# On the risk of convex-constrained least squares estimators under misspecification 

BILLY FANG* and ADITYANAND GUNTUBOYINA**<br>Department of Statistics, University of California, Berkeley CA 94720, USA. E-mail: ${ }^{*}$ blfang@berkeley.edu; ** aditya@stat.berkeley.edu

We consider the problem of estimating the mean of a noisy vector. When the mean lies in a convex constraint set, the least squares projection of the random vector onto the set is a natural estimator. Properties of the risk of this estimator, such as its asymptotic behavior as the noise tends to zero, have been well studied. We instead study the behavior of this estimator under misspecification, that is, without the assumption that the mean lies in the constraint set. For appropriately defined notions of risk in the misspecified setting, we prove a generalization of a low noise characterization of the risk due to [Found. Comput. Math. 16 (2016) 965-1029] in the case of a polyhedral constraint set. An interesting consequence of our results is that the risk can be much smaller in the misspecified setting than in the well-specified setting. We also discuss consequences of our result for isotonic regression.

Keywords: convex constraint; isotonic regression; least squares; misspecification; statistical dimension

## 1. Introduction

In many statistical problems, it is common to model the observations $y_{1}, \ldots, y_{n} \in \mathbb{R}$ as $y_{i}=$ $\theta_{i}^{*}+\sigma z_{i}$ where $\theta_{1}^{*}, \ldots, \theta_{n}^{*}$ are unknown parameters of interest, $z_{1}, \ldots, z_{n}$ represent noise or error variables that have mean zero, and $\sigma>0$ denotes a scale parameter. In vector notation, this is equivalent to writing

$$
Y=\theta^{*}+\sigma Z,
$$

where $Y:=\left(y_{1}, \ldots, y_{n}\right), \theta^{*}:=\left(\theta_{1}^{*}, \ldots, \theta_{n}^{*}\right)$, and $Z:=\left(z_{1}, \ldots, z_{n}\right)$. A common instance of this model is the Gaussian sequence model, where the $z_{1}, \ldots, z_{n}$ are independent standard Gaussian random variables, in which case the model can be written as $Y \sim N\left(\theta^{*}, \sigma^{2} I_{n}\right)$, where $I_{n}$ is the $n \times n$ identity matrix.

A standard method of estimating $\theta^{*}$ from the observation vector $Y$ is to fix a closed convex set $\mathcal{C}$ of $\mathbb{R}^{n}$ and use the least squares estimator under the constraint given by $\theta \in \mathcal{C}$. Specifically, the least squares projection is

$$
\Pi_{\mathcal{C}}(x):=\underset{\theta \in \mathcal{C}}{\operatorname{argmin}}\|x-\theta\|^{2}
$$

(where $\|\cdot\|$ denotes the standard Euclidean norm in $\mathbb{R}^{n}$ ), and one estimates $\theta^{*}$ by

$$
\hat{\theta}(Y):=\Pi_{\mathcal{C}}(Y)
$$

When $\mathcal{C}$ is taken to be $\left\{X \beta:\|\beta\|_{1} \leq R\right\}$ for some deterministic $n \times p$ matrix $X$ and $R>0$, this estimator becomes LASSO in the constrained form as originally proposed by Tibshirani [11]. When $\mathcal{C}$ is taken to be $\left\{X \beta: \min _{j} \beta_{j} \geq 0\right\}$, this estimator becomes nonnegative least squares. Note that shape restricted regression estimators are special cases of nonnegative least squares for appropriate choices of $X$ (see, for example, [4]). Also, note that both sets $\left\{X \beta:\|\beta\|_{1} \leq R\right\}$ and $\left\{X \beta: \min _{j} \beta_{j} \geq 0\right\}$ are examples of polyhedral sets. Therefore in most applications, the constraint set $\mathcal{C}$ is polyhedral.

There exist many results in the literature studying the accuracy of $\hat{\theta}(Y)$ as an estimator for $\theta^{*}$. Most of these results make the assumption that $\theta^{*} \in \mathcal{C}$. In this paper, we shall refer to this assumption as the well-specified assumption. Essentially, the constraint set $\mathcal{C}$ can be taken to be a part of the model specification, and the assumption $\theta^{*} \in \mathcal{C}$ means that the true mean vector $\theta^{*}$ satisfies the model assumptions, that is, the model is well-specified.

Under the well-specified assumption, it is reasonable and common to measure the accuracy of $\hat{\theta}(Y)$ via its risk under squared Euclidean distance. More precisely, the risk of $\hat{\theta}(Y)$ is defined by

$$
R\left(\hat{\theta}, \theta^{*}\right):=\mathbb{E}_{\theta^{*}}\left\|\hat{\theta}(Y)-\theta^{*}\right\|^{2}
$$

where $\mathbb{E}_{\theta^{*}}$ refers to expectation taken with respect to the noise $Z$ in the model $Y=\theta^{*}+\sigma Z$.
Many results on $R\left(\hat{\theta}, \theta^{*}\right)$ in the well-specified setting are available in the literature. Of all the available results, let us isolate two results from [8] because of their generality. In the setting where $Z \sim N\left(0, I_{n}\right)$, Oymak and Hassibi [8] first proved the upper bound

$$
\begin{equation*}
\frac{1}{\sigma^{2}} R\left(\hat{\theta}, \theta^{*}\right) \leq \delta\left(T_{\mathcal{C}}\left(\theta^{*}\right)\right) \tag{1}
\end{equation*}
$$

where $T_{\mathcal{C}}\left(\theta^{*}\right)$ denotes the tangent cone of $\mathcal{C}$ at $\theta^{*}$, defined by

$$
\begin{equation*}
T_{\mathcal{C}}\left(\theta^{*}\right)=\operatorname{cl}\left\{\alpha\left(\theta-\theta^{*}\right): \alpha \geq 0, \theta \in \mathcal{C}\right\} \tag{2}
\end{equation*}
$$

("cl" denotes closure), and where $\delta\left(T_{\mathcal{C}}\left(\theta^{*}\right)\right.$ ) denotes the statistical dimension of the cone $T_{\mathcal{C}}\left(\theta^{*}\right)$. In general, the statistical dimension of a closed cone $T \subseteq \mathbb{R}^{n}$ is defined as

$$
\begin{equation*}
\delta(T):=\mathbb{E}\left\|\Pi_{T}(Z)\right\|^{2}, \tag{3}
\end{equation*}
$$

where the expectation is with respect to $Z \sim N\left(0, I_{n}\right)$. Many properties of the statistical dimension are covered by Amelunxen et al. [2].

In the case when the constraint set $\mathcal{C}$ is a subspace, the estimator $\hat{\theta}(Y)$ is linear and, in this case, it is easy to see that $\delta\left(T_{\mathcal{C}}\left(\theta^{*}\right)\right)$ is simply the dimension of $\mathcal{C}$, so that inequality (1) becomes an equality. For general closed convex sets, it is therefore reasonable to ask how tight inequality (1) is. It is not hard to construct examples of $\mathcal{C}$ and $\theta^{*} \in \mathcal{C}$ where inequality (1) is loose for fixed $\sigma>0$. However Oymak and Hassibi [8] proved remarkably that the upper bound in (1) is tight in the limit as $\sigma \downarrow 0$ (we shall refer to this in the sequel as the low $\sigma$ limit); that is, when $Z \sim N\left(0, I_{n}\right)$,

$$
\begin{equation*}
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} R\left(\hat{\theta}, \theta^{*}\right)=\delta\left(T_{\mathcal{C}}\left(\theta^{*}\right)\right) . \tag{4}
\end{equation*}
$$

In summary, Oymak and Hassibi [8] proved that $\sigma^{2} \delta\left(T_{\mathcal{C}}\left(\theta^{*}\right)\right)$ is a nice formula for the risk of $\hat{\theta}(Y)$ that is, in general, an upper bound which is tight in the low $\sigma$ limit.

We remark that although Oymak and Hassibi [8] state the results (1) and (4) for the specific case $Z \sim N\left(0, I_{n}\right)$, their proof automatically extends to the more general setting where $Z$ is an arbitrary zero mean random vector with $\mathbb{E}\|Z\|^{2}<\infty$ (the components $Z_{1}, \ldots, Z_{n}$ of $Z$ can be arbitrarily dependent), provided we generalize the definition (3) of statistical dimension by taking the expectation with respect to $Z$, without assuming $Z$ is standard Gaussian. We refer to this modification of the definition (3) as the generalized statistical dimension of the cone $T$. As a slight abuse of notation, we use the same notation $\delta(\cdot)$ for this more general concept, with the understanding that the expectation in the definition is with respect to the distribution of $Z$. By dropping the Gaussian assumption, the generalized statistical dimension loses much of the interpretability and nice geometric properties of the usual statistical dimension [2], but still serves as an abstract notion of the size of a cone $T$ with respect to a distribution $Z$.

This paper deals with the behavior of the estimator $\hat{\theta}(Y)$ when the assumption $\theta^{*} \in \mathcal{C}$ is violated. We shall refer to the situation when $\theta^{*} \notin \mathcal{C}$ as the misspecified setting. Note that, in practice, one can never know if the unknown $\theta^{*}$ truly lies in $\mathcal{C}$. It is therefore necessary to study the behavior of $\hat{\theta}(Y)$ under misspecification.

For the misspecified setting, one must first note that it is no longer reasonable to measure the performance of $\hat{\theta}(Y)$ by the risk $R\left(\hat{\theta}, \theta^{*}\right)$, simply because $\hat{\theta}(Y)$ is constrained to be in $\mathcal{C}$ and hence cannot be expected to be close to $\theta^{*}$ which is essentially unconstrained. There are two natural notions of accuracy of $\hat{\theta}(Y)$ in the misspecified setting, which we call the misspecified risk and the excess risk. The misspecified risk is defined as

$$
\begin{equation*}
M\left(\hat{\theta}, \theta^{*}\right):=\mathbb{E}_{\theta^{*}}\left\|\hat{\theta}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \tag{5}
\end{equation*}
$$

and the excess risk is defined as

$$
\begin{equation*}
E\left(\hat{\theta}, \theta^{*}\right):=\mathbb{E}_{\theta^{*}}\left\|\hat{\theta}(Y)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{C}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2} \tag{6}
\end{equation*}
$$

The misspecified risk, $M\left(\hat{\theta}, \theta^{*}\right)$, is motivated by the observation that, in the misspecified case, the estimator $\hat{\theta}(Y)$ is really estimating $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ so it is natural to measure its squared distance from $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$. On the other hand, the excess risk, $E\left(\hat{\theta}, \theta^{*}\right)$, measures the squared distance of the estimator from $\theta^{*}$ relative to the squared distance of $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ from $\theta^{*}$. We refer the reader to [3] and Section 2 for some background and basic properties on these notions of accuracy under misspecification. For example, it can be shown that $M\left(\hat{\theta}, \theta^{*}\right)$ is always less than or equal to $E\left(\hat{\theta}, \theta^{*}\right)$ (see (12)). It is easy to see that both of these risk measures equal $R\left(\hat{\theta}, \theta^{*}\right)$ in the well-specified case that is,

$$
R\left(\hat{\theta}, \theta^{*}\right)=M\left(\hat{\theta}, \theta^{*}\right)=E\left(\hat{\theta}, \theta^{*}\right), \quad \text { when } \theta^{*} \in \mathcal{C}
$$

An analogue to inequality (1) for the case of misspecification has been proved by Bellec [3], Corollary 2.2, who showed that

$$
\begin{equation*}
\frac{1}{\sigma^{2}} M\left(\hat{\theta}, \theta^{*}\right) \leq \frac{1}{\sigma^{2}} E\left(\hat{\theta}, \theta^{*}\right) \leq \delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right) \tag{7}
\end{equation*}
$$

Again, although this was originally stated for $Z \sim N\left(0, I_{n}\right)$, it holds for arbitrary zero mean random vectors $Z$ with $\mathbb{E}\|Z\|^{2}<\infty$. Note the similarity between the right-hand sides of the inequalities (1) and (7). The only difference is that the tangent cone at $\theta^{*}$ is replaced by the tangent cone at $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ in the case of misspecification. Moreover, in the well-specified setting, the above inequality (7) reduces to (1).

It is now very natural to ask if the second inequality in (7) is tight in the low $\sigma$ limit. One might guess that this should be the case given the result (4) for the well-specified setting. However, it turns out that (7) is not sharp in the low $\sigma$ limit. The main contribution of this paper is to provide an exact formula for the low $\sigma$ limit of $M\left(\hat{\theta}, \theta^{*}\right)$ and $E\left(\hat{\theta}, \theta^{*}\right)$ when $\mathcal{C}$ is polyhedral. Specifically, in Theorem 3.1, we prove that if the noise $Z$ is zero mean with $\mathbb{E}\|Z\|^{2}<\infty$ and if $\mathcal{C}$ is polyhedral, then

$$
\begin{equation*}
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} M\left(\hat{\theta}, \theta^{*}\right)=\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} E\left(\hat{\theta}, \theta^{*}\right)=\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}\right) \tag{8}
\end{equation*}
$$

where $v^{\perp}:=\left\{u \in \mathbb{R}^{n}:\langle u, v\rangle=0\right\}$ for vectors $v \in \mathbb{R}^{n}$. As we remarked earlier, in most applications, the constraint set $\mathcal{C}$ is polyhedral.

Because the set $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$ is a subset of $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$, the right hand side of (8) is never larger than $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right.$ ). Under the assumption that the polyhedron $\mathcal{C}$ has a nonempty interior, along with a mild condition on the noise $Z$, it can be proved that the right hand side of (8) is strictly smaller than $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right.$ ) when $\theta^{*} \notin \mathcal{C}$ (an even stronger statement is proved in Lemma 3.4), which then implies that $\lim _{\sigma \downarrow 0} \sigma^{-2} M\left(\hat{\theta}, \theta^{*}\right)<$ $\lim _{\sigma \downarrow 0} \sigma^{-2} R\left(\hat{\theta}, \Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$. This inequality is more interpretable in the following form:

$$
\begin{equation*}
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left\|\hat{\theta}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2}<\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} \mathbb{E}_{\Pi_{\mathcal{C}}\left(\theta^{*}\right)}\left\|\hat{\theta}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \quad \text { whenever } \theta^{*} \notin \mathcal{C} \tag{9}
\end{equation*}
$$

(Note that the nonempty interior assumption is essential; for instance, when $\mathcal{C}$ is a proper subspace of $\mathbb{R}^{n}$, the inequality (9) becomes equality.) Inequality (9) can be qualitatively understood as follows. The left-hand side above corresponds to misspecification where the data are generated from $\theta^{*} \notin \mathcal{C}$ while the right hand side corresponds to the well-specified setting where the data are generated from $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$. Note that in both cases, the estimator $\hat{\theta}(Y)$ is really estimating $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ so it is natural to compare the squared expected distance to $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ in both situations. The interesting aspect is that (in the low $\sigma$ limit) the expected squared distance is smaller in the misspecified setting compared to the well-specified setting. To the best of our knowledge, this fact has not been noted in the literature previously at this level of generality, although it has been noticed in certain instances of specific problems such as isotonic density estimation and regression (see, for example, [6]).

Our main result, Theorem 3.1, is stated and proved in Section 3 where some intuition is also provided for the exact form of the low $\sigma$ limit in misspecification. The low $\sigma$ limit can be explicitly computed in certain specific situations. In Section 4, we specialize to the Gaussian model $Z \sim N\left(0, I_{n}\right)$ and study in detail the examples when $\mathcal{C}$ is the nonnegative orthant and when $\mathcal{C}$ is the monotone cone (this latter case corresponds to isotonic regression).

In Section 5, we explore issues naturally related to Theorem 3.1. In Section 5.1, we consider the situation when $\mathcal{C}$ is not polyhedral. It seems hard to characterize the low $\sigma$ misspecification
limits in this case but it is possible to compute them when $\mathcal{C}$ is the unit ball. It is interesting to note that the low $\sigma$ limits of $M\left(\hat{\theta}, \theta^{*}\right)$ and $E\left(\hat{\theta}, \theta^{*}\right)$ are different in this case (in sharp contrast to the polyhedral situation). In Section 5.2, we deal with the risks when $\sigma$ is large. Under some conditions, it is possible to write a formula for the large $\sigma$ limits of $M\left(\hat{\theta}, \theta^{*}\right)$ and $E\left(\hat{\theta}, \theta^{*}\right)$; see Proposition 5.3. In Section 5.3, we deal with the maximum normalized risks:

$$
\begin{equation*}
\sup _{\sigma>0} \frac{1}{\sigma^{2}} M\left(\hat{\theta}, \theta^{*}\right) \quad \text { and } \quad \sup _{\sigma>0} \frac{1}{\sigma^{2}} E\left(\hat{\theta}, \theta^{*}\right) \tag{10}
\end{equation*}
$$

In the well-specified setting, inequalities (1) and (4) together imply that the maximum normalized risk equals $\delta\left(T_{\mathcal{C}}\left(\theta^{*}\right)\right)$. However in the misspecified setting, the quantities (10) lie between $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}\right)$ and $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right)$. It seems hard to write down an exact formula for the quantities (10) but we present some simulation evidence in Section 5.3 to argue that they can be strictly between $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}\right)$ and $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right)$.

We conclude with an Appendix that contains technical lemmas and proofs of the various intermediate results throughout the paper.

## 2. Background and notation

In this short section, we shall set up some notation and also recollect standard results in convex analysis that will be used in the remainder of the paper.

For $x \in \mathbb{R}^{n}$ and $r>0$, we denote by $B_{r}(x):=\left\{u \in \mathbb{R}^{n}:\|u-x\| \leq r\right\}$ the closed ball of radius $r$ centered at $x$. For $v \in \mathbb{R}^{n}$, let $v^{\perp}:=\left\{u \in \mathbb{R}^{n}:\langle u, v\rangle=0\right\}$ denote the hyperplane with normal vector $v$. For $\theta_{0} \in \mathcal{C}$, let $F_{\mathcal{C}}\left(\theta_{0}\right):=\left\{\theta-\theta_{0}: \theta \in \mathcal{C}\right\}$ be the result of re-centering the set $\mathcal{C}$ about $\theta_{0}$. Also recall the definition of the tangent cone (2) and note that $T_{\mathcal{C}}\left(\theta_{0}\right)=\operatorname{cl}\left\{\alpha x: x \in F_{\mathcal{C}}\left(\theta_{0}\right), \alpha>\right.$ $0\}$.

If $A$ is an $m \times n$ matrix and $J \subseteq\{1, \ldots, m\}$, we let $a_{j}$ denote the $j$ th row of $A$, and let $A_{J}$ denote the matrix obtained by combining the rows of $A$ indexed by $J$.

A polyhedron refers to a set of the form $\left\{x \in \mathbb{R}^{n}: A x \leq b\right\}$ for some $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^{n}$ where the inequality $\leq$ is interpreted coordinate-wise, i.e. $\left\langle a_{j}, x\right\rangle \leq b_{j}$ for $j=1, \ldots, m$. We will assume that no two pairs $\left(a_{j}, b_{j}\right)$ and $\left(a_{k}, b_{k}\right)$ are scalar multiples of each other. A polyhedral cone is a set of the form $\left\{x \in \mathbb{R}^{n}: A x \leq 0\right\}$ for some $A \in \mathbb{R}^{m \times n}$. Again, we will assume that no two rows of $A$ are scalar multiples of each other. A face of a polyhedron refers to any subset obtained by setting some of the polyhedron's linear inequality constraints to equality instead.

In the remainder of this section, we shall collect some standard results above convex projections that will be used in the paper. These results can be found in a standard reference such as [5]. Recall that $\Pi_{\mathcal{C}}(x)$ denotes the projection of a vector $x \in \mathbb{R}^{n}$ on a closed convex set $\mathcal{C}$. It is well known that $\Pi_{\mathcal{C}}(x)$ is the unique vector in $\mathcal{C}$ satisfying the optimality condition

$$
\begin{equation*}
\left\langle z-\Pi_{\mathcal{C}}(x), x-\Pi_{\mathcal{C}}(x)\right\rangle \leq 0, \quad \forall z \in \mathcal{C} \tag{11}
\end{equation*}
$$

Consequently, we have the following Pythagorean inequality

$$
\|z-x\|^{2}=\left\|z-\Pi_{\mathcal{C}}(x)\right\|^{2}+\left\|\Pi_{\mathcal{C}}(x)-x\right\|^{2}+2\left|z-\Pi_{\mathcal{C}}(x), \Pi_{\mathcal{C}}(x)-x\right\rangle
$$

$$
\geq\left\|z-\Pi_{\mathcal{C}}(x)\right\|^{2}+\left\|\Pi_{\mathcal{C}}(x)-x\right\|^{2}
$$

Plugging in $z=\Pi_{\mathcal{C}}(y)$ and $x=\theta^{*}$ shows that the misspecified error is upper bounded by the excess error, that is,

$$
\begin{equation*}
\left\|\Pi_{\mathcal{C}}(y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \leq\left\|\Pi_{\mathcal{C}}(y)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{C}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2}, \quad \forall y \in \mathbb{R}^{n} \tag{12}
\end{equation*}
$$

If instead we plug in $z=\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ to (11), we have $\left\langle\Pi_{\mathcal{C}}(x)-\Pi_{\mathcal{C}}\left(\theta^{*}\right), x-\Pi_{\mathcal{C}}(x)\right\rangle \geq 0$, which implies

$$
\begin{aligned}
& \left\|\Pi_{\mathcal{C}}(x)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{C}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2} \\
& \quad=-\left\|\Pi_{\mathcal{C}}(x)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2}+2\left\langle\Pi_{\mathcal{C}}(x)-\Pi_{\mathcal{C}}\left(\theta^{*}\right), \Pi_{\mathcal{C}}(x)-\theta^{*}\right\rangle \\
& \quad \leq-\left\|\Pi_{\mathcal{C}}(x)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2}+2\left\langle\Pi_{\mathcal{C}}(x)-\Pi_{\mathcal{C}}\left(\theta^{*}\right), x-\theta^{*}\right\rangle \\
& \quad \leq\left\|x-\theta^{*}\right\|^{2}
\end{aligned}
$$

Combining this with (12), we see that for $Y=\theta^{*}+\sigma Z$ we have

$$
\begin{equation*}
0 \leq\left\|\Pi_{\mathcal{C}}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \leq\left\|\Pi_{\mathcal{C}}(Y)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{C}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2} \leq \sigma^{2}\|Z\|^{2} \tag{13}
\end{equation*}
$$

In the special case where $\mathcal{C}$ is a cone, the optimality condition (11) implies that $\Pi_{\mathcal{C}}(x)$ is the unique vector in $\mathcal{C}$ satisfying

$$
\begin{equation*}
\left\langle\Pi_{\mathcal{C}}(x), x-\Pi_{\mathcal{C}}(x)\right\rangle=0 \quad \text { and } \quad\left\langle z, x-\Pi_{\mathcal{C}}(x)\right\rangle \leq 0, \quad \forall z \in \mathcal{C} \tag{14}
\end{equation*}
$$

## 3. Main theorem: Low noise limit for polyhedra

Our main result below provides a precise characterization of the low $\sigma$ limits of the risks (5) and (6) (normalized by $\sigma^{2}$ ) in the misspecified setting (i.e., when $\theta^{*} \notin \mathcal{C}$ ) for polyhedral $\mathcal{C}$. An implication of this result is that the low $\sigma$ limit can be much smaller than the upper bound (7) of [3].

Theorem 3.1 (Low noise limit of risk for polyhedra). Let $\mathcal{C} \subseteq \mathbb{R}^{n}$ be a closed convex set, and let $Y=\theta^{*}+\sigma Z$ where $\theta^{*} \in \mathbb{R}^{n}$ is not necessarily in $\mathcal{C}$, and $Z$ is zero mean with $\mathbb{E}\|Z\|^{2}<\infty$. Suppose the following "locally polyhedral" condition holds.

$$
\begin{align*}
& T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \quad \text { is a polyhedral cone } \quad \text { and } \\
& T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap B_{r^{*}}(0)=F_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap B_{r^{*}}(0) \quad \text { for some } r^{*}>0 . \tag{15}
\end{align*}
$$

Then,

$$
\begin{equation*}
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} M\left(\hat{\theta}, \theta^{*}\right)=\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} E\left(\hat{\theta}, \theta^{*}\right)=\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}\right) \tag{16}
\end{equation*}
$$

Note again that $\delta(\cdot)$ denotes the generalized statistical dimension induced by the noise $Z$, and reduces to the usual statistical dimension [2] when $Z \sim N\left(0, I_{n}\right)$.

We remark that from the proof of Theorem 3.1 below, one can show the following stronger non-asymptotic result. Let $\mathcal{K}:=\Pi_{\mathcal{C}}\left(\theta^{*}\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$ so that the low noise limit (16) is $\delta(\mathcal{K})$. Then for any $\sigma>0$,

$$
\begin{align*}
&\left|\frac{1}{\sigma^{2}} M\left(\hat{\theta}, \theta^{*}\right)-\delta(\mathcal{K})\right| \leq 2 \mathbb{E}\left[\|Z\|^{2} \mathbf{1}_{\left\{\|Z\|>r^{*} / \sigma\right\}}\right] \\
&\left|\frac{1}{\sigma^{2}} E\left(\hat{\theta}, \theta^{*}\right)-\delta(\mathcal{K})\right| \leq 2 \mathbb{E}\left[\|Z\|^{2} \mathbf{1}_{\left\{\|Z\|>r^{*} / \sigma\right\}}\right] \tag{17}
\end{align*}
$$

Note that the error term on the right-hand side (17) tends to zero as $\sigma \downarrow 0$, due to $\mathbb{E}\|Z\|^{2}<\infty$ and the dominated convergence theorem, which recovers the result (16) of the theorem.

We remark that the "locally polyhedral" condition (15) essentially states that $\mathcal{C}$ looks like a polyhedron in a neighborhood around $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$. As established in the following lemma, it automatically holds if $\mathcal{C}$ is a polyhedron, so one can replace any mention of condition (15) with " $\mathcal{C}$ is a polyhedron" for the sake of readability. We provide some remarks on the case when $\mathcal{C}$ is not polyhedral in Section 5.1.

Lemma 3.2. Let $\mathcal{C}$ be a polyhedron. Then the locally polyhedral condition (15) holds for any $\theta^{*} \in \mathbb{R}^{n}$.

Next, the following lemma establishes that the set $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$ that appears in the limit (16) is a face of the tangent cone $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$.

Lemma 3.3. Let $\theta^{*} \in \mathbb{R}^{n}$ and let $\mathcal{C} \subseteq \mathbb{R}^{n}$ be a closed convex set satisfying the locally polyhedral condition (15). Let $A \in \mathbb{R}^{m \times n}$ be such that $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)=\{u: A u \leq 0\}$. Then there exists some subset $J \subseteq\{1, \ldots, m\}$ such that

$$
T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}=\left\{u: A_{J} u=0, A_{J^{c}} u \leq 0\right\}
$$

Thus, $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$ is a face of $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$.
Both the above lemmas are proved in Appendix A.
If $\theta^{*} \in \mathcal{C}$, then we have $\Pi_{\mathcal{C}}\left(\theta^{*}\right)=\theta^{*}$, and Theorem 3.1 reduces to the result (4) of [8]: the excess risk and the misspecified risk become the same, and the common limit is the statistical dimension of $T_{\mathcal{C}}\left(\theta^{*}\right)$. We must remark here that the result of Oymak and Hassibi [8] holds for non-polyhedral $\mathcal{C}$ as well. We discuss the non-polyhedral setting further in Section 5.1.

Theorem 3.1 states that in the misspecified case $\theta^{*} \notin \mathcal{C}$, the low sigma limit still involves the tangent cone $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$, but one needs to intersect it with the hyperplane $\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$ before taking the statistical dimension. Due to the optimality condition (11) characterizing $\Pi_{\mathcal{C}}$, the tangent cone lies entirely on one side of the hyperplane, so the hyperplane does not intersect the interior of the tangent cone. Therefore, the interior of the tangent cone $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$ does not contribute to the low $\sigma$ limit of the risk under misspecification. This makes sense because when
$\theta^{*} \notin \mathcal{C}$ and $\sigma$ is small, the observation vector $Y$ is outside $\mathcal{C}$ with high probability so that $\hat{\theta}(Y)$ lies on the boundary of $\mathcal{C}$. In fact, in the proof of the theorem we essentially show that for any $\sigma>0$, under the event $\left\{\|Z\| \leq r^{*} / \sigma\right\}$ (which has probability tending to 1 as $\sigma \downarrow 0$ ) we have

$$
\Pi_{\mathcal{C}}(Y)=\Pi_{\Pi_{\mathcal{C}}\left(\theta^{*}\right)+T_{\mathcal{C}}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}(Y)
$$

In general, the intersection $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$ can be anything from $\{0\}$ to the full tangent cone $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$ and so the low sigma limit can be anything between 0 and $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right)$. The case when the limit equals zero corresponds to the situation where $\theta^{*}$ lies in the interior of the preimage of $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ under the map $\Pi_{\mathcal{C}}$ so that every point in some neighborhood of $\theta^{*}$ is projected onto the same point $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ (see, Figure 1 (a), for an example).

The following lemma (proved in Appendix A), provides mild conditions under which the intersection $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$ has strictly smaller generalized statistical dimension than the full tangent cone $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$.

Lemma 3.4. Let $\mathcal{C} \subseteq \mathbb{R}^{n}$ be a polyhedron with nonempty interior. Then

$$
\sup _{\theta^{*} \notin \mathcal{C}: \Pi_{\mathcal{C}}\left(\theta^{*}\right)=\theta_{0}} \delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}\right)<\delta\left(T_{\mathcal{C}}\left(\theta_{0}\right)\right)
$$

for every $\theta_{0} \in \mathcal{C}$, provided the random vector $Z$ has nonzero probability of lying in the interior of $T_{\mathcal{C}}\left(\theta_{0}\right)$.

As mentioned already, Lemma 3.4 combined with the main result Theorem 3.1 implies the risk gap (9). In summary, under the nonempty interior assumption, if we think of the low $\sigma$ limit as a function of $\theta^{*}$, we see that as $\theta^{*}$ approaches $\mathcal{C}$ from the outside there is a "jump" when $\theta^{*}$ enters $\mathcal{C}$. This "jump" phenomenon is not unique to the polyhedral case. In Section 5.1, we discuss a non-polyhedral example that also exhibits this jump phenomenon.

Theorem 3.1 suggests something that may seem nonintuitive: if $\theta^{*} \notin \mathcal{C}$ and we use the estimator $\hat{\theta}(Y)=\Pi_{\mathcal{C}}(Y)$, the risk when $Y=\theta^{*}+\sigma Z$ is smaller than the risk when $Y=\Pi_{\mathcal{C}}\left(\theta^{*}\right)+\sigma Z$. As mentioned already, in the case $Y=\theta^{*}+\sigma Z$ the estimator is actually estimating $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$, not $\theta^{*}$. Moreover, the risks (5) and (6) measure error relative to $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ rather than to $\theta^{*}$. Furthermore, the intuition is that in the low $\sigma$ limit, the estimator $\hat{\theta}(Y)$ in the misspecified setting is a projection onto a much smaller set than in the well-specified setting (essentially, a face of a tangent cone instead of the full tangent cone), so more of the original noise in $Y$ is eliminated. This qualitatively explains why having $Y$ generated from $\theta^{*}$ outside $\mathcal{C}$ allows the estimator to estimate $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ better than if $Y$ were generated from $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ instead.

Finally, we observe that in the misspecified setting, there is a gap between Bellec's upper bound $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right)(7)$ and the low $\sigma$ risk limit, unlike in the well-specified setting where the result (4) implies that the normalized risk increases to the upper bound in the low $\sigma$ limit. The upper bound, which is constant in $\sigma$, can become very loose as $\sigma \downarrow 0$. However, in Section 5.3 we shown a few examples where the normalized risk is close to the upper bound for some $\sigma$, as well as examples where the normalized risk remains much smaller than the upper bound for all $\sigma>0$.

### 3.1. Proof of Theorem 3.1

We establish one key lemma (proved in Appendix A) before proving Theorem 3.1. It is a deterministic result that contains the core of the argument: roughly, if we have a polyhedral cone $\mathcal{T}$ and any $\theta^{*} \in \mathbb{R}^{n}$ satisfying $\Pi_{\mathcal{T}}\left(\theta^{*}\right)=0$, then any point $u$ sufficiently near $\theta^{*}$ will have its projection $\Pi_{\mathcal{T}}(u)$ lying in the hyperplane with normal direction $\theta^{*}$.

Lemma 3.5 (Key lemma). Fix $\theta^{*} \in \mathbb{R}^{n}$, and let $\mathcal{T}$ be a closed convex set such that the recentered set $\left\{\theta-\Pi_{\mathcal{T}}\left(\theta^{*}\right): \theta \in \mathcal{T}\right\}$ is a polyhedral cone. Then there exists $r>0$ such that

$$
\begin{equation*}
\Pi_{\mathcal{T}}(u)-\Pi_{\mathcal{T}}\left(\theta^{*}\right) \in\left(\theta^{*}-\Pi_{\mathcal{T}}\left(\theta^{*}\right)\right)^{\perp}, \quad \forall u \in B_{r}\left(\theta^{*}\right) \tag{18}
\end{equation*}
$$

With this lemma, along with some standard results collected in Section 2, we can proceed with proving Theorem 3.1.

Proof of Theorem 3.1. We first prove

$$
\begin{equation*}
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} M\left(\hat{\theta}, \theta^{*}\right)=\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}\right) \tag{19}
\end{equation*}
$$

For any $r>0$, we can write

$$
\begin{align*}
\frac{1}{\sigma^{2}} M\left(\hat{\theta}, \theta^{*}\right)= & \frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left\|\Pi_{\mathcal{C}}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \mathbf{1}_{\left\{Y \in B_{r}\left(\theta^{*}\right)\right\}}\right]  \tag{20}\\
& +\frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left\|\Pi_{\mathcal{C}}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \mathbf{1}_{\left\{Y \notin B_{r}\left(\theta^{*}\right)\right\}}\right]
\end{align*}
$$

We claim the second term on the right-hand side vanishes as $\sigma \downarrow 0$ (regardless of the value of $r>0$ ). Since the projection $\Pi_{\mathcal{C}}$ is non-expansive [5],

$$
\begin{aligned}
0 & \leq \frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left\|\Pi_{\mathcal{C}}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \mathbf{1}_{\left\{Y \notin B_{r}\left(\theta^{*}\right)\right\}}\right] \\
& \leq \frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left\|Y-\theta^{*}\right\|^{2} \mathbf{1}_{\left\{Y \notin B_{r}\left(\theta^{*}\right)\right\}}\right] \\
& =\mathbb{E}\left[\|Z\|^{2} \mathbf{1}_{\{\|Z\|>r / \sigma\}}\right] .
\end{aligned}
$$

Then, the dominated convergence theorem implies the right-hand side tends to zero as $\sigma \downarrow 0$, because $\mathbb{E}\|Z\|^{2}<\infty$ and the random variable $\|Z\|^{2} \mathbf{1}_{\{\sigma\|Z\|>r\}}$ converges to zero pointwise.

Thus, it remains to show

$$
\begin{equation*}
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left\|\Pi_{\mathcal{C}}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \mathbf{1}_{\left\{Y \in B_{r}\left(\theta^{*}\right)\right\}}\right]=\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}\right) \tag{21}
\end{equation*}
$$

for some $r>0$.

We define the re-centered tangent cone

$$
\mathcal{T}:=\left\{\Pi_{\mathcal{C}}\left(\theta^{*}\right)+u: u \in T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right\} .
$$

We claim there exists some $r>0$ such that

$$
\begin{equation*}
\Pi_{\mathcal{C}}(u)=\Pi_{\mathcal{T}}(u), \quad \forall u \in B_{r}\left(\theta^{*}\right) . \tag{22}
\end{equation*}
$$

Indeed, note that the locally polyhedral condition (15) implies the existence of some $r^{*}>0$ such that

$$
\begin{equation*}
\mathcal{C} \cap B_{r^{*}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)=\mathcal{T} \cap B_{r^{*}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \tag{23}
\end{equation*}
$$

Since both projections $\Pi_{\mathcal{C}}$ and $\Pi_{\mathcal{T}}$ are continuous [5] at $\theta^{*}$, there exists some $r>0$ such that the image of $B_{r}\left(\theta^{*}\right)$ under both projections lies in $B_{r^{*}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$. In fact, since the projections are each non-expansive [5], we may take $r=r^{*}$. Thus the local equality (22) of the projections follows from the locally polyhedral condition (23).

By combining this argument with Lemma 3.5, we have shown there exists some $r>0$ that satisfies not only (22), but also (18). With this value of $r$, the equality (22) implies that replacing each instance of $\mathcal{C}$ with $\mathcal{T}$ in (21) does not change either side, since $\Pi_{\mathcal{C}}(Y)=\Pi_{\mathcal{T}}(Y), \Pi_{\mathcal{C}}\left(\theta^{*}\right)=$ $\Pi_{\mathcal{T}}\left(\theta^{*}\right)$, and

$$
T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)=T_{\mathcal{C} \cap B_{r^{*}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)=T_{\mathcal{T} \cap B_{r^{*}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)=T_{\mathcal{T}}\left(\Pi_{\mathcal{T}}\left(\theta^{*}\right)\right)
$$

by the equality (23) and the definition of the tangent cone. Thus, it remains to prove

$$
\begin{equation*}
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left\|\Pi_{\mathcal{T}}(Y)-\Pi_{\mathcal{T}}\left(\theta^{*}\right)\right\|^{2} \mathbf{1}_{\left\{Y \in B_{r}\left(\theta^{*}\right)\right\}}\right]=\delta(\mathcal{K}) \tag{24}
\end{equation*}
$$

where $\mathcal{K}:=T_{\mathcal{T}}\left(\theta^{*}\right) \cap\left(\theta^{*}-\Pi_{\mathcal{T}}\left(\theta^{*}\right)\right)^{\perp}$.
Since $r$ satisfies (18), some re-centering yields

$$
\begin{equation*}
\Pi_{\mathcal{T}}(Y)-\Pi_{\mathcal{T}}\left(\theta^{*}\right)=\Pi_{T_{\mathcal{T}}\left(\theta^{*}\right)}\left(Y-\Pi_{\mathcal{T}}\left(\theta^{*}\right)\right)=\Pi_{\mathcal{K}}\left(Y-\Pi_{\mathcal{T}}\left(\theta^{*}\right)\right) \tag{25}
\end{equation*}
$$

in the event $\left\{Y \in B_{r}\left(\theta^{*}\right)\right\}$.
For $W:=\left(\theta^{*}-\Pi_{\mathcal{T}}\left(\theta^{*}\right)\right)^{\perp}$, we claim

$$
\Pi_{\mathcal{K}}=\Pi_{\mathcal{K}} \circ \Pi_{W}
$$

In fact this holds for any subspace $W$ and closed convex $\mathcal{K} \subseteq W$, by the Pythagorean theorem:

$$
\begin{aligned}
\Pi_{\mathcal{K}}(x) & =\underset{u \in \mathcal{K}}{\operatorname{argmin}}\|x-u\|^{2} \\
& =\underset{u \in \mathcal{K}}{\operatorname{argmin}}\left\{\left\|x-\Pi_{W}(x)\right\|^{2}+\left\|\Pi_{W}(x)-u\right\|^{2}\right\} \\
& =\Pi_{\mathcal{K}}\left(\Pi_{W}(x)\right)
\end{aligned}
$$

Applying this to (25) yields

$$
\begin{aligned}
& \Pi_{\mathcal{T}}(Y)-\Pi_{\mathcal{T}}\left(\theta^{*}\right) \\
& \quad=\Pi_{\mathcal{K}}\left(Y-\Pi_{\mathcal{T}}\left(\theta^{*}\right)\right) \\
& \quad=\Pi_{\mathcal{K}}\left(\Pi_{W}\left(\theta^{*}+\sigma Z-\Pi_{\mathcal{T}}\left(\theta^{*}\right)\right)\right)
\end{aligned}
$$

$$
=\Pi_{\mathcal{K}}\left(\Pi_{W}(\sigma Z)\right) \quad \Pi_{W} \text { is linear, } \Pi_{W}\left(\theta^{*}-\Pi_{\mathcal{T}}\left(\theta^{*}\right)\right)=0
$$

$$
=\Pi_{\mathcal{K}}(\sigma Z)=\sigma \Pi_{\mathcal{K}}(Z) \quad \mathcal{K} \text { is a cone }
$$

in the event $\left\{Y \in B_{r}\left(\theta^{*}\right)\right\}$. By plugging this into the left-hand side of equation (24), we have

$$
\lim _{\sigma \downarrow 0} \mathbb{E}_{\theta^{*}}\left[\left\|\Pi_{\mathcal{K}}(Z)\right\|^{2} \mathbf{1}_{\left\{Y \in B_{r}\left(\theta^{*}\right)\right\}}\right]=\mathbb{E}\left\|\Pi_{\mathcal{K}}(Z)\right\|^{2}=\delta(\mathcal{K})
$$

where the first equality follows by dominated convergence $\left(\left\|\Pi_{\mathcal{K}}(Z)\right\|^{2} \leq\|Z\|^{2}\right.$ and $\mathbb{E}\|Z\|^{2}<$ $\infty$ ). This verifies the desired equality (24) and concludes the proof of the first low $\sigma$ limit (19).

We now prove the other equality

$$
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} M\left(\hat{\theta}, \theta^{*}\right)=\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} E\left(\hat{\theta}, \theta^{*}\right)
$$

We claim

$$
\begin{equation*}
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left(\left\|\Pi_{C}(Y)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{C}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2}\right) \mathbf{1}_{\left\{Y \notin B_{r}\left(\theta^{*}\right)\right\}}\right]=0 \tag{26}
\end{equation*}
$$

for any $r>0$. Applying some basic properties (13) of the projection $\Pi_{\mathcal{C}}$ yields

$$
0 \leq \frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left(\left\|\Pi_{\mathcal{C}}(Y)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{C}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2}\right) \mathbf{1}_{\left\{Y \notin B_{r}\left(\theta^{*}\right)\right\}}\right] \leq \mathbb{E}\left[\|Z\|^{2} \mathbf{1}_{\{\|Z\| \geq r / \sigma\}}\right]
$$

so applying the dominated convergence theorem as before leads to the limit (26).
Thus, it suffices to prove

$$
\begin{align*}
& \lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left\|\Pi_{\mathcal{C}}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \mathbf{1}_{\left\{Y \in B_{r}\left(\theta^{*}\right)\right\}}\right] \\
& \quad=\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left(\left\|\Pi_{\mathcal{C}}(Y)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{C}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2}\right) \mathbf{1}_{\left\{Y \in B_{r}\left(\theta^{*}\right)\right\}}\right] \tag{27}
\end{align*}
$$

for some $r>0$. We choose $r$ as before so that (18) and (22) both hold. By the same reasoning as before, we can replace each instance of $\mathcal{C}$ with $\mathcal{T}$ without changing anything. Furthermore, the condition (18) implies we have $\left\langle\Pi_{\mathcal{T}}(Y)-\Pi_{\mathcal{T}}\left(\theta^{*}\right), \theta^{*}-\Pi_{\mathcal{T}}\left(\theta^{*}\right)\right\rangle=0$ in the event $\left\{Y \in B_{r}\left(\theta^{*}\right)\right\}$, so the Pythagorean inequality (12) becomes equality:

$$
\begin{equation*}
\left\|\Pi_{\mathcal{T}}(Y)-\Pi_{\mathcal{T}}\left(\theta^{*}\right)\right\|^{2} \mathbf{1}_{\left\{Y \in B_{r}\left(\theta^{*}\right)\right\}}=\left(\left\|\Pi_{\mathcal{T}}(Y)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{T}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2}\right) \mathbf{1}_{\left\{Y \in B_{r}\left(\theta^{*}\right)\right\}} \tag{28}
\end{equation*}
$$

Therefore, the equality (27) holds, which concludes the proof of Theorem 3.1.

## 4. Examples

In this section, we assume the Gaussian noise model $Z \sim N\left(0, I_{n}\right)$, or equivalently $Y \sim$ $N\left(\theta^{*}, \sigma^{2} I_{n}\right)$. Thus, $\delta(\cdot)$ denotes the usual statistical dimension [2], where $Z$ in the definition (3) is a standard Gaussian vector.

### 4.1. Nonnegative orthant

We now apply Theorem 3.1 to the nonnegative orthant $\mathbb{R}_{+}^{n}:=\left\{u \in \mathbb{R}^{n}: u_{i} \geq 0, \forall i\right\}$. In Figure 1 we provide visualizations of the geometry of the main theorem when applied to this constraint set.

Corollary 4.1 (Nonnegative orthant). Let $Y \sim N\left(\theta^{*}, \sigma^{2} I\right)$ where $\theta^{*} \in \mathbb{R}^{n}$. Let $n_{+}:=$ $\sum_{i=1}^{n} \mathbf{1}_{\left\{\theta_{i}^{*}>0\right\}}$ and $n_{0}:=\sum_{i=1}^{n} \mathbf{1}_{\left\{\theta_{i}^{*}=0\right\}}$ denote the number of positive components and number of zero components of $\theta^{*}$ respectively. Then the normalized excess risk (6) and normalized mispecified risk (5) of the least squares estimator $\hat{\theta}(Y):=\Pi_{\mathbb{R}_{+}^{n}}(Y)$ with respect to $\mathbb{R}_{+}^{n}$ both tend to

$$
\frac{n_{0}}{2}+n_{+}
$$

as $\sigma \downarrow 0$.
Proof. By Theorem 3.1, it suffices to prove that the statistical dimension term in (16) is $\frac{n_{0}}{2}+n_{+}$. Note that for $y \in \mathbb{R}^{n}, \Pi_{\mathbb{R}_{+}^{n}}(y)=\max \{y, 0\}$ is obtained by taking the component-wise maximum of $y$ with 0 . Consequently,

$$
T_{\mathbb{R}_{+}^{n}}\left(\Pi_{\mathbb{R}_{+}^{n}}\left(\theta^{*}\right)\right)=\left\{u \in \mathbb{R}^{n}: u_{i} \geq 0 \text { if }\left(\Pi_{\mathbb{R}_{+}^{n}}\left(\theta^{*}\right)\right)_{i}=0\right\}=\left\{u \in \mathbb{R}^{n}: u_{i} \geq 0 \text { if } \theta_{i}^{*} \leq 0\right\} .
$$


(a) $\theta=(1,-1) ; \delta=1$

(b) $\theta=(0,-1) ; \delta=1 / 2$

(c) $\theta=(-1,-1) ; \delta=0$

Figure 1. $\mathbb{R}_{+}^{2}$ is marked by the gray area. The intersection $T_{\mathbb{R}_{+}^{2}}\left(\Pi_{\mathbb{R}_{+}^{2}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathbb{R}_{+}^{2}}\left(\theta^{*}\right)\right)^{\perp}$ [translated to be centered at $\Pi_{\mathbb{R}_{+}^{2}}\left(\theta^{*}\right)$ ] is marked by the bold lines in the first two examples, and the bold point in the third example. Each sub-caption states the statistical dimension $\delta=\delta\left(T_{\mathbb{R}_{+}^{2}}\left(\Pi_{\mathbb{R}_{+}^{2}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathbb{R}_{+}^{2}}\left(\theta^{*}\right)\right)^{\perp}\right)$.

Also,

$$
\left(\theta^{*}-\Pi_{\mathbb{R}_{+}^{n}}\left(\theta^{*}\right)\right)^{\perp}=\left\{u \in \mathbb{R}^{n}: \sum_{i: \theta_{i}^{*}<0} \theta_{i}^{*} u_{i}=0\right\}
$$

The intersection is thus

$$
\begin{aligned}
& T_{\mathbb{R}_{+}^{n}}\left(\Pi_{\mathbb{R}_{+}^{n}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathbb{R}_{+}^{n}}\left(\theta^{*}\right)\right)^{\perp} \\
& \quad=\left\{u \in \mathbb{R}^{n}: \begin{array}{l}
u_{i} \geq 0 \quad \text { if } \theta_{i}^{*}=0 \\
u_{i}=0 \\
\text { if } \theta_{i}^{*}<0
\end{array}\right\} \cong \mathbb{R}^{n_{+}} \times \mathbb{R}_{+}^{n_{0}} \times\{0\}^{n-n_{+}-n_{0}}
\end{aligned}
$$

The result follows by noting $\delta(\mathbb{R})=1$ and $\delta\left(\mathbb{R}_{+}\right)=1 / 2$ and by using the fact that $\delta\left(T_{1} \times T_{2}\right)=$ $\delta\left(T_{1}\right)+\delta\left(T_{2}\right)$ for any two cones $T_{1}$ and $T_{2}$ [2].

Remark 4.2. For $\theta^{*} \in \mathbb{R}^{n}$ let $n_{+}$and $n_{0}$ be as defined in Corollary 4.1. Then the low $\sigma$ limit for the corresponding well-specified problem $Y \sim N\left(\Pi_{\mathbb{R}_{+}^{n}}\left(\theta^{*}\right), \sigma^{2} I\right)$ is $\frac{n-n_{+}}{2}+n_{+}$since all negative components of $\theta^{*}$ are sent to zero by $\Pi_{\mathbb{R}_{+}^{n}}$. This is larger than the low $\sigma$ limit for the misspecified problem $Y \sim N\left(\theta^{*}, \sigma^{2} I\right)$ because $n-n_{+} \geq n_{0}$, with strict inequality if $\theta^{*} \notin \mathbb{R}_{+}^{n}$.

### 4.2. Consequences for isotonic regression

This section details interesting consequences of Theorem 3.1 for isotonic regression under misspecification. Let

$$
\mathcal{S}^{n}:=\left\{u \in \mathbb{R}^{n}: u_{1} \leq \cdots \leq u_{n}\right\}
$$

be the monotone cone. We call elements of $\mathcal{S}^{n}$ nondecreasing.
By a block, we refer to a set of the form $\{k, k+1, \ldots, l\}$ for two nonnegative integers $k \leq$ l. Consider a partition of $\{1, \ldots, n\}$ into blocks $I_{1}, \ldots, I_{m}$ listed in increasing order (i.e., the maximum entry of $I_{i}$ is strictly smaller than the minimum entry of $I_{j}$ for $i<j$ ). Let $\left|I_{j}\right|$ denote the cardinality of $I_{j}$ and note that $\sum_{j=1}^{m}\left|I_{j}\right|=n$ as $I_{1}, \ldots, I_{m}$ form a partition of $\{1, \ldots, n\}$. Let $\mathcal{S}_{\left|I_{1}\right|, \ldots,\left.\right|_{m} \mid}$ denote the induced block monotone cone defined as

$$
\begin{equation*}
\mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}:=\left\{u \in \mathcal{S}^{n}: u \text { is constant on each of the blocks } I_{1}, \ldots, I_{m}\right\} . \tag{29}
\end{equation*}
$$

For example,

$$
\mathcal{S}_{2,3,2}=\left\{u \in \mathbb{R}^{2+3+2}: u_{1}=u_{2} \leq u_{3}=u_{4}=u_{5} \leq u_{6}=u_{7}\right\} .
$$

Theorem 3.1 implies the following result, which we prove in Section B.3.
Proposition 4.3 (Isotonic regression). Let $Y \sim N\left(\theta^{*}, \sigma^{2} I\right)$ where $\theta^{*} \in \mathbb{R}^{n}$. Let $\left(J_{1}, \ldots, J_{K}\right)$ be the partition of $\{1, \ldots, n\}$ into blocks such that $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)$ is constant on each $J_{k}$ with respective values $\mu_{1}<\cdots<\mu_{K}$. For each $k \in\{1, \ldots, K\}$, there exists a unique finest partition
$\left(I_{1}^{k}, \ldots, I_{m_{k}}^{k}\right)$ of $J_{k}$ into blocks such that for all $j \in\left\{1, \ldots, m_{k}\right\}$, the mean of the components of $\theta^{*}$ on each $I_{j}^{k}$ equals $\mu_{k}$; that is,

$$
\begin{equation*}
\frac{1}{\left|I_{j}^{k}\right|} \sum_{i \in I_{j}^{k}} \theta_{i}^{*}=\mu_{k}, \quad 1 \leq j \leq m_{k} \tag{30}
\end{equation*}
$$

Then the common low $\sigma$ limit of the normalized excess risk (6) and normalized misspecified risk (5) of the isotonic least squares estimator $\hat{\theta}(Y):=\Pi_{\mathcal{S}^{n}}(Y)$ equals

$$
\sum_{k=1}^{K} \delta\left(\mathcal{S}_{\left|I_{1}^{k}\right|, \ldots,\left|I_{m_{k}}^{k}\right|}\right)
$$

It is clear from the above proposition that the low $\sigma$ behavior of the isotonic estimator under misspecification crucially depends on the statistical dimension of the block monotone cone $\mathcal{S}_{\left|I_{1}^{k}\right|, \ldots,\left|I_{m_{k}}^{k}\right|}$. [We remark again that throughout this section we only deal with the usual statistical dimension, where the noise $Z$ in the definition (3) is standard Gaussian.] Here, we provide two simple properties of the block monotone cone (29), each of which implies that when the block sizes are equal, the statistical dimension is simply that of $\mathcal{S}^{m_{k}}$. The first result provides a direct connection to weighted isotonic regression.

Lemma 4.4 (Weighted isotonic regression). Let $z \in \mathbb{R}^{n}$ and let $I_{1}, \ldots, I_{m}$ be a partition of $\{1, \ldots, n\}$ into blocks. Let $\bar{z}_{I_{j}}:=\frac{1}{\left|I_{j}\right|} \sum_{i \in I_{j}} z_{i}$. Then $\Pi_{\mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}}(y)$ is the vector that is constant on the blocks $I_{1}, \ldots, I_{m}$ with constant values $x_{1}^{*}, \ldots, x_{m}^{*}$, where $x^{*}=\left(x_{1}^{*}, \ldots, x_{m}^{*}\right)$ is

$$
x^{*}=\underset{x \in \mathcal{S}^{m}}{\operatorname{argmin}} \sum_{j=1}^{m}\left|I_{j}\right|\left(x_{i}-\bar{z}_{I_{j}}\right)^{2} .
$$

In other words, the values on the constant blocks of $\Pi_{\mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}}(z)$ can be found by weighted isotonic regression of $\left(\bar{z}_{I_{1}}, \ldots, \bar{z}_{I_{m}}\right) \in \mathbb{R}^{m}$ with weights $\left|I_{1}\right|, \ldots,\left|I_{m}\right|$.

Consequently, when $\left|I_{1}\right|=\cdots=\left|I_{m}\right|$, the statistical dimension of the block monotone cone is

$$
\delta\left(\mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}\right)=\sum_{j=1}^{m} \frac{1}{j} .
$$

The next lemma shows $\mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}$ is isometric to a particular cone in the lower-dimensional space $\mathbb{R}^{m}$.

Lemma 4.5 (Block monotone cone isometry). The block monotone cone $\mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|} \subseteq \mathbb{R}^{n}$ is isometric to

$$
\begin{equation*}
\left\{v \in \mathbb{R}^{m}: \frac{v_{1}}{\sqrt{\left|I_{1}\right|}} \leq \cdots \leq \frac{v_{m}}{\sqrt{\left|I_{m}\right|}}\right\} \subseteq \mathbb{R}^{m}, \tag{31}
\end{equation*}
$$

and thus both sets have the same statistical dimension. In particular, if $\left|I_{1}\right|=\cdots=\left|I_{m}\right|$, then the statistical dimension of the block monotone cone is

$$
\delta\left(\mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}\right)=\sum_{j=1}^{m} \frac{1}{j}
$$

Both lemmas are proved in Section B.1. Note that for the case $\left|I_{1}\right|=\cdots=\left|I_{m}\right|=1$, both lemmas reduce to the statement of the statistical dimension of the monotone cone $\mathcal{S}^{n}$ [2], Eq. D. 12 . More generally, when the $m$ blocks have equal size, the statistical dimension of the associated block monotone cone is the same as that of the monotone cone $\mathcal{S}^{m}$. In Section B.2, we discuss what Lemma 4.5 suggests for the completely general case when the block sizes are arbitrary.

By combining either of these two lemmas with Proposition 4.3, we immediately obtain an explicit expression for the low $\sigma$ limits in a special case. For $m \geq 1$, we denote the harmonic number $\sum_{j=1}^{m}(1 / j)$ by $H_{m}$.

Corollary 4.6 (Isotonic regression with equal sub-block sizes). Consider the setting of Proposition 4.3. In the special case where

$$
\begin{equation*}
\left|I_{1}^{k}\right|=\cdots=\left|I_{m_{k}}^{k}\right| \quad \text { for each } k \in\{1, \ldots, K\} \tag{32}
\end{equation*}
$$

the common low $\sigma$ limit has the following explicit expression:

$$
\sum_{k=1}^{K} H_{m_{k}}=\sum_{k=1}^{K} \sum_{j=1}^{m_{k}} \frac{1}{j}
$$

See the examples to follow (as well as Section B.2) for further discussion about how the statistical dimension of $\mathcal{S}_{\left|I_{1}^{k}\right|, \ldots,\left|I_{m_{k}}^{k}\right|}$ behaves in general, when the special condition (32) does not hold.

In Table 1, we demonstrate how to apply this theorem to various cases of $\theta^{*}$. In the "partition of $\theta^{*}$ " column, we use square brackets to partition the components of $\theta^{*}$ into $K$ blocks according to the constant pieces $\mu_{1}<\cdots<\mu_{K}$ of $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)$, and then within the $k$ th group use parentheses to further partition the components into $m_{k}$ sub-blocks each with common mean $\mu_{k}$.

We now discuss in detail what Proposition 4.3 states for certain cases of $\theta^{*}$.

1. Comparison with well-specified case. In the well-specified case where $\theta^{*} \in \mathcal{S}^{n}$, we have $\theta_{j}^{*}=\mu_{k}$ for all $j \in J_{k}$ and $k \in\{1, \ldots, K\}$, so the finest partition of each $J_{k}$ is the partition into singleton sets. Then $m_{k}=\left|J_{k}\right|$ for each $k$, and moreover $\left|I_{j}^{k}\right|=1$ for all valid $k$ and $j$. Thus, Proposition 4.3 implies that both low $\sigma$ limits are

$$
\sum_{k=1}^{K} H_{\left|J_{k}\right|}:=\sum_{k=1}^{K} \sum_{j=1}^{\left|J_{k}\right|} \frac{1}{j}
$$

Table 1. Examples of how to compute the limit in Proposition 4.3 in the special case (32)

| $\theta^{*}$ | $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)$ | Partition of $\theta^{*}$ | $m_{1}, \ldots, m_{K}$ | $\sum_{k=1}^{K} H_{m_{k}}$ |
| :--- | :--- | :--- | :--- | :--- |
| $(0,0,0,0,0,0)$ | $(0,0,0,0,0,0)$ | $[(0),(0),(0),(0),(0),(0)]$ | 6 | $H_{6}=2.45$ |
| $(1,-1,1,-1,1,-1)$ | $(0,0,0,0,0,0)$ | $[(1,-1),(1,-1),(1,-1)]$ | 3 | $H_{3}=1.8 \overline{3}$ |
| $(5,3,1,-1,-3,-5)$ | $(0,0,0,0,0,0)$ | $[(5,3,1,-1,-3,-5)]$ | 1 | $H_{1}=1$ |
| $(-1,-1,-1,-1$, | $(-1,-1,-1,-1$, | $[(-1),(-1),(-1),(-1)]$, | 4,2 | $H_{4}+H_{2}=3.58 \overline{3}$ |
| $2,2)$ | $2,2)$ | $[(2),(2)]$ |  |  |
| $(0,-2,1,-3,2,2)$ | $(-1,-1,-1,-1,2,2)[(0,-2),(1,-3)],[(2),(2)] 2,2$ | $H_{2}+H_{2}=3$ |  |  |
| $(0,0,-2,-2,3,1)$ | $(-1,-1,-1,-1,2,2)[(0,0,-2,-2)],[(3,1)]$ | 1,1 | $H_{1}+H_{1}=2$ |  |

This is precisely the upper bound (7) for the monotone cone as computed by Bellec [3], Prop. 3.1, so we recover the low $\sigma$ limit (4). Computations for the well-specified examples $\theta^{*}=(0,0,0,0,0,0)$ and $\theta^{*}=(-1,-1,-1,-1,2,2)$ appear in Table 1 .

Now, consider the misspecified problem $Y \sim N\left(\theta^{*}, \sigma^{2} I_{n}\right)$ with $\theta^{*} \notin \mathcal{S}^{n}$, and compare the statement of Proposition 4.3 with the corresponding statement for the well-specified problem $Y \sim N\left(\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right), \sigma^{2} I\right)$. In both cases, the partition of $\{1, \ldots, n\}$ into $\left(J_{1}, \ldots, J_{K}\right)$ is the same. However, we showed above that in the well-specified problem, the sub-partition of each $J_{k}$ consists of singletons, whereas for the misspecified problem we may get nontrivial partitions $\left(I_{1}^{k}, \ldots, I_{m_{k}}^{k}\right)$. Noting the inclusion $\mathcal{S}_{\left|I_{1}^{k}\right|, \ldots,\left|I_{m_{k}}^{k}\right|} \subseteq \mathcal{S}^{J_{k} \mid}$ for each $k$ and comparing (4.3) for the two cases yields

$$
\sum_{k=1}^{K} \delta\left(\mathcal{S}_{\left|I_{1}^{k}\right|, \ldots,\left|I_{m_{k}}^{k}\right|}\right) \leq \sum_{k=1}^{K} \delta\left(\mathcal{S}^{\left|J_{k}\right|}\right)
$$

which shows that in general the misspecified low $\sigma$ limit is smaller than the corresponding well-specified limit.
2. Decreasing sequences. Suppose $\theta^{*}$ is nonincreasing and nonconstant i.e., $\theta^{*} \in\left(-\mathcal{S}^{n}\right) \backslash$ $\mathcal{S}^{n}$. Then $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)$ is constant (see [10] for various properties of $\Pi_{\mathcal{S}^{n}}$ ), so $K=1$ and $\mu_{1}=\frac{1}{n} \sum_{i=1}^{n} \theta_{i}^{*}$. We also claim $m_{1}=1$. Indeed, if $m_{1}>1$ then there exists some $j<n$ such that $\mu=\frac{1}{j} \sum_{i=1}^{j} \theta_{i}^{*}=\frac{1}{n-j} \sum_{i=j+1}^{n} \theta_{i}^{*}$. However, the fact that $\theta^{*}$ is nonincreasing and nonconstant implies $\frac{1}{j} \sum_{i=1}^{j} \theta_{i}^{*}>\frac{1}{n-j} \sum_{i=j+1}^{n} \theta_{i}^{*}$, a contradiction. Thus, Proposition 4.3 implies that both low $\sigma$ limits are

$$
\begin{equation*}
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} M\left(\hat{\theta}, \theta^{*}\right)=\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} E\left(\hat{\theta}, \theta^{*}\right)=1 . \tag{33}
\end{equation*}
$$

(In fact, by combining the above argument with the proof of Proposition 4.3, we have shown that the intersection $T_{\mathcal{S}^{n}}\left(\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right)^{\perp}$ is simply the subspace of constant sequences.) On the other hand, since $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)$ is constant, the low $\sigma$ limit in the well-specified setting $Y \sim N\left(\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right), \sigma^{2} I_{n}\right)$ is $\sum_{j=1}^{n} \frac{1}{j} \asymp \log n$, which is much larger.

The logarithmic term appears here in the well-specified case due to the well-known spiking effect of isotonic regression (documented, for example, by Pal [9], Wu et al. [12] and Zhang [13]). Indeed, the isotonic estimator is inconsistent near the end points which leads to the logarithm term in the risk. However, in the misspecified case when $\theta^{*}$ is nonincreasing and nonconstant, a combination of the proof of Theorem 3.1 (in particular Lemma 3.5) with the fact that $T_{\mathcal{S}^{n}}\left(\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right)^{\perp}$ is the subspace of all constant sequences implies $\hat{\theta}(Y)$ is a constant sequence with probability increasing to 1 as $\sigma \downarrow 0$, in which case the constant value must be the sample mean $\bar{Y}:=\frac{1}{n} \sum_{i=1}^{n} Y_{i}$. Alternatively, one can rephrase the geometric argument in Lemma 3.5 more simply in this example; when $\sigma$ is small, $Y$ is near $\theta^{*}$ and thus is also nondecreasing with high probability, in which case $\hat{\theta}(Y)$ is constant, due to the properties of the projection $\Pi_{\mathcal{S}^{n}}$. Hence, in this situation the estimator is essentially a linear projection onto a one-dimensional subspace (hence the low noise limit of 1) and does not suffer from any spiking at the endpoints, and consequently there are no logarithmic terms in the risk in the misspecified case in the low sigma limit.

Computations for the specific example when $\theta^{*}=(5,3,1,-1,-3,-5)$ appear in Table 1.
3. In the first half of Table 1, we consider three choices for $\theta^{*}$ that project to $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)=$ $(0,0,0,0,0,0)$. Here $K=1$ and the sub-block sizes $\left|I_{1}^{1}\right|, \ldots,\left|I_{m_{1}}^{1}\right|$ are equal in each case (namely, the common block size is 1,2 , and 6 respectively), so we are in the special case (32). Thus, the limit is $\sum_{j=1}^{m_{1}} \frac{1}{j}$ where $m_{1}$ is the number of sub-blocks. We see that for the misspecified $\theta^{*}$ the low $\sigma$ limits are smaller.

One can heuristically interpret Theorem 3.1 for the example, $\theta^{*}=(1,-1,1,-1,1,-1)$ as follows. With probability increasing to 1 as $\sigma \downarrow 0$, the estimator $\hat{\theta}(Y)$ is nondecreasing and piecewise constant on three equally sized blocks, so the low $\sigma$ limit is the same as if we were estimating $(0,0,0)$ in $\mathcal{S}^{3}$.
4. Similarly in the second half of Table 1 , we consider three $\theta^{*}$ that project to $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)=$ $(-1,-1,-1,-1,2,2)$. Here $K=2$ but, since the low $\sigma$ limit decomposes, we can simply consider each constant piece separately. Again, we see that the more sub-blocks $I_{i}^{j}$, the higher the statistical dimension, with the well-specified case having the most sub-blocks (all singletons).
5. Beyond equi-sized sub-blocks. The concrete examples we have considered so far have been in the special case (32). In a few other cases we can still provide the low $\sigma$ limit. (See also Section B. 2 for further discussion.)
(a) If $K=1$ and $m_{1}=2$, then the low $\sigma$ limit is $\delta\left(S_{\left|I_{1}^{1}\right|,\left|I_{2}^{1}\right|}\right)$. By Lemma 4.5, this is the same as the statistical dimension of the half space $\left\{u \in \mathbb{R}^{2}: u_{1} / \sqrt{\left|I_{1}\right|} \leq u_{2} / \sqrt{\left|I_{2}\right|}\right\}$, which is 1.5 . However, when $m_{1}>2$, it is difficult to compute $\delta\left(S_{\left|I_{1}^{1}\right|, \ldots \mid I_{m_{1}}^{1}}\right)$ unless we are in the special case $\left|I_{1}^{1}\right|=\cdots=\left|I_{m_{1}}^{1}\right|$.
(b) In some other extreme cases, we can get an approximation. For example, if

$$
\theta^{*}=(0, \underbrace{1, \ldots, 1}_{(n-2) / 2}, \underbrace{-1, \ldots,-1}_{(n-2) / 2}, 0),
$$

then $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)=(0, \ldots, 0)$, so the low $\sigma$ limit is $\delta\left(\mathcal{S}_{1, n-2,1}\right)$. Lemma 4.5 shows that this is the same as the statistical dimension of $\left\{u \in \mathbb{R}^{3}: u_{1} \leq u_{2} / \sqrt{n-2} \leq u_{3}\right\}$. As
$n \rightarrow \infty$ tends to this set tends to $\left\{u \in \mathbb{R}^{3}: u_{1} \leq 0 \leq u_{3}\right\}$ which has statistical dimension $1+\frac{1}{2}+\frac{1}{2}=2$. Thus, $\delta\left(\mathcal{S}_{1, n-2,1}\right) \rightarrow 2$ as $n \rightarrow \infty$. We used simulations to verify that the low $\sigma$ limit is indeed near 2 even for $n=20$.

### 4.3. An example in convex regression

In the previous section, we showed (33) that when $\theta^{*}$ is nonincreasing and nonconstant, the low noise limits for isotonic regression are 1 . In this section, we prove the analogous result in the setting of convex regression.

We fix design points $x_{1}<\cdots<x_{n}$. In this section, we use the shorthand $x:=\left(x_{1}, \ldots, x_{n}\right)$. The relevant cone for convex regression on this design is

$$
\begin{equation*}
\mathcal{K}_{x}:=\left\{\theta \in \mathbb{R}^{n}: \frac{\theta_{i}-\theta_{i-1}}{x_{i}-x_{i-1}} \leq \frac{\theta_{i+1}-\theta_{i}}{x_{i+1}-x_{i}}, i=2, \ldots, n-1\right\} . \tag{34}
\end{equation*}
$$

The constraints naturally arise when considering $\theta_{i}=f\left(x_{i}\right)$ for some convex function $f$.
Unlike in the case of isotonic regression, the statistical dimension of tangent cones of $\mathcal{K}_{x}$ is not known exactly. However, upper bounds are available and the best known upper bound is due to Bellec [3], Prop. 4.2, who proved that for every $\theta \in \mathcal{K}_{x}$,

$$
\begin{equation*}
\delta\left(T_{\mathcal{K}_{x}}(\theta)\right) \leq 8 q \log \frac{e n}{q} \tag{35}
\end{equation*}
$$

where $q$ is the number of affine pieces in $\theta$ (i.e., one plus the number of inequalities in the definition (34) of $\mathcal{K}_{x}$ that hold with strict inequality).

Because finding an exact formula for $\delta\left(T_{\mathcal{K}_{x}}\left(\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)\right)\right)$ seems intractable, it is unlikely that an exact formula for the statistical dimension of $T_{\mathcal{K}_{x}}\left(\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)\right)^{\perp}$ can be found. As a result, in contrast to the case of isotonic regression, it might not be possible to precisely determine the limits in (16) for most $\theta^{*} \in \mathbb{R}^{n}$ in convex regression. However in the special case where $\theta^{*}$ is concave and nonlinear, it is indeed possible to characterize the limits in (16). This is the content of the next result.

Proposition 4.7. Suppose $\theta^{*} \in\left(-\mathcal{K}_{x}\right) \backslash \mathcal{K}_{x}$ is concave and nonlinear. Then the common low $\sigma$ limit of the normalized excess risk (6) and normalized misspecified risk (5) of the convex least squares estimator $\hat{\theta}(Y):=\Pi_{\mathcal{K}_{x}}(Y)$ is 2.

We prove this result in Appendix C. The intuition behind this result is completely analogous to the aforementioned example in isotonic regression. Here, when $\sigma$ is small, $Y$ is concave and nonlinear with high probability, and gets projected onto the subspace span $\{(1, \ldots, 1), x\}$ of affine vectors on $x$, which has dimension 2 .

Note that for the setting of Proposition 4.7, the analogous well-specified problem is $Y \sim$ $N\left(\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right), \sigma I_{n}\right)$, which by the upper bound (7) (and the bound (35)) has low noise limit $\delta\left(T_{\mathcal{K}_{x}}\left(\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)\right)\right) \leq 8 \log ($ en $)$. So, in this situation, the misspecified low noise limits, 2 , are much smaller than the corresponding well-specified limit, $\asymp \log n$ (assuming the upper bound (35) is tight). Recall the analogous situation (33) for decreasing $\theta^{*}$ in isotonic regression, where the misspecified and well-specified limits are 1 and $\asymp \log n$, respectively.

## 5. Further discussion

### 5.1. Generalizing Theorem 3.1 to the non-polyhedral case

Note that Theorem 3.1 requires the condition (15) that is, that $\mathcal{C}$ is locally a polyhedron near $\Pi_{\mathcal{C}}\left(\theta^{*}\right)$. Here we comment on the situation when $\mathcal{C}$ is non-polyhedral. Although non-polyhedral convex sets can be approximated by polyhedra, the low $\sigma$ limit magnifies the local geometry of the set and ignores the goodness of such an approximation. As a stark counterexample, consider any closed convex $\mathcal{C} \subseteq \mathbb{R}^{2}$ with nonempty interior, and let $Z \sim N\left(0, I_{n}\right)$. For any polygon in $\mathbb{R}^{2}$, Theorem 3.1 implies that the low $\sigma$ limits are either $0,1 / 2$, or 1 because in $\mathbb{R}^{2}$ the intersection of a convex cone with a line intersecting the origin is either the origin, a ray, or a line. Thus, for a sequence of polygons approximating $\mathcal{C}$ the sequence of corresponding low $\sigma$ limits need not even have a limit, never mind the matter of two different sequences of polygonal approximations having a common limit. Therefore, the low $\sigma$ limit for general $\mathcal{C}$ cannot be found using a polyhedral approximation.

In order to understand how the low $\sigma$ limits behave for general $\mathcal{C}$, we consider the following specific example. Let $\mathcal{C}:=\left\{\theta \in \mathbb{R}^{n}:\|\theta\| \leq 1\right\}$ be the unit ball so that $\Pi_{\mathcal{C}}(x)=\frac{x}{\max \{\|x\|, 1\}}$. Also let $\theta^{*}:=(r, 0, \ldots, 0)$ for some $r>1$ so that $\Pi_{\mathcal{C}}\left(\theta^{*}\right)=(1,0, \ldots, 0)$. By rotational symmetry of $\mathcal{C}$, the case of any general $\theta^{*} \notin \mathcal{C}$ can be reduced to this case.

In the corresponding well-specified case $Y \sim N\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right), \sigma^{2} I_{n}\right.$ ), the result (4) of Oymak and Hassibi [8] implies that the normalized misspecified risk (5) and the normalized excess risk (6) are equal in the low $\sigma$ limit with common value

$$
\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right)=n-\frac{1}{2}
$$

since the tangent cone is the half space $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)=\left\{x \in \mathbb{R}^{n}: x_{1} \leq 0\right\}$.
However, in the misspecified case, we observe some new phenomena that do not occur for polyhedra.

Proposition 5.1 (Low noise limits for the ball). Let $\mathcal{C}:=\left\{\theta \in \mathbb{R}^{n}:\|\theta\|_{2} \leq 1\right\}, \theta^{*} \notin \mathcal{C}$, and $Y \sim N\left(\theta^{*}, \sigma^{2} I_{n}\right)$. For the estimator $\hat{\theta}(Y)=\Pi_{\mathcal{C}}(Y)$, we have

$$
\begin{align*}
& \lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} M\left(\hat{\theta}, \theta^{*}\right)=\frac{n-1}{\left\|\theta^{*}\right\|^{2}}  \tag{36a}\\
& \lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} E\left(\hat{\theta}, \theta^{*}\right)=\frac{n-1}{\left\|\theta^{*}\right\|} \tag{36b}
\end{align*}
$$

The proof involves direct computation and appears in Appendix D.
We now highlight some of the interesting behavior. In the polyhedral case, both limits were equal; in the proof of Theorem 3.1 (in particular Lemma 3.5) we showed that with probability increasing to 1 (in the low $\sigma$ limit), $Y$ would be projected onto the hyperplane $\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$, producing the orthogonality required for the Pythagorean inequality (12) to become an equality.

In the general case, the Pythagorean inequality is not tight, and we explicitly see from this example that even in the low noise limit the the excess risk can be strictly larger than the misspecified risk.

Note that in contrast to the corresponding well-specified case $Y \sim N\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right), \sigma^{2} I_{n}\right)$ which has limit $n-\frac{1}{2}$, the misspecified limits $\frac{n-1}{\left\|\theta^{*}\right\|^{2}}$ and $\frac{n-1}{\left\|\theta^{*}\right\|}$ both tend to $n-1$ as $\left\|\theta^{*}\right\| \downarrow 1$, so there is a "jump" in the limits between the misspecified and well-specified setting. This is also a feature of Theorem 3.1 when the polyhedron $\mathcal{C}$ has nonempty interior, as we discussed earlier (see Lemma 3.4).

This example shows that Theorem 3.1 does not hold for nonpolyhedral constraint sets $\mathcal{C}$, as the two normalized risks are not equal in this particular example of the unit ball, and moreover neither limit equals

$$
\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}\right)=\delta\left(\left\{u:\left\langle u, \theta^{*}\right\rangle \leq 0\right\} \cap\left(\theta^{*}\right)^{\perp}\right)=\delta\left(\left(\theta^{*}\right)^{\perp}\right)=n-1 .
$$

The intuition for Theorem 3.1 is that, in the polyhedral case, the projections of $Y$ largely end up in some face of the polyhedron $\mathcal{C}$, which can be approximated by a lower-dimensional cone, for which the statistical dimension is well defined. When $\mathcal{C}$ is not polyhedral, the generalization of this "face" is hard to conceptualize and is likely not well approximated by a cone, so a statistical dimension can not be even applied. Indeed, for general $\mathcal{C}$ such as the ball, tangent cones are extremely poor approximations for the set. Contrary to this drawback, the result (4) of Oymak and Hassibi [8] shows that tangent cones are good enough for the well-specified setting. However for the misspecified setting, we expect that any general result for the low $\sigma$ limits does not involve a statistical dimension of some cone, since the surface of $\mathcal{C}$ is the essential object of interest and cannot be approximated by some cone except in special settings like the polyhedral case.

As mentioned already, Theorem 3.1 shows that in the misspecified setting, the upper bound (7), which holds for all $\sigma$, is not tight in the low $\sigma$ limit. One might ask whether a better upper bound for all $\sigma$ can be achieved, but Figure 2 shows that for some values of $\sigma$ the risks can be close to the upper bound, represented by the solid horizontal line. We observed this behavior in other examples (see also Figure 5): the risks can be close to the upper bound for some moderate values of $\sigma$, and then converge to the strictly smaller low $\sigma$ limit. Replacing the upper bound (7), which is constant in $\sigma$, with a $\sigma$-dependent upper bound would be an interesting result, but it would have to be extremely dependent on the geometry of the set $\mathcal{C}$. In the following sections, we further discuss the normalized risks as a function of $\sigma$.

### 5.2. High noise limit

Although not interesting in its own right, the high noise limit of the normalized risks can help characterize the maximum risks (10), as we discuss in Section 5.3 below. Proofs for this section appear in Appendix E.

For a closed convex set $\mathcal{C}$, we define the core cone

$$
\begin{equation*}
K_{\mathcal{C}}:=\bigcap_{\theta \in \mathcal{C}} T_{\mathcal{C}}(\theta) . \tag{37}
\end{equation*}
$$



Figure 2. Empirical estimates of the normalized misspecified risk (•) and normalized excess risk ( $\mathbf{\Delta})$ plotted against $\log _{10}(\sigma)$, for the ball $\mathcal{C}=\left\{\theta \in \mathbb{R}^{n}:\|\theta\| \leq 1\right\}$ in the case $n=3$ with $\theta^{*}=(1+\epsilon, 0,0)$ and $\epsilon \in\{0.01,0.1,1\}$. The solid horizontal line represents the upper bound $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right)=n-\frac{1}{2}=2.5$ guaranteed by (7). The dotted lines and dashed lines are the predicted low $\sigma$ limits $\frac{n-1}{(1+\epsilon)^{2}}$ and $\frac{n-1}{1+\epsilon}$, respectively. The dash-dot line is the high $\sigma$ limit 0 .

Recall the notation for the re-centered set $F_{\mathcal{C}}\left(\theta_{0}\right)=\left\{\theta-\theta_{0}: \theta \in \mathcal{C}\right\}$ where $\theta_{0} \in \mathcal{C}$. For a vector $v \in \mathbb{R}^{n}$ we let $\mathbb{R}_{+} v:=\{\alpha u: \alpha \geq 0\}$. We have the following equivalent characterizations of the core cone.

Lemma 5.2 (Characterizations of the core cone). Let $\mathcal{C} \subseteq \mathbb{R}^{n}$ be a closed convex set. For any $\theta_{0} \in \mathcal{C}$,

$$
K_{\mathcal{C}} \stackrel{(\mathrm{i})}{=}\left\{v: \mathbb{R}_{+} v \subseteq F_{\mathcal{C}}\left(\theta_{0}\right)\right\} \stackrel{(\mathrm{ii})}{=} \bigcap_{\sigma>0} \frac{F_{\mathcal{C}}\left(\theta_{0}\right)}{\sigma} .
$$

Additionally, the inclusion $K_{\mathcal{C}} \subseteq T_{\mathcal{C}}(\theta)$ holds for any $\theta \in \mathcal{C}$. Iffurthermore $F_{\mathcal{C}}\left(\theta_{0}\right)$ is a cone, then the equality $K_{\mathcal{C}}=T_{\mathcal{C}}(\theta)$ holds if and only if $\theta_{0}-\left(\theta-\theta_{0}\right) \in \mathcal{C}$; in particular, taking $\theta=\theta_{0}$ shows that $K_{\mathcal{C}}=T_{\mathcal{C}}\left(\theta_{0}\right)=F_{\mathcal{C}}\left(\theta_{0}\right)$.

Thus, up to a translation, the core cone can either be viewed as the result of shrinking $\mathcal{C}$ radially toward $\theta_{0} \in \mathcal{C}$, or as the largest cone centered at $\theta_{0} \in \mathcal{C}$ that is contained in $\mathcal{C}$. An interesting point is that $\theta_{0} \in \mathcal{C}$ can be chosen arbitrarily. An example is given in Figure 3 .

Furthermore, in the case when $\mathcal{C}$ is a cone, the core cone $K_{\mathcal{C}}$ simply equals this cone $\mathcal{C}$, and we can characterize which tangent cones are the "smallest" in the sense that they equal the intersection (37) of all tangent cones.

The following result shows that under a boundedness condition, the core cone characterizes both high $\sigma$ limits.

(a) $\mathcal{C}=\{(x, y): x \geq 0, y \geq \max (x, 1)\}$

(b) $K_{\mathcal{C}}=\{(x, y): 0 \leq x \leq y\}$

Figure 3. Example of a closed convex set $\mathcal{C}$ and its core cone $K_{\mathcal{C}}$. By the definition (37) and Lemma 5.2, we can equivalently think of the core cone as (1) the intersection of all tangent cones of $\mathcal{C}$, (2) the result of radially shrinking $\mathcal{C}$ toward any given point $\theta_{0} \in \mathcal{C}$, or (3) the collection of all rays that lie entirely in $\mathcal{C}$ with a given initial point $\theta_{0} \in \mathcal{C}$.

Proposition 5.3 (High noise limit). Let $\mathcal{C}$ be a closed convex set. Let $\theta^{*} \in \mathbb{R}^{n}$ and $Y:=\theta^{*}+\sigma Z$ where $Z$ is a zero mean random vector with $\mathbb{E}\|Z\|^{2}<\infty$. If the condition

$$
\begin{equation*}
\sup _{x \in \mathbb{R}^{n}}\left(\left\|\Pi_{F_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)}(x)\right\|^{2}-\left\|\Pi_{K_{\mathcal{C}}}(x)\right\|^{2}\right)<\infty \tag{38}
\end{equation*}
$$

holds, then

$$
\lim _{\sigma \rightarrow \infty} \frac{1}{\sigma^{2}} M\left(\hat{\theta}, \theta^{*}\right)=\lim _{\sigma \rightarrow \infty} \frac{1}{\sigma^{2}} E\left(\hat{\theta}, \theta^{*}\right)=\delta\left(K_{\mathcal{C}}\right)
$$

The main hurdle in applying Proposition 5.3 is verifying the condition (38). The following result covers two cases where it is easy to verify the condition.

Corollary 5.4 (Orthant and bounded sets). Let $\theta^{*} \in \mathbb{R}^{n}$ and $Y \sim N\left(\theta^{*}, \sigma^{2} I_{n}\right)$.

- If $\mathcal{C}=\mathbb{R}_{+}^{n}$ is the nonnegative orthant, then the high $\sigma$ limits are $\delta\left(\mathbb{R}_{+}^{n}\right)=n / 2$.
- Let $\mathcal{C}$ be a closed convex set. $K_{\mathcal{C}}=\{0\}$ if and only if $\mathcal{C}$ is bounded, in which case both high $\sigma$ limits are 0 .

Figure 2 and Figure 5 illustrate the result of this corollary.
Verifying the condition (38) for more general $\mathcal{C}$ is more difficult. We believe it might hold for polyhedral cones with any $\theta^{*}$, in which case Proposition 5.3 would imply that the high $\sigma$ limits are $\delta(\mathcal{C})$. For example, we are unable to verify that the monotone cone $\mathcal{C}=\mathcal{S}^{n}$ satisfies the condition (38), but our simulations in Figure 4 agree with the implication of Proposition 5.3 that the high noise limit is $\delta\left(\mathcal{S}^{n}\right)=H_{n}$. An interesting feature of the examples presented thus far is that the high $\sigma$ limits (including the veracity of (38)) do not depend on $\theta^{*}$.

Remark 5.5. More generally, suppose $\mathcal{C}$ is a general cone. By applying Lemma 5.2 with $\theta_{0}=0$ and $\theta=\Pi_{\mathcal{C}}\left(\theta^{*}\right)$, we observe that the core cone $K_{\mathcal{C}}$ is $\mathcal{C}$, and moreover $\mathcal{C} \subseteq T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right.$ ), with equality if and only if $-\Pi_{\mathcal{C}}\left(\theta^{*}\right) \in \mathcal{C}$. Thus, if the condition (38) holds, then Proposition 5.3 implies the high $\sigma$ limits are $\delta(\mathcal{C})$, and moreover Lemma 5.2 implies that these limits equal Bellec's upper bound (7), $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right.$ ), if and only if $\theta^{*}$ satisfies $-\Pi_{\mathcal{C}}\left(\theta^{*}\right) \in \mathcal{C}$.


Figure 4. Empirical estimates of the normalized misspecified risk ( $\bullet$ ) and normalized excess risk ( $\mathbf{\Delta}$ ) plotted against $\log _{10}(\sigma)$, for the monotone cone $\mathcal{C}=\mathcal{S}^{n}$ in the case $n=6$, for two different values of $\theta^{*}$. The solid horizontal lines represent the upper bound $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right.$ ) guaranteed by ( 7 ), which are $\delta\left(\mathcal{S}^{6}\right)=H_{6}=2.45$ and $\delta\left(\mathcal{S}^{4} \times \mathcal{S}^{2}\right)=H_{4}+H_{2}=3.58 \overline{3}$, respectively. The dashed lines are the predicted low $\sigma$ limits, which are 1 and 2, respectively. (See Section 4.2 for details on the computations.) The dash-dot line is the predicted high $\sigma \operatorname{limit} \delta\left(\mathcal{S}^{6}\right)=H_{6}=2.45$ (see discussion below Corollary 5.4).

However, the condition (38) does not hold for all $\mathcal{C}$. One can verify numerically that the epigraph $\mathcal{C}:=\left\{u \in \mathbb{R}^{2}: u_{2} \geq u_{1}^{2}\right\}$, whose core cone is $K_{\mathcal{C}}=\left\{\left(0, u_{2}\right): u_{2} \geq 0\right\}$, does not satisfy (38). Simulations also show that the high $\sigma$ limits are larger than $\delta\left(K_{\mathcal{C}}\right)=1 / 2$. In general, it is unclear exactly when the core cone does or does not characterize the high $\sigma$ limits.

### 5.3. Maximum normalized risk

Our low and high $\sigma$ limit results Theorem 3.1 and Proposition 5.3 provides an incomplete characterization of the maximum normalized risks (10). As mentioned already in (7), $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right)$ is an upper bound for both suprema.

In the well-specified case $\theta^{*} \in \mathcal{C}$, both suprema reduce to the usual normalized risk $\sigma^{-2} R\left(\hat{\theta}, \theta^{*}\right)$; moreover the upper bound becomes $\delta\left(T_{\mathcal{C}}\left(\theta^{*}\right)\right)$, and is attained as $\sigma \downarrow 0$ by the result (4) of Oymak and Hassibi [8].

However, in the misspecified case we have shown in Theorem 3.1 that in general the low $\sigma$ limit does not attain the upper bound (7). Moreover, simulations show that in some cases even


Figure 5. Empirical estimates of the normalized misspecified risk ( $\bullet$ ) and normalized excess risk ( $\mathbf{\Delta}$ ) plotted against $\log _{10}(\sigma)$, for the orthant $\mathcal{C}:=\mathbb{R}_{+}^{3}$ and $\theta^{*}=(1,1,-\epsilon)$ with $\epsilon \in\{0.01,0.1,1\}$. The solid horizontal line represents the upper bound $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right)=n-\frac{1}{2}$ guaranteed by (7). The dashed line is the common low $\sigma$ limit $n-1$ (see Corollary 4.1). The dash-dot line is the high $\sigma \operatorname{limit} \delta\left(\mathbb{R}_{+}^{n}\right)=3 / 2$.
the suprema do not attain the upper bound; see Figure 2 and Figure 5. We see that for some cases the suprema are close to the upper bound, but for others it is much smaller.

Of course, if one can show that either the low $\sigma$ limit or the high $\sigma$ limit is equal to the upper bound $\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right)$, then we know the upper bound is attained either as $\sigma \downarrow 0$ or $\sigma \rightarrow \infty$, respectively. However, in the settings of Theorem 3.1 and Proposition 5.3, this seldom happens. As discussed already, if $\mathcal{C}$ is polyhedral with nonempty interior, then the low $\sigma$ limit is strictly smaller than the upper bound. If Proposition 5.3 applies, then $K_{\mathcal{C}}=\bigcap_{\theta \in \mathcal{C}} T_{\mathcal{C}}(\theta) \subseteq T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$ shows that the high $\sigma$ limit is typically strictly smaller than the upper bound; for the special case where $\mathcal{C}$ is a cone, see Remark 5.5 for a necessary and sufficient condition for the high $\sigma$ limit to equal the upper bound.

Thus, in most cases the suprema are attained at some moderate values of $\sigma$, but it is difficult to provide a characterization of these maximizing values $\sigma$, as well as the value of the suprema and whether they are close to the upper bound or not. The plots suggest that as $\theta^{*}$ gets closer to $\mathcal{C}$, the suprema get closer to the upper bound as well.

## Appendix A: Proofs of lemmas in Section 3

The next lemma is a technical device for representing the largest face of a polyhedral cone that lies in a particular hyperplane. It is useful for proving Lemma 3.3 and Lemma 3.5.

Lemma A. 1 (Largest face in hyperplane). Let $\mathcal{K}=\{u: A u \leq 0\} \subseteq \mathbb{R}^{n}$ be a polyhedral cone, where $A \in \mathbb{R}^{m \times n}$ has distinct rows. For each $y \in \mathbb{R}^{n}$, consider the subsets $J \subseteq\{1, \ldots, m\}$ satis-
fying

$$
\begin{equation*}
\left\{u: A_{J} u=0\right\} \subseteq\left(y-\Pi_{\mathcal{K}}(y)\right)^{\perp} . \tag{39}
\end{equation*}
$$

We let $J_{y}$ denote the smallest such subset.
This subset $J_{y}$ characterizes a face of $\mathcal{K}$ in the following way.

$$
\begin{equation*}
\mathcal{K} \cap\left(y-\Pi_{\mathcal{K}}(y)\right)^{\perp}=\left\{u: A_{J_{y}} u=0, A_{J_{y}^{c}} u \leq 0\right\} . \tag{40}
\end{equation*}
$$

Proof. The optimality condition for a projection onto a cone (14) implies $\left\langle y-\Pi_{\mathcal{K}}(y), u\right\rangle \leq 0$ for all $u \in \mathcal{K}$. If $\mathcal{K}$ contains both $u$ and $-u$, then this implies $u \in\left(y-\Pi_{\mathcal{K}}(y)\right)^{\perp}$. Thus, for $J=\{1, \ldots, m\}$, (39) holds because $\left\{u: A_{J} u=0\right\} \subseteq \mathcal{K}$. This shows the existence of subsets $J$ that satisfy (39).

Next, note that if $J$ and $J^{\prime}$ both satisfy (39), then $J \cap J^{\prime}$ does as well, because

$$
\left\{u: A_{J \cap J^{\prime}} u=0\right\}=\left\{u+v: A_{J} u=A_{J^{\prime}} v=0\right\} \subseteq\left(y-\Pi_{\mathcal{K}}(y)\right)^{\perp}
$$

So, letting $J_{y}$ be the intersection of all $J$ satisfying (39) yields the unique subset of minimal size.
The $\supseteq$ inclusion in (40) follows immediately from $\left\{u: A_{J_{y}} u=0\right\} \subseteq\left(y-\Pi_{\mathcal{K}}(y)\right)^{\perp}$. For the other inclusion, suppose $v \in \mathcal{K} \cap\left(y-\Pi_{\mathcal{K}}(y)\right)^{\perp}$. Then $A v \leq 0$, so it remains to verify $A_{J_{y}} v=0$. That is, if $J \subseteq\{1, \ldots, m\}$ denotes the indices $j$ for which $\left\langle a_{j}, v\right\rangle=0$, we want to show $J_{y} \subseteq J$; furthermore, this reduces to showing $J$ satisfies (39), by minimality of $J_{y}$.

Any $u$ satisfying $A_{J} u=0$ can be rewritten as $u=v+w$ for some $w$ also satisfying $A_{J} w=0$. There exists some $c>0$ such that both $v+c w$ and $v-c w$ are in $\mathcal{K}$ because all the linear constraints outside of $J$ are strict inequalities at $v$. Then, the optimality condition for the projection onto a cone, yields $\left\langle v+c w, y-\Pi_{\mathcal{K}}(y)\right\rangle \leq 0$ and $\left\langle v-c w, y-\Pi_{\mathcal{K}}(y)\right\rangle \leq 0$. Since $v \in\left(y-\Pi_{\mathcal{K}}(y)\right)^{\perp}$, this yields $w \in\left(y-\Pi_{\mathcal{K}}(y)\right)^{\perp}$ and thus $u \in\left(y-\Pi_{\mathcal{K}}(y)\right)^{\perp}$, which verifies that $J$ satisfies (39).

Proof of Lemma 3.2. By definition there exist an integer $m$, matrix $A \in \mathbb{R}^{m \times n}$, and vector $b \in \mathbb{R}^{m}$ such that $\mathcal{C}:=\left\{u \in \mathbb{R}^{n}: A u \leq b\right\}$. Fix $\theta^{*} \in \mathbb{R}^{n}$ and let $\theta_{0}:=\Pi_{\mathcal{C}}\left(\theta^{*}\right)$. We will show

$$
T_{\mathcal{C}}\left(\theta_{0}\right)=\left\{u: A_{J} u \leq 0\right\}
$$

where $J=\left\{j:\left\langle a_{j}, \theta_{0}\right\rangle=b_{j}\right\}$. Then $T_{\mathcal{C}}\left(\theta_{0}\right)$ is a polyhedral cone.
If $u \in T_{\mathcal{C}}\left(\theta_{0}\right)$, then for some $r^{*}>0$ we have $\theta_{0}+r^{*} u \in \mathcal{C}$. Thus, $b_{J} \geq A_{J}\left(\theta_{0}+r u\right)=b_{J}+$ $r^{*} A_{J} u$ which implies $A_{J} u \leq 0$.
Conversely, suppose $u$ satisfies $A_{J} u \leq 0$. Choose $r^{*}>0$ so that $r^{*}\left\langle a_{j}, u\right\rangle \leq b_{j}-\left\langle a_{j}, \theta_{0}\right\rangle$ for all $j \notin J$. This is possible because $b_{j}>\left\langle a_{j}, \theta_{0}\right\rangle$ for each $j \notin J$. Then $\theta_{0}+r^{*} u \in \mathcal{C}$ so $u \in T_{\mathcal{C}}\left(\theta_{0}\right)$.

Finally, we need to prove the second part of the locally polyhedral condition (15), which will follow if we show $T_{\mathcal{C}}\left(\theta_{0}\right) \cap B_{r^{*}}(0) \subseteq F_{\mathcal{C}}\left(\theta_{0}\right)$ for some $r^{*}>0$. If $u \in T_{\mathcal{C}}\left(\theta_{0}\right)$, then $A_{J} u \leq 0=$ $b_{J}-A_{J} \theta_{0}$, so it suffices to find some $r^{*}$ such that $A_{J^{c}} u \leq b_{J^{c}}-A_{J^{c}} \theta_{0}$ for any $u \in B_{r^{*}}(0)$. For each $j \notin J$, we have $\left\langle a_{j}, \theta_{0}\right\rangle<b_{j}$ so there exists some $r^{*}>0$ such that all $\theta \in B_{r^{*}}\left(\theta_{0}\right)$ satisfy $\left\langle a_{j}, \theta\right\rangle<b_{j}$ for all $j \notin J$. Taking $u=\theta-\theta_{0}$ concludes the proof.

Proof of Lemma 3.3. Let $\mathcal{T}:=\left\{u+\Pi_{\mathcal{C}}\left(\theta^{*}\right): u \in T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)\right\}$. Using the locally polyhedral condition (15) and continuity [5] of $\Pi_{\mathcal{C}}$ and $\Pi_{\mathcal{T}}$, we have $\Pi_{\mathcal{T}}\left(\theta^{*}\right)=\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ (e.g., see the verification of (25)), and thus translating yields $\Pi_{T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)}\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)=0$. Applying Lemma A. 1 with $\mathcal{K}=T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right), y=\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)$, and $\Pi_{\mathcal{K}}(y)=0$ concludes the proof.

Proof of Lemma 3.4. Fix $\theta_{0} \in \mathcal{C}$. For any $\theta^{*} \notin \mathcal{C}$ such that $\Pi_{\mathcal{C}}\left(\theta^{*}\right)=\theta_{0}$, Lemma 3.2 implies the locally polyhedral condition (15) holds, and thus Lemma 3.3 establishes that $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap$ $\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$ is a face of the tangent cone $T_{\mathcal{C}}\left(\theta_{0}\right)$.

Since the tangent cone has finitely many faces, the supremum is actually a maximum over the statistical dimensions of finitely many such lower-dimensional faces. Thus it remains to show

$$
\delta\left(T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}\right)<\delta\left(T_{\mathcal{C}}\left(\theta_{0}\right)\right)
$$

for each $\theta^{*} \notin \mathcal{C}$ such that $\Pi_{\mathcal{C}}\left(\theta^{*}\right)=\theta_{0}$.
The set $\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$ is a hyperplane (not all of $\mathbb{R}^{n}$ ) because $\theta^{*} \notin \mathcal{C}$. Using the fact that the tangent cone $T_{\mathcal{C}}\left(\theta_{0}\right)$ has nonempty interior (because it contains the translation $F_{\mathcal{C}}\left(\theta_{0}\right)$ of $\mathcal{C}$ ), we see that the intersection $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$ is a face that that lies in a strictly lower-dimensional subspace of $\mathbb{R}^{n}$, and is therefore strictly smaller than the full cone $T_{\mathcal{C}}\left(\theta_{0}\right)$. Thus, we just need to show $\delta\left(T^{\prime}\right)<\delta(T)$ for any polyhedral cone $T$ with nonempty interior in $\mathbb{R}^{n}$, and any face $T^{\prime}$ of $T$ that lies in a strictly lower-dimensional subspace of $\mathbb{R}^{n}$.

For a point $x \in \mathbb{R}^{n}$ and a set $S \subseteq \mathbb{R}^{n}$ let $d(x, S):=\inf _{\theta \in S}\|x-\theta\|$. Note that the Moreau decomposition for cones [2], Sec. B, implies $\left\|\Pi_{\mathcal{K}}(x)\right\|=d\left(x, \mathcal{K}^{\circ}\right)$ for any $x \in \mathbb{R}^{n}$ and any cone $\mathcal{K}$, where $\mathcal{K}^{\circ}:=\left\{u \in \mathbb{R}^{n}:\langle u, \theta\rangle \leq 0, \forall \theta \in \mathcal{K}\right\}$ denotes the polar cone of $\mathcal{K}$. Since $T^{\circ} \subseteq\left(T^{\prime}\right)^{\circ}$, we have

$$
d\left(x,\left(T^{\prime}\right)^{\circ}\right) \leq d\left(x, T^{\circ}\right), \quad \forall x \in \mathbb{R}^{n}
$$

Thus, if we show the random vector $Z$ has nonzero probability of being in the set

$$
\mathcal{A}:=\left\{x \in \mathbb{R}^{n}: d\left(x,\left(T^{\prime}\right)^{\circ}\right)<d\left(x, T^{\circ}\right)\right\}=\left\{x \in \mathbb{R}^{n}:\left\|\Pi_{T^{\prime}}(x)\right\|<\left\|\Pi_{T}(x)\right\|\right\}
$$

then we immediately have the desired strict inequality

$$
\delta\left(T^{\prime}\right)=\mathbb{E} d\left(Z,\left(T^{\prime}\right)^{\circ}\right)<\mathbb{E} d\left(Z, T^{\circ}\right)=\delta(T)
$$

To prove the above claim that $\mathbb{P}(Z \in \mathcal{A})>0$, we show below that the interior of $T$ is contained in $\mathcal{A}$; then our assumption on $Z$ will conclude the proof.

Let $x$ be in the interior of $T$. Then $x \in T \backslash T^{\prime}$. Moreover, if we let $U$ be the smallest linear subspace of $\mathbb{R}^{n}$ containing $T^{\prime}$, then $x \notin U$ as well. Note the the Pythagorean theorem implies

$$
\begin{equation*}
\left\|\Pi_{T}(x)\right\|^{2}=\|x\|^{2}=\left\|\Pi_{U}(x)\right\|^{2}+\left\|x-\Pi_{U}(x)\right\|^{2}>\left\|\Pi_{U}(x)\right\|^{2} \tag{41}
\end{equation*}
$$

We also have

$$
\Pi_{T^{\prime}}(x)=\underset{\theta \in T^{\prime}}{\operatorname{argmin}}\|\theta-x\|^{2}=\underset{\theta \in T^{\prime}}{\operatorname{argmin}}\left\{\left\|\theta-\Pi_{U}(x)\right\|^{2}+\left\|\Pi_{U}(x)-x\right\|^{2}\right\}=\Pi_{T^{\prime}}\left(\Pi_{U}(x)\right),
$$

so combining this with the above inequality (41) and the optimality condition (14) for the projection of $\Pi_{U}(x)$ onto the cone $T^{\prime}$, we have

$$
\begin{aligned}
\left\|\Pi_{T^{\prime}}(x)\right\|^{2} & =\left\|\Pi_{T^{\prime}}\left(\Pi_{U}(x)\right)\right\|^{2} \\
& =\left\|\Pi_{U}(x)\right\|^{2}-\left\|\Pi_{U}(x)-\Pi_{T^{\prime}}\left(\Pi_{U}(x)\right)\right\|^{2} \leq\left\|\Pi_{U}(x)\right\|^{2}<\left\|\Pi_{T}(x)\right\|^{2}
\end{aligned}
$$

and thus $x \in \mathcal{A}$.
Proof of Lemma 3.5. The lemma holds immediately if $\theta^{*} \in \mathcal{T}$, so we assume $\theta^{*} \notin \mathcal{T}$.
By translating, we may without loss of generality assume $\Pi_{\mathcal{T}}\left(\theta^{*}\right)=0$ so that the cone is centered at 0 and can be written as $\mathcal{T}=\{u: A u \leq 0\}$ for some number of constraints $m$ and some matrix $A \in \mathbb{R}^{m \times n}$. The objective then reduces to

$$
\Pi_{\mathcal{T}}(y) \in\left(\theta^{*}\right)^{\perp}, \quad \text { for all } y \in B_{r}\left(\theta^{*}\right)
$$

For any $y \in \mathbb{R}^{n}$ let $J_{y} \subseteq\{1, \ldots, m\}$ be as defined in Lemma A. 1 for our polyhedral cone $\mathcal{T}$; it characterizes the largest face of $\mathcal{T}$ that lies in $\left(\theta^{*}\right)^{\perp}$. We claim there exists $r>0$ such that

$$
\begin{equation*}
\left\{u: A_{J_{y}} u=0\right\} \subseteq\left(\theta^{*}\right)^{\perp}, \quad \forall y \in B_{r}\left(\theta^{*}\right) \tag{42}
\end{equation*}
$$

If not, then there exists a sequence of points $y_{k} \notin \mathcal{T}$ converging to $\theta^{*}$ such that $\left\{u: A_{J_{y_{k}}} u=0\right\} \nsubseteq$ $\left(\theta^{*}\right)^{\perp}$ for all $k$. Since there are finitely many distinct subsets $J_{y_{k}}$, we may take a subsequence and without loss of generality assume it is common subset $J=J_{y_{k}}$ for all $k$, and $\left\{u: A_{J} u=0\right\} \nsubseteq$ $\left(\theta^{*}\right)^{\perp}$. By the definition (39) of $J_{y_{k}}$, any $u$ satisfying $A_{J} u=0$ also satisfies $\left\langle y_{k}-\Pi_{\mathcal{T}}\left(y_{k}\right), u\right\rangle=$ 0 . By continuity of $\Pi_{\mathcal{T}}$ and taking $k \rightarrow \infty$, we have $\left\langle\theta^{*}, u\right\rangle=0$ as well, a contradiction.

Finally, since the optimality condition (11) for $\Pi_{\mathcal{T}}$ implies $\left\langle\Pi_{\mathcal{T}}(y), y-\Pi_{\mathcal{T}}(y)\right\rangle=0$ for any $y \in \mathbb{R}^{n}$, (40) implies $\Pi_{\mathcal{T}}(y) \in\left\{u: A_{J_{y}} u=0\right\}$. Combining this with (42) concludes the proof.

## Appendix B: Proofs for Section 4.2 (isotonic regression)

## B.1. Proofs of block monotone cone lemmas

Proof of Lemma 4.4. The first claim follows from decomposing the squared Euclidean distance into blocks.

$$
\begin{aligned}
\min _{v \in \mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}}\|v-z\|^{2} & =\min _{x \in \mathcal{S}^{m}} \sum_{j=1}^{m} \sum_{i \in I_{j}}\left(x_{j}-z_{i}\right)^{2}=\min _{x \in \mathcal{S}^{m}} \sum_{j=1}^{m} \sum_{i \in I_{j}}\left(\left(x_{j}-\bar{z}_{I_{j}}\right)^{2}+\left(\bar{z}_{I_{j}}-y_{i}\right)^{2}\right) \\
& =\sum_{j=1}^{m} \sum_{i \in I_{j}}\left(z_{i}-\bar{z}_{I_{j}}\right)^{2}+\min _{x \in \mathcal{S}^{m}} \sum_{j=1}^{m}\left|I_{j}\right|\left(x_{j}-\bar{z}_{I_{j}}\right)^{2} .
\end{aligned}
$$

Let $Z$ and $Z^{\prime}$ be standard Gaussian in $\mathbb{R}^{n}$ and $\mathbb{R}^{m}$, respectively. If $\left|I_{1}\right|=\cdots=\left|I_{m}\right|=r$, then the first claim implies

$$
\begin{aligned}
\delta\left(\mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}\right) & :=\mathbb{E}\left\|\Pi_{\mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}}(Z)\right\|^{2} \stackrel{(\mathrm{i})}{=} r \mathbb{E}\left\|\Pi_{\mathcal{S}^{m}}\left(Z^{\prime} / \sqrt{r}\right)\right\|^{2} \\
& \stackrel{(\mathrm{ii})}{=} \mathbb{E}\left\|\Pi_{\mathcal{S}^{m}}\left(Z^{\prime}\right)\right\|^{2}=: \delta\left(\mathcal{S}^{m}\right) \\
& =\sum_{j=1}^{m} \frac{1}{j},
\end{aligned}
$$

where (i) is due to $Z^{\prime} / \sqrt{r} \stackrel{d}{=}\left(\bar{Z}_{I_{1}}, \ldots, \bar{Z}_{I_{m}}\right)$, and (ii) is due to $\Pi_{\mathcal{C}}(c x)=c \Pi_{\mathcal{C}}(x)$ for a cone $\mathcal{C}$ and $c>0$ (e.g., [3], Section 1.6). The statistical dimension of $\mathcal{S}^{m}$ is proved by Amelunxen et al. [2], Sec. D.4.

Proof of Lemma 4.5. We use two useful properties of the statistical dimension of any cone $\mathcal{C}$ [2], Prop. 3.1.

- Rotational invariance: for any orthogonal transformation $Q$, we have $\delta(Q \mathcal{C})=\delta(\mathcal{C})$.
- Invariance under embedding: $\delta\left(\mathcal{C} \times\{0\}^{k}\right)=\delta(\mathcal{C})$.

Thus it suffices to provide an orthogonal transformation $Q$ such that $Q \mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}$ is an embedding of the cone (31) into $\mathbb{R}^{n}$.

Let $e_{i}$ denote the $i$ th standard basis vector in $\mathbb{R}^{n}$. Let the last element of each block be denoted $k_{j}:=\max I_{j}$ for $1 \leq j \leq m$, with $k_{0}=0$ for convenience. The block monotone cone $\mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}$ is defined by the following constraints for $u \in \mathbb{R}^{n}$.

$$
\begin{align*}
\left\langle e_{i}-e_{i+1}, u\right\rangle \leq 0, & i \in\left\{k_{1}, \ldots, k_{m}\right\}  \tag{43a}\\
\left\langle e_{i}-e_{i+1}, u\right\rangle=0, & i \in\{1, \ldots, n-1\} \backslash\left\{k_{1}, \ldots, k_{m}\right\} . \tag{43b}
\end{align*}
$$

Let us focus on an arbitrary block $I_{j}$. Consider the $\left|I_{j}\right| \times\left|I_{j}\right|$ matrix

$$
\tilde{A}_{j}=\left[\begin{array}{ccccc}
1 & & & & \\
-1 & 1 & & & \\
& -1 & 1 & & \\
& & \ddots & \ddots & \\
& & & -1 & 1
\end{array}\right]
$$

Because $\tilde{A}_{j}$ is full rank, the ${\underset{Q R}{\sim}}^{\text {QR }}$ decomposition implies there exists an $\left|I_{j}\right| \times\left|I_{j}\right|$ orthogonal matrix $\tilde{Q}_{j}$ such that $\tilde{R}_{j}:=\tilde{Q}_{j} \tilde{A}_{j}$ is upper triangular with positive diagonal entries, and this decomposition is unique.

The block diagonal matrix $Q$ with blocks $\tilde{Q}_{1}, \ldots, \tilde{Q}_{m}$ is an $n \times n$ orthogonal matrix. Let $A$ and $R$ also be block diagonal, each constructed similarly using the $\tilde{A}_{j}$ and the $\tilde{R}_{j}$ respectively, so that $U=Q A$. We consider $Q S_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}$. We use the fact that if $v=Q u$ then $\langle b, u\rangle \leq 0 \Longleftrightarrow$ $\langle Q b, v\rangle \leq 0$ to rewrite the constraints (43a) and (43b). The following hold for each $j=1, \ldots, m$.

- Note that the $i$ th column of $A$ is $a_{i}=e_{i}-e_{i+1}$ when $k_{j-1}<i<k_{j}$. For these $i$, the equality constraints (43b) after the transformation become $0=\left\langle Q\left(e_{i}-e_{i+1}\right), v\right\rangle=\left\langle r_{i}, v\right\rangle$ where $r_{i}$ is the $i$ th column of $R$. Since $\tilde{R}_{j}$ is upper triangular with nonzero diagonal entries (because $\tilde{A}_{j}$ is full rank), induction on $i=k_{j-1}+1, \ldots, k_{j}-1$ implies

$$
v_{i}=0, \quad k_{j-1}<i<k_{j} .
$$

- When $j<m$, we have $e_{k_{j}}=a_{k}$ and $e_{k_{j}+1}=a_{k_{j}}+a_{k_{j}+1}+\cdots+a_{k_{j+1}}$ Thus for $j<m$ the inequality constraint $\left\langle e_{k_{j}}-e_{k_{j}+1}, u\right\rangle \leq 0$ becomes

$$
0 \geq\left\langle Q\left(e_{k_{j}}-e_{k_{j}+1}\right), v\right\rangle=\left\langle r_{k_{j}}-r_{k_{j}+1}-r_{k_{j}+2}-\cdots-r_{k_{j+1}}, v\right\rangle=\left\langle r_{k_{j}}-r_{k_{j+1}}, v\right\rangle,
$$

where the last equality is due to $\left\langle r_{i}, v\right\rangle=0$ for $k_{j}<i<k_{j+1}$, by the previous point. Since $\tilde{R}_{j}$ and $\tilde{R}_{j+1}$ are each upper triangular, the inequality reduces to $r_{k_{j}, k_{j}} v_{k_{j}} \leq r_{k_{j+1}, k_{j+1}} v_{k_{j+1}}$, where $r_{k, k}$ denotes the $k$ th diagonal entry of $R$. Lemma B. 1 (proved below) computes these diagonal elements and yields

$$
\frac{v_{k_{j}}}{\sqrt{\left|I_{j}\right|}} \leq \frac{v_{k_{j+1}}}{\sqrt{\left|I_{j+1}\right|}}
$$

Therefore, we have shown that $Q \mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}$ consists of all vectors satisfying

$$
\frac{v_{k_{1}}}{\sqrt{\left|I_{1}\right|}} \leq \frac{v_{k_{2}}}{\sqrt{\left|I_{2}\right|}} \leq \cdots \leq \frac{v_{k_{m}}}{\sqrt{\left|I_{m}\right|}} \quad \text { and } \quad v_{i}=0, \quad \forall i \in\{1, \ldots, n\} \backslash\left\{k_{1}, \ldots, k_{m}\right\}
$$

We have thus verified the claim that $Q \mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}$ is an embedding of (31) into $\mathbb{R}^{n}$.
When the blocks all have equal size $r$, the cone (31) becomes the monotone cone $\mathcal{S}^{m}$, whose statistical dimension is $\sum_{j=1}^{m} \frac{1}{j}$ [2], Sec. D.4.

Lemma B.1. Consider the $n \times n$ matrix

$$
A=\left[\begin{array}{ccccc}
1 & & & & \\
-1 & 1 & & & \\
& -1 & 1 & & \\
& & \ddots & \ddots & \\
& & & -1 & 1
\end{array}\right]
$$

There exists a unique orthogonal matrix $Q$ and a unique upper triangular matrix $R$ with positive diagonal entries such that $A=Q R$. The bottom-right entry of $R$ is $r_{n, n}=1 / \sqrt{n}$.

Proof. Let $q_{i}$ be the $i$ th column of $Q$. The last column $q_{n}$ is orthogonal to the span of the first $n-1$ columns of $A$, so $q_{n}$ is either $(1, \ldots, 1) / \sqrt{n}$ or its negative. The positivity constraint on the diagonal entries of $R$ implies the former, and thus $r_{n, n}=\left\langle q_{n}, e_{n}\right\rangle=1 / \sqrt{n}$.

## B.2. Statistical dimension of the block monotone cone in general

In Lemma 4.5, we provided an expression for the statistical dimension of the block monotone cone $\mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}$ when the block sizes were equal. In general, the statistical dimension can be higher or lower than $\sum_{j=1}^{m} \frac{1}{j}$. Consider the following examples for $m=3$.

Lemma 4.5 implies $\mathcal{S}_{n-2,1,1}$ has the same statistical dimension as $\left\{v \in \mathbb{R}^{3}: v_{1} / \sqrt{n-2} \leq v_{2} \leq\right.$ $\left.v_{3}\right\}$. As $n \rightarrow \infty$ this latter cone approaches $\left\{v \in \mathbb{R}^{3}: 0 \leq v_{2} \leq v_{3}\right\}$ which has statistical dimension $1+\left(\frac{1}{8} \cdot 2+\frac{1}{2} \cdot 1\right)=\frac{7}{4}=1.75$, which is smaller than $\sum_{j=1}^{3} \frac{1}{j}=\frac{11}{6}=1.8 \overline{3}$.

On the other hand, $\mathcal{S}_{1, n-2,1}$ has the same statistical dimension as $\left\{v \in \mathbb{R}^{3}: v_{1} \leq v_{2} / \sqrt{n-2} \leq\right.$ $\left.v_{3}\right\}$. As $n \rightarrow \infty$ this latter cone approaches $\left\{v \in \mathbb{R}^{3}: v_{1} \leq 0, v_{3} \geq 0\right\}$ which has statistical dimension $1+\frac{1}{2}+\frac{1}{2}=2$, which is larger than $1.8 \overline{3}$.

We suspect that the approach used to prove the statistical dimension of $\mathcal{S}^{n}$ [2], Sec. D.4, which uses the theory of finite reflection groups, cannot be generalized for $\mathcal{S}_{\left|I_{1}\right|, \ldots,\left|I_{m}\right|}$, due to the asymmetry of (31). However, using a result of [7], it is possible to show that the average statistical dimension among all block monotone cones with a given [unordered] set of $m$ block sizes is $H_{m}$ [1], Prop. 6.6.

## B.3. Proof of Proposition 4.3

When applying Theorem 3.1, it is useful to characterize $\mathcal{S}^{n}$ and its tangent cones using conic generators. If $T \subseteq \mathbb{R}^{n}$ is a cone and there exist $x_{1}, \ldots, x_{p} \in T$ such that

$$
T=\left\{\sum_{i=1}^{p} \alpha_{i} x_{i}: \alpha_{i} \geq 0, \forall i\right\}
$$

then we call $x_{1}, \ldots, x_{p}$ the conic generators of $T$, and write

$$
T=\operatorname{cone}\left\{x_{1}, \ldots, x_{p}\right\}
$$

Lemma B.2. Let $\theta^{*} \in \mathbb{R}^{n}$ and let $\mathcal{C} \subseteq \mathbb{R}^{n}$ be closed and convex. If the tangent cone $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$ is generated by $x_{1}, \ldots, x_{p} \in \mathbb{R}^{n}$, that is, $T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)=\operatorname{cone}\left\{x_{1}, \ldots, x_{p}\right\}$, then

$$
T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}=\operatorname{cone}\left(\left\{x_{1}, \ldots, x_{p}\right\} \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}\right)
$$

Proof of Lemma B.2. The inclusion $\supset$ is immediate, so it remains to prove the inclusion $\subseteq$. Note that the optimality condition (11) implies $\left\langle\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right), x\right\rangle \leq 0$ for any $x \in T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$. In particular, if $v \in T_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$, then $v$ can be written as the conical combination $v=\sum_{i=1}^{p} \alpha_{i} x_{i}$ with $\alpha_{i} \geq 0$, and we have

$$
0=\left\langle\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right), v\right\rangle=\sum_{i=1}^{p} \alpha_{i} \underbrace{\left\langle\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right), x_{i}\right\rangle}_{\leq 0} .
$$

Thus, if a generator $x_{i}$ is not in the hyperplane $\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$, then $\alpha_{i}=0$, so $x_{i}$ does not contribute in the conical combination of $v$. Thus, $v$ can be written as a conical combination of generators in $\left(\theta^{*}-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)^{\perp}$.

We are now ready to prove Proposition 4.3.
Proof of Proposition 4.3. By Theorem 3.1, it suffices to prove that the statistical dimension term is $\sum_{k=1}^{K} \delta\left(\mathcal{S}_{\left|I_{1}^{k}\right|, \ldots, \mid I_{m_{k}}^{k}} \mid\right.$.

For $p \geq 1$ let

$$
M_{p}:=\left[\begin{array}{cccc}
-1 & -1 & \cdots & -1 \\
1 & 1 & \cdots & 1 \\
& 1 & \cdots & 1 \\
& & \ddots & \vdots \\
& & & 1
\end{array}\right] \in \mathbb{R}^{(p+1) \times p}
$$

The rows of $M_{p}$ are the conic generators of $\mathcal{S}^{p}$.
Suppose first that $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)$ is constant, so that $K=1$ and $J_{1}=\{1, \ldots, n\}$. Then $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)=$ $\left(\mu_{1}, \mu_{1}, \ldots, \mu_{1}\right)$ where $\mu_{1}:=\frac{1}{n} \sum_{i=1}^{n} \theta_{i}^{*}$; this follows directly by minimizing $\sum_{i=1}^{n}\left(\theta_{i}^{*}-\mu_{1}\right)^{2}$ with respect to $\mu_{1}$.

The finest partition $\left(I_{1}^{1}, \ldots, I_{m_{1}}^{1}\right)$ of $J_{1}$ into blocks satisfying (30) can be constructed greedily as follows. Begin populating $I_{1}^{1}$ with the elements of $\{1, \ldots, n\}$ in order, stopping as soon as the mean of the elements of $I_{1}^{1}$ is $\mu_{1}$. Then begin populating $I_{2}^{1}$ with the remaining elements in order, again stopping when the mean of the elements in $I_{2}^{1}$ is $\mu_{1}$. Continue in this manner until the last element $n$ is placed in a subset $I_{m_{1}}^{1}$. The mean of the elements of this last subset $I_{m_{1}}^{1}$ is $\mu_{1}$ as well, since the mean of all components of $\theta^{*}$ is $\mu_{1}$. Thus this partition satisfies (30). To establish uniqueness, note that if some other partition of $J_{1}$ satisfies (30), then our partition ( $I_{1}^{1}, \ldots, I_{m_{1}}^{1}$ ) must be a refinement, due to the greedy construction.

Because $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)$ is constant, the tangent cone there is $T_{\mathcal{S}^{n}}\left(\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right)=\mathcal{S}^{n}$ [3], Prop. 3.1, which is generated by the rows of $M_{n}$. In order to use Lemma B.2, we need to determine which rows of $M_{n}$ are in the hyperplane $\left(\theta^{*}-\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right)^{\perp}$. We already know the mean of the components of $\theta^{*}-\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)$ is zero, so the first two rows are in the hyperplane.

We claim that exactly $m_{1}-1$ of the remaining $n-1$ rows of $M_{n}$ also lie in the hyperplane. Explicitly, if $\left(I_{1}^{1}, \ldots, I_{m_{1}}^{1}\right)$ is without loss of generality assumed to be sorted in increasing order, then the remaining rows of $M_{n}$ that lie in the hyperplane are the indicator vectors for

$$
\begin{equation*}
\bigcup_{j=u}^{m_{k}} I_{j}^{1}, \quad 2 \leq u \leq m_{k} \tag{44}
\end{equation*}
$$

No other rows of $M_{n}$ can be in the hyperplane, else there would exist a finer partition of $J_{1}$.
So, Lemma B. 2 implies $T_{\mathcal{S}^{n}}\left(\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right)^{\perp}$ is the cone generated by $(-1, \ldots,-1),(1, \ldots, 1)$, and the indicator vectors of the subsets (44), otherwise known as the cone of nondecreasing vectors that are piecewise constant on the blocks $I_{1}^{1}, \ldots, I_{m_{1}}^{1}$. Its statisti-
cal dimension is denoted by $\delta\left(\mathcal{S}_{\left|I_{1}^{1}\right|, \ldots,\left|I_{m_{1}}^{1}\right|}\right)$. This concludes the proof in the case when $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)$ is constant.

We now turn to the general case where $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)$ is piecewise constant with values $\mu_{1}<\cdots<$ $\mu_{K}$ on $J_{1}, \ldots, J_{K}$, respectively. We claim

$$
\begin{equation*}
\mu_{k}=\frac{1}{\left|J_{k}\right|} \sum_{i \in J_{k}} \theta_{i}^{*} . \tag{45}
\end{equation*}
$$

Since $\mathcal{S}^{n}$ is a cone, the projection satisfies $\left\langle\theta^{*}-\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right), x\right\rangle \leq 0$ for all $x \in \mathcal{S}^{n}$, with equality if $x=\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ (e.g., [3], Section 1.6). Letting $x_{1}, \ldots, x_{n+1}$ be the conic generators of $\mathcal{S}^{n}$ (the rows of $M_{n}$ ), we have $\Pi_{\mathcal{C}}\left(\theta^{*}\right)=\sum_{i=1}^{n+1} \alpha_{i} x_{i}$ for some coefficients $\alpha_{i} \geq 0$. Then,

$$
0=\left\langle\theta^{*}-\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right), \Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\rangle=\sum_{i=1}^{n+1} \alpha_{i} \underbrace{\left\langle\theta^{*}-\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right), x_{i}\right\rangle}_{\leq 0},
$$

which implies $\left\langle\theta^{*}-\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right), x_{i}\right\rangle=0$ if $\alpha_{i}>0$. Consequently, if $\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)$ changes value from component $j-1$ to $j$, then $\sum_{i=j}^{n}\left[\theta_{i}^{*}-\left(\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right)_{i}\right]=0$. Thus (45) holds.

By Proposition 3.1 of [3], the tangent cone is

$$
T_{\mathcal{S}^{n}}\left(\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right)=\mathcal{S}^{n_{1}} \times \cdots \times \mathcal{S}^{n_{K}}
$$

which is generated by the rows of the block diagonal matrix

$$
A:=\left[\begin{array}{lll}
M_{n_{1}} & & \\
& \ddots & \\
& & M_{n_{K}}
\end{array}\right]
$$

To find which rows of $A$ are in the hyperplane $\left(\theta^{*}-\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right)^{\perp}$, we can treat each block $M_{n_{k}}$ separately and repeat the above argument. Doing so shows that $T_{\mathcal{S}^{n}}\left(\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{S}^{n}}\left(\theta^{*}\right)\right)^{\perp}$ is the cone of vectors that are piecewise constant on $\left(I_{1}^{1}, \ldots, I_{m_{1}}^{1}, \ldots, I_{1}^{K}, \ldots, I_{m_{K}}^{K}\right)$ and are increasing within each of the blocks $\left(J_{1}, \ldots, J_{K}\right)$. The statistical dimension of this cone is $\sum_{k=1}^{K} \delta\left(\mathcal{S}_{\left|I_{1}^{k}\right|, \ldots,\left|I_{m_{k}}^{k}\right|}\right)$.

## Appendix C: Proofs for Section 4.3 (convex regression)

Recall we fix a design $x_{1}<\cdots<x_{n}$ and use the shorthand $x:=\left(x_{1}, \ldots, x_{n}\right)$.
We collect a few technical lemmas before proving Proposition 4.7.
Lemma C. 1 (Conic generators for $\mathcal{K}_{x}$ ). A set of conic generators for $\mathcal{K}_{x}$ are $\pm(1, \ldots, 1), \pm x$, and $g^{(2)}, \ldots, g^{(n-1)}$ where $g_{i}^{(j)}:=\max \left\{x_{i}-x_{j}, 0\right\}$.

Proof. If $\theta \in \mathcal{K}_{x}$, then

$$
\theta=\left(\theta_{2}-\frac{\theta_{2}-\theta_{1}}{x_{2}-x_{1}} x_{2}\right)(1, \ldots, 1)+\frac{\theta_{2}-\theta_{1}}{x_{2}-x_{1}} x+\sum_{j=2}^{n-1}\left(\frac{\theta_{i+1}-\theta_{i}}{x_{i+1}-x_{i}}-\frac{\theta_{i}-\theta_{i-1}}{x_{i}-x_{i-1}}\right) g^{(j)}
$$

Note that the coefficients for the $g^{(j)}$ are nonnegative by the definition of $\mathcal{K}_{x}$.
Lemma C.2. If $\theta_{0} \in \operatorname{span}\{(1, \ldots, 1), x\}$, then $T_{\mathcal{K}_{x}}\left(\Pi_{\mathcal{K}_{x}}\left(\theta_{0}\right)\right)=\mathcal{K}_{x}$.
Proof. If $\theta \in \mathcal{K}_{x}$, then for any $\alpha>0$, we have $\alpha\left(\theta-\theta_{0}\right) \in \mathcal{K}_{x}$, since subtracting an affine function from a convex function yields a convex function, and scaling by a positive constant does not affect convexity. Since $\mathcal{K}_{x}$ is closed, we have $T_{\mathcal{K}_{x}}\left(\Pi_{\mathcal{K}_{x}}\left(\theta_{0}\right)\right) \subseteq \mathcal{K}_{x}$.

Conversely, if $v \in \mathcal{K}_{x}$, then $\theta_{0}+v \in \mathcal{K}_{x}$ since adding an affine function to a convex function yields a convex function. Thus, $v \in T_{\mathcal{K}_{x}}\left(\Pi_{\mathcal{K}_{x}}\left(\theta_{0}\right)\right)$.

Lemma C.3. Suppose $\theta^{*} \in-\mathcal{K}_{x}$ is concave. Then $\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)$ is the projection of $\theta^{*}$ on $\operatorname{span}\{(1, \ldots, 1), x\}$, that is, it is obtained by a simple linear regression of $\theta^{*}$ on $x$.

Proof. By definition, there exists a concave function $f:\left[x_{1}, x_{n}\right] \rightarrow \mathbb{R}$ such that $\theta_{i}^{*}=f\left(x_{i}\right)$ for all $i$. For any $\theta \in \mathcal{K}_{x}$, there exists a convex function $g:\left[x_{1}, x_{n}\right] \rightarrow \mathbb{R}$ such that $\theta_{i}=g\left(x_{i}\right)$ for all $i$.

If $g$ is not affine on $\left[x_{1}, x_{n}\right]$, then it intersects $f$ in exactly 0,1 , or 2 points. In any of these cases, there exists an affine function $h$ such that $\min \{f, g\} \leq h \leq \max \{f, g\}$ on $\left[x_{1}, x_{2}\right]$. (One can note that such an $h$ must touch all intersection points of $f$ and $g$, if there are any.)

If we let $\theta^{\prime}$ be defined by $\theta_{i}^{\prime}:=h\left(x_{i}\right)$, then by construction we have $\left\|\theta^{\prime}-\theta^{*}\right\|^{2}<\left\|\theta-\theta^{*}\right\|^{2}$. (The strict inequality appears since equality would only occur if $f$ and $g$ are both affine, but we assumed $g$ is not affine.) Therefore, $\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)$ must be linear, and the result of the lemma follows immediately.

We are now ready to prove the proposition.
Proof of Proposition 4.7. By Lemma C.3, $\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)$ is obtained by a linear regression of $\theta^{*}$ on $x$. Thus, $T_{\mathcal{K}_{x}}\left(\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)\right)=\mathcal{K}_{x}$ by Lemma C.2.

By Lemma B.2, the intersection $T_{\mathcal{K}_{x}}\left(\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)\right)^{\perp}=\mathcal{K}_{x} \cap\left(\theta^{*}-\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)\right)^{\perp}$ is the cone generated by the conic generators of $\mathcal{K}_{x}$ that are orthogonal to $\theta^{*}-\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)$.

We already know $\pm(1, \ldots, 1)$ and $\pm x$ are orthogonal to $\theta^{*}-\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)$, by Lemma C.3. We claim each $g^{(j)}$ for $j=2, \ldots, n-2$ is not orthogonal to $\theta^{*}-\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)$. Otherwise, we would have $\left\langle g^{(j)}-\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right), \theta^{*}-\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)\right\rangle=0$ and then

$$
\left\|g^{(j)}-\theta^{*}\right\|^{2}=\left\|g^{(j)}-\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)\right\|^{2}+\left\|\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2}<\left\|\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2}
$$

where $g^{(j)} \neq \Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)$ by Lemma C.3. This contradicts the optimality of $\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)$.
As mentioned above, we therefore have $T_{\mathcal{K}_{x}}\left(\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)\right) \cap\left(\theta^{*}-\Pi_{\mathcal{K}_{x}}\left(\theta^{*}\right)\right)^{\perp}=\operatorname{span}\{(1, \ldots, 1)$, $x\}$, which has [statistical] dimension 2. Applying Theorem 3.1 concludes the proof of the proposition.

## Appendix D: Proof of Proposition 5.1

Let $r:=\left\|\theta^{*}\right\|$. By rotating the problem, we may without loss of generality assume $\theta^{*}=$ $(r, 0, \ldots, 0)$.

Let $E:=\left\{Y \in B_{(r-1) / 2}\left(\theta^{*}\right)\right\}$. Then we have $E \subseteq\{Y \notin \mathcal{C}\}$, so under the event $E$ we have $\hat{\theta}(Y)=Y /\|Y\|$. Noting $\|Y\|^{2}=\left\|\theta^{*}+\sigma Z\right\|^{2}=r^{2}+2 \sigma r Z_{1}+\sigma^{2}\|Z\|^{2}$, we have

$$
\begin{aligned}
\frac{1}{\sigma^{2}}\left\|\hat{\theta}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2}= & \frac{1}{\sigma^{2}}\left(\frac{r+\sigma Z_{1}}{\sqrt{r^{2}+2 \sigma r Z_{1}+\sigma^{2}\|Z\|^{2}}}-1\right)^{2} \\
& +\frac{\sum_{i=2}^{n} Z_{i}^{2}}{r^{2}+2 \sigma r Z_{1}+\sigma^{2}\|Z\|^{2}}
\end{aligned}
$$

The second term converges to $r^{-2} \sum_{i=2}^{n} Z_{i}^{2}$ as $\sigma \downarrow 0$. We show the first term vanishes as $\sigma \downarrow 0$. Defining $g(\sigma):=\left\|\theta^{*}+\sigma Z\right\|$, we have

$$
\begin{aligned}
g(\sigma) & =\sqrt{r^{2}+2 \sigma r Z_{1}+\sigma^{2}\|Z\|^{2}}, \\
g^{\prime}(\sigma) & =\frac{r Z_{1}+\sigma\|Z\|^{2}}{g(\sigma)}, \\
g^{\prime \prime}(\sigma) & =\frac{\|Z\|^{2}}{g(\sigma)}-\frac{\left(r Z_{1}+\sigma\|Z\|^{2}\right) g^{\prime}(\sigma)}{g(\sigma)^{2}} .
\end{aligned}
$$

Moreover, we have $g(0)=r, g^{\prime}(0)=Z_{1}$, and $g^{\prime \prime}(0)=\left(\|Z\|^{2}-Z_{1}^{2}\right) / r$. Then by L'Hôpital's rule,

$$
\begin{aligned}
\lim _{\sigma \downarrow 0} \frac{1}{\sigma}\left(\frac{r+\sigma Z_{1}}{\sqrt{r^{2}+2 \sigma r Z_{1}+\sigma^{2}\|Z\|^{2}}}-1\right) & =\lim _{\sigma \downarrow 0} \frac{r+\sigma Z_{1}-g(\sigma)}{\sigma g(\sigma)} \\
& =\lim _{\sigma \downarrow 0} \frac{Z_{1}-g^{\prime}(\sigma)}{g(\sigma)+\sigma g^{\prime}(\sigma)}=\frac{Z_{1}-Z_{1}}{r+0}=0 .
\end{aligned}
$$

Note $\mathbf{1}_{E} \rightarrow 1$ almost surely as $\sigma \downarrow 0$. Thus, $\sigma^{-2}\left\|\hat{\theta}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \mathbf{1}_{E} \rightarrow r^{-2} \sum_{i=2}^{n} Z_{i}^{2}$ almost surely. By the upper bound (7), we may use the dominated convergence theorem to get

$$
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left\|\hat{\theta}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \mathbf{1}_{E}\right]=\frac{1}{r^{2}} \sum_{i=2}^{n} \mathbb{E} Z_{i}^{2}=\frac{n-1}{r^{2}}
$$

To conclude the proof of the first limit (36a), note that

$$
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left\|\hat{\theta}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \mathbf{1}_{E^{c}}\right]=0
$$

which holds by the argument used in the proof of Theorem 3.1 (e.g., see the second term in (20)).

A similar proof holds for the second limit (36b). Let $E$ and $g(\sigma)$ be the same as before. Then

$$
\begin{aligned}
\frac{1}{\sigma^{2}} & \left(\left\|\hat{\theta}(Y)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{C}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2}\right) \\
& =\frac{1}{\sigma^{2}}\left(\frac{r+\sigma Z_{1}}{\sqrt{r^{2}+2 \sigma r Z_{1}+\sigma^{2}\|Z\|^{2}}}-r\right)^{2}+\frac{\sum_{i=2}^{n} Z_{i}^{2}}{r^{2}+2 \sigma r Z_{1}+\sigma^{2}\|Z\|^{2}}-\frac{(r-1)^{2}}{\sigma^{2}} \\
& =\frac{1}{\sigma^{2}}\left[\left(\frac{r+\sigma Z_{1}}{g(\sigma)}-r\right)^{2}-(r-1)^{2}\right]+\frac{\sum_{i=2}^{n} Z_{i}^{2}}{r^{2}+2 \sigma r Z_{1}+\sigma^{2}\|Z\|^{2}}
\end{aligned}
$$

Again, the second term tends to $r^{-2} \sum_{i=2}^{n} Z_{i}^{2}$ as $\sigma \downarrow 0$. To handle the first term, we use L'Hôpital's rule again. Let

$$
\begin{aligned}
h(\sigma) & :=\frac{r+\sigma Z_{1}}{g(\sigma)}-r \\
h^{\prime}(\sigma) & =\frac{Z_{1}}{g(\sigma)}-\frac{\left(r+\sigma Z_{1}\right) g^{\prime}(\sigma)}{g(\sigma)^{2}}, \\
h^{\prime \prime}(\sigma) & =-\frac{Z_{1} g^{\prime}(\sigma)}{g(\sigma)^{2}}+2 \frac{\left(r+\sigma Z_{1}\right) g^{\prime}(\sigma)^{2}}{g(\sigma)^{3}}-\frac{Z_{1} g^{\prime}(\sigma)+\left(r+\sigma Z_{1}\right) g^{\prime \prime}(\sigma)}{g(\sigma)^{2}} .
\end{aligned}
$$

Recalling the limits $g(0)=r, g^{\prime}(0)=Z_{1}$, and $g^{\prime \prime}(0)=\left(\|Z\|^{2}-Z_{1}^{2}\right) / r$, we have $h(\sigma) \rightarrow-(r-$ $1), h^{\prime}(\sigma) \rightarrow 0$, and

$$
h^{\prime \prime}(0)=-\frac{Z_{1}^{2}}{r^{2}}+2 \frac{r Z_{1}^{2}}{r^{3}}-\frac{Z_{1}^{2}+\|Z\|^{2}-Z_{1}^{2}}{r^{2}}=\frac{Z_{1}^{2}-\|Z\|^{2}}{r^{2}}
$$

Then, L'Hôpital's rule allows us to compute the limit of the first term.

$$
\begin{aligned}
\lim _{\sigma \downarrow 0} \frac{1}{\sigma^{2}}\left[\left(\frac{r+\sigma Z_{1}}{g(\sigma)}-r\right)^{2}-(r-1)^{2}\right] & =\lim _{\sigma \downarrow 0} \frac{h(\sigma)^{2}-(r-1)^{2}}{\sigma^{2}} \\
& =\lim _{\sigma \downarrow 0} \frac{h(\sigma) h^{\prime}(\sigma)}{\sigma} \\
& =\lim _{\sigma \downarrow 0}\left(h^{\prime}(\sigma)^{2}+h(\sigma) h^{\prime \prime}(\sigma)\right) \\
& =\frac{(r-1)\left(\|Z\|^{2}-Z_{1}^{2}\right)}{r^{2}} .
\end{aligned}
$$

Combining terms yields

$$
\frac{1}{\sigma^{2}}\left(\left\|\hat{\theta}(Y)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{C}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2}\right) \mathbf{1}_{E} \rightarrow \frac{(r-1)\left(\|Z\|^{2}-Z_{1}^{2}\right)+\sum_{i=2}^{n} Z_{i}^{2}}{r^{2}}=\frac{\sum_{i=2}^{n} Z_{i}^{2}}{r}
$$

so again by dominated convergence with the upper bound (7), we have

$$
\frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left(\left\|\hat{\theta}(Y)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{C}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2}\right) \mathbf{1}_{E}\right] \rightarrow \frac{n-1}{r}
$$

To conclude the proof of (36b), note that

$$
\frac{1}{\sigma^{2}} \mathbb{E}_{\theta^{*}}\left[\left(\left\|\hat{\theta}(Y)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{C}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2}\right) \mathbf{1}_{E^{c}}\right] \rightarrow 0
$$

which was proved in the proof of Theorem 3.1 (see (26)).

## Appendix E: Proofs for Section 5.2

The following lemma shows that the left-hand side of (38) is nonnegative.
Lemma E.1. For any $\theta_{0} \in \mathcal{C}$,

$$
\left\|\Pi_{F_{\mathcal{C}}\left(\theta_{0}\right)}(x)\right\|^{2} \geq\left\|\Pi_{K_{\mathcal{C}}}(x)\right\|^{2}
$$

Proof of Lemma E.1. Because $K_{\mathcal{C}}$ is a cone, we have $\left\langle x, \Pi_{K_{\mathcal{C}}}(x)\right\rangle=\left\|\Pi_{K_{\mathcal{C}}}(x)\right\|^{2}$. Since $K_{\mathcal{C}} \subseteq$ $F_{\mathcal{C}}\left(\theta_{0}\right)$, the optimality condition for $\Pi_{F_{\mathcal{C}}\left(\theta_{0}\right)}(x)$ implies $\left\langle x-\Pi_{F_{\mathcal{C}}\left(\theta_{0}\right)}(x), \Pi_{K_{\mathcal{C}}}(x)\right\rangle \leq 0$ and thus

$$
\left\|\Pi_{K_{\mathcal{C}}}(x)\right\|^{2} \leq\left\langle\Pi_{F_{\mathcal{C}}\left(\theta_{0}\right)}(x), \Pi_{K_{\mathcal{C}}}(x)\right\rangle \leq\left\|\Pi_{F_{\mathcal{C}}\left(\theta_{0}\right)}(x)\right\|\left\|\Pi_{K_{\mathcal{C}}}(x)\right\|
$$

Thus, $\left\|\Pi_{F_{\mathcal{C}}\left(\theta_{0}\right)}(x)\right\| \geq\left\|\Pi_{K_{\mathcal{C}}}(x)\right\|$ and $M_{\theta_{0}} \geq 0$.
Proof of Lemma 5.2. We first prove the equalities (i) and (ii).
(i) Let $v \in\left\{u: \mathbb{R}_{+} u \subseteq F_{\mathcal{C}}\left(\theta_{0}\right)\right\}$ and let $\theta \in \mathcal{C}$. For any $c>0$, we have $\theta_{0}+c v \in \mathcal{C}$, and convexity implies $\theta+\alpha\left(\theta_{0}+c v-\theta\right) \in \mathcal{C}$ for all $\alpha \in[0,1]$. For large $c$, we have $\| \theta_{0}+$ $c v-\theta \|>1$ and thus $\theta+\frac{\theta_{0}+c v-\theta}{\left\|\theta_{0}+c v-\theta\right\|} \in \mathcal{C}$. Taking $c \rightarrow \infty$ and using the fact that $\mathcal{C}$ is closed yields $\theta+\frac{v}{\|v\|} \in \mathcal{C}$ and thus $v \in T_{\mathcal{C}}(\theta)$. Since $\theta$ was arbitrary, we have $v \in K_{\mathcal{C}}$.

Conversely, suppose $v \in K_{\mathcal{C}}$. Let $c^{*}:=\sup \left\{c>0: \theta_{0}+c v \in \mathcal{C}\right\}$. The supremum is over a nonempty set because $v \in T_{\mathcal{C}}\left(\theta_{0}\right)$. Suppose for sake of contradiction that $c^{*}<\infty$. Since $\mathcal{C}$ is closed, $\theta_{0}+c^{*} v \in \mathcal{C}$. Thus $v \in T_{\mathcal{C}}\left(\theta_{0}+c^{*} v\right)$ which implies $\theta_{0}+\left(c^{*}+\alpha\right) v \in \mathcal{C}$ for some $\alpha>0$, contradicting the definition of $c^{*}$. Thus $c^{*}=\infty$ and $\theta_{0}+c v \in \mathcal{C}$ for all $c>0$.
(ii) Both sides can be expressed as the set of $v \in \mathbb{R}^{n}$ satisfying $\theta_{0}+\sigma v \in \mathcal{C}$ for all $\sigma>0$.

We now prove the second part of the lemma. The definition (37) implies $K_{\mathcal{C}} \subseteq T_{\mathcal{C}}(\theta)$ for any $\theta \in \mathcal{C}$.

Now, assume $F_{\mathcal{C}}\left(\theta_{0}\right)$ is a cone. If the reverse inclusion $T_{\mathcal{C}}(\theta) \subseteq F_{\mathcal{C}}(\theta)$ holds, then $\theta_{0}-\theta \in$ $T_{\mathcal{C}}(\theta)=F_{\mathcal{C}}(\theta)$ so $\theta_{0}-\left(\theta-\theta_{0}\right) \in \mathcal{C}$. Conversely, suppose $\theta_{0}-\left(\theta-\theta_{0}\right) \in \mathcal{C}$. If $v \in T_{\mathcal{C}}(\theta)$, then $\theta+c v \in \mathcal{C}$ for some $c>0$. By convexity, $\theta_{0}+c v / 2 \in \mathcal{C}$, so $v \in F_{\mathcal{C}}\left(\theta_{0}\right)$. Thus, $T_{\mathcal{C}}(\theta) \subseteq F_{\mathcal{C}}(\theta)$.

Proof of Proposition 5.3. We use $Y$ instead of $\theta^{*}+\sigma Z$ throughout the proof, but note that $Y$ depends on $\sigma$.

Without loss of generality, we can translate the problem so that $\Pi_{\mathcal{C}}\left(\theta^{*}\right)=0$.
In view of (13), we may use the dominated convergence theorem on $\sigma^{-2}\left\|\Pi_{\mathcal{C}}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2}$, so

$$
\begin{aligned}
& \lim _{\sigma \rightarrow \infty} \frac{1}{\sigma^{2}} \mathbb{E}\left\|\Pi_{\mathcal{C}}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \\
& \quad=\mathbb{E} \lim _{\sigma \rightarrow \infty} \frac{1}{\sigma^{2}}\left\|\Pi_{\mathcal{C}}(Y)-\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right\|^{2} \quad \text { dom. conv. with } \mathbb{E}\|Z\|^{2} \\
& \quad=\mathbb{E} \lim _{\sigma \rightarrow \infty} \frac{1}{\sigma^{2}}\left\|\Pi_{\mathcal{C}}(Y)\right\|^{2} \\
& \stackrel{(\mathrm{i})}{=} \mathbb{E}\left\|\Pi_{K_{\mathcal{C}}}(Z)\right\|^{2}=\delta\left(K_{\mathcal{C}}\right)
\end{aligned}
$$

where we verify the equality (i) below.
Similarly, (13) allows us to use the dominated convergence theorem again for the excess risk.

$$
\begin{aligned}
& \lim _{\sigma \rightarrow \infty} \frac{1}{\sigma^{2}}\left(\mathbb{E}\left\|\Pi_{\mathcal{C}}(Y)-\theta^{*}\right\|^{2}-\left\|\Pi_{\mathcal{C}}\left(\theta^{*}\right)-\theta^{*}\right\|^{2}\right) \\
& \quad=\mathbb{E} \lim _{\sigma \rightarrow \infty} \frac{1}{\sigma^{2}}\left(\left\|\Pi_{\mathcal{C}}(Y)-\theta^{*}\right\|^{2}-\left\|\theta^{*}\right\|\right) \quad \text { dom. conv. with } \mathbb{E}\|Z\|^{2} \\
& \quad=\mathbb{E} \lim _{\sigma \rightarrow \infty} \frac{1}{\sigma^{2}}\left(\left\|\Pi_{\mathcal{C}}(Y)\right\|^{2}-2\left\langle\Pi_{\mathcal{C}}(Y), \theta^{*}\right\rangle\right) \\
& \stackrel{(\mathrm{i})}{=} \mathbb{E}\left\|\Pi_{K_{\mathcal{C}}}(Z)\right\|^{2}=\delta\left(K_{\mathcal{C}}\right)
\end{aligned}
$$

It remains to verify (i) and (ii).
(i)

$$
\begin{aligned}
& \left|\frac{1}{\sigma^{2}}\left\|\Pi_{\mathcal{C}}(Y)\right\|^{2}-\left\|\Pi_{K_{\mathcal{C}}}(Z)\right\|^{2}\right| \\
& \quad \leq \frac{1}{\sigma^{2}}\left|\left\|\Pi_{\mathcal{C}}(Y)\right\|^{2}-\left\|\Pi_{K_{\mathcal{C}}}(Y)\right\|^{2}\right|+\left|\frac{1}{\sigma^{2}}\left\|\Pi_{K_{\mathcal{C}}}(Y)\right\|^{2}-\left\|\Pi_{K_{\mathcal{C}}}(Z)\right\|^{2}\right| \\
& \quad \leq \frac{c}{\sigma^{2}}+\left|\left\|\Pi_{K_{\mathcal{C}}}\left(\theta^{*} / \sigma+Z\right)\right\|^{2}-\left\|\Pi_{K_{\mathcal{C}}}(Z)\right\|^{2}\right| \quad \text { Lemma E.1; } K_{\mathcal{C}} \text { is a cone } \\
& \xrightarrow{\sigma \rightarrow \infty} 0 . \quad x \mapsto\left\|\Pi_{K_{\mathcal{C}}}(x)\right\|^{2} \text { is continuous. }
\end{aligned}
$$

(ii) We already showed $\left\|\Pi_{\mathcal{C}}(Y)\right\|^{2} / \sigma^{2} \rightarrow\left\|\Pi_{K_{\mathcal{C}}}(Z)\right\|^{2}$, so it suffices to show the cross term vanishes. Indeed, we have $\left\|\Pi_{\mathcal{C}}(Y)\right\| / \sigma \rightarrow\left\|\Pi_{K_{\mathcal{C}}}(Z)\right\|$, so

$$
\frac{1}{\sigma^{2}}\left|\left\langle\Pi_{\mathcal{C}}(Y), \theta^{*}\right\rangle\right| \leq \frac{1}{\sigma^{2}}\left\|\Pi_{\mathcal{C}}(Y)\right\|\left\|\theta^{*}\right\| \xrightarrow{\sigma \rightarrow \infty} 0
$$

Proof of Corollary 5.4. We begin with the first claim. Since $\mathcal{C}=\mathbb{R}_{+}^{n}$ is a cone, we have $K_{\mathcal{C}}=$ $\mathbb{R}_{+}^{n}$. Provided we verify (38), the result follows from Proposition 5.3. Let $\theta:=\Pi_{\mathcal{C}}\left(\theta^{*}\right)$ and fix $x \in \mathbb{R}^{n}$. Then some casework yields

$$
\left\|\Pi_{F_{\mathcal{C}}(\theta)}(x)\right\|^{2}-\left\|\Pi_{K_{\mathcal{C}}}(x)\right\|^{2}=\sum_{i=1}^{n} \max \left\{x_{i},-\theta_{i}\right\}^{2}-\sum_{i=1}^{n} \max \left\{x_{i}, 0\right\}^{2} \leq \sum_{i=1}^{n} \theta_{i}^{2}=\|\theta\|^{2}=: c
$$

We now turn to the second claim. If $\mathcal{C}$ is bounded, then by Lemma 5.2, $K_{\mathcal{C}}=\left\{u: \mathbb{R}_{+} u \subseteq\right.$ $\left.F_{\mathcal{C}}\left(\theta_{0}\right)\right\}=\{0\}$ for any $\theta_{0} \in \mathcal{C}$.

Conversely, suppose $\mathcal{C}$ is unbounded and fix $\theta_{0} \in \mathcal{C}$. Let

$$
U_{r}:=\left\{v \in S^{n-1}: \theta_{0}+c v \notin \mathcal{C} \text { for some } c \in(0, r)\right\}
$$

This set is open: if $\left(v_{n}\right)$ is a sequence in $U_{r}^{c}$ converging to $v$, then the fact that $\mathcal{C}$ is closed implies $\theta_{0}+r v_{n} \in \mathcal{C}$ for all $n$, and consequently $\theta_{0}+r v \in \mathcal{C}$ and finally $v \in U_{r}^{c}$.

If $\bigcup_{r>0} U_{r}$ is an open cover of the compact set $S^{n-1}$, then $S^{n-1} \subseteq U_{r}$ for some $r>0$, which implies $\mathcal{C} \subseteq B_{r}\left(\theta_{0}\right)$, a contradiction. Thus, some direction $v \in S^{n-1}$ does not lie in $\bigcup_{r>0} U_{r}$, that is, $\theta_{0}+c v \in \mathcal{C}$ for all $c \geq 0$. This implies $c v \in K_{\mathcal{C}}$ for all $c \geq 0$.

We now apply Proposition 5.3. If $\mathcal{C}$ is bounded, then so is $F_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$. Choosing $c$ large enough so that $F_{\mathcal{C}}\left(\Pi_{\mathcal{C}}\left(\theta^{*}\right)\right)$ lies in the ball of radius $c$ suffices to satisfy (38). Then Proposition 5.3 implies that the high $\sigma$ limits are $\delta\left(K_{\mathcal{C}}\right)=0$.

## Acknowledgments

We thank the reviewers for their helpful comments and suggestions. We also thank Dennis Amelunxen for an informative email correspondence and to Bodhisattva Sen for helpful discussions.

Billy Fang was supported by NSF GRFP.
Adityanand Guntuboyina was supported by NSF CAREER Grant DMS-1654589.

## References

[1] Amelunxen, D. and Lotz, M. (2017). Intrinsic volumes of polyhedral cones: A combinatorial perspective. Discrete Comput. Geom. 58 371-409. MR3679941
[2] Amelunxen, D., Lotz, M., McCoy, M.B. and Tropp, J.A. (2014). Living on the edge: Phase transitions in convex programs with random data. Inf. Inference 3 224-294. MR3311453
[3] Bellec, P.C. (2018). Sharp oracle inequalities for least squares estimators in shape restricted regression. Ann. Statist. 46 745-780. MR3782383
[4] Groeneboom, P. and Jongbloed, G. (2014). Nonparametric Estimation Under Shape Constraints: Estimators, Algorithms and Asymptotics. Cambridge Series in Statistical and Probabilistic Mathematics 38. New York: Cambridge Univ. Press. MR3445293
[5] Hiriart-Urruty, J.-B. and Lemaréchal, C. (2001). Fundamentals of Convex Analysis. Grundlehren Text Editions. Berlin: Springer. Abridged version of Convex analysis and minimization algorithms. I [Springer, Berlin, 1993; MR1261420 (95m:90001)] and II [ibid.; MR1295240 (95m:90002)]. MR1865628
[6] Jankowski, H. (2014). Convergence of linear functionals of the Grenander estimator under misspecification. Ann. Statist. 42 625-653. MR3210981
[7] Klivans, C.J. and Swartz, E. (2011). Projection volumes of hyperplane arrangements. Discrete Comput. Geom. 46 417-426. MR2826960
[8] Oymak, S. and Hassibi, B. (2016). Sharp MSE bounds for proximal denoising. Found. Comput. Math. 16 965-1029. MR3529131
[9] Pal, J.K. (2008). Spiking problem in monotone regression: Penalized residual sum of squares. Statist. Probab. Lett. 78 1548-1556. MR2453808
[10] Robertson, T., Wright, F.T. and Dykstra, R.L. (1988). Order Restricted Statistical Inference. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. Chichester: Wiley. MR0961262
[11] Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. J. Roy. Statist. Soc. Ser. B 58 267-288. MR1379242
[12] Wu, J., Meyer, M.C. and Opsomer, J.D. (2015). Penalized isotonic regression. J. Statist. Plann. Inference 161 12-24. MR3316548
[13] Zhang, C.-H. (2002). Risk bounds in isotonic regression. Ann. Statist. 30 528-555. MR1902898
Received June 2017 and revised February 2018

