

# Parameter Restrictions for the Sake of Identification: Is There Utility in Asserting That Perhaps a Restriction Holds?

Paul Gustafson

*Abstract.* Statistical modeling can involve a tension between assumptions and statistical identification. The law of the observable data may not uniquely determine the value of a target parameter without invoking a key assumption, and, while plausible, this assumption may not be obviously true in the scientific context at hand. Moreover, there are many instances of key assumptions which are untestable, hence we cannot rely on the data to resolve the question of whether the target is legitimately identified. Working in the Bayesian paradigm, we consider the grey zone of situations where a key assumption, in the form of a parameter space restriction, is scientifically reasonable but not incontrovertible for the problem being tackled. Specifically, we investigate statistical properties that ensue if we structure a prior distribution to assert that *maybe* or *perhaps* the assumption holds. Technically this simply devolves to using a mixture prior distribution putting just some prior weight on the assumption, or one of several assumptions, holding. However, while the construct is straightforward, there is very little literature discussing situations where Bayesian model averaging is employed across a mix of fully identified and partially identified models.

*Key words and phrases:* Bayesian model averaging, Bayes risk, large-sample theory, partial identification.

## 1. INTRODUCTION

### 1.1 Background

A tension can arise in applied statistical modeling. To fully identify a parameter of interest, “strong” model assumptions may be required. Of course it is not prudent to invoke an assumption without a solid rationale. Without identification, however, an uncomfortable truth ensues: even an infinite amount of data would not reveal the true value of the target parameter. (Though the data may help some. We shall see situations of *partial* or *set* identification, where an infinite amount of data could rule out *some a priori* plausible values of the target.) To further muddy the waters, often an identifying assumption is not testable empirically. So we cannot necessarily rely on the data to resolve the situation.

Consider a context where an identifying assumption or parameter *restriction* is scientifically plausible but not incontrovertible. One possibility, at least within the

Bayesian paradigm, is to specify a prior distribution asserting that “large” violations of the restriction are unlikely. For instance, say the restriction is  $\lambda = 0$ . A prior  $\lambda \sim N(0, \tau^2)$ , for a suitably small choice of  $\tau$ , could encode such information. Some work in this spirit includes Scharfstein, Daniels and Robins (2003), Gustafson and Greenland (2006), Gustafson (2007), Keele and Quinn (2017), Franks, D’Amour and Feller (2020). “Bayesian-like” approaches, often referred to as *probabilistic bias analysis*, have also been considered (see, for instance, Greenland, 2003, 2005, Lash et al., 2014, Fox, MacLehose and Lash, 2022).

Assigning a prior distribution which probabilistically limits the magnitude of violation for an identifying restriction is not the only way to proceed. Arguably a more direct encoding of “plausible but not incontrovertible” results from a mixture prior distribution, giving some weight to the restriction being met exactly, and the remaining weight to it being violated (without necessarily making a stringent judgement about the magnitude of violation). While mixture prior distributions are ubiquitous in Bayesian hypothesis testing and model selection proce-

---

Paul Gustafson is Professor, Department of Statistics, University of British Columbia, Vancouver, Canada (e-mail: [gustaf@stat.ubc.ca](mailto:gustaf@stat.ubc.ca)).

dures, there is scant literature on their use in the context of identifying restrictions.

## 1.2 Motivating Example: Prevalence Estimation with Missing Data

Consider the HIV surveillance study described by Verstraeten et al. (1998), which also motivated the methodological developments in Vansteelandt et al. (2006). Blood draws were taken from a sample of  $n = 787$  members of the target population. Let  $Y$  indicate the HIV test result (0 for negative, 1 for positive), let  $R$  indicate the observation of the result (1 for observed, 0 for missing), and let  $p_{ry} = \Pr(R = r, Y = y)$ . Hence, the target parameter, HIV prevalence, is  $\psi = \Pr(Y = 1) = p_{01} + p_{11}$ . The study data are summarized by  $c_{10} = 699$  negative tests ( $R = 1, Y = 0$ ),  $c_{11} = 52$  positive tests ( $R = 1, Y = 1$ ), and  $c_{0+} = 36$  missing test results ( $R = 0$ ).

Consider inferring  $\psi$  without invoking any identifying restriction, that is, allowing there may be nonignorable missingness (NIM). (See Daniels and Hogan (2008) or Little and Rubin (2019) for full discussions of NIM.) A weakly informative prior specification is  $(p_{00}, p_{01}, p_{10}, p_{11}) \sim \text{Dirichlet}(1, 1, 1, 1)$ , that is, a uniform distribution. By reparameterizing to  $(s, p_{0+}, p_{10}, p_{11})$ , where  $p_{0+} = p_{00} + p_{01}$  and  $s = p_{01}/(p_{00} + p_{01})$ , the posterior distribution is characterized by  $(p_{0+}, p_{10}, p_{11}) \sim \text{Dirichlet}(2 + c_{0+}, 1 + c_{10}, 1 + c_{11})$  and independently  $s \sim \text{Unif}(0, 1)$ . This induces the marginal posterior distribution on the target  $\psi = sp_{0+} + p_{11}$ . The impact of not having an identifying restriction is clear. Since a posteriori  $s \sim \text{Unif}(0, 1)$  for any dataset, the posterior distribution of  $\psi$  will not reduce to a point-mass in the infinite limit of further data collection. For the data at hand, the posterior distribution of  $\psi$  is depicted in Figure 1. Note that R code to reproduce this figure, and indeed all the empirical work that follows, is available as Gustafson (2023) or via <https://github.com/paulgstf/param-restrict-maybe>.

Alternately, say we believe the missing-at-random (MAR) assumption is justified, that is, presuming  $R$  to

be independent of  $Y$ . A weakly informative prior specification is to let  $p_{ry} = (1 - \gamma)^{1-r} \gamma^r (1 - \psi)^{1-y} \psi^y$ , with  $\gamma = \Pr(R = 1)$  and  $\psi = \Pr(Y = 1)$  independently and identically distributed as  $\text{Unif}(0, 1)$ . This yields the marginal posterior distribution of the target  $\psi$  as simply  $\text{beta}(1 + c_{11}, 1 + c_{10})$ . Clearly this posterior would concentrate to a point-mass in the infinite limit of further data collection. For the data at hand, this posterior is also depicted in Figure 1.

An investigator believing that the MAR assumption *might* hold could apply *Bayesian model averaging* (BMA) (see, for instance, Kass and Raftery, 1995, Hoeting et al., 1999, Wasserman, 2000). Applied here, the prior is simply taken as a mixture of the two specifications above. The conjugate prior specifications give the Bayes factor contrasting the two specifications (with MAR in the numerator versus NIM in the denominator) as:

$$(1) \quad b = \frac{(n+2)(n+3)}{6(c_{0+}+1)(n-c_{0+}+1)},$$

which evaluates to  $b = 3.73$  for the present data. And the posterior odds favoring MAR are  $b$  times the prior odds. For instance, if the mixture prior gives equal weights of 0.5 to each of the NIM and MAR specifications, then for the present data the mixture posterior gives weight 0.211 to the NIM posterior and weight 0.789 to the MAR posterior. This model-averaged posterior for the target is also depicted in Figure 1.

With the NIM and MAR prior specifications being only weakly informative, a Bayes factor of nearly four might seem surprising, as MAR is well known to be an untestable assumption. Similarly, the strong dependence of (1) on the proportion of data missing is curious. However, oftentimes Bayes factors are sensitive to within-model prior specifications. In Appendix A, we elaborate on this point.

The findings above beg more general questions. While there is a rich literature on BMA, averaging across models which do and don't identify the target parameter has not been addressed. Thus we seek to understand statistical performance when model averaging is used to declare the *a priori* supposition that *perhaps* an identifying restriction holds.

## 2. THEORY AND METHODS

### 2.1 Framework

Let  $M_0$  be the “unrestricted” statistical model at hand, which lacks full identification. The initial declaration of  $M_0$  is presumed to use parameters which are as scientifically meaningful as possible. In Bayesian terms, it is likely natural to specify a prior distribution in terms of this initial “scientific” parameterization. We presume, however, that  $M_0$  can then be reparameterized using what Gustafson (2015) refers to as a *transparent* parameteri-




FIG. 1. Posterior distribution of HIV prevalence without an identifying restriction (allowing nonignorable missingness, NIM), with the missing-at-random (MAR) restriction, and the Bayesian model averaged (BMA) synthesis of the two distributions.

zation. This yields parameters  $\theta = (\phi, \lambda)$ , such that the distribution of the observable data  $D$  depends on  $\theta$  only through  $\phi$ . Or, more directly in Bayesian terms,  $D$  is conditionally independent of  $\lambda$ , given  $\phi$ . The lack of full identification is thus clear: the data inform  $\phi$  (indeed we presume fully, in the sense that the  $(D|\phi)$  model supports consistent estimation of  $\phi$ ). However, for any dataset, the posterior conditional distribution of  $(\lambda|\phi)$  is the same as the prior conditional distribution. To connect with the motivating example of Section 1.2,  $(p_{00}, p_{01}, p_{10}, p_{11})$  is a scientifically interpretable parameterization, but, to understand the lack of identification,  $\phi = (p_{0+}, p_{10}, p_{11})$  and  $\lambda = s = p_{01}/p_{0+}$  is a transparent parameterization. Moreover, the Dirichlet(1, 1, 1, 1) prior distribution is declared in the scientific parameterization, hence the prior distribution on  $\theta = (\phi, \lambda)$  follows by change-of-variables.

In what follows, the primary (scalar) target of inference is expressed as  $\psi = g(\phi, \lambda)$ . Provided that  $g()$  varies non-trivially with  $\lambda$ , the target is not fully identified. (This is the case in the motivating example, with  $\psi = sp_{0+} + p_{11}$ .)

Let  $\theta \in \Theta_0$  denote the  $M_0$  parameter space, with the marginal spaces thus implied being  $\phi \in \Phi_0$  and  $\lambda \in \Lambda_0$ . Importantly, many partially identified models arising in practice are such that  $\phi$  and  $\lambda$  are *not* variation independent, that is,  $\Theta_0$  is *not* the Cartesian product space of  $\Phi_0$  and  $\Lambda_0$ . Or, more plainly, the possible values of  $\lambda$  depend on  $\phi$ , and vice-versa. This can arise in the move from a scientific parameterization (likely with a Cartesian parameter space) to a transparent parameterization. In such situations, direct learning about  $\phi$  from the observed data may induce some indirect learning about  $\lambda$ , via the support of  $\lambda$  depending on  $\phi$ . In non-Bayesian settings, this manifests as bounds on  $\lambda$  (and in turn bounds on the target parameter) that are functions of  $\phi$ , and therefore are fully estimable (see, for instance, Manski, 2003). In the Bayesian setting, this manifests as necessarily having prior dependence between  $\phi$  and  $\lambda$ , so that the posterior marginal distribution of  $\lambda$  can depend on the data, even when the posterior conditional distribution of  $(\lambda|\phi)$  does not. This situation is described at length in Gustafson (2015).

We consider one or more (in general  $J$ ) sub-models or restrictions of  $M_0$ , each of which (i), has some level of *a priori* scientific credence, and (ii), identifies the target. We label the  $j$ -th such sub-model as  $M_j$ , for  $j = 1, \dots, J$ , and can regard the restriction as  $\theta \in \Theta_j \subset \Theta_0$ . Since the sub-model identifies the target, over  $\Theta_j$ ,  $g(\phi, \lambda)$  does not vary with  $\lambda$ . More pragmatically, if  $M_j$  holds, then  $\phi$ , which is estimable, uniquely determines the target  $\psi$ .

As already mentioned, the BMA framework lets the user postulate that *perhaps* one of the identifying restrictions holds. Let  $\pi_j()$  be the  $j$ -th prior density for  $\theta$  over  $\Theta_0$ . This will generally be a “degenerate” density for  $j \geq 1$ , since each identifying restriction forces  $\theta$  onto a lower-dimensional subset of  $\Theta_0$ . Fortunately, the BMA

framework is particularly well-suited to fairly compromising between a larger model and a nested sub-model (see, for instance, Kass and Raftery, 1995). Indeed, the conceptual symmetry with which a smaller and larger model are treated in a Bayes factor is generally a point of appeal.

The BMA prior distribution is a mixture, which we can represent as:

$$(2) \quad \pi_{\text{MIX}}(\theta) = \sum_{j=0}^J w_j \pi_j(\theta),$$

where  $w_j = \text{Pr}(M_j)$  is the prior probability that model  $M_j$  is correct, hence  $\sum_{j=0}^J w_j = 1$ . To be clear,  $(w_1, \dots, w_J)$  are hyperparameters which must be specified by the user. For instance, if  $J = 1$ , then a setting of  $w = (0.5, 0.5)$  gives equal credence to the identifying restriction holding or not holding.

Upon receipt of data  $D$ , standard Bayesian updating from the mixture prior (2) gives the posterior distribution as a mixture:

$$\pi_{\text{MIX}}(\theta|D) = \sum_{j=0}^J \tilde{w}_j(D) \pi_j(\theta|D).$$

Here  $\pi_j(\theta|D)$  is the standard “within-model” posterior distribution of parameters, that is, based on  $\pi_j(\theta|D) \propto f(D|\theta)\pi_j(\theta)$ . And  $\tilde{w}_j(D)$  is the posterior probability that  $M_j$  is correct. In the usual fashion,

$$\tilde{w}_j(D) = \frac{w_j f_j(D)}{\sum_{k=0}^J w_k f_k(D)},$$

where  $f_j(D) = \int_{\Theta_0} f(D|\theta)\pi_j(\theta) d\theta$  is the marginal density of the data under the  $j$ -th model.

While the framework above is standard, its implications are unexplored when the list of candidate models is a mix of partially and fully identified models. To large extent, the ensuing statistical behavior will be seen to be specific to the particular choice of partially identified model and identified sub-models, and to the nature of the specified within-model priors. However, to foreshadow the examples which follow, there are situations where:

- *Scenario 1.* For every dataset,  $\tilde{w}_1(D)/\tilde{w}_0(D) = w_1/w_0$ , that is, the data have zero ability to discriminate between the base model and an identified sub-model.
- *Scenario 2.* Let  $D_n$  denote a dataset of  $n$  observations arising under a particular parameter value  $\theta = \theta^\dagger$ , and let  $w_1^*(\theta^\dagger)$  be the large-sample limit of  $\tilde{w}_1(D_n)$ . It can happen that the restriction is fully testable, in that  $w_1^*(\theta^\dagger) = I_{\Theta_1}(\theta^\dagger)$ , that is, in the large-sample limit one correctly proves or refutes that the restriction holds.
- *Scenario 3.* It can happen that the range of  $w_1^*(\cdot)$  over  $\Theta_0$  is an interval  $[c, 1)$ , for some  $c \in (0, w_1)$ . Since the limiting posterior weight on  $M_1$  is never zero or one,

the restriction can never be fully refuted or fully supported. Also, since  $c$  is positive, there are not parameter values under which we get arbitrarily close to complete refutation of the restriction. In the other direction, however, there are parameter values under which we get arbitrarily close to full support for the restriction.

- *Scenario 4.* For data generated under some  $\theta^\dagger$  values in  $\Theta_0 - \Theta_1$ ,  $w_1^*(\theta^\dagger) = 0$ , while for other such values  $w_1^*(\theta^\dagger) \in (0, 1)$ . That is, a falsely asserted restriction may or may not be fully refuted.

Thus we find a richness in the variety of behaviors that can be encountered. Indeed, Scenarios 3 and 4 might surprise some, since they point to nuance beyond a binary construct of untestable versus testable restrictions.

### 2.2 Elucidation of Structure

Towards understanding the posterior model weights, note that, for  $j = 0, \dots, J$ ,

$$\begin{aligned}
 f_j(D) &= \int_{\Theta_0} f(D|\theta)\pi_j(\theta) d\theta \\
 (3) \quad &= \int_{\Theta_0} f(D|\phi)\pi_j(\phi, \lambda) d\lambda d\phi \\
 &= \int_{\Phi_0} f(D|\phi)\pi_j(\phi) d\phi.
 \end{aligned}$$

The marginal prior distribution of  $\phi$  thus plays a key role.

As was introduced informally in Section 2.1, let  $D_n$  be  $n$  independent datapoints generated under true parameter values  $(\phi, \lambda) = (\phi^\dagger, \lambda^\dagger)$ . (When helpful, we use the ‘dagger’ notation to emphasize specific, fixed parameter values giving rise to the observable data.) Importantly, note there is a true value of  $\lambda$ , even though this doesn’t influence the distribution of the observable data. We presume sufficient regularity for  $(D_n|\phi)$  such that standard asymptotic behavior based on Fisher information applies to maximum likelihood estimation of  $\phi$  (see, for instance, Casella and Berger, 2002). Additionally, we presume sufficient smoothness of marginal prior densities  $\pi_j(\phi)$  for each  $j$ , such that the asymptotic normality of the MLE for  $\phi$  extends to asymptotic normality of the posterior  $\pi_j(\phi|D_n)$  (see, for instance, Chen, 1985).

As introduced in Section 2.1, let  $w_j^* = \lim_{n \rightarrow \infty} \tilde{w}_j(D_n)$  be the limiting posterior weight on the  $j$ -th model, in the sense of almost sure convergence with respect to the distribution of  $(D_n|\phi^\dagger)$ . Toward characterizing  $w^*$ , from (3) we have, for  $j \neq k$ ,

$$\begin{aligned}
 (4) \quad \frac{f_j(D_n)}{f_k(D_n)} &= \frac{\int_{\Phi_0} \{\pi_j(\phi)/\pi_k(\phi)\} f(D_n|\phi)\pi_k(\phi) d\phi}{f_k(D_n)} \\
 &= E_k \{ \pi_j(\phi)/\pi_k(\phi) | D_n \},
 \end{aligned}$$

where the subscript  $k$  reminds us that this posterior expectation is taken with respect to prior  $\pi_k$ . Now presume that  $(\phi^\dagger, \lambda^\dagger) \in \Theta_k$  (so that (4) is a posterior expectation

arising in a “right model” scenario), and presume that the prior density ratio  $\pi_j(\phi)/\pi_k(\phi)$  is finite and continuous at  $\phi^\dagger$ . Then (4) leads to

$$\begin{aligned}
 (5) \quad \frac{w_j^*}{w_k^*} &= \left\{ \lim_{n \rightarrow \infty} \frac{f_j(D_n)}{f_k(D_n)} \right\} \left( \frac{w_j}{w_k} \right) \\
 &= \left\{ \frac{\pi_j(\phi^\dagger)}{\pi_k(\phi^\dagger)} \right\} \left( \frac{w_j}{w_k} \right).
 \end{aligned}$$

As a further technical note, should the prior density ratio  $\pi_j(\phi)/\pi_k(\phi)$  not be finite at  $\phi^\dagger$ , we can interchange  $j$  and  $k$  before appealing to (5). Or (5) could be made more formal by replacing the ratio of prior densities with the Radon–Nikodym derivative of one prior distribution with respect to the other.

Just as we can characterize the limiting posterior weights on models, we can also characterize the limiting values of within-model posterior means. Generically, let

$$\hat{\psi}(D; \pi_I) = \int g(\phi, \lambda)\pi_I(\phi, \lambda|D) d\phi d\lambda$$

be the posterior mean of the target  $\psi$  arising from the prior specification  $\pi_I$  and data  $D$ , where the  $I$  subscript reminds us this is the investigator’s choice of prior distribution. It is easy to verify that without an identifying restriction we have the large-sample limit of the posterior mean (again in the almost-sure sense with respect to the law of  $(D_n|\phi^\dagger)$ ) given as

$$\begin{aligned}
 (6) \quad \psi_0^* &\doteq \lim_{n \rightarrow \infty} \hat{\psi}(D_n; \pi_0) \\
 &= \lim_{n \rightarrow \infty} \int g(\phi, \lambda)\pi_0(\lambda|\phi)\pi_0(\phi|D_n) d\lambda d\phi \\
 &= \int g(\phi^\dagger, \lambda)\pi_0(\lambda|\phi^\dagger) d\lambda.
 \end{aligned}$$

We can similarly define limiting posterior means when the investigator imposes an identifying restriction, i.e.,  $\psi_j^* \doteq \lim_{n \rightarrow \infty} \hat{\psi}(D_n; \pi_j)$  for  $j > 0$ . Specifically, when the investigator invokes the  $j$ -th restriction, inference is based on an identified model parameterized by  $\phi$ . If the restriction is valid, the posterior mean of  $\psi$  will converge to the correct value of the target, that is, formally  $\psi_j^* = g_j(\phi^\dagger)$ , where  $g_j(\phi)$  is defined as the constant value of  $g(\phi, \lambda)$  arising for any  $\lambda$  such that  $(\phi, \lambda) \in \Theta_j$ . If the restriction is not valid, then the limiting posterior mean will converge to  $g_j(\cdot)$  evaluated at the “pseudo-true” value of  $\phi$  which minimizes the Kullback–Leibler divergence between the actual distribution of  $D$  and the modelled distribution of  $(D|\phi)$ , as per the wrong-model asymptotic theory of White (1982).

Assembling the pieces thus far, the large-sample limit of the BMA posterior mean is

$$(7) \quad \psi_{\text{MIX}}^* \doteq \sum_{j=0}^J w_j^* \psi_j^*.$$

### 2.3 Quantifying Performance

Regarding the posterior mean  $\hat{\psi}(D; \pi_I)$  as an estimator (in this case with a Bayesian motivation) of  $\psi$ , we quantify estimation accuracy starting with the usual frequentist mean-squared error (MSE), that is, averaging over repeated sampling of  $D$  for a fixed  $\theta$  to obtain  $\text{MSE}(\theta) = E_{\theta}[\{\hat{\psi}(D; \pi_I) - \psi(\theta)\}^2]$ . (Even though the distribution of  $(D|\theta)$  depends on  $\theta$  only through  $\phi$ , the MSE generally depends on *all* of  $\theta$ , since the target parameter  $\psi$  will be a function of both  $\phi$  and  $\lambda$ .) Then we average  $\text{MSE}(\theta)$  across different values of  $\theta$ , to reflect aggregate performance across different scenarios of what the truth might be. This averaging is weighted according to a density  $\pi_N(\cdot)$  over  $\Theta_0$ . We find it useful heuristically to label  $\pi_N(\cdot)$  as *Nature's* prior distribution (hence the subscript  $N$ ), as distinct from  $\pi_I$  which the investigator uses in performing the Bayesian analysis of the data. For data  $D_n$  with sample size  $n$ , then, the *average* mean-squared error (AMSE) is defined as:

$$(8) \text{AMSE}_n(\pi_N, \pi_I) \doteq E_{\pi_N} E_{\theta}[\{\hat{\psi}(D_n; \pi_I) - \psi(\theta)\}^2].$$

Note that we quite deliberately keep both the ‘A’ and the ‘M’ in ‘AMSE,’ to remind us that two averagings are in play: the mean (M) of squared error across  $D_n$  given  $\theta$ , distinct from the averaging (A) of this result across  $\theta$ .

Standard decision-theoretic terminology would see (8) referred to as the Bayes risk of the Bayesian estimator formed under  $\pi_I$ , with respect to  $\pi_N$ . Indeed, standard arguments tell us that for fixed  $\pi_N$ , (8) is minimized by taking  $\pi_I = \pi_N$ , and that this is the minimum achievable across *all* estimators, not just those arising as Bayesian estimators from some choice of prior (see, for instance, Berger, 1985, Section 4.4).

We focus on the large-sample limit, letting  $\text{AMSE}(\cdot) = \lim_{n \rightarrow \infty} \text{AMSE}_n(\cdot)$ . The large-sample limit of the posterior mean,  $\psi^*$ , depends on the true value of  $\phi$  and the chosen prior  $\pi_I$ , e.g.,  $\psi^*$  is given by (6) if  $\pi_I = \pi_0$ , or by (7) if  $\pi_I = \pi_{\text{MIX}}$ . More formally then

$$(9) \text{AMSE}(\pi_N, \pi_I) \doteq E_{\pi_N}[\{\psi^*(\phi; \pi_I) - \psi(\theta)\}^2].$$

For interpretability we will report the square root of AMSE, i.e.,  $\text{RAMSE} = \text{AMSE}^{1/2}$ .

The “two prior” framework above (Nature and investigator) deserves elaboration. The frequentist MSE of the posterior mean is being averaged, when (Nature’s) weighting over  $\Theta_0$  in forming this average may differ from the (investigator’s) weighting over  $\Theta_0$  in determining the posterior given data. While not seen commonly, this approach has precedents. For instance, Wang and Gelfand (2002) use the idea in a scheme to determine an appropriate sample size (using “simulation” and “fitting” in place of “Nature” and “investigator”). Heuristically, we can think of the averaging with respect to  $\pi_N$  as averaging

across all the scenarios Nature *might* present to us. Or, in a game-theory spirit, we can think of Nature as choosing the true (but unrevealed) state of the world by drawing from  $\pi_N$ . A different slant is espoused in Gustafson and Greenland (2009), whereby a quantity like (8) reflects the long-run performance of a research lab tasked with investigating a sequence of *different* scientific relationships, with  $\pi_N$  describing variation across these relationships.

The two-prior approach can naturally quantify the pros and cons of asserting that perhaps an identified sub-model holds. Particularly, we can compute  $\text{AMSE}(\pi_N, \pi_I)$  for the four combinations arising from  $\pi_I \in \{\pi_0, \pi_{\text{MIX}}\}$  and  $\pi_N \in \{\pi_0, \pi_{\text{MIX}}\}$ . Necessarily,  $\text{AMSE}(\pi_{\text{MIX}}, \pi_{\text{MIX}})$  is lower than (or possibly equal to)  $\text{AMSE}(\pi_{\text{MIX}}, \pi_0)$ . The extent to which it is lower reflects the benefit of making a maybe assertion in an appropriate context, since by setting  $\pi_N = \pi_{\text{MIX}}$  we study average-case performance across a mix of scenarios, some with, and some without, one of the identifying restrictions being true.

Conversely,  $\text{AMSE}(\pi_0, \pi_{\text{MIX}})$  is guaranteed to be higher than (or possibly equal to)  $\text{AMSE}(\pi_0, \pi_0)$ , and the extent to which it is higher reflects the risk of inappropriately making a maybe assertion. That is, across scenarios where none of the identifying restrictions hold, making the maybe assertion increases average-case MSE. Asserting maybe in a context where all the identified sub-models are actually *a priori* implausible can be regarded as a form of cheating. Thus the extent to which  $\text{AMSE}(\pi_0, \pi_{\text{MIX}})$  exceeds  $\text{AMSE}(\pi_0, \pi_0)$  quantifies the reduction in estimation performance resulting from such cheating.

Note that the AMSE comparisons just outlined could be applied in a fully identified setting with  $\psi$  identified under  $M_0$  (and consequently also identified under each sub-model  $M_j$ ,  $j = 1, \dots, J$ ). For a finite sample size  $n$ , the four  $\text{AMSE}_n$  values determined by (8) would retain the interpretations given above, quantifying both the value and the risk of speculating that the sub-models are *a priori* plausible. However, (9) would no longer be relevant. All four  $\text{AMSE}_n$  values would tend to zero as  $n$  tends to infinity, since both  $\hat{\psi}(D_n; \pi_0)$  and  $\hat{\psi}(D_n; \pi_{\text{MIX}})$  would consistently estimate  $\psi$ , under any values of  $\theta \in \Theta_0$ . Conversely, in the situation we study, with  $\psi$  not identified under  $M_0$ , consistent estimation does not arise. Hence the positive large-sample limits of the AMSE governed by (9) are fundamental descriptors of the situation.

### 3. EXAMPLE 1: PREVALENCE ESTIMATION WITH MISSING DATA

We return to the motivating example of Section 1.2. Here  $M_0$  is the NIM specification which does not fully identify the target, whereas  $M_1$  is the MAR specification which does. The prior specifications  $\pi_0(\cdot)$  and  $\pi_1(\cdot)$  are those given in Section 1.2. The posterior mean of the disease prevalence  $\psi$  has large-sample limit  $\psi_0^* =$

$p_{11}^\dagger + p_{0+}^\dagger/2$  when the investigator allows NIM ( $M_0$ ), but  $\psi_1^* = p_{11}^\dagger/p_{1+}^\dagger$  when the investigator presumes MAR ( $M_1$ ).

The prior marginal densities of  $\phi = (p_{0+}, p_{10}, p_{11})$  are readily obtained. Expressed as densities for  $(p_{0+}, p_{11})$  over the lower triangle of the unit square (rather than densities over the probability simplex), we have  $\pi_0(p_{0+}, p_{11}) = 6p_{0+}$  for the NIM specification, and  $\pi_1(p_{0+}, p_{11}) = (1 - p_{0+})^{-1}$ . Consequently,

$$(10) \quad \frac{w_1^*/w_0^*}{w_1/w_0} = \frac{1}{6p_{0+}^\dagger(1 - p_{0+}^\dagger)}.$$

This limiting Bayes factor could also be deduced by direct inspection of (1). Note that (10) places us in Scenario 3 of the taxonomy given earlier in Section 2.1.

Based on the prior specifications and the limiting posterior forms, and taking equal prior model weights  $w = (0.5, 0.5)$ , we have the following analytic expressions:

$$\begin{aligned} \text{AMSE}(\pi_0, \pi_0) &= E_{\pi_0} \left\{ \frac{(p_{01} - p_{00})^2}{4} \right\} \\ &= \frac{1}{40}, \end{aligned}$$

and

$$\begin{aligned} \text{AMSE}(\pi_{\text{MIX}}, \pi_0) &= (1/2)E_{\pi_0} \left\{ \frac{(p_{01} - p_{00})^2}{4} \right\} \\ &\quad + (1/2)E_{\pi_1} \left\{ \frac{(p_{01} - p_{00})^2}{4} \right\} \\ &= \frac{1}{2} \left( \frac{1}{40} + \frac{1}{36} \right). \end{aligned}$$

Moreover,

$$\text{AMSE}(\pi_0, \pi_{\text{MIX}}) = E_{\pi_0} \{w(p)^2\},$$

and

$$\begin{aligned} \text{AMSE}(\pi_{\text{MIX}}, \pi_{\text{MIX}}) &= \frac{1}{2} [E_{\pi_0} \{w(p)^2\} \\ &\quad + E_{\pi_1} \{w(p)^2\}], \end{aligned}$$

where

$$\begin{aligned} w(p) &= \frac{6p_{0+}p_{1+}}{1 + 6p_{0+}p_{1+}} \left( p_{11} + \frac{p_{0+}}{2} \right) \\ &\quad + \frac{1}{1 + 6p_{0+}p_{1+}} \left( \frac{p_{11}}{p_{1+}} \right) - p_{+1}. \end{aligned}$$

Given the nonlinearity of  $w()$ , the needed expectations do not have analytic forms. They are easily evaluated numerically, however, say via Monte Carlo draws from  $\pi_0()$  and  $\pi_1()$ .

Table 1 gives the four AMSE values. Positively, appropriate use of the maybe assertion leads to a 16% reduction in RAMSE, that is, we see this reduction when Nature's

prior is the mixture, so we are averaging performance across some scenarios where the assertion holds and others when it does not. Conversely, inappropriate use of the maybe assertion leads to a 12% increase in RAMSE. This is seen when Nature's prior is  $\pi_0$  alone, so we are averaging performance exclusively across NIM scenarios.

#### 4. EXAMPLE 2: ESTIMATING AN AVERAGE RISK DIFFERENCE FROM STRATIFIED DATA

Consider three binary variables,  $(C, X, Y)$ , with  $X$  an exposure variable,  $Y$  an outcome variable, and  $C$  a potential confounder. The target parameter is presumed to be the average risk difference,  $\psi = E\{E(Y|X = 1, C) - E(Y|X = 0, C)\}$ , which can be motivated from a causal inference perspective. Say the available data, however, are sampled conditional on  $C$ . That is, pre-specified numbers of realizations,  $n_0$  and  $n_1$ , are drawn from  $(Y, X|C = 0)$  and  $(Y, X|C = 1)$ , respectively.

Let  $\phi = (\phi_0, \phi_1)$ , where  $\phi_c$  comprises the  $(X, Y|C = c)$  cell probabilities, for  $c = 0, 1$ . (So each  $\phi_c$  is a probability vector, with four elements constrained to sum to one.) Also, let  $\lambda = \text{Pr}(C = 1)$ . Then  $(\phi, \lambda)$  is a transparent parameterization, with the likelihood  $f(D|\phi)$  factoring into  $\phi_0$  and  $\phi_1$  terms based on two independent multinomial samples. The target parameter can be expressed as  $g(\phi, \lambda) = (1 - \lambda)v(\phi_0) + \lambda v(\phi_1)$ , where  $v()$  returns the risk difference from a single set of cell probabilities. For the partially identified model  $M_0$  we specify a prior with independencies of the form  $\pi_0(\phi, \lambda) = \pi_0(\phi_0)\pi_0(\phi_1)\pi_0(\lambda)$ . For future reference, let  $\bar{\lambda}$  and  $\sigma_\lambda^2$  be the mean and variance of  $\lambda$  under the specified prior.

The first identifying restriction considered is that the prevalence of  $C$  is exactly known from external sources, that is,  $M_1$  is taken to be the sub-model of  $M_0$  defined by  $\lambda = \tilde{\lambda}$ , where  $\tilde{\lambda}$  is user-specified. We express this model via the prior specification  $\pi_1(\phi, \lambda) = \pi_1(\phi)\delta_{\tilde{\lambda}}(\lambda)$ , where  $\delta_x()$  is the Dirac delta function, that is, the 'density' of

TABLE 1

*RAMSE values for different combinations of Nature's prior and the investigator's prior, in Example 1. Values in the first column are exact. Values in the second column are computed as Monte Carlo averages based on  $10^5$  draws from each of  $\pi_0()$  and  $\pi_1()$ . The impact of the investigator using  $\pi_{\text{MIX}}$  rather than  $\pi_0$  is a 11.8% increase in RAMSE when this use is not warranted, but a 16.0% decrease in RAMSE when it is warranted. The Monte Carlo standard errors for these two percentage changes are 0.3 percentage points and 0.2 percentage points, respectively*

Nature's prior	Investigator's prior	
	$\pi_0$	$\pi_{\text{MIX}}$
$\pi_0$	0.158	0.177
$\pi_{\text{MIX}}$	0.162	0.136

a point-mass at  $x$ . We take  $\pi_1(\phi) = \pi_0(\phi)$  as before, but now  $\lambda = \tilde{\lambda}$  is taken as known.

The second identifying restriction considered is that of no interaction on the risk difference scale. Hence, we obtain  $M_2$  from  $M_0$  by constraining  $\phi$  according to  $v(\phi_0) = v(\phi_1)$ . An appropriate prior specification for the marginal  $\pi_2(\phi)$  would be the distribution induced from  $\pi_0(\phi)$  by conditioning on  $v(\phi_0) = v(\phi_1)$ . For  $M_2$  we can leave  $\pi_2(\lambda|\phi)$  unspecified, since it will play no role whatsoever in the posterior distribution of  $\psi$  or in the posterior weight of  $M_2$  relative to the other two models.

From the forms of  $\pi_j(\phi)$ , for  $j = 0, 1, 2$ , we have a very simple characterization of the limiting posterior model weights:

$$(11) \quad (w_0^*, w_1^*, w_2^*) = \begin{cases} (0, 0, 1) & \text{if } v(\phi_0^\dagger) = v(\phi_1^\dagger); \\ \left(\frac{w_0}{w_0 + w_1}, \frac{w_1}{w_0 + w_1}, 0\right) & \text{otherwise.} \end{cases}$$

Specifically,  $M_1$  is completely untestable compared to  $M_0$ , in the sense that  $f_0(D) = f_1(D)$  for every dataset. On the other hand  $M_2$  is completely testable compared to  $M_0$ , as  $v(\phi_0) = v(\phi_1)$  is either proven or disproven in the limit of infinite sample size. Referring back to the scenarios given in Section 2.1,  $M_1$  falls in the first scenario, while  $M_2$  falls in the second.

To describe within-model inference, under  $M_0$  the posterior distribution of  $(\psi|\phi)$  is a location-scale shift of the prior distribution of  $\lambda$  [with location  $v(\phi_0)$  and scale  $v(\phi_1) - v(\phi_0)$ ]. The limiting posterior mean of  $\psi$  is then  $v(\phi_0^\dagger) + \{v(\phi_1^\dagger) - v(\phi_0^\dagger)\}\bar{\lambda}$ . Whereas, relatedly, under  $M_1$  the limiting posterior mean is  $v(\phi_0^\dagger) + \{v(\phi_1^\dagger) - v(\phi_0^\dagger)\}\tilde{\lambda}$  (matching the target  $\psi^\dagger$  if  $M_1$  holds). If  $M_2$  holds, then asymptotically the  $M_2$  posterior will concentrate at the correct value of  $v(\phi_0^\dagger) = v(\phi_1^\dagger)$ . The limiting posterior mean under  $M_2$  when  $M_2$  does not hold is governed by standard wrong-model asymptotic theory (e.g., see White (1982)). We will not need to determine this limit, however, as (11) discredits  $M_2$  when it is wrong.

Armed with (11) and the limiting posterior means of  $\psi$  under each model, we proceed to determine AMSE for the four combinations of Nature’s prior and the investigator’s prior based on  $\pi_0()$  or  $\pi_{\text{MIX}}()$ . Letting  $k = \text{Var}\{v(\phi_1) - v(\phi_0)\} = 2 \text{Var}\{v(\phi_c)\}$  under  $\pi_0()$ , direct calculation gives

$$\begin{aligned} \text{AMSE}(\pi_0; \pi_0) &= k\sigma_\lambda^2, \\ \text{AMSE}(\pi_0; \pi_{\text{MIX}}) &= k\left[\sigma_\lambda^2 + \{w_2/(w_1 + w_2)\}^2(\tilde{\lambda} - \bar{\lambda})^2\right], \\ \text{AMSE}(\pi_{\text{MIX}}; \pi_0) &= k\{w_0\sigma_\lambda^2 + w_1(\tilde{\lambda} - \bar{\lambda})^2\}, \\ \text{AMSE}(\pi_{\text{MIX}}; \pi_{\text{MIX}}) &= k\left\{w_0\sigma_\lambda^2 + \frac{w_0w_1}{w_0 + w_1}(\tilde{\lambda} - \bar{\lambda})^2\right\}. \end{aligned}$$

TABLE 2

RAMSE values for different combinations of Nature’s prior and the investigator’s prior in Example 2. The impact of the investigator using  $\pi_{\text{MIX}}$  rather than  $\pi_0$  is a 1.9% increase when this use is not warranted, but a 3.5% decrease when it is warranted

Nature’s prior	Investigator’s prior	
	$\pi_0$	$\pi_{\text{MIX}}$
$\pi_0$	0.0464	0.0473
$\pi_{\text{MIX}}$	0.0288	0.0278

These expressions are sufficiently simple that one can “read off” the extents to which  $\text{AMSE}(\pi_0; \pi_{\text{MIX}})$  exceeds  $\text{AMSE}(\pi_0; \pi_0)$  and  $\text{AMSE}(\pi_{\text{MIX}}; \pi_{\text{MIX}})$  is reduced compared to  $\text{AMSE}(\pi_{\text{MIX}}; \pi_0)$ . Note that both these gaps collapse to zero if  $\tilde{\lambda}$  (the prior mean of  $\lambda$  under model  $M_0$ ) equals  $\bar{\lambda}$  (the presumed value of  $\lambda$  under Model  $M_1$ ).

To give a concrete example, say external information leads to the specifying  $\tilde{\lambda} = 0.15$  for the identifying restriction under  $M_1$ . And the same source suggests the prior  $\lambda \sim \text{Beta}(4, 18)$  under Model  $M_0$ , yielding the prior mode at 0.15, with prior mean of  $\bar{\lambda} = 4/22$  and prior variance of  $\sigma_\lambda^2 = (4/22)(18/22)(1/23) \approx 0.00647$ . Using a uniform prior on cell probabilities, that is,  $\phi_c \sim \text{Dirichlet}(1, 1, 1, 1)$ , gives  $k = 1/3$ . Presuming equal prior weights on the models,  $(w_0, w_1, w_2) = (1/3, 1/3, 1/3)$ , the resulting RAMSE values are given in Table 2. These reveal a low-stakes setting. Imposing the maybe restrictions when not warranted only carries a 1.9% increase in RAMSE (first row). While imposing them when warranted only carries a 3.5% reduction in RAMSE (second row).

**5. EXAMPLE 3: ESTIMATING A RISK DIFFERENCE WITH POSSIBLE OUTCOME MISCLASSIFICATION**

In a similar spirit to Example 2, say we are interested in the risk difference  $\text{Pr}(Y = 1|X = 1) - \text{Pr}(Y = 1|X = 0)$ , for binary variables  $X$  and  $Y$ . However, the available data are  $(X, Y^*)$  realizations, where  $Y^*$  may be an imperfect surrogate for  $Y$ , while  $Y$  itself is latent. Specifically, say  $Y^*$  and  $X$  are conditionally independent given  $Y$  (the so-called “nondifferential” assumption), with  $\lambda_y = \text{Pr}(Y^* = Y|Y = y)$ , that is,  $\lambda_0$  and  $\lambda_1$  are respectively, the specificity and sensitivity of the surrogate. Letting  $\omega_x = \text{Pr}(Y = 1|X = x)$ , the target of inference is  $\psi = \omega_1 - \omega_0$ . While  $(\omega, \lambda)$  is an interpretable parameterization, a transparent parameterization is  $(\phi, \lambda)$ , where  $\phi = (\phi_0, \phi_1)$ , with  $\phi_x = \text{Pr}(Y^* = 1|X = x) = (1 - \omega_x)(1 - \lambda_0) + \omega_x\lambda_1$ . With respect to this parameterization, the target is  $\psi = (\phi_1 - \phi_0)/(\lambda_0 + \lambda_1 - 1)$ .

As a prior specification for  $M_0$ , we let  $\omega$  follow a uniform distribution over  $(0, 1)^2$ , and independently let  $\lambda$  follow a uniform distribution over  $(a_0, 1) \times (a_1, 1)$ . Thus hyperparameters  $a_0$  and  $a_1$  are worst-case assertions about

the magnitude of outcome misclassification. (We presume  $a_0 > 0.5$  and  $a_1 > 0.5$  for technical reasons, which rules out placing any prior mass on a negative dependence between  $Y$  and  $Y^*$ .) The prior independence between  $\omega$  and  $\lambda$  is intuitive; prior assertions about  $(Y|X)$  prevalences are not tied to prior assertions about the quality of  $Y^*$  as a surrogate for  $Y$ . (And prior independence of  $\phi$  and  $\lambda$  would not be possible, since the support of  $\lambda$  depends on  $\phi$ .)

Upon moving to the transparent parameterization, following the related model formulation in Gustafson, Le and Saskin (2001), the  $M_0$  prior density transforms to

$$(12) \quad \pi_0(\phi, \lambda) = \frac{I_A(\phi, \lambda)}{(1 - a_0)(1 - a_1)(\lambda_0 + \lambda_1 - 1)^2}.$$

Here  $A$  is the intersection of three sets, imposing the restrictions that: (i),  $\phi \in (0, 1)^2$ ; (ii),  $\lambda \in (a_0, 1) \times (a_1, 1)$ ; and (iii),  $\phi$  and  $\lambda$  are compatible with each other in that  $(1 - \lambda_0) < \phi_x < \lambda_1$ , for  $x = 0, 1$ .

The identifying restriction considered is the assertion of perfect outcome classification, that is,  $M_1$  is the sub-model of  $M_0$  defined by  $\lambda = (1, 1)'$ . Under  $M_1$  we take the prior on  $\phi \sim \text{Unif}\{(0, 1)^2\}$  *a priori*. This is consistent with  $M_0$  in that it matches the  $(\phi|\lambda = (1, 1)')$  prior; see Appendix B for further discussion of prior specifications which match in this sense.

For a datastream arising under  $\phi = \phi^\dagger$ , we have the limiting Bayes factor  $\pi_1(\phi^\dagger)/\pi_0(\phi^\dagger)$  as per (5). Clearly,  $\pi_1(\phi) = 1$ , while marginalizing (12) gives

$$(13) \quad \pi_0(\phi) = \frac{1}{(1 - a_0)(1 - a_1)} [\log b_0(\phi) + \log b_1(\phi) - \log\{b_0(\phi) + b_1(\phi) - 1\}],$$

where  $b_0(\phi) = \max\{a_0, 1 - \phi_0, 1 - \phi_1\}$  and  $b_1(\phi) = \max\{a_1, \phi_0, \phi_1\}$ . The bivariate density (13) over the unit square has a “tabletop” shape. It is constant over the square defined by  $\min\{\phi_0, \phi_1\} > 1 - a_0$  and  $\max\{\phi_0, \phi_1\} < a_1$ . Whereas the density falls off away from this square, with  $\pi_0(\phi) \rightarrow 0$  as  $\min\{\phi_0, \phi_1\} \rightarrow 0$  or  $\max\{\phi_0, \phi_1\} \rightarrow 1$ .

The shape of (13) has clear implications for what the data can say about the identifying restriction. Say we are working with even prior models odds,  $w = (0.5, 0.5)'$ . Then, for a datastream generated under a value of  $\phi^\dagger$  within the tabletop mentioned above,

$$\frac{w_1^*}{w_0^*} = \{(1 - a_0)(1 - a_1)\} / [\log\{a_0 a_1 / (a_0 + a_1 - 1)\}].$$

This provides a bound on the extent of criticism the data can provide *against* the identified sub-model. On the other hand, there are values of  $\phi^\dagger$  outside the tabletop for which  $w_1^*/w_0^*$  is arbitrarily large. Thus, circumstances exist under which the data can provide strong support *for* the identified sub-model. (Though not as strong as if  $M_0$  were identified, in which case either  $w^* = (0, 1)'$  or  $w^* = (1, 0)'$ .) Referring back to taxonomy in Section 2.1, the restriction  $M_1$  falls into the third scenario.




FIG. 2. Distribution of  $w_1^*$ , the limiting posterior weight on  $M_1$ , in Example 3. The panels correspond to parameter ensembles for which  $M_0$  (left) and  $M_1$  (right) are true. Hyperparameters are  $a = (0.85, 0.85)'$  and  $w = (0.5, 0.5)'$ .

To better convey the ability of data to discriminate between  $M_0$  and  $M_1$ , we examine the  $w^*$  arising for two ensembles of  $\phi^\dagger$  values, generated by sampling from  $\pi_0()$  and  $\pi_1()$ , respectively. The  $M_0$  hyperparameter setting is  $(a_0, a_1) = (0.85, 0.85)$  for  $\pi_0()$ , and equal *a priori* model weights,  $(w_0, w_1) = (0.5, 0.5)$ , are prescribed. The two ensembles of the limiting posterior model weight  $w_1^*$  appear in Figure 2. As might be expected, there is relatively more mass at the lower boundary for  $w_1^*$  when  $M_1$  is false. (For the present hyperparameter values this lower boundary is 0.416, representing a very limited scope for criticism of  $M_1$ .) As per the discussion above, the right tail of the  $w_1^*$  values approaches one when  $M_1$  is true, that is, there are occasional circumstances where the data can strongly support the identifying restriction.

In terms of limiting posterior means for the target, when the investigator presumes  $M_1$ , the limit is simply  $\psi_1^* = \phi_1^\dagger - \phi_0^\dagger$ . For  $M_0$ , numerical integration is required to obtain  $\psi_0^*$  for a given  $\phi^\dagger$ . Specifically,  $\psi_0^* = \int g(\phi^\dagger, \lambda)\pi_0(\lambda|\phi^\dagger) d\lambda$ , where the prior conditional density for  $(\lambda|\phi)$  is derived from the joint density (12).

The RAMSE values, given in Table 3, reveal higher stakes than manifested in Example 2. Imposing the maybe restriction when not warranted incurs a penalty of a 45%

TABLE 3

RAMSE values for different combinations of Nature’s prior and the investigator’s prior in Example 3. Each value is computed as a Monte Carlo average, using 20,000 realizations from each of  $\pi_0()$  and  $\pi_1()$ . The impact of the investigator using  $\pi_{\text{MIX}}$  rather than  $\pi_0$  is a 44.8% increase in RAMSE when this use is not warranted, but a 24.9% decrease in RAMSE when it is warranted. Monte Carlo standard errors for these two percentage changes are 0.9 percentage points and 0.3 percentage points, respectively

Nature’s prior	Investigator’s prior	
	$\pi_0$	$\pi_{\text{MIX}}$
$\pi_0$	0.0270	0.0391
$\pi_{\text{MIX}}$	0.0462	0.0347



increase in RAMSE (first row). Whereas imposing the maybe restriction when warranted produces a 25% reduction in RAMSE (second row).

**6. EXAMPLE 4: ESTIMATING AN AVERAGE RISK DIFFERENCE WITH POSSIBLE EXPOSURE MISCLASSIFICATION**

Our final, and most complex, example blends elements of Examples 2 and 3. As per Example 2, consider the joint distribution of  $(C, X, Y)$ , where  $C$  is a binary confounding variable,  $X$  is a binary exposure variable, and  $Y$  is a binary outcome variable. The target of inference is again presumed to be the average risk difference,  $\psi = E\{\Pr(Y = 1|X = 1, C) - \Pr(Y = 1|X = 0, C)\}$ . However, the exposure  $X$  may be subject to misclassification. Hence, the observable variables are  $(C, X^*, Y)$ , where the surrogate exposure  $X^*$  may differ from  $X$ . We confine attention to a simple form of exposure misclassification. We take it as known that the misclassification is nondifferential, in the sense that  $X^*$  is conditionally independent of  $(C, Y)$  given  $X$ . Moreover, we presume  $\Pr(X^* = 0|X = 0) = 1$ , that is, the classification scheme has perfect *specificity*. However, the *sensitivity*,  $\lambda = \Pr(X^* = 1|X = 1)$ , may be less than one. Some recent work on “unidirectional” misclassification of this form includes Xia and Gustafson (2016, 2018).

**6.1 Convenience Prior and Posterior**

To organize prior specification and posterior computation across multiple models, it is helpful to specify a *convenience* prior distribution which in turn produces a convenience posterior distribution with a simple form. Appropriate model-specific prior specifications and posterior calculations can be expressed as tweaks to the convenience analysis.

Let  $\mathcal{P}_d$  be the space of probability vectors over  $d$  mutually distinct and exhaustive outcomes, and let  $\phi \in \mathcal{P}_8$  be the cell probabilities describing the distribution of  $(C, X^*, Y)$ . We specify the convenience prior density on  $\theta = (\phi, \lambda)$  to be  $\pi_*(\phi, \lambda) = \pi_*(\phi)\pi_*(\lambda)$ , where  $\pi_*(\phi)$  is the Dirichlet(1, ..., 1) density, while  $\pi_*(\lambda)$  is the Uniform( $b, 1$ ) density, where hyperparameter  $b$  is an *a priori* specified lower bound on the sensitivity of the exposure classification. For illustration, we take  $b = 0.5$  throughout.

**6.2 Model  $M_0$**

In the absence of any identifying restrictions, a first thought is that the convenience prior might be employed as the actual prior. This is not actually possible, however, since we do *not* have a Cartesian parameter space for  $(\phi, \lambda)$ . For a given sensitivity  $\lambda$ , let  $s_\lambda()$  map from the  $(C, X, Y)$  cell probabilities to the  $(C, X^*, Y)$  cell probabilities. Hence  $s_\lambda()$  maps from  $\mathcal{P}_8$  to its image  $\mathcal{S}_\lambda \subset \mathcal{P}_8$ .

This map is invertible, and for a given  $\phi \in \mathcal{P}_8$  it is possible to numerically determine whether  $\phi \in \mathcal{S}_\lambda$ , and, if so, compute  $s_\lambda^{-1}(\phi)$ . Thus an obvious adaptation of the convenience prior is to take the Model  $M_0$  prior as:

$$(14) \quad \pi_0(\phi, \lambda) = \frac{\pi_*(\phi, \lambda)I\{\phi \in \mathcal{S}_\lambda\}}{\Pr_*\{\phi \in \mathcal{S}_\lambda\}}.$$

Thus, we simply truncate the convenience prior to only those  $\phi$  and  $\lambda$  pairs which are compatible with each other.

It transpires that  $\phi \in \mathcal{S}_\lambda$  if and only if  $\lambda$  exceeds a threshold  $t(\phi)$ , that is, the cell probabilities for the observables  $(C, X^*, Y)$  imply a lower bound on the sensitivity of  $X^*$  as a surrogate for  $X$ . (Note that this bound may or may not exceed the *a priori* lower bound  $b$ .) Combined with the specification of  $\pi_*(\lambda)$  as Unif( $b, 1$ ), (14) can be simplified to

$$\pi_0(\phi, \lambda) = \frac{\pi_*(\phi, \lambda)I[\lambda > \max\{b, t(\phi)\}]}{(1 - b)^{-1}E_*[1 - \max\{b, t(\phi)\}]}.$$

Importantly, this marginalizes to

$$(15) \quad \pi_0(\phi) = \frac{\pi_*(\phi)[1 - \max\{b, t(\phi)\}]}{E_*[1 - \max\{b, t(\phi)\}]}.$$

In terms of the parameter of interest, say that  $\tilde{g}()$  maps from the  $(C, X, Y)$  cell probabilities to the target parameter, that is,  $\tilde{g}()$  returns  $E\{E(Y|X = 1, C) - E(Y|X = 0, C)\}$ . In the parameterization at hand then, the target parameter is  $g(\phi, \lambda) = \tilde{g}(s_\lambda^{-1}(\phi))$ .

**6.3 Model  $M_1$**

In a similar spirit to Example 3, the first identifying restriction considered is simply that the surrogate  $X^*$  is in fact perfect. So model  $M_1$  is the sub-model of  $M_0$  corresponding to  $\lambda = 1$ . An obvious prior specification is

$$\pi_1(\phi, \lambda) = \pi_*(\phi)\delta_1(\lambda).$$

This simply marginalizes to

$$(16) \quad \pi_1(\phi) = \pi_*(\phi).$$

Under this restriction  $\phi$  is identically the  $(C, X, Y)$  cell probabilities. Hence the target parameter is simply expressed as  $g(\phi) = \tilde{g}(\phi)$ .

**6.4 Model  $M_2$**

In the spirit of Example 2, Model  $M_2$  renders the target identifiable via asserting an absence of interaction on the risk difference scale, so that  $E(Y|X = 1, C = c) - E(Y|X = 0, C = c)$  does not depend on  $c$ . Letting  $\tilde{\phi}$  parameterize the so restricted  $(C, X, Y)$  cell probabilities (so that  $\tilde{\phi}$  has only six degrees of freedom), the resultant  $(C, X^*, Y)$  cell probabilities can be expressed as  $h(\tilde{\phi}, \lambda)$ . Here the map  $h()$  is invertible. However, its image, which we denote as  $\mathcal{H}$ , is a strict subset of  $\mathcal{P}_8$ . While it does not seem possible to express  $h^{-1}()$  in closed form, for a given

$\phi \in \mathcal{P}_8$  one can numerically determine whether  $\phi \in \mathcal{H}$ , and, if so, compute  $h^{-1}(\phi)$ .

A sensible prior construction for  $M_2$  thus takes the form

$$\pi_2(\phi, \lambda) = \frac{\pi_*(\phi)I_{\mathcal{H}}(\phi)}{\Pr_*(\phi \in \mathcal{H})} \delta_{m(\phi)}(\lambda),$$

where  $m(\phi)$  is the unique sensitivity value implied by  $\phi$ . Clearly this prior marginalizes to

$$(17) \quad \pi_2(\phi) = \frac{\pi_*(\phi)I_{\mathcal{H}}(\phi)}{\Pr_*(\phi \in \mathcal{H})}.$$

**6.5 Limiting Inference Across Models**

The limiting posterior model weights are governed by (15), (16), and (17), giving

$$(18) \quad \frac{w_0^*/w_1^*}{w_0/w_1} = \frac{1 - \max\{b, t(\phi^\dagger)\}}{1 - E_*[\max\{b, t(\phi)\}]},$$

and

$$(19) \quad \frac{w_2^*/w_1^*}{w_2/w_1} = \begin{cases} \{\Pr_*(\phi \in \mathcal{H})\}^{-1} & \text{if } \phi^\dagger \in \mathcal{H}, \\ 0 & \text{otherwise.} \end{cases}$$

Via (18) the data can offer some mild discrimination between  $M_0$  and  $M_1$ , in that a positive and finite limiting Bayes factor arises. Moreover, this limit is not one (except for a set of  $\phi^\dagger$  values with probability zero under all three model-specific priors). (Of course this is indeed mild. If  $M_0$  were identified, the limiting Bayes factor would be zero or infinity.) Via (19), when  $M_2$  is false, enough data may or may not discredit it. That is, the true model (either  $M_0$  or  $M_1$ ) may give rise to  $\phi^\dagger \notin \mathcal{H}$ , in which case  $M_2$  is discredited asymptotically. However, any of the three models can yield  $\phi^\dagger \in \mathcal{H}$ , in which case (19) offers no discrimination. Note that if  $M_2$  is not discredited, then it necessarily receives more support than  $M_1$ , i.e., (19) exceeds one. Linking back to the Section 2.1 taxonomy,  $M_1$  and  $M_2$  are instances of Scenarios 3 and 4, respectively.

**6.6 Performance**

To further convey the ability to discriminate between the three models, we examine  $w^*$  values arising for three ensembles of  $\phi^\dagger$  values, where the ensembles are randomly drawn from  $\pi_0(\phi)$ ,  $\pi_1(\phi)$ , and  $\pi_2(\phi)$  respectively. Equal prior weights  $w = (1/3, 1/3/1/3)$  are employed. The results are plotted in Figure 3. For the ensembles corresponding to  $M_0$  and  $M_1$ ,  $w_2^* = 0$  for large majorities of points. So  $M_2$  is often, but not always, fully discredited when it is false. On the other hand, when  $M_2$  is not discredited (the minorities of scenarios arising under  $M_0$  and  $M_1$ , but *all* the scenarios arising under  $M_2$ ), it receives strong support. The  $w_2^*$  values range from 0.72 to 0.92.

Focusing on the  $w_2^* = 0$  cases generated under  $M_0$  and  $M_1$ , we see a modest tendency for  $w_1^* = 1 - w_0^*$  to be smaller when  $M_0$  is true and larger when  $M_1$  is true. There




FIG. 3. Distribution of the limiting posterior weights ( $w_0^*$ ,  $w_1^*$ ,  $w_2^*$ ), in Example 4, under  $M_0$  (upper-left),  $M_1$  (upper-right), and  $M_2$  (lower-left). In each instance, an ensemble of 100  $\phi^\dagger$  points are simulated from  $\pi_j(\cdot)$ . Points in the upper panels are jittered with a small amount of random noise, in order to better see the distribution of those points with  $w_0^* + w_1^* = 1$ ,  $w_2^* = 0$ .

is also an asymmetry. The strongest evidence in favour of  $M_0$  corresponds to a value of  $w_0^* = 1 - w_1^*$  well below one (approximately  $w_0^* = 0.73$ ), and this extreme can be reached under either  $M_0$  or  $M_1$ . Conversely, there are narrow circumstances under which  $w_1^* = 1 - w_0^*$  is very close to one. As was the case in Example 3, there is more scope to support the identifying restriction  $M_1$  than there is to criticize it.

The RAMSE values are given in Table 4. They show fairly high stakes attached to invoking the maybe assumptions. RAMSE is increased by 31% in terms of average performance across scenarios where neither restriction holds. Whereas RAMSE is reduced by 40% in terms

TABLE 4  
RAMSE values for different combinations of Nature’s prior and the investigator’s prior in Example 4. Each value is computed as a Monte Carlo average, using 1600 draws from each of  $\pi_0(\cdot)$ ,  $\pi_1(\cdot)$ ,  $\pi_2(\cdot)$ . The impact of the investigator using  $\pi_{\text{MIX}}$  rather than  $\pi_0$  is a 31.3% increase in RAMSE when this use is not warranted, but a 39.5% decrease in RAMSE when it is warranted. Monte Carlo standard errors for these two percentage changes are 4.0 percentage points and 2.1 percentage points, respectively

Nature’s prior	Investigator’s prior	
	$\pi_0$	$\pi_{\text{MIX}}$
$\pi_0$	0.0452	0.0594
$\pi_{\text{MIX}}$	0.0778	0.0470

of average performance when the mix of scenarios indeed has one of the restrictions holding in a portion of cases.

## 7. DISCUSSION

In current applied practice, when faced with a lack of full identification unless an identifying but untestable assumption is imposed, one generally imposes the restriction or not. In the terminology of this paper, and with only a single restriction contemplated, this is choosing hyperparameter settings of either  $w = (0, 1)$  or  $w = (1, 0)$ . To our knowledge, this has been the first investigation of intermediate settings, such as  $w = (0.5, 0.5)$ , with the interpretation of asserting that perhaps the restriction holds.

We certainly do not claim a free lunch arises from saying perhaps. But we do think it a worthwhile analysis strategy to consider, simply because many applied problems do arise where an identifying restriction is both plausible, but not incontrovertible, in the scientific context at hand. Moreover, the finding that sometimes the limiting posterior model weights differ from the prior model weights adds intrigue, and additional promise, to the use of perhaps assertions.

In general, to specify a prior distribution is to implicitly choose the estimator with optimal average-case behavior, where the average is with respect to the joint distribution of parameters and data arising from the specified prior. In situations with plausible identifying restrictions, we have seen the stakes associated with prior assertions can be quite high. In one direction,  $\text{AMSE}(\pi_0, \pi_{\text{MIX}})$  can be substantially higher than  $\text{AMSE}(\pi_0, \pi_0)$ . Purely wishful thinking that perhaps an identifying restriction holds can come with a steep penalty. Equally, however,  $\text{AMSE}(\pi_{\text{MIX}}, \pi_{\text{MIX}})$  can be substantially lower than  $\text{AMSE}(\pi_{\text{MIX}}, \pi_0)$ . Realistic assessment that perhaps an identifying restriction holds comes with a reward.

An admitted limitation of the work is the presumption that  $M_0$  is not violated. Put plainly, we admit we might lack identification, but in doing so we presume no model misspecification. This issue is moot in Examples 1 and 2, since  $M_0$  does not place restrictions on the underlying  $(R, Y)$  distribution in the former instance, or the underlying  $(C, X, Y)$  distribution in the latter. In Examples 3 and 4, however,  $M_0$  involves a risk of misspecification, particularly through assumptions about misclassification mechanisms and *a priori* bounds on misclassification parameters.

Finally, returning to the posterior updating of model weights, we have seen identifying restrictions that are empirically untestable, in the sense that  $w_1^*/w_0^*$  is neither zero nor infinite, that is, even an infinite amount of data would neither definitively prove nor definitively disprove the identifying assumption. In some such cases, however, the positive and finite value of  $w_1^*/w_0^*$  does vary with the underlying parameter values  $\theta^\dagger$  (via dependence on  $\phi^\dagger$ ).

Commensurately, in finite samples the posterior model weights can depend on the data. Thus, a nuanced situation of partial learning about the plausibility of a restriction can result. Consequently, the pros and cons of making perhaps suppositions *a priori* are not easily and generally intuited. For a given problem, however, we have shown how to elucidate the structure of the inference. Consequently, the risks and rewards of giving some prior credence to one or more identifying restrictions can be quantified.

## APPENDIX A: SENSITIVITY TO HYPERPARAMETERS

Here we elaborate on sensitivity to the prior specification, in the context of the motivating example explored in Section 1.2 and Section 3. Say we retain the  $M_0$  within-model prior  $p \sim \text{Dirichlet}(1, 1, 1, 1)$ . Under the MAR assumption ( $M_1$ ), however, we generalize from  $\gamma, \psi \stackrel{\text{iid}}{\sim} \text{Beta}(1, 1)$  to  $\gamma, \psi \stackrel{\text{iid}}{\sim} \text{Beta}(k, 1)$  for some choice of hyperparameter  $k$ . (See Appendix B for a rationale for the choice  $k = 2$ .)

In general, regardless of identification issues, Bayes factors are known to be sensitive to the choice of prior specifications within the two models being compared. This is certainly the case in the present situation. The ‘‘conjugate integration’’ that gave the Bayes factor (1) generalizes to

$$(A.1) \quad b_k = \frac{k^2(n+3)!\Gamma(c_{11}+k)}{6\Gamma(n+k+1)c_{11}!(c_{0+}+1)(n-c_{0+}+1)}.$$

For the dataset from Section 1.2, the Bayes factor is indeed very sensitive to increasing  $k$  above the base value of  $k = 1$  (e.g.,  $b_{1.5} = 2.17$ ,  $b_2 = 1.00$ , compared to  $b_1 = 3.73$ ). In the other direction, there is less sensitivity to decreasing  $k$  (e.g.,  $b_{0.75} = 4.14$ ,  $b_{0.5} = 3.62$ ).

Importantly, alternative values of the hyperparameter  $k$  also result in a Bayes factor which responds sensitively to the observed data, a key point in the earlier developments. The case  $k = 2$  is instructive, as it yields a simple expression for (A.1), namely  $b_2 = (2/3)\{(c_{11}+1)/(c_{1+}+1)\}/\{c_{0+}/n\}$ , with the latter two terms being, approximately, the sample proportion of  $Y = 1$  amongst responders, and the sample proportion of non-responders. So we see again a strong dependence on the data, the lack of identification notwithstanding. However, the nature of this dependence is rather different than for  $k = 1$ .

## APPENDIX B: PRIOR SPECIFICATIONS FOR NESTED MODELS

With Bayesian model averaging, there is no overarching prescription for how respective within-model prior specifications should relate when one model is nested within another. Following our notation of  $M_1$  being nested within  $M_0$ , say  $\pi_0()$  is already specified. Kass and Raftery (1995) point out that specifying  $\pi_1()$  via either marginalizing or conditioning has intuitive appeal (though there aren’t

strictures on other possibilities). To be more explicit, presume an initial parameterization of  $M_0$  of the form  $\theta = (\theta_a, \theta_b)$ , with  $M_1$  defined by  $\theta_b = 0$ . Then marginalizing leads to  $\pi_1(\theta_a) = \int \pi_0(\theta_a, \theta_b) d\theta_b$ , while conditioning leads to  $\pi_1(\theta_a) = \pi_0(\theta_a | \theta_b = 0)$ . Many of our examples use the conditioning route. Particularly, this is true in Example 1 (with  $\theta_b = p_{11} - (p_{01} + p_{11})(p_{10} + p_{11})$ ), Example 2 (with  $\theta_b = \lambda - \hat{\lambda}$  and  $\theta_b = \nu(\phi_1) - \nu(\lambda_0)$  for the two nested models, respectively), and Example 3 (with  $\theta_b = \lambda = (1, 1)'$ ).

Kass and Raftery (1995), citing an unpublished communication chain with J. Dickey and L.J. Savage, raise an interesting point about the conditioning route being subject to the ‘‘Borel paradox.’’ This is indeed somewhat more than a theoretical curiosity. For instance, consider Example 1. Starting from  $\pi_0(\cdot)$  of the form  $p \sim \text{Dir}(1, 1, 1, 1)$ , instantiating the MAR assumption ‘additively’ by conditioning on  $p_{11} - (p_{01} + p_{11})(p_{10} + p_{11}) = 0$  results in  $\gamma, \psi \stackrel{\text{iid}}{\sim} \text{Beta}(1, 1)$ , as used in Sections 1.2 and 3. However, instead instantiating MAR ‘multiplicatively’ by conditioning on  $p_{11}(p_{01} + p_{11})^{-1}(p_{10} + p_{11})^{-1} - 1 = 0$  leads to  $\gamma, \psi \stackrel{\text{iid}}{\sim} \text{Beta}(2, 1)$ , as used for sensitivity analysis in Appendix A.

## ACKNOWLEDGMENTS

The author would like to thank the anonymous referees, an Associate Editor and the Editor for their constructive comments that improved the quality of this paper.

## FUNDING

The author was supported by the Natural Sciences and Engineering Research Council of Canada, Discovery Grant RGPIN-2019-03957.

## SUPPLEMENTARY MATERIAL

**Computer Code** (DOI: [10.1214/23-STS885SUPP](https://doi.org/10.1214/23-STS885SUPP); .zip). R code to produce all the empirical results in the paper is available at: <https://github.com/paulgstf/param-restrict-maybe>.

## REFERENCES

- BERGER, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. *Springer Series in Statistics*. Springer, New York. MR0804611 <https://doi.org/10.1007/978-1-4757-4286-2>
- CASELLA, G. and BERGER, R. L. (2002). *Statistical Inference*, 2nd ed. Duxbury, N. Scituate.
- CHEN, C. F. (1985). On asymptotic normality of limiting density functions with Bayesian implications. *J. Roy. Statist. Soc. Ser. B* **47** 540–546. MR0844485
- DANIELS, M. J. and HOGAN, J. W. (2008). *Missing Data in Longitudinal Studies: Strategies for Bayesian Modeling and Sensitivity Analysis*. Chapman & Hall, CRC Press, New York. <https://doi.org/10.1201/9781420011180>

- FOX, M. P., MACLEHOSE, R. F. and LASH, T. L. (2022). *Applying Quantitative Bias Analysis to Epidemiologic Data*, 2nd ed. Springer, Berlin. <https://doi.org/10.1007/978-3-030-82673-4>
- FRANKS, A. M., D’AMOUR, A. and FELLER, A. (2020). Flexible sensitivity analysis for observational studies without observable implications. *J. Amer. Statist. Assoc.* **115** 1730–1746. MR4189753 <https://doi.org/10.1080/01621459.2019.1604369>
- GREENLAND, S. (2003). The impact of prior distributions for uncontrolled confounding and response bias: A case study of the relation of wire codes and magnetic fields to childhood leukemia. *J. Amer. Statist. Assoc.* **98** 47–54. MR1977199 <https://doi.org/10.1198/01621450338861905>
- GREENLAND, S. (2005). Multiple-bias modelling for analysis of observational data. *J. Roy. Statist. Soc. Ser. A* **168** 267–306. MR2119402 <https://doi.org/10.1111/j.1467-985X.2004.00349.x>
- GUSTAFSON, P. (2007). Measurement error modelling with an approximate instrumental variable. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **69** 797–815. MR2368571 <https://doi.org/10.1111/j.1467-9868.2007.00611.x>
- GUSTAFSON, P. (2015). *Bayesian Inference for Partially Identified Models: Exploring the Limits of Limited Data*. *Monographs on Statistics and Applied Probability* **141**. CRC Press, Boca Raton, FL. MR3642458
- GUSTAFSON, P. (2023). Supplement to ‘‘Parameter restrictions for the sake of identification: Is there utility in asserting that perhaps a restriction holds?’’ <https://doi.org/10.1214/23-STS885SUPP>
- GUSTAFSON, P. and GREENLAND, S. (2006). The performance of random coefficient regression in accounting for residual confounding. *Biometrics* **62** 760–768. MR2247204 <https://doi.org/10.1111/j.1541-0420.2005.00510.x>
- GUSTAFSON, P. and GREENLAND, S. (2009). Interval estimation for messy observational data. *Statist. Sci.* **24** 328–342. MR2757434 <https://doi.org/10.1214/09-STS305>
- GUSTAFSON, P., LE, N. D. and SASKIN, R. (2001). Case-control analysis with partial knowledge of exposure misclassification probabilities. *Biometrics* **57** 598–609. MR1855698 <https://doi.org/10.1111/j.0006-341X.2001.00598.x>
- HOETING, J. A., MADIGAN, D., RAFTERY, A. E. and VOLINSKY, C. T. (1999). Bayesian model averaging: A tutorial (with comments by M. Clyde, David Draper and E. I. George, and a rejoinder by the authors). *Statist. Sci.* **14** 382–417. MR1765176 <https://doi.org/10.1214/ss/1009212519>
- KASS, R. E. and RAFTERY, A. E. (1995). Bayes factors. *J. Amer. Statist. Assoc.* **90** 773–795. MR3363402 <https://doi.org/10.1080/01621459.1995.10476572>
- KEELE, L. and QUINN, K. M. (2017). Bayesian sensitivity analysis for causal effects from  $2 \times 2$  tables in the presence of unmeasured confounding with application to presidential campaign visits. *Ann. Appl. Stat.* **11** 1974–1997. MR3743285 <https://doi.org/10.1214/17-AOAS1048>
- LASH, T. L., FOX, M. P., MACLEHOSE, R. F., MALDONADO, G., MCCANDLESS, L. C. and GREENLAND, S. (2014). Good practices for quantitative bias analysis. *Int. J. Epidemiol.* **43** 1969–1985. [https://doi.org/10.1007/978-3-030-82673-4\\_13](https://doi.org/10.1007/978-3-030-82673-4_13)
- LITTLE, R. J. A. and RUBIN, D. B. (2019). *Statistical Analysis with Missing Data*, Vol. 793, 3rd ed. Wiley, New York.
- MANSKI, C. F. (2003). *Partial Identification of Probability Distributions*. *Springer Series in Statistics*. Springer, New York. MR2151380
- SCHARFSTEIN, D. O., DANIELS, M. J. and ROBINS, J. M. (2003). Incorporating prior beliefs about selection bias into the analysis of randomized trials with missing outcomes. *Biostatistics* **4** 495–512. <https://doi.org/10.1093/biostatistics/4.4.495>

- VANSTEELENDT, S., GOETGHEBEUR, E., KENWARD, M. G. and MOLENBERGHS, G. (2006). Ignorance and uncertainty regions as inferential tools in a sensitivity analysis. *Statist. Sinica* **16** 953–979. [MR2281311](https://doi.org/10.1006/jmps.1999.1278)
- VERSTRAETEN, T., FARAH, B., DUCHATEAU, L. and MATU, R. (1998). Pooling sera to reduce the cost of HIV surveillance: A feasibility study in a rural Kenyan district. *Trop. Med. Int. Health* **3** 747–750. <https://doi.org/10.1046/j.1365-3156.1998.00293.x>
- WANG, F. and GELFAND, A. E. (2002). A simulation-based approach to Bayesian sample size determination for performance under a given model and for separating models. *Statist. Sci.* **17** 193–208. [MR1925941 https://doi.org/10.1214/ss/1030550861](https://doi.org/10.1214/ss/1030550861)
- WASSERMAN, L. (2000). Bayesian model selection and model averaging. *J. Math. Psych.* **44** 92–107. [MR1770003 https://doi.org/10.1006/jmps.1999.1278](https://doi.org/10.1006/jmps.1999.1278)
- WHITE, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica* **50** 1–25. [MR0640163 https://doi.org/10.2307/1912526](https://doi.org/10.2307/1912526)
- XIA, M. and GUSTAFSON, P. (2016). Bayesian regression models adjusting for unidirectional covariate misclassification. *Canad. J. Statist.* **44** 198–218. [MR3507780 https://doi.org/10.1002/cjs.11284](https://doi.org/10.1002/cjs.11284)
- XIA, M. and GUSTAFSON, P. (2018). Bayesian inference for unidirectional misclassification of a binary response trait. *Stat. Med.* **37** 933–947. [MR3760458 https://doi.org/10.1002/sim.7555](https://doi.org/10.1002/sim.7555)