

The Semiparametric Case-Only Estimator

Eric J. Tchetgen Tchetgen* and James Robins

Departments of Epidemiology and Biostatistics, Harvard School of Public Health,
677 Huntington Avenue, Boston, Massachusetts 02115, U.S.A.

**email:* etchetge@hsph.harvard.edu

SUMMARY. We propose a semiparametric case-only estimator of multiplicative gene–environment or gene–gene interactions, under the assumption of conditional independence of the two factors given a vector of potential confounding variables. Our estimator yields valid inferences on the interaction function if either but not necessarily both of two unknown baseline functions of the confounders is correctly modeled. Furthermore, when both models are correct, our estimator has the smallest possible asymptotic variance for estimating the interaction parameter in a semiparametric model that assumes that at least one but not necessarily both baseline models are correct.

KEY WORDS: Gene–environment independence; Gene–environment interaction; Generalized odds ratio; Double robustness; Local efficiency.

1. Introduction

The case-only design has recently become a popular approach for making inferences on the statistical interaction on the risk ratio scale between the effects of a genetic factor G and an environmental factor E (or a second genetic factor G') on the risk of a dichotomous disease status Y . Henceforth, we always refer to the second factor as an environmental factor to simplify the exposition. For dichotomous E and G , the crude case-only estimator of the interaction parameter is the empirical marginal odds ratio (i.e., crude odds ratio) between gene and environment among cases, and thus data on unaffected individuals are not required (Piegorsch, Weinberg, and Taylor, 1994). The validity of this estimator relies crucially on the assumption that gene and environment are independent in the population from which cases arose (Albert et al., 2001). Efficiency considerations have also contributed to the appeal of the case-only estimator; when the disease is rare, the case-only estimator is well known to be nearly efficient even when data on unaffected individuals are available as, for example, in a case–control study with controls sampled from subjects who remain unaffected at the end of the study. In contrast, under gene–environment independence, the standard prospective logistic regression estimator of gene–environment interaction can be less efficient. This is expected because it does not make use of the assumed independence (Prentice and Pyke, 1979; Breslow, Robins, and Wellner, 2000).

Although the case-only design allows for consistent estimation of G–E interaction on the risk ratio scale under the assumption of G–E independence, it does not allow for the estimation of the main effects of either G or E on the risk of disease, without which a meaningful interpretation of interactions may be difficult. In spite of this limitation, qualitative prior knowledge concerning main effects of both G and E can provide an appropriate background context for meaningful interpretation of the interaction parameters. For instance, Van-

derweele, Hernandez-Diaz, and Hernan (2009) recently discuss situations where a case-only estimate of a statistical interaction together with prior qualitative knowledge of main effects may be interpreted causally as indicating the presence of a mechanistic interaction between the genetic factor and the environmental exposure.

To guarantee that the gene–environment independence assumption holds, it is often necessary to condition on additional covariates \mathbf{Z} (Piegorsch et al., 1994; Gatto et al., 2004). Furthermore, it is often necessary to condition on covariates \mathbf{W} that confound the association between disease and the gene and/or the environmental factor. As covariates included in \mathbf{Z} may very well overlap with those included in \mathbf{W} , we simplify notation by letting \mathbf{L} denote $(\mathbf{Z}^T, \mathbf{W}^T)^T$ and further assume that conditional on \mathbf{L} , G and E are independent and unmeasured confounding is absent. Furthermore, as in Chatterjee and Carroll (2005), we assume that the genetic factor has finite support $\{g_0, \dots, g_{T-1}\}$ and thus we make use of the vector $\{I(G = g_1), \dots, I(G = g_{T-1})\}$, which we again denote by G to simplify notation, with the vector G of $T - 1$ zeroes corresponding to the level g_0 . The main objective of this article is to provide a semiparametric framework for making inferences under a case-only design on the generalized conditional log-odds ratio function

$$\theta(G, E, g_0, e_0, \mathbf{L}) = \log \left\{ \frac{h(G, E | \mathbf{L}, Y = 1) h(g_0, e_0 | \mathbf{L}, Y = 1)}{h(G, e_0 | \mathbf{L}, Y = 1) h(g_0, E | \mathbf{L}, Y = 1)} \right\} \quad (1)$$

between G and E given \mathbf{L} among cases ($Y = 1$), where $h(G, E | \mathbf{L}, Y = 1)$ is the retrospective density of G and E given \mathbf{L} and $Y = 1$, and (g_0, e_0) is a user-specified point in the sampling space that generally corresponds to an interpretable reference level for G and E such as $(0, 0)$ in the simple case where

both variables are binary. In light of the previous discussion, we formalize our model by assuming that

- (a.1) cases follow a semiparametric log-binomial model for the density of Y given G, E, \mathbf{L} with

$$\begin{aligned} \log(P(Y = 1 | G, E, \mathbf{L}; \psi_0)) \\ = m_G(G, \mathbf{L}) + \gamma(G, E, \mathbf{L}; \psi_0) + m_E(E, \mathbf{L}) + m_{\mathbf{L}}(\mathbf{L}), \end{aligned}$$

where $\gamma(G, E, \mathbf{L}; \psi)$ is a known function and ψ_0 is an unknown p -dimensional parameter. The term $\gamma(G, E, \mathbf{L}; \psi)$ is the generalized interaction function for the effects of E and G on Y within levels of \mathbf{L} , on the exponential scale, and takes the value 0 if either $G = g_0$, $E = e_0$, or $\psi = 0$, so that $\psi_0 = 0$ encodes the null hypothesis of no interaction on the risk ratio scale. The terms $m_E(E, \mathbf{L})$ and $m_G(G, \mathbf{L})$ represent main effects for environment and genetic status, respectively, within levels of \mathbf{L} . We assume that $m_E(e_0, \mathbf{L}) = m_G(g_0, \mathbf{L}) = 0$ and that there exists a function $B(\mathbf{L})$ such that $\max_{G, E} [\max\{|m_G(G, \mathbf{L})|, |m_E(E, \mathbf{L})|, |\gamma(G, E, \mathbf{L}; \psi_0)|\}] \leq B(\mathbf{L})$ for almost all \mathbf{L} . The functions $m_G(G, \mathbf{L})$ and $m_E(E, \mathbf{L})$ are otherwise unrestricted, whereas the function $m_{\mathbf{L}}(\mathbf{L})$ is restricted only by the condition $\max_{G, E} P(Y = 1 | G, E, \mathbf{L}; \psi_0) < 1$ for almost all \mathbf{L} . Note that, with this model, $\exp\{\gamma(G, E, \mathbf{L}; \psi_0)\}$ is the L -specific generalized relative risk interaction function, i.e., multiplicative G–E interaction function (see Vansteelandt et al., 2008, for further detail on this model);

- (a.2) E and G are conditionally independent given \mathbf{L} so that the joint conditional density of gene and environment given \mathbf{L} in the overall population, factorizes as $f_{G, E | \mathbf{L}}(G, E | \mathbf{L}) = f_{G | \mathbf{L}}(G | \mathbf{L}) \times f_{E | \mathbf{L}}(E | \mathbf{L})$, with $f_{G | \mathbf{L}}(G | \mathbf{L})$ and $f_{E | \mathbf{L}}(E | \mathbf{L})$ unrestricted conditional density functions;
- (a.3) the population marginal density $f_{\mathbf{L}}(\mathbf{L})$ of \mathbf{L} is unrestricted; and
- (a.4) the observed data (E, G, \mathbf{L}) are randomly sampled from the density $h_{G, E, \mathbf{L} | Y=1}(E, G, \mathbf{L} | Y = 1)$.

Under these assumptions, we show that

$$\gamma(G, E, \mathbf{L}; \psi_0) = \theta(G, E, g_0, e_0, \mathbf{L}); \quad (2)$$

in other words, the generalized interaction function between E and G within levels of \mathbf{L} coincides with the generalized odds ratio function, relating genetic and environmental factors within levels of \mathbf{L} among cases only.

In Section 2, we show that assumptions (a.1)–(a.4) induce a semiparametric model for the observed data characterized by the sole restriction (2). As our interest is in ψ_0 , we would prefer that the nuisance functions $m_E(E, \mathbf{L})$, $m_G(G, \mathbf{L})$, $f_{G | \mathbf{L}}(G | \mathbf{L})$, $f_{E | \mathbf{L}}(E | \mathbf{L})$, and $f_{\mathbf{L}}(\mathbf{L})$ remain unrestricted. However, because of the curse of dimensionality, this is not possible, when as we assume throughout, \mathbf{L} is highly multivariate and/or includes two or more continuous covariates. Thus, we adopt the practical yet flexible approach of Chen (2007) and Tchetgen Tchetgen, Robins, and Rotnitzky (2010). Specifically, we construct a so-called double robust (DR) estimator of ψ_0 that is regular and asymptotically linear (RAL), and thus consistent and asymptotically normal (CAN), in the semiparametric union model that

assumes that one but not necessarily both of the following statements hold: (i) a parametric model $\rho_G(G, \mathbf{L}; \alpha)$ indexed by α , for the conditional density function $\rho_G(G, \mathbf{L}) = f_{G | E, Y=1, \mathbf{L}}(G | E = e_0, Y = 1, \mathbf{L})$ is correct; or (ii) a parametric model $\rho_E(E, \mathbf{L}; \eta)$ indexed by η , for the conditional density function $\rho_E(E, \mathbf{L}) = f_{E | G, Y=1, \mathbf{L}}(E | G = g_0, Y = 1, \mathbf{L})$ is correct. It is because of this remarkable property that our estimator is called doubly robust, to reflect the fact that it provides the analyst with two separate chances for getting the correct answer. At the intersection submodel where both (i) and (ii) hold, our estimator is locally semiparametric efficient in the union model, in the sense that it is a semiparametric estimator in the union model whose asymptotic variance attains the semiparametric variance bound for the union model at the intersection submodel. We should emphasize that estimation of the functions $m_{\mathbf{L}}(\mathbf{L})$ and $f_{\mathbf{L}}(\mathbf{L})$ is not required by our semiparametric approach. Also, under a rare disease assumption, case-only estimation (in particular, our proposed doubly robust case-only approach) may also be used to evaluate the parameters of a gene–environment generalized interaction function operating on the logit scale; because when the outcome is rare within all strata defined by (E, G, \mathbf{L}) , assumption (a.1) may often serve as a practical approximation to the alternate assumption (a*.1) that cases arise prospectively according to a semiparametric logistic model $\logit\{P(Y = 1 | G, E, \mathbf{L}; \psi_0)\} = m_G(G, \mathbf{L}) + \gamma(G, E, \mathbf{L}; \psi_0) + m_E(E, \mathbf{L}) + m_{\mathbf{L}}(\mathbf{L})$; in which case our results would also hold approximately in this latter model. However, under the logistic model of (a*.1), it is well known that departures from the rare disease assumption can result in a severely biased estimate of interaction parameters (Schmidt and Schaid, 1999). Furthermore, as we argue later in the “Discussion” section, even when by common standards, the disease is considered rare, case-only test statistics and $(1 - \alpha)$ confidence intervals for interaction parameters on the odds ratio scale may substantially deviate from their nominal levels. This is in contrast to the case of gene–environment interactions on the risk ratio scale, where the case-only estimator remains valid whether or not the disease is rare (Schmidt and Schaid, 1999; Yang et al., 1999).

This article is organized as follows: In Section 2, we formally derive equation (2) that we use to construct RAL estimators of ψ_0 in the union model. In Section 3, results of simulations illustrate the finite-sample efficiency and robustness of our new estimators. In Section 4, we describe a simple specification test for detecting which of the models (i) and (ii) is correctly specified under the union model. In Section 5, the method is illustrated with the data from an Israeli study of the interactions between reproductive risk factors and BRCA1/2 in their effects on the risk of ovarian cancer. Some closing remarks are provided in Section 6.

In the following, to simplify notation, we suppose $g_0 = 0$ and $e_0 = 0$ throughout, so that $\gamma(G, 0, \mathbf{L}; \psi) = \gamma(0, E, \mathbf{L}; \psi) = \gamma(G, E, \mathbf{L}; 0) = 0$.

2. Our Estimators and Their Properties

The following theorem motivates our proposed approach.

THEOREM 1: *Assumptions (a.1)–(a.4) imply that the observed data E, G, \mathbf{L} follow a semiparametric model with sole restriction given by equation (2).*

Proof. By assumption (a.3) the conditional density $h_{\mathbf{L}|Y=1}(\mathbf{L}|Y=1)$ is unrestricted, thus it is sufficient to show that under assumptions (a.1), (a.2), and (a.4), the conditional density $h_{E,G|\mathbf{L},Y=1}(E,G|Y=1,\mathbf{L};\psi_0)$ of G and E given \mathbf{L} and $Y=1$ can be written as

$$\begin{aligned} & h_{G,E|\mathbf{L},Y=1}(G,E|\mathbf{L},Y=1;\psi_0) \\ &= \frac{\exp\{m_G(G,\mathbf{L}) + \gamma(G,E,\mathbf{L};\psi_0) + m_E(E,\mathbf{L})\} \times f_{G|\mathbf{L}}(G|\mathbf{L}) \times f_{E|\mathbf{L}}(E|\mathbf{L})}{\int \exp\{m_G(g,\mathbf{L}) + \gamma(g,e,\mathbf{L};\psi_0) + m_E(e,\mathbf{L})\} \times f_{G|\mathbf{L}}(g|\mathbf{L}) \times f_{E|\mathbf{L}}(e|\mathbf{L}) d\mu(e,g)} \\ &= \frac{\exp\{\gamma(G,E,\mathbf{L};\psi_0)\} \times f_{G|E,Y=1,\mathbf{L}}(G|E=0,Y=1,\mathbf{L}) \times f_{E|G,Y=1,\mathbf{L}}(E|G=0,Y=1,\mathbf{L})}{\int \exp\{\gamma(g,e,\mathbf{L};\psi_0)\} \times f_{G|E,Y=1,\mathbf{L}}(g|E=0,Y=1,\mathbf{L}) \times f_{E|G,Y=1,\mathbf{L}}(e|G=0,Y=1,\mathbf{L}) d\mu(e,g)}, \end{aligned}$$

where the first equality follows from the definition of the conditional joint density

$$\begin{aligned} & h_{E,G,\mathbf{L}|Y=1}(E,G,\mathbf{L}|Y=1;\psi_0) \\ &= \frac{\exp\{m_G(G,\mathbf{L}) + \gamma(G,E,\mathbf{L};\psi_0) + m_E(E,\mathbf{L}) + m_{\mathbf{L}}(\mathbf{L})\} \times f_{G|\mathbf{L}}(G|\mathbf{L}) \times f_{E|\mathbf{L}}(E|\mathbf{L}) \times f_{\mathbf{L}}(\mathbf{L})}{\int \exp\{m_G(g,\mathbf{L}) + \gamma(g,e,\mathbf{L};\psi_0) + m_E(e,\mathbf{L}) + m_{\mathbf{L}}(\mathbf{L})\} \times f_{G|\mathbf{L}}(g|\mathbf{L}) \times f_{E|\mathbf{L}}(e|\mathbf{L}) \times f_{\mathbf{L}}(\mathbf{L}) d\mu(g,e,\mathbf{L})} \end{aligned}$$

and the second equality follows from assumption (a.2) and the cancellation of all functions of only \mathbf{L} appearing in both the denominator and the numerator. Therefore, $h_{G,E|\mathbf{L},Y=1}(G,E|\mathbf{L},Y=1;\psi_0) = \frac{\exp\{\gamma(G,E,\mathbf{L};\psi_0)\} \rho_G(G,\mathbf{L}) \rho_E(E,\mathbf{L})}{\int \exp\{\gamma(g,e,\mathbf{L};\psi_0)\} \rho_G(g,\mathbf{L}) \rho_E(e,\mathbf{L}) d\mu(e,g)}$ is a conditional density of $G, E|\mathbf{L}, Y=1$ indexed by the conditional odds ratio function $\theta(G, E, g_0, e_0, \mathbf{L}) = \gamma(G, E, \mathbf{L}; \psi_0)$. Because $\{m_G(G, \mathbf{L}), f_{G|\mathbf{L}}(G|\mathbf{L})\}$ are unrestricted, and $\{m_E(E, \mathbf{L}), f_{E|\mathbf{L}}(E|\mathbf{L})\}$ are unrestricted by assumptions (a.1) and (a.2), $\rho_G(G, \mathbf{L}) = f_{G|e_0, Y=1, \mathbf{L}}(G|E=e_0, Y=1, \mathbf{L})$ and $\rho_E(E, \mathbf{L}) = f_{E|g_0, Y=1, \mathbf{L}}(E|G=g_0, Y=1, \mathbf{L})$ are unrestricted baseline conditional densities satisfying $\int \exp\{\gamma(g, e, \mathbf{L}; \psi)\} \rho_G(g, \mathbf{L}) \rho_E(e, \mathbf{L}) d\mu(g, e) < \infty$ for almost all \mathbf{L} .

Consider working models $\rho_G(G, \mathbf{L}; \alpha)$ and $\rho_E(E, \mathbf{L}; \eta)$ for density functions $\rho_G(G, \mathbf{L})$ and $\rho_E(E, \mathbf{L})$, respectively, then an individual's conditional likelihood (given \mathbf{L} and $Y=1$) contribution for this model is $h_{G,E|\mathbf{L},Y=1}(G,E|\mathbf{L},Y=1;\psi, \alpha, \eta) =$

$$\frac{\exp\{\gamma(G, E, \mathbf{L}; \psi)\} \rho_G(G, \mathbf{L}; \alpha) \rho_E(E, \mathbf{L}; \eta)}{\int \exp\{\gamma(g, e, \mathbf{L}; \psi)\} \rho_G(g, \mathbf{L}; \alpha) \rho_E(e, \mathbf{L}; \eta) d\mu(g, e)}, \quad (3)$$

where α and η are variation-independent parameters. Now, we could in principle obtain likelihood-based inferences on ψ_0 by taking one of the following three approaches: (1) we could obtain the maximum likelihood estimator (MLE) $(\hat{\psi}_{\text{mle}}, \hat{\eta}_{\text{mle}}, \hat{\alpha}_{\text{mle}})$ in model (3), which is consistent provided the model is correct. However, because we can never know that the intersection submodel (where both $\rho_G(\cdot, \cdot; \alpha)$ and $\rho_E(\cdot, \cdot; \eta)$ hold) applies, this approach may be too restrictive; and thus we may prefer the alternate approach (2) that obtains the conditional MLE (CMLE) $(\hat{\psi}_{\text{cmle},1}, \hat{\alpha}_{\text{cmle},1})$ based on the conditional likelihood of the density model for G given E, \mathbf{L} , and $Y=1$, given by $h_{G|E,\mathbf{L},Y=1}(G|E,\mathbf{L},Y=1;\psi, \alpha) = \frac{\exp\{\gamma(G,E,\mathbf{L};\psi)\} \rho_G(G,\mathbf{L};\alpha)}{\int \exp\{\gamma(g,e,\mathbf{L};\psi)\} \rho_G(g,\mathbf{L};\alpha) d\mu(g,e)}$, which does not depend on $\rho_E(E, \mathbf{L})$ unrestricted, but requires that $\rho_G(G, \mathbf{L}; \alpha)$ is correct; or alternately, we could obtain (3) the

conditional mle $(\hat{\psi}_{\text{cmle},2}, \hat{\eta}_{\text{cmle},2})$ by maximizing the conditional likelihood for the density model of E given G, \mathbf{L} , and $Y=1$, given by $h_{E|G,\mathbf{L},Y=1}(E|G,\mathbf{L},Y=1;\eta, \psi) = \frac{\exp\{\gamma(G,E,\mathbf{L};\psi)\} \rho_E(E,\mathbf{L};\eta)}{\int \exp\{\gamma(g,e,\mathbf{L};\psi)\} \rho_E(e,\mathbf{L};\eta) d\mu(g,e)}$, which leaves $\rho_G(G, \mathbf{L})$

unrestricted but requires that $\rho_E(E, \mathbf{L}; \eta)$ is correct. Although more robust than the first approach, the analyst will

generally not know which, if any, of the two models (2) or (3) is correct. For this reason, we construct a CAN estimator that is guaranteed to be consistent for the interaction function if at least one, but not necessarily both of the baseline models $\rho_E(E, \mathbf{L}; \eta)$ or $\rho_G(G, \mathbf{L}; \alpha)$ is correct.

To proceed, let $\epsilon(\psi, \alpha) = G - E(G|E, \mathbf{L}, Y=1; \psi, \alpha)$ where $E(\cdot|\cdot, \mathbf{L}, Y=1; \psi, \alpha, \eta)$ denotes conditional expectations with respect to $h_{G,E|\mathbf{L},Y=1}(G,E|\mathbf{L},Y=1;\psi, \eta, \alpha)$. Define $\Psi(E, \mathbf{L}; \psi, \alpha) = E\{\epsilon(\psi, \alpha)^{\otimes 2} | E, \mathbf{L}; \psi, \alpha\}$ and for a user-supplied $p \times (T-1)$ -dimensional function $\mathbf{K} = \mathbf{k}(E, \mathbf{L})$, let $\tilde{U}(\mathbf{k}; \psi, \alpha, \eta) = \{\mathbf{K} - \tilde{E}(\mathbf{K}|\mathbf{L}; \psi, \alpha, \eta)\} \times \epsilon(\psi, \alpha)$ be a function of p -dimensions, where $\tilde{E}\{\mathbf{K}|\mathbf{L}; \psi, \alpha, \eta\} = E\{\mathbf{k}(E, \mathbf{L}) \times \Psi(E, \mathbf{L}; \psi, \alpha) | \mathbf{L}; \psi, \alpha, \eta\} \times E\{\Psi(E, \mathbf{L}; \psi, \alpha) | \mathbf{L}; \psi, \alpha, \eta\}^{-1}$. By theorem 2 of Tchetgen Tchetgen et al. (2010), the estimator $\hat{\psi}(\mathbf{k})$ is RAL in the union model, where $\hat{\psi}(\mathbf{k})$ is the solution to

$$\sum_{i=1}^n \tilde{U}_i\{\mathbf{k}; \psi, \hat{\alpha}(\psi), \hat{\eta}(\psi)\} = 0,$$

and

$$\hat{\alpha}(\psi) = \arg \max_{\alpha} \sum_{i=1}^n \log\{h_{G|E,\mathbf{L},Y=1}(G_i|E_i, \mathbf{L}_i, Y_i=1; \psi, \alpha)\},$$

is the profile MLE of α at a fixed ψ ,

$$\hat{\eta}(\psi) = \arg \max_{\eta} \sum_{i=1}^n \log\{h_{E|G,\mathbf{L},Y=1}(E_i|G_i, \mathbf{L}_i, Y_i=1; \psi, \eta)\}$$

is the profile MLE of η at a fixed ψ . Furthermore, Tchetgen Tchetgen et al. (2010) also show that the estimator $\hat{\psi}_{\text{eff}} = \hat{\psi}(\hat{\mathbf{k}}_{\text{eff}})$ is locally semiparametric efficient in the union model at the intersection submodel, where $\hat{\mathbf{k}}_{\text{eff}}(E, \mathbf{L}) = \frac{\partial \{\gamma^T(E, \mathbf{L}; \psi)\}}{\partial \psi} \Big|_{\psi=\hat{\psi}}$, with $\gamma(E, \mathbf{L}; \psi)$ the $(T-1) \times 1$ vector with j th component equal to $\gamma(g_j, E, \mathbf{L}, \psi)$, $1 \leq j \leq T-1$; and $\hat{\psi}$ any preliminary estimator of ψ_0 that is consistent at the intersection submodel (this could be the MLE or either CMLES). Furthermore, these authors prove that under the union model

Table 1
Simulation results for $\psi_0 = 0$

		$n = 200$				$n = 1000$			
		$\hat{\psi}_{\text{cmle},1}$	$\hat{\psi}_{\text{cmle},2}$	$\hat{\psi}_{\text{mle}}$	$\hat{\psi}_{\text{eff}}$	$\hat{\psi}_{\text{cmle},1}$	$\hat{\psi}_{\text{cmle},2}$	$\hat{\psi}_{\text{mle}}$	$\hat{\psi}_{\text{eff}}$
$\alpha_{\text{true}}, \eta_{\text{true}}$	Bias	0.004	0.008	0.004	0.007	0.007	0.005	0.005	0.007
	Variance	0.123	0.113	0.125	0.113	0.023	0.022	0.021	0.023
	Coverage ^a	0.956	0.958	0.955	0.955	0.944	0.956	0.953	0.951
$\alpha_{\text{true}}, \eta_{\text{false}}$	Bias	0.004	-0.240	-0.239	0.009	0.007	-0.245	0.245	0.007
	Variance	0.123	0.114	0.114	0.134	0.023	0.020	0.020	0.023
	Coverage	0.956	0.891	0.895	0.945	0.944	0.584	0.587	0.949
$\alpha_{\text{false}}, \eta_{\text{true}}$	Bias	-0.245	0.008	-0.227	0.007	-0.245	0.005	-0.233	0.005
	Variance	0.108	0.113	0.098	0.125	0.020	0.022	0.019	0.023
	Coverage	0.875	0.958	0.882	0.952	0.593	0.956	0.618	0.953
$\alpha_{\text{false}}, \eta_{\text{false}}$	Bias	-0.245	-0.240	-0.229	-0.229	-0.245	-0.245	-0.252	-0.255
	Variance	0.108	0.114	0.108	0.107	0.020	0.020	0.021	0.020
	Coverage	0.875	0.891	0.880	0.894	0.593	0.584	0.563	0.550

^aCoverage of Wald-type confidence intervals using the second derivative of the score equations to estimate the variance of CMLES and of the MLE, and $\hat{\Sigma}$ to estimate the variance of $\hat{\psi}_{\text{eff}}$.

and regardless of whether or not the intersection submodel holds, we have $\sqrt{n}\{\hat{\psi}(\hat{\mathbf{k}}_{\text{eff}}) - \psi_0\}$ is asymptotically normal, with variance consistently estimated by

$$\hat{\Sigma} = n^{-1} \sum_{i=1}^n \times \left(\left[n^{-1} \sum_{j=1}^n \frac{\partial}{\partial \psi} M_j \{(\psi, \hat{\eta}(\hat{\psi}_{\text{eff}}), \hat{\alpha}(\hat{\psi}_{\text{eff}}); \hat{\mathbf{k}}_{\text{eff}}) |_{\psi=\hat{\psi}_{\text{eff}}}\}^{-1} \right]^{\otimes 2} \right) \times M_i \{ \hat{\psi}_{\text{eff}}, \hat{\eta}(\hat{\psi}_{\text{eff}}), \hat{\alpha}(\hat{\psi}_{\text{eff}}); \hat{\mathbf{k}}_{\text{eff}} \}$$

where

$$\begin{aligned} M(\psi, \eta, \alpha; \mathbf{k}) &= \tilde{U}(\mathbf{k}; \psi, \eta, \alpha) \\ &- E \left\{ \frac{\partial}{\partial \eta} \tilde{U}(\mathbf{k}; \psi, \eta, \alpha) \right\} E^{-1} \left\{ \frac{\partial}{\partial \eta} C(\psi, \eta) \right\} C(\psi, \eta) \\ &- E \left\{ \frac{\partial}{\partial \alpha} \tilde{U}(\mathbf{k}; \psi, \eta, \alpha) \right\} E^{-1} \left\{ \frac{\partial}{\partial \alpha} B(\psi, \alpha) \right\} B(\psi, \alpha), \end{aligned} \quad (4)$$

$C(\psi, \eta) = \frac{\partial}{\partial \eta} \log\{h_E |_{G, \mathbf{L}, Y=1}(E_i | G_i, \mathbf{L}_i, Y_i = 1; \psi, \eta)\}$ and $B(\psi, \alpha) = \frac{\partial}{\partial \alpha} \log\{h_G |_{E, \mathbf{L}, Y=1}(G_i | E_i, \mathbf{L}_i, Y_i = 1; \psi, \alpha)\}$ are the scores for η and α , respectively; which can be used to obtain Wald-type confidence intervals for components of ψ_0 . Alternatively, inferences on ψ_0 can be based on the nonparametric bootstrap.

An alternate approach that is approximately DR locally efficient was proposed by Chen (2007), however, his method requires the use of a computationally intensive alternating conditional expectations (ACE) algorithm to approximate the efficient score. Because in our setting, the efficient score can be written in a closed form, we do not need Chen's approximate approach.

3. A Simulation Study

We perform a simulation study where we study the finite-sample performance of $\hat{\psi}_{\text{cmle},1}$, $\hat{\psi}_{\text{cmle},2}$, $\hat{\psi}_{\text{mle}}$, and $\hat{\psi}_{\text{eff}}$

in 1000 data samples (of sample size $n = 200, 1000$) consisting of variables $\mathbf{L} = \{(L_1, L_2), G, E\}$ generated by repeatedly sampling L_1 from a Bernoulli (1/2) density, L_2 from a normal(0, 0.75²); and dichotomous pairs (G, E) from the conditional probability mass function (3) with logit $\{\rho_G(1, \mathbf{L}; \alpha)\} = \mathbf{L}'_G \alpha = (1, L_1, L_2^2, L_1 L_2^3) \alpha$, logit $\{\rho_E(1, \mathbf{L}; \eta)\} = \mathbf{L}'_E \eta = (1, L_1, L_2^2, L_1 L_2^3) \eta$, where $\alpha' = (0.1, -1, 1, -1.1)$ and $\eta' = (0.5, 0.5, -1.25, -1)$. Two values of $\psi = 0, 1$ are considered. The first three rows of results labeled $(\alpha_{\text{true}}, \eta_{\text{true}})$ in Tables 1 and 2 correspond to an analysis that correctly specifies both baseline functions. The next three rows labeled $(\alpha_{\text{true}}, \eta_{\text{false}})$ correspond to an analysis that misspecifies $\rho_E(E, \mathbf{L}; \eta)$ by using $\mathbf{L}' = (1, L_1, L_2, L_1 L_2)$ in the regression model instead of \mathbf{L}'_E , the following three rows labeled $(\alpha_{\text{false}}, \eta_{\text{true}})$ report an analysis that misspecifies $\rho_G(G, \mathbf{L}; \alpha)$ by using \mathbf{L} in the regression model instead of \mathbf{L}'_G , and the final three rows labeled $(\alpha_{\text{false}}, \eta_{\text{false}})$ correspond to an analysis where both baseline functions are incorrect. Under the correct model specification for the joint conditional density of E and G give \mathbf{L} and $Y = 1$ (corresponding to the first three rows of results in Tables 1 and 2), all four estimators have finite-sample bias of comparable but negligible size. $\hat{\psi}_{\text{eff}}$ is surprisingly efficient relative to the other estimators, including the MLE that suggests that in this simulation setting, little efficiency loss is incurred by our estimator in exchange for potential gain in robustness. Also under $(\alpha_{\text{true}}, \eta_{\text{true}})$, close to nominal type 1 error rates are achieved by all estimators in the null case $\psi_0 = 0$, and nearly nominal coverage rates are obtained in the case of $\psi_0 = 1$. Next, we find that as predicted by theory, $\hat{\psi}_{\text{cmle},1}$ and $\hat{\psi}_{\text{cmle},2}$ are both severely biased under misspecification of their respective baseline function as nicely illustrated in the fourth row of the first table for $n = 1000$, where the magnitude of the bias of $\hat{\psi}_{\text{cmle},2}$ is over 20-fold that of $\hat{\psi}_{\text{eff}}$, which remains negligible. Under model misspecification, the MLE is both biased and yields CIs with incorrect coverage as illustrated by rows corresponding to settings $(\alpha_{\text{true}}, \eta_{\text{false}})$, $(\alpha_{\text{false}}, \eta_{\text{true}})$, and $(\alpha_{\text{false}}, \eta_{\text{false}})$ in both tables. Interestingly, the bias of the MLE is always close to that of the misspecified CMLE in rows corresponding to $(\alpha_{\text{true}}, \eta_{\text{false}})$,

Table 2
Simulation results for $\psi_0 = 1$

		$n = 200$				$n = 1000$			
		$\hat{\psi}_{\text{cmle},1}$	$\hat{\psi}_{\text{cmle},2}$	$\hat{\psi}_{\text{mle}}$	$\hat{\psi}_{\text{eff}}$	$\hat{\psi}_{\text{cmle},1}$	$\hat{\psi}_{\text{cmle},2}$	$\hat{\psi}_{\text{mle}}$	$\hat{\psi}_{\text{eff}}$
$\alpha_{\text{true}}, \eta_{\text{true}}$	Bias	0.025	0.025	0.023	0.026	0.008	0.009	0.008	0.009
	Variance	0.140	0.124	0.120	0.143	0.030	0.027	0.027	0.030
	Coverage ^a	0.971	0.971	0.972	0.963	0.951	0.938	0.935	0.948
$\alpha_{\text{true}}, \eta_{\text{false}}$	Bias	0.025	-0.234	-0.236	0.028	0.008	-0.253	-0.254	0.002
	Variance	0.140	0.136	0.135	0.161	0.030	0.026	0.025	0.030
	Coverage	0.971	0.892	0.893	0.942	0.951	0.627	0.620	0.945
$\alpha_{\text{false}}, \eta_{\text{true}}$	Bias	-0.245	0.025	-0.230	0.037	-0.268	0.009	-0.247	0.006
	Variance	0.125	0.124	0.112	0.151	0.023	0.027	0.021	0.029
	Coverage	0.897	0.971	0.906	0.958	0.593	0.938	0.613	0.942
$\alpha_{\text{false}}, \eta_{\text{false}}$	Bias	-0.245	-0.234	-0.251	-0.264	-0.268	-0.253	-0.251	-0.252
	Variance	0.125	0.136	0.126	0.126	0.023	0.026	0.024	0.024
	Coverage	0.897	0.892	0.879	0.916	0.593	0.627	0.580	0.617

^aCoverage of Wald-type confidence intervals using the second derivative of the score equations to estimate the variance of CMLEs and of the MLE, and $\hat{\Sigma}$ to estimate the variance of $\hat{\psi}_{\text{eff}}$.

and $(\alpha_{\text{false}}, \eta_{\text{true}})$, with the bias, variance, and coverage of $\hat{\psi}_{\text{cmle},2}$ under $(\alpha_{\text{true}}, \eta_{\text{false}})$ always close to those of $\hat{\psi}_{\text{cmle},1}$ under $(\alpha_{\text{false}}, \eta_{\text{true}})$ at both sample sizes and values of ψ_0 . Also in the union model, we note that $\hat{\psi}_{\text{eff}}$ is surprisingly nearly as efficient as the correctly specified CMLE with corresponding close to nominal coverage rates, as illustrated by the eighth and ninth rows of Table 2. Increasing sample size has the intended effect of reducing finite-sample bias and variance of consistent estimators ($\hat{\psi}_{\text{cmle},1}, \hat{\psi}_{\text{eff}}$) under $(\alpha_{\text{true}}, \eta_{\text{false}})$, and ($\hat{\psi}_{\text{cmle},2}, \hat{\psi}_{\text{eff}}$) under $(\alpha_{\text{false}}, \eta_{\text{true}})$ for both null and alternative values of ψ_0 . As expected, the bias of misspecified models is little affected by increasing sample size. Finally, no estimator gives valid results when both baseline functions are incorrect.

4. A Specification Test Under the Union Model

Next we describe a simple specification test to detect which of the two baseline models $\rho_E(E, \mathbf{L}; \eta)$ and $\rho_G(G, \mathbf{L}; \alpha)$ is correct under the union model. The approach is based on the following observation: if $\rho_G(G, \mathbf{L}; \alpha)$ is correct, the corresponding CMLE of ψ should be close (within sampling variability) to the doubly robust estimator whether or not $\rho_E(E, \mathbf{L}; \eta)$ holds, as both estimators should be consistent for the truth. In contrast, if $\rho_G(G, \mathbf{L}; \alpha)$ is incorrect, the limiting value of the corresponding CMLE should differ (beyond sampling variability) from that of the doubly robust estimator, because the former is biased, whereas the latter is consistent under the union model.

Thus, our test statistic is the suitably standardized (by its covariance) difference between the estimates of a subset of the CMLE vector $\hat{\psi}_{\text{cmle},1}$ obtained under model $\rho_G(G, \mathbf{L}; \alpha)$ to those obtained from the proposed doubly robust method. This strategy essentially uses an idea due to Hausman (1978) and further developed by Newey (1985). In the simple case where ψ_0 is scalar, the statistic takes on the simple form $T = n(\hat{\psi}_{\text{dr}} - \hat{\psi}_{\text{cmle},1})^2 / (\hat{\sigma}_{\text{dr}}^2 - \hat{\sigma}_{\text{cmle},1}^2)$, where $\hat{\sigma}_{\text{dr}}^2$ and $\hat{\sigma}_{\hat{\psi},1}^2$ are consistent estimates of the variance of $n^{1/2}(\hat{\psi}_{\text{dr}} - \psi_0)$ and $n^{1/2}(\hat{\psi}_{\text{cmle},1} - \psi_0)$, respectively, and $\hat{\sigma}_{\text{dr}}^2 - \hat{\sigma}_{\text{cmle},1}^2$ is a consistent estimate of the variance of $n^{1/2}(\hat{\psi}_{\text{dr}} - \hat{\psi}_{\text{cmle},1})$ under the null hypothesis that

both estimators are consistent for ψ_0 . This estimator of the variance is of the simple form of a difference by virtue of $\hat{\psi}_{\text{cmle},1}$ being an efficient estimator under the null semiparametric model that specifies $\{\gamma(G, E, \mathbf{L}, \psi), \rho_G(G, \mathbf{L}; \alpha)\}$ (Hausman, 1978; Newey, 1985). Thus, under the null hypothesis, T has a χ_1^2 asymptotic distribution, and under the alternative it follows a noncentral χ_1^2 asymptotic distribution, with noncentrality parameter solely determined by the direction of the asymptotic bias of $\hat{\psi}_{\text{cmle},1}$ (Newey, 1985). To guarantee that T is always positive, an analytical estimate of the variance of the numerator based on influence function arguments may be used, or alternately, the nonparametric bootstrap estimator of the variance of the numerator may replace the difference estimator in the denominator of the test statistic. A similar approach yields a specification test for $\rho_E(E, \mathbf{L}; \eta)$ under the union model. Outside of the union model, both CMLEs and the doubly robust estimator are converging to distinct values all of which are biased. Our specification test may still be used. However, the finite-sample power of the test may be low under those rare alternatives where the two estimates being compared have similar bias, and thus are both wrong, and yet close to each other so that the proposed test statistic is unlikely to reject. The multivariate version of the test statistic is easily deduced from the previous description.

5. A Data Example

In this section, we illustrate the use of our methodology in an analysis of data from a population-based case-control study based on all ovarian cancer patients identified in Israel between 1 March 1994 and 30 June 1999 (Modan et al., 2001). Two controls per case were selected from the central population registry matching on age within 2 years, area of birth and place, and length of residence. Blood samples were collected on both cases and controls and used to test for the presence of mutation in two major breast and ovarian cancer susceptibility genes BRCA1 and BRCA2. Additional data were collected on reproductive and gynecological history such as parity, number of years of oral contraceptive use, and gynecological surgery. The main objective of the study was

to examine the interplay of the BRCA1/2 genes and known reproductive/gynecological risk factors of ovarian cancer. To test for interactions between reproductive risk factors and BRCA1/2 in their effects on the risk of ovarian cancer, the authors performed the unadjusted case-only analysis of interaction of Piegorsch et al. (1994) under an assumption that genetic variants and environment factor are unconditionally independent in the population. Chatterjee and Carroll (2005) reanalyzed these data using a fully parametric logistic regression model for disease given the gene, environment, and confounding factors, under the additional conditional independence assumption of gene and environment given a subset of measured covariates. However, their results are not directly comparable to ours for two reasons; first, because BRCA1/2 have high penetrance for the risk of ovarian cancer, the disease may not be rare within levels of G, E , and \mathbf{L} ; as a result, the interaction parameter on the logistic scale may differ from that of the risk ratio scale (in other words, assumption (a.1*) stated in the Introduction may not hold). Second, even were the disease rare within all strata of G, E , and \mathbf{L} , they assume a fully parametric model for the disease outcome regression model that may result in biased estimates of interactions if their specified working models for the functions $m_G(G, \mathbf{L}), m_E(E, \mathbf{L}), m_L(\mathbf{L})$ are incorrect. As a consequence, our estimated interaction parameter is not directly comparable to that obtained by Chatterjee and Carroll (2005). In our reanalysis, we illustrate the case-only method developed in this article by ignoring the data on controls. Specifically, we use 832 cases who did not have bilateral oophorectomy were interviewed for risk factor information and successfully tested for BRCA1/2 mutations. Our aim is to estimate the interaction between the dichotomous variable representing a person's BRCA1/2 mutation status and her use of oral contraceptives and parity. We dichotomize oral contraceptive use as use for over 6 years versus use for 6 years or less, whereas parity is dichotomized as having less than two children versus having two or more children. Thus, we estimate (ψ_1, ψ_2) in the interaction model

$$\begin{aligned} \gamma(G, E, \mathbf{L}; \psi) &= \psi_1 \times I(\text{BRCA1/2} = \text{yes}) \times I(\text{OC use} > 6 \text{ years}) \\ &+ \psi_2 \times I(\text{BRCA1/2} = \text{yes}) \times I(\text{no. of children} > 1). \end{aligned}$$

We assume conditional independence of gene and environmental factors given \mathbf{L} , consisting of age (categorical defined by decades), ethnic background (Ashkenazi or non-Ashkenazi), the presence of personal history of breast cancer, a history of gynecological surgery, and family history of breast or ovarian cancer (no cancer versus one breast cancer in the family versus one ovarian cancer or two or more breast cancer cases in the family). We specified a logistic model $\text{logit}[\rho_G\{1, \mathbf{L}; \boldsymbol{\alpha} = (\alpha_0, \boldsymbol{\alpha}_1^T)^T\}] = \alpha_0 + \mathbf{L}'\boldsymbol{\alpha}_1$ for BRCA1/2 and a polytomous logistic model $\log\{\rho_E(e_k, \mathbf{L}; \boldsymbol{\eta}) / \rho_E(e_0, \mathbf{L}; \boldsymbol{\eta})\} = \eta_{0k} + \mathbf{L}'\boldsymbol{\eta}_{1,k}, k = 1, 2, 3$, where $\boldsymbol{\eta}_k = (\eta_{0k}, \boldsymbol{\eta}_k^T)^T$, $\boldsymbol{\eta} = (\boldsymbol{\eta}_1^T, \boldsymbol{\eta}_2^T, \boldsymbol{\eta}_3^T)^T$, and $e_0 = I(\text{OC use} \leq 6 \text{ years}) \times I(\text{no. of children} \leq 1)$, $e_1 = I(\text{OC use} > 6 \text{ years}) \times I(\text{no. of children} \leq 1)$, $e_2 = I(\text{OC use} \leq 6 \text{ years}) \times I(\text{no. of children} > 1)$, $e_3 = I(\text{OC use} > 6 \text{ years}) \times I(\text{no. of children} > 1)$. We then used these models to obtain four estimates $(\hat{\psi}_{1,\text{cmle},1}, \hat{\psi}_{2,\text{cmle},1}) = (0.99, 0.44)$ (bootstrap SE = 0.43, 0.21), $(\hat{\psi}_{1,\text{cmle},2}, \hat{\psi}_{2,\text{cmle},2}) =$

$(1.08, 0.44)$ (bootstrap SE = 0.42, 0.21), $(\hat{\psi}_{1,\text{mle}}, \hat{\psi}_{2,\text{mle}}) = (1.05, 0.43)$ (bootstrap SE = 0.40, 0.19), and $(\hat{\psi}_{1,\text{eff}}, \hat{\psi}_{2,\text{eff}}) = (1.04, 0.46)$ (bootstrap SE = 0.42, 0.21) of the interaction parameters. All estimates indicate strong positive interactions between BRCA1/2 and oral contraceptive use and parity. These results suggest that after adjusting for background variables, the well-known protective effects of oral contraceptive use, and parity on ovarian cancer risk among BRCA1/2 noncarriers (Modan et al., 2001) may no longer apply to BRCA1/2 carriers. All point estimates are close to each other, providing convincing evidence that our models for both baseline functions are a reasonable fit to the data (in fact, the χ_2^2 specification test statistic using bootstrap covariance estimates yields p -values > 0.99 in comparing each CMLE to the doubly robust estimate). Interestingly, all estimators have very similar uncertainty, suggesting that as we also found in our simulation study, we do not necessarily give up much efficiency for the important gain in robustness afforded by the doubly robust approach.

Finally, despite not being directly comparable to our results, we apply the standard logistic regression approach using data both on cases and controls to estimate the interaction between the genetic variant and OC use/parity indicators in a regression model that also included main effects for BRCA1/2 indicator, OC use indicator, parity indicator, and \mathbf{L} . We obtained the following imprecise estimates $(\hat{\psi}_{1,\text{logistic}}, \hat{\psi}_{2,\text{logistic}}) = (0.32, 0.005)$ (SE = 1.12, 0.80), which, in light of the improved accuracy of our proposed locally efficient approach, further illustrates the inefficiency of logistic regression for case-control data under gene-environment independence. The first three of our estimators can be obtained from standard software. The doubly robust locally efficient estimator is implemented in **IML, SAS 9.1.3**, and can be downloaded from the first author's website. The data can freely be obtained from Nilanjan Chatterjee's website at the Biostatistics Branch of the National Cancer Institute.

6. Discussion

We have developed a doubly robust locally efficient estimator of interactions within a semiparametric case-only framework. The proposed methodology is a flexible and efficient extension of the original crude and adjusted case-only estimators. We recommend its use particularly in settings where inferences on gene-environment or gene-gene interactions are sought, genetic information is only collected in cases, the factors of interest are known to be conditionally independent, and, as in most observational studies, it is necessary to further adjust for high-dimensional confounders. However, if either the conditional independence assumption does not hold or it does hold but the disease is not rare and data on the unaffected are available, the approach of Vansteelandt et al. (2008) should be preferred to that presented herein. This is because the case-only estimator is no longer valid in the first setting, whereas in the second situation it is still valid, but does not make use of all of the available information and is therefore not efficient.

As mentioned in the Introduction, when the target of inference is a parameter ψ_0 for a gene-environment interaction operating on the odds ratio scale, as in the logistic regression given in (a*.1) and assumption (a.2) is satisfied, seriously flawed inferences may still result from a case-only

approach, even if by common standards, the disease is considered to be rare within levels of (G, E, L) in the population. To illustrate this point, consider the following rare disease asymptotic analysis. Suppose for simplicity that there are no covariates \mathbf{L} and (G, E) are both dichotomous, so that at sample size n , the observed data are generated under the model $\logit\{P(Y = 1 | G, E; \psi_0, \beta_0, p_n)\} = \log\{p_n / (1 - p_n)\} + \beta_0(G + E)$ such that $p_n \ll 1$ is a sequence of positive numbers, which converges to zero as n goes to infinity, $\psi_0 = 0$ and $m_G(\cdot) = m_E(\cdot) = \beta_0 \times \cdot$. A Taylor series expansion around the limit point $\lim_{n \rightarrow \infty} p_n = 0$ shows that the case-only estimand, i.e., the log-odds ratio relating G to E in cases only ($Y = 1$), is to first-order equal to $r(\beta_0)p_n$ where $r(\beta_0) = 2\{\exp(2\beta_0) - 1\} - \{\exp(\beta_0) - 1\}$. This observation implies that because the case-only estimator $\hat{\psi}$ is asymptotically unbiased for the case-only estimand, the case-only estimator $\hat{\psi}$, viewed as an estimator of the log prospective odds ratio parameter $\psi_0 = 0$, will have asymptotic bias equal to the large sample limit of $n^{1/2}r(\beta_0)p_n$. This in turn leads to the conclusion that whenever $\beta_0 \neq 0$ and p_n converges to zero at a rate no faster than $n^{-1/2}$, a case-only hypothesis test of the null that $\psi_0 = 0$, via say, either checking whether a case-only Wald-type $(1 - \alpha)$ -confidence interval contains zero, or whether the value of a Wald-type case-only test statistic rejects under the null, will in general have incorrect coverage and type I error, respectively. The practical implication of this asymptotic analysis is that the rare disease assumption may not be useful in settings where disease prevalence although small, is roughly of the order of magnitude $n^{-1/2}$ of the standard error of the case-only estimates, or of a larger order of magnitude (say $n^{-1/4} > n^{-1/2}$).

In this article, we have assumed that \mathbf{L} includes all factors that one needs to adjust for in the gene-environment conditional independence model and in the disease-risk model. However, as pointed out by a referee, the case-only design adopted in this article may not always be appropriate when there are variables \mathbf{V} listed in \mathbf{L} such that G and E are independent given \mathbf{L} , but G and E are not independent given $\mathbf{L} \setminus \mathbf{V}$; and the disease risk model of primary interest do not condition on \mathbf{V} (see Chatterjee and Chen, 2007, for an example). The current methods would still apply in such a situation if \mathbf{V} were also known not to be associated with disease given $\mathbf{L} \setminus \mathbf{V}$, which implies but is not implied by the following assumption: $\gamma(G, E, \mathbf{L}; \psi_0)$ does not depend on V . Thus, the G-E conditional interaction parameter ψ_0 would have the desired marginal interpretation, whereas the baseline densities $\rho_E(E, \mathbf{L})$ and $\rho_G(G, \mathbf{L})$ would remain functions of the entire vector \mathbf{L} .

ACKNOWLEDGEMENTS

The authors are grateful to the editor, the associate editor, and two referees for their constructive and insightful comments.

REFERENCES

- Albert, P., Ratnasinghe, D., Tangrea, J., and Wacholder, S. (2001). Limitations of the case-only design for identifying gene-environment interactions. *American Journal of Epidemiology* **154**, 687–693.
- Breslow, N. E., Robins, J. M., and Wellner, J. A. (2000). On the semi-parametric efficiency of logistic regression under case-control sampling. *Bernoulli* **6**, 447–455.
- Chatterjee, N. and Carroll, R. J. (2005). Semiparametric maximum likelihood estimation exploiting gene-environment independence in case-control studies. *Biometrika* **92**, 399–418.
- Chatterjee, N. and Chen, Y. H. (2007). Maximum likelihood inference on a mixed conditionally and marginally specified regression model for genetic epidemiologic studies with two phase sampling. *Journal of the Royal Statistical Society, Series B* **69**, 123–142.
- Chen, Y. H. (2007). A semi-parametric odds ratio model for measuring association. *Biometrics* **63**, 413–421.
- Gatto, N., Campbell, U., Rundle, A., and Ahsa H. (2004). Further development of the case-only design for assessing gene-environment interaction: Evaluation of and adjustment for bias. *International Journal of Epidemiology* **33**, 1014–1024.
- Hausman, J. A. (1978). Specification tests in econometrics. *Econometrica* **46**, 1251–1272.
- Modan, M. D., Hartge, P., Hirsh-Yechezkel, G., Chetrit, A., Lubin, F., Beller, U., Ben-Baruch, G., Fishman, A., Menczer, J., Struewing, J. P., Tucker, M. A., Ebbers, S. M., Friedman, E., Piura, B., Wacholder, S., for the National Israel Ovarian Cancer Study Group. (2001). Parity, oral contraceptives and the risk of ovarian cancer among carriers and noncarriers of a BRCA1 or BRCA2 mutation. *New England Journal of Medicine* **345**, 235–240.
- Newey, W. K. (1985). Generalized method of moments specification testing. *Journal of Econometrics* **29**, 229–256.
- Piegorsch, W. W., Weinberg, C. R., and Taylor, J. A. (1994). Non-hierarchical logistic models and case-only designs for assessing susceptibility in population-based case-control studies. *Statistics in Medicine* **13**, 153–162.
- Prentice, R. L. and Pyke, R. (1979). Logistic disease incidence models and case-control studies. *Biometrika* **66**, 403–411.
- Schmidt, S. and Schaid, D. (1999). Potential misinterpretation of the case-only study to assess gene-environment interaction. *American Journal of Epidemiology* **150**, 878–885.
- Tchetgen Tchetgen, E., Robins, J., and Rotnitzky, A. (2010). On doubly robust estimation of a semi-parametric odds ratio model. *Biometrika* **97**, 171–180.
- Vanderweele, T., Hernandez-Diaz, S., and Hernan, M. (2009). Case-only gene-environment interaction studies: When does association imply mechanistic interactions? Technical Report, Harvard School of Public Health, Boston, Massachusetts.
- Vansteelandt, S., VanderWeele, T. J., Tchetgen, E., and Robins, J. M. (2008). Multiply robust inference for statistical interactions. *Journal of the American Statistical Association* **103**, 1693–1704.
- Yang, Q., Khoury, M. J., Sun, F., and Flanders, D. (1999). Case-only design to measure gene-gene interaction. *Epidemiology* **10**, 167–170.

Received February 2009. Revised December 2009.
Accepted December 2009.