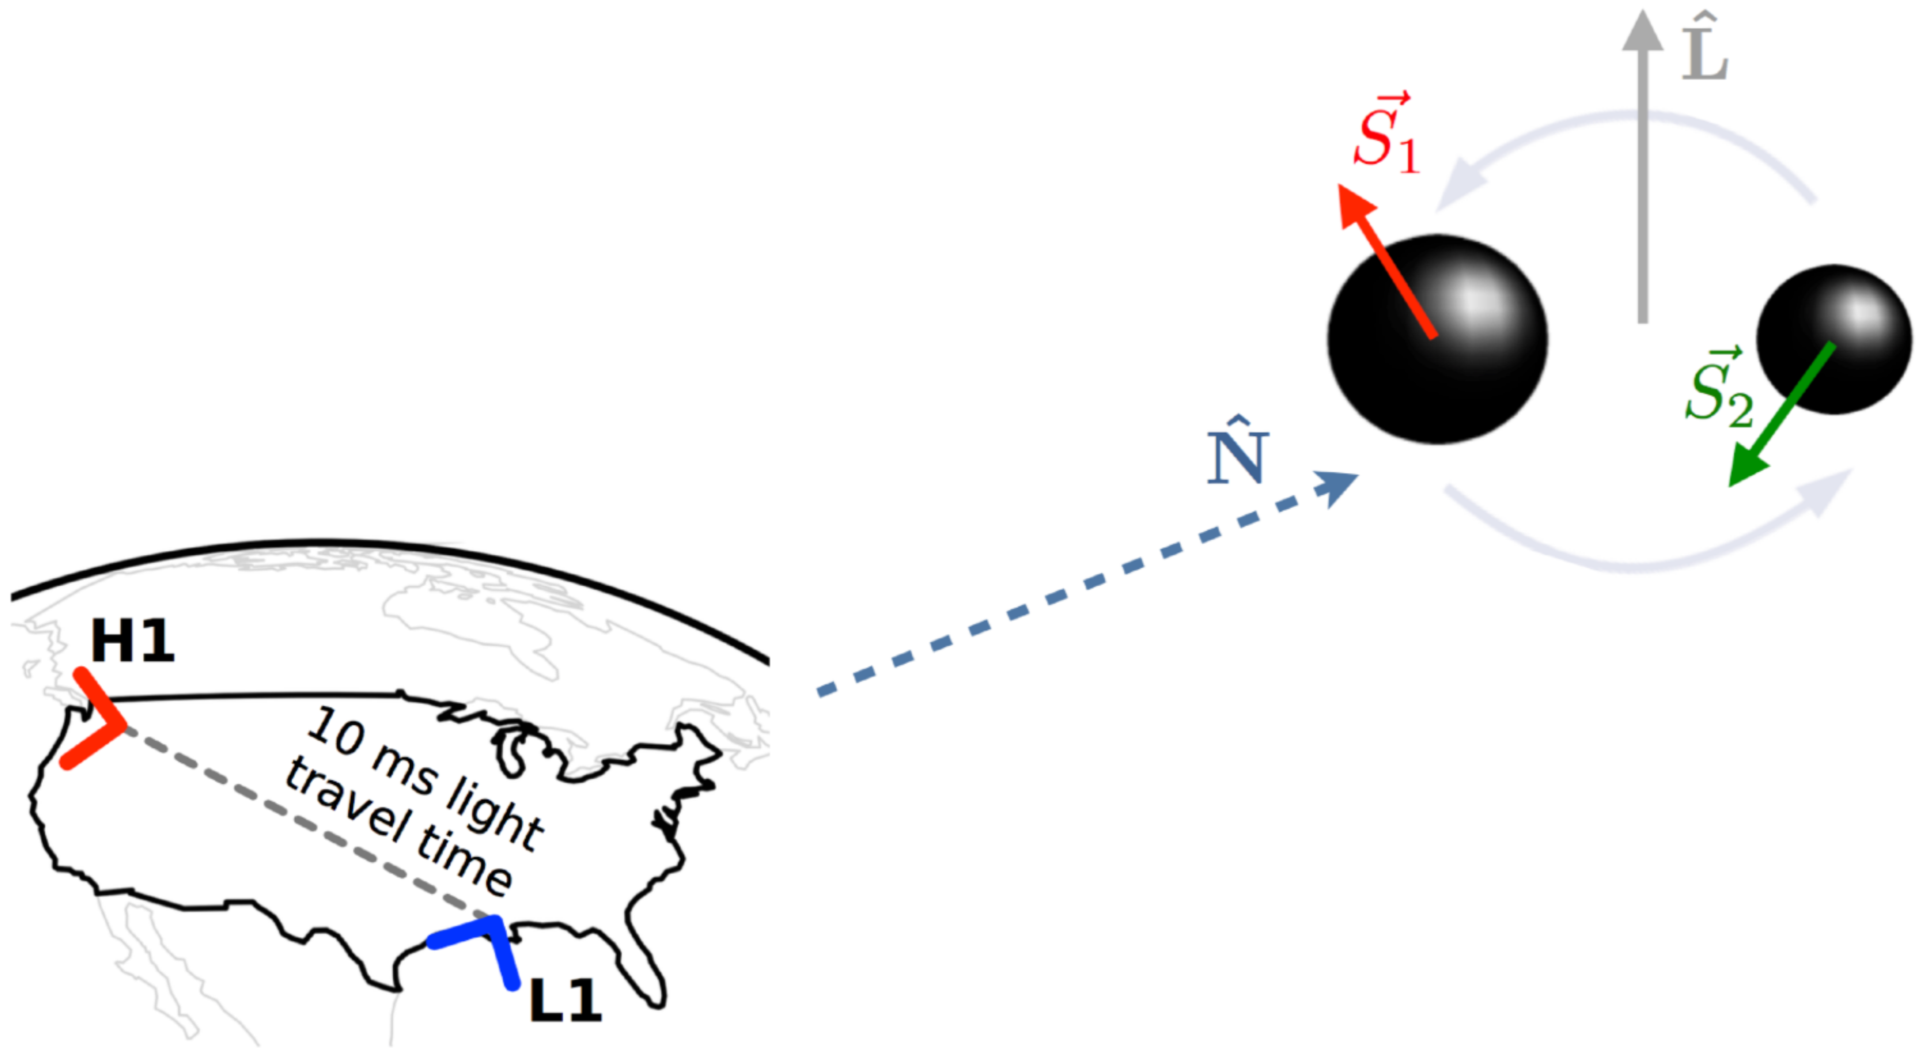

Credit: LIGO/Virgo

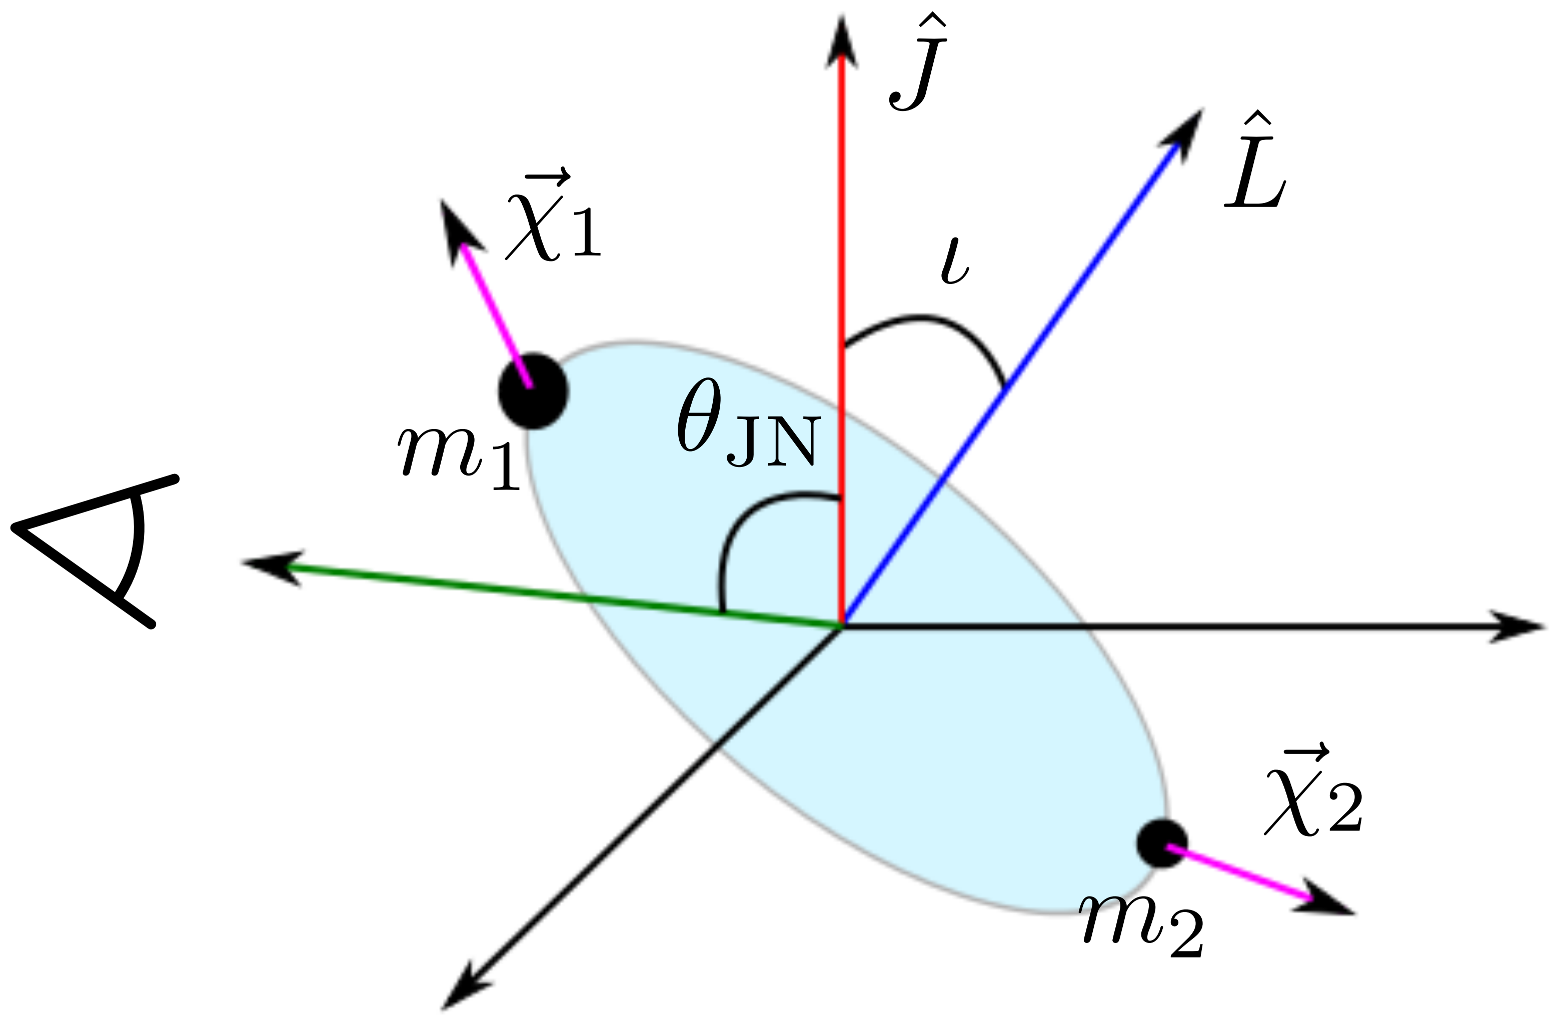

MEASURING BINARY PARAMETERS

- **Frequentist's approach:** Maximum likelihood estimation (MLE)
  - The template (**point estimator**) that best fits the data maximises the likelihood of observing that data given the parameters θ.
    - It approximates the true parameters except for the statistical error due to noise.
    - The statistical error is the variance of the estimator under different noise realisations.

- For a **Gaussian noise realisation**, the probability for the noise to have some realisation n_0 is:

$$p(n = n_0) \propto e^{-\langle n_0 | n_0 \rangle / 2}$$

- Conversely, the probability for measuring the data $d(t)$ is given by:

$$p(d) \propto e^{-\langle d-h | d-h \rangle / 2}$$

Gaussian likelihood

- The waveform that fits the data best, will maximise the probability!

MAXIMUM LIKELIHOOD ESTIMATION (MLE)

- Consider a data stream containing a signal with true parameters $\theta_{\rm true}$.
- The **MLE parameters** are then given by:

$$\hat{\theta}_{\rm MLE} = \arg\max_{\theta} p(d|h(\theta)) = \arg\min_{\theta} \langle d - h(\theta) | d - h(\theta) \rangle$$

- Recall that $\langle\, . \,|\, . \,\rangle$ defines a scalar product. It defines a metric on the signal manifold:

$$g_{ab} := \left\langle \frac{\partial \hat{h}}{\partial \theta^a} \middle| \frac{\partial \hat{h}}{\partial \theta^b} \right\rangle \equiv \rho^{-2} \Gamma_{ab}$$

Fisher information matrix (FIM)

- The metric can be used to compute the proper distance between any two waveforms with parameters θ^a and $\theta^a + d\theta^a$:

$$ds^2 = g_{ab} d\theta^a d\theta^b$$

- The MLE minimises the distance between the data and the signal surface.

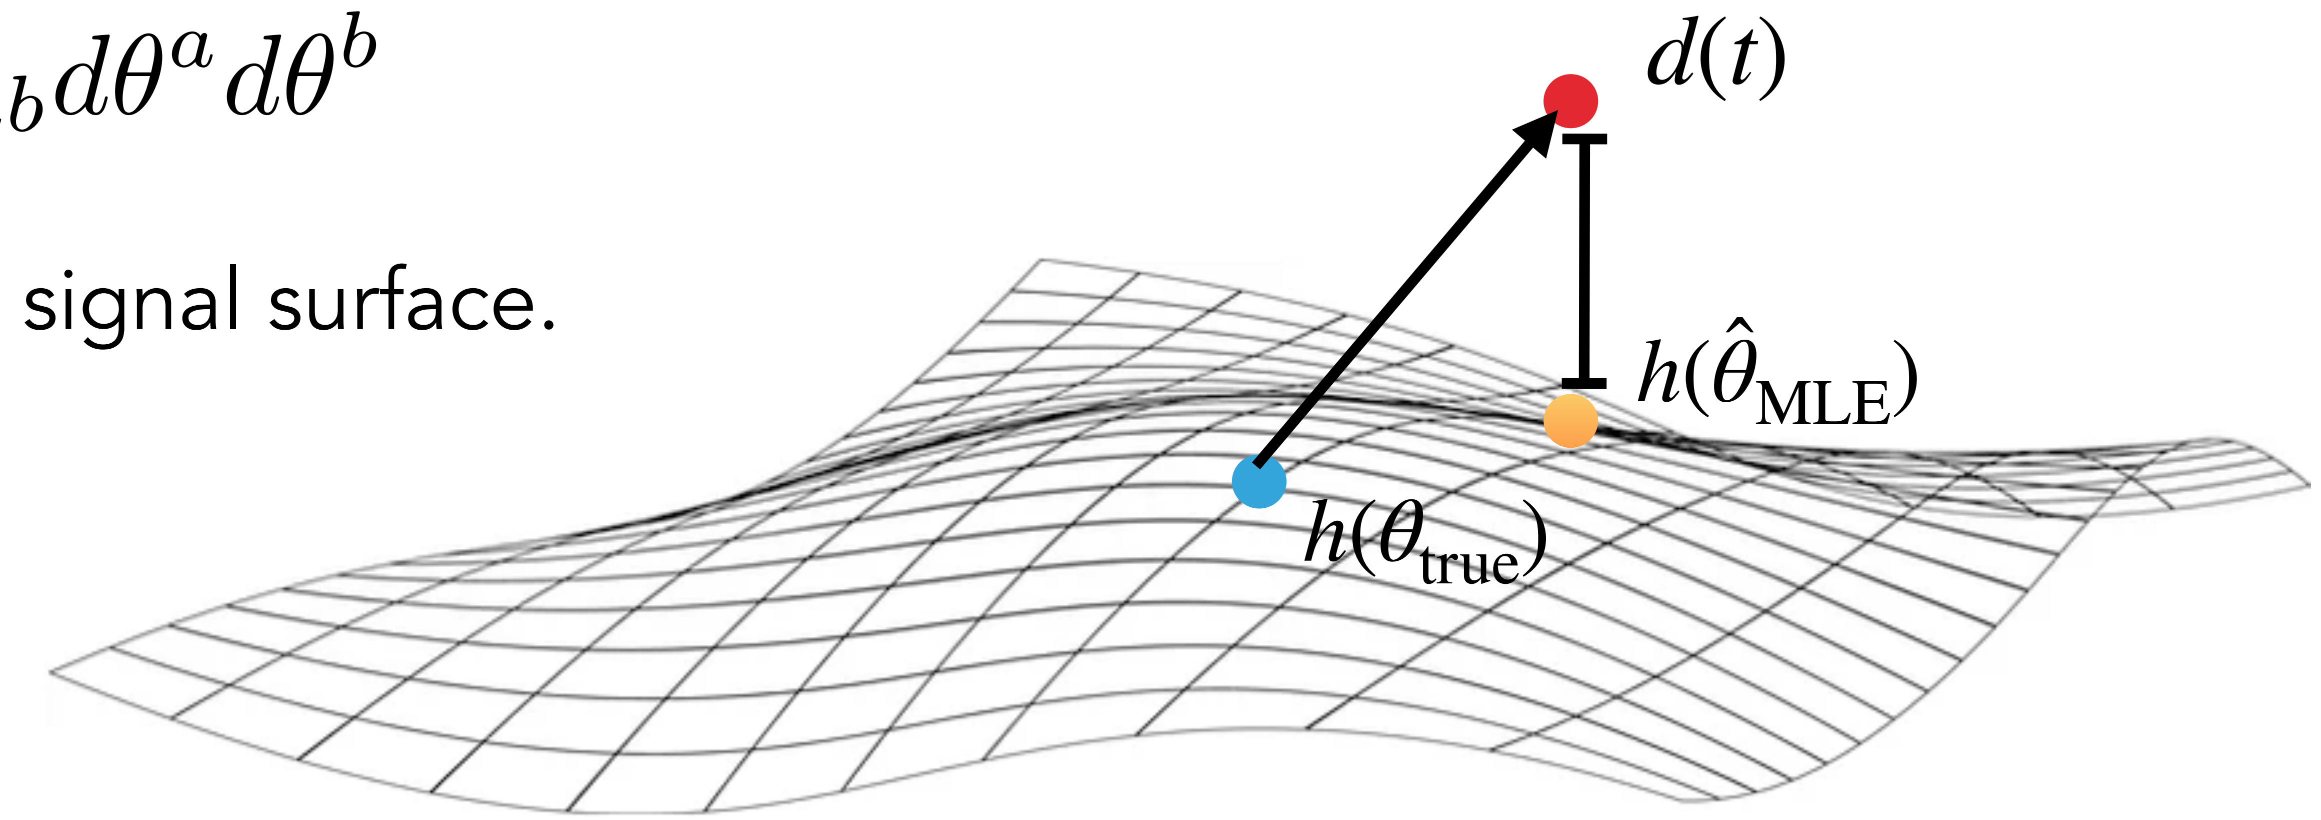

MLE UNCERTAINTIES

- Three options to calculate the **parameter uncertainty**:
- **Fisher information matrix (FIM):** quick & dirty; cannot handle correlations; only valid in high-SNR limit.
  - Near the ML parameters, the likelihood surface is approximately Gaussian, and the best-fit parameters will follow a multivariate Gaussian distribution.
- The Taylor expansion of the inner product in the likelihood around the MLE parameters yields:

$$\ln p(d) \approx \ln p(d|h(\hat{\theta}_{\mathrm{MLE}})) - \frac{1}{2}\Gamma_{ab}(\theta^a - \hat{\theta}^a_{\mathrm{MLE}})(\theta^b - \hat{\theta}^b_{\mathrm{MLE}})$$

$$p(\Delta\theta^a) = \mathcal{N} e^{-\Gamma_{ab}\Delta\theta^a\Delta\theta^b/2}$$

RMS error: $\sigma_i = \sqrt{(\Gamma^{-1})_{ii}}$

MLE UNCERTAINTIES

- 3 methods
  - **Likelihood ratio:** Expensive.

$$\Lambda(\theta) = \frac{p(d|h(\theta))}{p(d|h(\hat{\theta}_{\rm MLE})} = \exp\left(-\frac{\Delta\chi^2(\theta)}{2}\right)$$

$$\Delta\chi^2(\theta) = \langle d - h(\theta)|d - h(\theta)\rangle - \langle d - h(\hat{\theta}_{\rm MLE})|d - h(\hat{\theta}_{\rm MLE})\rangle$$

Define confidence regions from the χ^2

  - **Repeated noise realisations** $n_k(t)$: Expensive + requires reliable noise model.

- IMPORTANT: All of these error estimates are due to statistical uncertainty induced by the noise!

MEASURING BINARY PARAMETERS

- **The Bayesian approach:**
  - Rather than computing a point estimate, we compute the full posterior probability distribution of the signal parameters, which encodes the astrophysical information of the source.
  - Statistical errors are characterised by the spread of the posterior.

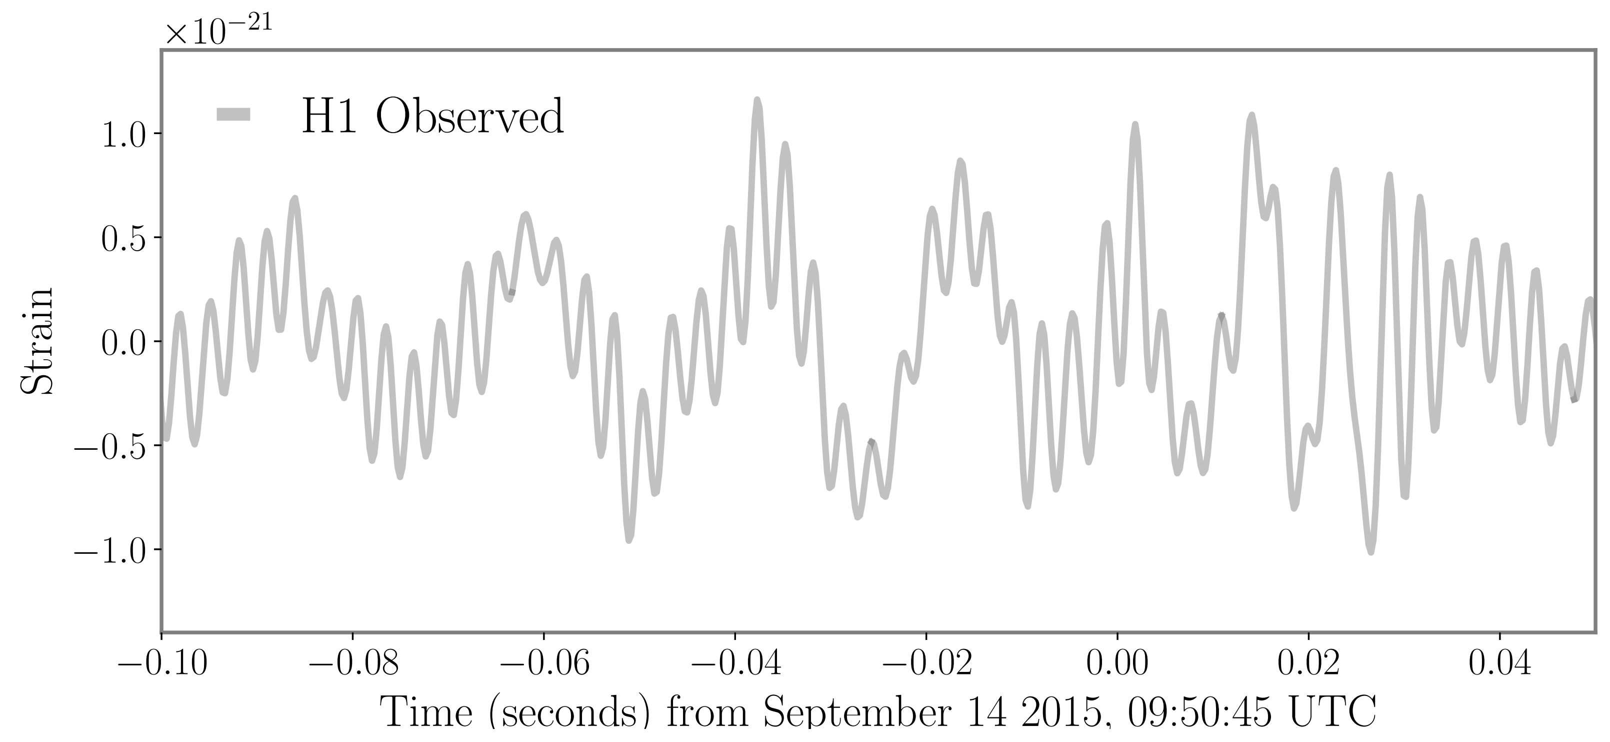

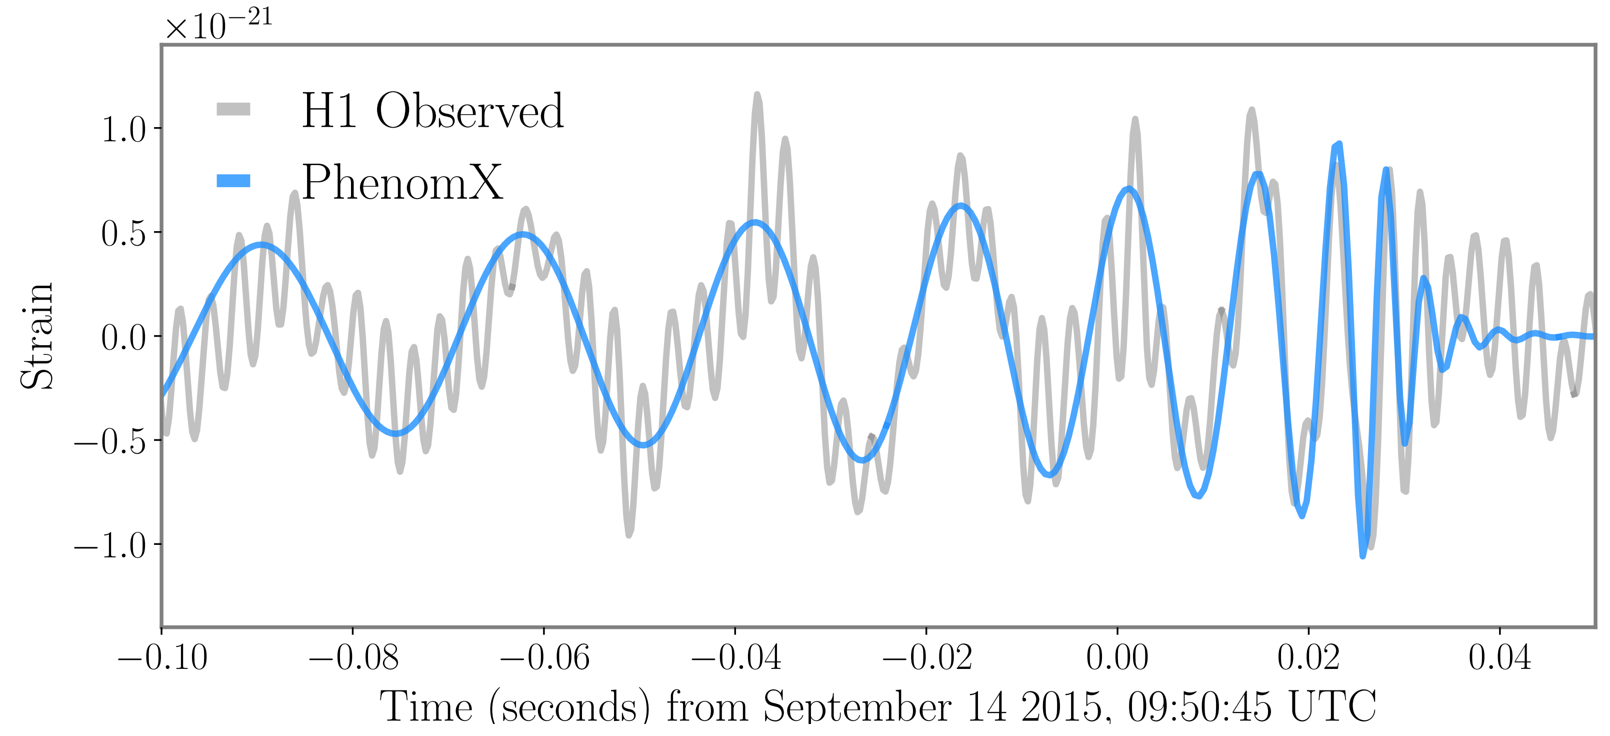

- Bayes' theorem:

Likelihood Priors

Posterior Probability
$$p(\theta \mid d) = \frac{\mathcal{L}(d \mid \theta)\pi(\theta)}{\mathcal{Z}}$$

Evidence

$$\mathcal{Z} = p(d) = \int d^n\theta \, p(d|\theta)p(\theta)$$

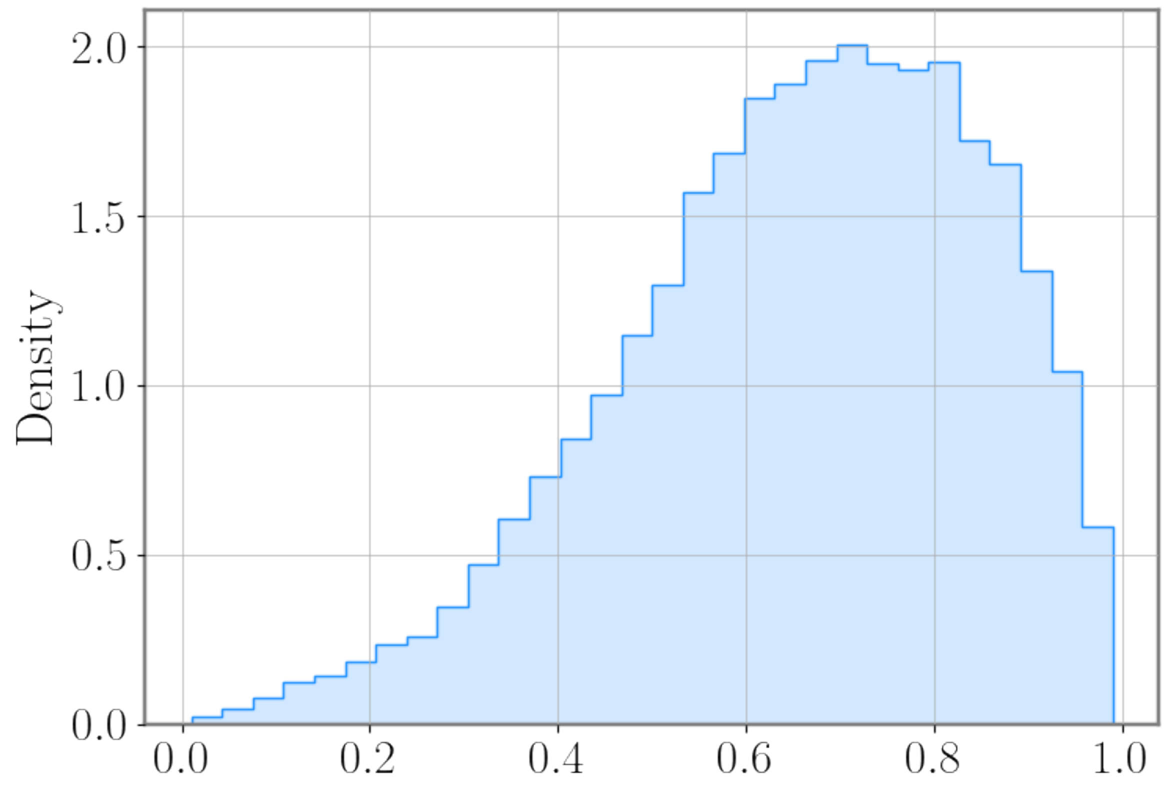

[Slide credit: G. Pratten]

THE LIKELIHOOD

- We are free to choose the likelihood.
- As before, we assume a likelihood associated with **stationary Gaussian noise** that has zero mean and a known variance:

waveform template

$$\mathcal{L}(d|\theta, H, S_n(f)) = \exp\left(\sum_i -\frac{2|\tilde{h}_i(\theta) - \tilde{d}_i|^2}{T S_n(f_i)} - \frac{1}{2}\log\left(\pi T S_n(f_i)/2\right)\right)$$

- Likelihood compares theoretical models against the data.
- For a network of GW detectors:

$$\mathcal{L}(d_{\mathrm{NW}}|\theta, H, S_{n\mathrm{NW}}(f)) = \prod_{i \in \mathrm{NW}} \mathcal{L}(d_i|\theta, H, S_{ni}(f))$$

NON-GAUSSIANITY

- Noise transients (glitches) impact the measurement:
  - Systematic errors in sky location
  - Systematic biases in source parameters

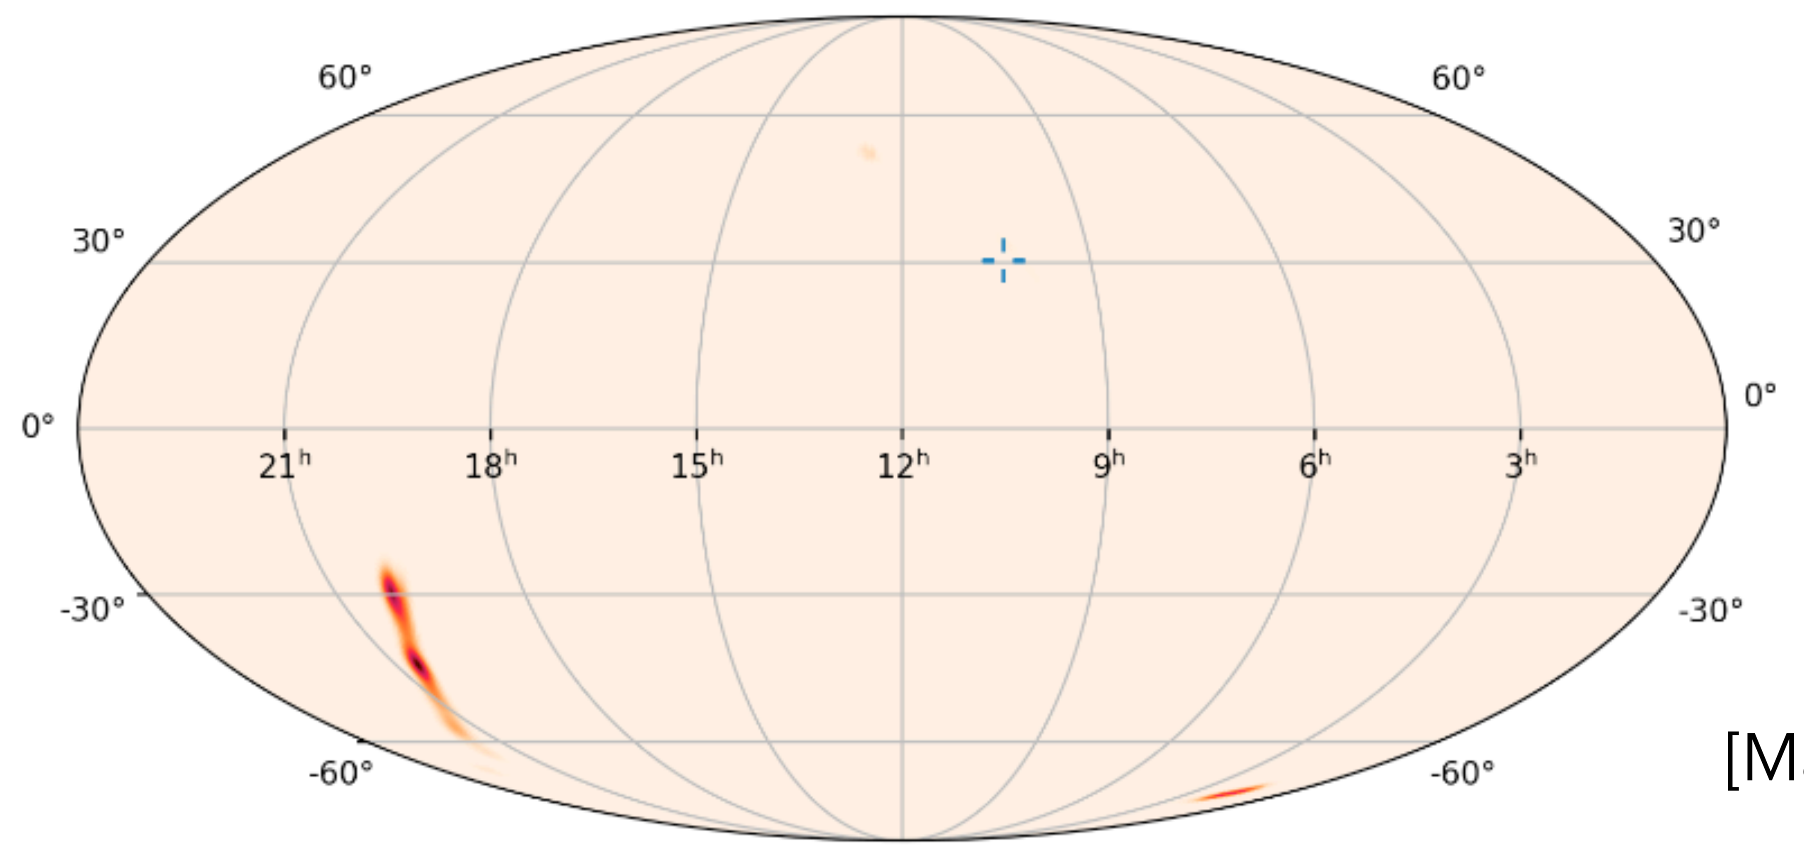

[Macas+, 2022]

(a) Sky localisation of a GW150914-like event injected at $t_0 + 30$ ms relative to the blip glitch central time t_0. The 90% credible area is 137 deg^2.

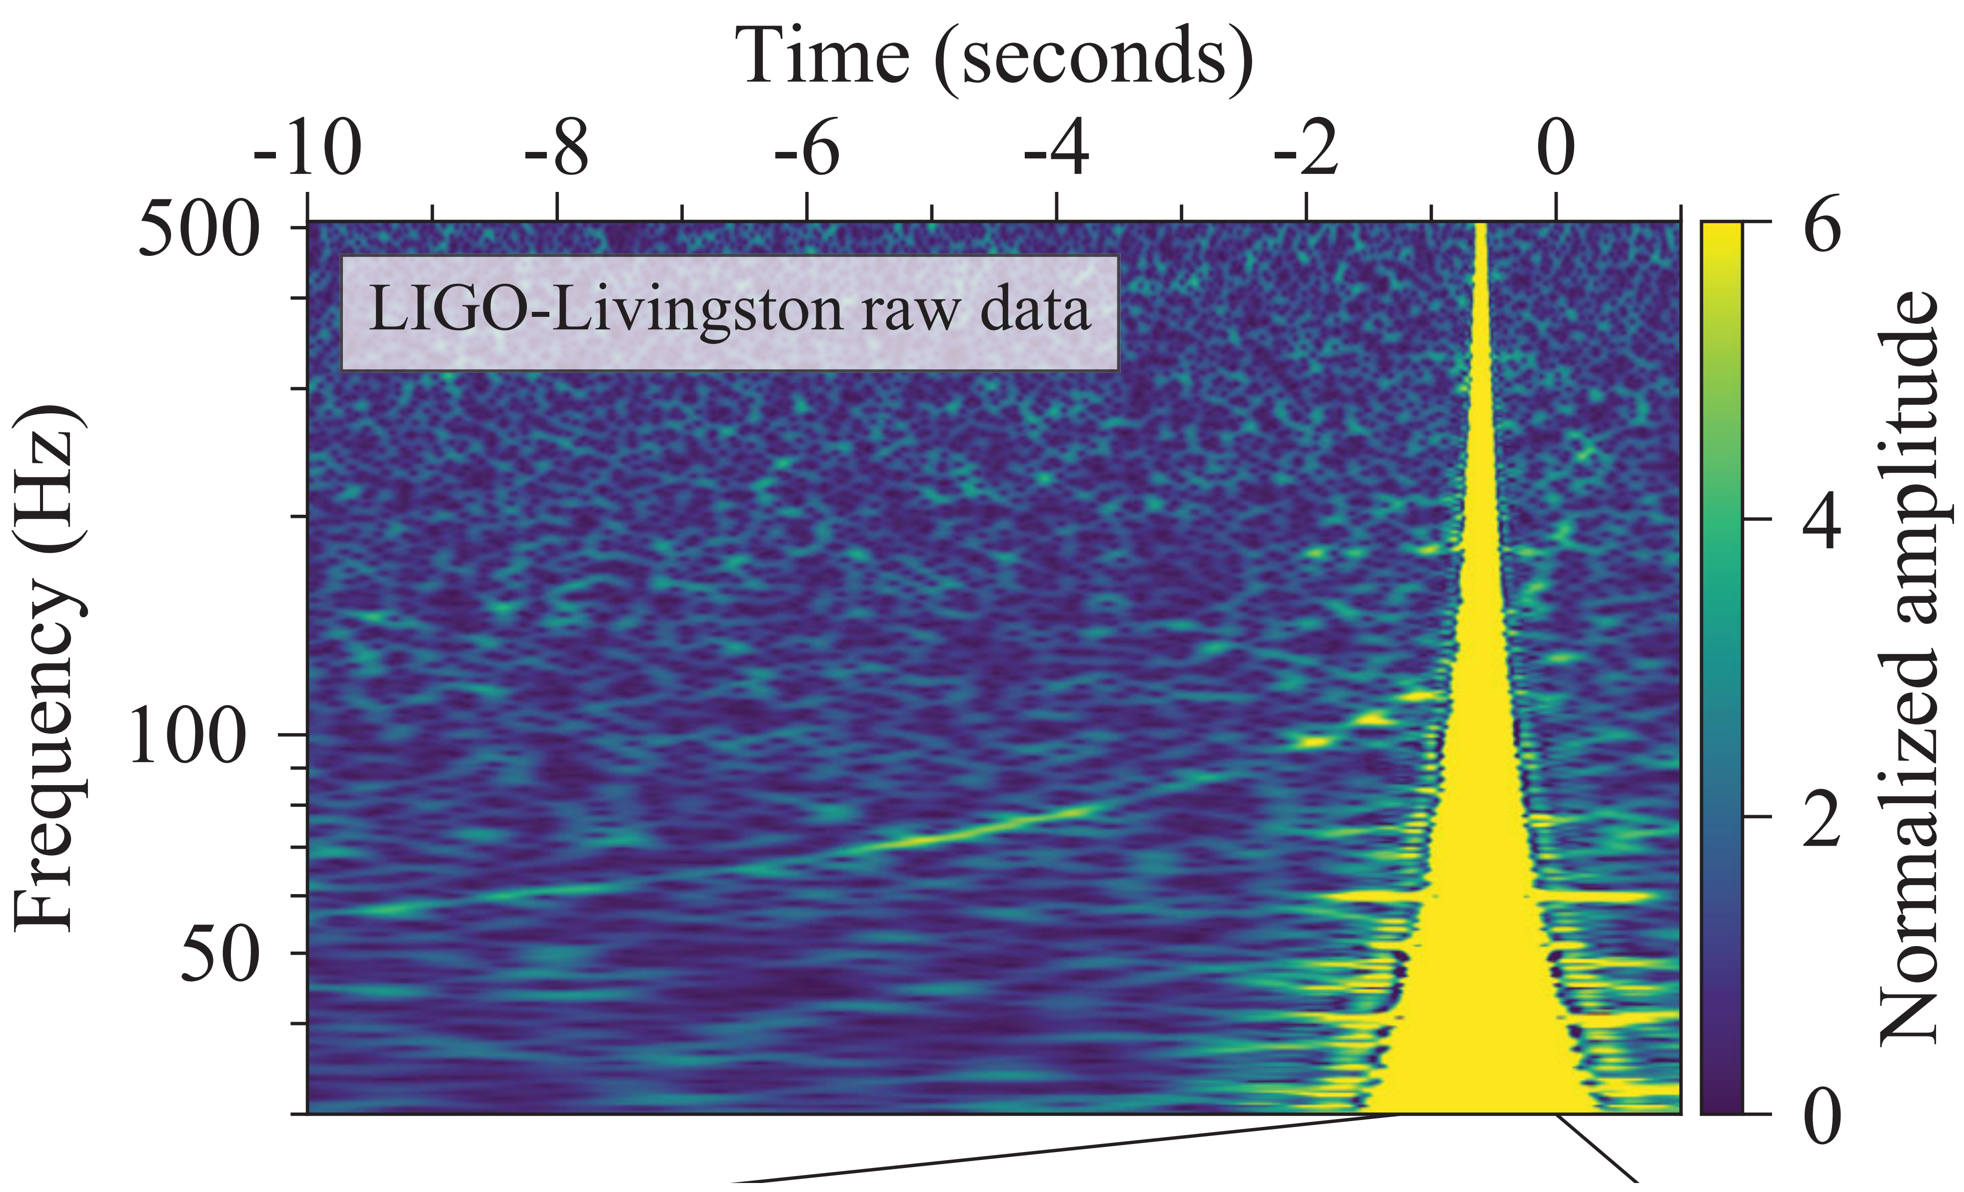

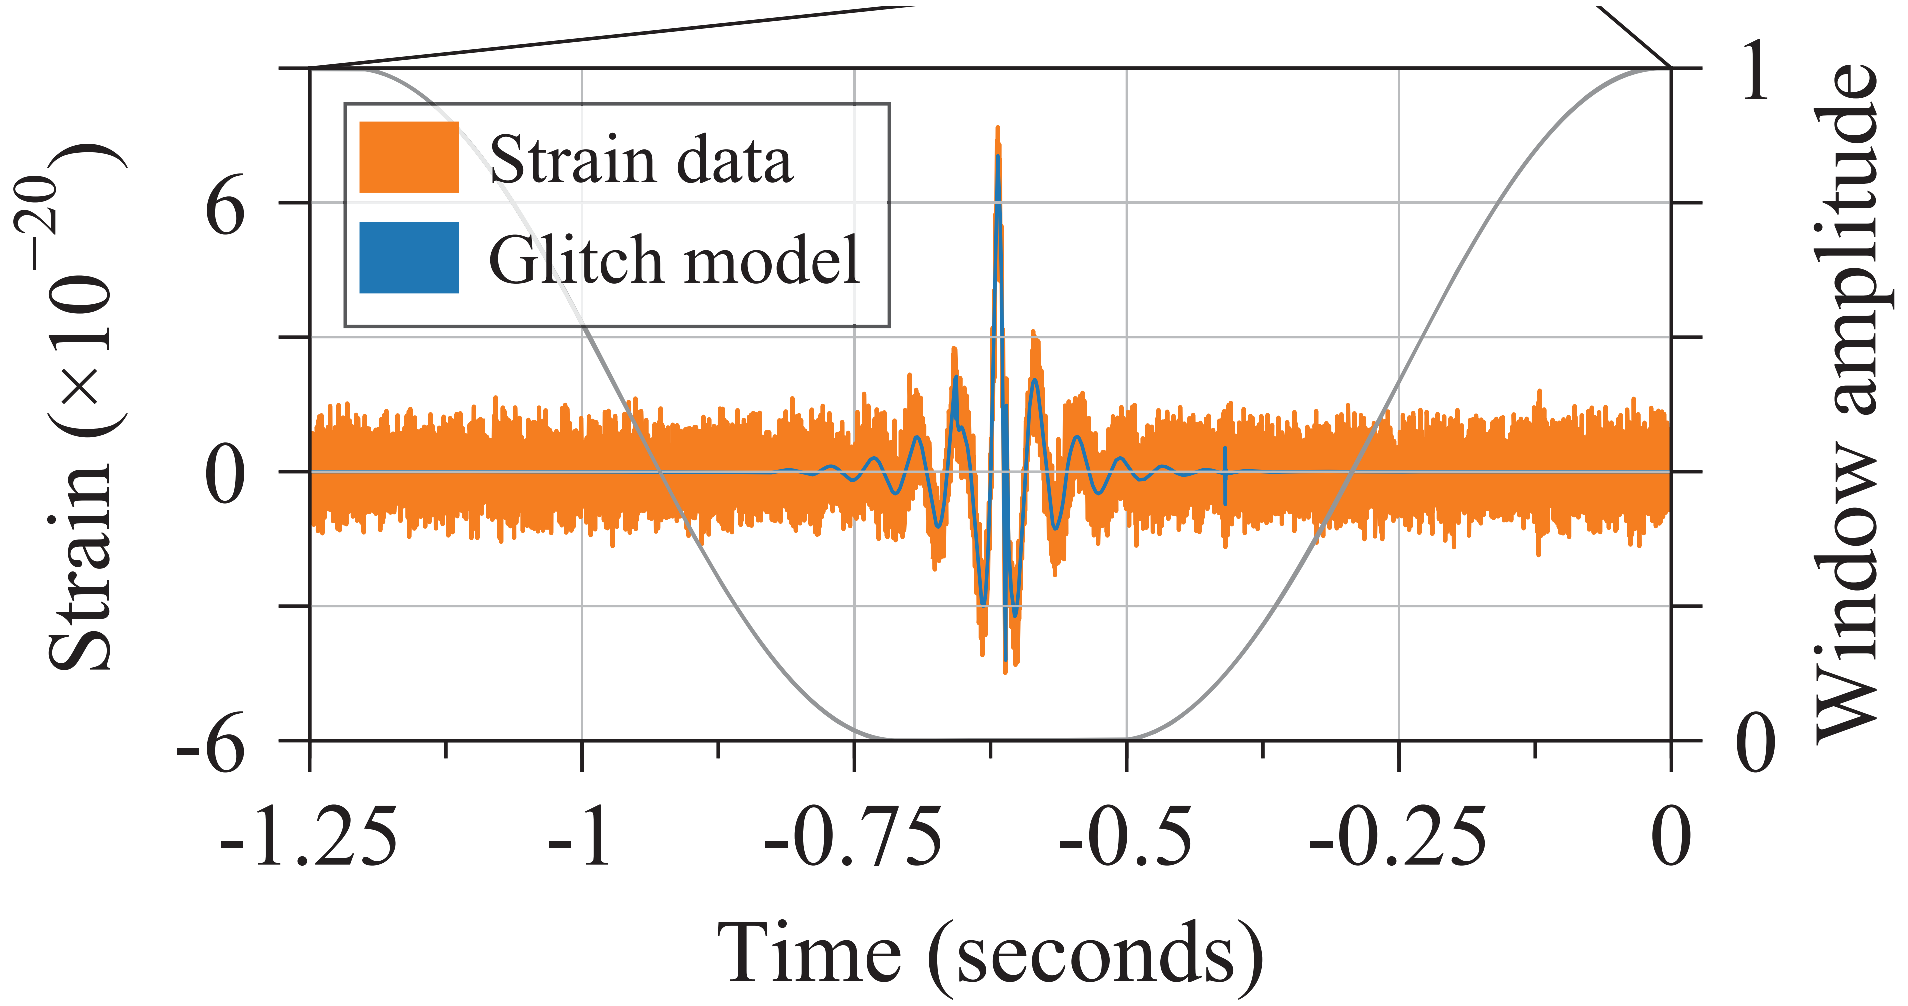

Example: GW170817 overlapped with a blip glitch in LIGO-Livingston

A BETTER NOISE MEASUREMENT

- **BayesWave** [Cornish+]: Obtain a PSD estimate plus uncertainty from the data segment that contains the signal via a transdimensional Bayesian analysis that simultaneously models the GW signal, the noise (PSD) and non-stationarities (glitches):

$$d(t) = h(t) + n(t) + g(t)$$

  - Uses Morlet-Gabor wavelets to represent the GW signal and glitches.
  - Uses a flexible spline model to model the PSD.
  - Uses reversible-jump MCMC to sample the posteriors for all three simultaneously.

- Pros:
  - Fully Bayesian, on-source, flexible, non-parametric, RJMCMC avoids overfitting

- Cons:
  - Computationally expensive, can struggle with long/complex signals, sensitive to priors, over- or undersmoothing of lines

STATIONARITY

- The noise floor in GW detectors varies over time
  - $S_n(f)$ changes between events
  - Compute the PSD for every event to take PSD drift into account

- Note: Uncertainty in the PSD estimation can also be taken into account via marginalisation by introducing extra parameters into the noise model

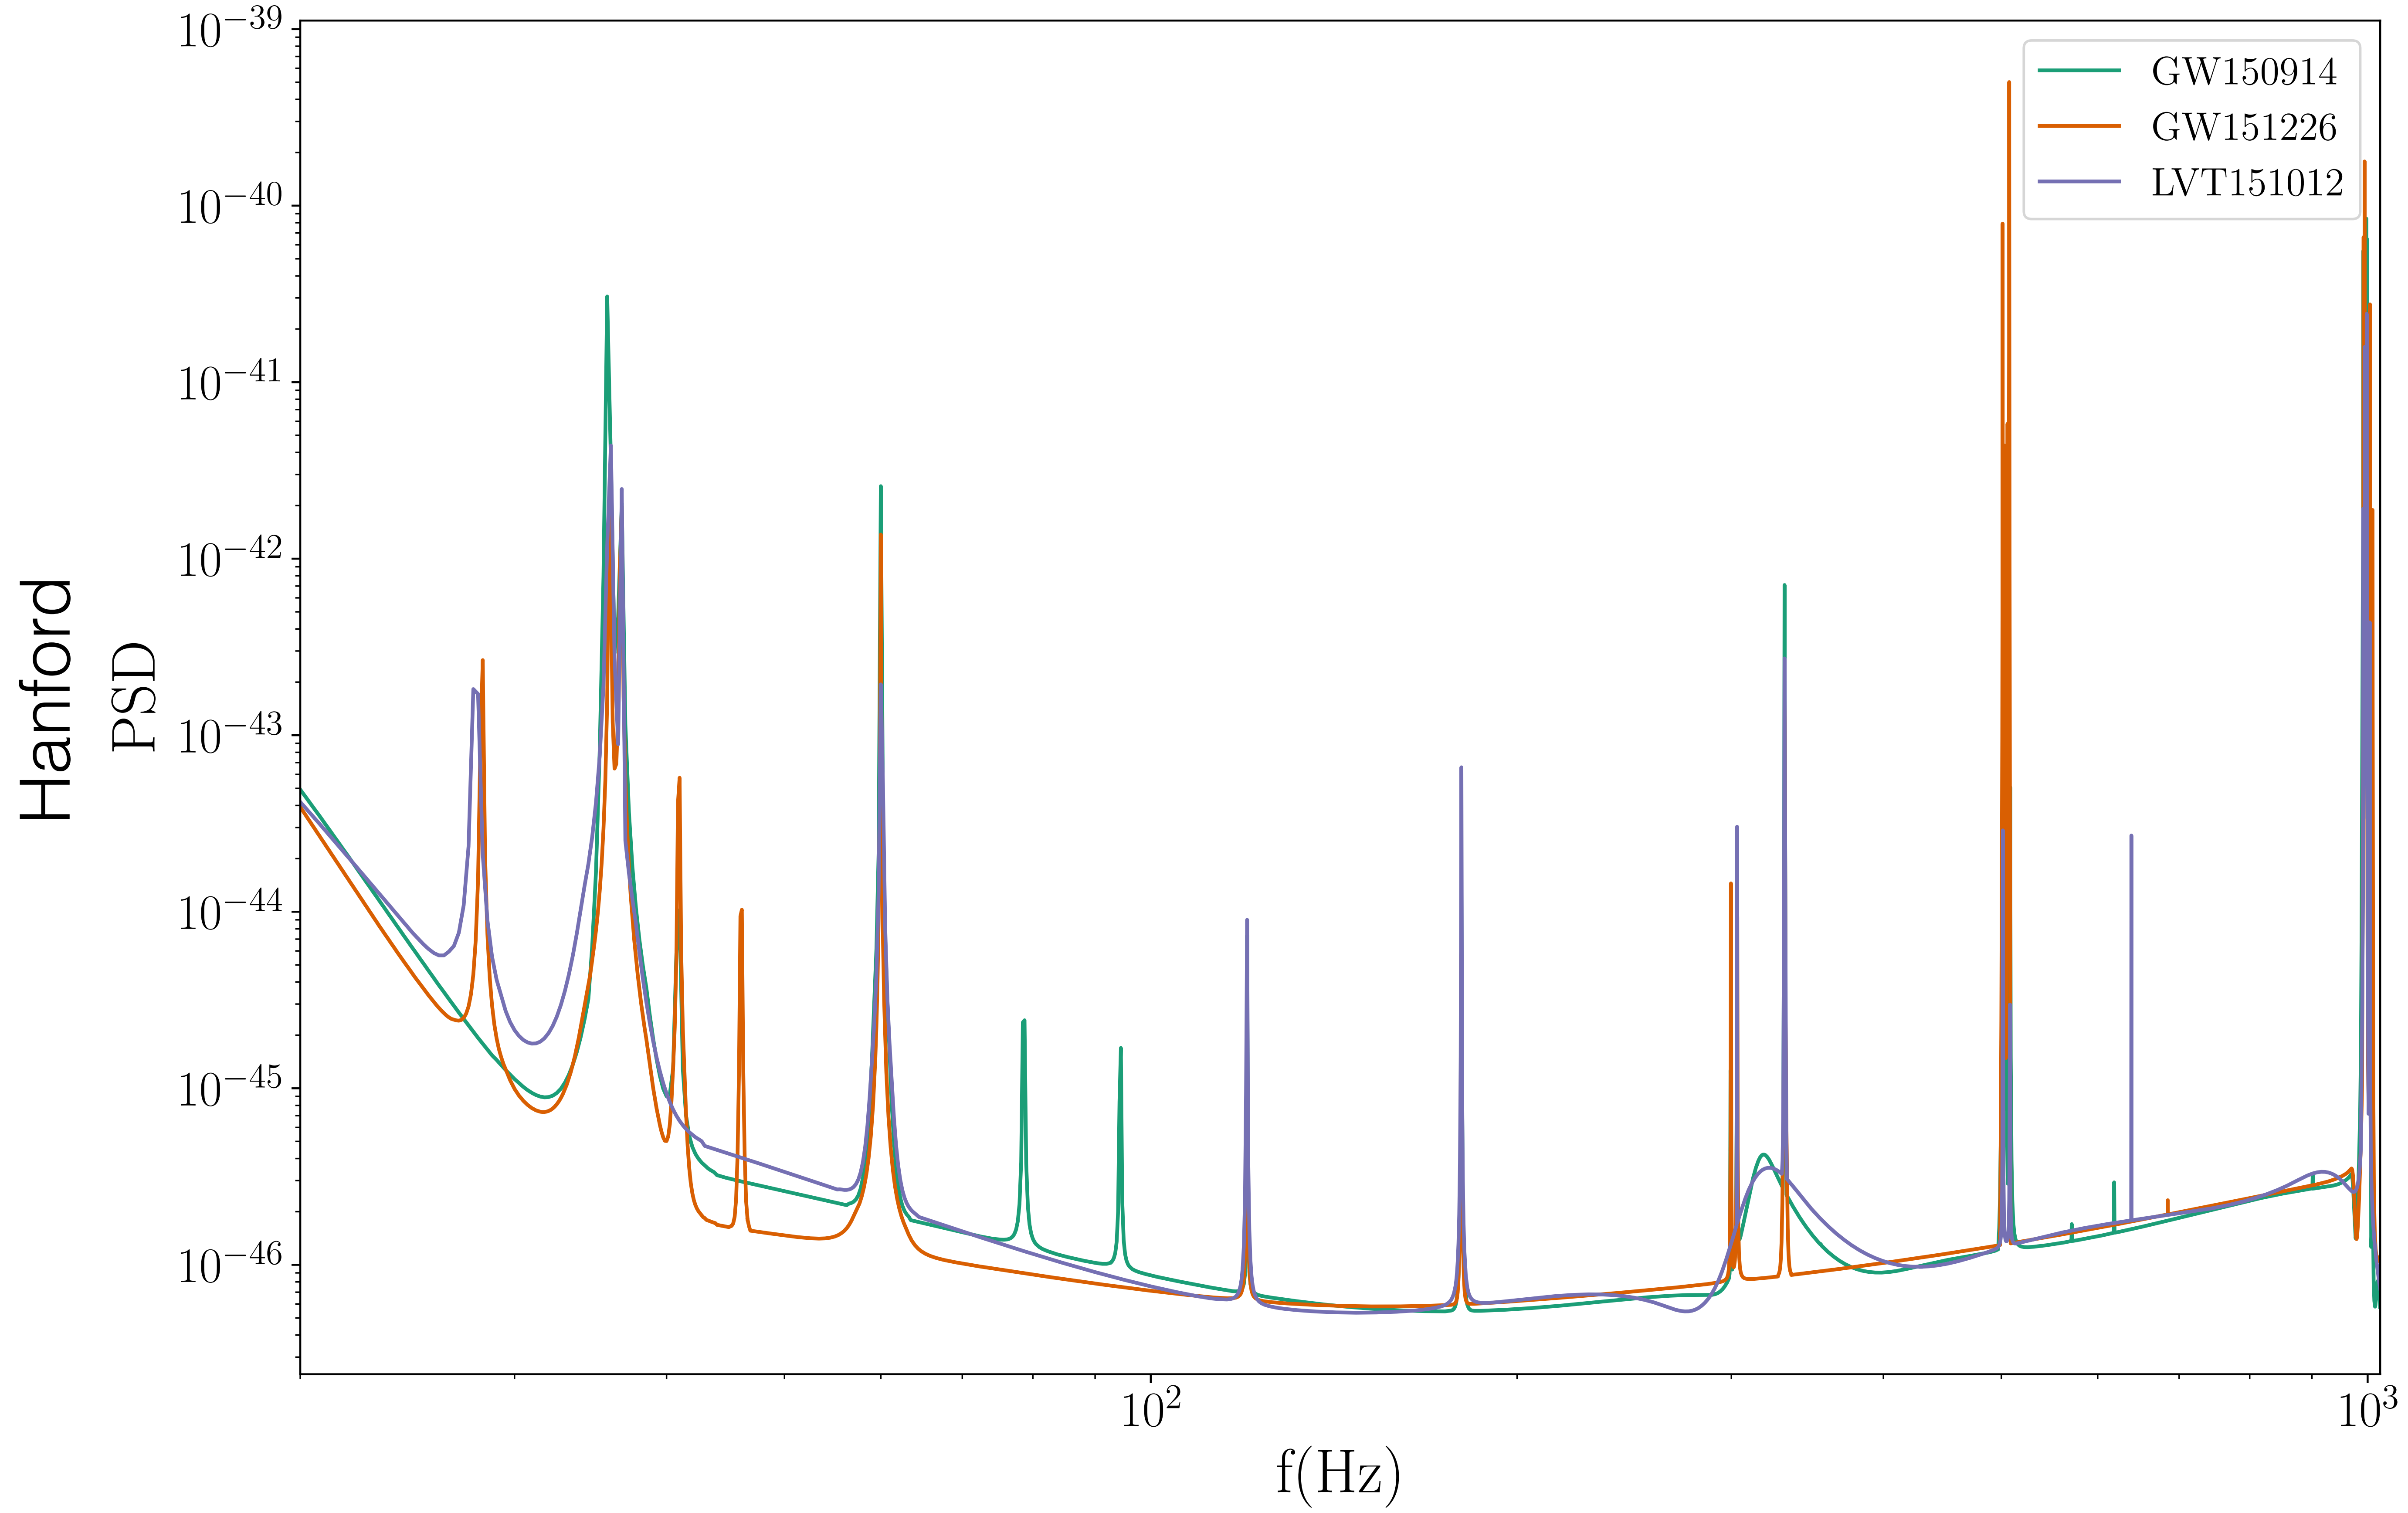

Credit: K. Chatziioannou

WHAT ABOUT CALIBRATION?

- Calibration quantifies the detector's response to incident GWs
  - Miscalibration results in biased strain data!
- Parameterised model for calibration uncertainty
  - Marginalised over during parameter estimation
- Calibration uncertainties are not the limiting factor in current GW observations
  - Statistical uncertainty from detector noise dominates

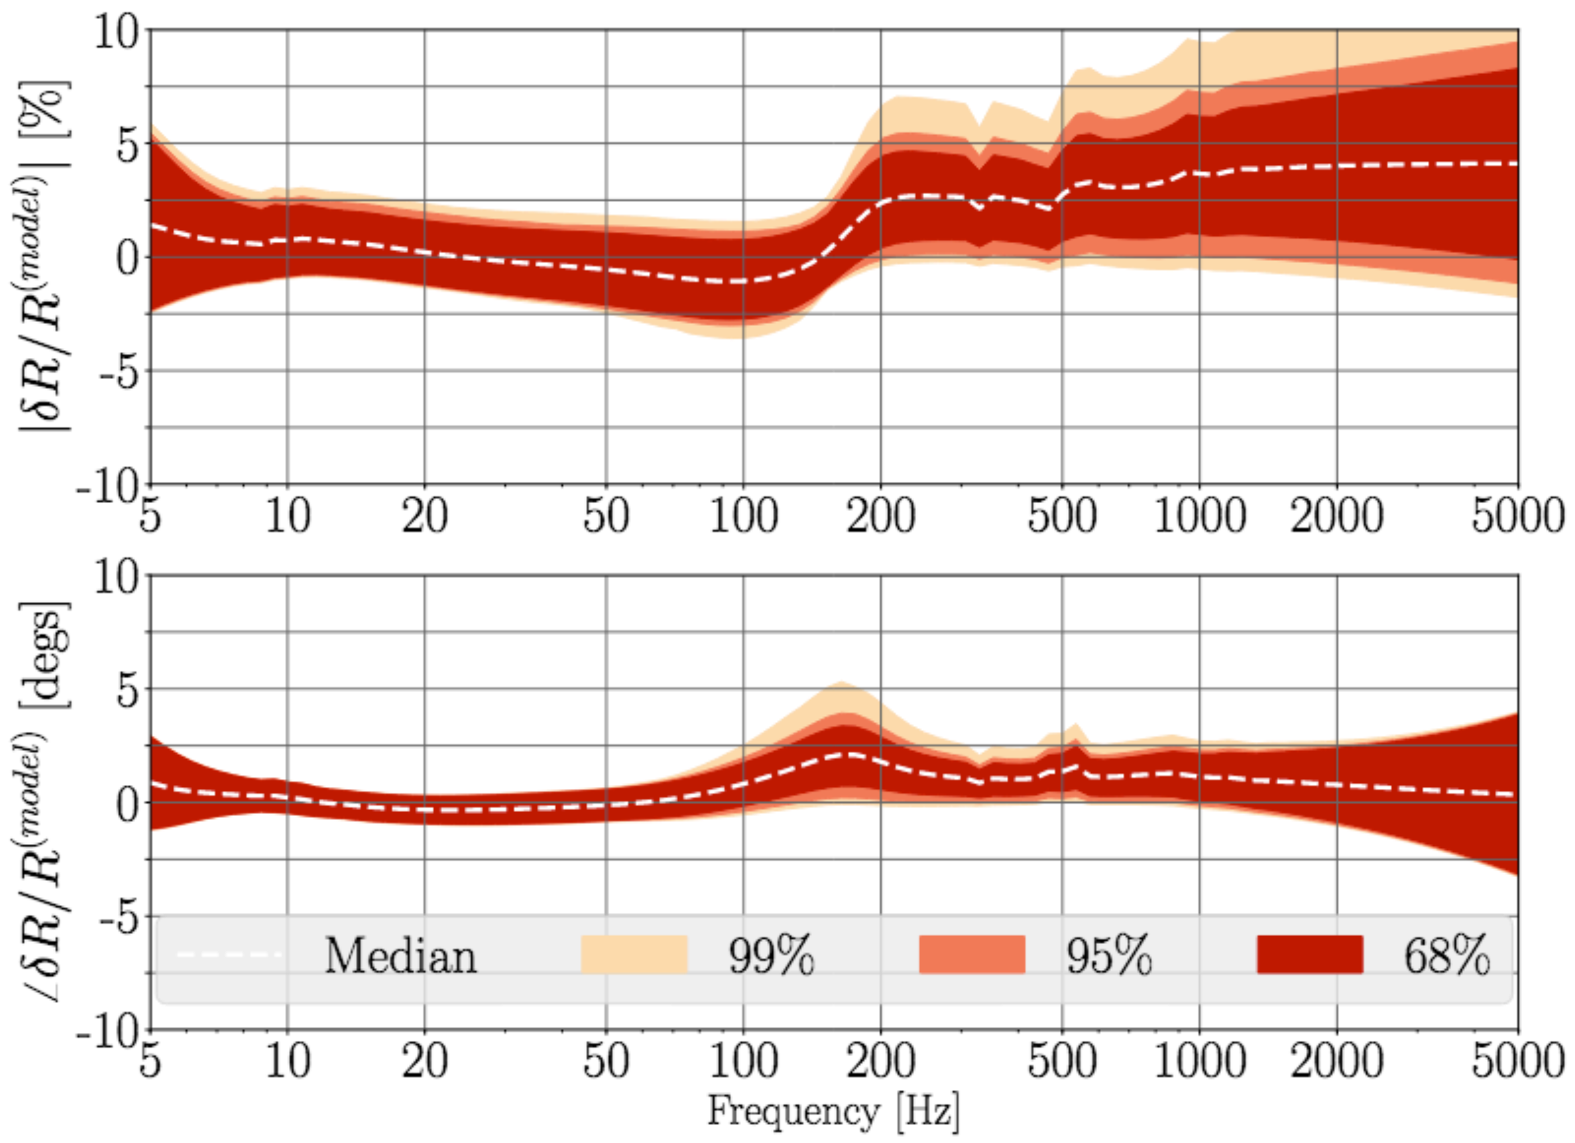

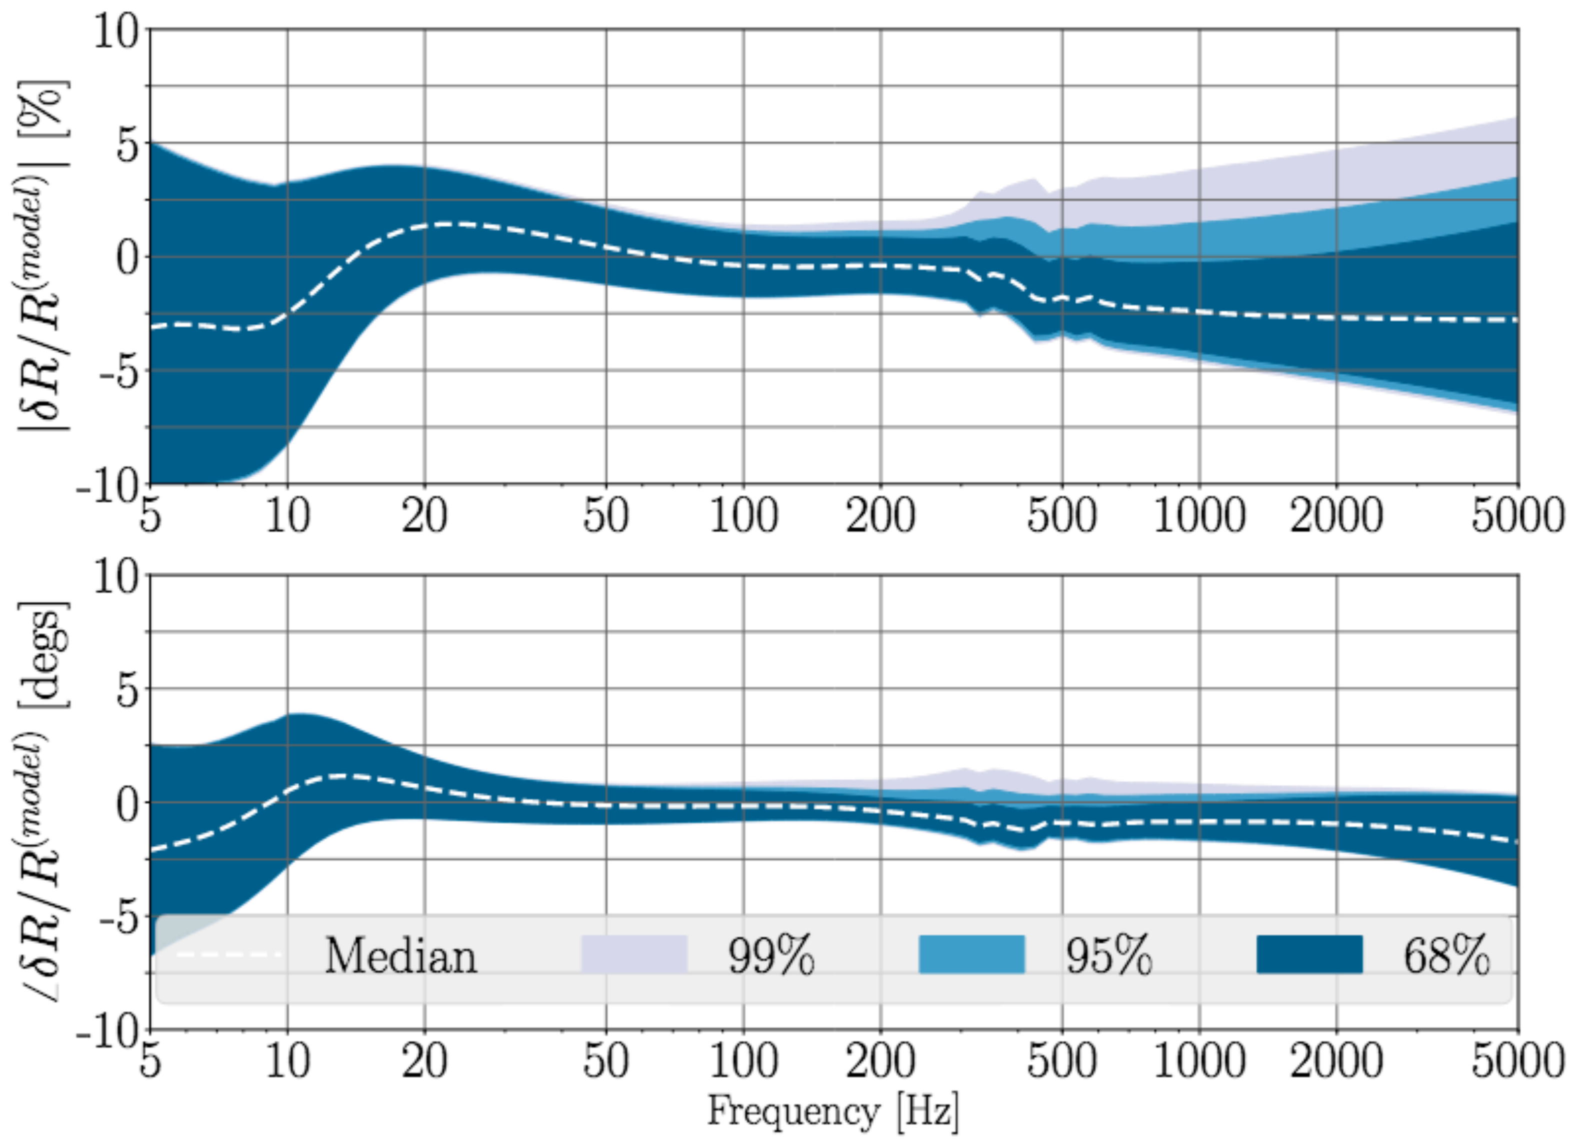

[Cahillane+, 2017]]

THE PRIOR

- Uninformative vs. astrophysically motivated priors
    - Population priors
- "Standard" GW priors:
    - Mass priors: often sample uniform in chirp mass $\mathcal{M}_c = \eta^{3/5}M$ & symmetric mass ratio $\eta = (m_1m_2)/M$
    - Distance: $p(D_L) \propto D_L^2$ (uniform in "luminosity" volume)
        - Or uniform in comoving volume
    - Spins magnitudes: uniform
    - Spin orientation: isotropic (i.e. uniform in $\cos(\theta)$)

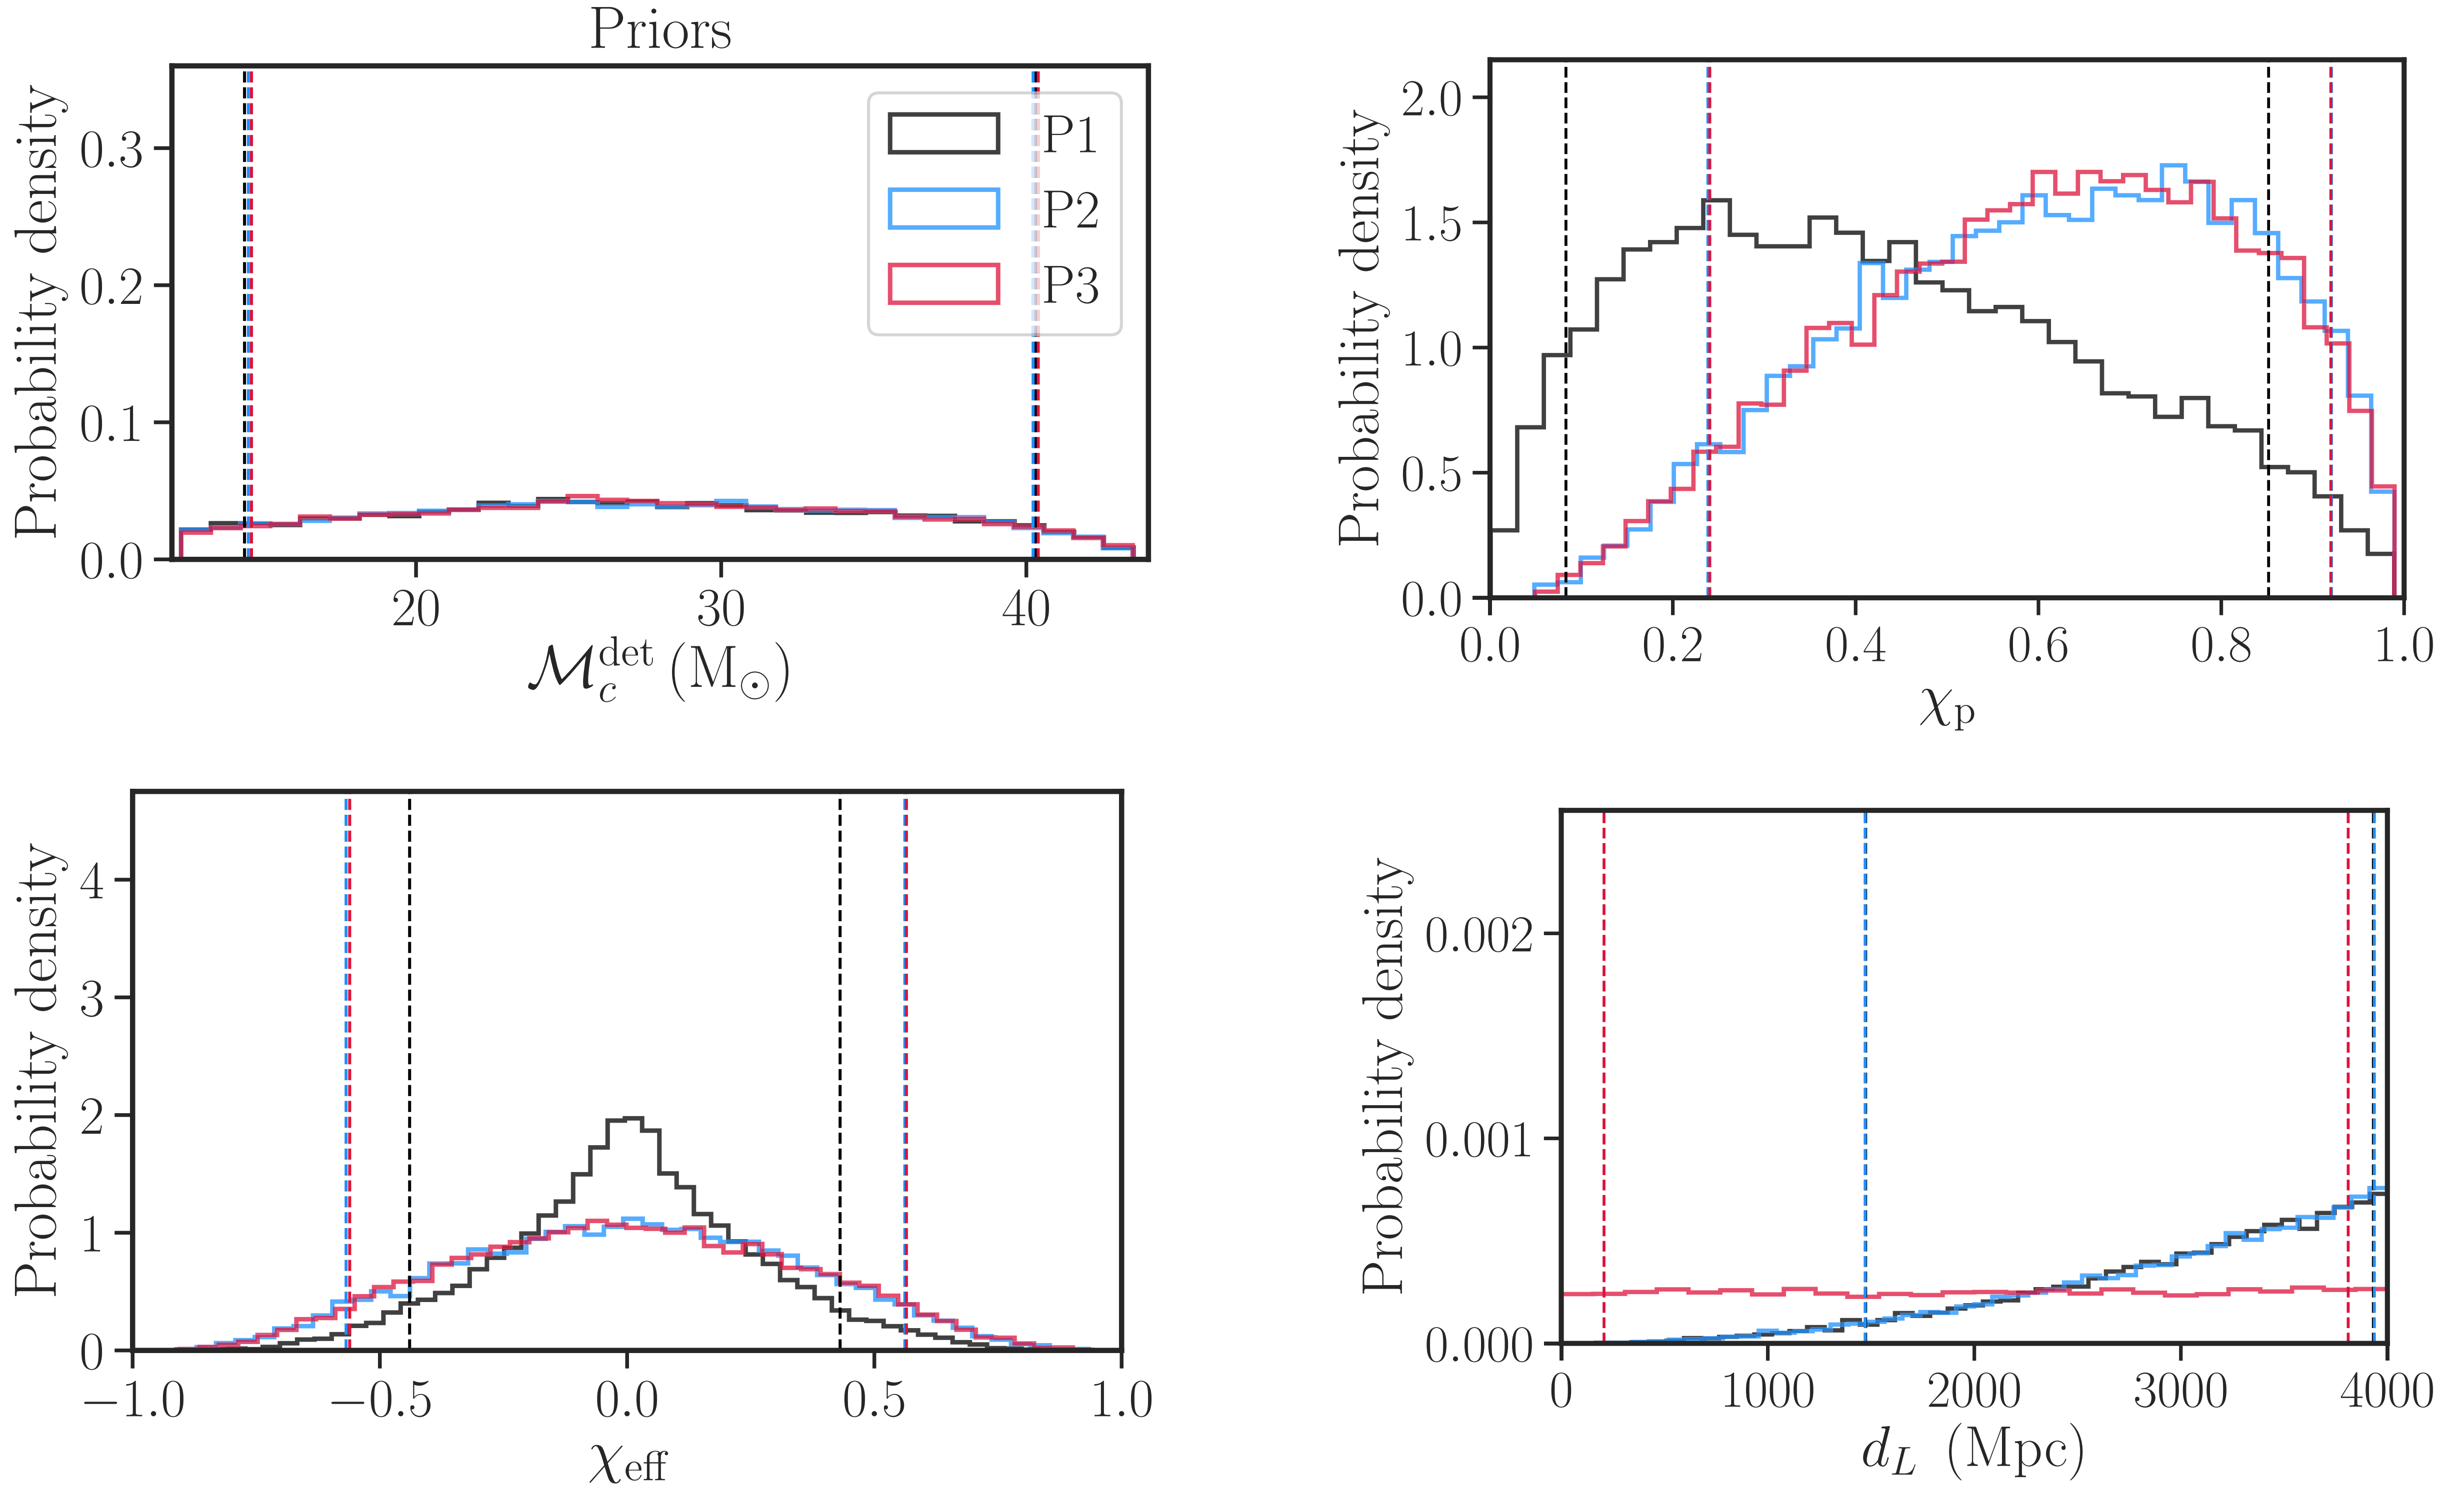

[LVC, GWTC-1]

SAMPLING THE POSTERIOR: MCMC VS. NESTED SAMPLING

MCMC

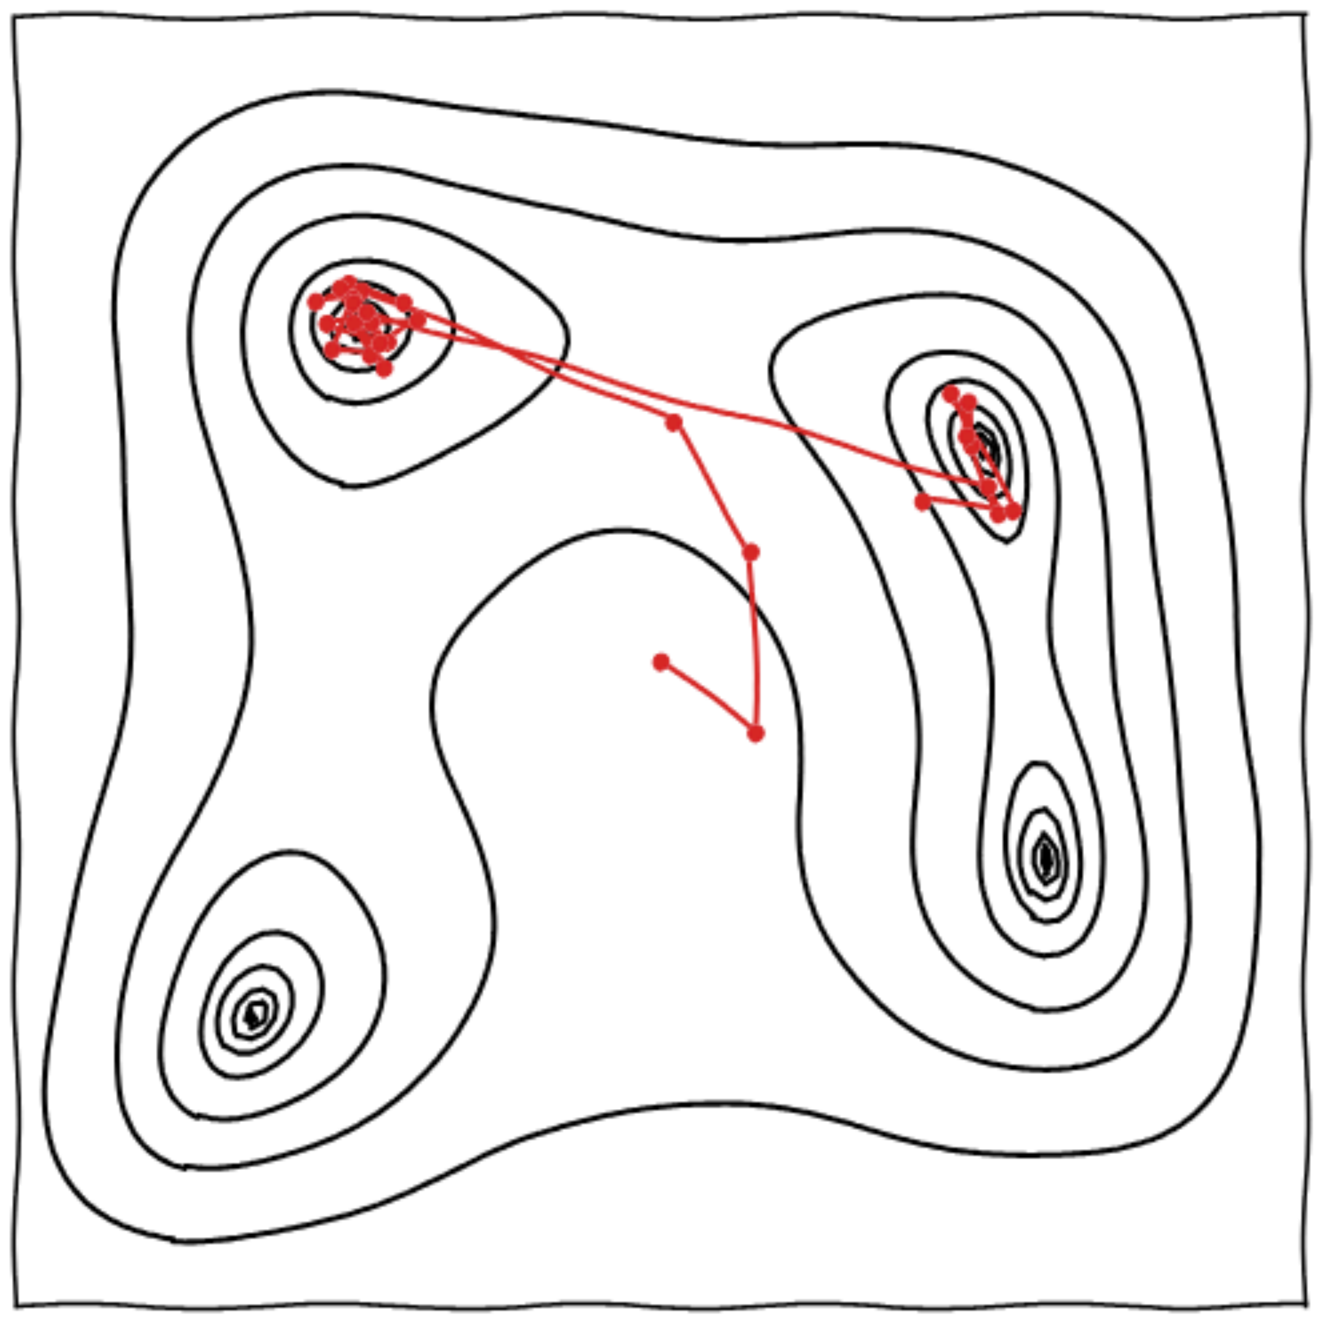

- Approximate the posterior distribution
- Starting with a “walker” drawn from the prior
- Performs a random walk to explore the posterior
- Next step is rejected or accepted depending on the likelihood ratio

- Can be very fast but jump proposal needs to be carefully tuned
- Can be problematic for multimodal distributions

[Adapted from slides by Will Handley]

SAMPLING THE POSTERIOR: MCMC VS. NESTED SAMPLING

- Compute the evidence integral by shrinking the prior volume X:

$$\mathcal{Z} = \int d^n\theta\, p(d|\theta)p(\theta) \Rightarrow \mathcal{Z} = \int_0^1 \mathcal{L}(X)dX$$

- Draw N live points uniformly distributed in prior space
- Start from the lowest likelihood point to draw an iso-likelihood surface
- Draw a new live point from the prior and keep if its likelihood is higher
- Accumulate the evidence to reconstruct the posterior

- Can be slow but well-suited for multimodal distributions

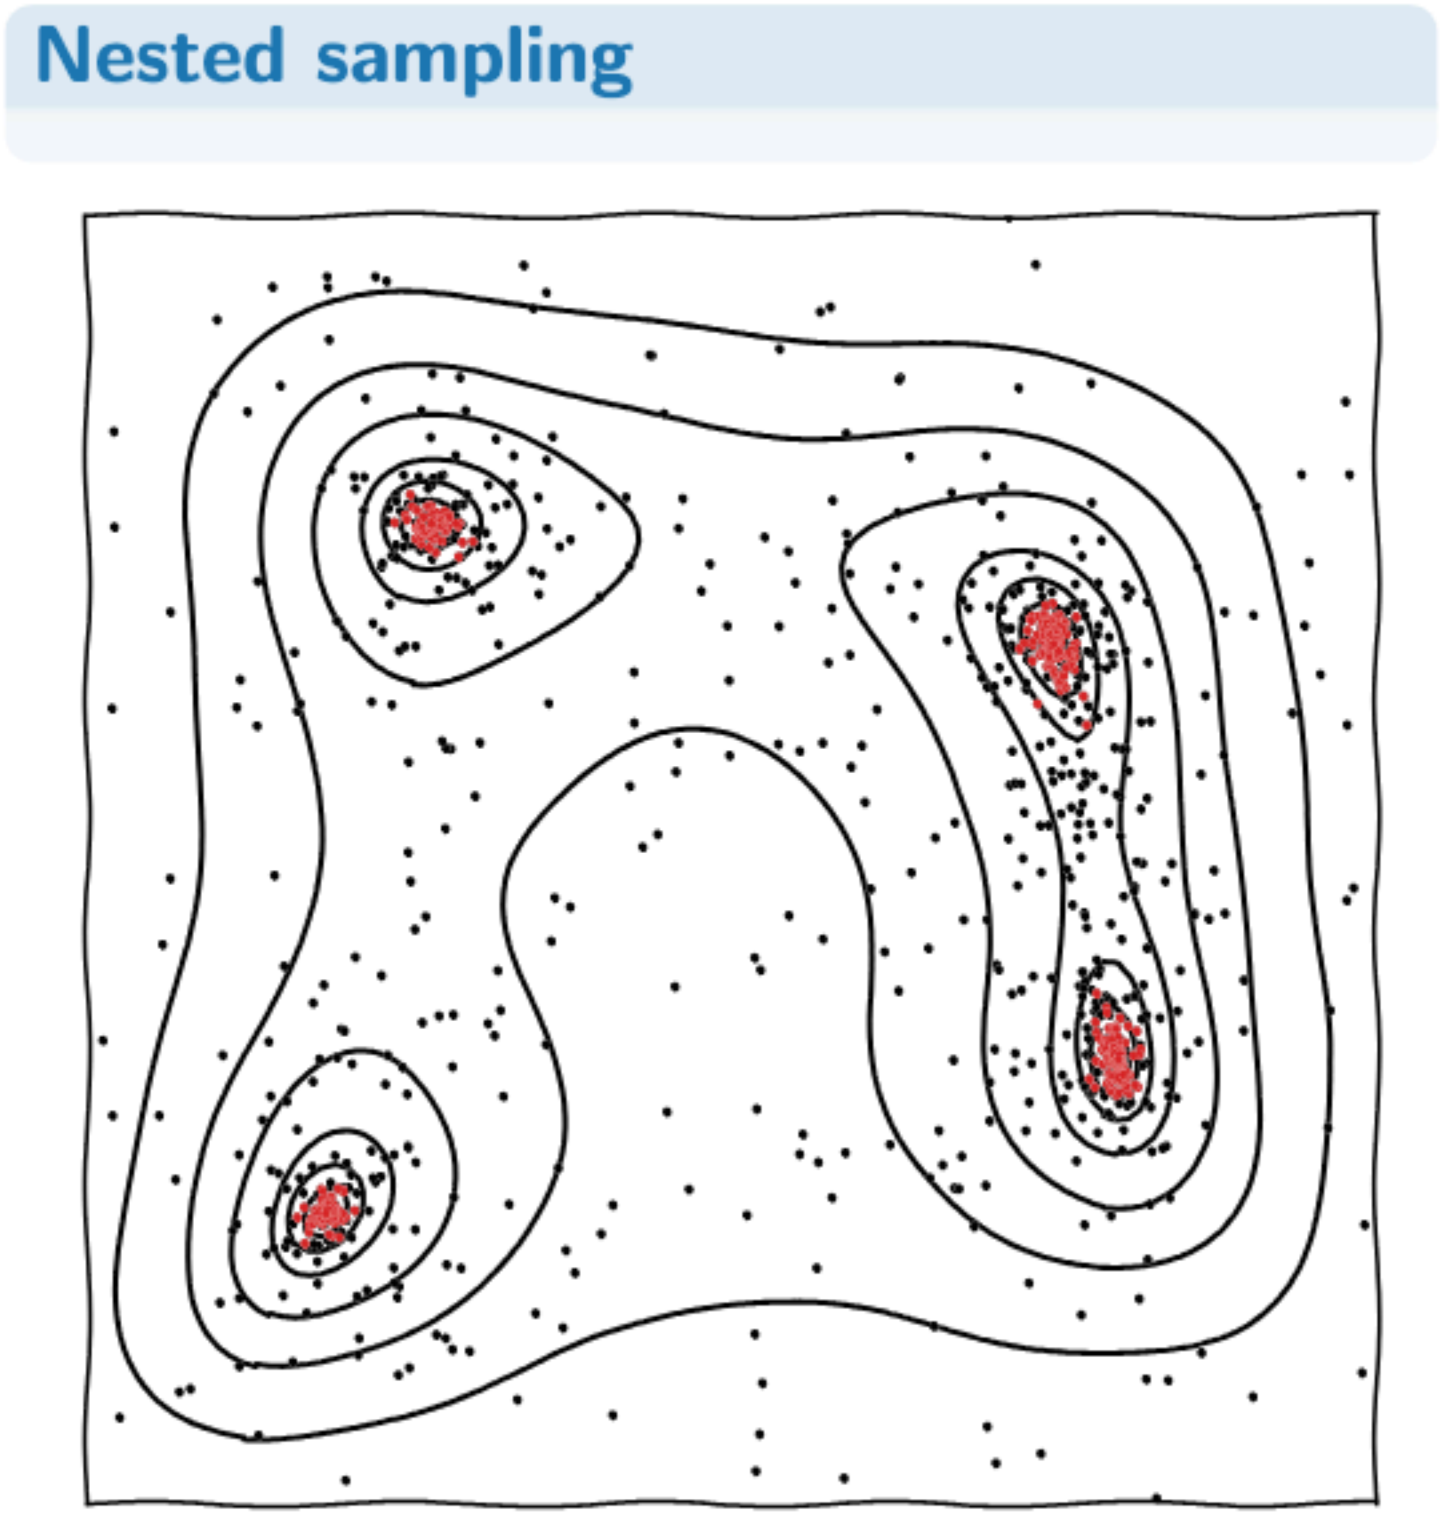

[Adapted from slides by Will Handley]

DO IT YOURSELF

- All LVK analysis software is publicly available: https://git.ligo.org/lscsoft/lalsuite
- Tutorials available from the Gravitational Wave Open Science Centre: https://gwosc.org/tutorials/
- Perform a matched filter search: https://pycbc.org/
- User-friendly Python inference package Bilby:
  - Source code: https://git.ligo.org/lscsoft/bilby
  - Documentation: https://lscsoft.docs.ligo.org/bilby/

WAVEFORMS

WHAT WE CAN MEASURE

- Waveform written as:

$$h(v) = \mathcal{A}(v)\, e^{i\phi(v)}$$

Phase

Amplitude

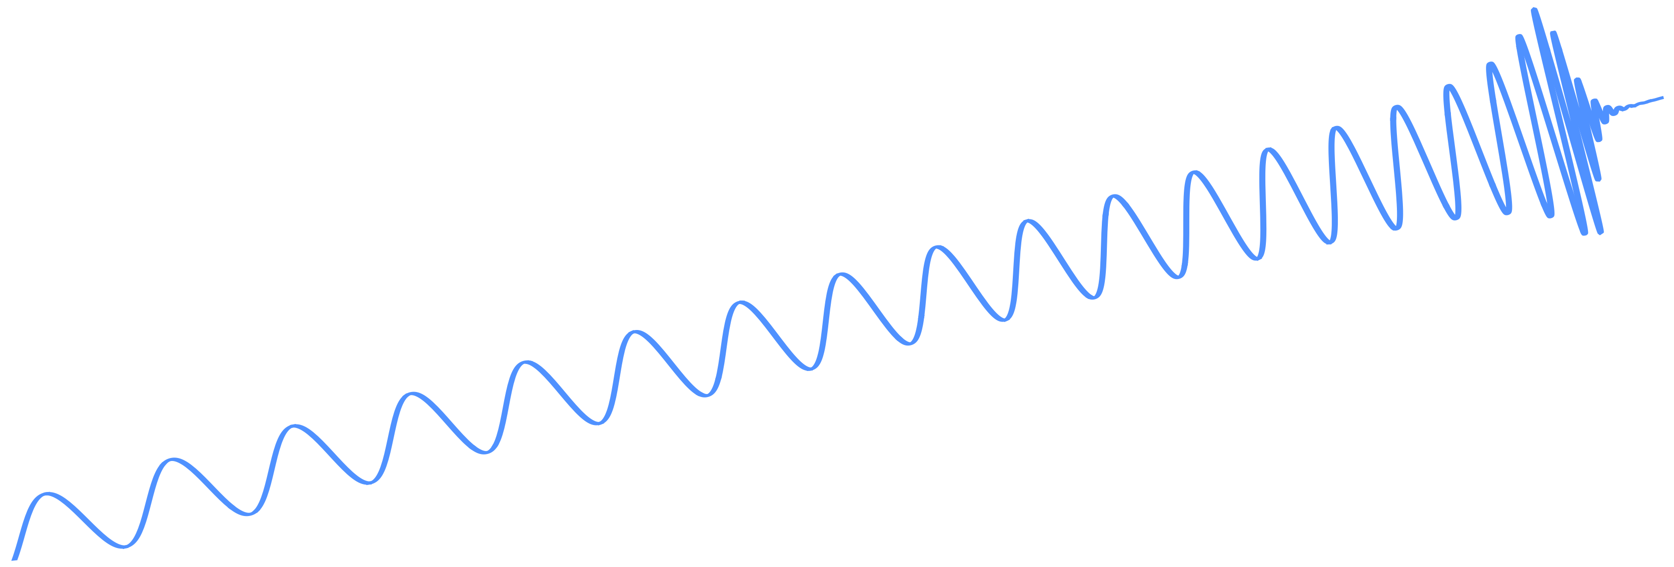

- The **GW phase** encodes the physics:

0PN 1PN 2.5PN[l] 5PN

$$\phi(v) = \left(\frac{v}{c}\right)^{-5}\left[\phi_0 + \phi_2\left(\frac{v}{c}\right)^2 + \cdots + \phi_{5l}\ln\left(\frac{v}{c}\right)\left(\frac{v}{c}\right)^5 + \cdots + \phi_{10}\left(\frac{v}{c}\right)^{10} + \cdots\right]$$

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1+m_2)^{1/5}} \qquad q = \frac{m_2}{m_1}$$

Spin effects at 1.5PN

Tail terms ~ back reaction due to scattering

Tidal effects first enter here…

[Slide credit: G. Pratten]

WHAT CAN WE TYPICALLY MEASURE?

- Mass parameters:
  - Low-mass binaries: chirp mass
  - High-mass binaries: total mass

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}$$

- Effective spin parameters:
  - Effective inspiral spin:

$$\chi_{\rm eff} = \frac{(m_1 \vec{\chi}_1 + m_2 \vec{\chi}_2) \cdot \hat{L}_N}{M}$$

  - Effective precession spin:

$$\chi_p := \frac{\max(A_1 S_{1\perp}, A_2 S_{2\perp})}{A_1 m_1^2}$$

- Binary tidal deformability:

$$\tilde{\Lambda} = \frac{16}{13} \frac{(m_1 + 12 m_2) m_1^4 \Lambda_1 + (m_2 + 12 m_1) m_2^4 \Lambda_2}{(m_1 + m_2)^5}$$

- Improved detectors = higher SNRs will allows to better constrain individual parameters

ON THE IMPORTANCE OF THE SIGNAL MODEL

- For compact binary mergers, we need signal models that are valid for the inspiral, the merger and the ringdown (**IMR**) and ideally incorporate all the physics.
  - Spin precession, higher-order modes, eccentric/hyperbolic/scattering orbits, unequal masses, environmental effects etc.
  - Everything, everywhere all at once: Missing physics/inaccurate models lead to **systematic biases** in the inference!

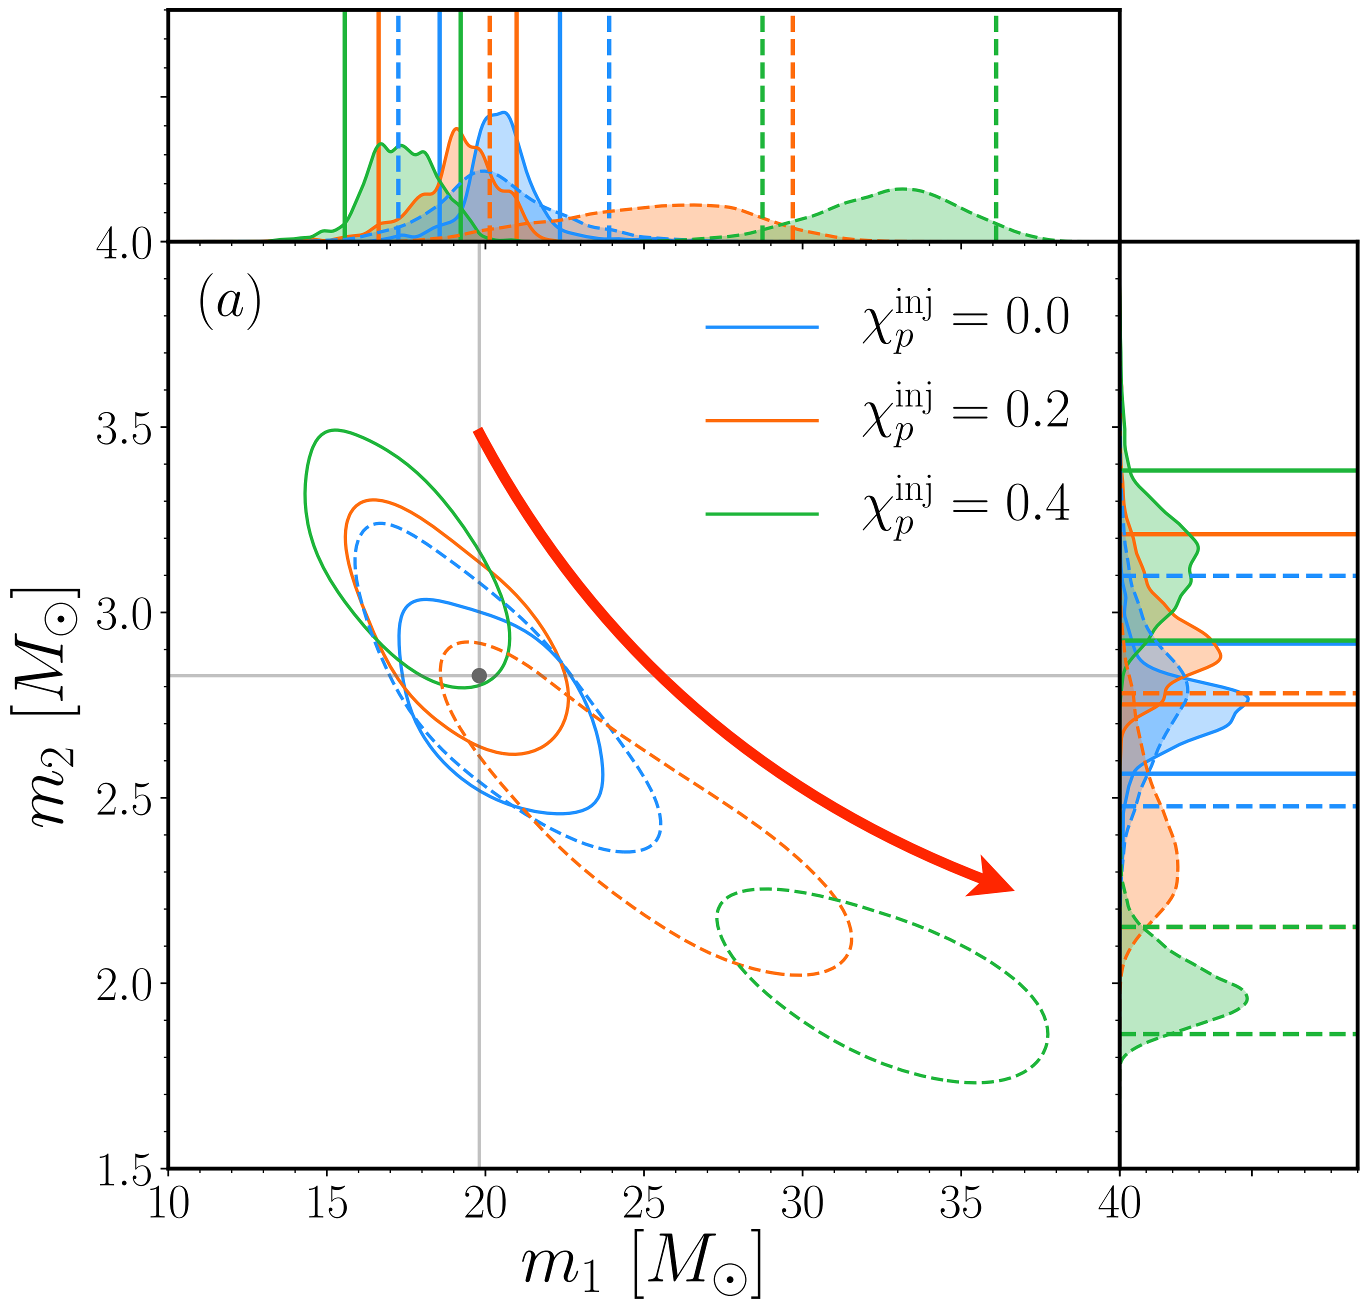

[Pratten+, 2020]

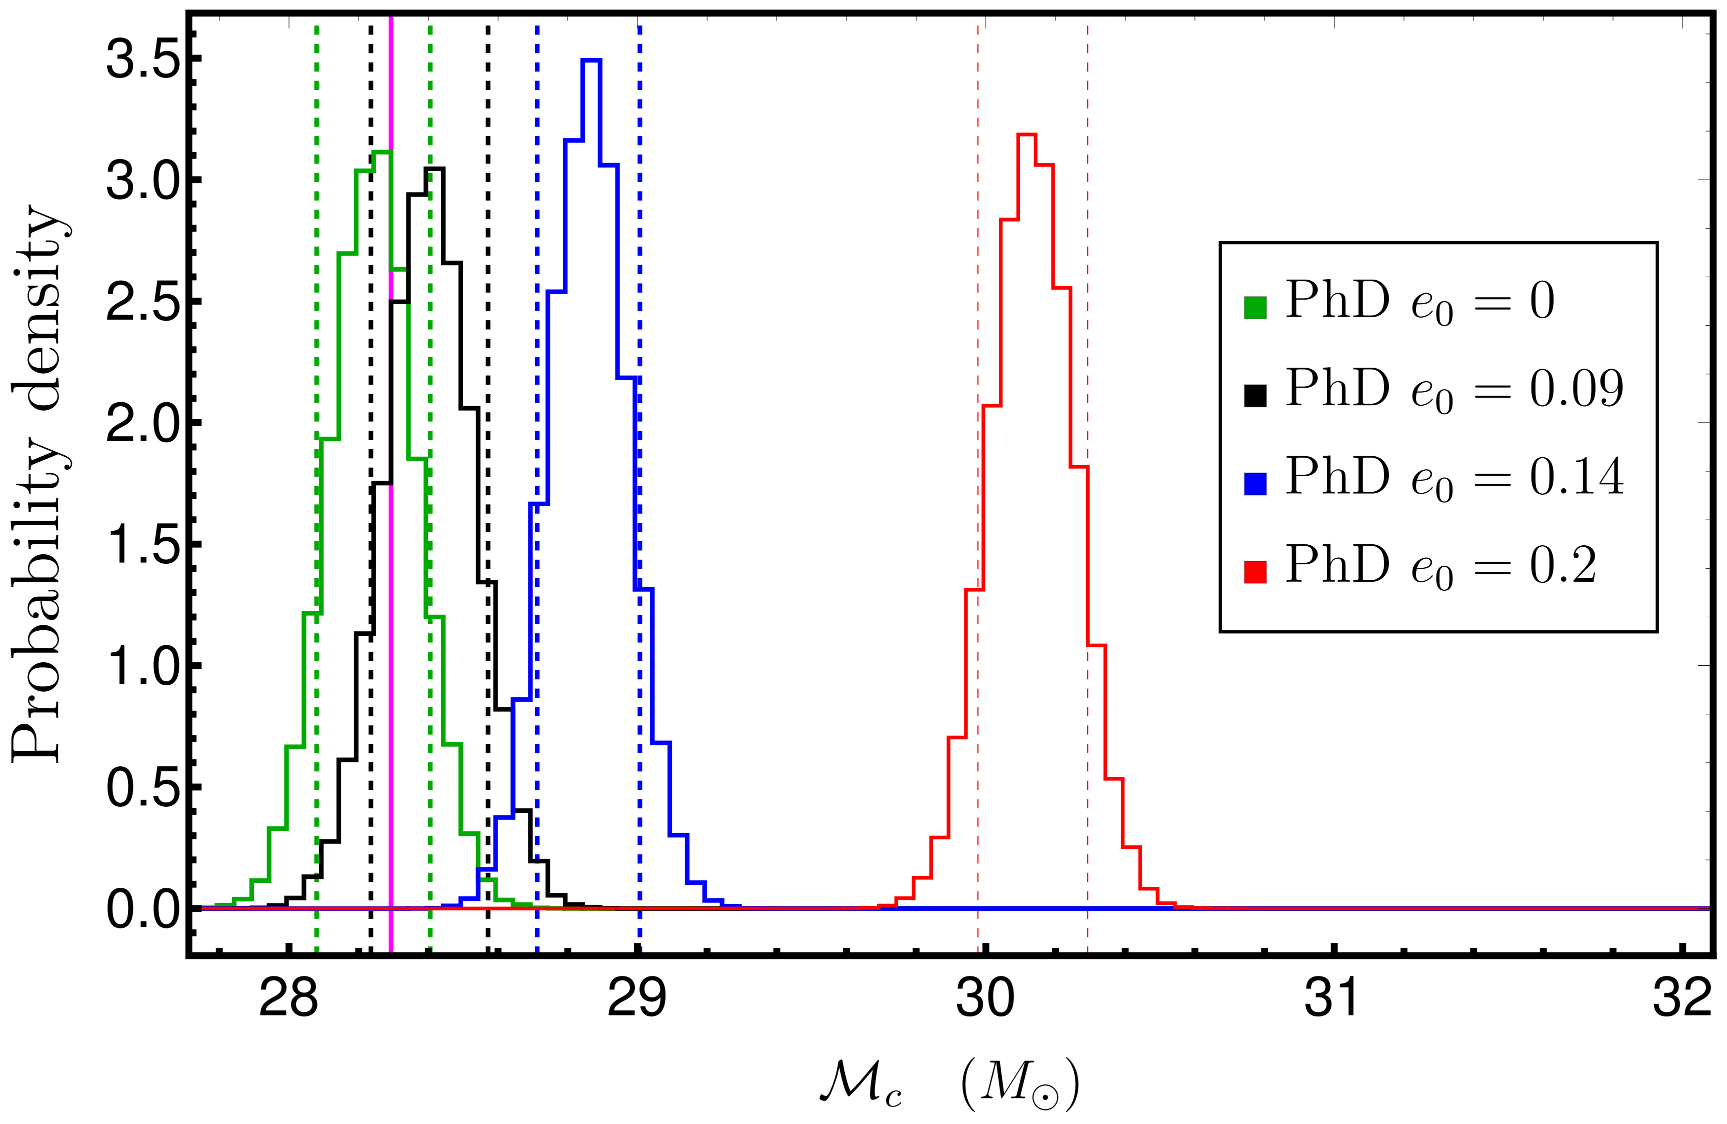

[Ramos-Buades+, 2020]

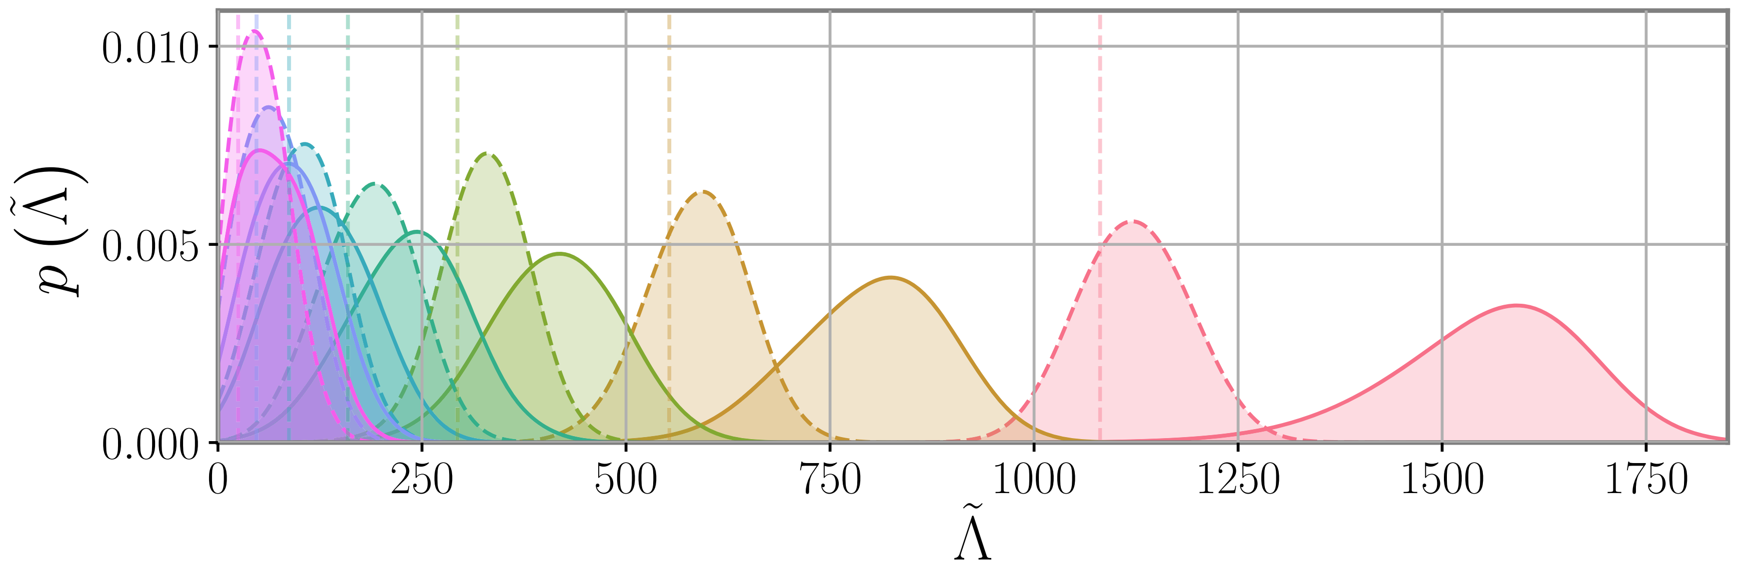

[Pratten+, 2022]

ON THE IMPORTANCE OF THE SIGNAL MODEL

- In practice, we use **different approaches** to construct full IMR waveform models:
  - **EOBNR:** NR-Calibrated Effective-One-Body models
  - **Phenom:** NR-Calibrated Phenomenological waveform models
- We also use "partial" waveform models for certain mass regions:
  - Pure PN waveform models for low-mass inspirals
  - **NR surrogate models** for high-mass mergers
  - (Calibrated) BH perturbation theory waveform for ringdown

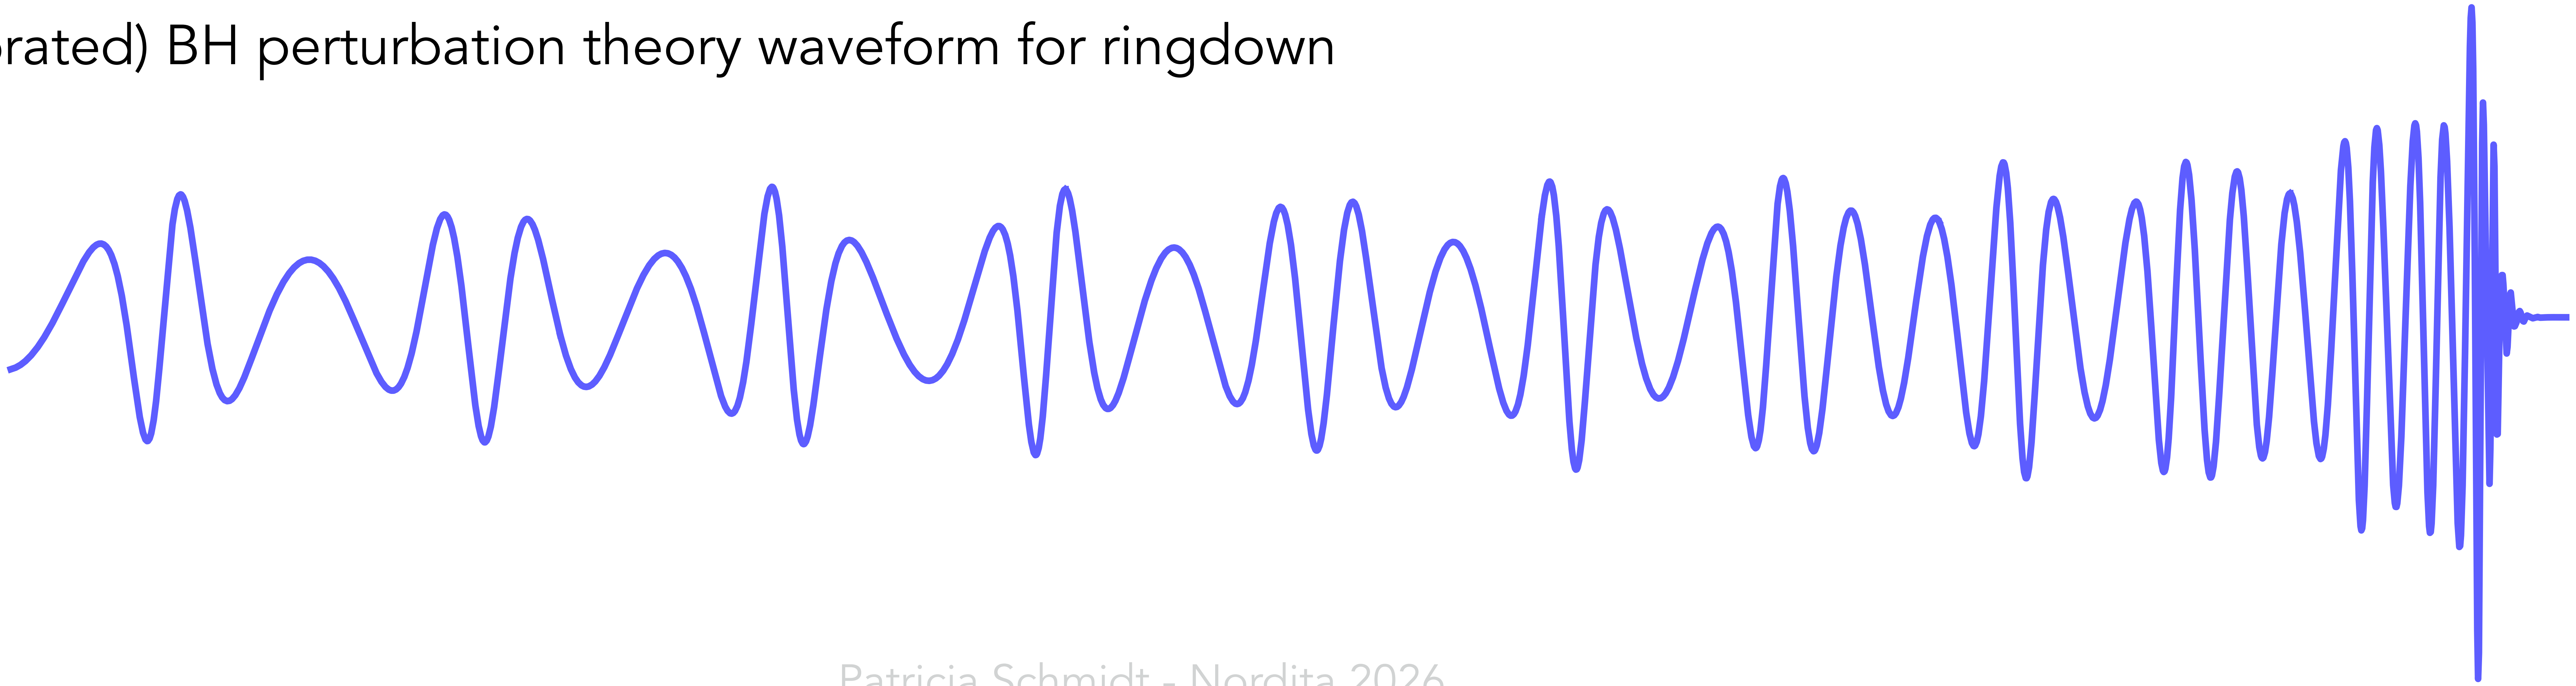

PHENOMENOLOGICAL WAVEFORM MODELS

- Original motivations:
  - Data analysis predominantly performed in the frequency domain. Fourier transforms are costly.
  - Traditional Bayesian inference requires 10^6 to 10^8 likelihood evaluations.
- Require **accurate and computationally efficient** waveform models.
- Phenom models combine:
  - Analytical information from PN, GSF, EOB and
  - Non-perturbative information from NR in the strong-field regime

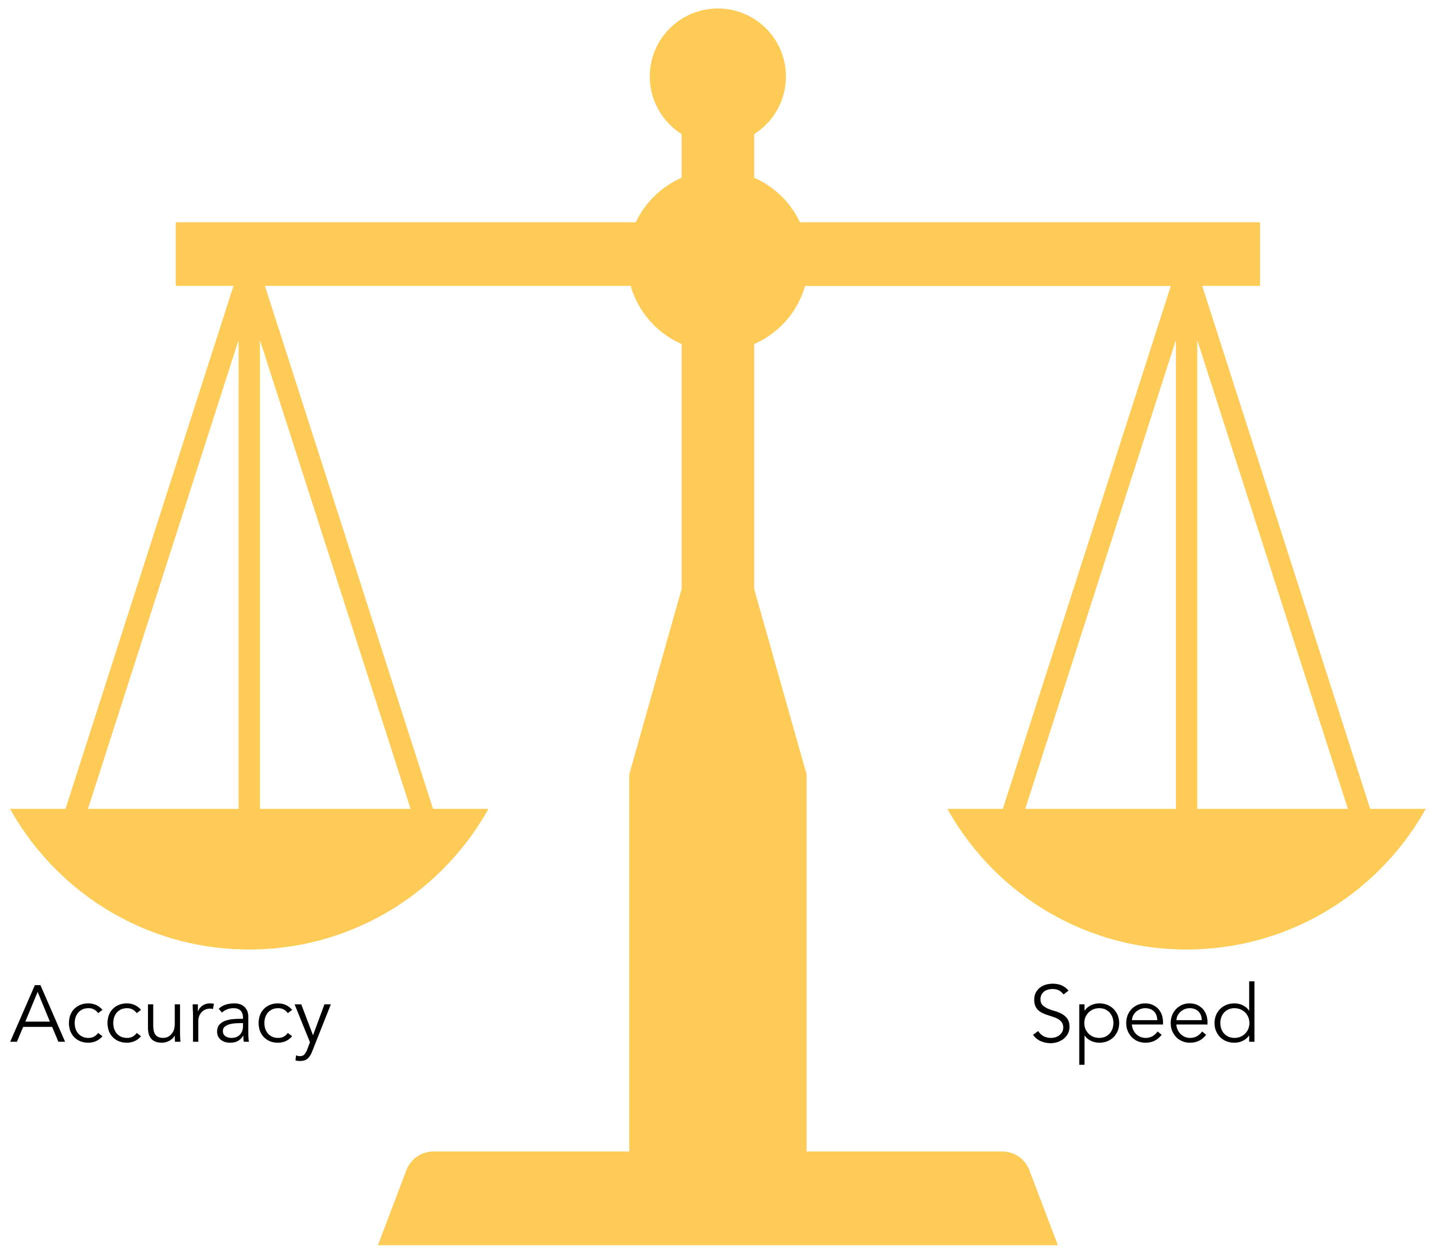

THE NEED FOR SPEED

- Timing comparisons between different flavours of Phenom waveform models and EOB models:

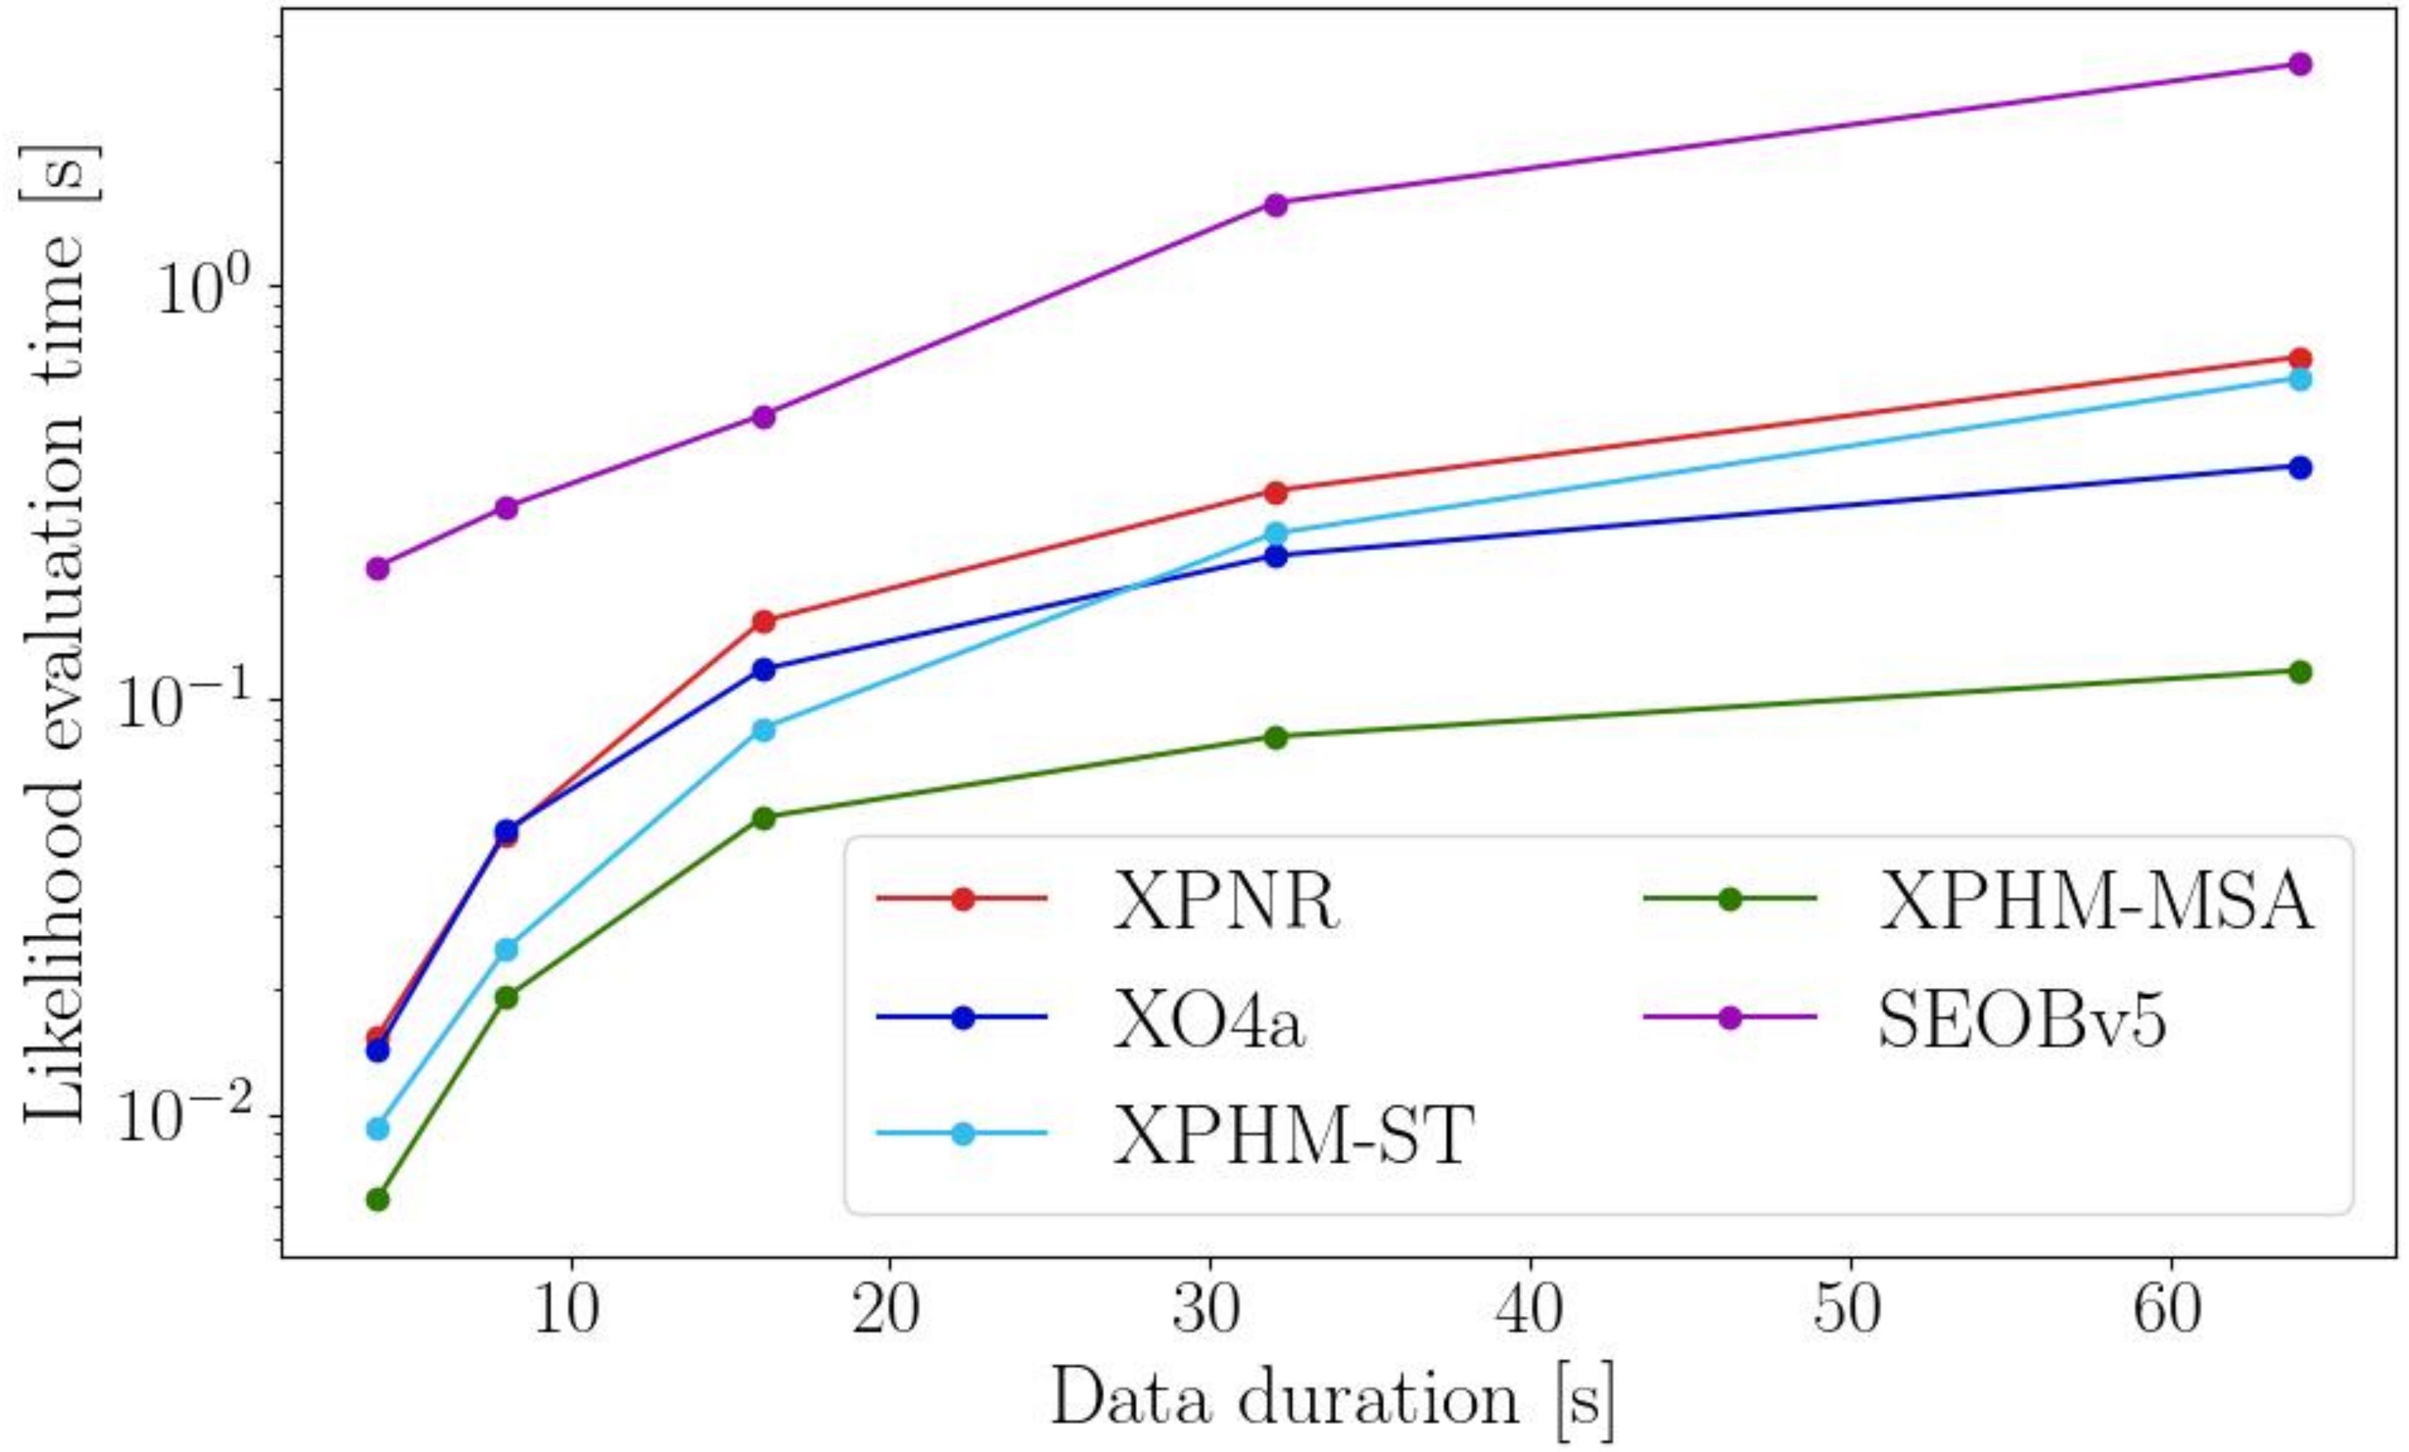

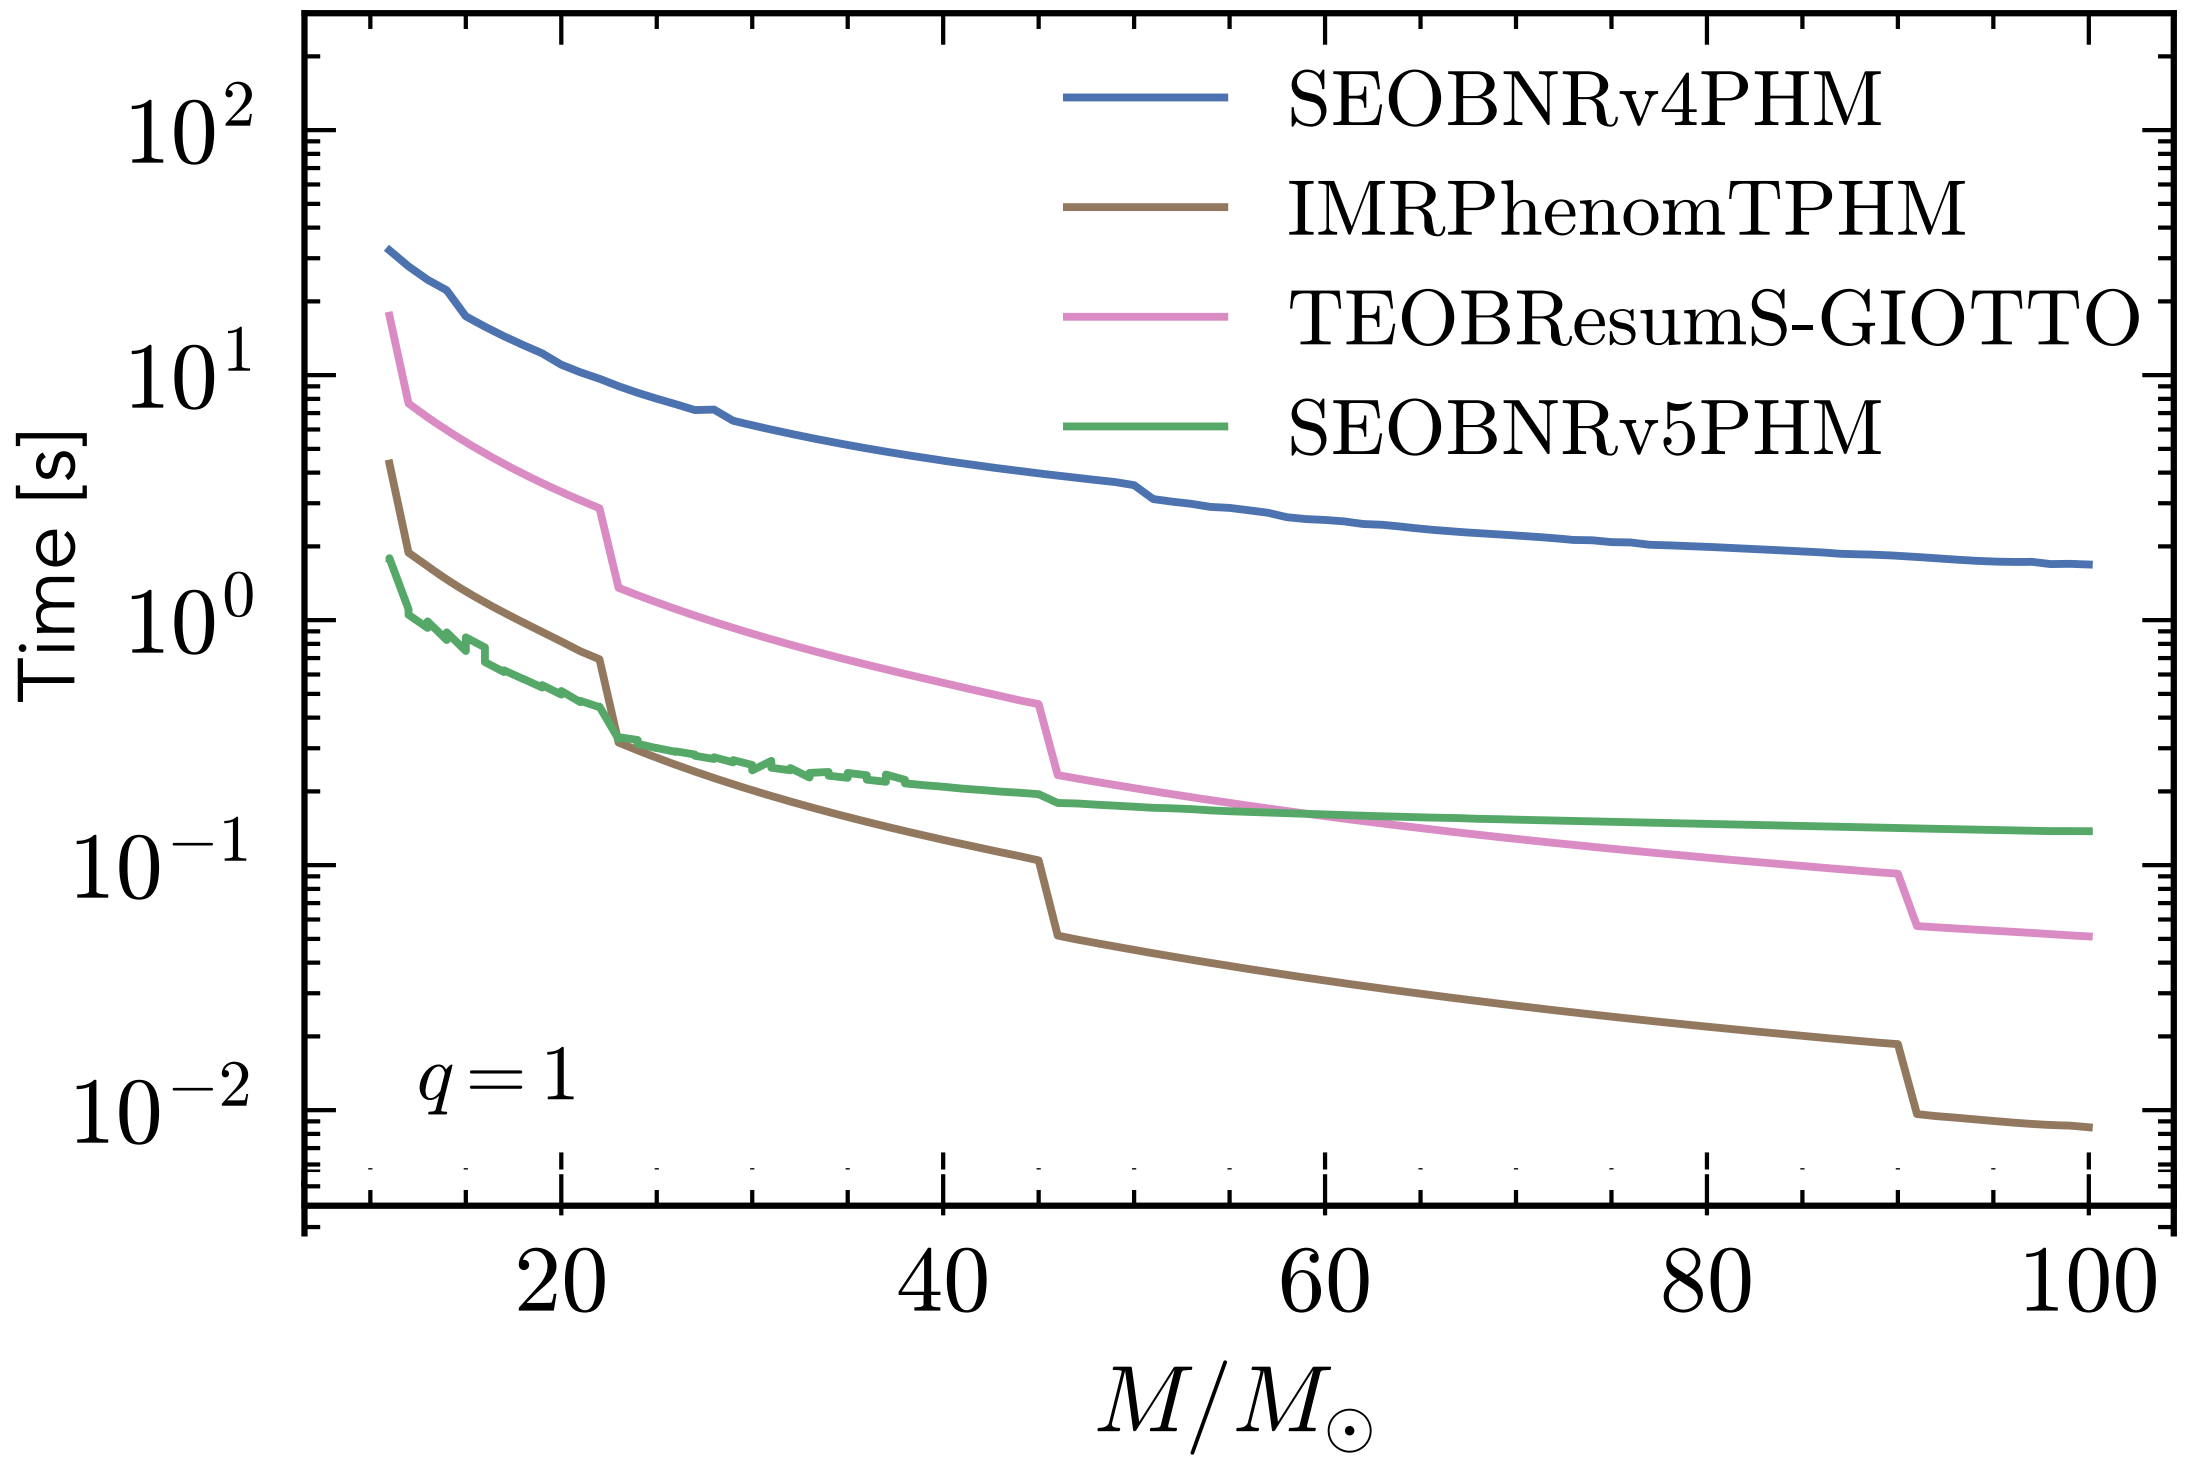

[Hamilton+, 2025]

[Ramos-Buades+, 2023]

PHENOMENOLOGICAL WAVEFORM MODELS

- Traditionally built in the frequency domain via a mode-by-mode ansatz:

$$\tilde{h}_{\ell m}(f;\theta) = \tilde{A}_{\ell m}(f;\theta)e^{-i\tilde{\phi}\ell m(f;\theta)}$$

- Example: IMRPhenomXAS [Pratten+]

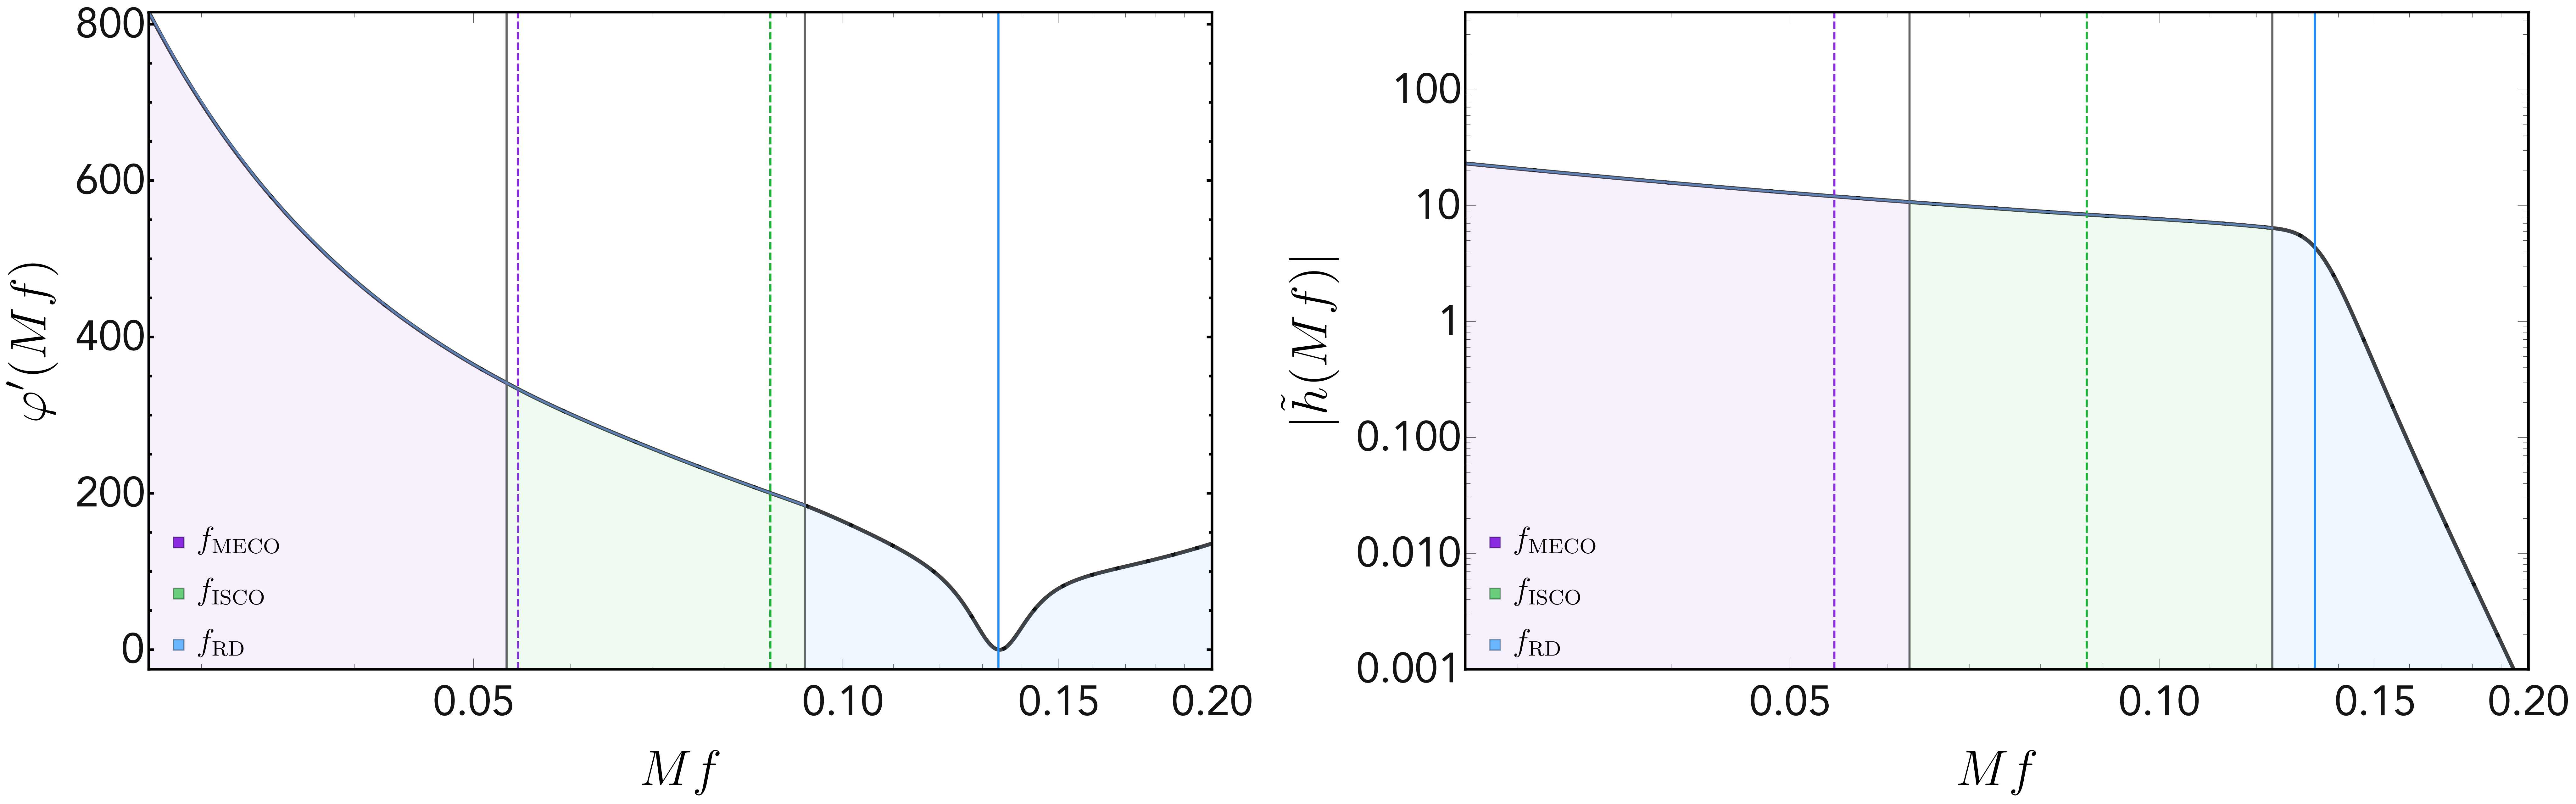

PHENOMENOLOGICAL WAVEFORM MODELS

- **Inspiral phase**: Combine known PN/EOB with with NR information

$$\phi_{\rm Ins} = \phi_{\rm TF2}(Mf;\vec{\theta}) + \frac{1}{\eta}\left(\sigma_0 + \sigma_1 f + \frac{3}{4}\sigma_2 f^{4/3} + \frac{3}{5}\sigma_3 f^{3/5} + \frac{1}{2}\sigma_4 f^2\right)$$

- **Intermediate region**: Use a polynomial ansatz to connect the inspiral to the ringdown

$$\eta\phi'_{\rm Int} = c_{\rm Int} + a_4 f^{-4} + a_2 f^{-2} + a_1 f^{-1} - \frac{4a_0 a_\phi}{(f - f_{\rm RD})^2 + (2f_{\rm damp})^2}$$

- **Ringdown**: Modelled as a deformed Lorentzian (inspired by perturbation theory)

$$\eta\phi'_{\rm RD} = c_{\rm RD} + \sum_i^n c_i f^{-p_i} + \frac{c_0 a_\varphi}{f_{\rm damp}^2 + (f - f_{\rm RD})^2}$$

PHENOMENOLOGICAL WAVEFORM MODELS

- Key ingredient: **Hybrid waveforms** = attach a PN/EOB inspiral to a set of NR simulations
  - Alignment (e.g. via least squares) such that $\phi_{\rm PN}(t_{\rm PN}) = \phi_{\rm NR}(t_{\rm NR}),\ \omega_{\rm PN}(t_{\rm PN}) = \omega_{\rm NR}(t_{\rm NR})$
  - "Blending" with transition functions, e.g.

$$\phi^{\rm hyb}(t) = \frac{\Delta t/2 - t}{\Delta t}\phi^{\rm PN}(t - t_{\rm PN}) + \frac{\Delta t/2 + t}{\Delta t}\phi^{\rm NR}(t - t_{\rm NR}) \ \text{ for } \ t \in [-\Delta t, \Delta t]$$

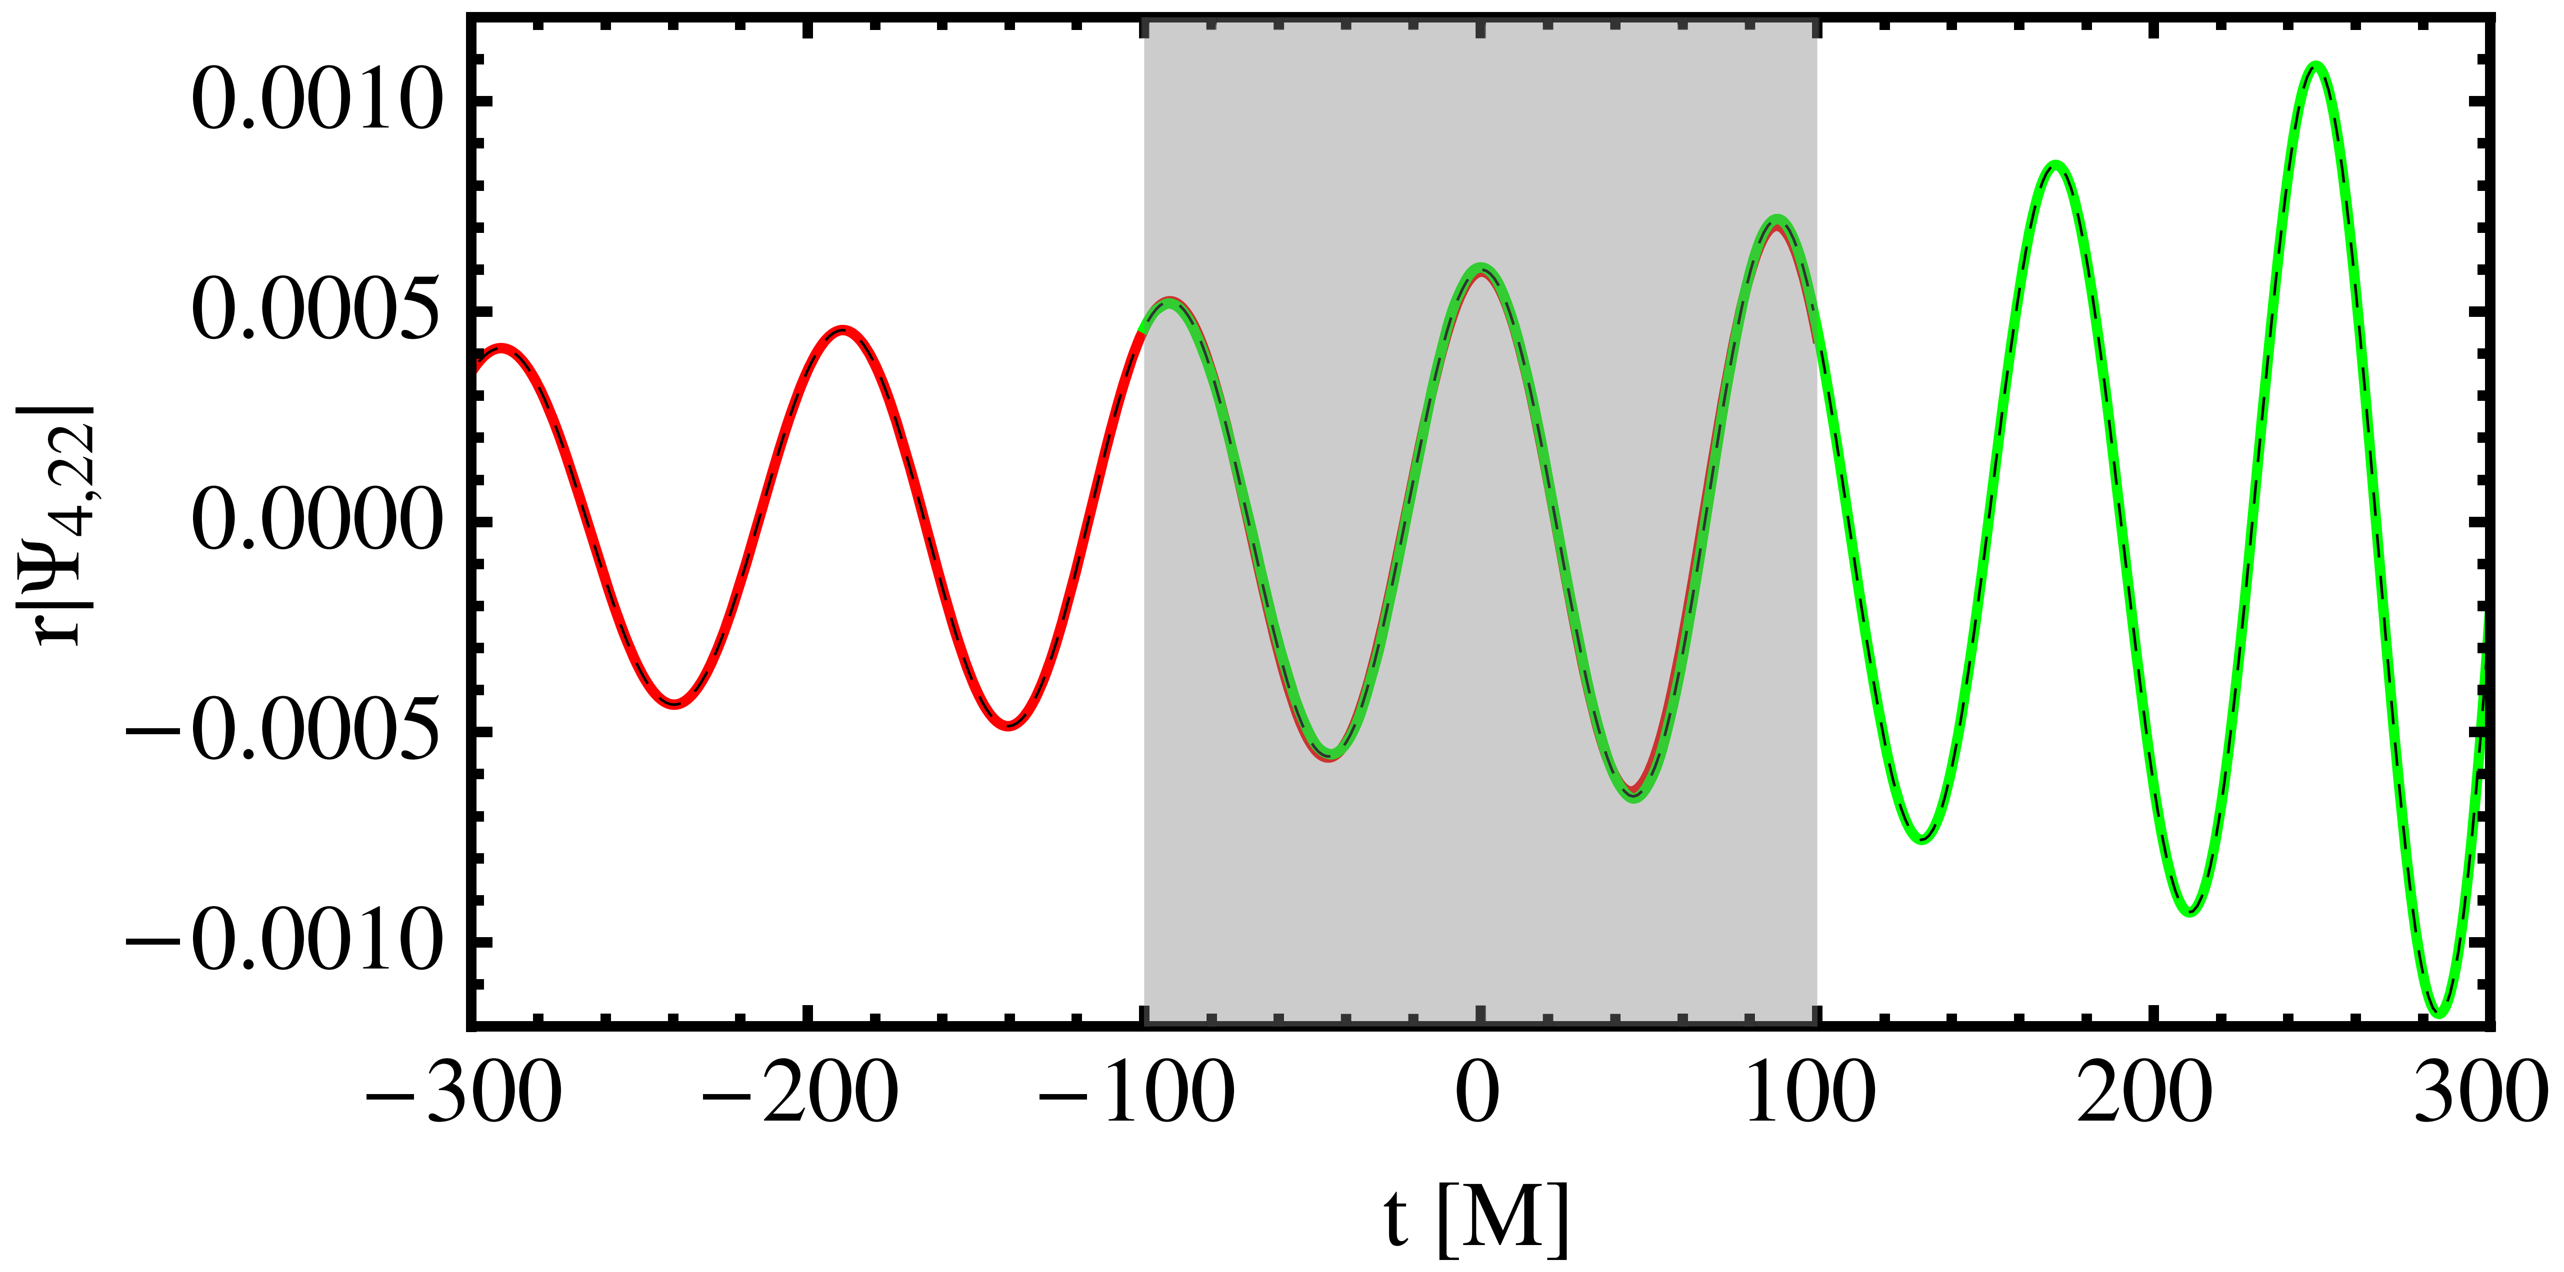

PHENOMENOLOGICAL WAVEFORM MODELS

- Need to fit the phenomenological coefficients:
  - Direct fits often problematic
  - Recent years: Hierarchical fitting and collocation point method

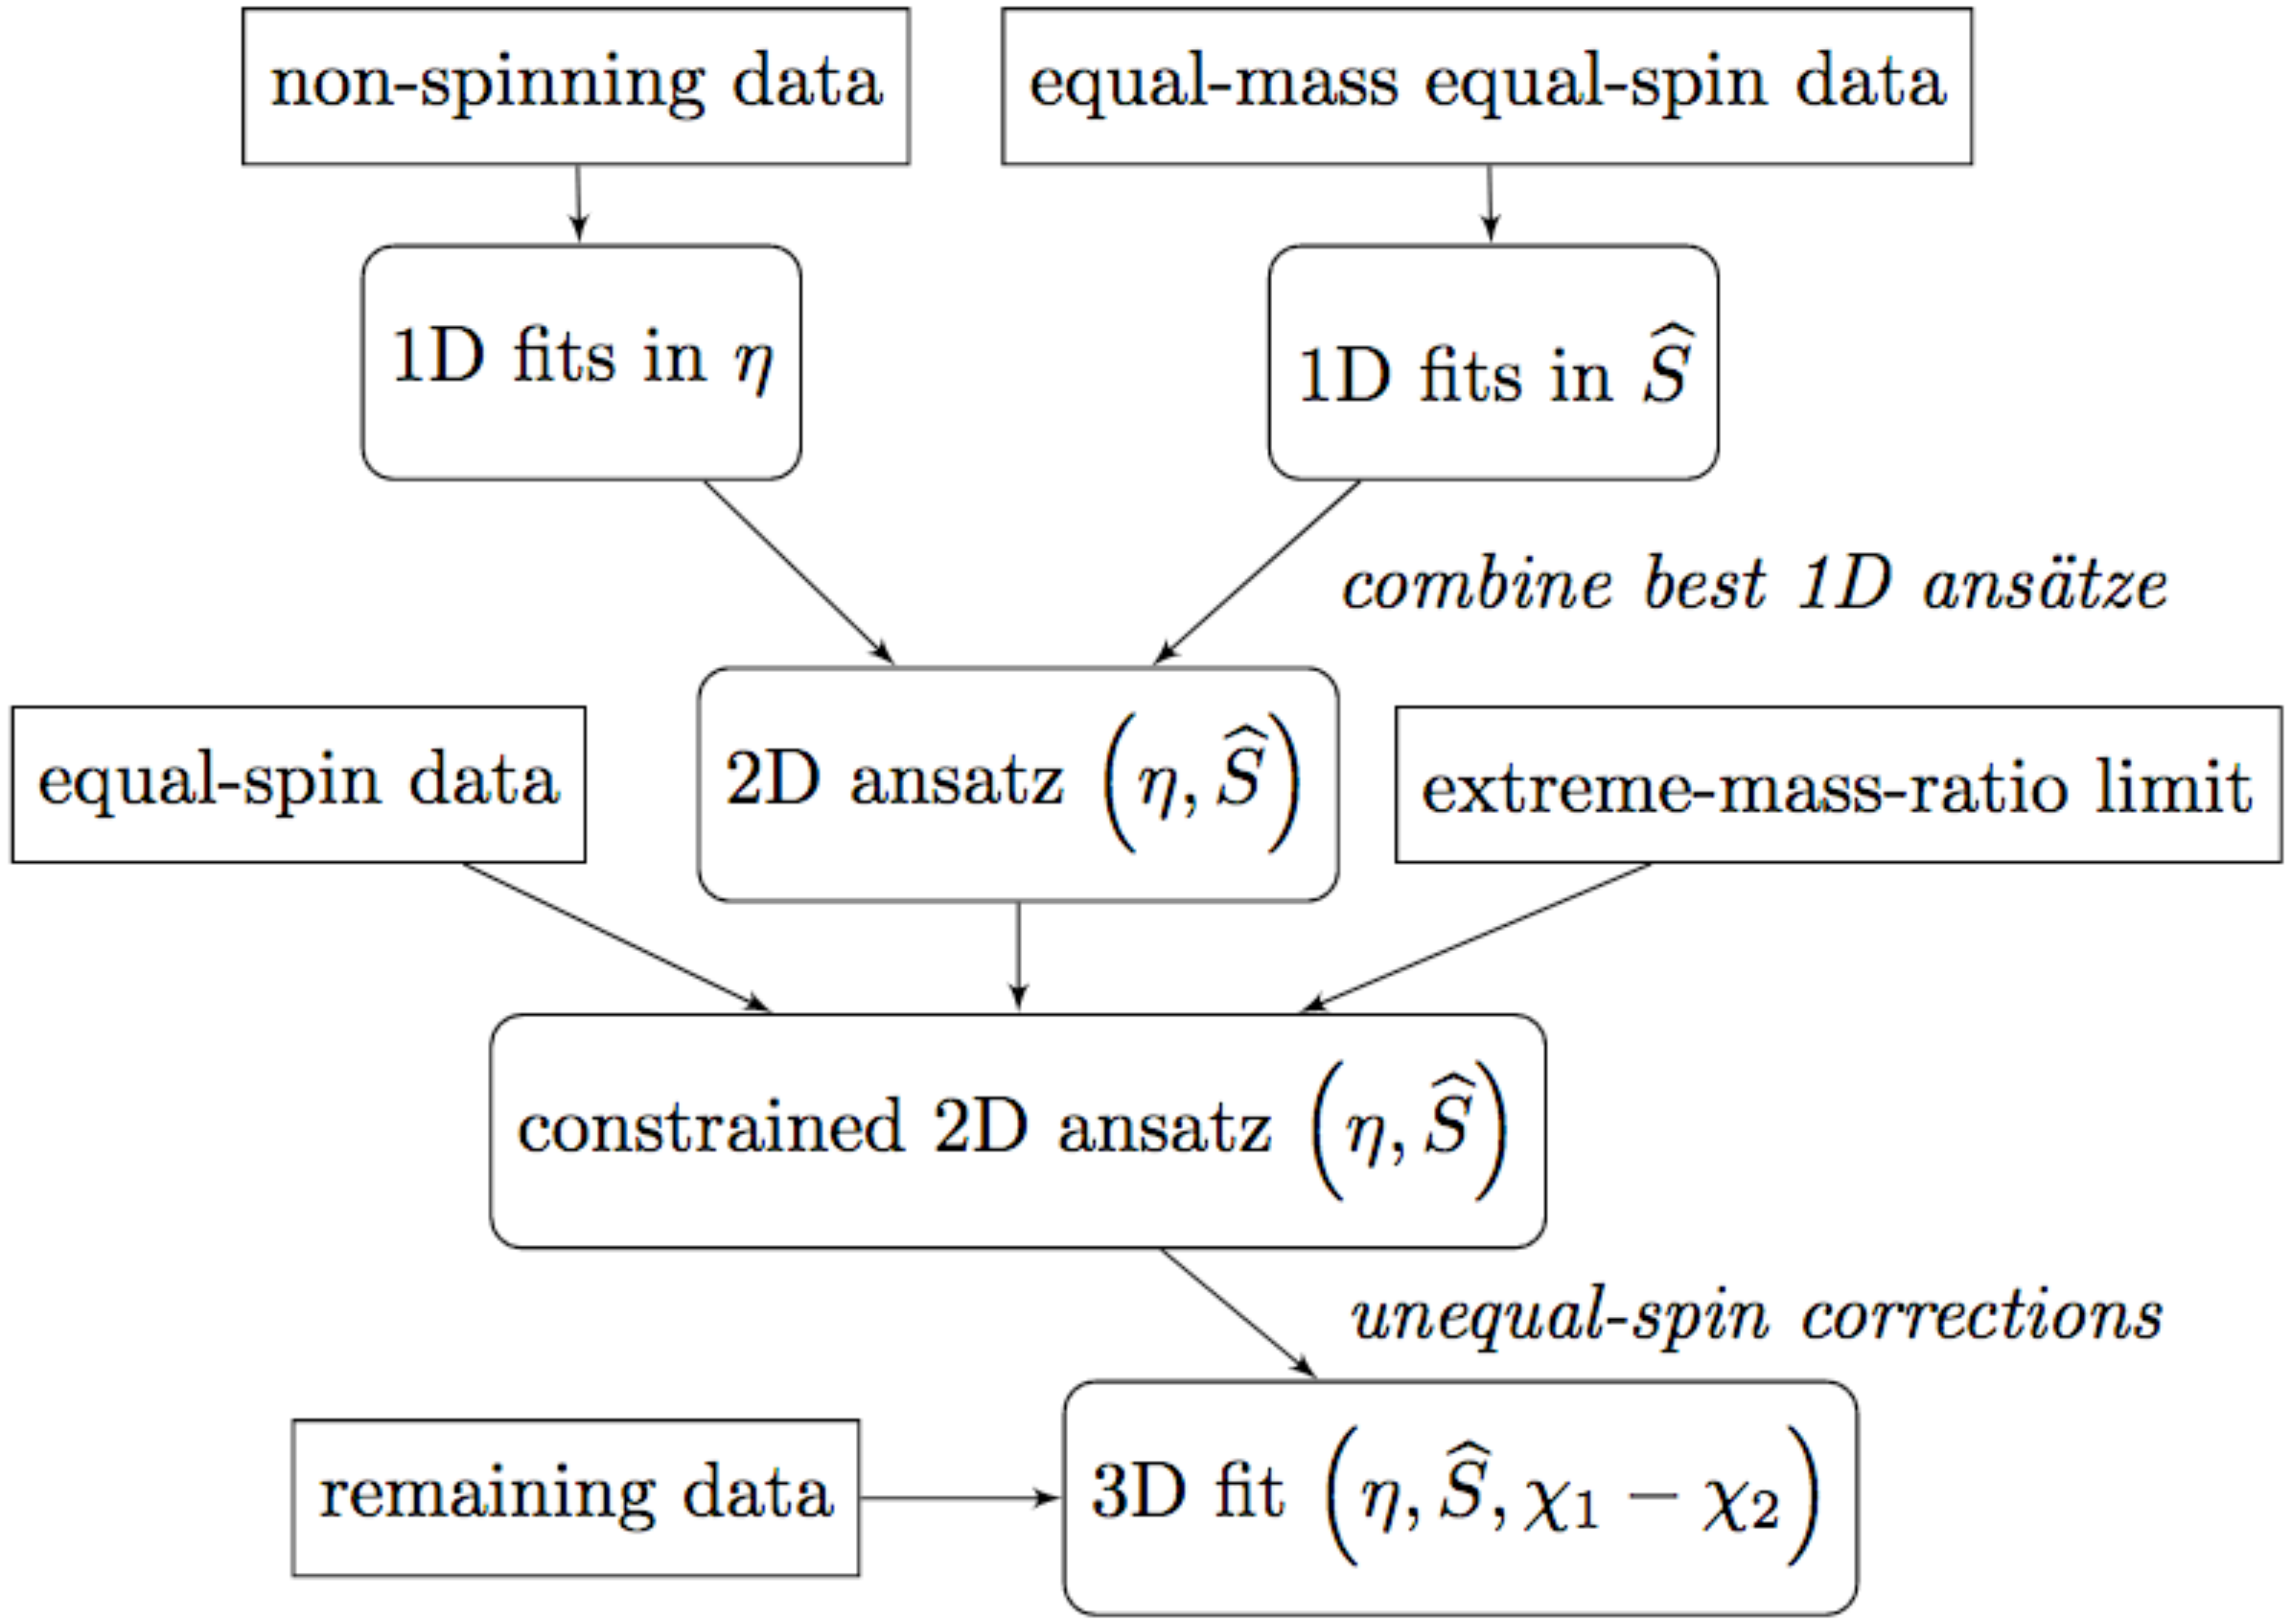

Example: IMRPhenomXAS

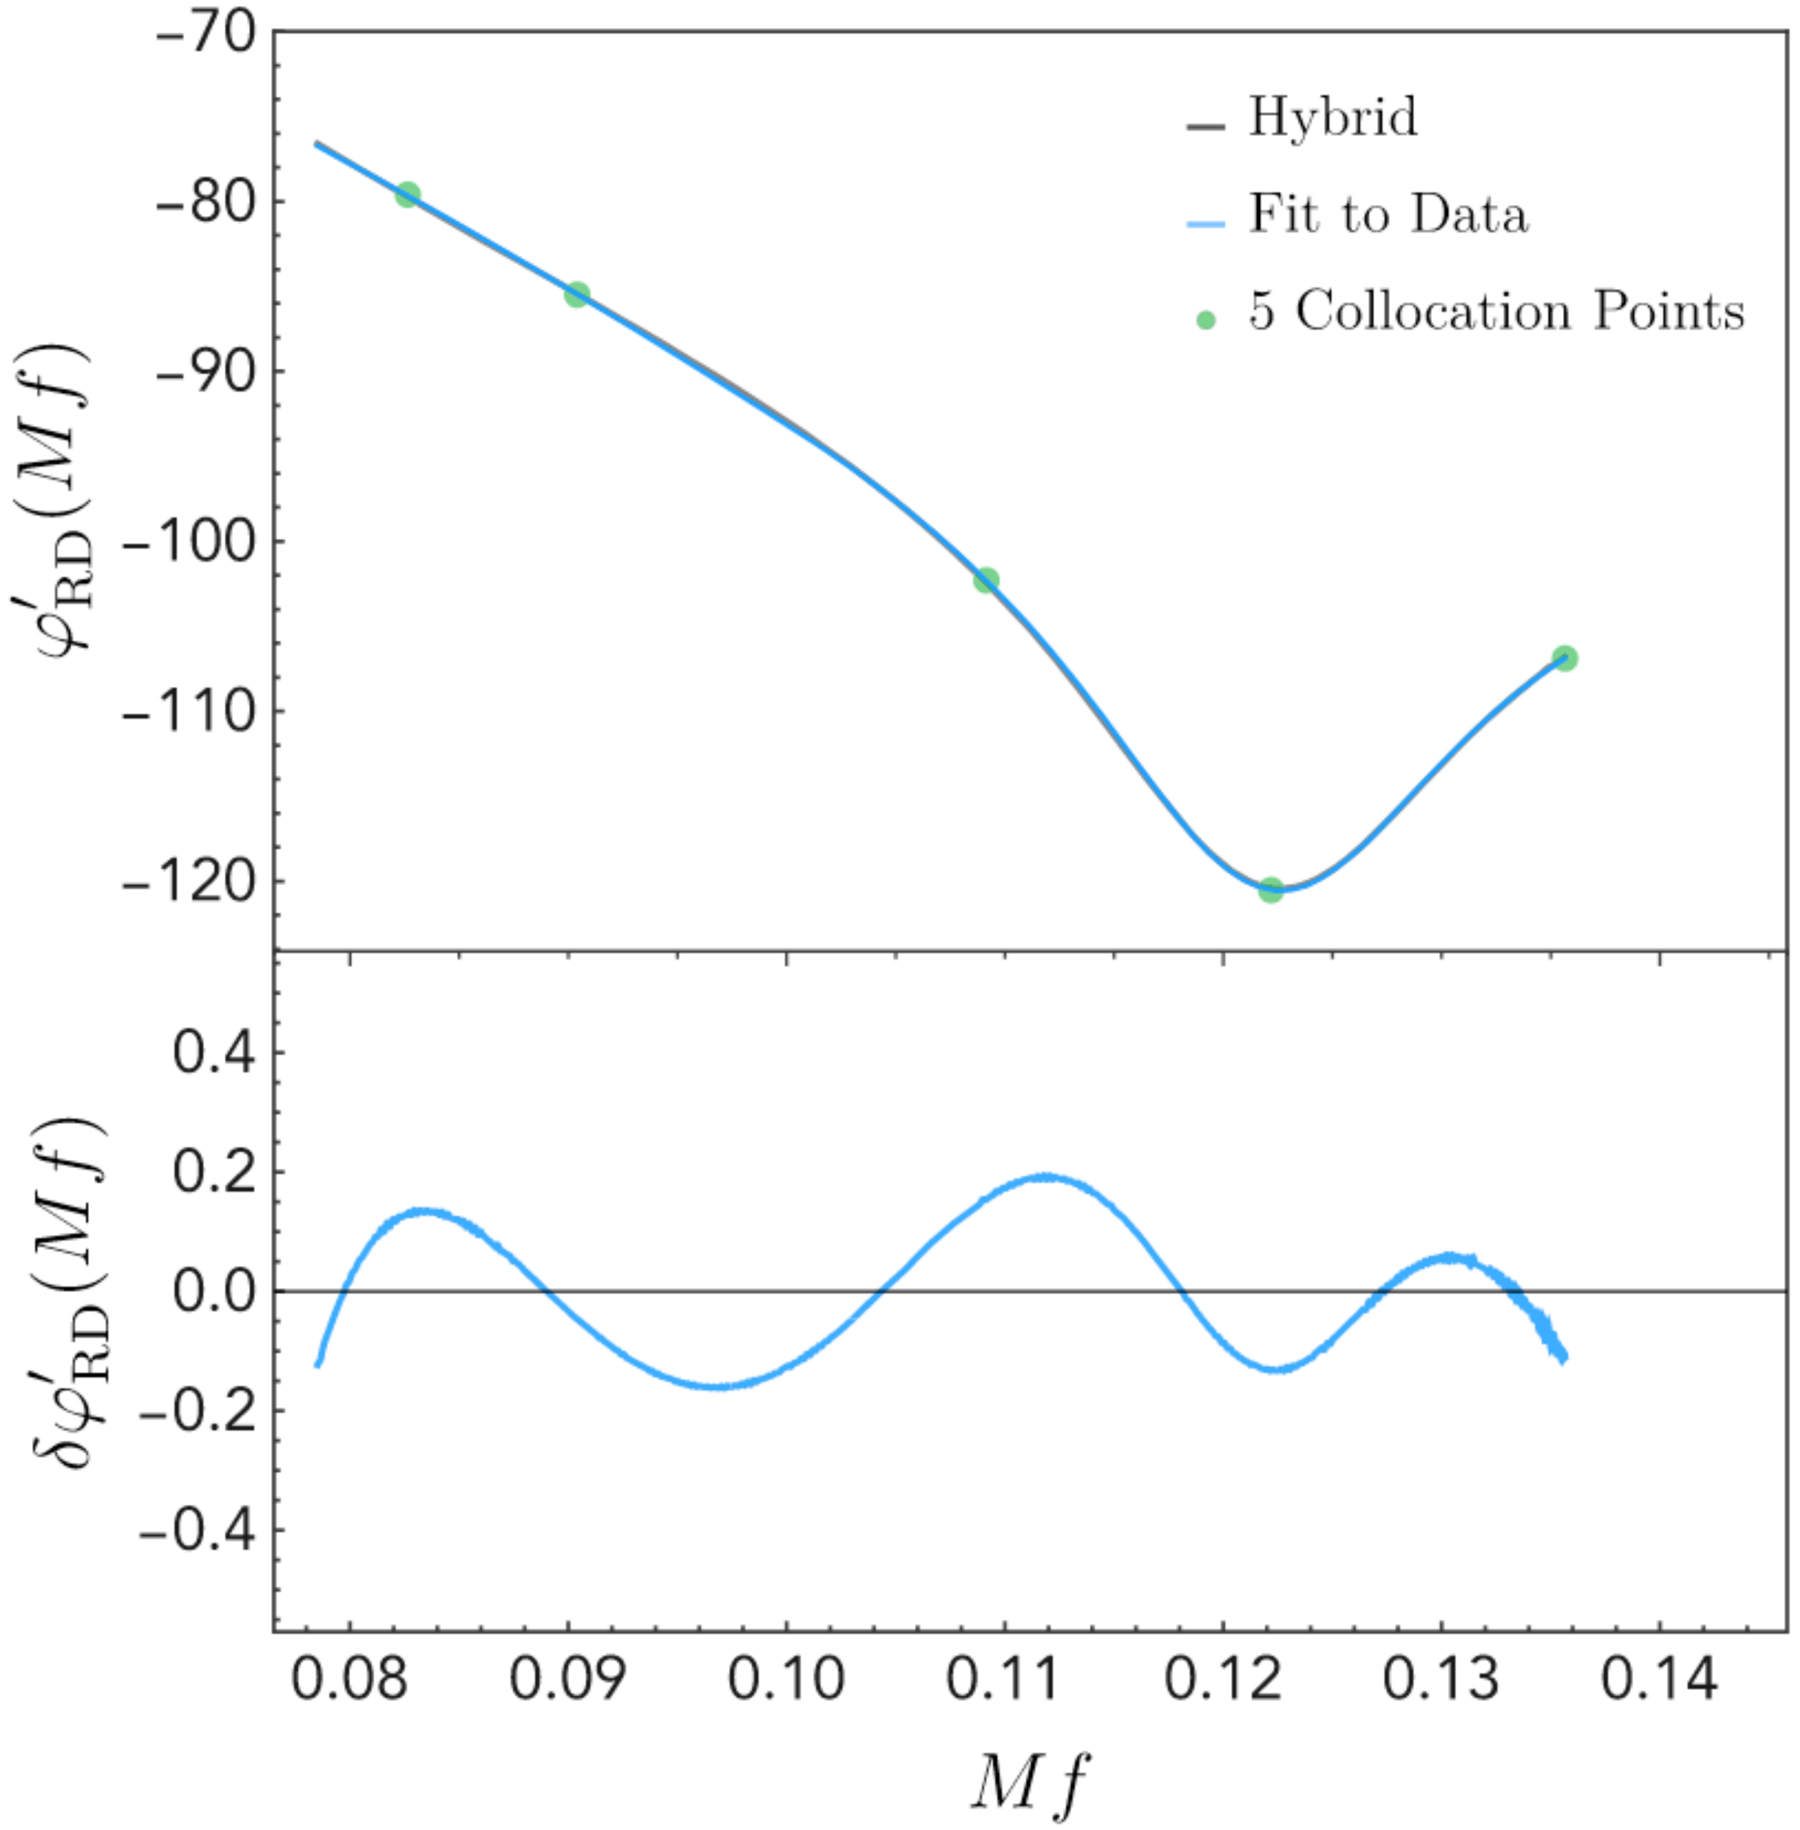

EFFECTIVE-ONE-BODY NR

- Calibration to NR in the strong-field is absolutely crucial!

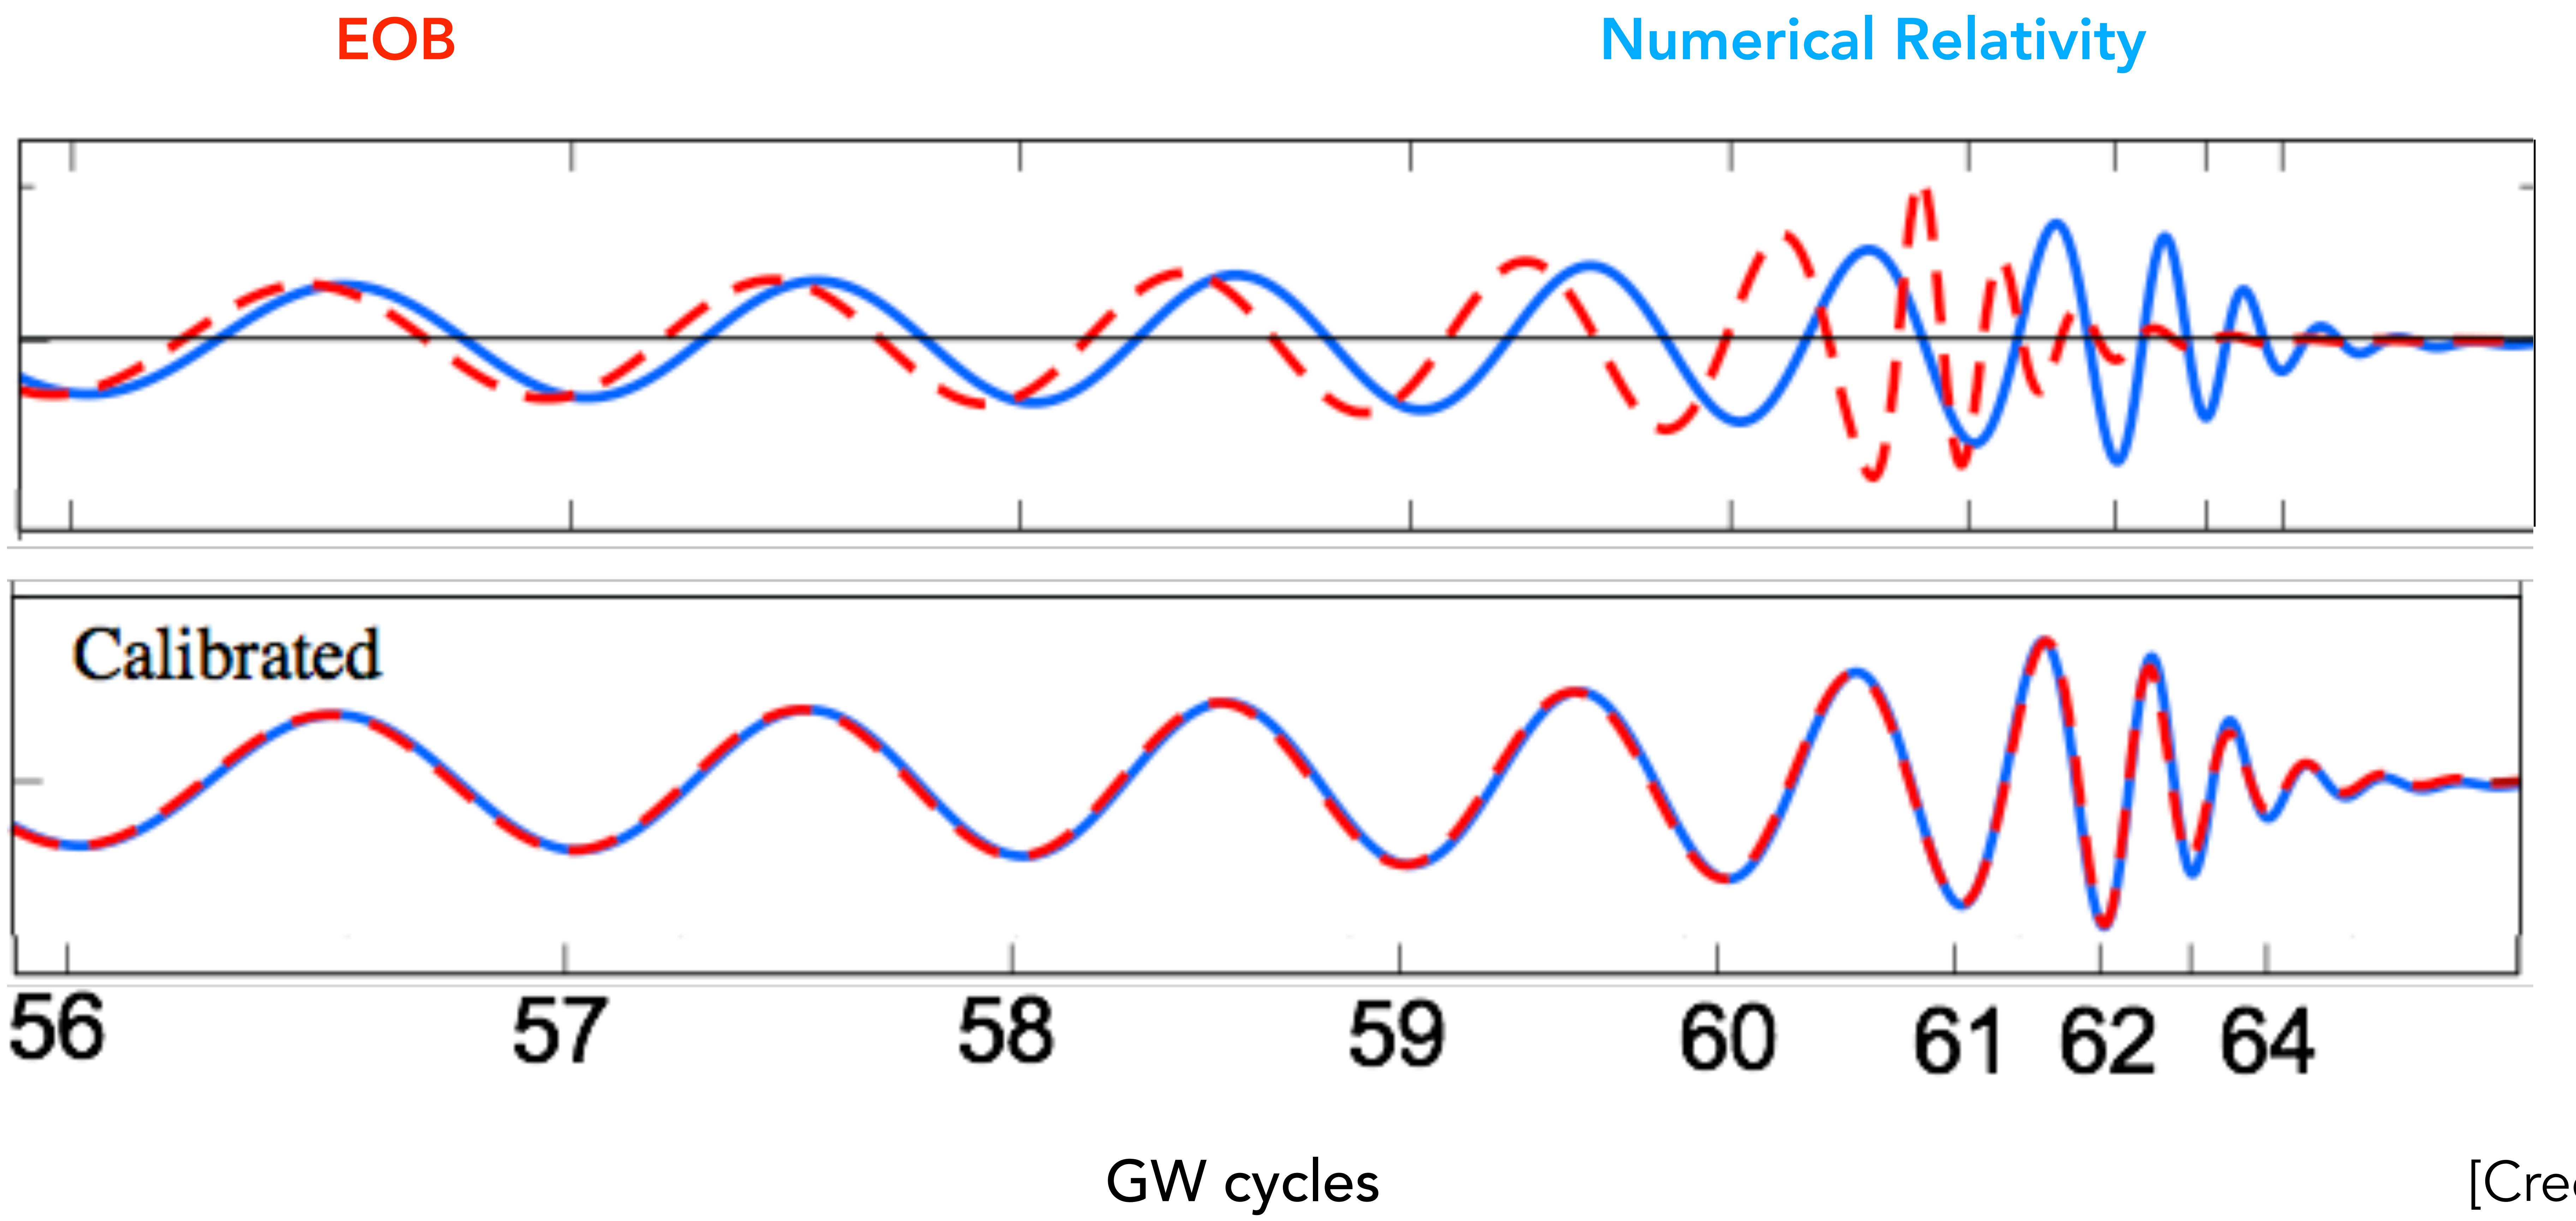

[Credit: Taracchini]

+ information from GSF!

NR SURROGATES

- NR simulations are computationally expensive
  - Limited, discrete coverage of the full compact binary parameter space
- Are pure NR waveform models achievable?
  - Reduced order modelling to compress information (= dimensional reduction)
  - **Interpolation** between discrete waveforms to build a continuous surrogate model

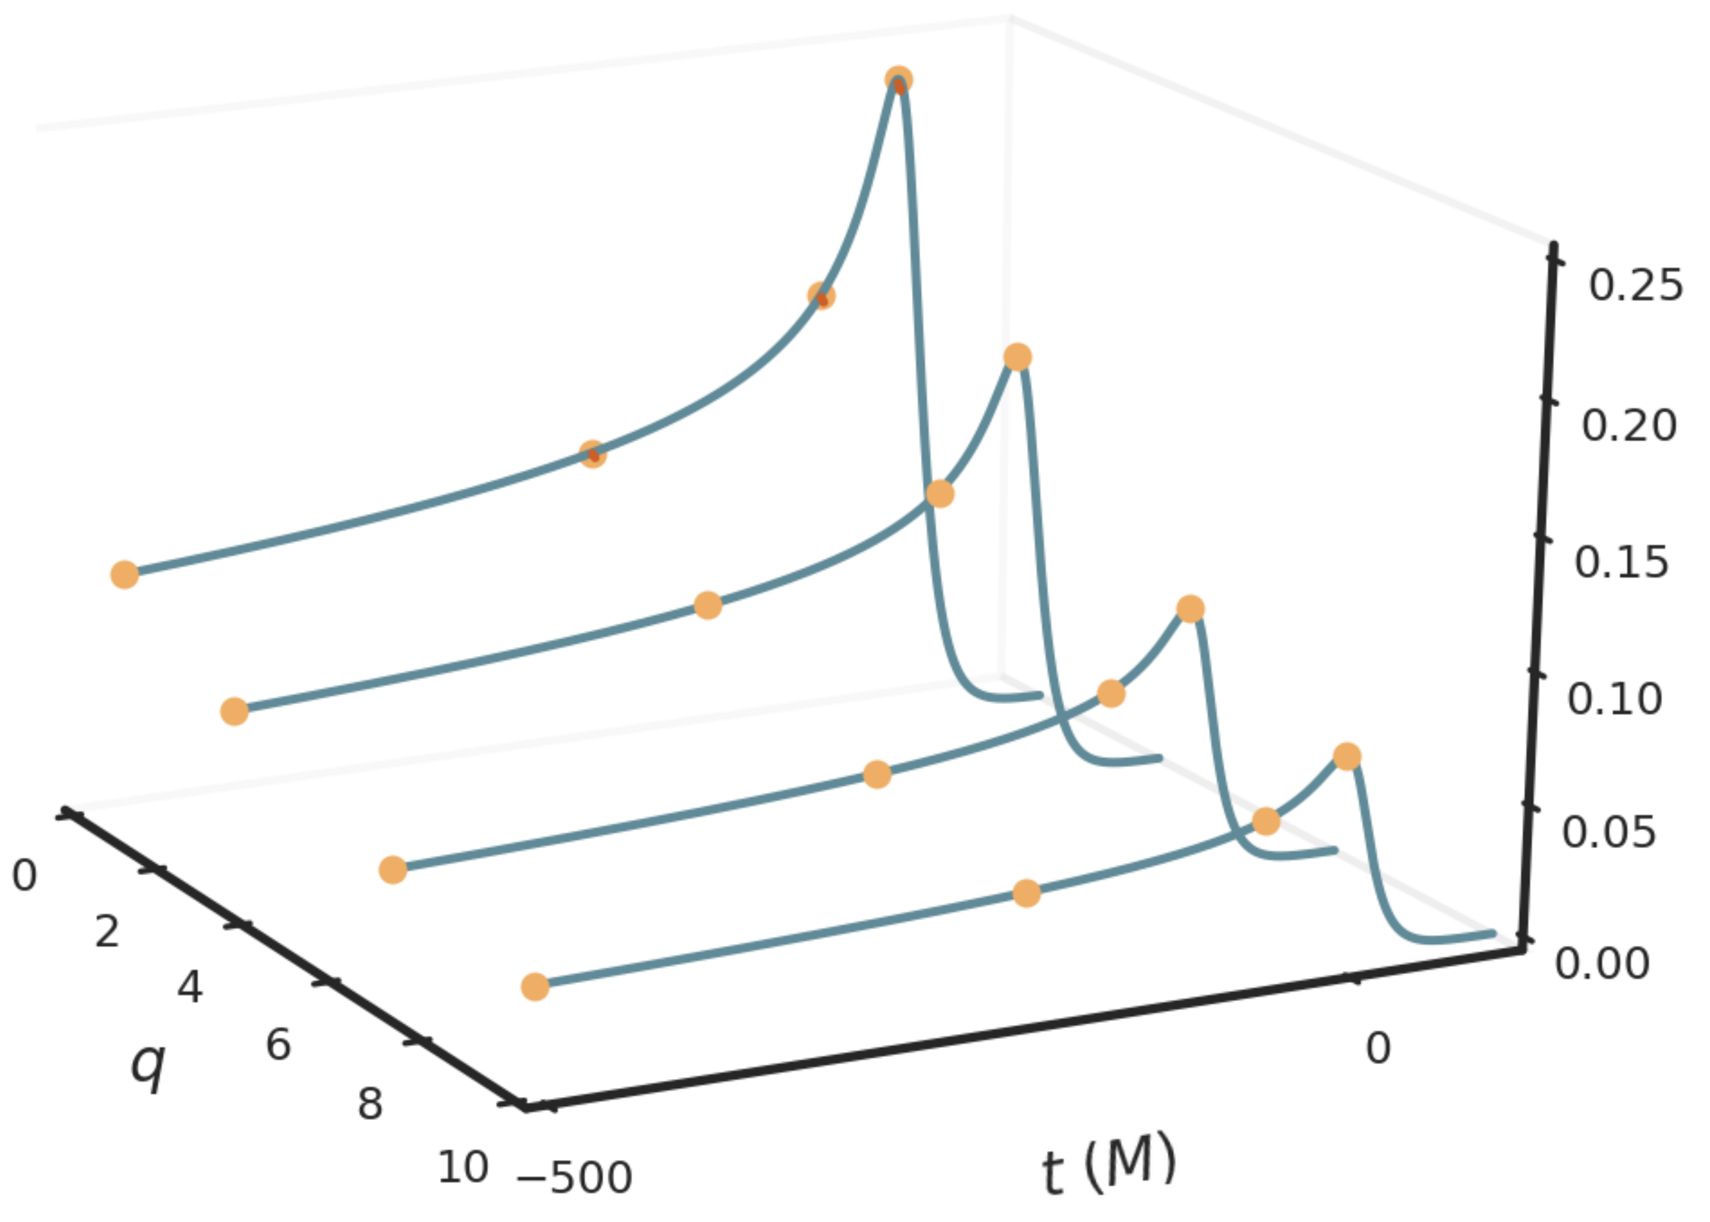

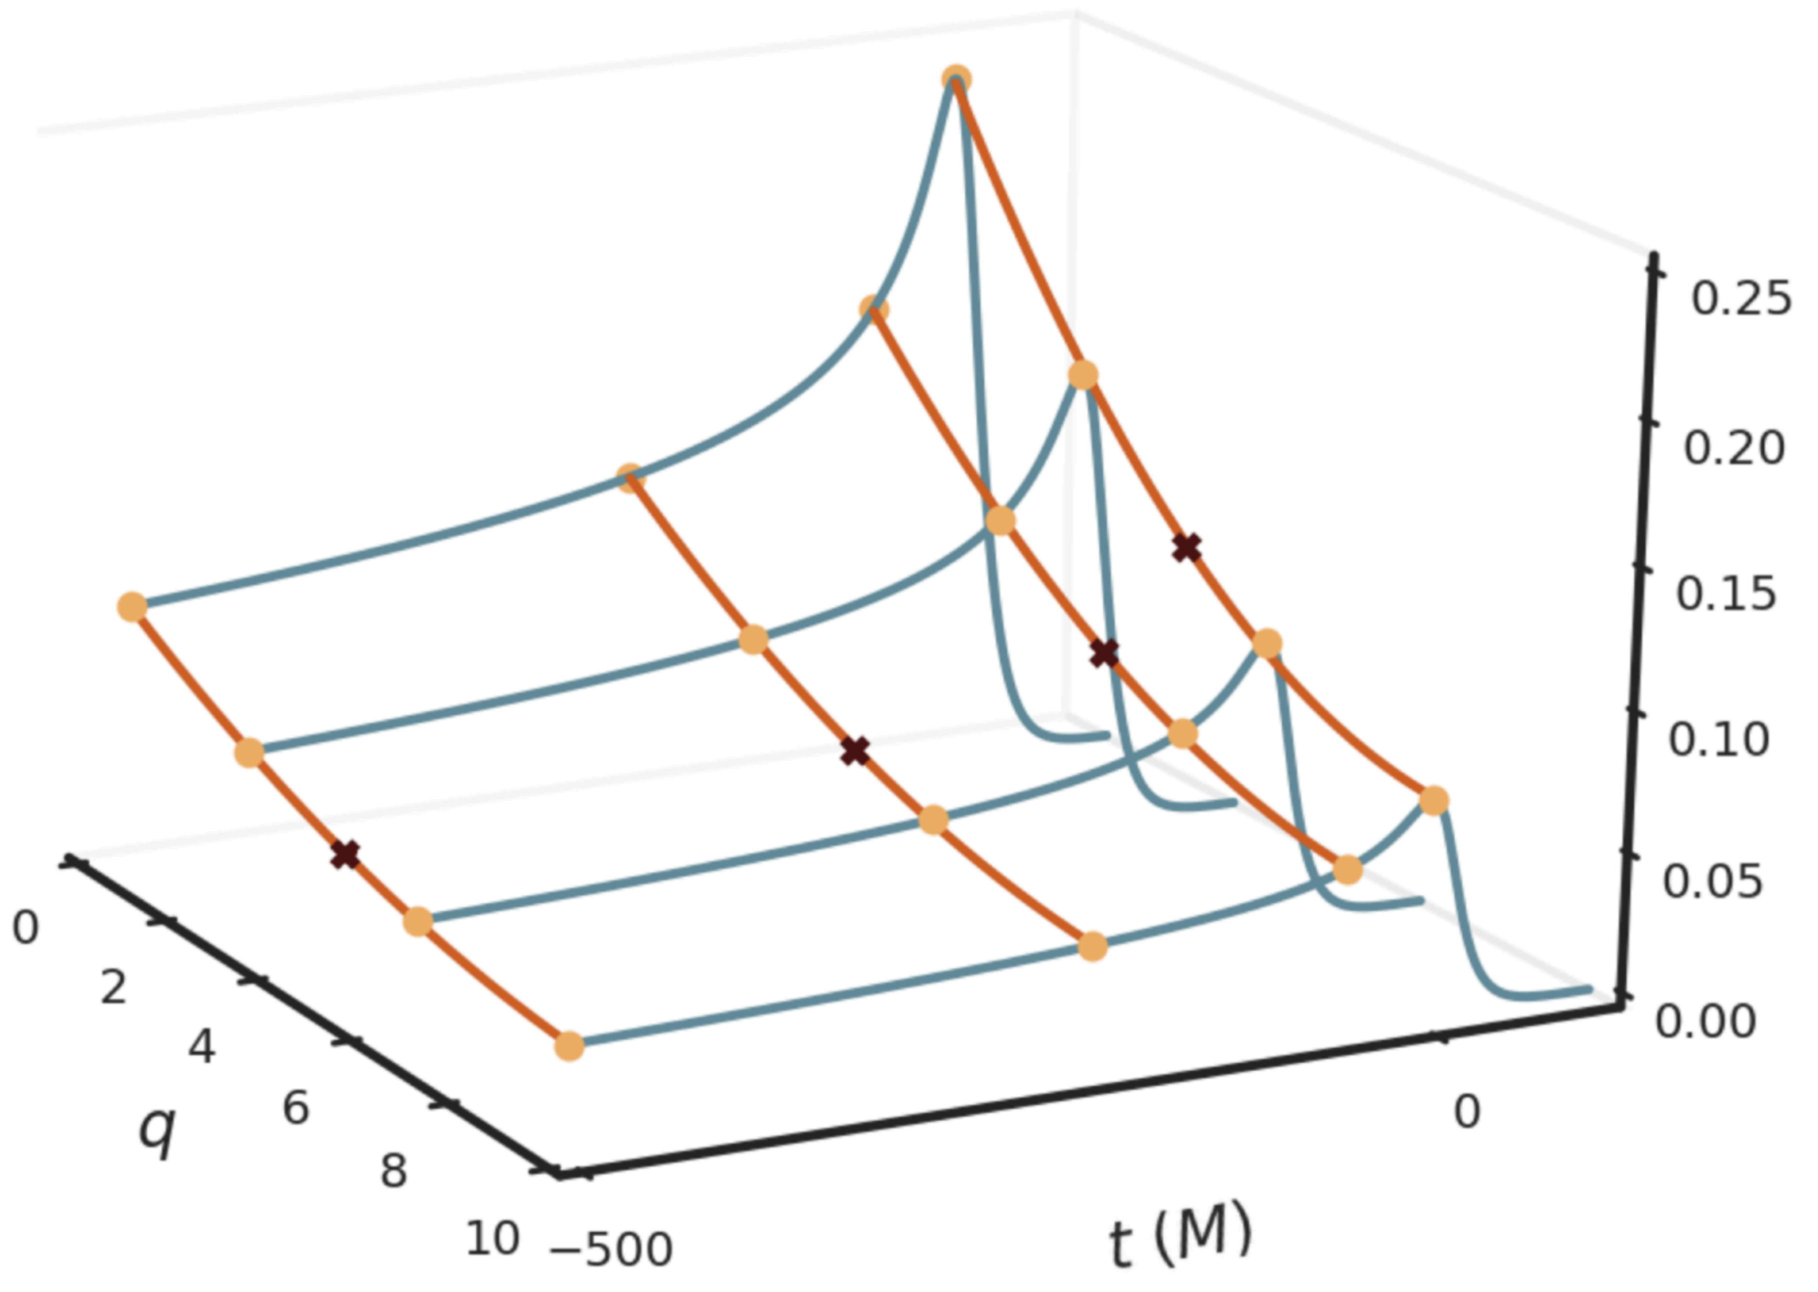

Image credit: Vijay Varma
https://github.com/vijayvarma392/SurrogateMovie

WAVEFORM ACCURACY & REQUIREMENTS

ACCURACY METRICS FOR WAVEFORM MODELS

- Noise-weighted inner product between two waveforms:

$$\langle h_1, h_2 \rangle = 4\,\mathfrak{Re} \int_{f_1}^{f_2} \frac{\tilde{h}_1(f)\,\tilde{h}_2^*(f)}{S_n(f)} df$$

- **Match/faithfulness** optimised over time and phase shift:

$$\mathcal{M}(h_1, h_2) = \max_{t_0, \phi_0} \frac{\langle h_1, h_2 \rangle}{\sqrt{\langle h_1, h_1 \rangle}\sqrt{\langle h_2, h_2 \rangle}}$$

Note: Mismatch

$1 - \mathscr{M}$

- For precession, additional (numerical) optimisation over polarisation:

$$\bar{\mathcal{M}}(h_1, h_2) = \max_{t_0, \phi_0, \psi_0} \frac{\langle h_1, h_2 \rangle}{\sqrt{\langle h_1, h_1 \rangle}\sqrt{\langle h_2, h_2 \rangle}}$$

ACCURACY

[Pürrer & Haster, 2020]

Key numbers:

$\mathcal{O}(10^{-3}) \to \rho \sim 100$

$\mathcal{O}(10^{-4}) \to \rho \sim 300$

$\mathcal{O}(10^{-5}) \to \rho \sim 1000$

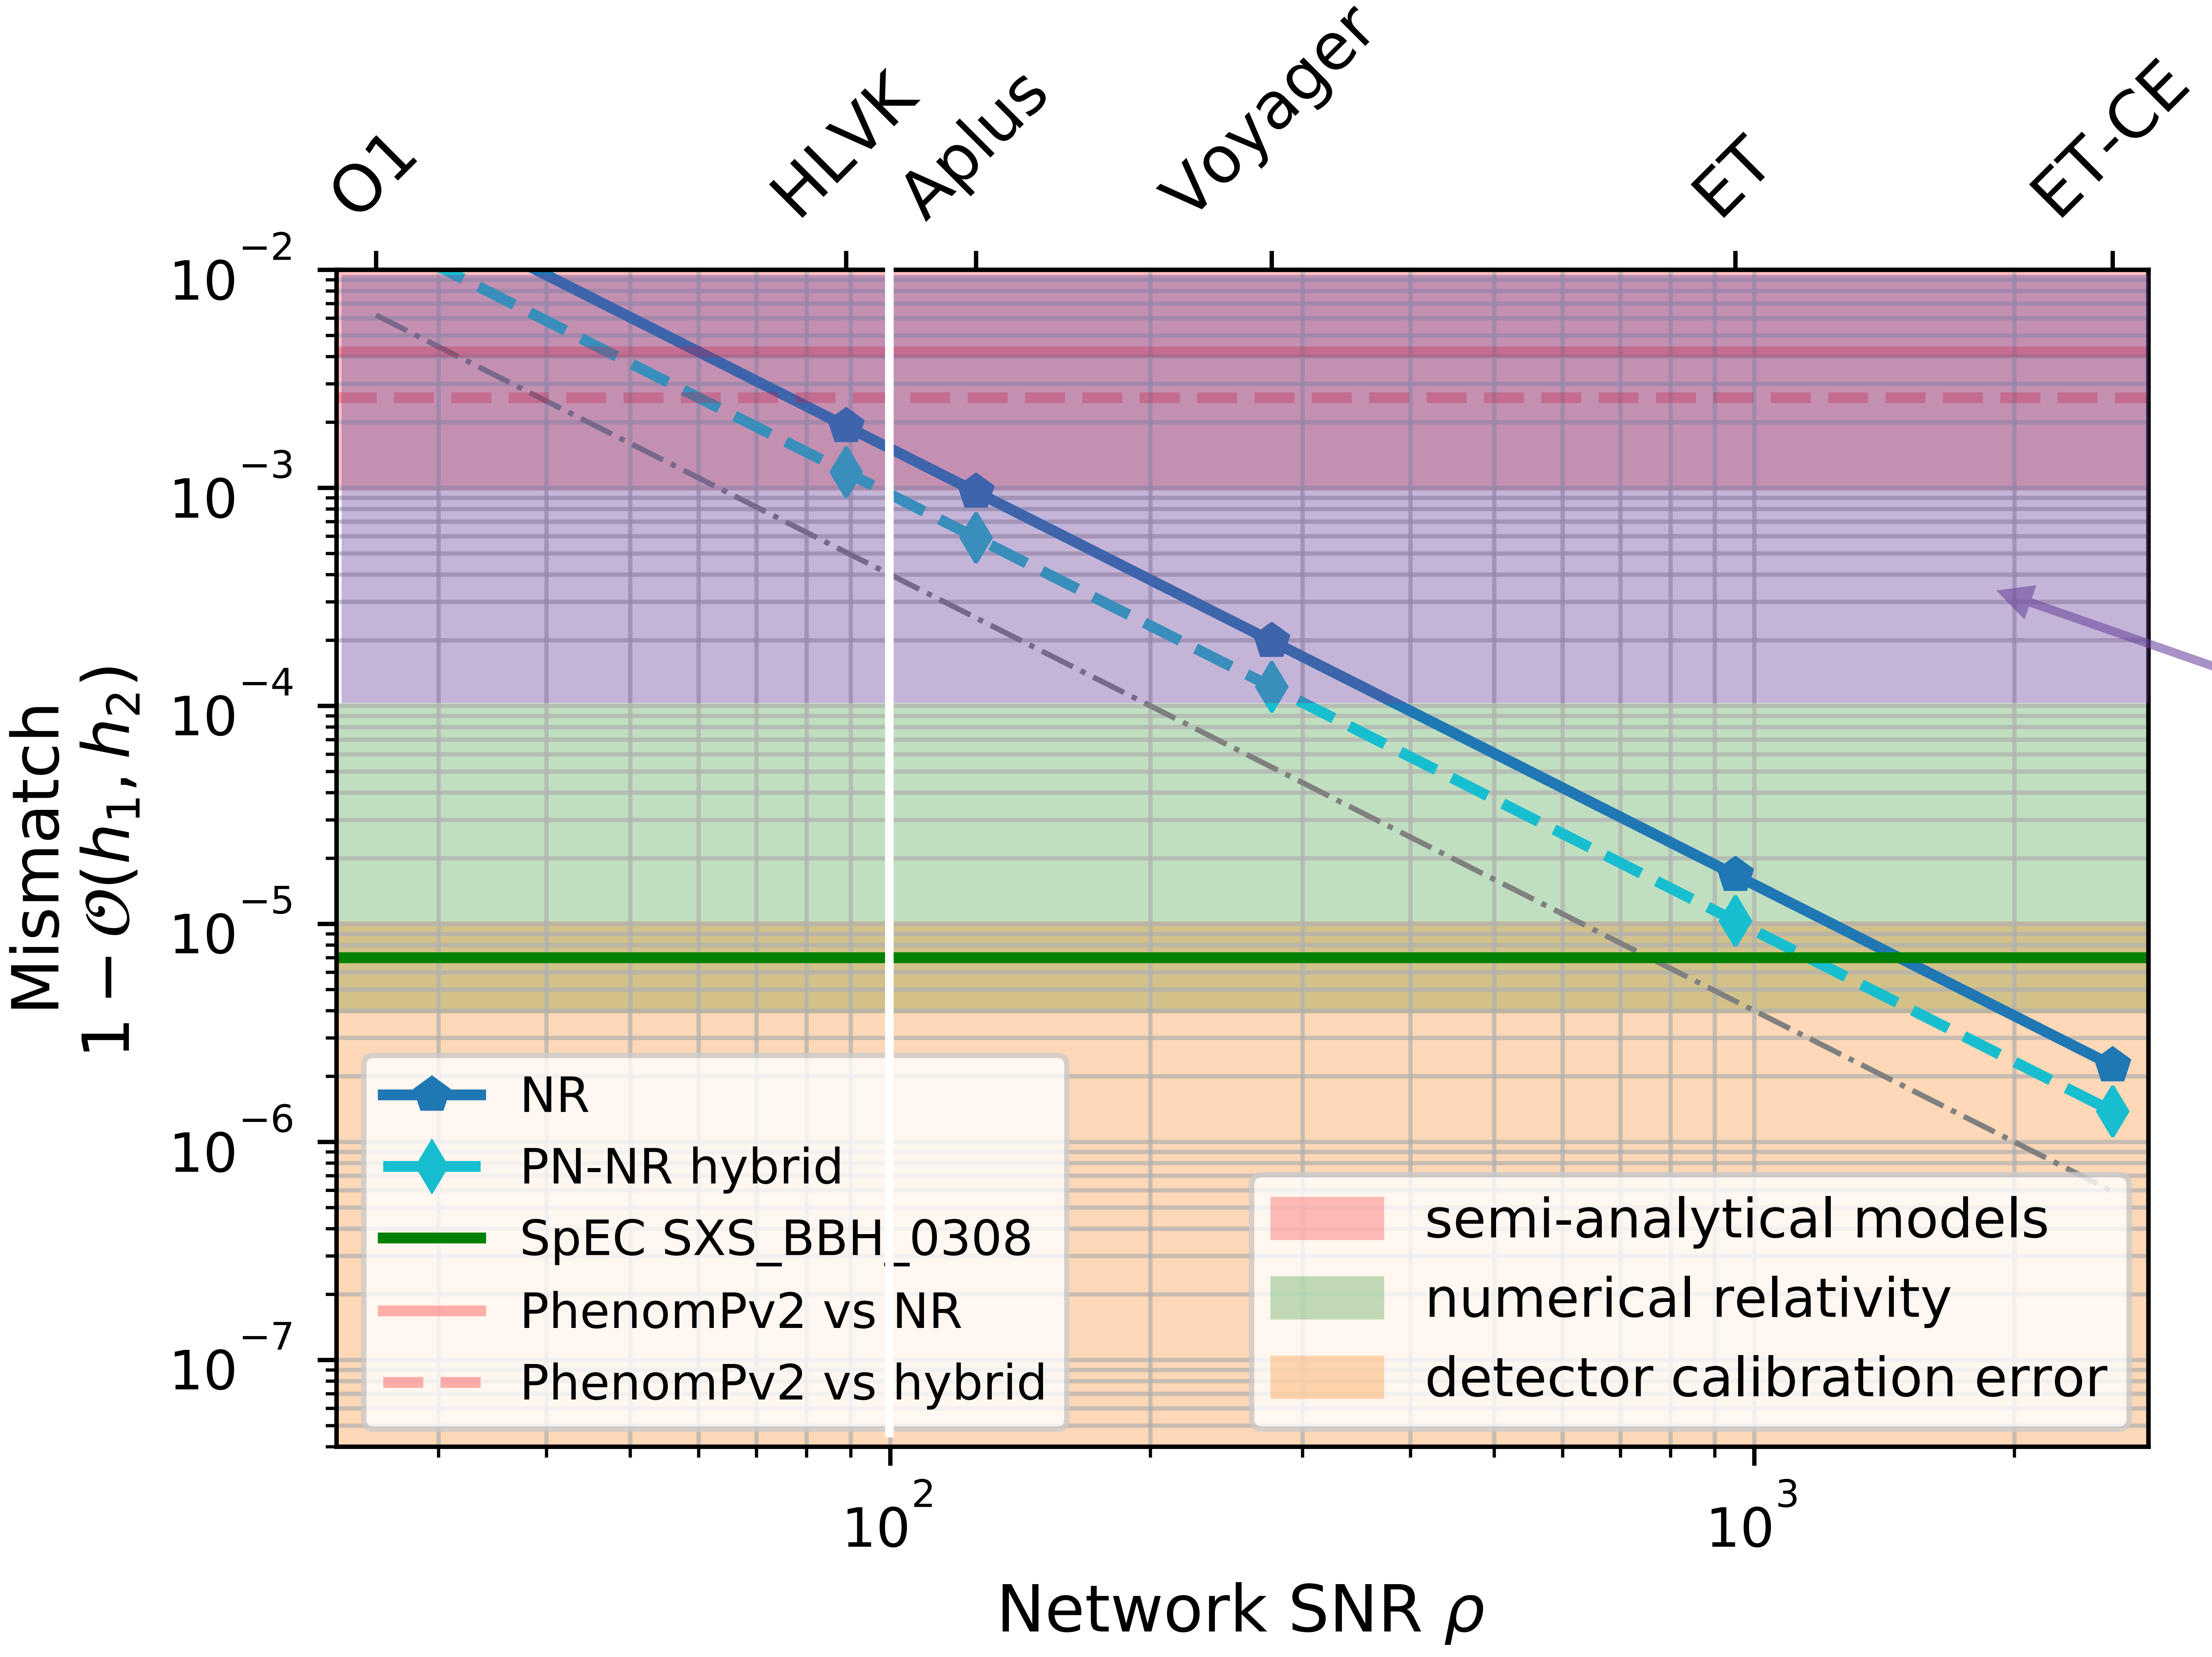

Current generation waveform models with precession and higher modes occupies ~ this band (parameter space dependent)

Requirements become more stringent for more challenging configurations

Configuration	$M_{\rm tot}^{\rm src}/M_\odot$	$\mathcal{M}^{\rm src}/M_\odot$	q	$\vec{\chi}_1$	$\vec{\chi}_2$	$\chi_{\rm eff}$	$\chi_{\rm p}$	$\theta_{\rm JN}$
SXS_BBH_0308	66.4555	28.7443	0.8143	$(-0.1407, 0.0225, 0.3053)$	$(-0.2209, 0.3075, -0.5580)$	-0.0822	0.2994	2.7454

MODEL ACCURACY: IMR

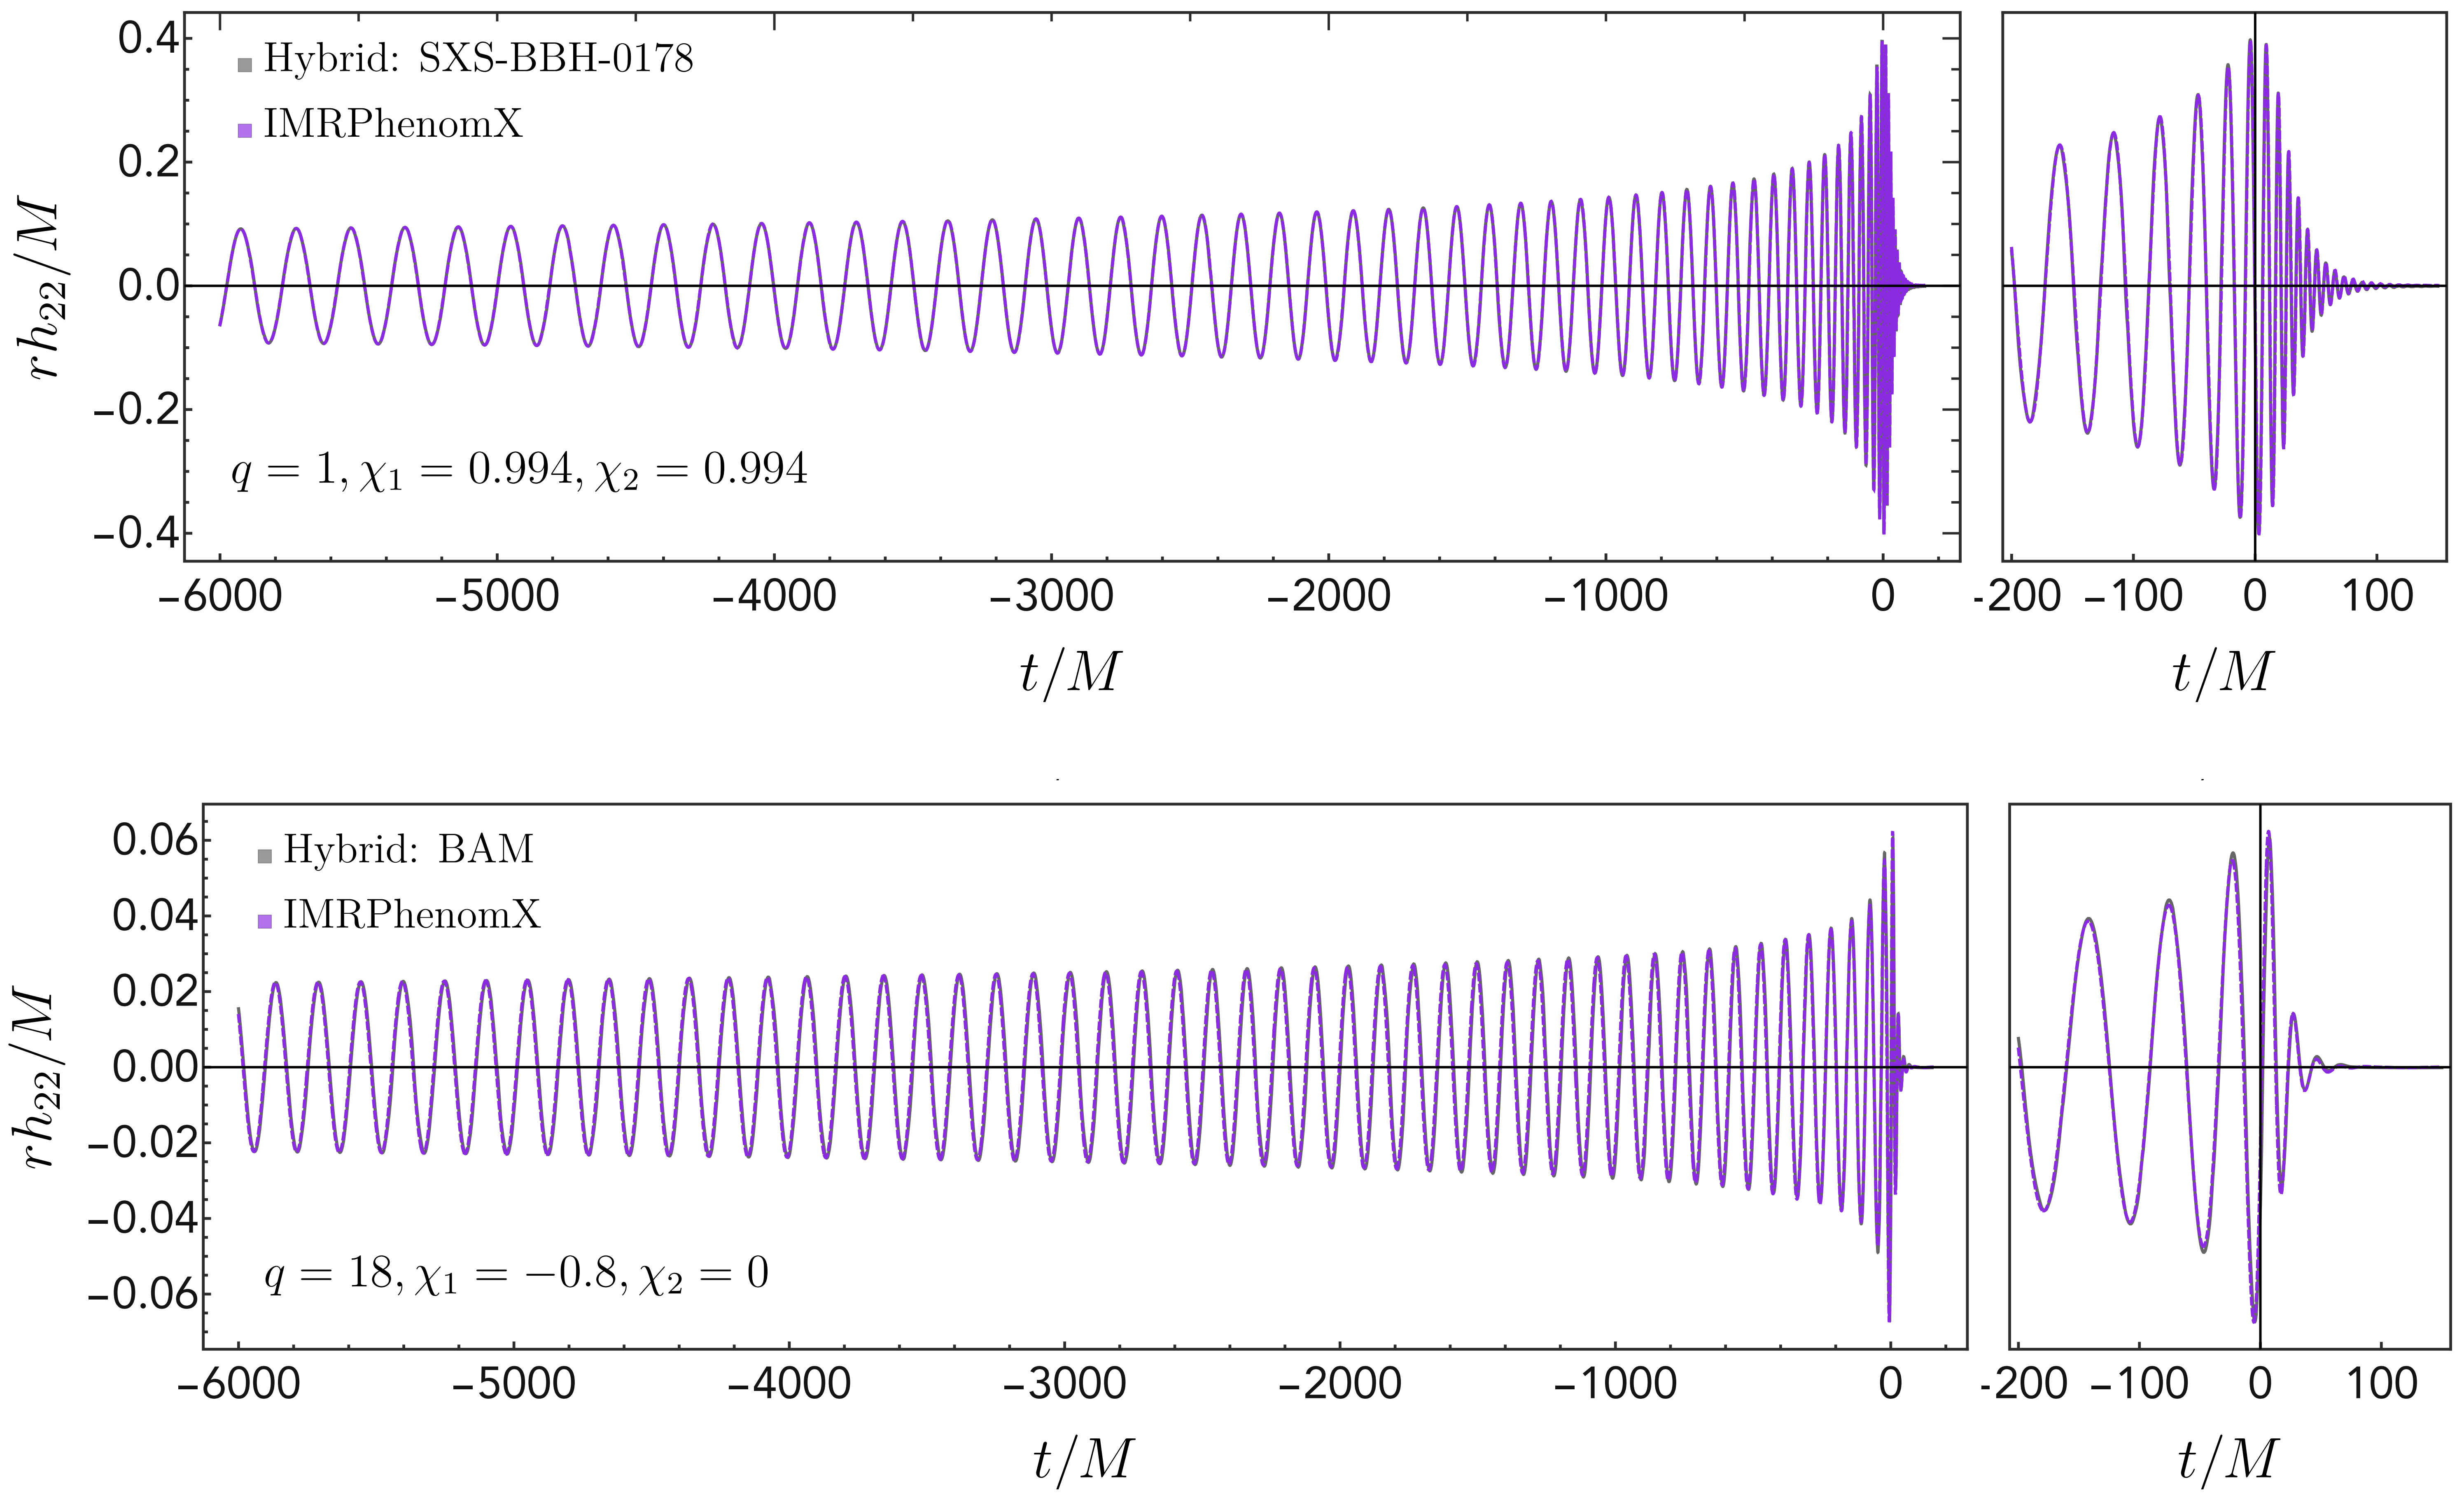

[Pratten+, 2020]

Mismatch against NR simulations, (2,2)-mode only

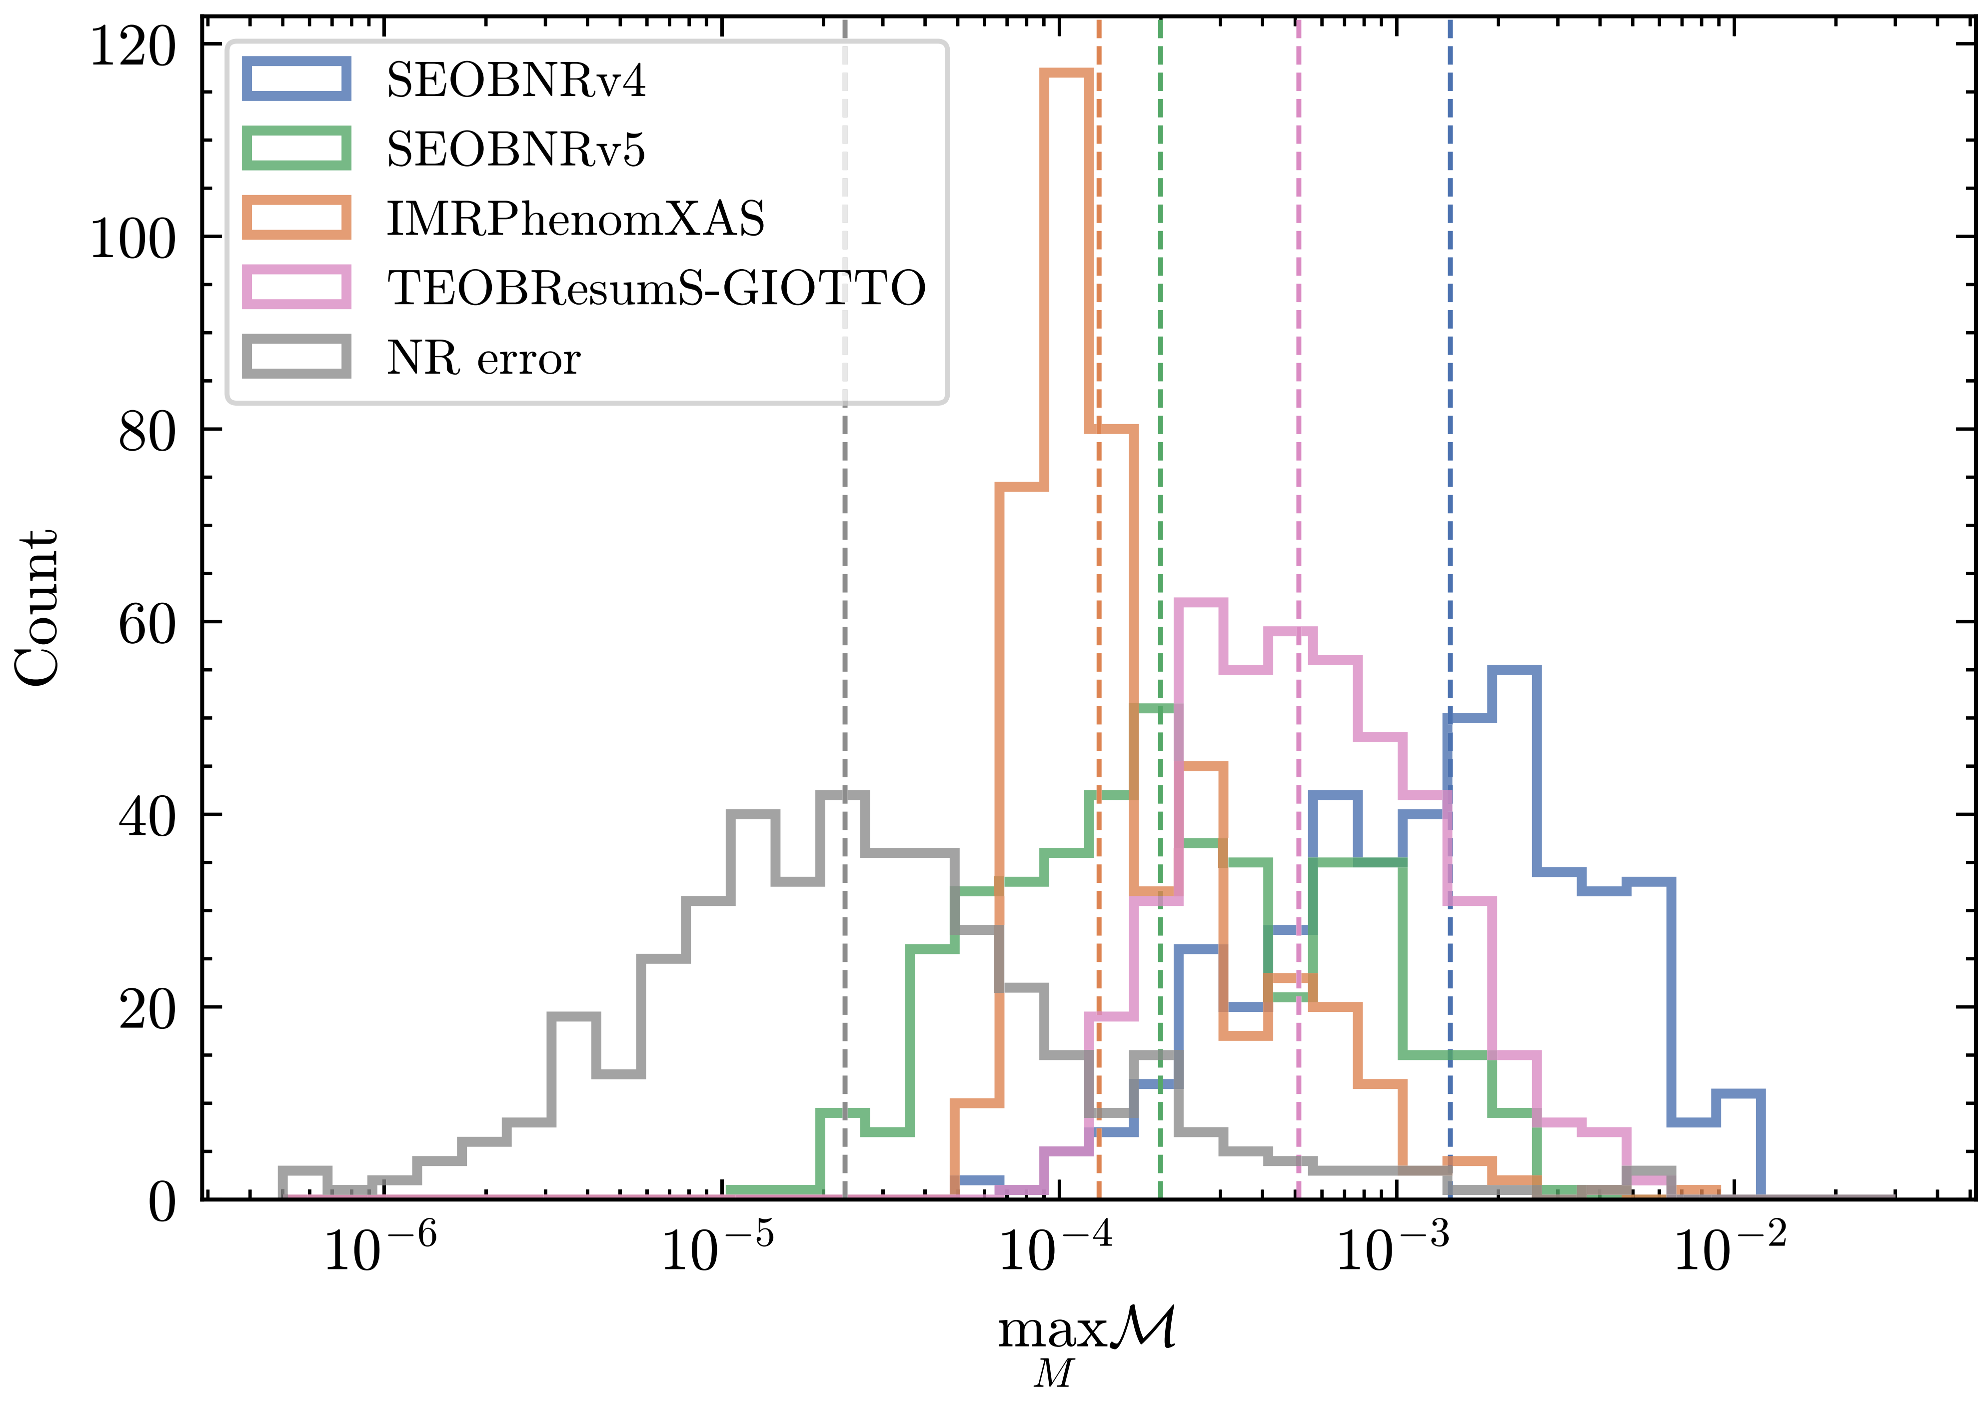

[Pompili+, 2023]

MODEL ACCURACY IMR WITH HIGHER ORDER MODES

- Mode decomposition:

$$h_+ - ih_\times \equiv \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} h_{\ell m}\, {}^{-2}Y_{\ell m}(\theta, \phi)$$

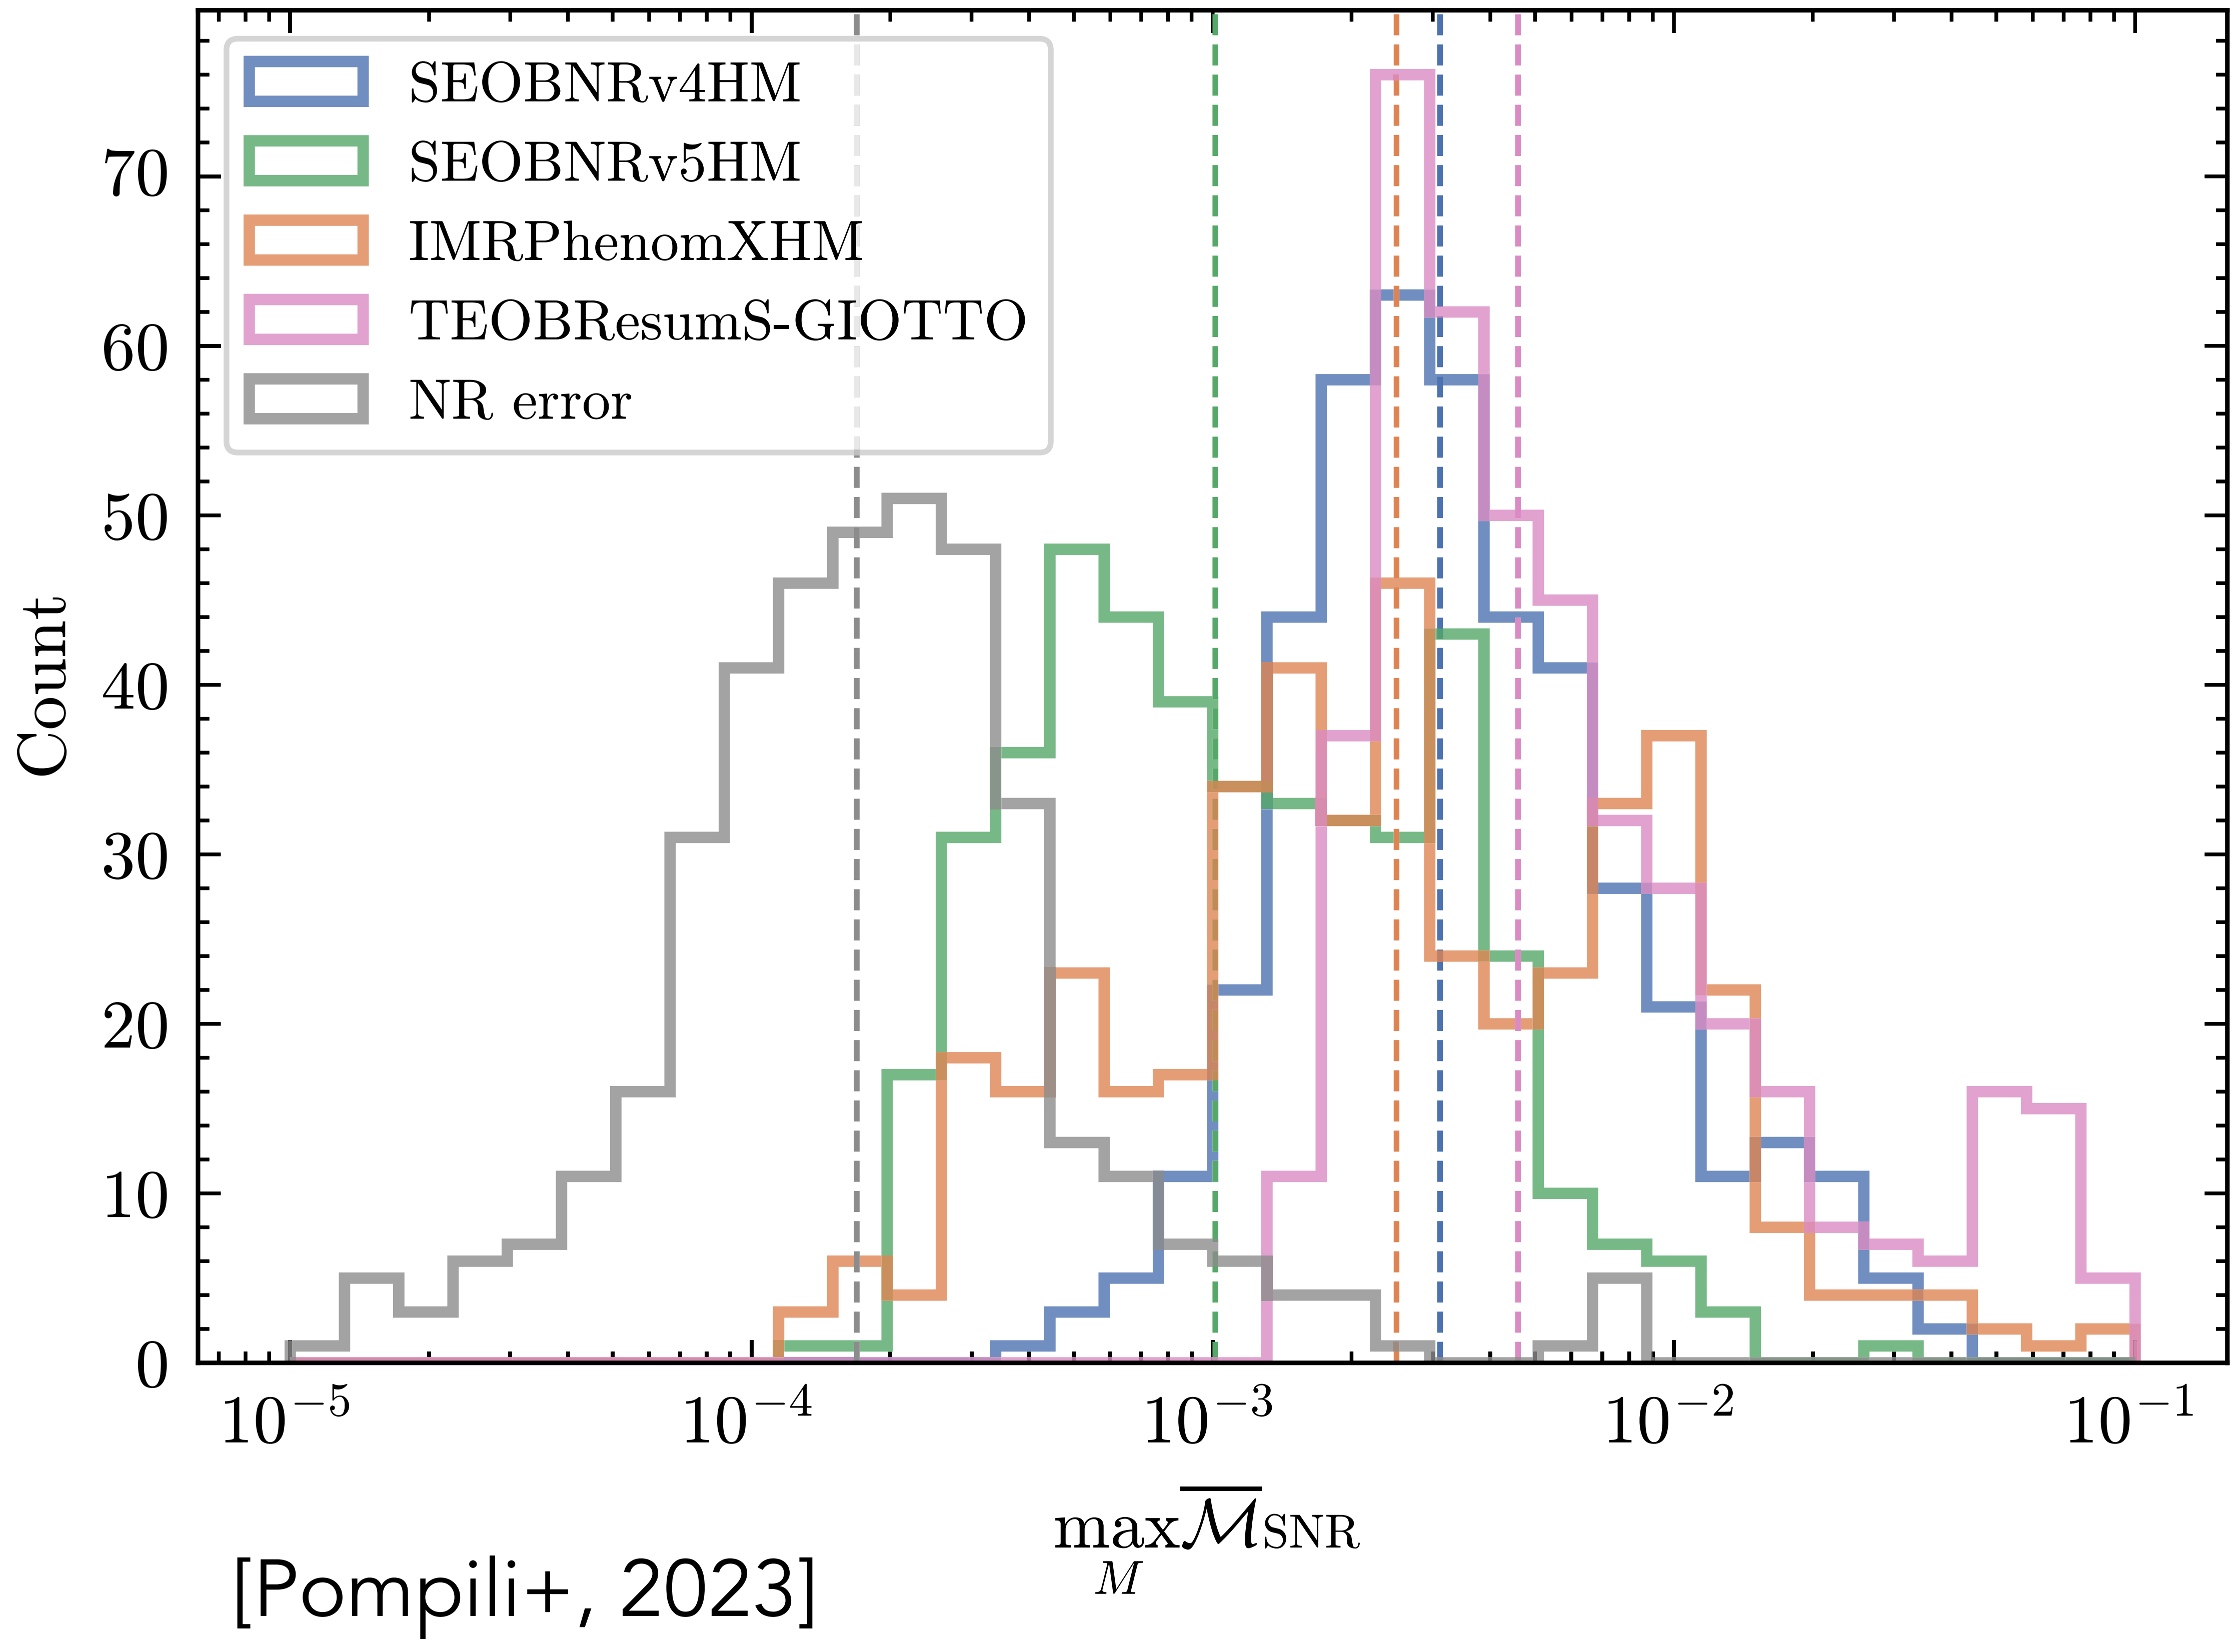

[Pompili+, 2023]

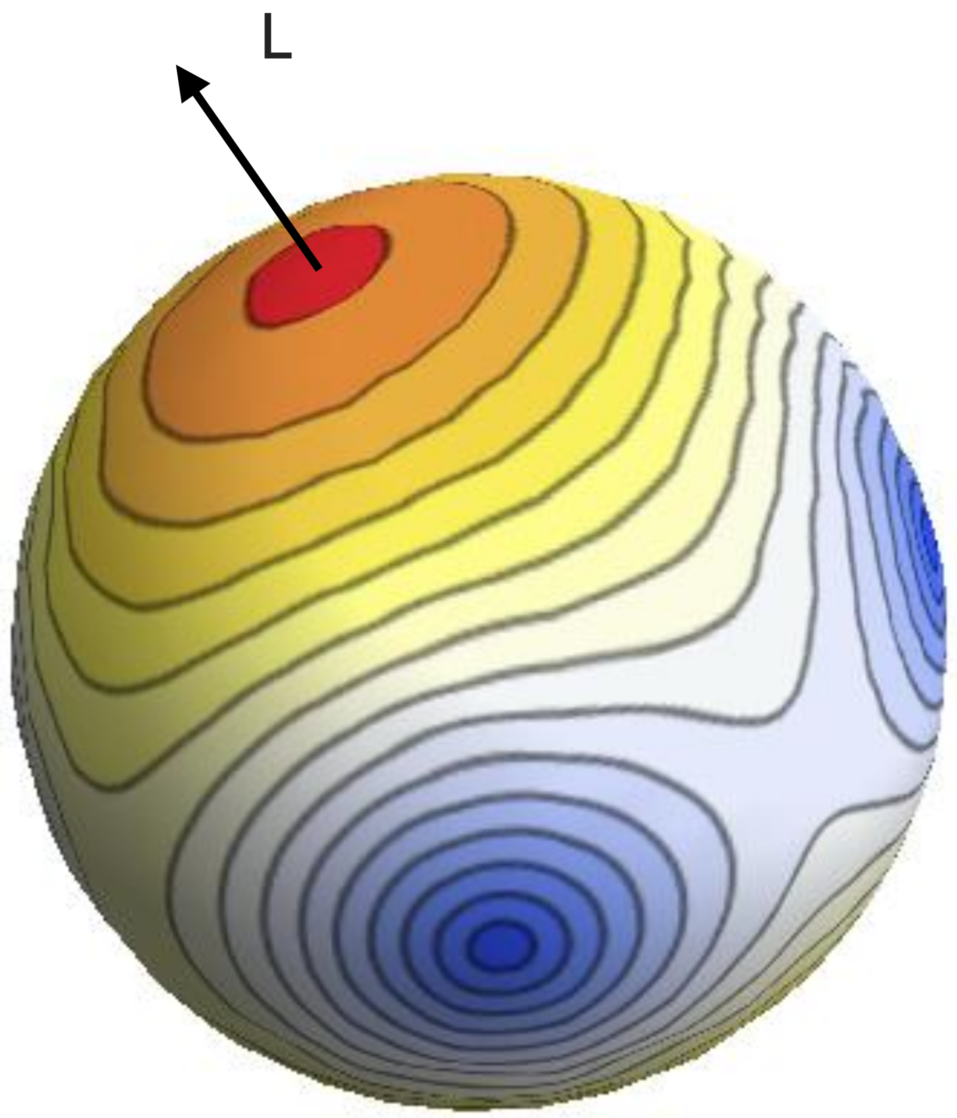

$${}^{-2}Y_{22}(\theta, \varphi) = \sqrt{\frac{5}{64\pi}}(1 + \cos\theta)^2 e^{2i\varphi}$$

$${}^{-2}Y_{21}(\theta, \varphi) = -\sqrt{\frac{5}{64\pi}} \sin\theta(1 + \cos\theta)e^{i\varphi}$$

MODEL ACCURACY IMR WITH SPIN PRECESSION

- Comparison of precessing Phenom and EOB models against NRSur7dq4:

$$\rho_{\mathrm{dist}} = \sqrt{\frac{D}{2(1-\mathcal{M})}}$$

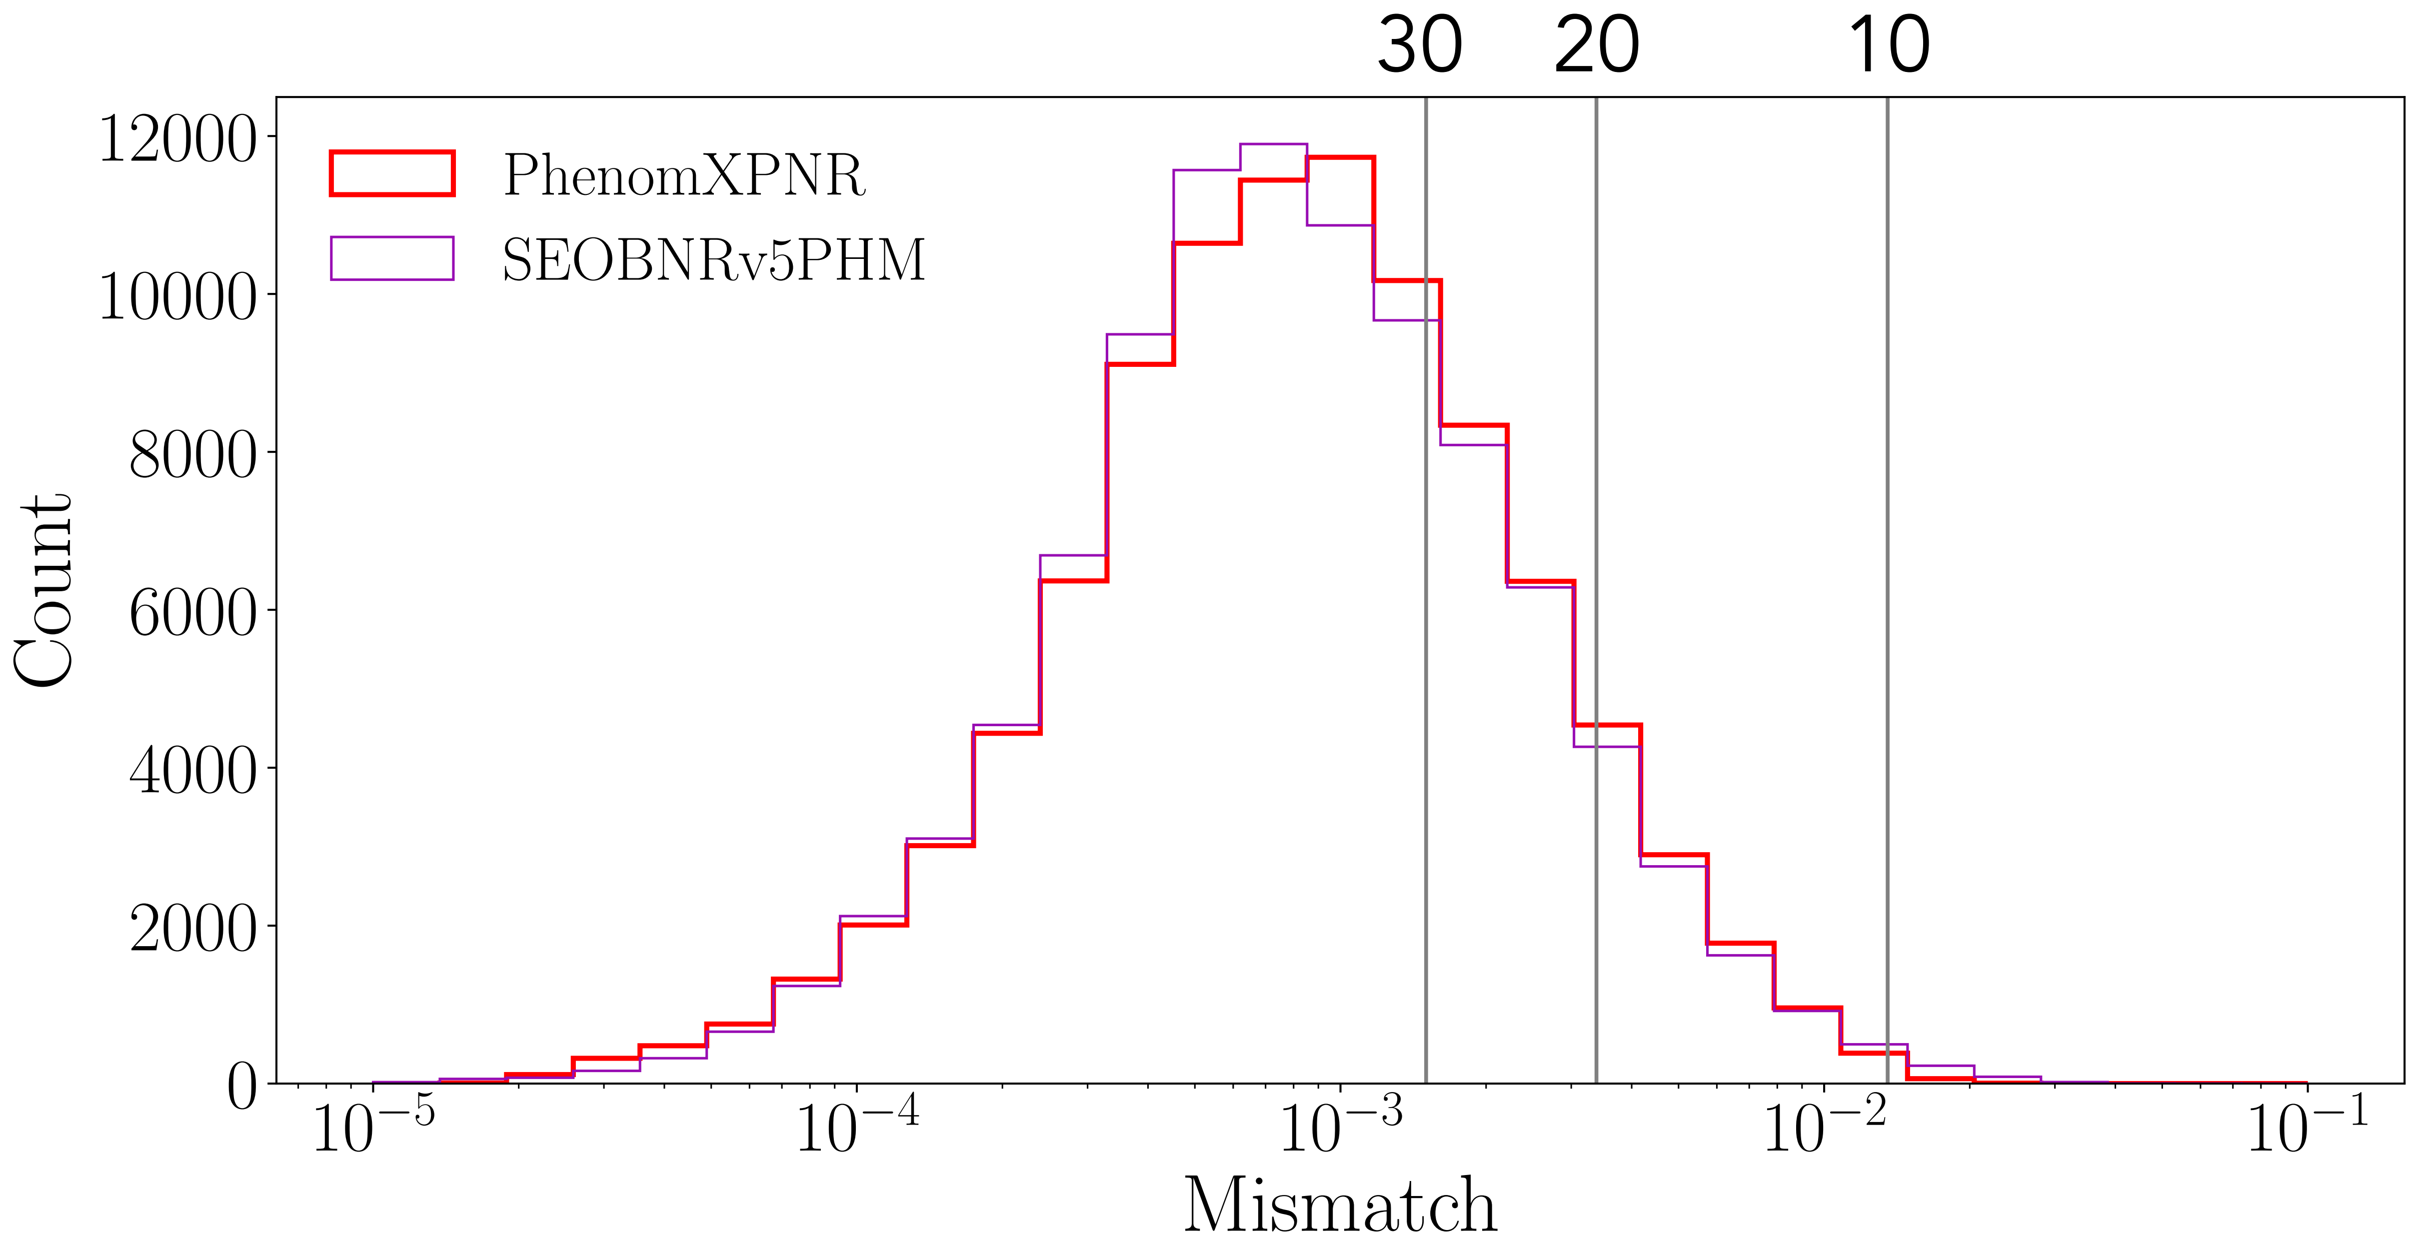

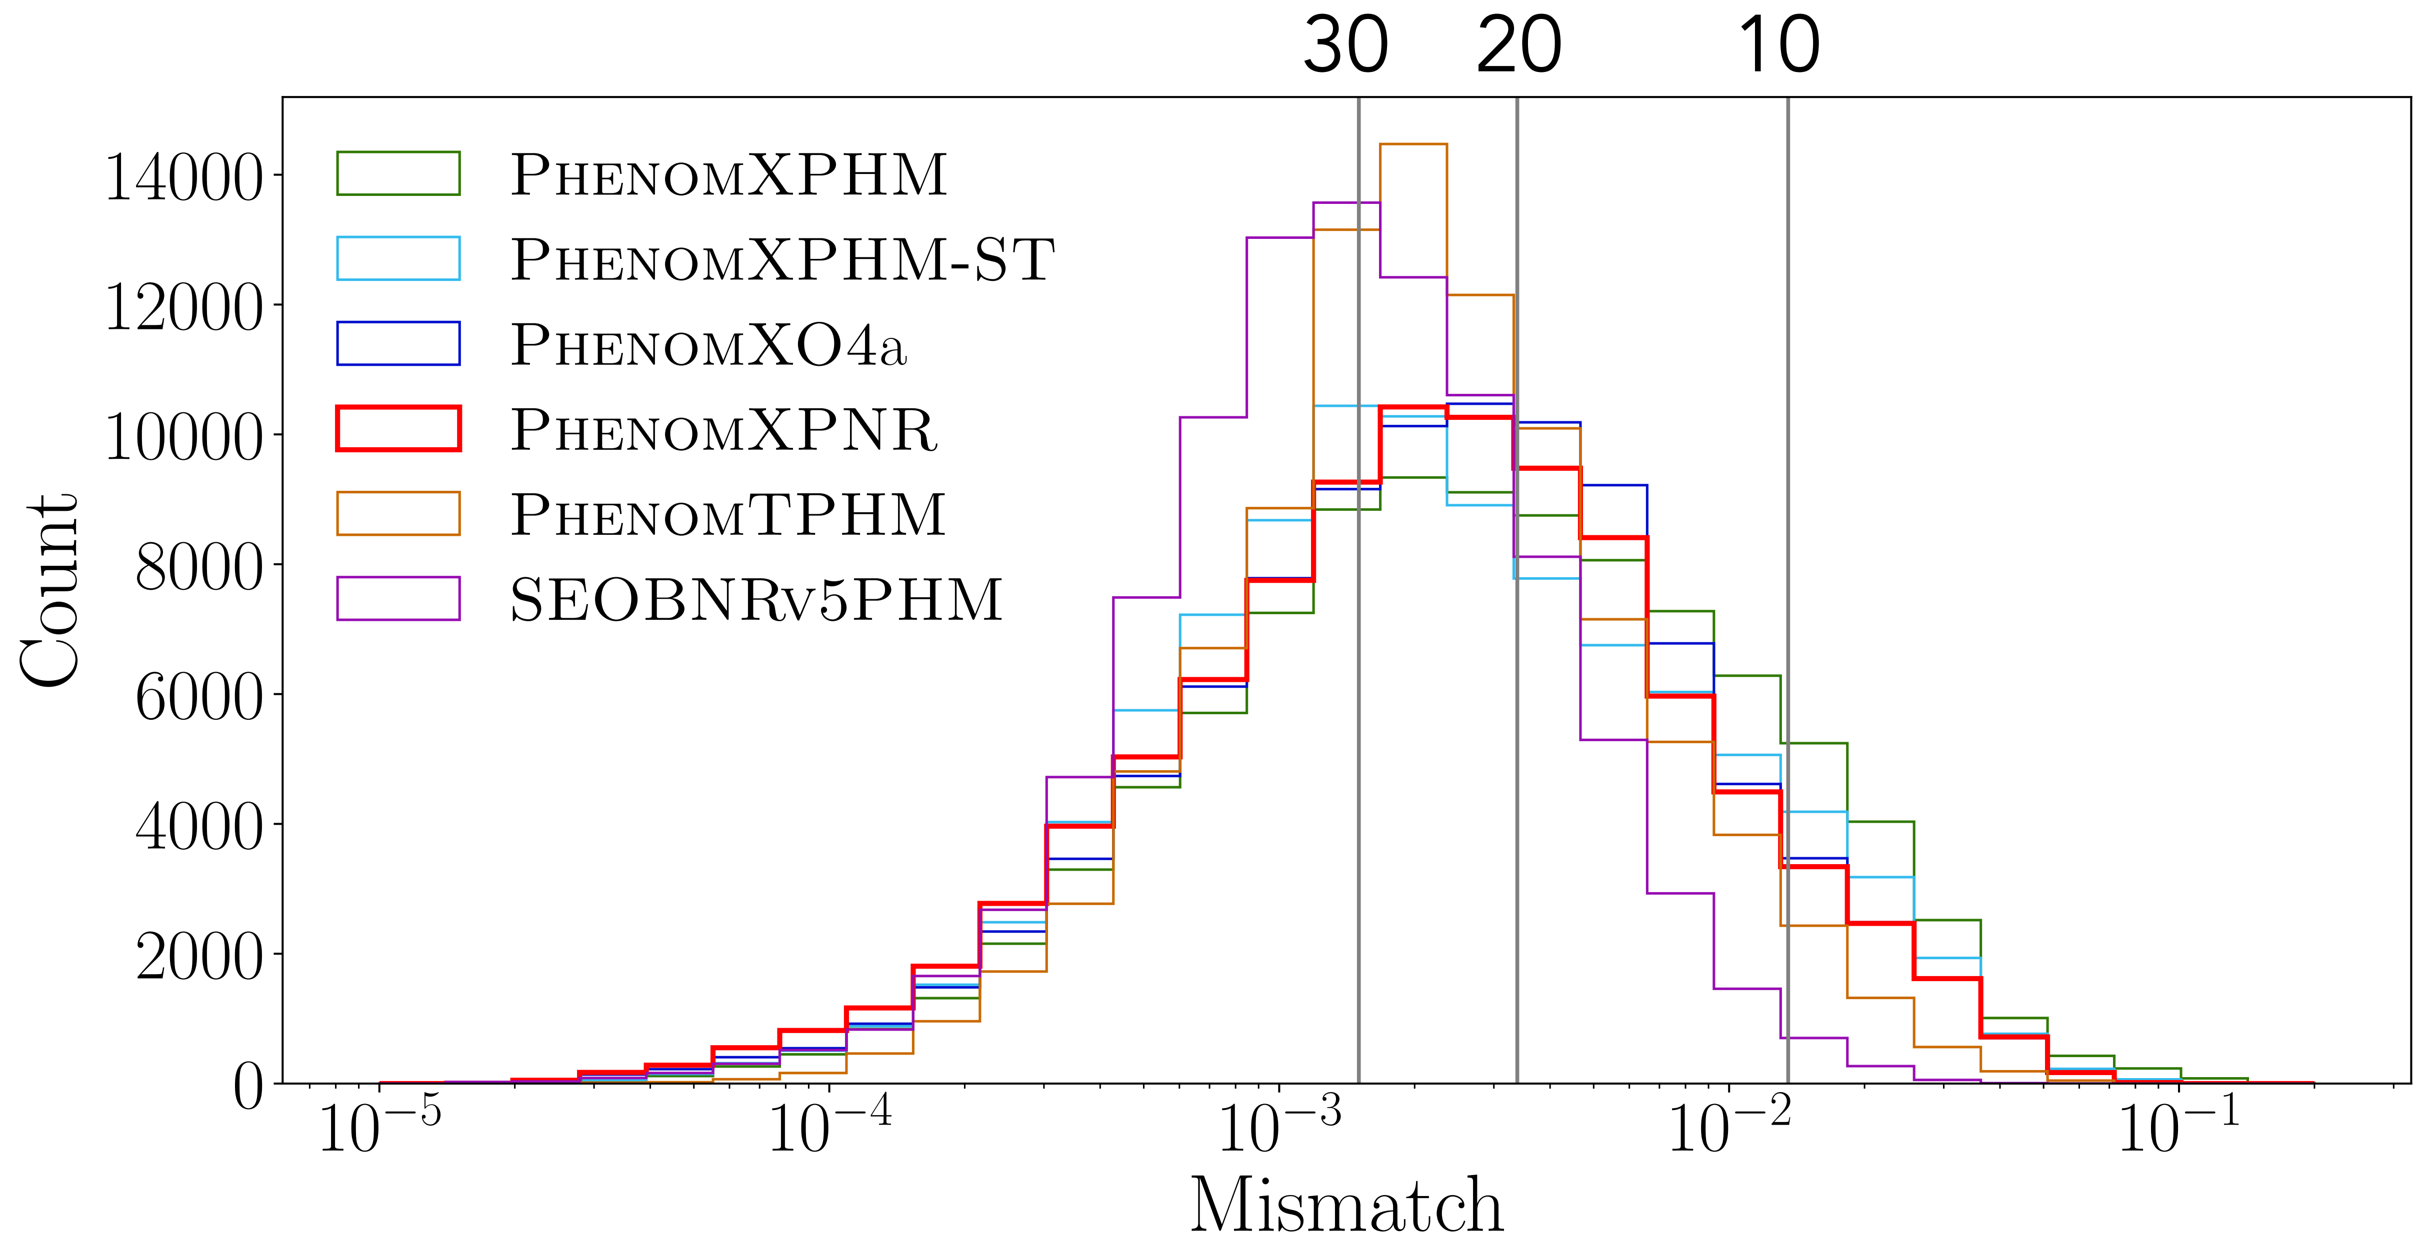

$\ell = 2$ multipoles only

[Hamilton+, 2025]

MODEL ACCURACY: NRSUR

Difference between NR resolutions

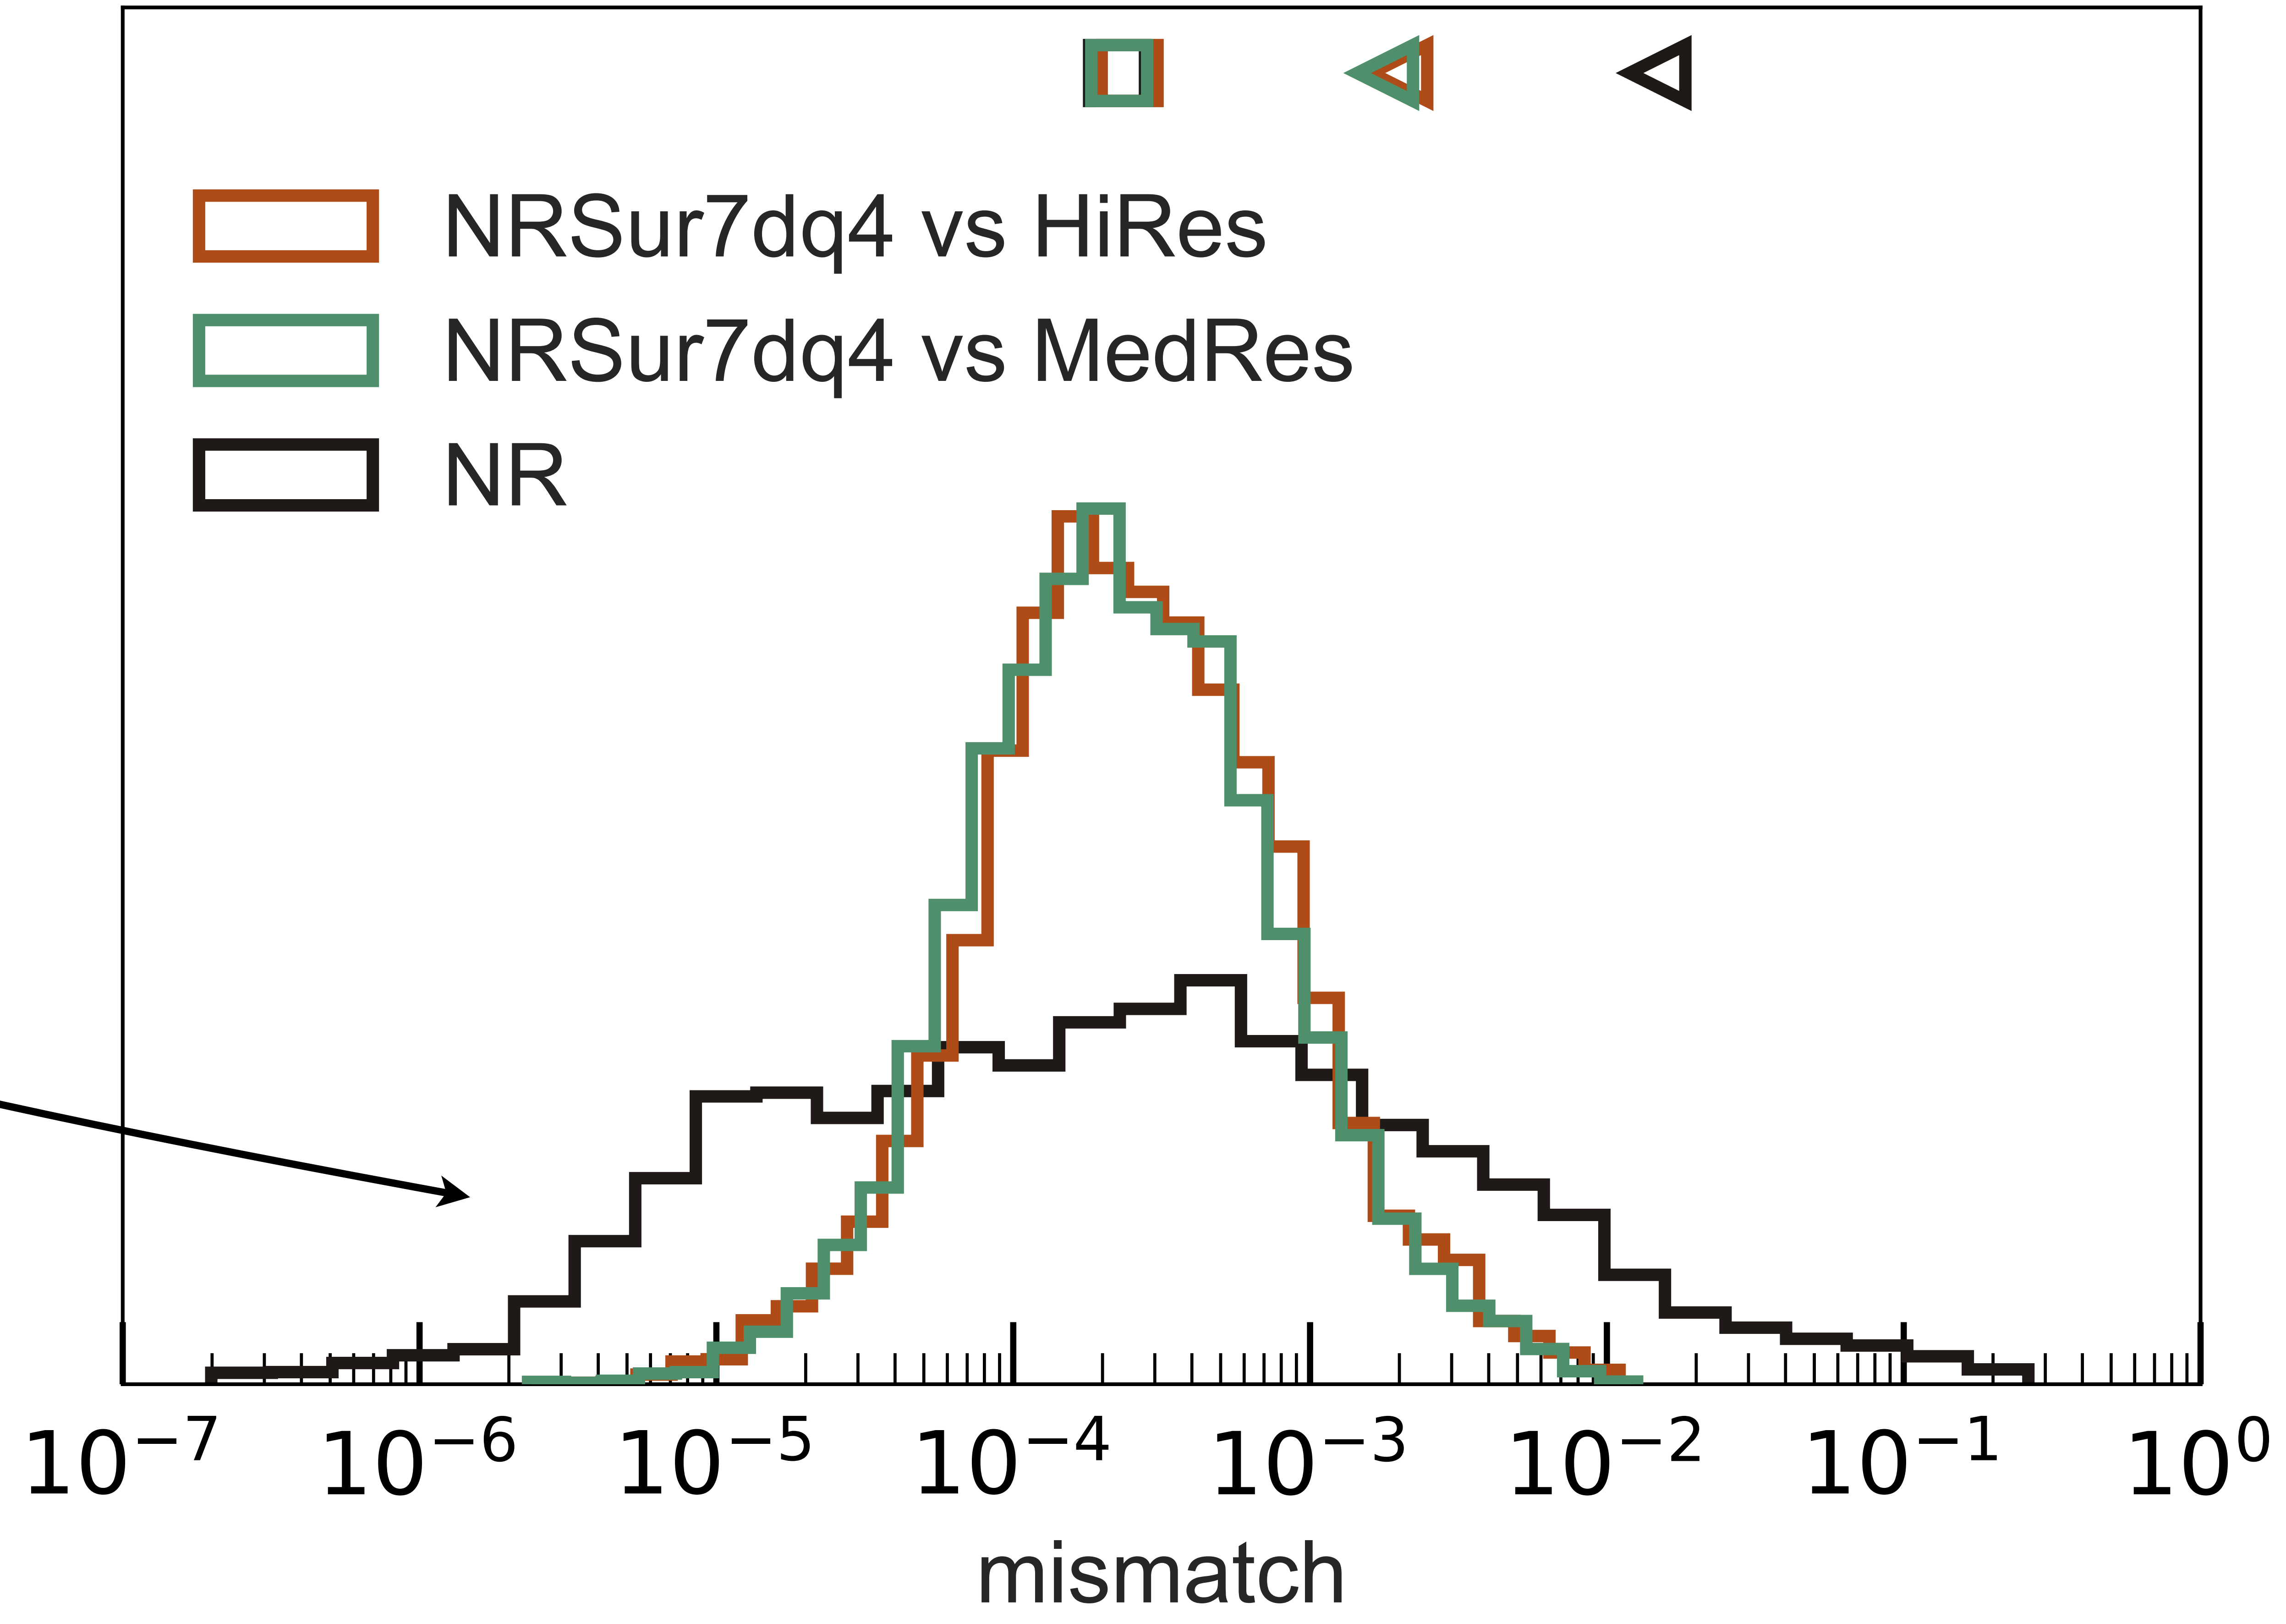

[Varma+ arXiv:1905.09300]

INSPIRAL-MERGER-RINGDOWN WAVEFORM MODELS CHEATSHEET

	Phenom	EOBNR	NR Surrogates
Domain	FD & TD	TD (&FD)	TD
Physics	Precession Higher modes Some tides Eccentricity (AS)	Precession Higher modes Some tides Eccentricity (AS)	Precession Higher modes Eccentricity (NS)
Efficiency	Fast	Medium to slow	Slow
Accuracy	Good	Good	Best
Caveats	Systematics due to approximations & missing physics	Systematics due to approximations & missing physics	Only for massive BBH (can be hybridised) Limited parameter space

WHAT PHYSICS IS INCLUDED?

- BBH models: Highly accurate for circular binaries with aligned spins
  - Including higher order multipoles
  - Verified up to moderate mass ratios + additional gravitational self-force information
  - Lack of simulations with large spins & long inspirals
  - Some eccentricity (lots of ongoing work & recent progress)

Precessing BBH: modelled only approximate

- Little to no calibration to NR

BNS: low-order analytic f-mode tides + some calibration to NR simulations

- No complete model including post-merger

NSBH: no precession; no complete model including post-merger

BEWARE OF SYSTEMATICS DUE TO MISSING PHYSICS!

GROWING EVIDENCE FOR NON-CIRCULARITY

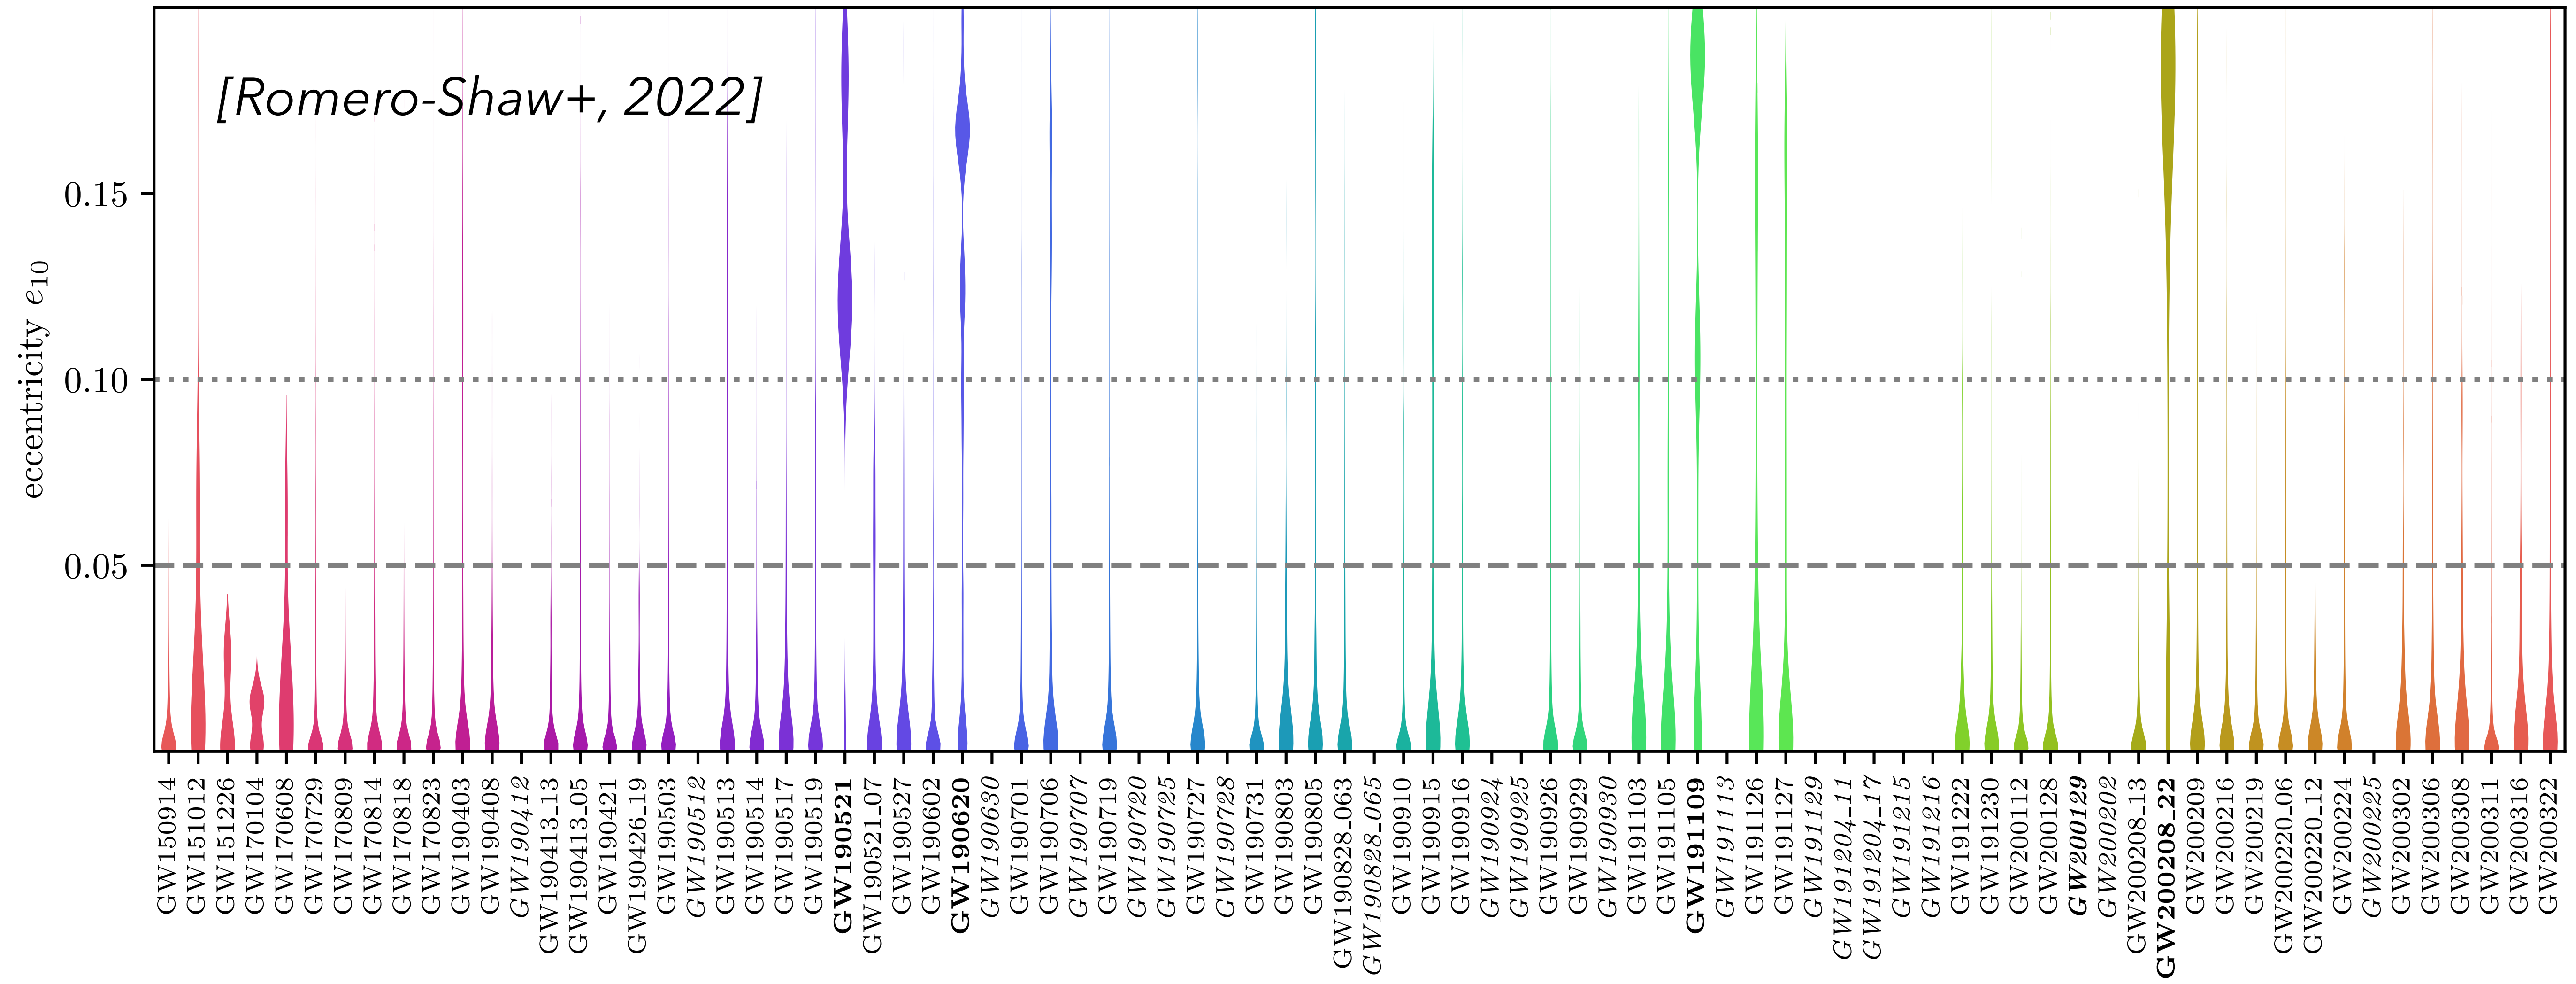

[Morras+, 2026]

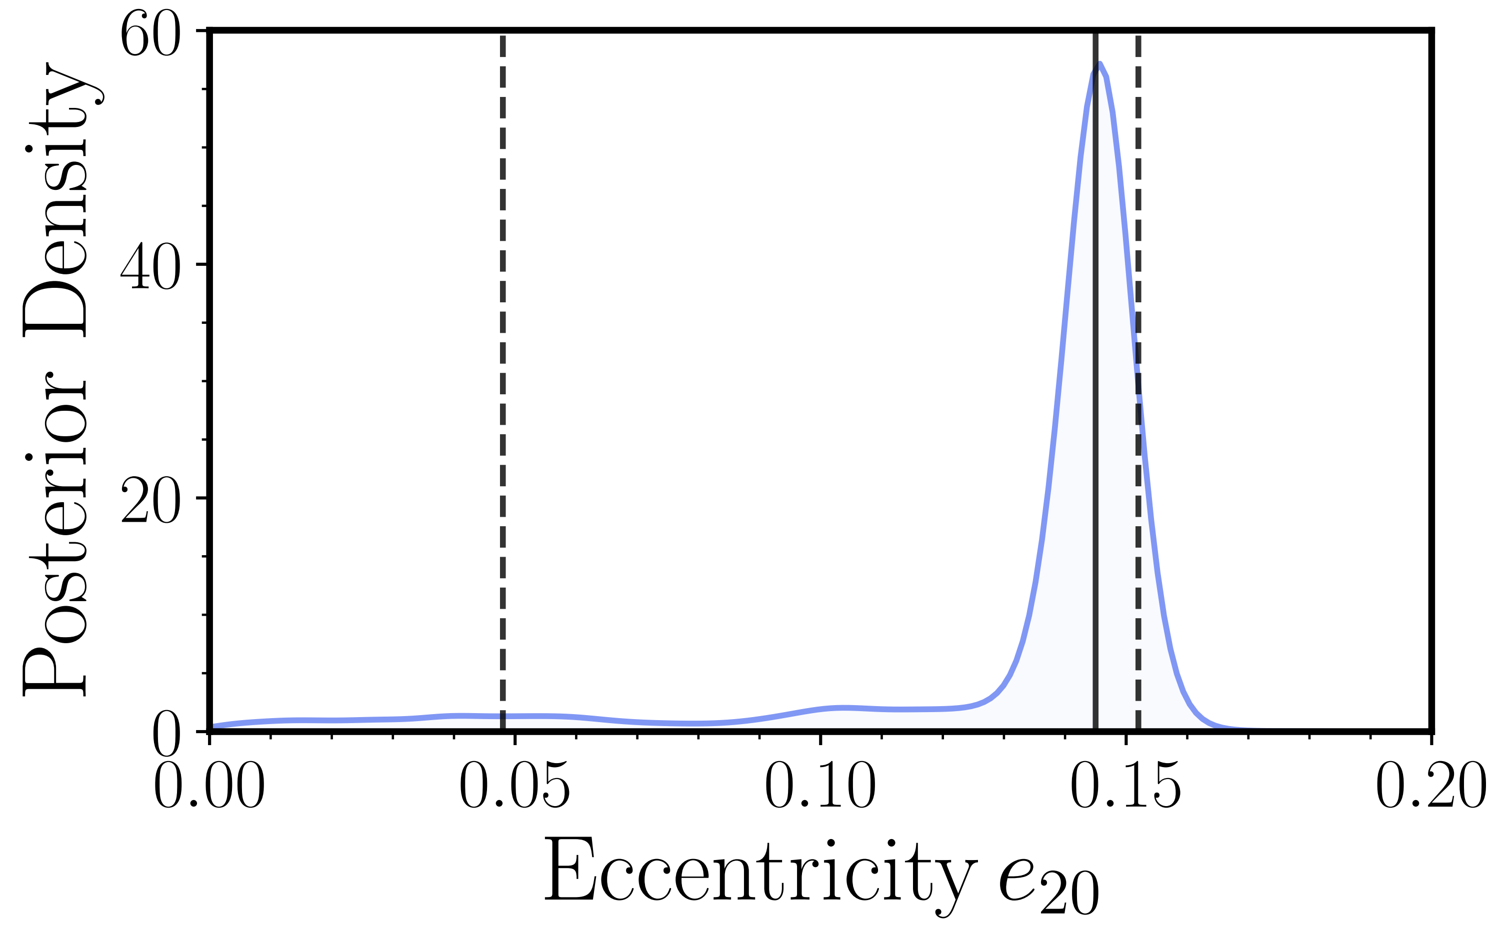

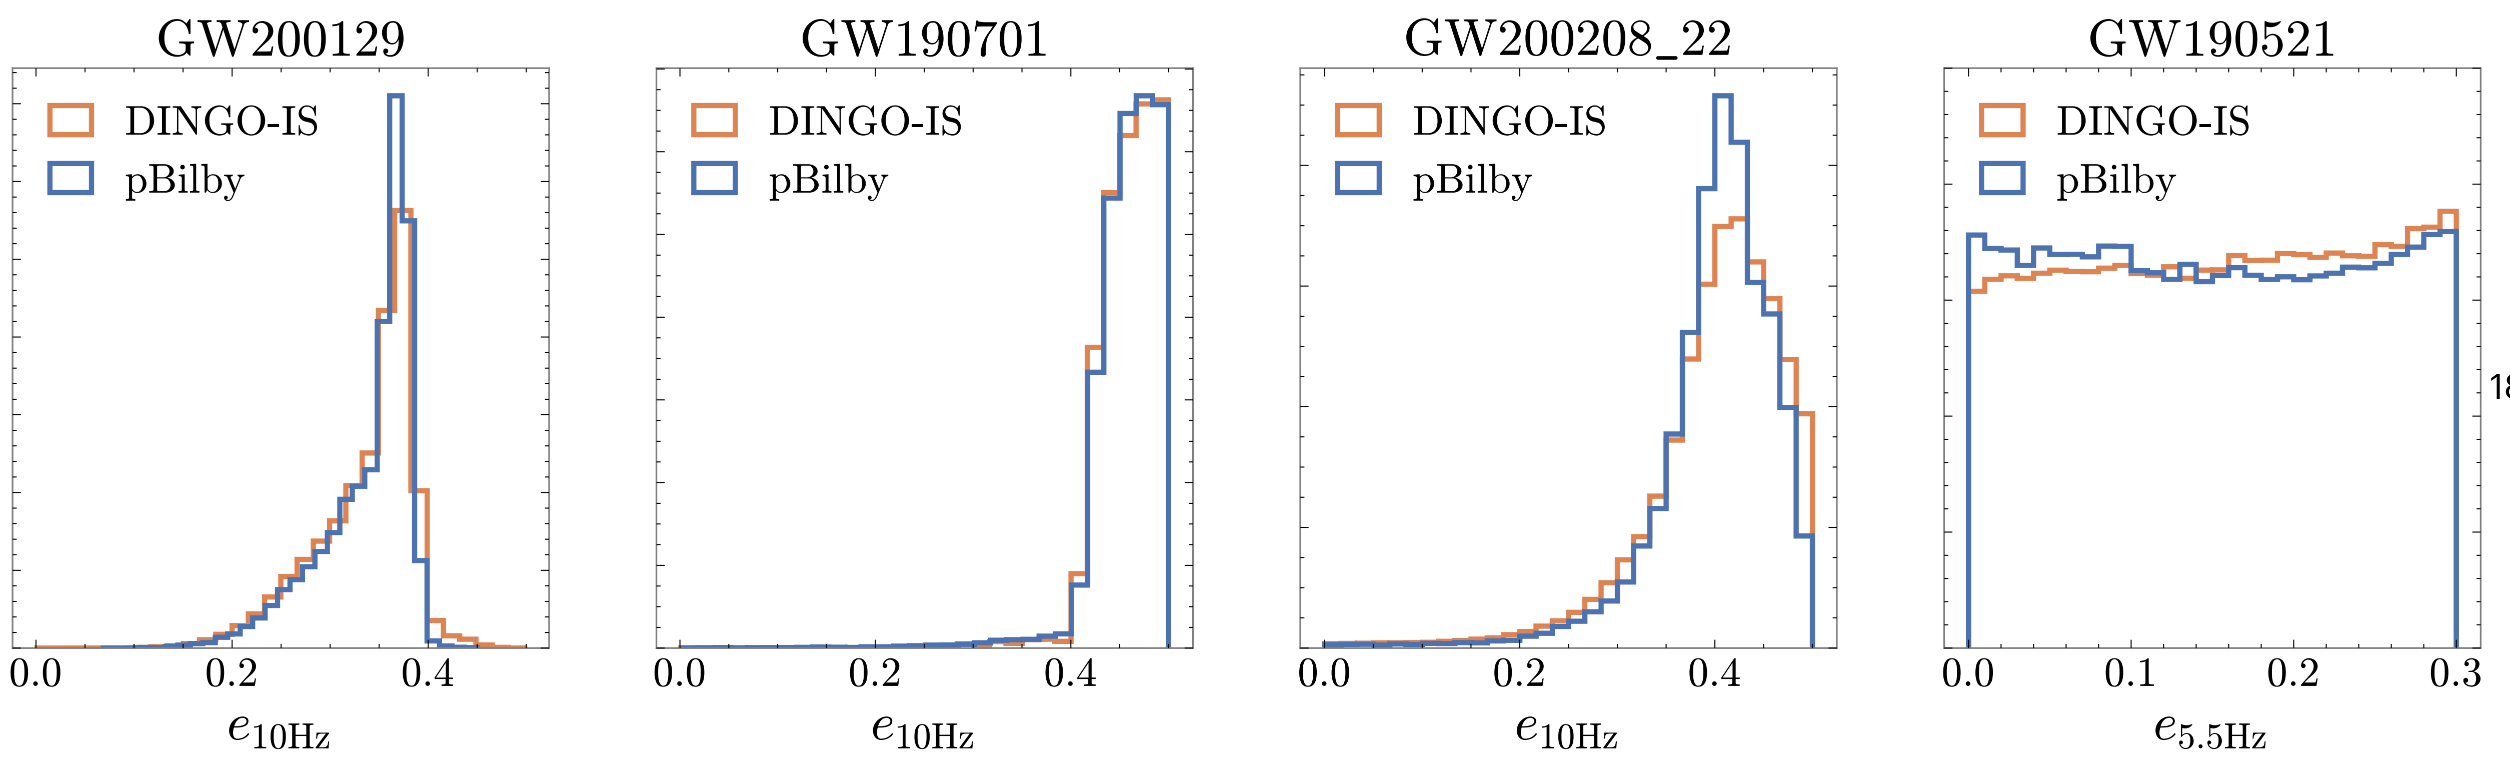

[Gupte+, 2024]

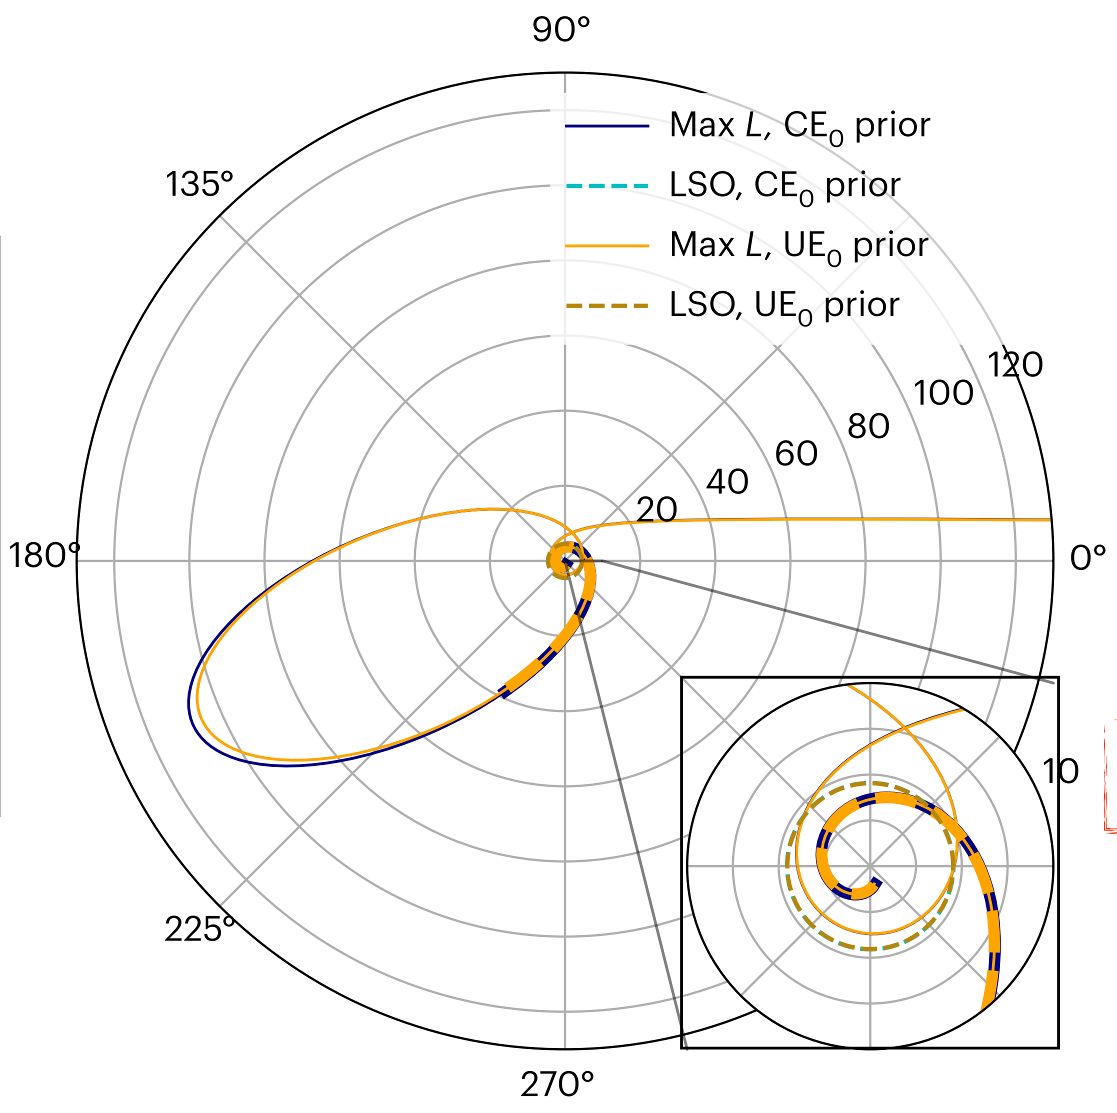

[Gamba+, 2023]

[Gayathri+, 2022]

Table 1 | Reconstructed properties of GW190521. For comparison we also show the properties obtained by LIGO-Virgo[20] using the NRSur7dq4 non-eccentric, precessing waveform model[61]. Error bars show 90% credible intervals

Parameters	This work	LIGO-Virgo
Primary mass ($M_\odot$)	78^{+9}_{-5}	95^{+29}_{-19}
Secondary mass ($M_\odot$)	78^{+9}_{-5}	69^{+22}_{-24}
Total mass ($M_\odot$)	155^{+17}_{-11}	164^{+39}_{-23}
Mass ratio[a]	1	$0.75^{+0.22}_{-0.35}$
Luminosity distance (Gpc)	$7.7^{+1.27}_{-1.65}$	$3.9^{+2.2}_{-2.0}$
Redshift	$1.13^{+0.15}_{-0.2}$	$0.68^{+0.28}_{-0.28}$
Eccentricity	$0.69^{+0.17}_{-0.22}$	0
Effective spin $\chi_{\rm eff}$	$0.27^{+0.56}_{-0.51}$	$0.03^{+0.31}_{-0.40}$
Precession spin $\chi_{\rm p}$	$0.66^{+0.28}_{-0.60}$	$0.67^{+0.26}_{-0.44}$
Sensitive volume[b] (Gpc^3)	46.5	25.3

[a]Reported here for maximum-likelihood waveform (see refs. [19,32,34] for the possibility of GW190521 having a low mass ratio). [b]Computed for the maximum-likelihood waveform, marginalizing over total mass (Methods).