TEVPA 2026 · TENDŌ, JAPAN · 3 SEPTEMBER 2026

Illuminating the Dark Sector

Confirming the 20 GeV Halo and Probing Its Mediator

T. R. Stenhouse

Department of Physics and Astronomy, University College London
Supervisors: Prof. Chamkaur Ghag · Prof. Frank F. Deppisch

Stenhouse, Ghag & Deppisch · arXiv:2607.08552 (v2 in preparation) | Stenhouse, Acar, Bashkanov, Deppisch, Ghag & Watts · arXiv:2608.29546

UCL

G. Bertone, D. Hooper & J. Silk, Phys. Rep. 405, 279 (2005) · G. Steigman, B. Dasgupta & J. F. Beacom, PRD 86, 023506 (2012) · J. L. Feng, Annu. Rev. Astron. Astrophys. 48, 495 (2010)

The problem of Dark Matter

- 85% of matter, but where is it?

Direct Detection

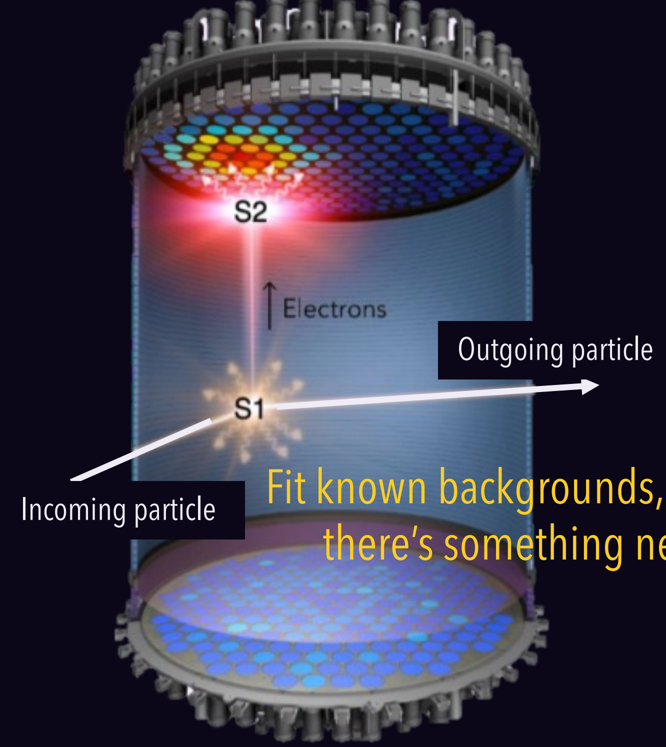

Fit known backgrounds, see if there's something new

S1+S2 gives characteristic signal

Indirect Detection

This is what we will focus on

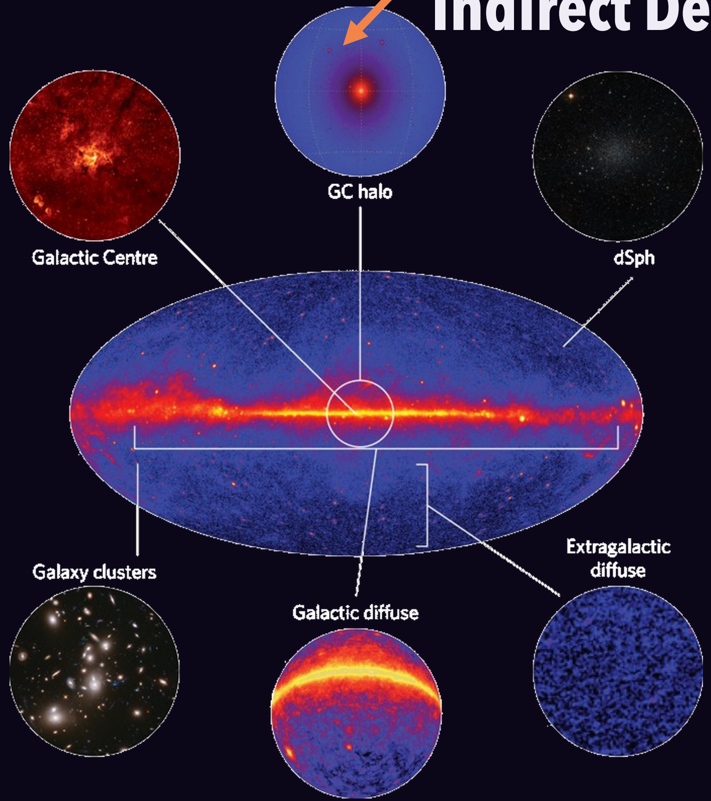

Fit known backgrounds, see if it needs something else

Collider Production

Measure what we can, see if there's something missing

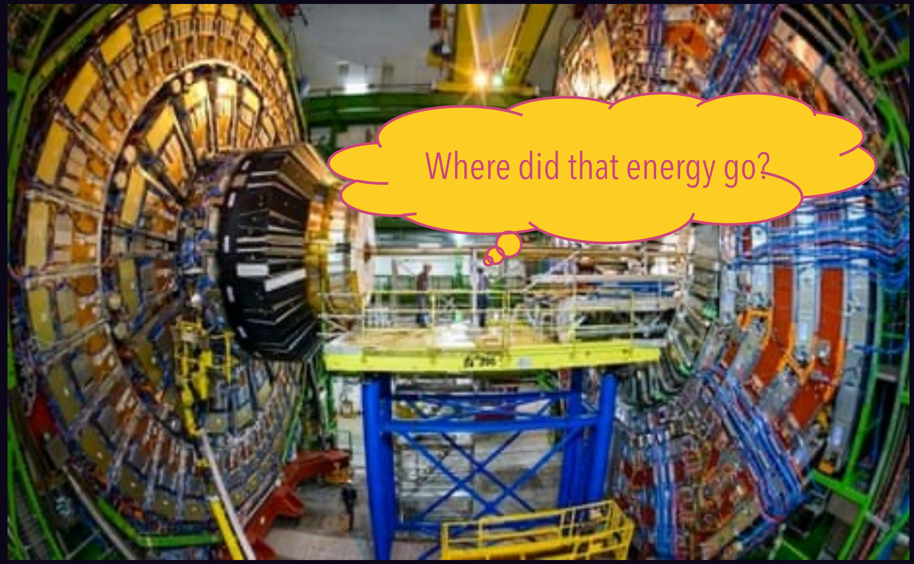

T. Totani, JCAP 11, 080 (2025) · L. Goodenough & D. Hooper, arXiv:0910.2998 · D. Hooper & L. Goodenough, Phys. Lett. B 697, 412 (2011) · K. N. Abazajian & M. Kaplinghat, PRD 86, 083511 (2012) · T. Daylan et al., Phys. Dark Univ. 12, 1 (2016) · F. Calore, I. Cholis & C. Weniger, JCAP 03, 038 (2015)

Where we started: Totani's halo-like excess

T. Totani, JCAP 11, 080 (2025) – 15 yr *Fermi*-LAT, $|\ell| \leq 60°$, $10° \leq |b| \leq 60°$. 10° cellwise binning. Residual after GALPROP + isotropic + Loop I + bubbles + point sources.

- The spherical, extended character is **not new**.
  - Hooper & Slatyer: Fermi bubble latitude bands → below $|b| \approx 20°$ approximately spherical component about Galactic centre, luminosity per volume tracking a squared generalised NFW.
  - Daylan et al: near-unit ellipticity out to at least 10°.
  - Calore, Cholis & Weniger: 10° lower limit on the extension.
- Totani's contribution is spectral rather than morphological: the high-latitude, **~20 GeV,** 13–19**σ** characterisation of a long-contested feature.
- NFW-compatible morphology, possibly slightly shallower toward the centre. Fitting prompt annihilation spectra: $m_\chi \sim 0.4$–0.8 TeV, $\langle\sigma v\rangle \sim (5$–$9)\times10^{-25}$ cm^3 s^{-1} ($b\bar{b}$ and W^+W^- together).

Figures here made using 17 yr data & pipeline of T.R.S. et al.

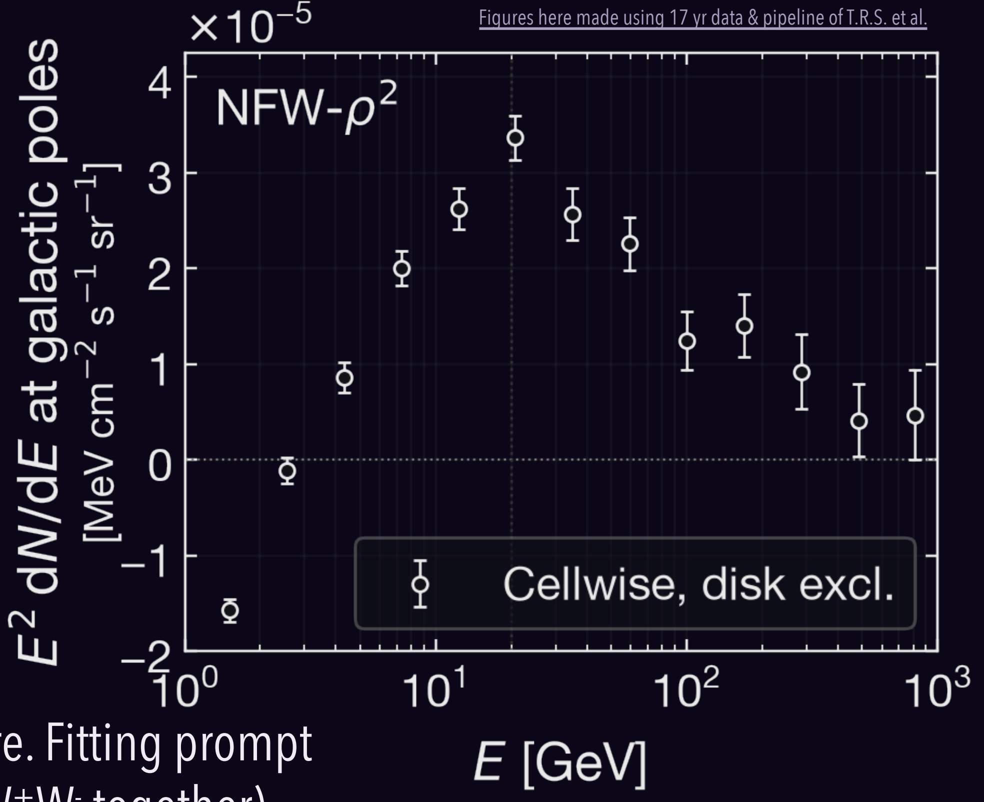

T. Totani, JCAP 11, 080 (2025) · L. Goodenough & D. Hooper, arXiv:0910.2998 · D. Hooper & L. Goodenough, Phys. Lett. B 697, 412 (2011) · K. N. Abazajian & M. Kaplinghat, PRD 86, 083511 (2012) · T. Daylan et al., Phys. Dark Univ. 12, 1 (2016) · F. Calore, I. Cholis & C. Weniger, JCAP 03, 038 (2015)

Where we started: Totani's halo-like excess

T. Totani, JCAP 11, 080 (2025) – 15 yr *Fermi*-LAT, $|\ell| \leq 60°$, $10° \leq |b| \leq 60°$. 10° cellwise binning. Residual after GALPROP + isotropic + Loop I + bubbles + point sources.

- The spherical, extended character is **not new**.
  - Hooper & Slatyer: Fermi bubble latitude bands → below $|b| \approx 20°$ approximately spherical component about Galactic centre, luminosity per volume tracking a squared generalised NFW.
  - Daylan et al: near-unit ellipticity out to at least 10°.
  - Calore, Cholis & Weniger: 10° lower limit on the extension.
- Totani's contribution is spectral rather than morphological: the high-latitude, **~20 GeV,** 13–19**σ** characterisation of a long-contested feature.
- NFW-compatible morphology, possibly slightly shallower toward the centre. Fitting prompt annihilation spectra: $m_\chi \sim 0.4$–0.8 TeV, $\langle\sigma v\rangle \sim (5\text{–}9)\times10^{-25}$ cm^3 s^{-1} ($b\bar{b}$ and W^+W^- together).

Figures here made using 17 yr data & pipeline of T.R.S. et al.

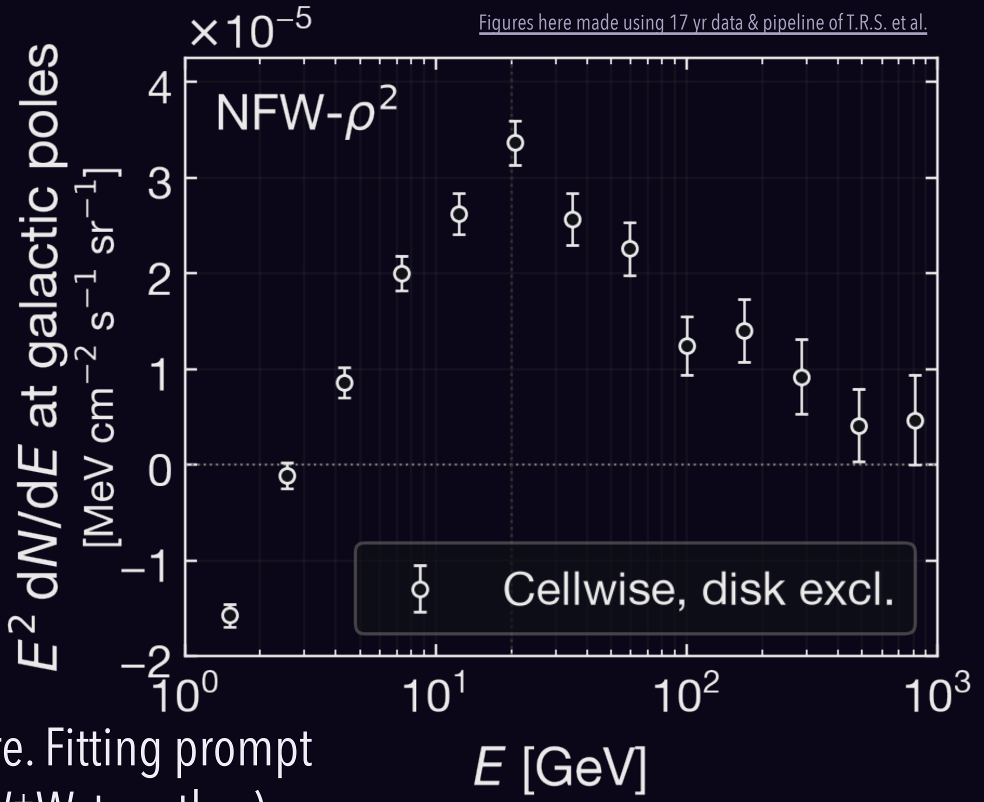

Is the excess just an artefact of 10° cellwise aggregation?

Two Independent Checks: Replicate & Extend

- Extended dataset: **17 yr** Pass 8 (Aug 2008 – Mar 2026), same 13 log bins, 1.16–10^{58} GeV.

(a) Replication

Independent pipeline, independent reduction, Poisson MCMC on $10° \times 10°$ cells – Totani's own method, re-implemented from scratch.

→

(b) Extension

Global pixel-level likelihood on native 0.125° maps. Adds energy-dependent LAT PSF forward folding (King profile, Pass 8) and bright 4FGL-DR4 source masking.

T. R. S., C. Ghag & F. F. Deppisch, arXiv:2607.08552

We recover it

- Cellwise and pixelwise agree on the **shape**, including the low-energy sign change.
- ρ^2 peak: **E ≈ 20.8 GeV**, E^2 dN/dE ≈ 4.0×10^{-5} MeV cm^3 s^{-1} sr^{-1} at poles.
- Pixelwise normalisation sits **~15-25%** above cellwise (≈20% for ρ^2) – a normalisation **offset**, not a shape difference.
  - Makes sense since cellwise spatial averaging naturally "smooths out" and washes down the cuspiness/gradient of a peaked halo profile.
- That offset propagates straight into the inferred ⟨**σ**v⟩ .

The feature isn't an artefact of 10° aggregation — PSF and masking barely move the high-latitude ρ² halo.

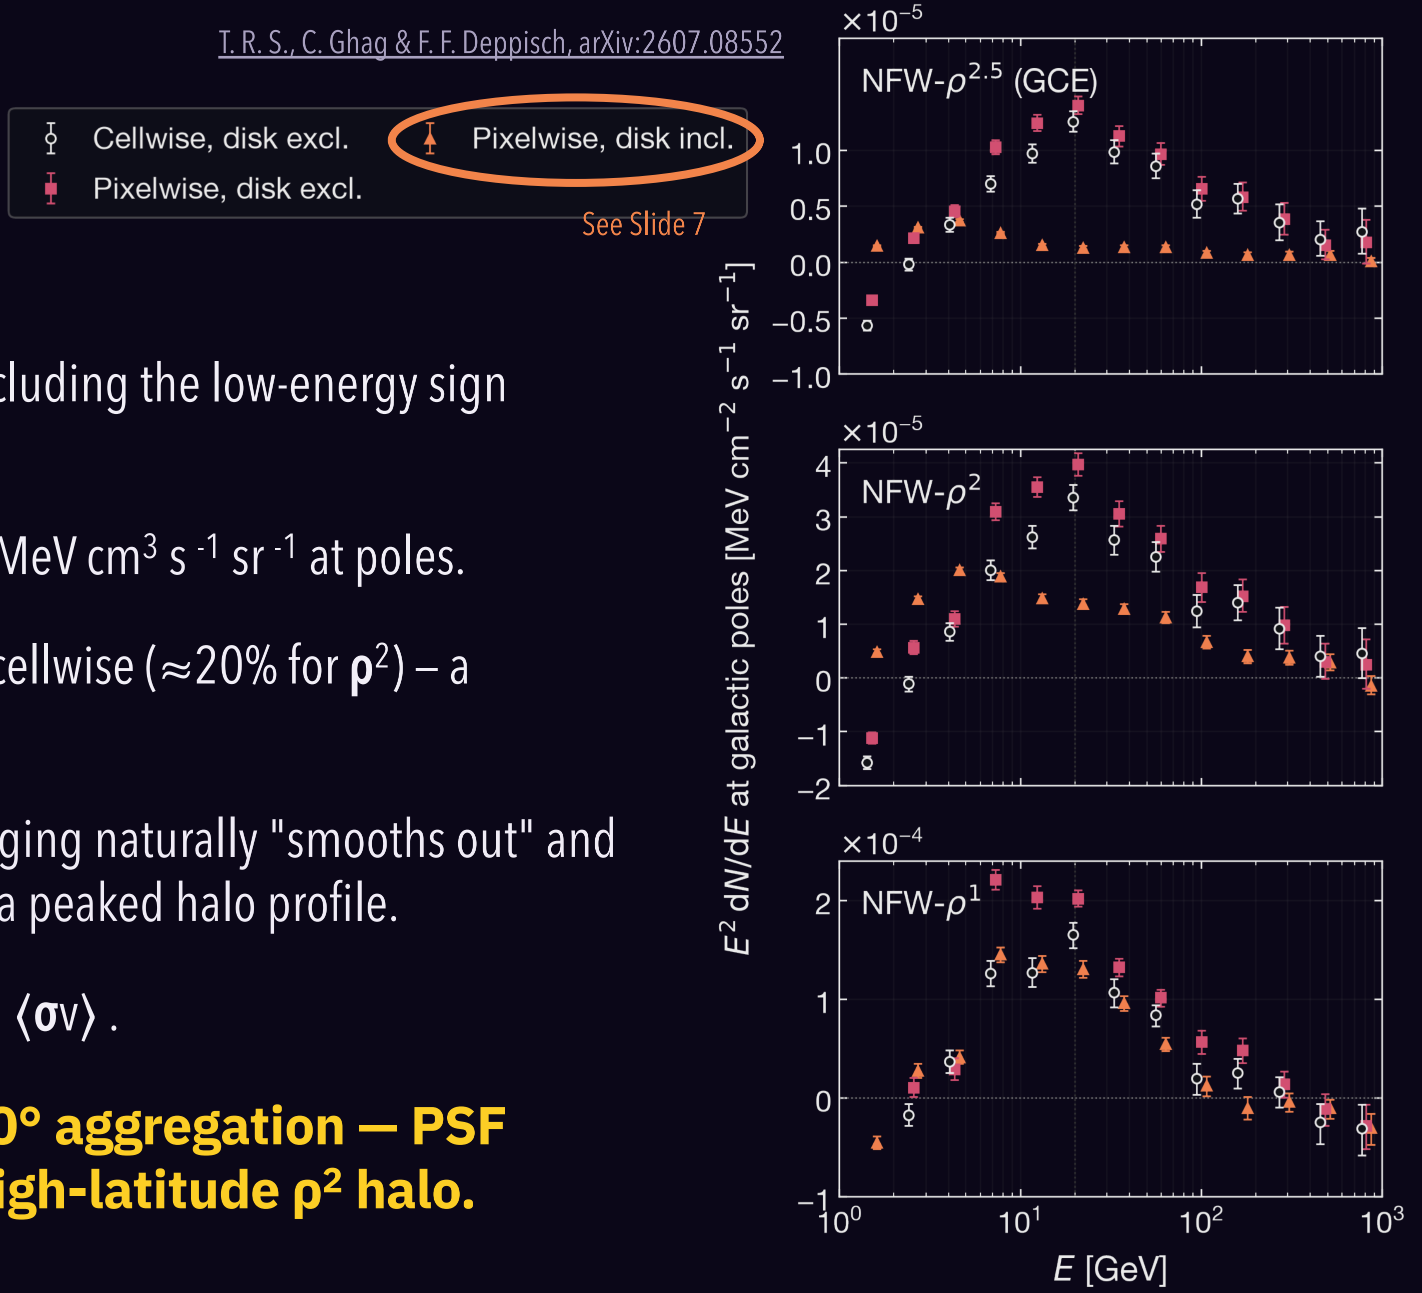

It’s structured, and it isn’t the GCE

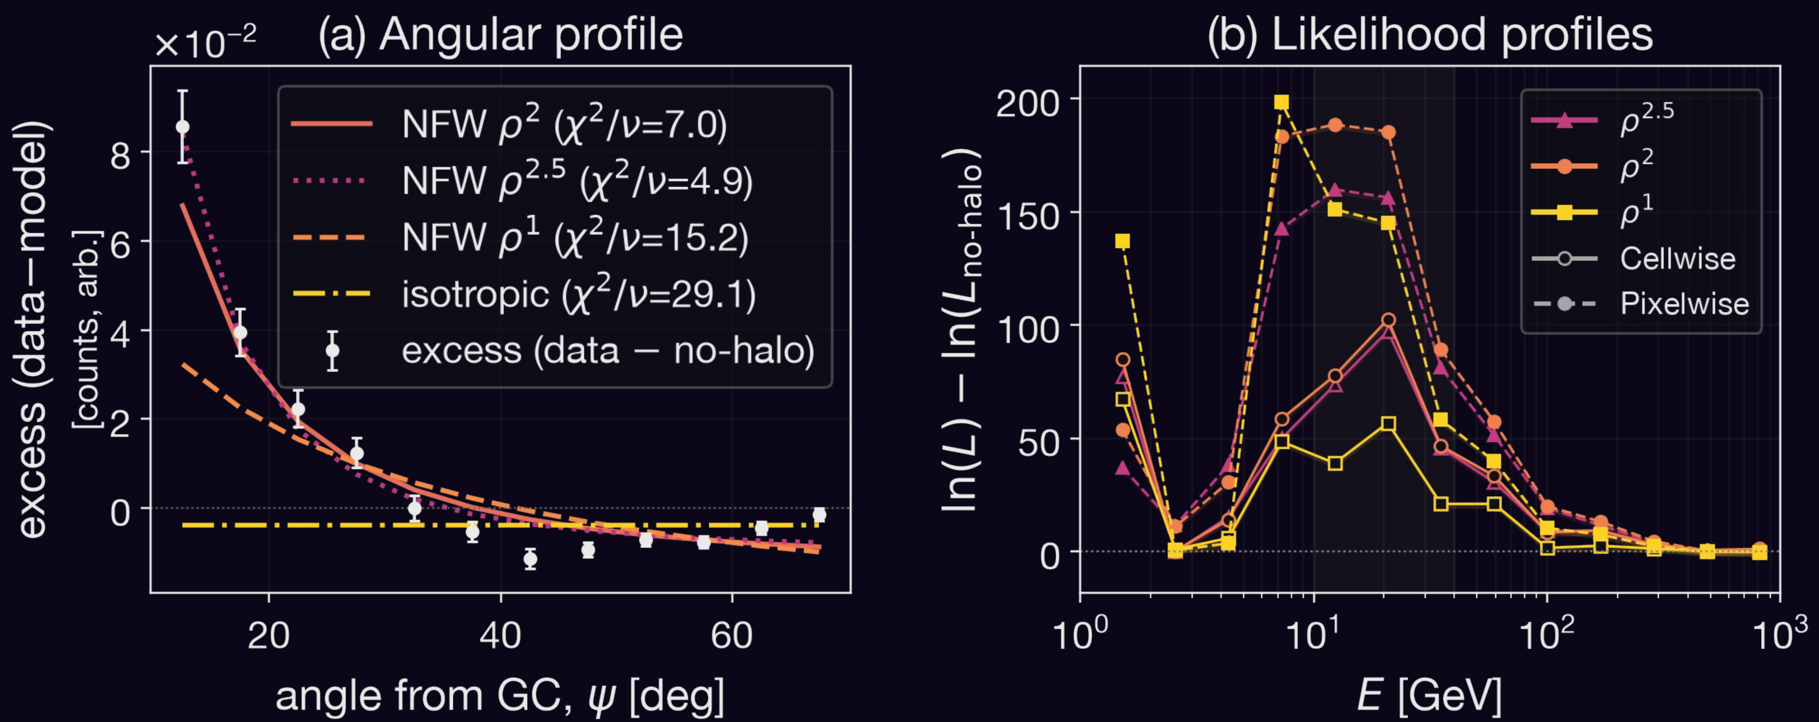

(i) Not isotropic, ρ^2 preferred by full pixel-level fit

Isotropic component is free in the fit. IGRB-like origin strongly disfavoured. $\boldsymbol{\rho}^2$ preferred over $\boldsymbol{\rho}^{2.5}$ by ≈220, over $\boldsymbol{\rho}^1$ by ≈200.

Honest caveat: no morphology is formally acceptable (best $\chi^2/\nu \approx 5$)

It's structured, and it isn't the GCE

T. R. S., C. Ghag & F. F. Deppisch, arXiv:2607.08552 · L. Goodenough & D. Hooper (2009) · R. Bartels, S. Krishnamurthy & C. Weniger, PRL 116, 051102 (2016) · F. Calore, I. Cholis & C. Weniger, JCAP 03, 038 (2015) · A. A. Abdo et al. (2PC), ApJS 208, 17 (2013)

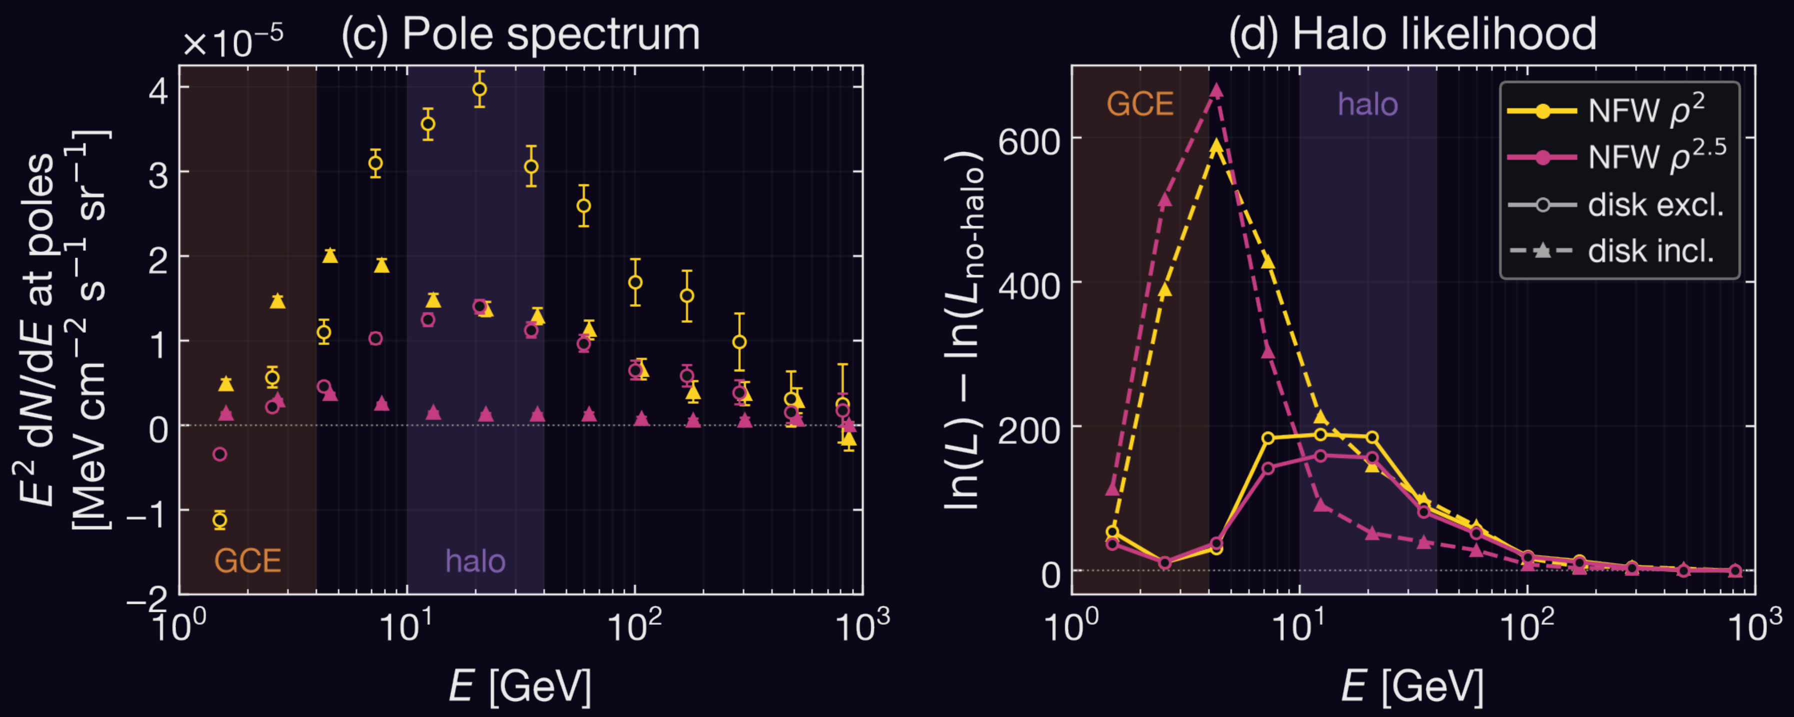

(ii) Not the GCE

Disk included ⟶ peak rolls to ~3 GeV; disk excluded ⟶ the 20 GeV halo. Likelihood in halo signal band (10–40 GeV) barely changes (≈920 vs ≈930). Spectrally, 20 GeV sits an order of magnitude above the GCE peak and well above the **E_c ~ 1–4 GeV** pulsar/MSP cutoff – **the two cannot be the same population.**

T. R. S., C. Ghag & F. F. Deppisch, arXiv:2607.08552 · M. Ackermann et al., ApJ 750, 3 (2012)

Could it just be cosmic ray electrons?

- Inject a hard population of **cosmic-ray electrons** in the inner Galaxy, propagate outward, compute the **IC emission** those electrons produce against ISRF

Syrovatskii Green's function, transport matched to the GALPROP model this analysis is built on, full Klein–Nishina – burst injection at **λ = 1.5/3/6 kpc**, **continuous**, **advected**, and a **CMB-only** ISRF bracket.

- Fit **alongside** halo: halo amplitude unchanged (ratio 1.000); injected coefficient consistent with zero. r(halo, injected) = −0.125 to −0.165 vs r(halo-ICS) =−0.48 in the same fit: the two are separable.

Angular profile alone does not discriminate (4.88, 5.09 vs 7.0 for ρ^2) – discrimination is in the full per-pixel likelihood.

- Forced to **replace** halo: disfavoured by 2ΔlnL ≈ 611–688; the 5 physical morphologies recover 5–21% of the measured halo peak, the CMB-only bracket 56%.

- The only configuration that dents the halo is the CMB-only bracket, requiring up to **4×10^68 erg, over two orders of magnitude** above the Galaxy's rest mass.

A cosmic-ray origin can be built. It just can't be made to fit.

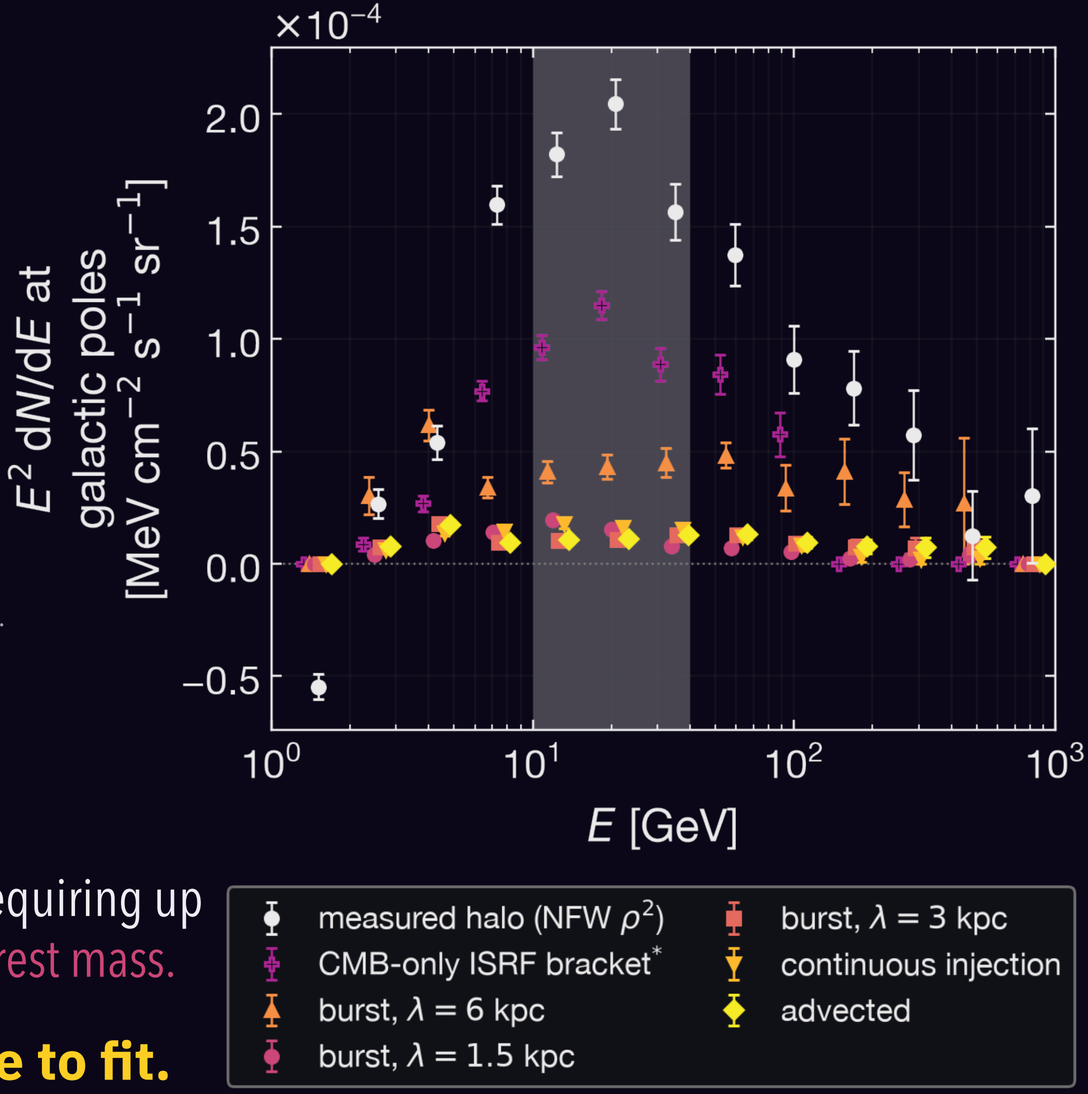

What actually limits this

- The fit is **systematics-limited**, not statistics-limited.
- Leading systematic: **template composition**.
  - GALPROP → GIEM-only: roughly 1/3 halo flux recovered. ICS/ISRF, PSF and masking variants shift things by ≲10%.
  - 14 GALPROP IEM models move halo normalisation **<5%,** some pull best-fit mass down (OB-star at $z_h = 6$: $m_\chi^{b\bar{b}} \approx 739 \to 369$ GeV)
- Two things it survives:
  - north/south split agrees to **1.4σ**;
  - injection-null test (no halo in the truth) recovers at most a 2σ amplitude, **~10% of observed peak**.

T. R. S., C. Ghag & F. F. Deppisch, arXiv:2607.08552 · M. Ackermann et al., ApJ 750, 3 (2012)

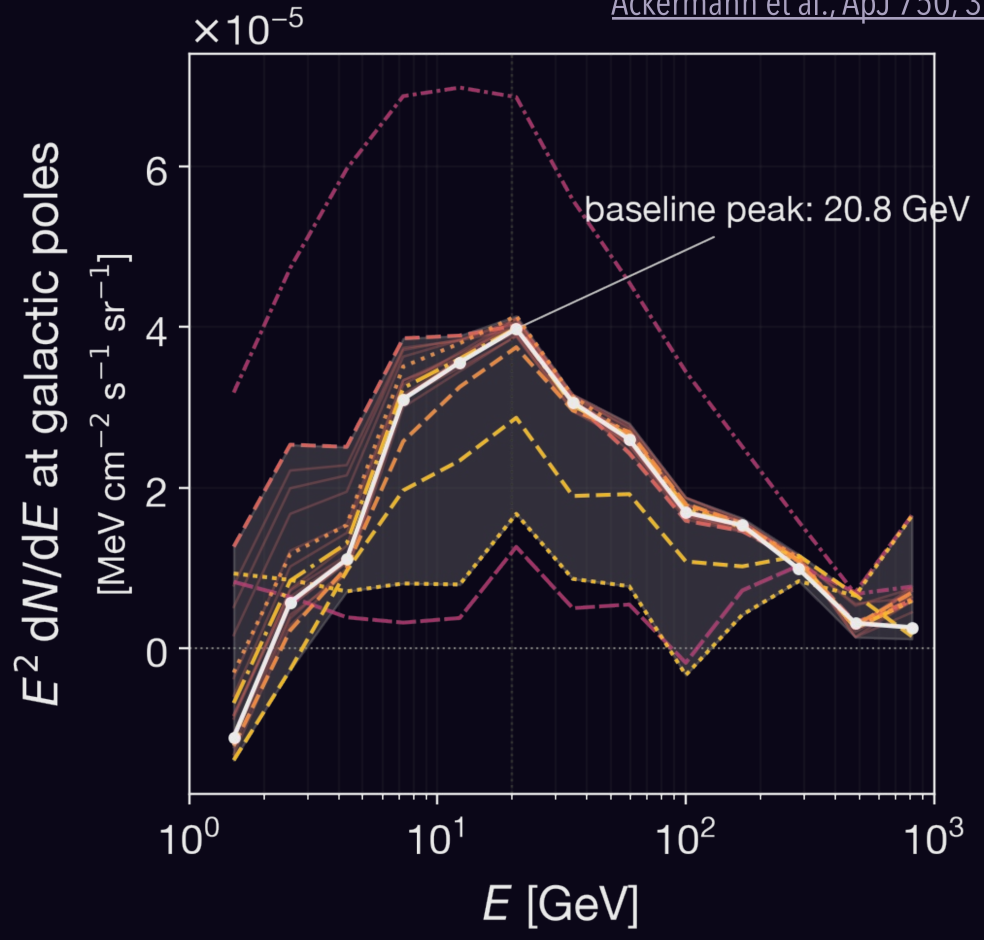

PSF and masking variants not shown; all within ≲10% of baseline
GIEM includes gas, ICS, bubbles and Loop I, so is fitted as a replacement for all of these.
Two stress tests are drawn but **excluded from the band.** GIEM with Loop I and bubbles retained (GIEM) → double-counts the same flux, suppresses the halo, pushes the band **down**. No bubbles → leaves large-scale flux with nowhere to go, so halo absorbs it (×1.831 amplitude)

If it is dark matter: mass and cross section

T. R. S., C. Ghag & F. F. Deppisch, arXiv:2607.08552 · M. Cirelli et al., JCAP 03, 051 (2011) PPPC4DMID

Best fits of prompt PPPC annihilation spectra to the global NFW-ρ^2 halo:

channel	ours	Totani (2025)
W^+W^-	0.56 TeV · 1.0×10^{-24} cm^3 s^{-1}	0.4–0.8 TeV · $(5–9)\times10^{-25}$ cm^3 s^{-1}
$b\bar{b}$	0.72 TeV · 9.8×10^{-25} cm^3 s^{-1}	
$\tau^+\tau^-$	70 GeV · $\chi^2/\nu \approx 14.6$ (disfavoured, as in his analysis)	40–80 GeV · $(7–9)\times10^{-26}$ cm^3 s^{-1}

- **$\chi^2/\nu \approx 7.5$ ($b\bar{b}$), 7.4 (W^+W^-): poor fit**
  - Not unexpected: the photon-rich bins below 5 GeV carry the statistical weight, so small fractional residuals in the diffuse model dominate the χ^2 sum.
  - **Not poor dark matter spectral modelling.** Mass well localised within the baseline foreground model.

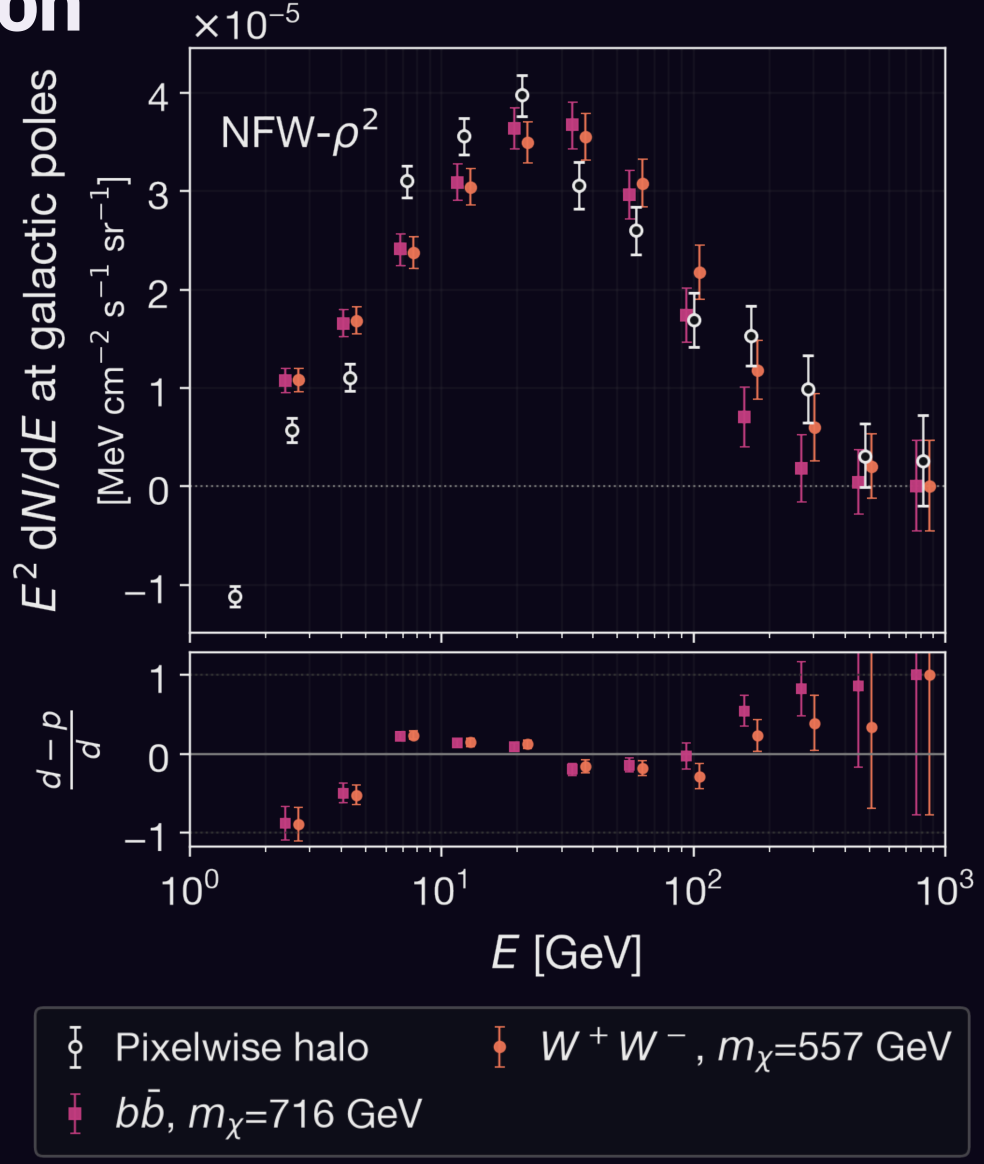

The dwarf problem

S. Hoof, A. Geringer-Sameth & R. Trotta, JCAP 02, 012 (2020) · A. McDaniel et al., PRD 109, 063024 (2024) · K. K. Boddy et al., PRD 102, 023029 (2020)

- $R \equiv \langle\sigma v\rangle_{halo} / \langle\sigma v\rangle_{dSph}^{UL}$. Two limit sets: Hoof+ (11 yr, 27 dwarfs), McDaniel+ (14 yr stacked).
- Baseline: **$R \approx 5.1$ ($b\bar{b}$), 4.1 (W^+W^-)** – consistent with Totani's "several times larger than the 95% containment region"
- Full systematic budget (foregrounds ⊕ halo J-factor ⊕ dwarf J-factors) **widens** the window to **$R_{b\bar{b}} \approx$ 1.6–9.8 about 5.1** & **$R_{WW} \approx$ 1.2–7.7 about 4.1**
  - Not excluded: lower edge within reach of the dwarf limit.
- **Mass freedom does not save you** – within the 95% CL mass window, min(R) stays above 4.

foreground range ×0.375–1.320 ($b\bar{b}$) ⊕ halo J-factor ≈48% ($\rho_\odot$ = 0.42±0.10) ⊕ dwarf J-factors ≈54%

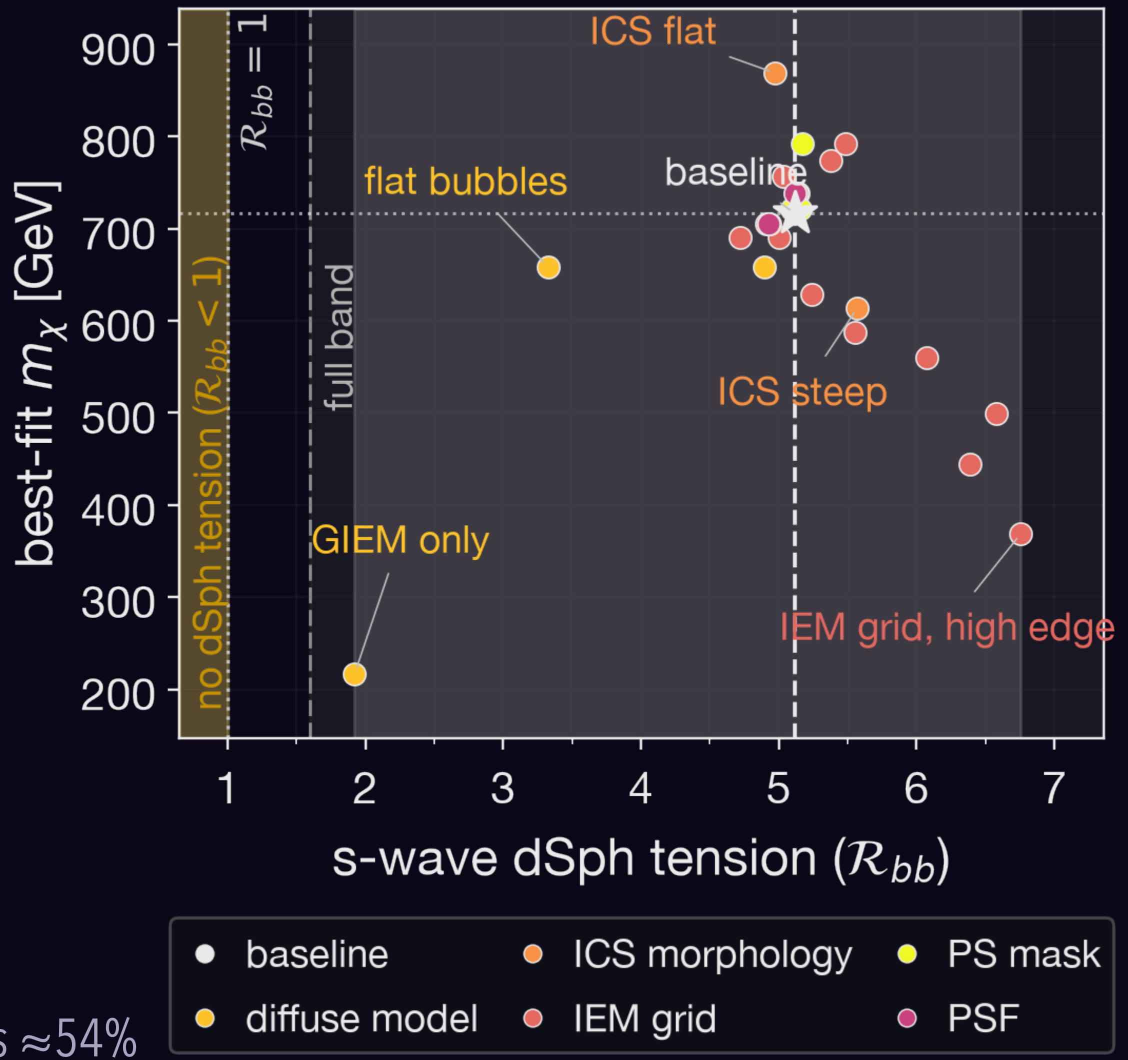

The dwarf problem

S. Hoof, A. Geringer-Sameth & R. Trotta, JCAP 02, 012 (2020) · A. McDaniel et al., PRD 109, 063024 (2024) · K. K. Boddy et al., PRD 102, 023029 (2020)

- $R \equiv \langle\sigma v\rangle_{halo} / \langle\sigma v\rangle_{dSph}^{UL}$. Two limit sets: Hoof+ (11 yr, 27 dwarfs), McDaniel+ (14 yr stacked).
- Baseline: **$R \approx 5.1$ ($b\bar{b}$), 4.1 (W^+W^-)** – consistent with Totani's "several times larger than the 95% containment region"
- Full systematic budget (foregrounds ⊕ halo J-factor ⊕ dwarf J-factors) **widens** the window to **$R_{b\bar{b}} \approx$ 1.6–9.8 about 5.1** & **$R_{WW} \approx$ 1.2–7.7 about 4.1**
  - Not excluded: lower edge within reach of the dwarf limit.
- **Mass freedom does not save you** – within the 95% CL mass window, min(R) stays above 4.

If it's DM, the tension has to be addressed in the particle model, not the fit.

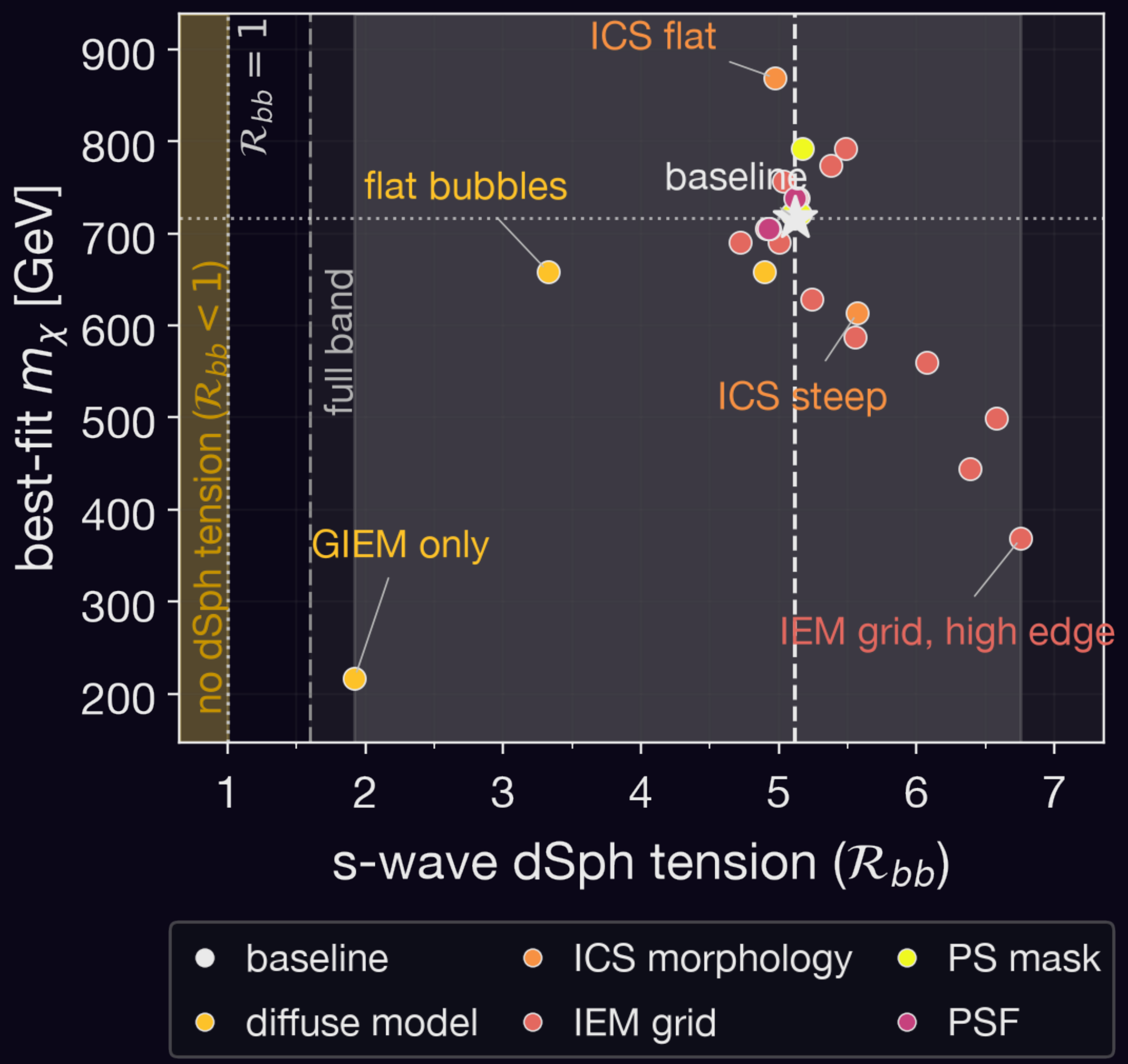

J. L. Feng, M. Kaplinghat & H.-B. Yu, PRD 82, 083525 (2010) · M. Ibe, H. Murayama & T. T. Yanagida, PRD 79, 095009 (2009) · K. K. Boddy et al., PRD 95, 123008 (2017) · S. Ando & K. Ishiwata, JCAP 05, 024 (2015) · T. Cohen et al., PRL 119, 021102 (2017) · C. Blanco & D. Hooper, JCAP 03, 019 (2019)

Two problems to contend with

- **relic abundance**
  $\langle\sigma v\rangle_{halo} \simeq 45 \times \langle\sigma v\rangle_{th}$ at baseline, ×4–49 across the suite a saturated, **non-resonant** enhancement is enough (benchmarks reach ~90 at 1 TeV)

- **dwarf ordering**
  $R \propto \langle S\rangle_{dwarf} / \langle S\rangle_{halo}$ and the dwarfs are colder, so R does not move: **only a resonance placed above threshold inverts it**, and that case is fine-tuned

Velocity dependent

Scenario	Dwarf tension ($b\bar{b}$)	Killed by
s-wave	$R \approx 5.1$ (1.6–9.8)	survives, edge of systematics
p-wave	$R \approx 0.02$ — over-resolves	**relic abundance**: under-predicts halo rate by ~7 dex
d-wave	$R \approx 10^{-4}$	same, ~12 dex
Decay (ρ^1, m_χ~1 TeV)	evades dwarfs	**IGRB**: needs $\tau \gtrsim 10^{28}$ s, off by >1 dex
Unsaturated Coulomb Sommerfeld	$R \approx 10^2$	unphysical limiting case, and Planck p_{ann}
Saturated Sommerfeld	**R unchanged from s-wave**	**reconciles relic abundance & halo rate. Planck p_{ann}: $b\bar{b}$ 10% below, W^+W^- 4% above**

J. L. Feng, M. Kaplinghat & H.-B. Yu, PRD 82, 083525 (2010) · M. Ibe, H. Murayama & T. T. Yanagida, PRD 79, 095009 (2009) · K. K. Boddy et al., PRD 95, 123008 (2017) · S. Ando & K. Ishiwata, JCAP 05, 024 (2015) · T. Cohen et al., PRL 119, 021102 (2017) · C. Blanco & D. Hooper, JCAP 03, 019 (2019)

Two problems to contend with

- **relic abundance**
  $\langle \sigma v \rangle_{halo} \simeq 45 \times \langle \sigma v \rangle_{th}$ at baseline, ×4–49 across the suite
  a saturated, **non-resonant** enhancement is enough
  (benchmarks reach ~90 at 1 TeV)

- **dwarf ordering**
  $R \propto \langle S \rangle_{dwarf} / \langle S \rangle_{halo}$ and the dwarfs are colder, so R does not move: **only a resonance placed above threshold inverts it**, and that case is fine-tuned

Velocity dependent

Scenario	Dwarf tension ($b\bar{b}$)	Killed by
s-wave	R ≈ 5.1 (1.6–9.8)	survives, edge of systematics
p-wave	R ≈ 0.02 — over-resolves	**relic abundance**: under-predicts halo rate by ~7 dex
d-wave	$R \approx 10^{-4}$	same, ~12 dex
Decay (ρ^1, m_χ~1 TeV)	evades dwarfs	**IGRB**: needs $\tau \gtrsim 10^{28}$ s, off by >1 dex
Unsaturated Coulomb Sommerfeld	$R \approx 10^2$	unphysical limiting case, and Planck p_{ann}
Saturated Sommerfeld	R unchanged from s-wave	reconciles relic abundance & halo rate. Planck p_{ann}: $b\bar{b}$ 10% below, W^+W^- 4% above
Resonant Sommerfeld / BW	**R < s-wave**	**only route that also fixes the dwarfs**

Light mediator – Jho 2512.24662 · Yoshimatsu 2602.23774 · Yamashita 2605.09979 · Nomura & Totani 2604.05016
Breit–Wigner – Murayama 2512.01404
Antiproton constraints – Wang & Duan (JCAP 05, 087, 2026

…and it costs you a structured dark sector

- The cost: this isn't a vanilla thermal WIMP. It needs **a dark sector** with enough structure to **enhance at Galactic velocities but not at freeze-out.**

the light-mediator realisation

- Two structurally different routes have been proposed for this excess:

Light mediator: Low-velocity Sommerfeld saturation
- Our model needs saturation **at or above v_{halo}**, sets floor **$m_\varphi \gtrsim 0.6$ GeV** at minimum coupling **$\alpha_\chi \approx 1.4 \times 10^{-2}$**.

Breit–Wigner resonance at ~$2m_\chi$:
- **No light state**. Heavy s-channel mediator tuned close to threshold

Light mediator – Jho 2512.24662 · Yoshimatsu 2602.23774 · Yamashita 2605.09979 · Nomura & Totani 2604.05016
Breit-Wigner – Murayama 2512.01404
Antiproton constraints – Wang & Duan (JCAP 05, 087, 2026

...and it costs you a structured dark sector

- The cost: this isn't a vanilla thermal WIMP. It needs **a dark sector** with enough structure to **enhance at Galactic velocities but not at freeze-out.**

annihilator
χ
$m_\chi \simeq 0.56 - 0.72$ TeV
thermal relic

low-velocity enhancement
Sommerfeld
two separable jobs
floor from $v_{sat} \gtrsim v_{halo}$

light mediator
ϕ
$m_\phi \gtrsim 0.6$ GeV
$\alpha_\chi \gtrsim 1.4 \times 10^{-2}$
heavier ϕ, larger α_χ

the light-mediator realisation

- Two structurally different routes have been proposed for this excess:

Light mediator: Low-velocity Sommerfeld saturation
- Our model needs saturation **at or above v_{halo}**, sets floor **$m_\varphi \gtrsim 0.6$ GeV** at minimum coupling $\alpha_\chi \approx 1.4\times10^{-2}$.
- Annihilation via a 4-body cascade (**cascade, $\chi\bar{\chi} \to \varphi\varphi \to$ SM**).
- At $m_\varphi \approx 0.6$ GeV (below $2m_p$ threshold), **antiproton production is kinematically closed.**

Heavier mediator → larger α_χ. If $m_\varphi \gtrsim 2m_p$ then antiprotons no longer kinematically closed.

Breit–Wigner resonance at ~$2m_\chi$:
- **No light state**. Heavy *s*-channel mediator tuned close to threshold
- BUT direct two-body decays ($\chi\bar{\chi} \to$ SM) **violate the AMS-02 antiproton bounds** by 15–18σ .

The preferred route **predicts a light mediator,** which we can look for.

A different handle: photon-DM scattering

A. Acar, C. Isaacson, M. Bashkanov & D. P. Watts, Phys. Lett. B 870, 139920 (2025) · T.R.S., A. Acar, M. Bashkanov, F. F. Deppisch, C. Ghag & D. P. Watts, arXiv:2608.29546 · C. Arina, A. Cheek, K. Mimasu & L. Pagani, EPJC 81, 223 (2021)

- The same vertices give $\bar{\chi}\chi \to \gamma\gamma$ (line searches) and $\gamma\gamma \to \bar{\chi}\chi$ (colliders). **Elastic $\gamma\chi \to \gamma\chi$ is their crossing** – spacelike transfer, both photons on-shell, linear rather than quadratic in ρ_χ
- We just computed the DM column towards the galactic centre ($J \sim 10^{22}$ GeV cm^{-2}) – region of "known" DM density, which all our measured photons crossed.
- Observable: optical depth $\tau(E_\gamma) = (J/m_\chi)\, \sigma_{tot}(E_\gamma, m_\chi)$, $J = \int \rho_\chi$. If DM couples to photons at all, a fraction of them scatter and the spectrum arrives attenuated and reshaped – energy dependently.
- Scan every gauge-invariant operator, dim-5 to dim-7 **– and the basis collapses.**
    - Anapole and charge radius couple through $\partial_\nu F^{\mu\nu}$, which vanishes identically for an on-shell photon. Attenuation is blind to them at any exposure. What survives is **dipoles and Rayleigh**.

B. Holdom, Phys. Lett. B 166, 196 (1986) · BaBar Collaboration, PRL 119, 131804 (2017) · NA64 Collaboration, PRL 131, 161801 (2023) · Y. Nomura & T. Totani, arXiv:2604.05016 (2026) · T.R.S., A. Acar, M. Bashkanov, F. F. Deppisch, C. Ghag & D. P. Watts, arXiv:2608.29546 · G. Lambiase et al., PRD 104, 023519 (2021) · Ning et al. (PandaX-4T), Nature 618, 47 (2023) · B. Kavanagh, P. Panci & R. Ziegler, JHEP 04, 089 (2019)

A different handle: photon-DM scattering

- **Dirac vs Majorana:** No Dipoles for self-conjugate DM, so the leading operator is dim-5 against dim-7: $\tau \propto E_\gamma^2$ vs E_γ^4–E_γ^6.

- If we see scattering: candidate mediator basis reduced before any data used. Kinetically mixed dark photon ⟶ matches the vanishing operators, invisible at any exposure (use BaBar and NA64). **Split electroweak doublet ⟶ visible**.

- **And the doublet's splitting closes the other two probes.**
$\delta \sim O(10^0$ keV), invoked to escape direct detection, leaves a transition moment $\bar{\chi}_2 \sigma^{\mu\nu} \chi_1 F_{\mu\nu}$.

δ against	which wins	result
halo kinetic budget, $\frac{1}{2}\mu_{\chi N} v^2 \simeq 31$ keV on xenon at 0.7 TeV	δ	direct detection closes at typical velocities, survives only in the fast tail
recombination temperature, $\delta/T \sim 4\times10^5$	δ	CMB annihilation bound closes **thermally**
photon energy, $E_\gamma \sim$ GeV	E_γ	attenuation has **no gate**

	μ [GeV^{-1}]
halo attenuation (this work)	$\lesssim 7.5$
CMB energy injection	1.5×10^{-3}
direct detection (PandaX-4T scaling est.)	9.5×10^{-7}
predicted doublet moment	$\mathbf{3.2\times10^{-7}}$

Attenuation is the only one of the three unsuppressed by the splitting — and the surviving corner is exactly where the model still lives.

First Operator-Resolved Sensitivity Estimates from the Galactic Halo

T.R.S., A. Acar, M. Bashkanov, F. F. Deppisch, C. Ghag & D. P. Watts, arXiv:2608.29546

- **Dirac vs Majorana discriminant:** leading operator dim-5 (Dirac), dim-7 (Majorana): $\tau \propto E_\gamma^2$ vs E_γ^4–E_γ^6 for $E_\gamma \ll m_\chi$.

- Reach: $\Lambda \simeq 0.32$ GeV (dim-5 Dirac dipoles), 0.21 GeV (dim-6 scalar Rayleigh), 0.79–1.06 GeV (dim-7 Rayleigh)

- Every contour sits in the EFT-invalid region across the cold-DM mass range.
Internal consistency statements of the analysis, not exclusions of open parameter space.

- The **structural results are the result**: from operator basis not exposure, hold at any sensitivity – including for instruments that don't exist yet.

The result is the method, not the number.

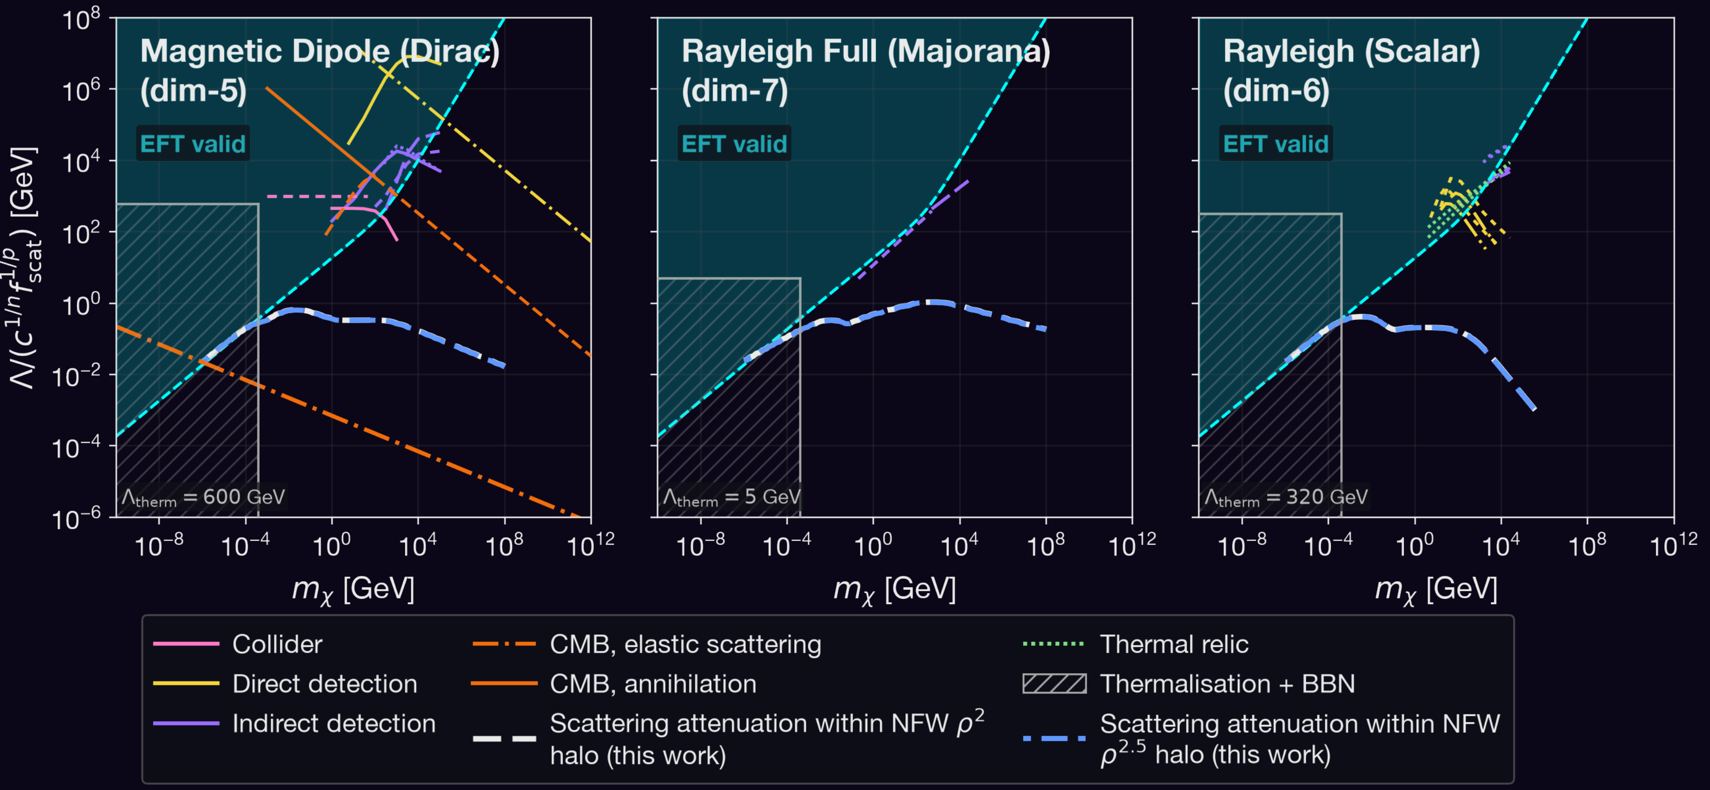

External constraints: Wilkinson, Lesgourgues & Boehm, JCAP 04, 026 (2014) · Lambiase et al. (2021) · LZ, PRL 131 (2023) · XENONnT, PRL 135, 221003 (2025) · Barducci et al., JHEP 06, 171 (2025)
Halo posterior: Stenhouse, Ghag & Deppisch, arXiv:2607.08552

Conclusions

1. The 20 GeV halo **survives** the move to pixel-level fitting with PSF folding and source masking. Peak 20.8 GeV, ~20% higher normalisation, same shape. Not an aggregation artefact.

2. It is **centrally concentrated, not isotropic**, and spectrally and spatially distinct from the GCE – the MSP cutoff sits an order of magnitude below its peak, and the halo correlates with the isotropic and point-source templates at −0.007 and −0.005.

3. The **mundane alternative fails**. An injected inner-Galaxy cosmic-ray population, across five physical morphologies, leaves the halo amplitude at ratio 1.000 and loses by $2\Delta \ln L \approx 611$–$688$ when asked to replace it.

4. What limits the measurement is template composition, not propagation physics. 14 GALPROP IEM models move the halo normalisation by <5%; how the flux divides between templates spans ×0.313 to ×1.831, the upper end an excluded stress test.

5. If it is DM: $m_\chi \approx 0.56$–0.72 TeV, $R_{b\bar{b}} = 5.1$ [1.6–9.8] & $R_{WW} = 4.1$ [1.2–7.7]. Mass freedom doesn't close it, p/d-wave fail relic abundance, decay fails the IGRB. Only **low-velocity Sommerfeld enhancement saturating at or above v_{halo}**. Four constraints of entirely different origins (**thermal relic density, dSph limits, Planck p_{ann}, and AMS-02 antiprotons**) convergingly point to this **same structured dark sector**, defining a concrete **sub-GeV mediator target**.

6. Photon attenuation is the crossed channel of that annihilation, so it tests the same operators independently. The **basis collapses to dipoles and Rayleigh:** a pure dark-photon mediator is invisible, the **split doublet** is visible here (potential **multi-messenger** synergy) . Sensitivity estimates EFT-invalid for now **– the framework is the result**.

Q&A

Thank You!

Any questions?

BACKUP

Backup Slides

T. Totani, JCAP 11, 080 (2025) · M. Ackermann et al., ApJ 750, 3 (2012) · J. Ballet et al., 4FGL-DR4 (2023) · J. F. Navarro, C. S. Frenk & S. D. M. White, ApJ 462, 563 (1996)
Figures here made using 17 yr data & pipeline of T. Stenhouse et al.

The sky model – six templates

- Standard template family: GALPROP gas + ICS, isotropic, Loop I, Fermi bubbles, point sources.
- Halo template: **NFW profiles** ρ^1 (decay-like), ρ^2 (annihilation), $\rho^{2.5}$ (GCE)

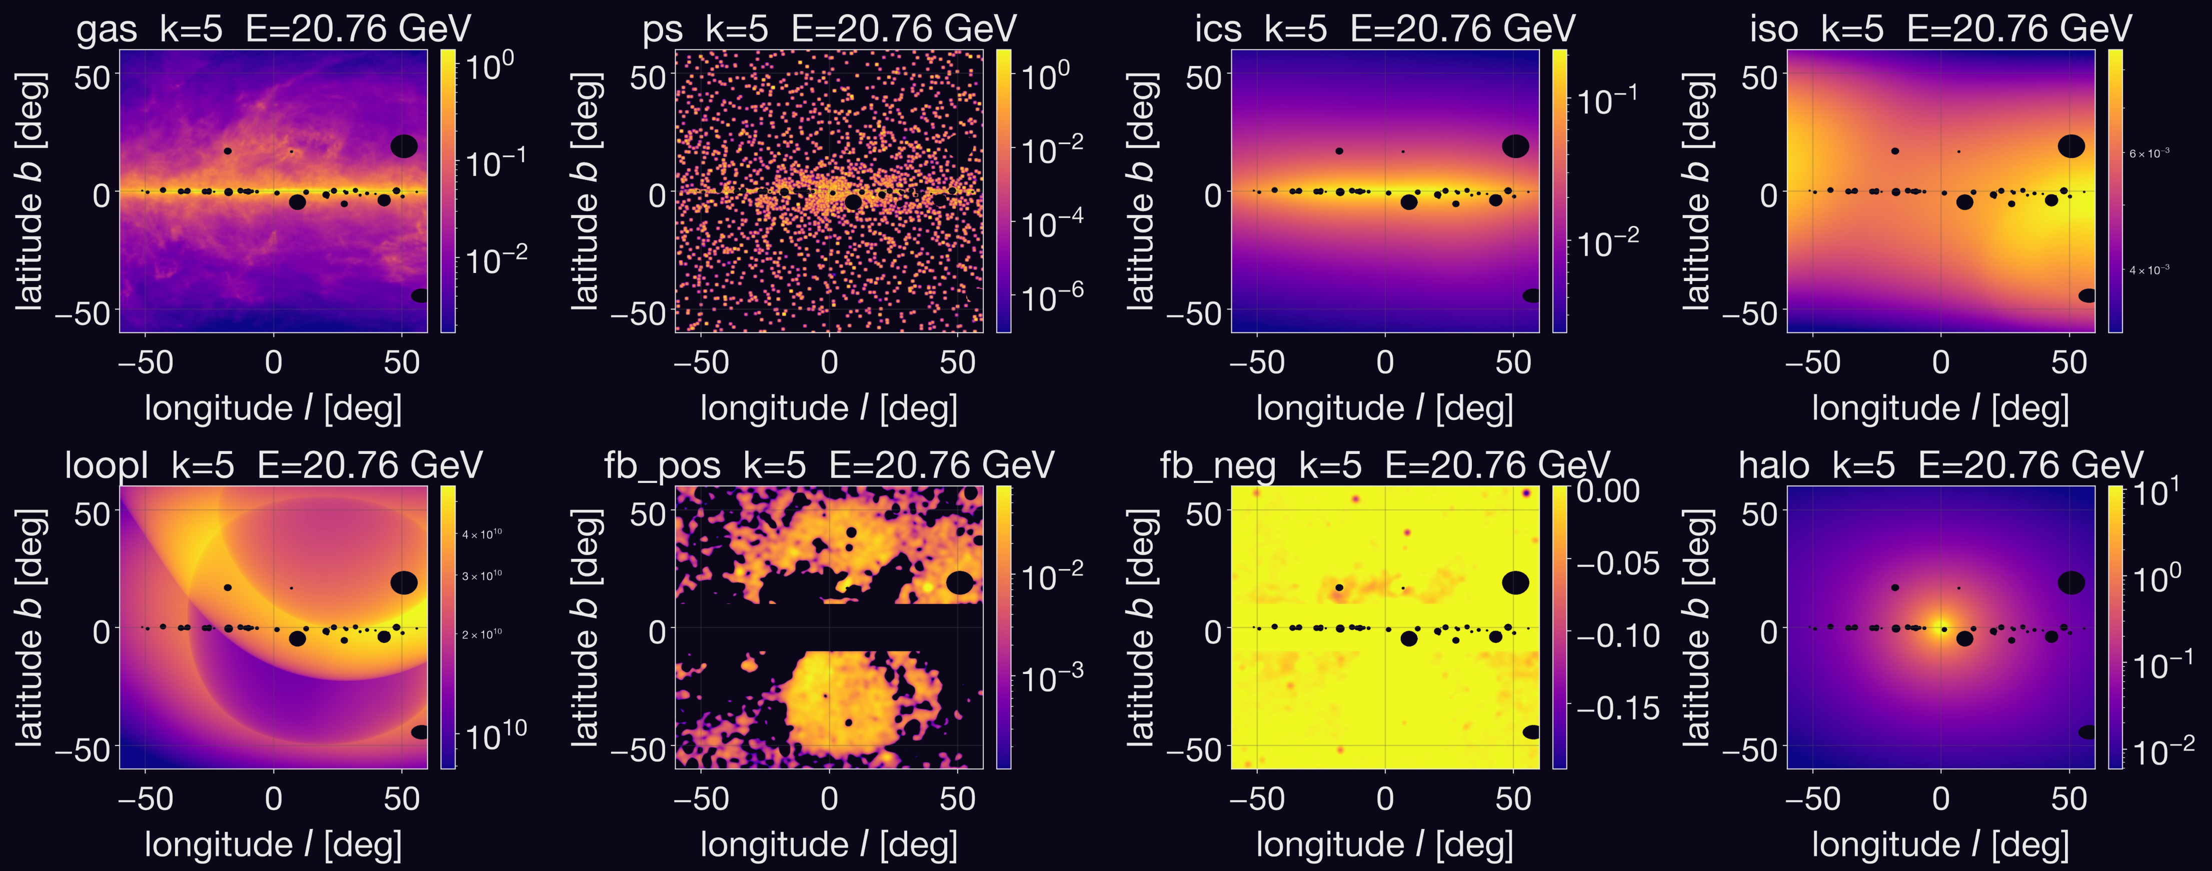

D. Finkbeiner, ApJ 614, 186 (2004) · G. Dobler & D. Finkbeiner, ApJ 680, 1222 (2008) · Planck Collab., A&A 554, A139 (2013) · G. Dobler et al., ApJ 717, 825 (2010) · M. Su, T. Slatyer & D. Finkbeiner, ApJ 724, 1044 (2010) · M. Ackermann et al., ApJ 793, 64 (2014) · P. Predehl et al., Nature 588, 227 (2020) · D. D. Dixon et al., New Astron. 3, 539 (1998) · K. C. Sarkar, A&A Rev. 32, 1 (2024)

Precedent: smooth residuals that were not diffuse

- WMAP template regression found a centrally concentrated **haze** with an anomalously hard spectrum (Finkbeiner 2004; Dobler & Finkbeiner 2008), later confirmed by Planck. The gamma-ray counterpart, the **Fermi haze**, was reported as smooth, hard and quasi-elliptical out to $|b| \approx 50°$ (Dobler 2010).
- Su et al. (2010) applied improved diffuse modelling and **resolved that apparently featureless component into the sharp-edged Fermi bubbles**. Since measured: edge width 3.4°, power law index 1.9 ± 0.2, exponential cutoff 110 ± 50 GeV, no significant spectral variation above $|b| = 10°$.
- The bubbles are now understood as the inner region of a larger outflow, traced in soft X-rays by eROSITA.
- And it is not confined to the haze. **Dixon et al. (1998) reported a statistically significant spherical gamma-ray halo in EGRET data**, later attributed to diffuse-model error.

The smoothness of a template-subtracted residual is not by itself evidence of a diffuse physical origin. That is exactly why we work at the native pixel scale.

Two Independent Checks: Replicate & Extend

- Extended dataset: **17 yr** Pass 8 (Aug 2008 – Mar 2026), same 13 log bins, 1.16–10^{58} GeV.

(a) Replication

Independent pipeline, independent reduction, Poisson MCMC on 10°×10° cells – Totani's own method, re-implemented from scratch.

→

(b) Extension

Global pixel-level likelihood on native 0.125° maps. Adds energy-dependent LAT PSF forward folding (King profile, Pass 8) and bright 4FGL-DR4 source masking.

- Extension uses fast deterministic global ML fit (L-BFGS-B algorithm) first. Reproduces MCMC halo spectrum and its statistical errors to within ~5% in the constraining signal bins at a fraction of the cost.
    - ML point estimates initialise MCMC's walkers (reduce burn-in).
    - Also used for systematics suite.
- All fitting uses exact Poisson likelihood.
    - At 10° counts are large, so distinction from Gaussian is immaterial.
    - At 0.125° it isn't. Only ~5.6% of pixels occupied near 20 .8 GeV peak, falling to ~0.011% in the top bin.

Energy dispersion and effective area

- We forward-fold the energy-dependent LAT PSF. **We do not forward-fold the energy dispersion.**
  - The Pass 8 resolution at the 20 GeV peak is $\sigma_E/E \approx 8\%$, against a logarithmic bin width of a factor 1.69, about 70%. Dispersion therefore redistributes counts well inside a single bin.
  - It neither shifts the 20 GeV peak nor biases the inferred mass at a level comparable to the foreground-modelling systematics.
- **The LAT effective-area systematic is also not propagated.** The instrument team quantifies it as an energy-dependent band at the few-per-cent level over the range used here.

It enters multiplicatively on the recovered flux, so linearly on $\langle\sigma v\rangle$ and on R, an order of magnitude below the ±48% and ±54% J-factor terms. It also cancels in part in R, since the dSph limits it is compared against carry the same calibration.

Both are disclosed in the paper as unpropagated floors rather than buried.

Exact Poisson Likelihood & Pixel Sparsity

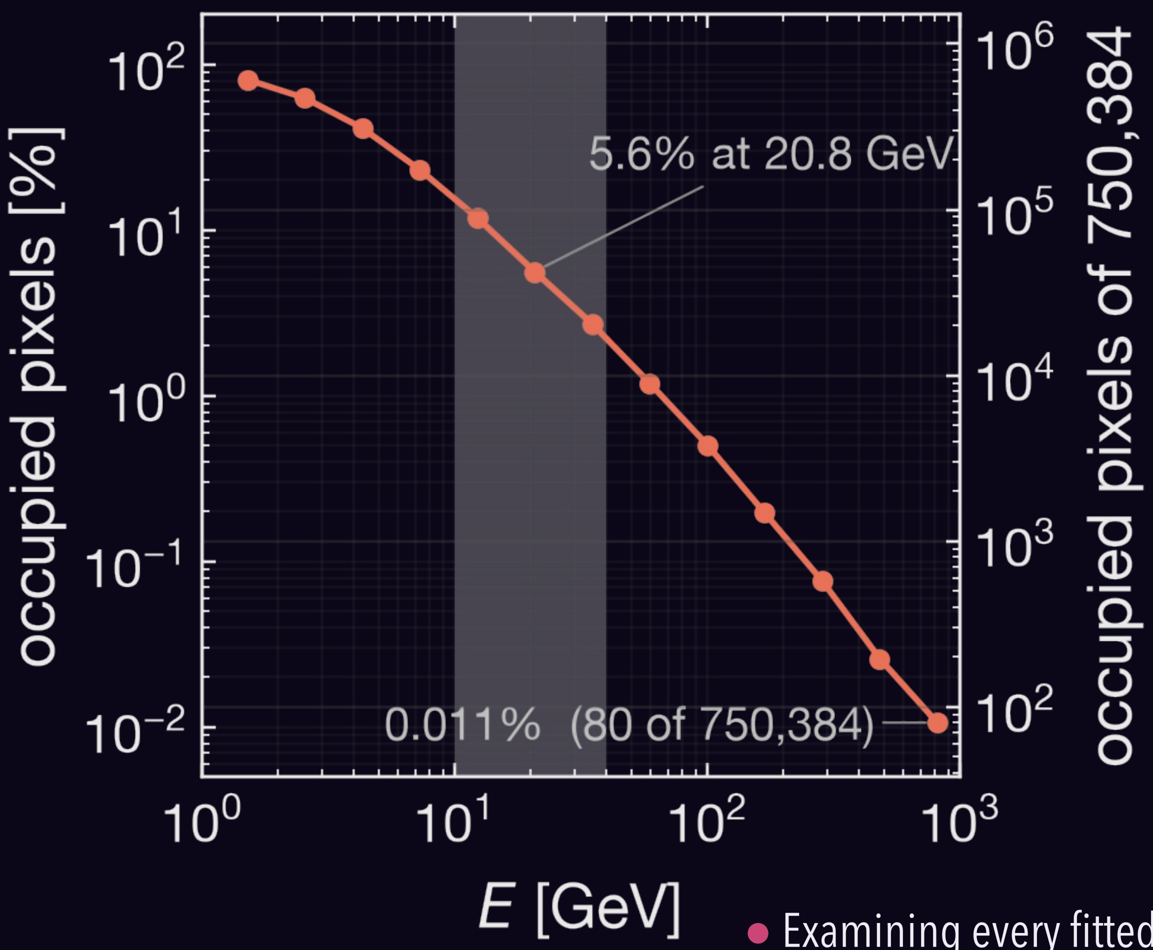

- The exact per-pixel Poisson log-likelihood formula:

$$\ln L_k = \sum_{j\in\text{fit}} \left[d_{jk} \ln \mu_{jk} - \mu_{jk} - \ln(d_{jk}!) \right]$$

evaluated on 750,384 fitted pixels, 0.125° CAR, 960 × 800 (disk-excl.)

  - Occupancy: **~5.6% of pixels occupied at the 20.8 GeV bin, ~0.011% in the top bin.** Counts fall roughly three and a half orders of magnitude across the band.
  - A Gaussian per-pixel error model is invalid where the overwhelming majority of pixels are empty. The Poisson likelihood is exact for empty pixels and requires no binning-up to rescue it.

- Model prediction is required positive at every occupied pixel. Non-positive fractions are 0.35% and 0.63% in the top two bins (2640 and 4754 pixels)

- Examining every fitted pixel in all 13 bins, the number at which the floor binds **and** the observed counts are non-zero is **zero**. Every floored pixel is empty, so the floor cannot shape the fit.

Production MCMC and convergence

- Affine-invariant ensemble sampler (emcee),
  - **64 walkers**, per energy bin.
  - Up to **60,000 steps**, 10,000 burn-in, thinned ×10.
  - Convergence criterion: **min N/τ > 50**.
  - Initialised at the deterministic solution. This shortens burn-in and does not enter as a prior.
  - Reports the maximum-likelihood sample and the marginalised quantiles p_{16} / p_{50} / p_{84}, with $\boldsymbol{\sigma} = ½(p_{84} − p_{16})$.

- **Errors marginalise over every other template amplitude in the same bin.** The halo error bar already contains the diffuse degeneracy.

- Deterministic cross-check. The fast **L-BFGS-B** fit reproduces the MCMC **spectra and per-bin errors to ~5%** across the well-populated bins of the signal band. It is validated before it is trusted to drive the systematics suite.

Resolution Dilution and the 15–25% Normalisation Offset

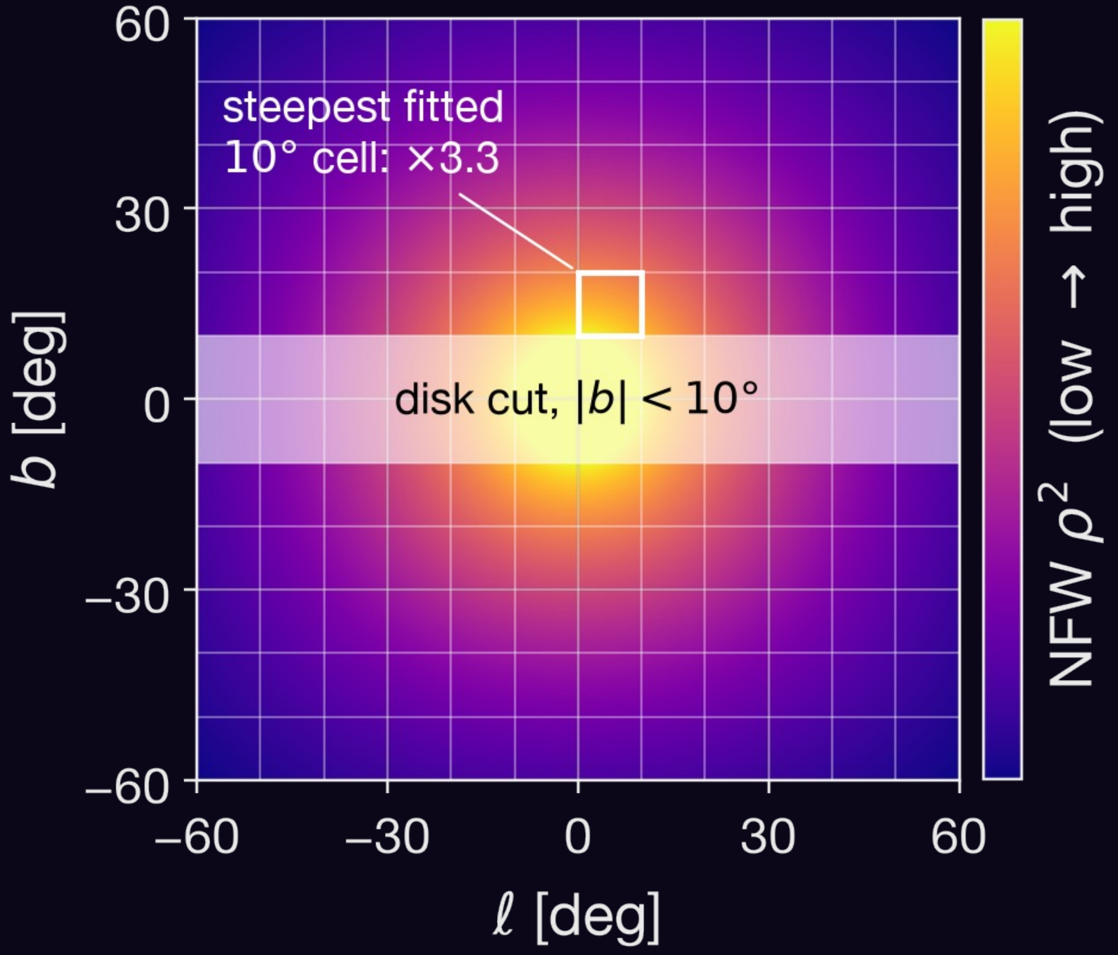

$$\langle \Phi_{\rm halo} \rangle_{\rm cell} = \frac{1}{\Omega_{\rm cell}} \int_{\rm cell} \Phi_{\rm halo}(\theta, \phi)\, d\Omega < \Phi_{\rm halo}(\text{pixel})$$

- The cellwise method averages intensity within independent $10^\circ \times 10^\circ$ cells, 10^0 deg^2 each, to guarantee high counts per cell.
- The NFW halo template has a genuine gradient across a cell at these latitudes. Averaging over 10^0 deg^2 dilutes the peak intensity.
- Fitting native 0.125° pixels preserves that gradient, so the recovered peak flux is higher.

Result: pixelwise sits 15–25% above cellwise, ≈20% for ρ^2.

Why $2\Sigma\Delta\ln L \approx 1.7\times10^3$ is not a significance

- It is a **template-level fit-improvement diagnostic**, not a calibrated significance, and we decline the Wilks reading. The templates are highly correlated and the diffuse model is incomplete.

- The DM prompt annihilation fit is **systematics-limited**: $\chi^2/\nu \approx$ **7.5 ($b\bar{b}$), 7.4 (W^+W^-)** at $\nu = 11$, driven by spatially coherent diffuse residuals in the photon-rich bins below 5 GeV, not by underestimated statistical errors.
    - Across the suite the range is 2.2–9.5 ($b\bar{b}$) and 1.9–9.5 (W^+W^-).

The residual sky, 10–40 GeV signal band

Left: no-halo fit, residual counts. **Right:** same fit with the NFW ρ^2 halo included. ROI $|\ell| \leq 60°$, $10° \leq |b| \leq 60°$, disk excluded.

- The halo template absorbs a centrally concentrated residual that the free isotropic term does not reproduce. Ring-averaged, it gives $\chi^2/\nu = 7.0$ for NFW ρ^2 against 29.1 for isotropic (See Slide 6).

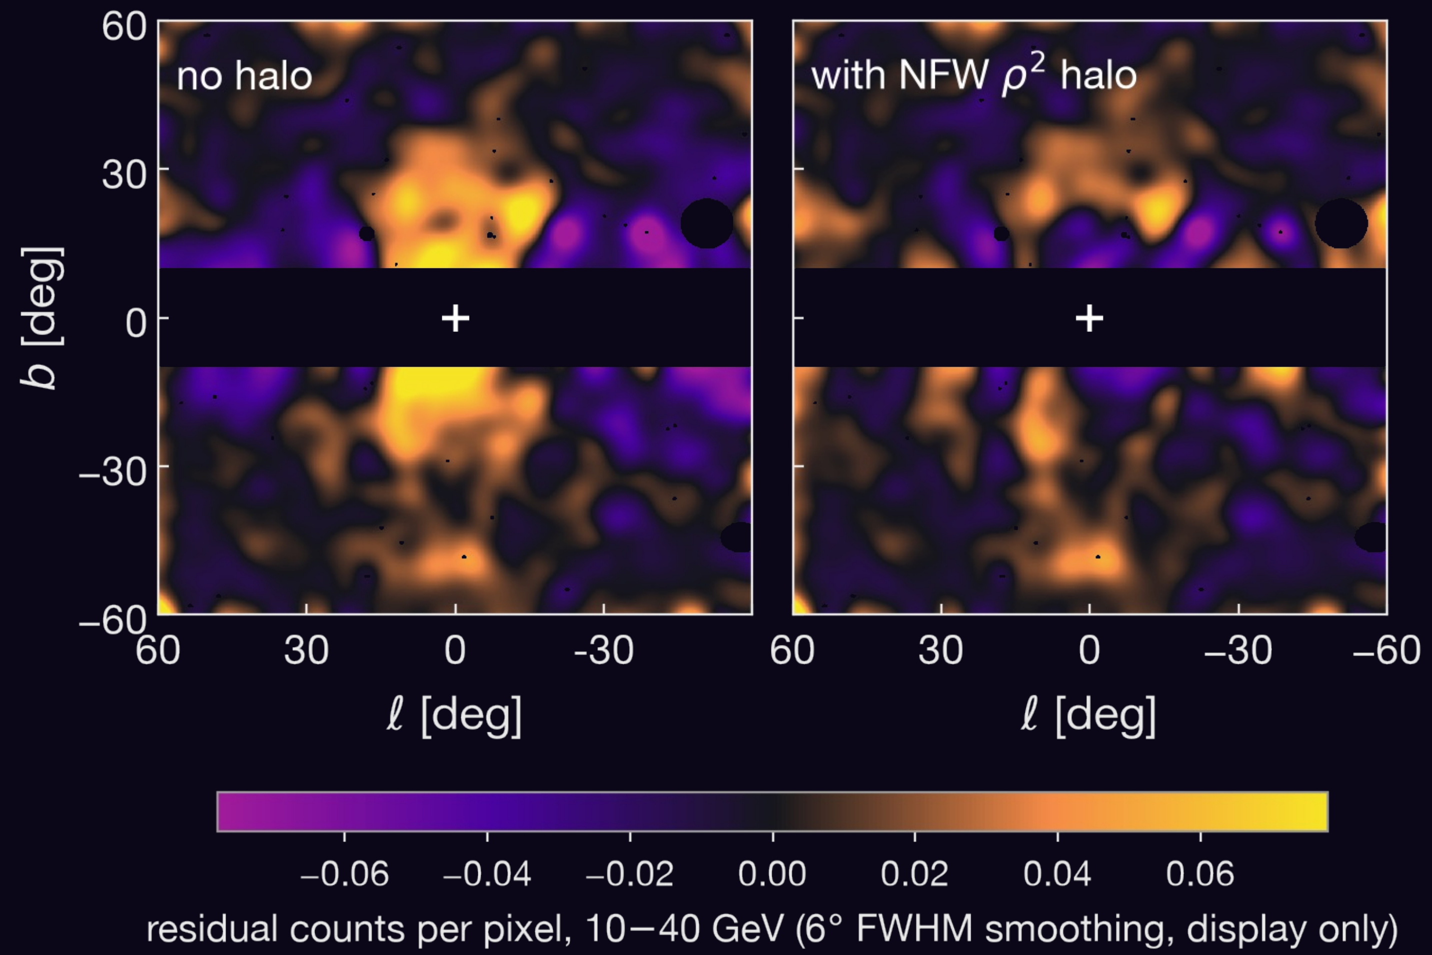

Morphology does not drift with energy

- Fitting the emissivity power ***p*** independently **per energy bin** across the signal band:
    - Constant-p fit across the band: **p = 2.8 ± 0.2**, χ^2 = 2.00 for 2 dof.
    - None of the three band intervals rails on the scanned grid.

bin	p
12.3 GeV	2.8 ± 0.3
20.8 GeV	2.7 ± 0.3
35.1 GeV	3.4 ± 0.5

- A cooling leptonic population loses energy as it propagates, so its emission shrinks with increasing photon energy. The recovered morphology does not. It is **consistent with a single constant power** across the band.

- The 482 GeV interval spans the entire scanned range of p, outside the signal band.

Could an injected IC component produce that ring profile?

Yes, on the ring profile alone. That is why the ring profile is **not the test**.

- Two injected morphologies give χ^2/ν of **4.88** (CMB-only bracket) and **5.09** (burst, $\lambda = 6$ kpc).
- NFW-ρ^2 gives **7.0**. NFW-$\rho^{2.5}$ gives **4.9**.

So on radial shape alone an extended injected-IC component is at least as good as the halo, and comparable to the cuspier profile.

What does discriminate: The joint spatial and spectral per-pixel likelihood with every template free.

- Fitted alongside the halo, the injected coefficient is **unconstrained** and **consistent with zero** in every signal-band bin, and the halo amplitude is unchanged at ratio 1.000.
- Correlation r(halo, injected) = −0.133, against halo to ICS −0.48 in the same fit. The two are separable.
- Forced to replace the halo, the injected population is disfavoured by **$2\Sigma\Delta\ln L \approx$ 611–688** and recovers at most **~30%** of the likelihood improvement the ρ^2 halo provides.

The 14-Model Interstellar Emission Model (IEM) Grid

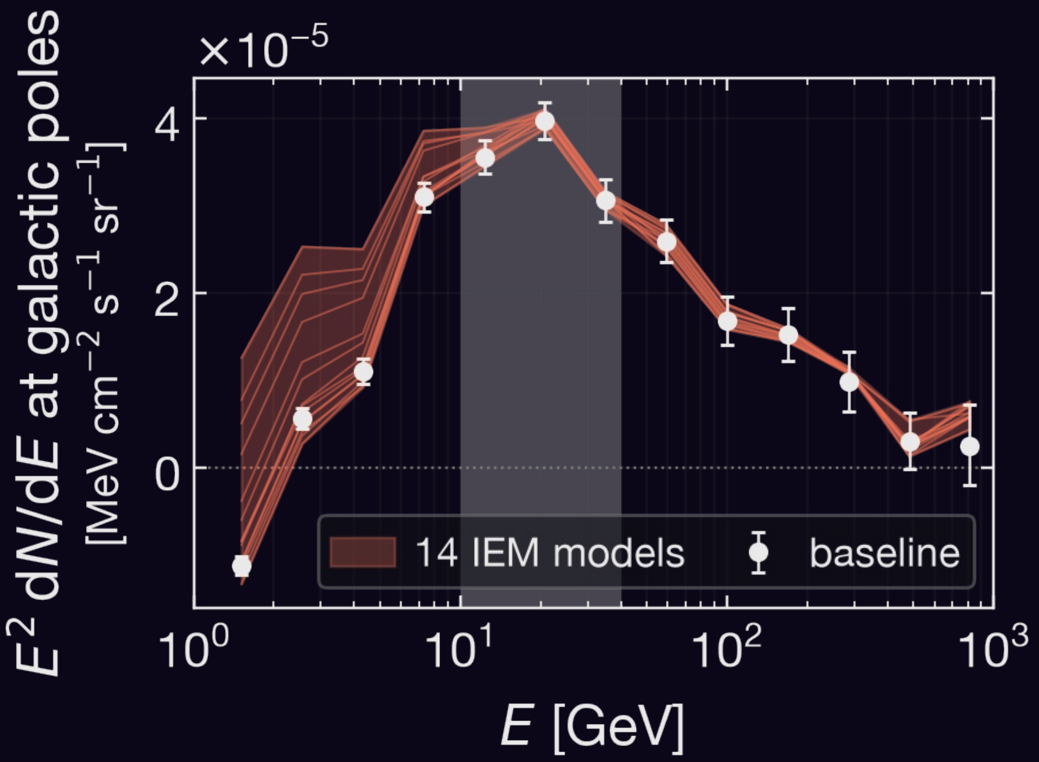

- 14 Ackermann et al. (2012) models, each LAT-normalised to local cosmic-ray data.
  - Gas and IC templates regenerated per model, and Fermi-bubble templates re-extracted self-consistently per model.

axis	values
CR source distribution	Lorimer, SNR, Yusifov, OB-star
Diffusion halo half-height z_h	4, 6, 8, 10 kpc
Radial boundary	20, 30 kpc
Spin temperature T_S	150 K, 10^5 K (optically thin)
E(B—V) magnitude cut	2, 5

Result: the halo normalisation moves ×0.970 to ×1.043, under 5%.

The bubble re-extraction is load-bearing. Retain the baseline bubble templates instead and the grid spans ×0.976 to ×1.169, and the envelope edge moves from 1.32 to 1.62.

No alternative CR source distribution in the published grid makes the foreground more halo-like.

The foreground systematics suite

20 configurations: 1 baseline, 3 template-composition, 2 ICS/ISRF, 14 IEM grid.

Template composition on halo flux: ×0.313 (GIEM-only) to ×1.831 (no bubbles).

variant	$R_{b\bar{b}}$	R_{WW}
baseline	5.1	4.1
GIEM-only	1.9	1.4
flat bubbles	3.3	2.6
no Loop I	4.9	–
ICS/ISRF	×0.97–1.09	×0.89–1.08

Foreground leg on the tension

channel	range
$b\bar{b}$	×0.375 to 1.320
W^+W^-	×0.341 to 1.314

Quadrature note. The foreground contribution is an extremal envelope across the suite, not a quadrature sum over variants. Quadrature enters only between that envelope and the two J-factor legs (halo 48%, dwarf 54%), which are independent of the diffuse model and of each other.

Why two variants are excluded

- GIEM with Loop I and bubbles **retained**: the GIEM is data-driven and already contains that large-scale structure, so keeping separate templates alongside it models the same flux twice and suppresses the halo by construction.
- **No bubbles**: dropping the bubble template without replacement leaves large-scale flux with nowhere to go, so the halo absorbs it. ×1.831 upward.

Note GIEM-only, which also removes Loop I and bubbles as separate templates is **in** the band and sets its lower edge.

Retaining the excluded GIEM variant would take the lower edge from $R_{b\bar{b}} \approx 1.6$ to ≈ 1.2, so the band as reported is the narrower choice at that end.

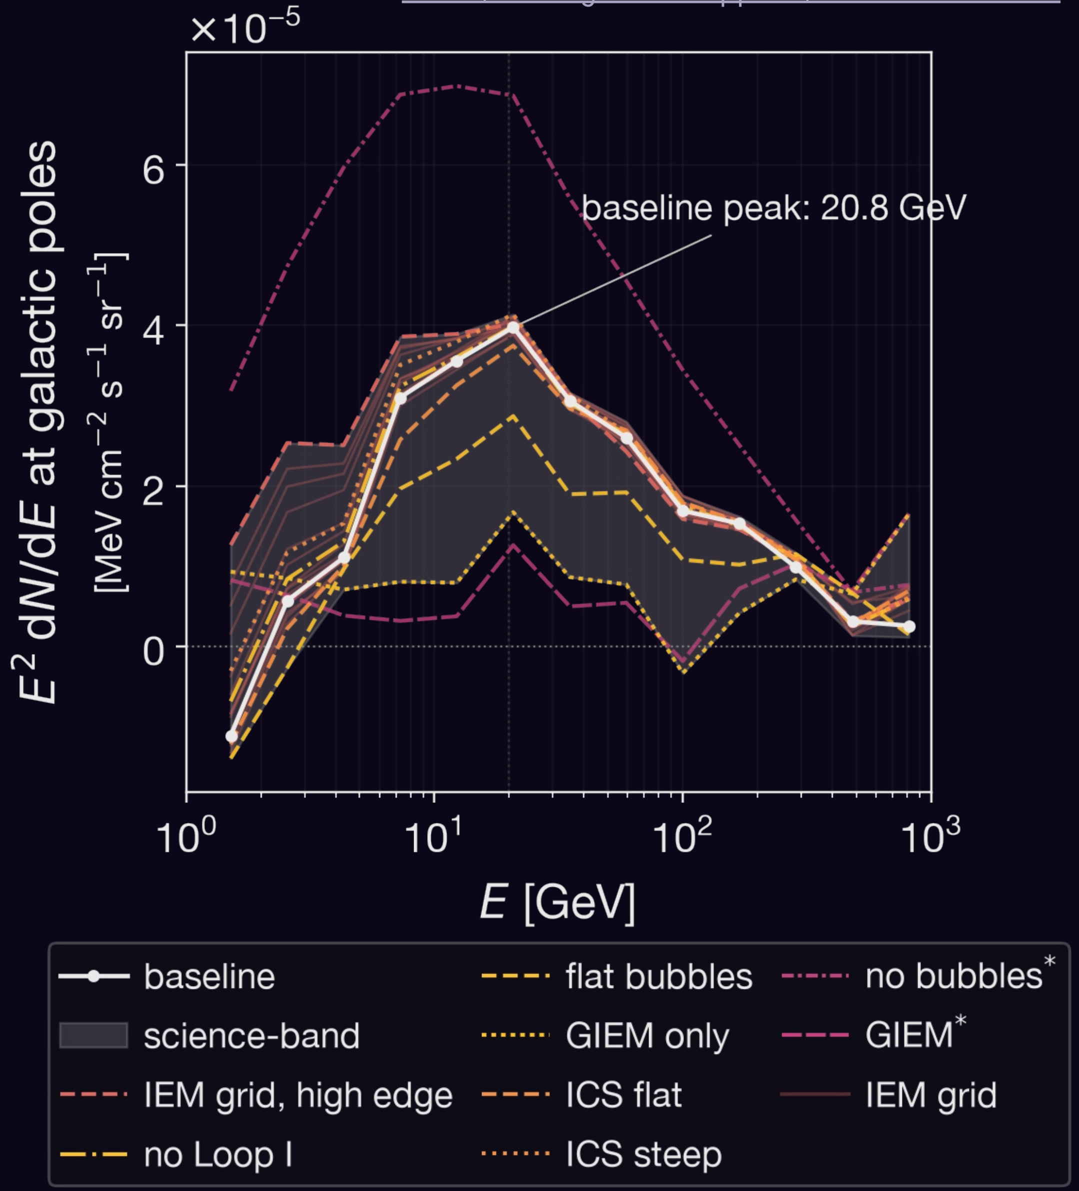

T. R. S., C. Ghag & F. F. Deppisch, arXiv:2607.08552

How stable is the best fit mass?

Not very:

- The extreme is OB-star at $z_h = 6$ kpc: **739 → 369 GeV** for $b\bar{b}$, with the **⟨σv⟩** ratio at 0.56 and the tension rising 32%. The halo **amplitude** in that same variant moves only ×1.036.
- Excluding the low-energy bins collapses the grid spread from **2.15 to 1.23**, and the seven Lorimer variants then span **614–644 GeV**.

- The **amplitude is stable to under 5%, the mass is not**. The mass is set by the low-energy shoulder, and that is where the foreground models disagree.

- The 1.5–4.4 over 10–40 GeV flux ratio is **0.05** at baseline against **0.57** for OB-star, and the difference sits at the inner ROI edge: **+0.38** at $|b| = 10$–20°, **+0.26** at 20–40°, **+0.06** at 40–60°.

set	$b\bar{b}$	W^+W^-
14-model IEM grid	369–793 GeV (factor 2.15)	293–587 GeV (factor 2.00)
Full foreground suite	217–869 GeV	161–644 GeV

T. R. S., C. Ghag & F. F. Deppisch, arXiv:2607.08552

Posterior template covariance

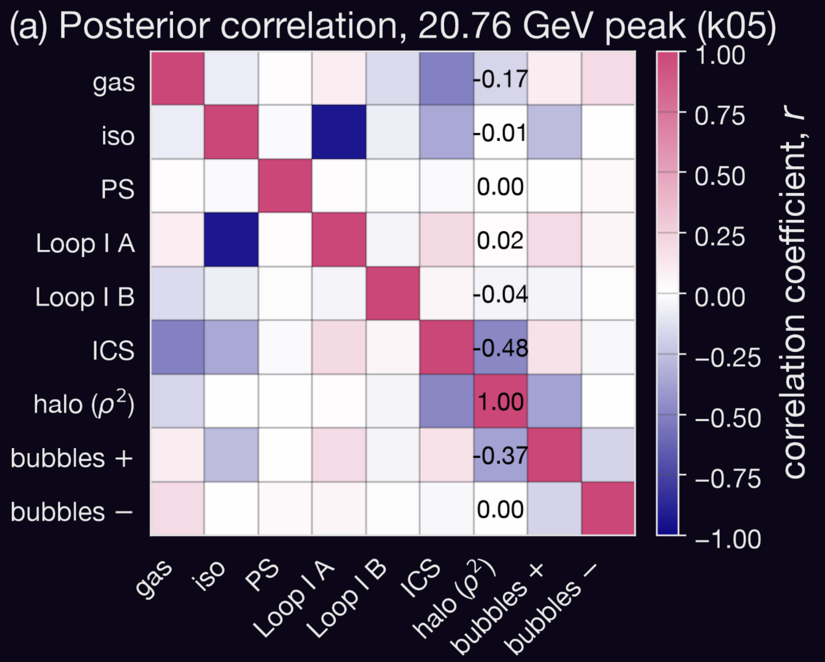

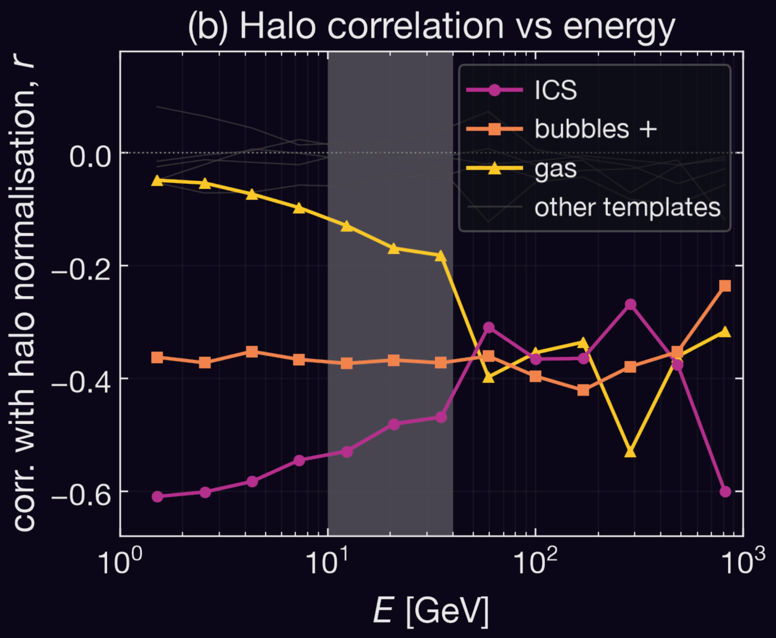

template	r (peak)	signal-band mean
inverse Compton	**−0.481**	−0.493
bubbles +	−0.368	−0.371
gas	−0.169	−0.160
Loop I B	−0.040	
Loop I A	+0.020	
isotropic	**−0.007**	
point sources	**−0.005**	
bubbles −	−0.002	

- Fit varies every template normalisation simultaneously, halo error bar **already** marginalises over gas, IC, Loop I and bubbles → pixel scale adds ability to report the covariance itself.
- Every significant correlation is **negative** – competing descriptions of the same flux, exactly as expected. Sum of squared halo correlations **0.397**: **Degenerate with the diffuse model, yes. Determined by it, no.**

The dwarf tension in both channels

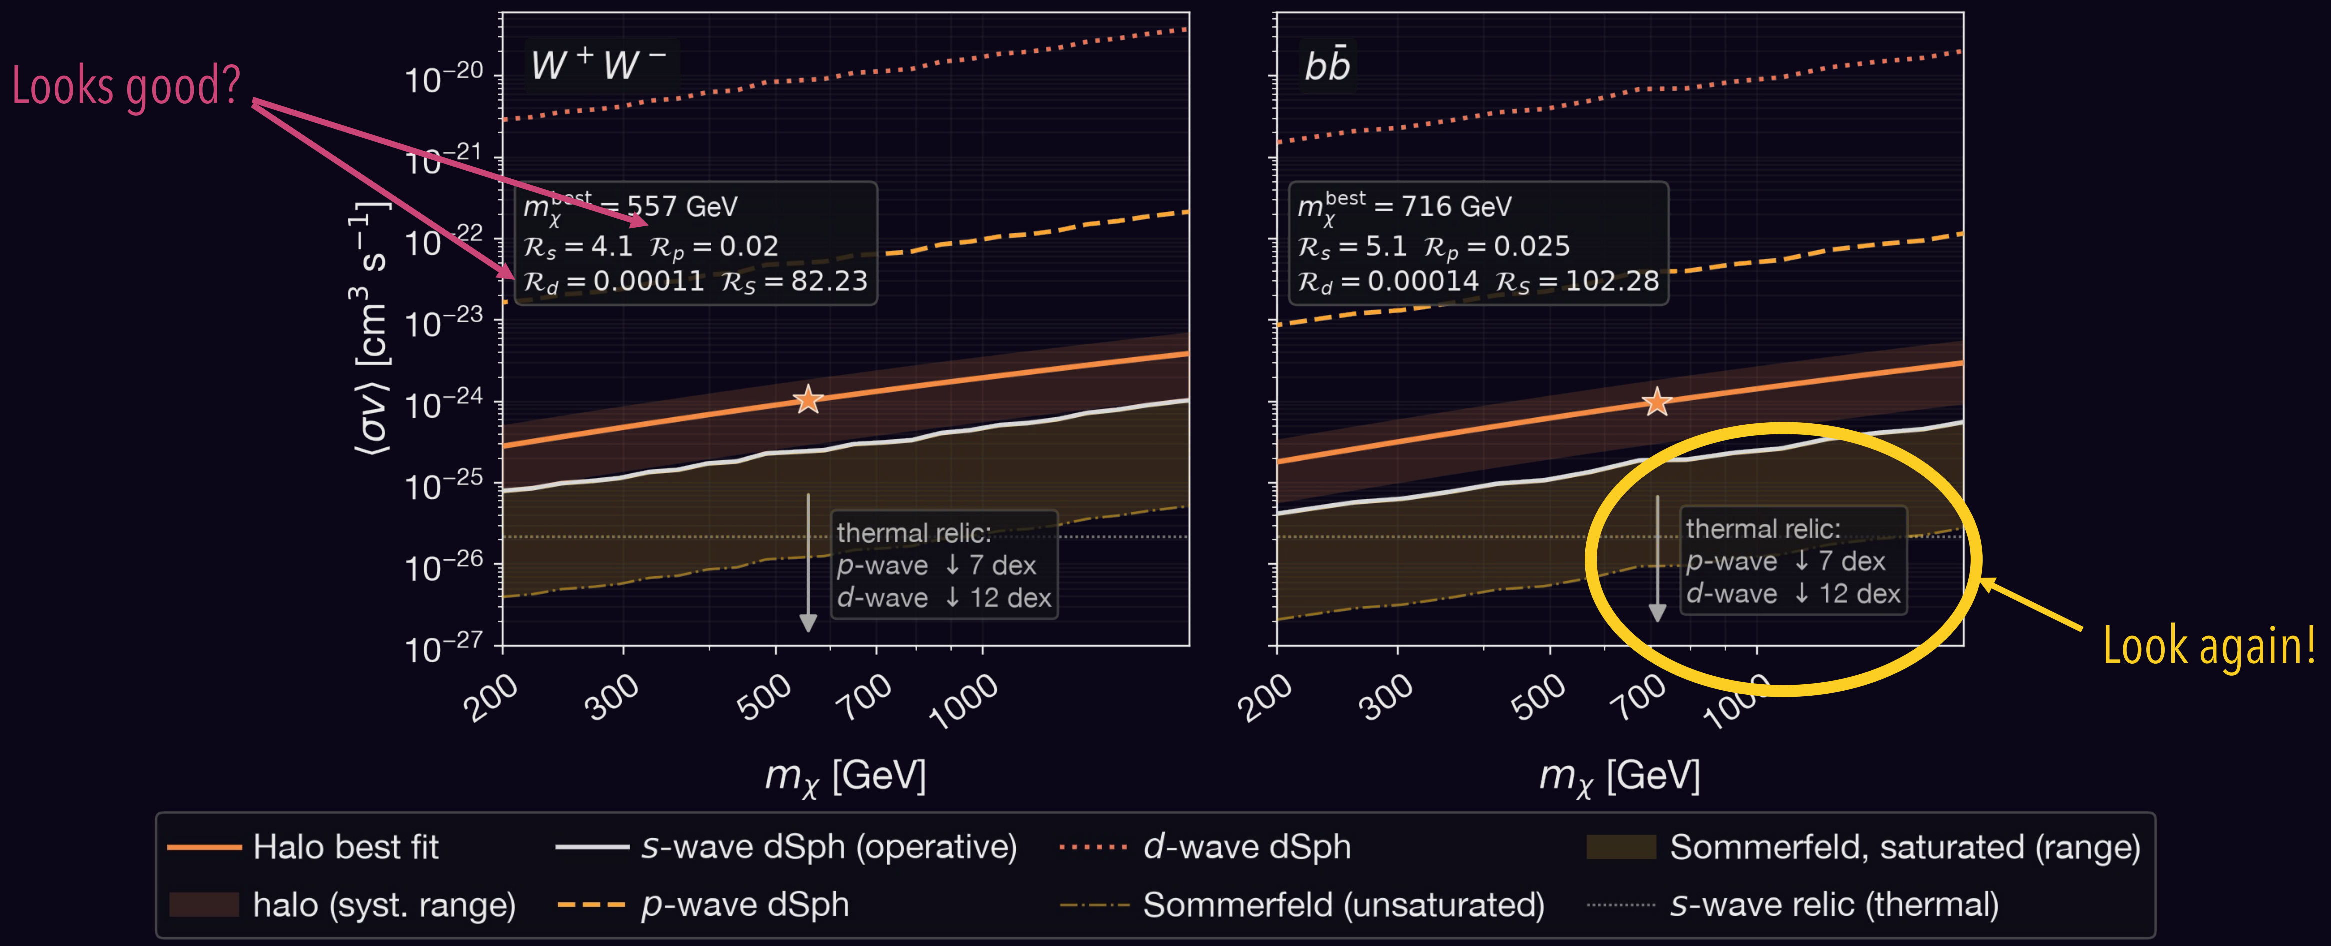

Updated dwarf limits: Circiello et al. (2026)

- Circiello et al., ApJL 1004, L33 (2026), arXiv:2605.25586. **Stronger limits** than Hoof's at our masses, 0.77 of it at 716 GeV, **tension increases**.

- Not adopted: release ships $b\bar{b}$ and $\tau^+\tau^-$ only.
  - With no W^+W^-, the band would sit half on one limit set and half on another **(inconsistent).** R_{WW} stays 4.07 [1.20, 7.75] on Hoof.

dSph limit set	dwarfs	R ($b\bar{b}$)	band
Hoof et al., as published – used in the paper	27	5.12	1.64 – 9.76
Circiello, measured-only	30	5.29	1.64 – 9.74
Circiello, benchmark – their fiducial	**42**	**6.47**	**1.96 – 11.90**
Circiello, inclusive	50	6.45	2.05 – 12.06

- Our own reconstruction from their per-dSph TS profiles, not a published limit curve. Sample dependence is small on the new data, 6.47 for 42 dwarfs against 6.45 for 50, which is their own claim holding up.

The AMS-02 antiproton constraint

- Wang & Duan (JCAP 05, 087, 2026) find that at Totani's parameters the implied antiproton flux **exceeds** AMS-02 by **15σ for $b\bar{b}$** and **18σ for W^+W^-**.
  - That constrains **direct** annihilation to those final states.
  - Surviving light-mediator models **cascade**, $\chi\bar{\chi} \to \phi\phi \to$ SM, which softens the antiproton spectrum. Argued to evade the bound for $m_\phi \sim 10$ GeV (Jho et al.).
  - That closure applies at the floor of the permitted range. **Heavier mediators satisfying the same saturation condition are not exempt**, as the O(10) GeV realisations show.

The CMB bound on Sommerfeld enhancement

T. R. S., C. Ghag & F. F. Deppisch, arXiv:2607.08552 · Planck Collab., A&A 641, A6 (2020) · T. R. Slatyer, PRD 93, 023527 (2016) · N. Arkani-Hamed et al., PRD 79, 015014 (2009) · J. L. Feng, M. Kaplinghat & H.-B. Yu, PRD 82, 083525 (2010)

- A monotonic 1/v Sommerfeld enhancement is ~10^5 times larger at recombination than in the halo today. That is where such models usually die.

- **We fold it in.** Planck bounds annihilation energy injection at $p_{ann} = f_{eff}\langle\sigma v\rangle/m_\chi < 3.2\times10^{-28}$ cm^3 s^{-1} GeV^{-1}. With $f_{eff} = 0.21$ ($b\bar{b}$) and 0.18 (W^+W^-) from Slatyer's tabulated deposition efficiencies: $\langle\sigma v\rangle < 1.1\times10^{-24}$ ($b\bar{b}$), 1.0×10^{-24} (W^+W^-).

The $b\bar{b}$ baseline 9.8×10^{-25} sits ~10% below its bound. The W^+W^- baseline 1.03×10^{-24} sits some 4% above its own, within the uncertainty on feff and on the halo normalisation.

  - An **unsaturated** boost of 45 is **excluded outright**, and **the CMB independently requires saturation at or above v_{halo}**, the same condition forced by dwarf ordering.

- The saturated non-resonant option survives at best with little to no headroom. Improved p_{ann} measurements test it directly. An above-threshold resonance relaxes the bound further.

An unrelated signal, the same structural requirement

- Fan et al., PRL 137, 071003 (2026): evidence for a **~43 GeV line in a stack of nearby galaxy clusters, absent in the inner Galaxy**. They note that a velocity-independent annihilation rate does not accommodate that ordering.
- Lacroix et al., JCAP 10 (2022) 021: a systematic classification across target classes finds that clusters can outshine other targets for s-wave annihilation on resonance.
- What the two share is the **structural requirement**: an annihilation rate that is **not monotonically increasing toward low velocity** – the same thing our above-threshold resonance needs.

That case differs from ours in signal, mass and target, and we draw no connection between the two.

What would falsify the halo result?

- **Planck p_{ann}, the sharpest near-term test.** The saturated non-resonant option sits ~10% below its bound for $b\bar{b}$ and ~4% above for W^+W^-. An improved p_{ann} measurement tests it directly.
- **Deeper dwarf limits.** R scales inversely with the dSph upper limit. Circiello et al. already move it up by roughly 30%.
- **A demonstrably better inner-Galaxy foreground.** Template composition spans ×0.313 to ×1.831. That is the systematic that sets both R and the mass, and it is the one that would move most.
- **The mediator.** The saturation floor $m_\varphi \gtrsim 0.6$ GeV is a concrete target. A dark sector with no state at or above it cannot supply the boost.

What would not falsify it: a null attenuation result. That reach sits some seven orders from the doublet prediction, so the photon channel tests the operator structure, not the rate.

The operator catalogue and the structural null

- All gauge-invariant operators of dimension 5 to 7 coupling scalar, Majorana or Dirac dark matter to the photon.

- The dimension-6 **anapole** ($\bar{\chi}\gamma^\mu\gamma^5\chi\,\partial^\nu F^{\mu\nu}$) and **Dirac charge radius** ($\bar{\psi}\gamma^\mu\psi\,\partial^\nu F^{\mu\nu}$) couple through $\partial^\nu F^{\mu\nu}$, whose momentum-space vertex is the transverse projector $-(k^2\varepsilon_\mu - k_\mu\, k\cdot\varepsilon)$. For an on-shell photon, $k^2 = 0$ and $k\cdot\varepsilon = 0$, so it **vanishes identically**.
  - Real-photon attenuation is blind to them at tree level, at any exposure, on any dataset. The real-photon basis collapses to **dipoles and Rayleigh**.

Peak reach, 90% CL, current Fermi-LAT exposure

operator	Λ	at m_χ
dim-5 Dirac dipoles	**0.32 GeV**	60 GeV
dim-6 scalar Rayleigh	**0.21 GeV**	0.22 GeV
dim-7 Rayleigh	**0.79–1.06 GeV**	0.12–0.58 TeV

- The dipole contour is flat at $\Lambda \simeq 0.29$ GeV below $m_\chi \sim 1$ GeV. In the $\omega \gg m_\chi$ regime $\sigma \propto \mu^4 m_\chi\, \omega$, so $\tau = (J/m_\chi)\sigma$ loses its mass dependence. It looks like a bug and is not.

- **Rescaling for a sub-dominant scatterer:** $\Lambda \to \Lambda\, f^{1/4}$ for dimension 5 and 6, $\Lambda \to \Lambda\, f^{1/6}$ for dimension 7. Any limit here can be read at any scattering fraction

EFT Cutoff Validity & The Unitarity Wedge

Validity condition: $\Lambda^2 \geq \max(s_{max}, |t_{max}|)$, evaluated at ω_{max} = 168.9 GeV, the log-centre of the highest bin surviving the 200 GeV cut. Not the cut itself.

- Every 90% CL contour sits below the wedge across the cold-DM mass range, clearing it at the CDM floor by 0.30 to 0.64 dex.
- The contours do cross into the wedge at low mass: the dipoles below $m_\chi \simeq$ **240 keV**, the scalar Rayleigh below **130 keV**, the dimension-7 Rayleigh operators below **12–19 keV**.
- Those crossings are where two shallow curves meet and are correspondingly sensitive to scan resolution, so they are quoted to two significant figures.

The low-mass window is not open parameter space.

Early-Universe Thermalisation & BBN Limits

- By crossing, the same operators drive **$\gamma\gamma \leftrightarrow \chi\bar{\chi}$**. At T = 1 MeV that rate exceeds the Hubble rate by roughly **10^{12}** for the dimension-5 and dimension-6 windows, and by **10^{3}** for dimension-7.
  - Anything in the 12–240 keV window therefore **thermalises** with the photon bath.
  - Big-bang nucleosynthesis forbids a thermalised species below **0.4 MeV**.
  - A species of mass 12–240 keV in thermal contact with the photon bath is excluded.

So the region where the contours are EFT-valid is already closed by cosmology. The low-mass contours are internal consistency statements of the attenuation analysis, not exclusions of open parameter space.

T.R.S., A. Acar, M. Bashkanov, F. F. Deppisch, C. Ghag & D. P. Watts, arXiv:2608.29546 · Y. Nomura & T. Totani, arXiv:2604.05016 · B. Holdom, PLB 166, 196 (1986) · T. Hambye & X.-J. Xu, JHEP 11, 156 (2021) · N. Weiner & I. Yavin, PRD 86, 075021 (2012)

The Inelastic "Splitting Gate" Kinematics

The higgsino-like electroweak doublet carries a one-loop magnetic dipole, $\mu \sim e\, g^2/(64\pi^2 m_\chi)$, from W and chargino loops.

An inelastic mass splitting of $\boldsymbol{\delta} \sim O(10^0$ keV) splits the Dirac state into two Majorana states $\boldsymbol{\chi_1}$ and $\boldsymbol{\chi_2}$. Diagonal moments vanish, $\bar{\chi}i\, \sigma^{\mu\nu} \chi i \equiv 0$, but the dipole survives as a transition moment $\bar{\chi}_2\, \sigma^{\mu\nu} \chi_1 F_{\mu\nu}$. Direct detection closes at typical halo velocities and survives only in the fast tail, ER $\gtrsim \boldsymbol{\delta}$.

One number against three energy scales, three different answers.

δ against	which wins	result
halo kinetic budget, $\frac{1}{2}\mu_{\chi N} v^2 \simeq 31$ keV on xenon at 0.7 TeV	**δ**	direct detection closes at typical velocities, survives only in the fast tail
recombination temperature, $\boldsymbol{\delta}/T \sim 4\times10^5$	**δ**	CMB annihilation bound closes **thermally**
photon energy, $E_\gamma \sim$ GeV	E_γ	attenuation has **no gate**

The two-insertion Compton amplitude passes through a virtual χ_2 and returns. **At $E_\gamma \gg \delta$ it cannot tell split from unsplit.**

T.R.S., A. Acar, M. Bashkanov, F. F. Deppisch, C. Ghag & D. P. Watts, arXiv:2608.29546 · G. Lambiase et al., PRD 104, 023519 (2021) · R. Wilkinson, J. Lesgourgues & C. Boehm, JCAP 04, 026 (2014) · XENONnT Collab., PRL 135, 221003 (2025) · Ning et al. (PandaX-4T), Nature 618, 47 (2023)

The Inelastic "Splitting Gate" Kinematics

- Attenuation trails the other two by **3.7 and 6.9 decades**. Direct detection sits within a factor of **~3** of the prediction, so the unsplit realisation is close to exclusion already. **That is why the splitting is invoked**, and the splitting is what hands the corner to attenuation.

- The splitting does not relieve the AMS-02 antiproton bound on W^+W^-. This is complementarity among the three photon-sector probes, not survival against everything.

bound	μ [GeV^{-1}]
halo attenuation (this work)	$\lesssim 7.5$
CMB energy injection	1.5×10^{-3}
direct detection (PandaX-4T, scaling est.)	9.5×10^{-7}
predicted doublet moment	$\mathbf{3.2\times10^{-7}}$

9.5×10^{-7} is our own scaling estimate, not a PandaX or LZ limit. Their published $\mu_\chi < 4.8\times10^{-10}\ \mu_B$ at 20–40 GeV is 1.42×10^{-7} GeV^{-1}, and $\sqrt{m}$-scaling to 0.7 TeV gives $(6.0–8.4)\times1^{-77}$. We quote 9.5×10^{-7}, erring loose.

What would actually improve the sensitivity

- **Exposure is the weakest lever.** $\Delta\chi^2 \propto \Lambda^{-8}$ for the dipole and scalar Rayleigh, Λ^{-12} for the dimension-7 family.
  - $\times 10$ in exposure buys $10^{1/8} \approx 1.3$ and $10^{1/12} \approx 1.2$, only 0.13 and 0.08 decades in Λ.
  - Even $\times 10^6$ moves the dimension-7 boundary by only $10^{6/12} \approx 3$, from $\Lambda \sim 0.79$–1.06 GeV to $\Lambda \approx 3.4$ GeV.

- Band is the lever that works. The validity wedge scales as $\Lambda \propto \sqrt{\omega_{max}}$ and the exclusion contour as $\Lambda_{90} \propto \omega_{max}^{k/p}$, with k the energy index of the cross section and p = 4 or 6. The gap closes as (½ − k/p) per decade of ω_{max}.
  - **Dipoles:** k/p = ½ exactly, so the gap is band-invariant. Only improved fractional precision helps.
  - **Rayleigh:** k/p = ⅔ to 1, so the gap closes. CTA and SWGO help this family.

Downward extends the lever arm that measures the energy index, which is the Dirac vs Majorana discriminant. MeV instruments matter for that, not for reach.

Polarisation is unexploited. The optical depth used throughout is unpolarised, which discards the linear polarisation of the scattered photons. That polarisation distinguishes spin-2 exchange from the scalar and dipole structures and is a handle no unpolarised measurement can reproduce.

LZ Collaboration, preprint (1 Sept 2026) · N. Anand, A. L. Fitzpatrick & W. C. Haxton, PRC 89, 065501 (2014) · P. W. Graham, H. Ramani & S. S. Y. Wong, PRD 111, 055030 (2025) · Y. Nomura & T. Totani, arXiv:2604.05016

LZ: one event at 248 keV in the extended window

LZ Collaboration, preprint 1 Sept 2026

One event, in a region where the known background expectation is low.

exposure	2.84 tonne-year
NR window	extended to ~270 keV
the event	248 ± 23 keV
max local	3.4σ
global	2.6σ, look-elsewhere
intervals	two-sided, lower limit off zero for some masses

Their two models are our split doublet

O_1^s inelastic, δ scanned at 1 TeV. **They say it resembles a Higgsino model.** Nomura & Totani link is ours, not theirs.

L^s_{10}, a magnetic moment interaction. **Same $\bar{\chi}\sigma^{uv}\chi$ bilinear as our dim-5 dipole.**

So it needs Dirac dark matter

$\bar{\chi}\sigma^{uv}\chi \equiv 0$ for self-conjugate χ. **Their operator, our discriminant.**

Not yet our number

Their Eq. 1 is a contact interaction, not photon-mediated. d^s_{10} is dimensionless, our μ is in GeV^{-1}. Their inelastic example is fixed at 1 TeV; their L^s_{10} limit scans mass and spans our best-fit mass range.

Why 248 keV is hard for the split state

For 0.7 TeV on xenon, $\frac{1}{2}\mu v^2$ runs from ~31 keV at typical halo velocity to **~350 keV at the escape tail**. An endothermic δ peaks the spectrum high.

What does this result mean?

- We report the 20 GeV halo as consistent with sub-TeV annihilation rather than as evidence for it, and the distinction is not a formality.

- The feature survives every test applied here, and the diffuse model remains the reason that survival is not yet a measurement of dark matter.

- Establishing more requires the interstellar-emission model to be marginalised over within the fit rather than bracketed after it:
    - the grid assembled here makes this tractable.
    - this would convert the leading systematic on this class of measurement from a reported envelope into a fitted uncertainty.

- Four constraints of different origin, the relic density, the dSph limits, the energy-injection bound and the antiprotons point to the same class of dark sector without yet agreeing on its mass scale, and the mediator scale of order 0.6 GeV that our own analysis requires is a definite target for future work.