
LASSO with Non-linear Measurements is Equivalent to One With Linear Measurements

Christos Thrampoulidis,
Department of Electrical Engineering
Caltech
cthrampo@caltech.edu

Ehsan Abbasi
Department of Electrical Engineering
Caltech
eabbasi@caltech.edu

Babak Hassibi
Department of Electrical Engineering
Caltech
hassibi@caltech.edu *

Abstract

Consider estimating an unknown, but structured (e.g. sparse, low-rank, etc.), signal $\mathbf{x}_0 \in \mathbb{R}^n$ from a vector $\mathbf{y} \in \mathbb{R}^m$ of measurements of the form $y_i = g_i(\mathbf{a}_i^T \mathbf{x}_0)$, where the \mathbf{a}_i 's are the rows of a known measurement matrix \mathbf{A} , and, $g(\cdot)$ is a (potentially unknown) nonlinear and random link-function. Such measurement functions could arise in applications where the measurement device has nonlinearities and uncertainties. It could also arise by design, e.g., $g_i(x) = \text{sign}(x + z_i)$, corresponds to noisy 1-bit quantized measurements. Motivated by the classical work of Brillinger, and more recent work of Plan and Vershynin, we estimate \mathbf{x}_0 via solving the Generalized-LASSO, i.e., $\hat{\mathbf{x}} := \arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2 + \lambda f(\mathbf{x})$ for some regularization parameter $\lambda > 0$ and some (typically non-smooth) convex regularizer $f(\cdot)$ that promotes the structure of \mathbf{x}_0 , e.g. ℓ_1 -norm, nuclear-norm, etc. While this approach seems to naively ignore the nonlinear function $g(\cdot)$, both Brillinger (in the non-constrained case) and Plan and Vershynin have shown that, when the entries of \mathbf{A} are iid standard normal, this is a good estimator of \mathbf{x}_0 up to a constant of proportionality μ , which only depends on $g(\cdot)$. In this work, we considerably strengthen these results by obtaining explicit expressions for $\|\hat{\mathbf{x}} - \mu\mathbf{x}_0\|_2$, for the *regularized* Generalized-LASSO, that are asymptotically *precise* when m and n grow large. A main result is that the estimation performance of the Generalized LASSO with non-linear measurements is *asymptotically the same* as one whose measurements are linear $y_i = \mu\mathbf{a}_i^T \mathbf{x}_0 + \sigma z_i$, with $\mu = \mathbb{E}\gamma g(\gamma)$ and $\sigma^2 = \mathbb{E}(g(\gamma) - \mu\gamma)^2$, and, γ standard normal. To the best of our knowledge, the derived expressions on the estimation performance are the first-known precise results in this context. One interesting consequence of our result is that the optimal quantizer of the measurements that minimizes the estimation error of the Generalized LASSO is the celebrated Lloyd-Max quantizer.

*This work was supported in part by the National Science Foundation under grants CNS-0932428, CCF-1018927, CCF-1423663 and CCF-1409204, by a grant from Qualcomm Inc., by NASA's Jet Propulsion Laboratory through the President and Directors Fund, by King Abdulaziz University, and by King Abdullah University of Science and Technology.

1 Introduction

Non-linear Measurements. Consider the problem of estimating an unknown signal vector $\mathbf{x}_0 \in \mathbb{R}^n$ from a vector $\mathbf{y} = (y_1, y_2, \dots, y_m)^T$ of m measurements taking the following form:

$$y_i = g_i(\mathbf{a}_i^T \mathbf{x}_0), \quad i = 1, 2, \dots, m. \quad (1)$$

Here, each \mathbf{a}_i represents a (known) measurement vector. The g_i 's are independent copies of a (generically random) link function g . For instance, $g_i(x) = x + z_i$, with say z_i being normally distributed, recovers the standard linear regression setup with gaussian noise. In this paper, we are particularly interested in scenarios where g is *non-linear*. Notable examples include $g(x) = \text{sign}(x)$ (or $g_i(x) = \text{sign}(x + z_i)$) and $g(x) = (x)_+$, corresponding to 1-bit quantized (noisy) measurements, and, to the censored Tobit model, respectively. Depending on the situation, g might be known or unspecified. In the statistics and econometrics literature, the measurement model in (1) is popular under the name *single-index model* and several aspects of it have been well-studied, e.g. [5, 6, 16, 17]¹.

Structured Signals. It is typical that the unknown signal \mathbf{x}_0 obeys some sort of *structure*. For instance, it might be sparse, i.e. only a few $k \ll n$, of its entries are non-zero; or, it might be that $\mathbf{x}_0 = \text{vec}(\mathbf{X}_0)$, where $\mathbf{X}_0 \in \mathbb{R}^{\sqrt{n} \times \sqrt{n}}$ is a matrix of low-rank $r \ll n$. To exploit this information it is typical to associate with the structure of \mathbf{x}_0 a properly chosen function $f: \mathbb{R}^n \rightarrow \mathbb{R}$, which we refer to as the *regularizer*. Of particular interest are *convex* and non-smooth such regularizers, e.g. the ℓ_1 -norm for sparse signals, the nuclear-norm for low-rank ones, etc. Please refer to [2, 7, 15] for further discussions.

An Algorithm for Linear Measurements: The Generalized LASSO. When the link function is *linear*, i.e. $g_i(x) = x + z_i$, perhaps the most popular way of estimating \mathbf{x}_0 is via solving the Generalized LASSO algorithm:

$$\hat{\mathbf{x}} := \arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2 + \lambda f(\mathbf{x}). \quad (2)$$

Here, $\mathbf{A} = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m]^T \in \mathbb{R}^{m \times n}$ is the known measurement matrix and $\lambda > 0$ is a regularizer parameter. This is often referred to as the ℓ_2 -LASSO or the square-root-LASSO [4] to distinguish from the one solving $\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda f(\mathbf{x})$, instead. Our results can be accustomed to this latter version, but for concreteness, we restrict attention to (2) throughout. The acronym LASSO for (2) was introduced in [31] for the special case of ℓ_1 -regularization; (2) is a natural generalization to other kinds of structures and includes the group-LASSO [34], the fused-LASSO [32] as special cases. We often drop the term ‘‘Generalized’’ and refer to (2) simply as the LASSO.

One popular, measure of estimation performance of (2) is the squared-error $\|\hat{\mathbf{x}} - \mathbf{x}_0\|_2^2$. Recently, there have been significant advances on establishing tight bounds and even *precise* characterizations of this quantity, in the presence of linear measurements [3, 11, 19, 25, 28, 30]. Such precise results have been core to building a better understanding of the behavior of the LASSO, and, in particular, on the exact role played by the choice of the regularizer f (in accordance with the structure of \mathbf{x}_0), by the number of measurements m , by the value of λ , etc.. In certain cases, they even provide us with useful insights into practical matters such as the tuning of the regularizer parameter.

Using the LASSO for Non-linear Measurements?. The LASSO is by nature tailored to a linear model for the measurements. Indeed, the first term of the objective function in (2) tries to fit $\mathbf{A}\mathbf{x}$ to the observed vector \mathbf{y} presuming that this is of the form $y_i = \mathbf{a}_i^T \mathbf{x}_0 + \text{noise}$. Of course, no one stops us from continuing to use it even in cases where $y_i = g(\mathbf{a}_i^T \mathbf{x}_0)$ with g being *non-linear*². But, the question then becomes: Can there be any guarantees that the solution $\hat{\mathbf{x}}$ of the Generalized LASSO is still a good estimate of \mathbf{x}_0 ?

The question just posed was first studied back in the early 80's by Brillinger [6] who provided answers in the case of solving (2) without a regularizer term. This, of course, corresponds to standard Least Squares (LS). Interestingly, he showed that when the measurement vectors are Gaussian, then the LS solution is a consistent estimate of \mathbf{x}_0 , up to a constant of proportionality μ , which only depends on the link-function g . The result is sharp, but only under the assumption that the number of measurements m grows large, while the signal dimension n stays fixed, which was the typical

¹ The single-index model is a classical topic and can also be regarded as a special case of what is known as *sufficient dimension reduction* problem. There is extensive literature on both subjects; unavoidably, we only refer to the directly relevant works here.

²Note that the Generalized LASSO in (2) does not assume knowledge of g . All that is assumed is the availability of the measurements y_i . Thus, the link-function might as well be unknown or unspecified.

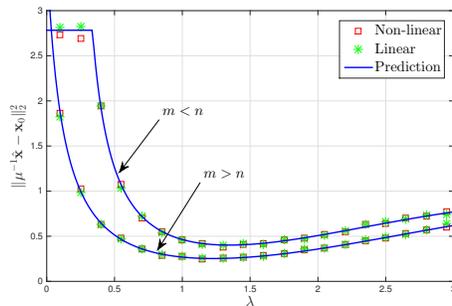


Figure 1: Squared error of the ℓ_1 -regularized LASSO with non-linear measurements (\square) and with corresponding linear ones (\ast) as a function of the regularizer parameter λ ; both compared to the asymptotic prediction. Here, $g_i(x) = \text{sign}(x + 0.3z_i)$ with $z_i \sim \mathcal{N}(0, 1)$. The unknown signal \mathbf{x}_0 is of dimension $n = 768$ and has $\lceil 0.15n \rceil$ non-zero entries (see Sec. 2.2.2 for details). The different curves correspond to $\lceil 0.75n \rceil$ and $\lceil 1.2n \rceil$ number of measurements, respectively. Simulation points are averages over 20 problem realizations.

setting of interest at the time. In the world of structured signals and high-dimensional measurements, the problem was only very recently revisited by Plan and Vershynin [20]. They consider a *constrained* version of the Generalized LASSO, in which the regularizer is essentially replaced by a constraint, and derive upper bounds on its performance. The bounds are not tight (they involve absolute constants), but they demonstrate some key features: i) the solution to the constrained LASSO $\hat{\mathbf{x}}$ is a good estimate of \mathbf{x}_0 up to the same constant of proportionality μ that appears in Brillinger’s result. ii) Thus, $\|\hat{\mathbf{x}} - \mu\mathbf{x}_0\|_2^2$ is a natural measure of performance. iii) Estimation is possible even with $m < n$ measurements by taking advantage of the structure of \mathbf{x}_0 .

1.1 Summary of Contributions

Inspired by the work of Plan and Vershynin [20], and, motivated by recent advances on the precise analysis of the Generalized LASSO with linear measurements, this paper extends these latter results to the case of non-linear measurements. When the measurement matrix \mathbf{A} has entries i.i.d. Gaussian (henceforth, we assume this to be the case without further reference), and the estimation performance is measured in a mean-squared-error sense, we are able to *precisely* predict the asymptotic behavior of the error. The derived expression accurately captures the role of the link function g , the particular structure of \mathbf{x}_0 , the role of the regularizer f , and, the value of the regularizer parameter λ . Further, it holds for all values of λ , and for a wide class of functions f and g .

Interestingly, our result shows in a very precise manner that in large dimensions, modulo the information about the magnitude of \mathbf{x}_0 , the LASSO treats non-linear measurements exactly as if they were scaled and noisy linear measurements with scaling factor μ and noise variance σ^2 defined as

$$\mu := \mathbb{E}[\gamma g(\gamma)], \quad \text{and} \quad \sigma^2 := \mathbb{E}[(g(\gamma) - \mu\gamma)^2], \quad \text{for } \gamma \sim \mathcal{N}(0, 1), \quad (3)$$

where the expectation is with respect to both γ and g . In particular, when g is such that $\mu \neq 0^3$, then,

the estimation performance of the Generalized LASSO with measurements of the form
 $y_i = g_i(\mathbf{a}_i^T \mathbf{x}_0)$ *is asymptotically the same as if the measurements were rather of the form*
 $y_i = \mu \mathbf{a}_i^T \mathbf{x}_0 + \sigma z_i$, *with* μ, σ^2 *as in (3) and* z_i *standard gaussian noise.*

Recent analysis of the squared-error of the LASSO, when used to recover structured signals from noisy *linear* observations, provides us with either precise predictions (e.g. [3, 29]), or in other cases, with tight upper bounds (e.g. [11, 19]). Owing to the established relation between non-linear and (corresponding) linear measurements, such results also characterize the performance of the LASSO in the presence of nonlinearities. We remark that some of the error formulae derived here in the general context of non-linear measurements, have not been previously known even under the prism of linear measurements. Figure 1 serves as an illustration; the error with non-linear measurements matches well with the error of the corresponding linear ones and both are accurately predicted by our analytic expression.

³This excludes for example link functions g that are even, but also some other not so obvious cases [12, Sec. 2.2]. For a few special cases, e.g. sparse recovery with binary measurements y_i [33], different methodologies than the LASSO have been recently proposed that do not require $\mu = 0$.

Under the generic model in (1), which allows for g to even be unspecified, \mathbf{x}_0 can, in principle, be estimated only up to a constant of proportionality [6, 17, 20]. For example, if g is unknown then any information about the norm $\|\mathbf{x}_0\|_2$ could be absorbed in the definition of g . The same is true when $g(x) = \text{sign}(x)$, even though g might be known here. In these cases, what becomes important is the *direction* of \mathbf{x}_0 . Motivated by this, and, in order to simplify the presentation, we have assumed throughout that \mathbf{x}_0 has unit Euclidean norm⁴, i.e. $\|\mathbf{x}_0\|_2 = 1$.

1.2 Discussion of Relevant Literature

Extending an Old Result. Brillinger [6] identified the asymptotic behavior of the estimation error of the LS solution $\hat{\mathbf{x}}_{LS} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{y}$ by showing that, when n (the dimension of \mathbf{x}_0) is fixed,

$$\lim_{m \rightarrow \infty} \sqrt{m} \|\hat{\mathbf{x}}_{LS} - \mu \mathbf{x}_0\|_2 = \sigma \sqrt{n}, \quad (4)$$

where μ and σ^2 are same as in (3). Our result can be viewed as a generalization of the above in several directions. First, we extend (4) to the regime where $m/n = \delta \in (1, \infty)$ and both grow large by showing that

$$\lim_{n \rightarrow \infty} \|\hat{\mathbf{x}}_{LS} - \mu \mathbf{x}_0\|_2 = \frac{\sigma}{\sqrt{\delta - 1}}. \quad (5)$$

Second, and most importantly, we consider solving the Generalized LASSO instead, to which LS is only a very special case. This allows versions of (5) where the error is finite even when $\delta < 1$ (e.g., see (8)). Note the additional challenges faced when considering the LASSO: i) $\hat{\mathbf{x}}$ no longer has a closed-form expression, ii) the result needs to additionally capture the role of \mathbf{x}_0 , f , and, λ .

Motivated by Recent Work. Plan and Vershynin consider a *constrained* Generalized LASSO:

$$\hat{\mathbf{x}}_{\text{C-LASSO}} = \arg \min_{\mathbf{x} \in \mathcal{K}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2, \quad (6)$$

with \mathbf{y} as in (1) and $\mathcal{K} \subset \mathbb{R}^n$ some known set (not necessarily convex). In its simplest form, their result shows that when $m \gtrsim D_{\mathcal{K}}(\mu \mathbf{x}_0)$ then with high probability,

$$\|\hat{\mathbf{x}}_{\text{C-LASSO}} - \mu \mathbf{x}_0\|_2 \lesssim \frac{\sigma \sqrt{D_{\mathcal{K}}(\mu \mathbf{x}_0)} + \zeta}{\sqrt{m}}. \quad (7)$$

Here, $D_{\mathcal{K}}(\mu \mathbf{x}_0)$ is the Gaussian width, a specific measure of complexity of the constrained set \mathcal{K} when viewed from $\mu \mathbf{x}_0$. For our purposes, it suffices to remark that if \mathcal{K} is properly chosen, and, if $\mu \mathbf{x}_0$ is on the boundary of \mathcal{K} , then $D_{\mathcal{K}}(\mu \mathbf{x}_0)$ is less than n . Thus, estimation is in principle possible with $m < n$ measurements. The parameters μ and σ that appear in (7) are the same as in (3) and $\zeta := \mathbb{E}[(g(\gamma) - \mu\gamma)^2 \gamma^2]$. Observe that, in contrast to (4) and to the setting of this paper, the result in (7) is non-asymptotic. Also, it suggests the critical role played by μ and σ . On the other hand, (7) is only an upper bound on the error, and also, it suffers from unknown absolute proportionality constants (hidden in \lesssim).

Moving the analysis into an asymptotic setting, our work expands upon the result of [20]. First, we consider the regularized LASSO instead, which is more commonly used in practice. Most importantly, we improve the loose upper bounds into *precise* expressions. In turn, this proves in an exact manner the role played by μ and σ^2 to which (7) is only indicative. For a direct comparison with (7) we mention the following result which follows from our analysis (we omit the proof for brevity). Assume \mathcal{K} is convex, $m/n = \delta \in (0, \infty)$, $D_{\mathcal{K}}(\mu \mathbf{x}_0)/n = \rho \in (0, 1]$ and $n \rightarrow \infty$. Also, $\delta > \rho$. Then, (7) yields an upper bound $C\sigma\sqrt{\rho/\delta}$ to the error, for some constant $C > 0$. Instead, we show

$$\|\hat{\mathbf{x}}_{\text{C-LASSO}} - \mu \mathbf{x}_0\|_2 \leq \sigma \frac{\sqrt{\rho}}{\sqrt{\delta - \rho}}. \quad (8)$$

Precise Analysis of the LASSO With Linear Measurements. The first *precise* formula predicting the limiting behavior of the LASSO reconstruction error was established in [3, 11]. The authors consider the ℓ_2^2 -LASSO with ℓ_1 -regularization and the analysis is based on the the Approximate Message Passing (AMP) framework [10]. A more general line of work, [19, 25, 29, 30] studies the problem using a recently developed framework termed Convex Gaussian Min-max Theorem (CGMT) [28], which is a *tight* version of a classical Gaussian comparison inequality by Gordon

⁴In [20, Remark 1.8], they note that their results can be easily generalized to the case when $\|\mathbf{x}_0\|_2 \neq 1$ by simply redefining $\tilde{g}(x) = g(\|\mathbf{x}_0\|_2 x)$ and accordingly adjusting the values of the parameters μ and σ^2 in (3). The very same argument is also true in our case.

[14]. The CGMT framework was initially used by Stojnic [25] to derive tight upper bounds on the constrained LASSO with ℓ_1 -regularization; [19] generalized those to general convex regularizers and also to the ℓ_2 -LASSO; the ℓ_2^2 -LASSO was studied in [30]. Those bounds hold for all values of SNR, but they become tight only in the high-SNR regime. A precise error expression was derived in [29] for the ℓ_2 -LASSO with ℓ_1 -regularization under a gaussianity assumption on the distribution of the non-zero entries of \mathbf{x}_0 . When measurements are linear, our Theorem 2.3 generalizes this assumption. Our main Theorem 2.2 provides error predictions for regularizers going beyond the ℓ_1 -norm, e.g. $\ell_{1,2}$ -norm, nuclear norm, which appear to be novel. When it comes to non-linear measurements, to the best of our knowledge, this paper is the first to derive asymptotically precise results on the performance of any LASSO-type program.

2 Results

2.1 Modeling Assumptions

Unknown structured signal. We let $\mathbf{x}_0 \in \mathbb{R}^n$ represent the unknown signal vector. We assume that $\mathbf{x}_0 = \bar{\mathbf{x}}_0 / \|\bar{\mathbf{x}}_0\|_2$, with $\bar{\mathbf{x}}_0$ sampled from a probability density $p_{\bar{\mathbf{x}}_0}$ in \mathbb{R}^n . Thus, \mathbf{x}_0 is deterministically of unit Euclidean-norm (this is mostly to simplify the presentation, see Footnote 4). Information about the structure of $\bar{\mathbf{x}}_0$ (and correspondingly of \mathbf{x}_0) is encoded in $p_{\bar{\mathbf{x}}_0}$. E.g., to study an $\bar{\mathbf{x}}_0$ which is sparse, it is typical to assume that its entries are i.i.d. $\bar{\mathbf{x}}_{0,i} \sim (1 - \rho)\delta_0 + \rho q_{\bar{\mathbf{x}}_0}$, where $\rho \in (0, 1)$ becomes the normalized sparsity level, $q_{\bar{\mathbf{x}}_0}$ is a scalar p.d.f. and δ_0 is the Dirac delta function⁵.

Regularizer. We consider *convex* regularizers $f : \mathbb{R}^n \rightarrow \mathbb{R}$.

Measurement matrix. The entries of $\mathbf{A} \in \mathbb{R}^{m \times n}$ are i.i.d. $\mathcal{N}(0, 1)$.

Measurements and Link-function. We observe $\mathbf{y} = \vec{g}(\mathbf{A}\mathbf{x}_0)$ where \vec{g} is a (possibly random) map from \mathbb{R}^m to \mathbb{R}^m and $\vec{g}(\mathbf{u}) = [g_1(u_1), \dots, g_m(u_m)]^T$. Each g_i is i.i.d. from a real valued random function g for which μ and σ^2 are defined in (3). We assume that μ and σ^2 are nonzero and bounded.

Asymptotics. We study a linear asymptotic regime. In particular, we consider a sequence of problem instances $\{\bar{\mathbf{x}}_0^{(n)}, \mathbf{A}^{(n)}, f^{(n)}, m^{(n)}\}_{n \in \mathbb{N}}$ indexed by n such that $\mathbf{A}^{(n)} \in \mathbb{R}^{m \times n}$ has entries i.i.d. $\mathcal{N}(0, 1)$, $f^{(n)} : \mathbb{R}^n \rightarrow \mathbb{R}$ is proper convex, and, $m := m^{(n)}$ with $m = \delta n$, $\delta \in (0, \infty)$. We further require that the following conditions hold:

- (a) $\bar{\mathbf{x}}_0^{(n)}$ is sampled from a probability density $p_{\bar{\mathbf{x}}_0}^{(n)}$ in \mathbb{R}^n with one-dimensional marginals that are independent of n and have bounded second moments. Furthermore, $n^{-1} \|\bar{\mathbf{x}}_0^{(n)}\|_2^2 \xrightarrow{P} \sigma_x^2 = 1$.
- (b) For any $n \in \mathbb{N}$ and any $\|\mathbf{x}\|_2 \leq C$, it holds $n^{-1/2} f(\mathbf{x}) \leq c_1$ and $n^{-1/2} \max_{\mathbf{s} \in \partial f^{(n)}(\mathbf{x})} \|\mathbf{s}\|_2 \leq c_2$, for constants $c_1, c_2, C \geq 0$ independent of n .

In (a), we used “ \xrightarrow{P} ” to denote convergence in probability as $n \rightarrow \infty$. The assumption $\sigma_x^2 = 1$ holds without loss of generality, and, is only necessary to simplify the presentation. In (b), $\partial f(\mathbf{x})$ denotes the subdifferential of f at \mathbf{x} . The condition itself is no more than a normalization condition on f .

Every such sequence $\{\bar{\mathbf{x}}_0^{(n)}, \mathbf{A}^{(n)}, f^{(n)}\}_{n \in \mathbb{N}}$ generates a sequence $\{\mathbf{x}_0^{(n)}, \mathbf{y}^{(n)}\}_{n \in \mathbb{N}}$ where $\mathbf{x}_0^{(n)} := \bar{\mathbf{x}}_0^{(n)} / \|\bar{\mathbf{x}}_0^{(n)}\|_2$ and $\mathbf{y}^{(n)} := \vec{g}^{(n)}(\mathbf{A}\mathbf{x}_0)$. When clear from the context, we drop the superscript (n) .

2.2 Precise Error Prediction

Let $\{\mathbf{x}_0^{(n)}, \mathbf{A}^{(n)}, f^{(n)}, \mathbf{y}^{(n)}\}_{n \in \mathbb{N}}$ be a sequence of problem instances that satisfying all the conditions above. With these, define the sequence $\{\hat{\mathbf{x}}^{(n)}\}_{n \in \mathbb{N}}$ of solutions to the corresponding LASSO problems for fixed $\lambda > 0$:

$$\hat{\mathbf{x}}^{(n)} := \min_{\mathbf{x}} \frac{1}{\sqrt{n}} \left\{ \|\mathbf{y}^{(n)} - \mathbf{A}^{(n)}\mathbf{x}\|_2 + \lambda f^{(n)}(\mathbf{x}) \right\}. \quad (9)$$

The main contribution of this paper is a precise evaluation of $\lim_{n \rightarrow \infty} \|\mu^{-1} \hat{\mathbf{x}}^{(n)} - \mathbf{x}_0^{(n)}\|_2^2$ with high probability over the randomness of \mathbf{A} , of \mathbf{x}_0 , and of g .

2.2.1 General Result

To state the result in a general framework, we require a further assumption on $p_{\bar{\mathbf{x}}_0}^{(n)}$ and $f^{(n)}$. Later in this section we illustrate how this assumption can be naturally met. We write f^* for the Fenchel’s conjugate of f , i.e., $f^*(\mathbf{v}) := \sup_{\mathbf{x}} \mathbf{x}^T \mathbf{v} - f(\mathbf{x})$; also, we denote the Moreau envelope of f at \mathbf{v} with index τ to be $e_{f, \tau}(\mathbf{v}) := \min_{\mathbf{x}} \left\{ \frac{1}{2} \|\mathbf{v} - \mathbf{x}\|_2^2 + \tau f(\mathbf{x}) \right\}$.

⁵Such models have been widely used in the relevant literature, e.g. [8,9,11]. In fact, the results here continue to hold as long as the marginal distribution of $\bar{\mathbf{x}}_0$ converges to a given distribution (as in [3]).

Assumption 1. We say Assumption 1 holds if for all non-negative constants $c_1, c_2, c_3 \in \mathbb{R}$ the point-wise limit of $\frac{1}{n} e_{\sqrt{n}(f^*)^{(n)}, c_3} (c_1 \mathbf{h} + c_2 \bar{\mathbf{x}}_0)$ exists with probability one over $\mathbf{h} \sim \mathcal{N}(0, \mathbf{I}_n)$ and $\bar{\mathbf{x}}_0 \sim p_{\bar{\mathbf{x}}_0}^{(n)}$. Then, we denote the limiting value as $F(c_1, c_2, c_3)$.

Theorem 2.1 (Non-linear=Linear). Consider the asymptotic setup of Section 2.1 and let Assumption 1 hold. Recall μ and σ^2 as in (3) and let $\hat{\mathbf{x}}$ be the minimizer of the Generalized LASSO in (9) for fixed $\lambda > 0$ and for measurements given by (1). Further let $\hat{\mathbf{x}}^{lin}$ be the solution to the Generalized LASSO when used with linear measurements of the form $\mathbf{y}^{lin} = \mathbf{A}(\mu \mathbf{x}_0) + \sigma \mathbf{z}$, where \mathbf{z} has entries i.i.d. standard normal. Then, in the limit of $n \rightarrow \infty$, with probability one,

$$\|\hat{\mathbf{x}} - \mu \mathbf{x}_0\|_2^2 = \|\hat{\mathbf{x}}^{lin} - \mu \mathbf{x}_0\|_2^2.$$

Theorem 2.1 relates in a very precise manner the error of the Generalized LASSO under non-linear measurements to the error of the same algorithm when used under appropriately scaled noisy linear measurements. Theorem 2.2 below, derives an asymptotically exact expression for the error.

Theorem 2.2 (Precise Error Formula). Under the same assumptions of Theorem 2.1 and $\delta := m/n$, it holds, with probability one,

$$\lim_{n \rightarrow \infty} \|\hat{\mathbf{x}} - \mu \mathbf{x}_0\|_2^2 = \alpha_*^2,$$

where α_* is the unique optimal solution to the convex program

$$\max_{\substack{0 \leq \beta \leq 1 \\ \tau \geq 0}} \min_{\alpha \geq 0} \beta \sqrt{\delta} \sqrt{\alpha^2 + \sigma^2} - \frac{\alpha \tau}{2} + \frac{\mu^2 \tau}{2\alpha} - \frac{\alpha \lambda^2}{\tau} F\left(\frac{\beta}{\lambda}, \frac{\mu \tau}{\lambda \alpha}, \frac{\tau}{\lambda \alpha}\right). \quad (10)$$

Also, the optimal cost of the LASSO in (9) converges to the optimal cost of the program in (10).

Under the stated conditions, Theorem 2.2 proves that the limit of $\|\hat{\mathbf{x}} - \mu \mathbf{x}_0\|_2$ exists and is equal to the *unique* solution of the optimization program in (10). Notice that this is a *deterministic* and *convex* optimization, which only involves three *scalar* optimization variables. Thus, the optimal α_* can, in principle, be efficiently numerically computed. In many specific cases of interest, with some extra effort, it is possible to yield simpler expressions for α_* , e.g. see Theorem 2.3 below. The role of the normalized number of measurement $\delta = m/n$, of the regularizer parameter λ , and, that of g , through μ and σ^2 , are explicit in (10); the structure of \mathbf{x}_0 and the choice of the regularizer f are implicit in F . Figures 1-2 illustrate the accuracy of the prediction of the theorem in a number of different settings. The proofs of both the Theorems are deferred to Appendix A. In the next sections, we specialize Theorem 2.2 to the cases of sparse, group-sparse and low-rank signal recovery.

2.2.2 Sparse Recovery

Assume each entry $\bar{x}_{0,i}, i = 1, \dots, n$ is sampled i.i.d. from a distribution

$$p_{\bar{X}_0}(x) = (1 - \rho) \cdot \delta_0(x) + \rho \cdot q_{\bar{X}_0}(x), \quad (11)$$

where δ_0 is the delta Dirac function, $\rho \in (0, 1)$ and $q_{\bar{X}_0}$ a probability density function with second moment normalized to $1/\rho$ so that condition (a) of Section 2.1 is satisfied. Then, $\mathbf{x}_0 = \bar{\mathbf{x}}_0 / \|\bar{\mathbf{x}}_0\|_2$ is ρn -sparse on average and has unit Euclidean norm. Letting $f(\mathbf{x}) = \|\mathbf{x}\|_1$ also satisfies condition (b). Let us now check Assumption 1. The Fenchel's conjugate of the ℓ_1 -norm is simply the indicator function of the ℓ_∞ unit ball. Hence, without much effort,

$$\begin{aligned} \frac{1}{n} e_{\sqrt{n}(f^*)^{(n)}, c_3} (c_1 \mathbf{h} + c_2 \bar{\mathbf{x}}_0) &= \frac{1}{2n} \sum_{i=1}^n \min_{|v_i| \leq 1} (v_i - (c_1 \mathbf{h}_i + c_2 \bar{x}_{0,i}))^2 \\ &= \frac{1}{2n} \sum_{i=1}^n \eta^2(c_1 \mathbf{h}_i + c_2 \bar{x}_{0,i}; 1), \end{aligned} \quad (12)$$

where we have denoted

$$\eta(x; \tau) := (x/|x|) (|x| - \tau)_+ \quad (13)$$

for the soft thresholding operator. An application of the weak law of large numbers to see that the limit of the expression in (12) equals $F(c_1, c_2, c_3) := \frac{1}{2} \mathbb{E} [\eta^2(c_1 h + c_2 \bar{X}_0; 1)]$, where the expectation is over $h \sim \mathcal{N}(0, 1)$ and $\bar{X}_0 \sim p_{\bar{X}_0}$. With all these, Theorem 2.2 is applicable. We have put extra effort in order to obtain the following equivalent but more insightful characterization of the error, as stated below and proved in Appendix B.

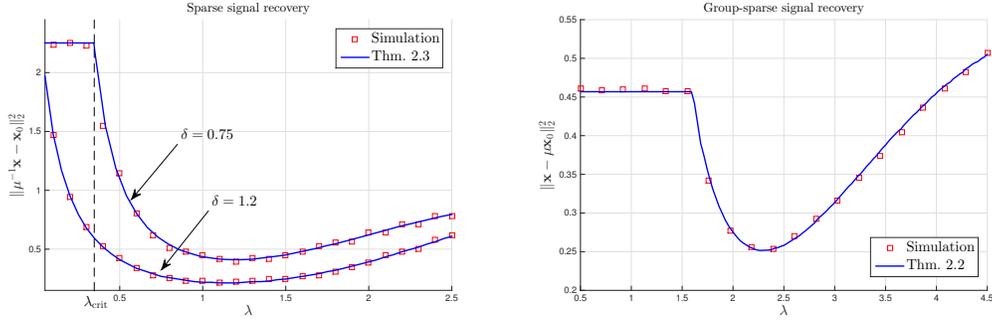


Figure 2: Squared error of the LASSO as a function of the regularizer parameter compared to the asymptotic predictions. Simulation points represent averages over 20 realizations. (a) Illustration of Thm. 2.3 for $g(x) = \text{sign}(x)$, $n = 512$, $p_{\bar{X}_0}(+1) = p_{\bar{X}_0}(-1) = 0.05$, $p_{\bar{X}_0}(+1) = 0.9$ and two values of δ , namely 0.75 and 1.2. (b) Illustration of Thm. 2.2 for \mathbf{x}_0 being group-sparse as in Section 2.2.3 and $g_i(x) = \text{sign}(x + 0.3z_i)$. In particular, \mathbf{x}_0 is composed of $t = 512$ blocks of block size $b = 3$. Each block is zero with probability 0.95, otherwise its entries are i.i.d. $\mathcal{N}(0, 1)$. Finally, $\delta = 0.75$.

Theorem 2.3 (Sparse Recovery). *If $\delta > 1$, then define $\lambda_{\text{crit}} = 0$. Otherwise, let $\lambda_{\text{crit}}, \kappa_{\text{crit}}$ be the unique pair of solutions to the following set of equations:*

$$\begin{cases} \kappa^2 \delta = \sigma^2 + \mathbb{E}[(\eta(\kappa h + \mu \bar{X}_0; \kappa \lambda) - \mu \bar{X}_0)^2], & (14) \\ \kappa \delta = \mathbb{E}[(\eta(\kappa h + \mu \bar{X}_0; \kappa \lambda) \cdot h)], & (15) \end{cases}$$

where $h \sim \mathcal{N}(0, 1)$ and is independent of $\bar{X}_0 \sim p_{\bar{X}_0}$. Then, for any $\lambda > 0$, with probability one,

$$\lim_{n \rightarrow \infty} \|\hat{\mathbf{x}} - \mu \mathbf{x}_0\|_2^2 = \begin{cases} \delta \kappa_{\text{crit}}^2 - \sigma^2 & , \lambda \leq \lambda_{\text{crit}}, \\ \delta \kappa_*^2(\lambda) - \sigma^2 & , \lambda \geq \lambda_{\text{crit}}, \end{cases}$$

where $\kappa_*^2(\lambda)$ is the unique solution to (14).

Figures 1 and 2(a) validate the prediction of the theorem, for different signal distributions, namely $q_{\bar{X}_0}$ being Gaussian and Bernoulli, respectively. For the case of compressed ($\delta < 1$) measurements, observe the two different regimes of operation, one for $\lambda \leq \lambda_{\text{crit}}$ and the other for $\lambda \geq \lambda_{\text{crit}}$, precisely as they are predicted by the theorem (see also [19, Sec. 8]). The special case of Theorem 2.3 for which $q_{\bar{X}_0}$ is Gaussian has been previously studied in [29]. Otherwise, to the best of our knowledge, this is the first precise analysis result for the ℓ_2 -LASSO stated in that generality. Analogous result, but via different analysis tools, has only been known for the ℓ_2^2 -LASSO as appears in [3].

2.2.3 Group-Sparse Recovery

Let $\bar{\mathbf{x}}_0 \in \mathbb{R}^n$ be composed of t non-overlapping blocks of constant size b each such that $n = t \cdot b$. Each block $[\bar{\mathbf{x}}_0]_i, i = 1, \dots, t$ is sampled i.i.d. from a probability density in \mathbb{R}^b : $p_{\bar{X}_0}(\mathbf{x}) = (1 - \rho) \cdot \delta_0(\mathbf{x}) + \rho \cdot q_{\bar{X}_0}(\mathbf{x}), \mathbf{x} \in \mathbb{R}^b$, where $\rho \in (0, 1)$. Thus, $\bar{\mathbf{x}}_0$ is ρt -block-sparse on average. We operate in the regime of linear measurements $m/n = \delta \in (0, \infty)$. As is common we use the $\ell_{1,2}$ -norm to induce block-sparsity, i.e., $f(\mathbf{x}) = \sum_{i=1}^t \|\bar{\mathbf{x}}_0\|_2$; with this, (9) is often referred to as group-LASSO in the literature [34]. It is not hard to show that Assumption 1 holds with $F(c_1, c_2, c_3) := \frac{1}{2b} \mathbb{E}[\|\bar{\eta}(c_1 \mathbf{h} + c_2 \bar{X}_0; 1)\|_2^2]$, where $\bar{\eta}(\mathbf{x}; \tau) = \mathbf{x} / \|\mathbf{x}\| (\|\mathbf{x}\|_2 - \tau)_+, \mathbf{x} \in \mathbb{R}^b$ is the vector soft thresholding operator and $h \sim \mathcal{N}(0, \mathbf{I}_b), \bar{X}_0 \sim p_{\bar{X}_0}$ and are independent. Thus Theorem 2.2 is applicable in this setting; Figure 2(b) illustrates the accuracy of the prediction.

2.2.4 Low-rank Matrix Recovery

Let $\bar{\mathbf{X}}_0 \in \mathbb{R}^{d \times d}$ be an unknown matrix of rank r , in which case, $\bar{\mathbf{x}}_0 = \text{vec}(\bar{\mathbf{X}}_0)$ with $n = d^2$. Assume $m/d^2 = \delta \in (0, \infty)$ and $r/d = \rho \in (0, 1)$. As usual in this setting, we consider nuclear-norm regularization; in particular, we choose $f(\mathbf{x}) = \sqrt{d} \|\mathbf{X}\|_*$. Each subgradient $\mathbf{S} \in \partial f(\mathbf{X})$ then satisfies $\|\mathbf{S}\|_F \leq d$ in agreement with assumption (b) of Section 2.1. Furthermore, for this choice of regularizer, we have

$$\begin{aligned} \frac{1}{n} e^{\sqrt{n}(f^*)^{(n)}, c_3} (c_1 \mathbf{H} + c_2 \bar{\mathbf{X}}_0) &= \frac{1}{2d^2} \min_{\|\mathbf{V}\|_2 \leq \sqrt{d}} \|\mathbf{V} - (c_1 \mathbf{H} + c_2 \bar{\mathbf{X}}_0)\|_F^2 \\ &= \frac{1}{2d} \min_{\|\mathbf{V}\|_2 \leq 1} \|\mathbf{V} - d^{-1/2}(c_1 \mathbf{H} + c_2 \bar{\mathbf{X}}_0)\|_F^2 = \frac{1}{2d} \sum_{i=1}^d \eta^2 \left(s_i \left(d^{-1/2}(c_1 \mathbf{H} + c_2 \bar{\mathbf{X}}_0) \right); 1 \right), \end{aligned}$$

where $\eta(\cdot; \cdot)$ is as in (13), $s_i(\cdot)$ denotes the i^{th} singular value of its argument and $\mathbf{H} \in \mathbb{R}^{d \times d}$ has entries $\mathcal{N}(0, 1)$. If conditions are met such that the empirical distribution of the singular values of (the sequence of random matrices) $c_1 \mathbf{H} + c_2 \bar{\mathbf{X}}_0$ converges asymptotically to a limiting distribution, say $q(c_1, c_2)$, then $F(c_1, c_2, c_3) := \frac{1}{2} \mathbb{E}_{x \sim q(c_1, c_2)} [\eta^2(x; 1)]$, and Theorem 2.1–2.2 apply. For instance, this will be the case if $d^{-1/2} \bar{\mathbf{X}}_0 = \mathbf{U} \mathbf{S} \mathbf{V}^t$, where \mathbf{U}, \mathbf{V} unitary matrices and \mathbf{S} is a diagonal matrix whose entries have a given marginal distribution with bounded moments (in particular, independent of d). We leave the details and the problem of (numerically) evaluating F for future work.

2.3 An Application to q -bit Compressive Sensing

2.3.1 Setup

Consider recovering a *sparse* unknown signal $\mathbf{x}_0 \in \mathbb{R}^n$ from scalar q -bit quantized linear measurements. Let $\mathbf{t} := \{t_0 = 0, t_1, \dots, t_{L-1}, t_L = +\infty\}$ represent a (symmetric with respect to 0) set of decision thresholds and $\boldsymbol{\ell} := \{\pm \ell_1, \pm \ell_2, \dots, \pm \ell_L\}$ the corresponding representation points, such that $L = 2^{q-1}$. Then, quantization of a real number x into q -bits can be represented as

$$\mathcal{Q}_q(x, \boldsymbol{\ell}, \mathbf{t}) = \text{sign}(x) \sum_{i=1}^L \ell_i \mathbf{1}_{\{t_{i-1} \leq |x| \leq t_i\}},$$

where $\mathbf{1}_S$ is the indicator function of a set S . For example, 1-bit quantization with level ℓ corresponds to $\mathcal{Q}_1(x, \ell) = \ell \cdot \text{sign}(x)$. The measurement vector $\mathbf{y} = [y_1, y_2, \dots, y_m]^T$ takes the form

$$y_i = \mathcal{Q}_q(\mathbf{a}_i^T \mathbf{x}_0, \boldsymbol{\ell}, \mathbf{t}), \quad i = 1, 2, \dots, m, \quad (16)$$

where \mathbf{a}_i^T 's are the rows of a measurement matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, which is henceforth assumed i.i.d. standard Gaussian. We use the LASSO to obtain an estimate $\hat{\mathbf{x}}$ of \mathbf{x}_0 as

$$\hat{\mathbf{x}} := \arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{A} \mathbf{x}\|_2 + \lambda \|\mathbf{x}\|_1. \quad (17)$$

Henceforth, we assume for simplicity that $\|\mathbf{x}_0\|_2 = 1$. Also, in our case, μ is known since $g = \mathcal{Q}_q$ is known; thus, is reasonable to scale the solution of (17) as $\mu^{-1} \hat{\mathbf{x}}$ and consider the error quantity $\|\mu^{-1} \hat{\mathbf{x}} - \mathbf{x}_0\|_2$ as a measure of estimation performance. Clearly, the error depends (besides others) on the number of bits q , on the choice of the decision thresholds \mathbf{t} and on the quantization levels $\boldsymbol{\ell}$. An interesting question of practical importance becomes how to optimally choose these to achieve less error. As a running example for this section, we seek optimal quantization thresholds and corresponding levels

$$(\mathbf{t}_*, \boldsymbol{\ell}_*) = \arg \min_{\mathbf{t}, \boldsymbol{\ell}} \|\mu^{-1} \hat{\mathbf{x}} - \mathbf{x}_0\|_2, \quad (18)$$

while keeping all other parameters such as the number of bits q and of measurements m fixed.

2.3.2 Consequences of Precise Error Prediction

Theorem 2.1 shows that $\|\mu^{-1} \hat{\mathbf{x}} - \mathbf{x}_0\|_2 = \|\hat{\mathbf{x}}^{\text{lin}} - \mathbf{x}_0\|_2$, where $\hat{\mathbf{x}}^{\text{lin}}$ is the solution to (17), but only, this time with a measurement vector $\mathbf{y}^{\text{lin}} = \mathbf{A} \mathbf{x}_0 + \frac{\sigma}{\mu} \mathbf{z}$, where μ, σ as in (20) and \mathbf{z} has entries i.i.d. standard normal. Thus, lower values of the ratio σ^2/μ^2 correspond to lower values of the error and the design problem posed in (18) is equivalent to the following simplified one:

$$(\mathbf{t}_*, \boldsymbol{\ell}_*) = \arg \min_{\mathbf{t}, \boldsymbol{\ell}} \frac{\sigma^2(\mathbf{t}, \boldsymbol{\ell})}{\mu^2(\mathbf{t}, \boldsymbol{\ell})}. \quad (19)$$

To be explicit, μ and σ^2 above can be easily expressed from (3) after setting $g = \mathcal{Q}_q$ as follows:

$$\mu := \mu(\boldsymbol{\ell}, \mathbf{t}) = \sqrt{\frac{2}{\pi}} \sum_{i=1}^L \ell_i \cdot \left(e^{-t_{i-1}^2/2} - e^{-t_i^2/2} \right) \quad \text{and} \quad \sigma^2 := \sigma^2(\boldsymbol{\ell}, \mathbf{t}) := \tau^2 - \mu^2, \quad (20)$$

where, $\tau^2 := \tau^2(\boldsymbol{\ell}, \mathbf{t}) = 2 \sum_{i=1}^L \ell_i^2 \cdot (Q(t_{i-1}) - Q(t_i))$ and $Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty \exp(-u^2/2) du$.

2.3.3 An Algorithm for Finding Optimal Quantization Levels and Thresholds

In contrast to the initial problem in (18), the optimization involved in (19) is explicit in terms of the variables ℓ and \mathbf{t} , but, is still hard to solve in general. Interestingly, we show in Appendix C that the popular Lloyd-Max (LM) algorithm can be an effective algorithm for solving (19), since the values to which it converges are stationary points of the objective in (19). Note that this is not a directly obvious result since the classical objective of the LM algorithm is minimizing the quantity $\mathbb{E}[\|\mathbf{y} - \mathbf{A}\mathbf{x}_0\|_2^2]$ rather than $\mathbb{E}[\|\mu^{-1}\hat{\mathbf{x}} - \mathbf{x}_0\|_2^2]$.

References

- [1] Per Kragh Andersen and Richard D Gill. Cox’s regression model for counting processes: a large sample study. *The annals of statistics*, pages 1100–1120, 1982.
- [2] Francis R Bach. Structured sparsity-inducing norms through submodular functions. In *Advances in Neural Information Processing Systems*, pages 118–126, 2010.
- [3] Mohsen Bayati and Andrea Montanari. The lasso risk for gaussian matrices. *Information Theory, IEEE Transactions on*, 58(4):1997–2017, 2012.
- [4] Alexandre Belloni, Victor Chernozhukov, and Lie Wang. Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika*, 98(4):791–806, 2011.
- [5] David R. Brillinger. The identification of a particular nonlinear time series system. *Biometrika*, 64(3):509–515, 1977.
- [6] David R Brillinger. A generalized linear model with” gaussian” regressor variables. *A Festschrift For Erich L. Lehmann*, page 97, 1982.
- [7] Venkat Chandrasekaran, Benjamin Recht, Pablo A Parrilo, and Alan S Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012.
- [8] David L Donoho and Iain M Johnstone. Minimax risk overl p-balls forl p-error. *Probability Theory and Related Fields*, 99(2):277–303, 1994.
- [9] David L Donoho, Iain Johnstone, and Andrea Montanari. Accurate prediction of phase transitions in compressed sensing via a connection to minimax denoising. *IEEE transactions on information theory*, 59(6):3396–3433, 2013.
- [10] David L Donoho, Arian Maleki, and Andrea Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
- [11] David L Donoho, Arian Maleki, and Andrea Montanari. The noise-sensitivity phase transition in compressed sensing. *Information Theory, IEEE Transactions on*, 57(10):6920–6941, 2011.
- [12] Alexandra L Garnham and Luke A Prendergast. A note on least squares sensitivity in single-index model estimation and the benefits of response transformations. *Electronic J. of Statistics*, 7:1983–2004, 2013.
- [13] Yehoram Gordon. Some inequalities for gaussian processes and applications. *Israel Journal of Mathematics*, 50(4):265–289, 1985.
- [14] Yehoram Gordon. *On Milman’s inequality and random subspaces which escape through a mesh in \mathbb{R}^n* . Springer, 1988.
- [15] Marwa El Halabi and Volkan Cevher. A totally unimodular view of structured sparsity. *arXiv preprint arXiv:1411.1990*, 2014.
- [16] Hidehiko Ichimura. Semiparametric least squares (sls) and weighted sls estimation of single-index models. *Journal of Econometrics*, 58(1):71–120, 1993.
- [17] Ker-Chau Li and Naihua Duan. Regression analysis under link violation. *The Annals of Statistics*, pages 1009–1052, 1989.
- [18] Whitney K Newey and Daniel McFadden. Large sample estimation and hypothesis testing. *Handbook of econometrics*, 4:2111–2245, 1994.
- [19] Samet Oymak, Christos Thrampoulidis, and Babak Hassibi. The squared-error of generalized lasso: A precise analysis. *arXiv preprint arXiv:1311.0830*, 2013.
- [20] Yaniv Plan and Roman Vershynin. The generalized lasso with non-linear observations. *arXiv preprint arXiv:1502.04071*, 2015.
- [21] Yaniv Plan, Roman Vershynin, and Elena Yudovina. High-dimensional estimation with geometric constraints. *arXiv preprint arXiv:1404.3749*, 2014.
- [22] R Tyrell Rockafellar. *Convex analysis*, volume 28. Princeton university press, 1997.

- [23] R Tyrrell Rockafellar and Roger J-B Wets. *Variational analysis*, volume 317. Springer Science & Business Media, 2009.
- [24] Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.
- [25] Mihailo Stojnic. A framework to characterize performance of lasso algorithms. *arXiv preprint arXiv:1303.7291*, 2013.
- [26] Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. High-dimensional error analysis of regularized m-estimators.
- [27] Christos Thrampoulidis, Samet Oymak, and Babak Hassibi. A tight version of the gaussian min-max theorem in the presence of convexity. *arXiv preprint arXiv:1408.4837*, 2014.
- [28] Christos Thrampoulidis, Samet Oymak, and Babak Hassibi. Regularized linear regression: A precise analysis of the estimation error. In *Proceedings of The 28th Conference on Learning Theory*, pages 1683–1709, 2015.
- [29] Christos Thrampoulidis, Ashkan Panahi, Daniel Guo, and Babak Hassibi. Precise error analysis of the lasso. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pages 3467–3471.
- [30] Christos Thrampoulidis, Ashkan Panahi, and Babak Hassibi. Asymptotically exact error analysis for the generalized ℓ_2^2 -lasso. In *Information Theory (ISIT), 2015 IEEE International Symposium on*, pages 2021–2025. IEEE, 2015.
- [31] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- [32] Robert Tibshirani, Michael Saunders, Saharon Rosset, Ji Zhu, and Keith Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108, 2005.
- [33] Xinyang Yi, Zhaoran Wang, Constantine Caramanis, and Han Liu. Optimal linear estimation under unknown nonlinear transform. *arXiv preprint arXiv:1505.03257*, 2015.
- [34] Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.

A Proof of Theorems 2.1 & 2.2

A.1 Theorem 2.2

We start with the proof of Theorem 2.2. Theorem 2.1 will follow as a direct corollary of this result.

Assume a sequence of problem instances as described in Section 2.1. To keep notation simple, we simply use $\|\mathbf{v}\|$ (rather than $\|\mathbf{v}\|_2$) for the Euclidean norm of \mathbf{v} and we shall also drop the superscript (n) when referring to elements of the sequence. Thus, we write

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{\sqrt{n}} \|\bar{\mathbf{g}}(\mathbf{A}\mathbf{x}_0) - \mathbf{A}\mathbf{x}\| + \frac{\lambda}{\sqrt{n}} f(\mathbf{x}), \quad (21)$$

but it is to be understood that the above actually produces a sequence of solutions $\hat{\mathbf{x}}^{(n)}$ indexed by n . Our goal is to characterize the nontrivial limiting behavior of $\|\hat{\mathbf{x}} - \mu\mathbf{x}_0\|$.

We start with a simple but useful change of variables $\mathbf{w} := \mathbf{x} - \mu\mathbf{x}_0$, to directly get a handle on the error vector \mathbf{w} . Then, (21) becomes:

$$\begin{aligned} \hat{\mathbf{w}} &:= \arg \min_{\mathbf{w}} \frac{1}{\sqrt{n}} \|\bar{\mathbf{g}}(\mathbf{A}\mathbf{x}_0) - \mu\mathbf{A}\mathbf{x}_0 - \mathbf{A}\mathbf{w}\| + \frac{\lambda}{\sqrt{n}} f(\mu\mathbf{x}_0 + \mathbf{w}) \\ &= \arg \min_{\mathbf{w}} \max_{\|\mathbf{u}\| \leq 1} \frac{1}{\sqrt{n}} (-\mathbf{u}^T \mathbf{A}\mathbf{w} + \mathbf{u}^T (\bar{\mathbf{g}}(\mathbf{A}\mathbf{x}_0) - \mu\mathbf{A}\mathbf{x}_0)) + \frac{\lambda}{\sqrt{n}} f(\mu\mathbf{x}_0 + \mathbf{w}), \end{aligned} \quad (22)$$

where the second line follows after using the fact $\|\mathbf{v}\| = \max_{\|\mathbf{u}\| \leq 1} \mathbf{u}^T \mathbf{v}$.

A.1.1 A Key Decomposition

The first key step in the proof is a trick adapted from the proofs of [20, Lem. 4.3] and [21, Thm. 1.3]. Until further notice, we condition on \mathbf{x}_0 . Also, we repeatedly make use of the assumption that $\|\mathbf{x}_0\| = 1$ without direct reference. The trick amounts to decomposing each measurement vector \mathbf{a}_i in its projection on the direction of \mathbf{x}_0 and its orthogonal complement. Denoting $\mathbf{P}^\perp = (\mathbf{I} - \mathbf{x}_0\mathbf{x}_0^T)$ for the projector onto the orthogonal complement of the span of \mathbf{x}_0 (recall $\|\mathbf{x}_0\|_2 = 1$), we have $\mathbf{a}_i^T = (\mathbf{a}_i^T \mathbf{x}_0) \mathbf{x}_0^T + \mathbf{a}_i^T \mathbf{P}^\perp$, or, in matrix form:

$$\mathbf{A} = (\mathbf{A}\mathbf{x}_0)\mathbf{x}_0^T + \mathbf{A}\mathbf{P}^\perp.$$

Then, (22) becomes:

$$\min_{\mathbf{w}} \max_{\|\mathbf{u}\| \leq 1} \frac{1}{\sqrt{n}} \left\{ -\mathbf{u}^T \mathbf{A}\mathbf{P}^\perp \mathbf{w} + \mathbf{u}^T (\bar{\mathbf{g}}(\mathbf{A}\mathbf{x}_0) - \mu\mathbf{A}\mathbf{x}_0 - (\mathbf{A}\mathbf{x}_0)\mathbf{x}_0^T \mathbf{w}) \right\} + \frac{\lambda}{\sqrt{n}} f(\mu\mathbf{x}_0 + \mathbf{w}). \quad (23)$$

Using the Gaussianity assumption on the entries of \mathbf{A} it is straightforward to show that $\mathbf{P}^\perp \mathbf{a}_i$ is *independent* of $\mathbf{a}_i^T \mathbf{x}_0$, for all $i = 1, \dots, m$. Also, conditioned on $\mathbf{a}_i^T \mathbf{x}_0$, $\mathbf{P}^\perp \mathbf{a}_i$ is independent of $(\bar{g}_i(\mathbf{a}_i^T \mathbf{x}_0) - \mu \mathbf{a}_i^T \mathbf{x}_0)$ since the latter only depends on \mathbf{a}_i through $\mathbf{a}_i^T \mathbf{x}_0$. Combining those, it follows that $\mathbf{P}^\perp \mathbf{a}_i$ is also *independent* of $(\bar{g}_i(\mathbf{a}_i^T \mathbf{x}_0) - \mu \mathbf{a}_i^T \mathbf{x}_0)$, [20, pg. 13]. Overall, $\mathbf{A}\mathbf{P}^\perp \mathbf{w}$ is independent of the rest terms in (23). This shows that the objective function of (23) is distributed identically even after replacing the $\mathbf{A}\mathbf{P}^\perp \mathbf{w}$ with $\mathbf{G}\mathbf{P}^\perp \mathbf{w}$, where \mathbf{G} is an independent copy of \mathbf{A} . After all these, (23) is identically distributed with the following:

$$\min_{\mathbf{w}} \max_{\|\mathbf{u}\| \leq 1} \frac{1}{\sqrt{n}} \left\{ -\mathbf{u}^T \mathbf{G}\mathbf{P}^\perp \mathbf{w} + \mathbf{u}^T (\mathbf{z}_e - (\mathbf{x}_0^T \mathbf{w})\mathbf{e}) \right\} + \frac{\lambda}{\sqrt{n}} f(\mu\mathbf{x}_0 + \mathbf{w}). \quad (24)$$

where \mathbf{G} and $\mathbf{e} := \mathbf{A}\mathbf{x}_0$ have entries i.i.d. standard normal and are independent of each other. Also, $\mathbf{z}_e := \bar{\mathbf{g}}(\mathbf{e}) - \mu\mathbf{e}$ for convenience.

A.1.2 Applying the CGMT

After the decomposition step in the previous section, we have transformed the initial problem to that of analyzing the (probabilistically) equivalent one in (24). In particular, we wish to evaluate the limiting behavior of $\|\hat{\mathbf{w}}\|$, i.e. the norm of the minimizer of the optimization in (24). The

analysis is possible thanks to the Convex Gaussian Min-max Theorem (CGMT) [28, Thm. 3], which is a stronger version of the classical result of Gordon [13] in the presence of additional convexity assumptions. According to the CGMT, the analysis of a Primary Optimization (PO) problem that is of the form

$$\min_{\mathbf{v} \in \mathcal{S}_v} \max_{\mathbf{u} \in \mathcal{S}_u} \mathbf{u}^T \mathbf{G} \mathbf{v} + \psi(\mathbf{v}, \mathbf{u}), \quad (25)$$

with \mathbf{G} being i.i.d. Gaussian, $\mathcal{S}_v, \mathcal{S}_u$ convex, compact sets and ψ a convex-concave function, can be carried out via analyzing a corresponding Auxiliary Optimization problem (AO) defined as

$$\min_{\mathbf{v} \in \mathcal{S}_v} \max_{\mathbf{u} \in \mathcal{S}_u} \|\mathbf{v}\| \mathbf{g}^T \mathbf{u} + \|\mathbf{u}\| \mathbf{h}^T \mathbf{v} + \psi(\mathbf{v}, \mathbf{u}). \quad (26)$$

Here, and onwards, \mathbf{g} and \mathbf{h} are i.i.d. standard Gaussian vectors of appropriate size. In (24) identify the bilinear term $\mathbf{u}^T \mathbf{G} \mathbf{P}^\perp \mathbf{w}$ and note that the rest of the objective function is convex in \mathbf{w} (recall that f is convex), and, linear (thus, concave) in \mathbf{u} . Overall, this is in the appropriate format of a (PO) problem as in (25) modulo the extra factor \mathbf{P}^\perp in the bilinear term. It is straightforward to show that this extra factor only requires a natural change of the corresponding terms in the (AO) problem as follows: $\|\mathbf{P}^\perp \mathbf{w}\| \mathbf{g}^T \mathbf{u} + \|\mathbf{u}\| \mathbf{h}^T \mathbf{P}^\perp \mathbf{w}$. With this minor modification, the CGMT continues to hold. A remaining technical caveat is that the minimization over \mathbf{w} in it appears unconstrained. For this, we assume that the minimizer of (24) satisfies $\|\hat{\mathbf{w}}\| \leq K_w$ for sufficiently large constant $K_w > 0$ independent of n . If our assumption is valid, then by the end of the proof we will have identified a quantity $\alpha_* > 0$ to which $\|\hat{\mathbf{w}}\|$ converges; If α_* turns out to be independent of the choice of K_w , then we may explicitly choose $K_w = 2\alpha_*$ (say) and α_* is the true limit; on the other hand, if α_* turns out to depend on K_w , this means that we could have chosen K_w arbitrarily large in the first place, and so the true limit diverges. Thus, assuming that $\|\hat{\mathbf{w}}\|$ the minimization in (24) is not affected by imposing the constraint $\|\mathbf{w}\| \leq K_w$. With these, we can write the corresponding (AO) problem as

$$\tilde{\mathbf{w}} = \arg \min_{\|\mathbf{w}\| \leq K_w} \max_{\|\mathbf{u}\| \leq 1} \frac{1}{\sqrt{n}} \{ \|\mathbf{P}^\perp \mathbf{w}\| \mathbf{g}^T \mathbf{u} - \|\mathbf{u}\| \mathbf{h}^T \mathbf{P}^\perp \mathbf{w} + \mathbf{u}^T (\mathbf{z}_e - (\mathbf{x}_0^T \mathbf{w}) \mathbf{e}) \} + \frac{\lambda}{\sqrt{n}} f(\mu \mathbf{x}_0 + \mathbf{w}). \quad (27)$$

We will see that analyzing this problem, with the goal of determining the converging value of the magnitude of its minimizer $\tilde{\mathbf{w}}$, is simpler than analyzing the (PO) (and certainly so of the one we started with in (22)). The CGMT essentially shows that $\|\tilde{\mathbf{w}}\|$ converges to the same value as $\|\hat{\mathbf{w}}\|$. Recall $\hat{\mathbf{w}}$ being the minimizer of the (PO) and the goal of Theorem 2.2 being to evaluate the converging value of its magnitude.

A.1.3 Analysis of the Auxiliary Optimization

The goal of this section is that of analyzing the (AO) problem in (27). In particular, we will prove that (i) the optimal cost of the (AO) problem converges to the optimal cost of the deterministic optimization in (10), which involves three scalar optimization variables α, β, τ , (ii) the max-min problem in (10) is *strongly* convex in α and jointly concave in β, τ , (iii) $\|\tilde{\mathbf{w}}\|$ converges to the unique optima α_* in (10). With these, the claim of the Theorem follows by [28, Thm. 1] (also, see [27, Cor. A.1]), as previously discussed.

The analysis requires several steps. The randomness in (27) is over $\mathbf{e}, \mathbf{g}, \mathbf{h}, \mathbf{x}_0$ and possibly the link function g ; at each step we condition on all but a subset of these and identify convergence of the objective function of the (AO) with respect to the remaining. Pointwise convergence (with respect to the involved optimization variables) needs to be turned into uniform convergence to guarantee that not only the objective function, but also the min/max value and the optimizer converge appropriately. (Strong) convexity of the objective will turn out to be crucial for the latter.

Introducing the Fenchel conjugate. To begin with, let us rewrite the (AO) problem above by expressing f in terms of its Fenchel conjugate, i.e.

$$f(\mathbf{x}) = \sup_{\bar{\mathbf{v}}} \bar{\mathbf{v}}^T \mathbf{x} - f^*(\bar{\mathbf{v}}) = \sup_{\bar{\mathbf{v}}} \sqrt{n} \mathbf{v}^T \mathbf{x} - f^*(\sqrt{n} \mathbf{v}). \quad (28)$$

Translating to our problem and after rescaling this gives,

$$n^{-1/2} f(\mu \mathbf{x}_0 + \mathbf{w}) = \sup_{\mathbf{v}} \mathbf{v}^T (\mu \mathbf{x}_0 + \mathbf{w}) - n^{-1/2} f^*(\sqrt{n} \mathbf{v}). \quad (29)$$

Now, from standard optimality conditions of (28), the optimal $\bar{\mathbf{v}}_*$ satisfies $\bar{\mathbf{v}}_* \in \partial f(\mathbf{x})$. Then, using condition (b) of Section 2.1, $\|\mathbf{v}_*\| = \mathcal{O}(\sqrt{n})$ for all \mathbf{x} such that $\|\mathbf{x}\| = \mathcal{O}(1)$. From this, and $\|\mathbf{w} + \mu\mathbf{x}_0\| = \mathcal{O}(1)$ we conclude that the optimal \mathbf{v}_* in (29) satisfies $\|\mathbf{v}_*\| \leq K_v < 0$ for sufficiently large constant K_v independent of n . Putting everything together, (27) is equivalent to

$$\min_{\|\mathbf{w}\| \leq K_w} \max_{\substack{\|\mathbf{u}\| \leq 1 \\ 0 \leq \|\mathbf{v}\| \leq K_v}} \frac{1}{\sqrt{n}} \mathbf{u}^T (\mathbf{z}_e - (\mathbf{x}_0^T \mathbf{w}) \mathbf{e}) - \|\mathbf{P}^\perp \mathbf{w}\| \mathbf{g} - \|\mathbf{u}\| \bar{\mathbf{h}}^T \mathbf{P}^\perp \mathbf{w} + \lambda \mathbf{v}^T (\mu \mathbf{x}_0 + \mathbf{w}) - \lambda \bar{f}^*(\mathbf{v}), \quad (30)$$

where we have also denoted $\bar{\mathbf{h}} := n^{-1/2} \mathbf{h}$ and $\bar{f}^*(\mathbf{v}) = n^{-1/2} f^*(\sqrt{n} \mathbf{v})$. Observe again that by condition (b) of Section 2.1, $\bar{f}^*(\mathbf{v}) = \max_{\mathbf{x}} \mathbf{x}^T \mathbf{v} - n^{-1/2} f(\mathbf{x}) = \mathcal{O}(1)$ since $\mathbf{v} = \mathcal{O}(1)$.

In order to somewhat simplify the exposition, we often omit explicitly carrying over the constraints $\|\mathbf{w}\| \leq K_w$, $\|\mathbf{v}\| \leq K_v$ until the very last step, but we often recall and actually make use of them.

Optimizing over the directions of \mathbf{u} and \mathbf{w} . Observe that the maximization over the direction of \mathbf{u} is easy in (30), which then becomes:

$$\min_{\mathbf{w}} \max_{0 \leq \beta \leq 1} \frac{1}{\sqrt{n}} \beta \|\mathbf{z}_e - (\mathbf{x}_0^T \mathbf{w}) \mathbf{e} - \|\mathbf{P}^\perp \mathbf{w}\| \mathbf{g}\| - \beta \bar{\mathbf{h}}^T \mathbf{P}^\perp \mathbf{w} + \lambda \mathbf{v}^T (\mu \mathbf{x}_0 + \mathbf{w}) - \lambda \bar{f}^*(\mathbf{v}). \quad (31)$$

At this point, the form of the objective function suggests that it is possible to do the same trick over \mathbf{w} , i.e. fix its magnitude and optimize over only its direction. The caveat is that the minimization over \mathbf{w} in (31) is done only after the maximization over β and \mathbf{v} . What is more, the objective function is not be convex in \mathbf{w} ; thus, flipping the order of min-max operations that would resolve the issue is not directly justified by (say) Sion's minimax theorem [24].

However, [26] shows that the flipping is possible when dimensions are large. To provide some intuition on why this should be true, observe that the (PO) problem in (24) is itself convex and satisfies all conditions of Sion's minimax theorem; thus the order of min-max operations can be flipped. According to the CGMT, the (PO) and the (AO) are tightly related in an asymptotic setting. [26] uses this to translate the convexity properties of the (PO) to the (AO).⁶ Hence, we have:

$$\begin{aligned} & \max_{0 \leq \beta \leq 1} \min_{\mathbf{w}} \frac{\beta}{\sqrt{n}} \|\mathbf{z}_e + (\mathbf{x}_0^T \mathbf{w}) \mathbf{e} - \|\mathbf{P}^\perp \mathbf{w}\| \mathbf{g}\| - \beta \bar{\mathbf{h}}^T \mathbf{P}^\perp \mathbf{w} + \lambda \mathbf{v}^T (\mu \mathbf{x}_0 + \mathbf{w}) - \lambda \bar{f}^*(\mathbf{v}) \\ &= \max_{0 \leq \beta \leq 1} \min_{\substack{\alpha_1, \alpha_2 \geq 0 \\ \mathbf{v}}} \frac{\beta}{\sqrt{n}} \|\mathbf{z}_e + \alpha_2 \mathbf{e} - \alpha_1 \mathbf{g}\| - \max_{\substack{\|\mathbf{P}^\perp \mathbf{w}\| = \alpha_1 \\ \mathbf{x}_0^T \mathbf{w} = \alpha_2}} \{\beta \bar{\mathbf{h}}^T \mathbf{P}^\perp \mathbf{w} - \lambda \mathbf{v}^T (\mu \mathbf{x}_0 + \mathbf{w}) + \lambda \bar{f}^*(\mathbf{v})\}. \end{aligned}$$

By decomposing \mathbf{w} as $\mathbf{P}^\perp \mathbf{w} + (\mathbf{x}_0^T \mathbf{w}) \mathbf{x}_0$, it is not hard to perform the maximization over \mathbf{w} to equivalently write the last display above as:

$$\max_{0 \leq \beta \leq 1} \min_{\substack{\alpha_1, \alpha_2 \geq 0 \\ \mathbf{v}}} \frac{\beta}{\sqrt{n}} \|\mathbf{z}_e + \alpha_2 \mathbf{e} - \alpha_1 \mathbf{g}\| - \alpha_1 \|\beta \mathbf{P}^\perp \bar{\mathbf{h}} - \lambda \mathbf{P}^\perp \mathbf{v}\| + \lambda \mu \mathbf{v}^T \mathbf{x}_0 + \alpha_2 \lambda (\mathbf{v}^T \mathbf{x}_0) - \lambda \bar{f}^*(\mathbf{v}). \quad (32)$$

The randomness of \mathbf{e} , \mathbf{g} and g . Until further notice condition on $\bar{\mathbf{h}}$ and \mathbf{x}_0 . All randomness in (32) is now on the first term.

Consider β, \mathbf{v} fixed for now. For any pair α_1, α_2 by the WLLN, $m^{-1} \|\mathbf{z}_e + \alpha_2 \mathbf{e} - \alpha_1 \mathbf{g}\|^2 \xrightarrow{P} \mathbb{E}[(g(\gamma) - \mu\gamma + \alpha_2\gamma - \alpha_1\gamma')^2]$, where $\gamma, \gamma' \sim \mathcal{N}(0, 1)$ and independent. Recall, $\mathbb{E}[(g(\gamma) - \mu\gamma)^2] = \sigma^2$, $\mathbb{E}[(g(\gamma) - \mu\gamma)\gamma] = \mu - \mu = 0$ and $m/n = \delta$, to conclude that $n^{-1/2} \|\mathbf{z}_e + \alpha_2 \mathbf{e} - \alpha_1 \mathbf{g}\| \xrightarrow{P} \sqrt{\delta} \sqrt{\sigma^2 + \alpha_1^2 + \alpha_2^2}$, where convergence is point-wise in α_1, α_2 . Also, the objective function in (32)

⁶To be more precise, it is shown in [26] that the (AO) problem as it appears in (26) can be slightly modified to the following

$$\max_{\beta \geq 0} \min_{\mathbf{v} \in \mathcal{S}_{\mathbf{v}}} \max_{\substack{\mathbf{u} \in \mathcal{S}_{\mathbf{u}} \\ \|\mathbf{u}\|_2 = \beta}} \|\mathbf{v}\| \mathbf{g}^T \mathbf{u} + \|\mathbf{u}\| \bar{\mathbf{h}}^T \mathbf{v} + \psi(\mathbf{v}, \mathbf{u}),$$

without changing the conclusions of the CGMT. Please refer to [26] for a proof.

is jointly convex in $[\alpha_1, \alpha_2]$. Thus, point-wise convergence translates to uniform as in [1, Cor. II.1], from which, it follows that (for any β, \mathbf{v}) the minimum over α_1, α_2 in (32) converges to

$$\min_{\alpha_1, \alpha_2 \geq 0} \beta \sqrt{\delta} \sqrt{\sigma^2 + \alpha_1^2 + \alpha_2^2} - \alpha_1 \|\beta \mathbf{P}^\perp \bar{\mathbf{h}} - \lambda \mathbf{P}^\perp \mathbf{v}\| + \lambda \mu \mathbf{v}^T \mathbf{x}_0 + \alpha_2 \lambda (\mathbf{v}^T \mathbf{x}_0) - \lambda \bar{f}^*(\mathbf{v}). \quad (33)$$

Furthermore, the function $\sqrt{\sigma^2 + \alpha_1^2 + \alpha_2^2}$ is (by direct differentiation) jointly strongly convex over $[\alpha_1, \alpha_2]$; thus (33) has a unique minimizer. Then, we can apply the Argmin theorem [18, Thm. 2.7] to conclude that the optimal α_1, α_2 of (32) converge to the corresponding (unique) optima of (33).

Up to now, β, \mathbf{v} were assumed fixed and the convergence from (32) to (33) holds point-wise with respect to β, \mathbf{v} . The point-wise minimum of concave functions is still concave, thus, uniform convergence is indeed true by [18, Thm. 2.7]. Hence, (32) converges to

$$\max_{0 \leq \beta \leq 1} \min_{\alpha_1, \alpha_2 \geq 0} \beta \sqrt{\delta} \sqrt{\sigma^2 + \alpha_1^2 + \alpha_2^2} - \alpha_1 \|\beta \mathbf{P}^\perp \bar{\mathbf{h}} - \lambda \mathbf{P}^\perp \mathbf{v}\| + \lambda \mu \mathbf{v}^T \mathbf{x}_0 + \alpha_2 \lambda (\mathbf{v}^T \mathbf{x}_0) - \lambda \bar{f}^*(\mathbf{v}), \quad (34)$$

and the optimal α_1, α_2 of the former converge to the corresponding optima of the latter.

Merging α_1 and α_2 . It is important to note that $\alpha_1^2 + \alpha_2^2$ in (34) correspond exactly to the squared norm of the error. Here, we simplify (34) by introducing the quantity $\alpha_1^2 + \alpha_2^2$ as the minimization variable rather than sperately α_1 and α_2 . By first order optimality conditions in (34) we find

$$\alpha_1 \beta \sqrt{\delta} = \|\beta \mathbf{P}^\perp \bar{\mathbf{h}} - \lambda \mathbf{P}^\perp \mathbf{v}\| \sqrt{\alpha_1^2 + \alpha_2^2 + \sigma^2} \quad \text{and} \quad -\alpha_2 \beta \sqrt{\delta} = \lambda \mathbf{v}^T \mathbf{x}_0 \sqrt{\alpha_1^2 + \alpha_2^2 + \sigma^2}. \quad (35)$$

Substituting this in (34), the objective becomes (ignoring the terms that do not involve α_1 or α_2):

$$\beta \sqrt{\delta} \sqrt{\sigma^2 + \alpha_1^2 + \alpha_2^2} - \frac{\sqrt{\sigma^2 + \alpha_1^2 + \alpha_2^2}}{\beta \sqrt{\delta}} (\|\beta \mathbf{P}^\perp \bar{\mathbf{h}} - \lambda \mathbf{P}^\perp \mathbf{v}\|^2 + (\lambda \mathbf{v}^T \mathbf{x}_0)^2)$$

But, from (35) we find $\sqrt{\sigma^2 + \alpha_1^2 + \alpha_2^2} \sqrt{\|\beta \mathbf{P}^\perp \bar{\mathbf{h}} - \lambda \mathbf{P}^\perp \mathbf{v}\|^2 + (\lambda \mathbf{v}^T \mathbf{x}_0)^2} = \beta \sqrt{\delta} \sqrt{\alpha_1^2 + \alpha_2^2}$. Combining, we conclude that (34) can be written as

$$\max_{0 \leq \beta \leq 1} \min_{\alpha \geq 0} \beta \sqrt{\delta} \sqrt{\sigma^2 + \alpha^2} - \alpha \|\beta \mathbf{P}^\perp \bar{\mathbf{h}} - \lambda \mathbf{v}\| + \lambda \mu \mathbf{v}^T \mathbf{x}_0 - \lambda \bar{f}^*(\mathbf{v}), \quad (36)$$

where the new optimization variable α plays the role of $\sqrt{\alpha_1^2 + \alpha_2^2}$, thus it represents the norm of the error vector $\|\mathbf{w}\|$. We have also identified $\|\beta \mathbf{P}^\perp \bar{\mathbf{h}} - \lambda \mathbf{P}^\perp \mathbf{v}\|^2 + (\lambda \mathbf{v}^T \mathbf{x}_0)^2 = \|\beta \mathbf{P}^\perp \bar{\mathbf{h}} - \lambda \mathbf{v}\|^2$.

Introducing a new optimization variable. To get a better handle at it, we square the norm term in (36) at the expense of introducing a new scalar optimization variable. This is based on the following trick:

$$\sqrt{x} = \min_{\tau > 0} \frac{\tau}{2} + \frac{x}{2\tau}, \quad (37)$$

for any $x \geq 0$. Thus, (36) becomes

$$\max_{0 \leq \beta \leq 1} \min_{\alpha \geq 0} \beta \sqrt{\delta} \sqrt{\sigma^2 + \alpha^2} - \frac{\alpha \tau}{2} - \frac{\alpha}{2\tau} \|\beta \mathbf{P}^\perp \bar{\mathbf{h}} - \lambda \mathbf{v}\|^2 + \lambda \mu \mathbf{v}^T \mathbf{x}_0 - \lambda \bar{f}^*(\mathbf{v}), \quad (38)$$

where we have also flipped the order of min-max between α and τ . We could do this as in [22, Cor. 37.3.2] since the objective is convex in α and concave in τ , the constraint sets are both convex and both of them are bounded. To argue the boundedness, recall that $\alpha \leq K_{\mathbf{w}}$; for τ it suffices to combine optimality conditions of (37) and boundedness of \mathbf{v} , $\|\mathbf{v}\|_2 \leq K_{\mathbf{v}}$.

Optimizing over \mathbf{v} . Note that the objective in (38) is concave in \mathbf{v} , convex in α and the constraint sets are convex compact. Thus, as it might be expected by now, we use [22, Cor. 37.3.2] to flip the corresponding order of max-min. Also, after some simple algebra while using $\mathbf{P}^\perp \mathbf{x}_0 = 0$ and $\|\mathbf{x}_0\| = 1$, it can be shown that

$$\|\beta \mathbf{P}^\perp \bar{\mathbf{h}} - \lambda \mathbf{v}\|^2 - 2 \frac{\tau}{\alpha} \lambda \mu \mathbf{v}^T \mathbf{x}_0 = \|\lambda \mathbf{v} - (\beta \mathbf{P}^\perp \bar{\mathbf{h}} + \frac{\tau}{\alpha} \mu \mathbf{x}_0)\|^2 - \mu^2 \frac{\tau^2}{\alpha^2}.$$

Combining, we conclude with

$$\begin{aligned}
(38) &= \max_{\substack{0 \leq \beta \leq 1 \\ \tau > 0}} \min_{\alpha \geq 0} \beta \sqrt{\delta} \sqrt{\sigma^2 + \alpha^2} - \frac{\alpha \tau}{2} + \mu^2 \frac{\tau}{2\alpha} - \frac{\alpha \lambda^2}{\tau} \min_{\mathbf{v}} \left\{ \frac{1}{2} \left\| \mathbf{v} - \left(\frac{\beta}{\lambda} \mathbf{P}^\perp \bar{\mathbf{h}} + \frac{\tau}{\alpha \lambda} \mu \mathbf{x}_0 \right) \right\|^2 + \frac{\tau}{\lambda \alpha} \bar{f}^*(\mathbf{v}) \right\}, \\
&= \max_{\substack{0 \leq \beta \leq 1 \\ \tau > 0}} \min_{\alpha \geq 0} G(\alpha, \beta, \tau).
\end{aligned} \tag{39}$$

Here, $G(\alpha, \beta, \tau)$ is convex in α (see (36)) and jointly concave in β, τ . To see the latter it suffices to show that $\frac{\alpha \lambda^2}{\tau} \left\| \mathbf{v} - \frac{\beta}{\lambda} (\mathbf{P}^\perp \bar{\mathbf{h}} + \frac{\mu \tau}{\lambda \alpha} \mathbf{x}_0) \right\|^2$ is jointly convex over β, τ, \mathbf{v} (minimization over \mathbf{v} does not change the joint convexity over τ and β). Norm is separable over its entries, so we equivalently show that for scalars τ, β, v , the function $\frac{1}{\tau} (v - c_1 \beta - c_2 \tau)^2$ is jointly convex over $\tau > 0, \beta$; this is true as the perspective function of $(v - c_1 \beta - c_2 \tau)^2$. One more remark is in place here regarding the form of $G(\alpha, \beta, \tau)$: even-though α appears in the denominator in (39), the limit of $\alpha \rightarrow 0$ of the expression is finite using the continuity of the Moreau envelope [23]. Another way to see this is by noting that the objective in (39) is equivalent to that in (36). Hence, evaluating G at $\alpha = 0$ in the minimization in (39) subsumes $G(0, \beta, \tau) = \lim_{\alpha \rightarrow 0} G(\alpha, \beta, \tau)$. (Of course, this is also the case with the optimization problem (10) of Theorem 2.2.)

The randomness of $\bar{\mathbf{h}}$ and \mathbf{x}_0 . Fix β, τ, α , denote $c_1 = \frac{\beta}{\lambda}, c_2 = \frac{\tau \mu}{\alpha \lambda}, c_3 = \frac{\tau}{\alpha \lambda}$, and, consider

$$R(\bar{\mathbf{h}}, \mathbf{x}_0) := R(\alpha, \beta, \tau; \bar{\mathbf{h}}, \mathbf{x}_0) := -\frac{c_2}{2} + \frac{1}{c_2} \min_{\mathbf{v}} \left\{ \frac{1}{2} \left\| \mathbf{v} - c_1 \mathbf{P}^\perp \bar{\mathbf{h}} - c_2 \mathbf{x}_0 \right\|^2 + c_3 \bar{f}^*(\mathbf{v}) \right\}.$$

Recall from Assumption 1, and, from the modeling condition (a) in Section 2.1 that

$$A(\bar{\mathbf{h}}, \mathbf{x}_0) := R(\alpha, \beta, \tau; \bar{\mathbf{h}}, \mathbf{x}_0) := -\frac{c_2}{2} \frac{\|\bar{\mathbf{x}}_0\|^2}{n} + \frac{1}{c_2} \min_{\mathbf{v}} \left\{ \frac{1}{2} \left\| \mathbf{v} - c_1 \bar{\mathbf{h}} - c_2 \frac{\bar{\mathbf{x}}_0}{\sqrt{n}} \right\|^2 + c_3 \bar{f}^*(\mathbf{v}) \right\} \tag{40}$$

converges to $\{-\frac{c_2}{2} + \frac{\mu}{c_2} F(c_1, c_2, c_3)\}$ in probability. Also, recall $\bar{\mathbf{x}}_0 = \mathbf{x}_0 \|\bar{\mathbf{x}}_0\|$. Next, we show that for all constant $\zeta > 0$

$$|R(\bar{\mathbf{h}}, \mathbf{x}_0) - A(\bar{\mathbf{h}}, \mathbf{x}_0)| \leq \zeta \tag{41}$$

with probability approaching one in the limit of $n \rightarrow \infty$. Combining this with Assumption 1, will prove that $R(\bar{\mathbf{h}}, \mathbf{x}_0)$ converges in $\{-\frac{c_2}{2} + \frac{\mu}{c_2} F(c_1, c_2, c_3)\}$ in probability, as well.

Proof of (41): Fix any $\epsilon > 0$. We condition on the following events:

$$\begin{cases} |\bar{\mathbf{h}}^T \mathbf{x}_0| \leq \epsilon, \\ 1 - \epsilon \leq n^{-1/2} \|\bar{\mathbf{x}}_0\| \leq 1 + \epsilon. \end{cases} \tag{42}$$

Each one of the events occurs with probability approaching one as $n \rightarrow \infty$; the first follows since $\bar{\mathbf{h}} \sim \mathcal{N}(0, \frac{1}{n} \mathbf{I}_n)$ and $\|\mathbf{x}_0\| = 1$ and from standard tail bounds on Gaussians; the second is due to condition (a) of Section 2.1. Without loss of generality assume $R(\bar{\mathbf{h}}, \mathbf{x}_0) \geq A(\bar{\mathbf{h}}, \mathbf{x}_0)$, and let \mathbf{v}_* be optimal in (40), then

$$\begin{aligned}
|R(\bar{\mathbf{h}}, \mathbf{x}_0) - A(\bar{\mathbf{h}}, \mathbf{x}_0)| &\leq \frac{c_2}{2} \left(\frac{\|\bar{\mathbf{x}}_0\|^2}{n} - 1 \right) + \frac{1}{2c_2} \left\| \mathbf{v}_* - c_1 \mathbf{P}^\perp \bar{\mathbf{h}} - c_2 \mathbf{x}_0 \right\|^2 - \frac{1}{2c_2} \left\| \mathbf{v}_* - c_1 \bar{\mathbf{h}} - c_2 \frac{\bar{\mathbf{x}}_0}{\sqrt{n}} \right\|^2 \\
&= \frac{c_2}{2} \left(\frac{\|\bar{\mathbf{x}}_0\|^2}{n} - 1 \right) \\
&\quad + \left(\frac{c_1}{c_2} (\mathbf{x}_0^T \bar{\mathbf{h}}) \mathbf{x}_0 + \bar{\mathbf{x}}_0 \left(\frac{1}{\sqrt{n}} - \frac{1}{\|\bar{\mathbf{x}}_0\|} \right) \right)^T \left(\mathbf{v}_* - c_1 \bar{\mathbf{h}} - \frac{1}{2} c_2 \bar{\mathbf{x}}_0 \left(\frac{1}{\sqrt{n}} + \frac{1}{\|\bar{\mathbf{x}}_0\|} \right) + \frac{1}{2} c_1 (\mathbf{x}_0^T \bar{\mathbf{h}}) \mathbf{x}_0 \right) \\
&= -\frac{1}{2} \frac{c_1^2}{c_2} (\mathbf{x}_0^T \bar{\mathbf{h}})^2 + \frac{c_1}{c_2} (\mathbf{x}_0^T \bar{\mathbf{h}}) (\mathbf{x}_0^T \mathbf{v}_*) - c_1 (\mathbf{x}_0^T \bar{\mathbf{h}}) \frac{\|\bar{\mathbf{x}}_0\|}{\sqrt{n}} + (\mathbf{x}_0^T \mathbf{v}_*) \left(\frac{\|\bar{\mathbf{x}}_0\|}{\sqrt{n}} - 1 \right) \\
&\leq \frac{1}{2} \frac{c_1^2}{c_2} \epsilon^2 + \frac{c_1}{c_2} \|\mathbf{v}_*\| \epsilon + c_1 \epsilon (1 + \epsilon) + \|\mathbf{v}_*\| \epsilon
\end{aligned} \tag{43}$$

where the last line follows after bounding the absolute values of the summands using (42). Recall now that $\|\mathbf{v}_*\| \leq K_{\mathbf{v}} < \infty$. Also, note that c_1 and $\frac{c_1}{c_2}$ are also bounded constants. Then, for all $\zeta > 0$ in (41) we can find sufficiently small $\epsilon > 0$ such that the value of the last expression in the panel above is no larger than ζ , thus completing the proof of (41).

Thus, we have shown that $G(\alpha, \beta, \tau)$ in (39) converges pointwise to

$$H(\alpha, \beta, \tau) := \sqrt{\delta} \sqrt{\sigma^2 + \alpha^2} - \frac{\alpha p}{2} + \mu^2 \frac{\tau}{2\alpha} - \frac{\alpha \lambda^2}{\tau} F\left(\frac{\beta}{\lambda}, \frac{\tau \mu}{\alpha \lambda}, \frac{\tau}{\alpha \lambda}\right),$$

in the limit of $n \rightarrow \infty$. Note that H is strongly convex in α and jointly concave in β, τ since taking limits does not affect convexity properties (recall that G is convex-concave). With these, it follows as per [18, Thm. 2.7] that (i)

$$\max_{0 \leq \beta \leq 1, \tau > 0} \min_{\alpha \geq 0} G(\alpha, \beta, \tau) \xrightarrow{P} \max_{0 \leq \beta \leq 1, \tau > 0} \min_{\alpha \geq 0} H(\alpha, \beta, \tau), \quad (44)$$

and, (ii) $\alpha_*(\mathbf{h}, \mathbf{x}_0) \xrightarrow{P} \alpha_*$, where α_* the unique minimizer of the second optimization in (44). This completes the proof of the theorem.

A.2 Theorem 2.1

The theorem is a direct consequence of Theorem 2.2. In particular, Theorem 2.2 proves that the value α_* to which the error converges only depends on g through the parameters μ and σ^2 . Those are the same (by definition) for the non-linear and the linear case considered. Therefore, the errors are the same.

B Proof of Theorem 2.3

Specializing Theorem 2.2 to the setup of Section 2.2.2 we showed in the same section that $\|\hat{\mathbf{x}} - \mu \mathbf{x}_0\|$ converges in probability to the unique minimizer α_* of the following max-min problem:

$$\max_{\substack{0 \leq \beta \leq 1 \\ \tau > 0}} \min_{\alpha \geq 0} H(\alpha, \beta, \tau) := \beta \sqrt{\delta} \sqrt{\sigma^2 + \alpha^2} - \frac{\tau \alpha}{2} + \frac{\tau \mu^2}{2\alpha} - \frac{\alpha}{2\tau} \mathbb{E} \left[\eta^2 \left(\beta h + \frac{\mu \tau}{\alpha} \bar{X}_0; \lambda \right) \right], \quad (45)$$

where the expectation is over $h \sim \mathcal{N}(0, 1)$ and $\bar{X}_0 \sim p_{\bar{X}_0}$. Here, we prove Theorem 2.3 by analyzing the optimality conditions of (45). Recall as in the proof of Theorem 2.2 that H is jointly concave in β, τ and strongly convex in α .

B.1 First Order Optimality Conditions.

We begin with a lemma, which characterizes the first-order optimality conditions of (45).

Lemma B.1 (Optimality Conditions). *Consider the following pair of equations with respect to β and κ :*

$$\begin{cases} \beta^2 \kappa^2 \delta = \sigma^2 + \mathbb{E} [(\eta(\beta \kappa h + \mu \bar{X}_0; \kappa \lambda) - \mu \bar{X}_0)^2], \\ \beta \kappa \delta = \mathbb{E}[(\eta(\beta \kappa h + \mu \bar{X}_0; \kappa \lambda) \cdot h)]. \end{cases} \quad (46)$$

Also, define λ_{\min} to be the unique non-negative solution to the equation

$$(1 + x^2) \int_{-\infty}^{-x} e^{-z^2/2} dz - x e^{-x^2/2} = \delta \sqrt{\frac{\pi}{2}}.$$

With these, let $(\beta_*, \tau_*, \alpha_*)$ be optimal in (45). Then,

$$\alpha_*^2 = \beta_*^2 \kappa_*^2 \delta - \sigma^2 \quad \text{and} \quad \kappa_* = \frac{\sigma}{\sqrt{\beta_*^2 \delta - \tau_*^2}}. \quad (48)$$

such that,

- (i) If $\beta_* = 1$ and $\lambda > \lambda_{\min}$, then κ_* is the unique solution to (46) for $\beta = 1$,

(ii) If $\beta_* \in (0, 1)$, then κ_*, β_* are solutions to the pair of equation (46)-(47).

Proof. Consider the derivation of the objective function with respect to α , τ and β as follows

$$\frac{\partial}{\partial \alpha} H(\alpha, \tau, \beta) = \frac{\beta \alpha \sqrt{\delta}}{\sqrt{\alpha^2 + \sigma^2}} - \frac{\tau}{2} - \frac{1}{2\tau} \mathbb{E}[(\eta(\beta h + \frac{\mu\tau}{\alpha} \bar{X}_0, \lambda) - \frac{\mu\tau}{\alpha} \bar{X}_0)^2] \quad (49)$$

$$\frac{\partial}{\partial \tau} H(\alpha, \tau, \beta) = -\frac{\alpha}{2} + \frac{\alpha}{2\tau^2} \mathbb{E}[(\eta(\beta h + \frac{\mu\tau}{\alpha} \bar{X}_0, \lambda) - \frac{\mu\tau}{\alpha} \bar{X}_0)^2] \quad (50)$$

$$\frac{\partial}{\partial \beta} H(\alpha, \tau, \beta) = \sqrt{\delta} \sqrt{\alpha^2 + \sigma^2} - \frac{\alpha}{\tau} \mathbb{E}[\eta(\beta h + \frac{\mu\tau}{\alpha} \bar{X}_0, \lambda) \cdot h] \quad (51)$$

We will prove that the optimal α_*, τ_* and β_* are all strictly positive. First, suppose $\alpha_* = 0$ and $\tau_* > 0$. Then, the first term in (49) goes to zero while $-\frac{\tau}{2}$ stays negative and the final term is always non-positive. This shows $\frac{\partial}{\partial \alpha} H(\alpha, \tau, \beta) < 0$, which means that $\alpha_* = 0$ cannot be optimal in this case by convexity. Next, assume $\alpha_* = 0$ and the optimal over τ is approached as $\tau \rightarrow 0$. In this case, it can be shown that the expected value in (49) is strictly positive, thus the derivative remains negative. Thus, $\alpha_* > 0$, as promised. A similar argument shows strict positivity of (50) when $\tau \rightarrow 0$. Thus, $\tau_* > 0$. Finally, $\frac{\partial}{\partial \beta} H(\alpha, \tau, \beta) > 0$, by independence of h and \bar{X}_0 , showing $\beta_* > 0$.

The argument above shows that the derivatives are equal to zero at the optimal. For convenience define

$$P\left(\frac{\tau}{\alpha}\right) := \frac{\tau \mu^2}{2\alpha} - \frac{\alpha \lambda^2}{2\tau} \mathbb{E}\left[\eta^2\left(\frac{\beta}{\lambda} h + \frac{\mu\tau}{\lambda\alpha} \bar{X}_0; 1\right)\right].$$

Then equating the derivatives in (45) with respect to α and τ with zero gives

$$\frac{\beta \sqrt{\delta} \alpha}{\sqrt{\alpha^2 + \sigma^2}} - \frac{\tau}{2} - \frac{\tau}{\alpha^2} P'\left(\frac{\tau}{\alpha}\right) = 0, \quad (52a)$$

$$-\frac{\alpha}{2} + \frac{1}{\alpha} P'\left(\frac{\tau}{\alpha}\right) = 0. \quad (52b)$$

Here, P' is the derivative of $P(x)$ with respect to x . Any optimal $\beta_*, \tau_*, \alpha_*$ satisfies these. Then, it only takes multiplying (52b) by $\frac{\tau}{\alpha}$ and adding the result to (52a) to see that

$$\alpha_* = \frac{\tau_* \sigma}{\sqrt{\beta_*^2 \delta - \tau_*^2}}. \quad (53)$$

Next, substituting (53) in (52b) it can be shown that,

$$-\frac{\sigma^2}{2} + \frac{\sigma^2}{2\tau_*^2} \mathbb{E}[(\eta(\beta h + \frac{\sqrt{\beta_*^2 \delta - \tau_*^2}}{\sigma} \mu \bar{X}_0, \lambda) - \frac{\sqrt{\beta_*^2 \delta - \tau_*^2}}{\sigma} \mu \bar{X}_0)^2] = 0.$$

To reach this we have also used the following facts: $\eta(x; \lambda) \frac{\partial}{\partial x} \eta(x; \lambda) = \eta(x; \lambda)$, $\lambda \eta(\frac{x}{\lambda}; 1) = \eta(x; \lambda)$ and $\mathbb{E}[\bar{X}_0^2] = 1$ by assumption. Multiplying the result with $2\tau_*^2/\sigma^2$ and defining

$$\kappa := \frac{\sigma}{\sqrt{\beta_*^2 \delta - \tau_*^2}},$$

we conclude with,

$$\beta_*^2 \delta \kappa^2 - \sigma^2 = \mathbb{E}[(\eta(\beta \kappa h + \mu \bar{X}_0; \kappa \lambda) - \mu \bar{X}_0)^2], \quad (54)$$

which is same as (46). Also, with respect to the optimal κ_* it is easily seen by (53) that

$$\alpha_*^2 = \beta_*^2 \kappa_*^2 \delta - \sigma^2. \quad (55)$$

The derivative in (45) with respect to β gives

$$\begin{aligned} \frac{\partial}{\partial \beta} H(\alpha, \beta, \tau) &= \sqrt{\delta} \sqrt{\sigma^2 + \alpha^2} - \frac{\alpha}{\tau} \mathbb{E}[\eta(\beta h + \frac{\mu\tau}{\alpha} \bar{X}_0; \lambda) h] \\ &= \beta \delta \kappa - \kappa \mathbb{E}[\eta(\beta h + \frac{\mu \bar{X}_0}{\kappa}, \lambda) h] = \beta \delta \kappa - \mathbb{E}[\eta(\kappa \beta h + \mu \bar{X}_0; \lambda \kappa) h]. \end{aligned} \quad (56)$$

where we have also used (55). Note that the above is same as (47) and recall the constraint $0 \leq \beta \leq 1$ in (45) to conclude with the desired.

It only remains to show that the solution with respect to κ of (46) (eqv. of (54)) is unique when $\beta = 1$ and $\lambda \geq \lambda_{\min}$. For $\beta = 1$, (46) is the same as fixed point equation [3, Eqn. (1.9)], which in turn was shown to admit a unique solution for all $\lambda > \lambda_{\min}$ in [11] (see [3, Prop. 1.3]). \square

B.2 The Regions of Operation

We build up to the proof of Theorem 2.3 through a series of auxiliary lemmas. Through the lemmas, we identify two “regimes of operation” of the LASSO. The first, we call \mathcal{R}_{bad} , and it corresponds to values of λ for which the optimal β is in the open set $(0, 1)$. The second regime, is such that $\beta = 1$. If $\delta < 1$, we prove in Lemma B.5 that there exists a unique critical value λ_{crit} separating the two regimes in the sense that \mathcal{R}_{bad} extends from 0 to λ_{crit} . If on the other hand $\delta \geq 1$, then there is no \mathcal{R}_{bad} region (Lemma B.6).

First, we need a few useful definitions.

Definition B.1. For any $\lambda > 0$, we let $\alpha_*(\lambda)$, $\tau_*(\lambda)$ and $\beta_*(\lambda)$ be optimal solutions in (45). Apart from $\alpha_*(\lambda)$, the others are not necessarily unique at this point. Also, $\kappa_*(\lambda)$ is defined as in (48).

Definition B.2 (Bad Regime). We say that a value $\lambda > 0$ is in the bad regime \mathcal{R}_{bad} , denote $\lambda \in \mathcal{R}_{\text{bad}}$, if there exists $\beta_*(\lambda) \in (0, 1)$.

Definition B.3 (Critical Regime). We say that a value $\lambda_{\text{crit}} > 0$ is in the critical regime $\mathcal{R}_{\text{crit}}$, denote $\lambda_{\text{crit}} \in \mathcal{R}_{\text{crit}}$ if for some κ_{crit} , the pair $\lambda_{\text{crit}}, \kappa_{\text{crit}}$ solves:

$$\begin{cases} \kappa^2 \delta = \sigma^2 + \mathbb{E}[(\eta(\kappa h + \mu \bar{X}_0; \kappa \lambda) - \mu \bar{X}_0)^2], & (57) \\ \kappa \delta = \mathbb{E}[(\eta(\kappa h + \mu \bar{X}_0; \kappa \lambda) \cdot h)]. & (58) \end{cases}$$

As an immediate consequence of the definition above and the first order optimality conditions in Lemma B.1, we have

$$\beta_*(\lambda_{\text{crit}}) = 1, \quad \kappa_*(\lambda_{\text{crit}}) = \kappa_{\text{crit}} \quad \text{and} \quad \alpha_*(\lambda_{\text{crit}}) = \sqrt{\delta \kappa_{\text{crit}}^2 - \sigma^2}. \quad (59)$$

Also, the following lemma reveals the importance of λ_{crit} : all $\lambda < \lambda_{\text{crit}}$ are in \mathcal{R}_{bad} and the squared error is constant in that regime, i.e. $\alpha_*(\lambda) = \alpha_*(\lambda_{\text{crit}})$.

Lemma B.2 (Error in \mathcal{R}_{bad}). Let $\lambda_{\text{crit}} \in \mathcal{R}_{\text{crit}}$. Then, for all $0 < \lambda' < \lambda_{\text{crit}}$, it holds $\lambda' \in \mathcal{R}_{\text{bad}}$. Furthermore, $\beta_*(\lambda') = \lambda/\lambda_{\text{crit}}$, $\lambda' \kappa_*(\lambda') = \kappa_{\text{crit}} \lambda_{\text{crit}}$ and $\alpha_*(\lambda') = \alpha_*(\lambda_{\text{crit}})$.

Proof. Fix any $0 < \lambda' < \lambda_{\text{crit}}$. By definition, there exists κ_{crit} such that $\lambda_{\text{crit}}, \kappa_{\text{crit}}$ satisfy (57)-(58). Define $\beta' := \lambda/\lambda_{\text{crit}}$ and $\kappa' := \kappa_{\text{crit}}/\beta'$. It is then easy to see that β', κ' solve (46)-(47) (for $\lambda = \lambda'$ therein). Also, $\beta' < 1$ by definition. Thus, $\lambda' \in \mathcal{R}_{\text{bad}}$ and $\beta_*(\lambda') = \lambda/\lambda_{\text{crit}}$, $\kappa_*(\lambda') = \kappa_{\text{crit}} \lambda_{\text{crit}}/\lambda'$. Also, using (48) and (59), $\alpha_*(\lambda) = \sqrt{\delta \beta_*^2(\lambda') \kappa_*^2(\lambda') - \sigma^2} = \sqrt{\delta \kappa_*^2(\lambda_{\text{crit}}) - \sigma^2} = \alpha_*(\lambda_{\text{crit}})$. \square

It is thus important to identify the critical values of the regularizer parameter, i.e. all $\lambda_{\text{crit}} \in \mathcal{R}_{\text{crit}}$. Values in \mathcal{R}_{bad} are important towards this direction, since as shown in the next lemma, for any $\lambda \in \mathcal{R}_{\text{bad}}$ there must exist some $\lambda_{\text{crit}} > \lambda$.

Lemma B.3 ($\mathcal{R}_{\text{bad}} \rightarrow \lambda_{\text{crit}}$). Let $\lambda_1 \in \mathcal{R}_{\text{bad}}$, then there exists $\lambda_2 \in \mathcal{R}_{\text{crit}}$ with $\lambda_2 > \lambda_1$.

Proof. Let $\beta_1, \alpha_1, \kappa_1$ be optimal corresponding to λ_1 . Since $\lambda_1 \in \mathcal{R}_{\text{bad}}$, it holds $0 < \beta_1 < 1$. Then, from Lemma B.1, κ_1, β_1 solve (46)-(47). Starting from these and substituting $\lambda_2 := \lambda_1/\beta_1$ and $\kappa_2 := \kappa_1 \beta_1$ therein, it is not hard to see that this is equivalent with λ_2, κ_2 satisfying (57)-(58). Thus, $\lambda_2 \in \mathcal{R}_{\text{crit}}$. Also, clearly $\lambda_2 > \lambda_1$. \square

The lemma below is important since it shows that when $\delta < 1$ there exists a *unique* $\lambda_{\text{crit}} \in \mathcal{R}_{\text{crit}}$.

Lemma B.4 (Unique λ_{crit}). Suppose $\delta < 1$. The set of equations (57)-(58) has a unique pair of solutions (κ, λ) . Thus, there exists unique $\lambda_{\text{crit}} \in \mathcal{R}_{\text{crit}}$.

Proof. First, we show that there exists at most one $\lambda_{\text{crit}} \in \mathcal{R}_{\text{crit}}$. For the sake of contradiction assume two different pairs of solutions, say (κ_1, λ_1) and (κ_2, λ_2) . By definition, $\lambda_1, \lambda_2 \in \mathcal{R}_{\text{crit}}$. First, note that we cannot have $\lambda_1 = \lambda_2$, since if this was the case then from (59) we would also have $\kappa_1 = \kappa_2$. Henceforth, assume w.l.o.g. that $\lambda_1 < \lambda_2$. It follows from Lemma B.2 that $\lambda_1 \in \mathcal{R}_{\text{bad}}$ and also $\kappa_*(\lambda_1)\lambda_1 = \kappa_*(\lambda_2)\lambda_2$. Thus,

$$\kappa_*(\lambda_1) < \kappa_*(\lambda_2). \quad (60)$$

But also, again from Lemma B.2, $\alpha_*(\lambda_1) = \alpha_*(\lambda_2)$. Since, $\lambda_1, \lambda_2 \in \mathcal{R}_{\text{crit}}$, this implies when combined with (59) that $\kappa_*(\lambda_1) = \kappa_*(\lambda_2)$, which contradicts (60), completing the proof of this part.

Let us now prove that $\mathcal{R}_{\text{crit}}$ is non-empty. To begin with, we show that \mathcal{R}_{bad} is non-empty in this case. In particular, we show that λ_{min} defined in Lemma B.1 is in \mathcal{R}_{bad} . Since, $\delta < 1$, we have $\lambda_{\text{min}} > 0$. Suppose that $(\beta_*(\lambda_{\text{min}}) = 1, \kappa_*(\lambda_{\text{min}}))$ is optimal for some $\kappa_*(\lambda_{\text{min}})$, then, from first-order optimality conditions, $\kappa_*(\lambda_{\text{min}}), \lambda_{\text{min}}$ solves (46) for $\beta = 1$. But, then as in [3, pg. 16] $\kappa_*(\lambda_{\text{min}}) \rightarrow \infty$. Also, since $H(\alpha, \tau, \beta)$ is concave in β , the above imply that $\frac{\partial H}{\partial \beta} \Big|_{(\beta=0, \kappa \rightarrow \infty)} \geq 0$, or equivalently from (56),

$$\int_{\lambda_{\text{min}}}^{\infty} h(h - \lambda_{\text{min}})e^{-h^2/2} dh \leq \delta \sqrt{\frac{\pi}{2}}.$$

Recalling the definition of λ_{min} in Lemma B.1, it can be shown (using standard inequalities on tail functions of gaussians) that the inequality above is violated for all $0 < \delta < 1$. Hence, it must be $\beta_*(\lambda_{\text{min}}) < 1$. Also, $\beta_*(\lambda_{\text{min}}) > 0$ because of (53). Thus, $\lambda_{\text{min}} \in \mathcal{R}_{\text{bad}}$. To complete, the proof use Lemma B.3 with $\lambda_1 = \lambda_{\text{min}}$ to see that there exists $\lambda_2 \in \mathcal{R}_{\text{crit}}$. \square

Lemma B.5 ($\delta < 1$). *Suppose $\delta < 1$ and let $\lambda_{\text{crit}} \in \mathcal{R}_{\text{crit}}$. Furthermore, i) for all $\lambda \leq \lambda_{\text{crit}}$, $\alpha_*(\lambda) = \alpha_*(\lambda_{\text{crit}})$, and, ii) for all $\lambda > \lambda_{\text{crit}}$, $\kappa_*(\lambda)$ is the unique solution to (46) for $\beta = 1$.*

Proof. Existence and uniqueness of λ_{crit} is proved in Lemma B.4

i) For $\lambda \leq \lambda_{\text{crit}}$, the claim follows directly from Lemma B.2.

ii) Next, we show that for $\lambda \geq \lambda_{\text{crit}}$, there exists an optimal solution for which $\beta_*(\lambda) = 1$. This suffices since then $\kappa_*(\lambda)$ solves (46) for $\beta = 1$ (by first order optimality conditions), and, also, the solution is unique by [11], [3, Prop. 1.3] and the fact that $\lambda_{\text{min}} \leq \lambda_{\text{crit}} \leq \lambda$. To see that $\beta_*(\lambda) = 1$, we argue as follows. First, $\beta_*(\lambda) \notin (0, 1)$. Otherwise, $\lambda \in \mathcal{R}_{\text{bad}}$, thus, by Lemma B.3 there exists $\lambda' > \lambda \geq \lambda_{\text{crit}}$ such that $\lambda' \in \mathcal{R}_{\text{crit}}$, which contradicts the uniqueness of λ_{crit} . Hence, $\beta_*(\lambda) = 1$. \square

Lemma B.6 ($\delta > 1$). *Suppose $\delta > 1$, then for all $\lambda > 0$, $\kappa_*(\lambda)$ is the unique solution to (46) for $\beta = 1$.*

Proof. First, let us show that for $\lambda \rightarrow 0$, $\beta_*(\lambda) = 1$. Indeed for $\beta = 1$ and $\lambda \rightarrow 0$, (56) gives

$$\frac{\partial H}{\partial \beta} = \delta - \mathbb{E}\left[h + \frac{\mu}{\kappa} \bar{X}_0\right]h = \delta - 1 > 0.$$

Thus, from concavity of H with respect to β , we find that the unique optimal value for β is

$$\beta_*(\lambda \rightarrow 0) = 1. \quad (61)$$

Also, as in the proof of Lemma B.5, $\beta_*(\lambda \rightarrow \infty) = 1$. Thus, again similar to Lemma B.5, it suffices to prove that there exists no $\lambda \in \mathcal{R}_{\text{bad}}$. For the sake of contradiction, suppose that there exists $\lambda_1 \in \mathcal{R}_{\text{bad}}$. By Lemma B.3, there exists $\lambda_1 < \lambda_{\text{crit}} \in \mathcal{R}_{\text{crit}}$. But, then $\beta_*(\lambda \rightarrow 0) \rightarrow 0$, which contradicts (61). This completes the proof. \square

Proof. (of Theorem 2.3) The claim of the theorem is now a direct consequence of Lemmas B.5 and B.6 combined with (55). \square

C Proofs for Section 2.3

C.1 The LM Algorithm

The Lloyd-Max algorithm is an algorithm for finding the quantization threshold t_i and the representation points ℓ_i . Given real values $x \in \mathbb{R}$ sampled from some probability density $\phi(x)$ it looks for optimal sets $\hat{\mathbf{t}}, \hat{\boldsymbol{\ell}}$ that minimizes the mean-square-error (MSE) between x and their corresponding quantized values $Q_q(x; \boldsymbol{\ell}, \mathbf{t})$, i.e.

$$(\hat{\boldsymbol{\ell}}, \hat{\mathbf{t}}) := \arg \min_{\boldsymbol{\ell}, \mathbf{t}} \mathbb{E}_{x \sim \phi} [(x - Q_q(x; \boldsymbol{\ell}, \mathbf{t}))^2]. \quad (62)$$

The algorithm simply alternates between i) optimizing the threshold t_i for a given set of $\boldsymbol{\ell}$, and then ii) optimizing the levels ℓ_i for the new thresholds. It is well known that the converging points $\boldsymbol{\ell}^{LM}, \mathbf{t}^{LM}$ of the algorithm satisfy

$$\mathbf{t}_i^{LM} = \frac{\ell_i^{LM} + \ell_{i+1}^{LM}}{2} \quad i = 1, \dots, L-1, \quad (63a)$$

$$\ell_i^{LM} = \left(\int_{\mathbf{t}_{i-1}^{LM}}^{\mathbf{t}_i^{LM}} \phi(x) dx \right)^{-1} \left(\int_{\mathbf{t}_{i-1}^{LM}}^{\mathbf{t}_i^{LM}} x \phi(x) dx \right) \quad i = 1, \dots, L. \quad (63b)$$

Furthermore, they are stationary points of the objective function in (62).

C.1.1 Gaussian case

Assume that the values x are sampled from a standard gaussian distribution, i.e. $x \sim \mathcal{N}(0, 1)$ and $\phi(x) = (1/\sqrt{2\pi}) \exp(-x^2/2)$. Also, recall the definition of the parameters μ, σ^2 in (3); setting $g = Q_q$ therein, we find (also, to compare with (20))

$$\mu := \mu(\boldsymbol{\ell}, \mathbf{t}) = 2 \sum_{i=1}^L \ell_i \int_{\mathbf{t}_{i-1}}^{\mathbf{t}_i} x \phi(x) dx, \quad (64a)$$

$$\tau^2 := \tau^2(\boldsymbol{\ell}, \mathbf{t}) = 2 \sum_{i=1}^L \ell_i^2 \int_{\mathbf{t}_{i-1}}^{\mathbf{t}_i} \phi(x) dx, \quad (64b)$$

In this notation, the objective in (62) can be writthen as $\tau^2 - 2\mu + 1$. Thus, $\boldsymbol{\ell}^{LM}, \mathbf{t}^{LM}$ satisfy

$$(\tau^2)'|_{(\boldsymbol{\ell}^{LM}, \mathbf{t}^{LM})} = 2\mu'|_{(\boldsymbol{\ell}^{LM}, \mathbf{t}^{LM})}, \quad (65)$$

Here and onwards we use $(\tau^2)', \mu'$ to denote the gradient of τ^2 and μ with respect to the vector $[\boldsymbol{\ell}^T, \mathbf{t}^T]$. The gradients are evaluated at the point $(\boldsymbol{\ell}^{LM}, \mathbf{t}^{LM})$ in (65).

C.2 q-Bit Compressive Sensing

We prove that the LM algorithm is an efficient algorithm when the objective is minimizing the LASSO reconstruction error of a signal \mathbf{x}_0 to which we have access through q -bit quantized linear measurements $Q_q(\mathbf{a}_i^T \mathbf{x}; \boldsymbol{\ell}, \mathbf{t})$. It was shown in Section 2.3.2 that the problem can be posed as that of finding $\boldsymbol{\ell}_*, \mathbf{t}_*$ such that

$$(\mathbf{t}_*, \boldsymbol{\ell}_*) = \arg \min_{\mathbf{t}, \boldsymbol{\ell}} \frac{\sigma^2(\mathbf{t}, \boldsymbol{\ell})}{\mu^2(\mathbf{t}, \boldsymbol{\ell})} = \arg \min_{\mathbf{t}, \boldsymbol{\ell}} \frac{\tau^2(\mathbf{t}, \boldsymbol{\ell})}{\mu^2(\mathbf{t}, \boldsymbol{\ell})}. \quad (66)$$

The following Lemma proves the claim made in Section 2.3.3, i.e. the converging points of the LM algorithm are stationary points of the objective function in (66).

Lemma C.1. *The converging points of the LM algorithm, say $(\mathbf{t}^{LM}, \boldsymbol{\ell}^{LM})$ satisfy*

$$\begin{aligned} \frac{\partial}{\partial \ell_i} \left(\frac{\tau^2(\boldsymbol{\ell}, \mathbf{t})}{\mu^2(\boldsymbol{\ell}, \mathbf{t})} \right) \Big|_{(\boldsymbol{\ell}, \mathbf{t})=(\boldsymbol{\ell}^{LM}, \mathbf{t}^{LM})} &= 0, & i = 1, \dots, L, \\ \frac{\partial}{\partial t_i} \left(\frac{\tau^2(\boldsymbol{\ell}, \mathbf{t})}{\mu^2(\boldsymbol{\ell}, \mathbf{t})} \right) \Big|_{(\boldsymbol{\ell}, \mathbf{t})=(\boldsymbol{\ell}^{LM}, \mathbf{t}^{LM})} &= 0, & i = 0, \dots, L-1. \end{aligned} \quad (67)$$

Proof. Call $R(\mathbf{t}, \boldsymbol{\ell}) = \frac{\tau^2(\mathbf{t}, \boldsymbol{\ell})}{\mu^2(\mathbf{t}, \boldsymbol{\ell})}$. We denote $R' := R'(\mathbf{t}, \boldsymbol{\ell})$ for its gradient with respect to the vector $[\mathbf{t}^T, \boldsymbol{\ell}^T]$. It suffices to prove that $R'|_{(\boldsymbol{\ell}^{LM}, \mathbf{t}^{LM})} = 0$, or equivalently, that at the point $(\mathbf{t}, \boldsymbol{\ell}) = (\mathbf{t}^{LM}, \boldsymbol{\ell}^{LM})$ the following holds:

$$(\tau^2)' \mu^2 = 2\tau^2 \mu \mu'. \quad (68)$$

To see that this is the case, note that

$$\tau^2(\mathbf{t}^{LM}, \boldsymbol{\ell}^{LM}) = \mu(\mathbf{t}^{LM}, \boldsymbol{\ell}^{LM}) \quad (69)$$

This follows by direct substitution of (63) in (64). Then, (68) follows from (69) and (65). \square