

1 **A GEOMETRIC ANALYSIS OF**
2 **MODEL- AND ALGORITHM-INDUCED UNCERTAINTIES**
3 **FOR RANDOMIZED LEAST SQUARES REGRESSION***

4 JOCELYN T. CHI[†] AND ILSE C. F. IPSEN[‡]

5 **Abstract.** For full-rank least squares regression problems under a Gaussian linear model, we
6 analyze the uncertainties when the minimum-norm solution is computed by random row-sketching
7 and, in particular random row-sampling. Our expressions for the total expectation and variance of
8 the solution—with regard to both model- and algorithm-induced uncertainties—are exact; hold for
9 general sketching matrices; and make no assumptions on the rank of the sketched matrix. They show
10 that expectation and variance are governed by the rank-deficiency and spatial geometry induced by
11 the sketching process, rather than by structural properties of specific sketching or sampling methods.
12 In order to analyze the rank-deficient matrices from row-sketching, we introduce two projectors that
13 connect least squares problems of different dimensions.

14 From a deterministic perspective, our structural perturbation bounds imply that least squares
15 solutions are less sensitive to multiplicative perturbations than to additive perturbations. From a
16 probabilistic perspective, we show that the differences between the total bias and variance on the
17 one hand, and the model bias and variance on the other hand, are governed by two factors: (i) the
18 expected rank deficiency of the sketched matrix, and (ii) the expected difference between projectors
19 onto the spaces of the original and the sketched problems. Surprisingly, the matrix condition number
20 has far less impact on the statistical quantities than it has on numerical errors.

21 **Key words.** Condition number with respect to inversion, projector, multiplicative perturba-
22 tions, Moore Penrose inverse, expectation, variance, matrix valued random variable

23 **AMS subject classification.** 62J05, 62J10, 65F20, 65F22, 65F35, 68W20

24 **1. Introduction.** We consider the randomized solution of least squares regres-
25 sion problems under the Gaussian linear model, and analyze the effect of both: the
26 statistical noise in the model, as well as the error due to algorithmic randomization.
27 Our analysis extends the pioneering work [15, 16] through rigorous validation in a
28 general setting, and demonstrates that expectation and variance are governed by ge-
29 ometry rather than by structural properties of specific classes of sketching matrices:
30 What matters is the rank deficiency induced by the sketching process, and the failure
31 of the sketched matrix to reproduce the original column space.

32 **1.1. Problem setting.** We start with a regression problem under the Gaussian
33 linear model,

34 (1.1)
$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta}_0 + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_n),$$

35 where $\mathbf{X} \in \mathbb{R}^{n \times p}$ is a given design matrix with $\text{rank}(\mathbf{X}) = p$, $\boldsymbol{\beta}_0 \in \mathbb{R}^p$ is the true but
36 unknown parameter vector, and the noise vector $\boldsymbol{\epsilon} \in \mathbb{R}^n$ has a standard multivari-
37 ate normal distribution. For a fixed response vector $\mathbf{y} \in \mathbb{R}^n$, the unique maximum
38 likelihood estimator of $\boldsymbol{\beta}_0$ is the solution $\hat{\boldsymbol{\beta}}$ of the full-rank least squares problem¹

39 (1.2)
$$\min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{X}\boldsymbol{\beta} - \mathbf{y}\|_2.$$

*The work was supported in part by NSF grants DGE-1633587, DMS-1760374, and DMS-1745654.

[†]Department of Statistics, North Carolina State University, Raleigh, NC 27695, USA,
jtchi@ncsu.edu

[‡]Department of Mathematics, North Carolina State University, Raleigh, NC 27695,
ipsen@ncsu.edu

¹Here $\|\mathbf{x}\|_2 = \sqrt{\mathbf{x}^T \mathbf{x}}$ represents the Euclidean two-norm, and the superscript T the transpose.

40 Solution of this least squares problem via random row-sketching,

41 (1.3)
$$\min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{S}(\mathbf{X}\boldsymbol{\beta} - \mathbf{y})\|_2,$$

42 is an effective approach in the highly over-constrained case [5, 6, 16, 22, 28] where
 43 observations far outnumber covariates, that is, \mathbf{X} is tall and skinny with $n \gg p$. Here
 44 $\mathbf{S} \in \mathbb{R}^{r \times n}$ is a random sketching matrix with $r \leq n$, and the minimum norm solution
 45 is $\tilde{\boldsymbol{\beta}}$.

46 **1.2. Existing work.** Random sketching is a form of preconditioning and seems
 47 to have originated in [24]. By now, there are many variants which can be classified
 48 according to [26, Section 1]: Compression of rows [2, 5, 6, 13, 15, 16, 23, 28]; or columns
 49 [1]; or both [18]. Matrix concentration inequalities are used to derive probabilistic
 50 bounds for the error due to randomization [1, 6], and for the condition number of
 51 the sampled matrix [13]. From a practical perspective, bootstrapping can deliver fast
 52 error estimates [14].

53 Most of the randomized least squares work comes from theoretical computer sci-
 54 ence and numerical analysis and is mainly concerned with errors due to algorithmic
 55 randomization, while ignoring statistical noise in the model. The pioneering work
 56 [15, 16] is the first to quantify the total uncertainty from model-induced and algorithm-
 57 induced randomness. This being the first analysis of its kind, it started out with a
 58 few assumptions: the sampling matrices must preserve rank, and their expected value
 59 must be known; and the conditional expectation and variances must admit Taylor
 60 series. Thus, the resulting first-order expansions hold only approximately.

61 **1.3. Specific Contributions.** We extend the first-order expansions in [15, 16]
 62 as follows:

- 63 1. We derive *exact* expressions for the total expectation and variance of $\tilde{\boldsymbol{\beta}}$ with
 64 regard to model- and algorithm-induced uncertainties (Theorem 4.5). The ex-
 65 pressions hold for *general* random sketching matrices \mathbf{S} , regardless of whether
 66 they preserve rank, and include sketching matrices that perform projections
 67 prior to sampling.
- 68 2. In contrast to most deterministic and randomized analyses, our expressions
 69 are not limited to full-rank matrices. We analyse the rank-deficient matrices
 70 in (1.3) by supplementing the *hat matrix* $\mathbf{P}_x = \mathbf{X}\mathbf{X}^\dagger$, i.e. the orthogonal
 71 projector onto $\text{range}(\mathbf{X})$, with two new projectors:
 - 72 (a) *Comparison hat matrix* $\mathbf{P} = \mathbf{X}(\mathbf{S}\mathbf{X})^\dagger\mathbf{S}$ (Lemma 3.1).
 73 This projector makes it possible to compare the model problem (1.1)
 74 with the lower-dimensional sketched problem (1.3). The difference
 75 $\mathbf{P}\mathbf{P}^T - \mathbf{P}_x$ quantifies the deviation of \mathbf{P} from being an orthogonal
 76 projector onto $\text{range}(\mathbf{X})$.
 - 77 (b) *Bias projector* $\mathbf{P}_0 = (\mathbf{S}\mathbf{X})^\dagger(\mathbf{S}\mathbf{X})$ (Lemma 4.1).
 78 This projector captures the failure of \mathbf{S} to preserve rank. The differ-
 79 ence $\mathbf{I} - \mathbf{P}_0$ quantifies the rank deficiency of the sketched matrix $\mathbf{S}\mathbf{X}$.
- 80 3. For the model-induced uncertainty of $\tilde{\boldsymbol{\beta}}$, conditioned on the sampling ma-
 81 trix \mathbf{S} , we show (Theorem 4.2, Corollary 4.3):
 - 82 (a) The conditional bias increases with the rank deficiency of $\mathbf{S}\mathbf{X}$.
 - 83 (b) The difference between conditional variance and model variance in-
 84 creases with the deviation of \mathbf{P} from being an orthogonal projector
 85 onto $\text{range}(\mathbf{X})$.

86 Thus, unbiasedness is easier to achieve because it only requires $\mathbf{S}\mathbf{X}$ to have full

87 column rank. In contrast, recovering the model variance requires reproducing
 88 all of the original space range(\mathbf{X}).

89 4. For the *total* uncertainty in the solution $\tilde{\boldsymbol{\beta}}$ we show (Theorem 4.5, Corol-
 90 lary 4.6):

- 91 (a) The total bias increases with the expected rank deficiency of $\mathbf{S}\mathbf{X}$.
- 92 (b) The difference between total variance and model variance increases
 93 with two terms: the expected rank deficiency of $\mathbf{S}\mathbf{X}$; and the expected
 94 deviation of \mathbf{P} from being an orthogonal projector onto range(\mathbf{X}).

95 Thus, total expectation and variance are governed by the expected spatial
 96 geometry induced by the sketching process rather than by structural proper-
 97 ties of specific \mathbf{S} . However, the condition number of \mathbf{X} has far less impact
 98 than one would have expected based on numerical perturbation theory.

99 5. We show analogous results for norm-wise quantities (Theorem 4.8, Corol-
 100 lary 4.9). The total expectations of the *regression sum of squares* $\|\mathbf{P}\mathbf{y}\|_2^2$ and
 101 the *residual sum of squares* $\|(\mathbf{I}-\mathbf{P})\mathbf{y}\|_2^2$ depend on the norms of the projectors
 102 \mathbf{P} and $\mathbf{I}-\mathbf{P}$, amplified by the model variance σ^2 .

103 6. We present structural bounds that improve existing perturbation bounds
 104 (Corollary 3.5). They imply that the minimum norm solution $\tilde{\boldsymbol{\beta}}$ is less sensi-
 105 tive to multiplicative perturbations than to additive perturbations, because
 106 the dependence is only on the condition number, rather than on its square as
 107 in the case of additive perturbations.

108 The judicious design of numerical experiments that are representative and informative
 109 from both, numerical and statistical perspectives, is beyond this scope, and will be
 110 the subject of a separate paper.

111 **1.4. Overview.** After reviewing the computational models for least squares re-
 112 gression (Section 2), we adopt two perspectives:

- 113 1. Deterministic (Section 3): The matrix \mathbf{S} is fixed and the sketched problem
 114 (1.3) is a multiplicative perturbation of the deterministic problem (1.2), and
 115 we present structural perturbation bounds.
- 116 2. Probabilistic (Section 4): The matrix \mathbf{S} is a matrix-valued random variable,
 117 and (1.3) is a randomized algorithm for solving the model problem (1.1), and
 118 we derive expressions for expectation and variance with regard to the model-
 119 and algorithm-induced uncertainties.

120 A brief discussion of our results (Section 5) ends the main part of the paper. Proofs
 121 are relegated to the Appendix (Section A), as are specific examples to provide insight
 122 for the geometry of the probabilistic results (Section B).

123 **2. Models for Least Squares Regression.** Given is a fixed design matrix
 124 $\mathbf{X} \in \mathbb{R}^{n \times p}$ with rank(\mathbf{X}) = p . Since \mathbf{X} has full column rank, the Moore-Penrose
 125 inverse is a left inverse with

$$126 \quad (2.1) \quad \mathbf{X}^\dagger = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \quad \text{and} \quad \mathbf{X}^\dagger \mathbf{X} = \mathbf{I}_p.$$

127 The two-norm condition number of \mathbf{X} with regard to left inversion is

$$128 \quad \kappa_2(\mathbf{X}) \equiv \|\mathbf{X}\|_2 \|\mathbf{X}^\dagger\|_2.$$

129 We review the different incarnations of least squares regression: the Gaussian lin-
 130 ear model (Section 2.1), the traditional computation (Section 2.2), and algorithmic
 131 leveraging (Section 2.3).

132 **2.1. Gaussian linear model.** Let $\boldsymbol{\beta}_0 \in \mathbb{R}^p$ denote the true but generally un-
 133 known parameter vector, and let the response vector $\mathbf{y} \in \mathbb{R}^n$ satisfy the Gauss-Markov
 134 assumptions,

$$135 \quad (2.2) \quad \mathbf{y} = \mathbf{X}\boldsymbol{\beta}_0 + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_n).$$

136 The noise vector $\boldsymbol{\epsilon} \in \mathbb{R}^n$ has a multivariate normal distribution whose mean is the
 137 vector of all zeros, $\mathbf{0} \in \mathbb{R}^n$, and whose covariance is a multiple $\sigma^2 > 0$ of the identity
 138 matrix $\mathbf{I}_n \in \mathbb{R}^{n \times n}$.

139 **2.2. Traditional algorithm for least squares solution.** For a fixed $\mathbf{y} \in \mathbb{R}^n$
 140 solve

$$141 \quad (2.3) \quad \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{X}\boldsymbol{\beta} - \mathbf{y}\|_2.$$

142 Since \mathbf{X} has full column rank, (2.3) is well posed and has the unique solution

$$143 \quad (2.4) \quad \hat{\boldsymbol{\beta}} \equiv \mathbf{X}^\dagger \mathbf{y}.$$

144 The prediction and the least squares residual are, respectively

$$145 \quad \hat{\mathbf{y}} \equiv \mathbf{X}\hat{\boldsymbol{\beta}} \quad \text{and} \quad \hat{\boldsymbol{\epsilon}} \equiv \mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}} = \mathbf{y} - \hat{\mathbf{y}}.$$

146 In terms of the *hat matrix* [3, 10, 27],

$$147 \quad (2.5) \quad \mathbf{P}_x \equiv \mathbf{X}\mathbf{X}^\dagger = \mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \in \mathbb{R}^{n \times n},$$

148 which is the orthogonal projector onto $\text{range}(\mathbf{X})$ along $\text{null}(\mathbf{X}^T)$, the prediction and
 149 least squares residual can be expressed as

$$150 \quad (2.6) \quad \hat{\mathbf{y}} = \mathbf{P}_x \mathbf{y} \quad \text{and} \quad \hat{\boldsymbol{\epsilon}} = (\mathbf{I} - \mathbf{P}_x) \mathbf{y}.$$

151 **2.3. Random Row-Sketching.** From a deterministic perspective, this can be
 152 considered an extension of weighted least squares [8, Section 6.1] to rectangular weight-
 153 ing matrices.

154 Given a sketching matrix $\mathbf{S} \in \mathbb{R}^{r \times n}$ with $1 \leq r \leq n$, solve

$$155 \quad (2.7) \quad \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{S}(\mathbf{X}\boldsymbol{\beta} - \mathbf{y})\|_2,$$

156 which has the minimum norm solution

$$157 \quad (2.8) \quad \tilde{\boldsymbol{\beta}} \equiv (\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{y}.$$

158 This problem is generally ill-posed: Just because \mathbf{S} has $r > p$ rows, this does not
 159 imply $\text{rank}(\mathbf{S}) = p$; and even if \mathbf{S} does have full column rank, $\text{rank}(\mathbf{S}\mathbf{X}) < p$ is still
 160 possible.

161 By design, \mathbf{S} has fewer rows than \mathbf{X} . Hence the corresponding predictions $\hat{\mathbf{y}} = \mathbf{X}\hat{\boldsymbol{\beta}}$
 162 and $\mathbf{S}\mathbf{X}\tilde{\boldsymbol{\beta}}$ have different dimension and cannot be directly compared; neither can
 163 their residuals. To remedy this, we follow previous work [5, 6, 22], and compare the
 164 predictions with regard to the *original* matrix,

$$165 \quad (2.9) \quad \tilde{\mathbf{y}} \equiv \mathbf{X}\tilde{\boldsymbol{\beta}} \quad \text{and} \quad \tilde{\boldsymbol{\epsilon}} \equiv \mathbf{y} - \mathbf{X}\tilde{\boldsymbol{\beta}} = \mathbf{y} - \tilde{\mathbf{y}}.$$

166 Note that $\tilde{\boldsymbol{\epsilon}}$ is not a least squares residual; the least squares residual for (2.7) is
 167 $\mathbf{S}\mathbf{X}\tilde{\boldsymbol{\beta}} - \mathbf{S}\mathbf{y}$. However, we need $\tilde{\boldsymbol{\epsilon}}$ to assess the performance of $\tilde{\boldsymbol{\beta}}$ in the context of the
 168 original problem (2.3).

169 **3. Structural (deterministic perturbation) bounds.** Here \mathbf{S} is a fixed, gen-
 170 eral matrix; and \mathbf{SX} is interpreted as a perturbation of \mathbf{X} . We derive expressions for
 171 the quantities of interest from the perturbed problem (Section 3.1), followed by mul-
 172 tiplicative perturbation bounds (Section 3.2).

173 **3.1. The perturbed problem.** We derive expressions for the solution, predic-
 174 tion and residual of the lower-dimensional problem (2.7). In order to relate them to
 175 the higher-dimensional original problem (2.3), we introduce (Lemma 3.1) a *compari-*
 176 *son hat matrix* \mathbf{P} for (2.7), which corresponds to the *hat matrix* \mathbf{P}_x in (2.5) for the
 177 original problem (2.3). This makes it possible to express the solution, prediction, and
 178 residual of the perturbed problem in terms of the original problem (Theorem 3.3).

179 LEMMA 3.1 (Comparison hat matrix). *With the assumptions in Section 2,*

$$180 \quad \mathbf{P} \equiv \mathbf{X}(\mathbf{SX})^\dagger \mathbf{S}$$

181 *is an oblique projector where*

- 182 1. $\mathbf{P}_x \mathbf{P} = \mathbf{P}$.
- 183 2. $\mathbf{P} - \mathbf{P}_x$ reflects the difference between the spaces $\text{null}(\mathbf{P})$ and $\text{null}(\mathbf{P}_x)$.
- 184 3. $\mathbf{P}\mathbf{X} = \mathbf{X}$ if $\text{rank}(\mathbf{SX}) = p$.

185 *Proof.* See Section A.1 □

186 The name *comparison hat matrix* will become clear in Theorem 3.3, where \mathbf{P}
 187 assumes the duties of the *hat matrix* \mathbf{P}_x for the expressions in (2.9).

188 *Remark 3.2.* The following cases are possible.

- 189 • If $\mathbf{S} = \mathbf{I}_n$, then $\mathbf{P} = \mathbf{P}_x$.
- 190 • If $\text{rank}(\mathbf{SX}) = \text{rank}(\mathbf{X})$, then \mathbf{P} is an oblique version of the orthogonal pro-
 191 jector \mathbf{P}_x with $\text{range}(\mathbf{P}) = \text{range}(\mathbf{P}_x)$, but $\text{null}(\mathbf{P}) \neq \text{null}(\mathbf{P}_x)$ in general.
- 192 • If

$$193 \quad \text{rank}(\mathbf{P}) = \text{rank}(\mathbf{SX}) < \text{rank}(\mathbf{X}) = \text{rank}(\mathbf{P}_x) = p,$$

194 then \mathbf{P} projects only onto a subspace of $\text{range}(\mathbf{X})$.

195 The comparison hat matrix \mathbf{P} generalizes the oblique projector \mathbf{P}_u in [22, (11)],
 196 which was introduced to quantify *prediction efficiency* and *residual efficiency* of
 197 sketching algorithms in the statistical setting (2.2). This projector \mathbf{P}_u is defined
 198 if $\text{rank}(\mathbf{SX}) = p$, and equals $\mathbf{P}_u \equiv \mathbf{U}(\mathbf{SU})^\dagger \mathbf{S} = \mathbf{P}$, where \mathbf{U} is an orthonormal ba-
 199 sis for $\text{range}(\mathbf{X})$. However, if $\text{rank}(\mathbf{SX}) < \text{rank}(\mathbf{X})$, then \mathbf{P}_u is not sufficient in our
 200 context.

201 THEOREM 3.3 (Perturbed least squares problem). *With the assumptions in Sec-*
 202 *tion 2, the solution of (2.7) satisfies*

$$203 \quad \tilde{\boldsymbol{\beta}} = \mathbf{X}^\dagger \mathbf{P}\mathbf{y} = \hat{\boldsymbol{\beta}} + \mathbf{X}^\dagger (\mathbf{P} - \mathbf{P}_x)\mathbf{y}.$$

204 *The prediction $\tilde{\mathbf{y}} = \mathbf{X}\tilde{\boldsymbol{\beta}}$ and residual $\tilde{\mathbf{e}} = \mathbf{y} - \mathbf{X}\tilde{\boldsymbol{\beta}}$ satisfy*

$$205 \quad \tilde{\mathbf{y}} = \mathbf{P}\mathbf{y} = \hat{\mathbf{y}} + (\mathbf{P} - \mathbf{P}_x)\mathbf{y},$$

$$206 \quad \tilde{\mathbf{e}} = (\mathbf{I} - \mathbf{P})\mathbf{y} = \hat{\mathbf{e}} + (\mathbf{P}_x - \mathbf{P})\mathbf{y}.$$

207 *Proof.* See Section A.2. □

208 Theorem 3.3 shows that the relation between perturbed and original least squares
 209 problems is governed by $\mathbf{P} - \mathbf{P}_x$, which reflects the difference between the spaces
 210 $\text{null}(\mathbf{P})$ and $\text{null}(\mathbf{P}_x)$. The dependence on the sketching matrix is implicit, through
 211 the induced spaces.

212 With its explicit expressions for $\tilde{\boldsymbol{\beta}}$ that hold for general matrices \mathbf{S} without as-
 213 sumptions on $\text{rank}(\mathbf{S}\mathbf{X})$, Theorem 3.3 also strengthens the ground breaking result [16,
 214 Lemma 1], reproduced in the lemma below.

215 LEMMA 3.4 (Lemma 1 in [15] and [16]). *If, in addition to the assumptions in*
 216 *Section 2, the matrix \mathbf{S} in (2.7) has a single nonzero entry per row, the vector $\mathbf{w} \equiv$*
 217 *$\text{diag}(\mathbf{S}^T \mathbf{S}) \in \mathbb{R}^n$ has a scaled multinomial distribution with expected value $\mathbb{E}[\mathbf{w}] = \mathbf{1}$,*
 218 *satisfies $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X})$, and admits a Taylor series expansion of the solution*
 219 *$\tilde{\boldsymbol{\beta}}(\mathbf{w})$ of (2.7) around $\mathbf{w}_0 = \mathbf{1}$ with $\tilde{\boldsymbol{\beta}}(\mathbf{w}_0) = \hat{\boldsymbol{\beta}}$, then*

$$220 \quad \tilde{\boldsymbol{\beta}}(\mathbf{w}) = \hat{\boldsymbol{\beta}} + \mathbf{X}^\dagger \text{diag}(\hat{\boldsymbol{\epsilon}})(\mathbf{w} - \mathbf{1}) + R(\mathbf{w}),$$

221 where $R(\mathbf{w})$ is the remainder of the Taylor series expansion. The Taylor series ex-
 222 pansion is valid if $R(\mathbf{w}) = o(\|\mathbf{w} - \mathbf{w}_0\|_2)$ with high probability.

223 **3.2. Multiplicative perturbation bounds.** We consider (2.7) as a multiplica-
 224 tive perturbation of the original problem (2.3) and derive norm-wise relative pertur-
 225 bation bounds (Corollary 3.5), followed by comparisons to existing work.

226 COROLLARY 3.5. *With the assumptions in Section 2, let $0 < \theta < \pi/2$ be the angle*
 227 *between \mathbf{y} and $\text{range}(\mathbf{X})$.*

228 *The solution $\tilde{\boldsymbol{\beta}}$ of (2.7) satisfies*

$$229 \quad \frac{\|\tilde{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}\|_2}{\|\hat{\boldsymbol{\beta}}\|_2} \leq \kappa_2(\mathbf{X}) \frac{\|\mathbf{y}\|_2}{\|\mathbf{X}\|_2 \|\hat{\boldsymbol{\beta}}\|_2} \|\mathbf{P} - \mathbf{P}_x\|_2 \leq \kappa_2(\mathbf{X}) \frac{\|\mathbf{P} - \mathbf{P}_x\|_2}{\cos \theta}.$$

230 *The prediction $\tilde{\mathbf{y}} = \mathbf{X}\tilde{\boldsymbol{\beta}}$ satisfies*

$$231 \quad \frac{\|\tilde{\mathbf{y}} - \hat{\mathbf{y}}\|_2}{\|\hat{\mathbf{y}}\|_2} \leq \frac{\|\mathbf{P} - \mathbf{P}_x\|_2}{\cos \theta}.$$

232 *Proof.* This is a direct consequence of Theorem 3.3, and of [8, (5.3.16)] which
 233 implies

$$234 \quad \|\mathbf{y}\|_2 / (\|\mathbf{X}\|_2 \|\hat{\boldsymbol{\beta}}\|_2) \leq \|\mathbf{y}\|_2 / \|\mathbf{X}\hat{\boldsymbol{\beta}}\|_2 = 1 / \cos \theta.$$

235 For $\mathbf{S} = \mathbf{I}_n$, the bounds in Corollary 3.5 are zero and therefore tight. Corol-
 236 lary 3.5 implies that the sensitivity of the minimum norm least squares solution $\tilde{\boldsymbol{\beta}}$
 237 to multiplicative perturbations depends on the distance between the spaces $\text{null}(\mathbf{P})$
 238 and $\text{null}(\mathbf{P}_x)$, quantified by $\|\mathbf{P} - \mathbf{P}_x\|_2$. This distance is amplified, as expected, by
 239 the conditioning of \mathbf{X} is with regard to (left) inversion, and by the closeness of \mathbf{y} to
 240 $\text{range}(\mathbf{X})$. Corollary 3.5 is an absolute as well as a relative bound since $\|\mathbf{P}_x\|_2 = 1$.

241 In contrast to multiplicative perturbation bounds for eigenvalue and singular value
 242 problems [11, 12], we do not require \mathbf{S} to be nonsingular or square. Weighted least
 243 squares problems [8, Section 6.1] employ nonsingular diagonal matrices \mathbf{S} for regular-
 244 ization or scaling of discrepancies, and do not view them as a perturbation.

245 In contrast to additive bounds [8, Section 5.3.6], [9, Section 20.1], [25, (3.4)],
 246 there is no squaring of the condition number and no need for requiring $\text{rank}(\mathbf{S}\mathbf{X}) =$

rank(\mathbf{X}). This suggests that the minimum norm solution of (2.7) and its residual are less sensitive to multiplicative perturbations than to additive perturbations.

In contrast to existing structural bounds for randomized least squares algorithms [6, Theorem 1], such as the one in Lemma 3.6 below, the bound for $\hat{\boldsymbol{\beta}}$ in Corollary 3.5 is more general and tighter because it does not exhibit nonlinear dependencies on the perturbations.

LEMMA 3.6 (Theorem 1 in [6]). *In addition to Section 2, assume $\|\mathbf{P}_{\mathbf{x}\mathbf{y}}\|_2 \geq \gamma \|\mathbf{y}\|_2$ for some $0 < \gamma \leq 1$ and $\|\tilde{\boldsymbol{\epsilon}}\|_2 \leq (1 + \eta) \|\hat{\boldsymbol{\epsilon}}\|_2$. Then*

$$\frac{\|\tilde{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}\|_2}{\|\hat{\boldsymbol{\beta}}\|_2} \leq \kappa_2(\mathbf{X}) \sqrt{\gamma^{-2} - 1} \sqrt{\eta}.$$

4. Model-induced and randomized algorithm-induced uncertainty. Under the linear model (2.2), the computed solution $\hat{\boldsymbol{\beta}}$ has nice statistical properties [20, Chapter 6], as it is an unbiased estimator of $\boldsymbol{\beta}_0$ and it has minimal variance among all linear unbiased estimators. We show how this changes with the addition of algorithm-induced uncertainty.

After briefly reviewing the uncertainty induced by the linear model (Section 4.1); we derive the expectation and variance of $\tilde{\boldsymbol{\beta}}$, conditioned on the algorithm-induced uncertainty (Section 4.2), and from that the total expectation and variance (Section 4.3), followed by the derivation of the conditional and total expectations for the regression sum of squares and the residual sum of squares (Section 4.4).

4.1. Model-induced uncertainty. We view the model-induced randomness in (1.1) and (2.2) as a property of the response vector \mathbf{y} , so that

$$\mathbb{E}_{\mathbf{y}}[\boldsymbol{\epsilon}] = \mathbf{0}, \quad \text{Var}_{\mathbf{y}}[\boldsymbol{\epsilon}] = \sigma^2 \mathbf{I}_n.$$

As a consequence

$$(4.1) \quad \mathbb{E}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] = \boldsymbol{\beta}_0, \quad \text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] = \sigma^2 (\mathbf{X}^T \mathbf{X})^{-1} \in \mathbb{R}^{p \times p}.$$

This implies that the computed solution $\hat{\boldsymbol{\beta}}$ is an unbiased estimator for $\boldsymbol{\beta}_0$, and it signals the well-known dependence of the variance on the conditioning of \mathbf{X} [25, Section 5].

The difficulty in analyzing random row-sketching (2.7), coupled with general concern about first-order expansions like the ones in [15, 16], is the frequent occurrence of rank deficiency in the sketched matrix, that is, $\text{rank}(\mathbf{S}\mathbf{X}) < \text{rank}(\mathbf{X})$. In this case $(\mathbf{S}\mathbf{X})^\dagger$ cannot be expressed in terms of $\mathbf{S}\mathbf{X}$ as in (2.1).

One can derive bounds [1, Theorem 3.2], [13, Theorems 4.1 and 5.2] on the probability that $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X})$ for matrices \mathbf{S} that perform uniform sampling and leverage score sampling. However, such bounds are not useful here, because we need the expected values to run over *all* instances of $\mathbf{S}\mathbf{X}$.

We introduce a projector that quantifies the deviation of the columns of $\mathbf{S}\mathbf{X}$ from being linearly independent.

LEMMA 4.1 (Bias projector). *With the assumptions in Section 2,*

$$\mathbf{P}_0 \equiv (\mathbf{S}\mathbf{X})^\dagger (\mathbf{S}\mathbf{X}) \in \mathbb{R}^{p \times p}$$

is an orthogonal projector with

$$1. \quad \mathbf{P}\mathbf{X} = \mathbf{X}\mathbf{P}_0$$

288 2. $\mathbf{P}_0 = \mathbf{I}_p$ if $\text{rank}(\mathbf{S}\mathbf{X}) = p$.

289 As a consequence, $\mathbf{I}_p - \mathbf{P}_0$ quantifies the rank deficiency of $\mathbf{S}\mathbf{X}$.

290 *Proof.* See Section A.3. □

291 If $\text{rank}(\mathbf{S}\mathbf{X}) < p$, then \mathbf{P}_0 characterizes the subspace of $\text{range}(\mathbf{X})$ onto which \mathbf{P}
292 projects. The name *bias projector* will become apparent in Theorem 4.2, where \mathbf{P}_0
293 represents the bias in $\tilde{\boldsymbol{\beta}}$.

294 **4.2. Model-induced uncertainty, conditioned on algorithm-induced un-**
295 **certainty.** We determine the conditional expectation and variance for the solution
296 of (2.7), by assuming that the random sketching matrix \mathbf{S} is fixed at a specific value \mathbf{S}_0 .
297 The expectation conditioned on \mathbf{S} is abbreviated as

$$298 \quad \mathbb{E}_{\mathbf{y}} \left[\cdot \mid \mathbf{S} \right] \equiv \mathbb{E}_{\mathbf{y}} \left[\cdot \mid \mathbf{S} = \mathbf{S}_0 \right].$$

299 The exact expressions below for general matrices \mathbf{S} extend the first-order expres-
300 sions for specific sampling matrices in [16, Lemmas 2-6].

301 **THEOREM 4.2** (Model-induced uncertainty conditioned on \mathbf{S}). *With the assump-*
302 *tions in Section 2, the solution $\tilde{\boldsymbol{\beta}}$ of (2.7) has the conditional expectation*

$$303 \quad \mathbb{E}_{\mathbf{y}} \left[\tilde{\boldsymbol{\beta}} \mid \mathbf{S} \right] = \mathbf{P}_0 \boldsymbol{\beta}_0 = \boldsymbol{\beta}_0 - (\mathbf{I} - \mathbf{P}_0) \boldsymbol{\beta}_0,$$

304 where $\mathbb{E}_{\mathbf{y}} \left[\tilde{\boldsymbol{\beta}} \mid \text{rank}(\mathbf{S}\mathbf{X}) = p \right] = \boldsymbol{\beta}_0$; and the conditional variance

$$305 \quad \begin{aligned} \text{Var}_{\mathbf{y}} \left[\tilde{\boldsymbol{\beta}} \mid \mathbf{S} \right] &= \sigma^2 (\mathbf{X}^\dagger \mathbf{P}) (\mathbf{X}^\dagger \mathbf{P})^T \\ 306 \quad &= \text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] + \sigma^2 \mathbf{X}^\dagger (\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{x}}) (\mathbf{X}^\dagger)^T, \end{aligned}$$

307 with $\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{x}}$ representing the deviation of \mathbf{P} from being an orthogonal projector
308 onto $\text{range}(\mathbf{X})$.

309 *Proof.* See Section A.4. □

310 Theorem 4.2 shows that the conditional bias and variance of $\tilde{\boldsymbol{\beta}}$ depend on the
311 rank deficiency of $\mathbf{S}\mathbf{X}$, and the ability of \mathbf{P} to reproduce the original space $\text{range}(\mathbf{X})$.
312 The fixed sketching matrix \mathbf{S} is involved only implicitly, through the spaces induced
313 by the sketching process. Specifically, Theorem 4.2 shows:

- 314 1. The conditional bias of $\tilde{\boldsymbol{\beta}}$ is proportional to the deviation $\mathbf{I} - \mathbf{P}_0$ of $\mathbf{S}\mathbf{X}$ from
315 having full column rank. That is, the conditional bias becomes worse as the
316 rank deficiency increases. If $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X})$, then $\tilde{\boldsymbol{\beta}}$ is a conditional
317 unbiased estimator of $\boldsymbol{\beta}_0$, regardless of the specific sketching class to which
318 \mathbf{S} belongs.
- 319 2. The conditional variance is close to the model variance $\text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]$, if \mathbf{P} is close
320 to being an orthogonal projector onto $\text{range}(\mathbf{X})$. In the extreme case $\mathbf{S} = \mathbf{I}_n$,
321 the conditional variance is identical to the model variance.

322 The relevance of $\mathbf{I} - \mathbf{P}_0$ and $\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{x}}$ is further corroborated below.

323 **COROLLARY 4.3** (Relative differences between conditional and model uncertain-
324 ties). *With the assumptions in Theorem 4.2,*

$$325 \quad \begin{aligned} &\| \mathbb{E}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \mid \mathbf{S}] - \boldsymbol{\beta}_0 \|_2 \leq \| \mathbf{I} - \mathbf{P}_0 \|_2 \| \boldsymbol{\beta}_0 \|_2 \\ 326 \quad &\frac{\| \text{Var}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \mid \mathbf{S}] - \text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] \|_2}{\| \text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] \|_2} \leq \| \mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{x}} \|_2. \end{aligned}$$

327 *Proof.* See Section A.5. □

328 Corollary 4.3 implies that the relative differences to conditional unbiasedness
 329 and model variance are solely governed by the quantities $\mathbf{I} - \mathbf{P}_0$ and $\mathbf{P}\mathbf{P}^T - \mathbf{P}_x$,
 330 respectively. Somewhat surprisingly, the condition number of the model variance
 331 $\text{Var}_y[\tilde{\boldsymbol{\beta}}]$ in (4.1) is not explicitly present. Instead, the conditional bias of $\tilde{\boldsymbol{\beta}}$ increases
 332 with the rank deficiency of $\mathbf{S}\mathbf{X}$, while the relative difference between conditional and
 333 model variances increases with the deviation of \mathbf{P} from being an orthogonal projector
 334 onto $\text{range}(\mathbf{X})$. Thus, unbiasedness is easier to achieve because it only requires $\mathbf{S}\mathbf{X}$
 335 to have full column rank, while recovering the model variance requires reproducing
 336 all of $\text{range}(\mathbf{X})$.

337 The examples in Section B.2.1 illustrate the effect of rank deficiency in Theo-
 338 rem 4.2 and Corollary 4.3.

339 *Remark 4.4* (Sampling versus sketching). To confirm the importance of the
 340 induced spaces and the peripheral role of the particular structure of \mathbf{S} , we perform
 341 sketching by first applying row-mixing [1, Section 3.2] with a unitary transform $\mathbf{F} \in$
 342 $\mathbb{R}^{n \times n}$ prior to sampling,

$$343 \quad (4.2) \quad \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{S}(\mathbf{X}\boldsymbol{\beta} - \mathbf{y})\|_2 \quad \text{where} \quad \mathbf{S} \equiv \mathbf{S}_1\mathbf{F},$$

344 where $\mathbf{F}^T\mathbf{F} = \mathbf{F}\mathbf{F}^T = \mathbf{I}_n$, and $\mathbf{S}_1 \in \mathbb{R}^{p \times n}$ is a sampling matrix. The row-mixed
 345 problem

$$346 \quad \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{F}(\mathbf{X}\boldsymbol{\beta} - \mathbf{y})\|_2$$

347 is equivalent to the original problem (2.3), since it has the same solution, and the
 348 same comparison hat matrix and bias projector,

$$349 \quad \mathbf{X}(\mathbf{F}\mathbf{X})^\dagger \mathbf{F} = \mathbf{X}\mathbf{X}^\dagger = \mathbf{P}_x$$

$$350 \quad (\mathbf{F}\mathbf{X})^\dagger (\mathbf{F}\mathbf{X}) = \mathbf{X}^\dagger \mathbf{X} = \mathbf{I}_n.$$

351 Thus, any damaging effect on the conditional bias and variance comes from the possible
 352 rank deficiency and the spaces induced by the sampling process.

353 **4.3. Combined algorithm-induced and model-induced uncertainty.** We
 354 determine the total expectation and the total variance for the solution from (2.7)
 355 when \mathbf{S} is a random sketching matrix, that is, \mathbf{S} is a matrix-valued random variable.

356 The algorithm-induced uncertainty of the random matrix \mathbf{S} is represented by the
 357 expectation $\mathbb{E}_s[\cdot]$ and the variance $\text{Var}_s[\cdot]$, while the total mean and variance of the
 358 combined uncertainty are denoted by $\mathbb{E}[\cdot]$ and $\text{Var}[\cdot]$. The total mean is computed by
 359 conditioning on the algorithm-induced randomness

$$360 \quad (4.3) \quad \mathbb{E}[\cdot] = \mathbb{E}_s \left[\mathbb{E}_y \left[\cdot \mid \mathbf{S} \right] \right].$$

361 Since \mathbf{S} is a matrix-valued random variable, so are the projectors \mathbf{P} and \mathbf{P}_0 .

362 The exact expressions below for general random matrices \mathbf{S} extend the first order
 363 approximations for specific sampling matrices in [16, Lemmas 2-6].

364 **THEOREM 4.5** (Total uncertainty). *With the assumptions in Section 2, let \mathbf{S} be*
 365 *a random sketching matrix. The solution $\tilde{\boldsymbol{\beta}}$ of (2.7) has total expectation and variance*

$$366 \quad \mathbb{E}[\tilde{\boldsymbol{\beta}}] = \mathbb{E}_s[\mathbf{P}_0\boldsymbol{\beta}_0] = \boldsymbol{\beta}_0 + \mathbb{E}_s[\mathbf{P}_0 - \mathbf{I}]\boldsymbol{\beta}_0$$

$$367 \quad \text{Var}[\tilde{\boldsymbol{\beta}}] = \sigma^2 \mathbf{X}^\dagger \mathbb{E}_s[\mathbf{P}\mathbf{P}^T] (\mathbf{X}^\dagger)^T + \text{Var}_s[\mathbf{P}_0\boldsymbol{\beta}_0]$$

$$368 \quad = \text{Var}_y[\hat{\boldsymbol{\beta}}] + \sigma^2 \mathbf{X}^\dagger \mathbb{E}_s[\mathbf{P}\mathbf{P}^T - \mathbf{P}_x] (\mathbf{X}^\dagger)^T + \text{Var}_s[(\mathbf{P}_0 - \mathbf{I})\boldsymbol{\beta}_0],$$

369 *where*

$$\begin{aligned}
370 \quad \text{Var}_{\mathbf{s}}[\mathbf{P}_0\boldsymbol{\beta}_0] &= \mathbb{E}_{\mathbf{s}} \left[(\mathbf{P}_0\boldsymbol{\beta}_0) (\mathbf{P}_0\boldsymbol{\beta}_0)^T \right] - (\mathbb{E}_{\mathbf{s}}[\mathbf{P}_0\boldsymbol{\beta}_0]) (\mathbb{E}_{\mathbf{s}}[\mathbf{P}_0\boldsymbol{\beta}_0])^T \\
371 \quad &= \text{Var}_{\mathbf{s}}[(\mathbf{P}_0 - \mathbf{I})\boldsymbol{\beta}_0].
\end{aligned}$$

372 *Proof.* See Section A.6. □

373 Theorem 4.5 shows that total expectation and variance are governed by the rep-
374 resentation of spaces associated with the original problem (2.3) and the sketched
375 problem (2.7), rather than the specific class of sketching matrices over which $\mathbb{E}_{\mathbf{s}}$ and
376 $\text{Var}_{\mathbf{s}}$ range. Specifically,

- 377 1. The total bias of $\tilde{\boldsymbol{\beta}}$ is proportional to the expected deviation of the matrix-
378 valued random variable $\mathbf{S}\mathbf{X}$ from having full column rank. Note that the
379 expectation $\mathbb{E}_{\mathbf{s}}[\mathbf{P}_0]$ of a projector \mathbf{P}_0 is in general not a projector, as the
380 example in Section B.2.3 illustrates.
- 381 2. The total variance of $\tilde{\boldsymbol{\beta}}$ is proportional to the expected rank deficiency of
382 $\mathbf{S}\mathbf{X}$, plus the expected deviation of the matrix-valued random variable \mathbf{P}
383 from being an orthogonal projector onto $\text{range}(\mathbf{X})$.

384 COROLLARY 4.6 (Relative differences between total and model uncertainties).

385 *With the assumptions in Theorem 4.5,*

$$\begin{aligned}
386 \quad &\|\mathbb{E}[\tilde{\boldsymbol{\beta}}] - \boldsymbol{\beta}_0\|_2 \leq \|\mathbb{E}_{\mathbf{s}}[\mathbf{I} - \mathbf{P}_0]\|_2 \|\boldsymbol{\beta}_0\|_2 \\
387 \quad &\frac{\|\text{Var}[\tilde{\boldsymbol{\beta}}] - \text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]\|_2}{\|\text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]\|_2} \leq \|\mathbb{E}_{\mathbf{s}}[\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{x}}]\|_2 + \frac{\|\text{Var}_{\mathbf{s}}[(\mathbf{I} - \mathbf{P}_0)\boldsymbol{\beta}_0]\|_2}{\|\text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]\|_2}.
\end{aligned}$$

388 *Proof.* See Section A.7. □

389 Corollary 4.6 implies that the relative differences to unbiasedness and model vari-
390 ance are solely governed by the quantities $\mathbb{E}_{\mathbf{s}}[\mathbf{I} - \mathbf{P}_0]$ and $\mathbb{E}_{\mathbf{s}}[\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{x}}]$. Specifically,
391 the total bias of $\tilde{\boldsymbol{\beta}}$ increases with the expected rank deficiency of $\mathbf{S}\mathbf{X}$, while the
392 relative difference between total and model variances increases with the expected de-
393 viation of \mathbf{P} from being an orthogonal projector onto $\text{range}(\mathbf{X})$, and the expected
394 rank deficiency of $\mathbf{S}\mathbf{X}$.

395 The examples in Sections B.2.3-B.2.5 illustrate the effect of expected rank defi-
396 ciency in Theorem 4.5 and Corollary 4.6.

397 **4.4. Regression and residual sums of squares.** Two quantities from the
398 original least squares problem (2.3) play a key role in hypothesis testing, regression
399 diagnostics, and model selection metrics, such as the (adjusted) R^2 statistic, *Mal-*
400 *lows's* C_p , the *Akaike information criterion*, and the *Bayesian information criterion*
401 [7, 17, 20, 21].

- 402 • *Regression sum of squares*, i.e. the squared norm of the prediction,

$$403 \quad \text{SSR}_{\text{ols}} \equiv \mathbf{y}^T \mathbf{P}_{\mathbf{x}} \mathbf{y} = \mathbf{y}^T \mathbf{P}_{\mathbf{x}}^T \mathbf{P}_{\mathbf{x}} \mathbf{y} = \|\hat{\mathbf{y}}\|_2^2.$$

- 404 • *Residual sum of squares*, i.e. the squared norm of the least squares residual,

$$405 \quad \text{RSS}_{\text{ols}} = \mathbf{y}^T (\mathbf{I} - \mathbf{P}_{\mathbf{x}}) \mathbf{y} = \mathbf{y}^T (\mathbf{I} - \mathbf{P}_{\mathbf{x}})^T (\mathbf{I} - \mathbf{P}_{\mathbf{x}}) \mathbf{y} = \|\hat{\mathbf{e}}\|_2^2,$$

406 From $\hat{\mathbf{y}}^T \hat{\mathbf{e}} = 0$ follows

$$407 \quad \|\mathbf{y}\|_2^2 = \|\hat{\mathbf{y}}\|_2^2 + \|\hat{\mathbf{e}}\|_2^2 = \text{SSR}_{\text{ols}} + \text{RSS}_{\text{ols}},$$

408 which decomposes the observation into a portion that is explained by the model; and
 409 a portion that represents the error in the model. The corresponding quantities for
 410 random row-sketching are

$$411 \quad \text{SSR} \equiv \mathbf{y}^T \mathbf{P}^T \mathbf{P} \mathbf{y} = \|\tilde{\mathbf{y}}\|_2^2$$

$$412 \quad \text{RSS} \equiv \mathbf{y}^T (\mathbf{I} - \mathbf{P})^T (\mathbf{I} - \mathbf{P}) \mathbf{y} = \|\tilde{\mathbf{e}}\|_2^2.$$

413 They relate to their counter parts in the original problem (2.3) via the two-norm
 414 version of Theorem 3.3,

$$415 \quad \text{SSR} = \text{SSR}_{\text{ols}} + \mathbf{y}^T (\mathbf{P}^T \mathbf{P} - \mathbf{P}_{\mathbf{x}}) \mathbf{y}$$

$$416 \quad \text{RSS} = \text{RSS}_{\text{ols}} + \|(\mathbf{P} - \mathbf{P}_{\mathbf{x}}) \mathbf{y}\|_2^2.$$

417 Since RSS evaluates the solution $\tilde{\boldsymbol{\beta}}$ of (2.7) in the context of the original problem,
 418 $\tilde{\boldsymbol{\beta}}$ is not a minimizer of (2.3), so clearly $\text{RSS} \geq \text{RSS}_{\text{ols}}$. The difference between the
 419 quantities from random sketching and their deterministic counterparts is governed by
 420 the deviation of \mathbf{P} from being an orthogonal projector onto $\text{range}(\mathbf{X})$.

421 **THEOREM 4.7** (Model-induced uncertainty conditioned on \mathbf{S}). *With the assump-*
 422 *tions in Section 2,*

$$423 \quad \mathbb{E}_{\mathbf{y}}[\text{SSR} \mid \mathbf{S}] = \|\mathbf{P} \mathbf{X} \boldsymbol{\beta}_0\|_2^2 + \sigma^2 \|\mathbf{P}\|_F^2$$

$$424 \quad \mathbb{E}_{\mathbf{y}}[\text{RSS} \mid \mathbf{S}] = \|(\mathbf{I} - \mathbf{P}) \mathbf{X} \boldsymbol{\beta}_0\|_2^2 + \sigma^2 \|\mathbf{I} - \mathbf{P}\|_F^2.$$

425 *Proof.* See Section A.8. □

426 The total expectations follow immediately from Theorem 4.7.

427 **THEOREM 4.8** (Total uncertainty). *With the assumptions in Section 2,*

$$428 \quad \mathbb{E}[\text{SSR}] = (\mathbf{X} \boldsymbol{\beta}_0)^T \mathbb{E}_{\mathbf{s}}[\mathbf{P}^T \mathbf{P}] (\mathbf{X} \boldsymbol{\beta}_0) + \sigma^2 \text{trace}(\mathbb{E}_{\mathbf{s}}[\mathbf{P}^T \mathbf{P}])$$

$$429 \quad \mathbb{E}[\text{RSS}] = (\mathbf{X} \boldsymbol{\beta}_0)^T \mathbb{E}_{\mathbf{s}}[(\mathbf{I} - \mathbf{P})^T (\mathbf{I} - \mathbf{P})] (\mathbf{X} \boldsymbol{\beta}_0) + \sigma^2 \text{trace}(\mathbb{E}_{\mathbf{s}}[(\mathbf{I} - \mathbf{P})^T (\mathbf{I} - \mathbf{P})]).$$

430 At last we show that the difference between combined and model uncertainties
 431 is governed by the expected deviation of \mathbf{P} from being an orthogonal projector onto
 432 $\text{range}(\mathbf{X})$, and the expected deviation of $\mathbf{I} - \mathbf{P}$ from being an orthogonal projector
 433 onto $\text{range}(\mathbf{X})^\perp$, both amplified by the model variance σ^2 .

434 **COROLLARY 4.9** (Difference between total and model uncertainty). *With the*
 435 *assumptions in Section 2,*

$$436 \quad \mathbb{E}[\text{SSR}] - \mathbb{E}_{\mathbf{y}}[\text{SSR}_{\text{ols}}] = (\mathbf{X} \boldsymbol{\beta}_0)^T \mathbb{E}_{\mathbf{s}}[\boldsymbol{\Gamma}] (\mathbf{X} \boldsymbol{\beta}_0) + \sigma^2 \text{trace}(\mathbb{E}_{\mathbf{s}}[\boldsymbol{\Gamma}])$$

$$437 \quad \mathbb{E}[\text{RSS}] - \mathbb{E}_{\mathbf{y}}[\text{RSS}_{\text{ols}}] = (\mathbf{X} \boldsymbol{\beta}_0)^T \mathbb{E}_{\mathbf{s}}[\boldsymbol{\Gamma}_\perp] (\mathbf{X} \boldsymbol{\beta}_0) + \sigma^2 \text{trace}(\mathbb{E}_{\mathbf{s}}[\boldsymbol{\Gamma}_\perp]),$$

438 where we abbreviate

$$439 \quad \boldsymbol{\Gamma} \equiv \mathbf{P}^T \mathbf{P} - \mathbf{P}_{\mathbf{x}}, \quad \boldsymbol{\Gamma}_\perp \equiv (\mathbf{I} - \mathbf{P})^T (\mathbf{I} - \mathbf{P}) - (\mathbf{I} - \mathbf{P}_{\mathbf{x}}).$$

440 *Proof.* See Section A.9. □

441 **5. Discussion.** We considered the randomized solution of least squares regres-
 442 sion problems

$$443 \quad \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{S}(\mathbf{X} \boldsymbol{\beta} - \mathbf{y})\|_2$$

444 arising from a standard Gaussian linear model

$$445 \quad \mathbf{y} = \mathbf{X}\boldsymbol{\beta}_0 + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_n),$$

446 and analyzed the effect on the solution $\tilde{\boldsymbol{\beta}}$ of the combined uncertainties from algo-
447 rithmic randomization and statistical noise.

448 Our results show that the expectation and variance of $\tilde{\boldsymbol{\beta}}$ are governed by the
449 spatial geometry of the sketching process, rather than by structural properties of
450 specific sketching matrices. Surprisingly, the condition number $\kappa_2(\mathbf{X})$ with respect
451 to (left) inversion has far less impact on the statistical measures than it has on the
452 numerical errors. Even from the deterministic view of the sampled problem as a
453 multiplicative perturbation, the relative accuracy of $\tilde{\boldsymbol{\beta}}$ depends only on $\kappa_2(\mathbf{X})$ —rather
454 than on the larger factor $\kappa_2(\mathbf{X})^2$ typical for additive perturbations.

455 The natural next step is the illustration of our analytical results through numer-
456 ical experiments that are representative and informative from both, numerical and
457 statistical perspectives.

458 **Appendix A. Proofs.** We present the proofs for Sections 3 and 4.

459 Our results depend on projectors constructed from the possibly rank-deficient
460 matrix $\mathbf{S}\mathbf{X}$. In this case, the Moore-Penrose inverse cannot be expressed in terms
461 of the matrix $\mathbf{S}\mathbf{X}$ proper, so we rely on the four conditions [8, Section 5.5.2] that
462 uniquely characterize the Moore-Penrose inverse,

$$463 \quad (\mathbf{S}\mathbf{X})(\mathbf{S}\mathbf{X})^\dagger(\mathbf{S}\mathbf{X}) = \mathbf{S}\mathbf{X}, \quad ((\mathbf{S}\mathbf{X})(\mathbf{S}\mathbf{X})^\dagger)^T = (\mathbf{S}\mathbf{X})(\mathbf{S}\mathbf{X})^\dagger$$

$$464 \quad (\mathbf{S}\mathbf{X})^\dagger(\mathbf{S}\mathbf{X})(\mathbf{S}\mathbf{X})^\dagger = (\mathbf{S}\mathbf{X})^\dagger, \quad ((\mathbf{S}\mathbf{X})^\dagger(\mathbf{S}\mathbf{X}))^T = (\mathbf{S}\mathbf{X})^\dagger(\mathbf{S}\mathbf{X}).$$

465 **A.1. Proof of Lemma 3.1.** The Moore-Penrose conditions [8, Section 5.5.2]
466 imply $\mathbf{P}^2 = \mathbf{P}$ for the generally nonsymmetric matrix \mathbf{P} .

- 467 1. This follows from the Moore-Penrose conditions (A.1).
- 468 2. Use the fact [19, Problem 5.9.12] that $\text{null}(\mathbf{P}) = \text{null}(\mathbf{P}_x)$ if and only if
469 $\mathbf{P}\mathbf{P}_x - \mathbf{P} = \mathbf{0}$ and $\mathbf{P}_x\mathbf{P} - \mathbf{P}_x = \mathbf{0}$. With item 1, this implies $\mathbf{P}_x\mathbf{P} - \mathbf{P}_x =$
470 $\mathbf{P} - \mathbf{P}_x$. Thus $\mathbf{P} - \mathbf{P}_x$ can be interpreted as a measure for the distance
471 between $\text{null}(\mathbf{P})$ and $\text{null}(\mathbf{P}_x)$.
- 472 3. This follows from (2.1).

473 **A.2. Proof of Theorem 3.3.** The first expression for $\tilde{\boldsymbol{\beta}}$ follows from (2.1),
474 (2.8), and Lemma 3.1. The second expression follows from adding and subtracting in
475 the first expression the term $\hat{\boldsymbol{\beta}} = \mathbf{X}^\dagger \mathbf{y} = \mathbf{X}^\dagger \mathbf{P}_x \mathbf{y}$.

476 The first expression for $\tilde{\mathbf{y}}$ follows from (2.8) and Lemma 3.1. The second ex-
477 pression follows from adding and subtracting in the first expression the first term in
478 (2.6).

479 The first expression for $\tilde{\mathbf{e}}$ follows from (2.9), (2.8) and Lemma 3.1. The second
480 expression for $\tilde{\mathbf{e}}$ follows from adding and subtracting in the first expression the second
481 term in (2.6).

482 **A.3. Proof of Lemma 4.1.** The Moore-Penrose conditions (A.1) imply $(\mathbf{P}_0)^2 =$
483 \mathbf{P}_0 and $(\mathbf{P}_0)^T = \mathbf{P}_0$, confirming that \mathbf{P}_0 is an orthogonal projector.

- 484 1. This follows from Lemma 3.1.
- 485 2. If $\text{rank}(\mathbf{S}\mathbf{X}) = p$, then (2.1) implies that $(\mathbf{S}\mathbf{X})^\dagger$ is a left-inverse.

486 **A.4. Proof of Theorem 4.2.** The conditional expectation follows from Theo-
487 rem 3.3, (4.1), Lemma 4.1, and (2.1).

488 The definition of variance, Theorem 3.3, and the above imply

$$489 \quad \mathbb{V}\text{ar}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \mid \mathbf{S}] = \mathbb{E}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}}\tilde{\boldsymbol{\beta}}^T \mid \mathbf{S}] - \mathbb{E}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \mid \mathbf{S}] \mathbb{E}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \mid \mathbf{S}]^T$$

$$490 \quad = (\mathbf{X}^\dagger \mathbf{P}) \mathbb{E}_{\mathbf{y}}[\mathbf{y}\mathbf{y}^T] (\mathbf{X}^\dagger \mathbf{P})^T - (\mathbf{P}_0 \boldsymbol{\beta}_0)(\mathbf{P}_0 \boldsymbol{\beta}_0)^T.$$

491 The middle term in the first summand equals

$$492 \quad \mathbb{E}_{\mathbf{y}}[\mathbf{y}\mathbf{y}^T] = (\mathbf{X}\boldsymbol{\beta}_0)(\mathbf{X}\boldsymbol{\beta}_0)^T + \mathbf{X}\boldsymbol{\beta}_0 \mathbb{E}_{\mathbf{y}}[\boldsymbol{\epsilon}]^T + \mathbb{E}_{\mathbf{y}}[\boldsymbol{\epsilon}](\mathbf{X}\boldsymbol{\beta}_0)^T + \mathbb{E}_{\mathbf{y}}[\boldsymbol{\epsilon}\boldsymbol{\epsilon}^T]$$

$$493 \quad (\text{A.2}) \quad = (\mathbf{X}\boldsymbol{\beta}_0)(\mathbf{X}\boldsymbol{\beta}_0)^T + \sigma^2 \mathbf{I}_n.$$

494 To obtain the first expression for the conditional variance, insert (A.2) into the con-
495 ditional variance above, and apply Lemma 4.1 to cancel out the expressions with \mathbf{P}_0 .

496 For the second expression, use (2.1) and (2.5) to write the model variance in (4.1)
497 as

$$498 \quad \mathbb{V}\text{ar}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] = \sigma^2 \mathbf{X}^\dagger \mathbf{P}_{\mathbf{x}} (\mathbf{X}^\dagger)^T.$$

499 Then add and subtract this term in the first expression for the conditional variance.

500 If \mathbf{P} were an orthogonal projector onto $\text{range}(\mathbf{X})$, then $\mathbf{P}^T \mathbf{P} = \mathbf{P} = \mathbf{P}_{\mathbf{x}}$. Thus,
501 $\mathbf{P}^T \mathbf{P} - \mathbf{P}_{\mathbf{x}}$ represents the deviation of \mathbf{P} from being an orthogonal projector onto
502 $\text{range}(\mathbf{X})$.

503 **A.5. Proof of Corollary 4.3.** The bound for the conditional expectation fol-
504 lows from (4.1), and the second expression for the expectation in Theorem 4.2. The
505 second expression for the conditional variance in Theorem 4.2 implies

$$506 \quad \|\mathbb{V}\text{ar}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \mid \mathbf{S}] - \mathbb{V}\text{ar}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]\|_2 \leq \sigma^2 \|\mathbf{X}^\dagger\|_2 \|\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{x}}\|_2 \|(\mathbf{X}^\dagger)^T\|_2.$$

507 Now apply $\|\mathbf{M}\|_2 \|\mathbf{M}^T\|_2 = \|\mathbf{M}\mathbf{M}^T\|_2$, and $\mathbf{M}^\dagger (\mathbf{M}^\dagger)^T = (\mathbf{M}^T \mathbf{M})^{-1}$ for a full column-
508 rank matrix \mathbf{M} to deduce

$$509 \quad (\text{A.3}) \quad \sigma^2 \|\mathbf{X}^\dagger\|_2 \|(\mathbf{X}^\dagger)^T\|_2 = \|\mathbb{V}\text{ar}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]\|_2,$$

510 where $\|\mathbb{V}\text{ar}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]\|_2 \neq 0$ by assumption in Section 2.1.

511 **A.6. Proof of Theorem 4.5.** Apply the iterated expectation (4.3), followed by
512 Theorem 4.2 to obtain the mean,

$$513 \quad \mathbb{E}[\tilde{\boldsymbol{\beta}}] = \mathbb{E}_{\mathbf{s}} \left[\mathbb{E}_{\mathbf{y}} \left[\tilde{\boldsymbol{\beta}} \mid \mathbf{S} \right] \right] = \mathbb{E}_{\mathbf{s}}[\mathbf{P}_0 \boldsymbol{\beta}_0] = \mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] \boldsymbol{\beta}_0.$$

514 Insert this into the definition of the variance, and apply again (4.3),

$$515 \quad \mathbb{V}\text{ar}[\tilde{\boldsymbol{\beta}}] = \mathbb{E}[\tilde{\boldsymbol{\beta}}\tilde{\boldsymbol{\beta}}^T] - \mathbb{E}[\tilde{\boldsymbol{\beta}}] \mathbb{E}[\tilde{\boldsymbol{\beta}}]^T$$

$$516 \quad = \mathbb{E}_{\mathbf{s}} \left[\mathbb{E}_{\mathbf{y}} \left[\tilde{\boldsymbol{\beta}}\tilde{\boldsymbol{\beta}}^T \mid \mathbf{S} \right] \right] - (\mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] \boldsymbol{\beta}_0) (\mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] \boldsymbol{\beta}_0)^T.$$

517 Treat the first summand as in the proof of Theorem 4.2 in Section A.4 to deduce

$$518 \quad \mathbb{E}_{\mathbf{y}} \left[\tilde{\boldsymbol{\beta}}\tilde{\boldsymbol{\beta}}^T \mid \mathbf{S} \right] = \sigma^2 \mathbf{X}^\dagger \mathbf{P} \mathbf{P}^T (\mathbf{X}^\dagger)^T + (\mathbf{P}_0 \boldsymbol{\beta}_0)(\mathbf{P}_0 \boldsymbol{\beta}_0)^T.$$

519 Condition this on \mathbf{y} and then insert it into the above expression for the variance,

$$520 \quad \mathbb{V}\text{ar}[\tilde{\boldsymbol{\beta}}] = \sigma^2 \mathbf{X}^\dagger \mathbb{E}_{\mathbf{s}} \left[\mathbf{P} \mathbf{P}^T \right] (\mathbf{X}^\dagger)^T$$

$$521 \quad + \underbrace{\mathbb{E}_{\mathbf{s}} \left[(\mathbf{P}_0 \boldsymbol{\beta}_0) (\mathbf{P}_0 \boldsymbol{\beta}_0)^T \right] - (\mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] \boldsymbol{\beta}_0) (\mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] \boldsymbol{\beta}_0)^T}_{\mathbb{V}\text{ar}_{\mathbf{s}}[\mathbf{P}_0 \boldsymbol{\beta}_0]}.$$

522 The second expression for $\text{Var}_s[\mathbf{P}_0\boldsymbol{\beta}_0]$ follows from adding and subtracting

$$523 \quad \boldsymbol{\beta}_0\boldsymbol{\beta}_0^T - \boldsymbol{\beta}_0(\mathbb{E}_s[\mathbf{P}_0]\boldsymbol{\beta}_0)^T - (\mathbb{E}_s[\mathbf{P}_0]\boldsymbol{\beta}_0)\boldsymbol{\beta}_0^T,$$

524 in other words from $\boldsymbol{\beta}_0$ having zero variance.

525 **A.7. Proof of Corollary 4.6.** The bound for the total expectation follows
 526 from (4.1), and the second expression for the expectation in Theorem 4.5. The bound
 527 for the total variance follows from the second expression for the variance in Theo-
 528 rem 4.5, and from (A.3).

529 **A.8. Proof of Theorem 4.7.** We need the following auxiliary result about
 530 expectations of quadratic forms.

531 **LEMMA A.1.** *With the assumptions in Section 2, if $\mathbf{A} \in \mathbb{R}^{n \times n}$ is a constant*
 532 *matrix, then,*

$$533 \quad \mathbb{E}[\mathbf{y}^T \mathbf{A} \mathbf{y}] = (\mathbf{X}\boldsymbol{\beta}_0)^T \mathbf{A} (\mathbf{X}\boldsymbol{\beta}_0) + \sigma^2 \text{trace}(\mathbf{A}).$$

534 *Proof.* This follows from $\mathbf{y}^T \mathbf{A} \mathbf{y}$ being a real scalar, the circular commutativity
 535 of the trace, the interchangeability of the trace and expectation since both are sums,
 536 and (A.2) as follows,

$$537 \quad \begin{aligned} \mathbb{E}[\mathbf{y}^T \mathbf{A} \mathbf{y}] &= \mathbb{E}[\text{trace}(\mathbf{y}^T \mathbf{A} \mathbf{y})] = \mathbb{E}[\text{trace}(\mathbf{A} \mathbf{y} \mathbf{y}^T)] = \text{trace}(\mathbf{A} \mathbb{E}[\mathbf{y} \mathbf{y}^T]) \\ 538 \quad &= \text{trace}(\mathbf{A} (\mathbf{X}\boldsymbol{\beta}_0)(\mathbf{X}\boldsymbol{\beta}_0)^T + \sigma^2 \mathbf{A}) = (\mathbf{X}\boldsymbol{\beta}_0)^T \mathbf{A} (\mathbf{X}\boldsymbol{\beta}_0) + \sigma^2 \text{trace}(\mathbf{A}). \end{aligned}$$

539

□

540 **Proof of the Theorem.** The expression for SSR follows from $\tilde{\mathbf{y}} = \mathbf{P} \mathbf{y}$ in Theo-
 541 rem 3.3, Lemma A.1, and $\text{trace}(\mathbf{P}^T \mathbf{P}) = \|\mathbf{P}\|_F^2$. Analogously, the expression for RSS
 542 follows from $\tilde{\mathbf{e}} = (\mathbf{I} - \mathbf{P}) \mathbf{y}$.

543 **A.9. Proof of Corollary 4.9.** From (2.6) and Lemma A.1 follows

$$544 \quad \begin{aligned} \mathbb{E}_y[\mathbf{y}^T \mathbf{P}_x \mathbf{y}] &= (\mathbf{X}\boldsymbol{\beta}_0)^T \mathbf{P}_x (\mathbf{X}\boldsymbol{\beta}_0) + \sigma^2 \text{trace}(\mathbf{P}_x) \\ 545 \quad \mathbb{E}_y[\mathbf{y}^T (\mathbf{I} - \mathbf{P}_x) \mathbf{y}] &= \underbrace{(\mathbf{X}\boldsymbol{\beta}_0)^T (\mathbf{I} - \mathbf{P}_x) (\mathbf{X}\boldsymbol{\beta}_0)}_0 + \sigma^2 \text{trace}(\mathbf{I} - \mathbf{P}_x). \end{aligned}$$

546 Add and subtract these to the respective expressions in Theorem 4.8.

547 **Appendix B. Examples with uniform row sampling.** We start with a
 548 brief review of sketching matrices for least squares problems (Section B.1), before
 549 presenting examples that give insight into the results of Section 4 and the detrimental
 550 effects of rank deficiency (Section B.2).

551 **B.1. Random sketching matrices in least squares.** We present a few ex-
 552 amples of sketching matrices used by the randomized least squares solvers [1, 2, 5, 6,
 553 14, 15, 16, 18, 23].

554 *Uniform sampling with replacement.* This is the *EXACTLY(c)* algorithm [6, Al-
 555 gorithm 3] with uniform probabilities, which performs row-wise compression for direct
 556 methods for the solution of full column rank least squares in [6, Algorithm 3], see also
 557 the *BasicMatrixMultiplication Algorithm* [4, Fig. 2], [13, Algorithm 3.2], [14, Algo-
 558 rithms 1 and 2], and the *Uniform Sampling Estimator* [16, Section 2.2].

559 The probability of a particular instance of $\text{diag}(\mathbf{S}^T \mathbf{S})$, and therefore \mathbf{S} is given by
 560 a scaled multinomial distribution [16, Section 3.1].

Algorithm B.1 Uniform sampling with replacement**Input:** Integers $n \geq 1$ and $1 \leq r \leq n$ **Output:** Sampling matrix $\mathbf{S} \in \mathbb{R}^{r \times n}$ with $\mathbb{E}_{\mathbf{s}}[\mathbf{S}^T \mathbf{S}] = \mathbf{I}_n$ **for** $t = 1 : r$ **do** Sample k_t from $\{1, \dots, n\}$ with probability $1/n$,
 independently and with replacement**end for**

$$\mathbf{S} = \sqrt{\frac{n}{r}} (\mathbf{e}_{k_1} \quad \dots \quad \mathbf{e}_{k_r})^T$$

561 *Random orthogonal sketching.* This is used in *Blendenpik* [1, Algorithm 1] to
 562 compute randomized preconditioners for the iterative solution of full column rank
 563 least squares problems.

564 Here $\mathbf{S} = \mathbf{BTD} \in \mathbb{R}^{n \times n}$, where $\mathbf{D} \in \mathbb{R}^{n \times n}$ is a diagonal matrix whose diagonal
 565 elements are independent Rademacher random variables, equaling ± 1 with equal prob-
 566 ability; $\mathbf{T} \in \mathbb{R}^{n \times n}$ is a unitary matrix, such as a Walsh-Hadamard, discrete cosine, or
 567 discrete Hartley transform; and \mathbf{B} is a diagonal matrix whose diagonal elements are
 568 Bernoulli variables, equaling 1 with probability $\gamma p/n$ for some $\gamma > 0$, and 0 otherwise.

569 *Gaussian sketching.* This is used in to compute randomized preconditioners for the
 570 iterative solution of general least squares problems [18, Algorithms 1 and 2].

571 Here the elements of $\mathbf{S} \in \mathbb{R}^{r \times n}$ are independent $\mathcal{N}(0, 1)$ random variables. In
 572 Matlab: $\mathbf{S} = \text{randn}(r, n)$.

573 **B.2. Examples.** The purpose is to provide insight for Theorem 4.2, Corol-
 574 lary 4.3, Theorem 4.5 and Corollary 4.6 in a way that is easy to reproduce. For
 575 a small example matrix, we illustrate the effect of rank deficiency $\mathbf{S}\mathbf{X}$ (Section B.2.1);
 576 perform uniform sampling with replacement (Section B.2.2); compute the expecta-
 577 tions for \mathbf{P}_0 (Section B.2.3) and $\mathbf{P}\mathbf{P}^T$ (Section B.2.4); and put this into context with
 578 two matrices \mathbf{S} at opposite ends of sampling performance (Section B.2.5).

579 Our example is the full column-rank matrix

$$580 \quad \mathbf{X} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 0 \end{pmatrix} \in \mathbb{R}^{4 \times 2} \quad \text{with} \quad \mathbf{X}^\dagger = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix},$$

581 and $\text{rank}(\mathbf{X}) = 2$. The hat matrix (2.5) and its null space are

$$582 \quad \mathbf{P}_{\mathbf{x}} = \mathbf{X}\mathbf{X}^\dagger = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 1 & 0 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \text{null}(\mathbf{P}_{\mathbf{x}}) = \text{range} \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ -1 & 0 \\ 0 & 1 \end{pmatrix}$$

583 while the model variance (4.1) is

$$584 \quad (\text{B.1}) \quad \text{Var}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}}] = \sigma^2 (\mathbf{X}^T \mathbf{X})^{-1} = \sigma^2 \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 1 \end{pmatrix}.$$

585 **B.2.1. Effect of rank deficiency in Theorem 4.2 and Corollary 4.3.** We
 586 choose two different matrices \mathbf{S} with full row-rank $\text{rank}(\mathbf{S}) = 2$, one producing a full
 587 rank $\mathbf{S}\mathbf{X}$, and the other one a rank-deficient $\mathbf{S}\mathbf{X}$.

588 1. *Full column-rank $\mathbf{S}\mathbf{X}$.* The sketching matrix is

589
$$\mathbf{S} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \quad \text{where} \quad \mathbf{S}\mathbf{X} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = (\mathbf{S}\mathbf{X})^\dagger = \mathbf{I}_2,$$

590 $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X}) = 2$. The comparison hat matrix in Lemma 3.1 and the bias
591 projector in Lemma 4.1 are

592
$$\mathbf{P} = \mathbf{X}(\mathbf{S}\mathbf{X})^\dagger \mathbf{S} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{P}_0 = (\mathbf{S}\mathbf{X})^\dagger (\mathbf{S}\mathbf{X}) = \mathbf{I}_2.$$

593 Thus $\text{range}(\mathbf{P}) = \text{range}(\mathbf{X})$. The deviation of \mathbf{P} from being an orthogonal projector
594 onto $\text{range}(\mathbf{X})$ is

595
$$\mathbf{P}\mathbf{P}^T - \mathbf{P}_x = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \|\mathbf{P}\mathbf{P}^T - \mathbf{P}_x\|_2 = 1.$$

596 Thus, the solution $\tilde{\boldsymbol{\beta}}$ of (2.7) is an unbiased estimator, but with increased variance.
597 Specifically,

- 598 • \mathbf{P} is a projector onto $\text{range}(\mathbf{X})$, but it is not an orthogonal projector, since
599 \mathbf{P} is not symmetric.
- 600 • The conditional expectation of $\tilde{\boldsymbol{\beta}}$ is $\mathbb{E}_y[\tilde{\boldsymbol{\beta}} | \mathbf{S}] = \boldsymbol{\beta}_0$, since $\mathbf{P}_0 = \mathbf{I}_2$, and the
601 corresponding bound in Corollary 4.3 holds with equality.
- 602 • The conditional variance has increased compared to (B.1), because

603
$$\text{Var}_y[\tilde{\boldsymbol{\beta}} | \mathbf{S}] = \sigma^2 \mathbf{X}^\dagger \mathbf{P}\mathbf{P}^T (\mathbf{X}^\dagger)^T = \sigma^2 \mathbf{I}_2.$$

604 In the worst case, it has zero norm-wise relative accuracy since

605
$$\|\text{Var}_y[\tilde{\boldsymbol{\beta}} | \mathbf{S}] - \text{Var}_y[\hat{\boldsymbol{\beta}}]\|_2 / \|\text{Var}_y[\hat{\boldsymbol{\beta}}]\|_2 = \frac{1}{2} \leq \|\mathbf{P}\mathbf{P}^T - \mathbf{P}_x\|_2 = 1.$$

606 2. *Rank deficient $\mathbf{S}\mathbf{X}$.* The sketching matrix is

607
$$\mathbf{S} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad \text{where} \quad \mathbf{S}\mathbf{X} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = (\mathbf{S}\mathbf{X})^\dagger,$$

608 $\text{rank}(\mathbf{S}\mathbf{X}) = 1 < \text{rank}(\mathbf{X})$, and $\text{range}(\mathbf{P}) \subset \text{range}(\mathbf{X})$. The comparison hat matrix in
609 Lemma 3.1 and the bias projector in Lemma 4.1 are

610
$$\mathbf{P} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{P}_0 = (\mathbf{S}\mathbf{X})^\dagger (\mathbf{S}\mathbf{X}) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

611 The deviation of \mathbf{P} from being an orthogonal projector onto $\text{range}(\mathbf{X})$ is

612
$$\mathbf{P}\mathbf{P}^T - \mathbf{P}_x = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & -1 & 0 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \|\mathbf{P}\mathbf{P}^T - \mathbf{P}_x\|_2 = 1,$$

613 and the rank deficiency of $\mathbf{S}\mathbf{X}$ is represented by $\|\mathbf{I} - \mathbf{P}_0\|_2 = 1$. Thus, the solution $\tilde{\boldsymbol{\beta}}$
614 of (2.7) is a biased estimator with a conditional variance that is singular. Specifically,

- 615 • Although \mathbf{P} is a projector, it is not an orthogonal projector onto $\text{range}(\mathbf{X})$,
616 since \mathbf{P} is not symmetric and it projects only onto a lower-dimensional sub-
617 space of $\text{range}(\mathbf{X})$.
- 618 • The conditional expectation of $\tilde{\boldsymbol{\beta}}$ is $\mathbb{E}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} | \mathbf{S}] \neq \boldsymbol{\beta}_0$, since $\mathbf{P}_0 \neq \mathbf{I}_2$, and the
619 relative distance to unbiasedness can be maximal in the worst case, since
620 $\|\mathbb{E}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} | \mathbf{S}] - \boldsymbol{\beta}_0\|_2 \leq \|\boldsymbol{\beta}_0\|_2$.
- 621 • The conditional variance has become singular,

$$622 \quad \text{Var}_{\mathbf{y}} \left[\tilde{\boldsymbol{\beta}} \mid \mathbf{S} \right] = \sigma^2 \mathbf{X}^\dagger \mathbf{P} \mathbf{P}^\top (\mathbf{X}^\dagger)^T = \sigma^2 \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix},$$

623 with zero norm-wise relative accuracy, and the corresponding bound holds
624 with equality,

$$625 \quad \|\text{Var}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} | \mathbf{S}] - \text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]\|_2 / \|\text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]\|_2 = 1 = \|\mathbf{P} \mathbf{P}^\top - \mathbf{P}_{\mathbf{x}}\|_2.$$

626 **B.2.2. Uniform sampling with replacement.** Algorithm B.1 with $n = 4$ and
627 $r = 2$ produces a sampling matrix $\mathbf{S} \in \mathbb{R}^{2 \times 4}$, which has $n^2 = 16$ instances

$$628 \quad \mathbf{S}_{ij} = \sqrt{2} \begin{pmatrix} \mathbf{e}_i^T \\ \mathbf{e}_j^T \end{pmatrix}, \quad 1 \leq i, j \leq n,$$

629 each occurring with probability $1/n^2$. For instance,

$$630 \quad \mathbf{S}_{11} = \sqrt{2} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{S}_{42} = \sqrt{2} \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}.$$

631 The expectation of the Gram product is an unbiased estimator of the identity, since

$$632 \quad \mathbb{E}_{\mathbf{s}}[\mathbf{S}^T \mathbf{S}] = \sum_{i=1}^4 \sum_{j=1}^4 \frac{1}{16} \mathbf{S}_{ij}^T \mathbf{S}_{ij} = \sum_{i=1}^4 \sum_{j=1}^4 \frac{1}{16} (\mathbf{e}_i \mathbf{e}_i^T + \mathbf{e}_j \mathbf{e}_j^T) = \mathbf{I}_4.$$

633 **B.2.3. Expected rank deficiency in Theorem 4.5 and Corollary 4.6.** The
634 total expectation of $\mathbf{P}_0 \in \mathbb{R}^{2 \times 2}$ is

$$635 \quad \mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] = \sum_{i=1}^4 \sum_{j=1}^4 \frac{1}{16} (\mathbf{S}_{ij} \mathbf{X})^\dagger (\mathbf{S}_{ij} \mathbf{X}) = \mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] = \frac{1}{16} \begin{pmatrix} 12 & 0 \\ 0 & 7 \end{pmatrix}.$$

636 For instance, representative summands include

$$637 \quad (\mathbf{S}_{13} \mathbf{X})^\dagger = \sqrt{\frac{1}{2}} \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix}^\dagger = \sqrt{\frac{1}{2}} \begin{pmatrix} 1/2 & 1/2 \\ 0 & 0 \end{pmatrix}, \quad (\mathbf{S}_{13} \mathbf{X})^\dagger (\mathbf{S}_{13} \mathbf{X}) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix},$$

$$638 \quad (\mathbf{S}_{32} \mathbf{X})^\dagger = \sqrt{\frac{1}{2}} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}^\dagger = \sqrt{\frac{1}{2}} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = \sqrt{\frac{1}{2}} \mathbf{I}_2, \quad (\mathbf{S}_{32} \mathbf{X})^\dagger (\mathbf{S}_{32} \mathbf{X}) = \mathbf{I}_2,$$

$$639 \quad (\mathbf{S}_{44} \mathbf{X})^\dagger = \sqrt{\frac{1}{2}} \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}^\dagger = \mathbf{0}, \quad (\mathbf{S}_{44} \mathbf{X})^\dagger (\mathbf{S}_{44} \mathbf{X}) = \mathbf{0}.$$

640 Among the sketched matrices $\mathbf{S} \mathbf{X}$, 75 percent are rank deficient. The ones with full
641 column rank are $\mathbf{S}_{12} \mathbf{X}$, $\mathbf{S}_{21} \mathbf{X}$, $\mathbf{S}_{23} \mathbf{X}$, and $\mathbf{S}_{32} \mathbf{X}$. The expected rank deficiency of $\mathbf{S} \mathbf{X}$
642 equals

$$643 \quad \mathbb{E}_{\mathbf{s}}[\mathbf{I} - \mathbf{P}_0] = \frac{1}{16} \begin{pmatrix} 4 & 0 \\ 0 & 9 \end{pmatrix} \quad \text{with} \quad \|\mathbb{E}_{\mathbf{s}}[\mathbf{I} - \mathbf{P}_0]\|_2 = \frac{9}{16}.$$

644 Thus, the solution $\tilde{\boldsymbol{\beta}}$ of (2.7) is a biased estimator. Specifically,

- 645 • $\mathbb{E}_s[\mathbf{P}_0]$ is not a projector, since it is not idempotent.
- 646 • The total expectation of $\tilde{\boldsymbol{\beta}}$ equals $\mathbb{E}_s[\tilde{\boldsymbol{\beta}}] \neq \boldsymbol{\beta}_0$, since $\mathbb{E}_s[\mathbf{P}_0] \neq \mathbf{I}_2$. and the
- 647 relative distance to unbiasedness can be large, since $\|\mathbb{E}[\tilde{\boldsymbol{\beta}}] - \boldsymbol{\beta}_0\|_2 \leq \frac{9}{16} \|\boldsymbol{\beta}_0\|_2$.

648 **B.2.4. Expected deviation of \mathbf{P} from being an orthogonal projector in**
 649 **Theorem 4.5 and Corollary 4.6.** To the expectation of $\mathbf{P}\mathbf{P}^T \in \mathbb{R}^{4 \times 4}$, note that
 650 the trailing column of \mathbf{X} is zero, and

$$651 \quad \mathbf{P}\mathbf{P}^T = \mathbf{X}(\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{S}^T ((\mathbf{S}\mathbf{X})^\dagger)^T \mathbf{X}^T,$$

652 the trailing row and column of all instances of $\mathbf{P}\mathbf{P}^T$ and $\mathbb{E}_s[\mathbf{P}\mathbf{P}^T]$ are zero as well,
 653 and

$$654 \quad \mathbb{E}_s[\mathbf{P}\mathbf{P}^T] = \sum_{i=1}^4 \sum_{j=1}^4 \frac{1}{16} \mathbf{X}(\mathbf{S}_{ij}\mathbf{X})^\dagger \mathbf{S}_{ij} \mathbf{S}_{ij}^T ((\mathbf{S}_{ij}\mathbf{X})^\dagger)^T \mathbf{X}^T = \frac{1}{16} \begin{pmatrix} 11 & 0 & 11 & 0 \\ 0 & 7 & 0 & 0 \\ 11 & 0 & 11 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

655 Thus, $\mathbb{E}_s[\mathbf{P}\mathbf{P}^T]$ is not a projector since it is not idempotent, and the expected devi-
 656 ation of \mathbf{P} from being an orthogonal projector onto $\text{range}(\mathbf{X})$ can be larger than 50
 657 percent, since

$$658 \quad \mathbb{E}_s[\mathbf{P}\mathbf{P}^T - \mathbf{P}_x] = \frac{1}{16} \begin{pmatrix} 3 & 0 & 3 & 0 \\ 0 & -9 & 0 & 0 \\ 3 & 0 & 3 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad \text{with} \quad \|\mathbb{E}_s[\mathbf{P}\mathbf{P}^T - \mathbf{P}_x]\|_2 = \frac{9}{16}.$$

659 **B.2.5. Extreme examples.** We consider two more 4×2 matrices, both with
 660 orthogonal columns, but at the opposite ends in terms of the performance for uniform
 661 sampling in Section B.2.2.

662 *Columns of the Hadamard matrix.* With its mass spread uniformly spread, which
 663 is quantified by minimal coherence and uniform leverage scores [13, 16], this matrix
 664 is optimal for uniform row sampling,

$$665 \quad \mathbf{X} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \\ 1 & -1 \end{pmatrix} \in \mathbb{R}^{4 \times 2}, \quad \mathbf{P}_x = \mathbf{X}\mathbf{X}^\dagger = \frac{1}{2} \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}.$$

666 Half of the sketched matrices $\mathbf{S}\mathbf{X}$ have full column rank. The expectations for the
 667 projectors are

$$668 \quad \mathbb{E}_s[\mathbf{P}_0] = \frac{12}{16} \mathbf{I}_2, \quad \mathbb{E}_s[\mathbf{P}\mathbf{P}^T] = \frac{11}{16} \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}.$$

669 Thus the expected deviation of $\mathbf{S}\mathbf{X}$ from full column-rank rank, and the expected
 670 deviation of \mathbf{P} from being an orthogonal projector onto $\text{range}(\mathbf{X})$ are

$$671 \quad \|\mathbb{E}_s[\mathbf{I} - \mathbf{P}_0]\|_2 = \frac{4}{16}, \quad \|\mathbb{E}_s[\mathbf{P}\mathbf{P}^T - \mathbf{P}_x]\|_2 = \frac{3}{16},$$

672 and clearly lower, and therefore better than the respective ones in Sections B.2.3
 673 and B.2.4.

674 *Columns of the identity matrix.* With its concentrated mass spread, which is
 675 quantified by maximal coherence and widely differing leverage scores [13, 16], this
 676 matrix presents a worst case for uniform row sampling of 4×2 a full column-rank
 677 matrix,

$$678 \quad \mathbf{X} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{pmatrix} \in \mathbb{R}^{4 \times 2}, \quad \mathbf{P}_{\mathbf{x}} = \mathbf{X}\mathbf{X}^\dagger = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

679 Only two among the 16 sketched matrices $\mathbf{S}\mathbf{X}$ have full column rank, $\mathbf{S}_{12}\mathbf{X}$ and $\mathbf{S}_{21}\mathbf{X}$.
 680 The expectations for the projectors are

$$681 \quad \mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] = \frac{7}{16}\mathbf{I}_2, \quad \mathbb{E}_{\mathbf{s}}[\mathbf{P}\mathbf{P}^T] = \frac{7}{16} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

682 The expected deviations of $\mathbf{S}\mathbf{X}$ from full column-rank and of \mathbf{P} from being an orthog-
 683 onal projector onto $\text{range}(\mathbf{X})$ are

$$684 \quad \|\mathbb{E}_{\mathbf{s}}[\mathbf{I} - \mathbf{P}_0]\|_2 = \frac{9}{16}, \quad \|\mathbb{E}_{\mathbf{s}}[\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{x}}]\|_2 = \frac{9}{16},$$

685 thus clearly worse than those for the Hadamard matrix.

686 **Acknowledgements.** We are grateful to Dennis Boos and Chris Waddell for
 687 many helpful discussions.

688 REFERENCES

- 689 [1] H. AVRON, P. MAYMOUNKOV, AND S. TOLEDO, *Blendenpik: supercharging Lapack's least-*
 690 *squares solver*, SIAM J. Sci. Comput., 32 (2010), pp. 1217–1236.
 691 [2] C. BOUTSIDIS AND P. DRINEAS, *Random projections for the nonnegative least-squares problem*,
 692 Linear Algebra Appl., 431 (2009), pp. 760–771.
 693 [3] S. CHATTERJEE AND A. S. HADI, *Influential observations, high leverage points, and outliers in*
 694 *linear regression*, Statist. Sci., 1 (1986), pp. 379–416. With discussion.
 695 [4] P. DRINEAS, R. KANNAN, AND M. W. MAHONEY, *Fast Monte Carlo Algorithms for Matrices.*
 696 *I: Approximating Matrix Multiplication*, SIAM J. Comput., 36 (2006), pp. 132–157.
 697 [5] P. DRINEAS, M. W. MAHONEY, AND S. MUTHUKRISHNAN, *Sampling algorithms for l_2 regression*
 698 *and applications*, in Proceedings of the Seventeenth Annual ACM-SIAM Symposium on
 699 Discrete Algorithms, ACM, New York, 2006, pp. 1127–1136.
 700 [6] P. DRINEAS, M. W. MAHONEY, S. MUTHUKRISHNAN, AND T. SARLÓS, *Faster least squares*
 701 *approximation*, Numer. Math., 117 (2011), pp. 219–249.
 702 [7] J. FRIEDMAN, T. HASTIE, AND R. TIBSHIRANI, *The elements of statistical learning*, vol. 1,
 703 Springer series in statistics New York, 2001.
 704 [8] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The Johns Hopkins University
 705 Press, Baltimore, fourth ed., 2013.
 706 [9] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, sec-
 707 ond ed., 2002.
 708 [10] D. C. HOAGLIN AND R. E. WELSCH, *The Hat matrix in regression and ANOVA*, Amer. Statist.,
 709 32 (1978), pp. 17–22.
 710 [11] I. C. F. IPSEN, *Relative perturbation results for matrix eigenvalues and singular values*, in Acta
 711 Numerica 1998, vol. 7, Cambridge University Press, Cambridge, 1998, pp. 151–201.
 712 [12] I. C. F. IPSEN, *An overview of relative $\sin \Theta$ theorems for invariant subspaces of complex*
 713 *matrices*, J. Comput. Appl. Math., 123 (2000), pp. 131–153. Invited Paper for the special
 714 issue *Numerical Analysis 2000: Vol. III – Linear Algebra*.
 715 [13] I. C. F. IPSEN AND T. WENTWORTH, *The effect of coherence on sampling from matrices with*
 716 *orthonormal columns, and preconditioned least squares problems*, SIAM J. Matrix Anal.
 717 Appl., 35 (2014), pp. 1490–1520.

- 718 [14] M. E. LOPES, S. WANG, AND M. W. MAHONEY, *Error estimation for randomized least-squares*
719 *algorithms via the bootstrap*, in Proc. 35th International Conference on Machine Learning,
720 vol. 80, PMLR, 2018, pp. 3217–3226.
- 721 [15] P. MA, M. W. MAHONEY, AND B. YU, *A statistical perspective on algorithmic leveraging*, in
722 Proceedings of the 31st International Conference on International Conference on Machine
723 Learning, vol. 32 of ICML'14, JMLR.org, 2014, pp. I–91–I–99.
- 724 [16] P. MA, M. W. MAHONEY, AND B. YU, *A statistical perspective on algorithmic leveraging*, J.
725 Mach. Learn. Res., 16 (2015), pp. 861–911.
- 726 [17] C. L. MALLOWS, *Some comments on c_p* , Technometrics, 15 (1973), pp. 661–675.
- 727 [18] X. MENG, M. A. SAUNDERS, AND M. W. MAHONEY, *LSRN: a parallel iterative solver for*
728 *strongly over- or underdetermined systems*, SIAM J. Sci. Comput., 36 (2014), pp. C95–
729 C118.
- 730 [19] C. D. MEYER, *Matrix analysis and applied linear algebra*, Society for Industrial and Applied
731 Mathematics (SIAM), Philadelphia, PA, 2000.
- 732 [20] J. F. MONAHAN, *A primer on linear models*, Texts in Statistical Science Series, Chapman &
733 Hall/CRC, Boca Raton, FL, 2008.
- 734 [21] D. POSADA AND T. R. BUCKLEY, *Model selection and model averaging in phylogenetics: ad-*
735 *vantages of Akaike information criterion and Bayesian approaches over likelihood ratio*
736 *tests*, Systematic biology, 53 (2004), pp. 793–808.
- 737 [22] G. RASKUTTI AND M. W. MAHONEY, *A statistical perspective on randomized sketching for*
738 *ordinary least-squares*, J. Mach. Learn. Res., 17 (2016), pp. Paper No. 214, 31.
- 739 [23] V. ROKHLIN AND M. TYGERT, *A fast randomized algorithm for overdetermined linear least-*
740 *squares regression*, Proc. Natl. Acad. Sci. USA, 105 (2008), pp. 13212–13217.
- 741 [24] T. SARLÓS, *Improved Approximation Algorithms for Large Matrices via Random Projections*,
742 in 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06), IEEE,
743 Oct 2006, pp. 143–152.
- 744 [25] G. W. STEWART, *Collinearity and least squares regression*, Statist. Sci., 2 (1987), pp. 68–100.
745 With discussion.
- 746 [26] G.-A. THANEI, C. HEINZE, AND N. MEINSHAUSEN, *Random Projections For Large-Scale Re-*
747 *gression*, 2017, <https://arxiv.org/abs/1701.05325>.
- 748 [27] P. F. VELLEMAN AND R. E. WELSCH, *Efficient computing of regression diagnostics*, Amer.
749 Statist., 35 (1981), pp. 234–242.
- 750 [28] H. WANG, R. ZHU, AND P. MA, *Optimal Subsampling for Large Scale Logistic Regression*, J.
751 Amer. Stat. Assoc., 113 (2018), pp. 829–844.