

Published in final edited form as:

Genomics. 2009 January ; 93(1): 5–9. doi:10.1016/j.ygeno.2008.05.002.

Estimation and testing of gene-environment interactions in family-based association studies

Heather J. Cordell

Institute of Human Genetics Newcastle University, UK

Abstract

Gene-environment interactions are of interest in genetic association studies for several reasons. Firstly, the power to detect genetic effects may be substantially decreased if those effects differ according to environmental exposure, and if no account is taken of this interaction with environmental exposure in the analysis. Secondly, such interactions may indicate a phenomenon of genuine biological interest (whereby a particular genetic effect only operates in the presence of an environmental trigger, or vice versa), understanding of which can lead us to a greater understanding of possible mechanisms and pathways in disease progression. Here I discuss the testing and estimation of gene-environment interactions via the case/pseudocontrol and related approaches. As originally proposed, the case/pseudocontrol approach applies to case/parent trios with no missing genotype data. I discuss some recent extensions that allow larger pedigree structures with some missing genotype data, and present computer simulations to compare the performance of several competing approaches.

Keywords

missing data; imputation; conditional likelihood

Introduction

Complex genetic diseases are, by definition, believed to result from the interplay of numerous different genetic and environmental factors. Interactions between such factors could account for the relatively modest successes in the detection of disease-predisposing genetic variants for common, complex diseases [1]. If genetic and environmental factors interact to cause disease susceptibility, the power to detect such effects, even in the current generation of large-scale, well-powered, genome-wide studies [2], may be compromised unless one stratifies by, or in some other way takes account of, the other factor(s) involved. For this reason, there is a growing interest in modeling gene-gene and gene-environment interactions, both in candidate gene and genome-wide studies. An additional motivation for investigation of interactions is a belief that these represent a phenomenon of biological interest, the elucidation of which may help uncover possible mechanisms and pathways in disease progression.

From a statistical point of view, interaction simply signifies departure from a linear model describing how two predictors (x_1 and x_2 , say) predict an outcome variable (y , say). This is perhaps most easily understood when y represents some quantitative trait (such as height, for example) and x_1 and x_2 are binary indicator variables representing the presence or absence

of some predictive factors. If the data is well fitted by a linear model $y = \mu + \beta_1 x_1 + \beta_2 x_2$ (where β_1 and β_2 are regression coefficients, representing the effects due to variables x_1 , x_2 , and μ is some baseline mean trait value that is expected in the absence of either factor) then the variables x_1 and x_2 are said to not interact with regards to predicting y . If, on the other hand, the data is better represented by a model that includes additional terms, such as $y = \mu + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_1 x_2$, then we say that x_1 and x_2 interact. More complex linear models may be postulated when the predictor variables take on several different levels or are measured on a quantitative scale and/or the outcome variable is qualitative or dichotomous (such as indicating presence/absence of disease). For a disease outcome and case/control data, the usual approach is to model the log odds of disease $\ln[p/(1-p)]$ (where p represents the probability of an individual becoming diseased) as a linear function of the relevant predictor variables [3]. For example, we might model the log odds as $\ln[p/(1-p)] = \beta_0 + \beta_e x_1 + \beta_g x_2 + \beta_{ge} x_1 x_2$, where x_1 and x_2 are binary indicator variables representing presence or absence of environmental and genetic exposures respectively, β_e and β_g are regression coefficients representing the environmental and genetic main effects, and β_{ge} represents a gene-environment interaction term [4].

The concept of interaction as departure from a linear model for the main effects of two variables is visualised pictorially in Figure 1. Expected trait values are shown for two different levels of a binary environmental exposure ('low' or 'high') and a three level genotype (such as might occur at a single nucleotide polymorphism). Figure 1A shows a situation where the genetic and environmental variables do not interact with regards to prediction of the trait: at each genetic level the effect of changing environment is to shift the trait mean by a constant amount, while at each level of the environment, the shift in trait mean between the different genotypes also remains constant. Figure 1 B-D on the other hand show varying types of interaction: in Figure 1B the difference in trait mean between the different genotypes is seen to be much greater when environment is 'high' than when environment is 'low', and, equivalently, the effect on the trait mean of moving from the 'low' to the 'high' environment is seen to be stronger for individuals with genotype g_3 than for those with g_1 . In Figure 1C this pattern is taken to the extreme in that there is no effect of genotype on trait at all in the 'low' environment. In Figure 1D the effect of genotype on trait is seen to be reversed in the 'low' environment as compared to the 'high' environment, sometimes referred to as a crossover model.

The difficulties in biological interpretation of statistical interaction are well-known [5-7] and result partly from the fact that statistical interaction is not invariant to transformations of scale of the outcome variable. For example, simply taking a monotonic transformation $\text{trait} \rightarrow \text{trait}^3/50000$ can convert the diagram in Figure 1A to the one shown in Figure 1B. Therefore, two variables which do not interact with regards to how they predict trait as measured on the original scale, may well interact with regards to how they predict $\text{trait}^3/50000$. This complicates biological interpretation of statistical interaction, unless the required scale of measurement is 'obvious' (so that we know there is only one particular scale in which we are interested) or unless an interaction term would be required on every scale (as would be the case in Figure 1C).

In spite of difficulties in interpretation, consideration of interaction terms may be warranted on the grounds of increasing power, given that the trait will generally be modeled on some particular scale, and interaction effects on that scale may well exist. Inclusion of interaction terms also allows improved predictive value of a model. Kraft et al. [4] showed that if the focus of a study is the detection of genetic effects, an appealing procedure is to fit a model that includes main effects and interactions and conduct a joint test for marginal (genetic main effect) association and gene-environment interaction (i.e. perform a test of association *allowing for* interaction). In the context of the linear model $\ln[p/$

$(1-p) = \beta_0 + \beta_e x_1 + \beta_g x_2 + \beta_{ge} x_1 x_2$, this joint test is a 2 df test of $\beta_g = \beta_{ge} = 0$. Although not universally most powerful, the joint test is nearly optimal over a wide range of plausible penetrance models.

Family-based association studies are a popular alternative to case/control studies for the detection of genetic effects. Although limited by sample size (since families are generally harder to collect than unrelated cases and controls), families have some advantages over case/control samples, allowing the construction of tests that are generally robust to population stratification [8], and the examination of potentially more interesting effects such as those due to maternal genotype and/or imprinting [9-12]. A popular design is to collect cases and their parents, who may be analysed using methods based on the transmission of alleles from heterozygous parents to affected offspring [13]. A more general method, that allows the fitting of linear regression models similar to those used in case/control studies, is the case/pseudocontrol approach [12,14], which conditions on the observed parental genotypes and constructs sets of matched ‘controls’ for the affected offspring from the untransmitted parental genotypes. This approach builds on previous methods for testing and estimation of genotype relative risks at a single locus [15,16] by extending these methods to allow for haplotype associations, gene-gene and gene-environment interactions, maternal genotype and parent-of-origin effects. In this approach, any environmental variables possessed by the case are copied over to the pseudocontrols, with the result that we cannot assess main effects of environment, but we can assess genetic effects and gene-environment interactions. As originally proposed, the case/pseudocontrol approach deals with missing data (such as unknown haplotype configurations due to phase uncertainty) through a complex conditioning argument [12,14], similar to the ‘conditioning on sufficient’ statistic approach used in the FBAT [17,18] program. A more efficient approach is to model the full likelihood (rather than the likelihood conditional on parental genotypes) and to account for missing data through use of missing data likelihood [8] or multiple imputation [19] approaches. Extensions to nuclear families with more than one affected offspring, and to larger pedigrees, can be derived either by conditioning on the identity-by-descent of alleles of related individuals [20] [8] or by use of an empirical variance estimate [20] that is robust to genotype correlations among related individuals.

Results and Discussion

Table 1 shows the results of simulations of case/parent trios to assess the performance of several different methods for testing and estimation of haplotype effects that influence, in conjunction with a binary environmental exposure, a disease outcome. Under the null, all methods give correct type 1 error rates of approximately 5%. Correct type 1 error rates are also seen when testing for gene-environment interaction if the data is generated under a main effects model (so that no interaction effects exist). Regardless of the true disease model, all methods are found to give unbiased parameter estimation and correct 95% confidence interval coverage when both main genetic (haplotype) effects and haplotype-environment interactions are included in the regression analysis (i.e. a ‘joint analysis’ is performed) although note that the 95% confidence interval coverage for the MI-TDT method is slightly over-conservative (i.e. greater than 95%). When a ‘marginal analysis’ (including only haplotype effects in the regression model) is performed, unbiased parameter estimation is seen when the marginal model is in fact correct (e.g. under null and main effects models). For models where a gene-environment interaction term exists, so that the marginal analysis model is misspecified, the value of the parameter estimate $\exp(\beta_g)$ (which equals the relative risk for the 2-2 haplotype) is seen to be biased towards the (unmodeled) value of the interaction term $\exp(\beta_{ge})$.

In all situations, the three methods (MI-TDT, Unphased, Unph-Pa) that try to estimate or reconstruct missing haplotype data show considerably higher power than the case/pseudocontrol approach implemented in the pseudoc program, which uses only families and pseudocontrols in which the haplotype configurations can be reconstructed with certainty. Power is also higher for a specific test of the effect of haplotype 2-2 (a 1df test, or a 2df test if tested jointly together with a haplotype environment-interaction) than for a 3df (or 6df if tested jointly together with haplotype-environment interactions) test of any difference in risk conferred by the four possible haplotypes. This is not unexpected, given the fewer degrees of freedom in the haplotype specific test. In practice, this higher power would only be achievable if we had reason to believe a priori that it is haplotype 2-2 that confers the differential disease risk, and so wish to test only this specific haplotype.

Our results illustrate the observation by Kraft et al. [4] that performing a joint test of genotype and genotype interaction can be a powerful strategy for detection of genetic factors predisposing to disease. When gene-environment interaction effects exist, the joint test has higher power than either a marginal test (in which only genetic effects are included in the regression model) or a test of the gene-environment interaction term alone. When gene-environment interaction effects do not exist, the joint test loses only a little power compared to the optimal (in this situation) marginal test of genetic effects only.

In the simulations presented here, the power of the three approaches (MI-TDT, Unphased, Unph-Pa) that try to estimate or reconstruct missing data appears to be very similar and, on this basis, there is little to choose between them. However, these simulations were conducted assuming that the data arise from a homogeneous population. In the presence of population stratification, the performance (and in particular the type 1 error) of the methods considered here is likely to vary considerably. The ‘additional conditioning event’ approach implemented in pseudoc, like the conditioning on sufficient statistic approach implemented in FBAT [17,18], should provide complete robustness to population stratification, but, as we have seen, this comes at the expense of power. A locally optimal approach that also provides complete robustness to population stratification has been proposed by Allen and Satten [21]; however this too loses power [8] when there is a strong genetic effect and a high proportion of missing data. The missing data likelihood approach implemented in Unphased, particularly when parental association parameters are modeled, can considerably improve the power at the expense of only a small increase in type 1 error [8] when there is missing genotype data and population stratification.

The multiple imputation approach used here provides only partial protection from population stratification [19]. However, it does have some advantages over the missing data likelihood approach in terms of its flexibility. In the multiple imputation approach, once multiple ‘complete’ imputed data sets have been constructed, we are free to fit whatever models we wish via analysis in a standard statistical package such as Stata or R (using methods from the multiple imputation literature [22,23] to combine estimates across the different multiple imputed data sets, and to perform tests). This allows us to fit models and perform tests that have not been implemented in specialised genetic analysis packages, such as the joint test of genotype main effect and gene-environment interaction, or the 3df test of haplotype interaction, neither of which are currently implemented in the Unphased software.

A danger of methods that reconstruct missing data probabilistically is that one is tempted to treat the reconstructed data as it were actually observed. Unless the uncertainty in the reconstructed missing data is appropriately allowed for in the analysis, one runs the risk of overestimating the amount of information actually available in the observed sample [24], resulting in an inflation of type 1 error and anti-conservative confidence intervals. Standard statistical theory, borne out by our simulations, predicts that inference based on a proper

missing data likelihood, as is used in the Unphased software, should appropriately allow for the uncertainty in haplotype reconstruction. Similarly, multiple imputation under the alternative hypothesis, using an IP (imputation/posterior sampling) algorithm [25] to sample the full Bayesian posterior distribution of haplotype data given the observed genotype and phenotype data, as is used in the MI-TDT software [19], should appropriately allow for any uncertainty in haplotype reconstruction (uncertainty given the complete-data model parameters, as well as uncertainty about these unknown model parameters). This is achieved through the generation of multiple reconstructed (imputed) complete data sets in which uncertainty is reflected through multiple alternative (differing) reconstructions (as output from the IP algorithm, essentially a Gibbs sampler), which must then be combined appropriately to provide the final inference [22].

It is common in the multiple imputation literature to use a relatively small number (e.g. 3-10) of imputed data sets to perform this final step. In the genetic context described here, and in previous investigations [19,26], the use of 10 imputed data sets appeared to be adequate, provided the level of missing data was not too large (up to about 30% genotypes missing). As pointed out by Nicolae et al. [24], the amount of information in an incomplete data set depends not only on the relationship between the observed and missing data, but on the hypothesis and test statistic being evaluated. Investigation of the sensitivity of the multiple imputation procedure to the number of imputed data sets generated, as well as to the use of different hypothesis tests, would be an interesting topic for further investigation. In theory, the use of larger numbers of imputed data sets should provide inferences that are better calibrated [22] as well as providing a better reflection of the inherent uncertainty in the haplotype reconstruction.

Methods

Simulations were conducted to examine the performance of several different family-based methods for testing and estimation of haplotype effects that influence, in conjunction with a binary environmental exposure, a disease outcome. In each case, 1000 case/parent trios were simulated under a model that assumed disease risk could be influenced by a two-marker haplotype acting in conjunction with a possible environmental exposure. The two genetic markers were both diallelic, leading to four possible haplotypes, denoted 1-1, 1-2, 2-1, 2-2 (here i - j denotes the occurrence of allele i at locus 1 in coupling with allele j at locus 2). Four different disease models were considered: 1) a null model in which there were no genetic effects but the environmental exposure (assumed to be at 30% frequency in the population) multiplied the offspring disease risk by a factor of two; 2) a main effects model in which the environmental exposure multiplied the offspring disease risk by two and each 2-2 haplotype multiplied the offspring disease risk by 1.5; 3) a pure interaction model in which each 2-2 haplotype multiplied the offspring disease risk by 1.5 only in the presence of the environmental trigger; and 4) a model in which main effects of both genotype and environment and their interaction all contributed to disease risk. In each case, 15% of genotypes in the parents and offspring were randomly set to be missing, in addition to missing data generated in the form of undeterminable phase (haplotype) resolutions.

The methods evaluated were: 1) the case/pseudocontrol approach [14] as implemented in the Stata program pseudocc (part of the genassoc [27] package); 2) the multiple imputation approach [19] implemented in the MI-TDT program; and 3) the missing data likelihood-based approach [8] implemented in Unphased program. The case/pseudocontrol approach deals with the missing data through use of an additional conditioning event [14] that means we only use families and pseudocontrols for which the haplotype configuration is inferrable; in addition the implementation in pseudocc only uses families in which there are no missing genotypes. The multiple imputation approach uses an iterative procedure to repeatedly fill in

the missing data (missing genotypes or phase resolutions) as described in [19]. In order to allow the fitting of gene-environment interaction effects, imputation here was performed within classes defined by the offspring's environmental exposure. This resulted in the generation of ten complete data sets that were then analysed using the Stata program *mim* [28] in order to combine estimates and construct tests [22,23]. The likelihood approach implemented in *Unphased* deals with missing data through direct maximisation of a full missing data likelihood, modeling genetic association parameters separately in parents and offspring [8]. The default is to assume that there is no genetic association in the parents; we also considered use of the *-parentrisk* (*-pa*) option in *Unphased*, in which association parameters are estimated (separately) in the parents as well as in the offspring.

For each disease model, 1000 simulation replicates were used to examine the bias (the difference between the true and the expected values of the relative risks and log relative risks) for the genetic main effects and gene-environment interactions, the coverage of the 95% confidence intervals for these parameters (which, if a method is working correctly, should equal 95%) and the powers and type 1 errors i.e. the probability of declaring a significant result (at the 5% significance level) when the null hypothesis is false (for power) or true (for type 1 error).

Acknowledgments

Support for this work was provided by the Wellcome Trust (Grant reference 074524)

References

1. Culverhouse R, Suarez BK, Lin J, Reich T. A perspective on epistasis: limits of models displaying no main effect. *American Journal of Human Genetics*. 2002; 70:461–471. [PubMed: 11791213]
2. WTCCC. Genome-wide association study of 14,000 cases of seven common diseases and 3,000 shared controls. *Nature*. 2007; 447:661–678. [PubMed: 17554300]
3. McCullagh, P.; Nelder, JA. *Generalized Linear Models*. Chapman & Hall; 1989.
4. Kraft P, Yen YC, Stram DO, Morrison J, Gauderman WJ. Exploiting gene-environment interaction to detect genetic associations. *Human Heredity*. 2007; 63:111–119. [PubMed: 17283440]
5. Thompson WD. Effect modification and the limits of biological inference from epidemiologic data. *Journal of Clinical Epidemiology*. 1991; 44:221–232. [PubMed: 1999681]
6. Phillips PC. The language of gene interaction. *Genetics*. 1998; 149:1167–1171. [PubMed: 9649511]
7. Cordell HJ. Epistasis: what it means, what it doesn't mean, and statistical methods to detect it in humans. *Human Molecular Genetics*. 2002; 11:2463–2468. [PubMed: 12351582]
8. Dudbridge F. Likelihood-based association analysis for nuclear families and unrelated subjects with missing genotype data. *Human Heredity*. 2008 in press.
9. Weinberg CR, Wilcox AJ, Lie RT. A log-linear approach to case-parent-triad data: assessing effects of disease genes that act either directly or through maternal effects and that may be subject to parental imprinting. *American Journal of Human Genetics*. 1998; 62:969–978. [PubMed: 9529360]
10. Weinberg CR. Methods for detection of parent-of-origin effects in genetic studies of case-parents triads. *American Journal of Human Genetics*. 1999; 65:229–235. [PubMed: 10364536]
11. Sinsheimer JS, Palmer CG, Woodward JA. Detecting genotype combinations that increase risk for disease: maternal-fetal genotype incompatibility test. *Genetic Epidemiology*. 2003; 24:1–13. [PubMed: 12508251]
12. Cordell HJ, Barratt BJ, Clayton DG. Case/pseudocontrol analysis in genetic association studies: a unified framework for detection of genotype and haplotype associations, gene-gene and gene-environment interactions and parent-of-origin effects. *Genetic Epidemiology*. 2004; 26:167–185. [PubMed: 15022205]

13. Spielman RS, McGinnis RE, Ewens WJ. Transmission test for linkage disequilibrium: The insulin gene region and insulin-dependent diabetes mellitus. *American Journal of Human Genetics*. 1993; 52:506–516. [PubMed: 8447318]
14. Cordell HJ, Clayton DG. A unified stepwise regression procedure for evaluating the relative effects of polymorphisms within a gene using case/control or family data: application to HLA in type 1 diabetes. *American Journal of Human Genetics*. 2002; 70:124–141. [PubMed: 11719900]
15. Schaid DJ, Sommer SS. Genotype relative risks: methods for design and analysis of candidate-gene association studies. *American Journal of Human Genetics*. 1993; 53:1114–1126. [PubMed: 8213835]
16. Schaid DJ. General score tests for associations of genetic markers with disease using cases and their parents. *Genetic Epidemiology*. 1996; 13:423–449. [PubMed: 8905391]
17. Laird NM, Horvath S, Xu X. Implementing a unified approach to family based tests of association. *Genetic Epidemiology Suppl*. 2000; 19:S36–S42.
18. Rabinowitz D, Laird NM. A unified approach to adjusting association tests for population admixture with arbitrary pedigree structure and arbitrary missing marker information. *Human Heredity*. 2000; 50:211–223. [PubMed: 10782012]
19. Croiseau P, Genin E, Cordell HJ. Dealing with missing data in family-based association studies: a multiple imputation approach. *Human Heredity*. 2007; 63:229–238. [PubMed: 17347570]
20. Lake SL, Blacker DB, Laird NM. Family-based tests of association in the presence of linkage. *American Journal of Human Genetics*. 2000; 67:1515–1525. [PubMed: 11058432]
21. Allen AS, Satten GA. Inference on haplotype/disease association using parent-affected-child data: the projection conditional on parental haplotypes method. *Genetic Epidemiology*. 2007; 31:211–223. [PubMed: 17266114]
22. Schafer, JL. *Analysis of incomplete multivariate data*. Chapman & Hall/CRC; 1997.
23. Li KH, Raghunathan TE, Rubin DB. Large sample significance levels from multiply imputed data using moment-based statistics and an f reference distribution). *Journal of the American Statistical Association*. 1987; 86:1065–1073.
24. Nicolae DL, Meng X-L, Kong A. Quantifying the fraction of missing information for hypothesis testing in statistical and genetic studies. *Statistical Science*. 2008 in press.
25. Tanner MA, Wong WH. The calculation of posterior distributions by data augmentation (with discussion). *Journal of the American Statistical Association*. 1987; 82:528–550.
26. Cordell HJ. Estimation and testing of genotype and haplotype effects in case-control studies: comparison of weighted regression and multiple imputation procedures. *Genetic Epidemiology*. 2006; 30:259–75. [PubMed: 16496312]
27. Clayton DG. genassoc: Stata module for analysis of genetic association studies. Available at <http://www-gene.cimr.cam.ac.uk/clayton/software/stata/README.txt>
28. Galati, JC.; Royston, P.; Carlin, JB. Mim: Stata module to analyse and manipulate multiply imputed datasets. *Statistical Software Components*; Boston College Department of Economics: Feb. 2007 Available at <http://ideas.repec.org/c/boc/bocode/s456825.html>

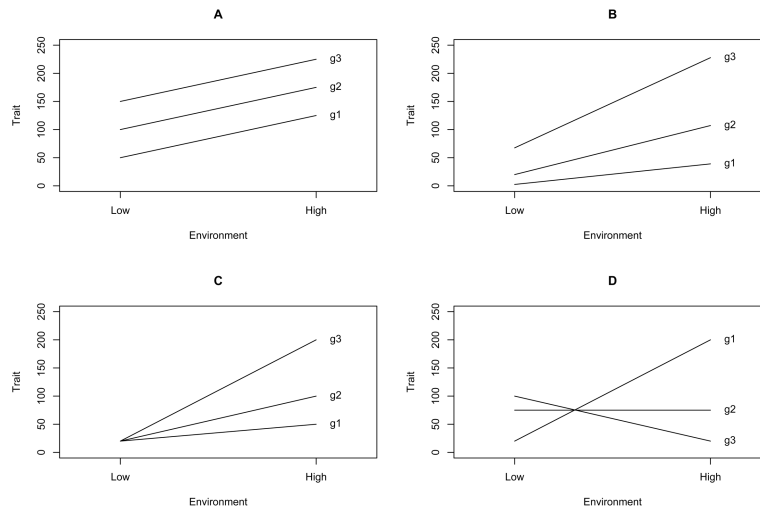


Figure 1. Visualisation of non-interaction and interaction models for the joint effects of genotype and environment on a quantitative trait. A. A non-interaction model. B. An interaction model. C. An extreme interaction model. D. A crossover interaction model.

Results of simulation study

Table 1

Simulation model	True values $\exp(\beta_g)$	Method ^a	Marginal analysis			Joint analysis									
			$\exp(\beta_g)$	Type I error or Power ^b	$\exp(\beta_g)$	Est ^e	Bias ^f	C β ^g	Est ^e	Bias ^f	C β ^g	Type I error or Power (joint test)	Type I error or Power (interaction)		
	$\exp(\beta_g)$		Est ^e	3df	1df	Est ^e	Bias ^f	C β ^g	Est ^e	Bias ^f	C β ^g	6df	2df	3df	1df
Null model	1.0	Pseudoccc	1.02	5.5	4.9	1.03	0.01	95.1	1.03	-0.01	94.3	3.7	4.2	3.8	5.7
		MI-TDT	1.01	5.0	5.2	1.02	0.01	97.0	1.01	0.00	97.1	5.3	5.5	4.9	5.5
		Unphased	1.01	5.5	5.6	1.02	0.01	94.9	1.01	0.00	94.8	-	-	-	5.3
		Unph-pa	1.01	5.5	5.7	1.02	0.01	94.9	1.01	0.00	94.6	-	-	-	5.4
Main effects	1.5	Pseudoccc	1.53	77.4	86.4	1.56	0.02	95.3	1.02	-0.01	95.7	62.8	78.6	5.1	4.3
		MI-TDT	1.51	98.1	99.8	1.52	0.00	97.3	1.01	0.00	97.7	94.5	99.5	4.1	4.2
		Unphased	1.52	99.5	99.8	1.53	0.01	94.6	1.01	0.00	95.5	-	-	-	4.5
		Unph-pa	1.52	99.5	99.8	1.53	0.01	94.6	1.01	0.00	95.7	-	-	-	4.4
Pure interaction	1.0	Pseudoccc	1.18	15.7	20.3	1.02	0.01	95.4	1.57	0.00	95.5	18.5	30.9	18.9	28.0
		MI-TDT	1.17	20.1	44.7	1.01	0.01	98.2	1.51	-0.01	97.5	42.6	73.1	34.5	58.9
		Unphased	1.18	34.0	45.9	1.02	0.01	95.4	1.52	0.00	94.3	-	-	-	61.2
		Unph-pa	1.18	33.9	45.9	1.02	0.01	95.6	1.52	0.00	94.2	-	-	-	61.3
Both (main plus interaction effects)	1.5	Pseudoccc	1.53	75.4	86.2	1.25	0.02	94.9	1.56	0.00	95.8	70.1	86.2	19.7	28.5
		MI-TDT	1.51	98.3	99.3	1.22	0.01	97.9	1.51	0.00	97.8	98.1	99.7	39.3	64.4
		Unphased	1.52	99.8	99.9	1.22	0.01	95.6	1.52	0.00	95.1	-	-	-	66.5
		Unph-pa	1.52	99.8	99.9	1.22	0.01	95.8	1.52	0.00	95.3	-	-	-	66.4

^aSee text for details. Unph-pa refers to the Unphased program [8] with parental and offspring association parameters estimated separately

^bProbability of reaching 5% significance level, for either a 3df test of any difference between the risks conferred by the four haplotypes, or a 1df test of the specific risk conferred by haplotype 2-2. Type I errors are shown in bold.

^cProbability of reaching 5% significance level, for a joint test of both the haplotype(s) and the haplotype by environment interaction(s). Type I errors are shown in bold.

^dProbability of reaching 5% significance level, for a test of the haplotype by environment interaction(s) only. Type I errors are shown in bold.

^eAverage value of the relative risk ($\exp(\beta_g)$ or $\exp(\beta_{ge})$) over 1000 simulation replicates

f_j Average difference between the log relative risk (β_j or β_{je}) and its true value over 1000 simulation replicates

f_{je} Coverage of the estimated 95% confidence interval for the log relative risk (β_j or β_{je}) over 1000 simulation replicates