Experimental implications of neutrino mass models for HNL searches

Jean-Loup Tastet · NuTs workshop · 2022-06-15

Plan

- Brief recap about HNLs [see also Xabi's talk last Friday]
- Searching for HNLs at the SHiP experiment
- Searching for HNLs at the ATLAS experiment at the LHC
- How to report collider limits in a generic way

Heavy Neutral Leptons

Limitations of the SM

Limitations of the SM

Observational limitations

Limitations of the SM

Observational limitations

- Massless neutrinos
  $\Longrightarrow$ no oscillations

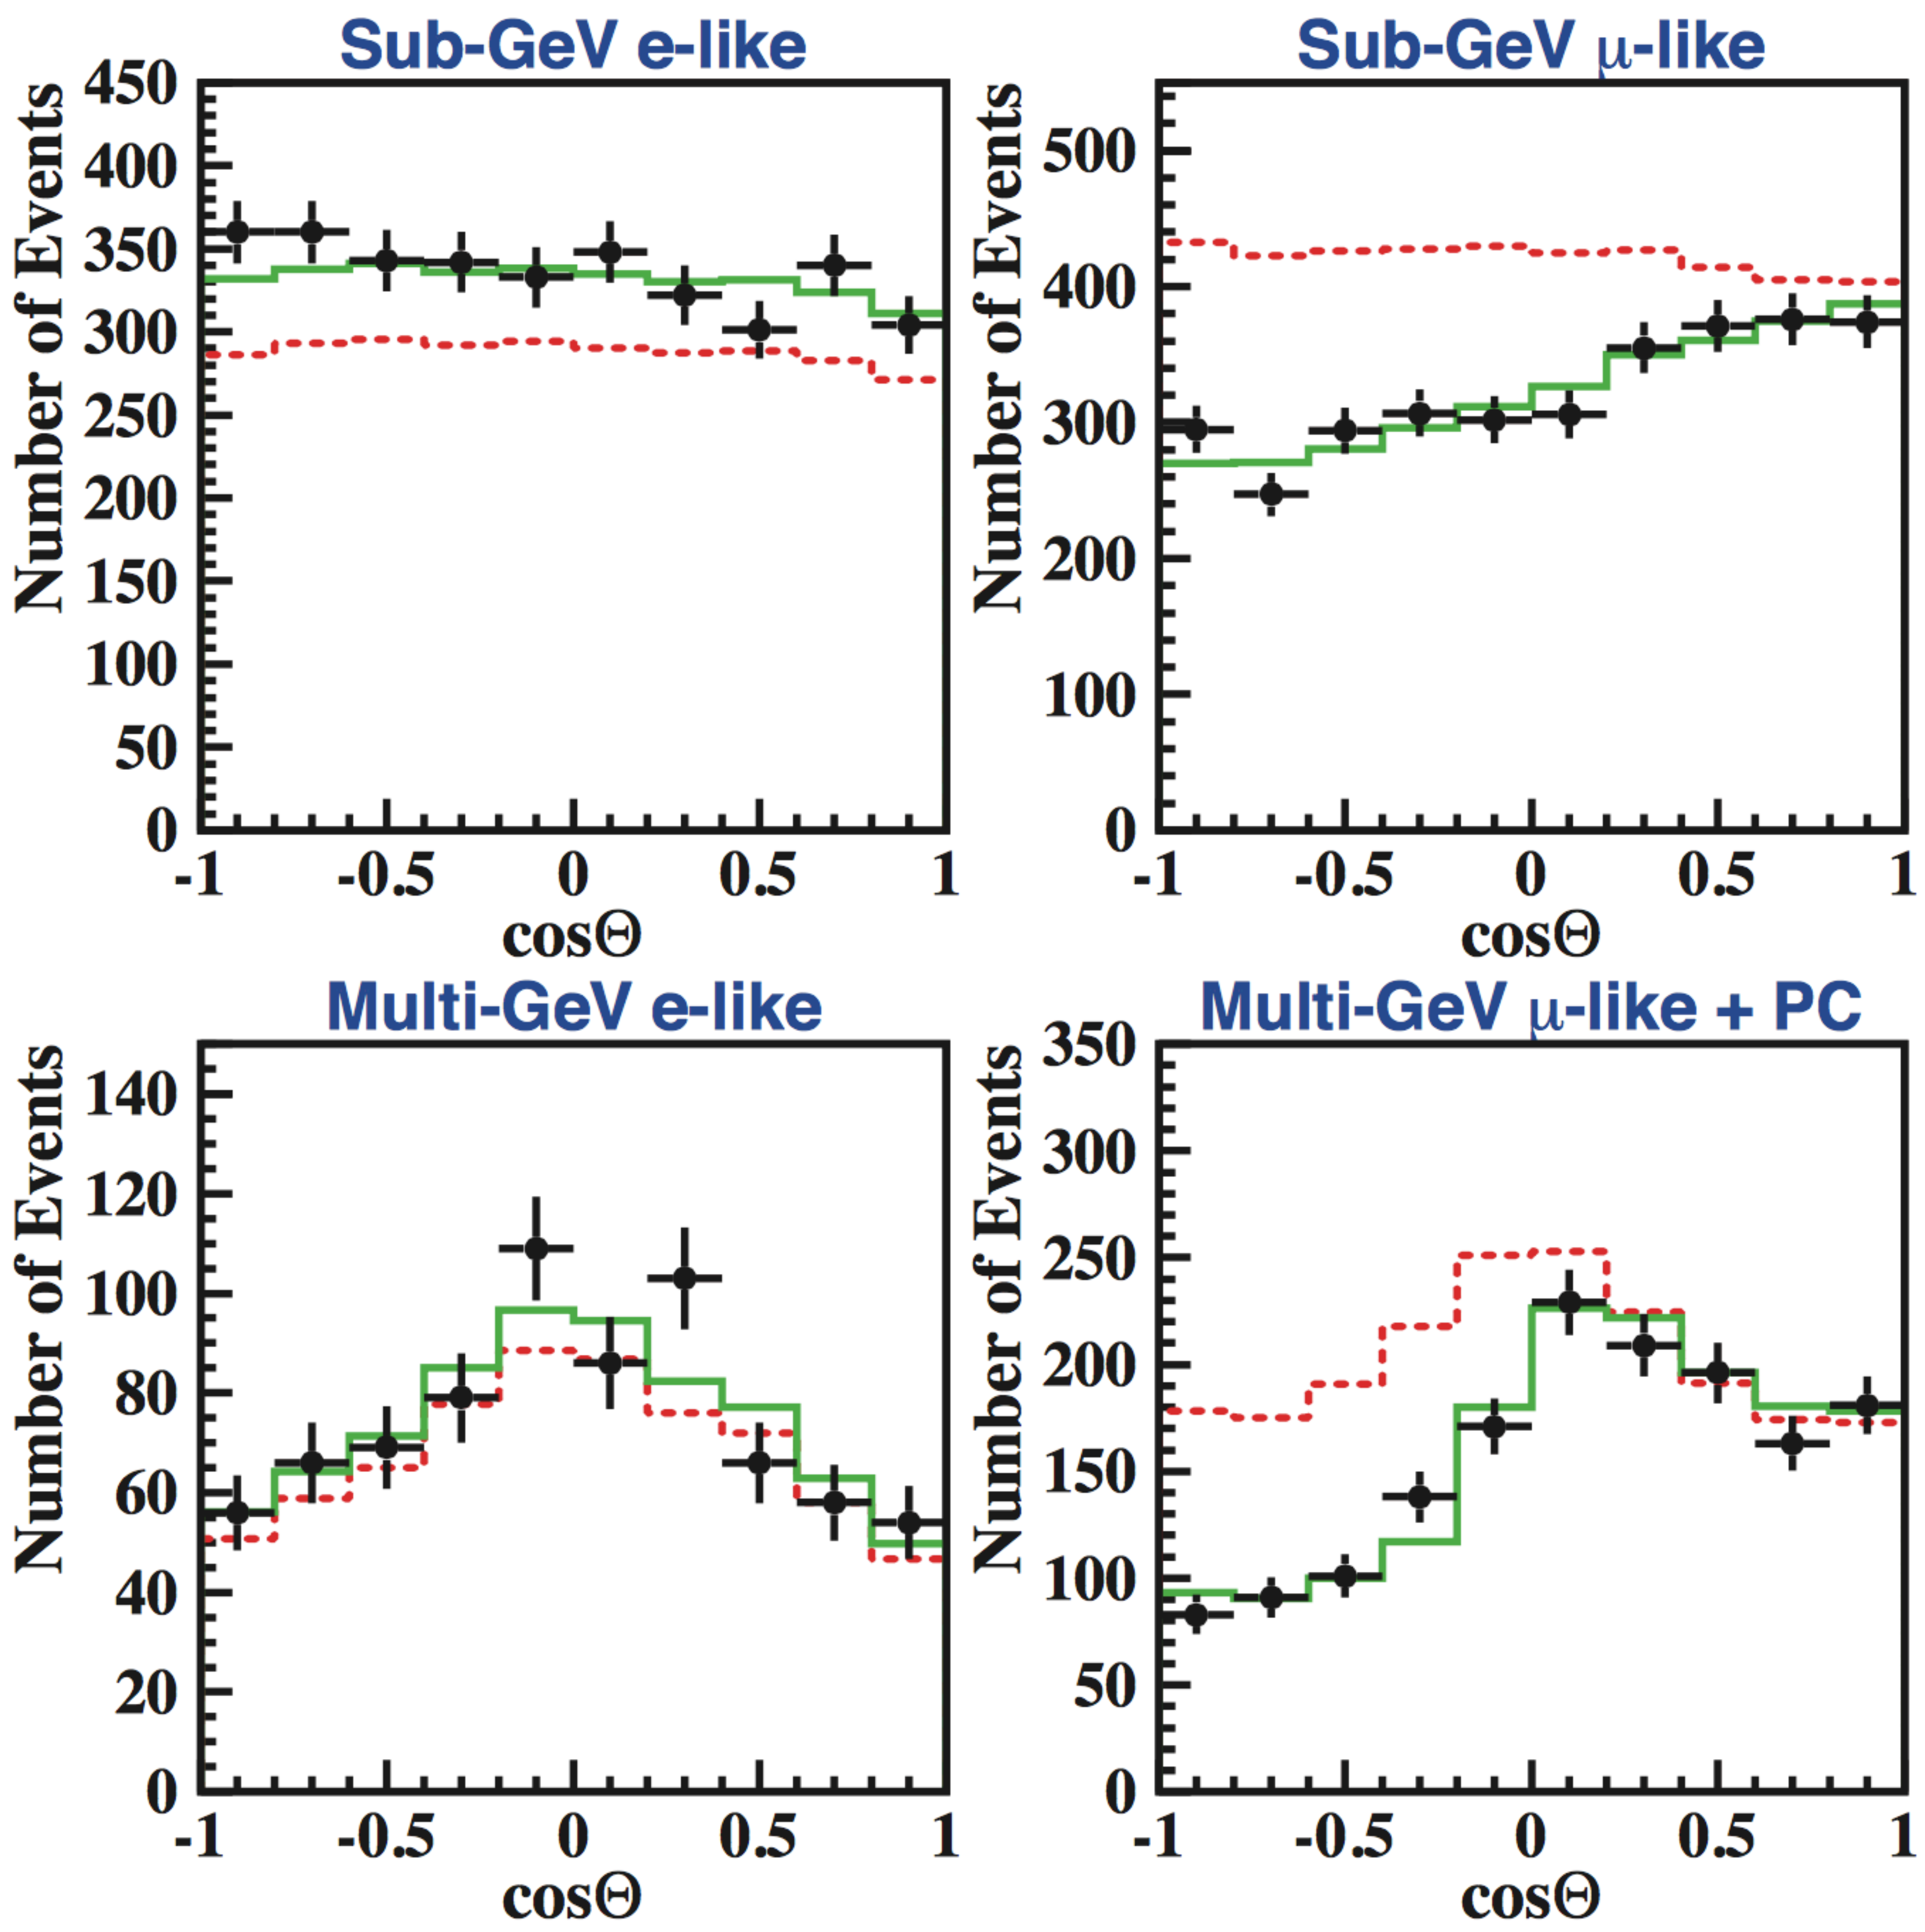

Limitations of the SM

Observational limitations

- Massless neutrinos
  $\Longrightarrow$ no oscillations
- No dark matter

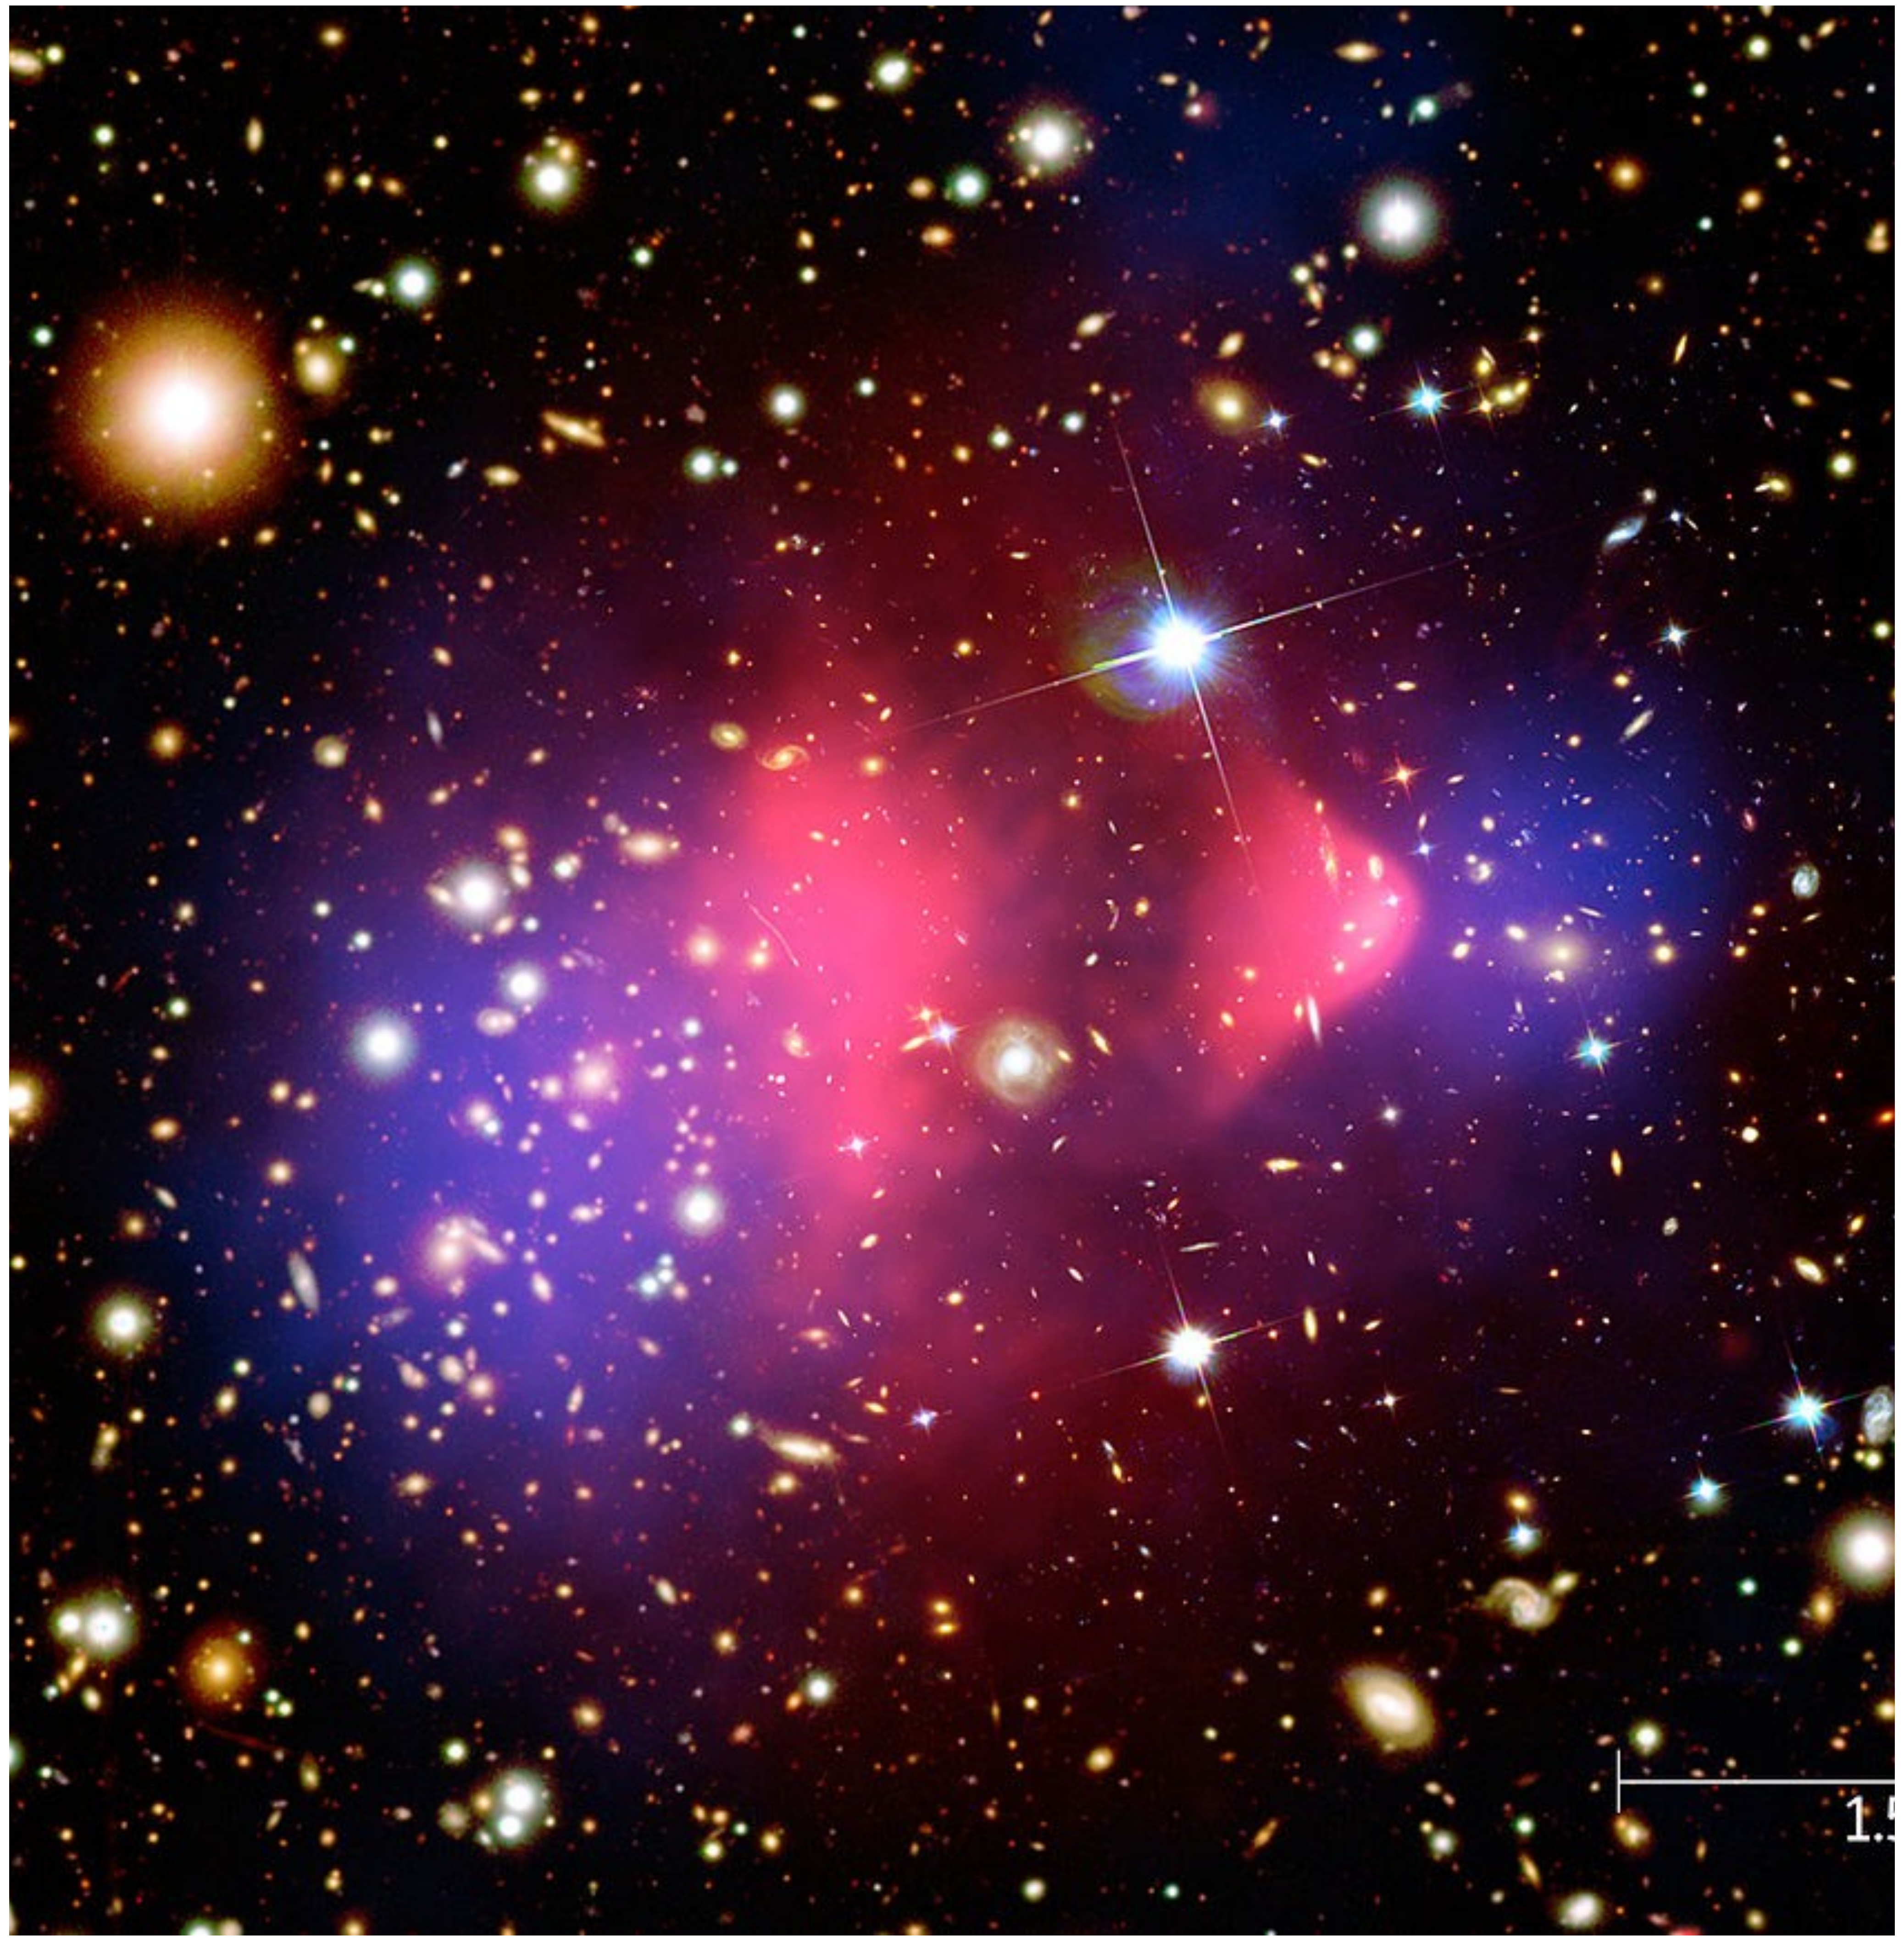

Limitations of the SM

Observational limitations

- Massless neutrinos
  $\implies$ no oscillations
- No dark matter
- No matter ($\eta = 0$)

Limitations of the SM

Limitations of the SM

Theoretical limitations

Limitations of the SM

Theoretical limitations

- Higgs naturalness

Limitations of the SM

Theoretical limitations

- Higgs naturalness
- Strong CP problem

Limitations of the SM

Theoretical limitations

- Higgs naturalness
- Strong CP problem
- Flavour puzzle

Limitations of the SM

Theoretical limitations

- Higgs naturalness
- Strong CP problem
- Flavour puzzle
- And more...

Limitations of the SM

[See e.g. de Gouvea, Hernandez, Tait: 1402.2658]

- Higgs naturalness

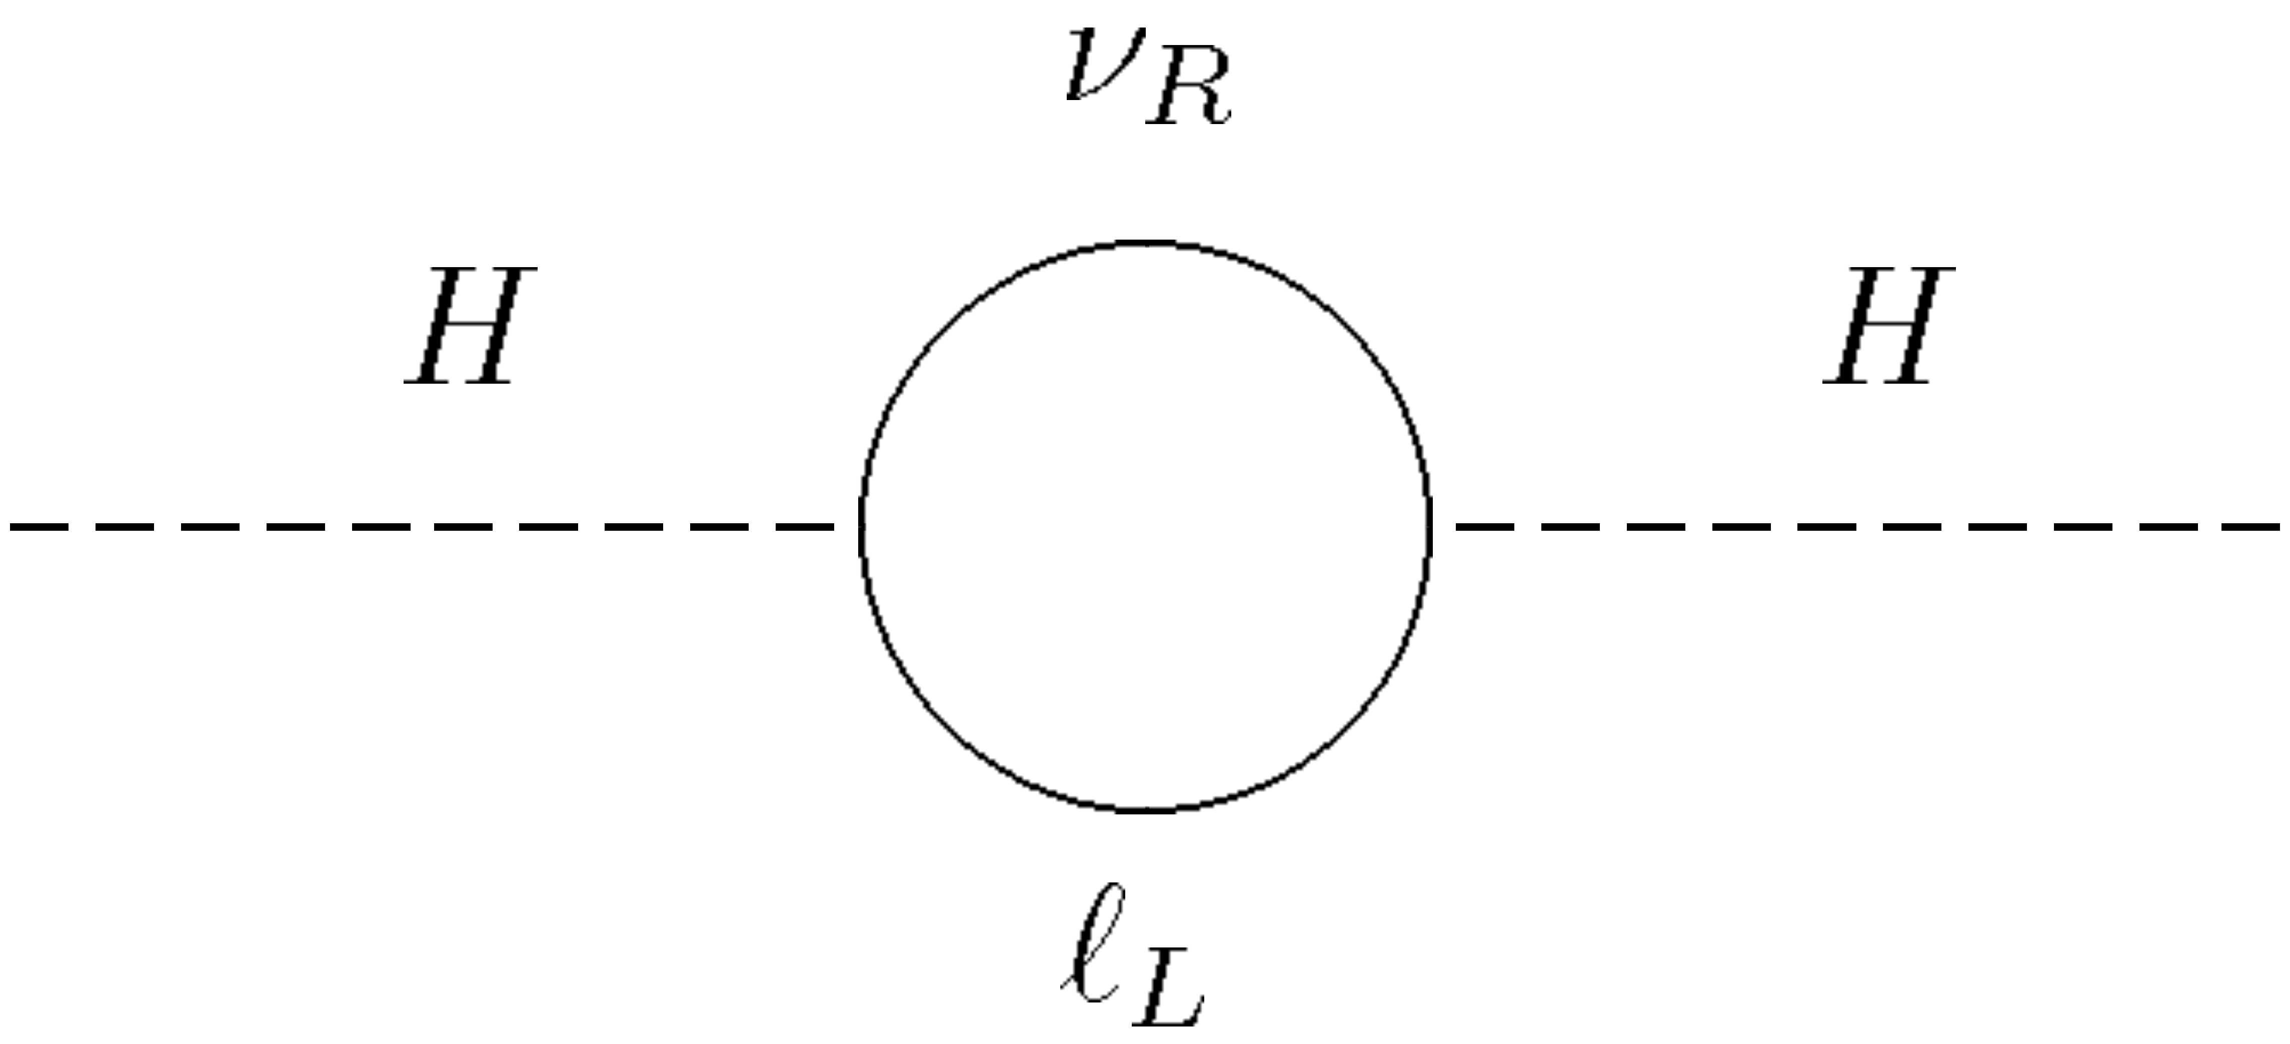

Threshold correction $\delta\mu^2 \sim \left(\frac{g^2}{16\pi^2}\right)^n M^2$

(n = loop order at which the new particle appears)

Possible solution:

New particles are light and/or feebly coupled to the Higgs

Heavy neutral leptons (HNLs)

Heavy neutral leptons (HNLs)

- No $SU(2)_L$ singlet ν_R in the SM

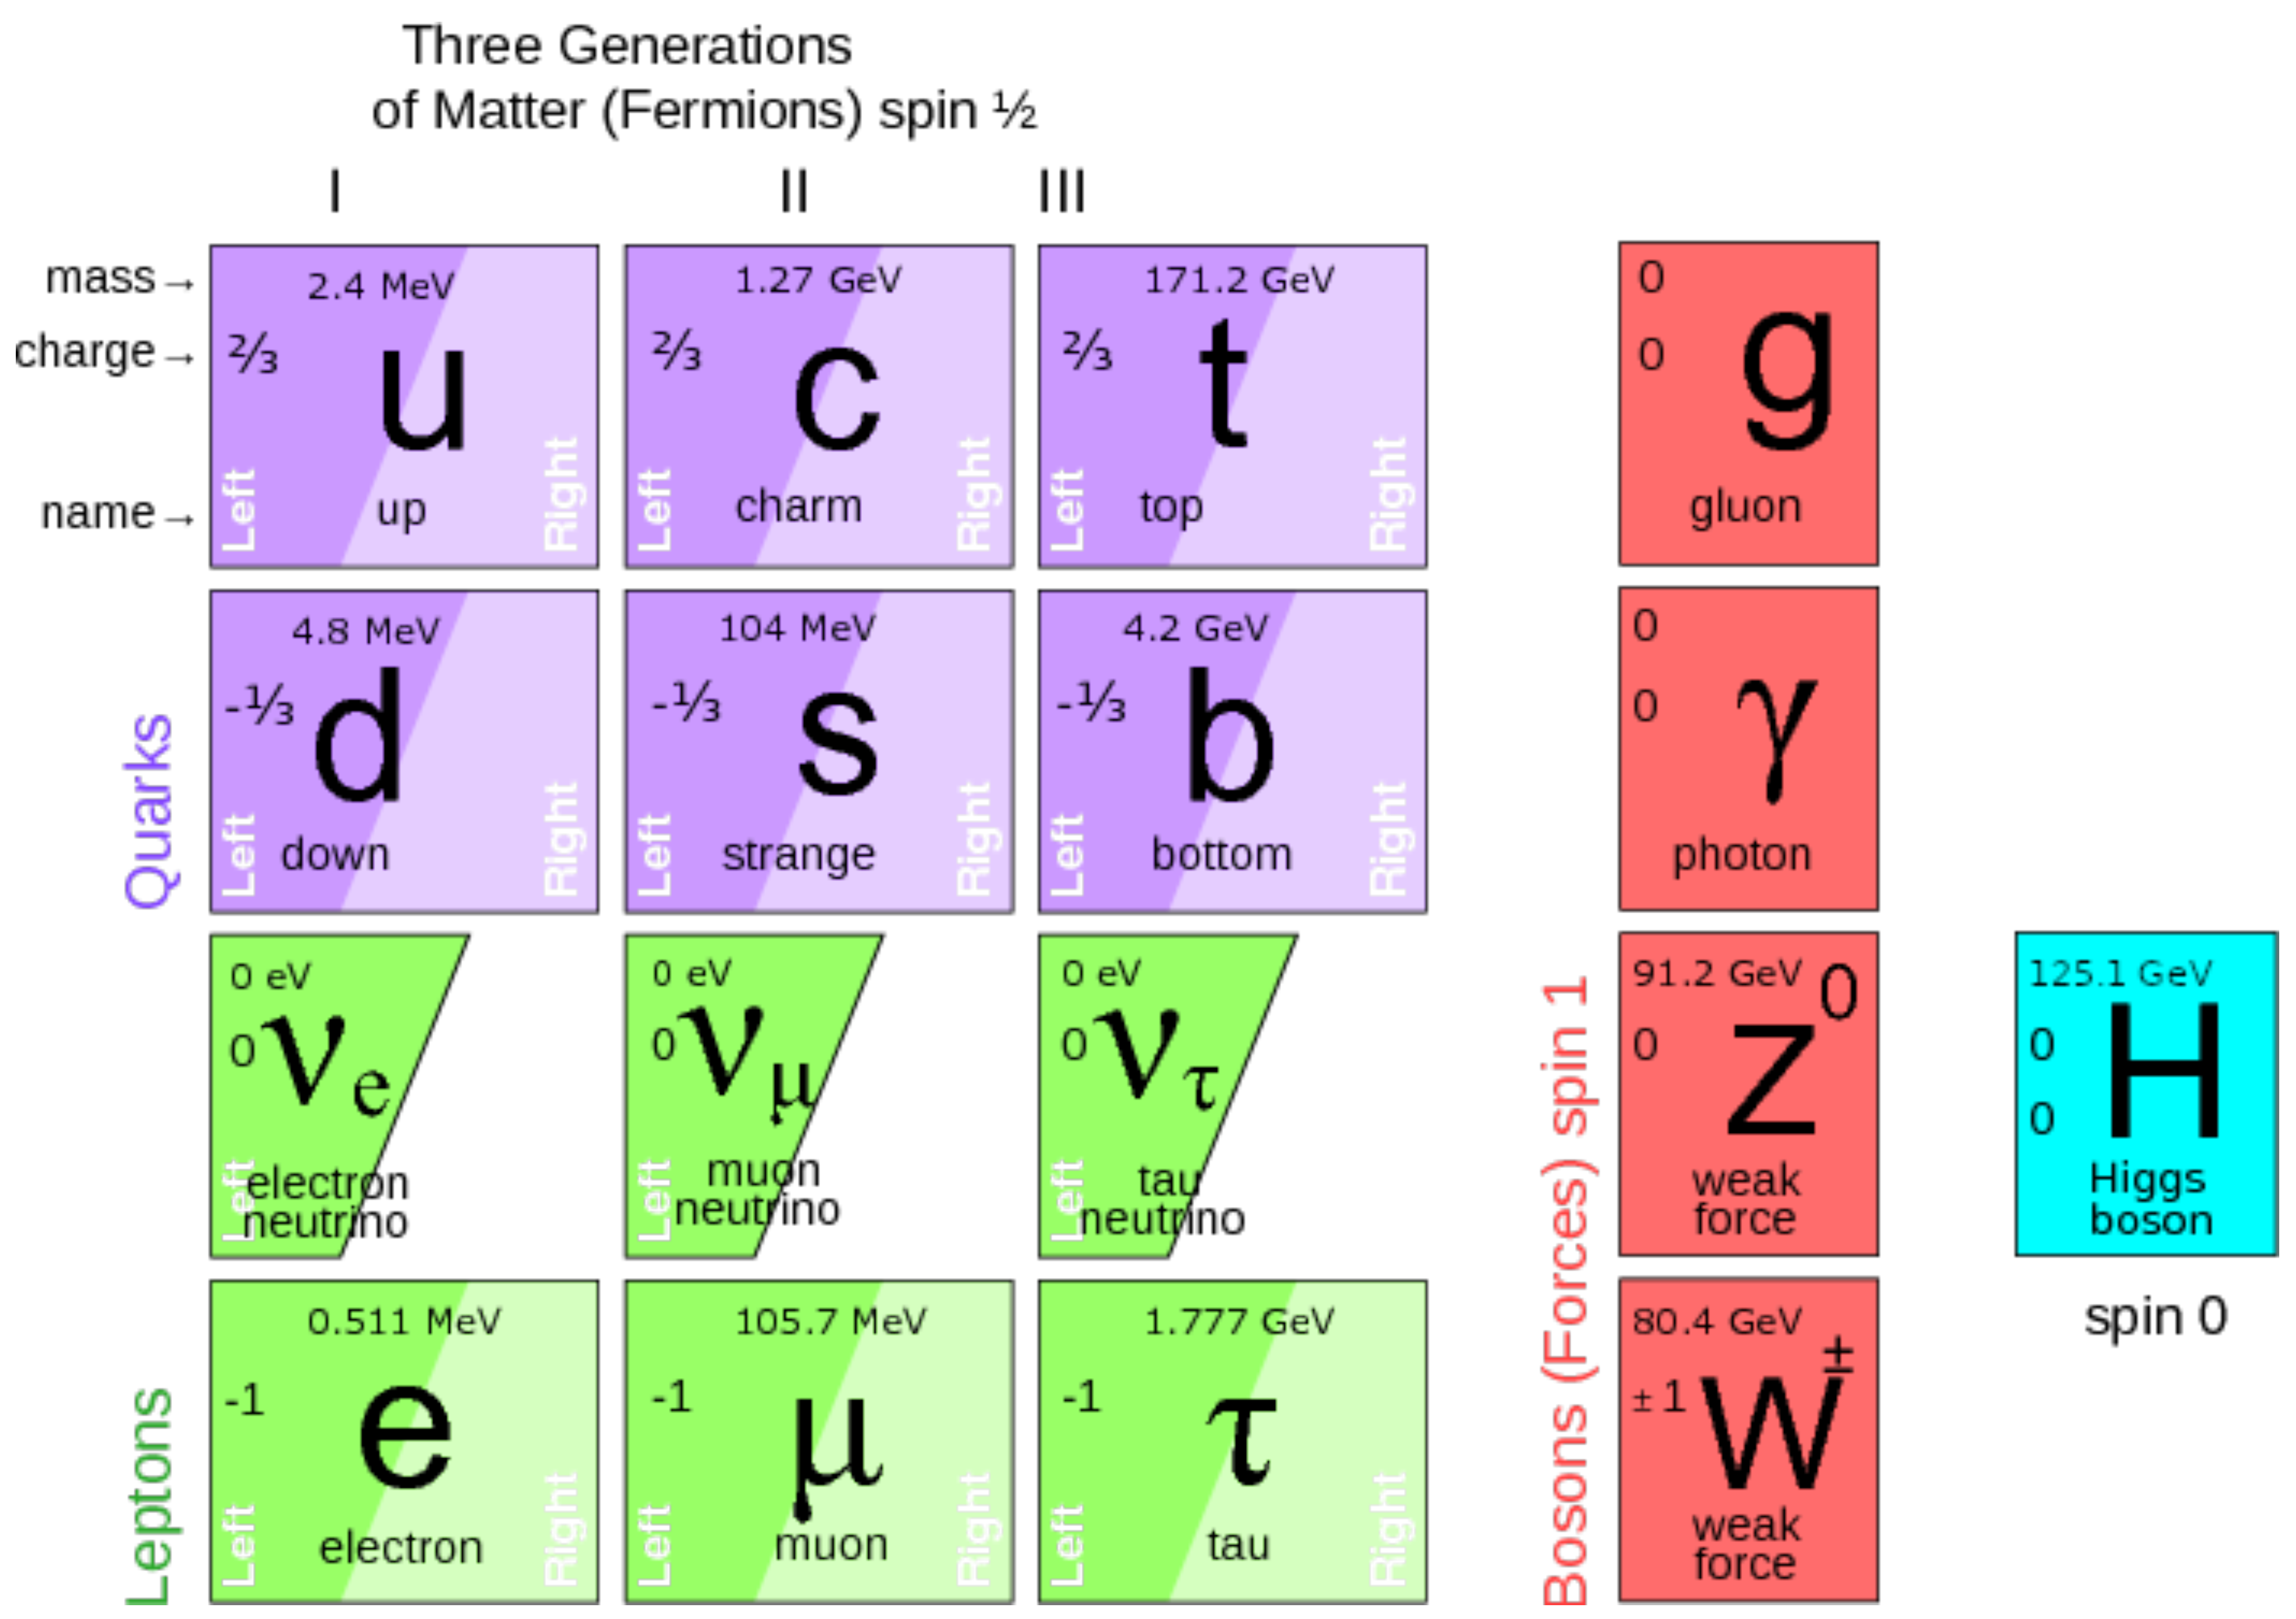

Heavy neutral leptons (HNLs)

- No $SU(2)_L$ singlet ν_R in the SM
- Simplest addition which can give a mass to neutrinos:

$$-(Y^{\nu}_{\alpha I})^{*}(L_{\alpha} \cdot \tilde{\phi}^{\dagger})\nu_{R,I} \longrightarrow (m_D)_{\alpha I}\nu_{L,\alpha}\nu_{R,I}$$

with the Dirac mass $m_D = \frac{v}{\sqrt{2}}(Y^{\nu}_{\alpha I})^{*}$

(where $\alpha = e, \mu, \tau$, $I = 1,2,\ldots,N_{\mathrm{HNL}}$)

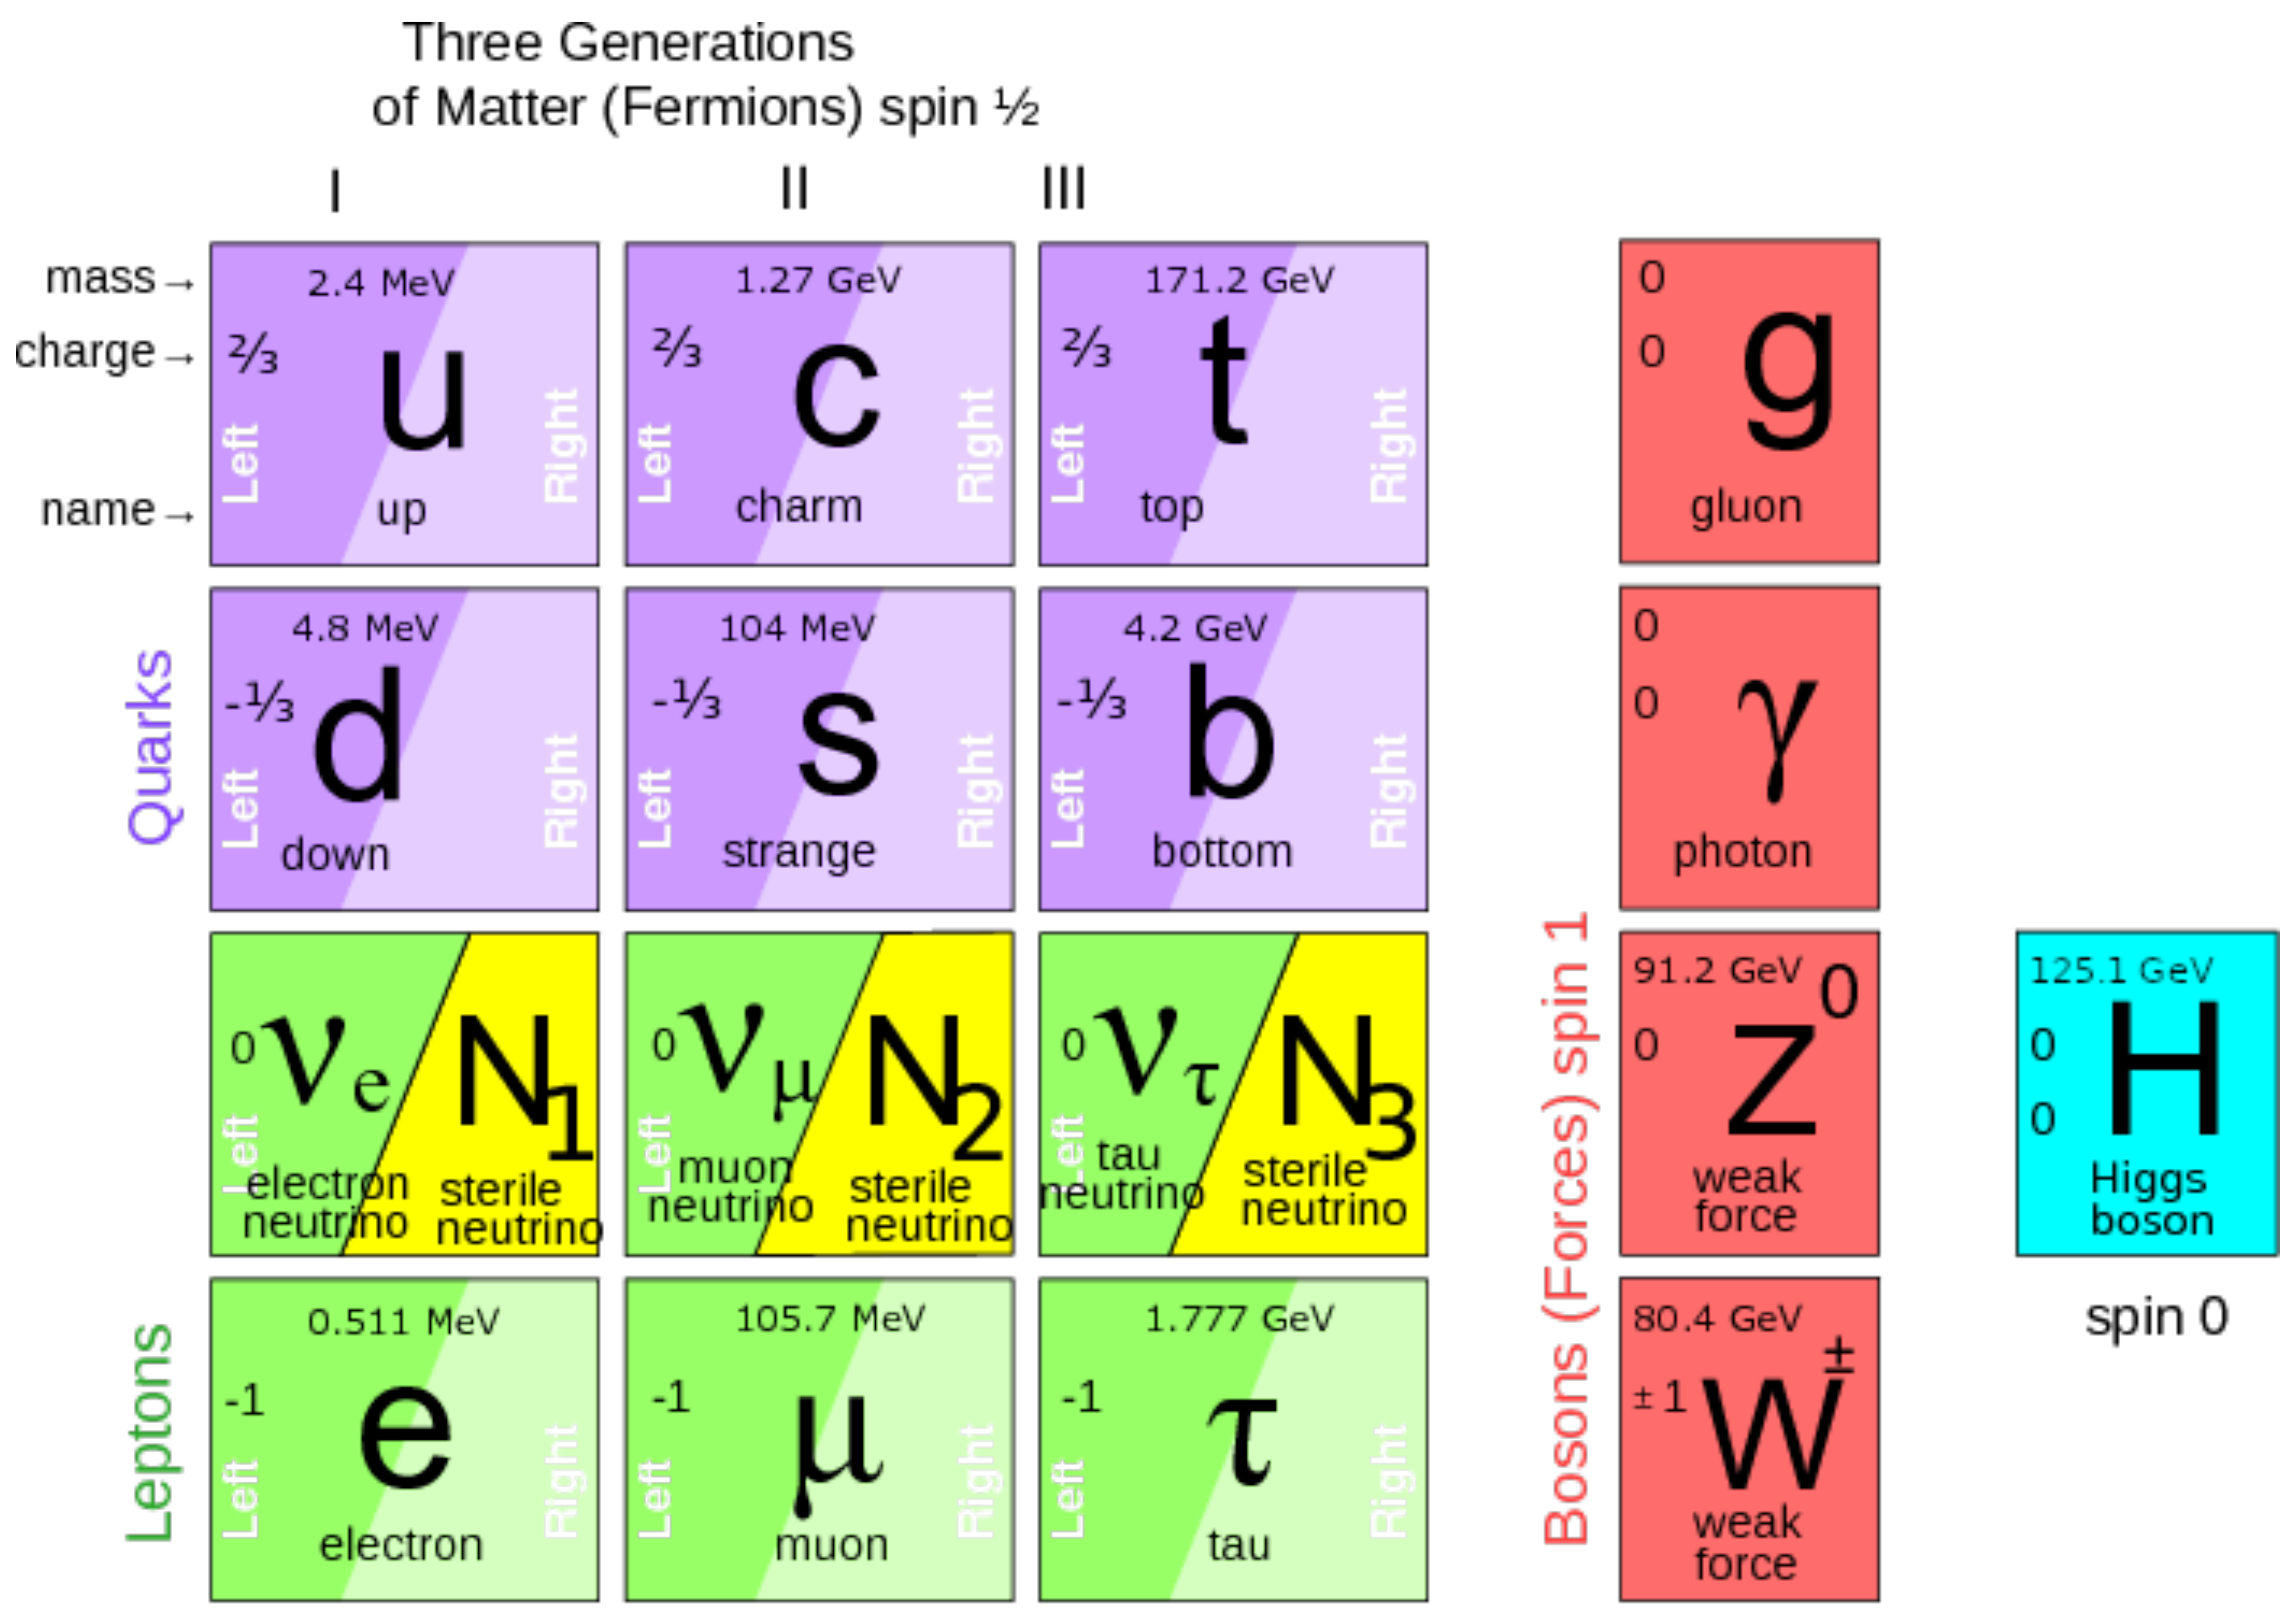

Heavy neutral leptons (HNLs)

- No $SU(2)_L$ singlet ν_R in the SM
- Simplest addition which can give a mass to neutrinos:
  $-(Y^\nu_{\alpha I})^*(L_\alpha \cdot \tilde{\phi}^\dagger)\nu_{R,I} \longrightarrow (m_D)_{\alpha I}\nu_{L,\alpha}\nu_{R,I}$
  with the Dirac mass $m_D = \frac{v}{\sqrt{2}}(Y^\nu_{\alpha I})^*$
  (where $\alpha = e, \mu, \tau$, $I = 1,2,\ldots,N_{\mathrm{HNL}}$)
- SM singlets can have a Majorana mass:
  $-\frac{M_I}{2}(\nu_{R,I}\nu_{R,I} + \nu^\dagger_{R,I}\nu^\dagger_{R,I})$

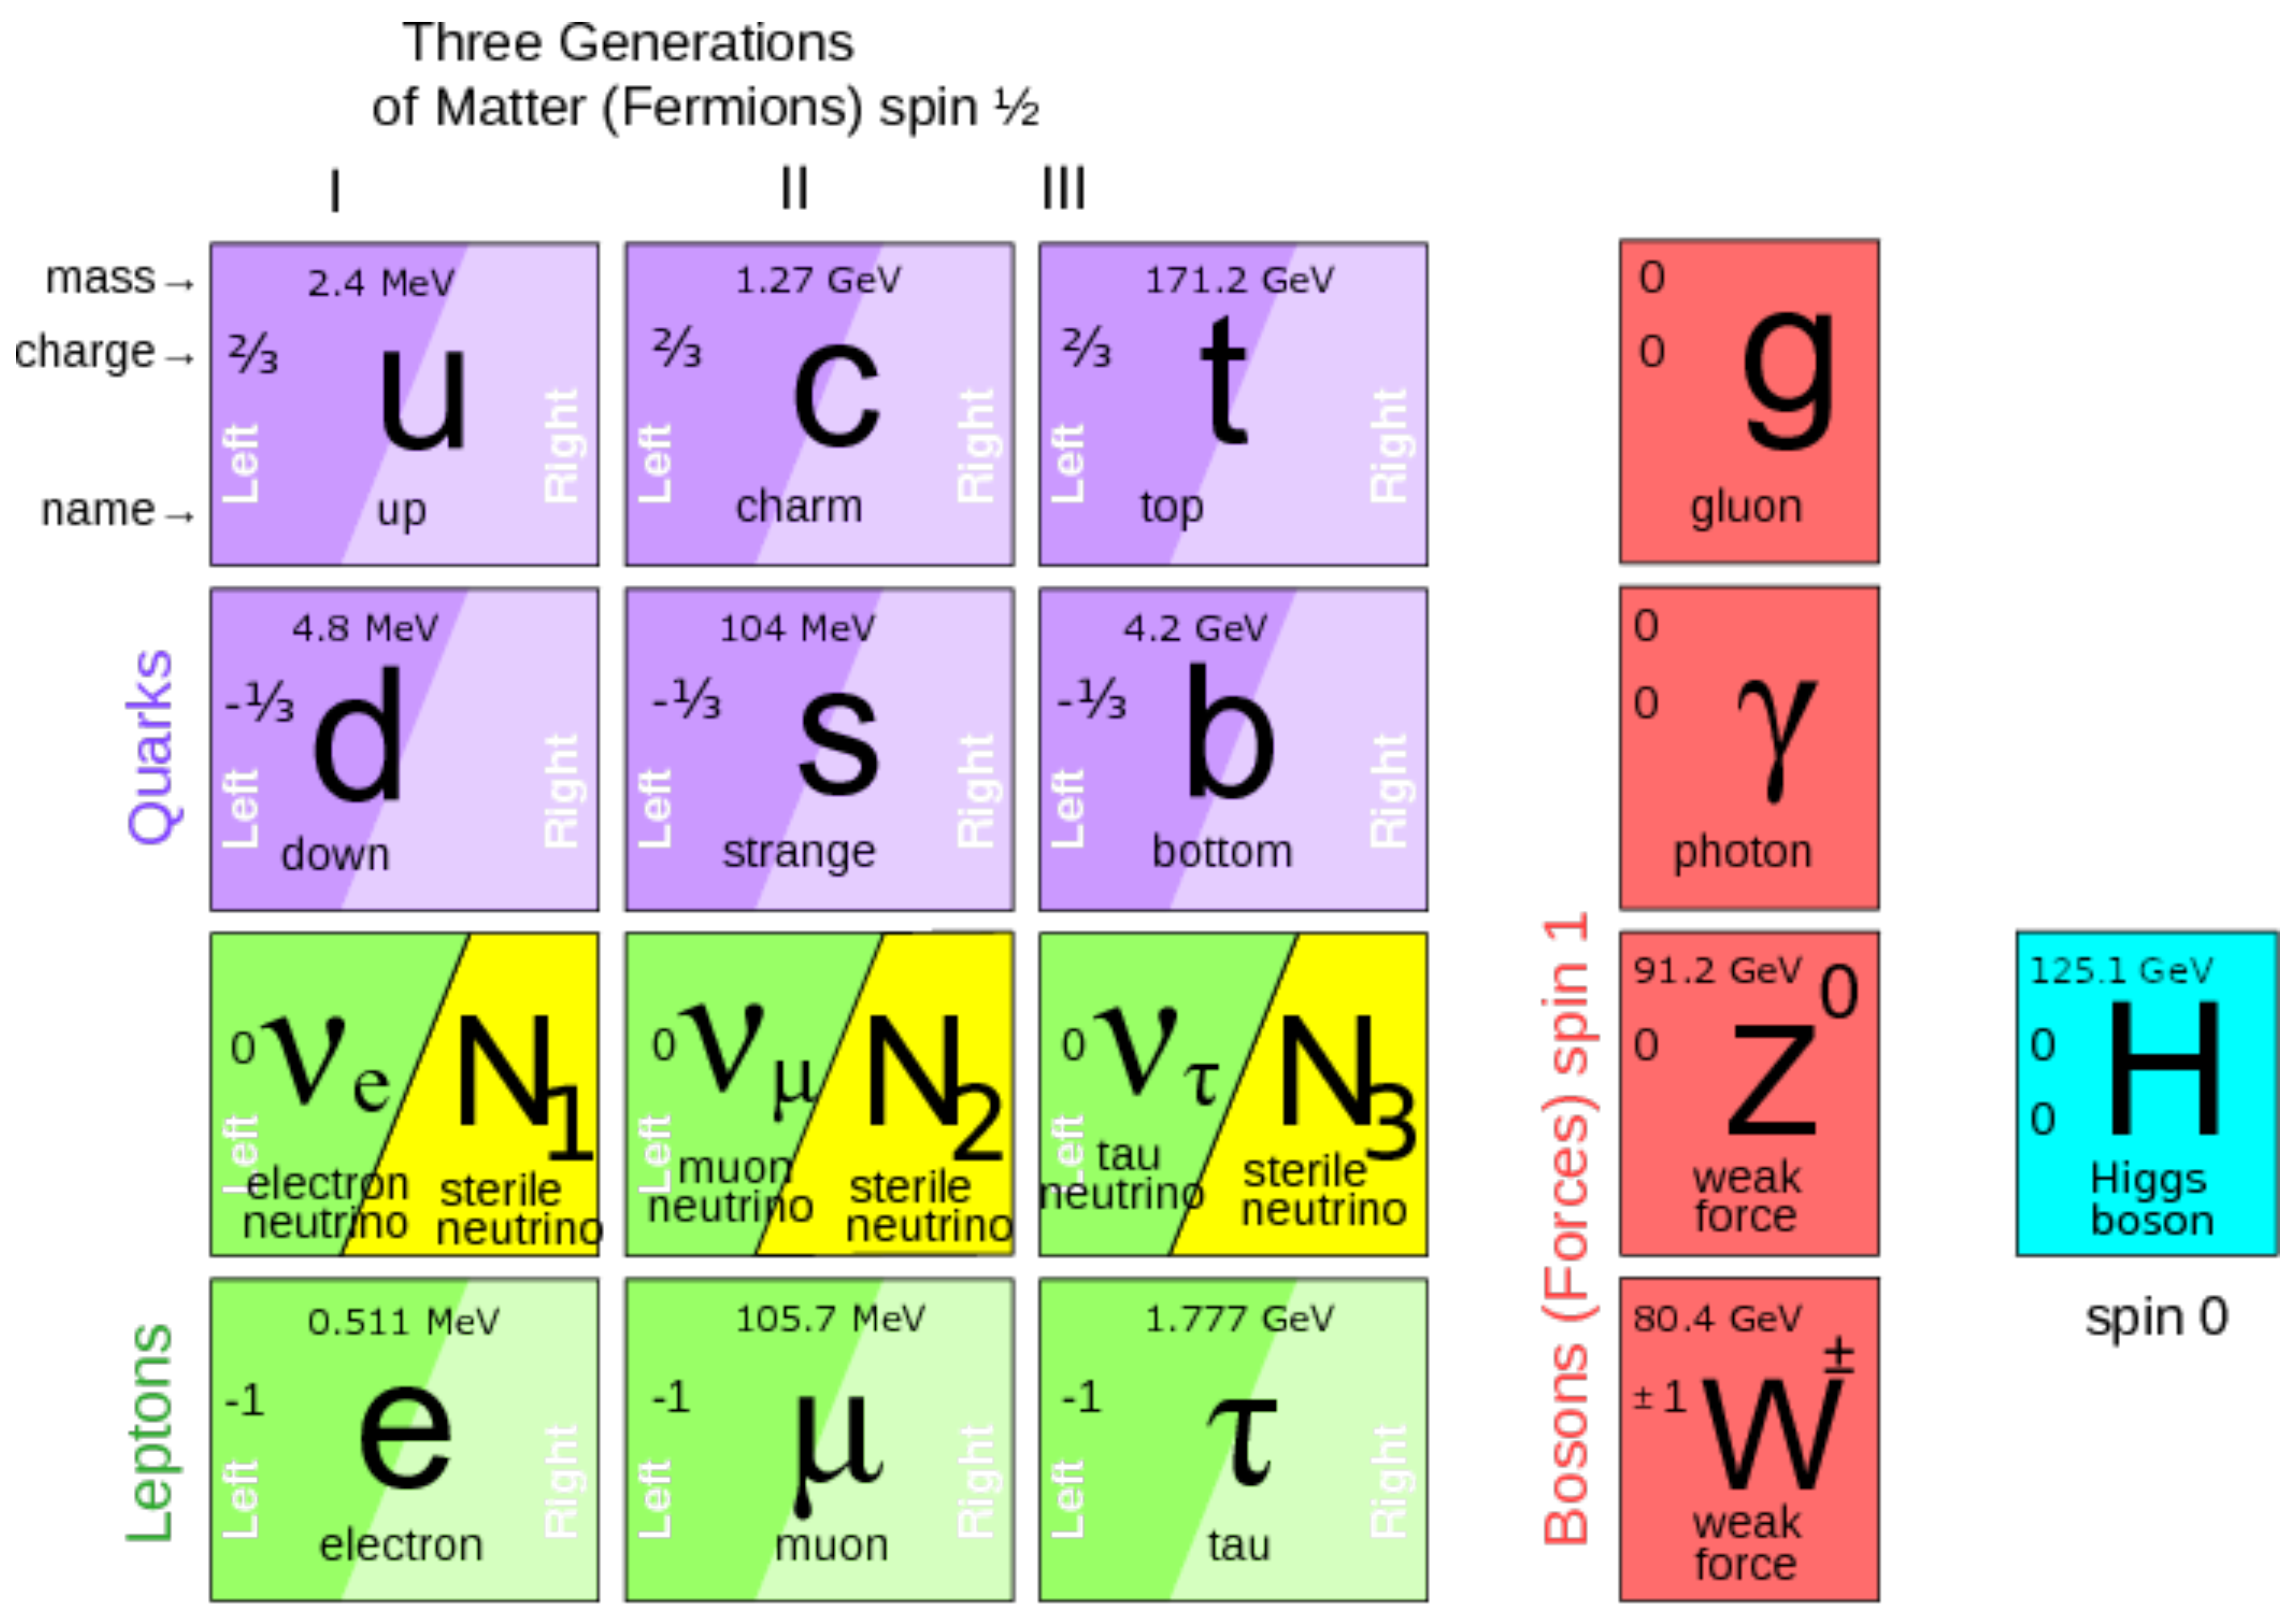

The type-I see-saw mechanism

Minkowski, Gell-Mann, Ramond, Slansky, Mohapatra, Senjanovic, Yanagida, Schechter, Valle, Shrock, ...

- Both mass terms are allowed: $-\frac{1}{2}\begin{pmatrix}\nu_L^T & \nu_R^T\end{pmatrix}\begin{pmatrix}0 & m_D^T \\ m_D & M_R\end{pmatrix}\begin{pmatrix}\nu_L \\ \nu_R\end{pmatrix} + \mathrm{h.c.}$

- Mass diagonalisation leads to mixing: $\nu_{L,\alpha} \cong U^{\mathrm{PMNS}}_{\alpha,i}\nu_i + \Theta_{\alpha,I}\nu_{R,I}$

- Neutrinos are light if HNLs are heavy, i.e. $M_R \gg m_D$ (or $\Theta \ll 1$)
  Their masses are given by the see-saw formula:

$$m^{\mathrm{light}}_{\alpha\beta} \approx -\sum_I \frac{(m_D)_{\alpha I}(m_D)_{\beta I}}{M_I} \approx -\sum_I M_I \Theta_{\alpha I}\Theta_{\beta I}$$

HNLs $\nu_{R,I}$ behave as:

- **Heavy neutrinos (Majorana or pseudo-Dirac)**
- **With the same interactions as light neutrinos $\nu_{L,\alpha}$**
- **But suppressed by a small mixing parameter $\Theta_{\alpha I}$**

HNLs $\nu_{R,I}$ behave as:

- **Heavy neutrinos (Majorana or pseudo-Dirac)**
- **With the same interactions as light neutrinos $\nu_{L,\alpha}$**
- **But suppressed by a small mixing parameter $\Theta_{\alpha I}$**

Prototypical example of a feebly interacting particle (FIP)

Will often be denoted by "N_I"

HNLs $\nu_{R,I}$ behave as:

- **Heavy neutrinos (Majorana or pseudo-Dirac)**
- **With the same interactions as light neutrinos $\nu_{L,\alpha}$**
- **But suppressed by a small mixing parameter $\Theta_{\alpha I}$**

Prototypical example of a feebly interacting particle (FIP)

Addressing the observational problems

- Small mass & mixing angle can make N_1 a metastable DM candidate.
- Two HNLs forming a quasi-Dirac pair can undergo CP-violating oscillations before decaying, potentially leading to successful baryogenesis.
  [Akhmedov, Rubakov, Smirnov ’98]
- Only 3 HNLs with the right parameters are sufficient to explain all the above.
  [Asaka, Shaposhnikov: hep-ph/0505013]

Type-I see-saw parameter space

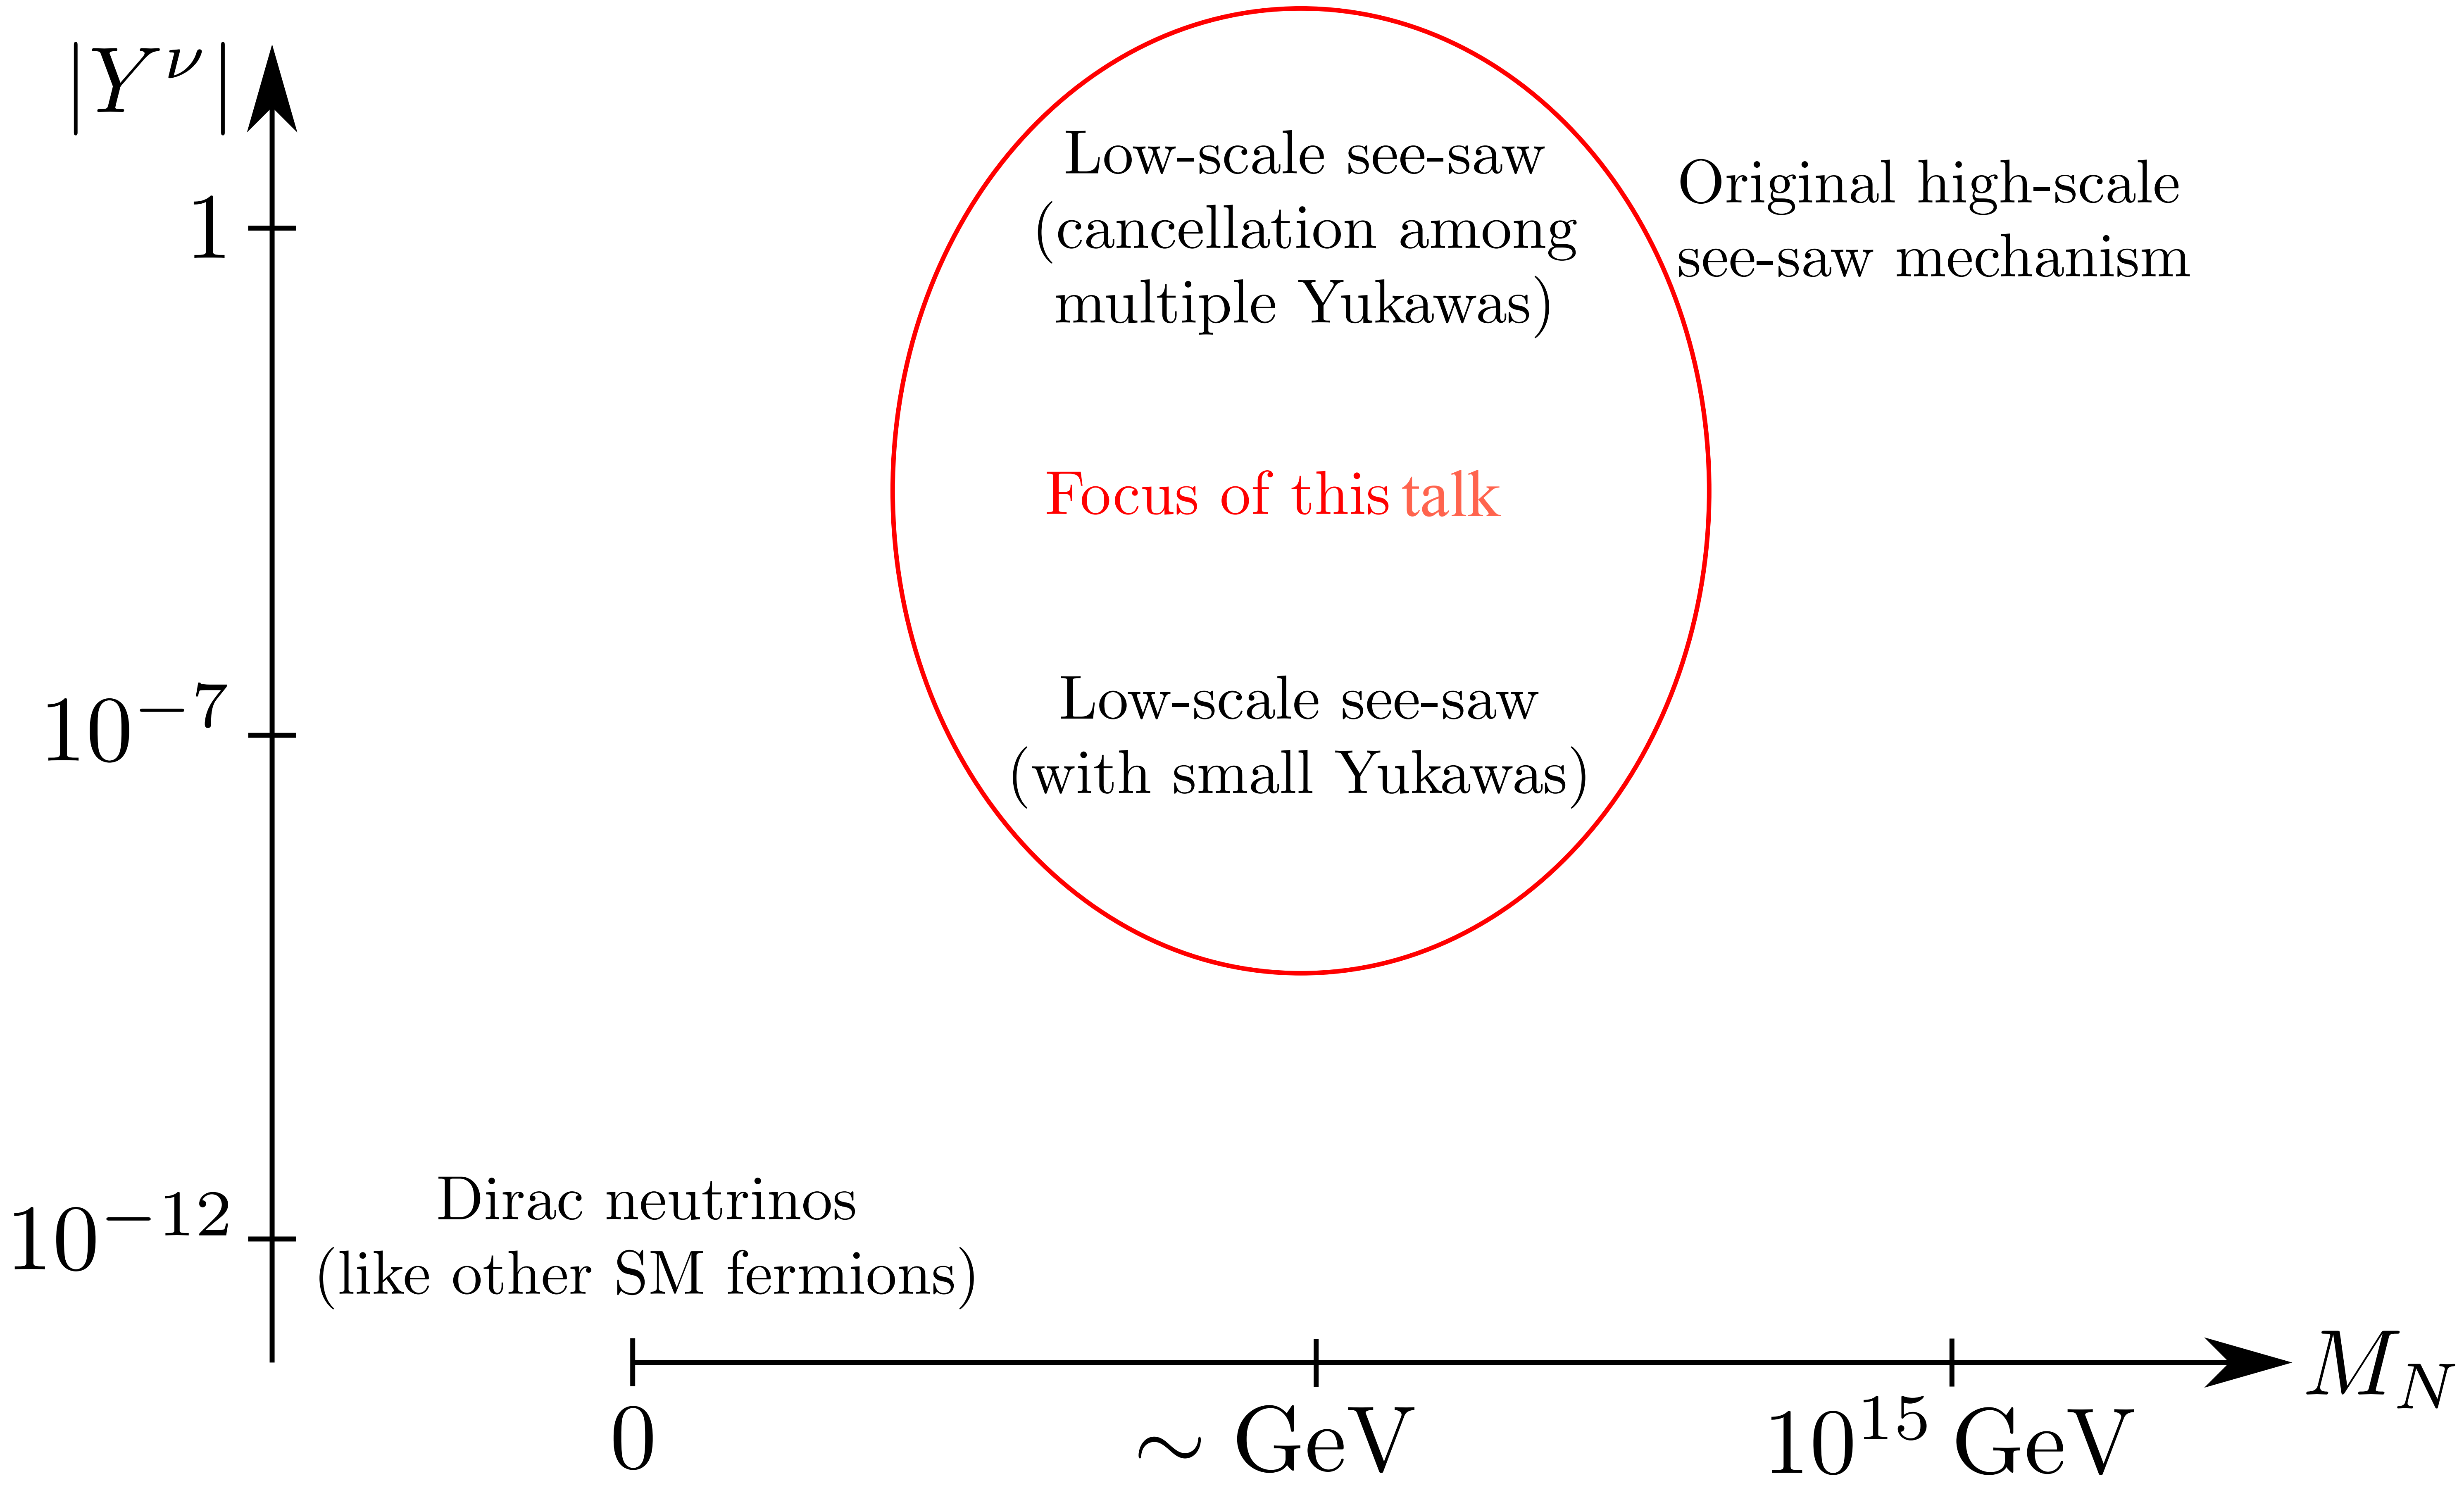

Type-I see-saw parameter space

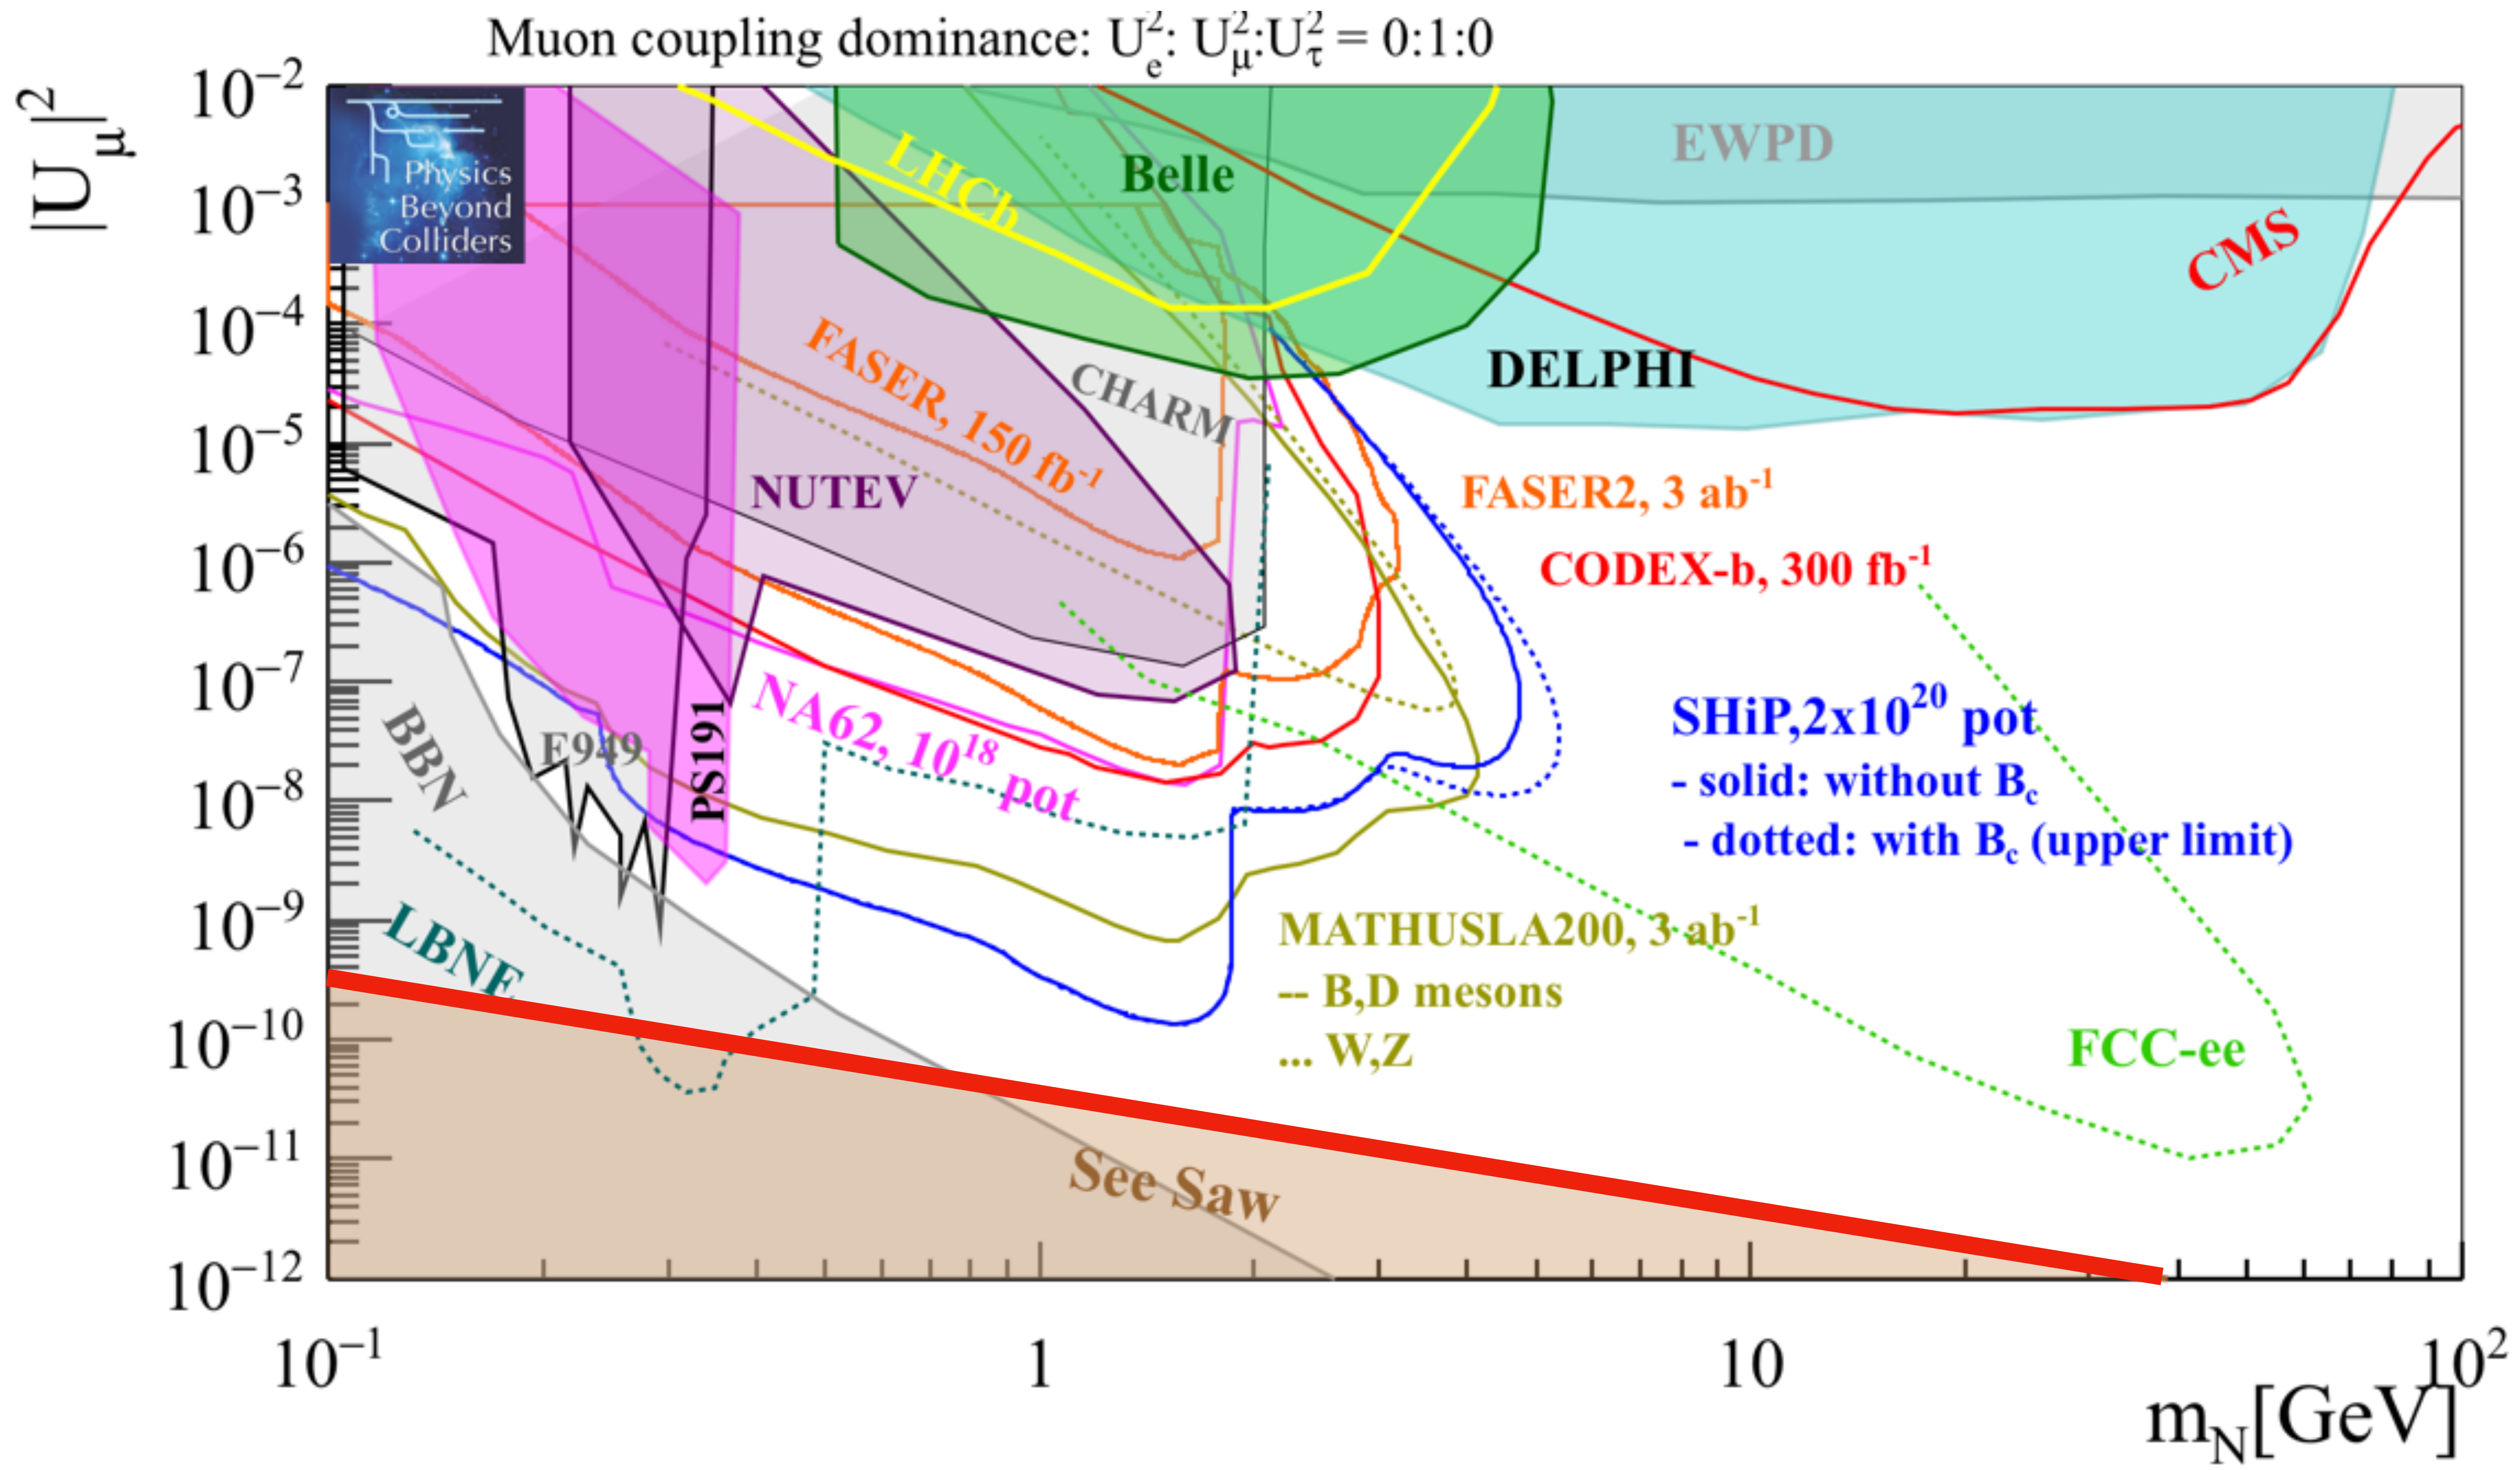

Type-I see-saw parameter space

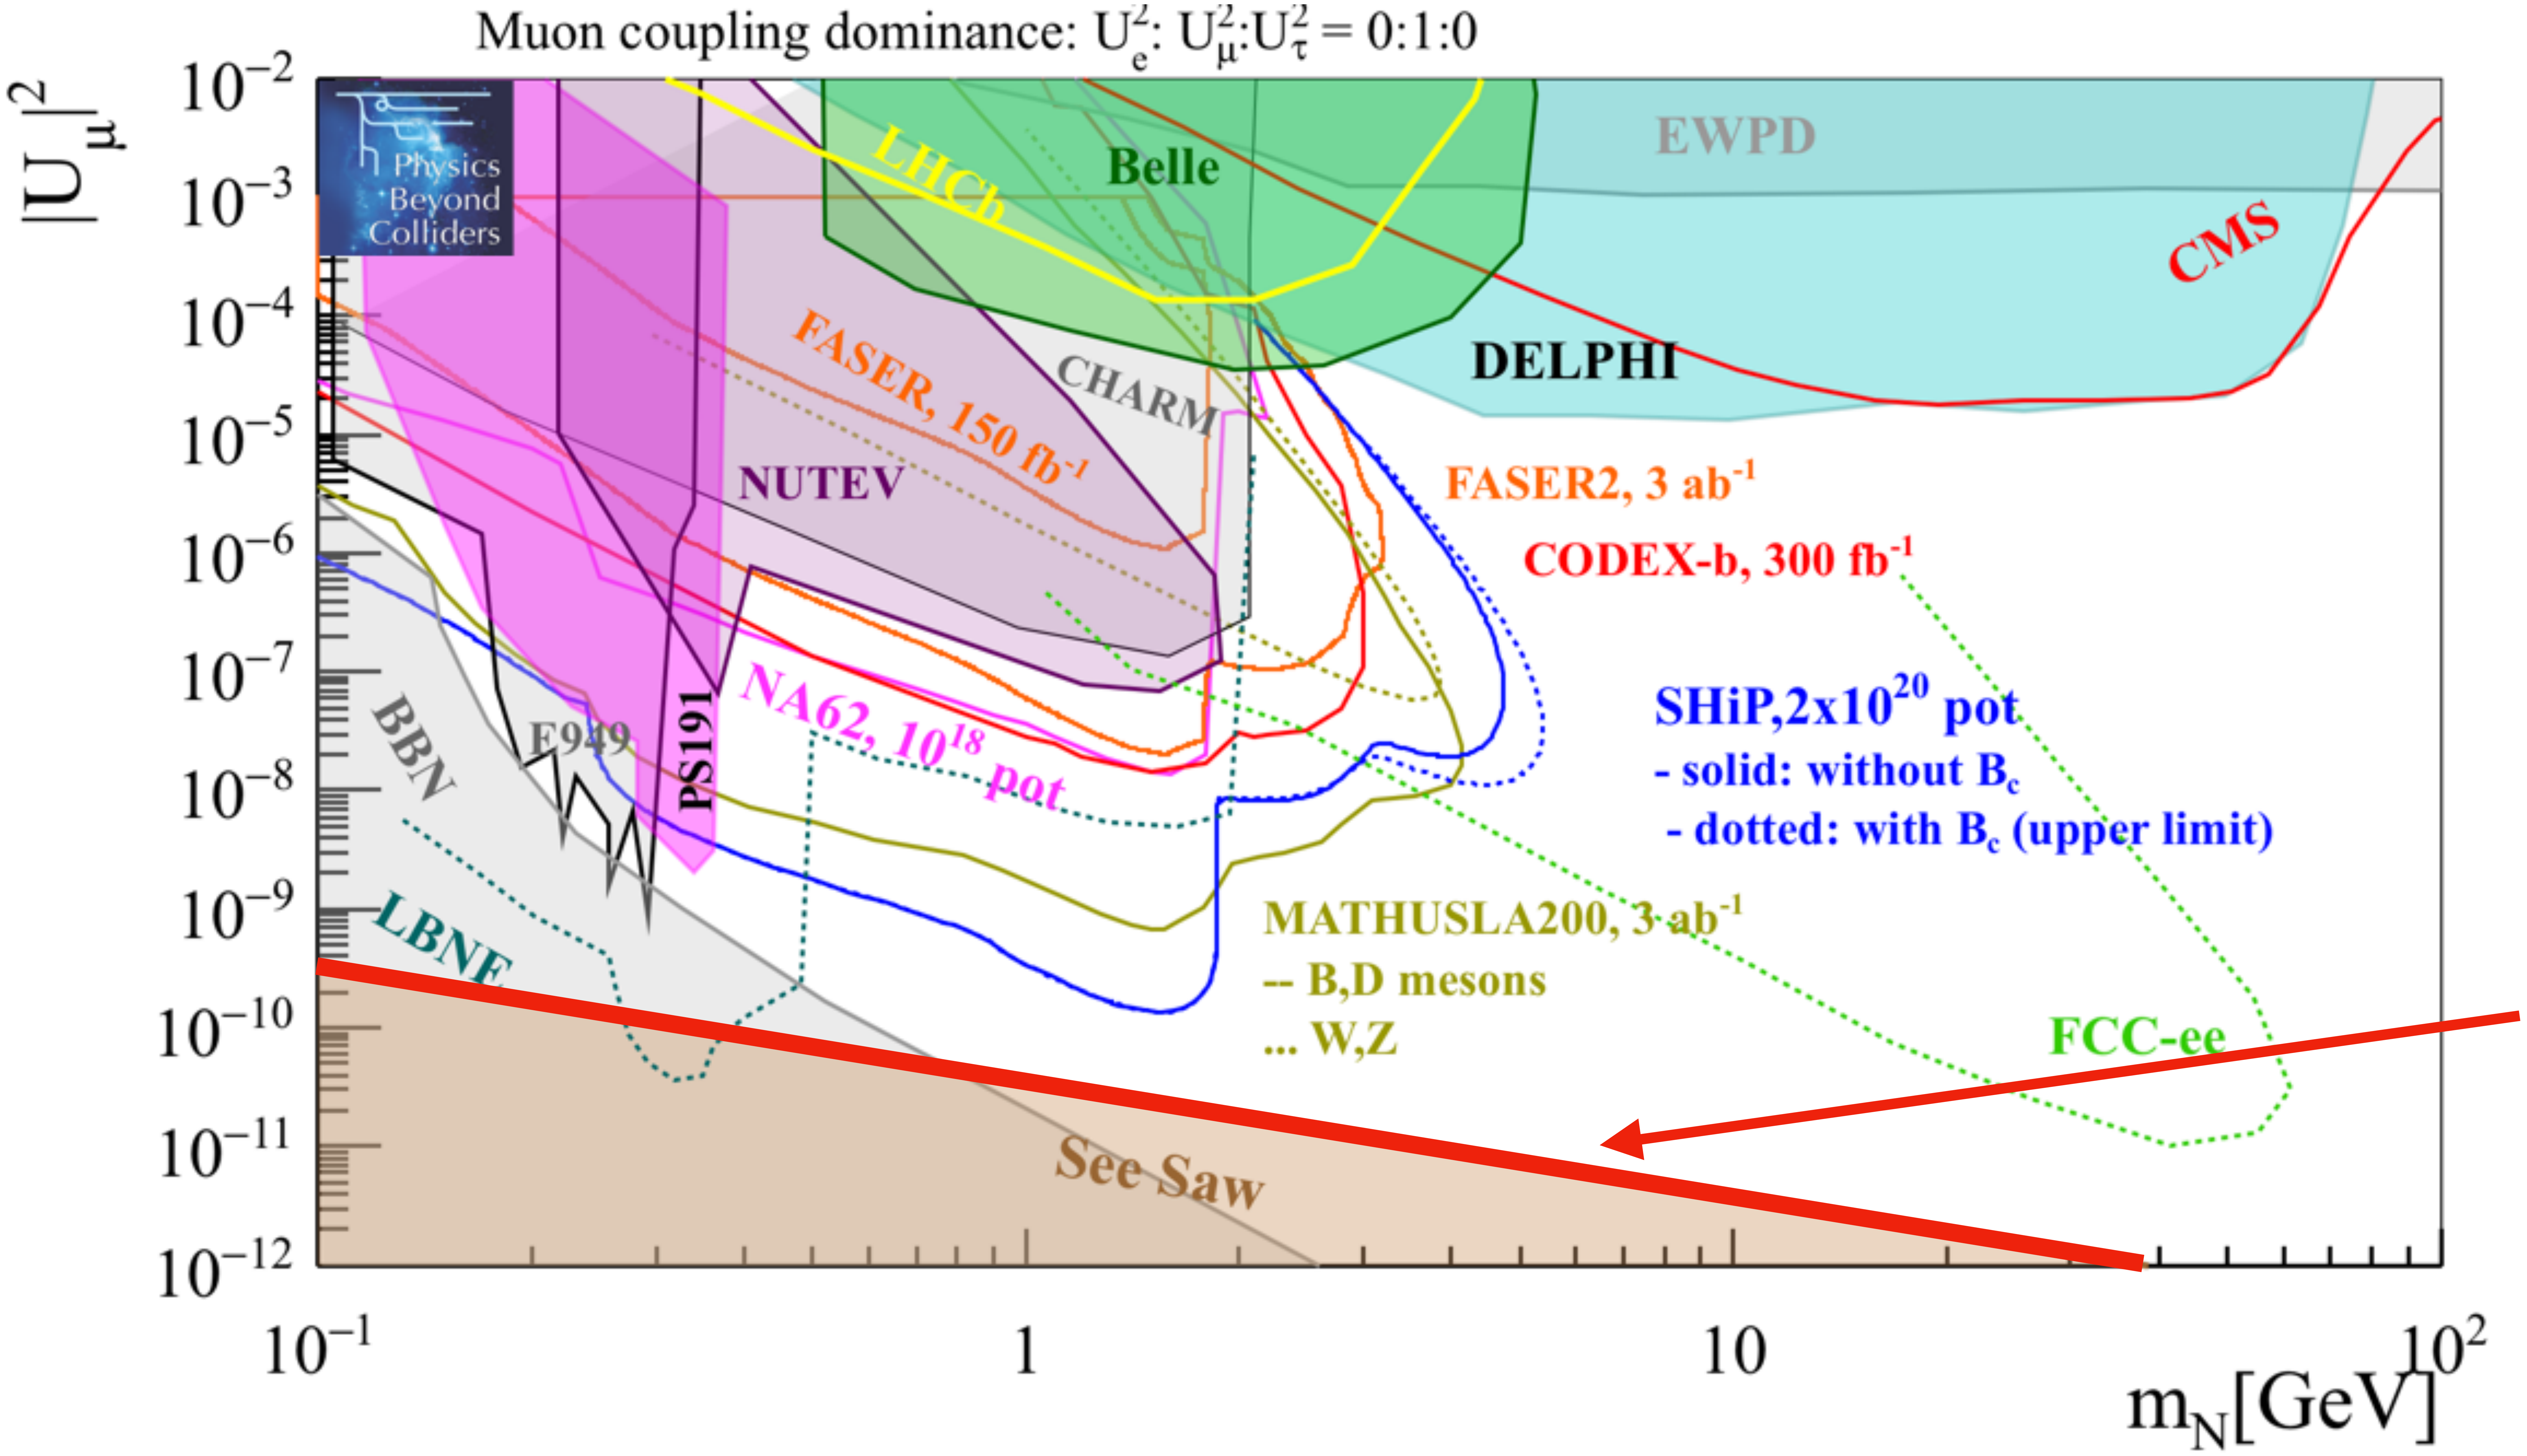

See-saw line:
"naive" expectation
$m^{\rm light} \sim m_N \times |\Theta|^2$
Mostly out of reach 😢

Type-I see-saw parameter space

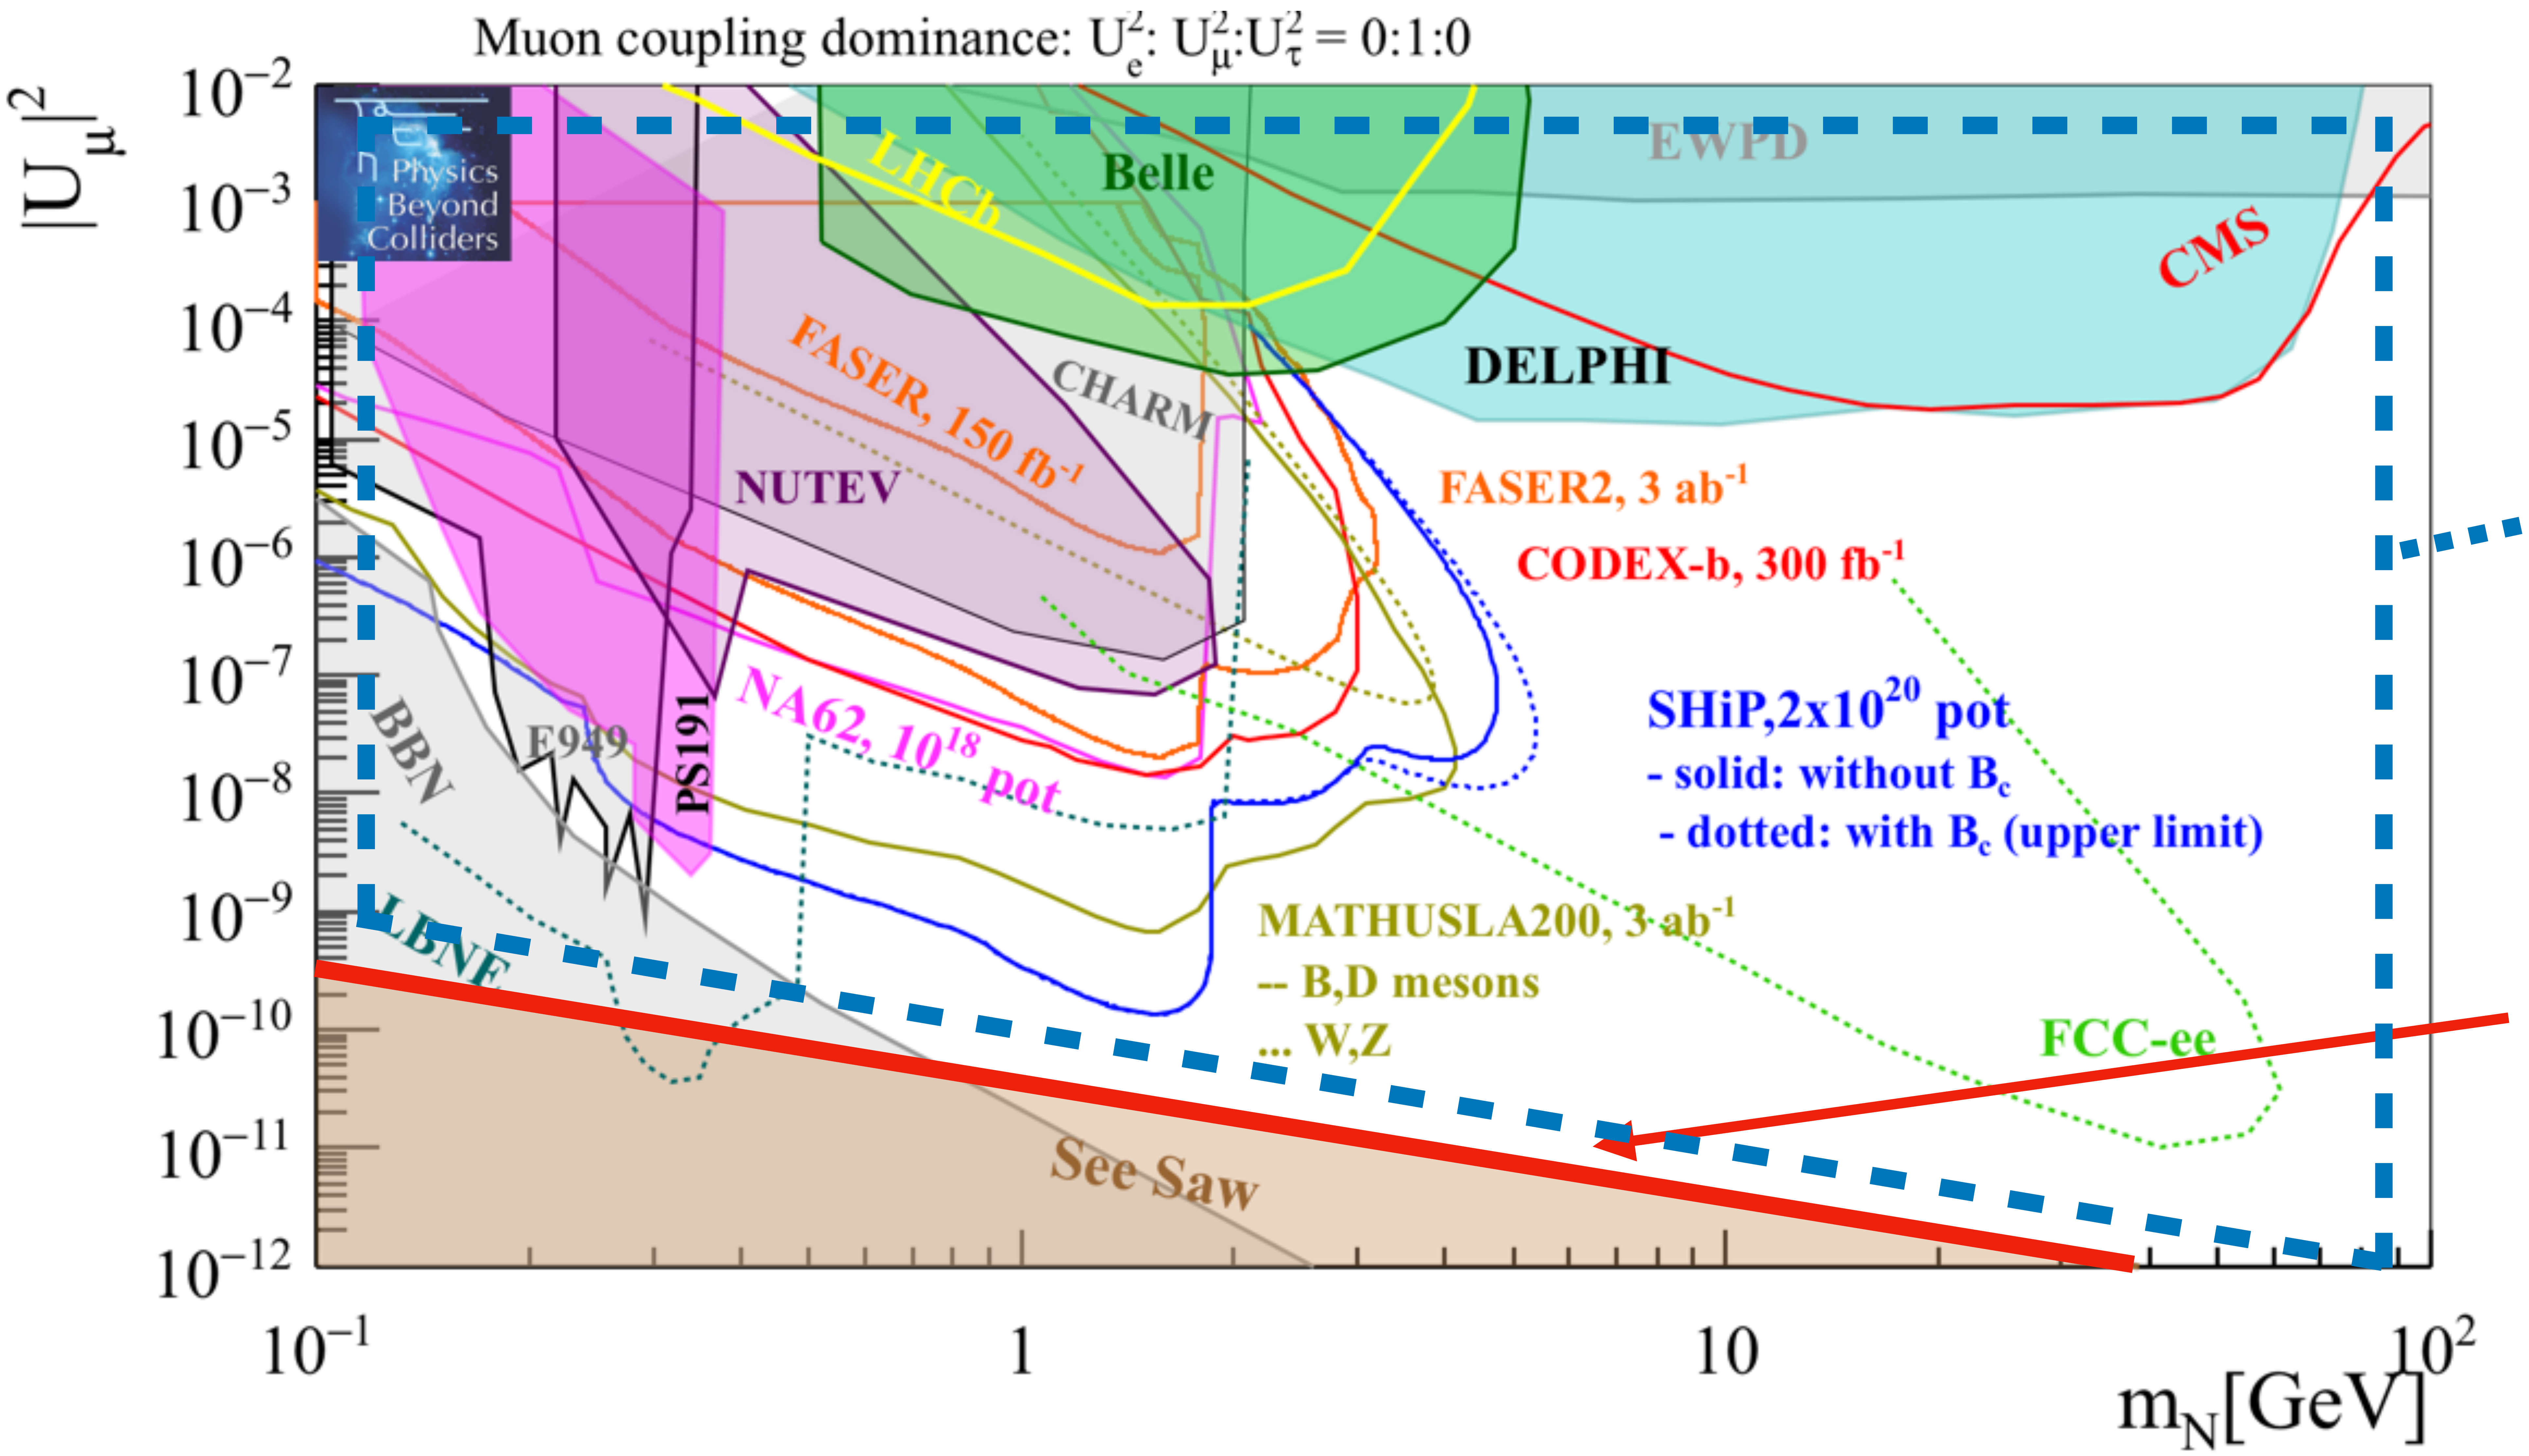

Cancellation required in $m^{\text{light}}_{\alpha\beta} \approx -\sum_I M_I \Theta_{\alpha I} \Theta_{\beta I}$

See-saw line:
"naive" expectation
$m^{\text{light}} \sim m_N \times |\Theta|^2$
Mostly out of reach 😢

Quasi-Dirac HNLs

- The (only?) natural way to implement the cancellation is to arrange HNLs in "quasi-Dirac" pairs, plus any number of "decoupled" HNLs with small Θ.
[Kersten, Smirnov: 0705.3221] [Moffat, Pascoli, Weiland: 1712.07611]

- Each pair corresponds to an approximately conserved lepton number.

- This can be implemented by a symmetry-protected see-saw:
E.g. linear see-saw, inverse see-saw, etc...

- In practice, the two HNLs must be nearly degenerate in mass, with mixing angles are related by a $\pm i$ phase:

$$\Theta_{\alpha 2} \approx \pm i \Theta_{\alpha 1} \Longrightarrow m_{\alpha\beta}^{\text{light}} \cong - \sum_{I=1,2} M_I \Theta_{\alpha I} \Theta_{\beta I} \lll M |\Theta|^2$$

Lepton number conservation & violation

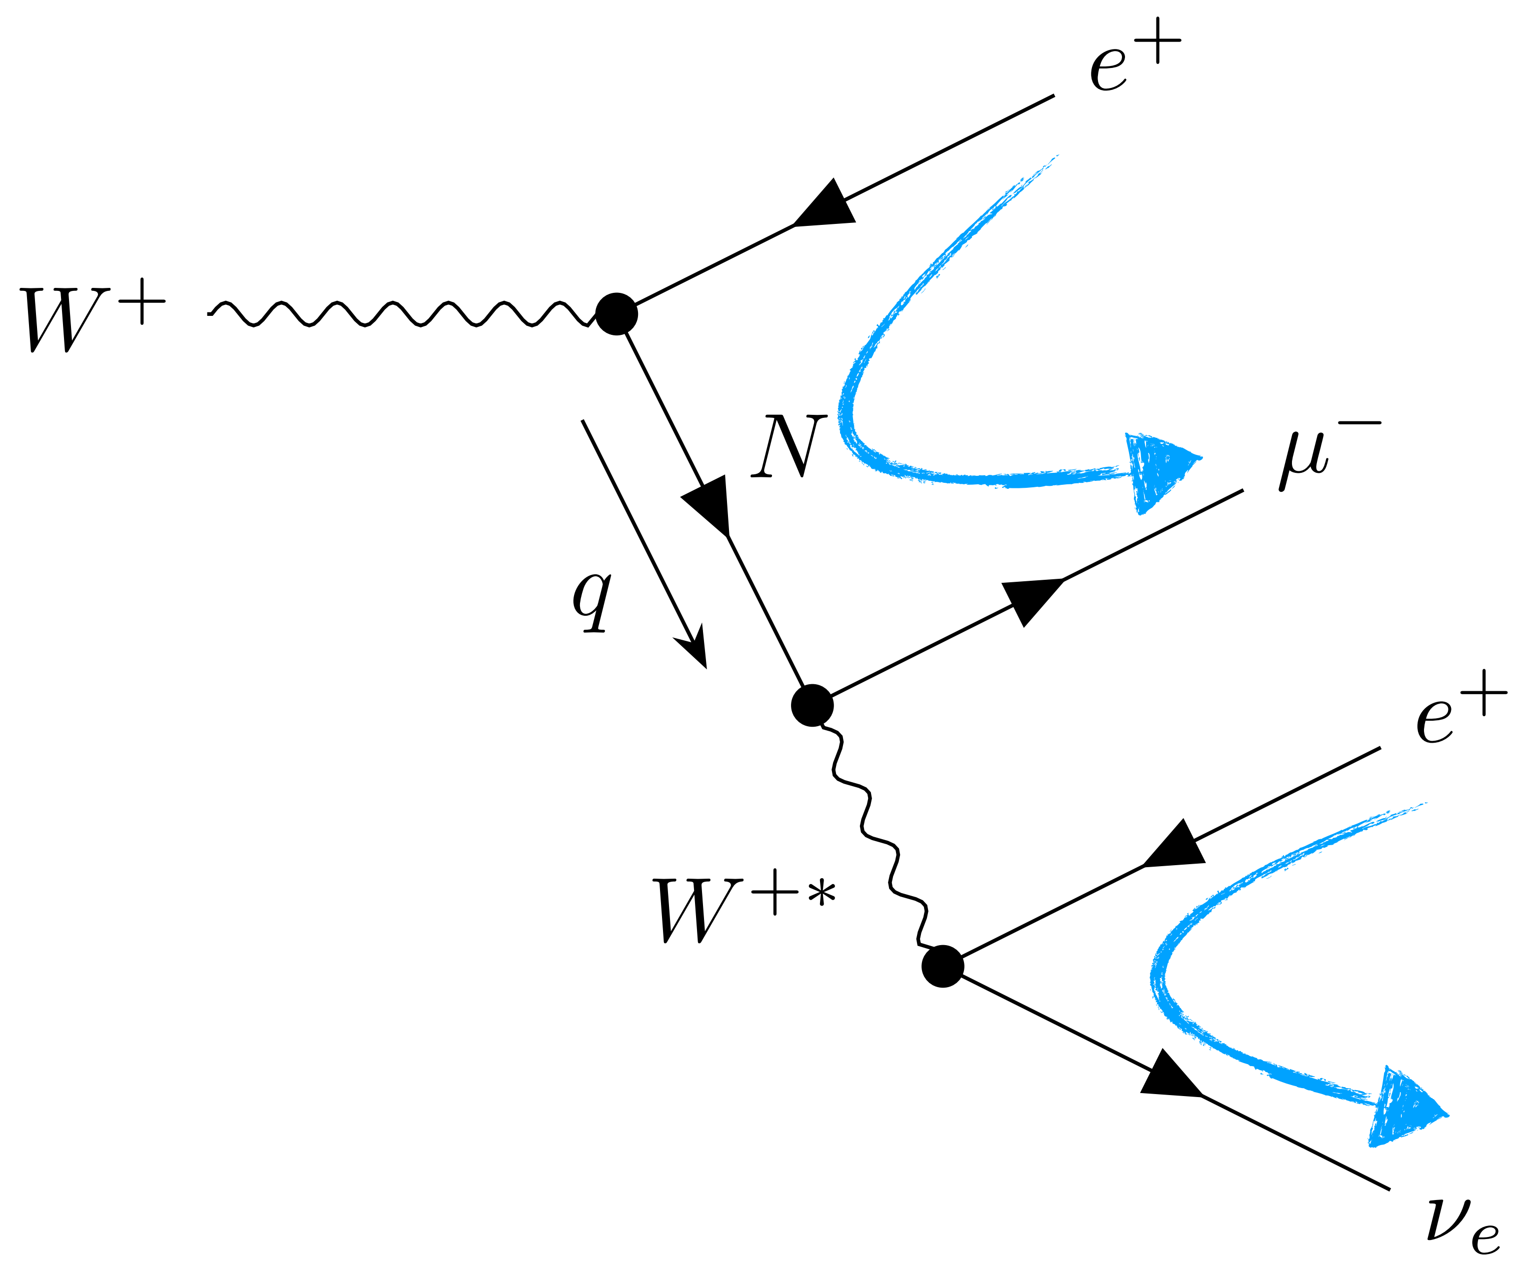

Lepton number conserving
(LNC)

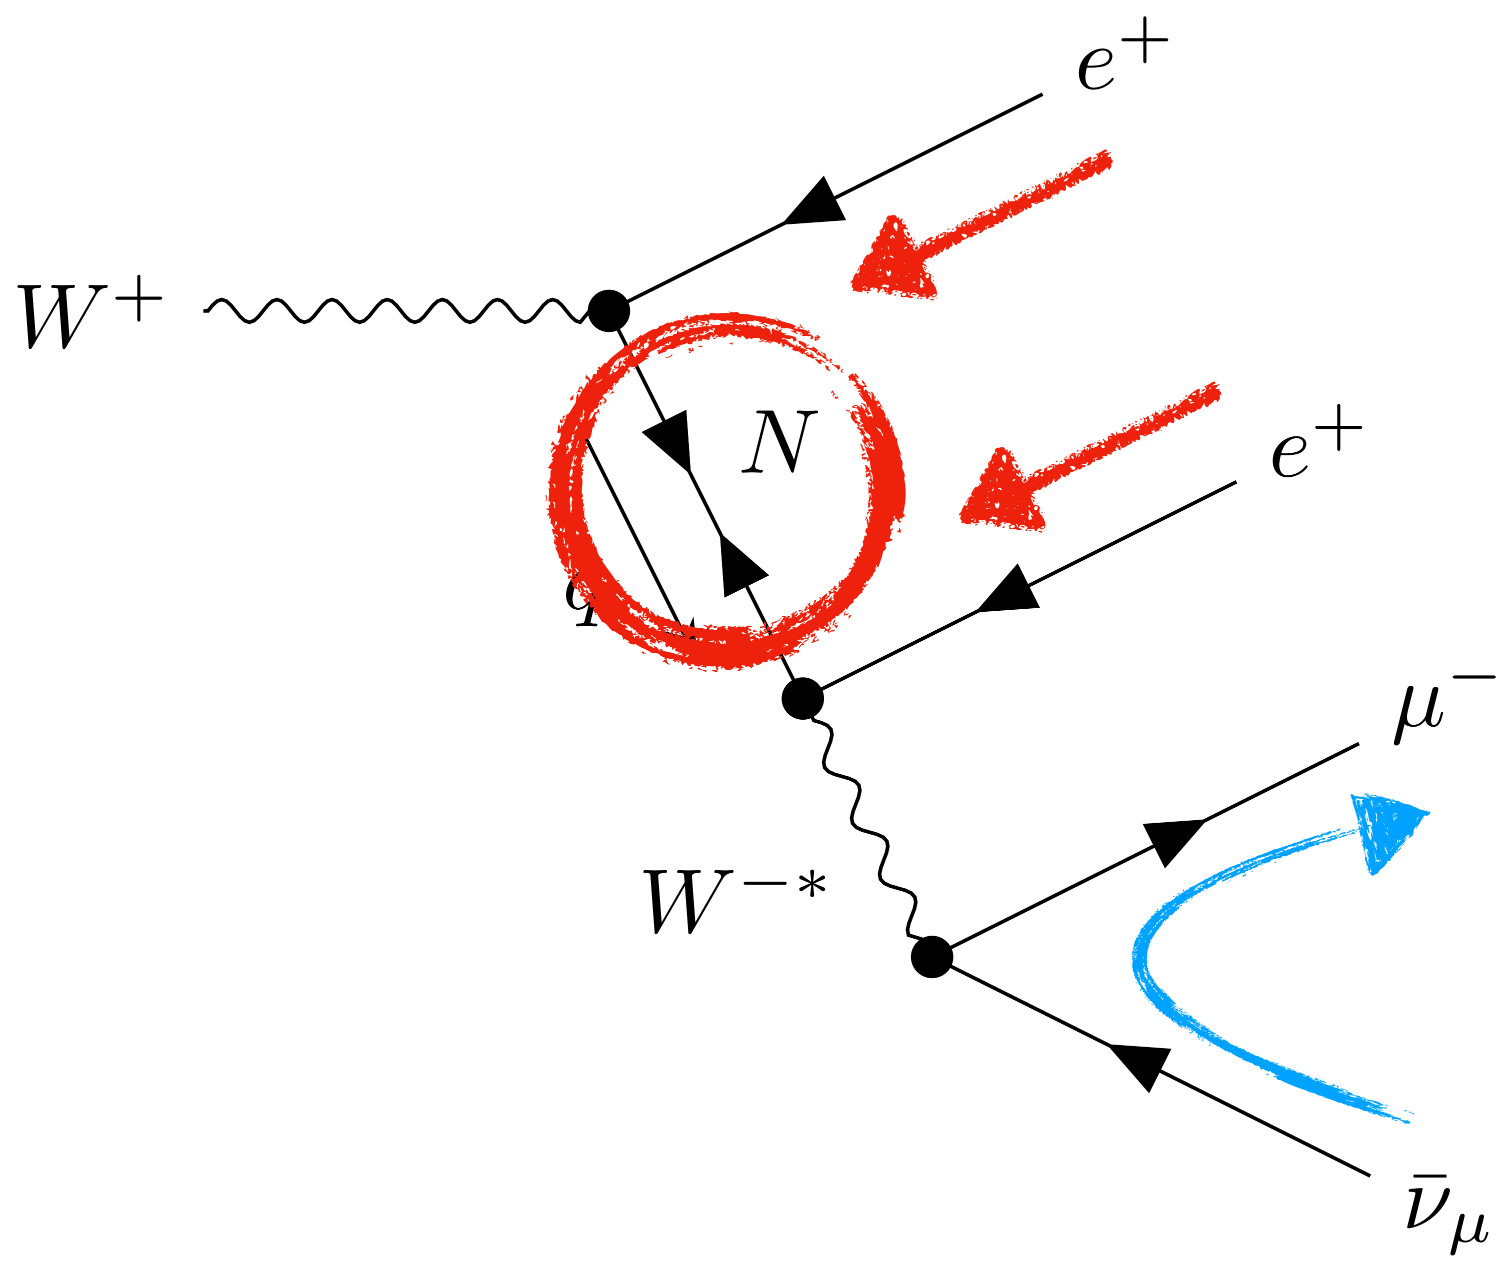

Lepton number violating
(LNV)

Quasi-Dirac HNLs & lepton number violation

- Quasi-Dirac limit: $M_2 = M_1, \Theta_{\alpha 2} = \pm i\Theta_{\alpha 1}$

$$\Longrightarrow m_{\alpha\beta}^{\text{light}} = -\sum_I M_I \Theta_{\alpha I}\Theta_{\beta I} = 0$$

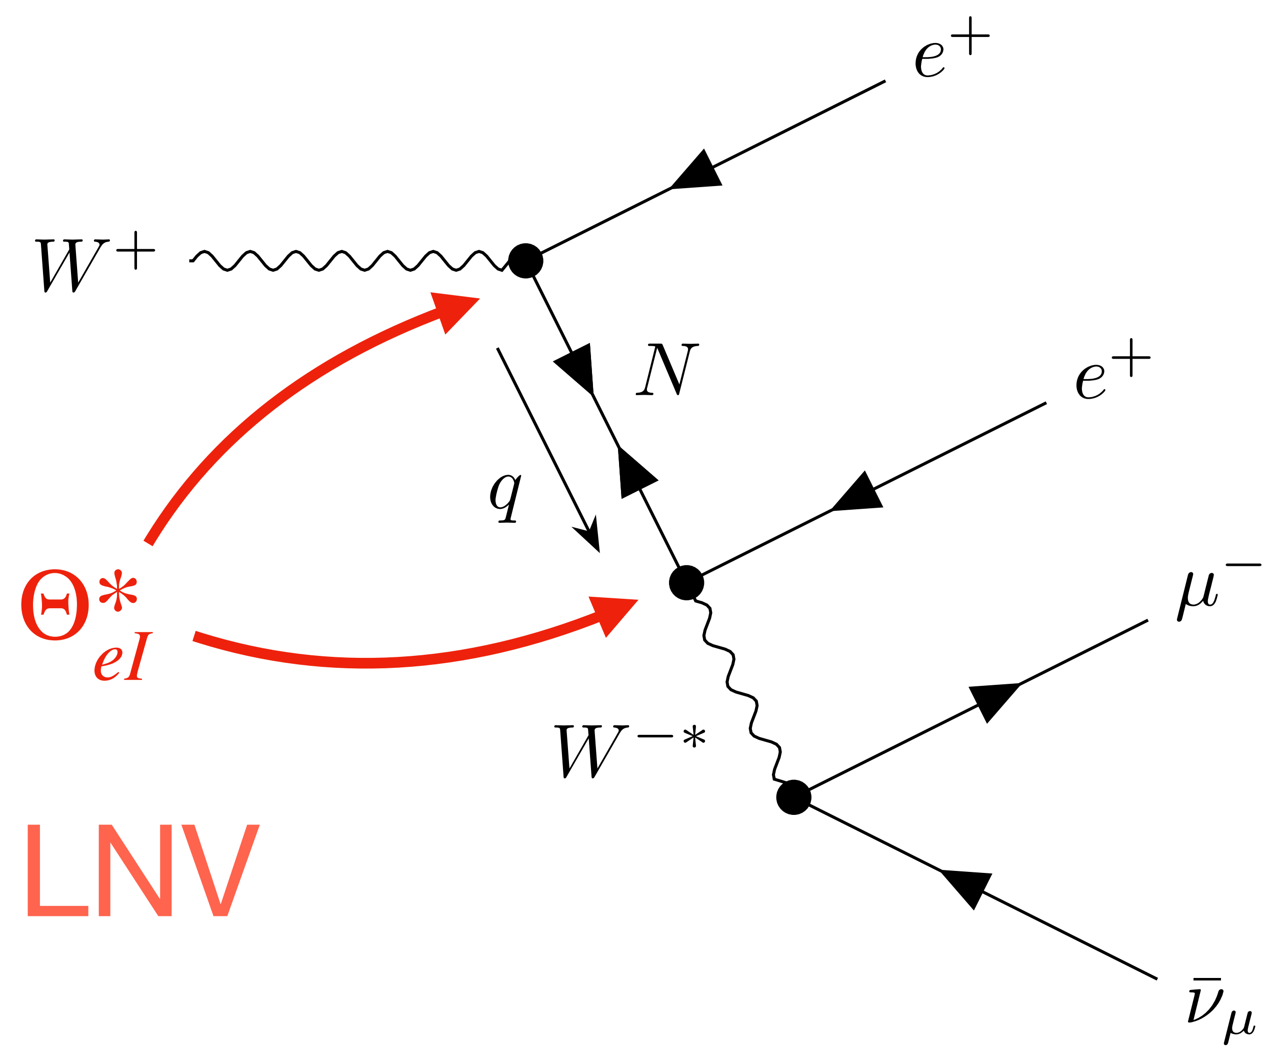

- **Problem:** amplitude of LNV processes $\propto \sum_I \Theta_{\alpha I}\Theta_{\beta I}$ (or complex conjugate)
- In the *exact* quasi-Dirac limit, amplitudes of LNV processes vanish.
- Is that still the case if the lepton number symmetry is only *slightly broken*?

Coherent HNL oscillations

- If the mass splitting $\delta M = M_2 - M_1$ is small enough, the two HNL mass eigenstates still contribute coherently to the same processes.

- However, they acquire an increasing phase difference as they propagate, which can be approximated as $e^{-i\delta M\tau}$ with $\tau = |x_{\text{prod}} - x_{\text{decay}}| > 0$ the proper time between the HNL production and decay. (A rigorous derivation requires a QFT treatment involving wave-packets)
[Beuthe: hep-ph/0109119], [Tastet: master thesis], [Antusch, Rosskopp: 2012.05763]

- This phase difference leads to HNL oscillations.
They share many similarities with heavy meson oscillations (e.g. $K^0 \leftrightarrow \overline{K}^0$).

- *Quantum-mechanical* coherence can be proven for a very large number of oscillations (in practice, it will be smoothed out by the finite energy resolution)

- Experimentally, there are **three** regimes of interest, depending on how δM^{-1} compares with the typical proper time scale τ probed by the experiment.

Coherent HNL oscillations: prompt decay

- If the HNL decays promptly, the relevant proper time scale is given by its expected lifetime (i.e.: does the HNL have time to oscillate before decaying?)

$\delta M \ll \Gamma$: Dirac-like

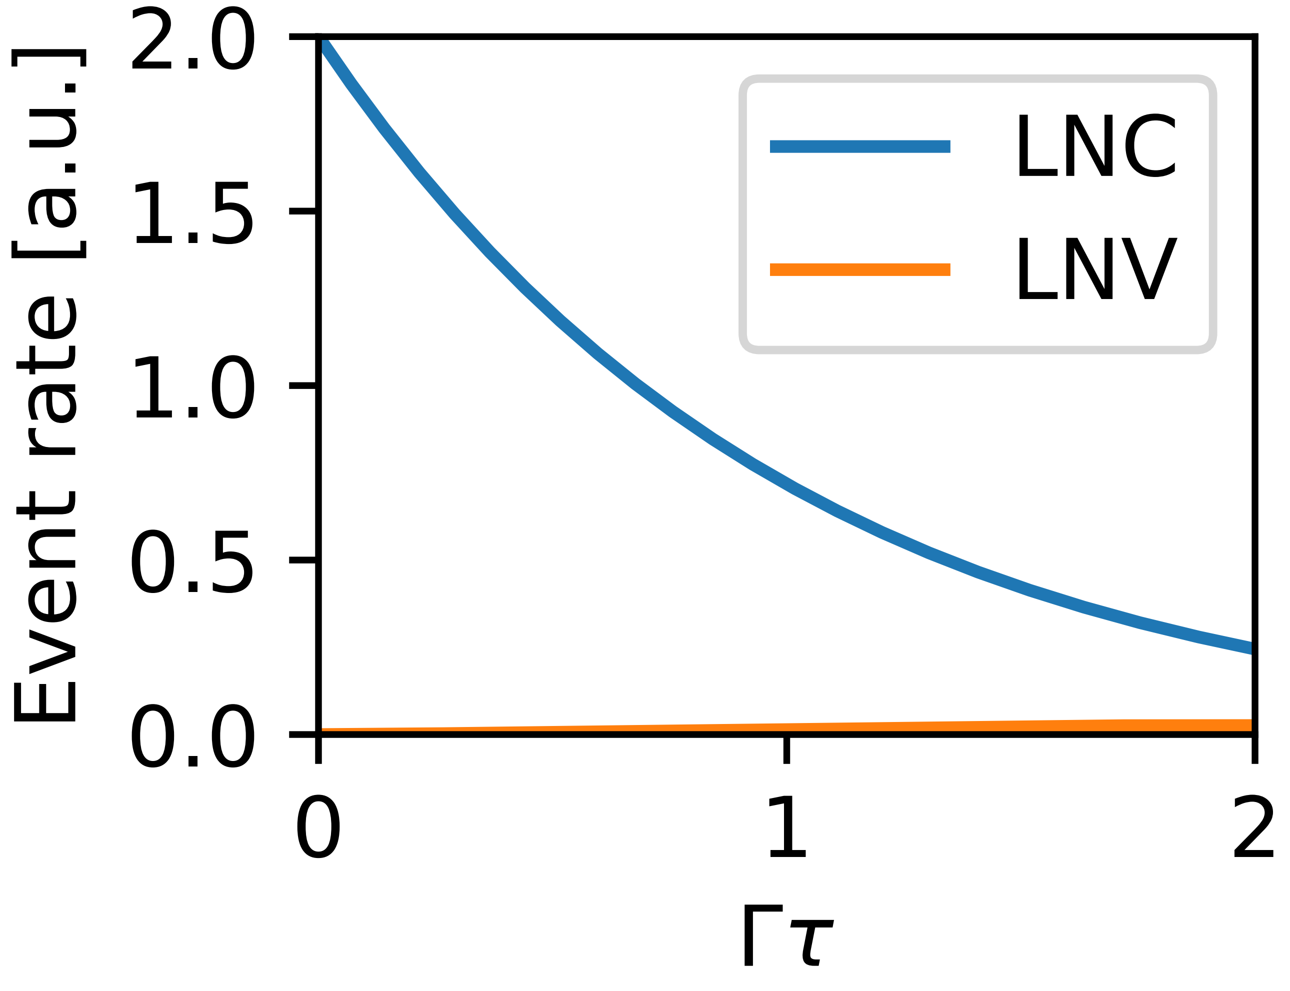

$\delta M \sim \pi\Gamma$: resolvable osc.

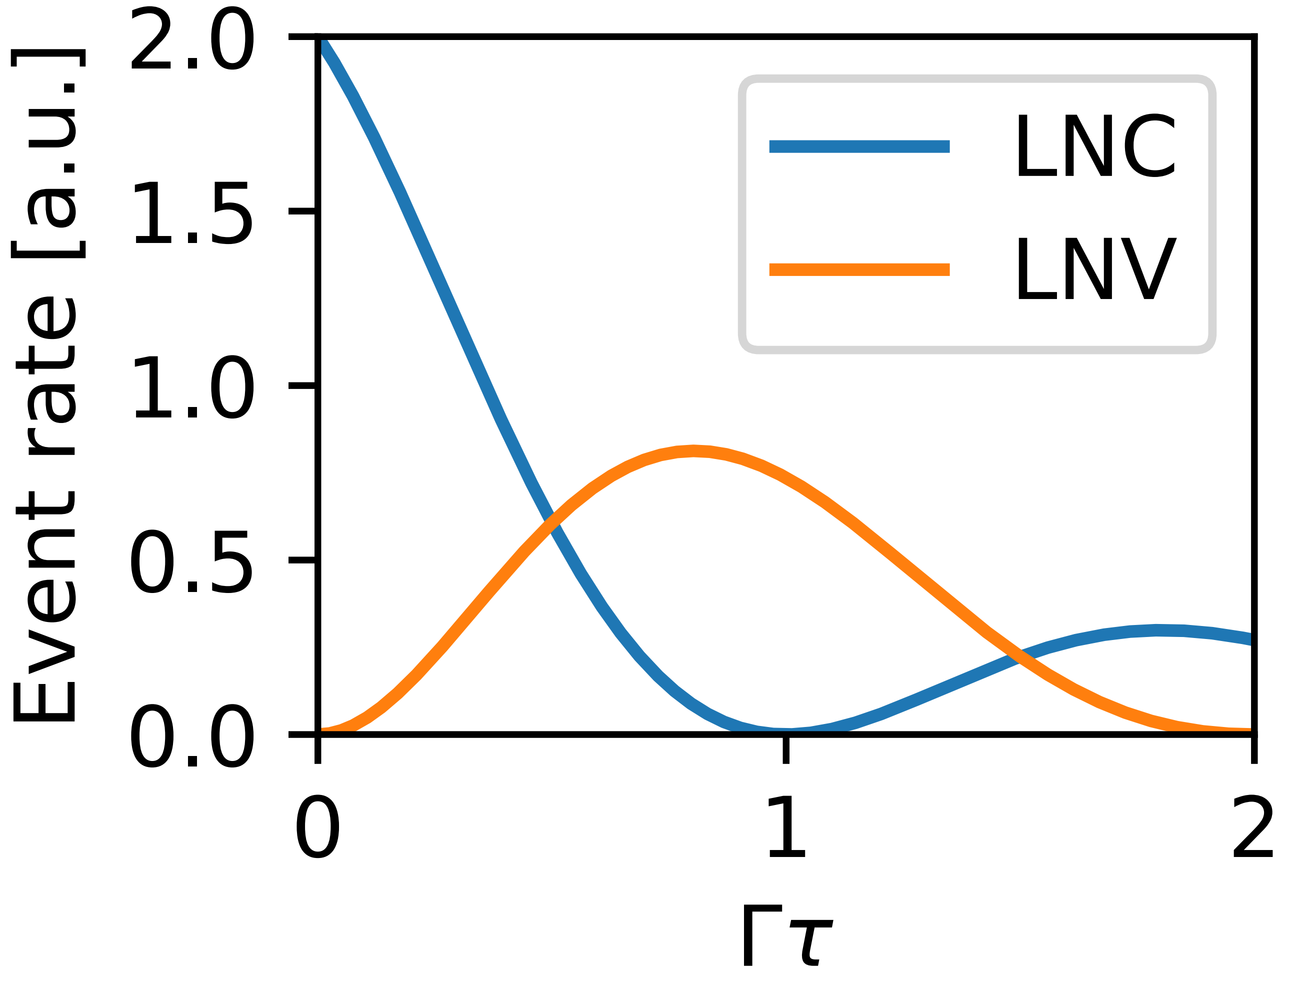

$\delta M \gg \Gamma$: Majorana-like

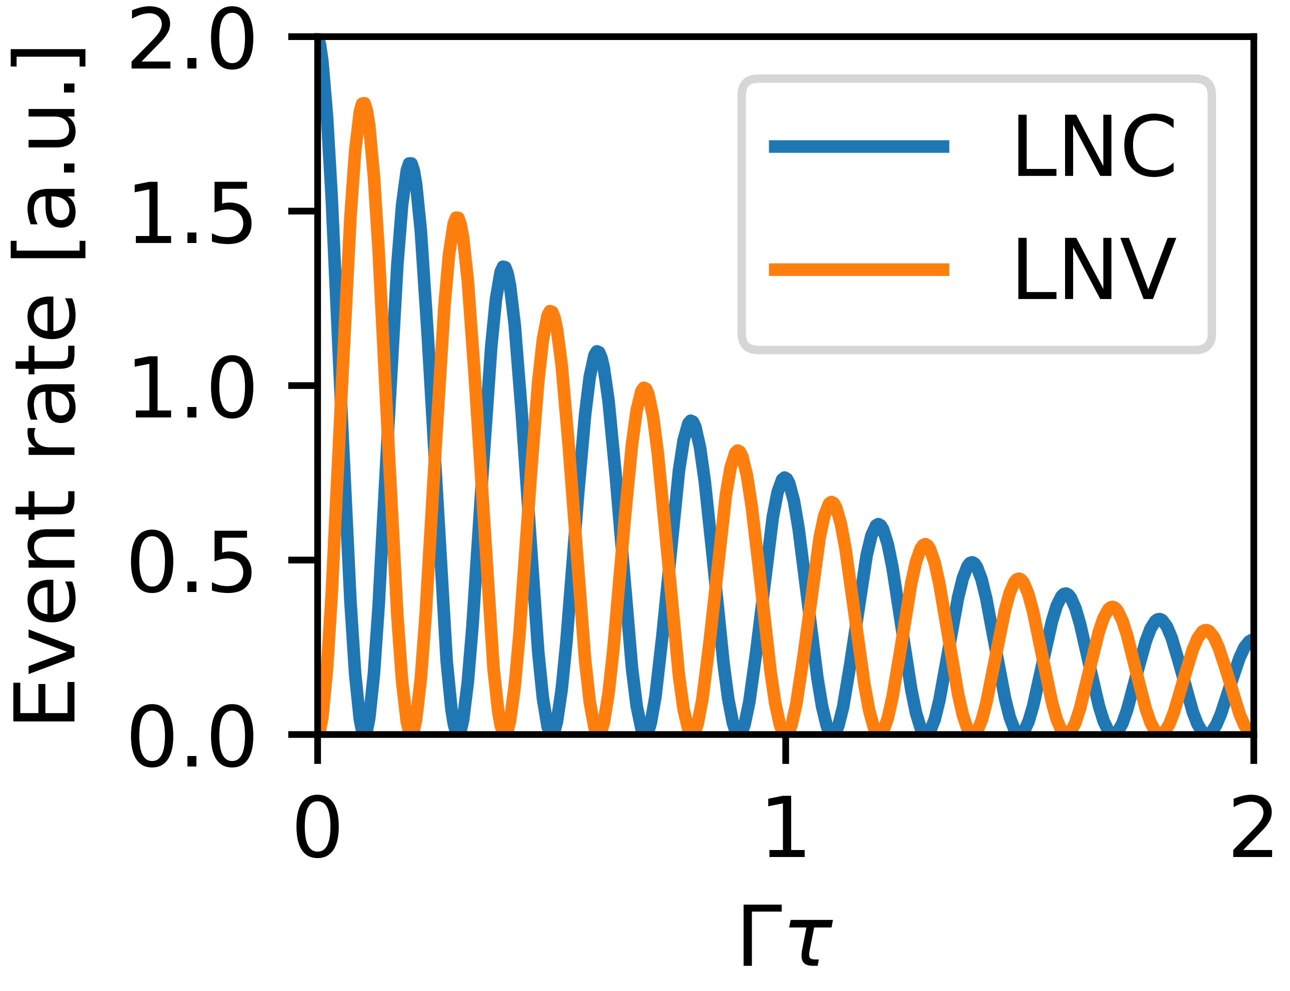

Coherent HNL oscillations: displaced decay

- If the HNL lives long enough to escape the detector, then the relevant length scale is the detector size *in the HNL frame*.
- E.g.: typical cut for a prompt search at ATLAS is ~1 mm *in the lab frame*.
- In the example, the HNL looks like a Dirac particle on a 1 mm scale, but like a Majorana particle if we use large-radius tracking over ~meter.

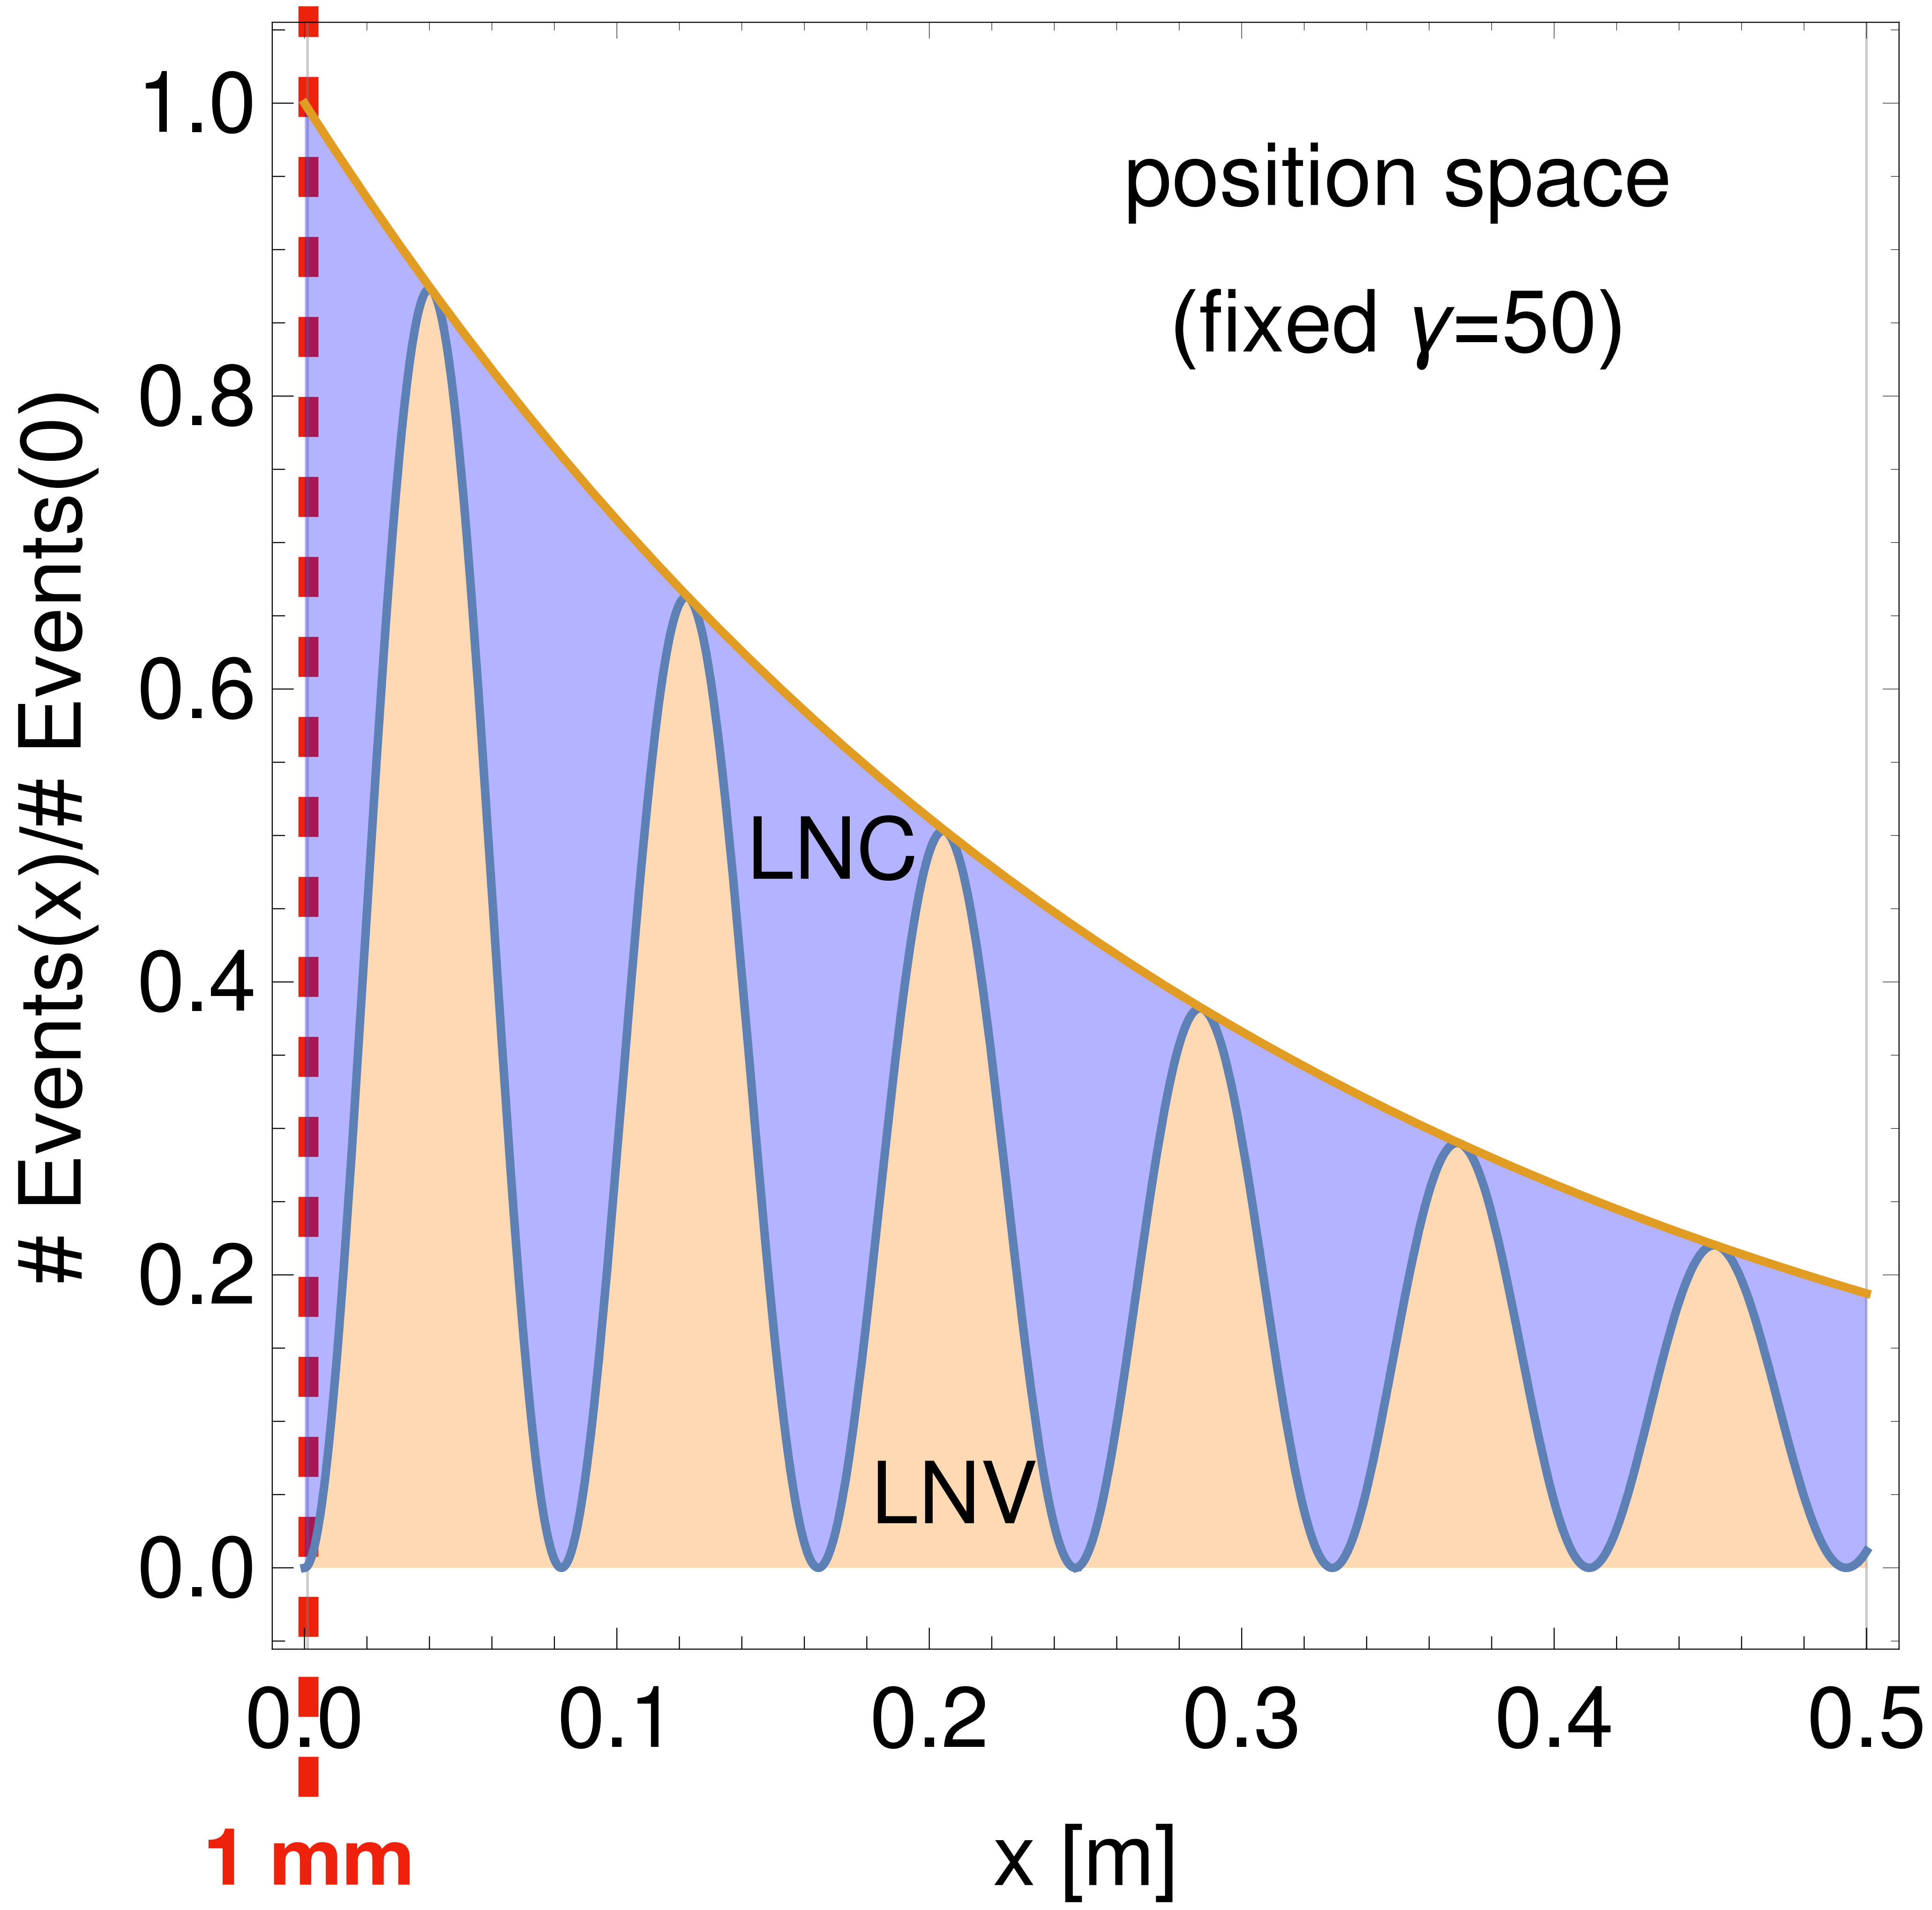

Plot from [Antusch, Cazzato, Fischer: 1709.03797]

Constraints from neutrino oscillation data

- So far we only considered constraints coming from the magnitude of m^{light}.
- However, the observed neutrino mass splittings and mixing angles also put strong (model-dependent) constraints on the HNL parameters.
- The Casas-Ibarra parametrisation allows one to obtain all the allowed Yukawa couplings Y^{ν} for a given set of HNL and neutrino parameters:

  $\Theta_{\alpha I} = i V^{\text{PMNS}}_{\alpha i} \sqrt{m_i} \Omega_{iI} \sqrt{M_I}^{-1}$ with Ω a complex orthogonal matrix (+ zeros)

  [Casas, Ibarra: hep-ph/0103065]
- The parametrisation itself is generic, and can easily produce "unnatural" sets of HNL parameters. Symmetry-protected seesaw models are more predictive.

Constraints from neutrino data for 2 HNLs

- Consider two quasi-Dirac HNLs
  (plus optionally one decoupled HNL)
  E.g. lepton numbers +1, -1 and 0.

- The lightest neutrino must be massless
  (up to loop corrections to its mass).

- **Vary** the light neutrino parameters within their uncertainties (as per NuFIT).

- **Scan** over all the free HNL parameters.

- Well above the see-saw line, the constraints do *not* depend on the HNL mass nor the total mixing angle.
  $\rightarrow$ Can be drawn on a *ternary plot*!

Constraints from neutrino data for 2 HNLs

- Consider two quasi-Dirac HNLs (plus optionally one decoupled HNL) E.g. lepton numbers +1, -1 and 0.
- The lightest neutrino must be massless (up to loop corrections to its mass).
- **Vary** the light neutrino parameters within their uncertainties (as per NuFIT).
- **Scan** over all the free HNL parameters.
- Well above the see-saw line, the constraints do *not* depend on the HNL mass nor the total mixing angle. → Can be drawn on a *ternary plot*!

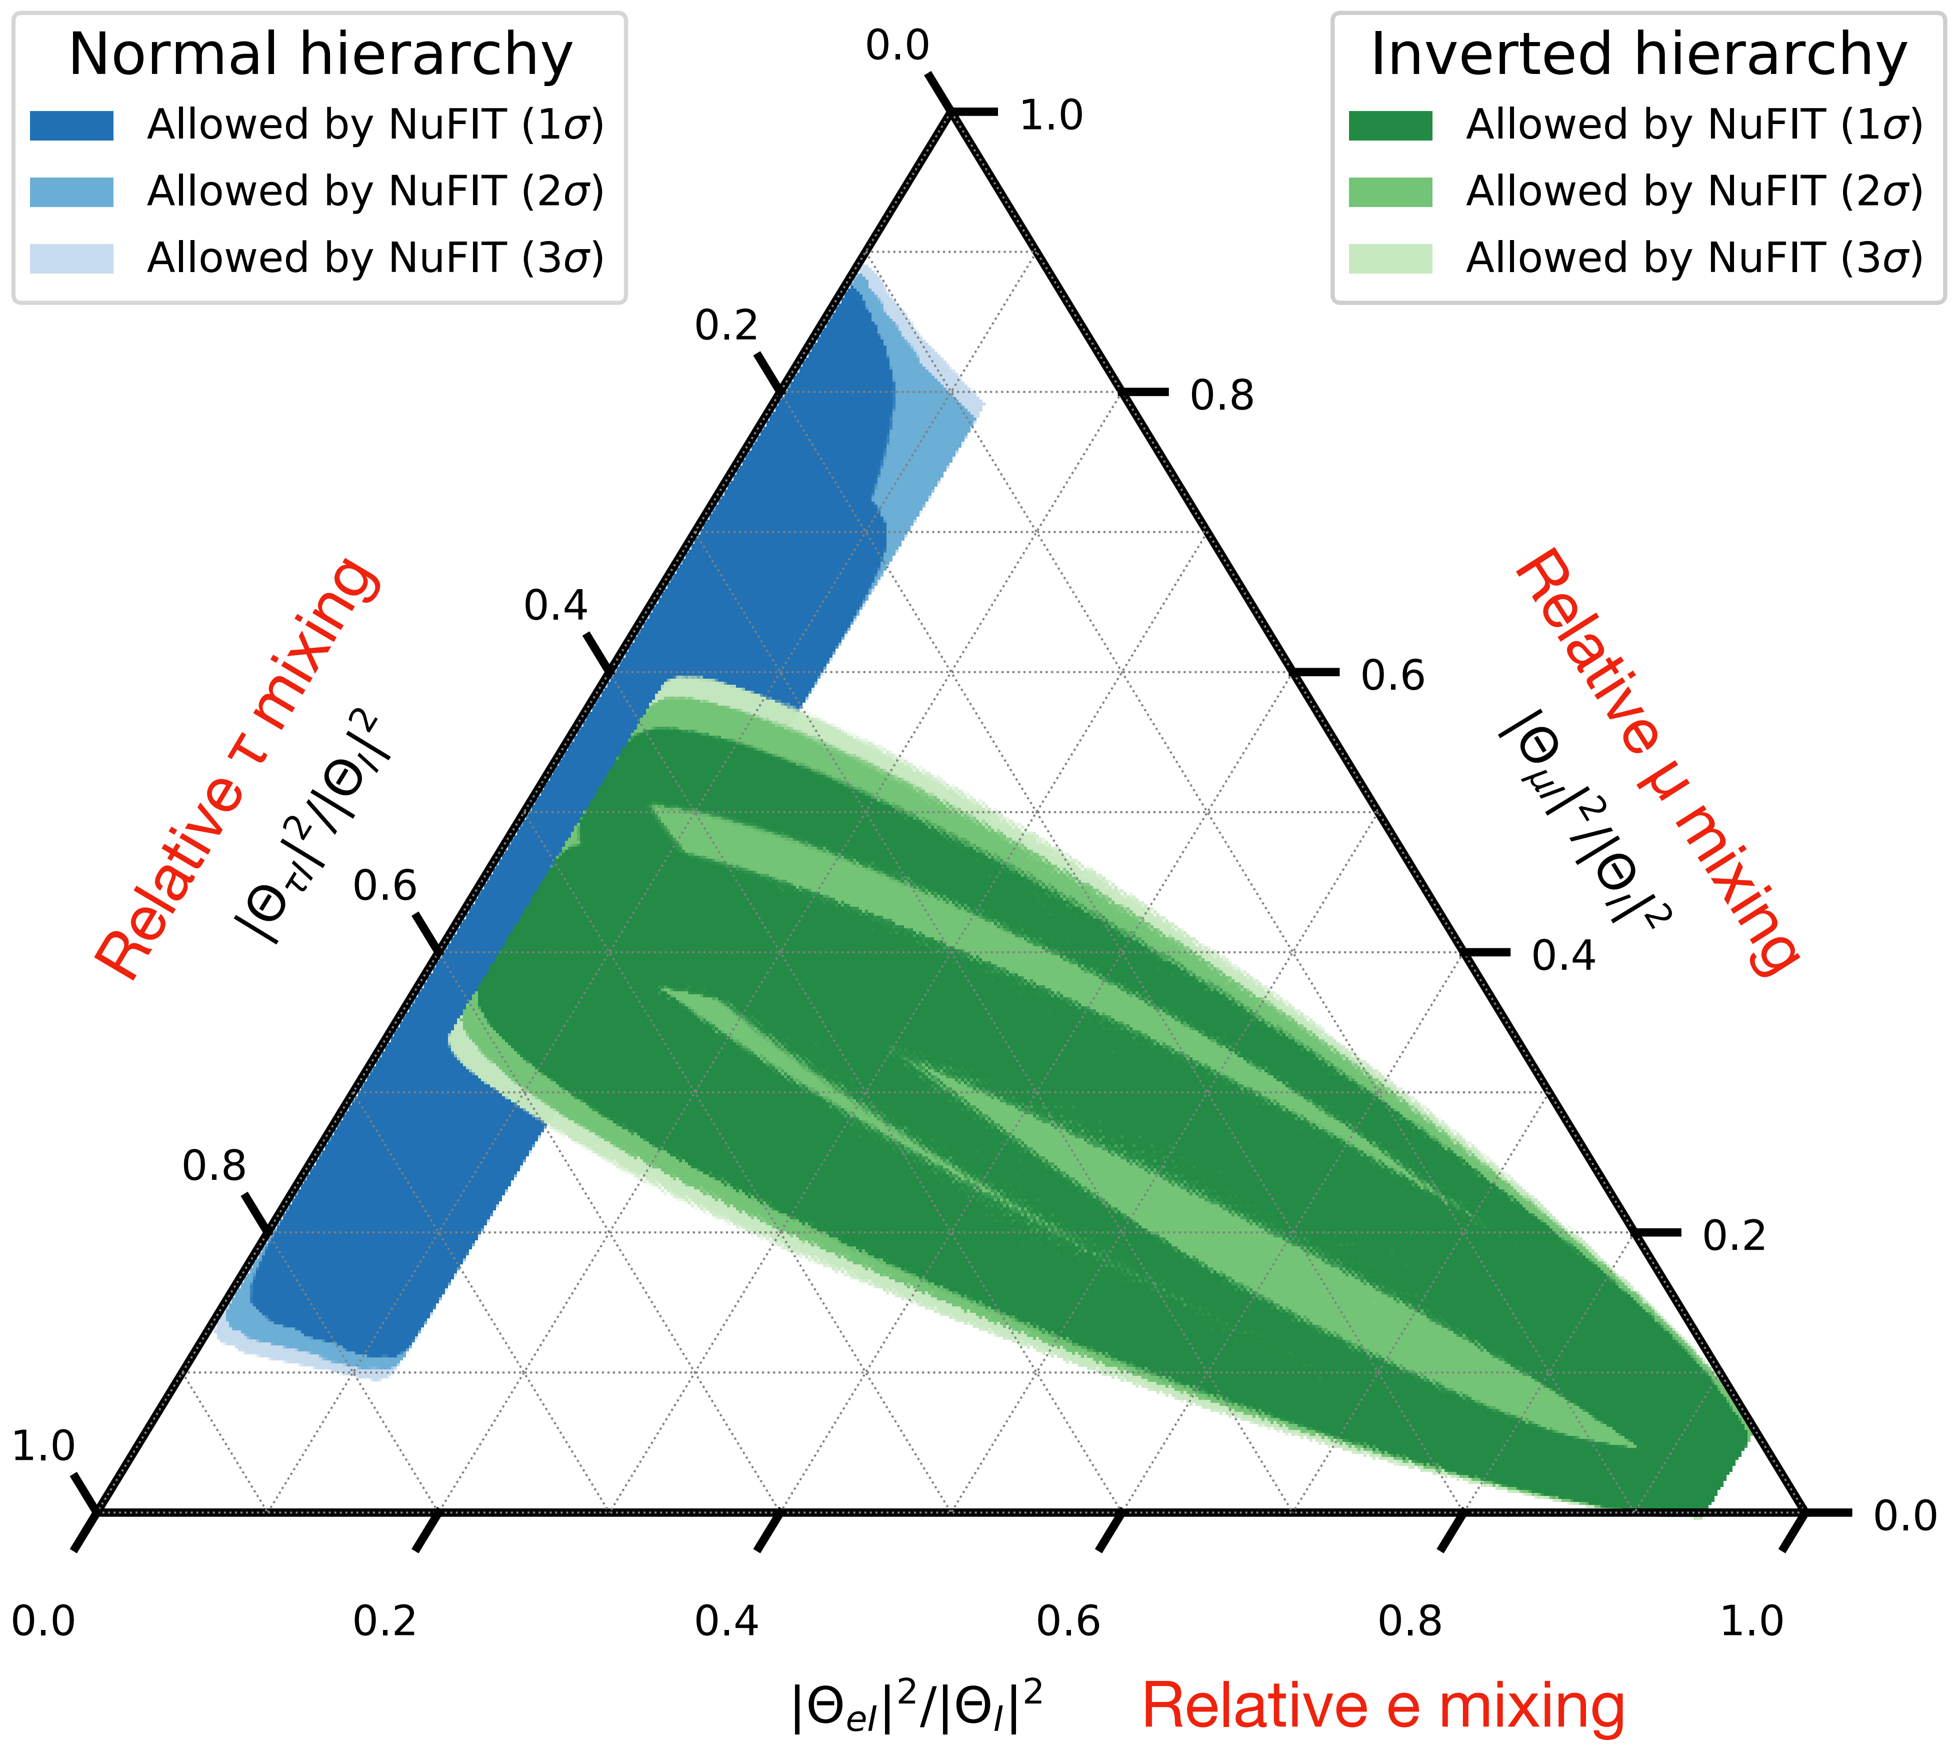

[Drewes, Garbrecht, Gueter, Klarić: 1609.09069]
[Caputo, Hernández, López-Pavón, Salvado: 1704.08721]

The SHiP experiment

Reminder

HNLs $\nu_{R,I}$ behave as:

- **Heavy neutrinos (Majorana or pseudo-Dirac)**
- **With the same interactions as light neutrinos $\nu_{L,\alpha}$**
- **But suppressed by a small mixing parameter $\Theta_{\alpha I}$**

Prototypical example of a feebly interacting particle (FIP)

Experimental requirements

Experimental requirements

Small mixing angles lead to:

- Suppressed production rate

Experimental requirements

Small mixing angles lead to:

- Suppressed production rate

Solution

- High intensity / luminosity
- Low background

Experimental requirements

Small mixing angles lead to:

- Suppressed production rate
- A possibly long lifetime:
  the particle may travel a long
  distance before decaying

Solution

- High intensity / luminosity
- Low background

Experimental requirements

Small mixing angles lead to:

- Suppressed production rate
- A possibly long lifetime: the particle may travel a long distance before decaying

Solution

- High intensity / luminosity
- Low background
- Displaced detector
- Large detector volume

The SHiP experiment (Search for Hidden Particles)

SHiP
Search for Hidden Particles

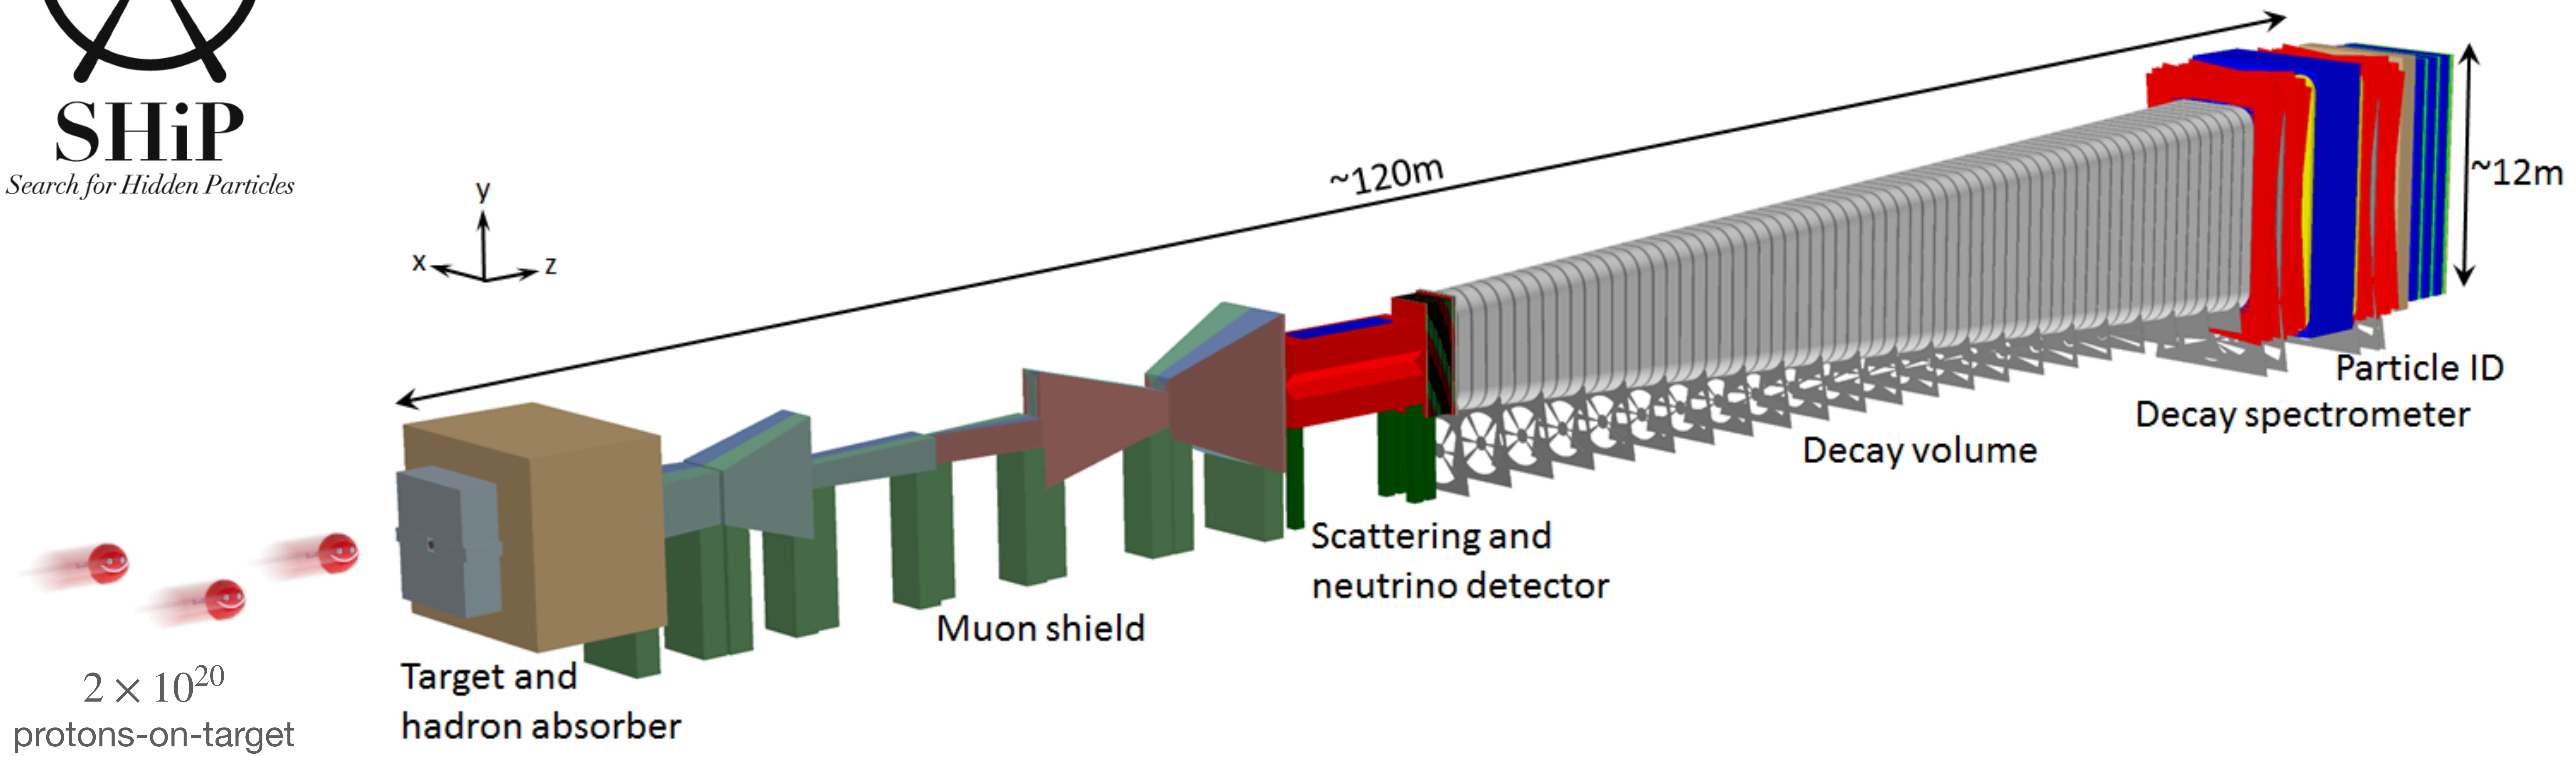

2×10^{20}
protons-on-target
/ 5 years
@ 400 GeV

[SHiP: 1504.04956, 1504.04855, 2112.01487]

The SHiP experiment (Search for Hidden Particles)

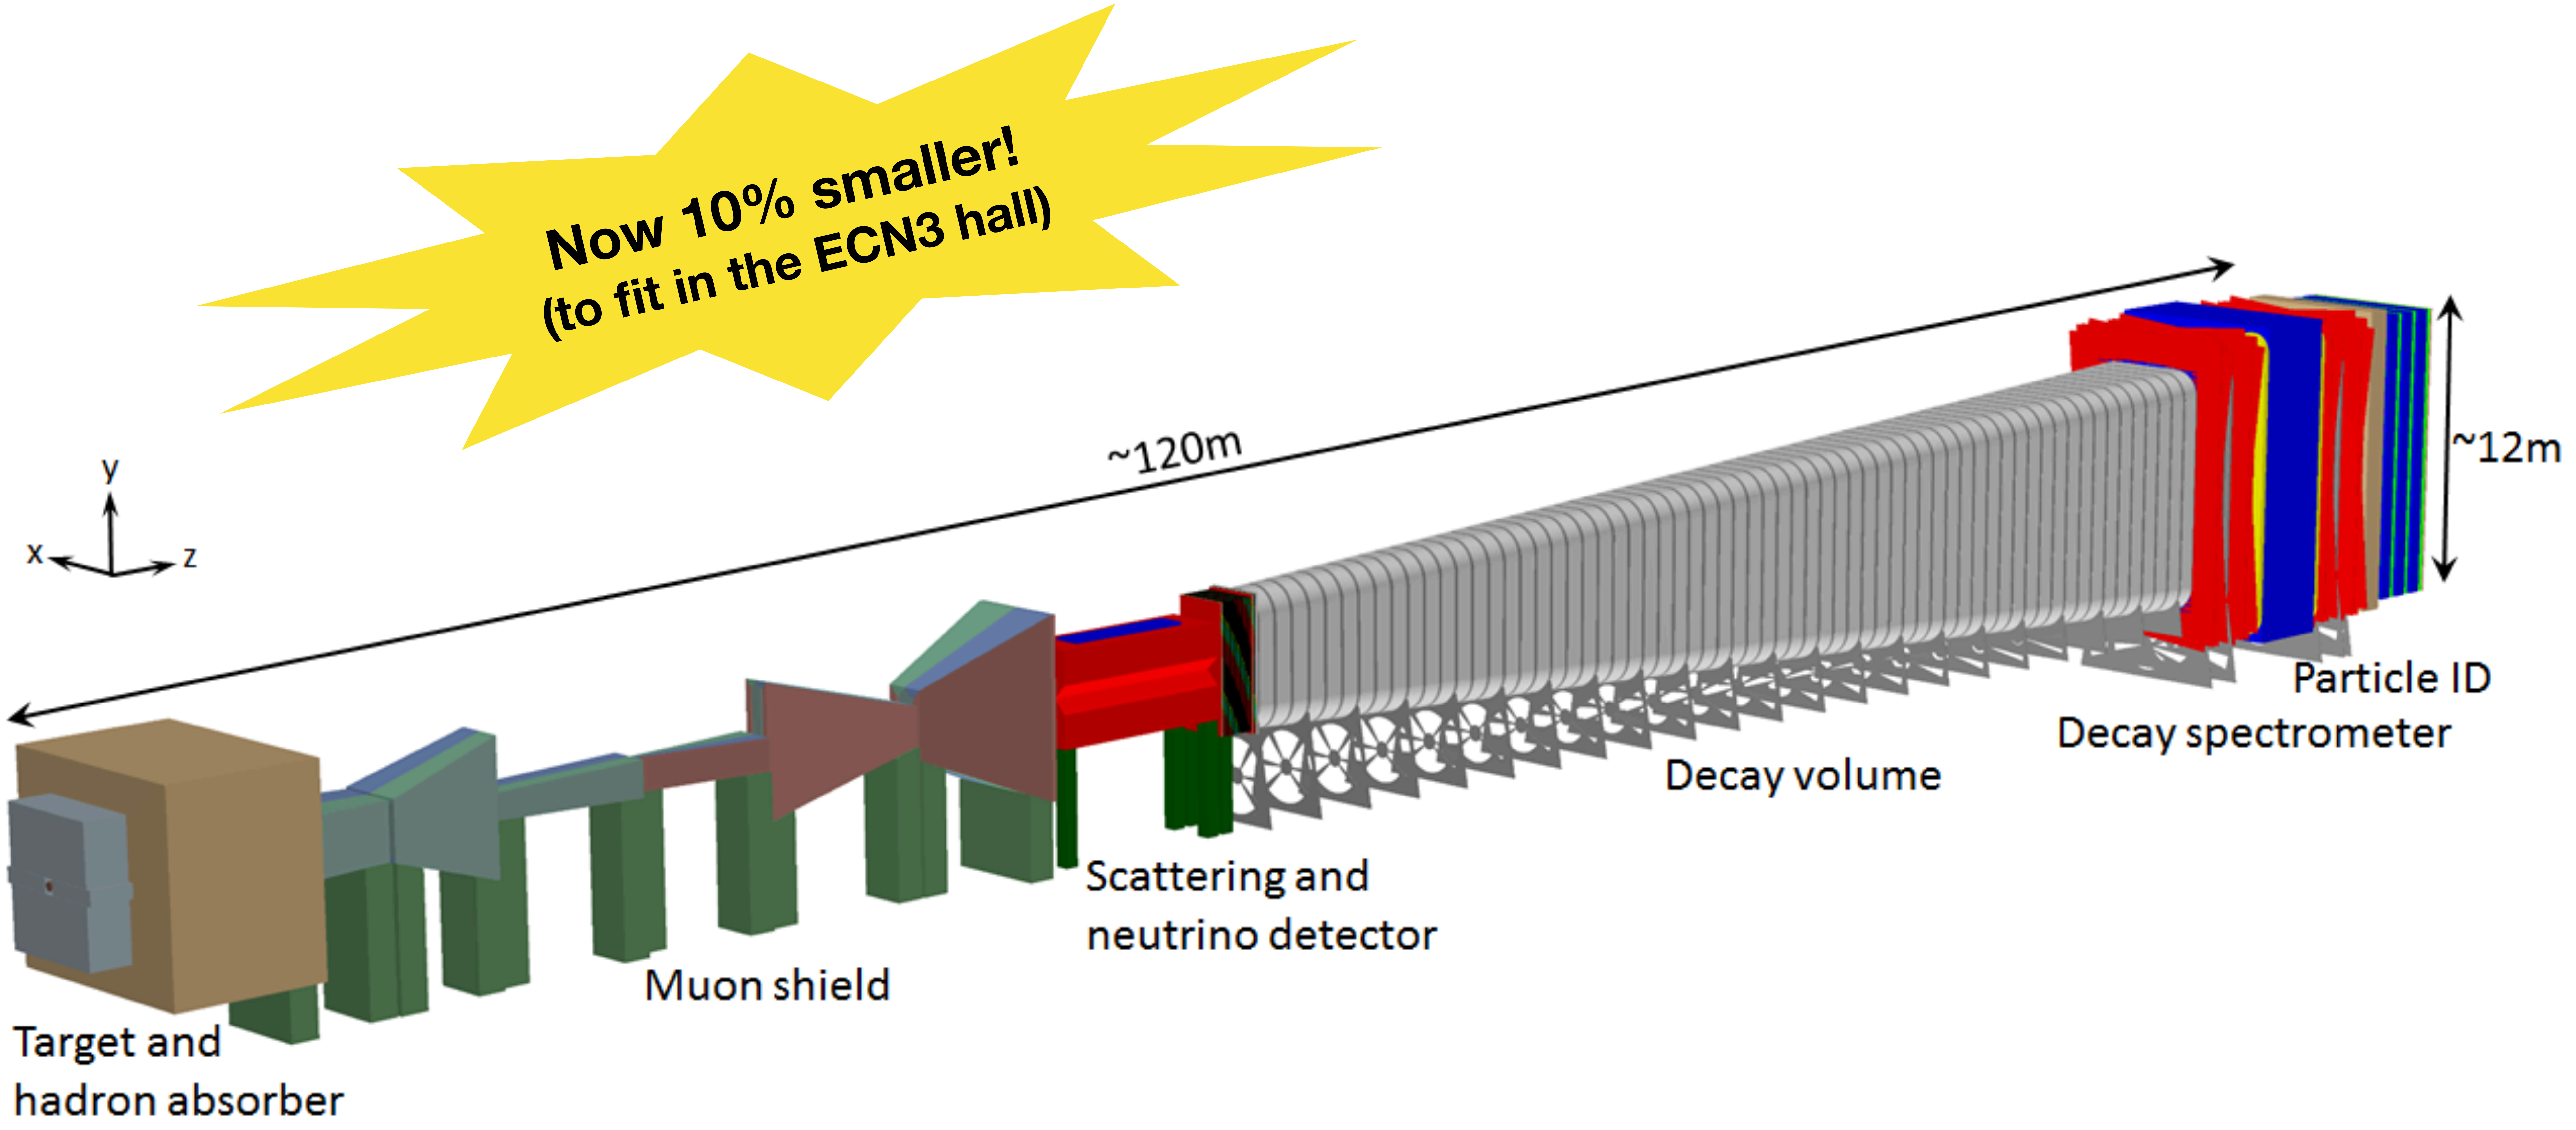

2×10^{20}
protons-on-target
/ 5 years
@ 400 GeV

[SHiP: 1504.04956, 1504.04855, 2112.01487]

The SHiP experiment (Search for Hidden Particles)

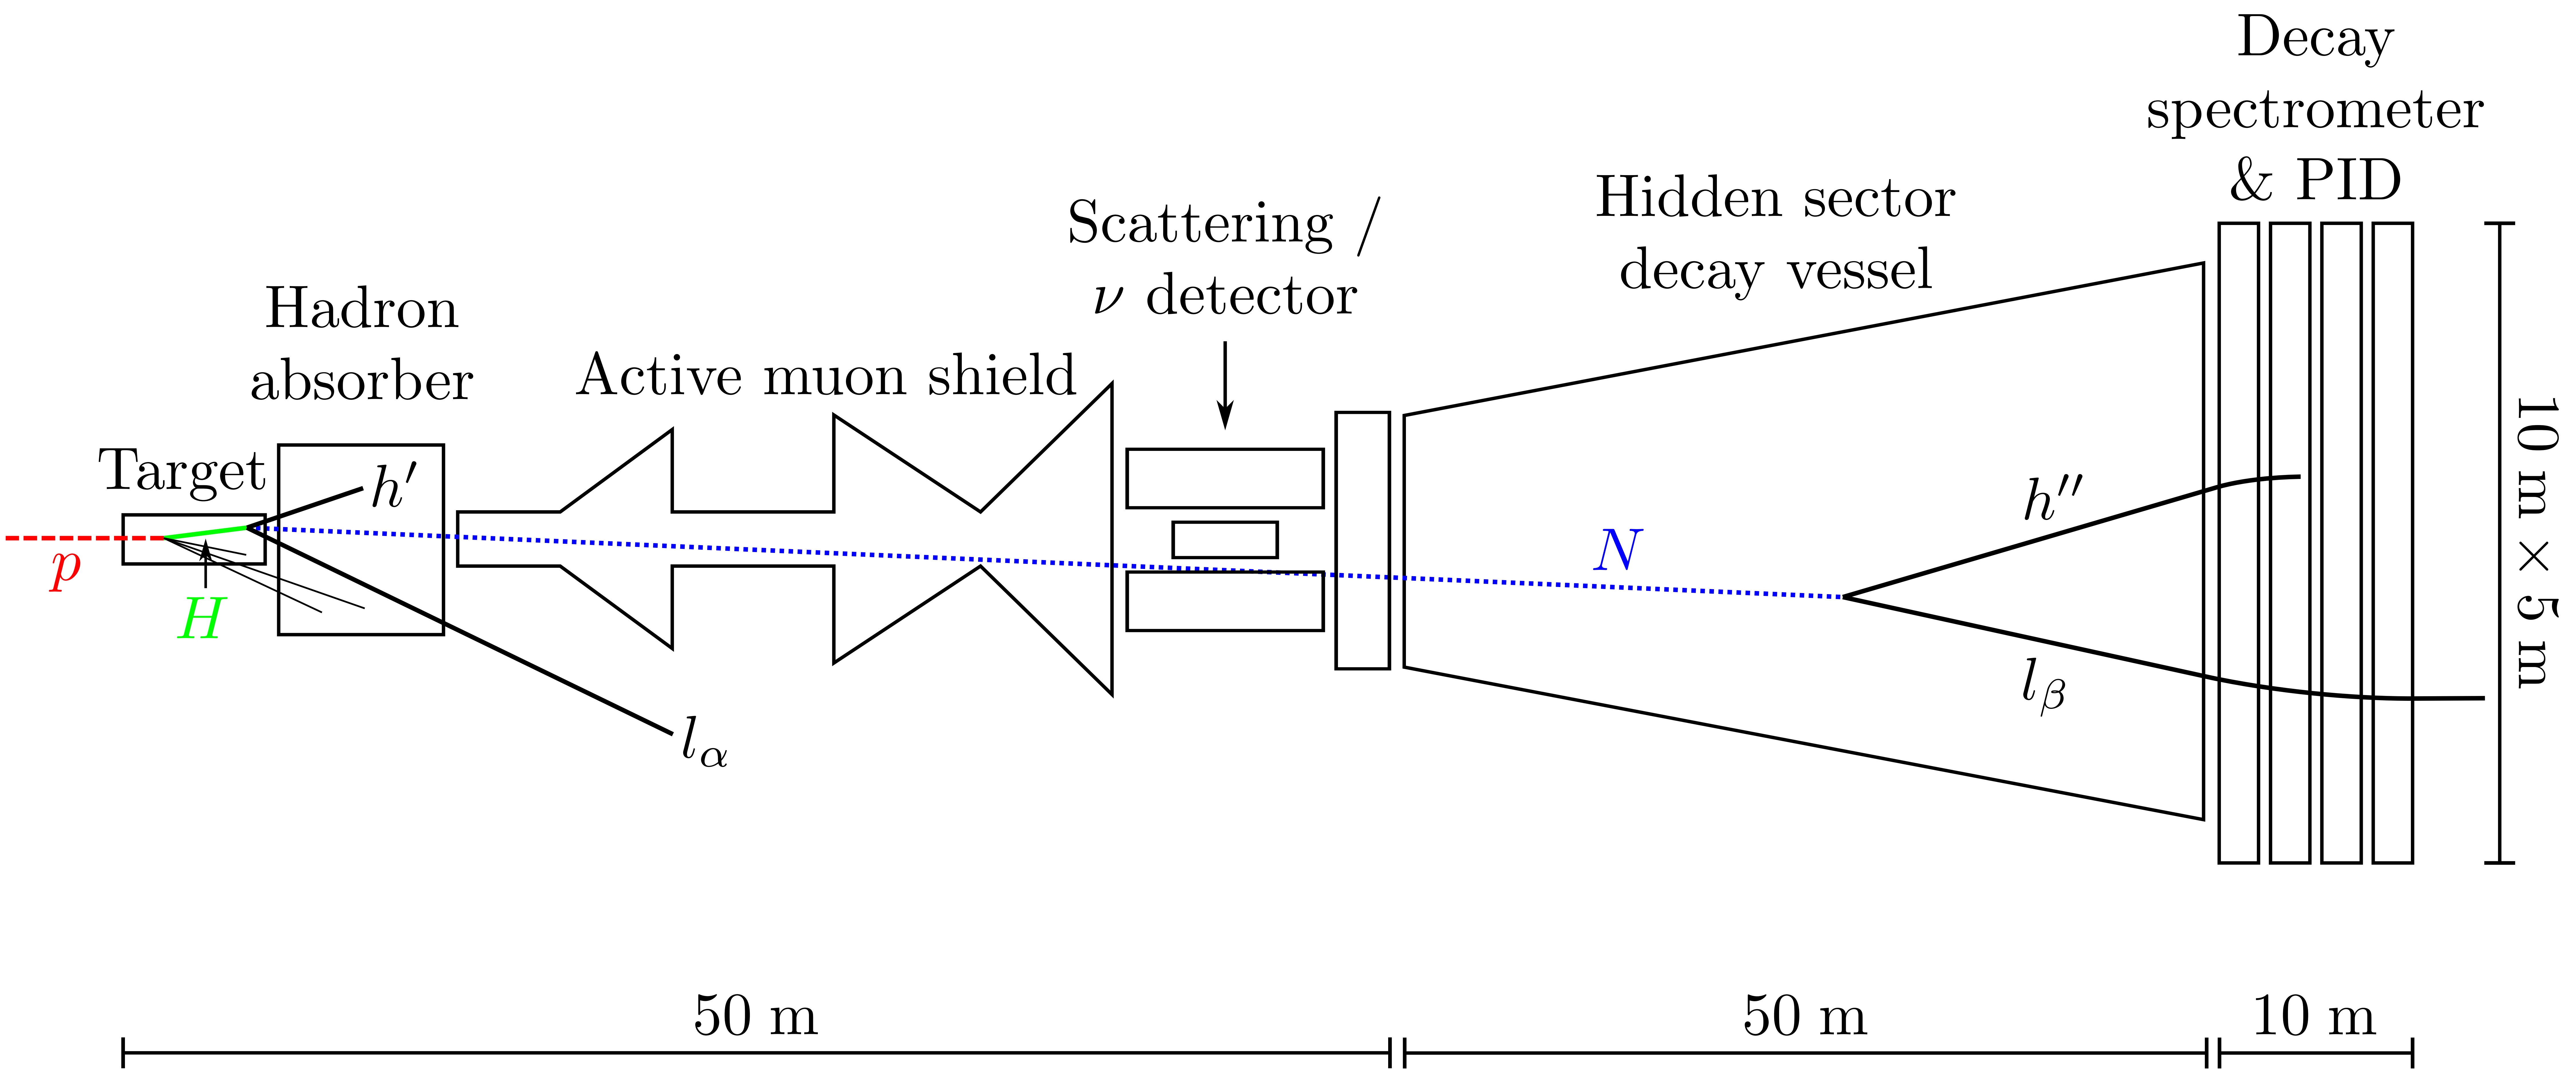

HNLs at SHiP

Event display ($N \to \pi^- \mu^+$)

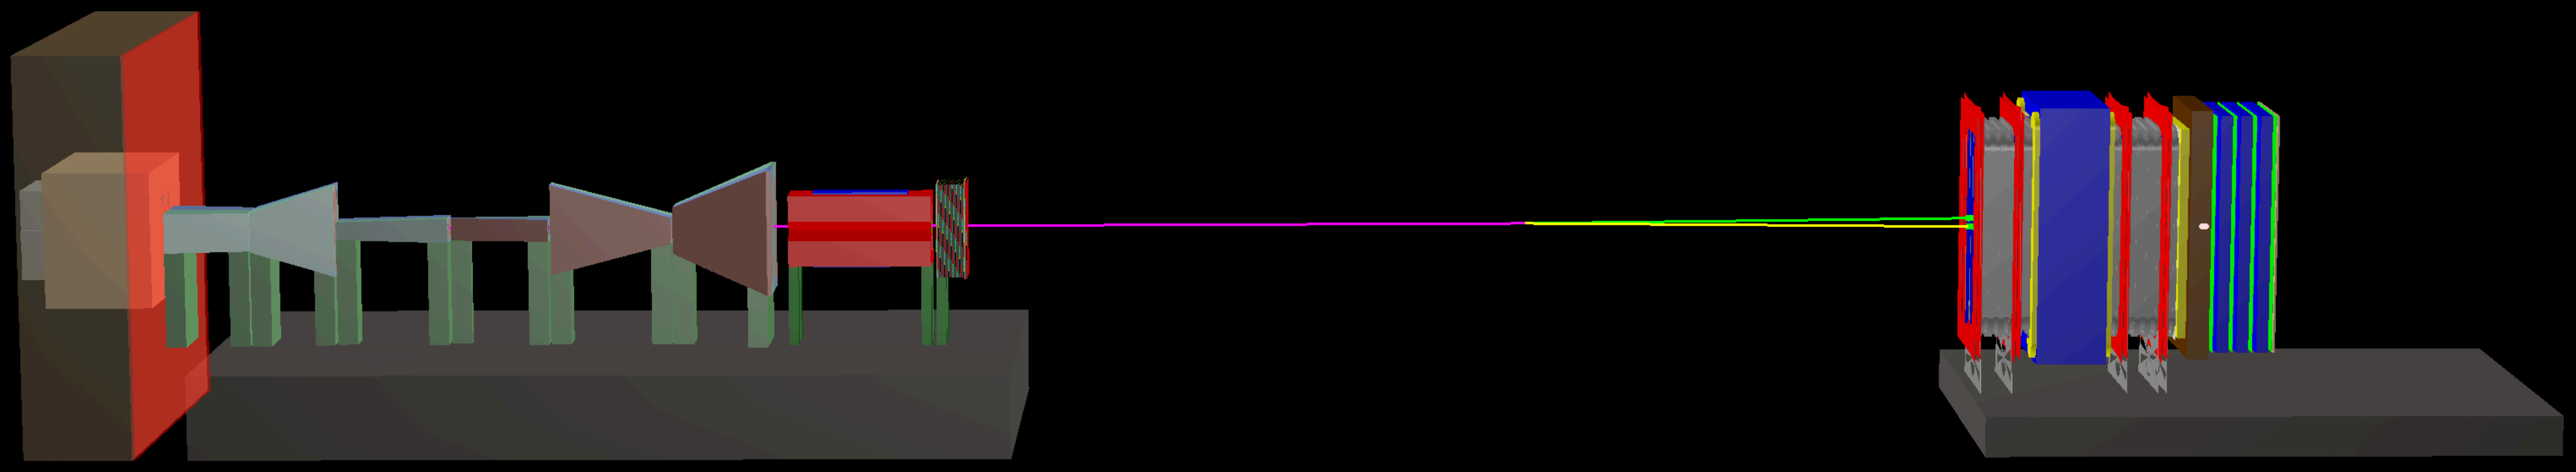

Meeting the experimental requirements

Feeble interactions lead to:

- Suppressed production rate

- A possibly long lifetime

Meeting the experimental requirements

Feeble interactions lead to:

- Suppressed production rate

- A possibly long lifetime

SHiP experiment

- 2×10^{20} POT / 5 yr
- 3×10^{17} D mesons, 3×10^{13} B mesons
- Large acceptance (detector close to target)
- 0.1–0.3 background events / 5 yr

Meeting the experimental requirements

SHiP experiment

Feeble interactions lead to:

- 2×10^{20} POT / 5 yr
- 3×10^{17} D mesons, 3×10^{13} B mesons
- Suppressed production rate
- Large acceptance (detector close to target)
- 0.1–0.3 background events / 5 yr
- Large 50m × 10m × 5m decay volume
- A possibly long lifetime
- Spectrometer with PID to reconstruct DV

Zero background, really?!

- Decay vessel is evacuated (< 1 mbar) to suppress interactions from neutrinos
- Active muon shield efficiently deflects muons away from the detector
- Surround veto around the decay vessel
- Requirement that the decay points back to the target (no MET) or near it (if MET).
- Timing coincidence to reduce combinatorial background

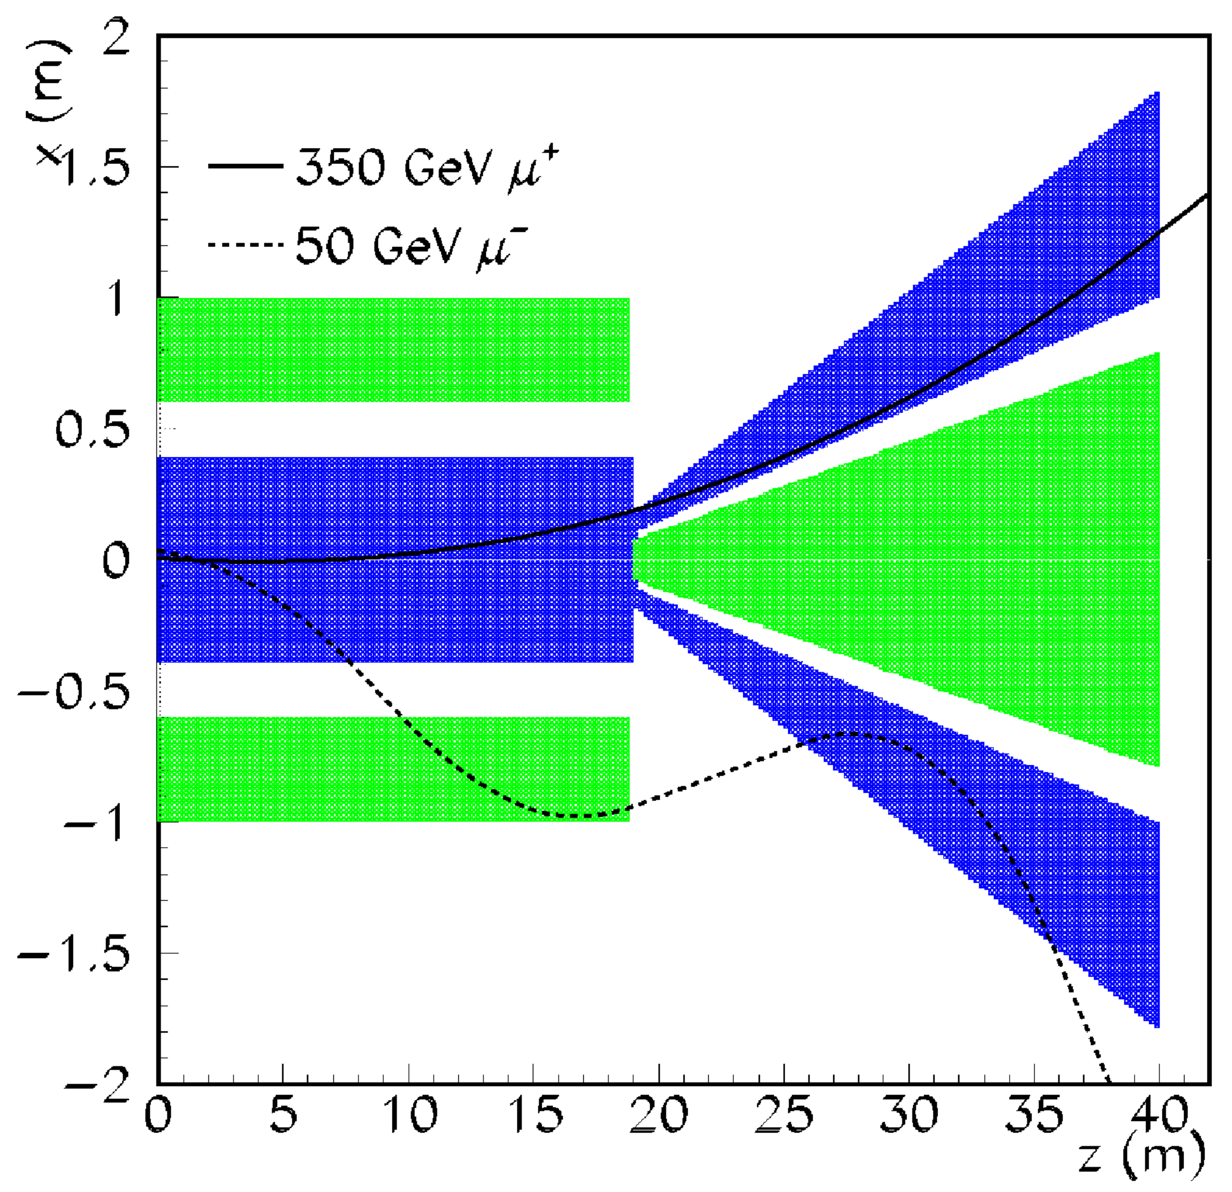

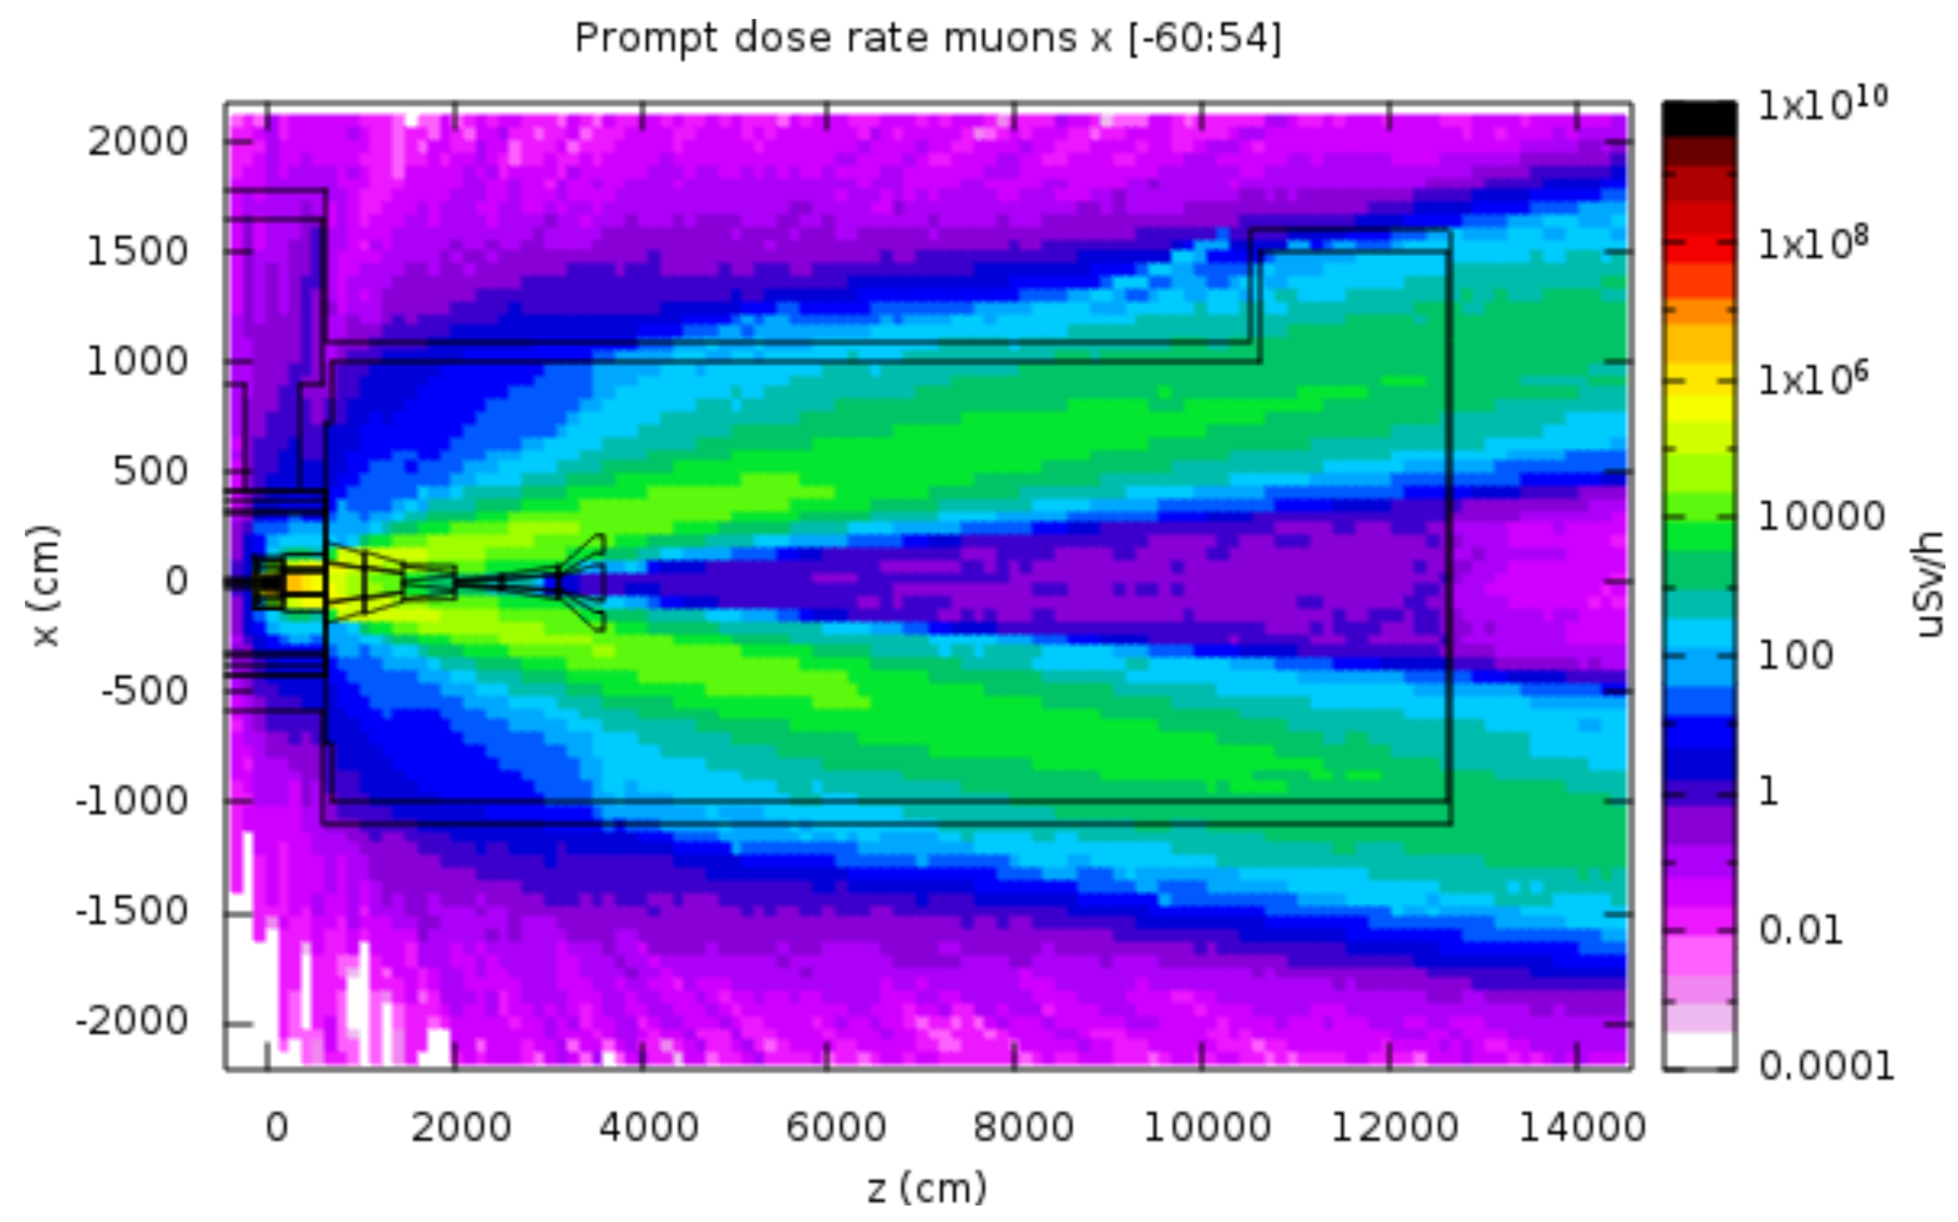

SHiP sensitivity to HNLs

[SHiP: 1811.00930]

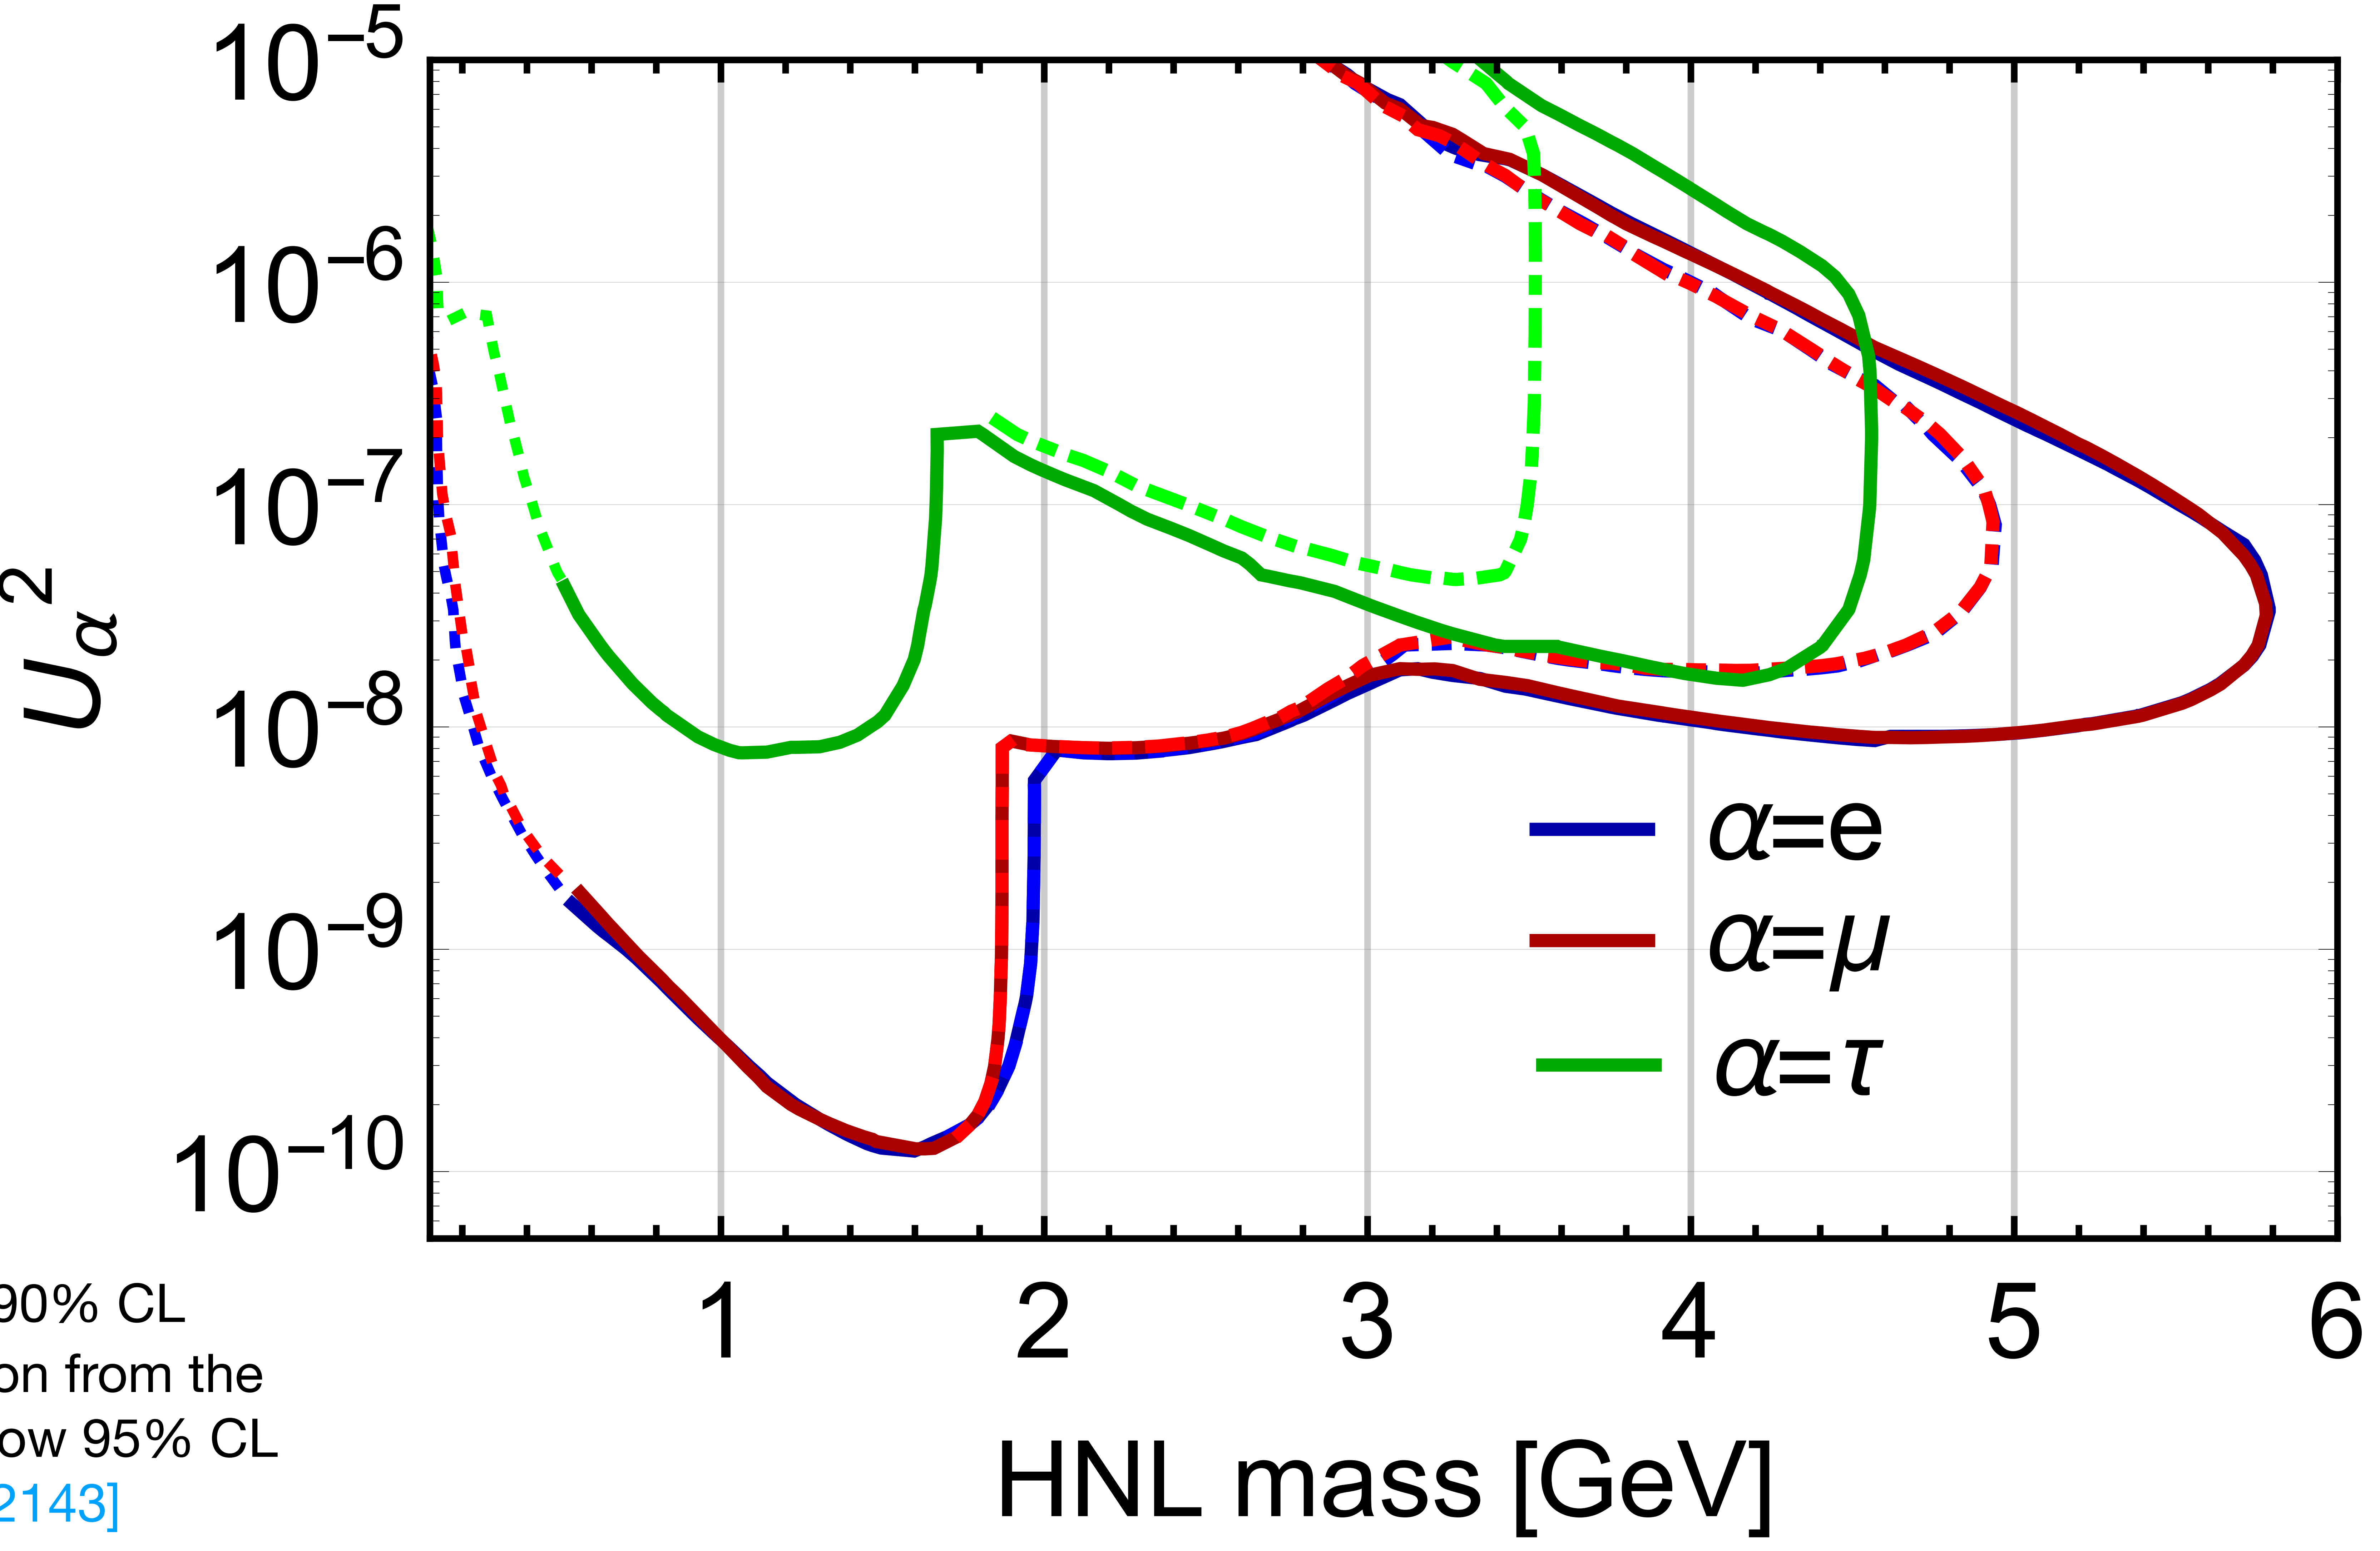

Solid line includes the B_c contribution (large uncertainty)

⚠️ Limits given for 90% CL
The recommendation from the FIPs workshop is now 95% CL
[FIPs 2020: 2102.12143]

SHiP sensitivity to HNLs

[SHiP: 1811.00930]

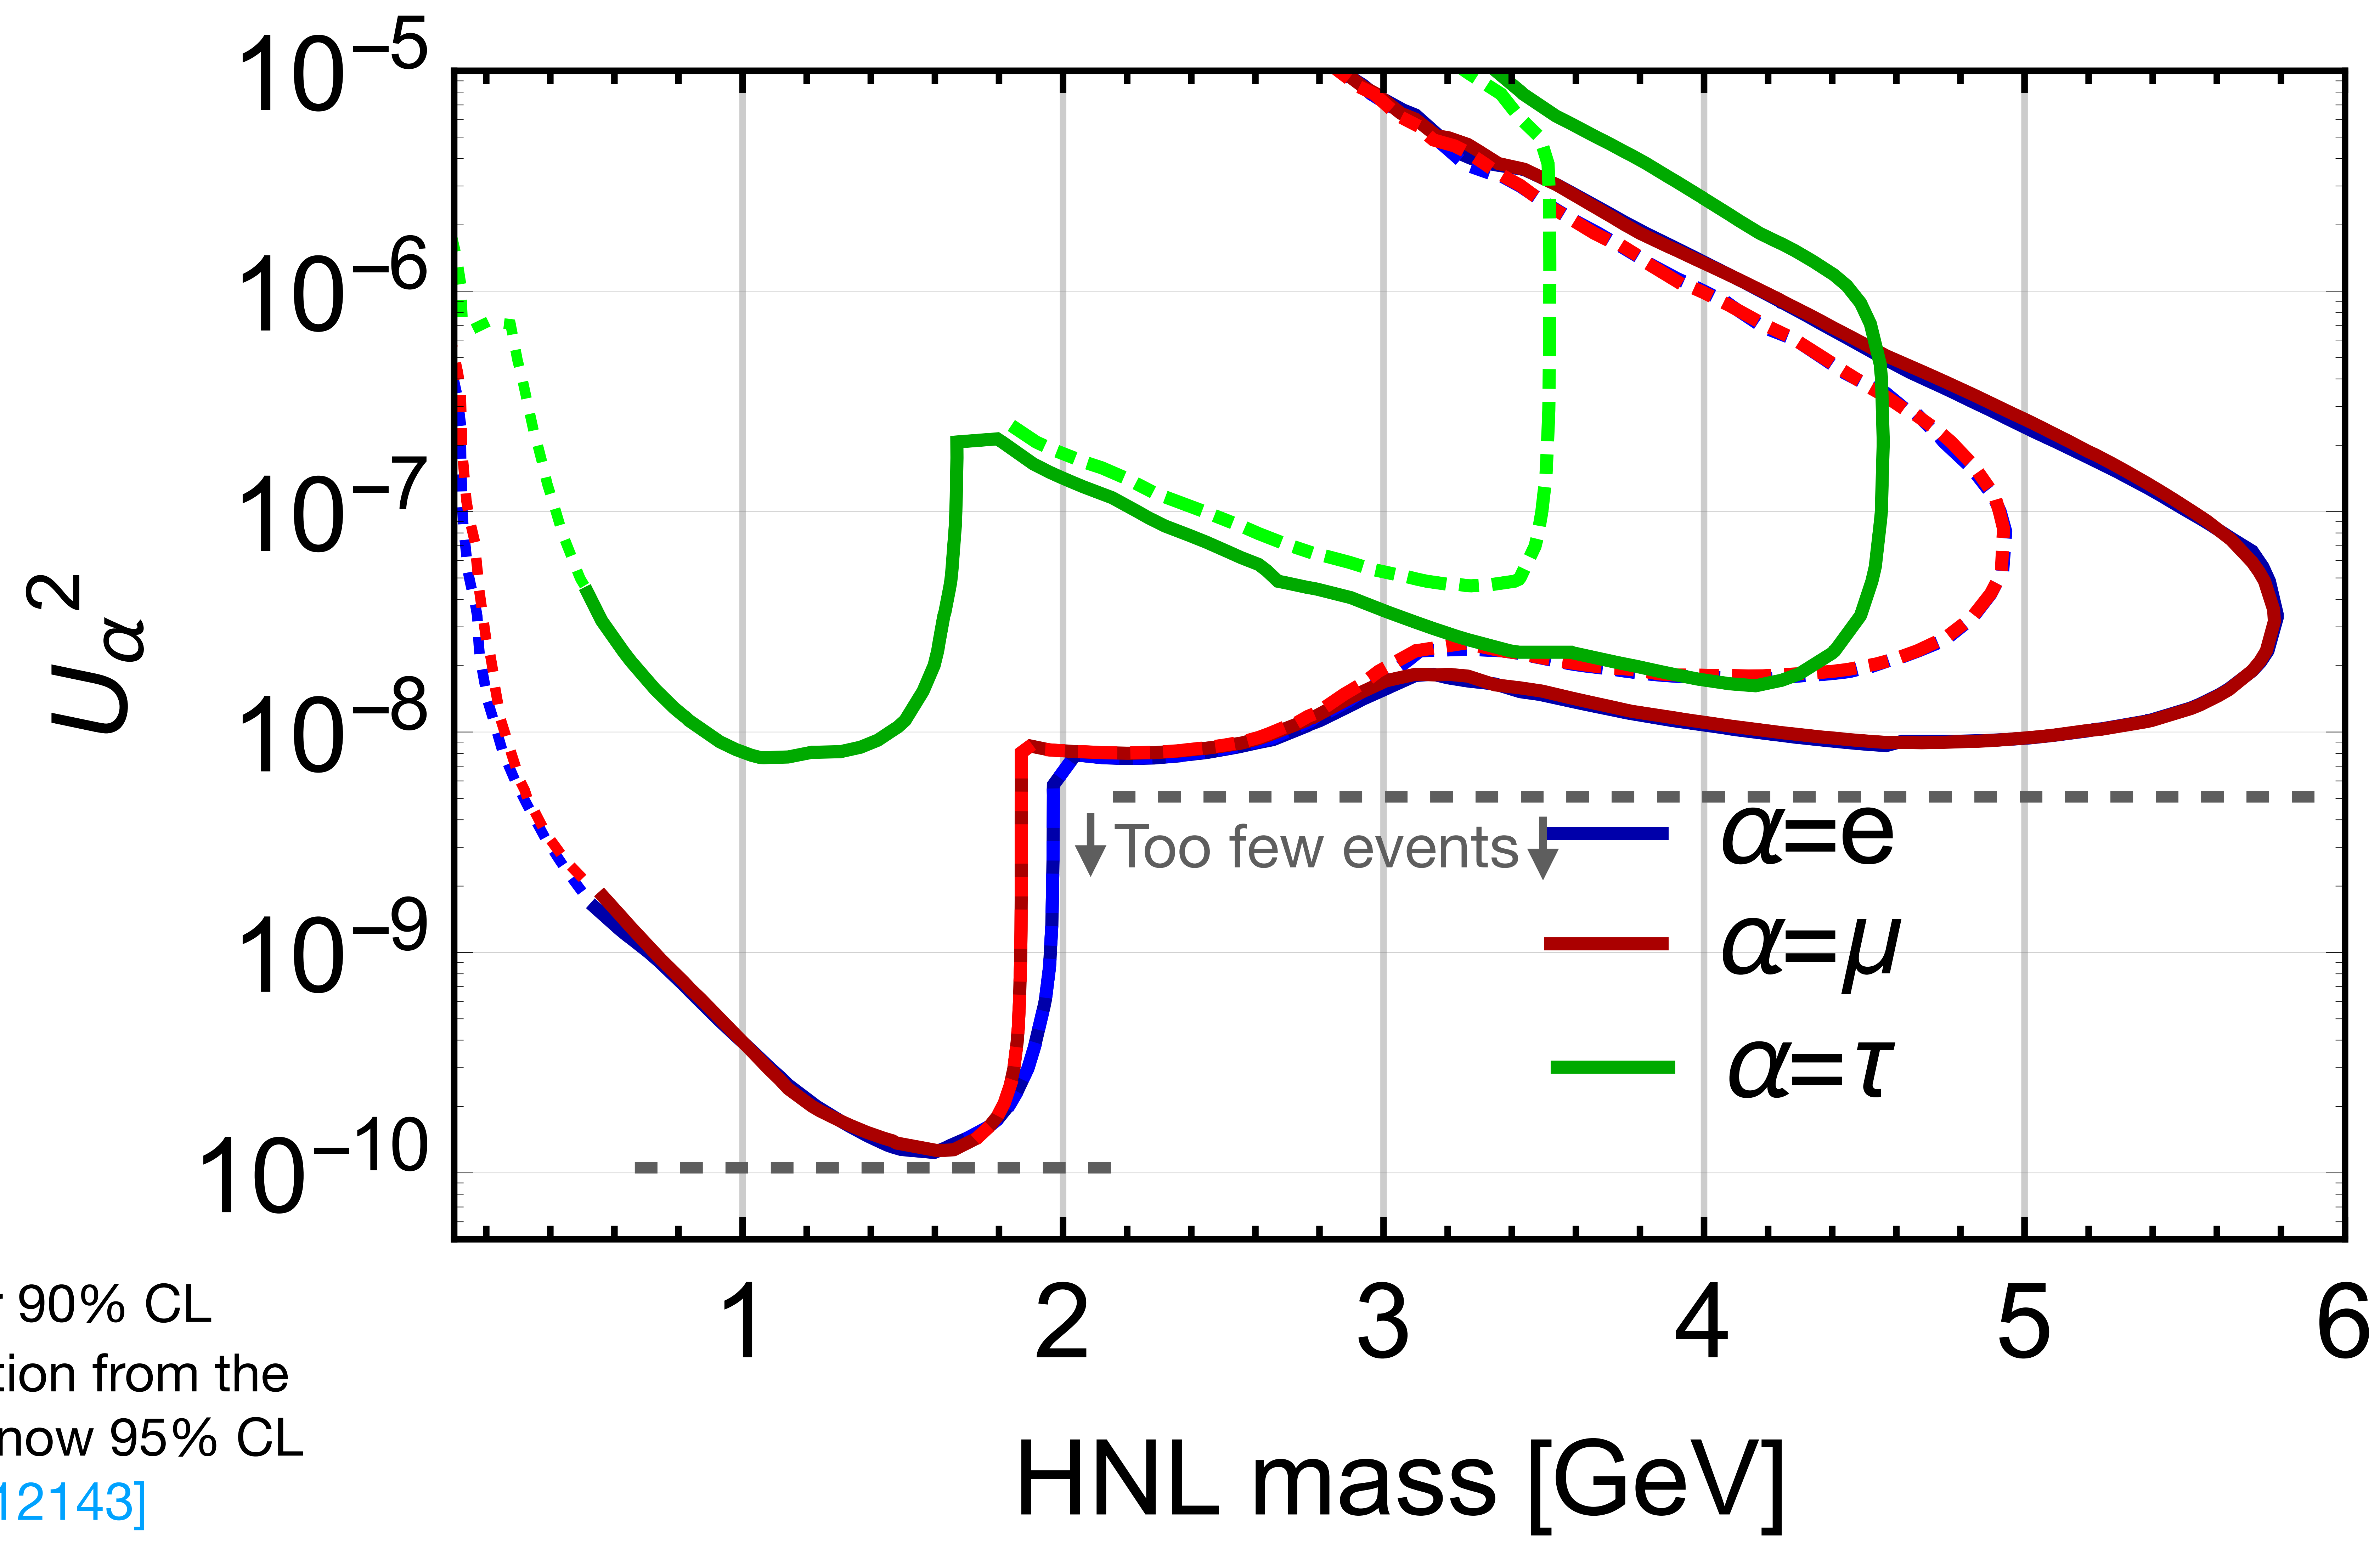

Solid line includes the B_c contribution (large uncertainty)

⚠️ Limits given for 90% CL
The recommendation from the FIPs workshop is now 95% CL
[FIPs 2020: 2102.12143]

SHiP sensitivity to HNLs

[SHiP: 1811.00930]

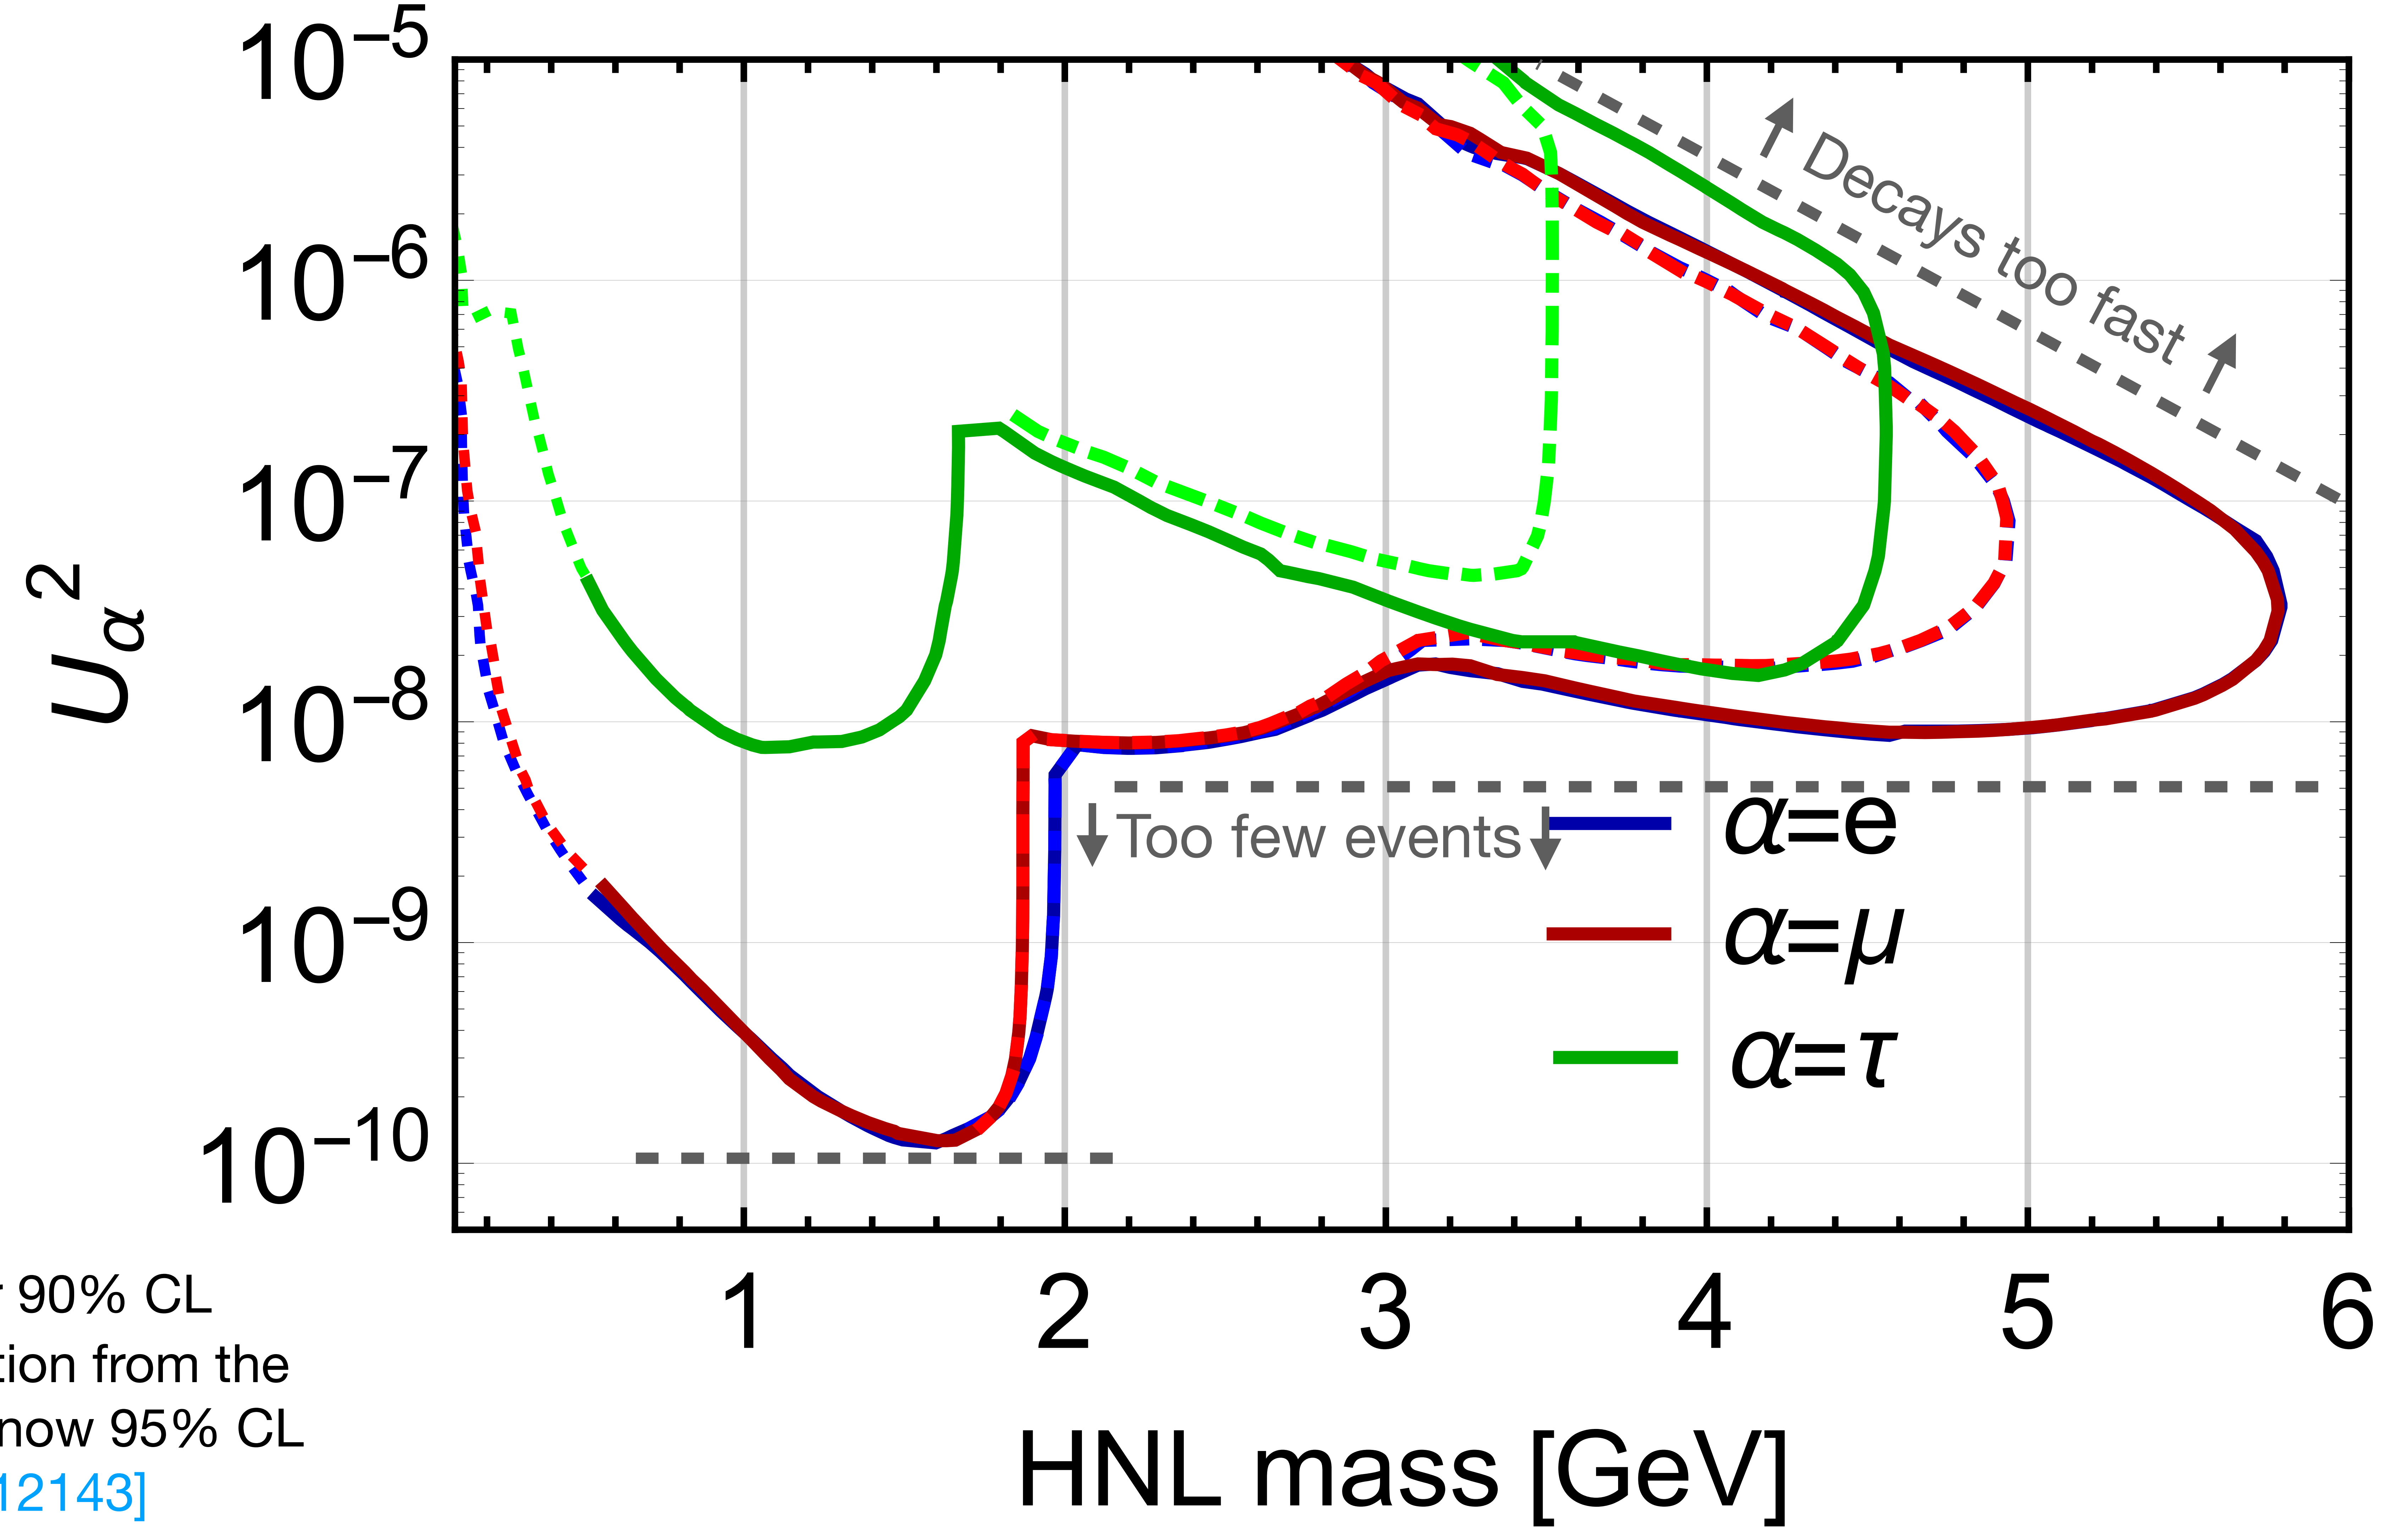

Solid line includes the B_c contribution (large uncertainty)

⚠️ Limits given for 90% CL
The recommendation from the FIPs workshop is now 95% CL
[FIPs 2020: 2102.12143]

SHiP sensitivity to HNLs

[SHiP: 1811.00930]

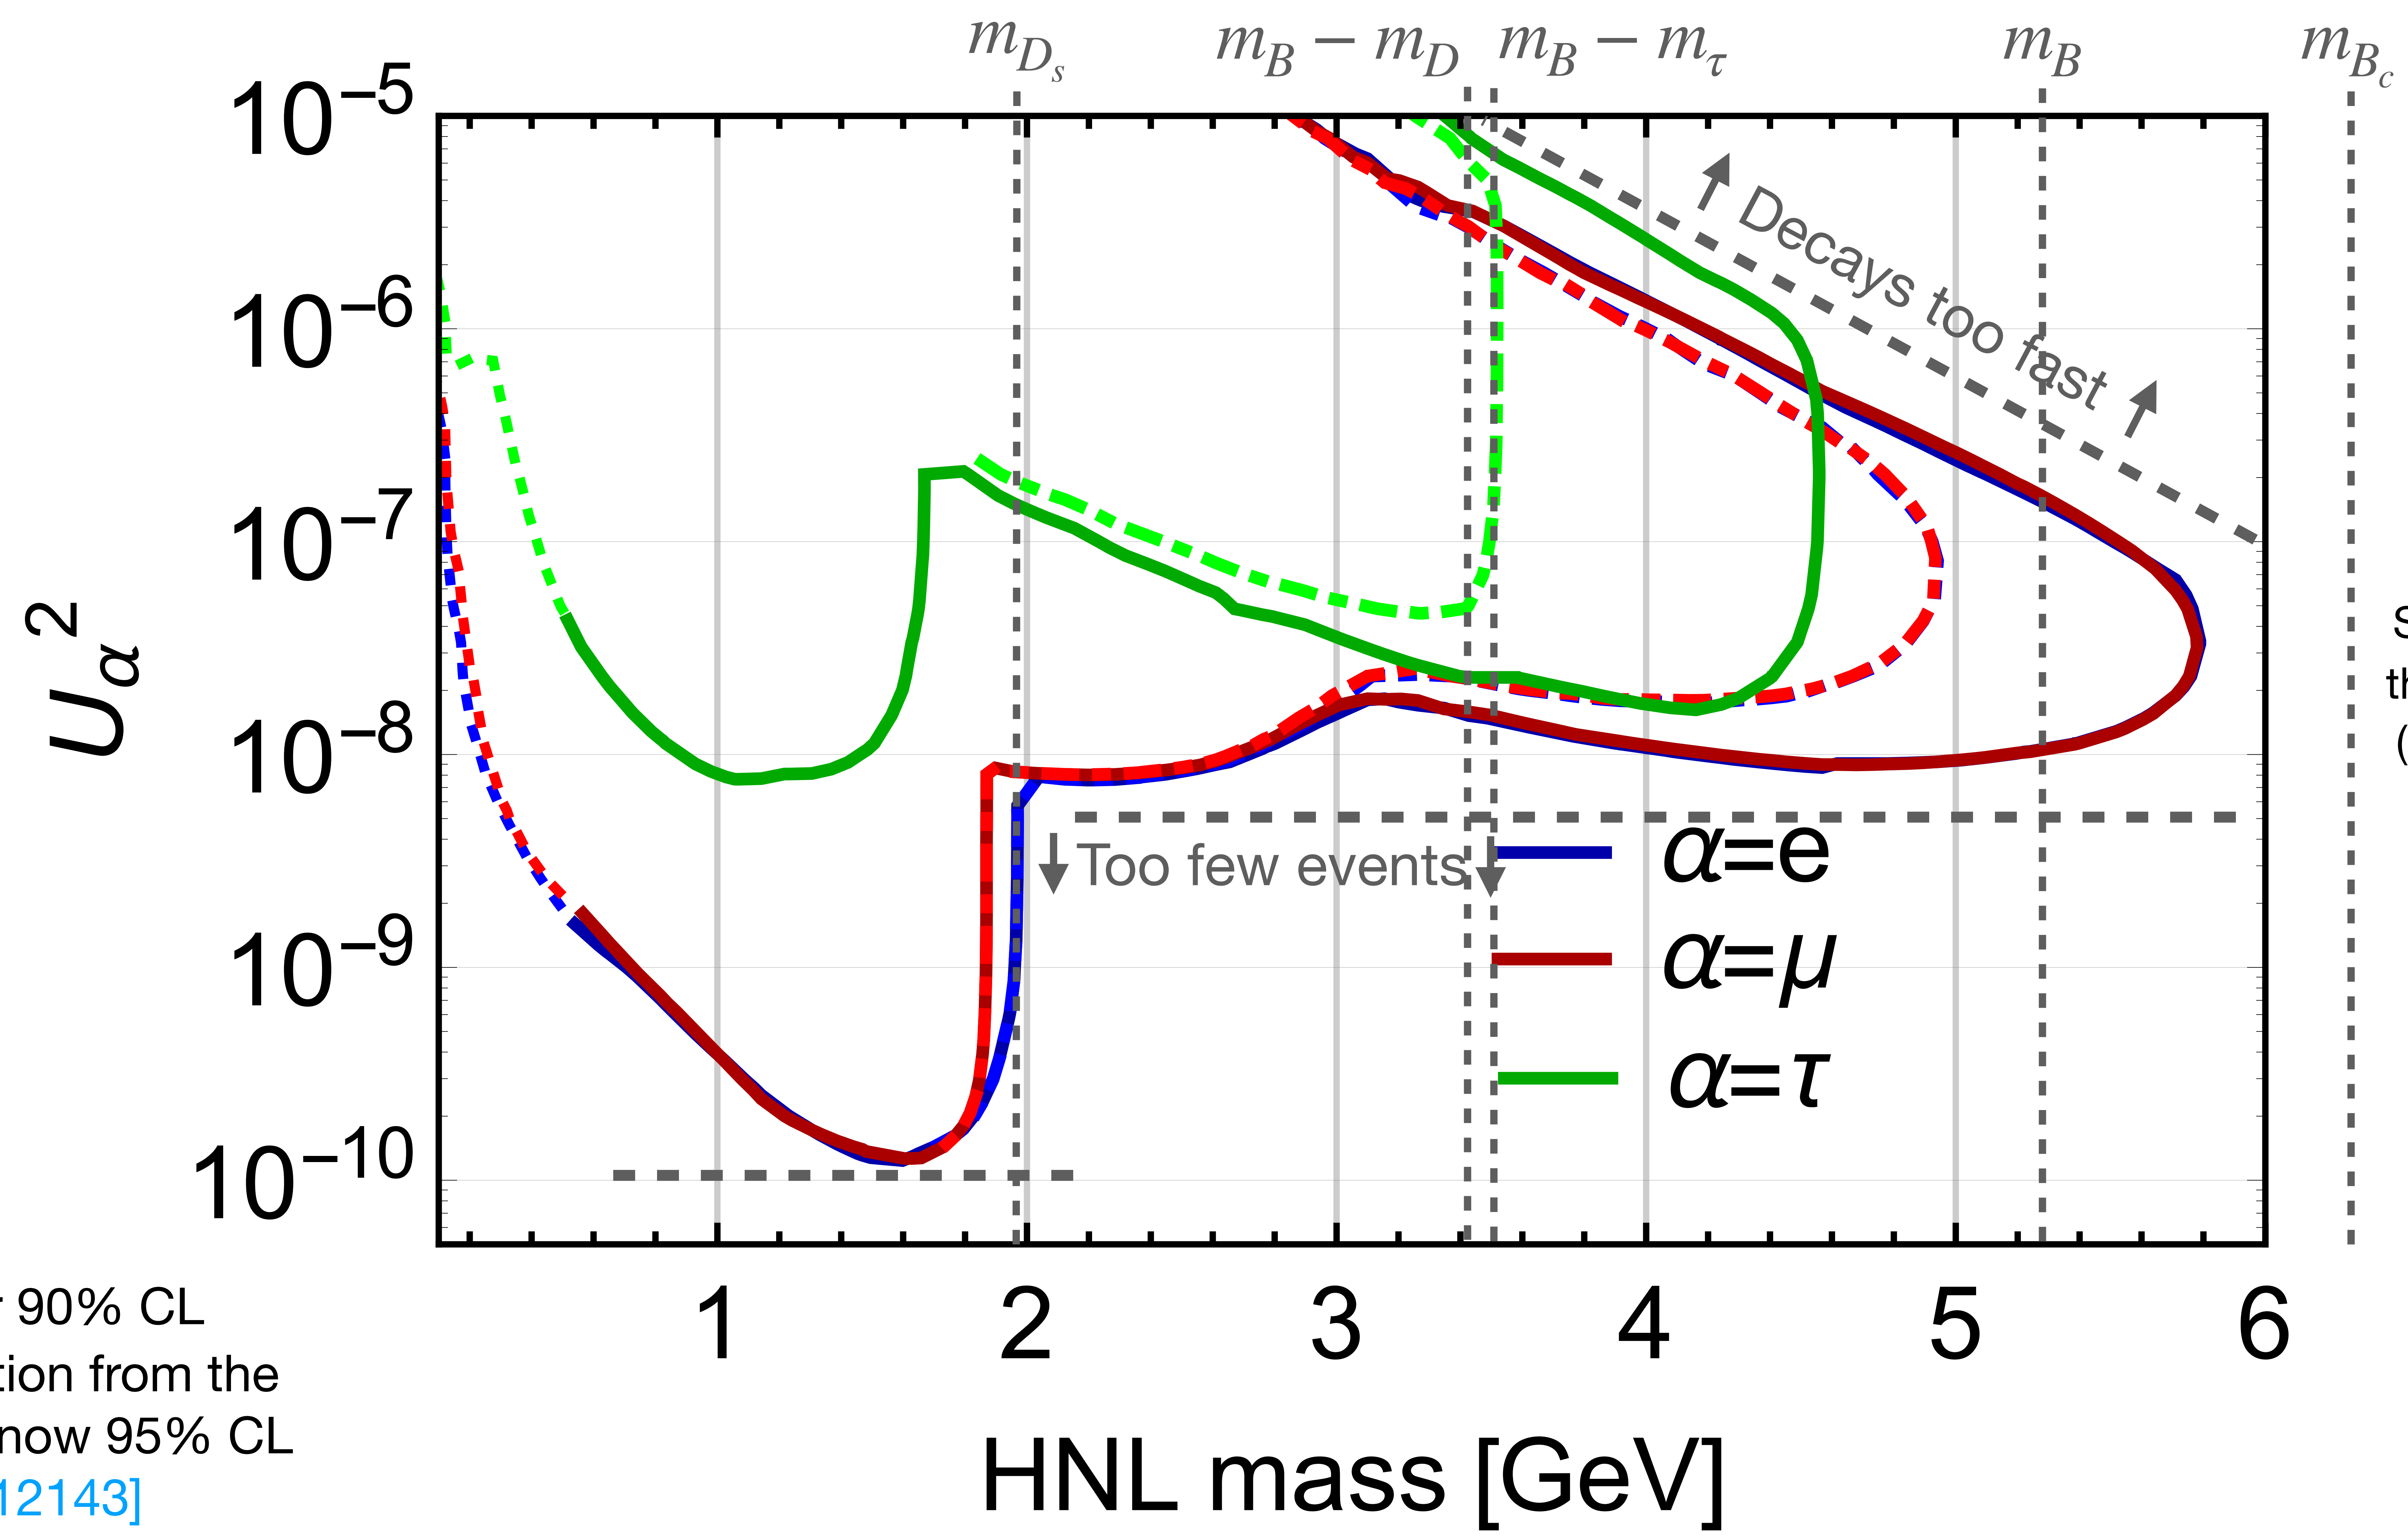

Solid line includes the B_c contribution (large uncertainty)

⚠️ Limits given for 90% CL
The recommendation from the FIPs workshop is now 95% CL
[FIPs 2020: 2102.12143]

Upper mass limit

Easier to see in log scale

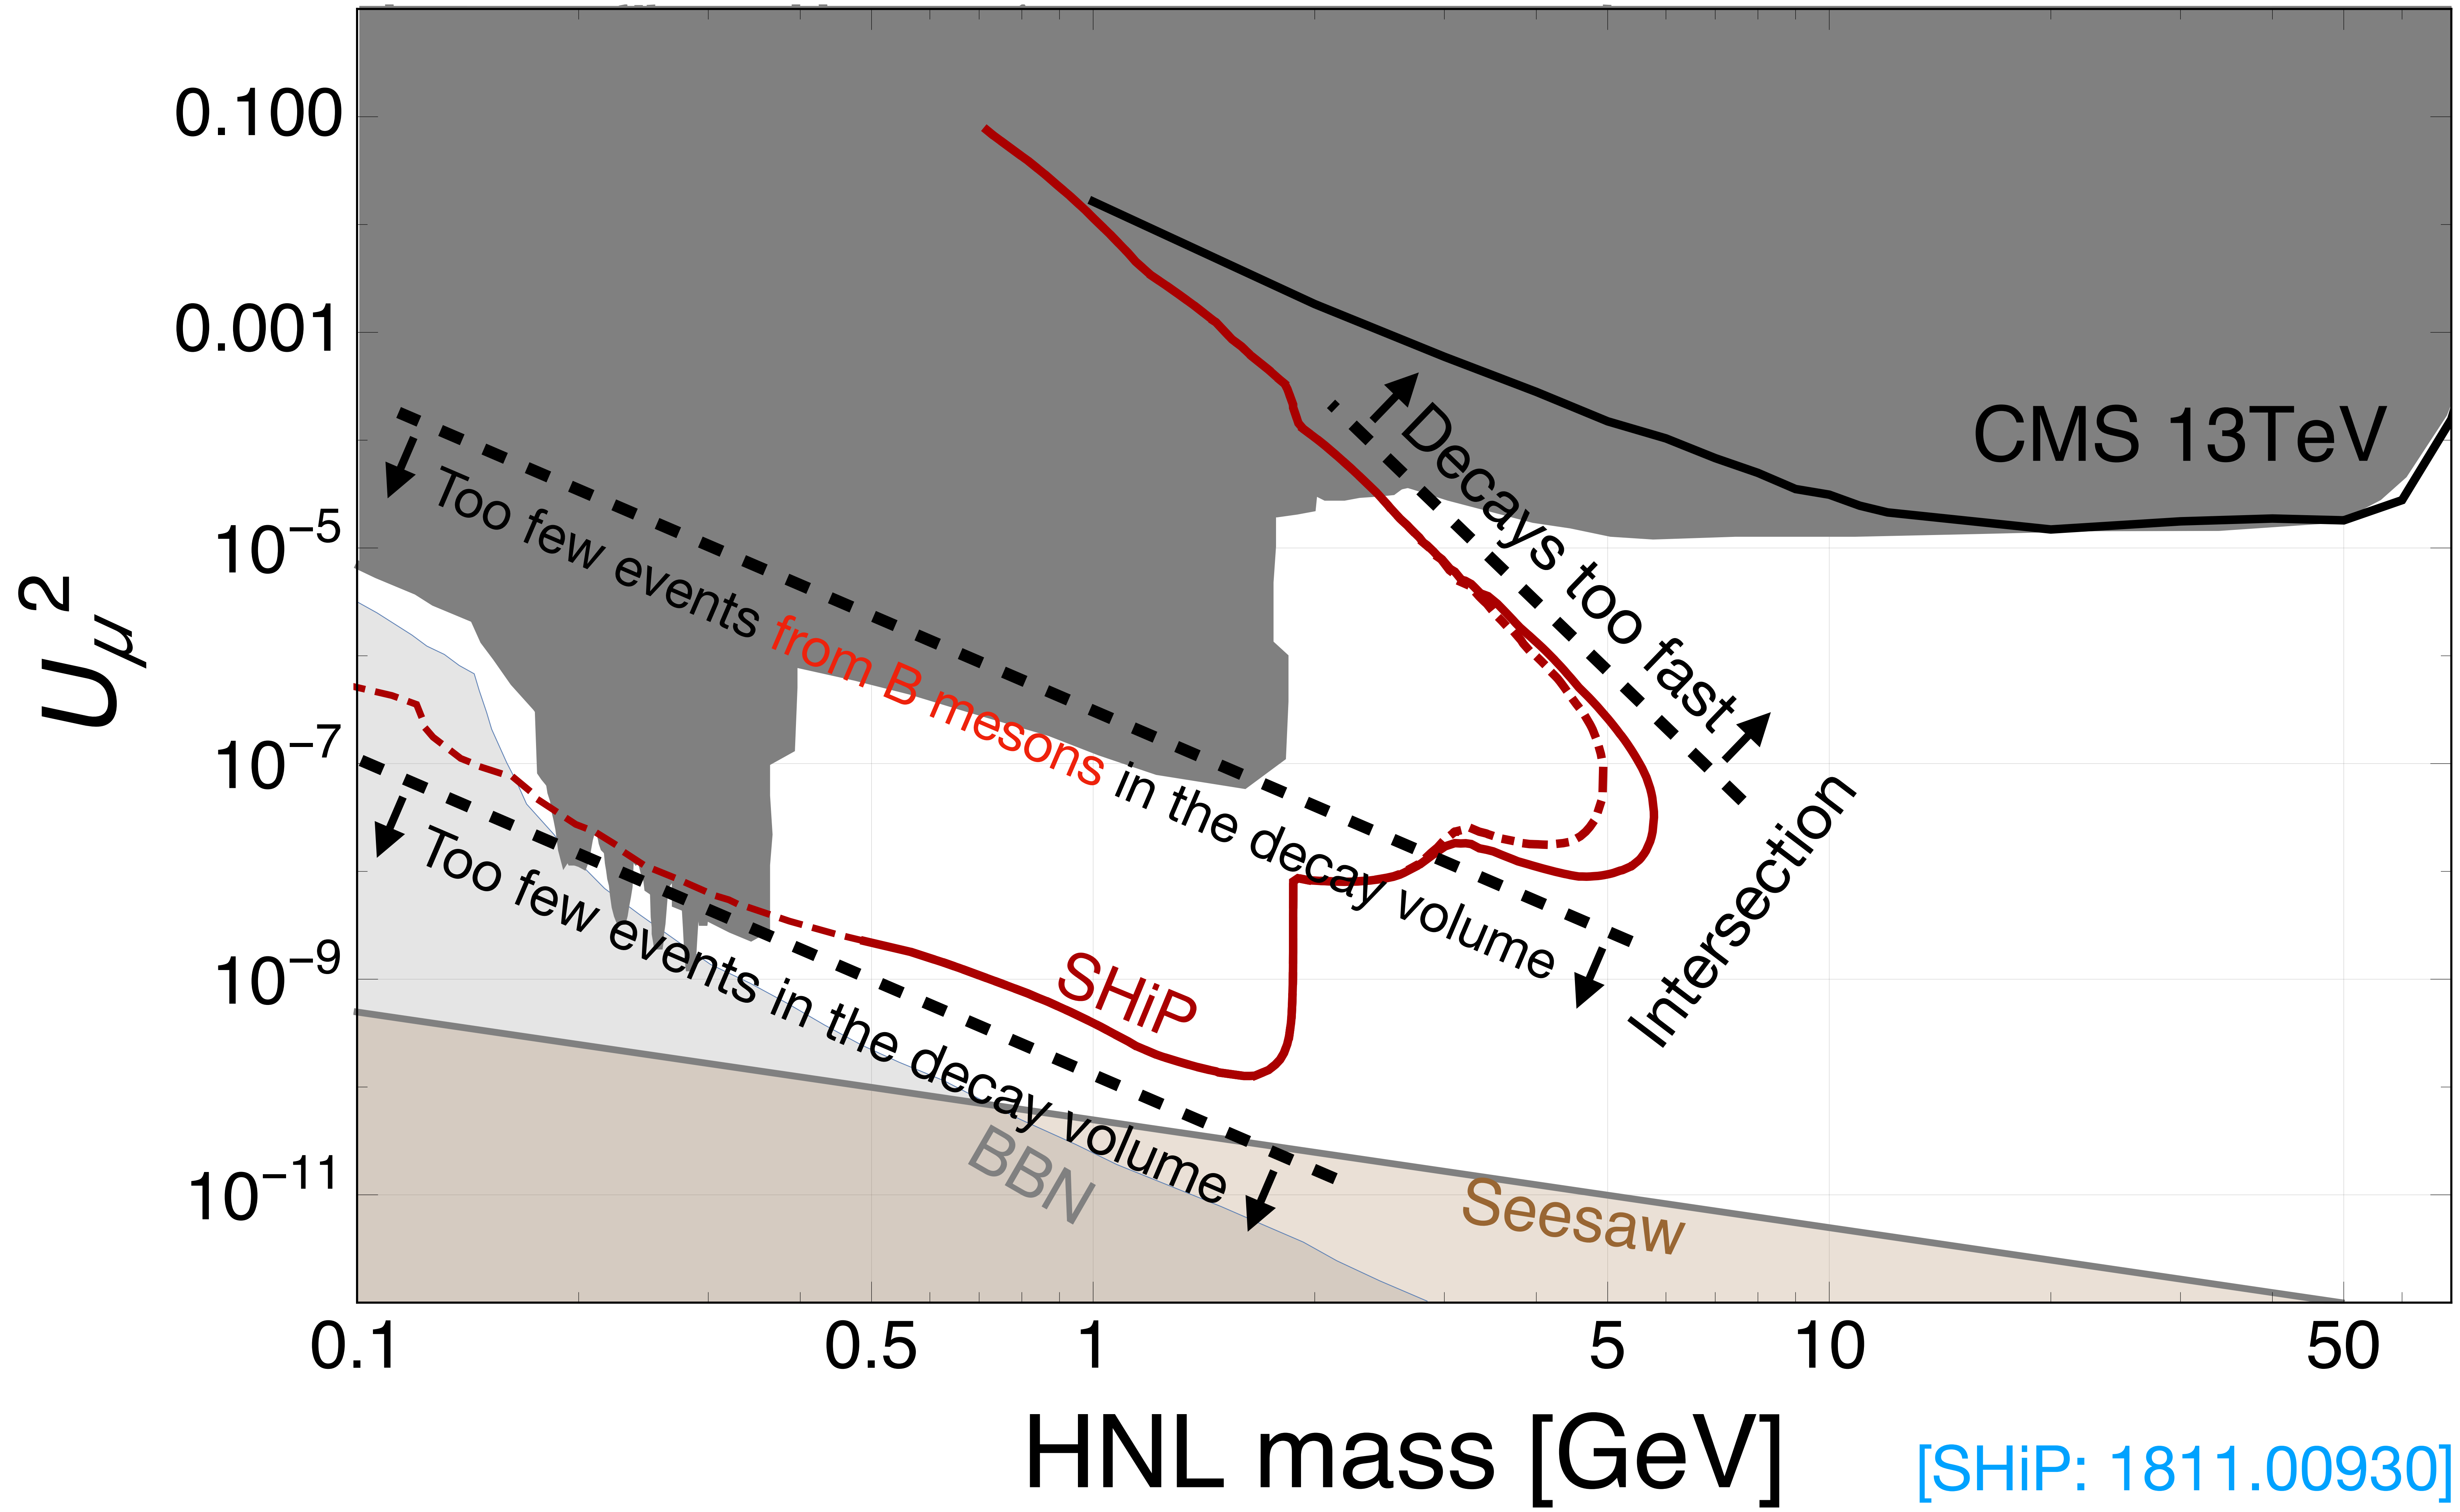

[SHiP: 1811.00930]

Remark:

Exclusion sensitivity $\neq$ discovery sensitivity

$\approx$ zero background

No events $\Longrightarrow$ < 2.3 exp. signal @ 95% CL

0.1 exp. bkg. (fully reconstructed)

2 events $\Longrightarrow$ 3.8 σ evidence

4 events $\Longrightarrow$ 5.4 σ discovery

Remark:

Exclusion sensitivity ≠ discovery sensitivity

$\approx$ zero background

No events $\Longrightarrow$ < 2.3 exp. signal @ 95% CL

0.3 exp. bkg. (partially reconstructed)

2 events $\Longrightarrow$ 2.9 σ evidence

5 events $\Longrightarrow$ 4.9 σ discovery

Remark:

Exclusion sensitivity ≠ discovery sensitivity

$\approx$ zero background

No events $\Longrightarrow$ < 2.3 exp. signal @ 95% CL

0.3 exp. bkg. (partially reconstructed)

2 events $\Longrightarrow$ 2.9 σ evidence

5 events $\Longrightarrow$ 4.9 σ discovery

+ the SHiP tracker can reconstruct the decays & measure the particle mass

SHiP sensitivity to HNLs (electron mixing)

[SHiP: 1811.00930]

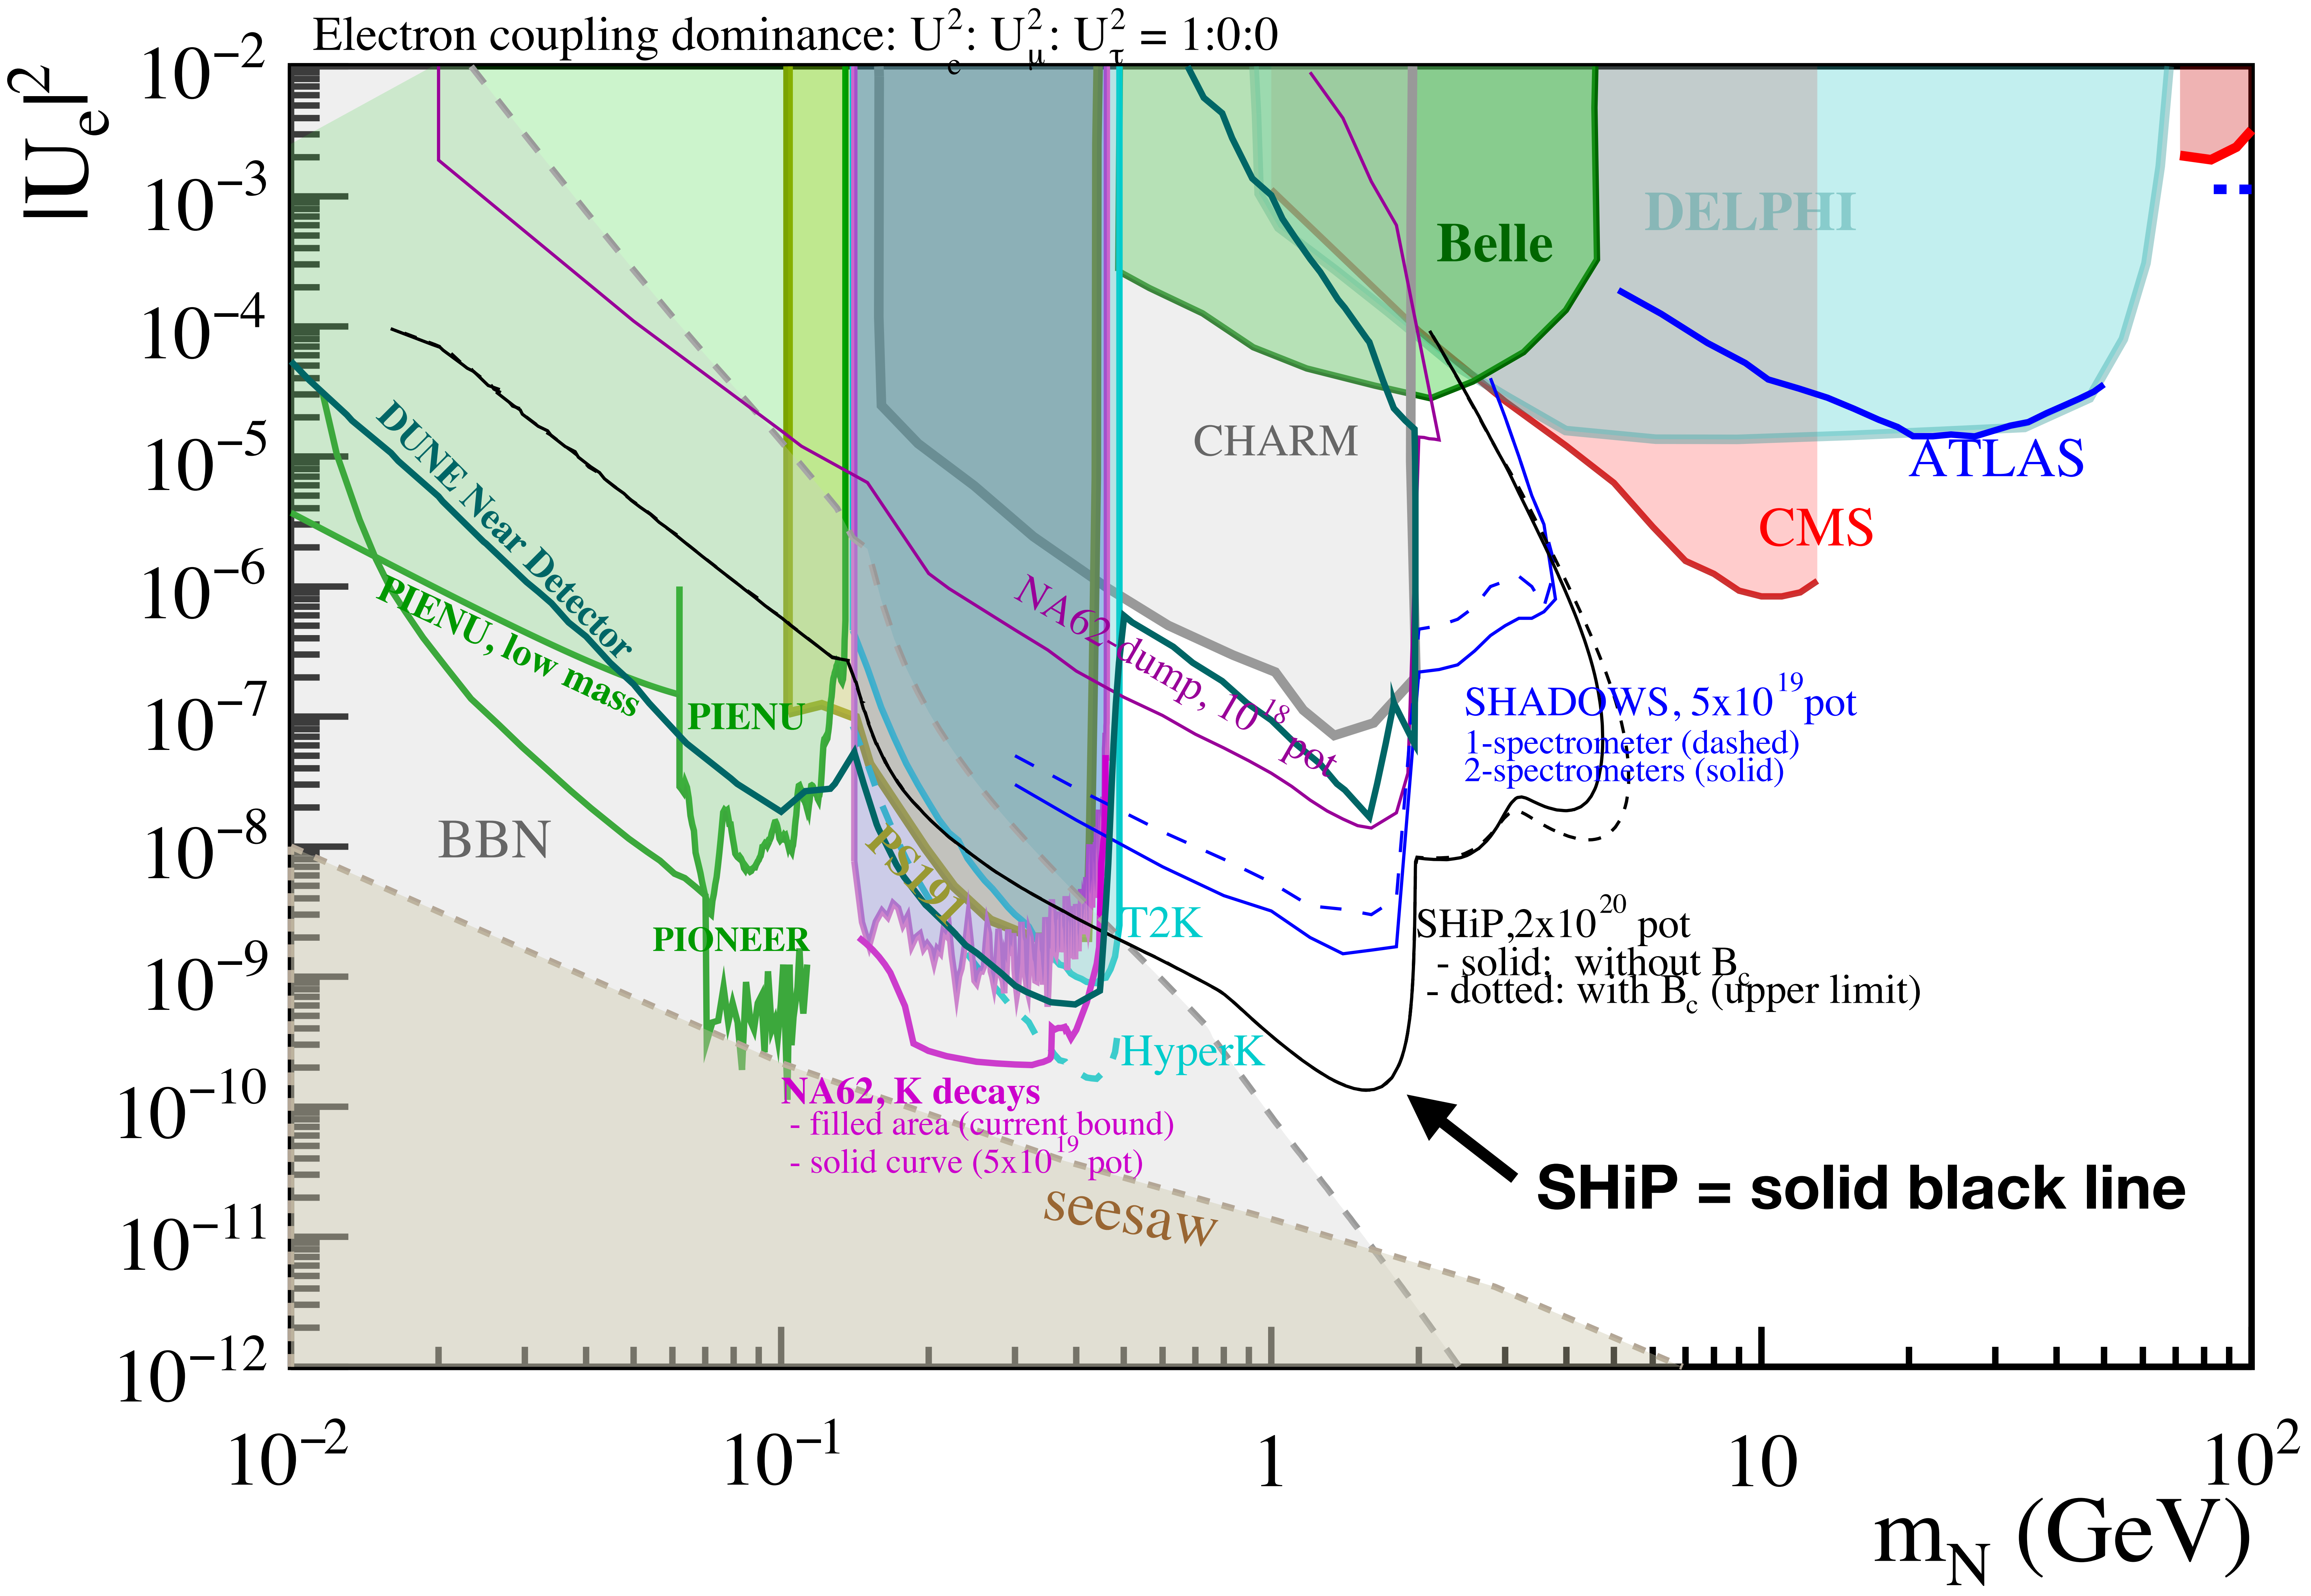

Plot from Snowmass paper [2203.08039] (45 authors)

SHiP sensitivity to HNLs (muon mixing)

[SHiP: 1811.00930]

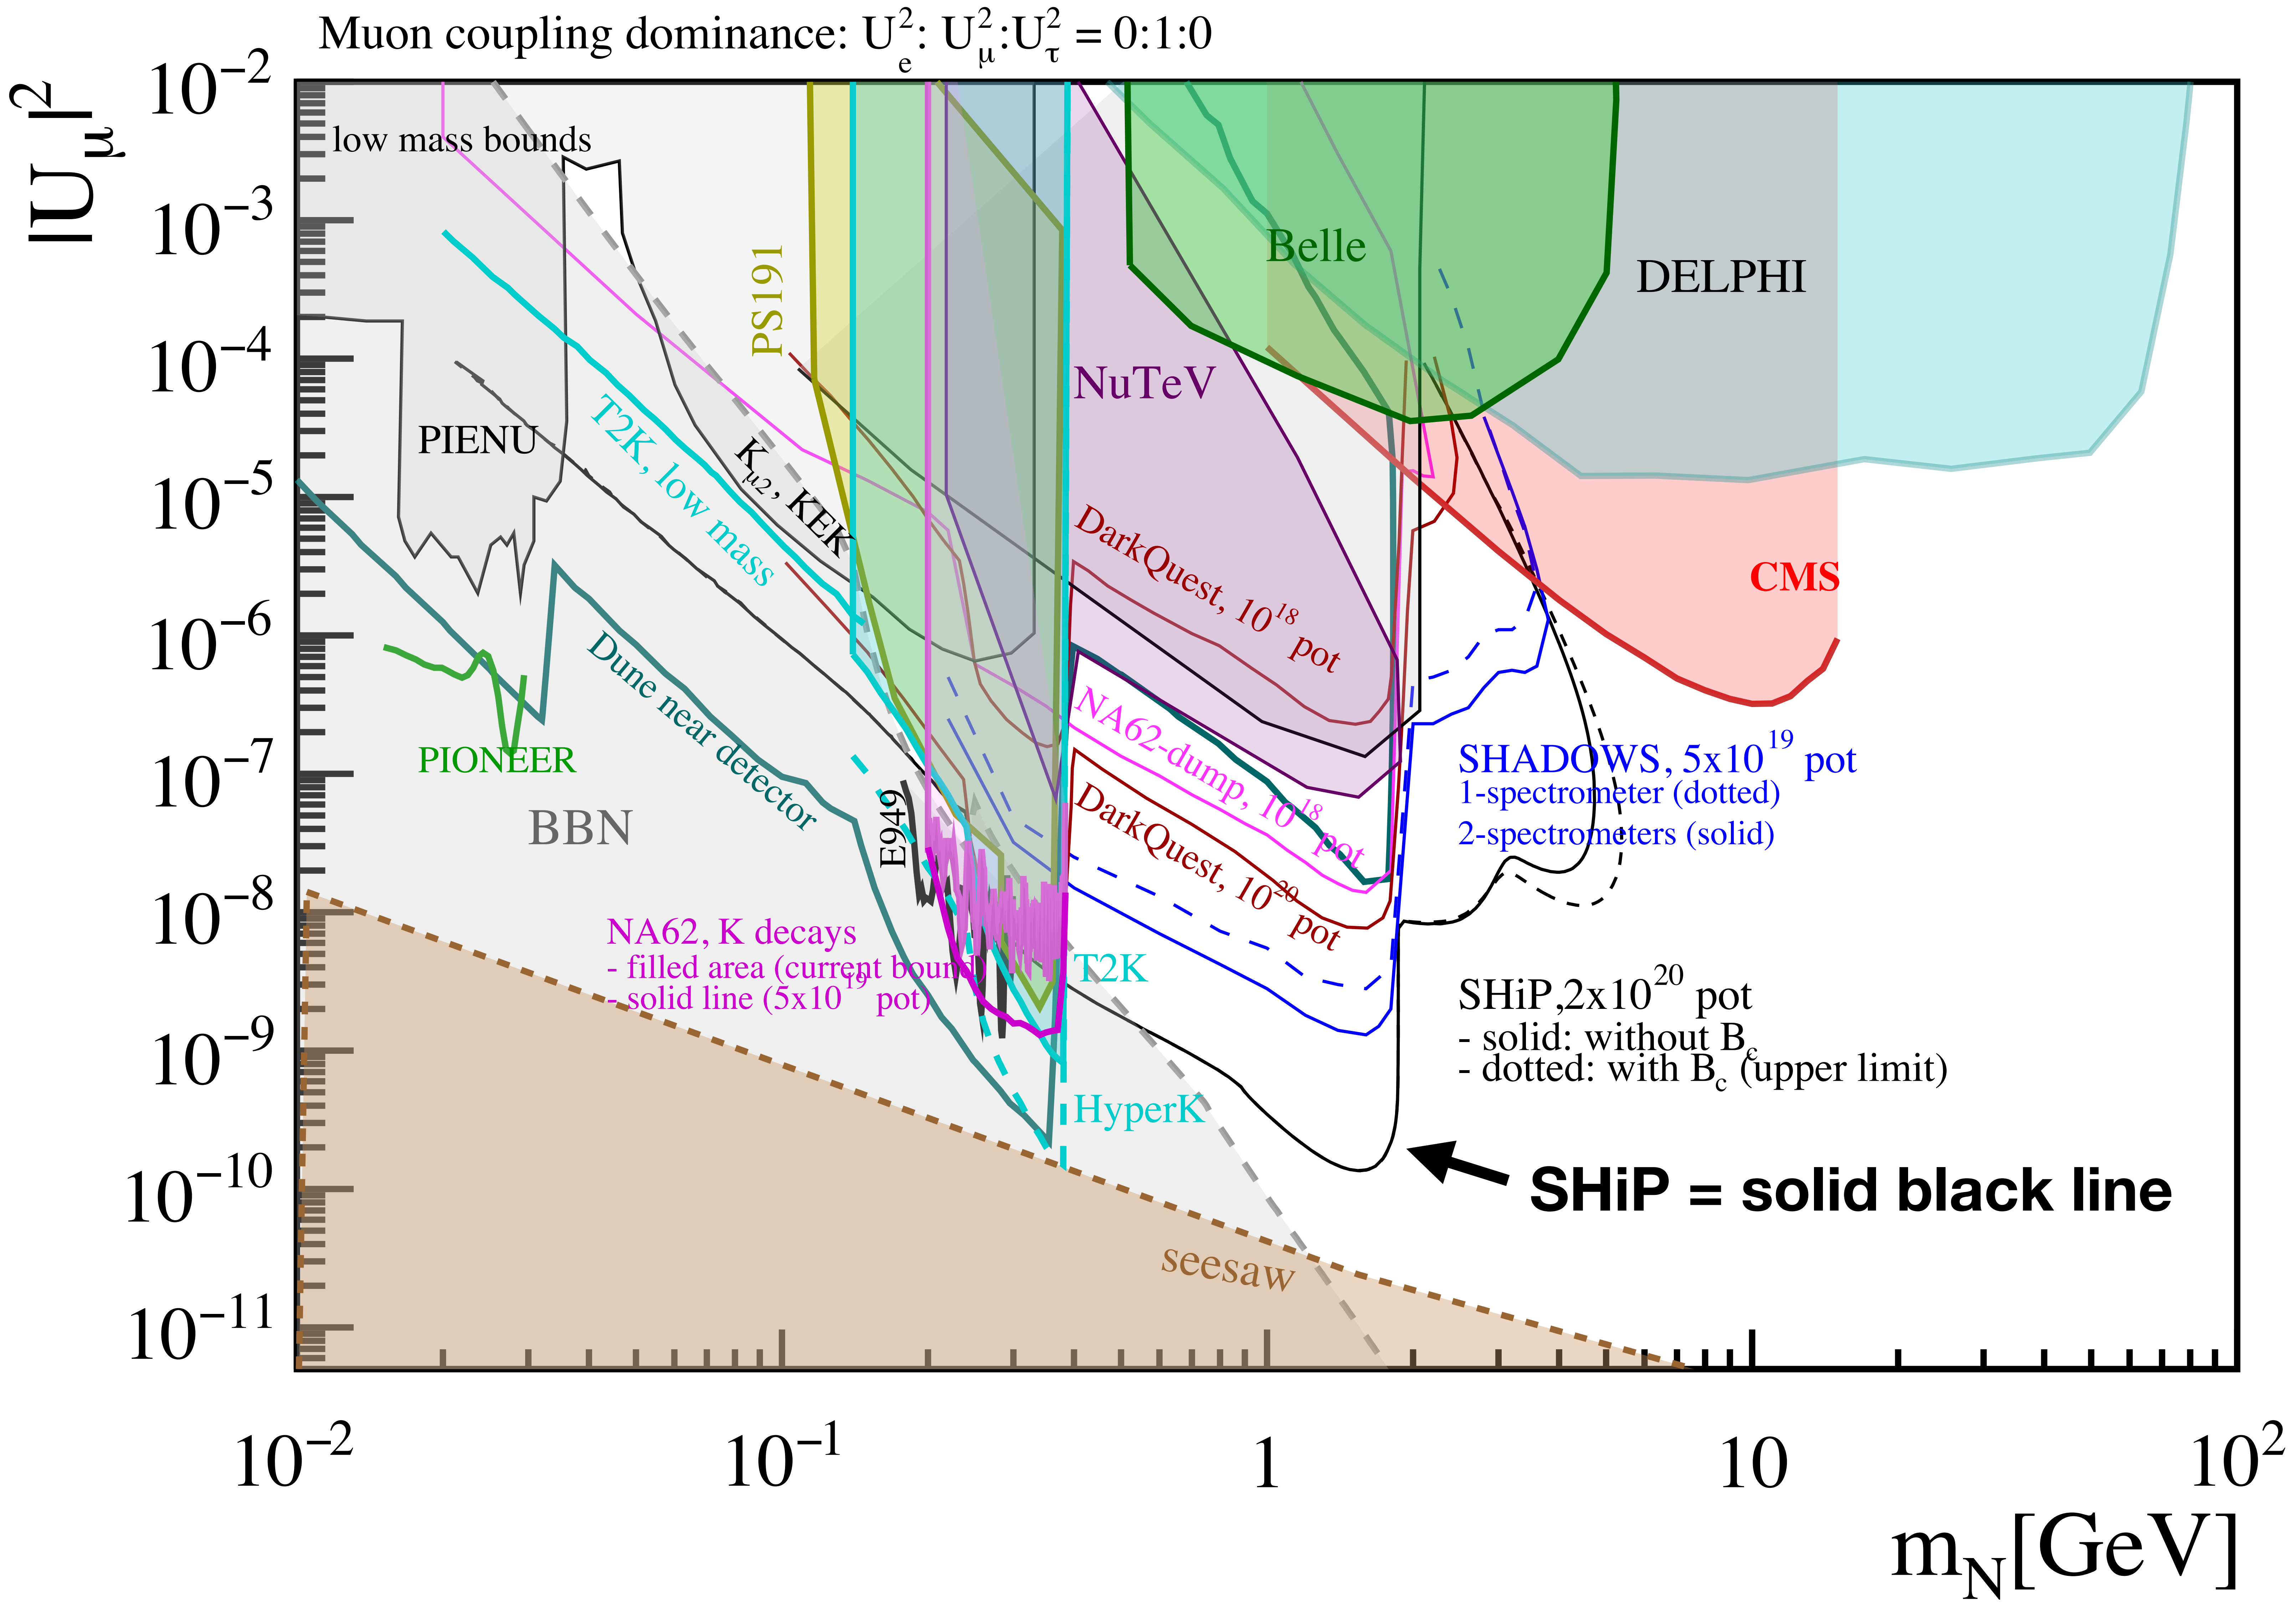

Plot from Snowmass paper [2203.08039] (45 authors)

SHiP sensitivity to HNLs (tau mixing)

[SHiP: 1811.00930]

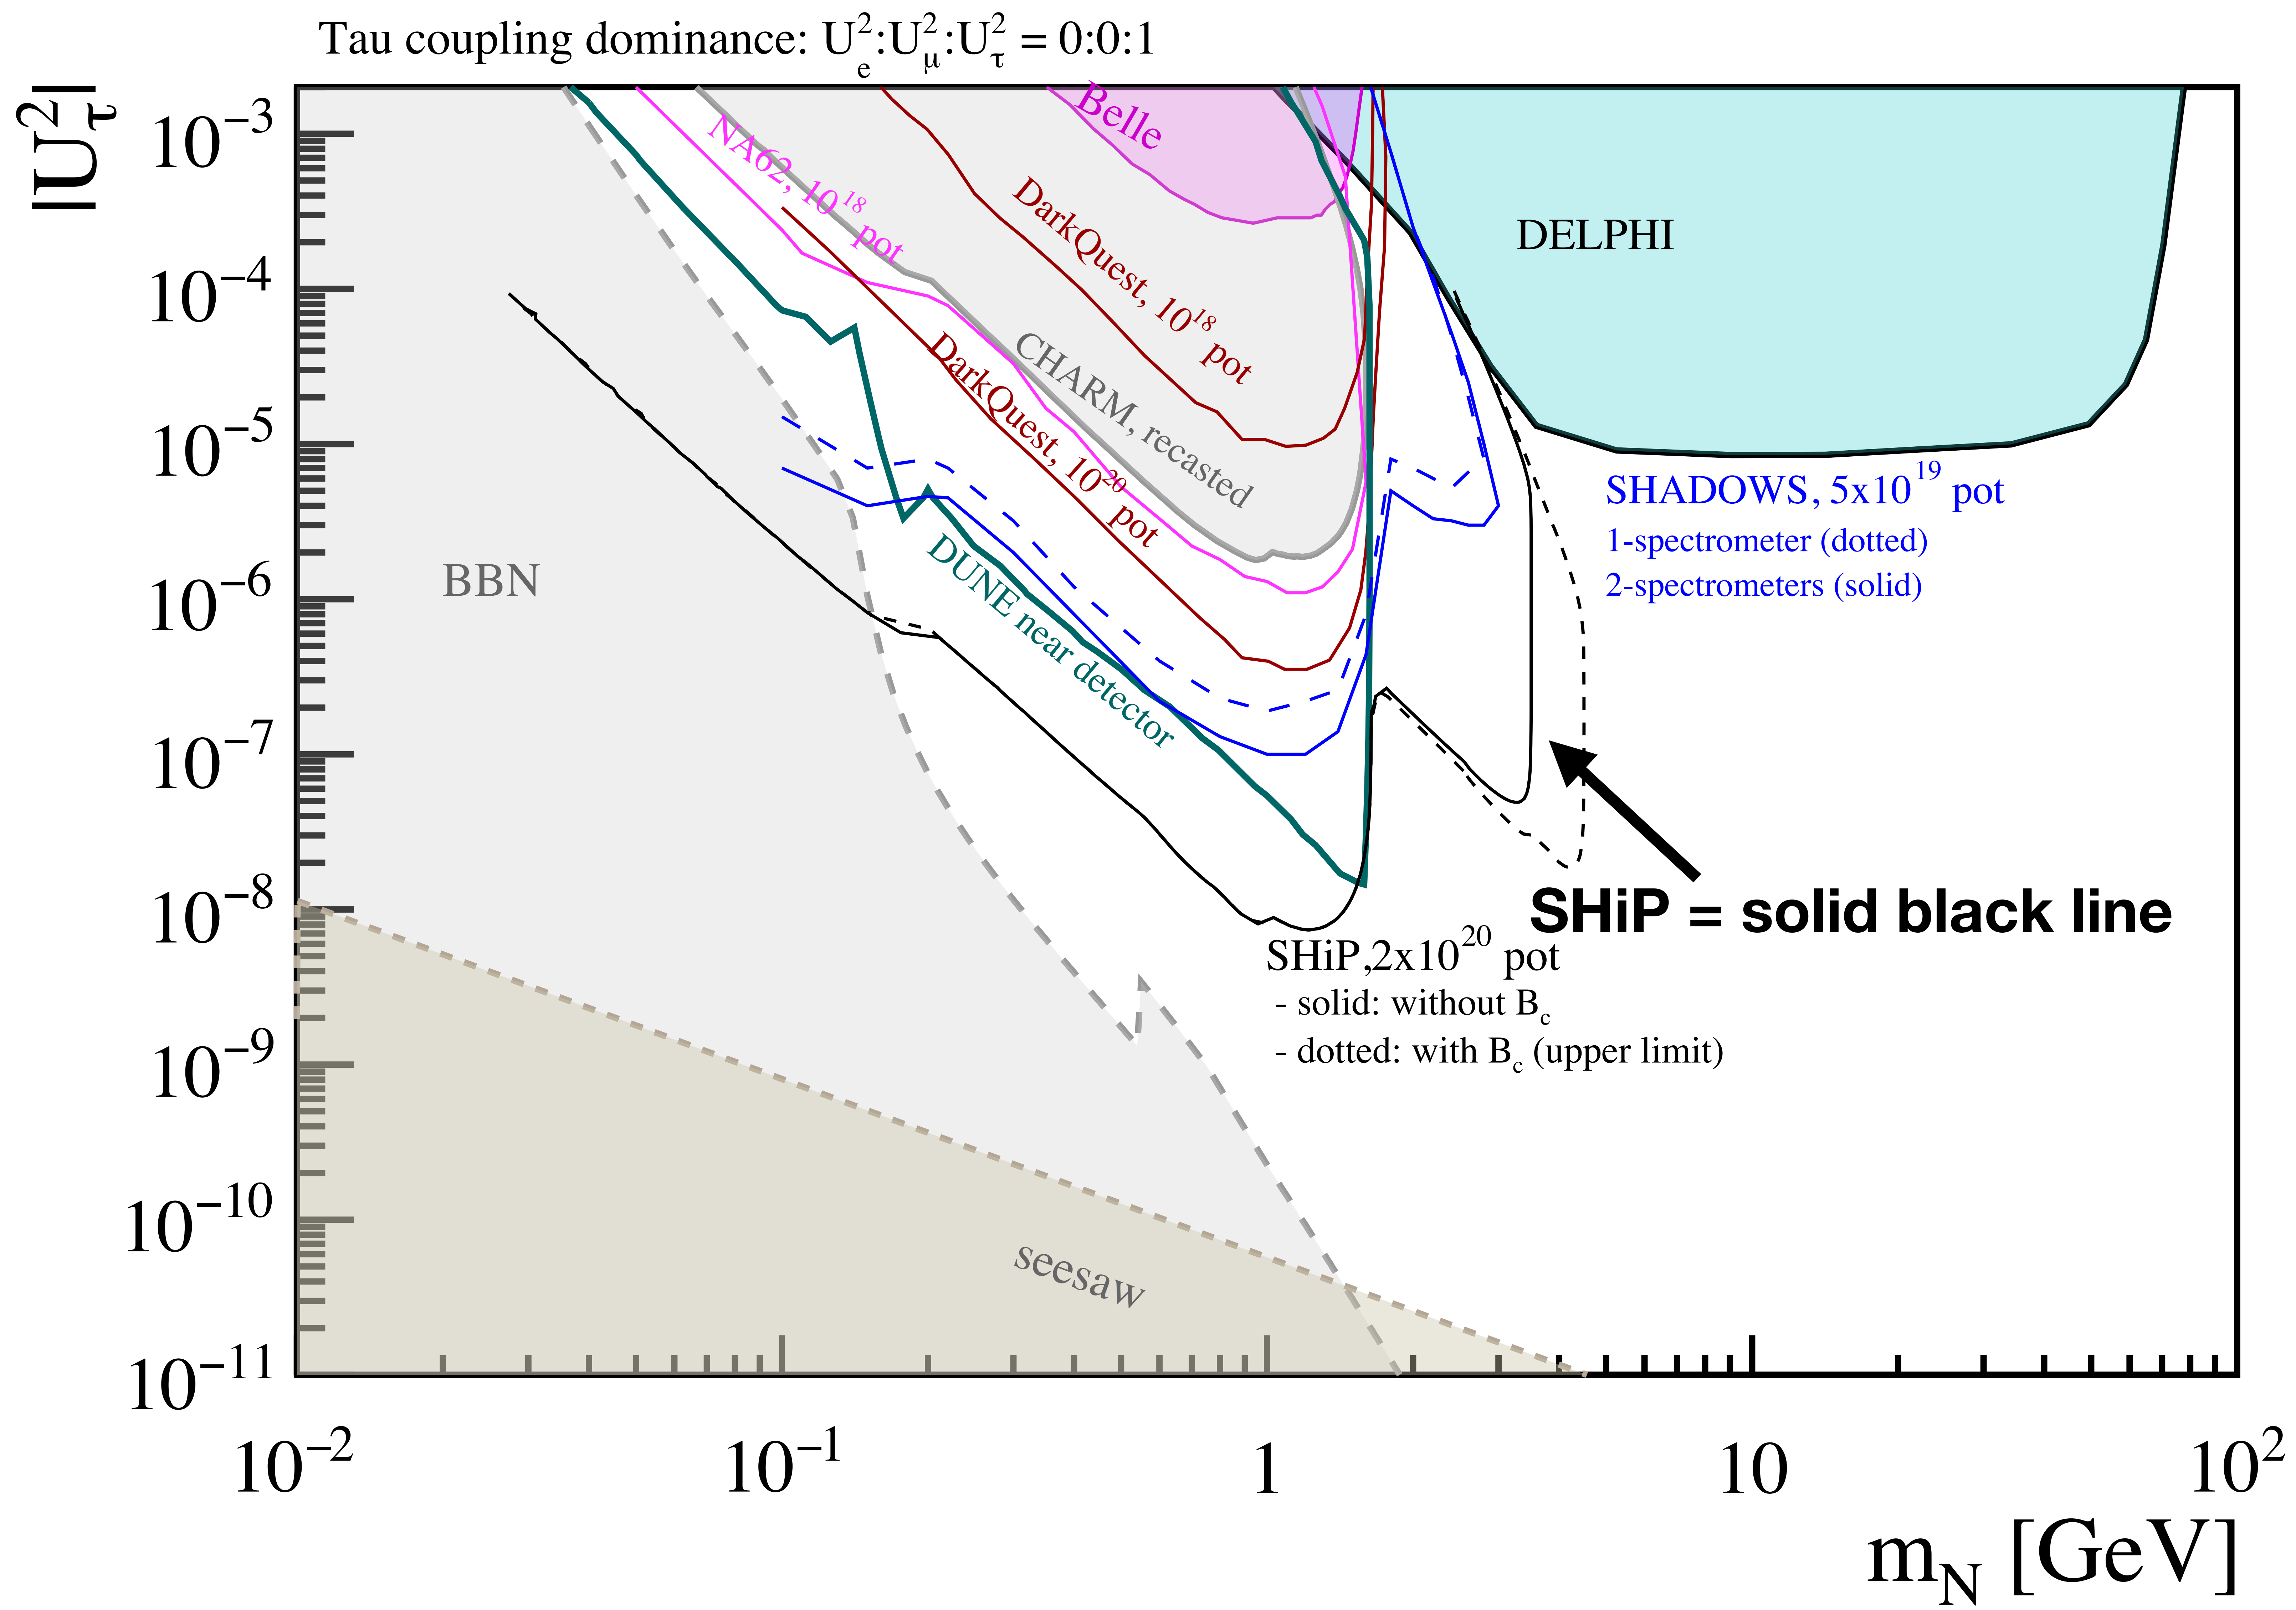

Plot from Snowmass paper [2203.08039] (45 authors)

SHiP sensitivity

Not just HNLs

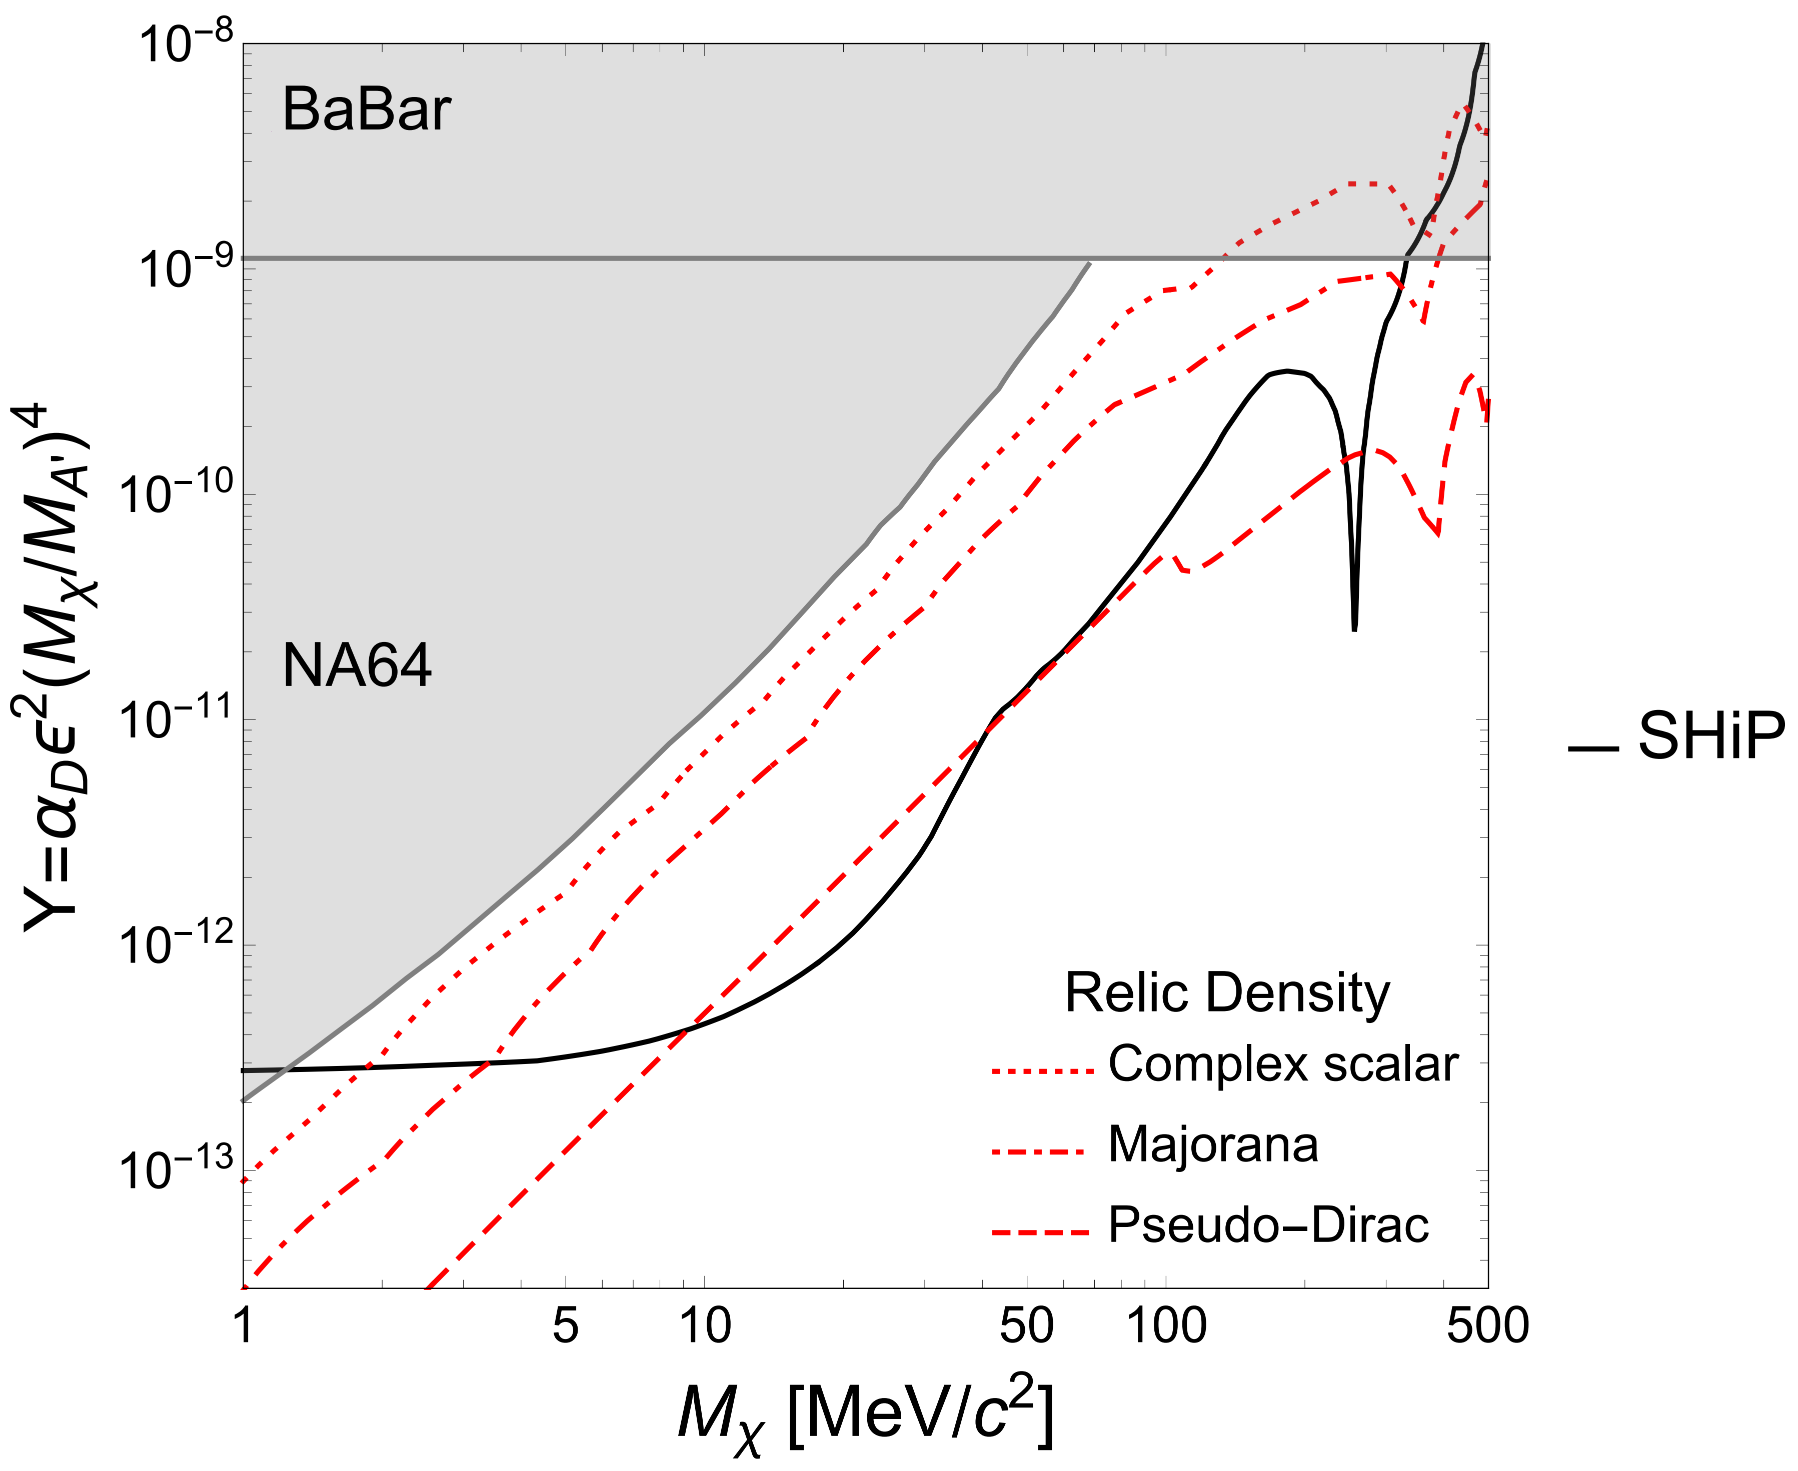

Light dark matter (in scattering detector)
[SHiP: 2010.11057]

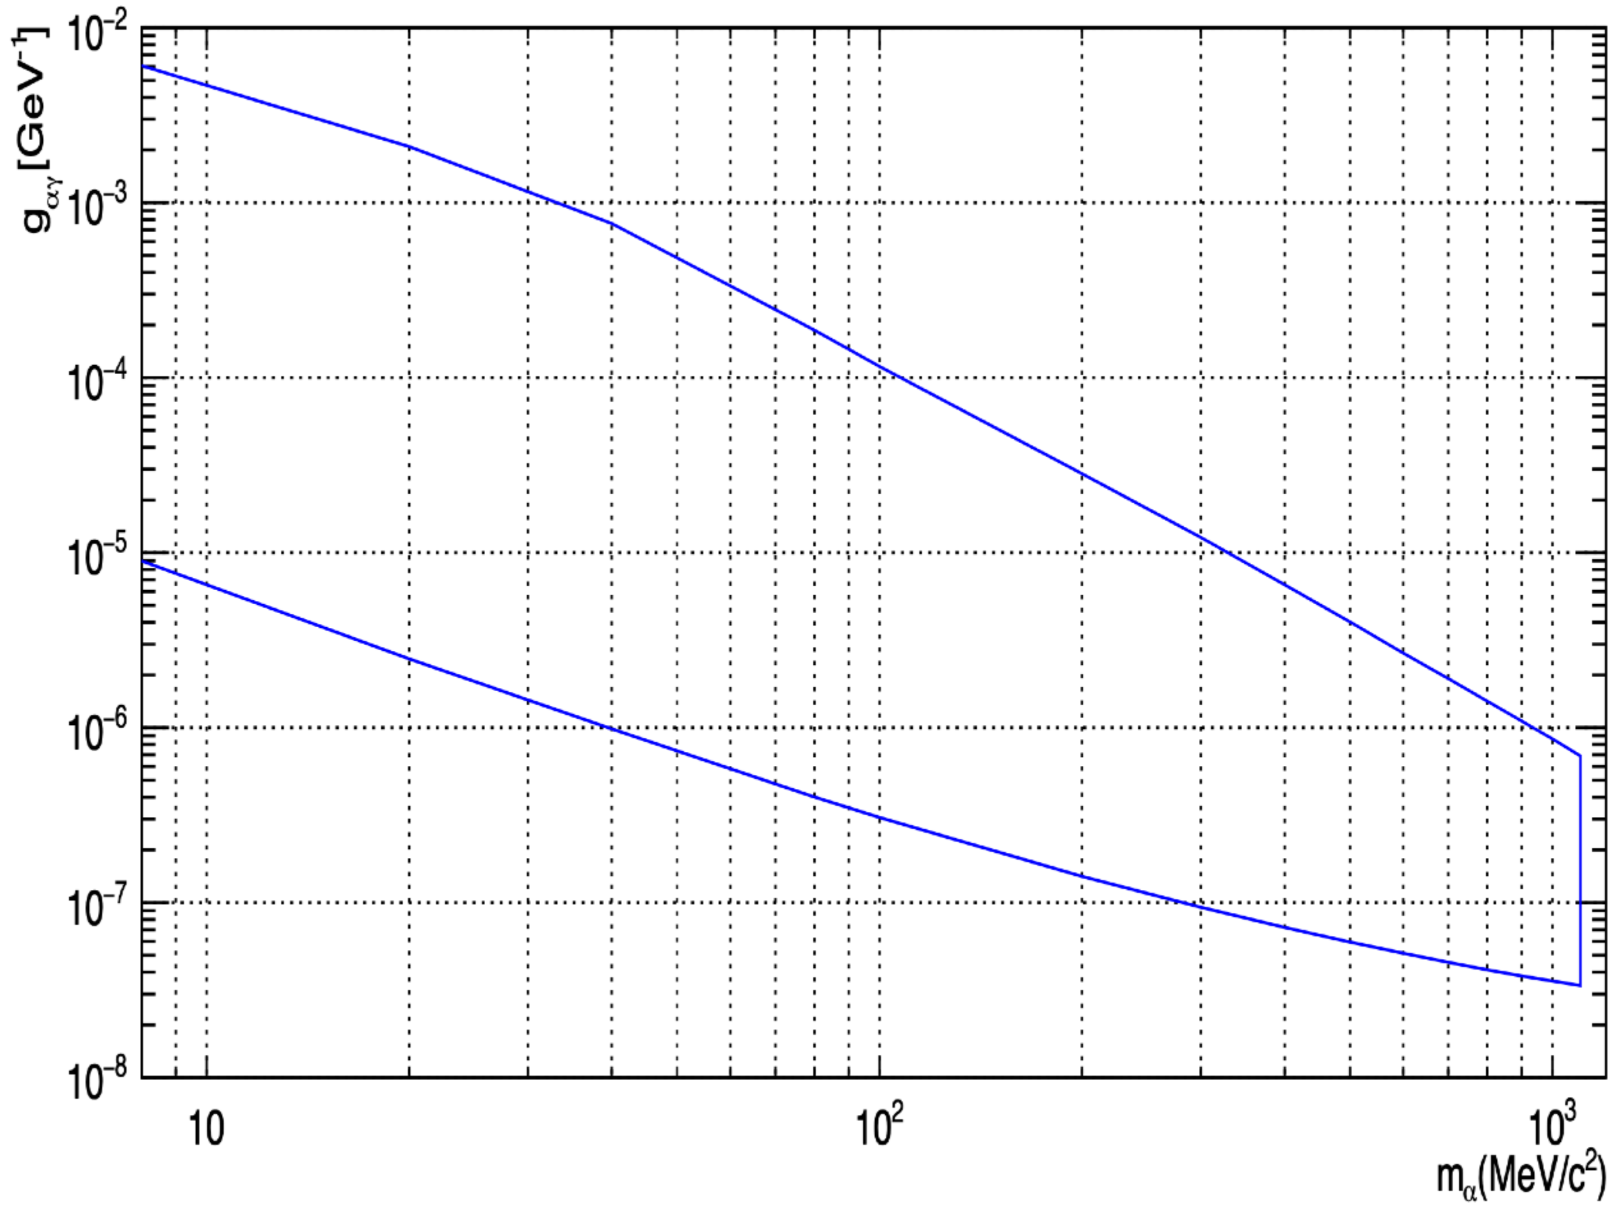

ALP [Akmete: master thesis]

Displaced decays (in vacuum vessel)

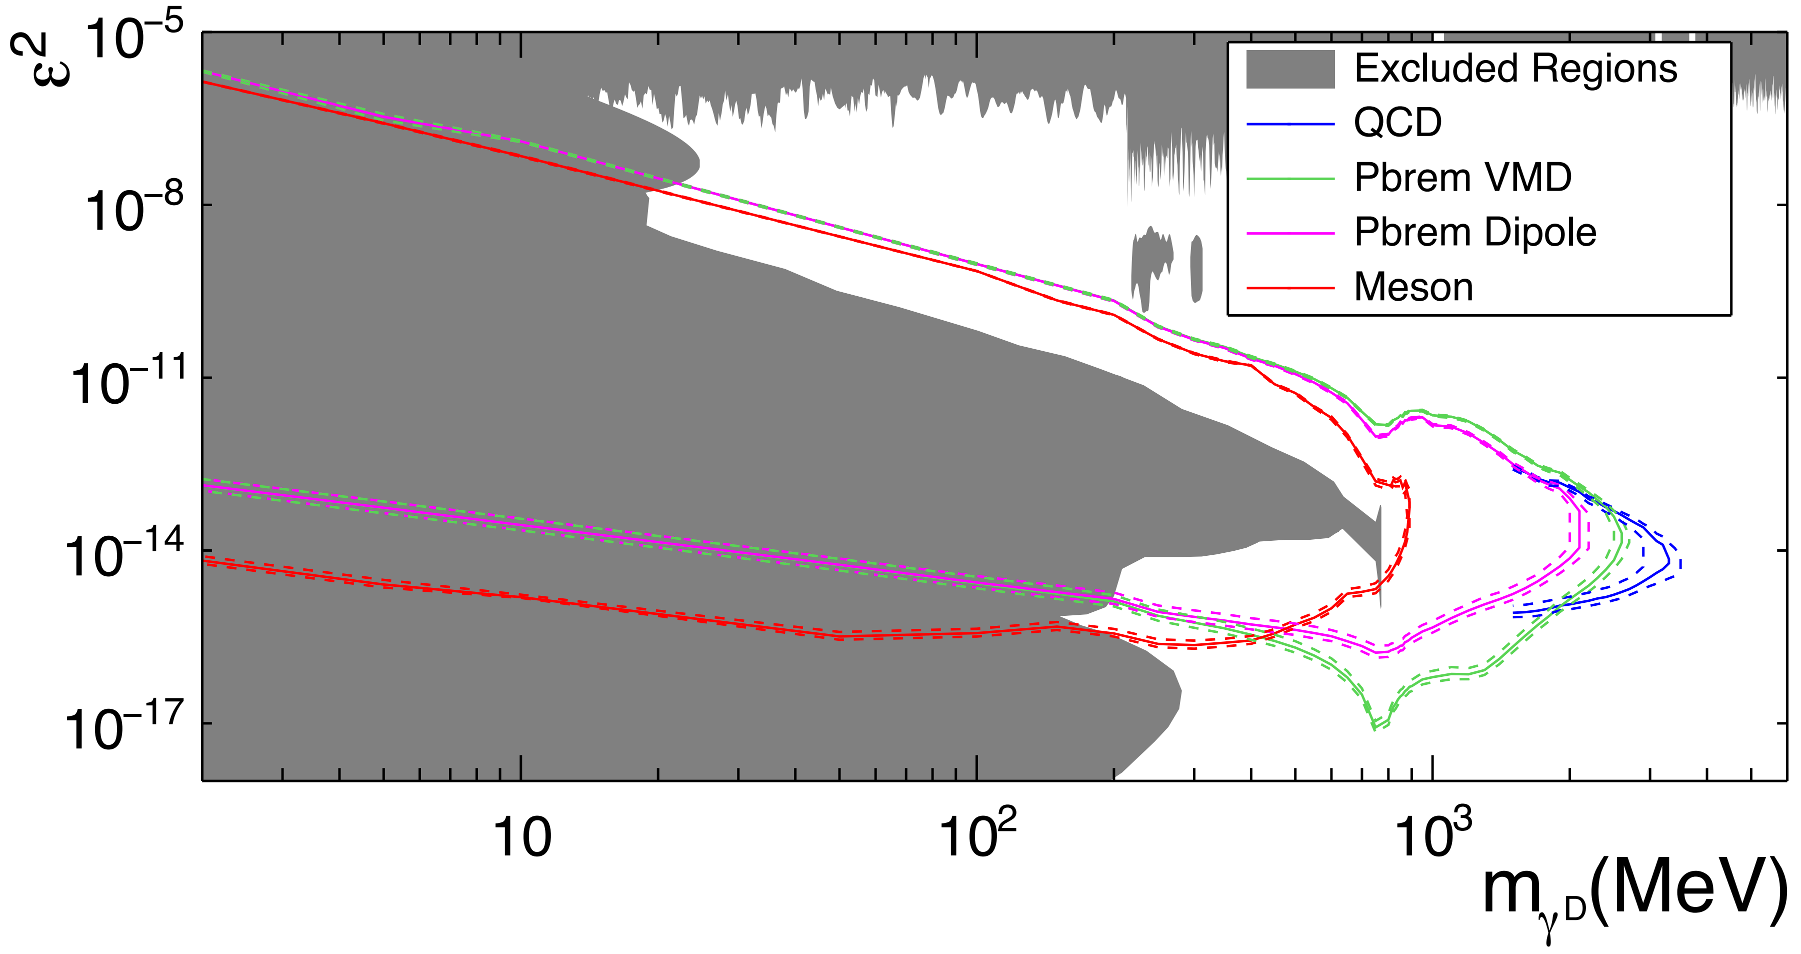

Dark photon [SHiP: 2011.05115]

SHiP sensitivity to HNL

For arbitrary flavour mixing patterns!

- Thanks to its zero-background environment, the 90% CL sensitivity of SHiP can be expressed as Number of events > 2.3 (or 3 events for 95% CL).
  (Here "sensitivity" = expected exclusion limit if no events are observed)

- For a fixed HNL mass, the number of events is proportional to:

$$\sum_{\text{production channels}} \underbrace{\mathrm{Br}(\text{meson} \to N + X \ldots)}_{\text{Linear in } |\Theta_\alpha|^2} \times \underbrace{P(\text{decay in detector})}_{\cong L_{\text{det}}/\gamma_N \tau_N \text{ if } \gamma_N \tau_N \gg L_{\text{det}}} \times \sum_{\text{decay channels}} \underbrace{\mathrm{Br}(N \to \text{SM final state})}_{\text{Linear in } |\Theta_\alpha|^2}$$

  Because each individual production or decay channel is proportional to $|\Theta_\alpha|^2$ for *some* flavour α.

- For long-lived HNLs, the number of events is bilinear in the mixing angles!

The sensitivity matrix

- The lower limit can be well approximated as:

  $|\Theta_\alpha|^2 N_{\alpha\beta} |\Theta_\beta|^2 > 2.3$ where $N_{\alpha\beta}$ is the sensitivity matrix.

  This allows drawing the 90% CL limit for any ratio $|\Theta_e|^2 : |\Theta_\mu|^2 : |\Theta_\tau|^2$.

- $N_{\alpha\beta}$ is the expected number of events if the HNL were produced only through flavour α and decayed only through flavour β, with the mixing angles set to 1.

- For HNLs with a lifetime comparable to the size of the detector, an exponential correction can be applied.

- A Mathematica notebook uploaded along with the paper allows you to compute the limits *for your favourite choice of parameters*! [SHiP: 1811.00930]

Probing HNL properties at SHiP

The Dirac / Majorana nature of HNLs

- Easiest way: measuring the change in lepton number: $|\Delta L| = 2 \Rightarrow$ Majorana
- This is not always possible, e.g. if a charged lepton or neutrino escapes
- At SHiP, the charged lepton produced with the HNL cannot be observed

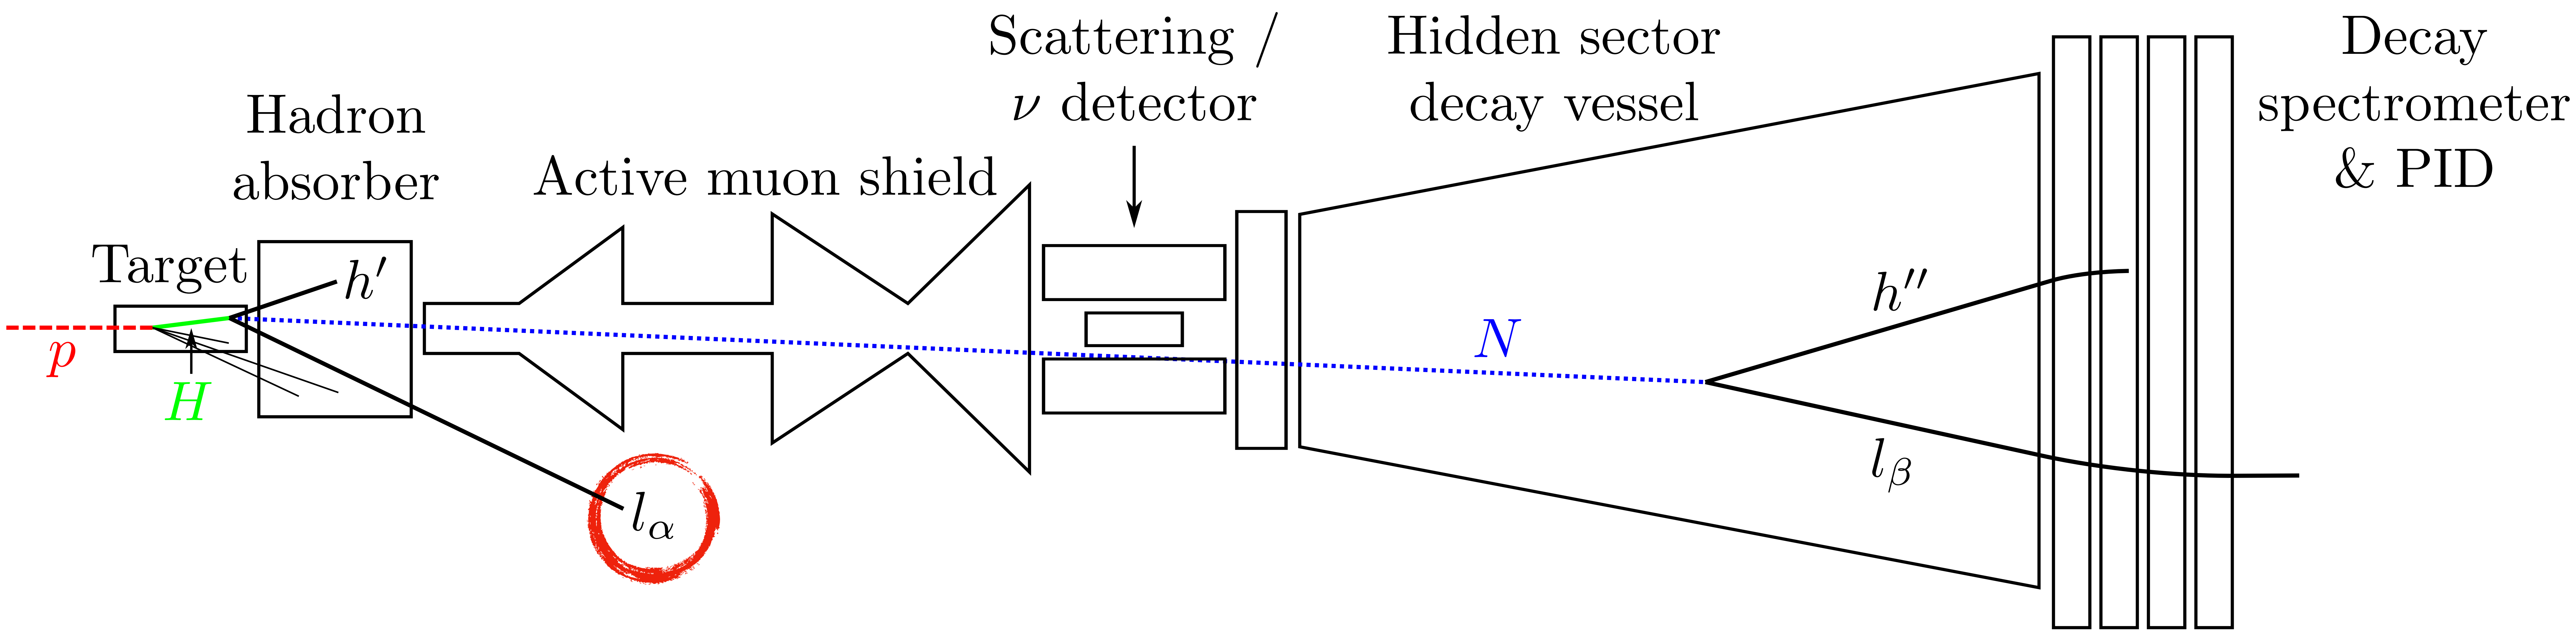

Dirac vs. Majorana at SHiP

Solution: spin correlations

- HNLs carry spin 1/2, which is affected by the production process
- Spin affects the decay kinematics, which carry information about LNC/LNV

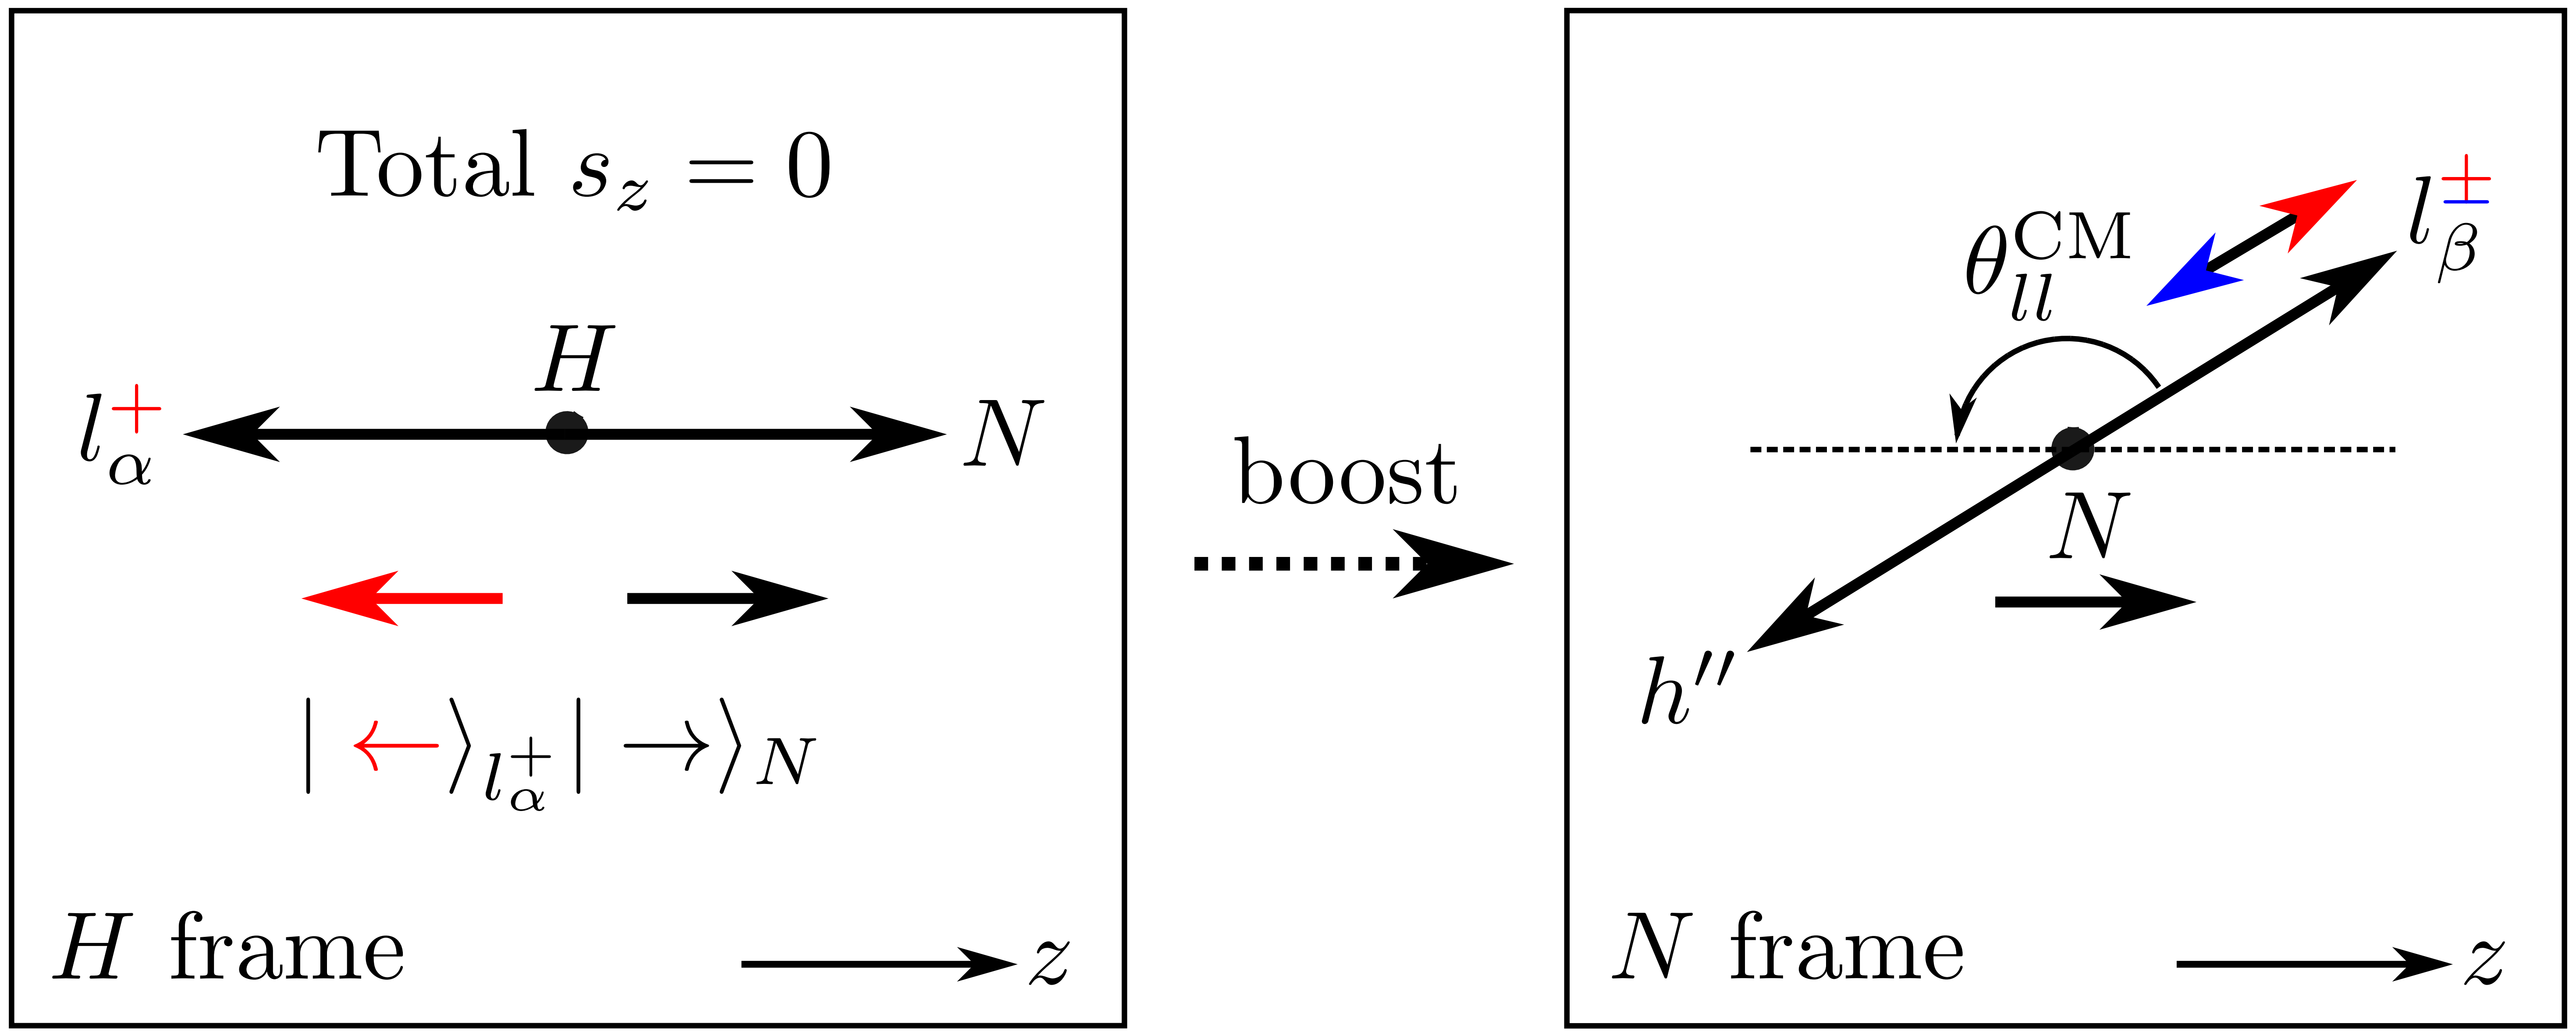

$$l_\beta^+ \langle \nearrow | \rightarrow \rangle_N = \sin\left(\frac{\theta_{ll}^{\mathrm{CM}}}{2}\right)$$

$$l_\beta^- \langle \swarrow | \rightarrow \rangle_N = \cos\left(\frac{\theta_{ll}^{\mathrm{CM}}}{2}\right)$$

$$\mathcal{P} \propto 1 \mp \cos\left(\theta_{ll}^{\mathrm{CM}}\right)$$

Dirac vs. Majorana at SHiP

Spin correlations in the lab frame

- We consider the decay $N \to \pi^{\mp} l^{\pm}$
- Clean cos/sin dependence in the frame of the parent meson
- Random boost to the HNL frame introduces some smearing (simulated using a toy Monte-Carlo with all relevant processes)
- However LNC/LNV can still be distinguished in the lab frame
- With enough events, we can distinguish Dirac (LNC) from Majorana (LNC+LNV)

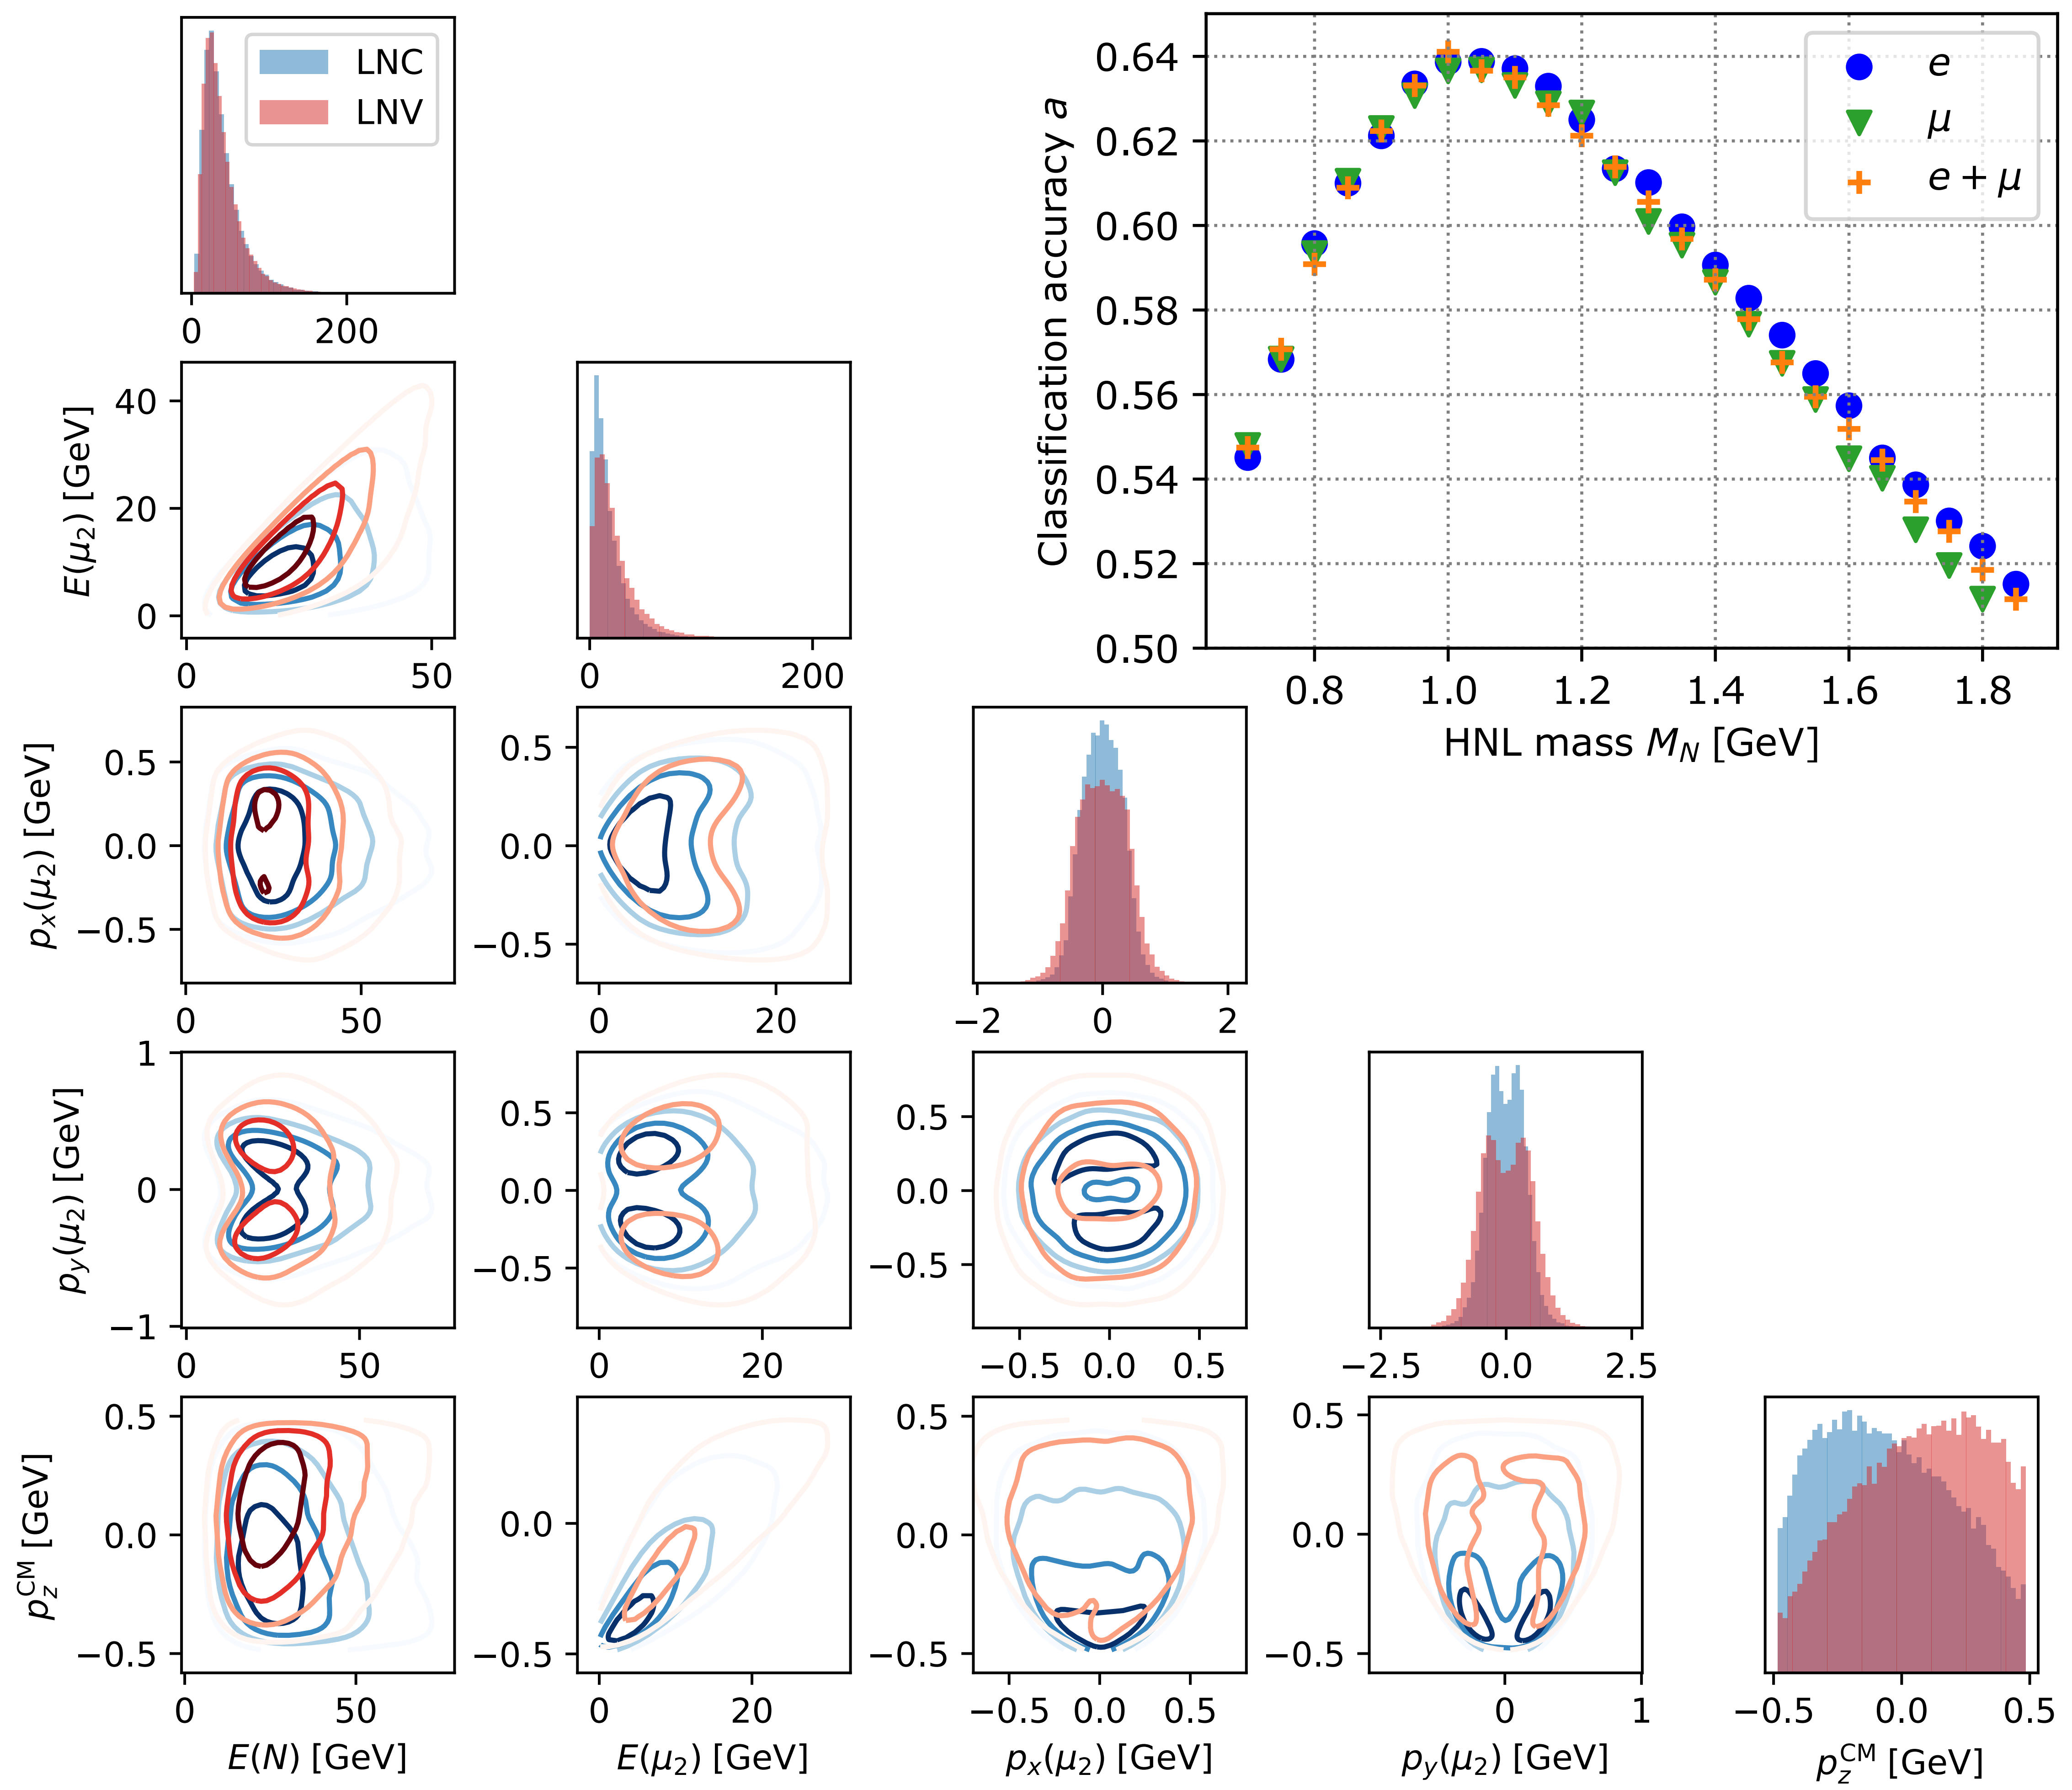

Dirac vs. Majorana at SHiP

Sensitivity [Tastet, Timiryasov: 1912.05520]

- Compute the number of events needed to exclude the "Dirac" or "LNC-only" hypothesis (i.e. discover that HNLs are Majorana)
- Convert it to a mixing angle $|\Theta|^2$
- There exist a currently-unconstrained region of parameter space where SHiP could both discover HNLs and show they are Majorana
- Similar studies have been done at colliders
  [Arbelaéz, Dib, Schmidt, Vasquez: 1712.08704]
  [Dib, Kim, Wang: 1703.01934]
  [Hernández, Jones-Pérez, Suarez-Navarro: 1810.07210]
  and more recently at DUNE
  [de Gouvêa, Fox, Kayser, Kelly: 2109.10358]

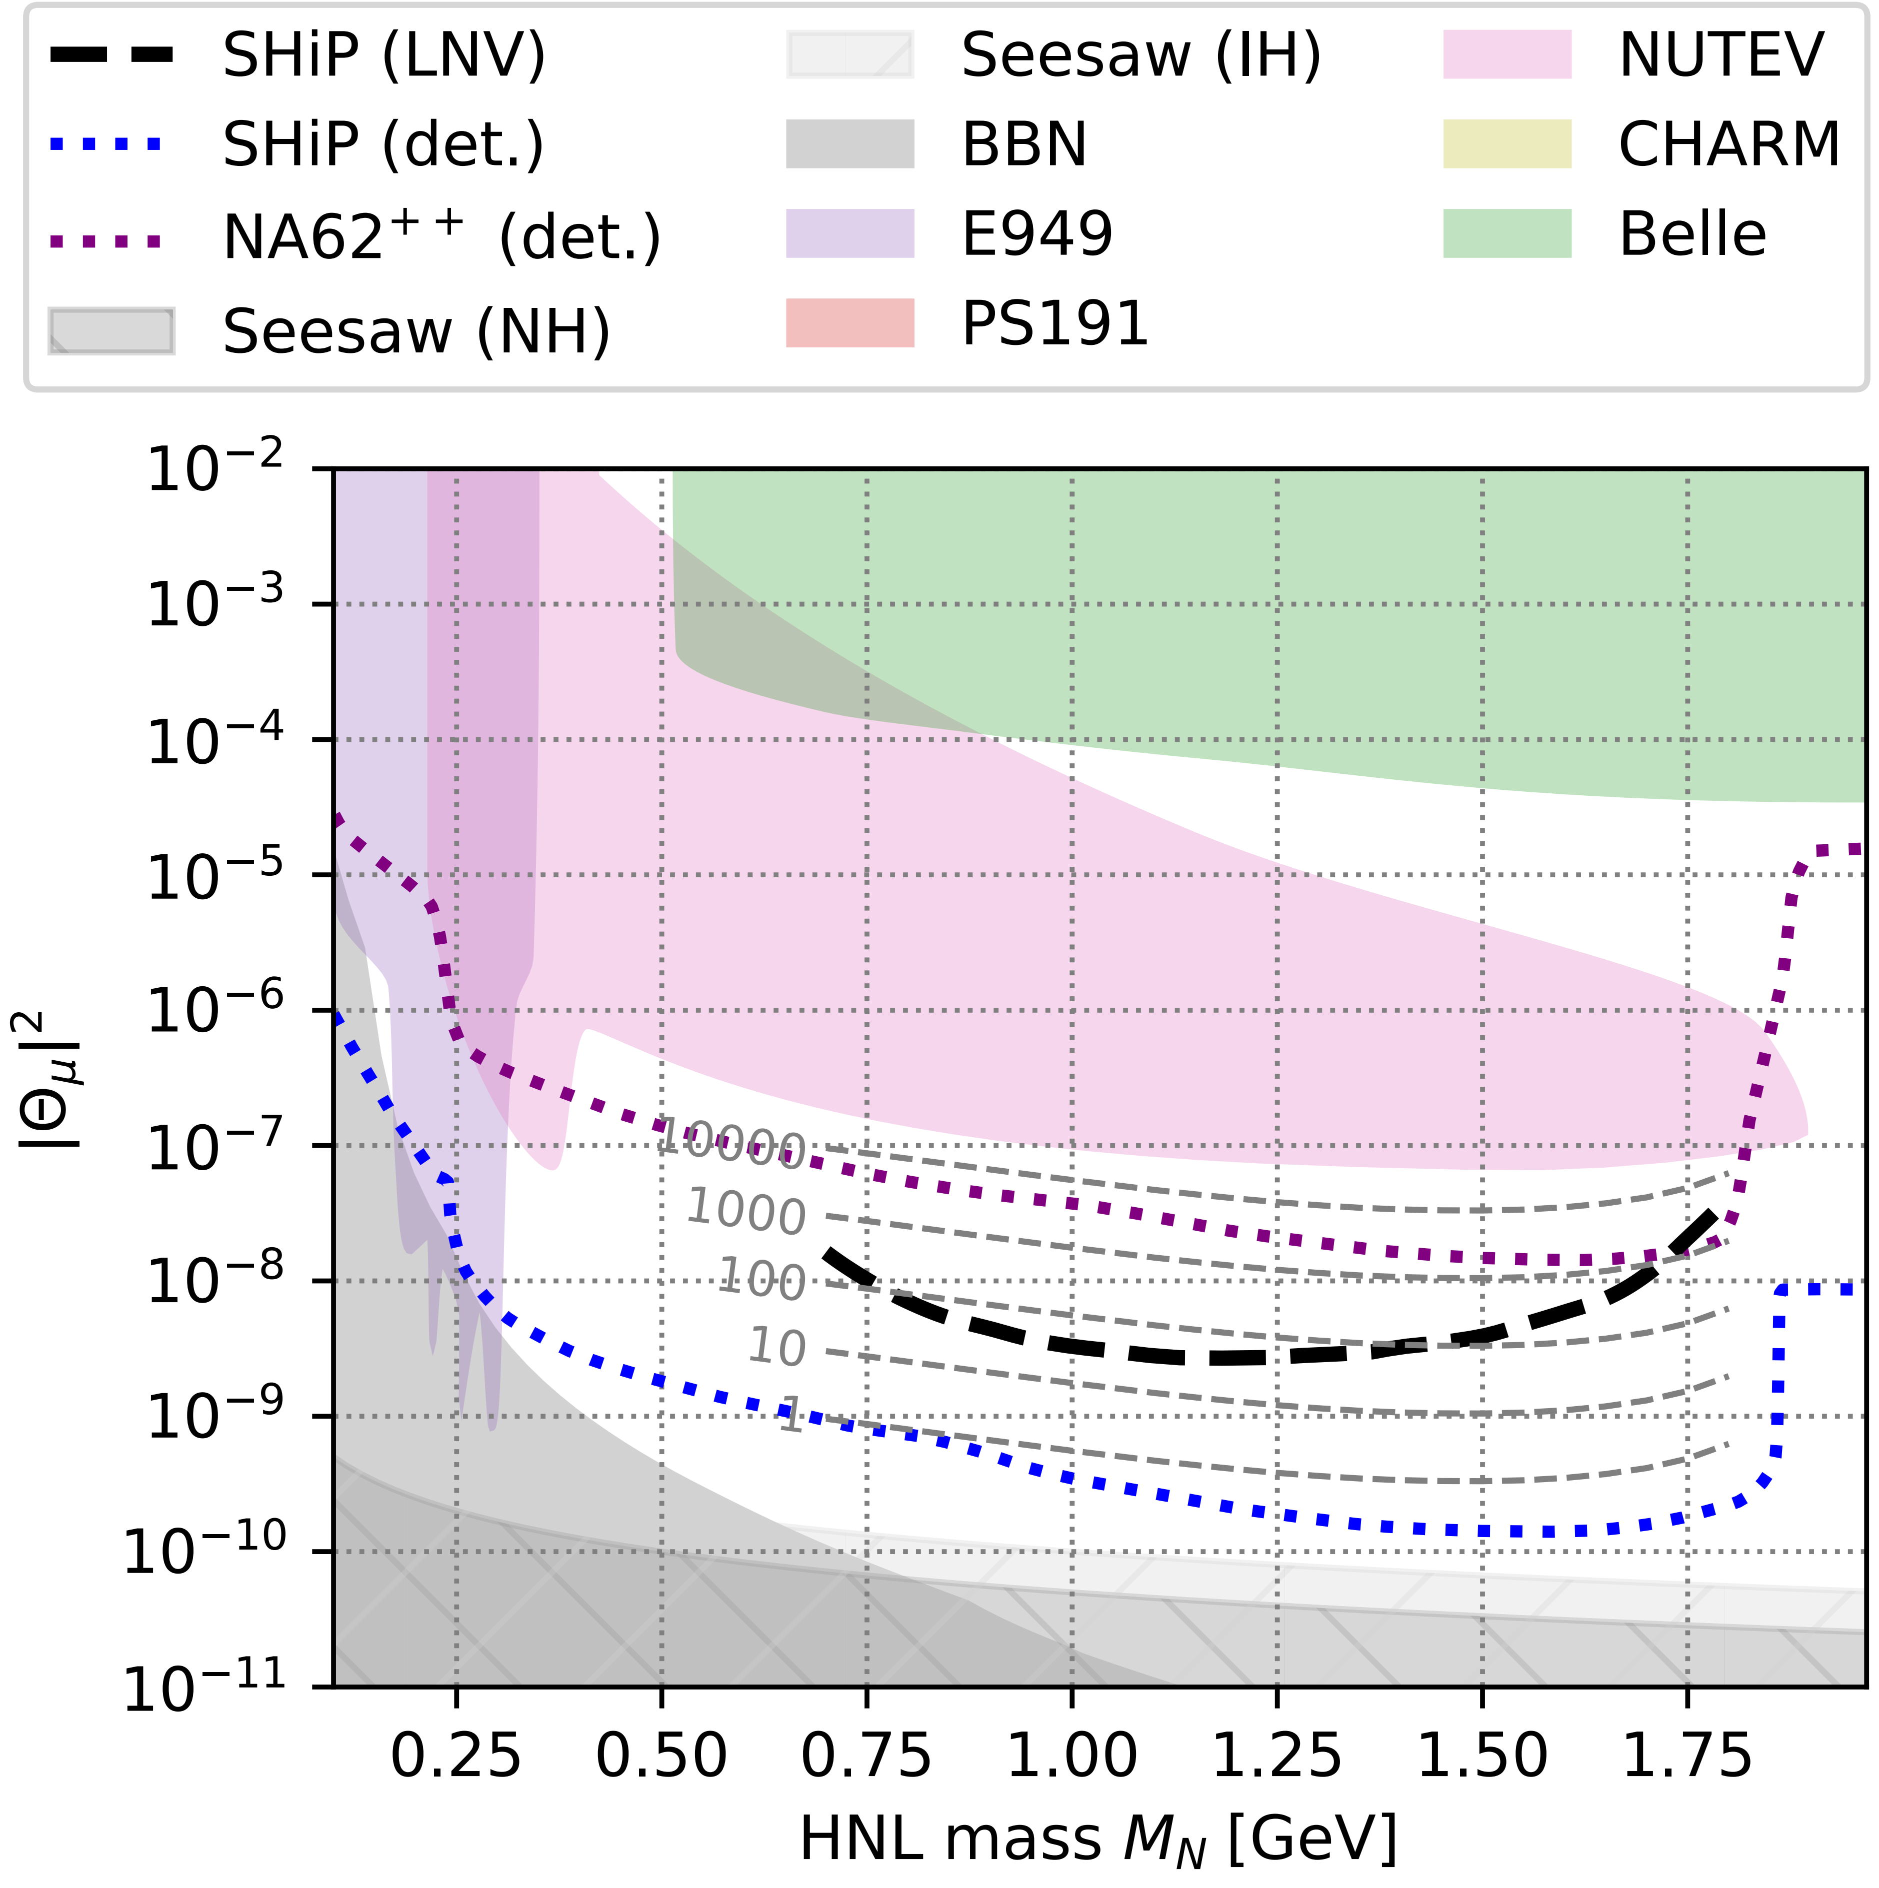

HNL oscillations at SHiP

What if $\delta M \sim \pi\Gamma$?

- Classify each event to obtain $\mathscr{P}$(LNV)
- Assume we have measured the HNL mass
- We roughly know its production vertex (within the target)
- We can precisely measure its decay vertex and its momentum

$\Longrightarrow$ We can compute its proper lifetime τ

- Bin events in τ, weight them by $\mathscr{P}$(LNV) and subtract the mean

This allows resolving the oscillation pattern!

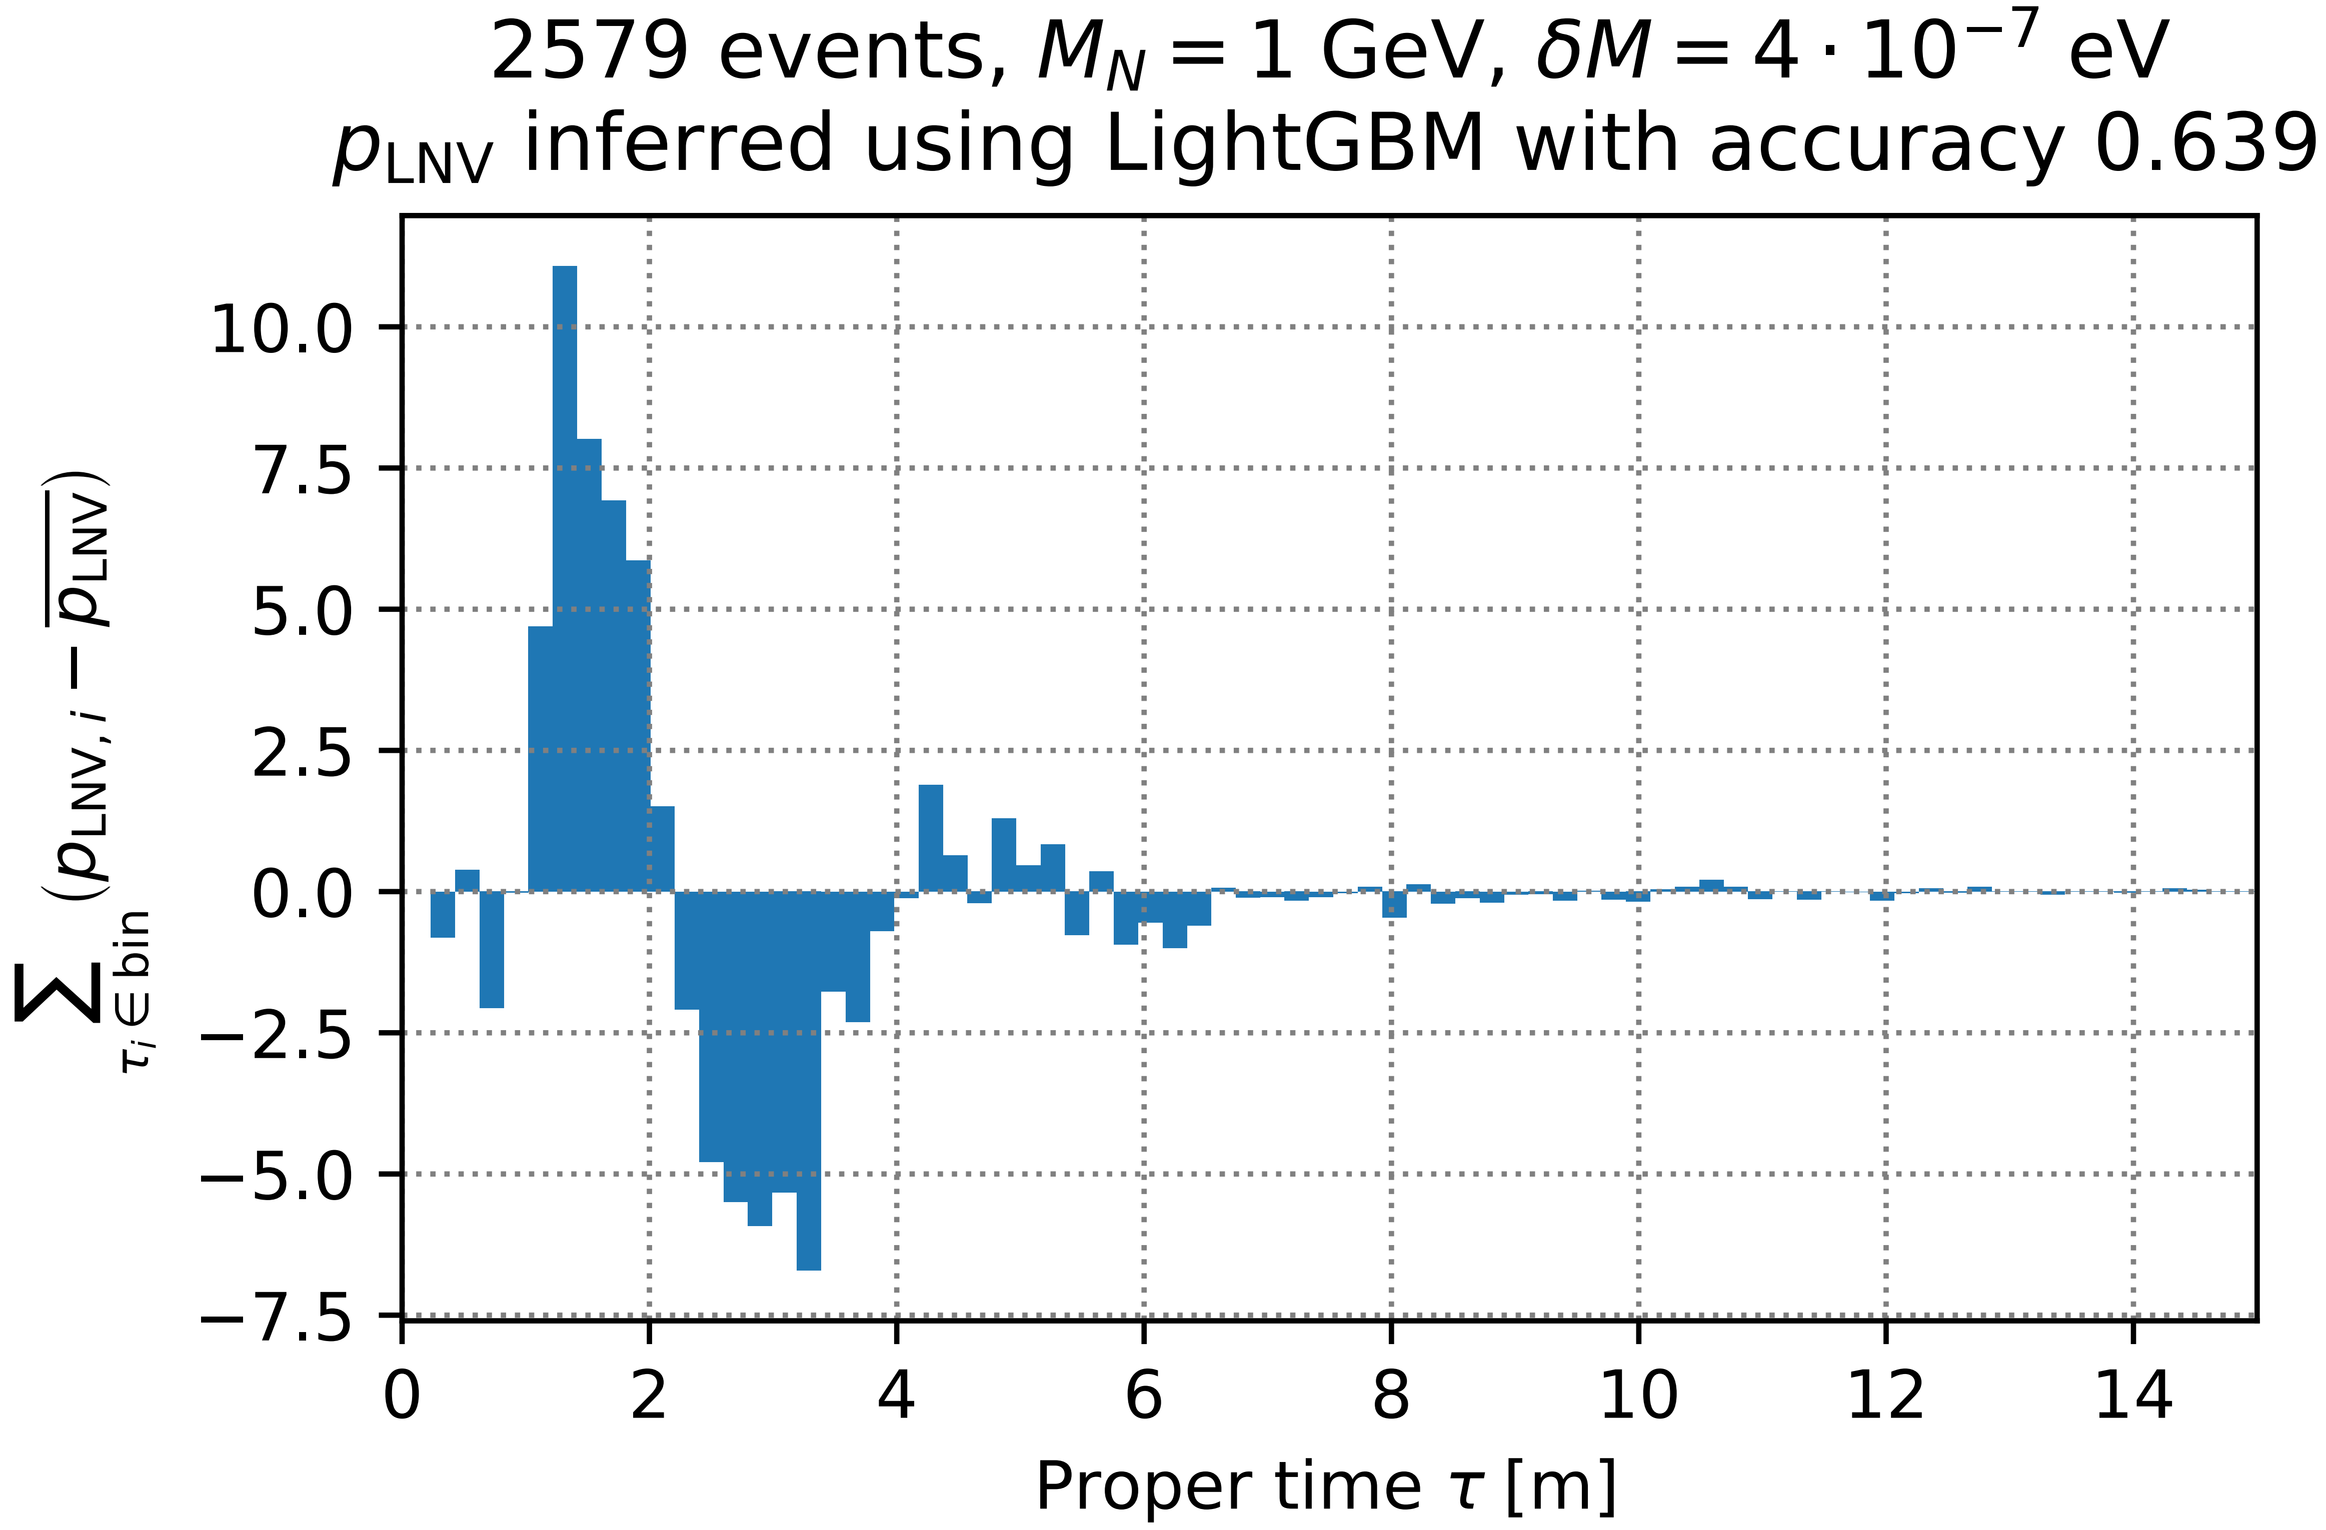

SHiP timeline

- Initially proposed in 2013, to be built in the North Area along with the Beam Dump Facility.
[SHiP: 1310.1762]

- Technical proposal and physics paper submitted to the SPS committee in 2015.
[SHiP: 1504.04956], [SHiP: 1504.04855]

- Comprehensive design study performed in 2019 [SHiP: CDS 2654870, CDS 2704147]

- European Strategy Update did not mention SHiP 😢

- Search for new locations in 2021-2022.

- Located the ECN3 hall, in use by the NA62 experiment until the end of Run 3.

- Possibility to build SHiP there after NA62 finishes. [cf. Alexey's talk last week]
Competes with the HIKE proposal (NA62 with 4× the statistics of the next run).

SHiP timeline

- Initially proposed in 2013, to be built in the North Area along with the Beam Dump Facility.
  [SHiP: 1310.1762]
- Technical proposal and physics paper submitted to the SPS committee in 2015.
  [SHiP: 1504.04956], [SHiP: 1504.04855]
- Comprehensive design study performed in 2019 [SHiP: CDS 2654870, CDS 2704147]
- European Strategy Update did not mention SHiP 😢
- Search for new locations in 2021-2022.
- Located the ECN3 hall, in use by the NA62 experiment until the end of Run 3.
- Possibility to build SHiP there after NA62 finishes. [cf. Alexey's talk last week]
  Competes with the HIKE proposal (NA62 with 4× the statistics of the next run).

Input needed from the community! A decision is expected by early next year.

Searching for HNLs at ATLAS

I. Prompt trilepton search

[ATLAS: 1905.09787]

$36.1\ \mathrm{fb}^{-1}$ at 13 TeV

$M_N \in [5,50]\,\mathrm{GeV}$

EUROPEAN ORGANISATION FOR NUCLEAR RESEARCH (CERN)

ATLAS EXPERIMENT

Submitted to: JHEP

CERN-EP-2019-071
13th November 2019

Search for heavy neutral leptons in decays of W bosons produced in 13 TeV pp collisions using prompt and displaced signatures with the ATLAS detector

The ATLAS Collaboration

The problems of neutrino masses, matter–antimatter asymmetry, and dark matter could be successfully addressed by postulating right-handed neutrinos with Majorana masses below the electroweak scale. In this work, leptonic decays of W bosons extracted from 32.9 fb^{-1} to 36.1 fb^{-1} of 13 TeV proton–proton collisions at the LHC are used to search for heavy neutral leptons (HNLs) that are produced through mixing with muon or electron neutrinos. The search is conducted using the ATLAS detector in both prompt and displaced leptonic decay signatures. The prompt signature requires three leptons produced at the interaction point (either $\mu\mu e$ or $ee\mu$) with a veto on same-flavour opposite-charge topologies. The displaced signature comprises a prompt muon from the W boson decay and the requirement of a dilepton vertex (either $\mu\mu$ or μe) displaced in the transverse plane by 4–300 mm from the interaction point. The search sets constraints on the HNL mixing to muon and electron neutrinos for HNL masses in the range 4.5–50 GeV.

Search signature

- HNLs produced in W boson decays.
- *Promptly* decaying to *leptons*.
  (i.e. the HNL displacement is not resolved)
- Two channels are considered:
  - Electron channel: $e^{\pm}e^{\pm}\mu^{\mp}$
  - Muon channel: $\mu^{\pm}\mu^{\pm}e^{\mp}$
- Final states with two leptons of opposite sign and same flavour (OSSF) are not considered due to background from $Z^{(*)}$.
- Sensitive to both LNC and LNV, but one neutrino *escapes* with the information on the final lepton number!

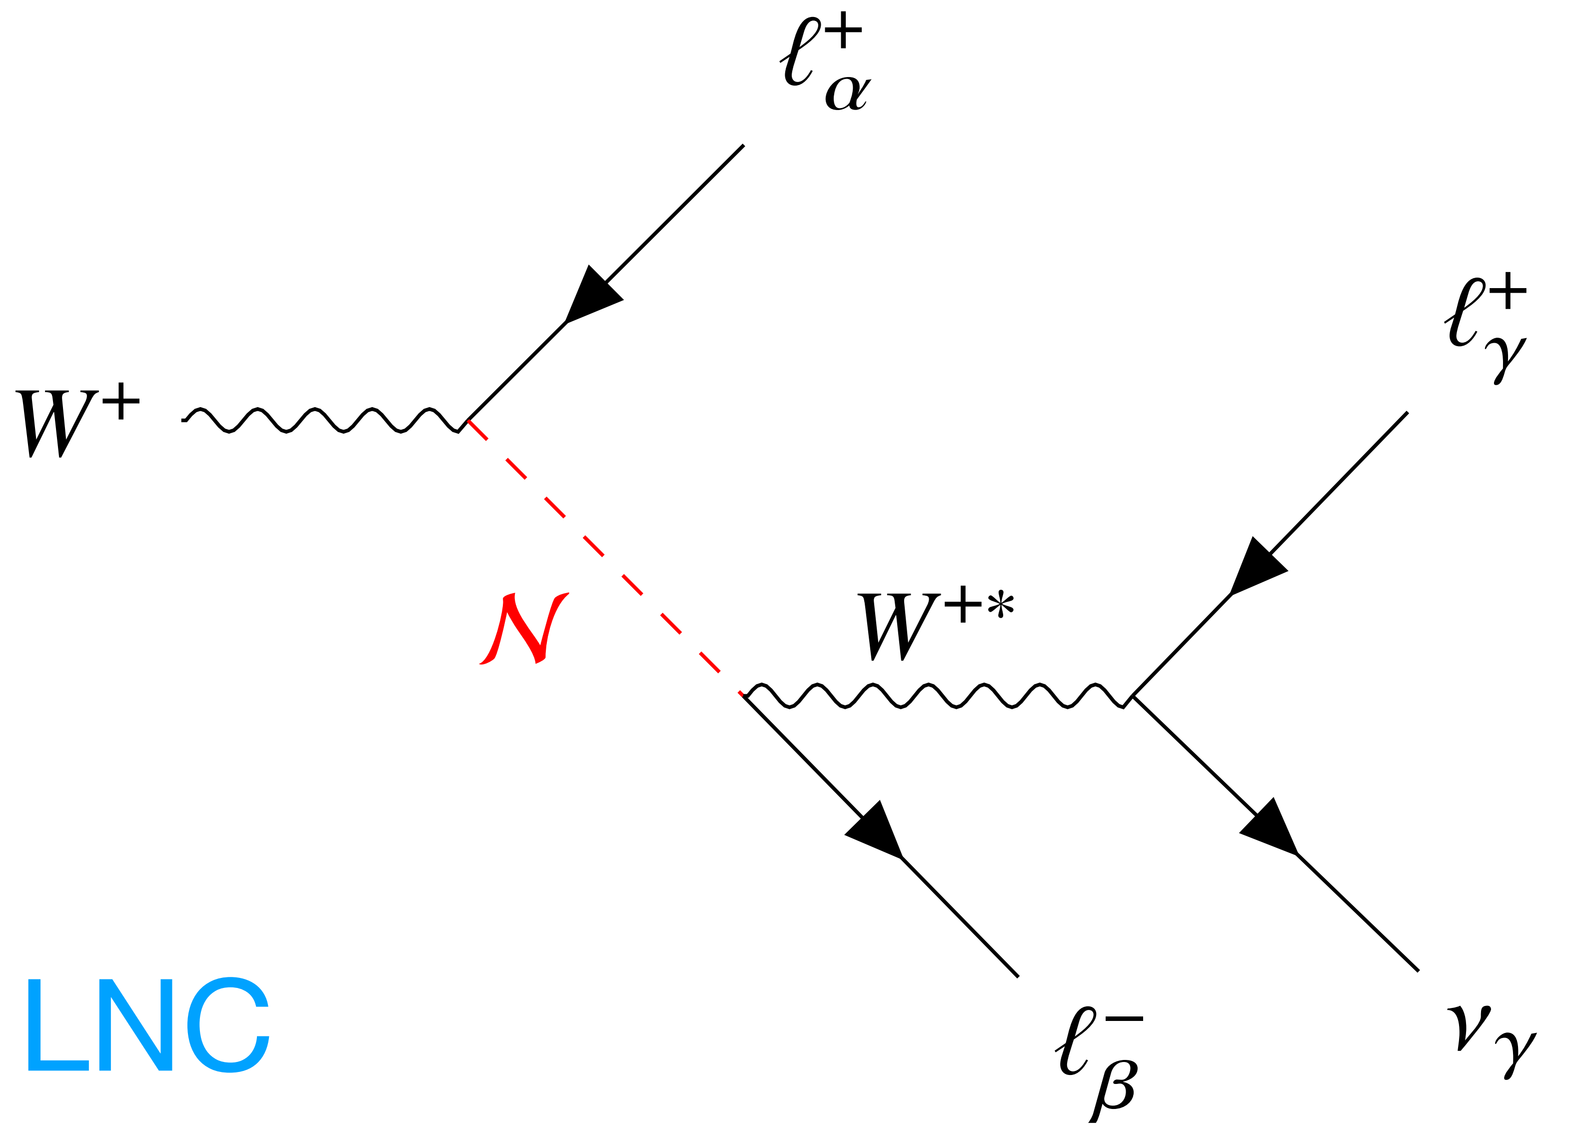

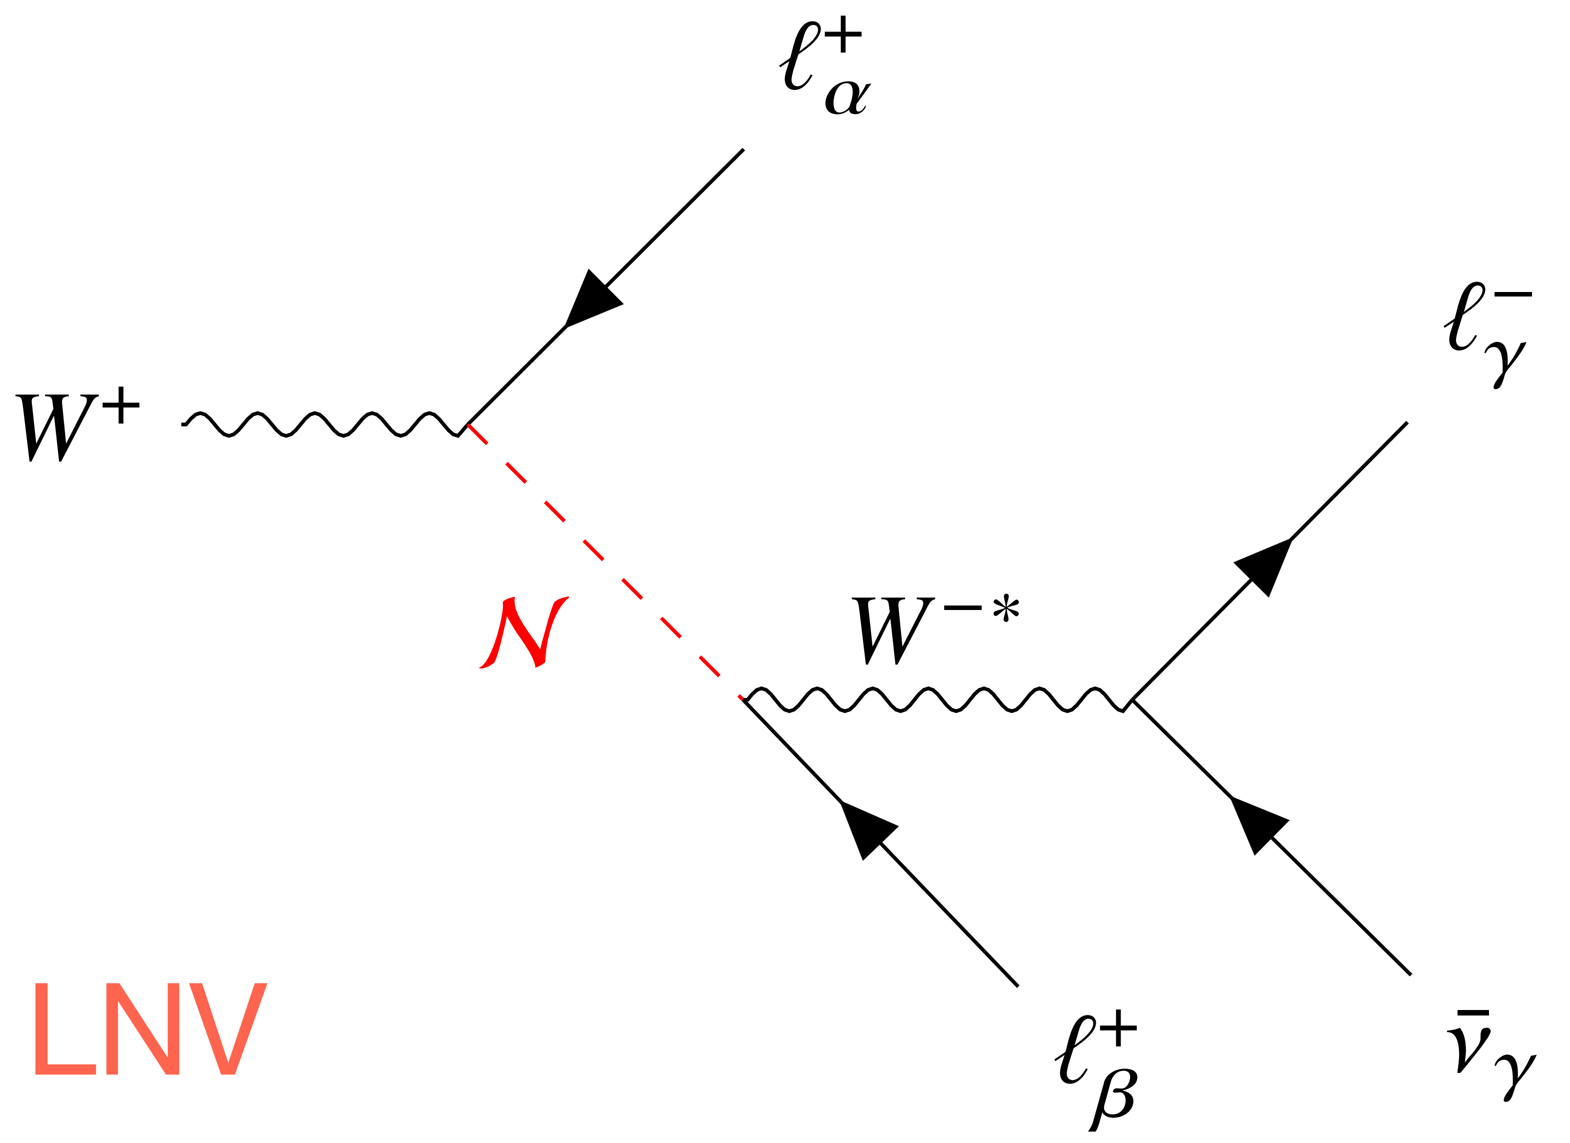

Cutflow

No OSSF

"Minimal" detector cuts

"Softest" available triggers

Avoid Z pole in e channel

Further cuts with almost no effect on signal

Muon channel	Electron channel
exactly $\mu^{\pm}\mu^{\pm}e^{\mp}$ signature	exactly $e^{\pm}e^{\pm}\mu^{\mp}$ signature
$p_T(\mu) > 4$ GeV $p_T(e) > 7$ GeV (2015), 4.5 GeV (2016)	
leading muon $p_T > 23$ GeV subleading muon $p_T > 14$ GeV	leading electron $p_T > 27$ GeV subleading electron $p_T > 10$ GeV
	$m(e, e) < 78$ GeV
$40 < m(\ell, \ell, \ell') < 90$ GeV b-jet veto $E_T^{miss} < 60$ GeV	

Backgrounds

- Contrary to SHiP, this search is **not** background-free.
- Fake leptons are "non-prompt" leptons from jets or from pileup.
- The main background components are "multi-fakes" and $t\bar{t}$ with a fake lepton. They are (in principle) reducible backgrounds.
- Multi-fakes are random crossings of multiple fake lepton tracks. They can be estimated by randomly shuffling lepton tracks in data. *Estimation regions* are used for this purpose.
- The normalisation factors μ_{mf} and μ_{tt} are also estimated from data, to correct the normalisation of the multi-fake estimation and $t\bar{t}$ Monte-Carlo sample. *Validation regions* are used for this purpose.

Results

Muon channel $\mu^{\pm}\mu^{\pm}e^{\mp}$

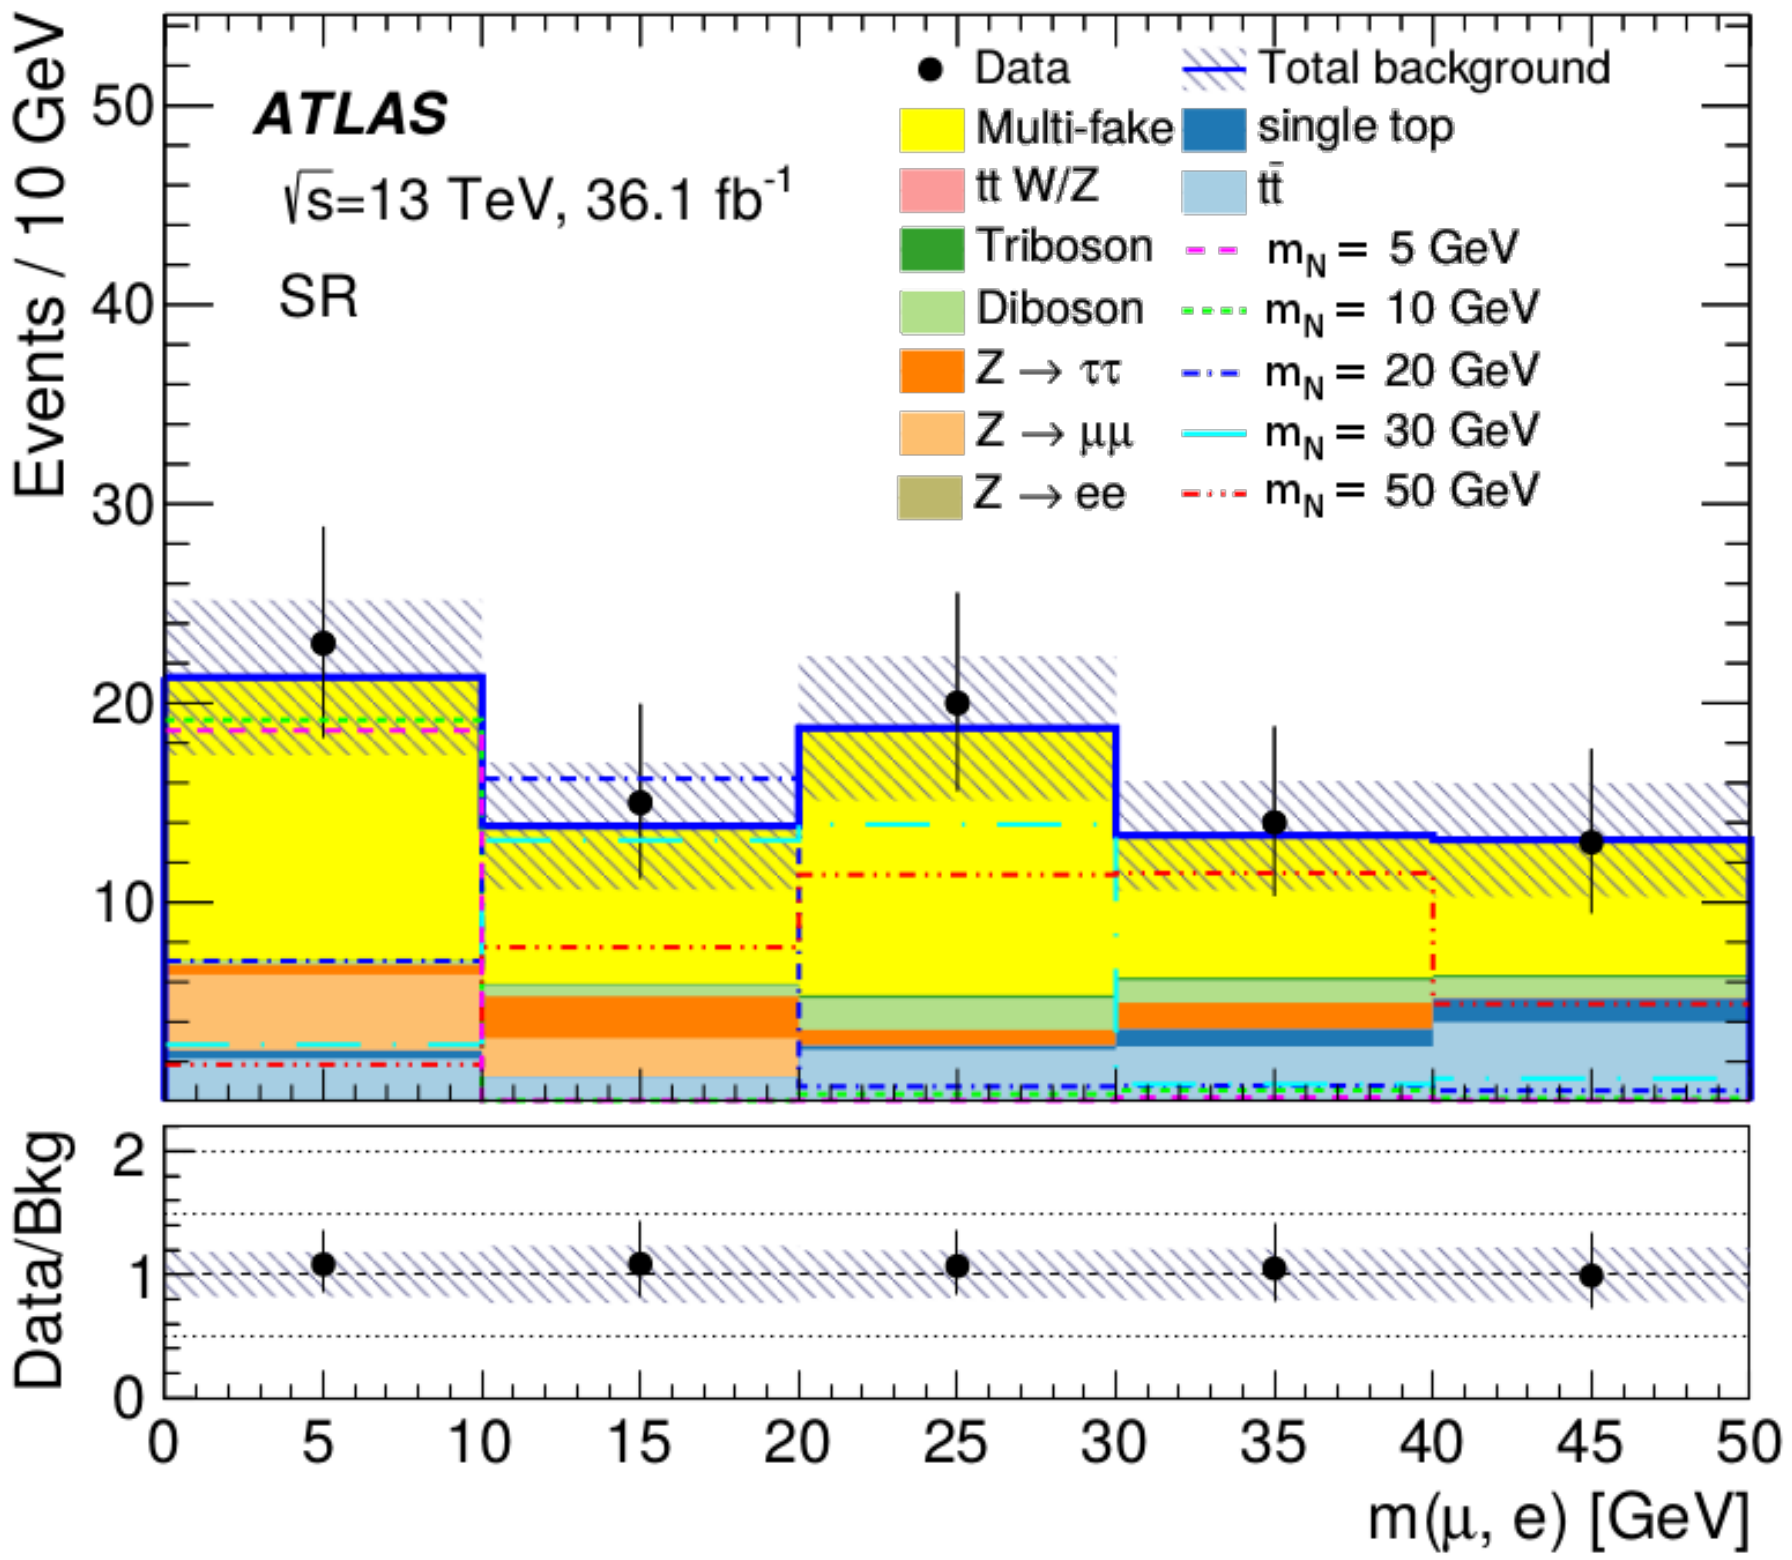

Electron channel $e^{\pm}e^{\pm}\mu^{\mp}$

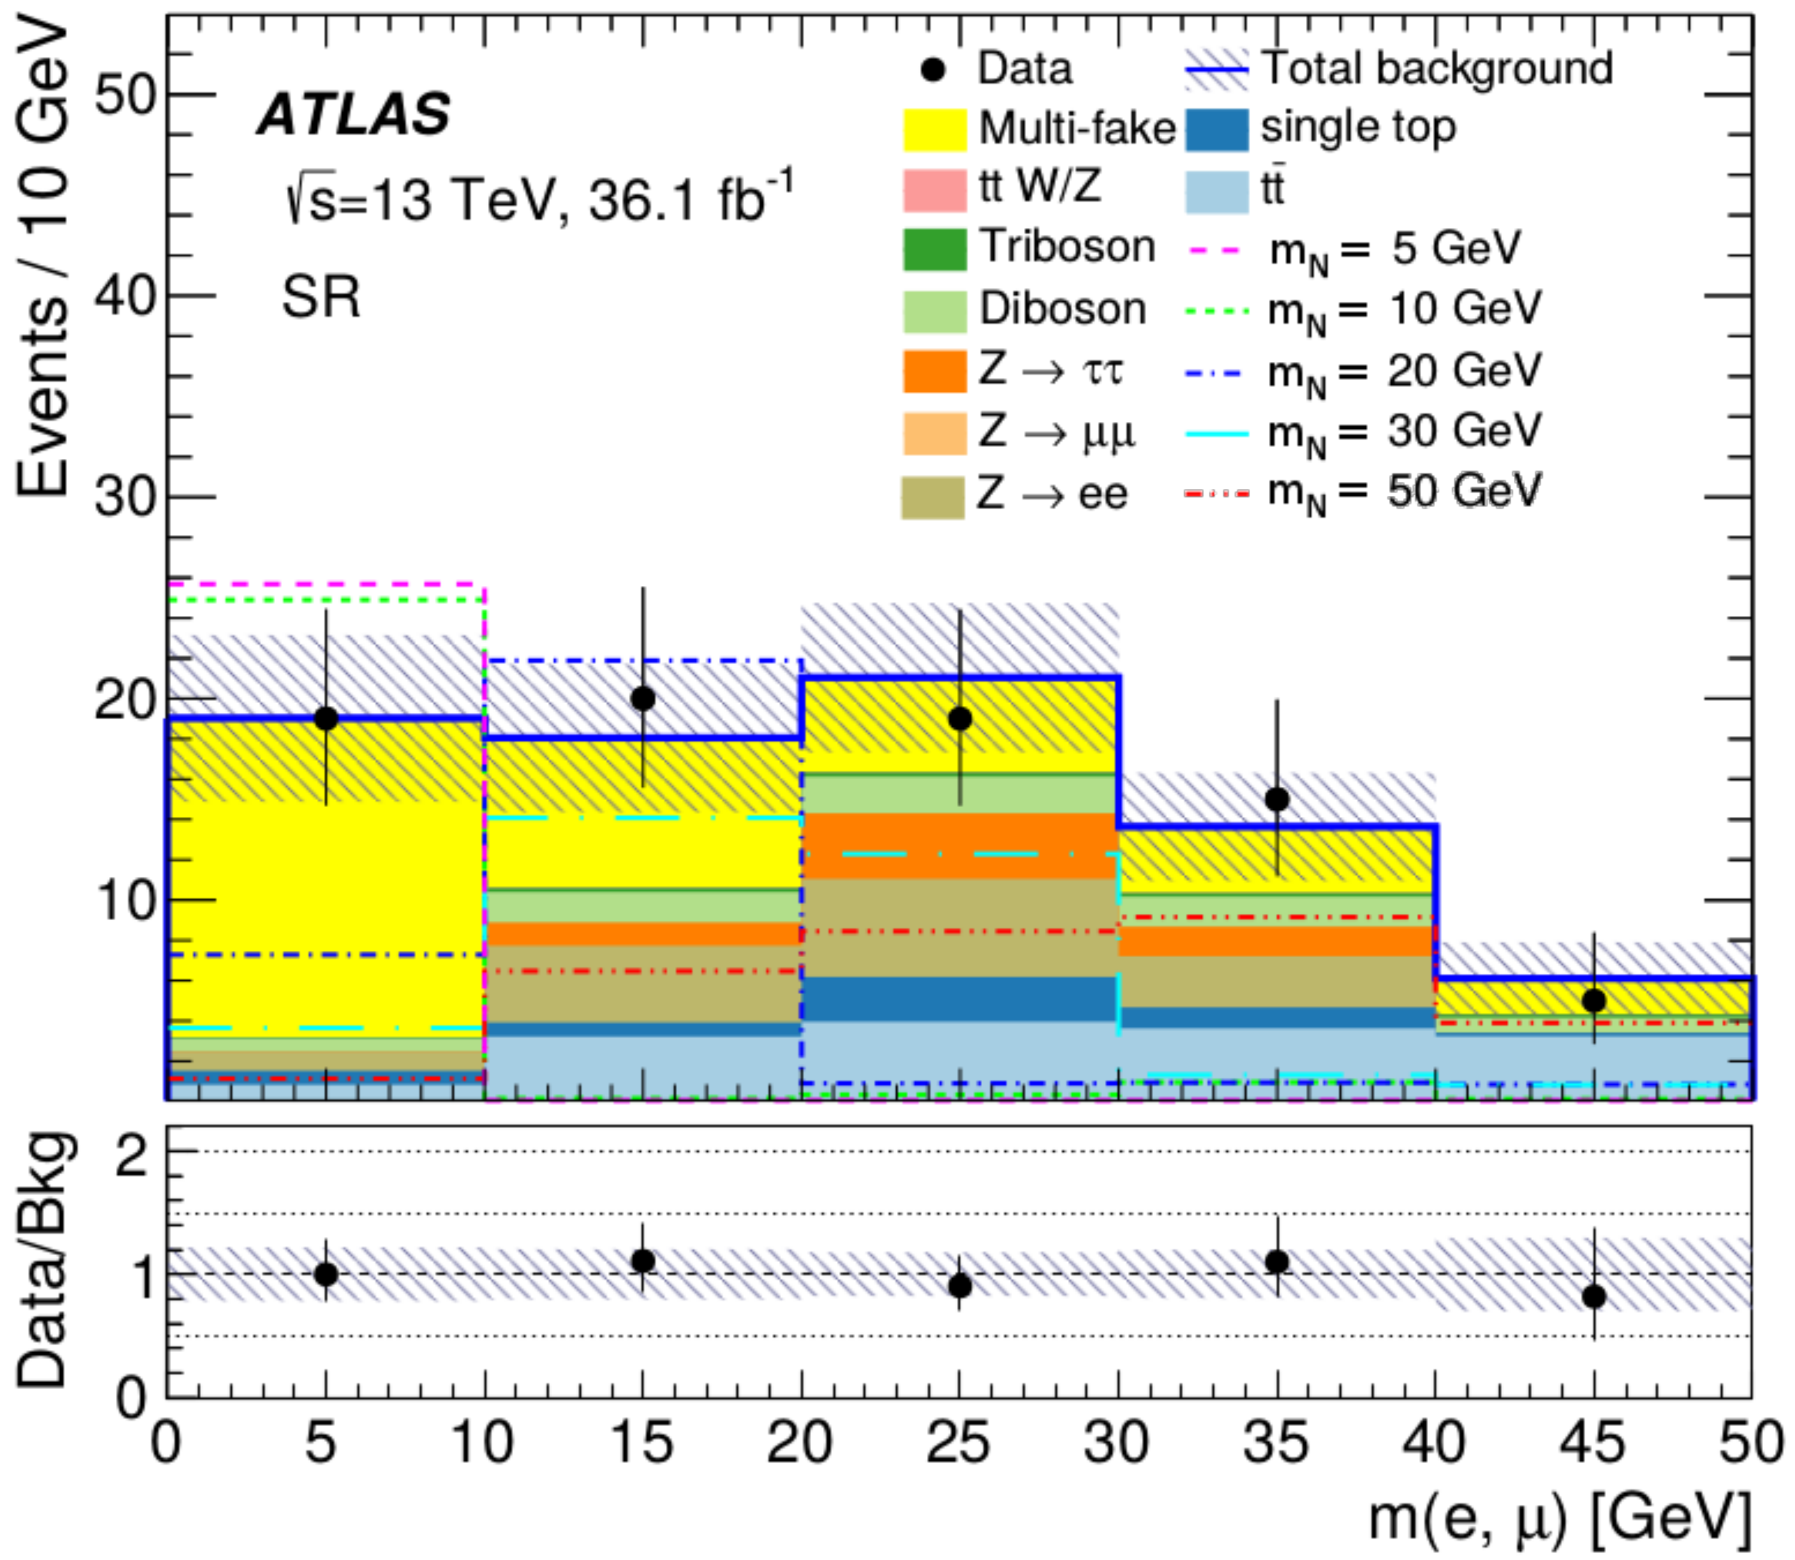

Results

No excess observed

Muon channel $\mu^{\pm}\mu^{\pm}e^{\mp}$

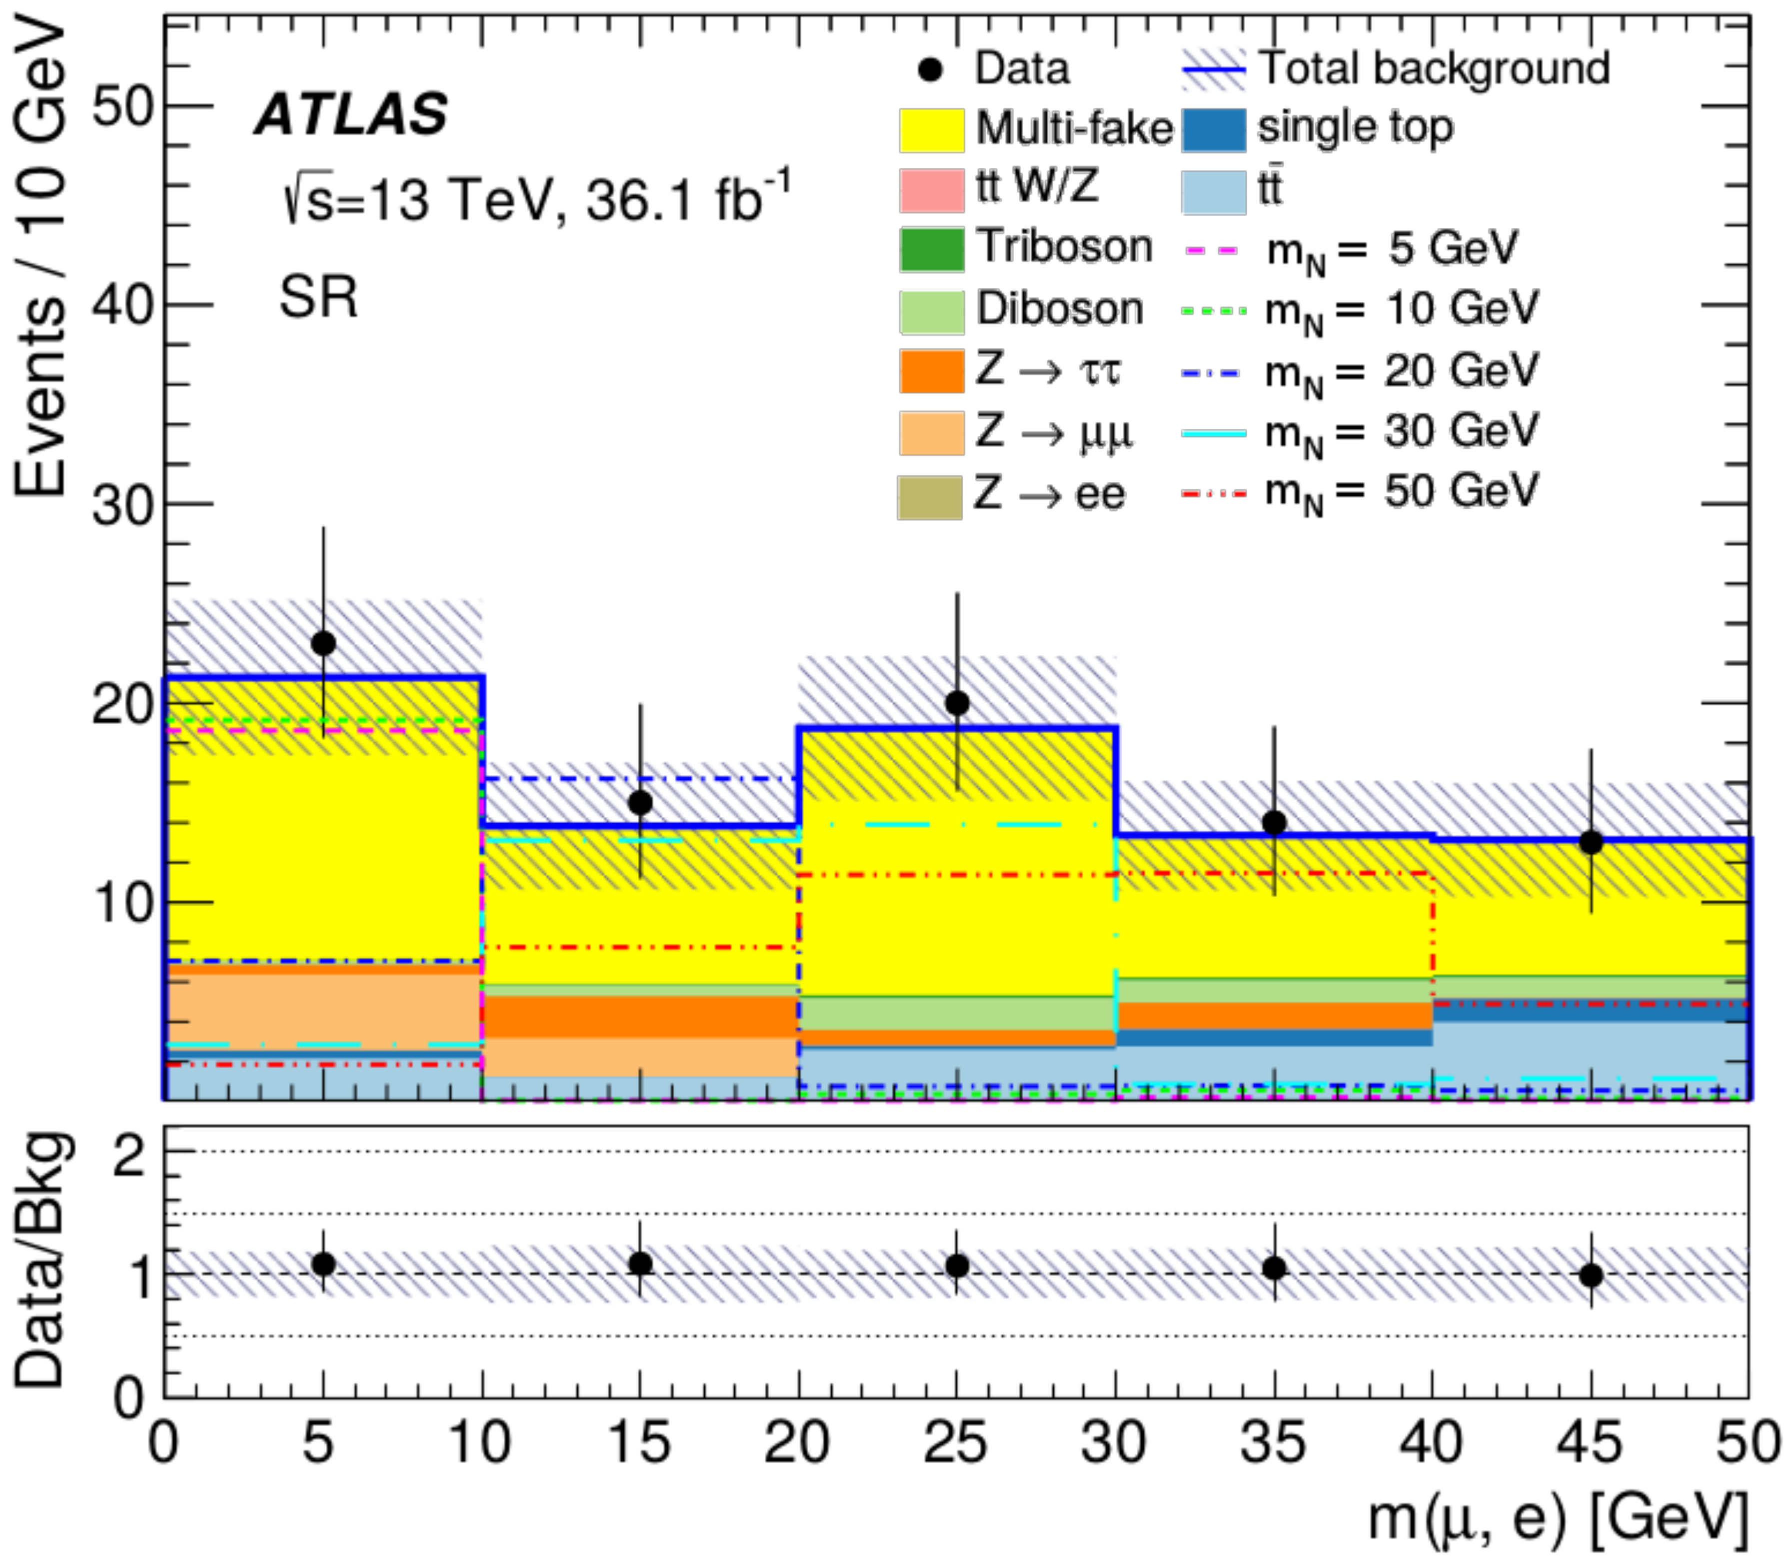

Electron channel $e^{\pm}e^{\pm}\mu^{\mp}$

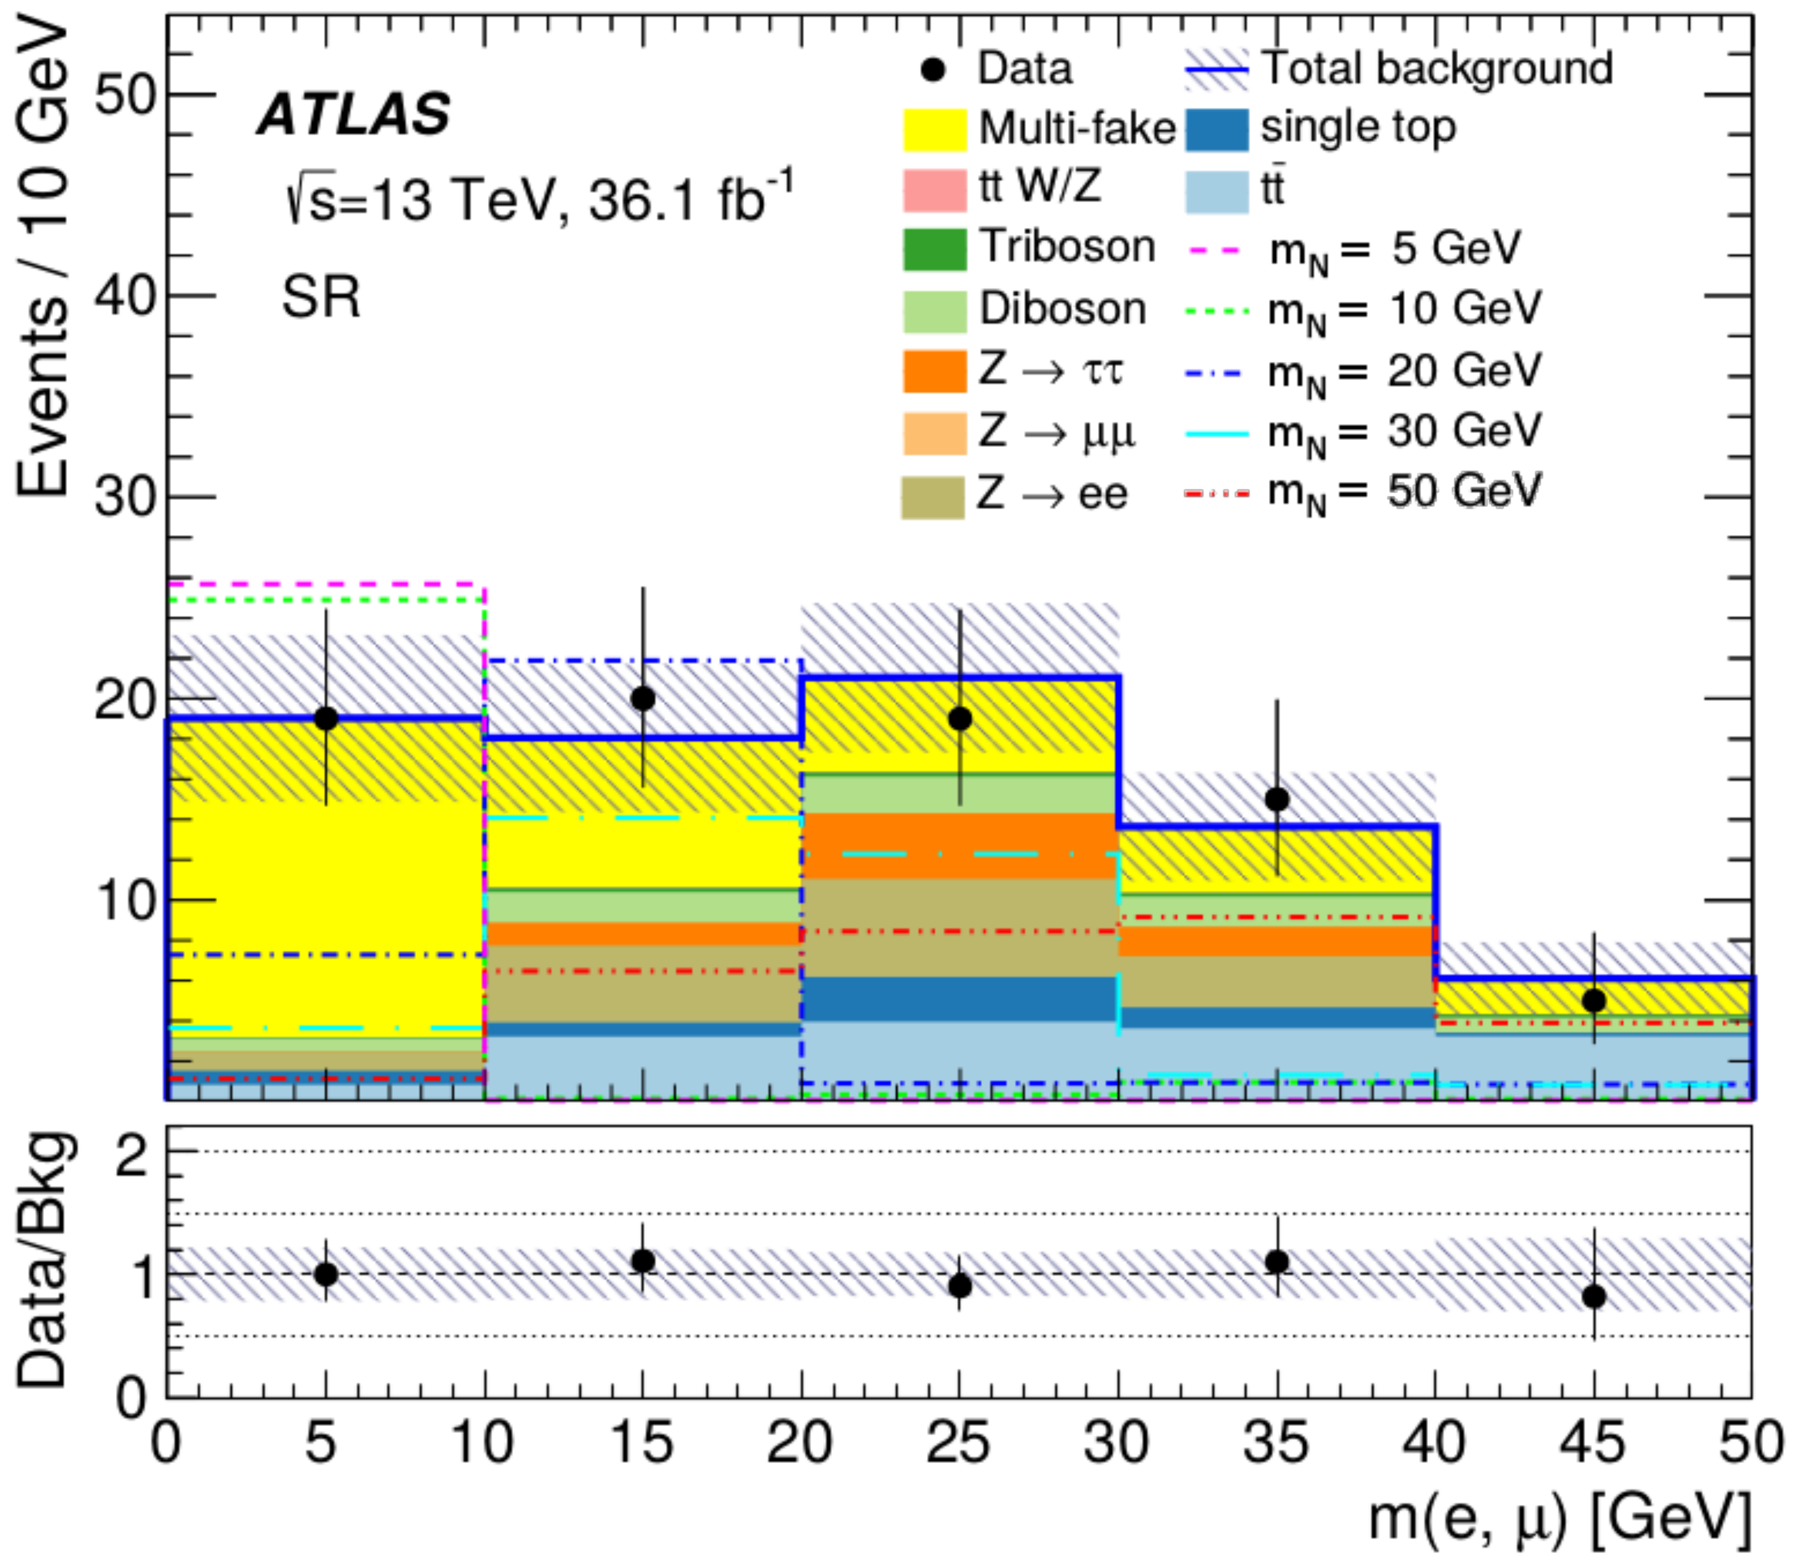

Exclusion limits

95% CL

- One Majorana HNL is assumed to mix with a single lepton flavour.
- Electron channel constrains $|\Theta_e|^2$
- Muon channel constrains $|\Theta_\mu|^2$

1 Majorana HNL mixing with electron flavour

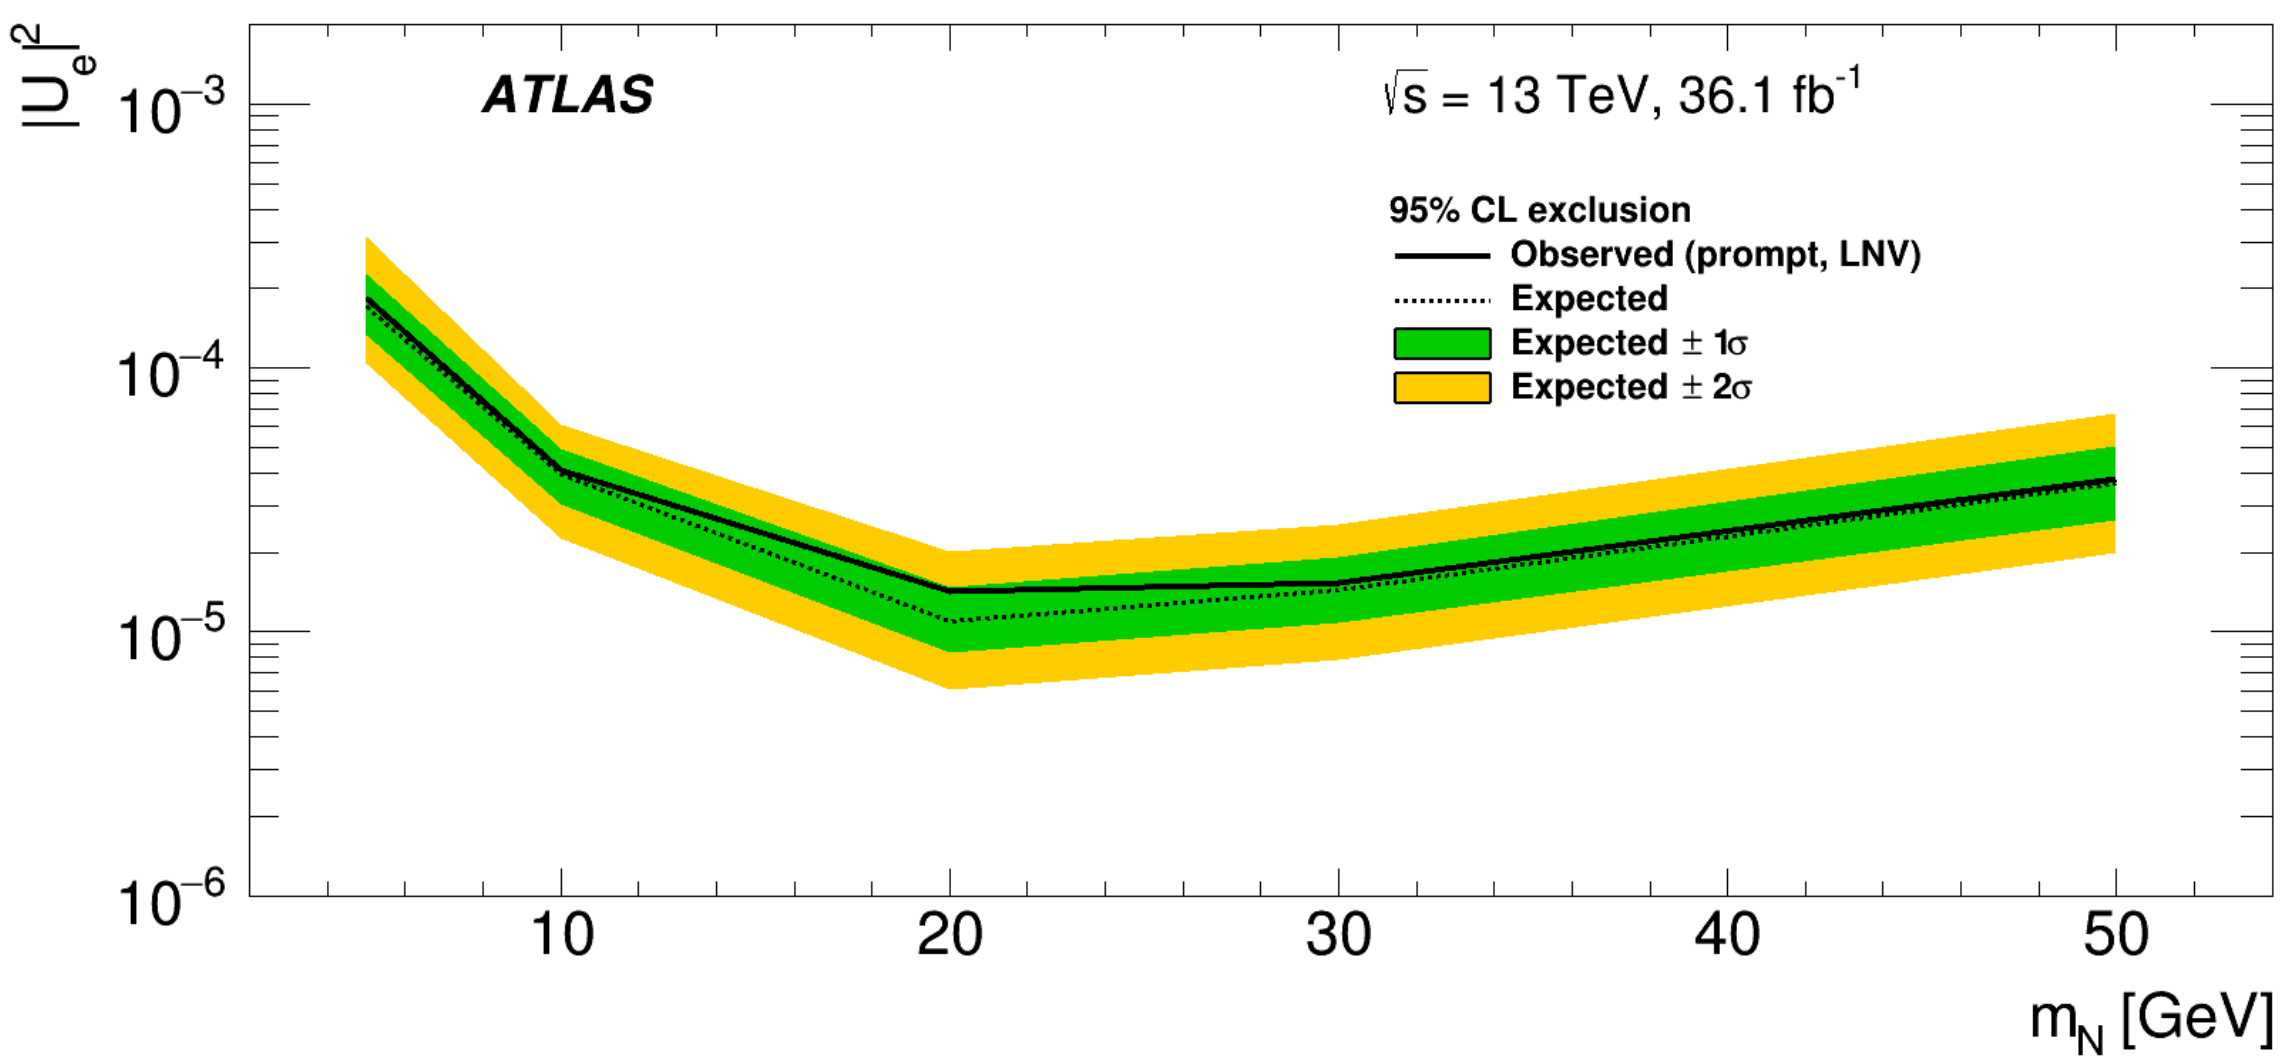

1 Majorana HNL mixing with muon flavour

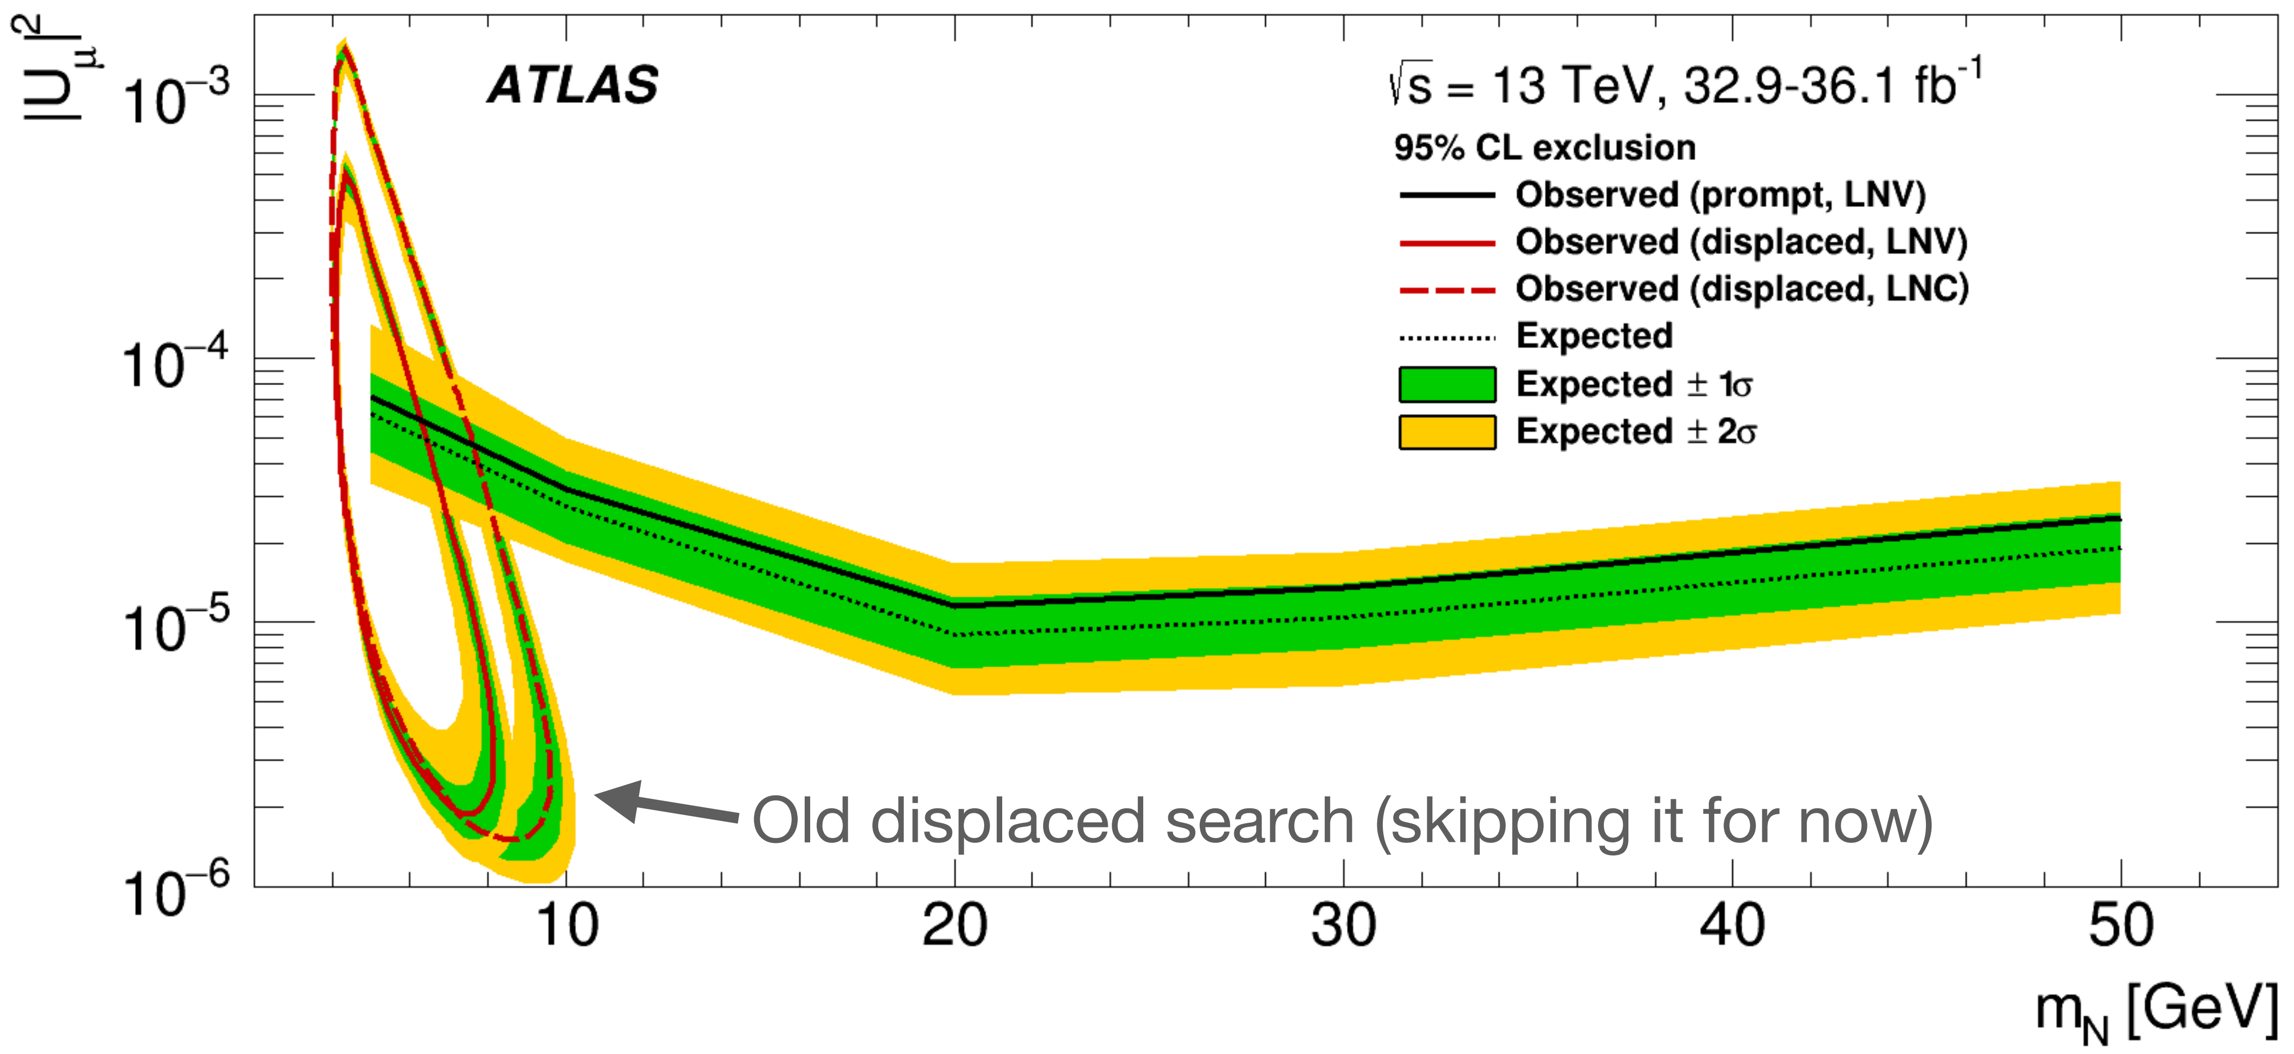

"Great! We have limits on 1 Majorana HNL mixing with a single flavour.
Where does it fit in our ternary plot?"

Oops #1

- One HNL can only give mass to one neutrino:

  → The mass is $m^\nu = M_N |\Theta|^2$.

- HNLs above our limits lead to a mass of around 1 MeV.

- This benchmark was never really a "realistic" model of neutrino masses.

Oops #1

- One HNL can only give mass to one neutrino:

  → The mass is $m^\nu = M_N |\Theta|^2$.

- HNLs above our limits lead to a mass of around 1 MeV.
- This benchmark was never really a "realistic" model of neutrino masses.

Ok, let's try with 2 quasi-Dirac HNLs...

Oops #2

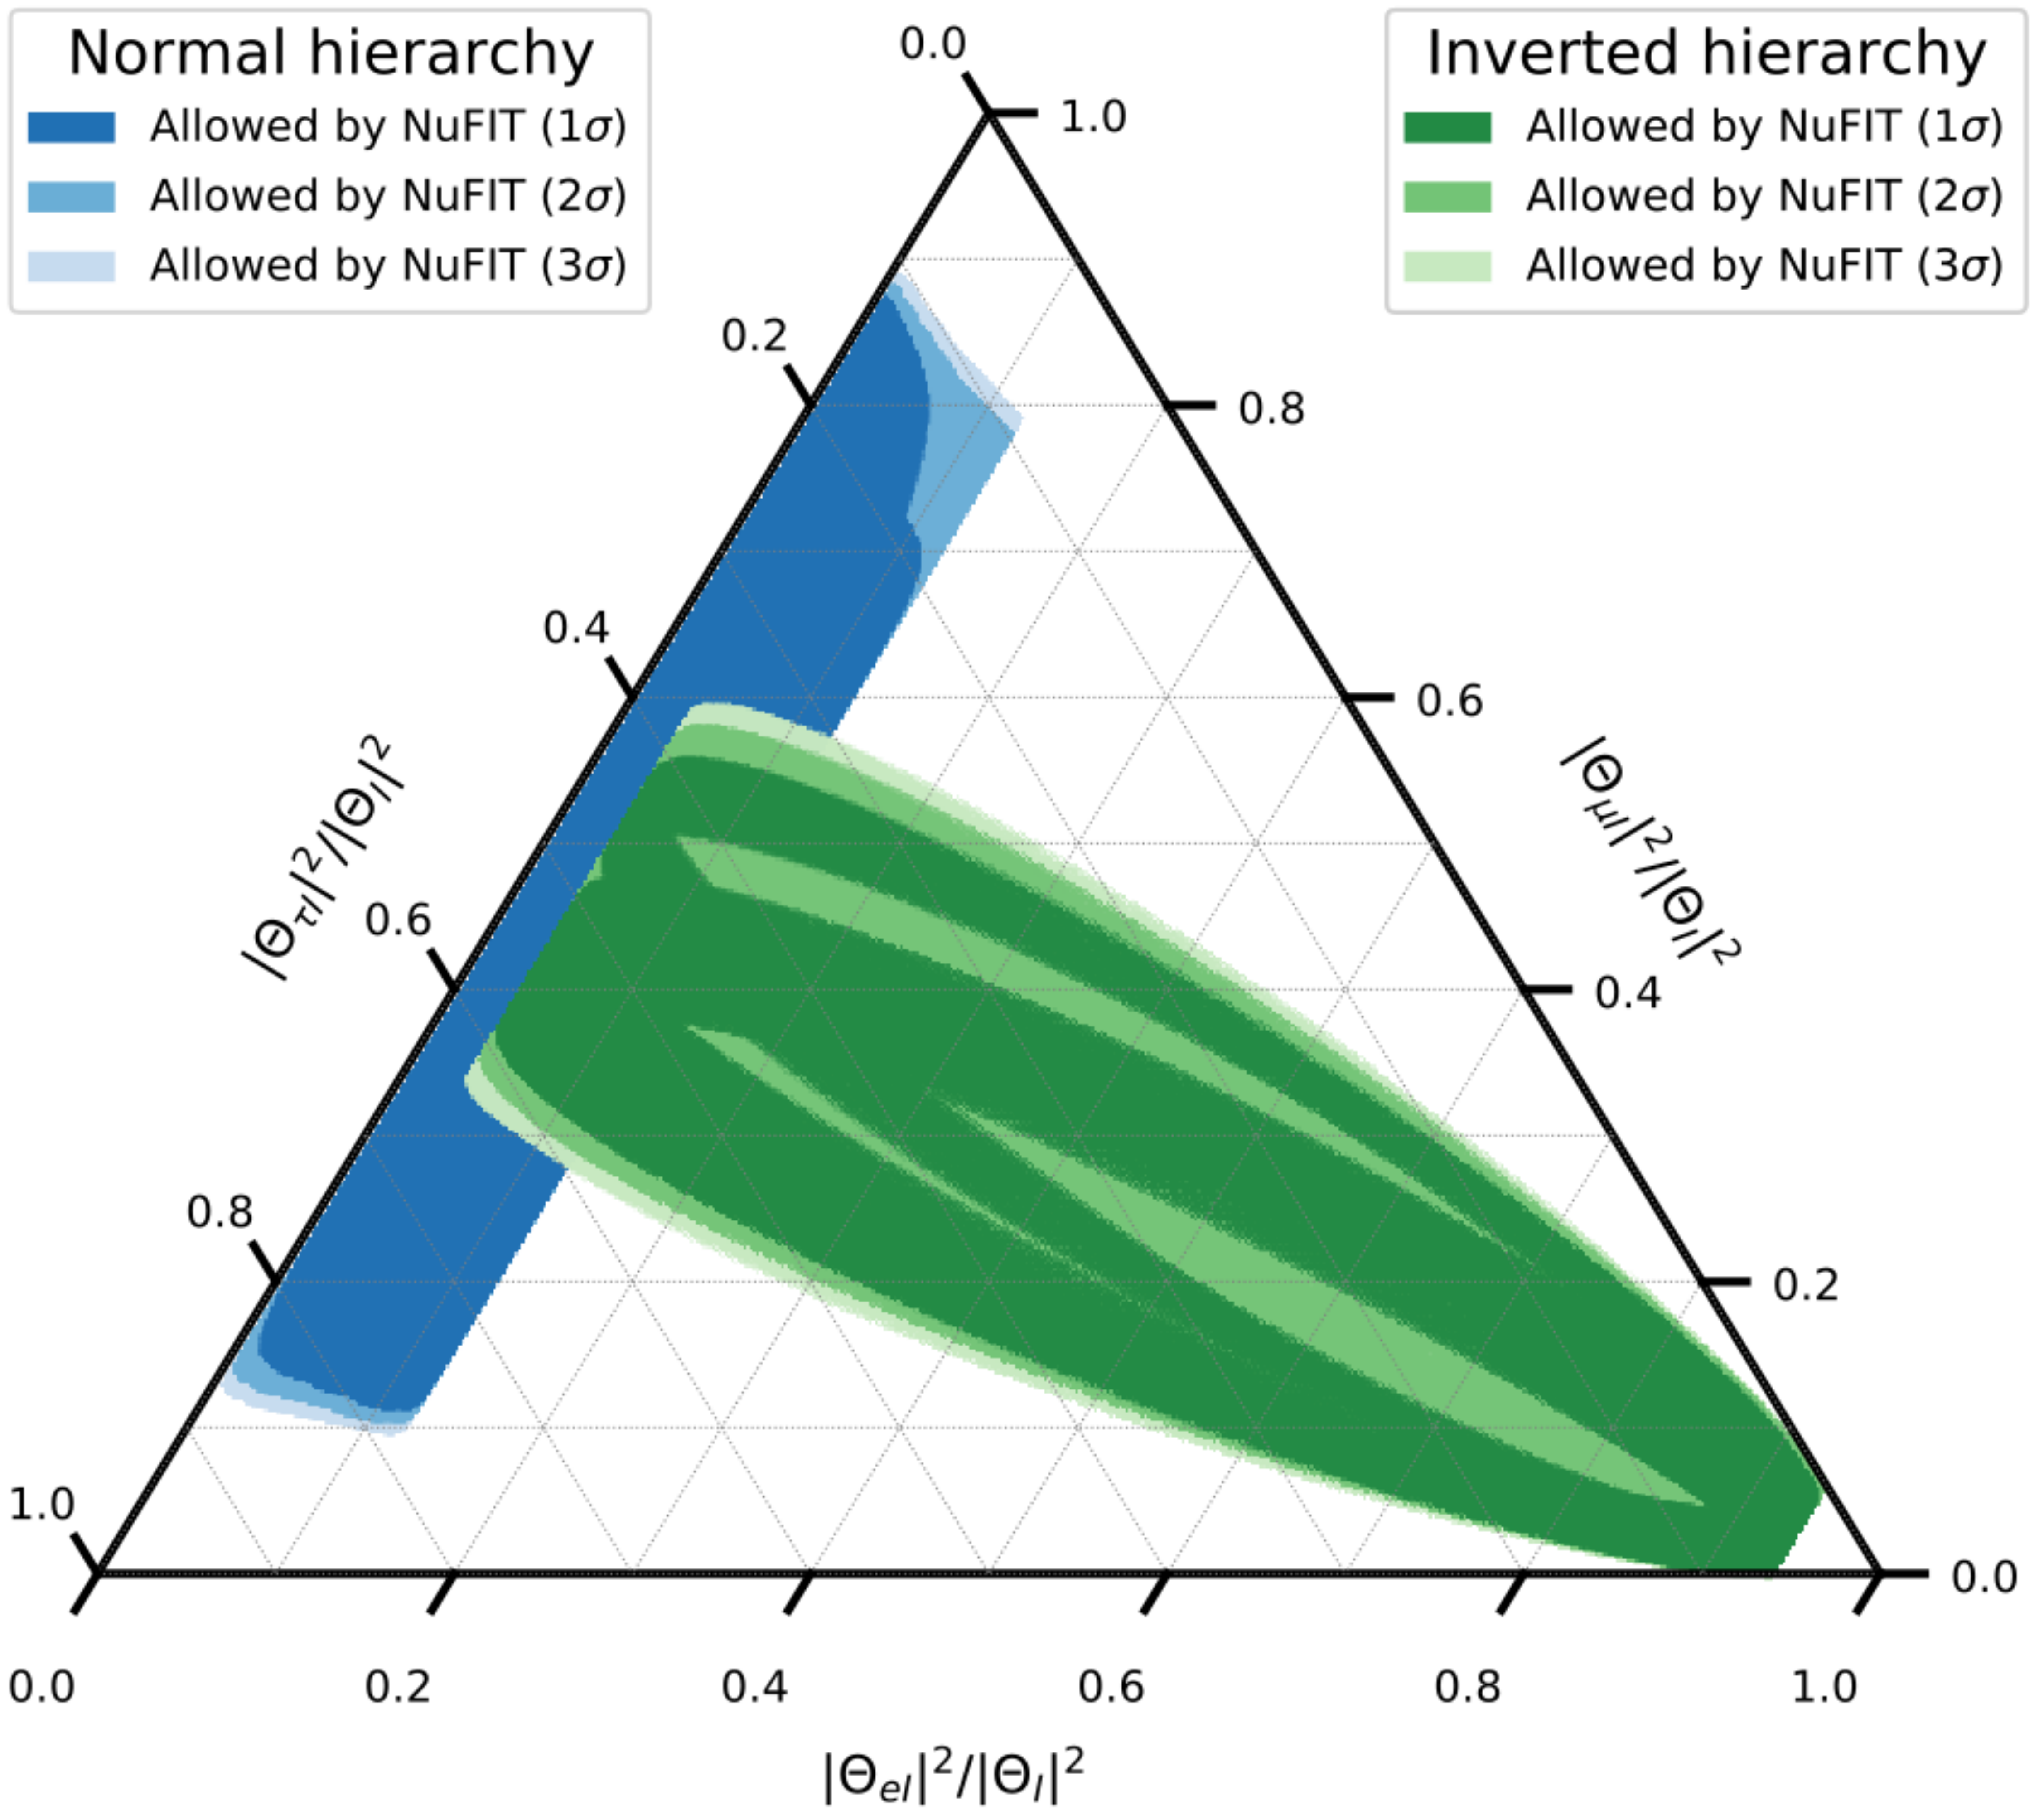

Normal hierarchy
Allowed by NuFIT (1σ)
Allowed by NuFIT (2σ)
Allowed by NuFIT (3σ)
Inverted hierarchy
Allowed by NuFIT (1σ)
Allowed by NuFIT (2σ)
Allowed by NuFIT (3σ)
0.0
0.2
0.4
0.6
0.8
1.0
$|\Theta_{\tau I}|^2/|\Theta_I|^2$
$|\Theta_{\mu I}|^2/|\Theta_I|^2$
$|\Theta_{eI}|^2/|\Theta_I|^2$

Oops #2

Even with two quasi-Dirac HNLs, single-flavour mixing is not compatible with neutrino oscillation data!

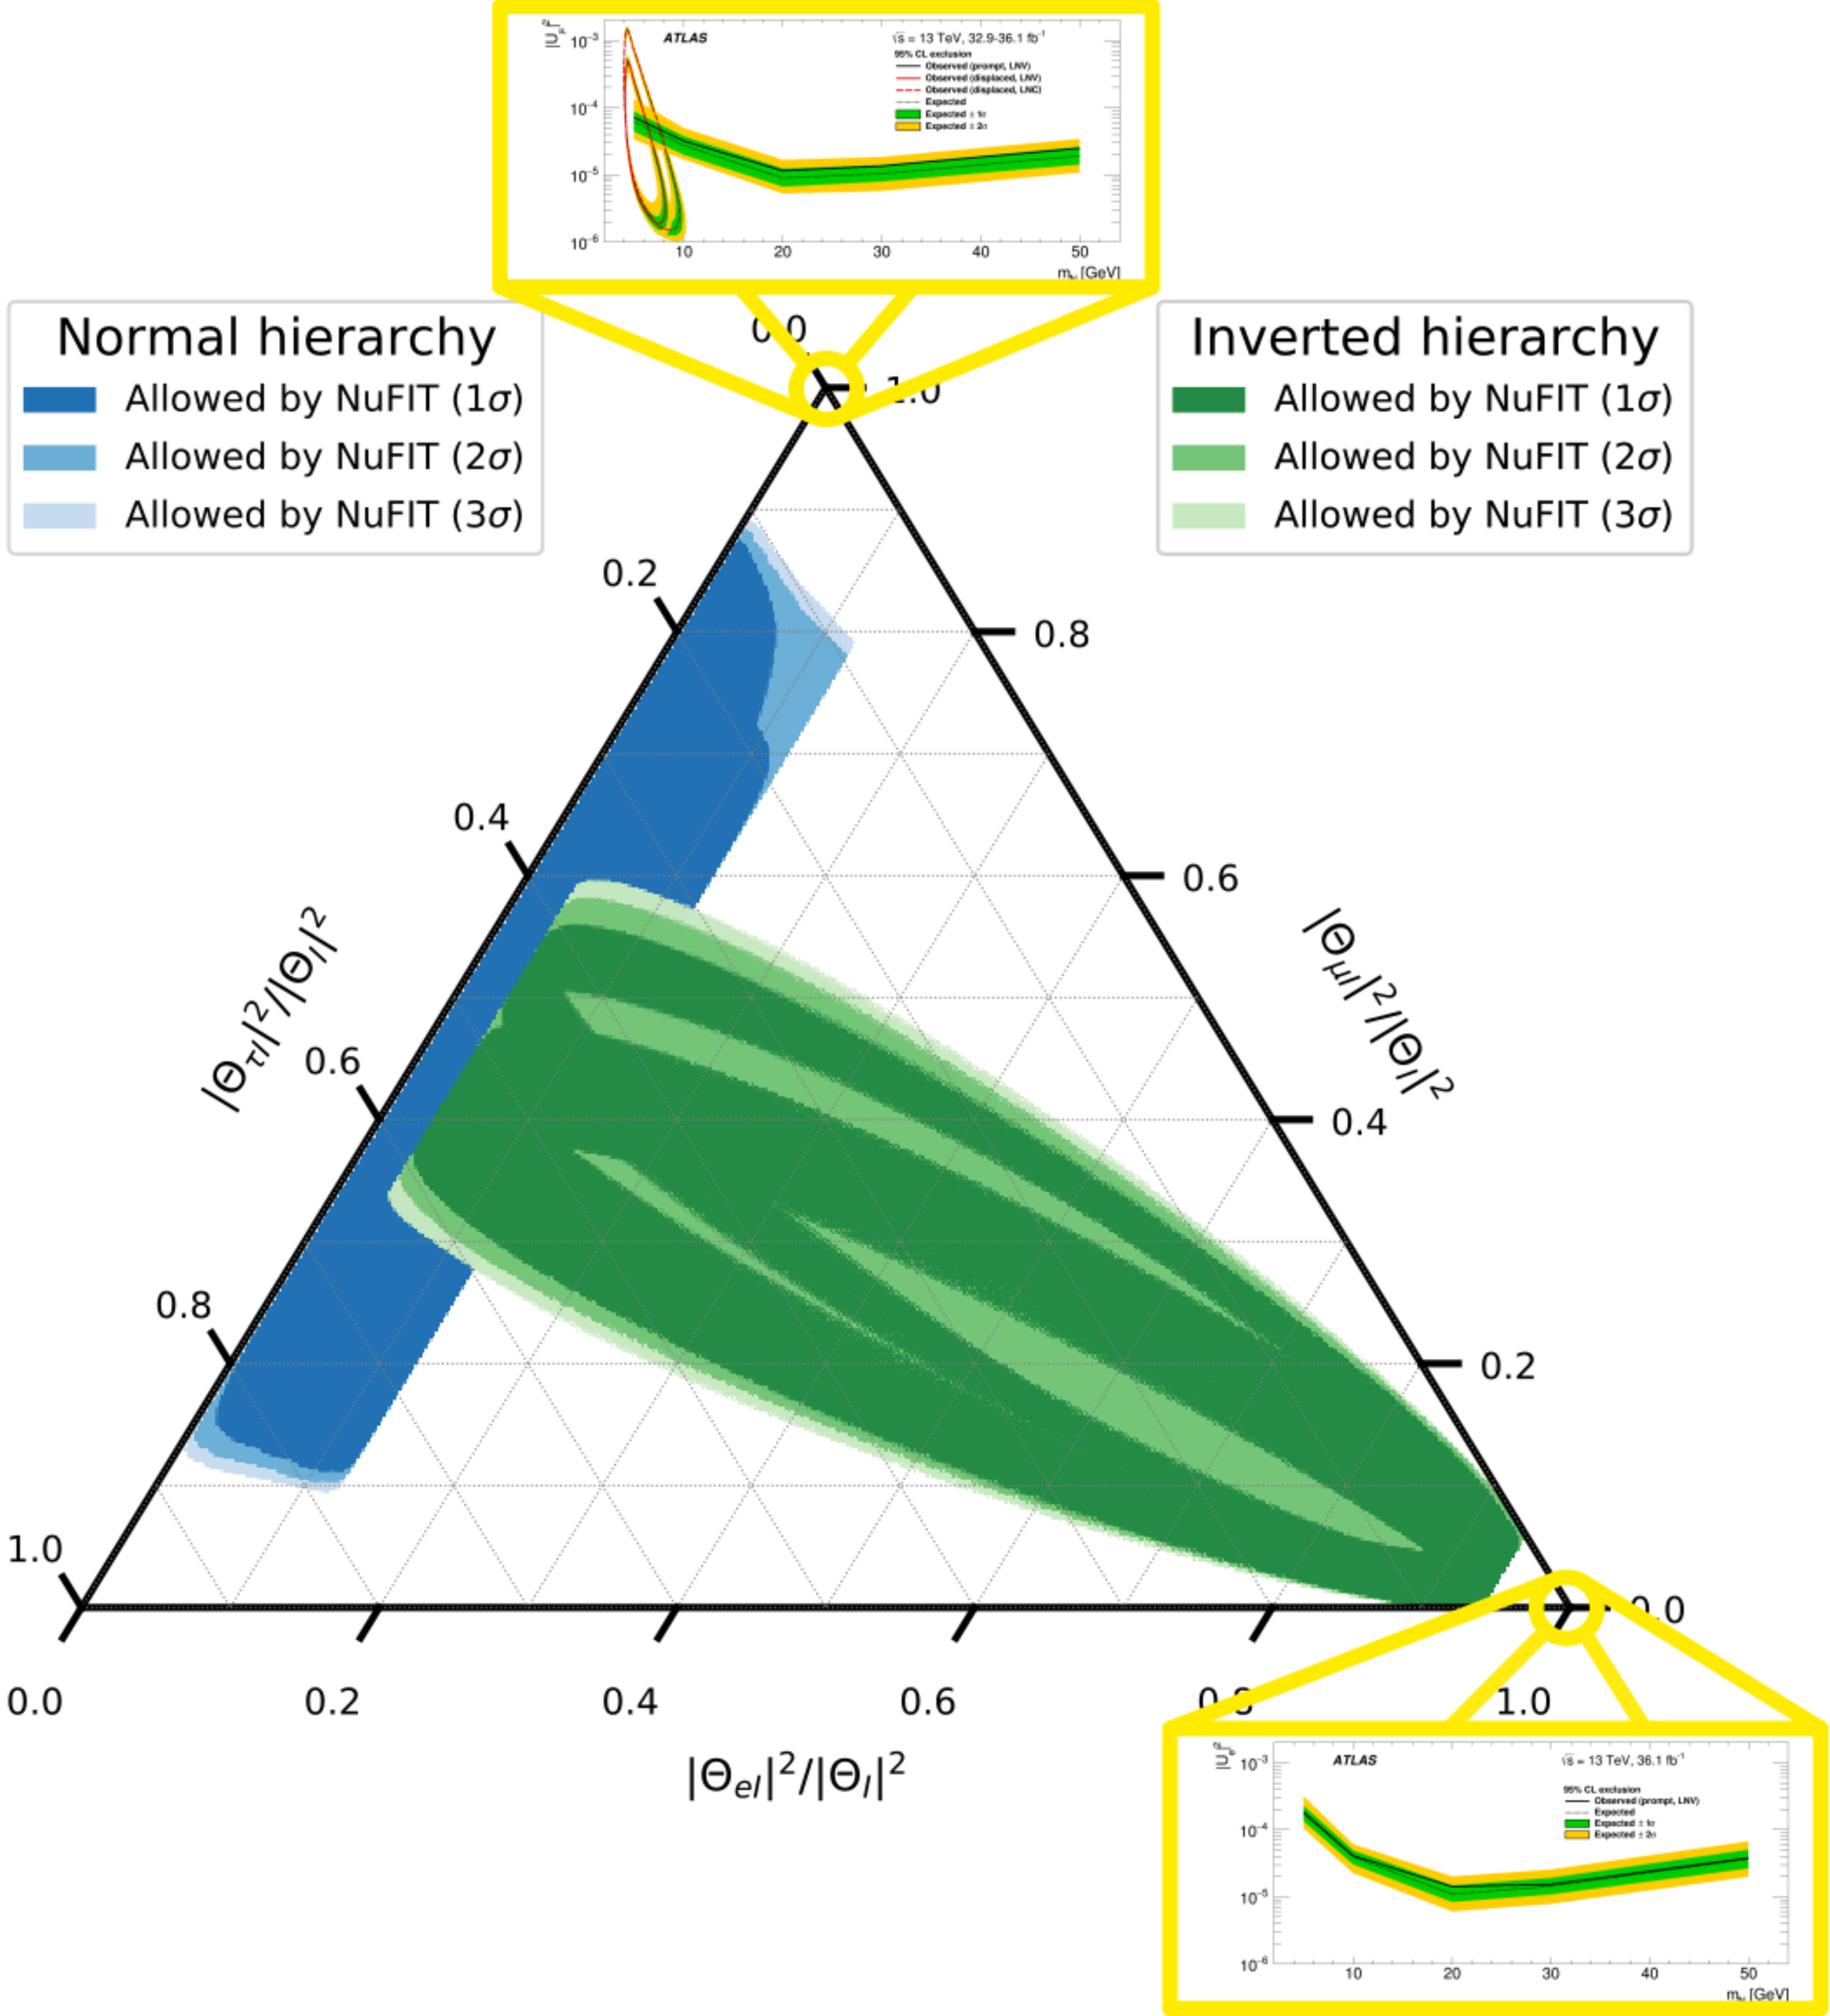

The need for realistic models

- It is **understood** that the simplified benchmarks used by experiments to report their limits are not meant to represent real models of neutrino masses.
- They are **useful** to consistently compare the limits between experiments.
- **However**, I will argue that:
    - They may give a misleading impression of how well the parameter space is *truly* constrained.
    - They are not very useful when trying to perform *global parameter scans*.

Blind spot #1

Lepton number conserving processes

- Take the muon channel as example: Final state = $\mu^+\mu^+e^-$ or c.c.
- Suppose the HNL mixes only with the muon flavour: $\Theta_e = 0$.
- The LNC process is suppressed because it requires both Θ_e and Θ_μ to be non-zero $\Rightarrow$ no sensitivity!
- There is an unsuppressed LNC diagram, but it contributes to the vetoed final state $\mu^+\mu^-e^+$.

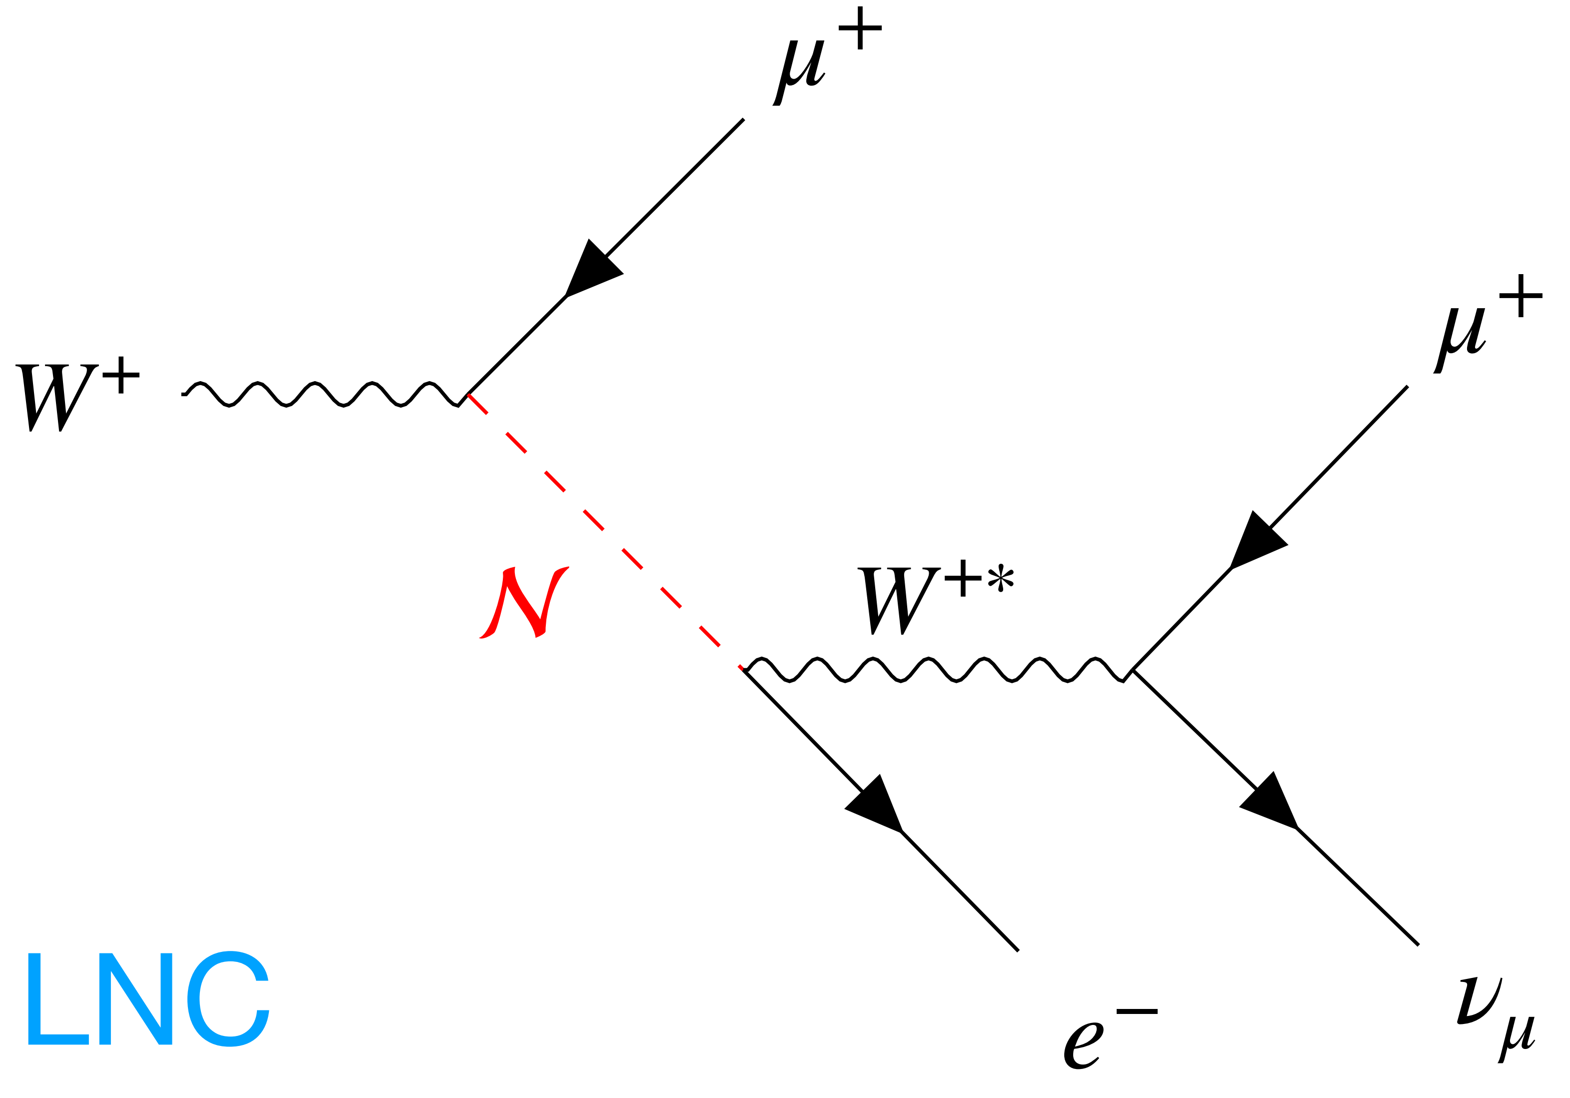

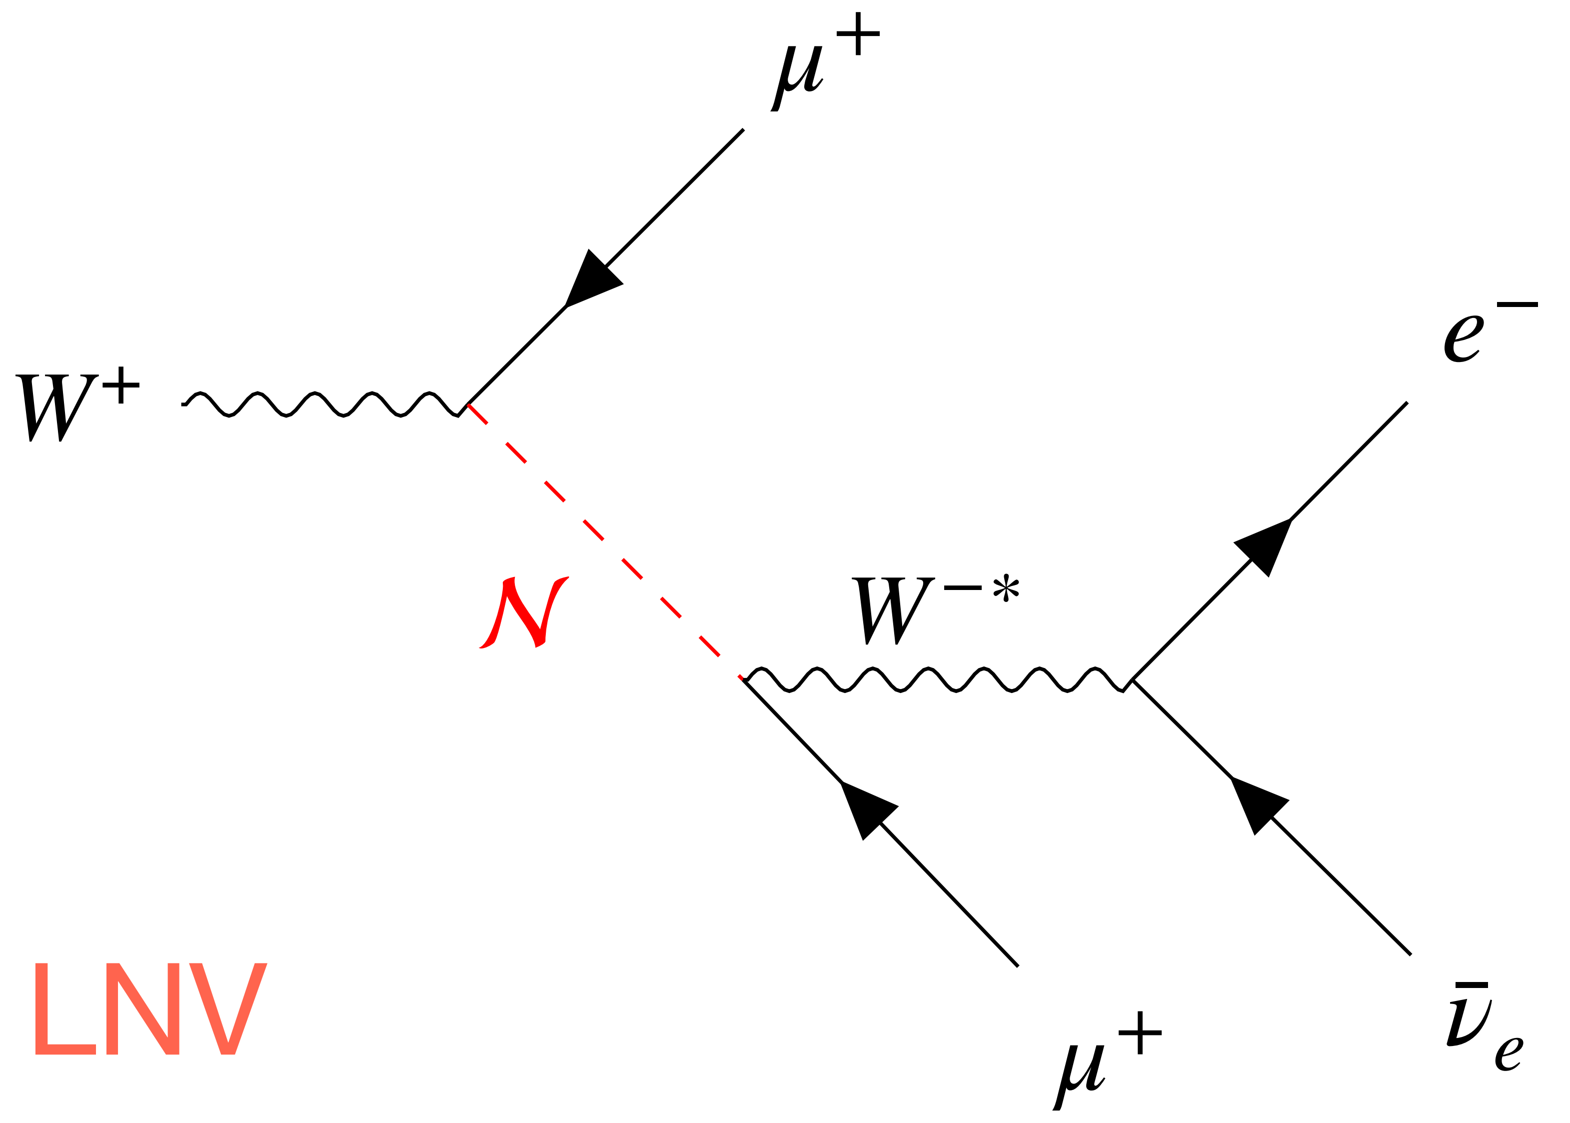

Blind spot #1

Lepton number conserving processes

- Take the muon channel as example: Final state = $\mu^+\mu^+e^-$ or c.c.
- Suppose the HNL mixes only with the muon flavour: $\Theta_e = 0$.
- The LNC process is suppressed because it requires both Θ_e and Θ_μ to be non-zero $\Rightarrow$ no sensitivity!
- There is an unsuppressed LNC diagram, but it contributes to the vetoed final state $\mu^+\mu^-e^+$.

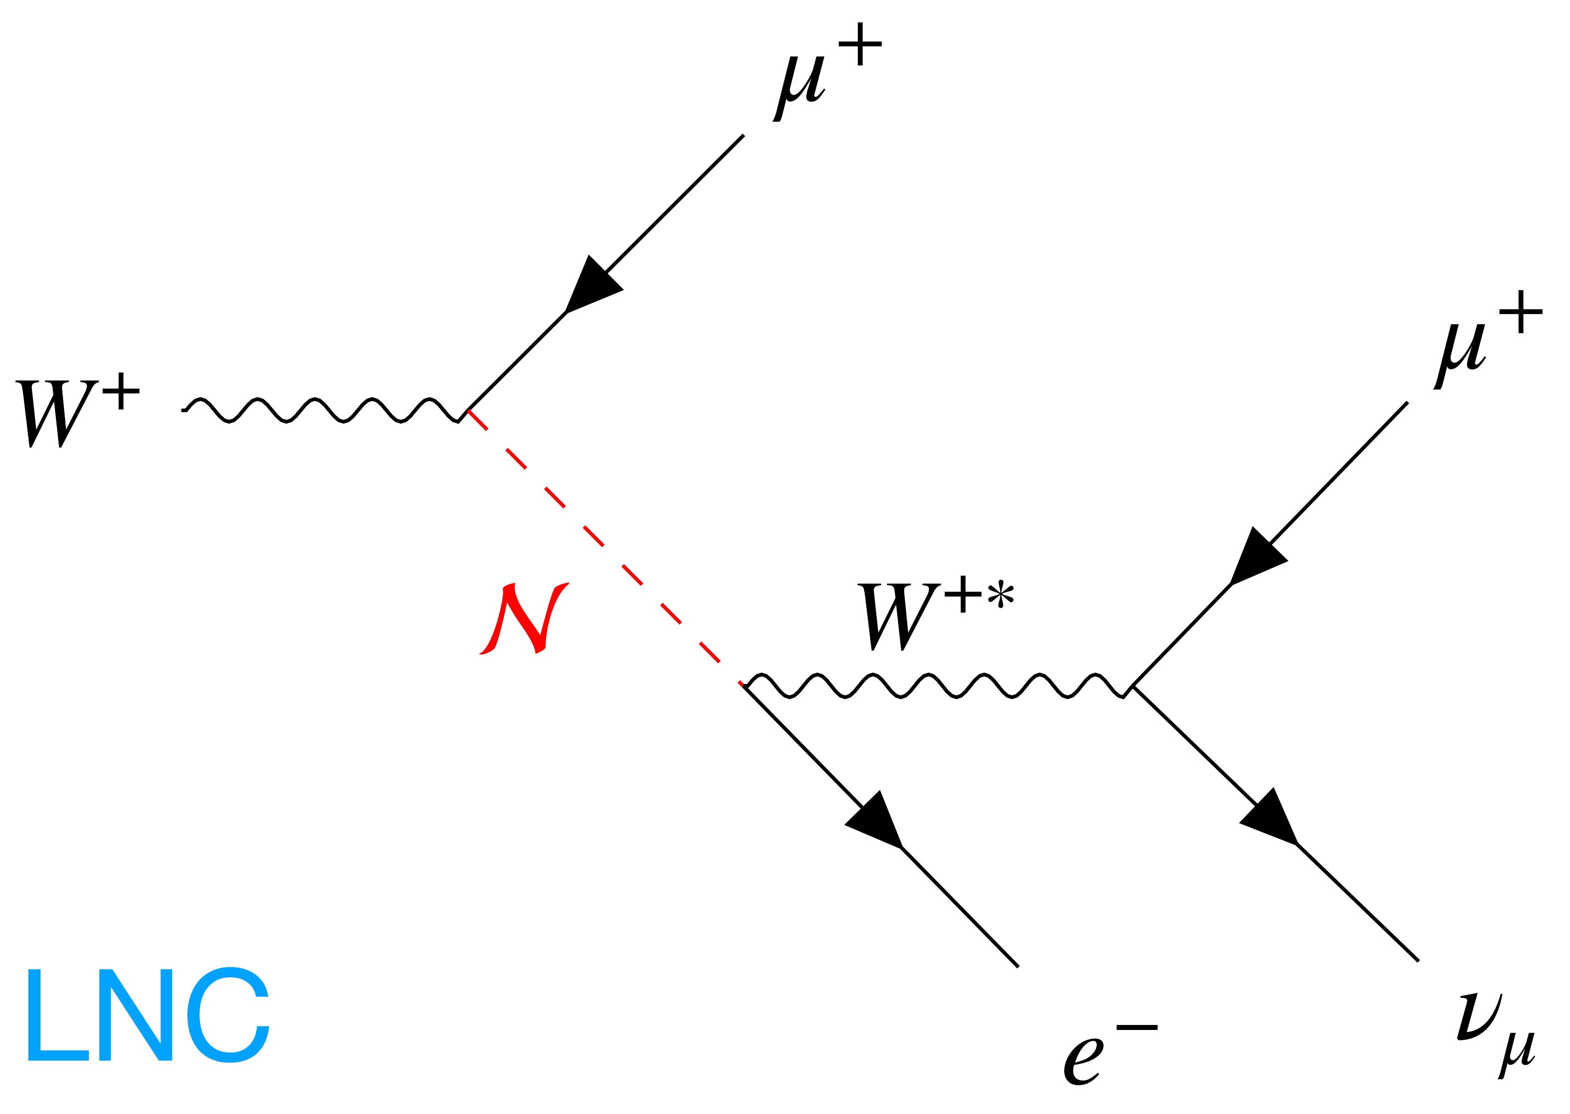

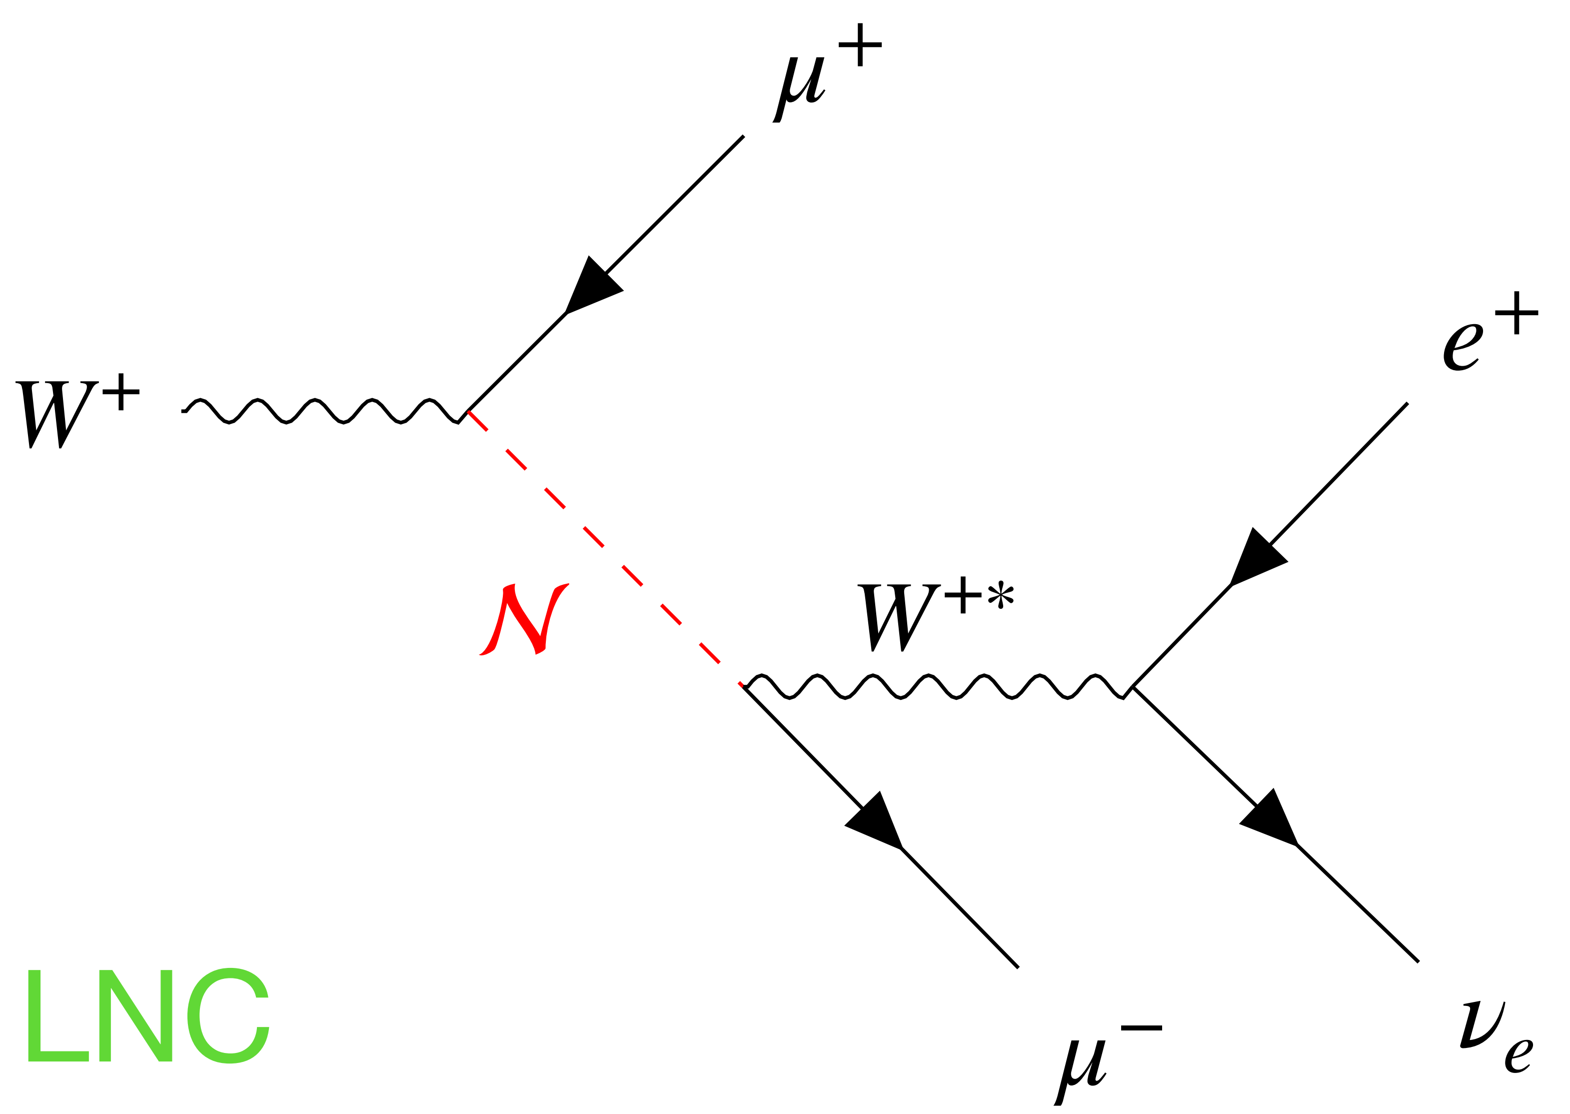

Blind spot #1

Lepton number conserving processes

- Take the muon channel as example: Final state = $\mu^+\mu^+e^-$ or c.c.
- Suppose the HNL mixes only with the muon flavour: $\Theta_e = 0$.
- The LNC process is suppressed because it requires both Θ_e and Θ_μ to be non-zero ⇒ no sensitivity!
- There is an unsuppressed LNC diagram, but it contributes to the vetoed final state $\mu^+\mu^-e^+$.

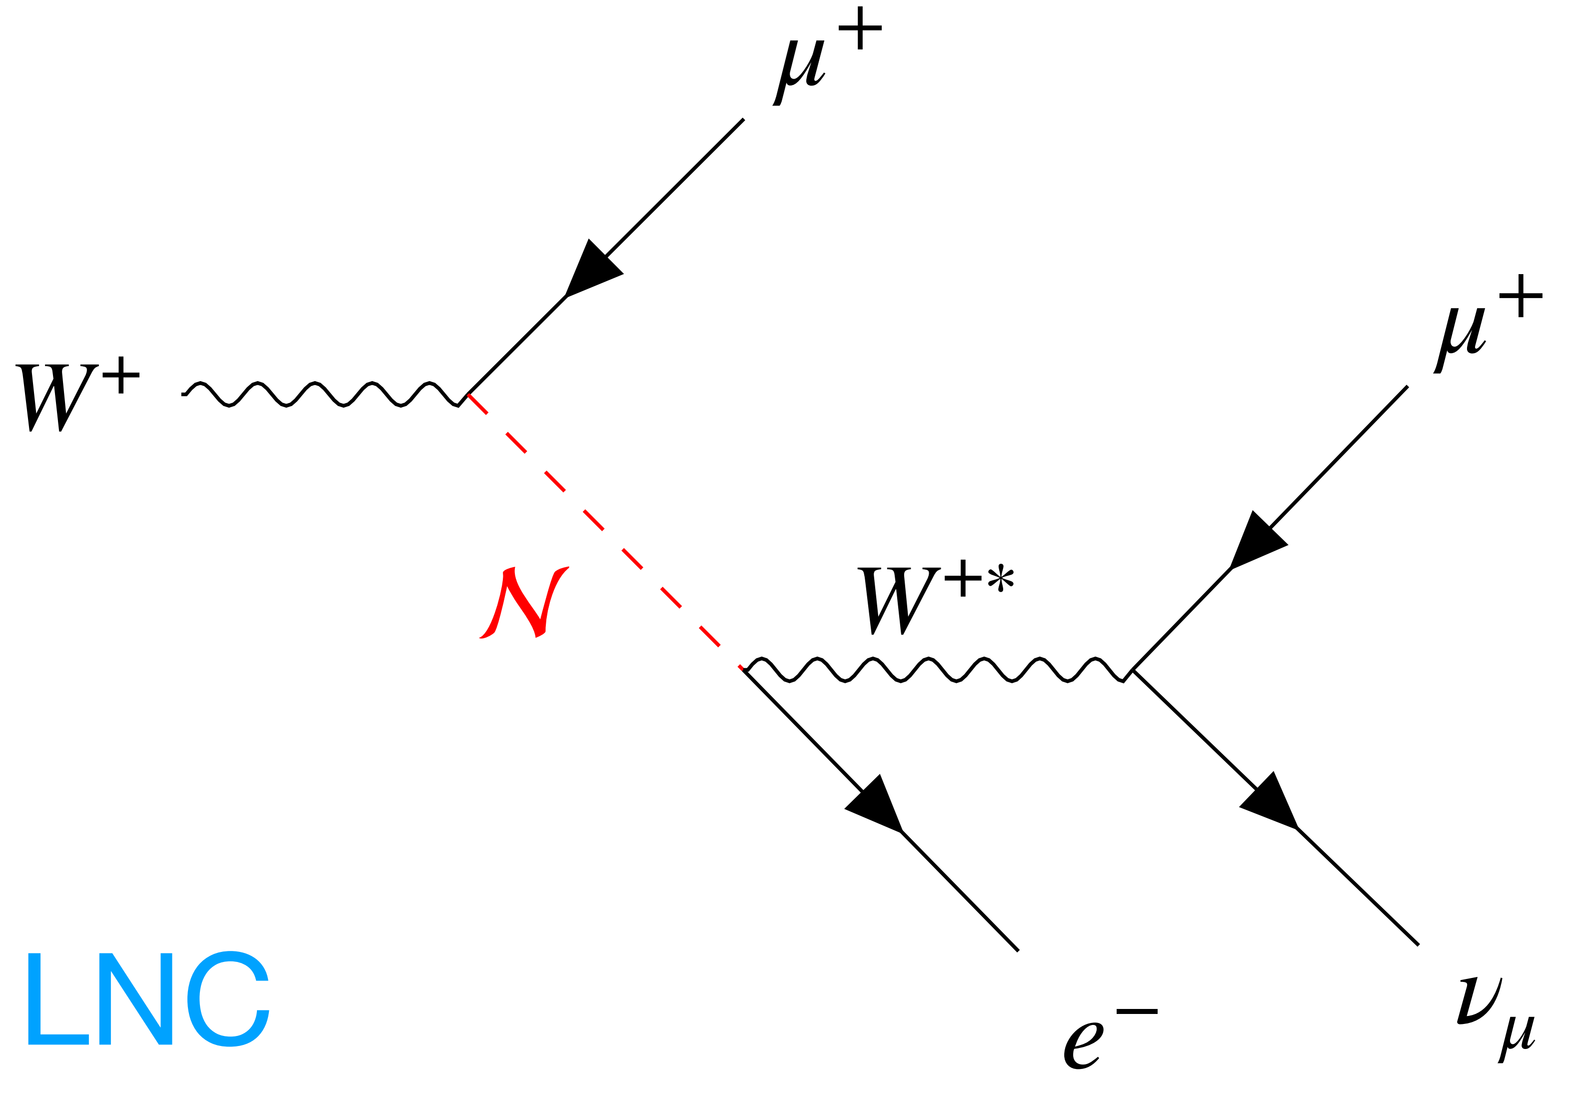

LNC

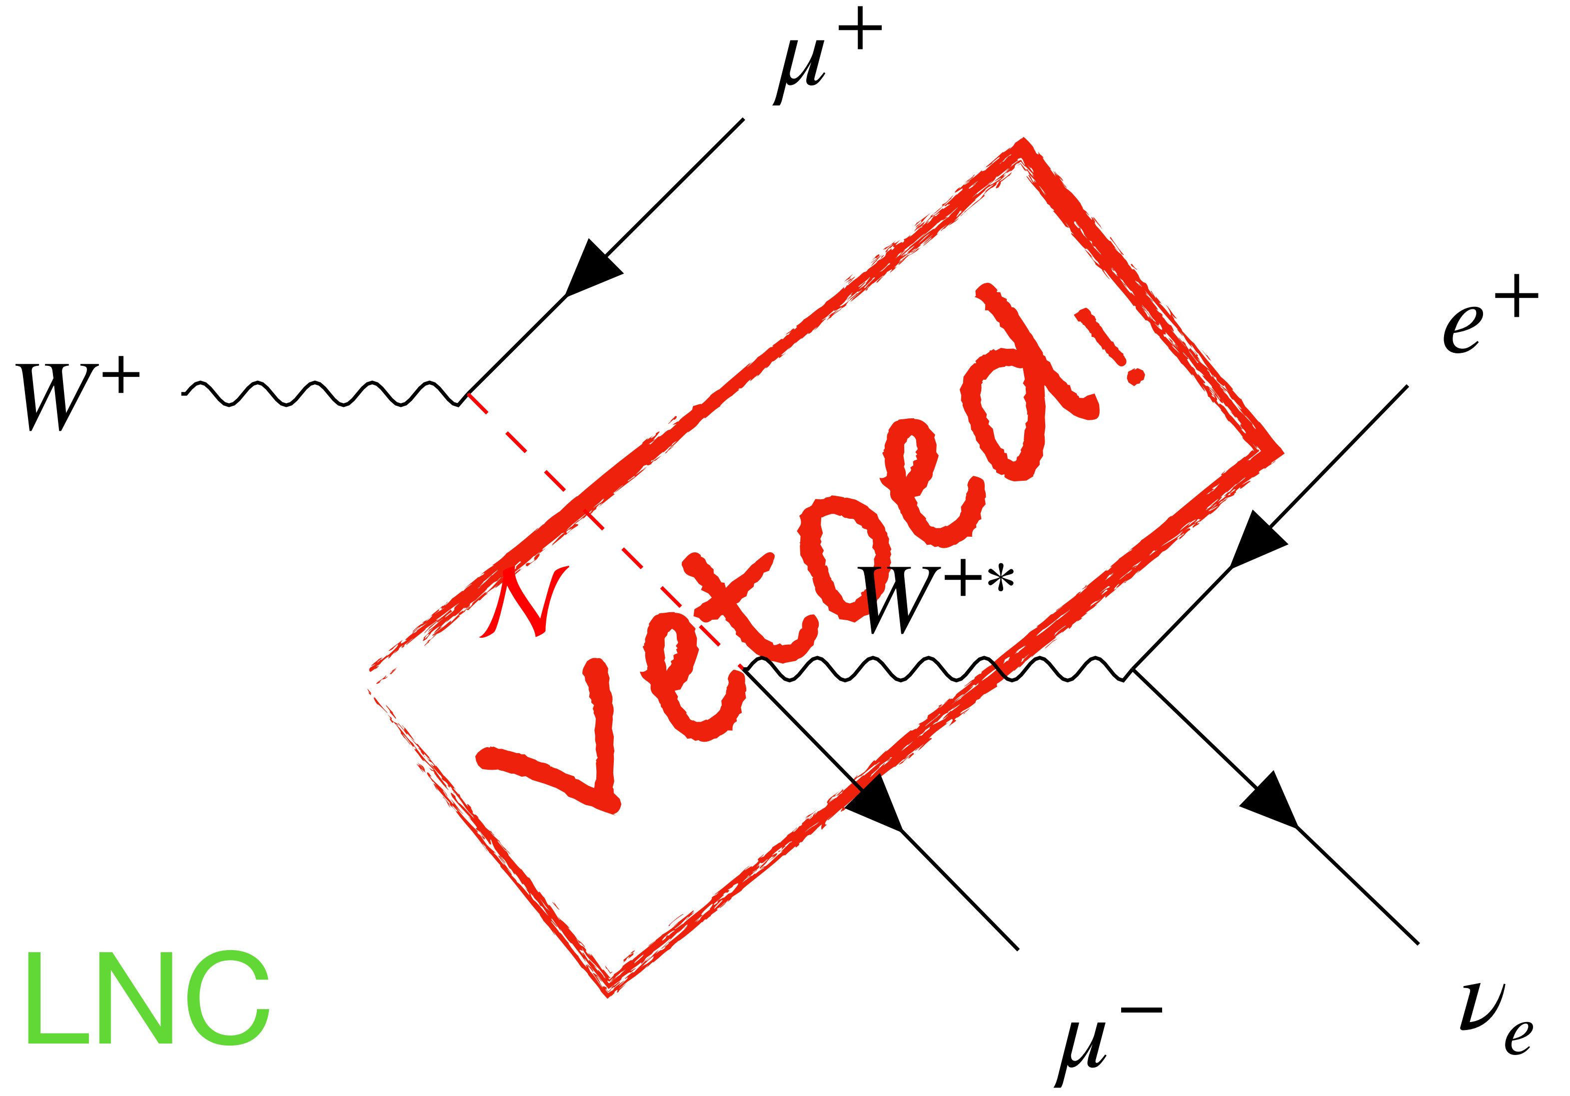

LNC

Blind spot #2

Decays to other channels [see also Abada, Bernal, Losada, Marcano: 1807.10024]

- Turning on more than one mixing angle opens new decay channels, potentially increasing the total HNL width and reducing the branching ratios of the processes being sought for.

- In order to quantify this statement, we need to perform a **reinterpretation** of the prompt trilepton search.

Beyond minimal benchmarks

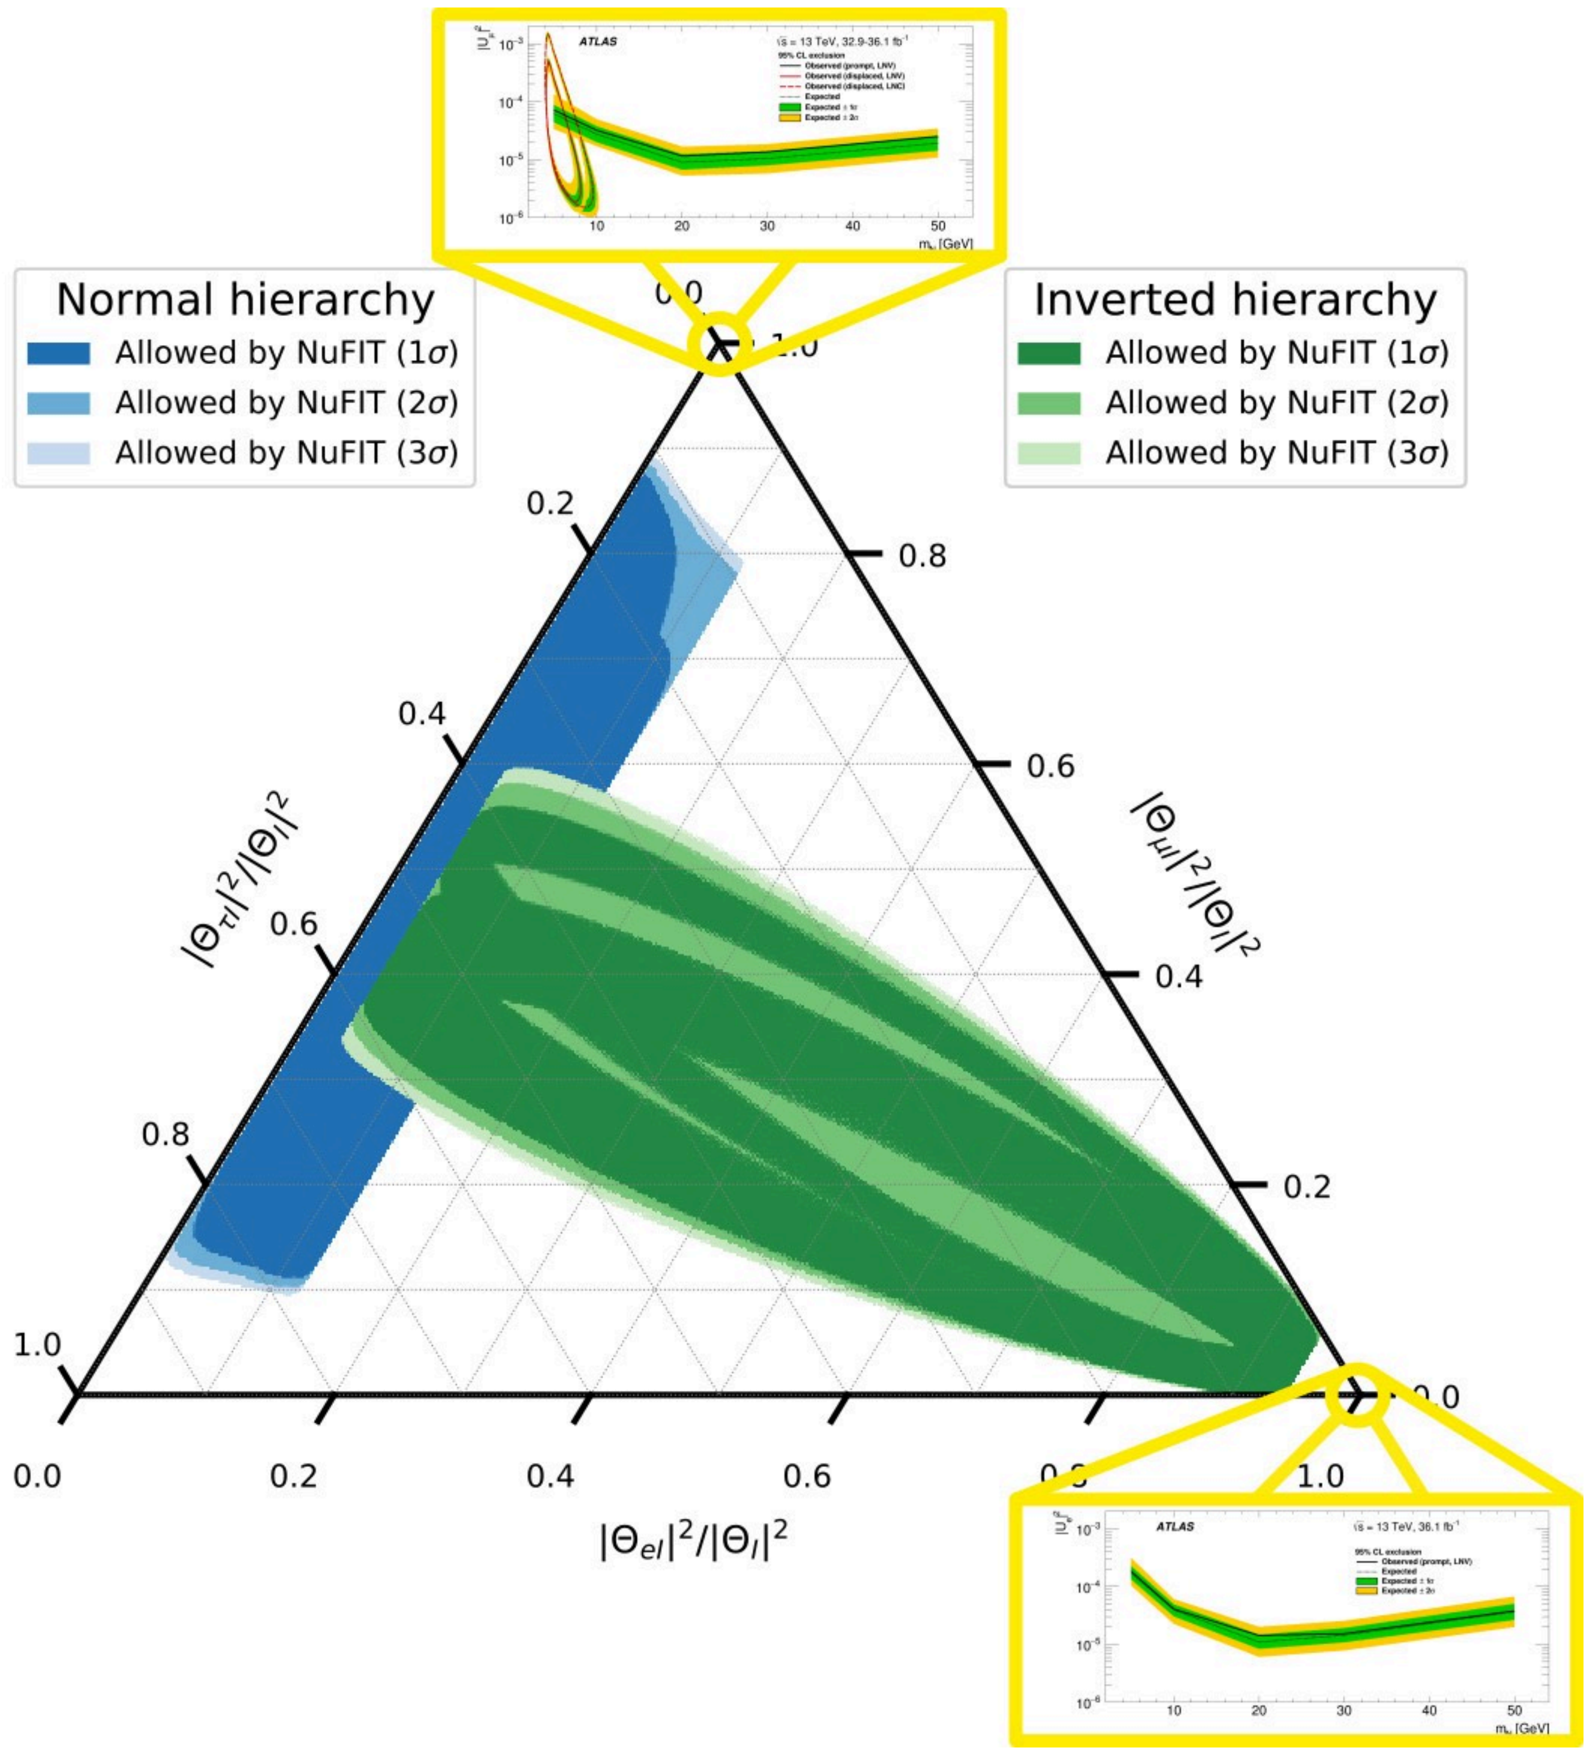

Single-flavour mixing

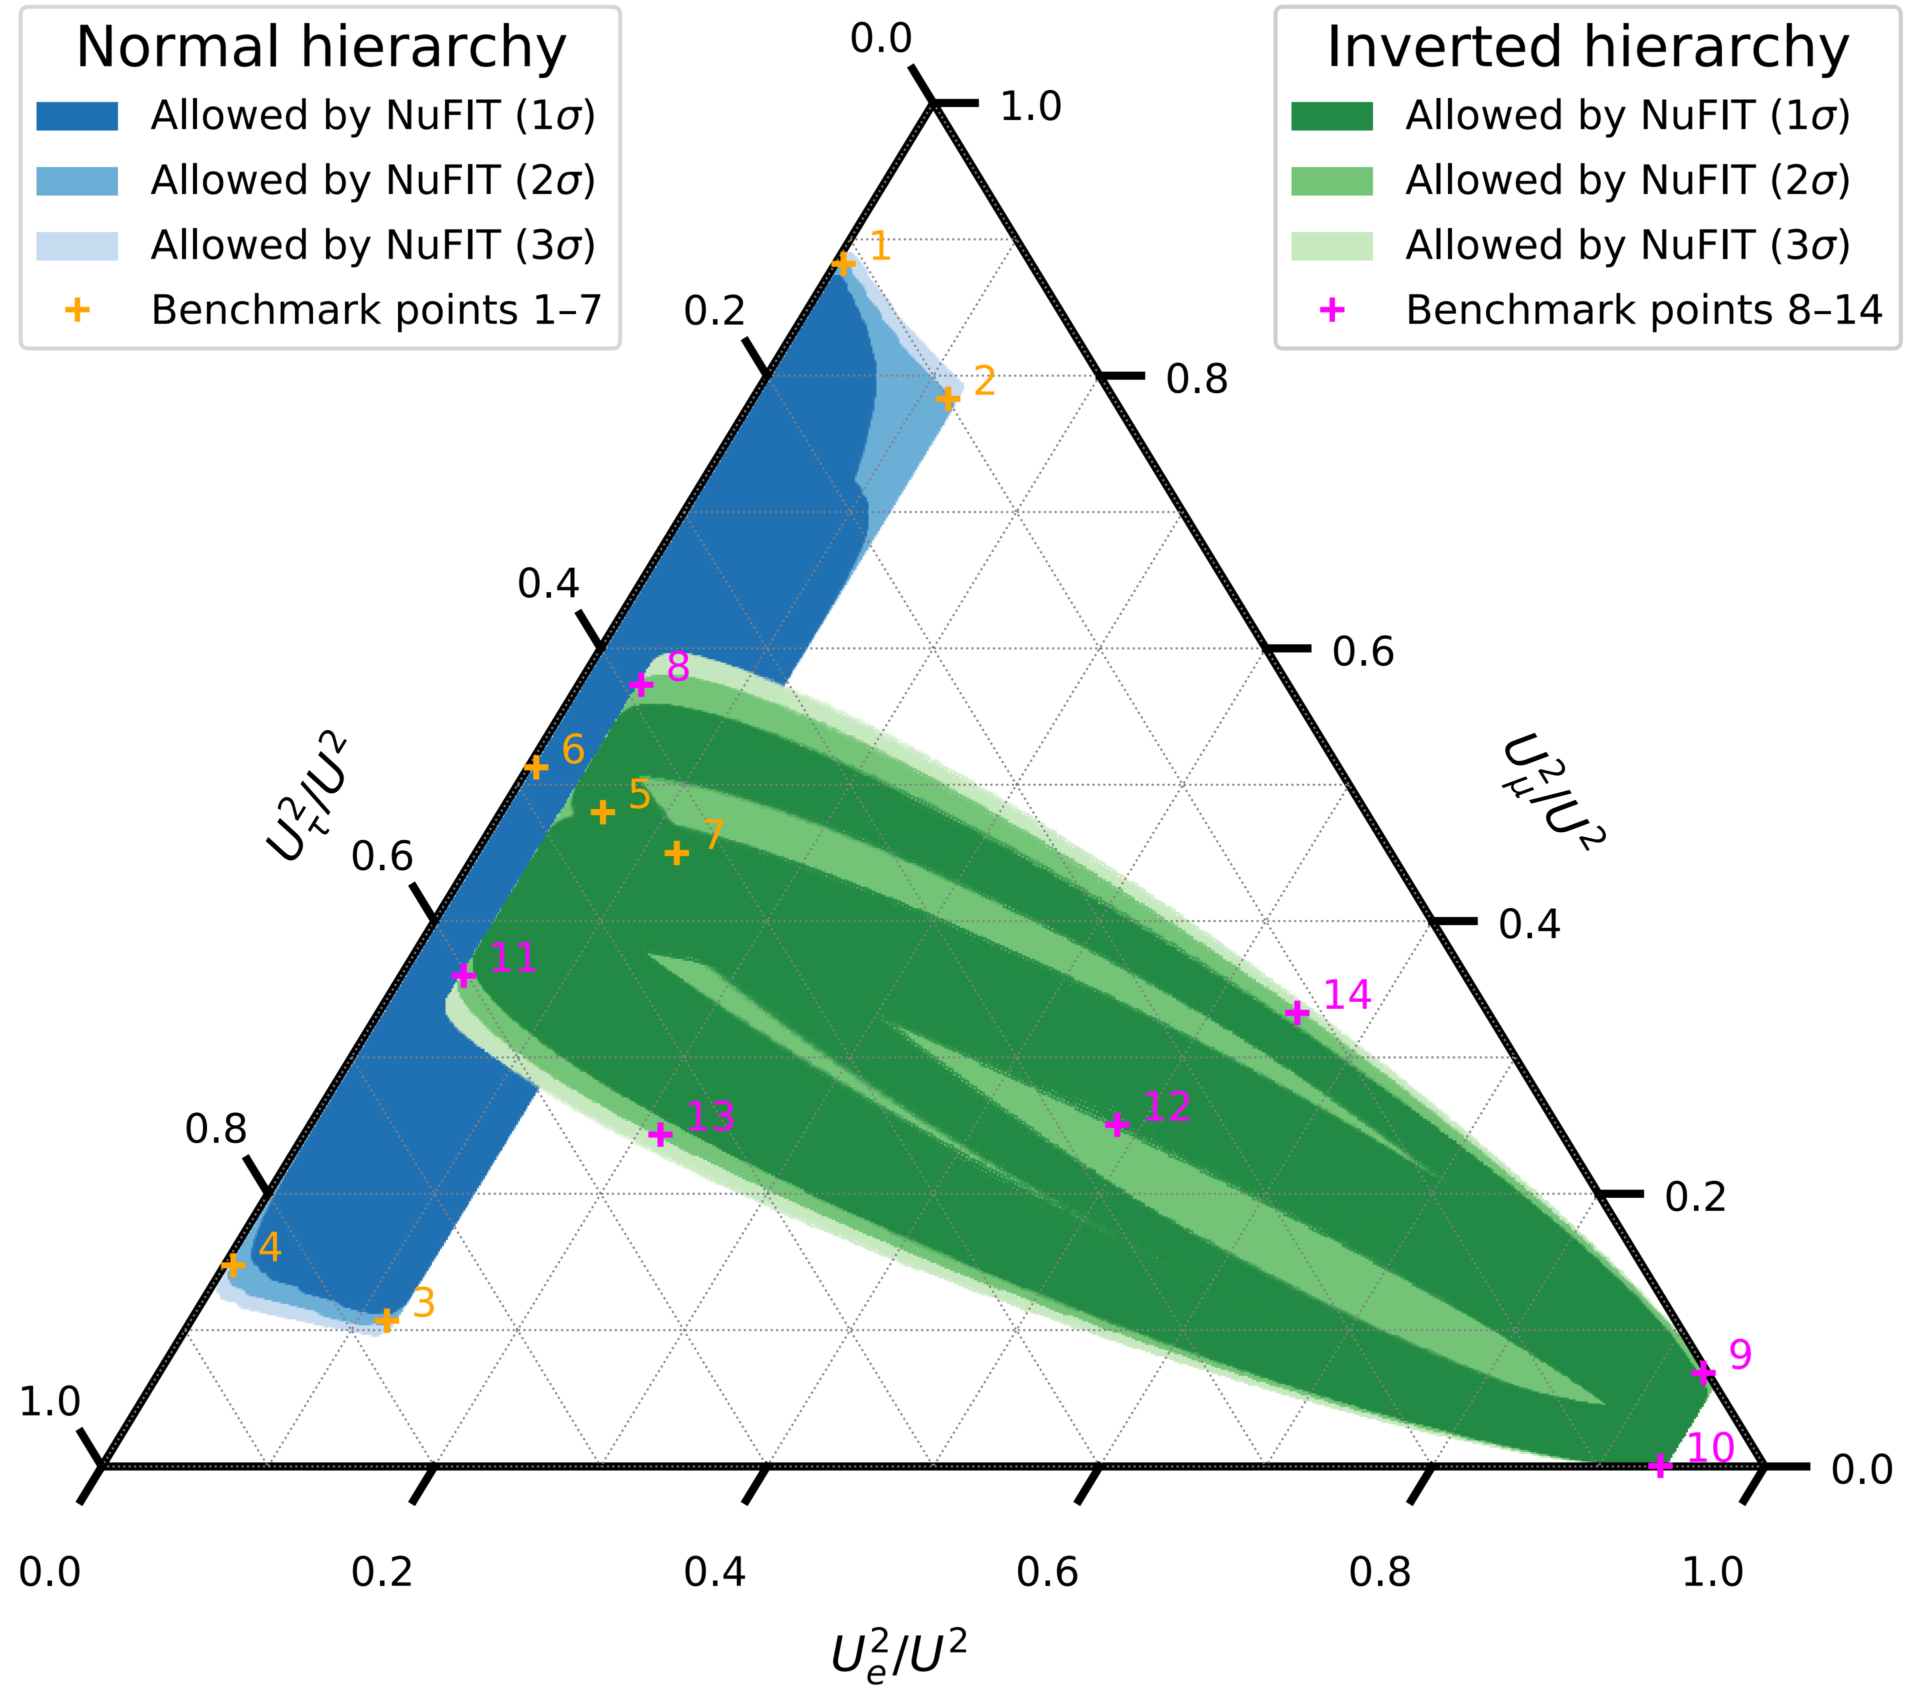

Non-minimal benchmarks

Reinterpretation procedure

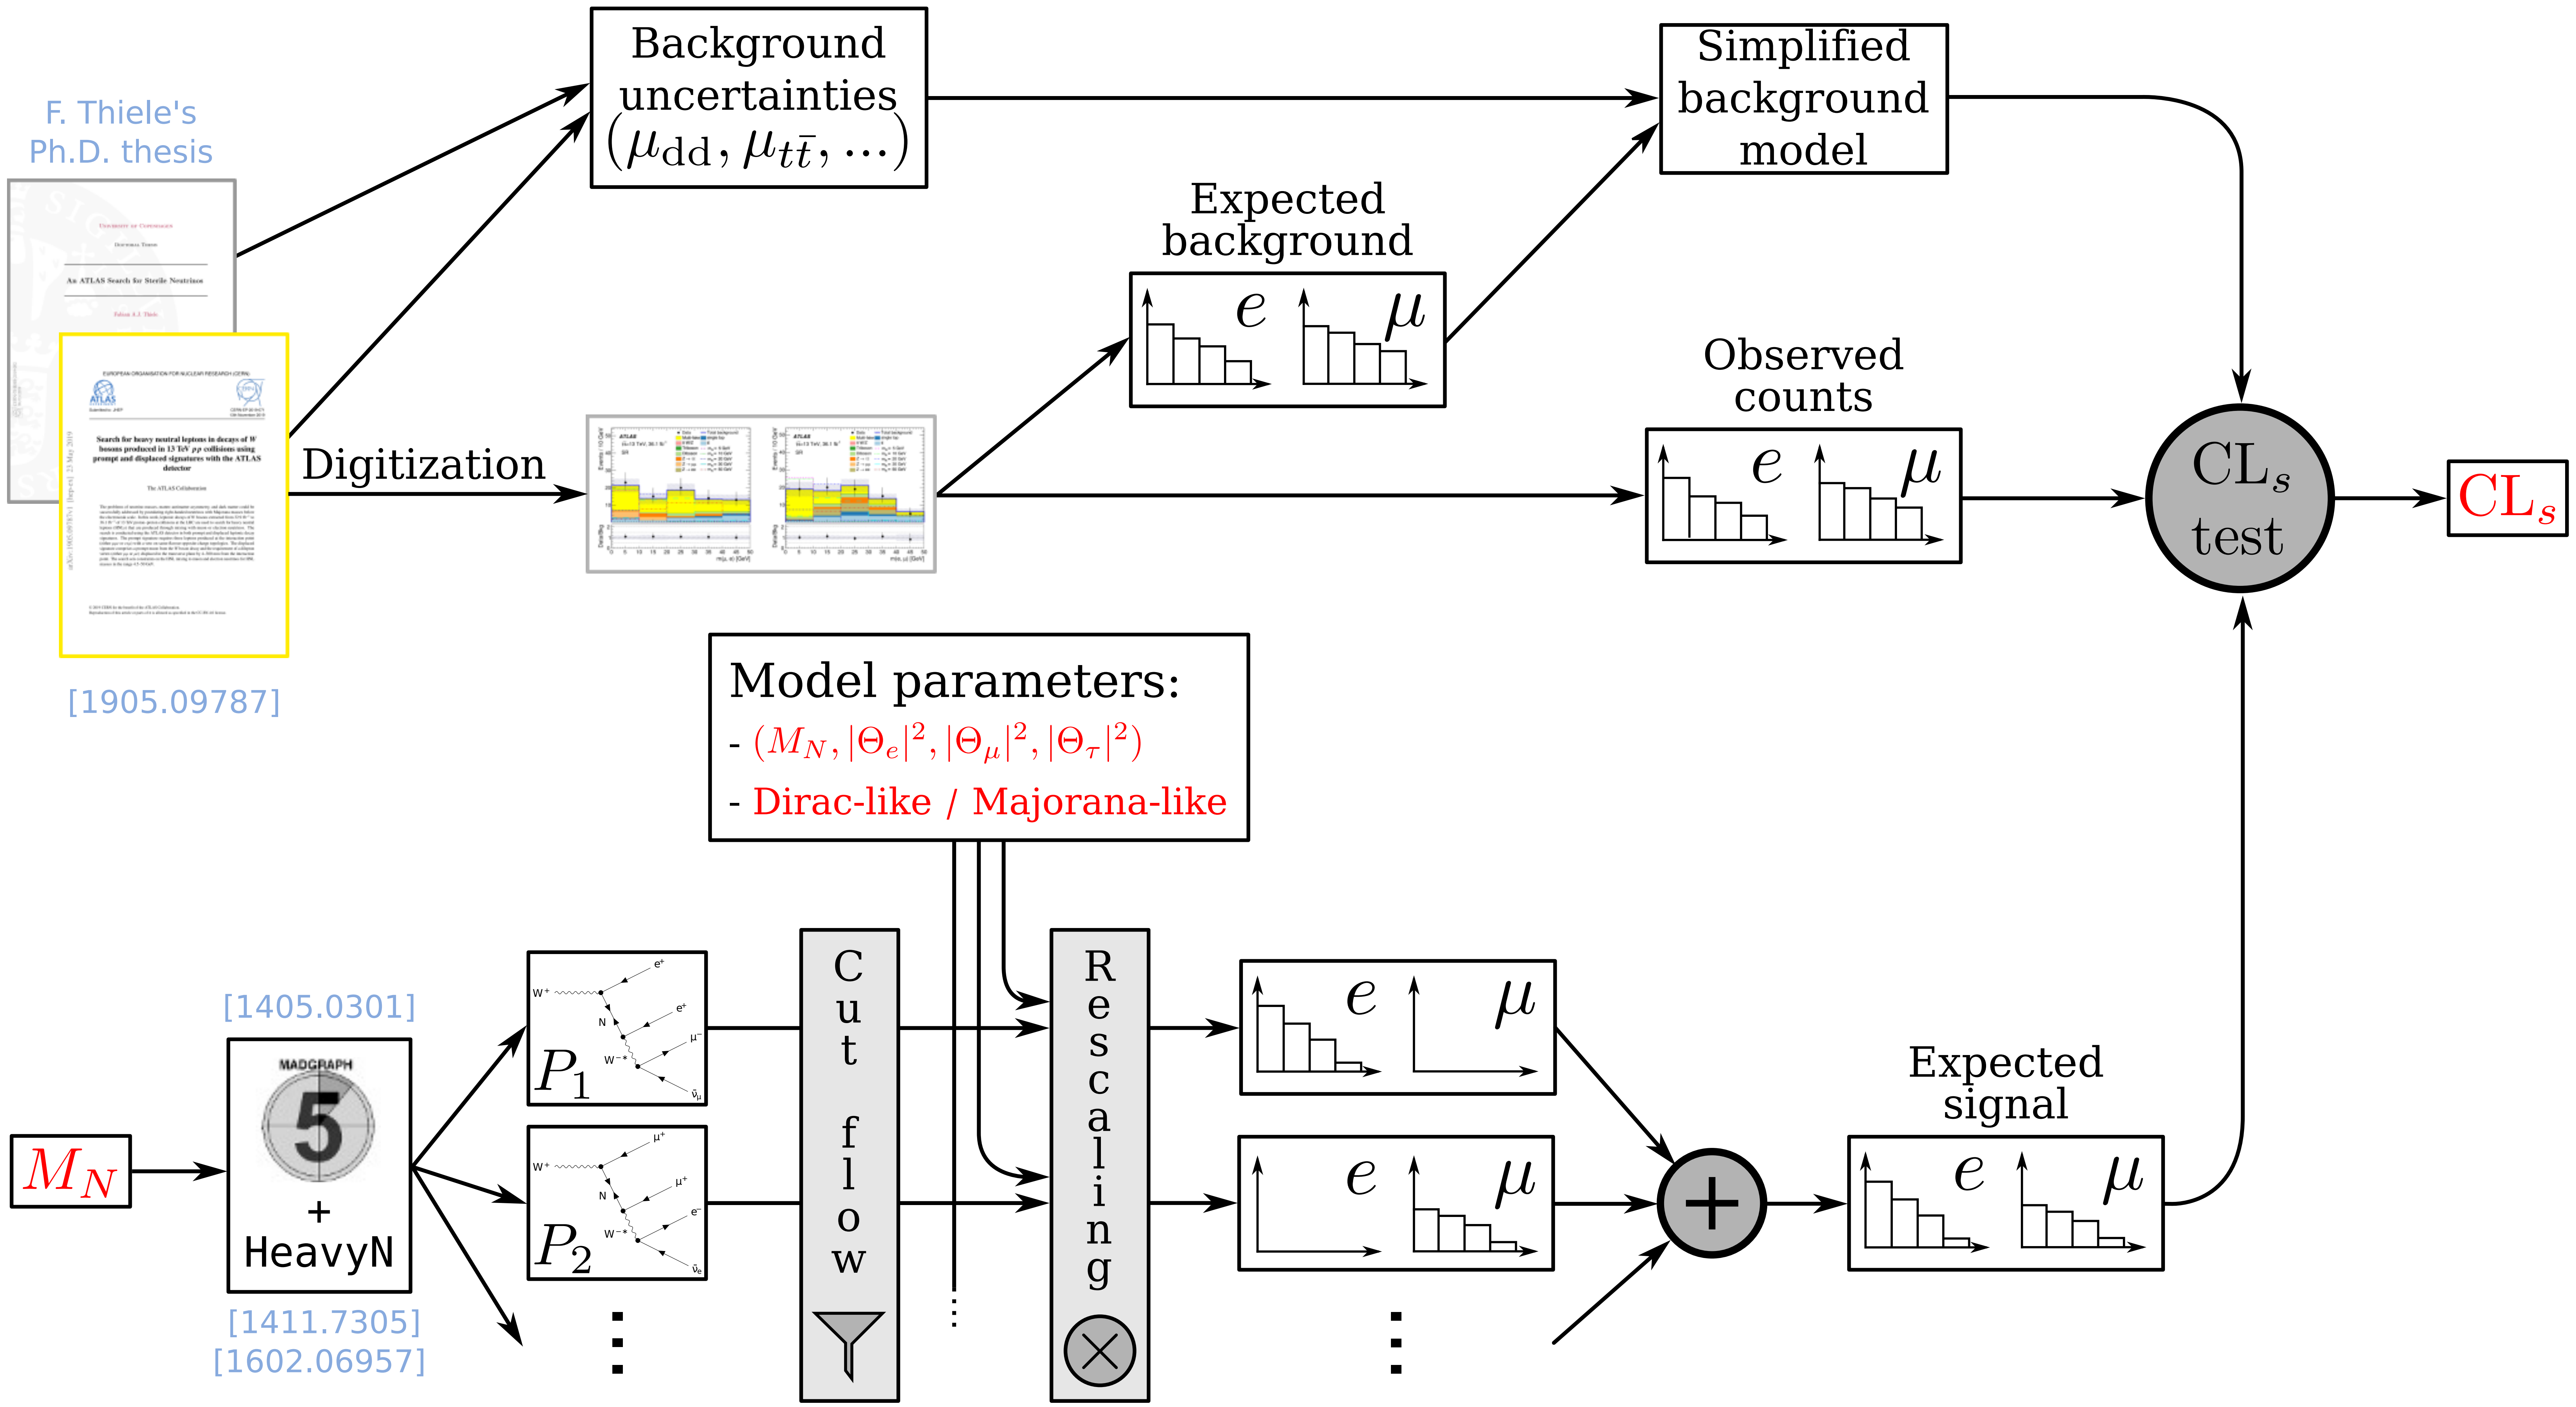

Signal efficiency validation

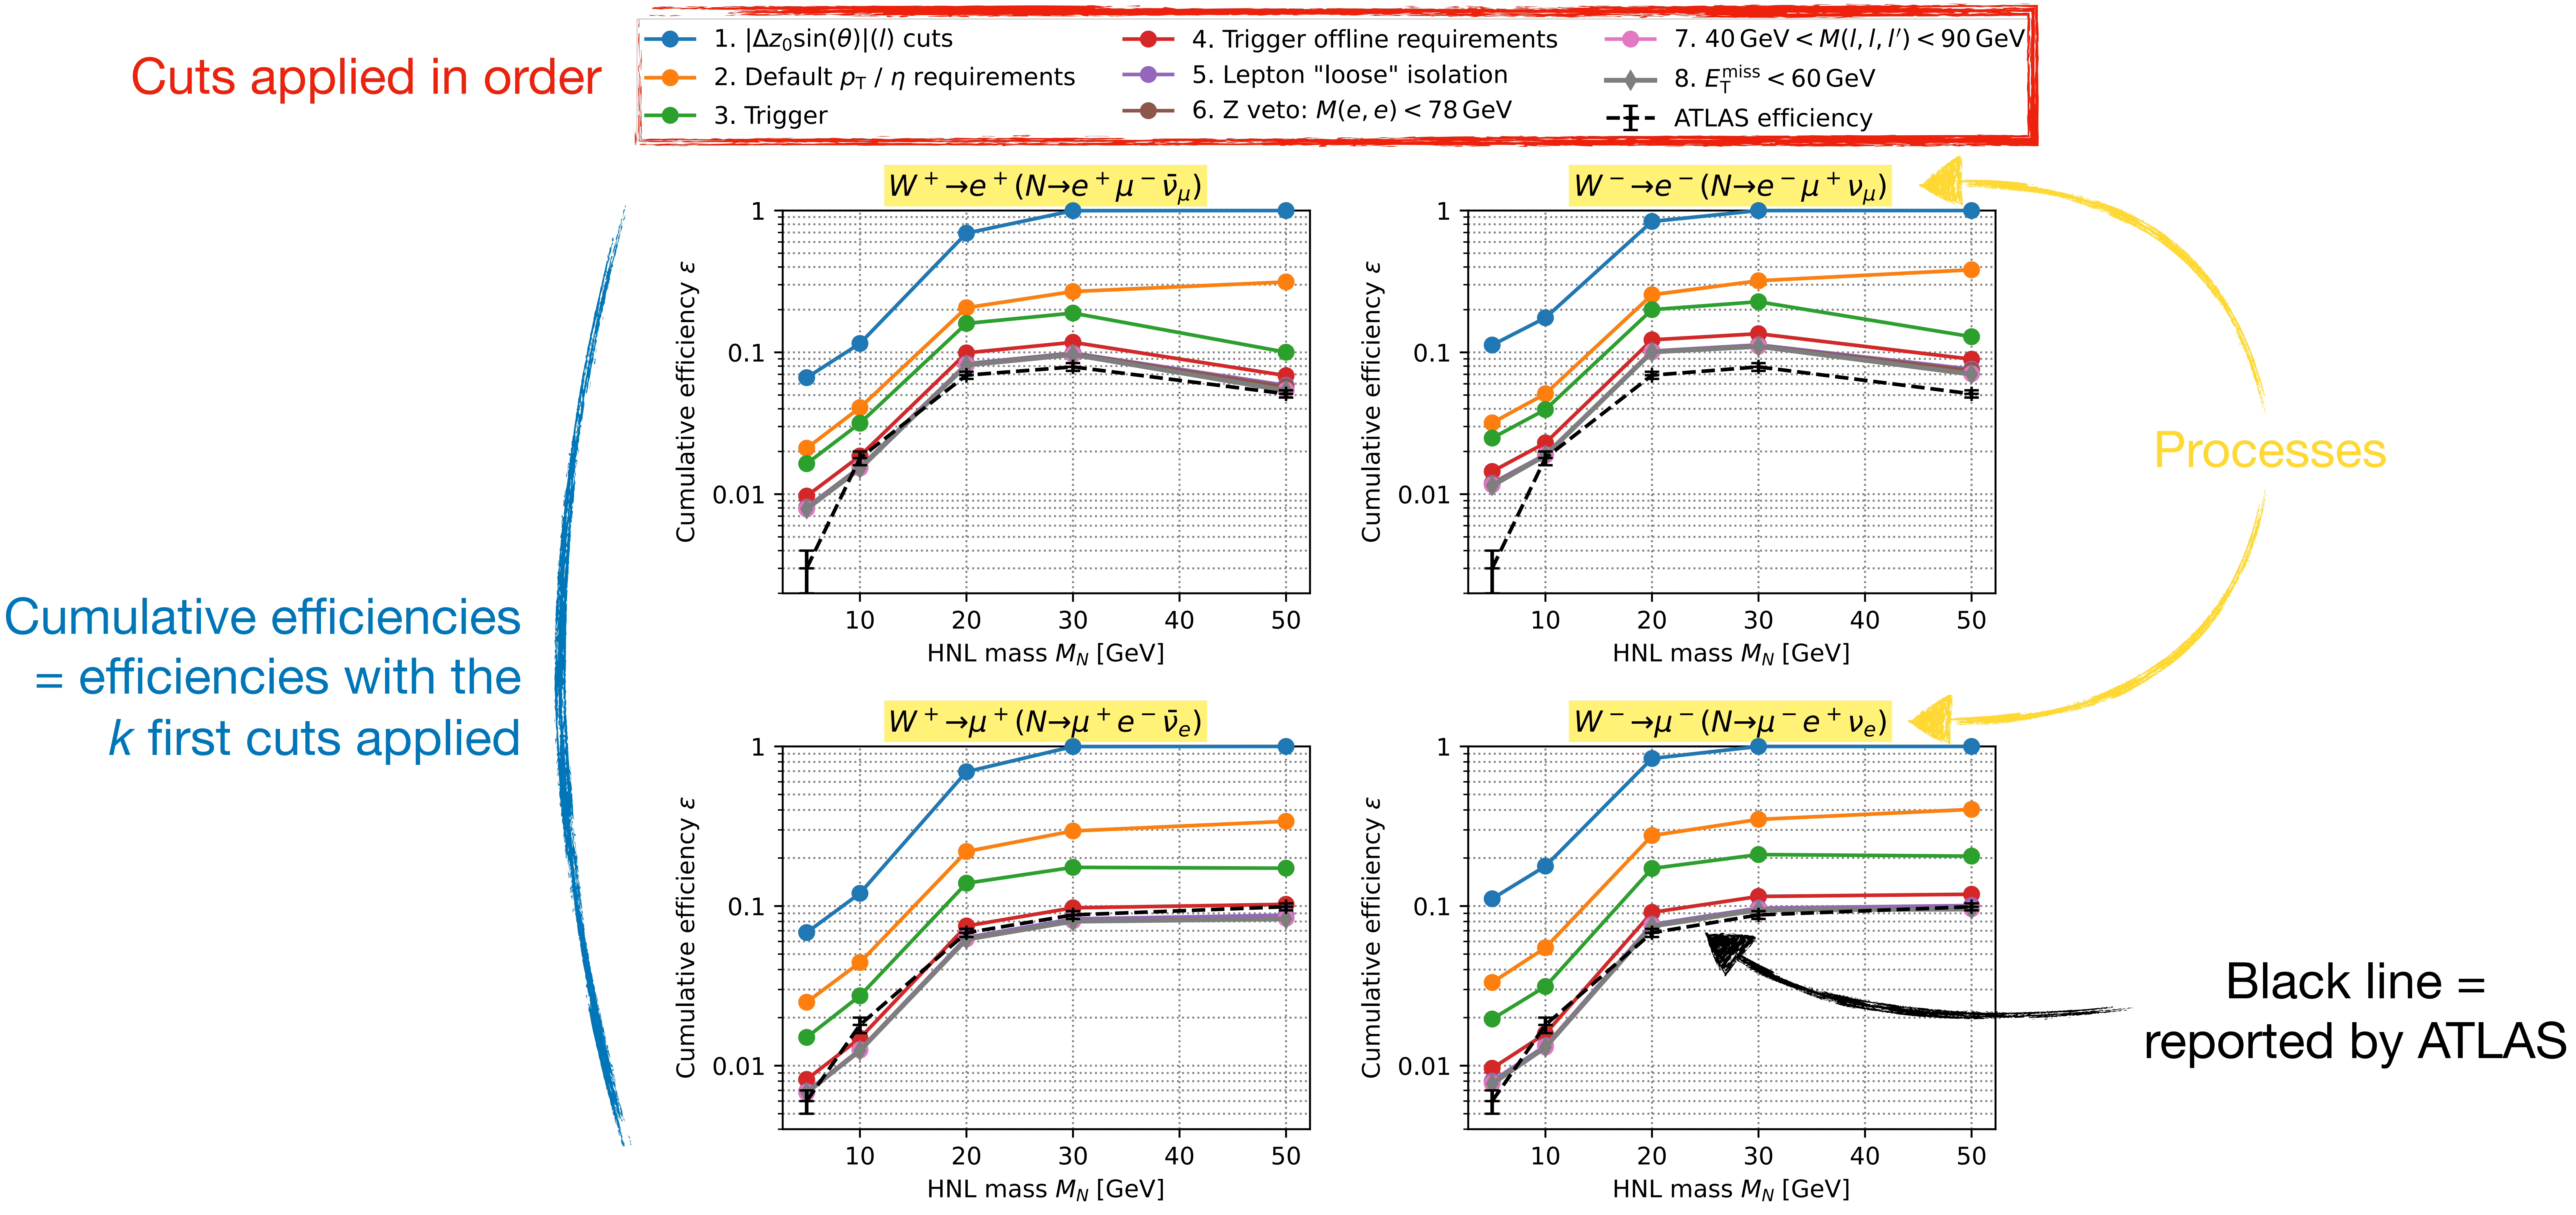

CL_s method validation

- Use the same numerical values as ATLAS
- Tests our modelling of background and systematic uncertainties
- Far from perfect, but **good enough** for this reinterpretation
  (ours/theirs = 0.64 (worst 0.42))
- Individual background components can be extracted from the available data, but not their uncertainties
  Difficult to estimate their correlations

Both signal regions fitted simultaneously

Gaussian constraint term assuming 1 dominant background component

Poisson statistics in each bin i

$$\mathcal{L}(x|H_{s+b}) = \mathcal{P}(\mu_{\mathrm{tot}}|\mathcal{N}(1,\sigma_{\mathrm{tot}})) \times \prod_{i=1}^{10} \mathcal{P}(x_i|\mathrm{Pois}(\mu_{\mathrm{tot}} b_i + s_i))$$

Likelihood

Background normalisation best fit

where $\mu_{\mathrm{tot}} = \dfrac{\sum_i (x_i - s_i)}{\sum_i b_i}$

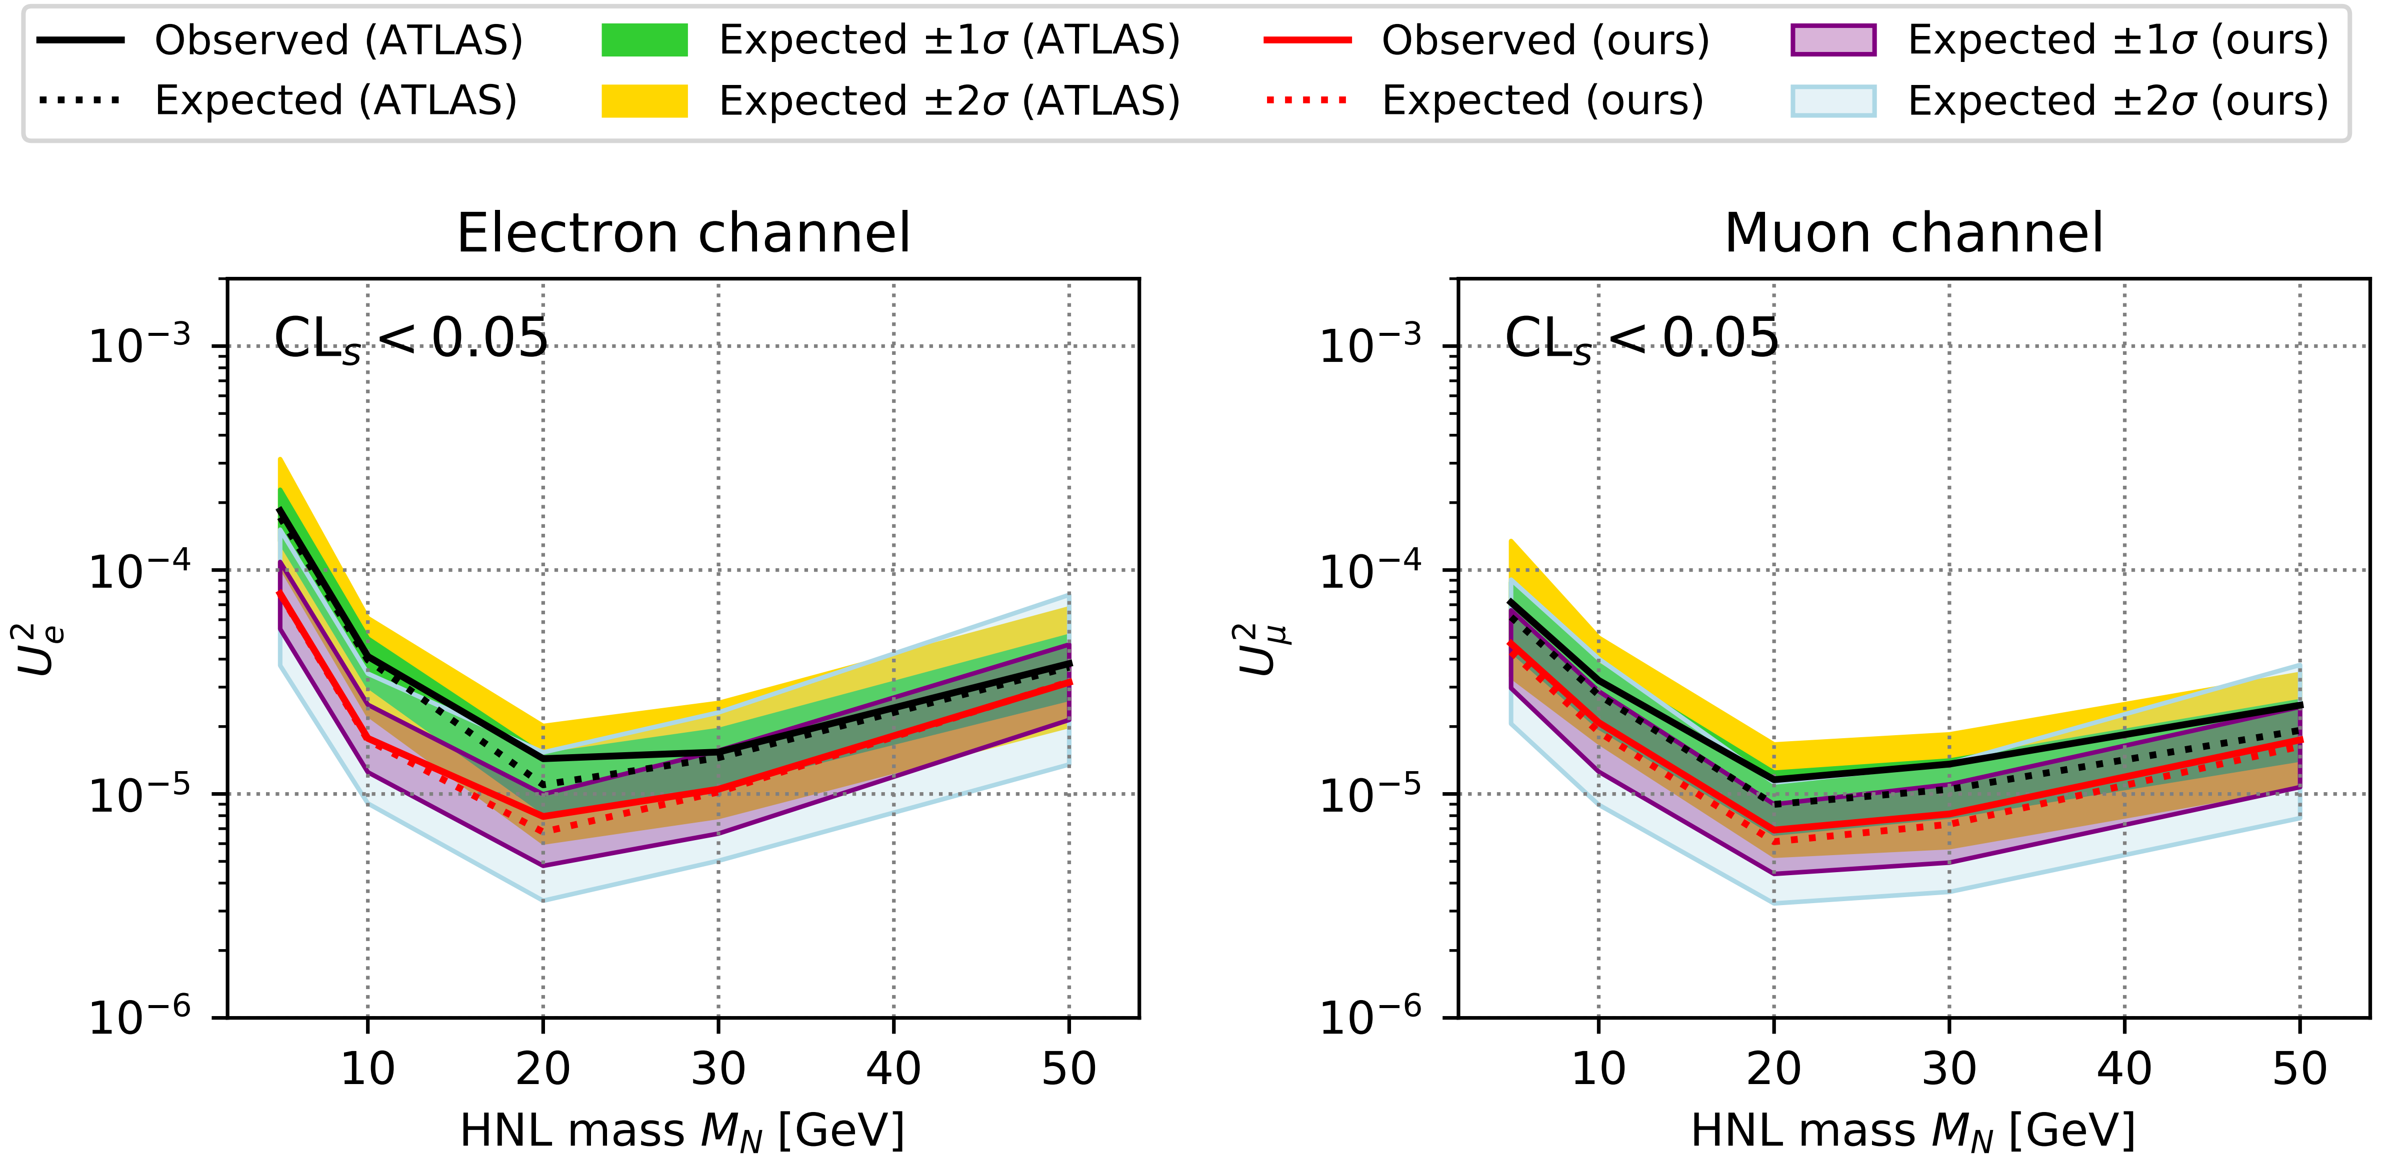

Reinterpretation of limits [Tastet, Ruchayskiy, Timiryasov: 2107.12980]

How to read the results → Decompose 4d parameter space into 2d + 2d

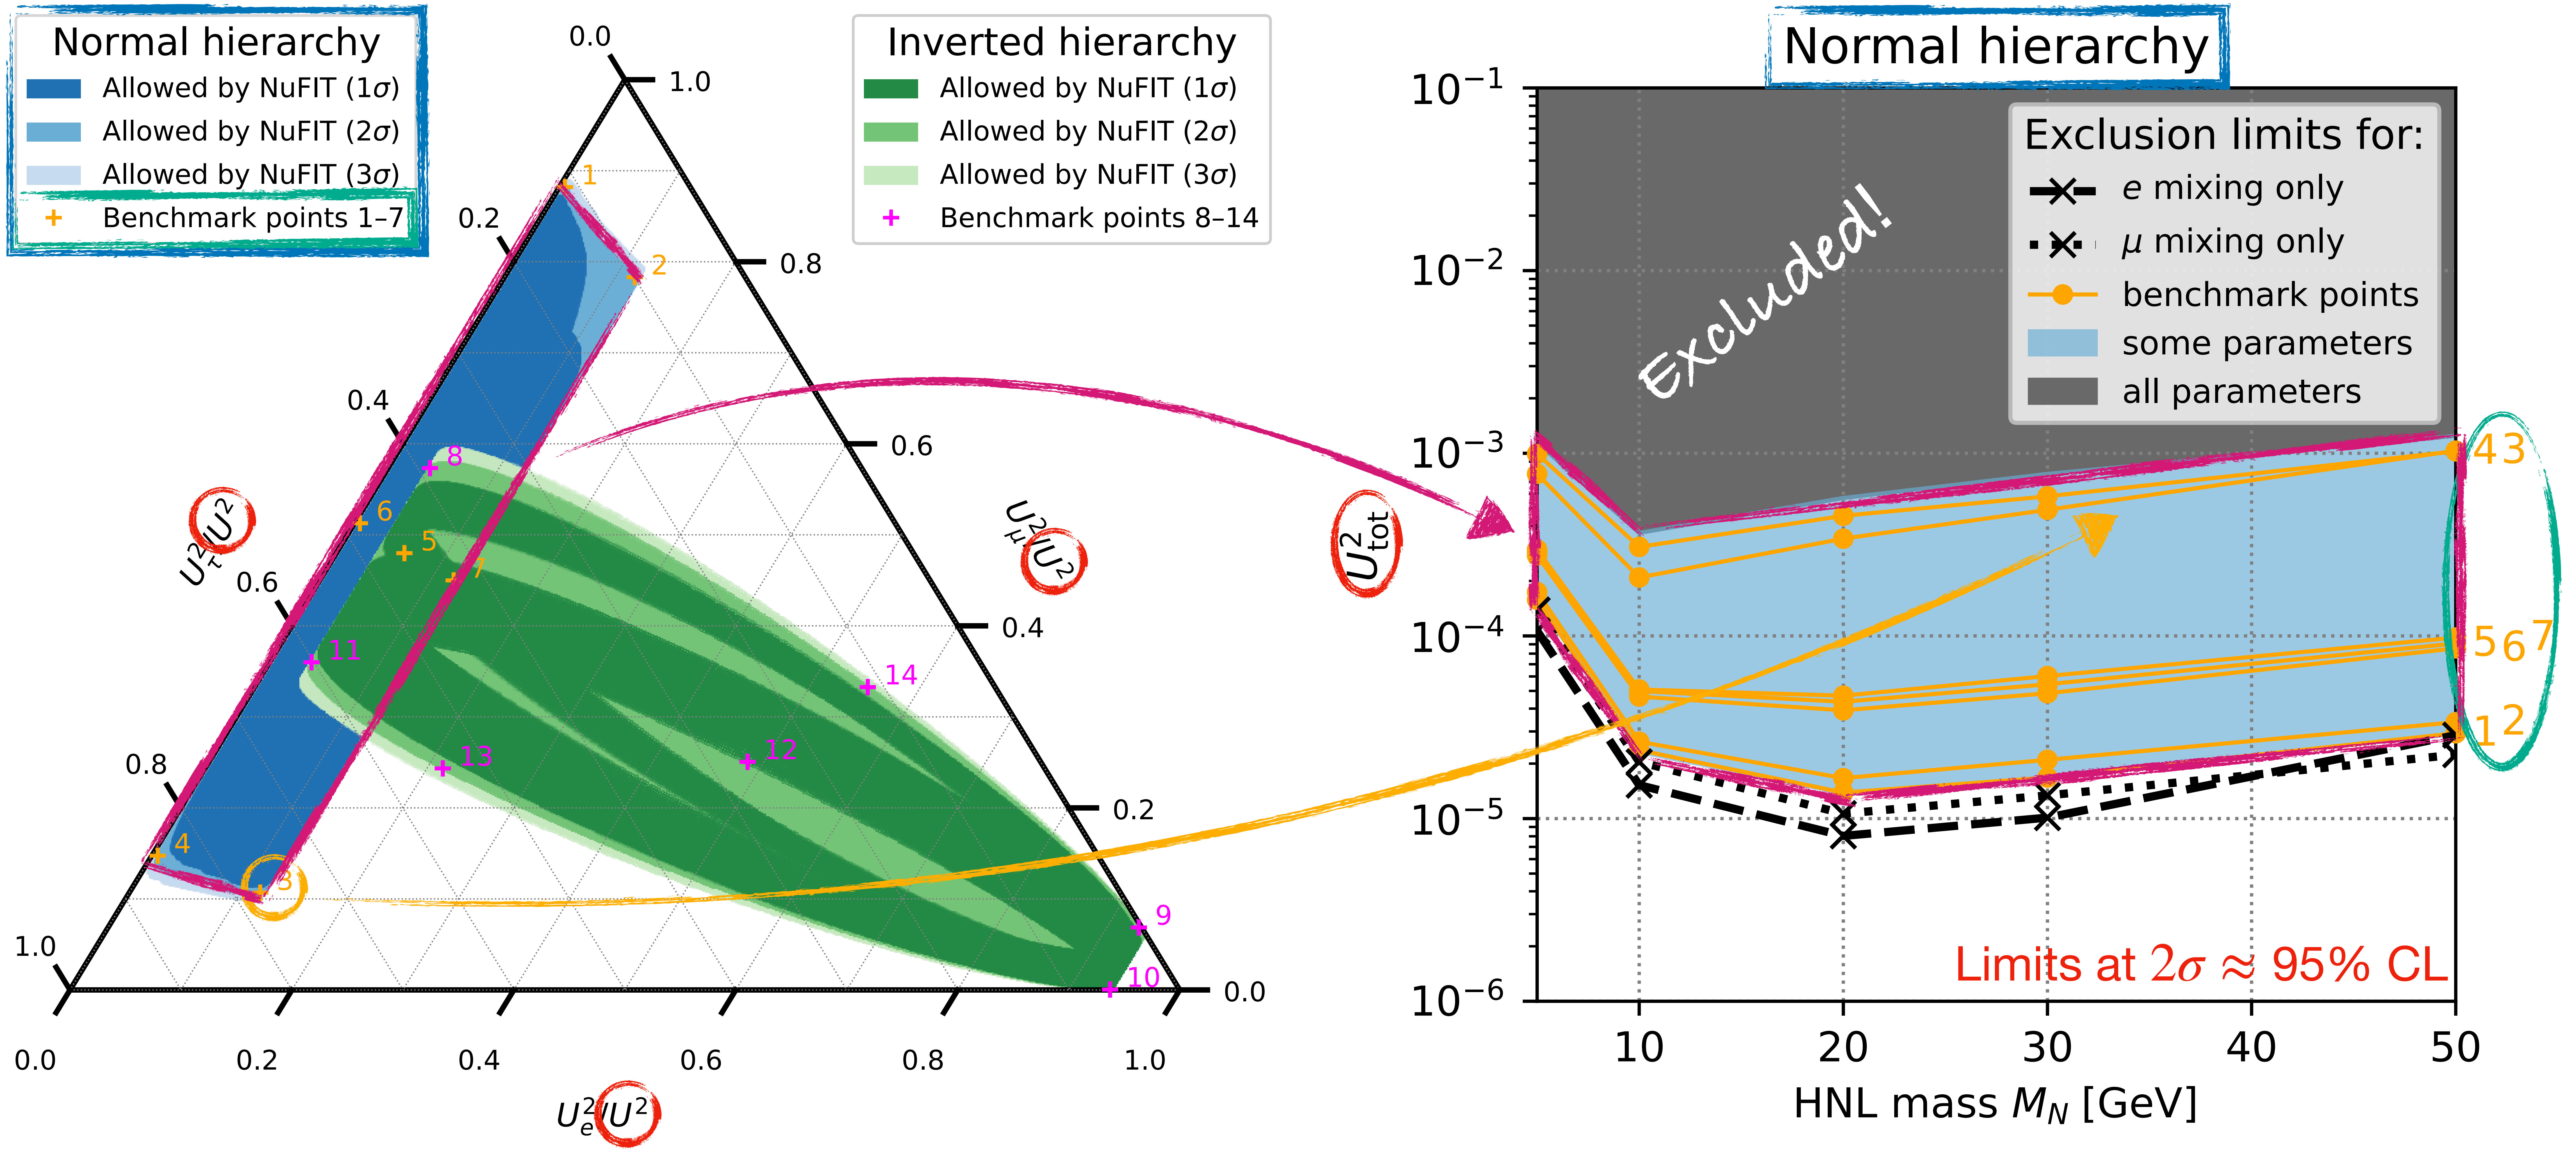

Reinterpretation of limits [Tastet, Ruchayskiy, Timiryasov: 2107.12980]

Majorana-like HNLs

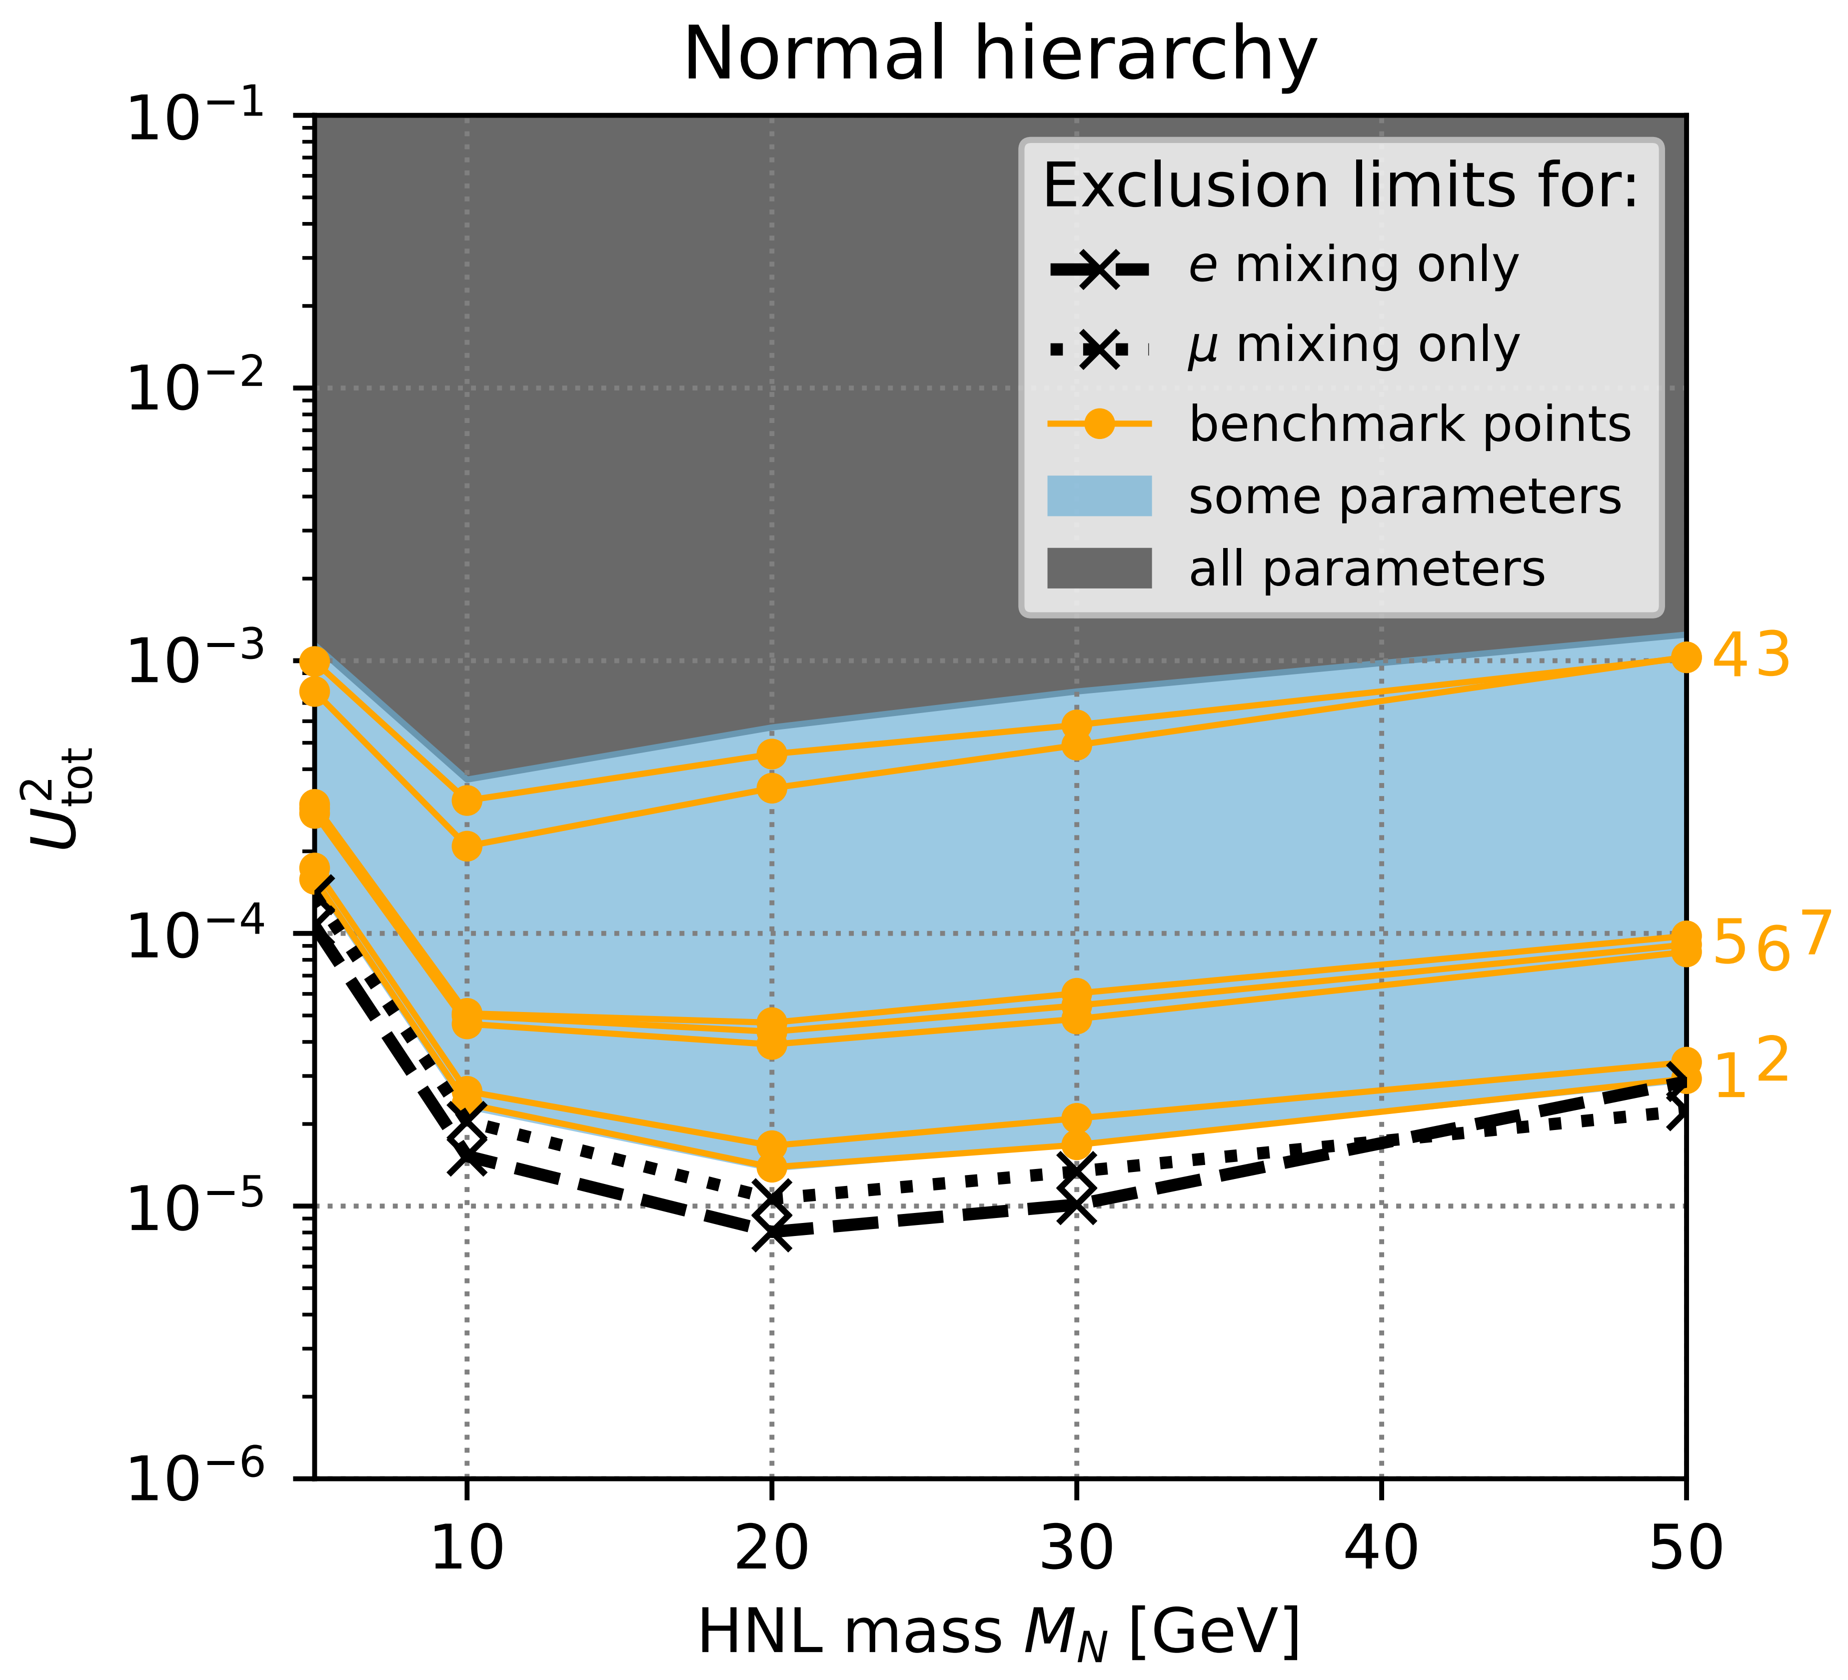

- Total mixing U^2_{tot} used for consistency
- Recast limits almost always weaker than single-flavour (up to 1 order of magnitude)
- Weakest limits ↔ largest τ mixing
  Smaller BR in signal channels
  Many HNLs produced in other channels!
  ⇒ Search for τ's to close the blind spots!
- Marginalise over allowed combinations of mixing angles to set an absolute limit
- Similar results for the inverted hierarchy

Reinterpretation of limits [Tastet, Ruchayskiy, Timiryasov: 2107.12980]

Dirac-like HNLs

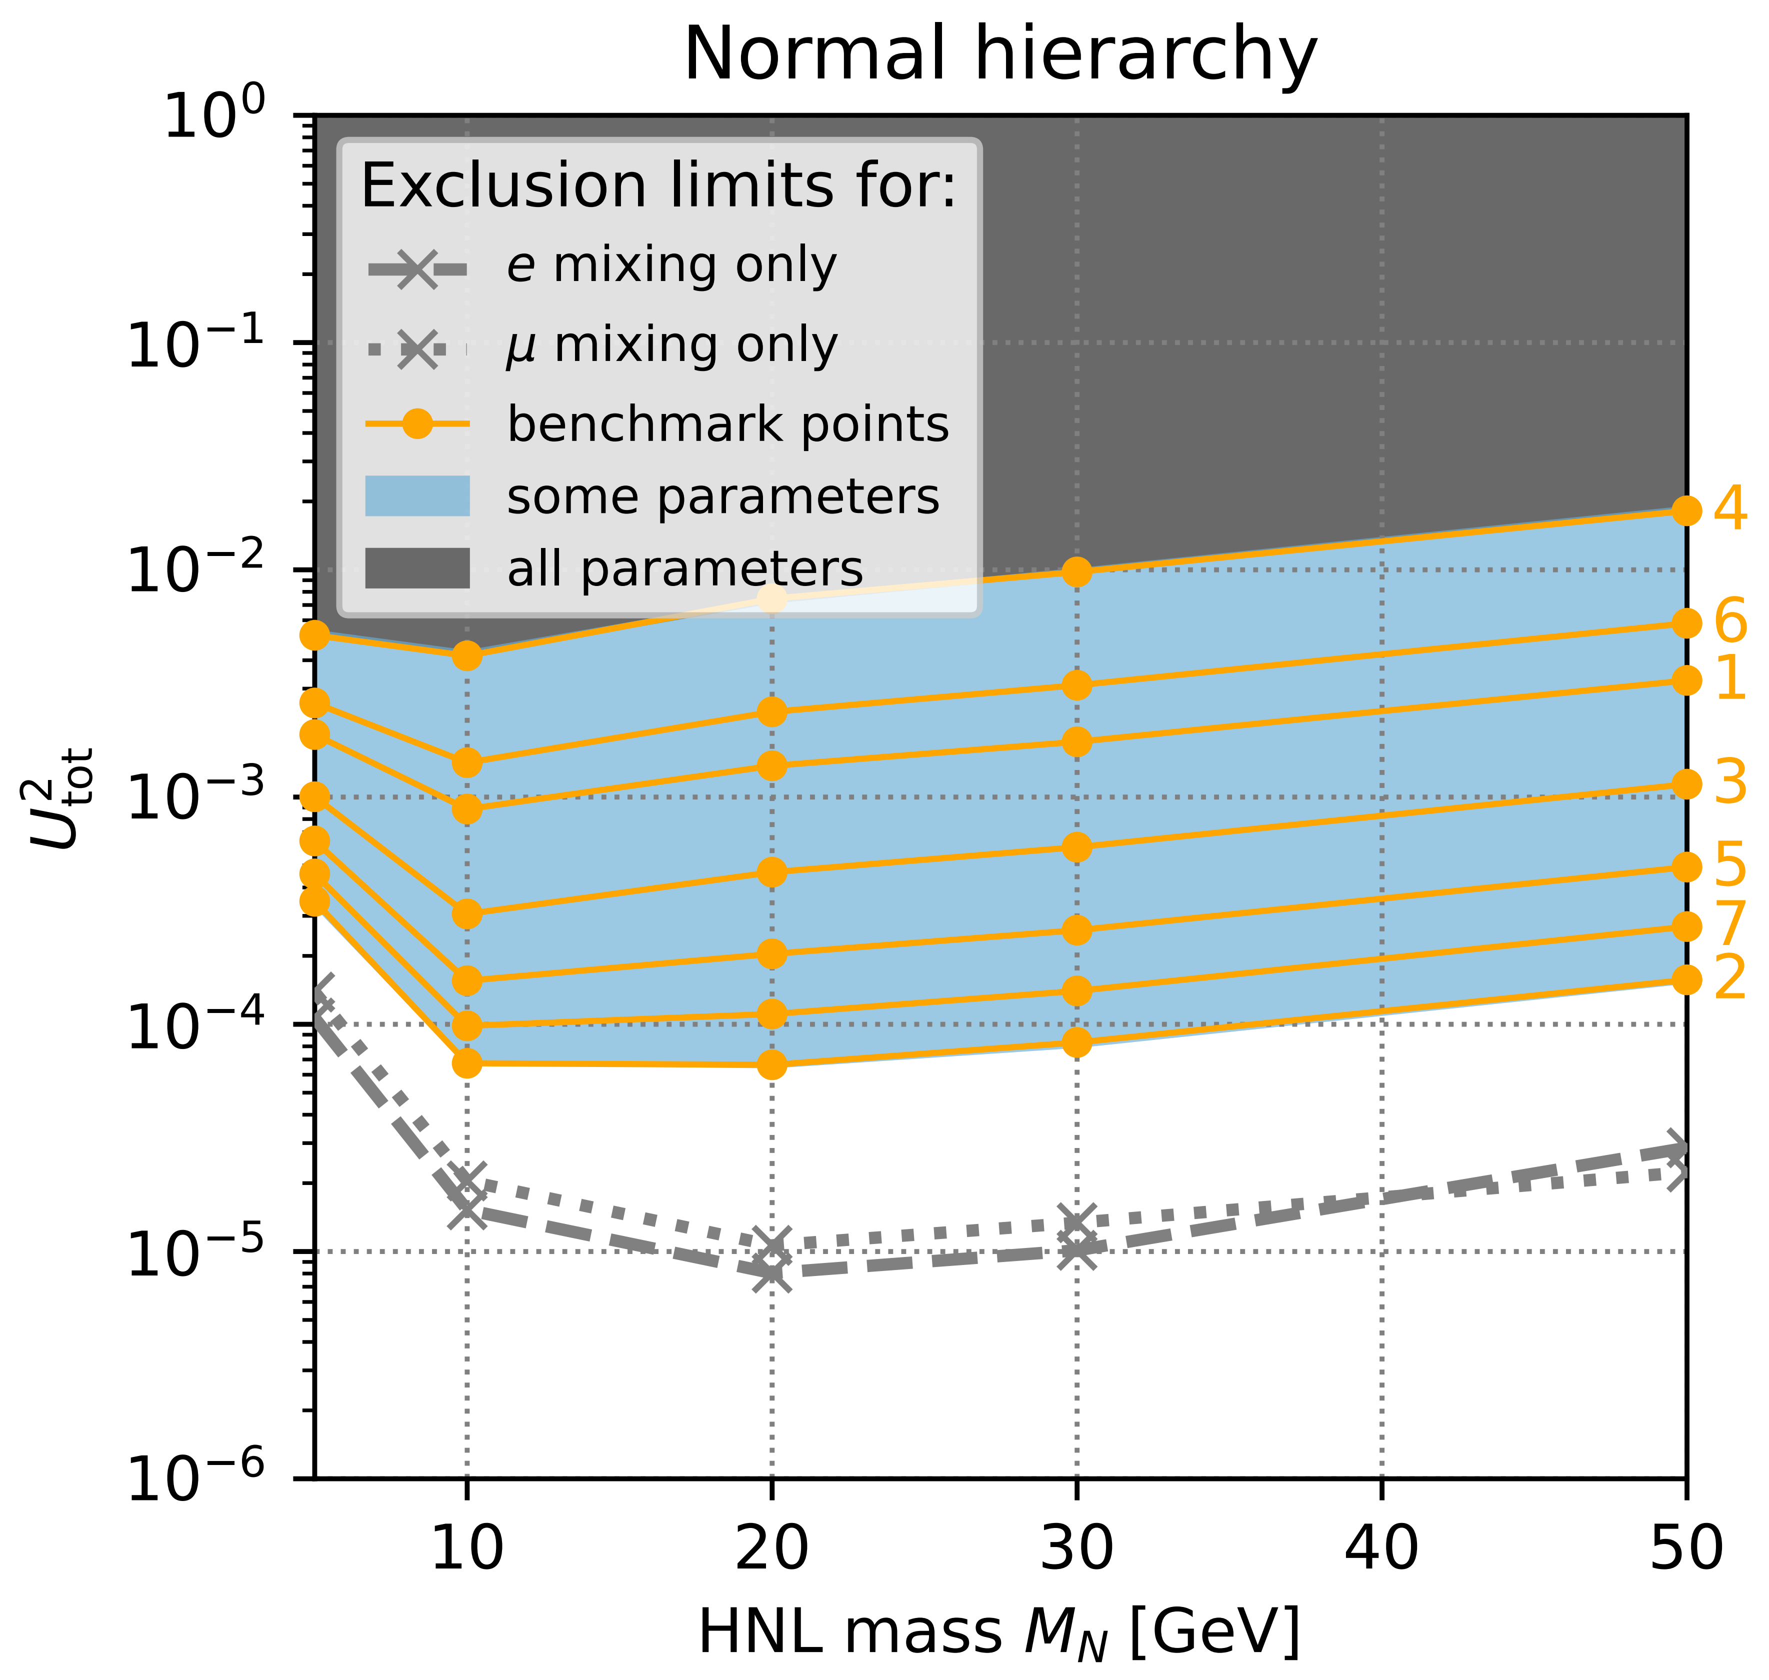

- Previously: no sensitivity for single-flavour
- Limits weaker by up to 3 orders of magnitude vs. original benchmarks (weakest limits when a mixing is suppressed)
- There exist allowed models 3 orders of magnitude above the reported limit
- Increased variance between benchmarks $\Longrightarrow$ weaker marginalised limit

ATLAS reinterpretation

Inverted hierarchy

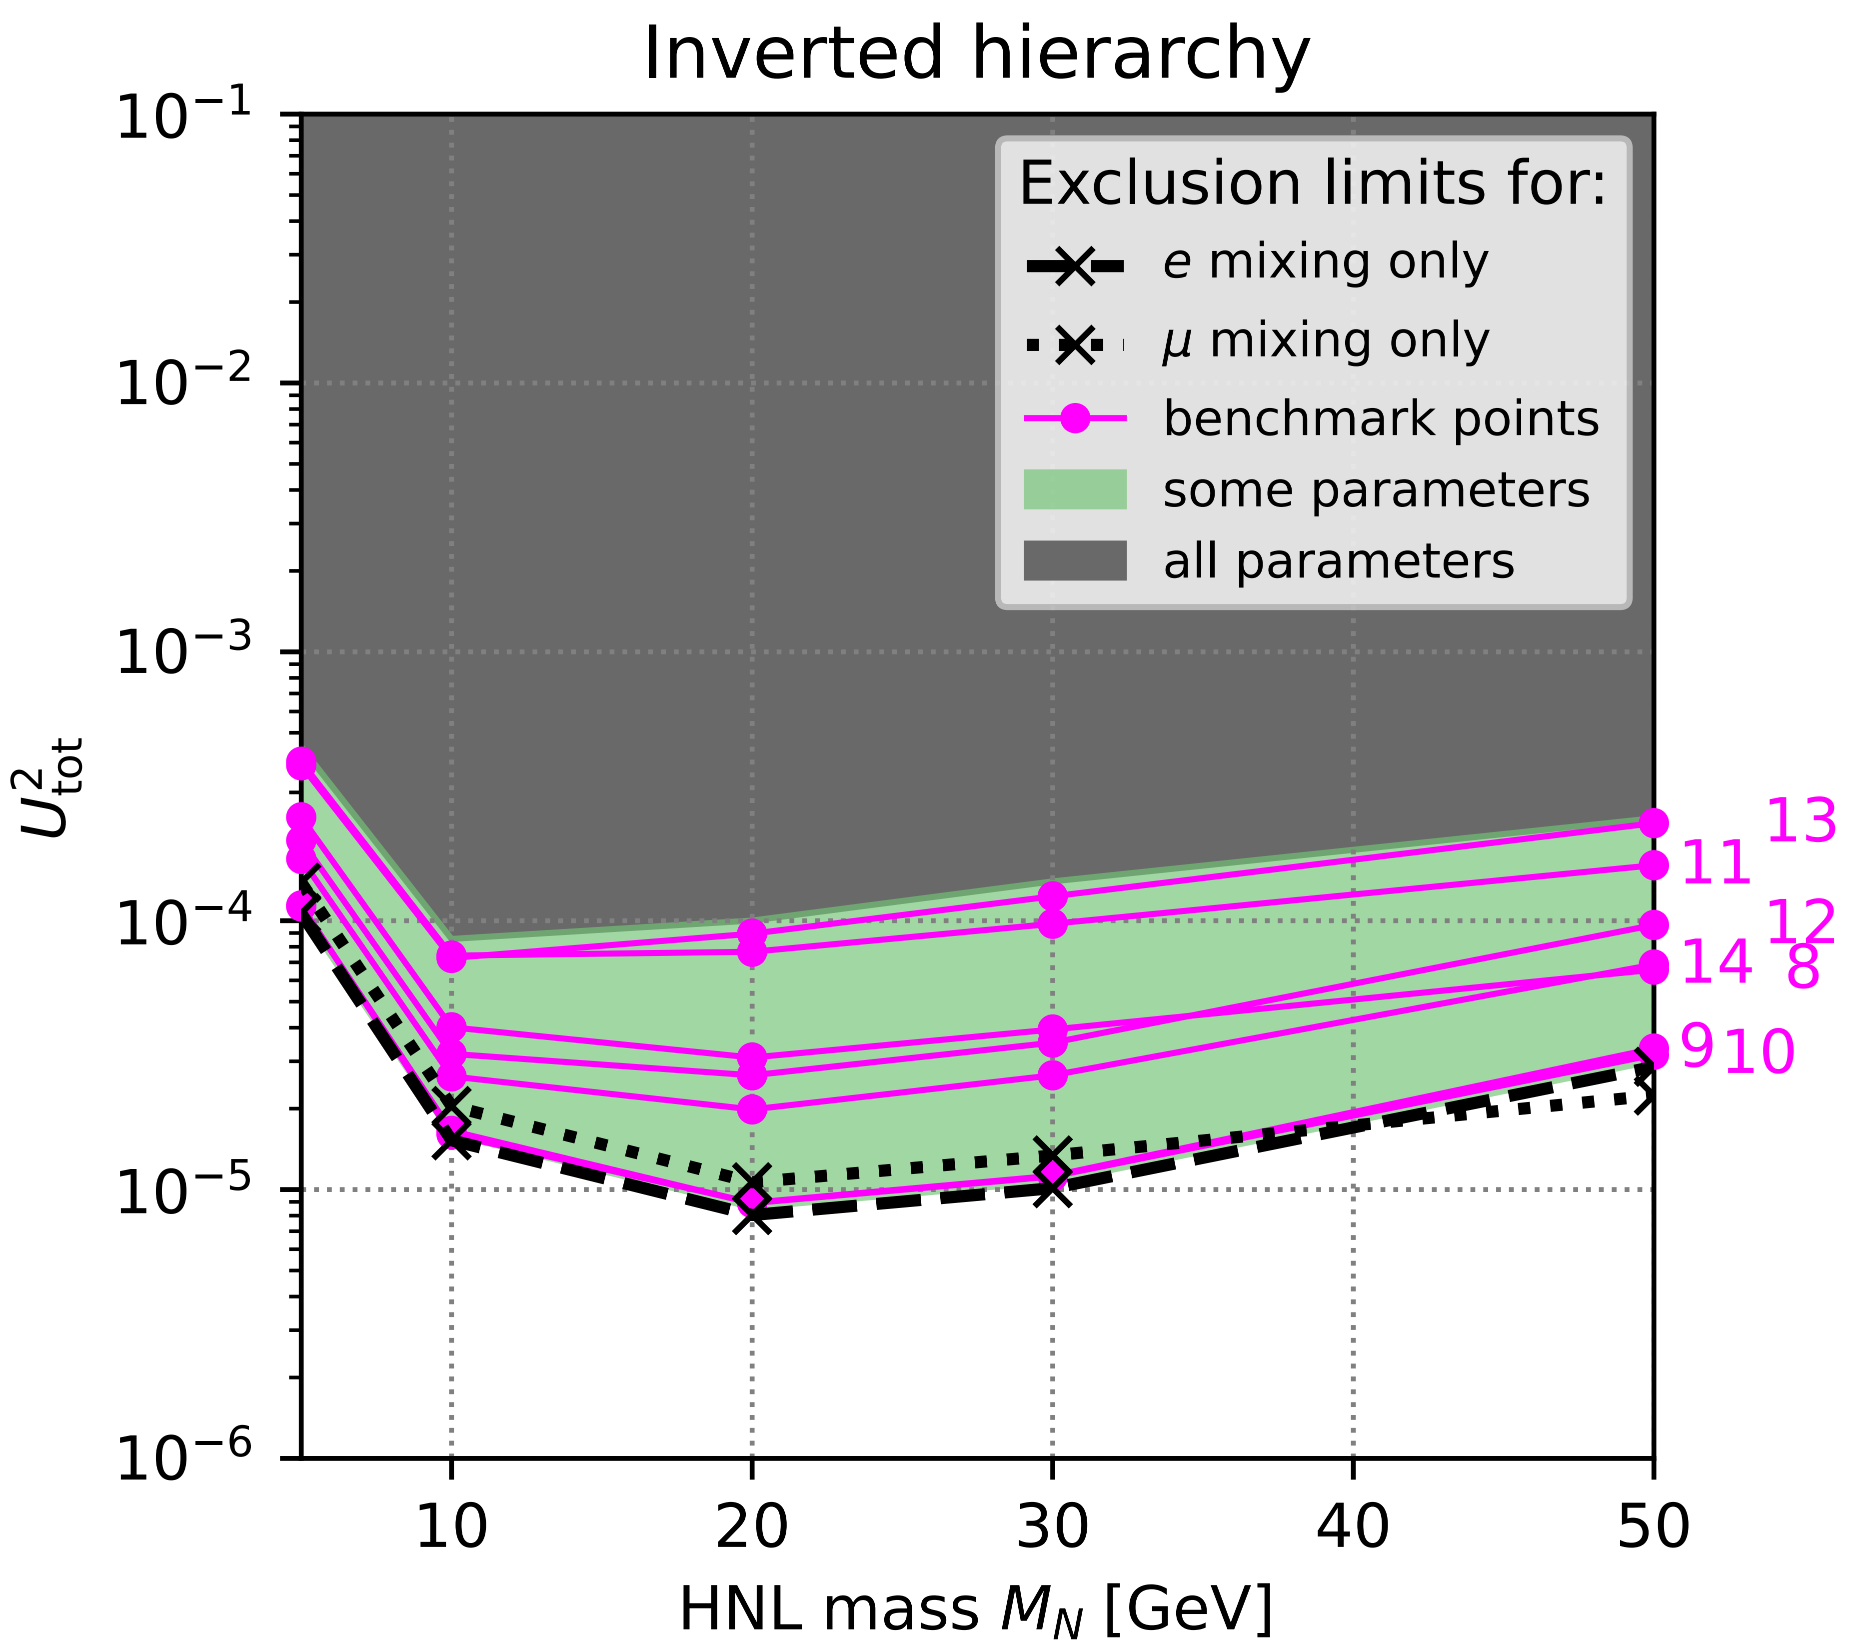

Majorana-like HNLs

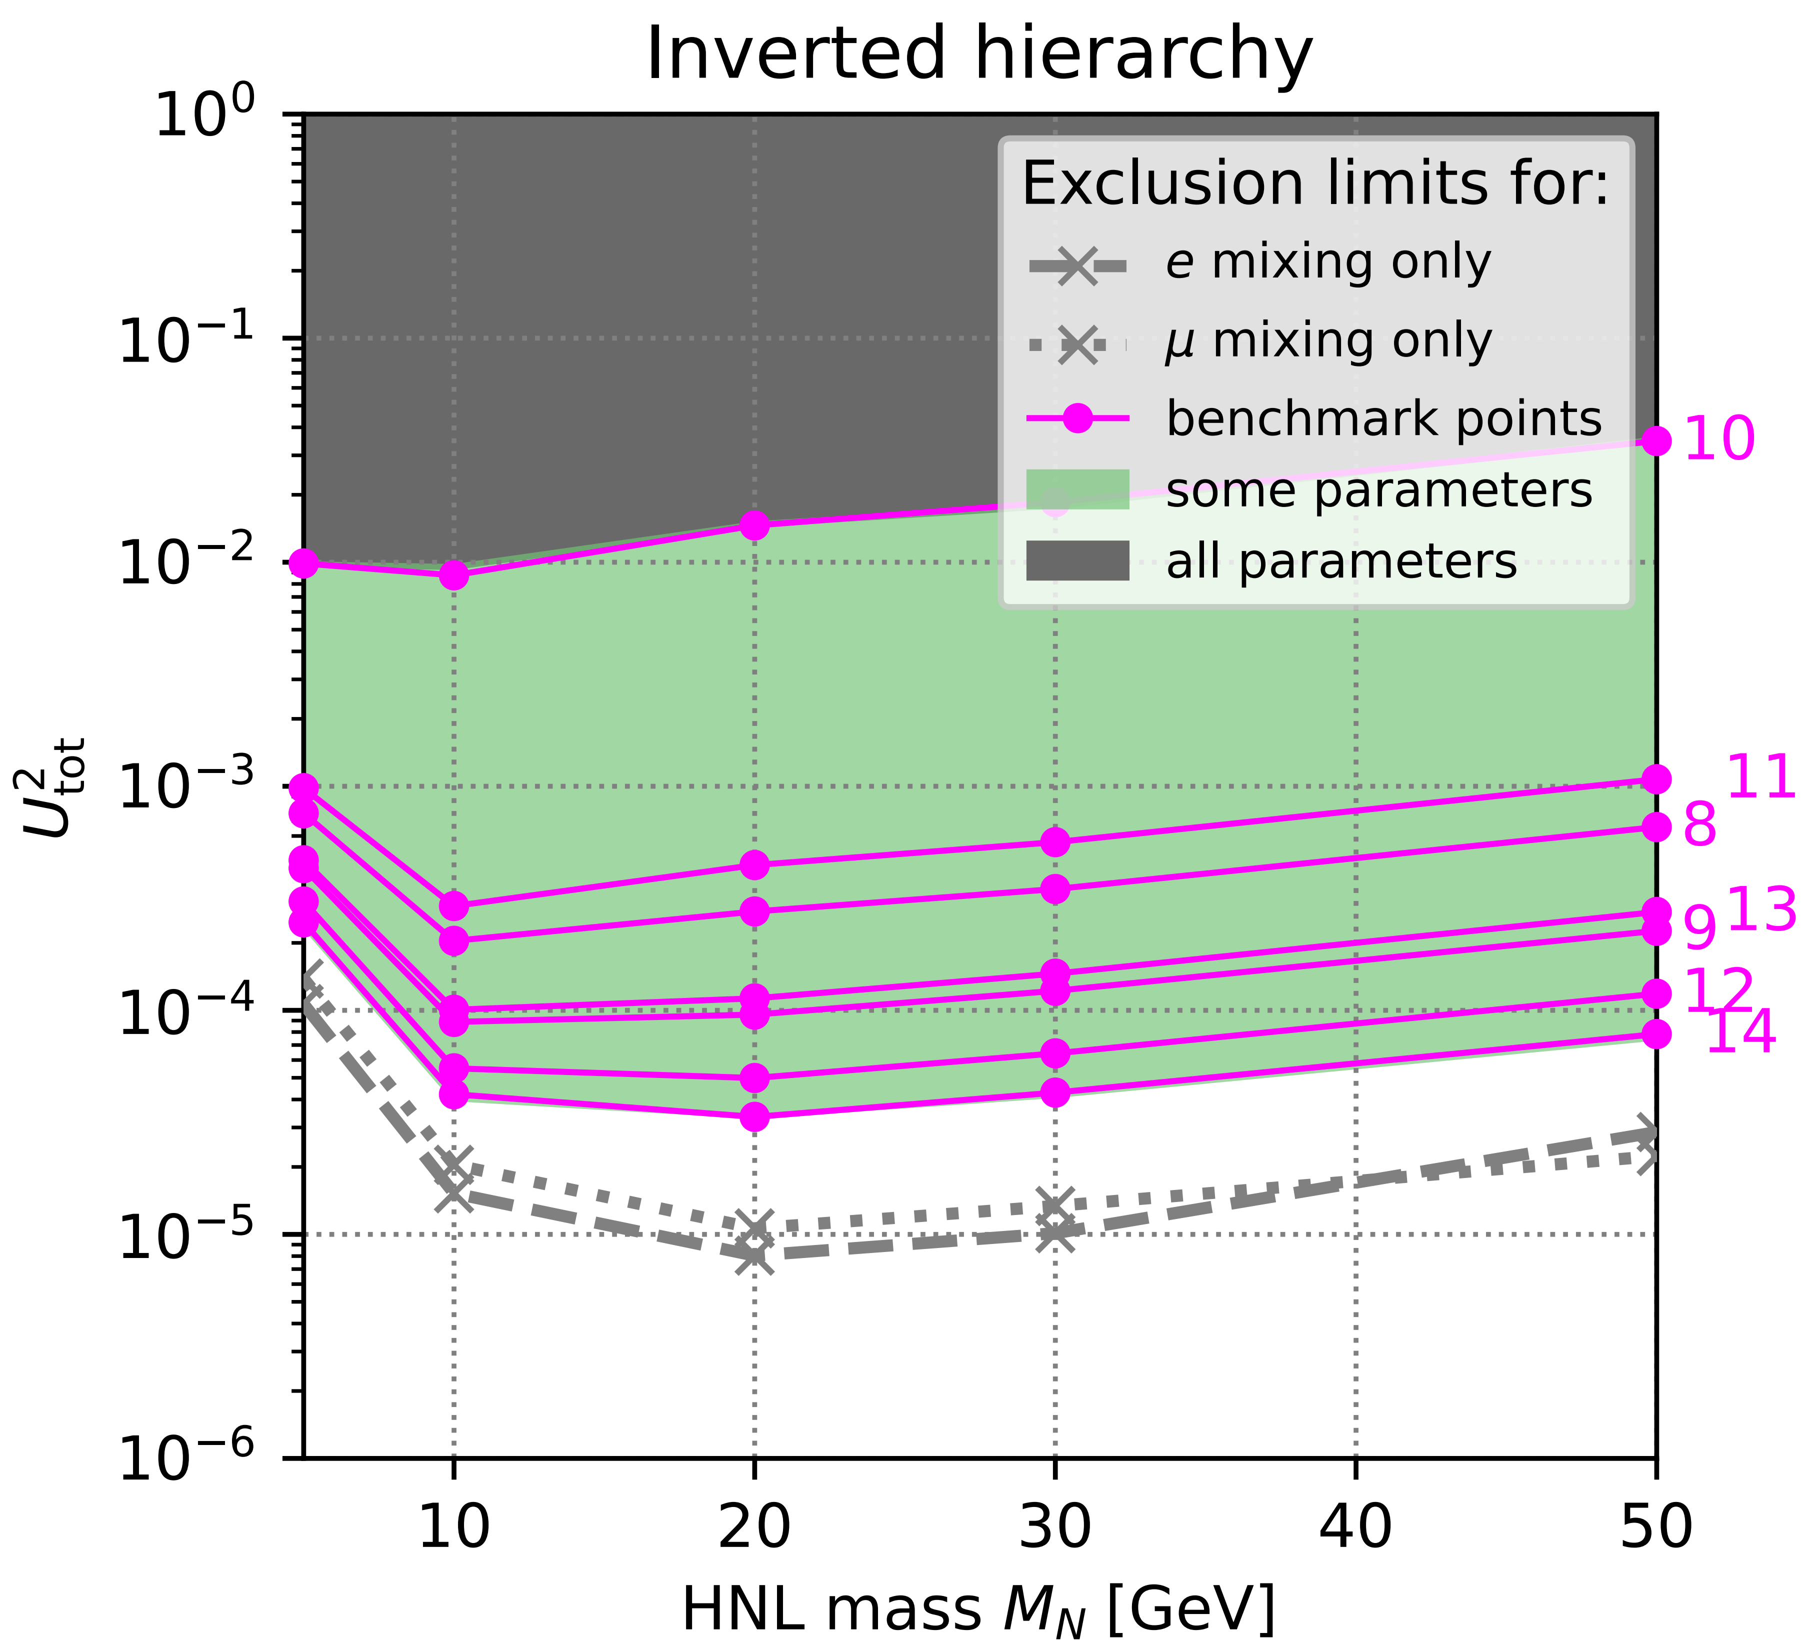

Dirac-like HNLs

"Anti-blind spot" #1

Constraints from flavour mixing pattern

Indirect limit on tau mixing!

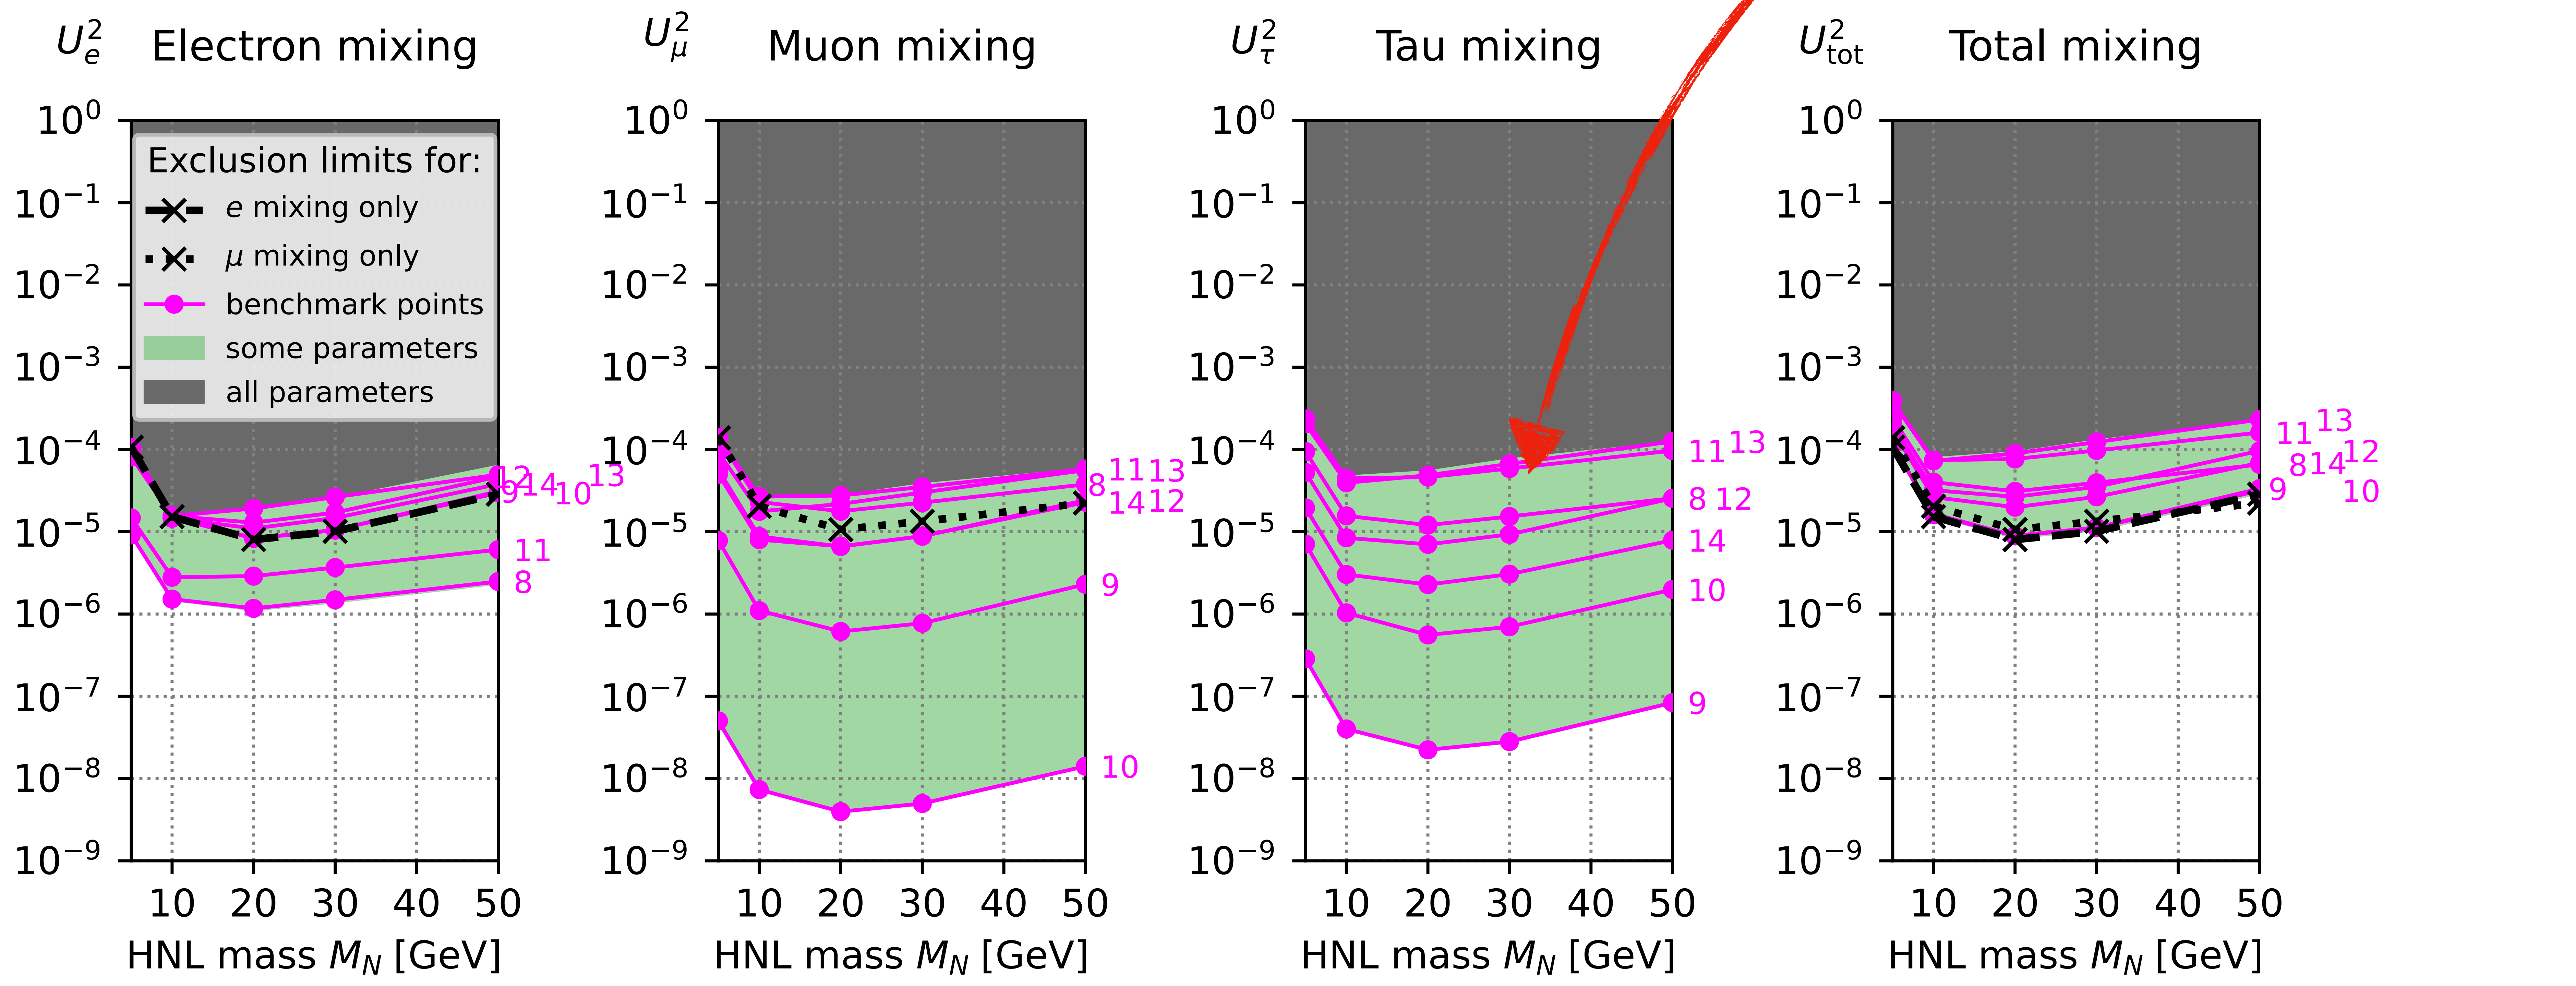

Fixing the ratio $|\Theta_e|^2 : |\Theta_\mu|^2 : |\Theta_\tau|^2$ leads to stronger constraints on the individual mixing angles.

"Anti-blind spot" #1

Constraints from flavour mixing pattern

Indirect limit on tau mixing!

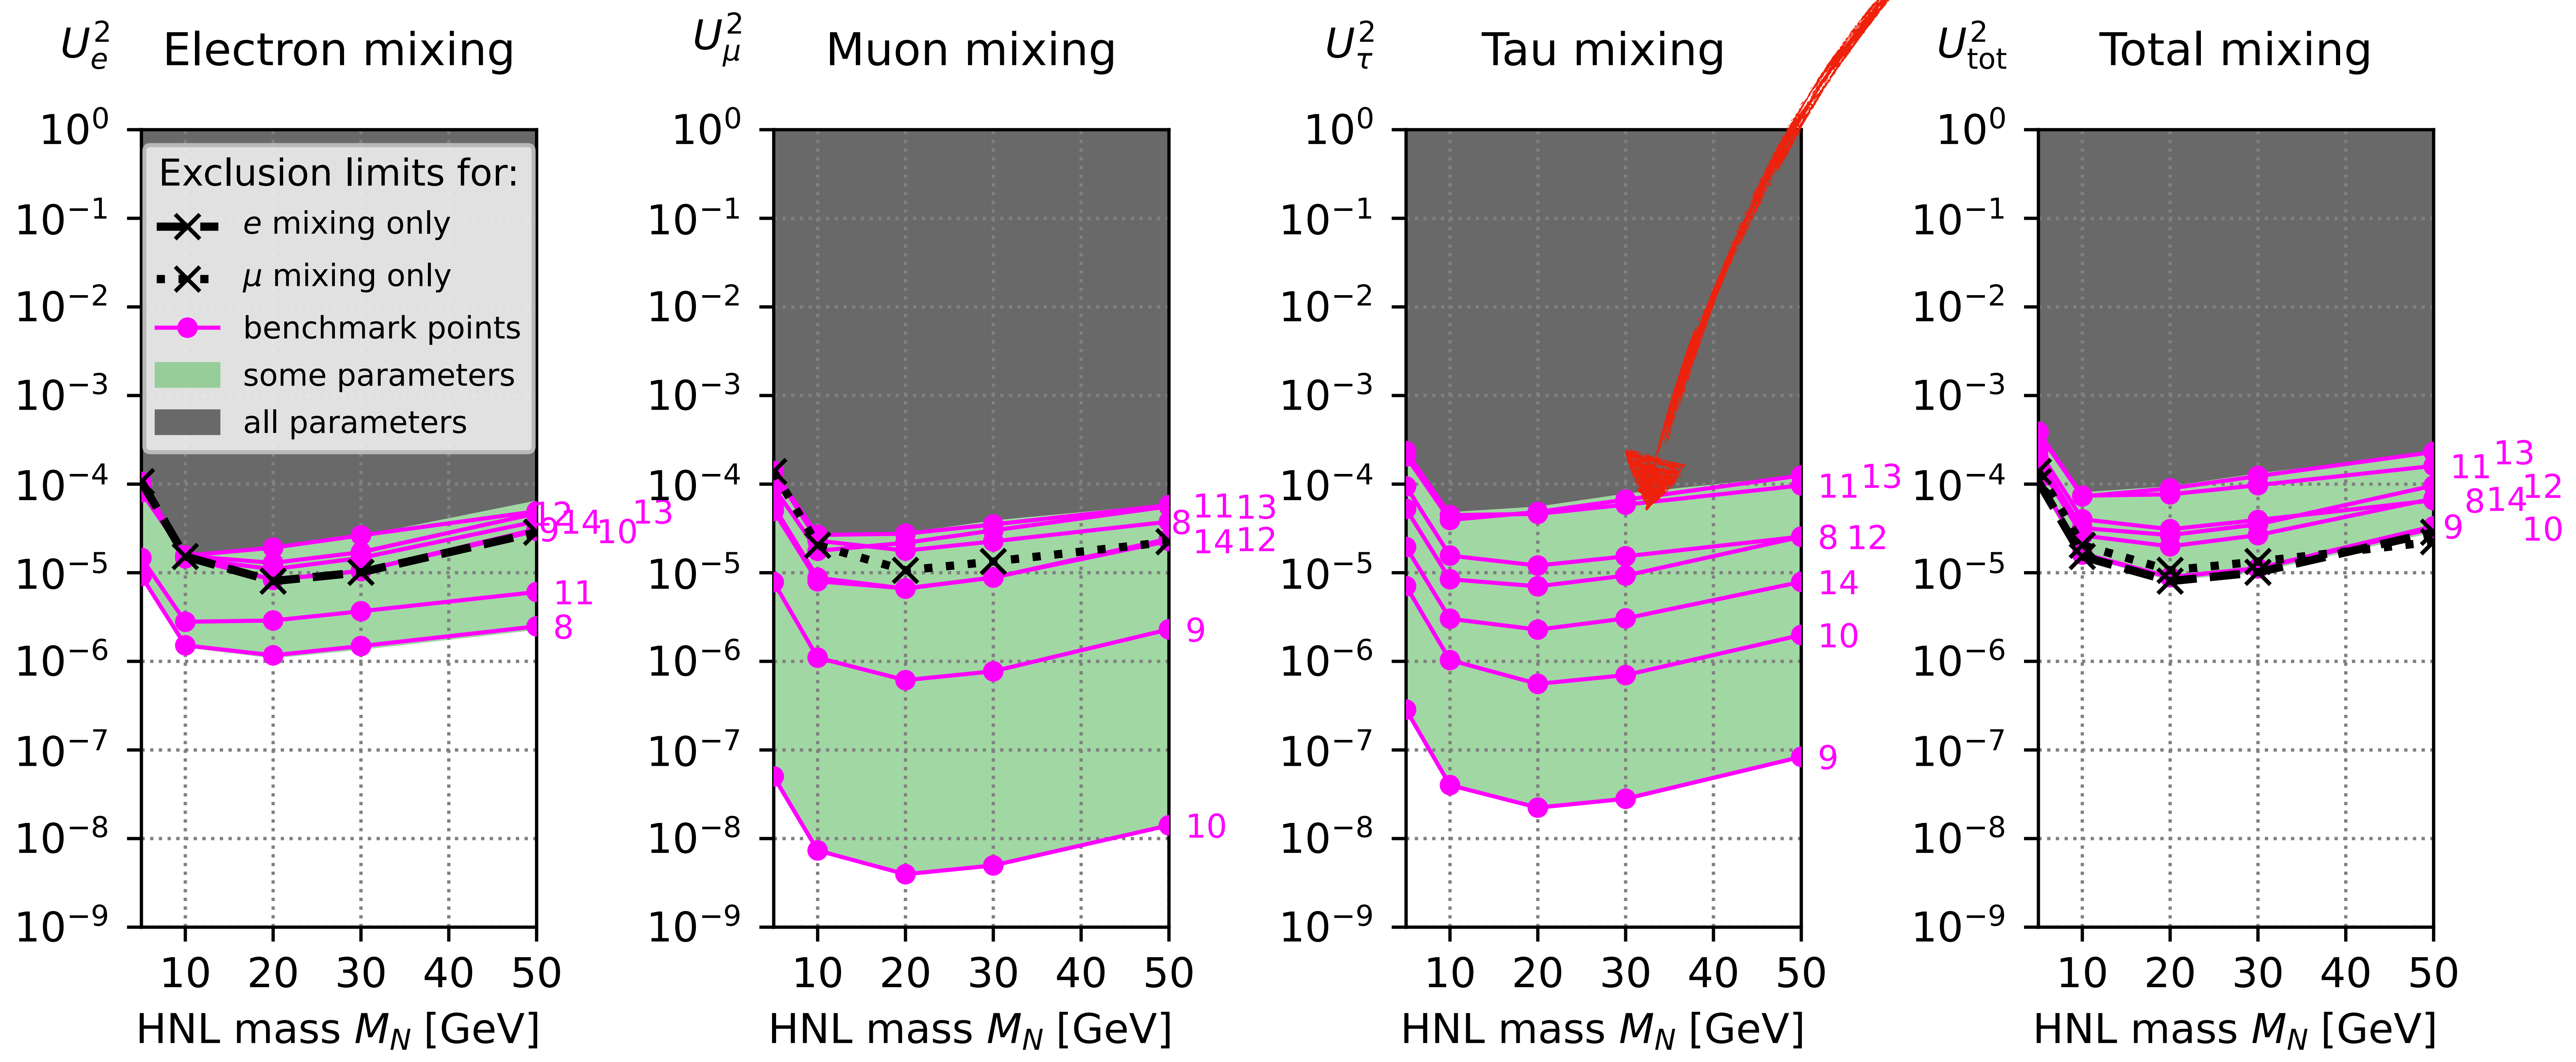

Fixing the ratio $|\Theta_e|^2 : |\Theta_\mu|^2 : |\Theta_\tau|^2$ leads to stronger constraints on the individual mixing angles.

Good for global scans & Bayesian analyses!

II. Displaced trilepton search

[ATLAS: 2204.11988]

$139\,\mathrm{fb}^{-1}$ at 13 TeV

$M_N \in [2.5, 15]\,\mathrm{GeV}$

EUROPEAN ORGANISATION FOR NUCLEAR RESEARCH (CERN)

ATLAS EXPERIMENT

Submitted to: PRL

CERN-EP-2022-017
April 27, 2022

Search for heavy neutral leptons in decays of W bosons using a dilepton displaced vertex in $\sqrt{s}$ = 13 TeV pp collisions with the ATLAS detector

The ATLAS Collaboration

A search for a long-lived, heavy neutral lepton ($\mathcal{N}$) in 139 fb^{-1} of $\sqrt{s}$ = 13 TeV pp collision data collected by the ATLAS detector at the Large Hadron Collider is reported. The $\mathcal{N}$ is produced via $W \to \mathcal{N}\mu$ or $W \to \mathcal{N}e$ and decays into two charged leptons and a neutrino, forming a displaced vertex. The $\mathcal{N}$ mass is used to discriminate between signal and background. No signal is observed, and limits are set on the squared mixing parameters of the $\mathcal{N}$ with the left-handed neutrino states for the $\mathcal{N}$ mass range 3 GeV $< m_{\mathcal{N}} <$ 15 GeV. For the first time, limits are given for both single-flavor and multiflavor mixing scenarios motivated by neutrino flavor oscillation results for both the normal and inverted neutrino-mass hierarchies.

Search signature

- HNLs produced in W boson decays.
- Travelling between 4 and 300 mm in the radial direction before decaying to *leptons*.
  Requires using large radius tracking (LRT)
  (up to 300 mm, in addition to standard tracking up to 10 mm)
- All combinations of e and μ flavours allowed.
  The displace vertex must be neutral.
- Neutral-current contribution for $N \to e^+e^-\nu$ and $N \to \mu^+\mu^-\nu$ decays.
- Sensitive to both LNC and LNV, but one neutrino *escapes* with the information on the final lepton number!

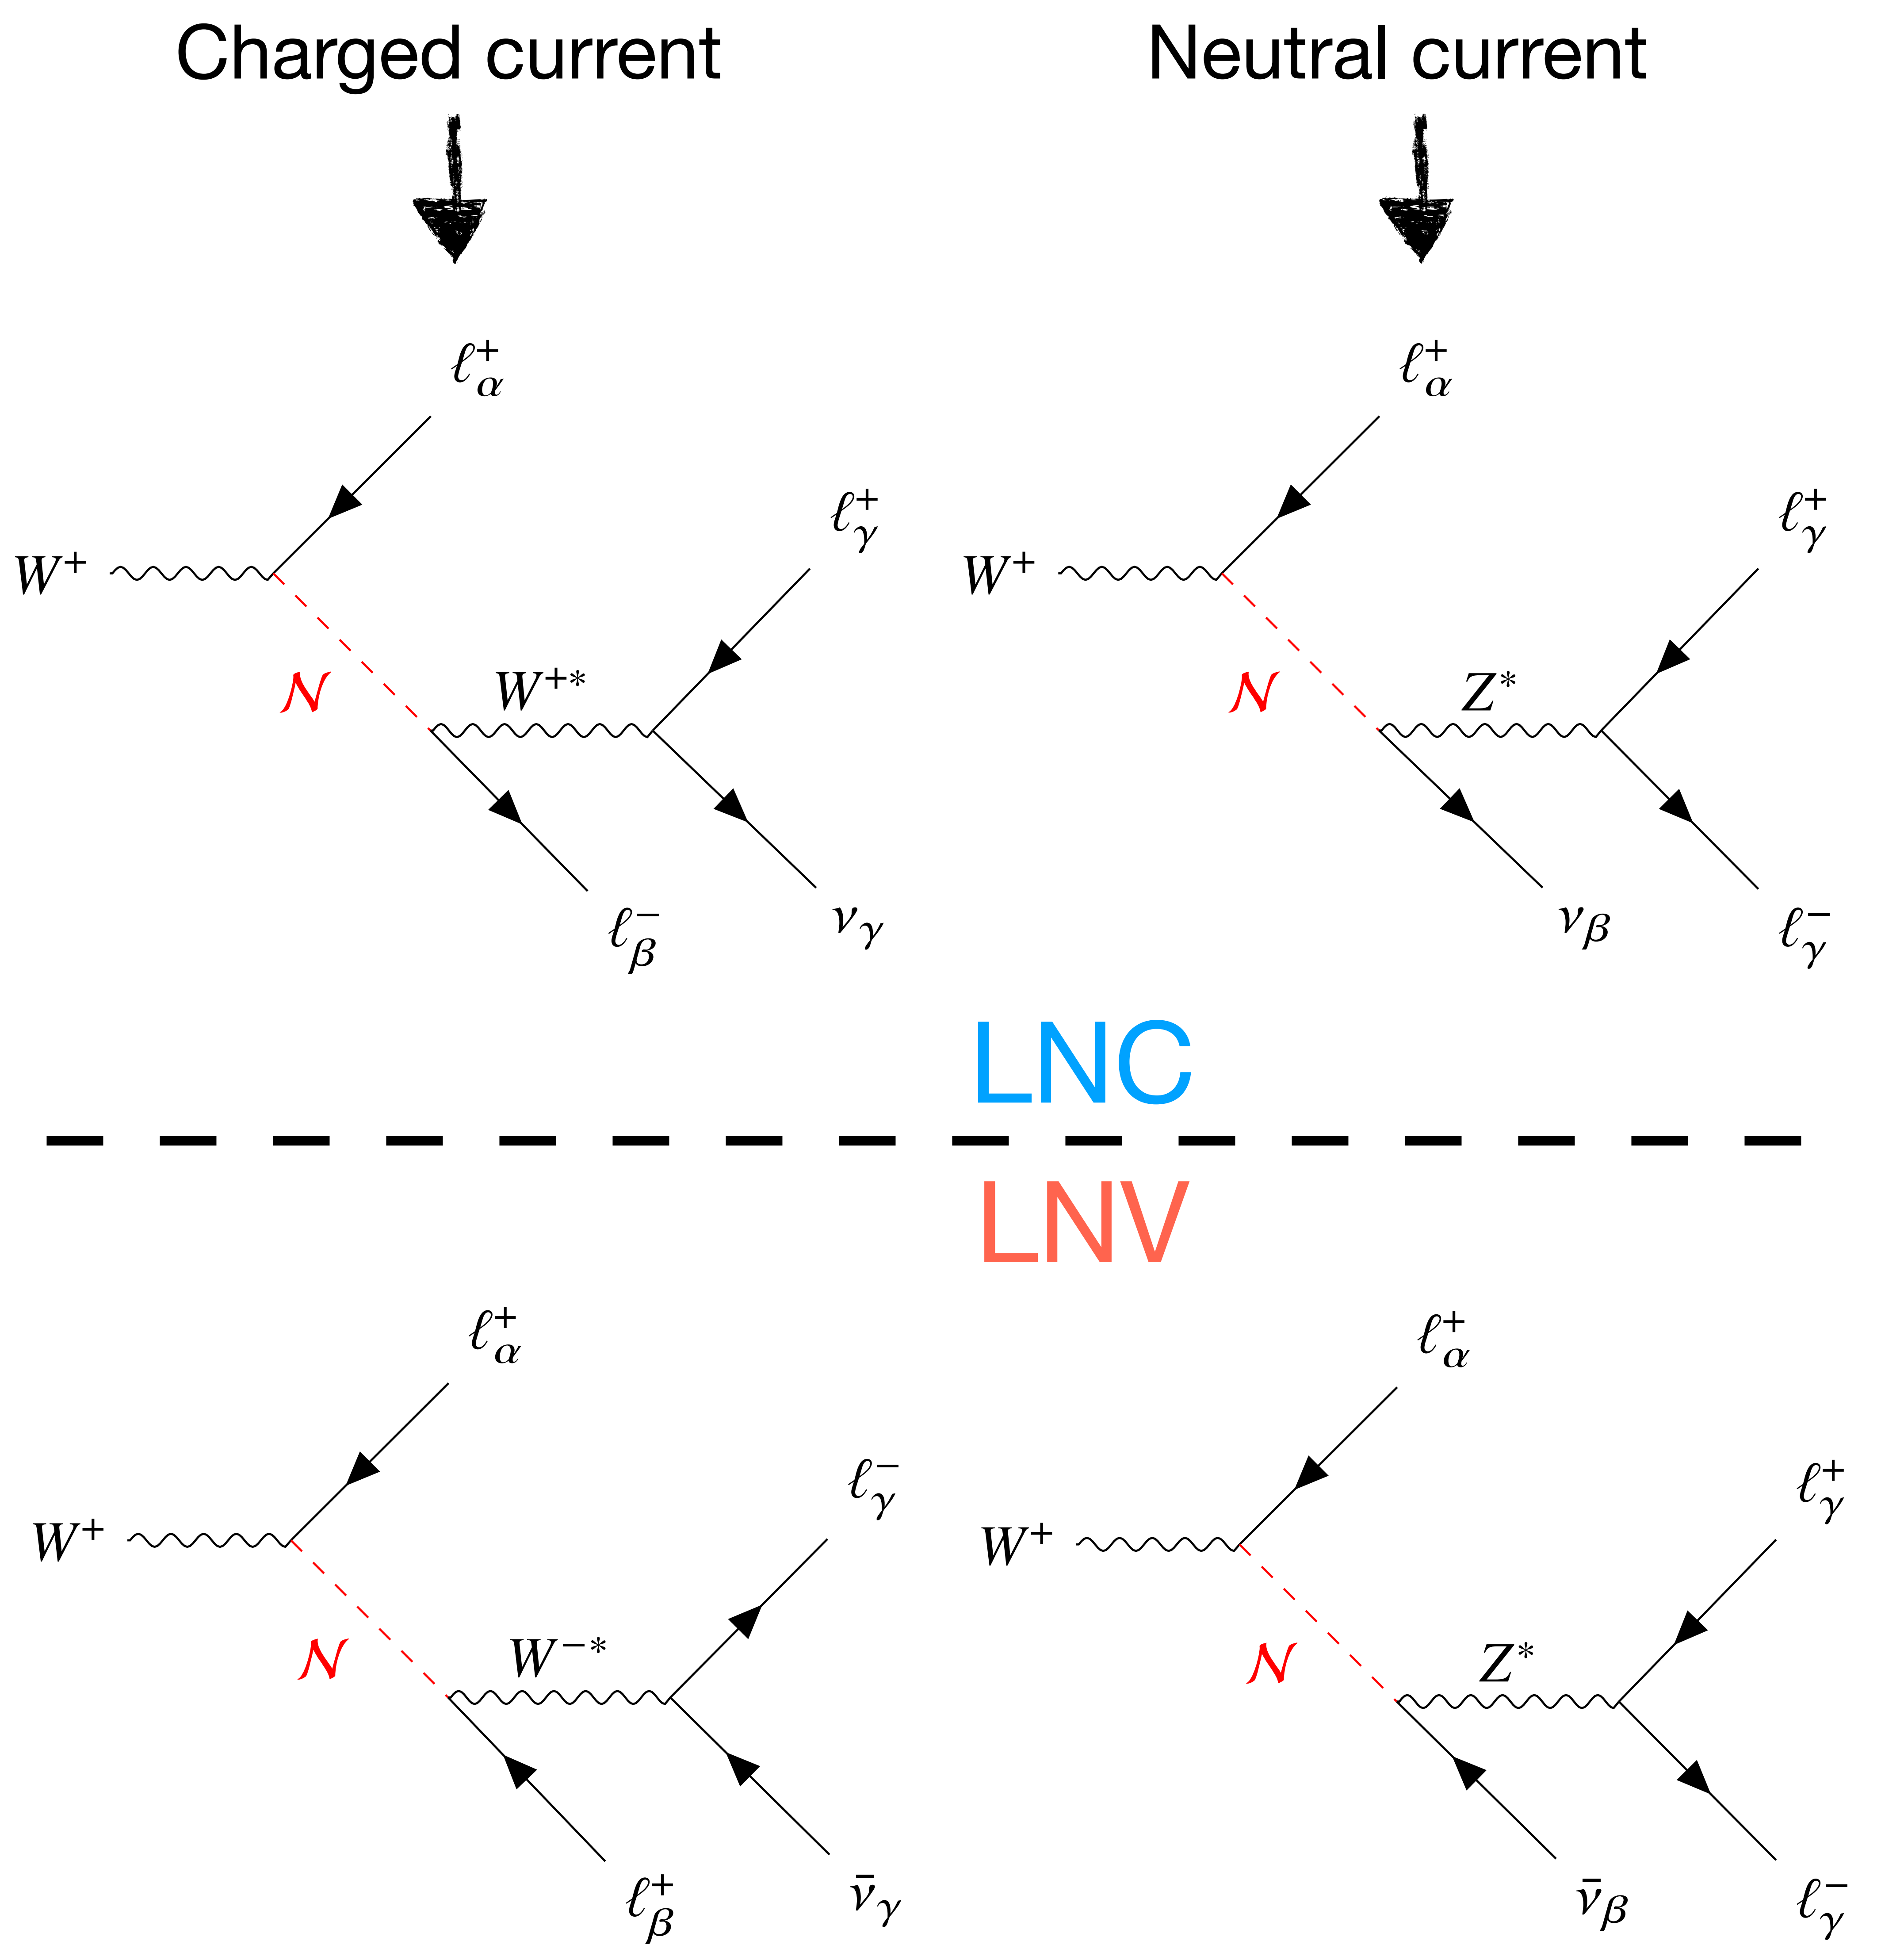

Background & cuts

- Require a prompt lepton and a displaced vertex (DV) formed by 2 charged leptons.
- Dominant background: random crossings of lepton tracks.
  Estimated by randomly shuffling lepton tracks from data.
- Other backgrounds: detector interactions, decays of metastable SM particles, $Z \to l^+l^-$ decays and cosmic muons. Can be **reduced** with simple cuts.
- Triggers & other cuts are rather **minimal**.
- The HNL mass can be reconstructed up to a twofold ambiguity, and is used to **bin** events (it will produce a peak over the background if HNLs are present).
- Control region: $m_{\mathrm{HNL}} \in [20,50]\,\mathrm{GeV}$ (too heavy to be long-lived).
- Validation region for data-driven background modelling: DV without prompt lepton.

Novelties of this search

- Improved modelling of spin correlations in the signal samples.
  $\Rightarrow$ Different efficiencies for LNC / LNV.

- Interpretation in terms of two quasi-Dirac HNLs, both in the Majorana ($\delta M \tau \gg 2\pi$) and Dirac ($\delta M \tau \ll 2\pi$) limit.
  The two benchmarks from the FIPs 2020 report are used. [FIPs 2020: 2102.12143]

- This is **in addition** to the interpretation in terms of the usual 1 HNL mixing with one flavour.

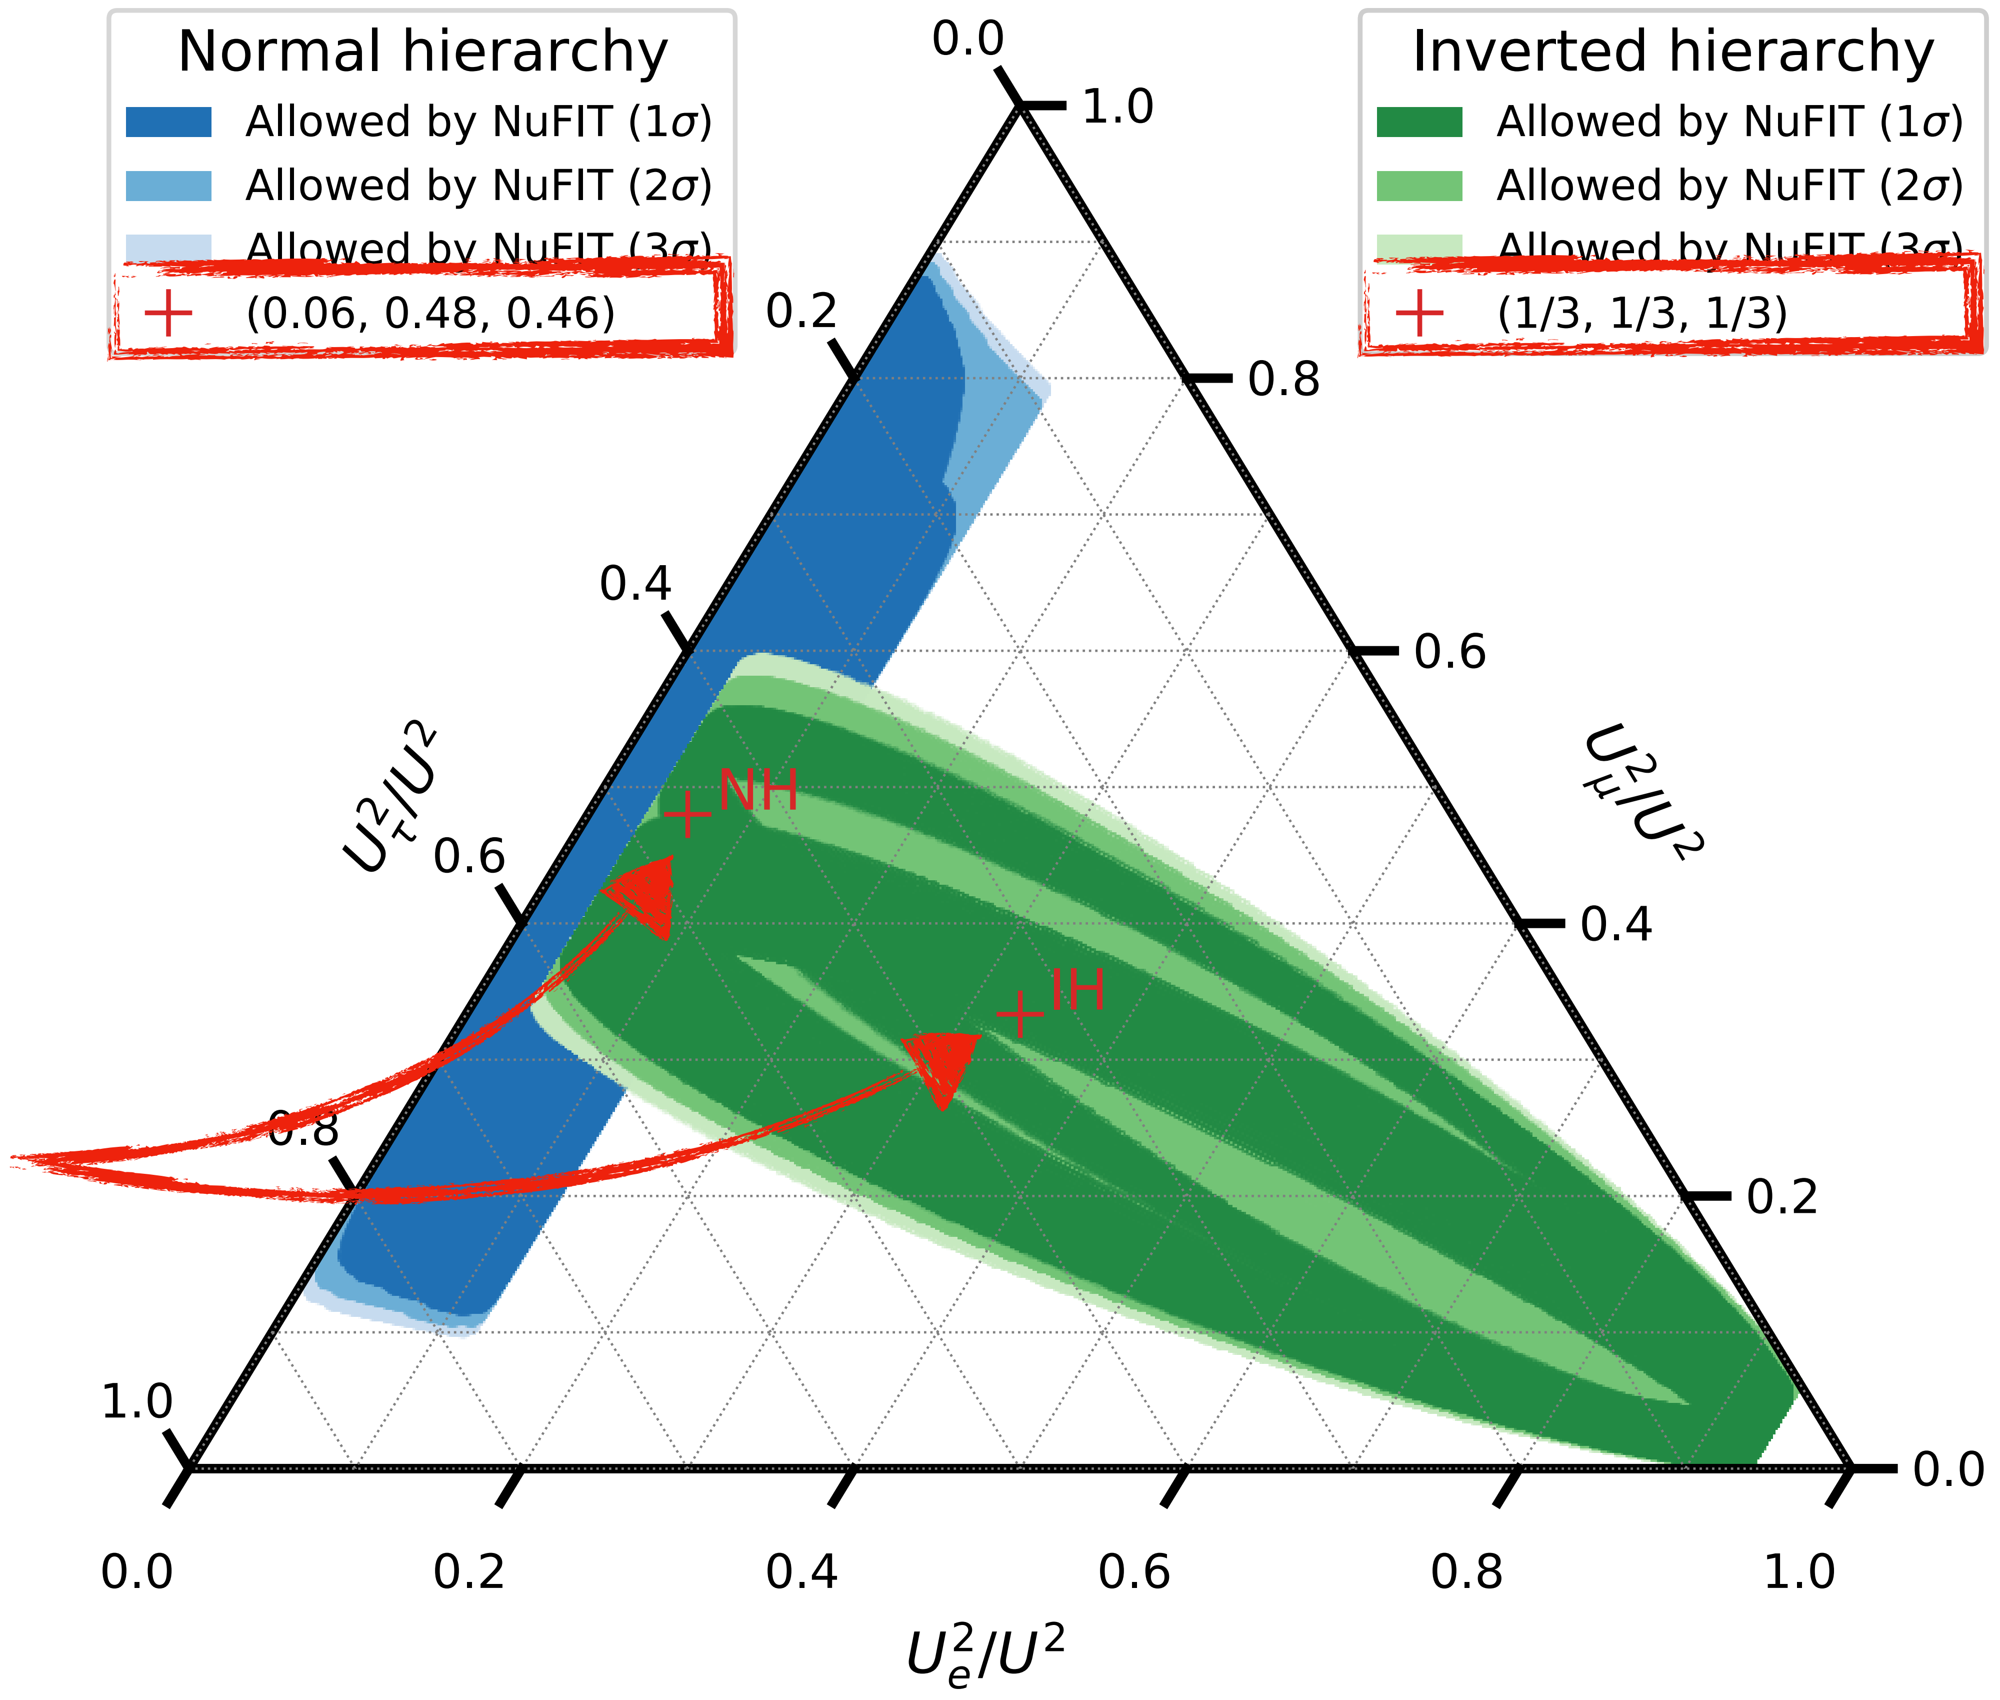

Results

No excess observed

New benchmarks!

"1 Majorana HNL mixing with a single flavour" is still there

Good parameter space coverage thanks to the inclusion of all flavour combinations
⟹ only small differences between benchmarks

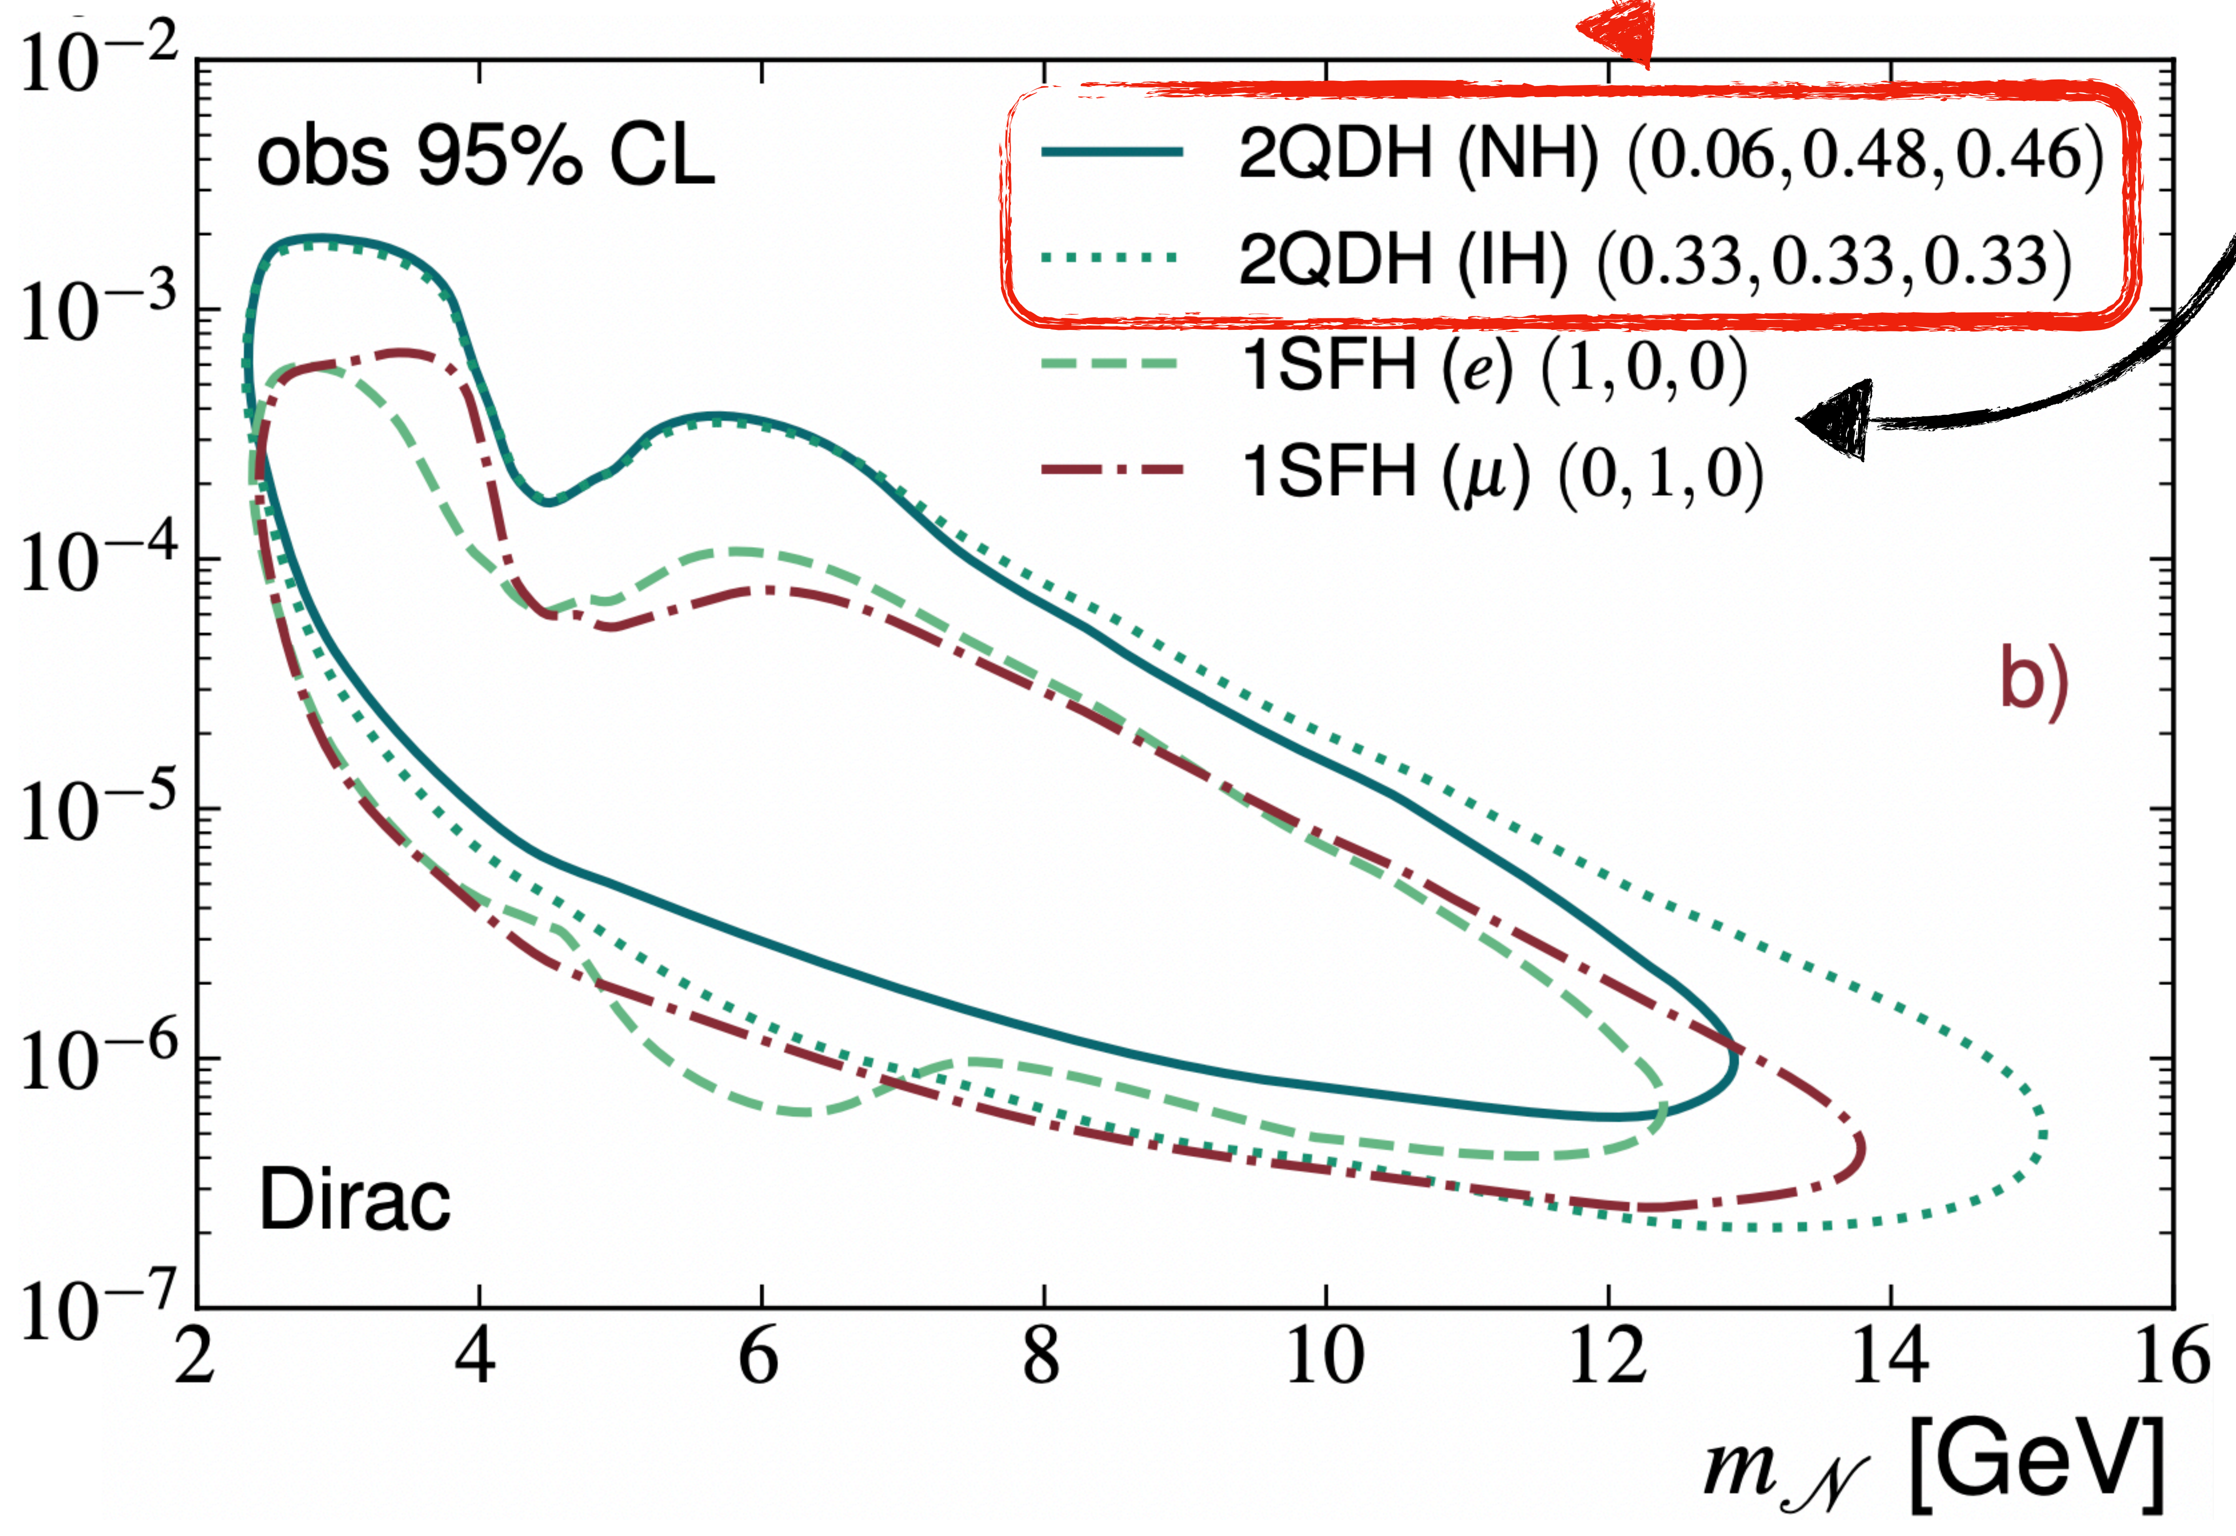

Majorana-like HNLs

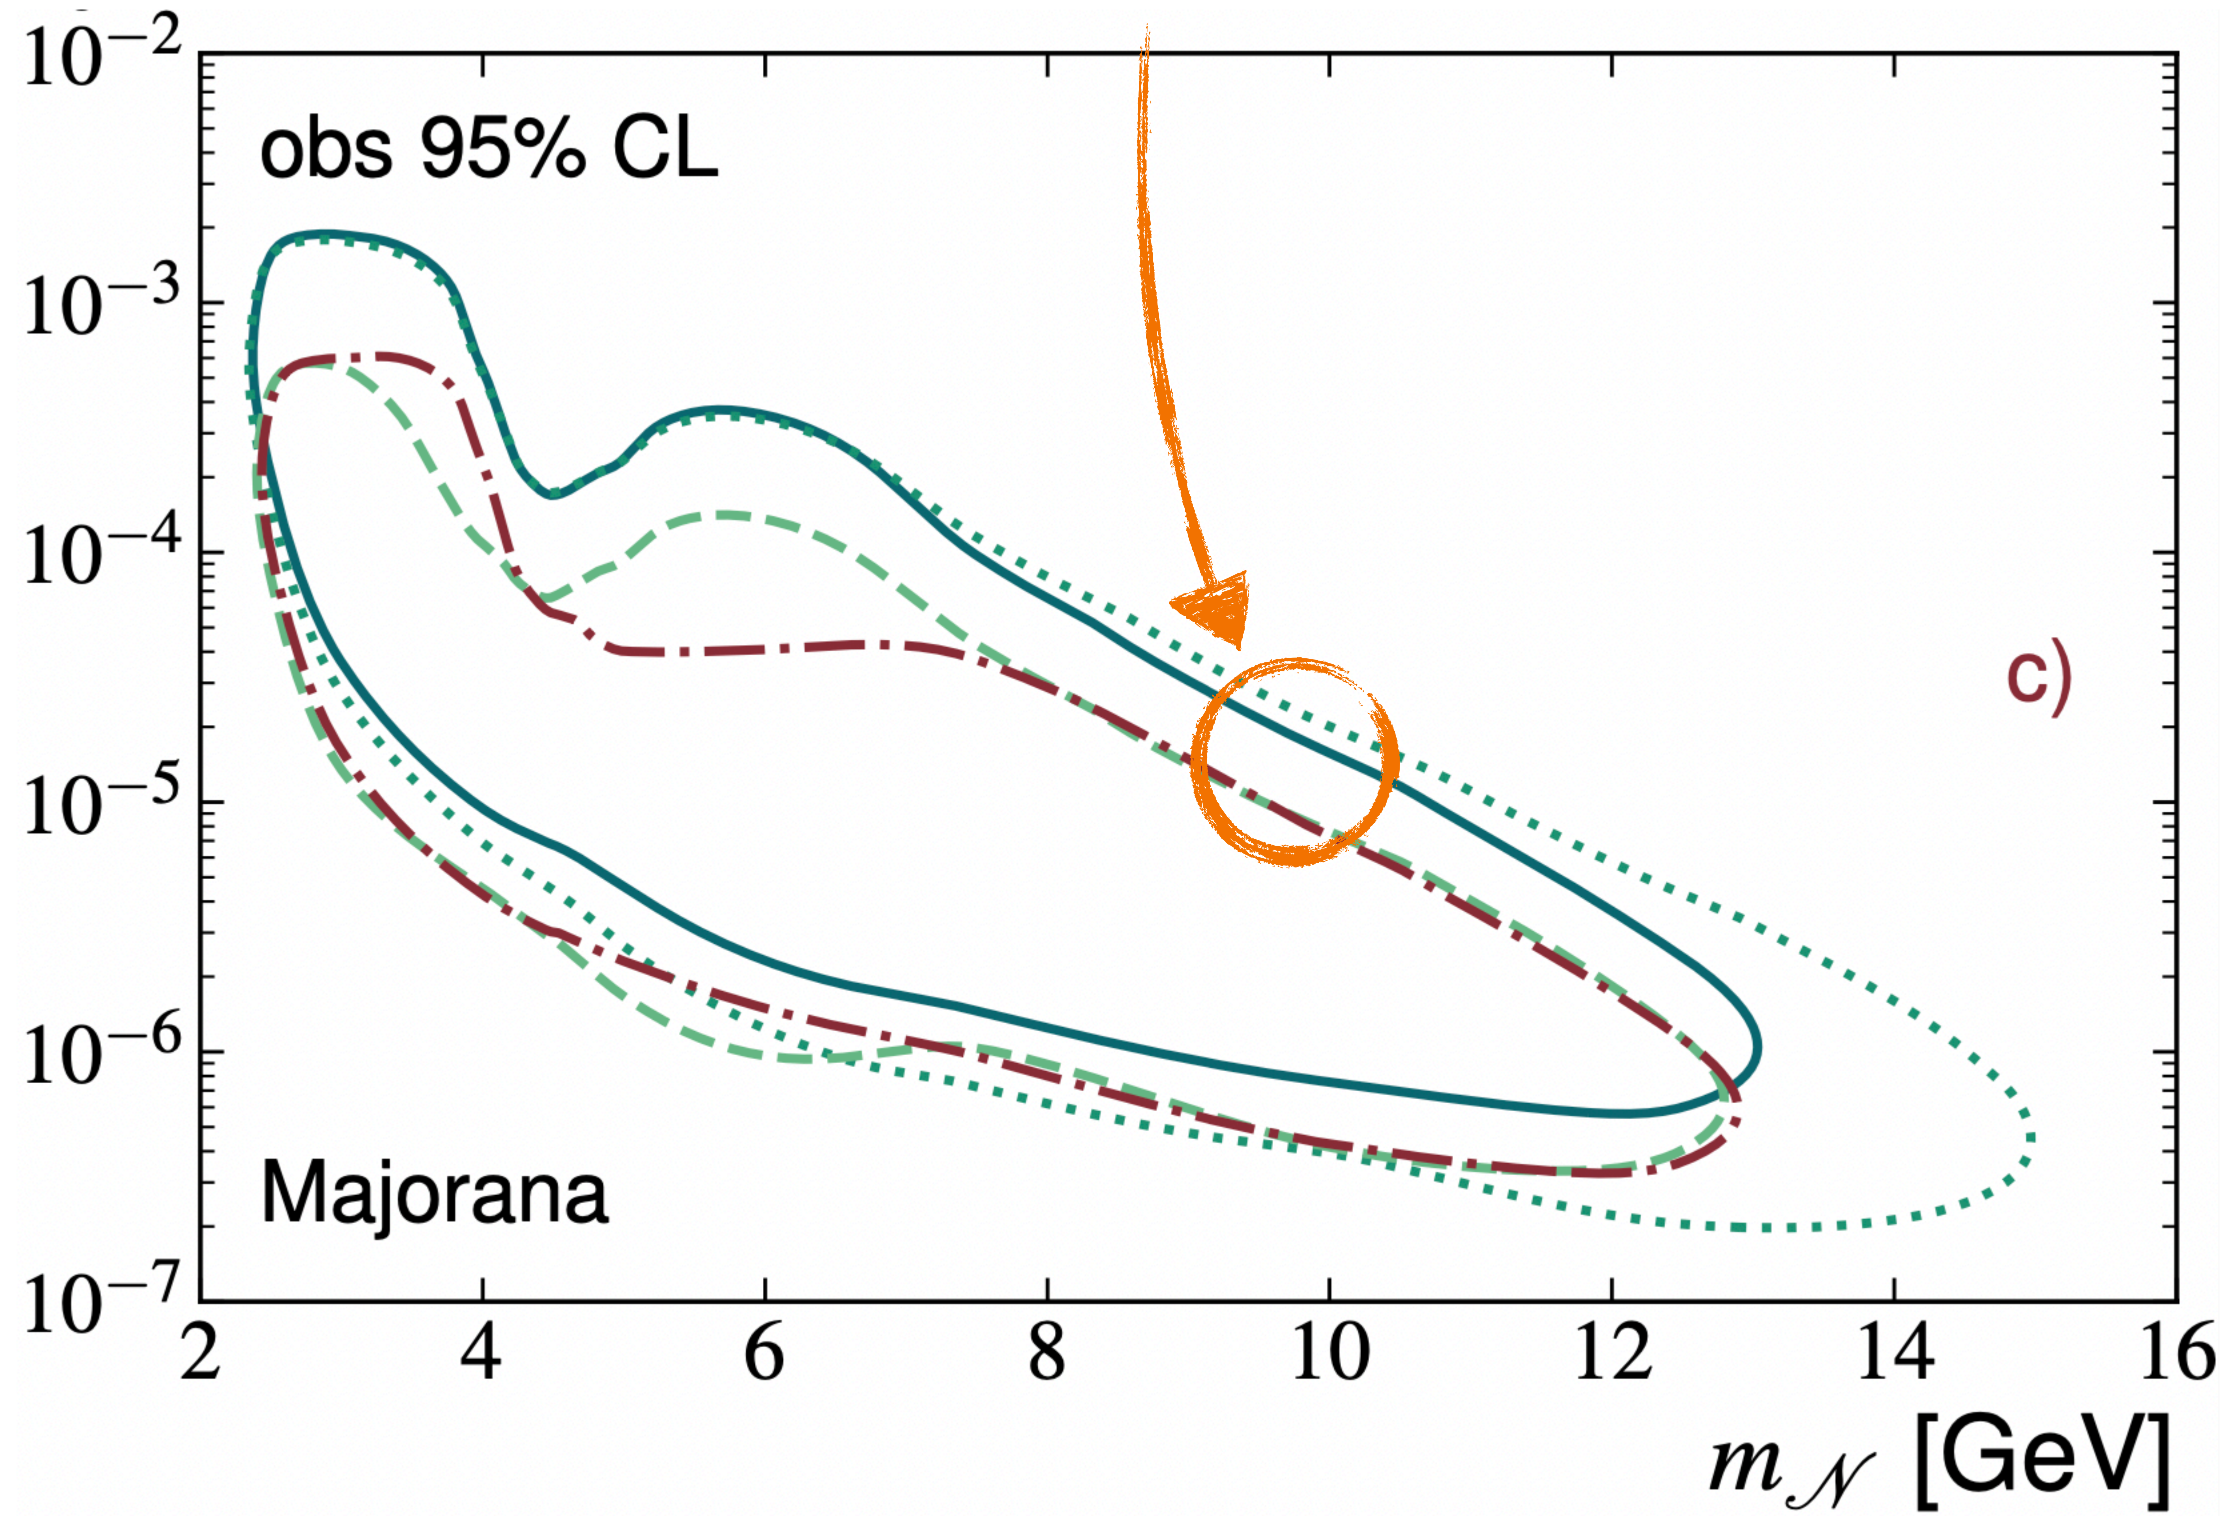

Dirac-like HNLs

Partial conclusion

- The new benchmarks address our first blind spot.
- We can be confident that the search has sensitivity to both "Majorana-like" and "Dirac-like" HNLs for reasonable choices of mixing angles.
- But what about **global parameter scans**? Or Bayesian analyses?
- Can we construct some **sensitivity matrix** like at SHiP?

Generic limits

General considerations

- LHC searches are **not** background-free.
- They typically have significant **systematic uncertainties**.
- Limits are computed using some complex statistical procedure (e.g. CL_s).
- Any accurate reinterpretation will require the collaboration to publish their background model, ideally in a machine-readable format (e.g. pyhf): Full or simplified likelihood, background correlation matrix.
  [Cf. LHC Reinterpretation Forum: 2003.07868]
- However, we can generalise the sensitivity matrix approach in order to **exactly** extrapolate the signal to an arbitrary choice of mixing angles.

Scaling properties of the signal

- HNL always nearly on-shell due to its small width → narrow-width approx.
- Cross-section for a given process:

$$\sigma_{\text{process}} = \sigma_{\text{prod}} \times \text{Br}_{\text{decay}}$$
$$\propto |\Theta_{\text{prod}}|^2 |\Theta_{\text{decay}}|^2 / \Gamma_{\text{total}}$$

- The total width Γ_{total} still depends on all the mixing angles.

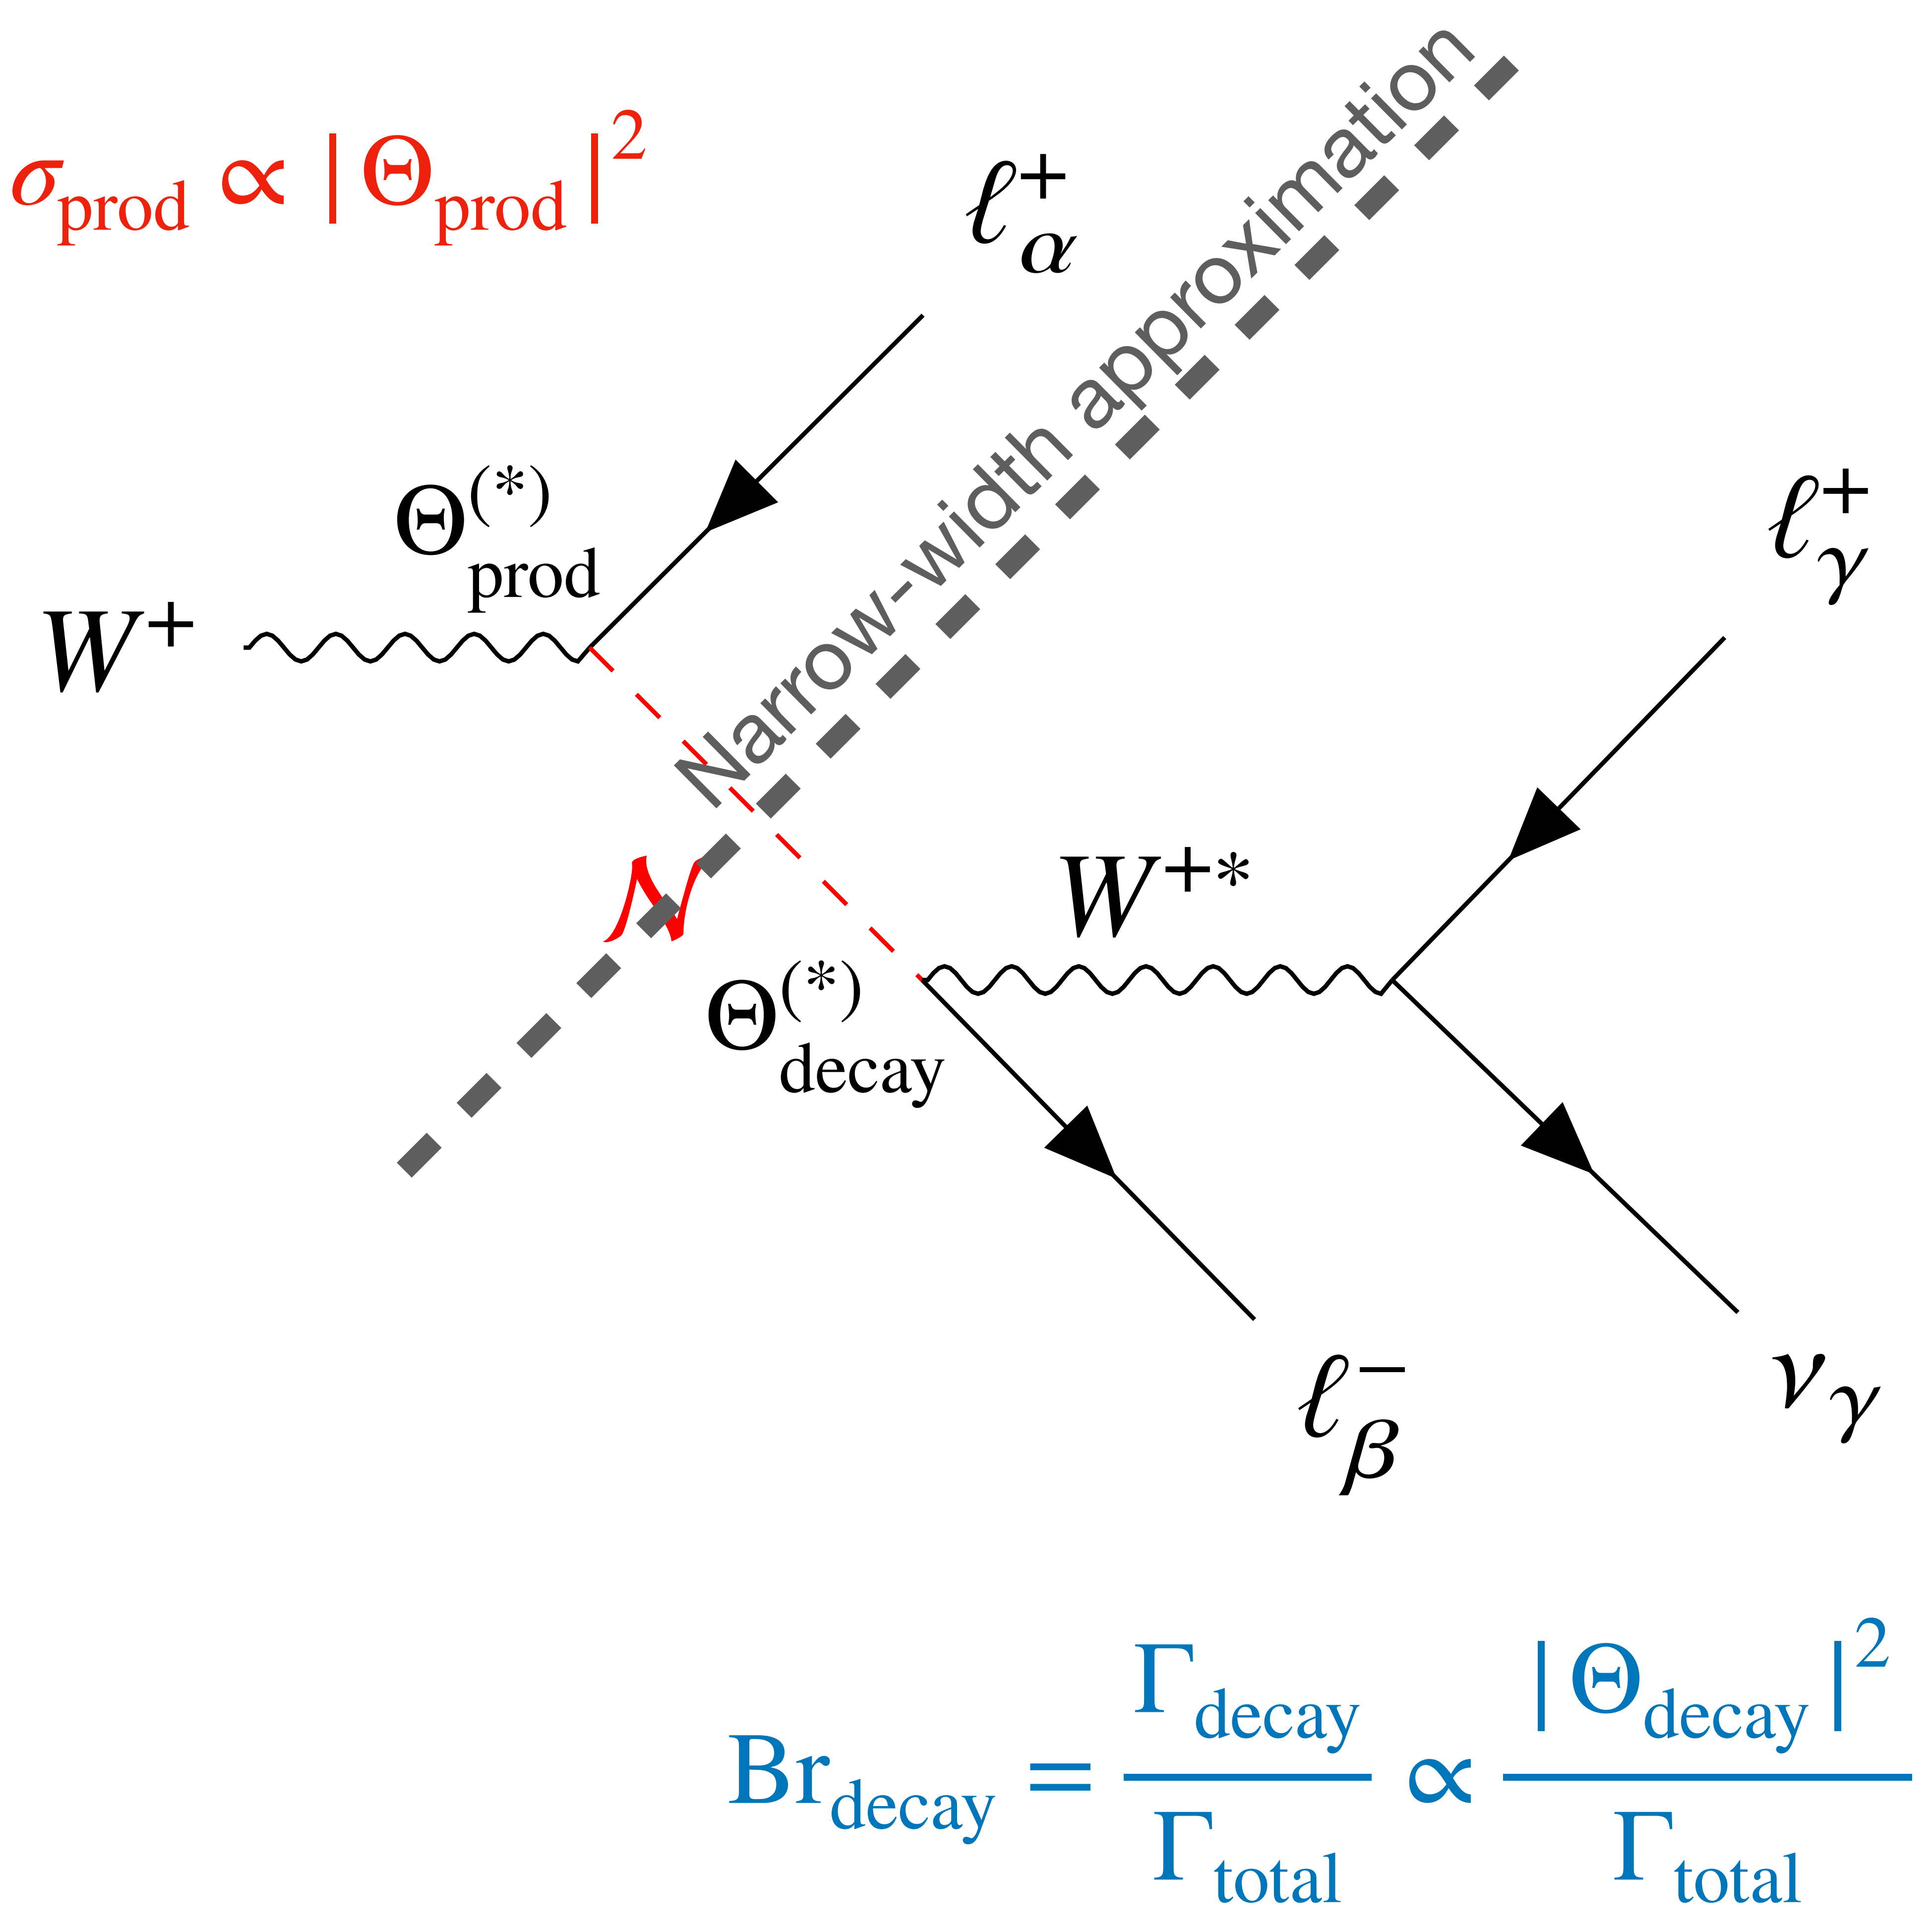

Summing over channels

- For HNLs, all diagrams contributing to the same final state involve the same mixing angles → no need to worry about interference!
- The total number of events (before considering efficiencies) is then:

$$N_{\text{total}}(M_N, \Theta_e, \Theta_\mu, \Theta_\tau) = L_{\text{int}} \times \frac{|\Theta_\alpha|^2 \Sigma_{\alpha\beta}(M_N) |\Theta_\beta|^2}{\Gamma_{\text{total}}(M_N, \Theta_e, \Theta_\mu, \Theta_\tau)/\Gamma_{\text{ref}}}$$

with the cross-section matrix $\Sigma_{\alpha\beta}$ the sum of the cross-sections of all processes mediated by flavour α at the HNL production vertex and β at its decay vertex, computed for **unit mixing angles** and assuming a small **reference width** Γ_{ref}.

Total HNL width

- The total width is the sum of partial widths, each mediated by one and only one mixing angle Θ_α.

  Therefore $\Gamma_{\text{total}}(M_N, \Theta_e, \Theta_\mu, \Theta_\tau) = |\Theta_\alpha|^2 \hat{\Gamma}_\alpha(M_N)$

  with $\hat{\Gamma}_\alpha$ the sum of the partial widths mediated by flavour α, computed for a **unit mixing angle**.

- Putting everything together:

$$N_{\text{total}}(M_N, \Theta_e, \Theta_\mu, \Theta_\tau) = L_{\text{int}} \times \frac{|\Theta_\alpha|^2 \Sigma_{\alpha\beta}(M_N) |\Theta_\beta|^2}{|\Theta_\gamma|^2 \hat{\Gamma}_\gamma(M_N)/\Gamma_{\text{ref}}}$$

Efficiencies & binning

Some further complications

- For displaced HNLs the signal will depend on the HNL lifetime τ_N through the experimental efficiencies.
- Let's temporarily treat the HNL lifetime τ_N as an independent parameter $\neq \Gamma_{\text{total}}^{-1}$.
- Let $\varepsilon_{P,b}(M_N, \tau_N)$ be the signal efficiency in bin *b* for process *P*, for an HNL with mass M_N and lifetime τ_N.
- The total number of events in bin *b* is then:

$$N_b = L_{\text{int}} \times \sum_P \varepsilon_{P,b}(M_N, \tau_N) \times \sigma_P(M_N, \Theta_e, \Theta_\mu, \Theta_\tau)$$

Quasi-Dirac HNLs

(Note that "2 Dirac-like HNLs" = "1 Dirac HNL" up to a rescaling of Θ by $\sqrt{2}$)

Nature	$c_P,\ P \in \text{LNC}$	$c_P,\ P \in \text{LNV}$	$c_\Gamma = \Gamma_N/\Gamma_{\text{Maj.}}$
One Majorana HNL (reference)	**1**	**1**	**1**
One Dirac HNL	1	0	1/2
Quasi-Dirac pair: **Majorana-like**	2	2	1
Quasi-Dirac pair: **Dirac-like**	4	0	1

- If HNLs are quasi-Dirac, it is enough to compute the cross-sections and width for one Majorana HNL, as long as we correct the cross-sections and total width with the following multiplicative factors:

$$N_b = L_{\text{int}} \times \sum_P \varepsilon_{P,b}(M_N, \tau_N) \times c_P \times \sigma_P(M_N, \Theta_e, \Theta_\mu, \Theta_\tau)$$

$$\Gamma_{\text{total}}(M_N, \Theta_e, \Theta_\mu, \Theta_\tau) = c_\Gamma |\Theta_\alpha|^2 \hat{\Gamma}_\alpha(M_N)$$

Putting it together

- Reordering the sum to arrange processes by flavours:

$$N_b(M_N, \tau_N, \Theta_e, \Theta_\mu, \Theta_\tau) = \frac{|\Theta_\alpha|^2 S_{b,\alpha\beta}(M_N, \tau_N) |\Theta_\beta|^2}{|\Theta_\gamma|^2 \hat{\Gamma}_\gamma(M_N)}$$

Only non-trivial part that we need from experiments

with the signal matrix $S_{b,\alpha\beta}(M_N, \tau_N) = L_{\text{int}} \times \sum_P \epsilon_{P,b}(M_N, \tau_N) \times \frac{c_P}{c_\Gamma} \times \hat{\sigma}_P(M_N, \tau_N) \times \Gamma_{\text{ref}}$

where the sum runs over processes *P* mediated by flavours α at production and β at decay, and $\hat{\sigma}_P$ is the cross-section computed for unix mixing angles and assuming the (small) reference width Γ_{ref}.

- The efficiencies $\epsilon_P(M_N, \tau_N)$ are typically computed on a $M_N \times \tau_N$ grid.
  To compute $S_{b,\alpha\beta}(M_N, \tau_N)$ at the physical lifetime $\Gamma^{-1}_{\text{total}}(M_N, \Theta_e, \Theta_\mu, \Theta_\tau)$, the efficiencies can be interpolated in τ_N.

Interpolation of efficiencies

Example from the reinterpretation of the prompt search

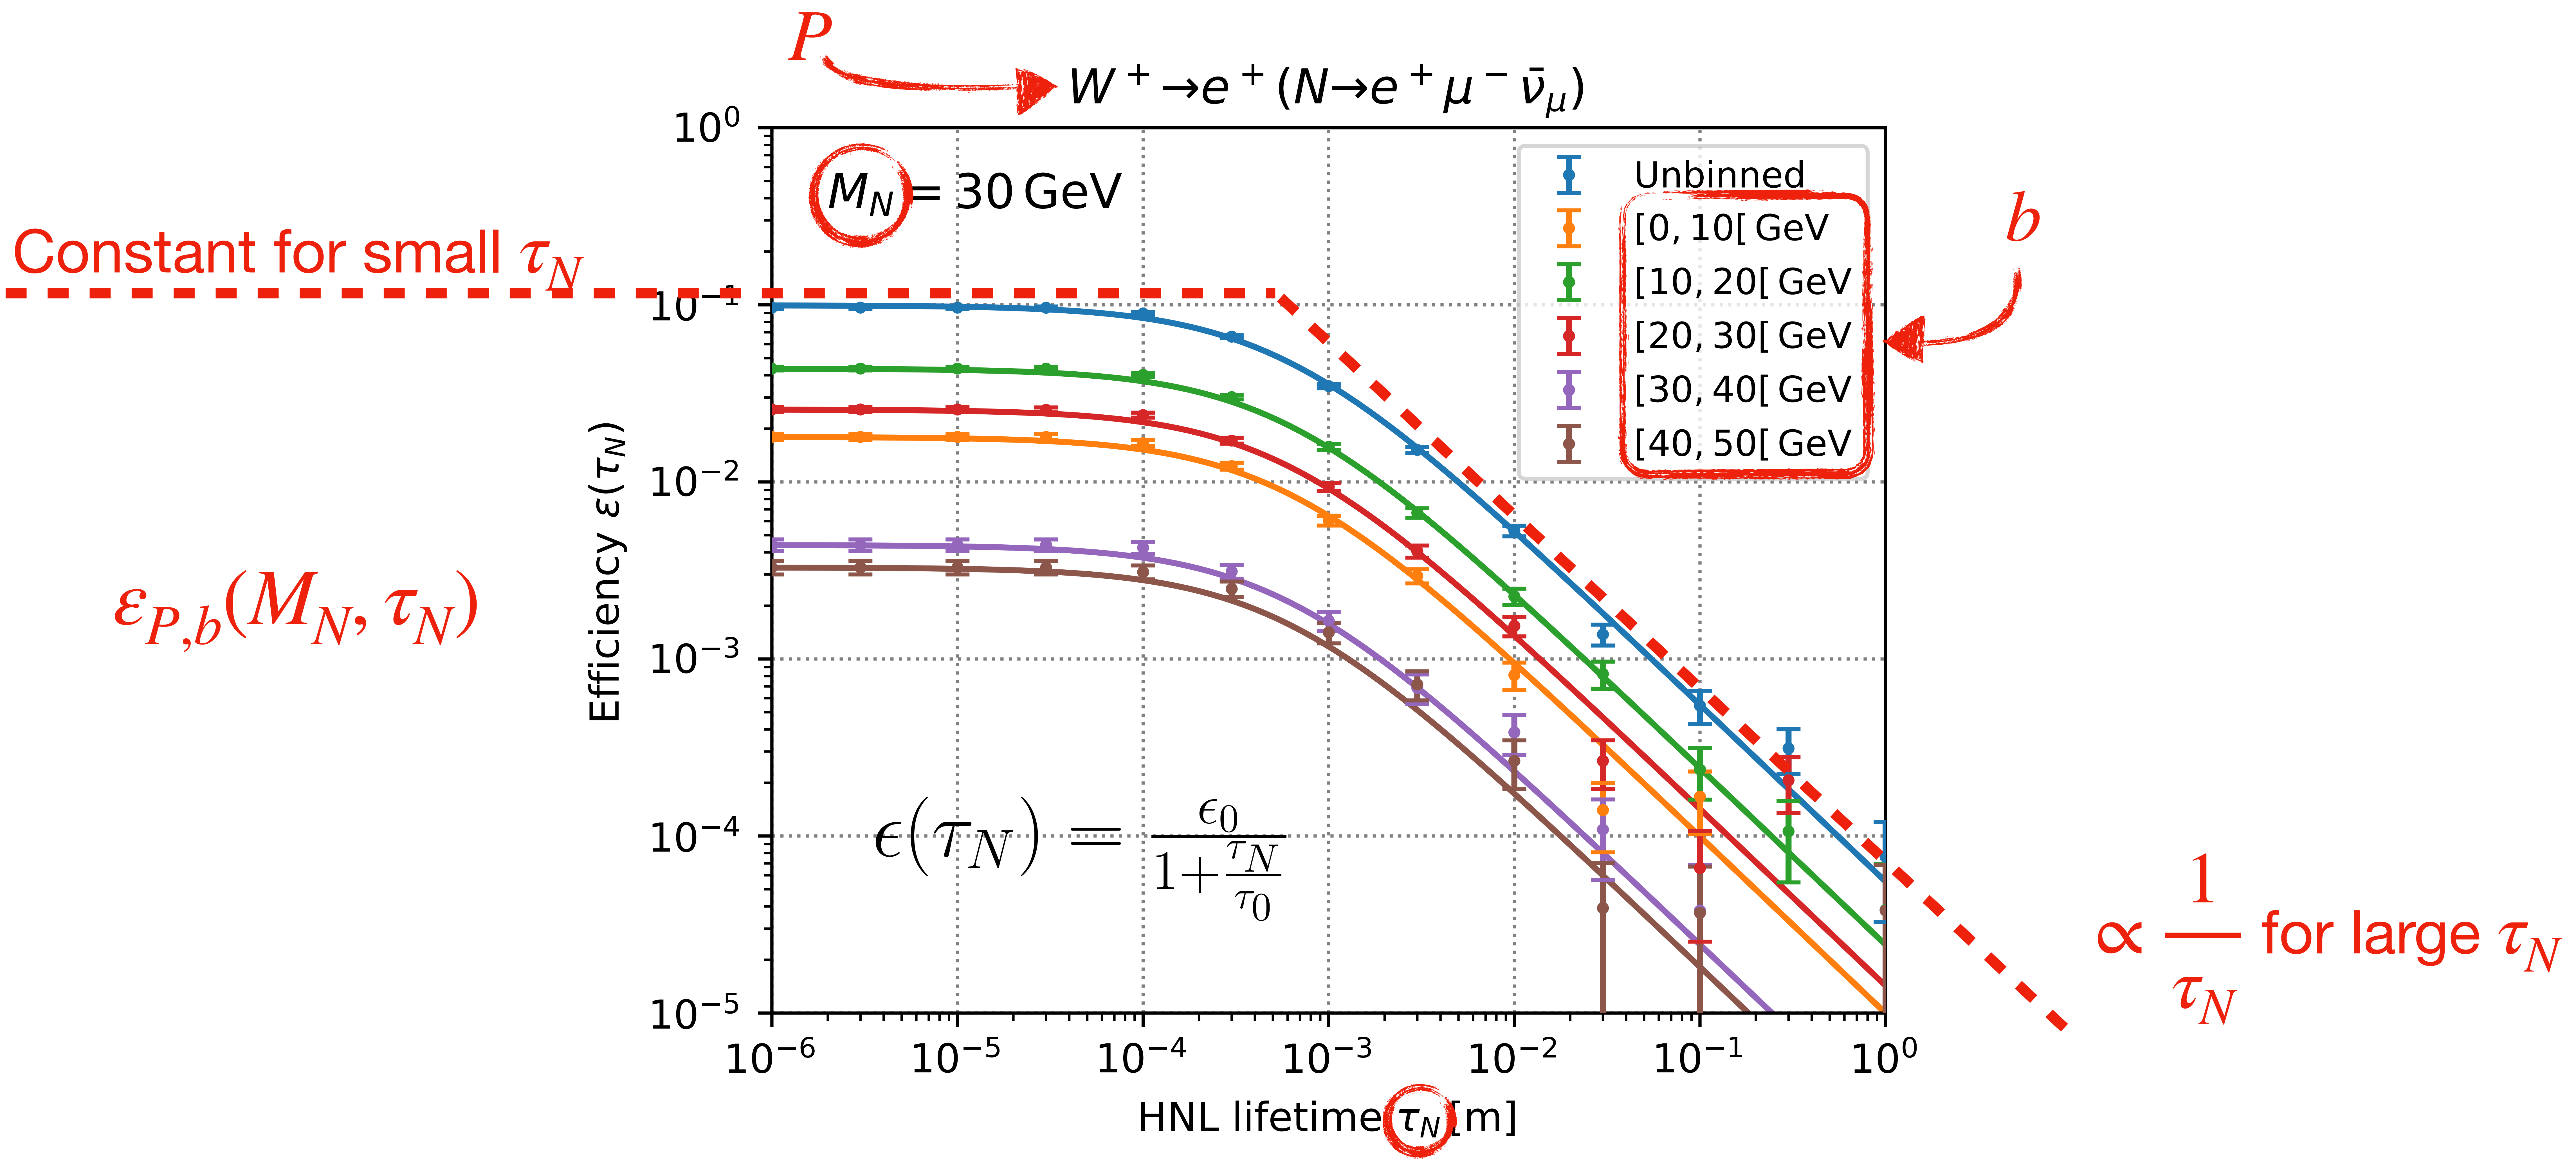

Wrapping up

- Although slightly more complicated, the sensitivity matrix approach can be generalised to work at the LHC.
- For this, we need:
    - the per-process, per-bin signal efficiencies $\varepsilon_{P,b}(M_N, \tau_N)$ for each process that contributes to the search signature, on a grid of HNL masses × lifetimes.
    - the likelihood, or a good approximation thereof,
    - the observed counts.

For more details see [Tastet, Ruchayskiy, Timiryasov: 2107.12980]

Conclusion

Conclusion

- **Heavy Neutral Leptons** are a well motivated extension of the Standard Model.
- They can be searched for both at dedicated experiments and at the LHC.
- The "one Majorana HNL mixing with one flavour" benchmark cannot explain ν masses.
- More realistic benchmarks such as the ones recommended by the FIPs 2020 report can give us a **qualitative** idea of how well the parameter space is covered.
- For **quantitative** applications such as global scans or Bayesian analyses, we need to be able to accurately interpret the search results for arbitrary model parameters.
- To this end the sensitivity matrix approach can be generalised. It requires that experiments report their likelihood as well as per-process, per-bin efficiencies.