

CHOOSING AMONG REGULARIZED ESTIMATORS IN EMPIRICAL ECONOMICS: THE RISK OF MACHINE LEARNING

Alberto Abadie and Maximilian Kasy*

Abstract—Many settings in empirical economics involve estimation of a large number of parameters. In such settings, methods that combine regularized estimation and data-driven choices of regularization parameters are useful. We provide guidance to applied researchers on the choice between regularized estimators and data-driven selection of regularization parameters. We characterize the risk and relative performance of regularized estimators as a function of the data-generating process and show that data-driven choices of regularization parameters yield estimators with risk uniformly close to the risk attained under the optimal (unfeasible) choice of regularization parameters. We illustrate using examples from empirical economics.

I. Introduction

APPLIED economists often confront problems that require estimation of a large number of parameters. Examples include (a) estimation of causal (or predictive) effects for a large number of treatments such as neighborhoods or cities, teachers, workers and firms, or judges; (b) estimation of the causal effect of a given treatment for a large number of subgroups; and (c) prediction problems with a large number of predictive covariates or transformations of covariates. The statistics and machine learning literature provides a host of estimation methods, such as ridge, lasso, and pretest, which are particularly well adapted to high-dimensional problems. In view of the variety of available methods, applied researchers face the question of which of these procedures to adopt in any given situation. This paper provides guidance on this choice based on the risk (mean squared error) properties of a class of regularization-based machine learning methods.

A practical concern that often motivates the adoption of machine learning procedures is the potential for overfitting in high-dimensional settings. To avoid overfitting, most machine learning procedures for “supervised learning” (i.e., regression and classification methods for prediction) involve two key features: regularized estimation and data-driven choice of regularization parameters. These features are also

central to nonparametric estimation methods in econometrics, such as kernel or series regression.

Main Takeaways. For empirical practitioners, we emphasize two main takeaway messages of this paper. First, no one method for regularization is universally optimal. Different methods work well in different types of settings, and we provide guidance based on theoretical considerations, simulations, and empirical applications. Second, a choice of tuning parameters using cross-validation or Stein’s unbiased risk estimate is guaranteed to work well in high-dimensional estimation and prediction settings under mild conditions. For econometric theorists, our main contribution are our results on uniform loss consistency for a general class of regularization procedures. In seeming contrast to some results in the literature, data-driven tuning performs uniformly well in high-dimensional settings. Further interesting findings include that lasso is surprisingly robust in its performance across many settings, and flexible regularization methods such as nonparametric empirical Bayes dominate in very high-dimensional settings but are dominated by more parsimonious regularizers in more moderate dimensions.

Setup. In this paper, we consider the canonical problem of estimating the unknown means, μ_1, \dots, μ_n , of a potentially large set of observed random variables, X_1, \dots, X_n . After some transformations, our setup covers applications (a) to (c) already mentioned and many others. Our results do not cover cases where the sample size is smaller than the dimension of the parameter of interest.

We consider componentwise estimators of the form $\hat{\mu}_i = m(X_i, \lambda)$, where λ is a nonnegative regularization parameter. Typically, $m(x, 0) = x$, so that $\lambda = 0$ corresponds to the unregularized estimator $\hat{\mu}_i = X_i$. Positive values of λ typically correspond to regularized estimators, which shrink toward 0, $|\hat{\mu}_i| \leq |X_i|$. The value $\lambda = \infty$ typically implies maximal shrinkage: $\hat{\mu}_i = 0$ for $i = 1, \dots, n$. Shrinkage toward zero is a convenient normalization but is not essential. Shifting X_i by a constant to $X_i - c$, for $i = 1, \dots, n$, results in shrinkage toward c .

The Risk Function of Regularized Estimators. Our paper is structured according to the two mentioned features of

Received for publication December 4, 2017. Revision accepted for publication September 17, 2018. Editor: Yuriy Gorodnichenko.

*Abadie: MIT; Kasy: Harvard University.

We thank Gary Chamberlain, Ellora Derenoncourt, Jiaying Gu, Jérémy L’Hour, José Luis Montiel Olea, Jann Spiess, Stefan Wager, and seminar participants at several institutions for helpful comments and discussions.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/rest_a_00812.

machine learning procedures: regularization and data-driven choice of regularization parameters. We first focus on regularized estimation and study the risk properties (mean squared error, averaged across components i) of regularized estimators with fixed and with oracle-optimal regularization parameters. We show that for any given data-generating process, there is an (infeasible) risk-optimal regularized componentwise estimator. This estimator has the form of the posterior mean of μ_I given X_I , where μ_I is drawn uniformly at random from the empirical distribution of μ_1, \dots, μ_n . The optimal regularized estimator is useful to characterize the risk properties of machine learning estimators. The risk function of any regularized estimator can be expressed as a function of the distance between that regularized estimator and the optimal one.

Instead of conditioning on μ_1, \dots, μ_n , one can consider the case where each (X_i, μ_i) is a realization of a random vector (X, μ) with distribution π and a notion of risk that is integrated over the distribution of both X and μ in the population. For this alternative definition of risk, we derive results analogous to those of the previous paragraph.

We next turn to a family of parametric models for π . We consider models that allow for a probability mass at zero in the distribution of μ , corresponding to the notion of sparsity, while conditional on $\mu \neq 0$, the distribution of μ is normal around some grand mean. For these parametric models, we derive analytic risk functions, assuming oracle-optimal (risk-minimizing) choices of λ . We focus our attention on three estimators: ridge, lasso, and pretest. When the point-mass of true zeros is small, ridge tends to perform better than lasso or pretest. When there is a sizable share of true zeros, the ranking of the estimators depends on the other characteristics of the distribution of μ : (a) if the nonzero parameters are smoothly distributed in a vicinity of zero, ridge still performs best; (b) if most of the distribution of non-zero parameters assigns large probability to a set well separated from zero, pretest estimation tends to perform well; and (c) lasso tends to do comparatively well in intermediate cases that fall somewhere between (a) and (b), and overall is remarkably robust across the different specifications. This characterization of the relative performance of ridge, lasso, and pretest is consistent with the results that we obtain for the empirical applications discussed later in the paper.

Data-Driven Choice of Regularization Parameters. The second part of the article turns to the feature of machine learning estimators and studies the data-driven choice of regularization parameters. We consider choices of regularization parameters based on the minimization of a criterion function that estimates risk. Ideally, a machine learning estimator evaluated at a data-driven choice of the regularization parameter would have a risk function that is uniformly close to the risk function of the infeasible estimator using an oracle-optimal regularization parameter (which minimizes true risk). We show that this type of uniform consistency can be achieved

under fairly mild conditions whenever the dimension of the problem under consideration is large, and risk is defined as mean squared error averaged across components i . We further provide fairly weak conditions under which machine learning estimators with data-driven choices of the regularization parameter, based on Stein's unbiased risk estimate (SURE) and on cross-validation (CV), attain uniform risk consistency. In addition to allowing data-driven selection of regularization parameters, uniformly consistent estimation of the risk of shrinkage estimators can be used to select among alternative shrinkage estimators on the basis of their estimated risk in empirical settings.

Applications. We illustrate our results in the context of three applications taken from the empirical economics literature. The first application uses data from Chetty and Hendren (2018) to study the effects of locations on intergenerational earnings mobility of children. The second application uses data from the event-study analysis in Della Vigna and La Ferrara (2010), who investigate whether the stock prices of weapon-producing companies react to changes in the intensity of conflicts in countries under arms trade embargoes. The third application considers nonparametric estimation of a Mincer equation using data from the Current Population Survey (CPS), as in Belloni and Chernozhukov (2011). The presence of many neighborhoods in the first application, many weapon-producing companies in the second one, and many series regression terms in the third one makes these estimation problems high-dimensional.

These examples showcase how simple features of the data-generating process affect the relative performance of machine learning estimators. They also illustrate the way in which consistent estimation of the risk of shrinkage estimators can be used to choose regularization parameters and select among different estimators in practice. For the estimation of location effects in Chetty and Hendren (2018), we find estimates that are not overly dispersed around their mean and no evidence of sparsity. In this setting, ridge outperforms lasso and pretest in terms of estimated mean squared error. In the setting of the event-study analysis in Della Vigna and La Ferrara (2010), our results suggest that a large fraction of values of parameters are closely concentrated around zero, while a smaller but nonnegligible fraction of parameters are positive and substantially separated from zero. In this setting, pretest dominates. Similarly, to the result for the setting in Della Vigna and La Ferrara (2010), the estimation of the parameters of a Mincer equation in Belloni and Chernozhukov (2011) suggests a sparse approximation to the distribution of parameters. In this setting, however, shrinkage at the tails of the distribution is still helpful, and lasso dominates ridge and pretest.

Road Map. The rest of this paper is structured as follows. Section II introduces our setup. Section IIA discusses a series of examples from empirical economics. Section III provides

characterizations of the risk function of regularized estimators. Section IV turns to data-driven choices of regularization parameters. We show uniform risk consistency results for Stein's unbiased risk estimate and for cross-validation. Section V reports simulation results. Section VI discusses several empirical applications. Section VII concludes. An online appendix contains proofs and supplemental materials, including a review of the substantial literature on statistical decision theory and machine learning this article builds on.

II. Setup

Throughout this paper, we consider the following setting. We observe a realization of an n -vector of real-valued random variables, $\mathbf{X} = (X_1, \dots, X_n)'$, where the components of \mathbf{X} are mutually independent with finite mean μ_i and finite variance σ_i^2 , for $i = 1, \dots, n$. Our goal is to estimate μ_1, \dots, μ_n .

In many applications, the X_i arise as preliminary least squares estimates of the coefficients of interest, μ_i . Consider, for instance, a randomized controlled trial where randomization of treatment assignment is carried out separately for n nonoverlapping subgroups. Within each subgroup, the difference in the sample averages between treated and control units, X_i , has mean equal to the average treatment effect for that group in the population, μ_i . Further examples are discussed in section IIA.

Componentwise estimators. We restrict our attention to componentwise estimators of μ_i ,

$$\hat{\mu}_i = m(X_i, \lambda),$$

where $m : \mathbb{R} \times [0, \infty] \rightarrow \mathbb{R}$ defines an estimator of μ_i as a function of X_i and a nonnegative regularization parameter, λ . The parameter λ is common across the components i and might depend on the vector \mathbf{X} . We study data-driven choices $\hat{\lambda}$ in section IV, focusing in particular on Stein's unbiased risk estimate (SURE) and cross-validation (CV).

Popular choices for $m(x, \lambda)$ are ridge, lasso, and pretest. They are defined as follows:


$$\text{Ridge: } m_R(x, \lambda) = \underset{m \in \mathbb{R}}{\operatorname{argmin}} (x - m)^2 + \lambda m^2 = \frac{1}{1 + \lambda} x$$

$$\text{Lasso: } m_L(x, \lambda) = \underset{m \in \mathbb{R}}{\operatorname{argmin}} (x - m)^2 + 2\lambda|m| = 1(x < -\lambda) \\ (x + \lambda) + 1(x > \lambda)(x - \lambda)$$

$$\text{Pretest: } m_{PT}(x, \lambda) = \underset{m \in \mathbb{R}}{\operatorname{argmin}} (x - m)^2 + \lambda^2 1(m \neq 0) = \\ 1(|x| > \lambda)x,$$

where $1(A)$ denotes the indicator function, which equals 1 if A holds and 0 otherwise. Figure 1 plots $m_R(x, \lambda)$, $m_L(x, \lambda)$ and $m_{PT}(x, \lambda)$ as functions of x . For reasons apparent in figure 1, ridge, lasso, and pretest estimators are sometimes referred to as linear shrinkage, soft thresholding, and hard thresh-

FIGURE 1.—SHRINKAGE ESTIMATORS



This graph plots $m_R(x, \lambda)$, $m_L(x, \lambda)$, and $m_{PT}(x, \lambda)$ as functions of x . The regularization parameters are $\lambda = 1$ for ridge, $\lambda = 2$ for lasso, and $\lambda = 4$ for pretest.

olding, respectively.¹ The functions $m_R(x, \lambda)$, $m_L(x, \lambda)$ and $m_{PT}(x, \lambda)$ all fall between the 45 degree line and the flat line at zero. That is, they produce shrinkage toward zero. As we discuss below, the problem of determining the optimal choice among these estimators in terms of minimizing mean squared error is equivalent to the problem of determining which of these estimators best approximates a certain optimal estimation function.

Let $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n)'$ and $\hat{\boldsymbol{\mu}} = (\hat{\mu}_1, \dots, \hat{\mu}_n)'$, where, for simplicity, we leave the dependence of $\hat{\boldsymbol{\mu}}$ on λ implicit in our notation. Let P_1, \dots, P_n be the distributions of X_1, \dots, X_n , and let $\mathbf{P} = (P_1, \dots, P_n)$.

Loss and risk. We evaluate estimates based on the squared error loss function, or compound loss,

$$L_n(\mathbf{X}, m(\cdot, \lambda), \mathbf{P}) = \frac{1}{n} \sum_{i=1}^n (m(X_i, \lambda) - \mu_i)^2,$$

where L_n depends on \mathbf{P} via $\boldsymbol{\mu}$. We write $m(\cdot, \lambda)$ here and below to emphasize that we are evaluating the function m , mapping X_i into $\hat{\mu}_i$, for fixed λ . We will use expected loss to rank estimators. There are different ways of taking this expectation, resulting in different risk functions, and the distinction between them is conceptually important.

Componentwise risk fixes P_i and considers the expected squared error of $\hat{\mu}_i$ as an estimator of μ_i :

$$R(m(\cdot, \lambda), P_i) = E[(m(X_i, \lambda) - \mu_i)^2 | P_i].$$

The expectation on the right-hand side averages over the sampling distribution of X_i , given P_i . We use Bayesian

¹Pretest is also known as Hodge's estimator, in settings where $\lambda \rightarrow 0$ at a rate slower than σ_i . See Leeb and Pötscher (2006).

notation, conditioning on the parameter P_i , to make this explicit.

Compound risk averages componentwise risk over the empirical distribution of P_i across the components $i = 1, \dots, n$. Compound risk is given by the expectation of compound loss L_n given \mathbf{P} :

$$\begin{aligned} R_n(m(\cdot, \lambda), \mathbf{P}) &= E[L_n(\mathbf{X}, m(\cdot, \lambda), \mathbf{P}) | \mathbf{P}] \\ &= \frac{1}{n} \sum_{i=1}^n E[(m(X_i, \lambda) - \mu_i)^2 | P_i] \\ &= \frac{1}{n} \sum_{i=1}^n R(m(\cdot, \lambda), P_i). \end{aligned}$$

Finally, *integrated* (or *empirical Bayes*) risk considers P_1, \dots, P_n to be themselves draws from some population distribution, Π . This induces a joint distribution, π , for (X_i, μ_i) . Throughout the paper, we often use a subscript π to denote characteristics of the joint distribution of (X_i, μ_i) . Integrated risk refers to loss integrated over π or, equivalently, componentwise risk integrated over Π :

$$\begin{aligned} \bar{R}(m(\cdot, \lambda), \pi) &= E_\pi[L_n(\mathbf{X}, m(\cdot, \lambda), \mathbf{P})] \\ &= E_\pi[(m(X_i, \lambda) - \mu_i)^2] \\ &= \int R(m(\cdot, \lambda), P_i) d\Pi(P_i). \end{aligned} \quad (1)$$

Notice the similarity between compound risk and integrated risk: they differ only by replacing an empirical (sample) distribution by a population distribution. For large n , the difference between the two vanishes, as we will demonstrate in section IV.

Regularization parameter. Throughout, we use $R_n(m(\cdot, \lambda), \mathbf{P})$ to denote the risk function of the estimator $m(\cdot, \lambda)$ with fixed (nonrandom) λ , and similarly for $\bar{R}(m(\cdot, \lambda), \pi)$. In contrast, $R_n(m(\cdot, \hat{\lambda}_n), \mathbf{P})$ is the risk function taking into account the randomness of $\hat{\lambda}_n$, where the latter is chosen in a data-dependent manner, and similarly for $\bar{R}(m(\cdot, \hat{\lambda}_n), \pi)$.

For a given \mathbf{P} , we define the “oracle” selector of the regularization parameter as the value of λ that minimizes compound risk,

$$\lambda^*(\mathbf{P}) = \operatorname{argmin}_{\lambda \in [0, \infty]} R_n(m(\cdot, \lambda), \mathbf{P}),$$

whenever the argmin exists. We use $\lambda_R^*(\mathbf{P})$, $\lambda_L^*(\mathbf{P})$ and $\lambda_{PT}^*(\mathbf{P})$ to denote the oracle selectors for ridge, lasso, and pretest, respectively. Analogously, for a given π , we define

$$\bar{\lambda}^*(\pi) = \operatorname{argmin}_{\lambda \in [0, \infty]} \bar{R}(m(\cdot, \lambda), \pi) \quad (2)$$

whenever the argmin exists, with $\bar{\lambda}_R^*(\pi)$, $\bar{\lambda}_L^*(\pi)$, and $\bar{\lambda}_{PT}^*(\pi)$ for ridge, lasso, and pretest, respectively. In section III, we

characterize compound and integrated risk for fixed λ and for the oracle-optimal λ . In section IV, we show that data-driven choices $\hat{\lambda}_n$ are, under certain conditions, as good as the oracle-optimal choice, in a sense to be made precise.

A. Empirical Examples

Our setup describes a variety of settings often encountered in empirical economics, where X_1, \dots, X_n are unbiased or close to unbiased but noisy least squares estimates of a set of parameters of interest, μ_1, \dots, μ_n . As mentioned in section I, examples include (a) studies estimating causal or predictive effects for a large number of treatments such as neighborhoods, cities, teachers, workers, firms, or judges; (b) studies estimating the causal effect of a given treatment for a large number of subgroups; and (c) prediction problems with a large number of predictive covariates or transformations of covariates.

Large number of treatments. Examples in the first category include Chetty and Hendren (2018), who estimate the effect of geographic locations on intergenerational mobility for a large number of locations. Chetty and Hendren use differences between the outcomes of siblings whose parents move during their childhood in order to identify these effects. The problem of estimating a large number of parameters also arises in the teacher value-added literature when the objects of interest are individual teachers’ effects (see, e.g., Chetty, Friedman, & Rockoff, 2014). In labor economics, estimation of firm and worker effects in studies of wage inequality has been considered in Abowd, Kramarz, and Margolis (1999). Another example within the first category is provided by Abrams, Bertrand, and Mullainathan (2012), who estimate differences in the effects of defendant’s race on sentencing across individual judges.

Treatment for large numbers of subgroups. Within the second category, which consists of estimating the effect of a treatment for many subpopulations, our setup can be applied to the estimation of heterogeneous causal effects of class size on student outcomes across many subgroups. For instance, Project STAR (Krueger, 1999) involved experimental assignment of students to classes of different sizes in 79 schools. Causal effects for many subgroups are also of interest in medical contexts or for active labor market programs, where doctors or policymakers have to decide on treatment assignment based on individual characteristics. In some empirical settings, treatment impacts are individually estimated for each sample unit. This is often the case in empirical finance, where event studies are used to estimate reactions of stock market prices to newly available information. For example, Della Vigna and La Ferrara (2010) estimate the effects of changes in the intensity of armed conflicts in countries under arms trade embargoes on the stock market prices of arms-manufacturing companies.

Prediction with many regressors. The third category is prediction with many regressors (but no more than the number of observations). This category fits in the setting of this paper after orthonormalization of the regressors, such that their sample second-moment matrix is the identity. Prediction with many regressors arises, in particular, in macroeconomic forecasting. Stock and Watson (2012), in an analysis complementing this paper, evaluate various procedures in terms of their forecast performance for a number of macroeconomic time series for the United States. Regression with many predictors also arises in series regression, where series terms are transformations of a set of predictors. Series regression and its asymptotic properties have been widely studied in econometrics (see, e.g., Newey, 1997). Wasserman (2006, sections 7.2–7.3) provides an illuminating discussion of the equivalence between the many means model studied in this paper and nonparametric regression estimation. For that setting, X_1, \dots, X_n and μ_1, \dots, μ_n correspond to the estimated and true regression coefficients on an orthogonal basis of functions. Application of lasso and pretesting to series regression is discussed, for instance, in Belloni and Chernozhukov (2011). Appendix A.1 further discusses the relationship between the many means model and prediction models.

In section VI, we return to three of these applications, revisiting the estimation of location effects on intergenerational mobility, as in Chetty and Hendren (2018), the effect of changes in the intensity of conflicts in arms embargo countries on the stock prices of arms manufacturers, as in Della Vigna and La Ferrara (2010), and nonparametric series estimation of a Mincer equation, as in Belloni and Chernozhukov (2011).

III. The Risk Function

We now turn to our first set of formal results, which pertain to the mean squared error of regularized estimators. Our goal is to guide the researcher's choice of estimator by describing the conditions under which each of the alternative machine learning estimators performs better than the others.

We first derive a general characterization of the mean squared error of regularized estimators. This characterization is based on the geometry of estimating functions m as depicted in figure 1. It is a priori not obvious which of these functions is best suited for estimation. We show that for any given data-generating process, there is an optimal function m_P^* that minimizes mean squared error. Moreover, we show that the mean squared error for an arbitrary m is equal, up to a constant, to the L^2 distance between m and m_P^* . A function m thus yields a good estimator if it is able to approximate the shape of m_P^* well.

In section IIIB, we discuss analytic characterizations for the componentwise risk of ridge, lasso, and pretest estimators, imposing the additional assumption of normality. Summing or integrating componentwise risk over some distribution for (μ_i, σ_i) delivers expressions for compound and integrated risk.

In section IIIC, we turn to a specific parametric family of data-generating processes where each μ_i is equal to zero with probability p , reflecting the notion of sparsity, and is otherwise drawn from a normal distribution with some mean μ_0 and variance σ_0^2 . For this parametric family indexed by (p, μ_0, σ_0) , we discuss analytic risk functions and visual comparisons of the relative performance of alternative estimators. This allows us to identify key features of the data-generating process that affect the relative performance of alternative estimators.

A. General Characterization

Recall the setup introduced in section II, where we observe n jointly independent random variables X_1, \dots, X_n , with means μ_1, \dots, μ_n . We are interested in the mean squared error for the compound problem of estimating all μ_1, \dots, μ_n simultaneously. In this formulation of the problem, μ_1, \dots, μ_n are fixed unknown parameters.

Let I be a random variable with a uniform distribution over the set $\{1, 2, \dots, n\}$ and consider the random component (X_I, μ_I) of (X, μ) . This construction induces a mixture distribution for (X_I, μ_I) (conditional on P),

$$(X_I, \mu_I)|P \sim \frac{1}{n} \sum_{i=1}^n P_i \delta_{\mu_i},$$

where $\delta_{\mu_1}, \dots, \delta_{\mu_n}$ are Dirac measures at μ_1, \dots, μ_n . Based on this mixture distribution, define the conditional expectation and the average conditional variance:

$$m_P^*(x) = E[\mu_I | X_I = x, P] \quad \text{and} \quad v_P^* = E[\text{var}(\mu_I | X_I, P) | P].$$

Theorem 1 characterizes the compound risk of an estimator in terms of the average squared discrepancy relative to m_P^* , which implies that m_P^* is optimal (lowest mean squared error) for the compound problem.

Theorem 1 (*characterization of risk functions*). *Consider the many means model of section II, where $X_I|P$ has distribution P_i with expectation μ_i . If $\sup_{\lambda \in [0, \infty]} E[(m(X_I, \lambda))^2 | P] < \infty$, the compound risk function R_n of $\hat{\mu}_i = m(X_i, \lambda)$ can be written as*

$$R_n(m(\cdot, \lambda), P) = v_P^* + E[(m(X_I, \lambda) - m_P^*(X_I))^2 | P],$$

which implies

$$\lambda^*(P) = \underset{\lambda \in [0, \infty]}{\operatorname{argmin}} E[(m(X_I, \lambda) - m_P^*(X_I))^2 | P]$$

whenever $\lambda^*(P)$ is well defined.

The proof of this theorem and all further results are in the online appendix.

The statement of this theorem implies that the risk of componentwise estimators is equal to an irreducible part v_P^* , plus the L^2 distance of the estimating function $m(\cdot, \lambda)$ to

the infeasible optimal estimating function m_p^* . A given data-generating process \mathbf{P} maps into an optimal estimating function m_p^* , and the relative performance of alternative estimators m depends on how well they approximate m_p^* .

We can easily write m_p^* explicitly because the conditional expectation defining m_p^* is a weighted average of the values taken by μ_i . Suppose, for example, that $X_i \sim N(\mu_i, 1)$ for $i = 1 \dots n$. Let ϕ be the standard normal probability density function. Then,

$$m_p^*(x) = \frac{\sum_{i=1}^n \mu_i \phi(x - \mu_i)}{\sum_{i=1}^n \phi(x - \mu_i)}.$$

Theorem 1 conditions on the empirical distribution of μ_1, \dots, μ_n , which corresponds to the notion of compound risk. Replacing this empirical distribution by the population distribution π , so that

$$(X_i, \mu_i) \sim \pi,$$

results analogous to those in theorem 1 are obtained for the integrated risk and the integrated oracle selectors in equations (1) and (2). That is, let

$$\bar{m}_\pi^*(x) = E_\pi[\mu_i | X_i = x] \quad \text{and} \quad \bar{v}_\pi^* = E_\pi[\text{var}_\pi(\mu_i | X_i)],$$

and assume $\sup_{\lambda \in [0, \infty]} E_\pi[(m(X_i, \lambda) - \mu_i)^2] < \infty$. Then

$$\bar{R}(m(\cdot, \lambda), \pi) = \bar{v}_\pi^* + E_\pi[(m(X_i, \lambda) - \bar{m}_\pi^*(X_i))^2],$$

and

$$\bar{\lambda}^*(\pi) = \underset{\lambda \in [0, \infty]}{\text{argmin}} E_\pi[(m(X_i, \lambda) - \bar{m}_\pi^*(X_i))^2]. \quad (3)$$

The proof of these assertions is analogous to the proof of theorem 1. m_p^* and \bar{m}_π^* are optimal componentwise estimators or “shrinkage functions” in the sense that they minimize the compound and integrated risk, respectively.


B. Componentwise Risk

The characterization of the risk of componentwise estimators in the previous section relies only on the existence of second moments. Explicit expressions for compound risk and integrated risk can be derived under additional structure. We now consider a setting in which the X_i are normally distributed:

$$X_i \sim N(\mu_i, \sigma_i^2).$$

This is a particularly relevant scenario in applied research, where the X_i are often unbiased estimators with a normal distribution in large samples, as in examples (a) to (c) in sections I and IIA. For concreteness, we focus on the three widely used componentwise estimators introduced in section

FIGURE 2.—COMPONENTWISE RISK FUNCTIONS



This figure displays componentwise risk, $R(m(\cdot, \lambda), P_i)$, as a function of μ_i for componentwise estimators, where $\sigma_i^2 = 2$. “mle” refers to the maximum likelihood (unregularized) estimator, $\hat{\mu}_i = X_i$, which has risk equal to $\sigma_i^2 = 2$. The regularization parameters are $\lambda = 1$ for ridge, $\lambda = 2$ for lasso, and $\lambda = 4$ for pretest, as in figure 1.

II—ridge, lasso, and pretest—whose estimating functions m were plotted in figure 1. Lemma A.1 in the appendix provides explicit expressions for the componentwise risk of these estimators.

Figure 2 plots the componentwise risk functions in lemma A.1 as functions of μ_i (with $\lambda = 1$ for ridge, $\lambda = 2$ for lasso, and $\lambda = 4$ for pretest). It also plots the componentwise risk of the unregularized maximum likelihood estimator, $\hat{\mu}_i = X_i$, which is equal to σ_i^2 . As figure 2 suggests, componentwise risk is large for ridge when $|\mu_i|$ is large. The same is true for lasso, except that risk remains bounded. For pretest, componentwise risk is large when $|\mu_i|$ is close to λ .

Notice that these functions are plotted for a fixed value of the regularization parameter. If λ is chosen optimally, then the componentwise risks of ridge, lasso, and pretest are no greater than the componentwise risk of the unregularized maximum likelihood estimator $\hat{\mu}_i = X_i$, which is σ_i^2 . The reason is that ridge, lasso, and pretest nest the unregularized estimator (as in the case $\lambda = 0$).

C. Spike and Normal Data-Generating Process

If we take the expressions for componentwise risk derived in lemma A.1 and average them over some population distribution of (μ_i, σ_i^2) , we obtain the integrated, or empirical Bayes, risk. For parametric families of distributions of (μ_i, σ_i^2) , this might be done analytically. We do so now, considering a family of distributions that is rich enough to cover common intuitions about data-generating processes but simple enough to allow for analytic expressions. Based on these expressions, we characterize scenarios that favor the relative

performance of each of the estimators considered in this paper.

We consider a family of distributions for (μ_i, σ_i) such that (a) μ_i takes value 0 with probability p and is otherwise distributed as a normal with mean value μ_0 and standard deviation σ_0 , and (b) $\sigma_i^2 = \sigma^2$. Proposition A.1 in the appendix derives the optimal estimating function \bar{m}_π^* , as well as integrated risk functions of ridge, lasso, and pretest for this family of distributions.

Even under substantial sparsity (i.e., if p is large), the optimal shrinkage function, \bar{m}_π^* , never shrinks all the way to zero (unless, of course, $\mu_0 = \sigma_0 = 0$ or $p = 1$). This could in principle cast some doubts about the appropriateness of thresholding estimators, such as lasso or pretest, which induce sparsity in the estimated parameters. However, as we will see below, despite this stark difference between thresholding estimators and \bar{m}_π^* , lasso and, to a certain extent, pretest are able to approximate the integrated risk of \bar{m}_π^* in the spike and normal model when the degree of sparsity in the parameters of interest is substantial.

Visual representations. Figure 3 plots the minimal integrated risk function of the different estimators. Each of the four subplots in figure 3 is based on a fixed value of $p \in \{0, 0.25, 0.5, 0.75\}$, with μ_0 and σ_0^2 varying along the bottom axes, and $\sigma^2 = 1$. For each value of the triple (p, μ_0, σ_0) , the first three panels of figure 3 report minimal integrated risk of each shrinkage estimator (ridge, lasso, and pretest) minimized over $\lambda \in [0, \infty]$. As a benchmark, the fourth panel of figure 3 reports the risk of the optimal shrinkage function, \bar{m}_π^* , simulated over 10 million repetitions. Figure 4 maps the regions of parameter values over which each of the three estimators (ridge, lasso, or pretest) performs best in terms of integrated risk.

Figures 3 and 4 provide some useful insights into the performance of shrinkage estimators. With no true zeros, ridge performs better than lasso or pretest. A clear advantage of ridge in this setting is that in contrast to lasso or pretest, ridge allows shrinkage without shrinking some observations all the way to zero. As the share of true zeros increases, the relative performance of ridge deteriorates for pairs (μ_0, σ_0) away from the origin. Intuitively, linear shrinkage imposes a disadvantageous trade-off on ridge. Using ridge to heavily shrink toward the origin in order to fit potential true zeros produces large expected errors for observations with μ_i away from the origin. As a result, ridge performance suffers considerably unless much of the probability mass of the distribution of μ_i is tightly concentrated around zero. In the absence of true zeros, pretest performs particularly poorly unless the distribution of μ_i has much of its probability mass tightly concentrated around zero, in which case, shrinking all the way to zero produces low risk. However, in the presence of true zeros, pretest performs well when much of the probability mass of the distribution of μ_i is located in a set that is well separated from zero, which facilitates the detection of true

zeros. Intermediate values of μ_0 coupled with moderate values of σ_0 produce settings where the conditional distributions $X_i|\mu_i = 0$ and $X_i|\mu_i \neq 0$ greatly overlap, inducing substantial risk for pretest estimation. The risk performance of lasso is particularly robust. It outperforms ridge and pretest for values of (μ_0, σ_0) at intermediate distances to the origin and uniformly controls risk over the parameter space. This robustness of lasso may explain its popularity in empirical practice. Despite the fact that, unlike optimal shrinkage, thresholding estimators impose sparsity, lasso—and, to a certain extent, pretest—are able to approximate the integrated risk of the optimal shrinkage function over much of the parameter space.

All in all, the results in figures 3 and 4 for the spike and normal case support the adoption of ridge in empirical applications where there are no reasons to presume the presence of many true zeros among the parameters of interest. In empirical settings where many true zeros may be expected, figures 3 and 4 show that the choice among estimators in the spike and normal model depends on how well separated the distributions $X_i|\mu_i = 0$ and $X_i|\mu_i \neq 0$ are. Pretest is preferred in the well-separated case, while lasso is preferred in the non-separated case.

IV. Data-Driven Choice of Regularization Parameters

In section IIIC, we adopted a parametric model for the distribution of μ_i to study the risk properties of regularized estimators under an oracle choice of the regularization parameter, $\bar{\lambda}^*(\pi)$. In this section, we return to a nonparametric setting and show that it is possible to consistently estimate $\bar{\lambda}^*(\pi)$ from the data, X_1, \dots, X_n , under some regularity conditions on π . We consider estimates $\hat{\lambda}_n$ of $\bar{\lambda}^*(\pi)$ based on Stein's unbiased risk estimate and based on cross-validation. The resulting estimators $m(X_i, \hat{\lambda}_n)$ have risk functions that uniformly close to those of the infeasible estimators $m(X_i, \bar{\lambda}^*(\pi))$. The asymptotic sequences we consider assume that (X_i, μ_i) are i.i.d. draws from the distribution π , and uniformity is over all distributions π with bounded fourth moments.

The uniformity part of this statement is important and not obvious. Absent uniformity, asymptotic approximations might misleadingly suggest good behavior, while in fact, the finite sample behavior of proposed estimators might be quite poor for plausible sets of data-generating processes. Notice also that our definition of compound loss averages (rather than sums) component-wise loss. For large n , any given component i thus contributes little to compound loss or risk, and uniform risk consistency is to be understood accordingly.

A. Uniform Loss and Risk Consistency

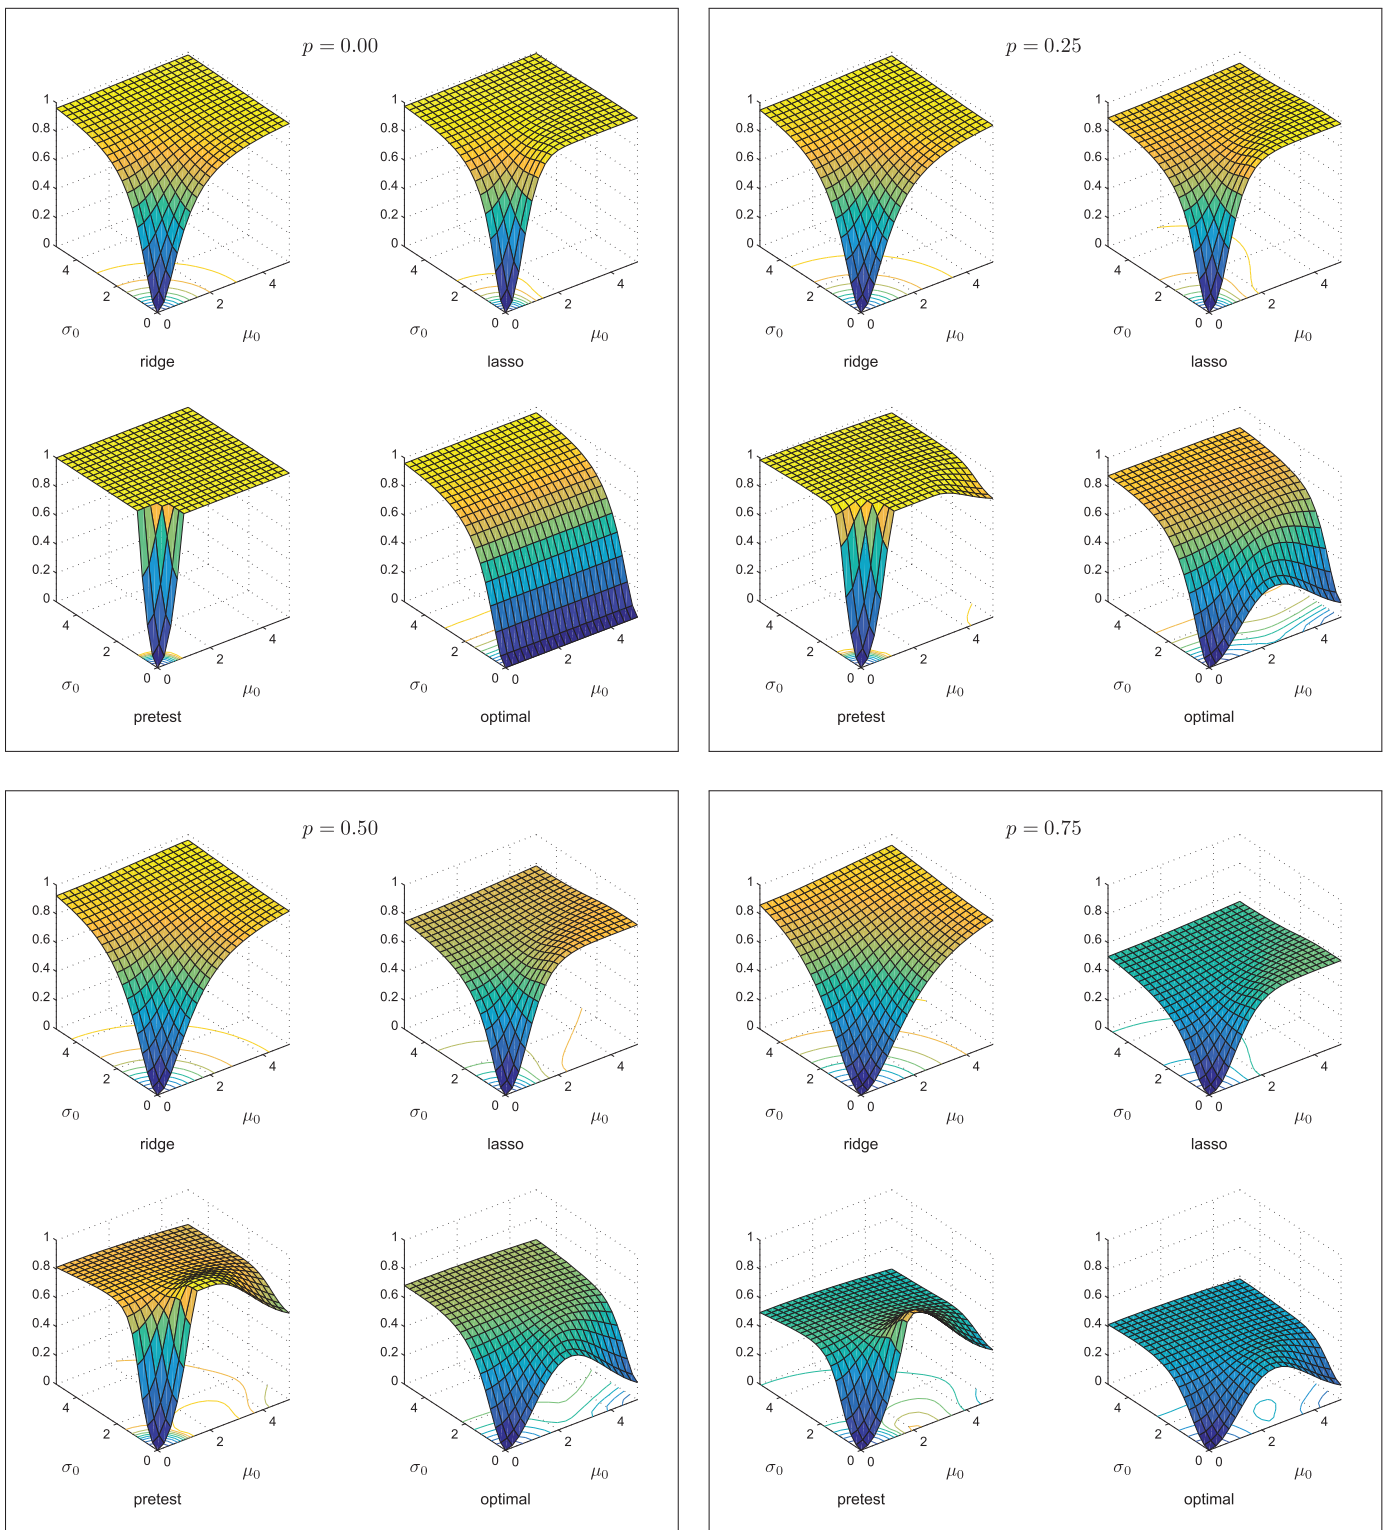
For the remainder of the paper, we adopt the following shorthand notation:

Compound loss: $L_n(\lambda) = L_n(X, m(\cdot, \lambda), \mathbf{P})$

Compound risk: $R_n(\lambda) = R_n(m(\cdot, \lambda), \mathbf{P})$

Empirical Bayes or Integrated Risk: $\bar{R}_\pi(\lambda) = \bar{R}(m(\cdot, \lambda), \pi)$

FIGURE 3.—RISK FOR ESTIMATORS IN SPIKE AND NORMAL SETTING




A color version of this figure is available online at https://doi.org/10.1162/rest_a_00812.

We now consider estimators $\hat{\lambda}_n$ of $\bar{\lambda}^*(\pi)$ that are obtained by minimizing some empirical estimate of the risk function \bar{R}_π (possibly up to a constant that depends only on π). The resulting $\hat{\lambda}_n$ is then used to obtain regularized estimators of

the form $\hat{\mu}_i = m(X_i, \hat{\lambda}_n)$. We will show that for large n , the compound loss, the compound risk, and the integrated risk functions of the resulting estimators are uniformly close to the corresponding functions of the same estimators evaluated

FIGURE 4.—BEST ESTIMATOR IN SPIKE AND NORMAL SETTING



This figure compares integrated risk values attained by ridge, lasso, and pretest for different parameter values of the spike and normal specification in section IIIC. Circles (blue in the digital version) are placed at parameter values for which ridge minimizes integrated risk, (green) crosses at values for which lasso minimizes integrated risk, and (red) dots are parameter values for which pretest minimizes integrated risk. A color version of this figure is available online at https://doi.org/10.1162/rest_a_00812.

at oracle-optimal values of λ . As $n \rightarrow \infty$, the differences between L_n , R_n , and \bar{R}_π vanish, so compound loss optimality, compound risk optimality, and integrated risk optimality become equivalent.

Let \mathcal{Q} be a set of probability distributions for (X_i, μ_i) . The following theorem establishes our key result for this section.

Theorem 2 (uniform loss consistency). Assume

$$\sup_{\pi \in \mathcal{Q}} P_\pi \left(\sup_{\lambda \in [0, \infty]} |L_n(\lambda) - \bar{R}_\pi(\lambda)| > \epsilon \right) \rightarrow 0, \quad \forall \epsilon > 0. \quad (4)$$

Assume also that there are functions, $\bar{r}_\pi(\lambda)$, \bar{v}_π , and $r_n(\lambda)$ (of (π, λ) , π , and $(\{X_i\}_{i=1}^n, \lambda)$, respectively) such that $\bar{R}_\pi(\lambda) = \bar{r}_\pi(\lambda) + \bar{v}_\pi$, and

$$\sup_{\pi \in \mathcal{Q}} P_\pi \left(\sup_{\lambda \in [0, \infty]} |r_n(\lambda) - \bar{r}_\pi(\lambda)| > \epsilon \right) \rightarrow 0, \quad \forall \epsilon > 0. \quad (5)$$

Then,

$$\sup_{\pi \in \mathcal{Q}} P_\pi \left(\left| L_n(\hat{\lambda}_n) - \inf_{\lambda \in [0, \infty]} L_n(\lambda) \right| > \epsilon \right) \rightarrow 0, \quad \forall \epsilon > 0,$$

where $\hat{\lambda}_n = \operatorname{argmin}_{\lambda \in [0, \infty]} r_n(\lambda)$.

Theorem 2 provides sufficient conditions for uniform loss consistency over $\pi \in \mathcal{Q}$, namely, that (a) the supremum of the difference between the loss, $L_n(\lambda)$, and the empirical Bayes risk, $\bar{R}_\pi(\lambda)$, vanishes in probability uniformly over $\pi \in \mathcal{Q}$ and (b) that $\hat{\lambda}_n$ is chosen to minimize a uniformly consistent estimator, $r_n(\lambda)$, of the risk function, $\bar{R}_\pi(\lambda)$ (possibly up to a constant \bar{v}_π). Under these conditions, the difference between loss $L_n(\hat{\lambda}_n)$ and the infeasible minimal loss $\inf_{\lambda \in [0, \infty]} L_n(\lambda)$ vanishes in probability uniformly over $\pi \in \mathcal{Q}$.

The sufficient conditions given by this theorem, as stated in equations (4) and (5), are rather high level. We now give more primitive conditions for these requirements to hold. In sections IVB and IVC, we propose suitable choices of $r_n(\lambda)$ based on Stein's unbiased risk estimator (SURE) and cross-validation (CV) and show that equation (5) holds for these choices of $r_n(\lambda)$.

Theorem 3 provides a set of conditions under which equation (4) holds, so the difference between compound loss and integrated risk vanishes uniformly. Aside from a bounded moment assumption, the conditions in theorem 3 impose some restrictions on the estimating functions, $m(x, \lambda)$. Lemma 1 shows that those conditions hold, in particular, for ridge, lasso, and pretest estimators.

Theorem 3 (uniform L^2 -convergence). *Suppose that:*

1. $m(x, \lambda)$ is monotonic in λ for all x in \mathbb{R}
2. $m(x, 0) = x$ and $\lim_{\lambda \rightarrow \infty} m(x, \lambda) = 0$ for all x in \mathbb{R}
3. $\sup_{\pi \in \mathcal{Q}} E_\pi[X^4] < \infty$
4. For any $\epsilon > 0$ there exists a set of regularization parameters $0 = \lambda_0 < \dots < \lambda_k = \infty$, which may depend on ϵ , such that

$$E_\pi[(|X - \mu| + |\mu|) |m(X, \lambda_j) - m(X, \lambda_{j-1})|] \leq \epsilon$$

for all $j = 1, \dots, k$ and all $\pi \in \mathcal{Q}$.

Then,

$$\sup_{\pi \in \mathcal{Q}} E_\pi \left[\sup_{\lambda \in [0, \infty]} \left(L_n(\lambda) - \bar{R}_\pi(\lambda) \right)^2 \right] \rightarrow 0. \quad (6)$$

Notice that the finiteness of $\sup_{\pi \in \mathcal{Q}} E_\pi[X^4]$ is equivalent to the finiteness of $\sup_{\pi \in \mathcal{Q}} E_\pi[\mu^4]$ and $\sup_{\pi \in \mathcal{Q}} E_\pi[(X - \mu)^4]$ via Jensen's and Minkowski's inequalities.

Lemma 1. *If $\sup_{\pi \in \mathcal{Q}} E_\pi[X^4] < \infty$, then equation (6) holds for ridge and lasso. If, in addition, X is continuously distributed with a bounded density, then equation (6) holds for pretest.*

Theorem 2 provides sufficient conditions for uniform loss consistency. The following corollary shows that under the same conditions we obtain uniform risk consistency, that is, the integrated risk of the estimator based on the data-driven choice $\hat{\lambda}_n$ becomes uniformly close to the risk of the oracle-optimal $\hat{\lambda}^*(\pi)$.² For the statement of this corollary, recall that $\bar{R}(m(\cdot, \hat{\lambda}_n), \pi)$ is the integrated risk of the estimator $m(\cdot, \hat{\lambda}_n)$ using the stochastic (data-dependent) $\hat{\lambda}_n$.

Corollary 1 (uniform risk consistency). *Under the assumptions of theorem 3,*

$$\sup_{\pi \in \mathcal{Q}} \left| \bar{R}(m(\cdot, \hat{\lambda}_n), \pi) - \inf_{\lambda \in [0, \infty]} \bar{R}_\pi(\lambda) \right| \rightarrow 0. \quad (7)$$

²An analogous result holds for uniform compound risk consistency.

In this section, we have shown that approximations to the risk function of machine learning estimators based on oracle-knowledge of λ are uniformly valid over $\pi \in \mathcal{Q}$ under mild assumptions. It is worth pointing out that such uniformity is not a trivial result. This is made clear by comparison to an alternative approximation, sometimes invoked to motivate the adoption of machine learning estimators, based on oracle knowledge of true zeros among μ_1, \dots, μ_n (see, e.g., Fan & Li, 2001). As shown in appendix A.2, assuming oracle knowledge of zeros does not yield a uniformly valid approximation.

B. Stein's Unbiased Risk Estimate

Theorem 2 provides sufficient conditions for uniform loss consistency using a general estimator r_n of risk. We now establish that our conditions apply to a particular estimator of r_n , known as Stein's unbiased risk estimate (SURE), first proposed by Stein (1981). SURE leverages the assumption of normality to obtain an elegant expression of risk as an expected sum of squared residuals plus a penalization term.

SURE as originally proposed requires that m be piecewise differentiable as a function of x , which excludes discontinuous estimators such as the pretest estimator $m_{PT}(x, \lambda)$. We provide a generalization in lemma 2 that allows for discontinuities. This lemma is stated in terms of integrated risk; with the appropriate modifications, the same result holds verbatim for compound risk.

Lemma 2 (SURE for piecewise differentiable estimators with discontinuities). *Suppose that $\mu \sim \vartheta$ and $X|\mu \sim N(\mu, 1)$. That is, the marginal density of X , f_π , is the convolution of ϑ with the standard normal distribution. Consider an estimator $m(X)$ of μ , and suppose that $m(x)$ is differentiable everywhere in $\mathbb{R} \setminus \{x_1, \dots, x_J\}$ but might be discontinuous at $\{x_1, \dots, x_J\}$. Let ∇m be the derivative of m (defined arbitrarily at $\{x_1, \dots, x_J\}$), and let $\Delta m_j = \lim_{x \downarrow x_j} m(x) - \lim_{x \uparrow x_j} m(x)$ for $j \in \{1, \dots, J\}$. Assume that $E_\pi[(m(X) - X)^2] < \infty$, $E_\pi[\nabla m(X)] < \infty$, and $(m(x) - x)\phi(x - \mu) \rightarrow 0$ as $|x| \rightarrow \infty$ ϑ -a.s. Then,*

$$\begin{aligned} \bar{R}(m(\cdot), \pi) &= E_\pi[(m(X) - X)^2] \\ &+ 2 \left(E_\pi[\nabla m(X)] + \sum_{j=1}^J \Delta m_j f_\pi(x_j) \right) - 1. \end{aligned}$$

The result of this lemma yields an objective function for the choice of λ of the general form we considered in section IVA, with $\bar{v}_\pi = -1$ and

$$\begin{aligned} \bar{r}_\pi(\lambda) &= E_\pi[(m(X, \lambda) - X)^2] \\ &+ 2 \left(E_\pi[\nabla_x m(X, \lambda)] + \sum_{j=1}^J \Delta m_j(\lambda) f_\pi(x_j) \right), \quad (8) \end{aligned}$$

where $\nabla_x m(x, \lambda)$ is the derivative of $m(x, \lambda)$ with respect to its first argument, and $\{x_1, \dots, x_J\}$ may depend on λ . The expression in equation (8) can be estimated using its sample analog,

$$r_n(\lambda) = \frac{1}{n} \sum_{i=1}^n (m(X_i, \lambda) - X_i)^2 + 2 \left(\frac{1}{n} \sum_{i=1}^n \nabla_x m(X_i, \lambda) + \sum_{j=1}^J \Delta m_j(\lambda) \hat{f}(x_j) \right), \quad (9)$$

where $\hat{f}(x)$ is an estimator of $f_\pi(x)$. This expression can be thought of as a penalized least squares objective function. The following are explicit expressions for the penalty for the cases of ridge, lasso, and pretest:

$$\text{Ridge: } \frac{2}{1 + \lambda}$$

$$\text{Lasso: } \frac{2}{n} \sum_{i=1}^n 1(|X_i| > \lambda)$$

$$\text{Pretest: } \frac{2}{n} \sum_{i=1}^n 1(|X_i| > \lambda) + 2\lambda(\hat{f}(-\lambda) + \hat{f}(\lambda))$$

The lasso penalty was previously derived in Donoho and Johnstone (1995). The result in lemma 2 applies not only to ridge, lasso, and pretest, but also to any other component-wise estimator of the normal means model.

To apply the uniform risk consistency in theorem 2, we need to show that equation (5) holds. That is, we have to show that $r_n(\lambda)$ is uniformly consistent as an estimator of $\bar{r}_\pi(\lambda)$. The following lemma provides the desired result.

Lemma 3. *Assume the conditions of theorem 3, and let $\bar{r}_\pi(\lambda)$ and $r_n(\lambda)$ be as in equations (8) and (9), respectively. Then equation (5) holds for $m(\cdot, \lambda)$ equal to $m_R(\cdot, \lambda)$, $m_L(\cdot, \lambda)$. If, in addition,*

$$\sup_{\pi \in \mathcal{Q}} P_\pi \left(\sup_{x \in \mathbb{R}} |x| \hat{f}(x) - |x| f_\pi(x) > \epsilon \right) \rightarrow 0 \quad \forall \epsilon > 0,$$

then equation (5) holds for $m(\cdot, \lambda)$ equal to $m_{PT}(\cdot, \lambda)$.

Identification of \bar{m}_π^* . Under the conditions of lemma 2, the optimal regularization parameter $\bar{\lambda}^*(\pi)$ is identified. In fact, under the same conditions, the stronger result holds that \bar{m}_π^* as defined in section IIIA is identified as well (see, e.g., Brown, 1971; Efron, 2011). The next lemma states the identification result for \bar{m}_π^* :

Lemma 4. *Under the conditions of lemma 2, the optimal shrinkage function is given by*

$$\bar{m}_\pi^*(x) = x + \nabla \log(f_\pi(x)).$$

Several nonparametric empirical Bayes estimators (NPEB) that target $\bar{m}_\pi^*(x)$ have been proposed (see Brown & Greenshtein, 2009; Jiang & Zhang, 2009; Efron, 2011; Koenker & Mizera, 2014). In particular, Jiang and Zhang (2009) derive asymptotic optimality results for nonparametric estimation of \bar{m}_π^* and provide an estimator based on the EM algorithm. The estimator proposed in Koenker and Mizera (2014), which is based on convex optimization techniques, is particularly attractive in terms of computational properties and because it sidesteps the selection of a smoothing parameters (see, e.g., Brown & Greenshtein, 2009). Both estimators, in Jiang and Zhang (2009) and Koenker and Mizera (2014), use a discrete distribution over a finite number of values to approximate the true distribution of μ . In sections V and VI, we use the Koenker-Mizera estimator to visually compare the shape of this estimated $\bar{m}_\pi^*(x)$ to the shape of ridge, lasso, and pretest estimating functions and to assess the performance of ridge, lasso, and pretest relative to the performance of a nonparametric estimator of \bar{m}_π^* .

C. Cross-Validation

A popular alternative to SURE is cross-validation, which chooses tuning parameters to optimize out-of-sample prediction. In this section, we investigate data-driven choices of the regularization parameter in a panel data setting, where multiple observations are available for each value of μ in the sample.

For $i = 1, \dots, n$, consider i.i.d. draws, $(x_{1i}, \dots, x_{ki}, \mu_i, \sigma_i)$, of a random variable $(x_1, \dots, x_k, \mu, \sigma)$ with distribution $\pi \in \mathcal{Q}$. Assume that the components of (x_1, \dots, x_k) are i.i.d. conditional on (μ, σ^2) and that for each $j = 1, \dots, k$,

$$E[x_j | \mu, \sigma] = \mu, \quad \text{and} \quad \text{var}(x_j | \mu, \sigma) = \sigma^2.$$

Let

$$X_k = \frac{1}{k} \sum_{j=1}^k x_j \quad \text{and} \quad X_{ki} = \frac{1}{k} \sum_{j=1}^k x_{ji}.$$

For concreteness and to simplify notation, we consider an estimator based on the first $k - 1$ observations for each group $i = 1, \dots, n$,

$$\hat{\mu}_{k-1,i} = m(X_{k-1,i}, \lambda),$$

and use observations x_{ki} , for $i = 1, \dots, n$, as a holdout sample to choose λ . Similar results hold for alternative sample partitioning choices. The loss function and empirical Bayes risk function of this estimator are given by

$$L_{n,k}(\lambda) = \frac{1}{n} \sum_{i=1}^n (m(X_{k-1,i}, \lambda) - \mu_i)^2 \quad \text{and}$$

$$\bar{R}_{\pi,k}(\lambda) = E_\pi[(m(X_{k-1}, \lambda) - \mu)^2].$$

Consider the following cross-validation estimator:

$$r_{n,k}(\lambda) = \frac{1}{n} \sum_{i=1}^n (m(X_{k-1,i}, \lambda) - x_{ki})^2.$$

Lemma 5. Assume conditions 1 and 2 of theorem 3 and $E_\pi[x_j^2] < \infty$, for $j = 1, \dots, k$. Then,

$$E_\pi[r_{n,k}(\lambda)] = \bar{R}_{\pi,k}(\lambda) + E_\pi[\sigma^2].$$

That is, the cross-validation yields an (up to a constant) unbiased estimator for the risk of the estimating function $m(X_{k-1}, \lambda)$. The following theorem shows that this result can be strengthened to a uniform consistency result.

Theorem 4. Assume conditions 1 and 2 of theorem 3 and $\sup_\pi E_\pi[x_j^4] < \infty$, for $j = 1, \dots, k$. Let $\bar{v}_\pi = -E_\pi[\sigma^2]$,

$$\bar{r}_{\pi,k}(\lambda) = E_\pi[r_{n,k}(\lambda)] = \bar{R}_{\pi,k}(\lambda) - \bar{v}_\pi,$$

and $\hat{\lambda}_n = \operatorname{argmin}_{\lambda \in [0, \infty]} r_{n,k}(\lambda)$. Then, for ridge, lasso, and pretest,

$$\sup_{\pi \in \mathcal{Q}} E_\pi \left[\sup_{\lambda \in [0, \infty]} (r_{n,k}(\lambda) - \bar{r}_{\pi,k}(\lambda))^2 \right] \rightarrow 0, \text{ and}$$

$$\sup_{\pi \in \mathcal{Q}} P_\pi \left(\left| L_{n,k}(\hat{\lambda}_n) - \inf_{\lambda \in [0, \infty]} L_{n,k}(\lambda) \right| > \epsilon \right) \rightarrow 0, \quad \forall \epsilon > 0.$$

Cross-validation has advantages as well as disadvantages relative to SURE. On the positive side, cross-validation does not rely on normal errors, while SURE does. Normality is less of an issue if k is large, so X_{ki} is approximately normal. On the negative side, however, cross-validation requires holding out some of the data from the second-step estimation of μ once the value of the regularization parameter has been chosen in a first step. This affects the essence of the cross-validation efficiency results, which apply to estimators of the form $m(X_{k-1,i}, \lambda)$, rather than to feasible estimators that use the entire sample in the second step, $m(X_{ki}, \lambda)$. Finally, cross-validation imposes greater data availability requirements, as it relies on availability of data on repeated realizations, x_{1i}, \dots, x_{ki} , of a random variable centered at μ_i , for each sample unit $i = 1, \dots, n$. This may hinder the practical applicability of cross-validation selection of regularization parameters in the context considered in this paper.

D. Comparison with Leeb and Pötscher (2006)

Our results on the uniform consistency of estimators of risk such as SURE or CV appear to stand in contradiction to those of Leeb and Pötscher (2006). They consider the same setting as we do—estimation of normal means—and the same types of estimators, including ridge, lasso, and pretest. In this setting, Leeb and Pötscher (2006) show that no uniformly consistent estimator of risk exists for such estimators.

The apparent contradiction between our results and the results in Leeb and Pötscher (2006) is explained by the different nature of the asymptotic sequence adopted in this paper to study the properties of machine learning estimators, relative to the asymptotic sequence adopted in Leeb and Pötscher (2006) for the same purpose. In this paper, we consider the problem of estimating a large number of parameters, such as location effects for many locations or group-level treatment effects for many groups. This motivates the adoption of an asymptotic sequence along which the number of estimated parameters increases as $n \rightarrow \infty$. In contrast, Leeb and Pötscher (2006) study the risk properties of regularized estimators embedded in a sequence along which the number of estimated parameters stays fixed as $n \rightarrow \infty$ and the estimation variance is of order $1/n$. We expect our approximation to work well when the dimension of the estimated parameter is large; the approximation of Leeb and Pötscher (2006) is likely to be more appropriate when the dimension of the estimated parameter is small while sample size is large.

In the simplest version of the setting in Leeb and Pötscher (2006), we observe a $(k \times 1)$ vector X_n with distribution $X_n \sim N(\mu_n, I_k/n)$, where I_k is the identity matrix of dimension k . Let X_{ni} and μ_{ni} be the i -components of X_n and μ_n , respectively. Consider the componentwise estimator $m_n(X_{ni})$ of μ_{ni} . Leeb and Pötscher (2006) study consistent estimation of the normalized compound risk,

$$R_n^{LP} = nE \|m_n(X_n) - \mu_n\|^2,$$

where $m_n(X_n)$ is a $(k \times 1)$ vector with the i th element equal to $m_n(X_{ni})$ and the sequence μ_n is taken as fixed.

Adopting the reparameterization, $Y_n = \sqrt{n}X_n$ and $h_n = \sqrt{n}\mu_n$, we obtain $Y_n - h_n \sim N(0, I_k)$. Notice that for the maximum likelihood estimator, $m_n(X_n) - \mu_n = (Y_n - h_n)/\sqrt{n}$ and $R_n^{LP} = E \|m(Y_n) - h_n\|^2 = k$, so the risk of the maximum likelihood estimator does not depend on the sequence h_n and therefore can be consistently estimated. This is not the case for shrinkage estimators, however. Choosing $h_n = h$ for some fixed h , the problem becomes invariant in n , $Y_n \sim N(h, I_k)$. In this setting, it is easy to show that the risk of machine learning estimators, such as ridge, lasso, and pretest, depends on h , and therefore it cannot be estimated consistently. For instance, consider the lasso estimator, $m_n(x) = m_L(x, \lambda_n)$, where $\sqrt{n}\lambda_n \rightarrow c$ with $0 < c < \infty$, as in Leeb and Pötscher (2006). Then, lemma A.1 in the appendix implies that R_n^{LP} converges to a constant that depends on h . As a result, R_n^{LP} cannot be estimated consistently.³

Contrast the setting in Leeb and Pötscher (2006) to the one adopted in this paper, where we consider a high-dimensional setting, such that X and μ have dimension equal to n . The

³This result holds more generally outside the normal error model. Let $m_L(X_n, \lambda)$ be the $(n \times 1)$ vector with the i th element equal to $m_L(X_i, \lambda)$. Consider the sequence of regularization parameters $\lambda_n = c/\sqrt{n}$; then $m_L(x, \lambda_n) = m_L(\sqrt{n}x, c)/\sqrt{n}$. This implies $R_n^{LP} = E \|m_L(Y_n, c) - h\|^2$, which is invariant in n .

pairs (X_i, μ_i) follow a distribution π , which may vary with n . As n increases, π becomes identified, and so does the average risk, $E_\pi[(m_n(X_i) - \mu_i)^2]$, of any componentwise estimator, $m_n(\cdot)$.

Whether the asymptotic approximation in Leeb and Pötscher (2006) or ours provides a better description of the performance of SURE, CV, or other estimators of risk in actual applications depends on the dimension of μ . If this dimension is large, as is typical in the applications we consider in this paper, we expect our uniform consistency result to apply: a “blessing of dimensionality.” As Leeb and Pötscher (2006) demonstrated, however, precise estimation of a fixed number of parameters does not ensure uniformly consistent estimation of risk.

E. Mixed Estimators and Estimators of the Optimal Shrinkage Function

We have discussed criteria such as SURE and CV as means to select the regularization parameter, λ . In principle, these same criteria might also be used to choose among alternative estimators, such as ridge, lasso, and pretest, in specific empirical settings. Our uniform risk consistency results imply that such a mixed-estimator approach dominates each of the estimators that are being mixed, for n large enough. Going even further, one might aim to estimate the optimal shrinkage function, \bar{m}_π^* , using the result of lemma 4, as in Jiang and Zhang (2009), Koenker and Mizera (2014), and others. Under suitable consistency conditions, this approach will dominate all other componentwise estimators for large enough n (Jiang & Zhang, 2009). In practice, these results should be applied with some caution, as they are based on neglecting the variability in the choice of estimation procedure or in the estimation of \bar{m}_π^* . For small and moderate values of n , procedures with fewer, degrees of freedom may perform better in practice. We return to this issue in section V, where we compare the finite sample risk of the machine learning estimators considered in this paper (ridge, lasso, and pretest) to the finite sample risk of the nonparametric empirical Bayes (NPEB) estimator of Koenker and Mizera (2014).

V. Simulations

Designs. To gauge the relative performance of the estimators considered in this paper, we next report the results of a set of simulations that employ the spike and normal data-generating process of section IIIC. That is, we consider distributions π of (X, μ) such that μ is degenerate at zero with probability p and normal with mean μ_0 and variance σ_0^2 with probability $(1 - p)$. We consider all combinations of parameter values $p = 0.00, 0.25, 0.50, 0.75, 0.95$, $\mu_0 = 0, 2, 4$, $\sigma_0 = 2, 4, 6$, and sample sizes $n = 50, 200, 1000$.

Given a set of values μ_1, \dots, μ_n , the values for X_1, \dots, X_n are generated as follows. To evaluate the performance of estimators based on SURE selectors and of the NPEB esti-

mator of Koenker and Mizera (2014), we generate the data as

$$X_i = \mu_i + U_i, \quad (10)$$

where the U_i follow a standard normal distribution independent of other components. To evaluate the performance of cross-validation estimators, we generate $x_{ji} = \mu_i + \sqrt{k}u_{ji}$ for $j = 1, \dots, k$, where the u_{ji} are draws from independent standard normal distributions. As a result, the averages,

$$X_{ki} = \frac{1}{k} \sum_{j=1}^k x_{ji},$$

have the same distributions as the X_i in equation (10), which makes the comparison between the cross-validation estimators and the SURE and NPEB estimators a meaningful one. For cross-validation estimators, we consider $k = 4, 20$.

Estimators. The SURE criterion function employed in the simulations is the one in equation (9), where, for the pretest estimator, the density of X is estimated with a normal kernel and the bandwidth implied by Silverman’s rule of thumb.⁴ The cross-validation criterion function employed in the simulations is a leave-one-out version of the one considered in section IVC,

$$r_{n,k}(\lambda) = \sum_{j=1}^k \left(\frac{1}{n} \sum_{i=1}^n (m(X_{-ji}, \lambda) - x_{ji})^2 \right), \quad (11)$$

where X_{-ji} is the average of $\{x_{1i}, \dots, x_{ki}\} \setminus x_{ji}$. Notice that because the result in theorem 4 applies to each of the k terms on the right-hand side of equation (11), it also applies to $r_{n,k}(\lambda)$.

Results. Figures 5, 6, and 7 report average compound risk across 1,000 simulations for $n = 50, n = 200$ and $n = 1,000$, respectively. Each row corresponds to a particular value of (p, μ_0, σ_0) , and each column corresponds to a particular estimator or regularization criterion. The results are coded row by row on a continuous color scale that varies from dark (minimum row value; blue in the digital version) to light (maximum row value; yellow in the digital version).

Several clear patterns emerge from the simulation results. First, even for a dimensionality as modest as $n = 50$, the patterns in figure 3, which were obtained for oracle choices of regularization parameters, are reproduced in Figures 5 to 7 for the same estimators but using data-driven choices of regularization parameters. As in figure 3, among ridge, lasso, and pretest, ridge dominates when there is little or no sparsity in the parameters of interest, pretest dominates when the distribution of nonzero parameters is substantially separated from zero, and lasso dominates in the intermediate cases.

⁴See Silverman (1986, eq. 3.31).

FIGURE 5.—AVERAGE COMPOUND LOSS ACROSS 1,000 SIMULATIONS WITH $n = 50$

p	μ_0	σ_0	SURE			Cross-Validation ($k = 4$)			Cross-Validation ($k = 20$)			NPEB
			ridge	lasso	pretest	ridge	lasso	pretest	ridge	lasso	pretest	
0.00	0	2	0.80	0.89	1.02	0.83	0.90	1.12	0.81	0.88	1.12	0.94
0.00	0	4	0.95	0.98	1.02	0.96	0.98	1.08	0.94	0.97	1.08	1.15
0.00	0	6	0.97	0.99	1.01	0.97	0.99	1.05	0.97	0.99	1.07	1.21
0.00	2	2	0.89	0.96	1.01	0.90	0.95	1.06	0.89	0.95	1.09	0.93
0.00	2	4	0.96	0.99	1.01	0.96	0.98	1.06	0.96	0.98	1.09	1.13
0.00	2	6	0.97	0.99	1.01	0.99	1.00	1.06	0.97	0.98	1.07	1.21
0.00	4	2	0.95	1.00	1.01	0.95	0.99	1.02	0.95	1.00	1.04	0.93
0.00	4	4	0.97	1.00	1.01	0.99	1.01	1.06	0.97	0.99	1.07	1.15
0.00	4	6	0.99	1.00	1.02	0.99	1.00	1.05	0.99	1.00	1.07	1.21
0.25	0	2	0.75	0.78	1.01	0.77	0.79	1.12	0.77	0.78	1.08	0.85
0.25	0	4	0.92	0.90	1.00	0.94	0.90	1.07	0.92	0.88	1.05	1.04
0.25	0	6	0.97	0.93	0.99	0.97	0.92	1.02	0.96	0.92	1.02	1.09
0.25	2	2	0.87	0.88	1.01	0.87	0.86	1.06	0.87	0.86	1.07	0.88
0.25	2	4	0.93	0.90	0.99	0.94	0.89	1.04	0.95	0.90	1.04	1.03
0.25	2	6	0.97	0.93	0.98	0.98	0.93	1.03	0.97	0.93	1.02	1.09
0.25	4	2	0.94	0.95	0.99	0.95	0.95	1.03	0.95	0.95	1.04	0.92
0.25	4	4	0.97	0.94	0.99	0.97	0.93	1.03	0.97	0.93	1.02	1.04
0.25	4	6	0.98	0.94	0.98	0.98	0.93	1.02	0.98	0.93	1.00	1.09
0.50	0	2	0.67	0.64	0.94	0.69	0.64	0.96	0.67	0.62	0.90	0.69
0.50	0	4	0.89	0.75	0.92	0.91	0.76	0.92	0.89	0.75	0.89	0.82
0.50	0	6	0.95	0.80	0.90	0.95	0.79	0.87	0.96	0.78	0.84	0.84
0.50	2	2	0.80	0.72	0.96	0.82	0.72	0.96	0.81	0.72	0.93	0.73
0.50	2	4	0.92	0.77	0.94	0.93	0.76	0.90	0.90	0.75	0.87	0.83
0.50	2	6	0.96	0.80	0.92	0.95	0.77	0.83	0.95	0.78	0.82	0.86
0.50	4	2	0.91	0.82	0.95	0.92	0.81	0.90	0.92	0.81	0.87	0.75
0.50	4	4	0.94	0.80	0.93	0.94	0.79	0.87	0.94	0.78	0.83	0.81
0.50	4	6	0.97	0.81	0.93	0.97	0.79	0.83	0.96	0.78	0.79	0.85
0.75	0	2	0.51	0.43	0.61	0.51	0.42	0.57	0.50	0.41	0.57	0.46
0.75	0	4	0.77	0.50	0.59	0.80	0.51	0.58	0.78	0.50	0.57	0.52
0.75	0	6	0.88	0.54	0.55	0.90	0.54	0.55	0.88	0.53	0.52	0.51
0.75	2	2	0.66	0.49	0.65	0.67	0.49	0.63	0.67	0.49	0.62	0.47
0.75	2	4	0.81	0.53	0.59	0.86	0.54	0.58	0.82	0.52	0.56	0.51
0.75	2	6	0.90	0.56	0.54	0.91	0.56	0.53	0.90	0.55	0.52	0.51
0.75	4	2	0.84	0.59	0.64	0.85	0.57	0.60	0.84	0.58	0.58	0.49
0.75	4	4	0.88	0.56	0.57	0.89	0.55	0.53	0.89	0.55	0.52	0.50
0.75	4	6	0.92	0.57	0.53	0.92	0.55	0.49	0.92	0.56	0.51	0.50
0.95	0	2	0.18	0.15	0.17	0.17	0.12	0.15	0.18	0.13	0.19	0.17
0.95	0	4	0.37	0.19	0.17	0.37	0.17	0.17	0.37	0.18	0.20	0.18
0.95	0	6	0.49	0.21	0.16	0.51	0.19	0.16	0.49	0.19	0.19	0.16
0.95	2	2	0.26	0.17	0.18	0.27	0.16	0.18	0.27	0.17	0.23	0.17
0.95	2	4	0.40	0.19	0.17	0.43	0.18	0.16	0.40	0.18	0.20	0.17
0.95	2	6	0.53	0.21	0.15	0.53	0.19	0.15	0.53	0.20	0.18	0.16
0.95	4	2	0.44	0.21	0.18	0.45	0.20	0.18	0.45	0.20	0.22	0.18
0.95	4	4	0.51	0.21	0.16	0.51	0.19	0.17	0.52	0.20	0.19	0.17
0.95	4	6	0.57	0.21	0.15	0.58	0.19	0.14	0.57	0.20	0.18	0.16

A color version of this figure is available online at https://doi.org/10.1162/rest_a_00812.

Second, while the results in Jiang and Zhang (2009) suggest good performance of nonparametric estimators of \bar{m}_π^* for large n , the simulation results in Figures 5 and 6 indicate that the performance of NPEB may be substantially worse than the performance of the other machine learning estimators in the table for moderate and small n . In particular, the performance of the NPEB estimator suffers in the settings with low or no sparsity, especially when the distribution of

the nonzero values of μ_1, \dots, μ_n has considerable dispersion. This is explained by the fact that in practice, the NPEB estimator approximates the distribution of μ using a discrete distribution supported on a small number of values. When most of the probability mass of the true distribution of μ is also concentrated around a small number of values (i.e., when p is large or σ_0 is small), the approximation employed by the NPEB estimator is accurate and the performance of the

FIGURE 6.—AVERAGE COMPOUND LOSS ACROSS 1,000 SIMULATIONS WITH $n = 200$

p	μ_0	σ_0	SURE			Cross-Validation ($k = 4$)			Cross-Validation ($k = 20$)			NPEB
			ridge	lasso	pretest	ridge	lasso	pretest	ridge	lasso	pretest	
0.00	0	2	0.80	0.87	1.01	0.82	0.88	1.04	0.80	0.87	1.04	0.86
0.00	0	4	0.94	0.96	1.00	0.95	0.97	1.02	0.94	0.96	1.03	1.03
0.00	0	6	0.98	0.99	1.01	0.98	0.99	1.02	0.98	0.99	1.03	1.09
0.00	2	2	0.89	0.95	1.00	0.90	0.95	1.02	0.89	0.94	1.03	0.86
0.00	2	4	0.95	0.97	1.00	0.96	0.98	1.02	0.96	0.97	1.03	1.03
0.00	2	6	0.98	1.00	1.01	0.98	0.99	1.02	0.98	0.99	1.03	1.10
0.00	4	2	0.95	1.00	1.00	0.96	1.00	1.01	0.95	1.00	1.02	0.86
0.00	4	4	0.97	0.99	1.00	0.97	0.99	1.01	0.97	0.98	1.02	1.03
0.00	4	6	0.98	0.99	1.01	0.98	0.99	1.01	0.99	0.99	1.03	1.09
0.25	0	2	0.75	0.77	1.00	0.77	0.78	1.07	0.75	0.75	1.04	0.78
0.25	0	4	0.92	0.88	0.99	0.93	0.88	1.02	0.93	0.88	1.02	0.95
0.25	0	6	0.96	0.91	0.99	0.97	0.91	1.01	0.96	0.91	1.00	0.98
0.25	2	2	0.86	0.86	1.00	0.87	0.86	1.03	0.86	0.85	1.03	0.80
0.25	2	4	0.94	0.90	1.00	0.95	0.90	1.02	0.93	0.88	1.01	0.95
0.25	2	6	0.97	0.92	0.99	0.97	0.92	1.00	0.97	0.91	1.00	0.98
0.25	4	2	0.94	0.95	1.00	0.94	0.93	1.00	0.94	0.94	1.01	0.83
0.25	4	4	0.96	0.92	0.99	0.97	0.92	1.01	0.95	0.91	0.99	0.94
0.25	4	6	0.97	0.92	0.98	0.97	0.92	0.99	0.97	0.92	0.98	0.98
0.50	0	2	0.67	0.61	0.90	0.69	0.62	0.93	0.67	0.61	0.90	0.63
0.50	0	4	0.89	0.74	0.90	0.90	0.74	0.89	0.89	0.73	0.86	0.76
0.50	0	6	0.94	0.77	0.86	0.95	0.76	0.82	0.95	0.77	0.83	0.77
0.50	2	2	0.80	0.70	0.94	0.82	0.71	0.93	0.80	0.69	0.91	0.65
0.50	2	4	0.92	0.75	0.92	0.92	0.75	0.87	0.91	0.74	0.86	0.76
0.50	2	6	0.95	0.78	0.88	0.96	0.78	0.83	0.95	0.77	0.82	0.77
0.50	4	2	0.91	0.80	0.94	0.92	0.81	0.87	0.91	0.80	0.87	0.67
0.50	4	4	0.94	0.78	0.94	0.95	0.78	0.83	0.94	0.77	0.82	0.74
0.50	4	6	0.96	0.79	0.92	0.97	0.79	0.81	0.97	0.78	0.80	0.76
0.75	0	2	0.50	0.39	0.55	0.51	0.40	0.57	0.50	0.39	0.55	0.40
0.75	0	4	0.80	0.50	0.55	0.81	0.50	0.57	0.80	0.49	0.56	0.48
0.75	0	6	0.90	0.53	0.49	0.91	0.53	0.52	0.89	0.52	0.50	0.47
0.75	2	2	0.67	0.47	0.59	0.68	0.47	0.61	0.67	0.46	0.59	0.42
0.75	2	4	0.83	0.50	0.53	0.84	0.51	0.56	0.83	0.51	0.55	0.46
0.75	2	6	0.91	0.54	0.50	0.91	0.54	0.52	0.91	0.53	0.51	0.47
0.75	4	2	0.83	0.56	0.55	0.85	0.57	0.58	0.83	0.55	0.56	0.42
0.75	4	4	0.89	0.54	0.50	0.90	0.54	0.52	0.88	0.53	0.50	0.45
0.75	4	6	0.93	0.55	0.47	0.93	0.55	0.49	0.92	0.53	0.48	0.46
0.95	0	2	0.17	0.12	0.14	0.17	0.12	0.14	0.17	0.12	0.15	0.12
0.95	0	4	0.43	0.17	0.16	0.43	0.17	0.16	0.42	0.17	0.16	0.14
0.95	0	6	0.61	0.18	0.14	0.62	0.18	0.14	0.61	0.18	0.14	0.14
0.95	2	2	0.28	0.16	0.17	0.29	0.16	0.18	0.28	0.15	0.17	0.14
0.95	2	4	0.46	0.17	0.15	0.48	0.17	0.16	0.47	0.17	0.16	0.14
0.95	2	6	0.63	0.19	0.14	0.64	0.19	0.14	0.63	0.18	0.14	0.13
0.95	4	2	0.49	0.20	0.17	0.50	0.20	0.17	0.48	0.19	0.17	0.14
0.95	4	4	0.58	0.19	0.14	0.59	0.19	0.14	0.59	0.18	0.15	0.14
0.95	4	6	0.68	0.19	0.13	0.70	0.19	0.13	0.67	0.19	0.14	0.13

A color version of this figure is available online at https://doi.org/10.1162/rest_a_00812.

NPEB estimator is good. This is not the case, however, when the true distribution of μ cannot be closely approximated with a small number of values (i.e., when p is small and σ_0 is large). Lasso shows a remarkable degree of robustness to the value of (p, μ_0, σ_0) , which makes it an attractive estimator in practice. For large n , as in figures 7, NPEB dominates except in settings with no sparsity and a large dispersion in μ ($p = 0$ and σ_0 large).

VI. Applications

In this section, we apply our results to three data sets from the empirical economics literature. The first application, based on Chetty and Hendren (2018), estimates the effect of living in a given commuting zone during childhood on intergenerational income mobility. The second application, based on Della Vigna and La Ferrara (2010), estimates

FIGURE 7.—AVERAGE COMPOUND LOSS ACROSS 1,000 SIMULATIONS WITH $n = 1,000$

p	μ_0	σ_0	SURE			Cross-Validation ($k = 4$)			Cross-Validation ($k = 20$)			NPEB
			ridge	lasso	pretest	ridge	lasso	pretest	ridge	lasso	pretest	
0.00	0	2	0.80	0.87	1.01	0.81	0.87	1.01	0.80	0.86	1.01	0.82
0.00	0	4	0.94	0.96	1.00	0.95	0.97	1.01	0.94	0.96	1.00	0.97
0.00	0	6	0.97	0.98	1.00	0.98	0.98	1.00	0.97	0.98	1.01	1.02
0.00	2	2	0.89	0.94	1.00	0.90	0.95	1.00	0.89	0.94	1.01	0.82
0.00	2	4	0.95	0.97	1.00	0.96	0.97	1.00	0.95	0.97	1.01	0.98
0.00	2	6	0.97	0.98	1.00	0.98	0.99	1.00	0.97	0.98	1.01	1.02
0.00	4	2	0.95	1.00	1.00	0.96	1.00	1.00	0.95	0.99	1.00	0.82
0.00	4	4	0.97	0.99	1.00	0.97	0.99	1.01	0.97	0.99	1.01	0.97
0.00	4	6	0.98	0.99	1.00	0.98	0.99	1.00	0.98	0.99	1.01	1.02
0.25	0	2	0.75	0.76	1.00	0.76	0.77	1.02	0.75	0.75	1.01	0.74
0.25	0	4	0.92	0.88	0.99	0.93	0.88	1.00	0.92	0.87	1.00	0.89
0.25	0	6	0.97	0.91	0.99	0.97	0.91	0.99	0.96	0.91	0.99	0.92
0.25	2	2	0.86	0.85	1.00	0.87	0.86	1.01	0.86	0.84	1.01	0.76
0.25	2	4	0.94	0.89	1.00	0.94	0.89	1.00	0.94	0.89	1.00	0.89
0.25	2	6	0.97	0.91	0.99	0.97	0.91	0.99	0.97	0.91	0.99	0.92
0.25	4	2	0.94	0.94	0.99	0.94	0.94	0.99	0.94	0.93	0.99	0.79
0.25	4	4	0.96	0.92	0.99	0.96	0.91	0.99	0.96	0.91	0.99	0.88
0.25	4	6	0.98	0.92	0.99	0.98	0.92	0.98	0.97	0.92	0.98	0.91
0.50	0	2	0.67	0.60	0.87	0.68	0.61	0.90	0.67	0.60	0.87	0.60
0.50	0	4	0.89	0.73	0.85	0.90	0.73	0.86	0.89	0.72	0.85	0.71
0.50	0	6	0.95	0.77	0.81	0.95	0.77	0.82	0.95	0.76	0.81	0.72
0.50	2	2	0.80	0.70	0.90	0.81	0.71	0.90	0.80	0.69	0.89	0.62
0.50	2	4	0.91	0.74	0.85	0.92	0.75	0.85	0.91	0.74	0.84	0.70
0.50	2	6	0.95	0.77	0.80	0.96	0.78	0.81	0.95	0.77	0.80	0.71
0.50	4	2	0.91	0.80	0.87	0.92	0.80	0.84	0.91	0.80	0.84	0.63
0.50	4	4	0.94	0.77	0.88	0.95	0.78	0.81	0.94	0.77	0.80	0.68
0.50	4	6	0.96	0.78	0.87	0.97	0.78	0.79	0.96	0.78	0.78	0.70
0.75	0	2	0.50	0.38	0.54	0.51	0.40	0.56	0.50	0.38	0.54	0.38
0.75	0	4	0.80	0.49	0.53	0.81	0.50	0.55	0.80	0.48	0.53	0.44
0.75	0	6	0.90	0.52	0.49	0.91	0.53	0.51	0.90	0.52	0.49	0.43
0.75	2	2	0.67	0.46	0.57	0.68	0.47	0.59	0.67	0.46	0.58	0.40
0.75	2	4	0.83	0.50	0.52	0.85	0.51	0.55	0.83	0.50	0.53	0.44
0.75	2	6	0.91	0.53	0.48	0.92	0.53	0.50	0.91	0.52	0.48	0.43
0.75	4	2	0.83	0.55	0.53	0.85	0.56	0.55	0.83	0.55	0.54	0.39
0.75	4	4	0.89	0.53	0.49	0.90	0.54	0.51	0.89	0.52	0.49	0.41
0.75	4	6	0.93	0.54	0.46	0.94	0.54	0.48	0.93	0.53	0.47	0.42
0.95	0	2	0.17	0.11	0.14	0.17	0.12	0.14	0.17	0.11	0.14	0.11
0.95	0	4	0.44	0.16	0.15	0.45	0.16	0.16	0.44	0.16	0.15	0.13
0.95	0	6	0.63	0.18	0.13	0.65	0.18	0.14	0.64	0.17	0.14	0.12
0.95	2	2	0.28	0.15	0.16	0.29	0.15	0.18	0.29	0.14	0.17	0.12
0.95	2	4	0.49	0.16	0.14	0.50	0.17	0.16	0.50	0.16	0.15	0.12
0.95	2	6	0.66	0.18	0.13	0.67	0.18	0.14	0.66	0.18	0.13	0.12
0.95	4	2	0.50	0.19	0.16	0.51	0.19	0.17	0.50	0.19	0.16	0.12
0.95	4	4	0.61	0.18	0.14	0.62	0.18	0.14	0.61	0.18	0.14	0.12
0.95	4	6	0.72	0.18	0.13	0.73	0.19	0.13	0.71	0.18	0.13	0.12

A color version of this figure is available online at https://doi.org/10.1162/rest_a_00812.

changes in the stock prices of arms manufacturers following changes in the intensity of conflicts in countries under arms trade embargoes. The third application uses data from the 2000 U.S. census, previously employed in Angrist, Chernozhukov, and Fernández-Val (2006) and Belloni and Chernozhukov (2011), to estimate a nonparametric Mincer regression equation of log wages on education and potential experience. We use these three examples for illustrative pur-

poses only. The original studies provide in-depth analyses of the issues considered in this section.

For all applications, we normalize the observed X_i by their estimated standard error. Note that this normalization defines (a) the implied loss function, which is quadratic error loss for estimation of the normalized latent parameter μ_i , and (b) the class of estimators considered, which are componentwise shrinkage estimators based on the normalized X_i .

A. Neighborhood Effects: Chetty and Hendren (2018)

Chetty and Hendren (2018) use information on income at age 26 for individuals who moved between commuting zones during childhood to estimate the effects of location on income. Identification comes from comparing differently aged children of the same parents who are exposed to different locations for different durations in their youth. In the context of this application, X_i is the (studentized) estimate of the effect of spending an additional year of childhood in commuting zone i , conditional on parental income rank, on child income rank relative to the national household income distribution at age 26.⁵ In this setting, the point zero has no special role; it is just defined, by normalization, to equal the average of commuting zone effects. We therefore have no reason to expect sparsity or the presence of a set of effects well separated from zero. Our discussion in section III would thus lead us to expect that ridge will perform well, and this is indeed what we find.

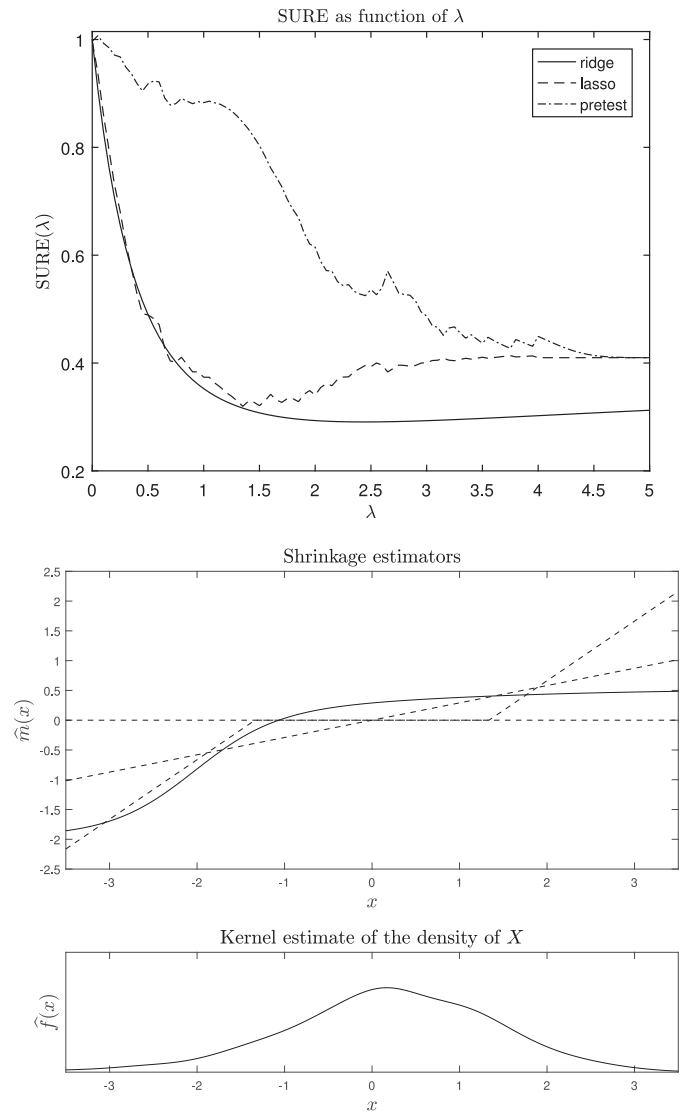
Figure 8 (top panel) reports SURE estimates of risk for ridge, lasso, and pretest estimators, as functions of λ . Among the three estimators, minimal estimated risk is equal to 0.29, and it is attained by ridge for $\hat{\lambda}_{R,n} = 2.44$. Minimal estimated risk for lasso and pretest are 0.31 and 0.41, respectively. The relative performance of the three shrinkage estimators reflects the characteristics of the example and, in particular, the very limited evidence of sparsity in the data.

The second panel of figure 8 shows the Koenker-Mizera NPEB estimator (solid line) along with the ridge, lasso, and pretest estimators (dashed lines) evaluated at SURE-minimizing values of the regularization parameters. The identity of the estimators can be easily recognized from their shape. The ridge estimator is linear, with positive slope equal to estimated risk, 0.29. Lasso has the familiar piecewise linear shape, with kinks at the positive and negative versions of the SURE-minimizing value of the regularization parameter, $\hat{\lambda}_{L,n} = 1.34$. Pretest is flat at 0 because SURE is minimized for values of λ higher than the maximum absolute value of X_1, \dots, X_n . The second panel shows a kernel estimate of the distribution of X .⁶ Among ridge, lasso, and pretest, ridge best approximates the optimal shrinkage estimator over most of the estimated distribution of X . Lasso comes a close second, as evidenced in the minimal SURE values for the three estimators, and pretest is way off. Despite substantial shrinkage, these estimates suggest considerable heterogeneity in the effects of childhood neighborhood on earnings. In addition, as expected given the nature of this application, we do not find evidence of sparsity in the location effects estimates.

⁵The data employed in this section were obtained from http://www.equality-of-opportunity.org/images/nbhds_online_data_table3.xlsx. We focus on the estimates for children with parents at the 25th percentile of the national income distribution among parents with children in the same birth cohort.

⁶To produce a smooth depiction of densities, for the panels reporting densities in this section we use the normal reference rule to choose the bandwidth. See Silverman (1986, equation 3.28).

FIGURE 8.—NEIGHBORHOOD EFFECTS



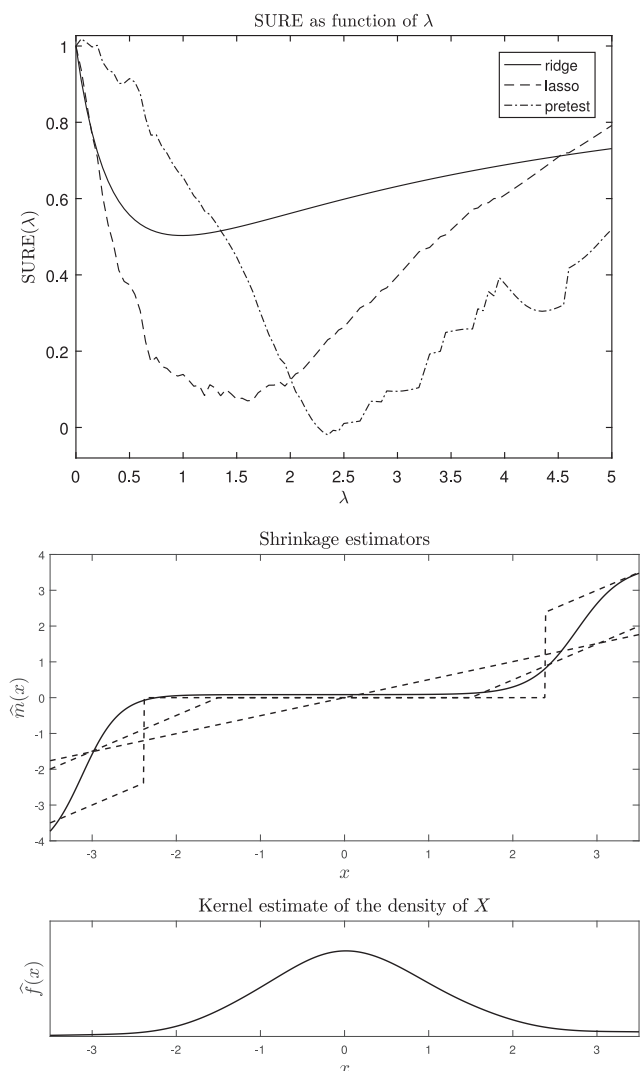
The middle panel shows the Koenker-Mizera NPEB estimator (solid line) along with the ridge, lasso, and pretest estimators (dashed lines) evaluated at SURE-minimizing values of the regularization parameters. The ridge estimator is linear, with positive slope equal to estimated risk, 0.29. Lasso is piecewise linear, with kinks at the positive and negative versions of the SURE-minimizing value of the regularization parameter, $\hat{\lambda}_{L,n} = 1.34$. Pretest is flat at 0 because SURE is minimized for values of λ higher than the maximum absolute value of X_1, \dots, X_n . The bottom panel shows a kernel estimate of the distribution of X .

B. Detecting Illegal Arms Trade: Della Vigna and La Ferrara (2010)

Della Vigna and La Ferrara (2010) use changes in stock prices of arms manufacturing companies at the time of large changes in the intensity of conflicts in countries under arms trade embargoes to detect illegal arms trade. In this section, we apply the estimators in section IV to data from the Della Vigna and La Ferrara study.⁷

⁷Della Vigna and La Ferrara (2010) divide their sample of arms manufacturers into two groups, depending on whether the company is headquartered in a country with a high or low level of corruption. They also divide the events of changes in the intensity of the conflicts in embargo areas in two groups, depending on whether the intensity of the conflict increased or decreased at the time of the event. For concreteness, we use the

FIGURE 9.—ARMS EVENT STUDY



The middle panel shows the Koenker-Mizera NPEB estimator (solid line) along with the ridge, lasso, and pretest estimators (dashed lines) evaluated at SURE-minimizing values of the regularization parameters. The ridge estimator is linear, with positive slope equal to estimated risk, 0.50. Lasso is piecewise linear, with kinks at the positive and negative versions of the SURE-minimizing value of the regularization parameter, $\hat{\lambda}_{L,n} = 1.50$. Pretest is discontinuous at $\hat{\lambda}_{PT,n} = 2.39$ and $-\hat{\lambda}_{PT,n} = -2.39$. The bottom panel shows a kernel estimate of the distribution of X .

In contrast to the location effects example in section VIA, in this application there are reasons to expect a certain amount of sparsity if changes in the intensity of the conflicts in arms embargo areas do not affect the stock prices of arms manufacturers that comply with the embargoes.⁸ In this case, our discussion in section III would lead us to expect that pretest might be optimal, which is again what we find. Figure 9 (top

214-event study estimates for events of increase in the intensity of conflicts in arms embargo areas and for companies in high-corruption countries. The data for this application are available at <http://eml.berkeley.edu/~sdellavi/wp/AEJDataPostingZip.zip>.

⁸In the words of Della Vigna and La Ferrara (2010), "If a company is not trading or trading legally, an event increasing the hostilities should not affect its stock price or should affect it adversely, since it delays the removal of the embargo and hence the re-establishment of legal sales. Conversely, if a company is trading illegally, the event should increase its stock price, since it increases the demand for illegal weapons."

panel) shows SURE estimates for ridge, lasso, and pretest. Pretest has the lowest estimated risk, for $\hat{\lambda}_{PT,n} = 2.39$,⁹ followed by lasso, for $\hat{\lambda}_{L,n} = 1.50$.

Figure 9 (middle panel) depicts the different shrinkage estimators and shows that lasso and, especially, pretest closely approximate the NPEB estimator over a large part of the distribution of X . The NPEB estimate suggests a substantial amount of sparsity in the distribution of μ . There is, however, a subset of the support of X around $x = 3$ where the estimate of the optimal shrinkage function implies only a small amount of shrinkage. Given the shapes of the optimal shrinkage function estimate and the estimate of the distribution of X , it is not surprising that the minimal values of SURE in figure 9 (top panel) for lasso and pretest are considerably lower than for ridge.

C. Nonparametric Mincer Equation: Belloni and Chernozhukov (2011)

In our third application, we use data from the 2000 U.S. Census in order to estimate a nonparametric regression of log wages on years of education and potential experience, similar to the example considered in Belloni and Chernozhukov (2011).¹⁰ We construct a set of 66 regressors by taking a saturated basis of linear splines in education, fully interacted with the terms of a sixth-order polynomial in potential experience. We orthogonalize these regressors and take the coefficients X_i of an OLS regression of log wages on these orthogonalized regressors as our point of departure. We exclude three coefficients of very large magnitude,¹¹ which results in $n = 63$. In this application, economics provides less intuition as to what distribution of coefficients to expect. Belloni and Chernozhukov (2011) argue that for plausible families of functions containing the true conditional expectation function, sparse approximations of the coefficients of series regression as induced by the lasso penalty have low mean squared error.

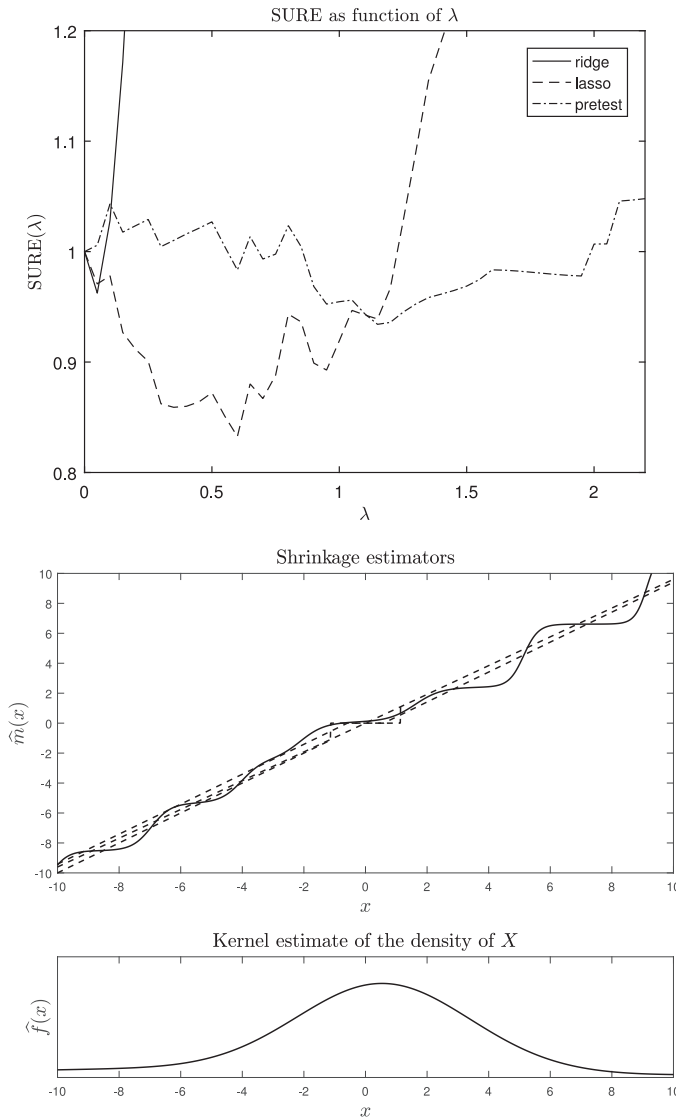
Figure 10 (top panel) reports SURE estimates of risk for ridge, lasso, and pretest. In this application, the estimated risk for lasso is substantially smaller than for ridge or pretest. The middle panel of figure 10 reports the three regularized estimators—ridge, lasso, and pretest—evaluated at the data-driven choice of regularization parameter, along with the Koenker-Mizera NPEB estimator. In order to visualize the differences between the estimates close to the origin, where most of the coefficients are, we report the value of the

⁹Notice that the pretest's SURE estimate attains a negative minimum value. This could be a matter of estimation variability, of inappropriate choice of bandwidth for the estimation of the density of X in small samples, or it could reflect misspecification of the model (in particular, normality of X given μ).

¹⁰The data for this application are available at <http://economics.mit.edu/files/384>.

¹¹The three excluded coefficients have values of 2938.04 (the intercept), 98.19, and -77.35 . The largest absolute value among the included coefficients is -21.06 . Most of the included coefficients are small in absolute value. About 40% of them have absolute values smaller than 1, and about 60% have absolute value smaller than 2.

FIGURE 10.—NONPARAMETRIC MINCER EQUATION



The middle panel shows the Koenker-Mizera NPEB estimator (solid line) along with the ridge, lasso, and pretest estimators (dashed lines) evaluated at SURE-minimizing values of the regularization parameters. The ridge estimator is linear, with positive slope equal to estimated risk, 0.996. Lasso is piecewise linear, with kinks at the positive and negative versions of the SURE-minimizing value of the regularization parameter, $\lambda = 0.59$. Pretest is discontinuous at $\hat{\lambda}_{PT,n} = 1.14$ and $-\hat{\lambda}_{PT,n} = -1.14$. The bottom panel shows a kernel estimate of the distribution of X .

estimates for $x \in [-10, 10]$. The bottom panel of figure 10 reports an estimate of the density of X . Locally, the shape of the NPEB estimate looks similar to a step function. This behavior is explained by the fact that the NPEB estimator is based on an approximation to the distribution of μ that is supported on a finite number of values. However, over the whole range of x in figure 10, the NPEB estimate is fairly linear. In view of this close-to-linear behavior of NPEB in the $[10, 10]$ interval, the very poor risk performance of ridge relative to lasso and pretest, as evidenced in figure 10 (top panel), may appear surprising. This is explained by the fact that in this application, some of the values in X_1, \dots, X_n fall exceedingly far from the origin. Linearly shrinking those values toward 0 induces severe loss. As a result, ridge attains minimal

risk for a close-to-zero value of the regularization parameter, $\hat{\lambda}_{R,n} = 0.04$, resulting in negligible shrinkage. Among ridge, lasso, and pretest, minimal estimated risk is attained by lasso for $\hat{\lambda}_{L,n} = 0.59$, which shrinks about 24% of the regression coefficients all the way to zero. Pretest induces higher sparsity ($\hat{\lambda}_{PT,n} = 1.14$, shrinking about 49% of the coefficients all the way to zero) but does not improve over lasso in terms of risk.

VII. Conclusion

The interest in adopting machine learning methods in economics is growing rapidly. Two common features of machine learning algorithms are regularization and data-driven choice of regularization parameters. We study the properties of such procedures. We consider, in particular, the problem of estimating many means μ_i based on observations X_i . This problem arises often in economic applications. In such applications, the “observations” X_i are usually equal to preliminary least squares coefficient estimates, like fixed effects.

Our goal is to provide guidance for applied researchers on the use of machine learning estimators. Which estimation method should one choose in a given application? And how should one choose regularization parameters? To the extent that researchers care about the squared error of their estimates, procedures are preferable if they have lower mean squared errors than the competitors do.

Based on our results, ridge appears to dominate the alternatives considered when the true effects μ_i are smoothly distributed, and there is no point mass of true zero. This is likely to be the case in applications where the objects of interests are the effects of many treatments, such as locations or teachers, and applications that estimate effects for many subgroups. Pretest appears to dominate if there are true zeros and non-zero effects are well separated from zero. This happens in economic applications when there are fixed costs for agents who engage in non-zero behavior. Lasso finally dominates for intermediate cases and appears to do well for series regression in particular.

Regarding the choice of regularization parameters, we prove a series of results that show that data-driven choices are almost optimal (in a uniform sense) for large-dimensional problems. This is the case, in particular, for choices of regularization parameters that minimize Stein’s unbiased risk estimate (SURE), when observations are normally distributed, and for cross-validation (CV), when repeated observations for a given effect are available. Although not explicitly analyzed in this paper, equation (3) suggests a new empirical selector of regularization parameters based on the minimization of the sample mean square discrepancy between $m(X_i, \lambda)$ and NPEB estimates of $\bar{m}_\pi^*(X_i)$.

There are, of course, some limitations to our analysis. First, we focus on a restricted class of estimators: those that can be written in the componentwise shrinkage form $\hat{\mu}_i = m(X_i, \hat{\lambda})$. This covers many estimators of interest for economists, most notably ridge, lasso, and pretest estimation. Many other estimators in the machine learning literature, such as random

forests or neural nets, do not have this tractable form. The analysis of the risk properties of such estimators constitutes an interesting avenue of future research. Finally, we focus on mean square error. This loss function is analytically quite convenient and amenable to tractable results. Other loss functions might be of practical interest, however, and might be studied using numerical methods. In this context, it is also worth emphasizing again that we were focusing on point estimation, where all coefficients μ_i are simultaneously of interest. This is relevant for many practical applications such as those discussed here. In other cases, however, one might instead be interested in the estimates $\hat{\mu}_i$ solely as input for a lower-dimensional decision problem or in (frequentist) testing of hypotheses on the coefficients μ_i . Our analysis of mean squared error does not directly speak to such questions.

REFERENCES

- Abowd, J. M., F. Kramarz, and D. N. Margolis, "High Wage Workers and High Wage Firms," *Econometrica* 67:2 (1999), 251–333.
- Abrams, D., M. Bertrand, and S. Mullainathan, "Do Judges Vary in Their Treatment of Race?" *Journal of Legal Studies* 41:2 (2012), 347–383.
- Angrist, J., V. Chernozhukov, and I. Fernández-Val, "Quantile Regression under Misspecification, with an Application to the U.S. Wage Structure," *Econometrica* 74:2 (2006), 539–563.
- Belloni, A., and V. Chernozhukov, "High Dimensional Sparse Econometric Models: An Introduction," (pp. 121–156), in P. Alquier, E. Gautier, and G. Stoltz, eds., *Inverse Problems and High-Dimensional Estimation: Stats in the Château Summer School, August 31 - September 4, 2009*, vol. 203 of *Lecture Notes in Statistics*, pp. 121–156 (Berlin: Springer, 2011).
- Brown, L. D., "Admissible Estimators, Recurrent Diffusions, and Insoluble Boundary Value Problems," *Annals of Mathematical Statistics* 42:3 (1971), 855–903.
- Brown, L. D., and E. Greenshtein, "Nonparametric Empirical Bayes and Compound Decision Approaches to Estimation of a High-Dimensional Vector of Normal Means," *Annals of Statistics* 37:4 (2009), 1685–1704.
- Chetty, R., J. N. Friedman, and J. E. Rockoff, "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood," *American Economic Review* 104:9 (2014), 2633–2679.
- Chetty, R., and N. Hendren, "The Impacts of Neighborhoods on Intergenerational Mobility: Childhood Exposure Effects and County-Level Estimates," *Quarterly Journal of Economics* 133:3 (2018), 1107–1162.
- Della Vigna, S., and E. La Ferrara, "Detecting Illegal Arms Trade," *American Economic Journal: Economic Policy* 2:4 (2010), 26–57.
- Donoho, D. L., and I. M. Johnstone, "Adapting to Unknown Smoothness via Wavelet Shrinkage," *Journal of the American Statistical Association* 90:432 (1995), 1200–1224.
- Efron, B., "Tweedie's Formula and Selection Bias," *Journal of the American Statistical Association* 106:496 (2011), 1602–1614.
- Fan, J. F., and R. Li, "Variable Selection via Nonconcave Penalized Likelihood and Its Oracle Properties," *Journal of the American Statistical Association* 96:456 (2001), 1348–1360.
- Jiang, W., and C.-H. Zhang, "General Maximum Likelihood Empirical Bayes Estimation of Normal Means," *Annals of Statistics* 37:4 (2009), 1647–1684.
- Koenker, R., and I. Mizera, "Convex Optimization, Shape Constraints, Compound Decisions, and Empirical Bayes Rules," *Journal of the American Statistical Association* 109:506 (2014), 674–685.
- Krueger, A. B., "Experimental Estimates of Education Production Functions," *Quarterly Journal of Economics* 114:2 (1999), 497–532.
- Leeb, H., and B. M. Pötscher, "Performance Limits for Estimators of the Risk or Distribution of Shrinkage-Type Estimators, and Some General Lower Risk-Bound Results," *Econometric Theory* 22:1 (2006), 69–97.
- Newey, W. K., "Convergence Rates and Asymptotic Normality for Series Estimators," *Journal of Econometrics* 79:1 (1997), 147–168.
- Silverman, B., *Density Estimation for Statistics and Data Analysis* (London: Chapman & Hall/CRC, 1986).
- Stein, C. M., "Estimation of the Mean of a Multivariate Normal Distribution," *Annals of Statistics* 9:6 (1981), 1135–1151.
- Stock, J. H., and M. W. Watson, "Generalized Shrinkage Methods for Forecasting Using Many Predictors," *Journal of Business and Economic Statistics* 30:4 (2012), 481–493.
- Wasserman, L., *All of Nonparametric Statistics* (New York: Springer, 2006).