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.

SLC45A4 dependency and statin-class selectivity across cancer cell lines: evidence from colorectal lines

Takayuki Sakuma

tsakuma@soka.ac.jp

Soka University https://orcid.org/0000-0003-0092-3895

Research Article

Keywords: DepMap, PRISM, statins, mevalonate pathway, causal inference, double machine learning, pharmacogenomics

Posted Date: July 21st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9218378/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.

SLC45A4 dependency and statin-class selectivity across cancer cell lines: evidence from colorectal lines

Takayuki Sakuma*
Soka University

July 2026

Abstract

We integrated DepMap CRISPR dependency profiles with PRISM viability responses to estimate five-fold cross-fitted, covariate-adjusted associations with selective sensitivity to lipophilic versus hydrophilic statins. UBIAD1 and MVK showed strong associations, while SLC45A4 exhibited a smaller, positive association that remained stable across cross-fitting seeds. These findings showed directional consistency across CTRP and PRISM secondary AUC outcomes. Using secondary dose–response data, the association with SLC45A4 was particularly pronounced in colorectal cell lines for both AUC and the fitted lower asymptote, whereas the corresponding colorectal interaction was near zero for the primary phenotype. Colorectal co-dependency and CCLE metabolomics showed patterns consistent with sterol- and polyamine-related pathways. Compared with 1,000 random genes, SLC45A4 lay in the empirical tail for raw coefficients. However, this pattern was less pronounced under a standardized z-score comparison that accounts for differing GeneEffect variance across genes. Specifically, the comparison was marginally significant for the primary outcome and farther into the tail for the secondary AUC. These analyses support SLC45A4 as a promising candidate for focused follow-up, including direct perturbation and pharmacological validation in colorectal cell lines.

Keywords: DepMap; PRISM; statins; SLC45A4; mevalonate pathway; pharmacogenomics.

1 Introduction

Public pharmacogenomic resources support systematic searches for molecular correlates of drug response. Earlier GDSC studies identified genomic markers of sensitivity across cancer cell-line panels [5, 6]. PRISM established a large repurposing screen with multiplexed viability readouts for thousands of compounds across hundreds of cell lines [1], while the Cancer Dependency Map (DepMap) provides genome-scale CRISPR dependency profiles and harmonized cell-line annotations [3]. Goncalves et al. integrated pharmacological and CRISPR screens to infer mechanism of action [4].

Several studies have taken related but distinct approaches to these public pharmacogenomic and functional-genomic resources. Covell reanalyzed public Cancer Genome Project and Cancer Cell Line Encyclopedia drug-screening data to evaluate genomic biomarkers and cross-dataset drug-response prediction, providing a precedent for hypothesis-generating biomarker development from public pharmacogenomic resources [7]. Bazaga et al. combined TCGA mutation and expression

*e-mail: tsakuma@soka.ac.jp.

data, DepMap gene-essentiality data, and gene-network features to rank candidate therapeutic targets across nine cancer types [8]. A separate methodological literature addresses heterogeneous and single-cell responses rather than bulk cell-line panels: Ohnmacht et al. studied pharmacogenomic associations using epithelial–mesenchymal-transition signatures [17], Berg et al. developed heterogeneity-sequencing to identify modulators of cellular responses [18], Fu and Hou analyzed heterogeneous effects in single-cell genetic perturbation screens [19], and Dong et al. focused on causal identification in single-cell perturbation experiments [20]. Statin-specific biomarker studies, meanwhile, have typically used smaller panels and single-assay settings; Raghu et al., for example, analyzed 14 NCI-60 lines [14], far fewer than the 578 cell lines and 4,518 compounds in the PRISM repurposing screen [1].

Covell's study addressed broad cross-drug biomarker mining, Bazaga et al. addressed genome-wide target ranking, and Goncalves et al.'s mechanism-mapping approach (introduced above) spans a broad compound and target space [4]. The present study asks a narrower question at PRISM/DepMap scale: whether gene dependency is associated with an explicitly constructed statin-class contrast after adjustment for lineage, growth, and shared mevalonate-pathway dependency. It focuses on one drug class and a small set of interpretable gene–phenotype associations, and limits its claims to covariate-adjusted association and hypothesis prioritization rather than mechanism identification or genome-wide discovery.

A central problem is covariate structure and drug responses and gene-dependency measurements vary by tissue of origin and correlate with intrinsic proliferation. Lineage or growth differences can therefore produce an apparent gene–drug relationship without a specific biological connection [5, 6, 9]. We estimate covariate-adjusted associations with a cross-fitted orthogonalized partial-linear model based on double/debiased machine learning (DML) [10]. The procedure residualizes both the statin phenotype and the focal gene-dependency measure against observed covariates before estimating their association.

Statins have been examined as potential anticancer agents, but cell-line and clinical responses are heterogeneous [11–15]. Lipophilic and hydrophilic statins differ in permeability and distribution. We therefore define a lipophilic–hydrophilic response contrast and ask whether gene dependency is associated with this class-selectivity phenotype after adjustment for lineage, intrinsic growth, and shared mevalonate-pathway dependency.

The primary coefficient table evaluates four genes: UBIAD1, MVK, SLC45A4, and RIN1. UBIAD1 and MVK provide mevalonate-pathway anchors; SLC45A4 is the candidate taken forward for lineage and multi-omic analyses; RIN1 is retained as an additional comparator. This four-gene panel was selected after exploratory analysis identified SLC45A4 as a candidate of interest; UBIAD1, MVK, and RIN1 were then added as pathway-anchor and comparator genes for this focused follow-up table. The supporting analyses address distinct questions. CTRP and PRISM secondary AUC outcomes test cross-assay directional consistency. Leave-one-drug-out checks assess dependence on individual statins. Random-gene and random-drug comparisons provide empirical reference distributions. Secondary dose–response, co-dependency, and metabolomics analyses examine whether the SLC45A4 pattern is concentrated in colorectal cell lines and biologically coherent.

2 Results

We first define the PRISM statin-selectivity phenotype, then report covariate-adjusted coefficients, cross-assay consistency, empirical reference analyses, and lineage-stratified and metabolomic follow-up for SLC45A4.

2.1 Statin selectivity phenotype in PRISM primary screens

We use PRISM primary replicate-collapsed log fold-change (LFC) viability profiles [1] and extract eight statins from the PRISM treatment annotations (Supplementary Section S1; Table 8). We group statins into lipophilic (atorvastatin, fluvastatin, lovastatin, mevastatin, pitavastatin, simvastatin) and hydrophilic (pravastatin, rosuvastatin) classes. For each cell line i, we define mean responses $Y_{\text{lipo}}(i)$ and $Y_{\text{hydro}}(i)$ and the selectivity phenotype $Y_\Delta(i) = Y_{\text{lipo}}(i) - Y_{\text{hydro}}(i)$; negative Y_Δ indicates stronger inhibition by lipophilic statins. We require $n_{\text{lipo}}(i) \geq 3$ and $n_{\text{hydro}}(i) \geq 1$ for minimal coverage, retaining $n = \mathbf{465}$ cell lines (full definition in Methods, Section 4.1).

Figure 1 shows the distribution of Y_Δ across $n = \mathbf{465}$ cell lines. Negative Y_Δ is common, consistent with lipophilic statins being more growth-inhibitory than hydrophilic statins in these in vitro assays. This wide spread motivates cell-line-specific analyses.

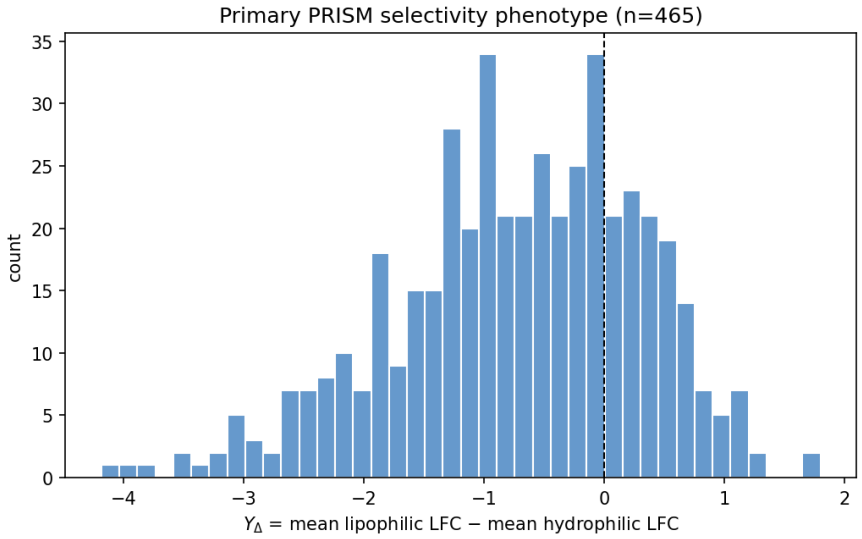


Figure 1: Distribution of the PRISM primary statin selectivity phenotype $Y_\Delta = Y_{\text{lipo}} - Y_{\text{hydro}}$ across $n = \mathbf{465}$ cell lines.

2.2 Primary covariate-adjusted associations for the four-gene panel

To estimate how gene dependency relates to statin selectivity after adjustment for lineage and intrinsic growth, we fit a partial-linear model and estimate the adjusted slope θ with cross-fitted orthogonalized regression [10]. The focal predictor is each gene-level DepMap dependency score (GeneEffect). Covariates include lineage indicators, inferred growth rate, and a mevalonate-core dependency score (Methods, Sections 4.1–4.2.3). In the main Results, the predictor is parameterized as $T = \text{GeneEffect}$ (more negative values indicate stronger dependency). Under this convention, $\hat{\theta} > 0$ means that stronger dependency is associated with a more negative Y_Δ (greater lipophilic vulnerability). Supplementary sensitivity analyses use $D = -\text{GeneEffect}$, which reverses the coefficient sign.

Table 1 reports the four-gene results. UBIAD1 and MVK show positive coefficients; stronger dependency is associated with more negative Y_Δ (greater vulnerability to lipophilic statins). SLC45A4 has a smaller positive association, while RIN1 is not statistically distinguishable from zero. The reported BH values are adjusted across these four coefficients only. Figure 2 shows the estimates and 95% normal-approximation intervals.

Table 1: Primary PRISM cross-fitted adjusted associations for Y_Δ. BH adjustment is across the four reported genes only and does not provide genome-wide FDR control.

Gene	$\hat{\theta}$	SE	p	q (BH)
UBIAD1	0.834	0.116	7.20e-13	2.88e-12
MVK	0.733	0.133	3.21e-08	6.42e-08
SLC45A4	1.043	0.453	2.13e-02	2.85e-02
RIN1	-0.334	0.340	3.26e-01	3.26e-01

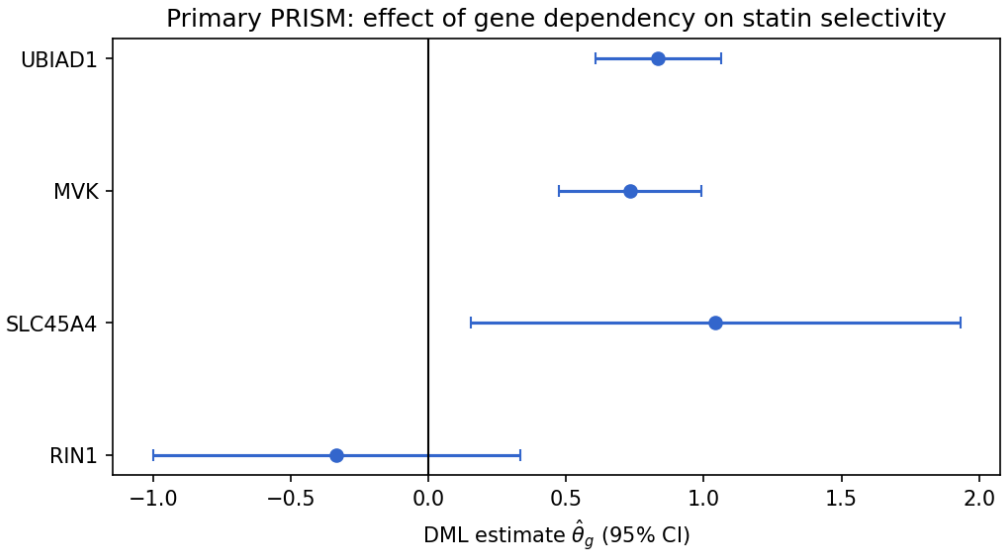


Figure 2: Primary cross-fitted adjusted coefficient estimates with 95% normal-approximation intervals.

2.3 Per-drug meta-analysis of the lipophilic–hydrophilic contrast

The primary phenotype in Section 2.2 averages log fold-change values across drugs before estimating θ, so a single pooled regression could in principle be dominated by one or two high-variance compounds or by scale differences across drugs. As an alternative that does not rely on this pooling step and targets the same lipophilic-versus-hydrophilic contrast as Y_Δ, we fit the cross-fitted adjustment specification separately for each of the eight statins (GeneEffect against each drug's own LFC, adjusting for lineage, growth, and the mevalonate-core score), then combined the eight drug-specific estimates within each class using a DerSimonian–Laird random-effects meta-analysis. This yields a class-level estimate for lipophilic statins ($k = 6$ drugs) and one for hydrophilic statins ($k = 2$ drugs).

Because the eight per-drug fits share the same cell lines, the same GeneEffect column, and the same covariates, their estimation errors are correlated. A standard Wald test assumes independence, understating the true uncertainty and overstating significance. For this reason we report inference from a cell-line-clustered bootstrap ($B = 1000$ resamples of cell lines, refitting the entire per-drug-and-meta-analysis pipeline on each resampled panel; Section 4.2.4). Table 2 reports the results. For UBIAD1 and MVK, the lipophilic-class estimate is large and the hydrophilic-class estimate is near zero, consistent with Table 1, and the bootstrap difference test is significant for both ($p = 5.7 \times 10^{-11}$

and $p = 6.7 \times 10^{-7}$). For SLC45A4, the lipophilic-class estimate is positive ($\hat{\theta}_{RE} = 1.020$, $k = 6$) and the hydrophilic-class estimate is smaller ($\hat{\theta}_{RE} = 0.099$, $k = 2$). The lipophilic-minus-hydrophilic difference is positive and agrees in sign with the pooled primary estimate ($\hat{\theta} = 1.043$, Table 1). Under the cell-line-clustered bootstrap, however, this difference is not significant at the conventional threshold ($\hat{\theta} = 0.921$, bootstrap SE= 0.498, $p = 0.065$, 95% CI -0.02 to 1.94). We therefore treat this analysis as directional support for the primary SLC45A4 result. RIN1 shows a similar pattern: its difference is negative but not significant under the bootstrap ($p = 0.15$), consistent with its non-significant primary coefficient in Table 1. A Wald test that incorrectly assumes independence between the eight per-drug fits shows RIN1 as nominally significant (Supplementary Section S6). The hydrophilic subgroup contains only two drugs and its class-level estimate is imprecise. This is part of reason why the bootstrap interval for SLC45A4 is wide.

Table 2: Per-drug DML estimates combined by DerSimonian–Laird random-effects meta-analysis within lipophilic ($k = 6$) and hydrophilic ($k = 2$) statin subgroups. This targets the same lipophilic–hydrophilic contrast as the primary Y_Δ phenotype (Table 1), estimated here from eight separate per-drug fits without first averaging the outcome across drugs. Because the eight per-drug fits share cell lines, covariates, and the GeneEffect column, SE, p, and q are from a cell-line-clustered bootstrap ($B = 1000$; Section 4.2.4); a Wald test would incorrectly treat the two subgroup meta-analyses as independent.

Gene	Lipophilic $\hat{\theta}_{RE}$	Hydrophilic $\hat{\theta}_{RE}$	Diff. (lipo−hydro)	SE	p	q (BH)
UBIAD1	0.817	0.016	0.801	0.122	5.68e-11	2.27e-10
MVK	0.680	-0.026	0.706	0.142	6.69e-07	1.34e-06
SLC45A4	1.020	0.099	0.921	0.498	6.45e-02	8.60e-02
RIN1	-0.645	-0.196	-0.449	0.313	1.51e-01	1.51e-01

2.4 Cross-assay consistency in AUC outcomes

To examine whether coefficient directions persist under alternative drug-response measurements, we repeat the same adjustment specification using CTRP_AUC and Repurposing_secondary_AUC accessed through the DepMap portal. The DepMap gene-dependency measurements and many cell lines overlap with the primary analysis; only the drug-response outcome is replaced by an AUC-based statin summary. These analyses assess cross-assay consistency and do not provide independent-sample replication. Each AUC outcome was oriented by its correlation with the primary phenotype over shared cell lines (Methods, Section 4.2.5). All recorded orientation factors were +1, so no numerical sign reversal was applied, but the sign-match statistic remains descriptive (Supplementary Section S2; Table 9). The displayed coefficients have the same direction as the primary estimates, and UBIAD1 remains statistically distinguishable from zero across the reported specifications (Table 3; Figs. 3–4).

Table 3: Cross-assay consistency between the primary phenotype and AUC-based outcomes.

Dataset	k (lipo thr)	Gene	n	theta (primary)	q (primary)	theta (external)	q (external)	Sign match
CTRP_AUC	1	MVK	421	0.701	1.69e-07	0.882	8.31e-08	1
CTRP_AUC	1	RIN1	421	-0.708	4.96e-02	-1.038	2.57e-02	1
CTRP_AUC	1	SLC45A4	421	1.488	2.02e-03	0.895	1.25e-01	1
CTRP_AUC	1	UBIAD1	421	0.847	1.08e-13	1.130	2.91e-16	1
CTRP_AUC	3	MVK	361	0.616	4.69e-05	0.835	4.49e-06	1
CTRP_AUC	3	RIN1	361	-0.379	3.35e-01	-1.059	3.41e-02	1
CTRP_AUC	3	SLC45A4	361	1.513	4.89e-03	1.068	9.39e-02	1
CTRP_AUC	3	UBIAD1	361	0.888	4.23e-13	1.252	3.18e-18	1
Repurposing_secondary_AUC	1	MVK	397	0.728	4.85e-07	0.023	6.61e-03	1
Repurposing_secondary_AUC	1	RIN1	397	-0.797	4.68e-02	-0.054	1.05e-02	1
Repurposing_secondary_AUC	1	SLC45A4	397	0.883	6.52e-02	0.046	7.57e-02	1
Repurposing_secondary_AUC	1	UBIAD1	397	0.880	9.49e-14	0.040	7.04e-10	1
Repurposing_secondary_AUC	3	MVK	342	0.470	3.13e-03	0.007	4.02e-01	1
Repurposing_secondary_AUC	3	RIN1	342	-1.151	3.44e-03	-0.076	5.00e-04	1
Repurposing_secondary_AUC	3	SLC45A4	342	0.684	1.59e-01	0.067	1.61e-02	1
Repurposing_secondary_AUC	3	UBIAD1	342	0.634	4.87e-07	0.027	2.05e-04	1

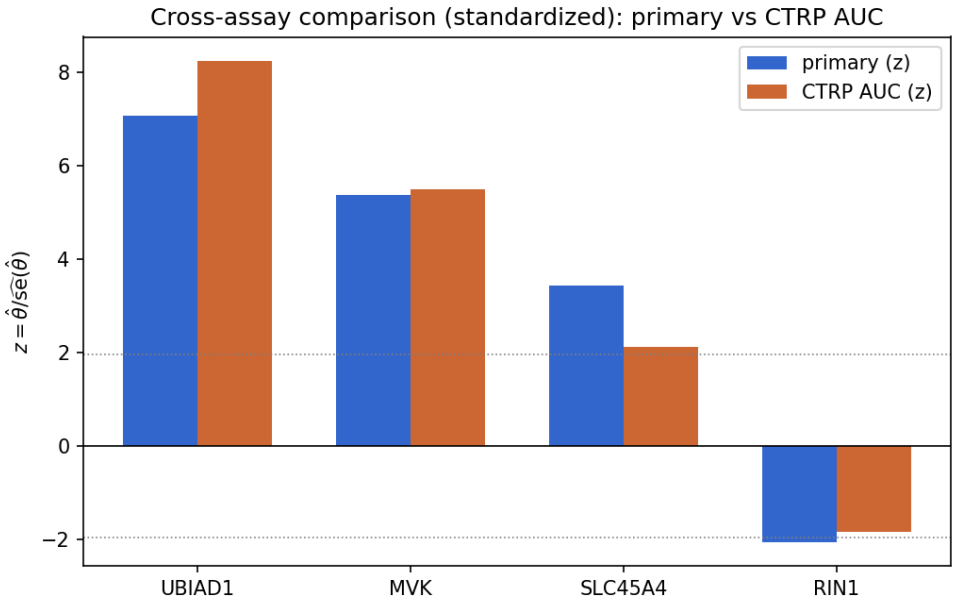


Figure 3: Cross-assay comparison in CTRP AUC, shown as z-statistics ($\hat{\theta}/\widehat{\mathrm{se}}$) instead of raw coefficients. The primary phenotype (PRISM LFC) and CTRP AUC are on different numeric scales and are not directly comparable as raw coefficients. Dotted lines mark $|z| = 1.96$.

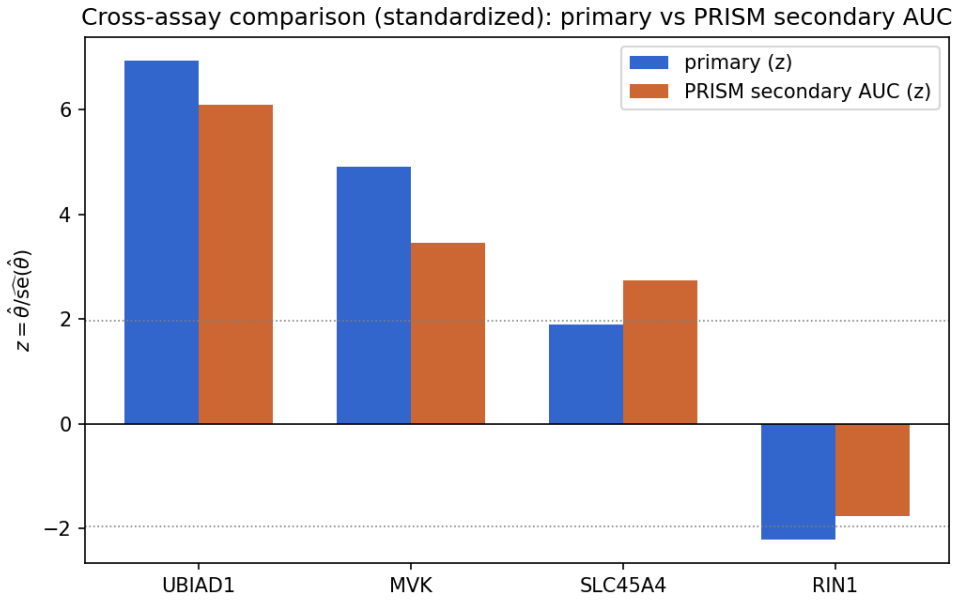


Figure 4: Cross-assay comparison in PRISM secondary AUC, shown as z-statistics for the same scale reason as Fig. 3.

2.5 LODO robustness and negative controls

We assess whether the estimates are driven by a single compound (leave-one-drug-out; LODO) and whether comparable associations arise from arbitrary drug groupings or from generic genes (negative controls). All three checks reuse the same cross-fitted adjustment specification and differ only in outcome or predictor construction (Methods, Section 4.2.5).

LODO analyses indicate that the results are not driven by a single statin. For UBIAD1 and MVK, the sign of $\hat{\theta}$ was stable and the min–max range across folds was small (Figs. 5–6, Table 4).

Table 4: LODO summary by gene (external datasets).

Dataset	Gene	n	n drugs	theta (all)	theta min (LOO)	theta max (LOO)	Sign match rate	q (all)
CTRP_AUC	MVK	421	3	0.897	0.805	1.023	1.00	5.73e-08
CTRP_AUC	RIN1	421	3	-1.070	-1.327	-0.705	1.00	2.08e-02
CTRP_AUC	SLC45A4	421	3	0.884	0.827	1.327	1.00	1.29e-01
CTRP_AUC	UBIAD1	421	3	1.123	1.057	1.235	1.00	4.89e-16
Repurposing_secondary_AUC	MVK	397	6	0.023	0.018	0.026	1.00	6.78e-03
Repurposing_secondary_AUC	RIN1	397	6	-0.055	-0.066	-0.051	1.00	9.38e-03
Repurposing_secondary_AUC	SLC45A4	397	6	0.044	0.004	0.062	1.00	8.50e-02
Repurposing_secondary_AUC	UBIAD1	397	6	0.041	0.032	0.046	1.00	2.89e-10

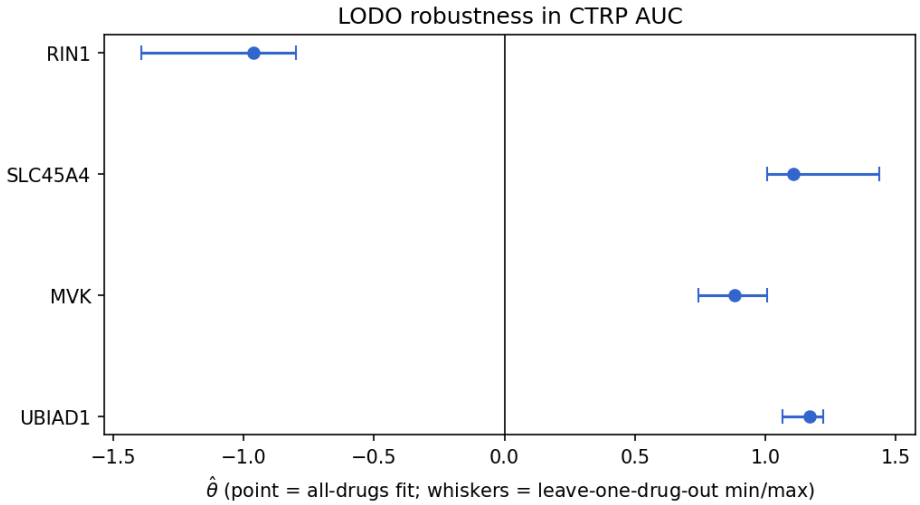


Figure 5: LODO robustness in CTRP AUC: min/max over leave-one-drug-out folds (within-outcome comparison; not affected by the cross-scale issue in Fig. 3).

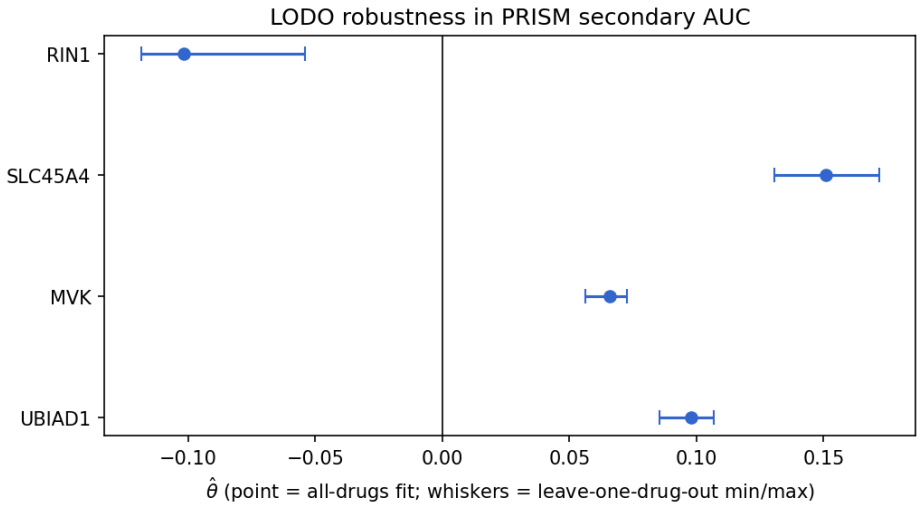


Figure 6: LODO robustness in PRISM secondary AUC: min/max over leave-one-drug-out folds.

Additional sensitivity analyses for the primary selectivity phenotype. Because the selectivity phenotype depends on having both lipophilic and hydrophilic statin observations, we also conducted stricter-coverage and leave-one-drug-out sensitivity analyses. Under a stricter coverage requirement of at least five lipophilic statins and both hydrophilic statins ($n = 437$), we re-estimated the effect using dependency strength $D = -\mathsf{GeneEffect}$ (so that more negative $\hat{\theta}$ corresponds to stronger dependency associated with greater vulnerability). Under this parameterization, the SLC45A4 estimate remains negative ($\hat{\theta} = -0.980$, $p = 0.043$; equivalently $+0.980$ under the GeneEffect convention of Table 1) and the sign remains stable under leave-one-lipophilic-out and leave-one-hydrophilic-out definitions (Table 20, Figs. 14–15).

Gene-level empirical reference analysis. As the primary empirical reference check, we compared SLC45A4 with 1000 randomly sampled genes analyzed with the same procedure. The SLC45A4 estimate falls in the tail of the resulting null distribution for the raw coefficient, for both primary selectivity ($p^{\theta}_{\text{emp}} = 0.030$) and secondary AUC ($p^{\theta}_{\text{emp}} = 0.019$). Because GeneEffect variance differs across genes (SLC45A4's is smaller than the other three genes'), we also report the standardized (z-score) comparison, which is more directly comparable across genes. This is borderline for the primary outcome ($p^{z}_{\text{emp}} = 0.051$) and closer to conventional significance for secondary AUC ($p^{z}_{\text{emp}} = 0.031$) (Table 21, Figs. 16–17). With $N = 1000$ draws, this check provides the primary gene-level empirical reference used in the Discussion.

A separate drug-level analysis compares the same genes with $B = 10000$ random bundles of other drugs under the primary covariate specification (Supplementary Section S6, Table 22). UBIAD1, MVK, and RIN1 are more extreme than nearly all random bundles ($p_{\text{emp}} \leq 0.001$), and SLC45A4 is more extreme than 99.7% of them ($p_{\text{emp}} = 0.003$). This analysis targets a combined eight-drug meta-analytic signal rather than the lipophilic–hydrophilic contrast, so we treat it as a supporting empirical reference. As reported above (Section 2.3), the per-drug meta-analysis targets the lipophilic–hydrophilic contrast through a different estimation route and agrees in direction, though it does not independently reach conventional significance for SLC45A4 under cell-line-clustered inference.

2.6 Lineage-stratified dose–response parameters: SLC45A4 signal in colorectal

To examine SLC45A4 beyond the primary four-gene table, we analyzed PRISM secondary dose–response curve parameters for lipophilic statins and fitted lineage-stratified regressions using dependency strength $D = -\mathsf{GeneEffect}$. These analyses are post-primary and exploratory.

In colorectal cell lines with secondary dose–response profiles ($n = 21$), stronger SLC45A4 dependency is associated with greater lipophilic-statin sensitivity; the largest fitted coefficients occur for AUC and the lower asymptote (lower_limit) (Table 5). For AUC, the multivariable estimate is $\hat{\theta} = -1.021$ ($p = 1.22 \times 10^{-4}$, $q = 5.84 \times 10^{-3}$); for lower_limit, $\hat{\theta} = -3.640$ ($p = 4.25 \times 10^{-4}$, $q = 1.02 \times 10^{-2}$), using a t-distribution appropriate to this sample size and BH-adjusting across the full 8-lineage-by-6-metric family tested (48 tests per model; Methods). Both remain significant at $q < 0.05$ after this correction, though the p-values are less extreme than a normal approximation would suggest, and the small colorectal sample makes the estimates potentially sensitive to individual cell lines. Lung estimates are near zero (AUC $\hat{\theta} = -0.027$, $q = 0.95$; lower_limit $\hat{\theta} = -0.549$, $q = 0.80$). This contrast shows heterogeneity in the fitted associations but is not a formal multiplicity-adjusted interaction test in the sense of jointly modeling an SLC45A4-dependency×lineage interaction term. Lineage-by-lineage results for AUC, lower_limit, and EC50 are reported in Supplementary Section S3 (Tables 10–12). All six curve metrics were included in the 48-test BH correction.

Within the displayed colorectal regressions, the association is largest for AUC and the fitted lower asymptote (lower_limit), whereas EC50 shows little association. AUC and lower_limit summarize the overall and maximal depth of growth inhibition; EC50 summarizes the concentration at a midpoint response. The pattern is therefore more consistent with variation in depth of inhibition than with a general potency shift, subject to the exploratory nature of the analysis. Figures 7 and 8 show lineage-specific estimates for AUC and the lower limit, and Fig. 9 summarizes the colorectal results across all six curve metrics produced by the dose-response fitting pipeline (AUC, lower_limit, EC50, fitted R^2, slope, and the upper asymptote). We focus interpretation on AUC, lower_limit, and EC50 because they have the clearest sensitivity and potency readings. All six metrics were inspected, however, so the BH-adjusted correction family (Methods) spans all six metrics across

the 8 tested lineages (48 tests per model), not only the three discussed in the main narrative.

Table 5: Lineage-stratified associations between SLC45A4 dependency strength ($D = -\mathsf{GeneEffect}$) and PRISM secondary lipophilic-statin dose–response parameters. "Adj." additionally controls for D_{UBIAD1}, D_{MVK}, and D_{HMGCR}. HC3 standard errors are in parentheses. Because n is small (as low as 21), p-values use a t-distribution with the residual degrees of freedom of each model, appropriate at this sample size (df= $n-2$ for the univariate model, $n-5$ for the multivariable model). q is BH-adjusted across all 8 tested lineages $\times$ all 6 curve metrics inspected (48 tests per model; Methods, Section 4.2.6), not only the 3 metrics shown in this table.

Lineage	Metric	n	$\hat{\theta}_{\text{uni}}$ (SE)	p_{uni} (t)	q_{uni}	$\hat{\theta}_{\text{adj}}$ (SE)	p_{adj} (t)	q_{adj}
colorectal	AUC	21	-0.676 (0.169)	7.80e-04	1.76e-02	-1.021 (0.203)	1.22e-04	5.84e-03
colorectal	lower_limit	21	-2.727 (0.659)	5.62e-04	1.76e-02	-3.640 (0.823)	4.25e-04	1.02e-02
colorectal	EC50	21	0.975 (2.663)	0.718	0.826	0.213 (3.734)	0.955	0.986
lung	AUC	80	-0.037 (0.141)	0.792	0.848	-0.027 (0.145)	0.851	0.950
lung	lower_limit	80	-0.549 (0.480)	0.256	0.577	-0.549 (0.514)	0.290	0.800
lung	EC50	80	1.502 (0.985)	0.131	0.469	1.528 (1.019)	0.138	0.712
skin	AUC	29	0.674 (0.517)	0.203	0.508	0.228 (0.524)	0.667	0.861
skin	lower_limit	29	2.417 (1.495)	0.118	0.469	1.314 (1.524)	0.397	0.800
skin	EC50	29	1.418 (4.322)	0.745	0.826	0.181 (4.117)	0.965	0.986

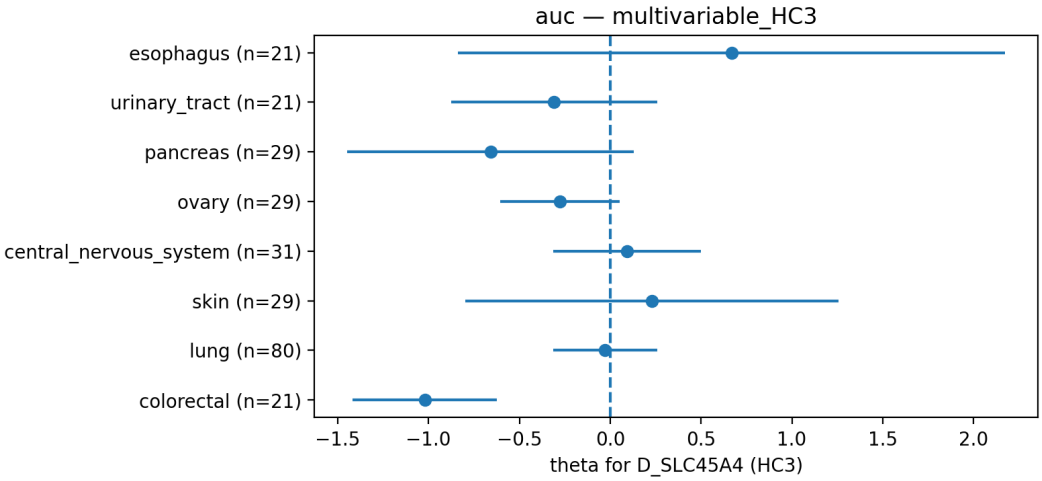


Figure 7: Lineage-stratified SLC45A4 association for lipophilic-statin AUC.

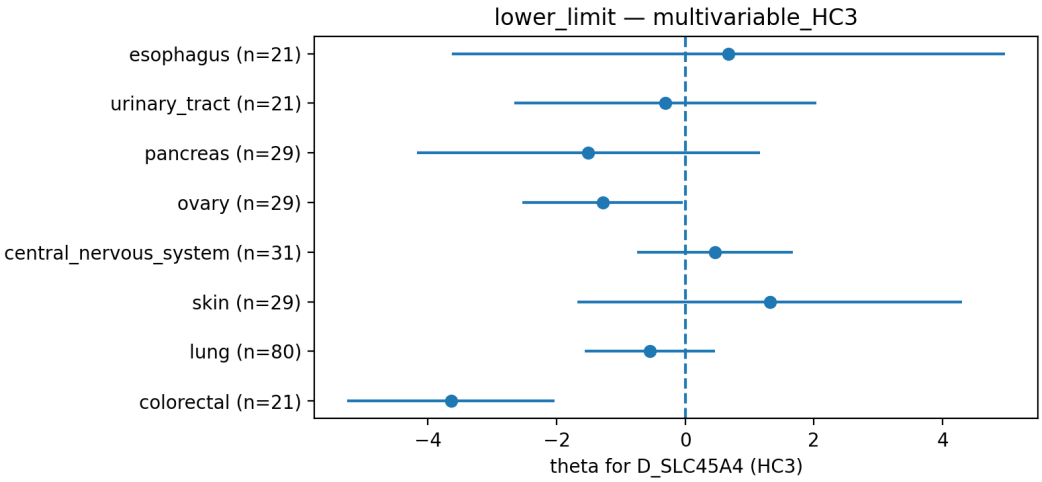


Figure 8: Lineage-stratified SLC45A4 association for lipophilic-statin lower_limit.

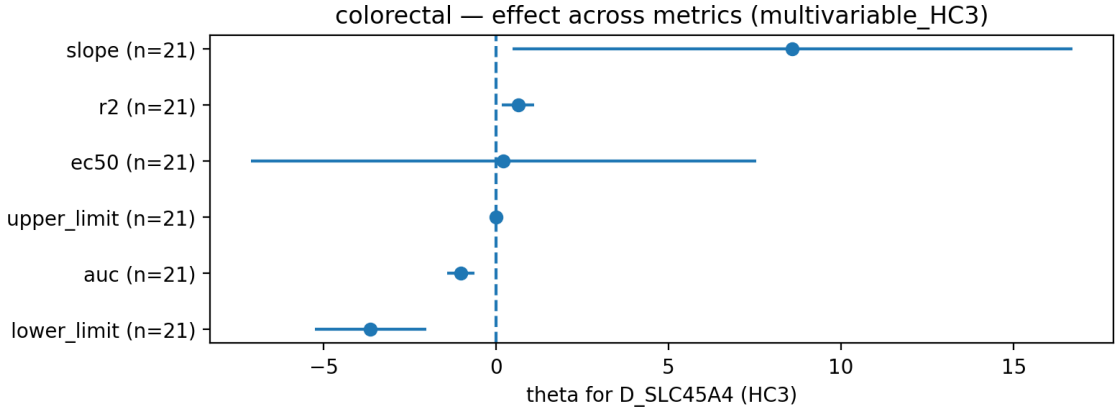


Figure 9: Colorectal dose–response parameter decomposition: estimated SLC45A4 effects across all six curve metrics produced by the dose-response fitting pipeline (AUC, lower_limit, EC50, fitted R^2, slope, and upper_limit; Methods, Section 4.2.6).

2.7 Phenotype-focused co-dependency supports sterol biosynthesis in phenotyped colorectal lines

To test whether the phenotyped colorectal subset reflects a distinct dependency pattern, we compared SLC45A4 co-dependency correlations calculated in (i) all colorectal lines with available Gene-Effect ($n = 58$) versus (ii) the phenotyped colorectal subset overlapping the secondary dose–response data ($n = 21$). The correlation structure is broadly similar (Fig. 10), but the resulting top200 co-dependency sets overlap only modestly (12 shared genes; Jaccard ≈ 0.031), which points to context dependence in SLC45A4 connectivity.

The phenotyped colorectal top 200 set is enriched for the study-defined mevalonate/cholesterol-biosynthesis set (Section 4.2.7; 4/21 core genes: FDFT1, SQLE, NSDHL, LSS; hypergeometric $p = 7.77 \times 10^{-5}$), whereas the all-colorectal top 200 set shows no overlap (Table 6).

A broader pathway enrichment analysis using Enrichr [21–23] ranks KEGG *Steroid biosynthesis* as the top pathway for the phenotyped colorectal top 200 (FDR = 0.00585; overlap genes: {FDFT1, SQLE, NSDHL, LSS}).

Table 6 summarizes these analyses. In the rows for study-defined sterol-set enrichment (hypergeometric test), k denotes the number of study-defined sterol-core genes contained in the SLC45A4 top-200 co-dependency set. In the subset-permutation rows, the reported values are empirical tail probabilities for observing overlaps at least as large as the observed overlap under repeated random draws of colorectal subsets of the same size as the phenotyped subset ($N = 2000$).

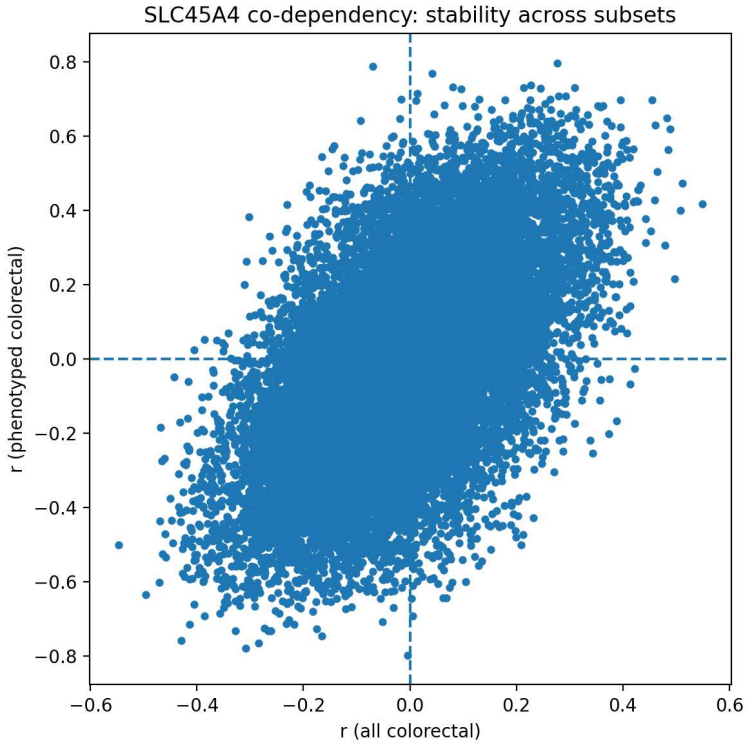


Figure 10: SLC45A4 co-dependency stability across colorectal subsets: Pearson correlations with other genes computed in all colorectal lines vs phenotyped colorectal subset.

Table 6: Sterol-related enrichment analyses in colorectal lines: study-defined-set enrichment, pathway enrichment, and subset-permutation empirical null. Here k is the overlap count between the SLC45A4 top200 co-dependency set and the study-defined sterol/mevalonate set (Section 4.2.7); empirical-null probabilities are based on $N = 2000$ subset permutations.

Analysis	Subset	n	Signal	Summary
Study-defined sterol-set enrichment (hypergeometric test)	phenotyped colorectal	21	$k = 4$	$p = 7.77 \times 10^{-5}$ (top200 $\cap$ core set)
Study-defined sterol-set enrichment (hypergeometric test)	all colorectal	58	$k = 0$	$p = 1.00$
KEGG pathway enrichment	phenotyped colorectal	21	Steroid biosynthesis	adj.$p = 5.85 \times 10^{-3}$; genes: SQLE, NSDHL, LSS, FDFT1
Empirical-null subset permutation	colorectal pool	58	P($k \geq 4$)	0.0055 (N=2000)
Empirical-null subset permutation	colorectal pool	58	P(all four sterol genes in top200)	0.0020

2.8 Subset-permutation empirical null argues against a random-subset explanation

Because the phenotyped colorectal subset is defined by data availability, we asked whether the observed sterol-gene overlap could arise simply because the analysis is restricted to a small subset of colorectal lines ($n = 21$ of $N = 58$). To address this question, we built an empirical null by repeatedly drawing random colorectal subsets of the same size, recomputing the SLC45A4 top 200 co-dependency set, and recounting how many study-defined mevalonate/cholesterol-core genes fell into that set.

In the observed phenotyped subset, the overlap is $k_{\text{obs}} = 4$ (FDFT1, SQLE, NSDHL, LSS). Across 2000 random subsets, overlaps this large were rare: $P(k \geq 4) = 0.0055$. The stricter event that *all* four sterol genes appeared in the top 200 set had probability 0.0020 (Table 6). Supplementary Section S4 reports the full subset-permutation null distribution and sterol-gene inclusion rates (Tables 13–14). Figure 11 visualizes this null distribution.

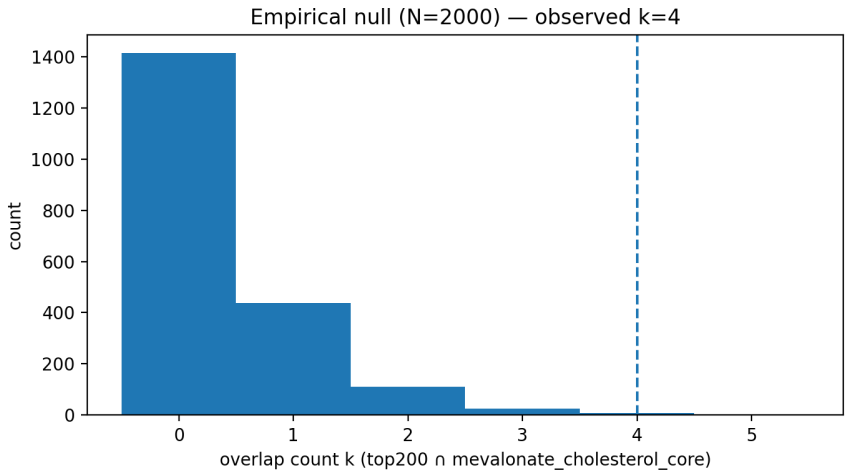


Figure 11: Empirical null for the sterol-core overlap count k under subset permutation. For each of 2000 random colorectal subsets of size 21, we recomputed the SLC45A4 top 200 co-dependency set and counted how many study-defined sterol/mevalonate-core genes were included. Bars show how often each k occurred; the dashed line marks the observed value in the phenotyped subset ($k_{\text{obs}} = 4$). Only 0.55% of random subsets produced $k \geq 4$, arguing against a purely random-subset explanation.

2.9 CCLE metabolomics triangulation: colorectal-associated cholesteryl ester and polyamine signatures

Independent functional studies suggest that SLC45A4 can transport polyamines such as putrescine in specific biological contexts [24, 25]. This provides context for examining putrescine and related metabolites when interpreting metabolomic patterns associated with SLC45A4 dependency in colorectal lines.

To assess whether the association is also reflected in metabolomic profiles, we merged DepMap CRISPR gene-effect and sample-information tables from the 24Q4 public release with the CCLE metabolomics resource of Li et al. [3, 16]. This let us evaluate whether the SLC45A4-by-colorectal interaction pattern is accompanied by consistent metabolite abundance patterns. For each metabolite M, we estimated an interaction model:

$$M_i = \alpha + \beta_D D_i + \beta_{D\times CRC}\, D_i \cdot \mathbf{1}\{CRC_i\} + \gamma^\top \text{Lineage}_i + \varepsilon_i, \tag{1}$$

where D is the (sign-flipped) SLC45A4 dependency, $\mathbf{1}\{CRC_i\}$ is an indicator that equals 1 if cell line i is colorectal and 0 otherwise, and Lineage_i denotes a vector of lineage indicator variables. We report HC3-robust p-values and an empirical subset null. For the subset null, we repeatedly sampled non-colorectal subsets of size n_{CRC} and recomputed $|\hat{\beta}_{D\times S}|$ for a randomly generated indicator S. These pathway-guided features were selected for biological relevance rather than through a metabolome-wide discovery screen. The reported p-values are therefore not adjusted across all measured metabolites and are interpreted as exploratory.

Across the merged set ($n = 423$, $n_{\text{CRC}} = 26$), polyunsaturated cholesteryl ester species exhibited positive interaction coefficients for colorectal lineage (Table 7; Fig. 12A–D). For example, C20:5 CE (a cholesteryl ester with a 20-carbon, five-double-bond fatty-acyl chain) had $\hat{\beta}_{D\times CRC} = 1.51$ ($p = 1.0\times10^{-4}$) with subset-null $p_{\text{emp}} = 9.99\times10^{-4}$. Putrescine showed the opposite pattern: $\hat{\beta}_{D\times CRC} = -2.01$ ($p = 0.039$; $p_{\text{emp}} = 0.0015$), consistent with lower putrescine in colorectal lines with stronger

SLC45A4 dependency (Fig. 12E). In contrast, the interaction for the primary PRISM selectivity outcome was small and non-significant (Fig. 12F).

Within colorectal lines (metabolomics $n_{\mathrm{CRC}} = 33$ when not requiring PRISM overlap), putrescine also correlated negatively with SLC45A4 dependency (Spearman $\rho = -0.49$, $p = 0.011$), supported by leave-one-out and an in-lineage permutation null ($p_{\mathrm{perm}} = 0.0090$). However, putrescine and cholesteryl ester levels were not directly correlated within colorectal (Table 15; Fig. 13). Moreover, conditioning on putrescine did not attenuate the SLC45A4×CRC interaction for CE features (Table 16; empirical shrinkage $p = 1$). Taken together, these results indicate that SLC45A4 dependency in colorectal lines is accompanied by changes in both putrescine-related and cholesteryl-ester features. However, within-colorectal analyses and adjusted models do not support a simple sequential relationship in which one metabolite change alone explains the other.

Table 7: Metabolomic interaction signals between SLC45A4 dependency and colorectal lineage. "Feature" denotes the metabolite name in CCLE metabolomics; C20:5 CE etc denote cholesteryl ester species with fatty-acyl chain length:unsaturation, and CE_total denotes total cholesteryl ester abundance. p denotes the HC3-robust interaction test; p_{emp} denotes the non-colorectal subset-permutation empirical null. The features are pathway-guided, and the reported values are not adjusted across the full metabolomics feature space.

Feature	n	n_{CRC}	$\hat{\beta}_{D\times CRC}$	p	p_{emp}
C20:5 CE	423	26	1.513	1.02e-04	9.99e-04
C20:4 CE	423	26	1.484	8.59e-04	9.99e-04
CE_total	423	26	10.755	0.005	0.005
C22:6 CE	423	26	1.190	0.007	0.012
putrescine	423	26	-2.014	0.039	0.001

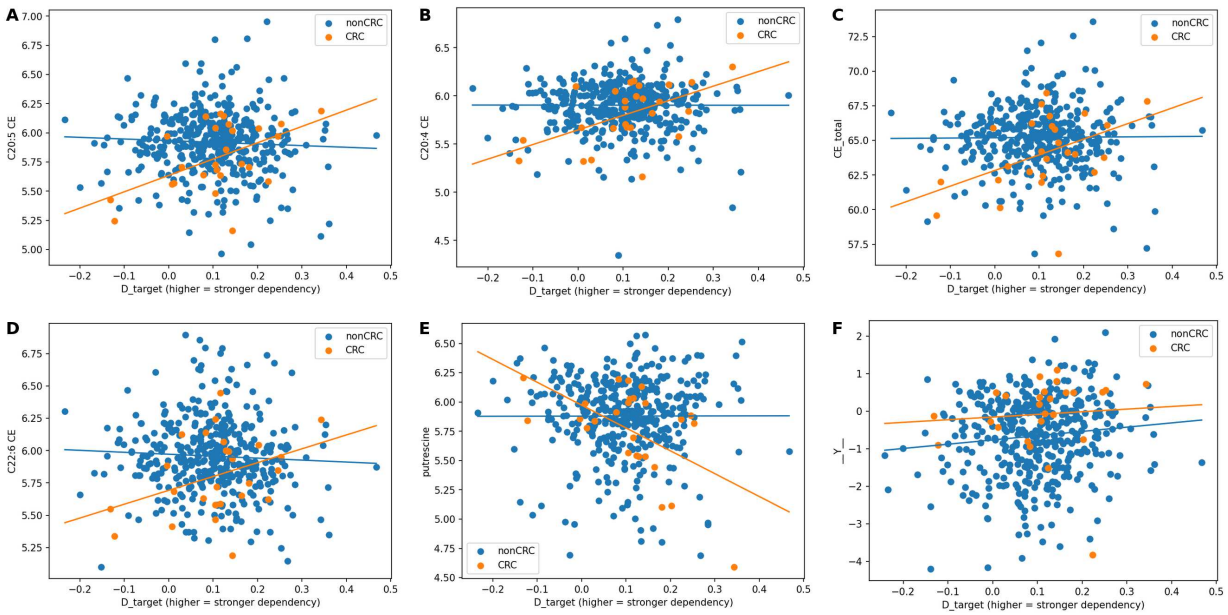


Figure 12: Interaction plots for selected CCLE metabolomics features and the primary PRISM selectivity outcome. Panels A–D: polyunsaturated cholesteryl esters (CE) and total CE show positive interactions with colorectal lineage; panel E: putrescine shows a negative interaction; panel F: the interaction for the primary PRISM selectivity outcome is near zero.

3 Discussion

This study estimates covariate-adjusted associations between DepMap gene dependency and a PRISM statin class-selectivity phenotype. UBIAD1 and MVK show the strongest and most stable primary coefficients, consistent with mevalonate-pathway involvement in differential statin response [12, 14]. SLC45A4 has a smaller primary coefficient than either mevalonate-pathway comparator. An alternative per-drug analysis (Section 2.3) reaches the same lipophilic-versus-hydrophilic contrast by fitting eight drug-specific models separately and combining them, without first averaging the outcome across drugs. It yields a consistent sign and comparable magnitude ($\hat{\theta}_{RE} = 1.020$ for the lipophilic class versus 0.099 for the hydrophilic class). Under a cell-line-clustered bootstrap that accounts for the shared structure across the eight per-drug fits, this difference falls just short of conventional significance ($p = 0.065$). SLC45A4 is best viewed as a candidate taken forward for focused follow-up.

The AUC analyses provide directional consistency across alternative response measurements. They do not constitute independent replication. The same DepMap dependency measurements are reused, cell-line samples overlap, and Repurposing_secondary_AUC is another PRISM assay. The EGFR and BRAF positive controls show that the integration pipeline can recover expected target–drug relationships in CTRP and GDSC, but they do not validate the SLC45A4 finding itself.

The colorectal interpretation arises from secondary dose–response, co-dependency, and metabolomic analyses. In the secondary dose–response data, SLC45A4 coefficients are largest for AUC and the fitted lower asymptote in 21 colorectal lines. In the phenotyped colorectal subset, the SLC45A4 co-dependency ranking is enriched for sterol-biosynthesis genes, and the subset-permutation analysis indicates that the observed overlap is uncommon under random colorectal subsets. CCLE metabolomics adds a parallel pattern involving polyunsaturated cholesteryl esters

and putrescine, although this pathway-guided analysis does not provide metabolome-wide error control. The primary PRISM selectivity interaction with colorectal lineage is near zero, however, so the colorectal label should not be read as a replicated primary lineage interaction.

Several design features limit inference. The primary BH family contains only four reported genes and does not control genome-wide discovery. The primary phenotype averages raw LFC values across drugs and can retain drug-specific scale differences. The per-drug meta-analysis (Section 2.3) partially addresses this by avoiding the averaging step, but its hydrophilic subgroup contains only two drugs and its class-level estimates are themselves imprecise; its cell-line-clustered bootstrap difference test for SLC45A4 does not reach conventional significance ($p = 0.065$). The external-outcome orientation was chosen using its correlation with the primary phenotype, although no sign reversal occurred. The random-drug-bundle analysis (Supplementary Section S6, rerun at $B = 10000$ using the primary covariate specification) tests a combined eight-drug signal, a different quantity from the lipophilic–hydrophilic contrast; we treat it as a corroborating check. We rely instead on the 1000-draw gene-level null (Table 21) as the primary empirical reference evidence, though it too is a single form of empirical reference and not a substitute for an experimentally randomized comparison. The colorectal dose–response analysis has a small sample, and although we correct for the tested lineage-by-metric family (Section 2.6), this is not a formal joint SLC45A4-dependency-by-lineage interaction model. Most importantly, cross-fitting and orthogonalization adjust observed covariates and reduce regularization bias; they do not make gene dependency random or remove unmeasured confounding.

The practical contribution is to prioritize a candidate for follow-up. A direct next test is to perturb SLC45A4 in colorectal models and measure lipophilic and hydrophilic statin dose–response curves under the same experimental protocol. Parallel measurements of sterol flux, cholesteryl esters, and putrescine would distinguish whether the statistical pattern reflects a reproducible biological state or an artifact of cross-resource integration.

4 Methods

4.1 Data sources and variables

4.1.1 DepMap gene dependency (GeneEffect)

DepMap quantifies gene dependency using genome-wide CRISPR knockout screens across many cancer cell lines. We use processed gene-level dependency scores (GeneEffect), where more negative values indicate stronger dependency. For each candidate gene, GeneEffect is the continuous focal predictor T in the model. The outcome Y is a statin selectivity phenotype derived from PRISM. The coefficient is interpreted as a covariate-adjusted cross-cell-line association. It does not represent the effect of experimentally changing gene dependency.

4.1.2 PRISM drug response and statin phenotypes

PRISM primary screens report log fold-change (LFC) viability relative to controls [1]. We extracted eight statins from PRISM primary treatment annotations and grouped them into lipophilic (atorvastatin, fluvastatin, lovastatin, mevastatin, pitavastatin, simvastatin) and hydrophilic (pravastatin,

rosuvastatin) classes (Table 8). For each cell line i, we computed:

$$Y_{\text{lipo}}(i) = \frac{1}{n_{\text{lipo}}(i)} \sum_{d \in \mathcal{D}_{\text{lipo}}(i)} \text{LFC}_{id}, \tag{2}$$

$$Y_{\text{hydro}}(i) = \frac{1}{n_{\text{hydro}}(i)} \sum_{d \in \mathcal{D}_{\text{hydro}}(i)} \text{LFC}_{id}, \tag{3}$$

$$Y_{\Delta}(i) = Y_{\text{lipo}}(i) - Y_{\text{hydro}}(i). \tag{4}$$

Negative Y_{Δ} indicates stronger growth inhibition by lipophilic statins relative to hydrophilic statins. We required $n_{\text{lipo}}(i) \geq 3$ and $n_{\text{hydro}}(i) \geq 1$ to ensure minimal coverage; this retained $n = \mathbf{465}$ cell lines (Fig. 1).

4.1.3 Covariates

Covariates are cell-line features associated with both dependency and drug response. Adjusting for them reduces associations attributable to broad lineage or growth structure.

We included the following covariates X: (i) one-hot encoded lineage (OncotreeLineage) from DepMap metadata, (ii) inferred growth rate from DepMap, and (iii) a *mevalonate-core dependency score*, defined as the first principal component (PC1) of GeneEffect values in a study-defined mevalonate gene set.

This pathway-level factor captures shared mevalonate-pathway dependency, so θ is conditional on the common pathway score. For MVK, the score was recomputed without MVK to avoid including the focal predictor in its own adjustment variable. This 13-gene set comprises HMGCR, HMGCS1, MVK, PMVK, MVD, IDI1, FDPS, GGPS1, FDFT1, SQLE, DHCR7, LSS, and CYP51A1.

4.1.4 CCLE metabolomics for orthogonal triangulation

To relate dependency signals to orthogonal biochemical readouts, we used the CCLE metabolomics resource described by Li et al. [16], downloaded through the DepMap/CCLE portal and public releases [2, 3]. This resource reports relative abundances for more than 225 metabolites across 928 cell lines. Metabolomics is not part of the primary adjustment model. We use metabolite associations as a separate triangulation analysis to contextualize the candidate-gene pattern (Results, Section 2.9). We focus on (i) polyamine-related metabolites (putrescine, ornithine, arginine) and (ii) sterol/lipid readouts, especially cholesteryl ester (CE) species and total CE. This is a pathway-guided triangulation analysis rather than a metabolome-wide screen; no multiplicity correction is applied across the full set of measured metabolites.

4.2 Estimation framework and robustness

4.2.1 Partial linear model

For a given candidate gene, let T_i denote its GeneEffect in cell line i and Y_i denote the statin phenotype (Y_{Δ} or related outcomes). We consider the partial linear model (PLM):

$$Y_i = \theta T_i + g(X_i) + \varepsilon_i, \qquad \mathbb{E}[\varepsilon_i \mid X_i, T_i] = 0, \tag{5}$$

where $g(\cdot)$ captures systematic covariate structure and θ is the adjusted slope relating dependency to statin selectivity. In DepMap, more negative T indicates stronger dependency. In PRISM,

more negative LFC indicates stronger growth inhibition. Therefore, if stronger dependency increases lipophilic statin sensitivity, both T and Y_Δ tend to decrease together. Under this PLM parameterization, that directional effect corresponds to a positive θ.

4.2.2 Interpretation and residual-support diagnostics

The parameter θ is interpreted as a covariate-adjusted association under the partial-linear model. Orthogonalization protects the slope estimate against first-order nuisance-model error under regularity conditions, but it does not make T randomized and does not rule out unmeasured common causes of dependency and drug response. The analysis also requires sufficient residual variation in T after adjustment; otherwise the slope is weakly supported by the observed data.

Residual-support assessment. In the complete-case strict-coverage sample used for these diagnostics ($n = 433$ after merging the outcome, focal dependency, and covariates; minimum coverage $n_{\text{lipo}} \geq 5$, $n_{\text{hydro}} \geq 2$), T is weakly predicted by X (out-of-fold $R^2 \approx 0$ for SLC45A4 and $R^2 < 0.06$ for the mevalonate comparators). Residual-variance ratios are near one (Table 23).

Summaries by major lineage and growth-rate quartile show overlapping SLC45A4 GeneEffect ranges (Tables 24–25; Fig. 19). Binary-split propensity-score summaries [26] likewise show limited fitted mass near 0 or 1 and high effective sample sizes (Table 26; Figs. 20–21). These checks indicate that the adjusted SLC45A4 slope is not driven by an almost deterministic predictor given the observed covariates; they do not address unmeasured confounding.

4.2.3 Double/debiased machine learning (DML)

The orthogonalized estimator allows separate nuisance models for $\ell(X) = \mathbb{E}[Y \mid X]$ and $m(X) = \mathbb{E}[T \mid X]$. Let:

$$m(X) = \mathbb{E}[T \mid X], \tag{6}$$

$$\ell(X) = \mathbb{E}[Y \mid X]. \tag{7}$$

Define residuals $V = T - m(X)$ and $U = Y - \ell(X)$. Under the partial-linear model, the adjusted slope satisfies the orthogonal moment:

$$\mathbb{E}\Big[(T - m(X))\big(Y - \ell(X) - \theta(T - m(X))\big)\Big] = 0. \tag{8}$$

Estimator with cross-fitting. We use K-fold cross-fitting [10]. Indices are split into folds $\mathcal{I}_1, \ldots, \mathcal{I}_K$. For each fold k, we fit auxiliary models $(\hat{m}^{(-k)}, \hat{\ell}^{(-k)})$ on data excluding $\mathcal{I}_k$ and compute out-of-fold residuals for $i \in \mathcal{I}_k$:

$$\tilde{T}_i = T_i - \hat{m}^{(-k)}(X_i), \qquad \tilde{Y}_i = Y_i - \hat{\ell}^{(-k)}(X_i). \tag{9}$$

The DML estimate solves:

$$\hat{\theta} = \frac{\frac{1}{n}\sum_{i=1}^{n} \tilde{T}_i \tilde{Y}_i}{\frac{1}{n}\sum_{i=1}^{n} \tilde{T}_i^2}. \tag{10}$$

For Table 1, we report the heteroskedasticity-robust (Neyman-orthogonal-score) standard error, $\widehat{\text{se}}(\hat{\theta}) = \sqrt{\sum_i \tilde{T}_i^2 \hat{e}_i^2} / \sum_i \tilde{T}_i^2$ with $\hat{e}_i = \tilde{Y}_i - \hat{\theta}\tilde{T}_i$. This is the standard variance estimator for the partialling-out score and does not assume constant residual variance across cell lines. For comparison, the homoskedastic residual-based variance is $\hat{\sigma}^2 / \sum_i \tilde{T}_i^2$, with a single pooled $\hat{\sigma}^2 =$

$\sum_i \hat{e}_i^2/(n-1)$. The two calculations agree closely in this dataset (Supplementary Section S6). We report two-sided normal-approximation p-values. As a stability check, we also repeat cross-fitting for SLC45A4 with four additional random seeds; $\hat{\theta}$ remains positive and $p < 0.05$ under the robust variance in every seed (Supplementary Section S6).

Implementation details. The primary implementation uses five-fold K-fold splitting with shuffling and random seed 0. Both nuisance models are ridge regressions with fixed penalty $\alpha = 1.0$; the penalty is not tuned by an inner cross-validation loop. Y and T are winsorized at their 1st and 99th percentiles before fitting. The primary BH values are computed across the four reported genes only [27].

4.2.4 Per-drug meta-analysis of the lipophilic–hydrophilic contrast

As an estimation route that does not depend on averaging LFC values across drugs before fitting, we applied the same cross-fitted partial-linear specification (Section 4.2.3) separately to each of the eight statins. Each fit used that drug's own LFC as Y and the focal gene's GeneEffect as T, with the same lineage, growth, and mevalonate-core covariates. This produced eight per-drug estimates $(\hat{\theta}_d, \widehat{\mathrm{se}}_d)$ for each focal gene. Within each statin class (lipophilic, $k = 6$; hydrophilic, $k = 2$), we combined the per-drug estimates using DerSimonian–Laird random-effects meta-analysis, which weights each drug by the inverse of its within-drug variance plus an estimated between-drug heterogeneity variance τ^2, giving class-level estimates $\hat{\theta}_{RE,\mathrm{lipo}}$ and $\hat{\theta}_{RE,\mathrm{hydro}}$ and their difference.

The eight per-drug fits are not independent studies. They share the same cell-line panel, the same GeneEffect column T, and the same covariates X, so their nuisance-model residuals are correlated even before accounting for any correlation between drug responses themselves. A Wald test that treats $\hat{\theta}_{RE,\mathrm{lipo}}$ and $\hat{\theta}_{RE,\mathrm{hydro}}$ as independent, $z = (\hat{\theta}_{RE,\mathrm{lipo}} - \hat{\theta}_{RE,\mathrm{hydro}})/\sqrt{\widehat{\mathrm{se}}^2_{RE,\mathrm{lipo}} + \widehat{\mathrm{se}}^2_{RE,\mathrm{hydro}}}$, therefore understates the true sampling variability of the difference. We instead obtain inference by a cell-line-clustered bootstrap. We resample cell lines with replacement (the full panel size, $B = 1000$ resamples, random seed 0). On each resampled panel, we rerun the entire pipeline – all eight per-drug DML fits, both subgroup meta-analyses, and their difference – to build an empirical distribution of $\hat{\theta}_{RE,\mathrm{lipo}} - \hat{\theta}_{RE,\mathrm{hydro}}$. We report the bootstrap standard deviation as the standard error and a two-sided p-value from the ratio of the observed difference to this bootstrap SE under a normal approximation. A percentile-based p-value (twice the smaller tail fraction below/above zero) and the percentile 95% interval agree closely with the normal-approximation value in all four genes (Supplementary Section S6). BH values are computed across the four reported genes. With only two hydrophilic drugs, the hydrophilic-class estimate is itself imprecise, which widens the bootstrap interval for the difference. We report this check as a supportive robustness analysis and do not treat it as a second confirmatory test with the same evidential weight as Table 1.

4.2.5 External AUC consistency, LODO, and empirical reference analyses

These analyses ask whether coefficient directions persist under alternative response measurements, whether a single drug drives the result, and how the observed statistics compare with empirical reference distributions. They are not treated as independent replication of the primary association.

External AUC outcomes. We extracted CTRP_AUC [28, 29] and Repurposing_secondary_AUC through the DepMap portal data slicer (release 24Q4) [2, 3]. We matched lipophilic statins by drug name and computed the mean AUC across available matched drugs. We selected the external-outcome orientation that maximized its correlation with the primary PRISM phenotype over shared

cell lines. All orientation factors equal +1, so this step did not numerically reverse an outcome; nevertheless, orientation and sign-match summaries are outcome-informed and descriptive. We then re-estimated the adjusted coefficients with the same covariate set.

Leave-one-drug-out (LODO). To assess whether associations are driven by a single compound, we repeated external estimation after removing each lipophilic statin column in turn. We summarize the range of $\hat{\theta}$ and sign stability across LODO folds.

Negative-control drug bundles. For the drug-level empirical reference analysis, we generated $B = 10000$ random bundles of the same size as the lipophilic statin set from compounds with the matched screening type (random seed 0). The analysis uses the same covariate specification as the primary analysis (lineage indicators, growth rate, and mevalonate-core PC1; Section 4.2). Table 22 reports its analysis-specific observed q values and +1-corrected empirical tail probabilities, $p_{\text{emp}} = (1 + \#\{q_b \leq q_{\text{obs}}\})/(B + 1)$. These q values are not the four-gene BH values in Table 1. This analysis tests a combined eight-drug meta-analytic signal, a different quantity from the lipophilic–hydrophilic contrast. The finer-resolution gene-level empirical null ($N = 1000$ random genes; Table 21) remains the primary empirical reference evidence in this manuscript; the random-drug-bundle check (Supplementary Section S6) and the per-drug meta-analysis (Section 4.2.4) are directionally consistent supporting checks; neither is an independent confirmation.

External positive-control checks. As a positive-control check on the external AUC matrices and integration procedure, we analyzed CTRP_AUC and the Sanger GDSC1/GDSC2 AUC matrices (all accessed via the DepMap portal). For each dataset, we formed drug-set outcomes as the median AUC across (i) reversible EGFR inhibitors (erlotinib, gefitinib) or a broader EGFR/HER set (adding afatinib and lapatinib when available), and (ii) BRAF inhibitors (vemurafenib/PLX-4032, PLX 4720, dabrafenib when available). We standardized outcomes within each dataset and estimated the DML effect of EGFR or BRAF GeneEffect using the same covariates as in the main analysis. We assessed robustness via within-lineage permutation of the residualized focal predictor and report empirical tail probabilities (Supplementary Section S8; Table 27).

4.2.6 Lineage-stratified secondary dose–response parameter analysis

To examine lineage heterogeneity and relate selectivity to standard dose–response summaries, we analyzed PRISM secondary dose–response parameters for lipophilic statins. For each cell line and statin, the curve fit provides six metrics: AUC, EC50, slope, fitted R^2, and the upper and lower asymptotes (upper_limit, lower_limit). We present the results focused on AUC, lower_limit, and EC50 because they have the clearest sensitivity and potency interpretations. All six metrics were inspected during this analysis, however, so the BH-adjusted family in Section 2.6 spans all six metrics across the 8 tested lineages (48 tests per model), not only the three metrics presented in the main text. Results for AUC, lower_limit, and EC50 are reported in Supplementary Section S3. All six curve metrics were included in the 48-test BH correction. We aggregated lipophilic-statin metrics within each cell line (median across lipophilic statins) and analyzed the association between dependency strength and each metric.

Because GeneEffect is oriented such that more negative values indicate stronger dependency, we define dependency strength as $D = -\text{GeneEffect}$ (larger D indicates stronger dependency). Within each lineage (restricted to lineages with at least 15 cell lines), we fit (i) a univariate regression of each metric on $D_{SLC45A4}$ and (ii) a multivariable regression additionally adjusting for D_{UBIAD1}, D_{MVK}, and D_{HMGCR}. We report HC3 standard errors. Given the small per-lineage sample sizes

(as low as $n = 21$), we compute p-values from a t-distribution using each model's residual degrees of freedom, appropriate at this sample size. We report BH-adjusted q-values across the full family of 8 tested lineages $\times$ all 6 curve metrics inspected (48 tests per model; Section 2.6).

4.2.7 Co-dependency ranking, enrichment, and empirical null

To test whether phenotype-defined subsets reflect distinct dependency programs, we computed gene–gene co-dependency as Pearson correlation between GeneEffect vectors across cell lines within a lineage. Intuitively, two genes are "co-dependent" when the same cell lines tend to be sensitive to knockout of both genes. We then tested whether top co-dependent genes were over-represented in known pathways, relative to chance under the tested-gene background.

For a target gene (SLC45A4), we ranked all genes by $|r|$ with the target and defined the "top200" co-dependency set.

We tested enrichment using a study-defined 21-gene mevalonate/cholesterol-biosynthesis set. The set comprises the 13-gene adjustment set listed in Section 4.1.3 (HMGCR, HMGCS1, MVK, PMVK, MVD, IDI1, FDPS, GGPS1, FDFT1, SQLE, DHCR7, LSS, and CYP51A1), together with *ACAT2* (acetyl-CoA acetyltransferase 2), *DHCR24*, *EBP*, *HSD17B7*, *MSMO1*, *NSDHL*, *SC5D*, and *TM7SF2*. The 13-gene adjustment set is a strict subset of this 21-gene enrichment set. The two sets serve different purposes: the 13-gene set defines the pathway-level adjustment covariate (Section 4.1.3), whereas the 21-gene set defines the enrichment target used here. The 21-gene enrichment set reflects established mevalonate/cholesterol pathway biology [12, 13] and was applied unchanged across multiple colorectal subsets during this analysis, including an earlier all-colorectal check that showed no top-200 overlap; it was not adjusted after observing the phenotyped-subset result reported here. The gene universe was defined by genes present in the GeneEffect matrix. This test quantifies how often random draws of the same size would contain at least as many pathway genes. As a complementary broader scan, we also ran Enrichr pathway enrichment on the top200 set [22, 23].

Because the phenotyped subset (cell lines with available dose–response parameters) is a non-random subset of colorectal lines, theoretical p-values alone can be misleading. We therefore built an *empirical null* by subset permutation. We repeatedly drew random colorectal subsets of the same size as the phenotyped subset. For each resampled subset, we recomputed the top200 co-dependency set for SLC45A4 and recalculated its overlap with the study-defined sterol/mevalonate set. We report empirical tail probabilities for (i) $k \geq k_{\text{obs}}$ overlaps and (ii) simultaneous inclusion of four sterol biosynthesis genes (FDFT1, SQLE, NSDHL, LSS) in the top200 set.

Data availability

All raw data used in this study are publicly available. DepMap 24Q4 gene-dependency profiles, cell-line annotations, growth rates, drug-sensitivity AUC matrices, and CCLE metabolomics data are available through the DepMap portal (`https://depmap.org/`). PRISM primary and secondary screening data are available through the DepMap Repurposing resource (`https://depmap.org/repurposing/`). CTRP and GDSC drug-sensitivity data were accessed through the DepMap portal. Dataset releases and identifiers are specified in the Methods and references.

Competing interests

The author declares no competing interests.

Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the author used generative AI and AI-assisted tools to assist with manuscript drafting, language refinement, code drafting and debugging, and discussion of the research. The author directed the project, reviewed, revised, and verified the manuscript and code, and takes full responsibility for the final manuscript.

References

[1] S. M. Corsello, J. A. Bittker, Z. Liu, et al. Discovering the anticancer potential of non-oncology drugs by systematic viability profiling. *Nature Cancer*, 1:235–248, 2020.

[2] DepMap Consortium. DepMap portal data downloads. Online resource, accessed March 2026. `https://depmap.org/portal/data_page/?tab=allData`.

[3] Broad DepMap. DepMap 24Q4 Public. Figshare+ dataset (2024). `https://doi.org/10.25452/figshare.plus.27993248.v1`.

[4] E. Goncalves, B. Segura-Cabrera, D. Pacini, et al. Drug mechanism-of-action discovery through the integration of pharmacological and CRISPR screens. *Molecular Systems Biology*, 16:e9405, 2020. doi:10.15252/msb.20199405.

[5] M. J. Garnett, E. J. Edelman, S. J. Heidorn, et al. Systematic identification of genomic markers of drug sensitivity in cancer cells. *Nature*, 483:570–575, 2012.

[6] F. Iorio, T. Knijnenburg, D. J. Vis, et al. A landscape of pharmacogenomic interactions in cancer. *Cell*, 166(3):740–754, 2016.

[7] D. G. Covell. Data mining approaches for genomic biomarker development: applications using drug screening data from the Cancer Genome Project and the Cancer Cell Line Encyclopedia. *PLOS ONE*, 10(7):e0127433, 2015. doi:10.1371/journal.pone.0127433.

[8] A. Bazaga, D. Leggate, and H. Weisser. Genome-wide investigation of gene-cancer associations for the prediction of novel therapeutic targets in oncology. *Scientific Reports*, 10:10787, 2020. doi:10.1038/s41598-020-67846-1.

[9] M. Hafner, M. Niepel, M. Chung, and P. K. Sorger. Growth rate inhibition metrics correct for confounders in measuring drug sensitivity. *Nature Methods*, 13(6):521–527, 2016. doi:10.1038/nmeth.3853.

[10] V. Chernozhukov, D. Chetverikov, M. Demirer, et al. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018.

[11] M.-F. Demierre, P. D. R. Higgins, S. B. Gruber, E. Hawk, and S. M. Lippman. Statins and cancer prevention. *Nature Reviews Cancer*, 5(12):930–942, 2005. doi:10.1038/nrc1751.

[12] P. J. Mullen, R. Yu, J. Longo, M. C. Archer, and L. Z. Penn. The interplay between cell signalling and the mevalonate pathway in cancer. *Nature Reviews Cancer*, 16(11):718–731, 2016. doi:10.1038/nrc.2016.76.

[13] D. Juarez and D. A. Fruman. Targeting the mevalonate pathway in cancer. *Trends in Cancer*, 7(6):525–540, 2021. doi:10.1016/j.trecan.2020.11.008.

[14] V. K. Raghu, C. H. Beckwitt, K. Warita, A. Wells, P. V. Benos, and Z. N. Oltvai. Biomarker identification for statin sensitivity of cancer cell lines. *Biochemical and Biophysical Research Communications*, 495(1):659–665, 2018. doi:10.1016/j.bbrc.2017.11.065.

[15] M. I. Abdullah, E. de Wolf, M. J. Jawad, and A. Richardson. The poor design of clinical trials of statins in oncology may explain their failure – lessons for drug repurposing. *Cancer Treatment Reviews*, 69:84–89, 2018. doi:10.1016/j.ctrv.2018.06.010.

[16] H. Li, S. Ning, M. Ghandi, et al. The landscape of cancer cell line metabolism. *Nature Medicine*, 25(5):850–860, 2019. doi:10.1038/s41591-019-0404-8.

[17] A. J. Ohnmacht, G. Avar, M. K. Schübel, et al. The pharmacogenomic assessment of molecular epithelial-mesenchymal transition signatures reveals drug susceptibilities in cancer cell lines. *bioRxiv*, 2024. doi:10.1101/2024.01.16.575190.

[18] K. Berg, L. Sakellaridi, T. Rummel, et al. Identifying modulators of cellular responses by heterogeneity-sequencing. *bioRxiv*, 2024. doi:10.1101/2024.10.28.620481.

[19] Z. Fu and L. Hou. Inference of heterogeneous effects in single-cell genetic perturbation screens. *bioRxiv*, 2024. doi:10.1101/2024.05.28.596224.

[20] M. Dong, B. Wang, J. Wei, et al. Causal identification of single-cell experimental perturbation effects with CINEMA-OT. *Nature Methods*, 20(11):1769–1779, 2023. doi:10.1038/s41592-023-02040-5.

[21] E. Y. Chen, C. M. Tan, Y. Kou, Q. Duan, Z. Wang, G. V. Meirelles, N. R. Clark, and A. Ma'ayan. Enrichr: interactive and collaborative HTML5 gene list enrichment analysis tool. *BMC Bioinformatics*, 14:128, 2013. doi:10.1186/1471-2105-14-128.

[22] M. V. Kuleshov, M. R. Jones, A. D. Rouillard, N. F. Fernandez, Q. Duan, Z. Wang, S. Koplev, S. L. Jenkins, K. M. Jagodnik, A. Lachmann, M. G. McDermott, C. D. Monteiro, G. W. Gundersen, and A. Ma'ayan. Enrichr: a comprehensive gene set enrichment analysis web server 2016 update. *Nucleic Acids Research*, 44(W1):W90–W97, 2016. doi:10.1093/nar/gkw377.

[23] Z. Xie, A. Bailey, M. V. Kuleshov, D. J. B. Clarke, J. E. Evangelista, S. L. Jenkins, A. Lachmann, M. L. Wojciechowicz, E. Kropiwnicki, K. M. Jagodnik, M. Jeon, and A. Ma'ayan. Gene set knowledge discovery with Enrichr. *Current Protocols*, 1(3):e90, 2021. doi:10.1002/cpz1.90.

[24] S. J. Middleton, S. Markússon, M. Åkerlund, et al. SLC45A4 is a pain gene encoding a neuronal polyamine transporter. *Nature*, 646(8084):404–412, 2025. doi:10.1038/s41586-025-09326-y.

[25] C. Colson, Y. Wang, J. Atherton, et al. SLC45A4 encodes a peroxisomal putrescine transporter that promotes GABA de novo synthesis. *Nature Communications*, 16:10198, 2025. doi:10.1038/s41467-025-62721-x.

[26] P. Rosenbaum and D. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.

[27] Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B*, 57(1):289–300, 1995.

[28] A. Basu, N. E. Bodycombe, J. H. Cheah, et al. An interactive resource to identify cancer genetic and lineage dependencies targeted by small molecules. *Cell*, 154(5):1151–1161, 2013.

[29] B. Seashore-Ludlow, M. G. Rees, J. H. Cheah, et al. Harnessing connectivity in a large-scale small-molecule sensitivity dataset. *Cancer Discovery*, 5(11):1210–1223, 2015.

A Supplementary material

This appendix provides supporting analyses, diagnostics, and robustness checks for the main text.

S1 Statin mapping and outcome definition

Table 8 makes the drug mapping explicit and reproducible (drug name → lipophilicity class). Because PRISM primary includes only two hydrophilic statins, this mapping also motivates stricter-coverage and leave-one-hydrophilic-out checks reported later (Table 20, Fig. 14).

Table 8: PRISM statins used in this study and their lipophilicity class. Every Broad ID and LFC column was checked directly against PRISM's own `primary_treatment_info.csv` metadata and matched verbatim.

Drug	Class	PRISM Broad ID	LFC column
atorvastatin	lipophilic	BRD-K69726342-238-02-4	BRD-K69726342-238-02-4::2.5::HTS
fluvastatin	lipophilic	BRD-K66296774-236-11-3	BRD-K66296774-236-11-3::2.63::HTS
lovastatin	lipophilic	BRD-K09416995-001-39-9	BRD-K09416995-001-39-9::2.5::HTS
mevastatin	lipophilic	BRD-K94441233-001-13-0	BRD-K94441233-001-13-0::2.6::HTS
pitavastatin	lipophilic	BRD-K75958547-238-01-0	BRD-K75958547-238-01-0::2.5::HTS
pravastatin	hydrophilic	BRD-K60511616-236-08-9	BRD-K60511616-236-08-9::2.5::HTS
rosuvastatin	hydrophilic	BRD-K82941592-238-04-5	BRD-K82941592-238-04-5::2.5::HTS
simvastatin	lipophilic	BRD-K22134346-001-22-3	BRD-K22134346-001-22-3::2.5::HTS

Notes: "PRISM Broad ID" corresponds to the PRISM metadata field `broad_id` (Broad Institute compound identifier). "LFC column" is the treatment column name in the PRISM primary replicate-collapsed log fold-change (LFC) matrix, typically formatted as `broad_id::dose::screen_id` (e.g., `BRD-...::2.5::HTS`).

S2 External outcome orientation

External AUC summaries can have a numerical orientation different from PRISM LFC. We selected the orientation that maximized correlation with the primary phenotype over shared cell lines. Table 9 records the resulting factors and correlations. Every factor is +1, so no numerical sign reversal occurred; the exercise remains descriptive because the primary outcome informed the orientation rule.

Table 9: Outcome correlations and orientation factors for the primary and AUC-based outcomes.

dataset	lipo_n_thr	n	sign	corr
CTRP_AUC	1	421	1.000	0.451
CTRP_AUC	3	361	1.000	0.488
Repurposing_secondary_AUC	1	397	1.000	0.746
Repurposing_secondary_AUC	3	342	1.000	0.771

Notes: `CTRP_AUC` denotes the Cancer Therapeutics Response Portal AUC summary accessed via DepMap. `Repurposing_secondary_AUC` denotes the PRISM Repurposing Secondary Screen AUC. `lipo_n_thr` is the minimum number of lipophilic statins required for the AUC average; `n` is the number of shared cell lines. "sign" is the multiplicative orientation factor chosen from the primary-outcome correlation; all factors equal +1.

S3 Lineage-stratified dose–response coefficients

The main text summarizes lineage-specific effects in compact figures. Tables 10–12 provide full coefficient tables across lineages for AUC, lower_limit, and EC50, showing where effects are near zero and where they are larger. The tables show heterogeneity in fitted associations across lineages; they do not constitute a multiplicity-adjusted interaction test.

Table 10: Lineage-stratified AUC estimates for SLC45A4 (multivariable HC3) in PRISM secondary dose–response parameters, all 8 tested lineages. p uses a t-distribution (df= $n-5$). q is BH-adjusted across all 8 lineages × 6 metrics (48 tests per model; Methods), not only the AUC/lower_limit/EC50 metrics shown in the three tables in this section.

Lineage	n	$\hat{\theta}$	SE	p (t)	q (BH)
colorectal	21	-1.021	0.203	1.22e-04	5.84e-03
pancreas	29	-0.658	0.402	1.14e-01	6.86e-01
ovary	29	-0.275	0.168	1.14e-01	6.86e-01
urinary tract	21	-0.309	0.290	3.01e-01	8.00e-01
esophagus	21	0.667	0.768	3.98e-01	8.00e-01
lung	80	-0.027	0.145	8.51e-01	9.50e-01
central nervous system	31	0.093	0.207	6.57e-01	8.61e-01
skin	29	0.228	0.524	6.67e-01	8.61e-01

Table 11: Lineage-stratified lower_limit estimates for SLC45A4 (multivariable HC3) in PRISM secondary dose–response parameters, all 8 tested lineages. p uses a t-distribution (df= $n-5$); q is BH-adjusted jointly with Table 10 across the full 48-test family as described in Methods.

Lineage	n	$\hat{\theta}$	SE	p (t)	q (BH)
colorectal	21	-3.640	0.823	4.25e-04	1.02e-02
ovary	29	-1.285	0.635	5.43e-02	4.35e-01
pancreas	29	-1.508	1.360	2.79e-01	8.00e-01
lung	80	-0.549	0.514	2.90e-01	8.00e-01
skin	29	1.314	1.524	3.97e-01	8.00e-01
central nervous system	31	0.463	0.617	4.59e-01	8.00e-01
urinary tract	21	-0.307	1.199	8.01e-01	9.16e-01
esophagus	21	0.673	2.193	7.63e-01	9.08e-01

Table 12: Lineage-stratified EC50 estimates for SLC45A4 (multivariable HC3) in PRISM secondary dose–response parameters, all 8 tested lineages. p uses a t-distribution (df= $n-5$); q is BH-adjusted jointly with Tables 10–11 across the full 48-test family, which also includes fitted R^2, slope, and upper_limit at each of the 8 lineages.

Lineage	n	$\hat{\theta}$	SE	p (t)	q (BH)
urinary tract	21	-2.307	3.080	4.65e-01	8.00e-01
pancreas	29	-0.946	2.279	6.82e-01	8.61e-01
central nervous system	31	-0.900	1.336	5.07e-01	8.11e-01
skin	29	0.181	4.117	9.65e-01	9.86e-01
colorectal	21	0.213	3.734	9.55e-01	9.86e-01
lung	80	1.528	1.019	1.38e-01	7.12e-01
ovary	29	1.788	2.000	3.80e-01	8.00e-01
esophagus	21	1.808	2.449	4.71e-01	8.00e-01

S4 Empirical null for the co-dependency enrichment signal

The phenotyped colorectal subset is defined by data availability, so apparent enrichment could be a subsetting artifact. Tables 13–14 report the empirical null from repeated random subsetting and the inclusion rates of sterol genes. Together, they quantify how unusual the observed sterol overlap is under a random-subset explanation.

Table 13: Empirical null distribution for cholesterol-core overlap k under subset permutation (2000 random colorectal subsets).

k	count	probability
0	1416	0.708
1	438	0.219
2	110	0.055
3	26	0.013
4	8	0.004
5	2	0.001

Table 14: Sterol biosynthesis genes: probability of inclusion in the SLC45A4 top200 co-dependency set under the subset-permutation null.

Gene	P(in top200 under null)
LSS	0.0310
FDFT1	0.0335
NSDHL	0.0435
SQLE	0.0620

S5 Within-colorectal metabolomics analyses: direct correlation and attenuation tests

The main text reports metabolomic signatures associated with SLC45A4 dependency in colorectal lines (higher CE and lower putrescine). Figure 13 and Table 15 test a stronger claim—a direct within-colorectal association between putrescine and CE features—and do not support it. Table 16 further shows that conditioning on putrescine does not systematically attenuate the dependency×CRC interaction for CE features. Together, these analyses support a coherent signature but not a single-step mediation chain.

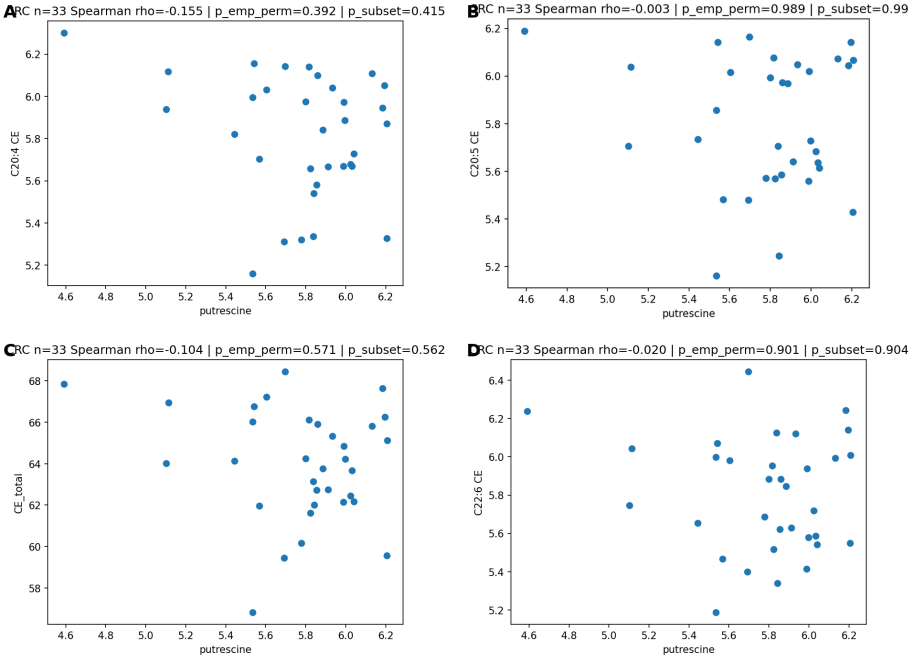


Figure 13: Within-colorectal scatterplots: putrescine vs. cholesteryl ester (CE) readouts.

Table 15: Within-colorectal correlation between putrescine and cholesteryl ester readouts. p_{perm} permutes putrescine within colorectal; p_{subset} is a non-colorectal subset null.

CE feature	n_{CRC}	ρ	p	p_{perm}	p_{subset}
C20:4 CE	33	-0.155	0.390	0.392	0.415
CE_total	33	-0.104	0.564	0.571	0.562
C22:6 CE	33	-0.020	0.913	0.901	0.904
C20:5 CE	33	-0.003	0.987	0.989	0.990

Table 16: Testing whether putrescine adjustment attenuates the SLC45A4×CRC interaction for CE features (empirical shrinkage null).

CE feature	$\hat{\beta}_{D\times CRC}$ (base)	$\hat{\beta}_{D\times CRC}$ (adj. PUT)	$\Delta\lvert\beta\rvert$	$p_{\text{shrink,emp}}$	$p_{\text{PUT}\times CRC}$
C20:5 CE	1.513	1.801	-0.288	1.000	0.354
C20:4 CE	1.484	1.437	0.047	0.287	0.967
CE_total	10.755	11.206	-0.451	1.000	0.888
C22:6 CE	1.190	1.286	-0.096	1.000	0.714

S6 Robustness, sensitivity, and negative-control analyses

Robust versus homoskedastic standard errors for the primary DML estimate. Table 1 reports the heteroskedasticity-robust (Neyman-orthogonal-score) standard error. Table 17 compares this estimate with the conventional homoskedastic residual-based standard error for all four genes. The two agree closely: SLC45A4's p-value moves from 0.0199 (homoskedastic) to 0.0213 (robust), and its BH q from 0.0265 to 0.0285; UBIAD1, MVK, and RIN1 are essentially unchanged. We also reran cross-fitting for SLC45A4 with four additional random seeds (Table 18); $\hat{\theta}$ stayed positive and the robust p-value stayed below 0.05 in every seed, ranging from 0.009 to 0.024.

Table 17: Homoskedastic versus heteroskedasticity-robust standard errors for the primary DML estimate (Table 1), random seed 0.

Gene	$\hat{\theta}$	SE (homo.)	p (homo.)	SE (robust)	p (robust)	q (robust)
UBIAD1	0.834	0.108	1.14e-14	0.116	7.20e-13	2.88e-12
MVK	0.733	0.125	4.40e-09	0.133	3.21e-08	6.42e-08
SLC45A4	1.043	0.448	1.99e-02	0.453	2.13e-02	2.85e-02
RIN1	-0.334	0.341	3.28e-01	0.340	3.26e-01	3.26e-01

Table 18: Cross-fitting seed stability for SLC45A4, robust standard errors, primary specification.

Seed	$\hat{\theta}$	SE (robust)	p (robust)
0	1.043	0.453	0.0213
1	1.205	0.458	0.0085
2	1.000	0.443	0.0240
3	1.095	0.441	0.0130
4	1.137	0.444	0.0104

Cell-line-clustered bootstrap for the per-drug analysis. Because the eight per-drug fits share the same cell-line panel, GeneEffect column, and covariates, an independence-assuming Wald test understates the sampling variability of the lipophilic-minus-hydrophilic difference. The manuscript therefore uses a cell-line-clustered bootstrap: $B = 1000$ resamples of cell lines (with replacement, full panel size, random seed 0), refitting all eight per-drug DML models and both subgroup meta-analyses on each resampled panel. Table 19 compares the naive Wald result with the bootstrap result for all four genes.

Table 19: Naive (independence-assuming) Wald test versus cell-line-clustered bootstrap ($B = 1000$) for the lipophilic-minus-hydrophilic difference in the per-drug meta-analysis. Point estimates ($\hat{\theta}$) are unchanged; only the inference changes.

Gene	$\hat{\theta}$ (diff.)	p (naive Wald)	SE (bootstrap)	p (bootstrap)
UBIAD1	0.801	1.3e-10	0.122	5.7e-11
MVK	0.706	7.6e-07	0.142	6.7e-07
SLC45A4	0.921	1.3e-02	0.498	6.5e-02
RIN1	-0.449	4.1e-03	0.313	1.51e-01

UBIAD1 and MVK remain highly significant under either method. The shared-structure correction increases their bootstrap SE relative to the naive Wald SE, but their effects are large enough that both remain far below conventional thresholds. SLC45A4 and RIN1 tell a different story. SLC45A4's naive Wald p-value (0.013) is roughly five times smaller than its bootstrap p-value (0.065); with the correlation between per-drug fits properly accounted for, the difference no longer reaches conventional significance, though it remains directionally consistent with the primary result. For RIN1, the naive Wald p-value is 0.004, whereas the bootstrap p-value is 0.15, consistent with its non-significant coefficient in Table 1. This contrast shows that the independence assumption is materially anti-conservative in this setting, so bootstrap-based inference is reported as primary throughout the manuscript.

Table 20 and Figs. 14–15 test whether the primary SLC45A4 selectivity estimate depends on outcome definition (minimum statin coverage) or on any single compound in the bundle. Table 21 and Figs. 16–17 compare the SLC45A4 estimate with 1000 randomly sampled genes analyzed under the same procedure. This provides an empirical reference for how extreme the observed estimate is relative to generic gene–drug associations.

Sign convention for $\hat{\theta}$. Table 1 parameterizes the focal predictor as $T = \mathsf{GeneEffect}$ (more negative values indicate stronger dependency). Under the PLM convention in Methods (Section 4.2.1), stronger dependency implying greater vulnerability corresponds to $\hat{\theta} > 0$. In the strict-sensitivity and negative-control analyses below, we use dependency strength $D = -\mathsf{GeneEffect}$ for interpretability; under $D = -T$, the same directional effect appears as $\hat{\theta} < 0$.

Table 20: Sensitivity of the SLC45A4 estimate in PRISM primary selectivity to strict statin coverage and leave-one-drug-out/leave-one-hydrophilic-out definitions. The "Full" analysis requires MIN_N_LIPO_PRIMARY= 5 and MIN_N_HYDRO_PRIMARY= 2 (minimum numbers of non-missing lipophilic/hydrophilic statin LFC columns per cell line used to compute Y_Δ). Estimates are parameterized using dependency strength $D = -\mathsf{GeneEffect}$ (sign inverted relative to Table 1).

Analysis	Dropped drug	n	$\hat{\theta}$	SE	p
Full (strict coverage)	–	437	-0.980	0.485	0.043
LODO (lipophilic)	atorvastatin	437	-1.221	0.528	0.021
LODO (lipophilic)	simvastatin	437	-1.007	0.479	0.036
LODO (lipophilic)	lovastatin	437	-1.178	0.534	0.027
LODO (lipophilic)	fluvastatin	437	-0.780	0.464	0.093
LODO (lipophilic)	pitavastatin	437	-0.712	0.449	0.113
LOHO (hydrophilic)	pravastatin	442	-1.075	0.508	0.034
LOHO (hydrophilic)	rosuvastatin	437	-1.009	0.502	0.044

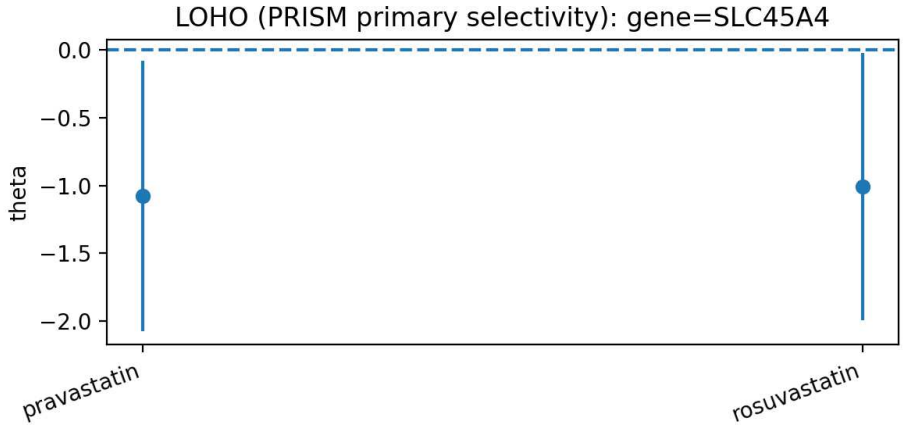


Figure 14: Leave-one-hydrophilic-out (LOHO) checks for SLC45A4 in the strict primary selectivity analysis.

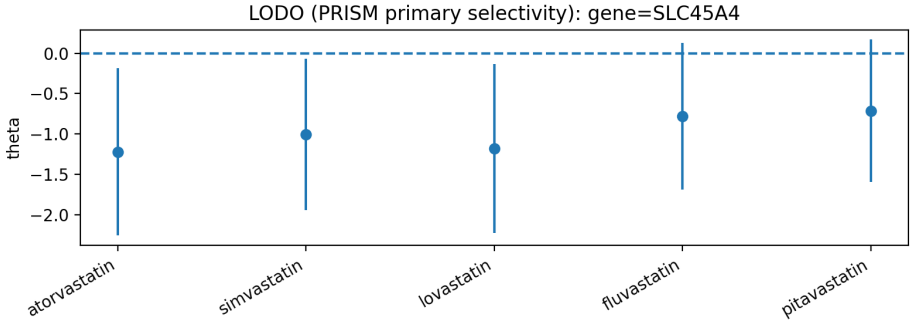


Figure 15: Leave-one-lipophilic-out (LODO) checks for SLC45A4 in the strict primary selectivity analysis.

Table 21: Gene-level negative-control comparison using 1000 randomly sampled genes analyzed with the same procedure. Effects are estimated using dependency strength $D = -\mathsf{GeneEffect}$; thus, a more negative $\hat{\theta}$ corresponds to stronger dependency associated with greater vulnerability. p^{θ}_{emp} is the fraction of negative controls with $\hat{\theta}$ more negative than SLC45A4; p^{z}_{emp} is defined analogously for z-scores.

Outcome	n	$\hat{\theta}$	z	p^{θ}_{emp}	p^{z}_{emp}
Primary selectivity (PRISM primary; strict)	437	-0.980	-2.02	0.030	0.051
Secondary AUC (PRISM secondary dose–response)	402	-0.168	-2.44	0.019	0.031

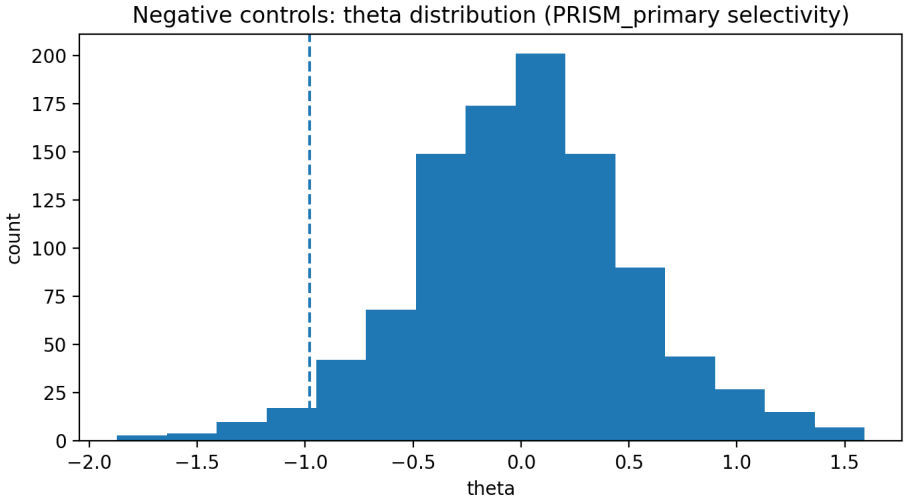


Figure 16: Negative-control histogram (primary selectivity): distribution of $\hat{\theta}$ over 1000 random-gene controls, with SLC45A4 overlaid.

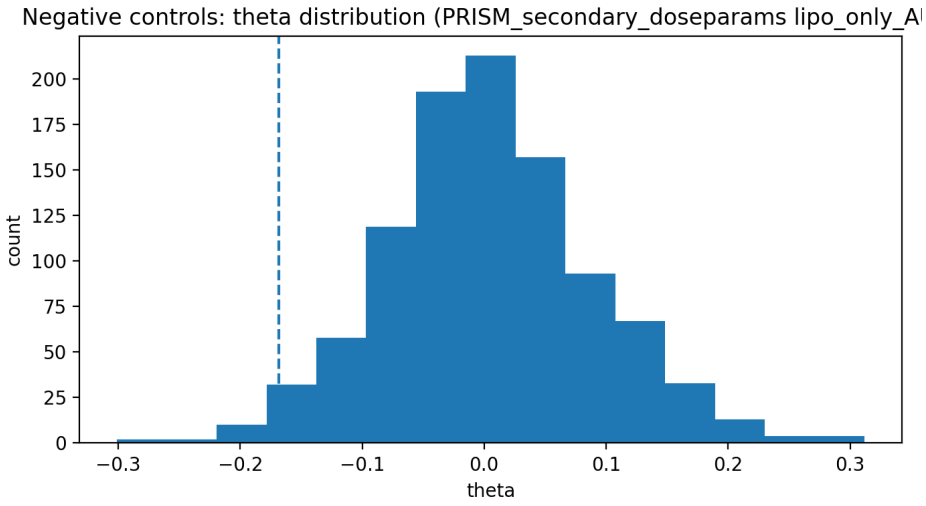


Figure 17: Negative-control histogram (secondary AUC): distribution of $\hat{\theta}$ over 1000 random-gene controls, with SLC45A4 overlaid.

Random-drug-bundle check ($B = 10000$). We compare each gene's coefficient with $B = 10000$ random bundles of other drugs, using the same covariate specification as the primary analysis (Table 22, Fig. 18; Methods, Section 4.2.5). This analysis tests a single combined eight-drug meta-analytic signal, a different quantity from the lipophilic–hydrophilic contrast, so its observed q values are not the four-gene q values in Table 1.

With $B = 10000$, UBIAD1, MVK, and RIN1 are more extreme than nearly all random bundles ($p_{\text{emp}} \leq 0.001$), while SLC45A4 is less extreme but still in the upper tail of the null distribution ($p_{\text{emp}} = 0.003$, stable in the range 0.003–0.005 across three random seeds). This pattern is consistent with the primary table, where SLC45A4 is also the weakest of the three genes distinguishable from zero. We treat this analysis as a corroborating check; the primary empirical reference evidence remains the gene-level null (Table 21).

Table 22: Random-drug-bundle empirical reference check at $B = 10000$ using the primary covariate specification. Observed q values are analysis-specific and are not the four-gene q values in Table 1.

Gene	Observed q^*	p_{emp}	B
MVK	2.99e-04	7.00e-04	10000
RIN1	3.53e-04	1.00e-04	10000
SLC45A4	5.36e-03	2.70e-03	10000
UBIAD1	7.06e-05	1.00e-03	10000

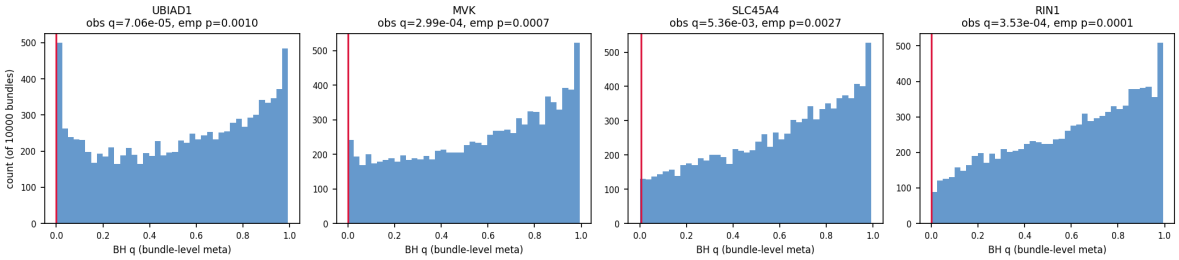


Figure 18: Null distributions of analysis-specific BH q over 10000 random drug bundles, by gene. Red lines mark each gene's observed analysis-specific q; titles report the observed q and the resulting empirical p.

S7 Residual-support assessment for the focal predictor

The adjusted slope requires residual variation in the focal predictor T after conditioning on X. We report (i) out-of-sample predictability of T from X, (ii) predictor ranges by lineage and growth bins, and (iii) propensity-score summaries [26] for binary splits of T. These diagnostics assess observed support. They cannot assess unmeasured confounding.

Table 23: Predictability of the focal gene-dependency measure from covariates: cross-validated R^2 and the residual-variance ratio. Values near zero R^2 and ratios near one indicate substantial residual variation in T given X.

gene	n	cv_R2	resid_var_ratio	resid_sd	D_sd
MVK	433	0.059	0.941	0.362	0.373
HMGCR	433	0.031	0.968	0.293	0.298
UBIAD1	433	-0.029	1.029	0.421	0.415
SLC45A4	433	-0.032	1.032	0.104	0.102

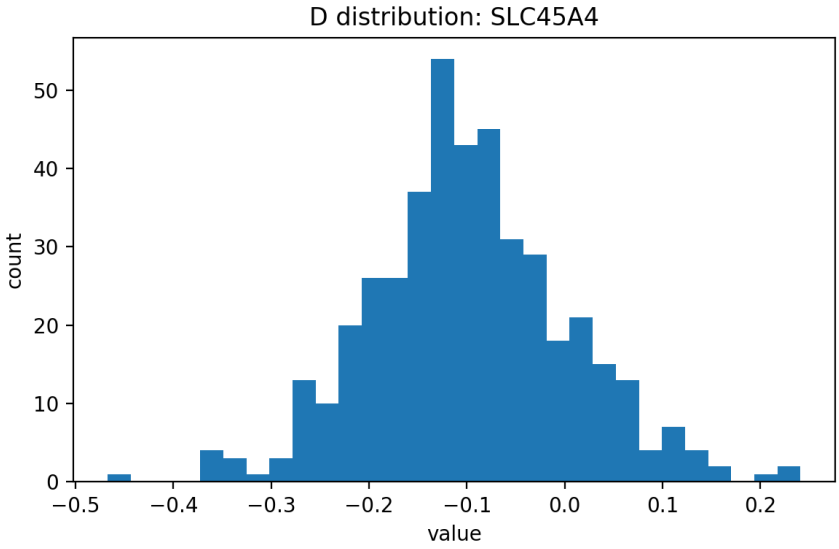


Figure 19: SLC45A4 gene-dependency distribution in the strict-coverage primary sample ($n = 433$).

Table 24: Support of SLC45A4 GeneEffect by major lineages (quantiles). The intersection of the min–max ranges across the listed lineages is non-empty, suggesting common support across lineages.

lineage	n	min	p05	p25	median	p75	p95	max
lung	77	-0.361	-0.244	-0.165	-0.105	-0.039	0.050	0.168
central_nervous_system	34	-0.467	-0.225	-0.131	-0.103	-0.044	0.010	0.033
ovary	32	-0.222	-0.205	-0.145	-0.089	-0.021	0.086	0.200
skin	28	-0.284	-0.216	-0.144	-0.106	-0.070	0.025	0.111
pancreas	28	-0.269	-0.233	-0.164	-0.099	-0.049	0.058	0.146
colorectal	25	-0.343	-0.251	-0.143	-0.116	-0.078	0.099	0.131
upper_aerodigestive	23	-0.264	-0.258	-0.160	-0.106	-0.045	0.078	0.139
breast	22	-0.352	-0.250	-0.189	-0.150	-0.084	0.074	0.091
esophagus	22	-0.308	-0.258	-0.177	-0.136	-0.091	0.027	0.040
urinary_tract	21	-0.273	-0.259	-0.197	-0.103	-0.026	0.064	0.066

Table 25: Support of SLC45A4 GeneEffect by inferred growth-rate quartiles (quantiles). Similar ranges across bins suggest that overlap is not driven by a single growth regime.

growth_bin	n	min	p05	p25	median	p75	p95	max
(0.553, 0.856]	109	-0.467	-0.241	-0.158	-0.092	-0.023	0.094	0.234
(0.856, 0.957]	108	-0.359	-0.247	-0.157	-0.114	-0.031	0.089	0.200
(0.957, 1.077]	108	-0.361	-0.252	-0.143	-0.095	-0.032	0.040	0.241
(1.077, 1.496]	108	-0.353	-0.271	-0.176	-0.109	-0.054	0.071	0.146

Table 26: Propensity-score overlap checks for binary splits of SLC45A4 GeneEffect (median split and q80–q20 split). frac_extreme denotes the fraction of samples with $\hat{e}(X)$ outside $[0.05, 0.95]$; ess is the effective sample size under inverse-propensity weights.

mode	n	q_low	q_high	ps_mean	ps_p05	ps_p95	frac_extreme	ess
median	433	0.500	0.500	0.501	0.309	0.772	0.000	401.689
q80_vs_q20	174	0.200	0.800	0.500	0.277	0.817	0.052	159.750

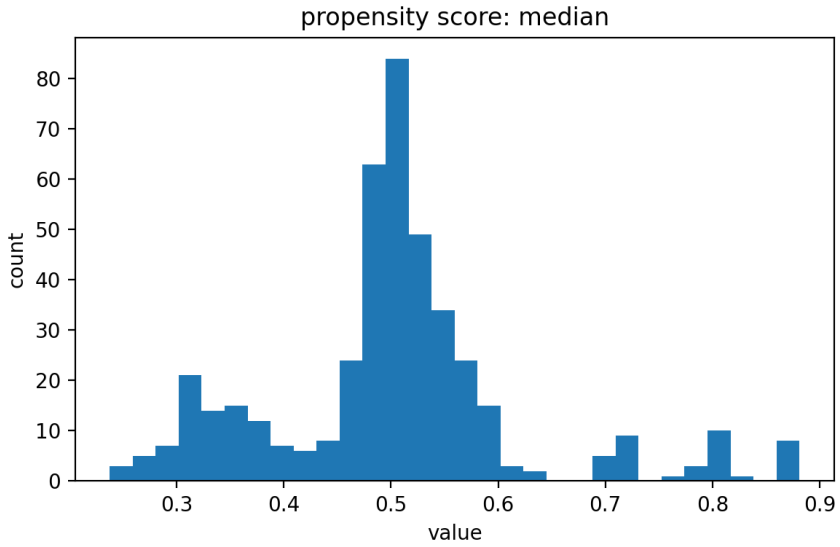


Figure 20: Propensity-score distribution for the median split of SLC45A4 GeneEffect.

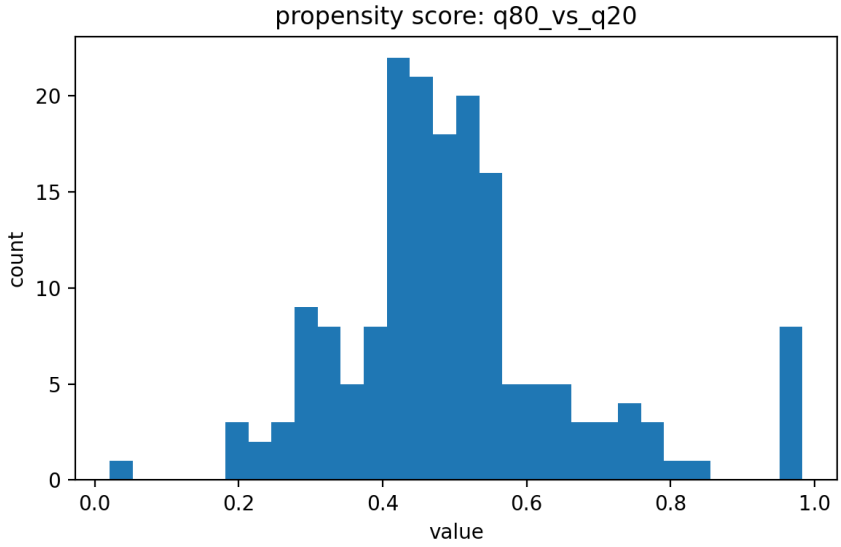


Figure 21: Propensity-score distribution for the q80 vs q20 split of SLC45A4 GeneEffect.

S8 External positive-control checks in AUC datasets

External AUC resources differ in assay protocol, numerical scale, and drug coverage. Accordingly, cross-assay disagreement can reflect assay and coverage differences. As a positive-control check, we test expected pharmacogenomic relationships: EGFR dependency with EGFR-inhibitor sensitivity and BRAF dependency with BRAF-inhibitor sensitivity. Empirical p-values use within-lineage permutation of the residualized focal predictor; with $B = 500$ permutations, the minimum attainable empirical value is $1/(B + 1) \approx 0.002$.

Table 27: Positive-control checks in CTRP and GDSC AUC datasets. Outcomes are drug-set median AUC values standardized within dataset. Focal predictors are target-gene GeneEffect values; p_{perm} is the two-sided within-lineage empirical value.

Dataset	Positive control axis	n	#drugs	$\hat{\theta}_{\text{DML}}$	SE	p	p_{perm}
CTRP	EGFR (rev. EGFR-TKIs)	599	2	1.012	0.200	4.44×10^{-7}	0.0020
GDSC1	EGFR (rev. EGFR-TKIs)	500	2	0.712	0.165	1.60×10^{-5}	0.0020
GDSC2	EGFR (rev. EGFR-TKIs)	444	2	1.221	0.231	1.22×10^{-7}	0.0020
CTRP	EGFR/HER (broad set)	604	4	1.086	0.123	9.20×10^{-19}	0.0020
GDSC1	EGFR/HER (broad set)	590	4	1.004	0.159	2.65×10^{-10}	0.0020
GDSC2	EGFR/HER (broad set)	498	3	0.833	0.135	5.89×10^{-10}	0.0020
CTRP	BRAF inhibitors	600	3	1.139	0.184	6.34×10^{-10}	0.0020
GDSC1	BRAF inhibitors	569	3	1.421	0.270	1.43×10^{-7}	0.0020
GDSC2	BRAF inhibitors	492	2	1.722	0.234	1.87×10^{-13}	0.0020