



AARHUS UNIVERSITY



Coversheet

This is the accepted manuscript (post-print version) of the article.

Contentwise, the post-print version is identical to the final published version, but there may be differences in typography and layout.

How to cite this publication

Please cite the final published version:

Caner, M., & Kock, A. B. (2016). Oracle Inequalities for Convex Loss Functions with Non-Linear Targets. *Econometric Reviews*, 35(8-10), 1377-1411. DOI: 10.1080/07474938.2015.1092797

Publication metadata

Title: Oracle Inequalities for Convex Loss Functions with Non-Linear Targets
Author(s): Caner, M., & Kock, A. B.
Journal: *Econometric Reviews*, 35(8-10), 1377-1411.
DOI/Link: <http://dx.doi.org/10.1080/07474938.2015.1092797>
Document version: Accepted manuscript (post-print)

General Rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- *Users may download and print one copy of any publication from the public portal for the purpose of private study or research.*
- *You may not further distribute the material or use it for any profit-making activity or commercial gain*
- *You may freely distribute the URL identifying the publication in the public portal*

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Oracle Inequalities for Convex Loss Functions with Non-Linear Targets

MEHMET CANER* ANDERS BREDAHL KOCK†

February 1, 2018

Abstract

This paper consider penalized empirical loss minimization of convex loss functions with unknown non-linear target functions. Using the elastic net penalty we establish a finite sample oracle inequality which bounds the loss of our estimator from above with high probability. If the unknown target is linear this inequality also provides an upper bound of the estimation error of the estimated parameter vector. These are new results and they generalize the econometrics and statistics literature. Next, we use the non-asymptotic results to show that the excess loss of our estimator is asymptotically of the same order as that of the oracle. If the target is linear we give sufficient conditions for consistency of the estimated parameter vector. Next, we briefly discuss how a thresholded version of our estimator can be used to perform consistent variable selection. We give two examples of loss functions covered by our framework and show how penalized nonparametric series estimation is contained as a special case and provide a finite sample upper bound on the mean square error of the elastic net series estimator.

Keywords and phrases: Empirical loss minimization, Lasso, Elastic net, Oracle inequality, Convex loss function, Nonparametric estimation, Variable selection.

JEL classification: C13, C21, C31.

*North Carolina State University, Department of Economics, 4168 Nelson Hall, Raleigh, NC 27695. Email: mcaner@ncsu.edu.

†Aarhus University and CREATES, Department of Economics, Fuglesangs Alle 4, 8210 Aarhus V, Denmark. Email: akock@creates.au.dk.

1 Introduction

Recently high-dimensional data sets have become increasingly available to researchers in many fields. In economics big data can be found in the analysis of consumer behavior based on scanner data from purchases. Furthermore, many macroeconomic variables are sampled rather infrequently leaving one with many variables compared to observations in models with many explanatory variables. Financial data is also of a high-dimensional nature with many variables and instruments being observed in small intervals due to high-frequency trading. Alternatively, models with many variables emerge when trying to control for non-linearities in a wage regression by including basis functions of the space in which the non-linearity is supposed to be found. Clearly, including more basis functions can result in better approximations of the non-linearity. However, this also results in a model with many variables, i.e. a high-dimensional model. For these reasons handling high-dimensional data sets has received a lot of attention in the econometrics and statistics literature in the recent years. In a seminal paper Tibshirani (1996) introduced the Lasso estimator which carries out variable selection and parameter estimation simultaneously. The theoretical properties of this estimator have been studied extensively since then in various papers and extensions such as the adaptive Lasso by Zou (2006), the bridge estimator by Huang et al. (2008), the sure independence screening by Fan and Lv (2008) or the square root Lasso by Belloni et al. (2011) have been proposed. For recent reviews see, e.g., Fan et al. (2011), Bühlmann and van de Geer (2011) or Belloni and Chernozhukov (2011).

In the econometrics literature Lasso-type estimators have also proven useful. For example Belloni et al. (2012) have established results in the context of instrumental variable estimation without imposing the hitherto much used assumption of sub-gaussianity by means of moderate deviation theorems for self-normalized random variables. Furthermore, they allow for heteroscedastic error terms which is pathbreaking and greatly widens the scope of applicability of their results.

Applications to panel data may be found in e.g. Kock (2013). The estimators have been studied in the context of GMM, factor models, and smooth penalties by, among others, Caner and Zhang (2013), Caner and Han (2013), Cheng and Liao (2013) and Fan and Li (2001). Within linear time series models oracle inequalities have been established by Kock and Callot (2013) and Negahban et al. (2012) have proposed a unified framework which is valid for regression as well as matrix estimation problems.

Most research has considered the linear regression model or other parametric models. In this

paper we shall focus on a very general setup. In particular, we will focus on penalized empirical loss minimization of convex loss functions with potentially non-linear target functions. van de Geer (2008) studied a similar setup for the Lasso which is a special case of our results for the elastic net. Furthermore, even though our main focus is on non-asymptotic bounds, we also present asymptotic upper bounds on the excess risk and estimation error (the latter in the case where the target is linear). We also show how our results can be used to give new non-asymptotic upper bounds on penalized series estimators with many series terms.

In particular, we

1. provide a finite sample oracle inequality for empirical risk minimization penalized by the elastic net penalty. This inequality is valid for convex loss functions and non-linear targets and contains an oracle inequality for the Lasso as a special case.
2. For the case where the target function is linear this oracle inequality can be used to establish finite sample upper bounds on the estimation error of the estimated parameter vector.
3. The finite sample inequality is used to establish asymptotic results. In particular, the excess risk of our estimator is of the same order as that of an oracle which trades off the approximation and estimation errors. When the target is linear we give sufficient conditions for consistency of the estimated parameter vector.
4. In the case where the target is linear we briefly explain how a thresholded version of our estimator can unveil the correct sparsity pattern.
5. We provide two examples of specific loss functions covered by our general framework. We verify in detail that the abstract conditions are satisfied in common settings. Then we show how nonparametric series estimation is contained as a special case of our theory and provide a finite sample upper bound on the mean square error of an elastic net series estimator. We explain why this series estimator may be more precise than classical series estimators.
6. We also note that when the loss function is quadratic the sample does not have to be identically distributed and our results are therefore also valid in the presence of heteroscedasticity.

We stress here that our main objective is to establish upper bounds on the performance of the elastic net. It is not our intention to promote either the Lasso or the elastic net, merely to analyze the properties of the latter. However, we shall make some brief comments on merits of the two

procedures when compared to each other. A clear ranking like the one in Hebiri and van de Geer (2011) is not available at this point. However, these authors only focus on quadratic loss for which a certain data augmentation trick facilitates the analysis.

We believe that the performance guarantees on the elastic net provided by this paper are useful for the applied researcher who increasingly faces high-dimensional data sets. The usefulness is enhanced by the fact that our results are valid for a wide range of loss functions and that heteroscedasticity is allowed for when the loss function is quadratic.

The paper is organized as follows. Section 2 puts forward the setup and notation. Section 3 introduces the main result, the oracle inequality for empirical loss minimization of convex loss functions penalized by the elastic net. Section 4 briefly discusses consistent variable selection by a thresholded version of the elastic net. Tuning parameter selection is handled in Section 5. Section 6 shows that the quadratic as well as the logistic loss are covered by our framework and provides an oracle inequality for penalized series estimators.

2 Setup and notation

We begin by setting the stage for general convex loss minimization. The setup is similar to the Lasso one in Section 6.3 in Bühlmann and van de Geer (2011). Let (Ω, \mathcal{F}, P) be a standard probability space. Consider a sample $\{Z_i\}_{i=1}^n = \{X_i, Y_i\}_{i=1}^n$ with $X_i \in \mathcal{X}$ and $Y_i \in \mathcal{Y} \subseteq \mathbb{R}$. Here, for the sake of exposition, \mathcal{X} can be thought of as a subset of \mathbb{R}^p for $p \geq 1$ but as we shall see below it can be much more general. Define $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ and let \mathbf{F} be a normed real vector space with norm $\|\cdot\|$. For each $f \in \mathbf{F}$ let $\rho_f : \mathcal{Z} \rightarrow \mathbb{R}$ be a loss function. More precisely, $f \in \mathbf{F}$ will be a function $f : \mathcal{X} \rightarrow \mathbf{R}$ and the corresponding norm will most often be the $L_2(P)$ -norm $(\int f^2(X)dP)^{1/2}$ (in fact $\|\cdot\|$ will be the $L_2(P)$ -norm in almost all our econometric examples¹). Furthermore, when v is a vector in \mathbb{R}^p , $\|v\|_1 = \sum_{j=1}^p |v_j|$ denotes the ℓ_1 -norm while $\|v\|_2 = \sqrt{\sum_{j=1}^p v_j^2}$ denotes the ℓ_2 -norm. Order symbols such as o , O , o_p and O_p are used with their usual meanings. Also, in accordance with the usual Landau notation, $g_1(n) \in \Omega(g_2(n))$ means that there exists a constant $a > 0$ such that $g_1(n) \geq ag_2(n)$ for n sufficiently large. $\Theta(g_2(n))$ denotes the intersection of $O(g_2(n))$ and $\Omega(g_2(n))$ and contains all functions $g_1(n)$ that are exactly of order $g_2(n)$. Finally, for any abstract set A , $|A|$ denotes its cardinality. Throughout the paper we shall assume:

Assumption 0: $\{Z_i\}_{i=1}^n = \{X_i, Y_i\}_{i=1}^n$ is an independent sample and the mapping $f \mapsto \rho_f(z)$

¹This choice of norm on \mathbf{F} is suitable when the sample is supposed to be i.i.d.

is convex for all $z \in \mathcal{Z}$.

The following examples provide illustrations of when the conditions in Assumption 0 are met:

Quadratic loss

Let $\mathcal{X} \subset \mathbb{R}^p$ and

$$Y_i = f^0(X_i) + \epsilon_i, \quad i = 1, \dots, n$$

where ϵ_i is some real error term and $f^0 \in \mathbf{F}$. Then, the standard case of quadratic loss is covered by the above setting upon choosing $\rho_f(x, y) = (y - f(x))^2$ which is clearly convex in $f(x)$. By letting \mathbf{F} only consist of linear functions $f(x) = \beta'x$ for some $\beta \in \mathbb{R}^p$ the case of linear least squares is covered. Non-linear least squares is covered by choosing $f(x) = g(\beta, x)$ for some parameter vector β . As we shall see in Section 6.1.1 this setup can also be used to obtain some new upper bounds on nonparametric series estimation.

Logistic loss

Let

$$Y_i^* = f^0(X_i) + \epsilon_i, \quad i = 1, \dots, n$$

where ϵ_i is independent of X_i and assumed to have a logistic distribution while $f^0 \in \mathbf{F}$. Assume that $Y_i = 1$ if $Y_i^* > 0$ and $Y_i = 0$ otherwise. Since ϵ_i has cdf $F(z) = \frac{e^z}{1+e^z}$ one gets

$$P(Y_i = 1 | X_i = x) = P(f^0(X_i) + \epsilon_i > 0 | X_i = x) = \frac{e^{f^0(x)}}{1 + e^{f^0(x)}}$$

Note that for $\mathcal{X} \subseteq \mathbb{R}^p$ and $f^0(x) = \beta^0'x$ for some parameter vector $\beta \in \mathbb{R}^p$ this is the usual expression for $P(Y_i = 1 | X_i = x)$ in the logit model. The above setting is more general, however, since it allows f^0 to be non-linear.

The log-likelihood function for a given $f \in \mathbf{F}$ is then given by (for $z = (x, y)$)

$$l(f | z_1, \dots, z_n) = \sum_{i=1}^n \left[y_i \log\left(\frac{e^{f(x_i)}}{1 + e^{f(x_i)}}\right) + (1 - y_i) \log\left(1 - \frac{e^{f(x_i)}}{1 + e^{f(x_i)}}\right) \right] = \sum_{i=1}^n \left[y_i f(x_i) - \log(1 + e^{f(x_i)}) \right]$$

Hence, a sensible loss function is the negative log-likelihood

$$\rho_f(x, y) = -y f(x) + \log(1 + e^{f(x)})$$

which is convex in $f(x)$.

Negative log-likelihood

The above two examples are both instances of the loss function being the negative of the log-likelihood. Hence, in a general setting with the negative of the log-likelihood being a convex function in $f(x)$ our results also apply. Again, a special case is $f(x) = \beta^t x$.

Returning to the general setup, denote by $P_n \rho_f = \frac{1}{n} \sum_{i=1}^n \rho_f(X_i, Y_i)$ and $P \rho_f = \frac{1}{n} \sum_{i=1}^n E \rho_f(X_i, Y_i)$ the empirical and population means of the loss function for a fixed $f \in \mathbf{F}$. We shall also denote these two quantities the empirical and population risk, respectively. Note also, that in the case of identically distributed variables the population mean reduces to the plain expectation $E \rho_f(X_1, Y_1)$. We define our target as the minimizer of the theoretical risk

$$f^0 := \operatorname{argmin}_{f \in \mathbf{F}} P \rho_f$$

where it is tacitly assumed that the minimizer exists and is unique for the $\|\cdot\|$ -norm on \mathbf{F} . Then, for any $f \in \mathbf{F}$, we define the excess population risk over the target as

$$\Xi(f) := P(\rho_f - \rho_{f^0})$$

Note that, by construction, $\Xi(f) \geq 0$ for all $f \in \mathbf{F}$. Since the joint distribution of (Y_i, X_i) is assumed to be unknown we shall consider empirical risk minimization instead of minimizing the population excess risk. Put differently, $P_n \rho_f$ is minimized. Furthermore, we will consider a linear subspace $\mathcal{F}_L = \{f_\beta(x) = \sum_{j=1}^p \beta_j \psi_j(x), \beta \in \Phi\}$ of \mathbf{F} where $\psi_j(x) : \mathcal{X} \rightarrow \mathbb{R}$ may be thought of as basis functions of \mathbf{F} . Of course, in the case where \mathcal{X} is a subset of \mathbb{R}^p , one could also think of $\psi_j(x)$ as being the j 'th coordinate projection. This is the choice we make whenever f^0 is assumed to be linear. In general, $\psi(x) = (\psi_1(x), \dots, \psi_p(x))'$ denotes a vector of (transformed) covariates. Φ is a convex subset of \mathbb{R}^p – in many cases we can even have $\Phi = \mathbb{R}^p$. In Section 6 we shall see an example of $\Phi = \mathbb{R}^p$ but also an example of Φ being a subset of \mathbb{R}^p . In case the target function f^0 is linear we will denote its parameter vector by β^0 .

As we shall see, it is possible to prove upper bounds on the excess risk of a penalized version of the empirical risk minimizer even when f^0 is non-linear while we only minimize over the linear subspace \mathcal{F}_L . This is non-trivial since the target belongs to a large set (\mathbf{F}) while we only minimize over a smaller set (\mathcal{F}_L). The following section gives an exact definition and discussion of our estimator.

2.1 The elastic net

Define

$$\hat{f} = f_{\hat{\beta}} = \operatorname{argmin}_{f_{\beta} : \beta \in \Phi} \left(P_n \rho_{f_{\beta}} + \lambda_1 \|\beta\|_1 + \lambda_2 \|\beta\|_2^2 \right) = \operatorname{argmin}_{f \in \mathcal{F}_L} \left(P_n \rho_f + \lambda_1 \|\beta\|_1 + \lambda_2 \|\beta\|_2^2 \right)$$

where λ_1 and λ_2 are positive constants. Hence, we are minimizing the empirical risk *plus* an elastic net penalty. This form of penalty was originally introduced by Zou and Hastie (2005) in the case of a linear regression model. The penalty is a compromise between the ℓ_1 -penalty of the plain Lasso and the squared ℓ_2 -loss in ridge regression. Ridge regression does not perform variable selection at all – all estimated coefficients are non-zero. On the other hand, if two variables are highly correlated, the Lasso has a tendency to include only one of these. The elastic net strikes a balance between these two extremes and hence performs particularly well in the presence of highly correlated variables. This benefit has been formalized by Hebiri and van de Geer (2011) in the case of quadratic loss. In particular they have shown that the elastic net behaves better with respect to certain restricted eigenvalue conditions than the plain Lasso.

2.2 Assumptions

We next turn to the assumptions needed to prove oracle inequalities for the elastic net. First, define $\mathbf{F}_{\text{local}} = \{f \in \mathbf{F} : \|f - f^0\|_{\infty} \leq \eta\}$ for some $\eta > 0$ where $\|f\|_{\infty} = \sup_{x \in \mathcal{X}} |f(x)|$ ². The margin condition requires that in $\mathbf{F}_{\text{local}} \subseteq \mathbf{F}$ the excess loss $\Xi(f)$ is bounded from below by a convex function of $\|f - f^0\|$.

Definition. *We say that the **margin condition** holds with strictly convex margin function $G(\cdot)$, if for all $f \in \mathbf{F}_{\text{local}}$ we have*

$$\Xi(f) \geq G(\|f - f^0\|).$$

In all examples we shall consider it can be shown that the margin condition holds for $G(u) = cu^2$ for some $c > 0$ such that for all $f \in \mathbf{F}_{\text{local}}$, $\Xi(f) \geq c\|f - f^0\|^2$. More generally, we present a sufficient condition for G to be quadratic in Section 6. The convex conjugate of $G(\cdot)$ will also play a role in the development of the oracle inequalities below. In particular, the following definition is taken from page 121 in Bühlmann and van de Geer (2011) and many more properties of convex conjugates can be found in Rockafellar (1997).

²Using the stronger $\|\cdot\|_{\infty}$ topology on \mathbf{F} to define $\mathbf{F}_{\text{local}}$ instead of $\|\cdot\|$ turns out to be useful when verifying that the margin condition is satisfied with a quadratic margin in Section 6.

Definition. Let G be a strictly convex function on $[0, \infty)$ with $G(0) = 0$. The **convex conjugate** H of G is defined as

$$H(v) = \sup_{u \geq 0} \{uv - G(u)\}, \quad v \geq 0,$$

Lemma 3 in the appendix establishes some properties of $H(v)$. Note also that if $G(u) = cu^2$, then $H(v) = v^2/(4c)$. Furthermore, from the definition of the convex conjugate

$$uv \leq G(u) + H(v). \tag{1}$$

which is also known as Fenchel's inequality. Next, for any subset S of $\{1, \dots, p\}$ and $\beta \in \mathbb{R}^p$ we define β_S such that $\beta_{S,j} = \beta_j 1_{\{j \in S\}}$ for $j = 1, \dots, p$. Letting $|S| = s$ denote the cardinality of S we may define

Definition. The **adaptive restricted eigenvalue condition** is satisfied with if

$$\phi^2(S) = \min \left\{ \frac{\|f_\beta\|^2}{\|\beta_S\|_2^2} : \beta \in \mathbb{R}^p \setminus \{0\}, \|\beta_{S^c}\|_1 \leq L_n \|\beta_S\|_2 \right\} > 0$$

where $L_n = 3(\sqrt{|S|} + \frac{2\lambda_2 \|\beta_S\|_2}{\lambda_1})$. As mentioned already, in many econometric examples one may choose $\|\cdot\|$ to be the $L_2(P)$ -norm. In this case, and if the covariates are also identically distributed (as will be assumed in our concrete examples in Section 6), $\|f_\beta\|^2 = \beta' \Sigma \beta$ where $\Sigma = E(\psi(X_1)\psi'(X_1))$. Hence,

$$\phi^2(S) = \min \left\{ \frac{\beta' \Sigma \beta}{\|\beta_S\|_2^2} : \beta \in \mathbb{R}^p \setminus \{0\}, \|\beta_{S^c}\|_1 \leq L_n \|\beta_S\|_2 \right\} \tag{2}$$

Since

$$\phi^2(S) \geq \min_{\beta \in \mathbb{R}^p} \frac{\beta' \Sigma \beta}{\|\beta_S\|_2^2} \geq \min_{\beta \in \mathbb{R}^p} \frac{\beta' \Sigma \beta}{\|\beta\|_2^2}$$

for all $S \in \{1, \dots, p\}$ the adaptive restricted eigenvalue condition is satisfied in particular when the smallest eigenvalue of the population covariance matrix is positive. However, since the minimum in (2) is taken over a subset of \mathbb{R}^p only, we may have $\phi(S) > 0$ even when Σ is singular. Note also that the adaptive restricted eigenvalue condition is used various guises in the literature and is similar to the eigenvalue conditions of Bickel et al. (2009) and Hebiri and van de Geer (2011).

Before defining what we understand by the oracle estimator, define $S_\beta = \{j : \beta_j \neq 0\}$ as the subset of $\{1, \dots, p\}$ containing the indices of the non-zero coefficients. Let $s_\beta = |S_\beta|$ denote the cardinality of this set. Then, letting Γ denote a collection of subsets of $\{1, \dots, p\}$, we define

Definition The *oracle estimator* is defined as

$$\beta^* = \operatorname{argmin}_{\beta: S_\beta \in \Gamma} \left\{ 3\Xi(f_\beta) + 2H \left(\frac{4\lambda_1 \sqrt{s_\beta} + 4\lambda_2 \|\beta\|_2}{\phi(S_\beta)} \right) \right\}. \tag{3}$$

Note that the definition of the oracle still leaves considerable freedom since Γ is defined by the user – a property which we shall utilize later when considering linear targets (see e.g. remark 2 after Theorems 1 below). In the case where Γ equals the power set of $\{1, \dots, p\}$ the oracle estimator may equivalently be written as

$$\beta^* = \operatorname{argmin}_{\beta \in \mathbb{R}^p} \left\{ 3\Xi(f_\beta) + 2H \left(\frac{4\lambda_1 \sqrt{s_\beta} + 4\lambda_2 \|\beta\|_2}{\phi(S_\beta)} \right) \right\}.$$

The definition of the oracle in (3) turns out to be convenient for technical reasons but it also has a useful interpretation as a tradeoff between approximation and estimation error: In the standard setting of a quadratic loss function with a linear target, i.e. $f^0(x) = x'\beta^0$, it is known that the squared ℓ_2 -estimation error of Lasso type estimators when estimating p parameters, of which s are non-zero, is of the order $\frac{s \log(p)}{n}$. In the case of quadratic loss in the beginning of this section one has $\Xi(f) = E(f(X_1) - f^0(X_1))^2$ if the sample is identically distributed and X_1 and ϵ_1 are independent. So G and hence H are quadratic in the definition of the margin condition. Choosing $\lambda_2 = \frac{\lambda_1 \sqrt{s}}{2\|\beta\|}$ and λ_1 of the order $\sqrt{\log(p)/n}$, which are both choices we shall adhere to in the sequel, one finds that $H(\cdot)$ is of the order $\frac{s \log(p)}{n}$. This is exactly the estimation error under quadratic loss and motivates coining $H(\cdot)$ the estimation error term. The term $\Xi(f_\beta)$ is referred to as the approximation error and (3) shows that the oracle trades of these two terms: a lower approximation error can be obtained by increasing s_β while this also implies estimating more parameters resulting in a higher estimation error.

Finally, letting $S^* = S_{\beta^*}$, $\phi_* = \phi(S^*)$, $s^* = |S^*|$ and $f^* = f_{\beta^*}$ we denote the oracle bound (value of the objective function minimized by the oracle) by

$$2\Delta^* := 3\Xi(f^*) + 2H \left(\frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi_*} \right).$$

The inequality in Theorem 1 below will be valid on a random set which we introduce next. In Theorem 2 we shall show that this set actually has a high probability by means of a suitable concentration inequality for suprema of empirical processes. Define the empirical process

$$\left\{ V_n(\beta) = \frac{1}{n} \sum_{i=1}^n [\rho_{f_\beta}(Z_i) - E\rho_{f_\beta}(Z_i)], \beta \in \mathbb{R}^p \right\}.$$

Next, we introduce a local supremum of the empirical process in incremental form

$$Z_M = \sup_{\|\beta - \beta^*\|_1 \leq M} |V_n(\beta) - V_n(\beta^*)|. \quad (4)$$

Then we define

$$M^* = \Delta^* / \lambda_0,$$

where λ_0 is a positive sequence and set

$$\tau = \{Z_{M^*} \leq \lambda_0 M^*\} = \{Z_{M^*} \leq \Delta^*\}.$$

The set τ is the one we shall work on in Theorem 1 below. Note in particular, that on τ , Z_M can not be larger than Δ^* which is the minimal value of the loss function of the oracle.

We are now ready to state our assumptions:

Assumption 1. *Assume the margin condition with strictly convex function $G(\cdot)$.*

Assumption 2. *Assume that $f^* \in \mathbf{F}_{\text{local}}$ and $f_\beta \in \mathbf{F}_{\text{local}}$ for all $\|\beta - \beta^*\|_1 \leq M^*$.*

Assumption 3. *Assume that adaptive restricted eigenvalue condition holds for S^* , i.e. $\phi(S^*) > 0$.*

As discussed above, the margin condition, Assumption 1, regulates the behavior of the excess risk function. When \mathbf{F} is equipped with the $L_2(P)$ we will see in Section 6 that the margin condition is actually often satisfied with $G(\cdot)$ being quadratic. Put differently, the margin condition is satisfied in many examples with a quadratic margin.

Assumption 2 is a technical condition which enables us to use the margin condition for f^* . The first part requires that the oracle is a good approximation to f^0 in the sup-norm. Of course the validity of this statement depends on how well linear combinations of elements in $\{\psi_j\}_{j=1}^p$ can approximate f^0 . The validity also depends on the choice of Γ in the definition of the oracle since the precise form of f^* depends on this. For concrete choices of \mathbf{F} one can make proper choices of bases that guarantee the desired degree of approximation. Note in particular, that it follows from remark 2 in Section 3, that when f^0 is linear one can choose Γ such that $f^* = f^0$ and so, a fortiori, $f^* \in \mathbf{F}_{\text{local}}$. The second part of Assumption 2 states that $f_\beta \in \mathbf{F}$ if β is close to β^* . This is rather innocent by the triangle inequality. We will give more detailed sufficient conditions for Assumption 2 in Section 6 for concrete econometric examples. In particular, if \mathbf{F} consists of sufficiently smooth functions³, we shall exhibit concrete choices of bases $\{\psi_j\}_{j=1}^\infty$ and collections of sets Γ such that f^* approximates f^0 to the desired degree. Assumption 3 has been discussed above and is valid when $\|\cdot\|$ is the $L^2(P)$ -norm and $\Sigma = E(\psi(X_1)\psi'(X_1))$ has full rank.

³To be concrete, we shall be considering a Hölder class of function to be defined precisely in Section 6.

3 An Oracle Inequality

In this section we extend Theorem 6.4 of Bühlmann and van de Geer (2011) from ℓ_1 penalty (Lasso) to $\ell_1 + \ell_2^2$ penalty (Elastic Net). This is not a trivial extension since the basic inequality used to establish the result has to be altered considerably. More precisely, the inequality that ties the estimator to the oracle has to be modified. The second difference is that we need to use the adaptive restricted eigenvalue condition which is different from the compatibility condition used in the ℓ_1 -case. Compared to the linear target with quadratic loss in Hebiri and van de Geer (2011) estimated by the elastic net our proof cannot benefit from the augmented regressors idea since this idea relies crucially on the loss function being quadratic. In the case of general convex loss function the proof technique is entirely different and we use the margin condition, Fenchel's inequality, and a careful definition of the oracle instead. We would like to stress that Theorem 1 below is purely deterministic in the sense that there are no probabilities attached to it. It is valid on the set τ to which we shall later attach a lower bound on its probability. It also provides a finite sample result – i.e. the result is valid for any sample size and not just asymptotically.

Theorem 1. *Suppose λ_1 satisfies $\lambda_1 \geq 8\lambda_0$. Then on the set τ , under Assumptions 1-3, we have*

$$\Xi(\hat{f}) + \lambda_1 \|\hat{\beta} - \beta^*\|_1 + \lambda_2 \|\hat{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq 4\Delta^* = 6\Xi(f^*) + 4H \left(\frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi_*} \right),$$

where $H(\cdot)$ is the convex conjugate of the function $G(\cdot)$ in the margin condition.

Note that Theorem 1 provides an upper bound, $4\Delta^*$, on the excess loss of \hat{f} in terms of the excess loss of the oracle f^* as well as an extra term $H(\cdot)$, the estimation error, which is hopefully not too big. We shall comment much more on this extra term in the sequel. Theorem 1 can also be used to give an upper bound on the ℓ_1 -estimation error. Due to its importance the theorem warrants some detailed remarks.

1. The result of Theorem 1 reduces to the result for the Lasso in Theorem 6.4 of Bühlmann and van de Geer (2011) when we set $\lambda_2 = 0$ except for the fact that our adaptive restricted eigenvalue condition is slightly stronger than their compatibility constraint. In that sense we generalize the oracle inequality of Bühlmann and van de Geer (2011). Their oracle inequality is, with ϕ_{**} being a compatibility constant (see p.157, Bühlmann and van de Geer (2011))

$$\Xi(\hat{f}) + \lambda_1 \|\hat{\beta} - \beta^*\|_1 \leq 6\Xi(f^*) + 4H \left(\frac{4\lambda_1 \sqrt{s^*}}{\phi_{**}} \right).$$

As mentioned, the only difference between their Theorem 6.4 and the result that can be deduced from our Theorem 1 is that $\phi_{**} \geq \phi_*$. However, we also carried out the proofs of Theorem 1 imposing the compatibility constraint instead of the adaptive restricted eigenvalue condition and Theorem 1 then reduced to Theorem 6.4 in (Bühlmann and van de Geer (2011)) upon setting $\lambda_2 = 0$. This result can be obtained from the authors on demand. The reason that the adaptive restricted eigenvalue condition is used in the general elastic net is that it gives sharper bounds than the compatibility condition in the general elastic net case. More precisely, if the compatibility condition is used in our case we get an extra $\sqrt{s^*}$ term in front of $4\lambda_2\|\beta^*\|_2$ in the function $H(\cdot)$.

2. Letting β^{BL} denote $\operatorname{argmin}_{\beta} \rho_{f_{\beta}}$, i.e. the best linear approximation, and setting $\Gamma = S_{BL} = \{j : \beta_j^{BL} \neq 0\}$ and choosing $\lambda_2 = \frac{\lambda_1 \sqrt{|\beta^{BL}|}}{2\|\beta^{BL}\|_2}$ it follows that

$$\beta^* = \operatorname{argmin}_{\beta \in S_{BL}} \left\{ 3\Xi(f_{\beta}) + 2H \left(\frac{6\lambda_1 \sqrt{|S_{BL}|}}{\phi(S_{BL})} \right) \right\}$$

Note how we have used our discretion in making a choice of Γ which will turn out to be useful below. Since the second term in the definition of β^* does not depend on β in this case it follows that β^* is the minimizer of $\Xi(f_{\beta})$ which itself also is the minimizer of $\rho_{f_{\beta}}$. Hence, $\beta^* = \beta^{BL}$ in this case. It follows that under the conditions of Theorem 1

$$\Xi(\hat{f}) + \lambda_1 \|\hat{\beta} - \beta^{BL}\|_1 + \lambda_2 \|\hat{\beta}_{S_{BL}} - \beta_{S_{BL}}^{BL}\|_2^2 \leq 6\Xi(f_{\beta^{BL}}) + 4H \left(\frac{6\lambda_1 \sqrt{|S_{BL}|}}{\phi(S_{BL})} \right) \quad (5)$$

So, in particular, Theorem 1 can be used to provide upper bounds on the ℓ_1 -distance of $\hat{\beta}$ to β^{BL} due to our freedom in defining the oracle β^* . If the target f^0 is also linear then clearly the best linear approximation equals the target implying that $\beta^{BL} = \beta^0$ and hence $f_{\beta^{BL}} = f^0$ and $\Xi(f_{\beta^{BL}}) = \Xi(f^0) = 0$. Using $S_{BL} = S^0 = \{j : \beta_j^0 \neq 0\}$, inequality (5) yields

$$\Xi(\hat{f}) + \lambda_1 \|\hat{\beta} - \beta^0\|_1 + \lambda_2 \|\hat{\beta}_{S^0} - \beta_{S^0}^0\|_2^2 \leq 4H \left(\frac{6\lambda_1 \sqrt{|S^0|}}{\phi(S^0)} \right). \quad (6)$$

Hence, in case the target is linear, (6) in particular yields an upper bound on the ℓ_1 -estimation error which does not depend on the excess loss of the oracle. We shall make use of this fact in Section 4 on variable selection. It is also worth pointing out that in practice one does not know S^0 and hence can't choose $\Gamma = S_{BL} = S^0$ in the development of (6). However, (6) is valid even without this knowledge since an even sharper upper bound follows from Theorem 1 by choosing $\Gamma = \{A \subset \{1, \dots, p\} : |A| \leq |S^0|\}$, i.e subsets of $\{1, \dots, p\}$ of cardinality at most $|S^0|$. This bound only relies on sparseness of the target and since S^0 is a member of Γ (6) follows a fortiori.

3. A key issue is to understand the effect of λ_2 on the excess risk. Clearly, the right hand side of Theorem 1 is increasing in λ_2 through $H(\cdot)$ (by Lemma 3 in the appendix $H(\cdot)$ is non-decreasing). However, the same is the case for the left hand side through its multiplication onto the squared ℓ_2 -error. This illustrates a tradeoff in the size of λ_2 .

4. Note also, that the very definition of the restricted set in the definition of the adaptive restricted eigenvalue condition also depends on λ_2 through L_n . In particular, increasing λ_2 increases the size of the set we are minimizing over in the definition of $\phi(S)$. This implies that choosing λ_2 too large may lead to $\phi(S) = 0$, or at least undesirably small values of $\phi(S)$. Note that choosing $\lambda_2 = \frac{\lambda_1 \sqrt{s^*}}{2\|\beta^*\|_2}$ results in $L_n = 6\sqrt{s^*}$. So in this case the size of the restricted set only depends on the cardinality of the oracle. Here it is worth noticing that the sparser the oracle (s^* small) the smaller will the restricted set be and the larger will ϕ_* be.

5. In many econometric examples the margin condition (Assumption 1) is satisfied with a quadratic margin resulting in $H(v) = v^2/4c$ for a positive constant c (as argued just after the definition of the convex conjugate). Setting $\lambda_2 = \frac{\lambda_1 \sqrt{s^*}}{2\|\beta^*\|_2}$ in Theorem 1 results in

$$\Xi(\hat{f}) + \lambda_1 \|\hat{\beta} - \beta^*\|_1 + \lambda_1 \frac{\sqrt{s^*}}{2\|\beta^*\|_2} \|\hat{\beta}_{S^*} - \beta_{S^*}\|_2^2 \leq 6\Xi(f^*) + \frac{36\lambda_1^2 s^*}{c\phi_*^2}.$$

Note that for $\lambda_2 = 0$, corresponding to a pure ℓ_1 -loss, Theorem 1 reduces to

$$\Xi(\hat{f}) + \lambda_1 \|\hat{\beta} - \beta^*\|_1 \leq 6\Xi(f^*) + \frac{16\lambda_1^2 s^*}{c\phi_*^2}.$$

Recall that Theorem 1 and the remarks following it are valid on the set $\tau = \{Z_{M^*} \leq \lambda_0 M^*\} = \{Z_{M^*} \leq \Delta^*\}$. As a consequence, we would like τ to have a large probability. This can be achieved by choosing λ_0 large. However, note that Theorem 1 supposes $\lambda_1 \geq 8\lambda_0$ such that the right hand side of Theorem 1 is also increasing in λ_0 . Put differently, there is a tradeoff between the tightness of the bound in Theorem 1 and the probability with which the bound holds. In the following we shall give a lower bound on the probability of τ which trades off these two effects.

Assume that $\rho_f(x, y) = \gamma(y, f(x)) + c(f)$ where $c(f)$ is a constant possibly depending on f . It will always be zero in our examples. We further assume that there exists a $D > 0$ such that

$$|\gamma(y, f_\beta(x)) - \gamma(y, f_{\tilde{\beta}}(x))| \leq D|f_\beta(x) - f_{\tilde{\beta}}(x)| \quad (7)$$

for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ and $\beta, \tilde{\beta} \in \Phi$. In other words $\gamma(\cdot, \cdot)$ is assumed to be Lipschitz continuous in its second argument over \mathcal{F}_L with Lipschitz constant D . The reason we only need Lipschitz continuity over \mathcal{F}_L is that it is used for a contraction inequality in connection with bounding the

local supremum of the empirical process Z_M in (4). Assume furthermore that

$$\frac{1}{n} \sum_{i=1}^n \max_{1 \leq j \leq p} E\psi_j^2(X_i) \leq 1, \quad (8)$$

and for a positive constant K

$$\max_{1 \leq j \leq p} \|\psi_j\|_\infty \leq K. \quad (9)$$

Note that (9) implies (8) when $K = 1$. Assuming $\max_{1 \leq j \leq p} \|\psi_j\|_\infty \leq K$ is rather innocent since many commonly used basis functions are bounded. As we shall see in Section 6 the most critical assumption in concrete examples is the Lipschitz continuity of the loss function. With this notation in place we state the following result which builds on Theorem 14.5 Bühlmann and van de Geer (2011) (see also Corollary A.1 in van de Geer (2008)).

Theorem 2. *Assume that (7)-(9) are valid and that $p \geq 2$. Assume furthermore, that $\log(p) \leq n$. Then, there exists a constant $d > 0$ such that choosing $\lambda_0 = dD\sqrt{\frac{\log(p)}{n}}$ yields*

$$P(\tau) \geq 1 - \left(\frac{1}{p}\right)$$

The assumption $p \geq 2$ is made for purely technical reasons and does not exclude any interesting problems. Similarly, $\log(p) \leq n$ still allows p to increase at an exponential rate in the sample size.⁴ From an asymptotic point of view Theorem 2 reveals that the measure of the set τ , on which the inequality in Theorem 1 is valid, tends to 1 as $p \rightarrow \infty$. In order to also cover the case of fixed p one can choose $\lambda_0 = dD\sqrt{\frac{\log(p)\log(n)}{n}}$ in Theorem 2 to obtain $P(\tau) \geq 1 - \left(\frac{1}{p}\right)^{\log(n)}$ by a slight modification of the proof of Theorem 2 which tends to one as $n \rightarrow \infty$ even for p fixed. But since Theorem 1 requires $\lambda_1 \geq 8\lambda_0$ this will yield a bigger upper bound in that Theorem since by Lemma 3 in the Appendix $H(\cdot)$ is a non-decreasing function. Combining Theorems 1 and 2 yields the following result.

Theorem 3. *Under the assumptions of Theorems 1 and 2 with the choice of $\lambda_0 = dD\sqrt{\frac{\log(p)}{n}}$ for some positive constant d it holds with probability at least $1 - \left(\frac{1}{p}\right)$*

$$\Xi(\hat{f}) + \lambda_1 \|\hat{\beta} - \beta^*\|_1 + \lambda_2 \|\hat{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq 4\Delta^* = 6\Xi(f^*) + 4H\left(\frac{4\lambda_1\sqrt{s^*} + 4\lambda_2\|\beta^*\|_2}{\phi_*}\right)$$

where $H(\cdot)$ is the convex conjugate of $G(\cdot)$.

⁴This is from an asymptotic point of view, though we wish to emphasize that the inequality in Theorem 2 holds for any given sample size (satisfying the conditions of the theorem).

Theorem 3 basically consists of the inequality in Theorem 1 with a lower bound attached to the measure of the set τ on which Theorem 1 is valid. In particular, the theorem reveals that the excess loss of, $\Xi(\hat{f})$, will not be much larger than the one of the oracle. The second term on the right hand side reflects the estimation error. Put differently, the excess loss of our estimator depends on the excess loss of the oracle as well as the distance to the oracle. From an asymptotic point of view Theorem 2 reveals that the measure of the set τ on which the inequality in Theorem 1 is valid tends to 1 as $p \rightarrow \infty$. In particular, we have the following result for the asymptotic excess loss of \hat{f} . To this end assume that $p \in O(\exp(n^a))$ and $|S^*| \in O(n^b)$ for some $a > 0$ and $b \geq 0$.

Corollary 1. *Assume that $\lambda_1 \leq L\lambda_0$ for some $L \geq 8$. Then, under the assumptions of Theorems 1 and 2, with $\lambda_2 = \lambda_1\sqrt{s^*}/2\|\beta^*\|_2$ and the choice of $\lambda_0 = dD\sqrt{\frac{\log(p)}{n}}$ for some positive constant d one has*

$$\limsup_{n \rightarrow \infty} \Xi(\hat{f}) \leq 6 \limsup_{n \rightarrow \infty} \Xi(f^*)$$

with probability approaching one if $a + b < 1$ and ϕ_* is bounded away from zero.

Corollary 1 shows that asymptotically the excess loss of \hat{f} will be of the same order as that of the oracle. This is useful since we saw in remark 2 above that we have considerable discretion in choosing f^* and hence in what we bound $\Xi(\hat{f})$ by from above in Corollary 1. In the case where f^0 is linear we know from Remark 2 above that we can choose Γ such that $f^* = f^0$ and hence $\Xi(f^*) = 0$. In this case Corollary 1 actually reveals that the excess loss of \hat{f} tends to zero. We next investigate the case of linear f^0 in more detail.

3.1 Linear target

In the case where the target function f^0 is linear Theorem 3 can be used to deduce the following result.

Corollary 2. *Assume that f^0 is linear.*

a) *Then, under the assumptions of Theorems 1 and 2*

$$\Xi(\hat{f}) \leq 4H \left(\frac{4\lambda_1\sqrt{|S^0|} + 4\lambda_2\|\beta^0\|_2}{\phi(S^0)} \right) \tag{10}$$

$$\|\hat{\beta} - \beta^0\|_1 \leq \frac{4}{\lambda_1} H \left(\frac{4\lambda_1\sqrt{|S^0|} + 4\lambda_2\|\beta^0\|_2}{\phi(S^0)} \right) \tag{11}$$

with probability at least $1 - \left(\frac{1}{p}\right)$.

b) If, furthermore, the margin is quadratic such that $H(v) = v^2/(4c)$ for some $c > 0$, and we choose $\lambda_2 = \frac{\lambda_1 \sqrt{|S^0|}}{2\|\beta^0\|_2}$ as well as $\lambda_1 \leq L\lambda_0$ for some $L \geq 8$, then, by (10) and (11)

$$\Xi(\hat{f}) \leq \frac{36 \lambda_1^2 |S^0|}{c \phi^2(S^0)} \leq \frac{36 (LdD)^2 \log(p) |S^0|}{c n \phi^2(S^0)} \quad (12)$$

$$\|\hat{\beta} - \beta^0\|_1 \leq \frac{36L \lambda_1 |S^0|}{c \phi^2(S^0)} \leq \frac{36LdD}{c} \sqrt{\frac{\log(p)}{n}} \frac{|S^0|}{\phi^2(S^0)} \quad (13)$$

with probability at least $1 - \left(\frac{1}{p}\right)$.

The bounds (10) and (11) bound the ℓ_1 -estimation error of the elastic net estimator for any type of loss function satisfying the conditions of Theorem 1. Note that there is no excess loss from the oracle entering in the upper bound. This is due to the fact that this is zero when the target is linear. The last two bounds in Corollary 2 specialize to the case where the quadratic margin condition is satisfied. We stress again, as we shall see later (see Section 6), that the margin condition is indeed quadratic in many econometric examples.

Furthermore, one sees from the above Corollary that the rate of convergence of the elastic net estimator in the ℓ_1 -norm is $\frac{\sqrt{n}}{\sqrt{\log(p)}} \frac{1}{|S^0|}$ provided that the adaptive restricted eigenvalue is bounded away from zero.

Furthermore, (13) can be used to deduce consistency of $\hat{\beta}$ for β^0 . As in Corollary 1 assume that $p = O(\exp(n^a))$ and $|S^0| = O(n^b)$ for some $a > 0$ and $b \geq 0$.

Corollary 3. *Assume that f^0 is linear and set $\lambda_2 = \frac{\lambda_1 \sqrt{|S^0|}}{2\|\beta^*\|_2}$ as well as $\lambda_1 \leq L\lambda_0$ for some $L \geq 8$. Let the quadratic margin condition be satisfied and assume furthermore that $\phi^2(S^0)$ is bounded away from 0. Then, under the assumptions of Theorems 1 and 2,*

$$\|\hat{\beta} - \beta^0\|_1 \xrightarrow{p} 0$$

if $a + 2b < 1$.

Corollary 3 shows that the elastic net can be consistent even when the dimension p increases at a subexponential rate in the sample size. Note, however, that the number of relevant variables, $|S^0|$ can not increase faster than the square root of the sample size (a can be put arbitrarily close to 0 to see this). Hence, even though the total number of variables can be very large, the number of relevant variables must still be quite low. This is in line with previous findings for the linear model in the literature. We also remark that this requirement is slightly stricter than the one needed when considering the excess loss in Corollary 1 (in that corollary we only needed $a + b < 1$). Also

note that the conditions in Corollary 3 are merely sufficient. For example one can let $\phi^2(S^0)$ tend to zero at the price of reducing the growth rate of p and $|S^0|$.

4 Variable Selection

In this section we briefly comment on how the results in Section 3 can be used to perform consistent variable selection in the case where f^0 is a linear function⁵. First note, that the results in Corollaries 2 and 3 can be used to provide rates of convergence of $\hat{\beta}$ for β^0 in the ℓ_1 -norm in the case of a linear target. If one furthermore assumes that $\min\{|\beta_j^0| : \beta_j^0 \neq 0\}$ is bounded away from zero by at least the rate of convergence of $\hat{\beta}$ it follows by standard arguments, see Lounici (2008) or Kock and Callot (2013), that no non-zero β^0 will be classified as such. Put differently, the elastic net possesses the *screening property*.

In order to remove all non-zero variables one may furthermore threshold the elastic net estimator by removing all variables with parameters below a certain threshold. Again standard arguments show that choosing the threshold of the order of the rate of convergence (details omitted) can yield consistent model selection asymptotically. Since thresholding is a generic technique which is not specific to our setup we shall not elaborate further on this at this stage.

One technical remark is in its place at this point. Since thresholding is done at the level of the individual parameter, what one really needs is an upper bound on the estimation error for each individual parameter. In other words, an upper bound on the sup-norm, $\max_{1 \leq j \leq p} |\hat{\beta}_j - \beta_j^0|$, is sought. However, our results in the previous section provide rates of convergence in the much stronger ℓ_1 -norm⁶. Of course one may simply use the ℓ_1 -rates of convergence to upper bound the sup-norm rates of convergence. But this is suboptimal. Alternatively, a strengthening of the adaptive restricted eigenvalue condition can yield rates of convergence in the ℓ_2 - or the sup-norm which are of a lower order of magnitude than the corresponding ℓ_1 -results. Upon request we can make results for the thresholded elastic net based on upper bounds on the ℓ_2 -norm rate of convergence available. We have omitted the results here since thresholding is a rather standard technique.

⁵If the target function is not linear we do not find it sensible to talk about consistent variable selection in a linear approximation of the target. Hence, this section restricts attention to the case where the target is linear.

⁶On finite-dimensional vector spaces these two norms are equivalent but here we are working in a setting where the dimension, p , tends to infinity.

5 Tuning Parameter Selection

Recently, Fan and Tang (2013) developed a method to select tuning parameters in high dimensional generalized linear models with more parameters than observations. Here we briefly describe their method with the terminology translated into our setting⁷. For a loss function of the form

$$\rho_\beta(Z_i) = -Y_i\beta'X_i + b(\beta'X_i) - c(Y_i, \xi), \quad (14)$$

where $b(\cdot)$ and $c(\cdot, \cdot)$ are known functions and ξ is a known scale parameter. The tuning parameter, λ , (as a generic tuning parameter) is chosen to minimize the Generalized Information Criterion (GIC)

$$GIC(\lambda) = 2 \left[\frac{1}{n} \sum_{i=1}^n \rho_{\hat{\beta}_\lambda}(Z_i) - \frac{1}{n} \sum_{i=1}^n \rho_{\text{saturated}}(Z_i) \right] + \frac{\log(\log(n)) \log(p)}{n} s_\lambda.$$

where $\rho_{\text{saturated}}$ corresponds to the saturated model as defined in Fan and Tang (2013) and $\hat{\beta}_\lambda$ is the elastic net minimizer corresponding to the penalty parameter λ . s_λ is the number of non-zero coefficients for a given value of λ .

Theorems 1-2 and Corollary 1 in Fan and Tang (2013) establish the consistency of this approach, i.e. $P(\text{supp}(\hat{\beta}_\lambda) = \text{supp}(\beta^0)) \rightarrow 1$ where $\text{supp}(v)$ indicates the support of the vector v , i.e. the location of its non-zero entries. Hence, GIC will, asymptotically, select the correct model (provided there exists a λ_0 for which $\text{supp}(\hat{\beta}_\lambda) = \text{supp}(\beta^0)$). A necessary condition for this to be meaningful is of course that the target is linear, i.e. $f^0(x) = \beta^0'x$ for some $\beta^0 \in \mathbb{R}^p$, but nothing prevents one from using GIC even when the target is non-linear. The theoretical merits of the procedure are, to our knowledge, unknown in that case.

At this point it is also worth mentioning that the quadratic loss is covered by (14) since

$$\frac{1}{2} (y - \beta'x)^2 = \frac{1}{2}y^2 + \frac{1}{2}(\beta'x)^2 - y(\beta'x)$$

such that we may choose $b(u) = \frac{1}{2}u^2$ and $c(y, \xi) = -\frac{1}{2}y^2$. Similarly, the logistic loss is covered since

$$\rho_f(y, x) = -y(\beta'x) + \log(1 + e^{\beta'x})$$

such that one may choose $b(u) = \log(1 + e^u)$ and $c(\cdot, \cdot) = 0$. This covers a large class of convex functions as in van de Geer (2008). Regarding the choice of threshold for variable selection we suggest using the procedure outlined in Caner and Knight (2013).

⁷Fan and Tang (2013) consider log-likelihood functions of generalized linear models. However, as our loss functions can often be written as the negative of the log-likelihood (as seen in Section 2), their setup applies to many of our examples.

6 Econometric Examples

In this section we present sufficient conditions for Assumptions 1-3 to be valid for concrete econometric examples. Recall that these Assumptions are sufficient for Theorem 1 to be valid. When necessary, we also comment on sufficient conditions for the assumptions underlying Theorem 2. As already argued in connection with this theorem the critical assumption is the Lipschitz continuity of the loss function while the boundedness assumption on the basis functions is rather innocuous. As a consequence, we will focus on verification of the Lipschitz continuity whenever this is not trivially satisfied. First, we present a general sufficient condition for loss functions to satisfy a quadratic margin condition. This condition is then exemplified on a couple of examples. In the following $\|\cdot\|$ denotes the $L_2(P)$ -norm on \mathbf{F} and the data is supposed to come from an i.i.d. sample.

Assume that the loss function is of the form $\rho_f(x, y) = \rho(f(x), y)$ such that it only depends on x through $f(x)$. By Doob's representation, see Lemma 1.13 in Kallenberg (2002), we define

$$l(f(X), X) := E[\rho(f(X), Y)|(X, f(X))]$$

Furthermore, by iterated expectations, it suffices to show that $l(f(x), x)$ satisfies the margin condition in order to verify that this is the case for $\rho_f(x, y) = \rho(f(x), y)$. The target function will be

$$f^0(x) = \operatorname{argmin}_{f \in \mathbf{F}} l(f(x), x)$$

which is the minimizer of the loss function and hence a natural choice. Next, fix an $x \in \mathcal{X}$ and assume that $l(a, x)$ is twice continuously differentiable in its first argument in an $\|\cdot\|_\infty$ -neighborhood of radius $\eta > 0$ around $f^0(x)$ with second derivative bounded from below by $2c > 0$, for $c > 0$. Then it follows by Lagrange's form of the remainder term in a Taylor series that for some \tilde{a} on the line segment joining $f(x)$ and $f^0(x)$

$$\begin{aligned} l(f(x), x) &= l(f^0(x), x) + l'_1(f^0(x), x)(f(x) - f^0(x)) + \frac{l''_{11}(\tilde{a}, x)}{2}(f(x) - f^0(x))^2 \\ &\geq l(f^0(x), x) + c(f(x) - f^0(x))^2 \end{aligned}$$

for all $f \in \mathbf{F}$ such that $|f(x) - f^0(x)| < \eta$, and $l'_1(\cdot, \cdot)$, $l''_{11}(\cdot, \cdot)$ represent the first and second order partial derivatives of l with respect to its first argument. Assuming this is valid for all $x \in \mathcal{X}$ implies that

$$l(f(x), x) - l(f^0(x), x) \geq c(f(x) - f^0(x))^2 \tag{15}$$

for all $f \in \mathbf{F}_{\text{local}} = \{f \in \mathbf{F} : \|f - f^0\|_\infty \leq \eta\}$. This yields, using that the $(X_i, Y_i)_{i=1}^n$ are identically distributed,

$$\Xi(f) = P\rho_f - P\rho_{f^0} = [E\rho_f(X_1, Y_1) - E\rho_{f^0}(X_1, Y_1)] = E[l(f(X_1, X_1) - l(f^0(X_1), X_1))] \geq c\|f - f^0\|^2$$

for all $f \in \mathbf{F}_{\text{local}}$ such that the margin condition is satisfied with $G(x) = cx^2$. Put differently, the above shows that it suffices to establish a lower bound on the second derivative of the conditional expectation of the loss function in order to show that the margin condition holds with quadratic margin.

We shall next use the above result to verify that some typical loss functions encountered in econometrics satisfy Assumptions 1-3.

6.1 Quadratic loss

Assume that the data is generated by the i.i.d. sequence

$$Y_i = f^0(X_i) + \epsilon_i \tag{16}$$

for $X_i \in \mathcal{X} \subseteq \mathbb{R}$.⁸ We show that the quadratic loss function

$$\rho(f(x), y) = (y - f(x))^2 \tag{17}$$

for $f \in \mathbf{F}$ can be encompassed by our general theory. The quadratic loss function is probably the most widely used loss in regression analysis. The main obstacle in fitting this type of loss into our general theory is that Theorems 2 and 3 rely on the loss function being Lipschitz continuous in order to lower bound the probability of the event τ . However, the quadratic loss is only locally Lipschitz continuous. As we shall see, this can still deliver an oracle inequality which holds with high probability if the covariates and the error terms don't have too heavy tails and the target function f^0 is bounded on compact subsets of \mathbb{R} . Before stating the first result note that for any $f \in \mathbf{F}$, if ϵ_1 is independent of X_1

$$\Xi(f) = E(Y_1 - f(X_1))^2 - E(Y_1 - f^0(X_1))^2 = E(f(X_1) - f^0(X_1))^2 \tag{18}$$

such that the excess loss reduces to the mean square error (MSE) in case of a quadratic loss function. This turns out to be particularly useful when discussing the relation of our procedure to

⁸It is not difficult to generalize this to \mathcal{X} being some subset of a normed space by modifying the definition of sub-gaussianity in footnote 9 below slightly.

nonparametric series estimation in Section 6.1.1 below. (18) also reveals that the margin condition is satisfied with a quadratic margin, even without using the technique from above, which implies that Assumption 1 is satisfied.

Remark: If the sample is merely assumed to be independently but not necessarily identically distributed the above calculations still go through with some small modifications. We explain how next. First, let the norm on \mathbf{F} be $\|f\| = \sqrt{\sum_{i=1}^n \frac{1}{n} E f^2(X_i)}$. Note that this reduces to the norm from the identically distributed case in case the variables are actually identically distributed. Assuming that ϵ_i is independent of X_i for all $i = 1, \dots, n$ one gets by the same arguments as above

$$\Xi(f) = \frac{1}{n} \sum_{i=1}^n E(Y_i - f(X_i))^2 - \frac{1}{n} \sum_{i=1}^n E(Y_i - f^0(X_i))^2 = \frac{1}{n} \sum_{i=1}^n E(f(X_i) - f^0(X_i))^2$$

which shows that the margin condition is satisfied with a quadratic margin function even in the case where the sample is not identically distributed. In particular, this means that the bounds to follow, including the nonparametric ones in Section 6.1.1, are valid in the presence of heteroscedasticity. For the sake of exposition we shall now return to the i.i.d. situation.

Lemma 1. *Assume that $\max_{1 \leq j \leq p} \|\psi_j\|_\infty \leq K$. Set $\lambda_2 = \frac{\lambda_1 \sqrt{s^*}}{2\|\beta^*\|_2}$ and take for some constant $L \geq 8$, $8\lambda_0 \leq \lambda_1 \leq L\lambda_0$ and suppose that*

$$\frac{36L^2\lambda_0 s^* K}{\eta\phi_*^2} \leq 1. \quad (19)$$

In addition, assume that $\|f^ - f^0\|_\infty \leq \eta/2$, with $\frac{3}{2}E(f^* - f^0)^2 \leq \frac{9\lambda_1^2 s^*}{\phi_*^2}$ and finally that the population covariance matrix of the covariates $\Sigma = E(\psi(X)\psi'(X))$ has full rank. Then,*

a) *Assumptions 1-3 are valid and on the set τ*

$$E(\hat{f} - f^0)^2 + \lambda_1 \|\hat{\beta} - \beta^*\|_1 + \lambda_2 \|\hat{\beta}_{S^*} - \beta_{S^*}\|_2^2 \leq 6E(f^* - f^0)^2 + \frac{36L^2\lambda_0^2 s^*}{\phi_*^2}. \quad (20)$$

b) *if, furthermore $\frac{1}{n} \sum_{i=1}^n \max_{1 \leq j \leq p} E\psi_j^2(X_i) \leq 1$, $\sup_{|x| \leq C_n} |f^0(x)| \leq F_{C_n} < \infty$ for all $C_n > 0$, $\Phi = \{\beta \in \mathbb{R}^p : \|\beta\|_1 \leq G < \infty\}$ and X_1, ϵ_1 are sub-gaussian⁹ one has for all $C_n > 0$ that (20) is valid with probability at least $1 - \left(\frac{1}{p}\right) - 2\alpha n \exp(-\delta C_n^2)$ for $\lambda_0 = dD_n \sqrt{\frac{\log(p)}{n}}$ with $D_n = 2(C_n + 2F_{C_n} + GK)$ and $d > 0$.*

Lemma 1 a) provides sufficient conditions for Assumptions 1-3, and a bound similar to the one in Theorem 1 which is valid on the set τ . Lemma 1 b) also provides a lower bound on the

⁹A real random variable V is said to be sub-gaussian if $P(|V| \geq x) \leq \alpha \exp(-\delta x^2)$ for some positive constants α and δ .

probability of the set τ in the case of a non-Lipschitz continuous loss function. Some of the assumptions of Lemma 1 may still seem rather high level but they are not very restrictive. We next give a concrete example of when they are satisfied. In particular, we shall assume that x has support in $[-1, 1]$. This can of course be achieved by a suitable (sigmoidal) transformation of the covariates. f^0 will be assumed to belong to a Hölder class of order $1/2 < r < \infty$ (or it is r -smooth in the terminology of Chen (2007)¹⁰). In this case one may choose \mathcal{F}_L to consist of p th degree polynomials, i.e. $\psi_j(x) = x^j$, $j = 1, \dots, p$, and $\Gamma = \{1, \dots, s^*\}$. Using other basis functions than polynomials is of course a possibility if one does not want to assume that the covariates are contained in $[-1, 1]$. However, this is a rather innocent assumption since it can be obtained by transforming the covariates. Alternatively, one may use a basis of bounded functions. In any case, it is not our purpose to promote one basis over another. We simply show that polynomials suffice for our theory to work.

The assumption $\max_{1 \leq j \leq p} \|\psi_j\|_\infty \leq K$ requires the basis functions ψ_j to be bounded. Since the covariates have support in $[-1, 1]$ it follows that $\|\psi_j\|_\infty = \|x^j\|_\infty = 1$ for all $j = 1, \dots, p$ and hence the assumption is satisfied with $K = 1$. As mentioned right after (9) this also implies $\frac{1}{n} \sum_{i=1}^n \max_{1 \leq j \leq p} E\psi_j^2(X_i) \leq 1$ which is needed in part b) of Lemma 1.

The approximation requirements $\|f^* - f^0\|_\infty \leq \eta/2$ and $\frac{3}{2}E(f^* - f^0)^2 \leq \frac{9\lambda_1^2 s^*}{\phi_*^2}$ state that the target should be approximated well by the oracle and imply that $f^* \in \mathbf{F}_{\text{local}}$. As explained previously, one may choose $f^* = f^0$ when f^0 is linear. Hence, the two approximation requirements are trivially satisfied in the linear case. In general the validity of these high level assumptions depend on the choice of \mathbf{F} and the existence of properly approximating basis functions $\{\psi_j\}_{j=1}^\infty$. As mentioned, we shall assume that the covariates are compactly supported. In general, assuming that f^0 belongs to a Hölder class of order $1/2 < r < \infty$, choosing \mathcal{F}_L to consist of p th degree polynomials, i.e. $\psi_j(x) = x^j$, $j = 1, \dots, p$, and $\Gamma = \{1, \dots, s^*\}$ it follows from page 5573 in Chen (2007) that $\|f^* - f^0\|_\infty \in O(s^{*-r})$. Hence, $\|f^* - f^0\|_\infty \leq \eta/2$ will be satisfied for any $\eta > 0$ as long as s is sufficiently large. Furthermore, since

$$E(f^* - f^0)^2 = \|f^* - f^0\|^2 \leq \|f^* - f^0\|_\infty^2 \in O(s^{*-2r}) \quad (21)$$

one also has that $\frac{3}{2}E(f^* - f^0)^2 \in O(s^{*-2r})$ such that $\frac{3}{2}E(f^* - f^0)^2 \leq \frac{9\lambda_1^2 s^*}{\phi_*^2}$ if s^* is at least of the order $\lambda_1^{-2/(2r+1)}$, i.e. $s^* \in \Omega(\lambda_1^{-2/(2r+1)})$, and ϕ_* is assumed to be bounded away from 0.

¹⁰Following Chen (2007), we say that a function f is r -smooth if it has derivatives up to order $\lfloor r \rfloor$ and the r -th derivative is Hölder continuous with an exponent of $r - \lfloor r \rfloor$. Here $\lfloor r \rfloor$ denotes the greatest integer strictly less than r . This is also often called a Hölder class of order r , see eg Tsybakov (2009).

The condition (19) is not restrictive either. Assuming that ϕ_* is bounded away from 0 it suffices to show that $\lambda_0 s^* \rightarrow 0$. Choosing $s^* \in \Theta(\lambda_1^{-2/(2r+1)})$, which is in accordance with $s^* \in \Omega(\lambda_1^{-2/(2r+1)}) = \Omega(\lambda_0^{-2/(2r+1)})$ (since λ_1 and λ_0 are of the same order), yields

$$\lambda_0 s^* \in \Theta(\lambda_0 \lambda_0^{-2/(2r+1)}) = \Theta(\lambda_0^{(2r-1)/2r}). \quad (22)$$

Hence, if $r > 1/2$, it is enough that $\lambda_0 \rightarrow 0$. Since the covariates are assumed to have compact support and f^0 is more than $\frac{1}{2}$ -smooth, and hence continuous, one has F_{C_n} is bounded and so $\lambda_0 \in \Theta(D_n \sqrt{\frac{\log(p)}{n}}) = \Theta(C_n \sqrt{\frac{\log(p)}{n}})$ (as seen next we think of C_n as an increasing sequence). For asymptotic considerations an obvious choice of C_n is $C_n = \sqrt{\frac{2}{\delta} \log(n)}$ since this ensures that $2\alpha n \exp(-\delta C_n^2)$ tends to zero in the lower bound on the probability with which (20) holds in part b) of Lemma 1. With this choice, $\lambda_0 \in \Theta(\sqrt{\log(n)} \sqrt{\frac{\log(p)}{n}})$ and it suffices that $\sqrt{\log(n)} \sqrt{\frac{\log(p)}{n}} \rightarrow 0$ which still permits p to increase at an almost exponential rate in n . Note also at this point that since $s^* \in \Theta(\lambda_1^{-2/(2r+1)})$, $\lambda_0 \in \Theta(\sqrt{\log(n)} \sqrt{\frac{\log(p)}{n}})$ implies $s^* \in \Theta([\frac{n}{\log(p)\log(n)}]^{1/(2r+1)})$. So the more smooth (r large) the function f^0 is the smaller can we afford to choose s^* since smooth functions are easier to approximate. This is of course desirable since this means few coefficients have to be estimated which results in a smaller estimation error.

Assuming Σ to have full rank is reasonably innocent as discussed in connection with the adaptive restricted eigenvalue condition in Section 2.

The additional assumptions of part b) of Lemma 1 are used to establish a lower bound on the probability of τ on which (20) is valid in the absence of Lipschitz continuity of the loss function. First, $\frac{1}{n} \sum_{i=1}^n \max_{1 \leq j \leq p} E\psi_j^2(X_i) \leq 1$, imposes a further boundedness assumption on the basis functions, which, as discussed above, is trivially satisfied since $K = 1$. The assumption $\sup_{|x| \leq C_n} |f^0(x)| \leq F_{C_n} < \infty$ requires f^0 to be locally bounded and is satisfied in particular if f^0 is continuous. Continuity of f^0 is of course ensured if the target is linear or r -smooth for $r > 0$ (the latter being covered by our working example). Hence, our theory covers the case of a linear target with quadratic loss. If, as in the discussion above, we continue to assume that the covariates are compactly supported F_{C_n} can even be chosen to be an absolute constant independent of the sample size (as already mentioned just after (22)). The assumption $\|\beta\|_1 \leq G$ is rather innocent since G can be chosen arbitrarily large¹¹. The sub-gaussianity assumption on the covariates and the error terms prevents them from having too heavy tails. This is quite common assumption in the literature on high-dimensional models. When the covariates are bounded they are, a fortiori,

¹¹It is also straightforward to generalize to the setting where G is a sequence depending on the sample size.

subgaussian. In conclusion, all conditions of Lemma 1 are valid when f^0 belongs to a Hölder class of order at least $1/2$ and the covariates have support $[-1, 1]$. In particular, choosing \mathcal{F}_L to consist of p th degree polynomials, i.e. $\psi_j(x) = x^j$, $j = 1, \dots, p$, and $\Gamma = \{1, \dots, s^*\}$ is sufficient.

Remark: Even though the sub-gaussianity assumptions on the covariates and the error terms is a standard one we would like to stress that our theory is also applicable for much heavier tails. A version of Lemma 1 b) can be developed even when we only have $E|X_1|^r$, $E|\epsilon_1|^r \leq \kappa < \infty$ for some $r \geq 2$. As long as the covariates and error terms possess enough moments Lemma 1 b) still applies. More precisely, a slight change of the last part of the proof of part b) yields that $1 - \left(\frac{1}{p}\right) - \frac{\kappa n}{C_n^r}$. So the price to pay for the increased generality is that the last term in the lower bound on the probability no longer tends to 0 exponentially fast in C_n such that C_n now has to be at least of the order $O(n^{1/r})$.

6.1.1 Relation to non-parametric series estimators

At this point it is worth pointing out that our theory encompasses non-parametric series estimation as a special case since $\hat{f}(x) = \sum_{j=1}^p \hat{\beta}_j \psi_j(x)$ can be seen as a series estimator of the unknown function $f^0(x)$. Note first that our results are mainly non-asymptotic while we are not aware of any finite sample results for classical series estimators. Apart from this, our estimator has another big advantage compared to the classical series estimator of e.g. Newey (1997): In conventional series estimators of the form $\sum_{j=1}^T \hat{\beta}_j \psi_j(x)$ the number of terms T in the series can usually only increase slower than the sample size such as $T \in o(n^{1/2})$ or $T \in o(n^{1/3})$ in Newey (1997). However, as we have already hinted at several times, the number of series terms p in our estimator can increase exponentially fast in asymptotic considerations. We shall be more precise about this in the discussion of Corollary 5 below. What we require instead is that the number of *non-zero* terms, s , in our series estimator increases slowly. As remarked already, though in a slightly different context, by Belloni and Chernozhukov (2011) this is still a considerable generalization since we do not require that it is the *first* s terms in $f^*(x) = \sum_{j=1}^p \beta_j^* \psi_j(x)$ that have the non-zero coefficients. Put differently, we do not assume any prior knowledge of the location of the non-zero coefficients of the oracle estimator. To be concrete, assume that $\mathbf{F} = L_2(P_{X_1})$ with the usual inner product $\langle \cdot, \cdot \rangle$ and P_{X_1} being the distribution of X_1 . This is a Hilbert space and hence there exists an orthonormal basis $\{g_j\}_{j=1}^\infty$ such that every $f \in \mathbf{F}$ can be written as

$$f(x) = \sum_{j=1}^{\infty} \langle f, g_j \rangle g_j(x)$$

Using the first s terms of this series representation as an approximation of f yields the approximation $\tilde{f}(x) = \sum_{j=1}^s \langle f, g_j \rangle g_j(x)$ with corresponding approximation error $\sum_{j=s+1}^{\infty} \langle f, g_j \rangle^2$. More generally, letting \mathcal{I}_s be any subset of cardinality s of the natural numbers, one could instead consider the approximation $\check{f}(x) = \sum_{j \in \mathcal{I}_s} \langle f, g_j \rangle g_j(x)$ with approximation error $\sum_{j \in \mathcal{I}_s^c} \langle f, g_j \rangle^2$. Clearly, choosing \mathcal{I}_s to consist of the s largest $\langle f, g_j \rangle^2$ yields the smallest approximation error. However, there is no reason to believe that the s largest terms in $\{\langle f, g_j \rangle^2\}_{j=1}^{\infty}$ are also the s first ones, in fact their location is generally unknown and hence our method is useful since it allows for $p \gg s$ candidate terms. Put differently, we choose \mathcal{I}_s from $\{A \subset \{1, \dots, p\} : |A| = s\}$ and clearly $\{1, \dots, s\} \in \{A \subset \{1, \dots, p\} : |A| = s\}$. This is considerably more general than simply choosing $\mathcal{I}_s = \{1, \dots, s\}$ since $s \ll p$.

Note also that Caner and Zhang (2013) have shown that the choice of dimension matters in structural function estimation via sieves. As we shall see now, as long as f^0 can be "well approximated" by some sparse linear combination $f^*(x) = \sum_{j=1}^p \beta_j^* \psi_j(x)$, part b) of Lemma 1 will provide a *finite* sample upper bound on the estimation error of \hat{f} (in an L_2 -sense). In particular, part b) of Lemma 1 yields

Corollary 4. *a) Let the assumptions of Lemma 1 be satisfied. Then*

$$E(\hat{f} - f^0)^2 \leq 6E(f^* - f^0)^2 + \frac{36L^2\lambda_0^2 s^*}{\phi_*^2}. \quad (23)$$

with probability at least $1 - \left(\frac{1}{p}\right) - 2\alpha n \exp(-\delta C_n^2)$.

b) Assume furthermore that the covariates have support $[-1, 1]$, f^0 is r -smooth for $1/2 < r < \infty$ (belongs to a Hölder class of order r as defined in footnote 10) and we choose \mathcal{F}_L to consist of polynomials of degree s^ . Then if Φ is compact, ϵ_1 sub-gaussian and ϕ_* bounded away from 0.*

$$E(\hat{f} - f^0)^2 \leq O(s^{*-2r}) + \frac{36L^2\lambda_0^2 s^*}{\phi_*^2}. \quad (24)$$

with probability at least $1 - \left(\frac{1}{p}\right) - 2\alpha n \exp(-\delta C_n^2)$.

c) Maintain the assumptions of b). Choosing $s^ \in \Theta\left(\lambda_0^{\frac{-2}{2r+1}}\right)$ yields*

$$E(\hat{f} - f^0)^2 \leq O\left(\lambda_0^{\frac{4r}{2r+1}}\right) = O\left(C_n^{\frac{4r}{2r+1}} \left[\frac{\log(p)}{n}\right]^{2r/(2r+1)}\right) \quad (25)$$

with probability at least $1 - \left(\frac{1}{p}\right) - 2\alpha n \exp(-\delta C_n^2)$.

d) Finally, if in addition one chooses $C_n = \sqrt{\frac{2}{\delta} \log(n)}$ ¹², one has

$$E(\hat{f} - f^0)^2 \in O\left(\log(n)^{\frac{2r}{2r+1}} \left[\frac{\log(p)}{n}\right]^{2r/(2r+1)}\right) \quad (26)$$

with probability at least $1 - \frac{1}{p} - \frac{2\alpha}{n}$.

(23) of Corollary 4 is remarkable since it provides a finite sample upper bound on the mean square error (MSE) of the Elastic Net series estimator \hat{f} of f^0 . The upper bound on the MSE depends on two terms. First, $\Xi(f^*)$, the approximation error stemming from approximating f^0 by f^* . The second term, $\frac{36L^2\lambda_0^2s^*}{\phi_*^2}$, can be interpreted as the estimation error of $\hat{\beta}$ in estimating the oracle parameter β^* . The larger the number of coefficients to be estimated, the larger will the estimation error be, but the approximation error will be smaller. In (24) we restrict attention to bounded covariates (which are, a fortiori, sub-gaussian) and smooth functions. This allows us to give an explicit order of magnitude on $E(f^* - f^0)^2$ as explained in the discussion after Lemma 1. Here it is worth mentioning that the order of approximation error is based on results of approximation by polynomials of degree s^* . However, our discussion prior to Corollary 4 indicates that it is not necessarily a linear combination of x, x^2, \dots, x^{s^*} which yields the best approximation and so it might be possible to further improve the bound in (24). But this is a topic in approximation theory which we do not feel is appropriate to pursue here. However, our estimator may perform much better in practice than the plain series estimator.

(25) follows from (24) by choosing s^* to minimize the order of the upper bound in (24). We remark here that this is exactly in accordance with the choice of s^* in connection with (21) and (22) which ensured the validity of the approximability conditions in Lemma 1 since λ_1 and λ_0 are of the same order. Note that (26) gives an explicit order of the estimation error in terms of p and n only. It reveals that the L_2 -distance between \hat{f} and f^0 tends to zero even when p is allowed to increase almost exponentially fast in n . We also remark that by utilizing the boundedness of the covariates in the proof of part b) of Lemma 1 the lower bound on the probability with which (24) and (25) are valid may be increased to $1 - \left(\frac{1}{p}\right) - \alpha n \exp(-\delta C_n^2)$. This can further be increased to $1 - \frac{1}{p}$ if the error terms are bounded as well. The corresponding lower bounds on the probability with which (26) is valid may similarly be increased to $1 - \frac{1}{p} - \frac{\alpha}{n}$ and $1 - \frac{1}{p}$, respectively.

¹²Where $\delta > 0$ is the subgaussianity constant of footnote 9.

6.1.2 Asymptotic results for quadratic loss

We now return to the general setting of Lemma 1. To illustrate the usefulness of Lemma 1 we remark that it can be used to establish the following asymptotic result. As in Corollary 3 we assume that $p \in O(\exp(n^a))$ and $s^* \in O(n^b)$ for some $a > 0$ and $b \geq 0$. Furthermore, we assume that $F_{C_n} \in O(n^{\tilde{d}})$ for some $\tilde{d} \geq 0$.

Corollary 5. *Let the assumptions of Lemma 1 be satisfied. Choose $C_n = \sqrt{\frac{2}{\delta} \log(n)}$ ¹³. Then, if ϕ_*^2 is bounded away from 0, one has with probability tending to one,*

a)

$$\limsup_{n \rightarrow \infty} E(\hat{f} - f^0)^2 \leq 6 \limsup_{n \rightarrow \infty} E(f^* - f^0)^2$$

if $a + b + 2\tilde{d} < 1$.

b)

$$|\hat{\beta} - \beta^0| = \|\hat{\beta} - \beta^0\|_1 \rightarrow 0$$

if the target f^0 is linear and $a + 2\tilde{d} < 1$.

The assumption $F_{C_n} \in O(n^{\tilde{d}})$ is not overly restrictive since even when $f^0(x) = \exp(\mu x^2)$ for some $\mu > 0$ one has that $F_{C_n} = \sup_{|x| \leq C_n} |f^0(x)| = \exp(2\frac{\mu}{\delta} \log(n)) = n^{2\frac{\mu}{\delta}}$. So the assumption of a polynomial growth of F_{C_n} can be satisfied even by functions increasing exponentially fast by choosing $\tilde{d} = 2\frac{\mu}{\delta}$. In the case where we only assume that the covariates and the error terms have bounded r th moments a similar argument shows that f^0 can increase at the rate of an r th degree polynomial. We have assumed in Corollary 5 that the assumptions of Lemma 1 are valid. Recall that we know this is the case if f^0 belongs to a Hölder class of order $r > 1/2$ and the covariates have support $[-1, 1]$. Then we may choose \mathcal{F}_L to consist of polynomials of degree p and $\Gamma = \{1, \dots, s\}$ to meet the assumptions of Lemma 1.

Part a) of Corollary 5 is similar to Corollary 1 but can not be deduced from it since the conditions of the latter Corollary are not satisfied in the case of quadratic loss. Similarly, part b) of Corollary 5 resembles Corollary 3 but can not be deduced from it since the conditions of the latter corollary are not satisfied in the case of quadratic loss. Recall as well that we also assume throughout Section 6.1 that x is one-dimensional which explains the equality in part b) above.

¹³Where $\delta > 0$ is the subgaussianity constant of footnote 9.

6.2 Logistic regression

Next we verify that the logistic loss satisfies Assumptions 1-3. This is done by verifying that the second derivative with respect to the first argument of the conditional loss function is bounded away from zero as discussed in the beginning of Section 6. Let $(Y_i, X_i)_{i=1}^n$ be i.i.d. As seen in Section 2 the loss in case of logistic regression is

$$\rho(f(x), y) = -yf(x) + \log(1 + \exp(f(x))). \quad (27)$$

Here y, x represent the values of the random variable Y_i and X_i . Then we have the following result.

Lemma 2. *Assume that $\epsilon_0 \leq \pi(x) \leq 1 - \epsilon_0$, for all $x \in \mathcal{X}$ for some $0 < \epsilon_0 \leq 1/2$, and $\max_{1 \leq j \leq p} \|\psi_j\|_\infty \leq K$. Set $\lambda_2 = \frac{\lambda_1 \sqrt{s^*}}{2\|\beta^*\|_2}$ and take for some constant $L \geq 8$, $8\lambda_0 \leq \lambda_1 \leq L\lambda_0$ and suppose that*

$$\frac{36L^2\lambda_0 s^* K}{\eta\phi_*^2 c} \leq 1.$$

for $c = \frac{\epsilon_0 e^{-\eta}}{2(1+e^\eta/\epsilon_0)^2}$. Assume that $\|f^* - f^0\|_\infty \leq \eta/2$, with $\frac{3}{2}\Xi(f^*) \leq \frac{9\lambda_1^2 s^*}{\phi_*^2 c}$ and finally that the population covariance matrix of the covariates $\Sigma = E(\psi(X)\psi'(X))$ has full rank. Then,

a) Assumptions 1-3 are valid and on the set τ ,

$$\Xi(\hat{f}) + \lambda_1 \|\hat{\beta} - \beta^*\|_1 + \lambda_2 \|\hat{\beta}_{S^*} - \beta_{S^*}\|_2^2 \leq 6\Xi(f^*) + \frac{36\lambda_1^2 s^*}{\phi_*^2 c}. \quad (28)$$

b) if, furthermore, $\frac{1}{n} \sum_{i=1}^n \max_{1 \leq j \leq p} E\psi_j^2(X_i) \leq 1$ one has that (28) is valid with probability at least $1 - \left(\frac{1}{p}\right)$ upon choosing $\lambda_0 = 2d\sqrt{\frac{\log(p)}{n}}$ (where d is the positive constant from Theorem 2).

Note that except for $\epsilon_0 \leq \pi(x) \leq 1 - \epsilon_0$ for all $x \in \mathcal{X}$ the assumptions of Lemma 2 are a subset of the ones in Lemma 1. Hence, we know that they are satisfied when the covariates have support $[-1, 1]$ and f^0 belongs to a Hölder class of order $r > 1/2$ since then we may choose $\psi_j(x) = x^j$, $j = 1, \dots, p$ and $\Gamma = \{1, \dots, s\}$ as discussed after Lemma 1.

(28) in Lemma 2 gives upper bounds on the excess risk as well as the estimation error in the logit model. Corollaries 1-3 can now be used to establish asymptotic results for these quantities.

By the definition of λ_0 it is not difficult to see that the second term on the right hand side tends to zero if $s^* \in o(n/\log(p))$ and ϕ_* is bounded away from zero. This in turn reveals that $\Xi(\hat{f})$ is of the same order of magnitude as the excess risk of the oracle $\Xi(f^*)$ which parallels part a) in Corollary 5. Mimicking the arguments in part b) of that Corollary one sees that

$$\|\hat{\beta} - \beta^0\|_1 \rightarrow 0$$

with probability tending to one if ϕ_*^2 is bounded away from 0 and the target f^0 is linear as long as $s^* \log(p)/n \rightarrow 0$.

7 Conclusion

This paper has established an oracle inequality for empirical loss minimization penalized by the elastic net penalty. This inequality is valid for convex loss functions and non-linear target functions and we stress again that this is a finite sample result. We have also seen that the results for a Lasso penalty are a special case of ours. For the case where the target is linear the oracle inequality can be used to deduce finite sample upper bounds on the estimation error of $\hat{\beta}$. The oracle inequality can also be used to show that the excess loss of our estimator is asymptotically of the same order as that of the oracle. Also, when the target is linear we give sufficient conditions for $\hat{\beta}$ to be consistent for β^0

Furthermore, we explain how to construct a thresholded elastic net estimator which can perform consistent variable selection when the truth is linear. To illustrate the generality of our framework we give two examples of settings which fit into our theory – the quadratic and the logistic loss. In the case of a quadratic loss function we allow for heteroscedastic error terms. In addition, we show how our results provide new insights into nonparametric series estimation which turns out to be encompassed as a special case of our theory. More precisely, we provide an upper bound on the mean square error of the elastic net series estimator.

Future avenues of research include, but are not limited to, proposing theoretically justified data driven methods for the choice of tuning parameters in the case where the truth is not linear. Also, bounds which are valid for dependent data are interesting for time series analysts. Finally, extending our results to allow for heteroscedasticity in other loss functions than the quadratic one is of interest for practical purposes.

Appendix

We start by proving Theorem 1. As a by-product of the proof we extend the basic inequality (6.28) of Bühlmann and van de Geer (2011) to cover the elastic net using a refined proof technique.

Proof of Theorem 1. The proof consists of three steps. In the first step we set up a basic inequality which we shall use in steps two and three. In the second and third steps we analyze two possibilities:

the penalties multiplied by estimation errors being greater than or equal to the oracle rate Δ^* , or these penalties multiplied by estimation errors being strictly less than the oracle rate, respectively. Throughout the appendix we shall let S_c denote the complement of the set S for any set S ¹⁴.

Step 1. This step uses the convexity of the loss function ρ_f . First, define

$$\tilde{\beta} = t\hat{\beta} + (1-t)\beta^*,$$

where

$$t = \frac{M^*}{M^* + \|\hat{\beta} - \beta^*\|_1}.$$

To simplify notation define $\tilde{\Xi} = \Xi(f_{\tilde{\beta}})$ and $\Xi^* = \Xi(f_{\beta^*})$ and set $\tilde{f} = f_{\tilde{\beta}}$, $\hat{f} = f_{\hat{\beta}}$, $f^* = f_{\beta^*}$. By the minimizing property of $\hat{\beta}$ one has that

$$P_n\rho_{\hat{\beta}} + \lambda_1\|\hat{\beta}\|_1 + \lambda_2\|\hat{\beta}\|_2^2 \leq P_n\rho_{f^*} + \lambda_1\|\beta^*\|_1 + \lambda_2\|\beta^*\|_2^2 \quad (29)$$

By the convexity of $\beta \mapsto \rho_\beta$ and the linearity of the P_n -integral it follows that

$$\begin{aligned} P_n\rho_{\tilde{f}} + \lambda_1\|\tilde{\beta}\|_1 + \lambda_2\|\tilde{\beta}\|_2^2 &\leq tP_n\rho_{\hat{f}} + (1-t)P_n\rho_{f^*} + t\lambda_1\|\hat{\beta}\|_1 + (1-t)\lambda_1\|\beta^*\|_1 + t\lambda_2\|\hat{\beta}\|_2^2 + (1-t)\lambda_2\|\beta^*\|_2^2 \\ &\leq P_n\rho_{f^*} + \lambda_1\|\beta^*\|_1 + \lambda_2\|\beta^*\|_2^2 \end{aligned} \quad (30)$$

where the second inequality follows from (29). Rearranging (30) yields

$$[P_n\rho_{\tilde{f}} - P_n\rho_{f^*}] + \lambda_1\|\tilde{\beta}\|_1 + \lambda_2\|\tilde{\beta}\|_2^2 \leq \lambda_1\|\beta^*\|_1 + \lambda_2\|\beta^*\|_2^2. \quad (31)$$

Next, note that

$$-[V_n(\tilde{\beta}) - V_n(\beta^*)] = -\left[\frac{1}{n}\sum_{i=1}^n(\rho_{\tilde{f}}(Z_i) - E\rho_{\tilde{f}}(Z_i)) - \frac{1}{n}\sum_{i=1}^n(\rho_{f^*}(Z_i) - E\rho_{f^*}(Z_i))\right]. \quad (32)$$

and recall that $\Xi^* = n^{-1}\sum_{i=1}^n(E\rho_{f^*} - E\rho_{f_0})$. Adding $-[V_n(\tilde{\beta}) - V_n(\beta^*)]$ and Ξ^* to both sides of (31)

$$-[V_n(\tilde{\beta}) - V_n(\beta^*)] + [P_n\rho_{\tilde{f}} - P_n\rho_{f^*}] + \Xi^* + \lambda_1\|\tilde{\beta}\|_1 + \lambda_2\|\tilde{\beta}\|_2^2 \leq -[V_n(\tilde{\beta}) - V_n(\beta^*)] + \Xi^* + \lambda_1\|\beta^*\|_1 + \lambda_2\|\beta^*\|_2^2. \quad (33)$$

Now note that by (32) and the definition of Ξ^* one gets

$$-[V_n(\tilde{\beta}) - V_n(\beta^*)] + [P_n\rho_{\tilde{f}} - P_n\rho_{f^*}] + \Xi^* = \tilde{\Xi}.$$

¹⁴It will be clear from the context of which set S is a subset.

Using this in (33) gives

$$\tilde{\Xi} + \lambda_1 \|\tilde{\beta}\|_1 + \lambda_2 \|\tilde{\beta}\|_2^2 \leq -[V_n(\tilde{\beta}) - V_n(\beta^*)] + \lambda_1 \|\beta^*\|_1 + \lambda_2 \|\beta^*\|_2^2 + \Xi^*. \quad (34)$$

Next, we bound $-[V_n(\tilde{\beta}) - V_n(\beta^*)]$. To do so we first note that

$$\|\tilde{\beta} - \beta^*\|_1 = \|t(\hat{\beta} - \beta^*)\|_1 = \frac{M^*}{M^* + \|\hat{\beta} - \beta^*\|_1} \|\hat{\beta} - \beta^*\|_1 \leq M^*.$$

Hence,

$$-[V_n(\tilde{\beta}) - V_n(\beta^*)] \leq \sup_{\|\beta - \beta^*\|_1 \leq M^*} |V_n(\beta) - V_n(\beta^*)| := Z_{M^*}.$$

So on $\tau = \{Z_{M^*} \leq \lambda_0 M^*\}$ one may rewrite (34) as

$$\tilde{\Xi} + \lambda_1 \|\tilde{\beta}\|_1 + \lambda_2 \|\tilde{\beta}\|_2^2 \leq \lambda_0 M^* + \lambda_1 \|\beta^*\|_1 + \lambda_2 \|\beta^*\|_2^2 + \Xi^*. \quad (35)$$

Now subtract $\lambda_1 \|\tilde{\beta}_{S^*}\|_1$ from both sides of (35). Using that $\|\tilde{\beta}\|_1 - \|\tilde{\beta}_{S^*}\|_1 = \|\tilde{\beta}_{S_c^*}\|_1$ and the continuity of the norm to conclude that $\|\beta^*\|_1 - \|\tilde{\beta}_{S^*}\|_1 \leq \|\beta^* - \tilde{\beta}_{S^*}\|_1 = \|\beta_{S^*}^* - \tilde{\beta}_{S^*}\|_1$ one gets

$$\tilde{\Xi} + \lambda_1 \|\tilde{\beta}_{S_c^*}\|_1 + \lambda_2 \|\tilde{\beta}\|_2^2 \leq \lambda_0 M^* + \lambda_1 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 + \lambda_2 \|\beta^*\|_2^2 + \Xi^*. \quad (36)$$

Furthermore, by continuity of the norm,

$$0 \leq \|\|\beta_{S^*}^*\|_2 - \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2\| \leq \|\tilde{\beta}_{S^*}\|_2 \leq \|\tilde{\beta}\|_2 \quad (37)$$

Squaring (37) yields,

$$\|\tilde{\beta}\|_2^2 \geq \|\tilde{\beta}_{S^*}\|_2^2 \geq \|\beta_{S^*}^*\|_2^2 + \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 - 2\|\beta_{S^*}^*\|_2 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2. \quad (38)$$

Using (38) on the left hand side of (36) and rearranging yields

$$\tilde{\Xi} + \lambda_1 \|\tilde{\beta}_{S_c^*}\|_1 + \lambda_2 \|\beta_{S^*}^*\|_2^2 + \lambda_2 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq \lambda_0 M^* + \lambda_1 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 + \lambda_2 \|\beta^*\|_2^2 + 2\lambda_2 \|\beta_{S^*}^*\|_2 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2 + \Xi^*.$$

Noting that $\|\beta_{S^*}^*\|_2 = \|\beta^*\|_2$, the inequality above simplifies to

$$\tilde{\Xi} + \lambda_1 \|\tilde{\beta}_{S_c^*}\|_1 + \lambda_2 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq \lambda_0 M^* + \lambda_1 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 + 2\lambda_2 \|\beta^*\|_2 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2 + \Xi^*. \quad (39)$$

We shall refer to (39) as the Basic Inequality. In the next steps we use the adaptive restricted eigenvalue condition and the margin condition to rewrite the Basic Inequality. We consider 2 cases, which will constitute steps 2 and 3. They depend on the behaviour of $\lambda_1 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 +$

$$2\lambda_2\|\beta^*\|_2\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2.$$

Step 2. We consider the case of

$$\lambda_1\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 + 2\lambda_2\|\beta^*\|_2\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2 \geq \Delta^* = \lambda_0 M^*. \quad (40)$$

while Step 3 considers the reverse inequality. First, it will be shown that $\|\tilde{\beta}_{S_c^*}\|_1 \leq L_n\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1$ which allows us to use the adaptive restricted eigenvalue condition. To this end note that from the definition of the oracle rate $\lambda_0 M^* = \Delta^* \geq \Xi^*$ since $H(\cdot) \geq 0$. Using this and $\tilde{\Xi} \geq 0$ in (39) yields

$$\lambda_1\|\tilde{\beta}_{S_c^*}\|_1 \leq 2\Delta^* + \lambda_1\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 + 2\lambda_2\|\beta^*\|_2\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2. \quad (41)$$

Use (40) to rewrite the right hand side of (41) as

$$\lambda_1\|\tilde{\beta}_{S_c^*}\|_1 \leq 3(\lambda_1\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 + 2\lambda_2\|\beta^*\|_2\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2). \quad (42)$$

Using $\lambda_1\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 \leq \lambda_1\sqrt{s^*}\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2$ in (42) one gets

$$\lambda_1\|\tilde{\beta}_{S_c^*}\|_1 \leq 3(\lambda_1\sqrt{s^*} + 2\lambda_2\|\beta^*\|_2)\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2$$

which implies

$$\|\tilde{\beta}_{S_c^*}\|_1 \leq L_n\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2, \quad (43)$$

where $L_n = 3(\sqrt{s^*} + 2\frac{\lambda_2}{\lambda_1}\|\beta^*\|_2)$. To bound the right hand side of (39) from above note that by the adaptive restricted eigenvalue condition

$$\begin{aligned} \lambda_1\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 + 2\|\beta^*\|_2\lambda_2\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2 &\leq \lambda_1\sqrt{s^*}\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2 + 2\lambda_2\|\beta^*\|_2\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 \\ &\leq (\lambda_1\sqrt{s^*} + 2\lambda_2\|\beta^*\|_2)\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2 \\ &\leq (\lambda_1\sqrt{s^*} + 2\lambda_2\|\beta^*\|_2)\frac{\|f_{\tilde{\beta}} - f_{\beta^*}\|}{\phi_*} \end{aligned} \quad (44)$$

Inserting (44) in (39) yields

$$\tilde{\Xi} + \lambda_1\|\tilde{\beta}_{S_c^*}\|_1 + \lambda_2\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq \lambda_0 M^* + (\lambda_1\sqrt{s^*} + 2\lambda_2\|\beta^*\|_2)\frac{\|f_{\tilde{\beta}} - f_{\beta^*}\|}{\phi_*} + \Xi^*. \quad (45)$$

Now add $\lambda_1\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1$ to both sides of (45)

$$\tilde{\Xi} + \lambda_1\|\tilde{\beta}_{S_c^*}\|_1 + \lambda_1\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 + \lambda_2\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq \lambda_0 M^* + \lambda_1\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 + (\lambda_1\sqrt{s^*} + 2\lambda_2\|\beta^*\|_2)\frac{\|f_{\tilde{\beta}} - f_{\beta^*}\|}{\phi_*} + \Xi^*. \quad (46)$$

and note that the second term on the right hand side of (46) can be bounded from above as follows using the adaptive restricted eigenvalue condition

$$\lambda_1 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 \leq \lambda_1 \sqrt{s^*} \frac{\|f_{\tilde{\beta}} - f_{\beta^*}\|}{\phi_*}. \quad (47)$$

Using (47) in (46) to get

$$\tilde{\Xi} + \lambda_1 \|\tilde{\beta}_{S_c^*}\|_1 + \lambda_1 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 + \lambda_2 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq \lambda_0 M^* + (2\lambda_1 \sqrt{s^*} + 2\lambda_2 \|\beta^*\|_2) \frac{\|f_{\tilde{\beta}} - f_{\beta^*}\|}{\phi_*} + \Xi^*. \quad (48)$$

Since we assume $f^* \in \mathbf{F}_{\text{local}}$, and since $f_{\tilde{\beta}} \in \mathbf{F}_{\text{local}}$ (because $\|\tilde{\beta} - \beta^*\| \leq M^*$ as shown in Step 1), we get using (1) and the margin condition that

$$\begin{aligned} (2\lambda_1 \sqrt{s^*} + 2\lambda_2 \|\beta^*\|_2) \frac{\|f_{\tilde{\beta}_{S^*}} - f_{\beta^*}\|}{\phi_*} &= \frac{\|f_{\tilde{\beta}_{S^*}} - f_{\beta^*}\|}{2} \left[\frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi_*} \right] \\ &\leq G\left(\frac{\|f_{\tilde{\beta}} - f_{\beta^*}\|}{2}\right) + H\left(\frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi_*}\right) \\ &\leq \tilde{\Xi}/2 + \Xi^*/2 + H\left(\frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi_*}\right), \end{aligned} \quad (49)$$

where we used (1) for the first inequality, with $u = \frac{\|f_{\tilde{\beta}_{S^*}} - f_{\beta^*}\|}{2}$ and $v = \frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi_*}$, and the second estimate uses the triangle inequality, convexity of $G(\cdot)$ and the margin condition. Substituting (49) into (48) gives

$$\tilde{\Xi} + \lambda_1 \|\tilde{\beta}_{S_c^*}\|_1 + \lambda_1 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 + \lambda_2 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq \lambda_0 M^* + 3\Xi^*/2 + H\left(\frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi_*}\right) + \tilde{\Xi}/2. \quad (50)$$

Next, note that $\|\tilde{\beta}_{S_c^*}\|_1 + \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 = \|\tilde{\beta} - \beta^*\|_1$ such that (50) can be written as

$$\tilde{\Xi} + \lambda_1 \|\tilde{\beta} - \beta^*\|_1 + \lambda_2 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq \lambda_0 M^* + 3\Xi^*/2 + H\left(\frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi_*}\right) + \tilde{\Xi}/2.$$

Now we use that $\Delta^* = \lambda_0 M^* = 3\Xi^*/2 + H\left(\frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi_*}\right)$ to rewrite the above inequality as

$$\tilde{\Xi} + \lambda_1 \|\tilde{\beta} - \beta^*\|_1 + \lambda_2 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq 2\Delta^* + \tilde{\Xi}/2.$$

The above can be rewritten as

$$\tilde{\Xi}/2 + \lambda_1 \|\tilde{\beta} - \beta^*\|_1 + \lambda_2 \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq 2\Delta^*. \quad (51)$$

The inequality (51) yields the desired oracle inequality but for $\tilde{\beta}$ instead of $\hat{\beta}$. However, it also follows from (51) (using $\Delta^* = \lambda_0 M^*$) that

$$\lambda_1 \|\tilde{\beta} - \beta^*\|_1 \leq 2\Delta^* = 2\lambda_0 M^*$$

which in turn yields (using $\lambda_1 \geq 4\lambda_0$)

$$\|\tilde{\beta} - \beta^*\|_1 \leq 2\frac{\lambda_0}{\lambda_1}M^* \leq M^*/2.$$

Next, note that by the definitions of $\tilde{\beta}$ and t

$$\tilde{\beta} - \beta^* = t\hat{\beta} + (1-t)\beta^* - \beta^* = t(\hat{\beta} - \beta^*) = \frac{M^*}{M^* + \|\hat{\beta} - \beta^*\|_1}(\hat{\beta} - \beta^*).$$

Hence,

$$\frac{M^*}{M^* + \|\hat{\beta} - \beta^*\|_1}\|\hat{\beta} - \beta^*\|_1 = \|\tilde{\beta} - \beta^*\|_1 \leq M^*/2$$

which upon rearranging yields $\|\hat{\beta} - \beta^*\|_1 \leq M^*$. But this means that all the above derivations are valid with $\hat{\beta}$ replacing $\tilde{\beta}$ by simply starting from (29) instead of (30). In particular, (51) yields

$$\Xi(\hat{f}) + 2\lambda_1\|\hat{\beta} - \beta^*\|_1 + 2\lambda_2\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq 4\Delta^* = 6\Xi(f^*) + 4H \left(\frac{4\lambda_1\sqrt{s^*} + 4\lambda_2\|\beta^*\|_2}{\phi_*} \right). \quad (52)$$

which implies the bound in Theorem 1.

Step 3. Here we consider the case

$$\lambda_1\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 + \lambda_2\|\beta^*\|_2\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2 < \Delta^*. \quad (53)$$

As in Step 2 we note that $\Delta^* \geq \Xi^*$ by the definition of Δ^* . Using also that $\lambda_0M^* = \Delta^*$ the Basic Inequality in (39) can be written as,

$$\tilde{\Xi} + \lambda_1\|\tilde{\beta}_{S_c^*}\|_1 + \lambda_2\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq \lambda_0M^* + \Delta^* + \Xi^* \leq \lambda_0M^* + 2\Delta^* = 3\Delta^*. \quad (54)$$

Then adding $\lambda_1\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1$ to both sides, and noting that $\|\tilde{\beta}_{S_c^*}\|_1 + \|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 = \|\tilde{\beta} - \beta^*\|_1$ one gets

$$\tilde{\Xi} + \lambda_1\|\tilde{\beta} - \beta^*\|_1 + \lambda_2\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq 3\Delta^* + \lambda_1\|\tilde{\beta}_{S^*} - \beta_{S^*}^*\|_1 \leq 4\Delta^*. \quad (55)$$

where the last inequality follows from (53). This is the desired oracle inequality but for $\tilde{\beta}$ instead of $\hat{\beta}$. However, it also follows from (55) that

$$\lambda_1\|\tilde{\beta} - \beta^*\|_1 \leq 4\Delta^*. \quad (56)$$

such that

$$\|\tilde{\beta} - \beta^*\|_1 \leq \frac{4\Delta^*}{\lambda_1} = \frac{4\lambda_0M^*}{\lambda_1} \leq M^*/2,$$

by $\lambda_0 M^* = \Delta^*$ and $\lambda_1 \geq 8\lambda_0$. By the same arguments as in step 2 it now follows that $\|\hat{\beta} - \beta^*\|_1 \leq M^*$. This implies that all the above arguments can be repeated with $\hat{\beta}$ replacing $\tilde{\beta}$. In particular, (55) yields

$$\Xi(\hat{f}) + \lambda_1 \|\hat{\beta} - \beta^*\|_1 + \lambda_2 \|\tilde{\beta}_{S^*} - \beta_{S^*}\|_2^2 \leq 4\Delta^*. \quad (57)$$

which is the desired inequality. \square

Proof of Theorem 2. Set

$$\zeta = D \left[4\Lambda\left(\frac{K}{3}, n, p\right) + \frac{tK}{3n} + \sqrt{\frac{2t}{n}} \sqrt{1 + 8\Lambda\left(\frac{K}{3}, n, p\right)} \right] \quad (58)$$

with $\Lambda\left(\frac{K}{3}, n, p\right) = \sqrt{\frac{2\log 2p}{n}} + \frac{K\log 2p}{3n}$. Then, for all $t > 0$ Bühlmann and van de Geer (2011) show (Theorem 14.5) that

$$P(Z_M \leq M\zeta) \geq 1 - \exp(-t).$$

Next note that there exist a constant $c > 0$ (whose value may change throughout the display below) such that

$$\Lambda\left(\frac{K}{3}, n, p\right) \leq c \left(\sqrt{\frac{\log(2)}{n}} + \sqrt{\frac{\log(p)}{n}} + \frac{\log(2)}{n} + \frac{\log(p)}{n} \right) \leq c \sqrt{\frac{\log(p)}{n}}$$

Hence, choosing $t = \log(p)$ implies that there exists a constant $\tilde{c} > 0$ (whose value may change throughout the display below) such that

$$\zeta \leq \tilde{c}D \left[\sqrt{\frac{\log(p)}{n}} + \frac{\log(p)}{n} + \sqrt{\frac{\log(p)}{n}} + \left(\frac{\log(p)}{n}\right)^{3/4} \right] \quad (59)$$

$$\leq \tilde{c}D \sqrt{\frac{\log(p)}{n}} := \lambda_0 \quad (60)$$

This implies

$$P(\tau) = P(Z_{M^*} \leq \lambda_0 M^*) \geq P(Z_{M^*} \leq \zeta M^*) \geq 1 - \exp(-\log(p)) = 1 - \left(\frac{1}{p}\right)$$

\square

Proof of Theorem 3. The bound in Theorem 1 is valid on the set τ . Theorem 2 provides the stated lower bound on the probability of τ . Combining these two results gives the theorem. \square

Lemma 3. Let $G(u)$ be a strictly convex function on $[0, \infty)$, with $G(0) = 0$. The convex conjugate $H(v) = \sup_{u \geq 0} [uv - G(u)]$, $v \geq 0$ satisfies

1. H is non-negative and non-decreasing.

2. H is convex.

3. H is right-continuous at 0.

Proof. The non-negativity of H follows from $H(v) \geq [0 \cdot v - G(0)] = 0$ for all $v \geq 0$. Let $0 \leq v_1 \leq v_2$. Then, since $[uv_1 - G(u)] \leq [uv_2 - G(u)]$ for all $u \geq 0$,

$$H(v_1) = \sup_{u \geq 0} [uv_1 - G(u)] \leq \sup_{u \geq 0} [uv_2 - G(u)] = H(v_2)$$

and so H is non-decreasing.

The convexity of H may be found in Theorem 12.2 of Rockafellar (1997). Here, for the sake of completeness, we give a more direct argument. For any $0 < \lambda < 1$ and $v_1, v_2 \geq 0$

$$\begin{aligned} H(\lambda v_1 + (1 - \lambda)v_2) &= \sup_{u \geq 0} [u(\lambda v_1 + (1 - \lambda)v_2) - G(u)] = \sup_{u \geq 0} [\lambda(uv_1 - G(u)) + (1 - \lambda)(uv_2 - G(u))] \\ &\leq \lambda \sup_{u \geq 0} (uv_1 - G(u)) + (1 - \lambda) \sup_{u \geq 0} (uv_2 - G(u)) = \lambda H(v_1) + (1 - \lambda)H(v_2) \end{aligned}$$

establishing the convexity of H .

To establish that H is right continuous at 0 note first that $H(0)$ is a lower bound for $\{H(x_n)\}_{n=1}^{\infty}$ for any sequence $x_n \downarrow 0$ since H is non-decreasing. Hence, $\{H(x_n)\}_{n=1}^{\infty}$ is a bounded non-increasing sequence and so it possesses a limit which furthermore satisfies $H(0) \leq \inf_n H(x_n) = \lim_n H(x_n)$. It suffices to show that $H(0) \geq \inf_{x>0} H(x) = \inf_n H(x_n)$ to conclude $H(0) = \inf_n H(x_n) = \lim_n H(x_n)$. We assume the converse to reach a contradiction., i.e. assume that $H(0) < \inf_{x>0} H(x)$. In particular, $H(0) < \inf_{0<\lambda<1} H((1 - \lambda)x)$ for all $x > 0$ such that there exists an $\epsilon > 0$ satisfying $\inf_{0<\lambda<1} H((1 - \lambda)x) = H(0) + \epsilon$. But by the convexity of H it holds for all $0 < \lambda < 1$ that

$$\lambda H(0) + (1 - \lambda)H(x) \geq H((1 - \lambda)x) \geq H(0) + \epsilon$$

By continuity of the left hand side in λ it follows that

$$H(0) = \lim_{\lambda \uparrow 1} [\lambda H(0) + (1 - \lambda)H(x)] \geq H(0) + \epsilon$$

which is a contradiction and so we can't have $H(0) < \inf_{x>0} H(x)$ and we conclude that H is right-continuous at 0. \square

Proof of Corollary 1. First, note that the probability with which the inequality in Theorem 3 is valid tends to one. Also, this inequality implies

$$\Xi(\hat{f}) \leq 6\Xi(f^*) + 4H \left(\frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi_*} \right)$$

Hence, it suffices to show that $\limsup_{n \rightarrow \infty} H \left(\frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi_*} \right) = \limsup_{n \rightarrow \infty} H \left(\frac{6\lambda_1 \sqrt{s^*}}{\phi_*} \right) = 0$. To this end, observe that with $\lambda_1 \in O(\lambda_0)$

$$\frac{6\lambda_1 \sqrt{s^*}}{\phi_*} \in O \left(\sqrt{\frac{n^a}{n} n^b} \right) = O \left(n^{a/2+b/2-1/2} \right) \subseteq o(1)$$

which yields the desired result by the-right continuity of H established in Lemma 3. \square

Proof of Corollary 2. The first two inequalities, (10) and (11), follow from Theorem 3 upon using the same reasoning as in remark 2 preceding Theorem 1. In particular, choose $\Gamma = S^0$ in the definition of the oracle. This implies $\Xi(f^*) = \Xi(f^0) = 0$ (as seen in remark 2). (12) and (13) follows from (10) and (11) under the given assumptions by simple insertion. \square

Proof of Corollary 3. First note that the probability with which inequality (13) is valid tends to one. It remains to be shown that the right hand side of (13) tends to zero. But under the stated conditions the right hand side is of order

$$O \left(\sqrt{\frac{n^a}{n} n^b} \right) = O \left(n^{a/2+b-1/2} \right) \subseteq o(1)$$

where the last inclusion follows from the assumption $a + 2b < 1$. \square

Proof of Lemma 1. a) *The Analysis of Assumption 1.* First, note that by the second derivative test $\rho(f(x), y)$ is a convex function in $f(x)$. Since, the second derivative is constant, and equal to 2 this also shows that the margin condition is satisfied with a quadratic margin and $2c = 2$, i.e. $c = 1$. This was of course already clear from (18) prior to Lemma 1. The analysis of assumption 2 is slightly more involved:

The Analysis of Assumption 2. We show that $f_\beta \in \mathbf{F}_{\text{local}}$ under the stated conditions. More precisely, we must show that $\|f_\beta - f^0\|_\infty \leq \eta$. Since

$$\|f_\beta - f^0\|_\infty \leq \|f_\beta - f^*\|_\infty + \|f^* - f^0\|_\infty \leq \|f_\beta - f^*\|_\infty + \eta/2$$

it suffices to show that $\|f_\beta - f^*\|_\infty \leq \eta/2$. To this end, note that

$$|f_\beta(x) - f^*(x)| = \left| \sum_{j=1}^p (\beta_j - \beta_j^*) \psi_j(x) \right| \leq \|\beta - \beta^*\|_1 \max_{1 \leq j \leq p} |\psi_j(x)|$$

which implies $\|f_\beta - f^*\|_\infty \leq M^*K$. Hence, it suffices to show that $M^*K \leq \eta/2$. To do so, recall that by using $H(v) = v^2/4c$ (with $c = 1$), $\Xi(f^*) = E(f^* - f^0)^2$ and $\lambda_2 = \frac{\lambda_1 \sqrt{s^*}}{2\|\beta^*\|_2}$

$$\begin{aligned}
M^* &= \frac{\Delta^*}{\lambda_0} = \frac{1}{\lambda_0} \left((3/2)E(f^* - f^0)^2 + H \left(\frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi(S^*)} \right) \right) \\
&= \frac{1}{\lambda_0} \left((3/2)E(f^* - f^0)^2 + \left(\frac{4\lambda_1 \sqrt{s^*} + 4\lambda_2 \|\beta^*\|_2}{\phi(S^*)} \right)^2 / 4 \right) \\
&= \frac{1}{\lambda_0} \left((3/2)E(f^* - f^0)^2 + \frac{9\lambda_1^2 s^*}{\phi^2(S^*)} \right) \\
&\leq \frac{1}{\lambda_0} \left(\frac{18\lambda_1^2 s^*}{\phi^2(S^*)} \right) \\
&\leq \frac{18L^2 \lambda_0 s^*}{\phi^2(S^*)}
\end{aligned}$$

such that $M^*K \leq \eta/2$ under the stated assumptions.

The Analysis of Assumption 3. The validity of Assumption 3 follows from the fact that Σ is assumed to have full rank. This is sufficient for Assumption 3 to be valid as argued just after (2) in Section 2.

Inequality (20) follows upon using $H(v) = v^2/4c$ with $c = 1$ as well as $\Xi(f) = E(f - f^0)^2$ for all $f \in \mathbf{F}$ in Theorem 1.

b) Next, we turn to part b) of the lemma. This result will be derived based on Theorem 2. Hence, we verify the assumptions of that theorem. The two boundedness conditions are valid by assumption. Next, we establish the local Lipschitz continuity. Note that on $\mathcal{A} = \{\max_{1 \leq i \leq n} |X_i| \vee |\epsilon_i| \leq C_n\}$

$$\begin{aligned}
\left| \frac{\partial \rho(f_\beta(X_i), Y_i)}{\partial f_\beta(X_i)} \right| &= 2|Y_i - f_\beta(X_i)| = 2(|\epsilon_i + f^0(X_i) - f_\beta(X_i)|) \leq 2 \left(|\epsilon_i| + F_{C_n} + \left| \sum_{j=1}^p \beta_j \psi_j(x) \right| \right) \\
&\leq 2 \left(|\epsilon_i| + F_{C_n} + \|\beta\|_1 \max_{1 \leq j \leq p} \|\psi_j\|_\infty \right) \leq 2(C_n + F_{C_n} + GK)
\end{aligned}$$

for all $i = 1, \dots, n$. So, on the set \mathcal{A} , the first derivative of the loss function is bounded and hence the loss function is Lipschitz continuous on this set. This implies that

$$P(\tau^c) \leq P(\tau^c \cap \mathcal{A}) + P(\mathcal{A}^c)$$

By the above arguments $\rho(f(x), y)$ is Lipschitz continuous on \mathcal{A} with Lipschitz constant $2(C_n + F_{C_n} + GK)$. Hence, by Theorem 2 $P(\tau^c \cap \mathcal{A}) \leq \left(\frac{1}{p}\right)$ and the Lipschitz constant D_n in the definition of $\lambda_0 = dD_n \sqrt{\frac{\log(p)}{n}}$ in Theorem 2 may be taken to be $2(C_n + F_{C_n} + GK)$. Next, through subgaussianity

of X_1, ϵ_1 , by a union bound it follows that $P(\mathcal{A}^c) \leq 2\alpha n \exp(-\delta C_n^2)$ for positive constants α and δ . This yields the stated lower bound on the probability of τ (on which inequality (20) is valid). \square

Proof of Corollary 4. It follows directly from Lemma 1 that

$$E(\hat{f} - f^0)^2 \leq 6E(f^* - f^0)^2 + \frac{36L^2\lambda_0^2 s^*}{\phi_*^2}. \quad (61)$$

with probability at least $1 - \left(\frac{1}{p}\right) - 2\alpha n \exp(-\delta C_n^2)$.

The assumptions of part b) are sufficient for Lemma 1 to be valid as seen in the discussion succeeding Lemma 1. (24) follows from (23) using (21). The estimate (25) follows by minimizing the order of the upper bound in (24) with respect to s^* . The equality in (25) uses that $\lambda_0 = dD_n \sqrt{\frac{\log(p)}{n}}$ with $D_n = 2(C_n + F_{C_n} + GK)$ and $d > 0$ and that F_{C_n} is bounded since f^0 is continuous and the covariates have compact support. (26) follows by insertion of C_n into 25. \square

Proof of Corollary 5. First note that the choice of C_n and $p \rightarrow \infty$ ensure that the probability with which inequality (20) is valid tends to one. To proof part a) it suffices to show that $\lambda_0^2 s^* \rightarrow 0$ (this follows from (20)) which is in turns implied by

$$(C_n^2 + F_{C_n}^2) \frac{\log(p)}{n} s^* \in O\left((\log(n) + n^{2\tilde{d}}) \frac{n^a}{n} n^b\right) \subseteq o(1)$$

where the first inclusion is by assumption and the second follows from $a + b + 2\tilde{d} < 1$.

Regarding part b), choosing $\Gamma = S^0$ in the definition of the oracle implies $\Xi(f^*) = \Xi(f^0) = 0$ (as seen in remark 2 after Theorem 1) since the best linear predictor of a linear target is just the target itself (in this case we of course also have $s^* = s^0 = 1$ and $\beta^* = \beta^0$). Hence, we deduce from (20) that

$$|\hat{\beta} - \beta^0| = \|\hat{\beta} - \beta^0\|_1 \leq \frac{9L^2\lambda_0}{2\phi_*^2}.$$

with probability tending to one for $\lambda_0 = dD_n \sqrt{\frac{\log(p)}{n}}$ with $D_n = 2(C_n + F_{C_n} + GK)$ and $d > 0$. So, it suffices to see that $\lambda_0 \rightarrow 0$ which is implied by

$$(C_n^2 + F_{C_n}^2) \frac{\log(p)}{n} \in O\left((\log(n) + n^{2\tilde{d}}) \frac{n^a}{n}\right) \subseteq o(1)$$

if $a + 2\tilde{d} < 1$. \square

Proof of Lemma 2. The proof is similar to the one of Lemma 6.8 in Bühlmann and van de Geer (2011). First, note that $\rho(f(x), y)$ is a convex function in $f(x)$ since it can be written as the sum

of convex functions: the first right hand side term in (27) of $\rho(f(x), y)$ is linear in $f(x)$, and the second term has positive second derivative. We start by showing that Assumptions 1-3 are satisfied.

Step 1. *The Analysis of Assumption 1.* We shall show that one may choose $G(x) = cx^2$ for some positive constant c to be defined precisely below. To do so we follow the general route laid out in the beginning of Section 6. Define

$$l(f(x), x) = E[\rho(f(X), Y)|(X, f(X)) = (x, f(x))] = -\pi(x)f(x) + \log(1 + \exp(f(x))) \quad (62)$$

where $\pi(x) = E(Y|(X, f(X)) = (x, f(x)))$. (62) is minimized with respect to $f \in \mathbf{F}$ at $f^0(x) = \log\left(\frac{\pi(x)}{1-\pi(x)}\right)$. Hence, $f^0(x) = \log\left(\frac{\pi(x)}{1-\pi(x)}\right)$. Note that the second order partial derivative of $l(f(x), x)$ with respect to $f(x)$ is

$$\frac{\partial^2 l(a, x)}{\partial a^2} \Big|_{a=f(x)} = \frac{\exp(f(x))}{1 + \exp(f(x))} \left(1 - \frac{\exp(f(x))}{1 + \exp(f(x))}\right) = \frac{\exp(f(x))}{(1 + \exp(f(x)))^2}. \quad (63)$$

So we must show that $\frac{\exp(f(x))}{(1+\exp(f(x)))}$ is bounded from below by a constant for $f \in \mathbf{F}_{\text{local}}$. To do so it suffices to bound $f(x)$ from above and below. To this end, note that for all $x \in \mathcal{X}$

$$f^0(x) - \|f - f^0\|_\infty \leq f(x) \leq f^0(x) + \|f - f^0\|_\infty$$

which implies that for all $f \in \mathbf{F}_{\text{local}}$

$$f^0(x) - \eta \leq f(x) \leq f^0(x) + \eta \quad (64)$$

Furthermore, since $f^0(x) = \log\left(\frac{\pi(x)}{1-\pi(x)}\right)$ and $\epsilon_0 \leq \pi(x) \leq 1 - \epsilon_0$, we get

$$\log\left(\frac{\epsilon_0}{1 - \epsilon_0}\right) \leq f^0(x) \leq \log\left(\frac{1 - \epsilon_0}{\epsilon_0}\right).$$

Together with (63) and (64) this implies that

$$\frac{\partial^2 l(a, x)}{\partial a^2} \Big|_{a=f(x)} \geq \frac{\frac{\epsilon_0}{1-\epsilon_0} e^{-\eta}}{\left(1 + \frac{1-\epsilon_0}{\epsilon_0} e^\eta\right)^2} \geq \frac{\epsilon_0 e^{-\eta}}{(1 + e^\eta/\epsilon_0)^2} > 0$$

for all $x \in \mathcal{X}$. Hence, one may use $2c = \frac{\epsilon_0 e^{-\eta}}{(1 + e^\eta/\epsilon_0)^2}$.

Step 2. *The Analysis of Assumption 2.* We show that $f_\beta \in \mathbf{F}_{\text{local}}$ under the stated conditions. More precisely, we must show that $\|f_\beta - f^0\|_\infty \leq \eta$. Since

$$\|f_\beta - f^0\|_\infty \leq \|f_\beta - f^*\|_\infty + \|f^* - f^0\|_\infty \leq \|f_\beta - f^*\|_\infty + \eta/2$$

it suffices to show that $\|f_\beta - f^*\|_\infty \leq \eta/2$. To this end, note that

$$|f_\beta(x) - f^*(x)| = \left| \sum_{j=1}^p (\beta_j - \beta_j^*) \psi_j(x) \right| \leq \|\beta - \beta^*\|_1 \max_{1 \leq j \leq p} |\psi_j(x)|$$

which implies $\|f_\beta - f^*\|_\infty \leq M^*K$. Hence, it suffices to show that $M^*K \leq \eta/2$. To do so, recall that, by using $H(v) = v^2/4c$, and $\lambda_2 = \frac{\lambda_1\sqrt{s^*}}{2\|\beta^*\|_2}$

$$\begin{aligned}
M^* &= \frac{\Delta^*}{\lambda_0} = \frac{1}{\lambda_0} \left((3/2)\Xi(f^*) + H \left(\frac{4\lambda_1\sqrt{s^*} + 4\lambda_2\|\beta^*\|_2}{\phi(S^*)} \right) \right) \\
&= \frac{1}{\lambda_0} \left((3/2)\Xi(f^*) + \left(\frac{4\lambda_1\sqrt{s^*} + 4\lambda_2\|\beta^*\|_2}{\phi(S^*)} \right)^2 / 4c \right) \\
&= \frac{1}{\lambda_0} \left((3/2)\Xi(f^*) + \frac{9\lambda_1^2 s^*}{\phi^2(S^*)c} \right) \\
&\leq \frac{1}{\lambda_0} \left(\frac{18\lambda_1^2 s^*}{\phi^2(S^*)c} \right) \\
&\leq \frac{18L^2\lambda_0 s^*}{\phi^2(S^*)c}
\end{aligned}$$

such that $M^*K \leq \eta/2$ under the stated assumptions.

Step 3. *The Analysis of Assumption 3.* The validity of Assumption 3 follows from the fact that Σ is assumed to have full rank. This is sufficient for Assumption 3 to be valid as argued just after (2) in Section 2.

Inequality (28) follows from Theorem 1 upon using $H(v) = v^2/4c$ with $c = \frac{\epsilon_0 e^{-\eta}}{2(1+e^\eta/\epsilon_0)^2}$.

Next, we turn to part b) of Lemma 2. It suffices to verify the assumptions of Theorem 2 since (28) is valid on the set τ . First, note that $\rho(f(x), y)$ is Lipschitz continuous in $f(x)$ for all $y \in \mathcal{Y}$ since

$$\left| \frac{\partial \rho(a, y)}{\partial a} \Big|_{a=f(x)} \right| = \left| -y + \frac{e^{f(x)}}{1 + f(x)} \right| \leq 2$$

and so D in 7 may be chosen to be 2. The two boundedness assumptions on the basis functions are valid by assumption. So, using $H(v) = v^2/4c$ in Theorem 3 yields

$$\Xi(\hat{f}) + \lambda_1 \|\hat{\beta} - \beta^*\|_1 + \lambda_2 \|\hat{\beta}_{S^*} - \beta_{S^*}^*\|_2^2 \leq 6\Xi(f^*) + 4H \left(\frac{4\lambda_1\sqrt{s^*} + 4\lambda_2\|\beta^*\|_2}{\phi_*} \right) \leq 6\Xi(f^*) + 36 \frac{\lambda_1^2 s^*}{\phi_*^2 c}$$

□

References

Belloni, A., D. Chen, V. Chernozhukov, and C. Hansen (2012). Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica* 80(6), 2369–2429.

- Belloni, A. and V. Chernozhukov (2011). High dimensional sparse econometric models: An introduction. *Inverse Problems and High-Dimensional Estimation*, 121–156.
- Belloni, A., V. Chernozhukov, and L. Wang (2011). Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika* 98(4), 791–806.
- Bickel, P. J., Y. Ritov, and A. B. Tsybakov (2009). Simultaneous analysis of lasso and dantzig selector. *The Annals of Statistics* 37(4), 1705–1732.
- Bühlmann, P. and S. van de Geer (2011). *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer-Verlag, New York.
- Caner, M. and X. Han (2013). Selecting the correct number of factors in approximate factor models: The large panel case with group bridge estimators. *Journal of Business and Economic Statistics* (conditionally accepted).
- Caner, M. and K. Knight (2013). An alternative to unit root tests: Bridge estimators differentiate between nonstationary versus stationary models and select optimal lag. *Journal of Statistical Planning and Inference*, 691–715.
- Caner, M. and H. H. Zhang (2013). Adaptive elastic net for generalized methods of moments. *Journal of Business & Economic Statistics* (just-accepted).
- Chen, X. (2007). Large sample sieve estimation of semi-nonparametric models. *Handbook of Econometrics* 6, 5549–5632.
- Cheng, X. and Z. Liao (2013). Select the valid and relevant moments: An information-based lasso for gmm with many moments, second version. Technical report, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.
- Fan, J. and R. Li (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association* 96(456), 1348–1360.
- Fan, J. and J. Lv (2008). Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 70(5), 849–911.
- Fan, J., J. Lv, and L. Qi (2011). Sparse high dimensional models in economics. *Annual review of economics* 3, 291–317.

- Fan, Y. and C. Y. Tang (2013). Tuning parameter selection in high dimensional penalized likelihood. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 75(3), 531–552.
- Hebiri, M. and S. van de Geer (2011). The smooth-lasso and other $l_1 + l_2$ -penalized methods. *Electronic Journal of Statistics* 5, 1184–1226.
- Huang, J., J. L. Horowitz, and S. Ma (2008). Asymptotic properties of bridge estimators in sparse high-dimensional regression models. *The Annals of Statistics* 36(2), 587–613.
- Kallenberg, O. (2002). *Foundations of modern probability*. Springer.
- Kock, A. and L. Callot (2013). Oracle inequalities for high dimensional vector autoregressions. *CREATES Research Paper*, available on *arXiv*.
- Kock, A. B. (2013). Oracle efficient variable selection in random and fixed effects panel data models. *Econometric Theory* 29, 115–152.
- Lounici, K. (2008). Sup-norm convergence rate and sign concentration property of lasso and dantzig estimators. *Electronic Journal of Statistics* 2, 90–102.
- Negahban, S. N., P. Ravikumar, M. J. Wainwright, and B. Yu (2012). A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers. *Statistical Science* 27(4), 538–557.
- Newey, W. K. (1997). Convergence rates and asymptotic normality for series estimators. *Journal of Econometrics* 79(1), 147–168.
- Rockafellar, R. T. (1997). *Convex Analysis*, Volume 28. Princeton University Press.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 267–288.
- Tsybakov, A. B. (2009). *Introduction to nonparametric estimation*. Springer.
- van de Geer, S. A. (2008). High-dimensional generalized linear models and the lasso. *The Annals of Statistics* 36(2), 614–645.
- Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American statistical association* 101(476), 1418–1429.

Zou, H. and T. Hastie (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67(2), 301–320.