

Decoding Gravitational Waves: Why Model Matters

Results from LVK stochastic data and direct physical-parameter sampling

Liliana Velasco-Sevilla


Yonsei University, South Korea
S. AbdusSalam, J. Kersten, A. Zafari, A. Ojaghi
PPC2026, Sidney, Sept. 3, 2026

September 3, 2026



How the project evolved

1. **Original project:** obtain limits on the physical parameters of BSM models undergoing first-order phase transitions from stochastic gravitational-wave data.



Credit: NASA + LV

How the project evolved

1. **Original project:** obtain limits on the physical parameters of BSM models undergoing first-order phase transitions from stochastic gravitational-wave data.
2. **A statistical issue emerged:** many Bayesian analyses quote apparently precise limits without systematically asking how strongly those limits depend on the assumed prior.
3. **First test:** are the conclusions robust when the Bayesian analysis is performed in terms of geometrical GW parameters?
4. **Second test:** what happens when the inference is performed directly in thermodynamic and dynamical phase-transition parameters?
5. **Final test:** does the conclusion change if we bypass the intermediate GW parameterisation and sample directly in the microscopic physical parameters of a concrete model? → **Yes!**

Bayesian Inference Framework

Goal: Estimate parameters θ given observed data $\mathcal{D}$

Bayes' Theorem:

$$p(\theta \mid \mathcal{D}) = \frac{\mathcal{L}(\mathcal{D} \mid \theta) p(\theta)}{p(\mathcal{D})}$$

- $p(\theta \mid \mathcal{D}) =$ **posterior** — what we want to learn
- $\mathcal{L}(\mathcal{D} \mid \theta) =$ **likelihood** — probability of data given model
- $p(\theta) =$ **prior** — knowledge before observing data
- $p(\mathcal{D}) =$ **evidence** — normalization (independent of θ)

For parameter estimation: $p(\theta \mid \mathcal{D}) \propto \mathcal{L}(\mathcal{D} \mid \theta) p(\theta)$

The Likelihood Function

Basic Idea: The likelihood quantifies model-data agreement

Gaussian measurement errors: For each measured quantity C_i with true value μ_i and uncertainty σ_i :

$$p(C_i | \mu_i) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp \left[-\frac{(C_i - \mu_i)^2}{2\sigma_i^2} \right]$$

If measurements are independent across frequencies and detector pairs:

$$\mathcal{L}(\mathcal{D} | \theta) = \prod_{I < J} \prod_f p(C_{IJ}(f) | \theta)$$

Taking the logarithm to simplify products $\rightarrow$ sums:

$$\ln \mathcal{L} = \text{const} - \frac{1}{2} \sum_{I < J} \sum_f \frac{[C_{IJ}(f) - \mu_{IJ}(f)]^2}{\sigma_{IJ}^2(f)}$$

Application: Stochastic GW Background

Model prediction: The expected cross-correlation for a GW background is

$$\mu_{IJ}(f) = \Gamma_{IJ}(f) \Omega_{\text{GW}}(f; \theta)$$

where:

- $\Gamma_{IJ}(f)$ = detector response function (geometry + sensitivity)
- $\Omega_{\text{GW}}(f; \theta)$ = GW energy density spectrum (model parameters)

Substituting into the Gaussian likelihood:


$$\mathcal{L}(\theta) \propto \exp \left[-\frac{1}{2} \sum_{I < J} \sum_f \frac{[C_{IJ}(f) - \Gamma_{IJ}(f) \Omega_{\text{GW}}(f; \theta)]^2}{\sigma_{IJ}^2(f)} \right]$$

This is a **chi-squared likelihood** — penalizes residuals weighted by measurement noise
LVK stochastic constraints on FOPT

Conventional use of Bayes Analysis

For example from: "Cosmological and High Energy Physics implications from gravitational-wave background searches in LIGO-Virgo-KAGRA's O1-O4a runs" [2510.26848v2]: for constraining GW parameters as if were coming from FOP where the predominant source are GW produced by sound waves.


Parameter	Prior
α_{PT}	$\text{LogUniform}[10^{-2}, 10^2]$
β/H_{PT}	$\text{LogUniform}[1, 10^2]$
T_{PT}/GeV	$\text{LogUniform}[10^5, 10^{10}]$
Ω_{ref}	$\text{LogUniform}[10^{-13}, 10^{-6}]$
v_w	fixed at 1



Conventional use of Bayes Analysis

Then a 95% CL of is claimed: “our analysis sets a 95th the amplitude of the CBC background, Ω_{ref} of 2.0×10^{-9} .. for the sound waves...”

Parameter	Prior
α_{PT}	$\text{LogUniform}[10^{-2}, 10^2]$
β/H_{PT}	$\text{LogUniform}[1, 10^2]$
T_{PT}/GeV	$\text{LogUniform}[10^5, 10^{10}]$
Ω_{ref}	$\text{LogUniform}[10^{-13}, 10^{-6}]$
v_w	fixed at 1



But wait a minute...


The underlying formula used has many assumptions and correlations

So, is it valid to treat the parameters independently?

$$h^2 \Omega_{\text{SW}}(f) \simeq \left(1.86 \times 10^{-5}\right) v_w \left(\frac{H_{\text{PT}}}{\beta}\right) \left(\frac{\alpha_{\text{PT}} \kappa_{\text{SW}}}{1 + \alpha_{\text{PT}}}\right)^2 \left(\frac{100}{g_*}\right)^{1/3} \left(\frac{f}{f_{\text{SW}}}\right)^3 \left[\frac{7}{4 + 3(f/f_{\text{SW}})^2}\right] \Upsilon, \quad (1)$$

where Υ is a suppression factor due to the finite lifetime of the sound waves.

For example from 2003.07360, Ellis, Lewicki, No.




What is being tested?

Central question

How much of an inferred constraint comes from the **likelihood**, and how much comes from the **prior volume**?

Our analysis therefore is aimed at **get robust constraints by studying robustness of priors** developing three levels:



- Geometrical sampling tests robustness at the geometrical level.
- Thermodynamical sampling tests robustness after direct thermodynamic sampling.
- Direct sampling of model defining couplings tests robustness directly in the $U(1)$ microscopic model.

Parameterisation of the GW Signal

Parameters entering the analysis

Geometrical

Describe the observed GW spectrum directly in frequency space.

$$\Omega_{\text{peak}}, \quad f_{\text{peak}}, \quad \Omega_{\text{ref}},$$

Thermodynamical

Characterise the underlying first-order phase transition.

$$\alpha_{\text{PT}}, \quad \beta/H_{\text{PT}}$$

$U(1)$ Model

Connect the transition to the underlying particle-physics model.

Transition parameters

$$v_w, \quad \kappa_{\text{sw}}, \quad \epsilon, \quad H_{\text{PT}}$$

$$\lambda, \quad g, \quad v_{\text{ev}}$$

Sound Waves: Double Broken Power Law

Sound waves and MHD turbulence dominate GW production for $\alpha < 1$. The spectrum as a double broken power law (DBPL):

$$\Omega_{\text{GW}}^{\text{DBPL}}(f) = \Omega_2 \frac{S(f)}{S(f_2)},$$

with

$$S(f) = \mathcal{N} \left(\frac{f}{f_1} \right)^{n_1} \left[1 + \left(\frac{f}{f_1} \right)^{a_1} \right]^{-\frac{n_1 - n_2}{a_1}} \\ \times \left[1 + \left(\frac{f}{f_2} \right)^{a_2} \right]^{-\frac{n_2 - n_3}{a_2}}.$$

- $\mathcal{N}$: normalization, $\int S(f) d \log f = 1$.

LVK stochastic constraints on FOPT

The asymptotic behaviour is

$$\Omega_{\text{GW}}^{\text{DBPL}} \propto \begin{cases} f^{n_1}, & f \ll f_1, \\ f^{n_2}, & f_1 \ll f \ll f_2, \\ f^{n_3}, & f \gg f_2. \end{cases}$$

Parameters

f_1	first break frequency
f_2	second break frequency
n_1	low- f slope
n_2	intermediate slope
n_3	high- f slope
a_1, a_2	transition sharpness
Ω_2	amplitude at f_2

GW observables and their dependence

Geometrical Parameters

Directly describe the observed GW spectrum.

$$\Omega_2(\Omega_{\text{peak}}) \quad f_{\text{peak}}$$

Thermodynamical Parameters

Characterise the underlying first-order phase transition.

$$\Omega_2(\alpha_{\text{PT}}, \beta/H_{\text{PT}}, T_{\text{PT}}, v_w, \kappa_{\text{sw}}, \dots)$$
$$f_{\text{peak}}(\beta/H_{\text{PT}}, T_{\text{PT}}, v_w, \dots)$$

$U(1)$ Parameters

Connect the GW signal to the particle-physics model.

$$\Omega_2(\lambda, g', v_\phi, m_\phi, \dots)$$
$$f_{\text{peak}}(\lambda, g, v_\phi, \dots)$$

Prior choices used for the Geometrical analysis



Prior	Idea	Characteristic choice
Jeffreys	Scale invariant	log-uniform
Detector-based	Concentrate prob. where LVK is sensitive	$\log_{10} f_p \sim \mathcal{N}(\log_{10} 25, 0.8)$
Model-based	Motivated by physical FOPT populations	$\pi(\Omega_p) \propto \Omega_p^{-0.7}$
Mixed	Low-amplitude pref. + detector information	$\pi(\Omega_p) \propto \Omega_p^{-3}$

For the broad O4a geometrical scan:

$$\Omega_{\text{ref}} \in [10^{-13}, 10^{-6}], \quad \Omega_p \in [10^{-13}, 10^{-5}], \quad f_p \in [10^{-2}, 10^3] \text{ Hz.}$$

O4a prior specification: Table 3 and Sec. 4.4 of the manuscript.

Geometrical priors







Figure: Posterior distributions for the double broken power law that models the sound waves for four different priors: Jeffreys, Detector-Based Alternative, Model-Based and Mixed, respectively from top to bottom, for the O1-O4a stochastic LVK search

What changes when the prior changes?

Parameter	Jeffreys	Model-Based	Mixed	Detector-Based	Range
$\log_{10} \Omega_{\text{ref}}$	$[-10.0, -8.8]$	$[-9.8, -9.0]$	$[-9.9, -9.1]$	$[-9.9, -9.1]$	1.2
$\log_{10} \Omega_p$	$[-5.0, -1.0]$	$[-4.0, -2.5]$	$[-4.5, -3.2]$	$[-4.2, -2.8]$	4.0
$\log_{10} f_p$	$[-1.0, 2.5]$	$[0.8, 1.5]$	$[1.0, 1.4]$	$[0.9, 1.3]$	3.5
$\log_{10} \Delta_w$	$[-3.0, -0.5]$	$[-2.5, -0.3]$	$[-2.5, -0.2]$	$[-2.5, -0.2]$	2.8

Stable

$$\Omega_{\text{ref}}$$

is comparatively stable across priors.

Prior sensitive

$$\Omega_p, \quad f_p$$

remain strongly prior dependent.

Better sensitivity does not imply prior robustness.

Table 9 and the discussion of Fig. 2.

Direct sampling of thermodynamic parameters

General form

First, we need to check if the parameters we are sampling are appropriate (α , β/H_{PT} , T_{PT} , v_w).

Mean bubble separation

$$R_* H_* = (8\pi)^{1/3} \max(v_w, c_s) \frac{H_{\text{PT}}}{\beta},$$

where $c_s = \frac{1}{\sqrt{3}}$.

If we do not want to assume anything about v_w : the combination $R_* H_*$ is suitable for sampling rather than β/H_{PT} .

Sound-wave contribution

$$f_1 \simeq 0.2 H_{*,0} (H_* R_*)^{-1},$$

$$f_2 \simeq 0.5 H_{*,0} \Delta_w^{-1} (H_* R_*)^{-1},$$

with

$$\Delta_w = \frac{v_{\text{shell}}}{\max(v_w, c_s)}.$$

Weak phase transitions: $\alpha < 0.1$

The integrated GW energy density is

$$h^2 \Omega_{\text{int}} = h^2 F_{\text{GW},0} A_{\text{sw}} K^2 (H_* \tau_{\text{sw}}) (H_* R_*).$$

with

$$A_{\text{sw}} \simeq 0.11,$$

$$K \simeq 0.6 \kappa \frac{\alpha}{1 + \alpha},$$

$$H_* \tau_{\text{sw}} \simeq \min \left(1, \frac{2H_* R_*}{\sqrt{3K}} \right).$$

Key dependences

$$\Omega_{\text{GW}} \longrightarrow \{R_* H_*, \Delta_w, \alpha, \kappa\}$$

Sampled parameters:

$$\{\log_{10} K, \log_{10}(H_* R_*), \log_{10} \Delta_w, \log_{10}(T_*/\text{GeV}), \log_{10} \epsilon\}.$$

O4a sound-wave + MHD analysis, comparing Jeffreys and Conditional priors.





Figure: Posterior distributions for the combination of sound waves modeled by the double broken power law and the MHD contribution based on the template for two different priors: Jeffreys and Conditional.

Physical interpretation

- The transition is no longer described only by Ω_p and f_p .
- The sampled parameters directly control the physical ingredients of the GW signal:

$$K, \quad H_* R_*, \quad \Delta_w, \quad T_*, \quad \epsilon.$$

- Strong correlations appear because several combinations of these quantities generate very similar spectra.
- In particular, T_* and $H_* R_*$ participate in the determination of the characteristic frequency.
- ϵ is intrinsically difficult to constrain because the MHD component is generally subdominant to the sound-wave contribution.

Main lesson

Direct physical sampling exposes degeneracies that are hidden when the analysis is performed only in geometrical GW variables $\rightarrow$ Only possibly to resolve with a definite model.

From GW parameters to microscopic physics: U(1) example

The Abelian Higgs model provides a direct test of whether the prior dependence survives all the way to physical couplings.

Microscopic parameterisation

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + |D_\mu\phi|^2 - V(\phi), \quad D_\mu = \partial_\mu - igA_\mu, \quad V(\phi) = -\mu^2|\phi|^2 + \frac{\lambda}{2}|\phi|^4.$$

$$\{g, \lambda, v\}$$

with

$$m_X = g v, \quad m_h = \sqrt{2\lambda} v.$$

Instead of

$$\{\Omega_{\text{ref}}, \Omega_p, f_p, \Delta_w\} \longrightarrow \{\alpha, \beta/H_*, T_*, v_w\},$$


we perform the inference directly through

$$\boxed{\{g, \lambda, v\} \longrightarrow V_{\text{eff}}(\phi, T) \longrightarrow \text{nucleation} \longrightarrow v_w \longrightarrow \Omega_{\text{GW}}(f)}$$


- The finite-temperature potential is evaluated for each point.
- Nucleation is solved only for points passing the first-order-transition filter.
- Wall dynamics are subsequently calculated for successful nucleation points.

Direct sampling of the $U(1)$ model

Abelian $U(1)$: marginalized posterior — $v_w = 0.5$, $\xi_{\text{shell}} = 0.1$, $\varepsilon = 0.05$



Abelian $U(1)$: marginalized posterior — $v_w = 0.5$, $\xi_{\text{shell}} = 0.1$, $\varepsilon = 0.05$



Results for the direct sampling of fundamental parameters of the $U(1)$. Linear prior (left) vs logarithmic prior (right).

How do we quantify “robustness”?

A visual comparison of posterior contours is useful, but we want quantitative diagnostics.

Maximin Robustness: Sensitivity to Prior Choice

For the same data, repeat the Bayesian inference using different priors. A robust result should not depend strongly on the prior choice.

Interpretation

$\rho_{\max} \simeq 1 \quad \Rightarrow \quad \text{robust to the prior choice}$

$\rho_{\max} \gg 1 \quad \Rightarrow \quad \text{strong p. dependence}$

1. Maximin width ratio

$$\rho_{\max} = \frac{W_{\max}}{W_{\min}}.$$

A large value means that the inferred uncertainty changes substantially when the prior is changed.

For each prior P_i , define the width of the 68% credible interval: $W_i = \theta_{84\%} - \theta_{16\%}$,

$$\rho_{\max} \simeq 1; \Rightarrow; \text{robust to prior choice}$$

$$\rho_{\max} \gg 1; \Rightarrow; \text{strong p. dependence} \quad (2)$$

2. Posterior-mean discrepancy

$$\Delta_{\text{rel}} = \left| 1 - \frac{E[\theta|y, \pi_1]}{E[\theta|y, \pi_2]} \right|.$$

Robustness to Prior Choice: Credible-Interval Widths

Parameter	$\rho_{\max}$	Criterion	Robustness	Status
Ω_{ref}	2.0	2.0–5.0	✓	Moderate
Ω_p	6.5	> 5.0	55	NOT ROBUST
f_p	7.8	> 5.0	55	NOT ROBUST
Δ_w	1.2	< 2.0	✓	Robust

Maximin robustness ratio

$$\rho_{\max} = \frac{\text{Width}_{\max}}{\text{Width}_{\min}}$$

Computed from the 68% credible intervals across different priors.

$\rho_{\max} < 2 \Rightarrow \text{robust}$ $2 \leq \rho_{\max} \leq 5 \Rightarrow \text{sensitive}$ $\rho_{\max} > 5 \Rightarrow \text{prior dominated}$

Robustness to Prior Choice: Posterior Means

Parameter	Jeffreys	MB	Mixed	DB	Status
$\log_{10} \Omega_{\text{ref}}$	-9.40	-9.40	-9.50	-9.50	Robust
$\log_{10} \Omega_p$	-3.00	-3.25	-3.85	-3.50	NOT ROBUST
$\log_{10}(f_p/\text{Hz})$	+0.75	+1.15	+1.20	+1.10	NOT ROBUST
$\log_{10} \Delta_w$	-1.75	-1.40	-1.35	-1.35	Robust

Posterior-mean ratio

$$\rho_{\text{mean}} = \frac{\mathbb{E}[\theta|y, \pi_1]}{\mathbb{E}[\theta|y, \pi_2]}$$

Relative discrepancy

$$\Delta_{\text{rel}} = |1 - \rho_{\text{mean}}|$$

$$\Delta_{\text{rel}} > 0.2 \Rightarrow \text{prior dependence}$$

Amplitude and peak frequency are strongly prior dependent, while Ω_{ref} and Δ_w remain comparatively stable.

Comprehensive Robustness Assessment

Parameter	$\rho_{\max}$	Classification
Ω_{ref}	2.0	MODERATE
Ω_p	6.5	NOT ROBUST
f_p	7.8	NOT ROBUST
Δ_w	1.2	ROBUST

Classification criteria

ROBUST

$$\rho_{\max} < 2$$

$$\rho_{\text{overlap}} > 0.7$$

MODERATE

$$2 \leq \rho_{\max} \leq 5$$

$$0.4 < \rho_{\text{overlap}} \leq 0.7$$

NOT ROBUST

$$\rho_{\max} > 5 \text{ or}$$

$$\rho_{\text{overlap}} < 0.4$$

Main conclusion

Ω_{ref} $\longrightarrow$ **relatively robust**
 Ω_p and f_p $\longrightarrow$ **strongly prior dependent**

Multiple-prior analyses are therefore essential for reliable inference of the physical GW parameters.

The $U(1)$ robustness test

Parameter	$\rho_{\max}$	Δ_{rel}	Interpretation
g	1.9	0.14	ROBUST
$\log_{10}(v/\text{GeV})$	2.1	0.02	width sensitive
λ	2.0	0.36	MODERATE
$\log_{10} \Omega_{\text{ref}}$	2.5	0.01	comparatively stable

Width test

$$\rho_{\max}$$

The allowed width can change substantially with the prior.

Point estimate

$$\Delta_{\text{rel}} > 0.3$$

for λ , hinting an unreliable prior-independent point estimate.

Tables 14 and 15. The manuscript uses $\rho_{\max} > 2$ and $\Delta_{\text{rel}} > 0.3$ as indicative of substantial prior sensitivity.

Conclusions and what remains to be done

1. Geometrical sampling is useful, but not sufficient.

Geometrical sampling results show that O4a produces narrower posteriors, while Ω_p and f_p remain strongly dependent on the adopted prior.

2. Moving to physical parameters reveals additional degeneracies.

Posterior results of U(1) parameters direct sampling demonstrates that the mapping from GW observables to thermodynamic quantities is statistically non-trivial.

3. The U(1) test points towards direct physical sampling as the preferred strategy for final physical limits.

If the final goal is a bound on g, λ, ν , the prior should ultimately be tested in that physical parameter space rather than inferred indirectly from a prior imposed on GW variables.

Next steps:

- repeat direct physical sampling with a larger and more systematic set of priors;
- extend the analysis beyond the Abelian benchmark;
- propagate theoretical uncertainties in the GW templates, nucleation and wall dynamics;
- quantify robustness simultaneously in geometric, thermodynamic and microscopic parameter spaces;
- establish converged, reproducible limits with respect to sampler settings and numerical uncertainties.

Bottom line

A Bayesian limit is meaningful only after its robustness against reasonable prior choices has been demonstrated.