

Cross-validation: what does it estimate and how well does it do it?

Stephen Bates*, Trevor Hastie† and Robert Tibshirani‡

March 31, 2021

Abstract

Cross-validation is a widely-used technique to estimate prediction error, but its behavior is complex and not fully understood. Ideally, one would like to think that cross-validation estimates the prediction error for the model at hand, fit to the training data. We prove that this is not the case for the linear model fit by ordinary least squares; rather it estimates the average prediction error of models fit on other unseen training sets drawn from the same population. We further show that this phenomenon occurs for most popular estimates of prediction error, including data splitting, bootstrapping, and Mallows’s C_p . Next, the standard confidence intervals for prediction error derived from cross-validation may have coverage far below the desired level. Because each data point is used for both training and testing, there are correlations among the measured accuracies for each fold, and so the usual estimate of variance is too small. We introduce a nested cross-validation scheme to estimate this variance more accurately, and show empirically that this modification leads to intervals with approximately correct coverage in many examples where traditional cross-validation intervals fail. Lastly, our analysis also shows that when producing confidence intervals for prediction accuracy with simple data splitting, one should *not* re-fit the model on the combined data, since this invalidates the confidence intervals.

1 Introduction

When deploying a predictive model, it is important to understand its prediction accuracy on future test points, so both good point estimates and accurate confidence intervals for prediction error are essential. Cross-validation (CV) is a widely-used approach for these two tasks, but in spite of its seeming simplicity, its operating properties remain opaque. Considering first estimation, it turns out to be challenging to precisely state the estimand corresponding to the cross-validation point estimate. In this work, we show that the the estimand of CV is not the accuracy of the model fit on the data at hand, but is instead the average accuracy over many hypothetical data sets. Specifically, we show that the CV estimate of error has larger mean squared error (MSE) when estimating the prediction error of the final model than when estimating the average prediction error of models across many unseen data sets for the special case of linear regression. Turning to confidence intervals for prediction error, we show that naïve intervals based on CV can fail badly, giving coverage far below the nominal level; we provide a simple example soon in Section 1.1. The source of this behavior is the estimation of the variance used to compute the width of the interval: it does not account for the correlation between the error estimates in different folds, which arises because each data point is used for both training and testing. As a result, the estimate of variance is too small and the intervals are too narrow. To address this issue, we develop a modification of cross-validation, *nested cross-validation* (NCV), that achieves coverage near the nominal level, even in challenging cases where the usual cross-validation intervals have miscoverage rates two to three times larger than the nominal rate.

*Depts. of Statistics and EECS, Univ. of California, Berkeley; stephenbates@berkeley.edu

†Depts. of Statistics and Biomedical Data Science, Stanford Univ.; hastie@stanford.edu

‡Depts. of Biomedical Data Science and Statistics, Stanford Univ.; tibs@stanford.edu




Figure 1: A plot of the true error of a model versus the CV estimates for 1000 replicates of the model from Section 1.1. The blue curve shows the average midpoint of the naïve CV confidence intervals. The green bands show the average 90% confidence interval for prediction error given by naïve CV. The red curves show the 5% and 95% quantiles from a quantile regression fit. To achieve nominal coverage, the green curves should approximate the red curves, but they are too narrow in this case.

1.1 A simple illustration

As a motivating example where naïve cross-validation confidence intervals fail, we consider a sparse logistic model

$$P(Y_i = 1 \mid X_i = x_i) = \frac{1}{1 + \exp\{-x_i^\top \theta\}} \quad i = 1, \dots, n, \quad (1)$$

with $n = 90$ observations of $p = 1000$ features, and a coefficient vector $\theta = c \cdot (1, 1, 1, 1, 0, 0, \dots)^\top \in \mathbb{R}^p$ with four nonzero entries of equal strength. The feature matrix $X \in \mathbb{R}^{n \times p}$ is comprised of independent and identically distributed (i.i.d.) standard normal variables, and we chose the signal strength c so that the Bayes misclassification rate is 20%. We estimate the parameters using ℓ_1 -penalized logistic regression with a fixed penalty level. In this case, naïve confidence intervals for prediction error are far too small: intervals with desired miscoverage of 10% give 31% miscoverage in our simulation. We visualize this in Figure 1. The intervals need to be made larger by a factor of about 1.6 to obtain coverage at the desired level in this case.

1.2 Related work

Cross-validation is used ubiquitously to estimate the prediction error of a model [1, 2, 3]. The enduring popularity of CV is due to the fact that it is a conceptually simple improvement over a one-time train-test split [4]. CV is part of a broader landscape of resampling techniques to estimate prediction error, with bootstrap-based techniques as the most common alternative [5, 6, 7, 8]. The other main category of prediction error estimates are based on analytic adjustments such as Mallows’s C_p [9], AIC [10], BIC [11], and general *covariance penalties* [12, 13]. The present work is primarily concerned with CV, but also addresses the properties of bootstrap, data splitting and covariance penalty methods.

In spite of CV’s apparent simplicity, the formal properties of this procedure are subtle; the seemingly basic question “what is cross-validation estimating?” has engendered considerable debate. Although the predictive accuracy of the model fit on the observed training data may seem like a natural estimand, it has been observed that the CV estimator tracks this quantity only weakly, suggesting that CV should instead be treated as an estimator of the average prediction error across training sets [14, 15, 16]. See also [17, 18] for a discussion about different potential estimands. In this work, we discuss this phenomenon in detail for the case of the linear model. Our main result uses a conditional independence argument to explain the aforementioned weak relationship between CV and the instance-specific error.

Turning to the question of inference, one important use of CV is to deliver confidence intervals for the prediction error (or, similarly, an estimate of the standard error) to accompany a point estimate. The second primary goal in this work is to provide such confidence intervals, which cannot be reliably created with naïve methods, as shown in our example in Figure 1. A fundamental prior result shows that there is no unbiased estimator of the standard error of the CV point estimate based on one instance of CV [19]. As a result, to obtain standard error estimates, one would either need to modify the CV procedure or make additional assumptions. Pursuing the former, [20] proposes a sampling schemes where the data is split in half, and CV is carried out within each half separately. This yields an estimate of standard error, but it will typically be much too conservative since the internal CV model fits each use a samples size that is less than half of the full sample. A related proposal due to [21] involves repeatedly performing leave-one-out CV with data sets of half of the original size, but this proposed estimator is not computationally feasible for most learning algorithms.

In a different direction, [20] and [22] propose alternative estimates of standard error, but these are based only on the sample size and higher moments of the errors and so do not address the source of the problem: a covariance term that we describe in Section 4.1. For bootstrap estimators, there are proposals to estimate the standard error of the (bootstrap) point estimates of prediction error with methods based on influence functions [5, 7], and this approach has been partially extended to CV for the special case of the area under the curve measure (AUC) of performance [23]. The CV proposal of [21] similarly involves leave-one-out resampling, which can be interpreted as an empirical estimate of the influence functions.

Accompanying these algorithmic proposals, there is some theoretical understanding of the asymptotic behavior of CV. [24] proves a central limit theorem (CLT) for a cross validation estimator, although it does not come with an estimate of the standard error. [25] provides a consistent estimator for the standard error in the special case of estimating the AUC. Further theoretical results establish the asymptotic normality of the CV estimate in more general cases [21, 26], with proofs relying on notions of algorithmic stability [27, 28, 29]. These works typically use the usual estimate of standard error, which we have seen fails badly in our example in Section 1.1 above. As a result, the asymptotic behavior that these works investigate does not explain the miscoverage phenomenon we observe in small samples—the subject of interest in this work.

Lastly, we note that CV is also used to compare predictive models, such as when selecting a model or a good value of a learning algorithm’s hyperparameters [30, 31, 32, 33, 34]. To this end, [35] suggests a bias-correction for the model selected by cross-validation, [36] shows how to return a confidence set for the best model parameters, and [33, 18] show that selecting the best value of the tuning parameter is a statistically easier goal for CV than estimating the prediction error, in some sense. While we expect that our proposed estimator would be of use for hyperparameter selection because it yields more accurate confidence intervals for prediction error, we do not pursue this problem further in the present work.

1.3 Our contribution

This work has two main thrusts. First, we study the choice of estimand for CV, giving results for the special case of the linear model. We prove a finite-sample conditional independence result (Theorem 1) with a supporting asymptotic result (Theorem 2) that together show that CV does not estimate the error of the specific model fit on the observed training set, but is instead estimating the average error over many training sets (Corollary 2 and Corollary 3). We also show that this holds for the other common estimates of prediction error: data splitting (Section 3.3), Mallows’s C_p (Section 3.4), and bootstrap (Appendix A). Second, we introduce a modified cross-validation scheme to give accurate confidence intervals for prediction error. We prove that our estimate for the MSE of the CV point estimate is unbiased (Theorem 3). Moreover, we validate our method with extensive numerical experiments, confirming that the coverage is consistently better than that of standard cross-validation (Section 5).

2 Setting and notation

We consider the supervised learning setting where we have features $X = (X_1, \dots, X_n) \in \mathcal{X}^n$ and response $Y = (Y_1, \dots, Y_n) \in \mathcal{Y}^n$, and we assume that the data points (X_i, Y_i) for $i = 1, \dots, n$ are i.i.d. from some distribution P . We wish to understand how well fitted models generalize to unseen data points, which we

formalize with a loss function

$$\ell(\hat{y}, y) : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$$

such that $\ell(\hat{y}, y) = 0$ for all y . For example, ℓ could be squared error loss, misclassification error, or deviance (cross-entropy). Now consider a class of models parameterized by θ . Let $\hat{f}(x, \theta)$ be the function that predicts y from $x \in \mathbb{R}^p$ using the model with parameters θ , which takes values in some space Θ . Let \mathcal{A} be a model-fitting algorithm that takes any number of data points and returns a parameter vector $\hat{\theta} \in \Theta$. Let $\hat{\theta} = \mathcal{A}(X, Y)$ be the fitted value of the parameter based on the observed data X and Y . We are interested in the *out-of-sample error* with this choice of parameters:

$$\text{Err}_{XY} := \mathbb{E} \left[\ell(\hat{f}(X_{n+1}, \hat{\theta}), Y_{n+1}) \mid (X, Y) \right],$$

where $(X_{n+1}, Y_{n+1}) \sim P$ is an independent test point from the same distribution. Notice Err_{XY} is a random quantity, depending on the training data. We denote the expectation of this quantity across possible training sets as

$$\text{Err} := \mathbb{E} [\text{Err}_{XY}].$$

We will discuss the relationship between these two quantities further in Section 3. We note that out-of-sample error is materially different from *in-sample-error* which is the focus of methods like the C_p and AIC statistics, and covariance penalties. These are discussed in Section 3.4.

In cross-validation, we partition the observations $\mathcal{I} = \{1, \dots, n\}$ into K disjoint subsets (*folds*) $\mathcal{I}_1, \dots, \mathcal{I}_K$ of size $m = n/K$ at random. Throughout this work, we will assume K divides n for convenience, and we will choose $K = 10$ in all of our numerical results. Consider the first fold, and let $\hat{\theta}^{(-1)} = \mathcal{A}((X_j, Y_j)_{j \in \mathcal{I} \setminus \mathcal{I}_1})$ be the model fit to only those points that are not in fold one. Then, let $e_i = \ell(\hat{f}(x_i, \hat{\theta}^{(-1)}), y_i)$ for each $i \in \mathcal{I}_1$. The errors e_i for points in other folds are defined analogously. We let

$$\widehat{\text{Err}}^{(\text{CV})} := \bar{e} = \frac{1}{n} \sum_{i=1}^n e_i \tag{2}$$

be the average error, which is the usual CV estimate of prediction error. If one desires a confidence interval for the prediction error, a straightforward approach is to compute the empirical standard deviation of the e_i divided by \sqrt{n} to get and estimate of the standard error:

$$\widehat{\text{SE}} := \frac{1}{\sqrt{n}} \cdot \sqrt{\frac{1}{(n-1)} \sum_{i=1}^n (e_i - \bar{e})^2}.$$

From here, we can create a confidence interval as

$$(\bar{e} - z_{1-\alpha/2} \cdot \widehat{\text{SE}}, \bar{e} + z_{1-\alpha/2} \cdot \widehat{\text{SE}}),$$

where $z_{1-\alpha/2}$ is the $1 - \alpha/2$ quantile of the standard normal distribution. We call these the *naïve cross-validation intervals* and they serve as our baseline approach. Importantly, we find that these naïve CV intervals are on average too small because the true sampling standard deviation of \bar{e} is larger than the naïve estimate $\widehat{\text{SE}}$ would suggest, so a better estimate of the standard error is needed.

As a final piece of notation for our asymptotic statements, for a reference sequence $a_1, a_2, \dots, O(a_m)$ denotes a sequence b_1, b_2, \dots such that the sequence $b_1/a_1, b_2/a_2, \dots$ has finite limit superior; $\Omega(a_m)$ denotes a sequence b_1, b_2, \dots such that the sequence $b_1/a_1, b_2/a_2, \dots$ has positive limit inferior; and $\Theta(a_m)$ denotes a sequence that satisfies both properties. Lastly, $o(a_m)$ denotes a sequence b_1, b_2, \dots such that the sequence $b_1/a_1, b_2/a_2, \dots$ converges to zero.

3 What prediction error are we estimating?

We next discuss targets of inference when assessing prediction accuracy. We discuss both Err and Err_{XY} , and also introduce an intermediate quantity Err_X . While cross-validation is our focus, our results hold identically for other estimates of prediction error: covariance penalties (Section 3.4), data splitting (Section 3.3), and bootstrap (Appendix A).




Figure 2: Possible targets of inference for cross-validation. Here, (X, Y) is the training data and Err_{XY} is the average error of the model fit on (X, Y) on a test data set of infinite size. From left to right, the random variables above are a constant, a function of X only, and a function of (X, Y) .

3.1 Err_X : a different target of inference

The two most natural quantities of interest to the analyst are Err_{XY} , the error of the model that was fit on our actual training set, and Err , the average error of the fitting algorithm run on the same-sized datasets drawn from the underlying distribution P . The former quantity is of the most interest to a practitioner deploying a specific model, whereas the latter may be of interest to a researcher comparing different fitting algorithms. While it may initially appear that the quantity Err_{XY} is easier to estimate—since it concerns the model at hand—it has been observed that the cross-validation estimate provides little information about Err_{XY} [14, 15, 16], a phenomenon sometimes called the *weak correlation* issue.

We now prove that CV has lower MSE for estimating Err than it does for Err_{XY} , for the special case of the linear model. In this sense, CV should be viewed as an estimate of Err rather than of Err_{XY} . In order to state this formally, for this section only, assume the homoskedastic linear model holds:

$$y_i = x_i^\top \theta + \epsilon_i \quad \text{where} \quad \epsilon_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2) \quad i = 1, \dots, n. \quad (3)$$

In this setting, a key quantity in our analysis is

$$\text{Err}_X := \mathbb{E}[\text{Err}_{XY} \mid X],$$

which falls between Err and Err_{XY} ; see Figure 2 for a visualization. This quantity is also considered by [37] in a high-dimensional regression setting, but to the best of our knowledge has not been considered in the literature on estimation of prediction error.

While our current focus is on cross-validation, the conclusions hold for a broad class of estimates of prediction error. In particular, we consider estimators of prediction error that satisfy the following property:

Definition 1 (Linearly invariant estimator). *We say that an estimator of prediction error $\widehat{\text{Err}}((X_1, Y_1), \dots, (X_n, Y_n), U)$ is linearly invariant if*

$$\widehat{\text{Err}}((x_1, y_1), \dots, (x_n, y_n), u) = \widehat{\text{Err}}((x_1, y_1 + x_1^\top \kappa), \dots, (x_n, y_n + x_n^\top \kappa), u). \quad (4)$$

Here κ is any p -vector and the random variable U is included to allow for randomized procedures like cross-validation, and without loss of generality it is taken to be $\text{unif}[0, 1]$ and independent of (X, Y) .

With OLS fitting, cross-validation satisfies this property:

Lemma 1. *When using OLS as the fitting algorithm and squared-error loss, the cross-validation estimate of prediction error, $\widehat{\text{Err}}^{(\text{CV})}$, is linearly invariant.*

Note that linear invariance is a deterministic property of an estimator and does not rely on any distributional assumptions.

Recall from classical linear regression theory that when using ordinary least squares (OLS), the estimated coefficient vector is independent of the residual sum of squares. This implies that the sum of squared residuals is independent of the true predictive error. It turns out that even further, the CV estimate of error (and all linearly invariant estimates of error) is independent of the true error, conditional on the feature matrix X .

Theorem 1. *Assume the homoskedastic Gaussian linear model (3) holds and that we use squared-error loss. Let $\widehat{\text{Err}}$ be a linearly invariant estimate of prediction error (such as $\widehat{\text{Err}}^{(\text{CV})}$ using OLS as the fitting algorithm). Then,*

$$\widehat{\text{Err}} \perp\!\!\!\perp \text{Err}_{XY} \mid X. \quad (5)$$




Figure 3: Left: mean squared error of the CV point estimate of prediction error relative to three different estimands: Err , Err_X , and Err_{XY} . Center: coverage of Err , Err_X , and Err_{XY} by the naïve cross-validation intervals in a homoskedastic Gaussian linear model. The nominal miscoverage rate is 10%. Each pair of points connected by a line represents 2000 replicates with the same feature matrix X . Right: 2000 replicates with the same feature matrix and the line of best fit (blue).

As a result, any linearly invariant estimator (such as cross-validation) has lower MSE as an estimate of Err_X than as an estimate of Err_{XY} :

Corollary 1. *Under the conditions of Theorem 1,*

$$\mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_{XY} \right)^2 \right] = \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X \right)^2 \right] + \underbrace{\mathbb{E} [\text{var}(\text{Err}_{XY} | X)]}_{\geq 0}.$$

We demonstrate this in an experiment in a simple linear model with $n = 100$ observations and $p = 20$ features, where the features are i.i.d. standard normal variables; see Figure 3. As predicted by Corollary 1, we see that the CV point estimate has lower MSE for Err_X than for Err_{XY} . Similarly, the naïve CV intervals cover Err_X more often than they cover Err_{XY} .

Remark 1. *Theorem 1 can be restated in an evocative way. Suppose we have two data sets (X, Y) and (X, Y') that share the same feature matrix X . Let Err_{XY} and $\text{Err}_{XY'}$ be the true errors of the model fit with these two data sets, respectively. Next, suppose we perform cross-validation on (X, Y) to get an estimate $\widehat{\text{Err}}_{XY}$ and do the same on (X, Y') to get an estimate $\widehat{\text{Err}}_{XY'}$. Then, (5) is equivalent to*

$$(\text{Err}_{XY}, \widehat{\text{Err}}_{XY}) \stackrel{d}{=} (\text{Err}_{XY}, \widehat{\text{Err}}_{XY'}).$$

This means that for the purpose of estimating Err_{XY} , we have no reason to prefer using the cross-validation estimate with (X, Y) to using the cross-validation estimate with a different data set (X, Y') , even though we wish to estimate the error of the model fit on (X, Y) .

3.2 Relationship with average error

The results of the previous section suggest that Err_X is a more natural target of inference than Err_{XY} . Next, we examine the relationship between Err and Err_X , showing that Err_X is close to Err , in that the variance of Err_X (which has mean Err) is small compared with the variance of Err_{XY} (which also has mean Err). Combined with the results of the previous section, this gives a formal statement that cross-validation is a better estimator for Err than for Err_{XY} .

To make this precise, consider the conditional variance decomposition of the variance of Err_{XY} ,

$$\text{var}(\text{Err}_{XY}) = \underbrace{\mathbb{E}_X [\text{var}(\text{Err}_{XY} | X)]}_{\text{var due to } Y | X} + \underbrace{\text{var}(\text{Err}_X)}_{\text{var due to } X}. \quad (6)$$




Figure 4: The relationship among various notions of prediction error in the proportional asymptotic limit (7). Recall that σ^2 is the Bayes error: the error rate of the best possible model. See Figure 5 for a simulation experiment demonstrating these rates. See Remark 3 for details about the rate of $\widehat{\text{Err}}$.

To quantify the relative contribution of the two terms in the right-hand side of (6), we will use a *proportional asymptotic limit*, where

$$n > p, \quad n, p \rightarrow \infty, \quad n/p \rightarrow \lambda > 1. \quad (7)$$

We use the proportional asymptotic limit rather than traditional p fixed, $n \rightarrow \infty$ asymptotics, because in the latter asymptotic regime, the difference between Err , Err_X , and Err_{XY} is asymptotic order lower than $1/\sqrt{n}$, so one always estimates these three targets with equal precision, and the analysis is less informative. See Appendix F for a complementary analysis in the traditional p fixed, $n \rightarrow \infty$ asymptotic regime and [33, 18] for a related discussion. By contrast, in the proportional asymptotic limit we will see that $\widehat{\text{Err}}^{(\text{CV})}$ is closer to Err and Err_X than to Err_{XY} .

Theorem 2. *Suppose the homoskedastic Gaussian linear model in (3) holds and that we use squared-error loss. In addition, assume that feature vectors $X_i \sim \mathcal{N}(0, \Sigma_p)$ for any full-rank Σ_p . Then, in the proportional asymptotic limit in (7), we have*

$$\mathbb{E}_X [\text{var}(\text{Err}_{XY} | X)] = \Theta(1/n)$$

and

$$\text{var}(\text{Err}_X) = \mathbb{E}(\text{Err}_X - \text{Err})^2 = \Theta(1/n^2),$$

as $n, p \rightarrow \infty$.

We summarize the asymptotic relationship among the various estimands in Figure 4. We see that the randomness caused by Y given X is of a larger order than that due to the randomness in X . This explains why in Figure 3, the coverage and MSE of cross-validation is similar when estimating either Err or Err_X , but is significantly different when estimating Err_{XY} . As a result, Err_X and Err_{XY} are asymptotically uncorrelated, and moreover, combining this with Theorem 1 shows that $\widehat{\text{Err}}^{(\text{CV})}$ is asymptotically uncorrelated with Err_{XY} , as stated next.

Corollary 2. *In the setting of Theorem 2,*

$$\text{cor}(\text{Err}_{XY}, \text{Err}_X) \rightarrow 0 \quad \text{as } n, p \rightarrow \infty.$$

Moreover, for any linearly invariant estimator $\widehat{\text{Err}}$ (such as $\widehat{\text{Err}}^{(\text{CV})}$ using OLS as the fitting algorithm),

$$\text{cor}(\text{Err}_{XY}, \widehat{\text{Err}}) \rightarrow 0 \quad \text{as } n, p \rightarrow \infty.$$

Notice that this is a marginal result, whereas the similar Theorem 1 is conditional on X . With respect to Figure 4, this result means that the fluctuations of Err_{XY} around Err are asymptotically uncorrelated with the fluctuations of $\widehat{\text{Err}}$ around Err . Combining Theorem 2 with Theorem 1, we conclude that CV has larger error for estimating Err_{XY} than for Err or Err_X :

Corollary 3. *In the setting of Theorem 2, let $\widehat{\text{Err}}$ be any linearly invariant estimator (such as $\widehat{\text{Err}}^{(\text{CV})}$ using OLS as the fitting algorithm). Suppose in addition that $\text{var}(\widehat{\text{Err}}) \rightarrow 0$ (an extremely weak condition satisfied*

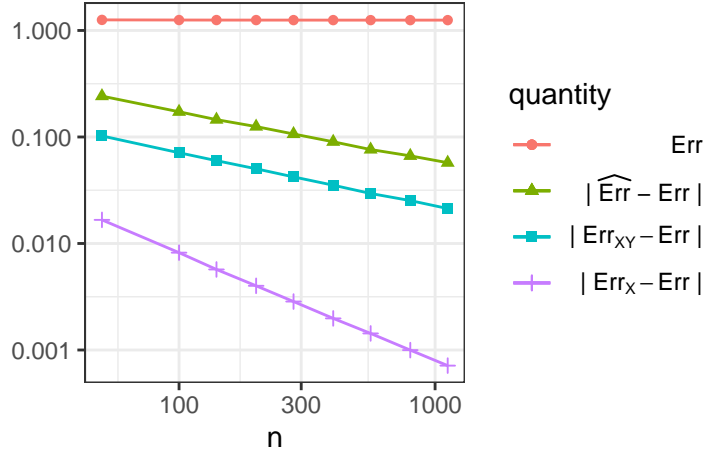


Figure 5: Simulation results demonstrating the asymptotic scaling presented in Figure 4. The fitted slopes of the lines (after log-transforming both axes) are 0.00, -0.46 , -0.50 , -1.01 , from top to bottom.

by any reasonable estimator). Then,

$$\begin{aligned} \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_{XY} \right)^2 \right] - \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X \right)^2 \right] &= \Omega(1/n), \\ \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_{XY} \right)^2 \right] - \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err} \right)^2 \right] &= \Omega(1/n), \quad \text{and} \\ \left| \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err} \right)^2 \right] - \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X \right)^2 \right] \right| &= o(1/n). \end{aligned}$$

The asymptotic theory perfectly predicts the experimental results presented in Figure 5; we see that even for moderate sample size, the scalings is exactly as anticipated. The main conclusion is that for a linearly invariant estimate of prediction error that has precision $1/\sqrt{n}$, our results show that asymptotically one has lower estimation error when estimating Err compared to Err_{XY} . Similarly, the correlation between a linearly invariant estimate and Err_{XY} goes to zero. These theoretical predictions are also corroborated by the experimental results presented in Figure D.5. Thus, cross-validation is estimating the average error Err more so than the specific error Err_{XY} .

Remark 2. Note that the results in this section apply both to K -fold cross-validation with fixed K , and leave-one-out cross validation where $K = n$. (Formally, the results assume only that one is using some sequence of linearly invariant estimators.)

Remark 3. Our focus here is on the possible estimands, and we have not explicitly proved that CV has precision $1/\sqrt{n}$ in the current regime. Nonetheless, we expect this to be the case for leave-one-out CV and K -fold CV with K growing. (It is not the case with fixed K due to the difference in sample size, but this is not the main topic of concern in this work.) Figure 5 shows that this scaling approximately holds for moderate n .

3.3 Data splitting

Perhaps the simplest way to estimate prediction error is to split the data into two disjoint sets, one for training and one for estimating the prediction accuracy. The previous results also shed light on the properties of data splitting. In particular, we will show that when estimating prediction error with data splitting, one should *not* re-fit the model on the full data, because this incurs additional variance that invalidates the confidence intervals for prediction error. We make this precise next.

3.3.1 Data splitting without refitting

We first consider *data splitting without refitting*. We partition the data into disjoint sets $\mathcal{I}^{(\text{train})} \cup \mathcal{I}^{(\text{out})} = \{1, \dots, n\}$, and fit the model on the training set $\hat{\theta}^{(\text{train})} = \mathcal{A}((X_i, Y_i)_{i \in \mathcal{I}^{(\text{train})}})$. As in cross-validation, we then let $e_i = \ell(\hat{f}(x_i, \hat{\theta}^{(\text{train})}), y_i)$ for $i \in \mathcal{I}^{(\text{out})}$. We can then estimate the prediction error as

$$\widehat{\text{Err}}^{(\text{split})} := \frac{1}{|\mathcal{I}^{(\text{out})}|} \sum_{i \in \mathcal{I}^{(\text{out})}} e_i \quad (8)$$

and give a valid estimate of its standard error as

$$\widehat{\text{SE}}^{(\text{split})} := \sqrt{\frac{1}{|\mathcal{I}^{(\text{out})}| - 1} \sum_{i \in \mathcal{I}^{(\text{out})}} (e_i - \widehat{\text{Err}}^{(\text{split})})^2}.$$

We define $\tilde{X} = (X_i)_{i \in \mathcal{I}^{(\text{train})}}$ and $\tilde{Y} = (Y_i)_{i \in \mathcal{I}^{(\text{train})}}$. Then $\widehat{\text{Err}}^{(\text{split})}$ is an unbiased estimate for

$$\text{Err}_{\tilde{X}\tilde{Y}} := \mathbb{E} \left[\ell \left(\hat{f}(X_{n+1}, \hat{\theta}^{(\text{train})}), Y_{n+1} \right) \mid (\tilde{X}, \tilde{Y}) \right], \quad (9)$$

where the expectation is only over a fresh test point (X_{n+1}, Y_{n+1}) . The advantage of this approach is that one can accurately estimate the prediction error and provide valid inference. A first disadvantage is the prediction error and standard error estimates are valid for the model trained only on the subset $\mathcal{I}^{(\text{train})}$. Secondly, the estimates of prediction error have reduced precision because they rely only on the subset $\mathcal{I}^{(\text{out})}$. Thus, confidence intervals for prediction error may be much wider than those from CV. We confirm this in experiments in Section 5.

3.3.2 Data splitting with refitting

A slightly different form of data splitting is also commonly used, which we will refer to as *data splitting with refitting*. In this approach, one follows the same steps as above to obtain the estimator $\widehat{\text{Err}}^{(\text{split})}$ from (8), but then conducts a final refitting step to obtain $\hat{\theta} = \mathcal{A}(X, Y)$, the model fit on the full data. Here, one deploys the model on the full data, $\hat{f}(\cdot, \hat{\theta})$ —the idea is that the model on the full data is superior to the model fit using only the training subset. Data-splitting with refitting is similar to cross-validation, and analogs of our results from Section 3 carry over to this setting.

Lemma 2. *The estimator $\widehat{\text{Err}}^{(\text{split})}$ is linearly invariant.*

Thus, the conclusions of Theorem 1, Corollary 1, Corollary 2, and Corollary 3 hold for $\widehat{\text{Err}}^{(\text{split})}$. In particular, $\widehat{\text{Err}}^{(\text{split})}$ has lower error for estimating Err than for estimating Err_{XY} , and $\widehat{\text{Err}}^{(\text{split})}$ is asymptotically uncorrelated with Err_{XY} . Moreover, the standard error estimate $\widehat{\text{SE}}^{(\text{split})}$ becomes invalid when refitting, and asymptotically it is too small, as stated next.

Proposition 1. *In the setting of Theorem 1,*

$$\mathbb{E} \left[\left(\widehat{\text{SE}}^{(\text{split})} \right)^2 \right] = \mathbb{E} \left[\left(\widehat{\text{Err}}^{(\text{split})} - \text{Err} \right)^2 \right] - \underbrace{\text{var}(\text{Err}_{\tilde{X}\tilde{Y}})}_{\geq 0} - \underbrace{(\text{Err} - \mathbb{E}[\text{Err}_{\tilde{X}\tilde{Y}}])^2}_{\geq 0}. \quad (10)$$

This is a non-asymptotic result. In the proportional asymptotic limit (7), for both sides of (10) above all terms are of order $1/n$, except for the final term on the right-hand side, which is of constant order.¹ This means that the standard error estimate derived from data splitting is too small, and confidence intervals based on this number will have coverage that is too small. Importantly, this behavior is *not* only due to the

¹It is of constant order when the test set $\mathcal{I}^{(\text{out})}$ is a constant fraction of n . This can be made smaller by choosing the ratio of $\mathcal{I}^{(\text{out})}$ to n converge to 0 at some rate, but then data splitting will not even achieve a $1/\sqrt{n}$ rate of precision.

difference in sample size used to fit $\hat{\theta}^{(\text{split})}$ and $\hat{\theta}$. The final term in (10) is the result of the difference in sample size, but even without this term, the estimate is too small asymptotically (due to the middle term on the right-hand side of (10)). The standard error estimate $\widehat{\text{SE}}^{(\text{split})}$ is similarly too small if one wishes to estimate $(\widehat{\text{Err}}^{(\text{split})} - \text{Err}_{XY})^2$; see Proposition 3 in Appendix C. In Figure D.8, we verify with an experiment that the data splitting intervals do not have coverage approaching the nominal level, even as n and p grow. (Naïve cross-validation has a similar miscoverage problem—see Figure D.6 and Figure D.7.)

In summary, data splitting with refitting has similar behavior to cross-validation; the point estimate is really an estimate of average prediction error and the standard error estimate is too small in the proportional asymptotic limit. Thus, data splitting resolves the inferential challenges of cross-validation only if one does not refit the model on the full data. In that case, one sacrifices efficiency both in the model fitting and in the precision of the prediction error estimates. This may be a viable approach when one has many samples available, but the loss of efficiency is likely to be too costly for data sets comprised of fewer than, say, a thousand samples. With limited data, another approach is necessary. We explore the loss of efficiency due to data splitting in simulation experiments in Section 5.

3.4 Connection with covariance penalties

For parametric models, there is an alternative theory of the estimation of prediction accuracy based on *covariance penalties*; see [12, 13, 17] for overviews of this approach. For the linear model with OLS and squared error loss, this approach specializes to the well-known Mallows C_p [9, 10] estimate of prediction error:

$$\widehat{\text{Err}}^{(C_p)} := \frac{1}{n} \sum_{i=1}^n (y_i - \hat{f}(x_i, \hat{\theta}))^2 + \frac{2p\hat{\sigma}^2}{n}.$$

The classical covariance penalty approach is focused on estimating *in-sample error*, the error for a fresh sample with the same features X :

$$\text{Err}_{\text{in}}(X) := \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (Y'_i - \hat{f}(X_i, \hat{\theta}))^2 \mid X \right],$$

where the expectation is only over Y_i and Y'_i for $i = 1, \dots, n$, and Y_i, Y'_i are independent draws from the distribution of $Y_i \mid X_i = x_i$. See Figure 6 for a visualization of how this relates our other notions of prediction error. The term *covariance penalty* comes from the following identity [13, 38]:

$$\text{Err}_{\text{in}}(X) = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (y_i - \hat{f}(x_i, \hat{\theta}))^2 \mid X \right] + \frac{2}{n} \sum_{i=1}^n \text{cov}(y_i, \hat{f}(x_i, \hat{\theta}) \mid X). \quad (11)$$

To extend this to the setting where the feature vector of future test points is also random (the setting of the present work), [17] introduce RC_p , which is a similar but slightly larger estimate of prediction error that accounts for the variability in the features of future test points:

$$\widehat{\text{Err}}^{(RC_p)} := \frac{1}{n} \sum_{i=1}^n (y_i - \hat{f}(x_i, \hat{\theta}))^2 + \frac{p\hat{\sigma}^2}{n} \left(2 + \frac{p+1}{n-p-1} \right).$$

When the linear model holds and the features are multivariate Gaussian, $\widehat{\text{Err}}^{(RC_p)}$ is an unbiased estimate of Err [17].

3.4.1 Estimand of C_p

We first discuss the estimation of prediction error with Mallows C_p , showing that (like cross-validation) it is a worse estimator for Err_{XY} than for Err . The results from Section 3.1 and Section 3.2 continue to hold for $\widehat{\text{Err}}^{(C_p)}$ and $\widehat{\text{Err}}^{(RC_p)}$, since they are linearly invariant:

Lemma 3. *The estimators $\widehat{\text{Err}}^{(C_p)}$ and $\widehat{\text{Err}}^{(RC_p)}$ are linearly invariant.*

		Test X	
		same	fresh
Training X	fixed	Err _{in}	Err _X
	random	Err _S	Err

Figure 6: A representation of the notions of error considered in this work. See [17] for a definition and discussion of Err_S.

This result is immediate from the fact the $\widehat{\text{Err}}^{(C_p)}$ and $\widehat{\text{Err}}^{(RC_p)}$ are functions only of the residuals of the OLS fit. Thus, the conclusions of Theorem 1, Corollary 1, Corollary 2, and Corollary 3 hold for $\widehat{\text{Err}}^{(C_p)}$. In particular, $\widehat{\text{Err}}^{(C_p)}$ has lower error for estimating Err and Err_X than for estimating Err_{XY}, and $\widehat{\text{Err}}^{(C_p)}$ is asymptotically uncorrelated with Err_{XY}. In summary, as before with cross-validation, Mallows's C_p is not able to estimate Err_{XY}, but is rather an estimate of Err or Err_X or Err (the latter two are close for large samples).

3.4.2 A decomposition of Err_X

Next, we develop a decomposition for Err_X. The results that we present are essentially implicit in [17], but in that work the results are averaged over X to instead obtain estimates for Err. From the definitions of Err_X and Err_{in}, we trivially have that

$$\text{Err}_X = \text{Err}_{\text{in}}(X) + \mathbb{E} \left[\mathbb{E} \left[(Y_{n+1} - \hat{f}(X_{n+1}, \hat{\theta}))^2 \mid (X, Y) \right] - \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (Y_i' - \hat{f}(X_i, \hat{\theta}))^2 \mid (X, Y) \right] \mid X \right]. \quad (12)$$

When the linear model holds, this simplifies as stated next.

Proposition 2. *For the linear model with the OLS fitting algorithm and squared error loss, assume in addition that the distribution of X_i has mean zero and covariance Σ of full rank. Then,*

$$\text{Err}_X = \text{Err}_{\text{in}}(X) + \frac{\sigma^2}{n} \left(\text{tr}(\hat{\Sigma}^{-1}\Sigma) - p \right),$$

where $\hat{\Sigma} = X^\top X/n$.

The second term in the sum can be either positive or negative. Roughly speaking, this term is smaller (more negative) if X is a good design that yields precise estimates of the regression coefficients, whereas this term is larger (more positive) if X yields less precise estimates. Note that we do not typically know the population covariance Σ , so this cannot be used as an estimator for prediction error. Instead, it is an expository decomposition relating Err_X with existing work about the estimation of prediction error. From this expression, we can read off the following results from [17].

Corollary 4 (Random-X covariance penalties [17]). *In the setting of Proposition 2, we have that*

$$\text{Err} \geq \mathbb{E}[\text{Err}_{\text{in}}].$$

Moreover, if the feature vector X_1 follows a multivariate Gaussian distribution, then

$$\mathbb{E}[\text{Err}_X] = \mathbb{E}[\text{Err}_{\text{in}}] + \frac{p\sigma^2}{n} \left(\frac{p+1}{n-p-1} \right).$$

The latter expression is the motivation for the RC_p penalty.

Remark 4. *The results of this section have focused on the case of squared error loss. Relation (12) of course holds for any loss function. Interestingly, the covariance formula (11) holds for any “q class” of error measures ([13]); these include misclassification rate and binomial deviance. These may lead to simplifications of (12) similar to what were achieved for squared-error loss.*

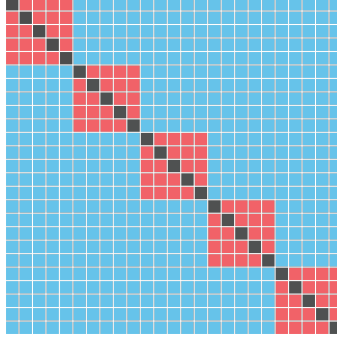


Figure 7: Covariance structure of the CV errors. Red entries correspond to the covariance between points in the same fold, and blue entries correspond to the covariance between points in different folds.

3.5 Bootstrap estimates of prediction error

Bootstrap estimates of prediction error are also linearly invariant, and so they are also estimates of the average prediction error. For brevity, we present these results in Appendix A.

4 Confidence intervals with nested cross-validation

In this section, we describe a modification of cross-validation that comes with an estimate of the MSE of the point estimate for prediction error. Our ultimate goal is then to use the estimated MSE to give confidence intervals for prediction error with approximately valid coverage.

4.1 Dependence structure of CV errors

Before developing our estimator for the cross-validation MSE, we pause here to build up intuition for why the naïve CV confidence intervals for prediction error can fail, as seen previously in our example in Section 1.1. The naïve CV intervals are too small, on average, because the true sampling variance of $\widehat{\text{Err}}^{(\text{CV})}$ is larger than the naïve estimate $\widehat{\text{SE}}$ would suggest. In particular, this estimate of the variance of the CV point estimate assumes that the observed errors e_1, \dots, e_n are independent. This is not true—the observed errors have less information than an independent sample since each point is used for both training and testing, which induces dependence among these terms.

In particular, the covariance matrix of the errors e_1, \dots, e_n has the block structure shown in Figure 7. As such, it is parameterized by only three numbers: $a_1 := \text{var}(e_1)$, $a_2 := \text{cov}(e_i, e_j)$ for i, j in the same fold, and $a_3 := \text{cov}(e_i, e_j)$ for i, j in different folds. It is easy to see that the variance of \bar{e} is

$$\text{var}(\bar{e}) = \frac{1}{n}a_1 + \frac{n/K - 1}{n}a_2 + \frac{n - n/K}{n}a_3; \quad (13)$$

see [19]. The constants a_2 and a_3 will typically be positive, in which case $\text{var}(\bar{e}) > a_1/n$, and so estimating the variance of \bar{e} as $\widehat{\text{SE}}^2$ results in an estimate that is too small. For example, in the setting from Section 1.1, the estimated variance is approximately a factor of 2.65 too small, so the naïve confidence intervals are too small by a factor of $\sqrt{2.65} \approx 1.6$. In order to correctly estimate the variance of \bar{e} , it would suffice to estimate a_1 , a_2 , and a_3 , but [19] proves surprising fact that there is no unbiased estimator of $\text{var}(\bar{e})$ based on a single run of cross-validation. Thus, estimating a_1 , a_2 , and a_3 cannot be done from a single run of cross-validation. Although this result suggests that something beyond the usual cross-validation will be required to give good estimates of the standard error of $\widehat{\text{Err}}^{(\text{CV})}$, it does not imply that it is impossible to get such estimates with other approaches.

To recap, the primary issue with the naïve cross-validation confidence intervals is that they rely on an independence assumption that is violated: they implicitly assume that $a_2 = 0$ and $a_3 = 0$. Thus, the usual estimate of the variance of $\widehat{\text{Err}}^{(\text{CV})}$ is too small, resulting in poor coverage. To remedy this issue, we develop

an estimator that empirically estimates the variance of $\widehat{\text{Err}}^{(\text{CV})}$ across many subsamples. Avoiding the faulty independence approximation leads to intervals with superior coverage. We turn to details of our proposed procedure next.

4.2 Our target of inference

From here onwards, our primary goal will be to estimate the mean-squared error (MSE) of cross-validation:

Definition 2. *For a sample of size n split into K folds, the cross-validation MSE is*

$$\text{MSE}_{K,n} := \mathbb{E} \left[\left(\widehat{\text{Err}}^{(\text{CV})} - \text{Err}_{XY} \right)^2 \right]. \quad (14)$$

The reader may wonder why we define the MSE with respect to Err_{XY} in view of the results from Section 3 that show that $\widehat{\text{Err}}^{(\text{CV})}$ is targeting Err , and we briefly comment on this. While the point estimate $\widehat{\text{Err}}^{(\text{CV})}$ is targeting Err in that the variation of Err_{XY} around Err is not accessible to CV (at least in the linear model), this does not preclude estimating the MSE defined as in (14). Due to a convenient decomposition in Lemma 4, we will be able to estimate this quantity without any modeling assumptions, whereas we do not know how to estimate (14) with Err in place of $\widehat{\text{Err}}$ without further assumptions. In any case, confidence intervals for Err_{XY} are often of the most interest to the analyst.

The MSE in (14) contains both a bias term and variance term, but the bias is typically small for cross-validation [5, 39, 7]. Thus, we can view the MSE as a slightly conservative version of the variance of the cross-validation estimator, i.e., the squared standard error. With this in mind, we will use an estimate of the MSE to construct confidence intervals for Err_{XY} .

We now give a generic decomposition of the mean-squared error of an estimate of prediction error, which will enable use to estimate $\text{MSE}_{K,n}$. Consider a single split of the data into a training set and holdout set, i.e., we partition $\mathcal{I} = \{1, \dots, n\}$ into $\mathcal{I}_{(\text{train})}$ and $\mathcal{I}_{(\text{out})}$ calling the training set (\tilde{X}, \tilde{Y}) . Using only (\tilde{X}, \tilde{Y}) , we use our fitting procedure to obtain estimated parameters $\hat{\theta}^{(\text{train})} = \mathcal{A}(\tilde{X}, \tilde{Y})$, and further assume we have some estimate of $\widehat{\text{Err}}_{\tilde{X}\tilde{Y}}$ of the prediction error $\text{Err}_{\tilde{X}\tilde{Y}}$ defined in (9). Here, $\widehat{\text{Err}}_{\tilde{X}\tilde{Y}}$ is any estimator of $\text{Err}_{\tilde{X}\tilde{Y}}$ based only on (\tilde{X}, \tilde{Y}) , such as cross-validation using only (\tilde{X}, \tilde{Y}) . Let $\{e_i^{(\text{out})}\}_{i \in \mathcal{I}_{(\text{out})}}$ be the losses of the fitted model $\hat{f}(\cdot, \hat{\theta}^{(\text{train})})$ on the holdout set, and let $\bar{e}^{(\text{out})}$ be their average. The MSE of $\widehat{\text{Err}}_{\tilde{X}\tilde{Y}}$ can be written as follows:

Lemma 4 (Holdout MSE identity). *In the setting above*

$$\underbrace{\mathbb{E} \left[\left(\widehat{\text{Err}}_{\tilde{X}\tilde{Y}} - \text{Err}_{\tilde{X}\tilde{Y}} \right)^2 \right]}_{\text{MSE}} = \underbrace{\mathbb{E} \left[\left(\widehat{\text{Err}}_{\tilde{X}\tilde{Y}} - \bar{e}^{(\text{out})} \right)^2 \right]}_{(a)} - \underbrace{\mathbb{E} \left[\left(\text{Err}_{\tilde{X}\tilde{Y}} - \bar{e}^{(\text{out})} \right)^2 \right]}_{(b)}. \quad (15)$$

The expectations above are over the complete data (X, Y) . The lemma follows from adding and subtracting $\text{Err}_{\tilde{X}\tilde{Y}}$ within term (a) then noting the cross-term is zero due to the independence of the training and test set.

This identity is of interest, since both (a) and (b) can be estimated from the data, which leads to an estimate of the MSE term. Specifically, (a) is an observed quantity and (b) is the variance of $\bar{e}^{(\text{out})}$, which can be estimated. Thus, for a single random split of the data we can give an unbiased estimate of the MSE of our estimate of prediction error. This suggests the following estimation strategy: we repeatedly split the data and then average together the observed (a) terms and estimates of (b). In the next section, we will pursue this strategy for the particular case where $\widehat{\text{Err}}_{\tilde{X}\tilde{Y}}$ is itself a cross-validation estimate based only on (\tilde{X}, \tilde{Y}) .




Figure 8: Visualization of nested CV. Using only the folds on left of the vertical line, we perform the usual cross-validation by holding out one fold at a time (the dark grey fold) and then fitting on the remaining folds (the light grey folds). The fresh holdout points (the blue fold) are never used in the inner CV step.

4.3 A nested CV estimate of MSE

Building from Lemma 4, we now turn to our proposed estimate of MSE. In our algorithm, we set aside one fold of data and run $(K - 1)$ -fold CV on the remaining data. We repeat this many times, which gives an estimate of (a) by taking the empirical mean of

$$\left(\widehat{\text{Err}}_{\bar{X}\bar{Y}} - \bar{e}^{(\text{out})}\right)^2$$

across the many splits and an estimate of (b) by taking the empirical variance of $\bar{e}^{(\text{out})}$ across the many splits. By (15), the difference of these two terms is an estimate of the MSE. We also get a point estimate of error by taking the empirical mean of $\widehat{\text{Err}}_{\bar{X}\bar{Y}}$ across the many splits. See Figure 8 for a visualization of the nested CV sample splitting. This procedure is described explicitly in Algorithm 1. We denote the resulting estimate of mean squared error by $\widehat{\text{MSE}}$ and the point estimate for prediction error $\widehat{\text{Err}}^{(\text{NCV})}$.

In view of (15), the estimator $\widehat{\text{MSE}}$ is targeting the MSE of $\widehat{\text{Err}}^{(\text{CV})}$ as an estimate of Err_{XY} , as we state next.

Theorem 3 (Estimand of nested CV). *For a nested CV with a sample of size n ,*

$$\mathbb{E} \left[\widehat{\text{MSE}} \right] = \text{MSE}_{K-1, n'},$$

where $n' = n(K - 1)/K$.

This result shows that $\widehat{\text{MSE}}$ obtained by nested CV is estimating the MSE of $(K - 1)$ -fold cross-validation on a sample of size $n(K - 1)/K$. Since nested CV uses an inner loop with samples of size $n(K - 1)/K$, we recommend re-scaling to obtain an estimate for a sample of size n by instead taking $(K - 1)/K \cdot \widehat{\text{MSE}}$ (although this re-scaled version is not guaranteed to be exactly unbiased for $\text{MSE}_{K, n}$). After adjusting the point estimate $\widehat{\text{Err}}^{(\text{NCV})}$ with a bias correction discussed next, we form our final confidence intervals as in (17).

Remark 5 (An interpretation of nested CV). *A simpler, perhaps more natural approach to this problem would be to borrow ideas from bootstrap calibration [e.g., 8, p. 263]. Specifically, one could use a nested CV scheme, and compute the standard normal confidence interval for error from the inner folds. Then one could check how well this interval covers $\bar{e}^{(\text{out})}$, and adjust the interval to achieve the desired coverage. The problem with this approach is that the left-out fold is finite, so that the interval is a prediction interval for $\bar{e}^{(\text{out})}$ rather than a confidence interval for the true underlying prediction error.*

NCV is similar in spirit to this: over repeated subsamples, we estimate the quantity (a) in (15), which leads to an empirical estimate of how much an interval around $\widehat{\text{Err}}_{\bar{X}\bar{Y}}$ must be widened in order to cover $\bar{e}^{(\text{out})}$. The latter is a random quantity, however, so this should be thought of as a calibrated prediction interval. In truth, we wish to cover not $\bar{e}^{(\text{out})}$, but its mean $\text{Err}_{\bar{X}\bar{Y}}$. We convert from a prediction interval for $\bar{e}^{(\text{out})}$ to a confidence interval for $\text{Err}_{\bar{X}\bar{Y}}$ by subtracting out the term (b).

Algorithm 1 Nested Cross-validation

Input: data (X, Y) , fitting algorithm \mathcal{A} , loss ℓ , number of folds, K , number of repetitions R

```
procedure NESTED_CROSSVAL( $X, Y$ ) ▷ primary algorithm
  es  $\leftarrow \square$  ▷ initialize empty vectors
  a_list  $\leftarrow \square$  ▷ (a) terms
  b_list  $\leftarrow \square$  ▷ (b) terms
  for  $r \in \{1, \dots, R\}$  do
    Randomly assign points to folds  $\mathcal{I}_1, \dots, \mathcal{I}_K$ 
    for  $k \in \{1, \dots, K\}$  do ▷ outer CV loop
       $e^{(\text{in})} \leftarrow \text{INNER\_CROSSVAL}(X, Y, \{\mathcal{I}_1, \dots, \mathcal{I}_K\} \setminus \mathcal{I}_k)$  ▷ inner CV loop
       $\hat{\theta} \leftarrow \mathcal{A}((X_i, Y_i)_{i \in \mathcal{I}_k})$ 
       $e^{(\text{out})} \leftarrow (\ell(\hat{f}(X_i, \hat{\theta}), Y_i))_{i \in \mathcal{I}_k}$ 
      a_list  $\leftarrow \text{append}(\mathbf{a\_list}, (\text{mean}(e^{(\text{in})}) - \text{mean}(e^{(\text{out})}))^2)$ 
      b_list  $\leftarrow \text{append}(\mathbf{b\_list}, \text{var}(e^{(\text{out})}))$ 
      es  $\leftarrow \text{append}(\mathbf{es}, e^{(\text{in})})$ 
   $\widehat{\text{MSE}} \leftarrow \text{mean}(\mathbf{a\_list}) - \text{mean}(\mathbf{b\_list})$  ▷ plug-in estimator based on (15)
   $\widehat{\text{Err}}^{(\text{NCV})} \leftarrow \text{mean}(\mathbf{es})$ 
  return:  $(\widehat{\text{Err}}^{(\text{NCV})}, \widehat{\text{MSE}})$  ▷ prediction error estimate and MSE estimate

procedure INNER_CROSSVAL( $X, Y, \{\mathcal{I}_1, \dots, \mathcal{I}_{K-1}\}$ ) ▷ inner cross-validation subroutine
   $e^{(\text{in})} \leftarrow \square$ 
  for  $k \in \{1, \dots, K-1\}$  do
     $\hat{\theta} \leftarrow \mathcal{A}((X_i, Y_i)_{i \in \mathcal{I}_1 \cup \dots \cup \mathcal{I}_{K-1} \setminus \mathcal{I}_k})$ 
     $e^{(\text{temp})} \leftarrow (\ell(\hat{f}(X_i, \hat{\theta}), Y_i))_{i \in \mathcal{I}_k}$ 
     $e^{(\text{in})} \leftarrow \text{append}(e^{(\text{in})}, e^{(\text{temp})})$ 
  return:  $e^{(\text{in})}$ 
```

Output: NESTED_CROSSVAL(X, Y)

4.4 Estimation of bias

The nested CV computations also yield a convenient estimate of the bias of the NCV point estimate of error, $\widehat{\text{Err}}^{(\text{NCV})}$. The NCV estimate of prediction error is unbiased for Err for the procedure with a reduced sample size of $n(K-2)/K$, but this will be typically be slightly biased upwards for Err with the full sample size. This discrepancy can be estimated by running both the usual K -fold CV and nested CV. An unbiased estimate for the difference in Err at a sample of size $n(K-2)/K$ (the sample size used in nested CV) to a sample of size $n(K-1)/K$ (the sample size used in standard CV) is

$$\widehat{\text{Err}}^{(\text{NCV})} - \widehat{\text{Err}}^{(\text{CV})}.$$

Now, since we expect that prediction error scales as $a + b/n$ in n for some unknown constants a and b (the parametric rate), an estimator for the difference in in Err at a sample of size n to Err at a sample of size $n(K-2)/K$ (the sample size of each model used in nested CV) is then:

$$\widehat{\text{bias}} := \left(1 + \left(\frac{K-2}{K}\right)\right) \left(\widehat{\text{Err}}^{(\text{NCV})} - \widehat{\text{Err}}^{(\text{CV})}\right). \quad (16)$$

The left term in the sum, “1”, accounts for the bias when going from size $n(K-2)/K$ to size $n(K-1)/K$ and the right term in the sum, accounts for the bias going from a sample size from $n(K-1)/K$ to size n ; this scaling due to the form of $a + b/n$. Combining this with the estimate of MSE in the previous section

leads to the following confidence intervals:

$$\left(\widehat{\text{Err}}^{(\text{NCV})} - \widehat{\text{bias}} - q_{1-\alpha/2} \cdot \sqrt{\frac{K-1}{K}} \cdot \sqrt{\widehat{\text{MSE}}}, \widehat{\text{Err}}^{(\text{NCV})} - \widehat{\text{bias}} + q_{1-\alpha/2} \cdot \sqrt{\frac{K-1}{K}} \cdot \sqrt{\widehat{\text{MSE}}} \right). \quad (17)$$

In practice, we also restrict $\sqrt{(K-1)/K} \cdot \sqrt{\widehat{\text{MSE}}}$ to fall between $\widehat{\text{SE}}$ (the estimated standard error if one had n independent points) and $\sqrt{K} \cdot \widehat{\text{SE}}$ (the estimated standard error if one had only n/K independent points). This is a minor implementation detail to help prevent implausible values of $\widehat{\text{MSE}}$ from arising.

5 Simulation experiments

We now explore the coverage of nested CV in a variety of settings. In each case, we will report the coverage of naïve CV (CV), nested CV (NCV), and data splitting with refitting (DS), where the nominal miscoverage rate is 10% (5% miscoverage in each tail). We also report on the width of the intervals, expressed relative to the width of the standard CV intervals. (We wish to produce intervals that are as narrow as possible while maintaining correct coverage.) We use 10-fold CV (the number of folds has little impact; see Appendix D.3) and NCV, with 200 random splits for the latter; see Appendix D.1 for the runtime of each experiment. For classification examples we use binary loss, and form confidence intervals for CV, NCV, and data splitting after taking the binomial variance-stabilizing transformation, described in detail in Appendix E. For regression examples, we use squared error loss.

For data splitting, we use 80% of the samples for training and 20% for estimating prediction error. Note that the data splitting without refitting intervals are the same as the data splitting with refitting intervals; the difference is that they are intended to cover different quantities. To make this comparable to CV and nested CV, we report on the coverage of Err and Err_{XY} here, which corresponds to data splitting with refitting. Data splitting without refitting (which seeks to cover the quantity in (9)) will typically have better coverage; we observed relatively accurate coverage in the classification examples and worse coverage in the regression examples, but do not explicitly report these results herein.

Scripts reproducing these experiments are available at https://github.com/stephenbates19/nestedcv_experiments.

5.1 Classification

5.1.1 Low-dimensional logistic regression

We consider the logistic regression data generating model 1 with $n = 100$ observations and $p = 20$ features, sampled as i.i.d. standard Gaussian variables. Due to the rotational symmetry of the features, the only parameter that affects behavior is the signal strength, and we explore models with Bayes error of either 33% or 23%. Here, we use (un-regularized) logistic regression as our fitting algorithm. We report the results in Table 1, finding that nested CV gives coverage much closer to the nominal level. Moreover, the point estimates have slightly less bias. We report the size of the NCV intervals relative to their CV counterparts per instance in Figure D.1.

5.1.2 High-dimensional sparse logistic regression

We return to the high-dimensional logistic regression model introduced in Section 1.1, generalizing slightly. We consider $n \in \{90, 200\}$ with $p = 1000$ features. The feature matrix has standard normal entries with an autoregressive covariance pattern such that adjacent columns have covariance ρ . In each case, we take $k = 4$ nonzero entries of the covariance matrix and use sparse logistic regression. We report on the results in Table 2 and give the width² in Figure 9. Again, NCV gives intervals with coverage much closer to the nominal level.

²The width in Figure 9 is reported relative to the version of cross-validation that holds out 2 folds at a time, since this is what is computed internally during NCV. In table Table 2 and elsewhere, we instead report widths relative to the usual K -fold CV.

Setting		Width		Point estimates				Miscoverage					
Bayes Error	Target	NCV	DS	Err	CV	NCV	DS	CV		NCV		DS	
								Hi	Lo	Hi	Lo	Hi	Lo
33.2%	Err _{XY}	1.23	2.23	39.1%	39.6%	39.0%	40.1%	10%	8%	3%	5%	7%	6%
"	Err	"	"	"	"	"	"	9%	8%	3%	4%	6%	5%
22.5%	Err _{XY}	1.47	2.25	28.7%	30.4%	28.1%	33.3%	11%	3%	4%	1%	16%	4%
"	Err	"	"	"	"	"	"	10%	2%	5%	0%	15%	3%

Table 1: Performance of cross-validation (CV), nested cross-validation (NCV), and data splitting with re-fitting (DS) in the low-dimensional logistic regression model from Section 5.1.1. Each row is a setting with a different signal strength, indexed by the Bayes error: the error of the true model. The nominal total error rate is 10%, i.e., 5% above and below. A “Hi” miscoverage is one where the confidence interval is too large and the point estimate falls below the interval; conversely for a “Lo” miscoverage. The standard error in each coverage estimate reported is about 0.5%. The “Target” column indicates the target of coverage—the intervals are always generated identically, but we report the coverage of both Err and Err_{XY}.




Figure 9: Size of the nested CV intervals relative to the size of the naïve CV intervals in the high-dimensional sparse logistic regression experiment from Section 5.1.2.

Setting			Width		Point estimates				Miscoverage					
n	ρ	Target	NCV	DS	Bayes error	Err	CV	NCV	CV		NCV		DS	
									Hi	Lo	Hi	Lo	Hi	Lo
90	0	Err _{XY}	1.53	2.24	22%	41.3%	41.8%	41.1%	16%	12%	6%	7%	9%	7%
"	"	Err	"	"	"	"	"	"	17%	13%	6%	8%	11%	9%
200	0	Err _{XY}	1.66	2.26	22%	25.6%	26.7%	25.6%	14%	7%	3%	5%	9%	4%
"	"	Err	"	"	"	"	"	"	15%	7%	4%	6%	8%	5%
90	0.5	Err _{XY}	1.80	2.25	13%	25.6%	27.5%	28.6%	20%	10%	5%	8%	15%	4%
"	"	Err	"	"	"	"	"	"	20%	11%	7%	9%	14%	3%

Table 2: Performance of cross-validation (CV), nested cross-validation (NCV), and data splitting (DS) in the high-d logistic regression model from Section 5.1.2. The nominal (target) error rate is 10%, i.e., 5% above and below. Other details as in Table 1.

5.2 Regression

5.2.1 Low-dimensional linear model

We next consider an OLS example. We take $X \in \mathbb{R}^{n \times p}$ with $p = 20$ comprised of i.i.d. $\mathcal{N}(0, 1)$. Further, we generate a response from the standard linear model:

$$Y = X\theta + \epsilon$$

where ϵ is likewise i.i.d. $\mathcal{N}(0, 1)$. We use OLS to estimate θ . Note that by Lemma 1, the choice of θ does not affect the coverage rate of CV. The same argument shows that the choice of θ will not affect the coverage rate of nested CV, so we can take θ to be 0 without loss of generality. Similarly, both CV and NCV are unchanged when X is transformed by an full-rank linear operator, so the results in this section would remain unchanged for Gaussian features with any full-rank correlation structure. We report the coverage of nested CV in Figure 10. We find that this scheme works well and has good coverage for any n , overcovering somewhat for very small n . By contrast, naïve CV has poor coverage until n is 400. In Figure D.2 we report on the width of the NCV intervals relative to their CV counterparts—the usual ratio is not that large for samples sizes of $n = 100$ or greater.

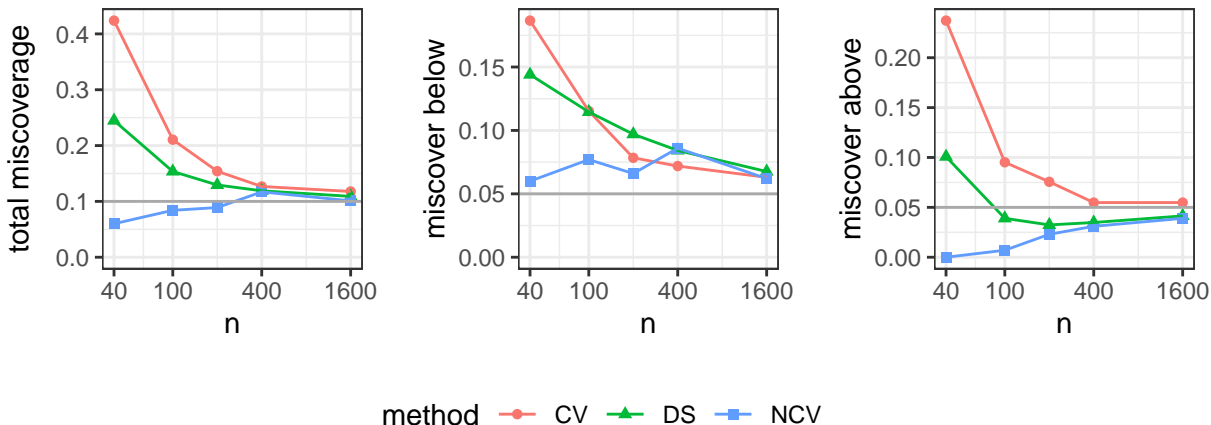


Figure 10: Coverage of CV, data splitting, and nested CV in the OLS case.

5.2.2 A high-dimensional sparse linear model

We continue as in the previous experiment, but with $n \in \{50, 100\}$ and $p = 500$. We choose θ to have 4 nonzero entries of equal strength such that

$$\frac{\text{var}(X\theta)}{\text{var}(\epsilon)} = 4.$$

Since $p > n$, we take the lasso estimator with a fixed penalty parameter. The parameter is chosen by minimizing the cross-validation estimate of prediction error on a single independent run, and then this value is fixed for the experimental replicates. We report on the results in Table 3, again finding the NCV has better coverage than CV, although both struggle when $n = 50$. The ratio the nested CV interval width to the naïve CV interval width is relatively stable across observations, see Figure D.3.

6 Real data examples

Lastly, we evaluate the nested CV procedure on real data sets from the UCI repository [40]. In each case, we repeatedly subsample a small number of observations, perform nested CV on the subsample, and then use

Setting		Width		Point estimates					Miscoverage					
n	Target	NCV	DS	Bayes Error	Err	CV	NCV	DS	CV		NCV		DS	
									Hi	Lo	Hi	Lo	Hi	Lo
50	Err _{XY}	1.27	2.33	1	2.37	2.61	2.37	3.08	14%	8%	9%	3%	8%	14%
"	Err	"	"	"	"	"	"	"	9%	7%	5%	1%	6%	16%
100	Err _{XY}	1.89	2.20	1	1.54	1.61	1.55	1.63	13%	14%	3%	2%	4%	11%
"	Err	"	"	"	"	"	"	"	14%	16%	4%	4%	5%	12%

Table 3: Performance of cross-validation (CV), nested cross-validation (NCV), and data splitting (DS) in the high-d linear regression model. The nominal (target) error rate is 10%, i.e., 5% above and below. Other details as in Table 1.

the many remaining observations to determine the accuracy of the fitted model. We consider the following data sets:

- **Communities and crimes (CC).** This data set is comprised of measurements of 1994 communities in the US. We predict the crime rate of each community, a real number normalized to be between 0 and 1, based on 99 demographic features of the community.
- **Crop mapping (crp).** This data set is comprised of optical radar measurements of cropland in Manitoba in 2012. We filter the data set to contain two classes, corn and oats, and then do binary classification based on 174 features. Here, we add a small amount of label noise so that the best possible classifier has a misclassification rate of about 5%.

We again use sparse linear or logistic regression as our fitting algorithm. The results are reported in Table 4. We find that nested CV generally has coverage that is much closer to the nominal rate than naïve CV. Data splitting has poor coverage in this case due to the small sample size, but is significantly better with $n = 100$ samples than with $n = 50$ samples.

Setting			Width		Point estimates				Miscoverage					
data	n	Target	NCV	DS	Err	CV	NCV	DS	CV		NCV		DS	
									Hi	Lo	Hi	Lo	Hi	Lo
CC	50	Err _{XY}	2.82	1.77	0.029	0.031	0.029	.034	4%	20%	1%	13%	1%	33%
"	"	Err	"	"	"	"	"	"	2%	22%	0%	12%	1%	37%
CC	100	Err _{XY}	1.46	2.05	0.023	0.024	0.023	.025	4%	13%	2%	7%	1%	24%
"	"	Err	"	"	"	"	"	"	2%	12%	1%	4%	1%	24%
crp	50	Err _{XY}	1.21	1.89	10.6%	10.7%	10.6%	10.8%	6%	8%	2%	6%	4%	31%
"	"	Err	"	"	"	"	"	"	7%	12%	3%	8%	2%	31%
crp	100	Err _{XY}	1.52	2.00	9.5%	9.7%	9.5%	9.4%	6%	6%	4%	5%	4%	15%
"	"	Err	"	"	"	"	"	"	8%	9%	5%	7%	4%	15%

Table 4: Performance of cross-validation (CV), nested cross-validation (NCV), and data splitting (DS) with the real data sets. The nominal (target) error rate is 10%, i.e., 5% above and below. Other details as in Table 1.

7 Discussion

We have made two main contributions. First, we discussed point estimates of prediction error via subsampling techniques. Our primary result is that common estimates of prediction error—cross-validation, bootstrap, data splitting, and covariance penalties—cannot be viewed as estimates of the prediction error of the final

model fit on the whole data. Rather, the estimate of prediction error is an estimate of the average prediction error of the final model across other hypothetical data sets from the same distribution. Note that the formal results here were all for the special case of the linear model using unregularized OLS for model-fitting. A major remaining question is to understand whether or not this holds with regularization. We expect that with a regularized model-fitting routine, cross-validation is closer to estimating Err_{XY} , and we look forward to future work explaining the behavior of cross-validation (and other estimates of prediction error) in that case.

Secondly, we discuss inference for cross-validation, deriving an estimator for the MSE of the CV point estimate, nested CV. The nested CV scheme has consistently superior coverage compared to naïve cross-validation confidence intervals, which makes it an appealing choice for providing confidence intervals for prediction error. Perhaps surprisingly, we also show that simple data splitting does not give valid confidence intervals (even asymptotically) when one refits the model. Thus, if one wants valid confidence intervals for prediction error, we can only recommend either data splitting *without* refitting the model (which is viable when one has ample data), or nested CV.

Nonetheless, we wish to be clear that nested CV is more computationally intensive than standard CV—we use about 1000 times more model fits per example because of the repeated splitting. For example, in the logistic regression example from Section 1.1, nested CV takes about one 10 seconds on a personal computer. Fortunately nested CV is embarrassingly parallel, so that this computation can be sped up considerably.

A fundamental open question is to understand under what conditions the standard CV intervals will be badly behaved, making the nested CV computations necessary. Roughly speaking, we expect the standard CV intervals to perform better when n/p is larger and when more regularization is used. In our experiments, we saw that even in the mundane linear model with $n/p = 10$, the miscoverage rate of standard CV was about 50% larger than the nominal rate. As n increases, however, the violation decreases. Moreover, the asymptotic results in [21] and [26] suggest that the violations in coverage disappear in the p fixed, $n \rightarrow \infty$ limiting regime. We look forward to more work theoretically explaining the miscoverage phenomenon in small samples or high-dimensions.

To conclude, we point out several additional future directions. First, one could adapt nested CV to cases with dependent data, as is done for standard CV [41]. Second, we note that the general leave-out style strategy of cross-validation can also be used to “fill-in” data for downstream use. Examples include *pre-validation* [42, 43] and *cross-fitting* [e.g., 44]). Further, confidence intervals for prediction accuracy are used to evaluate variable importance [45, 46]. We suspect that our nested cross-validation proposal could be adapted to improve the accuracy of these and related approaches. Lastly, cross-validation is often used to compare regression procedures (such as when selecting the value of a tuning parameter); see Section 1.2. We anticipate that nested CV can be extended to give valid confidence intervals for the difference in prediction error between two models.

An R package implementing nested CV is available at <https://github.com/stephenbates19/nestedcv>.

Acknowledgements

The authors would like to acknowledge Frank Harrell for a seminar and personal correspondence, alerting them to the miscoverage of cross-validation in the high-dimensional logistic regression model. We would like to thank Michael Celentano, Ryan Tibshirani, and Larry Wasserman for helpful discussions. We would especially like to thank Bradley Efron for his comments on this work and for his fundamental contributions to this general area over the last 40 years. S. B. was partially supported by a Ric Weiland Graduate Fellowship. T.H. was partially supported by grants DMS-2013736 And IIS 1837931 from the National Science Foundation, and grant 5R01 EB 001988-21 from the National Institutes of Health. R.T. was supported by the National Institutes of Health (5R01 EB001988-16) and the National Science Foundation (19 DMS1208164).

References

- [1] D. Allen, “The relationship between variable selection and data augmentation and a method of prediction,” *Technometrics*, vol. 16, pp. 125–7, 1974.
- [2] S. Geisser, “The predictive sample reuse method with applications,” *Journal of the American Statistical Association*, vol. 70, no. 350, pp. 320–328, 1975.
- [3] M. Stone, “Cross-validated choice and assessment of statistical predictions,” *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 36, no. 2, pp. 111–147, 1977.
- [4] A. Blum, A. T. Kalai, and J. Langford, “Beating the hold-out: bounds for k-fold and progressive cross-validation,” in *Proceedings of the Twelfth Annual Conference on Learning Theory*, 1999.
- [5] B. Efron, “Estimating the error rate of a prediction rule: Improvement on cross-validation,” *Journal of the American Statistical Association*, vol. 78, no. 382, pp. 316–331, 1983.
- [6] B. Efron, “How biased is the apparent error rate of a prediction rule?,” *Journal of the American Statistical Association*, vol. 81, no. 394, pp. 461–470, 1986.
- [7] B. Efron and R. Tibshirani, “Improvements on cross-validation: The .632+ bootstrap method,” *Journal of the American Statistical Association*, vol. 92, no. 438, pp. 548–560, 1997.
- [8] B. Efron and R. J. Tibshirani, *An Introduction to the Bootstrap*. Chapman & Hall/CRC, 1993.
- [9] C. L. Mallows, “Some comments on C_p ,” *Technometrics*, vol. 15, no. 4, pp. 661–675, 1973.
- [10] H. Akaike, “A new look at the statistical model identification,” *IEEE Transactions on Automatic Control*, vol. 19, no. 6, pp. 716–723, 1974.
- [11] G. Schwarz, “Estimating the Dimension of a Model,” *The Annals of Statistics*, vol. 6, no. 2, pp. 461 – 464, 1978.
- [12] C. M. Stein, “Estimation of the Mean of a Multivariate Normal Distribution,” *The Annals of Statistics*, vol. 9, no. 6, pp. 1135 – 1151, 1981.
- [13] B. Efron, “The estimation of prediction error,” *Journal of the American Statistical Association*, vol. 99, no. 467, pp. 619–632, 2004.
- [14] P. Zhang, “Assessing prediction error in non-parametric regression,” *Scandinavian Journal of Statistics*, vol. 22, no. 1, pp. 83–94, 1995.
- [15] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning, 2nd Ed.* Springer, New York, 2009.
- [16] W. A. Yousef, “A leisurely look at versions and variants of the cross validation estimator,” *arXiv preprint*, 2020. arXiv:1907.13413.
- [17] S. Rosset and R. J. Tibshirani, “From fixed-x to random-x regression: Bias-variance decompositions, covariance penalties, and prediction error estimation,” *Journal of the American Statistical Association*, vol. 115, no. 529, pp. 138–151, 2020.
- [18] S. Wager, “Cross-validation, risk estimation, and model selection: Comment on a paper by Rosset and Tibshirani,” *Journal of the American Statistical Association*, vol. 115, no. 529, pp. 157–160, 2020.
- [19] Y. Bengio and Y. Grandvalet, “No unbiased estimator of the variance of k-fold cross-validation,” *Journal of Machine Learning Research*, vol. 5, p. 1089–1105, Dec. 2004.
- [20] C. Nadeau and Y. Bengio, “Inference for the generalization error,” *Machine Learning*, vol. 52, no. 3, pp. 239–281, 2003.

- [21] M. Austern and W. Zhou, “Asymptotics of cross-validation,” *arXiv preprint*, 2020. arXiv:2001.11111.
- [22] M. Markatou, H. Tian, S. Biswas, and G. Hripcsak, “Analysis of variance of cross-validation estimators of the generalization error,” *Journal of Machine Learning Research*, vol. 6, p. 1127–1168, Dec. 2005.
- [23] W. A. Yousef, “Estimating the standard error of cross-validation-based estimators of classification rules performance,” *arXiv preprint*, 2019. arXiv:1908.00325.
- [24] S. Dudoit and M. J. van der Laan, “Asymptotics of cross-validated risk estimation in estimator selection and performance assessment,” *Statistical Methodology*, vol. 2, no. 2, pp. 131 – 154, 2005.
- [25] E. LeDell, M. Petersen, and M. van der Laan, “Computationally efficient confidence intervals for cross-validated area under the ROC curve estimates,” *Electronic Journal of Statistics*, vol. 9, no. 1, pp. 1583–1607, 2015.
- [26] P. Bayle, A. Bayle, L. Mackey, and L. Janson, “Cross-validation confidence intervals for test error,” in *Conference on Neural Information Processing Systems*, December 2020.
- [27] R. Kumar, D. Lokshitanov, S. Vassilvitskii, and A. Vattani, “Near-optimal bounds for cross-validation via loss stability,” in *Proceedings of the 30th International Conference on Machine Learning*, 2013.
- [28] S. Kale, R. Kumar, and S. Vassilvitskii, “Cross-validation and mean-square stability,” in *Proceedings of Innovations in Computer Science*, 2011.
- [29] A. Celisse and B. Guedj, “Stability revisited: new generalisation bounds for the leave-one-out,” *arXiv preprint*, 2016. arXiv:1608.06412.
- [30] J. Shao, “Linear model selection by cross-validation,” *Journal of the American Statistical Association*, vol. 88, no. 422, pp. 486–494, 1993.
- [31] T. G. Dietterich, “Approximate statistical tests for comparing supervised classification learning algorithms,” *Neural Computation*, vol. 10, no. 7, pp. 1895–1923, 1998.
- [32] Q.-S. Xu and Y.-Z. Liang, “Monte carlo cross validation,” *Chemometrics and Intelligent Laboratory Systems*, vol. 56, no. 1, pp. 1 – 11, 2001.
- [33] Y. Yang, “Consistency of cross validation for comparing regression procedures,” *The Annals of Statistics*, vol. 35, no. 6, pp. 2450 – 2473, 2007.
- [34] S. Arlot and A. Celisse, “A survey of cross-validation procedures for model selection,” *Statistics Surveys*, vol. 4, pp. 40–79, 2010.
- [35] R. J. Tibshirani and R. Tibshirani, “A bias correction for the minimum error rate in cross-validation,” *Annals of Applied Statistics*, vol. 3, no. 2, pp. 822–829, 2009.
- [36] J. Lei, “Cross-validation with confidence,” *Journal of the American Statistical Association*, 2019. doi:10.1080/01621459.2019.1672556.
- [37] T. Hastie, A. Montanari, S. Rosset, and R. J. Tibshirani, “Surprises in high-dimensional ridgeless least squares interpolation,” *arXiv preprint*, 2019. arXiv:1903.08560.
- [38] J. Ye, “On measuring and correcting the effects of data mining and model selection,” *Journal of the American Statistical Association*, vol. 93, no. 441, pp. 120–131, 1998.
- [39] B. Efron and G. Gong, “A leisurely look at the bootstrap, the jackknife, and cross-validation,” *The American Statistician*, vol. 37, no. 1, pp. 36–48, 1983.
- [40] D. Dua and C. Graff, “UCI machine learning repository,” 2017.
- [41] A. Rabinowicz and S. Rosset, “Cross-validation for correlated data,” *Journal of the American Statistical Association*, vol. 0, no. 0, pp. 1–14, 2020.

- [42] R. Tibshirani and B. Efron, “Pre-validation and inference in microarrays,” *Statistical Applications in Genetics and Molecular Biology*, vol. 1, pp. 1–20, August 2002.
- [43] H. Höfling and R. Tibshirani, “A study of pre-validation,” *The Annals of Applied Statistics*, vol. 2, no. 2, pp. 643 – 664, 2008.
- [44] W. K. Newey and J. R. Robins, “Cross-fitting and fast remainder rates for semiparametric estimation,” *arXiv preprint*, 2018. arXiv:1801.09138.
- [45] B. D. Williamson, P. B. Gilbert, M. Carone, and N. Simon, “Nonparametric variable importance assessment using machine learning techniques,” *Biometrics*, vol. 77, no. 1, pp. 9–22, 2021.
- [46] L. Zhang and L. Janson, “Floodgate: Inference for model-free variable importance,” *arXiv preprint*, 2020. arXiv:2007.01283.
- [47] L. Breiman, “Out-of-bag estimation,” *unpublished note*, 1996. <https://www.stat.berkeley.edu/breiman/OOBestimation.pdf>.
- [48] D. von Rosen, “Moments for the inverted wishart distribution,” *Scandinavian Journal of Statistics*, vol. 15, no. 2, pp. 97–109, 1988.
- [49] T. Groves and T. Rothenberg, “A note on the expected value of an inverse matrix,” *Biometrika*, vol. 56, pp. 690–691, 12 1969.
- [50] G. Yu, “Variance stabilizing transformations of poisson, binomial and negative binomial distributions,” *Statistics and Probability Letters*, vol. 79, no. 14, pp. 1621–1629, 2009.

A Results for bootstrap estimates of prediction error

In this section, we give results parallel to those in Section 3 for bootstrap estimates of prediction error. We consider two bootstrap estimators: the .632 estimator and the out-of-bag (OOB) error estimator [5, 7, 47]. In the interest of brevity, we reference the definitions of these two estimators from [7]: we will use $\widehat{\text{Err}}^{(\text{OOB})}$ to denote the out-of-bag estimator defined in equation (17) from [7] and $\widehat{\text{Err}}^{(.632)}$ to denote the .632 estimator defined in equation (24) from [7]. In both cases, we will assume that the underlying model fitting routine is OLS.

Lemma 5. *When using OLS as the fitting algorithm, $\widehat{\text{Err}}^{(.632)}$ and $\widehat{\text{Err}}^{(\text{OOB})}$ are linearly invariant.*

Thus, by the results in Section 3, we conclude that bootstrap should also be viewed as an estimator of Err or Err_X , rather than of Err_{XY} .

Remark 6. *We expect these properties also hold for most other resampling-based estimators of prediction error, in addition to those explicitly considered herein. For example, nested CV is linearly invariant.*

B Proofs

Proof of Lemma 1. Without loss of generality, consider point 1. We will show that the residual from CV on point 1 is the same when using either the data $(x_1, y_1), \dots, (x_n, y_n)$ or the data $(x_1, y'_1), \dots, (x_n, y'_n)$, where $y'_i = y_i + x_i^\top \kappa$.

Let $\mathcal{I}_1 \subset \{1, \dots, n\}$ be the indices in the same fold as observation 1. Let $\hat{\theta}_{(-1)}$ be the OLS estimate of θ based on only the points $(x_i, y_i)_{i \notin \mathcal{I}_1}$ —all the points not in the same fold as point 1—and let $\hat{\theta}'_{(-1)}$ be the

OLS estimate of θ based on only the points $(x_i, y_i)_{i \notin \mathcal{I}_1}$. Then

$$\begin{aligned}
\hat{\theta}_{(-1)} &= \operatorname{argmin}_{\theta} \sum_{i \notin \mathcal{I}_1} (y_i - x_i^\top \theta)^2 \\
&= \operatorname{argmin}_{\theta} \sum_{i \notin \mathcal{I}_1} (y_i + \kappa^\top x_i - x_i^\top (\theta + \kappa))^2 \\
&= -\kappa + \operatorname{argmin}_{\theta} \sum_{i \notin \mathcal{I}_1} (y'_i - x_i^\top \theta)^2 \\
&= \hat{\theta}'_{(-1)} - \kappa.
\end{aligned}$$

As a result,

$$y_1 - x_1^\top \hat{\theta}_{(-1)} = y'_1 - x_1^\top \hat{\theta}'_{(-1)}.$$

The CV estimate of prediction error is the mean of the squared residuals, so since by the preceding display the residuals are the same for either the original or shifted data, the CV estimate of prediction error is the same in each case. \square

Proof of Theorem 1. The true predictive error (Err_{XY}) is a function only of $\hat{\theta}$, the OLS estimate of θ based on the full sample $(x_1, y_1), \dots, (x_n, y_n)$. On the other hand, any linearly invariant $\widehat{\operatorname{Err}}$ is a function only of the residuals $Y - X\hat{\theta} = (I - X(X^\top X)^{-1}X^\top)Y$, by the invariance property. Since

$$\hat{\theta} \perp\!\!\!\perp (Y - X\hat{\theta}) \mid X,$$

from classical linear model results, the proof is complete. \square

Proof of Corollary 1.

$$\begin{aligned}
\mathbb{E}[(\widehat{\operatorname{Err}}^{(\text{CV})} - \operatorname{Err}_{XY})^2] &= \mathbb{E}[(\widehat{\operatorname{Err}}^{(\text{CV})} - \operatorname{Err}_X + \operatorname{Err}_X - \operatorname{Err}_{XY})^2] \\
&= \mathbb{E}[(\widehat{\operatorname{Err}}^{(\text{CV})} - \operatorname{Err}_X)^2] + \mathbb{E}[(\operatorname{Err}_X - \operatorname{Err}_{XY})^2] + \mathbb{E}[(\widehat{\operatorname{Err}}^{(\text{CV})} - \operatorname{Err}_X)(\operatorname{Err}_X - \operatorname{Err}_{XY})] \\
&= \mathbb{E}[(\widehat{\operatorname{Err}}^{(\text{CV})} - \operatorname{Err}_X)^2] + \mathbb{E}[(\operatorname{Err}_X - \operatorname{Err}_{XY})^2] \\
&\quad + \mathbb{E}\left[\mathbb{E}[(\widehat{\operatorname{Err}}^{(\text{CV})} - \operatorname{Err}_X)(\operatorname{Err}_X - \operatorname{Err}_{XY}) \mid X]\right] \\
&= \mathbb{E}[(\widehat{\operatorname{Err}}^{(\text{CV})} - \operatorname{Err}_X)^2] + \mathbb{E}[(\operatorname{Err}_X - \operatorname{Err}_{XY})^2]
\end{aligned}$$

Where the last equality follows from Theorem 1. The result follows by noting that the second term in the final line is $\mathbb{E}[\operatorname{var}(\operatorname{Err}_{XY} \mid X)]$. \square

Proof of Theorem 2. We first give an expression for Err_{XY} .

$$\begin{aligned}
\operatorname{Err}_{XY} &= \mathbb{E}_{X_{n+1}, Y_{n+1}} \left(X_{n+1}^\top \hat{\beta} - Y_{n+1} \right)^2 \\
&= \sigma^2 + \mathbb{E}_{X_{n+1}} \left(X_{n+1}^\top (\hat{\beta} - \beta) \right)^2 \\
&= \sigma^2 + \mathbb{E}_{X_{n+1}} \left[(\hat{\beta} - \beta)^\top X_{n+1} X_{n+1}^\top (\hat{\beta} - \beta) \right] \\
&= \sigma^2 + \left\| \Sigma^{1/2} (\hat{\beta} - \beta) \right\|^2 \\
&= \sigma^2 + \frac{\sigma^2}{n} \left\| \Sigma^{1/2} (X^\top X/n)^{-1/2} Z \right\|^2 \\
&= \sigma^2 + \frac{\sigma^2}{n} \left\| \Sigma^{1/2} \hat{\Sigma}^{-1/2} Z \right\|^2,
\end{aligned}$$

where $Z \sim \mathcal{N}(0, I)$ is a function only of the noise ϵ . Next, we decompose the variance with the conditional variance formula

$$\begin{aligned} \text{var}(\text{Err}_{XY}) &= \mathbb{E}[\text{var}(\text{Err}_{XY} | X)] + \text{var}(\mathbb{E}[\text{Err}_{XY} | X]) \\ &= \underbrace{\mathbb{E}[\text{var}(\text{Err}_{XY} | X)]}_{\text{var due to } Y | X} + \underbrace{\text{var}(\text{Err}_X)}_{\text{var due to } X}, \end{aligned}$$

which was previously stated in (6) of the main text. We next derive asymptotic rates for the two components of this sum.

Proof of first claim. We begin with the first term, which corresponds to the variance cause by the randomness in $Y | X$.

$$\begin{aligned} \mathbb{E}[\text{var}(\text{Err}_{XY} | X)] &= \mathbb{E} \text{var} \left(\frac{\sigma^2}{n} \left\| \Sigma^{1/2} \hat{\Sigma}^{-1/2} Z \right\|^2 \mid X \right) \\ &= \frac{\sigma^4}{n^2} \mathbb{E} \text{var}(\lambda_1^2 Z_1^2 + \dots + \lambda_p^2 Z_p^2) \\ &= \frac{2\sigma^4}{n^2} \mathbb{E}[\lambda_1^4 + \dots + \lambda_p^4] \end{aligned}$$

where $\lambda_1^2, \dots, \lambda_p^2$ are the eigenvalues of $\Sigma \hat{\Sigma}^{-1}$, a function of X . Thus, the proof is complete once we show that the expectation term in the final line is $\Theta(n)$, which we turn to next.

Notice that

$$\begin{aligned} \mathbb{E}[\lambda_1^4 + \dots + \lambda_p^4] &= \mathbb{E} \left[\text{tr} \left((\Sigma \hat{\Sigma}^{-1})(\Sigma \hat{\Sigma}^{-1}) \right) \right] \\ &= (c_1 + c_2)n^2 p + c_2 n^2 p^2 \quad (\text{by Corollary 3.1 of [48]}), \end{aligned}$$

where

$$c_1 := \frac{(n-p-2)}{(n-p)(n-p-1)(n-p-3)} \quad c_2 := \frac{1}{(n-p)(n-p-1)(n-p-3)}$$

We conclude

$$\mathbb{E}[\lambda_1^4 + \dots + \lambda_p^4] = \Theta(n).$$

This completes the proof of the first claim.

Proof of second claim. Now, we turn our attention to the second term, which corresponds to the variance caused by the randomness in X .

$$\begin{aligned} \text{var}(\text{Err}_X) &= \text{var} \left(\frac{\sigma^2}{n} \text{tr}(\Sigma \hat{\Sigma}^{-1}) \right) \\ &= \frac{\sigma^4}{n^2} \text{var} \left(\text{tr}(\Sigma \hat{\Sigma}^{-1}) \right) \\ &= \frac{\sigma^4}{n^2} \cdot \Theta(1) \quad \text{as } n, p \rightarrow \infty \end{aligned}$$

where the final equality comes from the second moments of the inverse Wishart distribution [48]. To elaborate on the last equality,

$$\begin{aligned} \text{var} \left(\text{tr}(\Sigma \hat{\Sigma}^{-1}) \right) &= p \cdot \text{var}((\Sigma \Sigma^{-1})_{11}) + p(p-1) \text{cov}((\Sigma \Sigma^{-1})_{11}, (\Sigma \Sigma^{-1})_{22}) \\ &= p \cdot \frac{2n^2}{(n-p-1)^2(n-p-3)} + p(p-1) \cdot \frac{2n^2}{(n-p)(n-p-1)^2(n-p-3)} \\ &= \Theta(1). \end{aligned}$$

□

Proof of Corollary 2. For the first claim,

$$\begin{aligned}
\text{cor}(\text{Err}_X, \text{Err}_{XY}) &= \frac{\text{cov}(\text{Err}_X, \text{Err}_{XY})}{\sqrt{\text{var}(\text{Err}_X)\text{var}(\text{Err}_{XY})}} \\
&= \frac{\text{cov}(\text{Err}_X, \mathbb{E}[\text{Err}_{XY} | X])}{\sqrt{\text{var}(\text{Err}_X)\text{var}(\text{Err}_{XY})}} \\
&= \frac{\text{cov}(\text{Err}_X, \text{Err}_X)}{\sqrt{\text{var}(\text{Err}_X)\text{var}(\text{Err}_{XY})}} \\
&= \sqrt{\frac{\text{var}(\text{Err}_X)}{\text{var}(\text{Err}_{XY})}} \rightarrow 0 \quad \text{as } n \rightarrow \infty.
\end{aligned}$$

The second equality above comes from the conditional covariance formula, conditioning on X .

For the second claim, note that $\widehat{\text{Err}} = g(\text{Err}_X, U)$ for an independent $U \sim \text{unif}[0, 1]$ for some function g (by Theorem 1). That is, $\widehat{\text{Err}}$ is a random function of Err_X . As a result,

$$\begin{aligned}
\text{cor}(\text{Err}_{XY}, \widehat{\text{Err}}) &= \text{cor}(\text{Err}_{XY}, g(\text{Err}_X, U)) \\
&= \frac{\text{cov}(\text{Err}_{XY}, g(\text{Err}_X, U))}{\sqrt{\text{var}(\text{Err}_{XY})\text{var}(g(\text{Err}_X, U))}} \\
&= \frac{\text{cov}(\text{Err}_X, \mathbb{E}[g(\text{Err}_X, U) | X])}{\sqrt{\text{var}(\text{Err}_{XY})\text{var}(g(\text{Err}_X, U))}} \\
&= \sqrt{\frac{\text{var}(\text{Err}_X)}{\text{var}(\text{Err}_{XY})}} \cdot \frac{\text{cov}(\text{Err}_X, \mathbb{E}[g(\text{Err}_X, U) | X])}{\sqrt{\text{var}(\text{Err}_X)\text{var}(g(\text{Err}_X, U))}} \\
&\leq \sqrt{\frac{\text{var}(\text{Err}_X)}{\text{var}(\text{Err}_{XY})}} \cdot \text{cor}(\text{Err}_X, \mathbb{E}[g(\text{Err}_X, U) | X]) \\
&\leq \sqrt{\frac{\text{var}(\text{Err}_X)}{\text{var}(\text{Err}_{XY})}} \rightarrow 0 \quad \text{as } n \rightarrow \infty.
\end{aligned}$$

The third equality above comes from the conditional covariance formula, conditioning on X . □

Proof of Corollary 3. For the first claim,

$$\begin{aligned}
\mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_{XY} \right)^2 \right] &= \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X + \text{Err}_X - \text{Err}_{XY} \right)^2 \right] \\
&= \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X \right)^2 + (\text{Err}_X - \text{Err}_{XY})^2 - 2 \cdot \left(\widehat{\text{Err}} - \text{Err}_X \right) \cdot (\text{Err}_X - \text{Err}_{XY}) \right] \\
&= \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X \right)^2 + (\text{Err}_X - \text{Err}_{XY})^2 \right] \quad (\text{Theorem 1}) \\
&= \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X \right)^2 \right] + \text{var}(\text{Err}_X) + \text{var}(\text{Err}_{XY}) - 2 \cdot \text{cov}(\text{Err}_X, \text{Err}_{XY}) \\
&\geq \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X \right)^2 \right] + \text{var}(\text{Err}_{XY}) - 2\sqrt{\text{var}(\text{Err}_{XY}) \cdot \text{var}(\text{Err}_X)} \\
&= \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X \right)^2 \right] + \frac{1}{n} \cdot \Omega(1).
\end{aligned}$$

For the third claim,

$$\begin{aligned}
\mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err} \right)^2 \right] - \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X \right)^2 \right] &= \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X + \text{Err}_X - \text{Err} \right)^2 \right] - \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X \right)^2 \right] \\
&= \text{var}(\text{Err}_X) + 2 \cdot \text{cov}(\widehat{\text{Err}} - \text{Err}_X, \text{Err}_X) \\
&= \frac{1}{n^2} \cdot O(1) + 2 \cdot \text{cov}(\widehat{\text{Err}}, \text{Err}_X)
\end{aligned}$$

so

$$\left| \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err} \right)^2 \right] - \mathbb{E} \left[\left(\widehat{\text{Err}} - \text{Err}_X \right)^2 \right] \right| \leq \frac{1}{n^2} \cdot O(1) + \frac{2}{n} \cdot \sqrt{\text{var}(\widehat{\text{Err}})} \cdot O(1).$$

The second claim follows by combining these two results. \square

Proof of Proposition 1.

$$\begin{aligned}
\mathbb{E} \left[\left(\widehat{\text{SE}}^{(\text{split})} \right)^2 \right] &= \mathbb{E} \left[\text{var}(\widehat{\text{Err}}^{(\text{split})} | \tilde{X}, \tilde{Y}) \right] && \text{def. of } \widehat{\text{Err}}^{(\text{split})} \\
&= \text{var}(\widehat{\text{Err}}^{(\text{split})}) - \text{var}(\mathbb{E}[\widehat{\text{Err}}^{(\text{split})} | \tilde{X}, \tilde{Y}]) \\
&= \mathbb{E} \left[\left(\widehat{\text{Err}}^{(\text{split})} - \mathbb{E}[\text{Err}_{\tilde{X}, \tilde{Y}}] \right)^2 \right] - \text{var}(\text{Err}_{\tilde{X}, \tilde{Y}}) \\
&= \mathbb{E} \left[\left(\widehat{\text{Err}}^{(\text{split})} - \text{Err} \right)^2 \right] - \text{var}(\text{Err}_{\tilde{X}, \tilde{Y}}) - (\text{Err} - \mathbb{E}[\text{Err}_{\tilde{X}, \tilde{Y}}])^2.
\end{aligned}$$

\square

Proof of Proposition 2. Let j be drawn uniformly on $\{1, \dots, n\}$. For $i = 1, \dots, n$, let Y'_i be an independent draw from the distribution of $Y_i | X_i$. Then

$$\begin{aligned}
\text{Err}_X &= \text{Err}_{\text{in}}(X) + \mathbb{E} \left[(Y_{n+1} - \hat{f}(X_{n+1}, \hat{\theta}))^2 - (Y'_j - \hat{f}(X_j, \hat{\theta}))^2 \right] \\
&= \text{Err}_{\text{in}}(X) + \mathbb{E} \left[(\hat{\theta} - \theta)^\top \Sigma (\hat{\theta} - \theta) \right] - \mathbb{E} \left[(\hat{\theta} - \theta)^\top \hat{\Sigma} (\hat{\theta} - \theta) \right] \\
&= \text{Err}_{\text{in}}(X) + \mathbb{E} \left[\mathbb{E} \left[(\hat{\theta} - \theta)^\top \Sigma (\hat{\theta} - \theta) | X \right] \right] - \frac{p\sigma^2}{n} \\
&= \text{Err}_{\text{in}}(X) + \frac{\sigma^2}{n} \text{tr}(\hat{\Sigma}^{-1} \Sigma) - \frac{p\sigma^2}{n}
\end{aligned}$$

\square

Proof of Corollary 4. The first claim follows from a Jensen-type inequality for matrices in [49]. The second claim is a result of the mean of the inverse-Wishart distribution. \square

Proof of Theorem 3. For the first part of the theorem, consider without loss of generality the first entry of the vector `a_list` in Algorithm 1. This is the term (a) in (15). Similarly, the first entry entry of `b_list` is unbiased for (b) in (15). The result follows. \square

Proof of Lemma 5. We consider $\widehat{\text{Err}}^{(\text{OOB})}$ and $\widehat{\text{Err}}^{(.632)}$. Consider a single bootstrap sample $\mathcal{I} \subset \{1, \dots, n\}$, a multiset of cardinality n . We will show that both the residuals and out-of-bag error are identical when using either the data $(x_1, y_1), \dots, (x_n, y_n)$ or the data $(x_1, y'_1), \dots, (x_n, y'_n)$, where $y'_i = y_i + x_i^\top \kappa$.

Let $\hat{\theta}$ be the OLS estimate for θ based on the bootstrap sample \mathcal{I} with points $(x_1, y_1), \dots, (x_n, y_n)$, and let $\hat{\theta}'$ be the corresponding estimate for the data $(x_1, y'_1), \dots, (x_n, y'_n)$. Then

$$\begin{aligned}\hat{\theta} &= \operatorname{argmin}_{\theta} \sum_{i \in \mathcal{I}} (y_i - x_i^\top \theta)^2 \\ &= \operatorname{argmin}_{\theta} \sum_{i \in \mathcal{I}} (y_i + \kappa^\top x_i - x_i^\top (\theta + \kappa))^2 \\ &= -\kappa + \operatorname{argmin}_{\theta} \sum_{i \in \mathcal{I}} (y'_i - x_i^\top \theta)^2 \\ &= \hat{\theta}' - \kappa.\end{aligned}$$

As a result

$$y_i - x_i^\top \hat{\theta} = y'_i - x_i^\top \hat{\theta}'$$

for $i = 1, \dots, n$. Both the out-of-bag and .632 bootstrap estimators are functions only of the quantities $y_i - x_i^\top \hat{\theta}$, across many different bootstrap samples, so the proof is complete. \square

C Additional technical results

Lemma 6. *In the setting of Proposition 1, with notation as in Section 3.3,*

$$(\operatorname{Err} - \mathbb{E}[\operatorname{Err}_{\tilde{X}\tilde{Y}}])^2 = \Theta(1).$$

Proof of Lemma 6. Using the notation from the proof of Theorem 2,

$$\begin{aligned}\operatorname{Err} &= \sigma^2 + \mathbb{E} \left[\frac{\sigma^2}{n} \left\| \Sigma^{1/2} \hat{\Sigma}^{-1/2} Z \right\|^2 \right] \\ &= \sigma^2 + \frac{\sigma^2}{n} \mathbb{E} \left[\operatorname{tr}(\Sigma \hat{\Sigma}^{-1}) \right] \\ &= \sigma^2 + \frac{\sigma^2 p}{n - p - 1}.\end{aligned}$$

The result follows. \square

Proposition 3. *In the setting of Theorem 2, with notation as in Section 3.3, suppose also that $\mathcal{I}^{(\text{train})} > p$ and $|\mathcal{I}^{(\text{out})}|/n \rightarrow c \in (0, 1)$.*

$$\mathbb{E} \left[\left(\widehat{\operatorname{SE}}^{(\text{split})} \right)^2 \right] = \mathbb{E} \left[\left(\widehat{\operatorname{Err}}^{(\text{split})} - \operatorname{Err}_{XY} \right)^2 \right] - \operatorname{var}(\operatorname{Err}_{\tilde{X}\tilde{Y}}) - \operatorname{var}(\operatorname{Err}_{XY}) - (\operatorname{Err} - \mathbb{E}[\operatorname{Err}_{\tilde{X}\tilde{Y}}])^2 + O(1/n^2). \quad (18)$$

In the proportional asymptotic limit (7), for both sides of (18) above all terms are of order $1/n$, except the $O(1/n^2)$ term on the right-hand side.

Proof of Proposition 3. From Proposition 1, we have

$$\mathbb{E} \left[\left(\widehat{\operatorname{SE}}^{(\text{split})} \right)^2 \right] = \mathbb{E} \left[\left(\widehat{\operatorname{Err}}^{(\text{split})} - \operatorname{Err} \right)^2 \right] - \operatorname{var}(\operatorname{Err}_{\tilde{X}\tilde{Y}}) - (\operatorname{Err} - \mathbb{E}[\operatorname{Err}_{\tilde{X}\tilde{Y}}])^2.$$

We now expand the first term:

$$\begin{aligned}
\mathbb{E} \left[\left(\widehat{\text{Err}}^{(\text{split})} - \text{Err} \right)^2 \right] &= \mathbb{E} \left[\left(\widehat{\text{Err}}^{(\text{split})} - \text{Err}_{XY} + \text{Err}_{XY} - \text{Err} \right)^2 \right] \\
&= \mathbb{E} \left[\left(\widehat{\text{Err}}^{(\text{split})} - \text{Err}_{XY} \right)^2 \right] + \text{var}(\text{Err}_{XY}) + \\
&\quad 2\mathbb{E} \left[\left(\widehat{\text{Err}}^{(\text{split})} - \text{Err}_{XY} \right) \cdot (\text{Err}_{XY} - \text{Err}) \right] \\
&= \mathbb{E} \left[\left(\widehat{\text{Err}}^{(\text{split})} - \text{Err}_{XY} \right)^2 \right] + \text{var}(\text{Err}_{XY}) + \\
&\quad 2 \cdot \text{cov} \left(\widehat{\text{Err}}^{(\text{split})} - \text{Err}_{XY}, \text{Err}_{XY} - \text{Err} \right) \\
&= \mathbb{E} \left[\left(\widehat{\text{Err}}^{(\text{split})} - \text{Err}_{XY} \right)^2 \right] - \text{var}(\text{Err}_{XY}) + \\
&\quad 2 \cdot \text{cov} \left(\widehat{\text{Err}}^{(\text{split})}, \text{Err}_{XY} \right)
\end{aligned}$$

Thus, it remains to show only that $\text{cov} \left(\widehat{\text{Err}}^{(\text{split})}, \text{Err}_{XY} \right) = o(1/n)$. To this end,

$$\text{cov} \left(\widehat{\text{Err}}^{(\text{split})}, \text{Err}_{XY} \right) = \text{cov} (\text{Err}_{\tilde{X}}, \text{Err}_X)$$

by the conditional covariance decomposition (conditioning on X) and applying Theorem 1. Applying Cauchy-Schwarz,

$$\text{cov} (\text{Err}_{\tilde{X}}, \text{Err}_X) \leq \sqrt{\text{var}(\text{Err}_X) \cdot \text{var}(\text{Err}_{\tilde{X}})} = O(n^{-2}),$$

where in the last equality we applied Theorem 2. □

D Further simulation results

D.1 Compute times

In Table D.5, we report on the runtime of CV and NCV for our experiments.

D.2 Additional details on experiments from Section 5

This section reports additional results from the set of experiments in Section 5.

D.3 Number of folds

We next investigate how the number of folds affects the *CV inflation*: the ratio of the true standard error of the point estimate compared to the CV estimate of standard error. We consider a linear model with $p = 20$ features that are sampled as i.i.d. standard Gaussians. The number of observations ranges from 50 to 400. We use OLS as the fitting algorithm, so as a result of Lemma 1, the results do not depend on the true coefficients θ . We report on the results in Figure D.4, where we find that the number of folds has minimal impact on the inflation, although more folds gives moderately better coverage for small n . We also find that even when n/p is as large as 20, there is appreciable CV inflation, and naïve cross-validation leads to intervals with poor coverage.

Section	Experiment	CV time (s)	NCV time (s)
5.1.1	Bayes Error 33.2%	0.02	12.2
"	Bayes Error 22.5%	0.02	13.5
5.1.2	$n = 90, \rho = 0$	0.16	6.3
"	$n = 200, \rho = 0$	0.2	17
"	$n = 90, \rho = 0.5$	0.2	11
5.2.1	$n = 40$	0.02	6.6
"	$n = 100$	0.01	7.1
"	$n = 200$	0.01	7.8
"	$n = 400$	0.02	10
"	$n = 1600$	0.02	22
5.2.2	$n = 50$	0.1	44
"	$n = 100$	0.1	46
6	CC, $n = 50$	0.02	15
"	CC, $n = 100$	0.02	15
"	crp, $n = 50$	0.2	22
"	crp, $n = 100$	0.2	24

Table D.5: Approximate computation times for one run of CV and NCV for each of the experimental settings.




Figure D.1: Size of the nested CV intervals relative to the size of the naïve CV intervals in the low-dimensional logistic regression example from Section 5.1.1.

D.4 CV in the proportional region

We next show some experimental results further exploring the regime from Section 3.2. In Figure D.5, we compare the accuracy of CV when covering Err compared to its accuracy when covering Err_{XY} , and we find that it has higher accuracy for the former, by a constant fraction as $n, p \rightarrow \infty$. We also see that Err_{XY} and $\widehat{\text{Err}}^{(\text{CV})}$ are essentially uncorrelated.

Next, we plot the coverage of CV in Figure D.6, and see that it is far from the nominal level, even as n and p grow and the intervals have oracle debiasing so that they are centered around the correct value. In Figure D.7, we show that these intervals have higher than the nominal miscoverage rate in both tails, confirming that the miscoverage rate is not due to bias—it is due to the intervals being too narrow.

D.5 Coverage of data splitting in proportional regime

Here, we record further results about data splitting in the setting from Section 3.3. In Figure D.8, we show that data splitting does not approach the nominal coverage rate, even when the intervals are debiased by an




Figure D.2: Width of nested CV intervals relative to the width of the naïve CV intervals in the experiment from Section 5.2.1




Figure D.3: Size of the nested CV intervals relative to the size of the naïve CV intervals in the high-dimensional sparse regression example from Section 5.2.2.

oracle to be centered at the correct value.

D.6 .632 bootstrap empirical influence function CIs

We next discuss standard error estimates with the bootstrap point estimates introduced previously in Section A. The OOB estimate of prediction error has an associated estimate of standard error, based on estimates of the empirical influence functions [5, 7]. In particular, from [7] we use the point estimate from (24) and the SE estimate from (36). To get an estimate of the standard error for the .632 estimator, we re-scale the estimated OOB standard error by the ratio of (24) to (17) of [7], as suggested therein. We investigate the coverage of these intervals on our real-data examples and report on the results in Table D.6. We find that the .632 confidence intervals are generally acceptable, with reasonable coverage. The intervals are typically, but not always, wider than the NCV intervals. The bootstrap point estimates are typically more biased than the NCV point estimates.

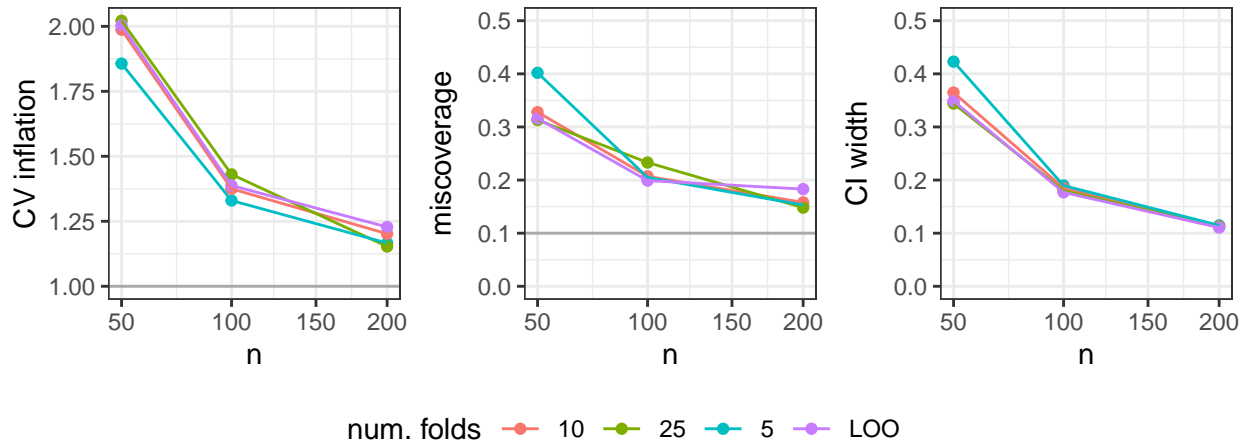


Figure D.4: The CV inflation and coverage of the naïve interval in the low-dimensional linear model with 20 features. The dark grey horizontal line in the middle panel gives the nominal miscoverage level. The right panel gives the average width of the confidence interval.




Figure D.5: Simulation results comparing the error of CV when estimating Err to its error when estimating Err_{XY} . Left: the mean absolute deviation between $\widehat{\text{Err}}^{(CV)}$ and Err or Err_{XY} . Right: $\text{cor}(\widehat{\text{Err}}^{(CV)}, \text{Err}_{XY})$

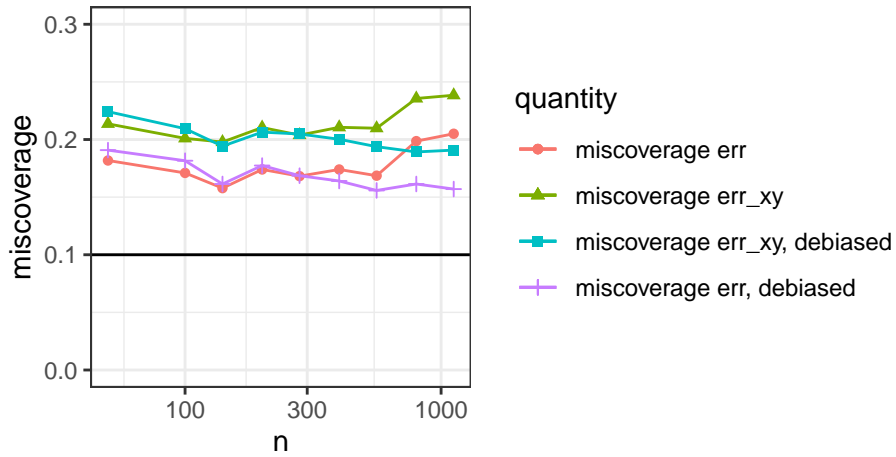


Figure D.6: Coverage of CV intervals with OLS in the proportional regime. Nominal coverage rate is 10%.




Figure D.7: Coverage of Err_{XY} of CV intervals with OLS and oracle debiasing in the proportional regime. Nominal coverage rate is 5% above and below.




Figure D.8: Coverage rate of data splitting with OLS. Nominal coverage rate is 10%.

Setting			Width and Point Estimates								Miscoverage							
Data	n	Target	Width			Err	CV	Estimate			CV		NCV		.632		OOB	
			NCV	.632	OOB			NCV	.632	OOB	Hi	Lo	Hi	Lo	Hi	Lo	Hi	Lo
CC	50	Err _{XY}	2.82	3.06	4.0	.029	.031	.029	.047	.067	4% 20%	1% 13%	0% 9%	2% 5%				
"	"	Err	"	"	"	"	"	"	"	"	2% 22%	0% 12%	0% 9%	1% 3%				
CC	100	Err _{XY}	1.46	1.82	2.22	.023	.024	.023	.025	.03	4% 13%	2% 7%	1% 8%	3% 4%				
"	"	Err	"	"	"	"	"	"	"	"	2% 12%	1% 4%	0% 6%	2% 2%				
crp	50	Err _{XY}	1.21	1.55	1.87	10.5%	10.7%	10.5%	12.0%	14.4%	7% 8%	2% 5%	1% 6%	1% 4%				
"	"	Err	"	"	"	"	"	"	"	"	6% 8%	3% 8%	1% 7%	2% 4%				
crp	100	Err _{XY}	1.10	1.20	1.32	9.5%	9.7%	9.6%	10.1%	11.1%	5% 7%	3% 6%	1% 8%	2% 5%				
"	"	Err	"	"	"	"	"	"	"	"	7% 7%	4% 6%	1% 8%	2% 5%				
S Lgstc	90	Err _{XY}	1.54	1.27	1.44	41.3%	41.8%	40.9%	40.9%	46.5%	17% 13%	7% 7%	4% 3%	19% 1%				
"	"	Err	"	"	"	"	"	"	"	"	17% 14%	7% 7%	2% 0%	11% 0%				

Table D.6: Performance of various methods with the real data sets. “S. Logistic” is the sparse logistic model from Section 1.1. Other details as in Table 1.

D.7 Number of Repeated Splits

In our experiments, we found that a large number (e.g., 200) of random splits of nested CV were needed to obtain stable estimates of the standard error. In Figure D.9, we show how the estimate of standard error relative to the naive CV estimate (*inflation*) converges as the number of repetitions increases for one example of the logistic regression model in Section 1.1.




Figure D.9: The nested CV inflation estimate after a large number of repeated splits.

E Variance stabilizing transformation for 0-1 loss

For classification settings with 0-1 loss, we can get better confidence intervals using a variance stabilizing transformation [e.g., 50]. If we observe an empirical miscoverage level of $\bar{e} \in [0, 1]$ from n samples, we form a confidence interval for $\sin^{-1}(\sqrt{\text{Err}})$ as

$$\left(\sin^{-1}(\sqrt{\bar{e}}) - z_{1-\alpha/2} \cdot \sqrt{\frac{1}{4n}}, \quad \sin^{-1}(\sqrt{\bar{e}}) + z_{1-\alpha/2} \cdot \sqrt{\frac{1}{4n}} \right),$$

and invert the transformation to get a confidence interval for Err (and analogously for Err_{XY} or other targets). With nested CV, we inflate the intervals by looking at the ratio between naïve CV and nested CV:

$$\left(\sin^{-1} \left(\sqrt{\widehat{\text{Err}}^{(\text{NCV})}} \right) - z_{1-\alpha/2} \cdot \frac{\sqrt{\widehat{\text{MSE}}}}{\widehat{\text{SE}}} \cdot \sqrt{\frac{1}{4n}}, \quad \sin^{-1} \left(\sqrt{\widehat{\text{Err}}^{(\text{NCV})}} \right) + z_{1-\alpha/2} \cdot \frac{\sqrt{\widehat{\text{MSE}}}}{\widehat{\text{SE}}} \cdot \sqrt{\frac{1}{4n}} \right).$$

(I.e., the transformed intervals are inflated by the amount that was estimated on the original scale.) Recall that $\sqrt{\widehat{\text{MSE}}}$ is the nested CV estimate of the width of the confidence interval, whereas $\widehat{\text{SE}}$ is the naïve CV estimate of the width of the confidence interval.

F A low-dimensional asymptotic analysis

In this section, we present a complementary asymptotic analysis to that of Section 3 in a more traditional asymptotic regime where p is fixed and $n \rightarrow \infty$. Here, we observed similar behavior; Err_X is close to Err , so we expect CV to have better accuracy for Err than for Err_{XY} . In this case, however, the behavior is all of order higher than $1/\sqrt{n}$. This means that for sufficiently large n , the difference between




Figure F.10: The relationship among various notions of prediction error in fixed p , $n \rightarrow \infty$ asymptotic limit. Recall that σ^2 is the Bayes error: the error rate of the best possible model

the various estimands is negligible compared to the variance of the CV estimator. Moreover, here Err and Err_{XY} approach the Bayes error at rate $1/n$, so asymptotically one has equal accuracy for estimating σ^2 , Err or Err_{XY} ; see [18] for a related discussion. Thus, in this section, the phenomena is only observable in higher-order asymptotics terms. See Figure F.10 for a visualization of the various rates.

Theorem 4. *Suppose the homoskedastic Gaussian linear model in (3) holds and that we use squared-error loss. In addition, assume that feature vectors $X_i \sim \mathcal{N}(0, \Sigma_p)$ for any full-rank Σ_p . Then, as $n \rightarrow \infty$ with p fixed, we have*

$$\mathbb{E}_X [\text{var}(\text{Err}_{XY} | X)] = \Theta(1/n^2)$$

and

$$\text{var}(\text{Err}_X) = \mathbb{E}(\text{Err}_X - \text{Err})^2 = \Theta(1/n^3).$$

The proof is included at the end of this section. This result can readily be extended beyond the case of Gaussian features, but we do not pursue this at this time. From this result, we can extract the following

Corollary 5. *In the setting of Theorem 4,*

$$\text{cor}(\text{Err}_{XY}, \text{Err}_X) \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Moreover, for any linearly invariant estimator $\widehat{\text{Err}}$ (such as $\widehat{\text{Err}}^{(\text{CV})}$ using OLS as the fitting algorithm),

$$\text{cor}(\text{Err}_{XY}, \widehat{\text{Err}}) \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Which means that CV not tracking Err_{XY} . The proof of this is as in the proof of Corollary 2, applying Theorem 4 in place of Theorem 2.

Proof of Theorem 4. The proof follows as in the proof of Theorem 2, noting that in this case

$$\mathbb{E}_X [\lambda_1^4 + \dots + \lambda_p^4] = \Theta(1)$$

and

$$\text{var} \left(\text{tr}(\Sigma \hat{\Sigma}^{-1}) \right) = \Theta(1/n).$$

□