Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Quantile QTL mapping reveals heterogeneous allelic effects in the aging human brain and Alzheimer's disease

Anjing Liu

al4225@cumc.columbia.edu

Center for Statistical Genetics, The Gertrude H. Sergievsky Center, Columbia University, New York, NY, USA https://orcid.org/0009-0000-0128-1877

Roulan Jiang

Department of Statistics and Data Science, Tsinghua University, Beijing, China

Ruixi Li

Center for Statistical Genetics, The Gertrude H. Sergievsky Center, Columbia University, New York, NY, USA https://orcid.org/0009-0000-3071-4678

Xuewei Cao

Center for Statistical Genetics, The Gertrude H. Sergievsky Center, Columbia University, New York, NY, USA

Zining Qi

Department of Human Genetics, The University of Chicago, Chicago, IL, USA

Ru Feng

Center for Statistical Genetics, The Gertrude H. Sergievsky Center, Columbia University, New York, NY, USA

Hao Sun

Graduate School of Biomedical Sciences, Icahn School of Medicine at Mount Sinai, New York, NY, USA

Masashi Fujita

Center for Translational & Computational Neuroimmunology, Columbia University, New York, NY, USA

Natacha Comandante-Lou

Center for Translational & Computational Neuroimmunology, Columbia University, New York, NY, USA

Chirag M Lakhani

New York Genome Center, New York, NY, USA

Jenny Empawi

Department of Biomedical Genetics, Boston University Chobanian & Avedisian School of Medicine, Boston, MA, USA

David A. Knowles

Department of Computer Science, Columbia University, New York, NY, USA

Xiaoling Zhang
Department of Biomedical Genetics, Boston University Chobanian & Avedisian School of Medicine, Boston, MA, USA
Kushal K. Dey
Computational and System Biology, Sloan Kettering Institute, Memorial Sloan Kettering Cancer Center, New York, NY, USA
Philip De Jager
Center for Translational & Computational Neuroimmunology, Columbia University, New York, NY, USA
David Bennett
Rush Alzheimer's Disease Center and Department of Neurological Sciences, Rush University Medical Center, Chicago, IL, USA
Tianying Wang
tianying.wang@colostate.edu

Department of Statistics, Colorado State University, Fort Collins, CO, USA
Gao Wang
wang.gao@columbia.edu

Center for Statistical Genetics, The Gertrude H. Sergievsky Center, Columbia University, New York, NY, USA

Research Article

Keywords:

Posted Date: November 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8219833/v2

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Quantile QTL mapping reveals heterogeneous allelic effects in the aging human brain and Alzheimer's disease

Anjing Liu[1,†], Roulan Jiang[2,†], Ruixi Li[1], Xuewei Cao[1,3], Zining Qi[1,4], Ru Feng[1,3], Hao Sun[1,5], Masashi Fujita[6,16], Natacha Comandante-Lou[6,16], Chirag M. Lakhani[7], Jenny Empawi[8], David A. Knowles[7,9,10,11], Xiaoling Zhang[8,12], Kushal K. Dey[3,13,14], Philip De Jager[6,15,16], David Bennett[17], The Alzheimer's Disease Functional Genomics Consortium , Tianying Wang[18,*], and Gao Wang[1,16,19,*]

[1] Center for Statistical Genetics, The Gertrude H. Sergievsky Center, Columbia University, New York, NY, USA
[2] Department of Statistics and Data Science, Tsinghua University, Beijing, China
[3] Computational and Systems Biology Program, Sloan Kettering Institute, Memorial Sloan Kettering Cancer Center, New York, NY, USA
[4] Department of Human Genetics, The University of Chicago, Chicago, IL, USA
[5] Graduate School of Biomedical Sciences, Icahn School of Medicine at Mount Sinai, New York, NY, USA
[6] Center for Translational & Computational Neuroimmunology, Columbia University, New York, NY, USA
[7] New York Genome Center, New York, NY, USA
[8] Department of Biomedical Genetics, Boston University Chobanian & Avedisian School of Medicine, Boston, MA, USA
[9] Department of Computer Science, Columbia University, New York, NY, USA
[10] Department of Systems Biology, Columbia University, New York, NY, USA
[11] Data Science Institute, Columbia University, New York, NY, USA
[12] Department of Biostatistics, Boston University School of Public Health, Boston, MA, USA
[13] Physiology, Biophysics and Systems Biology, Weill Cornell Medicine, New York, NY, USA
[14] Gerstner Sloan Kettering Graduate School of Biomedical Sciences, New York, NY, USA
[15] Taub Institute for Research on Alzheimer's Disease and the Aging Brain, Columbia University, New York, NY, USA
[16] Department of Neurology, Columbia University, New York, NY, USA
[17] Rush Alzheimer's Disease Center and Department of Neurological Sciences, Rush University Medical Center, Chicago, IL, USA
[18] Department of Statistics, Colorado State University, Fort Collins, CO, USA
[19] Department of Biostatistics, Columbia University, New York, NY, USA
[†] These authors contributed equally to this work
[*] Correspondence: Gao Wang, wang.gao@columbia.edu; Tianying Wang, tianying.wang@colostate.edu

Abstract

Molecular quantitative trait locus (QTL) maps connect genetic variants to brain gene expression and protein abundance, but a mean effect does not show where an allelic effect is largest across the adjusted molecular distribution. Using data from more than 1,500 donors, we mapped cis effects across conditional quantiles in 26 datasets spanning 14 molecular contexts in the aging human brain, including bulk expression, single-nucleus expression and protein abundance. Quantile QTL (qQTL) mapping identified associations missed by matched linear QTL mapping (qQTL-only, 6.5% of classified associations) and, among associations detected by both models, effects that varied across quantiles under a trend-based heterogeneity test (26.2%). Functional annotation enrichment and conditional trait heritability analyses supported qQTLs as regulatory signals alongside mean-effect QTL annotations. Variance and interaction QTL analyses provided complementary evidence for a dispersion effect at *PITRM1* and an exploratory *APOE* ε4 interaction at *TMEM106B*.

At the *PICALM* locus, the Alzheimer's disease (AD) risk allele was associated with lower expression in microglia and with a larger estimated reduction at lower than at upper conditional quantiles in astrocytes. Quantile-aware transcriptome-wide association analysis (qTWAS) identified AD candidate genes using genetic predictors constructed for different quantile bands of molecular traits. We then characterized these candidates using their association results across quantile bands, matched standard transcriptome-wide association results and external genetic evidence. This resource adds molecular and cellular context to the interpretation of AD risk loci and prioritizes candidate genes for studies of gene regulation in the aging brain.

Introduction

Genome-wide association studies (GWAS) have identified many Alzheimer's disease (AD) risk loci, including noncoding loci whose interpretation depends on regulatory evidence that connects variants to genes, molecular traits and cell contexts.[1–3] Molecular quantitative trait locus (QTL) maps provide this evidence by linking genetic variation to RNA expression, protein abundance and cell-type-specific molecular phenotypes across human tissues and aging brain cohorts.[4–9] For AD genetics, we need to understand whether a risk allele produces a broad shift in molecular abundance, has stronger effects in part of its conditional distribution or changes expression variability.

A mean-effect QTL summary estimates one allelic effect for each tested variant and molecular trait. It can identify a regulatory association, but it cannot distinguish a broad expression shift from an effect concentrated in lower or upper conditional quantiles. At AD loci, an allelic effect may shift expression across a tissue, act mainly in the lower part of the adjusted expression distribution, alter molecular dispersion or vary with a measured modifier. Sequence-to-function studies provide a complementary motivation by showing that regulatory prediction depends on cell type, disease state and personal genome context.[10,11] We therefore examined both whether a variant was associated with a molecular trait and where its estimated effect was strongest within that trait's conditional distribution.

Heterogeneity QTL methods characterize variation in genetic effects. Interaction QTL (iQTL) analyses test measured modifiers such as sex, age or cellular composition, whereas variance QTL (vQTL) analyses identify variants associated with molecular dispersion, consistent with effects that differ across unmeasured contexts.[12–15] These approaches establish that regulatory effects can vary by context, but they do not localize the largest allelic effect along the adjusted molecular distribution. Quantile regression provides that localization by estimating allelic effects at multiple conditional quantiles and has been applied to biomarker GWAS and expression QTL discovery.[16–19] In QTL mapping, these estimates show whether a variant effect is broad, concentrated in one tail or linked to dispersion.

Here, we map allelic effects across conditional quantiles in the aging human brain to refine AD locus interpretation and prioritize candidate genes. The analysis spans 14 molecular contexts encompassing bulk expression, single-nucleus expression and protein abundance.[5–9] We define a context as a tissue, cell type or molecular modality, and a dataset as a processed phenotype matrix for a given context. Our primary analysis comprised 26 datasets spanning 14 contexts, including six single-nucleus cell types from the Religious Orders Study and Rush Memory and Aging Project (ROSMAP), each available from three analysis sets: CUIMC1, MIT and Mega, a combined analysis of CUIMC1, MIT and a further ROSMAP single-nucleus dataset (CUIMC2) (Supplementary Table 1). After accounting for overlap across datasets and processing sources, the primary analysis included 1,515 unique donors. Of these, 901 were from ROSMAP, 430 from the Knight Alzheimer's Disease Research Center and 184 from the Mount Sinai Brain Bank (MSBB) proteomics dataset.

Across these datasets, quantile QTL (qQTL) mapping identifies cis effects that are broad, concen-

trated in lower or upper conditional quantiles, or linked to molecular dispersion. Variants detected only by qQTL mapping were more enriched in regulatory annotations than variants detected by mean-effect QTL mapping. qQTL annotations were also associated with brain and aging trait heritability after accounting for mean-effect QTL annotations. At *PICALM*, functional work links the AD risk allele to microglial regulation and reduced *PICALM* expression.[20] For the same AD risk allele, we estimated a larger reduction in *PICALM* expression at lower than at upper conditional quantiles in astrocytes. Quantile-aware transcriptome-wide association analysis (qTWAS) then uses quantile-specific regulatory evidence to prioritize AD candidate genes while evaluating their molecular context and external genetic support. This resource combines quantile effects, molecular context and external genetic evidence to support studies of regulatory variation at AD loci and prioritize candidate genes for further investigation.

Results

Quantile mapping reveals additional associations and heterogeneous effects in the aging brain

Quantile QTL (qQTL) mapping in the aging brain detected *cis* associations missed by linear regression and revealed allelic effects that varied across conditional quantiles of the molecular trait distribution. For each variant and molecular trait, quantile regression estimated $\beta(\tau)$ on the grid $\tau \in \{0.05, 0.10, ..., 0.95\}$, where $\beta(\tau)$ is the change in the τth conditional quantile of the trait per additional copy of the alternative (coded) allele (Fig. 1a). Significant qQTL associations were identified by combining quantile-specific rank-score tests and applying hierarchical multiple-testing correction to control the gene-level false discovery rate (FDR) at 5% (Methods). In comparison, standard linear QTL (xQTL) analysis estimated a single allelic effect β on the conditional mean and identified significant associations using the same hierarchical multiple-testing procedure. We classified the associations according to whether they were detected as qQTL, xQTL, or both. qQTL-only associations can arise from an allelic effect concentrated in one tail of the trait distribution, or from a scale-only effect with opposing changes in the lower and upper tails (Fig. 1b). Because such effects yield only a weak shift in the conditional mean, they are often missed by standard xQTL analysis but captured by qQTL analysis that evaluates association at each quantile. Among shared associations detected by both analyses, we further identified a subset with heterogeneous effects using a trend-based Wald test (Methods). These shared heterogeneous associations showed significant departures of $\beta(\tau)$ from a constant effect across quantiles, and the remaining shared associations had no detected heterogeneity.

This analysis identified 13,139,351 qQTL associations across 14 contexts (26 datasets), at 4,307,480 variant-gene pairs across 15,261 genes (Fig. 1c,d). The qQTL map represents each detected cis association as a curve of estimated allelic effects across quantiles.

Among classified qQTL associations, qQTL-only associations accounted for 6.47% and shared heterogeneous associations for 24.54% (Fig. 1e). The two classes together accounted for 31.01%.

Distributional heterogeneity also differed across molecular modalities. The shared heterogeneous fraction was highest in Knight Alzheimer's Disease Research Center (Knight-ADRC) protein QTLs (48.2% of shared associations, Supplementary Fig. 1). This protein pattern points to a modality in which tail effects may be especially visible, although modality comparisons also depend on sample size, normalization, measurement scale, and the abundance distribution of proteins.

At the median ($\tau = 0.5$), 48.2% of qQTL-only associations across the 26 datasets reached $P < 10^{-5}$, compared with 9.7% at $\tau = 0.1$ and 10.2% at $\tau = 0.9$. The smallest pooled P values occurred in the upper tail (Fig. 1f). The smallest P values in the pooled analysis came mainly from the MIT excitatory neuron (Exc) dataset. The quantile bin with the smallest P value varied across datasets (Supplementary Fig. 2). Central bins showed the strongest overall association evidence in 25 of 26 datasets, but the smallest P values fell in the upper tail bin in 13 datasets, the lower tail bin in 8 and a central bin in 5. Even within excitatory neurons, they fell in the lower tail bin for the Mega dataset and in the upper tail bin for the MIT dataset. Thus, the additional associations identified by qQTL mapping receive widespread support from central quantiles, alongside highly significant tail signals whose prominence varies across datasets.

We next examined whether tail significance was accompanied by larger effect magnitudes. Among associations significant at both the median ($\tau = 0.5$) and at least one endpoint ($\tau = 0.1$ or 0.9; quantile-specific $P < 10^{-5}$), the absolute endpoint effect exceeded the absolute effect at $\tau = 0.5$ in 76.5% of shared heterogeneous, 63.6% of qQTL-only and 57.9% of shared associations without detected heterogeneity, and the endpoint effect was a median of 1.23, 1.10 and 1.03 times the effect at $\tau = 0.5$, respectively (Fig. 1g). When both endpoints were significant, the endpoint with the larger absolute effect was used, and no minimum median effect size was imposed beyond significance at $\tau = 0.5$. These percentages and ratios describe the selected associations with significant effects at both the median and an endpoint.

We identified associations missed by mean-effect analysis and characterized the distributional profiles of estimated allelic effects across the aging brain QTL map. We next examined how these signals related to regulatory sequence and trait heritability.

Quantile QTLs identify regulatory variation linked to brain and aging traits

We tested whether qQTL-only variants were enriched in regulatory annotations and whether qQTL annotations were associated with trait heritability. The single-nucleus analyses used Mega data only. We compared qQTL-only variants with linear QTL variants using excess-of-overlap (EOO) enrichment across approximately 1,500 genomic and epigenomic annotations[21–24], with standard errors estimated by leaving one chromosome out at a time.

qQTL-only variants were enriched for regulatory annotations relative to variants detected by mean-effect QTL models. Enrichment was observed in cell-type-specific expression QTL annotations and chromatin interaction loops for bulk expression contexts, and in chromatin accessibility regions and transcription factor binding sites for protein contexts (Fig. 2a-c). These enrichment results support a regulatory interpretation of qQTL-only variants.

We then assessed whether the same qQTL annotations were associated with brain and aging trait heritability. Stratified linkage disequilibrium (LD) score regression (S-LDSC)[25–28] partitioned heritability across 86 GWAS, including 47 brain-related traits and 13 aging or neurodegenerative traits. Marginal models added one QTL annotation to the 97-category baseline model, and joint models added both the qQTL and mean-effect QTL annotations. qQTL annotations showed enrichment comparable to mean-effect QTL annotations in marginal models and were also associated with trait heritability in joint models (Fig. 2d; Supplementary Table 3). Because qQTL and mean-effect QTL annotations are correlated, this result supports a conditional link to trait heritability rather than a claim that qQTL annotations define an independent source of heritability.

In single-nucleus data, qQTL-only variants showed stronger enrichment in regulatory annotations from other cell types than in annotations from the assayed cell type. For example, astrocyte qQTL-only variants showed stronger enrichment in non-astrocyte transcription factor annotations than in astrocyte annotations in the annotation enrichment analysis (Fig. 2c). This cross-cell enrichment is consistent with shared regulatory programs or cell states across annotated cell types, but it can also reflect residual cell mixing, ambient RNA, doublets or imperfect cell state annotation.

qQTL-only variants showed regulatory annotation enrichment, and qQTL annotations were associated with trait heritability after accounting for mean-effect QTL annotations and the baseline model. We next compared the astrocyte quantile profile at *PICALM* with established evidence for microglial regulation to examine the cross-cell pattern at an individual AD locus.

The *PICALM* AD risk allele is associated with a larger reduction at lower astrocyte expression quantiles

At *PICALM*, the AD risk allele was associated with lower expression in both microglia and astrocytes. The reduction varied little across the examined quantiles in microglia but was more pronounced at lower than at upper conditional quantiles in astrocytes. *PICALM* is a well-established late-onset AD risk gene[1], and recent functional work reported[20] that the rs10792832 (chr11:86156833:A>G, GRCh38 reference genome) late-onset AD risk allele G weakens PU.1 binding at a microglial regulatory element, lowers *PICALM* expression and leads to lipid droplet accumulation with reduced uptake of amyloid-β and myelin debris. This microglial result provides the reference point for interpreting any astrocyte distributional effect at the same variant.

Under the qQTL allele coding used here, rs10792832 showed different regulatory effects in microglia and astrocytes. In microglia, the risk allele lowered *PICALM* expression by approximately 0.21 units on the normalized expression scale, with an estimated curve that was approximately flat across $\tau = 0.1$ to 0.9 (qQTL $P = 4.44 \times 10^{-47}$, Fig. 3a). The estimated effect of the same allele was smaller in magnitude in astrocytes than in microglia. Within astrocytes, the estimated effect was more negative at lower than at upper quantiles (qQTL $P = 5.27 \times 10^{-8}$, Fig. 3b). The estimated effect was $\beta = -0.048$ at $\tau = 0.2$, with a 95% confidence interval (CI) from −0.065 to −0.031. The estimate was $\beta = -0.025$ at $\tau = 0.8$ (95% CI, −0.040 to −0.010). An exploratory comparison gave an estimated contrast $\beta(0.8) - \beta(0.2)$ of approximately 0.023 (nominal Wald $P = 0.025$), consistent

with a larger estimated effect magnitude at lower quantiles.

Local regulatory annotation favors a microglial cis effect at the locus rather than an independent focal astrocyte cis effect. Within ±25 kb, the lead variant overlaps a microglial enhancer interval and several elements active in massively parallel reporter assays, and the region is densely covered by microglial and immune enhancers. The nearest astrocyte enhancer lies approximately 4 to 5 kb away, with no astrocyte enhancer or promoter overlapping the variant within ±5 kb (Fig. 3c). This pattern favors a direct microglial cis regulatory effect, although the absence of a focal astrocyte annotation in this window does not exclude astrocyte cis regulation. We therefore considered whether a downstream or shared relationship with microglial regulation could explain the stronger estimated effect at lower than at upper astrocyte expression quantiles.

Cross-cell analyses linked genotype-associated *PICALM* expression in microglia and astrocytes. In the forward decomposition, from genotype through microglial *PICALM* expression to astrocyte expression, the estimated indirect effect was −0.052 (95% confidence interval, −0.082 to −0.023), about 45% of the total effect. In the reverse analysis, through astrocyte expression to microglial expression, the estimated indirect effect was smaller in magnitude (−0.019, 95% confidence interval, −0.031 to −0.007) and was about 4% of the total effect (Fig. 3d, Supplementary Table 4). The forward and reverse estimates describe associations with different expression outcomes: astrocyte expression in the forward analysis and microglial expression in the reverse analysis. Their asymmetry supports a downstream or shared relationship between the two cell types, while leaving causality unresolved.

We used a simulation to ask whether a genotype effect on mean microglial expression could produce a larger effect magnitude at lower quantiles of astrocyte expression. In the simulation, a homogeneous genotype effect on microglial *PICALM* expression was coupled to a downstream relationship from microglia to astrocytes that changed distributional shape rather than only the mean. A pure location shift produced a downstream signal detectable by a mean-effect model. Adding a distributional component produced both shared QTL and qQTL-only associations in astrocytes, with effects concentrated in the lower tail and an average $\beta(\tau)$ curve matching the observed direction (Fig. 3e,f). This simulation illustrates a possible downstream relationship and does not establish mediation at *PICALM*.

PICALM illustrates how quantile profiles extend the interpretation of an established AD locus. Alongside the mean-effect evidence for microglial cis regulation, quantile analysis showed a larger estimated effect magnitude at lower than at upper conditional quantiles within the adjusted astrocyte expression distribution. Local annotation and cross-cell analyses support a downstream or shared interpretation of this profile, while the observational evidence leaves causality and mechanism unresolved. Other loci illustrate modifier associations and changes in molecular dispersion.

Modifier effects and expression variability at qQTL loci

We used *TMEM106B* and *PITRM1* to illustrate how interaction and variance QTL analyses complement qQTL profiles. *TMEM106B* shows an *APOE* ε4 modifier association in excitatory neurons,

whereas *PITRM1* shows a dispersion effect supported by qQTL and vQTL analyses without AD GWAS support in these data. A qQTL curve localizes the largest allelic effect within the adjusted molecular distribution. AD interpretation then depends on genetic support, local regulatory annotation, colocalization or fine-mapping, modifier tests and cross-cell evidence.

At *TMEM106B*[29], the genotype-by-*APOE* ε4 interaction in excitatory neurons was negative (β = −0.239, 95% CI, −0.34 to −0.14, P = 1.28 × 10^{-6}), consistent with a more negative *TMEM106B* genotype effect at higher ε4 dosage (Fig. 4g). In the targeted *APOE* ε4 interaction analysis, dosage was coded as 0, 1 or 2 copies (n = 541, 182 and 8, respectively), giving N = 731 after exclusion of six donors with missing dosage. The interaction remained significant after excluding the eight ε4 homozygotes, under carrier coding, and in permutation and bootstrap analyses. *TMEM106B* also had moderate AD GWAS support (P = 1.9 × 10^{-3}).

The estimated quantile profile at this locus suggested a more negative allelic effect at lower quantiles, but the contrast weakened in sensitivity analyses. In excitatory neurons, estimated allelic effects were about −0.042 to −0.040 at τ = 0.1 to 0.2 and approached zero at τ = 0.9 (−0.0009; Fig. 4f). The P value for the τ = 0.1 versus τ = 0.9 contrast increased from 0.047 to 0.23 after 1%/99% winsorization and to 0.29 after removing the 1% most influential donors, and the 95% confidence interval from 500 donor bootstrap samples included zero. The modifier interpretation therefore rests on the *APOE* ε4 interaction evidence, whereas the quantile analyses do not support a robust lower-tail interpretation.

PITRM1[30] illustrates a dispersion effect in anterior cingulate cortex (AC). The mean association for chr10:3147188:C>T (β = 0.108, P = 0.013) did not meet the matched discovery threshold. The variant showed strong support in both the variance and quantile tests (vQTL P = 3.97 × 10^{-14}, qQTL P = 1.72 × 10^{-7}; Fig. 4b,c). Its β(τ) changed sign across the distribution, positive in the lower tail and negative in the upper tail, with little effect at the median. Thus, an additional copy of the coded allele reduces the distance between the lower and upper conditional quantiles, indicating a change in dispersion that a mean-effect scan is poorly suited to represent. The variant did not show AD GWAS support in the datasets examined.

Of the variance QTLs (vQTLs) detected in the overlap comparison, 93.95% were also qQTLs (Fig. 4a). The detected vQTLs tended to involve both tails. Conversely, many qQTL effects confined to one tail did not meet the vQTL definition and would be missed by a variance test.

Beyond the *PICALM* example, *TMEM106B* supports an *APOE* ε4 modifier association despite uncertainty in its quantile contrast, whereas *PITRM1* shows molecular dispersion without AD GWAS support in these data. We next examined what quantile-based genetic predictors add to AD gene prioritization using qTWAS.

qTWAS prioritizes AD candidate genes across molecular contexts

Across four sets of AD GWAS summary statistics, qTWAS identified 357 distinct gene-dataset pairs representing 207 genes (Fig. 5a,b; per GWAS in Supplementary Fig. 6). Of these, 44 were qTWAS-only genes, defined as genes without a significant association in matched standard

transcriptome-wide association analysis (TWAS). We found at least suggestive standard TWAS support for 87.9% of all AD qTWAS genes ($P < 10^{-4}$; Supplementary Data 1, Supplementary Data 2) and for 19 of the 44 qTWAS-only genes, none of which passed the standard TWAS Bonferroni threshold. These results refine the prioritization of genes with suggestive standard TWAS support and nominate additional candidates for genetic and functional studies.

Quantile-aware transcriptome-wide association analysis (qTWAS)[31] links genetically predicted molecular traits to AD GWAS. We trained qTWAS weights for 115,063 gene-dataset pairs across 14 molecular contexts (20 datasets, Supplementary Table 2) and applied them to 4 AD GWAS summary-statistic sets from three meta-analyses (Bellenguez full and European Alzheimer & Dementia Biobank (EADB), Wightman, and Kunkle)[1–3]. Within each molecular trait dataset and GWAS, qTWAS used one Cauchy-combined P value per tested molecular trait[32], with Bonferroni correction over those traits. Standard TWAS[33,34] selected the minimum P value across prediction methods with results for each eligible gene and used a Bonferroni denominator equal to the number of tested molecular traits multiplied by the number of methods with results for that gene. Significant results were pooled across the four GWAS sets by taking their union. Genes detected only by qTWAS are treated as candidates rather than standalone evidence of distributional regulation.

Outside loci with significant associations, genomic inflation factors ranged from 0.97 to 1.06 across the four AD GWAS sets and from 0.90 to 1.20 across molecular datasets (Supplementary Fig. 9). In null GWAS simulations run through the qTWAS association test with fitted weights held fixed, the Cauchy-combined test had a genomic inflation factor of 0.88. The observed values were higher than this null reference, potentially reflecting sub-threshold polygenic signal. At a nominal threshold of $P < 10^{-4}$, the empirical rejection rate was 1.03×10^{-4}. The Bonferroni family-wise error rate per GWAS and dataset averaged 0.050, corresponding to an expected four GWAS-dataset analyses with at least one false-positive result across the 80 analyses under this null setting (Supplementary Fig. 10a,b). Weight permutation was evaluated for 19 qTWAS-only genes whose models contained more than six variants; 14 had a permutation $P < 0.05$ (Supplementary Fig. 10e).

qTWAS nominated 44 AD candidate genes beyond those detected by matched standard TWAS and characterized their association support across quantile bands. The unadjusted P values for their 62 gene-dataset-GWAS combinations are shown in Fig. 5c. In 44 of the 45 combinations with a standard TWAS P value, the qTWAS P value was smaller (median ratio about 940), and the standard TWAS P value was below 0.05 in 39 of them. Most qTWAS-only associations therefore had nominal support from standard TWAS but reached Bonferroni significance only with qTWAS.

These 44 qTWAS-only genes were defined using method-specific Bonferroni thresholds. At common cutoffs on the unadjusted P value of 10^{-7}, 10^{-6}, 10^{-5}, 10^{-4} and 10^{-3}, qTWAS identified 11, 16, 28, 101 and 472 AD genes, respectively, that were not detected by standard TWAS at the same cutoff (Fig. 5d). qTWAS therefore identified additional AD candidate genes at each common cutoff, while the size of this set depended on the threshold.

Association support for qTWAS-only genes defined by the primary Bonferroni analysis generally

extended across multiple quantile bands, and most genes remained significant after excluding extreme quantiles from the weights. The P values for the lower (τ = 0.01 to 0.33), middle (τ = 0.34 to 0.66) and upper (τ = 0.67 to 0.99) quantile bands were all below 10^{-4} in 32 of the 62 significant gene-dataset-GWAS combinations. Effect directions agreed across the three bands in 60 of the 62 combinations. After removing quantiles $\tau \leq 0.05$ and $\tau \geq 0.95$ from the band weights, 40 of the 44 qTWAS-only genes remained significant, and the four genes lost were near the significance threshold (Supplementary Fig. 10c).

Individual qTWAS-only associations also differed in which quantile band provided the strongest support. In 18 combinations, the smallest of the three band-level P values was at most one hundredth of the second-smallest. The smallest P value occurred in the lower band for 10 combinations, the middle band for 4 and the upper band for 4. For the four upper-band combinations, the corresponding P values in the middle band were 410 to about 2.5 million times those in the upper band. These combinations were *ARRB2* (Oli CUIMC1, Bellenguez), *SIX5* (monocyte, EADB), *TMEM163* (Mic Mega, Kunkle) and *NICN1* (posterior cingulate cortex, Bellenguez). For *ARRB2* and *TMEM163*, standard TWAS evidence was also weaker by factors of about 360,000 and 440, respectively. These comparisons describe how association support varies across quantile bands and prioritize candidates for follow-up.

At the gene level, replacing the quantile-varying weights with a single τ-invariant predictor on the same variants removed significance for 17 of the 44 qTWAS-only genes (39%), compared with 23 of the 163 qTWAS genes also detected by standard TWAS (14%; Fisher's exact $P = 8.7 \times 10^{-4}$; Supplementary Fig. 10d). Nevertheless, most qTWAS-only genes remained significant with the τ-invariant predictor on the same variants (27 of 44 genes). qTWAS therefore expands AD gene prioritization beyond standard TWAS and describes the quantile-band support for the additional associations. Most additional genes remained significant after excluding extreme quantiles. Although τ-invariant replacement removed significance more often among these genes than among genes also detected by standard TWAS, most remained significant.

We compared the 44 qTWAS-only genes with four external AD gene resources: the FunGen-xQTL AD loci gene tiers (T1–T5)[35,36], the ADSP Gene Verification Committee lists[37], a single-nucleus eQTL meta-analysis with colocalization and Mendelian randomization (SingleBrain)[38] and a single-nucleus TWAS[39] (Supplementary Data 1; Supplementary Fig. 11). Eighteen of the 44 genes were in at least one list. Eleven were supported by the FunGen-xQTL tiers, the ADSP lists or single-nucleus colocalization (*CCNT2*, *EPHA1*, *GPRC5B*, *IGHG1*, *KNOP1*, *LILRB1*, *PPM1E*, *RBCK1*, *TBC1D9B*, *TMEM163* and *TPM1*), and seven more were found only by the single-nucleus TWAS (*ANXA11*, *FGD4*, *IPMK*, *PPM1N*, *RTF2*, *WDR18* and *ZNF233*). The other 26 were not in the lists compared here and remain candidates for follow-up. Sample overlap is expected between the molecular training cohorts and the AD GWAS, and the four GWAS sets also share participants. Individual-level overlap was not quantified from the available summary statistics. Cross-GWAS agreement is therefore interpreted as consistency rather than independent replication.

The AD gene sets identified by qTWAS and standard TWAS were both enriched for pathways involved in lipid transport and immune processes (Fig. 5e)[40,41]. The terms shown were significantly

enriched in both the qTWAS and standard TWAS gene sets and had higher fold enrichment for qTWAS; the AD comparison included all 207 qTWAS genes and 447 standard TWAS genes. In a separate application to telomere length PC1, the same comparison recovered terms for major histocompatibility complex class I antigen presentation and telomere caps (Fig. 5f). These results do not establish pathway discovery specific to qTWAS, which would require enrichment among qTWAS-only or high-confidence genes relative to matched TWAS genes. The AD enrichment therefore supports the biological coherence of the prioritized gene list by recovering pathways already prominent in AD genetics.

The qTWAS candidate set combines molecular context, quantile profiles, matched TWAS results and external genetic evidence to guide regulatory follow-up at AD loci. The qTWAS-only genes comprise three groups: genes supported by the FunGen-xQTL tiers, the ADSP lists or single-nucleus colocalization, genes found only by the single-nucleus TWAS, and genes in none of the compared lists (Supplementary Data 1). These distinctions help prioritize candidate genes for further study according to their molecular context and available AD genetic support.

Discussion

Regulatory variation in the aging brain includes broad changes in molecular abundance, effects that differ across conditional quantiles and changes in expression variability. These patterns provide additional evidence for interpreting AD risk loci alongside mean-effect QTL results.

Across the genome, mean effect QTL mapping detected most qQTL associations. Quantile mapping identified additional associations and showed that allelic effects varied across quantiles in a subset of associations detected by both methods. Enrichment of qQTL-only variants in regulatory annotations and the conditional association of qQTL annotations with brain and aging trait heritability support the biological relevance of the analyzed annotation sets alongside mean-effect QTL annotations. These aggregate analyses do not establish the stability of individual tail-localized effects.

We used quantile profiles to refine the biological interpretation of AD loci. At *PICALM*, functional evidence links the AD risk allele to regulation in microglia. For the same allele, we estimated a larger reduction in *PICALM* expression at lower than at upper conditional quantiles in astrocytes. Local annotation favors a microglial cis effect, and cross-cell analyses link microglial and astrocyte *PICALM* expression. A downstream simulation illustrates how such a relationship could produce larger astrocyte effects at lower quantiles. Together, these findings support a downstream or shared interpretation of the astrocyte association, while leaving the cross-cell mechanism unresolved. The evidence at *TMEM106B* supports an *APOE* ε4 modifier association, whereas the *PITRM1* variant showed regulatory heterogeneity without AD GWAS support in the datasets examined.

qTWAS highlighted AD candidate genes across molecular contexts, including 44 genes not detected by matched standard TWAS. External genetic evidence distinguished candidates with prior

disease support from those requiring further study. Molecular context and quantile-band support characterize the associations, while matched TWAS evidence, consistency across AD GWAS and external AD gene resources help assess their genetic support. Most qTWAS-only genes remained significant with a τ-invariant predictor on the same variants. For genes detected only by qTWAS, evidence that the underlying variant effects differ across conditional quantiles can motivate follow-up on distributional regulation. These results expand the AD candidate gene set and guide regulatory follow-up according to molecular context and the available genetic support.

Allelic effect estimates at extreme conditional quantiles have higher sampling variability. Tail assignments therefore require calibration checks, sensitivity analyses for outliers, and assessment of stability under donor resampling. Tail labels based on estimated quantile effects are summaries of fitted associations, not direct observations of a cell state or subgroup. Annotation enrichment and S-LDSC analyses depend on well-defined variant sets, appropriate LD reference data and correlated annotation modeling. These checks define which qQTL effects are ready for biological follow-up.

Mechanistic interpretation remains a separate step. A lower tail effect can reflect variance effects, measured modifiers, latent cell states, downstream regulatory propagation or technical covariation. Annotation, targeted interaction tests, cross-cell decomposition and simulation can rank these explanations, but perturbation or other functional data are needed to establish mechanism. For qTWAS candidates, colocalization or credible set overlap can support a link to an AD locus. Functional evidence is needed to establish a candidate gene's biological role.

The aging brain qQTL map brings quantile effects together with molecular context and external genetic evidence to help interpret AD risk loci and prioritize candidate genes for studies of gene regulation.

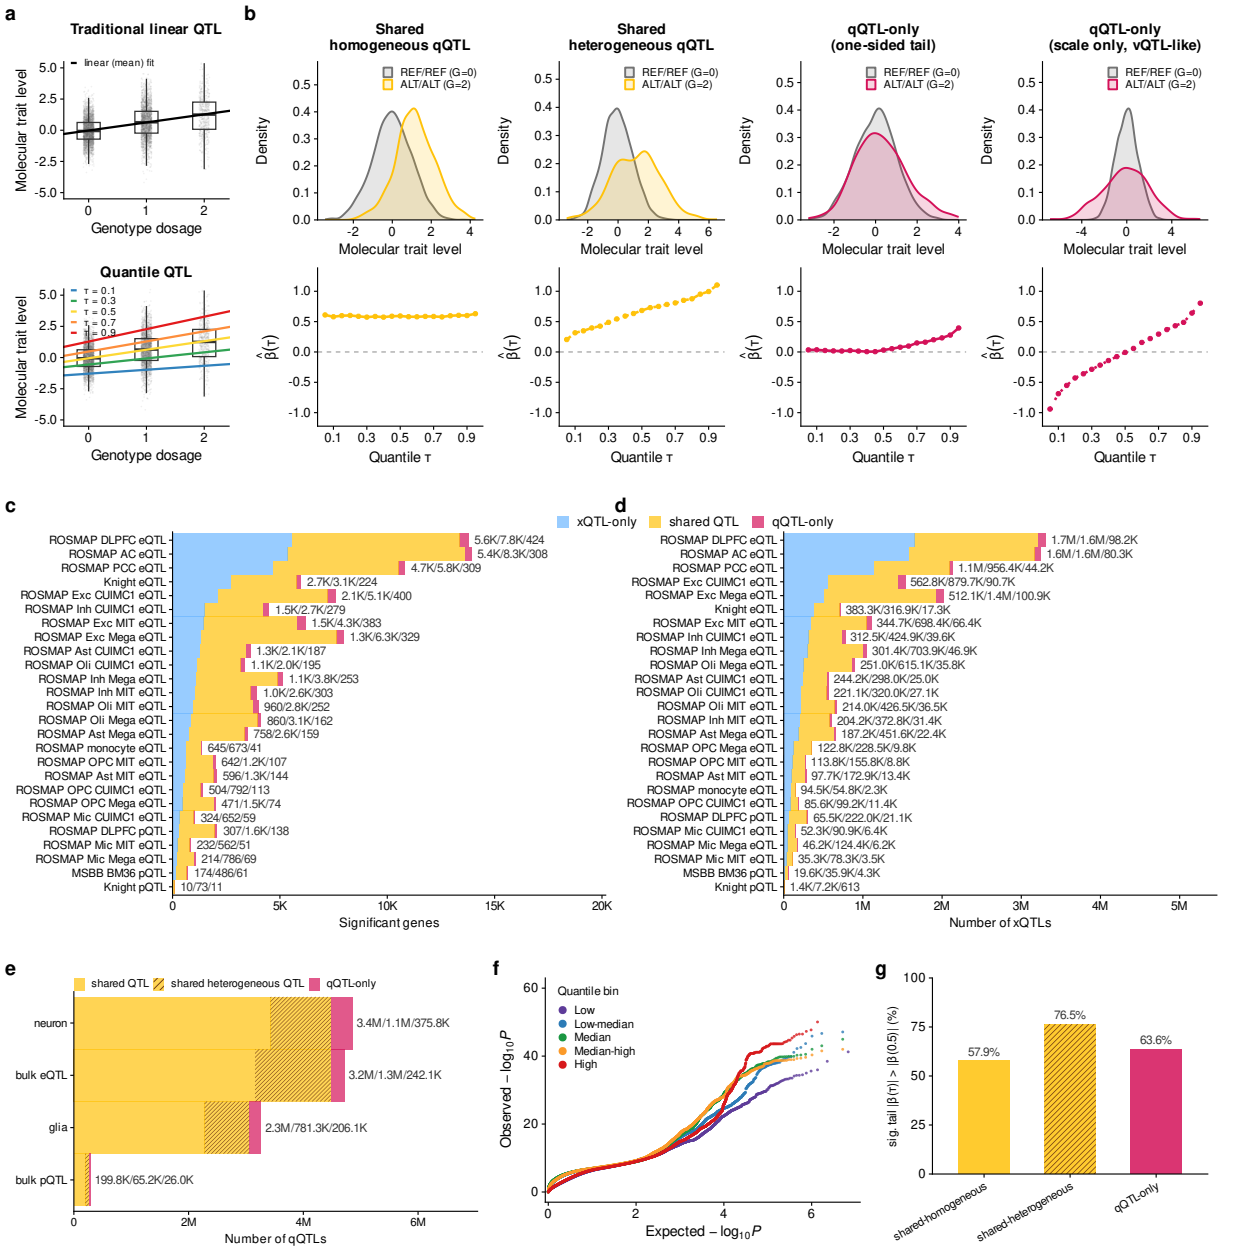

Figure 1. qQTL mapping represents detected cis associations as effect curves across quantiles, classifies associations as qQTL-only, shared without detected heterogeneity or shared heterogeneous, and shows that many heterogeneous effects are largest in lower or upper tails. a,b, Simulated examples illustrating mean-effect and quantile QTL models. **a,** Linear regression estimates a single effect β, whereas quantile regression estimates $\beta(\tau)$ at different quantiles. **b,** Trait densities for REF/REF and ALT/ALT genotypes (top) and quantile-specific effect estimates (bottom) illustrate shared homogeneous, shared heterogeneous, and two qQTL-only patterns: one-sided tail and scale-only effects. **c,d,** Counts of significant genes (**c**) and significant variant–gene pairs (**d**) for each dataset, stratified into qQTL-only (red), shared by qQTL and linear xQTL mapping (gold), and xQTL-only (blue). Labels give the xQTL-only/shared/qQTL-only counts. **e,** qQTL heterogeneity category by context grouped by neuron, bulk eQTL, glia, and bulk pQTL. qQTL-only associations are shown in red. Shared qQTL associations are shown in yellow, with hatched yellow indicating shared heterogeneous QTLs and solid yellow indicating shared homogeneous QTLs. Labels give the shared homogeneous/shared heterogeneous/qQTL-only counts. **f,** QQ plots of qQTL-only P values pooled into five quantile bins: Low (τ = 0.10–0.25), Low-median (τ = 0.30–0.40), Median (τ = 0.45–0.55), Median-high (τ = 0.60–0.70), and High (τ = 0.75–0.90). **g,** Percentage of associations for which the absolute estimated tail effect exceeds the absolute estimated median effect, by association class. Included associations have $P < 10^{-5}$ at the median (τ = 0.5) and at least one tail (τ = 0.1 or 0.9); when both tails meet this threshold, the tail with the larger absolute effect is selected.

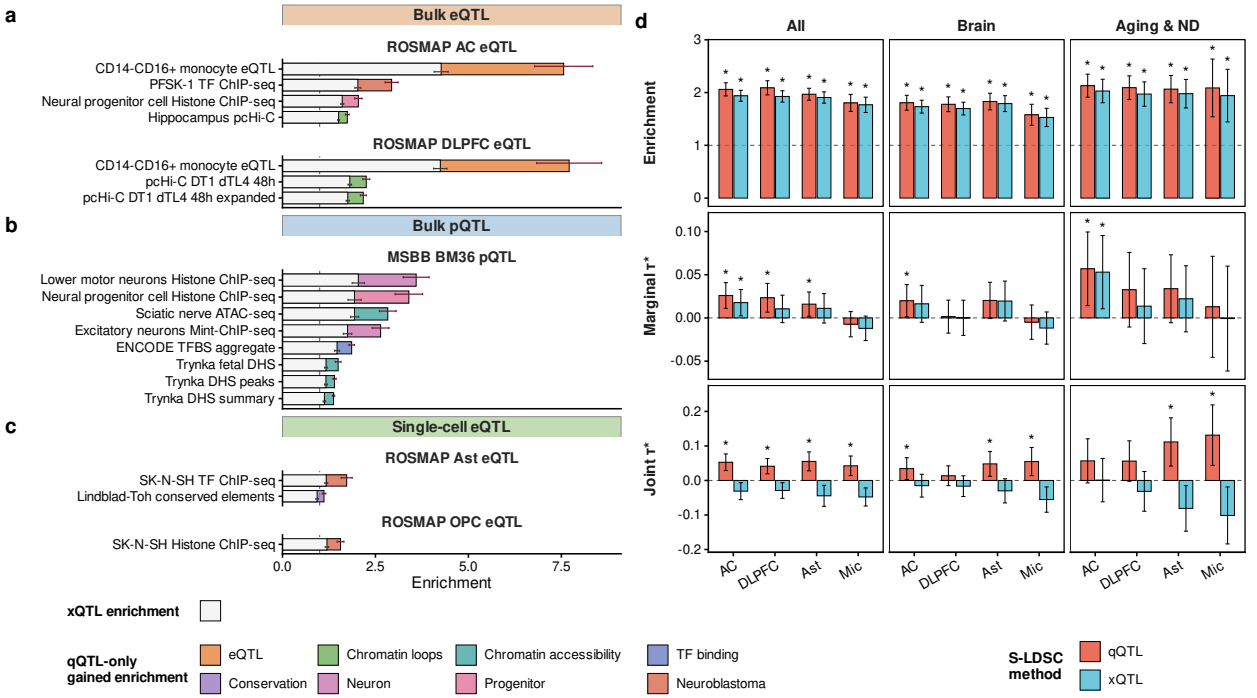

Figure 2. qQTL-only variants show stronger enrichment in selected regulatory annotations than variants detected by mean effect QTL mapping. qQTL annotations are associated with brain and aging trait heritability after accounting for mean effect QTL annotations. Panel d shows random-effects pooled estimates with 95% confidence intervals based on meta-analysis standard errors. The single-nucleus analyses used Mega data only. Cross-cell enrichment can reflect shared regulatory programs or cell states, as well as residual cell mixing, ambient RNA, doublets or imperfect cell state annotation. **a–c,** EOO (excess-of-overlap) enrichment analysis. Functional annotations showing significantly higher enrichment for qQTL-only variants relative to xQTL, analyzed separately for bulk eQTL (**a**), protein QTL (**b**), and single-cell eQTL (**c**). White bars, xQTL enrichment; coloured segments, additional enrichment of qQTL-only variants, coloured by annotation category. In (**c**), cross-cell enrichment differences may reflect residual cell mixing or ambient RNA in single-nucleus data. **d,** S-LDSC enrichment analysis. Comparison of qQTL versus xQTL stratified LD score regression enrichment and standardized effect size (τ^*), jointly modeled across annotations, across major brain regions (anterior cingulate cortex, dorsolateral prefrontal cortex) and brain cell types (astrocytes, microglia). Rows, heritability enrichment, marginal τ^* and joint τ^*; columns, all 86 traits, 47 brain-related traits and 13 aging or neurodegenerative (ND) traits; *$P < 0.05$, shown for positive estimates only (values in Supplementary Table 3).

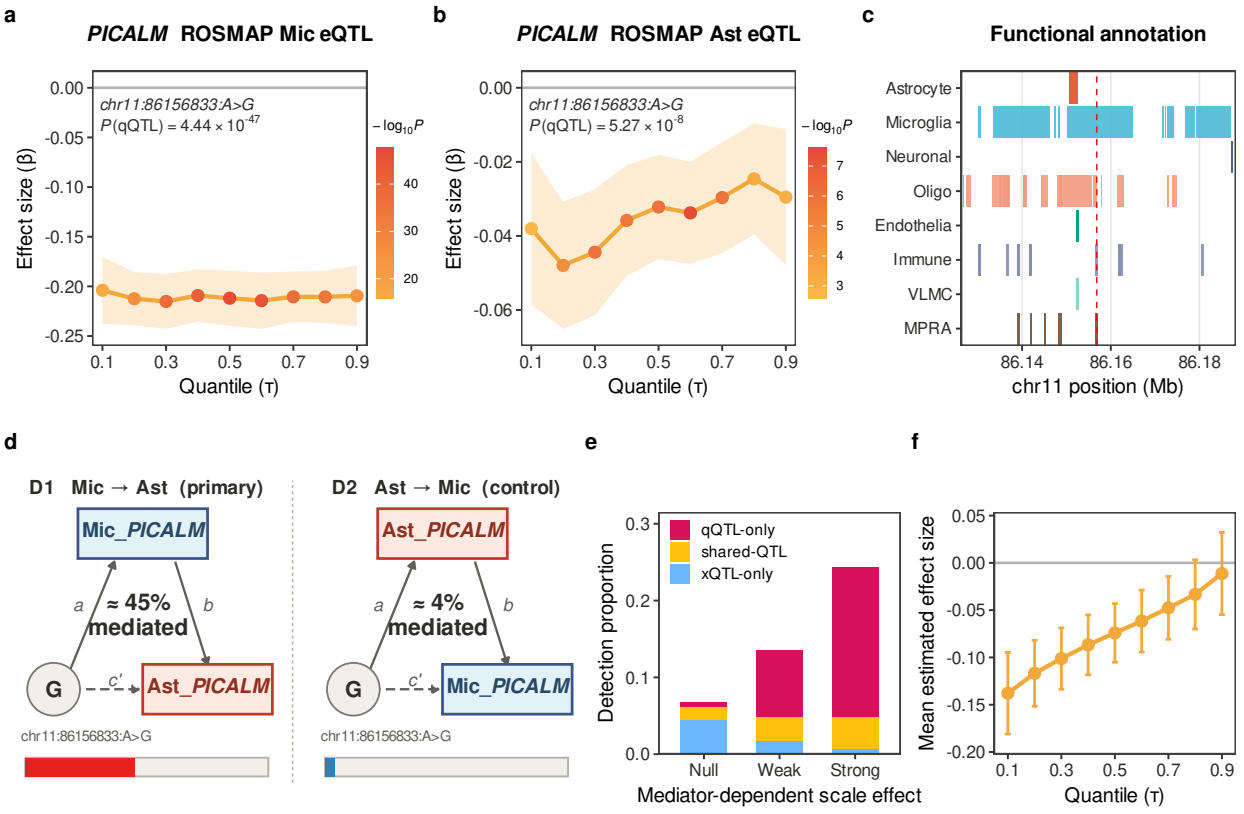

Figure 3. At *PICALM*, prior AD genetics and microglial functional evidence support a microglial cis effect. The estimated astrocyte qQTL profile has larger effect magnitudes at lower quantiles, with an exploratory lower-versus-upper quantile contrast. We examined a possible downstream or shared relationship between microglial *PICALM* regulation and the astrocyte profile using cross-cell decomposition, conditioning and downstream simulation analyses. **a,b,** qQTL of chr11:86156833:A>G on *PICALM* expression distinguishes an approximately homogeneous microglial effect (**a**) from a heterogeneous astrocyte effect (**b**). Points are coloured by $-\log_{10} P$ at each quantile. **c,** Local annotation is microglia/immune-centered and lacks direct focal astrocyte enhancer/promoter overlap. Dashed red line, position of the variant. **d,** Mediation-style directional comparison between G→Mic→Ast and G→Ast→Mic. The estimated indirect effect is larger in magnitude in the microglia-to-astrocyte direction than in the reverse direction. Bars and labels give the proportion mediated (path estimates in Supplementary Table 4). **e,** Mediation simulation under the weaker-signal setting ($a = -0.5$, $\beta_G = -0.1$, $\beta_M = -0.05$, $\gamma_0 = 0.4$, $n = 1000$, 1000 replicates) with null, weak, and strong mediator-dependent scale effect $\gamma_M \in \{0, -0.10, -0.15\}$ (Methods). The bar plot shows the proportion of xQTL-only, qQTL-only, and shared-QTL detections as the mediator-dependent distributional component is varied. **f,** Average estimated quantile-effect curve from the same simulation setting at mediator-dependent scale effect $\gamma_M = -0.1$. The curve shows negative effects across quantiles, with stronger negative effects at lower quantiles and weaker effects toward higher quantiles, qualitatively matching the astrocyte *PICALM* qQTL pattern.

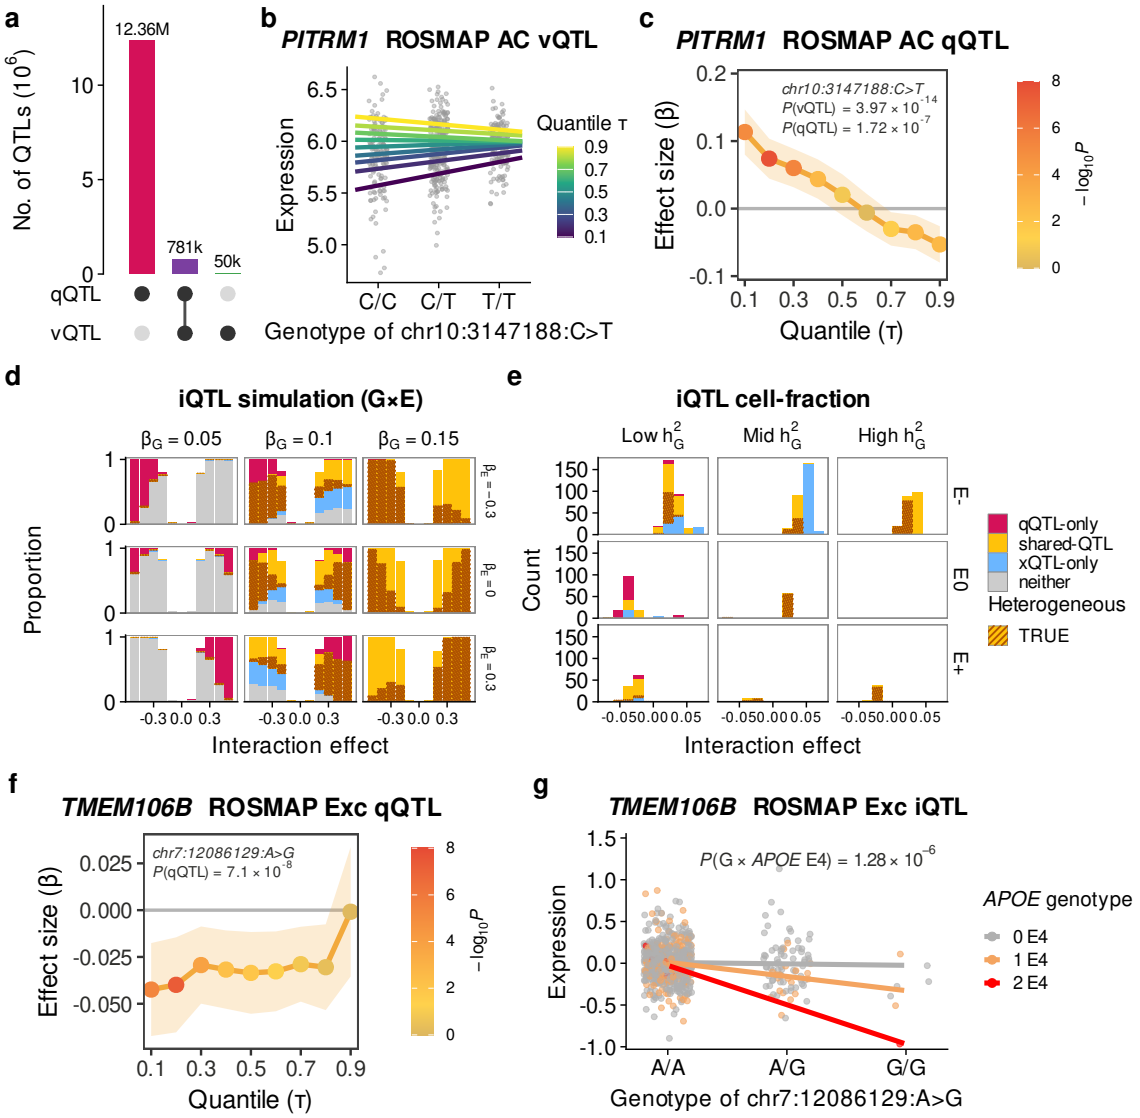

Figure 4. Analyses at selected loci illustrate different uses of qQTL profiles. Of the detected vQTLs in the overlap comparison, 93.95% were also qQTLs (**a**). *TMEM106B* shows an *APOE* ε4 modifier association, whereas the *PITRM1* variant showed qQTL and vQTL regulatory heterogeneity but did not show AD GWAS support in the datasets examined. AD interpretation depends on genetic or functional evidence rather than qQTL shape alone. **a–c,** vQTLs and their correspondence with qQTLs. **a,** Overlap between vQTL and qQTL. **b,c,** *PITRM1* expression in anterior cingulate cortex for variant chr10:3147188:C>T. **b,** Quantile regression fits across quantiles (τ = 0.1–0.9) stratified by genotype, showing differences in expression patterns across the distribution. **c,** Quantile regression coefficient profiles illustrating variation in effect sizes across τ. **d,e,** iQTLs captured by qQTL and linear interaction analyses. **d,** Simulation analysis of iQTL capture by qQTL and xQTL. Expression traits were simulated under a genotype-by-environment interaction model with $n = 5{,}000$. The simulated environmental factor E was generated independently of genotype as a binary factor and centered before constructing the interaction term. Columns correspond to fixed marginal genetic effects β_G, rows correspond to fixed environmental main effects β_E, and the x axis shows the interaction effect β_I. Bars show the proportion of simulated iQTLs classified as qQTL-only, xQTL-only, shared-QTL or neither at $P < 1 \times 10^{-6}$. Hatching denotes significant quantile-effect heterogeneity at $P_{\text{het}} < 0.05$. Only positive β_G settings are shown because the capture pattern was symmetric with respect to the sign of β_G. **e,** Capture of cell-fraction iQTLs by qQTL and xQTL in ROSMAP single cell data. Histograms show the estimated interaction effects β_I for cell-fraction iQTLs detected by qQTL, xQTL or both. iQTLs not detected by either method were omitted for visualization. Columns stratify cell-fraction iQTLs by marginal genetic variance explained, h^2_G: Low $h^2_G \leq 5\%$, Mid $5\% < h^2_G \leq 25\%$, and High $h^2_G > 25\%$. Rows stratify iQTLs by the sign class of the cell-fraction main effect, E−, E0 or E+. Interaction-effect estimates outside the plotted range were collapsed into the edge bins. Hatching marks shared qQTL/xQTL signals with significant quantile-effect heterogeneity ($P < 0.05$). **f,g,** *TMEM106B* expression in excitatory neurons for variant chr7:12086129:A>G. **f,** Quantile regression coefficient profiles across τ. **g,** *TMEM106B* expression by variant genotype and *APOE* ε4 dosage.

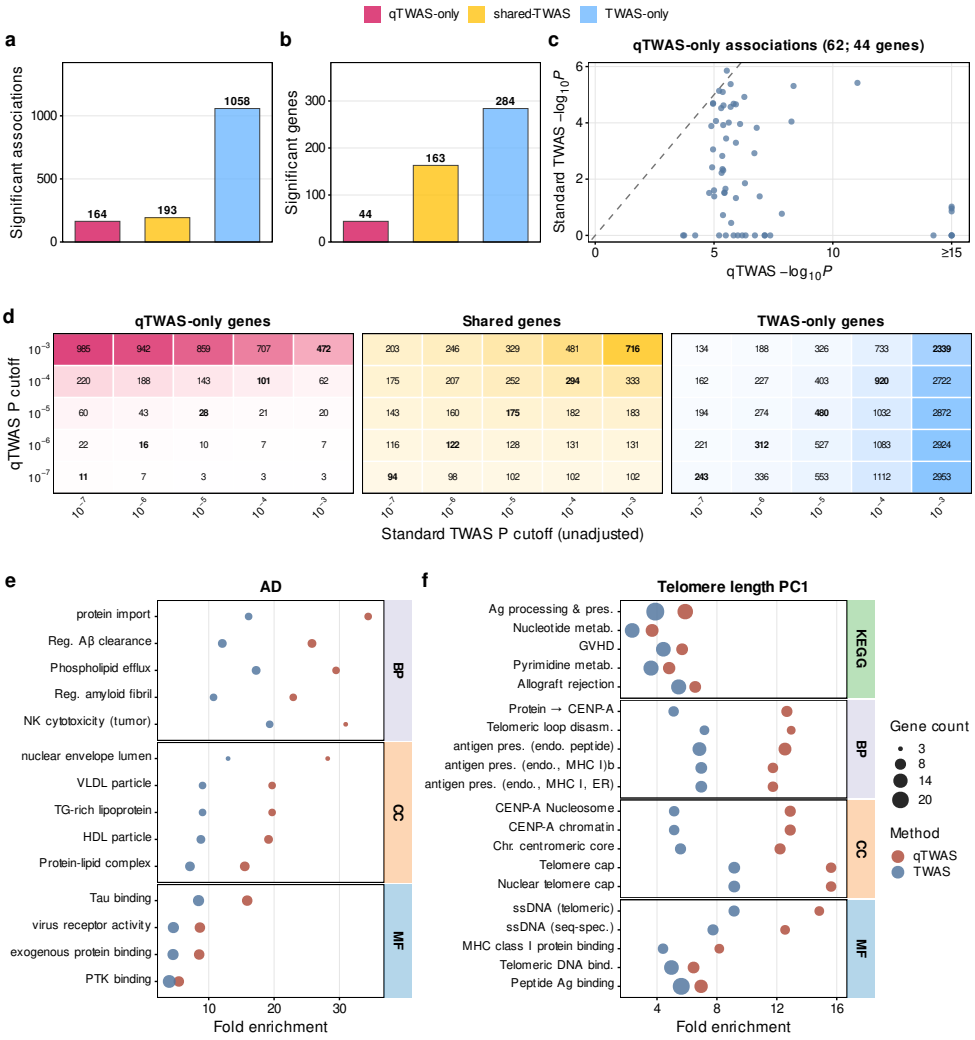

Figure 5. Most AD candidate genes identified by qTWAS had at least suggestive support from matched TWAS. We took the union of significant results across four Alzheimer's disease (AD) GWAS sets (Bellenguez 2022, EADB 2022, Kunkle 2019 and Wightman 2021). GWAS summary statistics for telomere length PC1 were analyzed separately. Significance was defined as Bonferroni-adjusted $P < 0.05$ across the molecular traits tested in each dataset (qTWAS, Cauchy-combined P; standard TWAS, minimum P across prediction methods with results for each eligible gene, with that method count also in the Bonferroni factor). Colocalization and credible-set evidence for the qTWAS-only candidates is in Supplementary Data 1. **a,b,** Significant gene-dataset pairs (**a**) and unique genes (**b**) from the union of the four AD GWAS sets, classified as qTWAS-only (red), shared by qTWAS and standard TWAS (gold) and TWAS-only (blue). **c,** Standard TWAS against qTWAS $-\log_{10} P$ (unadjusted) for the 62 qTWAS-only gene-dataset-GWAS combinations (44 genes); dashed line, equal P. Unlike **a**, which classifies each gene–dataset pair separately (164 qTWAS-only pairs), **c** includes only the 44 genes that are qTWAS-only at the gene level, that is, without a significant standard TWAS association in any dataset or AD GWAS set. All 62 combinations have Bonferroni-adjusted qTWAS $P < 0.05$. All combinations with an eligible standard TWAS result have adjusted standard TWAS $P \geq 0.05$. The 17 combinations without an eligible standard TWAS result are drawn at their assigned value of $P = 1$; qTWAS P values below 10^{-15} are drawn at the axis limit. **d,** Numbers of qTWAS-only, shared and TWAS-only AD genes when a cutoff on the unadjusted P is applied to each method and varied from 10^{-7} to 10^{-3}; a gene passes a cutoff if its smallest P across datasets and AD GWAS sets is below it. Bold numbers, the same cutoff for both methods. Unadjusted P is used because the two methods have different Bonferroni factors: the thresholds in **a–c** correspond to an unadjusted P of $0.05/N$ for qTWAS (median 9.8×10^{-6} across datasets and AD GWAS sets) and $0.05/(NK_g)$ for standard TWAS (median 1.5×10^{-6}; strictest 5.7×10^{-7}, hence the lowest cutoff of 10^{-7}). Here N is the number of molecular traits tested in the dataset for that GWAS, and K_g is the number of prediction methods with results for gene g. **e,f,** GO and KEGG terms enriched in both the qTWAS and the standard TWAS gene sets for AD (**e**) and telomere length PC1 (**f**). Terms shown are significant in both enrichments (Benjamini–Hochberg FDR < 0.05) and have higher fold enrichment for qTWAS than for standard TWAS; for each category, up to five terms with the largest difference (qTWAS minus TWAS) are shown. Input lists are all significant genes of each method (207 qTWAS and 447 standard TWAS genes for AD). Point size, number of genes; colour, method. GO categories: Biological Process (BP), Cellular Component (CC) and Molecular Function (MF); KEGG, curated pathway sets.

METHODS

Ethics statement

All human data used in this study are de-identified and were obtained from publicly available controlled-access resources under approved Data Use Agreements. All contributing studies obtained informed consent from participants and were approved by their respective institutional review boards (IRBs). Our analyses constitute secondary analysis of de-identified human data and are therefore exempt from additional IRB review.

Molecular QTL analysis setting and multiple testing

We first define the common molecular QTL testing setup. The molecular QTL and GWAS data sources are given in Section Molecular QTL and GWAS data sources. Consider a study with n independent individuals, indexed by $i = 1, \ldots, n$. Let Y_i denote the value of a molecular quantitative trait for individual i, such as the expression level of a gene. For a gene region containing m genetic variants, let G_{ij} denote the genotype of variant j for individual i, coded additively as the number of alternate (coded) alleles, so that $G_{ij} \in \{0, 1, 2\}$. Let $C_i = (C_{i1}, \ldots, C_{iq})^\top$ be a vector of q covariates for individual i. For each QTL type, let θ_j denote the genetic effect of variant j on a target feature of the conditional distribution of Y_i given genotype and covariates. We test $H_0 : \theta_j = 0$. Specifically, for linear QTL (xQTL), $\theta_j = \beta_j$ is the effect on the conditional mean; for quantile QTL (qQTL), θ_j collects the effects $\beta_j(\tau)$ across conditional quantiles; for variance QTL (vQTL), θ_j is the quantile-integrated effect on conditional variability; and for interaction QTL (iQTL), $\theta_j = \beta_I$ is the genotype-by-factor interaction coefficient. Detailed model specifications for qQTL, vQTL, and iQTL are provided below.

For *cis*-QTL identification, we restrict testing to variants within a TAD-boundary-enhanced *cis* window: the gene's start and end positions are extended by 1 Mb and combined with the generalized topologically associating domain and boundary region (TAD-B) containing the gene, retaining the outermost boundaries, and regions longer than the 99th percentile of the genome-wide region length (about 20 Mb) are winsorized[21,42,43]. We apply a common testing procedure consisting of variant-level quality control and hierarchical multiple-testing correction[44].

Quality control. We filter genetic variants by minor allele frequency (MAF > 0.05) to retain sufficiently common variants for stable inference in xQTL, qQTL, vQTL, and iQTL analyses. Missing genotypes are imputed using the mean dosage, and monomorphic variants (zero variance) are excluded[45].

Multiple-testing correction. We apply a two-stage hierarchical procedure following Huang et al.[46]. *Stage 1 (gene discovery).* For each gene g, let $\tilde{P}_{g,j}$ denote the Bonferroni-adjusted P value for variant j within the *cis* region of gene g, and let $\tilde{P}^*_g$ denote the Bonferroni-adjusted P value of the lead *cis* variant for gene g. We apply the Benjamini-Hochberg procedure to $\{\tilde{P}^*_g\}$ across all genes to control the gene-level false discovery rate (FDR) at 5%. *Stage 2 (variant discovery within significant genes).* Among genes passing Stage 1, we define a data-adaptive variant-level

significance threshold as $P^\star = \max_g \tilde{P}^*_g$, where the maximum is taken over the set of significant genes from Stage 1. We declare a variant-gene pair significant if $\tilde{P}_{g,j} < P^\star$. This hierarchical procedure controls the gene-level FDR at 5% while restricting variant-level discoveries to genes identified in Stage 1 using the Bonferroni-adjusted threshold $P^\star$ (see **Supplementary Note** Hierarchical multiple-testing correction for QTL discovery for details).

Quantile rank-score test for quantile QTL

We apply the quantile rank-score test[47] at multiple quantile levels to assess whether genotype is associated with conditional quantiles of the molecular trait and combine the resulting quantile-specific P values using the Cauchy combination method[16,32]. This combined test is sensitive to genetic effects on any part of the conditional distribution, even when they are not well summarized by a conditional mean model.

For a quantile level $\tau \in (0, 1)$, we model the conditional τth quantile of Y_i given genotype G_{ij} and covariates C_i via the linear quantile regression model[48]

$$Q_{Y_i}(\tau \mid G_{ij}, C_i) = \alpha_0(\tau) + C_i^\top \alpha(\tau) + G_{ij}\beta_j(\tau), \tag{1.1}$$

where $\alpha_0(\tau)$ is a quantile-specific intercept, $\alpha(\tau)$ is a vector of quantile-specific covariate effects, and $\beta_j(\tau)$ is the quantile-specific genetic effect of variant j. For fixed τ, we test $H_0 : \beta_j(\tau) = 0$ using a rank-score test [18,47].

The rank-score statistic at quantile level τ is

$$S_{j,\tau} = n^{-1/2} \sum_{i=1}^{n} G^*_{ij} \phi_\tau \left\{ Y_i - \widehat{\alpha}_0(\tau) - C_i^\top \widehat{\alpha}(\tau) \right\}, \tag{1.2}$$

where $\phi_\tau(u) = \tau - \mathbb{I}(u < 0)$ and $\mathbb{I}(\cdot)$ is the indicator function. Here, $(\widehat{\alpha}_0(\tau), \widehat{\alpha}(\tau))$ are estimated under H_0, so the null model is fitted once per gene and quantile level. The residualized genotype $G^*_j = (G^*_{1j}, \ldots, G^*_{nj})^\top$ is defined as

$$G^*_j = \left(I_n - \tilde{C}(\tilde{C}^\top \tilde{C})^{-1} \tilde{C}^\top \right) G_j,$$

where $G_j = (G_{1j}, \ldots, G_{nj})^\top$, $C =(C_1, \ldots, C_n)^\top$, and $\tilde{C} = [\mathbf{1}_n, C]$ includes an intercept.

Under H_0, the standardized statistic $S_{j,\tau}/\sigma_{j,\tau}$ is asymptotically standard normal, where

$$\sigma^2_{j,\tau} = \tau(1 - \tau) \frac{G^{*\top}_j G^*_j}{n}.$$

Let $P_{j,\tau}$ denote the corresponding P value at quantile level τ. For a prespecified grid of quantile levels $\mathcal{T} = \{\tau_1, \ldots, \tau_K\}$, we combine $\{P_{j,\tau} : \tau \in \mathcal{T}\}$ using the Cauchy combination method[32] to obtain a single P value for each variant. We then identify significant qQTL using the two-stage hierarchical multiple-testing correction described in Section Molecular QTL analysis setting and

multiple testing. The qTWAS analyses use quantile-regression prediction weights estimated after the variant filtering and harmonization steps described below.

Quantile TWAS for disease association testing

We use quantile-aware transcriptome-wide association analysis (qTWAS) to test whether genetic predictors based on different quantile bands of a molecular trait are associated with disease. qTWAS combines quantile-regression prediction weights with GWAS summary statistics. In contrast to standard TWAS, which uses weights for conditional mean expression, qTWAS estimates separate weights over quantile regions. We perform the following steps for each gene and molecular context; the gene index is suppressed.

Quantile weight estimation. For each context-gene pair, we first combine the quantile-specific rank-score P values of each variant across 19 quantile levels using the Cauchy combination method. We then apply the Benjamini-Hochberg procedure across variants within the gene and retain variants with FDR < 0.05 as the candidate set for qTWAS weight estimation. Candidate variants are drawn from the same TAD-B *cis* region as in QTL mapping (Section Molecular QTL analysis setting and multiple testing). These candidate variants are harmonized with an external Alzheimer's Disease Sequencing Project (ADSP) LD reference panel for TWAS, which includes 16,571 whole-genome sequenced individuals of European ancestry[49,50]; see Section Molecular QTL and GWAS data sources. Harmonization resolves minor-allele coding and strand discrepancies. We retain harmonized variants with MAF > 0.01 in the ADSP LD reference panel, and then apply hierarchical-clustering-based LD pruning at an R^2 threshold of 0.8.

Let p be the number of variants retained after harmonization and LD pruning. For individual i, let $G_{\mathsf{f},i} = (G_{\mathsf{f},i1}, \ldots, G_{\mathsf{f},ip})^\top$ denote the resulting vector of filtered genotypes. We estimate quantile-specific prediction weights by fitting a quantile regression model with all filtered variants at each of 99 equally spaced quantile levels $\tau \in \{0.01, 0.02, \ldots, 0.99\}$:

$$Q_{Y_i}(\tau \mid G_{\mathsf{f},i}, C_i) = \gamma_0(\tau) + C_i^\top \gamma(\tau) + G_{\mathsf{f},i}^\top \omega(\tau). \tag{1.3}$$

Here, $\gamma_0(\tau)$ is a quantile-specific intercept, $\gamma(\tau)$ denotes quantile-specific covariate effects, and $\omega(\tau) = (\omega_1(\tau), \ldots, \omega_p(\tau))^\top$ is the quantile-specific genetic effect vector. Let $\rho_\tau(u) = u(\tau - \mathbb{I}(u < 0))$ be the quantile check loss, where $\mathbb{I}(\cdot)$ is the indicator function. For each τ, we estimate $(\gamma_0(\tau), \gamma(\tau), \omega(\tau))$ by solving

$$(\widehat{\gamma}_0(\tau), \widehat{\gamma}(\tau), \widehat{\omega}(\tau)) = \arg\min_{\gamma_0(\tau), \gamma(\tau), \omega(\tau)} \frac{1}{n} \sum_{i=1}^{n} \rho_\tau\Big(Y_i - G_{\mathsf{f},i}^\top \omega(\tau) - C_i^\top \gamma(\tau) - \gamma_0(\tau)\Big). \tag{1.4}$$

The estimated quantile-specific weights $\widehat{\omega}(\tau)$ summarize the contribution of retained variants to predicted expression at each quantile, including effects concentrated in the lower tail or upper tail, as well as effects that vary more gradually across quantiles.

Quantile weight integration. To summarize effects over broader quantile ranges, we con-

struct weights for three fixed quantile regions. We assign the 99 quantile levels to three fixed intervals representing the lower, middle, and upper portions of the molecular trait distribution: $[0.01, 0.33], [0.34, 0.66], [0.67, 0.99]$.

A data-adaptive partition of the quantile range is described in Supplementary Note Quantile weight integration.

We define the integrated weight for region k as

$$\widehat{\omega}_k = \int_{\mathcal{A}_k} \widehat{\omega}(\tau) d\tau \approx \sum_l \tfrac{1}{2}\big(\widehat{\omega}(\tau_l) + \widehat{\omega}(\tau_{l+1})\big)(\tau_{l+1} - \tau_l), \tag{1.5}$$

where the sum runs over consecutive τ grid points in $\mathcal{A}_k$ and $\widehat{\omega}$ is padded with 0 at one grid step (0.01) beyond each end of the region. This trapezoidal rule approximates the signed integral of the quantile-specific weight function over each region. Because each increment is weighted by its corresponding $\Delta\tau$, the expression represents an integral over τ rather than a raw unweighted sum across grid points. Although the integrated weight vector can scale with region length, the region-level qTWAS statistic below is invariant to positive scalar rescaling of $\widehat{\omega}_k$. For each band, we integrate the signed weights across quantiles to obtain a single weight vector. We refer to the three fixed regions as A1, A2 and A3.

qTWAS association testing. For each quantile region k, we compute the qTWAS test statistic[31]

$$Z_k = \frac{\widehat{\omega}_k^\top z_{\text{GWAS}}}{\sqrt{\widehat{\omega}_k^\top \Sigma \widehat{\omega}_k}}, \tag{1.6}$$

where z_{GWAS} is the vector of GWAS z-scores for the p variants in the gene region, and Σ is the corresponding LD matrix estimated from the same ADSP reference panel used for harmonization. Under the null hypothesis that the genetically predicted expression in region k is not associated with disease, Z_k is asymptotically standard normal[31]. We combine the region-level P values across $k = 1, \ldots, K$ using the Cauchy combination method[32] to obtain a single gene-level P value. Within each molecular dataset and GWAS, we declare significance using a Bonferroni correction over the N molecular traits tested, which controls the family-wise error rate at 0.05 within that dataset and GWAS[51]. The same N is used for qTWAS and standard TWAS. Significant results are pooled across the four AD GWAS summary-statistic sets (Bellenguez full and EADB, Kunkle, Wightman) by taking their union, without a further correction for the number of GWAS-dataset combinations.

Comparison with standard TWAS. For each molecular dataset and GWAS, we compare qTWAS with standard TWAS on the same molecular traits. Standard TWAS results came from the FunGen-xQTL TWAS procedure[44,52]. Depending on the dataset and gene, results were available for three to eight prediction methods, drawn from SuSiE[53], mr.ash[54], LASSO[55], elastic net[56], BayesR[57], Bayesian LASSO[58] and, when fits across multiple contexts were available, mr.mash[59] and mvSuSiE[60]. A gene is eligible if its best model in cross-validation has $R^2 \geq 0.01$ and a cross-validation $P < 0.05$. For each eligible gene g, we take the minimum TWAS P value across all K_g prediction methods with results for that gene, regardless of whether each method meets the cross-

validation thresholds individually. We multiply this minimum P value by NK_g for the Bonferroni correction, where N is the number of molecular traits tested in that dataset and GWAS. We assign $P = 1$ for comparison to genes that fail the gene-level cross-validation screen or lack a standard TWAS result. A gene-dataset association is qTWAS-only, shared or TWAS-only according to which test is significant in at least one of the four AD GWAS, and genes are classified in the same way across datasets. Because the two tests have different Bonferroni denominators, we also compare them at common cutoffs on the unadjusted P value (10^{-7} to 10^{-3}, one per order of magnitude; qTWAS, Cauchy-combined P; standard TWAS, minimum P across prediction methods), classifying each AD gene by its smallest P value across datasets and AD GWAS (Fig. 5d). qTWAS-only combinations without an eligible standard TWAS result are drawn at their assigned value of $P = 1$ (Fig. 5c). A gene has suggestive standard TWAS support if its smallest standard TWAS P value across datasets and AD GWAS is below 10^{-4}.

Calibration of the qTWAS test. We compute the genomic inflation factor λ_{GC} of the qTWAS P values for each AD GWAS and molecular dataset, excluding gene-dataset pairs in LD blocks that contain a Bonferroni-significant qTWAS or standard TWAS association (Supplementary Fig. 9). To evaluate null calibration conditional on the fitted qTWAS weights, we simulate GWAS z-scores $z_0 \sim N(0, \Sigma_b)$ independently in each LD block b, using the LD matrix Σ_b from the ADSP R4 EUR reference panel used for association testing. We generate 100 replicates per block and use the same null draws for the weight sets corresponding to all four AD GWAS. The fitted qTWAS weights, tested variants, three quantile bands and Bonferroni denominator N are held fixed. We replace only the GWAS z-scores and apply the same region-level statistics, Cauchy combination and Bonferroni correction to all AD gene-dataset tests (Supplementary Fig. 10a,b).

Robustness of qTWAS-only associations. We recompute each qTWAS-only association after changing the weights, keeping the GWAS z-scores, LD and Bonferroni denominator fixed. We remove $\tau \leq 0.05$ and $\tau \geq 0.95$ from A1 and A3; replace the three region weights by one τ-invariant weight, the integral of $\widehat{\omega}(\tau)$ over all 99 quantile levels; and remove weight variants with GWAS $P < 5 \times 10^{-8}$, which applies only to genes with such a variant in their weights. For the τ-invariant predictor, we integrate the fitted weights $\widehat{\omega}(\tau)$ over the 99 quantile levels from 0.01 to 0.99 using the trapezoid rule, yielding one weight vector on the same variants without refitting. We substitute this vector for $\widehat{\omega}_k$ in Equation (1.6) and calculate a two-sided P value from the standard normal distribution. We apply the same Bonferroni correction over N tested molecular traits, without Cauchy combination because there is only one predictor. Finally, we permute the weights across variants[33], jointly for A1–A3, recompute the Cauchy-combined P value and define the permutation P value as the proportion of permutations with a P value at least as small as the observed one (10^4 permutations, extended to 10^6 when fewer than 10 exceed the observed value). We restrict the permutation test to models with more than six variants. Models with six or fewer variants have at most $6! = 720$ distinct permutations, too few to estimate the permutation P value precisely. We compare the fraction of genes that lose significance with the τ-invariant weight between qTWAS-only genes and genes also detected by standard TWAS using Fisher's exact test. We report each test separately and do not combine them into a robust set (Supplementary Fig. 10c–e). Supplementary Data 1 lists the genes that have an association passing trimming, lead-variant removal

and permutation (when it can be done), but that lose significance with the τ-invariant weight.

External genetic evidence for qTWAS-only genes. We compare qTWAS-only genes with four external AD gene resources. (i) The FunGen-xQTL AD loci gene tiers (data release of 2 November 2026)[35,36]. AD candidate loci combine 95% credible sets (CSs) from SuSiE-rRSS fine-mapping of each AD GWAS with 95% colocalization confidence sets (CoSs) from multi-GWAS or AD-xQTL colocalization; each contributing set has purity > 0.8, and overlapping sets across datasets and analyses are counted as one locus. xQTL CSs come from individual-level single-context SuSiE, regional fSuSiE and multi-context or multi-gene mvSuSiE fits, and each piece of evidence is assigned to the corresponding gene. An AD GWAS 95% CS and an xQTL 95% CS overlap if they share at least one variant. CoS support requires the AD GWAS and a cis molecular trait of the gene to be in the same ColocBoost event with a 95% CoS. Evidence is matched to the same gene and genetic signal. Tiers T1–T5 rely on localized cis-xQTL support, and each gene takes the first tier whose rule it meets: T1, 95% CS overlap and CoS support, with Mendelian randomization (MR) or causal TWAS (cTWAS) evidence; T2, 95% CS overlap or CoS support, with MR or cTWAS evidence; T3, 95% CS overlap or CoS support, with TWAS evidence; T4, CoS support and a corresponding xQTL CS; T5, other localized evidence, including CoS support alone or overlap between an AD GWAS 95% CS and an xQTL CS at 95% or 70% coverage. Gene-level TWAS, MR and cTWAS (including multi-context cTWAS) evidence follows the FunGen-xQTL regulome-wide association analysis[52]. Genes implicated only by trans-xQTL are reported for reference, because pleiotropic distal effects make locus-to-gene assignment ambiguous. We use tiers T1–T5 (353 genes across the 226 candidate loci). (ii) The ADSP Gene Verification Committee lists[37]: the TopGenes adjudication (best tier 1–2 or 3–7 across evaluations; tiers 1–2 are the GVC high-confidence loci with sufficient evidence for association) and the earlier NIAGADS list of 76 loci named by reported or closest gene and 20 adjudicated causal genes. (iii) SingleBrain[38] genes with colocalization (PP4 > 0.8) with any of the four AD GWAS, with or without Mendelian randomization support. (iv) Single-nucleus TWAS[39] genes with FDR ≤ 0.05 for AD (Bellenguez GWAS) in at least one cellular population. A gene is counted as supported by a resource if it appears in that list; absence from all lists is not a claim of novelty. We also record whether the gene has a SuSiE 95% credible set in the qTWAS dataset (ROSMAP eQTL only)[53]. Per-gene results, including the robustness checks, are given in Supplementary Data 1, and all significant AD qTWAS associations with raw and Bonferroni-adjusted P values in Supplementary Data 2.

Comparison of qQTL with linear QTL

We compare qQTL with linear QTL (xQTL), defined here as *cis* associations detected by mean-based linear regression as implemented in TensorQTL[61]. We identify xQTL and qQTL using the same *cis* window definition and the same set of tested variants, and we perform marginal *cis* association testing for both analyses. For xQTL discovery, we fit the standard linear model in Equation (S.1). For qQTL discovery, we apply the quantile rank-score test under the model in Equation (1.1), evaluating association across 19 equally spaced quantile levels $\tau \in \{0.05, 0.10, \ldots, 0.95\}$. Statistical significance for both xQTL and qQTL is determined using the

same hierarchical multiple-testing correction procedure (Section Molecular QTL analysis setting and multiple testing).

At the variant level, associations are classified into three categories:

- *qQTL-only*: significant in the qQTL analysis but not in the xQTL analysis,
- *xQTL-only*: significant in the xQTL analysis but not in the qQTL analysis,
- *shared QTL*: significant in both the qQTL and xQTL analyses.

At the gene level, we define qQTL genes and xQTL genes as genes with at least one significant qQTL association or xQTL association, respectively.

To further characterize *shared QTL*, we assess whether their genetic effects vary across quantile levels using a trend-based Wald heterogeneity test. For each variant j, we test the null hypothesis $H_0 : \beta_j(\tau_1) = \beta_j(\tau_2) = \cdots = \beta_j(\tau_K)$ over $K = 17$ quantile levels ($\tau = 0.10, 0.15, \ldots, 0.90$), where the genetic effect is constant across quantiles. To describe how $\beta_j(\tau)$ may depart from a constant, we consider three prespecified trend functions: $\Phi^{-1}(\tau)$ (Gaussian), $-\log(1-\tau) - 1$ (Lehmann) and $\log \tau + 1$ (inverse Lehmann). We fit the quantile regression in Equation (1.1) at these quantile levels to obtain $\widehat{\beta}_j(\tau_1), \ldots, \widehat{\beta}_j(\tau_K)$ and estimate their joint covariance. For each trend function, we use the estimated effects and their covariance to construct a Wald test targeting that trend shape, designed to achieve optimal power when $\beta_j(\tau)$ varies along it (Supplementary Note Trend-based Wald heterogeneity test). We then combine the three P values using the Cauchy combination method[32] to maintain high power across diverse patterns of heterogeneity. A shared QTL is labeled *shared heterogeneous* if the combined P value is below 0.05 (nominal).

Shared heterogeneous qQTL indicate detected quantile-dependent variation in the marginal genetic effect. The remaining shared QTL are classified as *shared without detected heterogeneity*; this category corresponds to the shared homogeneous category in summary tables. For descriptive summaries, *heterogeneous qQTL* (qQTL-het) denotes qQTL-only associations together with shared heterogeneous associations. This is a reporting set rather than a formal shape class for qQTL-only associations.

To quantify effect heterogeneity across quantiles, we compute a heterogeneity index $\log |\mathrm{sd}\{\widehat{\beta}_j(\tau)\}/\mathrm{mean}\{\widehat{\beta}_j(\tau)\}|$, where the standard deviation and mean are computed across the 19 quantile levels $\tau = 0.05, 0.10, \ldots, 0.95$. Larger values indicate greater variation in the genetic effect across quantiles relative to its mean, particularly when $\widehat{\beta}_j(\tau)$ changes sign and the mean effect is close to zero. We compare the distribution of this index among qQTL-only associations and shared QTL with and without detected heterogeneity, separately for bulk eQTL, protein QTL and cell-type-specific eQTL analyses (Supplementary Fig. 3).

Comparison of qQTL with variance and interaction QTL

We examined the relationship between qQTL and two major classes of nonlinear genetic effects, variance QTL (vQTL) and interaction QTL (iQTL), to understand how quantile-based association

signals relate to variance modulation and context-dependent regulation.

Variance QTL mapping and comparison with qQTL

We identify vQTL using the quantile integral linear model (QUAIL)[19], which aggregates information from multiple quantiles by integrating the difference $\beta_j(\tau) - \beta_j(1-\tau)$ between upper- and lower-quantile effects over $\tau \in (0.5, 1)$. The integral departs from zero when variant j shifts the upper and lower tails in opposite directions, which changes the conditional variance even when the conditional mean is unchanged. In contrast, qQTL analysis combines P values across individual quantiles to capture general distributional effects rather than variance alone. We therefore compare vQTL and qQTL across variant–molecular trait pairs within each context analyzed by both scans to evaluate their overlap. Following the QUAIL implementation, we obtain a P value for each variant using 19 quantile levels ($\tau = 0.05, 0.10, \ldots, 0.95$) and adjusting for covariates. The significance of vQTL is determined using the hierarchical multiple-testing procedure described in Section Molecular QTL analysis setting and multiple testing.

Overlap between interaction QTL and quantile QTL

iQTL mapping. Interaction QTLs (iQTLs) capture gene-gene (G×G) or gene-environment (G×E) effects through an explicit interaction between a genetic variant and an interacting factor[12,62]. Supplementary Note Interaction effects can generate marginal qQTL signals describes how unmodeled G×G and G×E interactions can generate marginal quantile shifts. We mapped iQTLs by fitting the linear model

$$Y_i = \beta_0 + \beta_G G_{ij} + \beta_E E_i + \beta_I (G_{ij} E_i) + C_i^\top \beta_C + \epsilon_i, \tag{1.7}$$

where Y_i is the molecular phenotype for individual i, G_{ij} is the additive genotype dosage at variant j, E_i is an interacting factor, and C_i denotes additional covariates. The coefficient β_G represents the genetic main effect at the mean of the interacting factor, β_E represents the main effect of the interacting factor, and β_I represents the genotype-by-factor interaction effect. A variant-gene pair was declared an iQTL if the interaction coefficient β_I was statistically significant.

We considered two classes of interacting factors: discrete factors included sex and *APOE* $\varepsilon 4$ dosage across all datasets, and continuous factors included estimated proportions of six major brain cell types. Before iQTL mapping, all interacting factors were mean-centered so that β_G represents the genetic effect at the mean of the interacting factor. Response variables were normalized using rank-based inverse normal transformation (INT) to reduce sensitivity to outliers, and variants with MAF ≤ 0.05 were excluded.

iQTL significance was determined using the hierarchical multiple-testing procedure described in Section Molecular QTL analysis setting and multiple testing, with gene-level FDR controlled at 0.05 for continuous factors and 0.25 for discrete factors. The more permissive threshold for discrete factors follows prior iQTL practice[12], which reported limited discovery of physiological-trait interactions relative to cell-type interactions, and was adopted to improve statistical power. Within

each gene, the Bonferroni step of this procedure used the effective number of independent variants estimated by eigenMT[63] instead of the number of tested variants.

Simulation comparing iQTL capture by marginal qQTL and xQTL scans. To evaluate how interaction effects are captured by marginal qQTL and xQTL scans, we simulated outcomes from Equation (1.7) using $n = 5{,}000$ individuals. Genotypes were generated as additive dosages $\{0, 1, 2\}$ under Hardy-Weinberg equilibrium with MAF $= 0.1$. The interacting factor was generated independently of genotype as a binary variable with event probability 0.5 and was mean-centered before constructing the interaction term. We varied the genetic main effect $\beta_G \in \{0.05, 0.10, 0.15\}$, the environmental main effect $\beta_E \in \{-0.3, 0, 0.3\}$, and the interaction effect $\beta_I \in \{-0.5, -0.4, \ldots, 0.5\}$. The covariates C_i comprised 45 independent standard normal variables, and errors were generated as $\epsilon_i \overset{\text{iid}}{\sim} N(0, 0.267)$; the covariate coefficients β_C and the error variance were estimated from a real snRNA expression phenotype (*TTN* in excitatory neurons). For each parameter setting, we generated 1,000 Monte Carlo replicates. In each replicate, we first applied INT to the simulated outcomes and then analyzed the transformed outcomes using iQTL, marginal qQTL, and marginal xQTL scans. Significance was assessed using $\alpha = 10^{-6}$, a fixed threshold chosen to be slightly more stringent than the gene-specific eigenMT thresholds observed in the real-data analysis, which ranged from approximately 6.1×10^{-6} to 3.8×10^{-5}.

For visualization and interpretation of the real-data results, we further stratified detected cell-fraction iQTLs by the direction of the cell-fraction main effect and the strength of the marginal genetic contribution. Specifically, iQTLs were labeled E0 if $P_E \geq 0.05$, E+ if $P_E < 0.05$ and $\beta_E > 0$, and E- if $P_E < 0.05$ and $\beta_E < 0$, where P_E is the nominal P value for the cell-fraction main effect in Equation (1.7). The strength of the marginal genetic contribution was quantified by the single-variant variance contribution

$$h^2_G = \hat{\beta}^2_G 2\text{MAF}_j(1 - \text{MAF}_j)/s^2_Y,$$

where $\hat{\beta}_G$ is the estimated genetic main effect from Equation (1.7), MAF_j is the minor-allele frequency of variant j, and s^2_Y is the sample variance of the phenotype used in the iQTL scan. Associations were then classified into low, intermediate and high marginal genetic effect groups using $h^2_G \leq 0.05$ (5%), $0.05 < h^2_G \leq 0.25$ (5% to 25%) and $h^2_G > 0.25$ (> 25%), respectively.

PICALM mediation analysis and downstream heterogeneity simulation

We used mediation analysis to assess whether microglial *PICALM* expression explained part of the association between genotype and astrocyte *PICALM* expression[64]. For individual i, let G_i denote the dosage of the alternate G allele at rs10792832 (chr11:86156833:A>G, GRCh38), identified as the AD risk-increasing allele in the fine-mapping table of Bellenguez et al.[1]. Among the 737 ROSMAP individuals included in the mediation analysis, the genotype counts for G-allele dosages of 0, 1, and 2 were 88, 348, and 301, respectively, corresponding to a G-allele frequency of 0.644 and a minor allele frequency of 0.356.

We residualized INT-normalized *PICALM* expression in each cell type against its corresponding

covariate matrix and denoted the resulting microglial and astrocyte residuals by M_i and A_i, respectively. Genotype and astrocyte expression were available for 737 ROSMAP individuals, with microglial expression available for 733. For the primary direction $G \to M \to A$, we jointly fitted the following two-equation mediation model using *lavaan*[65] and full information maximum likelihood (`missing = "fiml"`, `fixed.x = FALSE`) to accommodate missing expression data[66]:

$$M_i = \alpha_M + aG_i + \epsilon_{M,i},$$
$$A_i = \alpha_A + bM_i + c'G_i + \epsilon_{A,i}.$$

Here, a denotes the association from genotype to microglial expression, b denotes the association from microglial to astrocyte expression conditional on genotype, and c' denotes the remaining genotype-astrocyte association after conditioning on microglial expression. We defined the indirect effect as ab, the total effect as $c' + ab$, and the proportion mediated as $ab/(c' + ab)$. The proportion mediated was used as a descriptive summary and was interpreted together with the indirect effect and its confidence interval. Primary standard errors and confidence intervals were computed using the delta method. A nonparametric bootstrap[67] and a complete-case sensitivity analysis reproduced the mediation estimates qualitatively (Supplementary Note Robustness of the *PICALM* mediation analysis). To compare the two directions, we repeated the analysis for the reverse path $G \to A \to M$ with astrocyte expression as the mediator and microglial expression as the outcome.

Downstream heterogeneity simulation for qQTL effects

We performed a downstream heterogeneity simulation to illustrate how a homogeneous genotype effect on a mediator can generate a heterogeneous downstream qQTL. Specifically, we considered the following data-generating model:

$$M_i = aG_i + \gamma_0\epsilon_{M,i},$$
$$Y_i = \beta_G G_i + \beta_M M_i + (\gamma_0 + \gamma_M M_i)\epsilon_{Y,i},$$

where $\epsilon_{M,i}$ and $\epsilon_{Y,i}$ are independent standard normal random variables. In this model, a controls the homogeneous genotype-to-mediator effect, and γ_M controls the mediator-dependent scale effect. To illustrate the source of quantile heterogeneity, a simple heuristic calculation replacing M_i by $E(M_i \mid G_i) = aG_i$ gives

$$\beta(\tau) \approx \beta_G + a\beta_M + a\gamma_M\Phi^{-1}(\tau),$$

where $\Phi^{-1}(\tau)$ is the standard normal quantile function. Thus, when $\gamma_M \neq 0$, the mediator-dependent scale term makes $\beta(\tau)$ vary with τ, thereby inducing a heterogeneous marginal qQTL.

For the displayed simulation, genotype dosage $G_i \in \{0, 1, 2\}$ was generated under Hardy-Weinberg equilibrium with MAF $= 0.2$. We fixed $a = -0.5$, $\beta_G = -0.1$, $\beta_M = -0.05$, and $\gamma_0 = 0.4$, and varied $\gamma_M \in \{0, -0.1, -0.15\}$ to represent null, weak, and strong mediator-dependent scale effects. For each value of γ_M, we generated $1{,}000$ simulated datasets with $n = 1{,}000$ individuals,

applied the same marginal xQTL/qQTL classification procedure, and summarized the resulting detection proportions at $\alpha = 10^{-6}$.

Molecular QTL and GWAS data sources

Molecular QTL from FunGen-xQTL atlas

We analyze molecular quantitative trait loci (*mol*QTL) from the Alzheimer's Disease Sequencing Project (ADSP) Functional Genomics Consortium (FunGen-AD), which aggregates aging and Alzheimer's disease cohorts across brain regions, cell types, and molecular modalities[43,44]. The FunGen-xQTL atlas provides harmonized genotype data and molecular phenotypes processed for cross-cohort comparability. Molecular phenotypes are aligned, quantified, and normalized using shared procedures. Genotype data undergo sample- and variant-level quality control, ancestry inference, relatedness filtering, imputation, and harmonization to GRCh38, dbSNP 156, and ENSEMBL v103. For the qQTL analysis, we use the molecular contexts across cohorts summarized in Supplementary Table 1. The primary analysis comprises 26 datasets spanning 14 molecular contexts. Bulk data come from ROSMAP[7,9,68], the Knight Alzheimer's Disease Research Center[8,69] and the Mount Sinai Brain Bank[6,70]. Single-nucleus data for six ROSMAP cell types come from three analysis sets: CUIMC1[5,71], MIT[72,73] and Mega, a combined analysis of CUIMC1, MIT and a further ROSMAP single-nucleus dataset (CUIMC2)[74] in which each donor is included once. We counted each donor once across datasets and processing sources. The primary analysis includes 1,515 unique donors (901 ROSMAP, 430 Knight Alzheimer's Disease Research Center and 184 Mount Sinai Brain Bank (MSBB) proteomics). MSBB bulk eQTL datasets are reported in the Supplement only (Supplementary Figs. 4 and 5). qTWAS weights are trained for 14 molecular contexts (20 datasets, using the Mega and CUIMC1 single-nucleus data; Supplementary Table 2).

All *mol*QTL analyses use a common covariate structure that includes fixed covariates, genotype principal components, and phenotype-derived hidden factors. Genotype PCs explaining 70% of the genotype variance are retained, and hidden factors are obtained by principal component analysis of the molecular phenotypes, with the number of factors determined automatically[75]. We compute *mol*QTL using the same univariate *cis* association model implemented in TensorQTL[61], yielding harmonized eQTL and pQTL summary statistics that are directly comparable across tissues and cell types.

GWAS summary statistics for Alzheimer's disease and aging

We use publicly available GWAS summary statistics from European-ancestry samples for Alzheimer's disease (AD) from three large case-control meta-analyses that represent complementary study designs and sample compositions: (1) the primary AD meta-analysis (111,326 cases and 677,663 controls) and the European Alzheimer & Dementia Biobank (EADB) subset (20,301 cases and 21,839 controls) from Bellenguez et al.[1]; (2) a meta-analysis that includes both clinical AD cases and proxy AD cases (90,338 total cases, including 46,613 proxy cases, and 1,036,225 controls) from Wightman et al.[2]; and (3) a clinical and autopsy-confirmed AD

meta-analysis (35,274 cases and 59,163 controls) from Kunkle et al.[3]. These datasets provide AD genetic association summary statistics across complementary ascertainment strategies.

We also consider leukocyte telomere length, a molecular marker of cellular aging. The GWAS summary statistics come from a UK Biobank analysis of $462{,}666$ participants of European ancestry. The GWAS phenotype is telomere length PC1, the first principal component of combined qPCR- and TelSeq-derived measurements[76].

For qTWAS analyses, we construct a custom LD reference panel from the Alzheimer's Disease Sequencing Project (ADSP), consisting of 16,571 whole-genome sequenced individuals of European ancestry from ADSP Release 4[50].

Functional annotation and disease heritability enrichment

EOO enrichment analysis of functional annotations

To quantify functional differences between two sets of significant variants, we perform an excess-of-overlap (EOO) enrichment analysis[77]. EOO evaluates whether a target set of variants is over-represented in a functional annotation relative to the background prevalence of that annotation in the same genomic context.

For a functional annotation A and a target set of significant variants B, let $a = |A \cap B|$ denote the number of variants in B that overlap A, and let $|A|$ and $|B|$ denote the sizes of A and B, respectively. We define the EOO statistic as

$$\mathrm{EOO}(A, B) = \frac{a/|B|}{|A|/T}, \tag{1.8}$$

where T is the total number of tested variants in the same genomic context. Values greater than 1 indicate enrichment of A among the target variants, whereas values less than 1 indicate lower overlap than expected under the genomic-context background. We estimate standard errors by leaving out one chromosome at a time, and we summarize enrichment within each molecular context.

We applied EOO to a panel of approximately 1,500 genomic annotations grouped into 21 functional categories. The panel included annotations curated from FILER[22] along with baselineLD v2.2 annotations[25,28], STARR-seq annotations[23], ENCODE[24], Roadmap Epigenomics[78], chromatin profiling of human neural stem cells[79], human iPSC-derived cortical neuron ATAC-seq (GEO accession GSE70823), and progesterone-receptor and related transcription-factor ChIP-seq studies in decidualized human endometrial stromal cells[80,81]. For each molecular context, we defined two variant sets from marginal *cis* association results after hierarchical multiple-testing control (hierarchical FDR < 0.05): a qQTL-only set (significant in the qQTL analysis but not in the xQTL analysis) and an xQTL set (significant in the xQTL analysis). The tested variants in the same genomic context serve as the background for each EOO calculation. To test whether enrichment differed between the qQTL-only and xQTL sets within the same context, we performed a two-sided z test comparing the EOO estimates and controlled the false discovery rate across annotations using the Benjamini-Hochberg procedure[82] at 0.05. For each context, we report annotations with significant differences

and highlight the five annotations showing the largest enrichment differences between qQTL-only and xQTL sets.

S-LDSC heritability analysis of qQTL and xQTL annotations

To evaluate the contribution of qQTL and xQTL signals to complex trait heritability, we construct binary SNP annotations from marginal *cis* association results after hierarchical FDR control at 0.05. For each molecular context, we define two annotations: (1) a *qQTL* annotation, indicating SNPs that are themselves significant in at least one qQTL variant-gene association, and (2) an *xQTL* annotation, indicating SNPs that are themselves significant in at least one xQTL variant-gene association.

We apply stratified LD score regression (S-LDSC)[27] jointly with the 97-category baselineLD v2.2 model[25,28], using approximately 8 million autosomal SNPs and LD scores computed from the ADSP reference panel. For each annotation a, S-LDSC estimates the per-SNP heritability coefficient τ_a from the regression

$$E[\chi_j^2] = N \sum_a \tau_a \ell_{j,a} + N a_0 + 1,$$

where N is the GWAS sample size and $\ell_{j,a}$ is the LD score of SNP j with respect to annotation a, and a_0 measures confounding such as population stratification. For each annotation and molecular context, we report two summary measures. First, we report the marginal enrichment

$$\text{Enrichment}_a = \frac{h^2_{g,a}/h^2_g}{M_a/M}, \tag{1.9}$$

where $h^2_{g,a}$ is the heritability explained by SNPs in annotation a, h^2_g is the total SNP heritability, M_a is the number of SNPs in annotation a, and M is the total number of SNPs. Second, we report the standardized effect size

$$\tau^*_a = \frac{\tau_a sd(a)}{\sum_j \text{Var}(\beta_j)/M}, \tag{1.10}$$

where $sd(a)$ is the standard deviation of the annotation across SNPs, β_j denotes the per-variant effect size, and $h^2_g = \sum_j \text{Var}(\beta_j)$ so that $\sum_j \text{Var}(\beta_j)/M$ is the mean per-SNP heritability. The quantity τ^*_a can be interpreted as the proportional change in per-SNP heritability associated with a one-standard-deviation increase in annotation a, conditional on the other annotations in the model.

Marginal models add one QTL annotation to the 97-category baseline model. Joint models add both the qQTL and the xQTL annotation of the same context, and we report the τ^* of each annotation conditional on the other. For each trait, standard errors are obtained using the LDSC block jackknife (200 blocks). We pool trait-specific annotation estimates within each GWAS trait group using DerSimonian-Laird random-effects meta-analysis[83]. For τ^*, P values come from the meta-analysis z-score. For heritability enrichment, P values come from a separate random-effects meta-analysis of the enrichment test statistic. Error bars show 95% confidence intervals (pooled estimate ± 1.96 SE). Estimates for every dataset and trait group are given in Supplementary Table 3. Full methodological details are provided in Supplementary Note Stratified LD Score Regression

(S-LDSC) Enrichment Analysis.

KEGG and Gene Ontology enrichment analysis

To assess whether predefined gene sets are preferentially involved in known biological pathways or functional categories, we perform over-representation analyses using the KEGG pathway database[40] and the Gene Ontology (GO) resource[41]. For a gene universe $\mathcal{U}$, let $\mathcal{G} \subset \mathcal{U}$ denote a target gene set and let $\mathcal{T} \subset \mathcal{U}$ denote the genes annotated to a given KEGG pathway or GO term. We report fold enrichment, defined as the ratio between the observed fraction of genes in $\mathcal{G}$ that belong to $\mathcal{T}$ and the corresponding fraction in the background universe:

$$\text{FoldEnrichment}(\mathcal{G}, \mathcal{T}) = \frac{|\mathcal{G} \cap \mathcal{T}|/|\mathcal{G}|}{|\mathcal{T}|/|\mathcal{U}|}. \tag{1.11}$$

Values greater than 1 indicate over-representation of $\mathcal{T}$ in $\mathcal{G}$, whereas values less than 1 are interpreted descriptively as lower representation rather than as a formally tested lower-tail result. We assess statistical significance using a one-sided hypergeometric over-representation test that evaluates whether the observed overlap $|\mathcal{G} \cap \mathcal{T}|$ exceeds the expectation under random sampling from $\mathcal{U}$. We control multiple testing across pathways or GO terms using the Benjamini-Hochberg procedure, and we report terms with FDR-adjusted P values below 0.05.

For KEGG, we use the human pathway collection (organism = “hsa”) to evaluate enrichment of signaling, metabolic, and disease-related pathways. For Gene Ontology, we perform enrichment separately for the three GO domains, Biological Process (BP), Cellular Component (CC), and Molecular Function (MF). Gene identifiers are harmonized to Entrez Gene IDs prior to analysis, and genes without valid mappings are excluded from the testing universe. We use clusterProfiler[84] (`enrichKEGG`, and `enrichGO` with `org.Hs.eg.db`) with its default background of all genes annotated in each database. The target gene sets are the significant qTWAS genes and the significant standard TWAS genes for AD (union of the four GWAS) and for telomere length PC1. Among terms significant in both sets, the figures show, for each category, the five terms with the largest positive difference in fold enrichment between the qTWAS and the TWAS gene set.

Statistics and reproducibility

No statistical method was used to predetermine sample size; all available donors that passed quality control were analyzed, and sample sizes for each dataset are given in Supplementary Table 1. The analyses are observational, so randomization and blinding do not apply. Unless stated otherwise, tests are two-sided, and multiple testing is handled as described for each analysis. Analyses used R, Python, TensorQTL, quantreg, QRank, QUAIL, qQTLR, pecotmr, LDSC, clusterProfiler, org.Hs.eg.db and lavaan.

References

[1] C. Bellenguez et al. “New insights into the genetic etiology of Alzheimer’s disease and related dementias”. In: *Nature genetics* 54.4 (2022), pp. 412–436. DOI: 10.1038/s41588-022-01024-z.

[2] D. P. Wightman et al. “A genome-wide association study with 1,126,563 individuals identifies new risk loci for Alzheimer’s disease”. In: *Nature genetics* 53.9 (2021), pp. 1276–1282. DOI: 10.1038/s41588-021-00921-z.

[3] B. W. Kunkle et al. “Genetic meta-analysis of diagnosed Alzheimer’s disease identifies new risk loci and implicates Aβ, tau, immunity and lipid processing”. In: *Nature genetics* 51.3 (2019), pp. 414–430. DOI: 10.1038/s41588-019-0358-2.

[4] The GTEx Consortium et al. “The GTEx Consortium atlas of genetic regulatory effects across human tissues”. In: *Science* 369.6509 (2020), pp. 1318–1330. DOI: 10.1126/science.aaz1776.

[5] M. Fujita et al. “Cell subtype-specific effects of genetic variation in the Alzheimer’s disease brain”. In: *Nature Genetics* 56.4 (2024), pp. 605–614. DOI: 10.1038/s41588-024-01685-y.

[6] M. Wang et al. “The Mount Sinai cohort of large-scale genomic, transcriptomic and proteomic data in Alzheimer’s disease”. In: *Scientific data* 5.1 (2018), pp. 1–16. DOI: 10.1038/sdata.2018.185.

[7] P. L. De Jager et al. “A multi-omic atlas of the human frontal cortex for aging and Alzheimer’s disease research”. In: *Scientific data* 5.1 (2018), pp. 1–13. DOI: 10.1038/sdata.2018.142.

[8] M. V. Fernandez et al. “Genetic and multi-omic resources for Alzheimer disease and related dementia from the Knight Alzheimer Disease Research Center”. In: *Scientific data* 11.1 (2024), p. 768. DOI: 10.1038/s41597-024-03485-9.

[9] D. A. Bennett et al. “Religious orders study and rush memory and aging project”. In: *Journal of Alzheimer’s disease* 64.s1 (2018), S161–S189. DOI: 10.3233/JAD-179939.

[10] A. Lal et al. “Decoding sequence determinants of gene expression in diverse cellular and disease states”. In: *Nature Methods* 23.6 (2026), pp. 1138–1151. DOI: 10.1038/s41592-026-03102-0.

[11] A. E. Spiro et al. “A scalable approach to investigating sequence-to-function predictions from personal genomes”. In: *Nature Methods* 23.7 (2026), pp. 1308–1312. DOI: 10.1038/s41592-026-03124-8.

[12] S. Kasela et al. “Interaction molecular QTL mapping discovers cellular and environmental modifiers of genetic regulatory effects”. In: *The American Journal of Human Genetics* 111.1 (2024), pp. 133–149. DOI: 10.1016/j.ajhg.2023.11.013.

[13] A. M. Hulse and J. J. Cai. “Genetic Variants Contribute to Gene Expression Variability in Humans”. In: *Genetics* 193.1 (2013), pp. 95–108. DOI: 10.1534/genetics.112.146779.

[14] R. W. Corty and W. Valdar. “vqtl: An R Package for Mean-Variance QTL Mapping”. In: *G3: Genes, Genomes, Genetics* 8.12 (2018), pp. 3757–3766. DOI: 10.1534/g3.118.200642.

[15] A. K. Sarkar et al. “Discovery and characterization of variance QTLs in human induced pluripotent stem cells”. In: *PLOS Genetics* 15.4 (2019), e1008045. DOI: 10.1371/journal.pgen.1008045.

[16] C. Wang et al. “Genome-wide discovery for biomarkers using quantile regression at biobank scale”. In: *Nature Communications* 15.1 (2024), p. 6460. DOI: 10.1038/s41467-024-50726-x.

[17] B. Sun and L. Chen. “Quantile regression for challenging cases of eQTL mapping”. In: *Briefings in Bioinformatics* 21.5 (2020), pp. 1756–1765. DOI: 10.1093/bib/bbz097.

[18] X. Song et al. “QRank: a novel quantile regression tool for eQTL discovery”. In: *Bioinformatics* 33.14 (2017), pp. 2123–2130. DOI: 10.1093/bioinformatics/btx119.

[19] J. Miao et al. “A quantile integral linear model to quantify genetic effects on phenotypic variability”. In: *Proceedings of the National Academy of Sciences* 119.39 (2022), e2212959119. DOI: 10.1073/pnas.2212959119.

[20] A. Kozlova et al. “PICALM Alzheimer’s risk allele causes aberrant lipid droplets in microglia”. In: *Nature* 646.8087 (2025), pp. 1178–1186. DOI: 10.1038/s41586-025-09486-x.

[21] E. McArthur and J. A. Capra. “Topologically associating domain boundaries that are stable across diverse cell types are evolutionarily constrained and enriched for heritability”. In: *The American Journal of Human Genetics* 108.2 (2021), pp. 269–283. DOI: 10.1016/j.ajhg.2021.01.001.

[22] P. P. Kuksa et al. “FILER: a framework for harmonizing and querying large-scale functional genomics knowledge”. In: *NAR Genomics and Bioinformatics* 4.1 (2022), lqab123. DOI: 10.1093/nargab/lqab123.

[23] C. D. Arnold et al. “Genome-wide quantitative enhancer activity maps identified by STARR-seq”. In: *Science* 339.6123 (2013), pp. 1074–1077. DOI: 10.1126/science.1232542.

[24] E. P. Consortium et al. “An integrated encyclopedia of DNA elements in the human genome”. In: *Nature* 489.7414 (2012), pp. 57–74. DOI: 10.1038/nature11247.

[25] S. Gazal et al. “Functional architecture of low-frequency variants highlights strength of negative selection across coding and non-coding annotations”. In: *Nature genetics* 50.11 (2018), pp. 1600–1607. DOI: 10.1038/s41588-018-0231-8.

[26] S. Gazal, C. Marquez-Luna, H. K. Finucane, and A. L. Price. “Reconciling S-LDSC and LDAK functional enrichment estimates”. In: *Nature genetics* 51.8 (2019), pp. 1202–1204. DOI: 10.1038/s41588-019-0464-1.

[27] H. K. Finucane et al. “Partitioning heritability by functional annotation using genome-wide association summary statistics”. In: *Nature genetics* 47.11 (2015), pp. 1228–1235. DOI: 10.1038/ng.3404.

[28] S. Gazal et al. “Linkage disequilibrium–dependent architecture of human complex traits shows action of negative selection”. In: *Nature genetics* 49.10 (2017), pp. 1421–1427. DOI: 10.1038/ng.3954.

[29] J.-i. Satoh et al. “TMEM106B expression is reduced in Alzheimer’s disease brains”. In: *Alzheimer’s Research & Therapy* 6.2 (2014), p. 17. DOI: 10.1186/alzrt247.

[30] D. Brunetti et al. “Defective PITRM1 mitochondrial peptidase is associated with Aβ amyloidotic neurodegeneration”. In: *EMBO Molecular Medicine* 8.3 (2016), pp. 176–190. DOI: 10.15252/emmm.201505894.

[31] T. Wang, I. Ionita-Laza, and Y. Wei. “A unified quantile framework for nonlinear heterogeneous transcriptome-wide associations”. In: *The Annals of Applied Statistics* 19.2 (2025), pp. 967–985. DOI: 10.1214/24-AOAS1999.

[32] Y. Liu and J. Xie. “Cauchy combination test: a powerful test with analytic p-value calculation under arbitrary dependency structures”. In: *Journal of the American Statistical Association* 115.529 (2020), pp. 393–402. DOI: 10.1080/01621459.2018.1554485.

[33] A. Gusev et al. “Integrative approaches for large-scale transcriptome-wide association studies”. In: *Nature Genetics* 48.3 (2016), pp. 245–252. DOI: 10.1038/ng.3506.

[34] E. R. Gamazon et al. “A gene-based association method for mapping traits using reference transcriptome data”. In: *Nature Genetics* 47.9 (2015), pp. 1091–1098. DOI: 10.1038/ng.3367.

[35] The Alzheimer’s Disease Functional Genomics (FunGen-AD) Consortium. “Broad and deep dissection of Alzheimer’s disease loci through integrative xQTL modeling”. In submission. 2026.

[36] The Alzheimer’s Disease Functional Genomics (FunGen-AD) Consortium. *AD loci explorer, data release 2 November 2026 (commit c7cfba3)*. https://xqtl.niagads.org/AD-loci-explorer. Accessed 19 November 2026. 2026.

[37] Y. Y. Leung et al. “Comprehensive adjudication identifies 108 high-confidence genetic signals for Alzheimer’s disease and related dementias”. In: *Genome Medicine* (2026). DOI: 10.1186/s13073-026-01802-y.

[38] B. Jang et al. “A meta-analysis of single-nucleus expression quantitative trait loci linking genetic risk to brain disorders”. In: *Nature Genetics* 58.4 (2026), pp. 737–747. DOI: 10.1038/s41588-026-02541-x.

[39] S. Venkatesh et al. “Single-nucleus transcriptome-wide association study of human brain disorders”. In: *Nature* 657 (2026), pp. 1016–1026.

[40] M. Kanehisa and S. Goto. “KEGG: kyoto encyclopedia of genes and genomes”. In: *Nucleic acids research* 28.1 (2000), pp. 27–30. DOI: 10.1093/nar/28.1.27.

[41] M. Ashburner et al. “Gene ontology: tool for the unification of biology”. In: *Nature genetics* 25.1 (2000), pp. 25–29. DOI: 10.1038/75556.

[42] A. D. Schmitt et al. “A Compendium of Chromatin Contact Maps Reveals Spatially Active Regions in the Human Genome”. In: *Cell Reports* 17.8 (2016), pp. 2042–2059. DOI: 10.1016/j.celrep.2016.10.061.

[43] J. Cifello et al. “A TAD-informed aging-brain xQTL atlas of multi-modal and cell-type-resolved regulatory variation”. In: *medRxiv* (2026). DOI: 10.64898/2026.05.21.26353713.

[44] StatFunGen. *xQTL analysis protocol*. https://statfungen.github.io/xqtl-protocol/. 2026.

[45] Z. Qi et al. “Novel Missing Data Imputation Approaches Enhance Quantitative Trait Loci Discovery in Multi-Omics Analysis”. In: *medRxiv* (2023). DOI: 10.1101/2023.11.29.23299181.

[46] Q. Q. Huang, S. C. Ritchie, M. Brozynska, and M. Inouye. “Power, false discovery rate and Winner’s Curse in eQTL studies”. In: *Nucleic acids research* 46.22 (2018), e133–e133. DOI: 10.1093/nar/gky780.

[47] C. Gutenbrunner, J. Jurečková, R. Koenker, and S. Portnoy. “Tests of linear hypotheses based on regression rank scores”. In: *Journal of Nonparametric Statistics* 2.4 (1993), pp. 307–331. DOI: 10.1080/10485259308832561.

[48] R. Koenker and G. Bassett Jr. “Regression quantiles”. In: *Econometrica: journal of the Econometric Society* 46 (1978), pp. 33–50. DOI: 10.2307/1913643.

[49] A. Kuzma et al. “NIAGADS: A data repository for Alzheimer’s disease and related dementia genomics”. In: *Alzheimer’s & Dementia* 21.6 (2025). DOI: 10.1002/alz.70255.

[50] Y. Y. Leung et al. “Alzheimer’s Disease Sequencing Project release 4 whole genome sequencing dataset”. In: *Alzheimer’s & Dementia* 21.5 (2025), e70237. DOI: 10.1002/alz.70237.

[51] C. Bonferroni. “Teoria statistica delle classi e calcolo delle probabilita”. In: *Pubblicazioni del R istituto superiore di scienze economiche e commericiali di firenze* 8 (1936), pp. 3–62.

[52] C. Liu et al. “A multi-context regulome-wide association atlas for genetic studies of aging brain disorders”. In: *medRxiv* (2026). DOI: 10.64898/2026.05.15.26353329.

[53] G. Wang, A. Sarkar, P. Carbonetto, and M. Stephens. “A simple new approach to variable selection in regression, with application to genetic fine mapping”. In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 82.5 (2020), pp. 1273–1300. DOI: 10.1111/rssb.12388.

[54] Y. Kim, W. Wang, P. Carbonetto, and M. Stephens. “A flexible empirical Bayes approach to multiple linear regression and connections with penalized regression”. In: *Journal of Machine Learning Research* 25.185 (2024), pp. 1–59.

[55] R. Tibshirani. “Regression shrinkage and selection via the lasso”. In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 58.1 (1996), pp. 267–288. DOI: 10.1111/j.2517-6161.1996.tb02080.x.

[56] H. Zou and T. Hastie. “Regularization and variable selection via the elastic net”. In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 67.2 (2005), pp. 301–320. DOI: 10.1111/j.1467-9868.2005.00503.x.

[57] G. Moser et al. “Simultaneous discovery, estimation and prediction analysis of complex traits using a Bayesian mixture model”. In: *PLOS Genetics* 11.4, e1004969 (2015). DOI: 10.1371/journal.pgen.1004969.

[58] T. Park and G. Casella. “The Bayesian lasso”. In: *Journal of the American Statistical Association* 103.482 (2008), pp. 681–686. DOI: 10.1198/016214508000000337.

[59] F. Morgante et al. “A flexible empirical Bayes approach to multivariate multiple regression, and its improved accuracy in predicting multi-tissue gene expression from genotypes”. In: *PLOS Genetics* 19.7, e1010539 (2023). DOI: 10.1371/journal.pgen.1010539.

[60] Y. Zou, P. Carbonetto, D. Xie, G. Wang, and M. Stephens. “Fast and flexible joint fine-mapping of multiple traits via the Sum of Single Effects model”. In: *Nature Genetics* 58.2 (2026), pp. 454–462. DOI: 10.1038/s41588-025-02486-7.

[61] A. Taylor-Weiner et al. “Scaling computational genomics to millions of individuals with GPUs”. In: *Genome biology* 20.1 (2019), p. 228. DOI: 10.1186/s13059-019-1836-7.

[62] D. A. Knowles et al. “Allele-specific expression reveals interactions between genetic variation and environment”. In: *Nature methods* 14.7 (2017), pp. 699–702. DOI: 10.1038/nmeth.4298.

[63] J. R. Davis et al. “An efficient multiple-testing adjustment for eQTL studies that accounts for linkage disequilibrium between variants”. In: *The American Journal of Human Genetics* 98.1 (2016), pp. 216–224. DOI: 10.1016/j.ajhg.2015.11.021.

[64] K. Imai, L. Keele, and T. Yamamoto. “Identification, Inference and Sensitivity Analysis for Causal Mediation Effects”. In: *Statistical Science* 25.1 (2010), pp. 51–71. DOI: 10.1214/10-STS321.

[65] Y. Rosseel. “lavaan: An R Package for Structural Equation Modeling”. In: *Journal of Statistical Software* 48.2 (2012), pp. 1–36. DOI: 10.18637/jss.v048.i02.

[66] C. K. Enders and D. L. Bandalos. “The relative performance of full information maximum likelihood estimation for missing data in structural equation models”. In: *Structural equation modeling* 8.3 (2001), pp. 430–457. DOI: 10.1207/s15328007sem0803_5.

[67] B. Efron and R. J. Tibshirani. *An Introduction to the Bootstrap*. Monographs on Statistics and Applied Probability 57. New York: Chapman and Hall, 1993.

[68] E. C. Johnson et al. “Large-scale proteomic analysis of Alzheimer’s disease brain and cerebrospinal fluid reveals early changes in energy metabolism associated with microglia and astrocyte activation”. In: *Nature medicine* 26.5 (2020), pp. 769–780.

[69] C. Yang et al. “Genomic atlas of the proteome from brain, CSF and plasma prioritizes proteins implicated in neurological disorders”. In: *Nature Neuroscience* 24.9 (2021), pp. 1302–1312. DOI: 10.1038/s41593-021-00886-6.

[70] C. Coleman et al. “Multi-omic atlas of the parahippocampal gyrus in Alzheimer’s disease”. In: *Scientific Data* 10.1 (2023), p. 602. DOI: 10.1038/s41597-023-02507-2.

[71] G. S. Green et al. “Cellular communities reveal trajectories of brain ageing and Alzheimer’s disease”. In: *Nature* 633.8030 (2024), pp. 634–645. DOI: 10.1038/s41586-024-07871-6.

[72] H. Mathys et al. “Single-cell atlas reveals correlates of high cognitive function, dementia, and resilience to Alzheimer’s disease pathology”. In: *Cell* 186.20 (2023), pp. 4365–4385. DOI: 10.1016/j.cell.2023.08.039.

[73] H. Mathys et al. “Single-cell multiregion dissection of Alzheimer’s disease”. In: *Nature* 632.8026 (2024), pp. 858–868. DOI: 10.1038/s41586-024-07606-7.

[74] N. Comandante-Lou et al. “PLXNB1 and other signaling drives a pathologic astrocyte state contributing to cognitive decline in Alzheimer’s disease”. In: *bioRxiv* (2025).

[75] H. J. Zhou, L. Li, Y. Li, W. Li, and J. J. Li. “PCA outperforms popular hidden variable inference methods for molecular QTL mapping”. In: *Genome biology* 23.1 (2022), p. 210. DOI: 10.1186/s13059-022-02761-4.

[76] O. S. Burren et al. “Genetic architecture of telomere length in 462,666 UK Biobank whole-genome sequences”. In: *Nature Genetics* 56.9 (2024), pp. 1832–1840. DOI: 10.1038/s41588-024-01884-7.

[77] X. Cao et al. “Integrative multi-omics QTL colocalization maps regulatory architecture in aging human brain”. In: *medRxiv* (2025). DOI: 10.1101/2025.04.17.25326042.

[78] A. Kundaje et al. “Integrative analysis of 111 reference human epigenomes”. In: *Nature* 518.7539 (2015), pp. 317–330. DOI: 10.1038/nature14248.

[79] J. Cotney et al. “The autism-associated chromatin modifier CHD8 regulates other autism risk genes during human neurodevelopment”. In: *Nature communications* 6.1 (2015), p. 6404. DOI: 10.1038/ncomms7404.

[80] E. C. Mazur et al. “Progesterone receptor transcriptome and cistrome in decidualized human endometrial stromal cells”. In: *Endocrinology* 156.6 (2015), pp. 2239–2253. DOI: 10.1210/en.2014-1566.

[81] H. S. Kaya et al. “Roles of progesterone receptor A and B isoforms during human endometrial decidualization”. In: *Molecular Endocrinology* 29.6 (2015), pp. 882–895. DOI: 10.1210/me.2014-1363.

[82] Y. Benjamini and Y. Hochberg. “Controlling the false discovery rate: a practical and powerful approach to multiple testing”. In: *Journal of the Royal statistical society: series B (Methodological)* 57.1 (1995), pp. 289–300. DOI: 10.1111/j.2517-6161.1995.tb02031.x.

[83] R. DerSimonian and N. Laird. “Meta-analysis in clinical trials”. In: *Controlled Clinical Trials* 7.3 (1986), pp. 177–188.

[84] G. Yu, L.-G. Wang, Y. Han, and Q.-Y. He. “clusterProfiler: an R package for comparing biological themes among gene clusters”. In: *OMICS: A Journal of Integrative Biology* 16.5 (2012), pp. 284–287.

Author contributions

G.W. and T.W. designed and supervised the research. A.L., R.J., T.W. and G.W. conceived analysis plans. A.L. implemented the qQTL/qTWAS computational protocol and analyzed the data, with input from R.L., X.C., R.F., H.S., M.F., N.C., C.L. and J.E. on various QTL contexts and modalities. R.J. conducted simulation studies and contributed to iQTL analysis. A.L., R.J., X.C., T.W. and G.W. wrote the manuscript, with input from all authors who critically reviewed the manuscript, suggested revisions as needed, and approved the final version.

Competing interests

The authors declare no competing interests.

Supplementary Information

SUPPLEMENTARY NOTES

Model specification for linear QTL

We consider testing for association between Y_i and G_{ij} for genetic variant j conditional on covariates C_i.

$$Y_i = \alpha_0 + C_i^\top \alpha + G_{ij}\beta_j + \varepsilon_i, \tag{S.1}$$

where α_0 is the intercept, α is the vector of covariate effects, and β_j is the genetic effect of variant j on the phenotype Y_i. Here, $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$ are independent and identically distributed error terms for $i = 1, \ldots, n$. To test for the association between variant j and the phenotype, we test the null hypothesis $H_0 : \beta_j = 0$ using TensorQTL[61].

Hierarchical multiple-testing correction for QTL discovery

For both the xQTL and qQTL analyses, we applied a two-stage hierarchical multiple-testing procedure, following the strategy of Huang et al.[46]. Let g index genes/isoforms and j index cis-acting variants within the predefined cis window of gene g. Denote by $P_{g,j}$ the nominal P value for testing the association between variant j and gene g under a given model (linear or quantile). Let m_g be the number of cis variants tested for gene g.

Gene/global-level testing. For each gene g, we first computed Bonferroni-adjusted P values

$$\tilde{P}_{g,j} = \mathsf{min}\{1, m_g P_{g,j}\},$$

and defined the gene-level statistic

$$\tilde{P}^*_g = \min_j \tilde{P}_{g,j},$$

which corresponds to the minimum Bonferroni-adjusted P value over all cis variants of gene g (i.e., the Bonferroni-adjusted P value of the lead cis variant). We then applied the Benjamini-Hochberg procedure to the collection $\{\tilde{P}^*_g\}$ across all genes to control the false discovery rate (FDR) at level $\alpha_{\mathsf{gene}} = 0.05$. This yields BH-adjusted gene-level P values, denoted $\{q_g\}$; genes with $q_g \le \alpha_{\mathsf{gene}}$ were declared to have at least one significant cis-QTL under the given model.

Variant/local-level thresholding within significant genes. Conditional on gene-level discoveries, we next defined a data-adaptive variant-level significance threshold. We set

$$P^\star = \max_{g: q_g \le \alpha_{\mathsf{gene}}} \tilde{P}^*_g,$$

the largest lead-variant Bonferroni-adjusted P value among all significant genes. A variant-gene pair (g, j) was finally declared significant (a QTL) for that model if $\tilde{P}_{g,j} < P^\star$.

In this way, gene-level FDR is controlled at α_{gene}, and variant-level discoveries are restricted to genes that pass the gene-level screen, with a common Bonferroni threshold $P^\star$.

Trend-based Wald heterogeneity test

We use a trend-based Wald test to assess whether the genetic effect of a shared QTL is heterogeneous across quantile levels. Since the test is applied to each variant separately, we omit the variant index j and write G_i and $\beta(\tau)$ for G_{ij} and $\beta_j(\tau)$. For K quantile levels $\tau_1, \ldots, \tau_K$, let $\boldsymbol{\beta} = (\beta(\tau_1), \ldots, \beta(\tau_K))^\top$ and let $\widehat{\boldsymbol{\beta}}$ denote the corresponding vector of estimates. The null hypothesis of a constant genetic effect across quantiles is

$$H_0 : \ \boldsymbol{\beta} = \bar{\beta}\mathbf{1}_K \quad \text{for some } \bar{\beta} \in \mathbb{R},$$

where $\mathbf{1}_K$ is the K-vector of ones and $\bar{\beta}$ represents a homogeneous genetic effect that need not be zero.

Wald test under a prespecified trend. A conventional Wald test for H_0 jointly evaluates $K-1$ contrasts of the form $\beta(\tau_k) - \beta(\tau_{k'})$[85], allowing general departures from a constant effect. Since the quantile-specific coefficients are evaluations of the same functional coefficient $\beta(\tau)$, we can exploit a prespecified trend in its departure from constancy to construct a more powerful test when that trend is appropriate. Consider the working model

$$\beta(\tau) = \bar{\beta} + \gamma h(\tau), \qquad \text{or equivalently,} \qquad \boldsymbol{\beta} = \bar{\beta}\mathbf{1}_K + \gamma \boldsymbol{h},$$

where $h(\tau)$ is a prespecified trend that is nonconstant over the selected quantile levels and $\boldsymbol{h} = (h(\tau_1), \ldots, h(\tau_K))^\top$. The null hypothesis then reduces to $H_0 : \gamma = 0$.

Under standard regularity conditions, $\widehat{\boldsymbol{\beta}}$ is approximately multivariate normal with mean $\boldsymbol{\beta}$ and covariance $\boldsymbol{\Sigma}$. For any nonzero contrast $\boldsymbol{c}$ satisfying $\boldsymbol{c}^\top \mathbf{1}_K = 0$, the homogeneous effect cancels, giving $\boldsymbol{c}^\top \boldsymbol{\beta} = \gamma \boldsymbol{c}^\top \boldsymbol{h}$. Thus, $\boldsymbol{c}^\top \widehat{\boldsymbol{\beta}}$ has mean zero and variance $\boldsymbol{c}^\top \boldsymbol{\Sigma} \boldsymbol{c}$ under H_0. The corresponding Wald statistic is

$$\frac{(\boldsymbol{c}^\top \widehat{\boldsymbol{\beta}})^2}{\boldsymbol{c}^\top \boldsymbol{\Sigma} \boldsymbol{c}}.$$

This statistic is asymptotically χ^2_1 distributed under H_0. For $\gamma \neq 0$, the corresponding noncentral χ^2_1 approximation has noncentrality parameter

$$\lambda(\boldsymbol{c}) = \frac{\gamma^2 (\boldsymbol{c}^\top \boldsymbol{h})^2}{\boldsymbol{c}^\top \boldsymbol{\Sigma} \boldsymbol{c}}.$$

At a fixed significance level, power increases with $\lambda(\boldsymbol{c})$, so the optimal contrast maximizes this quantity subject to $\boldsymbol{c}^\top \mathbf{1}_K = 0$. Define

$$\widetilde{\boldsymbol{h}} = \boldsymbol{h} - \mathbf{1}_K \frac{\mathbf{1}_K^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{h}}{\mathbf{1}_K^\top \boldsymbol{\Sigma}^{-1} \mathbf{1}_K},$$

which removes the $\boldsymbol{\Sigma}^{-1}$-weighted projection of $\boldsymbol{h}$ onto $\mathbf{1}_K$, so that $\mathbf{1}_K^\top \boldsymbol{\Sigma}^{-1}\widetilde{\boldsymbol{h}} = 0$. Since $\boldsymbol{c}^\top \boldsymbol{h} = \boldsymbol{c}^\top \widetilde{\boldsymbol{h}}$ for any contrast satisfying $\boldsymbol{c}^\top \mathbf{1}_K = 0$, the Cauchy–Schwarz inequality gives

$$(\boldsymbol{c}^\top \boldsymbol{h})^2 = (\boldsymbol{c}^\top \widetilde{\boldsymbol{h}})^2 \le (\boldsymbol{c}^\top \boldsymbol{\Sigma} \boldsymbol{c})(\widetilde{\boldsymbol{h}}^\top \boldsymbol{\Sigma}^{-1} \widetilde{\boldsymbol{h}}),$$

with equality if and only if $\boldsymbol{c} \propto \boldsymbol{\Sigma}^{-1}\widetilde{\boldsymbol{h}}$. Thus, the optimal contrast can be chosen as

$$\boldsymbol{c}_{\text{opt}} = \boldsymbol{\Sigma}^{-1}\widetilde{\boldsymbol{h}}.$$

Implementation of the trend-based Wald test. We estimate $\boldsymbol{\Sigma}$ using a sandwich covariance estimator for quantile regression[85]:

$$\widehat{\Sigma}_{kk'} = n^{-1}\{\min(\tau_k, \tau_{k'}) - \tau_k \tau_{k'}\}\, \boldsymbol{e}^\top \widehat{\boldsymbol{D}}_1(\tau_k)^{-1}\, \widehat{\boldsymbol{D}}_0\, \widehat{\boldsymbol{D}}_1(\tau_{k'})^{-1} \boldsymbol{e}, \qquad k, k' = 1, \ldots, K.$$

Here $\boldsymbol{x}_i = (1, C_i^\top, G_i)^\top$ is the design vector in Equation (1.1), $\boldsymbol{e} = (0, \ldots, 0, 1)^\top$ selects the genotype coefficient, $\widehat{\boldsymbol{D}}_0 = n^{-1}\sum_{i=1}^n \boldsymbol{x}_i \boldsymbol{x}_i^\top$ and $\widehat{\boldsymbol{D}}_1(\tau) = n^{-1}\sum_{i=1}^n \widehat{f}_i(\tau)\boldsymbol{x}_i \boldsymbol{x}_i^\top$. Following Powell[86], $\widehat{f}_i(\tau)$ estimates the individual-specific error density at zero in the τth quantile regression from the fitted residuals using a Gaussian kernel.

With $\widetilde{\boldsymbol{h}}$ evaluated using $\widehat{\boldsymbol{\Sigma}}$, the estimated optimal contrast $\widehat{\boldsymbol{c}} = \widehat{\boldsymbol{\Sigma}}^{-1}\widetilde{\boldsymbol{h}}$ yields the Wald statistic

$$W_h = \frac{(\widehat{\boldsymbol{c}}^\top \widehat{\boldsymbol{\beta}})^2}{\widehat{\boldsymbol{c}}^\top \widehat{\boldsymbol{\Sigma}}\, \widehat{\boldsymbol{c}}} = \frac{(\widetilde{\boldsymbol{h}}^\top \widehat{\boldsymbol{\Sigma}}^{-1} \widehat{\boldsymbol{\beta}})^2}{\widetilde{\boldsymbol{h}}^\top \widehat{\boldsymbol{\Sigma}}^{-1} \widetilde{\boldsymbol{h}}}.$$

Under H_0, W_h converges in distribution to χ_1^2.

To accommodate different patterns of heterogeneity, we consider three prespecified trend functions and combine their Wald test P values using the Cauchy combination[32] to obtain a single P value for the trend-based Wald heterogeneity test. The three trend functions are

$$h_1(\tau) = \Phi^{-1}(\tau), \qquad h_2(\tau) = -\log(1-\tau) - 1, \qquad h_3(\tau) = \log \tau + 1,$$

referred to as the Gaussian, Lehmann and inverse Lehmann trends, respectively, corresponding to scale effects under symmetric, right-skewed and left-skewed error distributions. Specifically, the Gaussian trend arises from the location-scale model

$$Y_i = \alpha_0 + C_i^\top \alpha + \bar{\beta} G_i + (\sigma + \gamma G_i)\varepsilon_i, \qquad \varepsilon_i \sim \mathcal{N}(0,1).$$

The corresponding quantile-specific genetic effect is $\beta(\tau) = \bar{\beta} + \gamma \Phi^{-1}(\tau)$, with deviations from $\bar{\beta}$ in opposite directions below and above the median.

Type I error in simulations. Phenotypes were generated from Equation (1.7) with $\beta_E = \beta_I = 0$ and $\beta_G \in \{0, 0.2\}$, giving constant genetic effects across quantiles under H_0. We considered $n = 500$, 1,000 and 5,000 and MAF $= 0.1$, 0.2 and 0.4, with 1,000 replicates per setting. Within each

replicate, genotypes, covariates and errors were generated as in the iQTL simulation (Methods), and the tests were applied directly to the original simulated phenotypes. Supplementary Fig. 8 presents the empirical type I error rates for the three trend-specific Wald tests and their Cauchy combination across all simulation settings. For each combination of sample size and MAF, rejection rates were identical for $\beta_G = 0$ and 0.2 because the contrasts eliminate the constant shift in estimated quantile effects induced by changing β_G. The Cauchy-combined test showed slightly inflated type I error at smaller sample sizes or lower MAF, while rejection rates were close to the nominal 5% level at larger sample sizes and higher MAF.

Quantile weight integration

Quantile-specific qTWAS weights were integrated over broader quantile regions using fixed and data-adaptive partitions. Main-text qTWAS results use the fixed regions.

1. **Fixed quantile regions** The fixed strategy partitions the quantile range into three predefined intervals representing lower, middle, and upper portions of the expression distribution:

$$\mathcal{A}_k = \begin{cases} \tau \in [0.01, 0.33] & k = 1 \\ \tau \in [0.34, 0.66] & k = 2 \\ \tau \in [0.67, 0.99] & k = 3 \end{cases}$$

2. **Data-adaptive quantile regions** The data-adaptive strategy determines the number and boundaries of regions from the correlation structure of $\widehat{\omega}(\tau)$ across quantiles. The partition is determined through modularity optimization, which compares observed correlations within clusters to expected correlations under random assignments[87]. Let B be a $T \times K$ binary matrix, where T is the number of quantile levels and K is the number of clusters. Here, $B_{ij} = 1$ if quantile i belongs to cluster j and $B_{ij} = 0$ otherwise. The modularity Q is defined as:

$$Q = \frac{m_+}{m} Q_+ - \frac{m_-}{m} Q_-,$$

where $m_+ = \sum_{i,j} W_{ij}^+$ and $m_- = \sum_{i,j} |W_{ij}^-|$ are the sums of positive and negative correlations, respectively, in the correlation matrix W, with $W_{ij}^+ = W_{ij}\mathbf{1}(W_{ij} > 0)$ and $W_{ij}^- = W_{ij}\mathbf{1}(W_{ij} < 0)$. The components Q_+ and Q_- are defined as:

$$Q_+ = \frac{1}{m_+} \sum_{i,j} (W_{ij}^+ - \frac{k_i^+ k_j^+}{m_+}) B_{ij}, \quad Q_- = \frac{1}{m_-} \sum_{i,j} (W_{ij}^- - \frac{k_i^- k_j^-}{m_-}) B_{ij},$$

where k_i^+ and k_i^- are the sums of positive and negative correlations for quantile i with all other quantiles.

Using hierarchical clustering with a correlation threshold $\rho = 0.8$, we identify the optimal number of clusters K by maximizing Q. If all correlations exceed ρ, a single cluster is formed; otherwise, K is determined by the maximum modularity score. The algorithm further ensures

continuity of quantiles within each identified cluster, resulting in regions $\{\mathcal{C}_1, \ldots, \mathcal{C}_K\}$.

For either strategy, the integrated weight for region k is approximated by the trapezoidal rule

$$\widehat{\omega}_k = \int_{\mathcal{A}_k} \widehat{\omega}(\tau) d\tau \approx \sum_l \frac{1}{2}\{\widehat{\omega}(\tau_l) + \widehat{\omega}(\tau_{l+1})\}(\tau_{l+1} - \tau_l),$$

where the sum runs over consecutive grid points in $\mathcal{A}_k$, with $\widehat{\omega}$ set to 0 one grid step (0.01) beyond each end of the region. This expression approximates the signed integral of the quantile-specific weight function over the region rather than an unweighted sum across grid points. The region-level qTWAS statistic is invariant to positive scalar rescaling of $\widehat{\omega}_k$, so the integral is used to preserve the signed direction and shape of the quantile-specific weight function over the region.

Interaction effects can generate marginal qQTL signals

G×G interactions. When E represents another genetic variant, the interaction term $\beta_I(G \times E)$ captures a G×G interaction, where the effect of G on gene expression depends on the genotype of E. For example, suppose G is near a transcription factor gene and E lies in a regulatory region of a target gene. If the minor allele of E enhances binding affinity, G may increase expression only when this allele is present. In a marginal qQTL scan that omits E, individuals carrying the interacting allele may therefore form a subgroup with elevated or more variable expression within particular genotype groups of G, provided that the interacting allele is sufficiently frequent (e.g., MAF > 0.2). This subgroup heterogeneity can shift tail quantiles of Y across genotypes of G (e.g., $\tau = 0.75$ or $\tau = 0.9$), making G detectable as a qQTL even without explicitly modeling E.

G×E interactions. When E represents a biological context or exposure (e.g., tissue type or drug exposure), the interaction term $\beta_I(G \times E)$ captures context-specific genetic regulation. For example, G may regulate expression only in specific tissues (e.g., liver versus brain). If the environmental factor varies across individuals in the population, a marginal qQTL scan for G averages over subgroups with different values of E. The genetic effect may then be concentrated in one part of the expression distribution; for instance, individuals with $G = 2$ and high E may drive higher expression in the upper tail (e.g., $\tau = 0.9$), even if the marginal mean effect is weak. Thus, G×E interactions can appear as qQTLs when variation in the interacting context creates genotype-specific tail shifts or variance changes in Y.

Same-direction β_E and β_I increase marginal qQTL capture. Although qQTL scans test for differences in conditional quantiles, a simple variance calculation helps explain why qQTL capture depends on the relative directions of β_E and β_I. Consider the G×E model after centering E_i, omitting covariates, and comparing genotype groups $G_{ij} = 0$ and $G_{ij} = 1$:

$$\begin{aligned} Y_i \mid (G_{ij} = 0) &= \beta_0 + \beta_E E_i + \epsilon_i, \\ Y_i \mid (G_{ij} = 1) &= \beta_0 + \beta_G + (\beta_E + \beta_I) E_i + \epsilon_i. \end{aligned}$$

If E_i is independent of G_{ij} and the residual variance is shared across genotype groups, then after marginalizing over E_i,

$$\mathsf{Var}(Y_i \mid G_{ij} = 1) - \mathsf{Var}(Y_i \mid G_{ij} = 0) = (\beta_I^2 + 2\beta_E\beta_I)\mathsf{Var}(E_i).$$

The quantity $\beta_I^2 + 2\beta_E\beta_I$ helps explain why iQTLs with same-direction β_E and β_I are more likely to be captured by qQTL scans. When $\beta_E\beta_I > 0$, the cross-product $2\beta_E\beta_I$ amplifies the genotype-dependent distributional spread induced by the unmodeled factor E_i, producing stronger tail-quantile differences between genotype groups and thereby increasing the detectability of the signal in a marginal qQTL scan. When $\beta_E\beta_I < 0$, the interacting-factor main effect and interaction effect partly offset each other, attenuating genotype-specific tail differences and reducing qQTL-only capture. This calculation motivates the simulation and empirical overlap analyses shown in Fig. 4d,e.

Robustness of the *PICALM* mediation analysis

We evaluated the robustness of the real-data *PICALM* mediation result using a nonparametric bootstrap and a complete-case sensitivity analysis. Both analyses used the same INT-normalized, covariate-residualized microglial (M_i) and astrocyte (A_i) *PICALM* expression residuals as the main FIML analysis. This included the primary direction ($G \rightarrow M \rightarrow A$) and the reverse direction ($G \rightarrow A \rightarrow M$).

Nonparametric bootstrap. We generated $B = 1{,}000$ bootstrap replicates by resampling individuals with replacement from the full $N = 737$ sample. Within each replicate, microglial and astrocyte expression were re-residualized against their respective covariate matrices before refitting the two-equation SEM (main text Section *PICALM* mediation analysis and downstream heterogeneity simulation) by FIML, so that the residualization step itself contributed to the resampling uncertainty. Percentile confidence intervals for all path estimates defined in Section *PICALM* mediation analysis and downstream heterogeneity simulation (a, b, c', indirect, total, and proportion mediated) were computed from the empirical distribution of the replicate estimates.

Complete-case sensitivity analysis. Independently of the FIML missing-data treatment, we repeated the analysis using nested ordinary least squares (OLS) regression restricted to the $N = 733$ individuals with complete microglial and astrocyte expression data (a Baron-Kenny conditional-independence test[88]). For the primary direction, the outcome was the astrocyte residual A_i and the mediator was the microglial residual M_i. For the reverse direction, these roles were exchanged. We write the outcome as Y_i and the mediator as Z_i within each direction. We fitted a marginal model,

$$Y_i = \gamma_0 + cG_i + \epsilon_i,$$

and a conditional model that additionally includes the mediator,

$$Y_i = \gamma_0' + c'G_i + bZ_i + \epsilon_i,$$

Here, b and c' denote the paths defined by the SEM in Section *PICALM* mediation analysis and downstream heterogeneity simulation, and c denotes the total (marginal) genotype-outcome association estimated without the mediator. Following the Baron-Kenny mediation convention, the prime in c' distinguishes this mediator-adjusted coefficient from the unadjusted c. The reduction in the genotype coefficient,

$$\Delta = c - c',$$

is the OLS analog of the SEM indirect effect ab, and the proportion mediated is defined as

$$\Delta / c.$$

Confidence intervals for Δ were obtained by a nonparametric bootstrap over individuals ($B = 1{,}000$), and a nested F test was used to test whether the conditional model explains significantly more variance in Y_i than the marginal model.

These sensitivity analyses were used to assess whether the bootstrap and complete-case results agreed with the primary FIML mediation estimates.

Stratified LD Score Regression (S-LDSC) Enrichment Analysis

We used stratified LD score regression (S-LDSC) to quantify how functional genomic annotations contribute to SNP heritability. S-LDSC models the per-SNP effect size variance as a linear function of annotation values. Let β_j denote the standardized effect of SNP j, and let $A_{j,a}$ represent the value of annotation a at SNP j (binary or continuous). The model assumes

$$\mathrm{Var}(\beta_j) = \sum_a A_{j,a} \tau_a,$$

where τ_a quantifies the per-SNP heritability contribution of annotation a conditional on all other annotations in the model. In LD score regression form, this induces the expectation

$$\mathbb{E}[\chi_j^2] = N \sum_a \tau_a \ell_{j,a} + N a_0 + 1,$$

where N is the GWAS sample size, $\ell_{j,a} = \sum_k A_{k,a} r_{jk}^2$ is the stratified LD score of SNP j with respect to annotation a, and a_0 measures confounding such as population stratification.

Heritability partitioning and enrichment. For annotation a, its contribution to SNP heritability is

$$h_{g,a}^2 = \tau_a \sum_j A_{j,a},$$

and the total SNP heritability is

$$h_g^2 = \sum_j \mathrm{Var}(\beta_j).$$

We defined heritability enrichment as

$$\text{Enrichment}_a = \frac{h^2_{g,a}/h^2_g}{\left(\sum_j A_{j,a}\right)/M},$$

where M is the total number of SNPs. Values greater than 1 indicate that annotation a explains more heritability per SNP than expected under a uniform model.

Standardized annotation effect τ^*_a. To compare effect sizes across heterogeneous annotations, we computed the standardized effect

$$\tau^*_a = \tau_a \frac{\text{sd}(A_{\cdot,a})}{h^2_g/M},$$

where $\text{sd}(A_{\cdot,a})$ is the standard deviation of annotation a. The quantity τ^*_a represents the proportional change in per-SNP heritability associated with a one-standard-deviation increase in the annotation, conditional on all other annotations.

Joint models. In joint models, the qQTL and the xQTL annotation of the same context are added together to the baseline model, and we report the τ^* of each annotation conditional on the other. No combined joint τ^* is computed.

Computation details. All S-LDSC analyses included the full baseline-LD model as covariates to adjust for effects of minor allele frequency, LD-related architecture, and broad regulatory features. SNP-level annotations were aligned to the LDSC reference SNP list, and consistent minor allele frequency and imputation-quality filters were applied to both GWAS summary statistics and LD reference data. Importantly, LD scores and allele frequencies were computed using the **ADSP (Alzheimer's Disease Sequencing Project) panel samples** as the LD reference, ensuring population-matched LD structure for all S-LDSC analyses. Standard errors for τ_a, τ^*_a and enrichment were obtained using a block jackknife across 200 genomic blocks; reported P values come from the meta-analysis across traits below.

Meta-analysis across traits. To obtain trait-shared annotation effects, we meta-analyzed annotation-level statistics across GWAS traits using DerSimonian-Laird random-effects meta-analysis[83]. For trait-specific estimates $\hat{\theta}_{a,t}$ with standard errors $\text{SE}_{a,t}$ and between-trait variance $\hat{\sigma}^2_a$ (DerSimonian-Laird estimator), the meta-analyzed effect is

$$\hat{\theta}_{a,\text{meta}} = \frac{\sum_t w_{a,t}\hat{\theta}_{a,t}}{\sum_t w_{a,t}}, \qquad w_{a,t} = \frac{1}{\text{SE}^2_{a,t} + \hat{\sigma}^2_a}, \qquad \text{SE}_{a,\text{meta}} = \left(\sum_t w_{a,t}\right)^{-1/2},$$

with meta-analysis z-score $z_{a,\text{meta}} = \hat{\theta}_{a,\text{meta}}/\text{SE}_{a,\text{meta}}$ and two-sided P value $2\Phi(-|z_{a,\text{meta}}|)$. Error bars show $\hat{\theta}_{a,\text{meta}} \pm 1.96\,\text{SE}_{a,\text{meta}}$. For enrichment, the error bars use the pooled enrichment and its

SE, whereas the P value comes from a separate random-effects meta-analysis of the enrichment test statistic, which tests whether per-SNP heritability is higher inside the annotation than outside it. Meta-analyses were performed both across all traits jointly and within biologically defined trait subsets.

Supplementary Tables

Supplementary Table 1. Molecular QTL cohorts, datasets and donors.

Cohort	Molecular phenotype	Context	Dataset (source)	Donors
ROSMAP	snRNA-seq	Ast	Mega / CUIMC1 / MIT	737 / 419 / 385
		Exc	Mega / CUIMC1 / MIT	737 / 419 / 386
		Inh	Mega / CUIMC1 / MIT	736 / 419 / 379
		Mic	Mega / CUIMC1 / MIT	733 / 419 / 377
		OPC	Mega / CUIMC1 / MIT	735 / 418 / 383
		Oli	Mega / CUIMC1 / MIT	737 / 419 / 387
	Bulk RNA-seq	DLPFC	–	784
		PCC	–	441
		AC	–	593
		Monocyte	–	226
	Proteomics	DLPFC protein	–	416
MSBB	Proteomics	BM36 protein	–	184
	Bulk RNA-seq	BM10 / BM22 / BM36 / BM44	–	274 / 254 / 230 / 256
Knight-ADRC	Bulk RNA-seq	Brain RNA	–	354
	Proteomics	Brain protein	–	412

Unique donors (main analysis: 14 contexts, 26 datasets)	Donors
ROSMAP	901
of which single-nucleus (Mega, CUIMC1 and MIT)	739
Knight-ADRC (336 donors in both RNA and protein)	430
MSBB (proteomics)	184
Total	**1,515**

Donors are those entering QTL mapping after genotype and phenotype quality control. Unique donors were counted by deduplicating genotype sample IDs within each cohort; ROSMAP, Knight-ADRC and MSBB are independent brain banks. Of the single-nucleus donors, 417 of 419 CUIMC1 donors and all 387 MIT donors are also in Mega. MSBB bulk RNA-seq (BM10, BM22, BM36 and BM44) is used only in Supplementary Figs. 4 and 5 and is not included in the total.

Supplementary Table 2. qTWAS cohorts, datasets and donors.

Cohort	Molecular phenotype	Context	Dataset (source)	Donors
ROSMAP	snRNA-seq	Ast	Mega / CUIMC1	737 / 419
		Exc	Mega / CUIMC1	737 / 419
		Inh	Mega / CUIMC1	736 / 419
		Mic	Mega / CUIMC1	733 / 419
		OPC	Mega / CUIMC1	735 / 418
		Oli	Mega / CUIMC1	737 / 419
	Bulk RNA-seq	DLPFC	–	784
		PCC	–	441
		AC	–	593
		Monocyte	–	226
	Proteomics	DLPFC protein	–	416
MSBB	Proteomics	BM36 protein	–	184
Knight-ADRC	Bulk RNA-seq	Brain RNA	–	354
	Proteomics	Brain protein	–	412

qTWAS weights were trained in 14 contexts and 20 datasets from 1,515 unique donors (ROSMAP 901, Knight-ADRC 430, MSBB proteomics 184); dropping the MIT source does not change the ROSMAP total because all MIT donors are in Mega.

Supplementary Table 3. S-LDSC standardized effect sizes (τ^*) of qQTL and linear QTL annotations.

Dataset	Trait group	Marginal τ*, qQTL	Marginal τ*, linear QTL	Joint τ*, qQTL	Joint τ*, linear QTL
ROSMAP Ast (Mega)	All	0.016 (0.007)*	0.011 (0.009)	0.055 (0.014)*	-0.044 (0.015)
	Brain	0.020 (0.011)	0.020 (0.012)	0.048 (0.018)*	-0.030 (0.018)
	Aging & ND	0.034 (0.020)	0.022 (0.020)	0.112 (0.036)*	-0.081 (0.034)
ROSMAP Exc (Mega)	All	0.008 (0.008)	0.013 (0.008)	0.012 (0.015)	-0.002 (0.016)
	Brain	0.021 (0.012)	0.029 (0.012)*	0.000 (0.020)	0.019 (0.020)
	Aging & ND	0.035 (0.021)	0.041 (0.024)	0.015 (0.047)	0.033 (0.055)
ROSMAP Inh (Mega)	All	-0.011 (0.007)	-0.009 (0.007)	0.017 (0.013)	-0.025 (0.013)
	Brain	-0.007 (0.009)	-0.007 (0.009)	0.019 (0.018)	-0.027 (0.018)
	Aging & ND	-0.003 (0.024)	-0.010 (0.019)	0.016 (0.048)	-0.025 (0.036)
ROSMAP Mic (Mega)	All	-0.007 (0.007)	-0.012 (0.007)	0.043 (0.014)*	-0.048 (0.013)
	Brain	-0.005 (0.010)	-0.012 (0.010)	0.055 (0.021)*	-0.055 (0.019)
	Aging & ND	0.013 (0.030)	-0.001 (0.031)	0.131 (0.045)*	-0.101 (0.042)
ROSMAP OPC (Mega)	All	0.010 (0.007)	0.010 (0.006)	0.037 (0.012)*	-0.030 (0.012)
	Brain	0.011 (0.008)	0.018 (0.009)*	0.016 (0.016)	0.002 (0.016)
	Aging & ND	0.015 (0.019)	0.013 (0.020)	0.041 (0.037)	-0.026 (0.034)
ROSMAP Oli (Mega)	All	0.015 (0.007)*	0.011 (0.007)	0.051 (0.013)*	-0.045 (0.013)
	Brain	0.026 (0.010)*	0.024 (0.010)*	0.027 (0.018)	-0.002 (0.018)
	Aging & ND	0.050 (0.021)*	0.046 (0.020)*	0.077 (0.040)	-0.027 (0.039)
ROSMAP DLPFC RNA	All	0.023 (0.008)*	0.010 (0.008)	0.041 (0.011)*	-0.029 (0.012)
	Brain	0.001 (0.010)	0.000 (0.010)	0.014 (0.015)	-0.017 (0.015)
	Aging & ND	0.033 (0.022)	0.014 (0.022)	0.056 (0.030)	-0.032 (0.029)
ROSMAP PCC RNA	All	0.024 (0.008)*	0.026 (0.008)*	0.030 (0.012)*	-0.008 (0.012)
	Brain	0.015 (0.009)	0.022 (0.011)*	0.013 (0.016)	0.006 (0.016)
	Aging & ND	0.027 (0.021)	0.035 (0.021)	0.021 (0.032)	0.014 (0.031)
ROSMAP AC RNA	All	0.026 (0.008)*	0.018 (0.008)*	0.053 (0.012)*	-0.031 (0.013)
	Brain	0.020 (0.010)*	0.016 (0.011)	0.034 (0.016)*	-0.015 (0.017)
	Aging & ND	0.057 (0.022)*	0.053 (0.022)*	0.057 (0.033)	0.001 (0.032)
ROSMAP monocyte RNA	All	-0.003 (0.008)	0.015 (0.009)	-0.009 (0.011)	0.008 (0.012)
	Brain	-0.002 (0.010)	0.015 (0.012)	-0.011 (0.013)	0.009 (0.015)
	Aging & ND	0.002 (0.022)	0.022 (0.023)	-0.029 (0.031)	0.021 (0.030)
ROSMAP DLPFC protein	All	0.072 (0.009)*	0.074 (0.009)*	0.041 (0.018)*	0.036 (0.018)*
	Brain	0.078 (0.012)*	0.081 (0.013)*	0.042 (0.024)	0.047 (0.025)
	Aging & ND	0.072 (0.028)*	0.075 (0.031)*	0.019 (0.053)	0.056 (0.068)
MSBB BM36 protein	All	0.018 (0.008)*	0.023 (0.007)*	0.004 (0.015)	0.018 (0.014)
	Brain	0.029 (0.012)*	0.032 (0.010)*	0.017 (0.019)	0.007 (0.017)
	Aging & ND	0.003 (0.024)	-0.001 (0.020)	0.062 (0.046)	-0.078 (0.034)
Knight-ADRC RNA	All	0.003 (0.008)	-0.003 (0.008)	0.028 (0.012)*	-0.039 (0.012)
	Brain	0.007 (0.010)	0.000 (0.010)	0.030 (0.015)*	-0.029 (0.017)
	Aging & ND	0.011 (0.021)	0.002 (0.020)	0.032 (0.039)	-0.027 (0.037)
Knight-ADRC protein	All	-0.013 (0.006)	-0.012 (0.007)	0.012 (0.024)	-0.034 (0.021)
	Brain	-0.012 (0.009)	-0.007 (0.010)	-0.044 (0.032)	0.024 (0.029)
	Aging & ND	-0.030 (0.018)	-0.026 (0.021)	-0.032 (0.068)	0.001 (0.057)

Values are τ* (s.e.), meta-analysed across GWAS traits by random-effects meta-analysis; per-trait s.e. from a 200-block jackknife. All: 86 GWAS; Brain: 47 brain-related traits; Aging & ND: 13 aging or neurodegenerative traits. Marginal: one QTL annotation with the 97-category baseline model; joint: qQTL and linear QTL annotations together with the baseline model. $\tau^*_C = \tau_C \times sd_C \times M / h^2_g$. *$P$ < 0.05, positive estimates only.

Supplementary Table 4. Path estimates for the *PICALM* mediation analysis under three inferential approaches.

Direction	Path	FIML (*N* = 737)	Bootstrap (*N* = 737)	Complete-case OLS (*N* = 733)
D1	*a*	−0.497 [−0.547, −0.447]	−0.456 [−0.502, −0.408]	-
	b	0.105 [0.047, 0.163]	0.081 [0.028, 0.136]	0.105 ($P = 4.3 \times 10^{-4}$)
	c′	−0.065 [−0.114, −0.015]	−0.067 [−0.106, −0.023]	−0.067 ($P = 0.008$)
	Indirect effect (*ab*/*Δ*)	−0.052 [−0.082, −0.023]	−0.037 [−0.062, −0.013]	−0.052 [−0.083, −0.023] ($P = 0.001$)
	Total effect (*c*)	−0.117 [−0.157, −0.077]	−0.104 [−0.137, −0.068]	−0.119 ($P = 1 \times 10^{-8}$)
	Proportion mediated	0.447 [0.156, 0.739]	0.365 [0.123, 0.694]	0.438
D2	*a*	−0.117 [−0.157, −0.077]	−0.102 [−0.140, −0.066]	*Not estimated*
	b	0.160 [0.072, 0.249]	0.130 [0.035, 0.227]	0.160 ($P = 4.3 \times 10^{-4}$)
	c′	−0.479 [−0.529, −0.428]	−0.442 [−0.492, −0.394]	−0.479 ($P = 4.1 \times 10^{-63}$)
	Indirect effect (*ab*/*Δ*)	−0.019 [−0.031, −0.007]	−0.013 [−0.025, −0.003]	−0.019 [−0.032, −0.008] ($P = 0.001$)
	Total effect (*c*)	−0.497 [−0.547, −0.447]	−0.455 [−0.505, −0.406]	−0.498 ($P = 8.2 \times 10^{-69}$)
	Proportion mediated	0.038 [0.013, 0.062]	0.029 [0.007, 0.057]	0.039

Note: FIML SEM uses delta-method confidence intervals. Bootstrap re-residualizes microglial and astrocyte expression within each of $B = 1{,}000$ replicates before refitting FIML. Complete-case OLS is a Baron-Kenny nested regression restricted to individuals with no missing expression ($N = 733$). D1, primary direction ($G \rightarrow M \rightarrow A$); D2, reverse direction ($G \rightarrow A \rightarrow M$). *a*, genotype-to-mediator path; *b*, mediator-to-outcome path conditional on genotype; *c*, total (marginal) genotype-outcome association without the mediator; *c′*, the same association after adjusting for the mediator (see Supplementary Methods for the Baron-Kenny *c* versus *c′* convention). Values in brackets are 95% confidence intervals. The Complete-case OLS column reports *P* values where a confidence interval was not computed for that path; the indirect-effect row reports the bootstrap *P* value for *Δ*. The Complete-case OLS approach does not separately estimate path *a* because $\Delta = c - c'$ is obtained directly from the marginal and conditional regressions without decomposing it into *a* and *b*.

Supplementary Figures

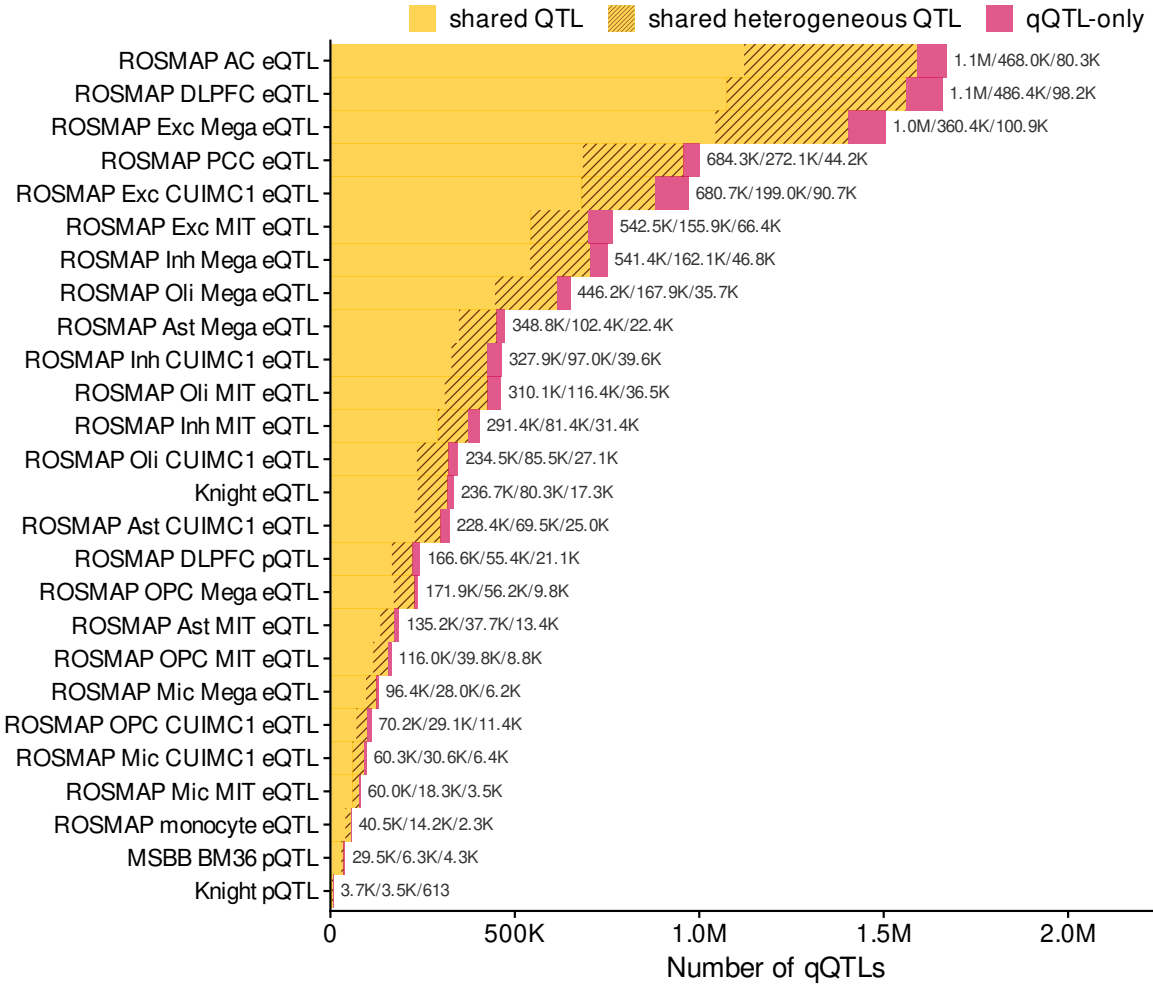

Supplementary Figure 1. qQTL heterogeneity category by dataset. qQTL-only associations are shown in red. Shared qQTL associations are shown in yellow, with hatched yellow indicating shared heterogeneous QTLs and solid yellow indicating shared homogeneous QTLs. Labels give the shared homogeneous/shared heterogeneous/qQTL-only counts.

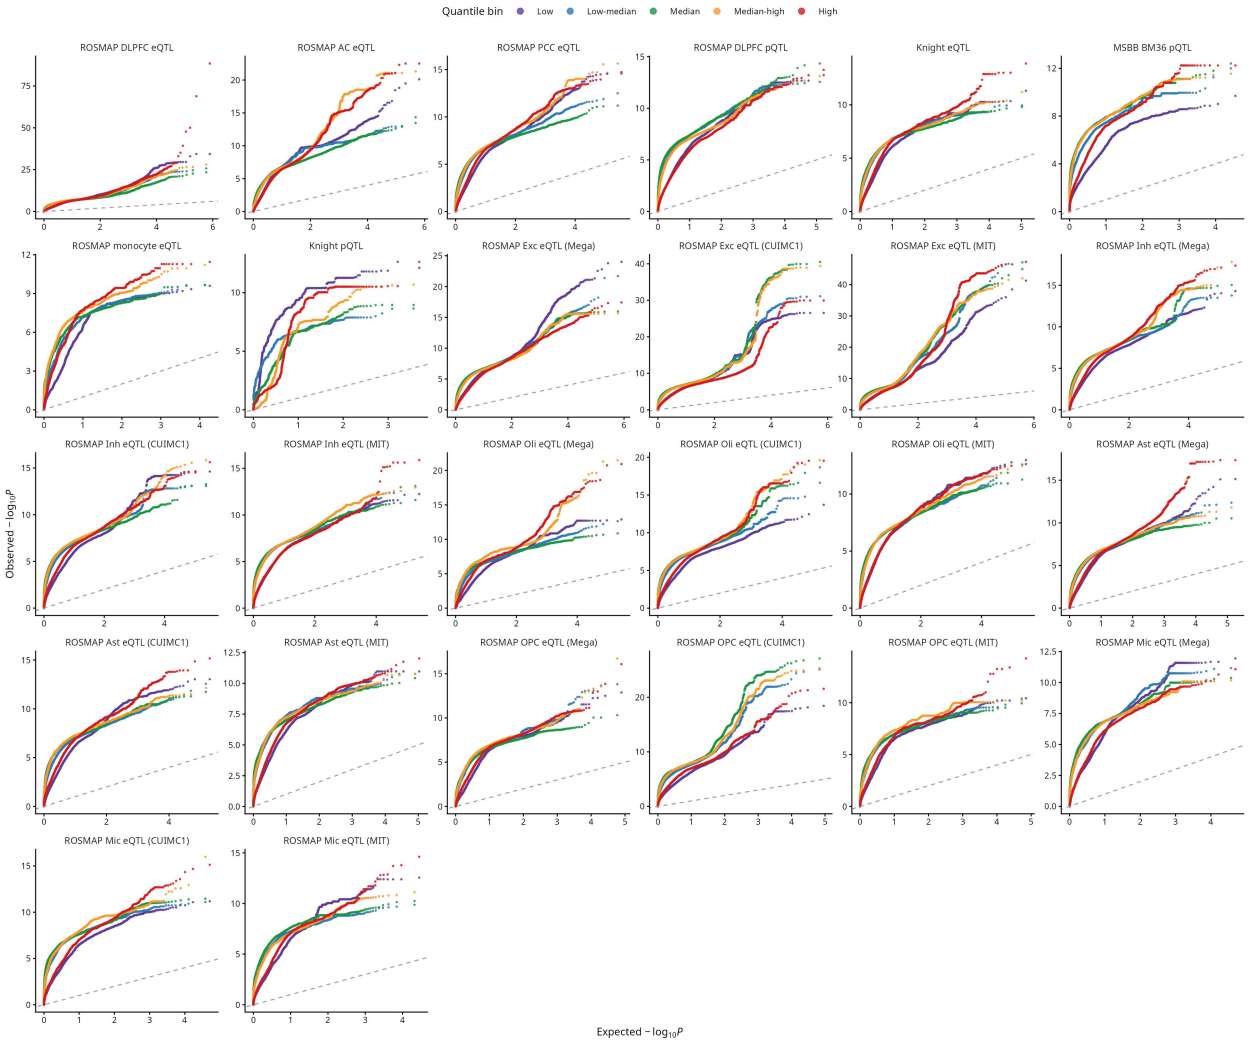

Supplementary Figure 2. QQ plots for qQTL-only associations in each of the 26 datasets. qQTL-only P values pooled into the five quantile bins of Fig. 1f (Low, τ = 0.10–0.25; Low-median, 0.30–0.40; Median, 0.45–0.55; Median-high, 0.60–0.70; High, 0.75–0.90). Dashed line, expected under the null.

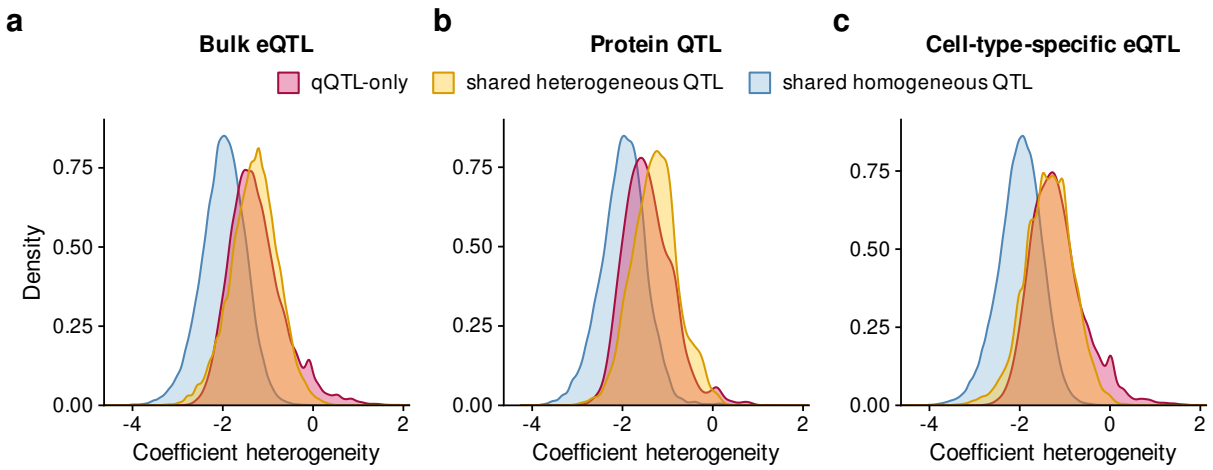

Supplementary Figure 3. Densities of the heterogeneity index for qQTLs. a–c, Density distributions of the heterogeneity index for qQTLs from bulk eQTL (**a**), protein QTL (**b**) and cell-type-specific eQTL (**c**) analyses. For each variant j, the heterogeneity index is defined as $\log|\text{sd}\{\widehat{\beta}_j(\tau)\}/\text{mean}\{\widehat{\beta}_j(\tau)\}|$, computed across the 19 quantile levels $\tau = 0.05, 0.10, \ldots, 0.95$, where $\widehat{\beta}_j(\tau)$ is the estimated quantile-specific genetic effect. Higher values correspond to greater variation in the genetic effect across quantiles relative to its mean. qQTL-only variants (pink) and shared heterogeneous QTLs (yellow) showed consistently elevated heterogeneity relative to shared homogeneous QTLs (blue).

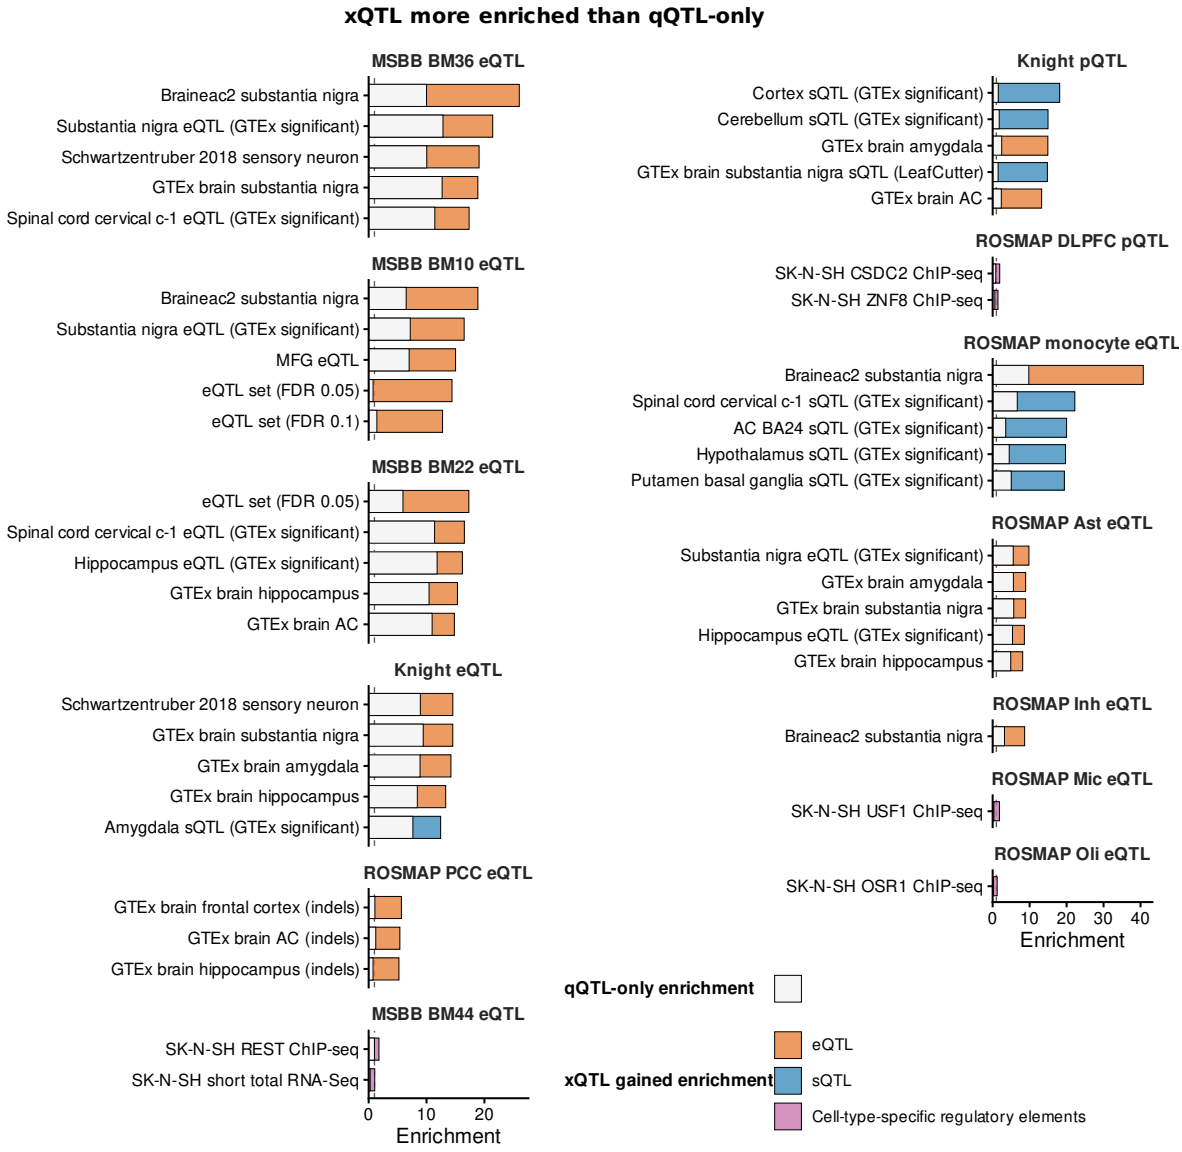

Supplementary Figure 4. Functional annotations preferentially enriched in xQTL variants relative to qQTL-only variants. Excess-of-overlap (EOO) enrichment analysis comparing xQTL and qQTL-only variants across bulk eQTL, protein QTL, and cell-type-specific eQTL datasets. For each molecular phenotype, we identified up to the five functional annotation categories showing the largest enrichment differences favoring xQTL variants over qQTL-only variants. Gray bars indicate the enrichment attributable to qQTL-only variants, whereas colored segments represent the additional enrichment observed for xQTL variants. Annotation categories are grouped into eQTL-related annotations (orange), sQTL-related annotations (blue), and cell-type-specific regulatory elements (purple). Across datasets, the annotations most preferentially enriched in xQTL variants were predominantly canonical eQTL and sQTL categories, consistent with standard xQTL mapping preferentially capturing established mean-level regulatory mechanisms.

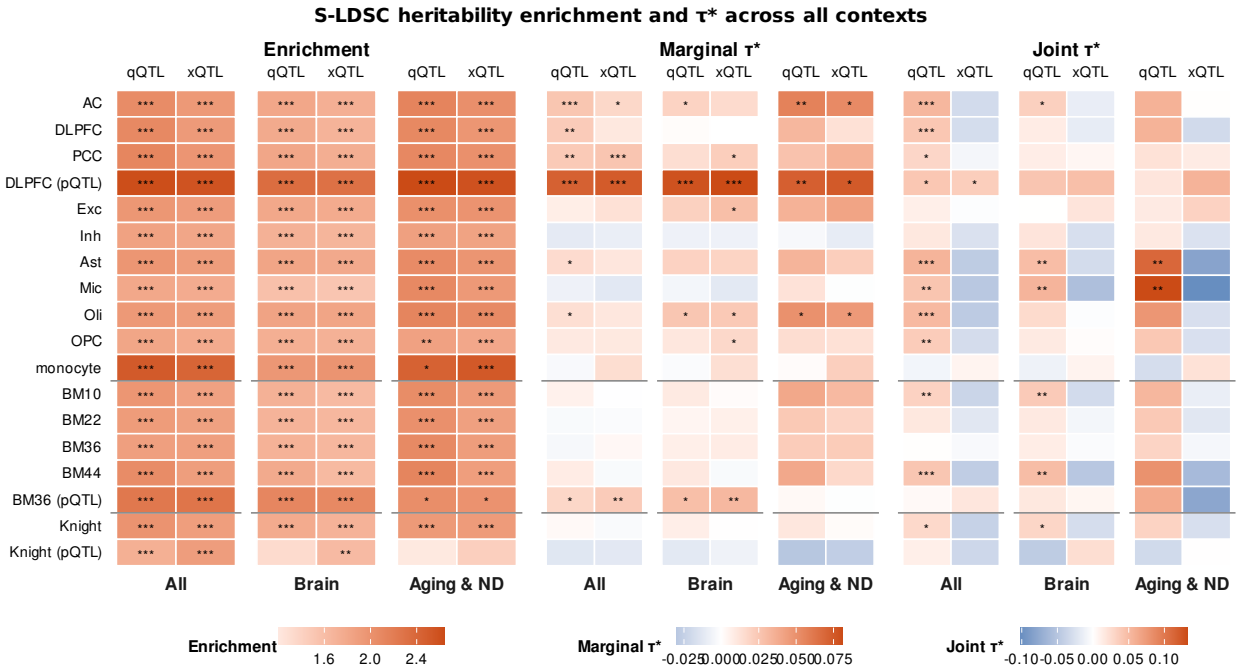

Supplementary Figure 5. Stratified LD score regression analysis across molecular contexts. S-LDSC heritability enrichment and standardized effect-size estimates (τ^*) for qQTL and conventional xQTL annotations across all molecular contexts. Rows correspond to individual datasets and cell types, grouped by cohort (ROSMAP, MSBB, and Knight), including bulk eQTL, protein QTL, and single-cell eQTL analyses. Columns show results for three GWAS trait groups: all 86 traits, 47 brain-related traits and 13 aging or neurodegenerative traits. Color intensity reflects the magnitude of enrichment or τ^*, and significance is indicated by asterisks (*$P \leq 0.05$, **$P \leq 0.01$, ***$P \leq 0.001$; positive estimates only). τ^* values, except for the MSBB bulk RNA datasets, are in Supplementary Table 3.

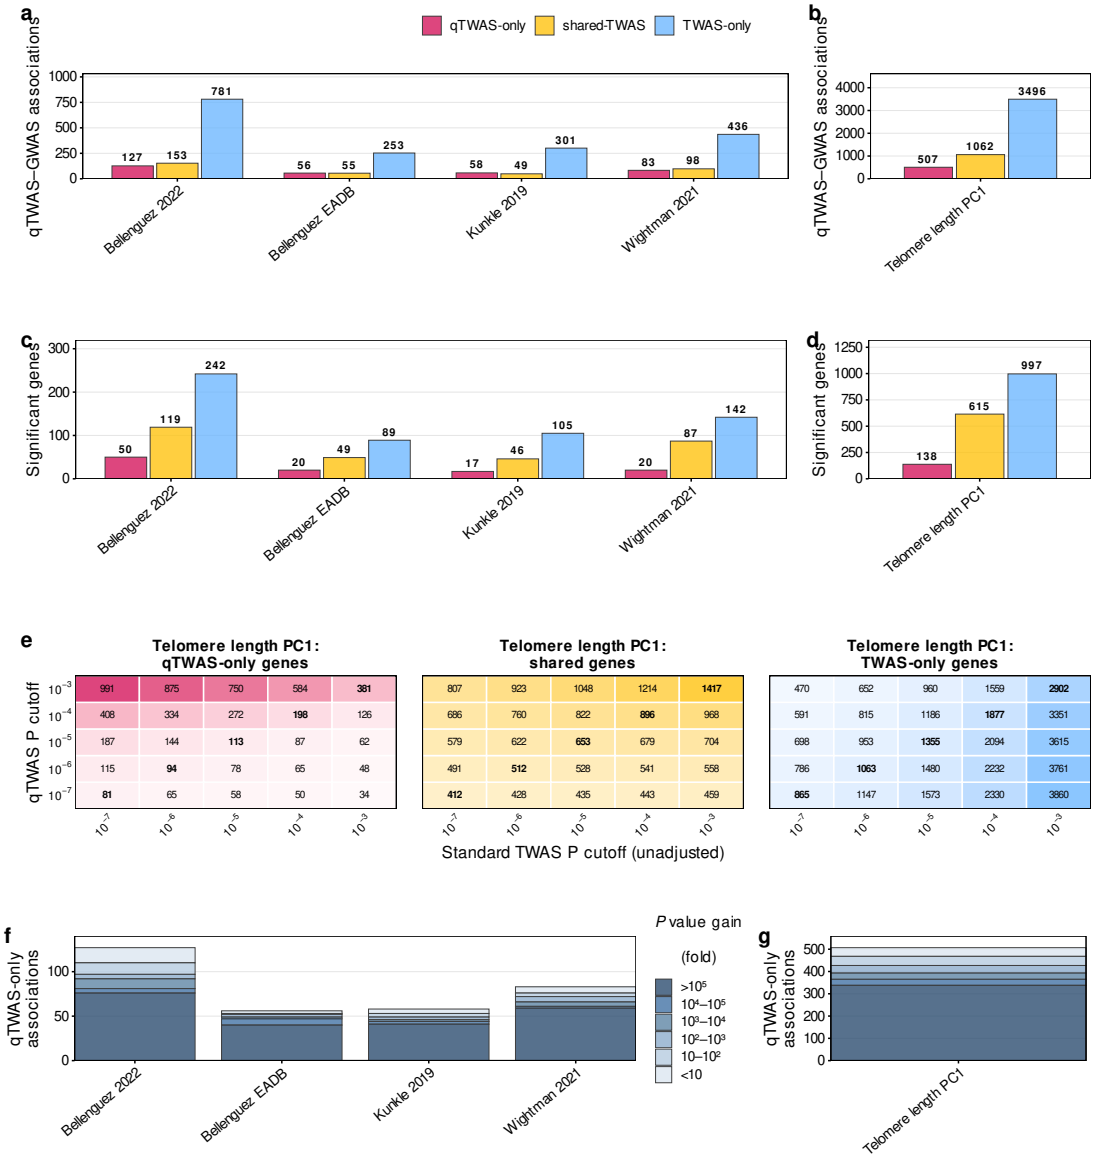

Supplementary Figure 6. qTWAS discovery on disease-relevant gene and context-specific associations for Alzheimer's disease and telomere length PC1. Analyses were performed across four Alzheimer's disease (AD) GWAS datasets (Bellenguez 2022, EADB 2022, Kunkle 2019, and Wightman 2021) and one telomere length PC1 GWAS. Significance was defined using a Bonferroni-adjusted threshold across molecular contexts. **a,b,** Numbers of significant gene-context associations identified by qTWAS-only (red), shared between qTWAS and standard TWAS (gold), and standard TWAS-only (blue) for AD GWAS datasets (**a**) and telomere length PC1 (**b**). **c,d,** Numbers of significant genes identified by qTWAS-only (red), shared between qTWAS and standard TWAS (gold), and standard TWAS-only (blue) for AD GWAS datasets (**c**) and telomere length PC1 (**d**). **e,** Telomere length PC1: number of qTWAS-only, shared and TWAS-only genes across cutoffs on the unadjusted qTWAS and standard TWAS P, as in Fig. 5d. **f,g,** Improvement in statistical significance achieved by qTWAS relative to standard TWAS for qTWAS-only associations in AD (**f**) and telomere length PC1 (**g**). Associations are grouped according to the fold-improvement in P value, measured as the ratio of TWAS P value to qTWAS P value. Darker shading denotes larger gains in statistical significance.

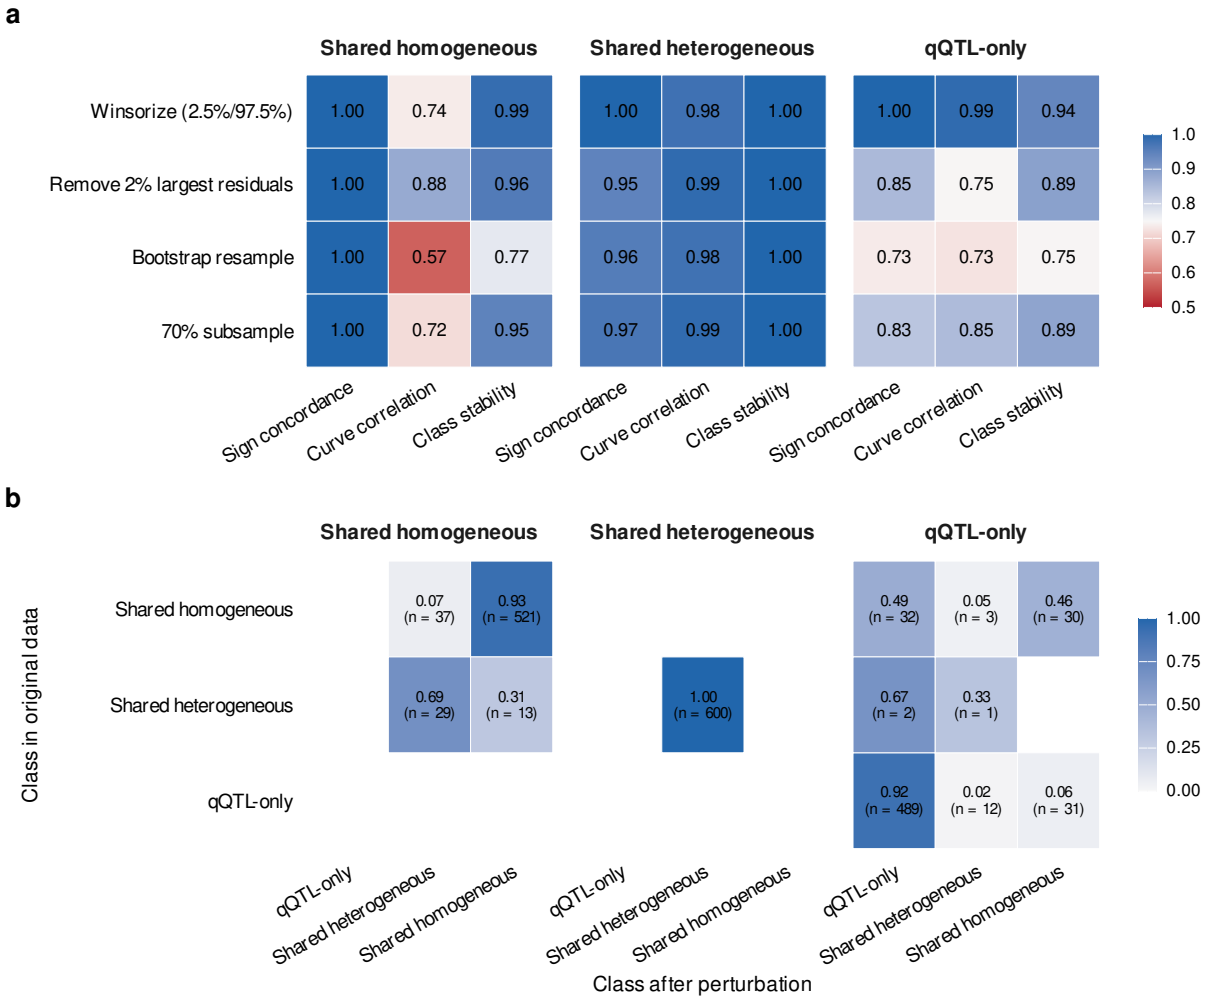

Supplementary Figure 7. Sign concordance, curve correlation, and class stability across 3 true scenarios $\times$ 4 perturbations. a, Heatmap of the three metrics (sign concordance, curve correlation, class stability under the trend-based Wald heterogeneity test with Cauchy combination) across the 3 true scenarios (homogeneous, heterogeneous, qQTL-only) and 4 perturbations (winsorize, influential-observation removal, bootstrap resample, 70% subsample), $n = 700$, 150 replicates per cell. Blue = high (stable/concordant), red = low. **b,** Class stability: proportion of simulated associations in each class in the original data (rows) assigned to each class after perturbation (columns), pooled over the four perturbations, for each true scenario; n, number of associations.

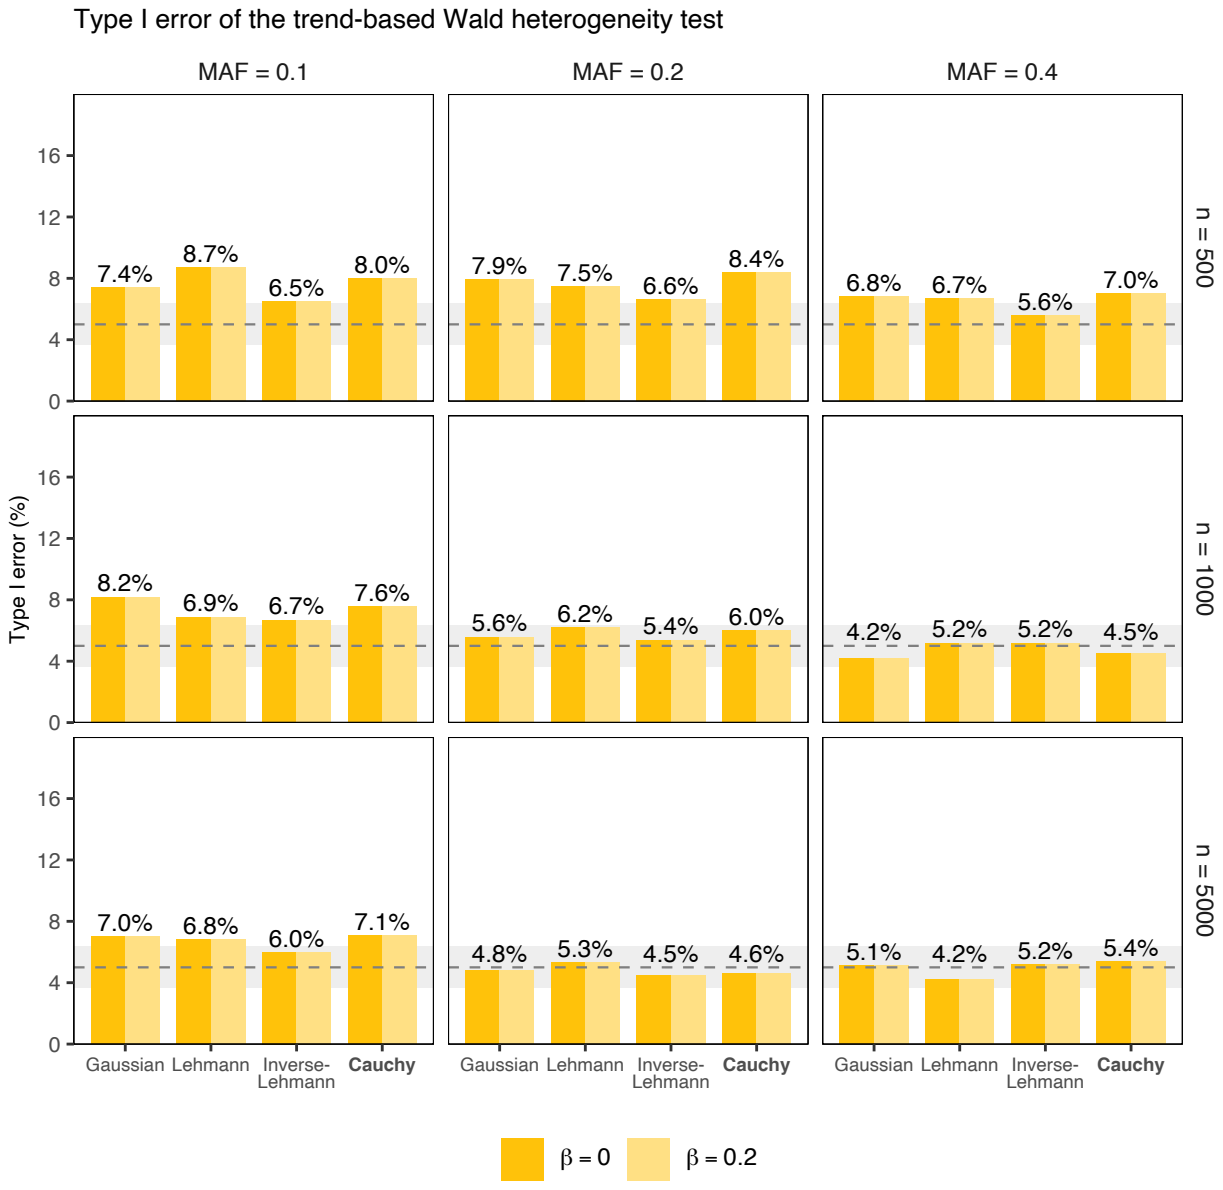

Supplementary Figure 8. Calibration of the trend-based Wald heterogeneity test. Empirical type I error at $\alpha = 0.05$ of the Wald tests for three trends (Gaussian, Lehmann and Inverse-Lehmann) and of their Cauchy combination ($K = 17$ quantile levels, τ = 0.10, 0.15, …, 0.90), for n = 500, 1,000 and 5,000 individuals (rows) and MAF = 0.1, 0.2 and 0.4 (columns). We simulated 1,000 replicates in each setting, with phenotypes generated under a constant genetic effect $\beta = 0$ (yellow) and $\beta = 0.2$ (light yellow); the trend-based Wald heterogeneity test is invariant to such a constant shift, so the two settings give identical results. Dashed line, nominal level of 5%; grey band, the 95% Monte Carlo range 3.65–6.35% ($0.05 \pm 1.96\sqrt{0.05 \times 0.95/1000}$).

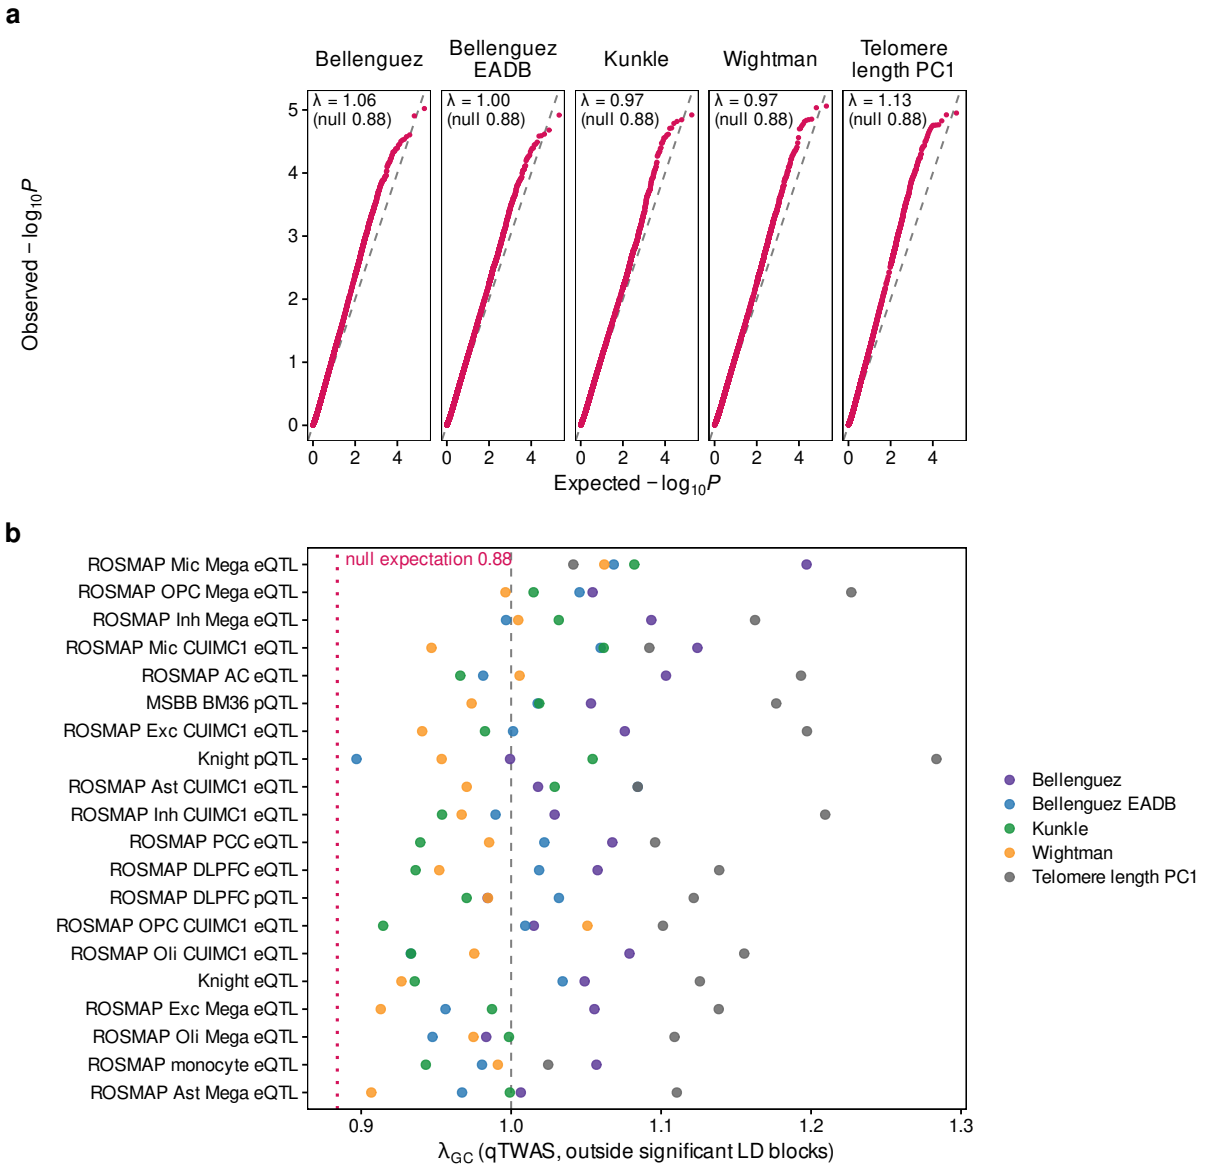

Supplementary Figure 9. qTWAS test statistics show modest genomic inflation outside significant loci. a, QQ plots of qTWAS P values for gene-dataset pairs whose LD block contains no Bonferroni-significant qTWAS or standard TWAS association for that GWAS set. Shown for the four AD GWAS summary-statistic sets and telomere length PC1. Points: qTWAS P values (Cauchy combination over three quantile bands). Insets give the genomic inflation factor λ_{GC} = χ^2_1 quantile at (1 − median P) / 0.455, and the null expectation for the Cauchy-combined statistic (0.88; Supplementary Fig. 10a). **b,** λ_{GC} per dataset for qTWAS, by GWAS set. The dashed line marks λ_{GC} = 1; the dotted red line, the null expectation 0.88.

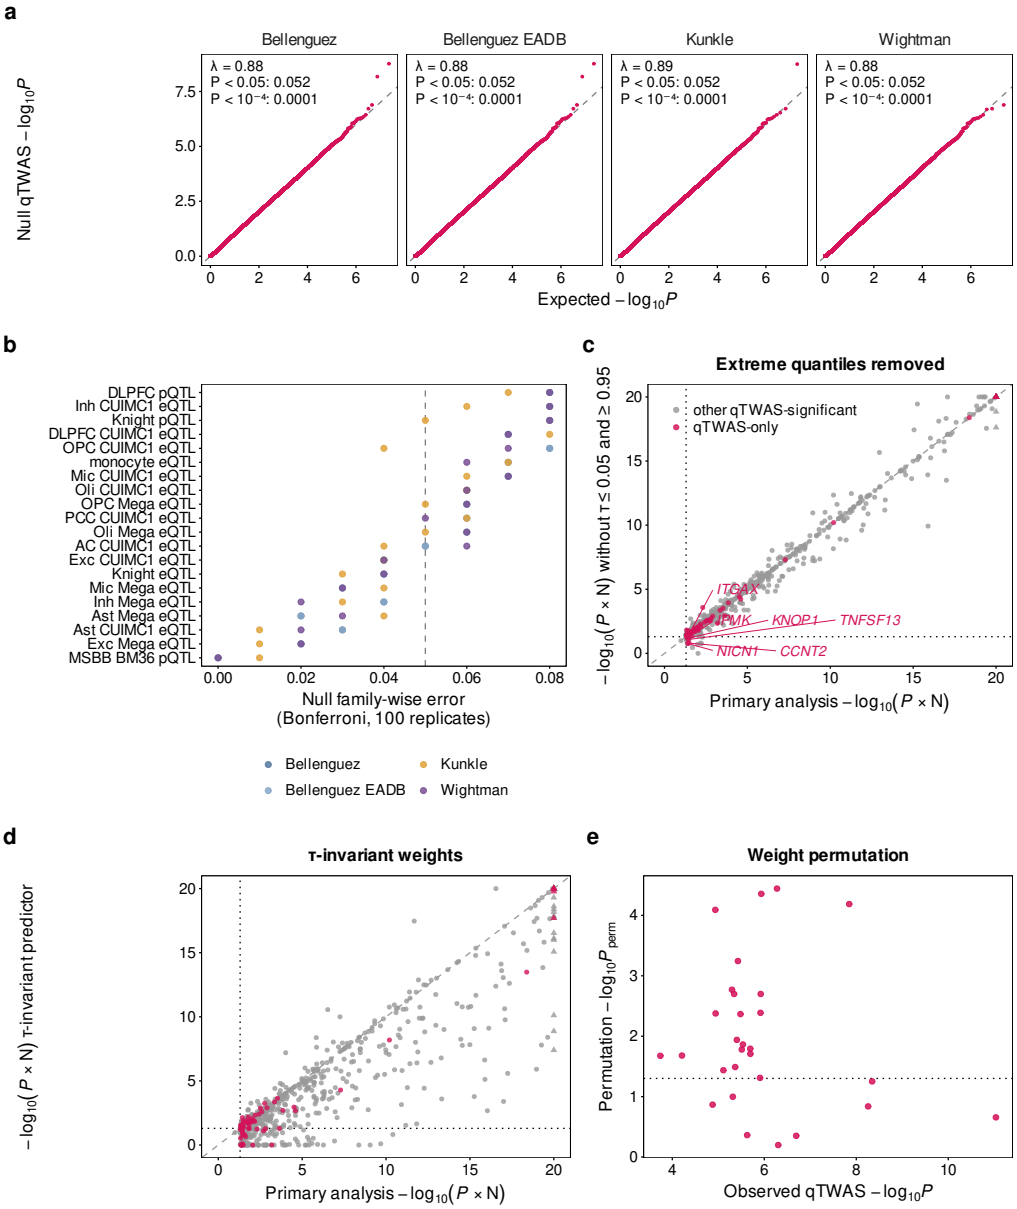

Supplementary Figure 10. qTWAS calibration and robustness of the qTWAS-only genes. a, QQ plots of Cauchy-combined qTWAS P values when null GWAS z-scores, simulated from the LD reference, are passed through the qTWAS procedure (100 replicates; all AD gene-dataset tests), with λ_{GC} and the fraction of tests below P = 0.05 and 10^{-4}. **b,** Bonferroni family-wise error rate for each AD GWAS and molecular dataset under the null; dashed line, 0.05. **c,** Bonferroni-adjusted P (P × N) after removing $\tau \leq 0.05$ and $\tau \geq 0.95$ from the band weights, against the primary-analysis P × N, for all qTWAS-significant AD associations; magenta, qTWAS-only; triangles, capped at 10^{-20}; dotted lines, 0.05. **d,** As in **c**, with the quantile-varying weights replaced by one τ-invariant predictor on the same variants. **e,** Weight-permutation P against the observed qTWAS P for the 29 qTWAS-only associations (19 genes) whose models have more than six variants; dotted line, P = 0.05.

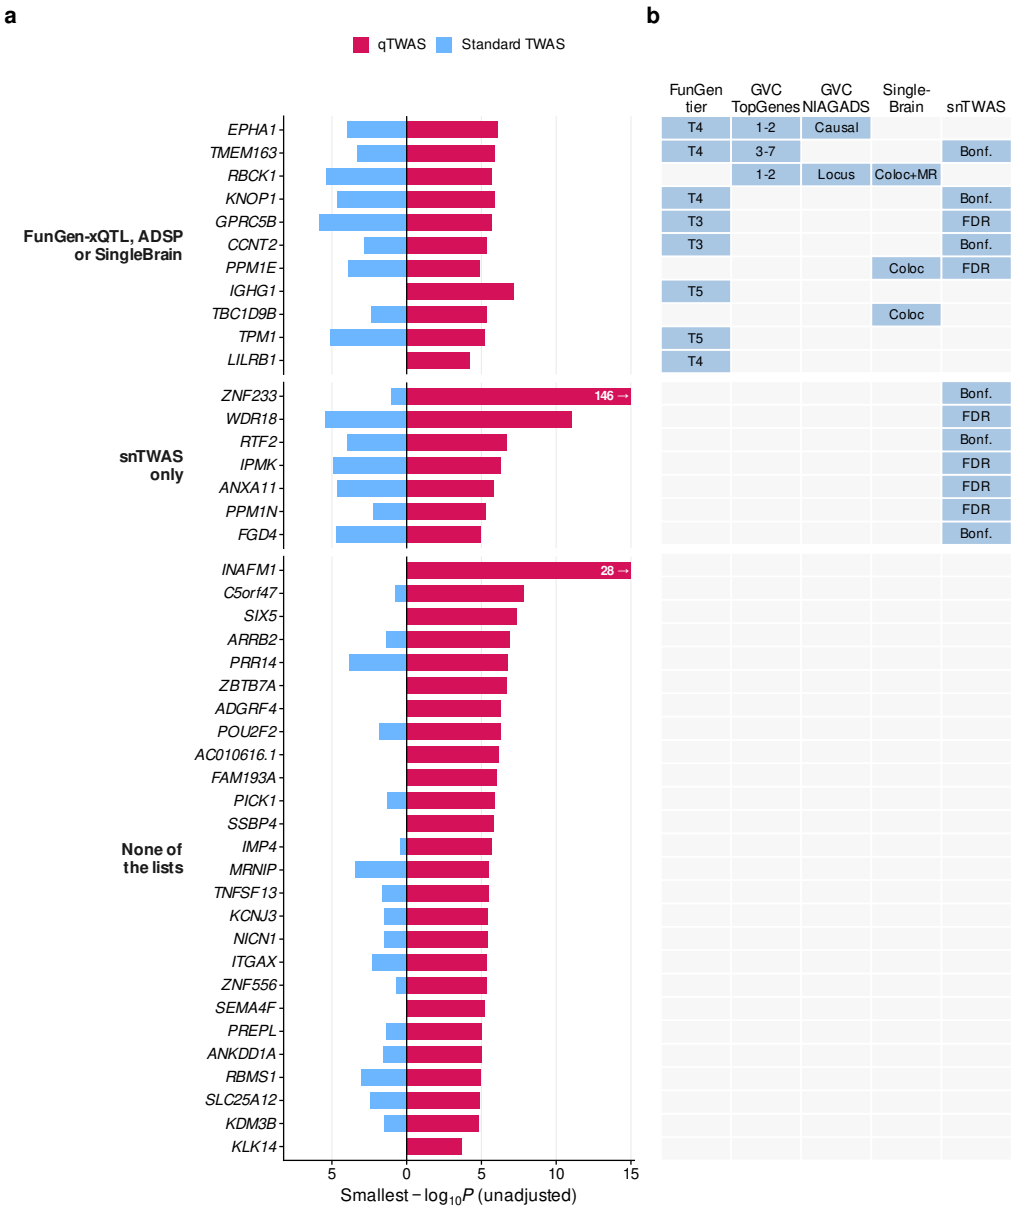

Supplementary Figure 11. External AD gene evidence for the 44 AD qTWAS-only genes. The 44 genes are significant by qTWAS (Bonferroni-adjusted $P < 0.05$) in at least one dataset and AD GWAS set and are not significant by matched standard TWAS in any dataset or AD GWAS set. They are grouped as in the Results: genes in the FunGen-xQTL AD loci tiers T1–T5, the ADSP Gene Verification Committee (GVC) lists or SingleBrain (11 genes); genes found only by the single-nucleus TWAS (7 genes); and genes in none of the compared lists (26 genes). Within each group, genes are ordered by the number of lists that include them and then by qTWAS P. **a,** Smallest unadjusted P ($-\log_{10}$ scale) across the 62 qTWAS-only gene–dataset–GWAS associations of each gene: qTWAS (Cauchy-combined P; magenta, right) and standard TWAS (minimum P across prediction methods; blue, left). qTWAS values above 15 are drawn at 15 with the value printed (*ZNF233*, 146; *INAFM1*, 28). Genes without a standard TWAS result have no blue bar. **b,** Membership in each external resource. A filled tile marks a gene that is in the list, and the text in the tile gives its evidence level. FunGen tier: FunGen-xQTL AD loci gene tier T1–T5[35,36] (Methods). GVC TopGenes: best ADSP GVC TopGenes tier across evaluations[37]; tiers 1–2 are high-confidence loci with sufficient evidence for association, and tiers 3–7 have weaker evidence. GVC NIAGADS: earlier ADSP GVC list (Causal, adjudicated causal gene; Locus, locus named by the reported or closest gene). SingleBrain: single-nucleus eQTL colocalization with AD GWAS (PP4 > 0.8) with (Coloc+MR) or without (Coloc) Mendelian randomization support[38]. snTWAS: single-nucleus TWAS association with AD[39] that is Bonferroni-significant within a cellular population (Bonf.) or FDR-significant only (FDR). Per-gene values are in Supplementary Data 1.

Supplementary References

[85] R. Koenker. *Quantile Regression*. Econometric Society Monographs. Cambridge: Cambridge University Press, 2005. DOI: 10.1017/CBO9780511754098.

[86] J. L. Powell. “Estimation of monotonic regression models under quantile restrictions”. In: *Nonparametric and Semiparametric Methods in Econometrics and Statistics*. Ed. by W. A. Barnett, J. L. Powell, and G. Tauchen. Cambridge: Cambridge University Press, 1991, pp. 357–384.

[87] H. Xie, X. Cao, S. Zhang, and Q. Sha. “Joint analysis of multiple phenotypes for extremely unbalanced case–control association studies using multi-layer network”. In: *Bioinformatics* 39.12 (2023), btad707. DOI: 10.1093/bioinformatics/btad707.

[88] R. M. Baron and D. A. Kenny. “The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations.” In: *Journal of personality and social psychology* 51.6 (1986), pp. 1173–1182. DOI: 10.1037/0022-3514.51.6.1173.