Jacopo Siniscalco

UCL

A new nuclear recoil model for low energy searches in the LUX-ZEPLIN experiment

Who am I?

Grad student @ UCL

Focus on LZ experiment

WS2024

Flamedisx

LZFlameFit

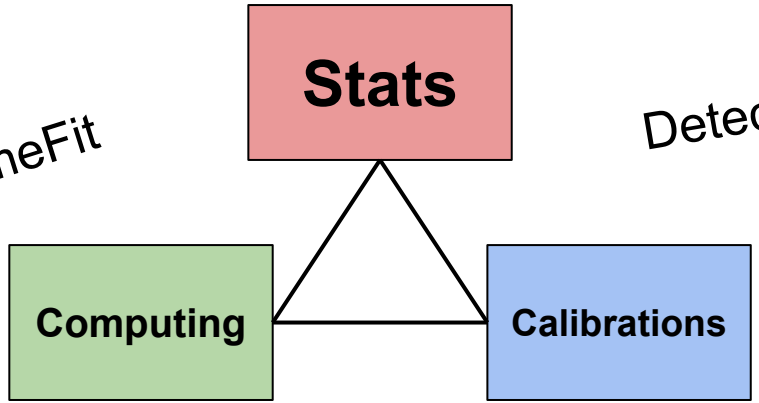

Detector tuning

Offline computing

NR model

DD_MCMC

Sims pipeline @ SLAC

LUX-ZEPLIN 101

- Dual-phase Xe **T**ime **P**rojection **C**hamber
- optimised for **WIMP** dark matter searches
- 7 T **L**iquid **Xe** target
- Top & bottom PMT arrays capture signals within TPC
- 1 mile underground at SURF, South Dakota

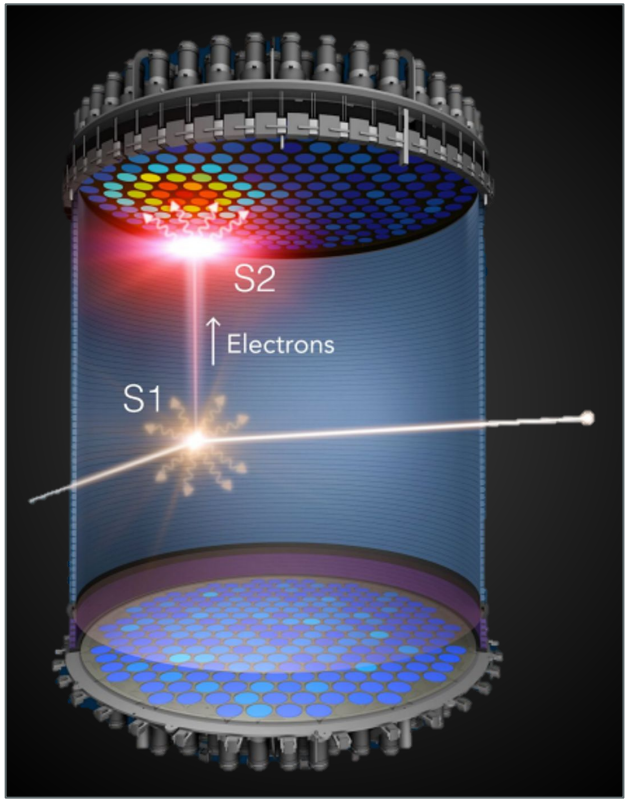

Dark matter & WIMPs

Weakly **I**nteracting **M**assive **P**articles

- GeV-TeV scale
- Theoretically well-motivated (eg. MSSM mass range)
- Detection technology is mature & scalable

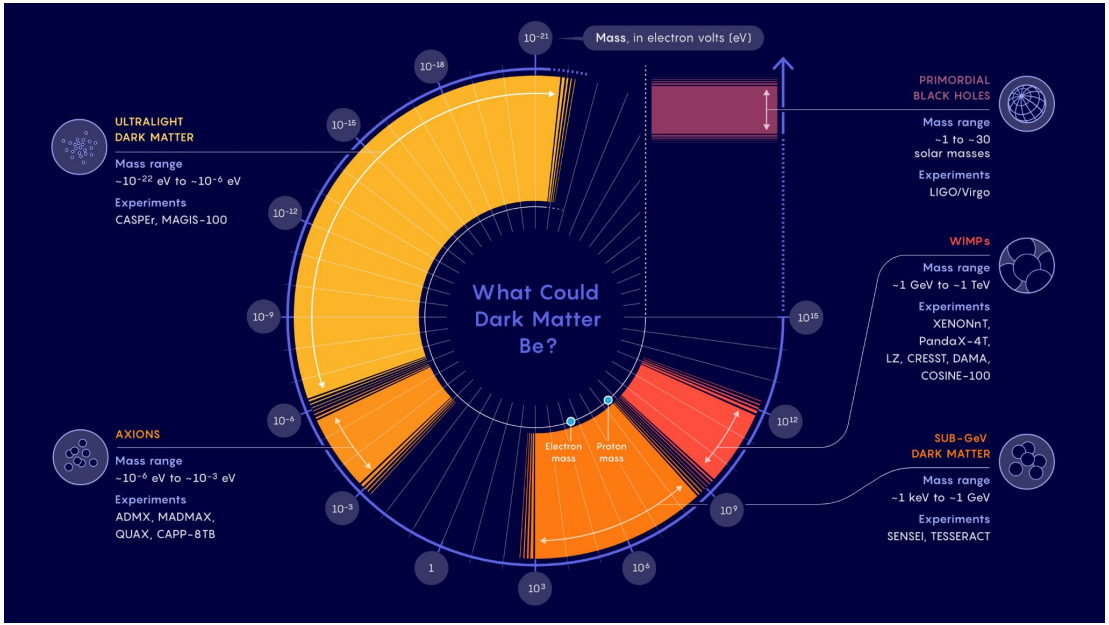

The Search for Dark Matter Is Dramatically Expanding, **Quanta Magazine, Nov 2020**

From Recoils to Signals

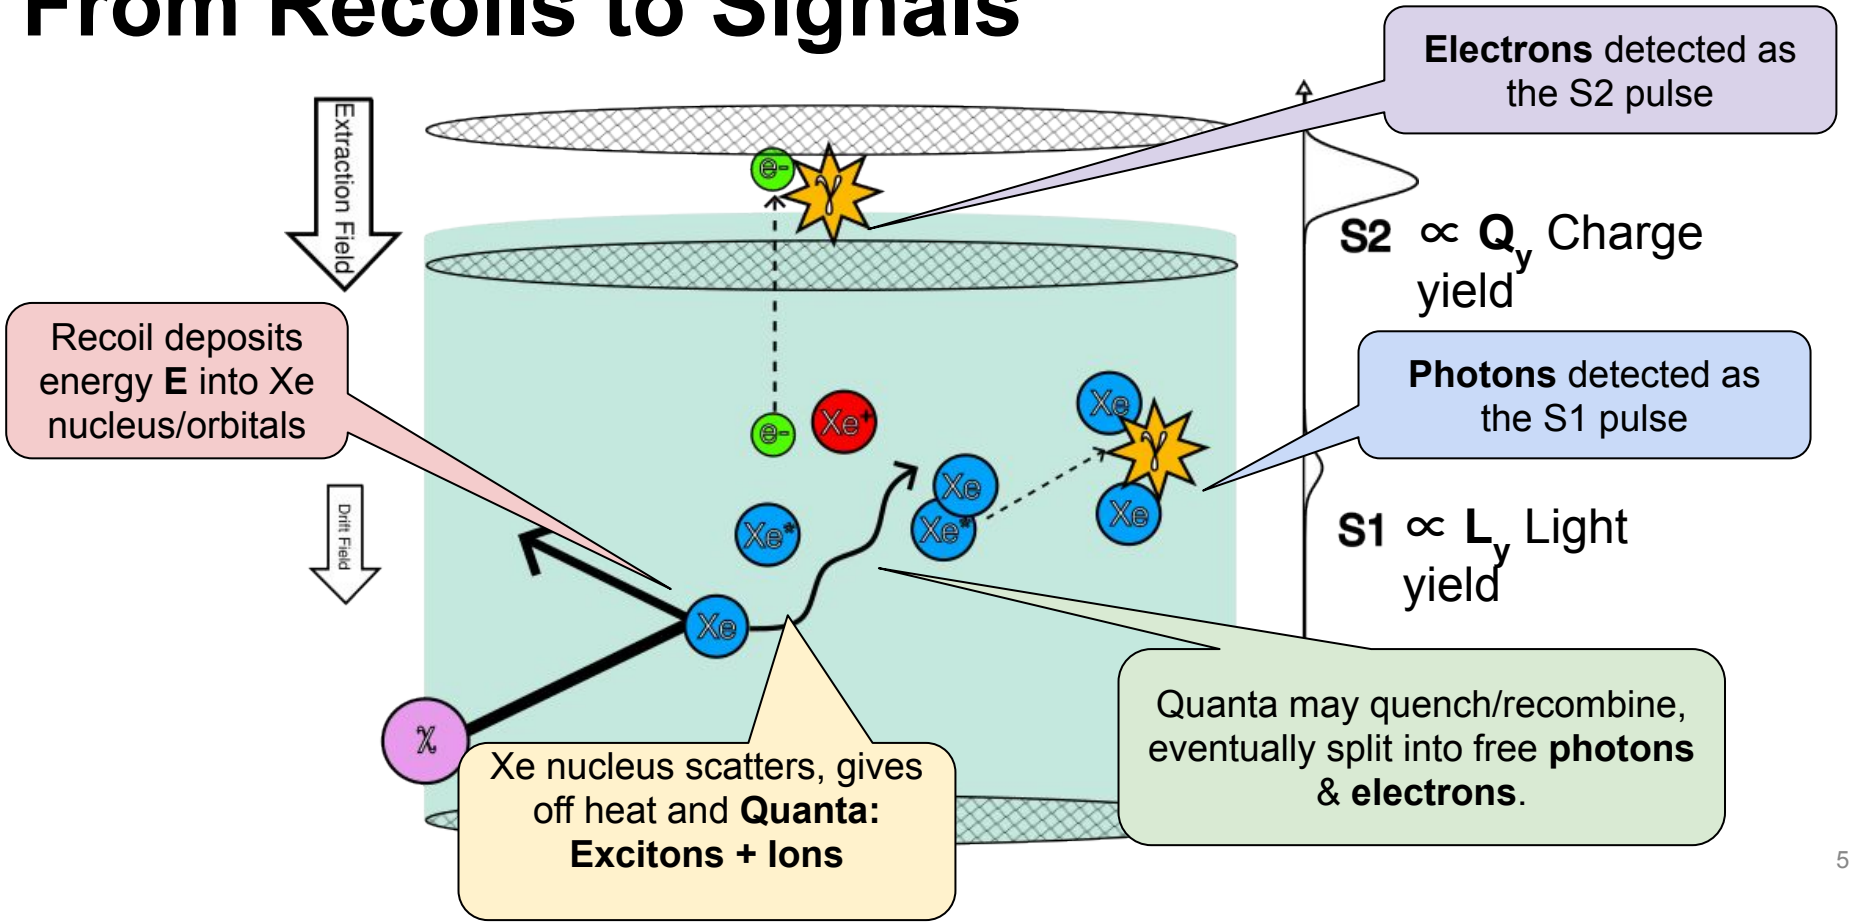

What LXe TPCs are good at

- Ratio of S1/S2 sensitive to the dynamics of the interaction
- Can trivially separate *electronic recoils* (ERs) from *nuclear recoils* (NRs)

"Natural" to envision model as predicting a set of (skew) **gaussian** bands

ERs
Betas
X-rays
gammas
Neutrinos

NRs
WIMPs
Neutrons
Neutrino CEvNS

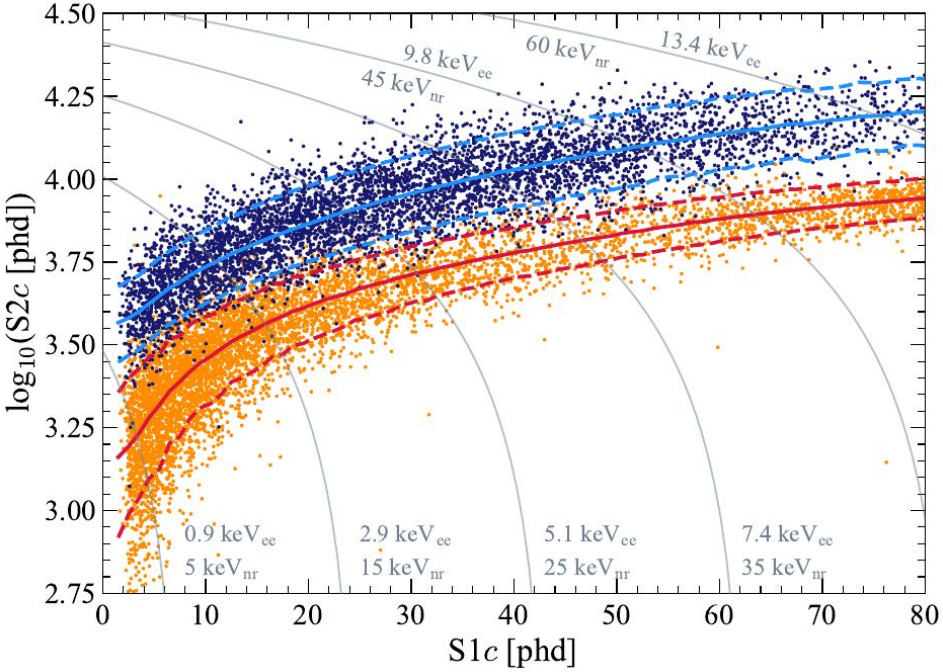

Recoil bands - the bread & butter

- Model tuning heavily focused on “fitting bands”
- Fewer degrees of freedom
- Likelihood is faster to compute & less sensitive to statistical fluctuations
- Can ignore systematics at threshold & endpoint

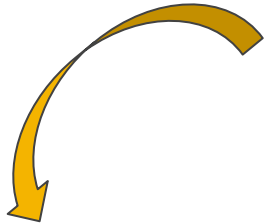

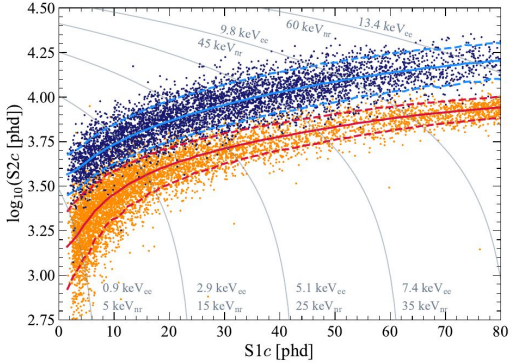

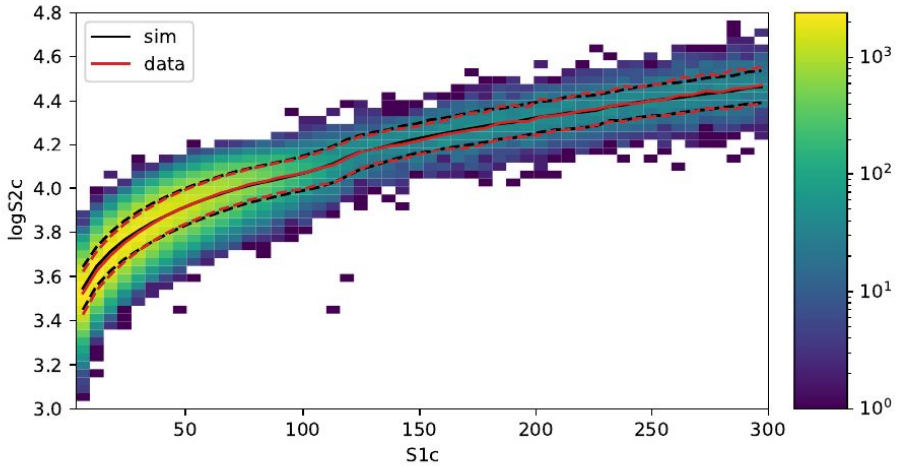

What did the response model need to provide?

- Biggest driver is the “ER leakage”
  - Proportion of ER band below some NR quantile
- No need to model low-E
- Can accept some error in yields vs energy - as long as bands are consistent

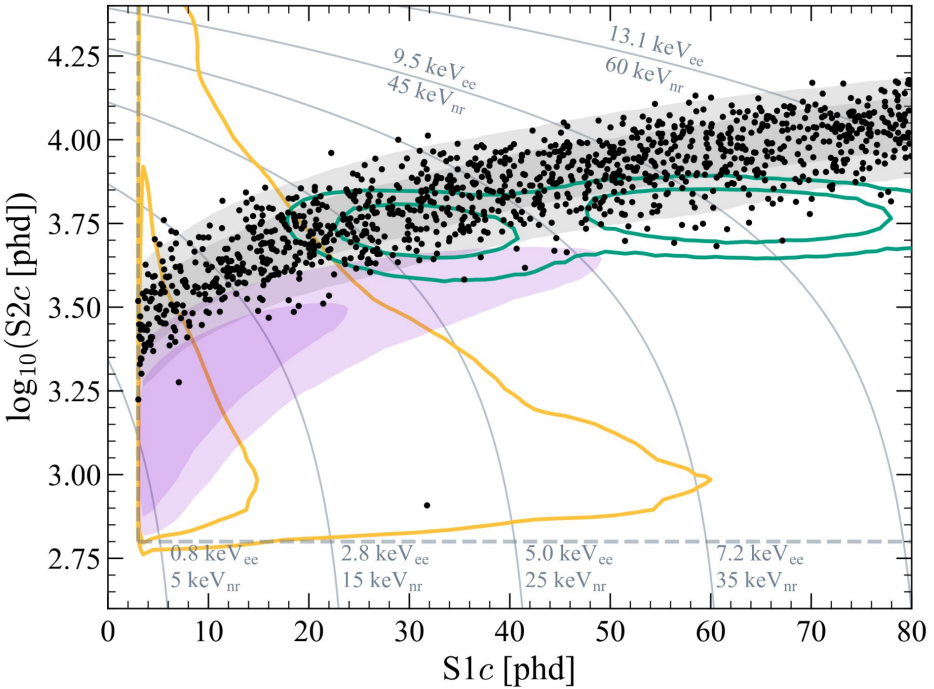

Previous searches

- WS2024 - restricted to WIMP mass > 9 GeV
- Fewer systematics
- Calibrations approach was very successful in this regime

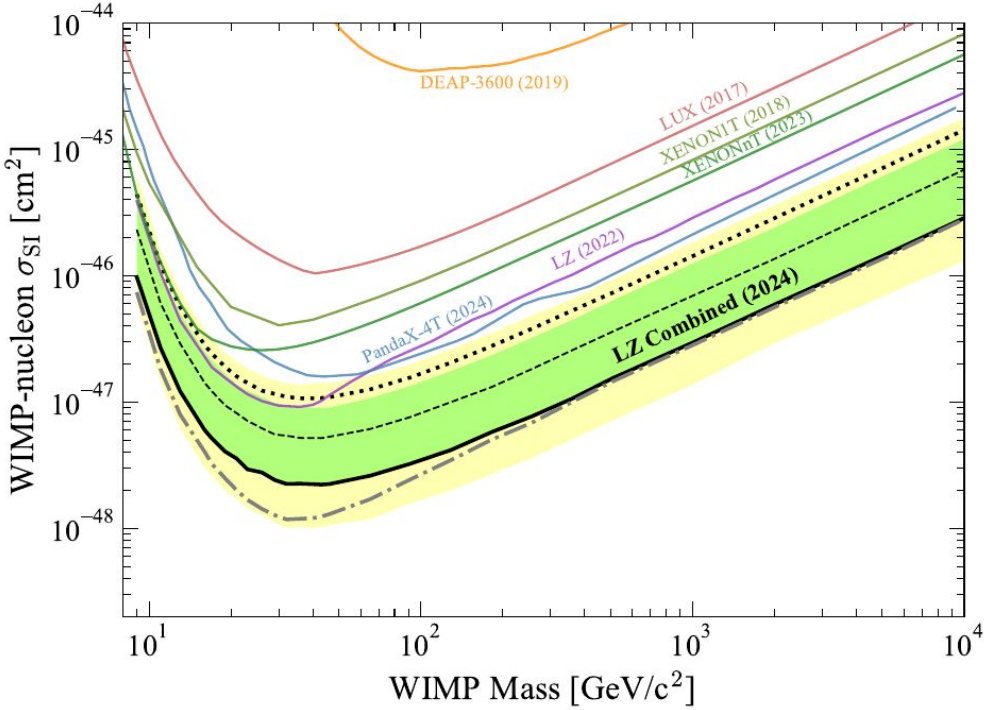

WS2025 - The low energy search

- Dec 8th - release WS2025 paper on arXiv
- World-leading limits at most masses for WIMPS ≤ 9 GeV - **L**ow **M**ass **W**IMPs
- For the first time, hit the “neutrino fog” - ^{8}B CEvNS
  - Main background is now an NR source

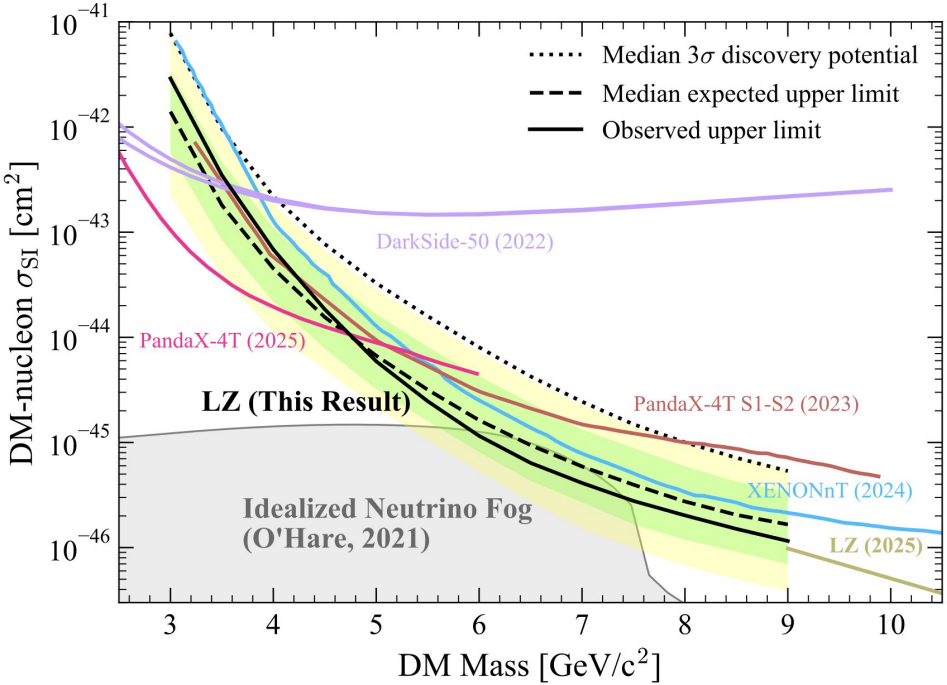

https://arxiv.org/abs/2512.08065

WS2025 - The low energy search

- WIMPS ≤ 9 GeV - recoil energy ~ few keV
- Detection threshold becomes important
  - Small changes in yield can have big effect
  - Become sensitive to physics below threshold due to fluctuations

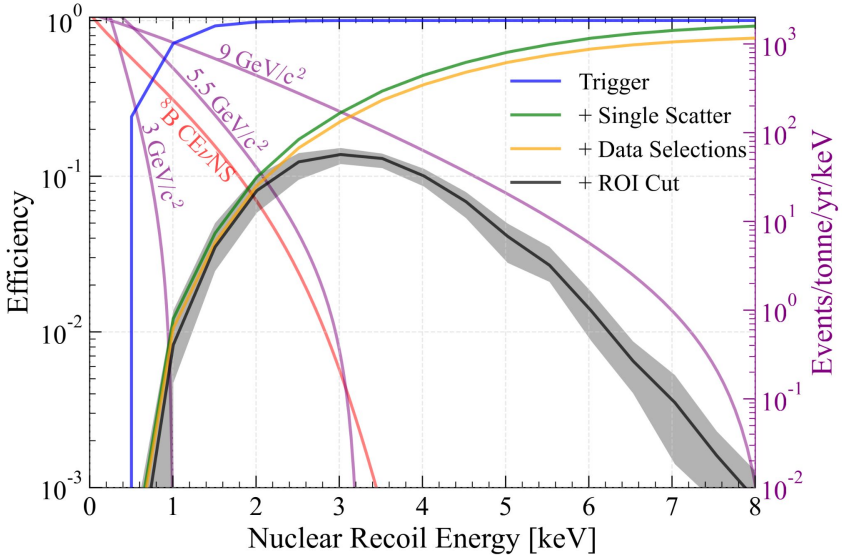

^{8}B CEνNS

- ^{8}B neutrinos expected to interact coherently with Xe nucleus
  - **C**oherent **E**lastic **neutrino** **N**uclear **S**cattering
- Validate SM prediction
- “Looks” like dark matter - we’re a DM experiment, we better see it!
- ^{8}B has a known flux - measuring in LZ requires understanding NR yields

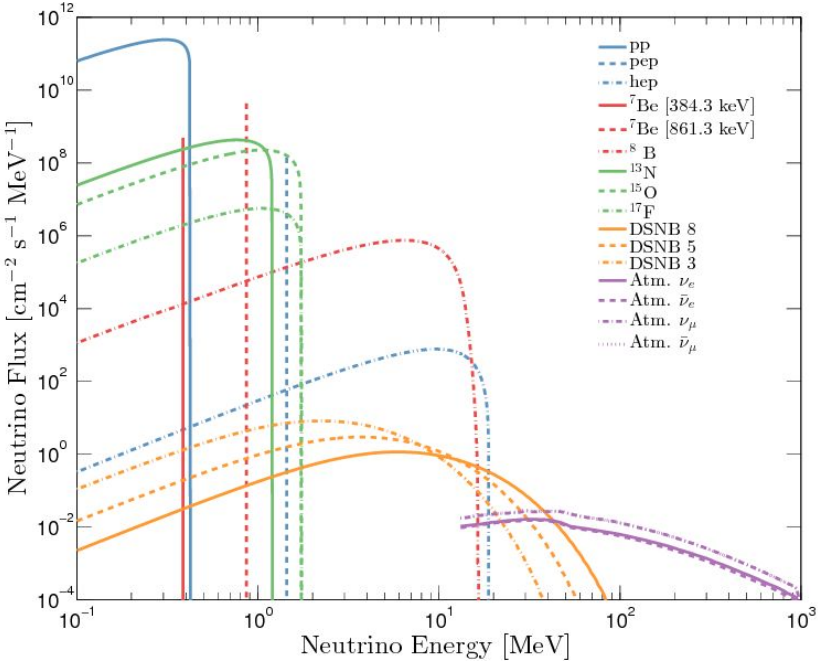

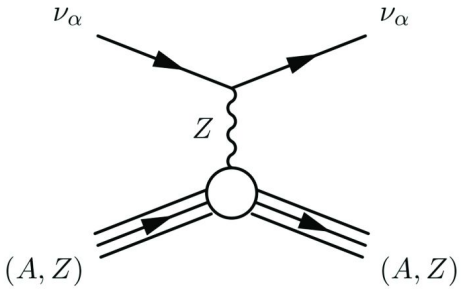

arXiv:1505.08061v2

Why the last result is different

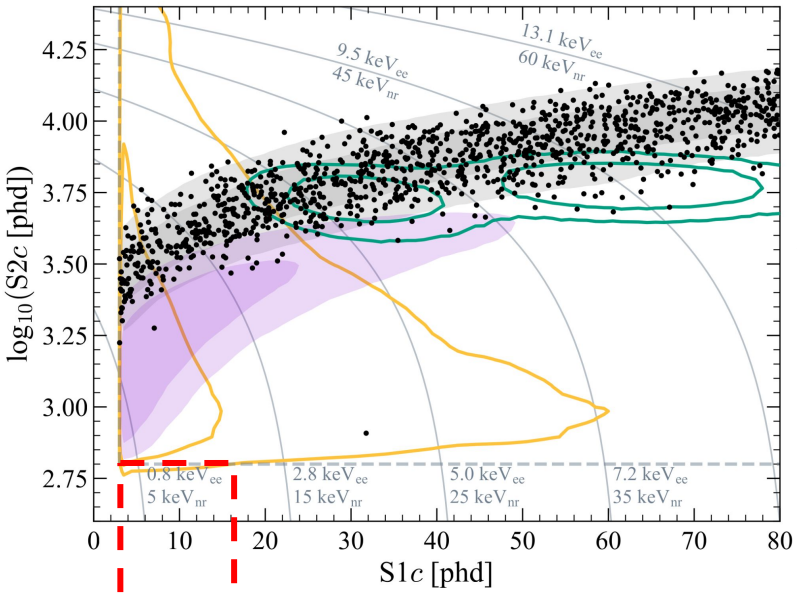

WS2024 **R**egion **O**f **I**nterest

WS2025 **R**egion **O**f **I**nterest

Why the last result is different

- No ER background
- No longer dealing with ER/NR bands
- Small changes in yields/fluctuations amplified by threshold
  - Becomes very important to quantify model uncertainty & include in final inference
- All of this, at low energy!

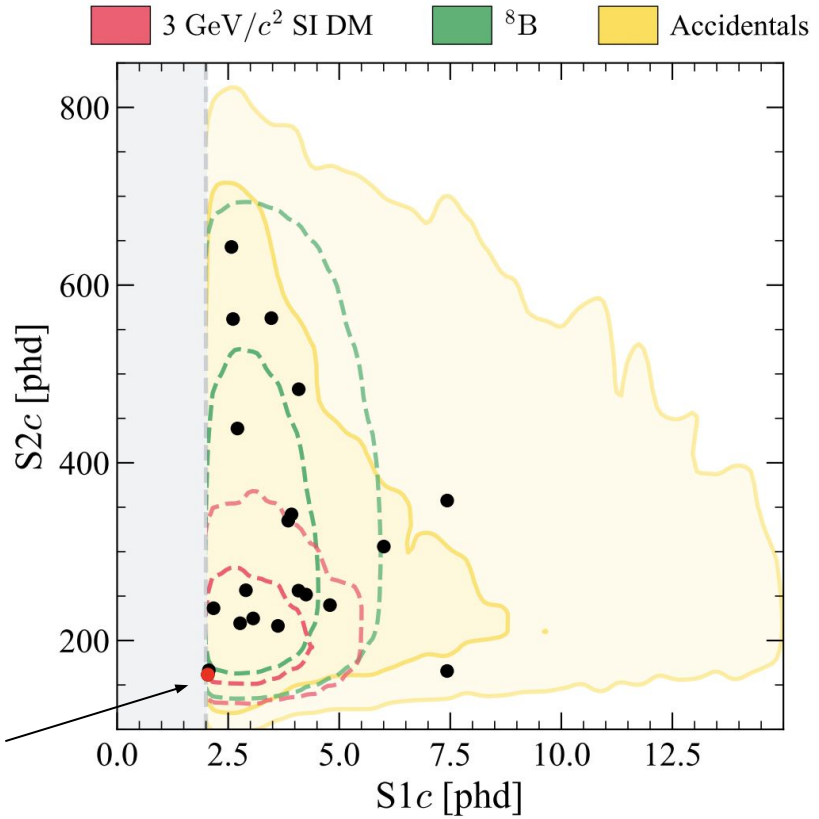

A change in paradigm

- Full 2D binned likelihood
  - More sensitive to systematics - energy spectrum, detection efficiencies, etc
  - Computationally expensive
- Extend to lowest energies
  - Detector systematics
  - endpoints of recoil energy becomes important
- Need to translate systematics to final model

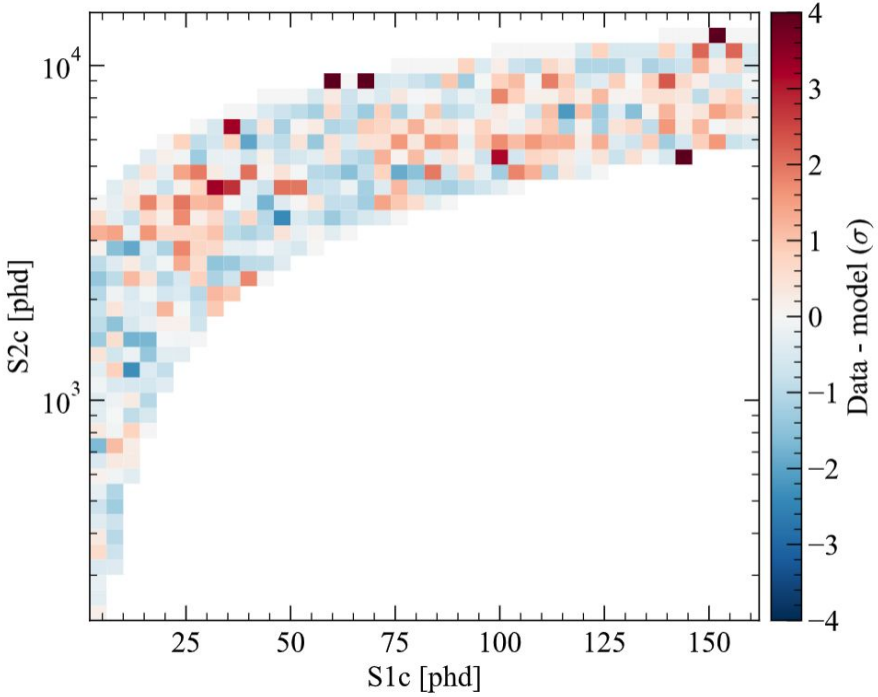

DD neutrons - the golden dataset

- “On demand” neutrons with Adelphi DD109 neutron generator
- Tunable luminosity
- monoenergetic neutrons
- Plasma production pulsed & synched to trigger

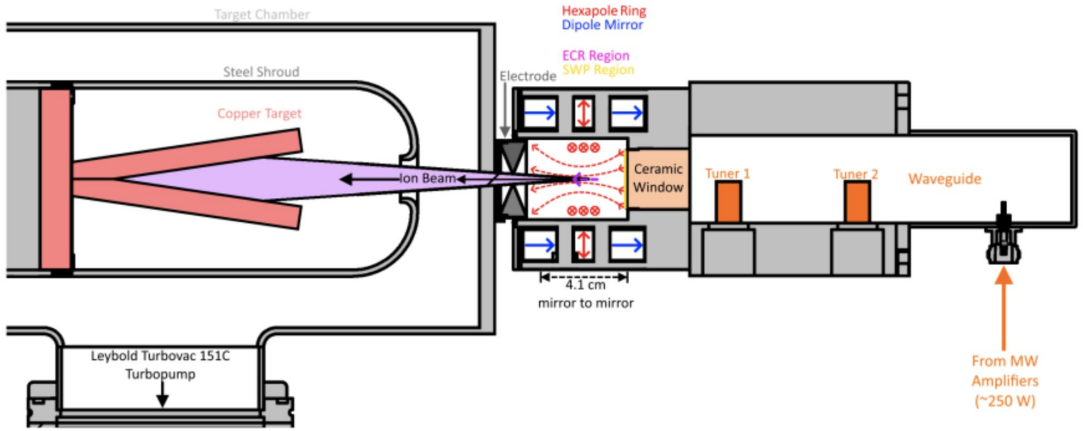

Deuterium-Deuterium Fusion

$^{2}H + {}^{2}H \rightarrow {}^{3}He + n\ (2.45\ MeV)$

C. Rhyne, Development of tagged low energy d- and h- based scintillator reflector neutron sources and characterizations of a d-d neutron generator for calibrations of the LZ detector (2024)

DD neutrons - the golden dataset

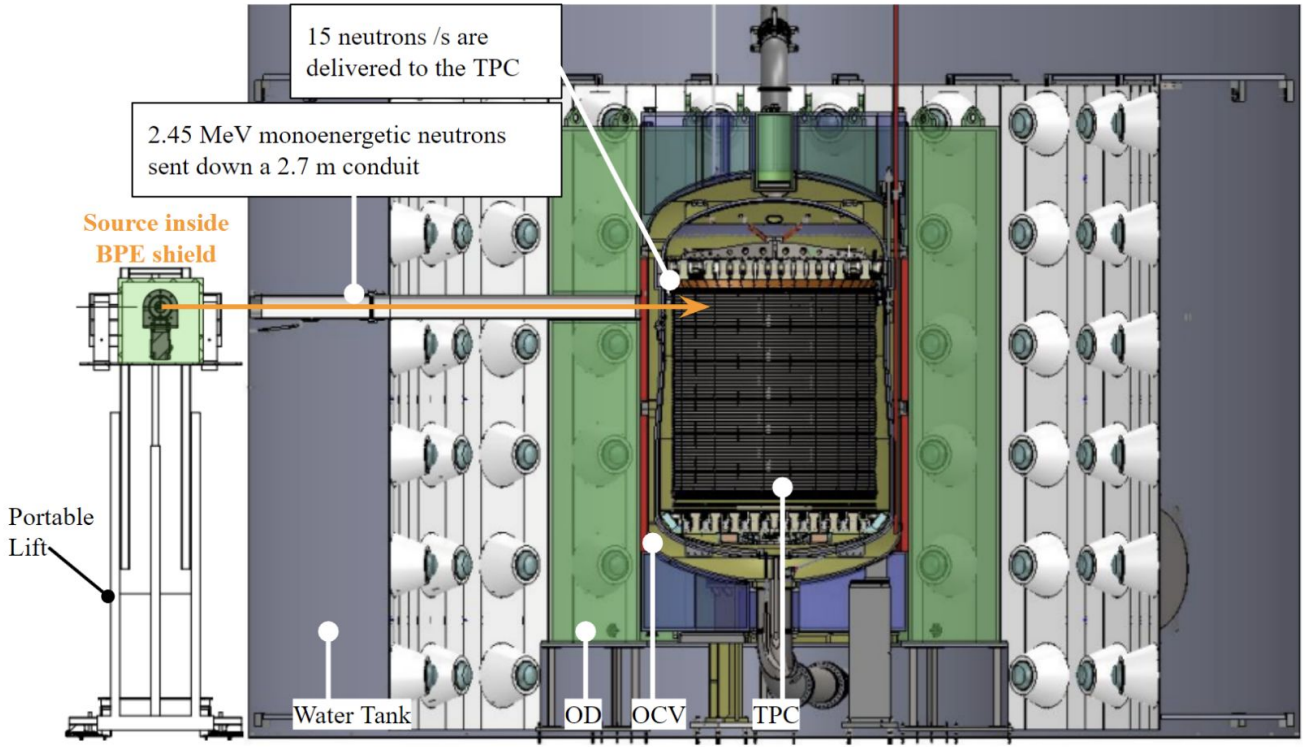

The clean data

20k events

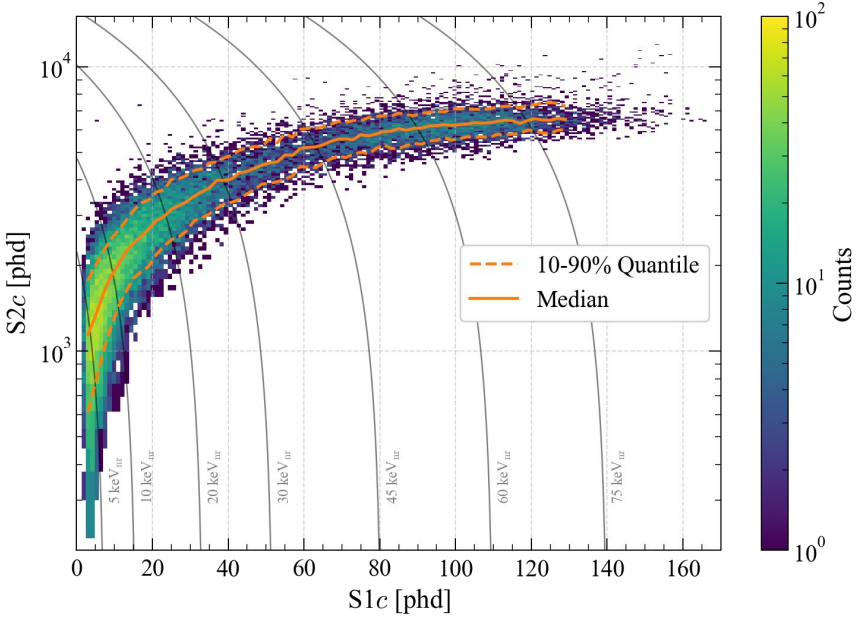

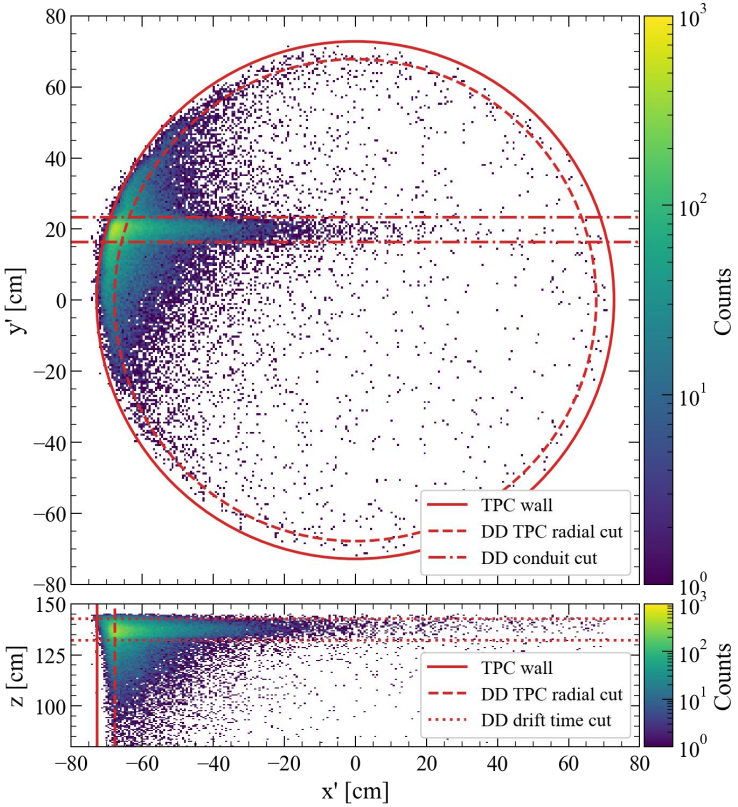

Model pipeline & incorporating systematics

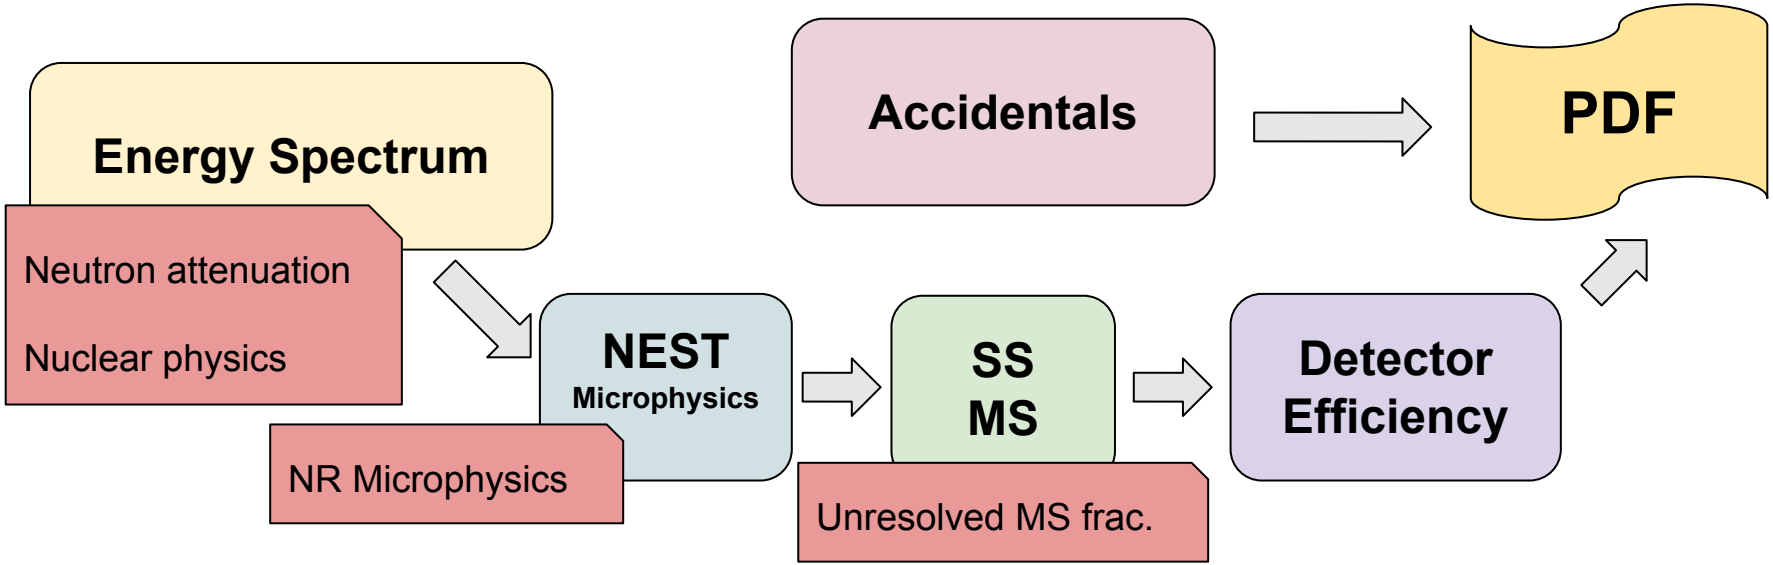

Not-so-monoenergetic neutrons

- DD neutrons not perfectly monoenergetic
- Neutrons may be attenuated by conduit/cryostat/DD apparatus

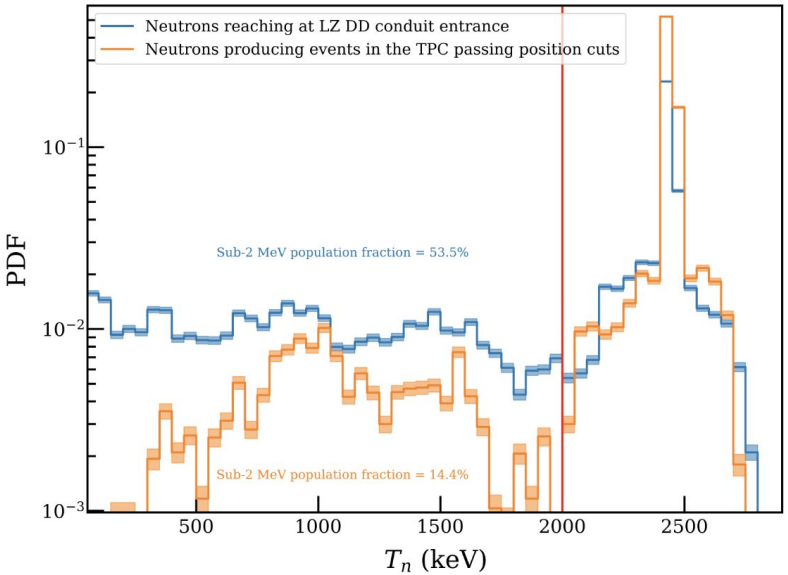

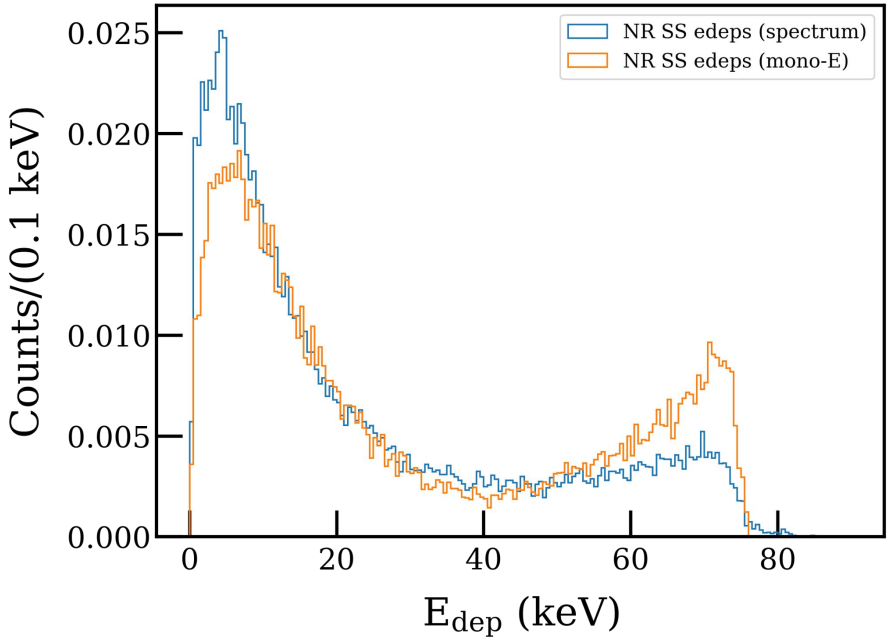

Recoil spectrum simulations

- Simulate collection of energy depositions in TPC
  - Sample from this collection to obtain spectrum
  - Include positional dependence of neutron attenuation
- Allow attenuation fraction to float
  - Best-fit ~13%
  - Uncertainty encoded into model output

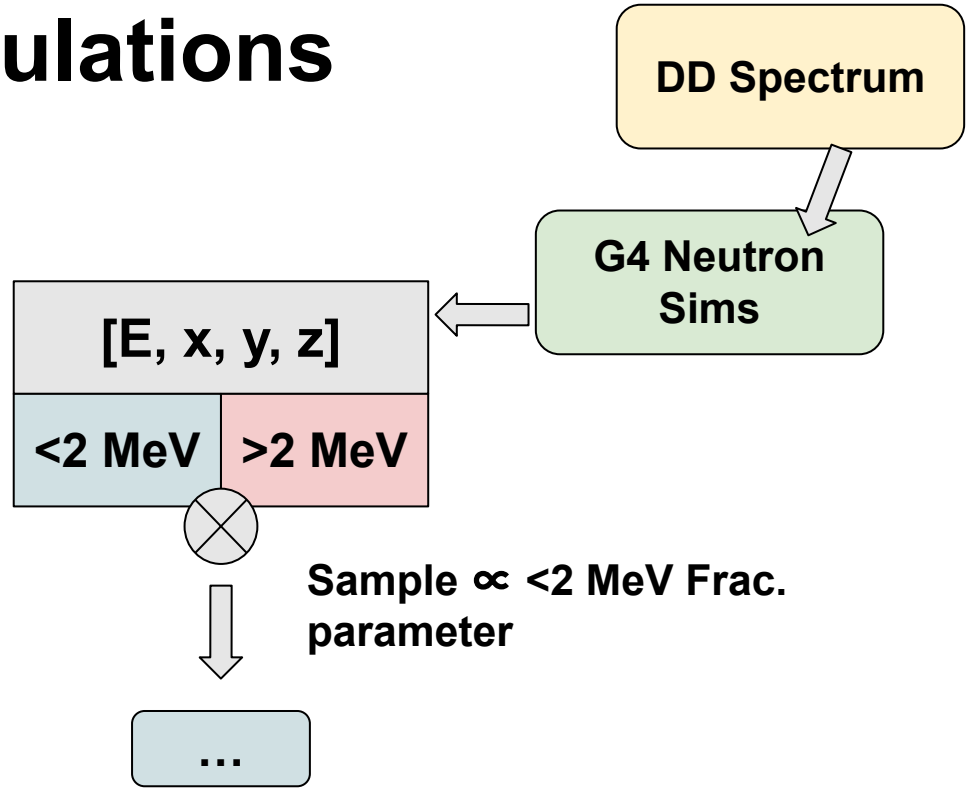

Nuclear physics

- LZ G4 sims package uses ENDF neutron physics
  - Allows to generate recoil spectrum in Xe
- Nontrivial uncertainties in elastic x-sections
- How to model?
  - Allow interpolating between ENDF & JENDL neutron libraries
  - best fit ~80% ± 10% ENDF

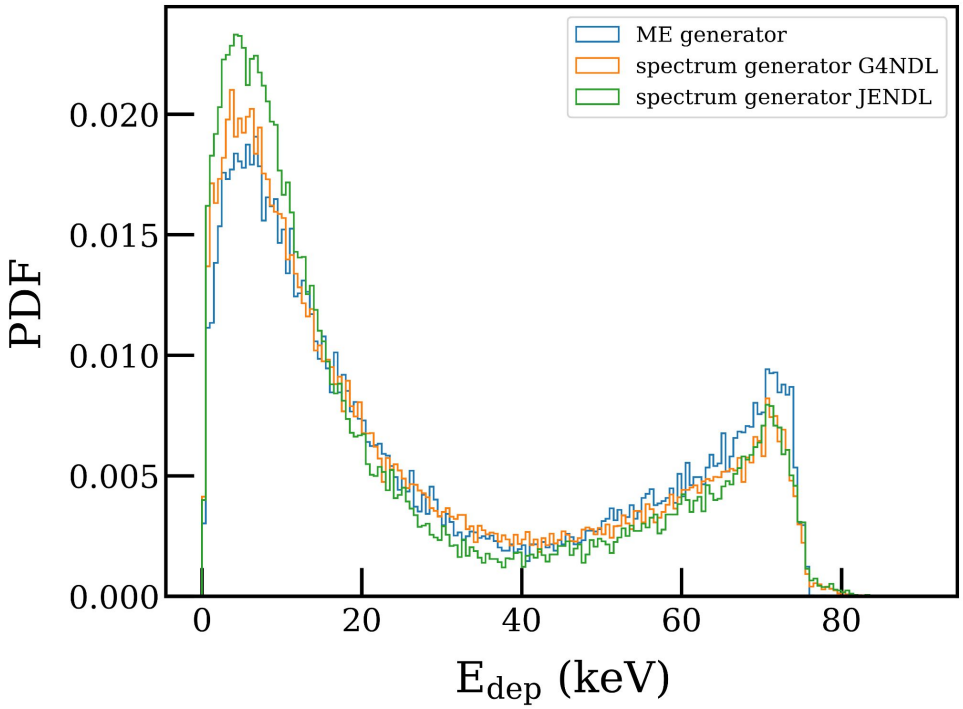

Noble Element Simulation Technique

- Physically informed empirical Xe response model
- Different models for different interactions
- NR response is modelled in stages

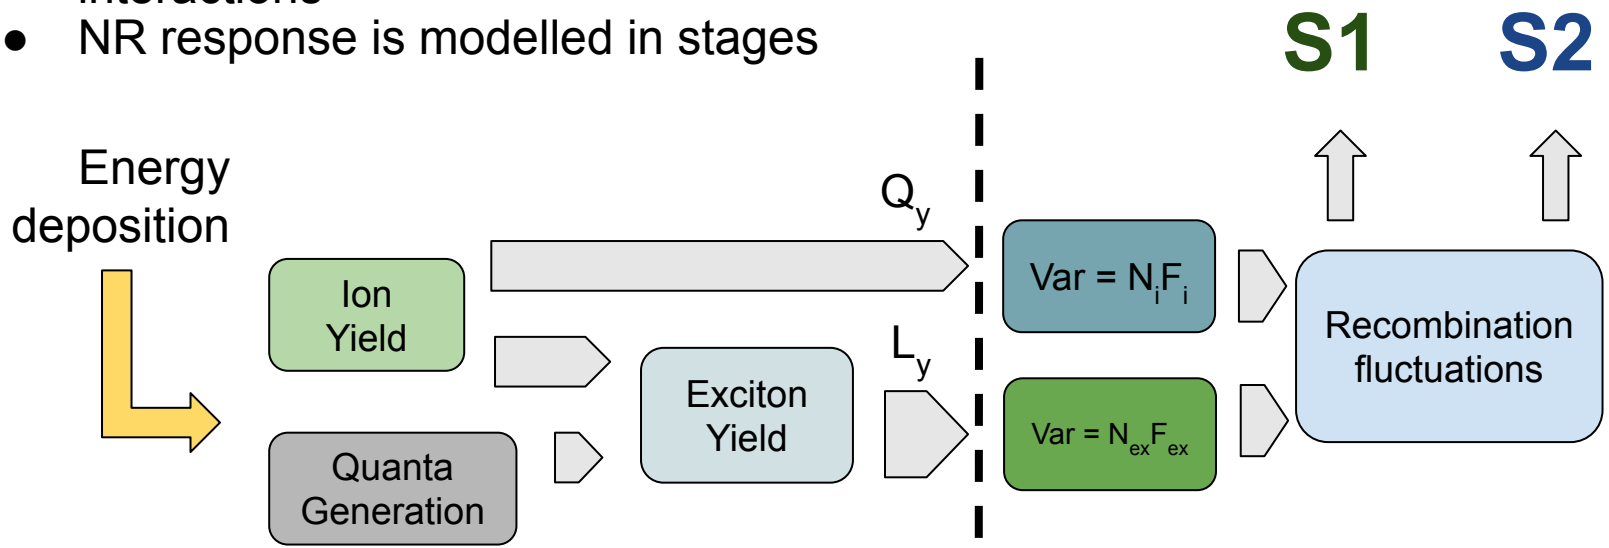

Microphysics model changes

From earlier analyses: key aspects of the fluctuation modelling inconsistent with data

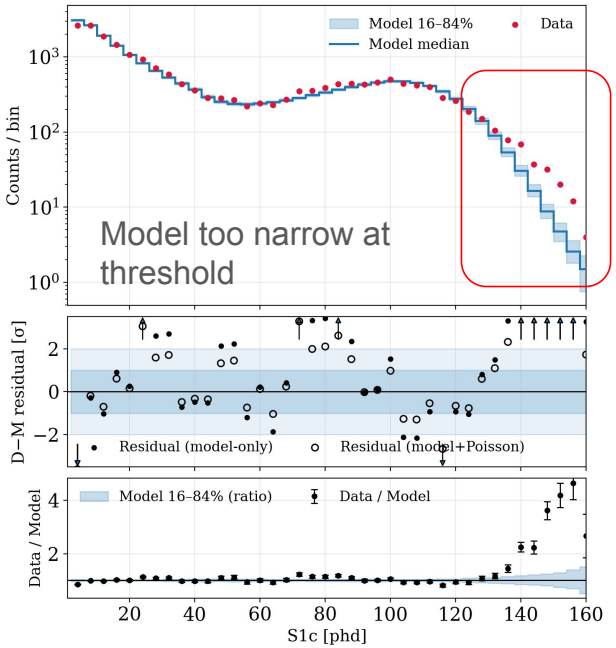

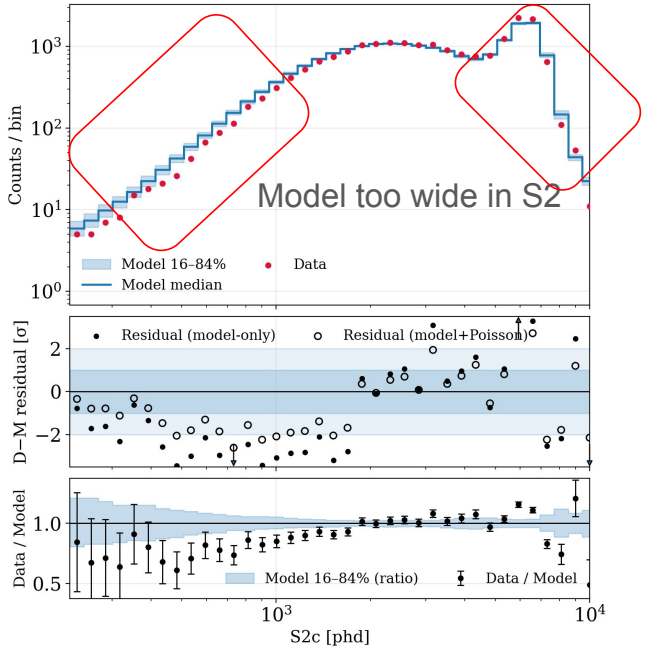

Model additions

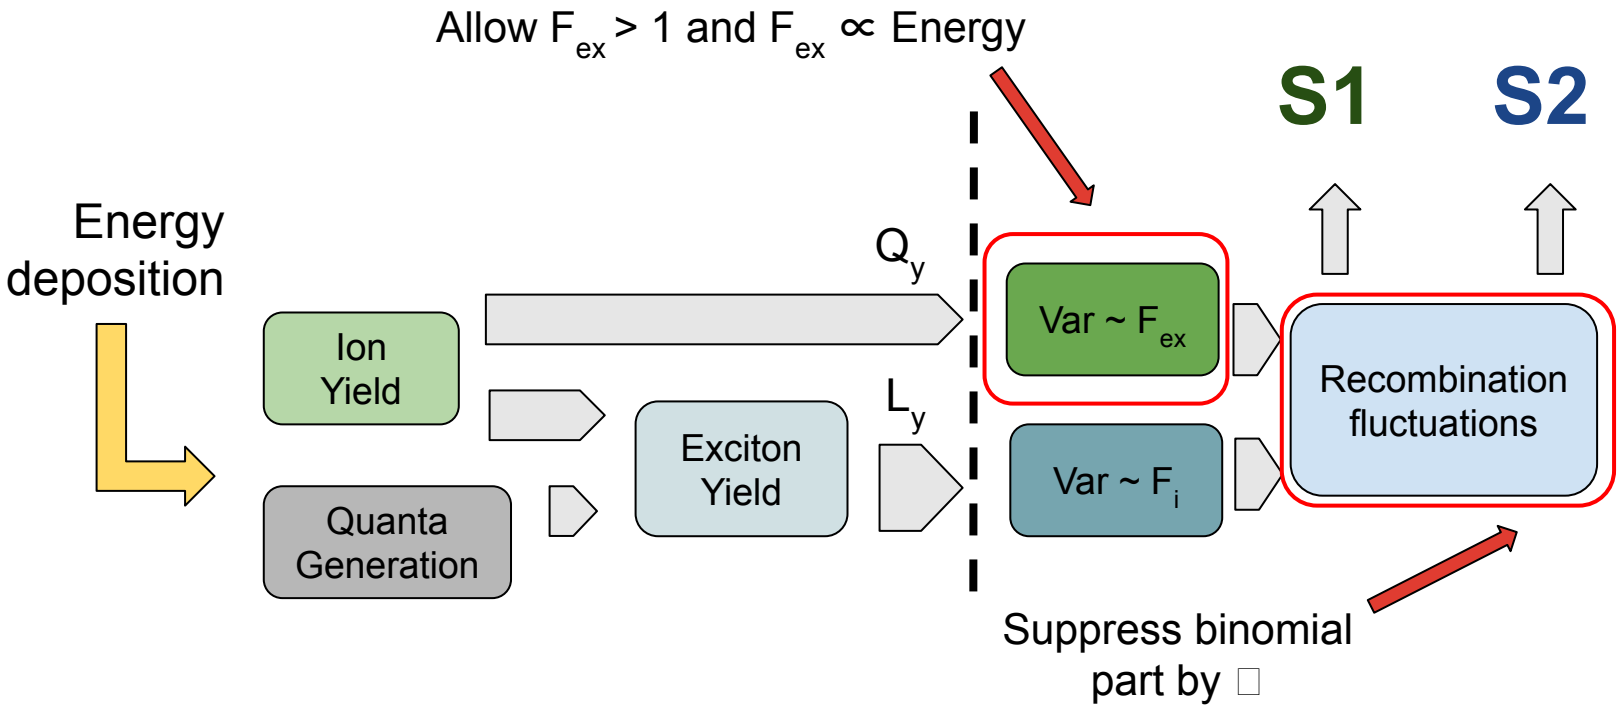

Fits with updated model

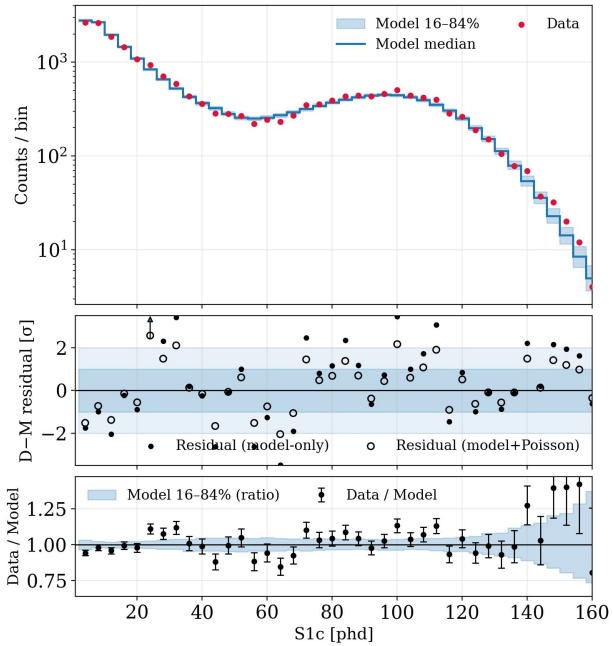

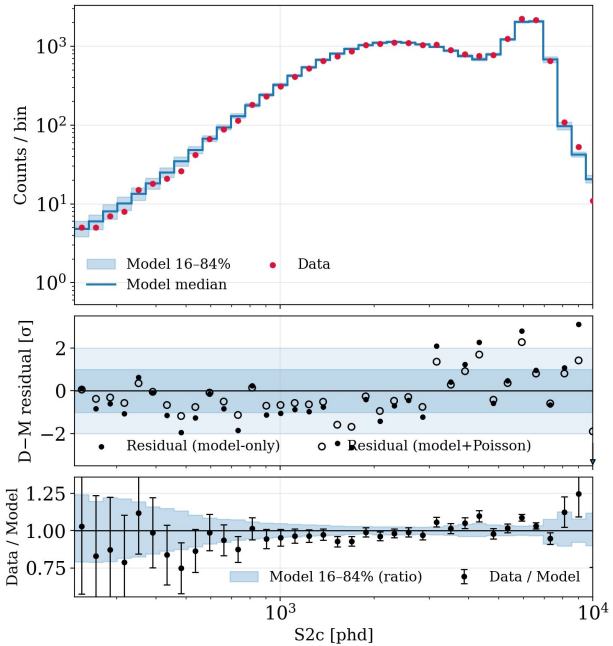

F_{ex} & low-E signals

- F_{ex} changes especially impactful to low-E search
- F_{ex} increases with E/N_{ex}
  - constant F_{ex} would bias model at low energy

- Still allow possibility of F_{ex} >1 at threshold

$Var(N_{ex}) = F_{ex}\ N_{ex}$

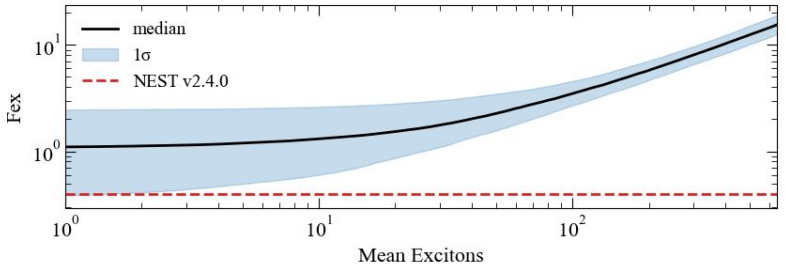

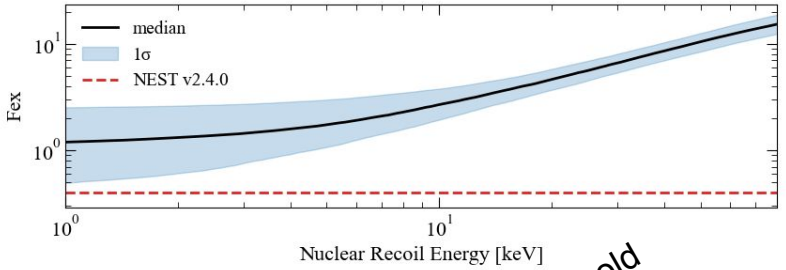

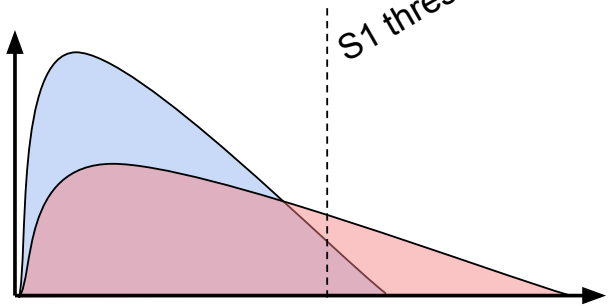

The Fantastic Four

- Potentially unknown physics at low energy
- $L_y = N_q - Q_y$: forced anticorrelation
- Allow extra freedom in yields at lowest energies

$$\text{yield} \times \left(1 - \frac{1}{1 + \left(\frac{E}{\text{scale}} \right)^{\text{shape}}} \right)$$

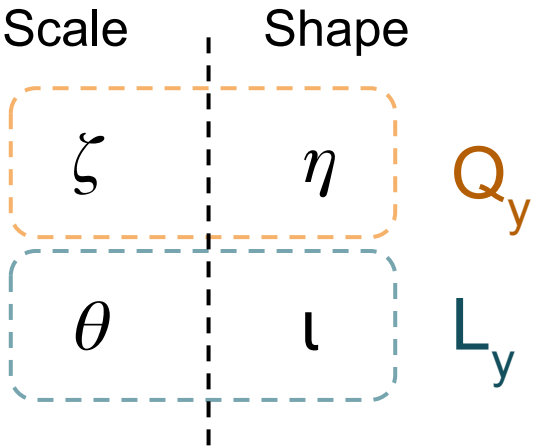

How low do we get signal?

Very hard to see low energy events:

- Geometric effects
- Pulse finding & classification inefficiencies
- S1 threshold effects

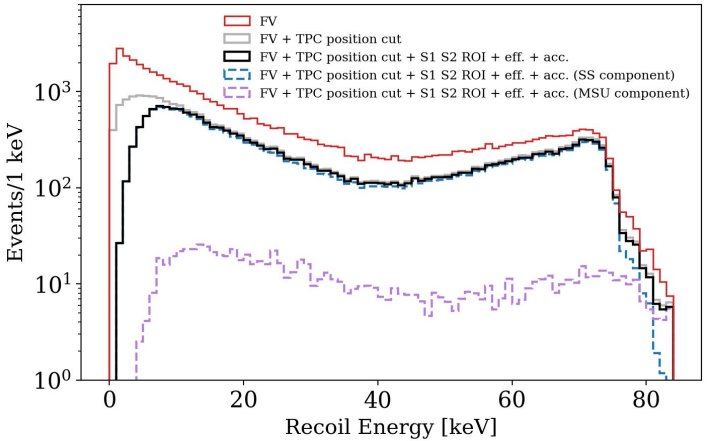

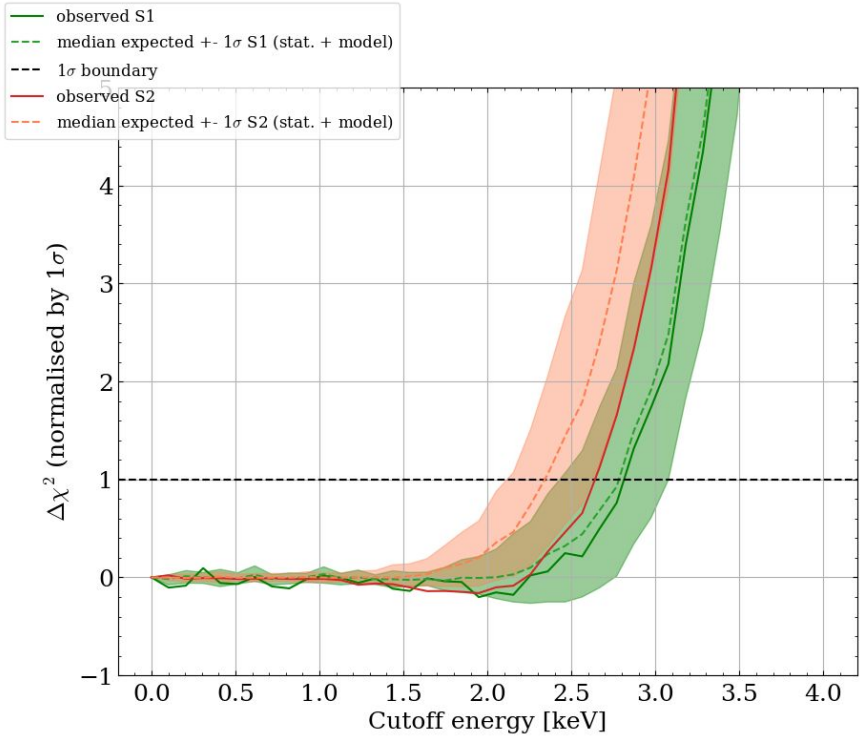

Including global data

- Use our data as best we can, augment with global data for the rest
- Implemented as priors in our fitting procedure
- Allows to constrain model as much as possible

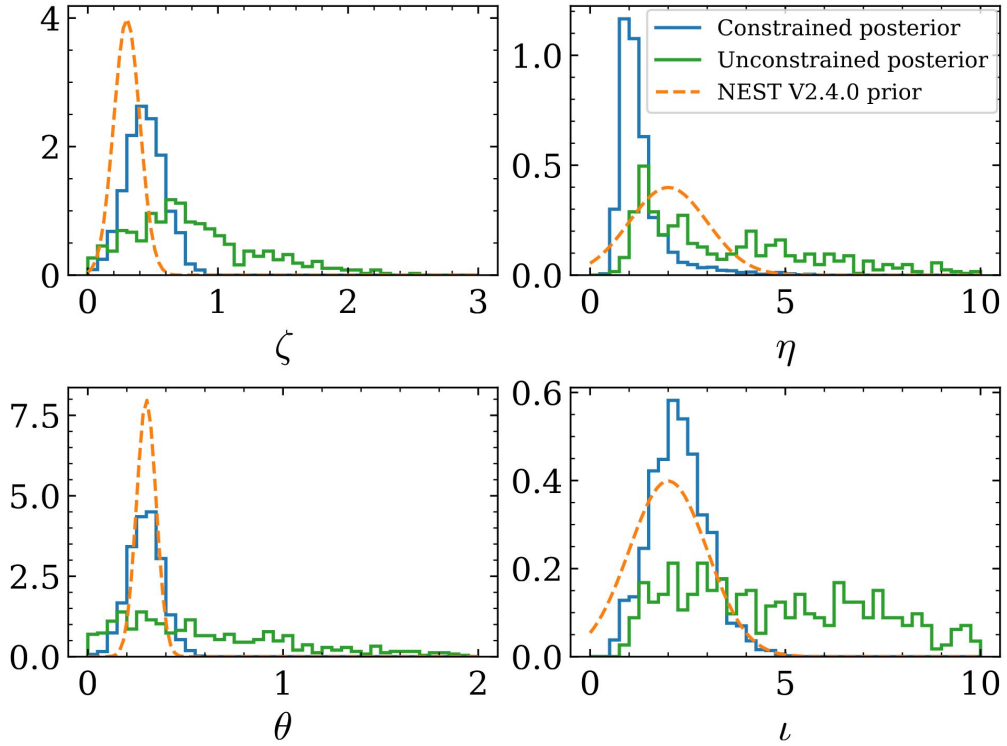

Multiple scatters

- MFP of neutrons in LXe is O(10 cm) - MS likely*
- If scatters are too close: Unresolved Multiple Scatter (**MSU**)
- Especially important for S2 tail
    - Needed to properly fit DD endpoint

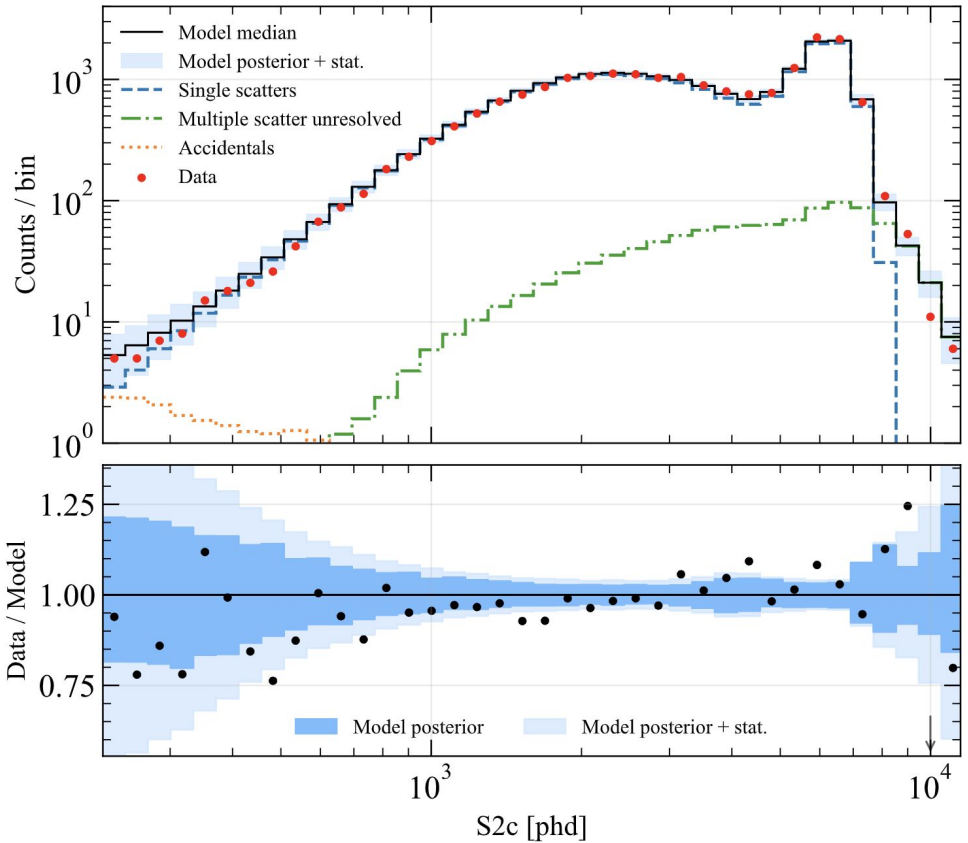

Accidentals

- “Background noise” of lone S1/S2s can get paired by event classifier - creates a fake event

DD trigger timing

Higher single S1/S2 rate during DD runs

Need to include in model

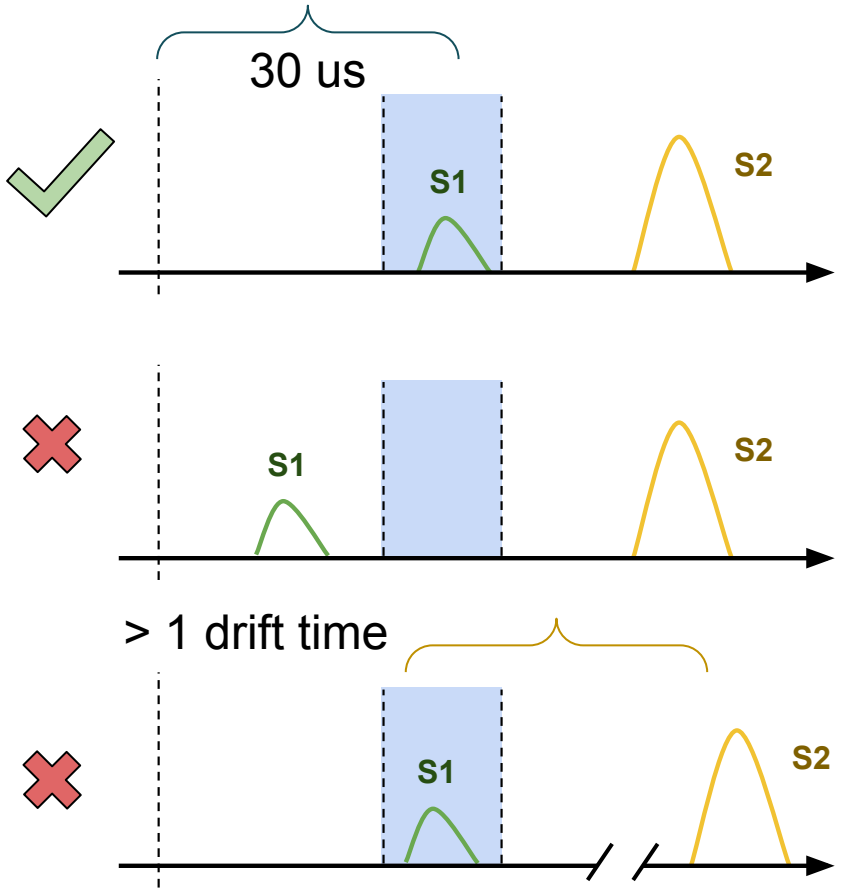

Accidentals Model

- Use sidebands to produce accidentals PDF - **data driven model**
- Estimate upper limit on accidentals rate < 3.7%
- Add to model PDF
- Additional tests run where accidentals rate was floated

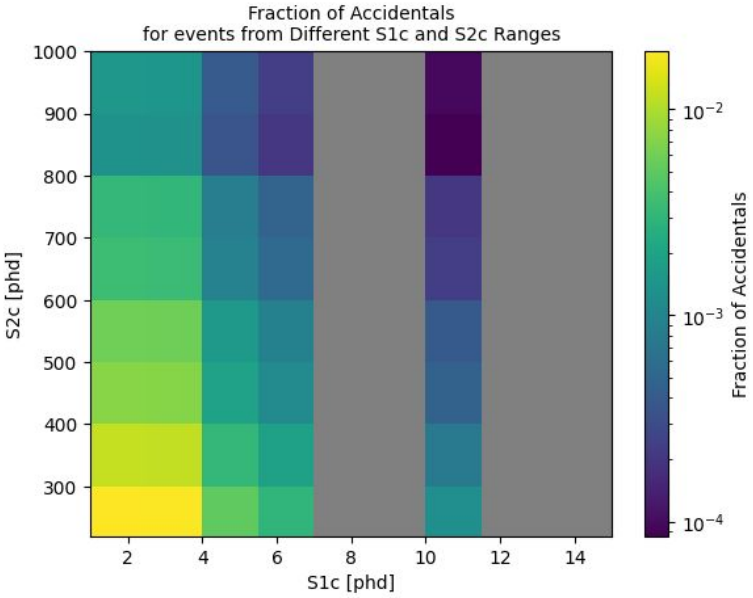

Systematics recap

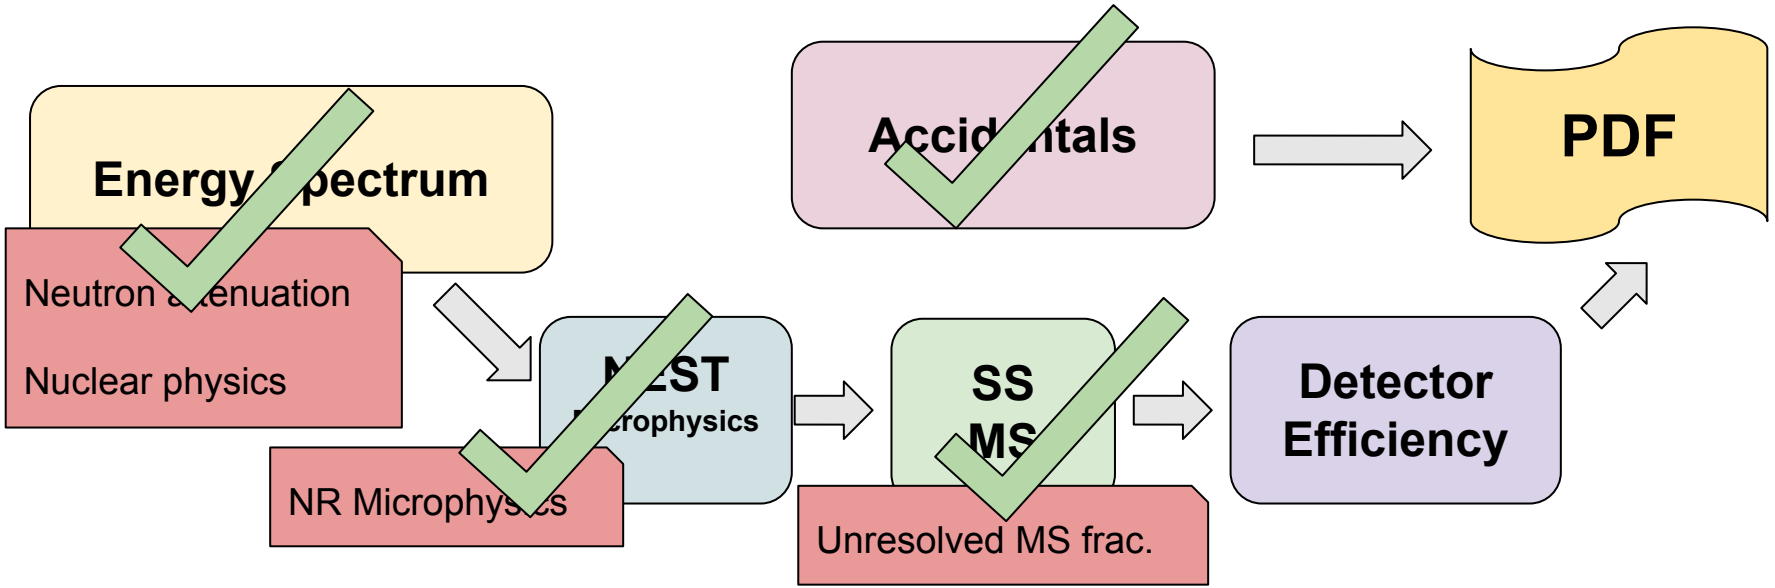

We have a great model. With 15 parameters…

In the past, could just grid-search

- Likelihood has some statistical error:
  - Breaks most “smart” optimizers
  - Really messes up Hessian - uncertainty estimate is hard
- Need to propagate systematics

- Option 1: make likelihood exact*
- Option 2: Markov Chain Monte Carlo

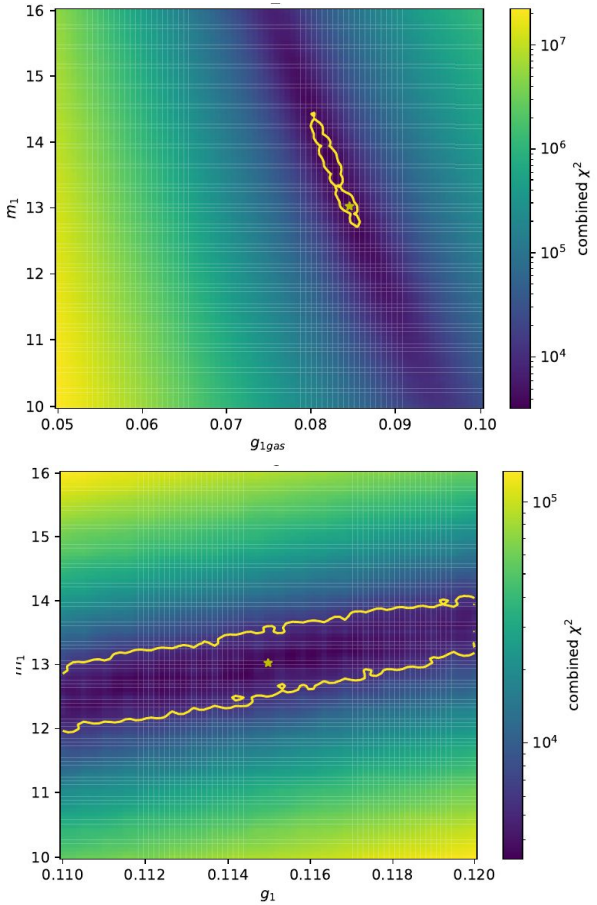

* Sounds crazy, but it’s being worked on: FlameNEST

Markov Chain Monte Carlo

Take a small step from each point (sampling algorithm)

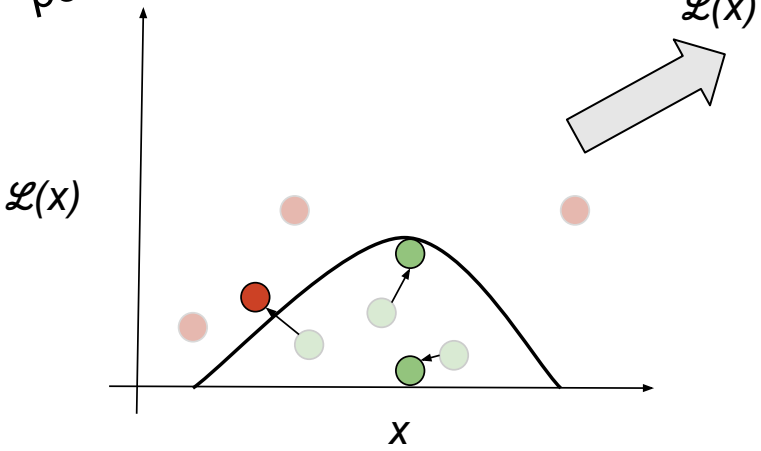

Start near a "good guess" for max $\mathcal{L}$

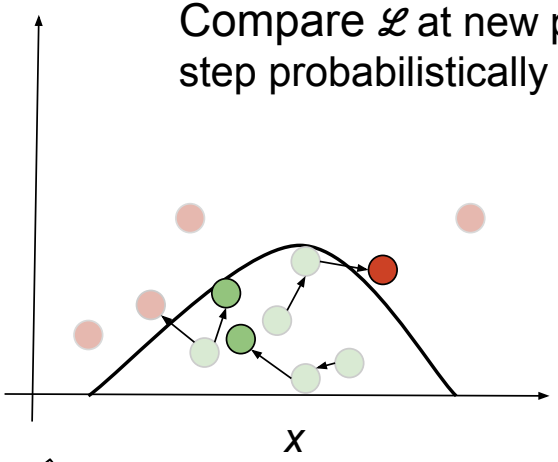

Compare $\mathcal{L}$ at new point. "Accept" the step probabilistically $\propto$ ratio of $\mathcal{L}$

Repeat, saving each step in a chain

$$\begin{pmatrix} x_0 & x_1 & x_1 & x_2 & x_3 & x_3 & x_4 & x_4 \\ x_0 & x_0 & x_1 & x_2 & x_2 & x_3 & x_4 & x_5 \\ x_0 & x_1 & x_2 & x_2 & x_2 & x_4 & x_5 & x_5 \end{pmatrix} \propto \mathcal{L}(x)$$

What it looks in practice

- Collections of “chains” of samples for each parameter
- subsequent steps not independent - quantify ACR & “thin” the chains
- Compute stat. quantities from samples, project them in various dimensions…

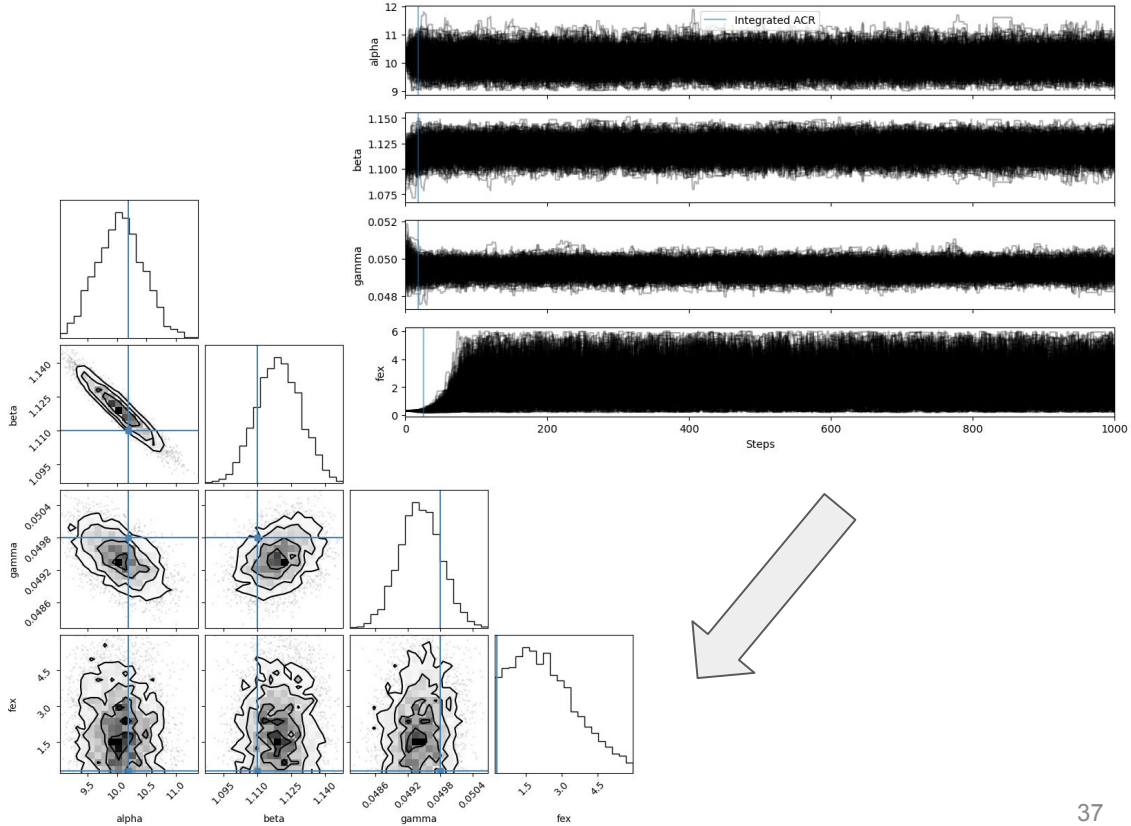

Here’s the model!!!

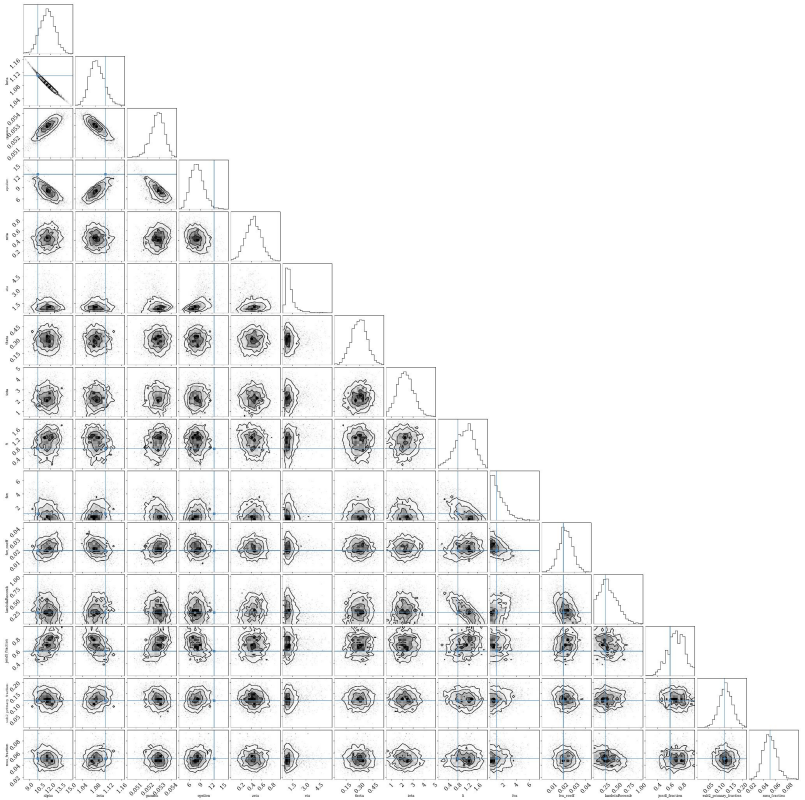

Summary statistics

- Compute summary stats of chains
- Color-code matrix of results based on correlation
- …Still useless for stats

High-E yields

Low-E yields

Fluctuations

Systematics

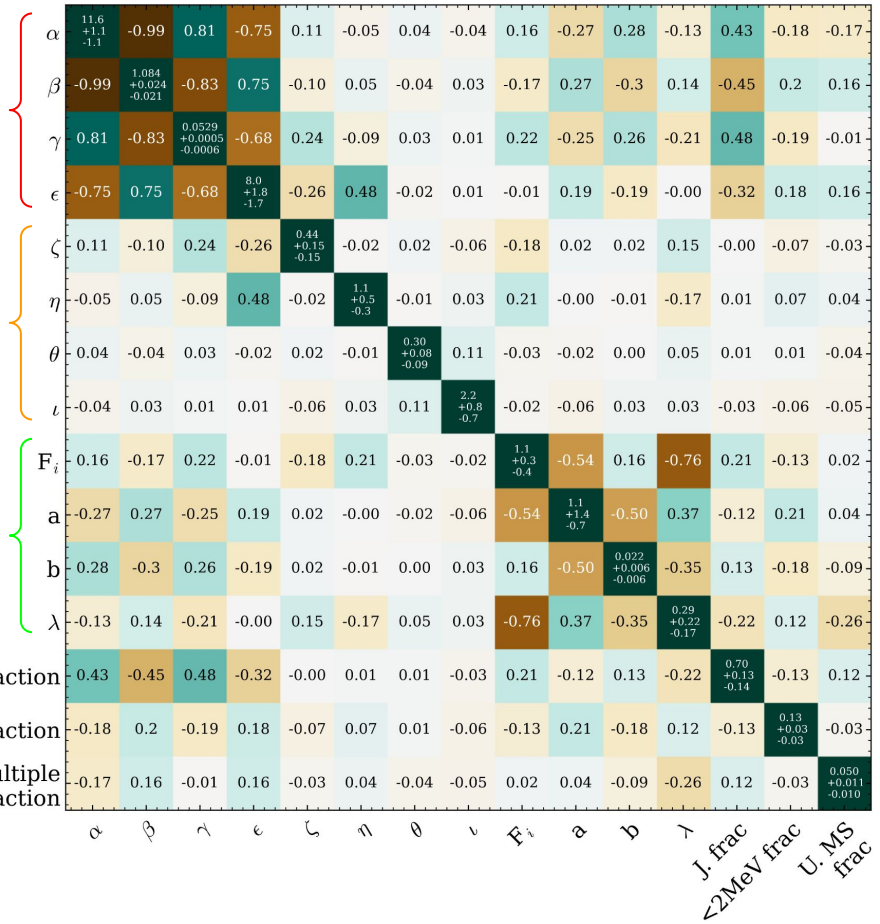

What does stats want?

Smallest set of parameters

- Template-based inference: need N-dimensional grid of templates
- Interpolation between parameter combinations is complex

Parameter priors should be Gaussian

- …Or at least have a nice analytic formula
- Need to include constraints in likelihood
- Need to sample priors for toy generation

Parameter priors should be uncorrelated

- Correlated priors harder to sample
- Stats software doesn't handle dependencies between parameters nicely

Translate to yield curves

Each line is the yield curve from a random sample

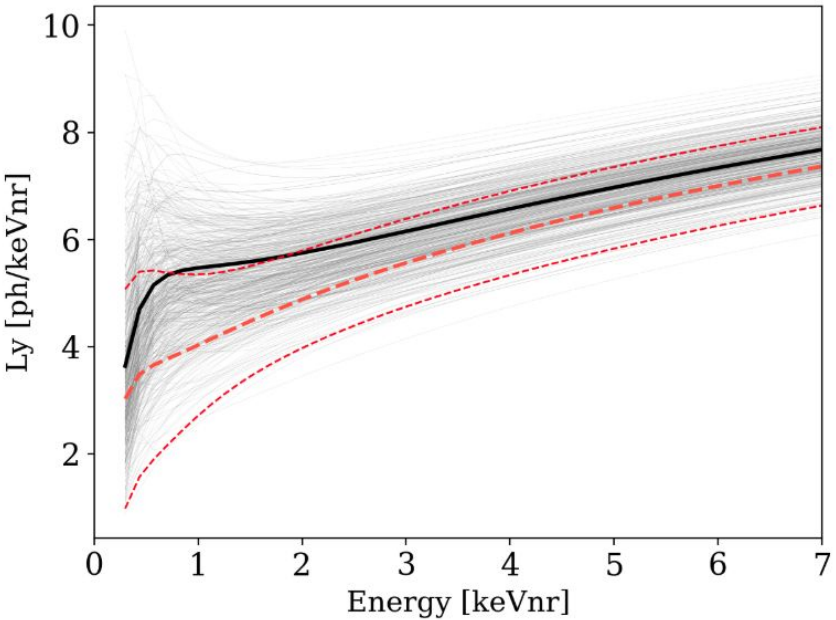

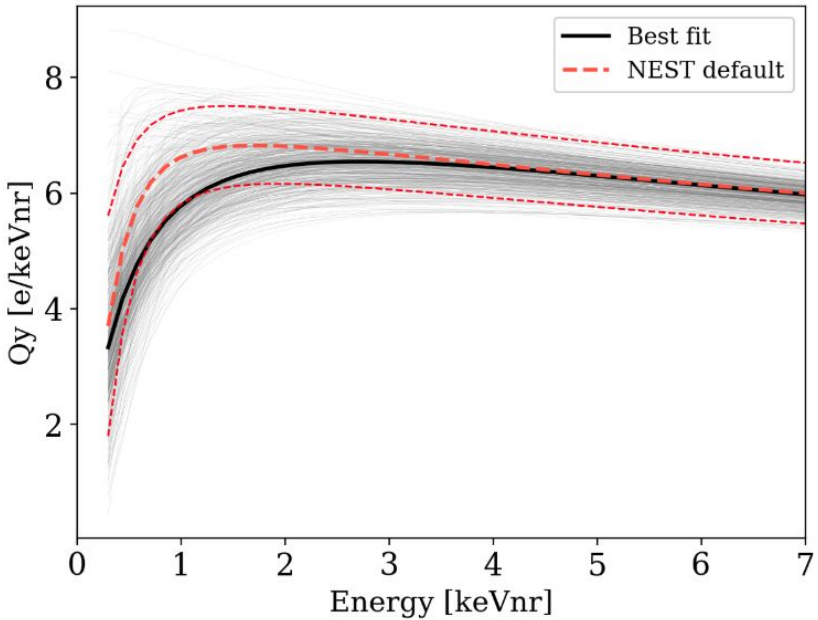

Eigendecomposition

With just 2 gaussian parameters, approximate all the model uncertainties in our yields

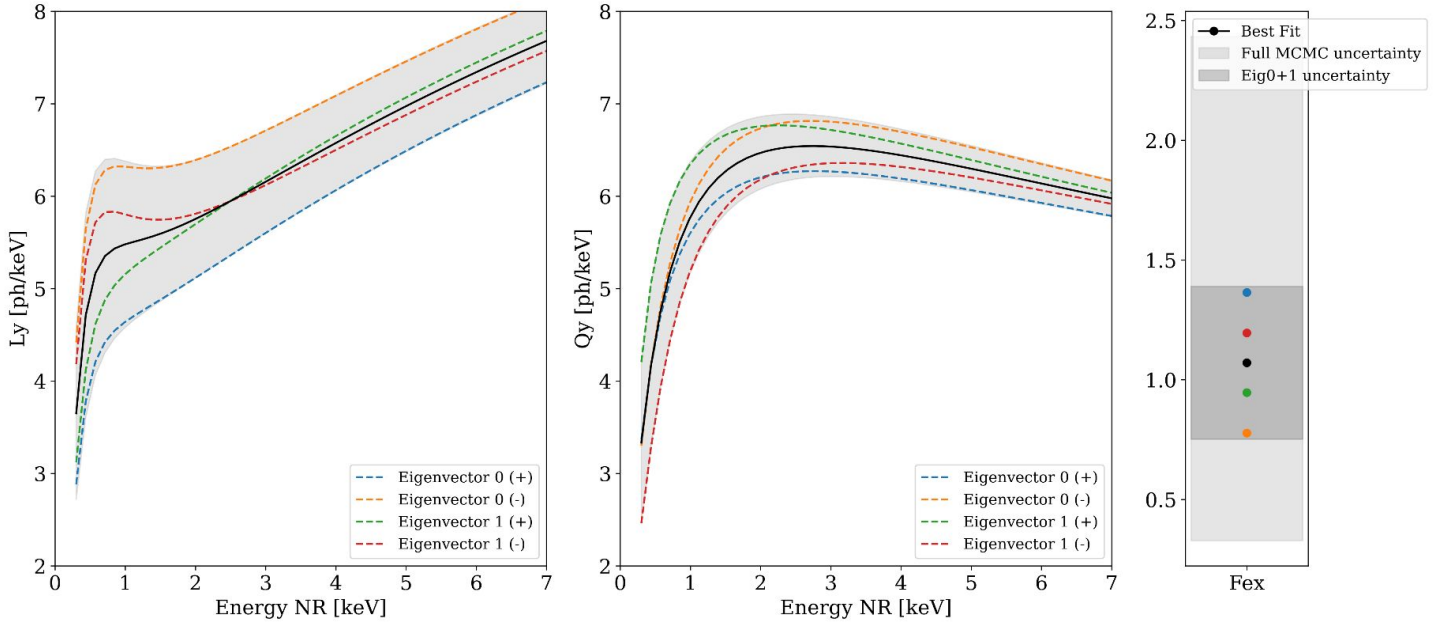

WS2025 LMW results

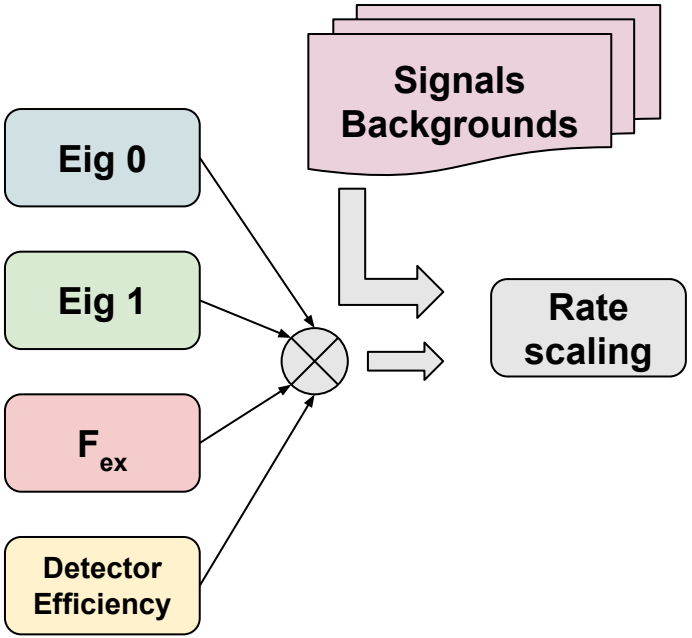

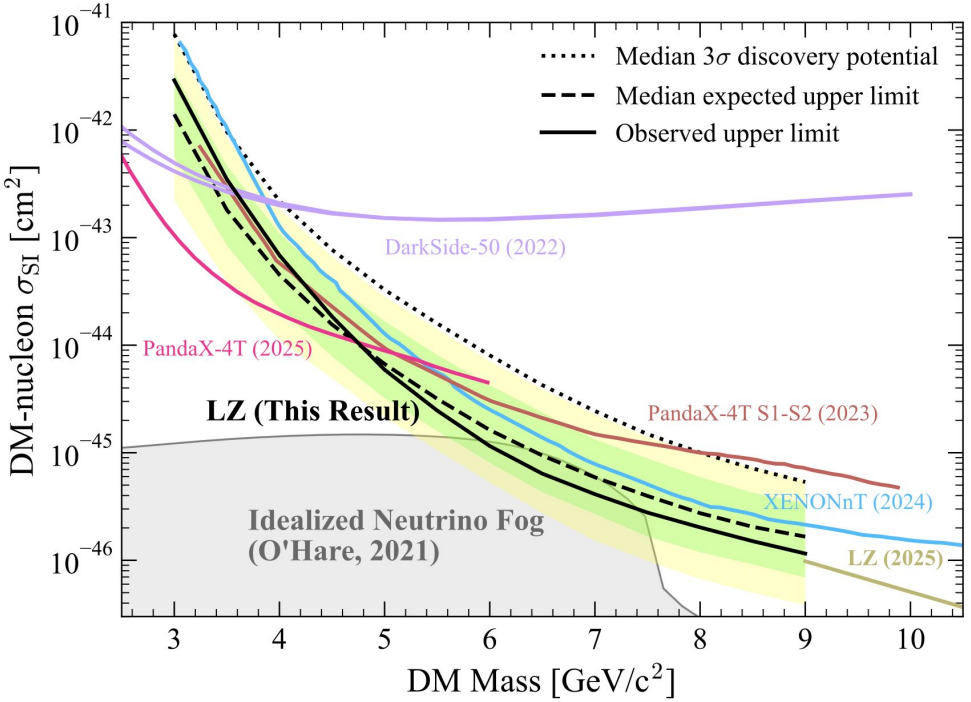

WS2025 ^{8}B evidence

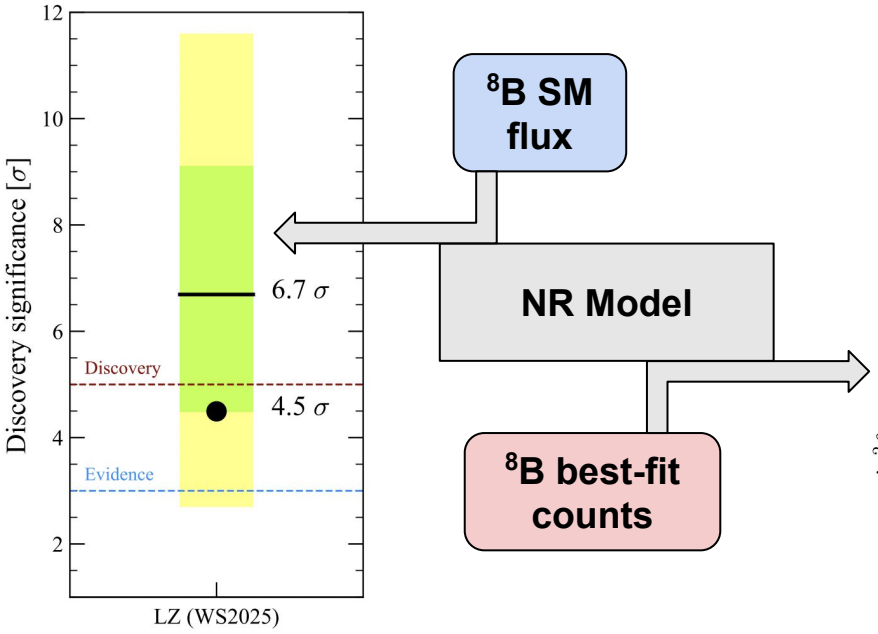

^{8}B SM flux

NR Model

^{8}B best-fit counts

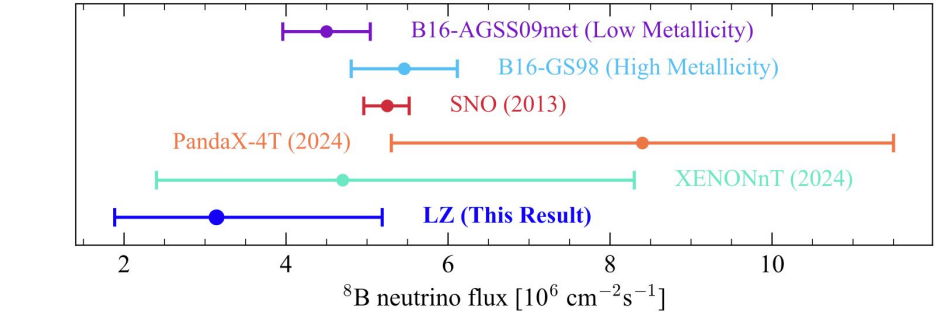

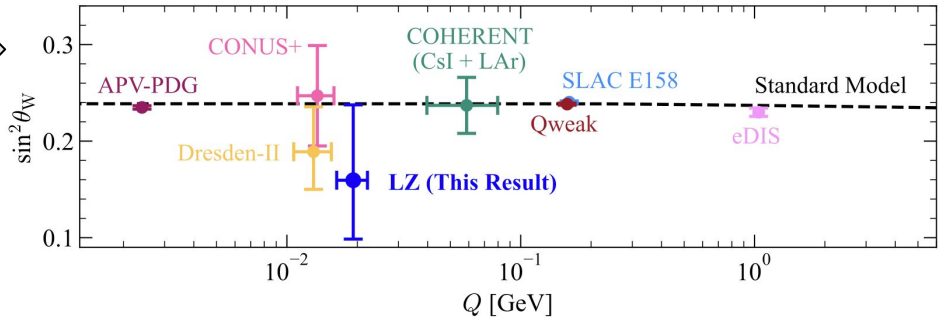

The end of the LMW journey

Completely new environment, requirements, systematics…

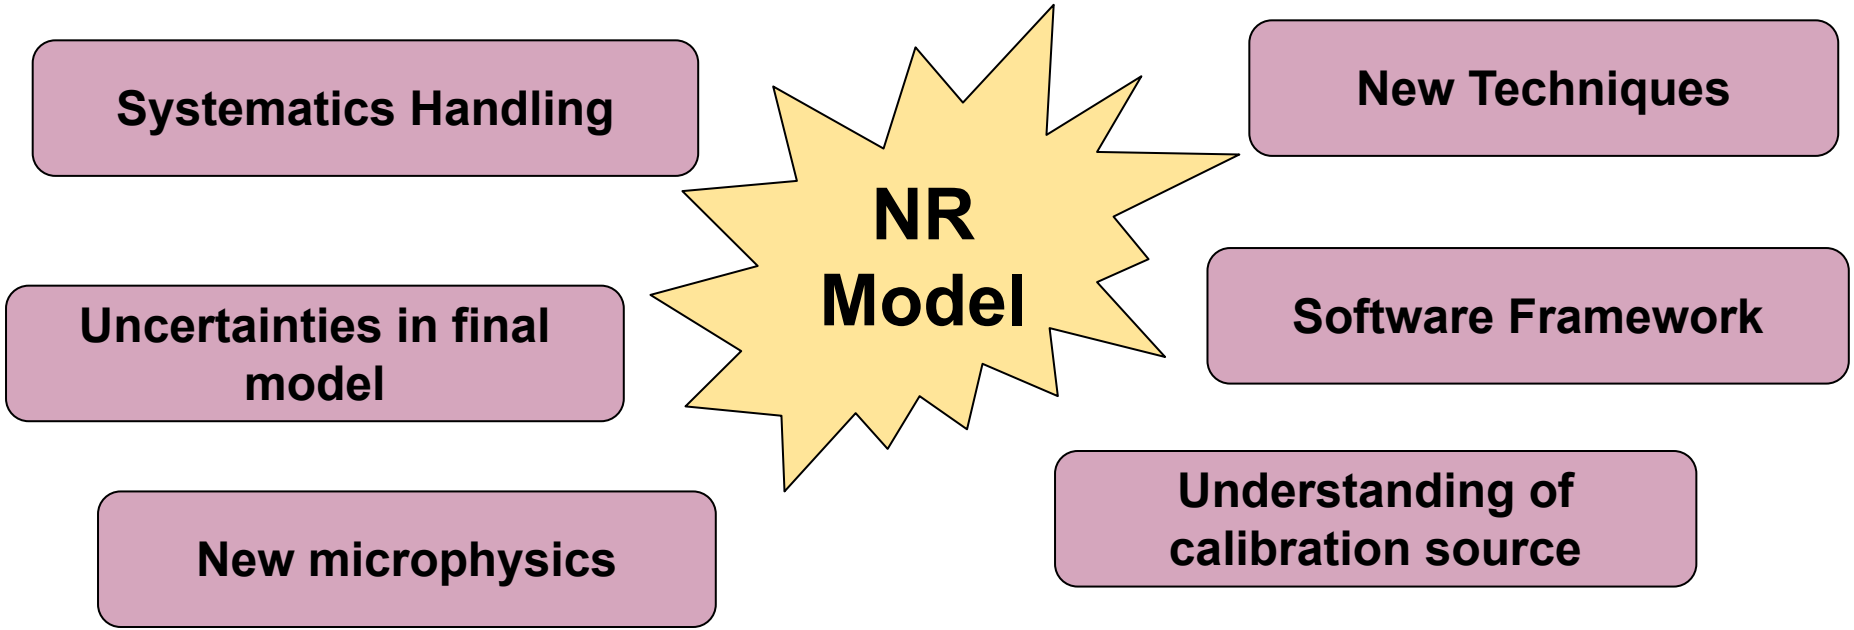

… Or is it just the first step?

We've demonstrated a whole set of tools & techniques to go beyond the old paradigms. Now let's use them.

- Many ways - proven and unproven - to push energy threshold lower
- Plenty of interesting datasets left: MS neutrons, reflector neutrons, AmBe…
- Fundamental questions to answer - why is Fex energy dependent? How can sub-binomial fluctuations happen?

Timeline to next-gen detectors might be long. Let's use this time wisely!

Thank you!

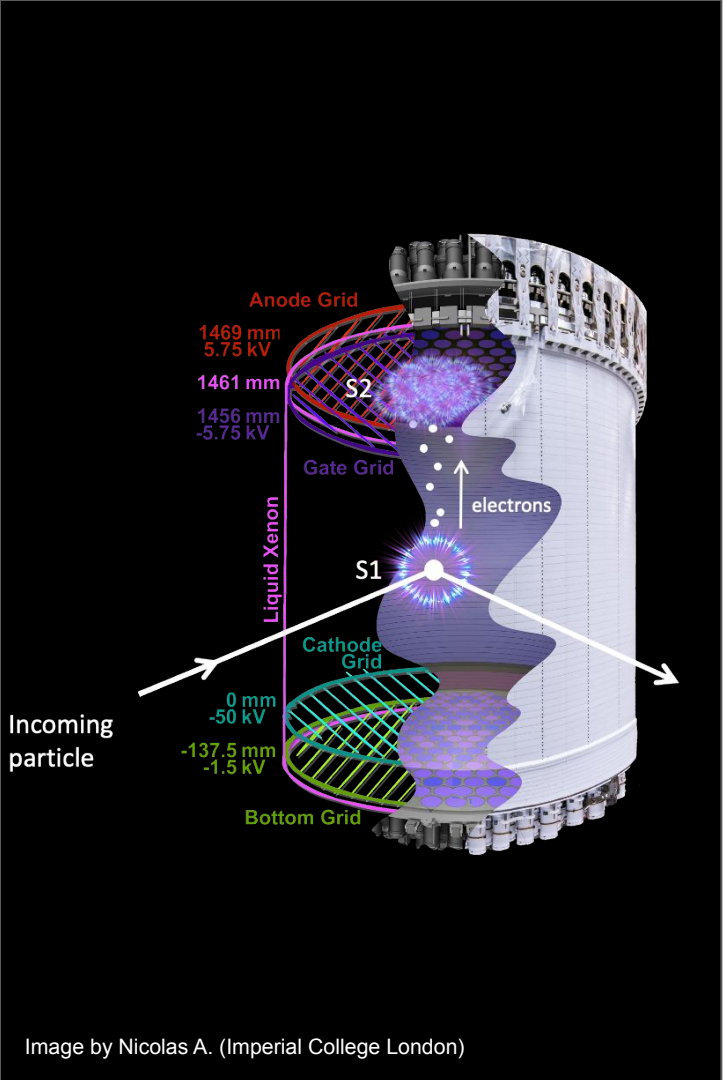

Image by Nicolas A. (Imperial College London)

Thanks to our sponsors and 38 participating institutions!

U.S. Department of Energy
Office of Science

Science and Technology Facilities Council

SANFORD UNDERGROUND RESEARCH FACILITY

Fundação para a Ciência e a Tecnologia

Institute for Basic Science

Swiss National Science Foundation

ARC CENTRE OF EXCELLENCE FOR DARK MATTER PARTICLE PHYSICS

APPENDIX

Comparing to other results

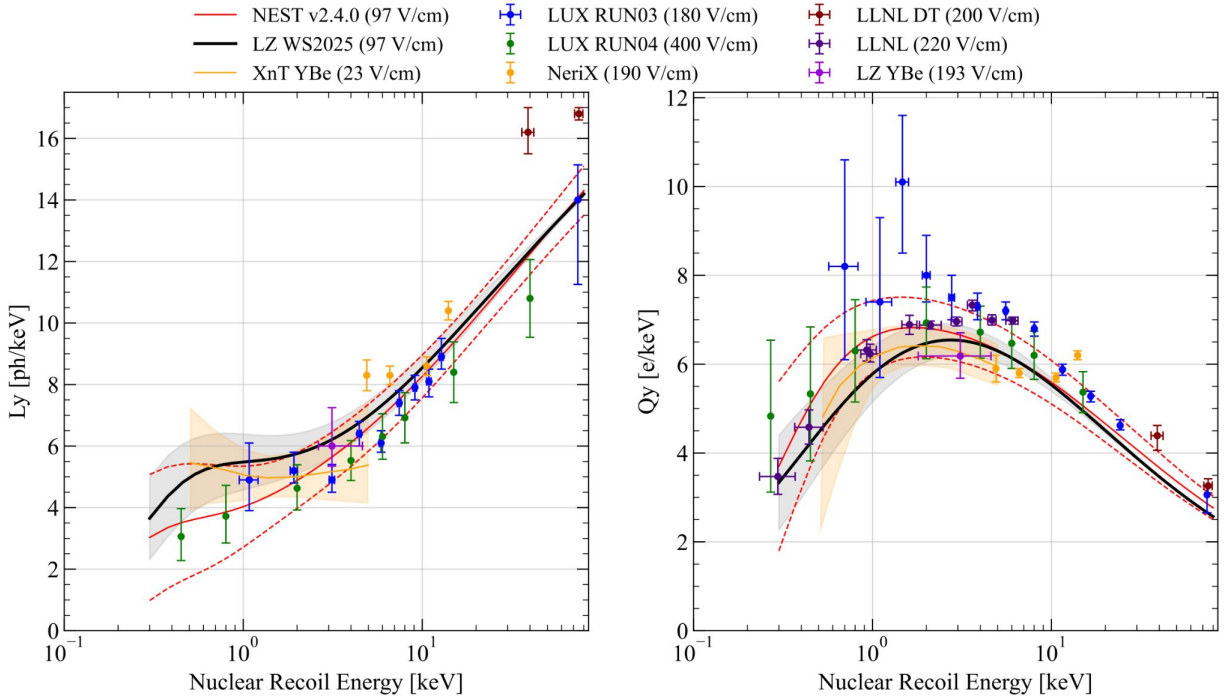

Eigendecomposition of yields - 1

$$C_i = \begin{pmatrix} L_{y0} & L_{y1} & L_{y2} & Q_{y0} & Q_{y1} & Q_{y2} \end{pmatrix}$$

Treat each L_y/Q_y bin as an independent RV

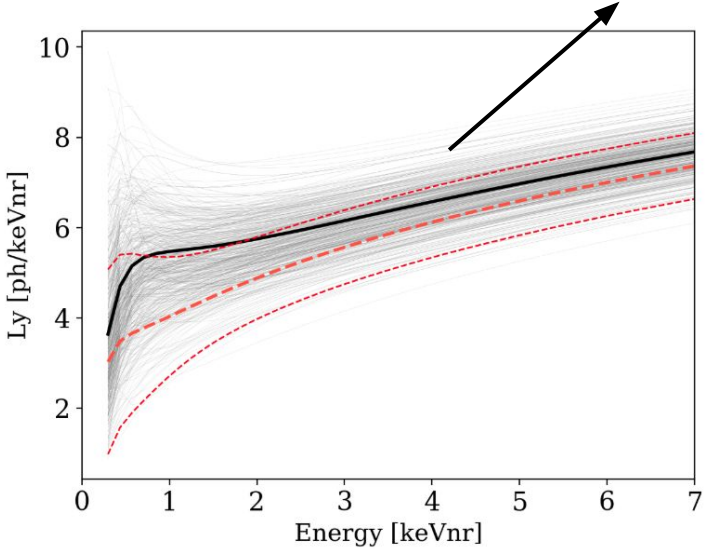

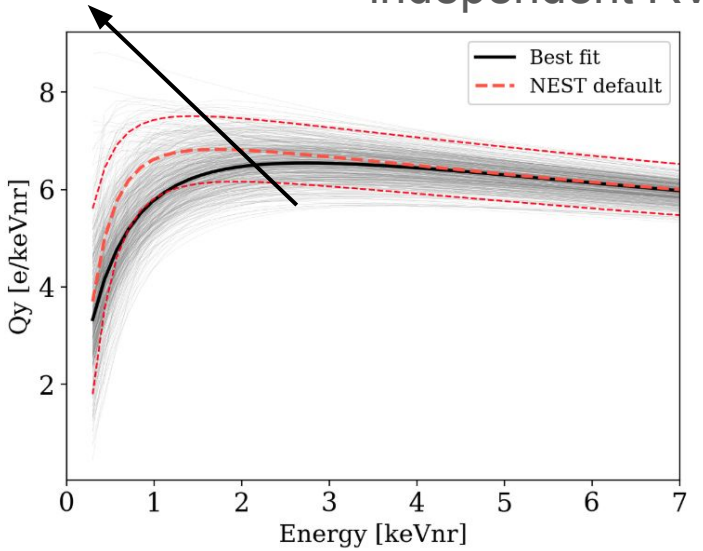

Eigendecomposition of yields - 2

$$C_0 = \begin{pmatrix} L_{y0} & L_{y1} & L_{y2} & Q_{y0} & Q_{y1} & Q_{y2} & F_{ex} \end{pmatrix}$$

$$C_1 = \begin{pmatrix} L_{y0} & L_{y1} & L_{y2} & Q_{y0} & Q_{y1} & Q_{y2} & F_{ex} \end{pmatrix}$$

$$C_2 = \begin{pmatrix} L_{y0} & L_{y1} & L_{y2} & Q_{y0} & Q_{y1} & Q_{y2} & F_{ex} \end{pmatrix}$$

$$\begin{pmatrix} \cdots \end{pmatrix}$$

Extract covariance

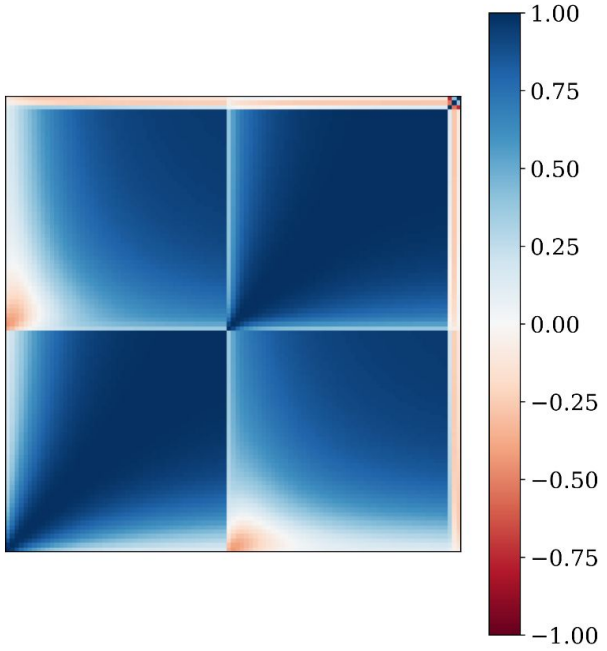

$$C_N = C_{mean} + \text{Multivar. Gauss}(C_{covariance})$$

Eigendecomposition of yields - 3

$C_{cov} = VDV^{T}$

$D = \mathrm{diag}(\lambda_i)$

$V =$

$$\begin{pmatrix} L_{y0\ \text{eigen}\ 0}, L_{y1\ \text{eigen}\ 0} \ldots, Q_{y3\ \text{eigen}\ 0} \\ L_{y0\ \text{eigen}\ 1}, L_{y1\ \text{eigen}\ 1} \ldots, Q_{y3\ \text{eigen}\ 1} \\ \ldots \end{pmatrix}$$

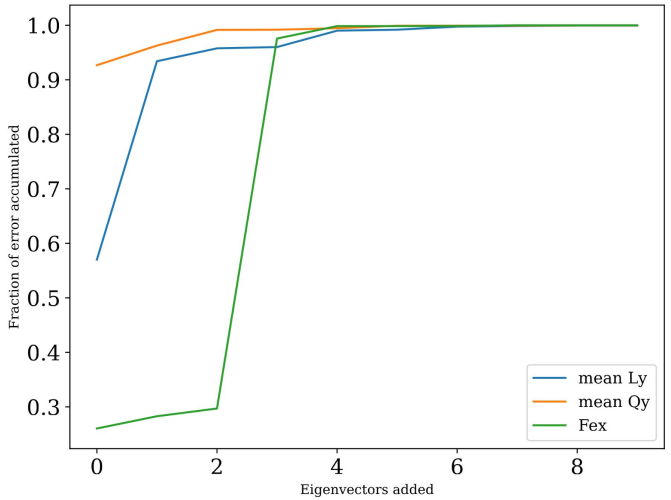

$C_N = C_{mean}$ + Multivar. Gauss(**1**) x V x sqrt(λ)

DD_MCMC

- 15 parameters in model, 2M steps, 50+ config settings, ~5 collaborators working on fits
- Needed standardisation + version control
- Custom python tool developed for this analysis
  - All configuration done in VC'd TOML files
  - Standardised QA tools run after each fit
  - Custom version of NESTpy with vectorisation support

```
[parameters.eta.gaussian]
mean = 2.0
sigma = 1.0

[parameters.theta]
[parameters.theta.gaussian]
mean = 0.3
sigma = 0.05

[parameters.iota]
[parameters.iota.gaussian]
mean = 2.0
sigma = 0.5

[parameters.fi]
[parameters.fi.flat]
start = 0.001
stop = 2.0
best = 0.8
```

Expected DD counts vs energy

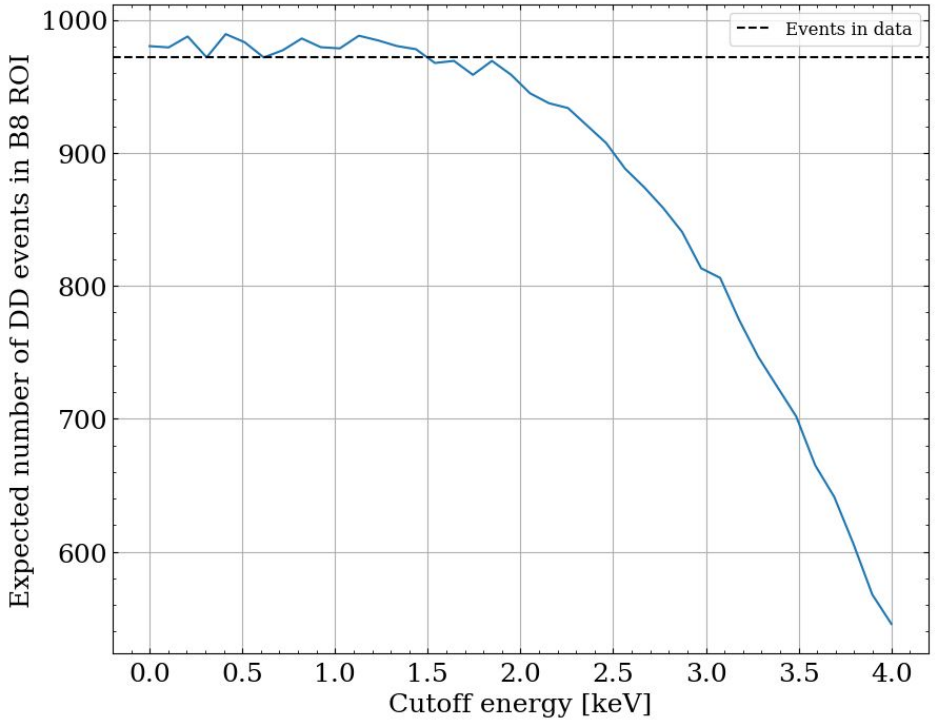

Polishing the gold - data quality cuts

Not a neutron?

- S1 timing selection
- Reconstructed energy selection
- Within ROI
- Conduit cut
- Drift time cut
- Rain cut

ACCIDENTALS!

Not in drift region?

- S2 width cut
- Drift time cut

Attenuated neutron?

- Conduit cut
- Drift time cut

Not a single scatter?

- S2 width cut
- Position reconstruction χ^2

PPP GoF

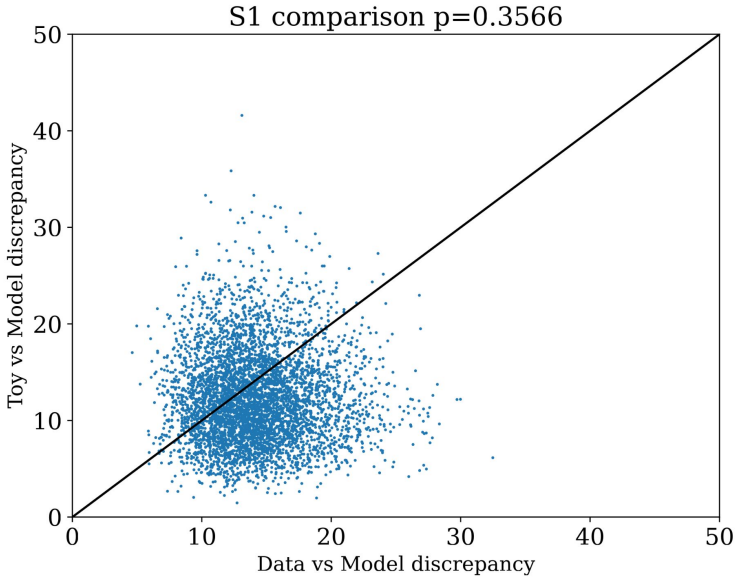
S1 comparison p=0.3566
Toy vs Model discrepancy
Data vs Model discrepancy
0
10
20
30
40
50

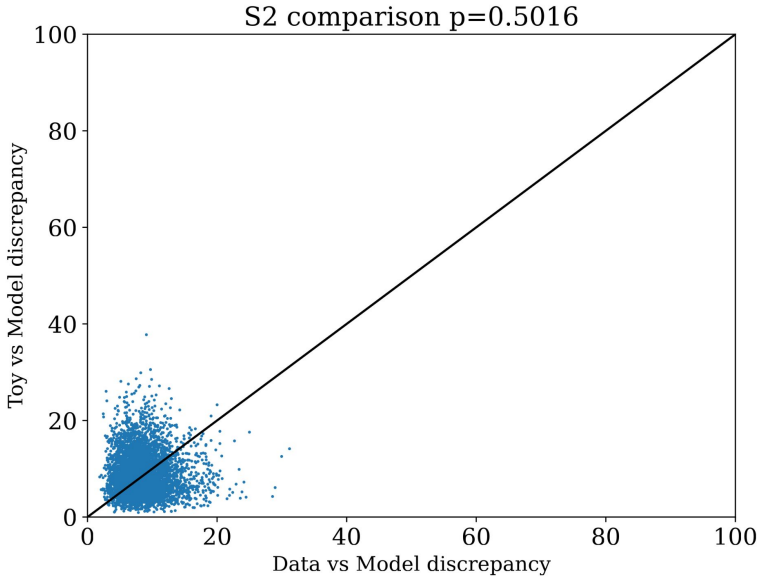
S2 comparison p=0.5016
Toy vs Model discrepancy
Data vs Model discrepancy
0
20
40
60
80
100

PPP GoF

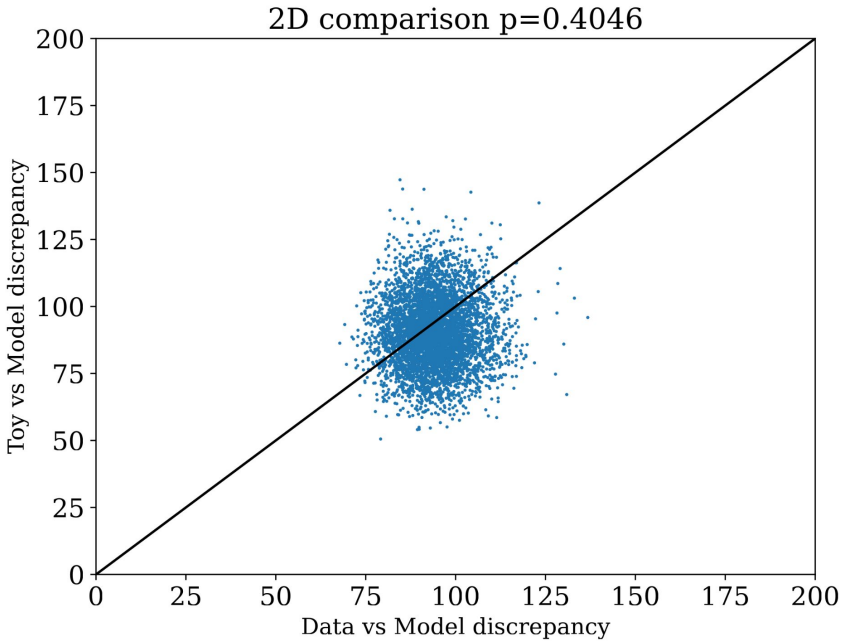
2D comparison p=0.4046
Toy vs Model discrepancy
Data vs Model discrepancy
0
25
50
75
100
125
150
175
200

Energy dists in S1/S2

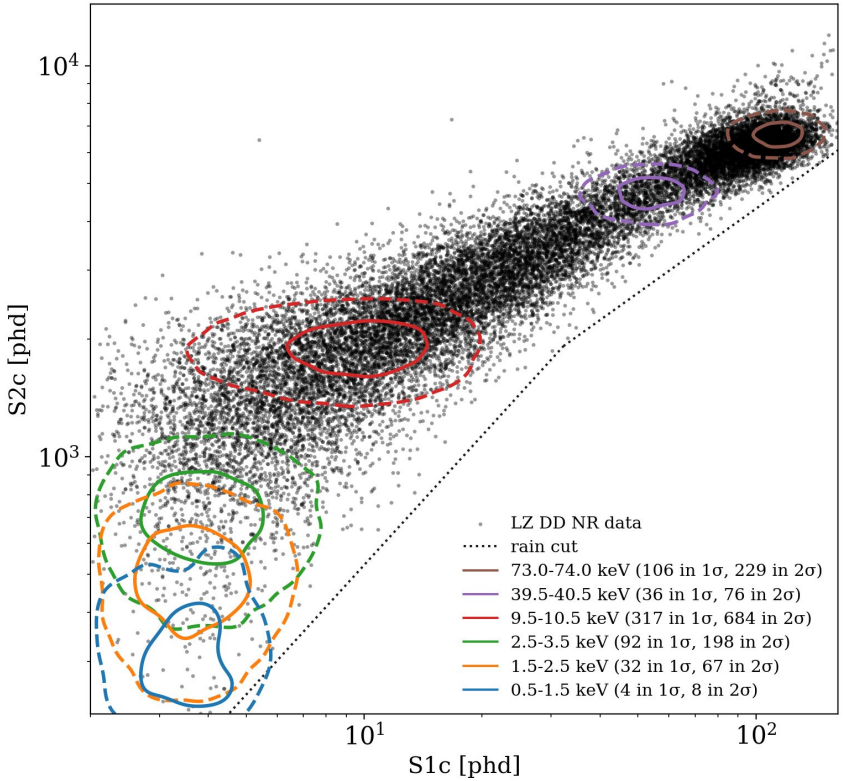

GoF projections

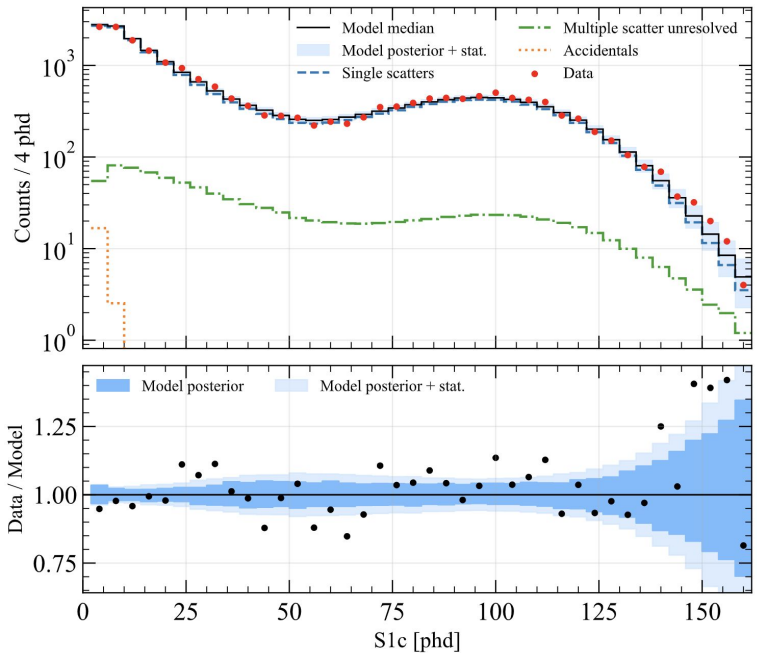

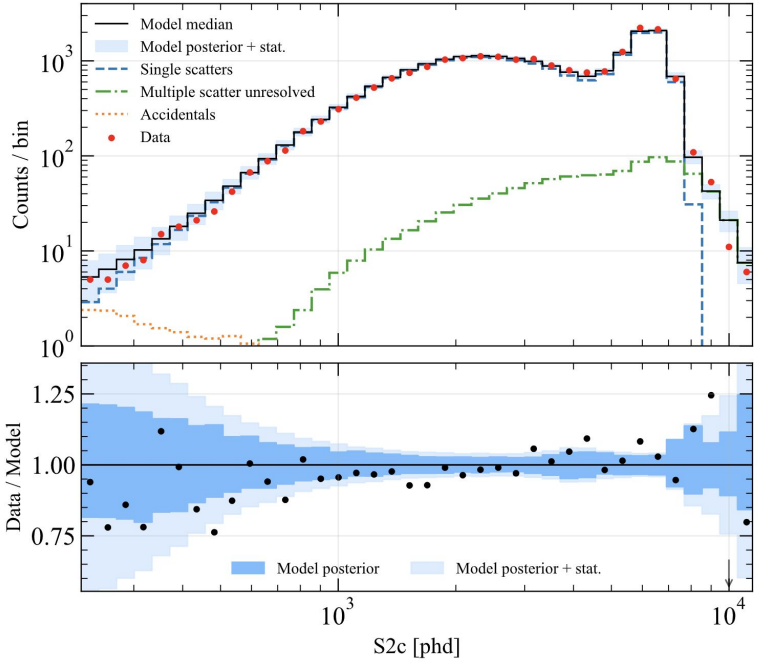

2D GoF

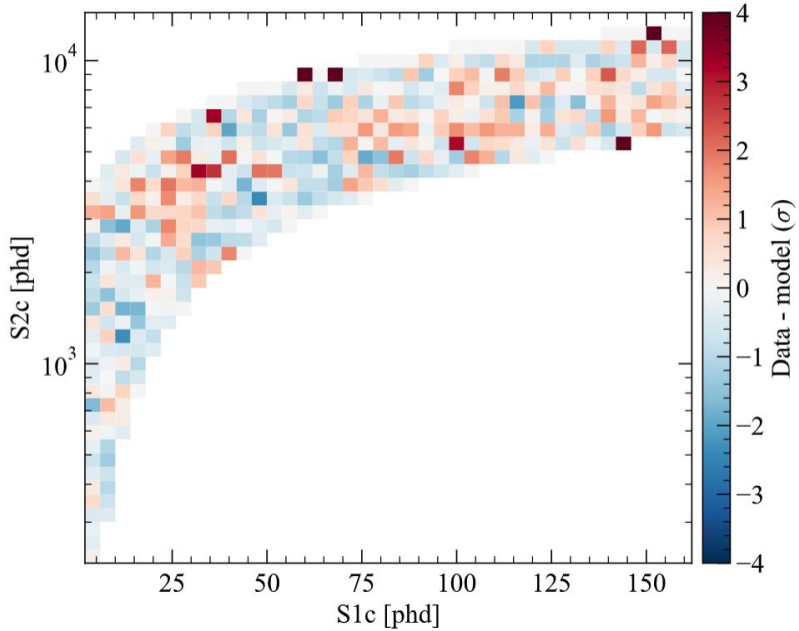
10^4
10^3
S2c [phd]
25
50
75
100
125
150
S1c [phd]
4
3
2
1
0
−1
−2
−3
−4
Data - model (σ)

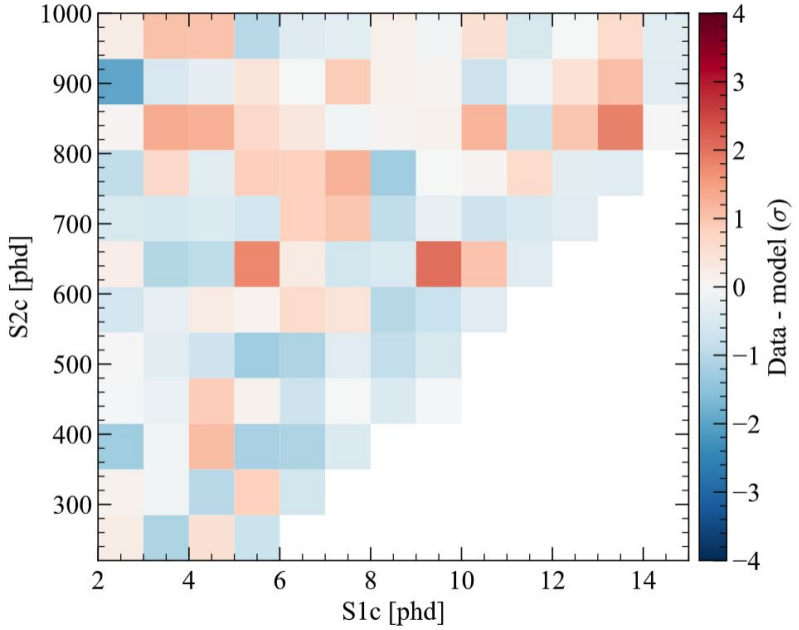
1000
900
800
700
600
500
400
300
S2c [phd]
2
4
6
8
10
12
14
S1c [phd]
4
3
2
1
0
−1
−2
−3
−4
Data - model (σ)