Longitudinal Λ Polarization in Tagged $Z \to s\bar{s}$ Events with the CEPC AGORA Detector

Full-simulation study toward strange-quark spin transfer

Shuo Han Minqia Li

Institute of High Energy Physics, CAS

20 July 2026

From inclusive LEP polarization to $s \to \Lambda$ spin transfer

Longitudinal Λ polarization is measured from the parity-violating decay angle:

$$\frac{1}{\Gamma}\frac{d\Gamma}{dc} = \frac{1}{2}\left(1 + \alpha_\Lambda P_\Lambda c\right), \qquad c \equiv \cos\theta^*.$$

c is the proton angle in the Λ rest frame relative to the reconstructed Λ flight direction; the $\bar{\Lambda}$ convention is charge-flipped.

ALEPH $z = p_\Lambda / p_{\mathrm{beam}} > 0.3$

$$P_\Lambda = -0.32 \pm 0.07$$

OPAL $x_E = 2E_\Lambda/\sqrt{s} > 0.3$

$$P_\Lambda = -0.329 \pm 0.055_{\mathrm{stat}} \pm 0.052_{\mathrm{syst}}$$

LEP inclusive hadronic Z data, high-momentum Λ

This study PrimeTag $Z \to s\bar{s}$ category, no high-x_E cut, full-simulation MC

Physics target: in tagged $Z \to s\bar{s}$ events, $D_{LL}^{s\to\Lambda}(z) \equiv P_\Lambda(z)/P_s^{\mathrm{SM}}$ quantifies how efficiently the primary s-quark polarization survives non-perturbative fragmentation into a Λ; $P_s^{\mathrm{SM}} \simeq -0.94$. This study establishes detector-level P_Λ extraction; s-jet association, z binning and prompt/feed-down corrections are still needed before extracting D_{LL}.

ALEPH, Phys. Lett. B 374 (1996) 319; OPAL, Eur. Phys. J. C 2 (1998) 49.

Three independent samples define observation, calibration and continuum

Sample	Events	Generation	Statistical role
Polarized $Z \to s\bar{s}$	200 000	WHIZARD 3.1.7 + PYTHIA6; StdHep decay-angle resampling	Simulated observation
Unpolarized $Z \to s\bar{s}$	198 000	Same generator chain; no resampling	Response and signal-mass template
Five-flavour $Z \to q\bar{q}$	250 000	WHIZARD + PYTHIA6; 50k per flavour	Tag migration and continuum

WHIZARD + PYTHIA6 → StdHep $\rightarrow$ StdHep $\Lambda \to p\pi$ decay-angle resampling* $\xrightarrow{\text{LCIO}}$ CEPCSW + Geant4 full simulation + reco $\rightarrow$ `CEPCAnalysisNtuple` jet tagging + V^0 variable construction $\rightarrow$ joint mass–angle likelihood

*Polarized sample only: decay directions are resampled with $1 + \alpha_\Lambda P_\Lambda \cos\theta^*$. The closure-test input $P_\Lambda = -0.7952$ uses $P_s^{\mathrm{SM}} \simeq -0.935$ times an SU(6)-motivated benchmark transfer $D_{LL} = 0.85$; it is not an SM prediction for P_Λ. Here $\alpha_\Lambda = 0.732$; vertices and parent–daughter links are preserved.

PrimeTag ZJet M10 has sizeable signed light-flavour migration

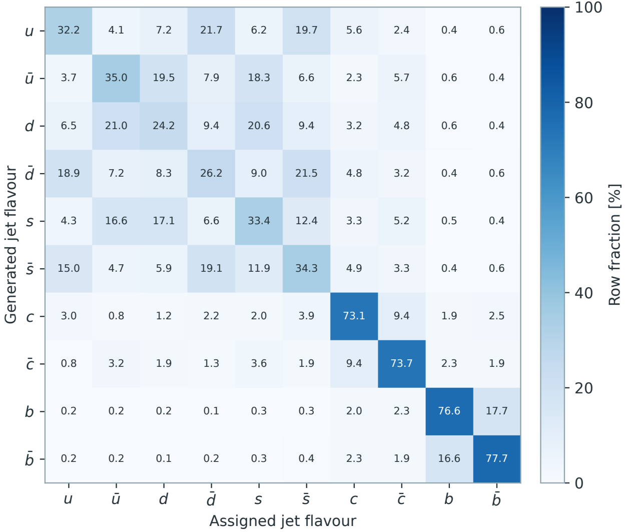

Generated \ Assigned	u	$\bar{u}$	d	$\bar{d}$	s	$\bar{s}$	c	$\bar{c}$	b	$\bar{b}$
u	32.2	4.1	7.2	21.7	6.2	19.7	5.6	2.4	0.4	0.6
$\bar{u}$	3.7	35.0	19.5	7.9	18.3	6.6	2.3	5.7	0.6	0.4
d	6.5	21.0	24.2	9.4	20.6	9.4	3.2	4.8	0.6	0.4
$\bar{d}$	18.9	7.2	8.3	26.2	9.0	21.5	4.8	3.2	0.4	0.6
s	4.3	16.6	17.1	6.6	33.4	12.4	3.3	5.2	0.5	0.4
$\bar{s}$	15.0	4.7	5.9	19.1	11.9	34.3	4.9	3.3	0.4	0.6
c	3.0	0.8	1.2	2.2	2.0	3.9	73.1	9.4	1.9	2.5
$\bar{c}$	0.8	3.2	1.9	1.3	3.6	1.9	9.4	73.7	2.3	1.9
b	0.2	0.2	0.2	0.1	0.3	0.3	2.0	2.3	76.6	17.7
$\bar{b}$	0.2	0.2	0.1	0.2	0.3	0.4	2.3	1.9	16.6	77.7

Model
ZJet M10
10 signed-flavour outputs

Working point
raw per-jet argmax
one jet tagged s and the other $\bar{s}$
no score threshold

Scope
250k full-simulation events
499,805 valid jets
rows normalized to 100%

Signed light-flavour performance is below the CEPC TDR benchmark; the cause is under investigation.

250k events; row-normalized by generated signed flavour; raw argmax. The TDR comparison is indicative because the configurations differ.

The double-strange category is 14.556% efficient and 60.56% pure

Generated process	N_{gen}	$N_{s\bar{s}\text{-tag}}$	ϵ_f [%]	Selected composition [%]
$Z \to s\bar{s}$	50 000	7 278	14.556 ± 0.158	**60.56**
$Z \to d\bar{d}$	50 000	2 895	5.790 ± 0.104	24.09
$Z \to u\bar{u}$	50 000	2 241	4.482 ± 0.093	14.46
$Z \to c\bar{c}$	50 000	135	0.270 ± 0.023	0.87
$Z \to b\bar{b}$	50 000	2	0.004 ± 0.003	0.02

Efficiency is selected/generated for each flavour; selected composition uses Standard Model Z partial-width weights. The leading backgrounds are $d\bar{d}$ and $u\bar{u}$. Uncertainties are binomial sample-statistical errors.

event level These tag metrics do not share the candidate-level denominator used in the V^0 cutflow.

Two-helix V^0 reconstruction provides a secondary-vertex benchmark

charged particle-flow object
first-hit track state preferred $\rightarrow$ opposite-charge helices
3D DOCA $\leq$ 5 mm $\rightarrow$ secondary vertex = midpoint
of closest helix points $\rightarrow$ propagate track momenta to vertex
store displacement and masses

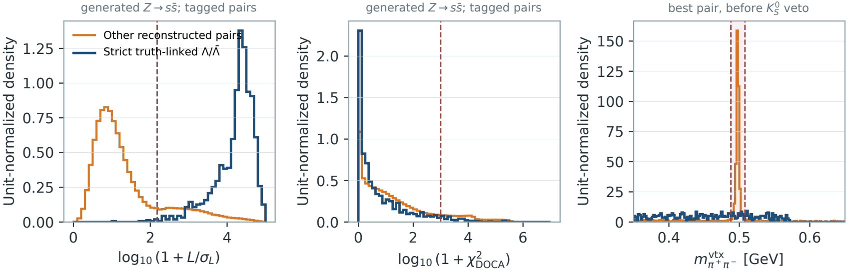

χ^2_{DOCA} = squared helix separation at closest approach, divided by the summed D_0 and Z_0 variances: a separation significance, not a vertex-fit χ^2/ndf.

K^0_S veto: $0.488 < m^{\mathrm{vtx}}_{\pi\pi} < 0.508$ GeV under the pion-mass hypothesis. A clear peak near $m_{K^0_S}$ (highest 2-MeV bin: 0.496–0.498 GeV) shows that the veto variable is resolved. Window: 2 542 other / 75 strict $\Lambda/\bar{\Lambda}$ pairs.

50k generated $Z \to s\bar{s}$ events; best-pair $m^{\mathrm{vtx}}_{\pi\pi}$ before the veto. `CEPCAnalysisNtuple`: opposite-charge pairs with 3D DOCA $\leq$ 5 mm; truth labels are diagnostic only.

The detailed V^0 cutflow exposes both retention and purity

Stage	All	Strict truth	Pair retention [%]	Purity [%]	Counting unit
All producer V0 candidates	4 309 926	7 653	100.00	0.18	track pairs
$L/\sigma_L > 150$	599 267	7 587	99.14	1.27	track pairs
$\chi^2_{\mathrm{DOCA}} < 1000$	211 488	7 151	93.44	3.38	track pairs
Double-strange event tag	33 214	1 506	19.68	4.53	track pairs
Minimum-χ^2_{DOCA} pair	6 287	751	9.81	11.95	track pairs
K^0_S veto	3 670	676	8.83	18.42	track pairs
Create $p\pi^-$ and $\bar{p}\pi^+$ hypotheses	7 340	676	–	–	mass hypotheses
$1.10 < m(p\pi) < 1.13$ GeV	942	674	–	71.55	mass hypotheses

Pair-level retention is relative to the 7 653 strict truth-linked pairs in the first row; it is not a generated-Λ efficiency.

Creating two charge hypotheses is a unit change, not a cut; pair and hypothesis retentions are therefore not compared. The mass range keeps 674/676 strict truth hypotheses.

50 000 generated $Z \to s\bar{s}$ events; truth information is used only for this diagnostic table.

The nominal mass–angle fit extracts about 2 666 signal candidates

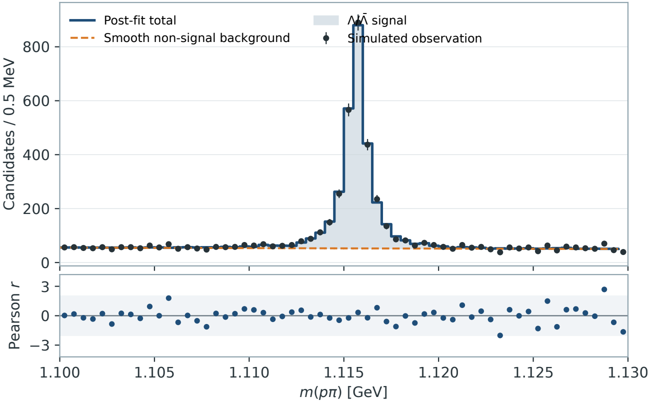

nominal model under validation

Fit all 10 angle bins × 60 mass bins together:

$$\nu_{ij} = S\,A_i(1 + \alpha_\Lambda P_\Lambda \kappa_i)M_j + B_i Q_{ij}.$$

S: signal normalization; A_i, $kappa_i$: $P = 0$ acceptance and angular response.
M_j: empirical $P = 0$ signal template.
$Q_{ij} \propto \exp(c_i x_j)$; M_j, Q_{ij} are mass-normalized.
B_i, c_i: free background yield and slope per angle bin.

Observed	5 821
Nominal post-fit signal	2 666
Post-fit background	3 155

In 1.110–1.120 GeV the nominal post-fit signal fraction is 71.14%; this is not a selection purity.

Mass-integrated projection of one joint fit, not a mass-only fit. The pseudo-observation includes a Poisson draw from a smoothed $Zq\bar{q}$ MC expectation; it is absorbed by the profiled smooth background. The two charge hypotheses from one track pair are treated as independent Poisson entries; their correlation is not yet calibrated.

All ten angle-bin mass fits come from the same joint likelihood

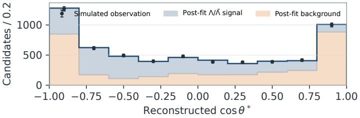

The mass-integrated sample is divided into ten reconstructed $\cos\theta^*$ bins. Each panel below is the mass projection for one bin; all ten enter the same joint likelihood.

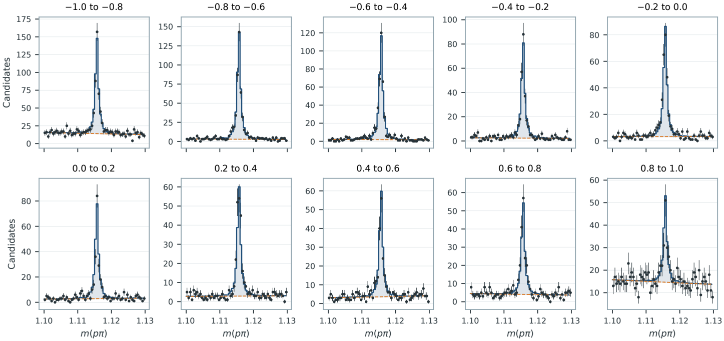

Fit status: converged; black is observation, blue total and orange non-signal.

Validation open: common mass template $p = 0.0040$; within-pair correlation not calibrated.

The post-fit signal yield carries the polarization dependence

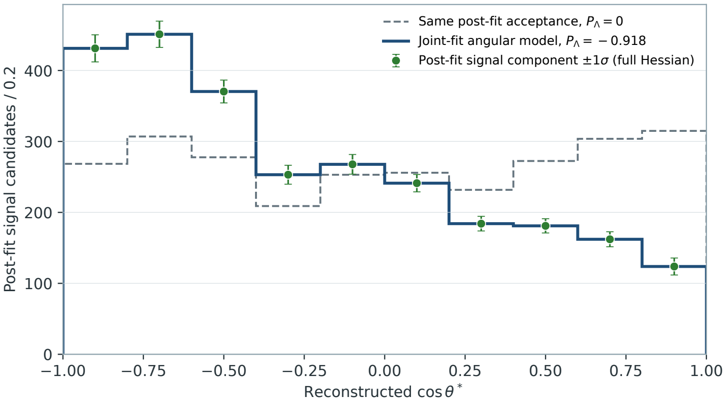

The signal yield in reconstructed angle bin i is

$$Y_i^{\mathrm{sig}} = S\,A_i\big(1 + \alpha_\Lambda P_\Lambda \kappa_i^{\mathrm{eff}}\big).$$

Green markers show the post-fit signal component; bars are correlated marginal 1σ uncertainties from the full fit Hessian. The blue line is the angular model at the fitted P_Λ.

Joint-fit result

$$P_\Lambda = -0.918,$$
$$\alpha_\Lambda P_\Lambda = -0.672.$$

P_Λ is profiled in the full 10×60 likelihood; these are not independent per-bin fits.

A_i and $W_i = 1 + \alpha_\Lambda P_\Lambda \kappa_i^{\mathrm{eff}}$ are post-fit response terms. The full-fit covariance is propagated to Y_i^{sig}; these are not independent per-bin fits.

Current nominal MC diagnostic: $P_\Lambda = -0.918$

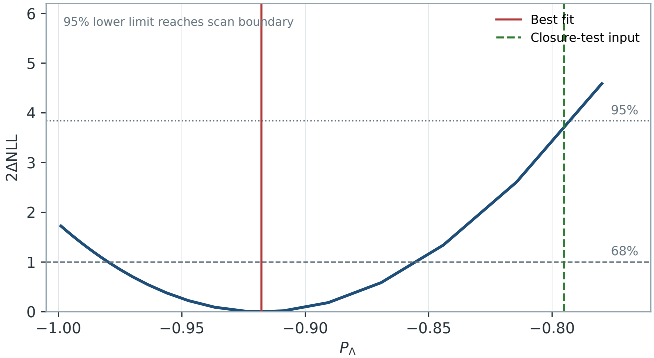

Conditional nominal fit

$P_\Lambda = -0.918$

68%: $[-0.980, -0.856]$

95%:

$[-0.999\ (\text{boundary}), -0.794]$

- closure input / generator-level check: -0.7952 / -0.7984 ± 0.0088;
- closure input: outside 68%, inside 95%; lower 95% endpoint at the boundary;
- the full detector-to-fit chain is not yet validated.

All nuisance parameters are re-minimized at fixed P_Λ. The interval is conditional on the nominal model; no lower 95% crossing is found before -0.999.

Λ reconstruction is demonstrated; polarization-model validation remains open

Established

- M10 tagging and the two-helix V^0 cutflow are quantified;
- a $\pi^+\pi^-$ peak near $m_{K^0_S}$ resolves the veto observable;
- the joint fit converges with about 2 666 post-fit signal candidates.

Before physics interpretation

- validate angle-dependent mass shapes and the joint pair-hypothesis model;
- validate selection/angular response and robust coverage;
- derive s-jet/feed-down corrections for $D_{LL}^{s\to\Lambda}(z)$.

Conditional nominal-model uncertainty: $\sigma(P_\Lambda) = 0.0628$; impacts are 0.037 from response/template MC and 0.018 from α_Λ (correlated). Other systematics are unestimated.

Possible sources of the closure-input–fit shift: angle-dependent mass shape ($p = 0.0040$), unmodeled correlation of the two hypotheses from one track pair, or incomplete selection/angular-response validation. These are hypotheses, not assigned uncertainties.

The current $P_\Lambda = -0.918$ is a conditional nominal-MC diagnostic. It is neither the inclusive LEP observable nor a final CEPC measurement.