

Analyzing the Effective and Component Spin Distributions of Binary Black Hole Mergers

ZOE KO, SIMONA MILLER, KATERINA CHATZIOANNOU

ABSTRACT

Gravitational wave (GW) observations of binary black hole (BBH) mergers provide measurements of BBH parameters such as mass and spin, which shed light on the evolutionary history of these systems. We explore the distribution of BBH spin on a population level through looking at different spin parametrizations. LIGO data currently provides strong constraints on effective spin, a mass weighted average of component spin projected in the direction of the angular momentum, while the component spins are weakly constrained for individual events. Through a Bayesian hierarchical inference approach, we explore whether we can differentiate synthetic populations with the same effective spin distributions but different component spin distributions. We explore these spin distributions for two different population sizes: one comparable to LIGO’s third observing run (O3) and one comparable to the predicted population size after the next observing run (O4). This study demonstrates how current models of spin distribution will improve with O4.

1. INTRODUCTION

LIGO made the first observation of a GW signal from a BBH merger on September 14, 2015 (Abbott et al. 2016a,b), marking the beginning of an era of discovery. Since then, the LIGO and Virgo Collaborations have confirmed the discovery of many tens of BBHs (Schmidt 2020). Each of these systems can be characterized by 15 parameters: so-called intrinsic parameters like the masses and spins of the two black holes, and extrinsic parameters like sky position and orientation of the binary. In this study, we focus on intrinsic parameters, which capture the formation and evolutionary history of BBH mergers. Specifically, we are interested in analyzing the spin parameters, using Bayesian inference techniques to carry out parameter estimation.

The spins of black holes in binary systems can be parameterized in two ways which are useful in the context of GW data: component spins and effective spin. Component spins are the individual spins of each black hole in the binary. The component spin parameterization includes spin magnitudes and tilt angles, as shown in Figure 1. The spin magnitude χ_i and the cosine of the tilt angle $\cos \theta_i$ can be extracted from LIGO data, but with current sensitivity they are poorly measured, as indicated by wide posterior distributions over the full allowed range of values. This makes it difficult to precisely extract specific magnitudes and the tilt angles. However, effective spin, χ_{eff} , is a well measured parameter in LIGO data that shows up at leading order in gravitational waveforms. Effective spin, is an average of the component spins in the direction of the angular momentum, weighted by each black hole’s mass:

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

Using effective spin and spin precession, the individual component spin parameters can be calculated. Spin precession, which also shows up at leading order, measures the amount of spin perpendicular to the direction of the angular momentum.

$$\chi_p = \max(\chi_1 \sin \theta_1, \frac{4q + 3}{4 + 3q} q \chi_2 \theta_2)$$

Figure 1 provides a visual for the component spins of a BBH.

Although effective spin is the best measured spin parameter, the component spins of BBH mergers are ultimately what we aspire to measure, as they provide strong insight into these system’s evolutionary origins. BBH mergers primarily form through two different channels: an isolated formation channel and a dynamical formation channel (Mapelli 2020). In the isolated formation channel, the system starts out as a binary with two Main Sequence (MS) stars. After its MS phase, one of the stars expands, allowing mass transfer to occur. The binary is then able to share a common envelope after that star collapses into a black hole. There is a strong gas drag from the shared envelope,

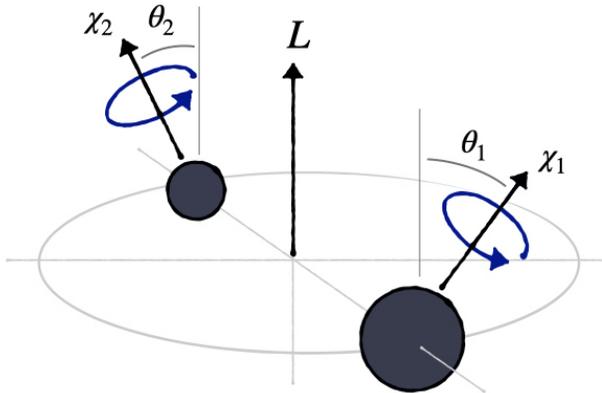


Figure 1. The component spins of a BBH. The effective spin is the mass-weighted average of component spins in the direction of angular momentum (L).

which makes the two bodies spiral into each other, transforming kinetic energy into thermal energy. This thermal energy heating up the envelope can lead to its ejection, finally allowing the system to evolve into a BBH merger. BBH mergers that formed through the isolated evolution channel are commonly believed to have spins that line up with the axis of the orbit. On the other hand, BBH mergers can form through a dynamical formation channel, which also starts as a binary with two MS stars. However, instead of evolving in an isolated environment to its final stage of a BBH, this system experiences a third-body encounter. One of the members of the binary is replaced by a stellar object in an environment such as a stellar cluster. Because this third body enters at a later time, BBHs that are formed dynamically have spins with random orientations. Therefore, gathering information on spin magnitudes and tilt angles of BBH systems can provide insight into these formation channels.

However, because of the difficulty of constraining component spin parameters, effective spin is more commonly used. BBHs formed in isolation have spins that are aligned, and therefore, the distribution of effective spins should be centered around a positive value. On the other hand, BBHs formed dynamically have spins that are randomly oriented, and therefore, the distribution of effective spins should be centered around zero. Using the data from the first and second runs of LIGO and Virgo, the effective spin of BBH systems was found to be very small, with the mean of the effective spin distribution, $\mu \sim 0$ with a narrow distribution (Miller et al. 2020). Miller et al. (2020) makes three hypotheses on the component spins based on the near-zero effective spin: the component spins are generally perpendicular to the binary’s orbital angular momentum, the component spins are generally anti-aligned, or the component spins are simply very small (Miller et al. 2020). While no conclusions were drawn regarding rates of isolated formation versus dynamical formation, the data points towards the existence of dynamically formed BBHs.

Using LIGO’s third observing run, Abbott et al. (2021) updated the posterior distributions of χ_{eff} and χ_p , finding χ_{eff} to be centered around 0.06, suggesting that spin-tilt misalignments do not fully cancel out. They found χ_p to be either centered around 0 with a broad distribution or centered around 0.2 with a narrow distribution.

While there is ongoing work studying the effective spin distribution in BBH mergers, analyses for component spin distributions remain largely inconclusive. This largely unexplored parameter space has the potential to reveal formation mechanisms in BBH systems.

Although component spin parameters at the individual event level are poorly measured, we currently have order(70) events and more to come in LIGO’s fourth observing run. Using existing data from past runs, we want to judge how informative the LIGO measurements are when analyzed collectively and see how these measurements could improve as we approach LIGO’s fourth observing run. Specifically, we aim to answer the question: if we simulate three different populations of BBH mergers with the same effective spin distribution but different component spin distributions, will we be able to measure that they are different populations?

2. METHODS

Existing data on BBH mergers that have been detected by LIGO yield relatively uninformative measurements of component spins, as indicated by wide posterior distributions on spin magnitudes and tilt angles (see Figure 3). We aimed to assess whether we can improve these measurements by conducting Bayesian hierarchical inference to obtain

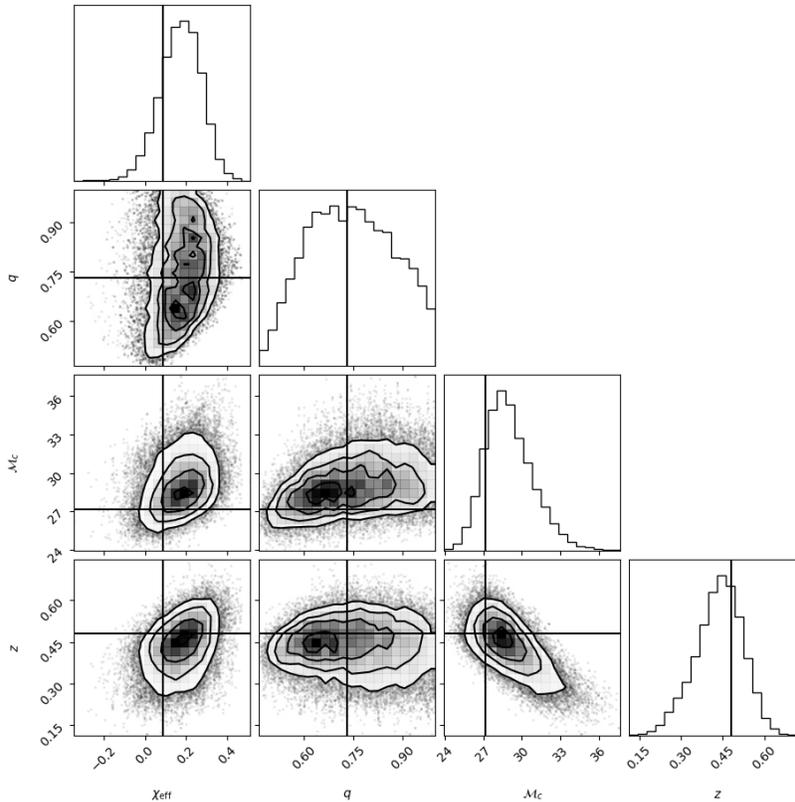


Figure 2. Example of an individual event posterior distribution showing leading order parameters for a synthetic BBH merger event. The black lines show the injected values.

population level posteriors on χ_i and $\cos\theta_i$. In addition to measuring the distribution of component spins across the BBH population, these population level posteriors can also be used to update the individual event posteriors. To do this, we simulated synthetic BBH merger populations, analyzed the effective and component spin distributions, and drew conclusions on the effectiveness of analyzing component spin distributions at the population level. We also experimented with different population sizes, running the signal injection and recovery using an synthetic dataset containing 70 events, 150 events, and then 300. This allowed us to observe the improvement in posteriors from using the current number of events we have to the estimated number of events we will have halfway through and fully after the next observing run. The overall method can be broken down into three steps: 1. Generate and inject synthetic BBH merger events into LIGO data. 2. Run parameter estimation for the individual events. 3. Conduct hierarchical inference at the population level.

2.1. Simulating and Injecting Synthetic BBH Events

We simulated three synthetic populations of BBH merger events with the same effective spin distribution but different spin precession distributions, which yields different component spin distributions. Effective spin measures how much total spin is in the direction of the angular momentum, while spin precession measures how much total spin is in the direction perpendicular to the angular momentum. GW data captures these two spin parameters to then derive component spin. Each population contains 50000 events, and the distribution of their spin parameters can be seen in Figure 4. This code was written by Simona Miller and is largely based off Callister et al. (2021).

After generating these three populations, we randomly drew n samples from each population and injected them into LIGO data, resulting in three mock catalogs, each of size n , containing synthetic BBH merger events. We ran the entire signal injection and recovery procedure with $n = 70$ to observe the posteriors for a catalog analogous to that of LIGO’s GWTC-3 and $n = 150$ and $n = 300$ to observe how the posteriors will improve once we obtain more data with O4.

2.2. Individual Parameter Estimation

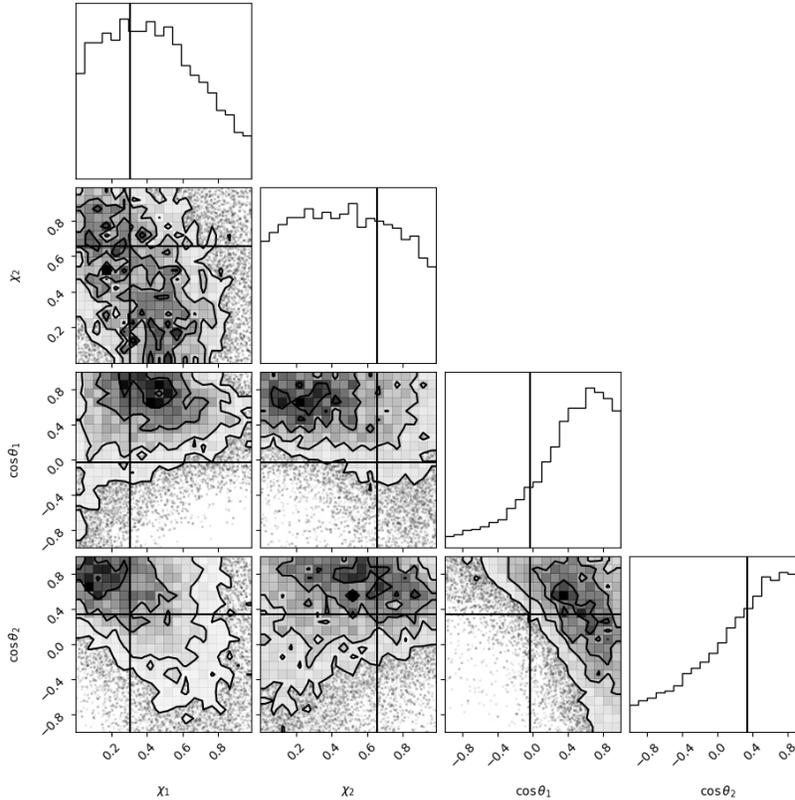


Figure 3. Example of an individual event posterior distribution showing wide posteriors for χ_i and $\cos \theta$ for a synthetic BBH merger event. The black lines show the injected values.

After generating three mock populations for a given population size n , we ran parameter estimation for each individual event in each population, using `bilby`, a Bayesian Statistics library designed for parameter estimation (Ashton et al. 2019). The `bilby` output yields posteriors on the leading order parameters for each GW event: effective spin, mass ratio, chirp mass, and redshift (see Figure 2). It also yields poorly constrained posteriors on component spin parameters (see Figure 3).

2.3. Hierarchical Inference: Population Level Parameter Estimation

Using the individual posteriors, we ran a Bayesian hierarchical modeling procedure to analyze the summation of these posteriors across the three different populations. This method takes in the `bilby` outputs and runs parameter estimation on the hyperparameters describing distributions in the overall population. If we are able to recover the original distributions shown in Figure 4, we can conclude that we are able to differentiate between BBH mergers with different component spin distributions, whereas if we are unable to differentiate between the three samples, we can conclude that we are unable to differentiate between BBH merger populations with different component spin distributions at current LIGO sensitivity.

We explore various underlying models to describe the component spin magnitude and tilt angle distributions, including Gaussian distributions, and binned distribution. We also explore modeling the spin magnitudes as a beta distribution and the cosine spin tilt angles as a mixture of aligned and isotropic subpopulations with and without a lower truncation bound, following the work of Collaboration et al. (2021); Callister et al. (2022). We are interested in the truncated model for spin tilt to explore whether or not there is a cutoff in the tilt distribution. If there is no truncation, the distribution is bounded by $\cos \theta = -1$ (spin aligned with angular momentum) and $\cos \theta = 1$ (spin aligned with angular momentum). However, if the lower bound is not consistent with -1, this would imply that spin tilt angles prefer aligned orientations. For example, if the truncation bound is 0, this would indicate that the data contains no anti-aligned systems.

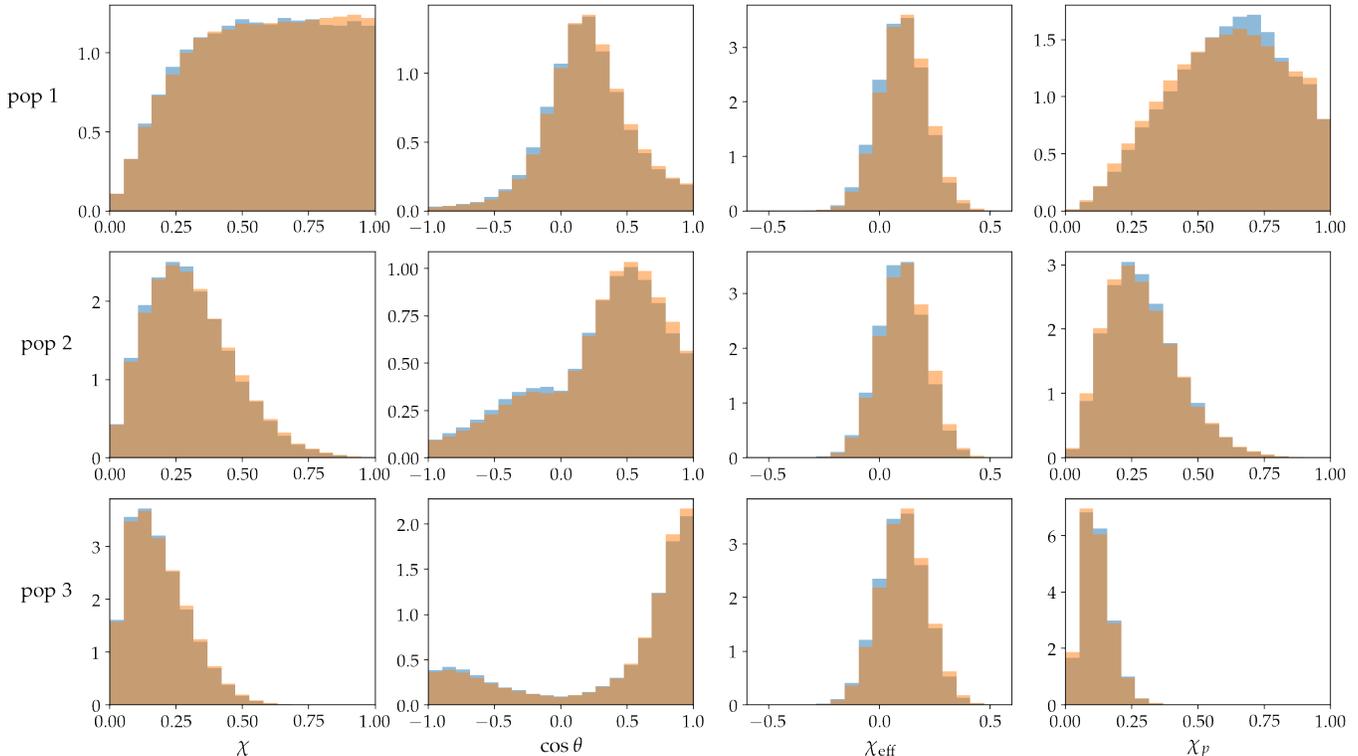


Figure 4. Three mock populations of BBH merger events with the same χ_{eff} distribution but different χ_p distribution, resulting in different component spin distributions. Population 1 has high amounts of spin precession, Population 2 has medium amounts of precession, and Population 3 has minimal precession. This results in three different component spin magnitude and tilt angle distributions. The blue histograms are the synthetic underlying distributions, while the orange histograms are the synthetic detected populations, which take into account selection effects. The similarity between the orange and blue histograms show that the spin of these populations do not contribute much to the detectability of a system. While this holds true for spin parameters, other parameters such as mass and redshift substantially impact selection effects.

After choosing a specific model for the hierarchical inference procedure, we followed the tutorial from Callister et al. (2022), running an MCMC sampler to explore the parameters. To inspect results, we first looked at the chains to confirm that all the walkers converged (see Figure 5). We then look at the corner plot of all the parameters of interest (see Figure 6). Finally, we make a trace plot corresponding to the spin magnitudes and tilt angles from Figure 6 (see Figure 7).

3. RESULTS

We explored a total of four different models for three mock catalogs of different sizes. We present the results for a mock catalog with 70 injected events, as the remaining two catalogs are still in progress.

3.1. Gaussian Model

We first used a Gaussian model to fit both the χ_i and $\cos\theta$ distributions for the three different populations. χ_i is drawn from a Gaussian distribution truncated between 0 and 1, while $\cos\theta$ is drawn from a Gaussian distribution truncated between -1 and 1.

$$p(\chi_i|\mu_\chi, \sigma_\chi) = N_{[0,1]}(\chi_i|\mu_\chi, \sigma_\chi)$$

$$p(\cos\theta|\mu_t, \sigma_t) = N_{[-1,1]}(\cos\theta|\mu_t, \sigma_t)$$

The results are shown in Figure 8.

The Gaussian distribution is the simplest model that we implemented, with only two hyperparameters, mean (μ) and standard deviation (σ). This model fits the $\cos\theta$ distribution of population 1 well, but the remaining fits for the

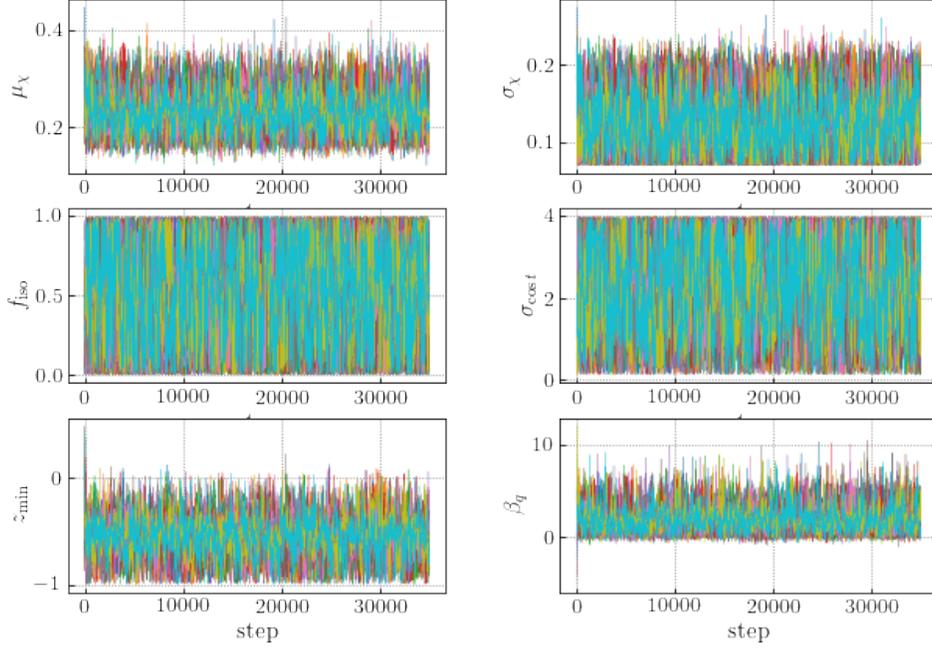


Figure 5. Example of parameter values for each walker at every step in the chain during parameter estimation. Each color corresponds to a different walker.

χ_i and $\cos \theta$ distributions are poor, indicating a poor model choice. However, despite the poor fits on χ_i and $\cos \theta$, χ_{eff} is well fit, demonstrating how effective spin is a well-constrained parameter of LIGO data.

3.2. Beta Plus Mixture Model

We also implemented the model that was used in O3A, a beta plus mixture model. The χ_i distribution is modeled by a beta distribution,

$$p(\chi_i|\alpha, \beta) = \frac{\chi_i^{1-\alpha}(1-\chi_i)^{\beta-1}}{c(\alpha, \beta)}$$

where α and β are both constrained to be greater than 1 to ensure a bounded distribution and $c(\alpha, \beta)$ is the normalization factor.

The $\cos \theta$ distribution is modeled as a mixture between an aligned and isotropic subpopulation,

$$p(\cos \theta|f_{iso}, \sigma_t) = \frac{f_{iso}}{2} + (1 - f_{iso})N_{[-1,1]}(\cos \theta|1, \sigma_t)$$

where f_{iso} is the fraction of events that make up the isotropic subpopulation. The second term in the equation is a truncated half-Gaussian that peaks at $\cos \theta = 1$ with width σ_t .

The results are shown in Figure 9. This model performs better than the Gaussian model, but the poor fit is especially evident in the $\cos \theta$ distribution of population 1. Once again, effective spin is well-constrained.

3.3. Beta Plus Truncated Mixture Model

We also implemented a beta plus truncated mixture model, where we introduced a lower truncation bound in the $\cos \theta$ distribution. The χ_i distribution remains the same as the beta plus mixture model, but the $\cos \theta$ distribution becomes:

$$p(\cos \theta|f_{isco}, \sigma_t, z_{min}) = \frac{f_{isco}}{1 - z_{min}} + (1 - f_{isco})N_{[z_{min},1]}(\cos \theta|1, \sigma_t)$$

This model allowed us to explore whether there is a preference toward more spin alignment. The results are shown in Figure 10.

The trace plots suggest the existence of a lower truncation bound in all three populations of the $\cos \theta$ distribution.

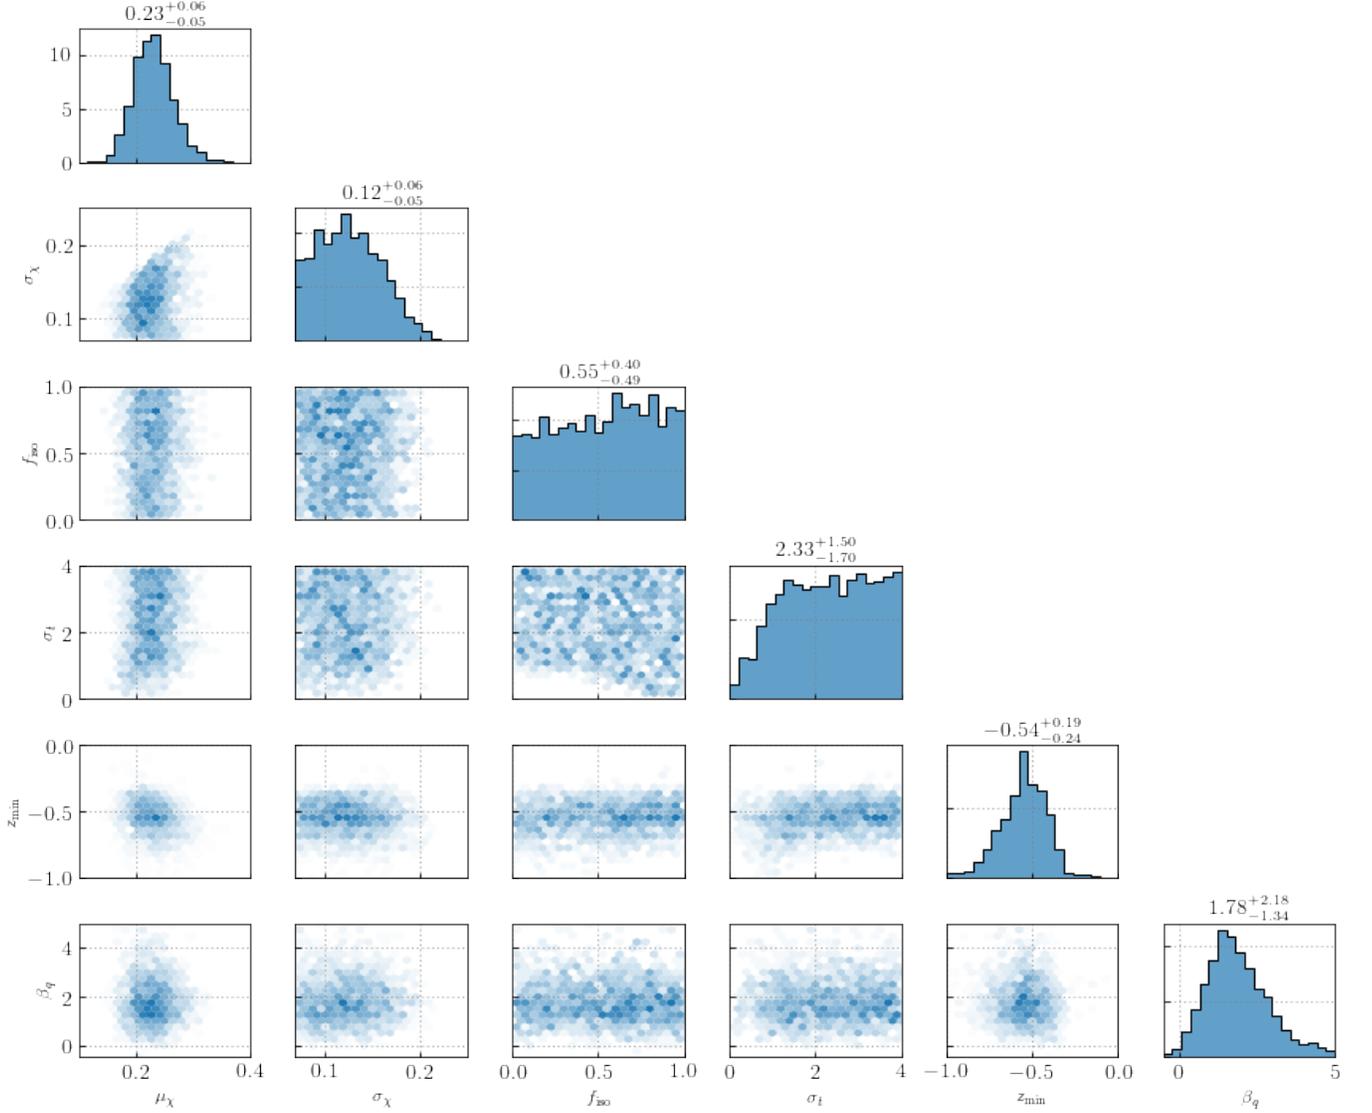


Figure 6. Corner plot showing marginalized posterior distributions of hyperparameters of interest for population 3, using a population size of $n = 10$.

3.4. Binned Model

Lastly, we explored the effects of implementing a binned model. This model does not assume any underlying distribution, allowing for the model to better fit the data. The results are shown in Figure 11.

The binned model was the most successful in recovering the underlying injected distributions for both the χ_i and $\cos\theta$ distributions. We can further evaluate how well our models are performing using posterior predictive plots. The posterior predictive plot for the binned model is shown in Figure 12. If the model performs perfectly, we expect the plot to follow $y = x$, which is plotted in red. The posterior predictive plots show the error in our fits as well as where the model overfits and underfits.

4. CONCLUSIONS

Through a Bayesian hierarchical inference approach, we conclude that we are able to successfully differentiate synthetic populations of BBH mergers with the same effective spin distribution but different component spin distributions, specifically through the binned model. We explored a mock catalog of 70 events, analogous to the size of the existing LIGO catalog. We plan to complete this project by also analyzing mock catalogs with 150 events and 300 events to quantify how current models of spin distributions will improve with LIGO's next observing run.

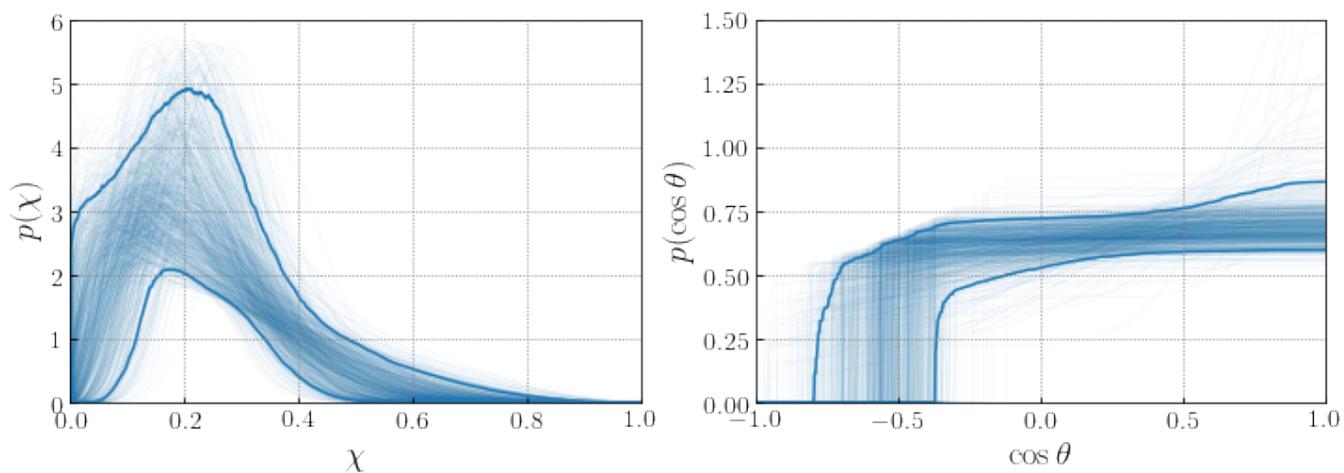


Figure 7. Trace plot showing the set of spin magnitude and tilt angle distributions for a single population. This trace plot assumes an underlying model of a beta plus truncated mixture. The solid blue curves show the 95th percentile.

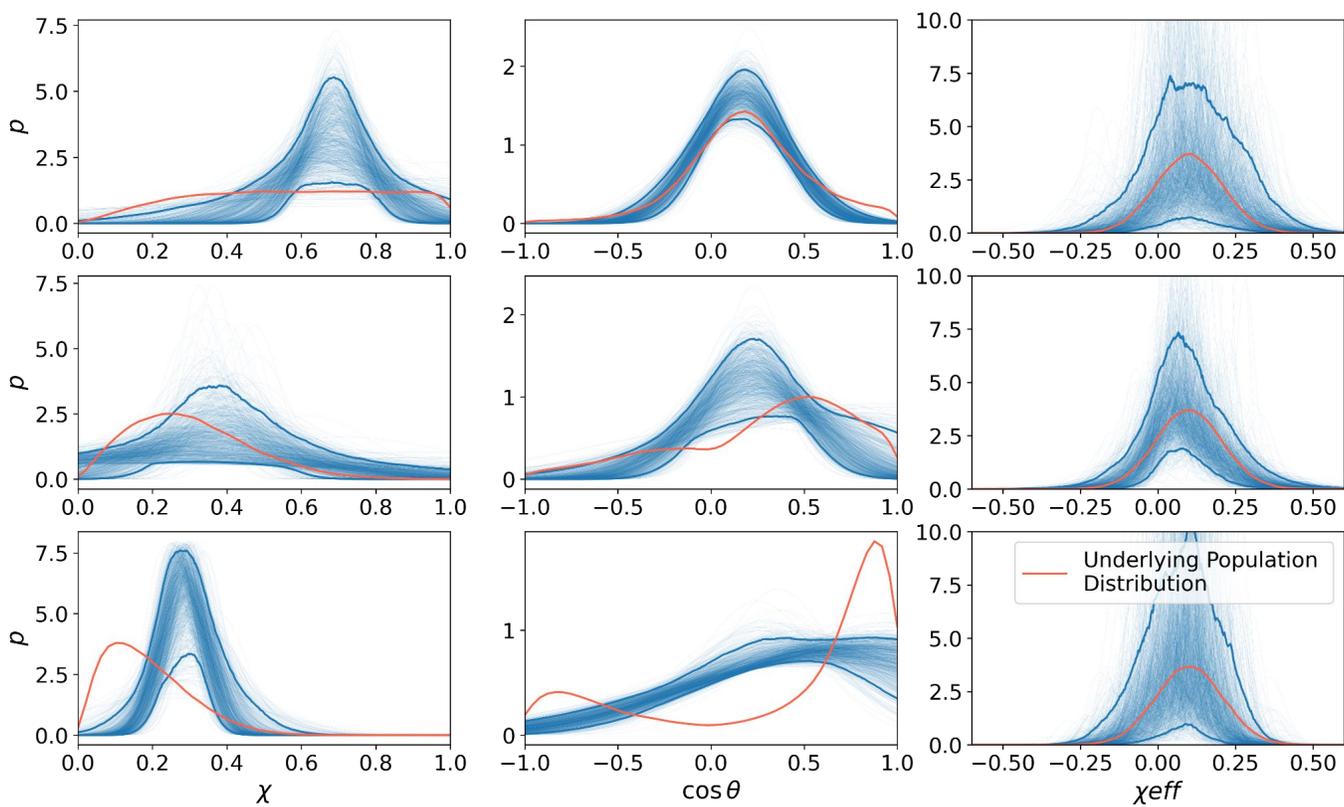


Figure 8. Trace plots showing the spin magnitude, tilt angle, and effective spin distributions for all three populations using a Gaussian model. The solid blue lines show the 95th percentile confidence interval. The red curves show the underlying distributions from Figure 4.

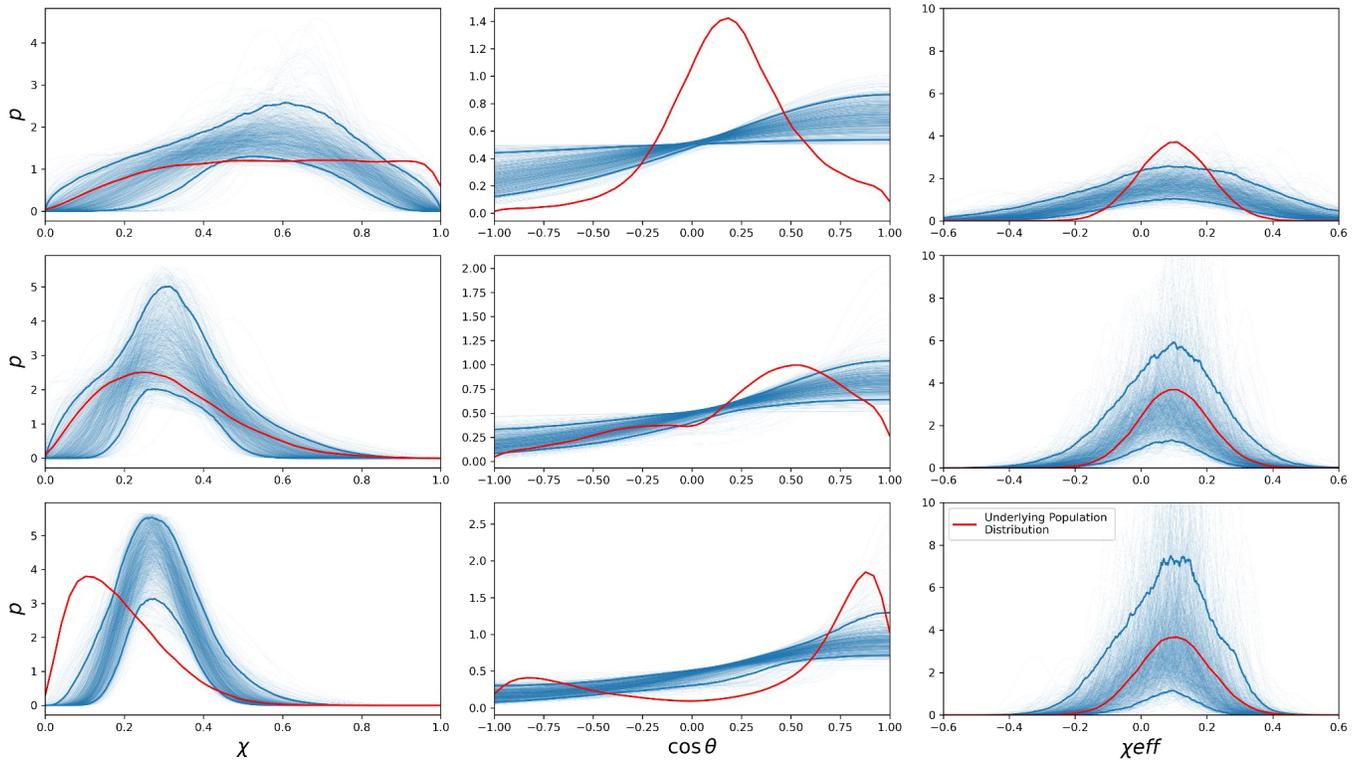


Figure 9. Trace plots showing the spin magnitude, tilt angle, and effective spin distributions for all three populations using a beta plus mixture model. The solid blue lines show the 95th percentile confidence interval. The red curves show the underlying distributions from Figure 4.

5. ACKNOWLEDGEMENTS

This work was made possible by the LIGO Summer Undergraduate Research Fellowships (SURF) program, run by Alan Weinstein and funded through the Research Experiences for Undergraduates (REU) Program of the National Science Foundation.

REFERENCES

- Abbott, B. P., Abbott, R., Abbott, T., et al. 2016a, *Physical review letters*, 116, 241102
- . 2016b, *Physical Review D*, 93, 122003
- Abbott, R., Abbott, T., Acernese, F., et al. 2021, arXiv preprint arXiv:2111.03634
- Ashton, G., Hübner, M., Lasky, P. D., et al. 2019, *The Astrophysical Journal Supplement Series*, 241, 27
- Callister, T. A., Haster, C.-J., Ng, K. K., Vitale, S., & Farr, W. M. 2021, *The Astrophysical Journal Letters*, 922, L5
- Callister, T. A., Miller, S. J., Chatziioannou, K., & Farr, W. M. 2022, arXiv preprint arXiv:2205.08574
- Collaboration, L. S., Collaboration, V., Collaboration, K. S., et al. 2021, arXiv preprint arXiv:2111.03634
- Mapelli, M. 2020, *Frontiers in Astronomy and Space Sciences*, 7, 38
- Miller, S., Callister, T. A., & Farr, W. M. 2020, *The Astrophysical Journal*, 895, 128

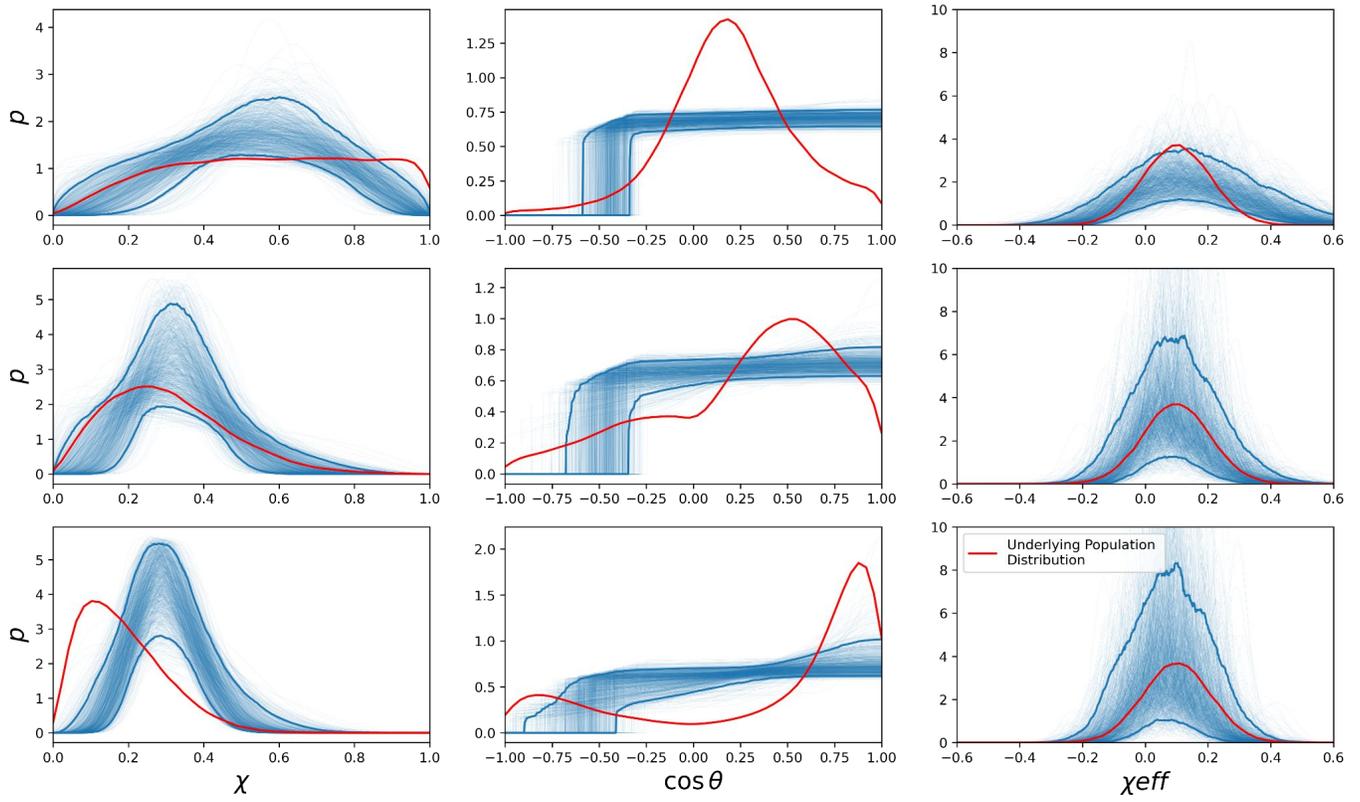


Figure 10. Trace plots showing the spin magnitude, tilt angle, and effective spin distributions for all three populations using a beta plus truncated mixture model. The solid blue lines show the 95th percentile confidence interval. The solid blue lines show the 95th percentile confidence interval. The red curves show the underlying distributions from Figure 4.

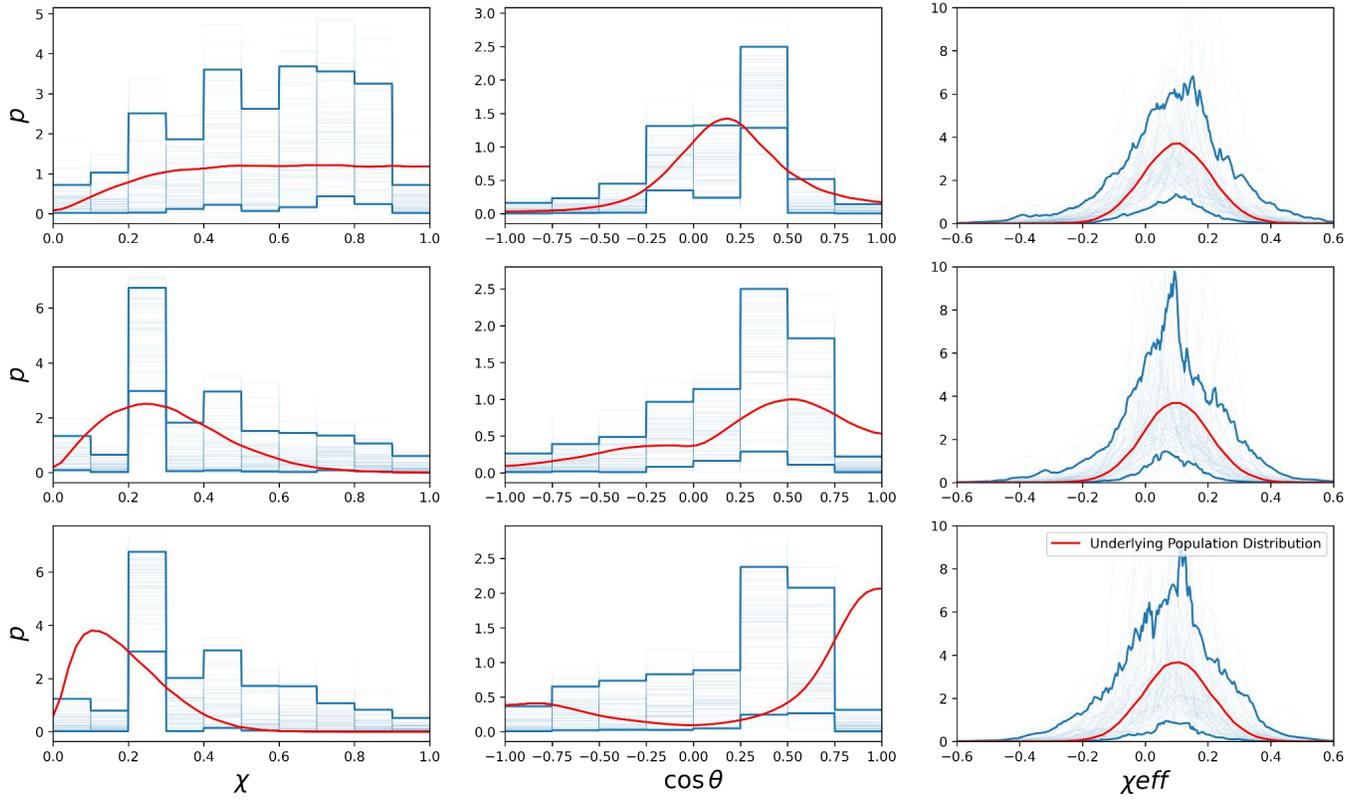


Figure 11. Trace plots showing the spin magnitude, tilt angle, and effective spin distributions for all three populations using a binned model. The red curves show the underlying distributions from Figure 4.

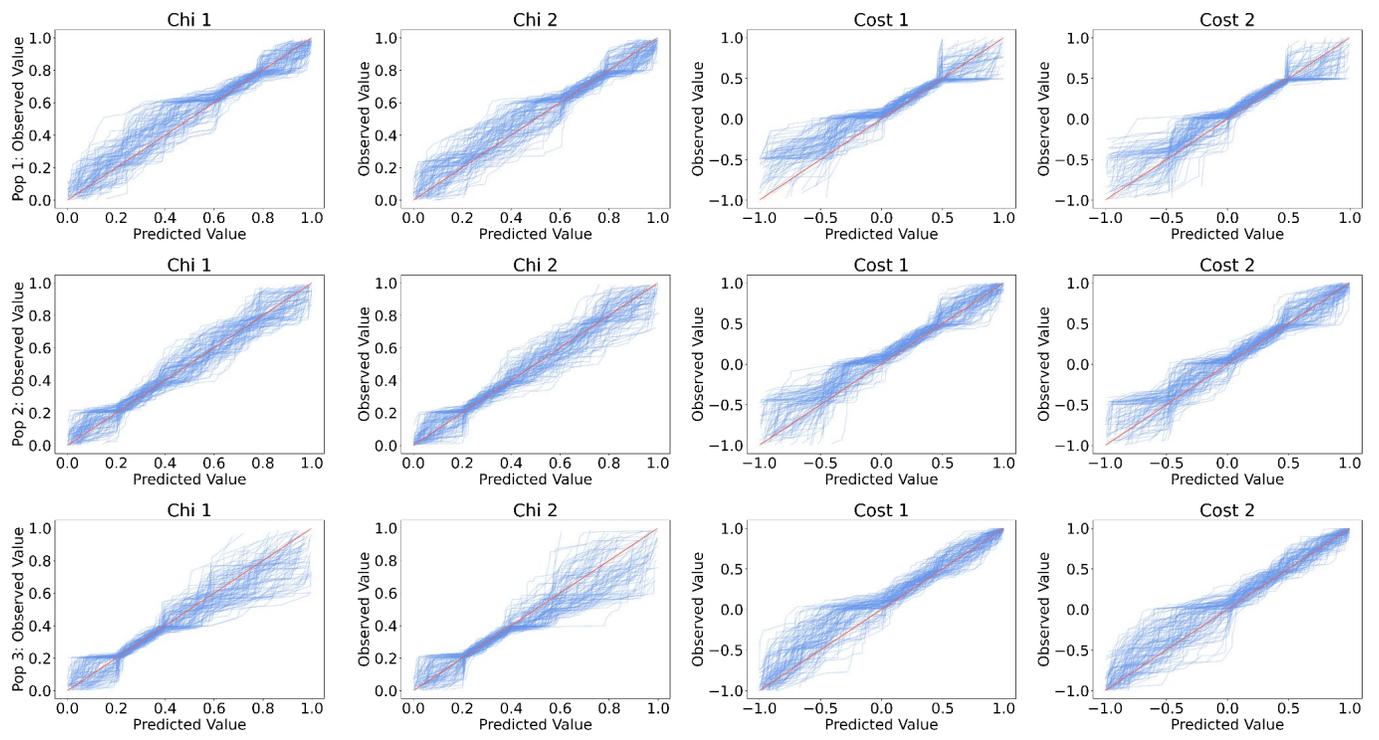


Figure 12. The posterior predictive plots for the four parameters of interest across the three populations for the binned model.