Prospects for inferring the population properties of compact binaries

Aditya Vijaykumar

The Future of Gravitational Wave Astronomy
28th October 2025

CITA
ICAT
Canadian Institute for Theoretical Astrophysics
L'institut Canadien d'astrophysique théorique

GW150914

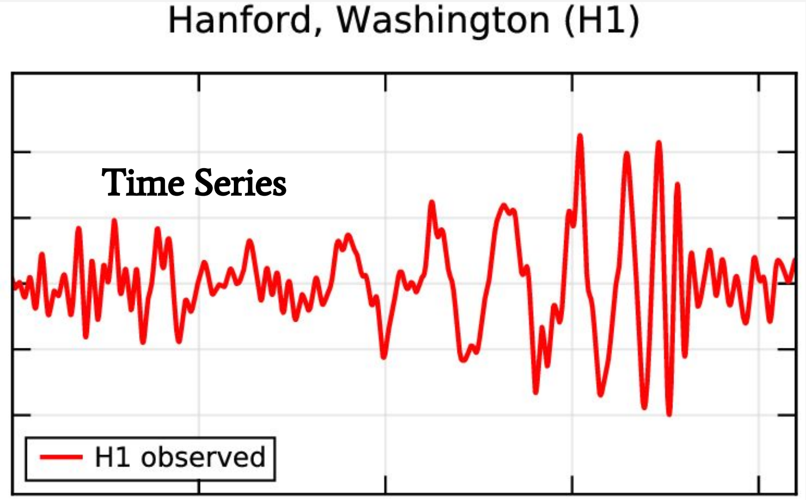

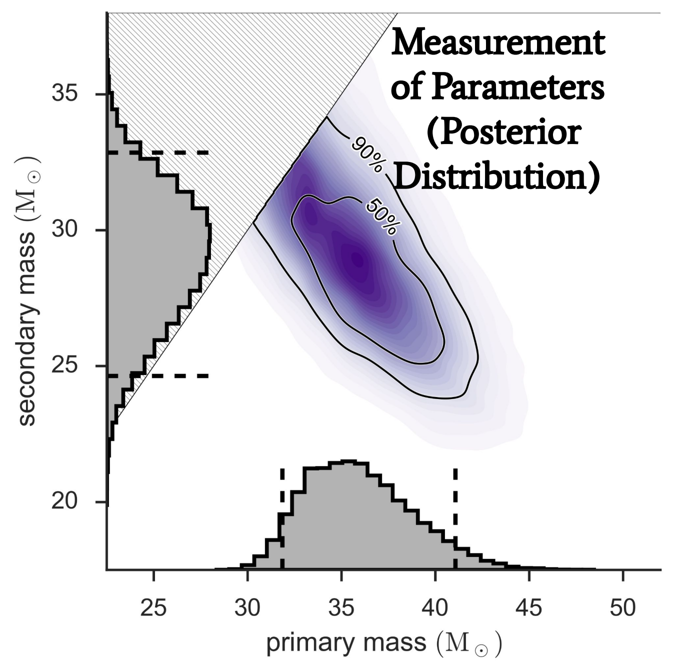

Abbott+ 2016, arXiv:1602.03837, arXiv:1602.03840

~200 events now!

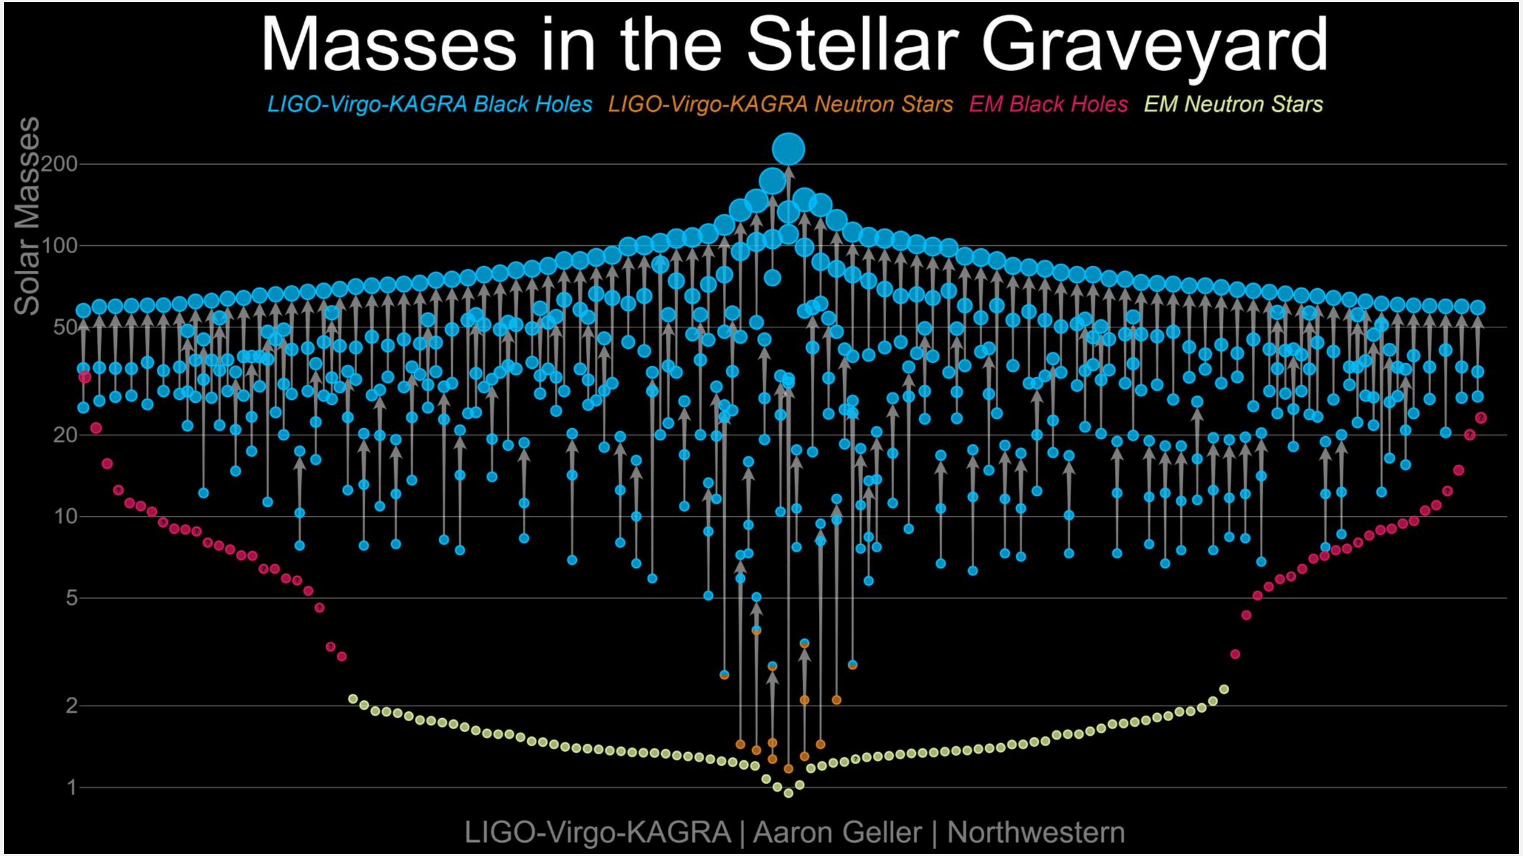

Population-level properties encode astrophysics

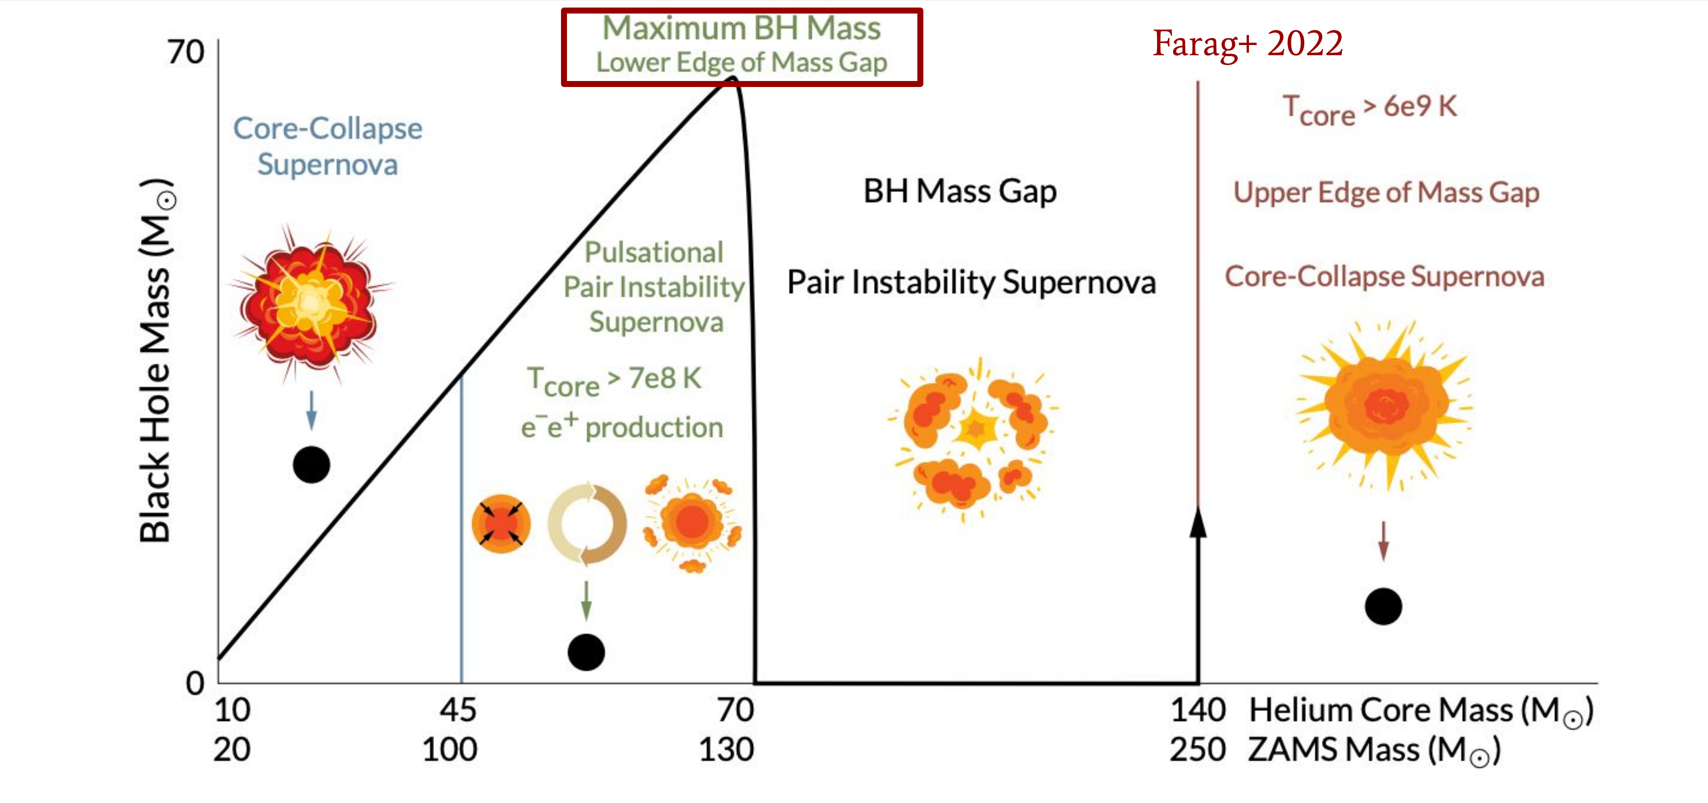

Population-level properties encode astrophysics

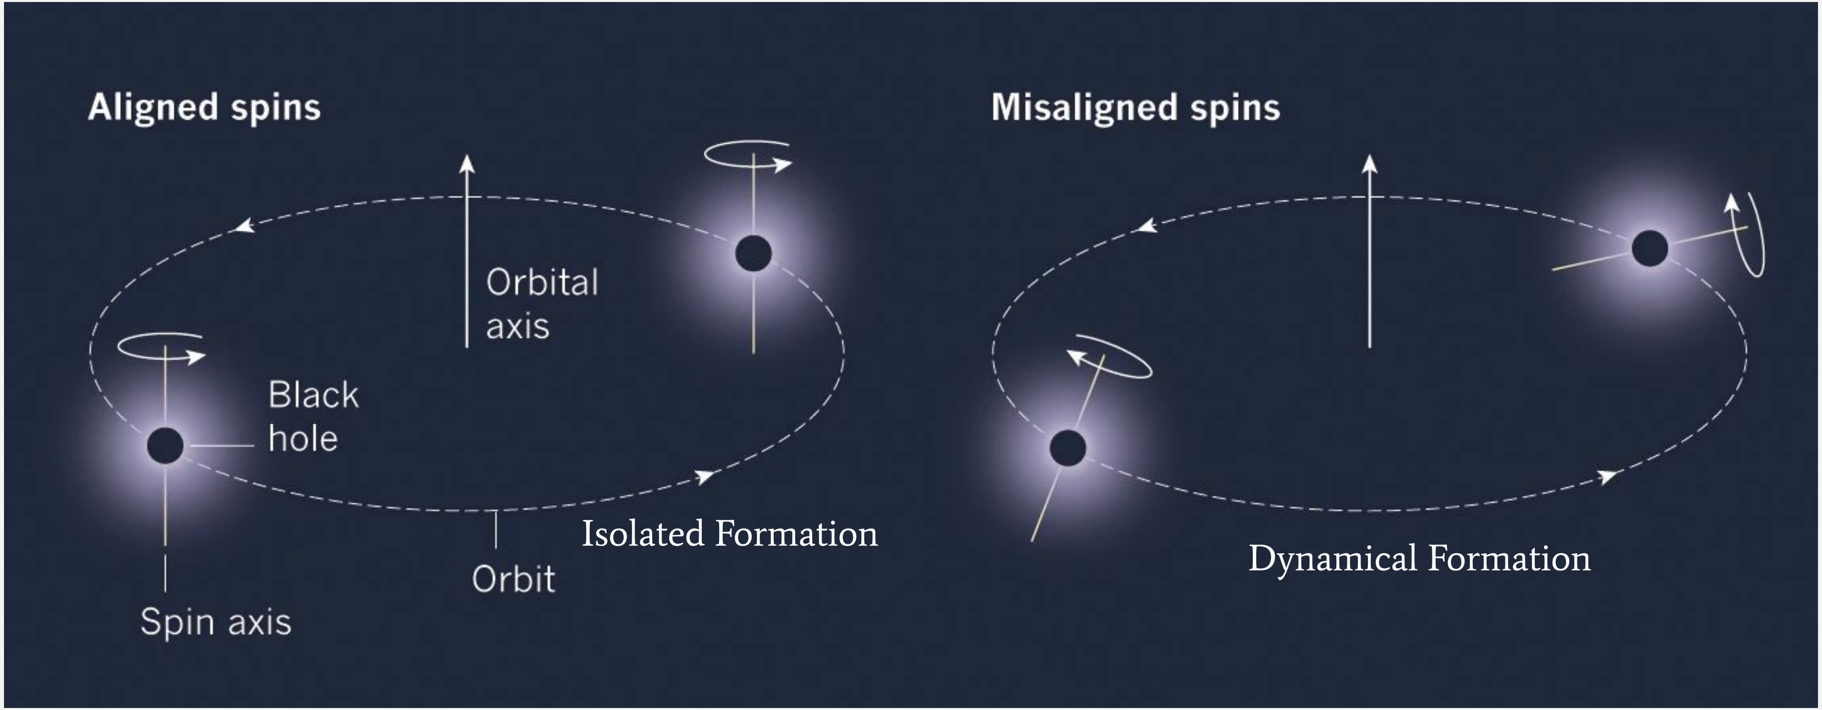

https://www.nature.com/articles/548397a/figures/1

The inference problem

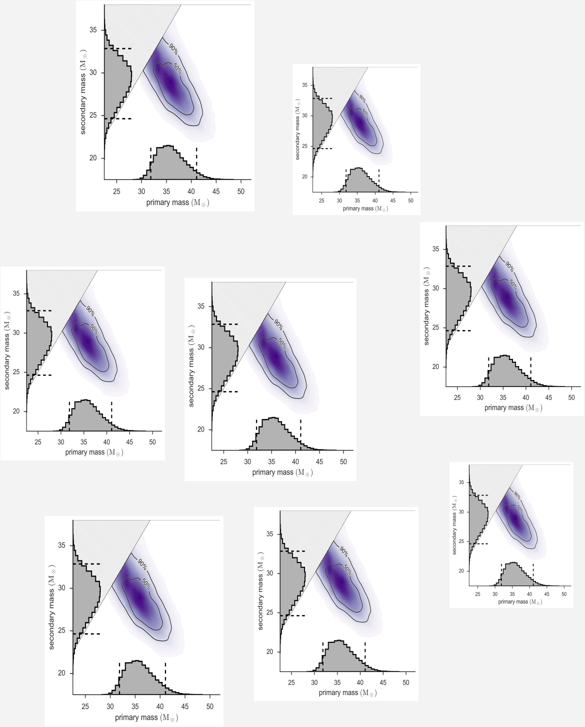

Many uncertain measurements from different sources

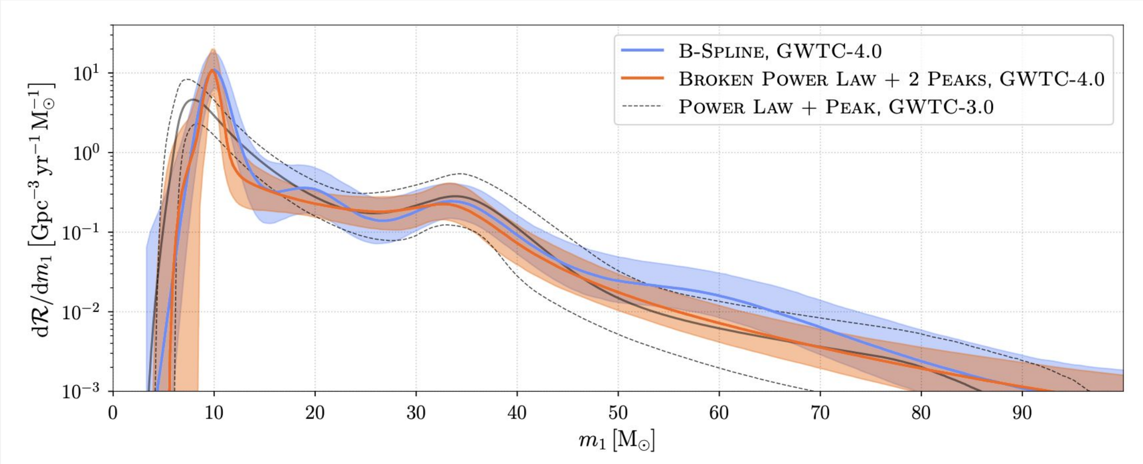

Measurements of population properties [LVK+, arXiv:2508.08083]

Selection Effects

$$A_{\rm GW} \propto \frac{\mathcal{M}^{5/6}}{D_L}$$

Inherently easier to detect massive binaries, and binaries that are close by.

...but we can correct for this!

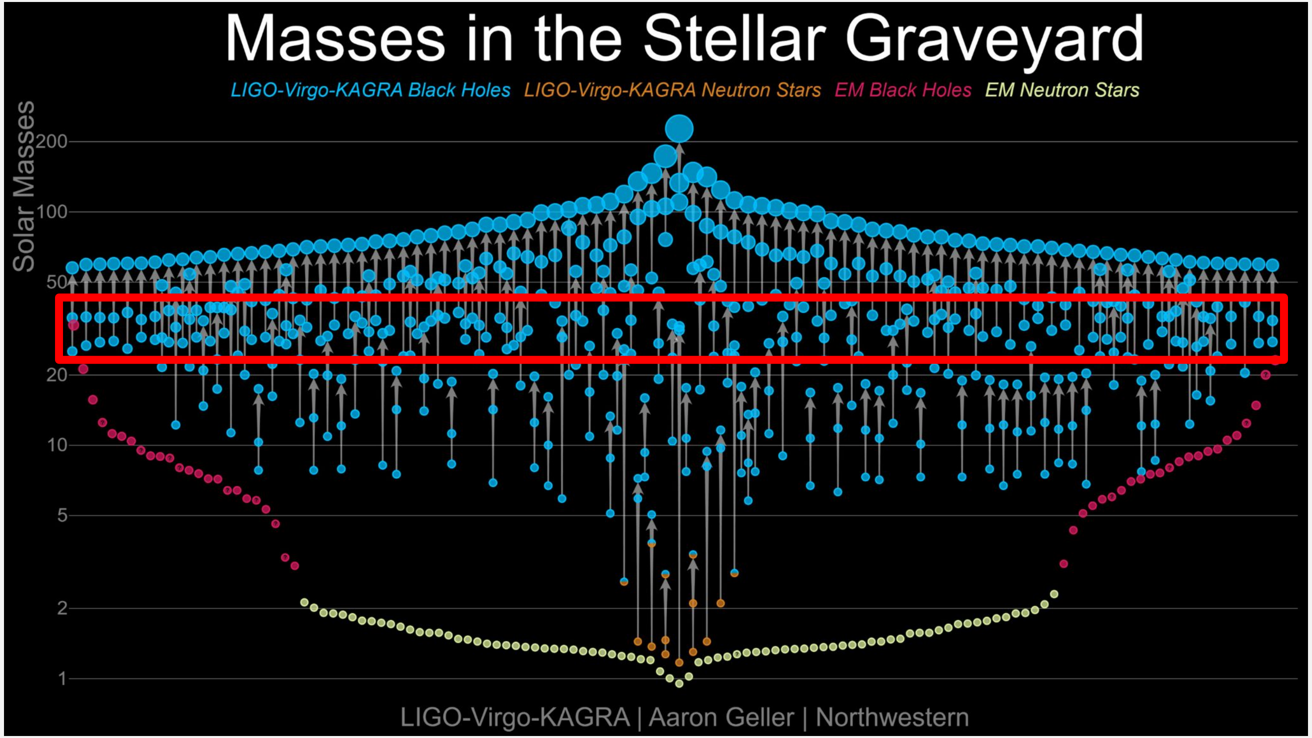

Population Model

Simple Model: The mass function of stars is a power law with exponent 2.3 with a M_{max}~300 Msun

What are the corresponding numbers for black holes?

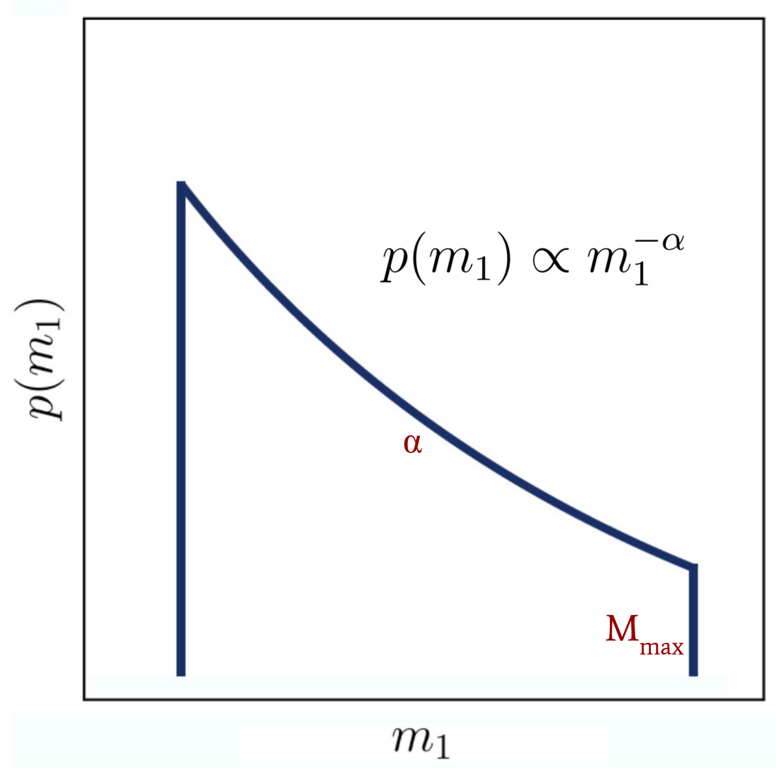

Abbott et al 2020 (GWTC-2 Population paper)

Population Model

$\Lambda = \{\alpha, M_{max}\}$

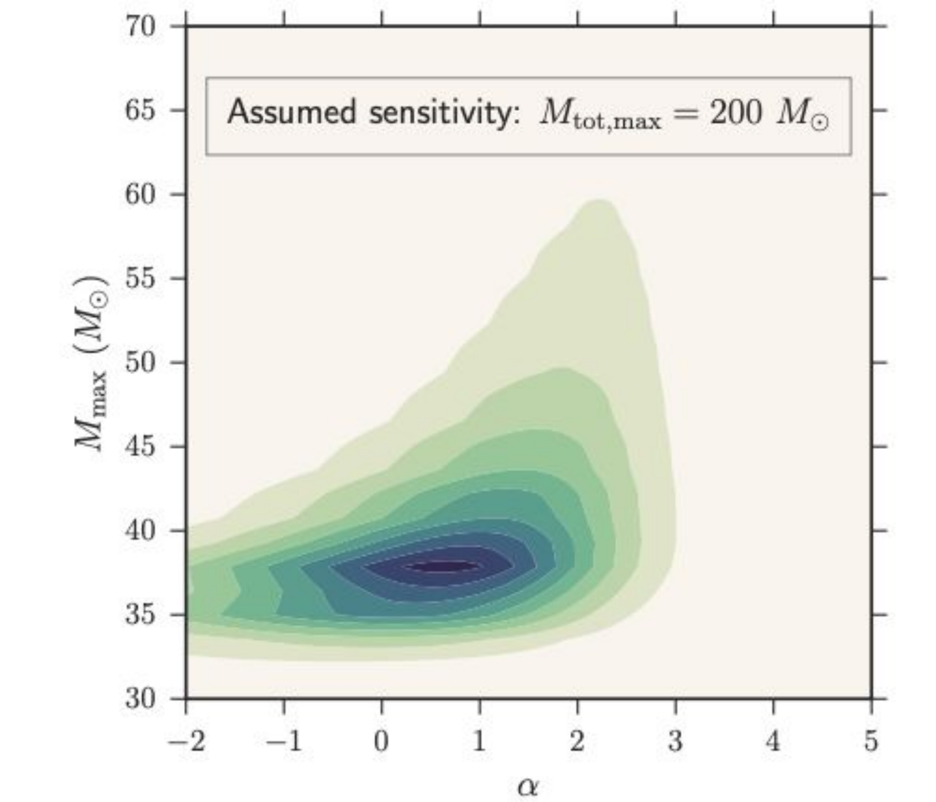

Fishbach and Holz 2017

Used only the first 6 reported GW detections

Mass Distribution

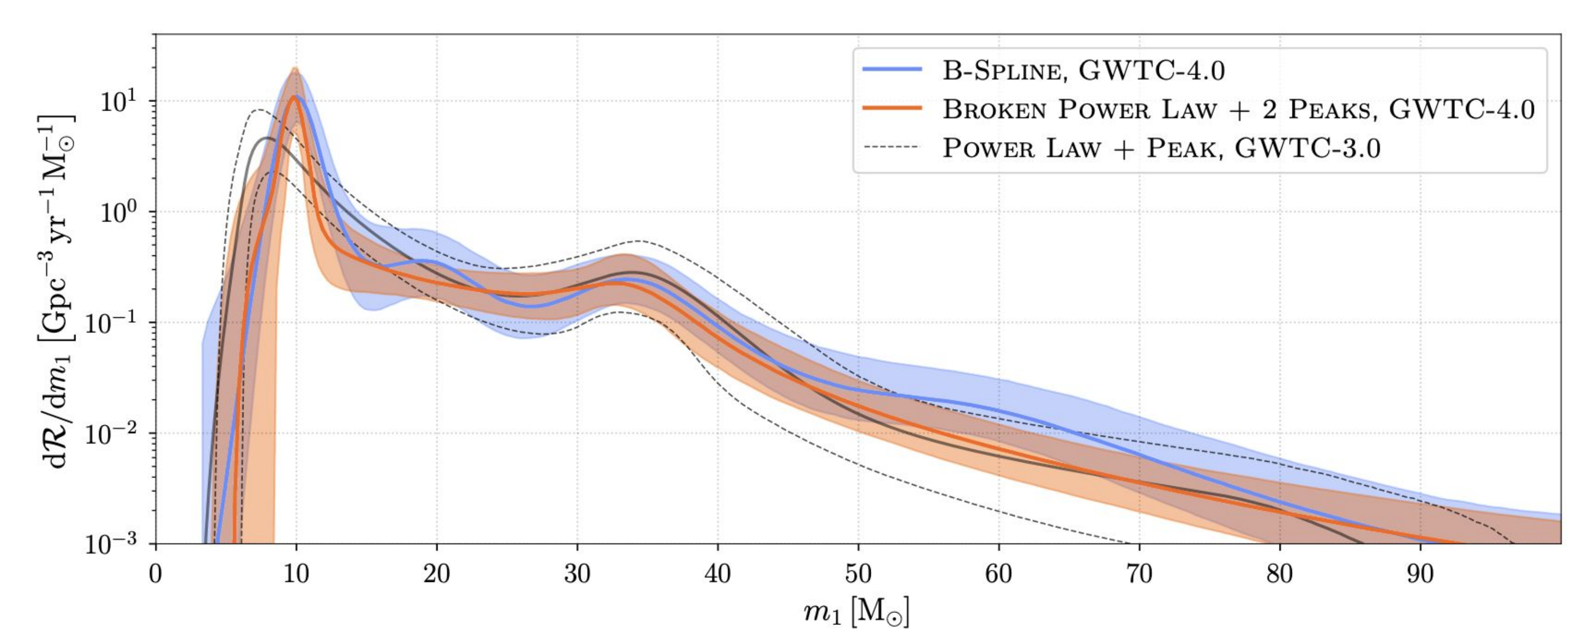

With more data, we now know that a simple power law is not a good fit to the data.

LVK+, arXiv:2508.08083

Effective Spin Distribution

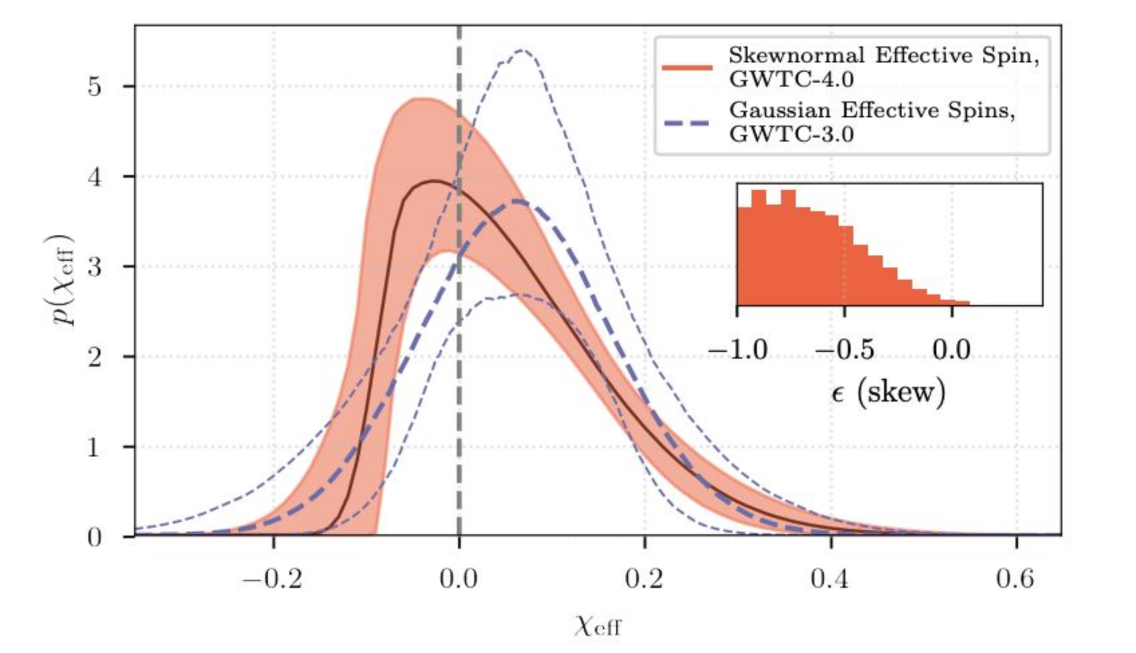

Distribution of effective spins skewed towards positive values, but significant support for negative spins.

LVK+, arXiv:2508.08083

Redshift Distribution

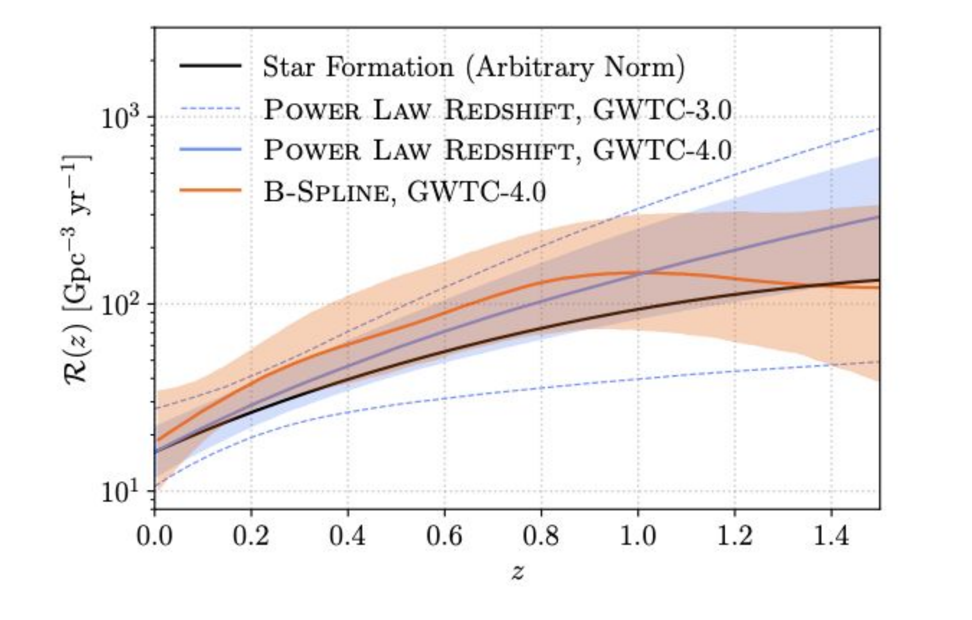

The rate of mergers per unit volume increases with redshift, consistent with star formation rate density.

LVK+, arXiv:2508.08083

Pair instability feature in the data?

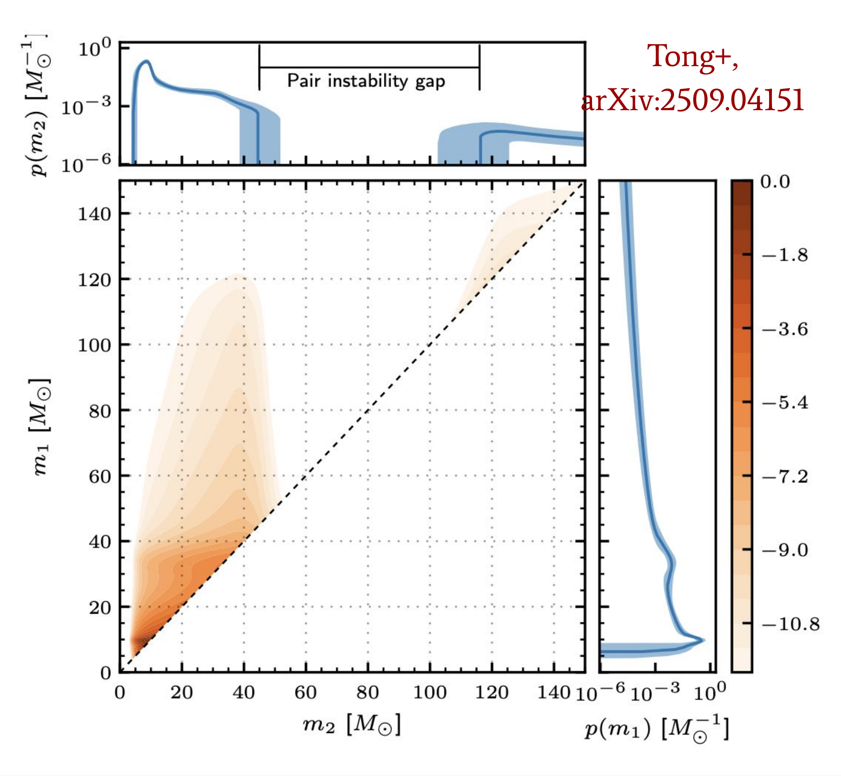

Tong+,
arXiv:2509.04151

The *secondary* mass distribution breaks after about 45 Msun, but there are primary masses above that value. Hierarchical mergers?

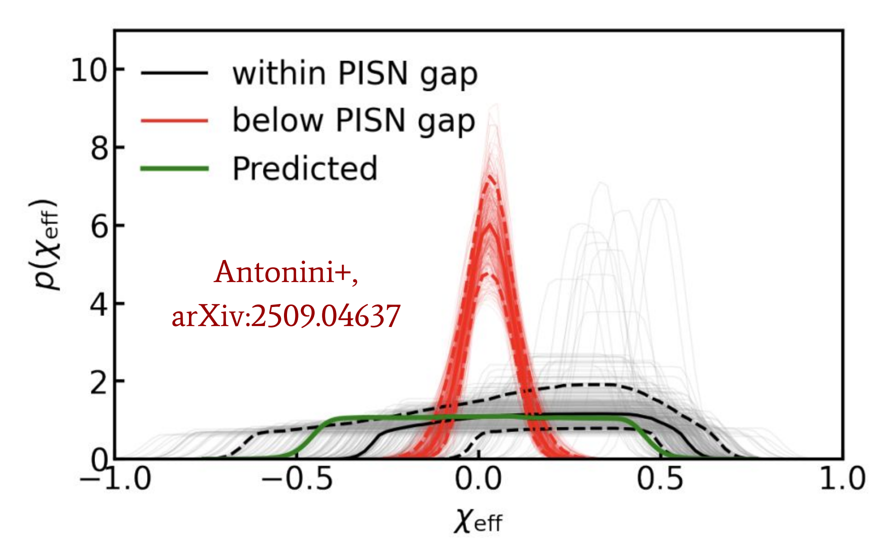

Antonini+,
arXiv:2509.04637

Challenges: What features should we look for in the data?

Difficult to disentangle formation channels: e.g., most formation channels can produce 30 Msun BHs. Need to look for robust features.

Hierarchical merger spins are a robust feature, peaking around 0.7.

More candidate features:

1. Do higher mass black holes have larger spins?
2. Does the 1G mass distribution shift to larger values as a function of redshift?
3. etc.

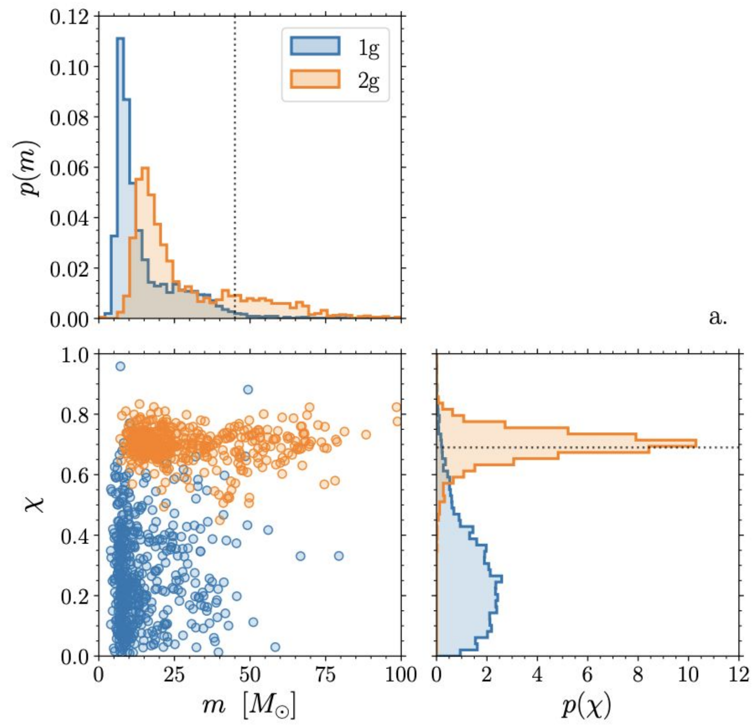

Gerosa & Fishbach 2021

Challenges: What features should we look for in the data?

Fraction of BBHs from Clusters

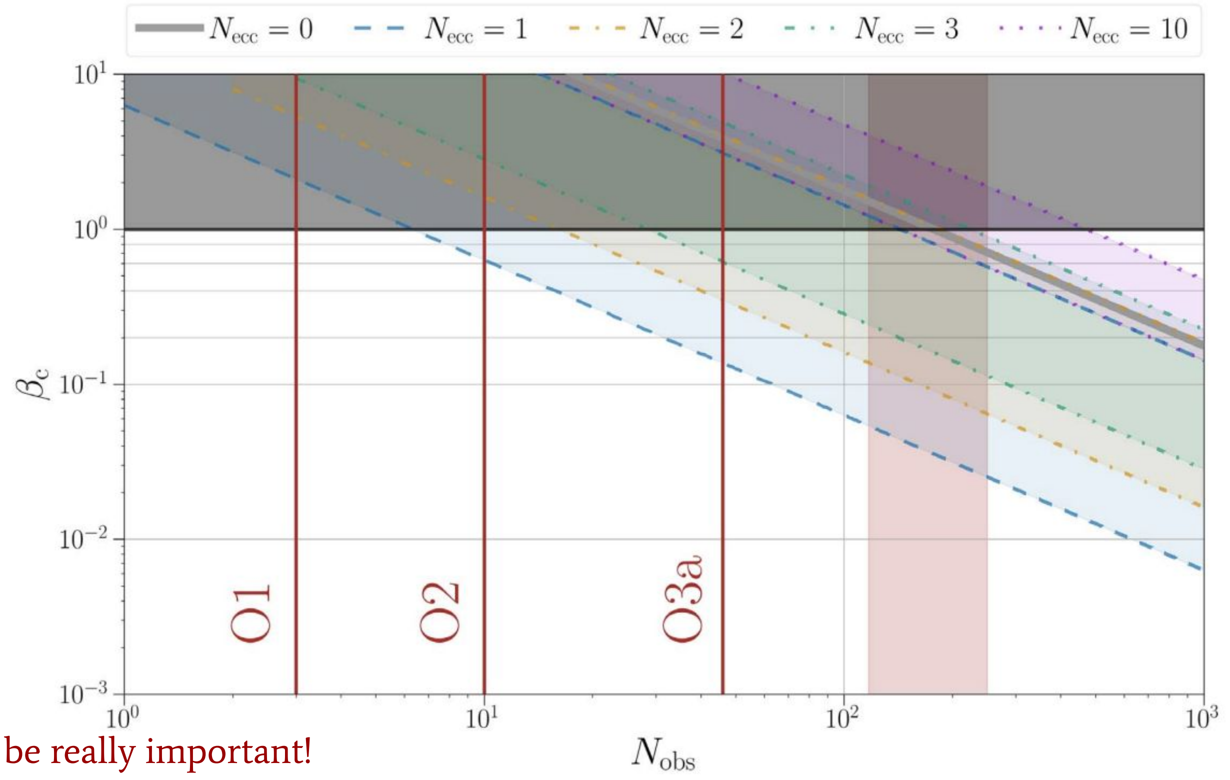

Zevin+ 2020

Eccentricity will be really important!

Challenges: How should we look for features in the data?

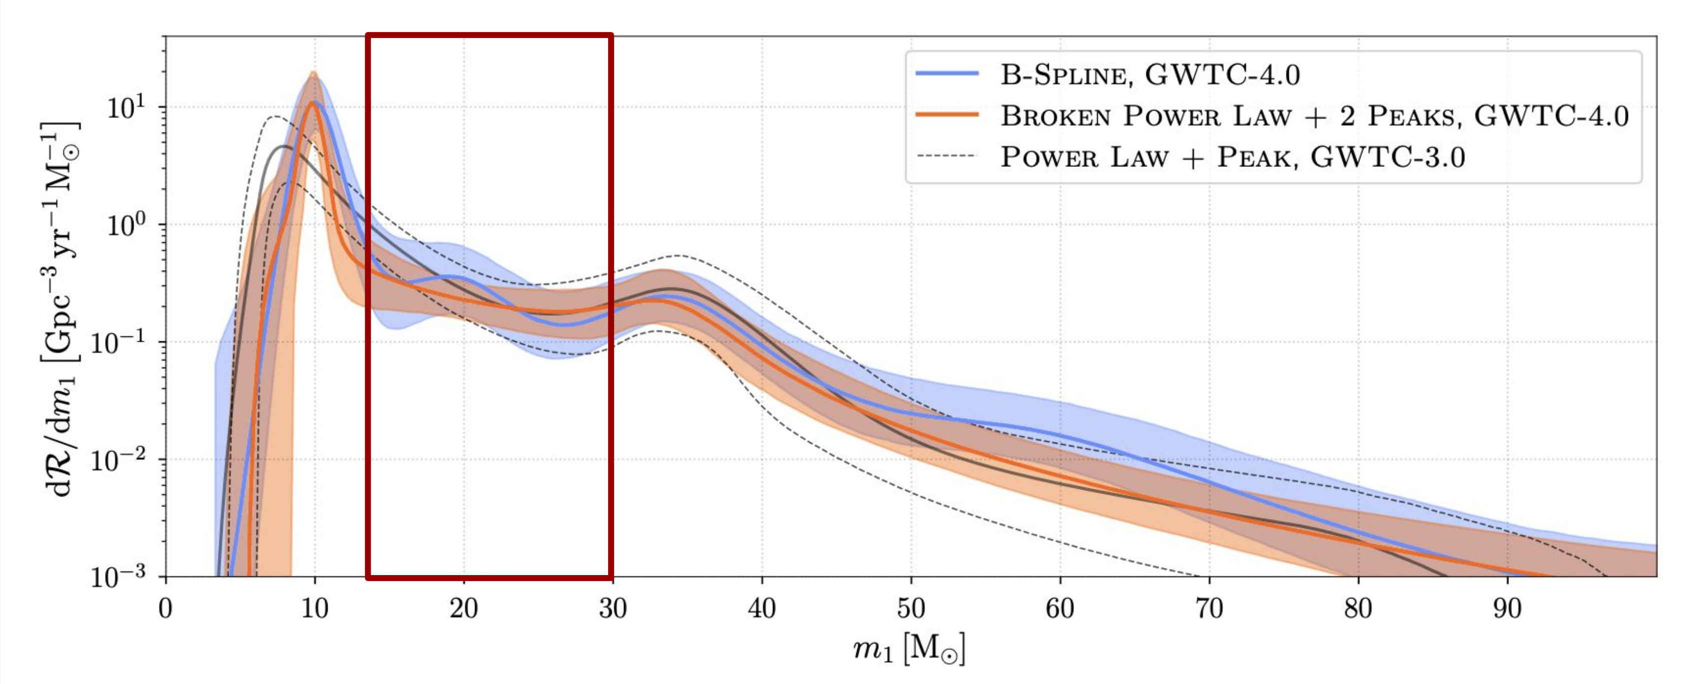

Flexible vs parametric models: Bias variance trade-off! Current flexible models are expensive in multiple dimensions. New methods?

Challenges: How should we do selection effects?

Compact Binary Coalescence Sensitivity Estimates with Injection Campaigns during the LIGO-Virgo-KAGRA Collaborations' Fourth Observing Run

Reed Essick,[1,2,3,a] Michael W. Coughlin,[4,5] Michael Zevin,[6,7,8] Deep Chatterjee,[9] Teagan A.

We describe the effort to characterize gravitational-wave searches and detector sensitivity to different types of compact binary coalescences during the LIGO-Virgo-KAGRA Collaborations' fourth observing run. We discuss the design requirements and example use cases for this data product, constructed from $> 4.33 \times 10^8$ injections during O4a alone. We also identify subtle effects with high confidence, like diurnal duty cycles within detectors. This paper accompanies a public data release of the curated injection set [1, 2], and the appendixes give detailed examples of how to use the publicly available data.

Selection effects are *expensive*, and requirements will only increase with time.

Summary

- Thinking beyond single-event parameters and viewing events as a population can help us constrain astrophysical properties

- Many challenges:
  - Modeling the population
  - Selection effects
  - Looking for robust features

- Synergies with other observational channels?
  - LISA?
  - Gaia BHs?
  - X-ray Binaries?

Thanks for Listening!

Get in Touch!

Website: adivijaykumar.github.io

Email: aditya@utoronto.ca

Extra Slides

Ingredients

$$p\left(\Lambda \middle| \{d\}\right) \propto p\left(\{d\} \middle| \Lambda\right) p(\Lambda)$$

Population parameters, could be:

Population Likelihood

Prior

1. Slope of the mass distribution
2. Minimum and maximum masses
3. Location of e.g. "peaks" in the mass distribution
4. Mean and variance of spin magnitudes
5. ...etc.

Population Likelihood

- The arrival of GWs at our detectors is a Poisson process, so we might think

$$p(N_{\rm ev}|N_{\rm exp}) = \frac{N_{\rm exp}^{N_{\rm ev}} e^{-N_{\rm exp}}}{N_{\rm ev}!}$$

N_{ev} is the number of events observed, and N_{exp} is the number of events expected given a population model parameterized by Λ.

Population Likelihood

- The arrival of GWs at our detectors is a Poisson process, so we might think

$$p(N_{\rm ev}|N_{\rm exp}) = \frac{N_{\rm exp}^{N_{\rm ev}} e^{-N_{\rm exp}}}{N_{\rm ev}!}$$

  N_{ev} is the number of events observed, and N_{exp} is the number of events expected given a population model parameterized by Λ.

- ...But it is *not* a homogeneous Poisson process: N_{exp} depends on the population model *and* the detector sensitivity to each member of the population
  - e.g. some population model might predict 5 Msun and 50 Msun BHs at the same rate, but selection effects mean we'll detect more 50 Msun BHs.

Population Likelihood

Assuming a inhomogeneous Poisson process, the population likelihood can be written as

Intrinsic rate

$$p(\{d\}|\{\theta\},\Lambda) \propto e^{-N_{\rm exp}(\Lambda)} \prod_i^{N_{\rm ev}} \int {\rm d}\theta_i \; p(d_i|\theta_i) \; [\mathcal{R}\, p(\theta_i|\Lambda)]$$

Single-event PE likelihood

Population model

For a full derivation, see Vitale et al 2020

$$N_{\rm exp}(\Lambda) = \int {\rm d}\theta \, {\rm d}t \; [\mathcal{R}\, p(\theta|\Lambda)] \; p({\rm det}|\theta)$$

Probability of detecting a signal with parameters θ. Is either 1 or 0.

Approximating the population likelihood

$$p(\{d\}|\{\theta\},\Lambda) \propto e^{-N_{\rm exp}(\Lambda)} \prod_{i}^{N_{\rm ev}} \int \mathrm{d}\theta_i \; p(d_i|\theta_i) \left[\mathcal{R}\, p(\theta_i|\Lambda)\right]$$

$$N_{\rm exp}(\Lambda) = \int \mathrm{d}\theta \, \mathrm{d}t \; \left[\mathcal{R}\, p(\theta|\Lambda)\right] \; p(\mathrm{det}|\theta)$$

Both these expressions involve integrals in high dimensions.
Solution: Convert integrals into Monte Carlo sums.

Approximating the population likelihood

$$\int \mathrm{d}x\, f(x)\, p(x) = \mathbb{E}_p\left[f(x)\right] = \frac{1}{N_{\mathrm{samps}}} \sum_{k=1}^{N_{\mathrm{samps}}} f(x_k)\Big|_{x_k \sim p(x)}$$

Integrals like these are expectation values of a function f(x) under some probability distribution p(x).

Can be approximated by drawing samples from p(x), and calculating an average of the f(x) over those samples

Approximating the population likelihood

$$N_{\mathrm{exp}}(\Lambda) = \int \mathrm{d}\theta\, \mathrm{d}t \left[\mathcal{R}\, p(\mathrm{det}|\theta) \frac{p(\theta|\Lambda)}{p_{\mathrm{draw}}(\theta|\Lambda_{\mathrm{draw}})}\right] p_{\mathrm{draw}}(\theta|\Lambda_{\mathrm{draw}})$$

$$N_{\mathrm{exp}}(\Lambda) = \frac{1}{N_{\mathrm{inj}}} \sum_{k=1}^{N_{\mathrm{found}}} \frac{p(\theta_k|\Lambda)}{p_{\mathrm{draw}}(\theta_k|\Lambda_{\mathrm{draw}})}$$

Choose a fiducial draw population distribution → simulate and inject event parameters → check if they are detected.

Approximate N_{exp} as a sum over “found” injections.

Approximating the population likelihood

$$\int d\theta_i \; p(d_i|\theta_i) \left[p(\theta_i|\Lambda) \right] = \int d\theta_i \; p(\theta_i|d_i) \left[\frac{p(\theta_i|\Lambda)}{\pi_{\rm PE}(\theta_i)} \right] = \frac{1}{N^i_{\rm samps}} \sum_{k=1}^{N^i_{\rm samps}} \left[\frac{p(\theta_i^k|\Lambda)}{\pi_{\rm PE}(\theta_i^k)} \right]$$

Similarly, the integral over individual event likelihoods can be approximated using single event posterior samples.

Approximating the population likelihood

We now have everything to calculate the likelihood!

$$p(\{d\}|\{\theta\}, \Lambda) \propto e^{-N_{\rm exp}(\Lambda)} \prod_{i}^{N_{\rm ev}} \int \mathrm{d}\theta_i \; p(d_i|\theta_i) \left[\mathcal{R}\, p(\theta_i|\Lambda)\right]$$

Approximate using simulated "injections"

Approximate using single-event PE samples

However, Monte Carlo sums have uncertainties, which become important especially when the population model has sharp features [Essick and Farr 2022, Talbot and Golomb 2023].

Typical strategy: set likelihood to 0 when variance in the Monte Carlo Estimator is above a threshold.