Bias Correction in Non-Differentiable Estimating Equations for Optimal Dynamic Regimes

Erica E. M. Moodie[1] Thomas S. Richardson[2]

November 26, 2006

[1]Department of Epidemiology, Biostatistics, and Occupational Health
McGill University
Montreal, QC H3A 1A2 CANADA

[2]Department of Statistics
University of Washington
Seattle, WA 98195

Authors' Footnote:

Erica E. M. Moodie is an Assistant Professor of Biostatistics in the Department of Epidemiology and Biostatistics, McGill University, Montreal, QC H3A 1A2, Canada (email: erica.moodie@mcgill.ca); Thomas S. Richardson is an Associate Professor in the Department of Statistics, University of Washington, Seattle, WA 98195 (email: thomasr@u.washington.edu).
Erica Moodie is currently supported by a NSERC University Faculty Award and a NSERC Discovery Grant. Thomas Richardson acknowledges support by NSF grant DMS 0505865 and NIH R01 AI032475.

Abstract

A dynamic regime is a function that takes treatment and covariate history and baseline covariates as inputs and returns a decision to be made. Robins (2004) proposed g-estimation using structural nested mean models for making inference about the optimal regime in a multi-interval trial that provide clear advantages over traditional parametric approaches.

Robins' g-estimation method always yields consistent estimators, but these can be asymptotically biased under a given structural nested mean model for certain longitudinal distributions of the treatments and covariates, termed exceptional laws. In fact, under the null hypothesis of no treatment effect, every distribution constitutes an exceptional law under structural nested mean models which allow for interaction of current treatment with past treatments or covariates. This paper provides an explanation of exceptional laws and describes a new approach to g-estimation which we call Zeroing Instead of Plugging In (ZIPI). ZIPI shares all of the asymptotic properties of recursive g-estimates at non-exceptional laws while providing substantial reduction in the bias at exceptional laws when decision rule parameters are not shared across intervals.

KEYWORDS: Dynamic treatment regimes; Optimal structural nested mean models; Asymptotic bias; g-estimation; Pre-test estimators.

1. INTRODUCTION

In a study aimed at estimating the mean effect of a treatment on a time-dependent outcome, dynamic treatment regimes are an obvious choice to consider. A dynamic regime is a function that has treatment and covariate history as well as baseline covariates as arguments; the value returned by the function is the treatment to be given. A recently proposed method of inference (Robins 2004) is g-estimation using structural nested mean models. This method is semi-parametric and provides consistent estimators under mis-specification of certain models, and is therefore clearly advantageous over traditional parametric approaches. However, g-estimation is asymptotically biased under specific data distributions which Robins has termed the *exceptional laws*. At exceptional laws, g-estimates of parameters for optimal structural nested mean models are asymptotically biased and exhibit non-regular behaviour.

The aim of this paper is to provide a means of reducing the asymptotic bias of g-estimates in the presence of exceptional laws. Exceptional laws are discussed in the optimal dynamic regimes literature only by Robins (2004); there, a method of detecting these laws and improving coverage of confidence intervals, but not point estimates, is proposed.

Structural nested mean models and g-estimation are discussed in section 2.1, followed by an detailed explanation of exceptional laws in g-estimation in section 2.2. Section 3.2 describes a method of reducing the asymptotic bias due to exceptional laws which we call *Zeroing Instead of Plugging In* (ZIPI), with theoretical properties derived in section 3.3. Sections 3.4 and 3.5 compare ZIPI to g-estimation and to the Iterative Minimization for Optimal Regimes (IMOR) method proposed by Murphy (2003). Section 4 concludes.

2. INFERENCE VIA G-ESTIMATION AND EXCEPTIONAL LAWS

2.1 Inference for optimal dynamic treatment regimes: structural nested mean models and g-estimation

We consider the estimation of treatment effect in a sequential study. All of the methods discussed in this paper are appropriate for a finite number of treatment intervals. We shall focus on the two-interval case where a single covariate is used to determine the optimal rule

at each interval. In particular, proofs shall be presented for two intervals only, however our development extends to the general case.

Notation Notation is adapted from Murphy (2003) and Robins (2004). Treatments are taken at K fixed times, $t_1, ..., t_K$. L_j are the covariates measured prior to treatment at the beginning of the j^{th} interval, while A_j is the treatment taken subsequent to having measured L_j. Y is the outcome observed at the end of the second interval; larger values of Y are favorable. Specific or observed values are denoted in lower-case. Denote a variable X_j at t_j and its history, $(X_1,X_2,...,X_j)$, by $\bar{X}_j$; denote the variable at t_j and its future, $(X_j,X_{j+1},...,X_K)$, by $\underline{X}_j$. Finally, denote a treatment at t_j that depends on history, $\bar{L}_j, \bar{A}_{j-1}$, by $D_j(\bar{L}_j, \bar{A}_{j-1})$.

In the two-interval case, treatments are taken at two fixed times, t_1 and t_2. L_1 and L_2 are the covariates measured prior to treatment at the beginning of the first and second intervals, respectively, i.e. at t_1 and t_2. In particular, L_1 represents baseline covariates and L_2 includes time-varying covariates which may depend on treatment received at t_1. A_j, $j = 1, 2$, is the treatment given subsequent to observing L_j. Thus, the data are, chronologically, (L_1, A_1, L_2, A_2, Y).

Throughout this paper, models shall be explained in terms of *potential outcomes*, the value of a covariate or the final outcome that would result if a subject were assigned to different treatments. In the two-interval case, we denote by $L_2(a_1)$ a subject's potential covariate at the beginning of the second interval if treatment a_1 is received by that subject, and $Y(a_1, a_2)$ denoting the potential end-of-study outcome if he follows regime (a_1, a_2).

Potential outcomes adhere to the *axiom of consistency*: $L_2(a_1) = L_2$ whenever treatment a_1 is actually received and $Y(a_1, a_2) = Y$ whenever (a_1, a_2) is received. That is, the actual and counterfactual covariates (or outcome) are equal when the regime in question is the regime actually received.

Assumptions To estimate the effect of a dynamic regime, we require the *stable unit treat-*

ment value assumption (Rubin 1978), which states that a subject's outcome is not influenced by other subjects' treatment allocation, and *no unmeasured confounders* (Robins 1997), which says that for any regime $\bar{a}_K$, $A_j \perp (\bar{L}_j(\bar{L}_{j-1}, \bar{a}_{j-2}), ..., Y(\bar{L}_K, \bar{a}_{K-1}))|\bar{L}_j, \bar{a}_{j-1}$, and always holds under sequential randomization. In two intervals, we have $A_1 \perp (L_2(a_1), Y(a_1, a_2))|L_1$ and $A_2 \perp Y(a_1, a_2)|(L_1, a_1, L_2)$. That is, treatment received in any interval is independent of any future potential outcome.

Without further assumptions the optimal regime may only be estimated from among the set of *feasible* regimes (Robins 1994); feasibility requires some subjects to have followed regime $(d_1(l_1), ..., d_K(\bar{l}_K, \bar{a}_{K-1}))$ so as to be able to perform non-parametric estimation.

Structural nested mean models and g-estimation Robins (2004, p. 209-214) produced a number of estimating equations for finding optimal regimes using structural nested mean models (SNMM) (Robins 1986). We use a subclass of SNMMs called the *optimal blip-to-zero* functions, denoted by $\gamma_j(\cdot)$, defined as the expected difference in outcome when using the "zero" regime (such as placebo or standard care) which we denote 0_j instead of treatment a_j at t_j, in subjects with treatment and covariate history $\bar{l}_j, \bar{a}_{j-1}$ who subsequently receive the optimal regime, components of which we denote $d_j^{opt}\big(\bar{l}_j, \bar{a}_{j-1}\big)$; what is optimal subsequent to t_j may depend on the treatment received at t_j. The treatment prescribed by the zero regime may also differ depending on past treatment. In the two-interval case, this gives blip functions

$$\begin{aligned}
\gamma_1(l_1, a_1) &= E\Big[Y\Big(a_1, d_2^{opt}\big(l_1, a_1, L_2(a_1)\big)\Big) - Y\Big(0_1, d_2^{opt}\big(l_1, 0_1, L_2(0_1)\big)\Big)\big|L_1 = l_1\Big], \\
\gamma_2(\bar{l}_2, \bar{a}_2) &= E\Big[Y(a_1, a_2) - Y(a_1, 0_2)\big|(\bar{L}_2, A_1) = (\bar{l}_2, a_1)\Big].
\end{aligned}$$

At the last (here, the second) interval there are no subsequent treatments, so the blip $\gamma_2(\cdot)$ is simply the expected difference in outcomes for having taken treatment a_2 as compared to the zero regime, 0_2, among people with history $\bar{l}_2, a_1$. At the first interval, the blip $\gamma_1(\cdot)$ is the expected (conditional) difference between the counterfactual outcome that would be

observed if treatment a_1 was given in the first interval, and optimal treatment was given in the second and the the counterfactual outcome that would be observed if no treatment was given in the first interval, and optimal treatment was given in the second interval.

Optimal blip functions are typically parameterized with parameters ψ. For example, a linear blip $\gamma_j(l_j, \bar{a}_j) = a_j(\psi_0 + \psi_1 l_j + \psi_2 l_j^2 + \psi_3 a_{j-1})$ implies that the expected effect of treatment a_j on outcome provided optimal treatment is subsequently given is quadratic in the covariate l_j and linear in the treatment received in the previous interval.

Robins (2004, p. 208) proposes finding the parameters ψ of the optimal blip-to-zero function via g-estimation. For two intervals, let

$$\begin{aligned} H_1(\psi) &= Y - \gamma_1(l_1; \psi) + \Big[\gamma_2(\bar{l}_2, \bar{a}_1, d_2^{opt}; \psi) - \gamma_2(\bar{l}_2, \bar{a}_1; \psi)\Big], \\ H_2(\psi) &= Y - \gamma_2(\bar{l}_2, \bar{a}_2; \psi). \end{aligned}$$

$H_1(\psi)$ and $H_2(\psi)$ equal, in expectation, the potential outcomes $Y\Big(0_1, d_2^{opt}\big(l_1, 0_1, L_2(0_1)\big)\Big)$ and $Y(a_1, 0_2)$, respectively. For the purpose of constructing an estimating procedure, we must specify a function $S_j(a_j) = s_j(a_j, \bar{l}_j, \bar{a}_{j-1}) \in \mathbb{R}^{dim(\psi_j)}$ which depends on variables that are thought to interact with treatment to influence outcome. For example, if the optimal blip at the second interval is linear, $\gamma_2(l_2, \bar{a}_2) = a_2(\psi_0 + \psi_1 l_2 + \psi_2 a_1 + \psi_3 l_2 a_1)$, a common choice for this function is $S_2(a_2) = \frac{\partial}{\partial \psi}\gamma_2(\bar{l}_2, \bar{a}_2) = a_2 \cdot (1, l_2, a_1, l_2 a_1)^T$ since the blip suggests that the effect of the treatment at t_2 on outcome is influenced by covariates at the start of the second interval and by treatment at t_1. Let

$$U_j(\psi) = (H_j(\psi) - E[H_j(\psi)|\bar{L}_j, \bar{A}_{j-1}])\{S_j(A_j) - E[S_j(A_j)|\bar{L}_j, \bar{A}_{j-1}]\}. \tag{1}$$

With $U(\psi) = \sum_{j=1}^{2} U_j(\psi)$, $E[U(\psi, s)] = 0$ is an unbiased estimating equation from which consistent estimates $\hat{\psi}$ of ψ may be found. Robins proves that estimates found by solving (1) are consistent and asymptotically normal under standard regularity conditions provided either $E[H_j(\psi)|\bar{L}_j, \bar{A}_{j-1}]$ or $p_j(a_j|\bar{L}_j, \bar{A}_{j-1})$ (the treatment model used to calculate

$E[S_j(A_j)|\bar{L}_j, \bar{A}_{j-1}]$) is correctly specified. As only one but not both of the nuisance models need be correct, this procedure is said to be *doubly-robust*. The estimators are not in general efficient.

Recursive, closed-form g-estimation Exact solutions to equation (1) can be found when optimal blips are linear in ψ and parameters are not shared between intervals. For details of the recursive procedure, see Robins (2004), Moodie, Richardson and Stephens (2006); an algorithm for the two-interval case is as follows:

1. Estimate the nuisance parameters α_2 in the treatment model $p_2(a_2|\bar{L}_2, A_1; \alpha_2)$.

2. Assume a linear model for the expected counterfactual, $E[H_2(\psi_2)|\bar{L}_2, A_1] = (\bar{L}_2, A_1)\varsigma_2^{\mathrm{T}}$; then the nuisance parameters ς_2 can be written as a function of the data and the unknown parameters, ψ_2.

3. Solve $U_2(\psi_2) = (H_2(\psi_2) - E[H_2(\psi_2)|\bar{L}_2, A_1])\{S_2(A_2) - E[S_2(A_2)|\bar{L}_2, A_1; \hat{\alpha}_2]\} = 0$ to find $\hat{\psi}_2$.

4. Estimate the nuisance parameters α_1 in $p_1(a_1|L_1; \alpha_1)$.

5. Plug $\hat{\psi}_2$ into $H_1(\psi_1, \psi_2)$ so that only ψ_1 is unknown. If $E[H_1(\psi, \hat{\psi}_2)|L_1]$ is assumed to be linear, it can be expressed in terms of ψ_1.

6. Solve $U_1(\psi_1) = (H_1(\psi_1, \hat{\psi}_2) - E[H_1(\psi_1, \hat{\psi}_2)|L_1])\{S_1(A_1) - E[S_1(A_1)|L_1; \hat{\alpha}_1]\} = 0$.

2.2 Asymptotic bias under exceptional laws

The cumulative distribution function of the observed longitudinal data is *exceptional* if, at some interval j, the optimal treatment decision depends on a non-empty subset of covariate and treatment history *and* there is a positive probability that the optimal rule is not unique (Robins 2004). Suppose, for example, that the blip function is linear, depending only on the current covariate and the previous treatment: $\gamma_j(\bar{L}_j, \bar{A}_j; \psi) = A_j(\psi_{j0} + \psi_{j1}L_j + \psi_{j2}A_{j-1})$. The distribution of (L_j, A_{j-1}) is exceptional if it puts positive probability on values of L_j and

A_{j-1} such that $\psi_{j0} + \psi_{j1}L_j + \psi_{j2}A_{j-1} = 0$. It is the combination of three factors that make a law exceptional: (i) the form of the blip model (whether it depends on past treatment or covariates), (ii) the true value of the parameters, and (iii) the distribution of treatments and covariates.

Robins (2004, p. 308) uses a simple scenario to explain the bias at exceptional laws. Suppose a random variable X_i is drawn from a distribution with mean η, and we wish to estimate $I[\eta > 0]\eta$. The MLE of $I[\eta > 0]\eta$ is a transformation of the sample mean, $\bar{X}'_j = I[\bar{X}_j > 0]\bar{X}_j$. In the usual frequentist vein, suppose that under repeated sampling a collection of sample means, $\bar{X}_1, \bar{X}_2, ...\bar{X}_k$ where $E[\bar{X}_j] = \eta$ for all j. The MLE may be biased: $E[\bar{X}'_j] \geq E[\bar{X}_j]$. If $\eta = 0$, approximately half of the sample means will be negative and half positive so that the *average* of the collection would be nearly 0. By truncating the negative statistics to zero to calculate $\bar{X}'_j$, we average over a collection of statistics that are all non-negative and ensure upward bias.

Now suppose the usual two-interval set-up, with linear optimal blips $\gamma_1(l_1, a_1) = a_1(\psi_{10} + \psi_{11}l_1)$ and $\gamma_2(l_2, a_1, a_2) = a_2(\psi_{20}+\psi_{21}l_2+\psi_{22}a_1+\psi_{23}l_2a_1)$. Let $\eta = \psi_{20}+\psi_{21}l_2+\psi_{22}a_1+\psi_{23}l_2a_1$. At the second interval, optimal treatment d_2^{opt} corresponds to $I[(\psi_{20}+\psi_{21}l_2+\psi_{22}a_1+\psi_{23}l_2a_1) > 0]$ (see, for example, Moodie et al. (2006)) so that:

$$\begin{aligned} H_1(\psi) &= Y - \gamma_1(l_1, a_1; \psi_1) + [\gamma_2(l_2, a_1, d_2^{opt}; \hat{\psi}_2) - \gamma_2(l_2, a_1, a_2; \hat{\psi}_2)] \\ &= Y - \gamma_1(l_1, a_1; \psi_1) + [(d_2^{opt} - a_2)(\hat{\psi}_{20} + \hat{\psi}_{21}l_2 + \hat{\psi}_{22}a_1 + \hat{\psi}_{23}l_2a_1)] \\ &= Y - \gamma_1(l_1, a_1; \psi_1) + I[\hat{\eta} > 0]\hat{\eta} - a_2\hat{\eta} \\ &\overset{E}{\geq} Y - \gamma_1(l_1, a_1; \psi_1) + I[\eta > 0]\eta - a_2\eta, \end{aligned}$$

where $\overset{E}{\geq}$ denotes "greater than or equal to in expectation". The quantity $\Big[\gamma_2(l_2, a_1, d_2^{opt}; \psi_2) - \gamma_2(l_2, a_1, a_2; \psi_2)\Big]$ in $H_1(\psi)$ – or more generally, the quantity $\sum_{m>j}\Big[\gamma_m(\bar{l}_m, (\bar{a}_{m-1}, d_m^{opt}); \psi_m) - \gamma_m(\bar{l}_m, \bar{a}_{m-1}; \psi_m)\Big]$ in $H_j(\psi_j)$ – corresponds to $I[\eta > 0]\eta$ which is upwardly biased as observed above. By using the upwardly-biased estimate of $I[\eta > 0]\eta$ in $H_1(\psi)$, some non-

mean zero value is added into the g-estimating equation for ψ_1. The estimating function no longer has expectation zero and hence is asymptotically biased.

Simulated example Consider a more realistic example, with randomly assigned treatments A_1, A_2 that take on the value 0 or 1 with equal probability and data: $L_1 \sim \mathcal{N}(0, 3)$; $L_2 \sim \text{round}(\mathcal{N}(0.75, 0.5))$; and $Y \sim \mathcal{N}(1, 1) - |\gamma_1(l_1, a_1)| - |\gamma_2(\bar{l}_1, a_1, a_2)|$ (see, for example, (Murphy 2003)). Note that the second interval covariate, L_2, is discrete. Suppose that the blip function is linear. For example, taking $\psi_1 = (1.6, 0.0)$ and ψ_2 = (-3.0,2.0,1.0,0.0), whenever $L_2 = A_1 = 1.0$, $\psi_{20} + \psi_{21}L_2 + \psi_{22}A_1 + \psi_{23}L_2A_1 = -3 + 2 \cdot 1 + 1 \cdot 1 + 0 \cdot 1 = 0$ even though not all components of ψ_2 equal zero. Table 1 summarizes the bias (absolute and relative to its standard error) and the coverage of Wald and score confidence intervals for n = 100, 250, and 1000. The bias in the first interval estimate of ψ_{11} is not negligible. Coverage is incorrect in smaller samples in part because of slow convergence the estimated covariance matrix but perhaps also to the exceptional law. Score intervals generally provide coverage closer to the correct level than Wald intervals.

Robins proves that under the scenario of an exceptional law such as the above, estimates $\hat{\psi}_1$ are $n^{1/2}$-consistent but are neither asymptotically normal nor asymptotically unbiased (Robins 2004, p. 313-315). For a graphical interpretation of asymptotic bias and consistency coinciding, we examine bias maps; these plots depict the change in bias of a parameter as sample size and another parameter are varied. Bias maps throughout this article focus on the absolute bias in ψ_{10} as a function of sample size and one of the second interval parameters, ψ_{20}, ψ_{21}, ψ_{22}, or ψ_{23}. Note that there are several combinations of parameters that lead to exceptional laws: as well as (-3.0, 2.0, 1.0, 0.0) with $p(A_1 = L_2 = 1) > 0$, exceptional laws occur with (-2.0, 2.0, 1.0, 0.0) and (-3.0, 3.0, 1.0, 0.0) if $p(A_1 = 0, L_2 = 1) > 0$.

Robins' consistency result gives, for a law and parameters that do not change with n and $\psi_{20} + \psi_{21}L_2 + \psi_{22}A_1 + \psi_{23}L_2A_1 \neq 0$, that estimates of ψ_{10}, ψ_{11} are not asymptotically biased and will converge to the true values. Consistency may be visualized by looking at a horizontal cross-section of any one of the bias maps in Figure 1: eventually, sample size will

be great enough that bias of first-interval estimates is smaller than some fixed number ϵ. However, under some data generating process where $\psi_{20}+\psi_{21}L_2+\psi_{22}A_1+\psi_{23}L_2A_1$ is allowed to *approach* zero (and so parameters "near" the situation of exceptional laws), asymptotic bias exists if sample size grows only proportional to the rate at which the exceptional law is approached. Thus, the proximity to the exceptional law and the sample size both determine the bias of the estimator (Figure 1). Contours of constant bias can be found along the lines on the bias map traced by plotting $\eta = kn^{1/2}$ against n (recall that η is a linear function of each of the second interval parameters). Note that the asymptotic bias is bounded.

In Figure 2, L_2 is not rounded so that $P(\psi_{20} + \psi_{21}L_2 + \psi_{22}A_1 + \psi_{23}L_2A_1 = 0) = 0$ and the law is not exceptional. No bias is observed asymptotically and there is virtually no small-sample bias. For a law to not be exceptional, it is sufficient for the blip function to contain a continuous covariate with a non-zero coefficient. However in finite samples, it is also necessary that the derivative of the blip function with respect to the continuous covariate L_j, $\frac{\partial}{\partial L_j}\gamma_j(\bar{L}_j, \bar{A}_j)$ is large whenever $\gamma_j(\bar{L}_j, \bar{A}_j)$ is in a neighbourhood of zero so as to be able to distinguish unique optimal regime from those that are not unique.

Asymptotic bias calculation It is useful in the ensuing discussion to express the blip as $\gamma_j(\bar{L}_j, \bar{A}_j; \psi_j) = A_j g_j(\bar{L}_j, \bar{A}_{j-1}; \psi_j)$. With this form of the blip function, the optimal rule – and the definition of exceptional laws – can be written in terms of the function $g_j(\bar{L}_j, \bar{A}_{j-1}; \psi_j)$. Then the optimal rule for binary treatment is defined by the indicator $I[g_j(\bar{L}_j, \bar{A}_{j-1}; \psi_j) > 0]$ and exceptional laws are said to exist when $P(g_j(\bar{L}_j, \bar{A}_{j-1}; \psi_j) = 0) > 0$ when $g_j(\bar{L}_j, \bar{A}_{j-1}; \psi_j) \neq g_j(\psi_j)$ so that $g_j(\bar{L}_j, \bar{A}_{j-1}; \psi_j)$ depends on at least one component of $\bar{L}_j, \bar{A}_{j-1}$. Robins (2004, p. 315) calculates the asymptotic bias of $\hat{\psi}_1(\psi_2)$ for linear blips in the special case where exceptional laws occur when $L_j = A_{j_1} = 1$. We generalize the result here:

Suppose a linear blip where $g_2(\bar{L}_2, A_1; \psi_2) \neq g_2(\psi_2)$ so that $g_2(\bar{L}_2, A_1; \psi_2) = 0$ whenever $\sum_{j=0}^{|\psi_2|-1} \psi_{2j} \cdot X_2 = 0$ where $|\psi_2|$ is the dimension of the vector ψ_2; X_2 denotes the matrix of covariates (in our simulations, $X_2 = (1, L_2, A_1, L_2A_1)$); and $\psi_1^\dagger$, $\psi_2^\dagger$ denote the true values

of ψ_1 and ψ_2. Let $\Delta_1^g(\hat{\psi}_2, \psi_2^\dagger) = U_1(\psi_1, \hat{\psi}_2) - U_1(\psi_1, \psi_2^\dagger)$. As indicated by the notation, $\Delta_1^g(\hat{\psi}_2, \psi_2^\dagger)$ does not depend on ψ_1 when blips are linear. For linear blip functions,

$$\begin{aligned}
\Delta_1^g(\hat{\psi}_2, \psi_2^\dagger) &= U_1(\psi_1, \hat{\psi}_2) - U_1(\psi_1, \psi_2^\dagger) \\
&= I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) = 0] \{S_1(A_1) - E[S_1(A_1)|L_1]\} \cdot \\
&\qquad \Big(I[g_2(\bar{L}_2, A_1; \hat{\psi}_2) > 0] - A_2\Big) \Big(g_2(\bar{L}_2, A_1; \hat{\psi}_2) - g_2(\bar{L}_2, A_1; \psi_2^\dagger)\Big) \\
&\qquad + I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) \neq 0] \{S_1(A_1) - E[S_1(A_1)|L_1]\} \cdot \\
&\qquad\quad \Big\{ \Big(I[g_2(\bar{L}_2, A_1; \hat{\psi}_2) > 0] - A_2\Big) g_2(\bar{L}_2, A_1; \hat{\psi}_2) - \\
&\qquad\qquad \Big(I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) > 0] - A_2\Big) g_2(\bar{L}_2, A_1; \psi_2^\dagger) \Big\} \\
&= I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) = 0] \{S_1(A_1) - E[S_1(A_1)|L_1]\} \cdot \\
&\qquad \Big(I[(\hat{\psi}_2 - \psi_2^\dagger) \cdot X > 0] - A_2\Big) (\hat{\psi}_2 - \psi_2^\dagger) \cdot X_2 + \\
&\qquad I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) \neq 0] \{S_1(A_1) - E[S_1(A_1)|L_1]\} \cdot \\
&\qquad\quad \Big\{ \Big(I[g_2(\bar{L}_2, A_1; \hat{\psi}_2) > 0] - A_2\Big) g_2(\bar{L}_2, A_1; \hat{\psi}_2) - \\
&\qquad\qquad \Big(I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) > 0] - A_2\Big) g_2(\bar{L}_2, A_1; \psi_2^\dagger) \Big\},
\end{aligned}$$

where $(\hat{\psi}_2 - \psi_2^\dagger) \cdot X_2 = g_2(\bar{L}_2, A_1; \hat{\psi}_2) - g_2(\bar{L}_2, A_1; \psi_2^\dagger)$. The second equality follows since we may add a zero (i.e., $g_2(\bar{L}_2, A_1; \psi_2^\dagger)$) inside the LHS of the inequality in the indicator $I[g_2(\bar{L}_2, A_1; \hat{\psi}_2) > 0]$. Since $\hat{\psi}_2$ is regular, a Taylor expansion gives

$$\begin{aligned}
n^{1/2}(\hat{\psi}_2 - \psi_2^\dagger) &= -n^{1/2}\mathbb{P}_n \left[\frac{\partial}{\partial \psi_2} U_2(\psi_2^\dagger)\right]^{-1} U_2(\psi_2^\dagger) + o_p(1) \\
&= \mathbb{P}_n (IF_{20}, IF_{21}, ..., IF_{2,|\psi_2|-1})^T + o_p(1) \\
&= \mathbb{P}_n IF_2 + o_p(1),
\end{aligned}$$

which can be substituted into $\Delta_1^g(\hat{\psi}_2, \psi_2^\dagger)$, giving

$$
\begin{aligned}
n^{1/2}\mathbb{P}_n\Delta_1^g(\hat{\psi}_2,\psi_2^\dagger) &= n^{1/2}\mathbb{P}_n\left\{I[g_2(\bar{L}_2,A_1;\psi_2^\dagger)=0]\left\{S_1(A_1)-E[S_1(A_1)|L_1]\right\}\right\}\cdot \\
&\quad \left(I\left[\left(n^{1/2}\mathbb{P}_n(IF_2\cdot X_2)\right)>0\right]-A_2\right)n^{1/2}\mathbb{P}_n(IF_2\cdot X_2)+ \\
&\quad n^{1/2}\mathbb{P}_n\left\{I[g_2(\bar{L}_2,A_1;\psi_2^\dagger)\neq 0]\left\{S_1(A_1)-E[S_1(A_1)|L_1]\right\}\right\}\cdot \\
&\quad \left(I[g_2(\bar{L}_2,A_1;\psi_2^\dagger)>0]-A_2\right)\cdot X_2 n^{1/2}\mathbb{P}_n(IF_2).
\end{aligned}
$$

Let $Z=(Z_1,Z_2)^T=(Z_1,Z_{20},Z_{21},...,Z_{2,|\psi_2|-1})^T$ be multivariate normal with mean zero and variance equaling that of $\left(U_1(\psi_1^\dagger,\hat{\psi}_2),IF_{20}X_{20},IF_{21}X_{21},...,\ IF_{2,|\psi_2|-1}X_{2,|\psi_2|-1}\right)^T$ where $X_2=(X_{20},...,X_{2,|\psi_2|-1})$, and use Z_{2+} to denote $\sum_{j=0}^{|\psi_2|-1}Z_{2j}$. A further Taylor expansion gives

$$
\begin{aligned}
n^{1/2}(\hat{\psi}_1-\psi_1^\dagger) &\approx -n^{1/2}\mathbb{P}_n\left[\frac{\partial}{\partial\psi_1}U_1(\psi_1^\dagger,\psi_2^\dagger)\right]^{-1}U_1(\psi_1^\dagger,\psi_2^\dagger) \\
&= -n^{1/2}\mathbb{P}_n\left[\frac{\partial}{\partial\psi_1}U_1(\psi_1^\dagger,\psi_2^\dagger)\right]\left(U_1(\psi_1^\dagger,\hat{\psi}_2)-\Delta_1^g(\hat{\psi}_2,\psi_2^\dagger)\right).
\end{aligned}
$$

Now $\mathbb{P}_nU_1(\psi_1^\dagger,\hat{\psi}_2)=\mathbb{P}_nZ_1=0$ since $U_1(\psi_1^\dagger,\hat{\psi}_2)$ is an unbiased estimating function, and so the asymptotic bias of $\hat{\psi}_1$ equals

$$
\begin{aligned}
&-n^{1/2}\mathbb{P}_n\left[\frac{\partial}{\partial\psi_1}U_1(\psi_1^\dagger,\psi_2^\dagger)\right]^{-1}\left\{n^{1/2}\mathbb{P}_n\left\{I[g_2(\bar{L}_2,A_1;\psi_2^\dagger)=0]\left\{S_1(A_1)-E[S_1(A_1)|L_1]\right\}\right\}\cdot\right. \\
&\quad \left(I\left[\left(n^{1/2}\mathbb{P}_n(Z_{2+})\right)>0\right]-A_2\right)n^{1/2}\mathbb{P}_n(Z_{2+})+n^{1/2}\mathbb{P}_n\Big\{I[g_2(\bar{L}_2,A_1;\psi_2^\dagger)\neq 0] \\
&\quad \left.\left\{S_1(A_1)-E[S_1(A_1)|L_1]\right\}\right\}\cdot\left(I[g_2(\bar{L}_2,A_1;\psi_2^\dagger)>0]-A_2\right)\cdot X_2n^{1/2}\mathbb{P}_n(Z_2)\Big\}
\end{aligned}
$$

As noted above, Z_2 has mean zero, so this can be further reduced to

$$
\begin{aligned}
&-n^{1/2}\mathbb{P}_n\left[\frac{\partial}{\partial\psi_1}U_1(\psi_1^\dagger,\psi_2^\dagger)\right]^{-1}\left\{n^{1/2}\mathbb{P}_n\left\{I[g_2(\bar{L}_2,A_1;\psi_2^\dagger)=0]\left\{S_1(A_1)-E[S_1(A_1)|L_1]\right\}\right\}\cdot\right. \\
&\quad \left.\left(I\left[\left(n^{1/2}\mathbb{P}_n(Z_{2+})\right)>0\right]-A_2\right)n^{1/2}\mathbb{P}_n(Z_{2+}).\right.
\end{aligned}
$$

When optimal regimes are unique, $I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) = 0] = 0$ and so the asymptotic bias is zero. When optimal regimes are not unique, $I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) = 0] = 1$ and the g-estimates have bounded asymptotic bias at exceptional laws equal to

$$-E\left[\frac{\partial}{\partial \psi_1} U_1(\psi_1^\dagger)\right]^{-1} E\left\{S_1(A_1) - E[S_1(A_1)|L_1]\right\} \cdot E[Z_{2+}|Z_{2+} > 0]P[Z_{2+} > 0].$$

2.3 Detecting exceptional laws

It is not necessary for all parameters to equal zero for a law to be exceptional, unlike, for instance, in weak identification of instrumental variables. Thus a simple F-test is not sufficient to check for exceptional laws. Robins (2004, p. 317) suggests the following steps:

1. Compute g-estimates of ψ_j for $j = 1, ..., K$ and Wald confidence intervals about each parameter.

2. At each interval j, compute $\left|d_j^{opt} \in \text{CI}\right|_i$ as the number of optimal rules possible for each subject i (each set of observed covariates) under all values ψ_j in the confidence set. E.g., if all values of ψ_j in the confidence set give the same optimal rule then $\left|d_j^{opt} \in \text{CI}\right|_i = 1$.

3. If the fraction of the n $\left|d_j^{opt} \in \text{CI}\right|_i$ values that is greater than 1, which will be denoted by $p_{op,j}$, is small (e.g. $p_{op,j} < 0.05$) then we are likely not near an exceptional law and inference based on Wald confidence sets is reliable for earlier intervals $m < j$.

The idea behind this approach is that if there are few instances for which $\left|d_j^{opt} \in \text{CI}\right|_i > 1$, then the confidence set is far away from values of ψ that, in combination with the distribution of covariate and treatment history, would produce an exceptional law. Thus the estimates $\hat{\psi}$ are likely to be regular so that Wald intervals are of the appropriate level. In simulation, Robins' method detects exceptional laws in all samples (Table 1) so that in all cases, score confidence intervals are recommended. In general, the score confidence intervals improve the coverage of the true parameters at both time intervals, particularly for smaller samples.

This method may save the analyst having to find a more computationally-intensive confidence set if score intervals are not recommended, but could itself be quite time-consuming. We suggest here some guidelines to ease this burden:

- Exceptional laws at the very first interval do not affect the regularity of estimates for ψ_j, $j = 2, 3, ..., K$. Hence it is sufficient to consider $\left|d_j^{opt} \in \text{CI}\right|_i$ only for the intervals $2, 3, ..., K$.

- If, at every interval, the Wald confidence interval for the parameter of at least one continuous variable (with no "spikes" or point-masses) excludes zero, the law is likely not exceptional.

- If at any interval, there is no effect of treatment (i.e., the confidence set of ψ_j includes the vector 0), the law is likely exceptional.

If a law is exceptional, Wald confidence intervals will not have the correct coverage, however score confidence intervals will still be uniformly of the correct level (Robins 2004, p. 222). This offers some comfort, though two issues remain. Construction of score intervals can be computationally intensive as they require searches over large volumes of the parameter space. Second, although score intervals provide better confidence sets, estimates have not been corrected and so asymptotic bias remains.

3. BIAS-CORRECTION AT EXCEPTIONAL LAWS: ZEROING INSTEAD OF PLUGGING IN (ZIPI)

3.1 Exceptional laws: an issue of subpopulations

In this section, we propose a modification to g-estimation that not only detects exceptional laws but reduces bias when exceptional laws exist. The method relies on the supposition that – at any interval j – the sample of study participants is drawn from two populations, $\mathcal{M}_j^0$ and $\mathcal{M}_j^1$, where those who are members of $\mathcal{M}_j^0$ are not affected by treatment at interval j and hence do not have a unique optimal treatment decision, while there is only a single

optimal treatment for all those belonging to $\mathcal{M}_j^1$. Let m_j^0 denote the members from the sample drawn from $\mathcal{M}_j^0$, and define m_j^1 similarly. The method that we propose attempts to determine an individual's membership in $\mathcal{M}_j^0$ or $\mathcal{M}_j^1$ and then treats the data collected on individuals from the two groups differently.

In practice, m_j^0 and m_j^1 are unknown and must be estimated; with finite samples, some degree of misclassification will always occur if type I error is held fixed. In this context, type I error is the error of falsely classifying an individual as having a unique rule when in fact he does not, i.e., $P(i \in \hat{m}_j^1 | i \in m_j^0)$, so the null hypothesis is that the law is exceptional. Correspondingly, type II error is the incorrect classification of an individual as having a non-unique optimal regime when in truth his optimal rule is unique: $P(i \in \hat{m}_j^0 | i \in m_j^1)$.

3.2 Zeroing Instead of Plugging In

As seen in section 2.2, bias enters the g-estimating equation of ψ_1 through the addition of the upwardly-biased estimate of $I[\eta > 0]\eta$ into H_1 when η is at or close to zero, for $\eta = g_j(\bar{L}_2, A_1; \psi_2)$. Our algorithm searches for individuals for whom it is likely that $g_j(\bar{L}_2, A_1; \psi_2^\dagger) = 0$ and then, instead of using the estimate $g_j(\bar{L}_2, A_1; \hat{\psi}_2)$ as in the usual recursive g-estimation, plugging in the "better guess" of zero for the estimate of $I[\eta > 0]\eta$. We call this *Zeroing Instead of Plugging In* (ZIPI). Specifically, ZIPI is performed by starting at the last interval, with $j = K$:

1. Compute $\hat{\psi}_j$ (when $j = K$, do this with the usual g-estimation formula).

2. Calculate the covariance matrix of $\hat{\psi}_j$ and use this to check whether a 95% (or 80%, 90%...) Wald confidence interval about $g_j(\bar{l}_{j,i}, \bar{a}_{j-1,i}; \psi_j^\dagger)$ contains zero for each individual in the sample.

3. Let $\mathcal{J}_i$ be the set of time intervals, t_j, for individual i where the confidence interval

about $g_m(\bar{l}_{m,i}, \bar{a}_{m-1,i}; \psi_m^\dagger)$ does not include 0. For individual i, set

$$H^0_{j-1,i} = Y_i - \gamma_{j-1}(\bar{l}_{j-1,i}, \bar{a}_{j-2,i}; \psi_{j-1}) + \sum_{m \in \mathcal{J}_i} [\gamma_m(\bar{l}_{m,i}, (\bar{a}_{m-1,i}, d^{opt}_{m,i}); \psi_m) - \gamma_m(\bar{l}_{m,i}, \bar{a}_{m-1,i}; \psi_{m-1})].$$

I.e., if at any interval m, $m > j$, the confidence interval about $g_m(\bar{l}_{m,i}, \bar{a}_{m-1,i}; \psi_m^\dagger)$ includes zero, $\gamma_m(\bar{l}_{m,i}, (\bar{a}_{m-1,i}, d^{opt}_{m,i}); \psi_m) - \gamma_m(\bar{l}_{m,i}, \bar{a}_{m-1,i}; \psi_{j-1})$ is set to 0 in $H_{\text{mod},j-1,i}$.

4. Using $H^0_{j-1,i}$ as defined above, find the ZIPI estimate for ψ_{j-1} by replacing $H^0_{j-1,i}$ with $H_{j-1,i}$ in equation (1), and repeat steps 2-3 above with j replaced by $j-1$.

Note that the last interval ZIPI estimate, $\hat{\psi}_K$, is equal to the usual recursive g-estimate and is always consistent and asymptotically normal. Table 2 compares g-estimation with ZIPI using different confidence levels to select individuals who are likely to have non-unique optimal rules, keeping to the simulation scenario used in all previous examples. The reduction in bias observed by using 0 instead of the estimate $I[g_j(\bar{L}_{2,i}, A_{1,i}; \hat{\psi}_2) > 0] g_j(\bar{L}_{2,i}, A_{1,i}; \hat{\psi}_2)$ is good at and far from exceptional laws, with performance dependent on the level of the confidence interval used to define $\mathcal{J}_i$. Coverage of the true value of ψ_1 is very nearly correct.

ZIPI falls into the class of pre-test estimators that are frequently used in a variety of statistical settings. Variable or model selection constitutes a common example. Even in these better-understood scenarios, the optimal choice of threshold is a complicated problem (see, for example, (Giles, Lieberman and Giles 1992; Kabaila and Leeb 2006)).

Unfortunately, ZIPI introduces bias where none exists using ordinary recursive g-estimation in areas of the parameter space that are moderately but not very close to the exceptional laws (Figure 3). As sample size increases, the magnitude of the bias and the region of "moderate proximity" to the exceptional law where bias is observed both decrease. The maximum bias observed using ZIPI is smaller than the maximum bias observed using usual g-estimation: 0.033 with 80% CI selection and 0.040 with 95% CI selection as compared to 0.060 with g-estimation (Figures 1,3-4).

3.3 Consistency and asymptotic bias of ZIPI estimates

ZIPI estimates that have all of the desirable asymptotic properties exhibited by ordinary g-estimates at non-exceptional laws, with the added benefit of those properties being more nearly true at exceptional laws also.

Theorem 3.1. *Under non-exceptional laws, ZIPI estimates converge to the usual recursive g-estimates and therefore are asymptotically consistent, unbiased, and normally distributed.*

The proof is trivial. Of even greater interest is the performance of ZIPI estimates under exceptional laws. We begin assuming that m_j^0 and m_j^1 are known for all j. The proof is given in the two-interval context; the generalization to K intervals follows readily.

Theorem 3.2. *When the distribution is exceptional and m_j^0, m_j^1 are known at every interval, ZIPI estimates are asymptotically consistent, unbiased, and normally distributed.*

Proof:

The usual g-estimation equation at the first interval can be decomposed into two components:

$$\mathbb{P}_n(U_1(\psi_1, \psi_2^\dagger)) = \frac{1}{n}\sum_{i=1}^{n} U_{1i}(\psi_1, \hat{\psi}_2) = \frac{1}{n}\sum_{i\in m_2^0} U_{1i}(\psi_1, \hat{\psi}_2) + \frac{1}{n}\sum_{i\in m_2^1} U_{1i}(\psi_1, \hat{\psi}_2).$$

For $i \in m_2^1$, the law is non-exceptional and g-estimates based on that fraction of the estimating equation where $g_2(\bar{l}_{2,i}, a_{1,i}; \psi_2) \neq 0$, i.e., based on $\frac{1}{n}\sum_{i\in m_2^1} U_{1i}(\psi_1, \hat{\psi}_2)$, are consistent and asymptotically unbiased. For $i \in m_2^0$, on the other hand, $E[U_{1i}(\psi_1, \hat{\psi}_2)] \neq 0$ because

$$H_{1,i}(\psi_1, \hat{\psi}_2) = Y_i - \gamma_1(l_{1,i}; \psi_1) + \left[(\hat{d}_{2,i}^{opt} - a_{2,i})(\hat{\psi}_{20} + \hat{\psi}_{21}l_{2,i} + \hat{\psi}_{22}a_{1,i} + \hat{\psi}_{23}l_{2,i}a_{1,i})\right]$$

is an upwardly biased estimator of $H_{1,i}(\psi_1, \psi_2^\dagger)$, as described in section 2.2 and Appendix 1 of Robins (2004). Now consider replacing $\hat{\psi}_{20}+\hat{\psi}_{21}l_{2,i}+\hat{\psi}_{22}a_{1,i}+\hat{\psi}_{23}l_{2,i}a_{1,i}$ with 0 and denote this by $H_{1,i}^0(\psi_1) = Y_i - \gamma_1(l_{1,i}; \psi_1)$, as it does not depend on $\hat{\psi}_2$. Then $|H_{1,i}(\psi_1, \psi_2^\dagger) - H_{1,i}^0(\psi_1)| = 0$, and $|H_{1,i}(\psi_1, \psi_2) - H_{1,i}^0(\psi_1)| \to 0$ as $n \to \infty$. Therefore $H_{\text{mod},1,i}^0(\psi_1)$ is a consistent and

asymptotically unbiased estimator of $H_{1,i}(\psi_1, \psi_2^\dagger)$. Let

$$E[U_{1i}^0(\psi_1, \hat{\psi}_2)] = \frac{1}{n}\Big\{ \sum_{i \in m^0} (H_{1,i}^0(\psi_1) - E[H_j(\psi)|\bar{L}_j, \bar{A}_{j-1}]) \cdot \{S_j(A_j) - E[S_j(A_j)|\bar{L}_j, \bar{A}_{j-1}]\}$$
$$+ (H_{1,i}(\psi_1) - E[H_j(\psi)|\bar{L}_j, \bar{A}_{j-1}]) \cdot \{S_j(A_j) - E[S_j(A_j)|\bar{L}_j, \bar{A}_{j-1}]\}\Big\}.$$

Then $E[U_{1i}^0(\psi_1, \hat{\psi}_2)] \to 0$ for $i \in m_2^0$ and estimates of ψ_1 using $H_1^0(\psi_1)$ (i.e., ZIPI estimates) are consistent and asymptotically unbiased for any ψ in the parameter space and at any law, exceptional or otherwise.

□

Of course, when using the ZIPI procedure in practice, m_j^0 and m_j^1 are unknown and must be estimated; with fixed type I error, some degree of misclassification will always occur, assuming both sets are non-empty. Recall type I error is the error of using the estimate $\hat{\psi}_2$ rather than 0 in $H_{1,i}(\psi_1)$ and type II error that of using a zero when a plug-in estimate would have been desired if m_j^0 and m_j^1 were known.

Theorem 3.3. *When the distribution law is exceptional and m_j^0, m_j^1 are estimated, ZIPI estimates have smaller asymptotic bias than the usual recursive g-estimates.*

Proof:

Misclassification arising from estimating m_j^0 and m_j^1 results in asymptotic bias in ZIPI estimates at exceptional laws; however the magnitude of the bias is less than that of g-estimates. If classification of individuals into m_j^0, m_j^1 is performed using $1-\epsilon$ confidence intervals, then with fixed type I error ϵ, as $n \to \infty$, the type II error will be driven to zero and so the number of people falsely classified as belonging to $\mathcal{M}_j^0$ will be driven to zero. However, $100\epsilon\%$ of the people with non-unique treatment rules will be falsely classified as belonging to $\mathcal{M}_j^1$.

Recall that we have assumed linear blips, so $\sum_{j=0}^{|\psi_2|-1} \psi_{2j} X_2 = 0$ when laws are exceptional and $\psi_1^\dagger$, $\psi_2^\dagger$ denote the true values of ψ_1, ψ_2.

Define $\Delta_1^z(\hat{\psi}_2, \psi_2^\dagger) = U_1^0(\psi_1, \hat{\psi}_2) - U_1(\psi_1, \psi_2^\dagger)$ to be the quantity in ZIPI estimation corresponding to $\Delta_1^g(\hat{\psi}_2, \psi_2^\dagger)$ in the usual g-estimation procedure. Assume that variables in the

covariates and treatment history are discrete so that, if $g_2(\bar{L}_{2,i}, A_{1,i}; \psi_2) \neq 0$, $g_2(\bar{L}_{2,i}, A_{1,i}; \psi_2)$ cannot be arbitrarily close to zero and assume further that sample size is large enough that power to detect whether $g_2(\bar{L}_{2,i}, A_{1,i}; \psi_2) \neq 0$ is effectively equal to 1. Thus, the only form of misclassification error in ZIPI is type I, i.e., individual i in with $\mathcal{M}_2^0$ will incorrectly be classified as $i \in \hat{m}_2^1$. (This assumption requires that the distribution of $\bar{L}_2, A_1$ is such that $|g_2(\bar{L}_2, A_1; \psi_2)| \geq \mu$ for some $\mu \geq 0$ for all individuals for whom the optimal rule at the second interval is unique). Then, asymptotically, $\Delta_1^z(\hat{\psi}_2, \psi_2^\dagger)$ consists of three components:

i. If $I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) = 0]$ and $i \in \hat{m}_j^0$, $\Delta_1^z(\hat{\psi}_2, \psi_2^\dagger) = 0$;

ii. If $I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) = 0]$ and $i \in \hat{m}_j^1$,
$\Delta_1^z(\hat{\psi}_2, \psi_2^\dagger) = \{S_1(A_1) - E[S_1(A_1)|L_1]\} \cdot \Big(I[(\hat{\psi}_2 - \psi_2^\dagger) \cdot X_2 > 0] - A_2\Big) (\hat{\psi}_2 - \psi_2^\dagger) \cdot X_2$;

iii. If $I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) \neq 0]$,
$\Delta_1^z(\hat{\psi}_2, \psi_2^\dagger) = \{S_1(A_1) - E[S_1(A_1)|L_1]\} \cdot \Big\{ \Big(I[g_2(\bar{L}_2, A_1; \hat{\psi}_2) > 0] - A_2\Big) g_2(\bar{L}_2, A_1; \hat{\psi}_2) -$
$\Big(I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) > 0] - A_2\Big) g_2(\bar{L}_2, A_1; \psi_2^\dagger)\Big\}$.

As observed in section 2.2, the asymptotic bias is proportional to $\Delta_1^z(\hat{\psi}_2, \psi_2^\dagger)$. In ZIPI, there is no difference between the estimating equation $U_1(\psi_1, \psi_2^\dagger)$ using the true value $\psi_2^\dagger$ and the estimating equation that uses zero instead of $g_2(\bar{L}_2, A_1; \hat{\psi}_2)$ which gives component (i); unfortunately, our classification procedure incorrectly fails to use zeros for $100\epsilon\%$ of the sample with $\mathcal{M}_1^0$, leaving component (ii). Therefore rewrite this as

$$\Delta_1^z(\hat{\psi}_2, \psi_2^\dagger) = \epsilon I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) = 0] \Delta_1(\hat{\psi}_2, \psi_2^\dagger) \; + \; I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) \neq 0] \Delta_1(\hat{\psi}_2, \psi_2^\dagger).$$

It follows directly that the asymptotic bias of ZIPI estimates at exceptional laws is

$$-\epsilon E\left[\frac{\partial}{\partial \psi_1} U_1(\psi_1^\dagger)\right]^{-1} E\left\{S_1(A_1) - E[S_1(A_1)|L_1]\right\} \cdot E[Z_{2+}|Z_{2+} > 0] P[Z_{2+} > 0].$$

□

In words: at exceptional laws, ZIPI estimates using fixed-level confidence intervals to

classify individuals are asymptotically biased however the magnitude of the bias is less than that of g-estimates. ZIPI estimates where m_j^0, m_j^1 are estimated lie between the unbiased estimates that would be found if m_j^0 and m_j^1 were known and the usual g-estimates.

In finite samples, both type I *and* type II error will be observed, which leads to additional bias in the ZIPI estimates. Whereas, asymptotically, $\Delta_1^z(\hat{\psi}_2, \psi_2^\dagger)$ can be written in terms of three components, when sample size is not sufficient for power to equal 1, $\Delta_1^z(\hat{\psi}_2, \psi_2^\dagger)$ must be expressed in terms of *four* components. We omit details for brevity, but note that if type II error is denoted by β, then $\Delta_1^z(\hat{\psi}_2, \psi_2^\dagger)$ can be approximated in finite samples by:

$$\epsilon I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) = 0]\Delta_1(\hat{\psi}_2, \psi_2^\dagger) \ + \ \beta I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) \neq 0]\Delta_1(\hat{\psi}_2, \psi_2^\dagger).$$

The approximate bias of ZIPI estimates in finite samples is then

$$\begin{aligned} &- \epsilon E\left[\frac{\partial}{\partial \psi_1} U_1(\psi_1^\dagger)\right]^{-1} E\left\{S_1(A_1) - E[S_1(A_1)|L_1]\right\} \cdot E[Z_{2+}|Z_{2+} > 0]P[Z_{2+} > 0] \\ &\quad - \beta E\left\{S_1(A_1) - E[S_1(A_1)|L_1]\right\} \cdot E\left\{\left(I[g_2(\bar{L}_2, A_1; \psi_2^\dagger) > 0] - A_2\right) g_2(\bar{L}_2, A_1; \psi_2^\dagger)\right\}. \end{aligned}$$

When covariates are integer-valued as in our simulations, there is good separation of $g_2(\bar{L}_2, A_1; \hat{\psi}_2)$ between those people with unique rules and those without; in this instance, simulations suggest that choosing a relatively large (e.g., 0.20) type I error provides a good balance between bias reduction *at* and *near* exceptional laws (Figures 3-4).

3.4 ZIPI versus simultaneous (non-recursive) g-estimation

When parameters are not shared between intervals, recursive estimation is the obvious choice. With ZIPI g-estimation proving both theoretically and in simulation to be asymptotically less biased than recursive g-estimation, it is the obvious choice. However the question then arises as to whether it is better than simultaneous (non-recursive) g-estimation when parameters are shared. Recall that Theorem 3.3 pertains to recursive, not simultaneous, estimation.

If parameters are believed to be shared, that is, $\psi_j = \psi_{j'}$ for some $j \neq j'$, recursive

methods may still be used by pretending otherwise, using a recursive form of g-estimation (the usual or ZIPI) then averaging the interval-specific estimates of ψ. Inverse covariance weighted recursive g-estimates have the advantage of being easier to compute than simultaneous estimates and are doubly robust under non-exceptional laws (Robins 2004, p. 221). On the other hand, simultaneous g-estimation can incorporate the stationarity information into the estimating equations and search for estimates of the (shared) parameters directly without needing to average estimates from different intervals.

Robins (2004, p. 221-222) recommends using simultaneous g-estimation at non-exceptional laws when there are few subjects who change treatment from one interval to the next since, with few changes of treatment at interval j, estimates $\hat{\psi}_j$ are extremely variable and so the inverse variance weighting will introduce bias in the estimating equation for ψ_m, $m < j$, in a manner not unlike the bias caused by exceptional laws.

If there is a moderate number of treatment changes from one interval to the next (recursive) ZIPI estimation may be useful for the purposes of bias reduction in large samples. Figure 5 compares ZIPI with 80% confidence intervals used to estimate membership in m_j^0 or m_j^1, recursive g-estimation, and simultaneous g-estimation for data as before: $L_1 \sim \mathcal{N}(0, 3)$; $L_2 \sim \text{round}(\mathcal{N}(0.75, 0.5))$; $Y \sim \mathcal{N}(1, 1) - |\gamma_1(L_1, A_1)| - |\gamma_2(L_2, A_1, A_2)|$ but we now impose stationarity so that $\gamma_1(L_1, A_1) = A_1(\psi_0 + \psi_1 L_1)$ and $\gamma_2(L_2, A_1, A_2) = A_2(\psi_0 + \psi_1 L_2 + \psi_{22} A_1 + \psi_{23} L_2 A_1)$. Note the sharing of ψ_0 and ψ_1 in blips of both intervals.

Simultaneous and recursive g-estimation are similar in terms of both bias and variability. Relative to simultaneous g-estimation, the recursive methods are easier to implement and do not suffer from the possibility of multiple roots under linear blip functions. At non-exceptional laws, ZIPI is asymptotically equivalent to the doubly-robust recursive g-estimation which, as noted above, is locally efficient.

These simulations demonstrate, however, that inverse-covariance weighted ZIPI estimates are not always less biased than simultaneous g-estimates even in quite large samples. ZIPI performs much better when $\psi_1 = 2$ as compared to when $\psi_1 = 3$. Parameter sharing implies that "biased" entries in the g-estimating equation make up a smaller proportion of the total

number of data points used to simultaneously estimate the shared decision rule parameters.

3.5 ZIPI versus Iterative Minimization for Optimal Regimes

As a final word, it is worthy of note that the asymptotic bias due to exceptional laws is present in Iterative Minimization for Optimal Regimes (IMOR) (Murphy 2003) method as well (Table 3). This is not surprising since IMOR and g-estimation are very closely related (Moodie et al. 2006). Nevertheless, when parameters are shared and the sample is not very large, IMOR may prove a better alternative than ZIPI for the reasons noted in the previous subsection.

3.6 Comparing ZIPI to the MLE in a simple, normal example

Recall the simple scenario of section 2.2: random variables X_i have mean η, and we wish to estimate $I[\eta > 0]\eta$. The MLE of $I[\eta > 0]\eta$ is $\bar{X}' = I[\bar{X} > 0]\bar{X}$, which is biased when η is near zero. We shall make this slightly more interesting by allowing $X_i \sim \mathcal{N}(\eta_i, 1)$.

For $i = 1, ..., n$, let L_i and A_i be binary variables in a 2×2 factorial design equally distributed among study participants and let $\eta_i = \psi_0 + \psi_1 L_i + \psi_2 A_i$. Then

$$\eta = E[\eta_i] = 0.25 \cdot \psi_0 + 0.25 \cdot (\psi_0 + \psi_1) + 0.25 \cdot (\psi_0 + \psi_2) + 0.25 \cdot (\psi_0 + \psi_1 + \psi_2).$$

For $(\psi_0, \psi_1, \psi_2) = (-1, 1, 1)$, $\eta = 0$, and $E[I[\eta_i > 0]\eta_i] = 0.75 \cdot 0 + 0.25 \cdot (\psi_0 + \psi_1 + \psi_2) = 0.25$. Let m^0 denote those members of the sample for whom $L_i \neq A_i$ so that $\eta_i = 0$, and m^1 denote those members for whom $L_i = A_i$ (and so $\eta_i \in \{-1, 1\}$).

For both maximum likelihood and ZIPI, we regress X on L and A in a simple linear model and to find $\hat{\eta}_i = \hat{\psi}_0 + \hat{\psi}_1 L_i + \hat{\psi}_2 A_i$ and estimate its standard error. We then compare the MLE, $\frac{1}{n}\sum_{i=1}^{n}\{I[\hat{\eta}_i > 0]\hat{\eta}_i\}$, with the ZIPI estimate, $\frac{1}{n}\sum_{i=1}^{n}\{I[\hat{\eta}_i^0 > 0]\hat{\eta}_i^0\}$, where $\hat{\eta}_i^0$ is equal to $\hat{\eta}_i$ whenever a $(1-\epsilon)\%$ confidence interval excludes 0 and is equal to zero otherwise. Equivalently, we may write the ZIPI estimate as $\frac{1}{n}\sum_{i \in \hat{m}^1}\{I[\hat{\eta}_i > 0]\hat{\eta}_i\} + \frac{1}{n}\sum_{i \in \hat{m}^0} 0$. Mean estimates and simulated standard errors are in Table 4. At all sample sizes and for all three levels of confidence intervals examined, ZIPI outperforms maximum likelihood.

4. DISCUSSION

In this paper, we have discussed the asymptotic bias of optimal decision rule parameter estimates in the presence of exceptional laws. We introduced time-saving suggestions to reduce the computational burden of trying to detect exceptional laws using Robins' method.

We introduced Zeroing Instead of Plugging In (ZIPI), a recursive g-estimation in which zeros are used rather than estimates in those individuals who are likely to have non-unique optimal treatments at particular intervals. ZIPI is consistent at all laws and though asymptotically biased at exceptional laws, its bias is smaller than that of recursive g-estimates when parameters are non-stationary. Therefore whenever parameters are not shared between intervals or the variability of parameters is not very large, ZIPI is to be preferred over recursive g-estimation if there is a possibility of laws being exceptional. Future work will include determination of the best choice of threshold in ZIPI.

Asymptotic bias due to exceptional laws cannot be avoided by using methods that do not require plug-in estimates such as simultaneous g-estimation or IMOR, even when parameters are stationary. When the sample is very large and there are moderate to large numbers of people receiving treatment in each interval, inverse-covariance weighted ZIPI estimates may provide a good alternative to either simultaneous g-estimation or IMOR, though this is not guaranteed. Any potential reduction in bias must always be weighed against the loss of efficiency that can result from using ZIPI.

We conclude by observing that one of g-estimation's primary points of appeal is that, unlike traditional approaches such as dynamic programming, estimates are guaranteed to be asymptotically unbiased under the null hypothesis of no treatment effect even if the distribution law is mis-specified provided the true law is not exceptional. However for a blip function that includes at least one component of treatment and covariate history, *under the hypothesis of no treatment effect, every distribution is an exceptional law* and therefore g-estimation is asymptotically biased under the null for all such blip functions/laws. In such circumstances, we may consider fitting a "null" blip model which does not include any

covariates and hence is not at an exceptional law.

REFERENCES

Giles, D. E. A., Lieberman, O., and Giles, J. A. (1992), "The Optimal Size of a Preliminary Test of Linear Restrictions in a Misspecified Regression Model," *Journal of the American Statistical Association*, 87, 1153–1157.

Kabaila, P., and Leeb, H. (2006), "On the Large-Sample Minimal Coverage Probability of Confidence Intervals After Model Selection," *Journal of the American Statistical Association*, 101, 619–629.

Moodie, E. E. M., Richardson, T. S., and Stephens, D. A. (2006), "Demystifying Optimal Dynamic Treatment Regimes," *Biometrics* (in press).

Murphy, S. A. (2003), "Optimal Dynamic Treatment Regimes," *Journal of the Royal Statistical Society, Ser. B*, 65, 331–366.

Robins, J. M. (1986), "A New Approach to Causal Inference in Mortality Studies with Sustained Exposure Periods - Application to Control of the Healthy Worker Survivor Effect," *Mathematical Modelling*, 7, 1393–1512.

—— (1994), "Correcting for Non-Compliance in Randomized Trials using Structural Nested Mean Models," *Communincations in Statistics*, 23, 2379–2412.

—— (1997), Latent Variable Modeling and Applications to Causality, Springer, New York, pp. 69–117.

—— (2004), Optimal Structural Nested Models for Optimal Sequential Decisions, in *Proceedings of the Second Seattle Symposium on Biostatistics*, eds. D. Y. Lin, and P. J. Heagerty, Springer, New York.

Rubin, D. B. (1978), "Bayesian Inference for Causal Effects: The role of randomization," *Annals of Statistics*, 6, 34–58.

Tables

Table 1: Summary statistics of g-estimates and Robins' (2004) proposed method of detecting exceptional laws from 1000 data-sets of sample sizes 100, 250 and 1000.

n	ψ	$\hat{\psi}$	|Bias|	$\lvert t\rvert^a$	rMSE	Cov.^{b}	Cov.^{c}
100							
	$\psi_{10} =$ 1.6	1.693	0.093	0.318	0.386	93.7	94.9
	$\psi_{11} =$ 0.0	0.001	0.001	0.007	0.093	92.4	94.0
	$\psi_{20} =$ -3.0	-2.997	0.003	0.023	0.660	89.0	93.5
	$\psi_{21} =$ 2.0	2.000	0.000	0.052	0.726	89.3	92.4
	$\psi_{22} =$ 1.0	1.007	0.007	0.009	0.969	91.5	92.5
	$\psi_{23} =$ 0.0	-0.016	0.016	0.053	1.103	91.1	92.3
	$p_{op,2}$: Mean (Range)^d: 76.0 (28.0, 100.0)						
250							
	$\psi_{10} =$ 1.6	1.657	0.057	0.306	0.245	95.1	94.2
	$\psi_{11} =$ 0.0	-0.001	0.001	0.027	0.057	94.2	94.7
	$\psi_{20} =$ -3.0	-3.001	0.001	0.001	0.424	92.5	94.5
	$\psi_{21} =$ 2.0	2.006	0.006	0.017	0.460	92.0	93.5
	$\psi_{22} =$ 1.0	1.004	0.004	0.003	0.615	93.5	93.6
	$\psi_{23} =$ 0.0	-0.001	0.001	0.008	0.693	94.2	94.7
	$p_{op,2}$: Mean (Range): 63.8 (30.4, 92.0)						
1000							
	$\psi_{10} =$ 1.6	1.629	0.029	0.321	0.121	95.5	95.4
	$\psi_{11} =$ 0.0	-0.001	0.001	0.026	0.028	96.1	95.2
	$psi_{20} =$ -3.0	-3.001	0.001	0.017	0.203	95.3	95.3
	$\psi_{21} =$ 2.0	2.002	0.002	0.007	0.225	94.2	94.7
	$\psi_{22} =$ 1.0	1.000	0.000	0.002	0.298	95.4	94.9
	$\psi_{23} =$ 0.0	-0.006	0.006	0.032	0.348	94.5	95.2
	$p_{op,2}$: Mean (Range): 35.7 (28.2, 70.2)						

[a] $|t| = (\hat{\psi} - \psi)/\text{SE}(\hat{\psi})$
[b] Coverage of 95% Wald confidence intervals
[c] Coverage of 95% score confidence intervals
[d] Estimated proportion of sample with non-unique optimal rules (see section 2.3)

Table 2: Parameter estimates under exceptional laws: summary statistics of g-estimation using Zeroing Instead of Plugging In (ZIPI) and the usual g-estimates from 1000 data-sets of sample sizes 250, 500 and 1000 using 80%, 90%, and 95% CIs for ZIPI. Second interval parameter estimates are omitted for all but 80% CI exclusions since the exclusions affect only first-interval parameters when using ZIPI (see section 3.2).

CI		ZIPI					Usual g-estimation				
exclusion	ψ	$\hat{\psi}$	SE	$\lvert t\rvert^a$	rMSE	Cov.^{b}	$\hat{\psi}$	SE	$\lvert t\rvert^a$	rMSE	Cov.^{b}
$n = 250$											
80%	$\psi_{10} =$ 1.6	1.631	0.167	0.174	0.229	94.3	1.659	0.179	0.322	0.243	95.7
	$\psi_{11} =$ 0.0	-0.001	0.042	0.031	0.057	94.3	-0.001	0.042	0.029	0.057	94.6
	$\psi_{20} =$ -3.0	-2.991	0.299	0.016	0.413	92.8	-2.991	0.299	0.016	0.413	93.3
	$\psi_{21} =$ 2.0	1.977	0.325	0.095	0.461	92.3	1.977	0.325	0.095	0.461	92.4
	$\psi_{22} =$ 1.0	0.992	0.444	0.020	0.606	94.2	0.992	0.444	0.020	0.606	94.3
	$\psi_{23} =$ 0.0	0.014	0.505	0.015	0.696	93.1	0.014	0.505	0.015	0.696	93.4
90%	$\psi_{10} =$ 1.6	1.622	0.165	0.126	0.227	94.3					
	$\psi_{11} =$ 0.0	-0.001	0.042	0.031	0.057	94.1					
95%	$\psi_{10} =$ 1.6	1.617	0.164	0.098	0.225	94.4					
	$\psi_{11} =$ 0.0	-0.001	0.042	0.030	0.057	94.2					
$n = 500$											
80%	$\psi_{10} =$ 1.6	1.615	0.118	0.108	0.164	92.6	1.636	0.126	0.281	0.172	94.2
	$\psi_{11} =$ 0.0	0.000	0.030	0.003	0.041	94.8	0.000	0.030	0.002	0.041	94.7
	$\psi_{20} =$ -3.0	-2.998	0.212	0.006	0.288	93.5	-2.998	0.212	0.006	0.288	94.1
	$\psi_{21} =$ 2.0	1.993	0.233	0.047	0.317	93.7	1.993	0.233	0.047	0.317	93.9
	$\psi_{22} =$ 1.0	0.996	0.314	0.013	0.432	93.9	0.996	0.314	0.013	0.432	94.3
	$\psi_{23} =$ 0.0	-0.002	0.360	0.018	0.495	93.3	-0.002	0.360	0.018	0.495	93.7
90%	$\psi_{10} =$ 1.6	1.609	0.117	0.062	0.161	93.2					
	$\psi_{11} =$ 0.0	0.000	0.030	0.003	0.041	94.9					
95%	$\psi_{10} =$ 1.6	1.605	0.116	0.032	0.160	93.3					
	$\psi_{11} =$ 0.0	0.000	0.030	0.002	0.041	94.9					
$n = 1000$											
80%	$\psi_{10} =$ 1.6	1.616	0.083	0.180	0.114	94.9	1.632	0.090	0.348	0.121	95.3
	$\psi_{11} =$ 0.0	-0.001	0.021	0.041	0.028	94.8	-0.001	0.021	0.042	0.028	95.1
	$\psi_{20} =$ -3.0	-2.995	0.150	0.025	0.207	93.4	-2.995	0.150	0.025	0.207	93.7
	$\psi_{21} =$ 2.0	1.997	0.165	0.030	0.225	93.8	1.997	0.165	0.030	0.225	94.0
	$\psi_{22} =$ 1.0	0.992	0.222	0.036	0.305	94.2	0.992	0.222	0.036	0.305	94.5
	$\psi_{23} =$ 0.0	0.003	0.255	0.003	0.348	94.8	0.003	0.255	0.003	0.348	95.0
90%	$\psi_{10} =$ 1.6	1.611	0.083	0.127	0.113	95.1					
	$\psi_{11} =$ 0.0	-0.001	0.021	0.042	0.028	94.8					
95%	$\psi_{10} =$ 1.6	1.608	0.082	0.089	0.112	95.1					
	$\psi_{11} =$ 0.0	-0.001	0.021	0.042	0.028	94.7					

[a] $|t| = (\hat{\psi} - \psi)/\text{SE}(\hat{\psi})$

[b] Coverage of 95% Wald confidence intervals

Table 3: Parameter estimates under exceptional laws: summary statistics of g-estimation using Zeroing Instead of Plugging In (ZIPI) and the IMOR from 1000 data-sets of sample sizes 250, 500 and 1000 using 80% CIs for ZIPI.

		ZIPI				IMOR			
n	ψ	$\hat{\psi}$	SE	rMSE	Cov.a	$\hat{\psi}$	SE	rMSE	Cov.a
250									
	$\psi_{10} =$ 1.6	1.628	0.167	0.233	93.3	1.643	0.193	0.259	95.9
	$\psi_{11} =$ 0.0	0.002	0.042	0.058	93.6	-0.002	0.043	0.059	92.7
	$\psi_{20} =$ -3.0	-3.017	0.296	0.409	92.4	-3.042	0.391	0.524	95.7
	$\psi_{21} =$ 2.0	2.011	0.325	0.451	92.1	2.040	0.381	0.516	94.1
	$\psi_{22} =$ 1.0	1.022	0.442	0.606	94.6	1.046	0.496	0.675	94.4
	$\psi_{23} =$ 0.0	-0.020	0.507	0.693	94.9	-0.047	0.516	0.699	95.1
500									
	$\psi_{10} =$ 1.6	1.625	0.118	0.160	95.0	1.640	0.137	0.180	96.7
	$\psi_{11} =$ 0.0	0.000	0.030	0.041	94.4	0.000	0.030	0.041	94.6
	$\psi_{20} =$ -3.0	-2.993	0.212	0.288	93.4	-2.997	0.276	0.369	95.5
	$\psi_{21} =$ 2.0	1.995	0.234	0.319	93.6	2.000	0.269	0.362	94.5
	$\psi_{22} =$ 1.0	0.994	0.313	0.419	95.7	1.000	0.349	0.470	96.0
	$\psi_{23} =$ 0.0	0.009	0.361	0.482	95.8	0.002	0.363	0.488	96.3
1000									
	$\psi_{10} =$ 1.6	1.608	0.083	0.116	92.7	1.619	0.097	0.129	95.3
	$\psi_{11} =$ 0.0	0.000	0.021	0.028	96.2	0.000	0.021	0.028	96.2
	$\psi_{20} =$ -3.0	-3.010	0.151	0.206	93.8	-3.020	0.197	0.265	95.9
	$\psi_{21} =$ 2.0	2.009	0.166	0.227	94.0	2.017	0.191	0.258	95.2
	$\psi_{22} =$ 1.0	1.015	0.222	0.305	94.0	1.023	0.249	0.342	94.0
	$\psi_{23} =$ 0.0	-0.013	0.256	0.347	94.0	-0.018	0.258	0.353	94.3

a Coverage of 95% Wald confidence intervals

Table 4: A comparison of the MLE and Zeroing Instead of Plugging In (ZIPI) over 1000 data-sets of various sample sizes where $X_i \sim \mathcal{N}(\eta_i, 1)$. With $\eta_i = \psi_0 + \psi_1 L_i + \psi_2 A_i$ and $\eta = E[\eta_i]$ we wish to estimate $I[\eta > 0]\eta$.

	MLE	ZIPI (80%)	ZIPI (90%)	ZIPI (95%)
n	Estimate (SE)	Estimate (SE)	Estimate (SE)	Estimate (SE)
20	0.328 (0.143)	0.275 (0.151)	0.252 (0.157)	0.228 (0.162)
50	0.290 (0.091)	0.263 (0.085)	0.255 (0.082)	0.249 (0.081)
100	0.285 (0.065)	0.266 (0.062)	0.259 (0.056)	0.254 (0.051)
250	0.268 (0.041)	0.257 (0.038)	0.253 (0.036)	0.250 (0.033)
500	0.265 (0.029)	0.257 (0.027)	0.254 (0.025)	0.253 (0.024)
1000	0.260 (0.020)	0.254 (0.019)	0.252 (0.017)	0.250 (0.016)

Figures: Titles and Artwork

Figure 1. Absolute bias of $\hat{\psi}_{10}$ as sample size varies at and near exceptional laws. In each figure, one of $\psi_{20}, \psi_{21}, \psi_{22}, \psi_{23}$ is varied while the remaining are fixed at one of -3.0, 2.0, 1.0, or 0.0, respectively. Note that there are several combinations of parameters that lead to exceptional laws: as well as (-3.0, 2.0, 1.0, 0.0) with $p(A_1 = L_2 = 1) > 0$, exceptional laws occur with (-2.0, 2.0, 1.0, 0.0) and (-3.0, 3.0, 1.0, 0.0) if $p(A_1 = 0,\ L_2 = 1) > 0$.

Figure 2. Lack of bias of $\hat{\psi}_{10}$ as sample size varies for non-exceptional laws. In each figure, one of $\psi_{20}, \psi_{21}, \psi_{22}, \psi_{23}$ is varied while the other three parameters are fixed at one of -3.0, 2.0, 1.0, or 0.0, respectively. L_2 is continuous Normal so that $p(A_1 = L_2 = 1) = p(A_1 = 0,\ L_2 = 1) = 0$ and its co-efficient in the blip function is sufficiently larger than zero so that no optimal rules appear to be non-unique in samples of size 200 or greater.

Figure 3. Absolute bias of $\hat{\psi}_{10}$ as sample size varies at and near exceptional laws using Zeroing Instead of Plugging In (ZIPI) when 95% confidence intervals suggest exceptional laws. In each figure, one of $\psi_{20}, \psi_{21}, \psi_{22}, \psi_{23}$ is varied while the other three parameters are fixed at one of -3.0, 2.0, 1.0, or 0.0, respectively.

Figure 4. Absolute bias of $\hat{\psi}_{10}$ as sample size varies at and near exceptional laws using Zeroing Instead of Plugging In (ZIPI) when 80% confidence intervals suggest exceptional laws. In each figure, one of $\psi_{20}, \psi_{21}, \psi_{22}, \psi_{23}$ is varied while the other three parameters are fixed at one of -3.0, 2.0, 1.0, or 0.0, respectively.

Figure 5. Absolute bias in $\hat{\psi}_0$ as n varies about exceptional laws assuming stationarity: inverse-covariance weighted ZIPI estimates with 80% CIs, inverse-covariance weighted recursive g-estimates, and simultaneous g-estimates [left, center, and right columns]. Rows: the true value of one of $\psi_0, \psi_1, \psi_{22}, \psi_{23}$ is varied; the rest are fixed at -3.0, 2.0, 1.0, or 0.0, respectively.

Figure 1:

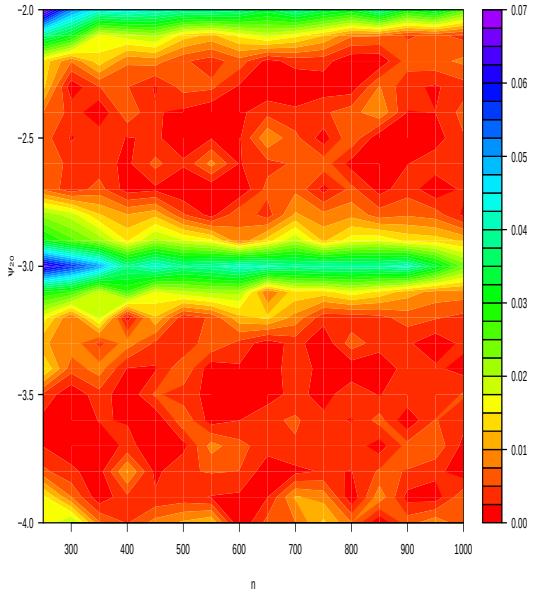

(a) Bias in $\hat{\psi}_{10}$, varying ψ_{20}

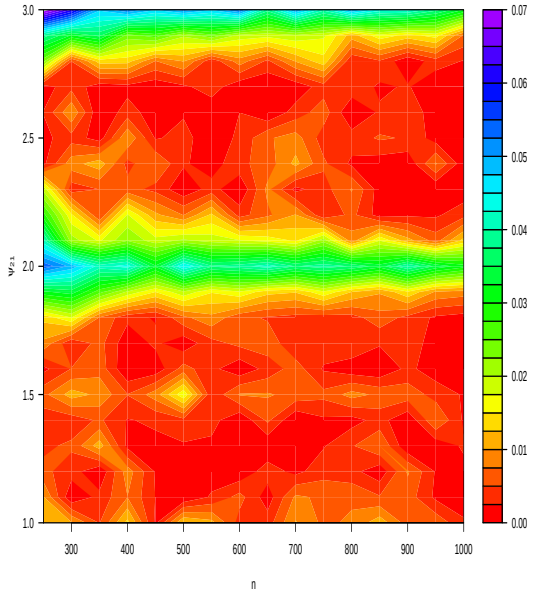

(b) Bias in $\hat{\psi}_{10}$, varying ψ_{21}

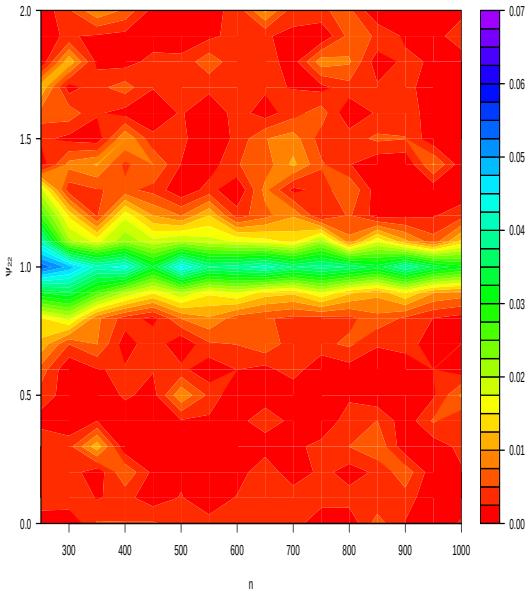

(c) Bias in $\hat{\psi}_{10}$, varying ψ_{22}

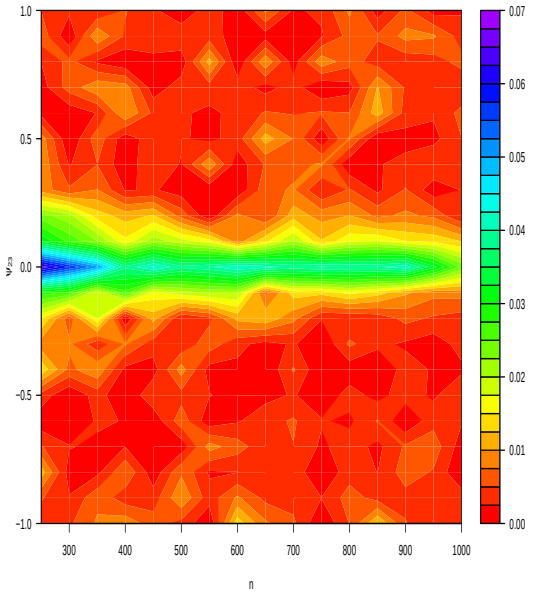

(d) Bias in $\hat{\psi}_{10}$, varying ψ_{23}

Figure 2:

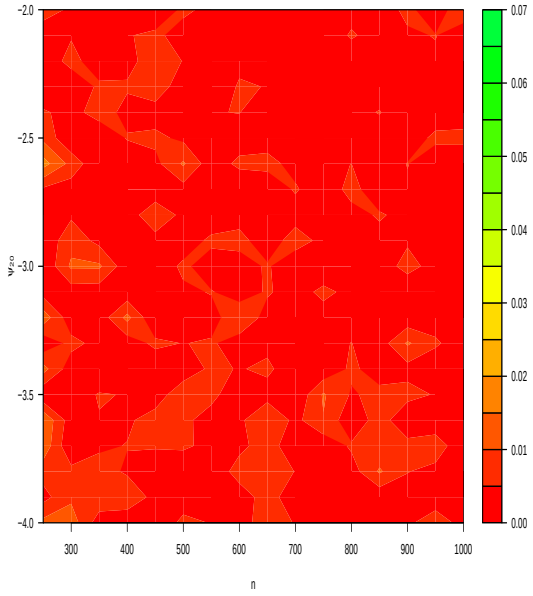

(a) Bias in $\hat{\psi}_{10}$, varying ψ_{20}

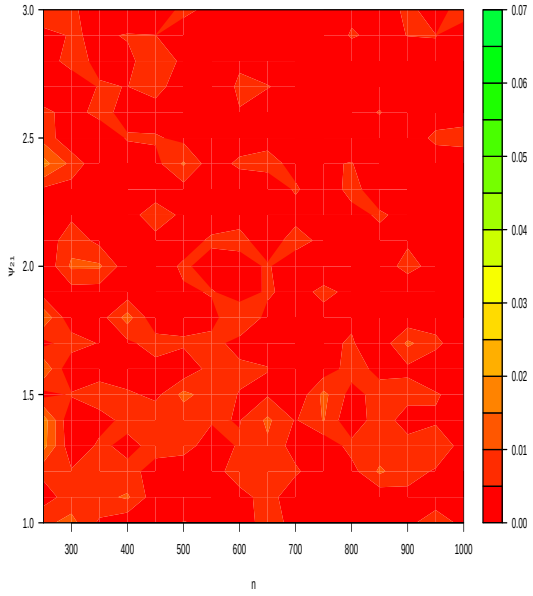

(b) Bias in $\hat{\psi}_{10}$, varying ψ_{21}

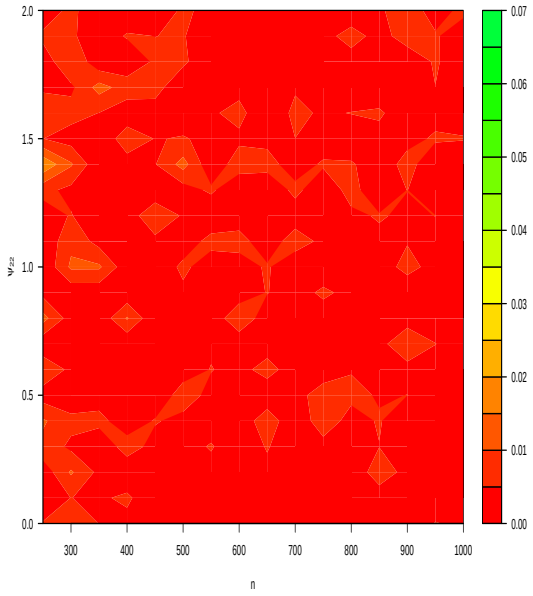

(c) Bias in $\hat{\psi}_{10}$, varying ψ_{22}

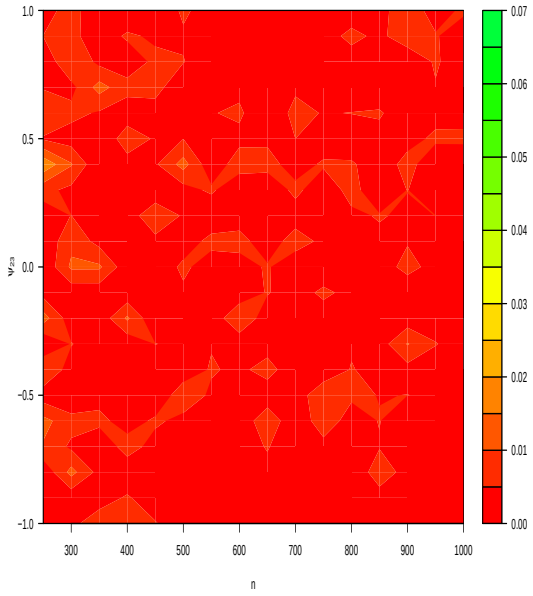

(d) Bias in $\hat{\psi}_{10}$, varying ψ_{23}

Figure 3:

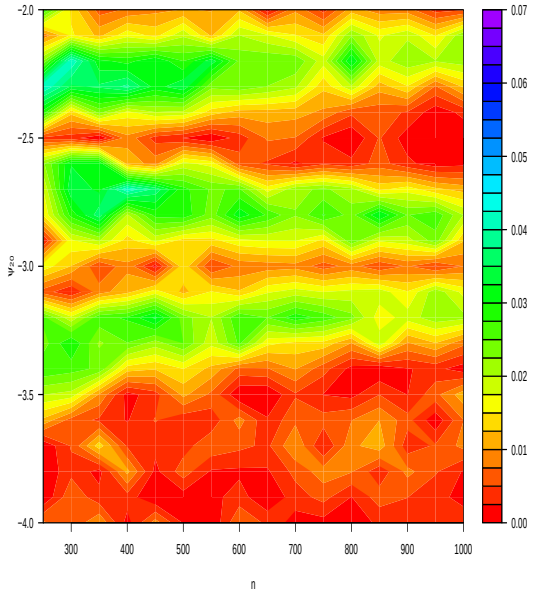

(a) Bias in $\hat{\psi}_{10}$, varying ψ_{20}

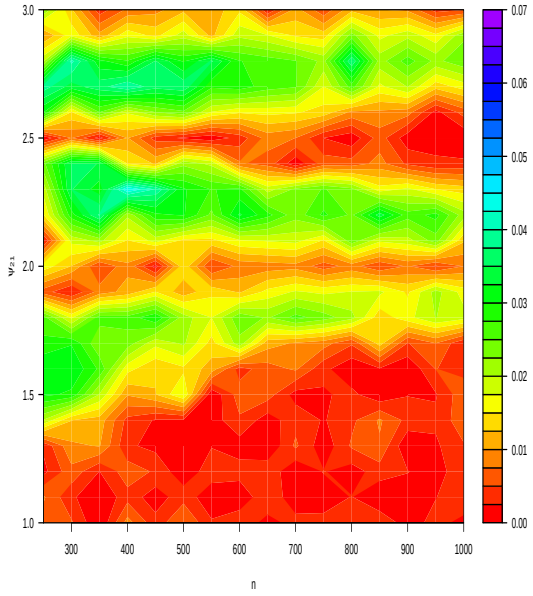

(b) Bias in $\hat{\psi}_{10}$, varying ψ_{21}

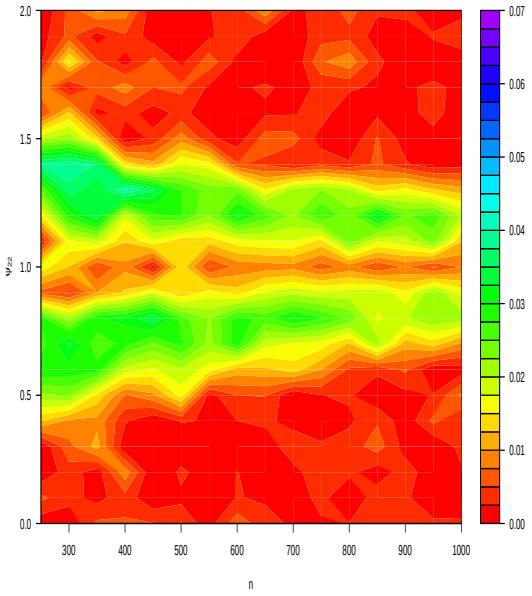

(c) Bias in $\hat{\psi}_{10}$, varying ψ_{22}

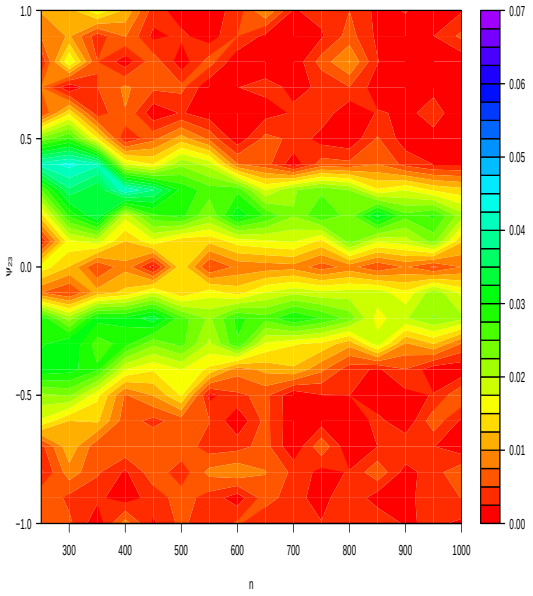

(d) Bias in $\hat{\psi}_{10}$, varying ψ_{23}

Figure 4:

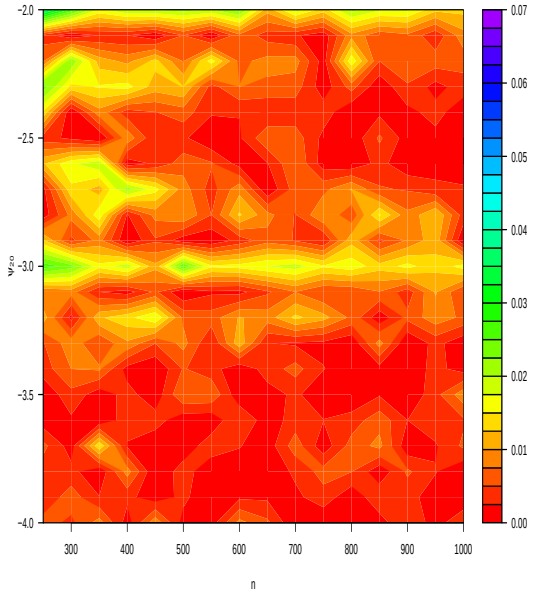

(a) Bias in $\hat{\psi}_{10}$, varying ψ_{20}

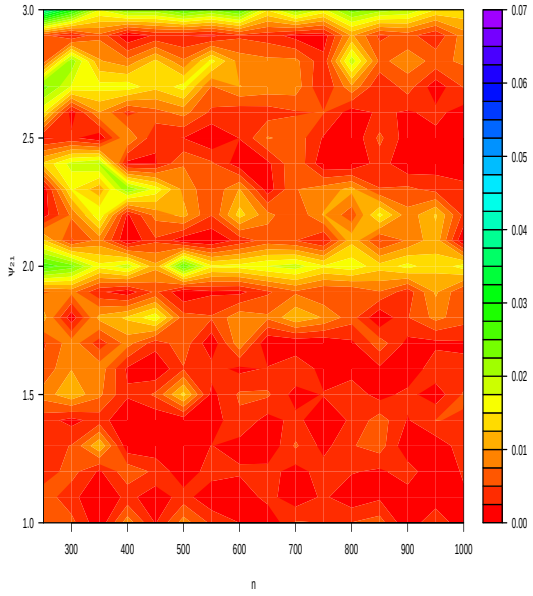

(b) Bias in $\hat{\psi}_{10}$, varying ψ_{21}

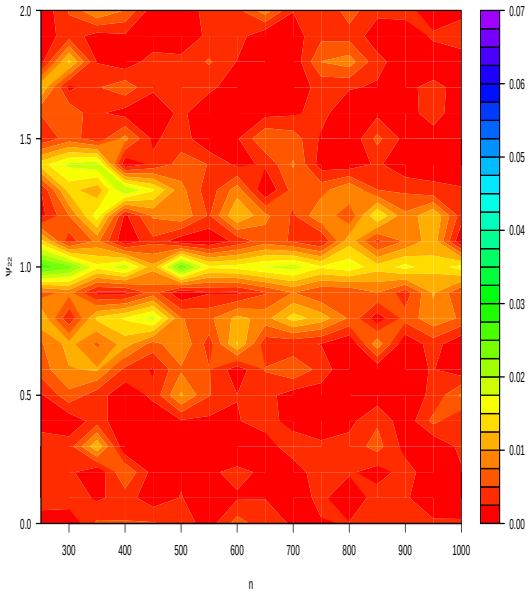

(c) Bias in $\hat{\psi}_{10}$, varying ψ_{22}

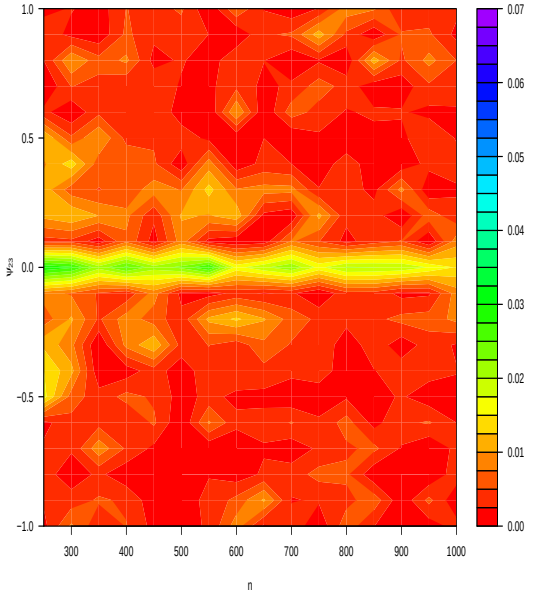

(d) Bias in $\hat{\psi}_{10}$, varying ψ_{23}
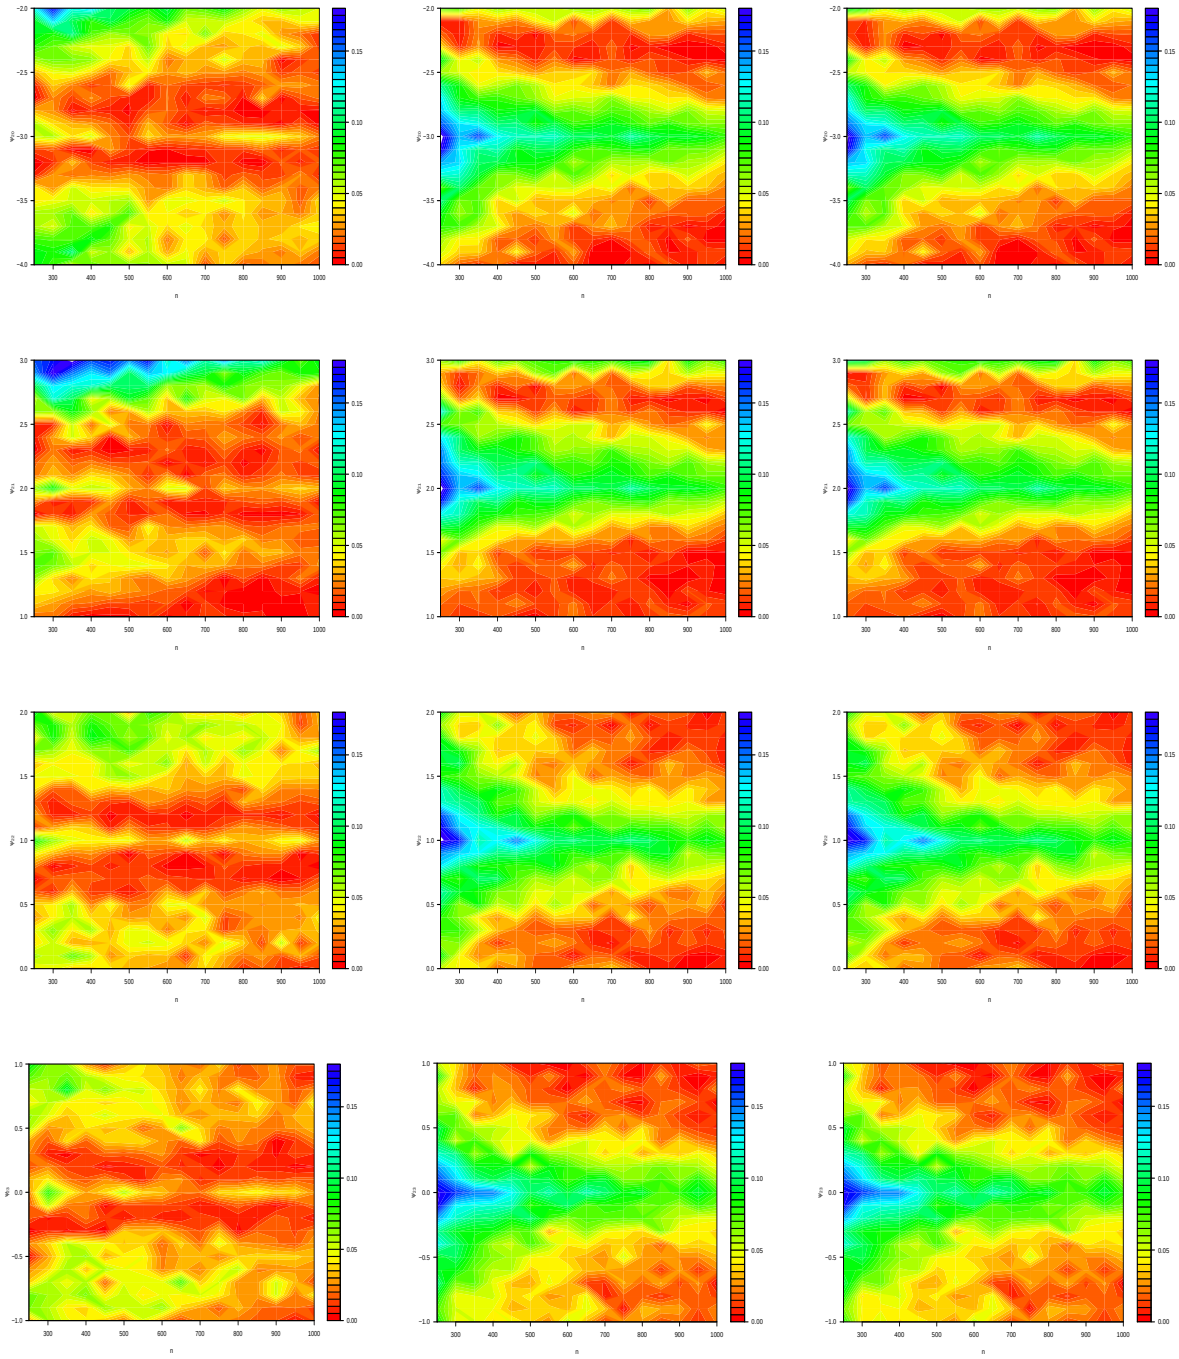
Figure 5: