

Randomized Least Squares Regression: Combining Model- and Algorithm-Induced Uncertainties*

Jocelyn T. Chi[†] and Ilse C. F. Ipsen[‡]

Abstract. We analyze the uncertainties in the minimum norm solution of full-rank regression problems, arising from Gaussian linear models, computed by randomized (row-wise sampling and, more generally, sketching) algorithms. From a deterministic perspective our structural perturbation bounds imply that least squares problems are less sensitive to multiplicative perturbations than to additive perturbations. From a probabilistic perspective, our expressions for the total expectation and variance with regard to both, model- and algorithm-induced uncertainties, are exact, hold for general sketching matrices, and make no assumptions on the rank of the sketched matrix. The relative differences between the total bias and variance on the one hand, and the model bias and variance on the other hand, are governed by two factors: (i) the expected rank deficiency of the sketched matrix, and (ii) the expected difference between projectors associated with the original and the sketched problems. A simple example, based on uniform sampling with replacement, illustrates the statistical quantities.

Key words. Condition number with respect to inversion, projector, multiplicative perturbations, Moore Penrose inverse, expectation, variance, matrix valued random variable

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

1. Introduction. We consider regression problems arising from the Gaussian linear model

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

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 unknown parameter vector, and the noise vector $\boldsymbol{\epsilon} \in \mathbb{R}^n$ has a multivariate normal distribution. For a fixed response vector $\mathbf{y} \in \mathbb{R}^n$, one can determine a unique maximum likelihood estimator of $\boldsymbol{\beta}_0$ by computing the unique solution $\hat{\boldsymbol{\beta}}$ of the least squares problem

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

Statistical quality measures include expectation and variance of $\hat{\boldsymbol{\beta}}$, and residual sum of squares $\|\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}\|_2^2$ [13, Section 7.2]; while roundoff errors from a numerically stable method are bounded in terms of the condition number of \mathbf{X} with respect to (left) inversion, and the least squares residual $\mathbf{X}\hat{\boldsymbol{\beta}} - \mathbf{y}$ [7, Chapter 5], [8, Chapter 20].

Randomized algorithms try to reduce the time complexity by first “compressing” or “preconditioning” the least squares problem. They can be classified according to [23, Section 1]: Compression of rows [2, 5, 6, 12, 15, 16, 21]; or columns [1]; or both [17]. We consider row compression

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

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

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

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

35 where $\mathbf{S} \in \mathbb{R}^{r \times n}$ is a random sampling or, more generally, sketching matrix with $r \leq n$,
 36 and the minimum norm solution is $\tilde{\boldsymbol{\beta}}$. Matrix concentration inequalities are used to derive
 37 probabilistic bounds for the error due to randomization [1, 6], and for the condition number
 38 of \mathbf{SX} [12]. From a practical perspective, bootstrapping can deliver fast error estimates [14].

39 The pioneering work [15, 16] was the first to combine the uncertainties from the Gaussian
 40 linear model with the algorithm-induced uncertainties from random sampling of rows. Here
 41 we extend the first-order expansions in [15, 16] in a number of ways.

42 1.1. Contributions.

- 43 1. Our main result presents *exact* expressions for the total expectation and variance of $\tilde{\boldsymbol{\beta}}$
 44 with regard to both, model- and algorithm-induced uncertainties (Theorem 4.5).
- 45 2. Our expressions hold for general random matrices \mathbf{S} , including sketching matrices that
 46 perform projections prior to sampling. Furthermore, our expressions also hold for rank
 47 deficient matrices \mathbf{SX} .
- 48 3. To compare least squares problems of different dimensions, we introduce the *compari-*
 49 *son hat matrix* $\mathbf{P} = \mathbf{X}(\mathbf{SX})^\dagger \mathbf{S}$, which reduces to the *traditional hat matrix* \mathbf{XX}^\dagger when
 50 \mathbf{S} is the identity (Lemma 3.1, Remark 3.2).
- 51 4. We quantify the relative change in the total uncertainty of $\tilde{\boldsymbol{\beta}}$ compared to that of the
 52 model problem (Corollary 4.6):
 - 53 (a) The total bias increases, in the relative sense, with the expected deviation of
 54 the random variable \mathbf{SX} from having full column rank.
 - 55 (b) The relative difference between total variance and model variance increases
 56 with two terms: the expected deviation of \mathbf{SX} from having full rank, plus the
 57 expected deviation of the random variable \mathbf{P} being an orthogonal projector
 58 onto $\text{range}(\mathbf{X})$.
- 59 5. We quantify the model-induced uncertainty of $\tilde{\boldsymbol{\beta}}$, conditioned on \mathbf{S} , compared to that
 60 of the model problem (Theorem 4.3, Corollary 4.4):
 - 61 (a) The bias increases, in the relative sense, with the deviation of \mathbf{SX} from having
 62 full column rank.
 - 63 (b) The variance changes, in the relative sense, with the deviation of \mathbf{P} from being
 64 an orthogonal projector onto $\text{range}(\mathbf{X})$.

65 Thus, unbiasedness is easier to achieve because it only requires \mathbf{SX} to have full column
 66 rank. In contrast, recovering the model variance requires reproducing all of $\text{range}(\mathbf{X})$.

- 67 6. Our structural bounds improve existing bounds, and imply that the minimum norm
 68 solution $\tilde{\boldsymbol{\beta}}$ and its residual are less sensitive to multiplicative perturbations than to
 69 additive perturbations (Corollary 3.5).

70 **1.2. Overview.** After reviewing the computational models for least squares regression
 71 (Section 2), we take adopt two perspectives:

- 72 1. Deterministic: The matrix \mathbf{S} is fixed and the sketched problem (1.3) is a multiplicative
 73 perturbation of the deterministic problem (1.2), and we present structural perturbation
 74 bounds (Section 3).
- 75 2. Probabilistic: The matrix \mathbf{S} is a matrix-valued random variable (1.3) and (1.3) is a
 76 randomized algorithm for solving the linear model (1.1), and we derive expressions for
 77 expectation and variance with regard to the model- and algorithm-induced uncertain-

78 ties (Section 4).

79 This is followed by a brief review of sketching matrices used in randomized least squares
80 solvers (Section 5); a simple example, designed to illustrate the bounds in a way that is easy
81 for readers to reproduce (Section 6); and finally the proofs (Appendix A).

82 **2. Models for Least squares Regression.** Given is a fixed design matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ with
83 $\text{rank}(\mathbf{X}) = p$. Since \mathbf{X} has full column rank, the Moore-Penrose inverse is a left inverse with

$$84 \quad (2.1) \quad \mathbf{X}^\dagger = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \quad \text{and} \quad \mathbf{X}^\dagger \mathbf{X} = \mathbf{I}_p.$$

85 We review the different incarnations of least squares regression: the Gaussian linear model
86 (Section 2.1), the traditional computation (Section 2.2), and the randomized algorithm (Sec-
87 tion 2.3).

88 **2.1. Gaussian linear model.** Let $\boldsymbol{\beta}_0 \in \mathbb{R}^p$ denote the true but generally unknown param-
89 eter vector, and let the response vector $\mathbf{y} \in \mathbb{R}^n$ satisfy the Gauss-Markov assumptions,

$$90 \quad (2.2) \quad \mathbf{y} = \mathbf{X}\boldsymbol{\beta}_0 + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_n).$$

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

93 **2.2. Traditional algorithm for least squares solution.** For a given \mathbf{y} solve

$$94 \quad (2.3) \quad \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{X}\boldsymbol{\beta} - \mathbf{y}\|_2,$$

95 where $\|\mathbf{x}\|_2 = \sqrt{\mathbf{x}^T \mathbf{x}}$ represents the two-norm and the superscript T the transpose.

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

$$97 \quad (2.4) \quad \hat{\boldsymbol{\beta}} \equiv \mathbf{X}^\dagger \mathbf{y}.$$

98 The prediction vector and the least squares residual vector are, respectively

$$99 \quad \hat{\mathbf{y}} \equiv \mathbf{X}\hat{\boldsymbol{\beta}} \quad \text{and} \quad \hat{\boldsymbol{\epsilon}} \equiv \mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}} = \mathbf{y} - \hat{\mathbf{y}}.$$

100 In terms of the so-called *hat matrix* [3, 9, 24],

$$101 \quad (2.5) \quad \mathbf{P}_x \equiv \mathbf{X}\mathbf{X}^\dagger = \mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \in \mathbb{R}^{n \times n},$$

102 which is the orthogonal projector onto $\text{range}(\mathbf{X})$ along $\text{null}(\mathbf{X}^T)$, we can write

$$103 \quad (2.6) \quad \hat{\mathbf{y}} = \mathbf{P}_x \mathbf{y} \quad \text{and} \quad \hat{\boldsymbol{\epsilon}} = (\mathbf{I} - \mathbf{P}_x) \mathbf{y}.$$

104 **2.3. Randomized algorithm for least squares solution.** A randomized algorithm based on
105 sketching, projecting or sampling of rows, is advantageous when \mathbf{X} contains many redundant
106 observations for a small set of variables, that is, $n \gg p$. From a deterministic perspective,
107 this can be considered an extension of weighted least squares [7, Section 6.1] to rectangular
108 weighting matrices.

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

$$110 \quad (2.7) \quad \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{S}(\mathbf{X}\boldsymbol{\beta} - \mathbf{y})\|_2,$$

111 which has the minimum norm solution

$$112 \quad (2.8) \quad \tilde{\boldsymbol{\beta}} \equiv (\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{y}.$$

113 Even if \mathbf{S} has $r > p$ rows, $\text{rank}(\mathbf{S}) < p$ is possible; and even if \mathbf{S} does have full column rank,
114 $\text{rank}(\mathbf{S}\mathbf{X}) < p$ is still possible. Thus (2.7) can have infinitely many solutions, and one way to
115 force uniqueness is to compute the solution of minimal two norm.

116 By design, \mathbf{S} has fewer rows than \mathbf{X} . Hence the corresponding predictions $\hat{\mathbf{y}} = \mathbf{X}\hat{\boldsymbol{\beta}}$ and
117 $\mathbf{S}\mathbf{X}\tilde{\boldsymbol{\beta}}$ have different dimensions and cannot be directly compared; neither can their residuals.
118 To remedy this, we follow previous work [5, 6, 20], and compare the predictions with regard
119 to the *original* matrix,

$$120 \quad (2.9) \quad \tilde{\mathbf{y}} \equiv \mathbf{X}\tilde{\boldsymbol{\beta}} \quad \text{and} \quad \tilde{\mathbf{e}} \equiv \mathbf{y} - \mathbf{X}\tilde{\boldsymbol{\beta}} = \mathbf{y} - \tilde{\mathbf{y}}.$$

121 **3. Structural (deterministic perturbation) bounds.** Here \mathbf{S} is a given, general matrix;
122 and $\mathbf{S}\mathbf{X}$ is interpreted as a perturbation of \mathbf{X} . After deriving expressions for the solution,
123 prediction and least squares residual of the perturbed problem (Section 3.1), we derive mul-
124 tiplicative perturbation bounds (Section 3.2), and discuss comparisons to existing work (Sec-
125 tion 3.3).

126 **3.1. The perturbed problem.** In analogy to the *hat matrix* $\mathbf{P}_\mathbf{x}$ in (2.5) for the original
127 problem (2.3) we introduce a *comparison hat matrix* \mathbf{P} for the perturbed problem (2.7), which
128 allows a clean comparison between two least squares problems of different dimensions.

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

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

is an oblique projector where

1. $\mathbf{P}_\mathbf{x}\mathbf{P} = \mathbf{P}$.
2. $\mathbf{P} - \mathbf{P}_\mathbf{x}$ reflects the difference between the spaces $\text{null}(\mathbf{P})$ and $\text{null}(\mathbf{P}_\mathbf{x})$.
3. $\mathbf{P}\mathbf{X} = \mathbf{X}$ if $\text{rank}(\mathbf{S}\mathbf{X}) = p$.

129

130 *Proof.* See Section A.1. ■

131 The name *comparison hat matrix* will become clear in Theorem 3.3, where \mathbf{P} assumes the
132 duties of the *hat matrix* $\mathbf{P}_\mathbf{x}$ in (2.9).

133 If $\mathbf{S} = \mathbf{I}_n$, then $\mathbf{P} = \mathbf{P}_\mathbf{x}$. In general,

$$134 \quad \text{rank}(\mathbf{P}) = \text{rank}(\mathbf{S}\mathbf{X}) \leq \text{rank}(\mathbf{X}) = \text{rank}(\mathbf{P}_\mathbf{x}) = p.$$

135 If $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X})$, then \mathbf{P} is an oblique version of $\mathbf{P}_\mathbf{x}$ with $\text{range}(\mathbf{P}) = \text{range}(\mathbf{P}_\mathbf{x})$, and
136 the only difference is in their nullspaces. If $\text{rank}(\mathbf{S}\mathbf{X}) < p$ then $\text{rank}(\mathbf{P}) = \text{rank}(\mathbf{S}\mathbf{X}) < p$, and
137 \mathbf{P} projects onto only a subspace of $\text{range}(\mathbf{X})$. The example in Section 6.1 illustrates this.

138 *Remark 3.2.* The comparison hat matrix \mathbf{P} generalizes the oblique projector $\mathbf{P}_{\mathbf{u}}$ in [20,
 139 (11)], which was introduced to quantify *prediction efficiency* and *residual efficiency* of sketch-
 140 ing algorithms in the statistical setting (2.2). This projector $\mathbf{P}_{\mathbf{u}}$ is defined if $\text{rank}(\mathbf{S}\mathbf{X}) = p$,
 141 and equals $\mathbf{P}_{\mathbf{u}} \equiv \mathbf{U}(\mathbf{S}\mathbf{U})^\dagger \mathbf{S}$, where \mathbf{U} is an orthonormal basis for $\text{range}(\mathbf{X})$. In this case we
 142 have $\mathbf{P}_{\mathbf{u}} = \mathbf{P}$. However, if $\text{rank}(\mathbf{S}\mathbf{X}) < \text{rank}(\mathbf{X})$, then $\mathbf{P}_{\mathbf{u}}$ is not sufficient in our context.

Theorem 3.3 (Perturbed least squares problem). *With the assumptions in Section 2, the solution of (2.7) satisfies*

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

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

$$\begin{aligned}\tilde{\mathbf{y}} &= \mathbf{P}\mathbf{y} = \hat{\mathbf{y}} + (\mathbf{P} - \mathbf{P}_{\mathbf{x}})\mathbf{y}, \\ \tilde{\mathbf{e}} &= (\mathbf{I} - \mathbf{P})\mathbf{y} = \hat{\mathbf{e}} + (\mathbf{P}_{\mathbf{x}} - \mathbf{P})\mathbf{y}.\end{aligned}$$

143

144 *Proof.* See Section A.2. ■

145 Theorem 3.3 shows that the relations between perturbed and original least squares prob-
 146 lems are governed by $\mathbf{P} - \mathbf{P}_{\mathbf{x}}$, which reflects the difference between $\text{null}(\mathbf{P})$ and $\text{null}(\mathbf{P}_{\mathbf{x}})$.
 147 Motivated by the ground breaking result [16, Lemma 1], reproduced in the lemma below,
 148 Theorem 3.3 strengthens it with explicit expressions for $\tilde{\boldsymbol{\beta}}$ that hold for general matrices \mathbf{S}
 149 and do not require assumptions on $\text{rank}(\mathbf{S}\mathbf{X})$.

150 **Lemma 3.4 (Lemma 1 in [15] and [16]).** *If, in addition to the assumptions in Section 2,*
 151 *the matrix \mathbf{S} in (2.7) has a single nonzero entry per row, the vector¹ $\mathbf{w} \equiv \text{diag}(\mathbf{S}^T \mathbf{S}) \in \mathbb{R}^n$ has*
 152 *a scaled multinomial distribution with expected value $\mathbb{E}[\mathbf{w}] = \mathbf{1}$, $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X})$, and a*
 153 *Taylor series expansion around $\mathbf{w}_0 = \mathbf{1}$ of the solution $\tilde{\boldsymbol{\beta}}(\mathbf{w})$ of (2.7) exists with $\tilde{\boldsymbol{\beta}}(\mathbf{w}_0) = \hat{\boldsymbol{\beta}}$,*
 154 *then*

$$155 \quad \tilde{\boldsymbol{\beta}}(\mathbf{w}) = \hat{\boldsymbol{\beta}} + \mathbf{X}^\dagger \text{diag}(\hat{\mathbf{e}})(\mathbf{w} - \mathbf{1}) + R(\mathbf{w}),$$

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

158 **3.2. Multiplicative perturbation bounds.** We consider the problem (2.7) as a multiplica-
 159 tive perturbation, and derive norm-wise relative error bounds for the solution, prediction, and
 160 least squares residual; and compare them to existing bounds.

161 The vector two-norm induces the matrix norm $\|\mathbf{X}\|_2$, and the two-norm condition number
 162 of the full column-rank matrix \mathbf{X} with regard to (left) inversion is

$$163 \quad \kappa_2(\mathbf{X}) \equiv \|\mathbf{X}\|_2 \|\mathbf{X}^\dagger\|_2 \geq 1.$$

¹For a matrix \mathbf{M} , $\text{diag}(\mathbf{M})$ represents the vector of diagonal elements.

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

The solution of (2.7) satisfies

$$\frac{\|\tilde{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}\|_2}{\|\hat{\boldsymbol{\beta}}\|_2} \leq \kappa_2(\mathbf{X}) \frac{\|\mathbf{y}\|_2}{\|\mathbf{X}\|_2 \|\hat{\boldsymbol{\beta}}\|_2} \|\mathbf{P} - \mathbf{P}_{\mathbf{x}}\|_2.$$

The least squares residual $\tilde{\mathbf{e}} = \mathbf{y} - \mathbf{X}\tilde{\boldsymbol{\beta}}$ satisfies

$$\frac{\|\tilde{\mathbf{e}} - \hat{\mathbf{e}}\|_2}{\|\hat{\mathbf{e}}\|_2} \leq \frac{\|\mathbf{P} - \mathbf{P}_{\mathbf{x}}\|_2}{\sin \theta}.$$

164

165

Proof. See Section A.3 ■

166

For $\mathbf{S} = \mathbf{I}_n$, the bounds in Corollary 3.5 are zero and tight since $\mathbf{P} = \mathbf{P}_{\mathbf{x}}$.

167

168

169

Remark 3.6 (Sensitivity to multiplicative perturbations). Corollary 3.5 implies that least squares solutions $\tilde{\boldsymbol{\beta}}$ are insensitive to multiplicative perturbations if \mathbf{X} is well conditioned with regard to inversion, and if \mathbf{y} is close to $\text{range}(\mathbf{X})$. The bound for $\tilde{\boldsymbol{\beta}}$ consists of two parts:

170

171

172

1. The perturbation $\|\mathbf{P} - \mathbf{P}_{\mathbf{x}}\|_2$ reflects the distance between the null spaces $\text{null}(\mathbf{P})$ and $\text{null}(\mathbf{P}_{\mathbf{x}})$. It is an absolute as well as a relative perturbation since $\|\mathbf{P}_{\mathbf{x}}\|_2 = 1$.
2. The amplifier can be bounded by [7, (5.3.16)]

173

$$\kappa_2(\mathbf{X}) \frac{\|\mathbf{y}\|_2}{\|\mathbf{X}\|_2 \|\hat{\boldsymbol{\beta}}\|_2} \leq \kappa_2(\mathbf{X}) \frac{\|\mathbf{y}\|_2}{\|\mathbf{X}\hat{\boldsymbol{\beta}}\|_2} = \frac{\kappa_2(\mathbf{X})}{\cos \theta}.$$

174

175

176

177

3.3. Comparison to existing work. In contrast to multiplicative perturbation bounds for eigenvalue and singular value problems [10, 11], we do not require \mathbf{S} to be nonsingular or square. Weighted least squares problems [7, Section 6.1] employ nonsingular diagonal matrices \mathbf{S} for regularization or scaling of discrepancies, and do not view them as a perturbation.

178

179

180

Remark 3.7 (Comparison to additive perturbations). Corollary 3.5 also implies that the minimum norm solution of (2.7) and its residual are less sensitive to multiplicative perturbations than to additive perturbations, which are reviewed below in Lemma 3.8.

181

182

183

In contrast to additive bounds [7, (5.3.12)], [8, (20.12)], the bound for the least squares residual $\tilde{\mathbf{e}}$ is not affected by $\kappa_2(\mathbf{X})$. Note that the $\sin \theta$ term in the denominator is also occurs in additive bounds [7, (5.3.12)] if the relative error is normalized by $\hat{\mathbf{e}}$ rather than \mathbf{y} .

184

185

186

187

In contrast to additive bounds [7, Section 5.3.6], [8, Section 20.1], [22, (3.4)], the bound for $\hat{\boldsymbol{\beta}}$ does not square the condition number and does not require $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X})$. This can be seen from Lemma 3.8 below, where the first summand corresponds to the bound for $\tilde{\boldsymbol{\beta}}$ in Corollary 3.5.

188

189

190

Lemma 3.8 (Theorem 5.3.1 in [7]). *With the assumptions in Section 2, let $\mathbf{X} + \mathbf{E}$ have $\text{rank}(\mathbf{X} + \mathbf{E}) = \text{rank}(\mathbf{X})$ and $\eta \equiv \|\mathbf{E}\|_2 / \|\mathbf{X}\|_2$.*

The solution $\bar{\boldsymbol{\beta}}$ to $\min_{\boldsymbol{\beta}} \|(\mathbf{X} + \mathbf{E})\boldsymbol{\beta} - \mathbf{y}\|_2$ satisfies

191

$$\frac{\|\bar{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}\|_2}{\|\hat{\boldsymbol{\beta}}\|_2} \leq \kappa_2(\mathbf{X}) \left(\frac{\|\mathbf{y}\|_2}{\|\mathbf{X}\|_2 \|\hat{\boldsymbol{\beta}}\|_2} + 1 \right) \eta + \kappa(\mathbf{X})^2 \frac{\|\hat{\mathbf{e}}\|_2}{\|\mathbf{X}\|_2 \|\hat{\boldsymbol{\beta}}\|_2} \eta + \mathcal{O}(\eta^2).$$

192 Compared to existing structural bounds for randomized least squares algorithms, which
 193 are reproduced in Lemma 3.9, the bound for $\tilde{\boldsymbol{\beta}}$ in Corollary 3.5 is more general and tighter in
 194 the sense that it does not exhibit nonlinear dependences on the perturbations.

195 **Lemma 3.9 (Theorem 1 in [6]).** *In addition to Section 2, also assume that $\|\mathbf{P}_x \mathbf{y}\|_2 \geq \gamma \|\mathbf{y}\|_2$
 196 for some $0 < \gamma \leq 1$ and $\|\tilde{\boldsymbol{\epsilon}}\|_2 \leq (1 + \eta) \|\hat{\boldsymbol{\epsilon}}\|_2$. Then*

$$197 \quad \frac{\|\tilde{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}\|_2}{\|\hat{\boldsymbol{\beta}}\|_2} \leq \kappa_2(\mathbf{X}) \sqrt{\gamma^{-2} - 1} \sqrt{\eta}.$$

198 **4. Model-induced and randomized algorithm-induced uncertainty.** Under the linear
 199 model (2.2), the computed solution $\hat{\boldsymbol{\beta}}$ has nice statistical properties [19, Chapter 6], as it
 200 is an unbiased estimator of $\boldsymbol{\beta}_0$ and it has minimal variance among all linear unbiased estima-
 201 tors. We show how this changes with the addition of algorithm-induced uncertainty.

202 After briefly reviewing the uncertainty induced by the linear model (Section 4.1); we
 203 derive the expectation and variance of $\tilde{\boldsymbol{\beta}}$, conditioned on the algorithm-induced uncertainty
 204 (Section 4.2). From that we derive the total expectation and variance (Section 4.3).

205 **4.1. Model-induced uncertainty.** We view the model-induced randomness in (2.2) as a
 206 property of the response vector \mathbf{y} . That is, the noise vector $\boldsymbol{\epsilon}$ in (2.2) has mean and covariance

$$207 \quad \mathbb{E}_{\mathbf{y}}[\boldsymbol{\epsilon}] = \mathbf{0}, \quad \text{Var}_{\mathbf{y}}[\boldsymbol{\epsilon}] = \sigma^2 \mathbf{I}_n.$$

208 The well-known statistical properties of (2.3) are reviewed below.

209 **Lemma 4.1 (Model-induced uncertainty for (2.3)).** *With the assumptions in Section 2, the
 210 response vector (2.2), and the least squares prediction (2.6) and solution (2.4) satisfy*

$$\begin{aligned} 211 \quad \mathbb{E}_{\mathbf{y}}[\mathbf{y}] &= \mathbf{X}\boldsymbol{\beta}_0, & \text{Var}_{\mathbf{y}}[\mathbf{y}] &= \sigma^2 \mathbf{I}_n \\ 212 \quad \mathbb{E}_{\mathbf{y}}[\hat{\mathbf{y}}] &= \mathbf{X}\boldsymbol{\beta}_0, & \text{Var}_{\mathbf{y}}[\hat{\mathbf{y}}] &= \sigma^2 \mathbf{P}_x \in \mathbb{R}^{n \times n} \\ 213 \quad \mathbb{E}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] &= \boldsymbol{\beta}_0, & \text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] &= \sigma^2 (\mathbf{X}^T \mathbf{X})^{-1} \in \mathbb{R}^{p \times p}. \end{aligned}$$

214 *Proof.* See Section A.4. ■

215 Lemma 4.1 asserts that the computed solution $\hat{\boldsymbol{\beta}}$ is an unbiased estimator for $\boldsymbol{\beta}_0$, and
 216 points to the well known dependence of the variance on the conditioning of \mathbf{X} [22, Section 5].

217 The difficulty in analyzing the sketched problem (2.7), coupled with general concern about
 218 the first-order expansions like the ones in [15, 16], is that there are instances of \mathbf{S} for which
 219 $\text{rank}(\mathbf{S}\mathbf{X}) < \text{rank}(\mathbf{X})$. In this case $(\mathbf{S}\mathbf{X})^\dagger$ cannot be expressed in terms of $\mathbf{S}\mathbf{X}$ as in (2.1), and
 220 the least squares problem (2.7) is ill-posed.

221 One can derive bounds [1, Theorem 3.2], [12, Theorems 4.1 and 5.2] on the probability
 222 that $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X})$ for matrices \mathbf{S} that perform uniform sampling and leverage score
 223 sampling. However, such bounds are not useful here, because expected values run over *all*
 224 instances of $\mathbf{S}\mathbf{X}$.

225 We introduce a quantity that signals the deviation of the columns of $\mathbf{S}\mathbf{X}$ from linear
 226 independence.

Lemma 4.2 (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 where

1. $\mathbf{P}\mathbf{X} = \mathbf{X}\mathbf{P}_0$
2. $\mathbf{P}_0 = \mathbf{I}_p$ if $\text{rank}(\mathbf{S}\mathbf{X}) = p$.
3. $\mathbf{I} - \mathbf{P}_0$ represents the deviation of $\mathbf{S}\mathbf{X}$ from full-rank.

227

228

Proof. See Section A.5. ■

229

The name *bias projector* will become apparent in Theorem 4.3, where \mathbf{P}_0 represents the bias in $\tilde{\boldsymbol{\beta}}$.

230

231

If $\mathbf{S} = \mathbf{I}_n$, then $\mathbf{P}_0 = \mathbf{P}_x$. If $\text{rank}(\mathbf{S}\mathbf{X}) = p$, then Lemma 4.2 recovers $\mathbf{P}\mathbf{X} = \mathbf{X}$ from Lemma 3.1, confirming that \mathbf{P} is a projector onto $\text{range}(\mathbf{X})$. However, if $\text{rank}(\mathbf{S}\mathbf{X}) < p$, then \mathbf{P}_0 characterizes the subspace of $\text{range}(\mathbf{X})$ onto which \mathbf{P} projects.

232

233

234

4.2. Model-induced uncertainty, conditioned on algorithm-induced uncertainty. We

235

determine the expectation and variance for the solution of (2.7) conditioned on \mathbf{S} . That is, we

236

assume that the random sketching matrix \mathbf{S} is fixed at a specific value \mathbf{S}_0 and use $\mathbb{E}_y \left[\cdot \mid \mathbf{S} \right]$

237

as an abbreviation for the conditional expectation $\mathbb{E}_y \left[\cdot \mid \mathbf{S} = \mathbf{S}_0 \right]$.

Theorem 4.3 (Model-induced uncertainty in (2.7), conditioned on \mathbf{S}). *With the assumptions in Section 2, the solution of (2.7) satisfies*

$$\begin{aligned} \mathbb{E}_y \left[\tilde{\boldsymbol{\beta}} \mid \mathbf{S} \right] &= \mathbf{P}_0 \boldsymbol{\beta}_0 = \boldsymbol{\beta}_0 + (\mathbf{I} - \mathbf{P}_0) \boldsymbol{\beta}_0 \\ \text{Var}_y \left[\tilde{\boldsymbol{\beta}} \mid \mathbf{S} \right] &= \sigma^2 \mathbf{X}^\dagger \mathbf{P} \mathbf{P}^T (\mathbf{X}^\dagger)^T \in \mathbb{R}^{p \times p}, \\ &= \text{Var}_y[\hat{\boldsymbol{\beta}}] + \sigma^2 \mathbf{X}^\dagger (\mathbf{P} \mathbf{P}^T - \mathbf{P}_x) (\mathbf{X}^\dagger)^T, \end{aligned}$$

where $\mathbf{P} \mathbf{P}^T - \mathbf{P}_x$ is the deviation of \mathbf{P} from being an orthogonal projector onto $\text{range}(\mathbf{X})$.

Furthermore $\mathbb{E}_y \left[\tilde{\boldsymbol{\beta}} \mid \text{rank}(\mathbf{S}\mathbf{X}) = p \right] = \boldsymbol{\beta}_0$.

238

239

Proof. See Section A.6. ■

240

The exact expressions for general sketching matrices \mathbf{S} in Theorem 4.3 extend the first-order expressions for specific sampling matrices in [16, Lemmas 2-6]. The examples in Section 6.1 illustrate the effect of rank deficiency of $\mathbf{S}\mathbf{X}$ on the quantities in Theorem 4.3.

241

242

243

Theorem 4.3 shows that the bias of $\tilde{\boldsymbol{\beta}}$ is proportional to the deviation $\mathbf{I} - \mathbf{P}_0$ of $\mathbf{S}\mathbf{X}$ from having full column rank. In other words, the bias becomes worse as the rank deficiency increases. If $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X})$, then $\tilde{\boldsymbol{\beta}}$ is an unbiased estimator of $\boldsymbol{\beta}_0$. Theorem 4.3 also implies that the conditional variance is close to the model variance if \mathbf{P} is close to being an orthogonal projector onto $\text{range}(\mathbf{X})$.

244

245

246

247

248

The relevance of $\mathbf{I} - \mathbf{P}_0$ and $\mathbf{P} \mathbf{P}^T - \mathbf{P}_x$ becomes clear in the relative differences below.

Corollary 4.4 (Relative difference between conditional and model quantities). *With the assumptions in Theorem 4.3,*

$$\frac{\|\text{Var}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} | \mathbf{S}] - \text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]\|_2}{\|\text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]\|_2} \leq \|\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{x}}\|_2.$$

If also $\boldsymbol{\beta}_0 \neq \mathbf{0}$, then

$$\frac{\|\mathbb{E}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} | \mathbf{S}] - \boldsymbol{\beta}_0\|_2}{\|\boldsymbol{\beta}_0\|_2} \leq \|\mathbf{I} - \mathbf{P}_0\|_2.$$

249

250

Proof. See Section A.7. ■

251

252

253

254

255

256

257

258

Corollary 4.4 implies that the relative differences to unbiasedness and model variance are solely governed by the quantities $\mathbf{I} - \mathbf{P}_0$ and $\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{x}}$, respectively. Both of them are absolute as well as relative measures since $\|\mathbf{I}\|_2 = \|\mathbf{P}_{\mathbf{x}}\|_2 = 1$. Specifically, the relative difference between conditional and model variance increases with the deviation of \mathbf{P} from being an orthogonal projector onto $\text{range}(\mathbf{X})$; and the bias of $\tilde{\boldsymbol{\beta}}$ increases, in the relative sense, with the deviation of $\mathbf{S}\mathbf{X}$ from full column rank.

Thus, unbiasedness is easier to achieve because it only requires $\mathbf{S}\mathbf{X}$ to have full column rank. In contrast, recovering the model variance requires reproducing all of $\text{range}(\mathbf{X})$.

259

260

261

262

263

264

265

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

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

266 (4.1)

$$\mathbb{E}[\cdot] = \mathbb{E}_{\mathbf{s}} \left[\mathbb{E}_{\mathbf{y}} \left[\cdot \mid \mathbf{S} \right] \right].$$

267

268

Since \mathbf{S} is a matrix-valued random variable, so are the projectors \mathbf{P} and \mathbf{P}_0 . Examples of \mathbf{S} can be found in Sections 5 and 6.

Theorem 4.5 (Total mean and variance for (2.7)). *With the assumptions in Section 2, let \mathbf{S} be a random sketching matrix. The solution of (2.7) satisfies*

$$\begin{aligned}\mathbb{E}[\tilde{\boldsymbol{\beta}}] &= \mathbb{E}_{\mathbf{S}}[\mathbf{P}_0]\boldsymbol{\beta}_0 = \boldsymbol{\beta}_0 + \mathbb{E}_{\mathbf{S}}[\mathbf{P}_0 - \mathbf{I}]\boldsymbol{\beta}_0 \\ \text{Var}[\tilde{\boldsymbol{\beta}}] &= \sigma^2 \mathbf{X}^\dagger \mathbb{E}_{\mathbf{S}}[\mathbf{P}\mathbf{P}^T] (\mathbf{X}^\dagger)^T + \text{Var}_{\mathbf{S}}[\mathbf{P}_0\boldsymbol{\beta}_0] \\ &= \text{Var}_{\mathbf{Y}}[\hat{\boldsymbol{\beta}}] + \sigma^2 \mathbf{X}^\dagger \mathbb{E}_{\mathbf{S}}[\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{X}}] (\mathbf{X}^\dagger)^T + \text{Var}_{\mathbf{S}}[(\mathbf{P}_0 - \mathbf{I})\boldsymbol{\beta}_0].\end{aligned}$$

where

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

269

270

Proof. See Section A.8. ■

271

272

273

274

275

276

277

278

279

280

281

Theorem 4.5 presents exact expressions for general random matrices \mathbf{S} , thereby extending the first order approximations for specific sampling matrices in [16, Lemmas 2-6], and shows:

1. The total bias of $\tilde{\boldsymbol{\beta}}$ is proportional to the expected deviation of the matrix-valued random variable $\mathbf{S}\mathbf{X}$ from having full column rank.

The expectation $\mathbb{E}_{\mathbf{S}}[\mathbf{P}_0]$ of a projector \mathbf{P}_0 is not a projector, in general, as the example in Section 6.3 illustrates.

2. The total variance of $\tilde{\boldsymbol{\beta}}$ is proportional to the expected deviation of $\mathbf{S}\mathbf{X}$ from full column rank, plus the expected deviation of the matrix-valued random variable \mathbf{P} from being an orthogonal projector onto $\text{range}(\mathbf{X})$.

The importance of the expected deviations of the projectors appears in the analog of Corollary 4.4 below.

Corollary 4.6. *With the assumptions in Theorem 4.5,*

$$\frac{\|\text{Var}[\tilde{\boldsymbol{\beta}}] - \text{Var}_{\mathbf{Y}}[\hat{\boldsymbol{\beta}}]\|_2}{\|\text{Var}_{\mathbf{Y}}[\hat{\boldsymbol{\beta}}]\|_2} \leq \|\mathbb{E}_{\mathbf{S}}[\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{X}}]\|_2 + \frac{\|\text{Var}_{\mathbf{S}}[(\mathbf{P}_0 - \mathbf{I})\boldsymbol{\beta}_0]\|_2}{\|\text{Var}_{\mathbf{Y}}[\hat{\boldsymbol{\beta}}]\|_2}.$$

If also $\boldsymbol{\beta}_0 \neq \mathbf{0}$, then

$$\frac{\|\mathbb{E}_{\mathbf{S}}[\tilde{\boldsymbol{\beta}}] - \boldsymbol{\beta}_0\|_2}{\|\boldsymbol{\beta}_0\|_2} \leq \|\mathbb{E}_{\mathbf{S}}[\mathbf{I} - \mathbf{P}_0]\|_2.$$

282

283

284

285

286

Corollary 4.6 implies that the bias of $\tilde{\boldsymbol{\beta}}$ increases, in the relative sense, with the expected deviation of $\mathbf{S}\mathbf{X}$ from full rank; and that the relative difference from total variance to model variance increases with (i) the expected deviation of \mathbf{P} being an orthogonal projector onto $\text{range}(\mathbf{X})$, plus (ii) the expected deviation of $\mathbf{S}\mathbf{X}$ from full rank.

287

288

5. Random sketching matrices in least squares. We present a few examples of sketching matrices from the randomized least squares solvers [1, 2, 5, 6, 14, 15, 16, 17, 21].

289 *Uniform sampling with replacement.* This is the *EXACTLY(c)* algorithm [6, Algorithm 3]
 290 with uniform probabilities, which is used for row-wise compression of direct methods for the
 291 solution of full column rank least squares in [6, Algorithm 3], see also the *BasicMatrixMul-*
 292 *tification Algorithm* [4, Fig. 2], [12, Algorithm 3.2], [14, Algorithms 1 and 2], and [16,
 293 UNIF].

Algorithm 5.1 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

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

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

296 *Random orthogonal sketching.* This is used in *Blendenpik* [1, Algorithm 1] to compute ran-
 297 domized preconditioners for the iterative solution of full column rank least squares problems.

298 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 elements
 299 are independent Rademacher random variables, equaling ± 1 with equal probability; $\mathbf{T} \in \mathbb{R}^{n \times n}$
 300 is a unitary matrix, such as a Walsh-Hadamard, discrete cosine, or discrete Hartley transform;
 301 and \mathbf{B} is a diagonal matrix whose diagonal elements are Bernoulli variables, equaling 1 with
 302 probability $\gamma p/n$ for some $\gamma > 0$, and 0 otherwise.

303 *Gaussian sketching.* This is used in to compute randomized preconditioners for the iterative
 304 solution of general least squares problems [17, Algorithms 1 and 2].

305 Here the elements of $\mathbf{S} \in \mathbb{R}^{r \times n}$ are independent $\mathcal{N}(0, 1)$ random variables. In Matlab:
 306 $\mathbf{S} = \text{randn}(r, n)$.

307 **6. Example.** We illustrate the projectors in Corollary 3.5, Theorem 4.3, Corollary 4.4,
 308 Theorem 4.5, and Corollary 4.6 in a way that is easy for readers to reproduce. For a small
 309 example matrix, we illustrate the effects of rank deficiency $\mathbf{S}\mathbf{X}$ (Section 6.1); perform uniform
 310 sampling with replacement (Section 6.2); compute the expectations for \mathbf{P}_0 (Section 6.3) and
 311 $\mathbf{P}\mathbf{P}^T$ (Section 6.4); and put this into context with two matrices at opposite ends of sampling
 312 performance (Section 6.5).

313 Consider the full column rank matrix

$$314 \quad \mathbf{X} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 0 \end{pmatrix} \in \mathbb{R}^{4 \times 2} \quad \text{with} \quad \mathbf{X}^\dagger = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix},$$

315 $\text{rank}(\mathbf{X}) = 2$,

$$316 \quad (6.1) \quad \mathbf{P}_x = \mathbf{X}\mathbf{X}^\dagger = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 1 & 0 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \text{Var}_y[\tilde{\boldsymbol{\beta}}] = \sigma^2(\mathbf{X}^T\mathbf{X})^{-1} = \sigma^2 \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 1 \end{pmatrix},$$

317 and

$$318 \quad \text{null}(\mathbf{P}_x) = \text{range}(\mathbf{I} - \mathbf{P}_x) = \text{range} \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ -1 & 0 \\ 0 & 1 \end{pmatrix}.$$

319 **6.1. Effects of rank deficiency.** We illustrate the effect of rank deficiency on the quantities
 320 in Corollary 3.5, Theorem 4.3 and Corollary 4.4 by choosing two different matrices \mathbf{S} with
 321 $\text{rank}(\mathbf{S}) = 2$.

322 **Full column rank \mathbf{SX} .** Here

$$323 \quad \mathbf{S} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \quad \text{where} \quad \mathbf{SX} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = (\mathbf{SX})^\dagger = \mathbf{I}_2,$$

324 $\text{rank}(\mathbf{SX}) = \text{rank}(\mathbf{X}) = 2$, and $\text{range}(\mathbf{P}) = \text{range}(\mathbf{X})$. This gives the projectors

$$325 \quad \mathbf{P}_0 = (\mathbf{SX})^\dagger(\mathbf{SX}) = \mathbf{I}_2, \quad \mathbf{P} = \mathbf{X}(\mathbf{SX})^\dagger\mathbf{S} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$

326 with

$$327 \quad \text{null}(\mathbf{P}) = \text{range}(\mathbf{I} - \mathbf{P}) = \text{range} \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

328 This shows:

- 329 • \mathbf{P} is not an orthogonal projector, since it is not symmetric.
- 330 • The solution $\tilde{\boldsymbol{\beta}}$ in Theorem 4.3 is an unbiased estimator.
- 331 • The conditional variance in Theorem 4.3 has increased compared to (6.1), since

$$332 \quad \text{Var}_y[\tilde{\boldsymbol{\beta}} \mid \mathbf{S}] = \sigma^2 \mathbf{X}^\dagger \mathbf{P} \mathbf{P}^T (\mathbf{X}^\dagger)^T = \sigma^2 (\mathbf{SX})^\dagger \mathbf{S} \mathbf{S}^T ((\mathbf{SX})^\dagger)^T = \sigma^2 \mathbf{I}_2.$$

333 **Rank deficient \mathbf{SX} .** Here

$$334 \quad \mathbf{S} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad \text{where} \quad \mathbf{SX} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = (\mathbf{SX})^\dagger,$$

335 $\text{rank}(\mathbf{SX}) = 1 < \text{rank}(\mathbf{X})$, and $\text{range}(\mathbf{P}) \subset \text{range}(\mathbf{X})$. This gives the projectors

$$336 \quad \mathbf{P}_0 = (\mathbf{SX})^\dagger(\mathbf{SX}) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad \mathbf{P} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$

337 with

$$338 \quad \text{null}(\mathbf{P}) = \text{range}(\mathbf{I} - \mathbf{P}) = \text{range} \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

339 This shows:

- 340 • The rank deficiency of \mathbf{SX} causes the dimension of $\text{null}(\mathbf{P})$ to increase.
- 341 • The solution $\tilde{\boldsymbol{\beta}}$ in Theorem 4.3 is a biased estimator since $\mathbf{P}_0 \neq \mathbf{I}_2$.
- 342 • The conditional variance in Theorem 4.3 has become singular, since

$$343 \quad \text{Var}_y \left[\tilde{\boldsymbol{\beta}} \mid \mathbf{S} \right] = \sigma^2 \mathbf{X}^\dagger \mathbf{P} \mathbf{P}^T (\mathbf{X}^\dagger)^T = \sigma^2 (\mathbf{SX})^\dagger \mathbf{S} \mathbf{S}^T ((\mathbf{SX})^\dagger)^T = \sigma^2 \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}^\dagger.$$

344 **6.2. Uniform sampling with replacement.** Algorithm 5.1 with $n = 4$ and $r = 2$ produces
 345 a sampling matrix $\mathbf{S} \in \mathbb{R}^{2 \times 4}$, which has $n^2 = 16$ instances

$$346 \quad \mathbf{S}_{ij} = \sqrt{2} \begin{pmatrix} \mathbf{e}_i^T \\ \mathbf{e}_j^T \end{pmatrix}, \quad 1 \leq i, j \leq n,$$

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

$$348 \quad \mathbf{S}_{11} = \sqrt{2} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{S}_{42} = \sqrt{2} \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}.$$

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

$$350 \quad \mathbb{E}_{\mathbf{s}}[\mathbf{S}^T \mathbf{S}] = \sum_{i=1}^4 \sum_{j=1}^4 \frac{1}{16} \mathbf{S}_{ij}^T \mathbf{S}_{ij} = \sum_{i=1}^4 \sum_{j=1}^4 \frac{1}{16} (\mathbf{e}_i \mathbf{e}_i^T + \mathbf{e}_j \mathbf{e}_j^T) = \mathbf{I}_4.$$

351 **6.3. Expected deviation of \mathbf{SX} from rank deficiency.** We compute the expectation of
 352 $\mathbf{P}_0 \in \mathbb{R}^{2 \times 2}$ in Theorem 4.5 and Corollary 4.6,

$$353 \quad \mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] = \sum_{i=1}^4 \sum_{j=1}^4 \frac{1}{16} (\mathbf{S}_{ij} \mathbf{X})^\dagger (\mathbf{S}_{ij} \mathbf{X}) = \mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] = \frac{1}{16} \begin{pmatrix} 12 & 0 \\ 0 & 7 \end{pmatrix}.$$

354 Representative summands include

$$355 \quad (\mathbf{S}_{13}\mathbf{X})^\dagger = \sqrt{\frac{1}{2}} \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix}^\dagger = \sqrt{\frac{1}{2}} \begin{pmatrix} 1/2 & 1/2 \\ 0 & 0 \end{pmatrix}, \quad (\mathbf{S}_{13}\mathbf{X})^\dagger(\mathbf{S}_{13}\mathbf{X}) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$$

$$356 \quad (\mathbf{S}_{32}\mathbf{X})^\dagger = \sqrt{\frac{1}{2}} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}^\dagger = \sqrt{\frac{1}{2}} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = \sqrt{\frac{1}{2}}\mathbf{I}_2, \quad (\mathbf{S}_{32}\mathbf{X})^\dagger(\mathbf{S}_{32}\mathbf{X}) = \mathbf{I}_2$$

$$357 \quad (\mathbf{S}_{44}\mathbf{X})^\dagger = \sqrt{\frac{1}{2}} \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}^\dagger = \mathbf{0}, \quad (\mathbf{S}_{44}\mathbf{X})^\dagger(\mathbf{S}_{44}\mathbf{X}) = \mathbf{0}.$$

358 Among the sketched matrices, 75 percent are rank deficient. The ones with full column rank
359 are $\mathbf{S}_{12}\mathbf{X}$, $\mathbf{S}_{21}\mathbf{X}$, $\mathbf{S}_{23}\mathbf{X}$, and $\mathbf{S}_{32}\mathbf{X}$. This shows

- 360 • $\mathbb{E}_s[\mathbf{P}_0]$ is not a projector, since it is not idempotent.
361 • The solution $\tilde{\boldsymbol{\beta}}$ in Theorem 4.5 is a biased estimator, since $\mathbb{E}_s[\mathbf{P}_0] \neq \mathbf{I}_2$.
362 • The relative difference of $\tilde{\boldsymbol{\beta}}$ from unbiasedness in Corollary 4.6 can exceed 50 percent,
363 since it is bounded by $\|\mathbb{E}_s[\mathbf{I} - \mathbf{P}_0]\|_2 = \frac{9}{16}$, where

$$364 \quad \mathbb{E}_s[\mathbf{I} - \mathbf{P}_0] = \frac{1}{16} \begin{pmatrix} 4 & 0 \\ 0 & 9 \end{pmatrix}.$$

365 **6.4. Expected deviation of \mathbf{P} from being an orthogonal projector.** We compute the
366 expectation of $\mathbf{P}\mathbf{P}^T \in \mathbb{R}^{4 \times 4}$ in Theorem 4.5 and Corollary 4.6. Since the trailing column of
367 \mathbf{X} is zero, and

$$368 \quad \mathbf{P}\mathbf{P}^T = \mathbf{X}(\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{S}^T ((\mathbf{S}\mathbf{X})^\dagger)^T \mathbf{X}^T,$$

369 the trailing row and columns of all instances of $\mathbf{P}\mathbf{P}^T$ and $\mathbb{E}_s[\mathbf{P}\mathbf{P}^T]$ are, too. Thus

$$370 \quad \mathbb{E}_s[\mathbf{P}\mathbf{P}^T] = \sum_{i=1}^4 \sum_{j=1}^4 \frac{1}{16} \mathbf{X}(\mathbf{S}_{ij}\mathbf{X})^\dagger \mathbf{S}_{ij} \mathbf{S}_{ij}^T ((\mathbf{S}_{ij}\mathbf{X})^\dagger)^T \mathbf{X}^T = \frac{1}{16} \begin{pmatrix} 11 & 0 & 11 & 0 \\ 0 & 7 & 0 & 0 \\ 11 & 0 & 11 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

371 This shows

- 372 • $\mathbb{E}_s[\mathbf{P}\mathbf{P}^T]$ is not a projector since it is not idempotent.
373 • The expected deviation of \mathbf{P} from being an orthogonal projector onto $\text{range}(\mathbf{X})$ in
374 Corollary 4.6 can exceed 50 percent, since it is bounded by $\|\mathbb{E}_s[\mathbf{P}\mathbf{P}^T - \mathbf{P}_x]\|_2 = \frac{9}{16}$,
375 where with the hat matrix \mathbf{P}_x in (6.1)

$$376 \quad \mathbb{E}_s[\mathbf{P}\mathbf{P}^T - \mathbf{P}_x] = \frac{1}{16} \begin{pmatrix} 3 & 0 & 3 & 0 \\ 0 & -9 & 0 & 0 \\ 3 & 0 & 3 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

377 **6.5. Extreme examples.** We consider two more 4×2 matrices, both with orthogonal
378 columns, but at the opposite ends in terms of the performance for uniform sampling in Sec-
379 tion 6.2.

380 **Columns of the Hadamard matrix.** With its mass spread uniformly spread, which is
 381 quantified by minimal coherence and uniform leverage scores [12, 16], this matrix is optimal
 382 for uniform sampling,

$$383 \quad \mathbf{X} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \\ 1 & -1 \end{pmatrix} \in \mathbb{R}^{4 \times 2}, \quad \mathbf{P}_x = \mathbf{X}\mathbf{X}^\dagger = \frac{1}{2} \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}.$$

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

$$386 \quad \mathbb{E}_s[\mathbf{P}_0] = \frac{12}{16}\mathbf{I}_2, \quad \mathbb{E}_s[\mathbf{P}\mathbf{P}^T] = \frac{11}{16} \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}.$$

387 The expected deviations of $\mathbf{S}\mathbf{X}$ from full column rank and of \mathbf{P} from being an orthogonal
 388 projector are clearly lower, thus better, than the respective ones in Sections 6.3 and 6.4,

$$389 \quad \|\mathbb{E}_s[\mathbf{I} - \mathbf{P}_0]\|_2 = \frac{4}{16}, \quad \|\mathbb{E}_s[\mathbf{P}\mathbf{P}^T - \mathbf{P}_x]\|_2 = \frac{3}{16}.$$

390 **Columns of the identity matrix.** With its concentrated mass spread, which is quantified
 391 by maximal coherence and widely differing leverage scores [12, 16], this matrix presents a
 392 worst case for a 4×2 matrix of full column rank.

$$393 \quad \mathbf{X} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{pmatrix} \in \mathbb{R}^{4 \times 2}, \quad \mathbf{P}_x = \mathbf{X}\mathbf{X}^\dagger = \frac{1}{2} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

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

$$396 \quad \mathbb{E}_s[\mathbf{P}_0] = \frac{7}{16}\mathbf{I}_2, \quad \mathbb{E}_s[\mathbf{P}\mathbf{P}^T] = \frac{7}{16} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

397 The expected deviations of $\mathbf{S}\mathbf{X}$ from full column rank and of \mathbf{P} from being an orthogonal
 398 projector are

$$399 \quad \|\mathbb{E}_s[\mathbf{I} - \mathbf{P}_0]\|_2 = \frac{9}{16}, \quad \|\mathbb{E}_s[\mathbf{P}\mathbf{P}^T - \mathbf{P}_x]\|_2 = \frac{9}{16},$$

400 thus clearly worse than those for the Hadamard matrix.

Appendix A. Proofs. We present the proofs for Sections 3.1 and 4.

Our results depend on projectors constructed from the possibly rank-deficient matrix \mathbf{SX} . In this case, the Moore-Penrose inverse cannot be expressed in terms of the matrix \mathbf{SX} proper, so we rely on the four conditions [7, Section 5.5.2] that uniquely characterize the Moore-Penrose inverse,

$$(A.1) \quad (\mathbf{SX})(\mathbf{SX})^\dagger(\mathbf{SX}) = \mathbf{SX}, \quad \left((\mathbf{SX})(\mathbf{SX})^\dagger \right)^T = (\mathbf{SX})(\mathbf{SX})^\dagger$$

$$(\mathbf{SX})^\dagger(\mathbf{SX})(\mathbf{SX})^\dagger = (\mathbf{SX})^\dagger, \quad \left((\mathbf{SX})^\dagger(\mathbf{SX}) \right)^T = (\mathbf{SX})^\dagger(\mathbf{SX}).$$

A.1. Proof of Lemma 3.1. The Moore-Penrose conditions (A.1) imply

$$\mathbf{P}^2 = \mathbf{X} \underbrace{(\mathbf{SX})^\dagger \mathbf{S} \mathbf{X} (\mathbf{SX})^\dagger}_{(\mathbf{SX})^\dagger} \mathbf{S} = \mathbf{X} (\mathbf{SX})^\dagger \mathbf{S} = \mathbf{P}.$$

Since $\mathbf{P}^2 = \mathbf{P}$, but \mathbf{P} is not symmetric in general, it is an oblique projector.

1. From (2.1) follows

$$\mathbf{P}_x \mathbf{P} = \mathbf{X} \mathbf{X}^\dagger \mathbf{P} = \underbrace{\mathbf{X} \mathbf{X}^\dagger \mathbf{X}}_{\mathbf{X}} (\mathbf{SX})^\dagger \mathbf{S} = \mathbf{X} (\mathbf{SX})^\dagger \mathbf{S} = \mathbf{P}.$$

2. Use the fact [18, Problem 5.9.12] that $\text{null}(\mathbf{P}) = \text{null}(\mathbf{P}_x)$ if and only if $\mathbf{P} \mathbf{P}_x - \mathbf{P} = \mathbf{0}$ and $\mathbf{P}_x \mathbf{P} - \mathbf{P}_x = \mathbf{0}$. For the latter, the above implies $\mathbf{P}_x \mathbf{P} - \mathbf{P}_x = \mathbf{P} - \mathbf{P}_x$. Thus we can interpret $\mathbf{P} - \mathbf{P}_x$ as a measure for the distance between $\text{null}(\mathbf{P})$ and $\text{null}(\mathbf{P}_x)$.

3. If $\text{rank}(\mathbf{SX}) = p$ then we can express the Moore-Penrose inverse as in (2.1),

$$\mathbf{P} \mathbf{X} = \mathbf{X} \underbrace{\left((\mathbf{SX})^T \mathbf{SX} \right)^{-1} (\mathbf{SX})^T}_{(\mathbf{SX})^\dagger} \mathbf{S} \mathbf{X} = \mathbf{X}.$$

A.2. Proof of Theorem 3.3. The first expression for the least squares solution follows

from (2.1), (2.8), Lemma 3.1, and

$$\tilde{\boldsymbol{\beta}} = \mathbf{X}^\dagger \mathbf{X} \tilde{\boldsymbol{\beta}} = \mathbf{X}^\dagger \mathbf{X} (\mathbf{SX})^\dagger \mathbf{S} \mathbf{y} = \mathbf{X}^\dagger \mathbf{P} \mathbf{y}.$$

Adding $\hat{\boldsymbol{\beta}} - \mathbf{X}^\dagger \mathbf{y} = \mathbf{0}$ from (2.4) to the above gives the second expression

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

where the last equality is due to (A.1) and

$$\mathbf{X}^\dagger = \mathbf{X}^\dagger \mathbf{X} \mathbf{X}^\dagger = \mathbf{X}^\dagger \mathbf{P}_x.$$

Regarding the least squares residual, from (2.9), the first expression for $\tilde{\boldsymbol{\beta}}$, (2.5) and Lemma 3.1 follows

$$\tilde{\mathbf{y}} = \mathbf{X} \tilde{\boldsymbol{\beta}} = \mathbf{X} \mathbf{X}^\dagger \mathbf{P} \mathbf{y} = \mathbf{P}_x \mathbf{P} \mathbf{y} = \mathbf{P} \mathbf{y}.$$

428 Adding $\hat{\mathbf{y}} - \mathbf{P}_x \mathbf{y} = \mathbf{0}$ from (2.6) gives

$$429 \quad \tilde{\mathbf{y}} = \hat{\mathbf{y}} + \mathbf{P} \mathbf{y} - \mathbf{P}_x \mathbf{y} = \hat{\mathbf{y}} + (\mathbf{P} - \mathbf{P}_x) \mathbf{y}.$$

430 As for the predictor, (2.9) and the above expression for $\tilde{\mathbf{y}}$ imply

$$431 \quad \tilde{\mathbf{e}} = \mathbf{y} - \tilde{\mathbf{y}} = (\mathbf{I} - \mathbf{P}) \mathbf{y}.$$

432 Adding and subtracting $\hat{\mathbf{e}} - (\mathbf{I} - \mathbf{P}_x) \mathbf{y} = \mathbf{0}$ from (2.6) gives

$$433 \quad \tilde{\mathbf{e}} = \hat{\mathbf{e}} + (\mathbf{I} - \mathbf{P}) \mathbf{y} - (\mathbf{I} - \mathbf{P}_x) \mathbf{y} = \hat{\mathbf{e}} + (\mathbf{P}_x - \mathbf{P}) \mathbf{y}.$$

434 **A.3. Proof of Corollary 3.5.** The bounds are a direct consequence of Theorem 3.3.

435 From [7, Theorem 5.3.1] follows that $\|\hat{\mathbf{e}}\|_2 / \|\mathbf{y}\|_2 = \sin \theta$. The assumption $\theta < \pi/2$ implies
 436 $\sin \theta < 1$, hence $\|\hat{\mathbf{e}}\|_2 < \|\mathbf{y}\|_2$ and therefore $\hat{\boldsymbol{\beta}} \neq \mathbf{0}$. The assumption $\theta > 0$ implies $\mathbf{y} \notin$
 437 $\text{range}(\mathbf{X})$, thus $\hat{\mathbf{e}} \neq \mathbf{0}$. Therefore we can divide by the appropriate quantities. In the bound
 438 for $\tilde{\mathbf{e}}$, write [7, Theorem 5.3.1]

$$439 \quad \|\mathbf{y}\|_2 / \|\hat{\mathbf{e}}\|_2 = 1 / \sin \theta.$$

440 **A.4. Proof of Lemma 4.1.** The linearity of the mean and (2.2) imply

$$441 \quad \mathbb{E}_{\mathbf{y}}[\mathbf{y}] = \mathbb{E}_{\mathbf{y}}[\mathbf{X} \boldsymbol{\beta}_0] + \mathbb{E}_{\mathbf{y}}[\boldsymbol{\epsilon}] = \mathbf{X} \boldsymbol{\beta}_0 + \mathbf{0} = \mathbf{X} \boldsymbol{\beta}_0$$

$$442 \quad \text{Var}_{\mathbf{y}}[\mathbf{y}] = \text{Var}_{\mathbf{y}}[\mathbf{X} \boldsymbol{\beta}_0 + \boldsymbol{\epsilon}] = \text{Var}_{\mathbf{y}}[\boldsymbol{\epsilon}] = \sigma^2 \mathbf{I}_n.$$

443 From (2.6), the above, and $(\mathbf{P}_x)^2 = \mathbf{P}_x$ follows

$$444 \quad \mathbb{E}_{\mathbf{y}}[\hat{\mathbf{y}}] = \mathbb{E}_{\mathbf{y}}[\mathbf{P}_x \mathbf{y}] = \mathbf{P}_x \mathbb{E}_{\mathbf{y}}[\mathbf{y}] = \mathbf{P}_x \mathbf{X} \boldsymbol{\beta}_0 = \mathbf{X} \boldsymbol{\beta}_0$$

$$445 \quad \text{Var}_{\mathbf{y}}[\hat{\mathbf{y}}] = \text{Var}_{\mathbf{y}}[\mathbf{P}_x \mathbf{y}] = \mathbf{P}_x \text{Var}_{\mathbf{y}}[\mathbf{y}] \mathbf{P}_x = \sigma^2 \mathbf{P}_x.$$

446 From the above, (2.4), and (2.1) follows

$$447 \quad \mathbb{E}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] = \mathbb{E}_{\mathbf{y}}[\mathbf{X}^\dagger \mathbf{y}] = \mathbf{X}^\dagger \mathbb{E}_{\mathbf{y}}[\mathbf{y}] = \mathbf{X}^\dagger \mathbf{X} \boldsymbol{\beta}_0 = \boldsymbol{\beta}_0$$

$$448 \quad \text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] = \text{Var}_{\mathbf{y}}[\mathbf{X}^\dagger \mathbf{y}] = \mathbf{X}^\dagger \text{Var}_{\mathbf{y}}[\mathbf{y}] (\mathbf{X}^\dagger)^T = \sigma^2 (\mathbf{X}^T \mathbf{X})^{-1}.$$

449 **A.5. Proof of Lemma 4.2.** The Moore-Penrose conditions (A.1) imply

$$450 \quad (\mathbf{P}_0)^2 = (\mathbf{S} \mathbf{X})^\dagger \underbrace{(\mathbf{S} \mathbf{X})(\mathbf{S} \mathbf{X})^\dagger (\mathbf{S} \mathbf{X})}_{\mathbf{S} \mathbf{X}} = (\mathbf{S} \mathbf{X})^\dagger (\mathbf{S} \mathbf{X}) = \mathbf{P}_0,$$

451 and $(\mathbf{P}_0)^T = \mathbf{P}_0$, confirming that \mathbf{P}_0 is an orthogonal projector.

452 1. Lemma 3.1 implies $\mathbf{P} \mathbf{X} = \mathbf{X} (\mathbf{S} \mathbf{X})^\dagger \mathbf{S} \mathbf{X} = \mathbf{X} \mathbf{P}_0$.

453 2. If $\text{rank}(\mathbf{S} \mathbf{X}) = p$, then $(\mathbf{S} \mathbf{X})^\dagger$ is a left-inverse, see (2.1), so that $\mathbf{P}_0 = \mathbf{I}_p$.

454 **A.6. Proof of Theorem 4.3.** The expectation follows from Theorem 3.3, Lemma 4.1,
455 Lemma 4.2, and (2.1),

$$456 \quad (\text{A.2}) \quad \mathbb{E}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \mid \mathbf{S}] = \mathbf{X}^\dagger \mathbf{P} \mathbb{E}_{\mathbf{y}}[\mathbf{y}] = \mathbf{X}^\dagger \underbrace{\mathbf{P} \mathbf{X}}_{\mathbf{X} \mathbf{P}_0} \boldsymbol{\beta}_0 = \mathbf{X}^\dagger \mathbf{X} \mathbf{P}_0 \boldsymbol{\beta}_0 = \mathbf{P}_0 \boldsymbol{\beta}_0.$$

457 From the definition of variance, Theorem 3.3, and the above follows

$$458 \quad \begin{aligned} \text{Var}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \mid \mathbf{S}] &= \mathbb{E}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \tilde{\boldsymbol{\beta}}^T \mid \mathbf{S}] - \mathbb{E}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \tilde{\boldsymbol{\beta}}^T \mid \mathbf{S}] \mathbb{E}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \tilde{\boldsymbol{\beta}}^T \mid \mathbf{S}]^T \\ 459 \quad &= \mathbf{X}^\dagger \mathbf{P} \mathbb{E}_{\mathbf{y}}[\mathbf{y} \mathbf{y}^T] \left(\mathbf{X}^\dagger \mathbf{P} \right)^T - (\mathbf{P}_0 \boldsymbol{\beta}_0)(\mathbf{P}_0 \boldsymbol{\beta}_0)^T. \end{aligned}$$

460 For the middle term in first summand, Lemma 4.1 implies

$$461 \quad \begin{aligned} \mathbb{E}_{\mathbf{y}}[\mathbf{y} \mathbf{y}^T] &= (\mathbf{X} \boldsymbol{\beta}_0)(\mathbf{X} \boldsymbol{\beta}_0)^T + \mathbf{X} \boldsymbol{\beta}_0 \mathbb{E}_{\mathbf{y}}[\boldsymbol{\epsilon}]^T + \mathbb{E}_{\mathbf{y}}[\boldsymbol{\epsilon}](\mathbf{X} \boldsymbol{\beta}_0)^T + \mathbb{E}_{\mathbf{y}}[\boldsymbol{\epsilon} \boldsymbol{\epsilon}^T] \\ 462 \quad (\text{A.3}) \quad &= (\mathbf{X} \boldsymbol{\beta}_0)(\mathbf{X} \boldsymbol{\beta}_0)^T + \sigma^2 \mathbf{I}_n, \end{aligned}$$

463 and when inserting this into the leading half of the first summand, one obtains as in (A.2)
464 that

$$465 \quad (\text{A.4}) \quad \mathbf{X}^\dagger \mathbf{P} \mathbf{X} \boldsymbol{\beta}_0 = \mathbf{P}_0 \boldsymbol{\beta}_0.$$

466 This gives the first expression for the conditional variance,

$$467 \quad \begin{aligned} \text{Var}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \mid \mathbf{S}] &= (\mathbf{P}_0 \boldsymbol{\beta}_0)(\mathbf{P}_0 \boldsymbol{\beta}_0)^T + \sigma^2 \mathbf{X}^\dagger \mathbf{P} \mathbf{P}^T (\mathbf{X}^\dagger)^T - (\mathbf{P}_0 \boldsymbol{\beta}_0)(\mathbf{P}_0 \boldsymbol{\beta}_0)^T \\ 468 \quad &= \sigma^2 \mathbf{X}^\dagger \mathbf{P} \mathbf{P}^T (\mathbf{X}^\dagger)^T. \end{aligned}$$

469 To obtain the second expression, multiply the model variance from Lemma 4.1 by $\mathbf{I} =$
470 $(\mathbf{X}^T \mathbf{X})(\mathbf{X}^T \mathbf{X})^{-1}$,

$$471 \quad \begin{aligned} \text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] &= \sigma^2 (\mathbf{X}^T \mathbf{X})^{-1} = \sigma^2 (\mathbf{X}^T \mathbf{X})^{-1} (\mathbf{X}^T \mathbf{X}) (\mathbf{X}^T \mathbf{X})^{-1} \\ 472 \quad &= \sigma^2 (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{P}_x \mathbf{X} (\mathbf{X}^T \mathbf{X})^{-1} = \sigma^2 \mathbf{X}^\dagger \mathbf{P}_x (\mathbf{X}^\dagger)^T, \end{aligned}$$

473 where the remaining equalities follow from $\mathbf{X} = \mathbf{P}_x \mathbf{X}$ in (2.5) and from (2.1). Now add
474 $\text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] - \sigma^2 \mathbf{X}^\dagger \mathbf{P}_x (\mathbf{X}^\dagger)^T = \mathbf{0}$ in the first expression for the variance.

475 If \mathbf{P} were an orthogonal projector onto $\text{range}(\mathbf{X})$, then $\mathbf{P}^T \mathbf{P} = \mathbf{P} = \mathbf{P}_x$. Thus, $\mathbf{P}^T \mathbf{P} - \mathbf{P}_x$
476 represents the deviation of \mathbf{P} from being an orthogonal projector onto $\text{range}(\mathbf{P}_x)$.

477 **A.7. Proof of Corollary 4.4.** The second expression for the variance in Theorem 4.3 and
478 submultiplicativity imply

$$479 \quad \begin{aligned} \|\text{Var}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \mid \mathbf{S}] - \text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]\|_2 &\leq \sigma^2 \|\mathbf{X}^\dagger\|_2 \|\mathbf{P} \mathbf{P}^T - \mathbf{P}_x\|_2 \|(\mathbf{X}^\dagger)^T\|_2 \\ 480 \quad &= \|\mathbf{P} \mathbf{P}^T - \mathbf{P}_x\|_2 \|\text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}]\|_2, \end{aligned}$$

481 where the equality follows from $\|\mathbf{M}\|_2 \|\mathbf{M}^T\|_2 = \|\mathbf{M} \mathbf{M}^T\|_2$, and for any full-column rank
482 matrix \mathbf{M} ,

$$483 \quad \mathbf{M}^\dagger (\mathbf{M}^\dagger)^T = (\mathbf{M}^T \mathbf{M})^{-1} \mathbf{M}^T \mathbf{M} (\mathbf{M}^T \mathbf{M})^{-1} = (\mathbf{M}^T \mathbf{M})^{-1}.$$

484 The second expression for the expectation in Theorem 4.3 and submultiplicativity imply

$$485 \quad \|\mathbb{E}_{\mathbf{y}}[\tilde{\boldsymbol{\beta}} \mid \mathbf{S}] - \boldsymbol{\beta}_0\|_2 \leq \|\mathbf{I} - \mathbf{P}_0\|_2 \|\boldsymbol{\beta}_0\|_2.$$

486 **A.8. Proof of Theorem 4.5.** The expectation follows from sequential conditioning (4.1)
 487 and Lemma 4.3,

$$488 \quad \mathbb{E}[\tilde{\boldsymbol{\beta}}] = \mathbb{E}_{\mathbf{s}} \left[\mathbb{E}_{\mathbf{y}} \left[\tilde{\boldsymbol{\beta}} \mid \mathbf{S} \right] \right] = \mathbb{E}_{\mathbf{s}}[\mathbf{P}_0 \boldsymbol{\beta}_0] = \mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] \boldsymbol{\beta}_0.$$

489 Insert this expression for the mean into the definition of the variance, and apply sequential
 490 conditioning (4.1),

$$491 \quad \begin{aligned} \text{Var}[\tilde{\boldsymbol{\beta}}] &= \mathbb{E}[\tilde{\boldsymbol{\beta}} \tilde{\boldsymbol{\beta}}^T] - \mathbb{E}[\tilde{\boldsymbol{\beta}}] \mathbb{E}[\tilde{\boldsymbol{\beta}}]^T \\ 492 \quad &= \mathbb{E}_{\mathbf{s}} \left[\mathbb{E}_{\mathbf{y}} \left[\tilde{\boldsymbol{\beta}} \tilde{\boldsymbol{\beta}}^T \mid \mathbf{S} \right] \right] - (\mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] \boldsymbol{\beta}_0) (\mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] \boldsymbol{\beta}_0)^T. \end{aligned}$$

493 From Theorem 3.3, (A.3) and (A.3) follows

$$494 \quad \begin{aligned} \mathbb{E}_{\mathbf{y}} \left[\tilde{\boldsymbol{\beta}} \tilde{\boldsymbol{\beta}}^T \mid \mathbf{S} \right] &= \mathbf{X}^\dagger \mathbf{P} \mathbb{E}_{\mathbf{y}}[\mathbf{y} \mathbf{y}^T] \mathbf{P}^T (\mathbf{X}^\dagger)^T \\ 495 \quad &= \mathbf{X}^\dagger \mathbf{P} (\sigma^2 \mathbf{I}_n + (\mathbf{X} \boldsymbol{\beta}_0)(\mathbf{X} \boldsymbol{\beta}_0)^T) \mathbf{P}^T (\mathbf{X}^\dagger)^T \\ 496 \quad &= \sigma^2 \mathbf{X}^\dagger \mathbf{P} \mathbf{P}^T (\mathbf{X}^\dagger)^T + (\mathbf{P}_0 \boldsymbol{\beta}_0)(\mathbf{P}_0 \boldsymbol{\beta}_0)^T. \end{aligned}$$

497 Conditioning this on \mathbf{S} gives

$$498 \quad \mathbb{E}_{\mathbf{s}} \left[\mathbb{E}_{\mathbf{y}} \left[\tilde{\boldsymbol{\beta}} \tilde{\boldsymbol{\beta}}^T \mid \mathbf{S} \right] \right] = \sigma^2 \mathbf{X}^\dagger \mathbb{E}_{\mathbf{s}}[\mathbf{P} \mathbf{P}^T] (\mathbf{X}^\dagger)^T + \mathbb{E}_{\mathbf{s}} \left[(\mathbf{P}_0 \boldsymbol{\beta}_0) (\mathbf{P}_0 \boldsymbol{\beta}_0)^T \right].$$

499 Put everything together to obtain the first expression for the variance,

$$500 \quad \begin{aligned} \text{Var}[\tilde{\boldsymbol{\beta}}] &= \sigma^2 \mathbf{X}^\dagger \mathbb{E}_{\mathbf{s}}[\mathbf{P} \mathbf{P}^T] (\mathbf{X}^\dagger)^T \\ 501 \quad &\quad + \underbrace{\mathbb{E}_{\mathbf{s}} \left[(\mathbf{P}_0 \boldsymbol{\beta}_0) (\mathbf{P}_0 \boldsymbol{\beta}_0)^T \right]}_{\text{Var}_{\mathbf{s}}[\mathbf{P}_0 \boldsymbol{\beta}_0]} - (\mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] \boldsymbol{\beta}_0) (\mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] \boldsymbol{\beta}_0)^T. \end{aligned}$$

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

$$503 \quad \boldsymbol{\beta}_0 \boldsymbol{\beta}_0^T - \boldsymbol{\beta}_0 (\mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] \boldsymbol{\beta}_0)^T - \mathbb{E}_{\mathbf{s}}[\mathbf{P}_0] \boldsymbol{\beta}_0 \mathbb{E}_{\mathbf{s}}[\boldsymbol{\beta}_0]^T.$$

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

506

REFERENCES

- 507 [1] H. AVRON, P. MAYMOUNKOV, AND S. TOLEDO, *Blendenpik: supercharging Lapack's least-squares solver*,
 508 SIAM J. Sci. Comput., 32 (2010), pp. 1217–1236.
 509 [2] C. BOUTSIDIS AND P. DRINEAS, *Random projections for the nonnegative least-squares problem*, Linear
 510 Algebra Appl., 431 (2009), pp. 760–771.
 511 [3] S. CHATTERJEE AND A. S. HADI, *Influential observations, high leverage points, and outliers in linear*
 512 *regression*, Statist. Sci., 1 (1986), pp. 379–416. With discussion.
 513 [4] P. DRINEAS, R. KANNAN, AND M. W. MAHONEY, *Fast Monte Carlo Algorithms for Matrices. I: Approx-*
 514 *imating Matrix Multiplication*, SIAM J. Comput., 36 (2006), pp. 132–157.

- 515 [5] P. DRINEAS, M. W. MAHONEY, AND S. MUTHUKRISHNAN, *Sampling algorithms for l_2 regression and ap-*
516 *plications*, in Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms,
517 ACM, New York, 2006, pp. 1127–1136.
- 518 [6] P. DRINEAS, M. W. MAHONEY, S. MUTHUKRISHNAN, AND T. SARLÓS, *Faster least squares approxima-*
519 *tion*, Numer. Math., 117 (2011), pp. 219–249.
- 520 [7] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The Johns Hopkins University Press, Balti-
521 more, fourth ed., 2013.
- 522 [8] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, second ed., 2002.
- 523 [9] D. C. HOAGLIN AND R. E. WELSCH, *The Hat matrix in regression and ANOVA*, Amer. Statist., 32
524 (1978), pp. 17–22.
- 525 [10] I. C. F. IPSEN, *Relative perturbation results for matrix eigenvalues and singular values*, in Acta Numerica
526 1998, vol. 7, Cambridge University Press, Cambridge, 1998, pp. 151–201.
- 527 [11] I. C. F. IPSEN, *An overview of relative $\sin \Theta$ theorems for invariant subspaces of complex matrices*, J.
528 Comput. Appl. Math., 123 (2000), pp. 131–153. Invited Paper for the special issue *Numerical Analysis*
529 *2000: Vol. III – Linear Algebra*.
- 530 [12] I. C. F. IPSEN AND T. WENTWORTH, *The effect of coherence on sampling from matrices with orthonormal*
531 *columns, and preconditioned least squares problems*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 1490–
532 1520.
- 533 [13] K. LANGE, *Numerical analysis for statisticians*, Statistics and Computing, Springer, New York, sec-
534 ond ed., 2010.
- 535 [14] M. E. LOPES, S. WANG, AND M. W. MAHONEY, *Error estimation for randomized least-squares algorithms*
536 *via the bootstrap*, in Proc. 35th International Conference on Machine Learning, vol. 80, PMLR, 2018,
537 pp. 3217–3226.
- 538 [15] P. MA, M. W. MAHONEY, AND B. YU, *A statistical perspective on algorithmic leveraging*, in Proceedings
539 of the 31st International Conference on International Conference on Machine Learning, vol. 32 of
540 ICML’14, JMLR.org, 2014, pp. I–91–I–99.
- 541 [16] P. MA, M. W. MAHONEY, AND B. YU, *A statistical perspective on algorithmic leveraging*, J. Mach.
542 Learn. Res., 16 (2015), pp. 861–911.
- 543 [17] X. MENG, M. A. SAUNDERS, AND M. W. MAHONEY, *LSRN: a parallel iterative solver for strongly over-*
544 *or underdetermined systems*, SIAM J. Sci. Comput., 36 (2014), pp. C95–C118.
- 545 [18] C. D. MEYER, *Matrix analysis and applied linear algebra*, Society for Industrial and Applied Mathematics
546 (SIAM), Philadelphia, PA, 2000.
- 547 [19] J. F. MONAHAN, *A primer on linear models*, Texts in Statistical Science Series, Chapman & Hall/CRC,
548 Boca Raton, FL, 2008.
- 549 [20] G. RASKUTTI AND M. W. MAHONEY, *A statistical perspective on randomized sketching for ordinary*
550 *least-squares*, J. Mach. Learn. Res., 17 (2016), pp. Paper No. 214, 31.
- 551 [21] V. ROKHLIN AND M. TYGERT, *A fast randomized algorithm for overdetermined linear least-squares re-*
552 *gression*, Proc. Natl. Acad. Sci. USA, 105 (2008), pp. 13212–13217.
- 553 [22] G. W. STEWART, *Collinearity and least squares regression*, Statist. Sci., 2 (1987), pp. 68–100. With
554 discussion.
- 555 [23] G.-A. THANEI, C. HEINZE, AND N. MEINSHAUSEN, *Random Projections For Large-Scale Regression*, 2017,
556 <https://arxiv.org/abs/1701.05325>.
- 557 [24] P. F. VELLEMAN AND R. E. WELSCH, *Efficient computing of regression diagnostics*, Amer. Statist., 35
558 (1981), pp. 234–242.