Bootstrapping gravitational EFTs at low impact parameter

Aninda Sinha

Centre for High Energy Physics, Indian Institute of Science, Bangalore, India
Quantum Horizons Alberta, Department of Physics and Astronomy,
University of Calgary, Canada

A stringy dispersion relation as a spectral microscope

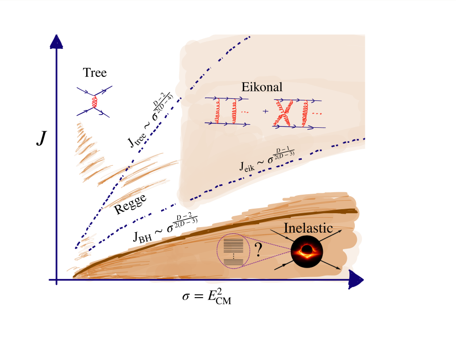

This talk grows out of a sequence of papers

THE PROGRAM

Crossing-symmetric dispersion relations

A. Sinha & A. Zahed, Crossing Symmetric Dispersion Relations in QFTs
Phys. Rev. Lett. **126** (2021) 181601; arXiv:2012.04877

String amplitudes as a field-theory guide

A. P. Saha & A. Sinha, Field Theory Expansions of String Theory Amplitudes
Phys. Rev. Lett. **132** (2024) 221601; arXiv:2401.05733

The stringy dispersion relation used today

F. Bhat, A. P. Saha & A. Sinha, A stringy dispersion relation for field theory
Phys. Rev. D **113** (2026) 066016; arXiv:2506.03862

Second half of this talk

D. Das & A. Sinha

Bootstrapping black holes at low impact parameter

arXiv:2607.05503

A numerical study of where the SDR extremal spectrum is supported.

The question is about the spectrum

THE PROBLEM

Once the universal large-impact-parameter eikonal carrier is supplied…

where does the remaining positive spectrum go?

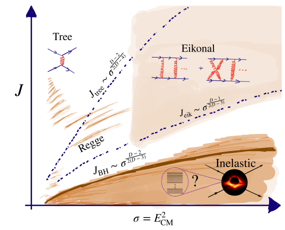

Six-dimensional impact-parameter ruler

$$\frac{b}{R_S} = 2\left(J+\frac{3}{2}\right)\left(\frac{16\pi^2}{3}\right)^{1/3}\sigma^{-2/3}\left(\frac{M_{\rm Pl}}{M_{\rm EFT}}\right)^{4/3}.$$

$$\sigma = E^2/M_{\rm EFT}^2.$$

Das–Sinha, arXiv:2607.05503; Giddings–Porto, Phys. Rev. D **81** (2010) 025002

The route and the answer

A MAP FOR THE NEXT 45 MINUTES

First 1/2

A new dispersion relation

→

Second 1/2

Use it to study "black holes" as intermediate states

The representation first; the spectrum it reveals second.

Caron-Huot–Van Duong, JHEP 05 (2021) 280; Caron-Huot–Mazáč–Rastelli–Simmons-Duffin, JHEP 07 (2021) 110 and JHEP 11 (2021) 164; de Rham–Tolley–Wang–Zhou, JHEP 01 (2026) 027.

Life before AI

Before AI,
hallucinations were called conjectures.

Dispersion relations turn analyticity into positive UV sums

THE GENERAL STRATEGY

Analyticity + crossing → **Cauchy integral** → **positive partial waves $\longrightarrow$ EFT inequalities**

$$\mathcal{M}_{\mathrm{low}} = \text{subtraction polynomial} + \sum_J \int d\sigma\, K_J(\sigma)\, \rho_J(\sigma), \qquad \rho_J \geq 0.$$

> **The bootstrap lives in the representation and in the kernel: the same amplitude can give very different numerical leverage.**

Adams–Arkani-Hamed–Dubovsky–Nicolis–Rattazzi, JHEP 10 (2006) 014; de Rham–Melville–Tolley–Zhou, Phys. Rev. D 96 (2017) 081702; Caron-Huot–Van Duong, JHEP 05 (2021) 280.

Gravity ruins the usual forward expansion

$$\mathcal{M}_{\rm low}(x,y) = -\frac{8\pi G_N x^2}{y} + g_2 x + g_3 y + \cdots .$$

$$s_1 + s_2 + s_3 = 0, \qquad x = -\sum_{i<j} s_i s_j, \qquad y = -s_1 s_2 s_3 .$$

$$\boxed{X \equiv \frac{g_2 M_{\rm EFT}^2}{8\pi G_N}, \qquad Y \equiv \frac{g_3 M_{\rm EFT}^4}{8\pi G_N}}$$

All later leaf plots use X horizontally and Y vertically; the numerical figures set $M_{\rm EFT} = 1$.

Massive EFT

Expand the fixed-t relation at $t = 0$. Every Wilson coefficient becomes a convergent positive moment.

Gravity

The graviton pole is singular at $t = 0$. The forward partial-wave expansion is no longer the right regulator.

A crossing-symmetric regulator avoids $t = 0$ while retaining the forward information.

Begin with the familiar fixed-t dispersion relation

WARM-UP

Fix $t < 0$ inside the analyticity domain. Our standing assumption throughout the next two- and three-channel derivations is

$$\boxed{\mathcal{M}(s,t) \longrightarrow 0 \quad (|s| \to \infty,\ t < 0 \text{ fixed}).}$$

We are therefore deriving the *unsubtracted* SDRs, for which the large arc vanishes:

$$0 = \frac{1}{2\pi i} \oint_{\Gamma} \frac{d\sigma}{\sigma - s}\, \mathcal{M}(\sigma, t).$$

Shrinking the large contour picks up the pole at $\sigma = s$ and the s-channel cut:

$$\boxed{\mathcal{M}(s,t) = \frac{1}{\pi} \int_{C_s} \frac{d\sigma}{\sigma - s}\, \mathcal{A}^{(s)}(\sigma, t).}$$

Good: the discontinuity is positive.

Missing: crossed-channel poles do not appear in this representation.

A lesser-known CSDR

The essential achievement

The Auberson–Khuri construction, revived in the fixed- a CSDR, makes all three channels manifest. It is a powerful framework for crossing and positivity.

At fixed

$$a = \frac{y}{x}, \qquad x = -(st + tu + us), \quad y = -stu,$$

the natural expansion contains terms $a^n x^m$.

Where locality becomes indirect

Returning to the EFT variables gives

$$a^n x^m = y^n x^{m-n}.$$

Terms with $n > m$ contain spurious negative powers of x. They cancel only after the locality, or null, constraints are imposed.

The original derivation therefore proceeds through a special parametrization, a dispersion relation in the new variable, the inverse map, and a final local subtraction.

> **Crossing is manifest, but the local representation is reached indirectly.**

Auberson–Khuri (1972); Mahoux–Roy–Wanders, NPB 70 (1974); Sinha–Zahed, PRL 126 (2021); Song, PRL 131 (2023); see also Elias Miró–Guerrieri–Gümüş–Zahed (2025).

Gravitational EFTs: Chang–Parra-Martinez (2025); Pasiecznik (2025); de Rham–Tolley–Wang–Zhou (2025); Peng–Rodina–Tokareva–Xu (2026).

One worldsheet, two field-theory channels

WORLDSHEET INTUITION

Two broad pair-of-pants pieces sewn into one four-point worldsheet

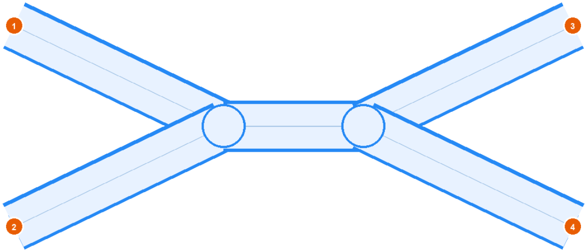

Plumbing fixture: cut, propagate, sew

WORLDSHEET INTUITION

Choose a long s-channel strip

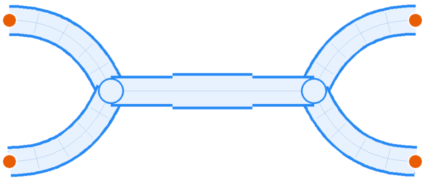

The long strip is the string propagator.

> **The same dispersion relation should also look like exchange channels plus a contact term.**

From fixed t to a moving crossed invariant

TWO CHANGES TO THE CONTOUR

Familiar fixed-t

$$0 = \frac{1}{2\pi i} \oint_{\Gamma_\infty} \frac{d\sigma}{\sigma - s} \mathcal{M}(\sigma, t).$$

The crossed invariant is held fixed. The Cauchy kernel exposes the s-channel pole, and the spectral data are organized as an s-channel discontinuity.

Let the crossed invariant move

$$0 = \frac{1}{2\pi i} \oint_{\Gamma_\infty} d\sigma \, H_\lambda^{(2)}(\sigma; s, t) \, \mathcal{M}(\sigma, f_\lambda(\sigma; s, t)),$$

$$\boxed{H_\lambda^{(2)} = \frac{1}{\sigma - s} + \frac{1}{\sigma - t} - \frac{1}{\sigma + \lambda}.}$$

The two physical poles are now symmetric. The auxiliary pole at $-\lambda$ preserves the $1/\sigma$ large-circle behaviour and will connect the fixed-t and fixed-s limits.

> **The kernel is chosen first. The moving function f_λ will be fixed by demanding that the two cuts agree.**

Bhat–Saha–Sinha, arXiv:2506.03862, Sec. 2

Three kernel poles and two amplitude cuts

SHRINK THE CONTOUR BEFORE SOLVING FOR F_λ

Kernel poles

$$f_\lambda(s) = t, \qquad f_\lambda(t) = s, \qquad f_\lambda(-\lambda) = \infty.$$

They contribute

$$\begin{aligned} \sigma = s &: \ \mathcal{M}(s,t), \\ \sigma = t &: \ \mathcal{M}(t,s) = \mathcal{M}(s,t), \\ \sigma = -\lambda &: \ -\mathcal{M}(-\lambda,\infty) = 0. \end{aligned}$$

The last equality uses the assumed large-energy falloff.

Amplitude cuts

I_1: the first argument σ crosses its physical cut.

I_2: the moving second argument $f_\lambda(\sigma; s,t)$ crosses its physical cut.

Shrinking the contour therefore gives

$$\boxed{2\mathcal{M}(s,t) = I_1 + I_2.}$$

> **Everything now reduces to one question: can the second cut be rewritten as the first?**

Bhat–Saha–Sinha, arXiv:2506.03862, Eqs. (2.1)–(2.7)

Choose the path so that the two cuts agree

THE CHANGE OF VARIABLES FIXES THE JACOBIAN

On the second cut, set

$$\tau = f_\lambda(\sigma; s, t).$$

Crossing exchanges the two arguments of the discontinuity. For the transformed integral to trace the first cut with the same weight, its measure must be unchanged.

$$\text{second cut} \xrightarrow{\tau=f_\lambda(\sigma)} \text{first cut}$$

The required measure identity is

$$\boxed{\frac{d\sigma}{d\tau}\, H_\lambda^{(2)}(\sigma; s, t) = H_\lambda^{(2)}(\tau; s, t).}$$

The channel interchange also requires

$$\tau = f_\lambda(\sigma), \qquad \sigma = f_\lambda(\tau).$$

Then crossing gives

$$I_2 = I_1,$$

and the contour relation becomes

$$2\mathcal{M}(s, t) = 2I_1.$$

> **The Jacobian condition is not an extra trick: it is exactly what makes the two cut integrals identical.**

Bhat–Saha–Sinha, arXiv:2506.03862, Eqs. (2.8)–(2.10)

The Jacobian determines the moving invariant

THE MÖBIUS MAP IS THE OUTPUT, NOT THE ANSATZ

Since

$$H_\lambda^{(2)}(\sigma; s, t) = \frac{d}{d\sigma} \log \frac{(\sigma - s)(\sigma - t)}{\sigma + \lambda},$$

the Jacobian equation first integrates to

$$\frac{(\sigma - s)(\sigma - t)}{\sigma + \lambda} = c\, \frac{(\tau - s)(\tau - t)}{\tau + \lambda}.$$

Interchanging σ and τ must describe the same change of variables; this fixes $c = 1$.

The quadratic equation has two roots:

$$\tau = \sigma \qquad \text{or} \qquad \boxed{\tau = f_\lambda(\sigma) = -\lambda + \frac{(s + \lambda)(t + \lambda)}{\sigma + \lambda}.}$$

The nontrivial root has exactly the required limits:

$$f_\lambda(s) = t, \qquad f_\lambda(t) = s,$$
$$f_\lambda(-\lambda) = \infty, \qquad f_\lambda(\infty) = -\lambda.$$

Only now do we notice the useful consequence

$$f_\lambda(f_\lambda(\sigma)) = \sigma.$$

Bhat–Saha–Sinha, arXiv:2506.03862, Eqs. (2.8)–(2.12)

Two cuts become one positive dispersion integral

THE TWO-CHANNEL SDR

$$2\mathcal{M}(s,t) = I_1 + I_2, \qquad I_2 = I_1,$$

hence

$$\boxed{\mathcal{M}(s,t) = \frac{1}{\pi}\int_{\sigma_0}^{\infty} d\sigma\, H_\lambda^{(2)}(\sigma;s,t)\,\mathcal{A}^{(s)}(\sigma, f_\lambda(\sigma;s,t))\,.}$$

$$\mathcal{A}^{(s)} \equiv \frac{1}{2i}\,\mathrm{Disc}_s\,\mathcal{M} = \mathrm{Im}\,\mathcal{M}(\sigma + i0, \hat{t}),$$

whose partial-wave coefficients obey

$$\rho_J(\sigma) \geq 0.$$

The familiar limits are built in:

$$\lambda = -t: \quad f_\lambda(\sigma) = t, \quad H_\lambda^{(2)} = \frac{1}{\sigma - s},$$

$$\lambda = -s: \quad f_\lambda(\sigma) = s, \quad H_\lambda^{(2)} = \frac{1}{\sigma - t}.$$

> **The auxiliary λ labels a continuous family of equal representations of the same amplitude.**

Bhat–Saha–Sinha, arXiv:2506.03862, Sec. 2

A simple example

A TWO-CHANNEL SDR WORKTHROUGH

Take the crossing-symmetric meromorphic amplitude

$$\mathcal{M}_\mu(s,t) = \frac{1}{(\mu - s)(\mu - t)}.$$

Its only s-channel pole sits at $\sigma = \mu$, with residue

$$\mathcal{A}^{(s)}_\mu(\sigma, \hat{t}) = \pi \frac{\delta(\sigma - \mu)}{\mu - \hat{t}}.$$

The SDR simply evaluates at $\sigma = \mu$:

$$\mathcal{M}_\mu(s,t) = \frac{H^{(2)}_\lambda(\mu; s, t)}{\mu - \hat{t}_\lambda(\mu)}.$$

Using the crossing map,

$$\mu - \hat{t}_\lambda(\mu) = \frac{(\mu + \lambda)^2 - (s + \lambda)(t + \lambda)}{\mu + \lambda},$$

while the kernel has precisely the same numerator:

$$H^{(2)}_\lambda(\mu; s, t) = \frac{(\mu + \lambda)^2 - (s + \lambda)(t + \lambda)}{(\mu - s)(\mu - t)(\mu + \lambda)}.$$

The same numerator appears above and below. Thus the apparent λ-dependence cancels before any sum is taken.

Elementary meromorphic check of the two-channel SDR; compare Bhat–Saha–Sinha, arXiv:2506.03862, Sec. 2

Veneziano check

THE SDR WARM-UP

$$\mathcal{V}(s,t) = \frac{\Gamma(-s)\Gamma(-t)}{\Gamma(-s-t)}.$$

The fixed-t relation gives the familiar s-pole expansion

$$\mathcal{V}(s,t) = \sum_{n=0}^{\infty} \frac{-1}{n!}\,\frac{(t+1)_n}{s-n}, \qquad \mathrm{Re}\,t < 0.$$

The t-channel poles are present, but only after analytic continuation.

The two-channel SDR instead gives

$$\mathcal{V}(s,t) = \sum_{n=0}^{\infty} \frac{-1}{n!}\left(\frac{1}{s-n} + \frac{1}{t-n} + \frac{1}{\lambda+n}\right) \times \left(1-\lambda+\frac{(s+\lambda)(t+\lambda)}{n+\lambda}\right)_n .$$

Every term displays the s- and t-channel poles. The auxiliary λ-dependence cancels only after the full sum.

Bhat–Saha–Sinha, arXiv:2506.03862, Eqs. (1.1), (1.2)

λ, not external t, controls Regge convergence

$$\mathcal{M}(s,t) \sim s^{\alpha(t)}, \qquad \alpha(t) = a_0 + a_1 t + a_2 t^2 + \cdots + a_n t^n.$$

Fixed t

$$\frac{\mathcal{M}(s,t)}{s^N} \longrightarrow 0 \quad \Longleftrightarrow \quad \mathrm{Re}\,\alpha(t) < N.$$

The external kinematics must lie in the decay domain.

Stringy dispersion relation

$$\widehat{t}_\lambda(\sigma) = -\lambda + O(1/\sigma), \qquad \mathcal{M}(\sigma, \widehat{t}_\lambda) \sim \sigma^{\alpha(-\lambda)}.$$

At large σ, the dispersive path approaches transfer $-\lambda$:

$$\mathrm{Re}\,\alpha(-\lambda) < N.$$

Veneziano: $\alpha(t) = t$. Fixed transfer needs $\mathrm{Re}\,t < 0$; the SDR needs only $\mathrm{Re}\,\lambda > 0$.

Choose λ in the convergence domain; the external kinematics can range much more widely. For gravity, $N = 2$.

Bhat–Saha–Sinha, arXiv:2506.03862, Secs. 2.1–2.2; the same asymptotic mechanism holds for the three-channel path

At fixed angle, the SDR truncation stays closer

A VENEZIANO CHECK AWAY FROM THE POLES

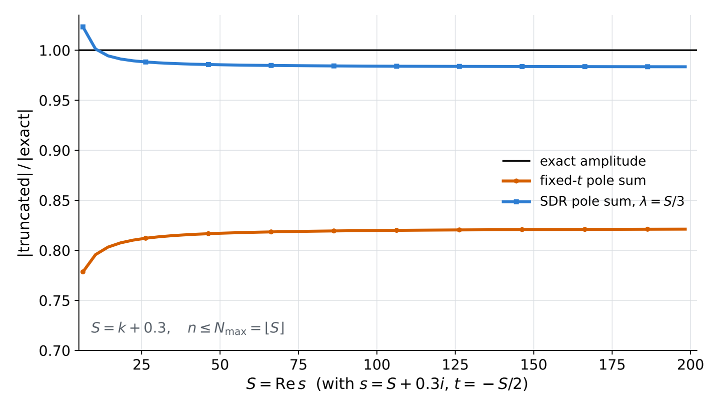

With the same levels $n \leq \lfloor S \rfloor$: fixed t, 82%; SDR with $\lambda = S/3$, 98%.

Direct evaluation of the two representations of the Euler beta function; $s = S + 0.3i$, $t = -S/2$

Fixed t fails in the unphysical fixed-angle quadrant

THE CONVERGENCE DOMAIN MATTERS

Move the Veneziano amplitude off its zero-width poles:

$$s = 12.3 + 0.3i, \qquad t = \frac{s}{2}.$$

A fixed-t summand behaves as

$$-\frac{(t+1)_n}{n!\,(s-n)} \sim \frac{n^{t-1}}{\Gamma(t+1)}.$$

Hence $\mathrm{Re}\,t > 0$ makes the partial sums run away.

By contrast, the SDR path obeys

$$\widehat{t}_\lambda(n) = -\lambda + O(n^{-1}),$$

and $\lambda = 1/2$ gives a convergent pole sum.

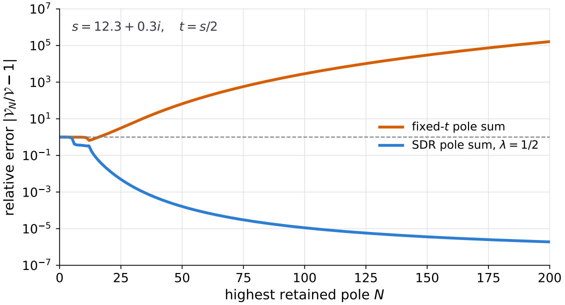

> **The $i\epsilon$ shift avoids poles but cannot repair fixed-t divergence; the SDR puts convergence under λ control.**

The three-channel stringy dispersion relation

THE SAME STRATEGY, NOW FULLY CROSSING-SYMMETRIC

The same contour and crossing argument, with a slightly more involved kinematic map, gives

$$\boxed{\mathcal{M}(s_1, s_2, s_3) = \frac{1}{\pi} \int_{C_1} d\sigma\, H_\lambda(\sigma; s_i)\, \mathcal{A}^{(s_1)}\big(\sigma, \hat{s}_2^{(\lambda)}(\sigma), \hat{s}_3^{(\lambda)}(\sigma)\big).}$$

Kinematics and physical cut

$$s_1 + s_2 + s_3 = 0,$$
$$x = -\sum_{i<j} s_i s_j, \qquad y = -s_1 s_2 s_3.$$

$\sigma \in C_1$, the physical s_1-channel cut, and
$\mathcal{A}^{(s_1)} \equiv \mathrm{Im}_{s_1} \mathcal{M}$ is the absorptive part on that cut.

Crossing-symmetric kernel

$$H_\lambda(\sigma; s_i) = \sum_{i=1}^{3} \frac{1}{\sigma - s_i} - \frac{1}{\sigma + \lambda}.$$

Moving crossed-channel coordinates

$$\hat{s}_2^{(\lambda)} + \hat{s}_3^{(\lambda)} = -\sigma,$$
$$\frac{\prod_{i=1}^{3}(s_i - \sigma)}{\sigma + \lambda} = \frac{\prod_{i=1}^{3}(s_i - \hat{s})}{\hat{s} + \lambda}.$$

The two nontrivial roots in $\hat{s}$ define $\hat{s}_{2,3}^{(\lambda)}(\sigma)$.

Auxiliary parameter

The auxiliary $\lambda > 0$ labels equivalent representations:
$\partial_\lambda \mathcal{M} = 0$.

Bhat–Saha–Sinha, arXiv:2506.03862, Sec. 3. See also Mahoux–Roy–Wanders, NPB 70 (1974); Sinha–Zahed, PRL 126 (2021); and Elias Miró–Guerrieri–Gümüş–Zahed, arXiv:2509.14170.

The SDR inside the geometric CSDR family

PROOF SUPPLIED BY OPENAI CODEX

Begin with an MRW linear leaf

$$y' = a(x' - x_0), \qquad x'(\sigma) = \frac{ax_0 - \sigma^3}{a - \sigma}.$$

The remaining roots of

$$z^3 - x'(\sigma)z + y'(\sigma) = 0.$$

give the two moving Mandelstam variables.
At the external point (x, y), choose

$$\boxed{a = -\lambda, \qquad x_0 = x + \frac{y}{\lambda}.}$$

The leaf is then the pencil of straight lines

$$\boxed{y' + \lambda x' = y + \lambda x,}$$

all passing through the same external point. Moreover,

$$x'(\sigma) - x = \frac{\sigma^3 - x\sigma + y}{\sigma + \lambda} = \frac{\prod_i(\sigma - s_i)}{\sigma + \lambda},$$

so its pulled-back Cauchy kernel is

$$\boxed{\frac{dx'}{d\sigma}\frac{1}{x'(\sigma) - x} = \sum_{i=1}^{3}\frac{1}{\sigma - s_i} - \frac{1}{\sigma + \lambda} = H_\lambda.}$$

> **The unsubtracted SDR is this special linear leaf. The geometric framework varies leaves to cover kinematics; the SDR varies λ through one fixed point, turning the same freedom into exact sum rules.**

SDR: Bhat–Saha–Sinha, arXiv:2506.03862. Earlier three-channel construction: Mahoux–Roy–Wanders, NPB 70 (1974); modern CSDR revival: Sinha–Zahed, PRL 126 (2021); general geometric framework: Elias Miró–Guerrieri–Gümüş–Zahed, arXiv:2509.14170. The identity displayed here follows by matching the unsubtracted pullbacks; subtractions require a corresponding matching of schemes.

Gravity requires the explicitly subtracted SDR

FROM THE DERIVATION TO THE BOOTSTRAP

$$\mathcal{M}(x,y) - \mathcal{M}(x_0,y_0) = -\frac{1}{\pi}\int_0^\infty d\sigma \left[\begin{aligned} &K_0(\sigma;x,y;\lambda)\,\mathcal{A}^{(s)}(\sigma,t_\lambda(x,y;\sigma)) \\ &\quad - K_0(\sigma;x_0,y_0;\lambda)\,\mathcal{A}^{(s)}(\sigma,t_\lambda(x_0,y_0;\sigma)) \end{aligned} \right],$$

$$(x_0,y_0) = (3\lambda^2,-2\lambda^3), \qquad K_0(\sigma;x,y;\lambda) = \frac{x-3\sigma^2}{y-x\sigma+\sigma^3} + \frac{1}{\sigma+\lambda}.$$

Allowed fixed-t growth

$\mathcal{M}(s,t) = O(|s|^{2-\epsilon})$**,** $\epsilon > 0$**, at fixed** $t < 0$**.**

If more subtractions are needed

Bhat–Saha–AS give the systematic higher-subtracted SDR construction.

This is the dispersive form projected onto g_2 and g_3 in the gravitational bootstrap.

Bhat–Saha–Sinha, arXiv:2506.03862; Das–Sinha, arXiv:2607.05503

Virasoro–Shapiro makes the gravitational pole explicit

THE SDR CHECK

$$\mathcal{M}_{\mathrm{VS}}(s_i) = x^2 \prod_{i=1}^{3} \frac{\Gamma(-s_i)}{\Gamma(1+s_i)}, \qquad s_1 + s_2 + s_3 = 0.$$

Its s_1-channel discontinuity is a discrete pole sum,

$$\mathcal{A}_{\mathrm{VS}}^{(s_1)} = -\pi \sum_{k=0}^{\infty} R_k(s_2, s_3)\, \delta(k - s_1).$$

Substitution into the three-channel SDR yields

$$\boxed{\begin{aligned} \mathcal{M}_{\mathrm{VS}} = {} & \frac{x^2}{y} + \frac{2x}{\lambda} + \frac{y}{\lambda^2} \\ & - \sum_{k=1}^{\infty} \left(\frac{1}{k - s_1} + \frac{1}{k - s_2} + \frac{1}{k - s_3} - \frac{1}{k + \lambda} \right) R_k\big(\hat{s}_2^{(\lambda)}, \hat{s}_3^{(\lambda)}\big). \end{aligned}}$$

The $k = 0$ state gives the explicit graviton pole. The remaining pole sum is all-channel, convergent away from its poles, and exactly λ-independent.

Bhat–Saha–Sinha, arXiv:2506.03862, Sec. 4; Virasoro–Shapiro check

One auxiliary parameter gives a whole family of constraints

FROM THE SDR TO THE BOOTSTRAP

Coefficient sum rules

$$\frac{\lambda\, K_2[\rho](\lambda)}{8\pi G_N} = 1 + 2\lambda X + \lambda^2 Y.$$

The same EFT data must be reproduced at every allowed λ.

Finite-amplitude differences

$$\Delta\mathcal{M}_{12} = \mathcal{M}(x_1, y_1) - \mathcal{M}(x_2, y_2).$$

The unknown subtraction constant cancels, so the same spectrum is tested at finite kinematics.

> **At finite resolution, sampled λ values are not repetitive data: together they constrain which positive partial-wave spectra can realize an EFT boundary.**

Bhat–Saha–Sinha, arXiv:2506.03862, Sec. 5; Caron-Huot–Van Duong, JHEP 05 (2021) 280; de Rham–Tolley–Wang–Zhou, JHEP 01 (2026) 027

Life after AI

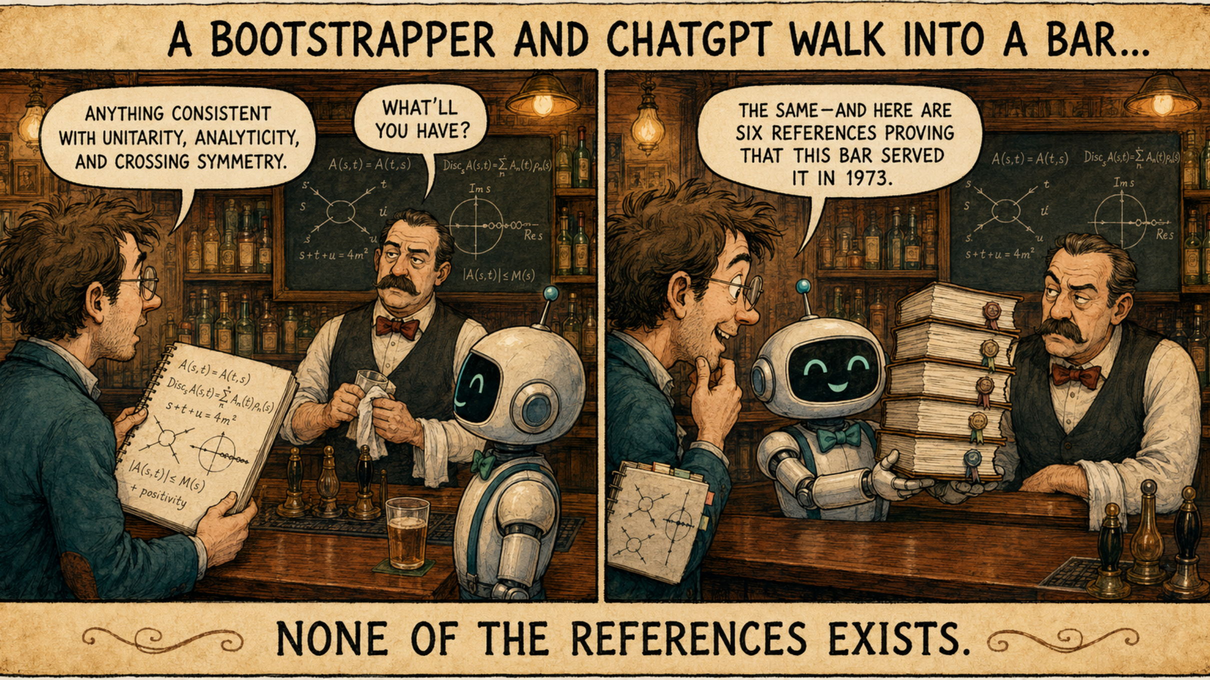

A BOOTSTRAPPER AND CHATGPT WALK INTO A BAR...
ANYTHING CONSISTENT WITH UNITARITY, ANALYTICITY, AND CROSSING SYMMETRY.
WHAT'LL YOU HAVE?
A(s,t) = A(t,s)
Disc_s A(s,t) = Σ_n A_n(t) ρ_n(s)
Im s
Re s
s+t+u = 4m²
|A(s,t)| ≤ M(s)
+ positivity
THE SAME—AND HERE ARE SIX REFERENCES PROVING THAT THIS BAR SERVED IT IN 1973.
NONE OF THE REFERENCES EXISTS.

Spectral Experiment

THE SPECTRAL MICROSCOPE

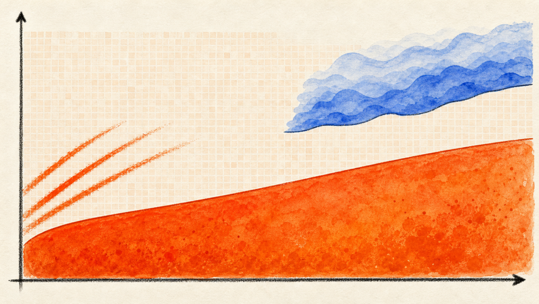

Blue cloud

Prescribed eikonal input.

Pale grid
Available residual variables.

Orange structures
Support selected by the LP.

Only the blue region is prescribed.

> **The LP can use any pale bin; the orange structures are its choice.**

Das–Sinha, arXiv:2607.05503; Häring–Zhiboedov, arXiv:2410.21499

Without the cap, the SDR wedge moves toward fixed-t

THE SPECTRAL MICROSCOPE

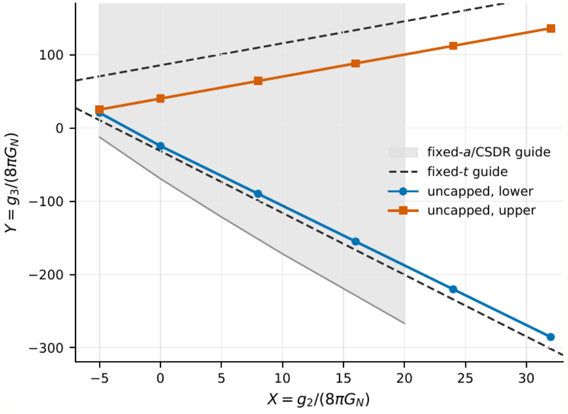

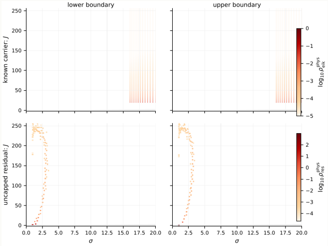

The complete eikonal carrier rotates the SDR wedge toward the fixed-t answer.

Gray comparison: fixed-a CSDR of Chang–Parra-Martinez and Peng–Rodina–Tokareva–Xu.

Das–Sinha, arXiv:2607.05503; Chang–Parra-Martinez, JHEP 08 (2025) 175; Peng–Rodina–Tokareva–Xu, arXiv:2604.15235

At weak gravity, the physical leaf hardly moves

THE SAME LEAVES IN TWO COORDINATE SYSTEMS

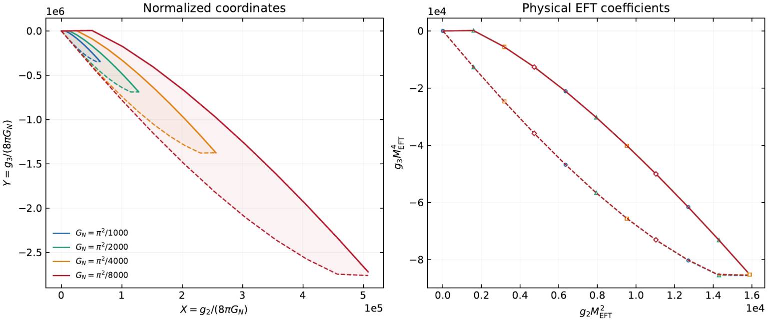

Normalized leaves expand as G_N falls; in physical g_2, g_3, they nearly coincide.

Das–Sinha, arXiv:2607.05503

The Giddings–Porto (GP) guide

A GEOMETRIC BENCHMARK, NOT AN INPUT TO THE OPTIMIZATION

In six dimensions

$$R_S(E) \sim (G_N E)^{1/3}.$$

A gravitational radius grows with energy.

Partial-wave kinematics

$$b \simeq \frac{2(J+3/2)}{E}, \qquad E = \sqrt{\sigma}.$$

Spin is an impact-parameter coordinate.

Rotating capture scale

$$b = \kappa R_J(E, J)$$
$$\Downarrow$$
$$J_{\mathrm{GP}} \propto G_N^{1/3} \sigma^{2/3}$$

The plotted $\kappa = 3$ curve fixes the order-one normalization.

The LP is never told this curve. It is a geometric ruler used only after the extremal spectrum has been found.

Giddings–Porto, Phys. Rev. D **81** (2010) 025002; D=6 impact-parameter kinematics used in Das–Sinha, arXiv:2607.05503

Turning on the physical cap makes a finite leaf

SCALE-REVERSED TOY MICROSCOPE: $M_{\mathrm{PL}} < M_{\mathrm{EFT}}$

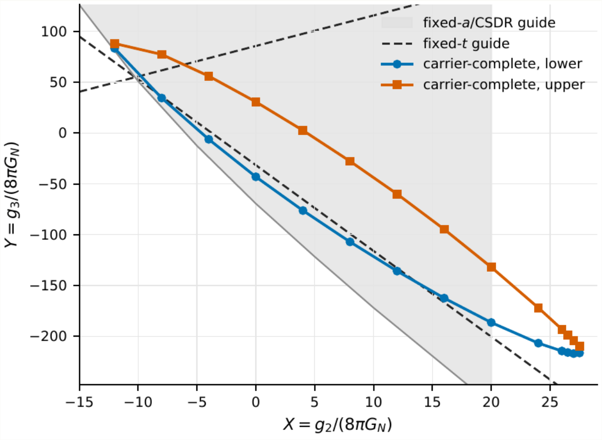

Das–Sinha, arXiv:2607.05503; Peng–Rodina–Tokareva–Xu, arXiv:2604.15235

Physical cap

$$0 < \rho_J(\sigma) < 2.$$

Deliberately scale-reversed

$$M_{\mathrm{Pl}} < M_{\mathrm{EFT}}.$$

This is a toy calculation, not a hierarchy-correct EFT. It pulls the black-hole radius into a tractable finite grid so we can ask where the extremal support goes.

The high-energy wedge has two distinct envelopes

COUPLING LADDER CROSSING $M_{\mathrm{PL}} = M_{\mathrm{EFT}}$

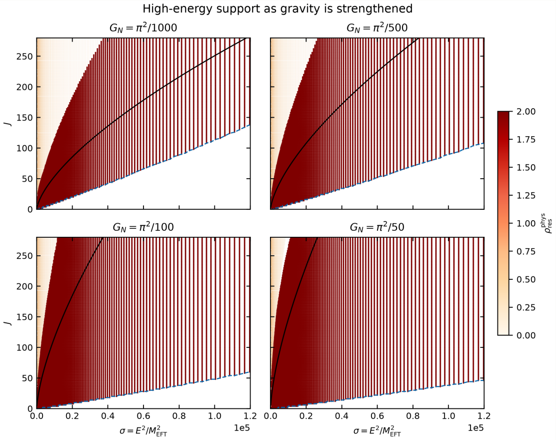

Outer envelope

$$J_{\mathrm{out}} \propto G_N^{1/3} \sigma^{2/3}$$

It has the rotating gravitational homogeneity.

Lower envelope

$$J_{\mathrm{low}} \simeq a(G_N)\sigma$$

A Regge-like boundary, not yet a reconstructed asymptotic trajectory.

Common high-energy coupling ladder, Das–Sinha, arXiv:2607.05503; Giddings–Porto, Phys. Rev. D **81** (2010) 025002

The spectrum reorganizes as we move along the leaf

SCALE-REVERSED TOY MICROSCOPE: $M_{\mathrm{PL}} < M_{\mathrm{EFT}}$

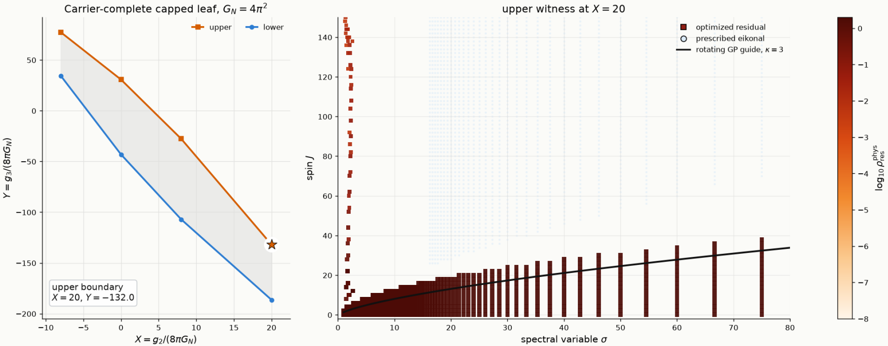

Blue is prescribed. Orange/red is selected by the optimizer. The empty region remains available.

Upper branch: the GP-bounded low-impact region grows. **Negative-X lower branch: it disappears.**

Carrier-complete 600 × 240 witnesses, Das–Sinha, arXiv:2607.05503; Giddings–Porto, Phys. Rev. D 81 (2010) 025002

Why the LP selects the low-impact band

TOY MICROSCOPE: $M_{\rm PL} < M_{\rm EFT}$

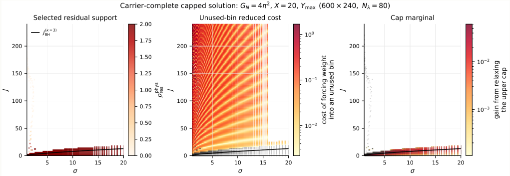

Force weight into a gap bin ⇒ $Y_{\max}$ decreases after global reoptimization.

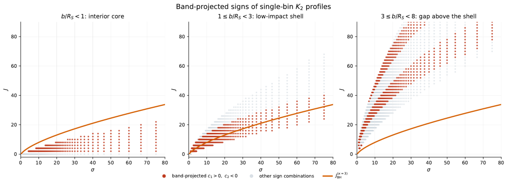

Single-bin profiles in the low-impact shell have the sign pattern needed by positive X.

> **The kernel profiles identify a promising shell before optimization; the reduced costs show that the LP actually uses that shell and avoids the gap.**

Reduced costs, cap marginals, and single-bin kernel profiles at $X = 20$, Das–Sinha, arXiv:2607.05503; Caron-Huot–Van Duong, JHEP 05 (2021) 280

Move the eikonal handover and watch a strand emerge

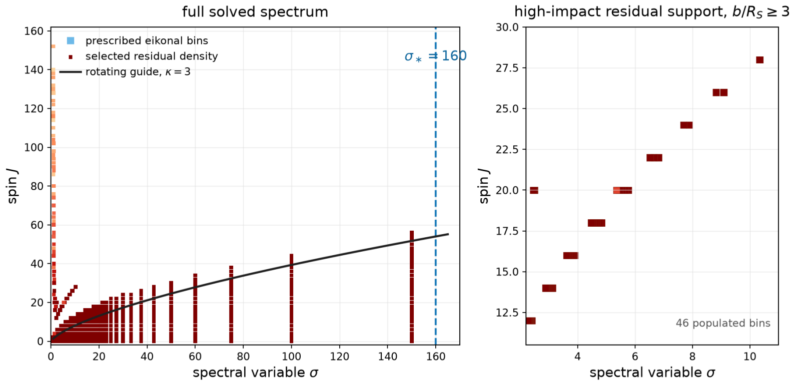

The analytic high-J tail remains in every frame. The emerging strand is organized Regge-like support, not yet a reconstructed continuous trajectory.

Carrier-handover control, Das–Sinha, arXiv:2607.05503

Conjecture/Hallucinations

HIERARCHY-CORRECT REGIME: $M_{\rm PL} > M_{\rm EFT}$

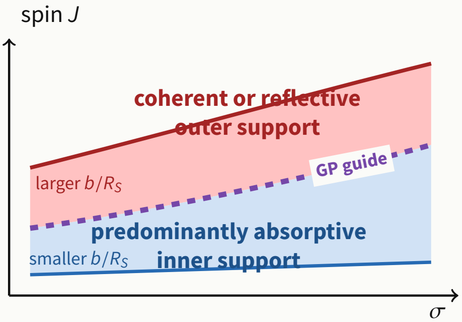

Above the GP guide
coherent or reflective

Below the GP guide
effectively absorptive

Horowitz–Polchinski transition?

Absorption of one elastic channel need not prove a horizon: unresolved horizonless states can also look absorptive.

Carrier-complete phase reconstruction will test this conjecture.

Conjecture motivated by the carrier-complete LP morphology; GP guide from Giddings–Porto, Phys. Rev. D **81** (2010) 025002.

What seems to survive all checks

SYNTHESIS

1	**Weak-gravity baseline**	The capped SDR problem has an organized microscopic branch even at strictly zero gravity.
2	**Coupling ladder**	The GP curve was not imposed, yet the outer envelope follows its $G_N^{1/3}\sigma^{2/3}$ homogeneity.
3	**Resolved strong-gravity toy**	A black-hole-scale band, Regge-like support, and an empty gap coexist.

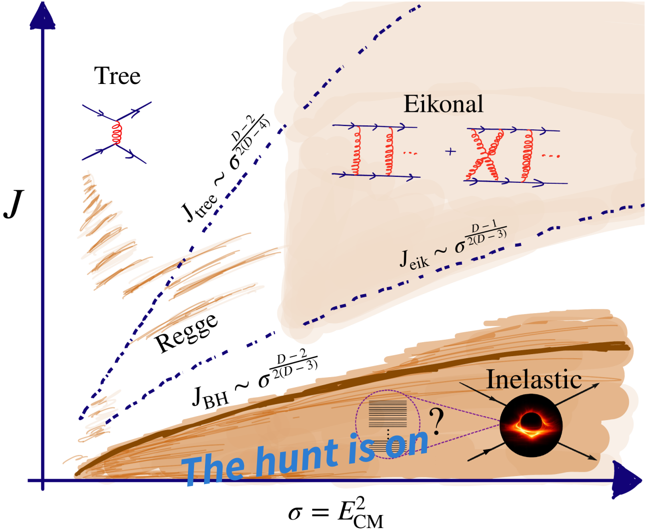

Tree
Eikonal
J
$J_{\text{tree}} \sim \sigma^{\frac{D-2}{2(D-4)}}$
$J_{\text{eik}} \sim \sigma^{\frac{D-1}{2(D-3)}}$
Regge
$J_{\text{BH}} \sim \sigma^{\frac{D-2}{2(D-3)}}$
Inelastic
?
The hunt is on
$\sigma = E^2_{\text{CM}}$

Thank you

Bootstrapping gravitational EFTs at low impact parameter

Supported by ANRF grant ANRF/ARG/2025/001338/PS.
Quantum Horizons Alberta.

Near-forward constraints select the bound and sharpen the band

BACKUP: WHAT DO THE CONSTRAINTS NEAR $\lambda = 0$ DO?

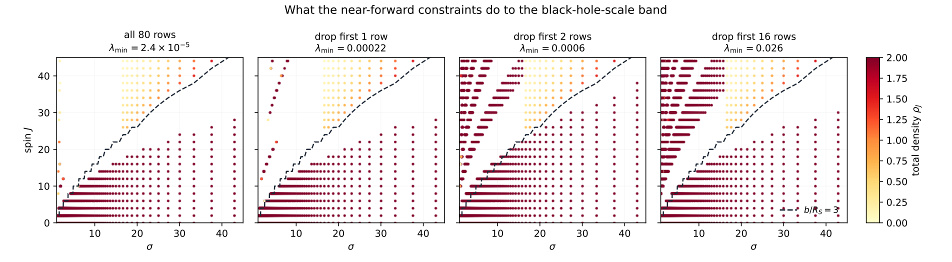

15.5% shift in Y after removing only the first row

98.4% of the near-cap support at $b/R_S < 3$ remains

$O(1)$ violation when the omitted row is tested afterwards

The near-forward rows are indispensable for the bound and suppress extra high-spin support, but the broader λ-family already organizes the black-hole-scale envelope.

Carrier-complete LP pilot at $G_N = 4\pi^2$**,** $X = 20$**,** $(N_\sigma, J_{\max}, N_\lambda) = (300, 160, 80)$

Controls expose what was input and what emerged

BACKUP: CARRIER CONTROLS

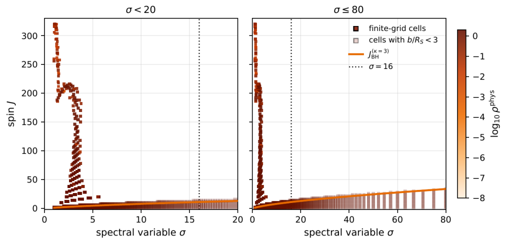

Remove the carrier

The pole is rebuilt by a low-mass, high-spin branch.

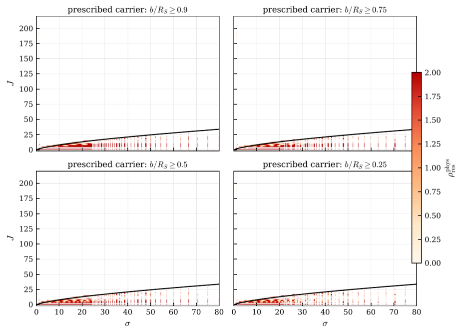

Move the carrier inward

The residual develops rapid density oscillations; a phase-proxy construction is only a proxy.

No-eikonal and inward-carrier controls, Das–Sinha, arXiv:2607.05503

Two complementary finite bootstrap problems

BACKUP: NUMERICAL SCALE

This SDR calculation

$N_\sigma = 2160$, $J = 0, 2, \ldots, 560$

606,960 **spectral variables**

200 **sampled** K_2 **sum rules**

plus 9 **FAD profiles when imposed**

sparse LP, solved directly with HiGHS

Fixed-t gravity bootstrap

200–400 **trial sum rules**

10^4–2×10^5 **sampled states**

positivity tested through $J = 400$

semidefinite functional search, solved with SDPB

The capped SDR formulation remains a sparse LP: no semidefinite solver or arbitrary-precision SDP machinery is required.

This work, representative 2160×560, $N_\lambda = 200$ witness; Caron-Huot–Li–Parra-Martinez–Simmons-Duffin, Phys. Rev. D 108 (2023) 026007, Appendix E

The weak band survives the strict zero-gravity null

HIERARCHY-CORRECT BASELINE: $M_{\mathrm{PL}} > M_{\mathrm{EFT}}$

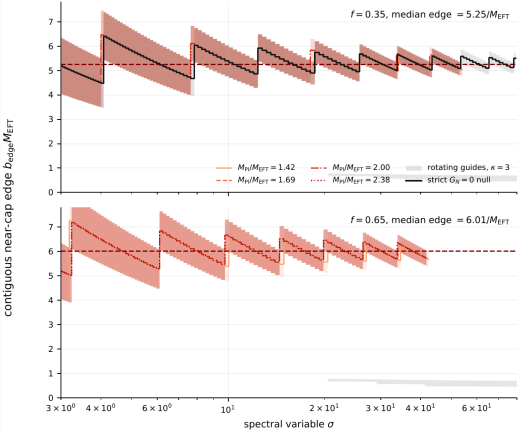

Matched physical-g_2, $G_N = 0$ comparison in Das–Sinha, arXiv:2607.05503

4.9%
RMS edge difference

0.988
support Jaccard
(populated spectral region is almost identical)

The weak band is an intrinsic baseline of the capped SDR problem on these grids.