# Mendelianization: Concentrating Polygenic Signal into a Single Causal Locus

Eric V. Strobl

## Abstract

Complex disorders such as depression and alcohol use involve numerous genetic variants, and implicated loci continue to grow with sample size. This proliferation hampers interpretability, as the mechanisms by which so many variants jointly contribute to pathophysiology remain unclear. In contrast, classical Mendelian diseases arise from a single causal locus and are easier to interpret. We thus introduce *Mendelianization* – an algorithm distinct from Mendelian randomization – that learns weighted combinations of outcomes so that each aggregated phenotype concentrates association at one locus. We prove that this locus is causal under four structural assumptions natural to genetic data. The method handles partial sample overlap, provides calibrated hypothesis tests, maps coefficients to interpretable scales, and quantifies the degree of Mendelianism using summary $z$-statistics alone. In experiments, Mendelianization enhances statistical power to detect Mendelian symptom profiles even in heterogeneous disorders like major depression, generalized anxiety, and alcohol use disorder. An R implementation is available at github.com/ericstrobl/Mendelianization.

**Key words:** causal inference, canonical correlation analysis, GWAS summary statistics, outcome learning

## Introduction

Complex diseases arise from diverse genetic and environmental risk factors. As a result, genome-wide association studies (GWAS) have identified hundreds of associated variants, with larger sample sizes only expanding the catalog.[35] However, the accumulation of associations often complicates rather than clarifies disease mechanism. Even after fine-mapping prioritizes putative causal variants within loci,[27] the genetic architecture remains dispersed across many regions and resists reductionist interpretation.[29,17] A major culprit is phenotypic heterogeneity, where complex disorders like depression manifest as different symptom constellations (e.g., low mood, anhedonia, fatigue, insomnia). Conventional GWAS collapse this diversity into a single diagnosis or composite, broadening the phenotype and inflating the number of hits without sharper insight.[8,35]

Existing multi-trait methods do not reduce this proliferation of associated loci. Structure-learning methods (e.g., Genomic-SEM,[10] PDR,[2] SHAHER,[31] FactorGo[37]) learn latent outcomes that preserve trait-variant covariance, typically yielding even more polygenic, diffuse signals rather than locus-specific effects. By contrast, power-seeking methods (e.g., HIPO,[21] metaCCA,[4] fastASSET[22]) combine traits to maximize a detection criterion such as a non-centrality statistic, identifying still even more variants rather than clarifying mechanism; their nominal $p$-values also tend to be anti-conservative because the outcome is learned on the same data used for testing. Meanwhile, causal inference methods either refine signals within loci (multi-trait fine-mapping[1,38]) or estimate exposure-outcome relations (multi-response MR[40] and DAG methods[39]), rather than trying to identify variant-to-outcome causal effects genome-wide.

We thus advance an alternative approach that leverages multiple outcomes to simplify – not expand – the set of associated variants by learning weighted outcome combinations that approximate "Mendelianized" traits whose genetic associations concentrate in a single causal locus (Figure 1). We specifically make the following **contributions**:

1. We introduce *Mendelianization*, a new framework – distinct from Mendelian randomization – that constructs composite outcomes which concentrate associations at a single locus via a substantially modified canonical correlation analysis (CCA).
2. We prove that this locus contains at least one causal variant under just four structural assumptions.

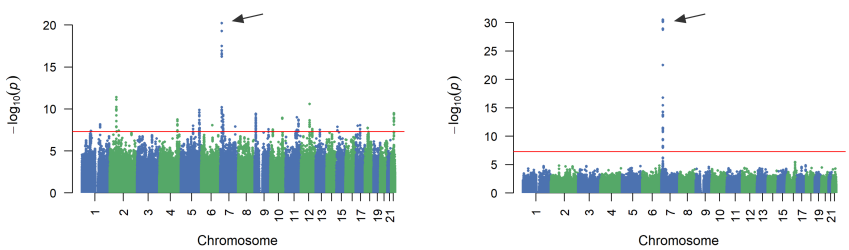


**Fig. 1. Main idea.** Standard outcomes – such as diagnosis or total symptom severity – associate with many loci and obscure underlying mechanisms (left). Mendelianization instead learns an outcome whose association converges on a single locus, yielding a solitary Manhattan plot tower (right). In practice, **multiple** such outcomes are inferred, one for each lead variant. The red line marks genome-wide significance ($5 \times 10^{-8}$).

3. We develop a corresponding fast hypothesis test that remains valid after outcome learning.
4. We operationalize Mendelianization at scale by addressing missing data, preserving coefficient interpretability, ensuring accurate inference, and introducing a continuous *Score of Mendelianism* to quantify achieved Mendelian properties.
5. We demonstrate that Mendelianization uncovers interpretable Mendelian-like traits in complex disorders using symptom-level summary $z$-statistics.[12]

Together, these advances provide a reductionist framework for identifying Mendelian-like symptom profiles dominated by a single causal locus.

# Methods

## One-Dimensional Canonical Correlation Analysis

One-dimensional CCA will serve as the underlying engine of Mendelianization. We thus consider summary statistics comprised of the partial correlation matrix $r \in \mathbb{R}^{q \times m}$ between $q$ variants $V$ and $m$ outcomes $Y$. We assume that all covariates like biological sex, age, and ancestry have already been partialed out from $V$ and $Y$. We focus on population-level properties for ease of presentation and address sample-level issues later. We thus have $r_{jk} = \mathrm{cor}(V_j, Y_k)$. Let $r_j$ correspond to the row of $r$ associated with the variant $V_j$.

For each $V_j$, we construct a composite outcome $Y\alpha_j$ by one-dimensional CCA, so that the correlation between $V_j$ and $Y\alpha_j$ is as large as possible:

$$\arg\max_{\alpha \neq 0} \mathrm{cor}(V_j, Y\alpha). \tag{1}$$

We will also show that this procedure implicitly reduces the correlations between loci not containing $V_j$ and $Y\alpha_j$, so that $Y\alpha_j$ resembles a Mendelian trait.

Without loss of generality, we assume that each variant and outcome has already been standardized to mean zero and unit variance. We can then simplify the above objective to:

$$\arg\max_{\alpha \neq 0} \frac{r_j \alpha}{\sqrt{\alpha^\top \Sigma \alpha}}, \tag{2}$$

where $\Sigma$ corresponds to the positive definite correlation matrix of $Y$. Notice that the numerator maximizes alignment with $V_j$, while the denominator minimizes the scale of $\alpha$ according to the geometry of $\Sigma$.

The solution to Expression (2) is the familiar CCA weight vector that whitens the outcomes and then aligns with $r_j$. In particular, up to an arbitrary scale factor, the optimal weights satisfy

$$\alpha \propto \Sigma^{-1} r_j^\top,$$

so we set $\alpha_j = A r_j^\top$ with $A = \Sigma^{-1}$ from here on. We provide a detailed derivation of this solution in Supplementary Materials A.

## Causal Analysis

We next ask when the CCA-derived composite $Y\alpha_j$ admits a causal interpretation. In particular, we show that, under a set of structural assumptions motivated by genetic data, the optimization problem in Expression (1) produces a composite outcome whose association pattern is concentrated within a single linkage disequilibrium (LD) block, and that this block must contain at least one causal variant for $Y\alpha_j$. We refer to this property as *Mendelianization*.

### *Structure of Latent Confounding*

Even after adjusting for observed covariates such as age, sex, and ancestry, genetic association studies typically retain residual confounding from processes such as batch effects, residual population structure, and dynastic effects.[35] These processes tend to act on many phenotypes at once and therefore induce shared structure across outcomes.

We formalize this by assuming:

**Assumption 1.** *(Low-dimensional latent confounding) For sufficiently large $m$, the remaining confounding effects lie in a subspace $U^m \subset \mathbb{R}^m$, and there exists an integer constant $r_0 \geq 0$ such that $\dim(U^m) \leq r_0$.*

**Remark.** Assumption 1 states that $V$ and $Y$ are influenced by only a small number of latent confounders, or by many confounders whose effects span a low-dimensional subspace. This matches empirical experience in large biobanks, where a handful of unmeasured processes generate broad, shared patterns across many traits.

Assumption 1 implies that each vector of variant-outcome correlations $r_j^m$ can be written as

$$r_j^m = r_j^{m\parallel} + r_j^{m\perp},$$

where $r_j^{m\parallel}$ lies in the confounding subspace $U^m$ and $r_j^{m\perp}$ lies in its orthogonal complement (with respect to $A^m = (\Sigma^m)^{-1}$). Informally, $r_j^{m\parallel}$ captures residual confounding, while $r_j^{m\perp}$ captures the causal component of the association between $V_j$ and the outcomes. We use this decomposition to study how Mendelianization behaves with and without latent confounding.

### *Without Latent Confounding*

We first consider the idealized case where latent confounding is absent, so $r_j^m = r_j^{m\perp}$ for all $j$. Let $r^m \alpha_j^m$ denote the vector of associations between each variant and the composite outcome defined by $\alpha_j^m = A^m r_j^{m\top}$. For each variant $V_k$,

$$\frac{1}{m} r_k^m \alpha_j^m = \frac{1}{m} r_k^m A^m r_j^{m\top} = \frac{1}{m} \langle r_k^m, r_j^m \rangle_{A^m},$$

where $\langle \cdot, \cdot \rangle_{A^m}$ denotes the inner product induced by $A^m$. When $k = j$, this quantity measures how strongly $V_j$ associates with its own learned outcome. When $k \neq j$, the same expression measures the average alignment between the causal pattern of $V_j$ and that of $V_k$ across the $m$ outcomes.

To obtain a single-locus "tower" in the Manhattan plot of $Y\alpha_j^m$, we require that this average alignment vanish for variants outside the LD block around $V_j$. Let $\mathcal{S}_j$ denote the set of variants in LD with $V_j$ (the locus or LD block containing $V_j$). We assume:

**Assumption 2.** *(Outcome diversity) We have*

$$\max_{k \notin \mathcal{S}_j} \frac{1}{m} \left| \langle r_k^{m\perp}, r_j^{m\perp} \rangle_{A^m} \right| \in o_m(1).$$

**Remark.** In the absence of confounding, $r_j^{m\perp}$ represents the causal pattern of effects of $V_j$ on the outcomes. Assumption 2 says that, as we add more non-redundant outcomes and whiten them, the causal patterns from different loci become increasingly distinct on average. Mixed positive and negative effects from variants outside $\mathcal{S}_j$ tend to cancel out when averaged across many outcomes, driving their

contribution to $Y\alpha_j^m$ toward zero. By contrast, variants in $\mathcal{S}_j$ share the same underlying causal mechanism for $Y\alpha_j^m$ and therefore retain strong alignment. This is plausible in high-dimensional phenotyping because adding more granular traits (e.g., subphenotypes or symptom items) typically increases the heterogeneity of causal effects across loci. In sum, the vector $r^m\alpha_j^m$ develops a sharp tower within $\mathcal{S}_j$ as $m$ grows, while associations at other loci approach zero.

### *With Latent Confounding*

In real data, latent confounding is present and can also create tall peaks in Manhattan plots. For example, the composite $Y\alpha_j^m$ might load on a pattern driven purely by confounding, such that

$$\frac{1}{m} r_j^m \alpha_j^m = \frac{1}{m} \langle r_j^{m\parallel}, r_j^{m\parallel} \rangle_{A^m}$$

even if no variant in $\mathcal{S}_j$ is truly causal for $Y\alpha_j^m$. To rule out such purely confounded towers, we impose additional assumptions on the strength of both causal and confounded components.

We first assume that the relevant signals do not vanish as we add more outcomes:

**Assumption 3.** *(Signal floors) For sufficiently large $m$, if $l \in \mathcal{S}_j$ causally influences $Y\alpha_j^m$, then there exists a constant $\beta_{\mathrm{c}} > 0$ such that*

$$\frac{1}{m} \|r_l^{m\perp}\|_{A^m}^2 \geq \beta_{\mathrm{c}}.$$

*Moreover, there exists a constant $\beta_{\mathrm{on}} > 0$ such that*

$$\frac{1}{m} \|r_l^{m\parallel}\|_{A^m}^2 \geq \beta_{\mathrm{on}}$$

*for every $l \in \mathcal{S}_j$ that does not cause $Y\alpha_j^m$.*

**Remark.** The first part of Assumption 3 requires that if $\mathcal{S}_j$ contains a causal variant for $Y\alpha_j$, its causal signal remains non-negligible even as we increase the number of outcomes. This is a minimal requirement for any causal effect to be detectable. The second part states that non-causal variants in the same LD block share a non-vanishing confounded component, because LD makes them inherit similar exposure to the same local confounders (such as local ancestry or technical artifacts). This is consistent with how confounding typically behaves within LD regions in GWAS.

We also formalize how low-dimensional confounding propagates across the genome:

**Assumption 4.** *(Isolation-spillover) If some $l \in \mathcal{S}_j$ causally influences $Y\alpha_j^m$, then*

$$\max_{k \notin \mathcal{S}_j} \frac{1}{m} \left| \langle r_k^{m\parallel}, r_j^{m\parallel} \rangle_{A^m} \right| \to 0.$$

*Conversely, if no $l \in \mathcal{S}_j$ causes $Y\alpha_j^m$, then there exists a variant $k \notin \mathcal{S}_j$ such that*

$$|\cos_{A^m}(r_j^{m\parallel}, r_k^{m\parallel})| \geq \gamma_j > 0 \quad \textit{and} \quad \frac{1}{m} \|r_k^{m\parallel}\|_{A^m}^2 \geq \beta_{\mathrm{off}} > 0$$

*for sufficiently large $m$.*

**Remark.** The first clause states that when $\mathcal{S}_j$ contains a causal variant for $Y\alpha_j$, the associated confounding pattern does not align strongly with distant loci, so any remaining confounding is effectively local to the LD block. The second clause describes the opposite scenario: if the tower at $\mathcal{S}_j$ is driven *purely* by low-dimensional confounding, then at least one other locus must share the same confounding pattern with non-negligible strength, leading to detectable "spillover" of confounded signal elsewhere in the genome. This matches empirical experience, where broad confounding (for example, uncorrected population structure) usually produces widespread inflation rather than a single isolated peak.

### *Asymptotic and Finite Mendelianism*

Under Assumptions 1–4, the CCA estimator enjoys a strong causal localization property:

**Theorem 1.** *(Asymptotic Mendelianism) Consider Assumptions 1–4 and suppose $\Sigma^m$ is positive definite for sufficiently large $m$. We have $\frac{1}{m} r_j^m \alpha_j^m > 0$ for sufficiently large $m$ and $\max_{k \notin \mathcal{S}_j} |\frac{1}{m} r_k^m \alpha_j^m| \to 0$ (i.e., Mendelianism is asymptotically achieved for $Y\alpha_j^m$ at $\mathcal{S}_j$) if and only if at least one variant $l \in \mathcal{S}_j$ causes $Y\alpha_j^m$.*

**Remark.** We provide all proofs in the Supplementary Materials. Theorem 1 establishes both directions: the composite outcome $Y\alpha_j^m$ becomes asymptotically Mendelian at the LD region $\mathcal{S}_j$ if and only if this region contains at least one causal variant for $Y\alpha_j^m$. Intuitively, under the four structural assumptions (low-dimensional latent confounding, outcome diversity, signal floors, and isolation-spillover), maximizing the correlation between $V_j$ and $Y\alpha$ forces the resulting composite outcome to localize its causal signal to a single LD block. This contrasts with common practice, which fixes a single outcome or factor and tests all variants, often yielding many significant loci scattered across the genome (Figure 1, left). By comparison, the estimator $\alpha_j^m = A^m r_j^{m\top}$ learns a variant-specific composite outcome for each $V_j$, and Theorem 1 shows that all causal variants for this learned outcome $Y\alpha_j^m$ concentrate in one LD region (Figure 1, right).

Beyond this localization guarantee, Theorem 1 also changes how we interpret CCA loadings. For each variant $V_j$, we estimate the canonical coefficients $\alpha_j^m$ and form the composite $Y\alpha_j^m$. **We then hold this composite fixed and evaluate association genome-wide** – not just at $V_j$ itself. Each $\alpha_j^m$ thus defines a Mendelianized trait that we scan across all variants. We interpret $Y\alpha_j^m$ as a Mendelian trait when the corresponding Manhattan plot shows a single tower confined to an LD region containing $V_j$.

Finally, we recover a stronger version when a finite outcome set is already sufficiently rich:

**Proposition 1.** *(Perfect Mendelianism at finite $m$) Consider Assumptions 5–8 (Supplementary Materials A) with a prespecified $m$ and suppose $\Sigma^m$ is positive definite. We have $\frac{1}{m} r_j^m \alpha_j^m > 0$ and $\max_{k \notin \mathcal{S}_j} |\frac{1}{m} r_k^m \alpha_j^m| = 0$ (i.e., Mendelianism is perfectly achieved for $\mathcal{S}_j$ at $m$) if and only if at least one variant $l \in \mathcal{S}_j$ causes $Y\alpha_j^m$.*

**Remark.** Proposition 1 shows that, under stronger assumptions, Mendelianization can in principle recover a classical single-locus architecture for the learned outcome at a fixed number of outcomes $m$. We emphasize Theorem 1 because it uses weaker, more realistic assumptions and allows Mendelianism to emerge asymptotically as we enrich the outcome set. This distinction is important for complex,

polygenic diseases, where perfect single-locus resolution may not be achievable at any fixed $m$.

## Fast Hypothesis Testing

We now test the null of no dependence between $V_j$ and $Y$, even though we learn $\alpha_j$ from the same data. We henceforth drop the superscript $m$ to simplify notation. The null and alternative hypotheses are:

$$\begin{aligned} \mathcal{H}_0 &: V_j \perp\!\!\!\perp Y, \\ \mathcal{H}_1 &: r_j \alpha_j > 0. \end{aligned} \tag{3}$$

A naive Pearson test would be anti-conservative, because it treats $\alpha_j$ as fixed even though we estimated it using the observed associations. To obtain a valid test, we construct a statistic whose null distribution accounts for the fact that $\alpha_j$ is data-driven.

Under $\mathcal{H}_0$, with i.i.d. samples and finite fourth moments,

$$\sqrt{n}\,\widehat{r}_j \coloneqq z_j \xrightarrow{d} \mathcal{N}(0, \Sigma), \tag{4}$$

where $n$ denotes sample size[1]. The canonical coefficients are proportional to $Az_j^\top$, which motivates the quadratic form $Q_j = z_j A z_j^\top$.

To understand the null distribution of $Q_j$, write the Cholesky decomposition $\Sigma = LL^\top$ and let $u \sim \mathcal{N}(0, I_m)$ so that $z_j = u^\top L^\top$. Substituting into $Q_j$ gives

$$Q_j = (Lu)^\top A(Lu) = u^\top (L^\top AL)u = u^\top u,$$

so that $Q_j \sim \chi^2_m$ under $\mathcal{H}_0$. We then compute the right-tail probability of $Q_j$ under the $\chi^2_m$ distribution to obtain the $p$-value.

The above $\chi^2$ result is asymptotic, and its finite-sample accuracy depends on how well we approximate $\Sigma$. In the next section, we show that a modified version of $\Sigma$ can be consistently estimated using millions of variants under realistic LD and sparsity conditions, so that large $n$ and large $q$ asymptotics work decisively in our favor.

## Operationalizing Mendelianization at Scale

Implementing the above hypothesis test requires access to $\Sigma$. However, biobanks rarely report $\Sigma$ in their summary statistics, and the outcomes needed to estimate $\Sigma$ are typically measured on partially overlapping sets of individuals. Scaling Mendelianization also raises two additional challenges beyond test calibration: redundant traits can obscure the interpretation of raw canonical coefficients, and the available traits may not be rich enough to approach near-perfect Mendelianism. We therefore (a) replace $\Sigma$ with a quantity that is estimable from summary $z$-statistics under partial overlap, (b) show how to estimate it at scale despite LD, (c) rescale coefficients for interpretability, and (d) introduce a metric to quantify how close a locus comes to Mendelian behavior.

### *From $\Sigma$ to $\Gamma$ Under Partial Overlap*

If all traits are fully observed and $V \perp\!\!\!\perp Y$, then $\mathbb{E}[\,z_j z_j^\top\,] = \Sigma$ with i.i.d. sampling, and we can consistently estimate $\Sigma$ using (approximately) independent variants. In practice, samples overlap only partially across traits, so $z$-statistics and correlations are computed from unequal sample sizes.

Let $M_{a,i} \in \{0,1\}$ indicate whether $Y_{a,i}$ is observed, and define $S_a = \{i : M_{a,i} = 1\}$, $n_a = |S_a|$, and $n_{ab} = |S_a \cap S_b|$. We assume that $V$ contains no missing values, and that each $Y_a$ has been centered so that $E(Y_a|M) = 0$ for the realized missingness pattern $M$. Let

$$\Sigma^{\text{obs}}_{ab} \coloneqq \mathbb{E}(Y_{a,i}Y_{b,i}\,|M_a = 1, M_b = 1, M_{-a,-b})\,.$$

denote the second-moment matrix conditional on both $Y_a$ and $Y_b$ being observed (i.e., over $S_a \cap S_b$). Then:

**Proposition 2.** *Assume i.i.d. samples across individuals, $V \perp\!\!\!\perp Y$, and $V \perp\!\!\!\perp M \mid Y$. Then*

$$\mathbb{E}\big[z_{ja}z_{jb} \mid M\big] = \frac{n_{ab}}{\sqrt{n_a n_b}}\,\Sigma^{\text{obs}}_{ab}\,.$$

**Remark.** The assumption $V \perp\!\!\!\perp M \mid Y$ is reasonable in biobank settings since missingness in $Y$ is largely driven by the presence or severity of symptoms themselves (as well as factors such as sex, age, and ancestry, which we already adjusted for). Proposition 2 also motivates replacing $\Sigma$ with the matrix $\Gamma$ in Expression (2), where

$$\Gamma_{ab} \coloneqq \frac{n_{ab}}{\sqrt{n_a n_b}}\Sigma^{\text{obs}}_{ab}\,.$$

Replacing $\Sigma$ with $\Gamma$ has a convenient interpretation. Define modified outcomes

$$Y'_a = \sqrt{\tfrac{n}{n_a}}\,M_a \odot Y_a,$$

which set missing entries of $Y_a$ to zero and rescale by the square root of the observed sample proportion. Then $E(Y'|M) = 0$ for the realized $M$, and hence $E(Y'_a) = 0$. Moreover:

**Proposition 3.** *Assume i.i.d. samples across individuals, $p_a = \mathbb{P}(M_a = 1) > 0$, and $p_b = \mathbb{P}(M_b = 1) > 0$. Then*

$$\mathbb{E}\left(\frac{1}{n}\sum_t Y'_{a,t}Y'_{b,t} \mid M\right) = \Gamma_{ab} \xrightarrow{a.s.} \frac{p_{ab}}{\sqrt{p_a p_b}}\Sigma^{\text{obs}}_{ab}$$

*and*

$$\mathbb{E}\left(\frac{1}{n}\sum_t Y'_{a,t}Y'_{b,t}\right) \to \frac{p_{ab}}{\sqrt{p_a p_b}}\Sigma^{\text{obs}}_{ab}$$

*as $n \to \infty$.*

**Remark.** $\Gamma_{ab}$ is therefore the *conditional* covariance of $Y'$ for the realized missingness pattern $M$, but it converges to the *unconditional* covariance of $Y'$ as sample size grows, so dependence on the specific pattern of missingness fades. As a result, we can maximize the (asymptotically) unconditional correlation in Expression (2) by working with $Y'$ instead of $Y$:

$$\max_{\alpha \neq 0} \frac{\widehat{r}'_j\,\alpha}{\sqrt{\alpha^\top \Gamma \alpha}}, \qquad \text{where } \widehat{r}'_j = \widehat{\text{cov}}(V_j, Y').$$

Importantly, this reparameterization targets $\Gamma$ directly and therefore **does not require missingness-at-random (MAR)**[14] for $\Sigma$.

[1] All arguments can be modified by substituting sample sizes with Fisher-effective sample sizes in generalized linear models with $o_p(1)$ remainders.

Theorem 1 then holds for $Y'$ at the population level with $\Sigma$ replaced by $\Gamma$. The scaling factors $\sqrt{n/n_a}$ used to define $Y'$ do not affect the optimization, since:

$$Y'\alpha = (Y \odot M)D^{1/2}\alpha = (Y \odot M)\gamma,$$

where $D = \mathrm{diag}(n/n_1, \ldots, n/n_m)$ and $\gamma = D^{1/2}\alpha$. Hence:

$$\max_{\alpha \neq 0} \widehat{\mathrm{cor}}(V_j, Y'\alpha) = \max_{\gamma \neq 0} \widehat{\mathrm{cor}}(V_j, (Y \odot M)\gamma).$$

Even better, we never need explicit access to $Y'$ because the sample covariance simplifies to:

$$\widehat{r}'_{ja} = \frac{1}{n}\sum_{s=1}^{n} V_{j,s} Y'_{a,s} = \frac{1}{n}\sqrt{\frac{n}{n_a}} \sum_{s: M_{a,s}=1} V_{j,s} Y_{a,s} = \frac{z_{ja}}{\sqrt{n}},$$

which implies:

$$\widehat{\mathrm{cor}}(V_j, Y'\alpha) = \frac{1}{\sqrt{n}} \frac{z_j \alpha}{\sqrt{\alpha^\top \Gamma \alpha}},$$

except that we replace $\frac{1}{n}\sum_t Y'_{a,t} Y'_{b,t}$ in the denominator with its conditional counterpart $\Gamma$, which has the same large-sample limit by Proposition 3. Thus, we only need the $z$-statistics $z_j$ and an estimate of $\Gamma$.

This reasoning motivates us to maximize:

$$\arg\max_{\alpha \neq 0} \frac{z_j \alpha}{\sqrt{\alpha^\top \Gamma \alpha}}, \tag{5}$$

which mirrors Expression (2) with $\Sigma$ replaced by $\Gamma$. We henceforth take $\alpha_j = \Omega z_j^\top$ as our solution, where $\Omega = \Gamma^{-1}$.

### *Estimating $\Gamma$ with LD*

Although $\Gamma$ is natural to target, empirical estimation is nontrivial. In principle, Proposition 2 shows that genome-wide $z$-statistics identify $\Gamma$ under the global null $V \perp\!\!\!\perp Y$. In practice, two main obstacles remain: LD induces dependence between nearby variants, and a subset of variants truly associate with $Y$. Fortunately, we can handle both these issues with the following result:

**Theorem 2.** *Consider the estimator: $\widehat{S}_{ab} = \frac{1}{q}\sum_{j=1}^{q} z_{ja} z_{jb}$. Let $G_q$ index the variants that are independent of $Y$, and $B_q$ index the variants that are dependent on $Y$. Let $C_1, C_1', C_2, L$ and $D$ denote finite constants. Assume:*

1. *(Bounded moments for $G_q$) For $j \in G_q$, we have $\mathbb{E}(z_{ja} z_{jb} | M) = \Gamma_{ab}$ and $\sup_{j \in G_q} \mathrm{Var}(z_{ja} z_{jb} | M) \leq C_1$.*
2. *(Bounded moments for $B_q$) For $j \in B_q$, we have $\mathbb{E}(z_{ja} z_{jb} | M) = \mu_j$ with $\sup_{j \in B_q} |\mu_j - \Gamma_{ab}| \leq L$ and $\sup_{j \in B_q} \mathrm{Var}(z_{ja} z_{jb} | M) \leq C_1'$.*
3. *(Local LD) There is an undirected dependency graph with edge set $E_q$ and maximum degree $D$ such that $|\mathrm{cov}(z_{ja} z_{jb}, z_{ka} z_{kb} | M)| \leq C_2$ when $(j,k) \in E_q$ and $\mathrm{cov}(z_{ja} z_{jb}, z_{ka} z_{kb} | M) = 0$ when $(j,k) \notin E_q$.*
4. *(Vanishing contamination) We have $\frac{|B_q|}{q} \to 0$ as $q \to \infty$.*

*Then, $\widehat{S}_{ab} \xrightarrow{p} \Gamma_{ab}$ conditional on $M$ as $q \to \infty$.*

**Remark.** The key condition is the local LD assumption. In practice, investigators routinely treat LD as local by bounding each variant's LD neighborhood (for example, within 1 Mb), which makes the dependency degree $D$ finite. Under this condition, we can show that the variance of $\widehat{S}_{ab}$ goes to zero, so $\widehat{S}_{ab} \xrightarrow{p} \Gamma_{ab}$ given the realized missingness pattern. Moreover, the vanishing contamination assumption reflects the sparse genetic architecture of most traits: after standard quality control, only a tiny fraction of variants are typically associated with any given outcome, so the non-null set $B_q$ grows more slowly than $q$. This matches empirical experience in large biobanks, where the vast majority of variants behave as null.

Taken together, these conditions imply a weak law of large numbers despite LD. In practice, with millions of variants, we can estimate each entry of $\Gamma$ to negligible error, so we treat $\Gamma$ as known in subsequent derivations. Consequently, under $\mathcal{H}_0$, $z_j \xrightarrow{d} \mathcal{N}(0, \Gamma)$, and we replace $\Sigma$ with $\Gamma$ and $A$ with $\Omega = \Gamma^{-1}$ so that $Q_j = z_j \, \Omega \, z_j^\top \sim \chi^2_m$ in the test of Expression (3). We will later show that this test works very well in real data.

### *Interpretable Coefficients*

We can now solve Expression (5) using an estimate of $\Gamma$. However, the raw coefficients correspond to:

$$\alpha_{ja} = \sqrt{n}\, \widehat{\mathrm{cov}}(V_j, R'_a)/\mathrm{var}(R'_a),$$

where $R'_a$ is the residual of $Y'_a$ after regressing it onto $Y'_{-a}$, so they depend on the residual variance $\mathrm{var}(R'_a)$, which varies across outcomes and complicates comparisons.

To address this, we report rescaled coefficients $\widehat{\alpha}'_{ja} := \widehat{\alpha}_{ja}/\sqrt{\Omega_{aa}}$ for interpretability:

**Proposition 4.** *We have $\widehat{\alpha}'_{ja} = \sqrt{n}\ \widehat{\mathrm{cor}}(V_j, R'_a)$ for every $Y_a \in Y$.*

Thus, $\widehat{\alpha}'_{ja}$ equals $\widehat{\mathrm{cor}}(V_j, R'_a)$ up to the common factor $\sqrt{n}$. The rescaled canonical coefficients are unitless, directly comparable across outcomes for a given variant $V_j$, and their signs match the allelic direction of effect on the unique component of $Y'_a$ (the residual $R'_a$ after adjusting for the other traits). In practice, $\widehat{\alpha}'_{ja}$ can be viewed as a standardized effect of $V_j$ on the part of trait $a$ that is not explained by the remaining traits in $Y$.

### *Quantifying the Degree of Mendelianism*

Near-perfect Mendelianism rarely occurs when the outcome set is small (for example, $m \leq 15$). Although Theorem 1 provides a yes/no criterion in the limit, applied analyses benefit from a continuous measure that captures how close $Y\alpha_j$ comes to Mendelian behavior.

We therefore fix the coefficients $\alpha_j$ learned at the lead variant $V_j$ and, for any variant $V_k$ (including $k = j$), compute

$$\mathcal{Z}_{k|j} = \frac{z_k \, \alpha_j}{\sqrt{\alpha_j^\top \Gamma \, \alpha_j}}.$$

Under $\mathcal{H}_{0,k} : r_k \alpha_j = 0$, we have $\mathcal{Z}_{k|j} \xrightarrow{d} \mathcal{N}(0,1)$ with i.i.d. samples and finite fourth moments, yielding the two-sided $p$-value $p_{k|j} = 2\Phi(-|\mathcal{Z}_{k|j}|)$[2]. We compute $\mathcal{Z}_{k|j}$ and $p_{k|j}$ for all $V_k \in V$ (i.e., genome-wide).

If we reject $\mathcal{H}_0$ in Expression (3) for $V_j$, we summarize how concentrated the signal is at the $j$-locus with the **Score of**

[2] The p-values are not always calibrated in practice, because the test assumes that $\alpha_j$ is fixed. We do not use this test for formal inference, but only for visualizing and summarizing the degree of Mendelianism.

**Mendelianism (SoM)**:

$$\mathcal{M}_j = \frac{\frac{1}{|\mathcal{S}_j|}\sum_{i\in\mathcal{S}_j} h(\mathcal{Z}_{i|j})}{\frac{1}{|\mathcal{S}_j|}\sum_{i\in\mathcal{S}_j} h(\mathcal{Z}_{i|j}) \;+\; \sum_{k\notin\mathcal{S}_j} h(\mathcal{Z}_{k|j})}, \qquad (6)$$

where $h(\mathcal{Z}) = \big(|\mathcal{Z}| - z_\alpha\big)_+$, and $z_\alpha$ is the $z$-score corresponding to a two-sided $p$-value (default $\alpha = 5 \times 10^{-8}$). We construct $\mathcal{S}_j$ using proximity-based clustering[23] around $V_j$ (single-linkage with 500 kb gap, padded by $\pm$500 kb), though alternatives such as LD-based clustering relative to $V_j$ are also possible.

**Remark.** The SoM lies in $[0, 1]$ and can be interpreted as the fraction of "excess" Mendelianized signal (above the genome-wide threshold $z_\alpha$) that lies within the LD region of interest. A value of $\mathcal{M}_j = 1$ indicates complete isolation of the signal to a single locus, while $\mathcal{M}_j$ approaches 0 when the signal lies entirely outside $\mathcal{S}_j$ or is highly polygenic. We average within $\mathcal{S}_j$ to avoid favoring larger loci and *do not* average outside $\mathcal{S}_j$ so that even small off-target regions remain visible rather than being diluted across millions of variants. In practice, we use $\mathcal{M}_j$ only as a ranking metric among loci that are already genome-wide significant; we do not use it for hypothesis testing or confidence intervals.

### *Summary*

Mendelianization corresponds to one-dimensional CCA, modified in five ways: (1) it uses $z$-statistics in place of raw correlation coefficients, (2) it overcomes missingness by normalizing with $\Gamma$ estimated over variants rather than with $\Sigma$ estimated over outcomes, (3) it replaces Wilks' $\Lambda$ with a $\chi^2_m$ test for significance, (4) it rescales canonical coefficients by the square roots of the diagonals of $\Omega$ for interpretability, and (5) it applies each significant composite outcome genome-wide by Theorem 1. Pseudocode and a formal time complexity analysis (Supplementary Materials B) show that Mendelianization scales cubically with the number of outcome variables but only linearly with the number of variants. This remains computationally practical in modern biobanks, where the number of outcomes rarely exceeds 100, whereas the number of variants often surpasses one million.

## Results

### Comparators

We compared Mendelianization against:

- **Meta-analytic Canonical Correlation Analysis (metaCCA)**:[4] performs CCA on summary statistics using Wilks' $\Lambda$, which assumes perfectly overlapping and equal sample sizes per trait. metaCCA also estimates $\Sigma$ from linear regression $\beta$ coefficients, so it cannot fully combine linear and logistic regression models whose coefficients differ in scale and null variances. Finally, metaCCA lacks causal guarantees, and its scale-dependent canonical coefficients are not directly comparable across traits.
- **Heritability Informed Power Optimization (HIPO)**:[21] learns multiple sets of outcome weights that increase statistical power by maximizing a standardized non-centrality parameter across – rather than within – loci. We compare against the most powerful, first set of HIPO weights.
- **Fast Association analysis based on SubSETs (fastASSET)**:[22] a multi-trait analysis method that first pre-screens traits with suggestive associations and adjusts for this selection. It then searches over all (two-sided) subsets of traits – allowing opposite effect directions – and combines evidence into a single, multiple testing adjusted omnibus $p$-value per variant.

Mendelianization is the only method that optimizes composite outcomes so that each learned outcome is associated with a single tower in the Manhattan plot that contains at least one causal variant. We finally also compared Mendelianization against **three ablated variants**: (i) using variant-specific outcomes (i.e., $\mathcal{Z}_{i|j} = \mathcal{Z}_{i|i}$) as opposed to fixed outcomes applied genome-wide, (ii) estimating $\Gamma$ using $\beta$ coefficients instead of $z$-statistics, (iii) a combination of both (i) and (ii).

### Metrics

Our primary evaluation metric was the **Score of Mendelianism (SoM)**, where higher values indicate a greater ability to isolate a single locus. The metric provides a fair basis for comparison because no algorithm, including Mendelianization, optimizes SoM. Given that Mendelianization and metaCCA have different coefficient interpretations, we computed SoM with variant-specific outcomes for metaCCA, and with fixed outcomes applied genome-wide for Mendelianization (the outcome learned at the lead $V_j$, applied to all variants). This ensures each method is evaluated in line with its intended use.

Beyond detecting a Mendelian-like pattern, rigorous applications also require valid, efficient inference. As a result, we evaluated Type I error control using the **signed Kolmogorov–Smirnov (KS) distance** between observed $p$-values above 0.05 and the uniform distribution on $[0.05, 1]$. Well-calibrated methods should yield nearly uniform $p$-values in this range, since most variants are uncorrelated with the target. For methods with positive (anti-conservative) KS-distances below 0.01, we then assessed the **statistical power** of the hypothesis tests for $p$-values below $10^{-4}$, since most variants with $p$-values below this threshold are likely to be associated with their learned outcome. Finally, we measured **run-time**.

### Simulations

We first conducted simulations using individual-level autosomal genotype data from European-ancestry participants in the 1000 Genomes Phase 3 dataset.[5] After regressing out the first five principal components, we partitioned the genome into 1,006 LD blocks using the optimal LD-splitting algorithm[20] (parameters: $r^2$ threshold for ignoring pairwise correlations = 0.02, window size = 500 Kb, maximum number of variants per block = 2,000, minimum number of variants per block = 200). We then randomly sampled 100,000 variants and, for each of 40 bootstrap replicates, randomly selected five causal variants $V_c$ from this set. For each replicate, we first generated 50 latent traits according to

$$\widetilde{Y} = V_c\zeta,$$

and then constructed 32 observed outcomes,

$$Y = \widetilde{Y}\tau + \varepsilon_Y,$$

where $\widetilde{Y}$ contained 50 variables, each element of $\zeta$ was drawn from a uniform distribution on $[-0.05, -0.01] \cup [0.01, 0.05]$, each entry of $\tau$ was drawn from a standard normal distribution, and $\varepsilon_Y$ also followed a standard normal distribution. This construction mimics realistic settings in which observed outcomes are noisy, imperfect measurements that do not map cleanly onto the underlying biology. For each outcome, we randomly drew the sample size from a uniform distribution between 50,000 and 100,000 individuals by

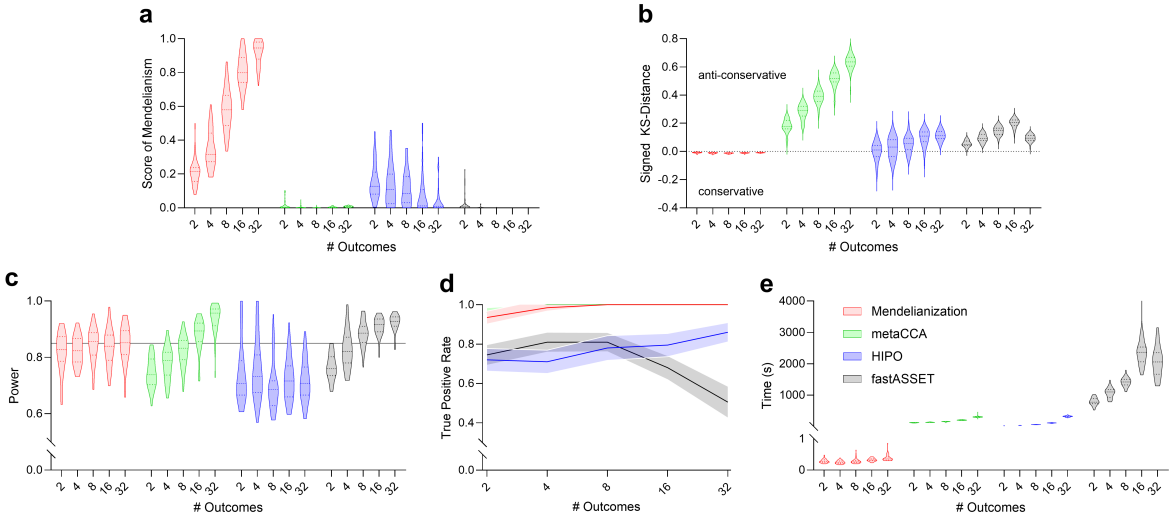


**Fig. 2. Main results from the synthetic data experiments.** All plots share the same legend. (a) Mendelianization achieved the highest SoM across all numbers of outcomes, with SoM increasing as the number of outcomes grew, consistent with Theorem 1. (b) Mendelianization was the only algorithm that produced well-calibrated $p$-values, in agreement with Theorem 2. (c,d) Mendelianization maintained statistical power of approximately 0.85 and a near-perfect true positive rate across all outcome settings. Panel (d) is displayed as a line plot because the mean true positive rate assumes only a discrete set of values. Error bars denote 95% confidence intervals for the mean. (e) Mendelianization completed in under one second for all runs.

bootstrap resampling individuals from 1000 Genomes. Within each LD block, we then performed a second-level bootstrap to disrupt cross-block correlations and to approximate GWAS-scale sampling variability given the relatively small 1000 Genomes sample size. We applied all methods to the resulting 40 simulated datasets for each outcome configuration (2, 4, 8, 16, and 32 outcomes).

We summarize the simulation results in Figure 2, using primarily violin plots. Mendelianization achieved the highest SoM across all numbers of outcomes (Figure 2 (a)). Consistent with Theorem 1, the SoM steadily increased toward one as the number of outcome variables increased, indicating that the algorithm identified increasingly Mendelian-like traits with more outcomes. In contrast, the other methods produced SoMs near zero, reflecting the fact that we generated the data from five randomly chosen causal variants rather than a single causal variant.

Mendelianization also yielded the most well-calibrated $p$-values for null variants, with calibration close to ideal, consistent with Theorem 2 (Figure 2 (b)). By comparison, most alternative algorithms showed progressively poorer calibration as the number of outcomes increased, indicating a failure to control the type I error rate under outcome learning. Mendelianization was further distinguished as the only method that consistently retained statistical power of approximately 0.85 across all outcome dimensions (Figure 2 (c)), whereas the power of competing approaches fluctuated substantially, in line with their unstable $p$-value calibration.

In terms of discovery performance, Mendelianization almost always recovered all five causal variants, and its true positive rate increased with the number of outcomes (Figure 2 (d)), demonstrating robust power even in high-dimensional outcome settings. Finally, Mendelianization completed in under one second for all runs, whereas other algorithms required approximately three- to four-fold longer computation times (Figure 2 (e)). Overall, these results indicate that Mendelianization identifies more Mendelian-like traits as the number of outcomes grows, maintains well-calibrated $p$-values with high power, and remains computationally efficient. Finally, ablation results confirmed the necessity of applying canonical coefficients genome-wide and using $z$-statistics to achieve maximal performance (Supplementary Figure 1).

## Depression and Generalized Anxiety

We next evaluated the algorithms using real European ancestry summary statistics from the Pan-UK Biobank.[3,12] The outcome variables comprised all 52 non-collinear items in the UK Biobank Mental Health Questionnaire assessing lifetime major depression and generalized anxiety. For all experiments, we restricted analyses to high-quality HapMap3 autosomal, non-MHC, biallelic SNPs with INFO>0.9 and MAF>1% in UK Biobank and, when available, gnomAD genome/exome.

Figure 3 summarizes the results. Mendelianization identified three loci that surpassed the significance threshold of $5 \times 10^{-8}$. These loci also attained perfect SoMs of one, whereas most loci identified by the other algorithms had scores close to zero (Figure 3 (a)). Applying each of the three learned composite $Y\alpha_j$ genome-wide and plotting the $p$-values associated with the $z$-statistics $\mathcal{Z}_{\cdot|j}$ yielded a single tower per locus, as implied by the SoMs and predicted by Theorem 1 (Figure 3 (b)). The interpretable coefficients further delineated three distinct patient profiles: (i) uncontrollable worry with recurrent depression, (ii) melancholic-like depression characterized by insomnia and worthlessness, and (iii) controllable worry accompanied by somatic complaints (Supplementary Figure 2). Together, these results demonstrate that Mendelianization achieved the greatest spikedness while yielding interpretable coefficient patterns.

The three loci discovered by Mendelianization correspond to biologically plausible causal regions. The chromosome 15 signal with lead variant rs8038726 lies in the 15q11-q13 $GABA_A$ receptor cluster (GABRB3, GABRA5, GABRG3), implicated in abnormalities of inhibitory signaling, anxiety and mood disorders, and sleep disturbance.[11,15] The cluster also mediates the pharmacological effects of benzodiazepines and related hypnotics[28] – consistent with a phenotype dominated

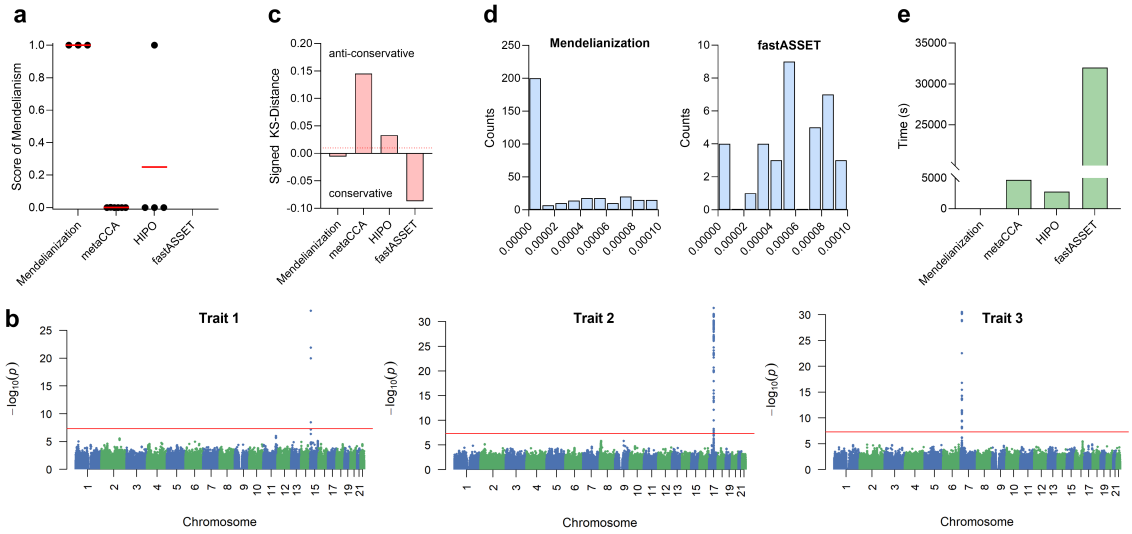


**Fig. 3. Main results for depression and generalized anxiety.** (a) Each dot denotes the SoM for the lead variant within a Manhattan-plot tower, and the red lines correspond to the mean SoMs. Only Mendelianization achieved perfect SoMs for all lead variants. No variants reached the genome-wide significance threshold under fastASSET. (b) For each of the three dots in the Mendelianization column of panel (a), the learned outcome corresponded to an essentially perfect Mendelian trait because it produced a Manhattan plot with a single tower, consistent with Theorem 1. (c) Only Mendelianization produced well-calibrated $p$-values; fastASSET was conservative. (d) Among methods with well-calibrated or conservative $p$-values, Mendelianization achieved the highest statistical power by a wide margin. (e) Mendelianization was approximately four orders of magnitude faster than alternative methods.

by uncontrollable worry and recurrent depression. The dense chromosome 17 block with lead variant rs2696455 falls within the 17q21.31 inversion containing CRHR1, whose variants have been associated with major depression, frequently via gene-environment interactions involving childhood maltreatment and through effects on hypothalamic-pituitary-adrenal (HPA) axis reactivity.[30,18,34] Melancholic depression, in particular, is the subtype most tightly linked to HPA-axis hyperactivity, elevated corticotropin-releasing hormone (CRH), and increased noradrenergic tone, with classic work demonstrating hypercortisolism and elevated CRH levels in melancholic compared with non-melancholic depression.[9] The chromosome 7 locus with lead variant rs1990622 overlaps TMEM106B, which multiple GWAS and pleiotropic analyses implicate in depression risk and symptomatology.[13,6] Experimental work has also linked this gene to an anxiety-like phenotype in mice,[19] broadly consistent with a profile characterized by controllable worry and prominent somatic complaints. Overall, these three loci localize to biologically plausible regions harboring putative causal variants in line with Theorem 1.

We next evaluated secondary performance metrics. Only Mendelianization produced well-calibrated $p$-values (Figure 3 (c)). FastASSET generated conservative $p$-values, whereas metaCCA and HIPO yielded anti-conservative estimates, indicating unreliable significance levels. We plot the individual histograms in Supplementary Figure 3. Among methods with calibrated or conservative behavior, Mendelianization achieved the greatest statistical power by far (Figure 3 (d)). The algorithm also required only 2 seconds to complete, while the other methods were slower by at least four orders of magnitude (Figure 3 (e)). Overall, Mendelianization uniquely combined well-calibrated inference, high power, and rapid computation. Finally, ablation results confirmed the necessity of all components of the algorithm (Supplementary Figure 4).

## Alcohol Use Disorder

We increased the difficulty in alcohol use disorder by running the algorithms with only 10 outcome variables corresponding to all individual items of the Alcohol Use Disorders Identification Test (AUDIT).[26] We again used European ancestry summary statistics from the Pan-UK Biobank with the same quality control. According to Theorem 1 and the simulation results, Mendelianization is unlikely to recover near-perfect Mendelian traits with only 10 outcomes, but we still expect it to find approximate Mendelian traits.

The results, summarized in Figure 4, support this expectation. Mendelianization did not achieve perfect SoMs, but every locus achieved higher scores than those from any competing method. Visualization of the per-outcome genome-wide Manhattan plots revealed a substantial, though not perfect, degree of Mendelianism (Figure 4 (b)). Each trait was associated with an interpretable set of canonical coefficients (Supplementary Figure 5).

The three loci have strong prior support for causal involvement in alcohol-related phenotypes. The chromosome 17 tower with lead variant rs41382552 again lies in the 17q21.31 inversion encompassing CRHR1–MAPT; CRHR1 variants associate with binge drinking and alcohol dependence[33] and interact with stress exposure,[24] consistent with a mechanism in which stress-axis signaling contributes causally to episodic binge drinking. The chromosome 2 signal with lead variant rs780093 maps to GCKR, and correlated variants robustly associate with alcohol consumption and AUDIT scores in large GWAS.[25] Experimental perturbations of the GCKR–glucokinase axis in C. elegans[32] and a GCKR P446L knock-in mouse model[16] further suggest that this pathway can causally influence alcohol-related behaviors. The chromosome 4 locus with lead variant rs11940694 overlaps KLB, encoding the FGF21 co-receptor; human GWAS and experimental manipulations of FGF21 pathways in rodents[36] and non-human primates[7] converge on a model in which this axis directly

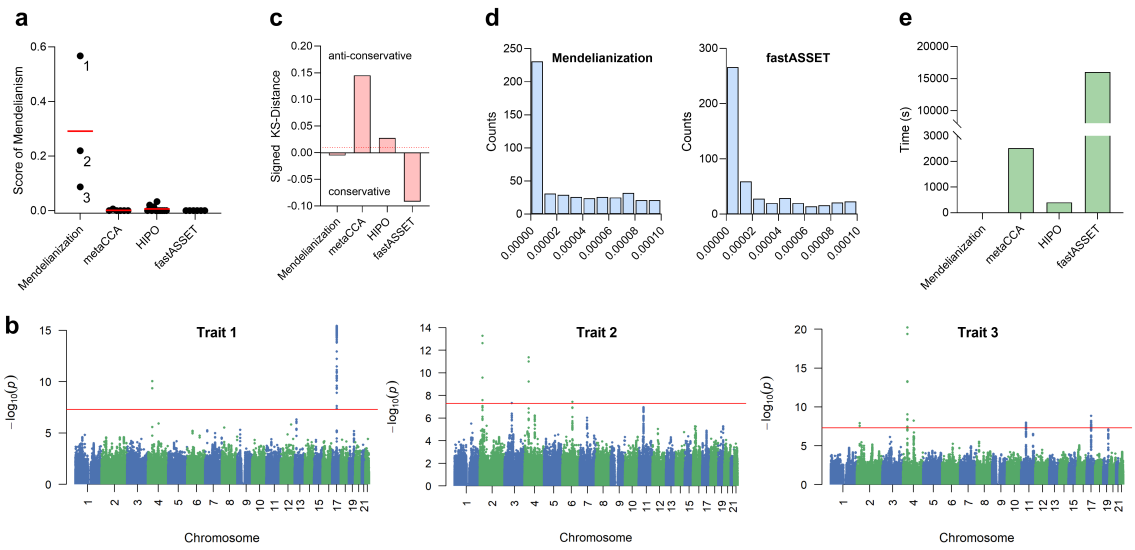


**Fig. 4. Main results for alcohol use disorder.** (a) Mendelianization achieved the highest SoMs for each lead variant. (b) The learned outcome for each point in the Mendelianization column of panel (a) approximated a Mendelian trait. (c) Only Mendelianization produced well-calibrated $p$-values. (d) Both Mendelianization and fastASSET attained high statistical power. (e) Mendelianization completed in 0.58 seconds.

regulates alcohol preference and intake. External data therefore support a locus-level causal role for CRHR1/17q21.31, GCKR, and KLB in distinct patterns of alcohol use.

As before, Mendelianization was the only method to yield well-calibrated $p$-values (Figure 4 (c); Supplementary Figure 6), while both Mendelianization and fastASSET achieved high statistical power (Figure 4 (d)). Finally, Mendelianization completed at least three orders of magnitude faster than competing methods (Figure 4 (e)), and ablation results confirmed the necessity of the individual components of the algorithm (Supplementary Figure 7). We conclude that Mendelianization was again the only method that simultaneously achieved high degrees of Mendelianism, well-calibrated inference, high power, and rapid computation.

## Discussion

Mendelianization learns composite outcomes whose associations asymptotically concentrate at a single locus. The locus contains at least one causal variant under just four structural assumptions – low-dimensional confounding, outcome diversity, signal floors, and isolation-spillover. The algorithm also handles partial sample-overlap, yields cross-outcome comparable canonical coefficients, and quantifies single-locus concentration via the SoM. In short, Mendelianization performs genome-wide, locus-level causal inference using a substantially modified one-dimensional CCA.

Despite these strengths, Mendelianization has limitations. First, it cannot resolve LD within a locus because it is a one-dimensional method; isolating variant-level causality generally requires fine-mapping – a separate class of techniques fundamentally distinct from Mendelianization. Second, Mendelianization cannot accommodate completely non-overlapping samples between outcome pairs, which poses challenges when distinct populations are assessed with different measures. Third, Mendelianization remains a linear method that may miss nonlinear confounding and causal effects. Future research should therefore explore combining Mendelianization with fine-mapping approaches, extending it to handle fully non-overlapping samples, and generalizing it to nonlinear settings.

In summary, Mendelianization provides a powerful framework for isolating causal loci to the extent permitted by available outcomes, advancing a reductionist approach that helps clarify – rather than complicate – the pathophysiological understanding of complex disease.

## Competing interests

No competing interest is declared.

## Author contributions statement

E.V.S. conceived, designed, and executed the project in its entirety.

## Acknowledgments

TBD

## References

1. Arvanitis, M., Tayeb, K., Strober, B.J., Battle, A.: Redefining tissue specificity of genetic regulation of gene expression in the presence of allelic heterogeneity. The American Journal of Human Genetics **109**(2), 223–239 (2022)
2. Ballard, J.L., O'Connor, L.J.: Shared components of heritability across genetically correlated traits. The American Journal of Human Genetics **109**(6), 989–1006 (2022)
3. Bycroft, C., Freeman, C., Petkova, D., Band, G., Elliott, L.T., Sharp, K., Motyer, A., Vukcevic, D., Delaneau, O., O'Connell, J., et al.: The uk biobank resource with deep phenotyping and genomic data. Nature **562**(7726), 203–209 (2018)

4. Cichonska, A., Rousu, J., Marttinen, P., Kangas, A.J., Soininen, P., Lehtimäki, T., Raitakari, O.T., Järvelin, M.R., Salomaa, V., Ala-Korpela, M., et al.: metacca: summary statistics-based multivariate meta-analysis of genome-wide association studies using canonical correlation analysis. Bioinformatics **32**(13), 1981–1989 (2016)
5. Consortium, .G.P., et al.: A global reference for human genetic variation. Nature **526**(7571), 68 (2015)
6. Fabbri, C., Pain, O., Hagenaars, S.P., Lewis, C.M., Serretti, A.: Transcriptome-wide association study of treatment-resistant depression and depression subtypes for drug repurposing. Neuropsychopharmacology **46**(10), 1821–1829 (2021)
7. Flippo, K.H., Trammell, S.A., Gillum, M.P., Aklan, I., Perez, M.B., Yavuz, Y., Smith, N.K., Jensen-Cody, S.O., Zhou, B., Claflin, K.E., et al.: Fgf21 suppresses alcohol consumption through an amygdalo-striatal circuit. Cell Metabolism **34**(2), 317–328 (2022)
8. Fried, E.I., Nesse, R.M.: Depression is not a consistent syndrome: An investigation of unique symptom patterns in the star* d study. Journal of Affective Disorders **172**, 96–102 (2015)
9. Gold, P., Chrousos, G.: Organization of the stress system and its dysregulation in melancholic and atypical depression: high vs low crh/ne states. Molecular Psychiatry **7**(3), 254–275 (2002)
10. Grotzinger, A.D., Rhemtulla, M., de Vlaming, R., Ritchie, S.J., Mallard, T.T., Hill, W.D., Ip, H.F., Marioni, R.E., McIntosh, A.M., Deary, I.J., et al.: Genomic structural equation modelling provides insights into the multivariate genetic architecture of complex traits. Nature Human Behaviour **3**(5), 513–525 (2019)
11. Hodges, L.M., Fyer, A.J., Weissman, M.M., Logue, M.W., Haghighi, F., Evgrafov, O., Rotondo, A., Knowles, J.A., Hamilton, S.P.: Evidence for linkage and association of gabrb3 and gabra5 to panic disorder. Neuropsychopharmacology **39**(10), 2423–2431 (2014)
12. Karczewski, K.J., Gupta, R., Kanai, M., Lu, W., Tsuo, K., Wang, Y., Walters, R.K., Turley, P., Callier, S., Shah, N.N., et al.: Pan-uk biobank genome-wide association analyses enhance discovery and resolution of ancestry-enriched effects. Nature Genetics pp. 1–10 (2025)
13. Li, Y., Dang, X., Chen, R., Teng, Z., Wang, J., Li, S., Yue, Y., Mitchell, B.L., Zeng, Y., Yao, Y.G., et al.: Cross-ancestry genome-wide association study and systems-level integrative analyses implicate new risk genes and therapeutic targets for depression. Nature Human Behaviour pp. 1–18 (2025)
14. Little, R.J., Rubin, D.B.: Statistical analysis with missing data. John Wiley & Sons (2019)
15. Luscher, B., Shen, Q., Sahir, N.: The gabaergic deficit hypothesis of major depressive disorder. Molecular Psychiatry **16**(4), 383–406 (2011)
16. Mehrhoff, E.A., Bunch, L., Bower, M., Fowler, N., Yang, E., Hendricks, L., Verner, J., Lee, H., Aki, C., Branney, M., Funke, A., Ehringer, M.A.: Alcohol-related behaviors in a mouse model containing the human GCKR SNP rs1260326 (p446l). Alcohol: Clinical and Experimental Research **48**(S1), 117–118 (2024). https://doi.org/10.1111/acer.15316, speaker abstract, 47th Annual Research Society on Alcoholism Scientific Meeting, Minneapolis, MN, USA
17. Mostafavi, H., Spence, J.P., Naqvi, S., Pritchard, J.K.: Systematic differences in discovery of genetic effects on gene expression and complex traits. Nature Genetics **55**(11), 1866–1875 (2023)
18. Nievergelt, C.M., Maihofer, A.X., Atkinson, E.G., Chen, C.Y., Choi, K.W., Coleman, J.R., Daskalakis, N.P., Duncan, L.E., Polimanti, R., Aaronson, C., et al.: Genome-wide association analyses identify 95 risk loci and provide insights into the neurobiology of post-traumatic stress disorder. Nature genetics **56**(5), 792–808 (2024)
19. Perneel, J., Lastra Osua, M., Alidadiani, S., Peeters, N., De Witte, L., Heeman, B., Manzella, S., De Rycke, R., Brooks, M., Perkerson, R.B., et al.: Increased tmem106b levels lead to lysosomal dysfunction which affects synaptic signaling and neuronal health. Molecular neurodegeneration **20**(1), 1–26 (2025)
20. Privé, F.: Optimal linkage disequilibrium splitting. Bioinformatics **38**(1), 255–256 (2022)
21. Qi, G., Chatterjee, N.: Heritability informed power optimization (hipo) leads to enhanced detection of genetic associations across multiple traits. PLoS Genetics **14**(10), e1007549 (2018)
22. Qi, G., Chhetri, S.B., Ray, D., Dutta, D., Battle, A., Bhattacharjee, S., Chatterjee, N.: Genome-wide large-scale multi-trait analysis characterizes global patterns of pleiotropy and unique trait-specific variants. Nature Communications **15**(1), 6985 (2024)
23. Quinlan, A.R., Hall, I.M.: Bedtools: a flexible suite of utilities for comparing genomic features. Bioinformatics **26**(6), 841–842 (2010)
24. Ray, L.A., Sehl, M., Bujarski, S., Hutchison, K., Blaine, S., Enoch, M.A.: The crhr1 gene, trauma exposure, and alcoholism risk: a test of g× e effects. Genes, Brain and Behavior **12**(4), 361–369 (2013)
25. Sanchez-Roige, S., Palmer, A.A., Fontanillas, P., Elson, S.L., 23andMe Research Team, t.S.U.D.W.G.o.t.P.G.C., Adams, M.J., Howard, D.M., Edenberg, H.J., Davies, G., Crist, R.C., et al.: Genome-wide association study meta-analysis of the alcohol use disorders identification test (audit) in two population-based cohorts. American Journal of Psychiatry **176**(2), 107–118 (2019)
26. Saunders, J.B., Aasland, O.G., Babor, T.F., De la Fuente, J.R., Grant, M.: Development of the alcohol use disorders identification test (audit): Who collaborative project on early detection of persons with harmful alcohol consumption-ii. Addiction **88**(6), 791–804 (1993)
27. Schaid, D.J., Chen, W., Larson, N.B.: From genome-wide associations to candidate causal variants by statistical fine-mapping. Nature Reviews Genetics **19**(8), 491–504 (2018)
28. Sigel, E., Ernst, M.: The benzodiazepine binding sites of gabaa receptors. Trends in Pharmacological Sciences **39**(7), 659–671 (2018)
29. Strobl, E.V., Gamazon, E.R.: Transcriptome-wide root causal inference. PLOS Computational Biology **21**(9), e1013461 (2025)
30. Su, X., Li, W., Lv, L., Li, X., Yang, J., Luo, X.J., Liu, J.: Transcriptome-wide association study provides insights into the genetic component of gene expression in anxiety. Frontiers in Genetics **12**, 740134 (2021)
31. Svishcheva, G.R., Tiys, E.S., Elgaeva, E.E., Feoktistova, S.G., Timmers, P.R., Sharapov, S.Z., Axenovich, T.I., Tsepilov, Y.A.: A novel framework for analysis of the shared genetic background of correlated traits. Genes **13**(10), 1694 (2022)
32. Thompson, A., Cook, J., Choquet, H., Jorgenson, E., Yin, J., Kinnunen, T., Barclay, J., Morris, A.P., Pirmohamed,

M.: Functional validity, role, and implications of heavy alcohol consumption genetic loci. Science Advances **6**(3), eaay5034 (2020)
33. Treutlein, J., Kissling, C., Frank, J., Wiemann, S., Dong, L., Depner, M., Saam, C., Lascorz, J., Soyka, M., Preuss, U., et al.: Genetic association of the human corticotropin releasing hormone receptor 1 (crhr1) with binge drinking and alcohol intake patterns in two independent samples. Molecular Psychiatry **11**(6), 594–602 (2006)
34. Tyrka, A.R., Price, L.H., Gelernter, J., Schepker, C., Anderson, G.M., Carpenter, L.L.: Interaction of childhood maltreatment with the corticotropin-releasing hormone receptor gene: effects on hypothalamic-pituitary-adrenal axis reactivity. Biological psychiatry **66**(7), 681–685 (2009)
35. Uffelmann, E., Huang, Q.Q., Munung, N.S., De Vries, J., Okada, Y., Martin, A.R., Martin, H.C., Lappalainen, T., Posthuma, D.: Genome-wide association studies. Nature Reviews Methods Primers **1**(1), 59 (2021)
36. Wang, T., Tyler, R.E., Ilaka, O., Cooper, D., Farokhnia, M., Leggio, L.: The crosstalk between fibroblast growth factor 21 (fgf21) system and substance use. iScience **27**(7) (2024)
37. Zhang, Z., Jung, J., Kim, A., Suboc, N., Gazal, S., Mancuso, N.: A scalable approach to characterize pleiotropy across thousands of human diseases and complex traits using gwas summary statistics. The American Journal of Human Genetics **110**(11), 1863–1874 (2023)
38. Zou, Y., Carbonetto, P., Xie, D., Wang, G., Stephens, M.: Fast and flexible joint fine-mapping of multiple traits via the sum of single effects model. bioRxiv pp. 2023–04 (2024)
39. Zuber, V., Cronjé, T., Cai, N., Gill, D., Bottolo, L.: Bayesian causal graphical model for joint mendelian randomization analysis of multiple exposures and outcomes. The American Journal of Human Genetics **112**(5), 1173–1198 (2025)
40. Zuber, V., Lewin, A., Levin, M.G., Haglund, A., Ben-Aicha, S., Emanueli, C., Damrauer, S., Burgess, S., Gill, D., Bottolo, L.: Multi-response mendelian randomization: Identification of shared and distinct exposures for multimorbidity and multiple related disease outcomes. The American Journal of Human Genetics **110**(7), 1177–1199 (2023)

# Supplementary Materials

## A. Derivaton of CCA Solution

Correlation and, likewise, the solution to Expression (1) is scale invariant, so we can impose the constraint $\alpha^\top \Sigma \alpha = 1$ and maximize the numerator $r_j \alpha$. We invoke the Cauchy-Schwarz inequality in the $\Sigma$-inner product:

$$r_j \alpha = \langle r_j \Sigma^{-1}, \alpha^\top \rangle_\Sigma \leq \| r_j \Sigma^{-1} \|_\Sigma \| \alpha^\top \|_\Sigma = \sqrt{r_j \Sigma^{-1} r_j^\top},$$

where the last equality follows because we imposed the constraint $\|\alpha^\top\|_\Sigma = \sqrt{\alpha^\top \Sigma \alpha} = 1$. The inequality becomes an equality when $\alpha = \frac{\Sigma^{-1} r_j^\top}{\sqrt{r_j \Sigma^{-1} r_j^\top}}$. Recall that the solution to Expression (2) is scale-invariant, so we set $\alpha = A r_j^\top$ with $A = \Sigma^{-1}$.

# A. Finite $m$ Assumptions

**Assumption 5.** *(Low-dimensional latent confounding) For a prespecified $m \in \mathbb{N}^+$, the confounding effects are contained in a subspace $U^m \subset \mathbb{R}^m$, and there exists an integer constant $r_0 \geq 0$ such that $\dim(U^m) \leq r_0$.*

**Assumption 6.** *(Outcome diversity) We have $\max_{k \notin \mathcal{S}_j} \frac{1}{m} |\langle r_k^{m\perp}, r_j^{m\perp} \rangle_{A^m}| = 0$ for a prespecified $m$.*

**Assumption 7.** *(Signal floor) For a prespecified $m$, if $l \in \mathcal{S}_j$ causes $Y\alpha_j$, then there exists a constant $\beta_{\mathrm{c}} > 0$ such that $\frac{1}{m} \|r_l^{m\perp}\|_{A^m}^2 \geq \beta_{\mathrm{c}}$. Moreover, there exists a constant $\beta_{\mathrm{on}} > 0$ such that $\frac{1}{m} \|r_l^{m\|}\|_{A^m}^2 \geq \beta_{\mathrm{on}}$ for every $l \in \mathcal{S}_j$ that does not cause $Y\alpha_j$.*

**Assumption 8.** *(Isolation-spillover) If some $l \in \mathcal{S}_j$ causes $Y\alpha_j$, we have $\max_{k \notin \mathcal{S}_j} \frac{1}{m} |\langle r_k^{m\|}, r_j^{m\|} \rangle_{A^m}| = 0$ for a prespecified $m$. On the other hand, if there does not exist $l \in \mathcal{S}_j$ that causes $Y\alpha_j$, then $\exists k \notin \mathcal{S}_j$ with $|\cos_{A^m}(r_j^{m\|}, r_k^{m\|})| \geq \gamma_j > 0$ and $\frac{1}{m} \|r_k^{m\|}\|_{A^m}^2 \geq \beta_{\mathrm{off}} > 0$ for a prespecified $m$.*

# B. Pseudocode & Time Complexity Analysis

We present the pseudocode for Mendelianization in Algorithm 1. The algorithm first computes $\widehat{S}$ in Line 1, which serves as a consistent estimator of $\Gamma$ under Theorem 2. Using this estimate, it then derives the raw canonical coefficients in Line 2 by solving Expression (5). These coefficients also form the basis of the $Q$ statistics in Expression (3) that enable hypothesis testing; recall that each $Q_j$ follows a $\chi^2_m$ distribution under the null, so the algorithm can readily compute $p$-values in Line 4. Next, Mendelianization computes SoMs for lead variants that surpass a genome-wide significance threshold (e.g., $5 \times 10^{-8}$) in Line 5. Finally, the algorithm converts the raw coefficients into their interpretable form $\alpha'$ by applying Proposition 4 in Line 6. Thus, although the theory of Mendelianization is intricate, the resulting algorithm is remarkably straightforward.

**Algorithm 1** Mendelianization

**Input:** multiple outcome $z$-statistics

**Output:** interpretable canonical coefficients $\alpha'$; test statistics $Q$ and $p$-values; Scores of Mendelianism $\mathcal{M}$

1: $\Gamma \leftarrow$ compute $\widehat{S}$ using $z$-statistics (Theorem 2)
2: $\alpha \leftarrow \Omega z^\top$
3: $Q \leftarrow \mathrm{diag}(z\alpha)$
4: $p \leftarrow$ right tail probability of each $Q_j$ under $\chi^2_m$
5: $\mathcal{M} \leftarrow$ compute SoMs (Equation (6); Theorem 1)
6: $\alpha' \leftarrow \alpha / \sqrt{\mathrm{diag}(\Omega)}$ (Proposition 4)

The algorithm also has a favorable time complexity. In particular, estimating the correlation matrix $\Gamma$ described in Line 1 costs $O(qm^2)$ time. In Line 2, inverting $\Gamma$ requires $O(m^3)$ time, and the matrix product $\Omega z^\top$ required to estimate $\alpha$ costs $O(m^2 q)$. Constructing all $Q_j$ statistics and evaluating their $p$-values in Lines 3 and 4 adds $O(mq)$, and computing $\alpha'$ in Line 6 requires another $O(mq)$. For the SoM step, suppose there are $s$ significant variants, and each constitutes its own tower to take the worst case. Computing each statistic $\mathcal{Z}_{j|k}$ requires $O(m)$ time for the numerator and $O(m^2)$ for the denominator in Equation (6). Moreover, the numerator must be evaluated genome-wide over $q$ variants, so an upper bound on the total cost of SoM amounts to $O(s(mq + m^2))$. Summing these components, the overall complexity of Mendelianization is $O(m^3 + m^2 q + sqm)$, where the simplification uses $q \gg m$. Thus, the algorithm scales linearly with the number of variants and cubically with the number of outcome variables. In practice, Mendelianization completes within seconds and is at least three orders of magnitude faster than the competing methods.

# C. Proofs

**Theorem 1.** *(Asymptotic Mendelianism) Consider Assumptions 1-4 and suppose $\Sigma^m$ is positive definite for sufficiently large $m$. We have $\frac{1}{m} r_j^m \alpha_j^m > 0$ for sufficiently large $m$ and $\max_{k \notin \mathcal{S}_j} |\frac{1}{m} r_k^m \alpha_j^m| \to 0$ (i.e., Mendelianism is asymptotically achieved for $Y\alpha_j^m$ at $\mathcal{S}_j$) if and only if at least one variant $l \in \mathcal{S}_j$ causes $Y\alpha_j^m$.*

*Proof* We first prove the forward direction by contrapositive. Suppose there does not exist a variant $l \in \mathcal{S}_j$ that causes $Y\alpha_j^m$. As a result, $j \in \mathcal{S}_j$ does not cause $Y\alpha_j^m$. We need to show $\frac{1}{m} r_j^m \alpha_j^m \to 0$ or $\max_{k \notin \mathcal{S}_j} |\frac{1}{m} r_k^m \alpha_j^m| > 0$ (or both) for $m$ sufficiently large. We will show that $\max_{k \notin \mathcal{S}_j} |\frac{1}{m} r_k^m \alpha_j^m| > 0$.

Assumption 1 allows us to write:

$$\frac{1}{m}\langle r_k^m, r_j^m\rangle_{A^m} = \frac{1}{m}\langle r_k^{m\parallel}, r_j^{m\parallel}\rangle_{A^m} + \frac{1}{m}\langle r_k^{m\perp}, r_j^{m\perp}\rangle_{A^m}. \tag{7}$$

Choose $M_1 > 0$. By Assumptions 3 and 4, $\exists M_2 \geq M_1$ such that $\forall m \geq M_2, \exists k \notin \mathcal{S}_j$ where:

$$\left|\langle r_k^{m\parallel}, r_j^{m\parallel}\rangle_{A^m}\right| \geq \|r_j^{m\parallel}\|_{A^m}\|r_k^{m\parallel}\|_{A^m}\gamma_j \geq m\sqrt{\beta_{\text{on}}}\sqrt{\beta_{\text{off}}}\gamma_j.$$

The reverse triangle inequality allows us to write:

$$\begin{aligned}\left|\langle r_k^m, r_j^m\rangle_{A^m}\right| &= \left|\langle r_k^{m\parallel}, r_j^{m\parallel}\rangle_{A^m} + \langle r_k^{m\perp}, r_j^{m\perp}\rangle_{A^m}\right| \\ &\geq \left|\left|\langle r_k^{m\parallel}, r_j^{m\parallel}\rangle_{A^m}\right| - \left|\langle r_k^{m\perp}, r_j^{m\perp}\rangle_{A^m}\right|\right| \geq \left|\langle r_k^{m\parallel}, r_j^{m\parallel}\rangle_{A^m}\right| - \left|\langle r_k^{m\perp}, r_j^{m\perp}\rangle_{A^m}\right|.\end{aligned}$$

Let $\varepsilon_m = \max_{k \notin \mathcal{S}_j} \frac{1}{m}|\langle r_k^{m\perp}, r_j^{m\perp}\rangle_{A^m}|$. Under Assumption 2, $\exists M_3 \geq M_2$ such that $\forall m > M_3$ we have $\varepsilon_m < \sqrt{\beta_{\text{on}}}\sqrt{\beta_{\text{off}}}\gamma_j$. As a result:

$$\frac{1}{m}|r_k^m \alpha_j^m| = \frac{1}{m}\left|\langle r_k^m, r_j^m\rangle_{A^m}\right| \geq \sqrt{\beta_{\text{on}}}\sqrt{\beta_{\text{off}}}\gamma_j - \varepsilon_m > 0,$$

where $\varepsilon_m = \max_{k \notin \mathcal{S}_j} \frac{1}{m}|\langle r_k^{m\perp}, r_j^{m\perp}\rangle_{A^m}|$. This completes the proof of sufficiency – i.e., asymptotic Mendelianism is sufficient to establish causality.

To prove the backward direction, assume that there exists at least one variant $l \in \mathcal{S}_j$ that causes $Y\alpha_j^m$. Then, we can invoke Assumption 1 and write Equation (7) for any $k$. We have two situations:

1. If $k \notin \mathcal{S}_j$, then $\max_{k \notin \mathcal{S}_j} \frac{1}{m}|\langle r_k^{m\perp}, r_j^{m\perp}\rangle_{A^m}| \to 0$ by Assumption 2. Moreover, $\frac{1}{m}|\langle r_k^{m\parallel}, r_j^{m\parallel}\rangle_{A^m}| \to 0$ by the first clause of Assumption 4. Hence, $\max_{k \notin \mathcal{S}_j} \frac{1}{m}|\langle r_k, r_j\rangle_{A^m}| \to 0$.
2. If $k \in \mathcal{S}_j$ and $k = j$, then we have two situations. If $j$ causes $Y\alpha_j^m$, then the first clause of Assumption 3 allows us to write:

$$\frac{1}{m}\langle r_j^m, r_j^m\rangle_{A^m} = \frac{1}{m}|\langle r_j^m, r_j^m\rangle_{A^m}| = \frac{1}{m}\|r_j^{m\perp}\|_{A^m}^2 + \frac{1}{m}\|r_j^{m\parallel}\|_{A^m}^2 \geq \frac{1}{m}\|r_j^{m\perp}\|_{A^m}^2 \geq \beta_{\text{c}} > 0,$$

for sufficiently large $m$. On the other hand, if $j$ does not cause $Y\alpha_j^m$, then the second clause of Assumption 3 allows us to write:

$$\frac{1}{m}\langle r_j^m, r_j^m\rangle_{A^m} = \frac{1}{m}\|r_j^{m\perp}\|_{A^m}^2 + \frac{1}{m}\|r_j^{m\parallel}\|_{A^m}^2 \geq \frac{1}{m}\|r_j^{m\parallel}\|_{A^m}^2 \geq \beta_{\text{on}} > 0,$$

for sufficiently large $m$.

We conclude that $\frac{1}{m} r_j^m \alpha_j^m > 0$ for sufficiently large $m$ and $\max_{k \notin \mathcal{S}_j} |\frac{1}{m} r_k^m \alpha_j^m| \to 0$, which completes the proof of necessity – i.e., asymptotic Mendelianism is necessary to establish causality. □

**Proposition 1.** *(Perfect Mendelianism at finite m) Consider Assumptions 5-8 with a prespecified $m$ and suppose $\Sigma^m$ is positive definite. We have $\frac{1}{m} r_j^m \alpha_j^m > 0$ and $\max_{k \notin \mathcal{S}_j} |\frac{1}{m} r_k^m \alpha_j^m| = 0$ (i.e., Mendelianism is perfectly achieved for $\mathcal{S}_j$ at $m$) if and only if at least one variant $l \in \mathcal{S}_j$ causes $Y\alpha_j^m$.*

*Proof* The proof almost exactly parallels that of Theorem 1, but we include it here for completeness. We first prove the forward direction by contrapositive. Suppose there does not exist a variant $l \in \mathcal{S}_j$ that causes $Y\alpha_j^m$. As a result, $j \in \mathcal{S}_j$ does not cause $Y\alpha_j^m$. We need to show $\frac{1}{m} r_j^m \alpha_j^m = 0$ or $\max_{k \notin \mathcal{S}_j} |\frac{1}{m} r_k^m \alpha_j^m| > 0$ (or both) for prespecified $m$. We will show that $\max_{k \notin \mathcal{S}_j} |\frac{1}{m} r_k^m \alpha_j^m| > 0$.

Assumption 5 allows us to write:

$$\frac{1}{m}\langle r_k^m, r_j^m\rangle_{A^m} = \frac{1}{m}\langle r_k^{m\parallel}, r_j^{m\parallel}\rangle_{A^m} + \frac{1}{m}\langle r_k^{m\perp}, r_j^{m\perp}\rangle_{A^m}. \tag{8}$$

By Assumptions 7 and 8, there exists $k \notin \mathcal{S}_j$ where:

$$\left|\langle r_k^{m\parallel}, r_j^{m\parallel}\rangle_{A^m}\right| \geq \|r_j^{m\parallel}\|_{A^m}\|r_k^{m\parallel}\|_{A^m}\gamma_j \geq m\sqrt{\beta_{\text{on}}}\sqrt{\beta_{\text{off}}}\gamma_j.$$

The reverse triangle inequality allows us to write:

$$\begin{aligned}\left|\langle r_k^m, r_j^m\rangle_{A^m}\right| &= \left|\langle r_k^{m\parallel}, r_j^{m\parallel}\rangle_{A^m} + \langle r_k^{m\perp}, r_j^{m\perp}\rangle_{A^m}\right| \\ &\geq \left|\left|\langle r_k^{m\parallel}, r_j^{m\parallel}\rangle_{A^m}\right| - \left|\langle r_k^{m\perp}, r_j^{m\perp}\rangle_{A^m}\right|\right| \geq \left|\langle r_k^{m\parallel}, r_j^{m\parallel}\rangle_{A^m}\right| - \left|\langle r_k^{m\perp}, r_j^{m\perp}\rangle_{A^m}\right|.\end{aligned}$$

Thus, under Assumption 6, we have:

$$\frac{1}{m}|r_k^m \alpha_j^m| = \frac{1}{m}\left|\langle r_k^m, r_j^m\rangle_{A^m}\right| \geq \sqrt{\beta_{\text{on}}}\sqrt{\beta_{\text{off}}}\gamma_j > 0.$$

This completes the proof of sufficiency – i.e., perfect Mendelianism is sufficient to establish causality.

To prove the backward direction, assume that there exists at least one variant $l \in \mathcal{S}_j$ that causes $Y\alpha_j^m$. Then, we can invoke Assumption 5 and write Equation (8) for any $k$. We have two situations:

1.If $k \notin \mathcal{S}_j$, then $\max_{k\notin\mathcal{S}_j} \frac{1}{m}|\langle r_k^{m\perp}, r_j^{m\perp}\rangle_{A^m}| = 0$ by Assumption 6. Moreover, $\frac{1}{m}|\langle r_k^{m\parallel}, r_j^{m\parallel}\rangle_{A^m}| = 0$ by the first clause of Assumption 8. Hence, $\max_{k\notin\mathcal{S}_j} \frac{1}{m}|\langle r_k, r_j\rangle_{A^m}| = 0$.
2.If $k \in \mathcal{S}_j$ and $k = j$, then we have two situations. If $j$ causes $Y\alpha_j^m$, then the first clause of Assumption 7 allows us to write:

$$\frac{1}{m}\langle r_j^m, r_j^m\rangle_{A^m} = \frac{1}{m}|\langle r_j^m, r_j^m\rangle_{A^m}| = \frac{1}{m}\|r_j^{m\perp}\|_{A^m}^2 + \frac{1}{m}\|r_j^{m\parallel}\|_{A^m}^2 \geq \frac{1}{m}\|r_j^{m\perp}\|_{A^m}^2 \geq \beta_{\mathrm{c}} > 0.$$

On the other hand, if $j$ does not cause $Y\alpha_j^m$, then the second clause of Assumption 7 allows us to write:

$$\frac{1}{m}\langle r_j^m, r_j^m\rangle_{A^m} = \frac{1}{m}\|r_j^{m\perp}\|_{A^m}^2 + \frac{1}{m}\|r_j^{m\parallel}\|_{A^m}^2 \geq \frac{1}{m}\|r_j^{m\parallel}\|_{A^m}^2 \geq \beta_{\mathrm{on}} > 0.$$

We conclude that $\frac{1}{m}r_j^m\alpha_j^m > 0$ and $\max_{k\notin\mathcal{S}_j}|\frac{1}{m}r_k^m\alpha_j^m| = 0$, which completes the proof of necessity – i.e., perfect Mendelianism is necessary to establish causality. □

**Proposition 2.** *Assume i.i.d. samples across individuals, $V \perp\!\!\!\perp Y$ and $V \perp\!\!\!\perp M|Y$. Then $\mathbb{E}\big[z_{ja}z_{jb} \mid M\big] = \frac{n_{ab}}{\sqrt{n_a n_b}}\,\Sigma_{ab}^{\mathrm{obs}}$.*

*Proof* Note that $z$-statistics are computed on observed samples, so we can write:

$$\begin{aligned}\mathbb{E}\big[z_{ja}z_{jb} \mid M\big] &= \mathbb{E}\left(\sqrt{n_a}\widehat{r}_{ja}\sqrt{n_b}\widehat{r}_{jb} \mid M\right) = \mathbb{E}\left(\sqrt{n_a}\widehat{r}_{ja}\sqrt{n_b}\widehat{r}_{jb} \mid M_a = 1, M_b = 1, M_{-a,-b}\right)\\ &= \mathbb{E}\left(\frac{\sqrt{n_a}}{n_a}\sum_{i\in S_a} V_{ji}Y_{ai}\frac{\sqrt{n_b}}{n_b}\sum_{k\in S_b} V_{jk}Y_{bk} \mid M_a = 1, M_b = 1, M_{-a,-b}\right).\end{aligned} \tag{9}$$

Cross-terms cancel by i.i.d. samples. Moreover, $V \perp\!\!\!\perp Y$ and $V \perp\!\!\!\perp M|Y$ imply $V \perp\!\!\!\perp Y|M$ by the contraction and weak union semi-graphoid axioms. As a result, we have:

$$\mathbb{E}\left(\frac{\sqrt{n_a}}{n_a}\sum_{i\in S_a} V_{ji}Y_{ai}\frac{\sqrt{n_b}}{n_b}\sum_{k\in S_b} V_{jk}Y_{bk} \mid \widetilde{M}\right) = \frac{n_{ab}}{\sqrt{n_a}\sqrt{n_b}}\mathbb{E}\left(V_j^2 \mid \widetilde{M}\right)\mathbb{E}\left(Y_aY_b \mid \widetilde{M}\right),$$

where $\widetilde{M} = \{M_a = 1, M_b = 1, M_{-a,-b}\}$. Recall that genotypes were standardized so that $\mathbb{E}(V_j) = 0$ and $\mathrm{Var}(V_j) = 1$. Thus:

$$\mathbb{E}\left(V_j^2 \mid \widetilde{M}\right) = \mathbb{E}_Y\left(\mathbb{E}\left(V_j^2 \mid Y, \widetilde{M}\right) \mid \widetilde{M}\right) = \mathbb{E}_Y\left(\mathbb{E}\left(V_j^2 \mid Y\right) \mid \widetilde{M}\right) = \mathbb{E}_Y(1 \mid \widetilde{M}) = 1.$$

We applied $V \perp\!\!\!\perp M|Y$ in the second equality, and $V \perp\!\!\!\perp Y$ and standardization in the third. Hence:

$$\frac{n_{ab}}{\sqrt{n_a}\sqrt{n_b}}\mathbb{E}\left(V_j^2 \mid \widetilde{M}\right)\mathbb{E}\left(Y_aY_b \mid \widetilde{M}\right) = \frac{n_{ab}}{\sqrt{n_a}\sqrt{n_b}}\mathbb{E}\left(Y_aY_b \mid \widetilde{M}\right) = \frac{n_{ab}}{\sqrt{n_a n_b}}\,\Sigma_{ab}^{\mathrm{obs}}.$$

We have established the equality $\mathbb{E}\big[z_{ja}z_{jb} \,\big|\, \widetilde{M}\big] = \mathbb{E}\big[z_{ja}z_{jb} \,\big|\, M\big] = \frac{n_{ab}}{\sqrt{n_a n_b}}\,\Sigma_{ab}^{\mathrm{obs}}$. □

**Proposition 3.** *Assume i.i.d. samples across individuals, $p_a = \mathbb{P}(M_a = 1) > 0$, and $p_b = \mathbb{P}(M_b = 1) > 0$. We have $\mathbb{E}\Big(\frac{1}{n}\sum_t Y'_{a,t}Y'_{b,t} \mid M\Big) = \Gamma_{ab} \overset{a.s.}{\to} \frac{p_{ab}}{\sqrt{p_a p_b}}\Sigma_{ab}^{\mathrm{obs}}$ and $\mathbb{E}\Big(\frac{1}{n}\sum_t Y'_{a,t}Y'_{b,t}\Big) \to \frac{p_{ab}}{\sqrt{p_a p_b}}\Sigma_{ab}^{\mathrm{obs}}$ as $n \to \infty$.*

*Proof* For the first statement, we may write:

$$\mathbb{E}\Big(\frac{1}{n}\sum_t Y'_{a,t}Y'_{b,t} \mid M\Big) = \frac{1}{\sqrt{n_a}\sqrt{n_b}}\mathbb{E}\left(\sum_t M_{at}Y_{at}M_{bt}Y_{bt} \mid M\right) = \frac{n_{ab}}{\sqrt{n_a}\sqrt{n_b}}\mathbb{E}\left(Y_aY_b|M_a = 1, M_b = 1\right) = \frac{n_{ab}}{\sqrt{n_a}\sqrt{n_b}}\Sigma_{ab}^{\mathrm{obs}} = \Gamma_{ab}.$$

Note that:

$$\frac{n_{ab}}{\sqrt{n_a}\sqrt{n_b}} = \frac{n\widehat{p}_{ab}}{\sqrt{n\widehat{p}_a}\sqrt{n\widehat{p}_b}} = \frac{\widehat{p}_{ab}}{\sqrt{\widehat{p}_a}\sqrt{\widehat{p}_b}} \overset{a.s.}{\to} \frac{p_{ab}}{\sqrt{p_a}\sqrt{p_b}}.$$

where $p_{ab} = \mathbb{P}(S_a \cap S_b)$, and almost sure convergence holds by the strong law of large numbers and the continuous mapping theorem. Hence, $\Gamma_{ab} \overset{a.s.}{\to} \frac{p_{ab}}{\sqrt{p_a}\sqrt{p_b}}\Sigma_{ab}^{\mathrm{obs}}$.

We apply the law of total expectation for the second statement:

$$\mathbb{E}\left(\frac{1}{n}\sum_t Y'_{a,t}Y'_{b,t}\right) = \mathbb{E}_M\left(\mathbb{E}\left(\frac{1}{n}\sum_t Y'_{a,t}Y'_{b,t} \mid M\right)\right) = \mathbb{E}\left(\frac{n_{ab}}{\sqrt{n_a}\sqrt{n_b}}\Sigma^{\text{obs}}_{ab}\right) = \mathbb{E}\left(\frac{n_{ab}}{\sqrt{n_a}\sqrt{n_b}}\right)\Sigma^{\text{obs}}_{ab}.$$

Moreover, we have $0 \leq \frac{n_{ab}}{\sqrt{n_a}\sqrt{n_b}} \leq 1$ by cosine similarity, so the sequence is uniformly bounded. We combine this with the almost sure convergence above to invoke the dominated convergence theorem:

$$\mathbb{E}\left(\frac{n_{ab}}{\sqrt{n_a}\sqrt{n_b}}\right) \to \frac{p_{ab}}{\sqrt{p_a}\sqrt{p_b}}.$$

It follows that:

$$\mathbb{E}\left(\frac{n_{ab}}{\sqrt{n_a}\sqrt{n_b}}\right)\Sigma^{\text{obs}}_{ab} \to \frac{p_{ab}}{\sqrt{p_a}\sqrt{p_b}}\Sigma^{\text{obs}}_{ab}.$$

Hence, $\mathbb{E}\left(\frac{1}{n}\sum_t Y'_{a,t}Y'_{b,t}\right) \to \frac{p_{ab}}{\sqrt{p_a}\sqrt{p_b}}\Sigma^{\text{obs}}_{ab}$. □

**Theorem 2.** *Consider the estimator: $\widehat{S}_{ab} = \frac{1}{q}\sum_{j=1}^{q} z_{ja}z_{jb}$. Let $G_q$ index the variants that are independent of $Y$, and $B_q$ index the variants that are dependent on $Y$. Let $C_1, C'_1, C_2, L$ and $D$ denote finite constants. Assume:*

1. *(Bounded moments for $G_q$) For $j \in G_q$, we have $\mathbb{E}(z_{ja}z_{jb}|M) = \Gamma_{ab}$ and $\sup_{j\in G_q}\text{Var}(z_{ja}z_{jb}|M) \leq C_1$.*
2. *(Bounded moments for $B_q$) For $j \in B_q$, we have $\mathbb{E}(z_{ja}z_{jb}|M) = \mu_j$ with $\sup_{j\in B_q}|\mu_j - \Gamma_{ab}| \leq L$ and $\sup_{j\in B_q}\text{Var}(z_{ja}z_{jb}|M) \leq C'_1$.*
3. *(Local LD) There is an undirected dependency graph with edge set $E_q$ and maximum degree $D$ such that $|\text{cov}(z_{ja}z_{jb}, z_{ka}z_{kb}|M)| \leq C_2$ when $(j,k) \in E_q$ and $\text{cov}(z_{ja}z_{jb}, z_{ka}z_{kb}|M) = 0$ when $(j,k) \notin E_q$.*
4. *(Vanishing contamination) We have $\frac{|B_q|}{q} \to 0$ as $q \to \infty$.*

*Then, $\widehat{S}_{ab} \xrightarrow{p} \Gamma_{ab}$ conditional on $M$ as $q \to \infty$.*

*Proof* Let $X_j = z_{ja}z_{jb}$. Consider the numerator of $\widehat{S}_{ab}$ and write:

$$\sum_{j=1}^{q}(X_j - \Gamma_{ab}) = \underbrace{\sum_{j\in G_q}(X_j - \Gamma_{ab})}_{N_G} + \underbrace{\sum_{j\in B_q}(X_j - \Gamma_{ab})}_{N_B}.$$

We have $\mathbb{E}(N_G|M) = 0$ by Assumption 1. We can also write:

$$|\mathbb{E}(N_B|M)| = |\sum_{j\in B_q}(\mu_j - \Gamma_{ab})| \leq L|B_q|.$$

The inequality follows by Assumption 2. Thus $|\mathbb{E}(N_B|M)|/q \leq L|B_q|/q \to 0$ by Assumption 4. We proved that $\mathbb{E}(\widehat{S}_{ab}|M) \to \Gamma_{ab}$.

With Assumption 3, we can write:

$$\begin{aligned}\text{Var}(\widehat{S}_{ab}|M) &= \sum_j \frac{1}{q^2}\text{Var}(X_j|M) + \sum_{(j,k)\in E_q}\frac{1}{q^2}\text{cov}(X_j, X_k|M)\\ &\leq \frac{\max(C_1, C'_1)q}{q^2} + \frac{C_2 Dq}{q^2}.\end{aligned}$$

The first and second ratios converge to zero. Hence, $\text{Var}(\widehat{S}_{ab}|M) \to 0$.

Let $m = \mathbb{E}(\widehat{S}_{ab}|M)$. We invoke the triangle inequality to write:

$$\mathbb{P}_{|M}(|\widehat{S}_{ab} - \Gamma_{ab}| > \varepsilon) = \mathbb{P}_{|M}(|(\widehat{S}_{ab} - m) + (m - \Gamma_{ab})| > \varepsilon) \leq \mathbb{P}_{|M}(|\widehat{S}_{ab} - m| > \varepsilon - |m - \Gamma_{ab}|),$$

where we have moved the conditioning notation to the subscripts to avoid confusion with absolute value notation. We previously showed that $|m - \Gamma_{ab}| \to 0$. Now apply Chebyshev's inequality when $|m - \Gamma_{ab}| < \varepsilon$:

$$\mathbb{P}_{|M}(|\widehat{S}_{ab} - \Gamma_{ab}| > \varepsilon) \leq \mathbb{P}_{|M}(|\widehat{S}_{ab} - m| > \varepsilon - |m - \Gamma_{ab}|) \leq \frac{\text{Var}(\widehat{S}_{ab}|M)}{(\varepsilon - |m - \Gamma_{ab}|)^2}.$$

We also already showed that $\text{Var}(\widehat{S}_{ab}|M) \to 0$, so $\widehat{S}_{ab} \xrightarrow{p} \Gamma_{ab}$ as $q \to \infty$. □

**Proposition 4.** *We have $\widehat{\alpha}'_{ja} = \sqrt{n}\,\widehat{\text{cor}}(V_j, R'_a)$ for every $Y_a \in Y$.*

*Proof* Choose any $Y_a \in Y$. Consider regressing $Y'_a$ on $Y'_{-a}$:

$$\beta_a = (\Gamma_{-a,-a})^{-1}\Gamma_{-a,a} = -\frac{\Omega_{-a,a}}{\Omega_{aa}}.$$

We thus have:

$$R'_a = Y'_a - Y'_{-a}\beta_a = Y'w_a,$$

where $w_a = \frac{\Omega e_a}{\Omega_{aa}}$ and $e_a$ denotes the $a^{\text{th}}$ canonical basis vector. We also have $\text{Var}(R'_a) = 1/\Omega_{aa}$. Now consider:

$$\widehat{\text{cov}}(V_j, R'_a) = \widehat{\text{cov}}(V_j, Y')w_a = \widehat{r}'_j\frac{\Omega e_a}{\Omega_{aa}} = \frac{(\widehat{r}'_j\Omega)_a}{\Omega_{aa}}.$$

Since $\text{Var}(V_j) = 1$ and $\text{Var}(R'_a) = 1/\Omega_{aa}$, we have:

$$\widehat{\text{cor}}(V_j, R'_a) = \frac{(\widehat{r}'_j\Omega)_a}{\sqrt{\Omega_{aa}}}.$$

We previously showed that $\widehat{r}'_j = \widehat{\text{cov}}(V_j, Y') = \frac{z_j}{\sqrt{n}}$. It thus follows that:

$$\widehat{\text{cor}}(V_j, R'_a) = \frac{(z_j\Omega)_a}{\sqrt{\Omega_{aa}}\sqrt{n}} = \frac{(\Omega z_j^\top)_a}{\sqrt{\Omega_{aa}}\sqrt{n}} = \frac{\alpha_{ja}}{\sqrt{\Omega_{aa}}\sqrt{n}} = \frac{\alpha'_{ja}}{\sqrt{n}}.$$

We conclude that $\alpha'_{ja} = \sqrt{n}\widehat{\text{cor}}(V_j, R'_a)$ for every $Y_a \in Y$ because we chose $Y_a$ arbitrarily. $\square$

## D. Additional Experimental Results

### Simulations

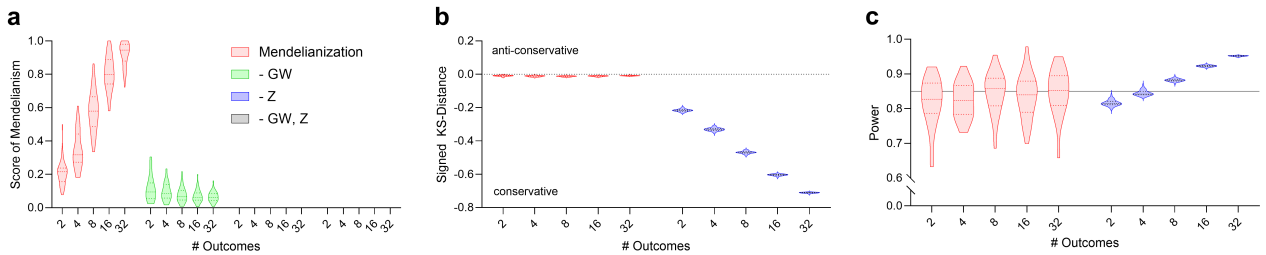


**Supplementary Fig. 1. Ablation results.** "-GW" denotes not applying each learned outcome genome-wide, "-Z" denotes estimating $\Gamma$ from $\beta$ coefficients instead of $z$-statistics, and "-GW, Z" denotes using both modifications. (a) Misapplying the CCA-derived outcomes (-GW) or using $\beta$ coefficients in place of $z$-statistics (-Z) reduced the SoM to essentially zero; for two of the ablations, the SoM violin plots collapse at the bottom and are therefore not visible. (b) Using $\beta$ coefficients produced highly conservative $p$-values. (c) The corresponding $p$-value histogram was strongly U-shaped, indicating severe miscalibration and yielding highly unstable power across outcome dimensions. Despite these deficiencies, the -GW and -Z variants still achieved near-perfect true positive rates and completed in under one second, comparable to Mendelianization.

### Depression and Generalized Anxiety

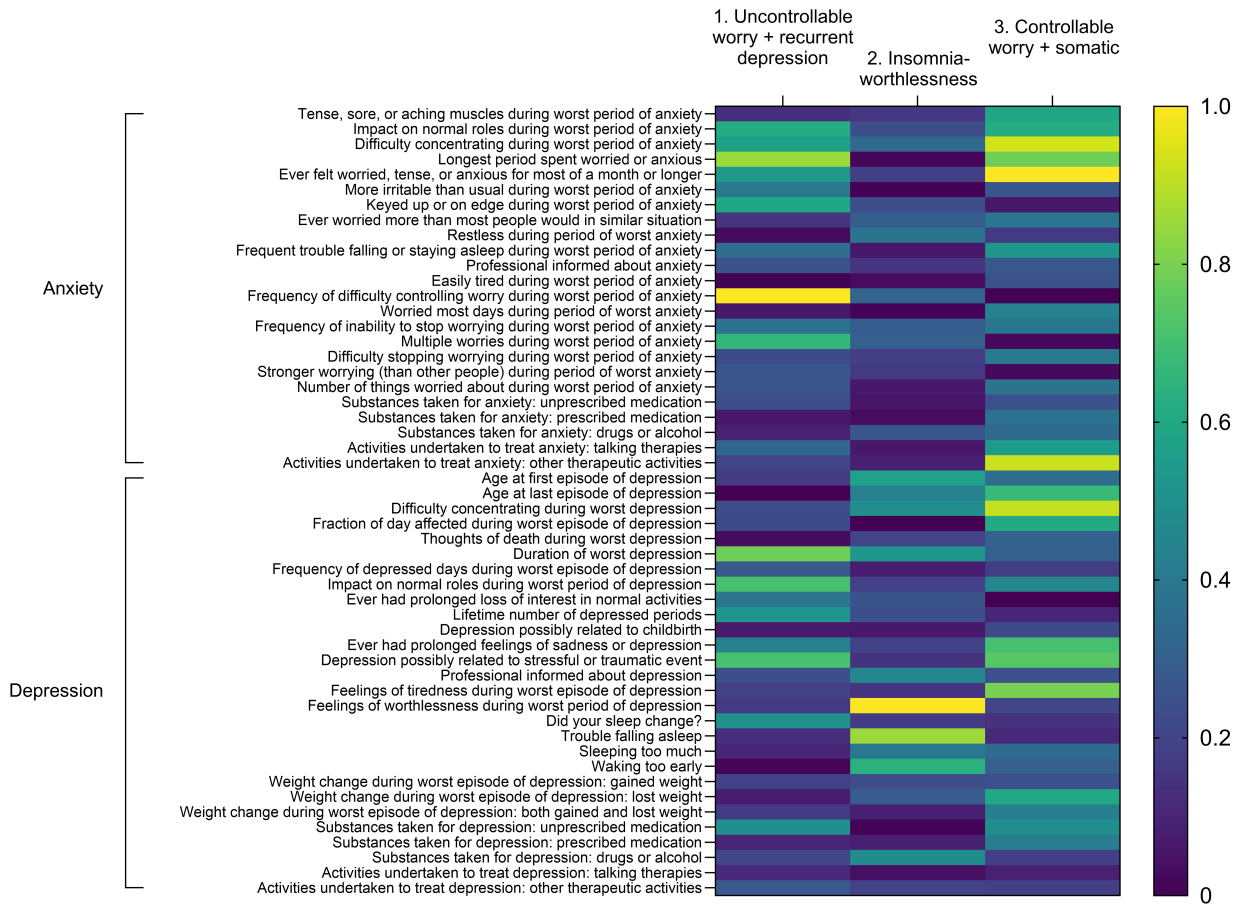


**Supplementary Fig. 2. Heatmap of $|\alpha'|$ values for the lead variants in depression and anxiety.** Each column displays the magnitudes of the interpretable coefficients for a lead variant, normalized to the range $[0, 1]$. The coefficient patterns resolve three phenotypic profiles: (i) uncontrollable worry with recurrent depression; (ii) melancholic-like depression characterized by insomnia, worthlessness, and cognitive slowing; and (iii) largely controllable worry accompanied by prominent somatic complaints. These profiles illustrate distinct modes of experiencing depression and anxiety, underscoring the disorders' heterogeneity.

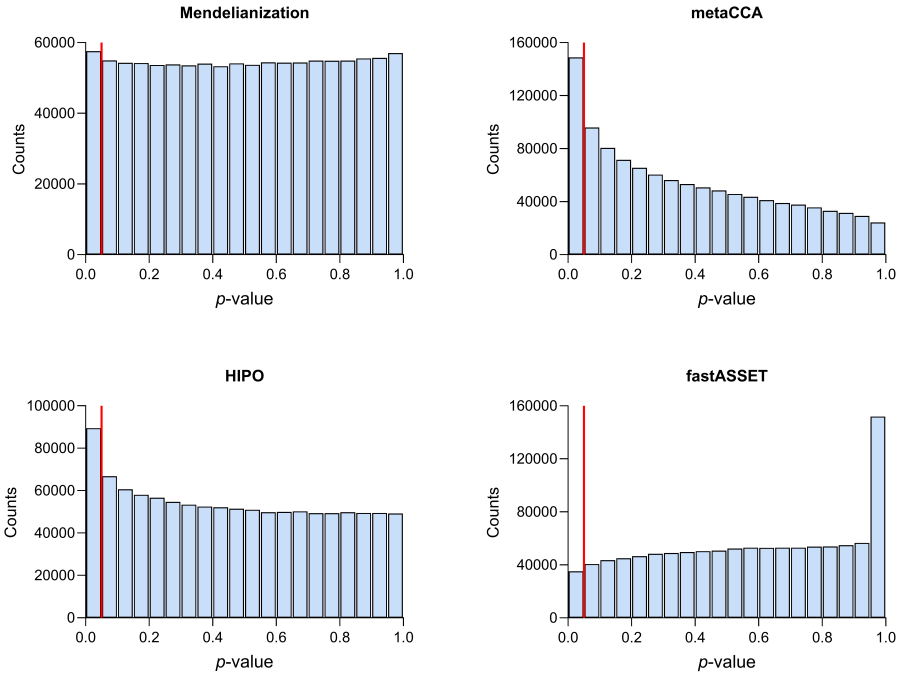


**Supplementary Fig. 3. Histograms of $p$-values across all variants.** Mendelianization was the only method to produce a near-uniform distribution for variants associated with $p$-values above 0.05 (red line). By contrast, metaCCA and HIPO yielded anti-conservative distributions with shrunken $p$-values, whereas fastASSET was overly conservative, producing $p$-values near one for many variants.

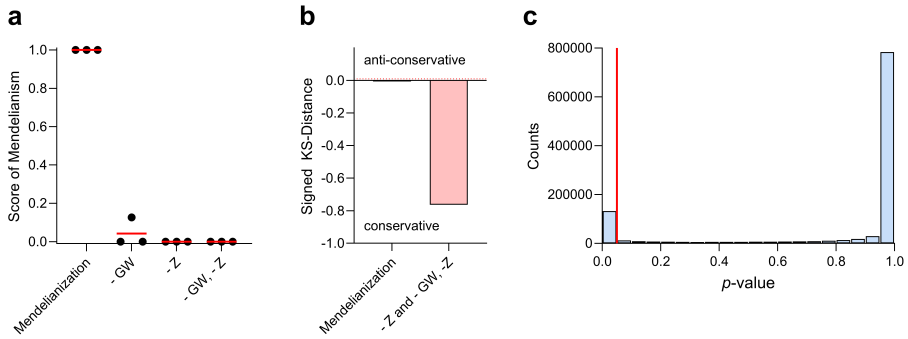


**Supplementary Fig. 4. Ablation results.** (a) Misapplying the learned outcomes or using $\beta$ coefficients in place of $z$-statistics (or both) decreased the SoM score to near zero. (b) The histogram obtained by $\beta$ coefficients was very conservative. (c) Inspection of the histogram derived from the $\beta$ coefficient estimates revealed highly miscalibrated $p$-values following a U-shape curve.

## Alcohol Use Disorder

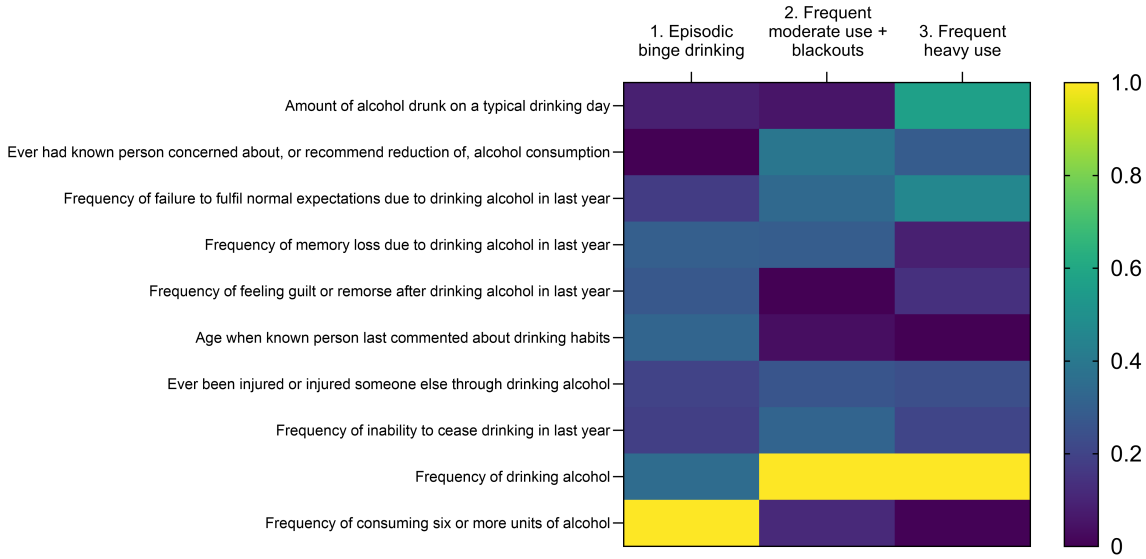


**Supplementary Fig. 5. Heatmap of $|\alpha'|$ values for the lead variants in alcohol use disorder.** The coefficient patterns correspond to three phenotypic profiles: (i) episodic binge drinking, (ii) frequent moderate use of alcohol with blackouts, and (iii) frequent heavy use without blackouts.

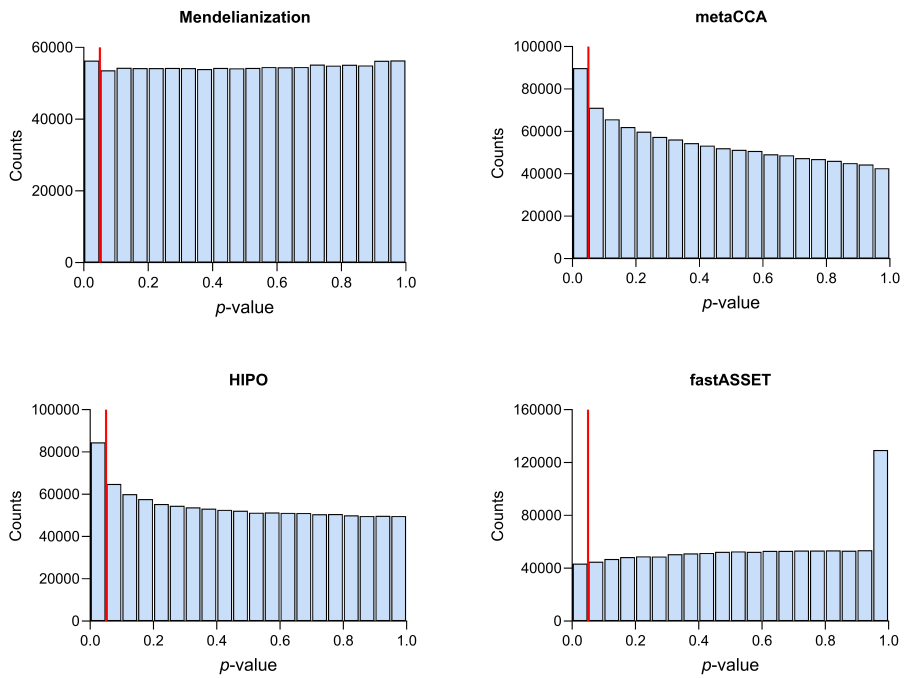


**Supplementary Fig. 6. Histograms of $p$-values across all variants.** Mendelianization was again the only method to produce a near-uniform distribution on $[0.05, 1]$.

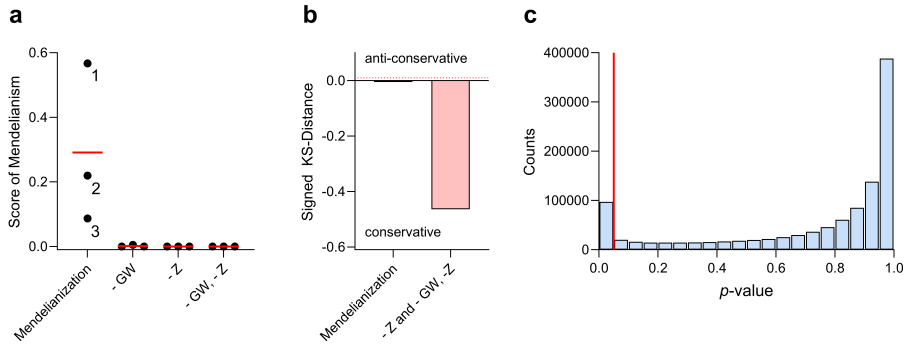


**Supplementary Fig. 7. Ablation results.** (a) Mendelianization again achieved the highest SoMs, and (b, c) estimating $\Gamma$ from $\beta$ coefficients yielded highly miscalibrated $p$-values.