



Munich Personal RePEc Archive

Finite-sample and asymptotic analysis of generalization ability with an application to penalized regression

Xu, Ning and Hong, Jian and Fisher, Timothy

School of Economics, University of Sydney, School of Economics,
University of Sydney, School of Economics, University of Sydney

10 September 2016

Online at <https://mpra.ub.uni-muenchen.de/73622/>

MPRA Paper No. 73622, posted 12 Sep 2016 11:14 UTC

Finite-sample and asymptotic analysis of generalization ability with an application to penalized regression *

Ning Xu

School of Economics, University of Sydney
n.xu@sydney.edu.au

Jian Hong

School of Economics, University of Sydney
jian.hong@sydney.edu.au

Timothy C.G. Fisher

School of Economics, University of Sydney
tim.fisher@sydney.edu.au

September 10, 2016

Abstract

In this paper, we study the generalization ability (GA)—the ability of a model to predict outcomes in new samples from the same population—of the extremum estimators. By adapting the classical concentration inequalities, we propose upper bounds for the empirical out-of-sample prediction error for extremum estimators, which is a function of the in-sample error, the severity of heavy tails, the sample size of in-sample data and model complexity. The error bounds not only serve to measure GA, but also to illustrate the trade-off between in-sample and out-of-sample fit, which is connected to the traditional bias-variance trade-off. Moreover, the bounds also reveal that the hyperparameter K , the number of folds in K -fold cross-validation, cause the bias-variance trade-off for cross-validation error, which offers a route to hyperparameter optimization in terms of GA. As a direct application of GA analysis, we implement the new upper bounds in penalized regression estimates for both $n \geq p$ and $n < p$ cases. We show that the \mathcal{L}_2 norm difference between penalized and un-penalized regression estimates can be directly explained by the GA of the regression estimates and the GA of empirical moment conditions. Finally, we show that all the penalized regression estimates are L_2 consistent based on GA analysis.

Keywords: generalization ability, upper bound of generalization error, penalized regression, bias-variance trade-off, lasso, high-dimensional data, cross-validation, \mathcal{L}_2 difference between penalized and unpenalized regression.

*The authors would like to thank Mike Bain, Colin Cameron, Peter Hall and Tsui Shengshang for valuable comments on an earlier draft. We would also like to acknowledge participants at the 12th International Symposium on Econometric Theory and Applications and the 26th New Zealand Econometric Study Group as well as seminar participants at Utah, UNSW, and University of Melbourne for useful questions and comments. Fisher would like to acknowledge the financial support of the Australian Research Council, grant DP0663477.

1 Introduction

Traditionally in econometrics, an estimation method is implemented on sample data in order to infer patterns in a population. Put another way, inference centers on generalizing to the population the pattern learned from the sample and evaluating how well the sample pattern fits the population. An alternative perspective is to consider how well a sample pattern fits another sample. In this paper, we study the ability of a model estimated from a given sample to fit new samples from the same population, referred to as the **generalization ability** (GA) of the model. As a way of evaluating the external validity of sample estimates, the concept of GA has been implemented in recent empirical research. For example, in the policy evaluation literature [Belloni et al., 2013, Gechter, 2015, Dolton, 2006, Blundell et al., 2004], the central question is whether any treatment effect estimated from a pilot program can be generalized to out-of-sample individuals. Similarly, for economic forecasting, Stock and Watson [2012] used GA as a criterion to pick optimal weight coefficients for model averaging predictors. Generally speaking, a model with higher GA will be more appealing for policy analysis or prediction.

With a new sample at hand, GA is easily measured using validation or cross-validation to measure the goodness of fit of an estimated model on out-of-sample data. Without a new sample, however, it can be difficult to measure GA *ex ante*. In this paper, we demonstrate how to quantify the GA of an in-sample estimate when only a single sample is available by deriving upper bounds on the empirical out-of-sample errors. The upper bounds on the out-of-sample errors depend on the sample size, an index of the complexity of the model, a loss function, and the distribution of the underlying population. As it turns out, the bounds serve not only as a measurement of GA, but also illustrate the trade-off between in-sample fit and out-of-sample fit. By modifying and adapting the bounds, we are also able to analyze the performance of K -fold cross-validation and penalized regression. Thus, the GA approach yields insight into the finite-sample and asymptotic properties of penalized regression as well as cross-validation.

As well as being an out-of-sample performance indicator, GA may also be used for model selection. Arguably, model selection is coming to the forefront in empirical work given the increasing prevalence of high-dimensional data in economics and finance. We often desire a smaller set of predictors in order to gain insight into the most relevant relationships between outcomes and covariates. Model selection based on GA not only offers improved interpretability of an estimated model, but, critically, it also improves the bias-variance trade-off relative to the traditional extremum estimation approach.

1.1 Traditional approach to the bias-variance trade-off

Without explicitly introducing the concept of GA, the classical econometrics approach to model selection focusses on the bias-variance trade-off, yielding methods such as the information criteria (IC), cross-validation, and penalized regression. For example, an IC may be applied to linear regression

$$Y = X\beta + u$$

where $Y \in \mathbb{R}^n$ is a vector of outcome variables, $X \in \mathbb{R}^{n \times p}$ is a matrix of covariates and $u \in \mathbb{R}^n$ is a vector of i.i.d. random errors. The parameter vector $\beta \in \mathbb{R}^p$ may be sparse in the sense that many of its elements are zero. Model selection typically involves using a score or penalty function that depends on the data [Heckerman et al., 1995], such as the Akaike information criterion [Akaike, 1973], Bayesian information criterion [Schwarz, 1978], cross-validation errors [Stone, 1974, 1977] or the mutual information score among variables [Friedman et al., 1997, 2000].

An alternative approach to model selection is penalized regression, implemented through the objective function:

$$\min_{b_\lambda} \frac{1}{n} (\|Y - Xb_\lambda\|_2)^2 + \lambda \|b_\lambda\|_\gamma \quad (1)$$

where $\|\cdot\|_\gamma$ is the \mathcal{L}_γ norm and $\lambda \geq 0$ is a penalty parameter. One way to derive the penalized regression estimates b_λ is through validation, summarized in Algorithm 1.

Algorithm 1: Penalized regression estimation under validation

-
1. Set $\lambda = 0$.
 2. Partition the sample into a training set T and a test set S . Standardize all variables (to ensure the penalized regression residual e satisfies $\mathbb{E}(e) = 0$ in T and S).
 3. Compute the penalized regression estimate b_λ on T . Use b_λ to calculate the prediction error on S .
 4. Increase λ by a preset step size. Repeat 2 and 3 until $b_\lambda = \mathbf{0}$.
 5. Select b_{pen} to be the b_λ that minimizes the prediction error on S .
-

As shown in Algorithm 1, validation works by solving the constrained minimization problem in eq. (1) for each value of the penalty parameter λ to derive a b_λ . When the feasible range of λ is exhausted, the estimate that produces the smallest out-of-sample error among all the estimated $\{b_\lambda\}$ is chosen as the penalized regression estimate, b_{pen} .

Note in eq. (1) that if $\lambda = 0$, the usual OLS estimator is obtained. The IC can be viewed as special cases with $\lambda = 1$ and $\gamma = 0$. The lasso [Tibshirani, 1996] corresponds to the case with $\gamma = 1$ (an \mathcal{L}_1 penalty). When $\gamma = 2$ (an \mathcal{L}_2 penalty), we have the familiar ridge estimator [Hoerl and Kennard, 1970]. For any $\gamma > 1$, we have the bridge estimator [Frank and Friedman, 1993], proposed as a generalization of the ridge.

A range of consistency properties have been established for the IC and penalized regression. Shao [1997] proves that various IC and cross-validation are consistent in model selection. Breiman [1995], Chickering et al. [2004] show that the IC have drawbacks: they tend to select more variables than necessary and are sensitive to small changes in the data. Zhang and Huang [2008], Knight and Fu [2000], Meinshausen and Bühlmann [2006], Zhao and Yu [2006] show that \mathcal{L}_1 -penalized regression is consistent in different settings. Huang et al. [2008], Hoerl and Kennard [1970] show the consistency of penalized regression with $\gamma > 1$. Zou [2006], Caner [2009], Friedman et al. [2010] propose variants of penalized regression in different scenarios and Fu [1998] compares different penalized regressions using a simulation study. Alternative approaches to model selection, such as combinatorial search algorithms may be computationally challenging to implement, especially with

high-dimensional data.¹

1.2 Major results and contribution

A central idea in this paper is that the analysis of GA is closely connected to the bias-variance trade-off. We show below that, loosely speaking, a model with superior GA typically achieves a better balance between bias and variance. Put another way, GA can be thought of as a way to understand the properties of model selection methods. By the same token, model selection can be thought of as a tool for GA: if the goal is to improve the GA of a model, model selection is necessary. From the joint perspective of GA and model selection, we unify the class of penalized regressions with $\gamma > 0$, and show that the finite-sample and asymptotic properties of penalized regression are closely related to the concept of GA.

The **first contribution** of this paper is to derive an upper bound for the prediction error on out-of-sample data based on the in-sample prediction error of the extremum estimator and to characterize the trade-off between in-sample fit and out-of-sample fit. As shown in Vapnik and Chervonenkis [1971a,b], McDonald et al. [2011], Smale and Zhou [2009], Hu and Zhou [2009], the classical concentration inequalities underlying GA analysis focus on the relation between the population error and the empirical in-sample error. In contrast, we quantify a bound for the prediction error of the extremum estimate from in-sample data on any out-of-sample data. The bound also highlights that the finite-sample and asymptotic properties of many penalized estimators can be framed in terms of GA. Classical methods to improve GA involve computing discrete measures of model complexity, such as the VC dimension, Radamacher dimension or Gaussian complexity. Discrete complexity measures are hard to compute and often need to be estimated. In contrast, we show that finite-sample GA analysis is easy to implement via validation or cross-validation and possesses desirable finite-sample and asymptotic properties for model selection.

A **second contribution** of the paper is to show that GA analysis may be used to choose the tuning hyper-parameter for validation (i.e., the ratio of training sample size to test sample size) or cross-validation (i.e., the number of folds K). Existing research has studied cross-validation for parametric and nonparametric model estimation [Hall and Marron, 1991, Hall et al., 2011, Stone, 1974, 1977]. In contrast, by adapting the classical error bound inequalities that follow from GA analysis, we derive the optimal tuning parameters for validation and cross-validation in a model-free setting. We also show how K affects the bias-variance trade-off for cross-validation: a higher K increases the variance and lowers the bias.

A **third contribution** of the paper is use GA analysis to derive the finite-sample and asymptotic properties, in particular that of \mathcal{L}_2 -consistency, for any penalized regression estimate. Various properties for penalized regression estimators have previously been established, such as probabilistic consistency or the oracle property [Knight and Fu, 2000, Zhao and Yu, 2006, Candes and Tao, 2007, Meinshausen and Yu, 2009, Bickel et al., 2009]. GA analysis reveals that similar properties can be established more generally for a wider class of estimates from penalized regression. We also

¹Chickering et al. [2004] point out that the best subset selection method is unable to deal with a large number of variables, heuristically 30 at most.

show that the \mathcal{L}_2 -difference between the OLS estimate and any penalized regression estimate can be quantified by their respective GAs.

Lastly, a **fourth contribution** of the paper is that our results provide a platform to extend GA analysis to time series, panel data and other non-i.i.d. data. The literature has demonstrated that the major tools of GA analysis can be extended to non-i.i.d. data: many researchers have generalized the VC inequality [Vapnik and Chervonenkis, 1971a,b]—one of the major tools in this paper to analyze i.i.d. data—to panel data and times series. Other studies show a number of ways to control for heterogeneity, which guarantees the validity of GA analysis. In addition, other tools used in this paper, such as the Glivenko–Cantelli theorem, the Hoeffding and von Bahr–Esseen bounds, have been shown to apply to non-i.i.d. data.² Hence, by implementing our framework with the techniques listed above, we can extend the results in this paper to a rich set of data types and scenarios.

The paper is organized as follows. In Section 2 we review the concept of GA, its connection to validation and cross-validation and derive upper bounds for the finite-sample GA of extremum estimates. In Section 3, we implement the results in the case of penalized regression and show that properties of penalized regression estimates can be explained and quantified by their GA. We also prove the \mathcal{L}_2 -consistency of penalized regression estimates for both $p \leq n$ and $p > n$ cases. Further, we establish the finite-sample upper bound for the \mathcal{L}_2 -difference between penalized and unpenalized estimates based on their respective GAs. In Section 4, we use simulations to demonstrate the ability of penalized regression to control for overfitting. Section 5 concludes with a brief discussion of our results. Proofs are contained in Appendix 1 and graphs of the simulations are in Appendix 2.

2 Generalization ability and the upper bound for finite-sample generalization errors

2.1 Generalization ability, generalization error and overfitting

In econometrics, choosing the best approximation to data often involves measuring a loss function, $Q(b|y_i, \mathbf{x}_i)$, defined as a functional that depends on the estimate b and the sample points (y_i, \mathbf{x}_i) . The population error functional is defined as

$$\mathcal{R}(b|Y, X) = \int Q(b|y, \mathbf{x}) dF(y, \mathbf{x})$$

where $F(y, \mathbf{x})$ is the joint distribution of y and \mathbf{x} . Without knowing the distribution $F(y, \mathbf{x})$ a priori, we define the empirical error functional as follows

$$\mathcal{R}_n(b|Y, X) = \frac{1}{n} \sum_{i=1}^n Q(b|y_i, \mathbf{x}_i).$$

For example, in the regression case, b is the estimated parameter vector and $\mathcal{R}_n(b|Y, X) = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$.

²See, for example, Yu [1993], Wellner [1981], Tang [2007].

When estimation involves minimizing the in-sample empirical error, we have the extremum estimator [Amemiya, 1985]. In many settings, however, minimizing the in-sample empirical error does not guarantee a reliable model. In regression, for example, often the R^2 is used to measure goodness-of-fit for in-sample data.³ However, an estimate with a high in-sample R^2 may fit out-of-sample data poorly, a feature commonly referred to as **overfitting**: the in-sample estimate is too tailored for the sample data, compromising its out-of-sample performance. As a result, in-sample fit may not be a reliable indicator of the general applicability of the model.

Thus, Vapnik and Chervonenkis [1971a] refer to the **generalization ability** (GA) of a model; a measure of how an extremum estimator performs on out-of-sample data. GA can be measured several different ways. In the case where X and Y are directly observed, GA is a function of the difference between the actual and predicted Y for out-of-sample data. In this paper, GA is measured by the out-of-sample empirical error functional.

Definition 1 (Subsamples, empirical training error and empirical generalization error).

1. Let (y, \mathbf{x}) denote a sample point from $F(y, \mathbf{x})$, where $F(y, \mathbf{x})$ is the joint distribution of (y, \mathbf{x}) . Given a sample (Y, X) , the **training set** $(Y_t, X_t) \in \mathbb{R}^{n_t \times p}$ refers to data used for the estimation of b and the **test set** $(Y_s, X_s) \in \mathbb{R}^{n_s \times p}$ refers to data *not* used for the estimation of b . Let $\tilde{n} = \min\{n_s, n_t\}$. The **effective sample size** for the training set, test set and the total sample, respectively, is n_t/p , n_s/p and n/p .
2. Let Λ denote the space of all models. The **loss function** for a model $b \in \Lambda$ is $Q(b|y_i, \mathbf{x}_i)$, $i = 1, \dots, n$. The **population error functional** for $b \in \Lambda$ is $\mathcal{R}(b|Y, X) = \int Q(b|y, \mathbf{x}) dF(y, \mathbf{x})$. The **empirical error functional** is $\mathcal{R}_n(b|Y, X) = \frac{1}{n} \sum_{i=1}^n Q(b|y_i, \mathbf{x}_i)$.
3. Let $b_{train} \in \Lambda$ denote an extremum estimator. The **empirical training error (eTE)** for b_{train} is $\min_b \mathcal{R}_{n_t}(b|Y_t, X_t) = \mathcal{R}_{n_t}(b_{train}|Y_t, X_t)$, where b_{train} minimizes $\mathcal{R}_{n_t}(b|Y_t, X_t)$. The **empirical generalization error (eGE)** for b_{train} is $\mathcal{R}_{n_s}(b_{train}|Y_s, X_s)$. The population error for b_{train} is $\mathcal{R}(b_{train}|Y, X)$.
4. For K -fold cross-validation, denote the training set and test set in the q th round, respectively, as (X_t^q, Y_t^q) and (X_s^q, Y_s^q) . In each round, the sample size for the training set is $n_t = n(K-1)/K$ and the sample size for the test set is $n_s = n/K$.

The most important assumptions for the analysis in this section of the paper are as follows.

Assumptions

- A1.** In the probability space (Ω, \mathcal{F}, P) , we assume \mathcal{F} -measurability of the loss $Q(b|y, \mathbf{x})$, the population error $\mathcal{R}(b|Y, X)$ and the empirical error $\mathcal{R}_n(b|Y, X)$, for any $b \in \Lambda$ and any sample point (y, \mathbf{x}) . All loss distributions have a closed-form, first-order moment.

³For regression, $R^2 = 1 - \mathcal{R}_n(b|Y, X) / (\text{TSS}/n)$ where $\mathcal{R}_n(b|Y, X) = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$ and $\text{TSS} = \sum_{i=1}^n (y_i - \bar{y})^2$.

- A2.** The sample (Y, X) is independently distributed and randomly chosen from the population. In cases with multiple random samples, both the training set and the test set are randomly sampled from the population. In cases with a single random sample, both the training set and the test set are randomly partitioned from the sample.
- A3.** For any sample, the extremum estimator $b_{train} \in \Lambda$ exists. The in-sample error for b_{train} converges in probability to the minimal population error as $n \rightarrow \infty$.

A few comments are in order for assumptions A1–A3. The loss distribution assumption A1 is merely to simplify the analysis. The existence and convergence assumption A3 is standard (see, for example, Newey and McFadden [1994]). The independence assumption in A2 is not essential because GA analysis is valid for both i.i.d. and non-i.i.d. data. While the original research in Vapnik and Chervonenkis [1974a,b] imposes the i.i.d. restriction on GA, subsequent work has generalized their results to cases where the data are dependent or not identically distributed.⁴ Others have shown that if heterogeneity is due to an observed random variable, the variable may be added to the model to control for the heterogeneity while if the heterogeneity is related to a latent variable, various approaches—such as the hidden Markov model, mixture modelling or factor modelling—are available for heterogeneity control.⁵ Either way, GA analysis is valid owing to the controls for heterogeneity. In this paper, due to the different measure-theory setting for dependent data, we focus on the independent case as a first step. In a companion paper [Xu et al., 2016], we specify the time series mixing type and the types of heterogeneity across individuals to generalize the results in this paper to time series and panel data. Lastly, given A1–A3, both the eGE and eTE converge to the population error:

$$\lim_{\tilde{n} \rightarrow \infty} \mathcal{R}_{n_t}(b_{train}|Y_t, X_t) = \lim_{\tilde{n} \rightarrow \infty} \mathcal{R}_{n_s}(b_{train}|Y_s, X_s) = \mathcal{R}(b_{train}|Y, X)$$

Typically two methods are implemented to compute the eGE of an estimate: validation and cross-validation. For validation when only one sample is available, the sample is randomly partitioned into a training set and a test set; if multiple samples are available, some are chosen as test sets and others as training sets. Either way, we use training set(s) for estimation and test set(s) to compute the eGE for the estimated model, yielding the validated eGE.

K -fold cross-validation may be thought of as ‘averaged multiple-round validation’. For cross-validation, the full sample is randomly partitioned into K subsamples or folds.⁶ One fold is chosen to be the test set and the remaining $K - 1$ folds comprise the training set. Following extremum estimation on the training set, the fitted model is applied to the test set to compute the eGE. The process is repeated K times, with each of the K folds getting the chance to play the role of the test set while the remaining $K - 1$ folds are used as the training set. In this way, we obtain K different estimates of the eGE for the fitted model. The K estimates of the eGE are averaged, yielding the

⁴See, for example, Yu [1994], Cesa-Bianchi et al. [2004], McDonald et al. [2011], Smale and Zhou [2009], Mohri and Rostamizadeh [2009], Kakade and Tewari [2009].

⁵See Michalski and Yashin [1986], Skrondal and Rabe-Hesketh [2004], Wang and Feng [2005], Yu and Joachims [2009], Pearl [2015].

⁶Typically, $K = 5, 10, 20, 40$ or N .

cross-validated eGE.

Cross-validation uses each data point in both the training and test sets. Cross-validation also reduces resampling error by running the validation K times over different training and test sets. Intuitively this suggests that cross-validation is more robust to resampling error and should perform at least as well as validation. In Section 3, we study the generalization ability of penalized extremum estimators in both the validation and cross-validation cases and discuss the difference between them in more detail.

2.2 The upper bound for the empirical generalization error

The traditional approach to model selection in econometrics is to use the AIC, BIC or HQIC, which involves minimizing the eTE and applying a penalty term to choose among alternative models. Based on a broadly similar approach, Vapnik and Chervonenkis [1971a,b] consider model selection from the perspective of generalization ability. Vapnik and Chervonenkis [1971a,b] posit there are essentially two reasons why a model estimated on one sample may have a weak generalization ability on another: the two samples may have different sampling errors, or the complexity of the model estimated from the original sample may have been chosen inappropriately.

To improve the generalization ability of a model, Vapnik and Chervonenkis [1971a,b] propose minimizing the upper bound of the population error of the estimate as opposed to minimizing the eTE. The balance between in-sample fit and out-of-sample fit is formulated by Vapnik and Chervonenkis [1974b] using the Glivenko-Cantelli theorem and Donsker's theorem for empirical processes. Specifically, the relation between $\mathcal{R}_n(b|Y, X)$ and $\mathcal{R}(b|Y, X)$ is summarized by the so-called **VC inequality** [Vapnik and Chervonenkis, 1974b] as follows.

Lemma 1 (The upper bound of the population error (the VC inequality)). *Under A1 to A3, the following inequality holds with probability $1 - \eta$, $\forall b \in \Lambda$, and $\forall n \in \mathbb{N}^+$,*

$$\mathcal{R}(b|Y, X) \leq \mathcal{R}_n(b|Y_t, X_t) + \frac{\sqrt{\varepsilon}}{1 - \sqrt{\varepsilon}} \mathcal{R}_n(b|Y_t, X_t) \quad (2)$$

where $\mathcal{R}(b|Y, X)$ is the population error, $\mathcal{R}_n(b|Y_t, X_t)$ is the training error from the model b , $\varepsilon = (1/n_t)[h \ln(n_t/h) + h - \ln(\eta)]$, and h is the **VC dimension**.

A few comments are in order for the VC inequality, eq. (2).

1. As shown in Figure 1, the RHS of eq. (2) establishes an upper bound for the population error based on the eTE and the VC dimension h . When the effective sample size for the training set (n_t/h) is very large, ε is very small, the second term on the RHS of (2) becomes small, and the eTE is close to the population error. In this case the extremum estimator has a good GA. However, if the effective sample size n_t/h is small (i.e., the model is very complicated), the second term on the RHS of (2) becomes larger. In such situations a small eTE does not guarantee a good GA, and overfitting becomes more likely.
2. The VC dimension h is a more general measure of model complexity than the number of parameters, p , which does not readily extend to nonlinear or non-nested models. While h

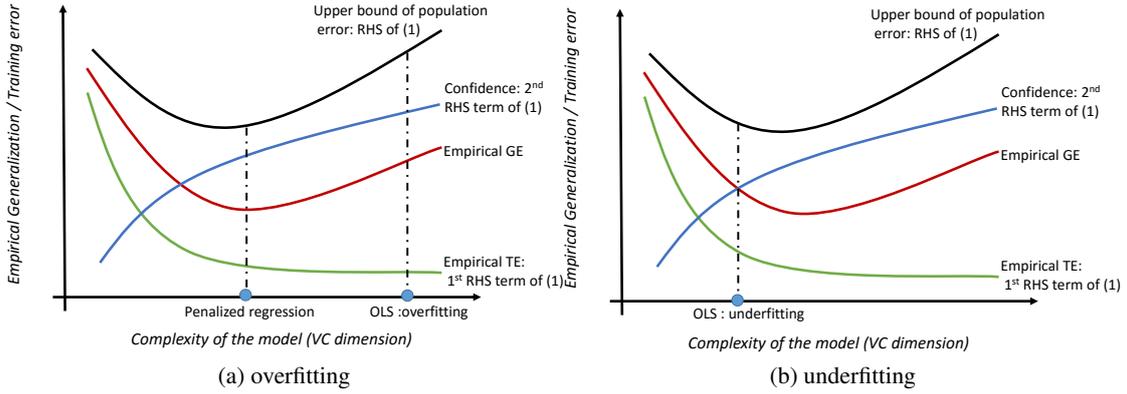


Figure 1: The VC inequality and eGE

reduces to p directly for generalized linear models, h can also be used to partially order the complexity of nonlinear or non-nested models by summarizing their geometric complexity.⁷ As a result, eq. (2) can be implemented as a tool for both nonlinear and non-nested model selection.

- Eq. (2) can be generalized to non-i.i.d. cases. While the VC inequality focuses on the relation between the population error and the eTE in the i.i.d. case, McDonald et al. [2011] generalizes the VC inequality for α - and β -mixing stationary time series while Smale and Zhou [2009] generalizes the VC inequality for panel data. Moreover, Michalski and Yashin [1986], Skrondal and Rabe-Hesketh [2004], Wang and Feng [2005], Yu and Joachims [2009], Pearl [2015] show that heterogeneity can be controlled by implementing the latent variable model or by adding the variable causing heterogeneity into the model, implying eq. (2) is valid.

Based on the VC inequality, Vapnik and Chervonenkis [1971a] propose that minimizing the RHS of (2), the upper bound of the population error, reduces overfitting and improves the GA of the extremum estimator. However, this may be hard to implement because it can be difficult to calculate the VC dimension for anything other than linear models. In practice, statisticians have implemented GA analysis by minimizing the eGE using validation or cross-validation. For example, cross-validation is used to implement many penalty methods, such as the lasso-type estimators, ridge regression or bridge estimators. Clearly, however, the eGE and the population error are not the same thing. Thus, the properties of the minimum eGE, such as its variance, consistency and convergence rate are of particular interest in the present context. By adapting and modifying eq. (2), we propose the following inequalities that analyze the relation between the eGE and the eTE in finite samples.

Theorem 1 (The upper bound of the finite-sample eGE for the extremum estimator). *Under AI to*

⁷In empirical processes, several other geometric complexity measures are connected to or derived from the VC dimension, such as the minimum description length (MDL) score, the Rademacher dimension (or complexity), Pollard's pseudo-dimension and the Natarajan dimension. Most of these measures, like the VC dimension, are derived and generalized from the Glivenko-Cantelli class of empirical processes.

A3, the following upper bound for the eGE holds with probability at least $\varpi(1 - 1/n_t)$, $\forall \varpi \in (0, 1)$.

$$\mathcal{R}_{n_s}(b_{train}|Y_s, X_s) \leq \frac{\mathcal{R}_{n_t}(b_{train}|Y_t, X_t)}{(1 - \sqrt{\varepsilon})} + \zeta, \quad (3)$$

where $\mathcal{R}_{n_s}(b_{train}|Y_s, X_s)$ is the eGE and $\mathcal{R}_{n_t}(b_{train}|Y_t, X_t)$ the eTE for the extremum estimator b_{train} , ε is defined in Lemma 1,

$$\zeta = \begin{cases} \sqrt[2]{\tau} (\mathbb{E}[Q(b_{train}|Y_s, X_s)]) / (\sqrt[2]{1 - \varpi} \cdot n_s^{1-1/\nu}) & \text{if } \nu \in (1, 2] \\ \frac{B}{n_s} \ln \sqrt{2/(1 - \varpi)} & \text{if } Q(\cdot) \in (0, B] \text{ and } B \text{ is bounded} \\ \text{var}[Q(b_{train}|y, \mathbf{x})] / (n(1 - \varpi)) & \text{if } \nu \in (2, \infty) \end{cases}$$

and

$$\tau \geq \sup \frac{[\int (Q(b|y, \mathbf{x}))^\nu dF(y, \mathbf{x})]^{1/\nu}}{\int Q(b|y, \mathbf{x}) dF(y, \mathbf{x})}.$$

A few comments follow from Theorem 1.

- (Upper bound of the finite-sample GA) eq. (3) establishes the upper bound of the eGE from any out-of-sample data of size n_s based on the eTE from any in-sample data of size n_t . Unlike the classical bound in Lemma 1, which captures the relation between the population error and the eTE, eq. (3) establishes inequalities to quantify the upper bound of the finite sample eGE. Usually, we need to use validation or cross-validation to measure the eGE of a model with new data. However, because the RHS of eq. (3) is directly computable it may be used as a measure of finite-sample eGE, avoiding the need for validation.
- (The eGE-eTE trade-off in model selection) eq. (3) also characterizes the trade-off between eGE and eTE for model selection in both the finite sample and asymptotic cases. In Figure 2b, the population eGE, population eTE and population error are expected to be identical in asymptotic case. Hence, minimizing eTE can directly lead to the true DGP in the population. In contrast, as illustrated in Figure 2a, in finite samples, an overcomplicated model with low n_t/h would have a small eTE for the data whereas eq. (3) show that the upper bound of the eGE on new data will be large. Hence, the overcomplicated model will *overfit* the in-sample data and typically have a poor GA. In contrast, an oversimplified model with high n_t/h , typically cannot adequately recover the DGP and the upper bound of the eGE will also be large. As a result, the oversimplified model will *underfit*, fitting both the in-sample and out-of-sample data poorly. Thus, the complexity of a model introduces a trade-off between the eTE and eGE in model selection.
- (GA and distribution tails) eq. (3) also shows how the tail of the error distribution affects the upper bound of the eGE. If the loss distribution $Q(\cdot)$ is bounded or light-tailed, the second term of eq. (3), ζ , is mathematically simple and converges to zero at the rate $1/n_s$. If the loss function is heavy-tailed and \mathcal{F} -measurable, ν , the highest order of the population moment

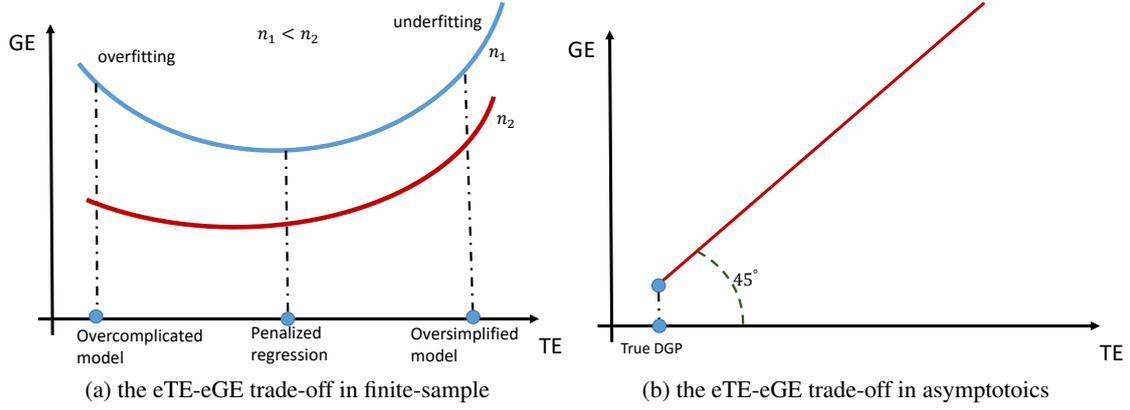


Figure 2: Schematic diagram of the trade-off between eGE and eTE

that is a closed-form for the loss distribution,⁸ can be used to measure the heaviness of the loss distribution tail, a smaller ν implying a heavier tail. In the case of a heavy tail, the second term of eq. (3), ζ , becomes mathematically complicated and its convergence rate decreases to $1/n_s^{1-1/\nu}$. Hence, eq. (3) shows that the heavier the tail of the loss distribution, the higher the upper bound of the eGE and the harder it is to control GA in finite samples. In the extreme case with $\nu = 1$, there is no way to adapt eq. (3).

Essentially, validation randomly partitions the data into a training set and a test set, yielding an estimate on the training set that is used to compute the eGE of the test set. Eq. (3) measures the upper bound the eGE on the test set from the model estimated on the training set with a given eTE and h . In other words, eq. (3) directly measure GA using validation. Furthermore, a similar bound to eq. (3) can be established for K -fold cross-validation.

Theorem 2 (The upper bound of the finite-sample eGE for the extremum estimator under cross-validation). *Under A1 to A3, the following upper bound for the eGE holds with probability at least $\varpi(1 - 1/K)$, $\forall \varpi \in (0, 1)$.*

$$\frac{1}{K} \sum_{j=1}^K \mathcal{R}_{n_s}(b_{train}|Y_s^j, X_s^j) \leq \frac{\frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_t}(b_{train}|Y_t^q, X_t^q)}{(1 - \sqrt{\varepsilon})} + \zeta_{cv}, \quad (4)$$

where $\mathcal{R}_{n_s}(b_{train}|Y_s^j, X_s^j)$ is the eGE of b_{train} in j th round of validations, $\mathcal{R}_{n_t}(b_{train}|Y_t^q, X_t^q)$ is the eTE of b_{train} in q th round of validations, and

$$\zeta_{cv} = \begin{cases} \sqrt[2]{\tau} \mathcal{R}(b_{train}|Y, X) / (\sqrt[2]{1 - \varpi} \cdot n_s^{1-1/\nu}) & \text{if } \nu \in (1, 2] \\ B \ln \sqrt{2/(1 - \varpi)} / n_s & \text{if } Q(\cdot) \in (0, B] \text{ and } B \text{ is bounded} \\ \text{var}[Q(b_{train}|y, \mathbf{x})] / (n_s^2(1 - \varpi)) & \text{if } \nu \in (2, \infty) \end{cases}$$

⁸It is closed-form because owing to A1, which guarantees closed-form, first-order moments for all loss distribution in the paper.

The errors generated by cross-validation are affected both by sampling randomness from the population and by sub-sampling randomness that arises from partitioning the sample into folds. Thus, the errors from cross-validation are potentially more volatile than the usual errors from estimation. Theorem 2 provides an upper bound for the average eGE under cross-validation, which offers a way to characterize the effect of sub-sampling randomness and suggests a method to approximate the GA from cross-validation. The following comments summarize the implications of eq. (4).

1. (*The upper bound of the eGE*) Similar to eq. (3), eq. (4) serves as the upper bound of the averaged eGE generated by cross-validation. Both equations show the eTE-eGE trade-off and reveal the effect of a heavy tail on GA.
2. (*Tuning the cross-validation hyperparameter K*) Eq. (4) characterizes how the hyperparameter K affects the averaged eGE from cross-validation (also called the cross-validation error in the literature). As explained previously, by randomly partitioning the sample into several folds, cross-validation results in sub-sampling randomness. Sub-sampling randomness may produce a different averaged eGE each time cross-validation is performed. Moreover, the partition will be completely different and hence the eGE will not be directly comparable when K changes. Unlike OLS, therefore, we cannot directly analyze the averaged eGE from cross-validation. Due to the influence of sub-sampling randomness, the bias-variance trade-off for cross-validation may only be established using its upper bound, as is illustrated in Figure 3.
 - (a) (*Large bias, small variance*) When K is small, n_t is smaller in each round of in-sample estimation. Hence, as shown in Figure 3a, the eTE in each round, $\mathcal{R}_{n_t}(b_{train}|Y_t^q, X_t^q)/(1 - \sqrt{\epsilon})$, is more biased from the population error. As shown in Figure 3b, the K -round averaged eTE, $\frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_t}(b_{train}|Y_t^q, X_t^q)/(1 - \sqrt{\epsilon})$, is more biased away from the true population error as K gets smaller. As a result, the RHS of eq. (4) suffers more from finite-sample bias. However, since small K implies that more data is used for eGE calculation in each round (n_s is not very small), in each round the eGE on the test set should not be very volatile. Thus, the K -round averaged eGE for cross-validation is not very volatile, which is shown by the fact that ζ_{cv} is not very large in eq. (4).
 - (b) (*Small bias, large variance*) When K is large, n_s is small and the test set in each round is small. Hence, with large K , the eGE in each round may be hard to bound from above, which implies that the averaged eGE from K rounds is more volatile. As a result, ζ_{cv} tends to be large. However, with large K , the RHS term $\frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_t}(b_{train}|Y_t^q, X_t^q)/(1 - \sqrt{\epsilon})$ tends to be closer to the true population error, so the averaged eGE suffers less from bias.

As shown in Figure 3b, the averaged eGE from cross-validation follows a typical bias-variance trade-off by value of K . If K is small, the averaged eGE is computationally cheap and less volatile but more biased away from the population error. As K gets larger, the

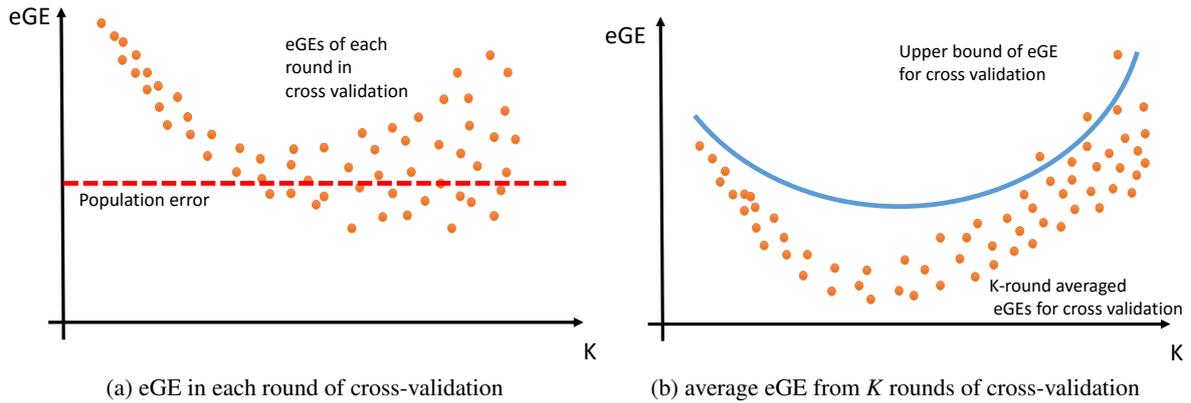


Figure 3: Representation of the bias-variance trade-off for cross-validation eGE

averaged eGE becomes computationally expensive and more volatile but less biased away from the population error. This result exactly matches the Kohavi et al. [1995] simulation study. More specifically, by tuning K to the lowest upper bound, we can find the K that maximizes the GA from cross-validation.

Theorems 1 and 2 establish the finite-sample and asymptotic properties of GA analysis for any extremum estimator. In finite-sample analysis, the results capture the trade-off between eGE and eTE, which can be used to measure the GA of an econometric model. In asymptotic analysis, eGE minimization is consistent. As a result, GA can be implemented as a criterion for model selection, and directly connects to the theoretical properties for model selection methods such as penalized extremum estimation, the various information criteria and maximum a posteriori (MAP) estimation. Minimizing eGE works especially well for penalized regression. As shown in Algorithm 1, penalized regression estimation returns a b_λ for each λ . Each value of λ generates a different model and a different eGE. Intuitively, Theorems 1 and 2 guarantee that the model with the minimum eGE from $\{b_\lambda\}$ has the best empirical generalization ability. In the next section, we study the finite-sample and asymptotic properties of eGE for all penalized regressions.

3 Finite-sample and asymptotic properties of eGE for penalized regression

Using the classical concentration inequalities in Section 2, we established the upper bound for the finite-sample eGE of the extremum estimator given any random sample of any size. We also revealed the trade-off between eTE and eGE for model selection and derived the properties of eGE under validation and cross-validation. In this section, we apply the framework and results from Section 2 to penalized regression.

3.1 Definition of penalized regression

Firstly, we formally define penalized regression and its two most popular variants: ridge regression (\mathcal{L}_2 -penalized regression) and the lasso (\mathcal{L}_1 -penalized regression).

Definition 2 (Penalized regression, \mathcal{L}_2 -eGE and \mathcal{L}_2 -eTE).

1. (*General form*) The general form of the objective function for penalized regression is as follows

$$\min_{b_\lambda} \frac{1}{n} (\|Y - Xb_\lambda\|_2)^2 + \lambda \text{Penalty}(\|b_\lambda\|_\gamma). \quad (5)$$

where $\text{Penalty}(\|\cdot\|_\gamma)$ stands for the penalty term, which is a function of the \mathcal{L}_γ norm of the b_λ .

2. (*b_λ and b_{pen}*) We denote b_λ to be the solution of eq. (5) given the value of the penalty parameter λ while b_{pen} is defined to be the model with the minimum eGE among all alternative $\{b_\lambda\}$, as in Algorithm 1 in Section 1.
3. (*Lasso and ridge*) The objective functions for lasso (\mathcal{L}_1 penalty) and ridge regression (\mathcal{L}_2 penalty), respectively, are

$$\min_{b_\lambda} \frac{1}{n} (\|Y - Xb_\lambda\|_2)^2 + \lambda \|b_\lambda\|_1, \quad (6)$$

and

$$\min_{b_\lambda} \frac{1}{n} (\|Y - Xb_\lambda\|_2)^2 + \lambda \|b_\lambda\|_2. \quad (7)$$

4. (*\mathcal{L}_2 error for regression*) the eTE and eGE for regression are defined in \mathcal{L}_2 form respectively as follows:

$$\begin{aligned} \mathcal{R}_{n_t}(b_{train}|Y_t, X_t) &= \frac{1}{n_t} \|Y_t - X_t b_{train}\|_2^2 \\ \mathcal{R}_{n_s}(b_{train}|Y_s, X_s) &= \frac{1}{n_s} \|Y_s - X_s b_{train}\|_2^2 \end{aligned}$$

The idea behind penalized regression is illustrated in Figure 4 where b_λ refers to the penalized regression estimates for some λ and b_{OLS} refers to the OLS estimates. As shown in Figure 4a, different \mathcal{L}_γ norms correspond to different boundaries for the estimation feasible set. For the \mathcal{L}_1 penalized regression (lasso), the feasible set is a diamond since each coefficient is equally penalized by the \mathcal{L}_1 norm. The feasible area shrinks under a $\mathcal{L}_{0.5}$ penalty. Hence, as shown in Figure 4a, given the same λ , the smaller is γ , the more likely b_λ is to be a corner solution. Hence, given the same λ , under the $\mathcal{L}_{0.5}$ penalty variables are more likely to be dropped than with the \mathcal{L}_1 or \mathcal{L}_2 penalty.⁹ In special cases when $\gamma = 0$ and λ is fixed at $2(\ln n_t)$, the \mathcal{L}_0 penalized regression is identical to the Akaike (Bayesian) information criterion.

⁹For $0 < \gamma < 1$, the penalized regression may be a non-convex programming problem. While general algorithms have not been found for non-convex optimization, Strongin and Sergeyev [2013], Yan and Ma [2001] and Noor [2008] have developed functioning algorithms. For $\gamma = 0$, the penalized regression becomes a discrete programming problem, which can be solved by Dantzig-type methods (see Candes and Tao [2007]).

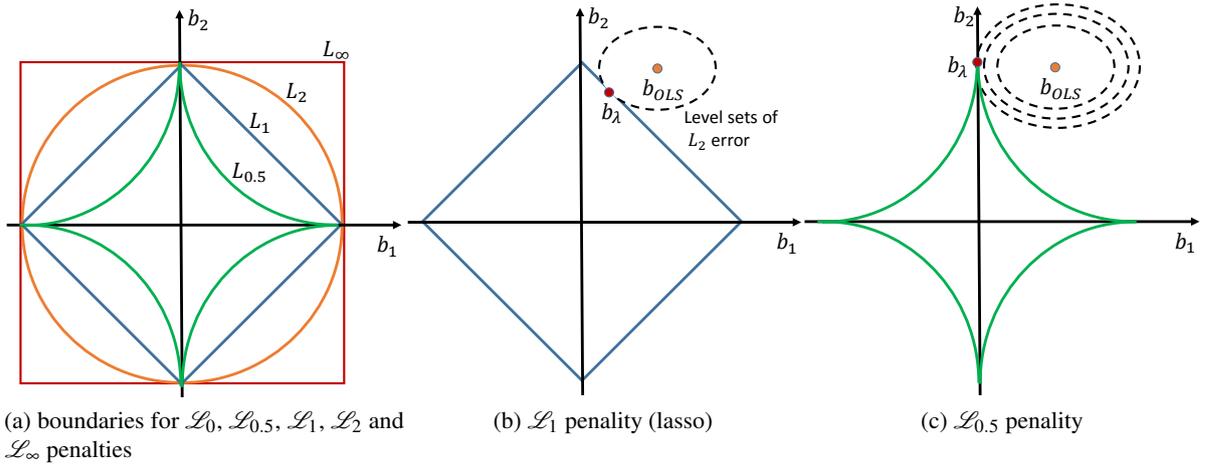


Figure 4: Illustration of various penalized regressions

The last important comment is that penalized regression primarily focuses on overfitting. By contrast, OLS minimizes the eTE without any penalty, typically causing a large eGE (as shown in Figure 1a). There is also the possibility that OLS predicts the data poorly, causing both the eTE and eGE to be large. The latter refers to underfitting and is shown in Figure 1b. We are more capable of dealing with overfitting than underfitting despite the fact that it is possible to quantify GA or the eGE.¹⁰ Typically overfitting in OLS is caused by including too many variables, which we can resolve by reducing p . However, underfitting in OLS is typically due to a lack of data (variables) and the only remedy is to collect additional relevant variables.

3.2 Schematics and assumptions for eGE minimization with penalized regression

As shown in Section 2, eGE minimization improves finite-sample GA, implying the estimator has a lower eGE on out-of-sample data. In this section, we implement the schematics of eGE minimization on penalized regression. We demonstrate: (1) specific error bounds for any penalized regression, (2) a general \mathcal{L}_2 consistency property for penalized regression estimates, (3) that the upper bound for the \mathcal{L}_2 difference between b_{pen} and b_{OLS} is a function of the eGE, the tail property of the loss distribution and sample exogeneity.

The classic route to derive asymptotic or finite-sample properties for regression is through analyzing the properties of the estimate in the space of the eTE. In contrast, to study how penalized regression improves GA or eGE and balances the in-sample and out-of-sample fit, we reformulate the asymptotic and finite-sample problems in the space of the eGE. We show that, under the framework of eGE minimization, a number of finite-sample properties of penalized regression can be explained by eGE or the finite-sample GA.

In asymptotic analysis, consistency is typically considered to be one of the most fundamental properties. To ensure that eGE minimization is a reliable estimation approach, we prove that the penalized regression, which is a specific form of the eGE minimizer, converges to the true DGP as

¹⁰See eq. (12) and (13).

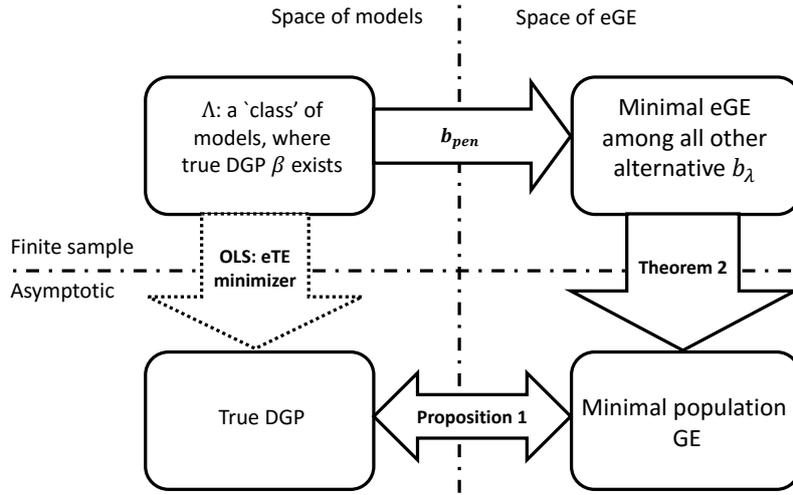


Figure 5: Outline of proof strategy

$n \rightarrow \infty$. Essentially, we show that penalized regression bijectively maps b_{pen} to the minimal eGE among $\{b_\lambda\}$ on the test set. To bridge between the finite sample and asymptotic results we need to show that if

- the true DGP β is bijectively assigned to the minimal eGE in population, and
- $\min_{b \in b_\lambda} \frac{1}{n_s} \sum_{i=1}^{n_s} \|Y_s - X_s b\|_2^2 \rightarrow \min_b \int \|y - \mathbf{x}^T b\|_2^2 dF(y, \mathbf{x})$,

then b_{pen} is consistent in probability or \mathcal{L}_2 , or

$$b_{pen} = \operatorname{argmin}\{\text{eGEs of } \{b_\lambda\}\} \xrightarrow{\mathbf{P} \text{ or } \mathcal{L}_2} \operatorname{argmin}_b \int \|y_s - \mathbf{x}_s^T b\|_2^2 dF(y, \mathbf{x}) = \beta.$$

At the outset, we stress that each variable in (Y, X) must be standardized before implementing penalized regression. Without standardization, as shown by [Tibshirani, 1996], the penalized regression may be influenced by the magnitude (units) of the variables. After standardization, of course, X and Y are unit- and scale-free.

To ensure the consistency of penalized regression, we require the following three additional assumptions.

Further assumptions

A4. The true DGP is $Y = X\beta + u$.

A5. $\mathbb{E}(u^T X) = \mathbf{0}$.

A6. No perfect collinearity in X .

The assumptions **A4** to **A6** restrict the true DGP β to be identifiable. Otherwise, there might exist another model that is not statistically different from the true DGP. The assumptions are quite standard for linear regression.

3.3 Necessary propositions for final results

Under assumptions **A1** to **A6**, we show that the true DGP is the most generalizable model, yielding Proposition 3.1.

Proposition 1 (Identification of β in the space of eGE). *Under assumptions **A1** to **A6**, the true DGP, $Y = X\beta + u$, is the one and only one offering the minimal eGE as $\tilde{n} \rightarrow \infty$.*

Proposition 1 states that there is a bijective mapping between β and the global minimum eGE in the population. If **A5** or **A6** are violated, there may exist variables in the sample that render the true DGP not to be the model with minimum eGE in population. As shown in Algorithm 1, penalized regression picks the model with the minimum eGE in $\{b_\lambda\}$ to be b_{pen} . As a result, we also need to prove that, when the sample size is large enough, the true DGP is included in $\{b_\lambda\}$, the list of models from which validation or cross-validation selects. This is shown in Proposition 3.2.

Proposition 2 (Existence of \mathcal{L}_2 consistency). *Under assumptions **A1** to **A6** and Proposition 1, there exists at least one $\tilde{\lambda}$ such that $\lim_{\tilde{n} \rightarrow \infty} \|b_{\tilde{\lambda}} - \beta\|_2 = 0$.*

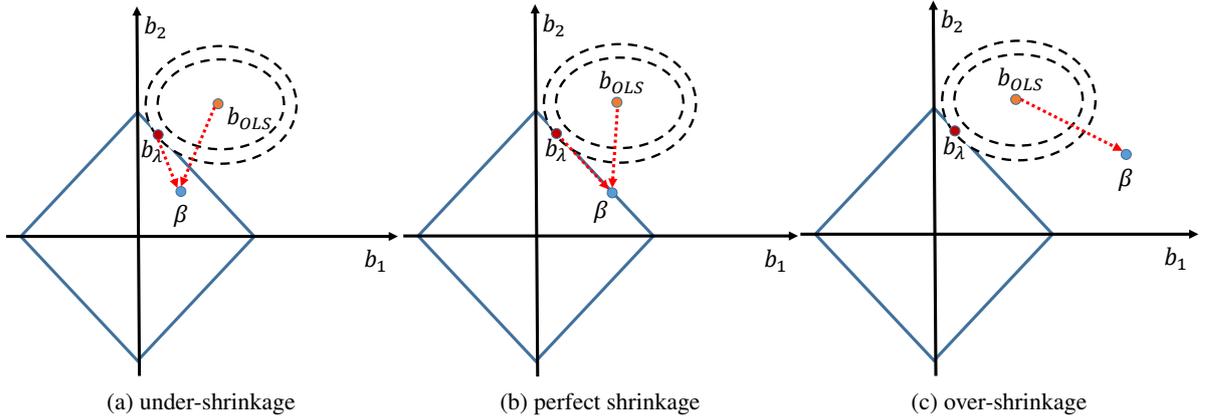


Figure 6: Various types of shrinkage under an \mathcal{L}_1 penalty

Since the penalized regression can be sketched as a constrained minimization of the empirical error, we can illustrate Proposition 1 and 2 in Figure 6 using lasso as the example of a penalized regression. In Figure 6, the parallelogram-shaped feasible sets are determined by the \mathcal{L}_γ penalty, b_λ refers to the solution of eq. (5), β refers to the true DGP, and b_{OLS} refers to the OLS estimates. Different values for λ imply different areas for the feasible set of the constrained minimization; the area of the feasible set gets smaller as the value of λ gets larger. Hence, one of three cases may occur: (i) as shown in Figure 6a, for a small value of λ , β lies in the feasible set (under-shrinkage) and offers the minimum eTE in the population; (ii) as shown in Figure 6b, for the oracle λ , β is located precisely on the boundary of the feasible set (perfect-shrinkage) and still offers the minimum eTE in the population; (iii) as shown in Figure 6c, for a large value of λ , β lies outside the feasible set (over-shrinkage). In cases (i) and (ii), the constraints become inactive as $\tilde{n} \rightarrow \infty$, so $\lim_{\tilde{n} \rightarrow \infty} b_\lambda = \lim_{\tilde{n} \rightarrow \infty} b_{OLS} = \beta$. However, in case (iii), $\lim_{\tilde{n} \rightarrow \infty} b_\lambda \neq \beta$. Therefore, tuning the penalty parameter λ is critical for the theoretical properties of penalized regression estimates.

3.4 Main results for penalized regression estimates

As shown above, intuitively the penalized regression estimate is expected to be consistent in some norm or measure as long as we can be sure that for a specific λ , β lies in the feasible set and offers the minimum eTE in population. However, in practice we may not know a priori which λ causes over-shrinkage and which does not, especially when the number of variables, p , is not fixed. As a result, we need a method such as cross-validation or validation to tune the value of λ . Thus, as a direct application of eGE minimization in Section 2, we use GA/eGE analysis to show that eGE minimization guarantees the model selected by penalized regression, b_{pen} , asymptotically converges in \mathcal{L}_2 to the true DGP.

In the following section we analyze the finite-sample and asymptotic properties of the penalized regression estimate in two scenarios: $n \geq p$ and $n < p$. In the case where $n \geq p$, OLS is feasible, so we take the OLS estimate for the unpenalized regression estimate. However, when $n < p$, OLS is not feasible, and we use forward stagewise regression (FSR) for the unpenalized regression estimate. Hereafter, b_{OLS} is the OLS estimate from the training set.

3.4.1 Case: $n \geq p$

Firstly, by specifying eq. (3) and (4) in the context of regression, we can establish the upper bound of the eGE, as shown in Lemma 2.

Lemma 2 (Upper bound of the eGE for the OLS estimate). *Under A1 to A6, if we assume the OLS error term $u \sim \text{Gaussian}(0, \text{var}(u))$,*

1. (**Validation**) *The following bound for the eGE holds with probability at least $\bar{\omega}(1 - 1/n_t)$ for b_{OLS} , $\forall \bar{\omega} \in (0, 1)$.*

$$\frac{1}{n_s} (\|e_s\|_2)^2 \leq \frac{\frac{1}{n_t} (\|e_t\|_2)^2}{(1 - \sqrt{\varepsilon})} + \frac{2(\text{var}(u))^2}{\sqrt{1 - \bar{\omega}} \cdot n_s}, \quad (8)$$

where e_s is the OLS eGE, e_t is the OLS eTE, and ε is defined in Lemma 1.

2. (**K-fold cross-validation**) *The following bound for the eGE holds with probability at least $\bar{\omega}(1 - 1/K)$ for b_{OLS} , $\forall \bar{\omega} \in (0, 1)$.*

$$\frac{1}{K} \sum_{j=1}^K \frac{(\|e_s^j\|_2)^2}{n/K} \leq \frac{n(K-1) \sum_{q=1}^K (\|e_t^q\|_2)^2}{K^2(1 - \sqrt{\varepsilon})} + \frac{2(\text{var}(u))^2}{\sqrt{1 - \bar{\omega}} \cdot (n/K)^2}, \quad (9)$$

where e_s^j is the eGE of OLS estimate in j th round, e_t^q is the eTE of OLS in q th round, while ε and ζ are defined in Lemma 1.

Eq. (8) and (9) show that the higher the variance of u in true DGP, the higher the upper bound of the eGE in validation and cross-validation. Based on eq. (9), the lowest upper bound of the cross-validation eGE is determined by minimizing the expectation of the RHS of eq. (9), yielding the way to find the optimal K as follows.

Corollary 1 (The optimal K for penalized regression). *Under A1 to A6 and eq. (9) from Lemma 2, if we assume the OLS error term $u \sim \text{Gaussian}(0, \text{var}(u))$, the optimal K for cross-validation in regression associated with the minimal expected upper bound of eGE, is defined:*

$$K^* = \underset{K}{\text{argmin}} \frac{\text{var}(u)}{1 - \sqrt{\varepsilon}} + \frac{4(\text{var}(u))^2}{\sqrt{1 - \overline{\omega}}(n/K)^2}$$

Similar to eq. (3), eq. (8) and (9) respectively measure the upper bound of the eGE for the OLS estimate using validation and cross-validation. In standard practice, neither validation nor cross-validation are implemented as part of OLS estimation and hence the eGE of the OLS estimate is rarely computed. Eq. (8) and (9) show that eGE can be computed without having to carry out validation or cross-validation.

In penalized regression, the penalty parameter λ can be tuned by validation or K -fold cross-validation. For $K \geq 2$, we have K different test sets for tuning λ and K different training sets for estimation. Based on eq. (8) and (9), an upper bound for the \mathcal{L}_2 predicted difference between b_{OLS} and b_{pen} can be established under validation and cross-validation, as shown below.

Proposition 3 (\mathcal{L}_2 distance between the penalized and unpenalized predicted values). *Under A1 to A6 and based on Lemma 2, Proposition 1, and Proposition 2,*

1. (**Validation**) *The following bound holds for the validated b_{pen} with probability $\overline{\omega}(1 - 1/n_t)$*

$$\frac{1}{n_s} (\|X_s b_{OLS} - X b_{pen}\|_2)^2 \leq \left(\frac{1}{n_t} \frac{\|e_t\|_2^2}{1 - \sqrt{\varepsilon}} - \frac{1}{n_s} \|e_s\|_2^2 \right) + \frac{4}{n_s} \|e_s^T X_s\|_\infty \|b_{OLS}\|_1 + \zeta \quad (10)$$

where b_{OLS} is the OLS estimate based on the training set and ζ is defined in Lemma 2.

2. (**K -fold cross validation**) *The following bound holds for the K -fold cross-validated b_{pen} with probability $\overline{\omega}(1 - 1/n_t)$*

$$\begin{aligned} \frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} (\|X_s^q b_{OLS}^q - X_s^q b_{pen}^q\|_2)^2 &\leq \left| \frac{1}{n_t} \frac{\frac{1}{K} \sum_{q=1}^K \|e_t^q\|_2^2}{1 - \sqrt{\varepsilon}} - \frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} \|e_s^q\|_2^2 \right| \\ &+ \frac{1}{K} \sum_{q=1}^K \frac{4}{n_s} \left\| (e_s^q)^T X_s^q \right\|_\infty \|b_{OLS}^q\|_1 + \zeta_{cv}. \end{aligned} \quad (11)$$

where e_t^q is the eTE of the OLS estimate in the q th round of cross-validation, e_s^q is the eGE of the OLS estimate in the q round of cross-validation and b_{OLS}^q is the OLS estimate in the q th round of cross-validation.

We can now derive the upper bound of $\|b_{OLS} - b_{pen}\|_2$, as shown in Theorem 3.

Theorem 3 (\mathcal{L}_2 distance between the penalized and unpenalized estimates). *Under A1 to A6 and based on Propositions 1, 2, and 3,*

1. (**Validation**) The following bound holds with probability $\bar{\omega}(1 - 1/n_t)$:

$$\|b_{OLS} - b_{pen}\|_2 \leq \sqrt{\left| \frac{1}{\rho n_t} \frac{\|e_t\|_2^2}{(1 - \sqrt{\varepsilon})} - \frac{1}{\rho n_s} \|e_s\|_2^2 \right|} + \sqrt{\frac{4}{\rho n_s} \|e_s^T X_s\|_\infty \|b_{OLS}\|_1} + \left(\frac{\zeta}{\rho}\right)^{\frac{1}{2}} \quad (12)$$

where ρ is the minimal eigenvalue of $X^T X$ and ζ is defined in Lemma 2.

2. (**K-fold cross validation**) The following bound holds with probability $\bar{\omega}(1 - 1/n_t)$:

$$\begin{aligned} \frac{1}{K} \sum_{q=1}^K (\|b_{OLS}^q - b_{pen}^q\|_2)^2 &\leq \left| \frac{1}{K} \sum_{q=1}^K \frac{1}{n_t \bar{\rho}} \frac{\|e_t^q\|_2^2}{1 - \sqrt{\varepsilon}} - \frac{1}{K} \sum_{q=1}^K \frac{1}{n_s \bar{\rho}} \|e_s^q\|_2^2 \right| \\ &\quad + \frac{1}{K} \sum_{q=1}^K \frac{4}{n_s \bar{\rho}} \|(e_s^q)^T X_s^q\|_\infty \|b_{OLS}^q\|_1 + \frac{\zeta}{\bar{\rho}} \end{aligned} \quad (13)$$

where $\bar{\rho}$ is defined as $\min \left\{ \rho_q \mid \rho_q \text{ is the minimal eigenvalue of } (X_s^q)^T X_s^q, \forall q \right\}$.

Some important remarks apply to Theorem 3. The LHS of eq. (12) essentially measures how much the penalized regression estimate differs from the OLS estimate under validation. The RHS of eq. (12) essentially captures the maximum \mathcal{L}_2 difference between b_{OLS} and b_{pen} . As shown in eq. (12), the maximum difference depends on the GA of the true DGP and the GA of the OLS model in several different forms.

- The first term on the RHS of eq. (12) (ignoring $1/\rho$) is the difference between the eGE from OLS and the upper bound of the population error for the OLS estimate, or, equivalently, *how far the GA of the OLS estimate is from its maximum*. The larger the GA of b_{OLS} , the less overfitting OLS generates, the closer the eGE of b_{OLS} is to the upper bound of the population error, and the smaller the first term on the RHS of eq. (12).
- The second term on the RHS of eq. (12) (ignoring $4/\rho$) measures the empirical endogeneity of the error term of the OLS estimate on the test set. On the training set $e_t^T X_s = \mathbf{0}$, but in general $e_s^T X_s \neq 0$ on the test set. Hence, $\frac{1}{n_s} \|e_s^T X_s\|_\infty \|b_{OLS}\|_1$ *measures the GA for the empirical moment condition of the OLS estimate on out-of-sample data*.¹¹ The more generalizable the OLS estimate, the closer $e_s^T X_s$ is to zero on out-of-sample data, and the smaller the second term on the RHS of eq. (12).
- The third term on the RHS of eq. (12) is affected by ζ , which *measures the heaviness of the tail in the loss distribution for the OLS estimate*. Similar to the comments of Theorem 1, the OLS loss distribution affects the GA of the OLS estimate. The heavier the loss tail, the more volatile the eGE on out-of-sample data, and the more difficult it is to bound the eGE for OLS.
- All three RHS terms in eq. (12) are affected by ρ , the minimum eigenvalue of $X_x^T X_x$, which can be thought of as a *measure of the curvature of the objective function for penalized regression*. The larger the minimum eigenvalue, the more convex the objective function. Put another way, it is easier to identify the true DGP β from the alternatives as n get larger.

¹¹Because we standardize the test and training data, the moment condition $\mathbb{E}(e_s) = 0$ holds directly.

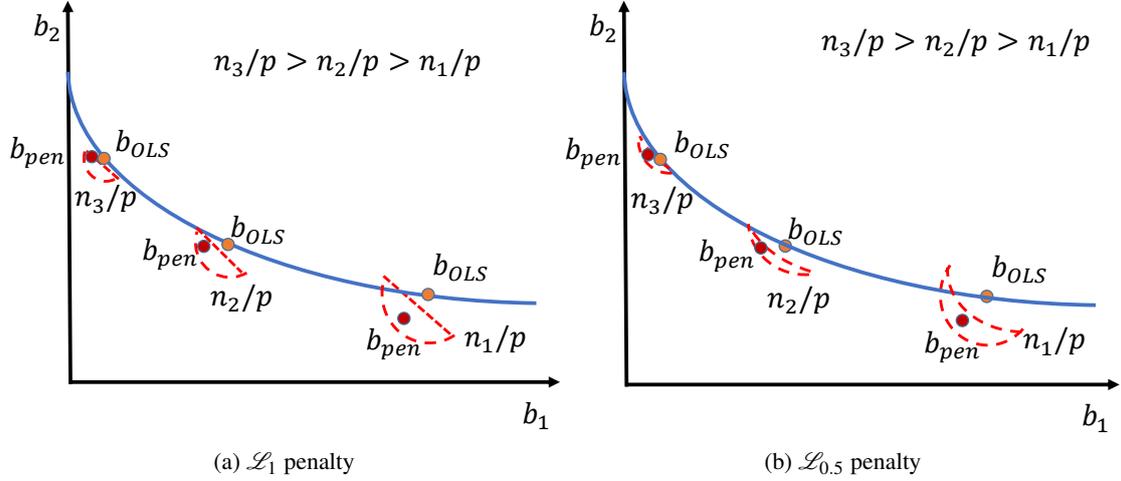


Figure 7: Illustration of the relation between b_{OLS} and b_{pen}

The interpretation of eq. (13) is similar to eq. (12) adjusting for cross-validation. Hence, the first term on the RHS of eq. (13) (ignoring $1/\bar{\rho}$) stands for how far away the *average* GA of OLS estimate is from its maximum in K rounds of validation. The second term on the RHS of eq. (13) (ignoring $4/\bar{\rho}$) indicates *on average* how generalizable the empirical moment condition of the OLS estimate is with out-of-sample data in K rounds of validation. Similarly, $\bar{\zeta}$ indicates *on average* the heaviness in the tail of the loss distribution in K rounds of validation. As a direct result of Theorem 3, the \mathcal{L}_2 consistency for the penalized regression estimate is established as follows.

Corollary 2 (\mathcal{L}_2 consistency for the penalized regression estimate when $n \geq p$). *Under A1 to A6 and Propositions 1, 2, and 3, if we assume the error term $u \sim \text{Gaussian}(0, \text{var}(u))$, then b_{pen} converges in the \mathcal{L}_2 norm to the true DGP if $\lim_{n \rightarrow \infty} p/\tilde{n} = 0$*

Theorem 3 and Corollary 2 are illustrated in Figure 7. Typically, due to the poor GA of the OLS estimate, the penalized regression estimate b_{pen} will not lie on the same convergence path as the OLS estimate. However, Theorem 3 shows that the deviation of b_{pen} from the convergence path is bounded. Furthermore, b_{pen} typically lies within an ε -ball centered on b_{OLS} whose radius is a function of the eGEs of the OLS estimate and the true DGP. Also, as shown in Figures 6a and 6b, b_{pen} always lies within the feasible set parameterized by $\lambda \|b\|_\gamma$. Hence, as shown in Figures 7a and 7b, b_{pen} typically is located in the small area at the intersection of the \mathcal{L}_γ feasible area and ε -ball. Unless the optimal λ from validation or cross-validation is 0, the OLS estimate will never be in the feasible area of the penalized regression estimate, which is why the intersection region is always below the b_{OLS} . As n/p increases, the ε -ball becomes smaller, the penalized regression estimate gets closer to the OLS estimate, and both converge to β .

3.4.2 Case: $n < p$

Typically, to ensure b_{OLS} can identify the true DGP β , we require that $(\|e\|_2)^2$ is strongly convex, or that the minimal eigenvalue of $X^T X$, ρ , is strictly larger than 0. However, if $p > n$, $\rho = 0$ and

the space of $(\|e\|_2)^2$ is flat in some direction. As a result, the $\|b_{OLS}\|_1$ is not closed-form, the true DGP β cannot be identified and Eqs. (12) and (13) are trivial.

To make the results above non-trivial, we need to ensure β is still identifiable even when $p > n$. Put another way, we need to ensure that the strong convexity of the space $(\|e\|_2)^2$ is still reserved for the $p > n$ case and β is still identifiable. This is guaranteed by the sparse eigenvalue condition [Bickel et al., 2009, Meinshausen and Yu, 2009, Zhang, 2010]—see the proof of Proposition 4 (below) in Appendix 1 for the details.

Regression can at most estimate $\max(p, n)$ coefficients. When $p > n$, penalized regression has to drop some variables to make it estimable, implying that $\gamma > 1$ does not apply to the $p > n$ case. Hence, for the $p > n$ case, we focus only on the \mathcal{L}_1 penalized regression, i.e., the lasso. As shown by Efron et al. [2004], Zhang [2010], lasso may be thought of as a forward stagewise regression (FSR) with an \mathcal{L}_1 norm constraint.¹² Hence, lasso regression can be treated as a form of GA/eGE control for FSR when $p > n$. As shown by Zhang [2010], even though FSR is a greedy algorithm that may result in overfitting in finite samples, it is still \mathcal{L}_2 consistent under the sparse eigenvalue condition

Thus, for $p > n$, we set the FSR estimate, b_{FSR} , to be the unpenalized regression estimate, and the \mathcal{L}_1 penalized regression estimate, b_{pen} , to be the penalized regression estimate. In Proposition 4 and Corollary 3, we show that the lasso preserves the properties and interpretations of the $n \geq p$ case by reducing the overfitting inherent in FSR.

Proposition 4 (\mathcal{L}_2 distance between \mathcal{L}_1 penalized and unpenalized FSR estimates). *Under A1 to A6 and based on Lemma 2, Propositions 1, and 2 and the sparse eigenvalue condition,*

1. (**Validation**) *The following bound holds with probability $\varpi(1 - 1/n_t)$*

$$\begin{aligned} \|b_{FSR} - b_{pen}\|_2 &\leq \sqrt{\left| \frac{1}{\rho_{re} n_t} \frac{\|e_t\|_2^2}{(1-\sqrt{\varepsilon})} - \frac{1}{\rho_{re} n_s} \|e_s\|_2^2 \right|} \\ &\quad + \sqrt{\frac{4}{\rho_{re} n_s} \|e_s^T X_s\|_\infty \|b_{FSR}\|_1} + \left(\frac{\varsigma}{\rho_{re}}\right)^{\frac{1}{2}} \end{aligned} \quad (14)$$

where ρ_{re} is the minimum of the sparse eigenvalues of $X^T X$.

2. (**K-fold cross-validation**) *The following bound holds with probability $\varpi(1 - 1/n_t)$*

$$\begin{aligned} \frac{1}{K} \sum_{q=1}^K \|b_{FSR}^q - b_{pen}^q\|_2^2 &\leq \left| \frac{1}{K} \sum_{q=1}^K \frac{1}{n_t \bar{\rho}_{re}} \frac{\|e_t^q\|_2^2}{1-\sqrt{\varepsilon}} - \frac{1}{K} \sum_{q=1}^K \frac{1}{n_s \bar{\rho}_{re}} \|e_s^q\|_2^2 \right| \\ &\quad + \frac{1}{K} \sum_{q=1}^K \frac{4}{n_s \bar{\rho}_{re}} \left\| (e_s^q)^T X_s^q \right\|_\infty \|b_{FSR}\|_1 + \frac{\varsigma}{\bar{\rho}_{re}} \end{aligned} \quad (15)$$

where $\bar{\rho}_{re}$ is defined as $\min \left[\rho_{re}^q \mid \rho_{re}^q \text{ is the minimal restricted eigenvalue of } (X_s^q)^T X_s^q, \forall q \right]$ and b_{FSR}^q is the FSR estimator in q th round of validation.

¹²The method of solving lasso by forward selection is the least angle regression (LARS). For details of LARS and its consistency, see Efron et al. [2004] and Zhang [2010].

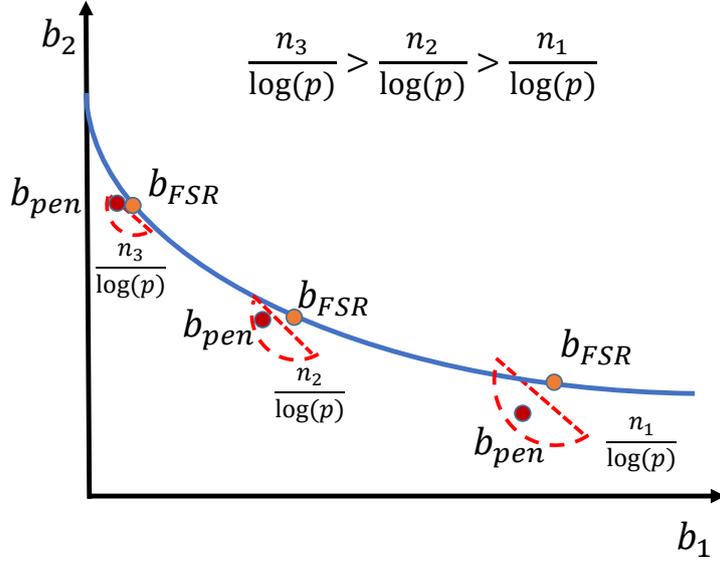


Figure 8: Illustration of the relation between the FSR estimate and the \mathcal{L}_1 penalized FSR estimate

Corollary 3 (\mathcal{L}_2 consistency for the \mathcal{L}_1 -penalized regression estimate when $n < p$). *Under A1 to A6 and based on Lemma 2, Propositions 1, and 2 and the sparse eigenvalue condition, b_{pen} converges in the \mathcal{L}_2 norm asymptotically to the true DGP if $\lim_{n \rightarrow \infty} \log(p)/\tilde{n} = 0$.*

The interpretation of Proposition 4 and Corollary 3 is very similar to the interpretation of Eqs. (12) and (13), which specifies the upper bound of \mathcal{L}_2 difference between b_{FSR} and b_{pen} as a function of GA of the FSR estimate and the population error. Hence, the interpretation for Proposition 4 and Corollary 3 can also be illustrated in Figure 8.

4 Simulation Study

In sections 2 and 3, we use eTE to measure in-sample fit and eGE to measure out-of-sample fit. However, to measure the GA and degree of overfitting, we need to compare the eTE and eGE to the total sum of square for the training set and test set, respectively. To summarize the in-sample and out-of-sample goodness of fit, we propose the generalized R^2 :

$$GR^2 = \left(1 - \frac{\mathcal{R}_{n_s}(b_{train}|Y_s, X_s)}{\text{TSS}(Y_s)}\right) \times \left(1 - \frac{\mathcal{R}_{n_t}(b_{train}|Y_t, X_t)}{\text{TSS}(Y_t)}\right) = R_s^2 \times R_t^2 \quad (16)$$

where R_s^2 is the R^2 for the test set and R_t^2 for the training set. If b_{train} is consistent, both $\mathcal{R}_{n_s}(b_{train}|Y_s, X_s)$ and $\mathcal{R}_{n_t}(b_{train}|Y_t, X_t)$ converge to the same limit in probability as $\tilde{n} \rightarrow \infty$.

Table 1: Schematic table and GR^2

$R_t^2 \backslash R_s^2$	high	low
high	high GR^2 , the ideal model	relatively low GR^2 , overfitting
low	relatively low GR^2	low GR^2 , underfitting

Clearly GR^2 considers both the in-sample fit and the out-of-sample fit. Intuitively, there are four scenarios for GR^2 . As summarized in Table 1, a model that fits both the training set and the test set well will have high R_t^2 and R_s^2 values and hence a high GR^2 . When overfitting occurs, the R_t^2 will be relatively high and the R_s^2 will be low, reducing the GR^2 . In contrast, the selected model may fit the training set and test set poorly, which is called underfitting. When underfitting occurs, the R_t^2 and R_s^2 will be low, reducing the GR^2 further. It is unlikely but possible that the model estimated on the training set fits the test set better (the R_s^2 is high while the R_t^2 low).

In the simulations, we illustrate the fact that penalized regression, by constraining the \mathcal{L}_γ norm of the regression estimate, obtains a superior GA compared with OLS or FSR. Penalized regression is less efficient at model selection when the norm of penalty term $\gamma > 1$, as illustrated in Figures 4 and 7. Thus, we focus on the \mathcal{L}_1 penalized or lasso-type regression.

For the simulations, we assume the outcome y is generated by the following DGP:

$$y = X'\beta + u = X_1^T\beta_1 + X_2^T\beta_2 + u$$

where $X = (x_1, \dots, x_p) \in \mathbb{R}^p$ is generated by a multivariate Gaussian distribution with zero mean, $\text{var}(x_i) = 1$, $\text{corr}(x_i, x_j) = 0.9, \forall i, j$, $\beta_1 = (2, 4, 6, 8, 10, 12)^T$ and β_2 is a $(p - 6)$ -dimensional zero vector. u is generated by a Gaussian distribution with zero mean. Here x_i does not cause x_j and there is no causal relationship between u and x_i . We set the sample size at 250, p is set at 200 or 500 and $\text{var}(u)$ at 1 or 5. Hence, we have four different cases. In each case, we repeat the simulation 50 times. In each simulation, we apply the Algorithm 1 to find the estimate of β and calculate its distance to the true value, the eGE, as well as our goodness-of-fit measure GR^2 . As a comparison, we also apply OLS in the $n \geq p$ cases and the FSR algorithm for the $n < p$ cases.

Boxplots (see Appendix 2) summarize the estimates of all the coefficients in β_1 (labeled b_1 to b_6) along with the four worst estimates among the coefficients in β_2 (labeled b_7 to b_{10}), where ‘worst’ refers to estimates with the largest bias. The lasso and OLS/FSR estimates and histograms of the GR^2 are reported for each case in Figures 9–12 (Appendix 2). Lastly, the distance between the estimates and the true values, the eGE, and the GR^2 (all averages across the 50 simulations) are reported in Table 1 for all four cases.

When $n > p$, as we can see from the boxplot in Figure (9) and (11), both lasso and OLS perform well; in the case $\text{var}(u) = 1$ both lasso and OLS perform better than when $\text{var}(u) = 5$. All the coefficient estimates are centered around the true values, and the deviations are relatively small. However, lasso outperforms OLS for the estimates of β_2 by having much smaller deviations. Indeed, a joint significance test (F test) fails to reject the null hypothesis that all coefficients in β_2 are zero for the OLS estimates. As shown in Figure (9) and (11), the GR^2 for the lasso is marginally larger than for OLS, but the differences are inconsequential.

When $n < p$, the regression model is not identified, OLS is infeasible, and we apply FSR. As shown in Figures (10) and (12), lasso still performs well and correctly selects the variables with non-zero coefficients. In contrast, although FSR also correctly identifies the non-zero coefficients, its biases and deviations are much larger than for the lasso. For the $p = 500$ case shown in Figures (10) and (12), the GR^2 for the FSR indicate that the FSR estimates are unreliable. Generally speaking,

overfitting is controlled well by lasso (all the GR^2 are close to 1) whereas the performance of FSR is mixed, as reflected by the deteriorating GR^2 as p increases. This suggests that, by imposing an \mathcal{L}_1 penalty on estimates, lasso mitigates the overfitting problem and that the advantage of lasso is likely to be more pronounced as p increases.

Table 2: Average bias, training error, generalization error, in-sample R^2 , out-of-sample R^2 , and GR^2 for lasso and OLS/FSR when $n = 250$

Measure	$var(u) = 1$		$var(u) = 5$	
	$p = 200$	$p = 500$	$p = 200$	$p = 500$
Bias				
b_{Lasso}	0.7923	0.8810	3.8048	4.1373
$b_{OLS/FSR}$	0.9559	11.7530	4.7797	13.7622
Training error				
Lasso	0.9167	0.8625	22.2476	21.1334
OLS/FSR	0.2164	832.9988	5.4097	1034.2636
Generalization error				
Lasso	1.1132	1.1478	27.8672	28.5125
OLS/FSR	5.2109	852.5822	134.8725	1070.6329
R^2 , in-sample				
Lasso	0.9994	0.9994	0.9866	0.9867
OLS/FSR	0.9999	0.4678	0.9967	0.3619
R^2 , out-of-sample				
Lasso	0.9993	0.9993	0.9830	0.9826
OLS/FSR	0.9967	0.4681	0.9181	0.3627
GR^2				
Lasso	0.9988	0.9987	0.9698	0.9695
OLS/FSR	0.9965	0.3659	0.9151	0.2935

Table 1 reinforces the impressions from the boxplots and histograms. When $p = 200$ OLS of course performs extremely well in terms of training error and more poorly in terms of generalization error while its GR^2 is very close to the lasso value. For $n < p$ what is noteworthy is the stable performance of the lasso relative to that of FSR. The training errors, generalization errors, and GR^2 are particularly poor for FSR, again illustrating the advantage of the lasso in avoiding overfitting.

5 Conclusion

In this paper, we study the performance of penalized and unpenalized extremum estimators from the perspective of generalization ability (GA). We propose inequalities for the extremum estimators, which bound empirical out-of-sample prediction errors by a function of the in-sample errors, sample sizes and model complexity. The inequalities serve not only as a measurement of GA, but also illustrate the trade-off between in-sample fit and out-of-sample fit, which in turn may be used for tuning parameter optimization, such as the number of folds K in cross-validation or n_t/n_s in validation. Moreover, the inequalities reveal that the GA of the extremum estimator is also affected by the heaviness in the tail of its error distribution. In Section 3, we apply the results of Section 2

to analyze the GA of penalized regression estimates for both $n \geq p$ and $n < p$ cases. We show that some finite-sample and asymptotic properties of the penalized estimators can be directly explained by the GA of the estimators. Furthermore, we use the new bounds to quantify the difference between penalized and corresponding unpenalized regression estimates in the \mathcal{L}_2 norm under both $n \geq p$ and $p > n$ cases.

Our work sheds new light not only on penalized regression, but also more generally on the applicability of GA for model selection and estimation, and in particular offering additional insight into the bias-variance trade-off. In this paper, we focus mainly on implementing penalized regression through minimization of eGE. However, other penalty methods could potentially fit into GA analysis, such as penalized MLE, functional regression, principle component analysis and decision trees. Furthermore, the results on penalized regression and GA analysis may be implemented on other empirical methods, like the EM algorithm, clustering, mixture and factor modeling, Bayes networks, and so on.

In providing a common property applying to all penalized regressions, our result in Section 3 necessarily reflect a worst-case scenario. Finer eGE bounds may well be derived if each specific penalized regression method were evaluated separately. Lastly, as an early attempt to incorporate GA analysis into econometrics, we focus on the i.i.d. case. However, it is clear that our framework has the potential to be generalized to non-i.i.d. data in settings like α - and β -mixing stationery time series data or dependent and non-identical panel data.

References

References

- Hirotsugu Akaike. Information theory and an extension of the maximum likelihood principle. In B. N. Petrov and F. Csaki, editors, *2nd International Symposium on Information Theory, Tsahkadsor, Armenia, USSR*, pages 267–281. Budapest: Akademiai Kiado, 1973.
- Takeshi Amemiya. *Advanced econometrics*. Harvard university press, 1985.
- Alexandre Belloni, Victor Chernozhukov, Ivan Fernández-Val, and Chris Hansen. Program evaluation with high-dimensional data. *arXiv preprint arXiv:1311.2645*, 2013.
- Peter J Bickel, Ya’acov Ritov, and Alexandre B Tsybakov. Simultaneous analysis of lasso and dantzig selector. *The Annals of Statistics*, 37:1705–1732, 2009.
- Richard Blundell, Monica Costa Dias, Costas Meghir, and John Reenen. Evaluating the employment impact of a mandatory job search program. *Journal of the European economic association*, 2(4): 569–606, 2004.
- Leo Breiman. Better subset regression using the nonnegative garrote. *Technometrics*, 37(4): 373–384, 1995.
- Emmanuel J. Candes and Terence Tao. The dantzig selector: statistical estimation when p is much larger than n . *The Annals of Statistics*, pages 2313–2351, 2007.
- Mehmet Caner. Lasso-type gmm estimator. *Econometric Theory*, 25(1):270–290, 2009.
- Nicolo Cesa-Bianchi, Alex Conconi, and Claudio Gentile. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.
- David M. Chickering, David Heckerman, and Christopher Meek. Large-sample learning of Bayesian networks is NP-hard. *Journal of Machine Learning Research*, 5:1287–1330, 2004.
- Peter Dolton. *The econometric evaluation of the New Deal for Lone Parents*. PhD thesis, Department of Economics, University of Michigan, 2006.
- Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.
- Ildiko E. Frank and Jerome H. Friedman. A statistical view of some chemometrics regression tools. *Technometrics*, 35(2):109–135, 1993.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. A note on the group lasso and a sparse group lasso. *arXiv preprint arXiv:1001.0736*, 2010.
- Nir Friedman, Dan Geiger, and Moises Goldszmidt. Bayesian network classifiers. *Machine Learning*, 29(2-3):131–163, 1997.

- Nir Friedman, Michal Linial, Iftach Nachman, and Dana Pe'er. Using Bayesian networks to analyze expression data. In *Proceedings of the Fourth Annual International Conference on Computational Molecular Biology*, RECOMB '00, pages 127–135, New York, NY, USA, 2000. ACM.
- Wenjiang J. Fu. Penalized regressions: the bridge versus the Lasso. *Journal of computational and graphical statistics*, 7(3):397–416, 1998.
- Michael Gechter. Generalizing the results from social experiments: Theory and evidence from mexico and india. *manuscript, Pennsylvania State University*, 2015.
- Peter Hall and James S Marron. Local minima in cross-validation functions. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 245–252, 1991.
- Peter Hall, Jeff Racine, and Qi Li. Cross-validation and the estimation of conditional probability densities. *Journal of the American Statistical Association*, 2011.
- David Heckerman, Dan Geiger, and David M. Chickering. Learning Bayesian networks: The combination of knowledge and statistical data. *Machine learning*, 20(3):197–243, 1995.
- Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- Arthur E. Hoerl and Robert W. Kennard. Ridge regression: Applications to nonorthogonal problems. *Technometrics*, 12(1):69–82, 1970.
- Ting Hu and Ding-Xuan Zhou. Online learning with samples drawn from non-identical distributions. *Journal of Machine Learning Research*, 10(Dec):2873–2898, 2009.
- Jian Huang, Joel L. Horowitz, and Shuangge Ma. Asymptotic properties of bridge estimators in sparse high-dimensional regression models. *The Annals of Statistics*, pages 587–613, 2008.
- Sham M Kakade and Ambuj Tewari. On the generalization ability of online strongly convex programming algorithms. pages 801–808, 2009.
- Keith Knight and Wenjiang Fu. Asymptotics for Lasso-type estimators. *Annals of statistics*, pages 1356–1378, 2000.
- Ron Kohavi et al. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Ijcai*, volume 14, pages 1137–1145, 1995.
- Daniel J McDonald, Cosma Rohilla Shalizi, and Mark Schervish. Generalization error bounds for stationary autoregressive models. *arXiv preprint arXiv:1103.0942*, 2011.
- Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the Lasso. *The Annals of Statistics*, pages 1436–1462, 2006.

- Nicolai Meinshausen and Bin Yu. Lasso-type recovery of sparse representations for high-dimensional data. *The Annals of Statistics*, pages 246–270, 2009.
- A. Michalski and A.I. Yashin. Structural minimization of risk on estimation of heterogeneity distributions. 1986. URL <http://pure.iiasa.ac.at/2785/1/WP-86-076.pdf>.
- Mehryar Mohri and Afshin Rostamizadeh. Rademacher complexity bounds for non-iid processes. pages 1097–1104, 2009.
- Whitney K. Newey and Daniel McFadden. Large sample estimation and hypothesis testing. *Handbook of econometrics*, 4:2111–2245, 1994.
- Muhammad Aslam Noor. Differentiable non-convex functions and general variational inequalities. *Applied Mathematics and Computation*, 199(2):623–630, 2008.
- Judea Pearl. Detecting latent heterogeneity. *Sociological Methods & Research*, page 0049124115600597, 2015.
- Gideon E. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6(2):461–464, 1978.
- Jun Shao. Asymptotic theory for model selection. *Statistica Sinica*, 7:221–242, 1997.
- Anders Skrondal and Sophia Rabe-Hesketh. *Generalized latent variable modeling: Multilevel, longitudinal, and structural equation models*. Crc Press, 2004.
- Steve Smale and Ding-Xuan Zhou. Online learning with markov sampling. *Analysis and Applications*, 7(01):87–113, 2009.
- James H Stock and Mark W Watson. Generalized shrinkage methods for forecasting using many predictors. *Journal of Business & Economic Statistics*, 30(4):481–493, 2012.
- M. Stone. Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society, Series B (Methodological)*, 36(2):111–147, 1974.
- M. Stone. An asymptotic equivalence of choice of model by cross-validation and akaike’s criterion. *Journal of the Royal Statistical Society, Series B (Methodological)*, 39(1):44–47, 1977.
- Roman G Strongin and Yaroslav D Sergeyev. *Global optimization with non-convex constraints: Sequential and parallel algorithms*, volume 45. Springer Science & Business Media, 2013.
- Yongqiang Tang. A hoeffding-type inequality for ergodic time series. *Journal of Theoretical Probability*, 20(2):167–176, 2007.
- Robert Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society, Series B (Methodological)*, 58:267–288, 1996.
- Joel A Tropp. Greed is good: Algorithmic results for sparse approximation. *Information Theory, IEEE Transactions on*, 50(10):2231–2242, 2004.

- Vladimir N. Vapnik and Alexey Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theoretical Probability and its Applications*, 16(2):264–280, 1971a.
- Vladimir N. Vapnik and Alexey Y. Chervonenkis. Theory of uniform convergence of frequency of appearance of attributes to their probabilities and problems of defining optimal solution by empiric data. *Avtomatika i Telemekhanika*, (2):42–53, 1971b.
- Vladimir N. Vapnik and Alexey Y. Chervonenkis. The method of ordered risk minimization, I. *Avtomatika i Telemekhanika*, (8):21–30, 1974a.
- Vladimir N. Vapnik and Alexey Y. Chervonenkis. On the method of ordered risk minimization, II. *Avtomatika i Telemekhanika*, (9):29–39, 1974b.
- Li-Wei Wang and Ju-Fu Feng. Learning gaussian mixture models by structural risk minimization. In *2005 International Conference on Machine Learning and Cybernetics*, volume 8, pages 4858–4863. IEEE, 2005.
- Jon A Wellner. A glivenko-cantelli theorem for empirical measures of independent but non-identically distributed random variables. *Stochastic Processes and their Applications*, 11(3): 309–312, 1981.
- Ning Xu, Jian Hong, and Timothy C.G. Fisher. Clustered model selection and clustered model averaging: Simultaneous heterogeneity control and modeling. School of Economics, The University of Sydney, 2016.
- Liexiang Yan and Dexian Ma. Global optimization of non-convex nonlinear programs using line-up competition algorithm. *Computers & Chemical Engineering*, 25(11):1601–1610, 2001.
- Bin Yu. Rates of convergence for empirical processes of stationary mixing sequences. *The Annals of Probability*, pages 94–116, 1994.
- Chun-Nam John Yu and Thorsten Joachims. Learning structural svms with latent variables. In *Proceedings of the 26th annual international conference on machine learning*, pages 1169–1176. ACM, 2009.
- Hao Yu. A glivenko-cantelli lemma and weak convergence for empirical processes of associated sequences. *Probability theory and related fields*, 95(3):357–370, 1993.
- Cun-Hui Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38:894–942, 2010.
- Cun-Hui Zhang and Jian Huang. The sparsity and bias of the Lasso selection in high-dimensional linear regression. *The Annals of Statistics*, 36:1567–1594, 2008.
- Tong Zhang. On the consistency of feature selection using greedy least squares regression. *Journal of Machine Learning Research*, 10:555–568, 2009.

Peng Zhao and Bin Yu. On model selection consistency of Lasso. *The Journal of Machine Learning Research*, 7:2541–2563, 2006.

Hui Zou. The adaptive Lasso and its oracle properties. *Journal of the American statistical association*, 101(476):1418–1429, 2006.

Appendix 1

Proof. Theorem 2.1. Since $b_{train} = \operatorname{argmin}_b \mathcal{R}_{n_t}(b|Y_t, X_t)$. Eq. (2) forms an upper bound for the generalization error with probability $1 - 1/n_t, \forall b$,

$$\mathcal{R}(b_{train}|Y, X) \leq \mathcal{R}_{n_t}(b_{train}|Y_t, X_t) (1 - \sqrt{\varepsilon})^{-1}$$

where $\mathcal{R}_{n_t}(b_{train}|Y_t, X_t)$ stands for the training error on (Y_t, X_t) , $\mathcal{R}(b_{train}|Y, X)$ stands for the true population error of b_{train} and $\varepsilon = (1/n_t) \{h \ln[(n_t/h)] + h - \ln(1/n_t)\}$.

To generalize eq. (2) to quantify the relation between eGE and and eTE, we need to think about if the loss function $Q(b_{train}|y, \mathbf{x})$ is of heavy tail, light tail or no tail.

no-tail If the loss function $Q(\cdot)$ is bounded between $[0, B]$, where $B \in (0, \infty)$. Based on Hoeffding's inequality [Hoeffding, 1963] for the extremum estimator b_{train} , the empirical process satisfies, $\forall \zeta \geq 0$,

$$\mathbb{P}\{|\mathcal{R}_{n_s}(b_{train}|X_s, Y_s) - \mathcal{R}(b_{train}|X, Y)| \leq \zeta\} \geq 1 - 2 \exp\left(-\frac{2n\zeta}{B}\right) \quad (17)$$

If we define $\varpi = 1 - 2 \exp(-2n^2\zeta/\sum_{i=1}^n B_i)$, then

$$\zeta = \frac{B}{n} \ln \sqrt{\frac{2}{1-\varpi}} \quad (18)$$

This implies, for any extremum estimator b_{train}

$$\mathbb{P}\{\mathcal{R}_{n_s}(b_{train}|X_s, Y_s) \leq \mathcal{R}(b_{train}|X, Y) + \zeta\} \geq \varpi. \quad (19)$$

Since both training and test sets are randomly sampled from population, eq. (2) could be modified and relaxed as follows: $\forall \zeta \geq 0, \forall \tau_1 \geq 0, \exists N_1 \in \mathbb{R}^+$ subject to

$$\mathbb{P}\left\{\mathcal{R}_{n_s}(b_{train}|X_s, Y_s) \leq \frac{\mathcal{R}_{n_t}(b_{train}|X_t, Y_t)}{1 - \sqrt{\varepsilon}} + \zeta\right\} \geq \varpi \left(1 - \frac{1}{n_t}\right) \quad (20)$$

light-tail Suppose the loss function $Q(\cdot)$ is unbounded but still \mathcal{F} -measurable, and specifically it possesses finite v^{th} moment when $v > 2$. Based on Chebyshev's inequality for the extremum estimator b_{train} , the empirical process satisfies, $\forall \zeta \geq 0$,

$$\mathbb{P}\{|\mathcal{R}_{n_s}(b_{train}|X_s, Y_s) - \mathcal{R}(b_{train}|X, Y)| \leq \zeta\} \geq 1 - \frac{\operatorname{var}(Q(b_{train}|y, \mathbf{x}))}{n\zeta} \quad (21)$$

If we define $\varpi = 1 - \frac{\operatorname{var}(Q(b_{train}|y, \mathbf{x}))}{n\zeta}$, then

$$\zeta = \frac{\operatorname{var}(Q(b_{train}|y, \mathbf{x}))}{n(1-\varpi)} \quad (22)$$

This implies, for any extremum estimator b_{train}

$$\mathbb{P}\{\mathcal{R}_{n_s}(b_{train}|X_s, Y_s) \leq \mathcal{R}(b_{train}|X, Y) + \zeta\} \geq \varpi. \quad (23)$$

Since both training and test sets are randomly sampled from population, eq. (2) could be modified and relaxed as follows: $\forall \zeta \geq 0, \forall \tau_1 \geq 0, \exists N_1 \in \mathbb{R}^+$ subject to

$$\mathbb{P}\left\{\mathcal{R}_{n_s}(b_{train}|X_s, Y_s) \leq \frac{\mathcal{R}_{n_t}(b_{train}|X_t, Y_t)}{1 - \sqrt{\varepsilon}} + \zeta\right\} \geq \varpi \left(1 - \frac{1}{n_t}\right) \quad (24)$$

heavy-tail Suppose the loss function $Q(\cdot)$ is unbounded but still \mathcal{F} -measurable, which has heavy tails with the property that, for $1 < \nu \leq 2, \exists \tau$, such that

$$\sup \frac{\sqrt[\nu]{\int [Q(b_{train}|y, \mathbf{x})]^\nu dF(y, \mathbf{x})}}{\int Q(b_{train}|y, \mathbf{x}) dF(y, \mathbf{x})} \leq \tau. \quad (25)$$

Based on the Bahr-Esseen inequality for the extremum estimator b_{train} , the empirical process satisfies, $\forall \zeta \geq 0$,

$$\begin{aligned} \mathbb{P}\{|\mathcal{R}(b_{train}|Y, X) - \mathcal{R}_{n_s}(b_{train}|Y_s, X_s)| \leq \zeta\} &\geq 1 - 2 \cdot \frac{\mathbb{E}[Q(b_{train}|y, \mathbf{x})^\nu]}{\zeta^\nu \cdot n_s^{\nu-1}} \\ &\geq 1 - 2\tau^\nu \cdot \frac{(\mathbb{E}[Q(b_{train}|y, \mathbf{x})])^\nu}{\zeta^\nu \cdot n_s^{\nu-1}} \end{aligned} \quad (26)$$

If we define

$$\varpi = 1 - 2\tau^\nu \cdot \frac{(\mathbb{E}[Q(b_{train}|y, \mathbf{x})])^\nu}{\zeta^\nu \cdot n_s^{\nu-1}} \quad (27)$$

then

$$\zeta = \frac{\sqrt[\nu]{2} \cdot \tau (\mathbb{E}[Q(b_{train}|y, \mathbf{x})])}{\sqrt[\nu]{1 - \varpi} \cdot n_s^{1-1/\nu}} \quad (28)$$

This implies, for any extremum estimator b_{train}

$$\mathbb{P}\{\mathcal{R}_{n_s}(b_{train}|X_s, Y_s) \leq \mathcal{R}(b_{train}|X, Y) + \zeta\} \geq \varpi. \quad (29)$$

Since both training and test sets are randomly sampled from population, eq. (2) could be modified and relaxed as follows: $\forall \zeta \geq 0, \forall \tau_1 \geq 0, \exists N_1 \in \mathbb{R}^+$ subject to

$$\mathbb{P}\left\{\mathcal{R}_{n_s}(b_{train}|X_s, Y_s) \leq \frac{\mathcal{R}_{n_t}(b_{train}|X_t, Y_t)}{1 - \sqrt{\varepsilon}} + \zeta\right\} \geq \varpi \left(1 - \frac{1}{n_t}\right) \quad (30)$$

□

Proof. Theorem 2.2. The upper bound of eGE for cross-validation could be established by adapting eq. (2) and (3). The proof of both Theorems are quite similar, expect for one fact : the eq. (2) and (3) measure the upper bound by one-shot eTE while eq. (4) by averaged multiple-round eTE,

$$\frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_t}(b_{train}|Y_t^q, X_t^q).$$

Hence, eq. (2) could be generalized as the following

$$\begin{aligned} \mathcal{R}(b_{train}|Y, X) &= \mathbb{E}[Q(b_{train}|y, \mathbf{x})] \\ &= \mathbb{E}\left[\frac{1}{n_t} \sum_i Q(b_{train}|y_t^i, \mathbf{x}_t^i)\right] \\ &= \mathbb{E}[\mathcal{R}_{n_t}(b_{train}|Y_t, X_t)] \\ &\leq \frac{\frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_t}(b_{train}|Y_t^q, X_t^q)}{1 - \sqrt{\varepsilon}}. \end{aligned} \quad (31)$$

Also, since

$$\begin{aligned} \mathcal{R}(b_{train}|Y, X) &= \mathbb{E}[Q(b_{train}|y, \mathbf{x})] \\ &= \mathbb{E}\left[\frac{1}{n_s} \sum_{i=1}^{n_s} Q(b_{train}|y_s^i, \mathbf{x}_s^i)\right] \\ &= \mathbb{E}[\mathcal{R}_{n_s}(b_{train}|Y_s, X_s)] \\ &= \mathbb{E}\left[\frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_s}(b_{train}|Y_s^q, X_s^q)\right], \end{aligned} \quad (32)$$

we can derive the following inequality

$$\begin{aligned} &\mathbb{P}\left\{\left|\mathbb{E}\left[\frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_s}(b_{train}|Y_s, X_s)\right] - \frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_s}(b_{train}|Y_s, X_s)\right| \leq \zeta_{cv}\right\} \\ &= \mathbb{P}\left\{\left|\mathcal{R}(b_{train}|Y, X) - \frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_s}(b_{train}|Y_s, X_s)\right| \leq \zeta_{cv}\right\} \\ &\geq \varpi \end{aligned} \quad (33)$$

where, similar to the proof of Theorem 2.1,

$$\zeta_{cv} = \begin{cases} (\sqrt[2]{2} \cdot \tau \cdot \mathcal{R}(b_{train}|Y, X)) / (\sqrt[2]{1 - \varpi} \cdot n_s^{1-1/\nu}) & \text{if } \nu \in (1, 2] \\ (B \ln \sqrt{2/(1 - \varpi)}) / n_s & \text{if } Q(\cdot) \in (0, B] \text{ and } B \text{ is bounded} \\ \text{var}(Q(b_{train}|y, \mathbf{x})) / ((n_s)^2 (1 - \varpi)) & \text{if } \nu \in (2, \infty) \end{cases} \quad (34)$$

As the result, we would have

$$\frac{1}{K} \sum_{j=1}^K \mathcal{R}_{n_s}(b_{train}|Y_s^j, X_s^j) \leq \frac{\frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_t}(b_{train}|Y_t^q, X_t^q)}{(1 - \sqrt{\varepsilon})} + \zeta_{cv} \quad (35)$$

and Theorem 2. □

Proof. Lemma 3.1. eq. (12) and (13) are the direct application of eq. (3) and (4). The only thing we need to be specific is the last term of the RHS, ζ and ζ_{cv} .

Since, for classical regression analysis, the error term u is distributed as zero-mean Gaussian, the loss function of OLS,

$$Q(b_{OLS}) \sim \text{var}(u) \cdot \chi^2(1).$$

Hence in eq. (12) the last RHS term

$$\zeta = \frac{2(\text{var}(u))^2}{n_s \sqrt{1 - \bar{\omega}}}$$

Furthermore, in cross-validation,

$$\mathcal{R}_{n_s}(b_{OLS}|Y, X) \sim \frac{\text{var}(u)}{n(K-1)/K} \cdot \chi^2(n(K-1)/K).$$

Hence, in eq. (13) the last RHS term

$$\zeta_{cv} = \frac{2(\text{var}(u))^2}{\sqrt{1 - \bar{\omega}}(n/K)^2}$$

By plunging the value of ζ and ζ_{cv} into eq. (3) and (4), we have eq. (12) and (13). □

Proof. Corollary 3.1. The optimal K or n_t/n_s could be obtained by finding the smallest expectation of the RHS for eq. (8) and (9).

Since, for classical regression analysis, the error term u is distributed as zero-mean Gaussian, the loss function of OLS,

$$Q(b_{OLS}) \sim \text{var}(u) \cdot \chi^2(1).$$

As a result,

$$\mathcal{R}_{n_s}(b_{OLS}|Y, X) \sim \frac{\text{var}(u)}{n(K-1)/K} \cdot \chi^2(n(K-1)/K).$$

Hence in the RHS of eq. (12),

$$\frac{1}{K} \sum_{j=1}^K \mathcal{R}_{n_t}(b_{train}|Y_t^j, X_t^j) \sim \frac{\text{var}(u)}{n(K-1)/K} \cdot \text{Gamma}(n(K-1)/2, 2/K)$$

Hence, the expectation of RHS for eq. (13) is equal to

$$\frac{\text{var}(u)}{1 - \sqrt{\varepsilon}} + \frac{4(\text{var}(u))^2}{\sqrt{1 - \overline{\omega}(n/K)^2}}$$

and

$$K^* = \underset{K}{\text{argmin}} \frac{\text{var}(u)}{1 - \sqrt{\varepsilon}} + \frac{4(\text{var}(u))^2}{\sqrt{1 - \overline{\omega}(n/K)^2}}$$

□

Proof. **Proposition 3.1.** Given A1–A4, the true DGP is

$$y_i = x_i^T \beta + u_i, \quad i = 1, \dots, n.$$

Proving that the true DGP has the highest generalization ability (the lowest eGE) is equivalent to proving, in a test set, that

$$\frac{\sum_{i=1}^n (y_i - x_i^T \beta)^2}{n} \leq \frac{\sum_{i=1}^n (y_i - x_i^T b)^2}{n}, \quad (36)$$

which is equivalent to proving that

$$\begin{aligned} 0 &\leq \frac{1}{n} \sum_{i=1}^n \left[(y_i - x_i^T b)^2 - (y_i - x_i^T \beta)^2 \right] \\ \iff 0 &\leq \frac{1}{n} \sum_{i=1}^n (y_i - x_i^T b + y_i - x_i^T \beta) (y_i - x_i^T b - y_i + x_i^T \beta) \\ \iff 0 &\leq \frac{1}{n} \sum_{i=1}^n (y_i - x_i^T b + y_i - x_i^T \beta) (x_i^T \beta - x_i^T b). \end{aligned}$$

Defining $\delta = \beta - b$, it follows,

$$\begin{aligned} 0 &\leq \frac{1}{n} \sum_{i=1}^n (2y_i - x_i^T b - x_i^T \beta) (x_i^T \delta) \\ \iff 0 &\leq \frac{1}{n} \sum_{i=1}^n (2y_i - x_i^T \beta + x_i^T \beta - x_i^T b - x_i^T \beta) (x_i^T \delta) \\ \iff 0 &\leq \frac{1}{n} \sum_{i=1}^n (2y_i - 2x_i^T \beta + x_i^T \delta) (x_i^T \delta) \\ \iff 0 &\leq \frac{1}{n} \sum_{i=1}^n (2u_i + x_i^T \delta) (x_i^T \delta) \end{aligned}$$

Hence, proving (36) is equivalent to proving

$$0 \leq \frac{1}{n} \sum_{i=1}^n (2u_i + x_i^T \delta) (x_i^T \delta)$$

Since $\mathbb{E}(X^T u) = \mathbf{0}$ (A2), it follows that

$$\frac{1}{n} \sum_{i=1}^n u_i \cdot x_i \xrightarrow{\mathbf{P}} \mathbf{0} \iff \frac{1}{n} \sum_{i=1}^n (u_i \cdot x_i^T) \beta \xrightarrow{\mathbf{P}} 0 \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n (u_i \cdot x_i^T) b \rightarrow 0$$

Hence, asymptotically

$$\frac{1}{n} \sum_{i=1}^n (2u_i + x_i^T \delta) (x_i^T \delta) = \frac{1}{n} \sum_{i=1}^n 2\delta u_i x_i^T + \frac{1}{n} \sum_{i=1}^n (x_i^T \delta)^2 \xrightarrow{\mathbf{P}} \mathbb{E} (x_i^T \delta)^2 \geq 0$$

□

Proof. Theorem 3.1. In the proof of Theorem 2.1 we defined $b_{train} = \operatorname{argmin}_b \mathcal{R}_{n_t}(b|Y_t, X_t)$, meaning that b_{train} is the extremum estimator without penalty on any training set. We also have $\bar{M} = R_{n_t}(b_{train}|Y_t, X_t) (1 - \sqrt{\varepsilon})^{-1}$. Theorem 2.1 shows the following bound holds $\forall b \in \Lambda$

$$\mathcal{R}_{n_s}(b|Y_s, X_s) \leq \mathcal{R}_{n_t}(b|Y_t, X_t) (1 - \sqrt{\varepsilon})^{-1} + \zeta$$

with probability at least $(1 - 1/n_t)\bar{\omega}$. Also, among all the $b \in \{b_\lambda\}$, b_{lasso} has the lowest eGE on the test set,

$$\mathcal{R}_{n_s}(b_{lasso}|Y_s, X_s) \leq \mathcal{R}_{n_s}(b|Y_s, X_s)$$

we have

$$\frac{1}{n_s} \|Y_s - X_s b_{lasso}\|_2^2 \leq \frac{1}{n_t} \|Y_t - X_t b_{train}\|_2^2 (1 - \sqrt{\varepsilon})^{-1} + \zeta$$

By defining $\Delta = b_{train} - b_{lasso}$, $Y_t - X_t b_{train} = e_t$ and $Y_s - X_s b_{train} = e_s$,

$$\begin{aligned} \frac{1}{n_s} \|Y_s - X_s b_{lasso}\|_2^2 &= \frac{1}{n_s} \|Y_s - X_s b_{train} + X_s \Delta\|_2^2 \\ &= \frac{1}{n_s} \|e_s + X_s \Delta\|_2^2 \\ &= \frac{1}{n_s} (e_s + X_s \Delta)^T (e_s + X_s \Delta) \\ &= \frac{1}{n_s} \left(\|e_s\|_2^2 + 2e_s^T X_s \Delta + \Delta^T X_s^T X_s \Delta \right) \end{aligned}$$

Hence,

$$\frac{1}{n_s} \|Y_s - X_s b_{lasso}\|_2^2 \leq \frac{1}{n_t} \|Y_t - X_t b_{train}\|_2^2 (1 - \sqrt{\varepsilon})^{-1} + \zeta$$

implies

$$\frac{1}{n_s} \|e_s\|_2^2 + \frac{2}{n_s} e_s^T X_s \Delta + \frac{1}{n_s} \Delta^T X_s^T X_s \Delta \leq \frac{\frac{1}{n_t} \|e_t\|_2^2}{1 - \sqrt{\varepsilon}} + \zeta.$$

It follows that

$$\frac{1}{n_s} \|X_s \Delta\|_2^2 \leq \left(\frac{1}{n_t} \frac{\|e_t\|_2^2}{1 - \sqrt{\varepsilon}} - \frac{1}{n_s} \|e_s\|_2^2 \right) - \frac{2}{n_s} e_s^T X_s \Delta + \zeta.$$

By the Holder inequality,

$$-e_s^T X_s \Delta \leq |e_s^T X_s \Delta| \leq \|e_s^T X_s\|_\infty \|\Delta\|_1.$$

It follows that

$$\frac{1}{n_s} \|X_s \Delta\|_2^2 \leq \left(\frac{1}{n_t} \frac{\|e_t\|_2^2}{1 - \sqrt{\varepsilon}} - \frac{1}{n_s} \|e_s\|_2^2 \right) + \frac{2}{n_s} \|e_s^T X_s\|_\infty \|\Delta\|_1 + \zeta.$$

Also, since $\|b_{lasso}\|_1 \leq \|b_{train}\|_1$

$$\begin{aligned} \|\Delta\|_1 &= \|b_{train} - b_{lasso}\|_1 \\ &\leq \|b_{lasso}\|_1 + \|b_{train}\|_1 \\ &\leq 2 \|b_{train}\|_1 \end{aligned}$$

As a result, we have

$$\frac{1}{n_s} \|X_s \Delta\|_2^2 \leq \left(\frac{1}{n_t} \frac{\|e_t\|_2^2}{1 - \sqrt{\varepsilon}} - \frac{1}{n_s} \|e_s\|_2^2 \right) + \frac{4}{n_s} \|e_s^T X_s\|_\infty \|b_{train}\|_1 + \zeta \quad (37)$$

where we denote \hat{y}_s to be the extremum estimator prediction on the test set calculated on the training set and \hat{y}_s^{lasso} to be the lasso prediction on the test set. It follows that (37) is the bound for $\mathbb{E} \left(\|\hat{y}_s - \hat{y}_s^{lasso}\|_2^2 \right)$, the expected difference between lasso prediction and extremum estimator prediction on the test set. The bound holds with probability $(1 - 1/n_t)\bar{\omega}$. \square

Proof. Corollary 2. We need to prove that the VC inequality and SRM also hold for lasso with cross-validation. If lasso is implemented by K -fold cross-validation, the sample is partitioned into K equal-sized folds. If $K = 2$, the theoretical result from K -fold cross-validation is identical to Theorem 3.1. Thus, we only discuss the case of $K \geq 3$ here.

For $K \geq 3$, we have K different test sets for λ -tuning and K different training sets for estimation. Denote the q^{th} training set as (X_t^q, Y_t^q) , the q^{th} test set as (X_s^q, Y_s^q) , the extremum estimator estimated from the k^{th} training set as b_{train}^k , the sample size for each test set as n_s and the sample size for each training set as n_t . Based on Theorem 1, for each test set, the following bound holds for k and $q \in [1, K]$ with probability at least $(1 - 1/n_t)\bar{\omega}_k$

$$\mathcal{R}_{n_s}(b_{train}^k | X_s^q, Y_s^q) \leq \mathcal{R}_{n_t}(b_{train}^k | X_t^q, Y_t^q) (1 - \sqrt{\varepsilon_k})^{-1} + \zeta_k.$$

Hence,

$$\begin{aligned} \frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_s}(b_{train}^k | X_s^q, Y_s^q) &\leq \mathcal{R}_{n_s}(\bar{b}_{train} | X_s^{q*}, Y_s^{q*}) \\ &\leq \mathcal{R}_{n_t}(\bar{b}_{train} | X_t^{q*}, Y_t^{q*}) (1 - \sqrt{\bar{\varepsilon}})^{-1} + \bar{\zeta} \end{aligned}$$

Since b_{lasso} minimizes $(1/K) \sum_{q=1}^K \mathcal{R}_{n_s}(b|X_s^q, Y_s^q)$,

$$\frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_s}(b_{lasso}|X_s^q, Y_s^q) \leq \frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_s}(b_{train}^k|X_s^q, Y_s^q), \quad \forall k \in [1, K]$$

It follows that

$$\frac{1}{K} \sum_{q=1}^K \mathcal{R}_{n_s}(b_{lasso}|X_s^q, Y_s^q) \leq \mathcal{R}_{n_t}(\bar{b}_{train}|X_t^{q^*}, Y_t^{q^*}) (1 - \sqrt{\bar{\epsilon}})^{-1} + \bar{\zeta}.$$

Denote $\mathcal{R}_{n_t}(b_{train}^{k^*}|X_t^{q^*}, Y_t)$ by \bar{e}_t and $\mathcal{R}_{n_s}(b_{train}^{k^*}|X_s^q, Y_s^{q^*})$ by \bar{e}_s . The above equation is equivalent to

$$\frac{1}{K} \sum_{q=1}^K \left(\frac{1}{n_s} \|Y_s^q - X_s^q b_{lasso}\|_2^2 \right) \leq \frac{\|\bar{e}_t\|_2^2}{n_t} \frac{1}{1 - \sqrt{\bar{\epsilon}}} + \bar{\zeta}.$$

By defining $\Delta = \bar{b}_{train} - b_{lasso}$ and $\bar{e}_s^q = Y_s^q - X_s^q \bar{b}_{train}$ we have

$$\begin{aligned} \frac{1}{n_s} \|Y_s^q - X_s^q b_{lasso}\|_2^2 &= \frac{1}{n_s} \|Y_s^q - X_s^q \bar{b}_{train} + X_s^q \Delta\|_2^2 \\ &= \frac{1}{n_s} \|\bar{e}_s^q + X_s^q \Delta\|_2^2 \\ &= \frac{1}{n_s} (\bar{e}_s^q + X_s^q \Delta)^T (\bar{e}_s^q + X_s^q \Delta) \\ &= \frac{1}{n_s} \left(\|\bar{e}_s^q\|_2^2 + 2 (\bar{e}_s^q)^T X_s^q \Delta + \Delta^T (X_s^q)^T X_s^q \Delta \right). \end{aligned}$$

Hence,

$$\frac{1}{K} \sum_{q=1}^K \left(\frac{1}{n_s} \|Y_s^q - X_s^q b_{lasso}\|_2^2 \right) \leq \frac{1}{n_t} \|Y_t^q - X_t^q \bar{b}_{train}\|_2^2 (1 - \sqrt{\bar{\epsilon}})^{-1} + \bar{\zeta}$$

implies

$$\frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} \|\bar{e}_s^q\|_2^2 + \frac{1}{K} \sum_{q=1}^K \frac{2}{n_s} (\bar{e}_s^q)^T X_s^q \Delta + \frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} \Delta^T (X_s^q)^T (X_s^q) \Delta \leq \frac{1}{1 - \sqrt{\bar{\epsilon}}} \frac{\|\bar{e}_t\|_2^2}{n_t} + \bar{\zeta}.$$

It follows that

$$\frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} \|X_s^q \Delta\|_2^2 \leq \frac{1}{n_t} \frac{\|\bar{e}_t\|_2^2}{1 - \sqrt{\bar{\epsilon}}} - \frac{1}{K} \sum_{q=1}^K \frac{\|\bar{e}_s^q\|_2^2}{n_s} - \frac{1}{K} \sum_{q=1}^K \frac{2}{n_s} (\bar{e}_s^q)^T X_s^q \Delta + \bar{\zeta}.$$

By the Holder inequality,

$$-1 \cdot (\bar{e}_s^q)^T X_s^q \Delta \leq |(\bar{e}_s^q)^T X_s^q \Delta| \leq \left\| (\bar{e}_s^q)^T X_s^q \right\|_{\infty} \|\Delta\|_1.$$

It follows that

$$\frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} \|X_s^q \Delta\|_2^2 \leq \left| \frac{1}{n_t} \frac{\|\bar{e}_t\|_2^2}{1 - \sqrt{\bar{\varepsilon}}} - \frac{1}{K} \sum_{q=1}^K \frac{\|\bar{e}_s^q\|_2^2}{n_s} \right| + \frac{1}{K} \sum_{q=1}^K \frac{2}{n_s} \left\| \left(\bar{e}_s^q \right)^T X_s^q \right\|_\infty \|\Delta\|_1 + \bar{\varsigma}.$$

Also, since $\|b_{lasso}\|_1 \leq \|\bar{b}_{train}\|_1$

$$\begin{aligned} \|\Delta\|_1 &= \|\bar{b}_{train} - b_{lasso}\|_1 \\ &\leq \|b_{lasso}\|_1 + \|\bar{b}_{train}\|_1 \\ &\leq 2 \|\bar{b}_{train}\|_1 \end{aligned}$$

Therefore, we have

$$\frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} \|X_s^q \Delta\|_2^2 \leq \left| \frac{1}{n_t} \frac{\|\bar{e}_t\|_2^2}{1 - \sqrt{\bar{\varepsilon}}} - \frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} \|\bar{e}_s^q\|_2^2 \right| + \frac{1}{K} \sum_{q=1}^K \frac{4}{n_s} \left\| \left(\bar{e}_s^q \right)^T X_s^q \right\|_\infty \|\bar{b}_{train}\|_1 + \bar{\varsigma}$$

This formula is the bound for $\mathbb{E}_k \left[\mathbb{E}_{(X_s^k, Y_s^k)} \left(\|\hat{y}_s - \hat{y}_s^{lasso}\|_2^2 \right) \right]$, the iterated expected difference between lasso prediction and extremum prediction on any (X_s^k, Y_s^k) . The bound holds with probability $(1 - 1/n_t)\bar{\omega}$. □

Proof. Theorem 3. (Consistency when $n \geq p$.) Under the Newey and McFadden [1994] condition, the extremum estimators is consistent. If $n \geq p$, the extremum estimator b_{train} is simply the OLS estimator. We prove that lasso tuned by validation is consistent for $n \geq p$.

As long as $(\tilde{n}/p) \rightarrow \infty$, $\mathcal{R}_{n_t}(b_{train}|Y_t, X_t) \xrightarrow{\mathbf{P}} \inf_b \mathcal{R}(b|Y, X)$ and $\mathcal{R}_{n_s}(b_{train}|Y_s, X_s) \xrightarrow{\mathbf{P}} \inf_b \mathcal{R}(b|Y, X)$, which means $(1/n_t) \|e_t\|_2^2$ and $(1/n_s) \|e_s\|_2^2$ all converge to the same limit. As a result,

$$\frac{(1/n_t) \|e_t\|_2^2}{(1/n_s) \|e_s\|_2^2} \xrightarrow{\mathbf{P}} 1.$$

Also $(4/n_s) \|e_s^T X_s\|_\infty \xrightarrow{\mathbf{P}} 0$, $\|b_{train}\|_1 \rightarrow \|\beta\|_1$ and $\varepsilon \rightarrow 0$. Also $\hat{y} \xrightarrow{\mathbf{P}} X\beta$ if $(n/p) \rightarrow \infty$. Hence $Xb_{lasso} \xrightarrow{\mathcal{L}_2} X\beta$.

For OLS, $(1/n) \|X_s \Delta\|_2^2 \geq \rho \|\Delta\|_2^2$, where ρ is the minimal eigenvalue for $X^T X$. Hence,

$$\begin{aligned} \rho \|\Delta\|_2^2 &\leq \frac{1}{n_s} \|X_s \Delta\|_2^2 \\ &\leq \left| \frac{1}{n_t} \frac{\|e_t\|_2^2}{(1 - \sqrt{\varepsilon})} - \frac{1}{n_s} \|e_s\|_2^2 \right| + \frac{4}{n_s} \|e_s^T X_s\|_\infty \|b_{train}\|_1 + \varsigma. \end{aligned}$$

It follows that

$$\rho \|\Delta\|_2^2 \leq \left| \frac{1}{n_t} \frac{\|e_t\|_2^2}{(1-\sqrt{\varepsilon})} - \frac{1}{n_s} \|e_s\|_2^2 \right| + \frac{4}{n_s} \|e_s^T X_s\|_\infty \|b_{train}\|_1 + \zeta.$$

By the Minkowski inequality, the above can be simplified to

$$\|b_{train} - b_{lasso}\|_2 \leq \sqrt{\left| \frac{1}{\rho n_t} \frac{\|e_t\|_2^2}{(1-\sqrt{\varepsilon})} - \frac{1}{\rho n_s} \|e_s\|_2^2 \right|} + \sqrt{\frac{4}{\rho n_s} \|e_s^T X_s\|_\infty \|b_{train}\|_1} + \left(\frac{\zeta}{\rho}\right)^{\frac{1}{2}}$$

Thus, the extremum estimator and the lasso estimator asymptotically converge to β . \square

Proof. Corollary 3. (Consistency when $n \geq p$.) If $n \geq p$, extremum estimation \bar{b}_{train} is the OLS estimator for the ‘worst case’. We prove that lasso tuned by cross-validation is consistent for $n \geq p$.

For OLS, $(1/n_s) \|X_s^q \Delta\|_2^2 \geq \rho_q \|\Delta\|_2^2$, where ρ_q is the minimal eigenvalue for $(X_s^q)^T (X_s^q)$. For cross-validated lasso,

$$\frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} \|X_s^q \Delta\|_2^2 \leq \left| \frac{1}{n_t} \frac{\|\bar{e}_t\|_2^2}{(1-\sqrt{\varepsilon})} - \frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} \|\bar{e}_s^q\|_2^2 \right| + \frac{1}{K} \sum_{q=1}^K \frac{4}{n_s} \left\| \left(\bar{e}_s^q\right)^T X_s^q \right\|_\infty \|\bar{b}_{train}\|_1 + \bar{\zeta}$$

implies that

$$\frac{1}{K} \sum_{q=1}^K \rho_q \|\Delta\|_2^2 \leq \left| \frac{1}{n_t} \frac{\|\bar{e}_t\|_2^2}{(1-\sqrt{\varepsilon})} - \frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} \|\bar{e}_s^q\|_2^2 \right| + \frac{1}{K} \sum_{q=1}^K \frac{4}{n_s} \left\| \left(\bar{e}_s^q\right)^T X_s^q \right\|_\infty \|\bar{b}_{train}\|_1 + \bar{\zeta}.$$

Denoting $\min_q \rho_q$ by ρ^* ,

$$\frac{\rho^*}{K} \sum_{q=1}^K \|\Delta\|_2^2 \leq \frac{1}{K} \sum_{q=1}^K \rho_q \|\Delta\|_2^2.$$

Hence,

$$\frac{\rho^*}{K} \sum_{q=1}^K \|\Delta\|_2^2 \leq \left| \frac{1}{n_t} \frac{\|\bar{e}_t\|_2^2}{(1-\sqrt{\varepsilon})} - \frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} \|\bar{e}_s^q\|_2^2 \right| + \frac{1}{K} \sum_{q=1}^K \frac{4}{n_s} \left\| \left(\bar{e}_s^q\right)^T X_s^q \right\|_\infty \|\bar{b}_{train}\|_1 + \bar{\zeta}$$

implies

$$\begin{aligned} \frac{1}{K} \sum_{q=1}^K \|\bar{b}_{train} - b_{lasso}\|_2^2 &\leq \left| \frac{1}{n_t \cdot \rho^*} \frac{\|\bar{e}_t\|_2^2}{(1-\sqrt{\varepsilon})} - \frac{1}{K} \sum_{q=1}^K \frac{1}{n_s \cdot \rho^*} \|\bar{e}_s^q\|_2^2 \right| \\ &\quad + \frac{1}{K} \sum_{q=1}^K \frac{4}{n_s \cdot \rho^*} \left\| \left(\bar{e}_s^q\right)^T X_s^q \right\|_\infty \|\bar{b}_{train}\|_1 + \frac{\bar{\zeta}}{\rho^*}. \end{aligned} \quad (38)$$

The equation above is the bound for

$$\mathbb{E}_k \left[\mathbb{E}_{(X_s^k, Y_s^k)} \left[\|\bar{b}_{train} - b_{lasso}\|_2^2 \right] \right].$$

As $n \rightarrow \infty$, the RHS of (38) converges to zero and $\bar{b}_{train} \xrightarrow{\mathcal{L}_2} b_{lasso}$. Since \bar{b}_{train} converges to β in \mathcal{L}_2 , as guaranteed by the asymptotic property of OLS, b_{lasso} also converges to β in \mathcal{L}_2 \square

Proof. Theorem 4. (Consistency when $n < p$.) In this proof we show that lasso and FSR both converge to the true DGP if lasso is tuned by validation. For regressions where $n < p$, the OLS estimator is not feasible because $X^T X$ is not of full rank and the traditional strong convexity condition fails. As a result, $(1/n) \|X_s \Delta\|_2^2 \geq \rho \|\Delta\|_2^2$ may not hold for all Δ , the extremum estimator may not converge to β , and the consistency result established in Theorem 3 may not be valid.

To solve this problem, we import the restricted eigenvalue condition from Bickel et al. [2009] and Meinshausen and Yu [2009].¹³ The restricted eigenvalue condition assumes that $(1/n) \|X_s \Delta\|_2^2 \geq \tilde{\rho} \|\Delta\|_2^2$ still holds for all $b \in \{b_\lambda\}$ [Bickel et al., 2009] and FSR estimators b_{train} [Zhang, 2009]. Also, in this scenario, the extremum estimator $\min_b (1/n_t) \|Y_t - X_t b\|_2^2$ can be implemented by forward selection of at most n variables that minimize the training error. As shown by Tropp [2004] and Zhang [2009], forward selection regression is consistent under the restricted eigenvalue condition.

As long as $n \rightarrow \infty$, $\mathcal{R}_{emp}(b_{train}|X_t^n, Y_t^n) \xrightarrow{\mathbf{P}} \inf_b \mathcal{R}(b|Y, X)$ and $\mathcal{R}_{emp}(b_{train}|X_s^n, Y_s^n) \xrightarrow{\mathbf{P}} \inf_b \mathcal{R}(b|Y, X)$, which means $(1/n) \|e_t\|_2^2$ and $(1/n) \|e_s\|_2^2$ all converge to the same limit. Thus, $(1/n) \|e_t\|_2^2 - (1/n) \|e_s\|_2^2 \xrightarrow{\mathbf{P}} 0$. Also $(4/n) \|e_s^T X_s\|_\infty \xrightarrow{\mathbf{P}} 0$, $\|b_{train}\|_1 \xrightarrow{\mathbf{P}} \|\beta\|_1$ and $\varepsilon \rightarrow 0$. Also $\hat{y}^* \xrightarrow{\mathbf{P}} X\beta$ if $(n/p) \rightarrow \infty$. Hence, $Xb_{lasso} \xrightarrow{\mathcal{L}_2} X\beta$.

For OLS, $(1/n) \|X_s \Delta\|_2^2 \geq \tilde{\rho} \|\Delta\|_2^2$, where ρ is the minimum restricted eigenvalue for $X^T X$. Similar to Theorem 3, equation (2) in the proof of Theorem 2 can be simplified to

$$\|b_{train} - b_{lasso}\|_2 \leq \sqrt{\left| \frac{1}{\tilde{\rho} n_t} \frac{\|e_t\|_2^2}{(1 - \sqrt{\varepsilon})} - \frac{1}{\tilde{\rho} n_s} \|e_s\|_2^2 \right|} + \sqrt{\frac{4}{\tilde{\rho} n_s} \|e_s^T X_s\|_\infty \|b_{train}\|_1} + \left(\frac{\varepsilon}{\tilde{\rho}}\right)^{\frac{1}{2}}.$$

Since Tropp [2004] and Zhang [2009] prove that Forward Stagewise regression is consistent, it follows that the extremum estimator and lasso estimator asymptotically converge to β . \square

Proof. Corollary 4. (Consistency when $n < p$.) In this proof we show that, under cross-validation, a very similar bound to Theorem 4 holds for lasso as well.

As above, denoting the $\min_q \tilde{\rho}_q$ by $\tilde{\rho}^*$,

$$\frac{\tilde{\rho}^*}{K} \sum_{q=1}^K \|\Delta\|_2^2 \leq \frac{1}{K} \sum_{q=1}^K \tilde{\rho}_q \|\Delta\|_2^2.$$

Hence,

$$\frac{\tilde{\rho}^*}{K} \sum_{q=1}^K \|\Delta\|_2^2 \leq \left| \frac{1}{n_t} \frac{\|\bar{e}_t\|_2^2}{1 - \sqrt{\varepsilon}} - \frac{1}{K} \sum_{q=1}^K \frac{1}{n_s} \left\| \bar{e}_s^q \right\|_2^2 \right| + \frac{1}{K} \sum_{q=1}^K \frac{4}{n_s} \left\| \left(\bar{e}_s^q \right)^T X_s^q \right\|_\infty \|\bar{b}_{train}\|_1 + \bar{\varsigma}$$

¹³Meinshausen and Yu [2009] develop a version of the restricted eigenvalue condition, which they call the sparse eigenvalue condition.

implies

$$\begin{aligned} \frac{1}{K} \sum_{q=1}^K \|\bar{b}_{train} - b_{lasso}\|_2^2 &\leq \left| \frac{1}{n_t \cdot \tilde{\rho}^*} \frac{\|\bar{e}_t\|_2^2}{1 - \sqrt{\bar{\epsilon}}} - \frac{1}{K} \sum_{q=1}^K \frac{1}{n_s \cdot \tilde{\rho}^*} \|\bar{e}_s^q\|_2^2 \right| \\ &\quad + \frac{1}{K} \sum_{q=1}^K \frac{4}{n_s \cdot \tilde{\rho}^*} \left\| \left(\bar{e}_s^q \right)^T X_s^q \right\|_{\infty} \|\bar{b}_{train}\|_1 + \frac{\bar{\zeta}}{\tilde{\rho}^*}. \end{aligned} \quad (39)$$

The equation above is the bound for

$$\mathbb{E}_k \left[\mathbb{E}_{(X_s^k, Y_s^k)} \left[\|\bar{b}_{train} - b_{lasso}\|_2^2 \right] \right].$$

As $n \rightarrow \infty$, the RHS of (39) converges to zero and $\bar{b}_{train} \xrightarrow{\mathcal{L}_2} b_{lasso}$. Since \bar{b}_{train} converges to β in \mathcal{L}_2 [Tropp, 2004, Zhang, 2009], b_{lasso} also converges to β in \mathcal{L}_2 \square

Forward Stagewise regression algorithm

-
1. Standardize Y and the variables $X_j, j = 1, \dots, p$
 2. Start the regression from $Y = u$
 3. Add the variable having the largest correlation with u into the regression and estimate $Y = Xb + u$
 4. Repeat 3, one variable at a time, until the maximum correlation between u and the most recent variable added to the model is less than some preset value.
-

Appendix 2

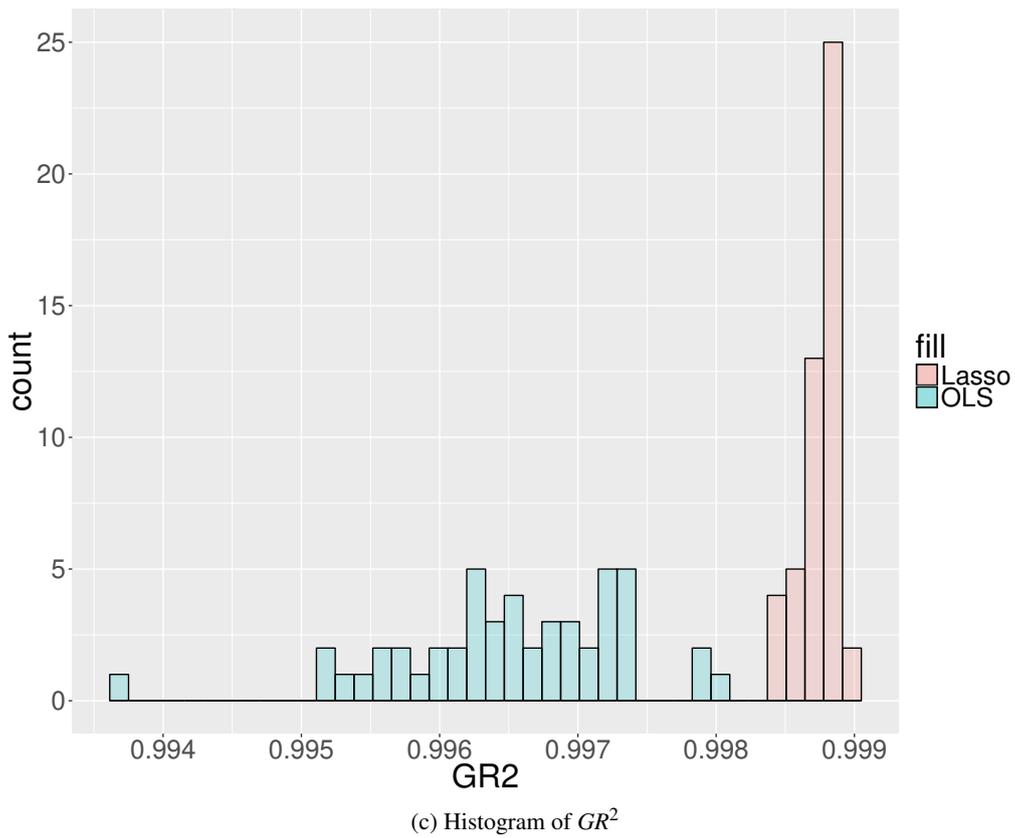
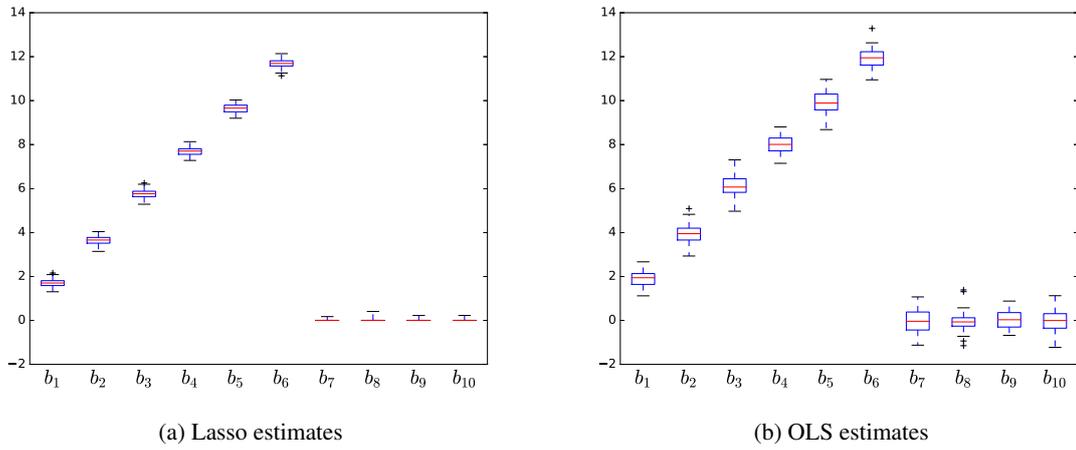


Figure 9: Boxplots of estimates and GR^2 for DGP $n = 250, p = 200, \text{var}(u) = 1$

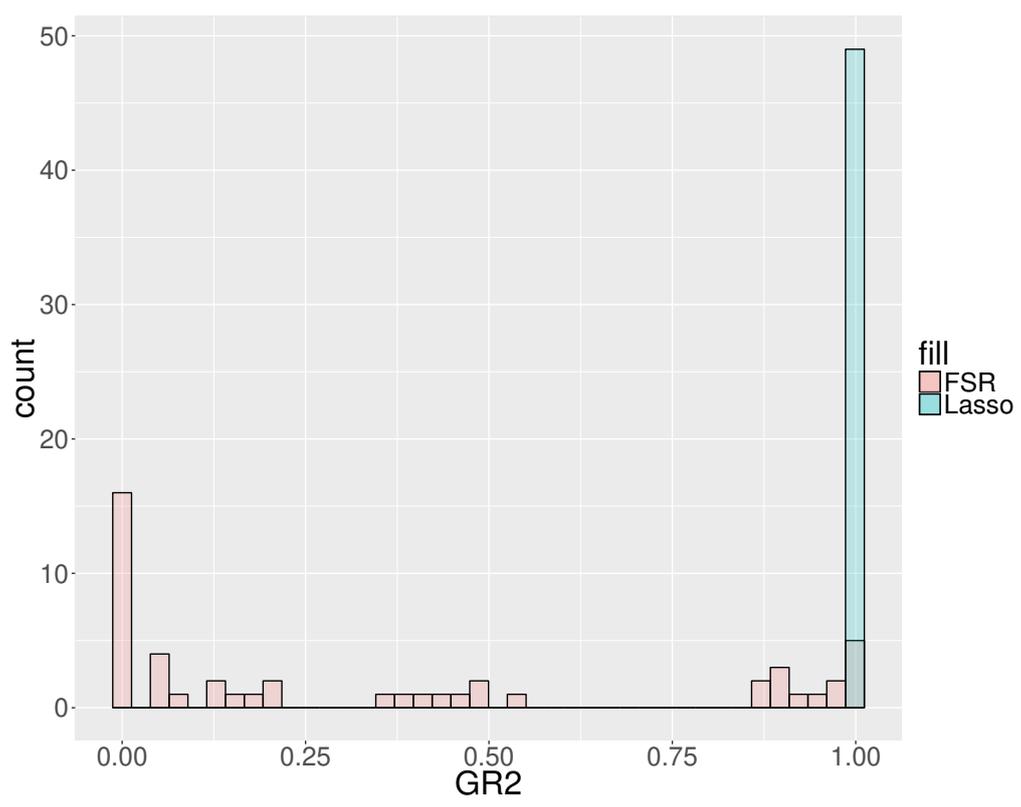
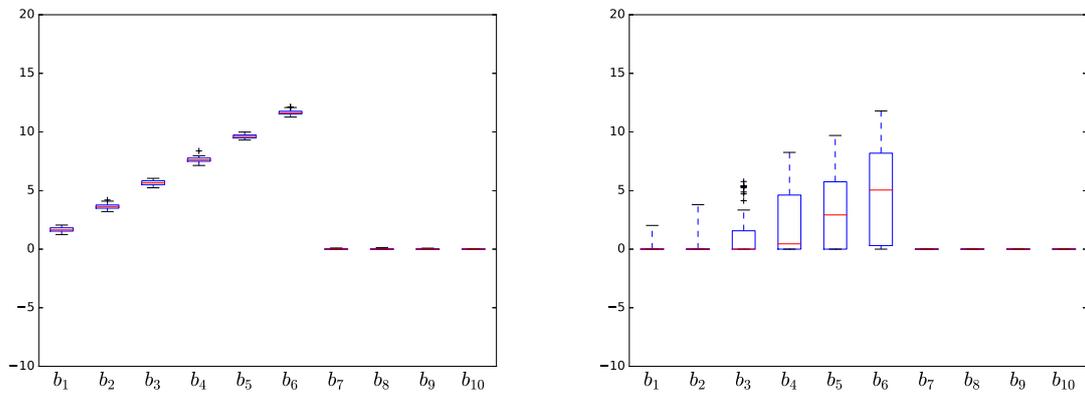
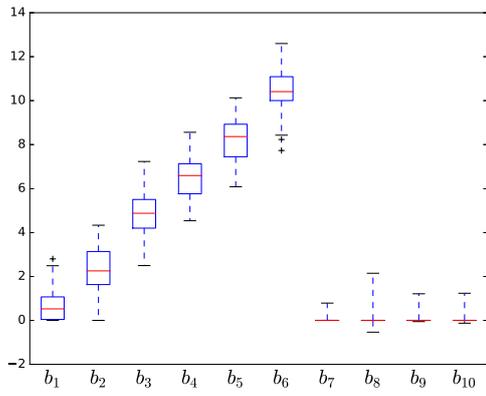
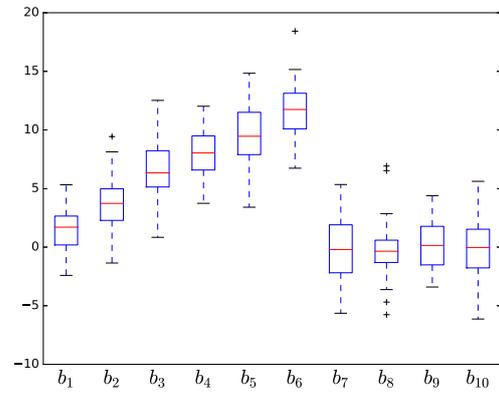


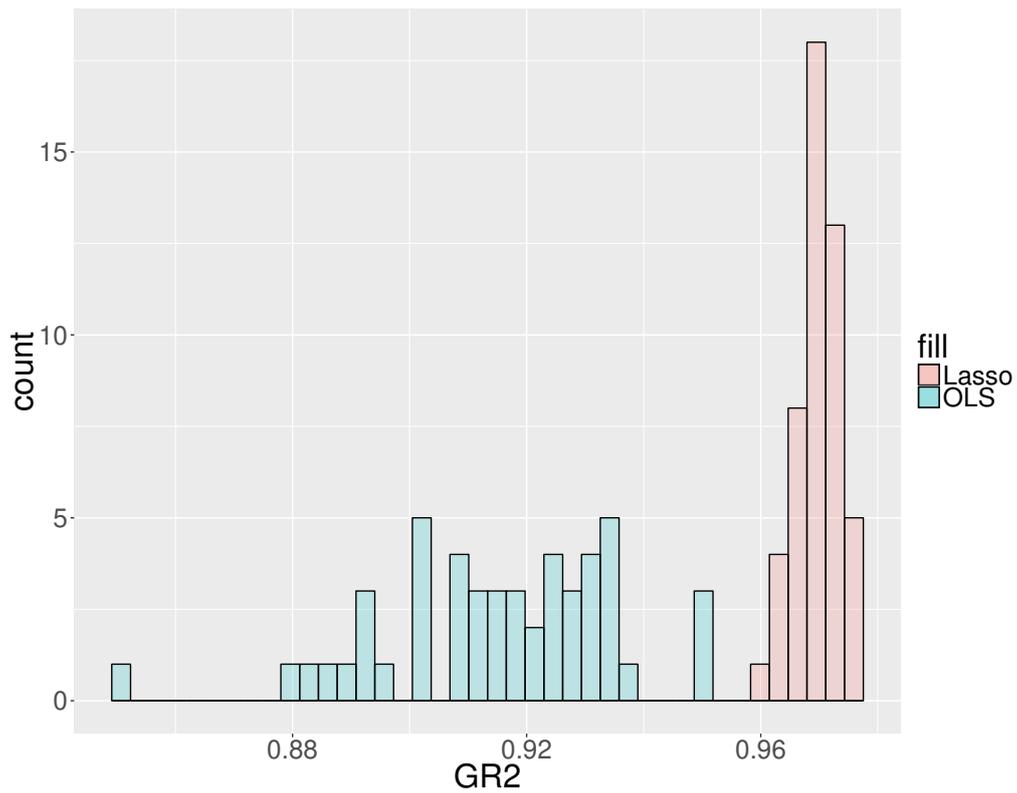
Figure 10: Boxplots of estimates and GR^2 for DGP $n = 250, p = 500, \text{var}(u) = 1$



(a) Lasso estimates

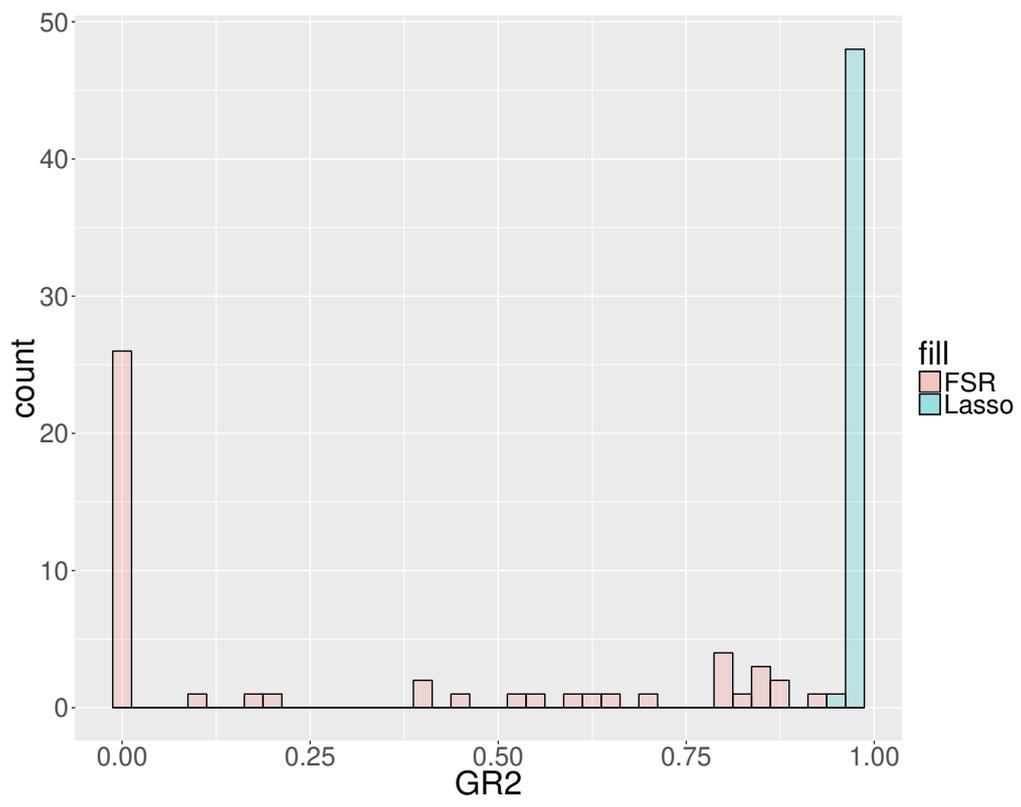
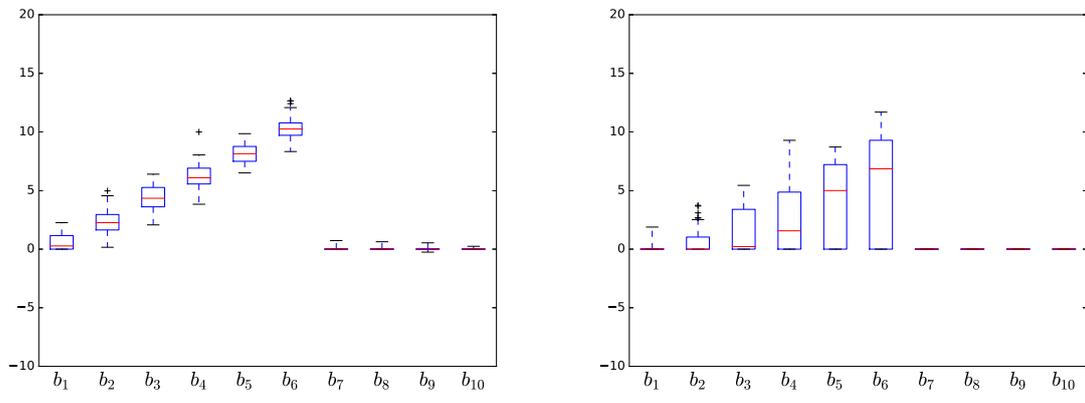


(b) OLS estimates



(c) Histogram of GR^2

Figure 11: Boxplots of estimates and GR^2 for DGP $n = 250$, $p = 200$, $\text{var}(u) = 5$



(c) Histogram of GR^2

Figure 12: Boxplots of estimates and GR^2 for DGP $n = 250, p = 500, \text{var}(u) = 5$