Decoding the cosmic dipole with likelihood-based and likelihood-free inference

Oliver Oayda

PhD Candidate
Sydney Institute for Astronomy
The University of Sydney

Supervised by
Geraint Lewis
Tara Murphy

December 16, 2025

MaxEnt @ The University of Auckland

December 2025

Simulation-based inference helps us understand tricky systematics and measure what we're actually interested in.

CatWISE2020

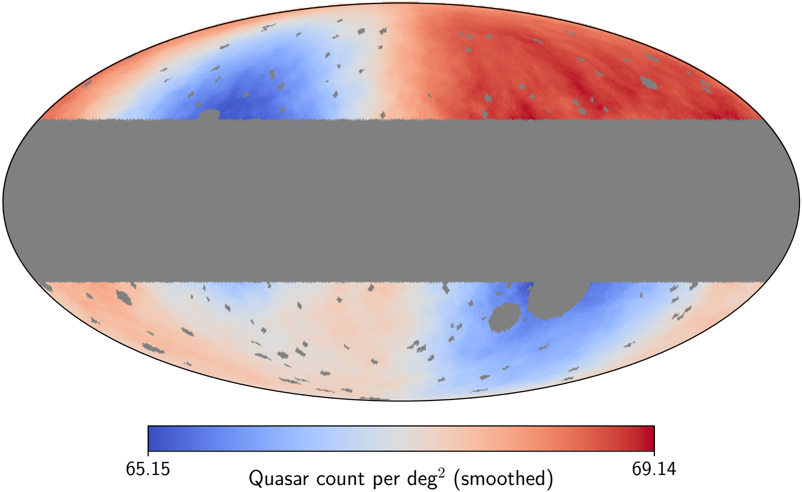

CatWISE2020 quasar sample.

The Kinematic Dipole

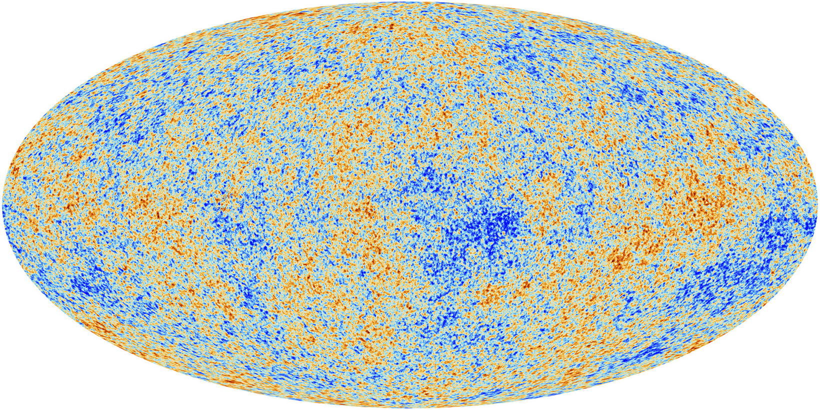

CMB temperature map (Planck satellite).

The Kinematic Dipole

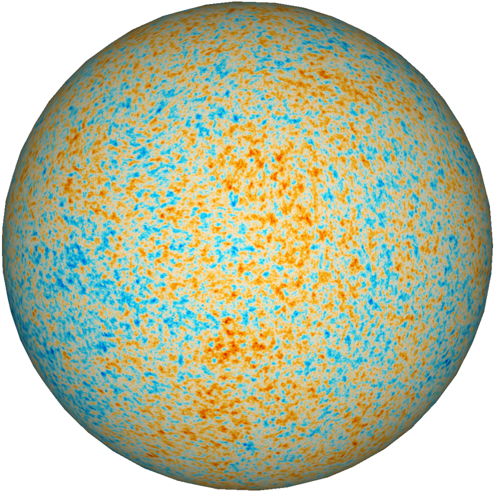

CMB as a sphere.

The Kinematic Dipole

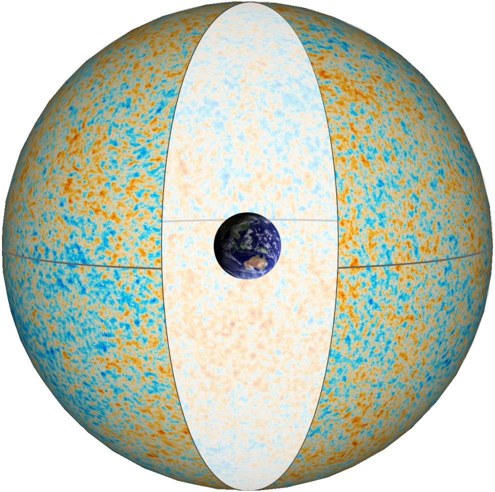

CMB as a sphere (Earth inside).

The Kinematic Dipole

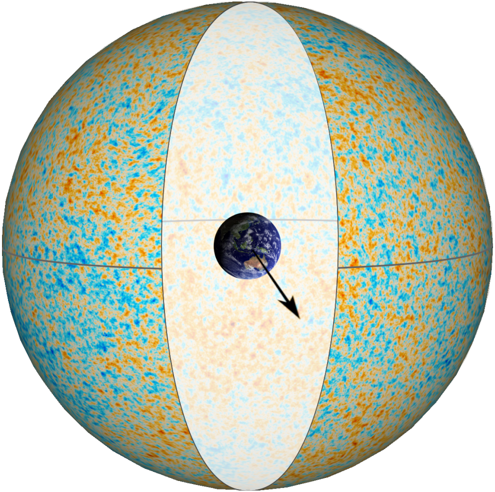

We're moving through the Universe!

The Kinematic Dipole

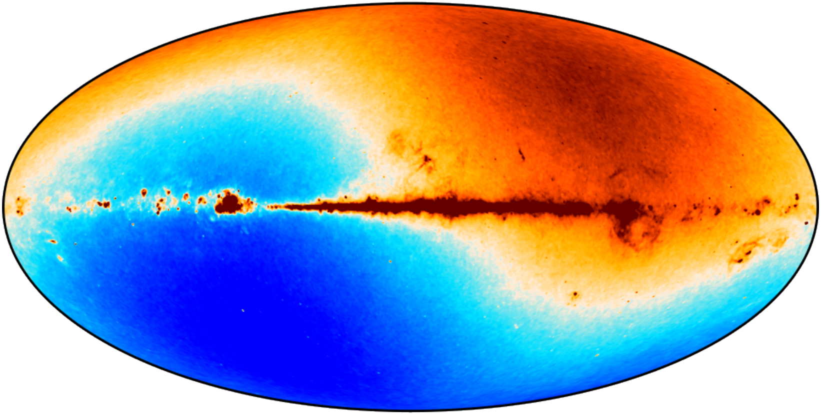

CMB temperature map (dipole included; BeyondPlanck).
★: dipole direction.

The Kinematic Dipole

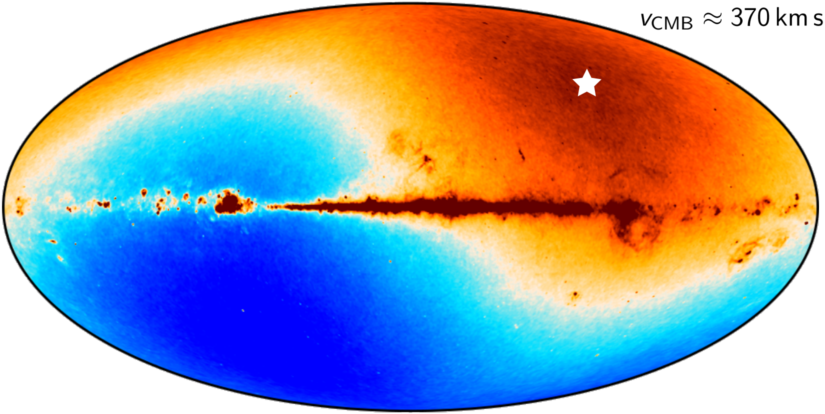

$v_{\mathrm{CMB}} \approx 370\,\mathrm{km\,s^{-1}}$

CMB temperature map (dipole included; BeyondPlanck).
★: dipole direction.

The Cosmic Dipole

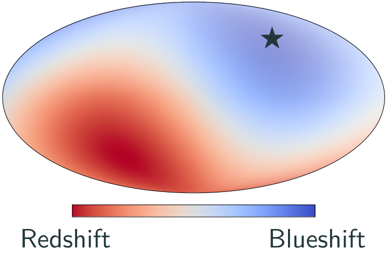

Our motion $\implies$ a dipole in source density.

The Cosmic Dipole

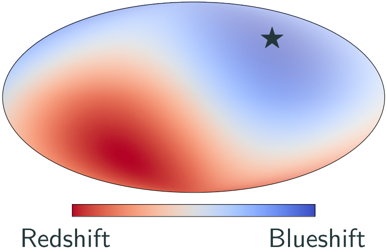

Our motion $\implies$ a dipole in source density.

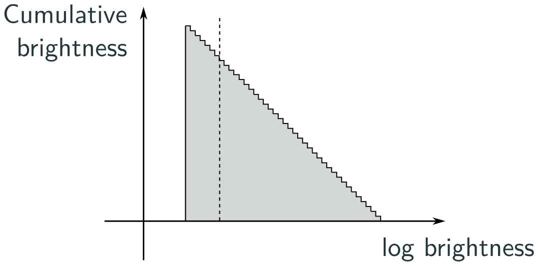

The Cosmic Dipole

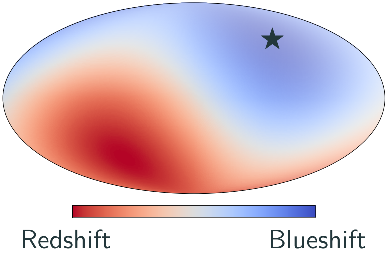

Our motion $\implies$ a dipole in source density.

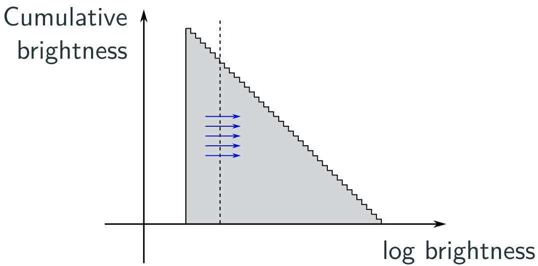

The Cosmic Dipole

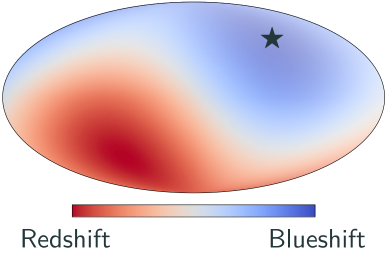

Our motion $\implies$ a dipole in source density.

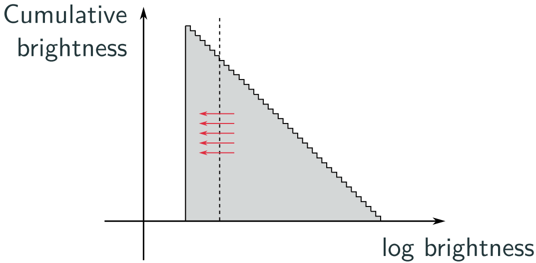

The Cosmic Dipole

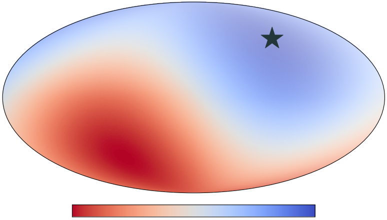

Our motion $\implies$ a dipole in source density.

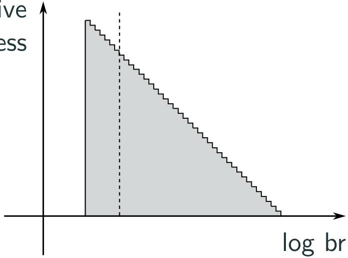

The Cosmic Dipole

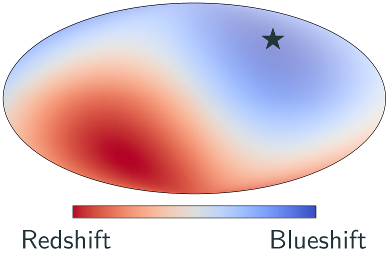

Our motion $\implies$ a dipole in source density.

$$\mathcal{D}_{\mathrm{CMB}} = [2 + x\,(1+\alpha)]\,\frac{v_{\mathrm{CMB}}}{c}.$$

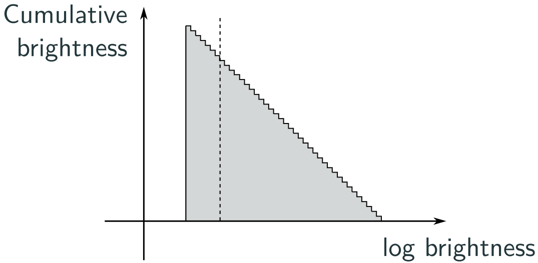

The Cosmic Dipole

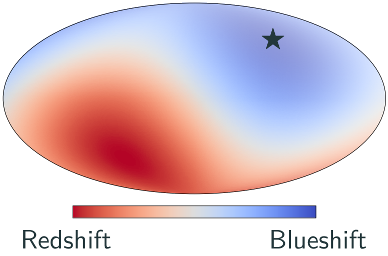

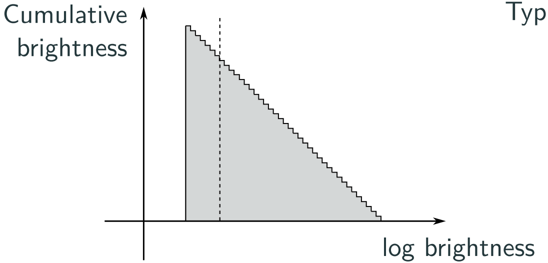

Our motion $\implies$ a dipole in source density.

$$\boxed{\mathcal{D}_{\text{CMB}} = [2 + x(1+\alpha)]\,\frac{v_{\text{CMB}}}{c}.}$$

Typical values: 0.004 – 0.007.

A *0.5% effect!*

Counting Galaxies

$$N_i = \overline{N}(1 + \mathcal{D}\cos\theta_i)$$

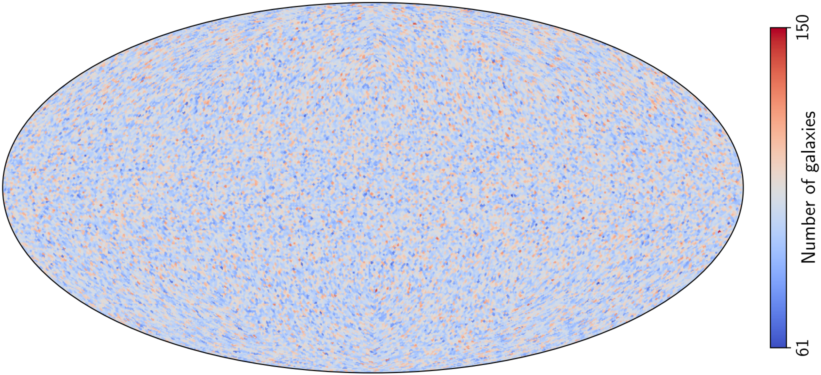

Simulated isotropic galaxy map (+ dipole).

Counting Galaxies

$$N_i = \overline{N}(1 + \mathcal{D}\cos\theta_i)$$

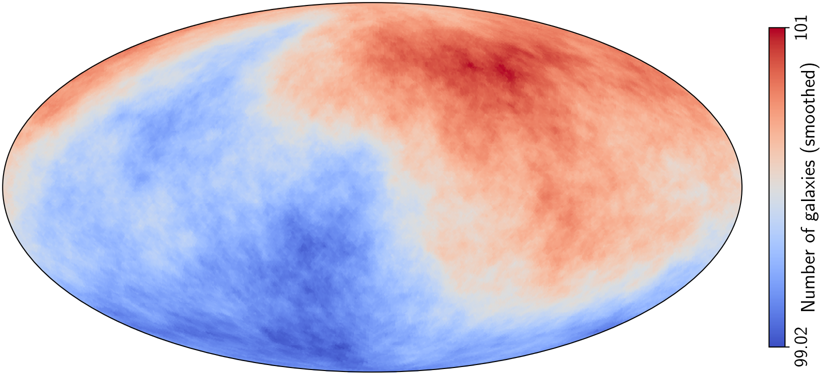

Simulated isotropic galaxy map (+ dipole).

The Amplitude Excess

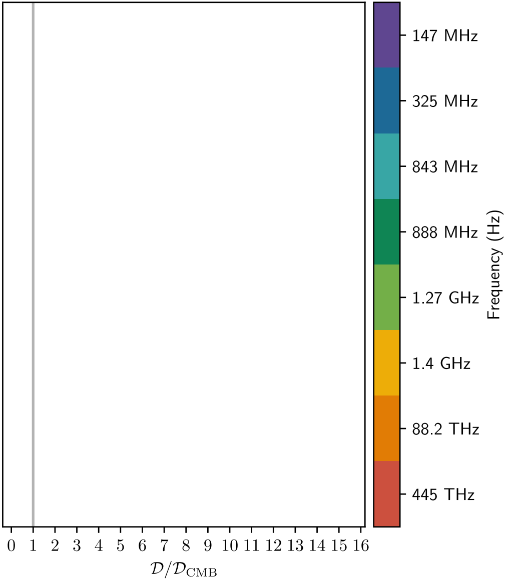

Cosmic dipole should be consistent with CMB dipole...

The Amplitude Excess

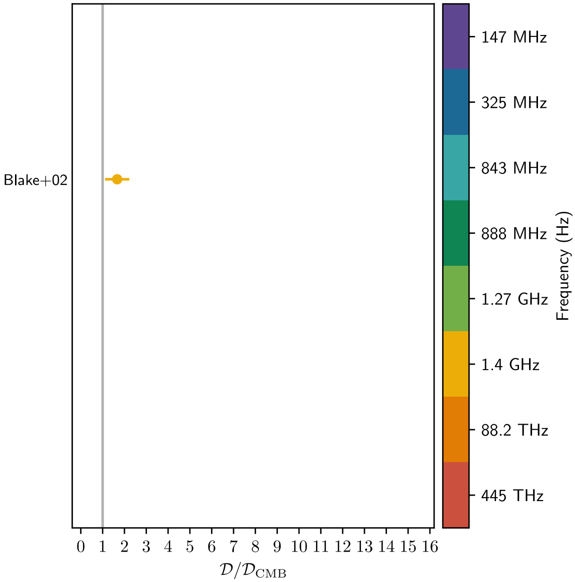

Cosmic dipole should be consistent with CMB dipole...

All is well!

The Amplitude Excess

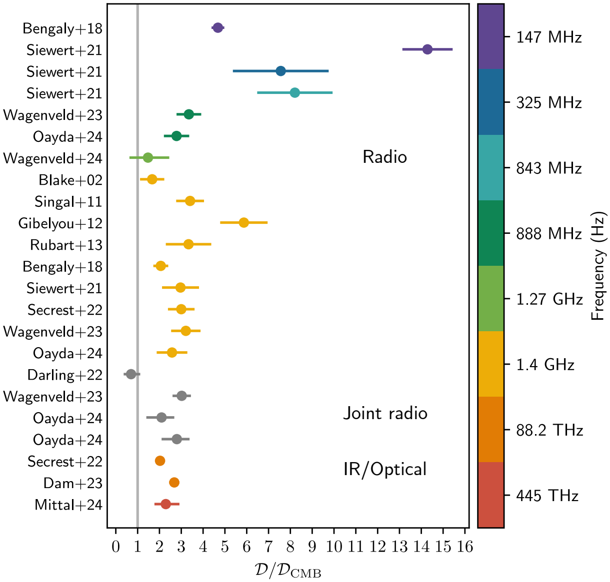

Cosmic dipole should be consistent with CMB dipole...

All is well!

Wait...

Quasars in CatWISE2020

The Astrophysical Journal Letters, 908:L51 (6pp), 2021 February 20
https://doi.org/10.3847/2041-8213/abdd40

A Test of the Cosmological Principle with Quasars

Nathan J. Secrest[1], Sebastian von Hausegger[2,3,4], Mohamed Rameez[5], Roya Mohayaee[3], Subir Sarkar[4], and Jacques Colin[3]

Quasars in CatWISE2020

The Astrophysical Journal Letters, 908:L51 (6pp), 2021 February 20
https://doi.org/10.3847/2041-8213/abdd40

CrossMark

A Test of the Cosmological Principle with Quasars

Nathan J. Secrest[1], Sebastian von Hausegger[2,3,4], Mohamed Rameez[5], Roya Mohayaee[3], Subir Sarkar[4], and Jacques Colin[3]

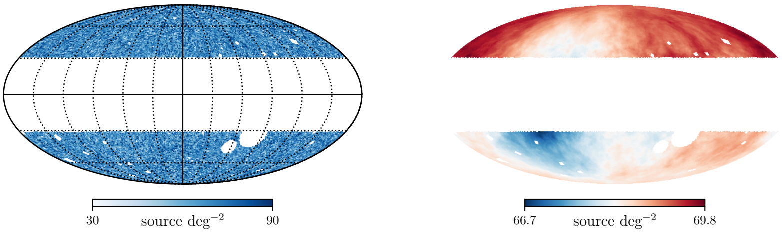

CatWISE2020 sample (counts + smoothed).

Quasars in CatWISE2020

The Astrophysical Journal Letters, 908:L51 (6pp), 2021 February 20
https://doi.org/10.3847/2041-8213/abdd40

CrossMark

A Test of the Cosmological Principle with Quasars

Nathan J. Secrest[1], Sebastian von Hausegger[2,3,4], Mohamed Rameez[5], Roya Mohayaee[3], Subir Sarkar[4], and Jacques Colin[3]

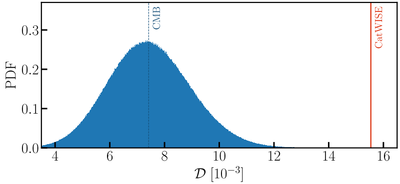

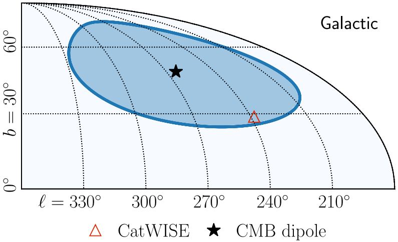

Null significance test (dipole amplitude + direction).

Quasars in CatWISE2020

The Astrophysical Journal Letters, 937:L31 (9pp), 2022 October 1
https://doi.org/10.3847/2041-8213/ac88c0

A Challenge to the Standard Cosmological Model

Nathan J. Secrest[1], Sebastian von Hausegger[2], Mohamed Rameez[3], Roya Mohayaee[2,4], and Subir Sarkar[2]

Quasars in CatWISE2020

The Astrophysical Journal Letters, 937:L31 (9pp), 2022 October 1
https://doi.org/10.3847/2041-8213/ac88c0

A Challenge to the Standard Cosmological Model

Nathan J. Secrest[1], Sebastian von Hausegger[2], Mohamed Rameez[3], Roya Mohayaee[2,4], and Subir Sarkar[2]

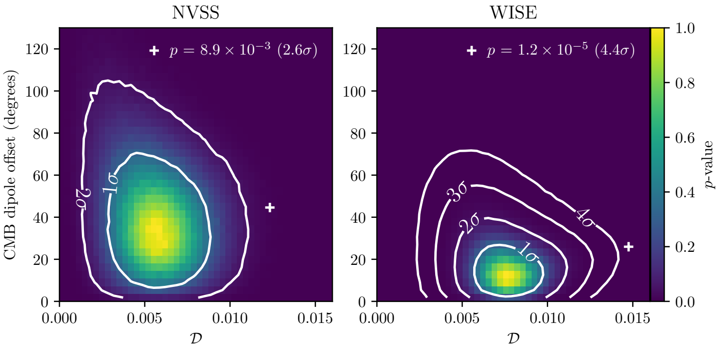

Null significance test (dipole amplitude + offset).

More to the story...

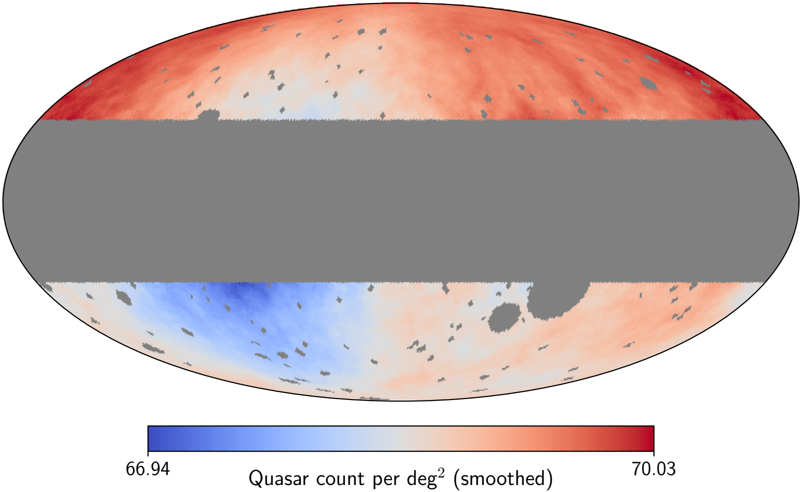

CatWISE quasar map from Secrest+21.

More to the story...

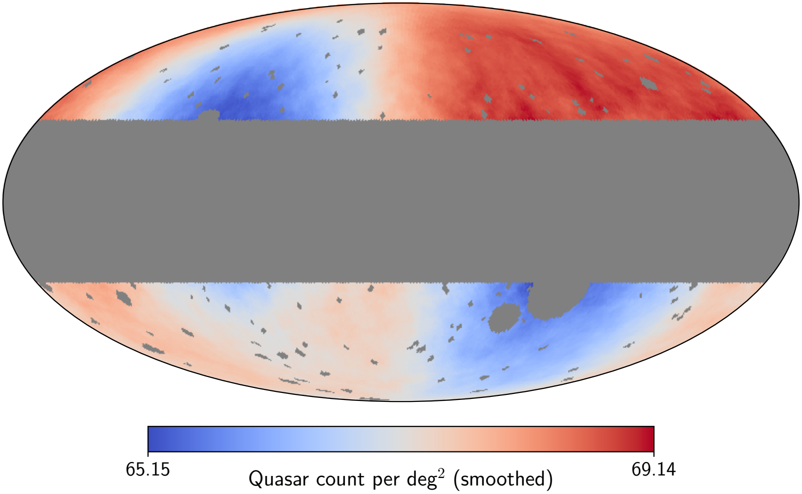

CatWISE quasar map, no linear weighting.

More to the story...

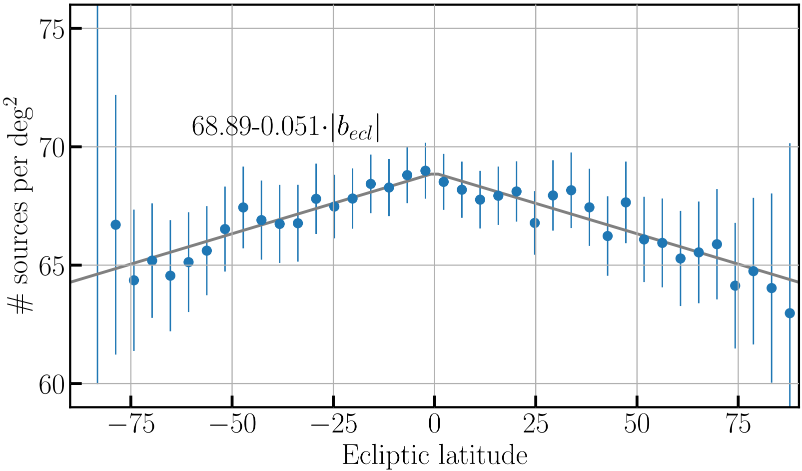

Linear fit to density vs. declination (Secrest+21).

Bayesian Treatment

Testing the cosmological principle with CatWISE quasars: a bayesian analysis of the number-count dipole

Lawrence Dam[1,2★], Geraint F. Lewis[1★] and Brendon J. Brewer[3]

- Frequentist: ecliptic bias is a correction/weighting to **D**.

Bayesian Treatment

Testing the cosmological principle with CatWISE quasars: a bayesian analysis of the number-count dipole

Lawrence Dam [1,2]★ Geraint F. Lewis [1]★ and Brendon J. Brewer[3]

- Frequentist: ecliptic bias is a correction/weighting to **D**.
- Bayesian: ecliptic bias is a parameter for our model $P(\mathbf{D}|\Theta, M)$.

Bayesian Treatment

Testing the cosmological principle with CatWISE quasars: a bayesian analysis of the number-count dipole

Lawrence Dam [1,2]★ Geraint F. Lewis [1]★ and Brendon J. Brewer[3]

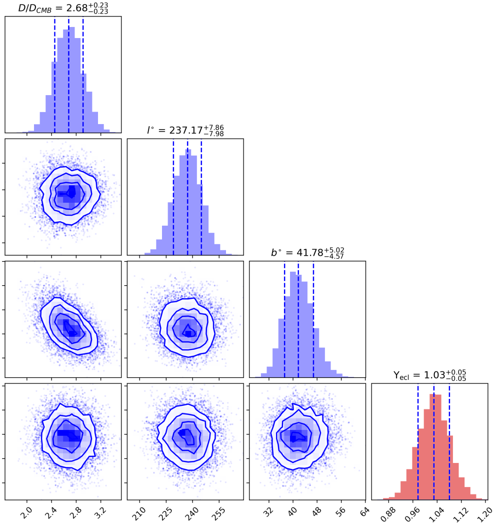

Bayesian Treatment

Testing the cosmological principle with CatWISE quasars: a bayesian analysis of the number-count dipole

Lawrence Dam[1,2★], Geraint F. Lewis[1★] and Brendon J. Brewer[3]

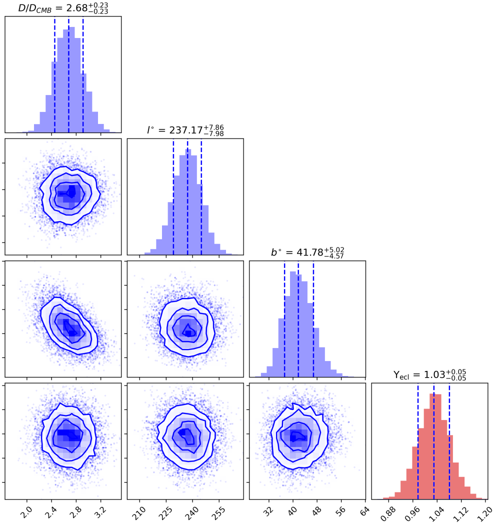

Model	ln B
Null (no dipole)	0.0
Free dipole, no bias	83.5
Free dipole	263.5

**Is it good enough to assume a linear fit —
— to make an ad hoc correction after the fact?**

WISE's Scanning Law

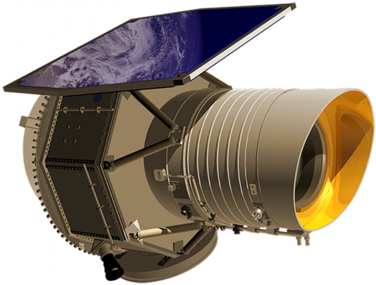

Render of the WISE satellite (NASA JPL).

WISE's Scanning Law

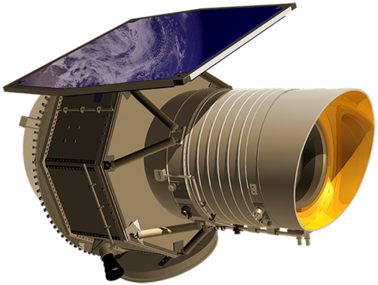

Render of the WISE satellite (NASA JPL).

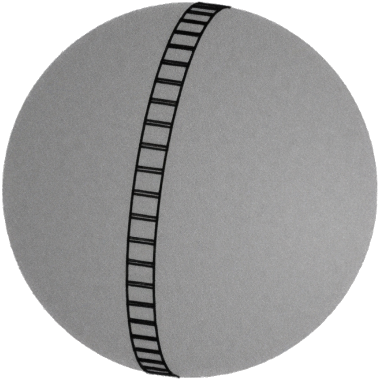

Frames seen over 1 orbit.

WISE's Scanning Law

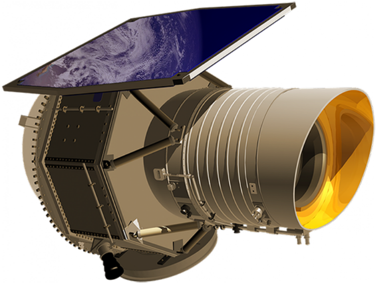

Render of the WISE satellite (NASA JPL).

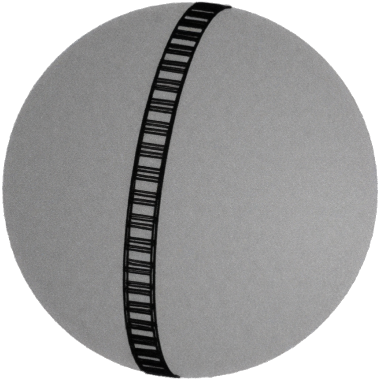

Frames seen over 2 orbits.

WISE’s Scanning Law

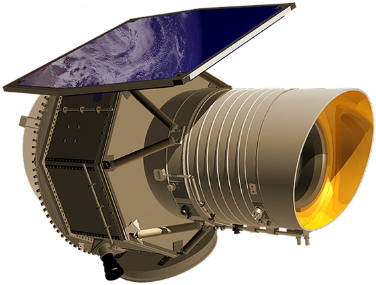

Render of the WISE satellite (NASA JPL).

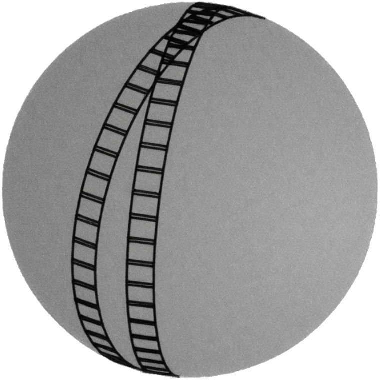

Frames seen over 2 orbits 20 days apart.

WISE's Scanning Law

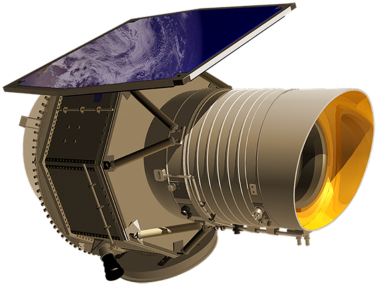

Render of the WISE satellite (NASA JPL).

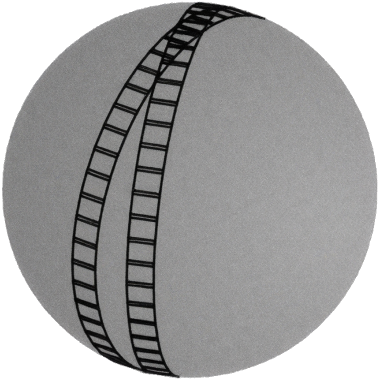

Frames seen over 2 orbits 20 days apart.

Obeys a scanning law over the survey's lifetime.

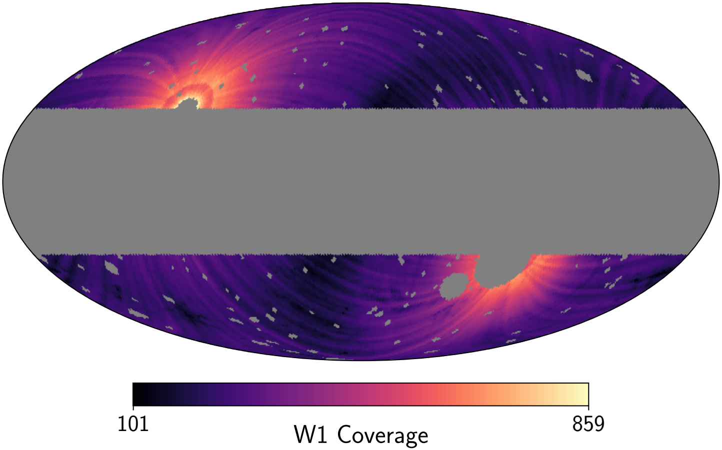

WISE coverage in W1 band for the Secrest+21 sample.

Oayda & Lewis (2025), submitted to MNRAS

Photometric Errors

Oayda & Lewis (2025), submitted to MNRAS

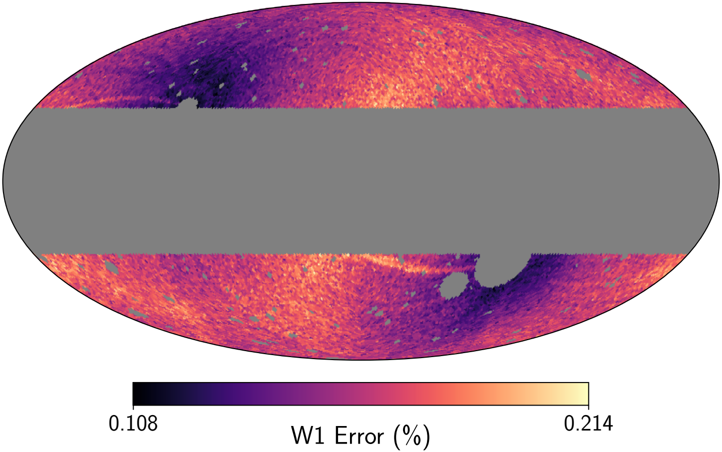

Source photometric error (%) in W1 band for the Secrest+21 sample.

Photometric Errors

Oayda & Lewis (2025), submitted to MNRAS

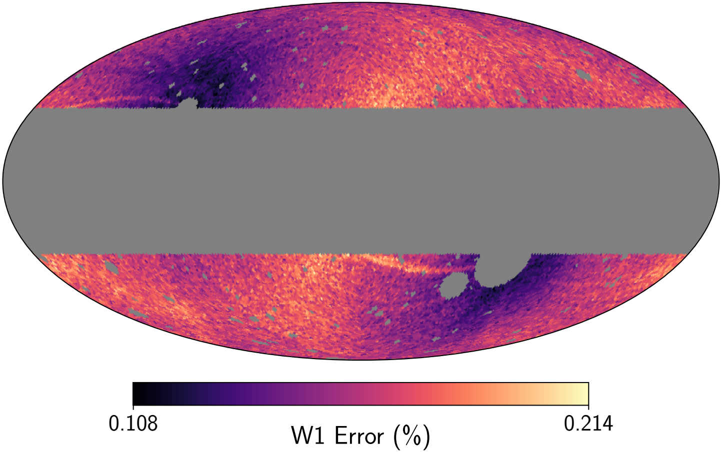

Source photometric error (%) in W1 band for the Secrest+21 sample.

We should hope $\sigma \propto \frac{1}{\sqrt{\mathrm{Cov}}}$ — photon counting.

Photometric Errors

Oayda & Lewis (2025), submitted to MNRAS

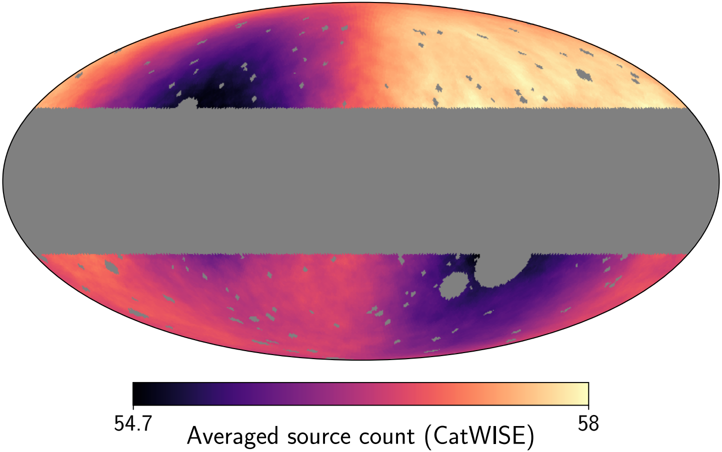

CatWISE quasar map, no linear weighting.

We should hope $\sigma \propto \frac{1}{\sqrt{\text{Cov}}}$ — photon counting.

Eddington Bias

$$S_i = S_i^{\text{true}}$$

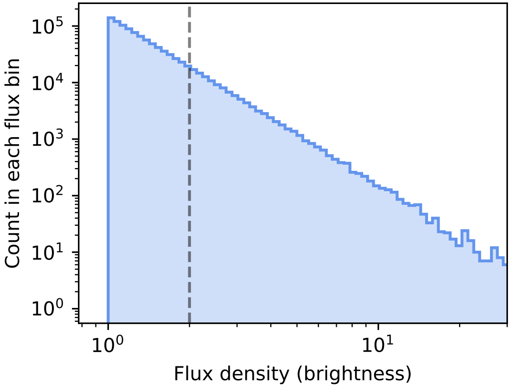

True distribution of flux densities.

Eddington Bias

$$S_i\,(>2) = S_i^{\text{true}}$$

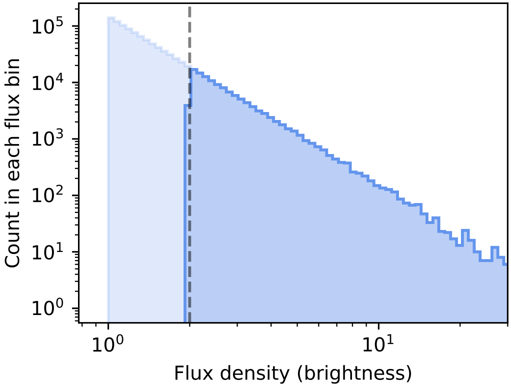

True distribution of flux densities with flux cut.

Eddington Bias

$$S_i\,(>2) = S_i^{\text{true}} + \Delta S \qquad \Delta S \sim \mathcal{G}(\mu = 0, \sigma = \sigma_{\text{const.}})$$

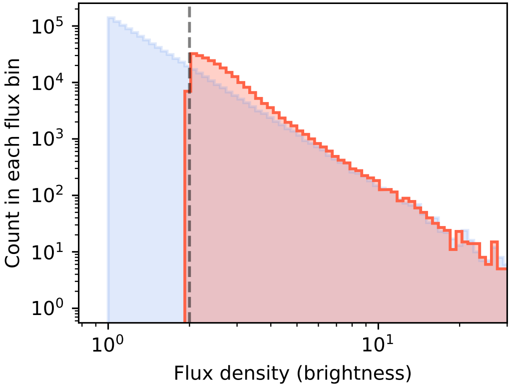

Noisy distribution of flux densities with flux cut.

Turning to Simulations

- This is just something we can simulate.

Turning to Simulations

- This is just something we can simulate.
    - Draw sources on isotropic sky, sample magnitudes.

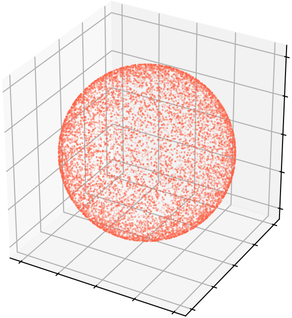

Turning to Simulations

- This is just something we can simulate.
  - Draw sources on isotropic sky, sample magnitudes.
  - Do special relativity.

$$(l^\circ, b^\circ) \to (l'^\circ, b'^\circ)$$
$$m_\nu \to m'_\nu$$
$$d\Omega \to d\Omega'$$

Turning to Simulations

- This is just something we can simulate.
    - Draw sources on isotropic sky, sample magnitudes.
    - Do special relativity.
    - Add photometric errors.

$$
\begin{aligned}
(l^\circ, b^\circ) &\rightarrow (l'^\circ, b'^\circ) \\
m_\nu &\rightarrow m'_\nu \\
d\Omega &\rightarrow d\Omega' \\
m'_\nu &\rightarrow m'_\nu + \Delta m'_\nu
\end{aligned}
$$

Turning to Simulations

- This is just something we can simulate.
  - Draw sources on isotropic sky, sample magnitudes.
  - Do special relativity.
  - Add photometric errors.
  - Make magnitude cuts.

$$9 < m'_{\mathrm{W1}} < 16.4$$
$$m'_{\mathrm{W1}} - m'_{\mathrm{W2}} > 0.8$$

Turning to Simulations

- This is just something we can simulate.
    - Draw sources on isotropic sky, sample magnitudes.
    - Do special relativity.
    - Add photometric errors.
    - Make magnitude cuts.
- Key point: what do we do with the errors?

Turning to Simulations

Oayda & Lewis (2025), submitted to MNRAS

- This is just something we can simulate.
  - Draw sources on isotropic sky, sample magnitudes.
  - Do special relativity.
  - Add photometric errors.
  - Make magnitude cuts.
- Key point: what do we do with the errors?

Magnitude-coverage → error lookup.

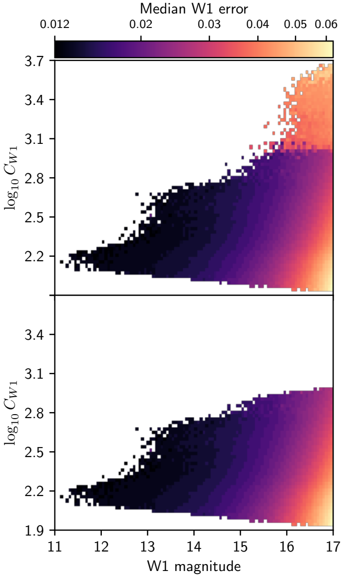

What Now...?

What Now...?

The classic explicit-likelihood approach:

$$P(\Theta|\mathbf{D}, M) = \frac{\mathcal{L}(\mathbf{D}|\Theta, M)\,\pi(\Theta|M)}{\mathcal{Z}(\mathbf{D}|M)}.$$

What Now...?

The classic explicit-likelihood approach:

$$P(\Theta|\mathbf{D}, M) = \frac{\mathcal{L}(\mathbf{D}|\Theta, M)\,\pi(\Theta|M)}{\mathcal{Z}(\mathbf{D}|M)}.$$

$$\mathcal{L} = \prod^{N_{\mathrm{pix.}}} \mathrm{Pois}(N_i|\lambda_i, M_{\mathrm{dipole}})$$

where $\lambda_i = \bar{N}(1 + \mathcal{D}\cos\theta)$.

What Now...?

Don't know $\mathcal{L}$? No problem!

What Now...?

Don't know $\mathcal{L}$? No problem!

Take data-generating process
$f_M : \Theta \to \mathbf{D}$.

What Now...?

Don't know $\mathcal{L}$? No problem!

Take data-generating process
$f_M : \Theta \to \mathbf{D}$.

Use neural network to learn $\mathcal{L}(\mathbf{D}|\Theta, M)$.

This is **S**imulation-**b**ased **I**nference.

Neural Likelihood Estimator

- We rely on a Neural Likelihood Estimator (NLE).

Neural Likelihood Estimator

- We rely on a Neural Likelihood Estimator (NLE).
- Learn a surrogate likelihood $q_\phi(x|\theta) \approx P(x|\theta) = \mathcal{L}$.

Neural Likelihood Estimator

- We rely on a Neural Likelihood Estimator (NLE).
- Learn a surrogate likelihood $q_\phi(x|\theta) \approx P(x|\theta) = \mathcal{L}$.
- Minimise D_{KL} when inside the support $P(\theta) > 0$.

$$\begin{aligned} \text{NLL} &= -\mathbb{E}_{P(x,\theta)}[\ln q_\phi] \\ &\sim \mathbb{E}_{P(\theta)}\left[D_{KL}(\mathcal{L} \,||\, q_\phi)\right] \end{aligned}$$

Neural Likelihood Estimator

- We rely on a Neural Likelihood Estimator (NLE).
- Learn a surrogate likelihood $q_\phi(x|\theta) \approx P(x|\theta) = \mathcal{L}$.
- Minimise D_{KL} when inside the support $P(\theta) > 0$.
- NLE & nested sampling to get $\mathcal{Z}$.

$$\begin{aligned}\mathrm{NLL} &= -\mathbb{E}_{P(x,\theta)}[\ln q_\phi] \\ &\sim \mathbb{E}_{P(\theta)}\left[D_{KL}(\mathcal{L} \,||\, q_\phi)\right]\end{aligned}$$

github.com/handley-lab/blackjax

Normalising Flows

- We use a series of Masked Autoregressive Flows (MAFs).

Normalising Flows

- We use a series of Masked Autoregressive Flows (MAFs).
- **Flow:** Transform a simple base distribution to a more complex one.

Normalising Flows

- We use a series of Masked Autoregressive Flows (MAFs).
- **Flow:** Transform a simple base distribution to a more complex one.

Composing individual bijections (from Papamakarios et al. 2019).

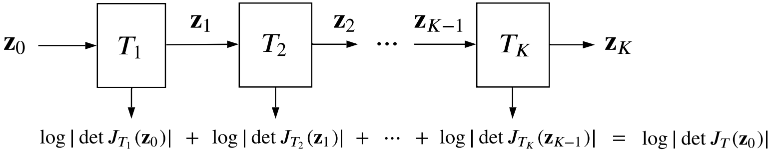

Normalising Flows

- We use a series of Masked Autoregressive Flows (MAFs).
- **Flow:** Transform a simple base distribution to a more complex one.

Composing individual bijections (from Papamakarios et al. 2019).

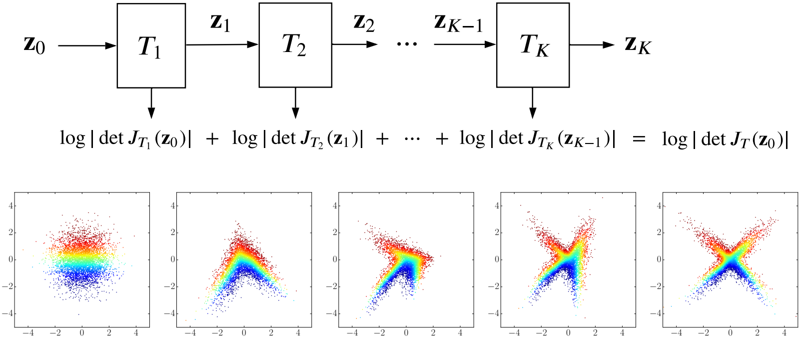

Challenge: Dimensionality

- Simulation output: 49152 pixels.
- Model a 49152-dim base distribution? No.
- We want:
  - A way to $\downarrow$ data dim.
  - Make sure inferences don't change.
- Why? Suppose $\mathbf{z} = f(\mathbf{x}_{\text{sim.}})$. Then $\mathcal{Z}(\mathbf{z}|M) \neq \mathcal{Z}(\mathbf{D}|M)$.

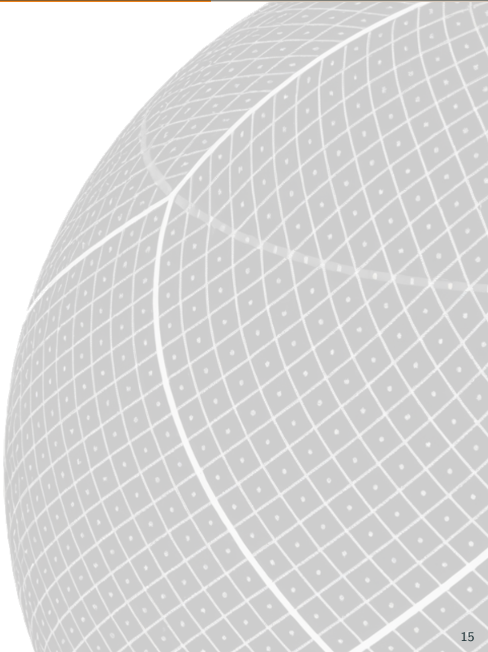

Downscaling

Aim

Check SBI likelihood is consistent with analytic likelihood.

Downscaling

Aim

Check SBI likelihood is consistent with analytic likelihood.

- Suppose $N_i \sim \text{Pois}(\lambda_i)$.

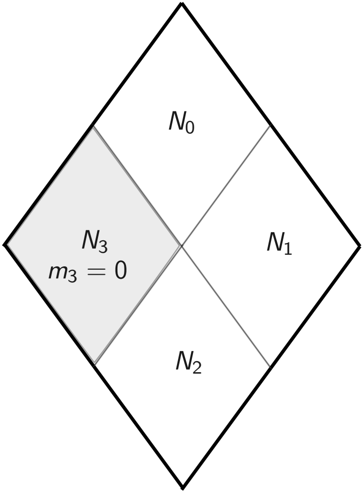

Downscaling

Aim

Check SBI likelihood is consistent with analytic likelihood.

- Suppose $N_i \sim \text{Pois}(\lambda_i)$.
- Then $N^{\text{super}} = \sum_{i=0}^{3} N_i$.

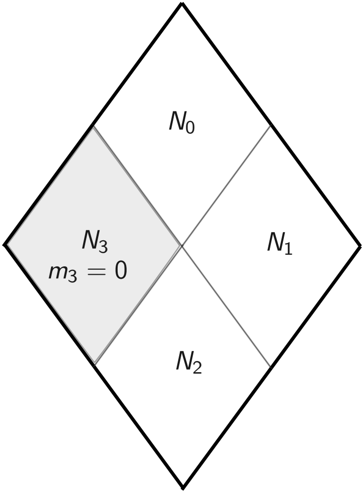

Downscaling

Aim

Check SBI likelihood is consistent with analytic likelihood.

- Suppose $N_i \sim \text{Pois}(\lambda_i)$.
- Then $N^{\text{super}} = \sum_{i=0}^{3} N_i$.
- $N_{\text{super}} \sim \text{Pois}(\lambda = \sum_{i=0}^{3} \lambda_i)$.

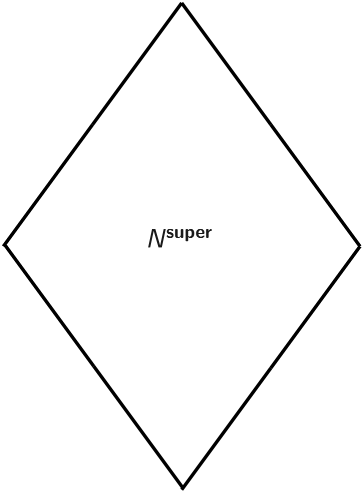

Downscaling

Aim

Check SBI likelihood is consistent with analytic likelihood.

- Suppose $N_i \sim \text{Pois}(\lambda_i)$.
- Then $N^{\text{super}} = \sum_{i=0}^{3} N_i$.
- $N_{\text{super}} \sim \text{Pois}(\lambda = \sum_{i=0}^{3} \lambda_i)$.
- Chain a series of downscales to get to a low resolution output.

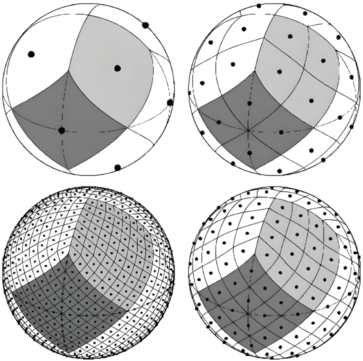

Consistency Check — Evidence

- NLE-computed $\ln \mathcal{Z}$ converges to truth.
- Bayes factors agree with truth (but large dispersion).

Estimated evidence vs. inference round (cf. truth).

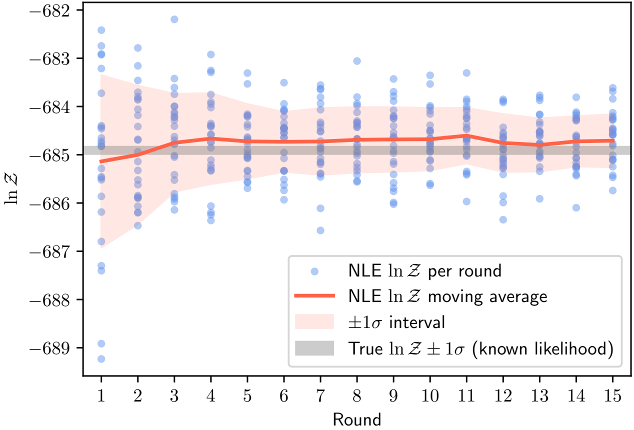

Oayda & Lewis (2025), submitted to MNRAS

Consistency Check — Evidence

- NLE-computed $\ln \mathcal{Z}$ converges to truth.
- Bayes factors agree with truth (but large dispersion).

Estimated Bayes factors versus truths.

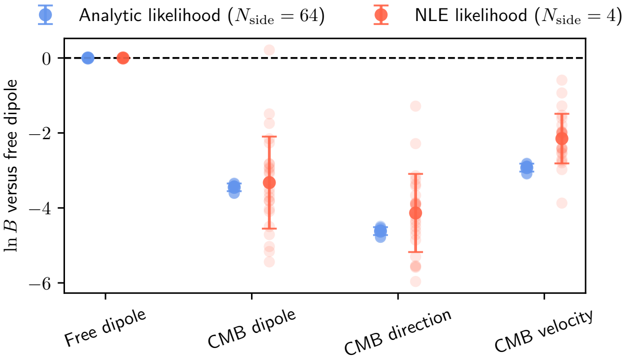

Consistency Check — Posteriors

SBI-derived posterior versus ground truth.

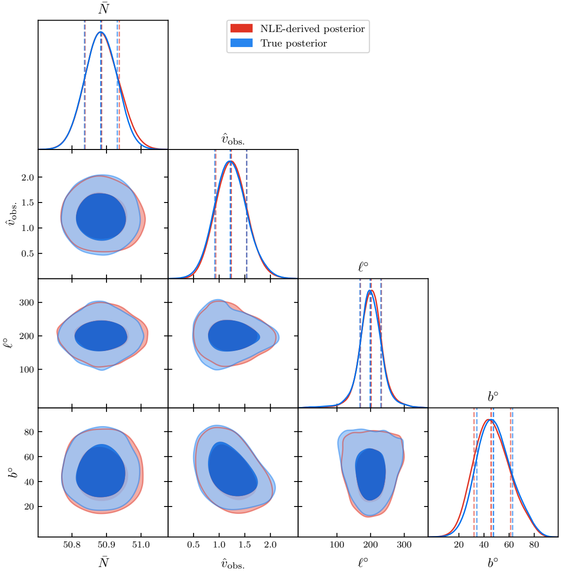

CatWISE2020... Again

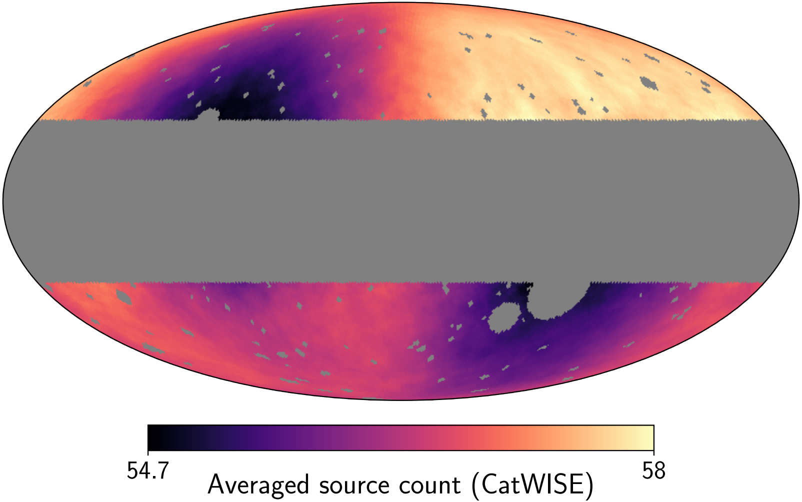

CatSIM Results

Model	$\ln B$
Dipole from Secrest+21	2.8 ± 0.9
Dipole from Dam+23	1.9 ± 0.8
Free dipole, extra error	0 ± 0
CMB direction, free velocity	-2.4 ± 0.9
CMB velocity & direction	-8.3 ± 0.9
Free dipole, no extra error	-113.6 ± 6.4

Table 1: Bayes factors w.r.t. fiducial model.

CatSIM Results

Oayda & Lewis (2025), submitted to MNRAS

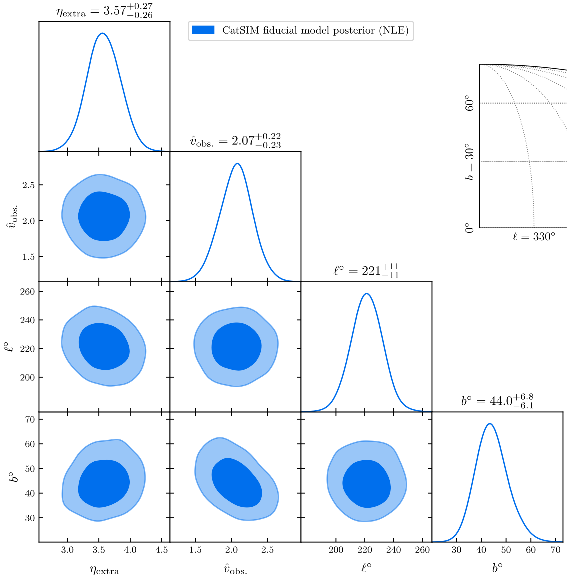

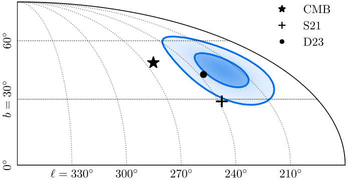

Posterior Predictive

Oayda & Lewis (2025), submitted to MNRAS

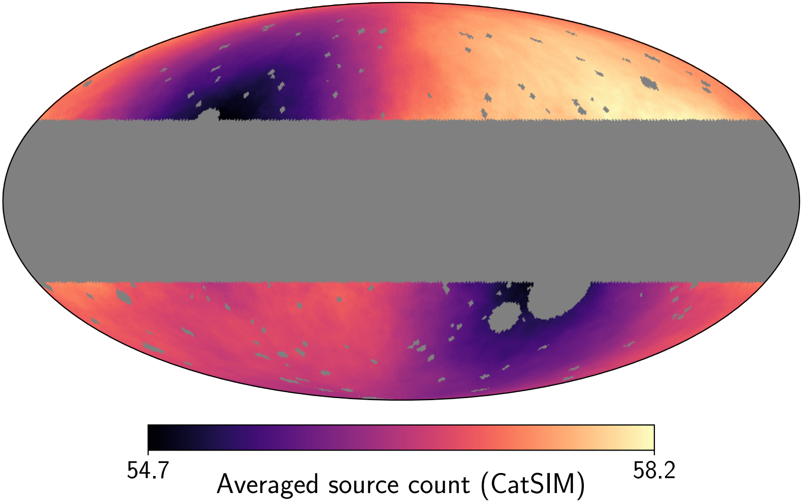

Posterior Predictive

Oayda & Lewis (2025), submitted to MNRAS

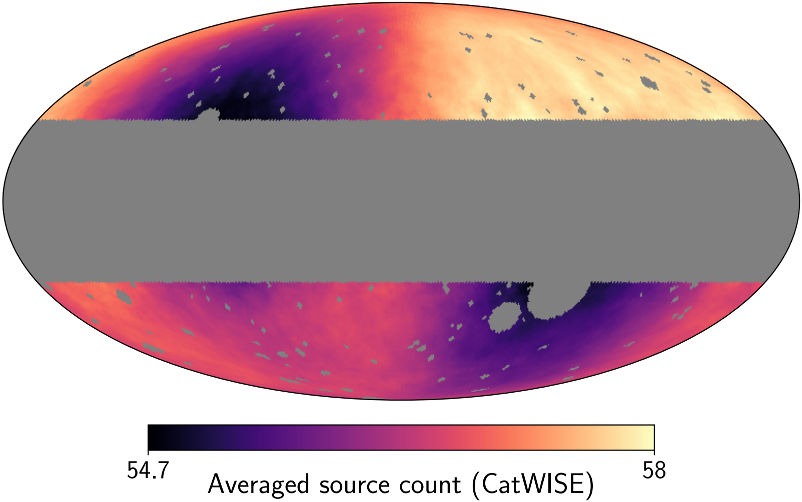

- **Describe and account for ecliptic bias** ✓

- **Describe and account for ecliptic bias** ✓
- **Measure cosmic dipole** ✓

- **Describe and account for ecliptic bias** ✓
- **Measure cosmic dipole** ✓
- **Leverage power of simulations and SBI** ✓