

Additional file 1 to “Genetic prioritization of *HSDL2* in suicide and intentional self-harm: multi-omics evidence linking lipid metabolism and neuronal vulnerability”

Ziyan Zhang^{1,3#}, Yutong Wu^{2,3#}, Fanxu Meng^{3,4}, Hanbing Ji^{1,3}, Hai Zhong⁵, Shucheng Si⁶, Yun Wei^{2,3}, Yilei Ge^{2,3}, Sijia Wu^{2,3}, Qingxin Luo^{2,3}, Le Wang^{1,3}, Tiemei Liu^{1,3}, Jiawei Xiu^{1,3}, Yiman Guo^{1,3}, Yue Li^{1,3}, Lei Hou^{2,3}, Hao Chen^{1,3}, Xiaoru Sun^{2,3}, Yuanyuan Yu^{2,3}, Qingzhen Hou^{1,3}, Shanshan Gao^{7*}, Fuzhong Xue^{2,3,8*}, Hongkai Li^{1,3*}

#Ziyan Zhang and Yutong Wu contributed equally to this work.

*Corresponding authors: Shanshan Gao, Fuzhong Xue and Hongkai Li

Author affiliations:

1. Department of Biostatistics, School of Public Health, Cheeloo College of Medicine, Shandong University, Jinan, 250012, China
2. Department of Medical Dataology, School of Public Health, Cheeloo College of Medicine, Shandong University, Jinan, 250012, China
3. National Institute of Health and Medical Big Data, Jinan, 250003, China
4. Department of Computer Science, Vrije Universiteit Amsterdam, 1081HV, Amsterdam, The Netherlands
5. Department of Radiology, the Second Qilu Hospital of Shandong University, Shandong University, Jinan, 250063, China
6. Research Center of Clinical Epidemiology, Peking University Third Hospital, Beijing, 100191, China
7. Center for Reproductive Medicine, the Second Hospital, Cheeloo College of Medicine, Shandong University, Jinan, Shandong, 250012, China
8. Qilu Hospital of Shandong University, Jinan, 250012, China

Corresponding author:

1. Hongkai Li,

E-mail: lihongkaiyouxiang@163.com

ORCID: 0000-0003-1848-937X

Address: School of Public Health, Shandong University, 44 Wenhua West Road, Jinan, Shandong province, China

2. Fuzhong Xue,

E-mail: xuefzh@sdu.edu.cn

ORCID: 0000-0003-0378-7956

Address: School of Public Health, Shandong University, 44 Wenhua West Road, Jinan, Shandong province, China

3. Shanshan Gao

E-mail: 202162019102@email.sdu.edu.cn

ORCID: 0000-0003-4355-8365

Address: Center for Reproductive Medicine, the Second Hospital, Cheeloo College of Medicine, Shandong University, Jinan, Shandong, 250012, China

Contents

Supplementary Note 1: Outcome definition.....	1
1.1 Primary outcome: FinnGen suicide or other intentional self-harm	1
1.2 Secondary outcome: PGC suicide attempt	1
1.3 Secondary outcome: UK Biobank ever attempted suicide	1
Supplementary Note 2: Molecular QTL resources and instrument construction	2
2.1 Cis-eQTL resources and thresholds (eQTLGen; GTEx v8)	2
2.2 Cis-pQTL resources and thresholds (UKB-PPP; ARIC)	2
2.3 Psychiatric liability adjustment and exploratory metabolomic mediation	3
2.4 Metabolomics GWAS and thresholds (CLSA)	4
2.5 Single-cell eQTL resources and thresholds	4
2.6 PheWAS summary statistics (UKB SAIGE)	5
2.7 Sample-overlap and population-structure sensitivity analyses	5
2.8 Protein knowledge resources (UniProtKB)	5
Supplementary Note 3: <i>HSDL2</i> Protein Sequence, Domain Architecture, and Structural Analyses ..	6
3.1 <i>HSDL2</i> protein sequence information	6
3.2 Molecular Docking	6
3.3 Structure-Based Virtual Screening	6
3.4 Molecular Dynamic Simulations	2
3.5 Comparative SDR-family pocket analysis	2
Supplementary Figures	4
Fig. S1. TWMR and PWMR results of TOP 55 prioritized genes.	4
Fig. S2. GO Cell Component enrichment analysis.	5
Fig. S3. Manhattan plot of Phenome-wide association results for <i>APP</i>	6
Fig. S4. Manhattan plot of Phenome-wide association results for <i>SI00A11</i>	7
Fig. S5. Molecular Docking Results of Cholic Acid in <i>HSDL2</i> SCP2 Domain.	8
Fig. S6. Molecular Docking Results of Benzylcinnamate in <i>HSDL2</i> SCP2 Domain.	9
Fig. S7. Molecular Docking Results of Benzyl4hydroxycinnamate in <i>HSDL2</i> SCP2 Domain.	10
Fig. S8. Molecular Docking Results of Taurine in <i>HSDL2</i> SCP2 Domain.	11
Fig. S9. Molecular Docking Results of Benzylcinnamate in <i>HSDL2</i> SDR Domain.	12
Fig. S10. Molecular Docking Results of Benzyl4hydroxycinnamate in <i>HSDL2</i> SDR Domain.	13
Fig. S11. Molecular Docking Results of Taurine in <i>HSDL2</i> SDR Domain.	14
Fig. S12. Distribution of docking scores for FDA-approved compounds.	15
Fig. S13. Molecular dynamics simulation of the cholic acid-bound <i>HSDL2</i> complex.	16
Fig. S14. Molecular dynamics simulation of the taurine-bound <i>HSDL2</i> complex.	17
Fig. S15. Comparative structural pocket analysis of <i>HSDL2</i> and representative SDR-family proteins.	18
Fig. S16. Balanced accuracy of XGBoost models with different feature-integration strategies.	19
Fig. S17. Sensitivity of XGBoost models with different feature-integration strategies.	20
Fig. S18. Specificity of XGBoost models with different feature-integration strategies.	21
Fig. S19. Negative predictive value of XGBoost models with different feature- integration strategies.	22
Fig. S20. Positive predictive value of XGBoost models with different feature- integration	

strategies.	23
Fig. S21. Precision-recall area under the curve of XGBoost models with different feature-integration strategies.	24
Fig. S22. Youden index of XGBoost models with different feature-integration strategies.	25
Fig. S23. SHAP-based feature contribution analysis for the proteomic baseline classification model.	26
Fig. S24. Sample-overlap sensitivity analysis for the UKBPPP- <i>HSDL2</i> -FinnGen MR association.	27
Fig. S25. Global sample-overlap sensitivity analysis across UKB-PPP-FinnGen MR comparisons.	28
Fig. S26. Overlap of FDR-significant TWMR genes across SSH, major depressive disorder and bipolar disorder.	29

Supplementary Note 1: Outcome definition

1.1 Primary outcome: FinnGen suicide or other intentional self-harm

Suicide or other intentional self-harm:

The primary outcome was suicide or other intentional self-harm (SSH) from FinnGen release 12, defined using the endpoint VWXY20_SUICI_OTHER_INTENTI_SELF_H. This endpoint captures both fatal and non-fatal intentional self-harm events identified through Finnish nationwide health registers, including hospital discharge records and cause-of-death registers. Cases were defined using ICD-10 codes X60-X84, corresponding to intentional self-harm, together with equivalent historical ICD-9 and ICD-8 codes when applicable under the original FinnGen phenotype definition. The FinnGen R12 SSH genome-wide association study included 11,538 cases and 488,810 controls. Controls were individuals without recorded diagnostic codes corresponding to intentional self-harm or suicide-related events. Because this endpoint integrates fatal suicidal events and non-fatal self-harm presentations, we treated it as a broad composite SSH phenotype. This broad definition increases statistical power but may also introduce phenotypic heterogeneity; therefore, additional suicidal-behaviour-related outcomes were analysed as sensitivity phenotypes.

1.2 Secondary outcome: PGC suicide attempt

We additionally evaluated suicide attempt using summary statistics from the Psychiatric Genomics Consortium (PGC). This outcome was based on genome-wide association data for suicide attempt, generally defined as a history of self-injurious behaviour with intent to die. Unlike register-based FinnGen phenotypes, PGC suicide attempt data integrate psychiatric and cohort-based sources, and therefore may differ in ascertainment strategy, clinical context and participant composition. The PGC suicide attempt outcome was included to assess whether the primary SSH-associated transcriptomic signals were also observed in an independent psychiatric-genetic framework. Because this dataset differs from FinnGen in phenotype ascertainment, sample composition and study design, concordance across outcomes was interpreted as supportive evidence for shared suicidal-behaviour-related genetic architecture, whereas discordance was not taken as evidence against the primary SSH association.

1.3 Secondary outcome: UK Biobank ever attempted suicide

We further used a UK Biobank suicide attempt phenotype based on self-reported lifetime history of attempted suicide and/or linked health-record information when available. This phenotype reflects whether an individual had ever attempted suicide before or during UK Biobank ascertainment, depending on the source and timing of the record. This outcome differs from the primary FinnGen SSH endpoint in two major respects. First, it is more strongly influenced by self-report and questionnaire-based ascertainment. Second, it captures lifetime suicide-attempt history rather than the broader composite of fatal and non-fatal intentional self-harm events used in FinnGen. We therefore used this phenotype as an additional sensitivity outcome to examine whether the *HSDL2*-related and transcriptome-wide signals were robust to a self-report-enriched suicidal behaviour definition

Supplementary Note 2: Molecular QTL resources and instrument construction

2.1 Cis-eQTL resources and thresholds (eQTLGen; GTEx v8)

For cis-eQTL data, we utilized the eQTLGen Consortium datasets and the Genotype-Tissue Expression Project version 8 (GTEx v8) and which provide large-scale summary statistics for gene expression across multiple tissues and blood samples.

eQTLGen Consortium: This consortium focuses specifically on blood-based gene expression. Its data are derived from a meta-analysis of 37 independent cohorts, comprising 31,684 individuals. The dataset includes cis-eQTL summary statistics for 19,250 genes, representing one of the largest resources for genetic regulation of expression in human blood. Our analysis employed its cis-eQTL findings. We selected strong genetic instrumental variables (IVs) restricted to cis-variants (± 1 Mb), adopting the significance threshold recommended by the eQTLGen Consortium ($P < 2.02 \times 10^{-5}$), followed by LD clumping ($r^2 < 0.001$, window = 100 kb).

Genotype-Tissue Expression Project (GTEx v8): The GTEx project provides a broad survey of gene expression across multiple human tissues. The version 8 (V8) release contains data from 838 post-mortem donors, with RNA-seq from 54 tissue sites and 2 cell lines, totaling 17,382 samples. It serves as a foundational resource for understanding tissue-specific genetic regulation of expression. Our study leveraged the cis-eQTL summary statistics generated from this diverse healthy tissue collection. To maximize power genome-wide and refer to previous relevant studies, we set the cis-eQTL screening threshold at $P < 1 \times 10^{-5}$ which corresponded to the corrected significance level across all gene-tissue pairs. And we set LD clumping ($r^2 < 0.001$, window = 100 kb).

We additionally screened all 13 GTEx v8 bulk brain tissues for *HSDL2* eQTL evidence. Although *HSDL2* was tested in these tissues, no brain tissue reached the prespecified GTEx significant eGene threshold ($qval < 0.05$), precluding robust instrument construction for GTEx brain tissue-specific MR;

2.2 Cis-pQTL resources and thresholds (UKB-PPP; ARIC)

For cis-pQTL data, we used two population-based resources: the UK Biobank Pharma Proteomics Project (UKB-PPP) and the Atherosclerosis Risk in Communities (ARIC) study, both offering genome-wide associations between genetic variants and circulating protein levels.

UK Biobank Pharma Proteomics Project (UKB-PPP): This project provides large-scale plasma proteomic data generated through collaboration between the UK Biobank and multiple biopharmaceutical companies. It comprises protein quantitative trait loci (pQTL) analyses for 2,923 plasma proteins measured using the Olink Explore 3072 platform in 54,219 UK Biobank participants. It represents one of the largest and most comprehensive pQTL resources to date, reporting 14,287 primary genetic associations, the majority of which were novel discoveries at the time of publication. The data also include secondary, independent associations, significantly expanding the catalog of genetic instruments for protein levels. We select cis-pQTLs screening threshold at $P < 5 \times 10^{-8}$ and LD clumping ($r^2 < 0.001$, window = 100 kb).

Atherosclerosis Risk in Communities (ARIC) Study: The ARIC study provides pQTL data derived from a multi-ethnic cohort. The analysis includes plasma protein measurements from 7,213 European American and 1,871 African American individuals, identifying associations for 2,004 and 1,618

proteins in the respective ancestry groups. This resource offers valuable insights into both shared and population-specific genetic regulation of the plasma proteome. We select cis-pQTLs screening threshold at $P < 5 \times 10^{-8}$ and LD clumping ($r^2 < 0.001$, window = 100 kb).

2.3 Psychiatric liability adjustment and exploratory metabolomic mediation

Direct *HSDL2* estimates are reported with beta estimates, odds ratios, 95% confidence intervals, P values, FDR-adjusted P values and conditional F statistics. Conditional F statistics were used to evaluate instrument strength in the multivariable setting.

Because suicidal behaviour is genetically and clinically related to psychiatric disorders, we performed multivariable MR analyses to evaluate whether the *HSDL2*-SSH association was explained by broader psychiatric liability. *HSDL2* and each psychiatric trait were included together in separate MVMR-IVW models. Psychiatric traits included major depressive disorder, bipolar disorder, schizophrenia, post-traumatic stress disorder, substance use disorder, anxiety disorder, attention-deficit/hyperactivity disorder, borderline personality disorder and anorexia nervosa. An additional model adjusted for a cross-disorder psychiatric liability factor. The psychiatric GWAS sources correspond to Additional file 1 Supplementary Table S2-1.

Supplementary Table S2-1. Psychiatric GWAS summary statistics used for multivariable Mendelian randomization adjustment.

Trait ID in analysis	Psychiatric phenotype	Dataset or release	DOI	Ancestry or phenotype definition	Sample size reported for the source or selected release
MDD_noUK BB	Major depressive disorder	PGC MDD2025	10.6084/m9.figshare.24204843	European ancestry; no 23andMe; no UK Biobank public subset	Source publication: 688,808 cases and 4,364,225 controls across ancestries; selected file contains per-SNP NCAS/NCON fields rather than one fixed sample size
BIP_noUKB	Bipolar disorder	PGC BIP2024	10.6084/m9.figshare.27216117	European ancestry; no 23andMe; no UK Biobank public subset	Source publication: 158,036 cases and approximately 2.8 million controls across ancestries; selected file is the EUR no-UKB/no-23andMe release
SCZ_EUR	Schizophrenia	PGC3 SCZ2022	10.6084/m9.figshare.19426775	European ancestry; autosomal public summary statistics	European ancestry public file: 53,386 cases and 77,258 controls
PTSD_EUR	Post-traumatic stress disorder	PGC PTSD2024	10.6084/m9.figshare.26349322	European ancestry case-control PTSD liability	European ancestry meta-analysis: 137,136 cases and 1,085,746 controls
SUD_EUR	Substance use disorder / addiction liability	PGC SUD2023	10.6084/m9.figshare.24268882	European ancestry addiction factor from multivariate genomic SEM	Not a single case-control GWAS; European analysis included 1,025,550 individuals

ANX	Anxiety disorders	PGC ANX2026	10.6084/m9.figshare.31 389910	European ancestry anxiety disorder GWAS; UTAH cohort removed in this public release	Selected release: 122,083 cases and 729,602 controls; published analysis: 122,341 cases and 729,881 controls
	Attention-deficit/hyperactivity disorder	PGC/iPSYC H/deCODE ADHD2022	10.6084/m9.figshare.22 564390	European ancestry meta-analysis	38,691 cases and 186,843 controls
BPD_noUK BB	Borderline personality disorder	PGC BPD2025	10.6084/m9.figshare.30 777305	European ancestry; no UK Biobank public subset	Selected autosomal no-UKBB release: 12,157 cases and 1,039,897 controls
AN	Anorexia nervosa	PGC AN2019	10.6084/m9.figshare.14 671980	European ancestry anorexia nervosa GWAS	16,992 cases and 55,525 controls
P_FACTOR	Cross-disorder psychiatric p factor	PGC CDG2025	10.6084/m9.figshare.30 359017	Cross-disorder latent psychiatric liability factor	Not a single case-control trait; PGC CDG3 analysis spans 14 psychiatric disorders and 1,056,201 cases

2.4 Metabolomics GWAS and thresholds (CLSA)

The Canadian Longitudinal Study on Aging (CLSA) comprises 51338 participants aged 45–85 years at baseline from all 10 Canadian provinces. Plasma metabolomics profiling was conducted using the Metabolon DiscoveryHD4 platform. GWAS of 1091 metabolites were performed in 8192 individuals of European ancestry. Summary-level statistics for these metabolites were obtained from the GWAS Catalog, corresponding to accession numbers GCST90199621–GCST90201000. We selected IVs using a significance threshold of $P < 5 \times 10^{-6}$ and stringent LD clumping ($r^2 < 0.01$, window = 5000 kb).

2.5 Single-cell eQTL resources and thresholds

To evaluate whether the *HSDL2*-SSH association showed brain cell-type-specific regulatory patterns, we used two independent single-nucleus brain eQTL resources. The Bryois et al. dataset was used as the discovery resource, and the SingleBrain resource was used as the validation resource.

For the discovery dataset, Bryois et al performed eQTL analysis using single-nuclei RNA-seq on brain samples from 192 genotyped donors and finally identified 7,607 cis-eQTL genes in total. The single-cell eQTL dataset contains eQTL summary statistics for all SNPs-gene pairs in 8 major brain cell types (Astrocytes, Endothelial cells, Excitatory neurons, Inhibitory neurons, Microglia, Oligodendrocytes, Oligodendrocyte precursor cells and Pericytes) derived from the prefrontal cortex, temporal cortex and white matter. The full eQTL summary statistics were publicly obtained from the Zenodo repository (accession no. 7276971). Due to small sample of single-cell data, we relaxed the cis-eQTL screening threshold at $P < 0.005$ and LD clumping ($r^2 < 0.001$, window = 100 kb).

For the validation analysis, we used SingleBrain, a large-scale single-nucleus eQTL meta-analysis generated by Jang et al and published online on 19 March 2026. SingleBrain integrated single-nucleus RNA-seq and genotype data from four independent human brain cohorts and included 5.8 million nuclei from 983 donors of European ancestry. This resource provides cis-eQTL summary statistics across seven major brain cell types and 28 subtypes, including astrocytes, endothelial cells, excitatory neurons, inhibitory neurons, microglia, oligodendrocytes and oligodendrocyte progenitor cells.

SingleBrain summary statistics were obtained from Zenodo (accession no. 14908182). In the validation analysis, we selected *HSDL2* cis-eQTL instruments using a genome-wide significance threshold of $P < 5 \times 10^{-8}$, followed by LD clumping at $r^2 < 0.001$ within a 100 kb window.

Cell-type labels were harmonized across the discovery and validation resources. Because pericyte eQTLs were available in the Bryois discovery dataset but not among the seven major SingleBrain validation cell types, cross-resource validation focused on comparable major cell types shared by both resources. For each cell type, Mendelian randomization estimates were generated using IVW when multiple instruments were available and Wald ratio when only one instrument was available. Robustness was assessed using heterogeneity, horizontal pleiotropy, Steiger directionality and leave-one-out analyses when instrument number permitted. Bayesian colocalization was then used to evaluate whether cell-type-specific *HSDL2* regulatory signals and the FinnGen SSH association shared local genetic architecture. For the SingleBrain validation analyses, pairwise conditional colocalization was additionally applied to relax the single-causal-variant assumption and to assess whether the cell-type-specific signal persisted after conditioning on local association structure.

2.6 PheWAS summary statistics (UKB SAIGE)

Zhou et al introduced the Scalable and Accurate Implementation of Generalized Mixed Model (SAIGE), a method designed to efficiently analyze large-scale genetic data while addressing challenges such as unbalanced case-control ratios and sample relatedness. The authors applied SAIGE to analyze ~1,400 UKBiobank Phecode binary phenotypes (400,000 White British samples) on 28 million imputed variants, resulting in a comprehensive repository of summary statistics for these traits. The UKB SAIGE GWAS database constitutes a valuable resource for the global genetics community and significantly advances the understanding of the genetic architecture underlying human diseases.

2.7 Sample-overlap and population-structure sensitivity analyses

Potential sample overlap between exposure and outcome datasets was evaluated using overlap-adjusted IVW sensitivity analyses under hypothetical overlap scenarios. We varied assumed exposure-outcome overlap proportions and phenotypic correlations to estimate whether sample overlap could materially change the *HSDL2* MR estimate. Additional simulations estimated overlap-induced bias and type I error inflation under different assumptions regarding observational bias.

To assess potential population-structure-related bias, we examined residual inflation in the FinnGen SSH GWAS using LD score regression diagnostics and performed allele-frequency-based sensitivity analyses. MR estimates were recalculated after excluding instruments with large allele-frequency discrepancies between exposure and outcome datasets or low minor allele frequency in FinnGen. These analyses were used to evaluate whether the *HSDL2* association was robust to population-specific allele-frequency differences and possible residual stratification.

2.8 Protein knowledge resources (UniProtKB)

Protein sequences and related functional annotations information were retrieved from the UniProt Knowledgebase (UniProtKB). UniProtKB is a comprehensive, high-quality, and freely accessible central resource for protein sequence and functional information, which integrates high-quality manually reviewed entries (UniProtKB/Swiss-Prot) and automatically annotated (UniProtKB/TrEMBL) data.

Supplementary Note 3: *HSDL2* Protein Sequence, Domain Architecture, and Structural Analyses

3.1 *HSDL2* protein sequence information

The full-length protein sequence of *HSDL2* was obtained from the UniProt database (UniProt ID: Q6YN16, *HSDL2_HUMAN*). Sequence Annotation for PDB 3KVO: The structural model corresponds to the SDR domain (red). The full sequence provided is: (MGSSHHHHHHSSGLVPRGSHMASMTGGQQMGRGSAMLPTGRLAGCTVFITGASRGIGKAIKALKAADGANIVIAAKTAQPHPKLLGTIYTAAEEIEAVGGKALPCIVDVRDEQQISA AVEKAIKKFGGIDILVNNASAI SLTNTLDTPTKRLDLMMNVNTRGTYLASKACIPYLKKS KVAHILNISPLNLPVWFKQHCA YTI AKYGMSMYVLGMAEEFKGEIAVNALWPKTAIHTAAMDMLGGPGIESQCRKVDIIADAAYSIFQKPKSFTGNFVIDENILKEEGIENFDVYAIKPGHPLQPDFFLDEYPEAVSKKVESTGAVPELACGRTRAPPPPLRSGC). The N-terminal segment (MGSSHHHHHHSSGLVPR GSHMASMTGGQQMGRGSA) represents a standard hexahistidine (His6) tag used for Ni-NTA affinity purification, while the C-terminal segment (LACGRTRAPPPPLRSGC) constitutes an artificial tag.

A high-precision 3D structural model was successfully constructed using AlphaFold3, with pLDDT >90, indicating high model reliability. Structural analysis revealed that the *HSDL2* protein is primarily composed of two distinct domains. The N-terminal region (residues 1-293) constitutes a Short-chain Dehydrogenase/Reductase (SDR) domain, which adopts a classic Rossmann fold (α/β structure) and accounts for approximately two-thirds of the protein. This domain contains the active site motif YxxxK (residues 130-134)/(168-172) and a glycine-rich region (TGASRGIG, residues 16-23). In contrast, the C-terminal region (residues 294-418) is an α -helix-rich SCP2 domain.

3.2 Molecular Docking

The AlphaFold 3-predicted *HSDL2* structure was used for molecular docking analyses. Protein preparation was performed using AutoDockTools (v1.5.7), including addition of polar hydrogens and assignment of docking-compatible parameters. Molecular docking was conducted using AutoDock Vina (v1.2.7). Candidate ligands were docked into the predefined binding pocket within the SDR-domain region of *HSDL2*.

Docking poses were clustered using an RMSD threshold of <2.0 Å. For each ligand, the representative pose with the lowest predicted binding affinity from the largest cluster was selected for downstream visualization and interpretation. Protein-ligand interaction maps were generated using PyMOL (v3.1.6.1) and BIOVIA Discovery Studio 2019. Docking scores were interpreted as computational estimates of binding plausibility rather than direct evidence of physiological binding.

3.3 Structure-Based Virtual Screening

To further assess the ligand-binding capacity of the predicted *HSDL2* pocket, structure-based virtual screening was performed using a curated library of FDA-approved compounds from DrugBank. After preprocessing, 2,504 compounds were retained for docking analysis. All compounds were docked into the same predefined *HSDL2* binding pocket used in the candidate ligand docking analysis.

Docking simulations were performed using AutoDock Vina with grid parameters centered on the predicted cavity: center_x = 30.012, center_y = -13.621, and center_z = 69.543. The resulting docking scores were used to rank compounds by predicted binding affinity. Because virtual screening is based on static structural docking, these results were considered hypothesis-generating and require experimental validation.

3.4 Molecular Dynamic Simulations

All-atom molecular dynamics simulations were performed to evaluate the dynamic stability of *HSDL2*-ligand complexes and to address the limitation of static AlphaFold 3 and docking models. Simulations were conducted using GROMACS 2024.3 for taurine-bound and bile acid-bound *HSDL2* systems. For each system, a 50-ns production simulation was performed after energy minimization and equilibration.

Ligand topologies were generated using the ACPYPE interface to Antechamber with the GAFF2 force field and AM1-BCC partial charges. Each protein-ligand complex was solvated in a cubic box with a minimum distance of 0.8 nm between the protein surface and the box edge, using the TIP3P water model. Counterions were added to neutralize the system.

Energy minimization was performed using the steepest descent algorithm until the maximum force was below 1000 kJ mol⁻¹ nm⁻¹. The system was then equilibrated in two steps: NVT equilibration at 300 K followed by NPT equilibration at 1 bar. Heavy-atom position restraints were applied during equilibration. Temperature coupling was performed using the V-rescale thermostat, and pressure coupling was performed using the Parrinello-Rahman barostat. Periodic boundary conditions were applied in all directions. Long-range electrostatic interactions were treated using the particle mesh Ewald method, and covalent bonds involving hydrogen atoms were constrained using the LINCS algorithm.

Trajectory analyses were performed on the production trajectories after equilibration. Structural stability was assessed using backbone RMSD, radius of gyration, residue-level RMSF, and protein-ligand distance profiles.

3.5 Comparative SDR-family pocket analysis

Because *HSDL2* contains a conserved short-chain dehydrogenase/reductase (SDR) domain, potential target specificity and off-target binding represent important considerations for downstream translational interpretation. To provide preliminary structural context, we compared *HSDL2* with representative human SDR-family proteins, including *HSD17B6* and *CBR3* (Additional file 1 Supplementary Table S2-2, Fig. S15). Although these proteins share SDR-domain architecture, they differ in catalytic activity, substrate preference, cellular localization and biological function. *HSDL2* is annotated as an orphan hydroxysteroid dehydrogenase-like protein with potential roles in lipid- and bile-acid-related metabolism, whereas *HSD17B6* is involved in broader steroid metabolism and *CBR3* functions primarily as a carbonyl/quinone reductase.

We further performed structural pocket analysis based on the predicted 3D structures of *HSDL2* and comparator SDR-family proteins (Supplementary Fig. S15 and Table 2). *HSDL2* showed a substantially smaller predicted pocket volume than *HSD17B6* and *CBR3*, suggesting a more structurally constrained binding environment. This smaller pocket may indicate a narrower ligand-accommodation space and provides preliminary structural support for potential specificity differences among SDR-family members. These findings should be interpreted as

hypothesis-generating structural evidence rather than direct proof of pharmacological selectivity. Although the smaller predicted *HSDL2* pocket may reduce the likelihood of accommodating a broad range of small molecules, it does not exclude off-target interactions or shared ligand classes across SDR-family proteins. Experimental selectivity profiling and functional perturbation assays will therefore be required to determine whether *HSDL2* can be selectively modulated in a therapeutically relevant context.

Supplementary Table S2-2. Comparison of representative SDR family proteins

Feature	<i>HSDL2</i> (Q6YN16)	<i>HSD17B6</i> (O14756)	<i>CBR3</i> (O75828)
Catalytic activity	No typical steroid dehydrogenase activity	Broad steroid/retinol activity	Carbonyl/quinone reductase activity
Subcellular localization	Peroxisome, mitochondria	ER/microsome membranes	Cytoplasm
Functional implication	Lipid and bile acid metabolism	Steroid metabolism	enobiotic metabolism

Supplementary Figures




Fig. S1. TWMR and PWMR results of TOP 55 prioritized genes.

Forest plot showing the causal estimates (OR and 95% CI) for the top 55 prioritized genes derived from TWMR (blue) and PWMR (orange) analyses.




Fig. S2. GO Cell Component enrichment analysis.

The dot plot displays significantly enriched pathways grouped by category, where the size of each point represents the gene count and the color gradient indicates statistical significance (FDR).




Fig. S3. Manhattan plot of Phenome-wide association results for *APP*.

Ordinate representation of the p value in phenome-wide MR results of *APP*. A dot represents a disease trait. The x-axis organizes phenotypes into 17 disease categories.




Fig. S4. Manhattan plot of Phenome-wide association results for *S100A11*.

Ordinate representation of the p value in phenome-wide MR results of *S100A11*. A dot represents a disease trait. The x-axis organizes phenotypes into 17 disease categories.




Fig. S5. Molecular Docking Results of Cholic Acid in *HSDL2* SCP2 Domain.

The 3D visualization shows the interaction between Cholic Acid (yellow stick) and the SCP2 domain of *HSDL2* (cyan ribbon). Key hydrogen bonds are indicated by green dashed lines.




Fig. S6. Molecular Docking Results of Benzylcinnamate in *HSDL2* SCP2 Domain.

The 3D visualization shows the interaction between Benzylcinnamate (yellow stick) and the SCP2 domain of *HSDL2* (cyan ribbon). Key hydrogen bonds are indicated by green dashed lines.




Fig. S7. Molecular Docking Results of Benzyl4hydroxycinnamate in *HSDL2* SCP2 Domain.

The 3D visualization shows the interaction between Benzyl4hydroxycinnamate (yellow stick) and the SCP2 domain of *HSDL2* (cyan ribbon). Key hydrogen bonds are indicated by green dashed lines.




Fig. S8. Molecular Docking Results of Taurine in *HSDL2* SCP2 Domain.

The 3D visualization shows the interaction between Taurine (yellow stick) and the SCP2 domain of *HSDL2* (cyan ribbon). Key hydrogen bonds are indicated by green dashed lines.




Fig. S9. Molecular Docking Results of Benzylcinnamate in *HSDL2* SDR Domain.

The 3D visualization shows the interaction between Benzylcinnamate (yellow stick) and the SDR domain of *HSDL2* (cyan ribbon). Key hydrogen bonds are indicated by green dashed lines.




Fig. S10. Molecular Docking Results of Benzyl4hydroxycinnamate in *HSDL2* SDR Domain.

The 3D visualization shows the interaction between Benzyl4hydroxycinnamate (yellow stick) and the SDR domain of *HSDL2* (cyan ribbon). Key hydrogen bonds are indicated by green dashed lines.




Fig. S11. Molecular Docking Results of Taurine in *HSDL2* SDR Domain.

The 3D visualization shows the interaction between Taurine (yellow stick) and the SDR domain of *HSDL2* (cyan ribbon). Key hydrogen bonds are indicated by green dashed lines.




Fig. S12. Distribution of docking scores for FDA-approved compounds.

Structure-based virtual screening of 2,504 FDA-approved compounds against the predicted SDR-domain pocket of *HSDL2*. The histogram shows the distribution of predicted AutoDock Vina docking scores, reported in kcal/mol. The dashed vertical line indicates the docking score of cholic acid, which was included as a pathway-relevant reference ligand. Most FDA-approved compounds showed docking scores clustered around approximately -6 to -5 kcal/mol, whereas cholic acid showed a more favorable predicted binding score than most screened compounds. Docking scores were interpreted as computational estimates of binding plausibility and do not provide direct evidence of physiological binding or pharmacological activity.




Fig. S13. Molecular dynamics simulation of the cholic acid-bound *HSDL2* complex.

Fifty-nanosecond all-atom molecular dynamics simulation of the *HSDL2*-cholic acid complex. A, Backbone root-mean-square deviation (RMSD) over simulation time, showing that the global protein conformation reached a stable trajectory after initial equilibration. B, Residue-level root-mean-square fluctuation (RMSF), indicating that most residues remained relatively rigid, with higher flexibility mainly observed at terminal regions. C, Radius of gyration (Rg), showing maintenance of global compactness during the production simulation. D, Center-of-mass (COM) distance between cholic acid and *HSDL2*, indicating that cholic acid remained at a larger but relatively stable protein-ligand distance, consistent with a peripheral or looser binding mode. RMSD, RMSF, Rg and COM distance are reported in nanometres.




Fig. S14. Molecular dynamics simulation of the taurine-bound *HSDL2* complex.

Fifty-nanosecond all-atom molecular dynamics simulation of the *HSDL2*-taurine complex. A, Backbone root-mean-square deviation (RMSD) over simulation time, showing stabilization of the global protein conformation after initial equilibration. B, Residue-level root-mean-square fluctuation (RMSF), indicating that most residues remained relatively stable, with increased flexibility mainly localized to terminal regions. C, Radius of gyration (Rg), showing preservation of global protein compactness during the production simulation. D, Center-of-mass (COM) distance between taurine and *HSDL2*, showing that taurine maintained a shorter and relatively stable protein-ligand distance throughout most of the simulation, consistent with closer association with the predicted binding pocket. RMSD, RMSF, Rg and COM distance are reported in nanometres.




Fig. S15. Comparative structural pocket analysis of *HSDL2* and representative SDR-family proteins.

Predicted pocket-volume comparison among *HSDL2* and representative human short-chain dehydrogenase/reductase (SDR)-family proteins, including *HSD17B6* and *CBR3*. *HSDL2* showed a smaller predicted pocket volume than *HSD17B6* and *CBR3*, suggesting a more structurally constrained binding environment. This analysis provides preliminary structural context for potential ligand-accommodation differences among SDR-family proteins, but should not be interpreted as direct evidence of target selectivity or absence of off-target binding.




Fig. S16. Balanced accuracy of XGBoost models with different feature-integration strategies.

Box plots show bootstrap-resampled balanced accuracy estimates for baseline SSH classification in the held-out UKB-PPP test set. Models compared included demographic covariates only, *HSDL2* plus demographic covariates, top 3 MR-prioritized proteins plus demographic covariates, top 55 MR-prioritized proteins plus demographic covariates, and top 55 proteins excluding *HSDL2* plus demographic covariates. Each point represents one bootstrap resample. The *HSDL2* plus demographic model showed the highest balanced accuracy among the compared feature sets.




Fig. S17. Sensitivity of XGBoost models with different feature-integration strategies.

Box plots show bootstrap-resampled sensitivity estimates for baseline SSH classification in the held-out UKB-PPP test set. Sensitivity was calculated using the classification threshold selected in the training set by the Youden index and then applied unchanged to the test set. The *HSDL2* plus demographic model showed higher sensitivity than the demographic-only model and the other protein-integrated models.




Fig. S18. Specificity of XGBoost models with different feature-integration strategies.

Box plots show bootstrap-resampled specificity estimates for baseline SSH classification in the held-out UKB-PPP test set. Specificity was evaluated using the training-derived Youden threshold. Compared with the demographic-only model, the *HSDL2* plus demographic model showed lower specificity, consistent with its higher sensitivity and the expected trade-off between sensitivity and specificity under the selected threshold.




Fig. S19. Negative predictive value of XGBoost models with different feature-integration strategies.

Box plots show bootstrap-resampled negative predictive value estimates for baseline SSH classification in the held-out UKB-PPP test set. All models showed high negative predictive values, reflecting the low baseline prevalence of SSH in the UKB-PPP analytic cohort. The *HSDL2* plus demographic model showed slightly higher negative predictive value than the demographic-only model.




Fig. S20. Positive predictive value of XGBoost models with different feature- integration strategies.

Box plots show bootstrap-resampled positive predictive value estimates for baseline SSH classification in the held-out UKB-PPP test set. Positive predictive values remained low across all models, consistent with the low event proportion and substantial case-control imbalance in baseline SSH status. These results indicate that the models should be interpreted as exploratory baseline discrimination tools rather than standalone clinical classifiers.




Fig. S21. Precision-recall area under the curve of XGBoost models with different feature-integration strategies.

Box plots show bootstrap-resampled precision-recall area under the curve (PR-AUC) estimates for baseline SSH classification in the held-out UKB-PPP test set. PR-AUC was used as a prevalence-sensitive performance metric because SSH cases were rare in the analytic cohort. Protein-integrated models, particularly the *HSDL2* plus demographic and top 3 protein plus demographic models, showed higher PR-AUC than the demographic-only model.




Fig. S22. Youden index of XGBoost models with different feature-integration strategies.

Box plots show bootstrap-resampled Youden index estimates for baseline SSH classification in the held-out UKB-PPP test set. The Youden index summarizes the combined contribution of sensitivity and specificity and was also used to define the training-derived classification threshold. The *HSDL2* plus demographic model showed the highest Youden index among the compared feature sets.




Fig. S23. SHAP-based feature contribution analysis for the proteomic baseline classification model.

SHAP summary plot showing the top contributing features in the XGBoost baseline SSH classification model incorporating MR-prioritized proteins and demographic covariates. Each point represents an individual participant, with the x-axis indicating the SHAP value and color representing the standardized feature value. Positive SHAP values indicate features contributing toward higher predicted SSH probability, whereas negative values indicate features contributing toward lower predicted SSH probability. Demographic and clinical variables contributed strongly to model output, while MR-prioritized proteins including *S100A11* and *HSDL2* also contributed to the model. SHAP values describe model contribution and should not be interpreted as evidence of causal effects.




Fig. S24. Sample-overlap sensitivity analysis for the UKBPPP-*HSDL2*-FinnGen MR association.

Sample-overlap sensitivity analysis for the Mendelian randomization association between genetically predicted *HSDL2* and FinnGen R12 suicide or other intentional self-harm. Panel A shows the expected overlap-induced bias under different assumed proportions of exposure-outcome sample overlap and different assumed ordinary least squares (OLS) bias values. Panel B shows the corresponding estimated type I error rates under the same scenarios. The dashed horizontal line in Panel B indicates the nominal type I error rate of 0.05. Across all hypothetical overlap scenarios, including complete overlap and high assumed OLS bias, the expected bias remained small and the type I error rate remained close to the nominal level, supporting limited influence of sample overlap on the *HSDL2*-FinnGen MR estimate.




Fig. S25. Global sample-overlap sensitivity analysis across UKB-PPP-FinnGen MR comparisons.

Global sample-overlap sensitivity analysis across 90 UKB-PPP-FinnGen exposure-file comparisons. Panel A shows the median expected overlap-induced bias under increasing assumed proportions of exposure-outcome sample overlap and assumed OLS bias values. Panel B shows the corresponding median type I error rates. Lines represent the median across all exposure-file comparisons, and shaded ribbons represent the interquartile range. The dashed horizontal line in Panel B indicates the nominal type I error rate of 0.05. Across the 90 comparisons, the expected bias remained small and the estimated type I error rate showed minimal inflation, even under conservative hypothetical overlap assumptions. These results suggest that potential sample overlap between UKB-PPP and FinnGen is unlikely to materially affect the overall MR findings.




Fig. S26. Overlap of FDR-significant TWMR genes across SSH, major depressive disorder and bipolar disorder.

Venn diagram showing the overlap of FDR-significant eQTLGen-based transcriptome-wide MR genes across the primary suicide or other intentional self-harm (SSH) outcome and two major psychiatric outcomes, major depressive disorder (MDD) and bipolar disorder (BIP). Numbers indicate the count of genes in each outcome-specific or shared set. The partial overlap suggests shared but non-identical transcriptomic genetic architecture across SSH and related psychiatric disorders, supporting the need to evaluate whether the *HSDL2*-SSH association is independent of broader psychiatric liability.