Retrospective Association Analysis of Longitudinal Binary Traits Identifies Important Loci and Pathways in Cocaine Use

Weimiao Wu,¹ Zhong Wang,² Ke Xu,^{3,4} Xinyu Zhang,^{3,4} Amei Amei,⁵ Joel Gelernter,^{3,4} Hongyu
Zhao,¹ Amy C. Justice,^{4,6} and Zuoheng Wang^{1,*}

¹Department of Biostatistics, Yale School of Public Health, New Haven, CT

²Baker Institute for Animal Health, Cornell University, Ithaca, NY

³Department of Psychiatry, Yale School of Medicine, New Haven, CT

⁴VA Connecticut Healthcare System, West Haven, CT

⁵Department of Mathematical Sciences, University of Nevada, Las Vegas, NV

⁶Department of Internal Medicine, Yale School of Medicine, New Haven, CT

**email:* zuoheng.wang@yale.edu

SUMMARY. Longitudinal phenotypes have been increasingly available in genome-wide association studies (GWAS) and electronic health record-based studies for identification of genetic variants that influence complex traits over time. For longitudinal binary data, there remain significant challenges in gene mapping, including misspecification of the model for the phenotype distribution due to ascertainment. Here, we propose L-BRAT, a retrospective, generalized estimating equations-based method for genetic association analysis of longitudinal binary outcomes. We also develop RGMMAT, a retrospective, generalized linear mixed model-based association test. Both tests are retrospective score approaches in which genotypes are treated as random conditional on phenotype and covariates. They allow both static and time-varying covariates to be included in the analysis. Through simulations, we illustrated that retrospective association tests are robust to ascertainment and other types of phenotype model misspecification, and gain power over previous association methods. We applied L-BRAT and RGMMAT to a genome-wide association analysis of repeated measures of cocaine use in a longitudinal cohort. Pathway analysis implicated association with opioid signaling and axonal guidance signaling pathways. Lastly, we replicated important pathways in an independent cocaine dependence case-control GWAS. Our results illustrate that L-BRAT is able to detect important loci and pathways in a genome scan and to provide insights into genetic architecture of cocaine use.

KEY WORDS: Ascertainment; Genome-wide association studies; Model misspecification; Robustness; Score test.

1. Introduction

Genome-wide association studies (GWAS) have successfully discovered many disease susceptibility loci and provided insights into the genetic architecture of numerous human complex diseases and traits. In some genetic epidemiological studies, longitudinally collected phenotype data are available. This is the case for many electronic health record (EHR)-based studies. As many of these studies continue to follow enrolled subjects (e.g. the UK Biobank (UKB) and the Million Veteran Program (MVP)), longitudinal phenotypes will be increasingly available with the passage of time, providing new data resources that require appropriate analytical tools for optimal analysis. Standard association tests that consider one time point or collapse repeated measurements into a single value such as an average do not capture the trajectory of phenotypic traits over time and may result in a loss of statistical power to detect genetic associations. In addition, the effects of time-varying covariates cannot be easily incorporated in such analyses. Recently, methodological developments for GWAS have proliferated to make full use of the available longitudinal data. For population cohorts, methods that account for dependence among observations from an individual include mixed effects models (Furlotte et al. 2012; Sikorska et al. 2013), generalized estimating equations (GEE) (Sitlani et al. 2015), growth mixture models (Das et al. 2011; Londono et al. 2013), and empirical Bayes models (Meirelles et al. 2013). Most of these approaches are prospective analyses and have been successfully applied to quantitative phenotypes.

As many diseases are rare, efficient designs, such as the case-control design, are commonly applied in epidemiological studies to recruit study subjects. Despite the enhanced efficiency in the study sample, non-random ascertainment can be a major source of model misspecification that may lead to inflated type I error and/or power loss in association analysis. The linear mixed model and the logistic mixed model do not perform well when the case-control ratio is

unbalanced in large-scale genetic association studies (Zhou et al. 2018). Prospective analysis in which a population-based model is used ignores ascertainment bias and can result in compromised statistical inference. Furthermore, in the ascertained sample, the prospective approach conditional on the genotype and covariates may lose information when the joint distribution of the genotype and covariates carries additional information on whether the phenotype is associated with the genotype (Jiang et al. 2015). In this regard, several retrospective association methods have been proposed for analyzing ascertained population-based case-control studies (Hayeck et al. 2015; Jiang et al. 2016), family-based studies of continuous traits (Jakobsdottir and McPeck 2013), family-based case-control studies (Zhong et al. 2016; Hayeck et al. 2017), and family-based longitudinal quantitative traits (Wu and McPeck 2018). Compared to prospective tests, retrospective tests conditional on the phenotype and covariates are more robust to misspecification of the trait model (Jiang et al. 2015).

To generalize case-control sampling, outcome-dependent sampling designs have become popular for binary data in longitudinal cohort studies (Schildcrout and Heagerty 2008; Schildcrout et al. 2018a,b). However, association tests for longitudinally measured binary data are less well developed in GWAS. Here, we propose L-BRAT, a retrospective, GEE-based method for genetic association analysis of longitudinal binary outcomes. It requires specification of the mean of the outcome distribution and a working correlation matrix for repeated measurements. L-BRAT is a retrospective score approach in which genotypes are treated as random conditional on the phenotype and covariates. Thus, it is robust to ascertainment and trait model misspecification. It allows both static and time-varying covariates to be included in the analysis. We note that GMMAT, a recently proposed prospective test using the logistic mixed model to control for population structure and cryptic relatedness in case-control studies (Chen et al. 2016), can be adapted for repeated binary data. For comparison, we also develop RGMMAT,

a retrospective, generalized linear mixed model (GLMM)-based association test for longitudinal binary traits.

We performed simulation studies to evaluate the type I error and power of L-BRAT and RGMMAT, and compared them to the existing prospective methods. The results demonstrate that the retrospective association tests have better control of type I error when the phenotype model is misspecified, and are robust to various ascertainment schemes. Moreover, they are more powerful than the prospective tests. Finally, we applied L-BRAT and RGMMAT to a genome-wide association analysis of repeated measurements of cocaine use in a longitudinal cohort, the Veterans Aging Cohort Study (VACS), and replicated the results using data from an independent cocaine dependence case-control GWAS.

2. Methods

Suppose a binary trait is measured over time on a study population of n individuals. We have their genome-wide measures of genetic variation. A set of covariates, static or dynamic, are also available. Let n_i be the number of repeated measures on individual i and $N = \sum_{i=1}^n n_i$ be the total number of observations. For individual i , let $\mathbf{X}_{ij}$ and Y_{ij} be the p -dimensional covariate vector, assumed to include an intercept, and the binary response at time t_{ij} , respectively. In this setting, individuals are permitted to have measurements at different time points and different number of observations. We let $\mathbf{Y}$ denote the outcome vector of length N , and let $\mathbf{X}$ denote the $N \times p$ covariate matrix. Here, we focus on the problem of testing for association between a genetic variant and the longitudinal binary outcomes. Let $\mathbf{G}$ denote the vector of genotypes for the n individuals at the variant to be tested, where $G_i = 0, 1$, or 2 is the number of minor alleles of individual i at the variant.

2.1. GEE-based Model

We consider a GEE approach in which the mean of the outcome distribution, given the genotype and covariates, is specified as

$$E(Y_{ij} | \mathbf{G}, \mathbf{X}) = \mu_{ij}, \quad \text{logit}(\mu_{ij}) = \mathbf{X}_{ij}^T \boldsymbol{\beta} + G_i \gamma, \quad i = 1, \dots, n; j = 1, \dots, n_i, \quad (1)$$

where $\boldsymbol{\beta}$ is a p -dimensional vector of covariate effects and γ is a scalar parameter of interest representing the effect of the tested variant. Writing in a matrix form, we have the mean model

$$E(\mathbf{Y} | \mathbf{G}, \mathbf{X}) = \boldsymbol{\mu}, \quad \text{logit}(\boldsymbol{\mu}) = \mathbf{X}\boldsymbol{\beta} + \mathbf{B}\mathbf{G}\gamma, \quad (2)$$

where $\mathbf{B}$ is an $N \times n$ matrix representing the measurement clustering structure, and its (l, i) th entry B_{li} is an indicator of the l th entry of $\mathbf{Y}$ being a measurement on individual i . Here, the vector $\mathbf{B}\mathbf{G}$ is the vertically expanded genotype vector that maps the genotype data $\mathbf{G}$ from the individual level to the measurement level. The covariance structure of $\mathbf{Y}$ is given by

$$\text{Var}(\mathbf{Y} | \mathbf{G}, \mathbf{X}) = \boldsymbol{\Gamma}^{1/2} \boldsymbol{\Sigma} \boldsymbol{\Gamma}^{1/2}, \quad (3)$$

where $\boldsymbol{\Gamma} = \text{diag}\{\mu_{1,1}(1 - \mu_{1,1}), \dots, \mu_{1,n_1}(1 - \mu_{1,n_1}), \dots, \mu_{n,1}(1 - \mu_{n,1}), \dots, \mu_{n,n_n}(1 - \mu_{n,n_n})\}$ is an N -dimensional diagonal matrix and $\boldsymbol{\Sigma}$ is an $N \times N$ correlation matrix. The covariance specification in Eq. (3) ensures that the variance of the dichotomous response Y_{ij} depends on its mean in a way that is consistent with the Bernoulli distribution. To apply the GEE method, a working correlation structure such as independent, exchangeable, and first-order autoregressive (AR(1)) must be specified. For a given within-cluster correlation matrix $\boldsymbol{\Sigma}(\tau)$, which may depend on some parameter τ , the estimating equations for the unknown parameters $(\boldsymbol{\beta}, \gamma)$ are written as

$$\mathbf{U} = \begin{bmatrix} \mathbf{U}(\boldsymbol{\beta}) \\ \mathbf{U}(\gamma) \end{bmatrix} = \begin{bmatrix} \mathbf{X}^T \boldsymbol{\Gamma}^{1/2} \boldsymbol{\Sigma}^{-1} \boldsymbol{\Gamma}^{-1/2} (\mathbf{Y} - \boldsymbol{\mu}) \\ (\mathbf{B}\mathbf{G})^T \boldsymbol{\Gamma}^{1/2} \boldsymbol{\Sigma}^{-1} \boldsymbol{\Gamma}^{-1/2} (\mathbf{Y} - \boldsymbol{\mu}) \end{bmatrix}.$$

To detect association between the genetic variant and the phenotype, we consider a score approach to test $H_0: \gamma = 0$ against $H_1: \gamma \neq 0$. The null estimate of $\boldsymbol{\beta}$, denoted by $\hat{\boldsymbol{\beta}}_0$, is the solution to a system of estimating equations $\mathbf{U}(\boldsymbol{\beta}) = 0$ under the constraint $\gamma = 0$, which can be

computed iteratively between a Fisher scoring algorithm for β and the method of moments for τ until convergence. Then, the score function for γ is

$$U_0 = U(\gamma)_{|\hat{\beta}_0, 0, \hat{\tau}_0} = (\mathbf{BG})^T \hat{\Gamma}_0^{1/2} \hat{\Sigma}_0^{-1} \hat{\Gamma}_0^{-1/2} (\mathbf{Y} - \hat{\mu}_0), \quad (4)$$

where $\hat{\mu}_0$, $\hat{\Gamma}_0$ and $\hat{\Sigma}_0$ are μ , Γ and Σ evaluated at $(\beta, \gamma, \tau) = (\hat{\beta}_0, 0, \hat{\tau}_0)$.

In the GEE approach, the prospective score statistic for testing $H_0: \gamma = 0$ takes the form

$$T_{GEE} = \frac{U_0^2}{\text{Var}_0(U_0 | \mathbf{G}, \mathbf{X})} = \frac{[(\mathbf{BG})^T \hat{\Gamma}_0^{1/2} \hat{\Sigma}_0^{-1} \hat{\Gamma}_0^{-1/2} (\mathbf{Y} - \hat{\mu}_0)]^2}{(\mathbf{BG})^T \mathbf{Q} \mathbf{BG}}, \quad (5)$$

where the null variance of U_0 is evaluated using a robust sandwich variance estimator, conditional on the genotype and covariates. Here $\mathbf{Q} = \mathbf{V} - \mathbf{VX}(X^T \mathbf{VX})^{-1} X^T \mathbf{V}$, where $\mathbf{V} = \hat{\Gamma}_0^{1/2} \hat{\Sigma}_0^{-1} \hat{\Gamma}_0^{-1/2} \widehat{\text{Cov}}(\mathbf{Y}) \hat{\Gamma}_0^{-1/2} \hat{\Sigma}_0^{-1} \hat{\Gamma}_0^{1/2}$ and the sample covariance of $\mathbf{Y}$, $\widehat{\text{Cov}}(\mathbf{Y})$, is estimated by $(\mathbf{Y} - \hat{\mu}_0)(\mathbf{Y} - \hat{\mu}_0)^T$. Under the null hypothesis, the T_{GEE} test statistic has an asymptotic χ_1^2 distribution.

In what follows, we introduce a new GEE-based association testing method, L-BRAT (Longitudinal Binary-trait Retrospective Association Test). The L-BRAT test statistic is also based on the score function U_0 in Eq. (4). In contrast to the prospective GEE score test, L-BRAT takes a retrospective approach in which the variance of U_0 is assessed using a retrospective model of the genotype given the phenotype and covariates. An advantage of the retrospective approach is that the analysis is less dependent on the correct specification of the phenotype model. We assume that under the null hypothesis of no association between the genetic variant and the phenotype, the quasi-likelihood model of $\mathbf{G}$ conditional on $\mathbf{Y}$ and $\mathbf{X}$ is

$$E_0(\mathbf{G} | \mathbf{Y}, \mathbf{X}) = 2p\mathbf{1}_n, \quad \text{Var}_0(\mathbf{G} | \mathbf{Y}, \mathbf{X}) = \sigma_g^2 \mathbf{\Phi}, \quad (6)$$

where p is the minor allele frequency (MAF) of the tested variant, $\mathbf{1}_n$ is a vector of length n with every element equal to 1, σ_g^2 is an unknown variance parameter, and $\mathbf{\Phi}$ is an $n \times n$ genetic

relationship matrix (GRM) representing the overall genetic similarity between individuals due to population structure. Because $\mathbf{B}\mathbf{1}_n = \mathbf{1}_N$, which is the first column of $\mathbf{X}$ that encodes an intercept, and $\hat{\mathbf{\Gamma}}_0^{1/2}\hat{\mathbf{\Sigma}}_0^{-1}\hat{\mathbf{\Gamma}}_0^{-1/2}(\mathbf{Y} - \hat{\boldsymbol{\mu}}_0)$, the N -dimensional vector of transformed null phenotypic residuals, is orthogonal to the column space of $\mathbf{X}$, then the null mean model of $\mathbf{G}$ in Eq. (6) ensures that

$$E_0(U_0 | \mathbf{Y}, \mathbf{X}) = E_0(\mathbf{A}^T \mathbf{G} | \mathbf{Y}, \mathbf{X}) = 2p\mathbf{A}^T \mathbf{1}_n = 0,$$

where $\mathbf{A} = \mathbf{B}^T \hat{\mathbf{\Gamma}}_0^{1/2} \hat{\mathbf{\Sigma}}_0^{-1} \hat{\mathbf{\Gamma}}_0^{-1/2} (\mathbf{Y} - \hat{\boldsymbol{\mu}}_0)$ is the individual-level transformed phenotypic residual vector of length n .

In model (6), the GRM Φ can be obtained using genome-wide data, given by

$$\Phi = \frac{1}{K} \sum_{k=1}^K \frac{(\mathbf{G}_{(k)} - 2\hat{p}_k)(\mathbf{G}_{(k)} - 2\hat{p}_k)^T}{2\hat{p}_k(1 - \hat{p}_k)},$$

where K is the total number of genotyped variants, $\mathbf{G}_{(k)}$ is the genotype vector at the k th variant, and $\hat{p}_k$ is the estimated MAF, for example, $\hat{p}_k = \bar{G}_k/2$, the sample MAF at the k th variant. For the variant of interest, let $\hat{p} = \bar{G}/2$ be its sample MAF. Under Hardy-Weinberg equilibrium, the variance of the genotype is estimated by $\hat{\sigma}_g^2 = 2\hat{p}(1 - \hat{p})$. Or we can use a more robust variance estimator (Jakobsdottir and McPeck 2013) given by

$$\tilde{\sigma}_g^2 = (n - 1)^{-1} \mathbf{G}^T \mathbf{W} \mathbf{G},$$

where $\mathbf{W} = \Phi^{-1} - \Phi^{-1} \mathbf{1}_n (\mathbf{1}_n^T \Phi^{-1} \mathbf{1}_n)^{-1} \mathbf{1}_n^T \Phi^{-1}$. Finally, the L-BRAT test statistic can be defined as

$$\text{L-BRAT} = \frac{U_0^2}{\text{var}_0(U_0 | \mathbf{Y}, \mathbf{X})} = \frac{(\mathbf{A}^T \mathbf{G})^2}{\text{var}_0(\mathbf{A}^T \mathbf{G} | \mathbf{Y}, \mathbf{X})} = \frac{(\mathbf{A}^T \mathbf{G})^2}{\hat{\sigma}_g^2 \mathbf{A}^T \Phi \mathbf{A}}. \quad (7)$$

Under regularity conditions, L-BRAT asymptotically follows a χ_1^2 distribution under the null hypothesis.

2.2. GLMM-based Model

The Generalized linear Mixed Model Association Test (GMMAT) was originally designed to use random effects in logistic mixed models to account for population structure and cryptic relatedness in case-control studies (Chen et al. 2016). To extend the GMMAT method for case-control analysis to repeated binary data, we consider the following logistic mixed model:

$$\text{logit}(\mu_{ij}) = \mathbf{X}_{ij}^T \boldsymbol{\beta} + G_i \gamma + a_i + r_{ij}, \quad i = 1, \dots, n; j = 1, \dots, n_i, \quad (8)$$

where $\mu_{ij} = P(Y_{ij} = 1 \mid G_i, \mathbf{X}_{ij}, a_i, r_{ij})$ is the probability of a binary response at time t_{ij} for individual i , conditional on his/her genotype, covariates, and random effects a_i and r_{ij} , $\boldsymbol{\beta}$ and γ are the same as defined in model (1), a_i is the individual random effect, and r_{ij} is the individual-specific time-dependent random effect. The two random effects were used to capture the correlation among repeated measures in gene-based association test for longitudinal traits (Wang et al. 2017). Here, a_i 's are assumed to be independent and $a_i \sim N(0, \sigma_a^2)$. The vector of time-dependent random effects $\mathbf{r}_i = (r_{i1}, \dots, r_{i,n_i})$ has a multivariate normal distribution, $\mathbf{r}_i \sim MVN(\mathbf{0}, \sigma_r^2 \mathbf{R}_i)$, where an AR(1) structure is assumed for the correlation matrix $\mathbf{R}_i$, in which τ is the unknown parameter. The binary responses Y_{ij} are assumed to be independent given the random effects a_i and r_{ij} .

To construct a score test for the null hypothesis $H_0: \gamma = 0$ vs. the alternative $H_1: \gamma \neq 0$, we use the penalized quasi-likelihood method (Breslow and Clayton 1993) to fit the null logistic mixed model and obtain the null estimates of $\boldsymbol{\beta}$, σ_a^2 , σ_r^2 and τ , denoted by $\hat{\boldsymbol{\beta}}_0$, $\hat{\sigma}_a^2$, $\hat{\sigma}_r^2$ and $\hat{\tau}_0$ (Chen et al. 2016). Similarly, the best linear unbiased predictor (BLUP) of random effects, $\hat{\mathbf{a}}$ and $\hat{\mathbf{r}}$, can be obtained. Then, the resulting score function for γ is

$$S_0 = S(\gamma)_{|\hat{\boldsymbol{\beta}}_0, 0, \hat{\sigma}_a^2, \hat{\sigma}_r^2, \hat{\tau}_0, \hat{\mathbf{a}}, \hat{\mathbf{r}}} = (\mathbf{B}\mathbf{G})^T (\mathbf{Y} - \hat{\boldsymbol{\mu}}_0), \quad (9)$$

where $\hat{\boldsymbol{\mu}}_0 = \text{logit}^{-1}(\mathbf{X}\hat{\boldsymbol{\beta}}_0 + \mathbf{B}\hat{\mathbf{a}} + \hat{\mathbf{r}})$ is a vector of fitted values under H_0 .

In GMMAT, the null variance of the score S_0 is evaluated prospectively (Chen et al. 2016), i.e., $\text{Var}_0(S_0 | \mathbf{G}, \mathbf{X}) = (\mathbf{BG})^T \mathbf{PBG}$, where $\mathbf{P} = \mathbf{\Psi}^{-1} - \mathbf{\Psi}^{-1} \mathbf{X}(\mathbf{X}^T \mathbf{\Psi}^{-1} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{\Psi}^{-1}$, and $\mathbf{\Psi} = \hat{\mathbf{\Gamma}}_0^{-1} + \hat{\sigma}_a^2 \mathbf{BB}^T + \hat{\sigma}_r^2 \hat{\mathbf{R}}$. Here $\hat{\mathbf{\Gamma}}_0$ and $\hat{\mathbf{R}}$ are $\mathbf{\Gamma}$ and $\mathbf{R}$ evaluated at $(\boldsymbol{\beta}, \gamma, \sigma_a^2, \sigma_r^2, \tau) = (\hat{\boldsymbol{\beta}}_0, 0, \hat{\sigma}_a^2, \hat{\sigma}_r^2, \hat{\tau}_0)$, where $\mathbf{\Gamma}$ is the same as defined in Eq. (3) and $\mathbf{R} = \text{diag}\{\mathbf{R}_1, \dots, \mathbf{R}_n\}$ is a block diagonal matrix. The GMMAT test statistic can be written as

$$T_{GMMAT} = \frac{S_0^2}{\text{Var}_0(S_0 | \mathbf{G}, \mathbf{X})} = \frac{[(\mathbf{BG})^T (\mathbf{Y} - \hat{\boldsymbol{\mu}}_0)]^2}{(\mathbf{BG})^T \mathbf{PBG}}. \quad (10)$$

Like L-BRAT, we can construct a retrospective test to assess the significance of the GLMM score function of Eq. (9), which we call RGMMAT, based on the quasi-likelihood model of $\mathbf{G}$ in Eq. (6). Thus, we define the RGMMAT statistic by

$$\text{RGMMAT} = \frac{S_0^2}{\text{Var}_0(S_0 | \mathbf{Y}, \mathbf{X})} = \frac{(\mathbf{C}^T \mathbf{G})^2}{\text{Var}_0(\mathbf{C}^T \mathbf{G} | \mathbf{Y}, \mathbf{X})} = \frac{(\mathbf{C}^T \mathbf{G})^2}{\hat{\sigma}_g^2 \mathbf{C}^T \boldsymbol{\Phi} \mathbf{C}}, \quad (11)$$

where $\mathbf{C} = \mathbf{B}^T (\mathbf{Y} - \hat{\boldsymbol{\mu}}_0)$ is the n -dimensional vector of phenotypic residuals at the individual level by summing over all time points for an individual, and the phenotypic residuals are obtained by fitting the null logistic mixed model. Both the GMMAT and RGMMAT test statistics are assumed to have χ_1^2 asymptotic null distributions.

3. Simulation Studies

We performed simulation studies to evaluate the type I error and power of the two retrospective tests we propose, and compared them to the prospective GEE and GMMAT methods. We also assessed sensitivity of L-BRAT and RGMMAT in the presence of model misspecification and ascertainment. In the simulations, we considered two different trait models and three different ascertainment schemes. Because both the L-BRAT and GEE methods require specification of a working correlation matrix, we implemented three working correlation structures: (1) independent, (2) AR(1), and (3) a mixture of exchangeable and AR(1).

To generate genotypes, we first simulated 10,000 chromosomes over a 1 Mb region using a coalescent model that mimics the linkage disequilibrium (LD) and recombination rates of the European population (Schaffner et al. 2005). We then randomly selected 1,000 non-causal single nucleotide polymorphisms (SNPs) with $MAF > 0.05$. In addition, we generated two causal SNPs that were assumed to influence the trait value with epistasis. In the type I error simulations, we tested association at the 1,000 non-causal SNPs. In each simulation setting, we generated 1,000 sets of phenotypes at five time points. Putting together, 10^6 replicates were used for the type I error evaluation. In the power simulations, we tested the first of the two causal SNPs and empirical power was assessed using 1,000 simulation replicates. In all tests considered, the genotypes at the untested causal SNP(s) were assumed to be unobserved.

3.1. Trait Models

We simulated two types of binary trait models at five time points, in which the two unlinked causal SNPs were assumed to act on the phenotype epistatically. The first type is a logistic mixed model, given by

$$Y_{ij} | \mathbf{X}_{ij}, G_{i(1)}, G_{i(2)}, a_i, r_{ij} \sim \text{Bernoulli}(\mu_{ij}),$$

$$\text{logit}(\mu_{ij}) = -2.5 + 0.2(j - 1) + 0.5X_{ij(1)} + 0.5X_{i(2)} + \theta 1_{\{G_{i(1)} > 0, G_{i(2)} > 0\}} + a_i + r_{ij},$$

where $X_{ij(1)}$ is a continuous, time-varying covariate generated independently from a standard normal distribution, $X_{i(2)}$ is a binary, time-invariant covariate taking values 0 or 1 with a probability of 0.5, $G_{i(1)}$ and $G_{i(2)}$ are the genotypes of individual i at the two causal SNPs, θ is a scalar parameter encoding the effect of the causal SNPs, $1_{\{G_{i(1)} > 0, G_{i(2)} > 0\}}$ is an indicator function that takes value 1 when individual i has at least one copy of the minor allele at both the causal SNPs, a_i and r_{ij} are the individual-level time-independent and time-dependent random effects, respectively. Here we assume $a_i \sim N(0, \sigma_a^2)$ and $\mathbf{r}_i = (r_{i1}, \dots, r_{i5}) \sim MVN(\mathbf{0}, \sigma_r^2 \mathbf{R})$, where $\mathbf{R}$ is

a 5×5 correlation matrix specified by the AR(1) structure with a correlation coefficient τ . The two causal SNPs are assumed to be unlinked with MAFs 0.1 and 0.5, respectively. The variance components are set to $\sigma_a^2 = \sigma_r^2 = 0.64$ and $\tau = 0.7$.

The second type of trait model we considered is a liability threshold model in which an underlying continuous liability determines the outcome value based on a threshold. Specifically, the phenotype Y_{ij} is given by

$$Y_{ij} = 1 \text{ if } L_{ij} > 0,$$

$$\text{with } L_{ij} = -2.0 + 0.2(j - 1) + 0.5X_{ij(1)} + 0.5X_{ij(2)} + \theta 1_{\{G_{i(1)} > 0, G_{i(2)} > 0\}} + a_i + r_{ij} + e_{ij},$$

where L_{ij} is the underlying liability for individual i at time t_{ij} , and $e_{ij} \sim N(0, \sigma_e^2)$ represents independent noise, with $\sigma_e^2 = 0.64$. All other parameters are the same as those in the logistic mixed model.

In both models, we included a time effect and assumed that the mean of the outcome increases with time. The effect of the causal SNPs was set to $\theta = 0.34$ in the type I error simulations. For the power evaluation, we considered a range of values for θ , where we set $\theta = 0.3, 0.32, 0.34, 0.36$, and 0.38 . At the given parameter values, the prevalence of the event of interest ranges from 12.8% to 27.7% over time. The proportion of the phenotypic variance explained by the two causal SNPs ranges from 0.69% to 1.10% in the logistic mixed model, and from 0.49% to 0.78% in the liability threshold model.

3.2. Sampling Designs

We considered three different sampling designs. In the “random” sampling scheme, the sample contains 2,000 individuals that were randomly selected from the population regardless of their phenotypes. Thus, ascertainment is population based. In the “baseline” sampling scheme, we sampled 1,000 case subjects and 1,000 control subjects according to their outcome value at

baseline only. In the “sum” sampling scheme, individuals were stratified into three strata (1, 2, and 3) based on a total count that sums each subject’s response over time, where samples in stratum 1 never experienced the event of interest, i.e., $\sum_j Y_{ij} = 0$, samples in stratum 2 sometimes experienced the event, i.e., $0 < \sum_j Y_{ij} < n_i$, and samples in stratum 3 always experienced the event, i.e., $\sum_j Y_{ij} = n_i$. Following the outcome-dependent sampling design for longitudinal binary data (Schildcrout et al. 2018b), we selected 100, 1,800, and 100 individuals from the three strata respectively to oversample subjects who have response variation over the course of the study.

3.3. *Simulation Results*

To assess type I error, we tested for association at unlinked and unassociated SNPs. Table 1 gives the empirical type I error of the L-BRAT, RGMMAT, GEE, and GMMAT tests, based on 10^6 replicates, at the nominal type I error level α , for $\alpha = 0.05, 0.01, 0.001$, and 0.0001 . In all simulations, the type I error of the two retrospective tests, L-BRAT and RGMMAT, exhibited no inflation at any of the nominal levels considered. In contrast, the prospective GEE tests, regardless of the choice of working correlation, had inflated type I error at most of the nominal levels in all settings. This is likely due to the fact that the asymptotic distribution of robust sandwich variance estimators used in GEE are not well calibrated. The inflated type I error was also reported in longitudinal GWAS with quantitative traits using GEE (Sitlani et al. 2015). In GMMAT, the type I error was much lower than the nominal level when $\alpha = 0.05, 0.01, 0.001$, and 0.0001 . These results demonstrate that the two retrospective tests, L-BRAT and RGMMAT, are robust to trait model misspecification and ascertainment, whereas GEE has type I error inflation and GMMAT is overly conservative. Overall, the choice of the working correlation matrix does not have much impact on the type I error of the L-BRAT method.

To compare power, we considered five effect sizes at the two causal SNPs, and tested association between the trait and the first causal SNP. Empirical power was calculated at the significance level 10^{-3} , based on 1,000 simulated replicates. Figure 1 demonstrates the power results for each method. In all the simulation settings, the retrospective tests consistently had higher power than the prospective tests. The L-BRAT association tests under three different working correlation structures had similar power. The RGMMAT method also achieved high power. In contrast, the prospective GEE methods had the lowest power in all settings except under the baseline sampling and the liability threshold model, in which GMMAT performed the worst in power. Overall, we found that the baseline sampling scheme generated the highest power under different trait models, while the sum sampling scheme had a power gain over the random sampling scheme under the logistic mixed model, but was less powerful under the liability threshold model. These results suggest that L-BRAT and RGMMAT outperform the prospective tests, and the power of L-BRAT is not sensitive to the choice of the working correlation structure.

4. Analysis of Longitudinal Cocaine Use Data from VACS

We illustrated the utility of our proposed methods by analyzing a GWAS dataset of cocaine use from VACS (Justice et al. 2006). VACS is a multi-center, longitudinal observational study of HIV infected and uninfected veterans whose primary objective is to understand the risk of alcohol and other substance abuse in individuals with HIV infection. We analyzed longitudinal cocaine use in patient surveys collected at six clinic visits on 2,470 participants. Among them, 69.8% are African Americans (AAs), 19.3% are European Americans (EAs), and 10.9% are of other races. We considered the responses at each visit as 0 if individuals had never tried cocaine or had not used cocaine in the last year, and as 1 if individuals had used cocaine in the last year.

The proportion of case subjects at each visit ranges from 13.7% ($n = 192$) to 24.3% ($n = 526$), and the missing rate at each visit ranges from 3.0% to 44.2%.

All samples were genotyped on the Illumina OmniExpress BeadChip. After data cleaning, there are 2,458 individuals available for genotype imputation. IMPUTE2 (Howie et al. 2009) was used for imputation using the 1000 Genomes Phase 3 data as a reference panel. We excluded subjects who did not meet either of the following criteria: (1) completeness (i.e., proportion of successfully imputed SNPs) $> 95\%$ and (2) empirical self-kinship < 0.525 (i.e., empirical inbreeding coefficient < 0.05). Based on the above criteria, 2,231 individuals were retained in the analysis, with 2,114 males and 117 females, of whom 1,557 are AAs, 431 are EAs, and 243 are of other races. There are 1,433 individuals who had never used cocaine during the study period, 639 individuals who sometimes used cocaine, i.e., exhibited response variation, and 159 individuals who had used cocaine at least once every year over the course of the study. SNPs that satisfied all of the following quality-control conditions were included in the analysis: (1) call rate $> 95\%$, (2) Hardy-Weinberg χ^2 statistic P-value $> 10^{-6}$, and (3) MAF $> 1\%$. All together there are a final set of 10,215,072 SNPs retained in the analysis.

4.1. Association Analysis

Genome-wide association testing for longitudinal cocaine use was performed using L-BRAT, RGMMAT, and the prospective GEE and GMMAT tests in the entire VACS sample. Sex, age at baseline, HIV status, and time were included as covariates in the analysis. The top ten principal components (PCs) that explained 89.4% of the total genetic variation were included as covariates to control for population structure. We considered two working correlation structures: independent and AR(1). For the L-BRAT and RGMMAT methods, the GRM was calculated using the LD pruned SNPs with MAF > 0.05 .

To compare the performance of longitudinal association analysis with that of univariate analysis on the summary metrics of cocaine use in VACS, we considered two alternative cocaine phenotypes: baseline and trajectories. Longitudinal cocaine use trajectories were obtained using a growth mixture model that clusters longitudinal data into discrete growth trajectory curves (Muthén 2004). We fit a logistic model with a polynomial function of time. The number of groups was chosen based on the Bayesian information criterion (BIC). Each individual was then assigned to the trajectory with the highest probability of membership. Figure 2 shows the four cocaine use trajectory groups identified in the VACS sample. They were labeled as mostly never (0, $n = 1,682$), moderate decrease (1, $n = 296$), elevated chronic (2, $n = 86$), and mostly frequent (3, $n = 167$). We used CARAT, a case-control retrospective association test (Jiang et al. 2016), for the analysis of cocaine use at baseline, adjusted for sex, age at baseline, and HIV status. Cumulative logit model was used to test for association between the four ordered cocaine use trajectory groups and each of the SNPs, with adjustment for sex, age at baseline, HIV status, and the top ten PCs.

None of the retrospective tests exhibited evidence of inflation in the quantile-quantile (Q-Q) plot (Web Figure 1). The genomic control inflation factors were 0.993 and 0.991 for the L-BRAT genome scan under the independent and AR(1) working correlation, respectively, and 0.984 for the RGMMAT analysis. The prospective GEE tests showed some evidence of deflation in the Q-Q plot. The genomic control factors were 0.938 and 0.937 for the GEE tests under the independent and AR(1) working correlation. The most conservative test was GMMAT, with a genomic control factor 0.802.

Table 2 reports the results for SNPs for which at least one of the longitudinal tests gives a P-value $< 2 \times 10^{-7}$. Among them, the L-BRAT tests produced the smallest P-values, RGMMAT

and the trajectory-based analysis had comparable results, while GEE, GMMAT, and CARAT generated much larger P-values. The Manhattan plot of the smallest P-value from these tests in the VACS cocaine use data is shown in Web Figure 2. Among the top SNPs listed in Table 2, there are two SNPs, rs551879660 and rs150191017, located at 3p12 and 13q12 respectively, that reach the genome-wide significance ($P = 2.00 \times 10^{-8}$ and 3.77×10^{-8} , respectively). Each of these SNPs was reported to have MAF $< 1\%$ in the 1000 Genomes (MAF = 0.68% and 0.98%, respectively). The MAFs of the two SNPs were 1.2% and 1.1% in the entire VACS sample, respectively, and were slightly higher in the AA sample (MAF = 1.6% and 1.5%, respectively). Although both SNPs have MAF $> 1\%$, given the small sample size of VACS, there is limited information on them. SNP rs150191017 is located 31.5 kb from the gene *AL161616.2* which was reported to be associated with venlafaxine treatment response in a generalized anxiety disorder GWAS (Jung et al. 2017). A cluster of five SNPs in LD, rs76386683, rs114386843, rs186274502, rs376616438, and rs187855416, located at 9q33, showed association with longitudinal cocaine use ($P = 1.85 \times 10^{-7} - 1.93 \times 10^{-7}$). They are near *OR1L4*, an olfactory receptor gene that was reported to be associated with major depressive disorder (Wong et al. 2017). A cluster of olfactory receptor genes between *OR3A1* and *OR3A2* that belong to the olfactory receptor gene family were identified in a recent GWAS of cocaine dependence and related traits (Gelernter et al. 2014). The other three SNPs, rs188222191, rs1014278, rs75132056, are located at 5q21 ($P = 1.28 \times 10^{-7}$, 1.43×10^{-7} and 8.92×10^{-8} , respectively), close to the gene *EFNA5*, which was identified in several GWAS to be associated with bipolar disorder and schizophrenia (Wang et al. 2010). There was also evidence of association with rs114629793 ($P = 8.65 \times 10^{-8}$). This SNP is in an intron of the gene encoding *PSD3*, located at 8p22. Recently, two schizophrenia GWAS have identified association between *PSD3* and

schizophrenia (Goes et al. 2015; Li et al. 2017b), and one study has shown that *PSD3* is associated with paliperidone treatment response in schizophrenic patients (Li et al. 2017a). Gene network analysis revealed that *PSD3* is one of the differentially expressed hub genes that involve dysfunction of brain reward circuitry in cocaine use disorder (Ribeiro et al. 2017).

We further analyzed the data separately in each population, adjusted for the top ten PCs obtained within the group, and then combined the results from the three groups by meta-analysis using the optimal weights for score statistics that have essentially the same power as the inverse variance weighting (Zhou et al. 2011). Web Table 1 gives the results in the 1,557 AA samples. All the top twelve SNPs listed in Table 2 had a P-value $< 5 \times 10^{-5}$ in at least one of the longitudinal tests in AAs. L-BRAT consistently gave the smallest P-values among all the longitudinal tests. The results from the three groups (AAs, EAs and other races) were combined by meta-analysis. The meta-analysis P-values were of the same order of magnitude as that obtained from the entire sample adjusted for population structure for each longitudinal test (Table 3). All the top twelve SNPs listed in Table 2 had a meta-analysis P-value $< 8 \times 10^{-7}$ in at least one of the longitudinal tests. Among them, the L-BRAT test with either an independent or AR(1) working correlation gave the smallest meta-analysis P-values.

4.2. Pathway Analysis

We then performed pathway analysis on the top SNPs for which at least one of the longitudinal tests had a P-value $< 5 \times 10^{-5}$ using the Ingenuity Pathway Analysis (IPA). We identified two significant canonical pathways that belong to the neurotransmitters and nervous system signaling. The first one is the opioid signaling pathway ($P = 1.41 \times 10^{-4}$, adjusted $P = 0.010$), which plays an important role in opioid tolerance and dependence. Studies have shown that chronic administration of cocaine and opioids are associated with changes in dopamine

transporters and opioid receptors in various brain regions (Le Merrer et al. 2009; Soderman and Unterwald 2009). The second significant pathway is the axonal guidance signaling pathway ($P = 2.54 \times 10^{-4}$, adjusted $P = 0.012$), which is critical for neural development. The mRNA expression levels of axon guidance molecules have been found to be altered in some brain regions of cocaine-treated rats, which may contribute to drug abuse-associated cognitive impairment (Bahi and Dreyer 2005; Jassen 2006). Each of the two pathways remained significant when we performed pathway analysis, using the same P-value cutoff value to select top SNPs, based on the L-BRAT results generated under the independence and AR(1) working correlation, respectively. In contrast, only the opioid signaling pathway was significant based on the results from the GEE analysis using the independent working correlation, and only the axonal guidance signaling pathway was significant based on the RGMMAT results, whereas neither of them remained significant based on the GMMAT results and that from the GEE analysis with an AR(1) working correlation. These results demonstrate that L-BRAT provides more informative association results to help identify biological relevant pathways.

4.3. *eQTL Enrichment Analysis*

Lastly, we performed an enrichment analysis to see whether the top SNPs in our analysis are more likely to regulate brain gene expression. We considered the local expression quantitative trait loci (*cis*-eQTLs) reported in 13 human brain regions from the Genotype-Tissue Expression (GTEx) project (GTEx Consortium 2013, 2017), including amygdala, anterior cingulate cortex, caudate, cerebellar hemisphere, cerebellum, cortex, frontal cortex, hippocampus, hypothalamus, nucleus accumbens, putamen, spinal cord, and substantia nigra. Fisher's exact test was used to assess the enrichment of eQTLs ($FDR < 0.05$) in the top 2,778 SNPs for which at least one of the longitudinal tests had a P -value $< 10^{-4}$ in the VACS sample. Among the 13 brain regions,

amygdala is the only region in which eQTLs showed significant enrichment in our top SNP list (odds ratio = 2.06, $P = 3.0 \times 10^{-5}$).

4.4. *Replication of Top Findings*

We used an independent cocaine dependence case-control GWAS from the Yale-Penn study (Gelernter et al. 2014) to replicate the top findings from our longitudinal analysis results in VACS. The summary statistics from the Yale-Penn cocaine dependence GWAS were obtained. Note that the lifetime cocaine dependence diagnosis was made using the Semi-Structured Assessment for Drug Dependence and Alcoholism (SSADDA) (Pierucci-Lagha et al. 2005) which is different from the outcome used in VACS, and there were no longitudinal phenotype measures in Yale-Penn. Nevertheless, we performed pathway analysis using the SNP summary statistics of Yale-Penn to replicate the two pathways identified in the VACS sample. Among the top 2,778 SNPs for which at least one of the longitudinal tests had a P -value $< 10^{-4}$, we were able to retrieve 2,602 SNP summary statistics from Yale-Penn. Pathway analysis was conducted on the top 84 SNPs that had a P -value < 0.05 . Although none of the top twelve SNPs in Table 2 had a P -value < 0.05 in the Yale-Penn AA sample, each of the two pathways remained significant: the opioid signaling pathways ($P = 5.67 \times 10^{-4}$, adjusted $P = 3.77 \times 10^{-3}$) and the axonal guidance signaling ($P = 2.89 \times 10^{-4}$, adjusted $P = 2.97 \times 10^{-3}$).

5. Discussion

Longitudinal data can be used in GWAS to improve power for identification of genetic variants and environmental factors that influence complex traits over time. In this study, we have developed L-BRAT, a retrospective association testing method for longitudinal binary outcomes. L-BRAT is based on GEE, thus it requires assumptions on the mean but not the full distribution of the outcome. Correct specification of the covariance of repeated measurements within each

individual is not required, instead, a working covariance matrix is assumed. The significance of the L-BRAT association test is assessed retrospectively by considering the conditional distribution of the genotype at the variant of interest, given phenotype and covariate information, under the null hypothesis of no association. Features of L-BRAT include the following: (1) it is computationally feasible for genetic studies with millions of variants, (2) it allows both static and time-varying covariates to be included in the analysis, (3) it allows different individuals to have measurements at different time points, and (4) it has correct type I error in the presence of ascertainment and trait model misspecification. For comparison, we also propose a retrospective, logistic mixed model-based association test, RGMMAT, which requires specification of the full distribution of the outcome. Random effects are used to model dependence among observations for an individual. Like L-BRAT, RGMMAT is a retrospective analysis in which genotypes are treated as random conditional on the phenotype and covariates. As a result, RGMMAT is also robust to misspecification of the model for the phenotype distribution.

Through simulation, we demonstrated that the type I error of L-BRAT was well calibrated under different trait models and ascertainment schemes, whereas the type I error of the prospective GEE method was inflated relative to nominal levels. In the GLMM-based methods, GMMAT, a prospective analysis, was overly conservative, whereas the retrospective version, RGMMAT, was able to maintain correct type I error. We further demonstrated that the two retrospective tests, L-BRAT and RGMMAT, provided higher power to detect association than the prospective GEE and GMMAT tests under all the trait models and ascertainment schemes considered in the simulations. The choice of the working correlation matrix in L-BRAT resulted in little loss of power. We applied L-BRAT and RGMMAT to longitudinal association analysis of cocaine use in the VACS data, where we identified six novel genes that are associated with cocaine use. Moreover, our pathway analysis identified two significant pathways associated with

longitudinal cocaine use: the opioid signaling pathway and the axonal guidance signaling pathway. We were able to replicate both pathways in a cocaine dependence case-control GWAS from the Yale-Penn study. Lastly, we illustrated that the top SNPs identified by our methods are more likely to be the amygdala eQTLs in the GTEx data. The amygdala plays an important role in the processing of memory, decision-making, and emotional responses, and contributes to drug craving that leads to addiction and relapse (Hyman and Malenka 2001; Warlow et al. 2017). These findings verify that L-BRAT is able to detect important loci in a genome scan and to provide novel insights into the disease mechanism in relevant tissues.

The L-BRAT and RGMMAT methods are designed for single-variant association analysis of longitudinally measured binary outcomes. However, single-variant association tests in general have limited power to detect association for low-frequency or rare variants in sequencing studies. We have previously developed longitudinal burden test and sequence kernel association test, LBT and LSKAT, to analyze rare variants with longitudinal quantitative phenotypes (Wang et al. 2017). Both tests are based on a prospective approach. To extend L-BRAT and RGMMAT to rare variant analysis with longitudinal binary data, we could consider either a linear statistic or a quadratic statistic that combines the retrospective score test at each variant in a gene region. In addition, the genetic effect in L-BRAT and RGMMAT is assumed to be constant. We could consider an extension to allow for time-varying genetic effect so that the fluctuation of genetic contributions to the trait value over time is well calibrated.

ACKNOWLEDGEMENTS

This study was supported by the National Institute of Health grants K01 AA023321 and R21 AA022870. The authors appreciate the support of the Veterans Aging Cohort Study and Yale Center for Genome Analysis.

REFERENCES

- Bahi, A., and Dreyer, J.-L. (2005). Cocaine-induced expression changes of axon guidance molecules in the adult rat brain. *Molecular and Cellular Neuroscience* **28**, 275–291.
- Breslow, N. E., and Clayton, D. G. (1993). Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association* **88**, 9–25.
- Chen, H., Wang, C., Conomos, M. P., Stilp, A. M., Li, Z., Sofer, T., et al. (2016). Control for population structure and relatedness for binary traits in genetic association studies via logistic mixed models. *American Journal of Human Genetics* **98**, 653–666.
- Das, K., Li, J., Wang, Z., Tong, C., Fu, G., Li, Y., et al. (2011). A dynamic model for genome-wide association studies. *Human Genetics* **129**, 629–639.
- Furlotte, N. A., Eskin, E., and Eyheramendy, S. (2012). Genome-wide association mapping with longitudinal data. *Genetic Epidemiology* **36**, 463–471.
- Gelernter, J., Sherva, R., Koesterer, R., Almasy, L., Zhao, H., Kranzler, H. R., et al. (2014). Genome-wide association study of cocaine dependence and related traits: FAM53B identified as a risk gene. *Molecular Psychiatry* **19**, 717–723.
- Goes, F. S., Mcgrath, J., Avramopoulos, D., Wolyniec, P., Pirooznia, M., Ruczinski, I., et al. (2015). Genome-wide association study of schizophrenia in Ashkenazi Jews. *American Journal of Medical Genetics, Part B: Neuropsychiatric Genetics* **168**, 649–659.
- GTEx Consortium (2013). The genotype-tissue expression (GTEx) project. *Nature Genetics* **45**, 580–585.
- GTEx Consortium (2017). Genetic effects on gene expression across human tissues. *Nature* **550**, 204–213.
- Hayeck, T. J., Loh, P.-R., Pollack, S., Gusev, A., Patterson, N., Zaitlen, N. A., et al. (2017).

- Mixed model association with family-biased case-control ascertainment. *American Journal of Human Genetics* **100**, 31–39.
- Hayeck, T. J., Zaitlen, N. A., Loh, P.-R., Vilhjalmsen, B., Pollack, S., Gusev, A., et al. (2015). Mixed model with correction for case-control ascertainment increases association power. *American Journal of Human Genetics* **96**, 720–730.
- Howie, B. N., Donnelly, P., and Marchini, J. (2009). A flexible and accurate genotype imputation method for the next generation of genome-wide association studies. *PLoS Genetics* **5**, e1000529.
- Hyman, S. E., and Malenka, R. C. (2001). Addiction and the brain: the neurobiology of compulsion and its persistence. *Nature Reviews Neuroscience* **2**, 695–703.
- Jakobsdottir, J., and McPeck, M. S. (2013). MASTOR: Mixed-model association mapping of quantitative traits in samples with related individuals. *American Journal of Human Genetics* **92**, 652–666.
- Jassen, A. K. (2006). Receptor regulation of gene expression of axon guidance molecules: implications for adaptation. *Molecular Pharmacology* **70**, 71–77.
- Jiang, D., Mbatchou, J., and McPeck, M. S. (2015). Retrospective association analysis of binary traits: overcoming some limitations of the additive polygenic model. *Human Heredity* **80**, 187–195.
- Jiang, D., Zhong, S., and McPeck, M. S. (2016). Retrospective binary-trait association test elucidates genetic architecture of Crohn disease. *American Journal of Human Genetics* **98**, 243–255.
- Jung, J., Tawa, E. A., Muench, C., Rosen, A. D., Rickels, K., and Lohoff, F. W. (2017). Genome-wide association study of treatment response to venlafaxine XR in generalized anxiety disorder. *Psychiatry Research* **254**, 8–11.

- Justice, A. C., Dombrowski, E., Conigliaro, J., Fultz, S. L., Gibson, D., Madenwald, T., et al. (2006). Veterans aging cohort study (VACS): overview and description. *Medical care* **44**, S13–S24.
- Le Merrer, J., Becker, J. A. J., Befort, K., and Kieffer, B. L. (2009). Reward processing by the opioid system in the brain. *Physiological Reviews* **89**, 1379–1412.
- Li, Q., Wineinger, N. E., Fu, D.-J., Libiger, O., Alphs, L., Savitz, A., et al. (2017a). Genome-wide association study of paliperidone efficacy. *Pharmacogenetics and Genomics* **27**, 7–18.
- Li, Z., Chen, J., Yu, H., He, L., Xu, Y., Zhang, D., et al. (2017b). Genome-wide association analysis identifies 30 new susceptibility loci for schizophrenia. *Nature Genetics* **49**, 1576–1583.
- Londono, D., Chen, K.-M., Musolf, A., Wang, R., Shen, T., Brandon, J., et al. (2013). A novel method for analyzing genetic association with longitudinal phenotypes. *Statistical Applications in Genetics and Molecular Biology* **12**, 241–261.
- Meirelles, O. D., Ding, J., Tanaka, T., Sanna, S., Yang, H.-T., Dudekula, D. B., et al. (2013). SHAVE: Shrinkage estimator measured for multiple visits increases power in GWAS of quantitative traits. *European Journal of Human Genetics* **21**, 673–679.
- Muthén, B. (2004). Latent variable analysis: growth mixture modeling and related techniques for longitudinal data. In *The SAGE Handbook of Quantitative Methodology for the Social Sciences*, ed. D. Kaplan, Thousand Oaks, CA: Sage., pp. 346–369.
- Pierucci-Lagha, A., Gelernter, J., Feinn, R., Cubells, J. F., Pearson, D., Pollastri, A., et al. (2005). Diagnostic reliability of the semi-structured assessment for drug dependence and alcoholism (SSADDA). *Drug and Alcohol Dependence* **80**, 303–312.
- Ribeiro, E. A., Scarpa, J. R., Garamszegi, S. P., Kasarskis, A., Mash, D. C., and Nestler, E. J. (2017). Gene network dysregulation in dorsolateral prefrontal cortex neurons of humans with

- cocaine use disorder. *Scientific Reports* **7**, 5412.
- Schaffner, S. F., Foo, C., Gabriel, S., Reich, D., Daly, M. J., and Altshuler, D. (2005).
Calibrating a coalescent simulation of human genome sequence variation. *Genome Research* **15**, 1576–1583.
- Schildcrout, J. S., and Heagerty, P. J. (2008). On outcome-dependent sampling designs for
longitudinal binary response data with time-varying covariates. *Biostatistics* **9**, 735–749.
- Schildcrout, J. S., Schisterman, E. F., Aldrich, M. C., and Rathouz, P. J. (2018a). Outcome-
related, auxiliary variable sampling designs for longitudinal binary data. *Epidemiology* **29**,
58–66.
- Schildcrout, J. S., Schisterman, E. F., Mercaldo, N. D., Rathouz, P. J., and Heagerty, P. J.
(2018b). Extending the case-control design to longitudinal data: stratified sampling based on
repeated binary outcomes. *Epidemiology* **29**, 67–75.
- Sikorska, K., Rivadeneira, F., Groenen, P. J. F., Hofman, A., Uitterlinden, A. G., Eilers, P. H. C.,
et al. (2013). Fast linear mixed model computations for genome-wide association studies
with longitudinal data. *Statistics in Medicine* **32**, 165–180.
- Sitlani, C. M., Rice, K. M., Lumley, T., Mcknight, B., Cupples, L. A., Avery, C. L., et al. (2015).
Generalized estimating equations for genome-wide association studies using longitudinal
phenotype data. *Statistics in Medicine* **34**, 118–130.
- Soderman, A. R., and Unterwald, E. M. (2009). Cocaine-induced mu opioid receptor occupancy
within the striatum is mediated by dopamine D2 receptors. *Brain Research* **1296**, 63–71.
- Wang, K.-S., Liu, X.-F., and Aragam, N. (2010). A genome-wide meta-analysis identifies novel
loci associated with schizophrenia and bipolar disorder. *Schizophrenia Research* **124**, 192–
199.
- Wang, Z., Xu, K., Zhang, X., Wu, X., and Wang, Z. (2017). Longitudinal SNP-set association

- analysis of quantitative phenotypes. *Genetic Epidemiology* **41**, 81–93.
- Warlow, S. M., Robinson, M. J. F., and Berridge, K. C. (2017). Optogenetic central amygdala stimulation intensifies and narrows motivation for cocaine. *The Journal of Neuroscience* **37**, 8330–8348.
- Wong, M.-L., Arcos-Burgos, M., Liu, S., Vélez, J. I., Yu, C., Baune, B. T., et al. (2017). The PHF21B gene is associated with major depression and modulates the stress response. *Molecular Psychiatry* **22**, 1015–1025.
- Wu, X., and McPeck, M. S. (2018). L-GATOR: Genetic association testing for a longitudinally measured quantitative trait in samples with related individuals. *American Journal of Human Genetics* **102**, 574–591.
- Zhong, S., Jiang, D., and McPeck, M. S. (2016). CERAMIC: Case-control association testing in samples with related individuals, based on retrospective mixed model analysis with adjustment for covariates. *PLoS Genetics* **12**, e1006329.
- Zhou, B., Shi, J., and Whittemore, A. S. (2011). Optimal methods for meta-analysis of genome-wide association studies. *Genetic Epidemiology* **35**, 581–591.
- Zhou, W., Nielsen, J. B., Fritsche, L. G., Dey, R., Gabrielsen, M. E., Wofford, B. N., et al. (2018). Efficiently controlling for case-control imbalance and sample relatedness in large-scale genetic association studies. *Nature Genetics*, **50**, 1335–1341.

SUPPORTING INFORMATION

Additional supporting information referenced in Section 4 may be found online in the Supporting Information section at the end of the article. R package implementing L-BRAT and RGMMAT can be found at <https://github.com/ZWang-Lab/LBRAT>.

Table 1. Empirical type I error of L-BRAT, RGMMAT, GEE, and GMMAT, based on 10^6 replicates. Rates that are significantly larger than the nominal levels are in bold. Texts in the brackets following test statistics denote the working correlation structure. Specifically, L-BRAT(ind) and GEE(ind) denote the L-BRAT and GEE tests with an independent working correlation; L-BRAT(AR1) and GEE(AR1) denote the L-BRAT and GEE tests with an AR(1) working correlation; L-BRAT(mix) and GEE(mix) denote the L-BRAT and GEE tests with a mixture of exchangeable and AR(1) working correlation structure.

Analysis Type	Test	Nominal Level	Logistic Mixed Model			Liability Threshold Model		
			Random	Baseline	Sum	Random	Baseline	Sum
Prospective	GEE(ind)	0.05	5.38×10^{-2}	5.08×10^{-2}	5.27×10^{-2}	5.36×10^{-2}	5.19×10^{-2}	5.38×10^{-2}
		0.01	1.18×10^{-2}	1.04×10^{-2}	1.13×10^{-2}	1.17×10^{-2}	1.07×10^{-2}	1.17×10^{-2}
		0.001	1.32×10^{-3}	1.16×10^{-3}	1.23×10^{-3}	1.37×10^{-3}	1.14×10^{-3}	1.37×10^{-3}
		0.0001	1.67×10^{-4}	1.28×10^{-4}	1.43×10^{-4}	1.34×10^{-4}	1.36×10^{-4}	1.76×10^{-4}
	GEE(AR1)	0.05	5.36×10^{-2}	5.02×10^{-2}	5.26×10^{-2}	5.34×10^{-2}	5.17×10^{-2}	5.37×10^{-2}
		0.01	1.16×10^{-2}	1.04×10^{-2}	1.12×10^{-2}	1.16×10^{-2}	1.06×10^{-2}	1.17×10^{-2}
		0.001	1.31×10^{-3}	1.13×10^{-3}	1.21×10^{-3}	1.36×10^{-3}	1.14×10^{-3}	1.36×10^{-3}
		0.0001	1.73×10^{-4}	1.19×10^{-4}	1.37×10^{-4}	1.32×10^{-4}	1.35×10^{-4}	1.78×10^{-4}
	GEE(mix)	0.05	5.34×10^{-2}	5.07×10^{-2}	5.26×10^{-2}	5.34×10^{-2}	5.19×10^{-2}	5.37×10^{-2}
		0.01	1.17×10^{-2}	1.04×10^{-2}	1.13×10^{-2}	1.16×10^{-2}	1.07×10^{-2}	1.17×10^{-2}
		0.001	1.29×10^{-3}	1.17×10^{-3}	1.22×10^{-3}	1.38×10^{-3}	1.14×10^{-3}	1.36×10^{-3}
		0.0001	1.70×10^{-4}	1.29×10^{-4}	1.37×10^{-4}	1.31×10^{-4}	1.30×10^{-4}	1.70×10^{-4}
	GMMAT	0.05	3.89×10^{-2}	3.53×10^{-2}	4.76×10^{-2}	4.80×10^{-2}	4.89×10^{-2}	4.91×10^{-2}
		0.01	6.07×10^{-3}	5.24×10^{-3}	9.08×10^{-3}	9.29×10^{-3}	9.51×10^{-3}	9.33×10^{-3}
		0.001	4.29×10^{-4}	3.74×10^{-4}	7.84×10^{-4}	8.63×10^{-4}	8.96×10^{-4}	8.33×10^{-4}
		0.0001	2.20×10^{-5}	2.20×10^{-5}	6.80×10^{-5}	6.30×10^{-5}	9.10×10^{-5}	8.80×10^{-5}
Retrospective	L-BRAT(ind)	0.05	4.93×10^{-2}	4.91×10^{-2}	4.98×10^{-2}	5.01×10^{-2}	4.99×10^{-2}	4.98×10^{-2}
		0.01	9.45×10^{-3}	9.60×10^{-3}	9.84×10^{-3}	9.90×10^{-3}	9.75×10^{-3}	9.55×10^{-3}
		0.001	8.30×10^{-4}	9.78×10^{-4}	9.24×10^{-4}	9.55×10^{-4}	9.45×10^{-4}	8.78×10^{-4}
		0.0001	7.20×10^{-5}	9.50×10^{-5}	8.20×10^{-5}	8.20×10^{-5}	9.40×10^{-5}	9.20×10^{-5}
	L-BRAT(AR1)	0.05	4.93×10^{-2}	4.88×10^{-2}	4.97×10^{-2}	4.99×10^{-2}	4.98×10^{-2}	4.97×10^{-2}
		0.01	9.48×10^{-3}	9.72×10^{-3}	9.78×10^{-3}	9.84×10^{-3}	9.76×10^{-3}	9.55×10^{-3}
		0.001	8.26×10^{-4}	9.62×10^{-4}	9.22×10^{-4}	9.17×10^{-4}	9.47×10^{-4}	8.48×10^{-4}
		0.0001	8.80×10^{-5}	9.60×10^{-5}	8.20×10^{-5}	7.10×10^{-5}	1.02×10^{-4}	8.90×10^{-5}
	L-BRAT(mix)	0.05	4.93×10^{-2}	4.91×10^{-2}	4.99×10^{-2}	5.01×10^{-2}	4.98×10^{-2}	4.98×10^{-2}
		0.01	9.57×10^{-3}	9.61×10^{-3}	9.86×10^{-3}	9.88×10^{-3}	9.79×10^{-3}	9.54×10^{-3}
		0.001	8.35×10^{-4}	9.86×10^{-4}	9.26×10^{-4}	9.57×10^{-4}	9.37×10^{-4}	8.78×10^{-4}
		0.0001	8.20×10^{-5}	1.01×10^{-4}	8.60×10^{-5}	7.40×10^{-5}	9.70×10^{-5}	8.90×10^{-5}
	RGMMAT	0.05	4.72×10^{-2}	4.91×10^{-2}	4.98×10^{-2}	4.93×10^{-2}	4.99×10^{-2}	4.98×10^{-2}
		0.01	8.76×10^{-3}	9.64×10^{-3}	9.85×10^{-3}	9.63×10^{-3}	9.78×10^{-3}	9.55×10^{-3}
		0.001	7.20×10^{-4}	9.52×10^{-4}	9.09×10^{-4}	9.12×10^{-4}	9.43×10^{-4}	8.75×10^{-4}
		0.0001	6.80×10^{-5}	8.90×10^{-5}	8.20×10^{-5}	7.70×10^{-5}	9.10×10^{-5}	9.30×10^{-5}

Table 2. SNPs with P-value $< 2 \times 10^{-7}$ in at least one of the longitudinal tests in the entire VACS sample. The smallest P-value among all tests at the given SNPs are in bold. ^a CARAT applied to cocaine use at baseline, ^b Cumulative logit model applied to the four ordered cocaine use trajectory group.

Chr	Gene Region	SNP	Position	MAF	GEE (ind)	GEE (AR1)	GMMAT	L-BRAT (ind)	L-BRAT (AR1)	RGMMAT	CARAT ^a (BL)	CL ^b (traj)
3	<i>NIPA2P2</i>	rs551879660	75,146,492	0.012	1.87×10^{-4}	7.14×10^{-4}	9.07×10^{-4}	2.00×10^{-8}	3.19×10^{-6}	4.13×10^{-5}	5.78×10^{-4}	3.35×10^{-5}
5	<i>EFNA5</i>	rs188222191	105,411,547	0.042	6.86×10^{-6}	1.65×10^{-5}	8.87×10^{-5}	1.28×10^{-7}	4.17×10^{-7}	2.69×10^{-6}	8.95×10^{-5}	2.72×10^{-5}
		rs1014278	105,471,506	0.057	1.02×10^{-5}	1.10×10^{-5}	1.24×10^{-4}	1.50×10^{-7}	1.43×10^{-7}	4.88×10^{-6}	5.94×10^{-5}	3.00×10^{-5}
		rs75132056	105,480,442	0.05	1.05×10^{-5}	2.42×10^{-5}	1.89×10^{-4}	8.92×10^{-8}	2.89×10^{-7}	8.55×10^{-6}	2.59×10^{-4}	2.31×10^{-5}
8	<i>PSD3</i>	rs114629793	18,403,754	0.012	3.12×10^{-4}	4.73×10^{-4}	1.44×10^{-4}	8.65×10^{-8}	3.60×10^{-7}	2.82×10^{-6}	5.12×10^{-4}	3.06×10^{-6}
9	<i>OR1L4</i>	rs76386683	125,467,023	0.012	1.48×10^{-4}	9.15×10^{-5}	2.86×10^{-4}	1.03×10^{-6}	1.93×10^{-7}	5.92×10^{-6}	4.80×10^{-4}	3.30×10^{-6}
		rs114386843	125,469,425	0.012	1.47×10^{-4}	9.05×10^{-5}	2.82×10^{-4}	1.01×10^{-6}	1.88×10^{-7}	5.78×10^{-6}	4.75×10^{-4}	3.22×10^{-6}
		rs186274502	125,471,416	0.012	1.47×10^{-4}	9.05×10^{-5}	2.82×10^{-4}	1.01×10^{-6}	1.88×10^{-7}	5.78×10^{-6}	4.75×10^{-4}	3.22×10^{-6}
		rs376616438	125,472,267	0.012	1.44×10^{-4}	8.95×10^{-5}	2.77×10^{-4}	9.79×10^{-7}	1.85×10^{-7}	5.62×10^{-6}	4.79×10^{-4}	3.20×10^{-6}
		rs187855416	125,474,459	0.012	1.44×10^{-4}	8.95×10^{-5}	2.77×10^{-4}	9.79×10^{-7}	1.85×10^{-7}	5.62×10^{-6}	4.79×10^{-4}	3.20×10^{-6}
11	<i>AP000851.1</i>	rs139780693	102,509,700	0.03	2.60×10^{-5}	1.04×10^{-5}	2.78×10^{-4}	5.83×10^{-7}	1.26×10^{-7}	1.35×10^{-5}	1.06×10^{-4}	2.00×10^{-6}
13	<i>AL161616.2</i>	rs150191017	31,962,649	0.011	4.26×10^{-5}	9.72×10^{-5}	7.32×10^{-5}	3.77×10^{-8}	3.09×10^{-7}	7.87×10^{-7}	3.74×10^{-4}	5.48×10^{-7}

Table 3. Meta-analysis results of the top twelve SNPs from Table 2 in the VACS data. The smallest P-value among all tests at the given SNPs are in bold.

Chr	Gene Region	SNP	Position	GEE (ind)	GEE (AR1)	GMMAT	L-BRAT (ind)	L-BRAT (AR1)	RGMMA T
3	<i>NIPA2P2</i>	rs551879660	75,146,492	1.81×10^{-4}	5.86×10^{-4}	8.98×10^{-4}	5.26×10^{-8}	6.41×10^{-6}	6.49×10^{-5}
5	<i>EFNA5</i>	rs188222191	105,411,547	7.57×10^{-6}	1.28×10^{-5}	1.80×10^{-4}	2.55×10^{-7}	5.52×10^{-7}	1.10×10^{-5}
		rs1014278	105,471,506	1.26×10^{-5}	8.44×10^{-6}	3.15×10^{-4}	1.03×10^{-6}	5.59×10^{-7}	2.44×10^{-5}
		rs75132056	105,480,442	1.31×10^{-5}	2.00×10^{-5}	4.24×10^{-4}	7.31×10^{-7}	1.27×10^{-6}	3.56×10^{-5}
8	<i>PSD3</i>	rs114629793	18,403,754	2.92×10^{-4}	4.31×10^{-4}	1.66×10^{-4}	1.79×10^{-7}	7.98×10^{-7}	6.83×10^{-6}
9	<i>ORIL4</i>	rs76386683	125,467,023	1.44×10^{-4}	8.78×10^{-5}	3.75×10^{-4}	2.32×10^{-6}	5.12×10^{-7}	1.46×10^{-5}
		rs114386843	125,469,425	1.42×10^{-4}	8.62×10^{-5}	3.68×10^{-4}	2.25×10^{-6}	4.97×10^{-7}	1.41×10^{-5}
		rs186274502	125,471,416	1.42×10^{-4}	8.62×10^{-5}	3.68×10^{-4}	2.25×10^{-6}	4.97×10^{-7}	1.41×10^{-5}
		rs376616438	125,472,267	1.39×10^{-4}	8.51×10^{-5}	3.60×10^{-4}	2.18×10^{-6}	4.86×10^{-7}	1.37×10^{-5}
		rs187855416	125,474,459	1.39×10^{-4}	8.51×10^{-5}	3.60×10^{-4}	2.18×10^{-6}	4.86×10^{-7}	1.37×10^{-5}
11	<i>AP000851.1</i>	rs139780693	102,509,700	1.15×10^{-5}	4.16×10^{-6}	1.07×10^{-4}	4.04×10^{-7}	6.05×10^{-8}	4.41×10^{-6}
13	<i>ALI61616.2</i>	rs150191017	31,962,649	3.55×10^{-5}	6.77×10^{-5}	1.26×10^{-4}	6.68×10^{-8}	5.80×10^{-7}	3.12×10^{-6}

Figure 1. Empirical power of L-BRAT, RGMMAT, GEE, and GMMAT. Power is based on 1,000 simulated replicates at five time points with $\alpha = 10^{-3}$. In the upper panel, the trait is simulated by the logistic mixed model, and in the lower panel, it is by the liability threshold model. Power results are demonstrated in samples of 2,000 individuals according to three different ascertainment schemes: random, baseline, and sum. This figure appears in color in the electronic version of this article.

Figure 2. Group-based cocaine use trajectories in VACS. Dashed lines represent the estimated trajectories, solid lines represent the observed mean cocaine use for each trajectory group. Time is the number of years since the baseline visit.