



On the role of parameterization in models with a misspecified nuisance component

Heather S. Battey^{a,1,2} and Nancy Reid^{b,1,2}

Contributed by Nancy Reid; received February 8, 2024; accepted July 23, 2024; reviewed by Emmanuel J. Candès and Edward I. George

The paper is concerned with inference for a parameter of interest in models that share a common interpretation for that parameter but that may differ appreciably in other respects. We study the general structure of models under which the maximum likelihood estimator of the parameter of interest is consistent under arbitrary misspecification of the nuisance part of the model. A specialization of the general results to matched-comparison and two-groups problems gives a more explicit and easily checkable condition in terms of a notion of symmetric parameterization, leading to a broadening and unification of existing results in those problems. The role of a generalized definition of parameter orthogonality is highlighted, as well as connections to Neyman orthogonality. The issues involved in obtaining inferential guarantees beyond consistency are briefly discussed.

causality | model structure | parameter orthogonality | symmetric parameterization | model formulation

Scientific practice in nearly every field relies on the use of mathematical models: a provisional base for describing how observable quantities materialize. Some models, such as Einstein's theory of gravitation, provide predictions that can be verified with such accuracy that they are often considered to be the truth. This is, however, an atypical situation: The processes underlying most observable data are much too complicated to be described by exact laws and often require a probabilistic element. A useful statistical model captures the essence of the data-generating mechanism in a way that can accommodate new data collected under slightly different conditions. When the primary purpose of a model is to provide insight and explanation, as is typical in many scientific contexts, stable aspects should ideally be capable of interpretation. For this, statistical models will usually have a relatively small number of scientifically interpretable parameters of interest. So-called nuisance parameters are used to complete the specification of the model.

In contexts where prediction is the main focus, such as arise in many applications of machine learning, model-free approaches are sometimes advocated. This is justified on the grounds that faithful representation of a data-generating process is both unnecessary for such a task and also unrealistically optimistic with very large or complex sets of data. For instance in commercial recommender systems, the criteria of speed and accuracy of predictions are key, and the identification of interpretable parameters unimportant. In sharp contrast are contexts in which the goal is to understand how a treatment or exposure influences the distribution of outcomes, having accounted, to the extent feasible, for the diversity of individuals. Breiman (1) described this separation succinctly as the "two cultures" of statistical thought. The development in this paper is closer to the view expressed by Cox (2) in the discussion of that work.

We consider sets of statistical models that share a common interpretation for the parameter of subject-matter interest but differ in their specification of the nuisance component. A fundamental question is whether reliable inference for the interest parameter is still achievable using standard approaches when the nuisance aspect is misspecified. An exemplar setting is in understanding the effect of treatment or exposure on the health outcomes of individuals. These individuals may exhibit considerable heterogeneity and complex interdependencies stemming, for instance, from certain shared genetic traits. Modeling such interdependencies is a formidable challenge and often tackled by introducing a parameter for each individual. If these are, postulated to have been drawn from a parametric distribution, the resulting doubly stochastic models are called random-effects models. Since the choice of parametric distribution for the individual-specific nuisance parameters is typically based on mathematical convenience, the prospect of misspecification seems likely.

The extensive literature on inference under model misspecification appears to have started with Cox (3), who considered two nonnested models for a vector random

Significance

Statistical models are often chosen based on a combination of scientific understanding and flexibility or mathematical convenience. While the aspects of core scientific relevance may be relatively securely specified in terms of interpretable interest parameters, the rest of the formulation is often chosen somewhat arbitrarily. In many statistical models this formulation includes so-called nuisance parameters, which are of no direct subject-matter concern, but are needed to complete the model or reflect the complexity of the data. This paper contributes to the foundations of statistics by studying the interplay between model structure and likelihood inference, under misspecification of the nuisance component.

Author affiliations: ^aDepartment of Mathematics, Imperial College, London SW7 2AZ, United Kingdom; and ^bDepartment of Statistical Sciences, University of Toronto, Toronto, ON M5G 1Z5, Canada

Author contributions: H.S.B. and N.R. designed research; performed research; and wrote the paper.

Reviewers: E.J.C., Stanford University; and E.I.G., University of Pennsylvania.

The authors declare no competing interest.

Copyright © 2024 the Author(s). Published by PNAS. This open access article is distributed under Creative Commons Attribution License 4.0 (CC BY).

¹H.S.B. and N.R. contributed equally to this work.

²To whom correspondence may be addressed. Email: h.battey@imperial.ac.uk or nancym.reid@utoronto.ca.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2402736121/-DCSupplemental>.

Published August 30, 2024.

variable Y , specifying joint density or mass functions m and \check{m} , say, where the latter is completely known up to a finite-dimensional parameter θ . With $\hat{\theta}$ the maximum likelihood estimator under model \check{m} , Cox (3) showed that when the true density function is m , $\hat{\theta}$ is asymptotically normal with expected value θ_m^0 and covariance matrix given by what is now known as the sandwich formula. The quantity θ_m^0 solves

$$\mathbb{E}_m[\nabla_{\theta} \log \check{m}(Y; \theta)]_{\theta=\theta_m^0} = 0,$$

where \mathbb{E}_m is expectation under m . Equivalently θ_m^0 minimizes with respect to θ the Kullback–Leibler divergence

$$\int m(y) \log \left(\frac{m(y)}{\check{m}(y; \theta)} \right) dy.$$

More rigorous discussions of the distribution theory and regularity conditions were provided by Huber (4), Kent (5), and White (6, 7).

Inferential guarantees for the quantity θ_m^0 follow directly under classical regularity conditions on the family of models. Our interest, by contrast, is in the estimation of the true value of an interest parameter, which is assumed to have a common interpretation in both the true distribution and in the fitted model. Our focus is on uncovering the structure for which standard likelihood theory retains its first-order theoretical guarantees, providing foundational insight into the limits of likelihood inference. Reference to the unsolved nature of this question is widespread; see, for instance, Evans and Didelez (8, section 5).

A different line of exploration, with connections to the double-robustness literature [e.g., Robins et al. (9) and Chernozhukov et al. (10)] is the so-called assumption-lean approach to modeling of Vansteelandt and Dukes (11), who center their analyses on model-free estimands rather than parameters of a given model. The ideal estimand coincides with a particular model parameter when the chosen model includes m and retains a degree of interpretability when the model is misspecified. The tension between estimands and model parameters has a long history, going back to Fisher and Neyman; see Cox (12).

In the discussion of ref. 11, Battey (13) conjectured a condition under which the maximum likelihood estimator for an interest parameter is consistent in spite of arbitrary misspecification of the nuisance part of the model. The intuition for that incomplete claim came from a highly involved calculation in a particular paired-comparisons model studied in Battey and Cox (14). Any role played by the assumed model for the random effects was unclear from that calculation and is clarified by the more general insight provided here.

In Section 1, we study the general structure for such a consistency result, and specialize the analysis to settings of common relevance, providing more explicit conditions on the parameterization for matched-pair and two-group problems. An important conclusion from this analysis is that the conditions for consistency can sometimes be checked without knowledge of the true model. Section 2 shows through this route that some key insights have been overlooked in a rather large literature on misspecified random effects distributions, particularly the role of parameterization. Section 2.C recovers two well-known results for generalized linear models as an illustration of the more general statements, and Section 2.D provides a recent example in the context of a marginal structural model, for which the more general results of Section 1 provide elucidation. While some of

the material in Section 1 is rather technical, we have aimed to provide intuition for the main results and have provided several examples in Section 2. The examples are deliberately detached from specific subject-matter details, in order that they be broadly useful for a range of scientific applications.

1. Consistency in General Misspecified Models

1.A. General Conditions for Consistency. Let m represent the density function for the outcomes, parameterized in terms of an interest parameter ψ with true value ψ^* . The assumed model, while sharing the same interpretable interest parameter, is misspecified in other ways. The joint density function for the observations under the assumed model has parameters $(\psi, \lambda) \in \Psi \times \Lambda$, and $\ell(\psi, \lambda) = \log \check{m}(y; \psi, \lambda)$ denotes the observed log-likelihood function for that model, viewed as a function of (ψ, λ) for observed data $y = (y_1, \dots, y_n)$. As above, maximization of $\ell(\psi, \lambda)$ gives estimates $(\hat{\psi}, \hat{\lambda})$; their dependence on y is suppressed in the notation. As functions of the random variable $Y = (Y_1, \dots, Y_n)$ the probability limit of the maximum likelihood estimator, as $n \rightarrow \infty$, is (ψ_m^0, λ_m^0) , the solution to

$$\mathbb{E}_m[\nabla_{(\psi, \lambda)} \ell(\psi_m^0, \lambda_m^0)] = 0. \quad [1.1]$$

We consider the model to be misspecified if there are no values of λ in Λ for which the true density m is recovered. This precludes the situation in which the true distribution belongs to a submodel of an assumed encompassing family, briefly discussed in Section 3.C.

The following example will be used repeatedly to illustrate key ideas and notation.

Example 1.1: Consider a matched comparison problem in which, for each of n twin pairs, one individual from each pair is chosen at random to receive a treatment, the other being the untreated control. Let Y_{i1} and Y_{i0} denote the time until some medical event of interest (e.g., recovery from illness) for the treated and untreated individual in the i th pair. A simple parametric model specifies the outcomes Y_{i1} and Y_{i0} as exponentially distributed of rates $\gamma_i \psi > 0$ and $\gamma_i / \psi > 0$, respectively. The pair-specific nuisance parameter γ_i captures, for instance, genetic differences between the pairs of twins. The parameter of interest ψ with true value ψ^* quantifies the effect of the treatment: ψ^2 being the multiplicative effect of the treatment on the instantaneous probability of recovery relative to baseline. Other parameterizations are possible, but the above symmetric parameterization will be shown to play an important role. Suppose the pair effects are modeled as gamma distributed of shape κ and rate ρ . Then, under the assumed model, the joint density function for the outcomes in any given pair at (y_1, y_0) is

$$\frac{\Gamma(\kappa + 2)\rho^\kappa}{\Gamma(\kappa)(\gamma_1 \psi + \gamma_0 / \psi + \rho)^{\kappa+2}}. \quad [1.2]$$

The true random effects distribution could have been quite different, so that the model Eq. 1.2 is misspecified. Nevertheless, the interpretation of the interest parameter ψ is stable over the different specifications. The notional nuisance parameter is $\lambda = (\kappa, \rho)$.

Remark 1.1: A different analysis is possible in this example, treating the pair effects $\gamma_1, \dots, \gamma_n$ as fixed arbitrary constants, which can be eliminated from the analysis by basing likelihood inference on the distribution of the ratios $Z_i = Y_{i1}/Y_{i0}$; this is discussed briefly in Section 3.B.

Definition 1.1 (parameter m -orthogonality): Let $\nabla_{\psi\lambda}^2 \ell(\psi, \lambda)$ denote the cross-partial derivative of the assumed log-likelihood function. The parameter ψ is said to be m -orthogonal to the notional parameter λ if $\mathbb{E}_m[\nabla_{\psi\lambda}^2 \ell(\psi, \lambda)] = 0$. This can hold globally for all (ψ, λ) or at particular values, the notation $\Psi \perp_m \Lambda$ indicating global m -orthogonality and $\Psi \perp_m \lambda$ and $\psi \perp_m \Lambda$ indicating, respectively, local m -orthogonality at λ for any ψ , and local m -orthogonality at ψ for any λ .

To gain some intuition for parameter m -orthogonality, consider first the much stronger condition, $\nabla_{\psi\lambda}^2 \ell(\psi, \lambda) = 0$ for all ψ and λ . This corresponds to what is called a cut in the parameter space and implies $\mathbb{E}_m[\nabla_{\psi\lambda}^2 \ell(\psi, \lambda)] = 0$ for any m so that the true density m trivially plays no role. In this strongest setting, parameter orthogonality is a purely geometric property of the log-likelihood function that holds only in certain parameterizations. It is effectively an absence of torsion, implying that, for any fixed λ , the corresponding log-likelihood function over Ψ is maximized at the same point. This parameter separation in the likelihood function is, however, limited to a relatively small number of models. The weaker condition $\mathbb{E}_m[\nabla_{\psi\lambda}^2 \ell(\psi, \lambda)] = 0$ only requires the absence of torsion to hold on average over hypothetical repeated draws from the true distribution. If the model is correctly specified, the usual definition of parameter orthogonality with respect to Fisher information is recovered.

Propositions 1.1 and 1.2 give two alternative general conditions for consistency of $\hat{\psi}$ in spite of arbitrary misspecification of the nuisance part of the model. Their proofs, and those of subsequent results, are in *SI Appendix*.

Proposition 1.1. *Let the observed log-likelihood function for the assumed model be strictly concave as a function of (ψ, λ) . Then, $\psi_m^0 = \psi^*$ if and only if $\mathbb{E}_m[\nabla_{\psi} \ell(\psi^*, \lambda_m^0)] = 0$. The latter condition is equivalent to $\mathbb{E}_m[\nabla_{\psi} \ell(\psi^*, \lambda)] = 0$ for all λ if and only if $\psi^* \perp_m \Lambda$.*

Remark 1.2: Local orthogonality at a particular value, λ' say, is not sufficient to ensure that $\mathbb{E}_m[\nabla_{\psi} \ell(\psi^*, \lambda')] = 0$ implies $\mathbb{E}_m[\nabla_{\psi} \ell(\psi^*, \lambda_m^0)] = 0$, as is clear from the proof of Proposition 1.1.

Remark 1.3: In Example 1.1, it was shown in ref. 14 that ψ is globally m -orthogonal to $\lambda = (\kappa, \rho)$ for any distribution over the random effects, and the consistency of the maximum likelihood estimator of ψ was established through an explicit calculation. An initial motivation for the present work was to understand how sensitive that conclusion was to various aspects of the model formulation.

The first part of Proposition 1.1, i.e., the requirement of an unbiased estimating equation $\mathbb{E}_m[\nabla_{\psi} \ell(\psi^*, \lambda_m^0)] = 0$, is almost immediate. Since λ_m^0 depends on the unknown m and is typically unavailable, verification of the orthogonality condition $\psi^* \perp_m \Lambda$ uniformly over m , in the event of its validity, appears to be the simplest way of establishing consistency. The simplest special cases in which orthogonality can be checked are those in which there is a parameter cut. Such examples include misspecification of the dispersion component of a generalized linear model, discussed in Section 2.C, and Gaussian linear mixed models with a misspecified distribution over the random effects. Robustness of inference in the latter case has often been observed empirically, e.g., Schielzeth et al. (15) without reference to any underlying structure. A more elaborate example with a parameter cut is the

causal model of ref. 8 outlined in Section 2.D. Example 1.1 obeys the weaker form $\mathbb{E}_m[\nabla_{\psi\lambda}^2 \ell(\psi, \lambda)] = 0$ from Proposition 1.1, as does the unbalanced two-group problem of Example 2.4 in Section 2.A, in which the pair-specific parameters are replaced by stratum-specific parameters. The relevant structure underpinning these paired and two-group examples is elucidated in Section 1.C, where some intuition is provided.

The unbalanced two-group problem of Example 2.3, on the other hand, violates the m -orthogonality assumption of Proposition 1.1 with respect to one of its two nuisance parameters. Nevertheless, Cox and Wong (16) showed that the maximum likelihood estimator of the interest parameter in Example 2.3 is consistent, suggesting that Proposition 1.1, while easier to verify, is too strong for some situations. Proposition 1.2 presents weaker conditions for consistency.

Suppose that the orthogonality condition $\psi^* \perp_m \Lambda$ from Proposition 1.1 fails, so that there exists at least one $\lambda \neq \lambda_m^0$ such that $\mathbb{E}_m[\nabla_{\psi} \ell(\psi^*, \lambda)] \neq 0$. Let

$$g_{\psi}(\psi^*, \lambda) := \mathbb{E}_m[\nabla_{\psi} \ell(\psi^*, \lambda)], \quad g_{\lambda}(\psi^*, \lambda) := \mathbb{E}_m[\nabla_{\lambda} \ell(\psi^*, \lambda)], \quad [1.3]$$

and partition the inverse of the Fisher information matrix $i := i(\psi^*, \lambda) := \mathbb{E}_m[-\nabla_{(\psi, \lambda)}^2 \ell(\psi^*, \lambda)]$ as

$$\begin{pmatrix} i^{\psi\psi} & i^{\psi\lambda} \\ i^{\lambda\psi} & i^{\lambda\lambda} \end{pmatrix} := \begin{pmatrix} i_{\psi\psi} & i_{\psi\lambda} \\ i_{\lambda\psi} & i_{\lambda\lambda} \end{pmatrix}^{-1}. \quad [1.4]$$

In Proposition 1.2 and its proof, we refer to $g_{\psi}(\psi^*, \lambda)$ and $g_{\lambda}(\psi^*, \lambda)$ in the shorthand g_{ψ} and g_{λ} respectively.

Proposition 1.2. *If $i^{\psi\psi} g_{\psi} + i^{\psi\lambda} g_{\lambda} = 0$ for all $\lambda \in \Lambda$, then $\psi_m^0 = \psi^*$. If ψ and λ are both scalar parameters, the condition reduces to $i_{\lambda\lambda} g_{\psi} = i_{\psi\lambda} g_{\lambda}$.*

That Proposition 1.2 is more general than Proposition 1.1 is seen on noting that the orthogonality condition of Proposition 1.1 implies $i^{\psi\lambda} = 0$ and $g_{\psi} = 0$, so that the condition of Proposition 1.2 also holds. The conditions of Propositions 1.1 and 1.2 are clearly not necessary. In particular, if λ_m^0 can be calculated, the route to establishing consistency or inconsistency of $\hat{\psi}$ is more direct. Any practical relevance, beyond general theoretical understanding, is to situations in which λ_m^0 is not calculable, typically because m is unknown. The conditions in Propositions 1.1 and 1.2, while dependent on m , are nevertheless checkable in some situations, as illustrated in the examples of subsequent sections.

A generalization involves a second nuisance parameter vector, $v \in \mathcal{N}$ say, such that $\psi^* \perp_m \mathcal{N}$ and $\lambda \perp_m \mathcal{N}$ in Proposition 1.1 or $\Lambda \perp_m \mathcal{N}$ in Proposition 1.2. It can be shown via a straightforward extension of the arguments leading to Propositions 1.1 and 1.2 that the conclusion of the latter is unchanged by this modification. Examples involving a second nuisance parameter are provided in Section 2.

The consistency demonstrated by Battey and Cox (14) for Example 1.1 arises as a special case of Proposition 1.1, while the argument of Cox and Wong (16) for an unbalanced doubly stochastic two-group problem turns out to be an application of Proposition 1.2 and Corollary 1.1 below. We return to this example in Section 2.B.

1.B. Parameter Orthogonality and Orthogonalization in Misspecified Models. In view of Propositions 1.1 and 1.2, a key question is when the assumed Fisher information matrix, with

expectation computed under an erroneous model, coincides with the true Fisher information matrix in some or all regions of the parameter space. Such coincidence implies in particular that parameter orthogonality established under a notional model is true more generally. The practical import is that otherwise $\Psi \perp_m \Lambda$ or its local analog would typically not be verifiable without knowledge of m . A partial answer is presented as *Proposition 1.3*, whose proof is immediate by the definition of sufficiency.

Proposition 1.3. *Suppose that the cross-partial derivative of the assumed log-likelihood function $\nabla_{\psi\lambda}^2 \ell(\psi, \lambda)$ depends on the data only through a sufficient statistic $S = (S_1, \dots, S_k)$, for some $1 \leq k \leq n$. Write $\check{y}_{\psi\lambda}(\psi, \lambda) = \mathbb{E}_{(\psi, \lambda)}[-\nabla_{\psi\lambda}^2 \ell(\psi, \lambda)]$, where $\mathbb{E}_{(\psi, \lambda)}$ means expectation under the assumed model. Provided that $\nabla_{\psi\lambda}^2 \ell(\psi, \lambda)$ is additive in S_1, \dots, S_k and $\mathbb{E}_m(S_j) = \mathbb{E}_{(\psi, \lambda)}(S_j)$ for all j , then $i_{\psi\lambda}(\psi, \lambda) = \check{y}_{\psi\lambda}(\psi, \lambda)$.*

Corollary 1.1. *Under the conditions of Proposition 1.3, a sufficient condition for parameter m -orthogonality $\Psi \perp_m \Lambda$ is $\check{y}_{\psi\lambda}(\psi, \lambda) = 0$ for all $(\psi, \lambda) \in \Psi \times \Lambda$, and similarly for local m -orthogonality.*

Remark 1.4: For Example 1.1, m -orthogonality can be established more directly via *Proposition 2.1*.

In general, a reparameterization of a model from parameters ϕ to $\theta = \theta(\phi)$ is called interest-respecting if the parameter of interest is common to both: if $\phi = (\psi, \xi)$ then $\theta = (\psi, \lambda(\psi, \xi))$. The relevance of interest-respecting reparameterizations in scientific work is that the more important interpretable aspects of the model are retained. In ref. 17, interest-respecting orthogonal reparameterizations were explored for improving inference in correctly specified models. The role in misspecified settings is amplified in view of *Proposition 1.1*, which concerns the first-order properties of the estimator.

In a slightly more explicit notation, suppose that ϕ is a parameterization for which the (r, s) th component satisfies $i_{rs}^{(\phi)}(\phi) = \check{y}_{rs}^{(\phi)}(\phi) = 0$. Consider an interest-respecting reparameterization from an initial parameter ϕ to a new parameter θ , with coordinates

$$\phi^r(\theta^1, \dots, \theta^p), \quad \theta^a(\phi^1, \dots, \phi^p), \quad (a, r = 1, \dots, p),$$

where $\theta^1 = \phi^1 = \psi$. An implication of *Corollary 1.1* is that the approach of Cox and Reid (17) can still be used in misspecified models, provided that the additional structure of *Proposition 1.3* is present in the parameterization that is orthogonal under the assumed model. Specifically, from a starting parameterization ϕ , an orthogonal parameterization θ is any solution in $\theta^2, \dots, \theta^p$ to the system of $p - 1$ differential equations

$$\check{y}_{ab}^{(\phi)} \frac{\partial \phi^a(\theta)}{\partial \theta^1} + \check{y}_{1b}^{(\phi)} = 0, \quad b = 2, \dots, p,$$

or in matrix notation

$$\frac{\partial \phi(\theta)}{\partial \theta^1} = -(\check{y}_{\cdot\cdot}^{(\phi)})^{-1}(\check{y}_{\cdot 1}^{(\phi)}),$$

where $\check{y}_{\cdot\cdot}^{(\phi)}$ is the Fisher information matrix under the assumed model without the row and column corresponding to ϕ^1 , and $\check{y}_{\cdot 1}^{(\phi)}$ is the excluded column.

Even under the assumptions of *Proposition 1.3*, with model misspecification, the asymptotic variance of $\hat{\psi}$ is not $i^{\psi\psi}$ but

is given by the sandwich formula of ref. 3. This means that inferential guarantees beyond consistency are not, in general, available using standard likelihood-based approaches; see Section 3.A for a brief discussion of some exceptions.

1.C. Symmetric Parameterizations and Induced Antisymmetry

For the matched-pair and two-group examples of Sections 2.A and 2.B, a fundamental role is played by symmetric parameterization in transformation models, which induces an antisymmetry on the log-likelihood derivative and thereby the m -orthogonality of *Proposition 1.1* under arbitrary misspecification of the model. The class of problems covered in this subsection exemplifies settings in which m -orthogonality can be straightforwardly verified without knowledge of m . A simplified explanation is that in transformation models, there is a convenient duality between the parameter space and the sample space which results in the cancelation of the relevant terms in $\mathbb{E}_m[\nabla_{\psi\lambda}^2 \ell(\psi, \lambda)]$ provided that the appropriate parameterization is used.

Let \mathcal{P} be a set of continuous probability measures on \mathcal{Y} . A transformation model under the action of G [e.g., Barndorff-Nielsen and Cox (18, p. 53)] is a subset of \mathcal{P} parameterized by $\gamma \in \Gamma$ and possibly other parameters suppressed in the notation, such that

$$\mathcal{P}_G := \{p \in \mathcal{P} : p(gE; g\gamma) = p(E; \gamma), \quad g \in G, E \in \mathcal{E}, \gamma \in \Gamma\},$$

where \mathcal{E} is the set of measurable events on the sample space \mathcal{Y} . The definition implies that G acts on both the sample space and the parameter space. In particular, the action of G on \mathbb{F} , say, is a continuous map $G \times \mathbb{F} \rightarrow \mathbb{F}$, $(g, x) \rightarrow gx$ for $x \in \mathbb{F}$, where in the transformation model, x represents either the data point $y \in \mathcal{Y}$ or the parameter value $\gamma \in \Gamma$. In general, \mathcal{Y} and Γ need not be equal, although this will often be the case. The identity element on \mathcal{Y} or on Γ is written $e = g^{-1}g = gg^{-1}$, the context leaving no ambiguity.

The relevant group actions for present purposes depend only on the interest parameter ψ , so that we may write, in a more explicit notation, $g = g_\psi \in G$. Two examples serve to illustrate the properties of the transformation models. In location models, $\mathcal{Y} = \Gamma = \mathbb{R}$, the group action is addition and $g_\psi\gamma = \gamma + \psi$, giving $gg^{-1} = g_\psi g_{-\psi} = e = g_0$. In scale models, $\mathcal{Y} = \Gamma = \mathbb{R}^+$, the group action is multiplication and $g_\psi\gamma = \gamma\psi$, giving $gg^{-1} = g_\psi g_{(1/\psi)} = e = g_1$.

Let U be a random variable with a distribution depending only on γ . This is best thought of as the random variable at baseline with respect to ψ , i.e., $\psi = 0$ in location models and $\psi = 1$ in scale models. Write the probability density function of U at u as $f_U(u; \gamma)du$. The following definition plays an important role in establishing consistency for a treatment parameter in matched-pair and two-groups problems.

Definition 1.2 (symmetric parameterization): Let Y_1 and Y_0 be independent random variables with probability measures in \mathcal{P}_G and density functions f_1 and f_0 , respectively. Their joint distribution is said to be parameterized ψ -symmetrically with respect to (ψ, γ) if $g = g_\psi \in G$ depends only on ψ and if the density functions f_1 and f_0 relate to f_U by

$$f_U(u; \gamma)du = f_1(gu; g\gamma)d(gu) = f_0(g^{-1}u; g^{-1}\gamma)d(g^{-1}u).$$

In other words, the joint density function, when expressed in terms of $u_1 = g^{-1}y_1$ and $u_0 = gy_0$, is symmetric in u_1 and u_0 :

$$f_1(y_1; g\gamma)f_0(y_0; g^{-1}\gamma)dy_1dy_0 = f_U(u_1; \gamma)f_U(u_0; \gamma)du_1du_0. \quad [1.5]$$

Definition 1.2 says that the random variables $U_1 = g^{-1}Y_1$ and $U_0 = gY_0$ are equal in distribution to U , which has a “standardized form.” The inverted commas express that f_U might not be a standardized form in the conventional sense, as it depends on γ and possibly other parameters suppressed in the notation. The symmetric parameterization can equivalently be written

$$\begin{aligned} f_1(y_1; g\gamma)dy_1 &= f_U(g^{-1}y_1; \gamma)J_1^+ dy_1 \\ f_0(y_0; g^{-1}\gamma)dy_0 &= f_U(gy_0; \gamma)J_0^+ dy_0, \end{aligned} \quad [1.6]$$

where

$$J_1^+ = |d(g^{-1}y_1)/dy_1|, \quad J_0^+ = |d(gy_0)/dy_0|,$$

satisfy $J_1^+ J_0^+ = 1$. In a location model, Eq. 1.6 becomes

$$\begin{aligned} f_1(y_1; \gamma + \psi)dy_1 &= f_U(y_1 - \psi; \gamma)dy_1 \\ f_0(y_0; \gamma - \psi)dy_0 &= f_U(y_0 + \psi; \gamma)dy_0, \end{aligned} \quad [1.7]$$

while in a scale model Eq. 1.6 becomes

$$\begin{aligned} f_1(y_1; \gamma\psi)dy_1 &= f_U(y_1/\psi; \gamma)(1/\psi)dy_1 \\ f_0(y_0; \gamma/\psi)dy_0 &= f_U(y_0\psi; \gamma)\psi dy_0. \end{aligned} \quad [1.8]$$

Distributions within the scale family are often parameterized in terms of rate, or inverse scale, which reverses the roles of g and g^{-1} relative to their actions in the scale parameterization. Example 1.1 is of this form.

In a location-scale model, the action of G is $g = g_{\psi_2} \circ g_{\psi_1}$, where g_{ψ_1} is multiplication by a positive scalar ψ_1 and g_{ψ_2} is addition of a scalar ψ_2 , its inverse being $g^{-1} = g_{\psi_2}^{-1} \circ g_{\psi_1}^{-1}$. Thus, a two-parameter symmetric parameterization in a location-scale model is in terms of $g\gamma = \psi_1\gamma + \psi_2$ and $g^{-1}\gamma = (\gamma - \psi_2)/\psi_1$. For the applications we have in mind, the interest parameter ψ represents a treatment effect. Thus, if the treatment is assumed to affect either the location or the scale but not both, a location-scale distribution can effectively be treated as either location or scale. Examples include the extreme-value distributions, the most common parametric families arising in renewal theory and used in survival modeling, and members of the elliptically symmetric family in one or more dimensions, which encompasses the Gaussian, Student-t, Cauchy, and logistic distributions.

Consider, as a function of ψ alone, the log-likelihood contribution of y_1 and y_0 , realizations of Y_1 and Y_0 . This is of the form

$$\begin{aligned} \ell(\psi; \gamma, y_1, y_0) &= \log L(\psi; \gamma, y_1, y_0) \\ &= \log f_1(y_1; g\gamma) + \log f_0(y_0; g^{-1}\gamma). \end{aligned}$$

Variables to the right of the semicolon in the log-likelihood function ℓ are treated as fixed (although arbitrary), together with any other parameters suppressed in the notation.

Definition 1.3 (antisymmetry): The symmetric parameterization of f_1 and f_0 is said to induce antisymmetry on the associated log-likelihood derivative with respect to ψ if $\nabla_{\psi}\ell(\psi; \gamma, y_1, y_0)$, when expressed in terms of $u_1 = g^{-1}y_1$ and $u_0 = gy_0$, satisfies $\nabla_{\psi}\ell(\psi; \gamma, u_1, u_0) = -\nabla_{\psi}\ell(\psi; \gamma, u_0, u_1)$.

Example 1.2 (Continuation of Example 1.1): In the exponential matched pair problem, the gamma distribution over the random effects is irrelevant for illustrating the symmetric parameterization and induced antisymmetry, as these definitions are conditional on

γ within a single pair. Since multiplication of rates corresponds to division of scale, it is natural, for consistency with Eq. 1.8, to define the group operation g as division by ψ rather than multiplication by ψ . Thus let $u_1 = g^{-1}y_1 = \psi y_1$ and $u_0 = gy_0 = y_0/\psi$. Conditionally on γ , the joint density function of (Y_1, Y_0) is

$$f_1(y_1; \psi^*, \gamma)f_0(y_0; \psi^*, \gamma)dy_1 dy_0 = \gamma^2 \exp\{-\gamma(u_1 + u_0)\} du_1 du_0,$$

which is symmetric in u_1 and u_0 . The log-likelihood derivative with respect to ψ is

$$\nabla_{\psi}\ell(\psi; u_0, u_1) = -\gamma(y_1 - y_0/\psi) = -\gamma(u_1/\psi - u_0/\psi),$$

where the right hand side is $-(-\gamma(u_0/\psi - u_1/\psi)) = -\nabla_{\psi}\ell(\psi; u_1, u_0)$. This shows that the log-likelihood derivative is antisymmetric in the sense of Definition 1.3.

The main transformation models arising in statistics are the location models and the scale models. For these, we show in Example 1.3 that the symmetric parameterization automatically induces antisymmetry on the associated log-likelihood derivative. We also show this for a rotation family on the circle in Example 1.4. We conjecture that antisymmetry according to Definition 1.3 is a necessary consequence of the ψ -symmetric parameterization for any transformation model, but in the absence of general group-theoretic proof, we present the following necessary and sufficient conditions for asymmetry, that may be checked on a case-by-case basis for more exotic groups. Examples 1.3 and 1.4 illustrate the application of Proposition 1.4.

Proposition 1.4. Suppose that the joint distribution of Y_1 and Y_0 is parameterized ψ -symmetrically with respect to (ψ, γ) in the sense of Definition 1.2. The parameterization induces antisymmetry on the log-likelihood derivative in the sense of Definition 1.3 if and only if

$$a(u_1, u_0) := \frac{\partial u_1}{\partial \psi} + \frac{\partial u_0}{\partial \psi} = -a(u_0, u_1),$$

and

$$c(u_1, u_0) := \left(\frac{\partial}{\partial \psi} \left| \frac{\partial u_1}{\partial y_1} \right| \right) \left| \frac{\partial u_0}{\partial y_0} \right| + \left(\frac{\partial}{\partial \psi} \left| \frac{\partial u_0}{\partial y_0} \right| \right) \left| \frac{\partial u_1}{\partial y_1} \right| = -c(u_0, u_1),$$

where $u_1 = g^{-1}y_1$ and $u_0 = gy_0$.

A clearer but more cumbersome expression of the conditions $a(u_1, u_0) = -a(u_0, u_1)$ and $c(u_1, u_0) = -c(u_0, u_1)$ makes the dependence of the partial derivatives on u_1 and u_0 explicit. In this more explicit notation, $a(u_1, u_0) = -a(u_0, u_1)$ amounts to $(\partial u_1/\partial \psi)(u_1) = -(\partial u_0/\partial \psi)(u_1)$ and vice versa.

Example 1.3: In a symmetric parameterization of a location model $\partial u_1/\partial \psi = -1 = -\partial u_0/\partial \psi$ and

$$\left(\frac{\partial}{\partial \psi} \left| \frac{\partial u_1}{\partial y_1} \right| \right) \left| \frac{\partial u_0}{\partial y_0} \right| = 0 = -\left(\frac{\partial}{\partial \psi} \left| \frac{\partial u_0}{\partial y_0} \right| \right) \left| \frac{\partial u_1}{\partial y_1} \right|.$$

Thus, $a(u_1, u_0) = 0 = -a(u_0, u_1)$ and $c(u_1, u_0) = 0 = -c(u_0, u_1)$. In a symmetric parameterization of a scale model $\partial u_1/\partial \psi = -u_1/\psi$, $\partial u_0/\partial \psi = u_0/\psi$, and

$$\left(\frac{\partial}{\partial \psi} \left| \frac{\partial u_1}{\partial y_1} \right| \right) \left| \frac{\partial u_0}{\partial y_0} \right| = -\frac{1}{\psi} = -\left(\frac{\partial}{\partial \psi} \left| \frac{\partial u_0}{\partial y_0} \right| \right) \left| \frac{\partial u_1}{\partial y_1} \right|,$$

so that $a(u_1, u_0) = -a(u_0, u_1)$ and $c(u_1, u_0) = 0 = -c(u_0, u_1)$. Thus both cases satisfy the condition of Proposition 1.4.

Example 1.4: For a rotation model on \mathbb{R}^p , $\mathcal{Y} \subset \mathbb{R}^p$, $\Gamma = [0, 2\pi)$ and the group action is matrix multiplication by a rotation matrix $g \in G \subseteq \text{SO}(p)$, where $\text{SO}(p)$ is the special orthogonal group of dimension p . The identity element is matrix multiplication by the p -dimensional identity matrix and $f_U(u; \gamma)$ is a distribution on the hypersphere, where γ is a rotation parameter. The von Mises–Fisher distribution is one such example. Consider the case of $p = 2$.

Let $v \in \mathbb{R}^2$ be a point on a circle of fixed arbitrary radius, either a location parameter of a probabilistic model or a data point on the circle. In a symmetric parameterization of a rotation model, $v_\gamma = g_\gamma v$ defines a (counterclockwise) rotation of angle $v_\gamma^\top v = \gamma$ and $v_1 = g_\psi v_\gamma$, $v_0 = g_\psi^{-1} v_\gamma$ define (counterclockwise and clockwise) rotations of angles $v_1^\top v_\gamma = \psi$ and $v_0^\top v_\gamma = -\psi$, respectively, where

$$g_\psi^{-1} = \begin{pmatrix} \cos \psi & -\sin \psi \\ \sin \psi & \cos \psi \end{pmatrix}^{-1} = \begin{pmatrix} \cos \psi & \sin \psi \\ -\sin \psi & \cos \psi \end{pmatrix} \\ = \begin{pmatrix} \cos(-\psi + 2\pi) & -\sin(-\psi + 2\pi) \\ \sin(-\psi + 2\pi) & \cos(-\psi + 2\pi) \end{pmatrix} \in G.$$

and $g_\psi g_\gamma = g_{\gamma+\psi} \in G$, $g_\psi^{-1} g_\gamma = g_{\gamma-\psi} \in G$. Thus, in Eq. 1.6 g_γ should be understood either as addition of angles $g_\gamma = \gamma + \psi$ or multiplication of rotation matrices $g_\gamma = g_\psi g_\gamma$.

Let $y_1 = (y_{11}, y_{12})^\top$, $y_0 = (y_{01}, y_{02})^\top$. Then

$$u_1 = g^{-1} y_1 = \begin{pmatrix} y_{11} \cos \psi + y_{12} \sin \psi \\ -y_{11} \sin \psi + y_{12} \cos \psi \end{pmatrix}, \\ u_0 = g y_0 = \begin{pmatrix} y_{01} \cos \psi - y_{02} \sin \psi \\ y_{01} \sin \psi + y_{02} \cos \psi \end{pmatrix}.$$

The Jacobian determinants J_1^+ and J_0^+ are both equal to $|\cos^2 \psi + \sin^2 \psi| = 1$. Thus, $c(u_1, u_0) = 0 = -c(u_0, u_1)$ in Proposition 1.4. Since

$$\frac{\partial u_1}{\partial \psi} = \begin{pmatrix} -y_{11} \sin \psi + y_{12} \cos \psi \\ -y_{11} \cos \psi - y_{12} \sin \psi \end{pmatrix} = g_\theta^{-1} u_1, \\ \frac{\partial u_0}{\partial \psi} = \begin{pmatrix} -y_{01} \sin \psi - y_{02} \cos \psi \\ y_{01} \cos \psi - y_{02} \sin \psi \end{pmatrix} = g_\theta u_0,$$

with $\theta = \pi/2$, the quantity $a(u_1, u_0)$ from Proposition 1.4 is

$$a(u_1, u_0) = \begin{pmatrix} \cos(-\pi/2) & -\sin(-\pi/2) \\ \sin(-\pi/2) & \cos(-\pi/2) \end{pmatrix} u_1 \\ + \begin{pmatrix} \cos(\pi/2) & -\sin(\pi/2) \\ \sin(\pi/2) & \cos(\pi/2) \end{pmatrix} u_0 = -a(u_0, u_1),$$

verifying the conditions for the ψ -symmetric parameterization of rotation families on the circle.

2. Examples

2.A. Matched Pairs. Let Y_{i1} and Y_{i0} be random variables corresponding to observations on treated and untreated individuals in a matched pair design, where $i = 1, \dots, n$ is the pair index. The treatment effect is represented by a parameter ψ while the pair effects are encapsulated in the pair-specific nuisance parameter γ_i , avoiding explicit modeling assumptions in terms of covariates. As noted in Example 1.1, two approaches to analysis treat the pair effects as fixed arbitrary constants, or as independent and

identically distributed random variables. In the latter case, it is common to model the distribution parametrically, typically producing efficiency gains in the treatment effect estimator if the parametric assumptions hold to an adequate order of approximation.

The relevant calculations are the same for every pair, and we therefore omit the pair index i from the notation. Let the true joint density function of (Y_1, Y_0) be given by

$$m(y_1, y_0) = \int f_1(y_1; \psi^*, \gamma) f_0(y_0; \psi^*, \gamma) f(\gamma) d\gamma, \quad [2.1]$$

with $f(\gamma)$ an unknown density function for the nuisance parameters, which are treated as independent and identically distributed across pairs.

Conditionally on γ , let $\ell(\psi; \gamma, y_1, y_0)$ denote the log-likelihood contribution for ψ from a single pair.

Proposition 2.1. Suppose that the assumed model over γ is parameterized by λ , producing a log-likelihood contribution $\ell(\psi, \lambda; y_1, y_0)$, assumed strictly concave. Suppose further that, conditionally on γ , Y_1 and Y_0 have a joint distribution that is parameterized ψ -symmetrically (Definition 1.2). Then, provided that the group induces antisymmetry on the log-likelihood derivative (Definition 1.3), it follows that $\psi^* \perp_m \Lambda$ and

$$0 = \mathbb{E}_m[\nabla_\psi \ell(\psi^*, \lambda)] \\ = \int_{\mathcal{Y}} \int_{\mathcal{Y}} (\nabla_\psi \ell(\psi^*, \lambda; y_1, y_0)) m(y_1, y_0) dy_1 dy_0, \quad [2.2]$$

for all $\lambda \in \Lambda$. Thus, $\psi_m^0 = \psi^*$ by Proposition 1.1.

Example 2.1 (Continuation of Examples 1.1 and 1.2): Example 1.2 establishes antisymmetry of the log-likelihood derivative for the symmetric parameterization of the exponential matched pairs problem. Consistency of $\hat{\psi}$ thus follows by Proposition 2.1 regardless of the true distribution and assumed model for the random effects γ . This considerably extends the result of ref. 14 and provides deeper insight into the structure of inference under model misspecification.

Remark 2.1: It is not too difficult to check following calculations similar to that in Appendix A.4 of ref. 14 that if we parameterize this model nonsymmetrically, for example, with rates $\gamma_i \theta$ and γ_i , the resulting estimator is not consistent if the gamma random effects assumption is erroneous.

Proposition 1.4 gives easily verifiable conditions for antisymmetry of the log-likelihood derivative, which holds for all location and scale families in Example 1.3. The following explicit calculation for the normal distribution mirrors Example 1.2.

Example 2.2 (normal matched pairs with arbitrary mixing):

Let Y_1 and Y_0 be normally distributed with unit variance and means $\gamma + \psi^*$ and $\gamma - \psi^*$, respectively, γ being treated as random. The conclusion is unchanged if an unknown variance parameter is present. Conditionally on γ , with $u_1 = y_1 - \psi^*$, $u_0 = y_0 + \psi^*$, the joint density function of (Y_1, Y_0) is

$$f_1(y_1; \psi^*, \gamma) f_0(y_0; \psi^*, \gamma) dy_1 dy_0 \\ = \exp[-\{(u_1 - \gamma)^2 + (u_0 - \gamma)^2\}/2] du_1 du_0,$$

which is symmetric in u_1 and u_0 by construction as a result of using a ψ -symmetric parameterization with respect to addition,

the relevant group action for location models. Conditionally on γ , the likelihood as a function of ψ is

$$L(\psi; y_0, y_1, \gamma) \propto \exp\{-\frac{1}{2}(y_1 - \gamma - \psi)^2 - \frac{1}{2}(y_0 - \gamma + \psi)^2\} \\ \propto \exp[-\frac{1}{2}\{(y_1 - \psi)^2 + (y_0 + \psi)^2\}], \quad [2.3]$$

where the proportionality symbol on the second line absorbs the multiplicative term $\exp\{-\gamma^2 + \gamma(y_1 + y_0)\}$, which does not depend on ψ . On taking logarithms and differentiating,

$$\nabla_{\psi} \ell(\psi; u_0, u_1) = y_1 - y_0 - 2\psi = u_1 - u_0 = -\nabla_{\psi} \ell(\psi; u_1, u_0),$$

verifying the antisymmetry condition of Definition 1.3.

For any assumed random effects distribution for γ with density function parameterized by λ , an application of Proposition 2.1 guarantees that the resulting maximum likelihood estimator $\hat{\psi}$ is consistent for ψ^* under arbitrary misspecification of the random effects distribution.

Remark 2.2: A separate point in connection with Example 2.2 is that, due to the separation in the log-likelihood function specific to this case (a parameter cut in the terminology of ref. 19), $\hat{\psi} = (Y_1 - Y_0)/2$ regardless of the assumed random effects distribution, and for an analysis from n pairs $\hat{\psi} = (S_1 - S_0)/2n$, where $S_j = \sum_{i=1}^n Y_{ij}$, where i is the pair index. For any finite-variance random effects distribution f , $\text{var}(Y_1) = \text{var}(Y_0) = 1 + \text{var}_f(\gamma)$ and $\text{cov}(Y_1, Y_0) = \text{var}_f(\gamma)$. Thus $\text{var}(\hat{\psi}) = 2/n$ irrespective of f and of the assumed (erroneous) model over the random effects.

Remark 2.3: In the context of stratum-specific random effects, Lindsay (20, Thm 4.2) established the limiting distribution of his proposed estimator of the interest parameter, valid whether or not the random effects distribution is misspecified, under the assumption that his estimator is consistent. The latter aspect, central to the present paper, was not discussed in that work.

2.B. Two-Group Problems. Section 2.A dealt with an experimental setting in which the matching ensures balance. A more realistic situation in an observational context is that observations on treated and untreated individuals are stratified into groups that are as similar as possible. This in general produces unbalanced strata, having a potentially different number of individuals in the treated and untreated groups for each stratum.

Inference is based on the sufficient statistics S_{j1} and S_{j0} within treatment groups and strata. Let $(Y_{ij1})_{i=1}^{r_{j1}}$ and $(Y_{ij0})_{i=1}^{r_{j0}}$ be observations within the j th stratum for treated and untreated individuals, respectively. Some examples fix ideas. If the observations $(Y_{ij1})_{i=1}^{r_{j1}}$ and $(Y_{ij0})_{i=1}^{r_{j0}}$ are normally distributed with means $\gamma_j + \psi^*$ and $\gamma_j - \psi^*$ and variance τ , the likelihood contribution to the j th stratum depends on the data only through $\sum_{i=1}^{r_{j1}} Y_{ij1}/r_{j1}$ and $\sum_{i=1}^{r_{j0}} Y_{ij0}/r_{j0}$, whose distributions are normal with means unchanged, and variances τ/r_{j1} and τ/r_{j0} , respectively. If, conditionally on γ_j , the individual observations are Poisson distributed counts of rates $\gamma_j\psi^*$ and γ_j/ψ^* , the sufficient statistics are sums of these counts, Poisson distributed of rates $r_{j1}\gamma_j\psi^*$ and $r_{j0}\gamma_j/\psi^*$. As a third example, if the originating variables are exponentially distributed of rates $\gamma_j\psi^*$ and γ_j/ψ^* , the sufficient statistics are gamma-distributed sums of shape and rate parameters $(r_{j1}, \gamma_j\psi^*)$ and $(r_{j0}, \gamma_j/\psi^*)$, respectively. The matched comparison setting of Section 2.A is a special case of this formulation with $r_{j1} = r_{j0} = 1$ for all j . The more general balanced case with $r_{j1} = r_{j0} = r$ is also implicitly covered by Propositions 1.4 and 2.1.

The situation when $r_{j1} \neq r_{j0}$ is more complicated as it implies that even if the distributions of S_{j1} and S_{j0} belong to a group family, the transformations required to express their distributions in the standardized form f_U depend on the different values of r_{j1} and r_{j0} , so that Definition 1.2 is violated. The following extension of Proposition 2.1, which assumes the same random effects formulation for the nuisance parameters γ_j , omits the subscripts for the strata as before.

Proposition 2.2. *Suppose that, conditionally on γ , S_1 and S_0 are independent random variables with probability measures in \mathcal{P}_G . Let $h_1 \in G$ and $h_0 \in G$, not depending on ψ or γ , be such that $T_1 = h_1 S_1$ and $T_0 = h_0 S_0$ have a joint distribution that is parameterized ψ -symmetrically in the sense of Definition 1.2 and such that the log-likelihood derivative, when expressed in terms of $u_1 = g^{-1}t_1$ and $u_0 = gt_0$, is antisymmetric in the sense of Definition 1.3. Then, provided that $\ell(\psi, \lambda)$ is strictly concave, $\psi^* \perp \Lambda$ and*

$$0 = \mathbb{E}_m[\nabla_{\psi} \ell(\psi^*, \lambda)] \\ = \int \int (\nabla_{\psi} \ell(\psi^*, \lambda; s_1, s_0)) m(s_1, s_0) ds_1 ds_0, \quad [2.4]$$

for all $\lambda \in \Lambda$. Thus, $\psi_m^0 = \psi^*$ by Proposition 1.1.

For discrete problems such as the two-group Poisson situation mentioned above, there is no natural group and Propositions 2.1 and 2.2 are therefore inapplicable. Direct verification of Proposition 1.2 is sometimes more fruitful, as illustrated in the following examples.

Example 2.3 (stratified two-group Poisson problem with unbalanced strata): Cox and Wong (16) considered misspecification of a random effects distribution in a stratified two-group problem. The distributions of counts S_{j1} and S_{j0} in the treated and untreated groups for stratum j are, conditionally on γ_j , Poisson with means $r_{j1}\gamma_j \exp(\theta^*)$ and $r_{j0}\gamma_j \exp(-\theta^*)$, respectively, where r_{j1} and r_{j0} reflect the number of patients at risk in each group and γ_j are stratum-specific nuisance parameters.

The problem can equivalently be parameterized in terms of $\psi = e^{\theta}$ but the above formulation was used by Cox and Wong (16). The result of Proposition 2.3 was noted there; we complete their proof in SI Appendix.

Proposition 2.3 [Cox and Wong (16)]. *Suppose that in Example 2.3, the nuisance parameters $(\gamma_j)_{j=1}^n$ are modeled as independent and identically distributed random variables with gamma density $\xi(\xi\gamma)^{\omega-1} \exp(-\xi\gamma)/\Gamma(\omega)$, parameterized in terms of the notionally orthogonal parameters $\nu = \xi/\omega$ and $\omega > 0$. Provided that there is a member of the assumed model that gives the same expectation $\mathbb{E}_{\nu, \omega}(\gamma_j) = 1/\nu$ as under the true random effects distribution for $(\gamma_j)_{j=1}^n$, $\hat{\theta}$ is consistent for θ^* .*

Example 2.4 (time to the first event in a two-group problem with unbalanced strata): We describe here the generating mechanism for an exponential analog of Example 2.3. Suppose the available data in Example 2.3 consist also of the failure times. If these are, for each individual, independent exponentially distributed of rates $\gamma_j \exp(\theta^*)$ and $\gamma_j \exp(-\theta^*)$ in the j th stratum, with r_{j1} and r_{j0} individuals at risk in the treated and untreated groups, then the times S_{j1} and S_{j0} to the first failure in each group are the group minima, exponentially distributed with rates $r_{j1}\gamma_j \exp(\theta^*)$ and $r_{j0}\gamma_j \exp(-\theta^*)$, respectively. By using only the minima across strata, information is sacrificed, which

would normally not be advised unless to check on modeling assumptions. After reparameterization, $\psi = e^\theta$, the problem satisfies the conditions of *Proposition 2.2* with $T_{j1} = r_{j1}S_{j1}$ and $T_{j0} = r_{j0}S_{j0}$. Thus, $\hat{\psi}$ is consistent for $\psi^* = \exp(\theta^*)$.

2.C. Generalized Linear Models. A simple example in which *Proposition 1.1* is directly applicable is misspecified dispersion in exponential-family regression problems. Suppose that the true density function belongs to the class of canonical exponential family regression models with unknown regression coefficient ψ on a set of known covariates x_1, \dots, x_n (each column vectors of the same dimension as ψ), and dispersion parameters ϕ_1, \dots, ϕ_n . The form of the exponential-family log-likelihood function with respect to ϕ_i is such that $\psi^* \perp_m \phi_i$ for all i . In particular this implies that if ϕ_i is modeled erroneously as $\phi_i(\lambda)$, perhaps depending on covariates other than or including x_i , $\psi^* \perp_m \Lambda$. Write the log-likelihood function of the assumed model as

$$\ell(\psi, \lambda) = \sum_{i=1}^n \phi_i^{-1}(\lambda) \{y_i x_i^T \psi - K(x_i^T \psi)\} + \sum_{i=1}^n b\{y_i; \phi_i(\lambda)\}. \quad [2.5]$$

We do not distinguish notationally between realizations of the random variables Y_1, \dots, Y_n and arbitrary evaluation points for their joint density function conditional on covariates. With outcomes treated as random, the estimating equation for ψ is unbiased in the sense that $\mathbb{E}_m[\nabla_\psi \ell(\psi^*, \lambda)] = 0$ for all $\lambda \in \Lambda$. Thus, $\hat{\psi}$ is consistent in the presence of arbitrary misspecification of the dispersion parameters.

In a continuation of this generalized linear models example, we assume this time a simplified specification for dispersion but allow for relevant covariates to be omitted. Let $Z = (X, W)$ denote the included and omitted covariates. The conditional probability density or mass function of $Y = (Y_1, \dots, Y_n)$ at $y = (y_1, \dots, y_n)$ is, with $\theta^* = (\psi^{*T}, \lambda^{*T})^T$.

$$m(y; z_1^T \theta^*, \dots, z_n^T \theta^*) = \exp \left[\phi^{-1} \left\{ \theta^{*T} \sum_{i=1}^n z_i y_i - \sum_{i=1}^n K(z_i^T \theta^*) \right\} \right] \prod_{i=1}^n b(y_i, \phi),$$

Let $\check{m}(y; x_1^T \psi, \dots, x_n^T \psi)$ be the assumed model, in which the outcome is modeled only in terms of X . The orthogonality condition $\psi^* \perp_m \Lambda$ is

$$\sum_{i=1}^n K''(x_i^T \psi^* + w_i^T \lambda) x_i w_i^T = 0, \quad [2.6]$$

for all λ , a highly restrictive condition. By *Proposition 1.1*, the additional conditions under which $\psi_m^0 = \psi^*$ are those that set $\mathbb{E}_m[\nabla_\psi \ell(\psi^*, \lambda)] = 0$ for all λ , where ℓ is the log-likelihood function under the assumed model. Since λ is not present under the assumed model,

$$\begin{aligned} \mathbb{E}_m[\nabla_\psi \ell(\psi^*, \lambda)] &= \sum_{i=1}^n \{K'(z_i^T \theta^*) - K'(x_i^T \psi^*)\} x_i \\ &= \sum_{i=1}^n K''(x_i^T \psi^* + \bar{\xi}_i) x_i w_i^T \lambda^*, \end{aligned}$$

where $\bar{\xi}_i$ is on a line between 0 and $w_i^T \lambda^*$ and the first equality is because $K'(z_i^T \theta^*) = \mathbb{E}_m(Y_i | Z_i = z_i)$. Since

$K''(x_i^T \psi^* + \bar{\xi}_i) \neq K''(x_i^T \psi^* + w_i^T \lambda)$ in general, the second condition $\mathbb{E}_m[\nabla_\psi \ell(\psi^*, \lambda)] = 0$ typically does not hold exactly even when the columns of X and W are orthogonal. An exception is the normal-theory linear model, for which $K(\zeta) = \zeta^2/2$ so that $K''(\zeta) = 1$. For an insightful discussion of noncollapsibility in logistic and other regression models, see ref. 21, and for some relevant approximations, see ref. 22.

Although both conclusions of this subsection are well established in the literature, their purpose is to illustrate application of *Proposition 1.1* to some well-known settings.

2.D. Marginal Structural Models. Evans and Didelez (8) presented a further example in which consistency arises as a special case of *Proposition 1.1*. Their paper introduced the idea of a frugal parameterization of a marginal structural model, constructed to complete the specification of a model respecting the interventional components of interest, but without loss or redundancy. In a causal system of the form, e.g., $Z \rightarrow X, Z \rightarrow Y, X \rightarrow Y$, the authors suggest a reparameterization of the joint distribution p_{ZXY} from $(p_Z, p_{X|Z}, p_{Y|ZX})$ to $(p_Z, p_{X|Z}, p_{Y|X}^*, \phi_{YZ|X}^*)$, where $p_{Y|X}^*$, with finite-dimensional parameter $\theta_{Y|X}^*$, is typically the unobservable interventional distribution of interest for $Y | \text{do}(X)$, and $\phi_{YZ|X}^*$ is whatever is necessary to complete the joint model. The quantity $p_{X|Z}$ is known as the propensity score. We refer to the original paper for a detailed discussion of these quantities, which in turn are expressed in terms of finite-dimensional parameters, say ψ and λ . The vector parameter ψ with true value ψ^* comprises $\phi_{YZ|X}^*$ and the main component of interest $\theta_{Y|X}^*$ which together specify $p_Z p_{Y|ZX}$, while λ characterizes the propensity score $p_{X|Z}$. A consequence of the above parameterization within the class of marginal structural models is Theorem 5.1 of ref. 8, which establishes consistency and asymptotic normality of the maximum likelihood estimator $\hat{\psi}$ of ψ^* , even if $p_{X|Z} = p_{X|Z}(\lambda)$ is misspecified. As noted in ref. 8, the conclusion arises as a consequence of the parameter cut between $p_Z p_{Y|ZX}$ and $p_{X|Z}$, and therefore between ψ and λ . The parameter cut implies the parameter orthogonality condition $\psi^* \perp_m \Lambda$, and Evans and Didelez (8) implicitly establish in their proof the remaining condition of *Proposition 1.1*.

3. Discussion

3.A. Stronger Inferential Guarantees. Even if consistency holds, inferential guarantees entail estimation of variance, a difficult problem due to failure of Bartlett's second identity. Such difficulties persist under the stronger condition of *Proposition 1.3*.

In the notation of Section 1, let

$$q = \mathbb{E}_m[(\nabla_{(\psi, \lambda)} \ell)(\nabla_{(\psi, \lambda)} \ell)^T],$$

$$\check{q} = \mathbb{E}_{(\psi, \lambda)}[(\nabla_{(\psi, \lambda)} \ell)(\nabla_{(\psi, \lambda)} \ell)^T],$$

both functions of ψ and λ . *Proposition 1.3*, Eq. 1.1 and consistency of $\hat{\psi}$ imply

$$\mathbb{E}_{(\psi, \lambda)}[\nabla_{(\psi, \lambda)} \ell(\psi^*, \lambda_m^0)] = 0,$$

and differentiation gives Bartlett's second identity under the assumed model:

$$-\check{r}(\psi^*, \lambda_m^0) + \check{q}(\psi^*, \lambda_m^0) = 0.$$

The relevant quantities for establishing the limiting distributions of standard test statistics are, however, i and q , where i

was defined in Section 1. These do not in general satisfy $-i(\psi^*, \lambda_m^0) + q(\psi^*, \lambda_m^0) = 0$, even with the additional structure of Proposition 1.3. An exception is Example 2.2 because of the particular simplicity of this case.

It can be shown using now standard arguments that, with the log-likelihood function constructed from n independent observations, provided that $\hat{\psi}$ is consistent for ψ^* ,

$$\sqrt{n}(\hat{\psi} - \psi^*) \rightarrow_d N(0, (i^{-1}qi^{-1})_{\psi\psi}), \quad [3.1]$$

the variance being the so-called sandwich formula, in which the constituent terms are evaluated at (ψ^*, λ_m^0) . This reduces to $i^{\psi\psi}(\psi^*, \lambda_m^0) = i^{\psi\psi}(\psi^*, \lambda^*)$ when the model is correctly specified, so that the usual asymptotic result is recovered. The standard refs. 3, 5–7, and 23 have ψ_m^0 in place of ψ^* in Eq. 3.1. Qualitatively similar results hold for the profile score and profile likelihood ratio statistics, the main point being that confidence set estimation is infeasible without knowledge of m unless $\check{q}(\psi^*, \lambda_m^0) = q(\psi^*, \lambda_m^0)$ to some adequate order of approximation. In general, this requires, not only that the expectations of the sufficient statistics are robust to misspecification but also that the expectations of any relevant squares and cross terms are stable.

Suppose, however, that the parameterization is orthogonal in the sense $\psi^* \perp_m \lambda_m^0$. Then,

$$(i^{-1}qi^{-1})_{\psi\psi} = i^{\psi\psi} q_{\psi\psi} i^{\psi\psi},$$

and the familiar result is also recovered if $q_{\psi\psi} i^{\psi\psi} = I_{\dim(\psi)}$, a weaker requirement than $q = i$ especially if ψ is a scalar. Under the additional structure of Proposition 1.3, inference based on the assumed model is unaffected by the misspecification.

3.B. Treating Incidental Parameters As Fixed or Random. In the context of Section 2.A, the lack of general inferential guarantees beyond consistency may well be an argument for treating the nuisance parameters as fixed. Although the conceptual distinction is consequential for the analysis, a formulation in which incidental parameters are treated as fixed and arbitrary, and one in which they are treated as independent and identically distributed from a totally unspecified distribution, are numerically indistinguishable. In this sense, the fixed-parameter formulation is essentially nonparametric for the nuisance component, with the distinguishing feature that it evades estimation of the infinite-dimensional nuisance parameter.

The group structure used to define the symmetric parameterization in Proposition 2.1 also allows elimination of nuisance parameters by suitable preliminary maneuvers when the nuisance parameters are treated as fixed. In the exponential matched pairs setting of Example 2.1, (14) compared inference based on the distribution of Y_{i1}/Y_{i0} to that based on modeling the pair effects by a gamma distribution. The random pair effects model is more efficient when the gamma distribution is correct, but efficiency degrades substantially under misspecification, compared to maximum likelihood estimation based on the distribution of ratios.

Proposition 2.1 hinges on a symmetric parameterization having been chosen, which is possible only under the group structure of Section 1.C. Outside of this setting, there are no guarantees of consistency of the maximum likelihood estimator with misspecified random effects distribution, while it may still be possible to eliminate the pair effects with exact or approximate conditioning arguments, as illustrated in the second example of ref. 14.

3.C. Overstratification and Other Encompassing Models. Two essentially equivalent ways to adjust for potential confounders are to stratify by observed explanatory features, leading to strata effects $(\gamma_j)_{j=1}^m$ as in Example 2.3, or to adjust for the explanatory features in a regression analysis. Both approaches are a form of conditioning by model formulation. It is arguably more relevant in this context to treat the strata effects as fixed rather than random. In De Stavola and Cox (24), this modification of Example 2.3 was considered, showing the extent to which overstratification decreases efficiency of the estimator. As is intuitively clear, no bias is incurred through overstratification, yet the conditions of Propositions 1.1 and 1.2 do not always hold, as illustrated by the following example. The explanation is that the true density is nested within the encompassing model, so that $\lambda_m^0 = \lambda^* = 0$. The overstratified model is therefore not misspecified according to the definition below Eq. 1.1.

To isolate the point at issue, we assume that any dispersion parameters are equal to 1, the conclusion of Proposition 3.1 being unchanged for arbitrary dispersion parameter by the discussion of Section 2.C.

Proposition 3.1. *Suppose that, conditional on observed explanatory features, the probability density or mass function of $Y = (Y_1, \dots, Y_n)$ at $y = (y_1, \dots, y_n)$ is of canonical exponential family regression form, that is,*

$$\begin{aligned} m(y; x_1^T \psi^*, \dots, x_n^T \psi^*) \\ = \exp \left\{ \sum_{i=1}^n y_i x_i^T \psi^* - \sum_{i=1}^n K(x_i^T \psi^*) \right\} \prod_{i=1}^n b(y_i). \end{aligned}$$

Suppose further that the analysis is overstratified so that under the assumed model the log-likelihood function is $\ell(\psi, \lambda) = s^T \psi + t^T \lambda - K(\psi, \lambda)$, where $s = \sum_{i=1}^n x_i y_i$, $t = \sum_{i=1}^n w_i y_i$ and $K(\psi, \lambda) = \sum_{i=1}^n K(x_i^T \psi + w_i^T \lambda)$. The conditions of Propositions 1.1 and 1.2 are in general violated, yet the maximum likelihood estimator $\hat{\psi}$ is consistent for ψ^ .*

Another type of encompassing model arises in random effects models when the assumed family of distributions for the random effects is rich enough to include the true distribution, for instance if the postulated mixture class is nonparametric [Kiefer and Wolfowitz (25)], as discussed in Section 3.B. This type of situation can be understood in terms of the convex geometry of mixture distributions [Lindsay (26)] and was specialized to the case of binary matched pairs with logistic probabilities by Neuhaus et al. (27). Their conditions on the mixing distributions with regard to the resulting cell probabilities effectively imply that the model is correctly specified according to the definition below Eq. 1.1.

3.D. Estimating Equations and Neyman Orthogonality. A generalization of the score Eq. 1.1 is to other estimating equations, often but not always obtained as the vector of partial derivatives of a convex loss function. The model need only be partially specified in terms of a known function b of the data, an interest parameter ψ and a nuisance parameter λ , such that the expectation when the assumed model is true satisfies $\mathbb{E}_m[b(\psi^*, \lambda^*; Y)] = \mathbb{E}_{(\psi^*, \lambda^*)}[b(\psi^*, \lambda^*; Y)] = 0$. If the model is misspecified the limiting solutions satisfy, in analogy with Eq. 1.1,

$$\mathbb{E}_m[b(\psi_m^0, \lambda_m^0; Y)] = 0.$$

A special choice of b , establishing a connection to the double robustness literature, is the Neyman-orthogonal score for ψ , defined as

$$s_N^*(\psi; \lambda, Y) = \ell_\psi(\psi, \lambda) - w^{*\top} \ell_\lambda(\psi, \lambda),$$

where ℓ_ψ and ℓ_λ are the partial derivatives of ℓ with respect to ψ and λ and $w^{*\top}$ is a matrix of dimension $\dim(\psi) \times \dim(\lambda)$ given by $w^{*\top} = i_{\psi\lambda}^* i_{\lambda\lambda}^{*-1}$ when the model is correctly specified. Here, $i_{\psi\lambda}^*$ and $i_{\lambda\lambda}^*$ are the components of the Fisher information matrix at the true parameter values (ψ^*, λ^*) . Consider instead

$$s_N(\psi; \lambda, Y) = \ell_\psi(\psi, \lambda) - w^\top \ell_\lambda(\psi, \lambda), \quad w^\top = i_{\psi\lambda} i_{\lambda\lambda}^{-1}, \quad [3.2]$$

where $i_{\psi\lambda}$ and $i_{\lambda\lambda}$ are as defined in Eq. 1.4. Write the condition $i^{\psi\psi} g_\psi + i^{\psi\lambda} g_\lambda = 0$ of Proposition 1.2 as

$$g_\psi + i_{\psi\psi} i^{\psi\lambda} g_\lambda = 0, \\ i_{\psi\psi} = (i^{\psi\psi})^{-1} = i_{\psi\psi} - i_{\psi\lambda} i_{\lambda\lambda}^{-1} i_{\lambda\psi}, \quad [3.3]$$

where $i_{\psi\psi}$ is the Fisher information for ψ at (ψ^*, λ) , computed under the true model, having adjusted for estimation of λ . On noting that $i^{\psi\lambda} = -i^{\psi\psi} i_{\psi\lambda} i_{\lambda\lambda}^{-1}$, we see that the information identity $i^{\psi\psi} g_\psi + i^{\psi\lambda} g_\lambda = 0$ of Proposition 1.2 is equivalent to requiring that the Neyman orthogonal score Eq. 3.2, with score adjustment w computed under the true distribution, has zero expectation under the same distribution for all $\lambda \in \Lambda$. Specification of w in Eq. 3.2 when the model is misspecified thus hinges on the conditions of Proposition 1.3 being satisfied.

Data, Materials, and Software Availability. There are no data underlying this work.

ACKNOWLEDGMENTS. This research was partially supported by the UK Engineering and Physical Sciences Research Council and the Natural Sciences and Engineering Research Council of Canada. We are grateful to the reviewers for their very helpful comments.

1. L. Breiman, Statistical modeling: The two cultures. *Stat. Sci.* **16**, 199–215 (2001).
2. D. R. Cox, Comment on Breiman's "Statistical modeling: The two cultures". *Stat. Sci.* **16**, 216–218 (2001).
3. D. R. Cox, "Tests of separate families of hypotheses" in *Proceedings of the Fourth Berkeley Symposium*, J. Neyman, Ed. (University of California Press, Berkeley, CA, 1961), vol. I, pp. 105–123.
4. P. J. Huber, "The behavior of maximum likelihood estimates under nonstandard conditions" in *Proceedings of the Fifth Berkeley Symposium*, L. M. Le Cam, J. Neyman, Eds. (University of California Press, Berkeley, CA, 1967), vol. I, pp. 221–233.
5. J. Kent, Robust properties of likelihood ratio tests. *Biometrika* **69**, 19–27 (1982).
6. H. White, Maximum likelihood estimation of misspecified models. *Econometrica* **50**, 1–25 (1982).
7. H. White, Regularity conditions for Cox's test of non-nested hypotheses. *J. Econ.* **19**, 301–318 (1982).
8. R. J. Evans, V. Didelez, Parameterizing and simulating from causal models (with discussion). *J. R. Stat. Soc. Ser. B* **86**, 535–568 (2024).
9. J. M. Robins, A. Rotnitzky, L. P. Zhao, Estimation of regression coefficients when some regressors are not always observed. *J. Am. Stat. Assoc.* **89**, 846–866 (1994).
10. V. Chernozhukov *et al.*, Double/debiased machine learning for treatment and structural parameters. *Econ. J.* **21**, 1–68 (2018).
11. S. Vansteelandt, O. Dukes, Assumption-lean inference for generalised linear model parameters (with discussion). *J. Roy. Stat. Soc. Ser. B* **84**, 657–685 (2022).
12. D. R. Cox, "Statistical causality: some historical remarks" in *Causality: Statistical perspectives and applications*, C. Berzuini, A. P. Dawid, L. Bernardinelli, Eds. (Wiley, Chichester, 2012), pp. 1–5.
13. H. S. Battey, Heather battey's contribution to the discussion of "Assumption-lean inference for generalised linear model parameters" by Vansteelandt and Dukes. *J. R. Stat. Soc. Ser. B* **84**, 696–698 (2022).
14. H. S. Battey, D. R. Cox, High dimensional nuisance parameters: An example from parametric survival analysis. *Inf. Geom.* **3**, 119–148 (2020).
15. H. Schielzeth *et al.*, Robustness of linear mixed-effects models to violations of distributional assumptions. *Methods Ecol. Evol.* **11**, 1141–1152 (2020).
16. D. R. Cox, M. Y. Wong, A note on the sensitivity of assumptions of a generalized linear mixed model. *Biometrika* **97**, 209–214 (2010).
17. D. R. Cox, N. Reid, Parameter orthogonality and approximate conditional inference (with discussion). *J. R. Stat. Soc. Ser. B* **49**, 1–39 (1987).
18. O. E. Barndorff-Nielsen, D. R. Cox, *Inference and Asymptotics* (Chapman & Hall, London, 1994).
19. O. E. Barndorff-Nielsen, *Information and Exponential Families in Statistical Theory* (Wiley, New York, 1978).
20. B. G. Lindsay, Using empirical partially Bayes inference for increased efficiency. *Ann. Stat.* **13**, 914–931 (1985).
21. R. Daniel, J. Zhang, D. Farewell, Making apples from oranges: Comparing noncollapsible effect estimators and their standard errors after adjustment for different covariate sets. *Biom. J.* **63**, 528–557 (2021).
22. D. R. Cox, N. Wermuth, An approximation to maximum likelihood estimates in reduced models. *Biometrika* **77**, 747–761 (1990).
23. D. R. Cox, Further results on tests of separate families of hypotheses. *J. R. Stat. Soc. Ser. B* **24**, 406–424 (1962).
24. B. L. de Stavola, D. R. Cox, On the consequences of overstratification. *Biometrika* **95**, 992–996 (2008).
25. J. G. Kiefer, J. Wolfowitz, Consistency of the maximum likelihood estimator in the presence of many incidental parameters. *Ann. Math. Stat.* **27**, 887–906 (1983).
26. B. G. Lindsay, The geometry of mixture likelihoods: General theory. *Ann. Stat.* **11**, 86–94 (1983).
27. J. M. Neuhaus, J. D. Kalbfleisch, W. W. Hauck, Conditions for consistent estimation in mixed-effects models for binary matched-pairs data. *Can. J. Stat.* **22**, 139–148 (1994).