Extracting new physics from extreme mass ratio inspirals in LISA data

with James Alvey, Lorenzo Speri, Christoph Weniger, Uddipta Bhardwaj, Davide Gerosa and Gianfranco Bertone

arXiv:2505.16795

Philippa (Pippa) Cole, University of Milan-Bicocca

UNIVERSITÀ DEGLI STUDI DI MILANO BICOCCA

erc
European Research Council
Established by the European Commission

What's special about extreme mass ratio inspirals?

- A binary black hole system with mass ratio $q = m_2/m_1 \leq 10^{-4}$
- Rich dynamics due to eccentric orbits, spin precession and thousands of excited harmonics
- Long duration gravitational wave signals - could remain in LISA band for years - opportunity to observe millions of cycles

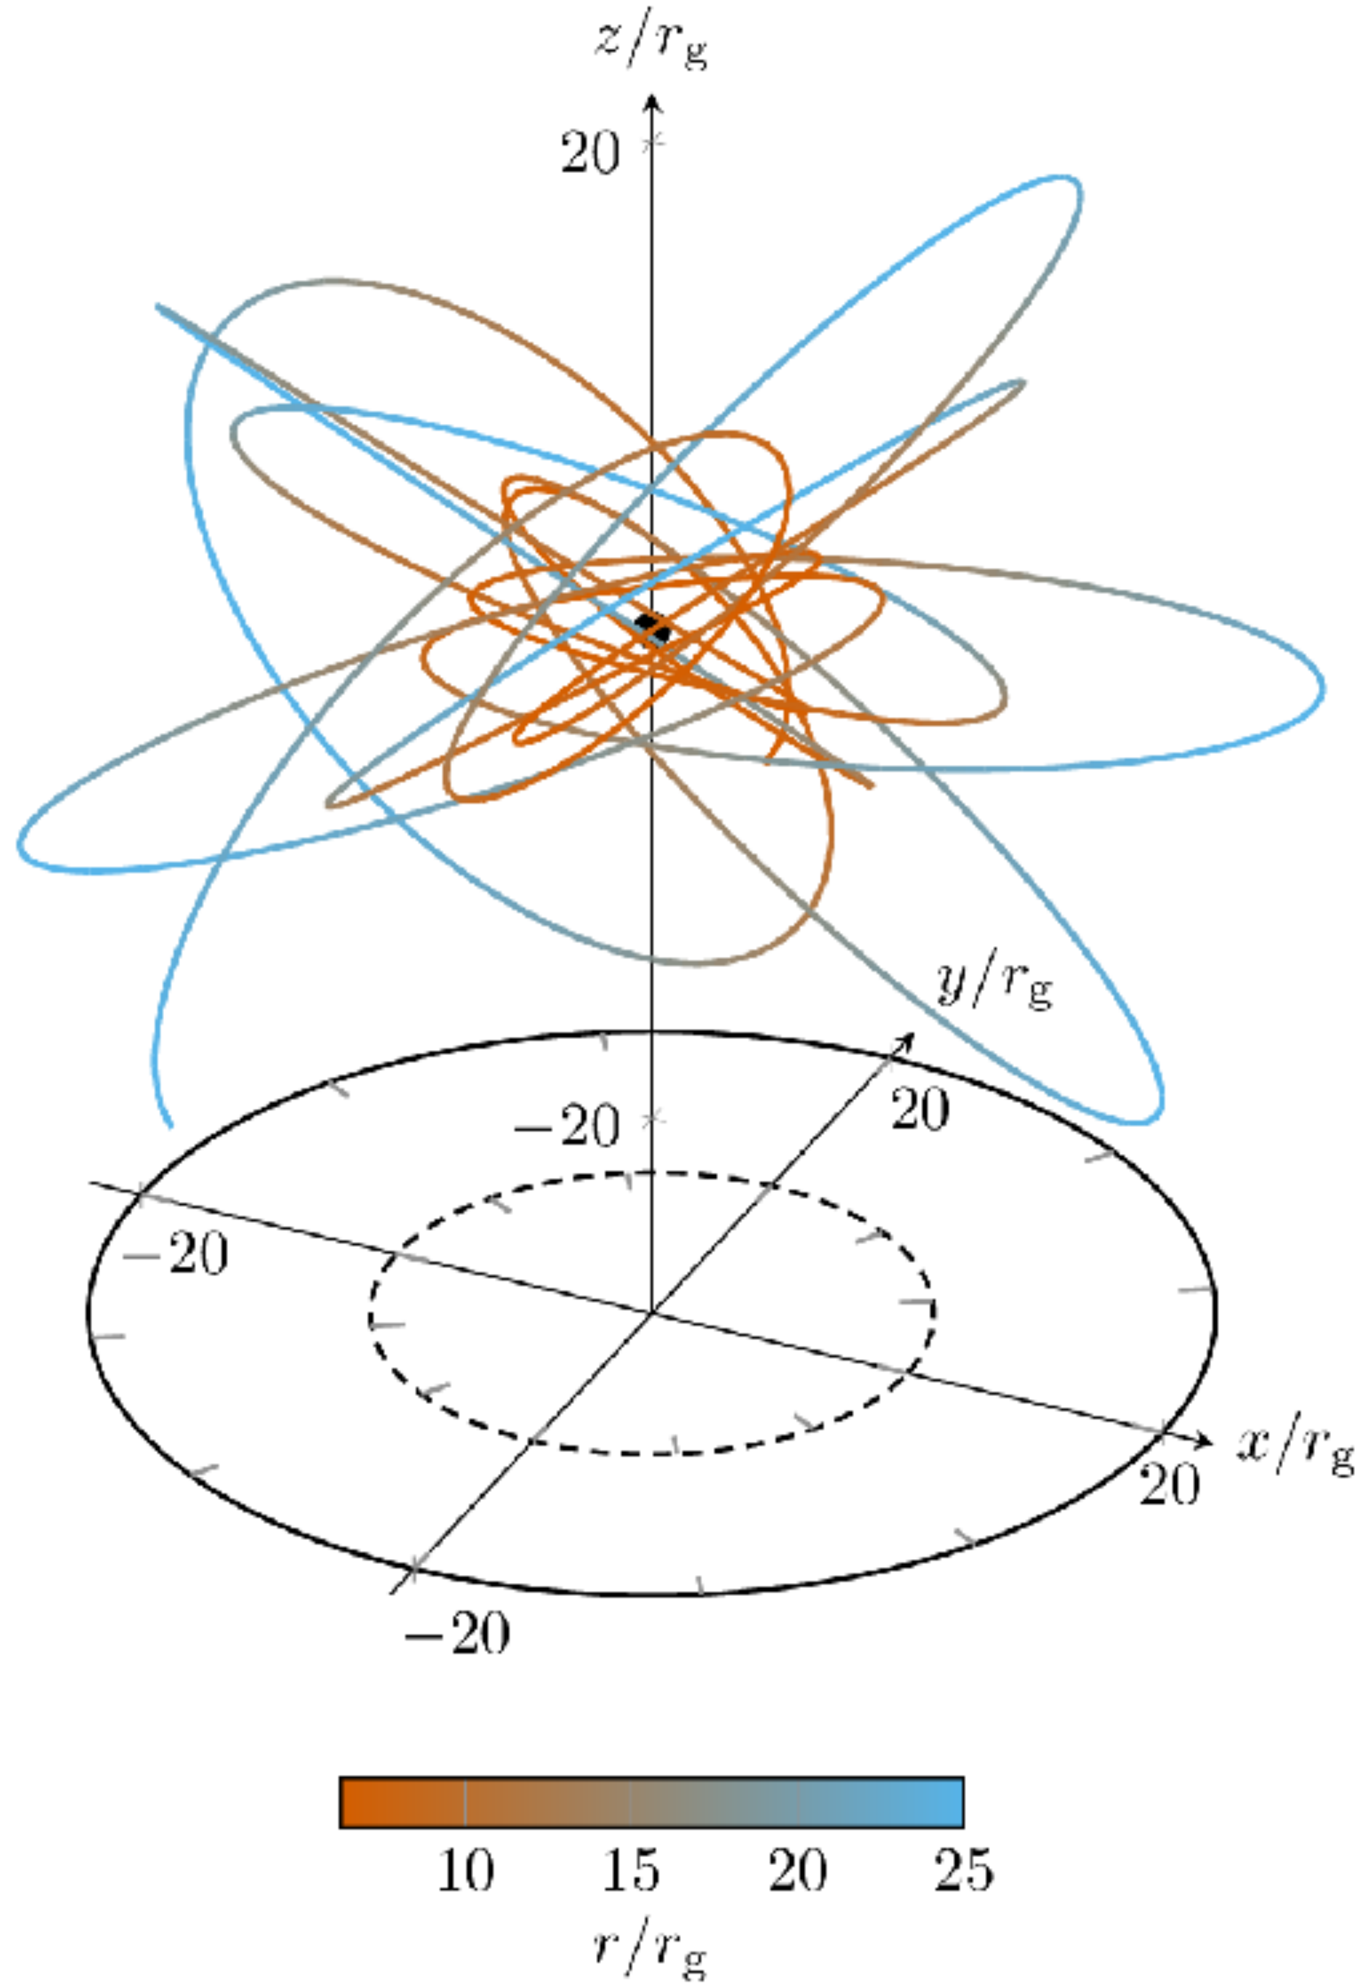

Berry et al. 2019

What’s special about extreme mass ratio inspirals?

- A binary black hole system with mass ratio $q = m_2/m_1 \leq 10^{-4}$
- Rich dynamics due to eccentric orbits, spin precession and thousands of excited harmonics
- Long duration gravitational wave signals - could remain in LISA band for years - opportunity to observe millions of cycles

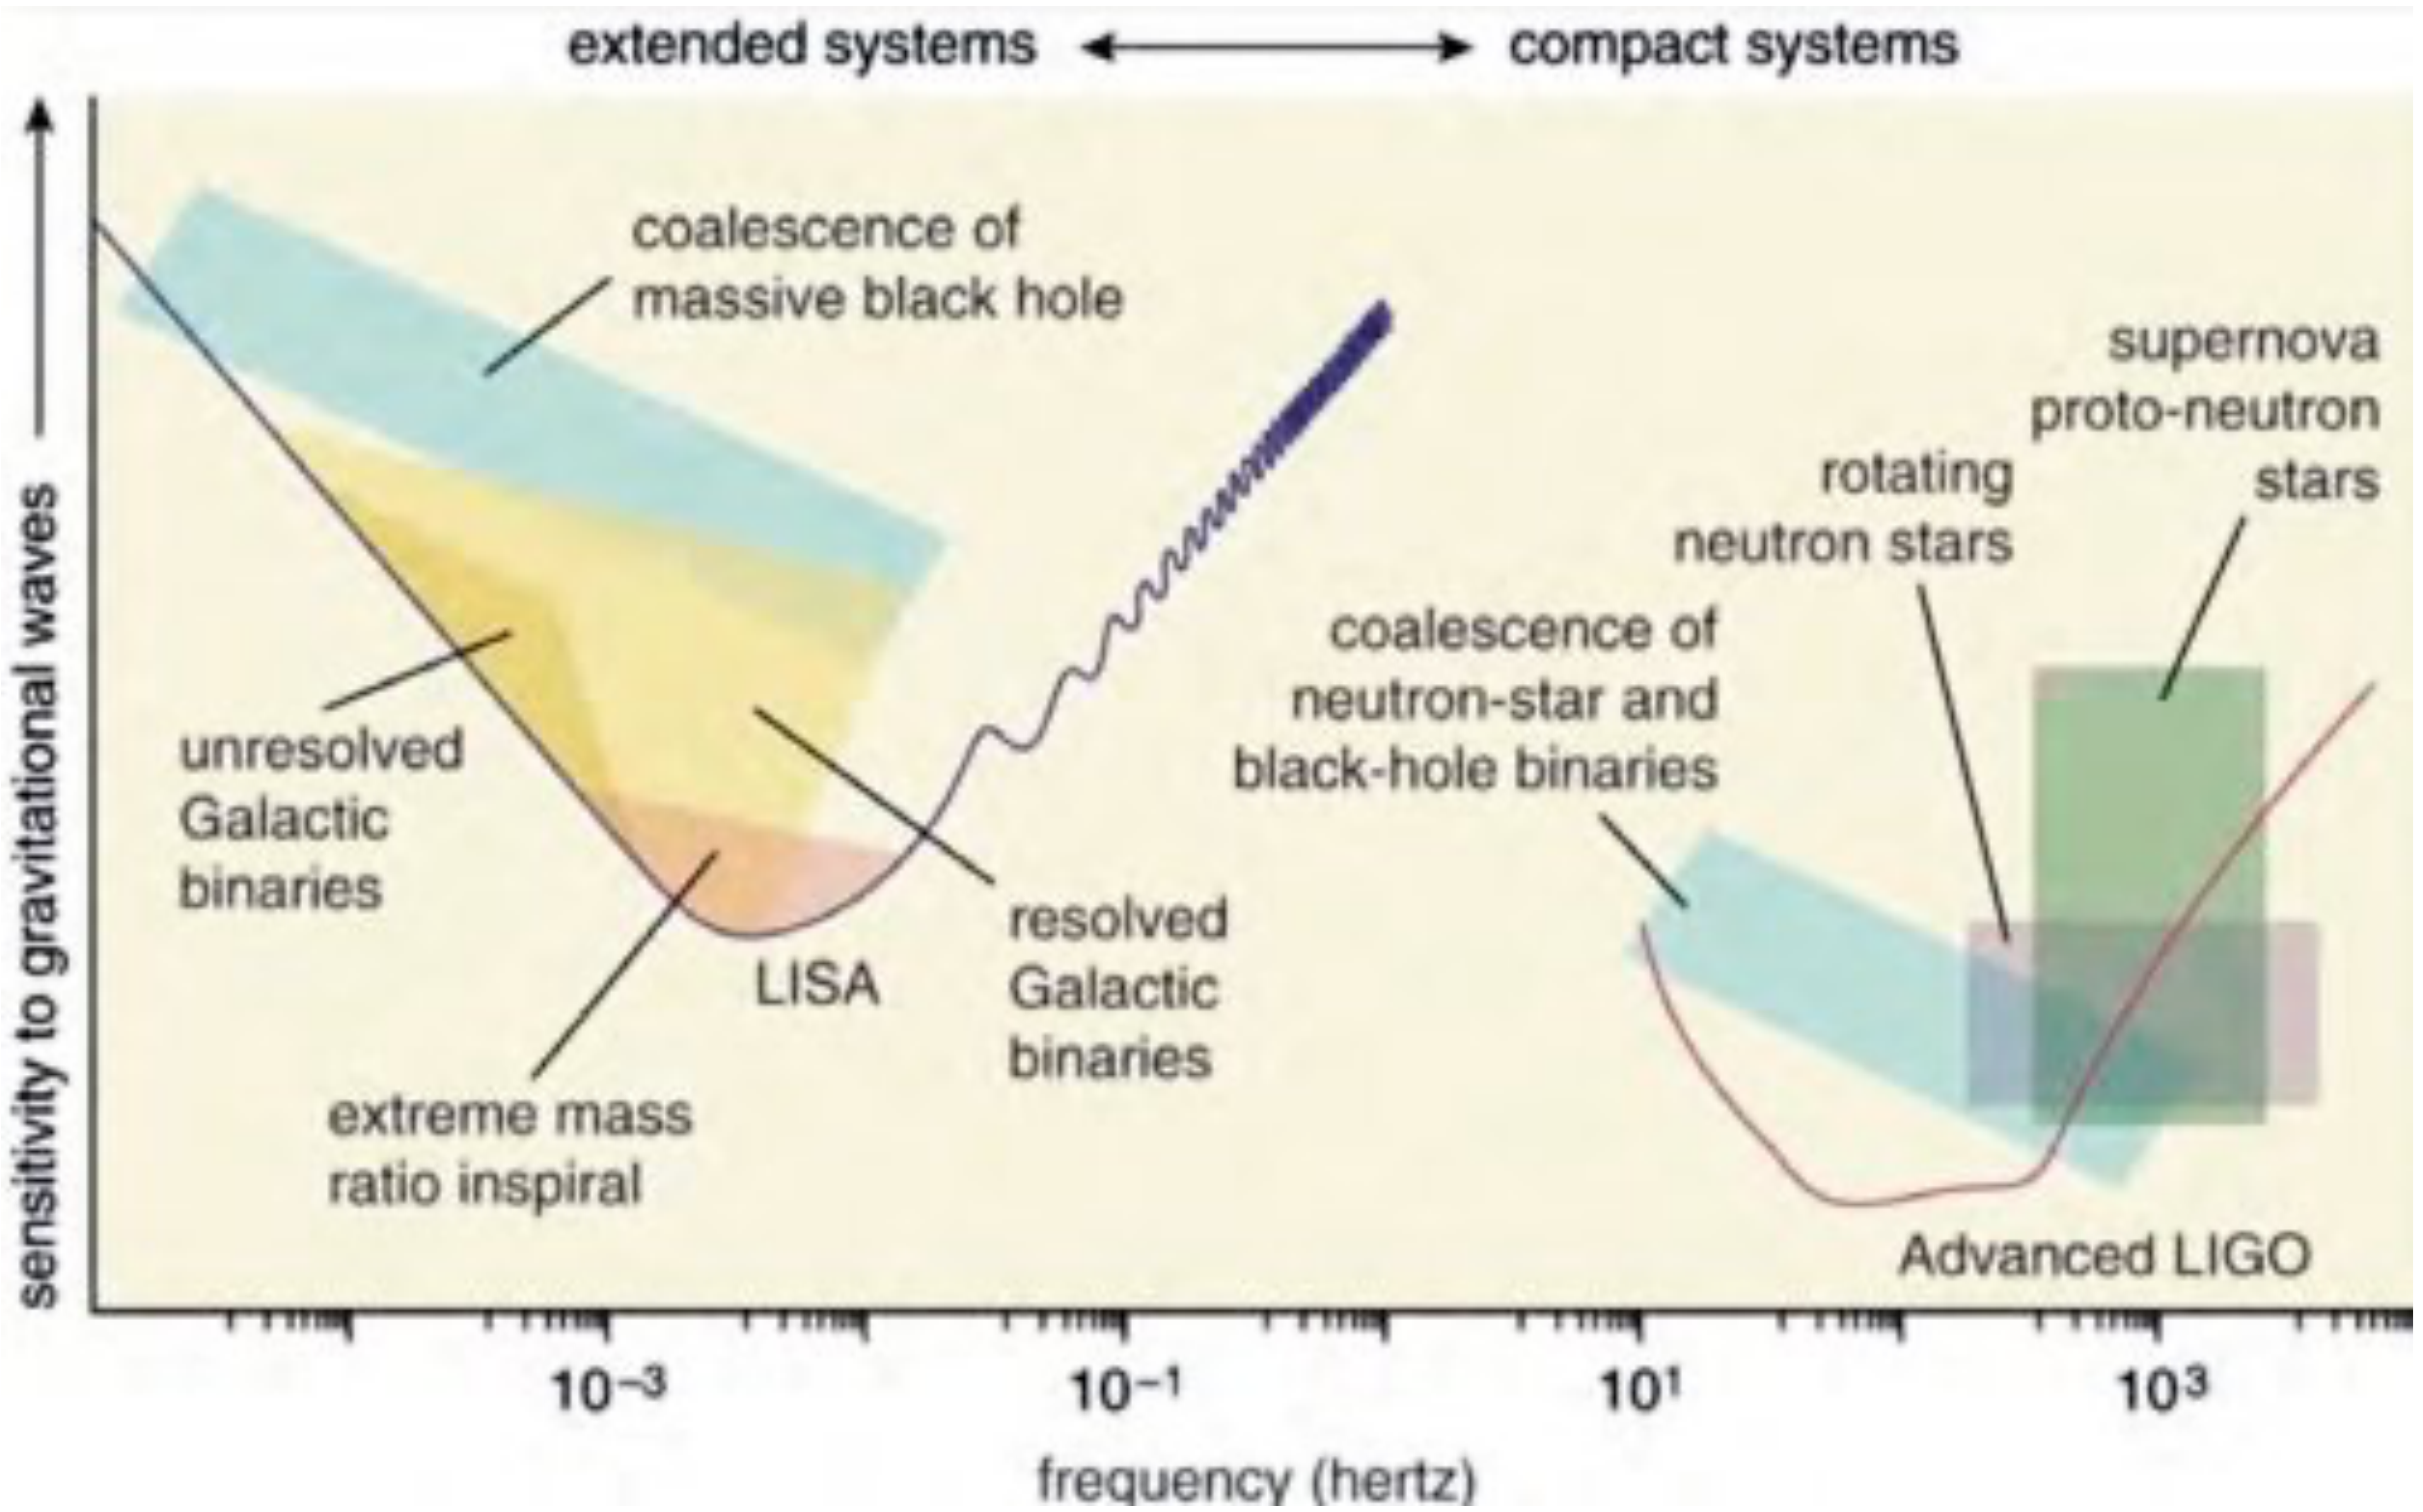

Key 2010

Many astrophysics and fundamental physics opportunities…

- Formation of intermediate mass black holes (Volonteri 2010)
- Environmental effects - dark matter, ultra-light bosons, accretion disks around primary object (Eda et al. 2013, Macedo et al. 2013, Barausse et al. 2014, Kavanagh et al. 2020, Khalvati et al. 2024, Baumann et al. 2022, Barsanti et al. 2023, Zhang et al. 2023, Tomaselli et al. 2024, Cole et al. 2023, Speri et al. 2023, Duque et al. 2024, Coppaoni et al. 2025…)
- Tests of General Relativity (Han & Chen 2019, Gupta et al. 2022, Speri et al. 2024, Kejriwal et al. 2024, A. Cardenas-Avendano & Sopuerta 2024)

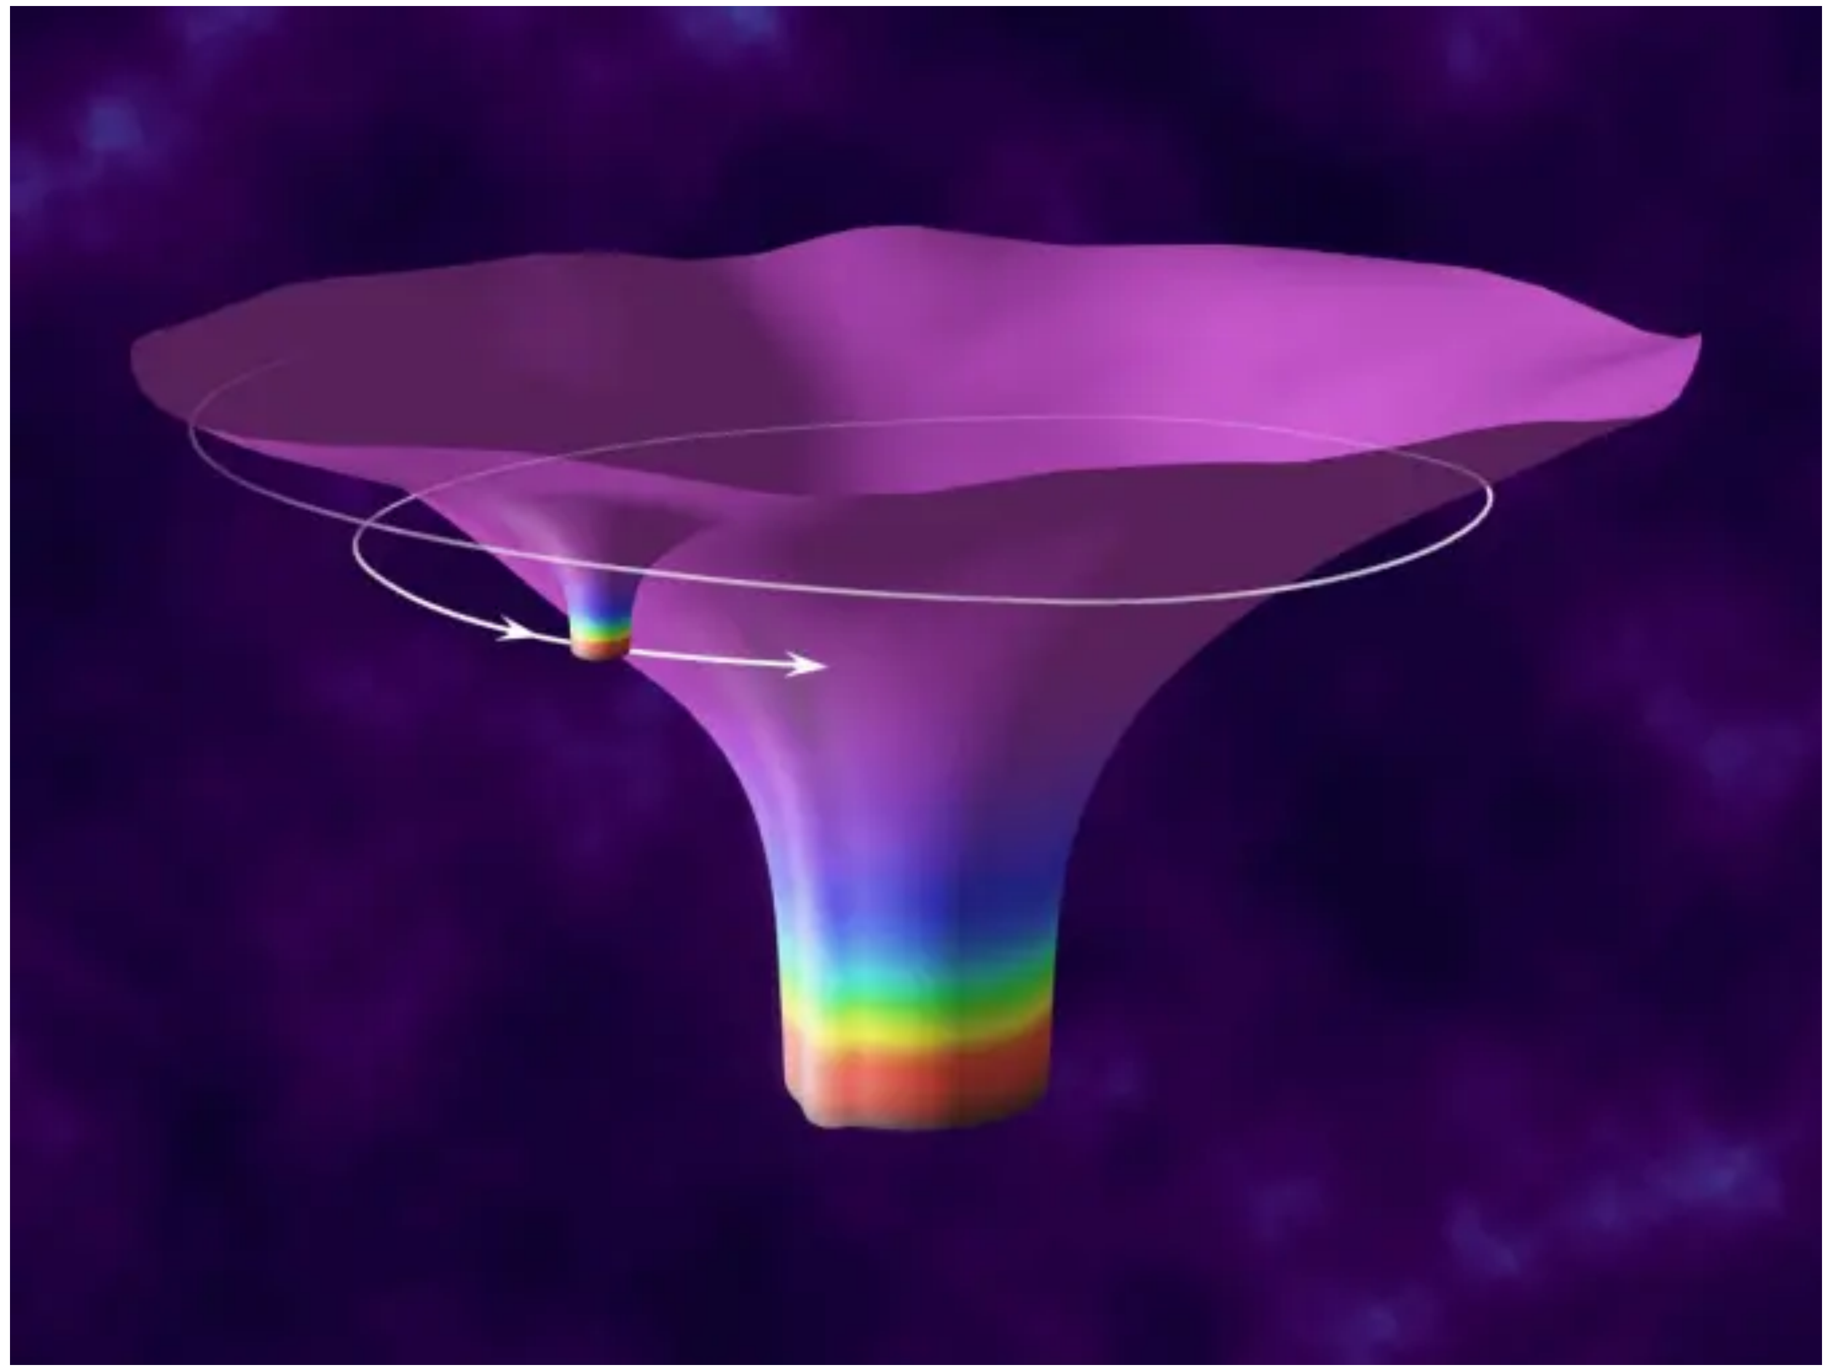

…if we can measure the parameters to very high precision

Credit: NASA

Even in vacuum, EMRIs have complicated waveforms that evolve significantly in time - enables precision measurement but expensive

Equal stellar-mass

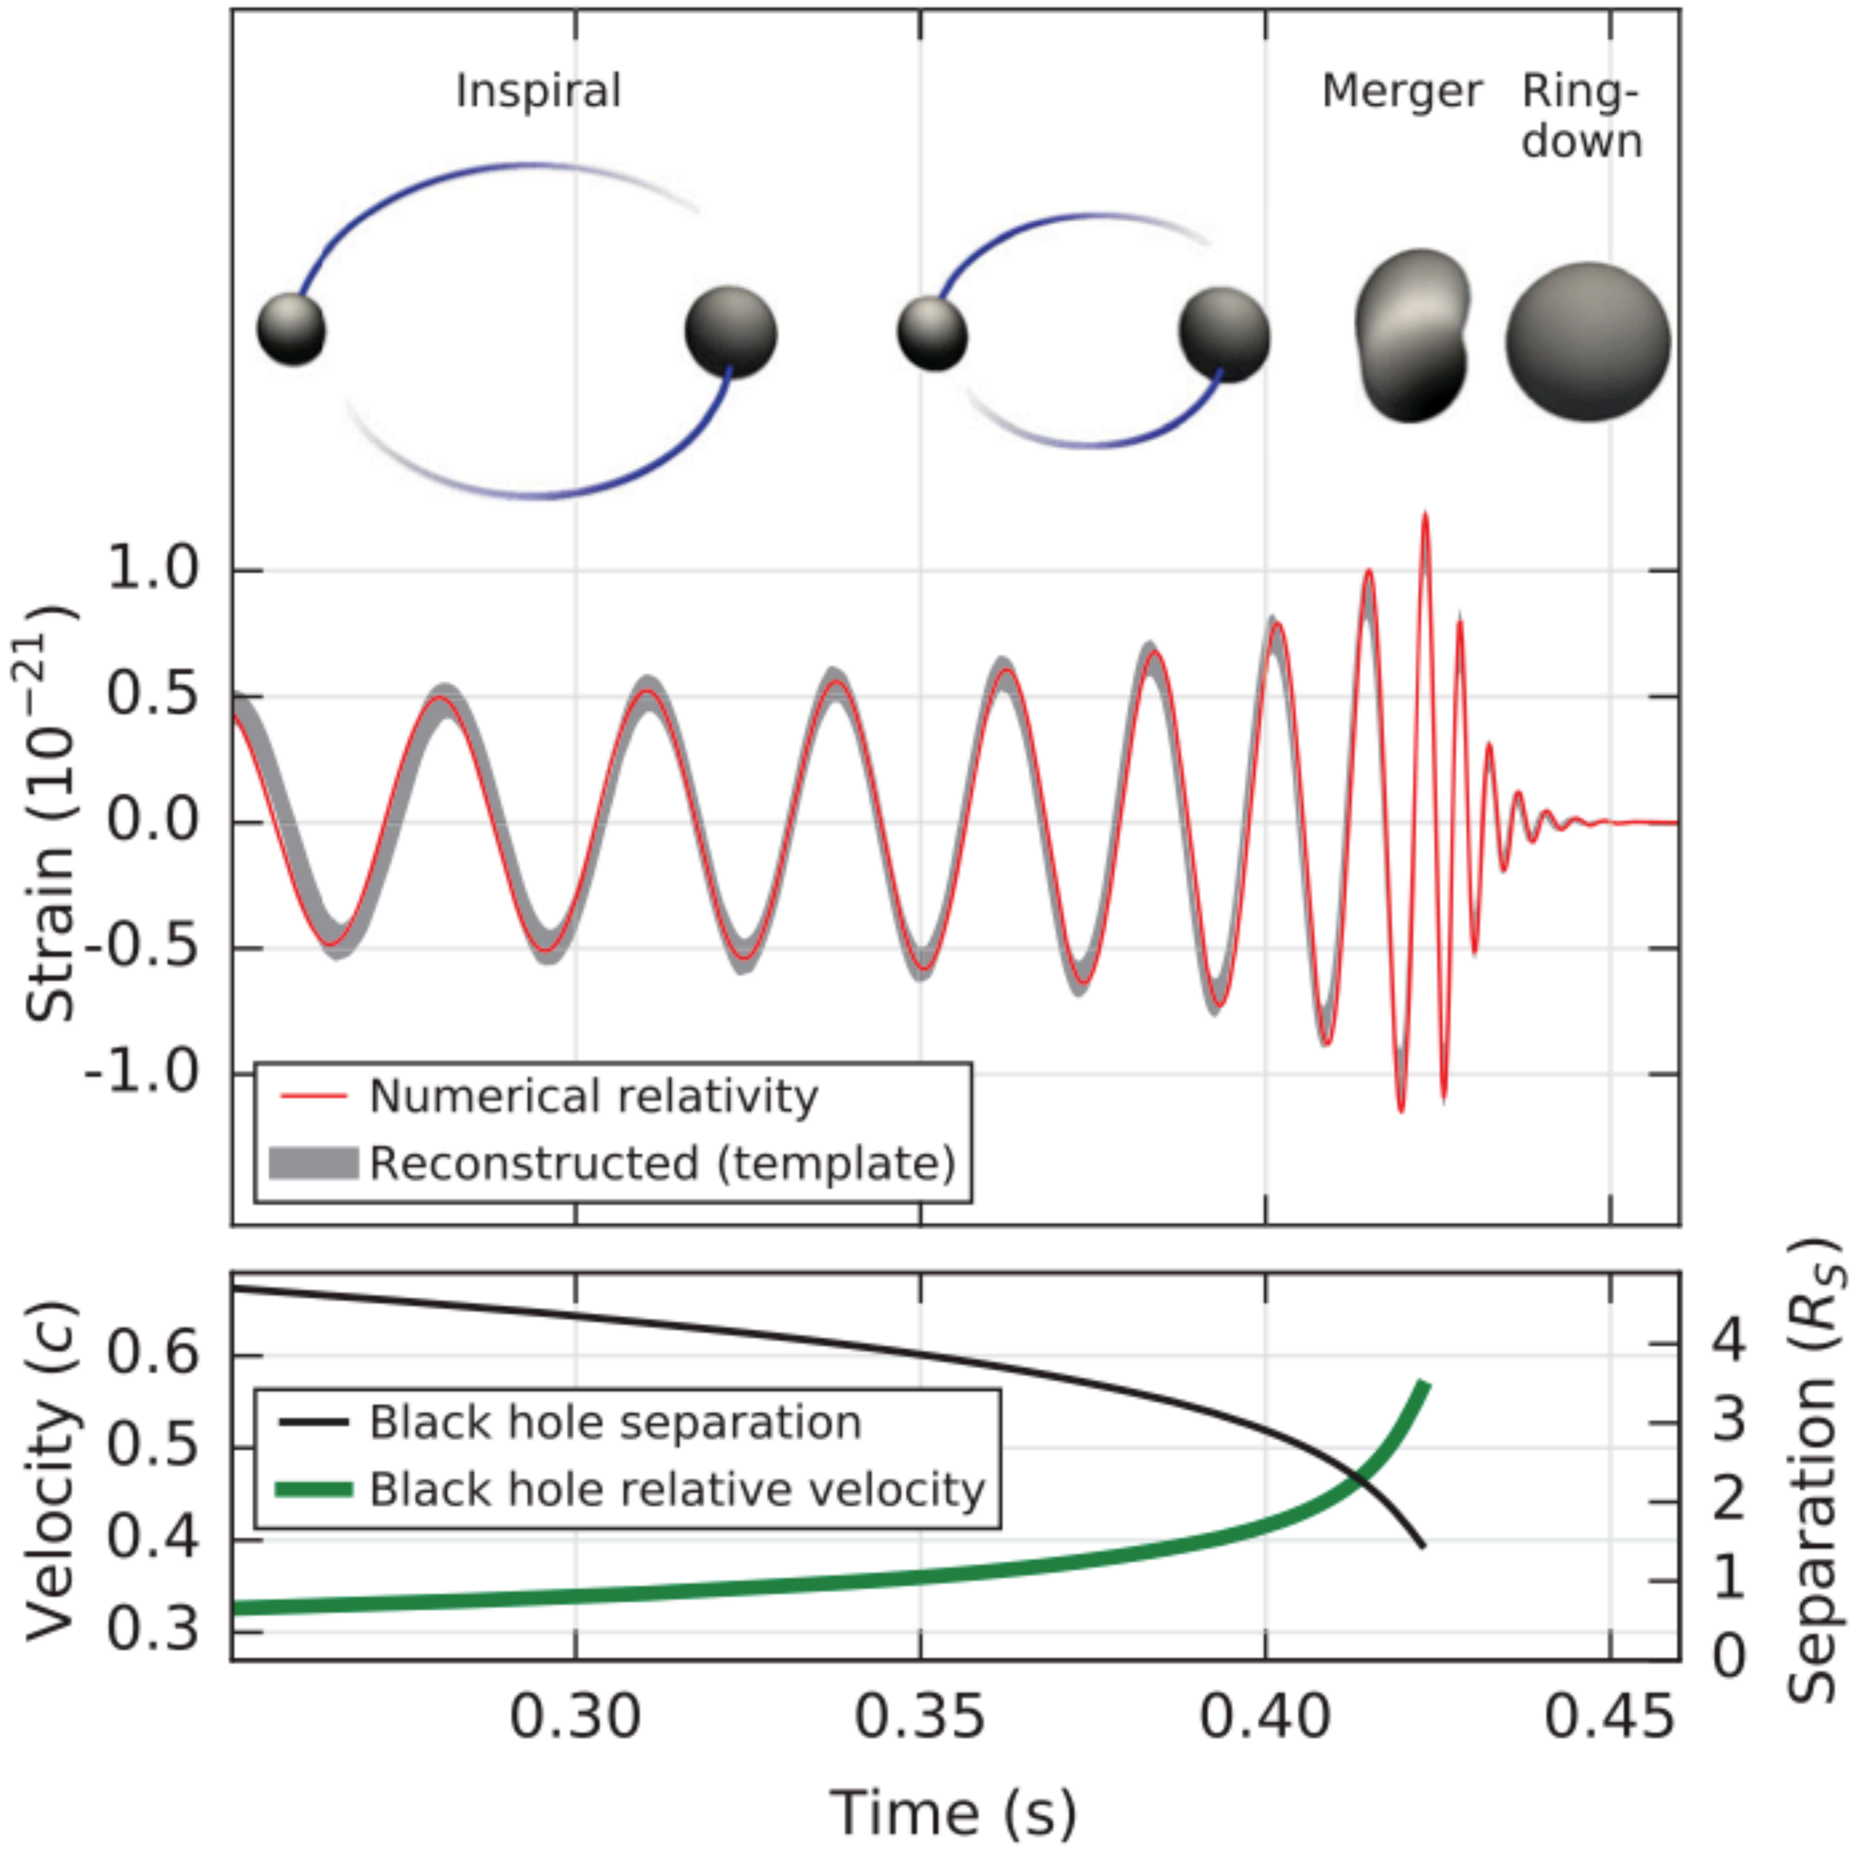

Abbott et al. arXiv:1602.03837

EMRIs

One year before ‘plunge’

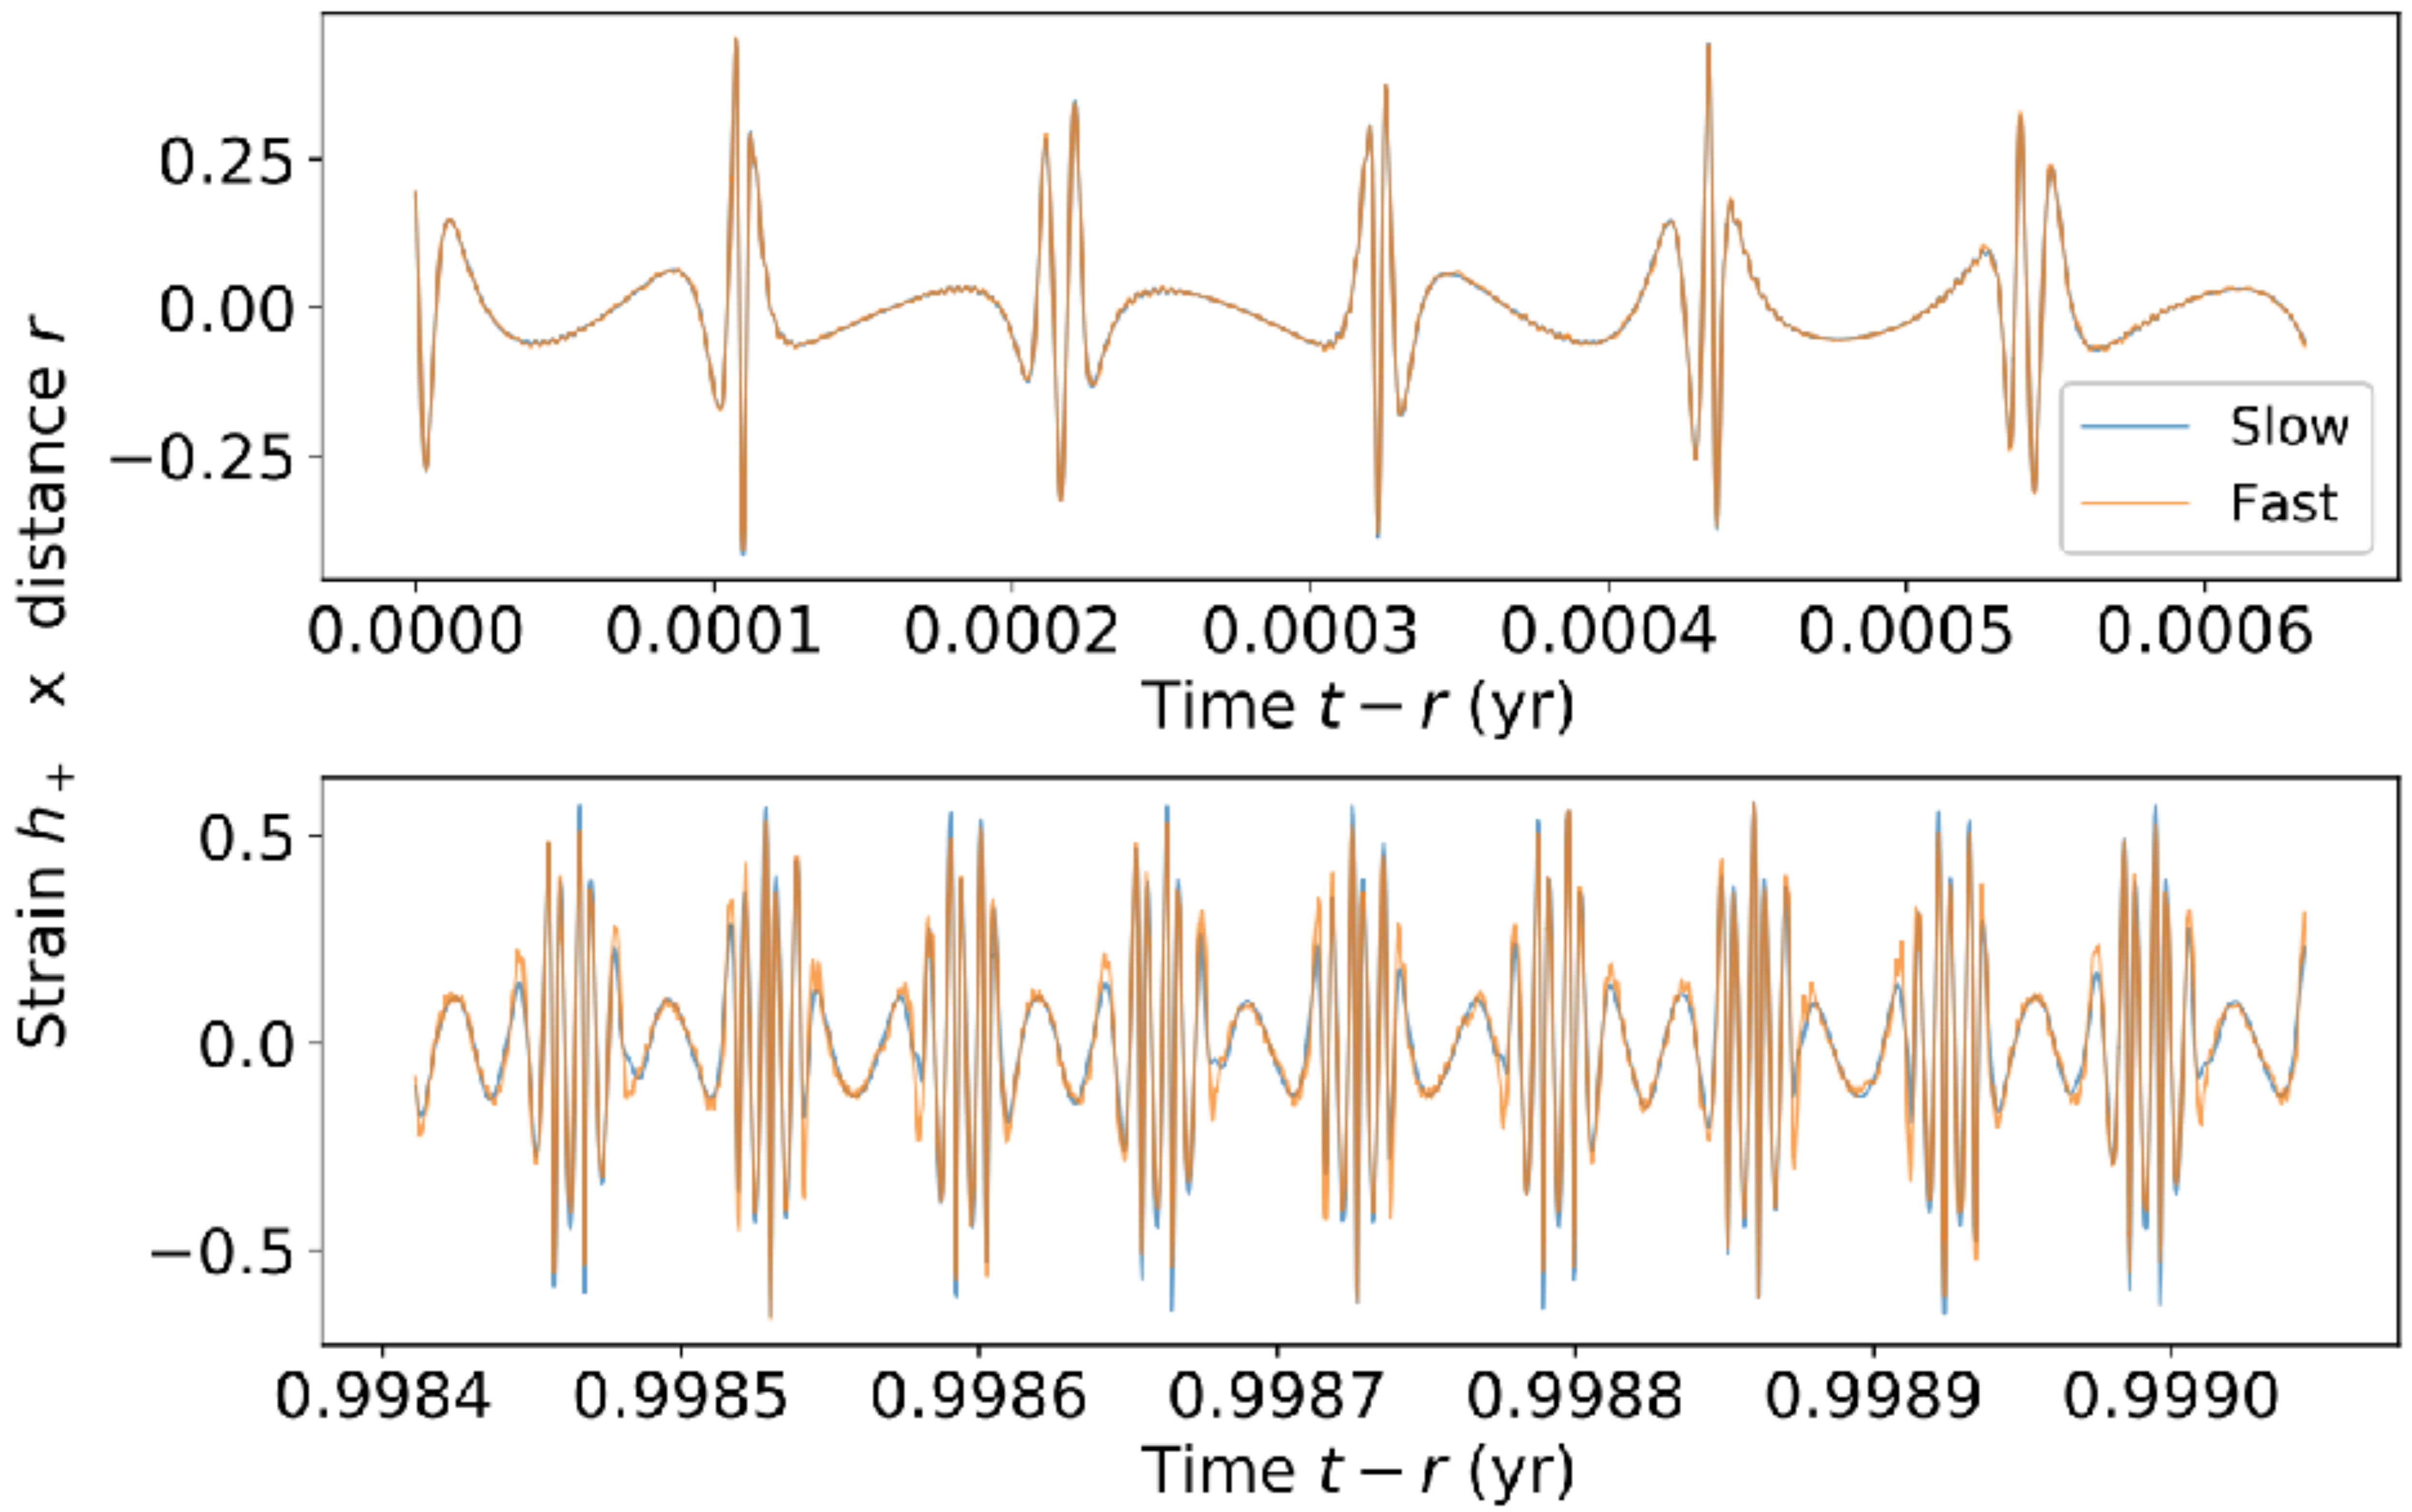

Just before ‘plunge’

Chua et al. arXiv:2008.06071

Potential for extremely precise parameter measurements

Liang et al. arXiv:2409.07957

- Some parameters can theoretically be measured to a precision of e.g. 0.01% (eccentricity), 0.003% (primary mass), 0.005% (secondary mass)
- Example of MCMC run initiated very close to the true injected values

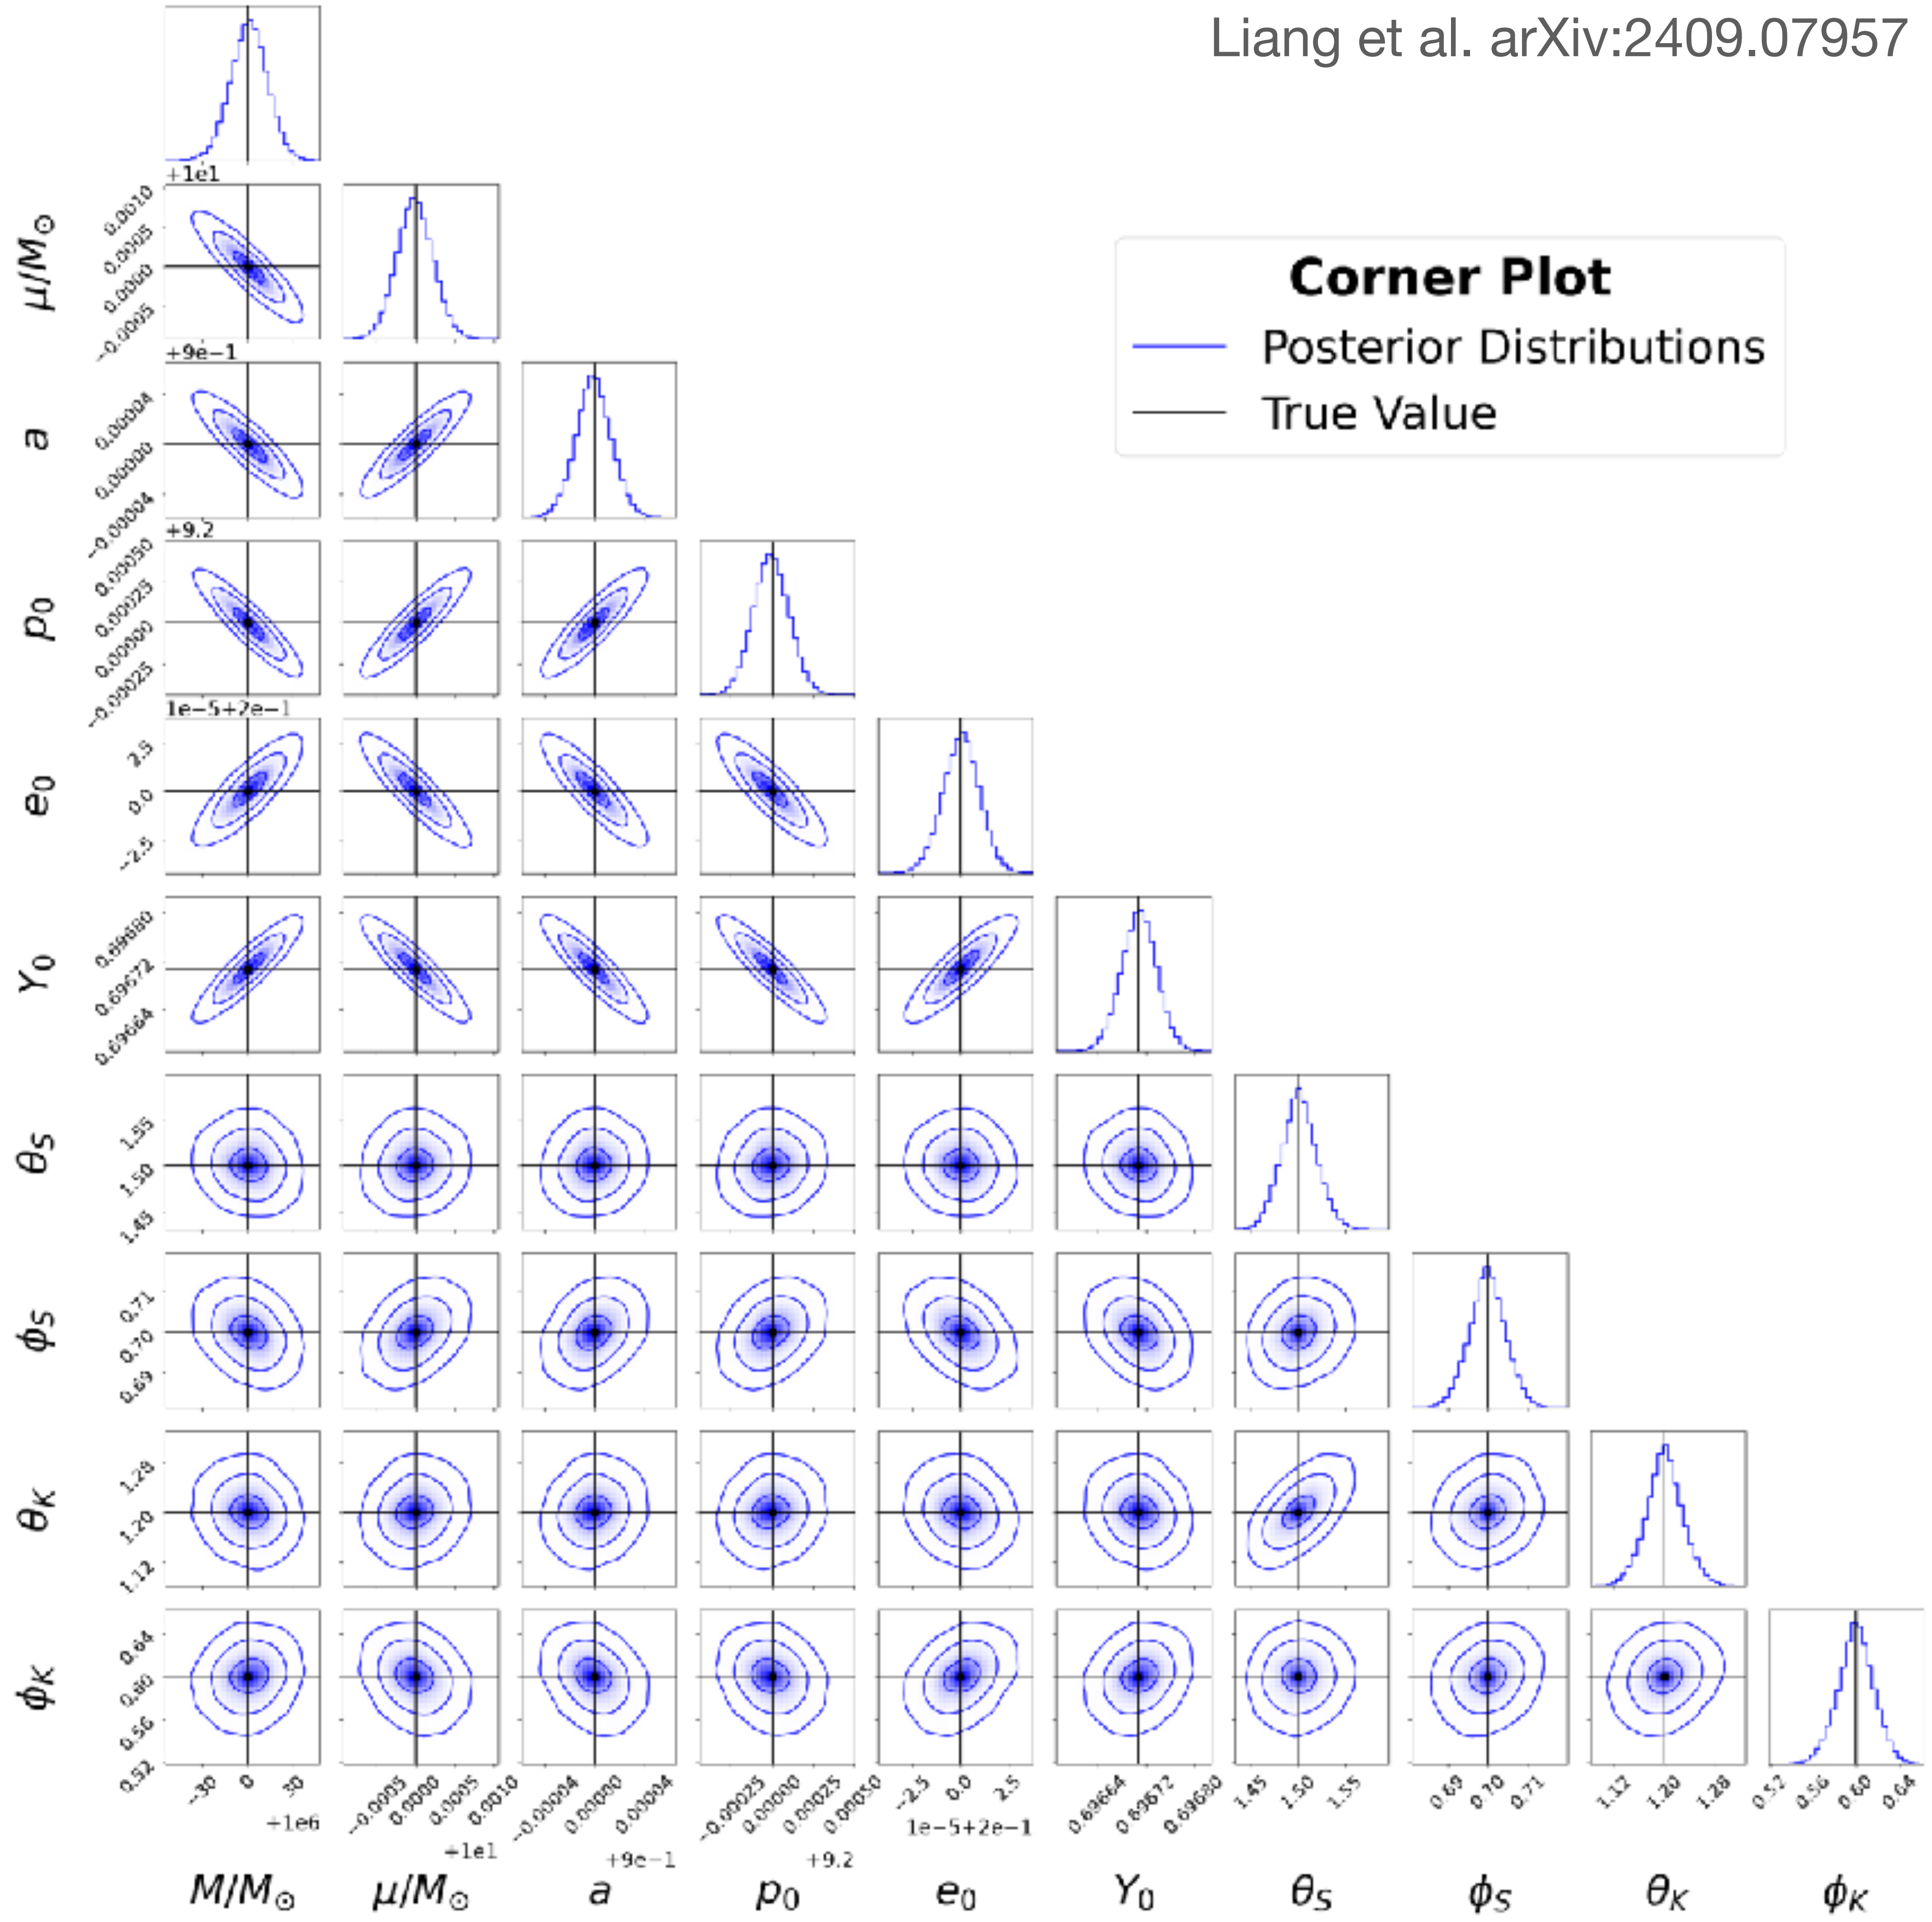

Potential for extremely precise parameter measurements

- Some parameters can theoretically be measured to a precision of e.g. 0.01% (eccentricity), 0.003% (primary mass), 0.005% (secondary mass)
- Example of MCMC run initiated very close to the true injected values
- Precision required for new physics inferences

Liang et al. arXiv:2409.07957

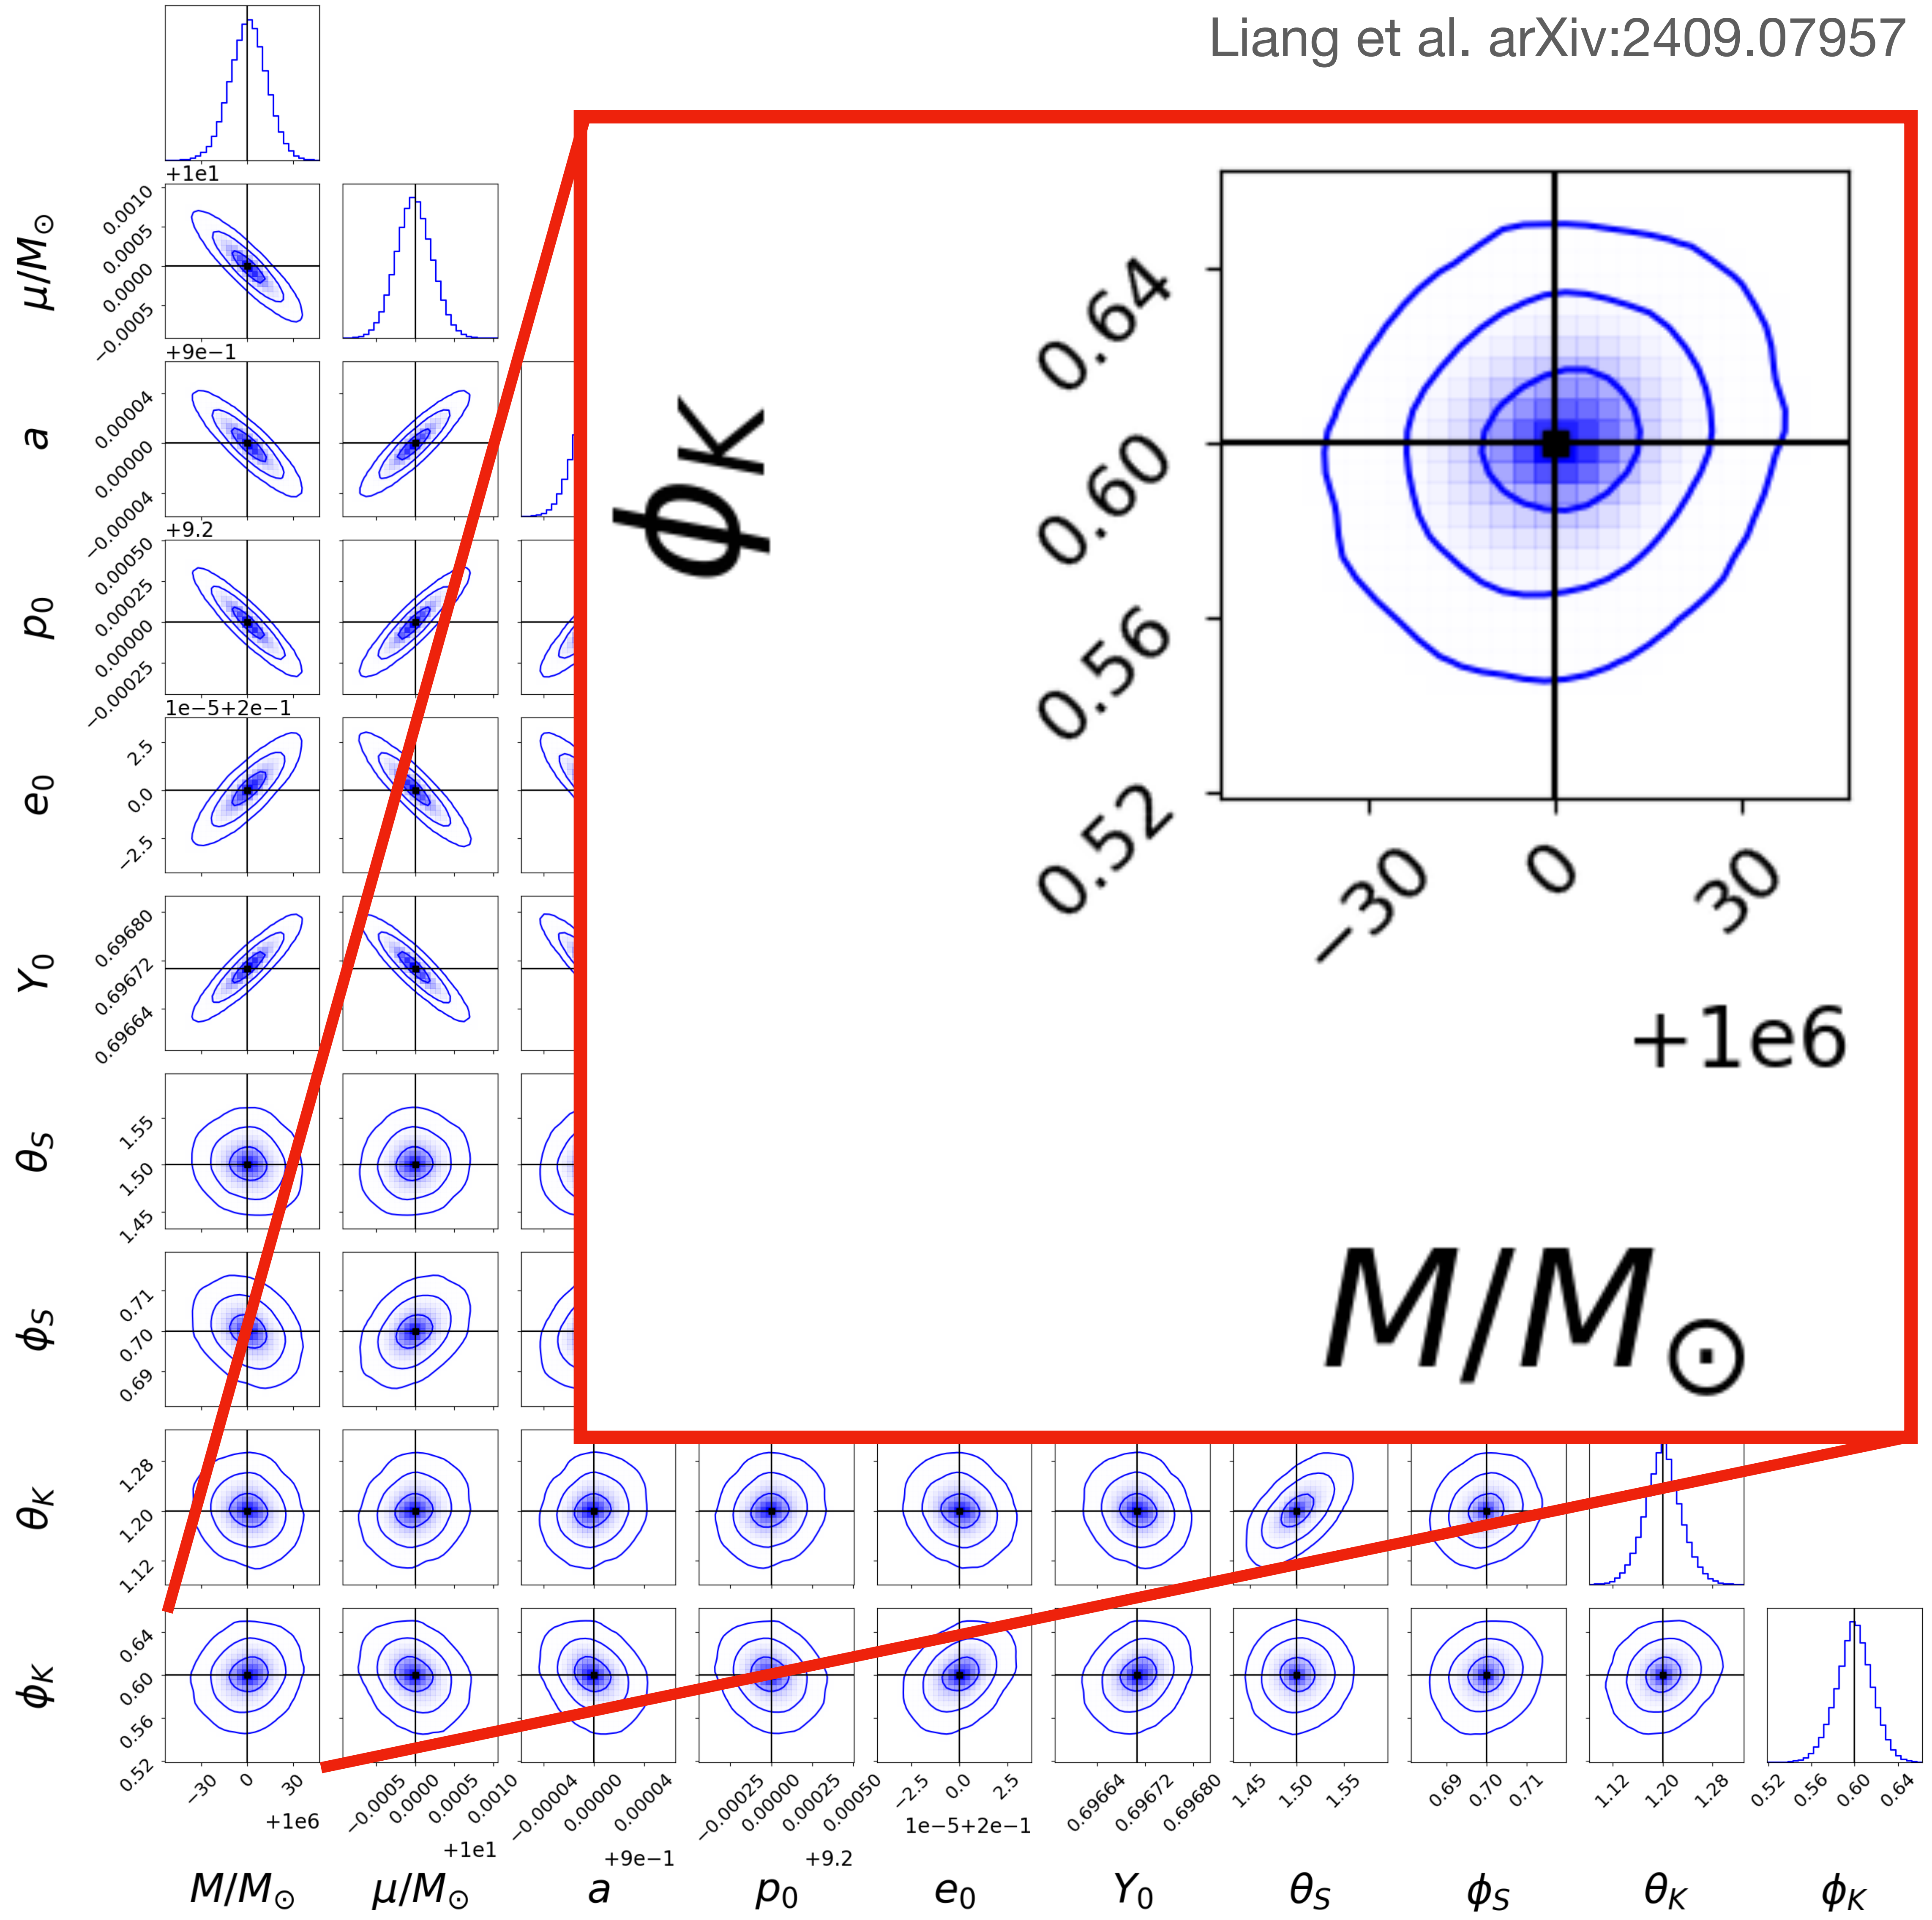

Vast and multi-modal parameter space

- However the EMRI parameter space is vast, making search and parameter estimation strategies extremely difficult
- There are also many degeneracies between parameters, so the parameter landscape is highly multi-modal
- Prior width on just the mass is a factor of 10^5 larger than posterior

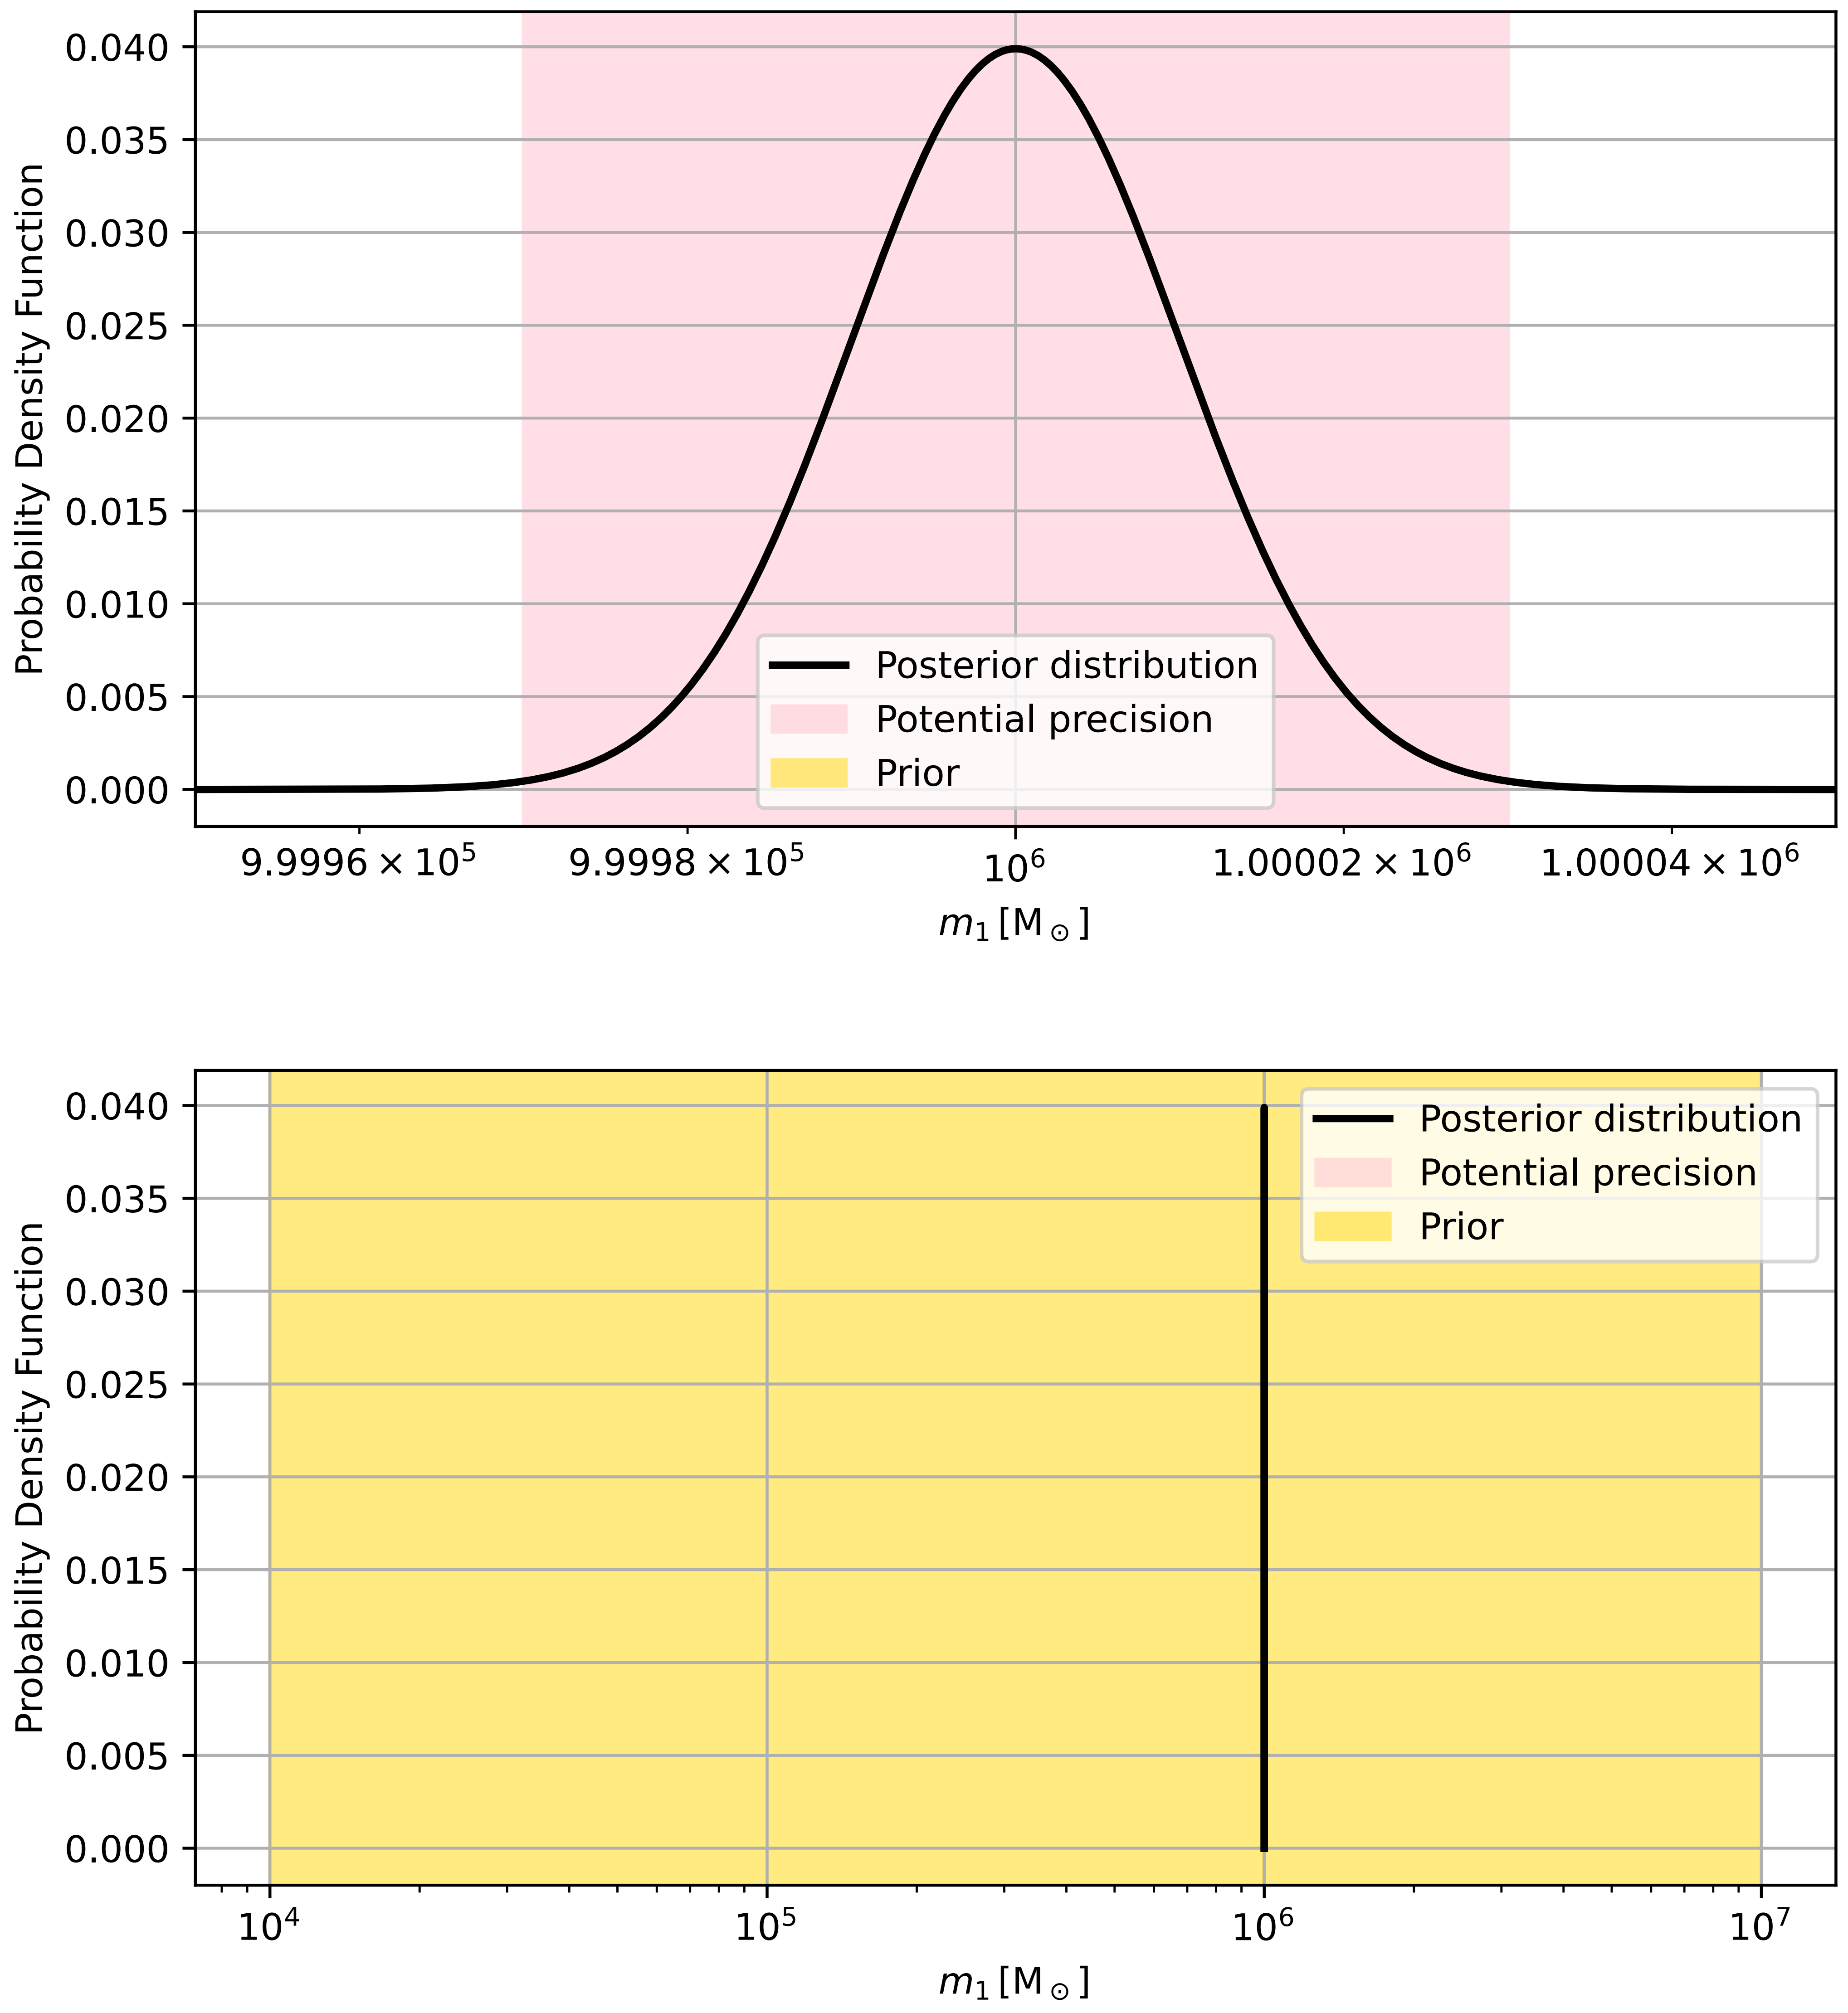

See e.g. Chua & Cutler arXiv:2109.14254

EMRIs in the milliHertz frequency band will be buried in non-stationary, non-Gaussian noise

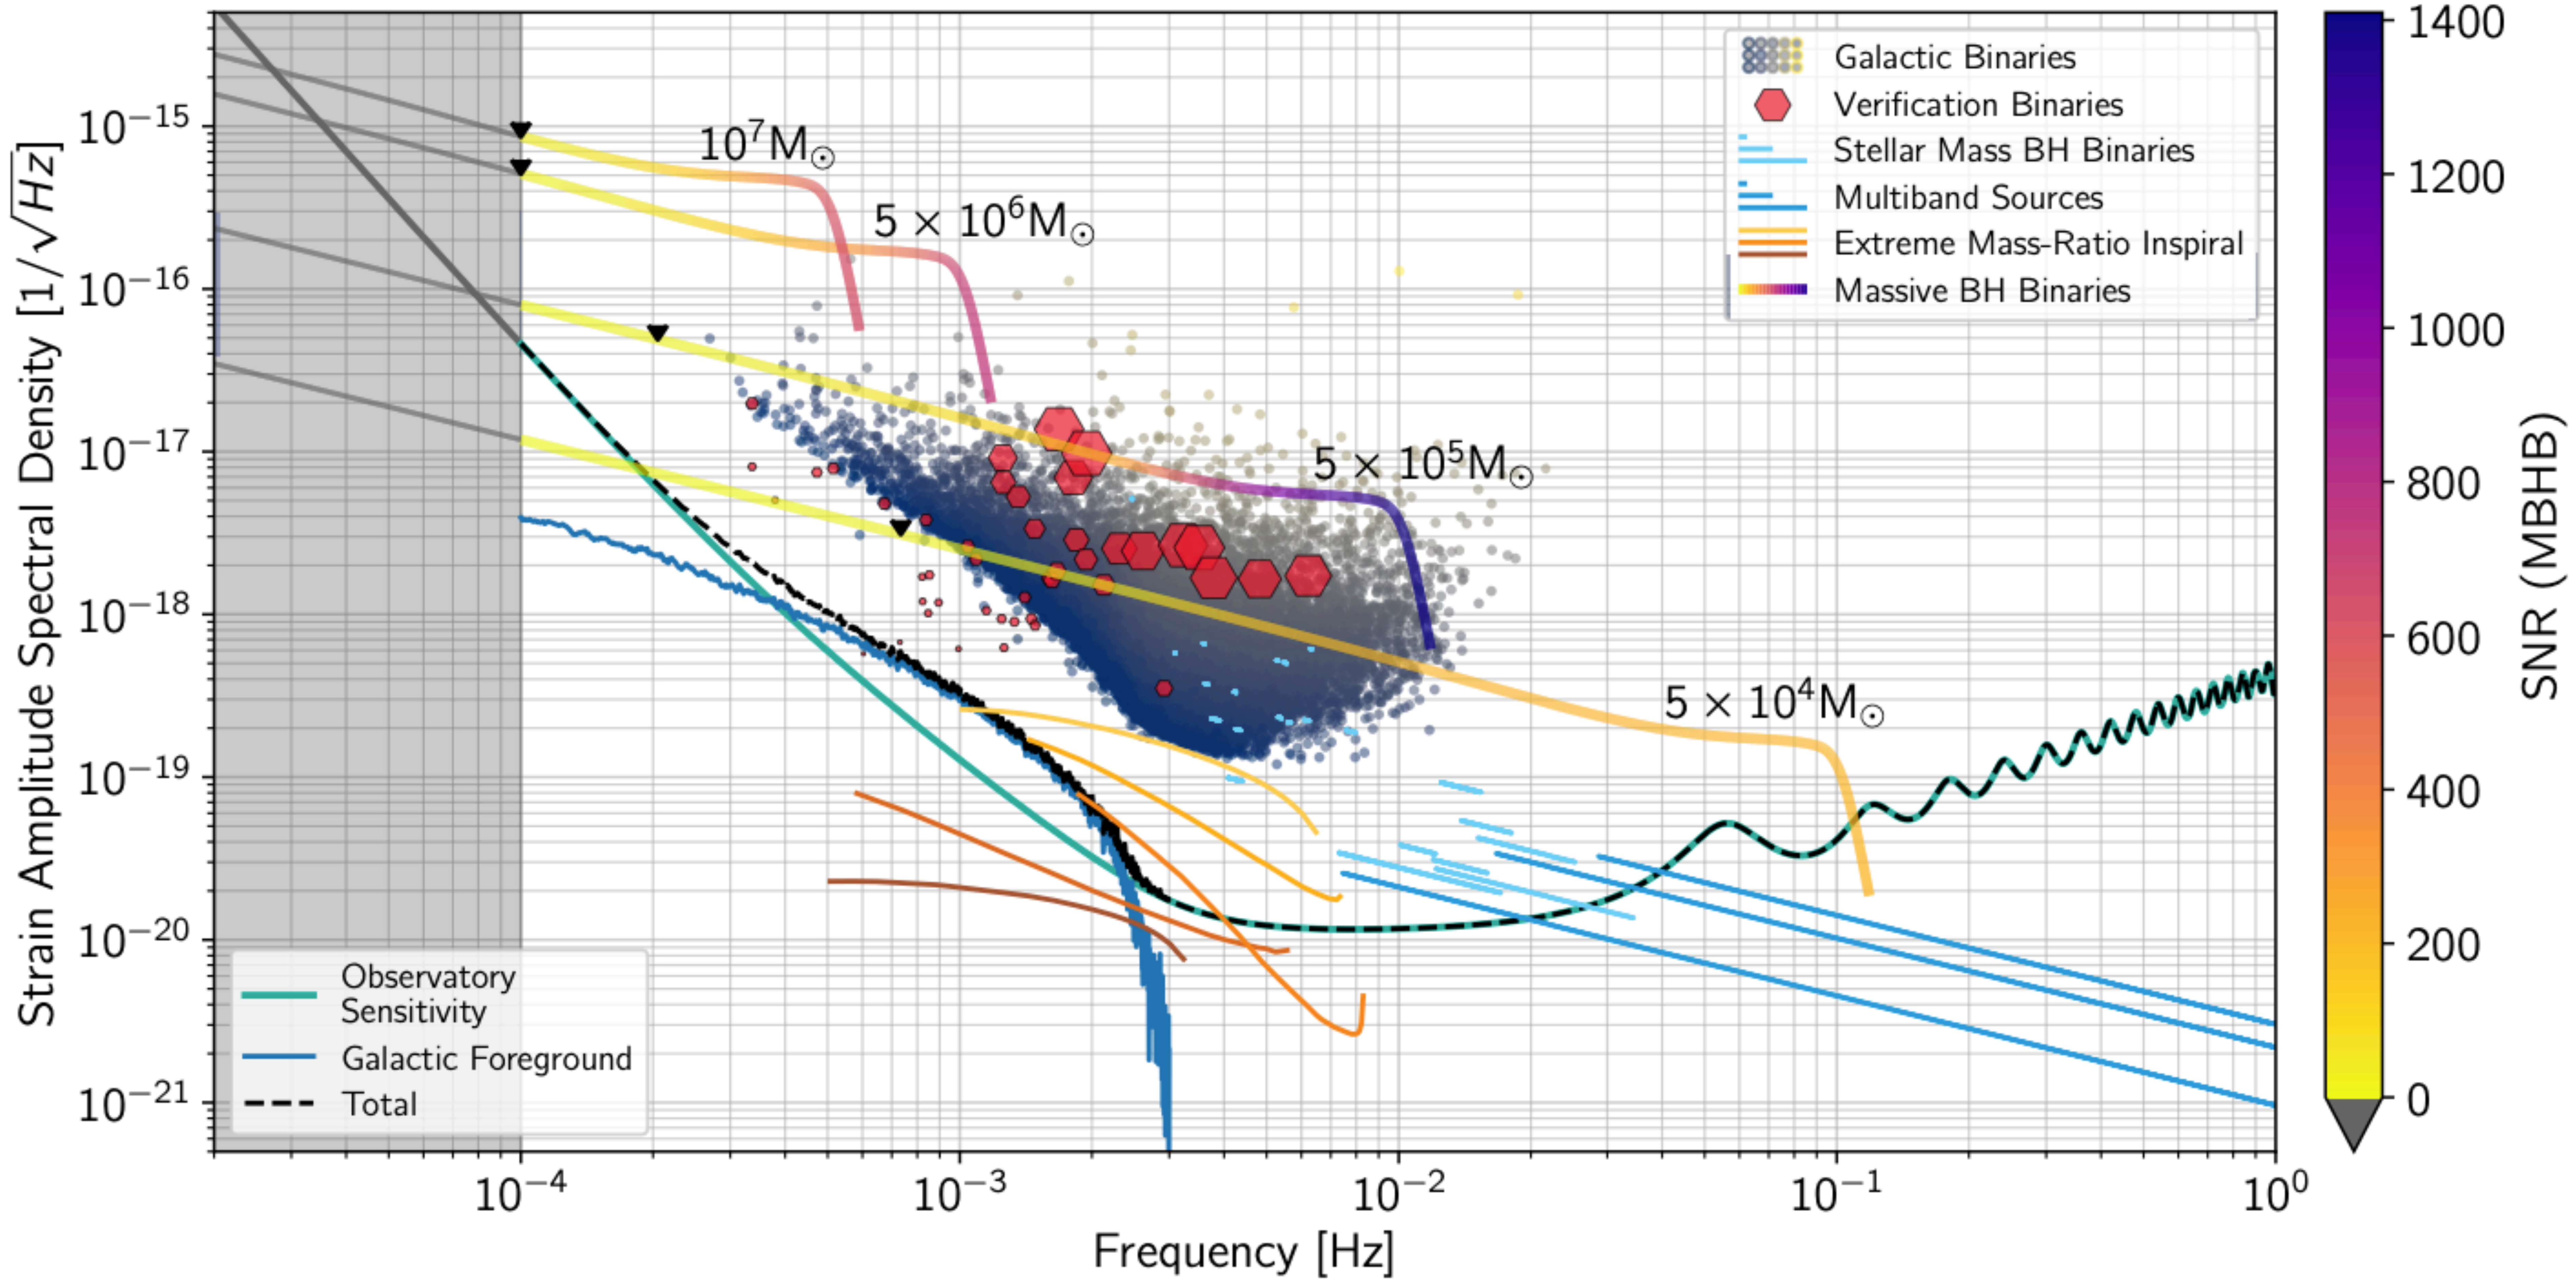

LISA Red Book: arXiv:2402.07571

How and why sequential simulation-based inference might help

1. Simulation-efficient way to narrow down a vast and multi-modal parameter space (important because EMRI simulations are also relatively expensive)

2. Future goal will be to cope with non-stationary, non-Gaussian noise (as well as many overlapping sources), where likelihood-based methods tend to become expensive or unfeasible

Sequential simulation-based inference at a glance

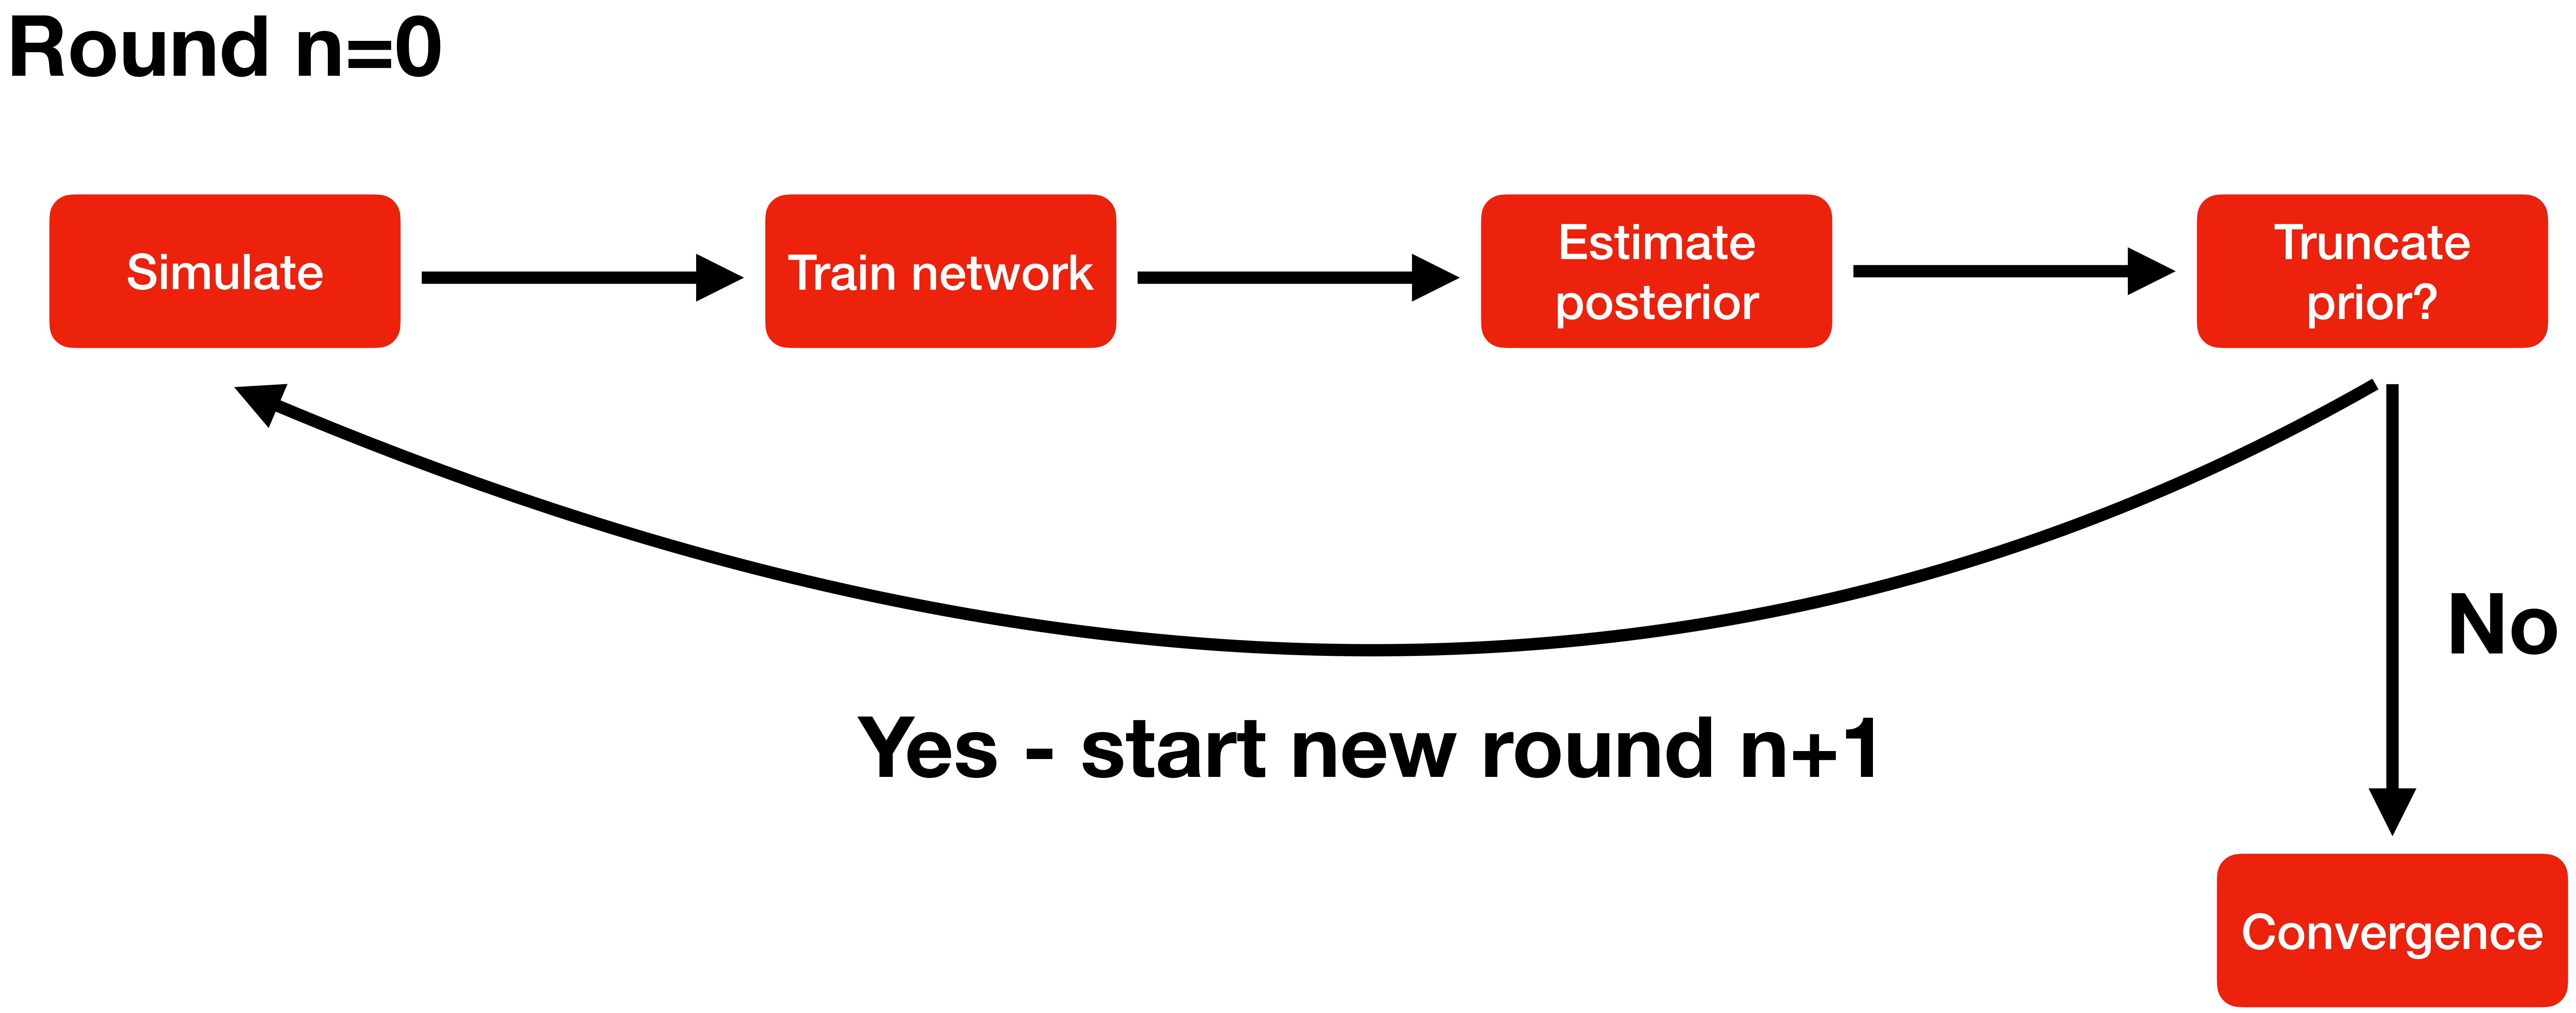

Bhardwaj et al. 2023

Simulation set-up

- Signal simulated with Fast EMRI Waveforms (`FEW`) code (version 1)
- Fed through `fastlisaresponse` which implements the LISA detector response in time domain and outputs the signal projected onto the ‘A’, ‘E’ and ’T’ time delay interferometry channels
- Add noise sampled from the power spectral density as computed by `pycbc` (analytical including confusion noise)

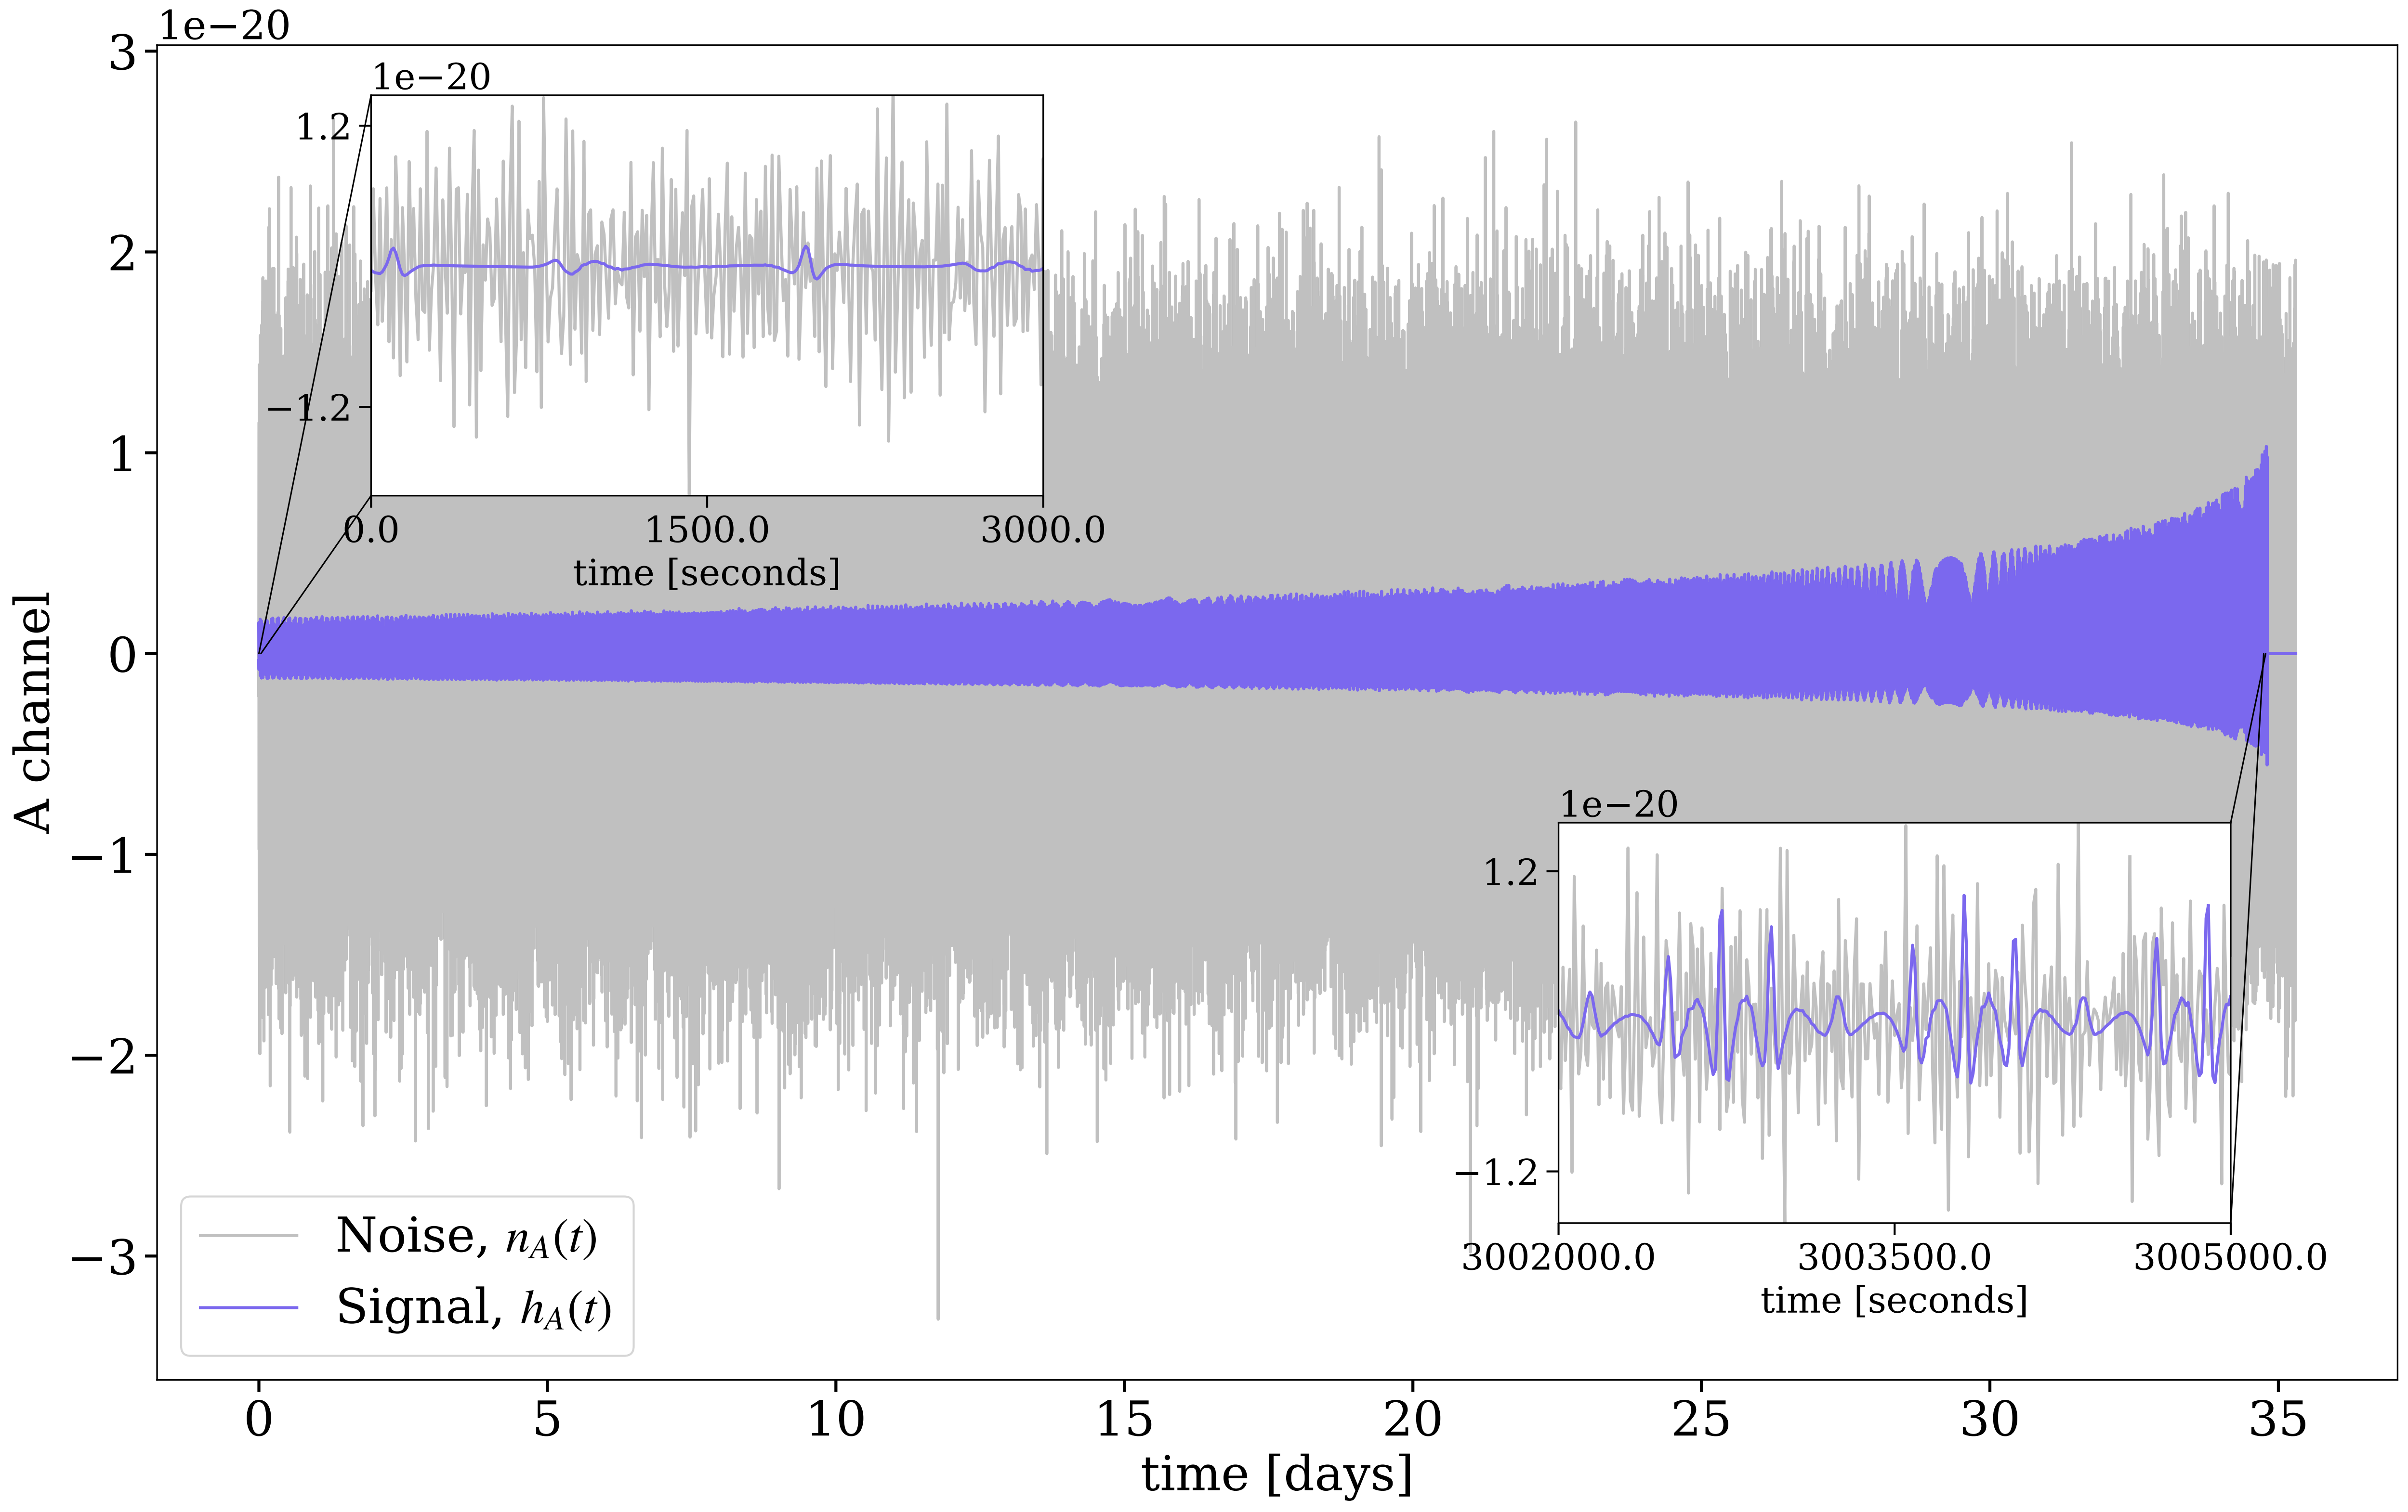

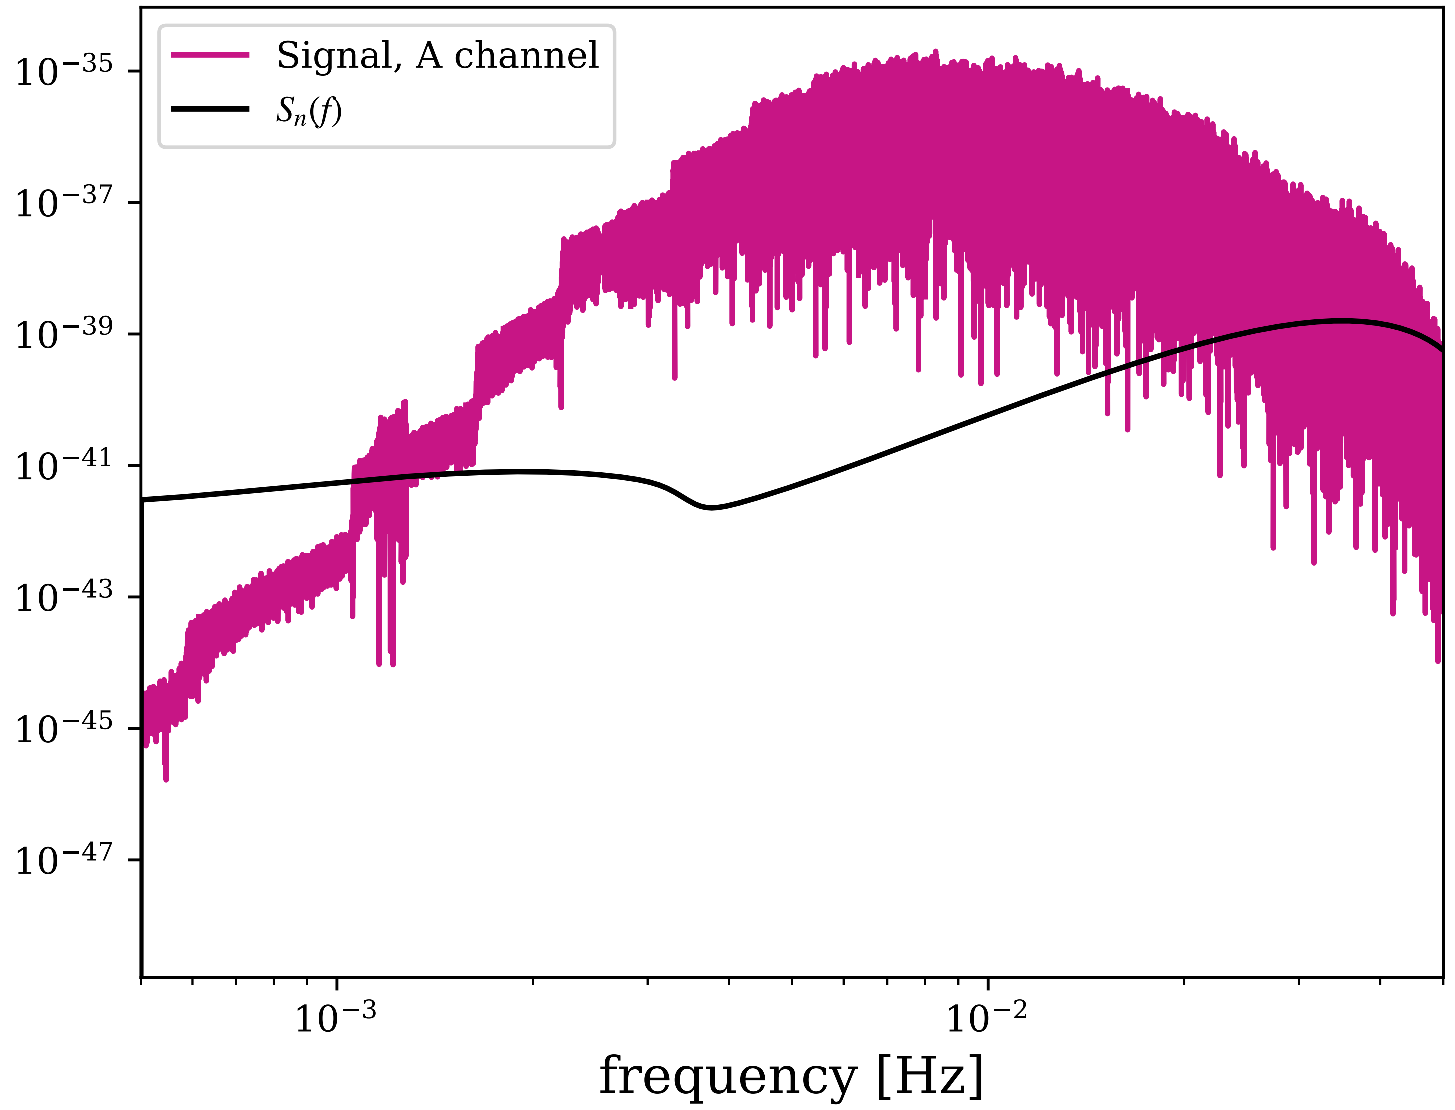

$dt = 10\,\mathrm{s}, T_{obs} = 0.1\,\mathrm{yr}, m_1 = 5.385 \times 10^5\,\mathrm{M}_\odot, m_2 = 50.55\,\mathrm{M}_\odot, p_0 = 10.35\; e_0 = 0.2927,\; \cos\theta_K = 0.0384, \phi_K = 5.212, d_L = 232.8\,\mathrm{Mpc}, \Phi_\phi = 2.966, \Phi_r = 2.014,\; \cos\theta_S = 0.4107, \phi_S = 3.124$

PEREGRINE

Gravitational Wave Parameter Inference with Neural Ratio Estimation

Training set-up

PEREGRINE-style approach - Truncated Marginal Neural Ratio Estimation

- Train a neural network to recognise whether parameter vector θ_k and signal x are drawn jointly $p(x, \theta_k)$ or marginally $p(x)p(\theta_k)$ - binary classification

- Minimise the binary-cross entropy loss function:

$$\mathscr{L}[\hat{\rho}_{k,\phi}] = -\int \left\{ p(x,\theta_k) \ln \sigma(\hat{\rho}_{k,\phi}(x,\theta_k)) + p(x)p(\theta_k) \ln \left[1 - \sigma(\hat{\rho}_{k,\phi}(x,\theta_k)) \right] \right\} dx\, d\theta_k \,.$$

- Optimal classifier $\hat{\rho}_{k,\phi}(x, \theta_k)$ is the log of the likelihood-to-evidence ratio - re-weight by prior samples to estimate the posterior

Likelihood / Jointly drawn samples / Posterior

$$r_k(\boldsymbol{x} \mid \boldsymbol{\vartheta}_k) := \frac{p(\boldsymbol{x} \mid \boldsymbol{\vartheta}_k)}{p(\boldsymbol{x})} = \frac{p(\boldsymbol{x}, \boldsymbol{\vartheta}_k)}{p(\boldsymbol{x})p(\boldsymbol{\vartheta}_k)} = \frac{p(\boldsymbol{\vartheta}_k \mid \boldsymbol{x})}{p(\boldsymbol{\vartheta}_k)}$$

Evidence / Marginally drawn samples / Prior

Bhardwaj et al. 2023
Miller et al. 2021

Effectively narrows down parameter space

- Parameter space for intrinsic parameters narrowed significantly via truncation between rounds
- Proposal distribution volume a million times smaller than prior volume
- Differential evolution (stochastic optimiser) performs less well (however see Strub et al. 2025)

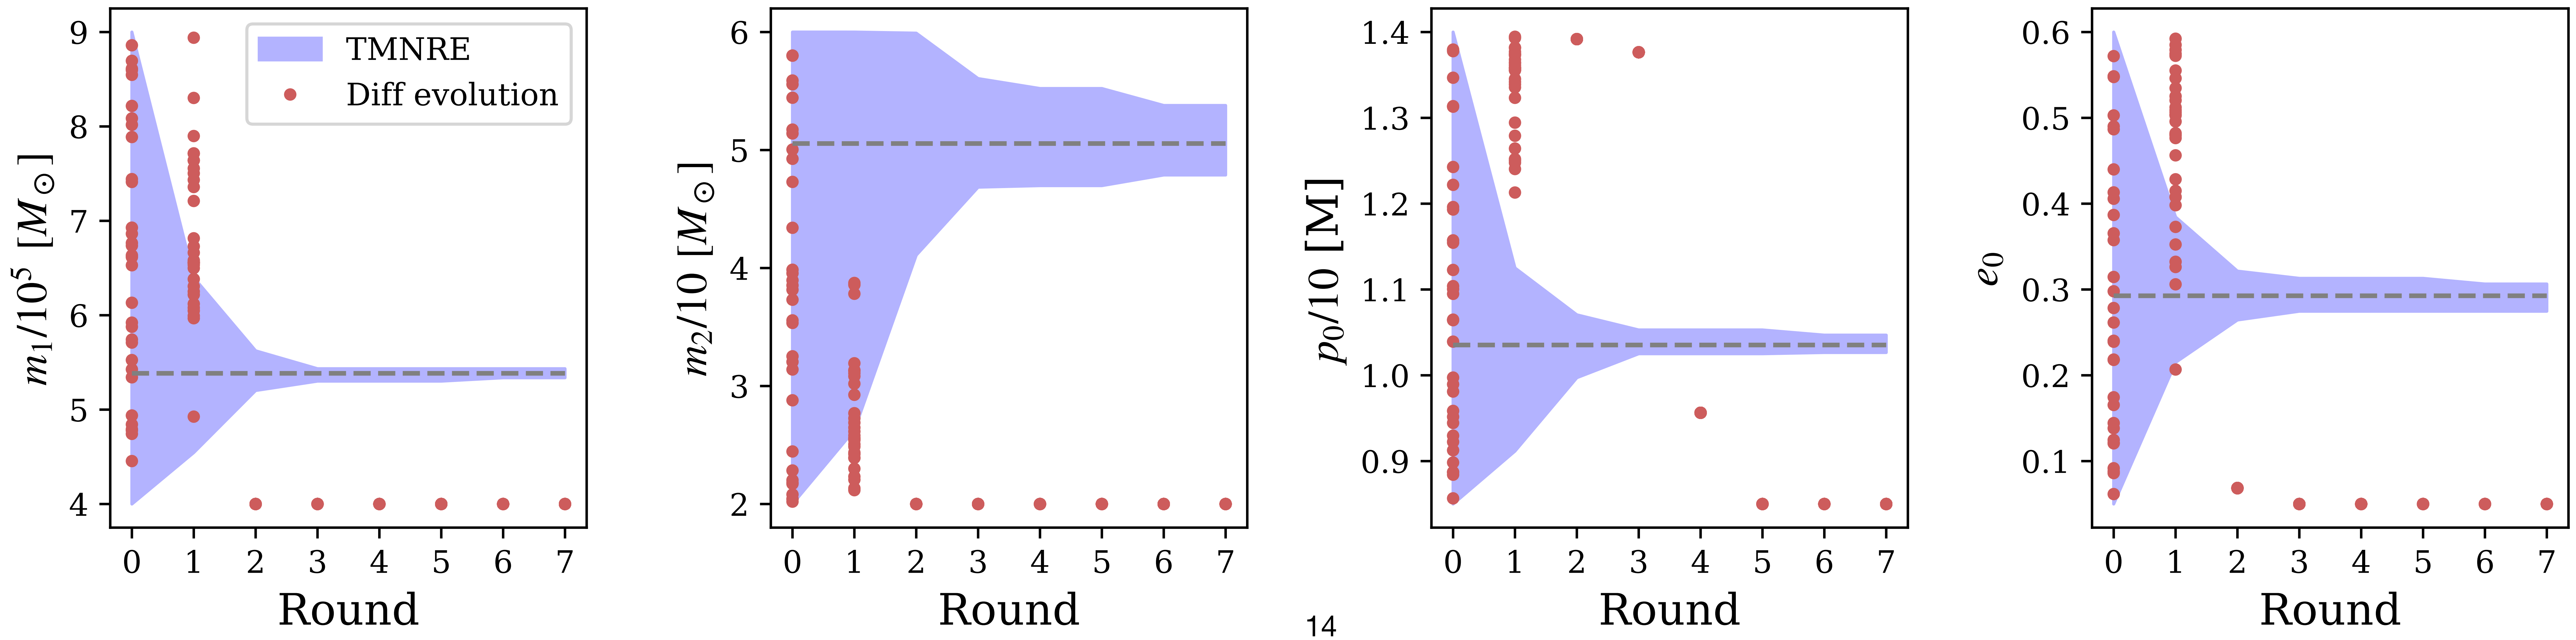

Posteriors/proposal distributions

- 6 sequential rounds, 150K simulations each round, approx 12 hours each round
- Injected values all within 2σ credible intervals
- The relative half-widths (2σ credible intervals) are 0.3% for m_1, 2% for m_2, 0.3% for p_0 and 2% for e_0.

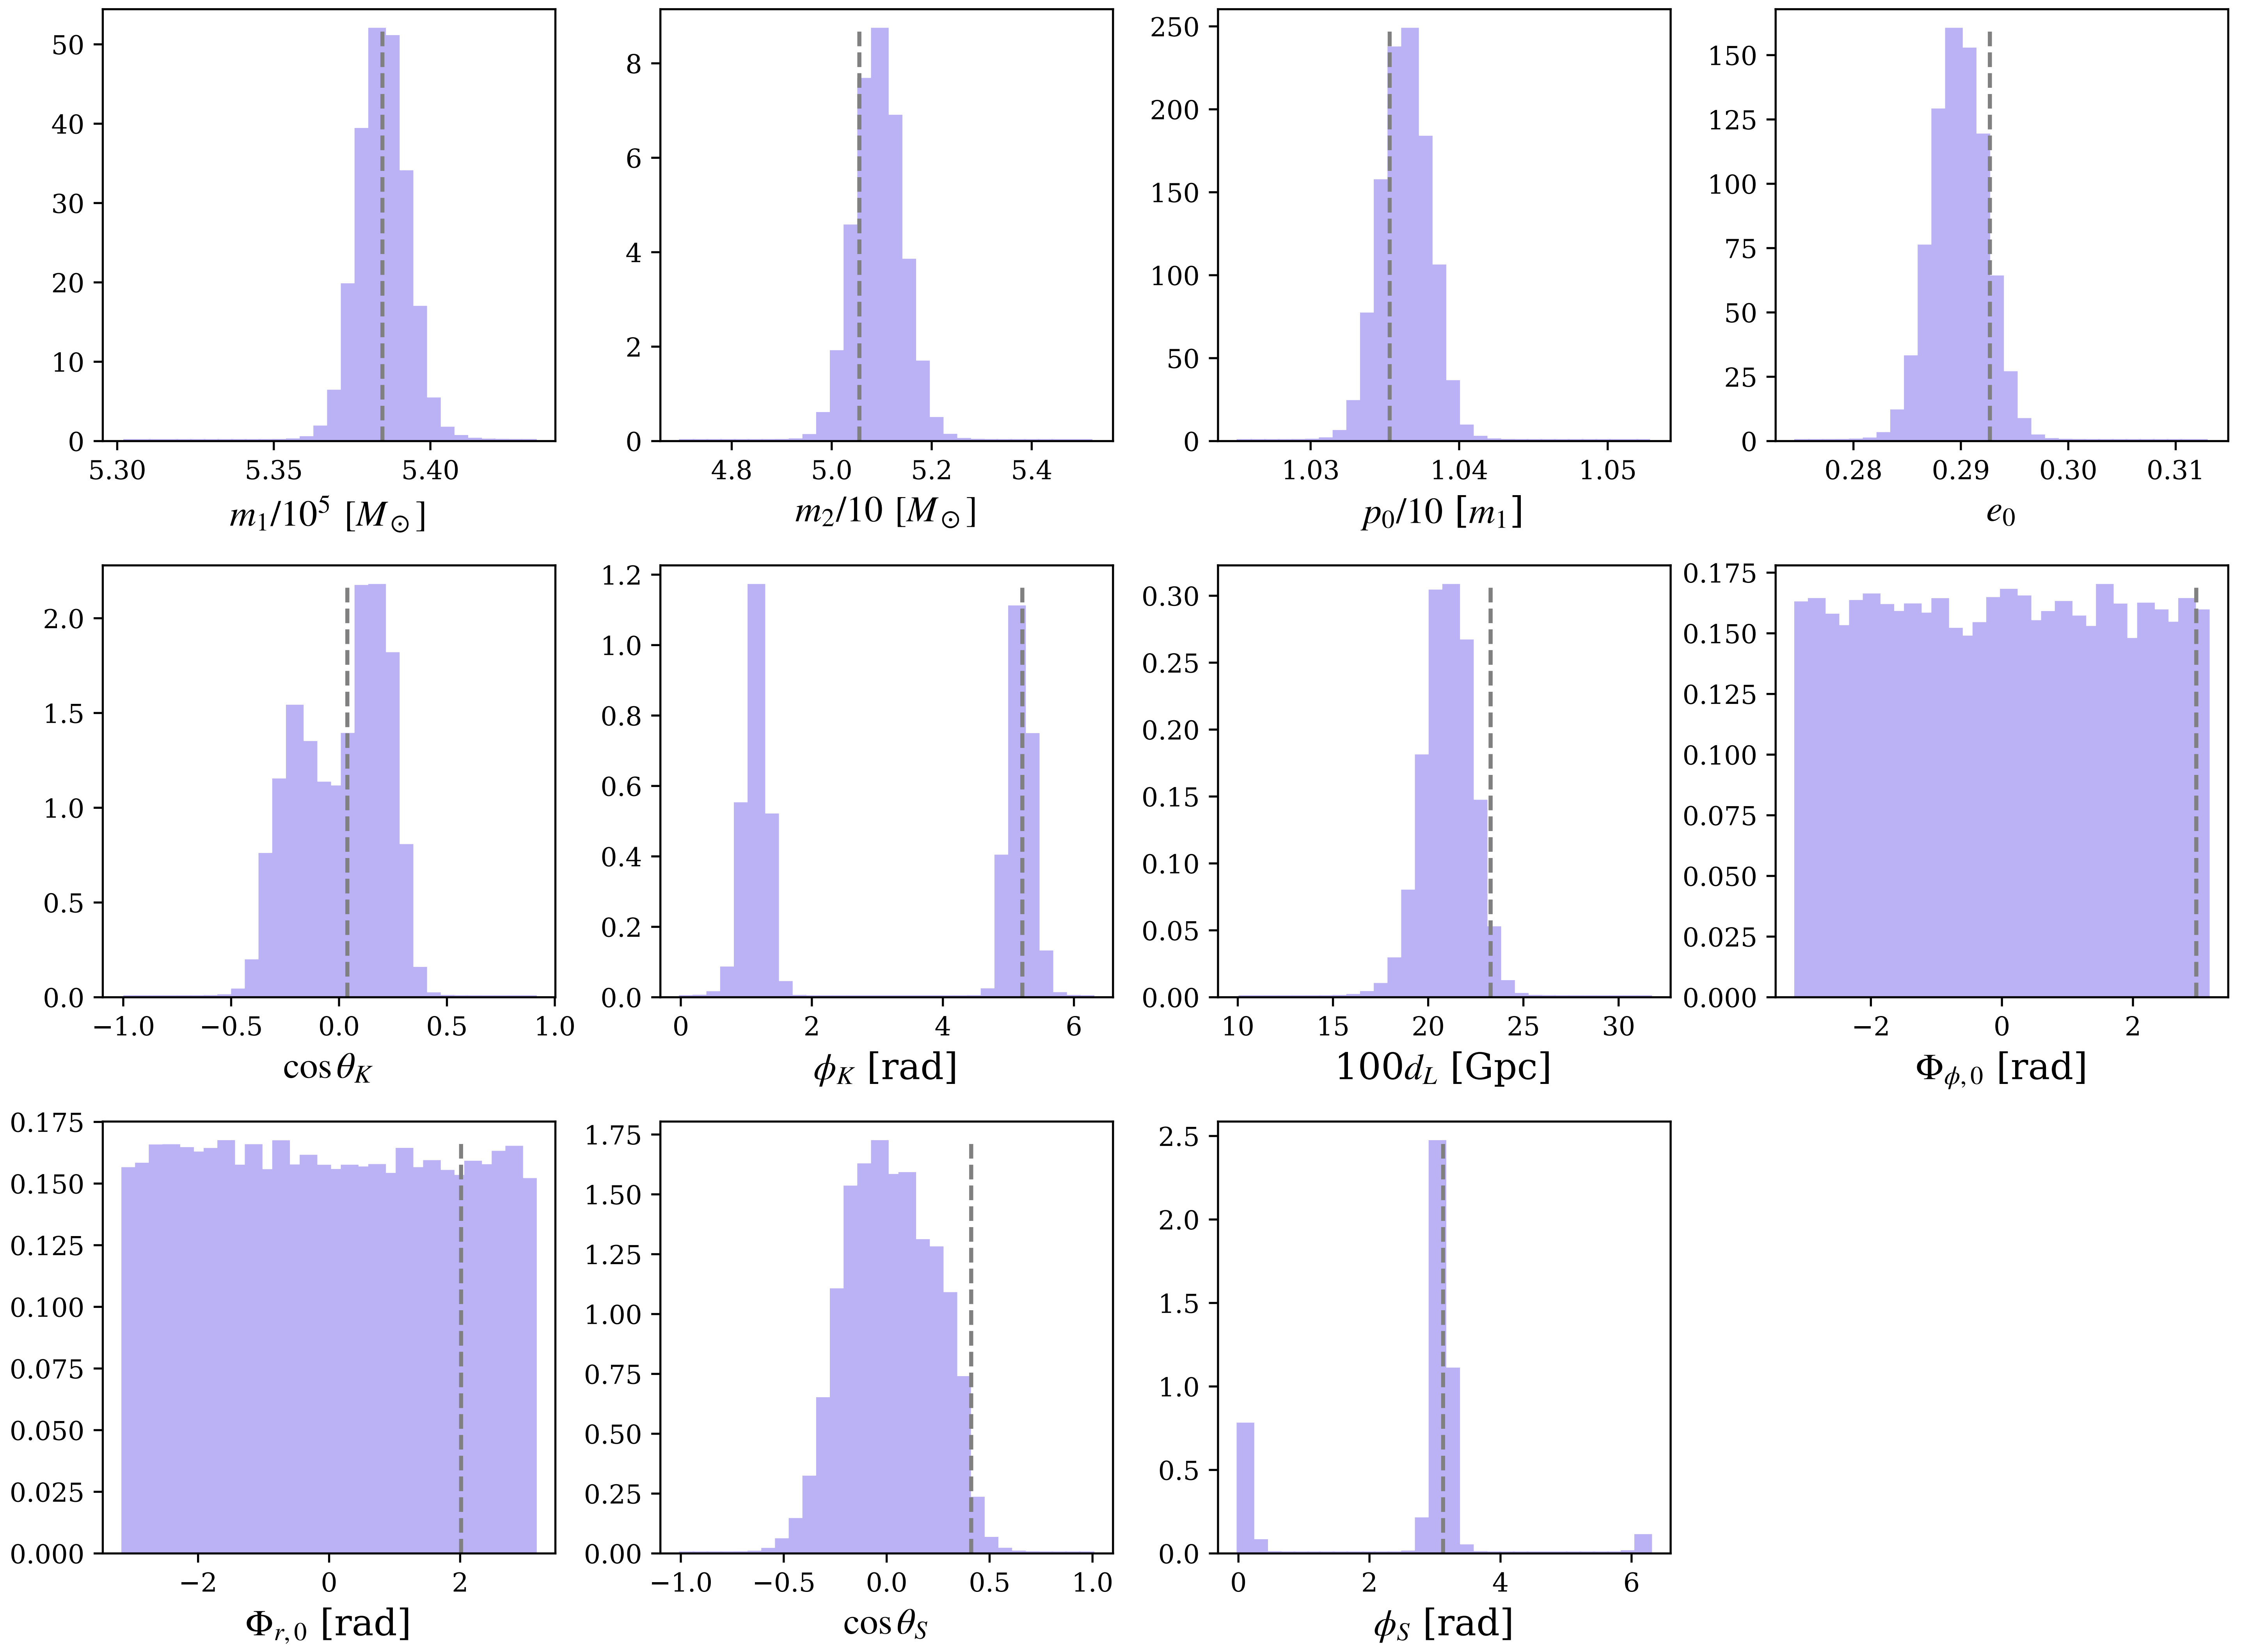

Posteriors/proposal distributions

- 6 sequential rounds, 150K simulations each round, approx 12 hours each round
- Injected values all within 2σ credible intervals
- The relative half-widths (2σ credible intervals) are 0.3% for m_1, 2% for m_2, 0.3% for p_0 and 2% for e_0.

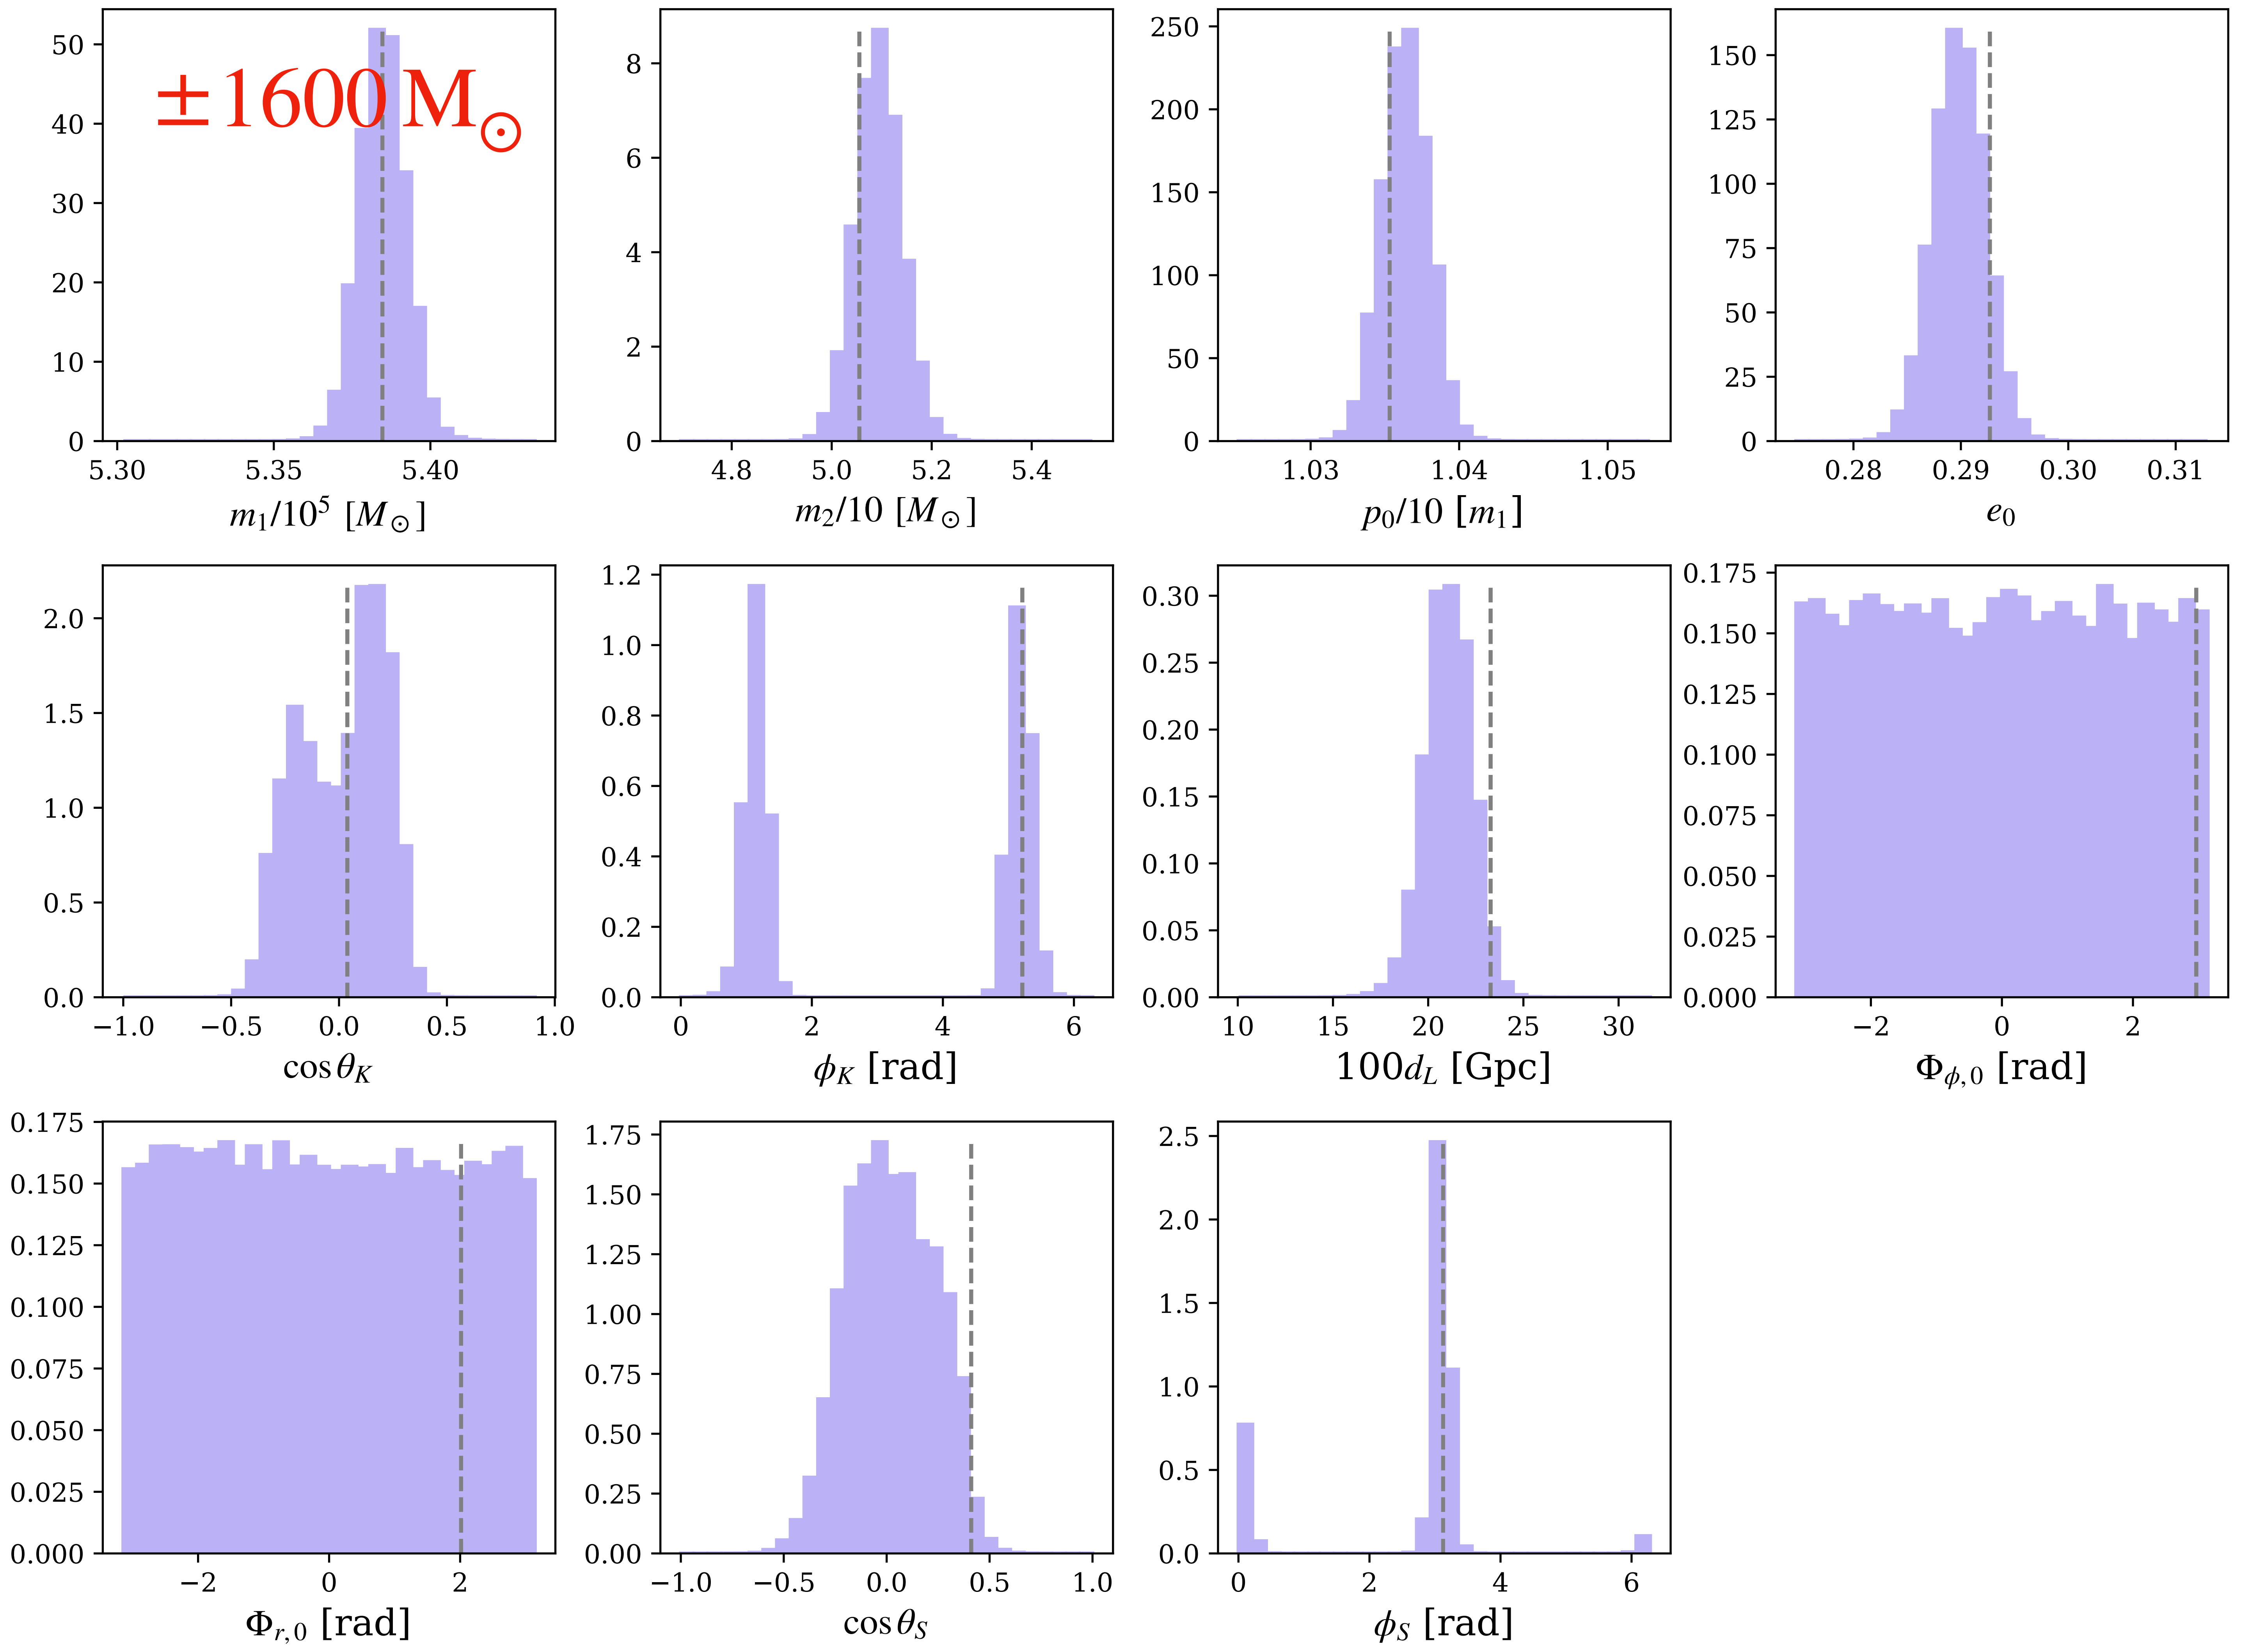

Comparison with MCMC

- Compare efficiency by choosing like-for-like starting states
- MCMC prior chosen to be the 6th round prior identified with TMNRE
- 32 walkers and 4687 steps ~ waveform evaluations approximately 150K
- 2 days wall-clock time
- Chains not converged

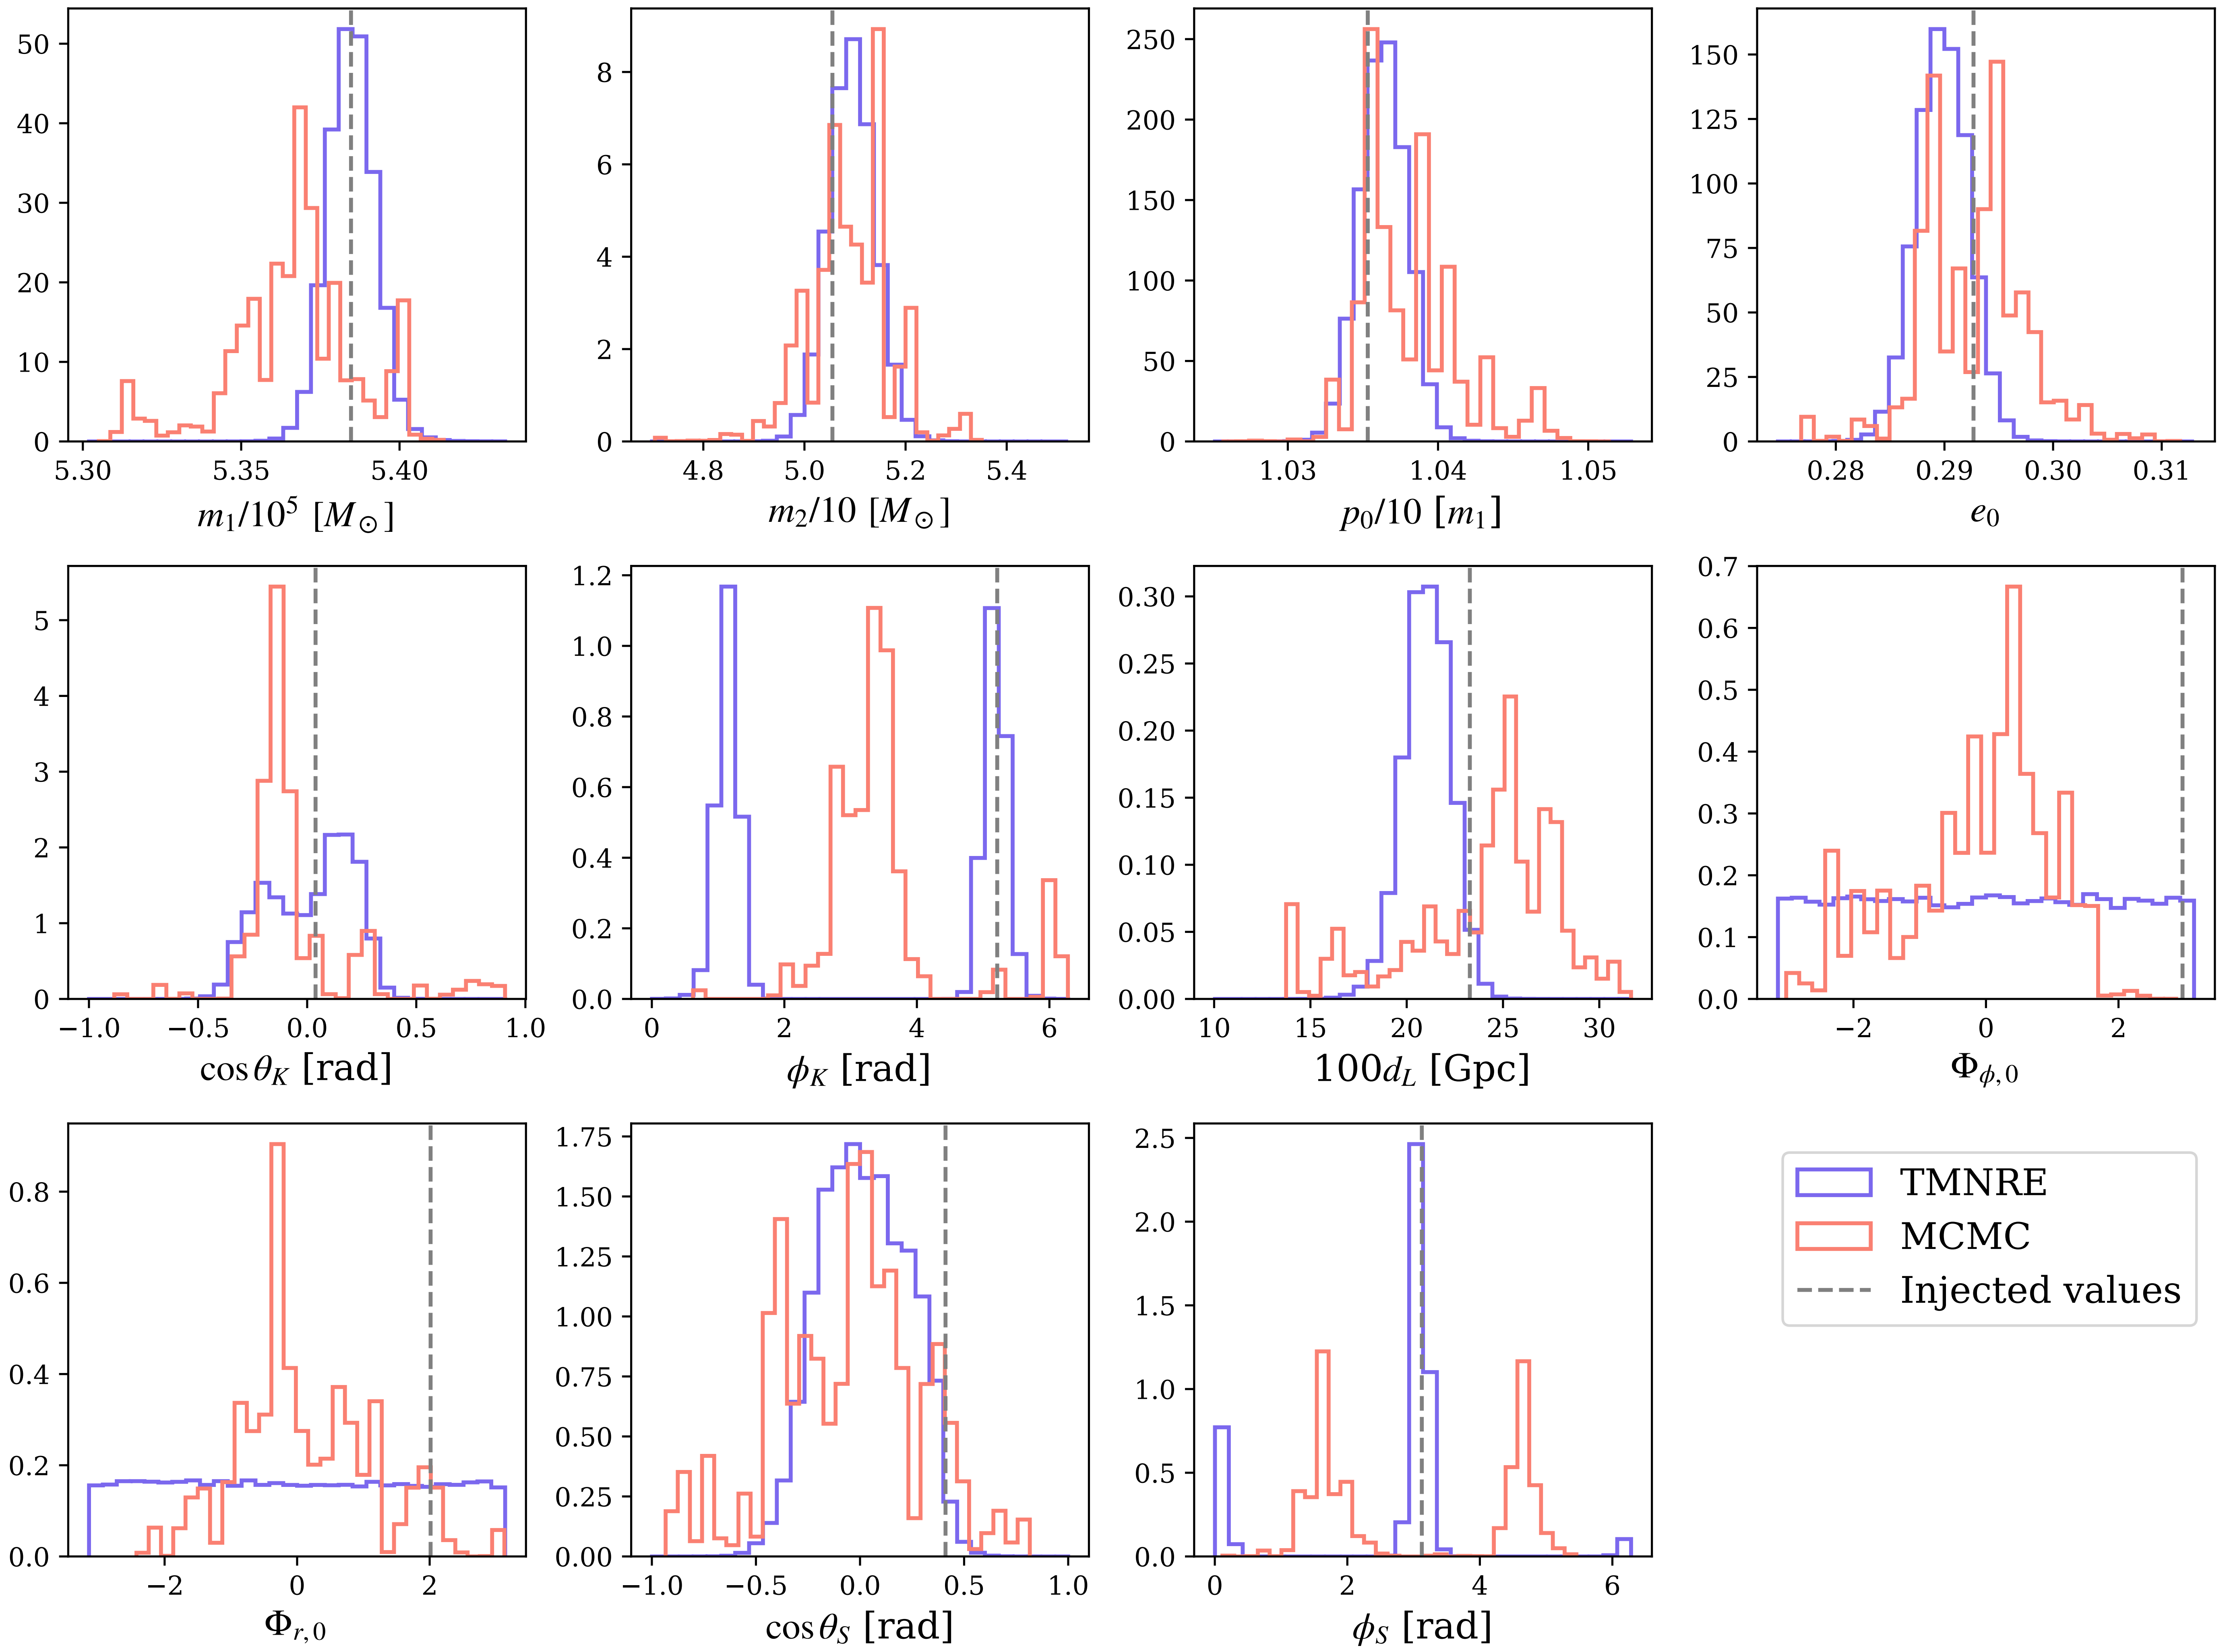

Conclusions

- TMNRE narrows down parameter space volume (by a factor of a million) very efficiently from wide priors
- Estimates proposal distributions that are better converged than MCMC when compared with same number of waveform evaluations and starting states
- Improvements required in order to achieve precise parameter measurements - use output as proposal distributions, compress data offline, non-uniform truncation methods
- Eventually tackle non-stationary, non-Gaussian noise and overlapping sources
- How and when do we fold back in environmental effects to the data analysis pipeline?

Many astrophysics and fundamental physics opportunities…

- Formation of intermediate mass black holes (Volonteri 2010)
- Environmental effects - dark matter, ultra-light bosons, accretion disks around primary object (Eda et al. 2013, Macedo et al. 2013, Barausse et al. 2014, Kavanagh et al. 2020, Khalvati et al. 2024, Baumann et al. 2022, Barsanti et al. 2023, Zhang et al. 2023, Tomaselli et al. 2024, Cole et al. 2023, Speri et al. 2023, Duque et al. 2024, Coppaoni et al. 2025…)
- Tests of General Relativity (Han & Chen 2019, Gupta et al. 2022, Speri et al. 2024, Kejriwal et al. 2024, A. Cardenas-Avendano & Sopuerta 2024)

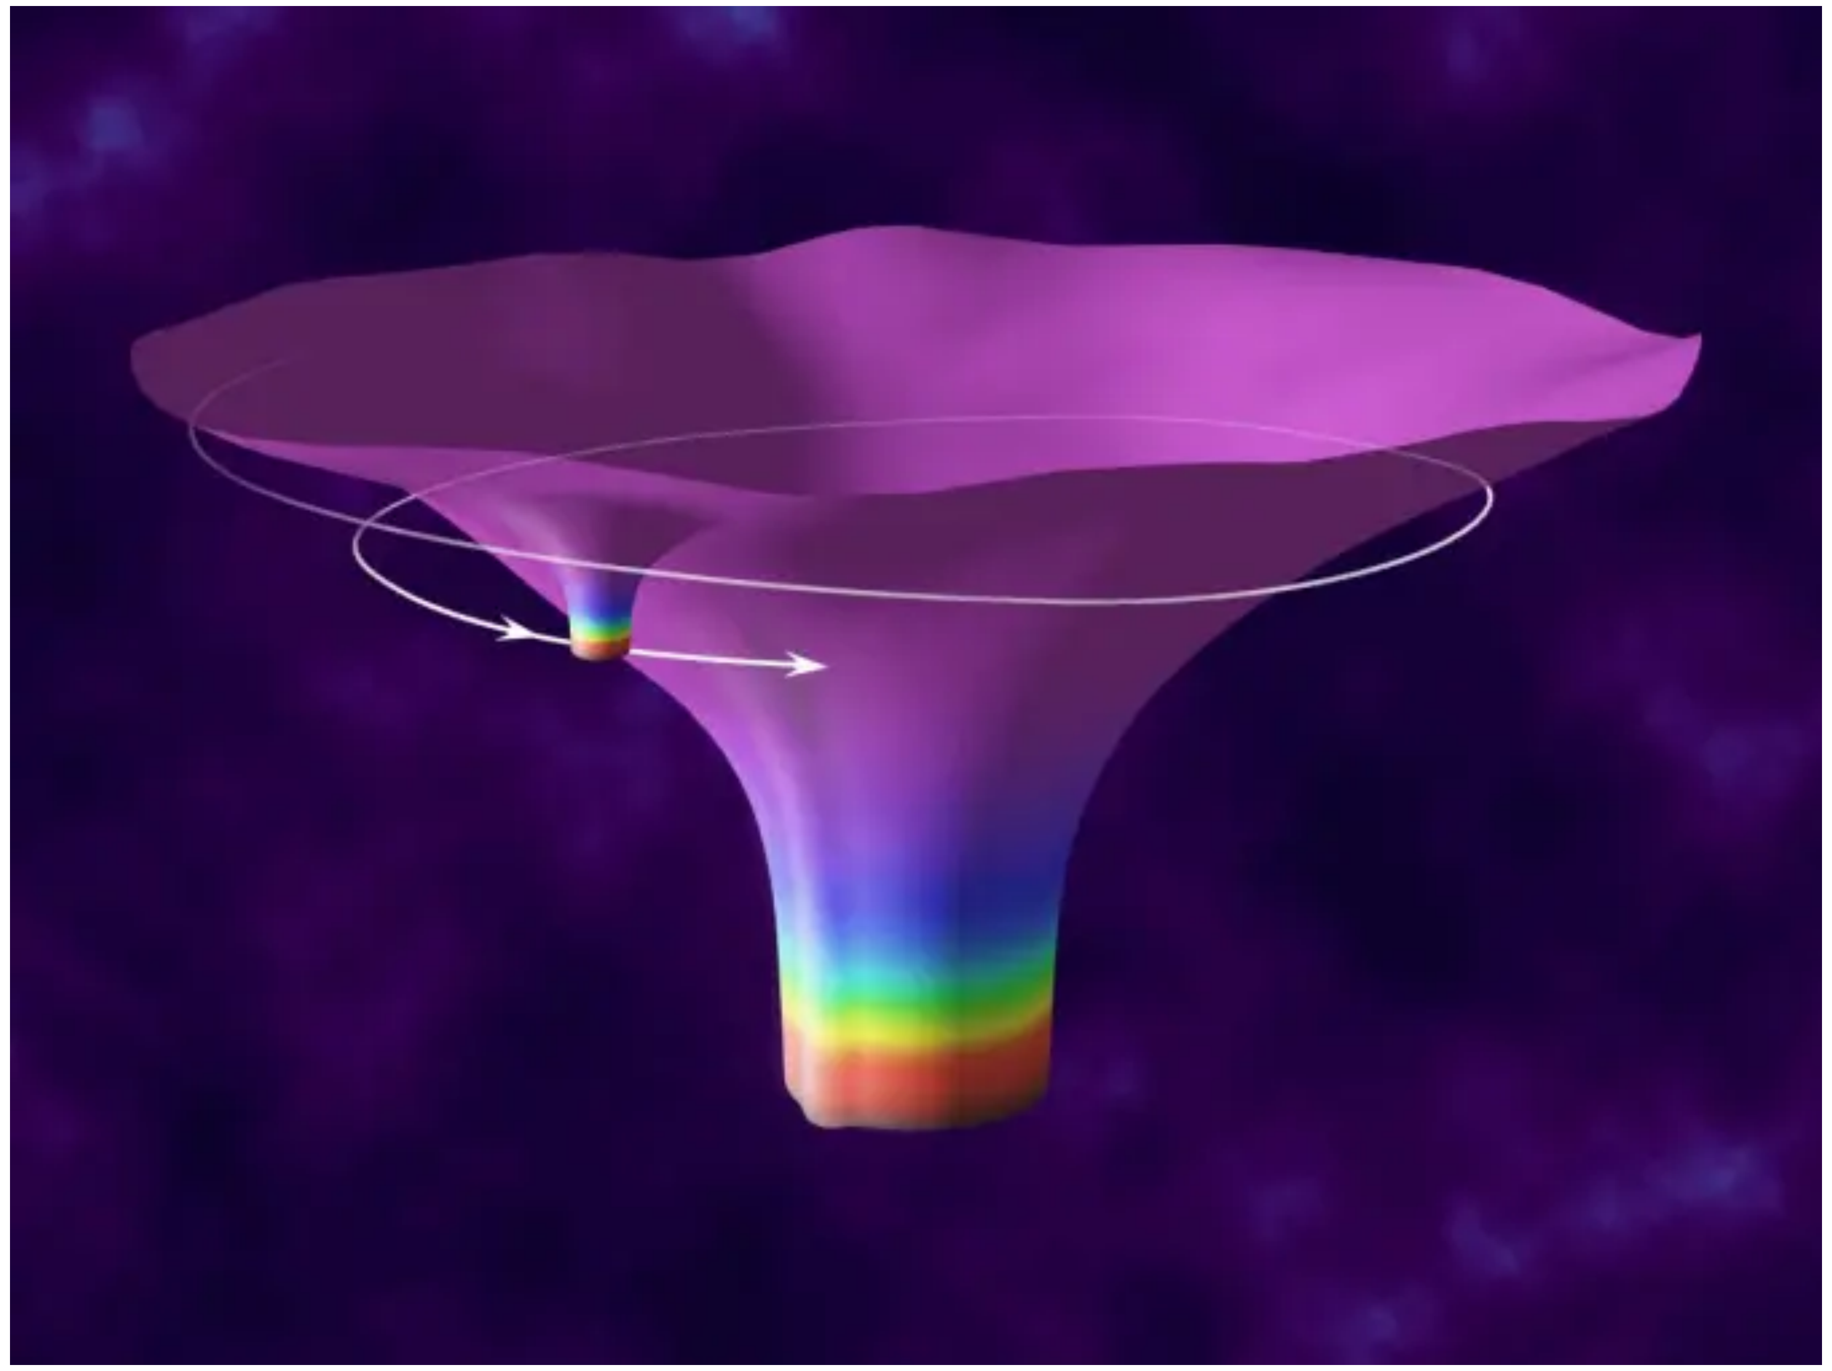

…if we can measure the parameters to very high precision

Credit: NASA

Dark dress

Cold, collisionless dark matter

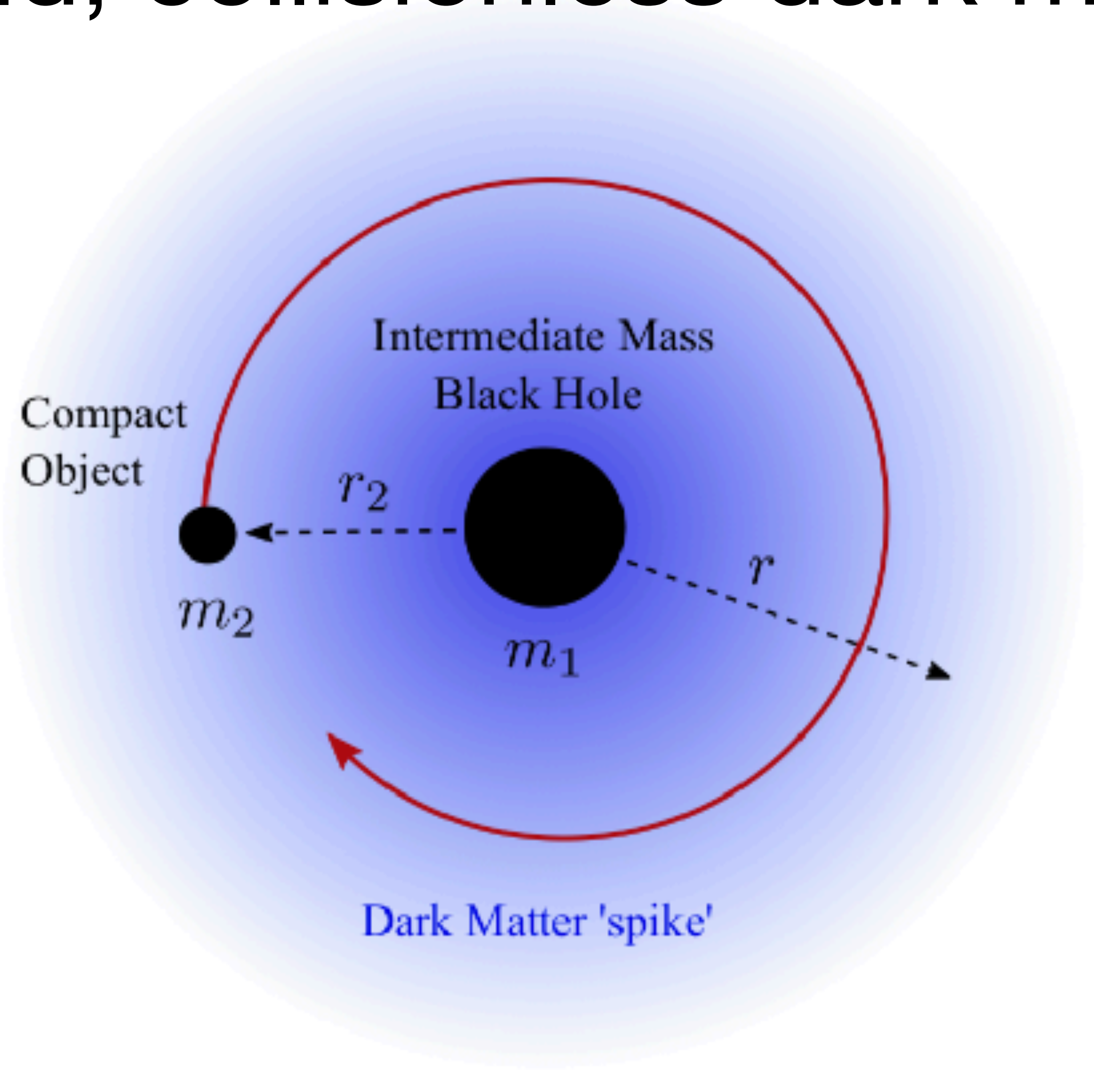

$$\rho(r) = \rho_6 \left(\frac{r_6}{r}\right)^{\gamma_s}$$

Eda et al. 2013, 2014
Gondolo, Silk 1999
Kavanagh et al. 2020
Coogan et al. 2021

Accretion disk

Baryonic matter

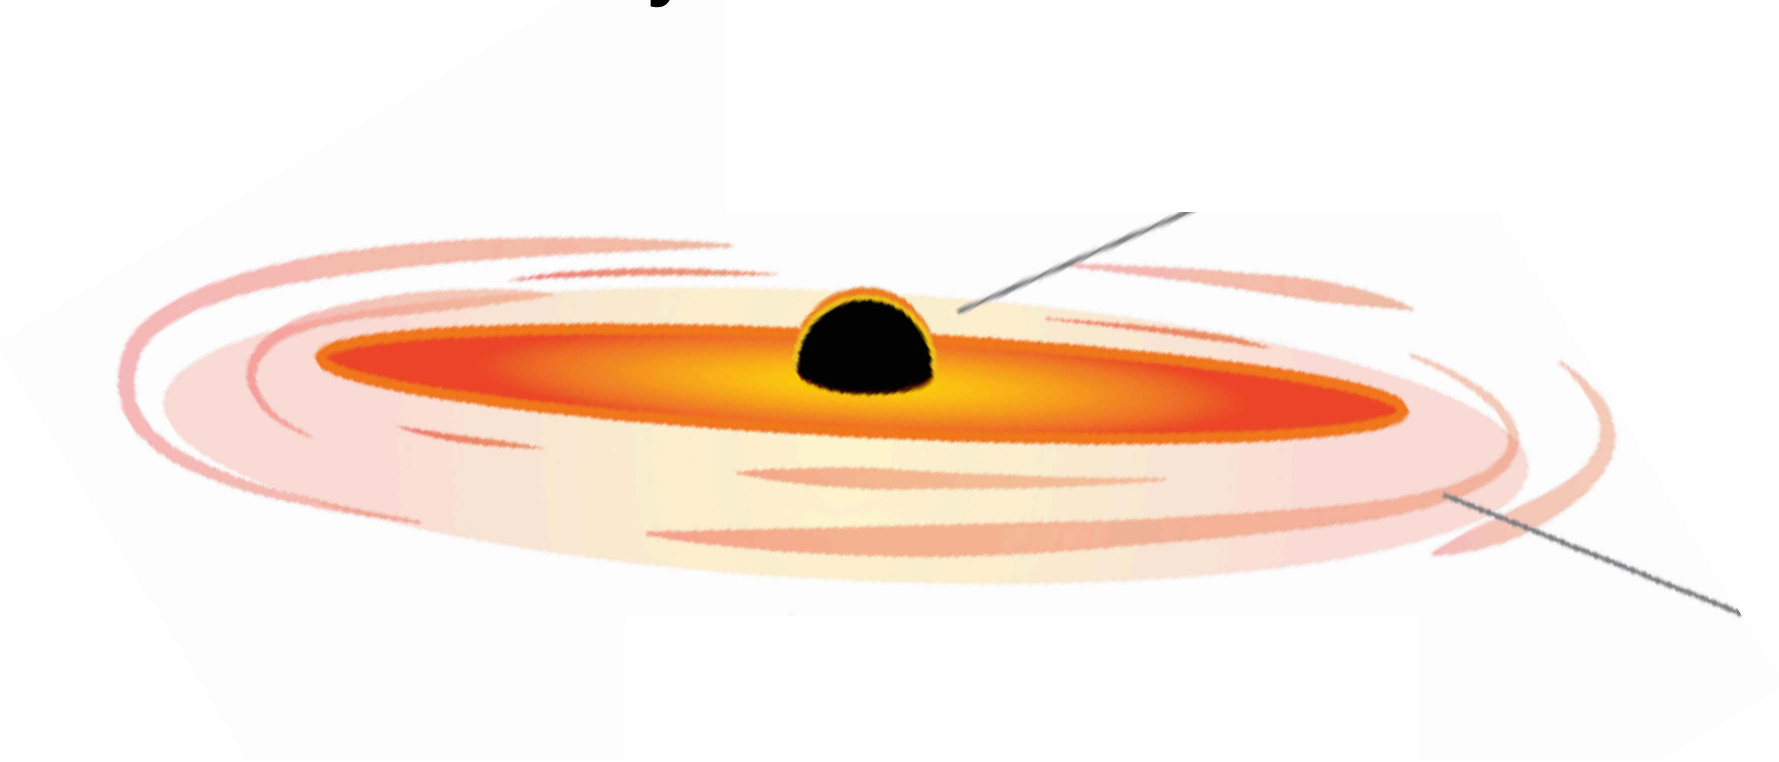

$$\Sigma(r) = \Sigma_0 \left(\frac{r}{r_0}\right)^{-1/2}$$

$$M = r/h$$

Goldreich & Tremaine 1980
Tanaka 2002
Derdzinski et al. 2020

Gravitational atom

Ultra-light bosons

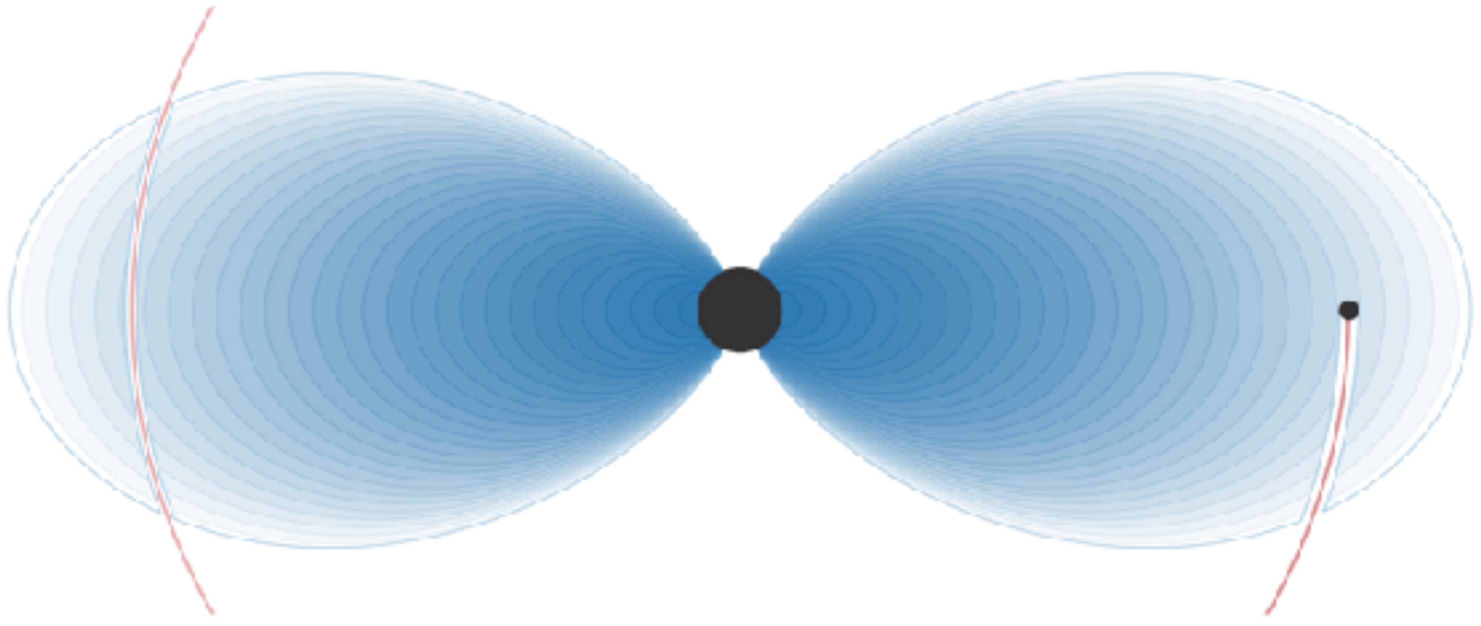

$$\rho(\vec{r}) = M_c |\psi(\vec{r})|^2$$

$$\alpha \equiv G m_1 \mu \ll 1$$

Mass of light scalar field
$(10^{-10} - 10^{-20}\,\mathrm{eV})$

Baumann et al. 2019
Arvanitaki & Dubovsky 2010
Bauman et al. 2021, 2022

Credit: Sophia Dagnello, NRAO/AUI/NSF

Binary loses energy due to gravitational waves, additional losses due to the environments

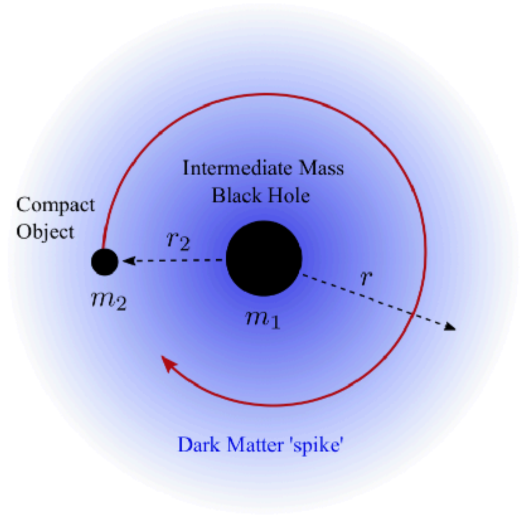

Dynamical friction

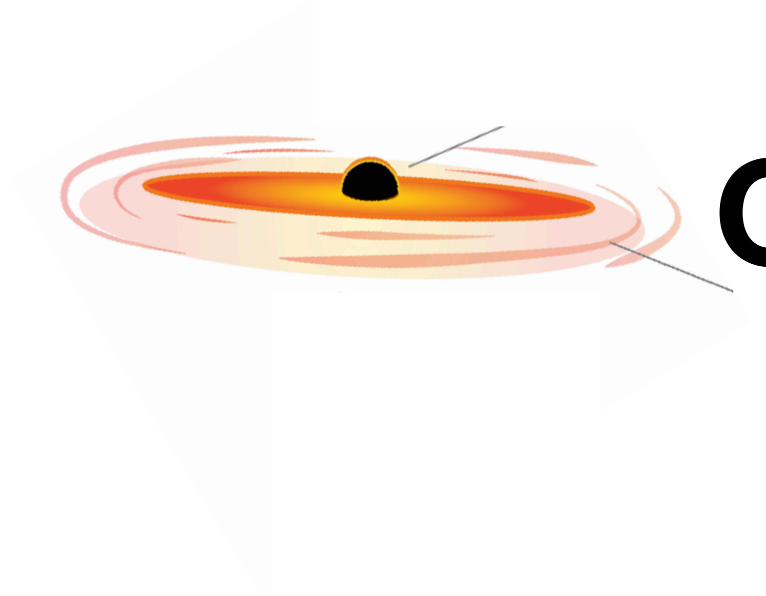

Gas torques

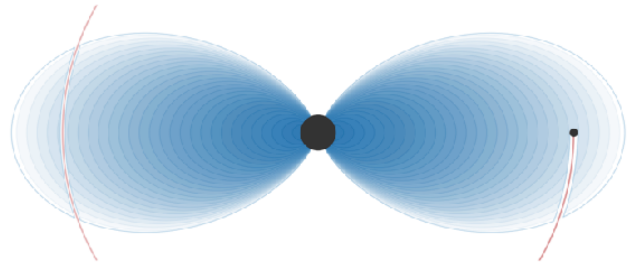

Ionisation

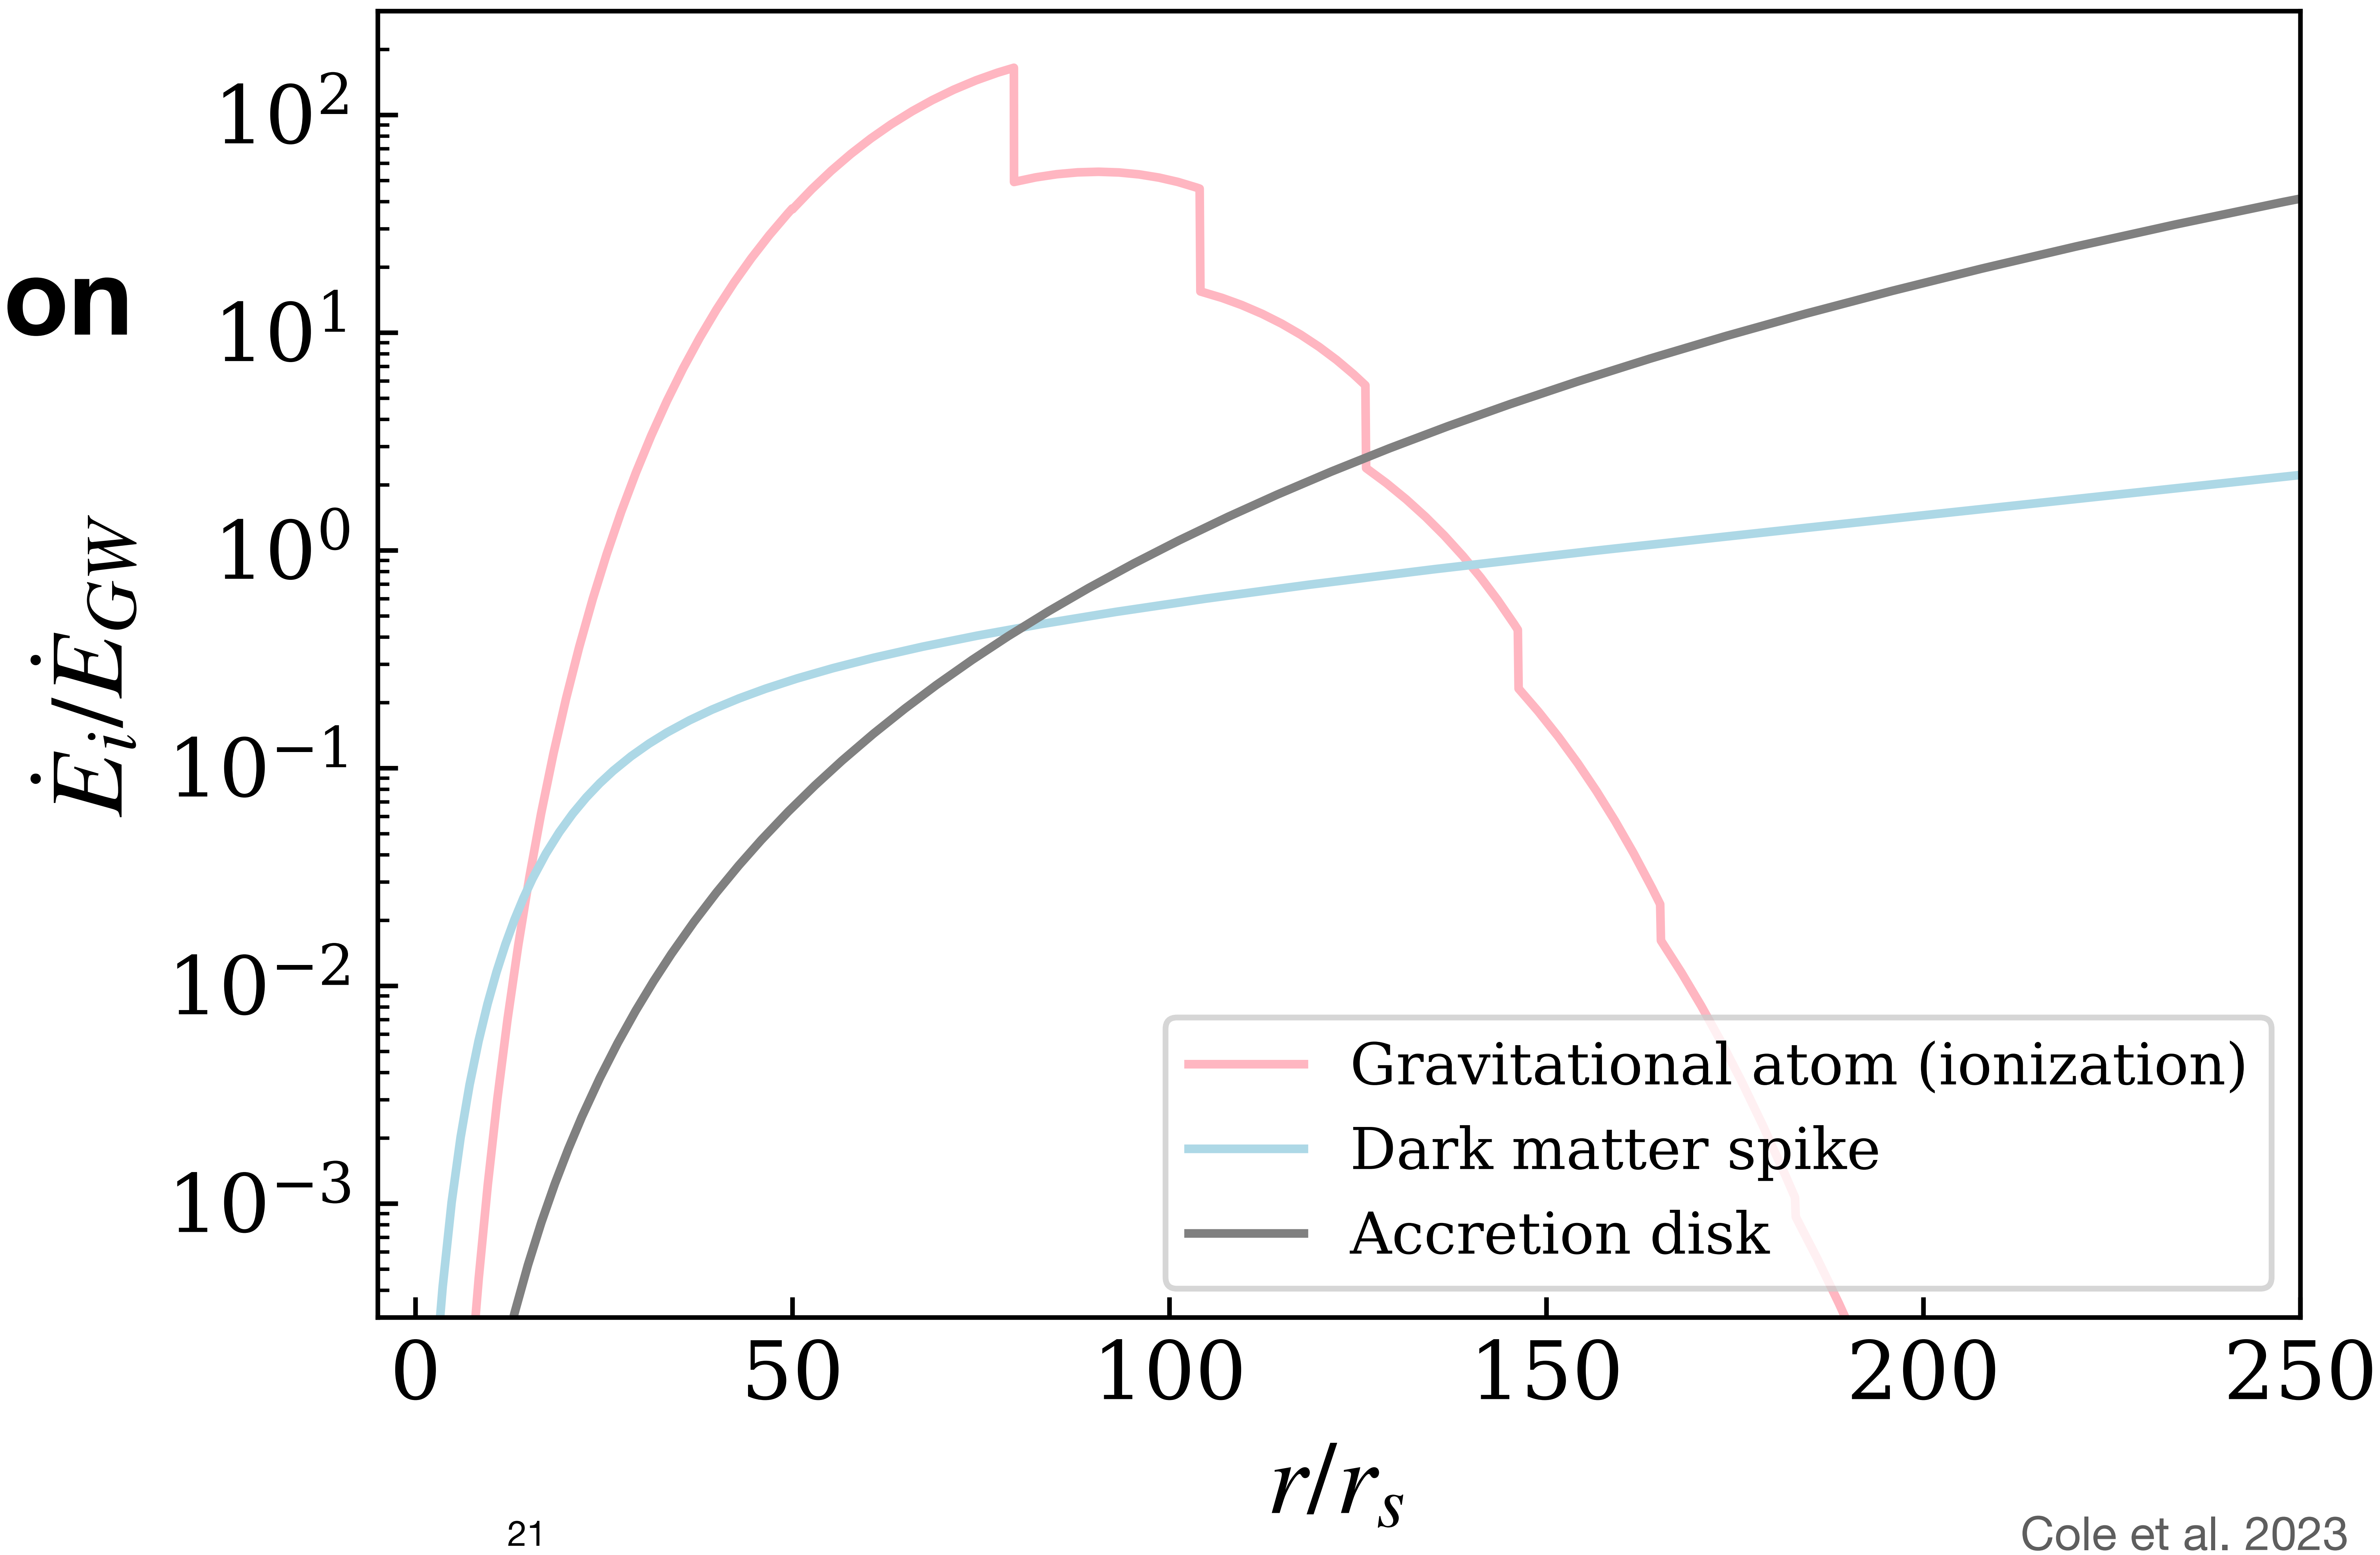

Cole et al. 2023

The information is within the waveforms to be able to measure the parameters of, and distinguish between, different environments.

Accretion disk

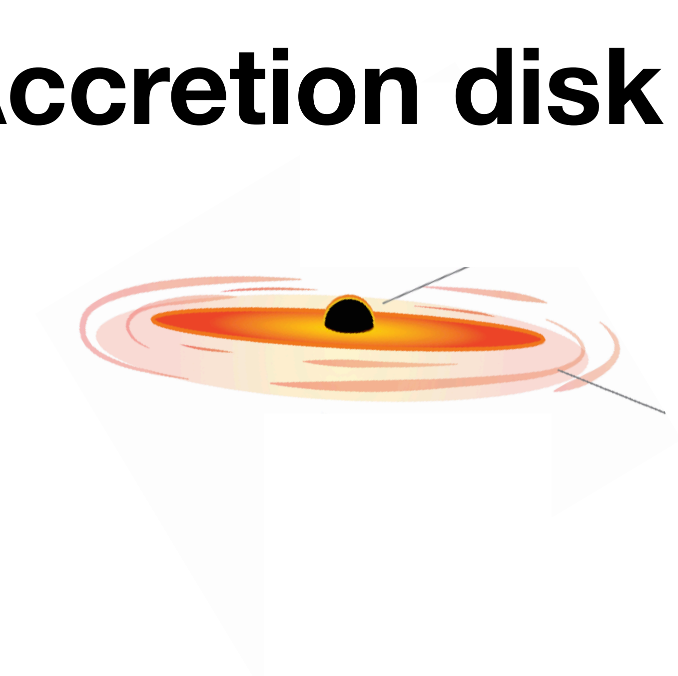

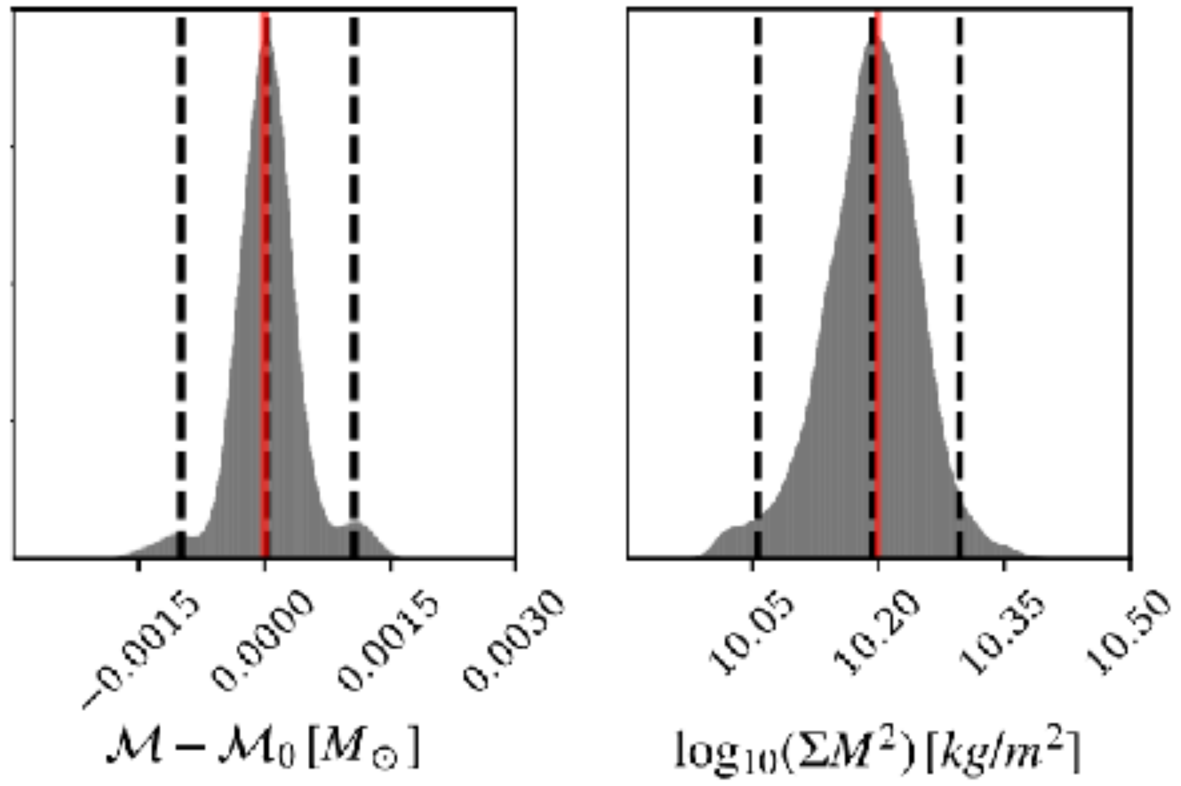

Dark dress

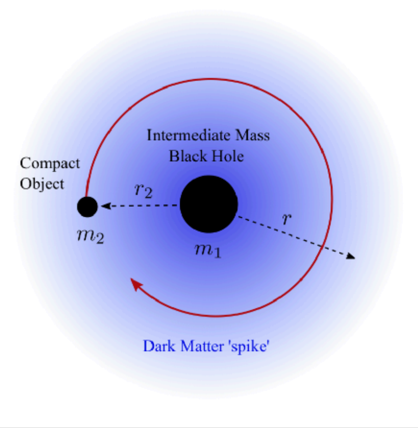

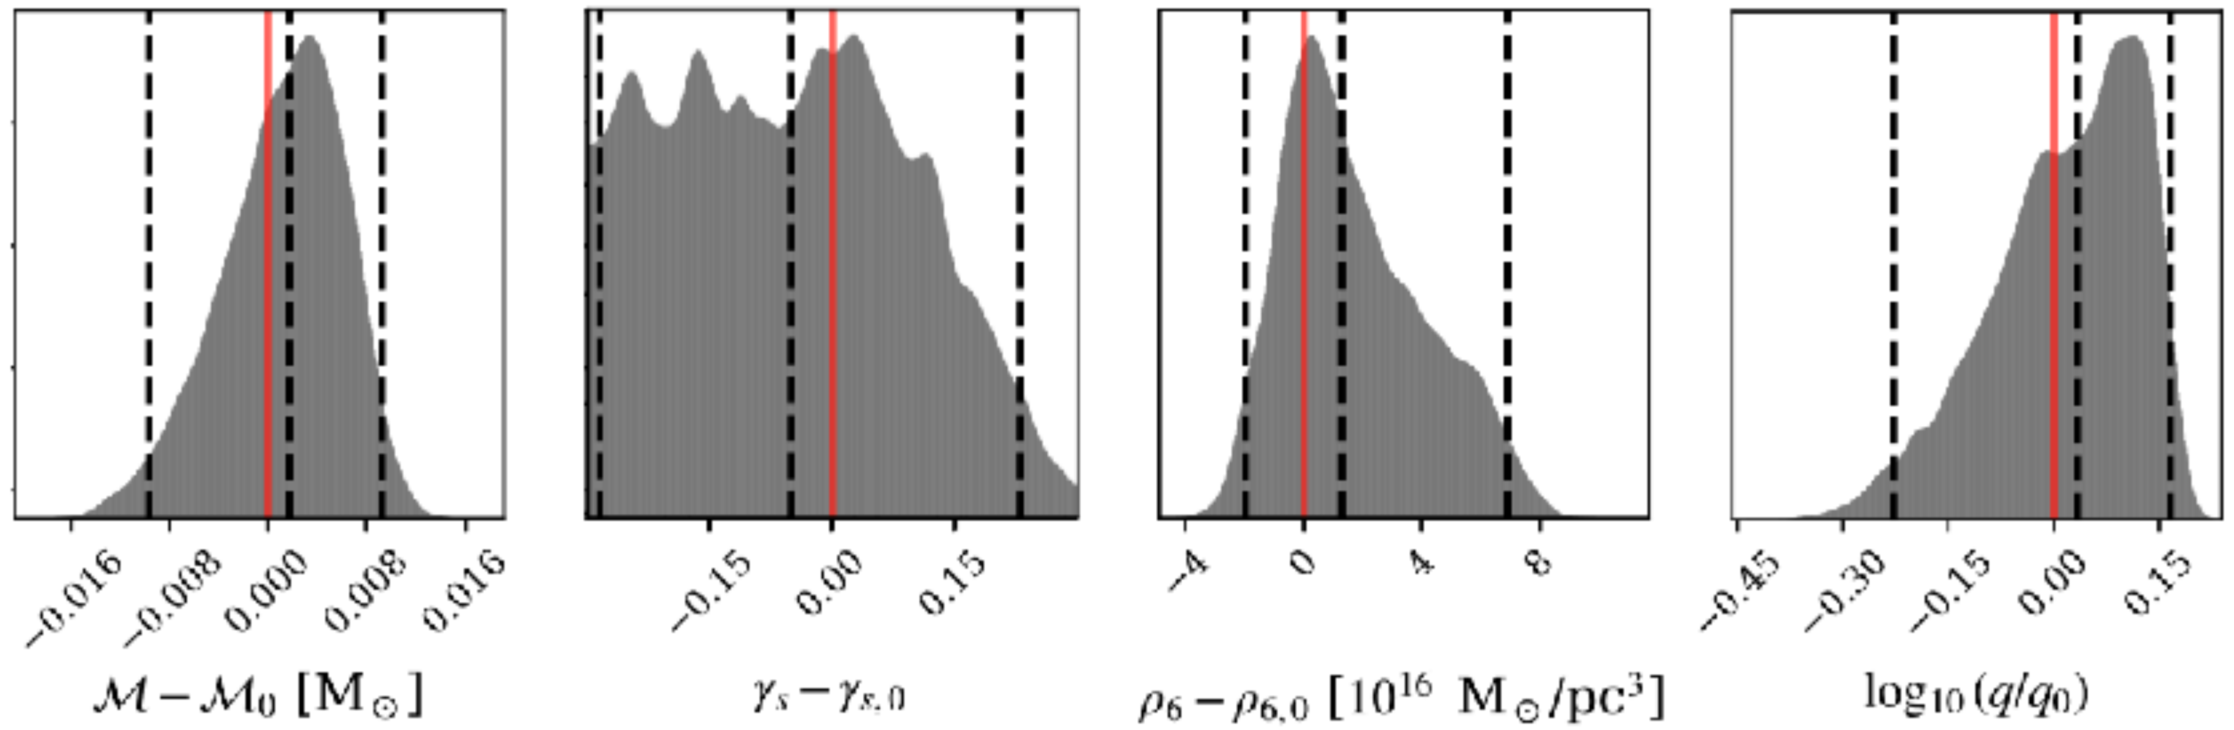

Gravitational atom

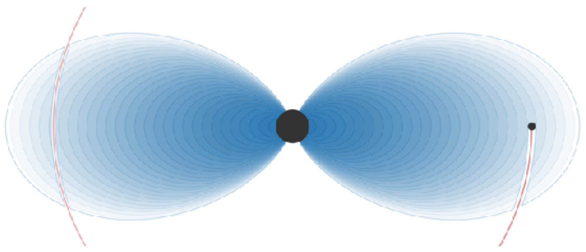

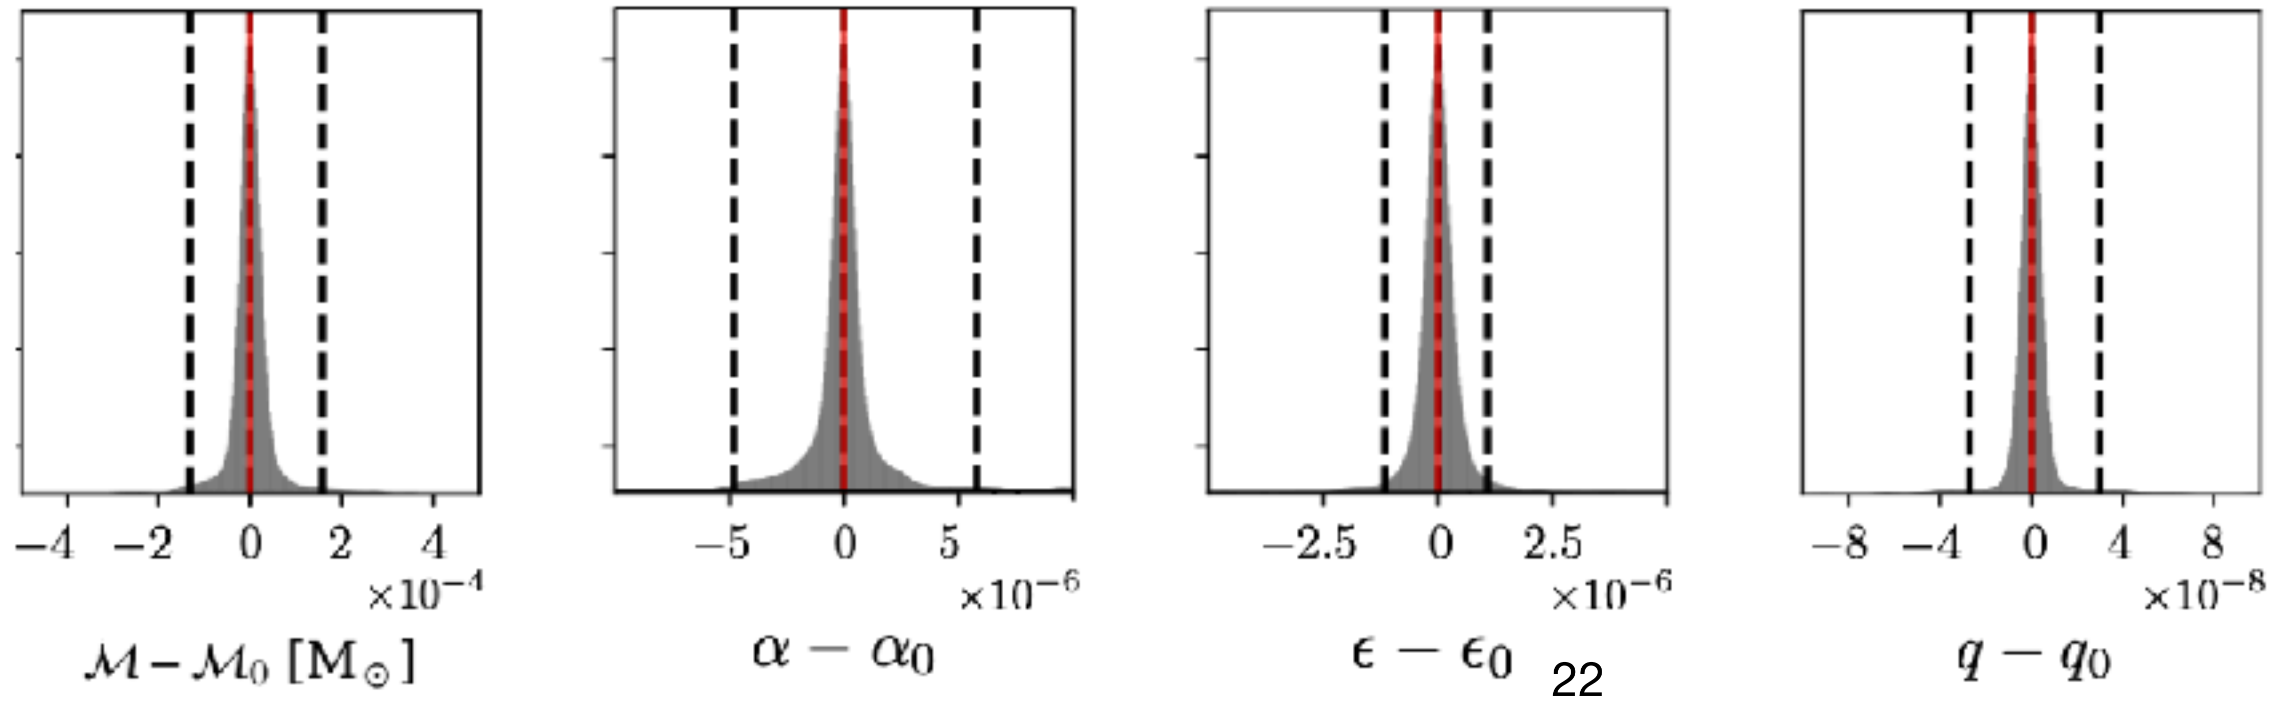

Cole et al. 2023 *Nat Astron* (2023). arXiv:2211.01362

Biases possible if these effects not included in parameter estimation

	Accretion disk signal	Dark dress signal	Gravitational atom signal
Vacuum template	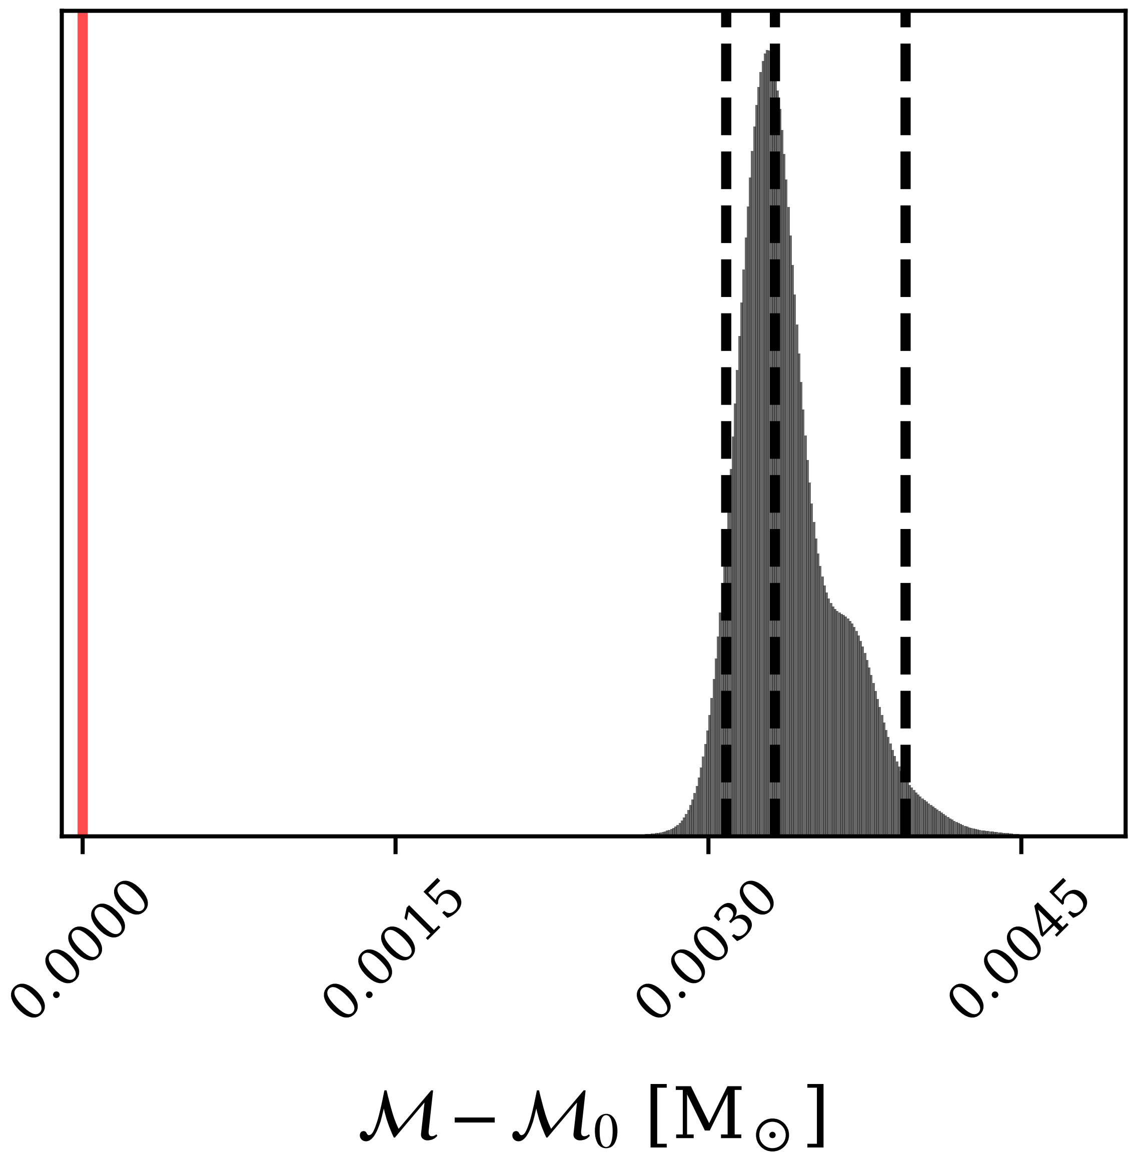	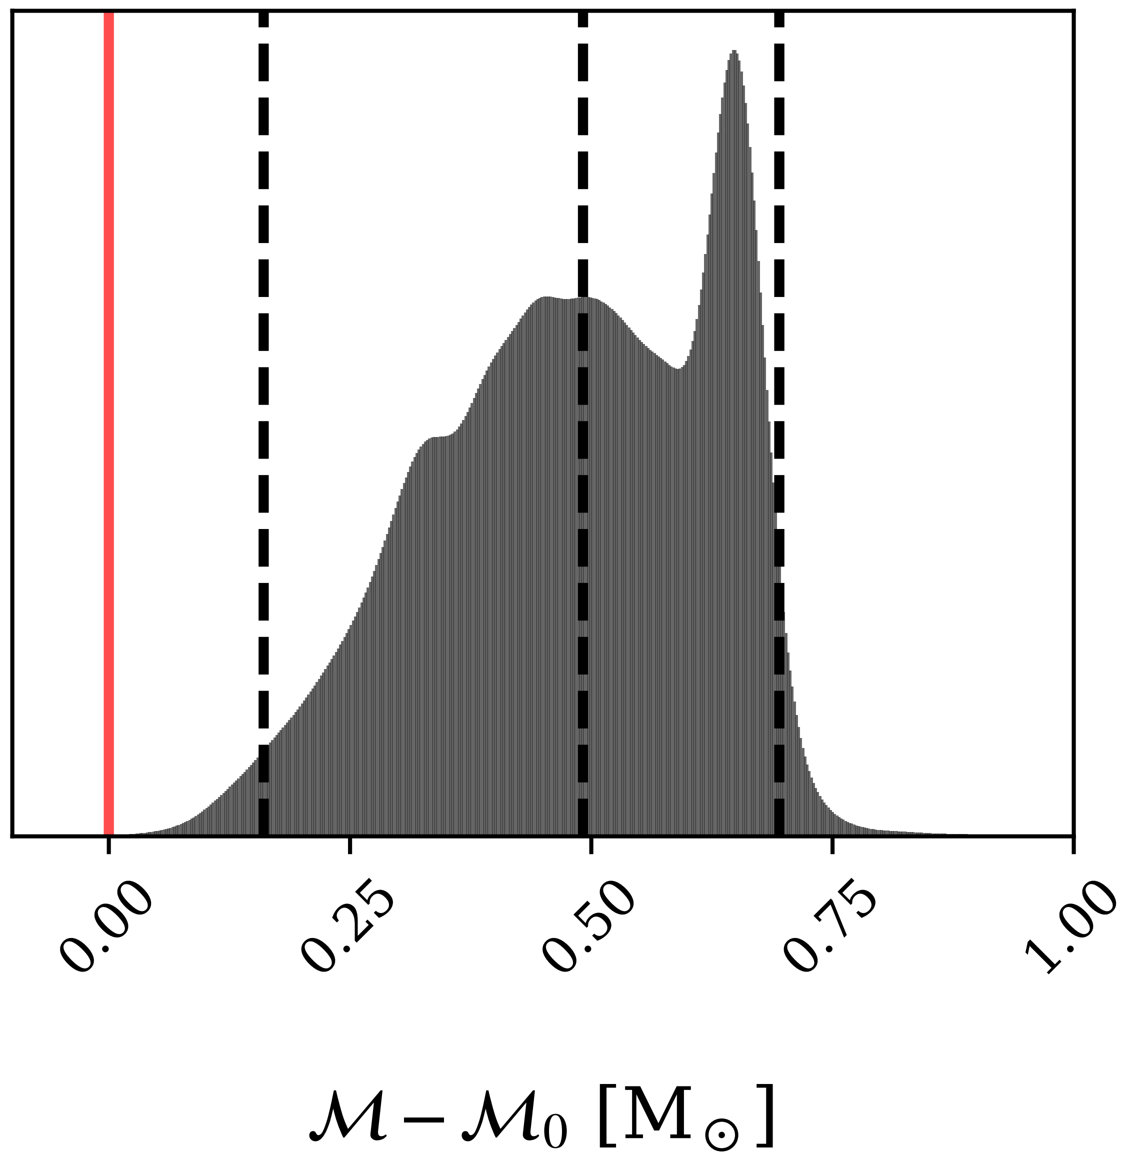	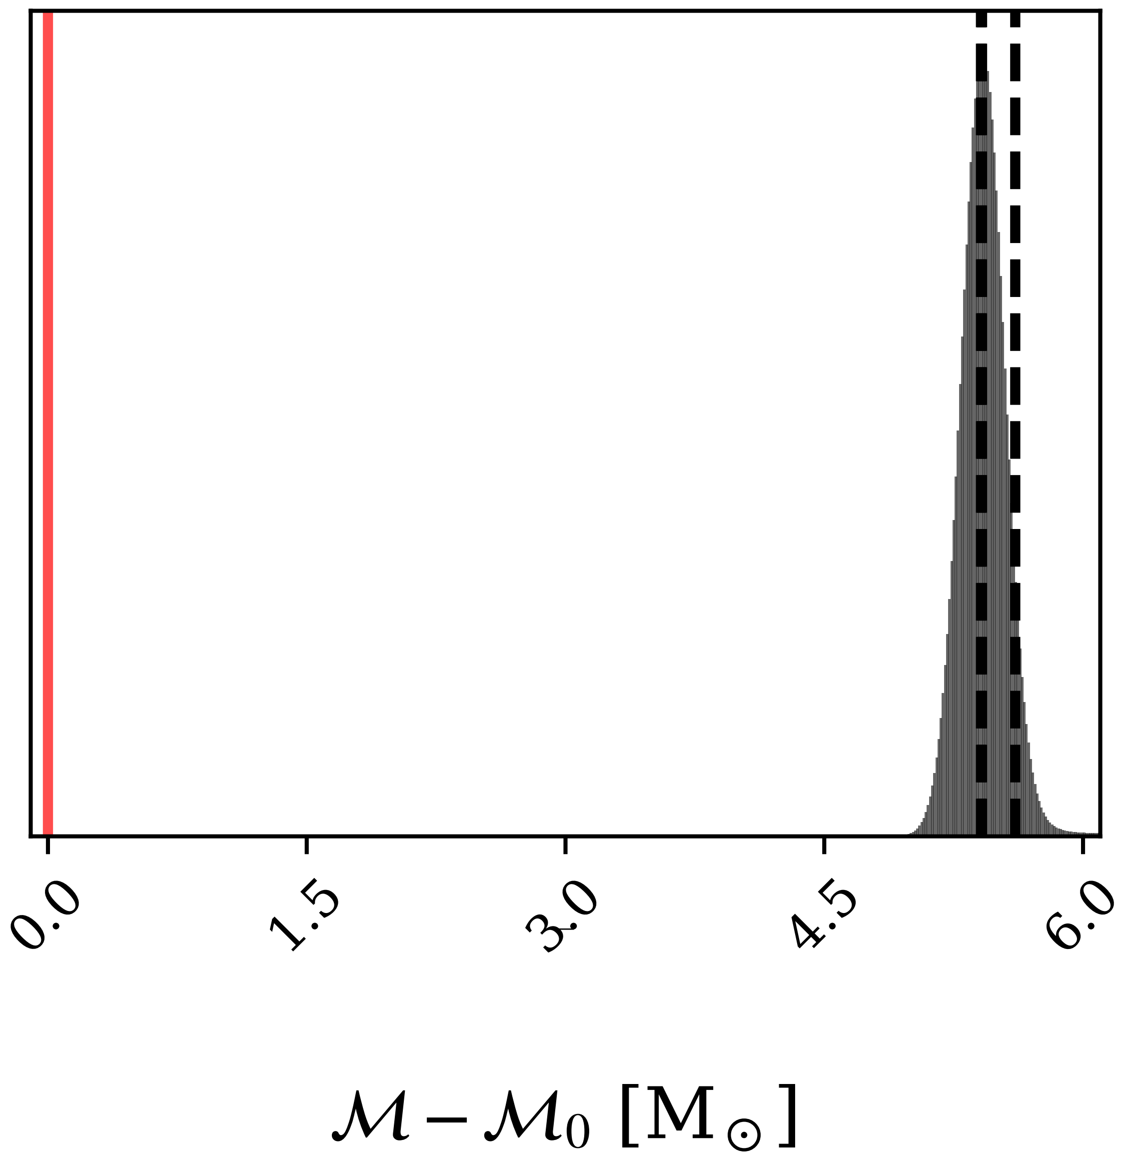

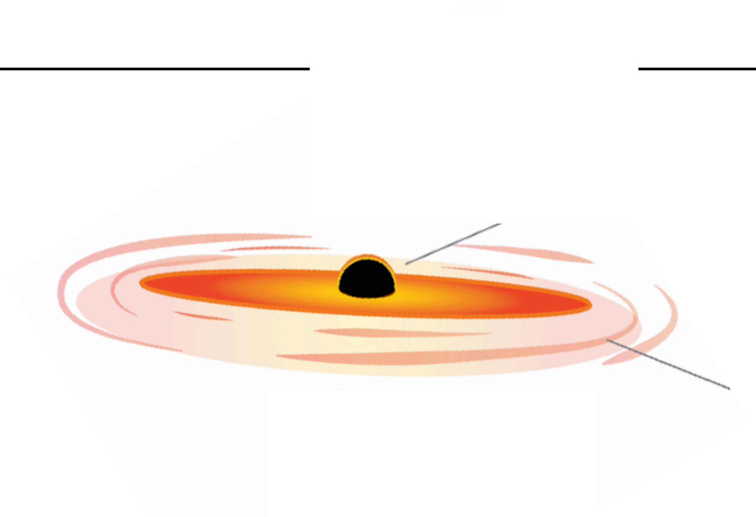

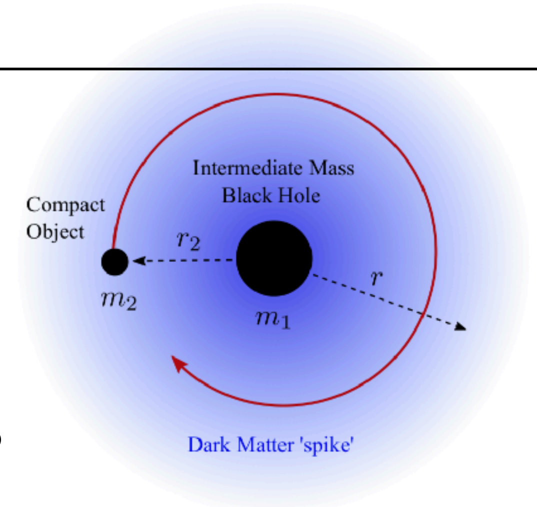

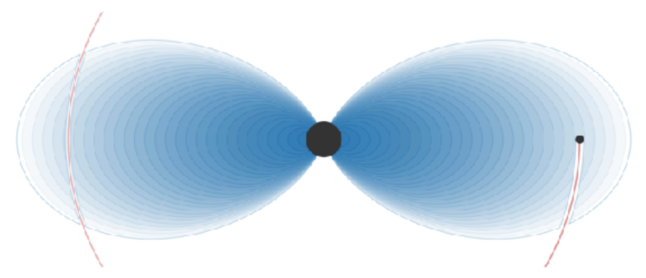

Cole et al. 2023

Confidently distinguish between different environments (and vacuum) with full Bayesian model comparison

Injection

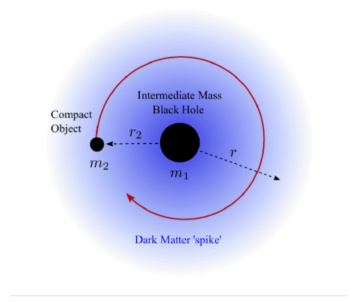 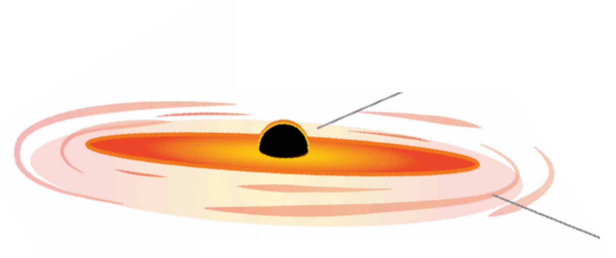 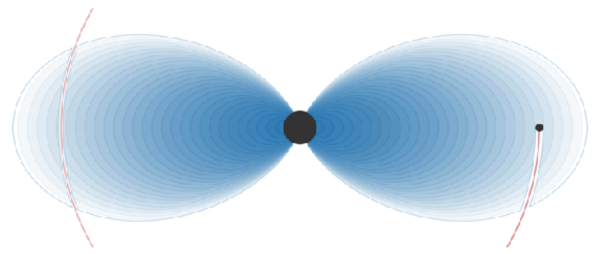

Recovery

$\log_{10} \mathscr{B}$	**Dark dress signal**	**Accretion disk signal**	**Gravitational atom signal**
Vacuum template	34	6	39
Dark dress template	-	3	39
Accretion disk template	17	-	33
Gravitational atom template	24	6	-

Cole et al. 2023

See also Hanneksula et al. Nature Astronomy, 2019

Also possible with accretion disk + dark matter simultaneously present

Measure non-zero density normalisations of both with correct template

Bayes factor prefers combined model significantly with respect to recovery with single environments

Nicole Grillo MSc thesis 2025

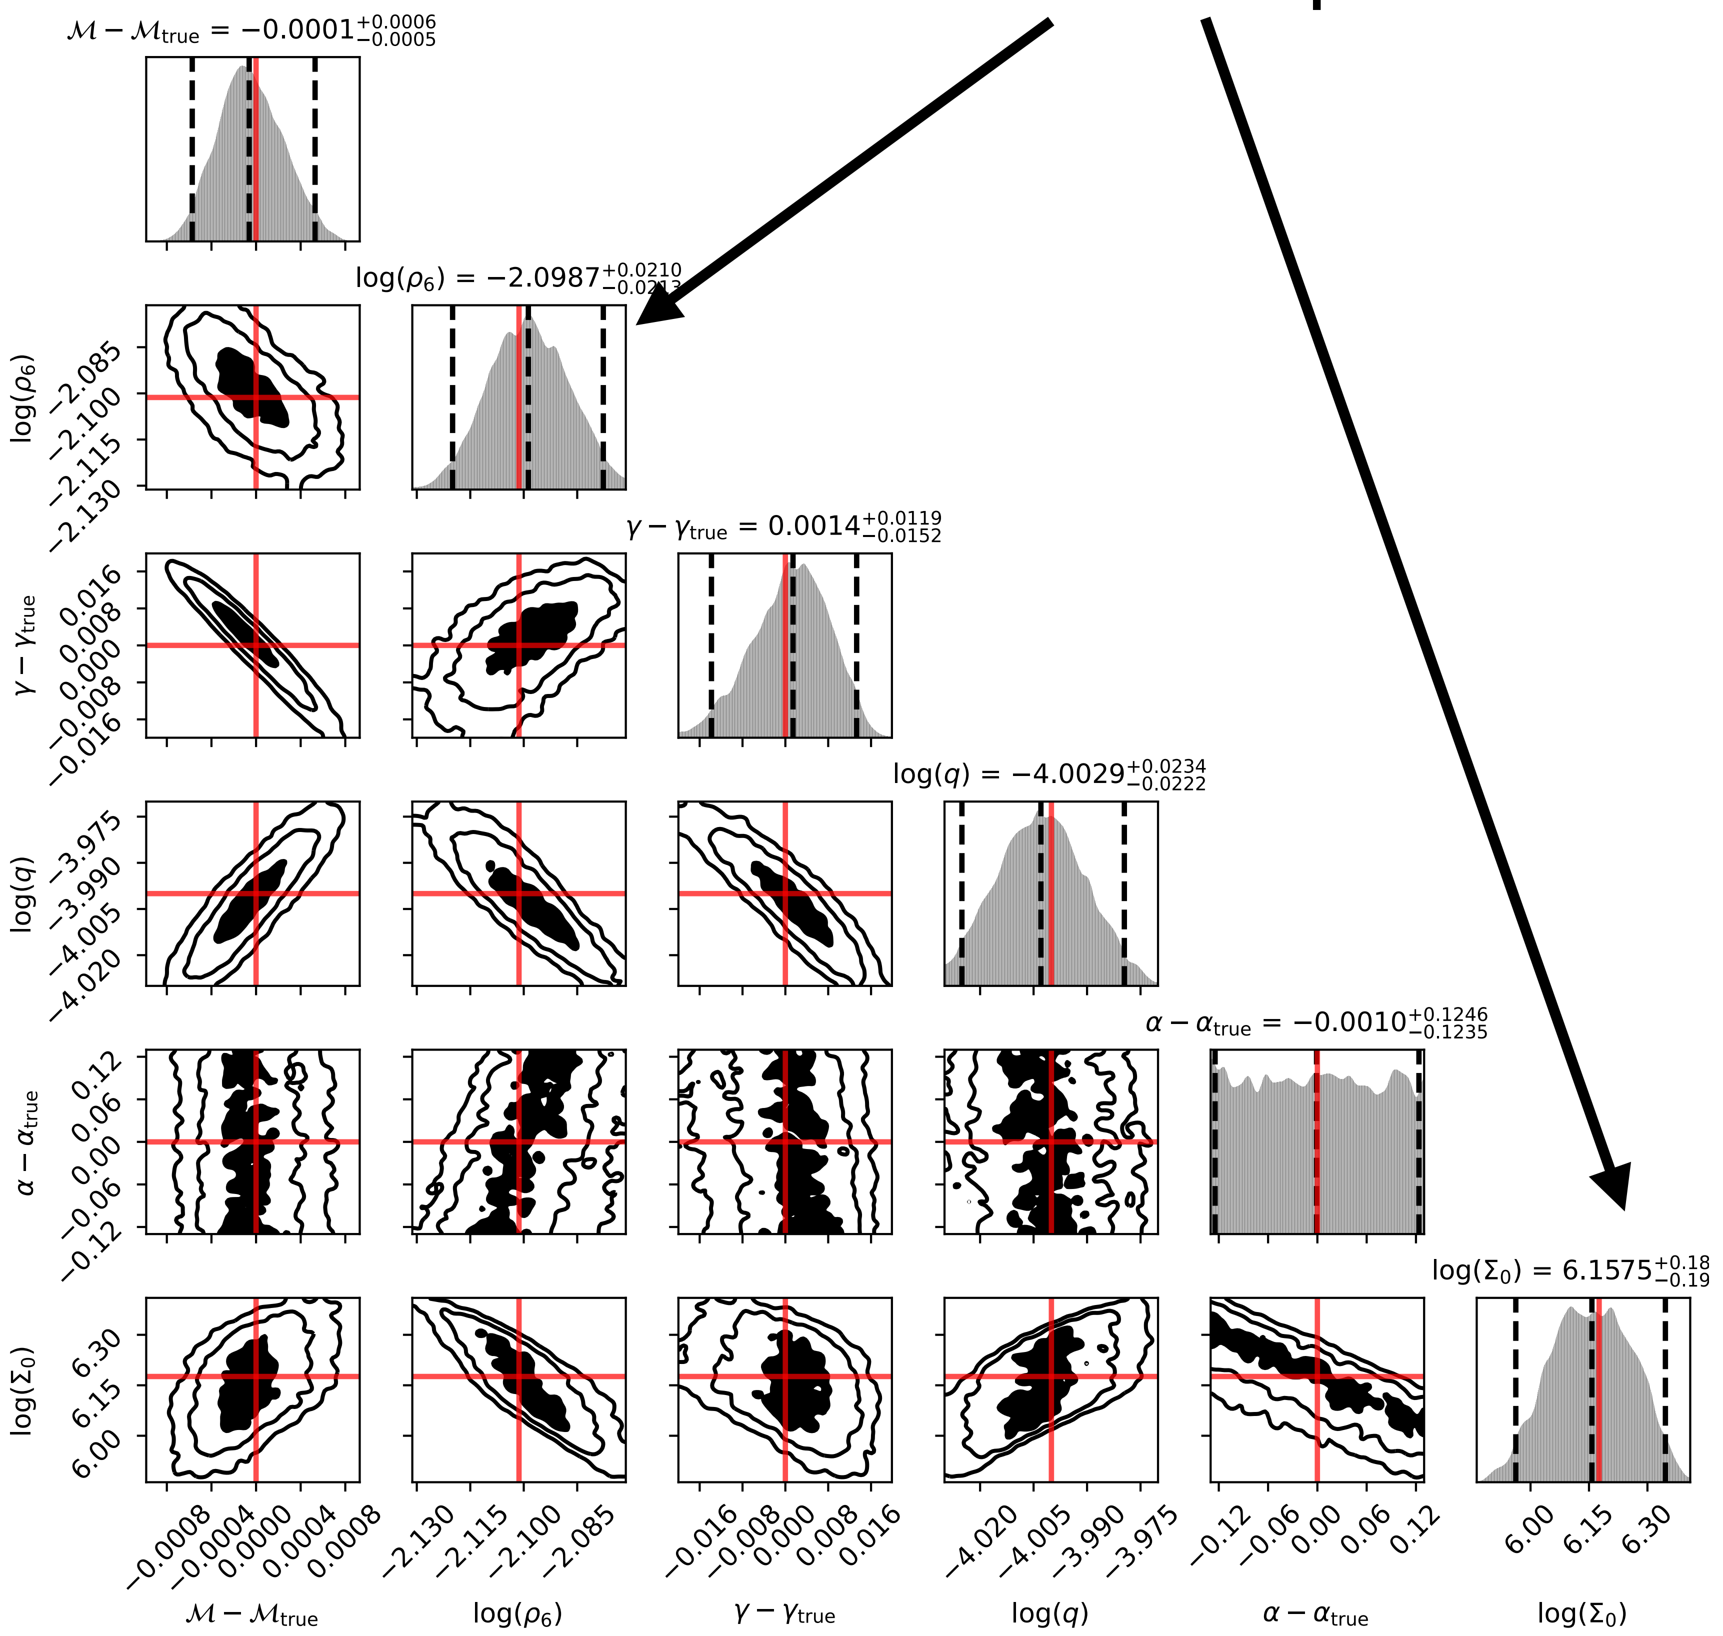

Recovery with correct combined model

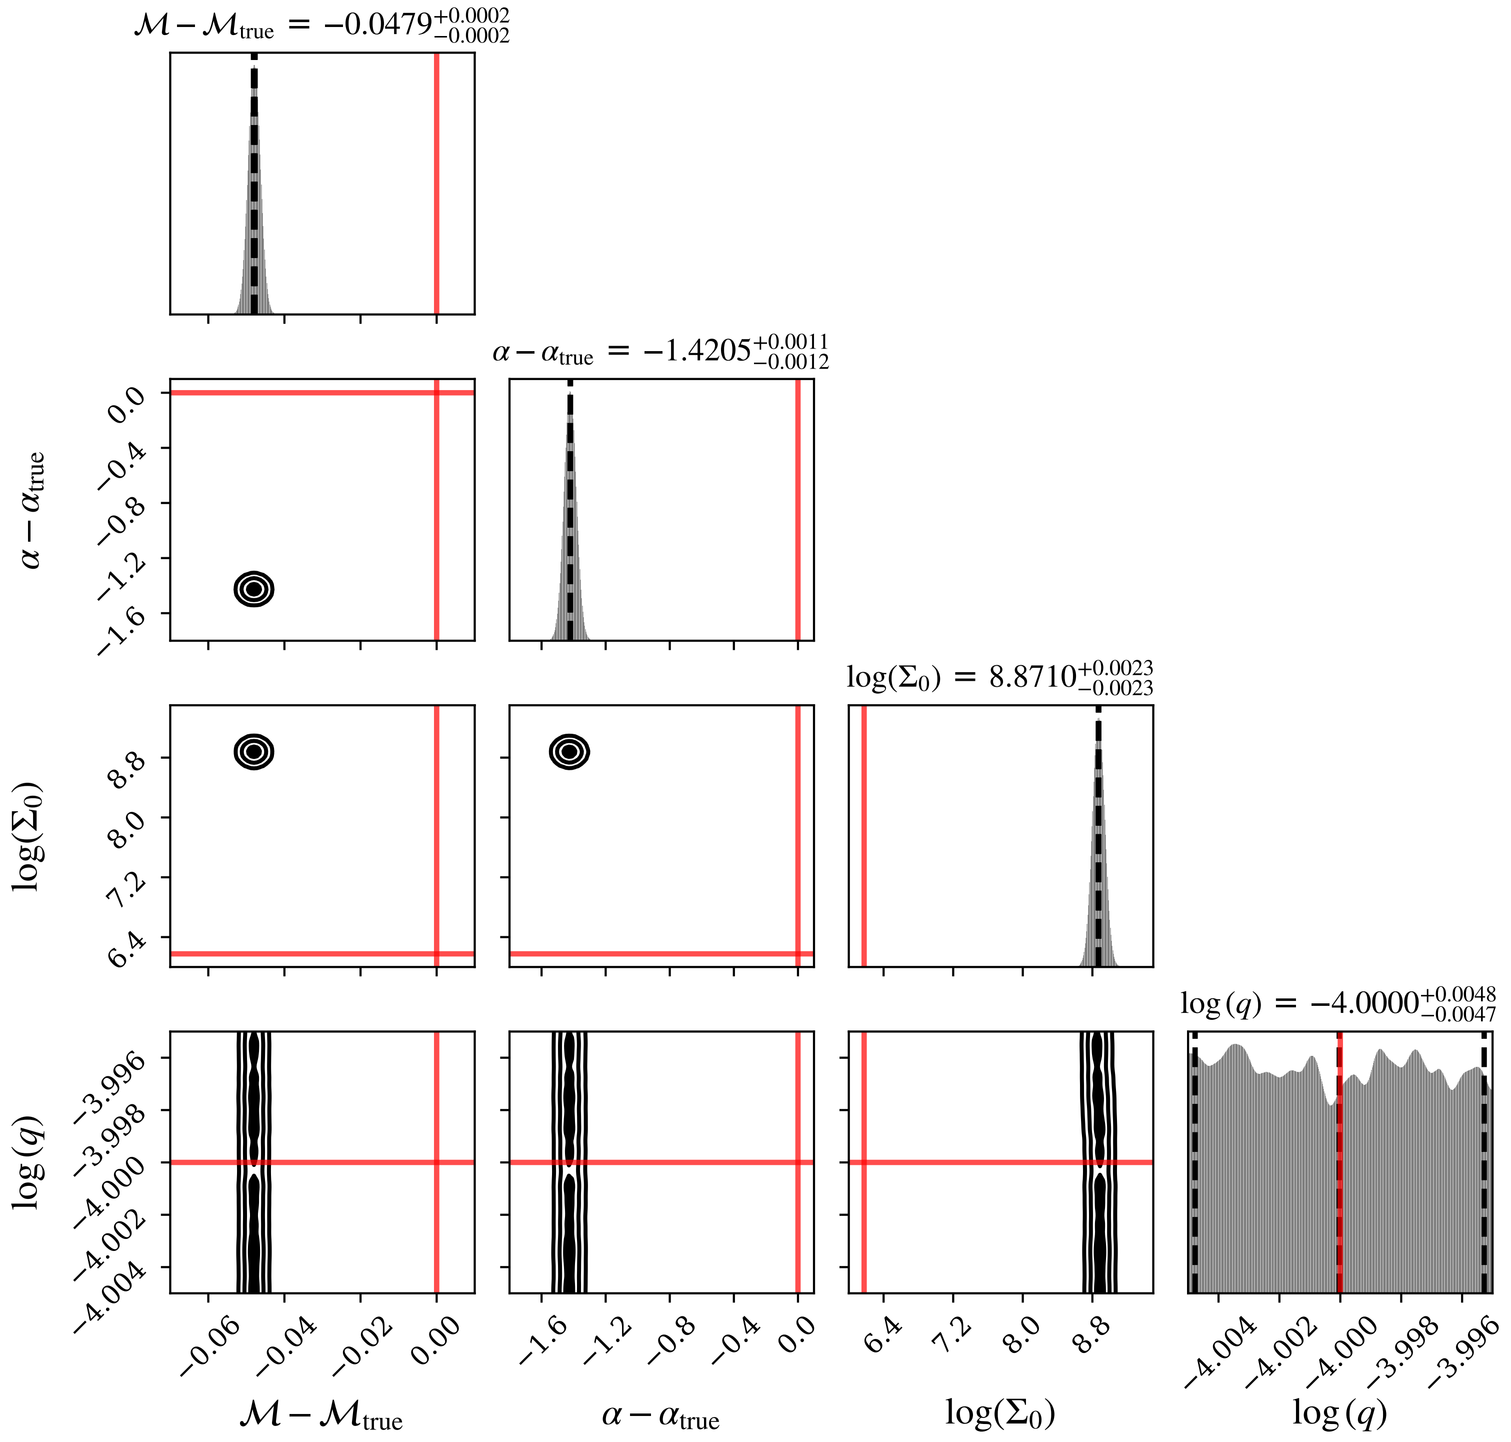

Recovery with accretion disk only

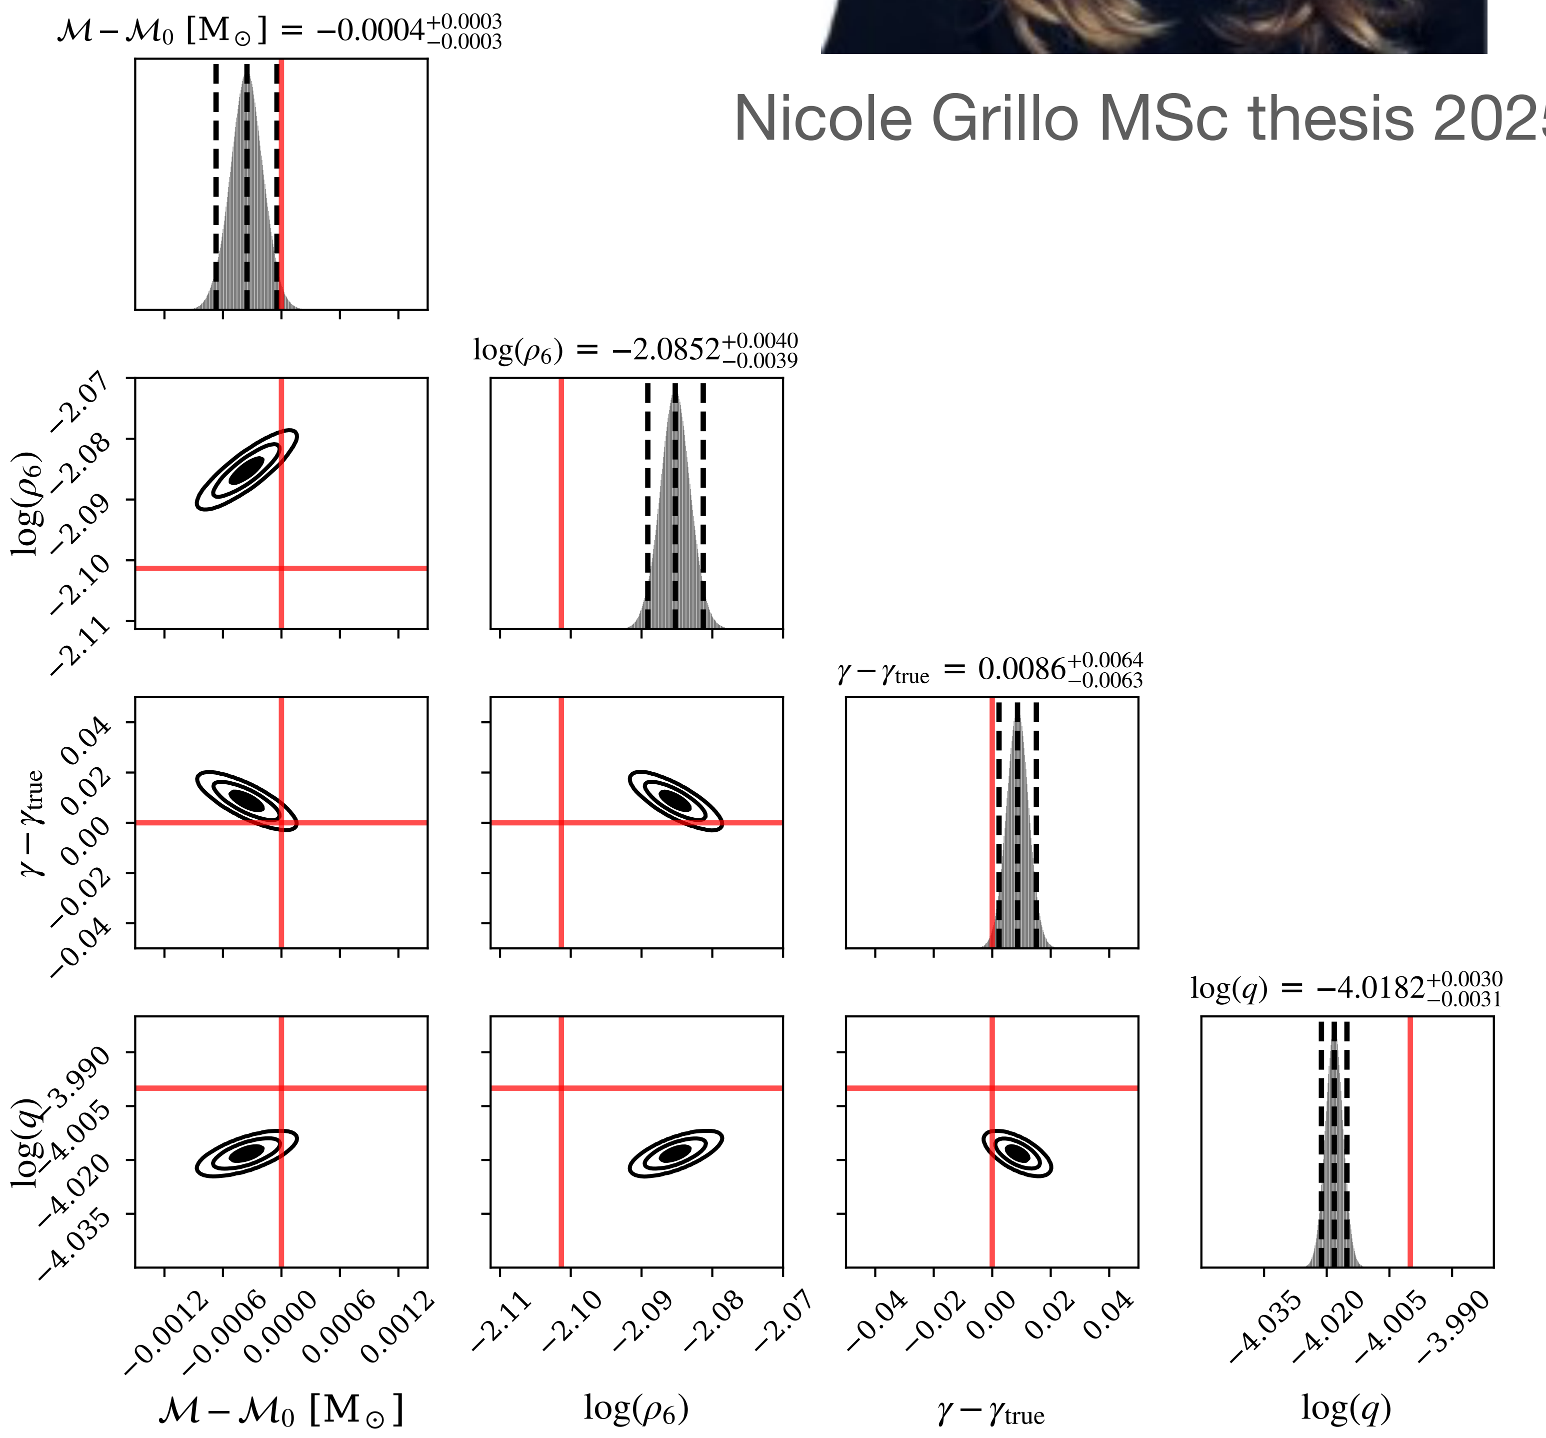

Recovery with dark matter spike only

Why should we care about environmental effects?

- We have a chance to learn about the environment itself (which could involve dark matter) via the dephasing in the waveform.
- If we search the data without including these effects we might miss the signal.
- If we do parameter estimation with the ‘wrong’ model, results will be biased.

See also Barausse, Cardoso, Pani 2011

Can we do this in practice?

Need to include (non-exhaustive list):

- Eccentricity
- Relativistic effects
- Detector response
- Noise
- Spins
- Higher order environmental effects
- Overlapping signals
- Other sources - global fit

Can we do this in practice?

“Running before you can walk”

Need to include (non-exhaustive list):

- Eccentricity
- Relativistic effects
- Detector response
- Noise
- Spins
- Higher order environmental effects
- Overlapping signals
- Other sources - global fit

} **running**

Can we do this in practice?

“Running before you can walk”

Need to include (non-exhaustive list):

- Eccentricity
- Relativistic effects
- Detector response
- Noise
- Spins
- Higher order environmental effects
- Overlapping signals
- Other sources - global fit

walking (Eccentricity, Relativistic effects, Detector response, Noise)

running (all of the above)

Schwarzschild vacuum EMRIs (produced with FEW v1) with detector response and noise

PEREGRINE
Gravitational Wave Parameter Inference with Neural Ratio Estimation

Training set-up

PEREGRINE-style approach - Truncated Marginal Neural Ratio Estimation

- Mappings are learnt marginally, e.g. 1d or 2d as opposed to full joint posterior
- Simulation efficient, and could be an interesting approach when nuisance parameters present/re-introduced

$$\mathscr{L}[\hat{\rho}_{k,\phi}] = -\int \left\{ p(x,\theta_k) \ln \sigma(\hat{\rho}_{k,\phi}(x,\theta_k)) + p(x)p(\theta_k) \ln \left[1 - \sigma(\hat{\rho}_{k,\phi}(x,\theta_k)) \right] \right\} dx\, d\theta_k\,.$$

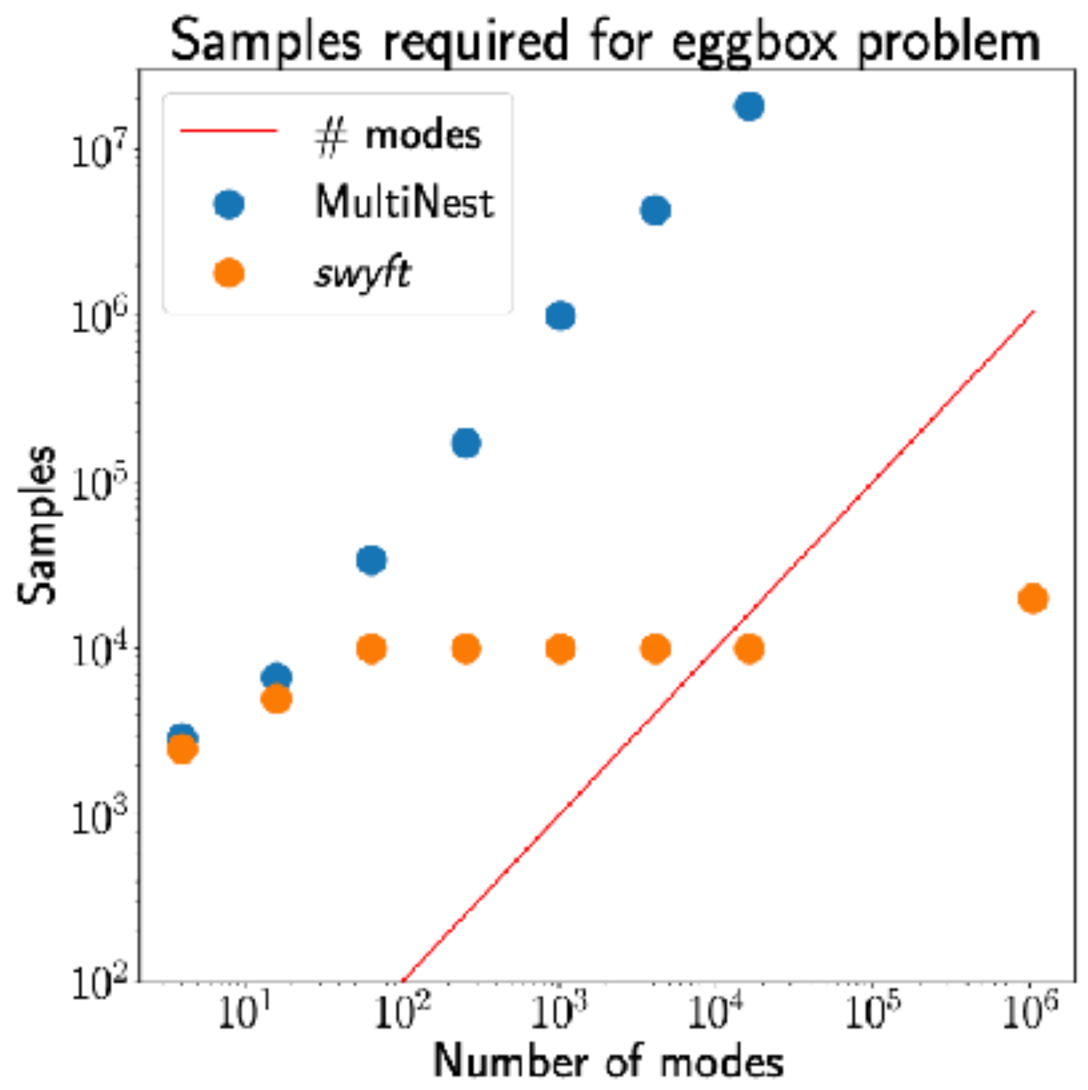

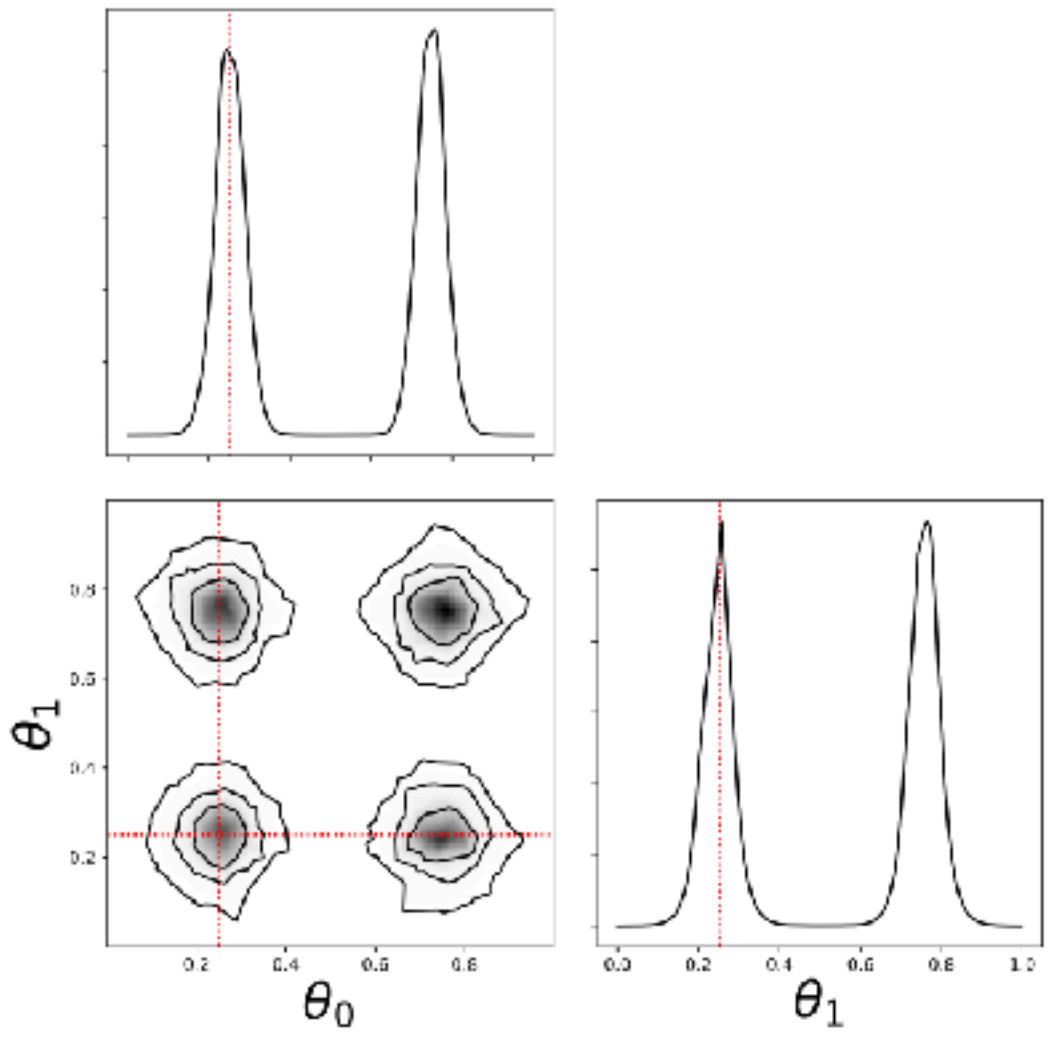

Miller et al. 2020, Miller et al. 2021

Differential evolution, with optimisation over SNR, performs better than the likelihood

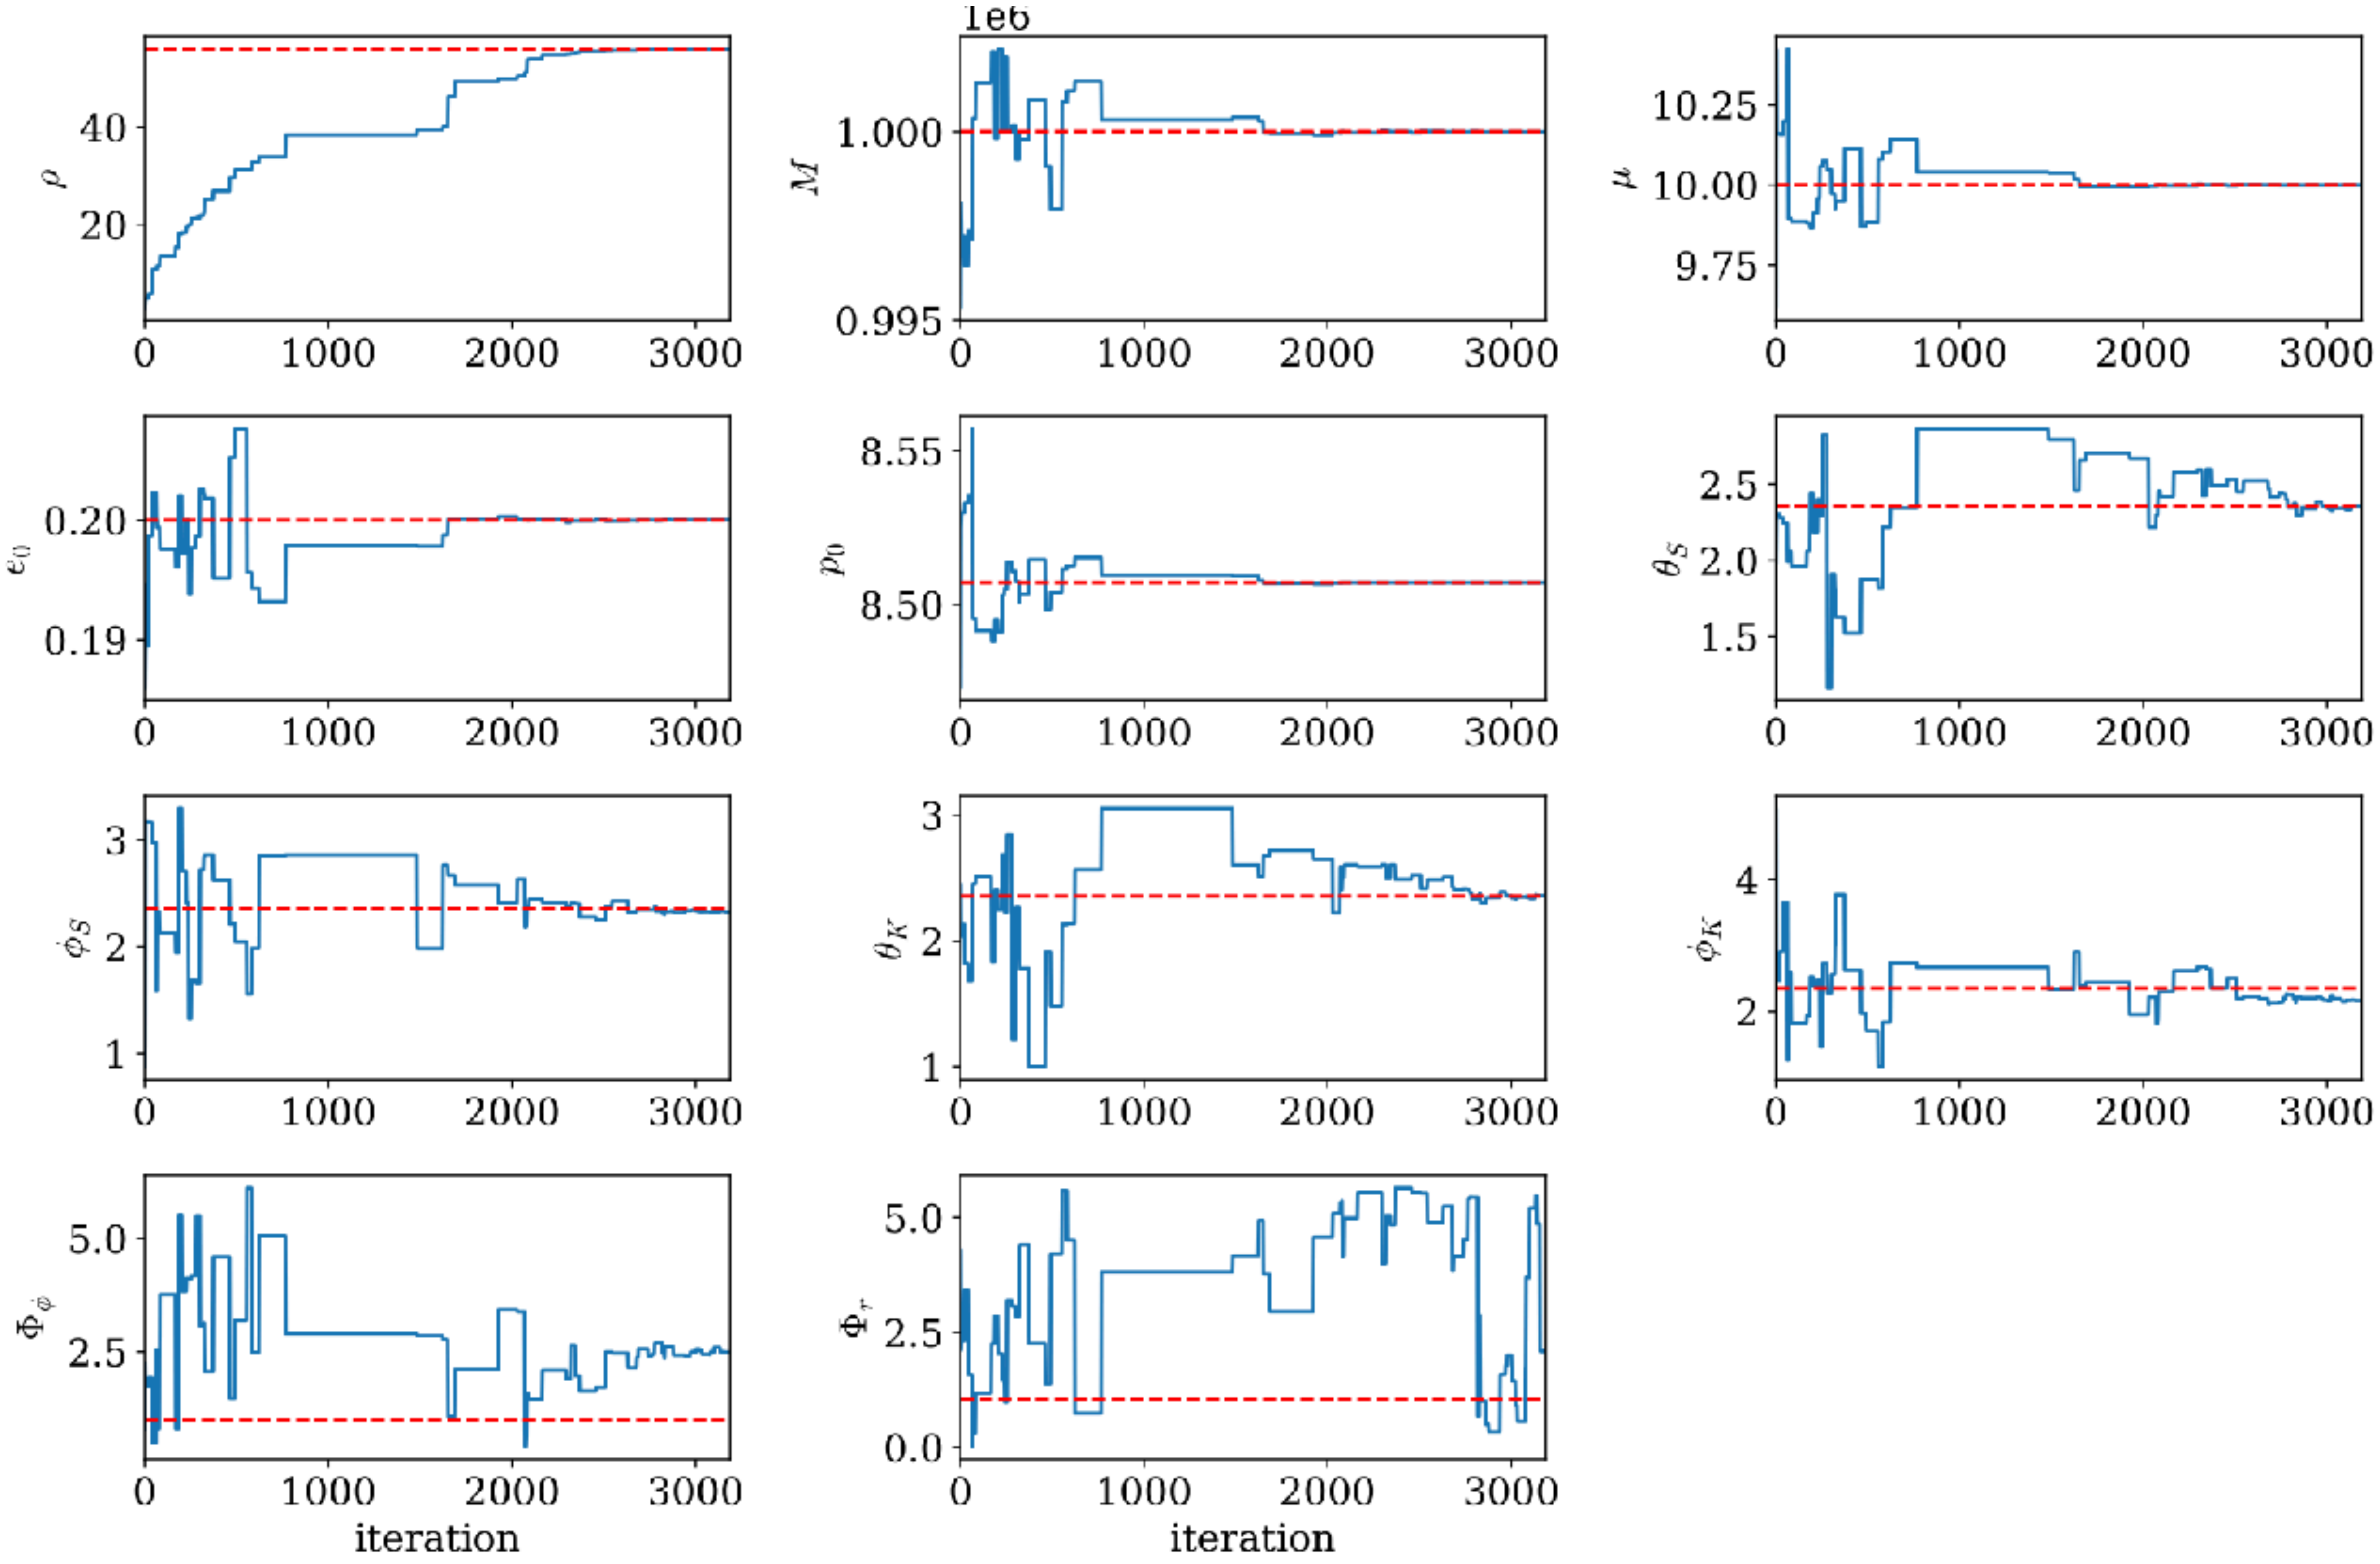

Strub et al. 2025

Narrowing down parameter space

- Works particularly well for intrinsic parameters
- Here larger primary mass, and even wider priors

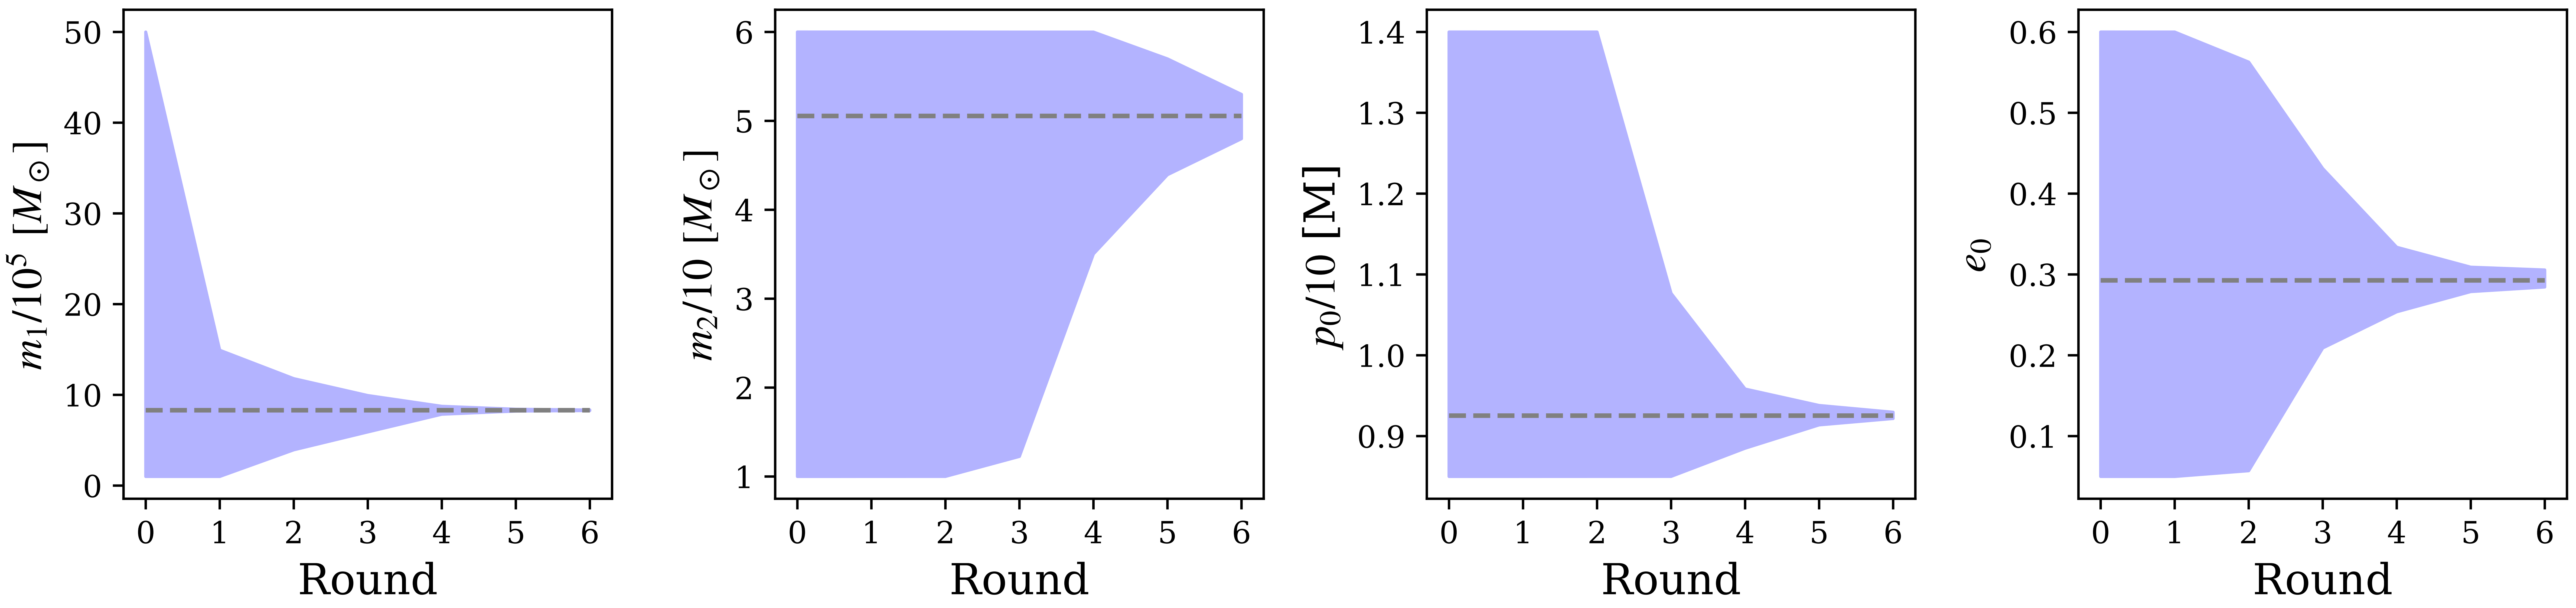

Training details

- Embedding network: linear transformation (Ax + b), normalisation across batch, ReLu ReLU(x)=max(0,x) with decreasing dimensions
- U-net (CNN): extracts both complex and simplistic features
- Stack the features from different channels
- Residual Net: Actual training of log-ratios, monitor binary cross-entropy loss function with AdamOptimizer

Training set-up

PEREGRINE-style approach

PEREGRINE
Gravitational Wave Parameter Inference with Neural Ratio Estimation

Bhardwaj et al. 2023
Miller et al. 2021

- 150K simulations per round
- Batch size = 128
- Initial learning rate = 10^{-4}
- Training:validation - 90:10
- Early stopping criterion: 7 epochs
- Utilise noise shuffling
- Bounding threshold = 10^{-5}
- Unet (CNN for compression -> reconstruction) -> Linear compression (MLP, each layer linear) -> Logratio estimator (ResNet): input 16 features, 11 parameters, dropout = 0.1
- O(20K) parameters in the logratio estimator (residual network)
- 2 residual blocks (4 hidden layers), 64 hidden features

2d marginals with TMNRE

- Intrinsic parameters consistent

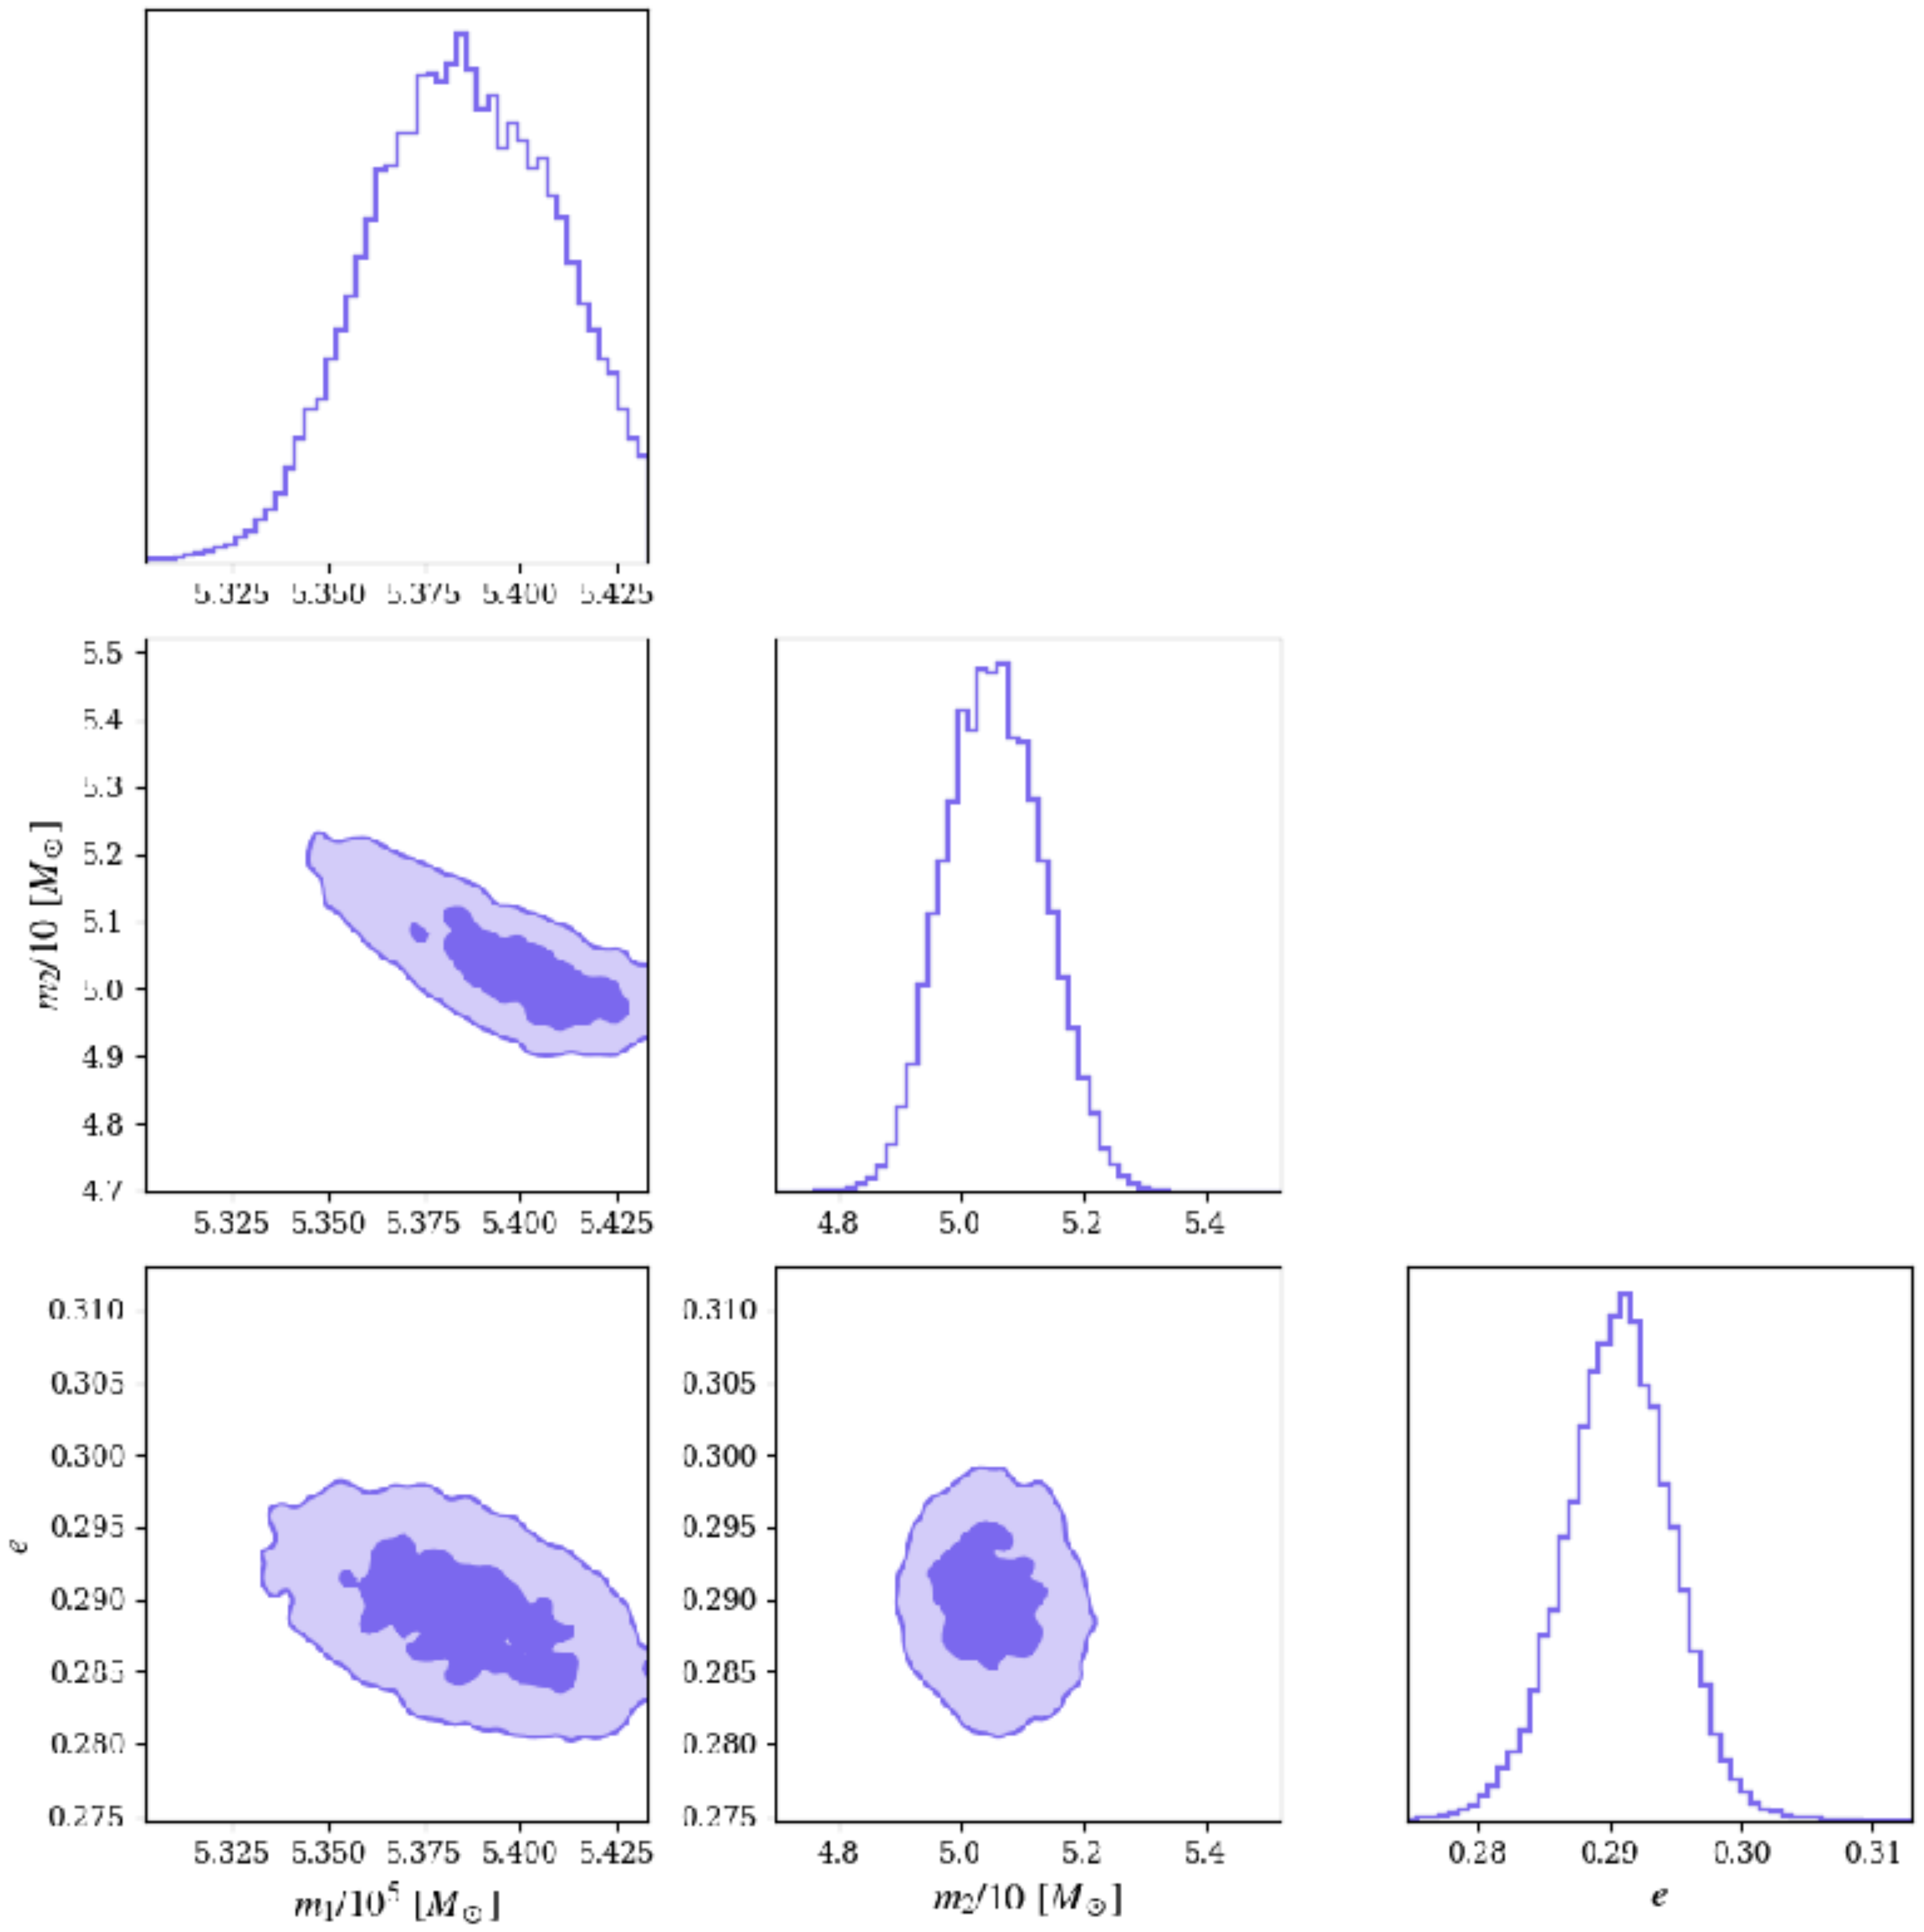

2d issues with TMNRE

- Some other 2d marginal distributions inconsistent with 1d distributions

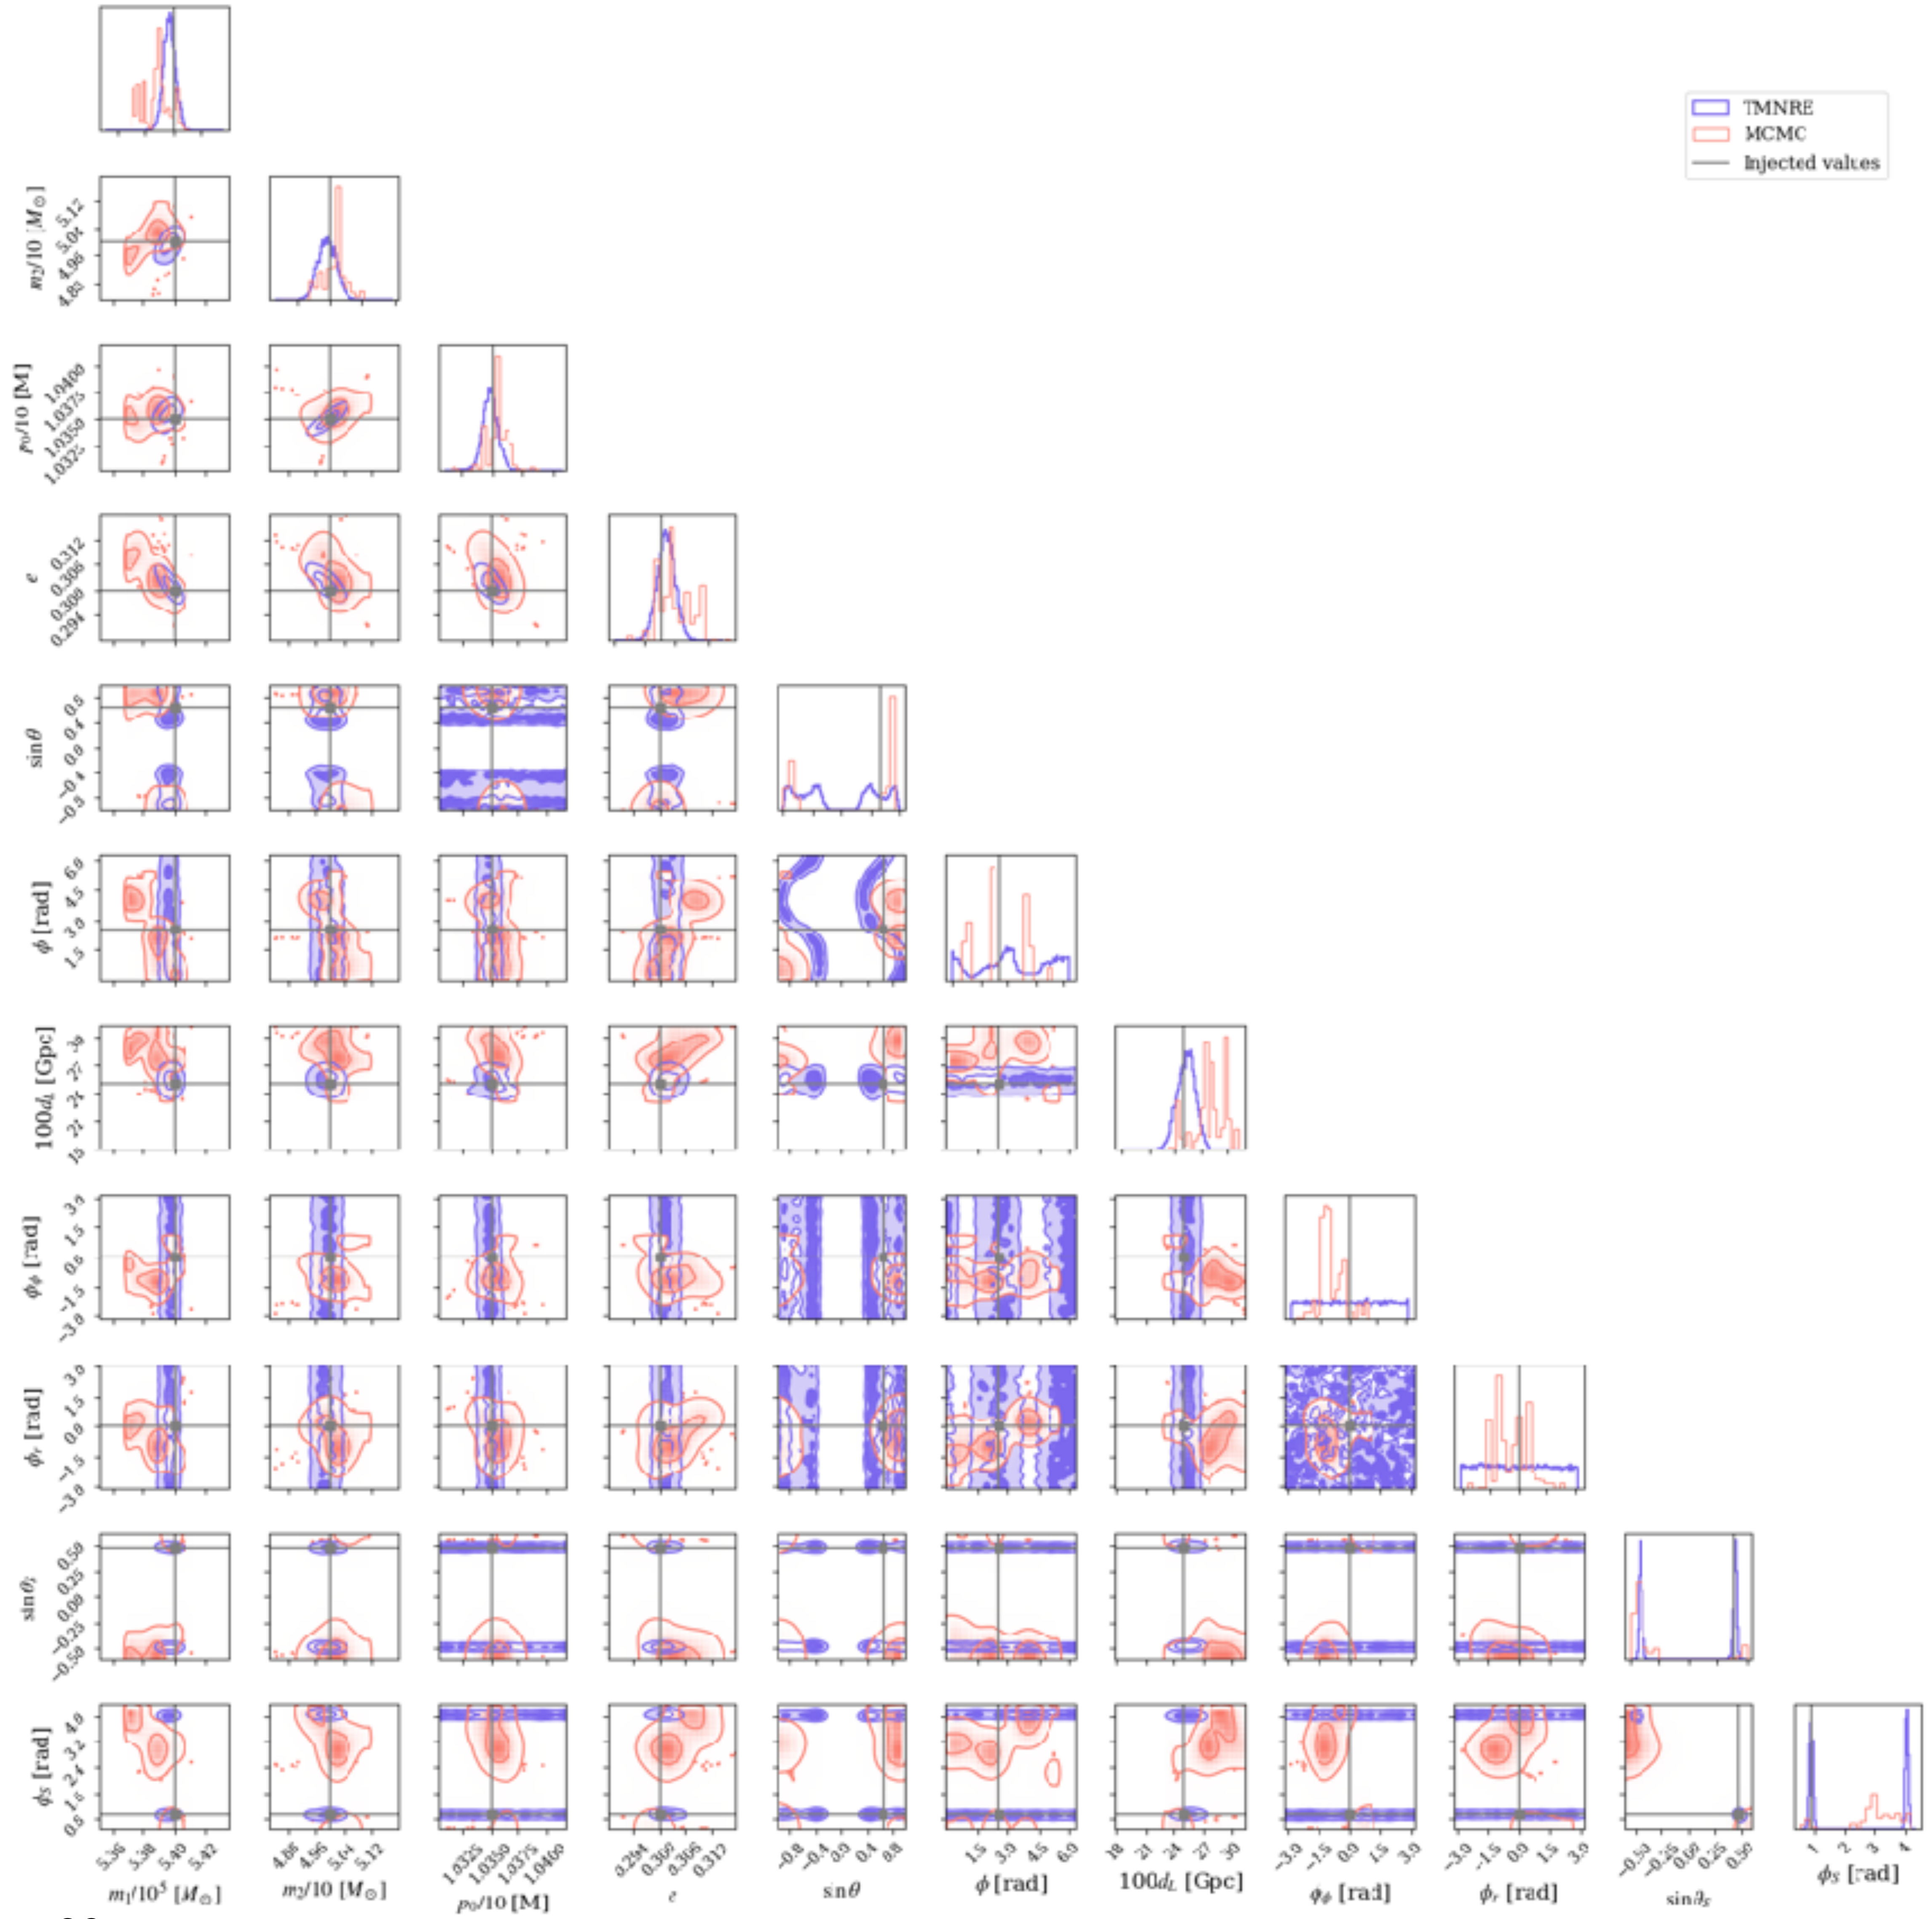

FEW details

- Phase and amplitude of each mode computed up to 1st order in gravitational self-force theory (expansion of the metric of the binary in powers of mass ratio).

- Modes summed over to produce adiabatic waveform $h(t) = h_+(t) - ih_\times(t)$ in time domain.