

On Hodges' superefficiency and merits of oracle property in model selection

Xianyi Wu¹ · Xian Zhou²

Received: 30 September 2016 / Revised: 15 April 2018 / Published online: 15 June 2018
© The Institute of Statistical Mathematics, Tokyo 2018

Abstract The oracle property of model selection procedures has attracted a large volume of favorable publications in the literature, but also faced criticisms of being ineffective and misleading in applications. Such criticisms, however, have appeared to be largely ignored by the majority of the popular statistical literature, despite their serious impact. In this paper, we present a new type of Hodges' estimators that can easily produce model selection procedures with the oracle and some other desired properties, but can be readily seen to perform poorly in parts of the parameter spaces that are fixed and independent of sample sizes. Consequently, the merits of the oracle property for model selection as extensively advocated in the literature are questionable and possibly overstated. In particular, because the mathematics employed in this paper are at an elementary level, this finding leads to new discoveries on the merits of the oracle property and exposes some overlooked crucial facts on model selection procedures.

Keywords Hodges' estimator · Model selection · Oracle property · Penalized maximum likelihood/least squares · Superefficiency

This work was partially supported by NSFC (Grant No. 71771089) and the 111 Project (Grant No. B14019).

✉ Xianyi Wu
xywu@stat.ecnu.edu.cn

¹ Department of Statistics and Actuarial Science, East China Normal University, Shanghai 200241, China

² Department of Actuarial Studies and Business Analytics, Macquarie University, Sydney, NSW 2109, Australia

1 Introduction

Model selection is undoubtedly an extensively employed technique in data analysis and has attracted a great deal of research interests in the literature. It has become increasingly popular and ubiquitous partly due to rapid advance in computational power. A large volume of literature on model selection has been published and widely used statistical software packages contain more or less routines for model selection. In particular, almost all textbooks on linear regressions (including those at undergraduate level) dedicate a whole chapter or a few sections to variable (model) selection.

Theoretically, there are two purposes for model selection in general: model identification and inference. The latter is often referred to as post-selection inference in practical data analyses (Berk et al. 2010, 2013). Typically in post-selection inference, model identification is an intermediate step in data analysis after that the analysts then perform further statistical analyses based on the selected model, pretending it is the “true” model.

This amounts to a two-step procedure in data analysis (Berk et al. 2010, 2013). Many textbooks on linear regression oriented in practical data analysis taught students how to use well-known techniques such as AIC (Akaike’s information criterion), BIC (Bayesian information criterion), adjusted R^2 , Mallows’ C_p to select “good” or “best” models, then draw statistical inference based on the selected models. Popular softwares, such as SAS and R, generally select a model by some criterion (such as AIC/BIC) and output the estimated coefficients of the selected explanatory variables as well as their corresponding p values computed as if the selected models truly represent the nature. Berk et al. (2010) provides an exemplified but limited list of remarkable research works of this type. More and interesting discussions on how model selection methods perform for model identification and inference can be found in Yang (2005, 2007).

In this paper, we focus on the parameter estimation for model selection based on an overall model with a fixed parameter space of dimension p (less than the sample size n). The models are selected by setting some of its parameters to zero. Thus, selecting a model corresponds to estimating some parameters by zero. Statistical inference is drawn from the nonzero estimates of the other parameters on the selected model. In this sense, the parameter estimation we consider covers both model identification and inference.

More recently, based on the idea of penalized maximum likelihood/least squares inherited from AIC and BIC, some researchers found a smart way to integrate this two-step analysis into a single-step procedure. It associates some cleverly designed penalties to the likelihood function (or squared errors in regression models) so that, by maximizing the penalized likelihood function (or minimizing the penalized squared errors in regression models), a part of parameters are estimated by zero and the others by nonzero quantities. This enables model identification and inference to be carried out together—referred to as a one-step procedure for model selection.

This new approach began with the famous LASSO algorithm (least absolute selection and shrinkage operator) proposed by Tibshirani (1996). Generally, let $l(Y, \theta)$ denote the log-likelihood of the parameters $\theta = (\theta_1, \theta_2, \dots, \theta_p)'$. Then the estimators of θ_i s are obtained by maximizing the penalized log-likelihood $PL(\theta) =$

$l(Y, \theta) + \sum_{i=1}^p f_i(\theta_i, \lambda_n)$, where Y is the sample of size n , $f_i(\theta_i, \lambda_n)$ the penalty associated with parameter θ_i , and λ_n a known tuning parameter. In parallel, the estimators for a linear regression model $Y = X_{n \times p}\theta + \varepsilon$ are obtained by minimizing penalized squares $PS(\theta) = (Y - X\theta)'(Y - X\theta) + 2 \sum_{i=1}^p f_i(\theta_i, \lambda_n)$. For particularly designed penalty f_i , maximizing $PL(\theta)$ produces such estimators that automatically estimate some θ_i by zero, so as to simultaneously select model and estimate the parameters of the selected model. Different designs of the penalties generate different selection estimation algorithms, such as LASSO (Tibshirani 1996), adaptive LASSO (Zou 2006), hard thresholding estimators, soft thresholding estimators, bridge regression estimators (Frank and Friedman 1993), SCAD (smoothly clipped absolute deviation penalty, Fan and Li 2001), elastic net method (Zou and Hastie 2005) and MCP (minimax concave penalty, Zhang 2010).

The first theoretical justification of the penalized maximum likelihood method appeared in Fan and Li (2001), who proved an appealing property of their SCAD estimators, referred to as the *oracle property*, as defined below.

Definition 1 Let $b(\theta) = \{i : \theta_i \neq 0, i = 1, 2, \dots, p\}$ and $\bar{b}(\theta) = \{1, 2, \dots, p\} - b(\theta)$ and rearrange the parameter vector as $\theta = (\theta'_{b(\theta)}, \theta'_{\bar{b}(\theta)})'$ without loss of generality.

An estimator $\hat{\theta}_n$ is said to possess the oracle property or an oracle estimator if

1. $\lim_{n \rightarrow \infty} \Pr_{\theta}(\hat{\theta}_n, \bar{b}(\theta) = 0) = 1$; and
2. $\sqrt{n}(\hat{\theta}_n, b(\theta) - \theta_{b(\theta)}) \xrightarrow{d} N(0, \mathcal{F}_{b(\theta)}^{-1})$, where \xrightarrow{d} indicates convergence in distribution and $\mathcal{F}_{b(\theta)}$ is the Fisher's information matrix of $\theta_{b(\theta)}$ knowing $\theta_{\bar{b}(\theta)} = 0$.

The oracle property defined above is described in limit sense as the sample size tends to infinity. It states that an oracle estimator is asymptotically equivalent to the ideal estimator with the active (nonzero) parameters only. This has extensively been promoted as a justification to using the single-step procedure SCAD estimation. The follow-ups of Fan and Li (2001) can be summarized in two aspects:

1. Many researchers have focused on finding model selection procedures for linear regression so as to produce oracle estimators that perform better in their finite sample simulations. Significant works in this line include adaptive LASSO by Zou (2006), adaptive group LASSO by Wang and Leng (2008), elastic net method by Zou and Hastie (2005), adaptive elastic net method by Zou and Zhang (2009), MCP by Zhang (2010), Orthogonalizing EM algorithm with nonconvex penalties by Xiong et al. (2017) and other variants by, e.g., Leng et al. (2014), Gefang (2014) and Wang et al. (2011), among a vast number of others.
2. More papers have aimed at extending the ideas growing up in linear regression to other models so as to derive estimators with the oracle property. Important examples include generalized linear models (Friedman et al. 2010; Geer 2008), Cox regression models (Fan and Li 2002; Zhang and Lu 2007), network exploration (Fan et al. 2009), additive models (Horowitz et al. 2006; Huang et al. 2010), panel data models (Fan and Li 2004 and Kock 2013) and propensity score models (Brookhart et al. 2006), among others.

What is worth special mention is the explosion of the literature in model selection with the oracle property under a huge number of statistical models, which have led

to thousands of citations to many of the papers mentioned above by following papers and books.

Such a popular property, however, is not universally accepted without criticisms. The Vienna school, led by H. Leeb and B. M. Pötscher, argued that the oracle property is “too good to be true” and seriously questioned the advertised merits of oracle model selection procedures in a series of papers, including [Leeb and Pötscher \(2008a,b\)](#), [Pötscher \(2009\)](#), [Pötscher and Leeb \(2009\)](#), [Pötscher and Schneider \(2009\)](#) and the references therein. They considered oracle estimators as a return of Hodges’ estimators ([Le Cam 1953](#)), which provided a typical counterexample to a conjecture by R. A. Fisher that the MLEs have minimum asymptotic variance (referred to as asymptotically efficient) and are known to perform poorly at some parameter values in finite sample size. Their arguments are summarized below.

1. Theoretically, the overall risk of a sparse estimator can be unbounded as $n \rightarrow \infty$.
2. Numerous Monte Carlo simulations were conducted to demonstrate that, in finite sample size, a number of oracle estimators perform quite poorly when the parameters take values in a vicinity of zero.

In particular, by replicating and extending Monte Carlo simulations of the performance of the SCAD estimator in Example 4.1 of [Fan and Li \(2001\)](#), [Leeb and Pötscher \(2008b\)](#) demonstrated that this estimator, when tuned to enjoy sparsity and oracle properties, can perform poorly in finite samples. Even if not tuned to sparsity, the SCAD estimator performs worse than the least squares estimator in parts of the parameter space. It is interesting to note that the simulation study in [Fan and Li \(2001\)](#) was conducted only at some points that happened to avoid the parts of the parameter space examined by [Leeb and Pötscher \(2008b\)](#). This phenomenon highlights the fact that simulations may produce results to support either side of a controversy and thus are unreliable to judge the goodness of an estimator. This is generally recognized in the statistical community but often overlooked. (More details will be discussed in “The power of simulations” in Sect. 5.) Consequently, [Leeb and Pötscher \(2008b\)](#) argued that the oracle property is highly misleading and cannot be relied upon to justify an estimator.

It is also worth to note that a procedure with oracle property is consistent for model selection. While consistency is an important and desirable property in limiting sense for large samples, it is not sufficient to justify the superiority of a procedure for finite (fixed) sample sizes. As indicated in [Yang \(2005, 2007\)](#), BIC type procedures are consistent and AIC type procedures are inconsistent but optimal in minimax rate of convergence. [Yang \(2005, 2007\)](#) also argued that the strength of AIC and BIC could not be shared. More introduction on optimal model selection in minimax rate of convergence can be found in the recent work of [Wang et al. \(2014\)](#).

In this paper, we attempt to address this controversy by revisiting the issue of Hodges’ superefficiency and theoretically analyzing the performance of oracle estimators in a class of generalized Hodges’ estimators without relying on numerical simulations. As simulations are subject to large variations and may produce different results in either side of a controversy, they are not capable of resolving the controversy convincingly. The theoretical analysis in this paper can avoid such drawbacks of the simulation approach and thus provide more convincing conclusions. In particular, because the mathematics employed in this paper are at a quite elementary level,

this finding leads to some new discoveries on the merits of the oracle property and exposes some overlooked crucial facts on model selection procedures. In addition, some significant but often ignored facts on asymptotic properties are also emphasized to warn the pitfall of justifying the merits of statistical procedures based on asymptotic measures that the oracle property relies on.

Our arguments proceed as follows.

1. Generally, asymptotic bias and variance do not reflect their versions in finite sample size. Thus, the asymptotic superefficiency and the oracle property do not necessarily lead to good performance of the estimators in any fixed sample size.
2. We introduce a new type of Hodges' estimators, referred to as *oracle Hodges' estimators*, which can reduce the asymptotic variance of any given estimator over certain multi-dimensional subsets of the parameter space and provide an easy way to obtain oracle model selection procedures.
3. By analyzing the performance of oracle Hodges' estimators in finite sample size, we are able to theoretically demonstrate that such estimators perform poorly at some (true) values of the parameters despite possessing the oracle property.
4. The performance of oracle Hodges' estimators indicates that the oracle property alone does not justify the use of oracle model selection procedures, and hence the merits of the oracle property as advocated in the literature are questionable. In particular, there has been no rigorous theory, in any case where the MLE/LSE are available, to prove that oracle estimators are better than MLE/LSE under certain commonly accepted criteria (such as smaller mean squared error) with fixed sample size.
5. Our results provide theoretical justification to support the view of [Leeb and Pötscher \(2008b\)](#) regarding the oracle model selection approach and clarified that the oracle property is not a simple return of classical Hodges' superefficiency as suggested by Leeb and Pötscher. Instead, the newly defined oracle Hodges' estimators, rather than their classical versions, have the same asymptotic behavior as the estimators from oracle model selection procedures.

In Sect. 2 next, the new oracle Hodges' estimators are defined and their asymptotic behaviors (superefficiency) are discussed. Section 3 introduces model selection methods derived from Hodges' estimators and discusses their connections to and differences from penalized maximum likelihood or least squares estimators. The performance of oracle Hodges' estimators is theoretically analyzed in Sect. 4 after a brief discussion on the connection of asymptotic biases and variances to their finite sample size versions. Some concluding remarks are discussed in Sect. 5. All proofs are put in Appendix.

2 Asymptotic efficiency and Hodges' estimators

The concept of asymptotic efficiency was introduced by Sir R. A. Fisher with the attempt to justify the goodness of MLEs and has great impacts on statistical inference in large samples. For a d -dimensional parameter $\theta = (\theta_1, \theta_2, \dots, \theta_d)'$, let $\hat{\theta}_n = (\hat{\theta}_{n1}, \hat{\theta}_{n2}, \dots, \hat{\theta}_{nd})'$ be any sequence of its estimators such that $r_n(\hat{\theta}_n - \theta) \xrightarrow{d} Z$ for some sequence of scalars $r_n \rightarrow \infty$ and a d -dimensional random vector Z with mean

zero and variance-covariance matrix V^{-1} for some matrix $V = (v_{ij})_{d \times d}$. A sequence of estimators $\{\hat{\theta}_n\}$ of θ such that $\sqrt{n}(\hat{\theta}_n - \theta) \xrightarrow{d} N(0, \Sigma(\theta))$, whatever is the true value of θ , is said to be *asymptotically efficient* (or *asymptotically optimal* in some literature) if $\Sigma(\theta) = \mathfrak{F}^{-1}(\theta)$, where $\mathfrak{F}(\theta)$ is the Fisher’s information matrix of the distribution. Namely, $\hat{\theta}_n$ is efficient if it is asymptotically unbiased (in the sense that the limiting random variable Z has zero mean) and “optimal” (in a certain sense based on the asymptotic variance) in the class of asymptotically unbiased estimators with order $n^{-1/2}$.

The general results state that, under certain regularity conditions, a sequence of roots of likelihood equations is asymptotically efficient. There have been a huge number of research efforts aimed at seeking asymptotically efficient estimators. The idea of measuring efficiency by asymptotic variance, however, appears not as successful as its counterpart in finite sample size, and the lower bounds defined by $\mathfrak{F}^{-1}(\theta)$ are not sufficiently low in the class of asymptotically unbiased estimators. The example in Sect. 2.1 below, whose prototype was made widely known by [Le Cam \(1953\)](#) under the name Hodges’ estimator, shows that, given any sequence of estimators $\hat{\theta}_n$ with $r_n(\hat{\theta}_n - \theta) \xrightarrow{d} Z$, one can construct another estimator sequence with asymptotic variance no more than that of $\hat{\theta}_n$ at any value of θ , and strictly less at certain values of θ . This phenomenon is referred to as (asymptotic) *superefficiency*. In spirit of this idea, we introduce variants of Hodges’ estimators in Sect. 2.2, which possess the oracle property (Definition 1) and can easily serve the purpose of finding oracle model selection procedures.

Let $\|x\| = \sqrt{\sum_{i=1}^d x_i^2}$ denote the Euclidean norm of $x = (x_1, x_2, \dots, x_d) \in \mathbb{R}^d$, $d(x, A) = \inf_{y \in A} \|x - y\|$ the distance between a point x and a subset A and $d(A, B) = \inf_{x \in A, y \in B} \|x - y\|$ the distance between two subsets of \mathbb{R}^d under the Euclidean norm.

2.1 Classical Hodges’ example of superefficiency

Let $\hat{\theta}_n$ be any sequence of estimators such that $r_n(\hat{\theta}_n - \theta) \xrightarrow{d} Z$ for a sequence of deterministic scalars $r_n \rightarrow \infty$, where the distribution of Z may depend on θ , whatever is the true value of θ .

Definition 2 (Hodges’ estimator) Let $\{a_n\}$ be a sequence of scalars and c any fixed point in the parameter space Θ . The Hodges’ estimator of θ is defined by

$$\check{\theta}_n(c) = cI(\|\hat{\theta}_n - c\| \leq a_n) + \hat{\theta}_n I(\|\hat{\theta}_n - c\| > a_n). \tag{1}$$

This estimator was initiated by Jr. Hodges with $r_n = \sqrt{n}$, $c = 0$ and $a_n = n^{-1/4}$ for one-dimensional θ (reported by [Le Cam 1953](#); see also [Lehmann and Casella 1998](#), p. 420, Example 2.5) and has been revisited many times by, e.g., [Leeb and Pötscher \(2005, 2008a, b\)](#) when they discussed the implications of consistent model selections. While (1) may be slightly generalized to

$$\check{\theta}_n(c) = ((1 - \alpha)c + \alpha\hat{\theta}_n)I(|\hat{\theta}_n - c| \leq a_n) + \hat{\theta}_n I(|\hat{\theta}_n - c| > a_n),$$

where $\alpha \in [0, 1]$ (cf., e.g., van der Vaart 1998 for the case with $r_n = \sqrt{n}$ and $c = 0$), we here take $\alpha = 0$ as this is sufficient for our purpose. The asymptotic distribution of $\check{\theta}_n(c)$ in Definition 2 is given in the following theorem.

Theorem 1 *For any sequence $\{a_n\}$ such that $a_n = o(1)$ and $r_n a_n \rightarrow \infty$ as $n \rightarrow \infty$, we have*

$$r_n(\check{\theta}_n(c) - \theta) \xrightarrow{d} ZI(\theta \neq c).$$

It is clear that, in terms of asymptotic variances, $\check{\theta}_n(c)$ is no worse than $\hat{\theta}_n$ at any θ and strictly better than $\hat{\theta}_n$ at $\theta = c$ because the asymptotic variance of $\check{\theta}_n(c)$ is zero at $\theta = c$. This example revealed an interesting phenomenon that, in terms of the asymptotic variance, any estimate can be improved at an arbitrary but fixed point in the parameter space.

2.2 Oracle Hodges' estimators

We next introduce a few closely linked variants of Hodges' estimators, which differ from the classical version in Definition 2, but we have kept Hodges' name for the new estimators because they retain the feature of superefficiency. We will refer to the new type of Hodges' estimators as *oracle Hodges' estimators* because they possess the oracle property in Definition 1, as will be shown via Theorems 2 to 5 below. The definition of oracle Hodges' estimators is presented in four versions: general version, continuous/smoothing version, MLE version and LSE version, where the first one is fundamental, the second is a refinement, and the other two are special cases with origins $\hat{\theta}_n$ being MLE and LSE, respectively.

2.2.1 General version

For any subset $b \subset \{1, 2, \dots, d\}$ and its complement $\bar{b} = \{1, 2, \dots, d\} - b$, rearrange the components of θ and $\hat{\theta}_n$ as $\theta' = (\theta'_b, \theta'_{\bar{b}})$ and $\hat{\theta}'_n = (\hat{\theta}'_{n,b}, \hat{\theta}'_{n,\bar{b}})$, respectively, such that $r_n(\hat{\theta}_n - \theta) \xrightarrow{d} Z$, where the current $Z = (Z'_b, Z'_{\bar{b}})'$ is also a rearrangement of the original Z in the same way as θ . Clearly, the mean of Z is still a zero vector but the covariance matrix changes to V^{-1} with

$$V = \begin{pmatrix} V_{bb} & V_{b\bar{b}} \\ V_{\bar{b}b} & V_{\bar{b}\bar{b}} \end{pmatrix} = \begin{pmatrix} (v_{ij})_{i,j \in b} & (v_{ij})_{i \in b, j \in \bar{b}} \\ (v_{ij})_{i \in \bar{b}, j \in b} & (v_{ij})_{i,j \in \bar{b}} \end{pmatrix}.$$

It is also easy to see that, if $b \neq \emptyset$, then the marginal vector $\hat{\theta}_{n,b}$ has an asymptotic distribution given by $r_n(\hat{\theta}_{n,b} - \theta_b) \xrightarrow{d} Z_b$ with mean zero and covariance matrix

$$\Delta_b = (V_{bb} - V_{b\bar{b}}V_{\bar{b}\bar{b}}^{-1}V_{\bar{b}b})^{-1}. \tag{2}$$

Note that the distribution of Z , and hence the variance matrix V^{-1} , may be related to parameters θ . The only requirement is that V^{-1} is a continuous function of θ . We use

\hat{V}^{-1} to denote any of consistent estimators of V^{-1} , e.g., obtained by substituting $\hat{\theta}_n$ for θ , so that the symbols \hat{V}_{bb} , $\hat{V}_{b\bar{b}}$ and so on are self-explained.

Let $c = (c_1, c_2, \dots, c_d)'$ be a known d -vector. For every nonempty and proper subset b of $\{1, 2, \dots, d\}$ (i.e., $\emptyset \neq b \neq \{1, 2, \dots, d\}$), denote

$$\check{\theta}_{n,b} = \hat{\theta}_{n,b} + \hat{V}_{bb}^{-1} \hat{V}_{b\bar{b}} (\hat{\theta}_{n,\bar{b}} - c_{\bar{b}}) \quad \text{and} \quad \check{\theta}_n(b) = (\check{\theta}'_{n,b}, c'_{\bar{b}})' \tag{3}$$

with the convention $\check{\theta}_{n,\{1,2,\dots,d\}} = \check{\theta}_n(\{1, 2, \dots, d\}) = \hat{\theta}_n$. Moreover, we redefine $b(\theta)$ and $\bar{b}(\theta)$ by

$$b(\theta) = \{j : j \in \{1, 2, \dots, d\}, \theta_j \neq c_j\} \quad \text{and} \quad \bar{b}(\theta) = \{1, 2, \dots, d\} - b(\theta). \tag{4}$$

The following definition introduces a sequence $\check{\theta}_n(c)$ of oracle Hodges' estimators in multi-dimensional case derived from $\hat{\theta}_n$.

Definition 3 (Oracle Hodges' estimators) Let $(a_{nj}) = (a_{n1}, \dots, a_{nd})$, $n = 1, 2, \dots$, denote a sequence of d -vectors with positive components. For every $n = 1, 2, \dots$, define two complementary random sets by

$$b_n(c) = \{j : |\hat{\theta}_{nj} - c_j| > a_{nj}\} \quad \text{and} \quad \bar{b}_n(c) = \{j : |\hat{\theta}_{nj} - c_j| \leq a_{nj}\} \tag{5}$$

and the corresponding oracle Hodges' estimator by

$$\check{\theta}_n(c) = \check{\theta}_n(b_n(c)) = \begin{cases} \hat{\theta}_n & \text{if } b_n(c) = \{1, 2, \dots, d\}, \\ (\check{\theta}'_{n,b_n(c)}, c'_{\bar{b}_n(c)})' & \text{if } b_n(c) \neq \emptyset, b_n(c) \neq \{1, 2, \dots, d\}, \\ c & \text{if } b_n(c) = \emptyset, \end{cases} \tag{6}$$

where $\check{\theta}_n(b_n(c))$ and $\check{\theta}_{n,b_n(c)}$ are obtained from the two equations in (3) by substituting $b_n(c)$ for b , and $b_n(c)$ is defined in (5).

For later reference, denote

$$\check{Z}_b = \begin{cases} V_{bb}^{-1} (V_{bb} & V_{b\bar{b}})Z, & \text{if } b \neq \emptyset, \\ 0, & \text{if } b = \emptyset. \end{cases} \tag{7}$$

The asymptotic properties of $\check{\theta}_n(c)$ defined in (6) are provided in the next theorem.

Theorem 2 *If the sequence of d -vectors $\{(a_{nj})\}$ satisfies*

$$\max_{1 \leq j \leq d} a_{nj} \rightarrow 0 \quad \text{and} \quad r_n \min_{1 \leq j \leq d} a_{nj} \rightarrow \infty \quad \text{as } n \rightarrow \infty, \tag{8}$$

then for any $b \subset \{1, 2, \dots, d\}$, $\theta_{\bar{b}} = c_{\bar{b}}$ implies $\lim_{n \rightarrow \infty} \Pr(\check{\theta}_{n,\bar{b}}(c) = \theta_{\bar{b}}) \rightarrow 1$ and

$$r_n(\check{\theta}_n(c) - \theta) \xrightarrow{d} \begin{pmatrix} \check{Z}_{b(\theta)} \\ 0 \end{pmatrix} \quad \text{under } \Pr_{\theta}, \tag{9}$$

whatever is the true value of θ , where $\check{Z}_{b(\theta)}$ is defined as in (7) with b replaced by $b(\theta)$ in (4).

2.2.2 Continuous/smoothing version

As a function of $\hat{\theta}_n$, the estimator $\tilde{\theta}_n(c)$ is not continuous at any point $\hat{\theta}_n = (\hat{\theta}_{n1}, \hat{\theta}_{n2}, \dots, \hat{\theta}_{nd})$ such that $|\hat{\theta}_{nj} - c_j| = a_{nj}$ for some $j \in \{1, 2, \dots, d\}$. Some authors think of the continuity as an important property (see, e.g., Fan and Li 2001) but others may disagree. If preferred, a continuous version of $\tilde{\theta}_n(c)$ can be easily achieved by the following procedure.

Let $(a_{nj}^{(1)})$ and $(a_{nj}^{(2)})$ be two sequences of d -vectors both satisfying conditions (8) on (a_{nj}) and $a_{nj}^{(1)} < a_{nj}^{(2)}$, $j = 1, 2, \dots, d$. A possible choice is $a_{nj}^{(1)} = r_n^{-1/2}$ and $a_{nj}^{(2)} = 2r_n^{-1/2}$, $j = 1, 2, \dots, d$. Define two sequences of oracle Hodges' estimators $\tilde{\theta}_n^{(1)}(c)$ and $\tilde{\theta}_n^{(2)}(c)$ by $(a_{nj}^{(1)})$ and $(a_{nj}^{(2)})$, respectively, as in Definition 3. Let $f_1(x), \dots, f_d(x)$ be any d continuous and increasing functions on $x \in \mathbb{R}^+$ such that $f_j(c_{nj} \pm a_{nj}^{(1)}) = c_{nj}$ and $f_j(c_{nj} \pm a_{nj}^{(2)}) = c_{nj} \pm a_{nj}^{(2)}$. Define

$$\tilde{\theta}_{nj}(c; f) = \begin{cases} c_{nj} & \text{if } |\hat{\theta}_{nj} - c_j| \leq a_{nj}^{(1)}, \\ f(\hat{\theta}_{nj}) & \text{if } a_{nj}^{(1)} \leq |\theta_{nj} - c_j| \leq a_{nj}^{(2)}, \\ \tilde{\theta}_{nj}^{(2)}(c) & \text{otherwise.} \end{cases} \tag{10}$$

Then $\tilde{\theta}_{nj}(c; f)$ is a continuous version of $\tilde{\theta}_n(c)$ such that

$$|\tilde{\theta}_{nj}^{(2)}(c) - c_{nj}| \leq |\tilde{\theta}_{nj}(c; f) - c_{nj}| < |\tilde{\theta}_{nj}^{(1)}(c) - c_{nj}|.$$

These inequalities ensure the following result.

Theorem 3 *The estimators $\tilde{\theta}_n(c; f)$ defined in (10) have the same asymptotic properties of $\tilde{\theta}_n^{(i)}(c)$, $i = 1, 2$, as presented in Theorem 2.*

2.2.3 MLE version

We next discuss the maximum likelihood estimation with the simplest i.i.d. case as an example. It is not difficult to extend the results to general situations.

The log-likelihood function of θ from i.i.d. X_1, \dots, X_n with a common density $f(x; \theta)$ is

$$l(\theta) = l(\theta; X_1, \dots, X_n) = \prod_{i=1}^n \log f(X_i; \theta).$$

It is well known that under certain regularity conditions, there exists a sequence of asymptotically efficient MLE $\hat{\theta}_n$, i.e., $\sqrt{n}(\hat{\theta} - \theta) \xrightarrow{d} N(0, \mathfrak{F}^{-1}(\theta))$, where

$$\mathfrak{F}(\theta) = -E_{\theta} \left[\frac{\partial^2}{\partial \theta \partial \theta'} \log f(X_1; \theta) \right]$$

is the Fisher’s information matrix. For any $b \subset \{1, 2, \dots, d\}$ and the corresponding rearrangement of $\theta = (\theta_b, \theta_{\bar{b}})$, $\mathfrak{F}(\theta)$ can be rewritten as

$$\begin{aligned} \mathfrak{F}(\theta) &= - \begin{pmatrix} E_{\theta} \left[\frac{\partial^2 \log f(X_1; \theta)}{\partial \theta_b \partial \theta'_b} \right] & E_{\theta} \left[\frac{\partial^2 \log f(X_1; \theta)}{\partial \theta_b \partial \theta'_{\bar{b}}} \right] \\ E_{\theta} \left[\frac{\partial^2 \log f(X_1; \theta)}{\partial \theta_{\bar{b}} \partial \theta'_b} \right] & E_{\theta} \left[\frac{\partial^2 \log f(X_1; \theta)}{\partial \theta_{\bar{b}} \partial \theta'_{\bar{b}}} \right] \end{pmatrix} \\ &= \begin{pmatrix} \mathfrak{F}_{bb'}(\theta) & \mathfrak{F}_{b\bar{b}'}(\theta) \\ \mathfrak{F}_{\bar{b}b'}(\theta) & \mathfrak{F}_{\bar{b}\bar{b}'}(\theta) \end{pmatrix}, \text{ say.} \end{aligned}$$

For any constant vector $c = (c_1, \dots, c_d)$, Theorem 2 yields the following immediate results.

Theorem 4 *If $\{\hat{\theta}_n\}$ is an efficient sequence of maximum likelihood estimators, then the oracle Hodges’ estimators $\tilde{\theta}(c)$ in Definition 3 have the following properties: For any $b \subset \{1, 2, \dots, d\}$, $\theta_{\bar{b}} = c_{\bar{b}}$ implies $\lim_{n \rightarrow \infty} \Pr(\hat{\theta}_{n,\bar{b}}(c) = \theta_{\bar{b}}) = 1$ and*

$$\sqrt{n}(\tilde{\theta}(c) - \theta) \xrightarrow{d} N \left(0, \begin{pmatrix} \mathfrak{F}_{b(\theta)b'(\theta)}^{-1}(\theta) & 0 \\ 0 & 0 \end{pmatrix} \right).$$

Clearly, $\mathfrak{F}_{bb'}(\theta)$ is the Fisher’s information matrix for parameter θ_b depending on the unknown $\theta_{\bar{b}}$. The asymptotic variance of $\sqrt{n}(\hat{\theta}_{n,b} - \theta_{n,b})$ is

$$(\mathfrak{F}^{-1}(\theta))_{bb'} = (\mathfrak{F}_{bb'}(\theta) - \mathfrak{F}_{b\bar{b}'}(\theta)\mathfrak{F}_{\bar{b}\bar{b}'}^{-1}(\theta)\mathfrak{F}_{\bar{b}b'}(\theta))^{-1} \geq \mathfrak{F}_{bb'}^{-1}(\theta) \tag{11}$$

with strict inequality if $b \neq \{1, 2, \dots, d\}$.

Note that $\mathfrak{F}_{b(\theta)b'(\theta)}$ is the Fisher’s information matrix of the marginal vector $\theta_{b(\theta)}$ knowing that $\theta_{\bar{b}(\theta)} = c_{\bar{b}(\theta)}$. Taking $b = b(\theta)$, Theorem 4 shows the superefficiency of $\tilde{\theta}_n(c)$ over the MLE $\hat{\theta}_n$ at any θ in the parameter space such that $b(\theta) \neq \{1, 2, \dots, d\}$ (i.e., $\theta_j = c_j$ for some $j \in \{1, 2, \dots, d\}$). This result covers such parametric models as linear regression with normally distributed errors and generalized linear regression.

2.2.4 LSE version

For a linear regression model $\mathbf{Y} = \mathbf{X}\beta + \varepsilon$ with $E[\varepsilon] = 0$ and $\text{Var}(\varepsilon) = \sigma^2 I_n$, we can generate the oracle Hodges’ estimator $\hat{\beta}(c)$ from the least square estimate $\hat{\beta}_{\text{LS}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$. Under certain regularity conditions (see, e.g., van de Vaart, 2000, Example 2.28), $\hat{\beta}_{\text{LS}}$ is asymptotically distributed as $\sqrt{n}(\hat{\beta}_{\text{LS}} - \beta) \xrightarrow{d} N(\mathbf{0}, \sigma^2 \Sigma_X^{-1})$, where $\Sigma_X = \lim_{n \rightarrow \infty} n^{-1}\mathbf{X}'\mathbf{X}$. By Theorem 2 again, the following result is obvious.

Theorem 5 *Given any fixed d -vector c and subset $b \subset \{1, 2, \dots, d\}$, $\beta_{\bar{b}} = c_{\bar{b}}$ implies that $\lim_{n \rightarrow \infty} \Pr(\hat{\beta}_{n,\bar{b}}(c) = \beta_{\bar{b}}) = 1$ and*

$$\begin{aligned} \sqrt{n}(\tilde{\beta}(c) - \beta) &\xrightarrow{d} N\left(0, \begin{pmatrix} \sigma^2 \Sigma_{b(\beta)b(\beta)}^{-1} & 0 \\ 0 & 0 \end{pmatrix}\right), \text{ where } \Sigma_{b(\beta)b(\beta)} \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \mathbf{X}'_{b(\beta)} \mathbf{X}_{b(\beta)}. \end{aligned} \tag{12}$$

This theorem shows that $\sqrt{n}(\tilde{\beta}_{b(\beta)}(c) - \beta_{b(\beta)})$ has the same asymptotic distribution as the oracle estimator $\hat{\beta}_{b(\beta)}^o = (\mathbf{X}'_{b(\beta)} \mathbf{X}_{b(\beta)})^{-1} \mathbf{X}'_{b(\beta)}(Y - \mathbf{X}_{\bar{b}(\beta)} c_{\bar{b}(\beta)})$ if the true value of β is $(\beta'_{b(\beta)}, c'_{\bar{b}(\beta)})'$.

2.3 Two remarks

We conclude this section with the following two remarks.

Remark 1 By Theorems 1, the classical version $\check{\theta}_n(c)$ has an asymptotic distribution given by $r_n(\check{\theta}_n(c) - \theta) \xrightarrow{d} ZI(\theta \neq c)$. That is, $\check{\theta}_n(c)$ can only improve the asymptotic variance of $\hat{\theta}_n$ at $\theta = c$ in the parameter space Θ , which is much more restrictive than the improvement achieved by the limit in (9) for $\tilde{\theta}_n(c)$ defined by (6). To see this, note that the variance of $\check{Z}_{b(\theta)}$ is

$$\begin{aligned} \check{V}_{b(\theta)} &= V_{b(\theta), b(\theta)}^{-1} (V_{b(\theta), b(\theta)} \quad V_{b(\theta), \bar{b}(\theta)}) V^{-1} \begin{pmatrix} V_{b(\theta), b(\theta)} \\ V_{b(\theta), \bar{b}(\theta)} \end{pmatrix} V_{b(\theta), b(\theta)}^{-1} \\ &= V_{b(\theta), b(\theta)}^{-1} \leq \Delta_{b(\theta)}, \end{aligned}$$

where $\Delta_{b(\theta)}$ is the asymptotic variance of $r_n(\hat{\theta}_{n, b(\theta)} - \theta_{b(\theta)})$ by (2), and the equality holds only when $\hat{\theta}_{n, b(\theta)}$ and $\hat{\theta}_{n, \bar{b}(\theta)}$ are asymptotically independent. Therefore, $\tilde{\theta}_n(c)$ can improve the asymptotic variance of $\hat{\theta}_n$ at any θ with $\bar{b}(\theta) \neq \emptyset$. Note also that $r_n(\hat{\theta}_n(c) - \theta) \rightarrow 0$ in probability at $\theta = c$. In terms of asymptotic variances, $\check{\theta}(c)$ improves $\hat{\theta}_n$ and $\tilde{\theta}_n(c)$ further improves $\check{\theta}(c)$. A further important feature of $\tilde{\theta}_n(c)$ is its ability to produce oracle model selection procedures due to its form of asymptotic variance, as we will show in the next subsection. In contrast, neither $\check{\theta}_n(c)$ nor $\hat{\theta}_n$ has such a capacity. These together highlight the significant differences between the new oracle Hodges' estimator $\tilde{\theta}_n(c)$ and the classical version $\check{\theta}(c)$.

Remark 2 Note that MLE and LSE and their versions of oracle Hodges' estimators are of root- n consistency under relevant regularity conditions. But this is not necessary for Definition 3. The general version of $\tilde{\theta}_n(c)$ and its continuous version do not require those regularity conditions, and they are not necessarily of root- n consistency. For example, let $Y_i = (Y_{i1}, \dots, Y_{id})', i = 1, 2, \dots$, be independent with identical uniform distributions over $\prod_{k=1}^d [-\theta_k, \theta_k], \theta_k > 0$, so that the MLE of $\theta = (\theta_1, \dots, \theta_d)$ is $\hat{\theta}_n = (\hat{\theta}_{n1}, \dots, \hat{\theta}_{nd})$ with $\hat{\theta}_{nk} = \max(|Y_{1k}|, \dots, |Y_{nk}|), k = 1, 2, \dots, d$. Then $n(\hat{\theta}_n - \theta) \xrightarrow{d} Z = (Z_1, \dots, Z_d)$ with mutually independent components Z_1, \dots, Z_d such that

$$\Pr_{\theta}(Z_k \leq x_k) = \begin{cases} e^{x_k/\theta_k} & \text{if } x_k < 0 \\ 1 & \text{if } x_k \geq 0 \end{cases}, \quad k = 1, 2, \dots, d.$$

In this case, the general version $\tilde{\theta}_n(c)$ in (6) and its continuous version $\tilde{\theta}_n(c, f)$ in (10) are still valid even though the regularity conditions of the likelihood function are not satisfied, but they are not of root- n consistency.

3 Model selection function and oracle property of $\tilde{\theta}_n(c)$

Definition 3 and Theorem 2 clearly indicate the following properties of $\tilde{\theta}_n(c)$:

1. $\tilde{\theta}_n(c)$ is a sparse estimate in the sense that some components of θ , say θ_j , may be estimated by component c_j of c .
2. $\lim_{n \rightarrow \infty} \Pr(\tilde{\theta}_{n, \tilde{b}(\theta)}(c) = c_{\tilde{b}(\theta)}) = 1$ for whatever true value of the parameter θ .
3. For any sequence of estimators $\hat{\theta}_n$, it is possible to define a new sequence $\tilde{\theta}_n(c)$ such that its asymptotic covariance matrix (with the same convergence rate as $\hat{\theta}_n$) is
 - equal to that of $\hat{\theta}_n$ if $\theta_j \neq c_j$ for all $j \in \{1, 2, \dots, d\}$, i.e., $b(\theta) = \{1, 2, \dots, d\}$,
 - positive definite and strictly less than that of $\hat{\theta}_n$ if $\emptyset \neq b(\theta) \neq \{1, 2, \dots, d\}$; in this case $\tilde{\theta}_n(c)$ is asymptotically more efficient than $\hat{\theta}_n$ because

$$\text{Cov} \begin{pmatrix} \check{Z}_{b(\theta)} \\ 0 \end{pmatrix} = \begin{pmatrix} V_{b(\theta)b'(\theta)}^{-1} & 0 \\ 0 & 0 \end{pmatrix} \leq V^{-1},$$

- zero at $\theta = c$, i.e., $b = \emptyset$.

Taking the center parameter $c = 0$, the estimator $\tilde{\theta}_n(0)$ obtained from any $\hat{\theta}_n$ provides a model selection procedure that removes any parameter θ_j estimated by $c_j = 0$ from $\theta = (\theta_1, \dots, \theta_d)$. This allows any model that omits some or all θ_j to be selected. In contrast, the classical Hodges' estimator $\check{\theta}(0)$ can only choose between two extreme models: the full model (corresponding to $\check{\theta}(0) \neq 0$) or the null model (corresponding to $\check{\theta}(0) = 0$), provided, with no loss of generality, that every element of the original $\hat{\theta}_n$ is nonzero.

The model selection methods derived from $\check{\theta}(0)$, $\tilde{\theta}(0)$, penalized maximum likelihood estimation (PMLE) and penalized least square estimation (PLSE) are discussed in more details below:

1. While both Hodges' estimators $\check{\theta}(0)$ and $\tilde{\theta}(0)$ can be applied as long as a good estimator $\hat{\theta}_n$ is available (r_n consistent for some constant sequence $r_n \rightarrow \infty$), PMLE can only be applied when likelihood functions are available and PLSE is limited to regression models, both under certain regularity conditions (see, e.g., Fan and Li for a set of regularity conditions) to produce root- n consistent estimators. This is demonstrated by the example discussed in Remark 2, where both $\check{\theta}(0)$ and $\tilde{\theta}(0)$ can be applied, but neither PMLE nor PLSE because the regularity conditions fail to hold.

The above comparison is limited to a fixed parameter space of dimension $p < n$. The case of $p > n$, with p varying with n , is not considered in this paper. It is noted

that PMLE/PLSE to a large extent are motivated by the need to deal with $p > n$ and may have certain advantages in such a case. The performance of PMLE/PLSE, and whether there are better estimators than PMLE/PLSE for $p > n$, may be interesting subjects for further research, which is, however, beyond the intention and scope of the present paper.

2. If $\check{\theta}_n(0)$ is used to select model by removing the parameters estimated by zero, then either all parameters are selected, or all excluded (cf. Remark 1), so that the resulting model selection does not possess oracle property. This highlights a major difference between the new type of Hodges' estimators $\tilde{\theta}_n(c)$ and the classical version $\check{\theta}_n(c)$.
3. Due to point (2) above, the model selection method based on $\check{\theta}(0)$ is limited to two candidate models only: the null model and the full model. In contrast, the other methods allow all submodels of the full model to be candidates with certain θ_i set to zero.
4. The classical Hodges' estimators are a special type of preliminary test estimators. It appears, however, that the finite sample behaviors of preliminary test estimators have not been adequately addressed in the literature either, despite the importance of this problem in statistics. Some exceptions can be found in, e.g., the book by [Judge and Bock \(1978\)](#) and the review paper of [Giles and Giles \(1993\)](#). Although PMLE/PLSE and oracle Hodges' estimators may look like a preliminary test estimation, there are fundamental differences:
 - Existing preliminary test procedures are essentially based on a single hypothesis that is either accepted or rejected as a whole.
 - The model selection procedures derived by PMLE/PLSE and oracle Hodges' estimation identify every single parameter in the multi-dimensional vector of parameters and judge if it is estimated by zero or some nonzero value. These are similar to estimation after multiple tests for a family of hypotheses (multiple tests are also known as multiple comparisons, see, e.g., [Hsu 1996](#)); it would be appropriate to call it preliminary-multiple-test estimation. In particular, the oracle Hodges' estimators provide an instance of such an estimation.
5. Properties (1)–(3) above ensure $\tilde{\theta}(0)$ model selection to possess the oracle property, provided the original estimator $\hat{\theta}_n$ is root- n consistent and efficient—which is the typical case where PMLE/PLSE can be employed to produce oracle model selection procedures and $\hat{\theta}_n$ is taken to be the efficient MLE/LSE. See, e.g., [Fan and Li \(2001\)](#) for a general discussion of the penalties producing oracle model selection procedures. As a result, Properties (1)–(3) are more general than oracle properties, and both $\tilde{\theta}(0)$ and PMLE/PLSE methods are sparse and have the same asymptotic behavior in the case of regular likelihood functions or regression models.
6. Note that, in particular, when V is a diagonal matrix, the model selection driven by $\tilde{\theta}_n(0)$ is a hard threshold model selection mentioned in [Fan and Li \(2001\)](#). If V is not diagonal, then by Definition 3, $\tilde{\theta}_j(0) = 0$ for $j \in \bar{b}(0)$ and $\tilde{\theta}_j(0)$ is obtained from $\hat{\theta}_{nj}$ for $j \notin \bar{b}(0)$ with adjustments by all estimators $\hat{\theta}_{n1}, \dots, \hat{\theta}_{nd}$, rather than solely from $\hat{\theta}_{nj}$. This also accounts for why the raw hard thresholding model selection does not have the oracle property.

7. [Leeb and Pötscher \(2005, 2008b\)](#) argued that SCAD and other sparse estimators are a return of Hodges’ estimators by examining their finite sample performance in the neighborhoods of $\theta = 0$. For any sparse estimator $\check{\theta}$ satisfying $\lim_{n \rightarrow \infty} \Pr_0(\check{\theta} = 0) = 1$, where \Pr_0 is the probability computed at $\theta = 0$, the maximum risk over a neighborhood of $\theta = 0$ tends to the maximum of the employed risk function as $n \rightarrow \infty$, and to infinity if the square loss is used. However, this section shows that SCAD or other oracle estimators are of the same asymptotic property with the new type of Hodges’ estimator $\check{\theta}_n(0)$, rather than a simple return of the classical form $\hat{\theta}_n(0)$ because $\hat{\theta}_n(0)$ does not have the oracle property.
8. PMLE maximizes the penalized likelihood functions. It is, however, generally unknown whether a PMLE can be expressed as a function of MLE $\hat{\theta}_{ML}$. In fact, PMLE is solved by direct maximization using numerical algorithms. In linear regression $\mathbf{Y} = \mathbf{X}\theta + \varepsilon$, because PLSE minimizes the penalized squares

$$\begin{aligned}
 &(\mathbf{Y} - \mathbf{X}\theta)'(\mathbf{Y} - \mathbf{X}\theta) + \sum_{i=1}^d f_i(\theta_i, \lambda_n) = \mathbf{Y}'(I - P)\mathbf{Y} \\
 &+ (\theta - \hat{\theta}_{LS})'\mathbf{X}'\mathbf{X}(\theta - \hat{\theta}_{LS}) + \sum_{i=1}^d f_i(\theta_i, \lambda_n),
 \end{aligned}$$

where P is the projection matrix onto the column space of \mathbf{X} , it is clear that PLSE is a function of the LSE $\hat{\theta}_{LS}$. Its analytical form, however, is also generally unavailable, and numerical methods are again needed to solve it. In contrast, if one takes $\hat{\theta}_n$ to be the MLE $\hat{\theta}_{ML}$ or LSE $\hat{\theta}_{LS}$ in these two cases, then both $\check{\theta}(c)$ and $\tilde{\theta}(c)$ can be expressed explicitly by $\hat{\theta}_{ML}$ or $\hat{\theta}_{LS}$ as in Definitions 2 and 3. Due to these explicit forms of $\check{\theta}(c)$ and $\tilde{\theta}(c)$, we can theoretically derive lower bounds for the performance of model selectors driven by $\check{\theta}(c)$ or $\tilde{\theta}(c)$ in finite sample size in the next section.

4 Performance of Hodges’ estimators

Hodges’ example gives a counterexample to the conjecture of R. A. Fisher that the MLE is asymptotically efficient with the smallest asymptotic variance (at rate \sqrt{n}). To overcome the difficulty thus caused, [Le Cam \(1953\)](#) and other researchers proved that the improvement in MLE can only occur in a subset of the parameter space with zero Lebesgue measure. This argument, however, did not provide any reason to rule out the use of Hodges’ estimators. An obvious question remains: if a Hodges’ estimator can outperform the MLE at even one point without paying any price, why not use it as a preferred one?

We now attempt to answer this question with the following two arguments:

1. Generally, the asymptotic bias and variance are not necessarily connected to their finite sample size versions; hence, a small asymptotic variance does not imply a small variance of an estimator even if the sample size is very large. To link the asymptotics to the finite sample size, a further condition of uniform integrability is required.

2. Even if the required uniform integrability is attached, the performance of Hodges' estimators $\tilde{\theta}_n(c)$ is still poor at the vicinity of $\theta = c$ due to a lack of uniformity in convergence over $\theta \in \Theta$.

For post-model-selection estimation in regression analysis, the lack of uniformity in convergence over parameters has been discussed by Yang (2005, 2007) and Leeb and Leeb and Pötscher (2008a, b), among others. In Sect. 4.1 below, we discuss this issue further with a more direct and elementary approach, which is made possible by the particular form of the oracle Hodges' estimators.

4.1 Uniform integrability and asymptotics

First we recall some misconception regarding asymptotic mean squared error (MSE) and asymptotic efficiency of an estimator. Given a normalized sequence $r_n(\hat{\theta}_n - \theta) \xrightarrow{d} Z$, the asymptotic mean and variance (hence MSE) are only the mean and variance of Z , but not the limit of $r_n E[\hat{\theta}_n - \theta]$ and $r_n^2 \text{Var}(\hat{\theta}_n - \theta)$ in general.

It is an easy exercise to construct examples in which a sequence of unbiased (biased) estimates might be asymptotically biased (unbiased), and even in the class of unbiased estimators, a sequence of estimates with smaller variances for every sample size n might have larger asymptotic variance, and vice versa. Consequently, the concepts of asymptotically unbiased estimate and asymptotic variance could be highly misleading when they are considered as analogies to unbiased estimate and variance in finite samples. Unfortunately, this lesson appears to have been overlooked by many statisticians for a long time, especially when one proves the oracle property of an estimator.


Let $\hat{\theta}_n$ be a sequence of estimates such that $r_n(\hat{\theta}_n - \theta) \xrightarrow{d} Z$, where $r_n \rightarrow \infty$ as $n \rightarrow \infty$ and Z is a random variable with mean zero and finite variance σ^2 . In this context, for $Z_n = \hat{\theta}_n - \theta$ and $Y_n = r_n(\hat{\theta}_n - \theta)$, we have the following two facts by Theorem 2.20 of van der Vaart (1998):

- (i) $E[Z_n] \rightarrow 0$ ($E[Y_n] \rightarrow 0$) if and only if $\{Z_n : n \geq 1\}$ ($\{Y_n : n \geq 1\}$) is uniformly integrable.
- (ii) $E[Z_n^2] \rightarrow 0$ ($E[Y_n^2] \rightarrow \sigma^2$) if and only if $\{Z_n^2 : n \geq 1\}$ ($\{Y_n^2 : n \geq 1\}$) is uniformly integrable.

Obviously, conditions like these have been completely neglected in the literature seeking estimators possessing the oracle property.

4.2 Performance of Hodges' estimators

To ensure the finite sample size quantities to approach their asymptotic versions as the sample size increases, additional conditions are required. In the remainder of this section, we assume that the required uniform integrability described above is satisfied, so that the bias and variance in finite sample size approach their asymptotic versions when the sample size is large. Even in such a case, however, both numerical and theoretical analyses below show that the overall performance of Hodges' estimators in finite sample does not match its limit as analyzed in this subsection below.



Legend: Blue curve corresponds to $n = 5$, purple to $n = 50$, and olive to $n = 500$.

Fig. 1 MSE (scaled by n) of Hodges' estimator

Due to the difficulty to obtain the exact MSE of Hodges' estimator in closed form, earlier arguments against the use of Hodges' estimators were largely based on numerical results. For example, a result for Hodges' estimate of the mean θ in the normal distribution $N(\theta, 1)$ is well known in the literature, see, e.g., [van der Vaart \(1998\)](#) and [Lehmann and Casella \(1998\)](#), which is recalled here. Let $\hat{\theta}_n = \bar{X} \sim N(\theta, 1/n)$. Taking $c = 0$ and $a_n = n^{-1/4}$ yields the original version of Hodges' estimator $\tilde{\theta}_n(0) = \bar{X}I(|\bar{X}| > n^{-1/4})$. Note that the sequences $\sqrt{n}(\hat{\theta}_n - \theta)$, $\sqrt{n}(\tilde{\theta}_n(0) - \theta)$, $n(\hat{\theta}_n - \theta)^2$ and $n(\tilde{\theta}_n(0) - \theta)^2$ are all uniformly integrable so that the asymptotic means and variances are equal to the respective limits of the means and variances in finite sample size. While the MSE (scaled by n) of \bar{X} is constant 1 for all θ 's, that of the Hodges' estimator $\tilde{\theta}_n(0)$ can only be numerically computed. Figure 1 shows the curves of the MSE of $\tilde{\theta}_n(0)$ in θ for sample sizes $n = 5, 50$ and 500 , which behave poorly in the vicinity of zero, particularly at large sample size ($n = 500$). This illustrates a much worse performance of the Hodges' estimator $\tilde{\theta}_n(0)$ than the MLE $\hat{\theta}_n$ near the center $c = 0$.

Although finding the exact closed form of the MSE (or other performance measures) of a Hodges' estimator is difficult, even in the simple case discussed just now, it turns out that some useful lower bounds of regular losses of Hodges' estimators $\tilde{\theta}_n(c)$ can be obtained to see the rationale behind their poor performance in the vicinity of the center point c as well as of the subsets of the parameter space with $\theta_j = c_j$ for some $j \in \{1, 2, \dots, d\}$. This seems to have been overlooked by the community.

The results presented below are from theoretical analyses on a general ground and distribution-free—they are valid regardless of the underlying distributions of the population.

4.3 Performance of classical Hodges' estimators

We first present the results for classical Hodges' estimators, which are in fact almost sure results.

Theorem 6 *Under the conditions of Theorem 1, given any $k > 0$, there exists a (deterministic) $N > 0$ such that for all $n > N$,*

$$r_n \|\check{\theta}_n(c) - \theta\| \geq k \text{ for all } \theta \text{ satisfying } k \leq r_n \|\theta - c\| \leq a_n r_n - k. \tag{13}$$

Furthermore, for θ_n such that $\|\theta_n - c\| = a_n/2$,

$$r_n \|\check{\theta}_n(c) - \theta_n\| \geq \frac{1}{2} r_n a_n \rightarrow \infty \text{ as } n \rightarrow \infty. \tag{14}$$

Proof Since $a_n \rightarrow 0$ and $a_n r_n \rightarrow \infty$ as $n \rightarrow \infty$ under the conditions of Theorem 1, the set $\{\theta : k \leq r_n \|\theta - c\| \leq a_n r_n - k\}$ is a nonempty ring when n is sufficiently large such that $a_n r_n > 2k$. For any $\check{\theta}_n(c)$ defined in (1), for any values the sample may take, as long as θ satisfies $k \leq r_n \|\theta - c\| \leq a_n r_n - k$, i.e., $k/r_n \leq \|\theta - c\| \leq a_n - k/r_n$, it is clear that

$$\begin{aligned} \|\check{\theta}_n(c) - \theta\| &= \|c - \theta\| I(\|\hat{\theta}_n - c\| \leq a_n) + \|\hat{\theta}_n - \theta\| I(\|\hat{\theta}_n - c\| > a_n) \\ &\geq \|c - \theta\| I(\|\hat{\theta}_n - c\| \leq a_n) + (\|\hat{\theta}_n - c\| - \|c - \theta\|) I(\|\hat{\theta}_n - c\| > a_n) \\ &\geq k/r_n I(\|\hat{\theta}_n - c\| \leq a_n) + (a_n - (a_n - k/r_n)) I(\|\hat{\theta}_n - c\| > a_n) \\ &= k/r_n. \end{aligned}$$

This proves the first assertion in (13). The second assertion (14) is obvious. □

The results of Theorem 6 have the following easy implications:

1. Formula (14) proves that even if the MSE of $\check{\theta}_n(c)$ (scaled by r_n) converges to that of the asymptotic distribution, the convergence is not generally uniform because

$$r_n^2 \max_{\theta \in \Theta} E_{\theta} [\|\check{\theta}_n(c) - \theta\|^2] \geq r_n^2 E_{\theta_n} [\|\check{\theta}_n(c) - \theta_n\|^2] \geq \frac{1}{4} r_n^2 a_n^2 \rightarrow \infty.$$

2. The same arguments also hold for more general loss functions $l(\hat{\theta}_n; \theta) = l(\|\hat{\theta}_n - \theta\|)$ with some nondecreasing function $l(u)$ in $u > 0$ satisfying $l(0) = 0$, so that the risk function scaled by $1/l(1/r_n)$ is $R_n(\hat{\theta}_n; \theta) = E_{\theta} [l(\|\hat{\theta}_n - \theta\|)/l(1/r_n)]$. Let $p = \min\{i \geq 1 : l^{(i)}(0+) \neq 0\}$. Then under the conditions of Theorem 1, for any $k > 0$, there exists $N > 0$ such that for all $n > N$,

$$\frac{l(\check{\theta}_n(c); \theta)}{l(1/r_n)} \geq \frac{l(k/r_n)}{l(1/r_n)} \rightarrow k^p \text{ as } n \rightarrow \infty$$

for all θ satisfying $k \leq r_n \|\theta - c\| \leq a_n r_n - k$. Furthermore, for all θ_n such that $\|\theta_n - c\| = a_n/2$,

$$\frac{l(\check{\theta}_n(c); \theta_n)}{l(1/r_n)} \geq \frac{l(a_n/2)}{l(1/r_n)} \rightarrow \infty \text{ as } n \rightarrow \infty.$$

The last formula also implies

$$\max_{\theta \in \Theta} R_n(\check{\theta}_n(c); \theta) \geq R_n(\check{\theta}_n(c); \theta_n) \geq \frac{l(a_n/2)}{l(1/r_n)} \rightarrow \infty \text{ as } n \rightarrow \infty.$$

3. Another way is to analyze a loss function $L(\hat{\theta}_n; \theta) = l(r_n \|\hat{\theta}_n - \theta\|)$ with a nondecreasing function $l(u)$ in $u > 0$ (cf. [Leeb and Pötscher 2008b](#)), which corresponds to a sequence of loss functions $l_n(u) = l(r_n u)$ ($r_n = \sqrt{n}$ in their paper), so that $R_n(\hat{\theta}_n; \theta) = E_\theta[L(\hat{\theta}_n; \theta)]$.

Then similarly under the conditions of Theorem 1, given any $k > 0$, there exists $N > 0$ such that $L(\check{\theta}_n(c); \theta) \geq l(k)$ for all $n > N$ and θ satisfying $k \leq r_n \|\theta - c\| \leq a_n r_n - k$. Moreover, for all θ_n such that $\|\theta_n - c\| = a_n/2$, $L(\check{\theta}_n(c); \theta_n) \geq l(r_n a_n/2) \rightarrow l(\infty)$ as $n \rightarrow \infty$. In particular, the last property indicates that

$$\max_{\theta \in \Theta} R_n(\check{\theta}_n(c); \theta) \geq E_{\theta_n}[L(\check{\theta}_n(c); \theta_n)] \rightarrow l(\infty). \tag{15}$$

This shows that even if $\lim_{n \rightarrow \infty} R_n(\check{\theta}_n(c); \theta) = R(Z; \theta) = E_\theta[l(Z; \theta)]$ for every θ pointwise, the maximum risk over θ may increasingly tend to $l(\infty)$. An example is $l(u) = I_{(z, \infty)}(u)$ for any fixed continuity point $z \in \mathbb{R}^+$ of the distribution of Z . The risk function of an estimator $\hat{\theta}_n$ with this $l(u)$ is $R_n(\hat{\theta}_n; \theta) = \Pr_\theta(r_n \|\hat{\theta}_n - \theta\| > z)$. Thus,

$$\lim_{n \rightarrow \infty} R_n(\check{\theta}_n(c); \theta) = \Pr_\theta(\|Z\| > z)I(\theta \neq c).$$

In contrast, $\Pr_\theta(r_n \|\check{\theta}_n(c) - \theta\| > z) = 1$ for all θ with $k \leq r_n \|\theta - c\| \leq a_n r_n - k$ and $z \in [0, k]$, regardless how large is n . Moreover, for any $z > 0$, if $\theta_n = \pm a_n/2$ and n is sufficiently large such that $a_n r_n > x$, then $\Pr_{\theta_n}(r_n \|\check{\theta}_n(c) - \theta_n\| > z) = 1$. Consequently,

$$\max_{\theta \in \Theta} R_n(\hat{\theta}_n; \theta) = \max_{\theta \in \Theta} \Pr_\theta(r_n \|\check{\theta}_n(c) - \theta\| > z) = 1.$$

4.4 Performance of the oracle Hodges' estimators

Now we turn to analyze the performance of the oracle Hodges' estimators $\check{\theta}_n(c)$ defined in Sect. 2.2. To simplify the exposition, we assume without loss of generality that the parameter space is $\Theta = \mathbb{R}^d$ and define the following subsets of Θ :

$$\begin{aligned} \Theta_{n1} &= \{\theta : \min_{1 \leq j \leq d} |\theta_{nj} - c_j| > a_{nj}\}, & \bar{\Theta}_{n1}^k &= \{\theta : d(\theta, \Theta_{n1}) \leq k/r_n\}, \\ \Theta_{n2} &= \{\theta : \min_{1 \leq j \leq d} |\theta_{nj} - c_j| = 0\}, & \bar{\Theta}_{n2}^k &= \{\theta : d(\theta, \Theta_{n2}) \leq k/r_n\}, \\ \bar{\Theta}_{n3}^k &= \Theta - \bar{\Theta}_{n1}^k(c) \cup \bar{\Theta}_{n2}^k(c), \end{aligned}$$

which are all dependent on the center point c . Under the uniform integrability of $\hat{\theta}_n$, Theorem 2 indicates that

$$\lim_{n \rightarrow \infty} r_n^2 E[\|\tilde{\theta}_n(c) - \theta\|^2] \leq \lim_{n \rightarrow \infty} r_n^2 E[\|\hat{\theta}_n - \theta\|^2]$$

with the strict inequality at certain values of θ . For fixed sample size n , however, we have the following contrary results, which are extensions of Theorem 6.

Theorem 7 *Under the conditions of Theorem 2, for any $k > 0$, there exists $N > 0$ such that*

$$r_n \|\tilde{\theta}_n(c) - \theta\| \geq k \text{ for all } \theta \in \bar{\Theta}_{n3}^k \quad \forall n > N, \tag{16}$$

and for any $\theta^{(n)} \in \bar{\Theta}_{n3}^k$ satisfying $\theta_j^{(n)} = c_j \pm a_{nj}/2$,

$$r_n \|\tilde{\theta}_n(c) - \theta^{(n)}\| \geq \frac{1}{2} r_n a_{nj} \rightarrow \infty \text{ as } n \rightarrow \infty. \tag{17}$$

Proof Again, we assume that N is sufficiently large such that $r_n \max_{1 \leq j \leq d} a_{nj} > k$ for all $n > N$. It can be easily shown that $d(\Theta_{n1}, \Theta_{n2}) = \min_{1 \leq j \leq d} a_{nj}$. Because for any $\theta_1 \in \bar{\Theta}_{n1}^k$ and $\theta_2 \in \bar{\Theta}_{n2}^k$,

$$d(\Theta_{n1}, \Theta_{n2}) \leq d(\Theta_{n1}, \theta_1) + d(\theta_1, \theta_2) + d(\theta_2, \Theta_{n2}) \leq \frac{2k}{r_n} + d(\theta_1, \theta_2),$$

the condition $\lim_{n \rightarrow \infty} r_n \min_{1 \leq j \leq d} a_{nj} = \infty$ implies

$$d(\bar{\Theta}_{n1}^k, \bar{\Theta}_{n2}^k) = \min_{\theta_1 \in \bar{\Theta}_{n1}^k, \theta_2 \in \bar{\Theta}_{n2}^k} d(\theta_1, \theta_2) \geq \min_{1 \leq j \leq d} a_{nj} - \frac{2k}{r_n} > 0$$

for sufficiently large n . Consequently, $\bar{\Theta}_{n1}^k \cap \bar{\Theta}_{n2}^k = \emptyset$ and $\bar{\Theta}_{n3}^k \neq \emptyset$.

By Definition 3 of $\tilde{\theta}_n(c)$, $\hat{\theta}_n \in \Theta_{n1}$ implies $\tilde{\theta}_n(c) = \hat{\theta}_n \in \Theta_{n1}$ and $\hat{\theta}_n \in \Theta - \Theta_{n1}$ implies $\tilde{b}_n(c) \neq \emptyset$, so that $\tilde{\theta}_n(c) \in \Theta_{n2}$ because $\tilde{\theta}_n, \tilde{b}_n(c)(c) - c_{\tilde{b}_n(c)} = 0$. That is, $\tilde{\theta}_n(c)$ takes values only in $\Theta_{n1} \cup \Theta_{n2}$. For any $\theta \in \bar{\Theta}_{n3}^k$, it is clear that $d(\theta, \Theta_{n1}^k \cup \Theta_{n2}^k) \geq k/r_n$. As a result, $\theta \in \bar{\Theta}_{n3}^k$ implies

$$\|\tilde{\theta}_n(c) - \theta\| \geq d(\theta, \Theta_{n1}^k \cup \Theta_{n2}^k) \geq \frac{k}{r_n}.$$

Thus, the assertion in (16) follows. The second assertion in (17) is easy to check. \square

The following extensions of Theorem 16 are minor modifications of the points presented earlier for classical Hodges' estimators.

1. Formulas (16) and (17) prove that, under the conditions in Theorem 2, for any $k > 0$, there exists $N > 0$ such that

$$r_n^2 E_\theta[\|\tilde{\theta}_n(c) - \theta\|^2] \geq k^2 \quad \text{for any } \theta \in \bar{\Theta}_{n3}^k \text{ if } n > N$$

and

$$\max_{\theta \in \Theta} r_n^2 E_\theta[\|\tilde{\theta}_n(c) - \theta\|^2] \geq \frac{1}{4} r_n^2 a_{nj}^2 \rightarrow \infty \text{ as } n \rightarrow \infty.$$

2. If we use the loss function $L(\hat{\theta}_n; \theta) = l(\|\hat{\theta}_n - \theta\|)$ with a nondecreasing function $l(u)$ in $u > 0$, so that the risk function scaled by $1/l(1/r_n)$ is $R_n(\hat{\theta}_n; \theta) = E_\theta[l(\|\hat{\theta}_n - \theta\|)/l(1/r_n)]$, then for sufficiently large n ,

$$\frac{l(\tilde{\theta}_n(c); \theta)}{l(1/r_n)} \geq \frac{l(k/r_n)}{l(1/r_n)} \rightarrow k^p \quad \text{as } n \rightarrow \infty \quad \text{for all } \theta \in \bar{\Theta}_{n3}^k;$$

and for $\theta^{(n)}$ with $\theta_j^{(n)} = c \pm a_{nj}/2$ for some j ,

$$\frac{l(\tilde{\theta}_n(c); \theta^{(n)})}{l(1/r_n)} \geq \frac{l(a_{nj}/2)}{l(1/r_n)} \rightarrow \infty \quad \text{as } n \rightarrow \infty.$$

The last formula also implies that

$$\max_{\theta \in \Theta} R_n(\tilde{\theta}_n(c); \theta) \geq R_n(\tilde{\theta}_n(c); \theta^{(n)}) \geq \frac{l(a_{nj}/2)}{l(1/r_n)} \rightarrow \infty \quad \text{as } n \rightarrow \infty.$$

3. If we analyze a loss function $L(\hat{\theta}_n; \theta) = l(r_n \|\hat{\theta}_n - \theta\|)$ with a nondecreasing function $l(u)$ in $u > 0$, as in [Leeb and Pötscher \(2008b\)](#), so that $R_n(\hat{\theta}_n; \theta) = E_\theta[L(\hat{\theta}_n; \theta)]$, then under the conditions of [Theorem 2](#), for any given $k > 0$, there exists $N > 0$ such that for all $n > N$, $L(\hat{\theta}_n(c); \theta) \geq l(k)$ for all $\theta \in \bar{\Theta}_{n3}^k$; and $L(\tilde{\theta}_n(c); \theta^{(n)}) \geq l(r_n a_{nj}/2) \rightarrow l(\infty)$ as $n \rightarrow \infty$ for $\theta^{(n)}$ with $\theta_j^{(n)} = c \pm a_{nj}/2$ for some j . The last property implies

$$\max_{\theta \in \Theta} R_n(\tilde{\theta}_n(c); \theta) \geq E_{\theta_n}[L(\tilde{\theta}_n(c); \theta_n)] \rightarrow l(\infty). \tag{18}$$

Because of this fact, even if $\lim_{n \rightarrow \infty} R_n(\tilde{\theta}_n(c); \theta) = R(Z; \theta) = E_\theta[l(Z; \theta)]$ for every θ pointwise, the maximum risk over θ may increasingly tend to $l(\infty)$.

Remark 3 We conclude this section by the following two points that highlight the difference between our work and those of others, such as [Leeb and Pötscher \(2008b\)](#).

1. [Leeb and Pötscher \(2008b\)](#) proved a result similar to [Eq. \(18\)](#), which is more general with an arbitrary sparse estimator (say, $\hat{\theta}_n$), but restricted to the regular case of $r_n = \sqrt{n}$, $c = 0$ and normally distributed Z . It also requires the condition that $P_{n,k/\sqrt{n}}$ is contiguous with respect to $P_{n,0}$, where $P_{n,\theta}$ is the distribution of $\hat{\theta}_n$, or a stronger condition that $P_{n,\theta}$ is locally asymptotically normal. In comparison, we obtained stronger results expressed in [\(16\)](#) and [\(17\)](#) for classical and oracle Hodges' estimators, which hold almost surely without such conditions as \sqrt{n} -consistency and contiguity.
2. [Leeb and Pötscher](#) revealed the erratic behavior of a sparse estimator only in the vicinity of 0, whereas we here showed that the erratic behavior of the oracle model selection procedure derived by a Hodges' estimator occurs not only in the vicinity of the point c but also in the vicinity of every subset of Θ with some component $\theta_j = c_j$.

5 Further discussions

To sum up, we have in this paper demonstrated that:

1. the oracle model selection procedures are not a simple return of the classical Hodges' estimators but more like *oracle Hodges' estimators* in asymptotic sense;
2. properly constructed oracle Hodges' estimators can easily generate oracle model selection procedures that satisfy the requirements of continuity or smoothing;
3. under the MSE criterion, the oracle Hodges' estimator $\hat{\theta}_n(c)$ does not outperform its origin $\hat{\theta}_n$ in finite sample size, despite having a smaller asymptotic variance than $\hat{\theta}_n$; and
4. Hodges' estimators possessing the oracle property can perform much worse than their origins under the minimax criterion.

Similar results to (3) and (4) can be found in Yang (2005, 2007) and the works by Leeb and his collaborators as mentioned above. The difference here is that the particular form of the oracle Hodges' estimators makes the proof quite obvious.

Points (3) and (4) above also provide an answer to the question why Hodges' estimators are not preferable to use even if they can improve the asymptotic efficiency at no cost. The key reasons behind this answer, as discussed earlier, are the disconnection between the performances of finite sample statistics and their asymptotics in certain situations (principally due to the lack of uniformity in integrability of the statistics), and the universal lack of uniformity in the convergence in the situations where even the convergence is guaranteed.

In addition, a few points worth for further discussions are listed below.

5.1 Uniformity in integrability and the convergence

By ignoring the uniformity in integrability and convergence, the widely adopted concept of asymptotically efficient estimation tends to place too much weight on the asymptotic distribution, leading to inappropriate use of asymptotic variance to measure the goodness of estimators. This creates the concept of superefficiency and supports Hodges estimators as superior, despite their poor performance in finite sample size. Because Hodges' estimators can be easily applied to generate model selection procedures possessing the oracle property, as demonstrated in point (2) above, our results also suggest that the oracle property itself is not a convincing reason to recommend the application of such oracle model selection procedures in theoretical studies or real-life applications.

This phenomenon was in fact recognized by statisticians decades ago. See, for example, Hájek (1971, p. 153)'s warning quoted below (which can also be found in Hušková et al. 1998, pp. 613–614):

“Having obtained an asymptotic result we are not usually able to tell how far it applies to particular cases with finite n Consequently, in applications we are guided by two epistemologically very different knowledge: (i) we have limit theorems giving some hope, but not assurance, of practical sample sizes; (ii) we work with some numerical experience, which we extend to cases that seem to

us to be similar. . . . Especially misinformative can be those limit results that are not uniform. Then the limit may exhibit some features that are not even approximately true for any finite n Superefficient estimates produced by L.J. Hodges (see [Le Cam 1953](#)) have their amazing properties only in the limit. For any finite n they behave quite poorly for some parameter values. These values, however, depend on n and disappear in the limit.”

Therefore, in the cases relying on asymptotic distribution of a statistic, uniform integrability and convergence of the statistic are important and desired properties to ensure the proximity between the finite sample version and its asymptotics, especially when the statistic does not have an analytical form and can only be obtained by numerical computations.

5.2 The power of simulations

Undoubtedly, with the development of contemporary computing facilities, Monte Carlo simulations provide more and more accessible and powerful tools in exploring properties of statistical inferences. This powerful tool, however, appears being applied excessively to justify certain theoretical properties that are difficult to prove mathematically. It is even a common practice that, instead of providing theoretical results on finite sample behaviors, many statisticians justify their statistical methods by finding and theoretically proving asymptotic properties of their statistical methods and then supporting their finite sample properties by Monte Carlo simulations.

Therefore, another point worth to mention is the need to avoid the pitfall of relying too much on simulations to judge the merits of a property or a method, particularly in the era of fast advancing computer technology and capacity, as demonstrated by [Leeb and Pötscher \(2008b\)](#) on the SCAD estimator in Example 4.1 of [Fan and Li \(2001\)](#). This pitfall is generally recognized in the statistical community, but appears often overlooked or ignored in a large volume of publications in the pursuit of finding new and exciting ideas and methods. Logically, simulation results can be convincing to counter a claim or conjecture made on a general ground, but not to support such a claim or conjecture, because it is impossible to exhaust all possible scenarios by simulations. An assertion on property or goodness of a statistical procedure should be theoretically examined whether it is to be claimed on a general ground, rather than relies on limited numerical simulations only.

5.3 Generality of the center point

The generality of the center point $c = (c_1, \dots, c_d)$ implies that, mathematically, one can make model selection not only on pre-specified θ_i but also on more general $\theta_i - c_i$, $1, \dots, d$. Take the oracle Hodges' estimators as an example, for any given constant $c = (c_1, \dots, c_d)$, one can make a model selection mathematically with the collection of candidate models defined by setting $\theta_i = 0$ for one or more $i \in \{1, \dots, d\}$. Similarly for PMLE/PLSE, the common practice is to use penalty $f_i(|\theta_i|, \lambda_n)$ to select the “best” with certain $\theta_i = 0$ from the candidates. By the same method and logic,

for any fixed $c = (c_1, \dots, c_d)$, it is equally feasible to use penalty $f_i(|\theta_i - c_i|, \lambda_n)$ to select the “best” model with $\theta_i = c_i$ for certain i . This raises an obvious question in practice: why does everyone only take $c = 0$ for model selection, not any other c ? As an example, if m analysts under a common supervisor independently analyze a same data set with parameter θ of interest using oracle Hodges' estimation, PMLE or PLSE with different points $c^{(1)}, \dots, c^{(m)}$, including zero or not, then they obtain m different estimates $\hat{\theta}^{(j)}$, $j = 1, \dots, m$. Clearly, all these estimators are of similar statistical properties and rationales, but there is no statistically sensible guidance for the supervisor to decide which one should be used. This question appears difficult to answer both logically and philosophically, and it raises a sobering question on the validity of the commonly adopted model selection procedures.

5.4 Model selection for predictions

In real applications of regression analysis, as well as in popular textbooks, variable selection is also discussed with a purpose of prediction, or equivalently, estimation of the regression function at certain points. In this paper, our efforts are focussed on the performance of the parameter estimators in model selection. This does not, however, point to a lack of generality of our results, due to the following two reasons:

1. In the case of linear regression with predictors represented as a p -vector \mathbf{x} (whether the intercept term 1 is included or not is irrelevant), taking θ as the vector of regression coefficients. Suppose that the regression function is to be estimated at d point $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_d$ under the performance measure $A = E\|\tilde{\mathbf{X}}\hat{\theta} - \tilde{\mathbf{X}}\theta\|^2 = \sum_{i=1}^d E(\tilde{\mathbf{x}}_i'\hat{\theta} - \tilde{\mathbf{x}}_i'\theta)^2$, where $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_d)'$ collects the data of the predictors at the d prediction points. Then the fundamental properties of positive definite matrices state that the order of different estimates under measure A can be derived from the order of the covariance matrices of the corresponding estimates, given the presence of the latter, in the sense that for two matrices C and D , $C > D$ if and only if $\mathbf{x}'C\mathbf{x} > \mathbf{x}'D\mathbf{x}$ for any dimension-compatible vector \mathbf{x} . As a result, the performance of the parameter estimators in our discussions is equivalent to that of the regression function estimator.
2. In case the regression function is possibly nonlinear in the parameters θ , including generalized linear models and general parametric nonlinear regression models, the similar arguments work asymptotically with the help of linear approximation of statistics (i.e., the commonly known delta method).

5.5 Open questions

Back to the controversy on the merits of the oracle property and the efforts to find oracle model selection procedures, the poor performance of Hodges' estimators and their oracle property appear to support the criticisms of the oracle estimators and cast serious doubts on the usefulness of oracle procedures. This further casts doubts in the validity and usefulness of prevailing model selection methods. We believe that the

following open and challenging questions need to be convincingly answered before a consensus can be reached one way or the other:

1. Without counting on the oracle property, are there any theoretical properties that ensure good performance of the popular model selection procedures, such as LASSO and SCAD? For example, is there any reasonable loss function beyond the regular ones discussed in this paper such that some of the popular model selection procedures can perform well?
2. Our arguments are limited to the case where the MLE/LSE exists so that $n > d$ is implicitly assumed. The cases of $d > n$, $d \gg n$, or more generally, the design matrices of reduced-ranks, have so far been discussed under the assumption of parsimony with few exception, which are significantly different from the classical model. We have not yet come up with a clear idea on how the oracle model selection performs under this significantly different setting in finite sample size. Further efforts are needed to examine various important topics regarding $d > n$ or $d \gg n$, including model selection, which will be in the agenda of our future research works.

These open questions call for further research efforts to investigate. Their answers will help resolve the controversy on the oracle property and thus point to the right direction of research on model selection.

A Appendix: Proofs of the theorems

A.1 Proof of Theorem 1

Proof Note first that $\hat{\theta}_n \xrightarrow{P} \theta$. For any $\theta \neq c$, the condition $a_n = o(1)$ implies that, for any $\varepsilon > 0$,

$$\begin{aligned} \Pr_{\theta}(r_n \|\check{\theta}_n(c) - \hat{\theta}_n\| > \varepsilon) &\leq \Pr_{\theta}(\|\hat{\theta}_n - c\| \leq a_n) \leq \Pr_{\theta}(\|\theta - c\| - \|\hat{\theta}_n - \theta\| \leq a_n) \\ &= \Pr_{\theta}(\|\hat{\theta}_n - \theta\| \geq \|\theta - c\| - a_n) \rightarrow 0 \quad \text{as } n \rightarrow \infty. \end{aligned}$$

Thus, $r_n(\check{\theta}_n(c) - \theta) = r_n(\check{\theta}_n(c) - \hat{\theta}_n) + r_n(\hat{\theta}_n - \theta) \xrightarrow{d} Z$. For $\theta = c$, thanks to $r_n a_n \rightarrow \infty$,

$$\Pr_c(r_n \|\check{\theta}_n(c) - c\| > \varepsilon) \leq \Pr_c(\check{\theta}_n \neq c) = \Pr_{\theta_0}(r_n \|\hat{\theta}_n - c\| > r_n a_n) \rightarrow 0. \quad (19)$$

This shows $r_n(\check{\theta}_n(c) - c) \xrightarrow{P} 0$. □

A.2 Proof of Theorem 2

Proof Because we are concerned with the asymptotic distribution of $\check{\theta}_{n,b}$ in this section, without loss of generality we can treat the easy case where V is known and $\check{\theta}_{n,b}$ is defined by

$$\check{\theta}_{n,b} = \hat{\theta}_{n,b} + V_{bb}^{-1} V_{b\bar{b}}(\hat{\theta}_{n,\bar{b}} - c_{\bar{b}}) \quad \text{and} \quad \check{\theta}_n(b) = (\check{\theta}'_{n,b}, c'_{\bar{b}})' \tag{20}$$

with the convention $\check{\theta}_{n,\{1,2,\dots,d\}} = \check{\theta}_n(\{1, 2, \dots, d\}) = \hat{\theta}_n$.

The first assertion is obvious, so we here only prove (9). With the definition of $b(\theta)$ in (4), it is clear that $\theta_{b(\theta)}$ is the sub-vector $(\theta_j : \theta_j \neq c_j)$ of θ and $\theta_{\bar{b}(\theta)} = (c_j : j \in \bar{b}(\theta)) = c_{\bar{b}(\theta)}$. Note that, by (20), $\check{\theta}_{n,b(\theta)} = \hat{\theta}_{n,b(\theta)}$ and hence $\check{\theta}_n(b(\theta)) = (\check{\theta}_{n,b(\theta)}, c_{\bar{b}(\theta)})$ are only pseudo-estimators that depend on the unknown parameters θ . However,

$$\begin{aligned} r_n(\check{\theta}_{n,b(\theta)} - \theta_{b(\theta)}) &= V_{b(\theta),b(\theta)}^{-1} (V_{b(\theta)b(\theta)} \quad V_{b(\theta),\bar{b}(\theta)}) r_n \begin{pmatrix} \hat{\theta}_{n,b(\theta)} - \theta_{b(\theta)} \\ \hat{\theta}_{n,\bar{b}(\theta)} - \theta_{\bar{b}(\theta)} \end{pmatrix} \\ &\stackrel{d}{\rightarrow} V_{b(\theta),b(\theta)}^{-1} (V_{b(\theta)b(\theta)} \quad V_{b(\theta),\bar{b}(\theta)}) Z = \check{Z}_{b(\theta)}, \end{aligned} \tag{21}$$

where the components of Z have been rearranged according to the order in θ and $\check{Z}_{b(\theta)}$ is derived from (7) by replacing b with $b(\theta)$.

For any θ , by comparing (6) and (20),

$$\begin{aligned} \Pr_{\theta} \left(r_n \|\tilde{\theta}_n(c) - \check{\theta}_n(b(\theta))\| > \varepsilon \right) &= \Pr_{\theta} \left(r_n \|\tilde{\theta}_n(c) - (\check{\theta}'_{b(\theta)}, c'_{\bar{b}(\theta)})'\| > \varepsilon \right) \\ &\leq \Pr_{\theta} \left(\tilde{\theta}_n(c) \neq (\check{\theta}'_{b(\theta)}, c'_{\bar{b}(\theta)})' \right) \\ &\leq \Pr_{\theta} (b_n(c) \neq b(\theta)). \end{aligned}$$

Since $\{b_n(c) \neq b(\theta)\} = \bigcup_{j \in b(\theta)} \{|\hat{\theta}_{nj} - c_j| \leq a_{nj}\} \bigcup_{j \in \bar{b}(\theta)} \{|\hat{\theta}_{nj} - c_j| > a_{nj}\}$,

$$\begin{aligned} &\Pr_{\theta} \left(r_n \|\tilde{\theta}_n(c) - \check{\theta}_n(b(\theta))\| > \varepsilon \right) \\ &\leq \Pr_{\theta} \left(\bigcup_{j \in b(\theta)} \{|\hat{\theta}_{nj} - c_j| \leq a_{nj}\} \bigcup_{j \in \bar{b}(\theta)} \{|\hat{\theta}_{nj} - c_j| > a_{nj}\} \right) \\ &\leq \sum_{j \in b(\theta)} \Pr_{\theta} \{|\hat{\theta}_{nj} - c_j| \leq a_{nj}\} + \sum_{j \in \bar{b}(\theta)} \Pr_{\theta} \{|\hat{\theta}_{nj} - c_j| > a_{nj}\} \\ &\leq \sum_{j \in b(\theta)} \Pr_{\theta} \{|\hat{\theta}_{nj} - \theta_j| \geq |\theta_j - c_j| - a_{nj}\} + \sum_{j \in \bar{b}(\theta)} \Pr_{\theta} \{r_n |\hat{\theta}_{nj} - c_j| > r_n a_{nj}\} \\ &\rightarrow 0 \quad \text{as } n \rightarrow \infty \end{aligned}$$

under the conditions on a_n . Thus, $r_n(\tilde{\theta}_n(c) - \check{\theta}_n(b(\theta))) = o_p(1)$. Combining this with (21), we get

$$r_n(\tilde{\theta}_n(c) - \theta) = r_n(\tilde{\theta}_n(c) - \check{\theta}_n(b(\theta))) + r_n(\check{\theta}_n(b(\theta)) - \theta) \stackrel{d}{\rightarrow} \begin{pmatrix} \check{Z}_{b(\theta)} \\ 0 \end{pmatrix} \tag{22}$$

under \Pr_θ . Next examine the case $b(\theta) = \emptyset$. Analogous to (19), under \Pr_c and the condition $r_n \min_{1 \leq j \leq d} a_{nj} \rightarrow \infty$,

$$\begin{aligned} \Pr_c(r_n \|\tilde{\theta}_n(c) - c\| > \varepsilon) &\leq \Pr_c(\tilde{\theta}_n \neq c) \\ &\leq \sum_{j=1}^d \Pr_c(r_n |\hat{\theta}_{nj} - c| > r_n a_{nj}) \rightarrow 0 \quad \text{as } n \rightarrow \infty. \end{aligned} \quad (23)$$

Hence $r_n |\tilde{\theta}_n(c) - c| \xrightarrow{d} 0$. The proof is then complete. \square

References

- Berk, R., Brown, L., Zhao, L. (2010). Statistical inference after model selection. *Journal of Quantitative Criminology*, 26(2), 217–236.
- Berk, R., Brown, L., Buja, A., Zhang, K., Zhao, L. (2013). Valid post-selection inference. *The Annals of Statistics*, 41(2), 802–837 (with an electronic supplement).
- Brookhart, M., Schneeweiss, S., Rothman, K., Glynn, R., Avorn, J., Strmer, T. (2006). Variable selection for propensity score models. *American Journal of Epidemiology*, 163(12), 1149–1156.
- Fan, J., Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96, 1348–1360.
- Fan, J., Li, R. (2002). Variable selection for Cox's proportional hazards model and frailty model. *Annals of Statistics*, 30, 74–99.
- Fan, J., Li, R. (2004). New estimation and model selection procedures for semiparametric modeling in longitudinal data analysis. *Journal of the American Statistical Association*, 99, 710–723.
- Fan, J., Feng, Y., Wu, Y. (2009). Network exploration via the adaptive LASSO and SCAD penalties. *The Annals of Applied Statistics*, 3(2), 521–541.
- Frank, I. E., Friedman, J. H. (1993). A statistical view of some chemometrics regression tools. *Technometrics*, 35, 109–148.
- Friedman, J., Hastie, T., Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1), 1–22.
- Gefang, D. (2014). Bayesian doubly adaptive elastic-net Lasso for VAR shrinkage. *International Journal of Forecasting*, 30(1), 1–11.
- Giles, J. A., Giles, D. E. A. (1993). Pre-test estimation and testing in econometrics: Recent developments. *Journal of Economic Surveys*, 7, 145–197.
- Hajek, J. (1971). Limiting properties of likelihoods and inference: In V.P. Godambe, D.A. Sprott (Eds.), *Foundations of statistical inference: Proceedings of the symposium on the foundations of statistical inference*, University of Waterloo, Ontario, March 31–April 9, 1970 (pp. 142–159). Holt, Rinehart & Winston, Toronto.
- Horowitz, J., Klemela, J., Mammen, E. (2006). Optimal estimation in additive regression models. *Bernoulli*, 12, 271–298.
- Huang, J., Horowitz, J., Wei, F. (2010). Variable selection in nonparametric additive models. *Annals of Statistics*, 38(4), 2282–2313.
- Hsu, J. (1996). *Multiple comparisons: Theory and methods*. New York: Chapman and Hall.
- Huřková, M., Beran, R., Dupac, V. (1998). *Collected works of Jaroslav Hájek: With commentary*. Chichester: Wiley.
- Judge, G. G., Bock, M. E. (1978). *The statistical implications of pre-test and stein-rule estimators in econometrics*. Amsterdam: North-Holland Publishing Co.
- Kock, A. (2013). Oracle efficient variable selection in random and fixed effects panel data models. *Econometric Theory*, 29(1), 115–152.
- Le Cam, L. (1953). On some asymptotic properties of maximum likelihood estimates and related Bayes estimates. *University of California Publications in Statistics*, 1, 277–330.
- Le Cam, L. (1960). Locally asymptotically normal families of distributions. *University of California Publications in Statistics*, 3, 37–98.

- Leeb, H., Pötscher, B. M. (2005). Model selection and inference: Facts and fiction. *Econometric Theory*, 21, 21–59.
- Leeb, H., Pötscher, B. M. (2008a). Model selection. In T. G. Anderson, R. A. Davis, J. P. Kreiss, T. Mikosch (Eds.), *The handbook of financial time series* (pp. 785–821). New York: Springer.
- Leeb, H., Pötscher, B. M. (2008b). Sparse estimators and oracle property, or the return of Hodges' estimator. *Journal of Econometrics*, 142, 201–211.
- Lehmann, E. L., Casella, G. (1998). *Theory of point estimation*. New York: Springer.
- Leng, C., Tran, M., Nott, D. (2014). Bayesian adaptive lasso. *Annals of the Institute of Statistical Mathematics*, 66(2), 221–244.
- Pötscher, B. M. (2009). Confidence sets based on sparse estimators are necessarily large. *Sankhyā: The Indian Journal of Statistics, Series A*, 71(1), 1–18.
- Pötscher, B. M., Leeb, H. (2009). On the distribution of penalized maximum likelihood estimators: The LASSO, SCAD, and thresholding. *Journal of Multivariate Analysis*, 100, 2065–2082.
- Pötscher, B. M., Schneider, U. (2009). On the distribution of the adaptive Lasso estimator. *Journal of Statistical Planning and Inference*, 139(8), 2775–2790.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B (Methodological)*, 58(1), 267–288.
- Van de Geer, S. (2008). High-dimensional generalized linear models and the lasso. *The Annals of Statistics*, 36, 614–645.
- van der Vaart, A. W. (1998). *Asymptotic statistics*. Cambridge: Cambridge University Press.
- Wang, H., Leng, C. (2008). A note on adaptive group lasso. *Computational Statistics and Data Analysis*, 52(12), 5277–5286.
- Wang, S., Nan, B., Rosset, S., Zhu, J. (2011). Random lasso. *The Annals of Applied Statistics*, 5(1), 468–485.
- Wang, Z., Paterlini, S., Gao, F., Yang, Y. (2014). Adaptive minimax regression estimation over sparse ℓ_q -hulls. *Journal of Machine Learning Research*, 15(May), 1675–1711.
- Xiong, S., Dai, B., Qian, P. Z. (2017). Achieving the oracle property of OEM with nonconvex penalties. *Statistical Theory and Related Fields*, 1(1), 28–36.
- Yang, Y. (2005). Can the strengths of AIC and BIC be shared? A conflict between model identification and regression estimation. *Biometrika*, 92(4), 937–950.
- Yang, Y. (2007). Prediction/Estimation with simple linear models: is it really that simple? *Econometric Theory*, 23(1), 1–36.
- Zhang, C. (2010). Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38, 894–942.
- Zhang, H., Lu, W. (2007). Adaptive Lasso for Cox's proportional hazards model. *Biometrika*, 94(3), 691–703.
- Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, 101(476), 1418–1429.
- Zou, H., Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B*, 67, 301–320.
- Zou, H., Zhang, H. (2009). On the adaptive elastic-net with a diverging number of parameters. *Annals of Statistics*, 37(4), 1733.