ESTIMATING TOTAL MEDIATION EFFECTS OF HIGH-DIMENSIONAL OMICS MEDIATORS IN CASE-CONTROL STUDIES

BY ZHIYU KANG[1], LI CHEN[2], PENG WEI[3], ZHICHAO XU[3], CHUNLIN LI[4,5,b], AND TIANZHONG YANG[1,a]

[1]*Division of Biostatistics and Health Data Science, University of Minnesota, Twin Cities,* [a]*yang3704@umn.edu*

[2]*School of Statistics, University of Minnesota, Twin Cities*

[3]*Department of Biostatistics, University of Texas MD Anderson Cancer Center*

[4]*Department of Statistics, University of Virginia,* [b]*chunlin@virginia.edu*

[5]*Department of Statistics, Iowa State University*

Mediation analysis helps uncover how exposures impact outcomes through intermediate variables. Traditional mean-based total mediation effect measures may suffer from the cancellation of opposite component-wise effects, and existing methods often lack the power to capture weak effects in high-dimensional mediators. Additionally, most high-dimensional mediation analysis methods have focused on continuous outcomes, with limited attention to binary outcomes, particularly in case-control studies. To fill this gap, we propose an R^2 total mediation effect measure within the liability framework that offers a clear and intuitive causal interpretation, provides additional insights beyond the mean-based measures, and is invariant to disease prevalence. We develop a cross-fitted, modified Haseman-Elston regression-based estimation procedure tailored for mediation analysis in case-control studies, which can also be applied to cohort studies. Our estimator remains consistent in the presence of non-mediators and weak effects, as demonstrated in extensive simulations. Theoretical justification for consistency is provided under mild conditions and without requiring exact mediator selection. In a case-control substudy of the Women's Health Initiative involving 2150 individuals, we found that many metabolites were mediators with weak effects in the path from BMI to coronary heart disease, and we estimated that 89% (95% CI: 57%–100%) of the BMI-explained variation in underlying CHD liability is mediated by the measured metabolomics. The proposed estimation procedure is implemented in the R package "r2MedCausal", available on GitHub.

1. Introduction. Mediation analysis is a powerful tool for exploring how an exposure affects an outcome through one or more intermediate variables, known as mediators. With recent advances in high-throughput technologies, high-dimensional mediation analysis of biological data has emerged as an area of scientific interest. In this paper, we focus on estimating the total mediation effect, which provides critical insights for downstream analyses. Specifically, our method is motivated by the Women's Health Initiative (WHI) nested case-control substudy, where we are interested in examining how obesity (quantitatively measured by the body mass index, BMI) influences the risk of incident coronary heart disease (CHD) through metabolomic profiles in postmenopausal women. Obesity (often defined as BMI $\geq 30kg/m^2$) is highly prevalent among adults in the US and a well-known risk factor for CHD (Eckel and Krauss, 1998). Higher BMI has been found to be associated with CHD by large epidemiological studies (Flint et al., 2010; Lyall et al., 2017; Global Burden of Metabolic Risk Factors for

Keywords and phrases: Causal mediation analysis, High-dimensional mediators, Omics mediators, Selection bias, Total mediation effect.

Chronic Diseases Collaboration et al., 2014) and causal analyses (Lyall et al., 2017; Silva, Fatumo and Nitsch, 2024). A recent randomized, double-blind, placebo-controlled clinical trial further corroborated this causal link (Packer et al., 2025), showing that a GLP-1 drug significantly reduced the risk of cardiovascular mortality and heart failure progression. With more and more weight-loss intervention strategies, understanding how BMI influences CHD risk through molecular phenotypes such as metabolomics can offer valuable insights.

Despite its importance, the mediation analysis of the WHI dataset poses a two-fold challenge. First, the existing methods fail to adequately capture the overall mediation effects, owing to the presence of many weak mediators, as suggested by our preliminary analysis in Section 6.1. A plausible reason is that many high-dimensional approaches rely on sparsity assumptions, which are not suitable for handling weak mediators with small and non-sparse effect sizes (Zhang et al., 2016; Gao et al., 2019). While some methods overcome this issue by considering transformed mediators, thus relaxing sparsity, they sacrifice interpretability (Huang and Pan, 2015). This highlights the need for a more robust, interpretable approach tailored to high-dimensional mediation analysis with weak mediators.

Second, existing mediation effect measures remain subject to ongoing improvement, especially for binary outcomes and the case-control design, despite their widespread relevance in the area (Clark-Boucher et al., 2023; Rijnhart et al., 2021). Indeed, quantifying the total mediation effect is nontrivial for high-dimensional settings. Most of the existing measures are variants of the mean-based product-of-coefficient (POC) measures, which are prone to cancellation in the presence of bidirectional mediation effects, a phenomenon that frequently occurs with omics mediators (Song et al., 2019; Yang et al., 2021). To address this issue, Yang et al. (2021) extended the variance-based R^2 measure originally proposed by Fairchild et al. (2009) to accommodate multiple and high-dimensional mediators under a mixed model framework. This method captures the non-zero total mediation effect by quantifying the variance in the continuous outcome that is attributable to both the exposure and the mediators, attempting to address bidirectional effects issues in high-dimensional mediator settings. Subsequently, this measure was extended to binary outcomes using a likelihood-based McFadden's R^2 for its stability and relative independence from disease prevalence (Chi et al., 2026). However, the McFadden's R^2 defined for binary outcomes appears to be less sensitive to disease prevalence but still varies with it. Additionally, its interpretability is limited, as it is more closely linked to measuring goodness-of-fit than to explaining variance. In another line of work, Song et al. (2019) proposed a L_2 norm measure of the mediation effect while simply treating binary outcomes as continuous. However, both the use of the L_2 norm and the treatment of binary outcomes as continuous lack theoretical support. Importantly, all of these measures lack rigorous causal interpretations. Thus, it underscores the need for developing a mediation measure that enables a unified, robust approach to mediation analysis for binary disease outcomes, across different disease prevalence and study designs.

In response to these challenges, we propose a novel R^2 mediation effect measure that provides a causal interpretation for binary outcomes. The novel R^2 mediation effect measure is defined on a latent variable scale, i.e., liability, which is invariant to disease prevalence as opposed to POC measures for binary outcomes, and maintains its connection with the classic R^2 defined for continuous outcomes. Thus, it offers a unified measure and a robust approach for mediation analysis for diseases with different prevalences. Unlike the traditional logistic model, the liability model requires ascertainment bias correction in a case-control study. To this end, we propose a two-stage estimation procedure to account for ascertainment bias in case-control studies and to address non-mediators and weak effects. We modify a heritability estimation method, phenotype-correlation genotype-correlation (PCGC) regression (Golan, Lander and Rosset, 2014), a generalization of Haseman-Elston regression, to estimate consistent variance components under case-control studies. It is worth noting that, although our procedure is designed for case-control studies, it is also applicable to cohort studies, thus unifying the data analysis across study designs. Extensive simulation studies showcased the robust

performance of our proposed method under varying disease prevalence, study design, weak effects, and model misspecification, consistently outperforming existing high-dimensional mediation methods. Furthermore, under mild conditions and without requiring exact mediator selection, we establish the consistency of the estimators. Applied to the motivating WHI case-control dataset, our method effectively captures a good amount of mediation effect of metabolites from BMI to CHD.

The rest of the paper is organized as follows. Section 2 reviews the existing mediation effect measures and highlights their limitations. Section 3 introduces our proposed $R^2_{\text{med;causal}}$ mediation measure and discusses its causal interpretation. Section 4 details the estimation procedure and addresses implementation considerations specific to case-control designs. Section 5 evaluates the performance of our method through extensive simulation studies, comparing it with existing approaches. Section 6 presents the application to WHI. Section 7 summarizes the contributions and limitations of our work and outlines future research directions. The theoretical proofs are provided in Section 1 of the Supplementary Material (Kang et al., 2025). Finally, an R package, "r2MedCausal", has been developed to implement the proposed estimation procedure.

2. Preliminaries. This section provides an overview of existing approaches and their limitations, which serve as the basis for the proposed approach.

A foundational framework in mediation analysis is established by Baron and Kenny (1986), who proposed the following structural equations for a continuous outcome, describing the relationships among an exposure X, a p-dimensional vector of mediators $\boldsymbol{M}$, and an outcome Y:

$$\boldsymbol{M} = \boldsymbol{\alpha}^{(c)} X + \boldsymbol{\Psi}^{(c)} \boldsymbol{C} + \boldsymbol{\xi}, \tag{1}$$

$$Y = \gamma^{(c)} X + \boldsymbol{\beta}^{(c)\top} \boldsymbol{M} + \boldsymbol{\theta}^{(c)\top} \boldsymbol{C} + \epsilon, \tag{2}$$

where $\boldsymbol{C}$ is a $q \times 1$ vector of covariates. The superscript $^{(c)}$ decorates coefficients corresponding to the continuous outcome setting. The coefficient $\boldsymbol{\alpha}^{(c)} = (\alpha_1^{(c)}, \alpha_2^{(c)}, \ldots, \alpha_p^{(c)})^\top$ captures the exposure-mediator relationship after adjusting for the covariates. Similarly, $\boldsymbol{\beta}^{(c)} = (\beta_1^{(c)}, \beta_2^{(c)}, \ldots, \beta_p^{(c)})^\top$ captures the mediator-outcome relationship after adjusting for both exposure and covariates. The parameter $\gamma^{(c)}$ is the direct effect of the exposure variable, and $\boldsymbol{\Psi}^{(c)}$ and $\boldsymbol{\theta}^{(c)}$ are the coefficients of covariates. The variables ϵ, $\boldsymbol{\xi} = (\xi_1, \xi_2, \ldots, \xi_p)^\top$ are error terms. In model (1)–(2), the total effect of the exposure X on the continuous outcome Y is decomposed into a direct effect and a mediation effect. The total effect $\boldsymbol{\alpha}^{(c)\top}\boldsymbol{\beta}^{(c)} + \gamma^{(c)}$ is typically estimated by regressing Y on X, after adjusting for covariates. The mediation effect can be estimated using either the difference-in-coefficients method, which subtracts the direct effect $\gamma^{(c)}$ from the total effect, or the POC approach, which computes the inner product $\boldsymbol{\alpha}^{(c)\top}\boldsymbol{\beta}^{(c)}$.

Most early development of mediation analysis was not formulated in a rigorous causal framework (Pearl, 2001; VanderWeele and Vansteelandt, 2010). In causal mediation analysis (with a continuous outcome), the natural indirect effect (NIE) and the natural direct effect (NDE) for an exposure contrast between x and x' are defined as:

$$\text{NIE} = \mathbb{E}[Y(x, M(x))] - \mathbb{E}[Y(x, M(x'))], \quad \text{NDE} = \mathbb{E}[Y(x, M(x'))] - \mathbb{E}[Y(x', M(x'))],$$

where $Y(x, M(x))$ is the potential outcome when both exposure and mediator take their natural values under $X = x$, $Y(x, M(x'))$ is the outcome when the exposure is set to x, but the mediator is fixed at the level it would take under $X = x'$, and $Y(x', M(x'))$ represents the outcome when both exposure and mediator take their natural values under $X = x'$. To identify NIE and NDE, the following assumptions are often imposed (VanderWeele, 2016).

ASSUMPTION 1 (No Unmeasured Confounding).
A1 No unmeasured confounding of the exposure-mediator effect.
A2 No unmeasured confounding of the mediator-outcome effect.
A3 No unmeasured confounding of the exposure-outcome effect.
A4 No mediator-outcome confounder that is itself affected by the exposure.

If Assumption 1 holds in model (1)–(2), then we have NIE $= \boldsymbol{\alpha}^{(c)\top}\boldsymbol{\beta}^{(c)}(x - x')$, which provides POC with a rigorous causal interpretation.

2.1. *The issues of POC and its alternatives.* Partly due to its causal interpretation, the POC measure and its variants, such as the proportion-mediated (PM) measure $\boldsymbol{\alpha}^\top\boldsymbol{\beta}/(\gamma + \boldsymbol{\alpha}^\top\boldsymbol{\beta})$ (MacKinnon et al., 2007), have been popular, including in high-dimensional mediation analysis of biological data (Zhang et al., 2016; Gao et al., 2019). However, only reporting POC-based measures can be misleading in high-dimensional settings where numerous bidirectional component-wise mediation effects (i.e., $\alpha_i\beta_i$ have different signs) are present (Song et al., 2019), potentially causing cancellation in the POC and PM measures. Therefore, a small (or near-zero) POC may not imply the absence of mediation, but rather that positive and negative indirect pathways offset in aggregate.

Recognizing these drawbacks, Song et al. (2019) proposes an alternative L_2 norm for summarizing the overall mediation effects: $\mathrm{IE}_{\mathrm{BAMA}} = \sum_{i=1}^{p}(\alpha_i\beta_i)^2$. While this modification no longer suffers from cancellation of bidirectional component-wise mediation effects, it lacks rigorous theoretical support for the causal interpretation. An alternative mediation effect measure is the R^2 measure, which was first proposed in a single-mediator model for continuous outcomes by Fairchild et al. (2009), and later extended by Yang et al. (2021) to high-dimensional mediation analysis:

$$R^2_{\mathrm{med}} = R^2_{Y,X} + R^2_{Y,\boldsymbol{M}} - R^2_{Y,\boldsymbol{M}X},$$

where $R^2_{Y,X} = \frac{\mathrm{var}(Y)-\mathrm{var}(Y|X)}{\mathrm{var}(Y)}$, $R^2_{Y,\boldsymbol{M}} = \frac{\mathrm{var}(Y)-\mathrm{var}(Y|\boldsymbol{M})}{\mathrm{var}(Y)}$, and $R^2_{Y,\boldsymbol{M}X} = \frac{\mathrm{var}(Y)-\mathrm{var}(Y|X,\boldsymbol{M})}{\mathrm{var}(Y)}$. The R^2 measure provides a robust, sensible variance-based measure for continuous outcomes. However, it does not have a clear causal interpretation.

2.2. *The lack of a unified, robust measure for binary outcomes.* The mediation framework is extended for binary outcomes. This is often done by replacing (2) with a logistic regression (MacKinnon et al., 2007). In this situation, the difference-in-coefficients method may encounter the noncollapsibility problem (Greenland, Pearl and Robins, 1999), making the POC approach preferable (VanderWeele, 2016). When Assumption 1 is satisfied and the outcome is rare, logistic regression can be used to estimate the causal mediation effect on an odds ratio scale as: $\mathrm{OR}^{\mathrm{NIE}} \approx \exp(\boldsymbol{\alpha}^\top\boldsymbol{\beta})$ (VanderWeele and Vansteelandt, 2010). However, when the outcome is not rare, odds ratios suffer from noncollapsibility, meaning that $\exp(\boldsymbol{\alpha}^\top\boldsymbol{\beta})$ does not provide an interpretable mediation effect estimate, as it depends on covariates adjustment. In such cases, log-binomial models are common to estimate the causal mediation effect on a risk ratio (RR) scale, given by $\mathrm{RR}^{\mathrm{NIE}} \approx \exp(\boldsymbol{\alpha}^\top\boldsymbol{\beta})$ (VanderWeele, 2016). Consequently, the preferred mediation effect scale varies by outcome prevalence, with odds ratios for rare outcomes and risk ratios for common ones. Moreover, the foregoing approaches are fundamentally based on the POC calculation, thus suffering from the same cancellation problem as described for continuous outcomes earlier.

Motivated by our analysis of WHI data, we aim to develop a new measure for mediation analysis of binary outcomes, which enjoys the following merits: (i) it has a causal interpretation; (ii) it overcomes the cancellation issue; and (iii) it offers a unified treatment for binary outcomes with different prevalence and study designs.

3. Modeling. For binary outcomes, the liability model provides a unified approach for analyzing both rare and common diseases. In addition, the variance-based mediation effect has a consistent interpretation for both continuous outcomes and binary outcomes under the liability model. It has been scarcely used in low-dimensional settings (MacKinnon et al., 2007; Nguyen et al., 2016) and recently adapted to high-dimensional settings (Gaynor, Schwartz and Lin, 2019; Sohn, Lu and Li, 2022). Thus, we utilize the liability-based mediation model, which is defined by the following structural equations:

$$\boldsymbol{M} = \boldsymbol{\alpha} X + \boldsymbol{\Psi} \boldsymbol{C} + \boldsymbol{\xi}, \tag{3}$$

$$l = \gamma X + \boldsymbol{\beta}^\top \boldsymbol{M} + \boldsymbol{\theta}^\top \boldsymbol{C} + \epsilon, \tag{4}$$

$$Y = \mathbf{1}(l > t), \tag{5}$$

where l is the underlying latent variable, Y is the observed binary outcome. The function $\mathbf{1}(\cdot)$ is an indicator function, with threshold t determining K, the prevalence of the trait in the population. The coefficients $\boldsymbol{\alpha}, \boldsymbol{\beta}, \gamma, \boldsymbol{\Psi}$, and $\boldsymbol{\theta}$ retain the same interpretations as in the continuous outcome model, now defined in relation to the latent liability. The error terms ϵ and $\boldsymbol{\xi}$ are assumed to be normally distributed.

In what follows, we perform analyses conditioned on $\boldsymbol{C}$ and treat $\boldsymbol{C}$ as fixed. Additionally, without loss of generality, X is assumed to be centered.

3.1. *Modeling non- and weak mediators.* In real-world data, the identities of true mediators are usually unknown, necessitating the inclusion of non-mediators in our mediation model. Potential mediators $\boldsymbol{M}$ are partitioned into true mediators $\boldsymbol{M}_{\mathcal{T}}$, and three types of non-mediators $(\boldsymbol{M}_{\mathcal{I}_1}, \boldsymbol{M}_{\mathcal{I}_2}, \boldsymbol{M}_{\mathcal{I}_3})$, as illustrated in Figure 1:

- $\boldsymbol{M}_{\mathcal{T}} = \{M_j | \alpha_j \neq 0 \text{ and } \beta_j \neq 0 \text{ for } j \in \mathcal{T}\}$;
- $\boldsymbol{M}_{\mathcal{I}_1} = \{M_j | \alpha_j = 0 \text{ and } \beta_j \neq 0 \text{ for } j \in \mathcal{I}_1\}$;
- $\boldsymbol{M}_{\mathcal{I}_2} = \{M_j | \alpha_j \neq 0 \text{ and } \beta_j = 0 \text{ for } j \in \mathcal{I}_2\}$;
- $\boldsymbol{M}_{\mathcal{I}_3} = \{M_j | \alpha_j = 0 \text{ and } \beta_j = 0 \text{ for } j \in \mathcal{I}_3\}$.

We denote the proportions of $\boldsymbol{M}_{\mathcal{T}}$, $\boldsymbol{M}_{\mathcal{I}_1}$, $\boldsymbol{M}_{\mathcal{I}_2}$, and $\boldsymbol{M}_{\mathcal{I}_3}$ as π_{11}, π_{01}, π_{10} and π_{00} respectively. The overall proportions of non-zero elements in $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are then given by $\pi_\alpha = \pi_{11} + \pi_{10}$ and $\pi_\beta = \pi_{11} + \pi_{01}$.

Existing approaches often assume the true mediators are sparse ($\pi_{11} \approx 0$) and try to identify mediators through either high-dimensional sparse linear regression (Zhang et al., 2016; Gao et al., 2019) or large-scale hypothesis testing (Huang and Pan, 2015). The success of selection roughly requires the signal strength $|\alpha_k| \wedge |\beta_k|$ to be larger than $\sqrt{\log(p)/n}$. However, such approaches may be plagued by the presence of dense *weak mediators* such that $|\alpha_k| \wedge |\beta_k| \lesssim \sqrt{\log(p)/n}$.

Moving beyond the sparse regime, we model the dense mediation effects based on the random-effects hypothesis (Dobriban and Wager, 2018). Heuristically, the results obtained under this random-effects hypothesis can be regarded as average-case analyses over dense parameters, offering a qualitatively different perspective for high-dimensional mediation analysis than the popular sparse regression approaches. Specifically, we consider a mixed-effect model under the working assumptions that

$$\alpha_k \sim \pi_\alpha N(0, \sigma_\alpha^2) + (1 - \pi_\alpha)\delta_0, \tag{6}$$

$$\beta_k \sim \pi_{11} N(0, \frac{\sigma_{11}^2}{p\pi_{11}}) + \pi_{01} N(0, \frac{\sigma_{01}^2}{p\pi_{01}}) + (1 - \pi_\beta)\delta_0, \quad k = 1, \ldots, p, \tag{7}$$

where δ_0 denotes a point mass at zero. The purpose of introducing the normalizing factor p^{-1} in (7) is to let the per-mediator contribution vanish and thus the aggregated mediation

variation remains finite (Dobriban and Wager, 2018). Consequently, it becomes impossible to accurately select mediators in this scenario, as we have $|\alpha_k| \wedge |\beta_k| \asymp 1/\sqrt{p} \ll 1/\sqrt{n}$; nevertheless, it remains feasible to estimate the aggregated mediation variation.

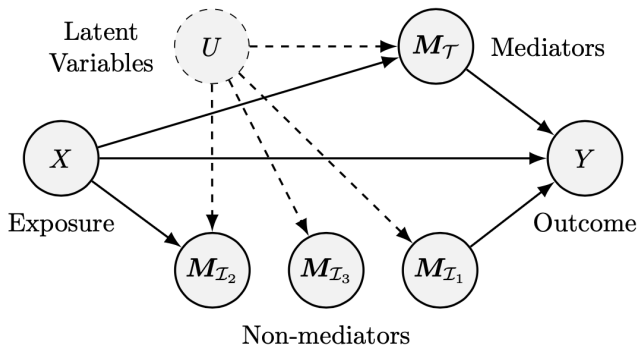

FIG 1. *Graph representation of a mediation model.*

3.2. *A novel R^2 measure with causal interpretation.* Aiming to achieve the robustness of rare disease assumption and ease of interpretation, we define a new R^2 measure for high-dimensional mediators under the liability scale for binary outcomes. Following Yang et al. (2021), under the structural equation model specified by (3)–(5), the R^2 measure would be defined as:

$$R^2_{\text{med}} = R^2_{l,X} + R^2_{l,\boldsymbol{M}} - R^2_{l,\boldsymbol{M}X}.$$

However, this definition of R^2_{med} does not inherently take into account the directional relationship between X and $\boldsymbol{M}$. Specifically, exchanging X and $\boldsymbol{M}$ would not change R^2_{med}, despite mediation analysis explicitly assuming directed causal pathways from X to Y and from X to $\boldsymbol{M}$ to Y. To resolve this issue and explicitly incorporate the causal direction, we introduce the do-operators in our R^2 measure within Pearl's Structural Causal Framework (Pearl, 2010). We formulate the new R^2 measure for binary outcomes as:

$$R^2_{\text{med;causal}} = R^2_{l,do(X)} + R^2_{l,do(\boldsymbol{M})} - R^2_{l,do(X,\boldsymbol{M})}. \tag{8}$$

Based on (3)–(5) and working model assumptions (6)–(7), we define the *causal coefficient of determination* as

$$R^2_{l,do(\boldsymbol{M})} = \frac{\text{var}_{\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{M}'}\left(\mathbb{E}\left(l \mid do(\boldsymbol{M}=\boldsymbol{M}')\right)\right)}{\text{var}_{\boldsymbol{\alpha},\boldsymbol{\beta},l}(l)}, \tag{9}$$

where $\boldsymbol{M}'$ is identically distributed as $\boldsymbol{M}$ but independent of all other variables in the model and the do-operator is a stochastic intervention of $\boldsymbol{M}$. The stochastic intervention $do(\boldsymbol{M}=\boldsymbol{M}')$ induces the so-called truncated factorization (Pearl, 2010), which removes all the incoming edges of $\boldsymbol{M}$ in Figure 1 and replace $\boldsymbol{M}$ with a new random vector $\boldsymbol{M}'$. Thus, $\mathbb{E}(l \mid do(\boldsymbol{M}=\boldsymbol{M}'))$ is the (counterfactual) mean of l conditioned on $\boldsymbol{M}'$ with respect to the truncation probability, had $\boldsymbol{M}$ been replaced by $\boldsymbol{M}'$. In causal mediation model (3)–(5), $\mathbb{E}(l \mid do(\boldsymbol{M}=\boldsymbol{M}')) = \boldsymbol{\beta}^\top\boldsymbol{M}' + \boldsymbol{\theta}^\top\boldsymbol{C} + \mathbb{E}(\gamma X + \epsilon) = \boldsymbol{\beta}^\top\boldsymbol{M}' + \boldsymbol{\theta}^\top\boldsymbol{C}$ (because we treat $\boldsymbol{C}$ as fixed and X is centered). As a result, $R^2_{l,do(\boldsymbol{M})} = \text{var}_{\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{M}'}(\boldsymbol{\beta}^\top\boldsymbol{M}')/\text{var}_{\boldsymbol{\alpha},\boldsymbol{\beta},l}(l) = \{\sigma^2_\alpha\sigma^2_{11}\,\text{var}(X) + (\sigma^2_{11}+\sigma^2_{01})p^{-1}\sum_{k=1}^p \text{var}(\xi_k)\}/\,\text{var}_{\boldsymbol{\alpha},\boldsymbol{\beta},l}(l)$. By introducing the stochastic

intervention, it ensures that $R^2_{l,do(\boldsymbol{M})}$ reflects the causal effect of $\boldsymbol{M}$ on l, in contrast to $R^2_{l,\boldsymbol{M}}$, which may be confounded by X. Similarly, we have

$$\begin{aligned} R^2_{l,do(X)} &= \frac{\operatorname{var}_{\boldsymbol{\alpha},\boldsymbol{\beta},X'}\left(\mathbb{E}\left(l \mid do(X=X')\right)\right)}{\operatorname{var}_{\boldsymbol{\alpha},\boldsymbol{\beta},l}(l)} \\ &= \frac{(\sigma_\alpha^2\sigma_{11}^2+\gamma^2)\operatorname{var}(X)}{\operatorname{var}_{\boldsymbol{\alpha},\boldsymbol{\beta},l}(l)}, \\ R^2_{l,do(X,\boldsymbol{M})} &= \frac{\operatorname{var}_{\boldsymbol{\alpha},\boldsymbol{\beta},X'',\boldsymbol{M}''}\left(\mathbb{E}\left(l \mid do((X,\boldsymbol{M})=(X'',\boldsymbol{M}''))\right)\right)}{\operatorname{var}_{\boldsymbol{\alpha},\boldsymbol{\beta},l}(l)} \\ &= \frac{(\sigma_{11}^2\sigma_\alpha^2+\gamma^2)\operatorname{var}(X)+(\sigma_{11}^2+\sigma_{01}^2)p^{-1}\sum_{k=1}^p \operatorname{var}(\xi_k)}{\operatorname{var}_{\boldsymbol{\alpha},\boldsymbol{\beta},l}(l)}, \end{aligned}$$

where X' and $(X'',\boldsymbol{M}'')$ are independent copies of X and $(X,\boldsymbol{M})$, respectively, drawn from the same distribution but independent of other variables in the model. Thus, we have

$$R^2_{\text{med;causal}} = \sigma_\alpha^2\sigma_{11}^2 \operatorname{var}(X)/\operatorname{var}_{\boldsymbol{\alpha},\boldsymbol{\beta},l}(l). \tag{10}$$

The proposed $R^2_{\text{med;causal}}$ can be interpreted as the proportion of variation in the underlying liability for the outcome that is *causally* explained by the exposure through the mediators.

As discussed in MacKinnon et al. (2007), the PM measure in traditional mediation analysis is given by $\boldsymbol{\alpha}^\top\boldsymbol{\beta}/(\gamma+\boldsymbol{\alpha}^\top\boldsymbol{\beta})$. In the causal mediation analysis framework, assuming a rare binary outcome, the PM measure can be expressed as $\text{OR}^{\text{NDE}}(\text{OR}^{\text{NIE}}-1)/(\text{OR}^{\text{NDE}}\text{OR}^{\text{NIE}}-1)$, where $\text{OR}^{\text{NDE}}=\exp(\gamma)$ and $\text{OR}^{\text{NIE}}=\exp(\boldsymbol{\alpha}^\top\boldsymbol{\beta})$ denote the NDE and NIE on the odds ratio scale (VanderWeele, 2016). However, neither measure is guaranteed to range from 0 to 1. Additionally, both rely on the sum of component-wise effects, suffering from the cancellation problem when effects are in different directions. As an alternative to the PM measure, we propose a new relative Q^2_{med} measure, which is bounded between 0 and 1:

$$Q^2_{\text{med}} = \frac{R^2_{\text{med;causal}}}{R^2_{\text{med;causal}}+\gamma^2\operatorname{var}(X)/\operatorname{var}(l)} = \frac{\sigma_\alpha^2\sigma_{11}^2}{\sigma_\alpha^2\sigma_{11}^2+\gamma^2}. \tag{11}$$

This quantity is the proportion of variation of liability causally explained by X indirectly through the mediators over the sum of the proportion of the variation explained indirectly and directly by X. It can be interpreted as the proportion of the exposure explained variance on the liability scale of the outcome that is accounted for by the mediators.

REMARK 1 (Relationship with causal entropy). The proposed causal coefficient of determination is connected to the existing concept of causal entropy (Simoes, Dastani and van Ommen, 2024). It can be shown that for a single mediator,

$$R^2_{l,do(\boldsymbol{M})} = \mathbb{E}_{m_0\sim P_{\boldsymbol{M}'}}\left[\frac{\operatorname{var}(l)-\operatorname{var}(l|do(\boldsymbol{M}=m_0))}{\operatorname{var}(l)}\right], \tag{12}$$

where $do(\boldsymbol{M})=m_0$ implies that we intervene $\boldsymbol{M}$ and set it to be m_0, ignoring the dependencies between M and its parent variables. This notation explicitly accounts for causality (direction) in our formula. It can be shown that (9) and (12) are equivalent for a single mediator. The connection between our definition and causal entropy can be seen directly through (12).

4. Estimation procedure. To estimate the proposed R^2 measure, we adapted the cross-fitting framework as described by Xu et al. (2024) to improve the robustness of mediation effect estimation. Cross-fitting mitigates bias introduced by the winner's curse during the variable screening step (Göring, Terwilliger and Blangero, 2001). The approach involves a two-stage procedure. First, the sample is randomly split into two subsamples $\mathcal{D}^{(1)}$ and $\mathcal{D}^{(2)}$. In Step 1, mediation selection and variance estimation of $\boldsymbol{\alpha}$ are performed on subsample $\mathcal{D}^{(1)}$ to get the filtered mediator index set $\hat{\mathcal{T}}^{(1)}$ and variance estimate $\hat{\sigma}_{\alpha}^{(1)2}$. In Step 2, the estimation of other parameters is conducted on subsample $\mathcal{D}^{(2)}$ based on the selected mediators $\hat{\mathcal{T}}^{(1)}$. The roles of the two subsamples are then reversed, allowing the whole sample to be used in the estimation process. Finally, the R^2 estimate is given by the mean of cross-fitted R^2 estimates $\hat{R}^{(1)2}_{\text{med;causal}}$ and $\hat{R}^{(2)2}_{\text{med;causal}}$. The procedure is visualized in Section 4 of the Supplementary Material (Kang et al., 2025), and the mediator selection and parameter estimation process are detailed in Subsections 4.1-4.3.

4.1. *Estimation of σ_{α}^2 and mediator selection.* The estimation of $\boldsymbol{\alpha}$ in our mediation analysis framework has been extensively studied as a secondary outcome analysis problem in case-control studies (Schifano, 2019). To estimate σ_{α}^2, we used the inverse probability weighting (IPW) method introduced by Horvitz and Thompson (1952). In this approach, each individual's contribution to the estimating equations is weighted by the reciprocal of their selection probability, effectively constructing a pseudo-sample that is more representative of the target population than the original case-control sample.

Let S_i denote the inclusion indicator for the i-th individual, and define the inverse weight as $w_i = 1/P(S_i = 1|Y_i, \boldsymbol{M}_i, X_i, \boldsymbol{C}_i)$. The IPW estimating equation for estimating each element of $\boldsymbol{\alpha}$ is given by (Sofer et al., 2017):

$$\sum_{i=1}^{n} U_{\text{IPW},i}(\alpha_k) = \sum_{i=1}^{n} S_i w_i h(X_i)[\boldsymbol{M}_{i,k} - \alpha_k X_i - \boldsymbol{\Psi}_k \boldsymbol{C}_i] = 0, \quad k = 1, \ldots, p,$$

where $\boldsymbol{\Psi}_k$ is the k-th row of coefficient $\boldsymbol{\Psi}$, $h(X_i)$ is a user-specified function ensuring that $\mathbb{E}(\partial U_{\text{IPW},i}/\partial \alpha_k)$ is invertible. Under the full-ascertainment assumption (Stene, 1977), $w_i = 1$ for case and $\frac{P(1-K)}{K(1-P)}$ for control, where K represents the population-level trait prevalence, and P denotes the case-control ratio in the study, both assumed to be known. Setting $h(X_i) = X_i$ simplifies the estimating equation into a weighted least squares regression problem.

To identify non-mediators and noise variables, Wald tests are conducted to detect $\boldsymbol{M}_{\mathcal{I}_1} = \{M_j | \alpha_j = 0 \text{ and } \beta_j \neq 0 \text{ for } j \in \mathcal{I}_1\}$ and $\boldsymbol{M}_{\mathcal{I}_3} = \{M_j | \alpha_j = 0 \text{ and } \beta_j = 0 \text{ for } j \in \mathcal{I}_3\}$. Wald test p-values are adjusted using a multiple testing correction procedure to control the false discovery rate (FDR) at a pre-specified threshold. The resulting filtered set of mediators, denoted by $\boldsymbol{M}_{\hat{\mathcal{T}}}$, is subsequently used in cross-fitted variance component estimation. Given that $\text{var}(\alpha) = \pi_\alpha \sigma_\alpha^2$, the variance estimator is obtained as $\hat{\sigma}_\alpha^2 = \widehat{\text{var}(\alpha)}/\hat{\pi}_\alpha$, where $\text{var}(\alpha)$ is empirically estimated from the IPW results, and π_α is estimated based on the ratio of the number of selected mediators and the total number of potential mediators.

4.2. *Estimation of γ.* The IPW approach can also be applied to estimate the parameter γ. By selecting an appropriate estimating equation, the IPW problem becomes a weighted probit regression problem with the same weight as described in Section 4.1. However, directly regressing the outcome Y on the exposure variable X and covariates $\boldsymbol{C}$ may lead to biased estimates, as it does not account for $\boldsymbol{M}$. To mitigate this bias, we include the top principal components (PCs) of the mediator matrix across individuals, denoted as $\boldsymbol{W}$, as additional covariates in the model. By incorporating these components, we can effectively control for the influence of $\boldsymbol{M}$ that may introduce bias in the estimation of γ (Lemma 3, Section 1 of the Supplementary Material (Kang et al., 2025)).

4.3. *Estimation of* σ_{11}^2. For the estimation of σ_{11}^2, we propose to use PCGC regression (Golan, Lander and Rosset, 2014). PCGC regression is a method of moments that estimates variance components under a liability model while accounting for the oversampling in case-control designs. As a generalization of the Haseman–Elston regression method (Haseman and Elston, 1972), PCGC regression models the expected correlation of the observed phenotype as a function of variance components and the correlation of its corresponding design matrix, which, in our context, corresponds to the mediator matrix across individuals. Unlike the restricted maximum likelihood (REML) method (Lee et al., 2011), which tends to underestimate variance components in case-control studies due to ascertainment bias, PCGC explicitly accounts for this bias by modeling the inclusion probability of individuals in the study. However, this approach implicitly assumes weak correlation among rows of the design matrix, an assumption that is likely violated by mediators due to their relationship with the exposure (Equation (3)). To account for this, we modify the PCGC regression framework and estimate σ_{11}^2 following the cross-fitting procedure as described in Section 4:

Step 1 (First subsample): Identify the mediators with $\alpha \neq 0$, i.e., $\boldsymbol{M}_{\mathcal{T}}$ and $\boldsymbol{M}_{\mathcal{I}_2}$ following Section 4.1 and denote the selected mediator set as $\hat{\mathcal{T}}$.

Step 2 (Second subsample): Estimate standardized residual $\hat{\boldsymbol{\xi}}_{\hat{\mathcal{T}}}$ by regressing scaled $\boldsymbol{M}$ on X, $\boldsymbol{C}$, $\boldsymbol{W}$ using the IPW approach. Apply PCGC regression with $\hat{\boldsymbol{\xi}}_{\hat{\mathcal{T}}}$ as the standardized design matrix and include X, $\boldsymbol{C}$ and $\boldsymbol{W}$ as covariates to estimate $\hat{\sigma}_{11}^2$.

After adjusting for covariates that drive the correlation between mediators, the resulting residuals $\hat{\boldsymbol{\xi}}_{\hat{\mathcal{T}}}$ in Step 2 would satisfy the weak correlation assumption in PCGC regression. We note that adjusting for $\boldsymbol{W}$ in the analysis not only provides us a consistent estimate of direct effect in high-dimensional settings (Section 4.2), but also potentially controls for unknown confounding in the same spirit as Xu et al. (2024). Additionally, ascertainment bias is appropriately addressed in both steps through the IPW approach for residual estimation and PCGC regression for variance estimation.

4.4. *Theoretical results on consistency.* In this section, we present the theoretical properties of the proposed method. For theoretical analysis, we assume that while the mediators may be correlated, there are no directed causal relationships among them. Instead, their correlation is driven by the exposure X, covariates $\boldsymbol{C}$, and unobserved/unknown latent factors $\boldsymbol{U}$ (Figure 1). This assumption is formalized as follows.

ASSUMPTION 2 (Conditional Parallel Mediators). There exist random variables $\boldsymbol{U}$ that have a global influence on $\boldsymbol{M}$, where $M_k \perp\!\!\!\perp M_j | X, \boldsymbol{U}, \boldsymbol{C}$ for any $k, j \in \{1, \ldots, p\}$.

Assumption 2 is commonly used in methods and applications to high-dimensional omics mediation analysis (Huang and Pan, 2015; Perera et al., 2022; Chi et al., 2026), as usually the omics data were measured at a single time point rather than sequentially, exemplified by our motivating dataset, WHI. An example for such unobserved/unknown latent variables $\boldsymbol{U}$ is batch effect (e.g., lab conditions, technicians, etc.), which frequently occurs in biological data and typically happens after sample collection (i.e., $\boldsymbol{U} \perp\!\!\!\perp S$). In our proposed estimation procedure, the random variable $\boldsymbol{U}$ is captured by PC analysis (Section 4.2). Moreover, our estimation method remains robust when Assumption 2 is violated, as empirically shown in Supplementary Material Section 2.8 (Kang et al., 2025).

Given Assumptions 1 and 2, Theorem 4.1 establishes the consistency of our proposed total mediation effect estimator $\widehat{R}^2_{\text{med;causal}}$ and relative measurement estimator $\widehat{Q}^2_{\text{med}}$ under mild regularity conditions.

THEOREM 4.1. *Under Assumptions 1-2 and Assumptions 3-10 in Section 1 of the Supplementary Material (Kang et al., 2025), and assume there exists $c_0 > 0$ such that $\frac{\text{TP}}{p} > c_0$, then for any $\epsilon > 0$, there exists $\delta > 0$, such that if $\frac{\text{TP}}{\text{TP+FP}} > 1 - \delta$ for p, n large enough, we have*

$$\lim_{p\to\infty, n\to\infty} \mathbb{P}(|\widehat{R}^2_{\text{med;causal}} - R^2_{\text{med;causal}}| < \epsilon) = 1,$$

and

$$\lim_{p\to\infty, n\to\infty} \mathbb{P}(|\widehat{Q}^2_{\text{med}} - Q^2_{\text{med}}| < \epsilon) = 1.$$

Let $A = \{\alpha \neq 0, \hat{\alpha} \neq 0\}$ and $B = \{\alpha = 0, \hat{\alpha} \neq 0\}$. Herein, TP and FP correspond to $|A|$ and $|B|$ in the initial filtering step of α (Section 4.1). A key assumption underlying this result is that the proportion of true mediators and non-mediators type $\boldsymbol{M}_{\mathcal{I}_2}$ among all potential mediators remains above a positive threshold as $p \to \infty$. This assumption ensures a sufficiently large number of mediators contribute to the total mediation effect. Another important assumption is the control of FDR when identifying non-mediators type $\boldsymbol{M}_{\mathcal{I}_1}$ and noise mediators $\boldsymbol{M}_{\mathcal{I}_3}$. We assume that the FDR for selecting variables impacted by exposure is stringent enough to minimize the inclusion of non-mediators type $\boldsymbol{M}_{\mathcal{I}_1}$ and noise mediators $\boldsymbol{M}_{\mathcal{I}_3}$. In practice, our simulations suggest that a reasonable choice, such as FDR $= 0.01$ using Benjamini-Hochberg (BH) procedure, yields robust performance across various sample sizes and mediator sparsity levels (Detailed empirical evaluation in Supplementary Section 2.7). Notably, the proposed estimator is shown to be consistent without requiring mediator selection consistency but only sufficiently low FDR in testing p marginal association (i.e., $\alpha = 0$). This is especially desirable in high-dimensional situations, as exact selection is rarely achieved.

5. Simulation studies.

5.1. *Simulation settings.* We conducted simulation studies to evaluate the consistency of the proposed estimators of $R^2_{\text{med;causal}}$ and Q^2_{med}. Four scenarios were considered, each varying the composition of mediator types while fixing the number of mediators $p = 2000$:

- Scenario 1 (S1): 180 $\boldsymbol{M}_{\mathcal{T}}$, 420 $\boldsymbol{M}_{\mathcal{I}_1}$, 420 $\boldsymbol{M}_{\mathcal{I}_2}$, 980 $\boldsymbol{M}_{\mathcal{I}_3}$;
- Scenario 2 (S2): 360 $\boldsymbol{M}_{\mathcal{T}}$, 240 $\boldsymbol{M}_{\mathcal{I}_1}$, 840 $\boldsymbol{M}_{\mathcal{I}_2}$, 560 $\boldsymbol{M}_{\mathcal{I}_3}$;
- Scenario 3 (S3): 360 $\boldsymbol{M}_{\mathcal{T}}$, 840 $\boldsymbol{M}_{\mathcal{I}_1}$, 240 $\boldsymbol{M}_{\mathcal{I}_2}$, 560 $\boldsymbol{M}_{\mathcal{I}_3}$;
- Scenario 4 (S4): 720 $\boldsymbol{M}_{\mathcal{T}}$, 480 $\boldsymbol{M}_{\mathcal{I}_1}$, 480 $\boldsymbol{M}_{\mathcal{I}_2}$, 320 $\boldsymbol{M}_{\mathcal{I}_3}$.

The binary outcome was generated under the liability threshold model, with population disease prevalence $K = 0.05$ and study disease prevalence $P = 0.5$ to mimic a case-control design with a 1:1 case-control ratio. Sample size n varies from 500 to 2000 with an increment of 500. We set a single exposure variable $X \sim N(0, 1)$ and its direct effect parameter $\gamma = 0.5$. We specified $\sigma^2_\alpha = 0.3$, the cumulative variance of the true mediators $\boldsymbol{M}_{\mathcal{T}}$ as $\sigma^2_{11} = 0.5$, and that of the non-mediators $\boldsymbol{M}_{\mathcal{I}_1}$ as $\sigma^2_{01} = 0.3$. The coefficients $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ were generated from the working model (6). To ensure Assumptions 1 and 2 hold, we generated $\xi_k = \alpha_0 U + \nu_k$ for $k = 1, \ldots, p$, where $U \sim N(0, 1)$, $\nu_k \sim N(0, 1)$ and $\alpha_0 \sim N(0, 0.3)$, with U, ν_k and α_0 mutually independent. The vector of $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_p)^\top$ was then used to generate mediators as $\boldsymbol{M} = \boldsymbol{\alpha} X + \boldsymbol{\xi}$. The error term for the outcome model was generated from $\epsilon \sim N(0, 1 - \sigma^2_{11} - \sigma^2_{01})$. Additionally, the simulation settings under the global null, i.e., no mediators and no mediation effects, are described in Section 2.1 of the Supplementary Material (Kang et al., 2025). For each scenario, we ran 100 replications across all sample sizes.

Under S1, we further compared our methods to competing methods: HIMA (Zhang et al., 2016), HDMA (Gao et al., 2019) and BAMA (Song et al., 2019) in terms of total mediation effect estimate and PM estimate. We varied the prevalence of disease K from 0.5 (cohort study) to 0.05 (case-control study) and the sample size from 500 to 2000. For the competing methods, we calculated mean-based mediation effects as recommended in their paper. For methods HIMA and HDMA, their IE measure is the inner product of $\hat{\boldsymbol{\alpha}}$ and $\hat{\boldsymbol{\beta}}$. Their proportion mediated, following the definition in traditional analysis, is $\hat{\boldsymbol{\alpha}}^\top\hat{\boldsymbol{\beta}}/(\hat{\gamma}+\hat{\boldsymbol{\alpha}}^\top\hat{\boldsymbol{\beta}})$. While for BAMA, the mediation effect measure is given by $\mathrm{IE}_{\mathrm{BAMA}}=\sum_{i=1}^{p}(\hat{\alpha}_i\hat{\beta}_i)^2$. The results are summarized in the following section. Furthermore, we calculated the corresponding $R^2_{\mathrm{med;causal}}$ and Q^2_{med} based on their component-wise estimates $\hat{\boldsymbol{\alpha}}$, $\hat{\boldsymbol{\beta}}$ and $\hat{\gamma}$. Relative bias was also reported for fair comparisons across different measures, defined as Bias/True Value.

5.2. *Simulation results.* Figure 2 and Table 1 summarize the estimation performance of two proposed measures in terms of bias, standard deviation, mean absolute deviation and mean square error. Overall, the proposed measures had minimal bias and error observed across different sample sizes and number/types of non-mediators. As expected, the performance metrics improved as the sample size increased. Furthermore, the violin plots (Figure 2) indicate the consistency of the proposed estimation methods as well as their robustness under various conditions. The performance of our proposed estimators under the global null is shown in Section 2.1 of the Supplementary Material (Kang et al., 2025), where all performance metrics are at the same scale as the Table 1, suggesting consistency in both null and alternative cases.

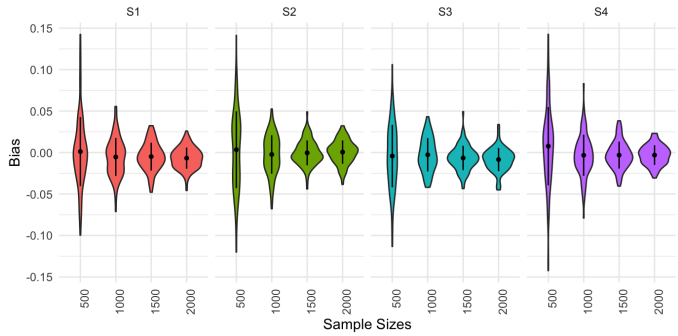

(a) Bias for $R^2_{\mathrm{med;causal}}$ estimation

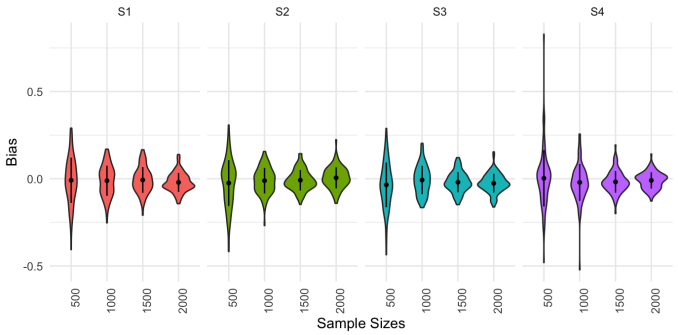

(b) Bias for Q^2_{med} estimation

FIG 2. *Violin plots of bias in a case-control study ($K=0.05$) under varying sample sizes and proportions of non-mediators across Scenarios S1-S4. The black dots represent the average bias, while the vertical lines correspond to the range of one standard deviation centered at the mean bias.*

TABLE 1

Performance of proposed estimators under different scenarios. Bias represents empirical bias, SD represents standard deviation, MAD is mean absolute deviation, and MSE is mean square error, all based on 100 repetitions.

		$R^2_{\text{med;causal}}$				Q^2_{med}			
Scenario	n	Bias	SD	MAD	MSE	Bias	SD	MAD	MSE
S1	500	0.0012	0.0415	0.0297	0.0017	-0.0088	0.13	0.101	0.0167
	1000	-0.0054	0.0229	0.0183	0.0005	-0.0114	0.0864	0.0714	0.0075
	1500	-0.0049	0.0168	0.0139	0.0003	-0.0071	0.0735	0.0592	0.0054
	2000	-0.0067	0.0128	0.0115	0.0002	-0.0209	0.0543	0.0469	0.0034
S2	500	0.0035	0.0464	0.0387	0.0021	-0.0248	0.132	0.103	0.0177
	1000	-0.0023	0.0232	0.0189	0.0005	-0.011	0.0734	0.0592	0.0055
	1500	-0.0003	0.0150	0.0116	0.0002	-0.0084	0.0583	0.0469	0.0034
	2000	0.0006	0.014	0.0111	0.0002	0.0053	0.0603	0.047	0.0036
S3	500	-0.0042	0.0377	0.0302	0.0014	-0.0351	0.128	0.104	0.0175
	1000	-0.0027	0.0201	0.0168	0.0004	-0.0072	0.0813	0.0666	0.0066
	1500	-0.0065	0.0147	0.0125	0.0003	-0.0197	0.0579	0.0492	0.0037
	2000	-0.0084	0.0139	0.0124	0.0003	-0.0258	0.054	0.0469	0.0035
S4	500	0.0077	0.0471	0.0359	0.0023	0.0025	0.162	0.108	0.0259
	1000	-0.0033	0.0248	0.0188	0.0006	-0.021	0.105	0.0771	0.0114
	1500	-0.0031	0.0162	0.013	0.0003	-0.018	0.0618	0.0505	0.0041
	2000	-0.0031	0.0118	0.0098	0.0001	-0.0095	0.0466	0.0371	0.0022

Table 2 presents a comparison of our method with HIMA, HDMA, and BAMA in terms of bias, standard deviation, and mean square error for estimating IE and PM under case-control designs ($K = 0.05$). Our proposed method consistently outperformed competing methods across all simulation scenarios. HIMA and HDMA often identified only a small number of mediators (typically fewer than 10), which indicated their limited ability to capture mediators with small effect sizes, resulting in higher bias and less reliable estimates. Moreover, neither HIMA nor HDMA accounts for ascertainment bias when estimating the exposure–mediator associations $\boldsymbol{\alpha}$, which likely contributes to their biased effect estimates under case-control designs. BAMA, on the other hand, explicitly assumes that mediator effects follow a normal distribution with either large or small variance, which may explain its improved performance over HIMA and HDMA in some simulation scenarios. However, the amount of bias was still higher than desired. Simulation results with $K = 0.2$ and 0.5 are presented in Section 2 of the Supplementary Material (Kang et al., 2025). As the prevalence of disease K increased, our method showed reduced efficiency in small sample sizes, indicated by the higher estimation variance. However, with larger sample sizes, the method's performance became more stable, yielding comparable results across different K values.

Computation times in both case-control and cohort studies are summarized in Section 2.3 of the Supplementary Material (Kang et al., 2025). No significant difference in running time was observed between cohort ($K = 0.5$) and case-control ($K = 0.05$) studies across all methods. However, the running time of other methods increased noticeably with larger

sample sizes. Overall, HIMA, HDMA, and BAMA exhibited approximately linear growth as n increased, while the proposed method has the slowest growth in running time. Particularly, BAMA, under a Bayesian sampling estimation framework, is the slowest approach. While HDMA was faster than BAMA, it remained slower than HIMA and the proposed method.

In Supplementary Sections 2.4–2.8 (Kang et al., 2025), we further assess the robustness of our method in the presence of a few strong mediators, disease prevalence misspecification, small total effects, varying FDR thresholds, and violation of the parallel mediation assumption. Overall, the proposed approach consistently demonstrated stable performance across these realistic scenarios.

TABLE 2

Comparison of the relative performance of IE and PM estimators in a case-control study ($K = 0.05$). Results are shown for both the proposed measures ($R^2_{\text{med;causal}}$ and Q^2_{med}) and the original method-specific measures: POC estimator for HIMA and HDMA, and $\sum_{i=1}^{p}(\hat{\alpha}_i\hat{\beta}_i)^2$ for BAMA when estimating IE; and $\hat{\boldsymbol{\alpha}}^\top\hat{\boldsymbol{\beta}}/(\hat{\gamma}+\hat{\boldsymbol{\alpha}}^\top\hat{\boldsymbol{\beta}})$ for HIMA and HDMA when estimating PM. Each cell presents the relative bias, with the standard deviation of the relative deviation shown in parentheses

		Proposed Measure				Method-Specific Measure		
Type	n	Proposed	HIMA	HDMA	BAMA	HIMA	HDMA	BAMA
IE	500	-0.02	0.24	1.12	-0.91	1.15	2.55	-0.92
		(0.09)	(1.74)	(1.71)	(0.01)	(0.61)	(2.34)	(0.01)
	1000	-0.02	0.41	1.71	-0.92	1.36	3.56	-0.93
		(0.18)	(1.12)	(1.14)	(0.01)	(0.53)	(2.27)	(0.01)
	1500	0	0.37	1.76	-0.92	1.38	3.40	-0.92
		(0.12)	(0.99)	(0.80)	(0.01)	(0.40)	(2.38)	(0.01)
	2000	-0.02	0.45	1.75	-0.92	1.49	3.80	-0.92
		(0.09)	(0.77)	(0.70)	(0.01)	(0.37)	(2.33)	(0.01)
PM	500	0.04	-0.44	0.24	-0.26	0.09	0.80	-
		(0.13)	(0.76)	(0.98)	(0.16)	(0.33)	(1.19)	
	1000	0.03	-0.28	0.78	-0.24	0.20	1.32	-
		(0.20)	(0.56)	(0.78)	(0.13)	(0.27)	(1.18)	
	1500	0.03	-0.29	0.78	-0.25	0.21	1.23	-
		(0.14)	(0.51)	(0.68)	(0.09)	(0.23)	(1.19)	
	2000	0.04	-0.22	0.88	-0.27	0.26	1.42	-
		(0.13)	(0.40)	(0.69)	(0.09)	(0.21)	(1.16)	

6. Application to the Women's Health Initiative (WHI) study. The WHI nested case-control substudy (dbGaP accession number: phs001334) examines the relationship between metabolomic profiles and incident CHD in postmenopausal women. The substudy consists of subjects from both the WHI Observational Study (WHI-OS) and the WHI Hormone Therapy Trials (WHI-HT). A total of 2,306 samples (681 case-control pairs from WHI-HT and 472 case-control pairs from WHI-OS) were analyzed for metabolomic profiling, representing 2,151 unique consented participants in the analysis. Baseline targeted metabolomic profiles were measured using liquid chromatography–mass spectrometry, yielding 371 metabolites after initial quality control. Cases were defined as postmenopausal women who experienced incident myocardial infarction (MI) or CHD death during follow-up. In WHI-HT, cases were those who developed MI or CHD death during the hormone therapy intervention period. In WHI-OS, cases were selected from previously identified incident CHD cases in the prior

biomarker study (BAA 22), excluding those who experienced a stroke before MI or CHD death. Controls were matched by age, race/ethnicity, enrollment period, and hysterectomy status, excluding those with prior cardiovascular diseases.

BMI is a well-established modifiable risk factor for CHD (Eckel and Krauss, 1998) and has widespread impact on metabolites (Dong et al., 2023), but its impact on CHD through metabolites remains less studied. Metabolites with more than 50% missing values were excluded from the analysis. For the remaining metabolites, missing values were imputed by assigning half of the lowest observed value (Paynter et al., 2018). Metabolomic profiles were log-transformed. We further removed an outlier by inspecting a plot of the top two principal components of metabolomic profiles. Smoking status, alcohol consumption, and education level were included as covariates, with missing values imputed using multiple imputation by chained equations (van Buuren and Groothuis-Oudshoorn, 2011). After these steps, the final dataset contained 2150 samples with 366 metabolite variables. The prevalence of disease in the population was estimated as a weighted sum of disease prevalence across ethnic groups, using data from the American Heart Association's Heart Disease and Stroke Statistics report (Tsao et al., 2022).

6.1. *Preliminary analysis: Evidence of weak effects.* In this section, we show that the data exhibited evidence of mediation signals in the paths of interest (BMI-metabolites-CHD as in Figure 3), particularly from weak mediators. Figure 3(a) presents a conceptual diagram of the hypothesized mediation structure in the WHI case-control study, where BMI influences CHD risk through high-dimensional metabolomics profiles. As a preliminary analysis, we first conducted exposure-mediator association tests to identify a subset of mediator candidates, using the significance threshold of 0.05 for p-values. Figure 3(b) displays a histogram of p-values for testing the effect of the subset of mediators on the outcome, conditional on BMI. Both the exposure–mediator and mediator–outcome tests were marginal tests and were adjusted for relevant covariates including smoking status, alcohol consumption, and education level. The red dashed line represents the estimated null density level, computed as the proportion of null mediators using Storey's method (Storey, 2002), multiplied by the expected uniform density under the null hypothesis. P-values above this level are more likely to originate from mediators with a non-zero effect on the outcome. The prevalence of such p-values, particularly for $p > 0.05$, suggests the presence of diffuse weak mediation signals. This pattern contrasts with the expected distribution under a sparse alternative, where only a few strong mediators produce a sharp concentration of p-values near zero, while the rest follow a uniform distribution under the null.

In addition, a total of 123 metabolites exhibited statistically significant associations in both marginal tests on exposure-mediator association and mediator-outcome association. However, HIMA and HDMA identified only 9 and 13 mediators, respectively, as they primarily capture mediators with strong effects. This contrast highlights the importance of capturing weaker mediation effects, which may collectively contribute to the mediation process for BMI and CHD, as well as other potential exposure-outcome relationships of interest, especially those with widespread impact on molecular phenotypes (Moore et al., 2014; Sebastiani et al., 2024).

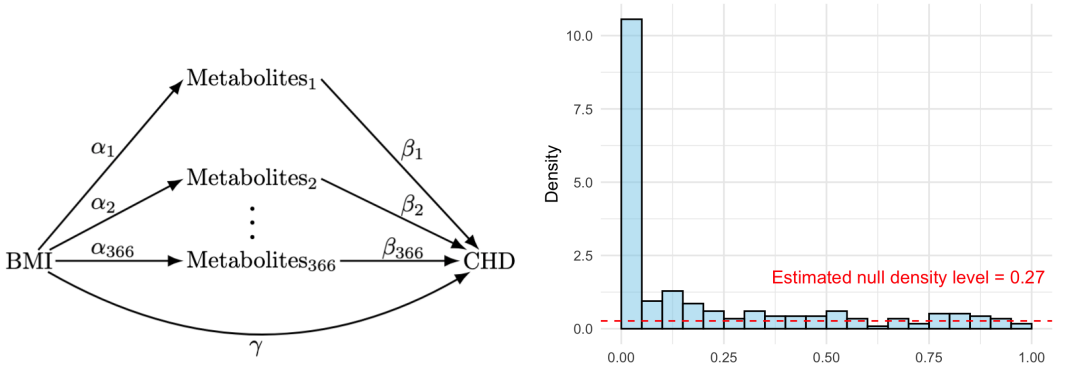

(a) Conceptual mediation diagram.

(b) Histogram of p-values for testing mediator-outcome associations.

FIG 3. *(a) Conceptual diagram illustrating the hypothesized mediation structure between BMI, metabolomics mediators, and CHD in the WHI case-control study. (b) Histogram of p-values for testing $\beta_i = 0$ in the outcome model, where β_i represents the effect of the i-th metabolite on CHD, conditional on BMI. The tests shown are restricted to mediator candidates that passed a prior test for exposure-mediator associations (i.e., with α_i significantly different from 0). The red dashed line represents the estimated density of null p-values using Storey's method.*

6.2. *Comparison and interpretation of total mediation effect estimates.* Table 3 summarizes the total mediation estimates of BMI on CHD through metabolites obtained from different methods. For HIMA, HDMA, and BAMA, empirical $R^2_{\text{med;causal}}$ values were calculated based on the component-wise mediation estimates. The confidence intervals (CIs) for the proposed estimates were calculated using the jackknife method, where each iteration involved omitting one individual from the dataset and recomputing the estimates. The standard error of the estimator was obtained from the variability of the jackknife replicates, and 95% CIs were then calculated using a normal approximation based on the jackknife standard error. For comparison, bootstrap-based CIs were used for HIMA and HDMA, while posterior credible intervals derived from the posterior distribution were reported for BAMA. The coverage probability and computational feasibility of confidence interval was evaluated through simulation and presented in Supplementary Material Section 2.9 (Kang et al., 2025). We also evaluated the coverage probability of jackknife-based CI procedure and observed conservative yet effective nominal coverage levels under the null in simulations (Details in Supplementary Section 2.2, Kang et al. (2025)).

Notably, evidence of opposing component-wise mediation effect estimates was observed across methods. HDMA method produced the smallest POC mediation effect estimate $\hat{\boldsymbol{\alpha}}^\top \hat{\boldsymbol{\beta}}$. Among the 13 significant mediators identified by HDMA, 7 exhibited positive effects while 6 exhibited negative effects, leading to substantial cancellation in the total estimate. In contrast, HIMA identified 9 significant mediators, 8 of which had positive effects, resulting in a POC mediation effect estimate that was approximately ten times larger than that of HDMA, despite having a much smaller $R^2_{\text{med;causal}}$ estimate. While HDMA had the largest $R^2_{\text{med;causal}}$ estimate, it included three potential non-mediators with p-values of $\alpha > 0.05$, i.e., these three non-mediators may be included due to an inflated type I error in estimating α, which arose from not accounting for case-control selection.

In contrast, our proposed method estimated the total mediation effect of BMI on CHD through the selected set of 366 metabolites as $R^2_{\text{med;causal}} = 0.0076$, with a 95% CI of (0.0033, 0.0119), indicating that approximately 0.76% of the variance in the underlying liability for CHD was causally explained by the metabolomic mediation pathway. Although the

estimated $R^2_{\text{med;causal}}$ value may appear small, it reflects a statistically significant total mediation effect of a single exposure on disease liability through thousands of high-dimensional metabolomic features. In fact, small absolute mediation effects are expected herein, given the limited total effect of BMI on CHD liability estimated as 0.0086. On the other hand, the proportion of the exposure's explained variance on the liability scale that is accounted for by metabolomic profiles, as quantified by Q^2_{med}, was estimated at 88.55%, with a 95% CI of (57.32%, 100%), suggesting that a substantial portion of BMI's effect on CHD liability was mediated through these profiles.

In high-dimensional omics studies, small $R^2_{\text{med;causal}}$ values on an absolute scale can reflect meaningful biological signal. Unlike existing approaches that focus on selecting mediators with relatively strong individual effects (e.g., HIMA and HDMA), our method aggregates contributions across weak mediators, yielding a $\hat{R}^2_{\text{med;causal}}$ that almost doubled the value reported by HIMA/much higher than BAMA, capturing the total mediation effect that is missed by existing methods (Table 3).

TABLE 3

Total mediation effect of metabolites from BMI to CHD in WHI nested case-control study. 95% confidence interval is presented in parenthesis.

	$\hat{R}^2_{\text{med;causal}}$	$\hat{\boldsymbol{\alpha}}^\top\hat{\boldsymbol{\beta}}$	$\sum_i(\hat{\alpha}_i\hat{\beta}_i)^2$	$\hat{Q}^2_{\text{med}}$	$\frac{\hat{\boldsymbol{\alpha}}^\top\hat{\boldsymbol{\beta}}}{\hat{\boldsymbol{\alpha}}^\top\hat{\boldsymbol{\beta}}+\hat{\gamma}}$
Proposed	0.0076 (0.0033, 0.0119)	-	-	0.8855 (0.5732, 1)	-
HIMA	0.0033 (0.0019, 0.0294)	0.09 (0.0376, 0.1744)	0.0031 (0.0022, 0.0438)	0.8962 (0.1961, 0.9989)	0.821 (0.2027, 5.6155)
HDMA	0.0174 (0.0284, 0.1506)	0.0082 (-0.3889, 0.4528)	0.0162 (0.0363, 0.2787)	0.4443 (0.1141, 0.9985)	0.0517 (-24.5906, 15.0723)
BAMA	0.0002 (0.0001, 0.0002)	0.0231 (0.0122, 0.0344)	0.0002 (0.0001, 0.0002)	0.4954 (0.0873, 0.9947)	0.7361 (0.02538, 2.0473)

7. Discussion. We have developed a novel variance-based total mediation effect measure for binary outcomes and high-dimensional mediators. Existing mediation measures for binary outcomes are typically mean-based and the choice of the measure often varies by disease prevalence. The key contributions of our work include (1) a new variance-based measure designed for binary outcomes, which addresses the challenge of bidirectional component-wise effects, invariant to disease prevalence, and provides a clear causal interpretation; (2) an efficient estimation procedure that accounts for ascertainment bias and effectively handles non-mediators and weak mediators. Many exposures, such as age, alcohol consumption, and BMI, have widespread effects on omics variables and may impact health outcomes through such omics mediators. The new measure and estimation method, motivated by real data, address the critical gaps in handling these broadly impactful exposures and realistic mediation patterns. To ensure theoretical validity, we demonstrate that under mild conditions, and without requiring exact mediator selection, which permits the inclusion of mediators with weak effects, our estimators converge reliably when the false discovery rate is appropriately controlled (see Section 1 of the Supplementary Material (Kang et al., 2025)). Through simulation studies, we demonstrated its robustness in estimating mediation effects under varying scenarios, and its applicability was further validated in the WHI study. The proposed method is implemented in the R package "r2MedCausal", available on GitHub at https://github.com/zhiyu-kang/r2MedCausal.

Despite its strengths, our proposed $R^2_{\text{med;causal}}$ estimation procedure has some limitations. First of all, our estimation method specifically targets capturing weak but non-negligible mediation effects. It does not take into account strong effects explicitly. In other words, when

only a few (≤ 5) sparse strong effects exist, existing high-dimensional mediation analysis methods with POC measures are more effective (Details in Supplementary Section 2.4, Kang et al. (2025)). Therefore, in practice, we recommend performing preliminary analysis as in Section 5.1 before using our method. Employing a multiple-step procedure may solve this issue (Li et al., 2025); however, it is not directly assessed in this work. Secondly, while PCGC was proposed to account for ascertainment bias in case-control studies, it may be less efficient than likelihood-based methods for cohort studies. We established its theoretical property under conditional parallel mediator assumption; however, we showed that our method was fairly robust to mild violation of this assumption (Supplementary Section 2.8). We recommend checking the residual distribution and consider pre-screening, increasing the number of principal components, or other batch effect correction methods if strong dependencies are observed. Thirdly, as a proportion estimator, Q^2_{med} becomes unstable when the denominator gets close to zero (e.g., total effect < 0.003; see Supplementary Section 2.6). In practice, we suggest pursuing mediation analysis only when sufficiently large total effect is observed, following Baron and Kenny (1986). Lastly, we used jackknife for CI estimation, which we have refined its computation in our R package; however, it could still be time-consuming with large sample sizes (Supplementary Section 2.9). In terms of further research directions, we are interested in extending our causal mediation framework to multi-omics or tensor data, where mediators originate from multiple sources on different scales.

Acknowledgments. Z. Kang and L. Chen contributed equally. C. Li and T. Yang jointly supervised the project. Correspondence: C. Li (chunlin@virginia.edu) and T. Yang (yang3704@umn.edu). T. Yang would like to thank Josey Sorenson for helping improve the interpretation of the proposed measures. The WHI program is funded by the National Heart, Lung, and Blood Institute, National Institutes of Health, U.S. Department of Health and Human Services through contracts HHSN268201600018C, HHSN268201600001C, HHSN268201600002C, HHSN268201600003C, and HHSN268201600004C. Metabolomic analyses were funded by NHLBI contract HHSN268201300008C. This manuscript was not prepared in collaboration with investigators of the WHI, has not been reviewed and/or approved by WHI, and does not necessarily reflect the opinions of the WHI investigators or the NHLBI. The WHI dataset used for the analyses described in this manuscript was obtained from dbGaP at https://www.ncbi.nlm.nih.gov/gap through accession number phs001334. The authors acknowledge the Minnesota Supercomputing Institute (MSI) at the University of Minnesota for providing resources that contributed to the research results reported in this paper. The authors also acknowledge the use of generative AI tools to assist in refining sentence structure, improving clarity, and ensuring conciseness in the manuscript. The authors would like to thank the referee, the Associate Editor and the Editor for the constructive feedback which greatly improved this work.

Funding. P. Wei was partially supported by the National Institutes of Health (NIH) grant R21HL170213. T. Yang was supported by St. Baldrick Career Award (1463960) and NIH grant R01AG074858. C. Li was partially supported by the National Science Foundation (NSF) grant DMS2515789.

Significance Statement. Modern biomedical studies often measure thousands of molecular features to understand how exposures influence disease risk. However, existing measures of total mediation are not well suited for high-dimensional settings, especially when many small effects act in opposite directions. We propose a new causally interpretable measure that summarizes the overall mediation effect for binary disease outcomes, together with a fast estimation procedure that applicable to case–control sampling and large numbers of mediators. An application to a metabolomics study demonstrates that the proposed method can

detect mediation signals missed by existing approaches. This work provides a practical tool for studying mediation effects in high-dimensional omics data.

SUPPLEMENTARY MATERIAL

Technical proofs
Contains all formal proofs of the main theoretical results and supporting lemmas, together with the full list of technical conditions used throughout the manuscript.

Additional simulation results
Presents extra numerical experiments evaluating estimator performance under a variety of settings.

PCGC algorithm
Gives a step-by-step description of the PCGC regression procedure when covariates are present.

Illustration of cross-fitting procedure
Provides a schematic diagram of the two-fold sample-splitting and cross-fitting steps.

r2MedCausal
Provides the R package "r2MedCausal" implementing the proposed method described in this manuscript. The package is also publicly available at https://github.com/zhiyu-kang/r2MedCausal.

REFERENCES

BARON, R. and KENNY, D. (1986). The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology* **51** 1173–1182.

CHI, S., LI, X., WEI, P. and HUANG, X. (2026). MELODY: Mediation Analysis in Logistic Regression for High-Dimensional Mediators and a Binary Outcome. *Genetic Epidemiology* **50** e70033.

CLARK-BOUCHER, D., ZHOU, X., DU, J., LIU, Y., NEEDHAM, B. L., SMITH, J. A. et al. (2023). Methods for mediation analysis with high-dimensional DNA methylation data: Possible choices and comparisons. *PLOS Genetics* **19** 1–26.

DOBRIBAN, E. and WAGER, S. (2018). High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics* **46** 247–279.

DONG, Q., SIDRA, S., GIEGER, C., WANG-SATTLER, R., RATHMANN, W., PREHN, C., ADAMSKI, J., KOENIG, W., PETERS, A., GRALLERT, H. et al. (2023). Metabolic signatures elucidate the effect of body mass index on type 2 diabetes. *Metabolites* **13** 227.

ECKEL, R. H. and KRAUSS, R. M. (1998). American Heart Association call to action: obesity as a major risk factor for coronary heart disease. *Circulation* **97** 2099–2100.

FAIRCHILD, A. J., MACKINNON, D. P., TABORGA, M. P. and TAYLOR, A. B. (2009). R2 effect-size measures for mediation analysis. *Behavior Research Methods* **41** 486–498.

FLINT, A. J., REXRODE, K. M., HU, F. B., GLYNN, R. J., CASPARD, H., MANSON, J. E., WILLETT, W. C. and RIMM, E. B. (2010). Body mass index, waist circumference, and risk of coronary heart disease: a prospective study among men and women. *Obesity Research & Clinical Practice* **4** e171–e181.

GAO, Y., YANG, H., FANG, R., ZHANG, Y., GOODE, E. L. and CUI, Y. (2019). Testing Mediation Effects in High-Dimensional Epigenetic Studies. *Frontiers in Genetics* **10** 1195.

GAYNOR, S. M., SCHWARTZ, J. and LIN, X. (2019). Mediation analysis for common binary outcomes. *Statistics in Medicine* **38** 512–529.

GOLAN, D., LANDER, E. S. and ROSSET, S. (2014). Measuring missing heritability: Inferring the contribution of common variants. *Proceedings of the National Academy of Sciences* **111** E5272–E5281.

GREENLAND, S., PEARL, J. and ROBINS, J. M. (1999). Confounding and Collapsibility in Causal Inference. *Statistical Science* **14** 29–46.

GÖRING, H. H. H., TERWILLIGER, J. D. and BLANGERO, J. (2001). Large Upward Bias in Estimation of Locus-Specific Effects from Genomewide Scans. *The American Journal of Human Genetics* **69** 1357–1369.

HASEMAN, J. and ELSTON, R. (1972). The investigation of linkage between a quantitative trait and a marker locus. *Behavior Genetics* **2** 3–19.

HORVITZ, D. G. and THOMPSON, D. J. (1952). A Generalization of Sampling Without Replacement From a Finite Universe. *Journal of the American Statistical Association* **47** 663–685.

HUANG, Y.-T. and PAN, W.-C. (2015). Hypothesis Test of Mediation Effect in Causal Mediation Model With High-Dimensional Continuous Mediators. *Biometrics* **72** 402–413.

KANG, Z., CHEN, L., WEI, P., XU, Z., LI, C. and YANG, T. (2025). Supplement to "Estimation of total mediation effect for a binary trait in a case-control study for high-dimensional omics mediators".

LEE, S. H., WRAY, N. R., GODDARD, M. and VISSCHER, P. M. (2011). Estimating missing heritability for disease from genome-wide association studies. *American Journal of Human Genetics* **88** 294–305.

LI, C., CHEN, L., PANKOW, J. S. and YANG, T. (2025). Measuring weak effects in high dimensional mediation analysis. *medRxiv* 2025–01.

LYALL, D. M., CELIS-MORALES, C., WARD, J., ILIODROMITI, S., ANDERSON, J. J., GILL, J. M., SMITH, D. J., NTUK, U. E., MACKAY, D. F., HOLMES, M. V. et al. (2017). Association of body mass index with cardiometabolic disease in the UK Biobank: a Mendelian randomization study. *JAMA cardiology* **2** 882–889.

MACKINNON, D. P., LOCKWOOD, C. M., BROWN, C. H., WANG, W. and HOFFMAN, J. M. (2007). The intermediate endpoint effect in logistic and probit regression. *Clinical Trials* **4** 499–513.

MOORE, S. C., MATTHEWS, C. E., SAMPSON, J. N., STOLZENBERG-SOLOMON, R. Z., ZHENG, W., CAI, Q. et al. (2014). Human metabolic correlates of body mass index. *Metabolomics* **10** 259–269.

NGUYEN, T. Q., WEBB-VARGAS, Y., KONING, I. M. and STUART, E. A. (2016). Causal mediation analysis with a binary outcome and multiple continuous or ordinal mediators: Simulations and application to an alcohol intervention. *Structural Equation Modeling: A Multidisciplinary Journal* **23** 368–383.

GLOBAL BURDEN OF METABOLIC RISK FACTORS FOR CHRONIC DISEASES COLLABORATION, LU, Y., HAJIFATHALIAN, K., EZZATI, M., WOODWARD, M., RIMM, E. B., DANAEI, G., D'ESTE, C. et al. (2014). Metabolic mediators of the effects of body-mass index, overweight, and obesity on coronary heart disease and stroke: a pooled analysis of 97 prospective cohorts with 1·8 million participants. *The Lancet* **383** 970–983.

PACKER, M., ZILE, M. R., KRAMER, C. M., BAUM, S. J., LITWIN, S. E., MENON, V., GE, J., WEERAKKODY, G. J., OU, Y., BUNCK, M. C. et al. (2025). Tirzepatide for heart failure with preserved ejection fraction and obesity. *New England Journal of Medicine* **392** 427–437.

PAYNTER, N. P., BALASUBRAMANIAN, R., GIULIANINI, F., WANG, D. D., TINKER, L. F., GOPAL, S. et al. (2018). Metabolic Predictors of Incident Coronary Heart Disease in Women. *Circulation* **137** 841–853.

PEARL, J. (2001). Direct and Indirect Effects. *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence* **411** 411–420.

PEARL, J. (2010). An Introduction to Causal Inference. *The International Journal of Biostatistics* **6** 7.

PERERA, C., ZHANG, H., ZHENG, Y., HOU, L., QU, A., ZHENG, C. et al. (2022). HIMA2: high-dimensional mediation analysis and its application in epigenome-wide DNA methylation data. *BMC Bioinformatics* **23** 296.

RIJNHART, J. J. M., LAMP, S. J., VALENTE, M. J., MACKINNON, D. P., TWISK, J. W. R. and HEYMANS, M. W. (2021). Mediation analysis methods used in observational research: a scoping review and recommendations. *BMC Medical Research Methodology* **21** 226.

SCHIFANO, E. D. (2019). A review of analysis methods for secondary outcomes in case-control studies. *Communications for Statistical Applications and Methods* **26** 103–129.

SEBASTIANI, P., MONTI, S., LUSTGARTEN, M. S., SONG, Z., ELLIS, D., TIAN, Q. et al. (2024). Metabolite signatures of chronological age, aging, survival, and longevity. *Cell Reports* **43** 114913.

SILVA, S., FATUMO, S. and NITSCH, D. (2024). Mendelian randomization studies on coronary artery disease: a systematic review and meta-analysis. *Systematic Reviews* **13** 29.

SIMOES, F. N. F. Q., DASTANI, M. and VAN OMMEN, T. (2024). Fundamental Properties of Causal Entropy and Information Gain. *arXiv* **abs/2402.01341**.

SOFER, T., CORNELIS, M. C., KRAFT, P. and TCHETGEN, E. J. T. (2017). Control function assisted ipw estimation with a secondary outcome in case-control studies. *Statistica Sinica* **27** 785–804.

SOHN, M. B., LU, J. and LI, H. (2022). A compositional mediation model for a binary outcome: Application to microbiome studies. *Bioinformatics* **38** 16–21.

SONG, Y., ZHOU, X., ZHANG, M., ZHAO, W., LIU, Y., KARDIA, S. L. R. et al. (2019). Bayesian Shrinkage Estimation of High Dimensional Causal Mediation Effects in Omics Studies. *Biometrics* **76** 700–710.

STENE, J. (1977). Assumptions for Different Ascertainment Models in Human Genetics. *Biometrics* **33** 523–527.

STOREY, J. D. (2002). A Direct Approach to False Discovery Rates. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **64** 479–498.

TSAO, C. W., ADAY, A. W., ALMARZOOQ, Z. I., ALONSO, A., BEATON, A. Z., BITTENCOURT, M. S. et al. (2022). Heart Disease and Stroke Statistics—2022 Update: A Report From the American Heart Association. *Circulation* **145** e153–e639.

VAN BUUREN, S. and GROOTHUIS-OUDSHOORN, K. (2011). mice: Multivariate Imputation by Chained Equations in R. *Journal of Statistical Software* **45** 1–67.

VANDERWEELE, T. J. (2016). Mediation Analysis: A Practitioner's Guide. *Annual Review of Public Health* **37** 17–32.

VANDERWEELE, T. J. and VANSTEELANDT, S. (2010). Odds Ratios for Mediation Analysis for a Dichotomous Outcome. *American Journal of Epidemiology* **172** 1339–1348.

XU, Z., LI, C., CHI, S., YANG, T. and WEI, P. (2024). Speeding up interval estimation for R2-based mediation effect of high-dimensional mediators via cross-fitting. *Biostatistics* **26** kxae037.

YANG, T., NIU, J., CHEN, H. and WEI, P. (2021). Estimation of total mediation effect for high-dimensional omics mediators. *BMC Bioinformatics* **22** 414.

ZHANG, H., ZHENG, Y., ZHANG, Z., GAO, T., JOYCE, B., YOON, G. et al. (2016). Estimating and testing high-dimensional mediation effects in epigenetic studies. *Bioinformatics* **32** 3150–3154.