The Annals of Statistics
2023, Vol. 51, No. 1, 211–232
https://doi.org/10.1214/22-AOS2247

BREAKING THE WINNER'S CURSE IN MENDELIAN RANDOMIZATION: RERANDOMIZED INVERSE VARIANCE WEIGHTED ESTIMATOR

BY XINWEI MA[1,a], JINGSHEN WANG[2,b] AND CHONG WU[3,c]

[1]*Department of Economics, University of California San Diego,* [a]*x1ma@ucsd.edu*
[2]*Division of Biostatistics, University of California Berkeley,* [b]*jingshenwang@berkeley.edu*
[3]*Department of Biostatistics, The University of Texas MD Anderson Cancer Center,* [c]*cwu18@mdanderson.org*

Developments in genome-wide association studies and the increasing availability of summary genetic association data have made the application of two-sample Mendelian Randomization (MR) with summary data increasingly popular. Conventional two-sample MR methods often employ the same sample for selecting relevant genetic variants and for constructing final causal estimates. Such a practice often leads to biased causal effect estimates due to the well-known "winner's curse" phenomenon. To address this fundamental challenge, we first examine its consequence on causal effect estimation both theoretically and empirically. We then propose a novel framework that systematically breaks the winner's curse, leading to unbiased association effect estimates for the selected genetic variants. Building upon the proposed framework, we introduce a novel rerandomized inverse variance weighted estimator that is consistent when selection and parameter estimation are conducted on the same sample. Under appropriate conditions, we show that the proposed RIVW estimator for the causal effect converges to a normal distribution asymptotically and its variance can be well estimated. We illustrate the finite-sample performance of our approach through Monte Carlo experiments and two empirical examples.

1. Introduction.

1.1. *Motivation*. Mendelian Randomization (MR) has become a popular tool for causal inference in observational studies. Its popularity partially ascribes to a fundamental limitation of observational data, where causation cannot be simply inferred from the association between an exposure and a disease due to the issue of unmeasured confounders [35]. Since genetic variants are randomly assorted from parents and are fixed at conception, they provide a source of exogenous variation in the exposure. MR mitigates unmeasured confounding bias by incorporating these genetic variants as instrumental variables (IVs) [22]. In MR, the most commonly used genetic variants are single nucleotide polymorphisms (SNPs).

The proliferation of genome-wide association study (GWAS) data not only improves the statistical power of MR analyses thanks to the increased sample sizes, but also facilitates efficient evaluations of causal effects without the need to share information at individual level [10]. Two-sample MR, a design strategy where genetic associations with the exposure and the outcome are taken from separate samples in the same population, yields valid causal estimates even if concomitant data on the SNPs, exposure, and outcome are not available for any individual, provided that each included SNP is a valid IV [8, 27]. To be precise, the included SNPs must be (i) associated with the exposure (relevance assumption), (ii) not associated with any unmeasured confounder (effective random assignment) and (iii) only associated with the outcome through the exposure (exclusion restriction).

Received May 2022; revised September 2022.
MSC2020 subject classifications. Primary 62E20, 62F10; secondary 62P10, 62P25.
Key words and phrases. Two-sample Mendelian randomization, inverse variance weighting, post-selection inference, instrumental variable, causal inference.

In practice, to make these assumptions plausible in the summary data setting where the whole-genome SNPs (about 10 million common variants) are harvested from GWAS, preselection procedures need to be conducted before carrying out the MR analysis. In particular, because the strength of each SNP can be tested empirically, practitioners routinely select SNPs that are strongly associated with the exposure to meet the relevance assumption. As it is well recognized in the literature [15, 21, 33], however, this preselection step can lead to biased causal effect estimates due to the well-known "winner's curse" phenomenon. This winner's curse bias, which is the result of using the same sample to select the SNPs and to form the final causal estimate, is also borne out in our simulation and empirical studies (see Tables 2 and 3).

1.2. *Our contributions*. Given the issues mentioned above, the primary contribution of this paper is a novel MR framework and a Rerandomized Inverse Variance Weighted (RIVW) estimator that break the winner's curse in the two-sample MR with summary data setting. In addition, we provide a systematic study of their theoretical properties along with comprehensive evaluations of their finite-sample performance in simulation studies and real data analyses. We further break down our contribution as follows.

On the statistical methodology side, we demonstrate that the winner's curse pushes the estimated effect of SNP on exposure away from zero, and such distortion leads to an underestimation bias in the classical two-sample Inverse Variance Weighted (IVW) estimator (Section 2). To break the winner's curse, we introduce pseudo SNP-exposure associations into the selection step, and then use Rao–Blackwellization to recover the correct center of the effect of selected SNPs on exposure (Section 3). Because Rao–Blackwellization [4] produces an estimator that is optimal by the mean-squared-error criterion among unbiased estimators, we make the fullest possible use of data from each SNP without incurring any selection bias from data re-use in SNP selection and parameter estimation. Moreover, we propose a novel RIVW estimator that not only eliminates winner's curse bias but also removes measurement error bias in the classical IVW approach.

On the theoretical side, we first establish an asymptotic normality result for the proposed RIVW estimator (Theorem 1). In particular, this result shows that our RIVW estimator is correctly centered despite the fact that the same data is used for SNP selection and causal effect estimation. In other words, the RIVW estimator is immune to winner's curse bias as well as measurement error bias. As part of this endeavor, we rigorously show that the Rao–Blackwellized SNP-exposure association is unbiased after selection, and we further provide an unbiased estimate of its variance (Section 4.2). The latter turns out to be crucial for correcting measurement error bias. Next, we propose a consistent standard error motivated by the regression interpretation of the RIVW estimator (Theorem 2). Different from other constructions in the literature, which can take different forms depending on if balanced horizontal pleiotropy is present, our standard error remains the same in both cases, alleviating the burden of choosing what type of assumptions should be used for practitioners (Corollary 1 and Section 5.1). Third, because our theoretical developments require bounding moments of distributions related to the truncated normal with a diverging truncation threshold, some of the theoretical results can be of independent interest (see the Supplementary Material [25] for details).

On the practical side, we illustrate the finite-sample performance of our approach through Monte Carlo experiments (Section 6) and two empirical examples (Section 7), demonstrating that the winner's curse leads to biased causal estimates and such a bias can be corrected by the RIVW estimator. In particular, when analyzing the causal association between body mass index (BMI) and COVID-19 illness, our approach, in accordance with existing literature, confirms that high BMI is a causal risk factor for COVID-19 severity. Moreover, as our

approach removes the downward bias caused by the winner's curse, we conjecture that the effect of high BMI on COVID-19 severity (odds ratio: 1.60) is higher than what the literature currently believes. Effectively communicating this fact to the general public and persuading people with obesity to take extra precautions may reduce the number of hospitalized COVID-19 patients and save lives (Section 7.3).

1.3. *Existing literature.* The winner's curse is well recognized in the Mendelian randomization and human genomics literature as a fundamental challenge. To avoid this issue, the most common practice is to use a third independent data set for instrument selection, known as "three-sample" MR designs. While appealing, finding a third independent sample might be difficult in practice, as practitioners tend to use the largest GWAS data (meta-analyzed results from most available cohorts) to improve power, leaving fewer options for the third independent data set. In two-sample MR analyses with summary data, MR estimators that break the winner's curse with rigorous statistical guarantees have been lacking.

Other than the issue of winner's curse bias, we also face the challenge that all GWAS associations are estimated, that is, the associations between the SNPs and the exposure/outcome are measured with errors. Such a measurement error issue is also known as weak instrument bias in the MR literature; see [1] and [33] for more discussions and additional references. Without accounting for this issue, classical MR estimators (such as the IVW estimator) are often biased downwards. Our proposed RIVW estimator formulated in (4) removes this measurement error bias, making it immune to both sources of biases in the two-sample MR with summary data framework. See [6, 41, 44] for other recent attempts to tackle measurement error bias.

By correcting winner's curse bias and measurement error bias, the proposed RIVW estimator improves upon the classical IVW approach, where the latter often serves as a benchmark in MR analyses [9]. For instance, in clinical studies (see [2, 20, 23, 28, 34, 38], for example), researchers use causal effect estimates obtained via IVW as their main results and validate their findings through sensitivity and robust MR analyses. This practice highlights an urgent need to provide an unbiased causal effect estimator of the IVW-type, which is our focus here. In this paper, we also demonstrate that the proposed RIVW estimator is robust to the presence of balanced horizontal pleiotropy, and we leave general robust MR analyses to future research.

Some existing approaches use all SNPs for MR analyses (i.e., without instrument selection). For example, [41] propose a debiased IVW estimator, which incorporates all SNPs into the estimator to avoid winner's curse bias. Likelihood-based approaches [43], which are also robust to measurement error bias, are arguably more complex than our current construction. Moreover, as mentioned by [43], the likelihood-based estimator may not deliver unique estimates for certain data generating processes. Our RIVW estimator, on the other hand, is easy to construct and does not require any numerical optimization.

1.4. *Two-sample MR with summary data framework.* In this section, we introduce two-sample MR framework with summary data, whereby genetic associations with the exposure and outcome are gleaned from independent samples to furnish a "two-sample" analysis.

Throughout this paper, we denote by β the causal effect of an exposure variable X on an outcome variable Y. In the presence of an unmeasured confounder U, Mendelian Randomization (MR) analyses incorporate p genetic variants $G_1, \ldots, G_p$ as instrumental variables (IVs). Here, genetic variant $G_j \in \{0, 1, 2\}$ represents the number of alleles of a SNP j assigned to an individual. Existing MR literature starts with the following model:

$$Y = \beta X + \theta_Y U + E_Y, \qquad X = \sum_{j=1}^{p} \gamma_j G_j + \theta_X U + E_X, \tag{1}$$

TABLE 1

Two sample MR with summary data illustration. $\widehat{\gamma}_j$ is obtained by regressing $(X_1^*, \ldots, X_{n_X}^*)$ on $(G_{1j}^*, \ldots, G_{n_X,j}^*)$, and $\widehat{\Gamma}_j$ follows from the regression of $(Y_1, \ldots, Y_{n_Y})$ on $(G_{1j}, \ldots, G_{n_Y,j})$. The corresponding standard errors are σ_{X_j} and σ_{Y_j}, respectively

	Individual-level data	Publicly available summary data
1. *Exposure data set*	$\{(X_i^*, G_{ij}^*)\}_{i=1}^{n_X}$	$\{(\widehat{\gamma}_j, \sigma_{X_j})\}_{j=1}^{p}$
2. *Outcome data set*	$\{(Y_i, G_{ij})\}_{i=1}^{n_Y}$	$\{(\widehat{\Gamma}_j, \sigma_{Y_j})\}_{j=1}^{p}$

where $\gamma_1, \ldots, \gamma_p, \beta, \theta_X, \theta_Y$ are unknown parameters. E_X and E_Y are mutually independent noise variables that are also independent with $(G_1, \ldots, G_p, U)$. We later discuss a more general model setup allowing for direct dependence between Y and G_j in Section 5.1.

Because sharing individual level data is usually impractical due to privacy constraints [10], two-sample MR with summary data attempts to conduct MR analyses using summarized estimates of SNP-exposure and SNP-outcome associations based on two large and independent data sets (see an illustration in Table 1). To be more precise, we have access to the marginal regression coefficients and standard errors for all SNPs, $\{(\widehat{\gamma}_j, \sigma_{X_j})\}_{j=1}^{p}$, which are obtained by regressing the exposure on each SNP separately on the exposure data set $\{(X_i^*, G_{ij}^*)\}_{i=1}^{n_X}$. Similarly, by regressing the outcome on each SNP separately on the outcome data set $\{(Y_i, G_{ij})\}_{i=1}^{n_Y}$, we obtain the marginal regression coefficients and the standard errors, $\{(\widehat{\Gamma}_j, \sigma_{Y_j})\}_{j=1}^{p}$.

To facilitate discussion, we start with the following assumption. We shall demonstrate in Section 5.1 that the proposed method remains valid under a more general balanced horizontal pleiotropy Assumption 1′ (see Corollary 1). In Section 5.3, we will further extend our proposed method to allow for correlated SNPs.

ASSUMPTION 1 (Measurement error model). (i) For any $j \neq j'$, the pairs, $(\widehat{\Gamma}_j, \widehat{\gamma}_j)$ and $(\widehat{\Gamma}_{j'}, \widehat{\gamma}_{j'})$ are mutually independent.

(ii) For each j,

$$\begin{bmatrix} \widehat{\Gamma}_j \\ \widehat{\gamma}_j \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \Gamma_j \\ \gamma_j \end{bmatrix}, \begin{bmatrix} \sigma_{Y_j}^2 & 0 \\ 0 & \sigma_{X_j}^2 \end{bmatrix} \right).$$

In addition, there exists some $\nu \to 0$, such that $\{\sigma_{Y_j}/\nu, \sigma_{X_j}/\nu : 1 \leq j \leq p\}$ are bounded and bounded away from zero.

The normality assumption is plausible and is widely adopted in the MR literature, because $\widehat{\Gamma}_j$ and $\widehat{\gamma}_j$ are estimated from modern GWAS with hundreds of thousands of samples, making the normal approximation highly accurate. To meet with the mutual independence assumption, one can conduct linkage disequilibrium (LD) pruning [29] on the exposure data set to pick independent SNPs.[1] Because classical LD pruning selects independent SNPs without using the GWAS summary statistics information, distribution of the selected SNP-exposure associations remains unchanged. In Section 5.3, we propose a revised sigma-based LD pruning method that improves the efficiency of causal effect estimation. Finally, to simplify the presentation, we assume that the estimation uncertainties (i.e., σ_{Y_j} and σ_{X_j}) are of the same order. While it is possible to generalize this condition in our theoretical analysis, it will unavoidably lead to more cumbersome notation.

[1] We note that this is a special case of linear instrumental variable models, in the sense that marginal association estimates are assumed to be independent. As a result, methods tailored to MR analyses (including ours) cannot be directly applied to more general instrumental variables models such as those discussed in [3, 13, 14, 16, 26, 45].

2. Winner's curse. In this section, we first introduce winner's curse bias in the two-sample MR with summary data framework. We then demonstrate how the winner's curse leads to the violation of Assumption 1. Lastly, we show empirically that the popular IVW estimator underestimates the true causal effect due to winner's curse bias (and measurement error bias).

In MR analyses, to meet the relevance assumption, practitioners often select a set of genetic instruments (i.e., SNPs) that are strongly associated with the exposure based on hard-thresholding individual z-scores:

$$\widetilde{\mathcal{S}}_\lambda = \left\{ j : \left| \frac{\widehat{\gamma}_j}{\sigma_{X_j}} \right| > \lambda,\ j = 1, \ldots, p \right\}, \tag{2}$$

where λ is a prespecified cutoff value.[2] This exercise casts "winner's curse" into MR analyses.

Heuristically, the winner's curse, as discussed here, is caused by the tendency of the estimates $\{\widehat{\gamma}_j, j \in \widetilde{\mathcal{S}}_\lambda\}$ to be more extreme than the underlying true association effects $\{\gamma_j, j \in \widetilde{\mathcal{S}}_\lambda\}$. It can be further understood from the "file drawer effect," coined in [32], which refers to the observation that statistically insignificant results are much less likely to be published, and hence they remain in file drawers. To put this classical problem in the MR context, $\widehat{\gamma}_j$ will be employed in the construction of the final causal effect estimator only when it is statistically significant. As a result, the distribution of $\widehat{\gamma}_j$ after surviving the selection step is a truncated Gaussian. Because the selection step will distort the distribution of $\widehat{\gamma}_j$, it leads to the violation of Assumption 1. Classical MR estimators (such as two-sample incorporating instrument selection may suffer from the winner's curse, rendering both estimation and statistical inference problematic.

Although there is a general awareness in the MR literature that data-driven instrument selection affects subsequent estimation and inference, it remains pervasive in practice to ignore the consequence of this selection step. This is primarily due to the fact that, in an attempt to increase sample size and statistical power, most GWAS tend to meta-analyze results from all available cohorts, making it challenging to find an independent (third) sample from the same population for instrument selection. To demonstrate the consequence of performing statistical inference with the same data set for instrument selection, we examine in a simulation study the performance of IVW estimators. The IVW estimator employs summary associations obtained from two independent GWAS with SNP selection, and it estimates β by a weighted regression of $\widehat{\Gamma}_j$ on $\widehat{\gamma}_j$:

$$\widehat{\beta}_{\lambda,\mathtt{IVW}} = \frac{\sum_{j \in \widetilde{\mathcal{S}}_\lambda} \widehat{\Gamma}_j \widehat{\gamma}_j / \sigma_{Y_j}^2}{\sum_{j \in \widetilde{\mathcal{S}}_\lambda} \widehat{\gamma}_j^2 / \sigma_{Y_j}^2}. \tag{3}$$

EXAMPLE 1. We generate 2000 Monte Carlo samples with the true causal effect $\beta = 0.2$ and $p = 200{,}000$ independent SNPs. The true SNP-exposure associations, γ_j follow the mixture distribution:

$$\begin{pmatrix} \gamma_j \\ \alpha_j \end{pmatrix} \sim \pi_x \rho \begin{pmatrix} \mathcal{N}(0, \varepsilon_x^2) \\ \delta_0 \end{pmatrix} + \pi_x (1-\rho) \begin{pmatrix} \mathcal{N}(0, \varepsilon_x^2) \\ \mathcal{N}(0, \tau^2) \end{pmatrix} + \pi_y \begin{pmatrix} \delta_0 \\ \mathcal{N}(0, \tau^2) \end{pmatrix} + (1 - \pi_x - \pi_y) \begin{pmatrix} \delta_0 \\ \delta_0 \end{pmatrix},$$

[2]The cutoff value, λ, is often chosen to be $\Phi^{-1}(1-\alpha/2)$, which is the $(1-\alpha/2)$th quantile of the standard normal distribution. A large cutoff value λ is often desired so that the selected SNPs meet the relevance assumption (e.g., corresponding to $\alpha = 5 \times 10^{-8}$ or 5×10^{-5}).

and the true SNP-outcome associations are $\Gamma_j = \beta\gamma_j + \alpha_j$. Here, α_j captures balanced horizontal pleiotropy (see Assumption 1′), $\mathcal{N}(0, \varepsilon^2)$ represents a normal distribution with mean 0 and variance ε^2, and δ_0 is the Dirac measure centred at zero. The first two components of the above mixture distribution correspond to relevant IVs (i.e., having nonzero associations with the exposure), and $1-\rho$ controls the fraction of relevant IVs that also exhibit pleiotropic effects. The third term represents IVs that are only associated with the outcome but not the exposure, while the last component in the mixture distribution corresponds to SNPs that have no association with either the outcome or the exposure. Similar data generating models have been widely used in the literature [30, 42]. For future reference, we define heritability as the proportion of the phenotypic variance explained by additive effects of all genetic variants [40]. Because the data has been standardized and the SNPs are assumed to be independent, the heritability of exposure (h_x^2) and outcome (h_y^2) can be calculated as

$$h_x^2 = \mathbb{V}\left[\sum_{j=1}^{p} \gamma_j G_j\right] / \mathbb{V}[X] = \sum_{j=1}^{p} \gamma_j^2 \mathbb{V}[G_j] \approx p\pi_x \varepsilon_x^2,$$

$$h_y^2 = \mathbb{V}\left[\sum_{j=1}^{p} \Gamma_j G_j\right] / \mathbb{V}[Y] = \sum_{j=1}^{p} \Gamma_j^2 \mathbb{V}[G_j] \approx \beta^2 h_x^2 + p\big(\pi_x(1-\rho) + \pi_y\big)\tau^2.$$

For simplicity, we set $\rho = 1$ in this example. Additional simulation evidence is available in Section 6 and in the Supplementary Material [25].

As the magnitude of winner's curse bias changes with the number of IVs around the cutoff value λ, we vary $\varepsilon_x^2 = \tau^2$ in the set $\{2\times10^{-5}, 3\times10^{-5}, 5\times10^{-5}, 1\times10^{-4}, 3\times10^{-4}, 5\times10^{-4}\}$ and $\pi_x = \pi_y$ in the set $\{0.005, 0.05\}$.[3] We set $\lambda = 5.45$ for the IVW estimator $\widehat{\beta}_{\lambda,\mathtt{IVW}}$. To generate the SNP-exposure effects $\widehat{\gamma}_j$ and SNP-outcome effects $\widehat{\Gamma}_j$, we set the sample size $n_X = n_Y = 100{,}000$. The standard deviations are set to be $\sigma_{X_j} = 1/\sqrt{n_X}$ and $\sigma_{Y_j} = 1/\sqrt{n_Y}$.

In Figure 1, we showcase the performance of (i) the three-sample IVW estimator as a benchmark, (ii) the two-sample IVW estimator with instrument selection and (iii) our proposed RIVW (to be discussed in Section 3), where we report the bias proportion (absolute of Monte Carlo bias divided by the true parameter $\beta = 0.2$) of each estimator. Note that for three-sample IVW, we use a third independent exposure GWAS data with sample size 100,000 for instrumental variable selection. See Section 6 for other implementation details.

Results in Figure 1 demonstrate that the bias of $\widehat{\beta}_{\lambda,\mathtt{IVW}}$ is larger if more SNP-exposure associations are around the selection cutoff λ. On the other hand, our proposed RIVW estimator is not sensitive to the change of such proportions, meaning that it is immune to winner's curse bias. Not surprisingly, the three-sample IVW estimator also does not suffer from winner's curse bias, although as we mentioned earlier, finding a large third independent sample from the same population for instrument selection might be difficult in practice. In addition, we still observe a small residual bias for the three-sample IVW estimator. This is caused by measurement error bias. We also provide mean F statistic (mF) for each simulation setting in Figure 1. Because instrument selection is employed by all three procedures, observed mean F statistic is reasonably large. As a result, we only see moderate measurement error bias. We discuss in Section 4.3 methods for correcting measurement error bias.

[3] Some combinations are ignored, because we require the heritability to be between 0 and 1, and that at least three IVs are selected for IVW with commonly used p-value threshold 5×10^{-8}.

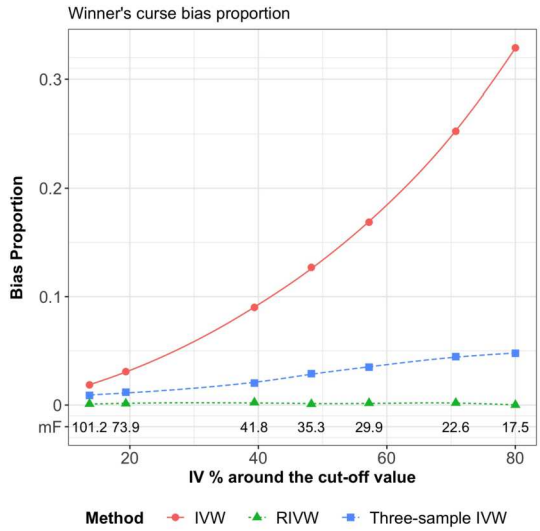

FIG. 1. *Illustration of winner's curse and measurement error bias following the simulation setup in Example 1. The IV proportion (x-axis) is calculated as the number of IVs with p-values lying between* 5×10^{-8} *and* 5×10^{-10} *divided by the number of selected IVs with p-value* $< 5 \times 10^{-8}$. *mF is mean F statistic for the instruments used in the three-sample IVW.*

3. Rerandomized IVW estimator. In this section, we first discuss how we lift the winner's curse after instrument selection and the intuition behind our method. We then introduce our rerandomized IVW (RIVW) estimator for estimating the true causal effect β. To streamline the presentation, theoretical results and their formal statements are presented in Section 4. Section 5 considers several extensions to our framework, including (i) allowing for balanced horizontal pleiotropy, (ii) a novel smoothed RIVW estimator constructed via reweighting the SNP-exposure associations and (iii) an attempt to incorporate correlated IVs.

3.1. *Winner's curse removal.* The fact that having a third independent sample for instrument selection lifts the winner's curse has a direct implication. Recall that $\widehat{\gamma}_j$ is the measured SNP-exposure association in the exposure data set, and suppose $\widehat{\gamma}'_j$ is another estimate of the SNP-exposure association (with standard error σ'_{X_j}) obtained from a third independent GWAS. Then the three-sample MR estimator is free of winner's curse bias, because the instrument selection event is *independent* of the SNP-exposure association estimation, that is,

$$\widehat{\gamma}_j \perp\!\!\!\perp \underbrace{\left|\frac{\widehat{\gamma}'_j}{\sigma'_{X_j}}\right| > \lambda}_{\substack{\text{IV selection event} \\ \text{on a third GWAS}}} \quad \Longrightarrow \quad \mathbb{E}\left[\widehat{\gamma}_j \mid \left|\frac{\widehat{\gamma}'_j}{\sigma'_{X_j}}\right| > \lambda\right] = \mathbb{E}[\widehat{\gamma}_j] = \gamma_j.$$

Therefore, the post-selection mean of $\widehat{\gamma}_j$ is not distorted by instrument selection, and the winner's curse does not arise. Following this line of reasoning, we can also lift the winner's curse in two-sample MRs, as long as we can *create* an unbiased estimator of γ_j that is independent of the instrument selection event within the same exposure data set. Carrying forward the heuristic of "creating independence," we propose a new two-sample MR framework that fully removes the winner's curse. We start by introducing a randomized instrument selection step that deviates from the hard-thresholding rule in (2):

STEP 1 (Randomized instrument selection). *For each SNP $j = 1, 2, \ldots, p$, generate a pseudo SNP-exposure association effect $Z_j \sim \mathcal{N}(0, \eta^2)$, and select SNP j if*

$$\left| \frac{\widehat{\gamma}_j}{\sigma_{X_j}} + Z_j \right| > \lambda.$$

Also, define the set of selected SNPs as

$$\mathcal{S}_\lambda = \{j : S_j > 0, j = 1, 2, \ldots, p\}, \quad \text{where } S_j = \left| \frac{\widehat{\gamma}_j}{\sigma_{X_j}} + Z_j \right| - \lambda.$$

Here, $\lambda > 0$ is a cutoff value, and η is a prespecified constant that reflects the noise level of the pseudo SNPs. Discussions on the choice of λ and η are given in Section 6.1.

It is worth mentioning that although we randomize $\widehat{\gamma}_j$ before checking whether it passes the threshold λ, such randomization will not affect the selection of strong SNPs (i.e., when $|\gamma_j/\sigma_{X_j}| \gg 0$). In other words, the pseudo effects we introduce merely randomize the selection around the cutoff value λ. The benefit of introducing randomized instrument selection will be apparent in the following bias removal step.

STEP 2 (Winner's curse removal by Rao–Blackwellization). *For each selected SNP $j \in \mathcal{S}_\lambda$, construct an unbiased estimator of γ_j as*

$$\widehat{\gamma}_{j,\mathtt{RB}} = \widehat{\gamma}_j - \frac{\sigma_{X_j}}{\eta} \frac{\phi(A_{j,+}) - \phi(A_{j,-})}{1 - \Phi(A_{j,+}) + \Phi(A_{j,-})}, \quad \text{where } A_{j,\pm} = -\frac{\widehat{\gamma}_j}{\sigma_{X_j}\eta} \pm \frac{\lambda}{\eta},$$

and $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal density and cumulative distribution functions.

To understand the above construction, consider the crude "initial estimator," $\widehat{\gamma}_{j,\mathtt{ini}} = \widehat{\gamma}_j - \frac{\sigma_{X_j}}{\eta^2} Z_j$. By standard calculations, it is unbiased for γ_j, and is also independent of the IV selection event $j \in \mathcal{S}_\lambda$. As a result, $\widehat{\gamma}_{j,\mathtt{ini}}$ is unbiased for γ_j both before and after IV selection:

$$\mathbb{E}[\widehat{\gamma}_{j,\mathtt{ini}} | j \in \mathcal{S}_\lambda] = \mathbb{E}[\widehat{\gamma}_{j,\mathtt{ini}}] = \gamma_j.$$

Then we employ Rao–Blackwellization to improve the initial estimator (i.e., to reduce its variance) by projecting it onto the sufficient statistic of the selection likelihood:

$$\widehat{\gamma}_{j,\mathtt{RB}} = \mathbb{E}[\widehat{\gamma}_{j,\mathtt{ini}} | \widehat{\gamma}_j, j \in \mathcal{S}_\lambda].$$

Implied by the Rao–Blackwell theorem, our $\widehat{\gamma}_{j,\mathtt{RB}}$ not only is unbiased for the true γ_j conditioning on the selection event, but also achieves the minimum variance for a given η. A rigorous justification of our proposal will be given in Section 4. Before proceeding, we note that the "initial estimator" is introduced only for theoretical discussions and to motivate our construction. Our approach neither employs this initial estimator nor requires constructing it explicitly. In a two-stage genome-wide association studies setting, a similar strategy based on Rao–Blackwellization has been adopted in [7].

EXAMPLE 2. To illustrate how instrument selection affects the distributions of $\widehat{\gamma}_j$, and to demonstrate the performance of our bias-corrected $\widehat{\gamma}_{j,\mathtt{RB}}$, we provide a simulation study in which we generate $\widehat{\gamma}_j/\sigma_{X_j}$ from the normal distribution $\mathcal{N}(\gamma_j/\sigma_{X_j}, 1)$. We vary the true normalized instrument effect, γ_j/σ_{X_j}, in the set $\{0.1\lambda,\ \lambda,\ 4\lambda\}$, representing weak, moderately strong, and strong instruments, respectively. We set $\lambda = \Phi^{-1}(1 - 5 \times 10^{-5}/2)$. From the three histograms in Figure 2, we observe that the naïve estimator is generally biased since it does not account for selection (unless the given SNP is a very strong instrument). On the other hand, our Rao–Blackwellized estimator remains correctly centered regardless of the strength of the instrument.

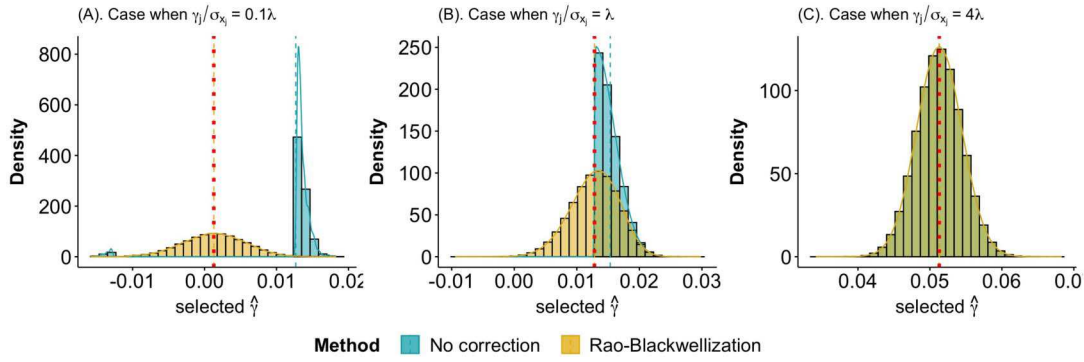

FIG. 2. *Comparison between $\widehat{\gamma}_{j,\mathrm{RB}}$ and $\widehat{\gamma}_j$ after SNP selection. Panel (A): the weak IV scenario with $\gamma_j/\sigma_{X_j} = 0.1\lambda$. Panel (B): the moderately strong IV scenario with $\gamma_j/\sigma_{X_j} = \lambda$. Panel (C): the strong IV scenario with $\gamma_j/\sigma_{X_j} = 4\lambda$. The true instrument effect, γ_j, is represented by red dotted lines.*

3.2. *Rerandomized IVW estimator.* Given the unbiased estimate of γ_j in the previous section, we propose our final estimator using $\widehat{\gamma}_{j,\mathrm{RB}}$. Because Mendelian randomization is the term applied to the random assortment of alleles at the time of gamete formation, and our approach further randomizes instrument/SNP selection, we name $\widehat{\beta}_{\mathrm{RIVW}}$ in the following the rerandomized IVW (RIVW) estimator.

STEP 3 (Rerandomized IVW estimator). *Construct the RIVW estimator as*

$$\widehat{\beta}_{\mathrm{RIVW}} = \frac{\sum_{j\in\mathcal{S}_\lambda} \widehat{\Gamma}_j \widehat{\gamma}_{j,\mathrm{RB}}/\sigma_{Y_j}^2}{\sum_{j\in\mathcal{S}_\lambda} (\widehat{\gamma}_{j,\mathrm{RB}}^2 - \widehat{\sigma}_{X_j,\mathrm{RB}}^2)/\sigma_{Y_j}^2}, \tag{4}$$

where

$$\widehat{\sigma}_{X_j,\mathrm{RB}}^2 = \sigma_{X_j}^2 \left(1 - \frac{1}{\eta^2}\frac{A_{j,+}\phi(A_{j,+}) - A_{j,-}\phi(A_{j,-})}{1-\Phi(A_{j,+})+\Phi(A_{j,-})} + \frac{1}{\eta^2}\left(\frac{\phi(A_{j,+})-\phi(A_{j,-})}{1-\Phi(A_{j,+})+\Phi(A_{j,-})}\right)^2\right),$$

and $A_{j,\pm}$ are defined in Step 2.

The specific form of our estimator deviates from the classical IVW estimator defined in (3). Such deviation is motivated by the consideration that $\widehat{\gamma}_{j,\mathrm{RB}}$ are random quantities measured with errors. We subtract from $\widehat{\gamma}_{j,\mathrm{RB}}^2$ a variance estimate, $\widehat{\sigma}_{X_j,\mathrm{RB}}^2$, to correct this measurement error bias.

We are now ready to present our final step, statistical inference in two-sample MR with summary data accounting for both instrument selection bias and measurement error bias.

STEP 4 (Statistical inference). *Let the estimated variance be*

$$\widehat{\mathrm{V}}_{\mathrm{RIVW}} = \frac{\sum_{j\in S_\lambda} (\widehat{\Gamma}_j \widehat{\gamma}_{j,\mathrm{RB}} - \widehat{\beta}_{\mathrm{RIVW}}(\widehat{\gamma}_{j,\mathrm{RB}}^2 - \widehat{\sigma}_{X_j,\mathrm{RB}}^2))^2/\sigma_{Y_j}^4}{(\sum_{j\in S_\lambda} (\widehat{\gamma}_{j,\mathrm{RB}}^2 - \widehat{\sigma}_{X_j,\mathrm{RB}}^2)/\sigma_{Y_j}^2)^2}, \tag{5}$$

then a level $1-\alpha$ confidence interval can be constructed as

$$\left[\widehat{\beta}_{\mathrm{RIVW}} - \Phi^{-1}\left(1-\frac{\alpha}{2}\right)\sqrt{\widehat{\mathrm{V}}_{\mathrm{RIVW}}}\,, \widehat{\beta}_{\mathrm{RIVW}} + \Phi^{-1}\left(1-\frac{\alpha}{2}\right)\sqrt{\widehat{\mathrm{V}}_{\mathrm{RIVW}}}\right].$$

Because our RIVW estimator resembles the "slope coefficient" obtained from a regression of $\widehat{\Gamma}_j$ on $\widehat{\gamma}_{j,\mathrm{RB}}$, the proposed variance estimator can be motivated from the use of "regression

residuals" (See Section 4.3 for detailed discussion). Thanks to the regression interpretation of the RIVW estimator, and because the presence of balanced horizontal pleiotropy simply adds centered noise to the SNP-outcome association estimates $\widehat{\Gamma}_j$, we shall demonstrate in Section 5.1 that our RIVW estimator remains asymptotically unbiased and normally distributed under the more general Assumption 1′, and our variance estimator is valid without further modification.

In Section 5.2, we consider an extension to our RIVW estimator, where we further "smooth" over the random IV selection, leading to an estimator that reweights the SNP-exposure association estimates. Finally, Section 5.3 provides a first attempt to correct LD clumping bias by generalizing our randomized instrument selection and Rao–Blackwellization to a correlated instrumental variables setting.

4. Theoretical investigations. In this section, we first provide theoretical justifications for our winner's curse removal procedure. We then demonstrate the statistical validity of our inference approach by showing that the RIVW estimator is asymptotically normally distributed under mild conditions (Theorem 1). Finally, we show that our variance estimator is consistent for the leading variance of the RIVW estimator (Theorem 2).

4.1. *Notation and assumptions.* Before discussing the theoretical results in detail, we revisit and introduce some notation and assumptions adopted in the paper. Recall that the jth genetic instrument/SNP is selected if and only if $S_j = |\frac{\widehat{\gamma}_j}{\sigma_{X_j}} + Z_j| - \lambda > 0$. The collection of selected instruments is denoted by $\mathcal{S}_\lambda = \{j : S_j > 0\}$. The asymptotic regime we consider involves $p \to \infty$ and $\lambda \to \infty$. This captures the phenomena of "many instruments/SNPs" in MR studies. On the other hand, the requirement that $n_X, n_Y \to \infty$ is implicit, as we follow the literature and assume that the estimated SNP-exposure and SNP-outcome associations are normally distributed (Assumption 1). We also introduce the following notation for probabilistic ordering. For two (sequences) of random variables, A and B, write $A \precsim_{\mathrm{p}} B$ if the ratio A/B is asymptotically bounded in probability. The strict relation, $A \prec_{\mathrm{p}} B$, implies that $A/B \xrightarrow{\mathrm{p}} 0$. Finally, $A \asymp_{\mathrm{p}} B$ indicates both $A \precsim_{\mathrm{p}} B$ and $B \precsim_{\mathrm{p}} A$.

To facilitate discussion, we introduce two additional quantities: the number of selected instruments as $p_\lambda = |\mathcal{S}_\lambda|$, and an overall measure of instrument strength after selection:

$$\kappa_\lambda = \frac{1}{p_\lambda} \sum_{j \in \mathcal{S}_\lambda} \left(\frac{\gamma_j}{\sigma_{X_j}} \right)^2.$$

We note that both p_λ and κ_λ are random, as they depend on the set of selected instruments. In addition, κ_λ is typically unknown, because it depends on the unobserved SNP-exposure associations, γ_j.

In addition to Assumption 1, we employ the following conditions.

ASSUMPTION 2 (Instrument selection). The cutoff value satisfies $\lambda \to \infty$.

ASSUMPTION 3 (No dominant instrument). The true instrument effect satisfies

$$\max_{j \in S_\lambda} \gamma_j^2 \Big/ \left(\sum_{j \in S_\lambda} \gamma_j^2 \right) \xrightarrow{\mathrm{p}} 0.$$

Assumption 2 requires the cutoff value to diverge, which is quite plausible given that λ is usually of the order $\sqrt{\log p}$ to account for multiple testing; see footnote 2. Assumption 3 requires that, after selection, no instrument has a "dominating effect." It helps rule out the

extreme scenario where only a handful of genetic variants are relevant. We will employ this condition when establishing the asymptotic normality of our estimator.[4]

4.2. *Properties of the Rao–Blackwellized estimator.* Our first result verifies that the proposed framework breaks the winner's curse: the construction, $\widehat{\gamma}_{j,\mathtt{RB}}$, is unbiased for γ_j conditional on the selection event $j \in \mathcal{S}_\lambda$.

LEMMA 1 (Instrument selection bias correction). *Under Assumption 1,* $\mathbb{E}[\widehat{\gamma}_{j,\mathtt{RB}} | j \in \mathcal{S}_\lambda] = \gamma_j$.

As an immediate result of this lemma, $\widehat{\gamma}_{j,\mathtt{RB}}$ is preferred to $\widehat{\gamma}_j$ since it helps lift the winner's curse by correcting instrument selection bias. To provide some intuition, the unbiasedness of $\widehat{\gamma}_{j,\mathtt{RB}}$ stems from properties of the initial estimator, $\widehat{\gamma}_{j,\mathtt{ini}}$, which we introduced in the previous section. In particular, the initial estimator is unbiased due to its independence from the selection event. Therefore, because $\widehat{\gamma}_{j,\mathtt{RB}} = \mathbb{E}[\widehat{\gamma}_{j,\mathtt{ini}} | \widehat{\gamma}_j, j \in \mathcal{S}_\lambda]$, iterated expectation implies $\mathbb{E}[\widehat{\gamma}_{j,\mathtt{RB}} | j \in \mathcal{S}_\lambda] = \mathbb{E}[\widehat{\gamma}_{j,\mathtt{ini}} | j \in \mathcal{S}_\lambda] = \gamma_j$. Finally, the specific expression of $\widehat{\gamma}_{j,\mathtt{RB}}$ can be found by explicit calculation, which is available in the Supplementary Material [25]. On a related note, $\widehat{\gamma}_{j,\mathtt{RB}}$ also has an advantage over the initial estimator: although the initial estimator is unbiased, $\widehat{\gamma}_{j,\mathtt{RB}}$ has a smaller variance thanks to Rao–Blackwellization.

As we will discuss in the next subsection, correcting measurement error bias requires the variance of the estimated SNP-exposure associations. In conventional MR analyses without instrument selection or when the selection step is carried out in an independent third sample, this information is directly available from $\sigma^2_{X_j}$. In our framework, however, both the selection step and the subsequent Rao–Blackwellization may affect the variance of the selected instrument effects. That is, $\mathbb{V}[\widehat{\gamma}_{j,\mathtt{RB}} | j \in \mathcal{S}_\lambda] \neq \sigma^2_{X_j}$ in general. We characterize this conditional variance in the following lemma.

LEMMA 2 (Variance of $\widehat{\gamma}_{j,\mathtt{RB}}$). *Under Assumption 1, conditioning on the selection event, the variance of* $\widehat{\gamma}_{j,\mathtt{RB}}$ *is*

$$\sigma^2_{X_j,\mathtt{RB}} = \sigma^2_{X_j}\Bigg(1 - \frac{1}{\eta \mathbb{P}[S_j > 0]} \int_{-\infty}^{\infty} y\phi(y)\big(\phi\big(B_{j,+}(y)\big) - \phi\big(B_{j,-}(y)\big)\big)\,\mathrm{d}y + \frac{1}{\eta^2 \mathbb{P}[S_j > 0]} \int_{-\infty}^{\infty} \phi(y) \frac{(\phi(B_{j,+}(y)) - \phi(B_{j,-}(y)))^2}{1 - \Phi(B_{j,+}(y)) + \Phi(B_{j,-}(y))}\,\mathrm{d}y\Bigg),$$

where we define $B_{j,\pm}(y) = -(\frac{\gamma_j}{\sigma_{X_j}\eta} + \frac{y}{\eta}) \pm \frac{\lambda}{\eta}$.

The above expression clearly demonstrates that the conditional variance depends on the true SNP-exposure associations, γ_j, in a nonlinear and complicated way. As a result, a direct plug-in estimator of the conditional variance is not expected to perform well.

To provide some intuition for our variance estimator in Step 3, we recall that the Rao–Blackwellized instrument effects are obtained as conditional expectations. Therefore, instead of characterizing its variance directly, one can employ the conditional variance decomposition formula, that is,

$$\sigma^2_{X_j,\mathtt{RB}} = \mathbb{V}[\widehat{\gamma}_{j,\mathtt{ini}} | j \in \mathcal{S}_\lambda] - \mathbb{E}\big[\mathbb{V}[\widehat{\gamma}_{j,\mathtt{ini}} | \widehat{\gamma}_j, j \in \mathcal{S}_\lambda] | j \in \mathcal{S}_\lambda\big].$$

[4]To better explain what Assumption 3 entails, consider the simulation setting in Example 1, and ignore the instrument selection step for simplicity. Then it is easy to show that $\max_{1 \le j \le p} \gamma_j^2 \precsim \varepsilon_x^2 \log(p\pi_x)$, while on the other hand $\sum_{j=1}^p \gamma_j^2 \asymp \varepsilon_x^2 p\pi_x$.

Due to independence, the first term on the right-hand side is simply the unconditional variance of the initial estimator: $\mathbb{V}[\widehat{\gamma}_{j,\mathtt{ini}}|j \in \mathcal{S}_\lambda] = \mathbb{V}[\widehat{\gamma}_{j,\mathtt{ini}}]$. The second term, however, is much more difficult to obtain. In fact, if one computes the outer expectation, then it reduces to the formula in Lemma 2. Our insight is that we define the variance estimator without explicitly computing the outer expectation. To be more precise, $\widehat{\sigma}^2_{X_j,\mathtt{RB}}$ is defined from

$$\widehat{\sigma}^2_{X_j,\mathtt{RB}} = \mathbb{V}[\widehat{\gamma}_{j,\mathtt{ini}}] - \mathbb{V}[\widehat{\gamma}_{j,\mathtt{ini}}|\widehat{\gamma}_j, j \in \mathcal{S}_\lambda].$$

Due to iterative expectation, the above is unbiased for $\sigma^2_{X_j,\mathtt{RB}}$ (conditional on the selection event), but is generally inconsistent. Fortunately, for measurement error bias correction, we only need to estimate the aggregate variance, $\sum_{j\in\mathcal{S}_\lambda} \sigma^2_{X_j,\mathtt{RB}}$. We further discuss this issue below.

4.3. *Measurement error bias correction.* The preceding section illustrates how our approach breaks the winner's curse. We now turn to discussing how our RIVW estimator also removes measurement error bias, which arises due to the randomness in $\widehat{\gamma}_j$ (or $\widehat{\gamma}_{j,\mathtt{RB}}$). The measurement error bias issue can be heuristically understood by viewing the IVW estimator as an estimated regression coefficient. In a hypothetical case where $\widehat{\gamma}_j$ are measured without any error (i.e., $\widehat{\gamma}_j = \gamma_j$, also referred to as no measurement error assumption in the literature, [5]), an "oracle" IVW estimator would regress $\widehat{\Gamma}_j$ on γ_j (after attaching a weight $1/\sigma_{Y_j}$ to them) and obtain the slope coefficient as an estimator for β. Such an oracle IVW estimator is unbiased for β. In practice, however, we only observe γ_j with a random error, which means that the slope estimate from regressing $\widehat{\Gamma}_j$ on $\widehat{\gamma}_j$ will be biased. Such a bias is also referred to as regression attenuation, or weak instrument bias in the MR literature [33].

Our RIVW removes measurement error bias following a similar logic to the strategy adopted by [41], who propose to replace $\widehat{\gamma}_j^2$ with $\widehat{\gamma}_j^2 - \sigma^2_{X_j}$ in the denominator of the conventional IVW estimator. Despite taking a similar form, we propose the use of a different variance estimator, $\widehat{\sigma}^2_{X_j,\mathtt{RB}}$ (Step 3), because the Rao–Blackwellization step can affect the variance of the instrument effects. That is, $\mathbb{V}[\widehat{\gamma}_{j,\mathtt{RB}}|j \in \mathcal{S}_\lambda] \neq \sigma^2_{X_j} = \mathbb{V}[\widehat{\gamma}_j]$. Although Lemma 2 provides a formula for the variance of the Rao–Blackwellized estimator, $\widehat{\gamma}_{j,\mathtt{RB}}$, the result is not directly amenable to implementation as the variance formula depends on the unknown γ_j nonlinearly. Nevertheless, thanks to the special form of the RIVW estimator, we only need to correctly estimate the "aggregated variance," $\sum_{j\in S_\lambda} \sigma^2_{X_j,\mathtt{RB}}$, to remove measurement error bias, as the next lemma demonstrates.

LEMMA 3 (Variance estimation). *Assume Assumptions 1 and 2 hold. Then* $\mathbb{E}[\widehat{\sigma}^2_{X_j,\mathtt{RB}}|\mathcal{S}_\lambda] = \sigma^2_{X_j,\mathtt{RB}}$, *and*

$$\left|\sum_{j\in S_\lambda} \widehat{\sigma}^2_{X_j,\mathtt{RB}} - \sum_{j\in S_\lambda} \sigma^2_{X_j,\mathtt{RB}}\right| \precsim_{\mathrm{p}} \sqrt{p_\lambda}\nu^2\lambda.$$

We will postpone the discussion on the above probabilistic order to the next subsection after we introduce the main asymptotic normality result.

4.4. *Properties of the RIVW estimator.* We are now in a position to describe the asymptotic behavior of our RIVW estimator.

THEOREM 1. (Asymptotic normality) *Assume Assumptions 1, 2 and 3 hold,* $p_\lambda \xrightarrow{\mathrm{p}} \infty$ *and* $\kappa_\lambda/\lambda^2 \xrightarrow{\mathrm{p}} \infty$. *Then conditional the selection event* $\mathcal{S}_\lambda$,

$$\mathsf{V}^{-\frac{1}{2}}_{\mathtt{RIVW}}(\widehat{\beta}_{\mathtt{RIVW}} - \beta) \xrightarrow{\mathcal{D}} \mathcal{N}(0,1),$$

where

$$\mathsf{V}_{\mathtt{RIVW}} = \frac{\mathbb{V}[\sum_{j\in S_\lambda} u_{j,\mathtt{RIVW}}/\sigma_{Y_j}^2 | S_\lambda]}{(\sum_{j\in S_\lambda} \gamma_j^2/\sigma_{Y_j}^2)^2} \asymp_{\mathrm{p}} \frac{1}{p_\lambda \kappa_\lambda},$$

$$u_{j,\mathtt{RIVW}} = \gamma_j(u_{Y_j} - \beta u_{X_j,\mathtt{RB}}) + (u_{X_j,\mathtt{RB}} u_{Y_j} - \beta(u_{X_j,\mathtt{RB}}^2 - \widehat{\sigma}_{X_j,\mathtt{RB}}^2)),$$

$u_{Y_j} = \widehat{\Gamma}_j - \Gamma_j$ *and* $u_{X_j,\mathtt{RB}} = \widehat{\gamma}_{j,\mathtt{RB}} - \gamma_j$.

We provide a decomposition below to demonstrate the intuition behind the asymptotic unbiasedness of our RIVW estimator and defer the formal proof to the Supplementary Material [25]. This decomposition also sheds light on the mechanism of our measurement error bias correction and how the error from variance estimation features in the asymptotic analysis.

To start, we rewrite our RIVW estimator as

$$\widehat{\beta}_{\mathtt{RIVW}} = \beta + \frac{\sum_{j\in\mathcal{S}_\lambda} u_{j,\mathtt{RIVW}}/\sigma_{Y_j}^2}{\sum_{j\in\mathcal{S}_\lambda}(\widehat{\gamma}_{j,\mathtt{RB}}^2 - \widehat{\sigma}_{X_j,\mathtt{RB}}^2)/\sigma_{Y_j}^2}.$$

Then from Lemmas 1 and 2, it should be clear that $u_{j,\mathtt{RIVW}}$ has a zero mean, which implies that our RIVW estimator is asymptotically unbiased. Lemma 3 further helps demonstrate that the denominator in our RIVW estimator converges to $\sum_{j\in\mathcal{S}_\lambda} \gamma_j^2/\sigma_{Y_j}^2$.

It is not surprising that the asymptotic variance of our RIVW estimator takes a complicated form, as the variance of $u_{j,\mathtt{RIVW}}$ involves higher-order moments of the Rao–Blackwellized instrument effects, $\widehat{\gamma}_{j,\mathtt{RB}}$. We employ the condition $\kappa_\lambda/\lambda^2 \xrightarrow{\mathrm{p}} \infty$ to simplify the analysis. In particular, this assumption implies that the leading variance stems from the first term in the expression of $u_{j,\mathtt{RIVW}}$. This condition seems mild and quite plausible in applications since instrument selection will help improve the overall IV strength. See the Supplemental Material for additional discussions and results on the probabilistic order of κ_λ.

The previous decomposition also motivates our regression residual based variance estimator. In particular, we note that the "error term," $u_{j,\mathtt{RIVW}}$, can also be written as $\widehat{\Gamma}_j\widehat{\gamma}_{j,\mathtt{RB}} - \beta(\widehat{\gamma}_{j,\mathtt{RB}}^2 - \widehat{\sigma}_{X_j,\mathtt{RB}}^2)$. As a result, in our standard error construction, we replace the unknown terms by their estimates. The following theorem establishes the consistency of our standard error.

THEOREM 2. (Consistent variance estimation) *Assume Assumptions* 1, 2 *and* 3 *hold*, $p_\lambda \xrightarrow{\mathrm{p}} \infty$ *and* $\kappa_\lambda/\lambda^2 \xrightarrow{\mathrm{p}} \infty$. *Then*

$$\frac{\widehat{\mathsf{V}}_{\mathtt{RIVW}}}{\mathsf{V}_{\mathtt{RIVW}}} \xrightarrow{\mathrm{p}} 1.$$

5. Extensions.

5.1. *Balanced horizontal pleiotropy.* Balanced pleiotropy (also known as systematic pleiotropy) refers to the scenario that the exclusion restriction assumption breaks down and the genetic variants have equal chances to affect the outcome either positively or negatively [17, 39]. For this extension, we consider the following model that is commonly adopted in the literature:

$$Y = \beta X + \sum_{j=1}^{p} \alpha_j G_j + \theta_Y U + E_Y, \quad \text{for each } j = 1, \ldots, p. \tag{6}$$

In the above, α_j captures the pleiotropic effect of SNP j on the outcome Y. Formally, we require the following.

ASSUMPTION 1 (Measurement error model under balanced horizontal pleiotropy). (i) For any $j \neq j'$, the pairs $(\widehat{\Gamma}_j, \widehat{\gamma}_j)$ and $(\widehat{\Gamma}_{j'}, \widehat{\gamma}_{j'})$ are mutually independent.

(ii) For each j,

$$\begin{bmatrix} \widehat{\Gamma}_j \\ \widehat{\gamma}_j \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \beta\gamma_j + \alpha_j \\ \gamma_j \end{bmatrix}, \begin{bmatrix} \sigma_{Y_j}^2 & 0 \\ 0 & \sigma_{X_j}^2 \end{bmatrix} \right).$$

In addition, there exists some $\nu \to 0$, such that $\{\sigma_{Y_j}/\nu, \sigma_{X_j}/\nu : 1 \leq j \leq p\}$ are uniformly bounded and bounded away from zero.

(iii) The pleiotropic effects, $\alpha_1, \ldots, \alpha_p$, are mutually independent and follow a distribution with mean 0, variance τ^2 and bounded third moment. In addition, τ/ν is bounded.

Because the pleiotropic effects follow a centered distribution, they do not introduce any additional bias into our RIVW estimator. In fact, as the reader will gather from checking our proof, the presence of balanced pleiotropic effects merely inflates the variance of the RIVW estimator. That is, our RIVW estimator remains asymptotically unbiased and normal under the more general Assumption 1′. Thanks to our regression residual based variance estimator in (5), the statistical inference procedure in Step 4 remains valid in this setting. In other words, our variance estimator can be directly applied to the balanced horizontal pleiotropy scenario without any modification. This feature allows practitioners to conduct statistical analysis without switching among various variance estimators under different assumptions on the existence of balanced horizontal pleiotropy. We also note that this balanced horizontal pleiotropy assumption can be viewed as a sufficient condition of Instrument Strength Independent on Direct Effect (InSIDE) assumption, requiring that the exposure effects of individual SNPs are independent of their pleiotropic effects on the outcome [11].

We summarize this conclusion in the following corollary.

COROLLARY 1. (Validity of the RIVW estimator under balanced horizontal pleiotropy) *The results in Theorems 1 and 2 continue to hold if Assumption 1 is replaced by 1′.*

5.2. *Smoothing over randomization.* Our RIVW estimator is constructed based on the selected instruments. While the selection step might be of interest, as it allows the researcher to further evaluate the validity of the selected instruments and the underlying causal mechanism, we do recognize that introducing the pseudo SNP-exposure effects, Z_j, into the procedure may inflate the variance of the RIVW estimator. In this subsection, we consider an extension where each instrument is weighted by its (conditional) probability of being selected. Formally, we define

$$\widehat{\omega}_j = \mathbb{P}[j \in \mathcal{S}_\lambda | \widehat{\gamma}_j] = 1 - \Phi(A_{j,+}) + \Phi(A_{j,-}),$$

and we consider the following smoothed RIVW (sRIVW) estimator:

$$\widehat{\beta}_{\mathtt{sRIVW}} = \frac{\sum_{j=1}^p \widehat{\Gamma}_j \widehat{\gamma}_{j,\mathrm{RB}} \widehat{\omega}_j / \sigma_{Y_j}^2}{\sum_{j=1}^p (\widehat{\gamma}_{j,\mathrm{RB}}^2 - \widehat{\sigma}_{X_j,\mathrm{RB}}^2) \widehat{\omega}_j / \sigma_{Y_j}^2}.$$

To conserve space, formal results, such as asymptotic normality and consistent variance estimation, are collected in the Supplementary Material [25]. In what follows, we intuitively demonstrate the properties of the sRIVW estimator.

To start, we consider the expectations of the numerator and denominator of the sRIVW estimator. By iterative expectation, it is straightforward to show that

$$\mathbb{E}[\widehat{\gamma}_{j,\mathrm{RB}} \widehat{\omega}_j] = \mathbb{E}[\widehat{\gamma}_{j,\mathrm{RB}} \mathbf{1}_{j \in \mathcal{S}_\lambda}] = \gamma_j \omega_j, \quad \omega_j = 1 - \Phi\left(\frac{\lambda - \frac{\gamma_j}{\sigma_{X_j}}}{\sqrt{1+\eta^2}} \right) + \Phi\left(\frac{-\lambda - \frac{\gamma_j}{\sigma_{X_j}}}{\sqrt{1+\eta^2}} \right),$$

and ω_j is the unconditional probability of the jth instrument being selected. As a result, the numerator of the sRIVW estimator has expectation $\beta \sum_{j=1}^{p} \gamma_j^2 \omega_j / \sigma_{Y_j}^2$. From Lemma 3 in the previous section, $\widehat{\sigma}_{X_j,\mathtt{RB}}^2$ is unbiased for $\sigma_{X_j,\mathtt{RB}}^2$ conditional on the selection event, which implies that the denominator of our sRIVW estimator has expectation $\sum_{j=1}^{p} \gamma_j^2 \omega_j / \sigma_{Y_j}^2$. As a result, we expect that our sRIVW approach also delivers accurate estimates of the causal parameter. Evidence on its finite sample performance is provided in Section 6.

To close this subsection, we provide in the following an estimator for the variance of the sRIVW estimator:

$$\widehat{\mathtt{V}}_{\mathtt{sRIVW}} = \frac{\sum_{j=1}^{p} (\widehat{\Gamma}_j \widehat{\gamma}_{j,\mathtt{RB}} \widehat{\omega}_j - \widehat{\beta}_{\mathtt{sRIVW}} (\widehat{\gamma}_{j,\mathtt{RB}}^2 - \widehat{\sigma}_{X_j,\mathtt{RB}}^2) \widehat{\omega}_j)^2 / \sigma_{Y_j}^4}{(\sum_{j=1}^{p} (\widehat{\gamma}_{j,\mathtt{RB}}^2 - \widehat{\sigma}_{X_j,\mathtt{RB}}^2) \widehat{\omega}_j / \sigma_{Y_j}^2)^2}.$$

5.3. *Correlated instrumental variables*. With correlated instruments, we consider two remedies to relax the independence assumption adopted in Assumption 1. The first remedy relies on a revised LD pruning procedure (referred to as sigma-based LD pruning), and our proposed RIVW estimator along with its variance estimator remains valid without further modification. As opposed to the classical clumping that sorts SNPs with their p-values, sigma-based LD pruning sorts the correlated genetic variants by their standard deviations, which are often assumed to be known in two-sample MR with summary data. We then take the first SNP (i.e., the SNP with the smallest standard deviation) and remove all SNPs in linkage disequilibrium with this first SNP. Our procedure goes on with the next SNP with the smallest standard deviation that has not been removed yet. In the end, the sigma-based LD pruning delivers a collection of independent instruments. In addition, because this revised sigma-based LD pruning procedure does not touch the information in the estimated SNP-exposure effect sizes, no additional selection bias is introduced. We compare the performance of this sigma-based LD pruning with classical LD pruning and p-value based LD clumping in the Supplementary Material [25].

We propose another remedy relying on a carefully crafted initial estimator that is independent with the classical p-value based LD clumping selection event. We then propose an updated Rao–Blackwellized estimator customized to LD clumping events when IVs are correlated. To conserve space in the main paper, we leave the detailed discussion and implementation of this estimator to the Supplementary Material [25].

6. Simulation.

6.1. *Choice of tuning parameters*. The RIVW estimator depends on the selection threshold (λ) as well as the level of randomization introduced via the pseudo SNPs (η). In this section, we discuss practical considerations in choosing these two tuning parameters. To begin with, the choice of the cutoff value λ is subjective but needs to be large enough for the selected SNPs to be informative IVs. Because the RIVW estimator lifts the winner's curse and removes measurement error bias, our framework allows for including more genetic instruments with moderate effects compared with traditional MR analyses. Therefore, we recommend $\lambda = \Phi^{-1}(1-\alpha/2)$ with $\alpha = 5 \times 10^{-5}$ as the default setting for our RIVW estimator.

In terms of choosing η, our theory only requires it to be bounded away from infinity and zero. The reason is quite straightforward: if η tends to zero, the pseudo SNPs are asymptotically negligible, which defeats their purposes; on the other hand, if η diverges to infinity, the pseudo SNPs mask the original signals $\widehat{\gamma}_j / \sigma_{X_j}$ and all instruments will be selected as a result. Rather than studying an optimal choice of η via theoretical derivations, extensive simulation evidence suggests that our procedure is not very sensitive to this choice. In Figure 3, we provide the Monte Carlo bias and standard deviation of our RIVW estimator for different values

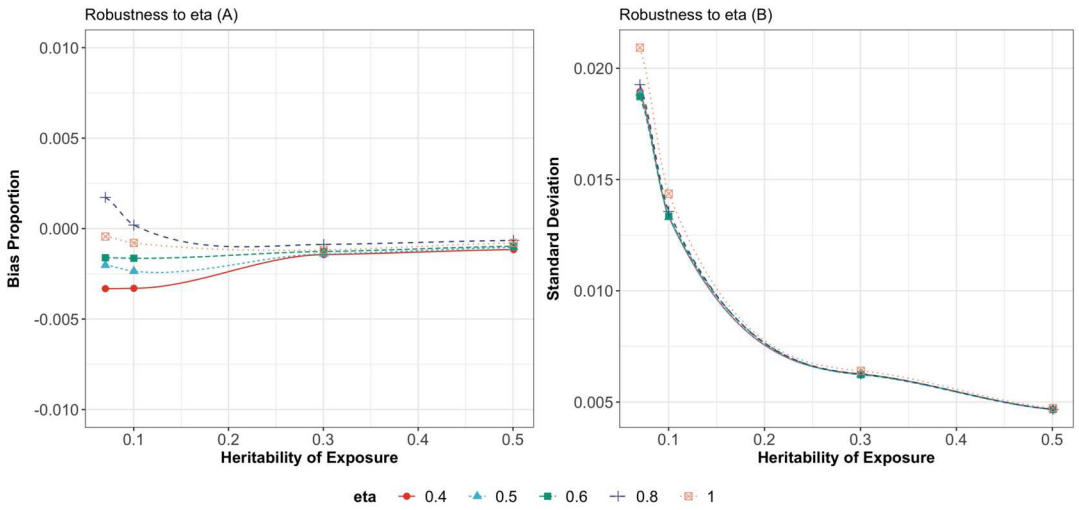

FIG. 3. *Monte Carlo samples are generated following the mixture model in Example* 1 *with* $\rho = 1$ *and* $\pi_x = \pi_y = 0.005$. *We vary* $\varepsilon_x^2 = \tau^2$ *in the set* $\{7 \times 10^{-5}, 1 \times 10^{-4}, 3 \times 10^{-4}, 5 \times 10^{-4}\}$.

of η in the set $\{0.4,\ 0.5,\ 0.6,\ 0.8,\ 1\}$. We observe that the bias is negligible across different choices of η. This is in line with our theoretical results, since the Rao–Blackwellization Step 2 removes instrument selection bias. Moreover, we observe that the standard deviation of the RIVW estimator is not sensitive to η either, as the five curves in panel (B) corresponding to different values of η almost coincide. We fix $\eta = 0.5$ as it leads to satisfactory performance for the RIVW estimator across a range of simulation settings.

6.2. *Simulations setup and results*. In this section, we compare the performance of our RIVW estimator with the classical two-/three-sample IVW, the debiased IVW (dIVW) [41], and the two-/three-sample RAPS estimator [44] in simulation studies. We adopt the simulation setup in Example 1. In particular, we set $\rho = 1$, $\pi_x = \pi_y$, and $\varepsilon_x^2 = \tau^2$ to reduce the number of free parameters. Additional simulation results are available in the Supplementary Material [25]. The parameter values used to generate Monte Carlo samples are given in corresponding tables. Our results are based on 2000 Monte Carlo repetitions.

Following common practices, the cutoff value λ is either 5.45 (corresponding to the significance threshold 5×10^{-8}) for the two-sample IVW estimators, or 0 (using all available SNPs) for the dIVW estimator. As three-sample MR involves a third independent sample for IV selection and are immune to winner's curse bias, we also consider the cutoff value $\lambda = 4.06$ (corresponding to the significant threshold 5×10^{-5}). Note that in three-sample MR analyses, we generate a new independent exposure data with the same sample size for IV selection.

We report our simulation results with six measures: "$\widehat{\beta}$" (average causal effect estimate across Monte Carlo samples), "monte SD" (Monte Carlo standard deviation), "SE" (average standard error across Monte Carlo samples), "CP" (average coverage probability of the 95% confidence interval), "length" (average length of the 95% confidence interval), and "# IVs" (number of selected IVs, averaged across Monte Carlo samples). Table 2 summarizes the performance of various estimators under different settings, which we discuss below.

For the two-sample IVW estimator, our simulation results confirm that it is biased toward zero due to winner's curse bias and measurement error bias, resulting in confidence intervals that undercover the true causal effect. In the Supplementary Material (Table S.8 of [25]), we further demonstrate that the two-sample IVW estimator with a liberal cutoff value ($\lambda = 4.06$) yields a larger bias than that with a stringent cutoff ($\lambda = 5.45$).

TABLE 2

Simulation results under different settings. We set the true effect $\beta = 0.2$. The columns are: *estimated effect size* ($\widehat{\beta}$), *Monte Carlo standard deviation* (*monte SD*), *average standard error* (*SE*), *empirical coverage probability of* 95% *confidence intervals* (*CP*), *average confidence interval length* (*length*) *and number of selected instruments* (*# IVs*). h_x^2 *and* h_y^2 *are the heritability of exposure and outcome*

	$\widehat{\beta}$	Monte SD	SE	CP	Length	# IVs
Low heritability: $\pi_x = \pi_y = 0.002$, $\varepsilon_x^2 = \tau^2 = 1 \times 10^{-4}$, $h_x^2 = 0.040$, $h_y^2 = 0.042$						
IVW ($\lambda = 5.45$)	0.182	0.023	0.023	0.865	0.090	40
dIVW ($\lambda = 0$)	0.209	0.125	0.123	0.959	0.484	200,000
RIVW ($\lambda = 4.06$, $\eta = 0.5$)	0.200	0.022	0.022	0.951	0.087	148
sRIVW ($\lambda = 4.06$, $\eta = 0.5$)	0.200	0.021	0.021	0.947	0.082	–
Three-sample IVW ($\lambda = 5.45$)	0.196	0.025	0.025	0.950	0.097	40
Three-sample IVW ($\lambda = 4.06$)	0.200	0.020	0.021	0.954	0.081	99
Three-sample RAPS ($\lambda = 5.45$)	0.200	0.026	0.028	0.962	0.108	40
Three-sample RAPS ($\lambda = 4.06$)	0.200	0.021	0.022	0.958	0.085	99
Three-sample dIVW ($\lambda = 5.45$)	0.201	0.026	0.026	0.954	0.101	40
Three-sample dIVW ($\lambda = 4.06$)	0.200	0.020	0.021	0.954	0.081	99
Medium heritability: $\pi_x = \pi_y = 0.01$, $\varepsilon_x^2 = \tau^2 = 1 \times 10^{-4}$, $h_x^2 = 0.200$, $h_y^2 = 0.208$						
IVW ($\lambda = 5.45$)	0.182	0.011	0.010	0.566	0.040	200
dIVW ($\lambda = 0$)	0.201	0.025	0.024	0.941	0.096	200,000
RIVW ($\lambda = 4.06$, $\eta = 0.5$)	0.200	0.010	0.009	0.944	0.037	509
sRIVW ($\lambda = 4.06$, $\eta = 0.5$)	0.200	0.009	0.009	0.947	0.036	–
Three-sample IVW ($\lambda = 5.45$)	0.195	0.011	0.011	0.917	0.043	200
Three-sample IVW ($\lambda = 4.06$)	0.200	0.009	0.009	0.956	0.036	452
Three-sample RAPS ($\lambda = 5.45$)	0.199	0.012	0.012	0.951	0.047	200
Three-sample RAPS ($\lambda = 4.06$)	0.200	0.009	0.010	0.955	0.037	452
Three-sample dIVW ($\lambda = 5.45$)	0.200	0.011	0.011	0.947	0.045	200
Three-sample dIVW ($\lambda = 4.06$)	0.200	0.009	0.009	0.956	0.036	452
High heritability: $\pi_x = \pi_y = 0.01$, $\varepsilon_x^2 = \tau^2 = 3 \times 10^{-4}$, $h_x^2 = 0.600$, $h_y^2 = 0.624$						
IVW ($\lambda = 5.45$)	0.193	0.005	0.004	0.688	0.017	655
dIVW ($\lambda = 0$)	0.200	0.010	0.009	0.928	0.035	200,000
RIVW ($\lambda = 4.06$, $\eta = 0.5$)	0.200	0.005	0.004	0.952	0.018	993
sRIVW ($\lambda = 4.06$, $\eta = 0.5$)	0.200	0.005	0.004	0.952	0.017	–
Three-sample IVW ($\lambda = 5.45$)	0.197	0.005	0.005	0.908	0.018	655
Three-sample IVW ($\lambda = 4.06$)	0.200	0.004	0.004	0.955	0.017	943
Three-sample RAPS ($\lambda = 5.45$)	0.200	0.005	0.005	0.959	0.019	655
Three-sample RAPS ($\lambda = 4.06$)	0.200	0.005	0.005	0.954	0.018	943
Three-sample dIVW ($\lambda = 5.45$)	0.200	0.005	0.005	0.960	0.018	655
Three-sample dIVW ($\lambda = 4.06$)	0.200	0.004	0.004	0.955	0.017	943

For the two-sample dIVW estimator, we have followed the recommendation of [41] and used all available IVs (i.e., $\lambda = 0$) to estimate β. While dIVW indeed provides accurate causal effect estimates when the average IV strength is relatively high (medium and high heritability cases), it exhibits larger variance among the considered estimators in the presence of too many weak instruments (e.g., in the low heritability situation).

For the three-sample IVW estimator, we observe that it tends to produce slightly biased causal effect estimates due to its failure to account for measurement error bias. By correcting measurement error bias, both the three-sample dIVW estimator and three-sample RAPS estimator yield accurate causal effect estimates. Furthermore, three-sample RAPS with liberal cutoff values ($\lambda = 4.06$) yields narrower confidence intervals than that with a stringent cutoff value ($\lambda = 5.45$). This indicates that a more liberal cutoff can be more preferable when both winner's curse bias and measurement error bias are fully removed.

By fully correcting measurement error bias and winner's curse bias, our RIVW estimator provides accurate causal effect estimates, and the resulting confidence intervals have coverage probabilities close to the nominal 95% level. Furthermore, because the RIVW estimator is able to accommodate a liberal cutoff value ($\lambda = 4.06$) and incorporate more IVs to estimate the causal effect, its confidence intervals have similar lengths to those produced by three-sample dIVW/RAPS. We also note that standard errors provided by our formula are close to the Monte Carlo standard deviations, validating our theoretical investigations (Theorem 2). Finally, the sRIVW estimator (proposed in Section 5.2) that smooths out the effect of IV selection yields similar results. We do not report the number of selected instruments for the sRIVW estimator, as it employs a reweighting strategy instead of IV selection.

We provide additional simulation evidence in the Supplementary Material [25] for a wider range of model specifications, such as different sample sizes for the exposure and outcome GWAS data (Tables S.5 and S.6), varying the sample size for each SNP (Table S.7), and $\beta = 0$ (i.e., no causal effect between the exposure and the outcome, Table S.9). In Section VIII.2 of the Supplementary Material, we also conduct simulation studies in the presence of balanced horizontal pleiotropy (by setting $1 - \rho > 0$ in Example 1), where we observe that the proposed RIVW estimator achieves high detection power and does not exhibit noticeable size distortion. In summary, the simulation results confirm that our RIVW estimator continues to deliver accurate causal effect estimation and valid statistical inference.

7. Real data applications. We conduct two real data analyses. First, the same-trait type analysis [44] demonstrates that the performance of the proposed RIVW estimator is in line with our simulation evidence and theoretical investigations. Second, the study of the relationship between BMI and COVID-19 severity provides new perspectives on the benefit of body weight management.

7.1. *Data harmonization*. We harmonize the data through the following steps. First, genetic variants that are not available in the outcome data set are excluded. Second, we select independent genetic variants that have no linkage disequilibrium (R-Squared < 0.001 with an extension of 10,000 Kb in the genome, which is suggested by the `twosampleMR` package (version 0.5.5; [18]) and has been widely adopted in applied MR studies) with other selected genetic variants through the revised sigma-based pruning procedure introduced in Section 5.3. Specifically, when genetic variants are in linkage disequilibrium, we select the variant with the smallest standard deviation of the SNP-exposure association. We employ this revised sigma-based pruning procedure because standard clumping, which is based on p-value comparisons, introduces a different type of selection bias; see [31] for related discussion, and Tables S.1 and S.2 in the Supplementary Material [25]. Third, by leveraging allele frequency information, we infer the strand direction of ambiguous SNPs and harmonize exposure-outcome data sets by using the `twosampleMR` package. We use the default setting with $\lambda = 4.06$ and $\eta = 0.5$ for our proposed RIVW estimator, and set $\lambda = 0$ and $\lambda = 5.45$ for the dIVW and IVW estimators, respectively.

7.2. *Same trait analyses*. To investigate the consequences of measurement error bias and winner's curse bias in real data applications, we use two separate GWAS data sets on the same trait as the exposure and outcome. The true causal effect is 1 and, therefore, such analyses provide opportunities to validate our proposed method. We conduct two same trait analyses: body mass index (BMI)–BMI analysis and high-density lipoprotein (HDL)–HDL analysis. In the BMI–BMI analysis, we use two BMI GWAS results from the UK Biobank (sample size 461,460, ID: ukb-b-19953; [12]) and from the GIANT consortium (sample size 339,224, ID: ieu-a-2; [24]), denoted by BMI-1 and BMI-2, respectively. In the HDL–HDL analysis, we use

TABLE 3

Same trait analyses results. The true causal effect is $\beta = 1$. $\widehat{\beta}$, SE, 95% CI, # IVs and F stand for the estimated effect size, its standard error, the 95% confidence interval, the number of selected IVs and the F statistic, respectively

	$\widehat{\beta}$	SE	95% CI	# IVs	F	$\widehat{\beta}$	SE	95% CI	# IVs	F
	BMI-1–BMI-2 analysis					HDL-1–HDL-2 analysis				
RIVW	1.030	0.024	0.983–1.077	913	21.9	1.022	0.036	0.952–1.092	766	26.8
dIVW	0.981	0.034	0.915–1.048	1811	3.7	0.912	0.048	0.819–1.006	1787	3.6
IVW	0.785	0.015	0.756–0.814	435	42.5	0.915	0.020	0.876–0.953	321	68.1
	BMI-2–BMI-1 analysis					HDL-2–HDL-1 analysis				
RIVW	1.194	0.060	1.076–1.312	401	18.1	1.039	0.048	0.944–1.134	380	20.9
dIVW	1.402	0.070	1.266–1.539	1810	2.1	0.887	0.050	0.789–0.984	1792	2.1
IVW	0.870	0.026	0.820–0.920	82	42.0	0.922	0.022	0.878–0.967	99	58.1

two HDL GWAS results from the UK Biobank (sample size 403,943, ID: ukb-b-109; [12]) and the Global Lipids Genetics Consortium (sample size 94,595, ID: ebi-a-GCST002223; [37]), denoted by HLD-1 and HDL-2, respectively.

The results are summarized in Table 3. Overall, the RIVW estimator provides accurate estimates of the causal effect $\beta = 1$, confirming its ability to remove both measurement error bias and winner's curse bias in the two-sample MR framework. In addition, RIVW yields smaller standard errors compared with the dIVW estimator, indicating that instrument selection may help improve estimation efficiency. The IVW estimator is biased downwards due to the two sources of biases. While the dIVW estimator (with $\lambda = 0$) generally performs well, we notice that the corresponding F statistic (average of $(\widehat{\gamma}_j/\sigma_{X_j})^2$ across instruments) is quite small, suggesting that many of the included IVs are not relevant for estimating the causal effect. We also observe that reversing the exposure and outcome GWAS leads to quite different numbers of instruments being selected. This is because the UK Biobank study (BMI-1 and HDL-1) has a much larger sample size. As a result, associations calculated from the UK Biobank study tend to have smaller standard errors on average. This should explain why more IVs are selected in the BMI-1–BMI-2 and HDL-1–HDL-2 analyses. Because both BMI-2–BMI-1 and HDL-2–HDL-1 analyses employ fewer IVs and the F statistics also become smaller, our RIVW estimator is more likely to produce noisy estimates in finite samples. Finally, compared with Table S.3 in the Supplementary Material [25], which employs standard pruning (random instrument selection without the use of any GWAS summary statistic), we note that the revised sigma-based pruning procedure leads to significant efficiency improvement.

7.3. *BMI and Covid-19 severity.* Epidemiological studies have reported associations between BMI and COVID-19 illness [36], and several MR analyses have confirmed that higher BMI is a causal risk factor (exposure) for COVID-19 severity [23, 28]. However, the causal effect estimated from these MR analyses may be biased due to measurement error bias and winner's curse bias. While identifying causal risk factors is important for the general public to combat COVID-19, providing accurate causal effect estimates is also important, as it both facilitates effective communication and offers solid evidence for public health policymaking. We apply our proposed method with BMI from the UK Biobank (sample size 461,460, ID: ukb-b-19953) as the exposure and COVID-19 severity from the covid-19hg (v6 leave out UK Biobank samples; [19]) as the outcome data. Detailed information on participating studies, quality control and analyses have been provided on the COVID-19 HGI website (http://www.covid19hg.org/results/). In brief, data from 20,980 hospitalized COVID-19 patients and 1,628,780 population controls were used. The hospitalized COVID-19 cases rep-

TABLE 4

BMI and COVID-19 severity. $\widehat{\beta}$, SE, 95% CI, # IVs and F stand for the estimated effect size (odds ratio), its standard error, the 95% confidence interval, the number of selected IVs and the F statistic, respectively

	$\widehat{\beta}$	SE	95% CI	# IVs	F
RIVW	1.60	0.142	1.32–1.88	1025	21.4
dIVW	1.49	0.163	1.17–1.81	1985	4.0
IVW	1.25	0.069	1.11–1.39	477	42.4

resent patients with (i) laboratory confirmed SARS-CoV-2 infection (RNA and/or serology based) and (ii) hospitalization due to COVID-related symptoms.

Table 4 summarizes our results. First, we confirm that higher BMI is a risk factor for COVID-19 severity, as all three methods yield odds ratios that are statistically significantly different from 1. Similar to the same trait analysis, our RIVW approach provides smaller standard errors compared with the dIVW estimator. Because the proposed RIVW estimator lifts the winner's curse and removes measurement error bias, it produces a larger estimated odds ratio. This larger estimate indicates that BMI is a critical risk factor (maybe even more important than we originally anticipated), and effectively communicating this finding may encourage the general public to take extra precautions and reduce COVID-19 related hospitalization.

8. Closing remarks. We introduced a new two-sample Mendelian randomization with summary data framework that breaks the winner's curse. We further propose the rerandomized inverse variance weighted estimator that eliminates both winner's curse bias and measurement error bias in the popular IVW estimator. To facilitate statistical hypothesis testing, we also provide valid standard error construction. Our RIVW estimator is easy to implement and enjoys rigorous theoretical guarantees. It also remains valid in the presence of balanced horizontal pleiotropy. We plan to further study winner's curse removal in the presence of correlated instruments, and extend our framework to robust MR analyses.

Acknowledgments. The authors' names appear in alphabetical order.

The authors would like to thank the Editor, the Associate Editor and three anonymous reviewers for their comments and suggestions that significantly improved the paper.

The authors also thank Matias D. Cattaneo, Wei Pan, Hongyu Zhao and Meng Zhuo (alphabetically ordered) for their valuable feedback and thoughtful discussions.

Funding. The authors acknowledge financial support from the National Science Foundation (DMS-2015325) and the National Institute of Health (1R03AG070669, R01MH125746). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

SUPPLEMENTARY MATERIAL

Supplement to "Breaking the winner's curse in Mendelian randomization: Rerandomized inverse variance weighted estimator" (DOI: 10.1214/22-AOS2247SUPP; .pdf). The Supplementary Material [25] contains all proofs, additional methodological and theoretical results and numerical evidence.

REFERENCES

[1] ANDREWS, I., STOCK, J. H. and SUN, L. (2019). Weak instruments in instrumental variables regression: Theory and practice. *Ann. Rev. Econ.* **11** 727–753.

[2] ANDREWS, S. J., GOATE, A. and ANSTEY, K. J. (2020). Association between alcohol consumption and Alzheimer's disease: A Mendelian randomization study. *Alzheimer's Dement.* **16** 345–353. https://doi.org/10.1016/j.jalz.2019.09.086

[3] BENJAMIN, D. J., CESARINI, D., CHABRIS, C. F., GLAESER, E. L., LAIBSON, D. I., GUDNASON, V. et al. (2012). The promises and pitfalls of genoeconomics. *Ann. Rev. Econ.* **4** 627–662.

[4] BLACKWELL, D. (1947). Conditional expectation and unbiased sequential estimation. *Ann. Math. Stat.* **18** 105–110. MR0019903 https://doi.org/10.1214/aoms/1177730497

[5] BOWDEN, J., DEL GRECO, F. M., MINELLI, C., DAVEY SMITH, G., SHEEHAN, N. and THOMPSON, J. (2017). Assessing the suitability of summary data for two-sample Mendelian randomization analyses using MR-egger regression: The role of the I2 statistic. *Int. J. Epidemiol.* **45** 1961–1974.

[6] BOWDEN, J., DEL GRECO, F. M., MINELLI, C., ZHAO, Q., LAWLOR, D. A., SHEEHAN, N. A. et al. (2019). Improving the accuracy of two-sample summary-data Mendelian randomization: Moving beyond the NOME assumption. *Int. J. Epidemiol.* **48** 728–742.

[7] BOWDEN, J. and DUDBRIDGE, F. (2009). Unbiased estimation of odds ratios: Combining genomewide association scans with replication studies. *Genet. Epidemiol.* **33** 406–418. https://doi.org/10.1002/gepi.20394

[8] BURGESS, S., BUTTERWORTH, A. and THOMPSON, S. G. (2013). Mendelian randomization analysis with multiple genetic variants using summarized data. *Genet. Epidemiol.* **37** 658–665. https://doi.org/10.1002/gepi.21758

[9] BURGESS, S., DAVEY SMITH, G., DAVIES, N. M., DUDBRIDGE, F., GILL, D., GLYMOUR, M. M. et al. (2020). Guidelines for performing Mendelian randomization investigations. *Wellcome Open Res.* **4**.

[10] BURGESS, S., SCOTT, R. A., TIMPSON, N. J., SMITH, G. D. and THOMPSON, S. G. (2015). Using published data in Mendelian randomization: A blueprint for efficient identification of causal risk factors. *Eur. J. Epidemiol.* **30** 543–552.

[11] BURGESS, S. and THOMPSON, S. G. (2017). Interpreting findings from Mendelian randomization using the MR-egger method. *Eur. J. Epidemiol.* **32** 377–389. https://doi.org/10.1007/s10654-017-0255-x

[12] BYCROFT, C., FREEMAN, C., PETKOVA, D., BAND, G., ELLIOTT, L. T., SHARP, K. et al. (2018). The UK biobank resource with deep phenotyping and genomic data. *Nature* **562** 203–209.

[13] CHAO, J. C. and SWANSON, N. R. (2005). Consistent estimation with a large number of weak instruments. *Econometrica* **73** 1673–1692. MR2156676 https://doi.org/10.1111/j.1468-0262.2005.00632.x

[14] CRUDU, F., MELLACE, G. and SÁNDOR, Z. (2021). Inference in instrumental variable models with heteroskedasticity and many instruments. *Econometric Theory* **37** 281–310. MR4246218 https://doi.org/10.1017/S026646662000016X

[15] GKATZIONIS, A. and BURGESS, S. (2019). Contextualizing selection bias in Mendelian randomization: How bad is it likely to be? *Int. J. Epidemiol.* **48** 691–701. https://doi.org/10.1093/ije/dyy202

[16] HANSEN, C. and KOZBUR, D. (2014). Instrumental variables estimation with many weak instruments using regularized JIVE. *J. Econometrics* **182** 290–308. MR3233473 https://doi.org/10.1016/j.jeconom.2014.04.022

[17] HEMANI, G., BOWDEN, J. and SMITH, G. D. (2018). Evaluating the potential role of pleiotropy in Mendelian randomization studies. *Hum. Mol. Genet.* **27** R195–R208. https://doi.org/10.1093/hmg/ddy163

[18] HEMANI, G., ZHENG, J., ELSWORTH, B., WADE, K. H., HABERLAND, V., BAIRD, D. et al. (2018). The MR-base platform supports systematic causal inference across the human phenome. *eLife*.

[19] INITIATIVE, C. H. G. (2021). Mapping the human genetic architecture of Covid-19 by worldwide meta-analysis. *Nature* **600** 472–277.

[20] LARSSON, S. C., TRAYLOR, M., MALIK, R., DICHGANS, M., BURGESS, S., MARKUS, H. S. and COSTREAM CONSORTIUM, ON BEHALF OF THE INTERNATIONAL GENOMICS OF ALZHEIMER'S PROJECT (2017). Modifiable pathways in Alzheimer's disease: Mendelian randomisation analysis. *BMJ* **359** j5375. https://doi.org/10.1136/bmj.j5375

[21] LAWLOR, D. A. (2016). Commentary: Two-sample Mendelian randomization: Opportunities and challenges. *Int. J. Epidemiol.* **45** 908–915.

[22] LAWLOR, D. A., HARBORD, R. M., STERNE, J. A. C., TIMPSON, N. and SMITH, G. D. (2008). Mendelian randomization: Using genes as instruments for making causal inferences in epidemiology. *Stat. Med.* **27** 1133–1163. MR2420151 https://doi.org/10.1002/sim.3034

[23] Leong, A., Cole, J. B., Brenner, L. N., Meigs, J. B., Florez, J. C. and Mercader, J. M. (2021). Cardiometabolic risk factors for Covid-19 susceptibility and severity: A Mendelian randomization analysis. *PLoS Med*. **18** e1003553. https://doi.org/10.1371/journal.pmed.1003553

[24] Locke, A. E., Kahali, B., Berndt, S. I., Justice, A. E., Pers, T. H., Day, F. R. et al. (2015). Genetic studies of body mass index yield new insights for obesity biology. *Nature* **518** 197–206.

[25] Ma, X., Wang, J. and Wu, C. (2023). Supplement to "Breaking the Winner's Curse in Mendelian Randomization: Rerandomized Inverse Variance Weighted Estimator." https://doi.org/10.1214/22-AOS2247SUPP

[26] Mikusheva, A. and Sun, L. (2022). Inference with many weak instruments. *Rev. Econ. Stud*. **89** 2663–2686. MR4478280 https://doi.org/10.1093/restud/rdab097

[27] Pierce, B. L. and Burgess, S. (2013). Efficient design for Mendelian randomization studies: Subsample and 2-sample instrumental variable estimators. *Amer. J. Epidemiol*. **178** 1177–1184. https://doi.org/10.1093/aje/kwt084

[28] Ponsford, M. J., Gkatzionis, A., Walker, V. M., Grant, A. J., Wootton, R. E., Moore, L. S. P. et al. (2020). Cardiometabolic traits, sepsis, and severe Covid-19: A Mendelian randomization investigation. *Circulation* **142** 1791–1793.

[29] Purcell, S., Neale, B., Todd-Brown, K., Thomas, L., Ferreira, M. A., Bender, D. et al. (2007). PLINK: A tool set for whole-genome association and population-based linkage analyses. *Am. J. Hum. Genet*. **81** 559–575.

[30] Qi, G. and Chatterjee, N. (2019). Mendelian randomization analysis using mixture models for robust and efficient estimation of causal effects. *Nat. Commun*. **10** 1–10.

[31] Robertson, D. S., Prevost, A. T. and Bowden, J. (2016). Accounting for selection and correlation in the analysis of two-stage genome-wide association studies. *Biostatistics* **17** 634–649. MR3604270 https://doi.org/10.1093/biostatistics/kxw012

[32] Rosenthal, R. (1979). The file drawer problem and tolerance for null results. *Psychol. Bull*. **86** 638–641.

[33] Sadreev, I. I., Elsworth, B. L., Mitchell, R. E., Paternoster, L., Sanderson, E. and Davies, N. M. (2021). Navigating sample overlap, winner's curse and weak instrument bias in Mendelian randomization studies using the UK biobank. *medRxiv*.

[34] Siedlinski, M., Jozefczuk, E., Xu, X., Teumer, A., Evangelou, E. and Schnabel, R. B. (2020). White blood cells and blood pressure: A Mendelian randomization study. *Circulation* **141** 1307–1317.

[35] Smith, G. D. and Ebrahim, S. (2004). Mendelian randomization: Prospects, potentials, and limitations. *Int. J. Epidemiol*. **33** 30–42. https://doi.org/10.1093/ije/dyh132

[36] Soeroto, A. Y., Soetedjo, N. N., Purwiga, A., Santoso, P., Kulsum, I. D., Suryadinata, H. et al. (2020). Effect of increased BMI and obesity on the outcome of Covid-19 adult patients: A systematic review and meta-analysis. *Diabetes Metab. Syndr. Clin. Res. Rev.* **14** 1897–1904.

[37] Willer, C. J., Schmidt, E. M., Sengupta, S., Peloso, G. M., Gustafsson, S. et al. (2013). Discovery and refinement of loci associated with lipid levels. *Nat. Genet*. **45** 1274–1283.

[38] Wu, C., Wu, L., Wang, J., Lin, L., Li, Y., Lu, Q. and Deng, H.-W. (2021). Systematic identification of risk factors and drug repurposing options for Alzheimer's disease. *Alzheimer's Dement*. **7** e12148.

[39] Xue, H., Shen, X. and Pan, W. (2021). Constrained maximum likelihood-based Mendelian randomization robust to both correlated and uncorrelated pleiotropic effects. *Am. J. Hum. Genet*. **108** 1251–1269.

[40] Yang, J., Zeng, J., Goddard, M. E., Wray, N. R. and Visscher, P. M. (2017). Concepts, estimation and interpretation of SNP-based heritability. *Nat. Genet*. **49** 1304–1310.

[41] Ye, T., Shao, J. and Kang, H. (2021). Debiased inverse-variance weighted estimator in two-sample summary-data Mendelian randomization. *Ann. Statist*. **49** 2079–2100. MR4319242 https://doi.org/10.1214/20-aos2027

[42] Zeng, J., De Vlaming, R., Wu, Y., Robinson, M. R., Lloyd-Jones, L. R., Yengo, L., Yap, C. X. et al. (2018). Signatures of negative selection in the genetic architecture of human complex traits. *Nat. Genet*. **50** 746–753.

[43] Zhao, Q., Chen, Y., Wang, J. and Small, D. S. (2019). Powerful three-sample genome-wide design and robust statistical inference in summary-data Mendelian randomization. *Int. J. Epidemiol*. **48** 1478–1492.

[44] Zhao, Q., Wang, J., Hemani, G., Bowden, J. and Small, D. S. (2020). Statistical inference in two-sample summary-data Mendelian randomization using robust adjusted profile score. *Ann. Statist*. **48** 1742–1769. MR4124342 https://doi.org/10.1214/19-AOS1866

[45] Zhu, Y. (2018). Sparse linear models and l_1-regularized 2SLS with high-dimensional endogenous regressors and instruments. *J. Econometrics* **202** 196–213. MR3778835 https://doi.org/10.1016/j.jeconom.2017.10.002