

Supplementary Information

Cross-trait fine-mapping and sequence modeling prioritize an AHSG splicing mechanism at an otosclerosis locus
Baptiste Dumoulin, MSc
Clinical Research Unit, CHU Brugmann, Place A. Van Gehuchten 4, 1020 Brussels, Belgium
Corresponding author: baptiste.dumoulin@chu-brugmann.be

Supplementary Information accompanying the revised, non-peer-reviewed preprint of this study (13 October 2026).
This supplement documents the primary disease–splicing analysis, historical-variant provenance, locus-boundary and LD sensitivities, secondary splice-signal colocalization, and associated computational checks. Section S17 adds a post hoc protein-QTL analysis and its focused 13 October 2026 model audit. Favorable and unfavorable multi-signal sensitivity results are retained; the protein evidence remains exploratory.

S1. Analysis hierarchy

hierarchy	definition	implementation
Primary molecular endpoint	AHSG liver exon-5-to-6 splicing event	GTE _x V11 SuSiE fine-mapping; GTE _x V8 targeted direction audit
Primary disease endpoint	Otosclerosis association at AHSG	Published three-biobank meta-analysis summary statistics
Primary cross-trait question	Do disease and exon-5-to-6 molecular signals converge on the same fine-mapped candidate set?	Descriptive five-variant credible-set concordance for the exact primary exon-5-to-6 event; complete primary regional statistics remain unavailable. Formal signal-level colocalisation is completed only for the distinct secondary exon-6-to-7 QTD000270 signal.
Primary sequence-model contrast	rs4917 C versus rs4918 C isolated substitutions	AlphaGenome isolated-variant controls
Secondary analyses	Exon-6-to-7 coloc; alternative phase; broader RNA tracks; systemic biomarker triangulation; Beethoven application	Secondary exon-6-to-7: formal signal-level SuSiE log-Bayes-factor colocalisation completed, with earlier coloc.abf sensitivity retained. Other secondary/exploratory analyses unchanged.

S2. Source GWAS design and comparison with present reanalysis

Rämö et al. analysed FinnGen, EstBB and UKBB using SAIGE mixed models and combined effect estimates by fixed-effect meta-analysis. Their published fine-mapping used FinnGen summary statistics with in-sample dosage LD, 3-Mb windows and up to ten effects. Project 2 retains a 350-kb meta-analysis re-fine-map using external European 1000 Genomes LD, adds a source-comparable 3-Mb summary-statistic boundary sensitivity centred on rs4917, repeats the 350-kb LD-aware SuSiE-RSS analysis using a large public UK Biobank British-ancestry LD matrix (N=337,545), and now adds formal signal-level colocalisation for the retained secondary QTD000270 AHSG exon-6-to-7 splice signal. The 3-Mb boundary analysis and UK Biobank LD sensitivity preserve the five-variant result. Exact in-sample LD for the full three-cohort meta-analysis and complete regional statistics for the exact primary exon-5-to-6 molecular event remain unavailable.

Source	Cases	Controls	Ancestry / selection	Association adjustment	Fine-mapping / role
FinnGen	1,563	249,281	Finnish ancestry; non-Finnish ancestry and twins/duplicates excluded	SAIGE v0.20 mixed model; age, sex, 10 PCs; genotyping batch; INFO >0.6	Source study fine-mapping used FinnGen summary statistics with in-sample dosage LD (LDStore2), 3-Mb windows, max 10 effects
Estonian Biobank	985	196,516	European ancestry; related individuals included	SAIGE v0.20 mixed model; first 10 PCs, birth year, sex	Contributed to fixed-effect meta-analysis; no present-study cohort-specific LD available in archive
UK Biobank	956	415,401	European ancestry assigned/refined under Pan-UKBB framework	SAIGE v0.20; age, sex, age*sex, age^2, age^2*sex, first 10 PCs; INFO >0.8 and AC >=20 retained	Contributed to fixed-effect meta-analysis; no present-study cohort-specific LD available in archive
Meta-analysis	3,504	861,198	Three biobank cohorts	Fixed-effect meta-analysis of cohort effect estimates; MAF <0.1% excluded within cohorts	Current project re-fine-maps meta-analysis summary statistics with European 1000G LD; large UKBB British-ancestry LD sensitivity added (N=337,545)

23 S3. Novelty benchmark against prior work

Component	Ramo et al. 2023	Related Beethoven susceptibility paper	Present study (revised preprint)
AHSG association	Yes	Uses/reports selected susceptibility signal	Uses source GWAS
Disease fine-mapping	FinnGen-only; in-sample dosage LD; 3-Mb windows; SuSiE	Not central	350-kb meta-analysis re-fine-map with 1000G EUR LD; source-comparable 3-Mb single-effect boundary sensitivity; large UKBB British-ancestry LD sensitivity (N=337,545)
Primary GTEx exon-5-to-6 SuSiE comparison	Not identified as the same cross-trait comparison	No	Yes
Multi-variant AlphaGenome haplotype modelling	No	No	Yes
rs4917 vs rs4918 isolated perturbation	No	No	Yes
Formal cross-trait colocalisation	No	No	Primary exon-5-to-6: not identifiable from the complete regional statistics available in retained public resources. Secondary exon-6-to-7 QTD000270 signal: formal signal-level analysis completed; PP.H4=0.950-0.999 across p12 sensitivity.
Beethoven application	No	Yes	Contextual only; not evidence establishing the mechanism

24 S4. Historical variant provenance and verification status

25 The following table consolidates the five recorded historical calls used only for contextual sequence scenarios. Retained

26 audit material includes direct source-VCF call-level re-audits for rs2077119, rs4917, rs4918 and rs9870756 and

27 callable-mask confirmation for rs2077119, rs4917 and rs9870756. A comparable direct source-VCF re-audit was not

retained for rs1071592, and callable-mask confirmation was not retained for rs4918 or rs1071592. These limitations do not affect the disease or molecular fine-mapping analyses and no historical phase or diagnosis is inferred.

Variant	GRCh37	GRCh38	REF/ALT	Historical GT	Observed alleles	DP/GQ
rs2077119	3:186330462	3:186612673	T/G	0/1	G/T	37/334
rs4917	3:186337713	3:186619924	T/C	0/1	C/T	17/111
rs4918	3:186338382	3:186620593	G/C	0/1	C/G	42/407
rs1071592	3:186338425	3:186620636	A/C	0/1	A/C	64/530
rs9870756	3:186344614	3:186626825	C/T	0/1	C/T	34/202

Verification status and cross-dataset notes

Variant	Primary-source verification status	Cross-dataset note
rs2077119	Call re-audited; callable mask confirmed; unphased; no read-level validation	GTEx/QTL and model alleles should be reconciled against source before submission
rs4917	Call re-audited; callable mask confirmed; unphased; no read-level validation	C is molecular-QTL/model ALT and disease protective effect allele in retained tables
rs4918	Call re-audited; callable-mask confirmation not retained; unphased; no read-level validation	C is molecular-QTL/model ALT and disease protective effect allele in retained tables
rs1071592	REPORTED in retained historical call table; direct source-VCF re-audit PENDING; callable-mask confirmation PENDING; unphased	Targeted V8 molecular-QTL row retained
rs9870756	Call re-audited; callable mask confirmed; unphased; no read-level validation	No significant-only V8 molecular-QTL row returned

S5. Reference-supported five-site haplotypes

Five-site alleles in order	Reference chromosomes (n=1,006)
G/C/C/C/C	444
T/C/C/C/C	236
T/T/G/A/C	123
T/T/G/A/T	107
T/T/G/C/C	86
G/T/G/A/T	5
G/T/G/A/C	3
G/T/G/C/C	2

Among the 50 reference individuals sharing Beethoven's five unphased heterozygous genotypes, 49 supported the primary configuration. This frequency is descriptive and is not a calibrated posterior for Beethoven's phase.

S6. AlphaGenome functional-model controls

Liver model track	rs4917 C alone	rs4918 C alone	Both coding C	Full H1-H2	Alt-phase C/C-T/G
polyA+ (+)	+0.009222	+0.020250	+0.026600	+0.026449	+0.020368
polyA+ (.)	+0.009194	+0.020355	+0.025690	+0.024959	+0.019491
total RNA (+)	+0.009064	+0.018572	+0.026219	+0.018247	+0.014308
Sequence scenario	Exon-6 acceptor polyA+	Exon-6 acceptor total	Junction 5-to-6 polyA+	Junction 5-to-7 polyA+	
Primary H1 / AHSG*1	0.878906	0.875000	11.195685	12.504654	
Primary H2 / AHSG*2	0.816406	0.812500	10.410587	12.734492	
rs4917 C alone	0.878906	0.875000	11.119974	12.299032	
rs4918 C alone	0.820312	0.812500	10.450480	12.906081	
Both coding C	0.878906	0.875000	11.202762	12.490958	
Alt. phase C/C	0.878906	0.875000	11.208670	12.500649	
Alt. phase T/G	0.820312	0.812500	10.443488	12.780004	

S7. Primary GTEx V8 targeted molecular-QTL audit

The targeted V8 API audit returned 208 liver samples for the exact candidate phenotype. Significant-only records cannot support formal regional colocalisation because non-selected variants are missing.

SNP	GRCh38 variant	ALT	NES	P value
rs2077119	chr3:186612673 T>G	G	+0.524568	4.88e-9
rs4917	chr3:186619924 T>C	C	+0.713992	1.15e-17
rs4918	chr3:186620593 G>C	C	+0.697230	6.02e-17
rs1071592	chr3:186620636 A>C	C	+0.654425	4.78e-12
rs9870756	chr3:186626825 C>T	T	No significant-only row	Not returned

38 S8. Primary cross-trait five-variant result

Variant	GTEx PIP	GTEx rank	Disease PIP (1% shrink)	Disease rank	GWAS beta	GWAS P
rs4917	0.306570	1	0.292346	1	-0.175586	3.35e-12
rs4918	0.165667	5	0.255668	2	-0.175125	3.85e-12
rs1900618	0.165669	4	0.253146	3	-0.175053	3.90e-12
rs7635884	0.175413	2	0.137453	4	-0.173098	7.13e-12
rs13073106	0.175413	2	0.070965	5	-0.170393	1.48e-11

39 Interpretation boundary: identical five-variant credible sets are descriptive cross-trait fine-mapping concordance. They
40 are not a formal shared-causal posterior, mediation test or proof that rs4917 is causal.

41 S9. Candidate-set LD sensitivity within 1000 Genomes Europe

Population	N	Minimum pairwise r2	Median pairwise r2	Maximum pairwise r2
EUR	503	0.986927	0.991269	1.000000
CEU	99	0.979423	1.000000	1.000000
FIN	99	0.977884	0.977884	1.000000
GBR	91	1.000000	1.000000	1.000000
IBS	107	0.974378	0.974378	1.000000
TSI	107	1.000000	1.000000	1.000000

42 This analysis demonstrates that near-complete LD among the five candidates persists in each constituent European 1000
43 Genomes subpopulation. It remains useful as a within-reference resolution check. The new S15 analysis provides the
44 more important external-reference upgrade: a 337,545-participant UK Biobank British-ancestry LD matrix reproduces
45 the same five-variant disease credible set and confirms the near-complete LD block.

46 S10. Secondary exon-6-to-7 colocalisation

Region	Keep palindromic	p12	sdY	SNPs	PP.H3	PP.H4
full	True	1e-06	1.0	1139	0.051494	0.948506
full	True	1e-05	1.0	1139	0.005400	0.994600
full	True	5e-05	1.0	1139	0.001085	0.998915
full	False	1e-06	1.0	985	0.040863	0.959137
full	False	1e-05	1.0	985	0.004242	0.995758
full	False	5e-05	1.0	985	0.000851	0.999149
75kb	True	1e-06	1.0	419	0.051494	0.948506
75kb	True	1e-05	1.0	419	0.005400	0.994600
75kb	True	5e-05	1.0	419	0.001085	0.998915
75kb	False	1e-06	1.0	365	0.040863	0.959137
75kb	False	1e-05	1.0	365	0.004242	0.995758
75kb	False	5e-05	1.0	365	0.000851	0.999149
100kb	True	1e-06	1.0	601	0.051494	0.948506
100kb	True	1e-05	1.0	601	0.005400	0.994600
100kb	True	5e-05	1.0	601	0.001085	0.998915
100kb	False	1e-06	1.0	524	0.040863	0.959137
100kb	False	1e-05	1.0	524	0.004242	0.995758
100kb	False	5e-05	1.0	524	0.000851	0.999149
full	True	1e-05	0.8	1139	0.005342	0.994658
full	True	1e-05	1.2	1139	0.005440	0.994560

47 This analysis is secondary because it evaluates a different intron from the prespecified exon-5-to-6 target. The earlier
48 single-causal-variant coloc.abf sensitivity is retained below; an additional analysis performs formal signal-level Bayes-
49 factor colocalisation using the official eQTL Catalogue SuSiE per-effect log-Bayes-factor matrix and the preferred UK
50 Biobank-LD disease SuSiE fit. Neither secondary analysis should be used as a substitute for formal primary-event
51 colocalisation.

S11. Reproducibility framework

The analysis archive includes machine-readable figure sources, reporting checklists, R scripts and audit records. These improve traceability but do not replace missing original source data. The precise susieR version used for early disease PIPs was not recorded, the complete raw AlphaGenome output archive was not retained for independent re-execution, and the entire original workflow was not independently rebuilt from individual-level data. For the post hoc protein analysis, however, a focused Google Colab R rerun on 13 October 2026 completed L=3 and L=10 models at 1% and 5% external-reference LD shrinkage. Saved fitted L=10 objects, posterior outputs, purity diagnostics and session logs were audited. Source-cohort LD matrices were not independently reconstructed, and assay-level biological validation remains outstanding.

The reproducibility record distinguishes independently checked source-input arithmetic from model-conditional posterior summaries. The present Supplementary Information reports model versions, allele-harmonization conventions, selected variant PIPs, fitted-model sensitivity, and negative findings. The optional focused R rerun script is provided without third-party datasets; to execute it, users must obtain and harmonize the original source statistics and public UK Biobank LD matrix under their respective access conditions. Saved R fit objects and diagnostic logs are retained in the author's internal archive but are not themselves evidence of causal mediation.

S12. Adapted STREGA reporting checklist

STREGA topic	Status	Manuscript location	Implementation note
Study design	Addressed	Methods: Study design and analysis hierarchy	Secondary computational synthesis; no new recruitment
Participants/source cohorts	Addressed	Methods + Supplementary source-cohort table	Cohort sizes, ancestry selection and covariates summarized from Ramo et al.
Population stratification	Addressed upstream	Methods: Source GWAS and LD framework	Original cohort methods reported; present analysis cannot redo individual-level PCA
Genotyping/imputation error	Partly addressed	Supplementary source-cohort table; variant provenance table	Four of five historical calls re-audited directly at VCF call level; callable mask confirmed for 3/5; rs1071592 direct source audit and rs4918/rs1071592 mask confirmation remain pending; no read-level validation claimed
Haplotype inference	Addressed	Methods: Reference-supported historical scenarios	Modern-panel phase frequencies are descriptive, not a posterior for Beethoven
Hardy-Weinberg equilibrium	Upstream / N/A locally	Supplementary checklist note	Handled in source cohort QC where applicable; not recalculated from summary statistics
Replication	Partly strengthened	Discussion: LD robustness and limitations	Large UK Biobank LD sensitivity (N=337,545) reproduces the exact five-variant disease credible set; exact FinnGen/EstBB/meta-analysis in-sample LD and independent molecular-QTL replication remain unavailable.
Variant selection	Addressed	Methods: AlphaGenome and candidate selection	Primary endpoints/contrasts separated from exploratory outputs

STREGA topic	Status	Manuscript location	Implementation note
Multiple testing / multiplicity	Addressed conceptually	Methods: Statistical interpretation	Model output rows are not independent tests; original GWAS multiplicity preserved; Bayesian analyses use posterior/prior sensitivity
Relatedness	Addressed upstream	Supplementary source-cohort table	SAIGE mixed models/kinship in source cohorts
Data volume / reproducibility	Addressed with caveats	Reproducibility section and archive	Manifest/checksums provided; full raw AlphaGenome output and clean frozen rerun remain pending
Limitations/generalizability	Addressed	Discussion: Limitations	Tissue indirectness, LD, external reference, historical phase and no wet-lab validation stated explicitly

S13. Current limitations and interpretation boundaries

The exact primary exon-5-to-6 molecular phenotype could not be subjected to formal regional colocalization because complete non-selected molecular summary statistics were unavailable in the retained public resources. Exact in-sample LD for the full three-cohort disease meta-analysis was also unavailable, although the candidate set was unchanged under the large UK Biobank LD sensitivity. Historical-call verification is incomplete for rs1071592 and callable-mask status was not retained for rs4918/rs1071592. No wet-lab splice assay was performed. These limits constrain nucleotide-level interpretation but do not alter the reported locus-level fine-mapping and secondary shared-signal results.


S14. Full-locus summary-statistic sensitivity (3 Mb)

The complete otosclerosis meta-analysis is deposited under GWAS Catalog accession GCST90129575. The uploaded GRCh38 source file contained 13,615,309 rows; extraction of chr3:185,119,924-188,119,924 (+/-1.5 Mb around rs4917 at chr3:186,619,924) retained 14,507 rows with no malformed/skipped records. Standardization identified 13,532 complete biallelic SNVs for the full-locus summary-statistic analysis. rs4917 remained the marginal lead ($P=3.35 \times 10^{-12}$), and the five strongest marginal variants were rs4917, rs4918, rs1900618, rs7635884 and rs13073106. Eleven SNVs reached $P < 5 \times 10^{-8}$, all within a 9,879-bp interval; the best association more than 50 kb from rs4917 was rs775702472 at chr3:187,647,186 ($P=2.35 \times 10^{-5}$). Wakefield single-effect approximate-Bayes-factor fine-mapping was repeated across the full 3-Mb SNV set using prior variances $W=0.01, 0.04, 0.09, 0.16$ and 0.25 . Every prior produced the same five-variant 95% credible set: rs4917, rs4918, rs1900618, rs7635884 and rs13073106. rs4917 ranked first throughout (posterior probability 0.285539-0.290043); at $W=0.04$ the five variants carried 0.998850 of total posterior mass. This materially reduces concern about a narrow-window artifact, but it is a single-effect summary-statistic sensitivity, not an LD-aware multi-signal analysis and not an independent/cohort-matched LD replication.

Prior W	95% CS size	95% CS	rs4917 posterior	Top-five posterior mass
0.01	5	same five*	0.285539	0.998419
0.04	5	same five*	0.289026	0.998850
0.09	5	same five*	0.289697	0.998919
0.16	5	same five*	0.289934	0.998942

Prior W	95% CS size	95% CS	rs4917 posterior	Top-five posterior mass
0.25	5	same five*	0.290043	0.998952

Supplementary Table S14. Full-locus 3-Mb Wakefield single-effect prior sensitivity. The analysis uses 13,532 biallelic SNVs from GCST90129575. *For every tested prior, the 95% credible set was rs4917, rs4918, rs1900618, rs7635884 and rs13073106. This is a locus-boundary sensitivity and does not replace LD-aware multi-signal fine-mapping.



Supplementary Figure S1. Marginal otosclerosis association across the source-comparable 3-Mb AHSG window. The dashed line marks $P=5 \times 10^{-8}$; the five variants shared with the GTEx V11 exon-5-to-6 credible set are labelled. The absence of a second genome-wide-significant marginal peak does not exclude additional conditional signals.

S15. UK Biobank LD robustness sensitivity

A public UK Biobank British-ancestry LD matrix distributed with PolyFun was used as a large cohort-relevant sensitivity. The source matrix is based on 337,545 participants and is indexed on GRCh37 coordinates; GWAS variants were matched by rsID and allele pair to avoid build ambiguity. The approximately 350-kb GRCh38 disease interval contained 1,806 eligible biallelic SNVs, of which 1,544 harmonised to the UK Biobank LD chunk. At all retained variants, the matrix A1/A2 order was opposite to the GWAS effect/other allele order, so z scores were sign-flipped to A1 orientation before SuSiE-RSS. This systematic flip changes direction only; posterior fine-mapping is evaluated after allele alignment.

The preferred 1% shrinkage SuSiE-RSS analysis reproduced the exact five-variant disease credible set: rs4917, rs4918, rs1900618, rs7635884 and rs13073106. rs4917 remained first (PIP 0.292281), compared with 0.292346 using 1000 Genomes EUR LD. Baseline, 5% shrinkage and L=1 analyses returned the same five-variant credible set. Across 1,475 variants present in both preferred analyses, 1000 Genomes and UK Biobank PIPs were strongly correlated (Pearson

108 $r=0.994718$; median absolute PIP difference 0.002344). Candidate LD remained extreme in UK Biobank: minimum
 109 pairwise r^2 0.991815, median 0.995639 and maximum 1.000000. This substantially reduces concern that the candidate
 110 set or LD-resolution limit is an artifact of the small 1000 Genomes panel, but it does not provide exact in-sample LD for
 111 the full three-cohort meta-analysis.

Metric	UK Biobank LD sensitivity
UKBB LD reference size	337,545
Eligible biallelic SNVs in ~350-kb GWAS window	1,806
Harmonised SNVs	1,544
Unmatched eligible SNVs	262
Preferred model	SuSiE-RSS L=10, 1% diagonal shrinkage
95% disease credible set	rs4917; rs4918; rs1900618; rs7635884; rs13073106
rs4917 rank / PIP	1 / 0.292281
Candidate minimum / median / maximum r^2	0.991815 / 0.995639 / 1.000000
1000G-vs-UKBB PIP correlation on 1,475 shared variants	Pearson $r=0.994718$

112 Supplementary Table S15a. Harmonisation and principal results for the UK Biobank LD robustness sensitivity. The
 113 public matrix is cohort-relevant to UK Biobank but is not exact in-sample LD for the full three-cohort otosclerosis
 114 meta-analysis.

Variant	GTEx V11 sQTL PIP	Disease PIP 1000G EUR	Disease PIP UKBB	UKBB rank
rs4917	0.306570	0.292346	0.292281	1
rs4918	0.165667	0.255668	0.255674	2
rs1900618	0.165669	0.253146	0.253167	3
rs7635884	0.175413	0.137453	0.143895	4
rs13073106	0.175413	0.070965	0.073912	5

115 Supplementary Table S15b. Five-variant comparison across GTEx V11 and the two disease LD references. UK Biobank
 116 and 1000 Genomes disease PIPs are from the preferred 1% shrinkage model.

117 **S16. Formal signal-level colocalization for the retained secondary AHSG splice event**

118 The eQTL Catalogue GTEx liver LeafCutter dataset QTD000270 connected-component output retained one AHSG
 119 molecular trait in the target region: chr3:186619940-186620586:clu_35430_+, with 1,316 regional association rows.
 120 Endpoint parsing confirms that this is the secondary exon-6-to-7 event; the exact primary exon-5-to-6 phenotype was
 121 not retained in the connected-component output. The official QTD000270 SuSiE log-Bayes-factor file provided 1,284
 122 variants for the retained trait, and 1,107 rsIDs overlapped the preferred UK Biobank-LD disease SuSiE-RSS fit. The
 123 signal-pair audit identified one active molecular effect led by rs7635884 and one active disease effect led by rs4917.
 124 Signal-level Bayes-factor colocalisation strongly favoured a shared regional signal under every prespecified p_{12} prior:
 125 $PP.H4=0.950052$ at $p_{12}=1e-6$, 0.994770 at $p_{12}=1e-5$ and 0.998950 at $p_{12}=5e-5$. Corresponding $PP.H3$ values
 126 decreased from 0.049948 to 0.001050. These results strengthen the evidence that an AHSG splice-regulatory signal and
 127 the otosclerosis association share a regional causal signal, but they concern the secondary exon-6-to-7 molecular event
 128 and must not be represented as formal colocalisation of the exact primary exon-5-to-6 phenotype. The earlier coloc.abf
 129 analysis of this secondary region remains a complementary single-causal-variant sensitivity.

Supplementary Table S16. Formal signal-level colocalisation of the retained QTD000270 AHSG exon-6-to-7 splice signal with the preferred UK Biobank-LD otosclerosis disease fine-map. $p_1=p_2=1e-4$; 1,107 variants were shared. eQTL lead rs7635884; disease lead rs4917.

p12	Shared variants	eQTL lead	Disease lead	PP.H3	PP.H4	Interpretation
1e-6	1,107	rs7635884	rs4917	0.049948	0.950052	Strong shared-signal support
1e-5	1,107	rs7635884	rs4917	0.005230	0.994770	Very strong shared-signal support
5e-5	1,107	rs7635884	rs4917	0.001050	0.998950	Very strong shared-signal support

Reproducibility note: the earlier secondary-colocalization analysis archive retained the selected molecular-trait status, the 1,284-variant molecular log-Bayes-factor matrix, 1,107 shared rsIDs, signal-pair posterior table, audit metadata and R session information (R 4.6.1; coloc 5.2.3; susieR 0.14.2). The optional classical coloc.abf rerun in the same Colab session failed because the initial script passed effect-allele frequency rather than minor-allele frequency on the disease side; that optional rerun is not used for the signal-level result reported above. The previously retained coloc.abf sensitivity remains separately documented in Supplementary Table S10.

S17. Exploratory AHSG disease-protein analysis and model sensitivity

Status and role. This post hoc secondary analysis was added after the first publicly posted preprint. It is not a prespecified primary endpoint, a formal demonstration of protein mediation, or a direct investigation of Beethoven's circulating AHSG concentration. Its results derive from the 9 October evidence addendum, original analysis notebooks, and a focused 13 October R rerun. All favorable and unfavorable model settings are reported, and the original disease-splicing analysis remains the primary endpoint.

S17.1 Data provenance and allele alignment

Otosclerosis: Rämö et al. (2023) three-biobank meta-analysis, 3,504 cases and 861,198 controls; rs4917 C versus T effect estimate $\log OR=-0.175586$ ($P=3.35 \times 10^{-12}$). Protein: UK Biobank Pharma Proteomics Project European-discovery Olink AHSG P02765 assay OID30706 (Sun et al. 2023), 33,985 participants; rs4917 C versus T Olink assay $\beta=+0.675802$, $SE=0.007918$. The released pQTL beta is not a concentration change expressed in g/L. All allele signs are harmonized; the pQTL regression and disease model are not on a common effect scale. A single-causal-signal analysis contained 1,659 regional harmonized variants. For LD-aware SuSiE, 1,321 distinct allele-pair variant keys were retained and matched to a UK Biobank British-ancestry LD reference derived from 337,545 participants. As detailed below, one rsID corresponds to two different allele pairs; variant-key identity must be maintained rather than collapsing this multi-allelic marker.

S17.2 Colocalization and sensitivity results

The reported coloc.abf calculation ($p_1=p_2=10^{-4}$; $p_{12}=10^{-5}$) was $PP.H4=0.997064$ over 1,659 variants. The 9 October addendum describes numerical reproduction using R and Python across 12 matched settings (largest $PP.H4$ discrepancy 4.44×10^{-16}); this verification was documented in the source addendum and not newly rerun here. The LD-aware results below compare 1% and 5% diagonal shrinkage at $p_{12}=10^{-5}$. L=1, 3 and 5 values refer to a principal disease-protein signal pair; L=10 values are maxima across potentially different signal pairs. They must not be compared as if pair identity were identical. The protein model was reported to saturate its effect-count cap. The sharp L=10 disagreement prohibits presenting the high constrained-model $PP.H4$ as a robustly established shared-causal mechanism.

Analysis	Variants / model	Reported PP.H4 (p12=10 ⁻⁵)	Interpretive boundary
Single-signal coloc.abf	1,659 regional variants	0.997064	Single-causal-signal assumption; R/Python numerical agreement reported in addendum
coloc.susie constrained	1,321 allele-pair keys; L=1, 3 or 5; 1%-5% LD shrinkage	Principal pair: 0.99709-0.99734	Prior-, LD- and signal-selection dependent
coloc.susie less constrained	1,321 allele-pair keys; L=10, 1% LD shrinkage	Maximum pair: 0.98686	Maximum across reported pairs; effect-cap saturation
coloc.susie less constrained	1,321 allele-pair keys; L=10, 5% LD shrinkage	Maximum pair: 0.00279	Major contradiction: not robust to LD shrinkage

Supplementary Table S17a. Post hoc disease-Olink AHSG comparison and explicit multi-signal sensitivity. These posterior probabilities are not frequencies of true causal mechanisms or estimates of the probability that Beethoven had otosclerosis.

S17.2a. Variant PIP reallocation and signal-pair identity sensitivity

An initial post-upload audit of 1,321 allele-aware L=10 PIP records and the completed per-signal posterior tables detected large model sensitivity. In that audit the highest L=10 PP.H4 values were 0.986858 at 1% shrinkage (rs13073106 protein lead) and 0.002788 at 5% (rs1029353); these maxima represent different signal pairs. A subsequent focused Colab R rerun on 13 October reproduced these conditional results and saved the fitted-model objects and diagnostics (Section S17.6).

Between L=10 shrinkage settings, disease PIPs remained similar (Pearson $r=0.99991$), but protein PIPs did not ($r=0.64951$). Protein PIP for rs13073106 shifted from 1.000 to approximately 0. Conversely, at L=3 and 5% shrinkage, rs1900618, rs4918 and rs4917 carried protein PIPs 0.567, 0.235 and 0.198; each was approximately zero under both L=10 fits. The protein model reached the allowed number of credible sets in every configuration. The new fit objects, logs and credible-set diagnostics are now available and confirm computational convergence, but the responsible model/LD or assay mechanism remains unresolved.

Variant	Protein PIP L=3, 5%	Protein PIP L=10, 1%	Protein PIP L=10, 5%
rs4917	0.198	0.000	0.000
rs1900618	0.567	0.000	0.000
rs4918	0.235	0.000	0.000
rs13073106	<0.001	1.000	0.000
rs35457250	1.000	1.000	1.000
rs140827890	1.000	1.000	1.000

Supplementary Table S17b. Protein PIPs from the supplied completed-model outputs, compared using allele-aware variant keys. PIPs are model-conditional and do not imply causation in Beethoven.

S17.3 Published contextual evidence, not independent causal replication

Published CHARGE data linked rs4917 T to approximately 0.066 g/L lower fetuin-A measured by ELISA (Jensen et al. 2017), and PREVEND linked rs4917 T to lower serum calcification propensity T50 (de Haan et al. 2022). These prior findings support the plausibility of systemic AHSG-linked mineralization biology, but rs4917 and nearby candidates are very highly correlated and those earlier results do not demonstrate that protein abundance mediates the otosclerosis association. Coding changes rs4917/rs4918 may influence Olink antibody binding (epitope effects); differing assays and protein forms limit cross-platform interpretation (Eldjarn et al. 2023).

S17.4 Independent arithmetic audit of harmonized input (13 October 2026)

The notebook-embedded 1,321-row harmonized summary-statistic input was independently decoded and checked before executing the focused rerun. It contained 1,321 unique allele-aware variant keys and 1,320 unique rsIDs; rs13094300 occurs twice with distinct allele pairs, requiring allele-aware joins. Signed-Z consistency and missing-data checks identified no arithmetic mismatch in the tested fields. The same input was subsequently used in the 13 October focused Colab R rerun. The full external UK Biobank LD was downloaded and used in Colab, but was not separately re-derived from primary genotypes or independently checked against in-sample LD in this document.

Audit item	Result	Scope / unresolved requirement
Harmonized allele-pair rows	1,321; all distinct	Full 1,321-row embedded file only; no external source re-download

Audit item	Result	Scope / unresolved requirement
Unique rsIDs	1,320	rs13094300 has two different allele pairs; maintain variant_key
Disease / protein Z orientation mismatches	0 / 0	Arithmetic test against signed GWAS and pQTL beta/SE
Missing/non-finite tested Z values	0	Does not test source-level missingness before filtering
LD-aware posterior fit	Focused R rerun completed	Four converged model configurations with saved L=10 RDS fits and QC logs; effect-cap saturation and LD sensitivity unresolved

Supplementary Table S17c. Independent arithmetic checks conducted for the revised preprint. This limited input audit is NOT an independent replication of the disease-protein colocalization posterior.

S17.5 Evidence synthesis figure and caveats

AHSG / otosclerosis: triangulation and LD-model sensitivity

Post-submission exploratory evidence — not a causal pathway demonstration

A Allele-aligned associations

Disease association rs4917 C: lower otosclerosis odds Râmă et al. 2023 / GWAS reanalysis
Liver splice signal Same five top fine-mapped variants GTEx V11; primary event not formally coloc-tested
Olink AHSG assay rs4917 C: higher measured AHSG UKB-PPP N = 33,985; assay may be epitope-sensitive
ELISA fetuin-A rs4917 T: lower fetuin-A (~0.066 g/L) CHARGE 2017; previously published
Serum calcification T50 rs4917 T: lower T50 PREVEND 2022; previously published

C/T contrasts are allele-oriented; effect scales differ across assays.

B SuSiE signal-pair PP.H4

	1% LD shrinkage	5% LD shrinkage
L = 1	99.71%	99.71%
L = 3	99.71%	99.71%
L = 5	99.73%	99.73%
L = 10	98.69%	0.28%

For $L \leq 5$: selected main disease-protein pair.

For $L = 10$: maximum over pairs; not directly equivalent.

All values conditional on model assumptions; $p_{12} = 10^{-5}$.

Supplementary Figure S2. Direction-aligned disease, splice, Olink, ELISA and serum T50 associations (left); disease-protein SuSiE sensitivity to model effect count and LD shrinkage (right). The $L \leq 5$ values summarize a selected principal signal pair, while the $L=10$ cells show maxima across possibly different pairs, so the latter are not pair-matched. This post hoc figure is descriptive and does not demonstrate causal mediation, otic-capsule biology, or a Beethoven-specific molecular mechanism.

S17.6 Focused 13 October R rerun: completed fits, purity and posterior sensitivity

The four disease/protein model pairs were executed in Google Colab using coloc 5.2.3, susieR 0.14.2, an allele-matched 1,321-variant input, effective disease $N=13,959.2$ and quantitative protein $N=33,985$. The public 337,545-participant UK Biobank British LD matrix was used with 1% or 5% diagonal shrinkage. Both traits reported convergence in all four settings; RDS objects were saved for both $L=10$ fits. Saved logs, R version information, credible-set purity files and normalized signal-pair posteriors were inspected. Convergence confirms computational completion, not scientific validity of signal decomposition. All protein fits had as many reported credible sets as their permitted effect count.

Supplementary Table S17d. Newly rerun $L=3$ and $L=10$ models (13 October 2026). PP.H4 is the H4 posterior of the highest-ranking disease-protein selected signal pair at $p_{12}=10^{-5}$, not a locus-level probability or the posterior for the same protein effect across all fits. Colocalization pairs change between model settings.

Model	LD shrink	Disease CS	Protein CS / cap	Top PP.H4*	Protein partner*
s01_L3	1%	1	3/3	0.997125	rs1900618

Model	LD shrink	Disease CS	Protein CS / cap	Top PP.H4*	Protein partner*
s05_L3	5%	1	3/3	0.997122	rs1900618
s01_L10	1%	1	10/10	0.986858	rs13073106
s05_L10	5%	1	10/10	0.002788	rs1029353

All four models selected one disease credible set. Among protein credible sets, the L=3 fits contained two singleton sets and one three-variant set; L=10 at 1% shrinkage contained nine singleton sets, whereas L=10 at 5% contained six singleton sets and three broad sets of 3–16 variants. The protein model did not demonstrate that the effect-count cap was sufficient. All 78 reported signal-pair posterior rows (four models × three priors × three or ten pairs) summed to one across H0–H4 to within 8.7×10^{-13} . No independently derived in-sample LD or independent-cohort biological replication was available. *Pairs are model-selected and not necessarily the same underlying association across conditions.

Supplementary Table S17e. Selected protein posterior inclusion probabilities (PIPs) across model settings, from the focused 13 October rerun. Printed values near zero are conditional posterior allocations and do not establish absence of biological effects.

Variant	L3 1%	L3 5%	L10 1%	L10 5%
rs4917	0.2003	0.1979	<0.0001	<0.0001
rs4918	0.2338	0.2346	<0.0001	<0.0001
rs1900618	0.5659	0.5675	<0.0001	<0.0001
rs13073106	<0.0001	<0.0001	1.0000	<0.0001
rs1029353	<0.0001	<0.0001	<0.0001	0.2102

In the L=10 fits, disease PIPs were highly correlated between 1% and 5% shrinkage ($r=0.99991$), but protein PIPs were less stable ($r=0.64951$). Under L=3, the strongest shared-protein signal was associated with rs1900618; under L=10 the highest PP.H4 pair involved rs13073106 (1% shrinkage) or rs1029353 (5% shrinkage). At $p_{12}=10^{-6}$, 10^{-5} and 5×10^{-5} , the L=10 / 5%-shrinkage maximal PP.H4 values were 0.000280, 0.002788 and 0.013788, respectively. This pattern is incompatible with presenting high protein colocalization as insensitive to reasonable model settings.

S17.7 Remaining limitations and next analytical steps

All four focused models converged, and the requested L=10 fitted objects and credible-set purity diagnostics were saved. Nevertheless, exact overlap of UK Biobank participants between disease and protein GWAS remains unknown. The external UK Biobank British-ancestry LD matrix is not exact in-sample LD for every contributing disease and protein cohort. Protein credible sets saturated their respective L caps; at L=10, 1% and 5% LD shrinkage produced sharply different variant assignments and colocalization posteriors. Olink cis-pQTLs near protein-altering variants may partly reflect differential antibody/epitope recognition rather than circulating protein quantity. The exact primary exon-5-to-6 molecular phenotype lacks complete regional inputs for formal colocalization, and no tissue-level experiment establishes an AHSG mechanism in Beethoven. Future analysis should assess alternative multi-effect settings and better-matched LD, orthogonal protein assays, source-sample overlap, and the biological meaning of candidate variants. These limitations do not directly overturn the reported primary disease–splicing candidate-set comparison, but they make the protein component explicitly exploratory.

Source references specific to S17: Sun BB et al. Nature 622:329–338 (2023), doi:10.1038/s41586-023-06592-6; Eldjarn GH et al. Nature 622:348–358 (2023), doi:10.1038/s41586-023-06563-x; Jensen MK et al. Hum Mol Genet 26:2156–2163 (2017), doi:10.1093/hmg/ddx091; de Haan A et al. Front Cardiovasc Med 8:809717 (2022), doi:10.3389/fcvm.2021.809717; Rämö JT et al. Nat Commun 14:157 (2023), doi:10.1038/s41467-022-32936-3.