

CONSISTENCY OF CROSS VALIDATION FOR COMPARING REGRESSION PROCEDURES

BY YUHONG YANG

University of Minnesota

Theoretical developments on cross validation (CV) have mainly focused on selecting one among a list of finite-dimensional models (e.g., subset or order selection in linear regression) or selecting a smoothing parameter (e.g., bandwidth for kernel smoothing). However, little has been known on consistency of cross validation when applied to compare between parametric and nonparametric methods or within nonparametric methods. We show that under some conditions, with an appropriate choice of data splitting ratio, cross validation is consistent in the sense of selecting the better procedure with probability approaching one.

Our results reveal interesting behaviors of cross validation. When comparing two models (procedures) converging at the same nonparametric rate, in contrast to the parametric case, it turns out that the proportion of data used for evaluation in CV does not need to be dominating in size. Furthermore, it can even be of a smaller order than the proportion for estimation while not affecting the consistency property.

1. Introduction. Cross validation (e.g., Allen (1974), Stone (1974), Geisser (1975)) is one of the most commonly used model selection criteria. Basically, based on a data splitting, part of the data is used for fitting each competing model (or procedure) and the rest of the data is used to measure the performance of the models, and the model with the best overall performance is selected. There are a few different versions of cross validation (CV) methods, including delete-one CV, delete- k ($k > 1$) CV and also generalized CV methods (e.g., Craven and Wahba (1979)).

Cross validation can be applied to various settings, including parametric

AMS 2000 subject classifications: Primary 62G07, 62B10; secondary 62C20

Keywords and phrases: consistency, cross validation, model selection

and nonparametric regressions. There can be different primary goals when applying a CV method: one mainly for identifying the best model/procedure among the candidates and another mainly for estimating the mean function or for prediction (see, e.g., Geisser (1975)). A number of theoretical results have been obtained, mostly in the areas of linear regression and in smoothing parameter selection for nonparametric regression. In linear regression, it has been shown that delete-1 and generalized CVs are asymptotically equivalent to AIC and they are all inconsistent in the sense that the probability of selecting the true model does not converge to 1 as n goes to ∞ (see, Li (1987)). In addition, interestingly, the analysis of Shao (1993) showed that in order for delete- k CV to be consistent, k needs to be dominatingly large in the sense that $k/n \rightarrow 1$ (and $n-k \rightarrow \infty$). Zhang (1993) proved that delete- k CV is asymptotically equivalent to the FPE criterion when $k \rightarrow \infty$. The readers are referred to Shao (1997) for more asymptotic results and references on model selection for linear regression. In the context of nonparametric regression, delete-1 CV for smoothing parameter selection leads to consistent regression estimators (e.g., Wong (1983) for kernel regression and Li (1984) for nearest neighbor method) and leads to asymptotically optimal or rate optimal choice of smoothing parameters and/or optimal regression estimation (see e.g., Speckman (1985) and Burman (1990) for spline estimation, Härdle, Hall and Marron (1988), Hall and Johnstone (1992) and references therein for kernel estimation). Györfi, et al (2002) gave risk bounds for kernel and nearest neighbor regression with bandwidth or neighbor size selected by delete-1 CV. See Opsomer, Wang and Yang (2000) for a review and references related to the use of CV for bandwidth selection for nonparametric

regression with dependent errors.

In real world applications of regression, in pursuing a better estimation accuracy, one may naturally consider the use of cross validation to choose between a parametric estimator and a nonparametric estimator (or at least to understand their relative performance). Similarly, when different types of nonparametric estimators are entertained as plausible candidates, cross validation is also applicable to choose one of them. Recently, a general CV methodology has been advocated by van der Laan, Dudoit, van der Vaart and their co-authors (e.g., van der Laan and Dudoit, 2003; van der Laan, Dudoit and van der Vaart, 2006; van der Vaart, Dudoit and van der Laan, 2006), which can be applied in other contexts (e.g., survival function estimation). Risk bounds for estimating the target function were derived and their implications on adaptive estimation and asymptotic optimality were obtained. When CV is used for the complementary purpose of identifying the best candidate, however, it is still unclear whether CV is generally consistent and if the data splitting ratio has a sensitive effect on consistency. For successful applications of CV in practice, a theoretical understanding of these issues is very much of interest.

In this paper, we address the aforementioned consistency issue and show that a voting-based cross validation is consistent for comparing general regression procedures when the data splitting ratio is properly chosen.

In the context of linear regression, Shao's result (1993) implies that with k/n not converging to 1, delete- k CV does not differentiate well between two correct models. However, in nonparametric regression, it turns out that this is not the case. In fact, as long as at least one of the competing procedures

converges at a nonparametric rate, the estimation size and evaluation size can be of the same order in data splitting, and sometimes the estimation size can even be the dominating one.

In the settings where theoretical properties of CV were investigated before, the best model (in the linear regression context) or the best smoothing parameter (in the case of kernel regression) exists in a natural way. When comparing two general estimators, the issue becomes more complicated. In this paper, we compare two estimators in terms of a loss function and the consistency in selection is established when one estimator is better than the other in that sense.

The paper is organized as follows. In Section 2, we set up the problem. The main result is presented in Section 3, followed by simulation results in Section 4. Concluding remarks are in Section 5. The proof of the main result is in Section 6.

2. Problem setup. Consider the regression setting:

$$Y_i = f(X_i) + \varepsilon_i, \quad 1 \leq i \leq n,$$

where $(X_i, Y_i)_{i=1}^n$ are independent observations with X_i iid taking values in a d -dimensional Borel set $\mathcal{X} \subset R^d$ for some $d \geq 1$, f is the true regression function and ε_i are the random errors with $E(\varepsilon_i|X_i) = 0$ and $E(\varepsilon_i^2|X_i) < \infty$ almost surely. The distribution of X_i is unknown.

Rates of convergence of various popular regression procedures have been well studied. Under the squared L_2 loss (and other closely related performance measures as well), for parametric regression (assuming that the true regression function has a known parametric form), estimators based on max-

imum likelihood (if available) or least squares usually converge at the rate n^{-1} . For nonparametric estimation, the convergence rates are slower than n^{-1} with the actual convergence rate depending on both the regression procedure and the smoothness property of the true regression function.

In applications, with many regression procedures available to be applied, it is often challenging to find the one with best accuracy. One often faces the issue: Should I go parametric or nonparametric? Which nonparametric procedure to use? As is well known, nonparametric procedures have more flexibility yet converge sub-optimally compared to estimation based on a correctly specified parametric model. Sometimes, for instance, there is a clear linear trend and naturally a simple linear model is a good candidate. It may be unclear, however, whether a nonparametric method can provide a better estimate to capture a questionable slight curvature seen in the data. For a specific example, for the rather famous Old Faithful Geyser data (Weisberg (2005)), there were several analyses related to the comparison of linear regression with nonparametric alternatives (see, e.g., Simonoff (1996), Hart (1997)).

For simplicity, suppose that there are two regression procedures, say δ_1 and δ_2 , that are considered. For example, δ_1 may be simple linear regression and δ_2 may be a local polynomial regression procedure (see, e.g., Fan and Gijbels (1996)). For another example, δ_1 may be a spline estimation procedure (see, e.g., Wahba (1990)) and δ_2 may be a wavelet estimation procedure (see, e.g., Donoho and Johnstone (1998)). Based on a sample $(X_i, Y_i)_{i=1}^n$, the regression procedures δ_1 and δ_2 yield estimators $\hat{f}_{n,1}(x)$ and $\hat{f}_{n,2}(x)$ respectively. We need to select the better one of them.

Though for simplicity we assumed that there are only two competing regression procedures, similar results hold when a finite number of candidate regression procedures are in competition. We emphasize that our focus in this work is not on tuning a smoothing parameter of a regression procedure (such as the bandwidth for kernel regression). Rather, our result is general and applicable also for the case when the candidate regression procedures are very distinct, possibly with different rates of convergence for estimating the regression function.

Cross validation is a natural approach to address the above model/procedure comparison issue. It has the advantage that it requires mild distributional assumptions on the data and it does not need to find characteristics such as degrees of freedom or model dimension for each model/procedure (which is not necessarily easy for complicated adaptive nonparametric procedures or even parametric procedures where model selection is conducted). To proceed, we split the data into two parts: the estimation data consist of $Z^1 = (X_i, Y_i)_{i=1}^{n_1}$ and the validation data consist of $Z^2 = (X_i, Y_i)_{i=n_1+1}^n$. Let $n_2 = n - n_1$. We apply δ_1 and δ_2 on Z^1 to obtain the estimators $\hat{f}_{n_1,1}(x)$ and $\hat{f}_{n_1,2}(x)$ respectively. Then we compute the prediction squared errors of the two estimators on Z^2 :

$$(2.1) \quad CV(\hat{f}_{n_1,j}) = \sum_{i=n_1+1}^n \left(Y_i - \hat{f}_{n_1,j}(X_i) \right)^2, \quad j = 1, 2.$$

If $CV(\hat{f}_{n_1,1}) \leq CV(\hat{f}_{n_1,2})$, δ_1 is selected and otherwise δ_2 is chosen. We call this a delete- n_2 CV. For investigating the consistency property of CV, in this work, we only consider the case when $\min(n_1, n_2) \rightarrow \infty$.

A more balanced use of the data is CV with multiple data splittings. A

voting-based version will be considered as well.

3. Consistency of cross validation. In this section, we first define the concept of consistency when two general regression procedures are compared and then state the conditions that are needed for our result. Note that there are two types of consistency results on CV. One is from the selection perspective (which concerns if the true model or the best candidate procedure is selected with probability tending to 1) and the other is in terms of convergence of the resulting estimator of the regression function (either in probability or in risk under a proper loss function). The former is our focus in this work. Note that in general, neither notion guarantees the other. Nevertheless, for a global loss function, consistency in selection usually implies consistency in estimation of the regression function at least in probability. In contrast, consistency in estimation can be far away from consistency in selection. For example, with two nested models, if both models are correct, then reasonable selection rules will be consistent in estimation but not necessarily so in selection. See Wegkamp (2003) for risk bounds for a modified CV (with an extra complexity penalty) for estimating the regression function.

3.1. *Definitions and conditions.* We first give some useful definitions for comparing estimators in terms of probability.

Let $L(\theta, \hat{\theta})$ be a loss function. Consider two estimation procedures δ and δ' for estimating a parameter θ . Let $\{\hat{\theta}_{n,1}\}_{n=1}^{\infty}$ and $\{\hat{\theta}_{n,2}\}_{n=1}^{\infty}$ be the corresponding estimators when applying the two procedures at sample sizes 1, 2, ... respectively.

Definition 1. Procedure δ (or $\{\hat{\theta}_{n,1}\}_{n=1}^{\infty}$, or simply $\hat{\theta}_{n,1}$) is asymptotically

better than δ' (or $\{\hat{\theta}_{n,2}\}_{n=1}^\infty$, or $\hat{\theta}_{n,2}$) under the loss function $L(\theta, \hat{\theta})$ if for every $0 < \epsilon < 1$, there exists a constant $c_\epsilon > 0$ such that when n is large enough,

$$(3.1) \quad P\left(L(\theta, \hat{\theta}_{n,2}) \geq (1 + c_\epsilon)L(\theta, \hat{\theta}_{n,1})\right) \geq 1 - \epsilon.$$

Remarks:

1. Suppose that $\hat{\theta}_{n,1}$ and $\tilde{\theta}_{n,1}$ are asymptotically equivalent under the loss function $L(\theta, \hat{\theta})$ in the sense that $L(\theta, \hat{\theta}_{n,1})/L(\theta, \tilde{\theta}_{n,1}) \rightarrow 1$ in probability. If $\hat{\theta}_{n,1}$ is asymptotically better than $\hat{\theta}_{n,2}$, then obviously $\tilde{\theta}_{n,1}$ is also asymptotically better than $\hat{\theta}_{n,2}$.

2. It seems clear that to evaluate a selection method that chooses between two procedures, the procedures need to be rankable. When two procedures are asymptotically equivalent, one may need to examine finer differences (e.g., higher order behaviors) to compare them. When two procedures can have “ties” in the sense that on a set with a non-vanishing probability the two procedures are identical or behave the same, it becomes tricky to define consistency of selection generally speaking.

The requirement in (3.1) is sensible. Obviously, if $L(\theta, \hat{\theta}_{n,2})/L(\theta, \hat{\theta}_{n,1}) \rightarrow \infty$ in probability, by the definition, δ is asymptotically better than δ' under L . The concept is also useful for comparing procedures that converge at the same order. In particular, it is worth pointing out that in the context of linear regression with an appropriate loss (e.g., global L_2 loss), when two correct models are compared, typically the estimator based on the one with a lower dimension is asymptotically better. As is expected, in the same context, an incorrect subset model yields an asymptotically worse estimator than any one based on a correct model.

We next define consistency of a model/procedure selection rule from the perspective of selecting the best (or better when there are only two candidates) model/procedure. Unless otherwise stated, consistency of CV refers to this notion in the rest of the paper.

Definition 2. Assume that one of the candidate regression procedures, say δ^* , is asymptotically better than the other candidate procedures. A selection rule is said to be consistent if the probability of selecting δ^* approaches 1 as $n \rightarrow \infty$.

This concept of consistency for a selection rule is more general than the definition that a selection rule is consistent if the true model (when existing and being considered) is selected with probability approaching 1. Obviously the latter does not apply for comparing two general regression procedures.

Let $\{a_n\}$ be a sequence of positive numbers approaching zero. The following simple definition concerns the rate of convergence in probability (cf. Stone (1980)).

Definition 3. A procedure δ (or $\{\hat{\theta}_n\}_{n=1}^\infty$) is said to converge exactly at rate $\{a_n\}$ in probability under the loss L if $L(\theta, \hat{\theta}_n) = O_p(a_n)$, and for every $0 < \epsilon < 1$, there exists $c_\epsilon > 0$ such that when n is large enough, $P(L(\theta, \hat{\theta}_n) \geq c_\epsilon a_n) \geq 1 - \epsilon$.

Clearly, the latter part of the condition in the above definition says that the estimator does not converge faster than the rate a_n .

Define the L_q norm

$$\|f\|_q = \begin{cases} (\int |f(x)|^q P_X(dx))^{1/q} & \text{for } 1 \leq q < \infty \\ \text{esssup}|f| & \text{for } q = \infty, \end{cases}$$

where P_X denotes the probability distribution of X_1 .

We assume that $\widehat{f}_{n,1}$ converges exactly at rate p_n in probability and $\widehat{f}_{n,2}$ converges exactly at rate q_n in probability under the L_2 loss. Note that p_n and q_n may or may not converge at the same rate.

Condition 0. (*Error variances*) The error variances $E(\varepsilon_i^2|X_i)$ are upper bounded by a constant $\bar{\sigma}^2 > 0$ almost surely for all $i \geq 1$.

This is a mild condition on the errors, which does not require them to be identically distributed. We also need to control the sup-norm of the estimators.

Condition 1. (*Sup-norm of the estimators*) There exists a sequence of positive numbers A_n such that for $j = 1, 2$, $\|f - \widehat{f}_{n,j}\|_\infty = O_p(A_n)$.

The condition almost always holds. But for our main result to be helpful, the constants A_n need to be suitably controlled.

We mention some useful sufficient conditions that imply Condition 1. One is that for $j = 1, 2$, $\|f - \widehat{f}_{n,j}\|_\infty$ is bounded in probability. Another stronger condition is that $E\|f - \widehat{f}_{n,j}\|_\infty$ is uniformly bounded in n for $j = 1, 2$. It clearly holds if with probability 1, $\|f - \widehat{f}_{n,j}\|_\infty$ is upper bounded by a constant $A > 0$, which is satisfied if the true regression function is bounded between two known constants and the estimators are accordingly restricted.

In order to have consistency in selection, we need that one procedure is better than the other.

Condition 2. (*One procedure being better*) Under the L_2 loss, either δ_1 is asymptotically better than δ_2 , or δ_2 is asymptotically better than δ_1 .

Clearly there are situations where neither of two competing procedures is asymptotically better than the other. In such a case, the concept of consistency is hard to define (or may be irrelevant). Note that under Condition

2, if δ_1 is asymptotically better than δ_2 , then we have $p_n = O(q_n)$. Clearly, if there exists a constant $C > 1$ such that with probability approaching one, $\|f - \hat{f}_{n,2}\|_2 \geq C \|f - \hat{f}_{n,1}\|_2$ or even $\|f - \hat{f}_{n,2}\|_2 / \|f - \hat{f}_{n,1}\|_2 \rightarrow \infty$ in probability, then Condition 2 is satisfied with δ_1 being better.

Another quantity, namely, $\|f - \hat{f}_{n,j}\|_4 / \|f - \hat{f}_{n,j}\|_2$ is involved in our analysis. Obviously, the ratio is lower bounded by 1 and there can not be any general positive upper bound. For various estimators, the ratio can be controlled. We need the following condition.

Condition 3. (*Relating L_4 and L_2 losses*) There exists a sequence of positive numbers $\{M_n\}$ such that for $j = 1, 2$, $\|f - \hat{f}_{n,j}\|_4 / \|f - \hat{f}_{n,j}\|_2 = O_p(M_n)$.

For many familiar infinite-dimensional classes of regression functions, the optimal rates of convergence under L_4 and L_2 are the same. If we consider an optimal estimator (in rate) under L_4 , then we can take M_n to be 1 for a typical f in such a function class. Parametric estimators typically have the ratio upper bounded in probability (i.e., $M_n = 1$) under some mild conditions.

For some nonparametric estimators, the sup-norm risk is often of only a slightly higher order than that under L_p for $p < \infty$. For example, for Hölder classes $\Sigma(\beta, L) = \{f : |f^{(m)}(x) - f^{(m)}(y)| \leq L|x - y|^\alpha\}$, where $m = \lfloor \beta \rfloor$ is an integer, $0 < \alpha \leq 1$ and $\alpha = \beta - m$, also $\|f\|_\infty$ is bounded, or for Sobolev classes, the rates of convergence under the sup-norm distance and L_p ($p < \infty$) are different only by a logarithmic factor (see, e.g., Stone (1982) and Nemirovski (2000)). If one takes an optimal or near optimal estimator under the L_∞ loss, Condition 3 is satisfied typically with M_n being a logarithmic

term.

3.2. *The main theorem.* Let $I^* = 1$ if δ_1 is asymptotically better than δ_2 and $I^* = 2$ if δ_2 is asymptotically better than δ_1 . Let $\hat{I}_n = 1$ if $CV(\hat{f}_{n_1,1}) \leq CV(\hat{f}_{n_1,2})$ and otherwise $\hat{I}_n = 2$.

Theorem 1. Under Conditions 0-3, if the data splitting satisfies 1) $n_2 \rightarrow \infty$ and $n_1 \rightarrow \infty$; 2) $n_2 M_{n_1}^{-4} \rightarrow \infty$; and 3)

$$\sqrt{n_2} \max(p_{n_1}, q_{n_1}) / (1 + A_{n_1}) \rightarrow \infty,$$

then the delete- n_2 CV is consistent, i.e., $P(\hat{I}_n \neq I^*) \rightarrow 0$ as $n \rightarrow \infty$.

Remarks:

1. The third requirement above is equivalent to $\sqrt{n_2} \max(p_{n_1}, q_{n_1}) \rightarrow \infty$ and $\sqrt{n_2} \max(p_{n_1}, q_{n_1}) / A_{n_1} \rightarrow \infty$. The latter has no effect, for example, when the estimators being compared by CV both converge in the sup-norm in probability. The effect of M_{n_1} is often more restrictive than A_{n_1} and sometimes can be more complicated to deal with.

2. Theorem 1 is stated under a single data generating distribution. Obviously, model selection becomes useful when there are various possible data generating mechanisms (e.g., corresponding to different parametric or non-parametric families of regression functions or different assumptions on the errors) that are potentially suitable for the data at hand. In the linear regression context, a number of finite-dimensional models is considered, and in the literature, a model selection rule is said to be consistent when the true model is selected with probability going to 1 no matter what the true model is (assumed to be among the candidates). Clearly our theorem can be used to get such a result when multiple scenarios of the data generating process

are possible. To that end, one just needs to find a data splitting ratio that works for all the scenarios being considered.

3. A potentially serious disadvantage of cross validation is that when two candidate regression procedures are hard to be distinguished, the forced action of choosing a single winner can substantially damage the accuracy of estimating the regression function. An alternative is to average the estimates. See Yang (2001, 2003) for references and theoretical results on combining models/procedures and simulation results that compare CV and a model combining procedure ARM.

It should be pointed out that the conclusion is sharp in the sense that there are cases in which the sufficient conditions on data splitting are also necessary. See Section 3.4 for details.

The “ideal” norm conditions are that $M_n = O(1)$ and $A_n = O(1)$. Since almost always, p_n and q_n are at least of order $n^{-1/2}$, the most stringent third requirement then is $n_2/n_1 \rightarrow \infty$.

Corollary 1. Under the same conditions in Theorem 1, if $M_n = O(1)$ and $A_n = O(1)$, then the delete- n_2 CV is consistent for each of the following two cases:

- (i). $\max(p_n, q_n) = O(n^{-1/2})$, with the choice $n_1 \rightarrow \infty$ and $n_2/n_1 \rightarrow \infty$;
- (ii). $\max(p_n, q_n)n^{1/2} \rightarrow \infty$, with any choice such that $n_1 \rightarrow \infty$ and $n_1/n_2 = O(1)$.

From the corollary, if the L_q norms of $\hat{f}_{n,1} - f$ and $\hat{f}_{n,2} - f$ behave “nicely” for $q = 2, 4$ and ∞ , then when at least one of the regression procedure converges at a nonparametric rate under the L_2 loss, any split ratio in CV works for consistency as long as both sizes tend to infinity and the estimation

size is no bigger (in order) than the evaluation size. Note that this splitting requirement is sufficient but not necessary. For example, if $\max(p_n, q_n) = O(n^{-1/4})$, then the condition $\sqrt{n_2} \max(p_{n_1}, q_{n_1}) \rightarrow \infty$ becomes $n_2/n_1^{1/2} \rightarrow \infty$ (i.e., $n_1 = o(n_2^2)$). Thus, for example, we can take $n_1 = n - \lfloor \sqrt{n} \log n \rfloor$ and $n_2 = \lfloor \sqrt{n} \log n \rfloor$, in which case the estimation proportion is dominating. This is in sharp contrast to the requirement of a much larger evaluation size ($n_2/n_1 \rightarrow \infty$) when both of the regression procedures converge at the parametric rate n^{-1} , which was discovered by Shao (1993) in the context of linear regression with fixed design. From Shao's results, one may expect that when p_n and q_n are of the same order, the condition $n_2/n_1 \rightarrow \infty$ may also be needed for consistency more generally. But, interestingly, our result shows that this is not the case. Thus there is a paradigm shift in terms of splitting proportion for CV when at least one of the regression procedure is nonparametric.

The result also suggests that in general delete-one (or delete a small fraction) is not geared towards finding out which candidate procedure is the better one.

Note that for various estimators based on local averaging or series expansion, we do not necessarily have the "ideal" norm requirement in Condition 3 met with $M_n = O(1)$, but may have $\|f - \hat{f}_n\|_4 \leq \|f - \hat{f}_n\|_\infty \leq a_n \|f - \hat{f}_n\|_2$ for some deterministic sequence a_n (possibly converging to ∞ at a polynomial order n^γ with $\gamma > 0$). Then for applying the theorem, we may need $n_2/n_1^{4\gamma} \rightarrow \infty$. This may or may not add any further restriction on the data splitting ratio in CV beyond the other requirements in the theorem, depending on the value of γ in relation to p_n, q_n and A_n .

3.3. *CV with multiple data splittings.* For Theorem 1, the data splitting in CV is done only once (and hence the name of *cross* validation may not be appropriate there). Clearly, the resulting estimator depends on the order of the observations. In real applications, one may do any of the following: 1) consider all possible splits with the same ratio (this is called *multifold CV*, see e.g., Zhang (1993)); 2) the same as in 1) but consider only a sample of all possible splits (this is called *repeated learning-testing*, see, e.g., Burman (1989)); 3) divide the data into r sub-groups and do prediction one at a time for each sub-group based on estimation using the rest sub-groups (this is called *r -fold CV*, see, Breiman, Friedman, Olshen, and Stone (1984)). When multiple splittings are used, there are two natural ways to proceed. One is to first average the prediction errors over the different splittings and then select the procedure that minimizes the average prediction error. Another is to count the number of times each candidate is preferred under the different splittings and then the candidate with the highest count is the overall winner. Such a voting is natural to consider for model selection. To make a distinction, we call the former (i.e., CV with averaging) CV-a and the latter (i.e., CV with voting) CV-v. When focusing on linear models with fixed design, theoretical properties for *multifold CV-a* or *r -fold CV-a* were derived under assumptions on the design matrix (see, e.g., Zhang (1993) and Shao (1993)). We next show that under the same conditions used for a single data splitting in the previous subsection, CV-v based on multiple data splittings is also consistent in selection when the observations are identically distributed.

Let π denote a permutation of the observations. Let $CV_\pi(\hat{f}_{n_{1,j}})$ be the

criterion value as defined in (2.1) except that the data splitting is done after the permutation π . If $CV_\pi(\hat{f}_{n_1,1}) \leq CV_\pi(\hat{f}_{n_1,2})$, then let $\tau_\pi = 1$ and otherwise let $\tau_\pi = 0$.

Let Π denote the set of all $n!$ permutations of the observations. If $\sum_{\pi \in \Pi} \tau_\pi \geq \frac{n!}{2}$, then we select δ_1 , and otherwise δ_2 is selected.

Theorem 2. Under the conditions in Theorem 1 and that the observations are independent and identically distributed, the CV-v procedure above is consistent in selection.

Proof of Theorem 2. Without loss of generality, assume that δ_1 is better than δ_2 . Let W denote the values of $(X_1, Y_1), \dots, (X_n, Y_n)$ (ignoring the orders). Under the i.i.d assumption on the observations, obviously, conditional on W , every ordering of these values has exactly the same probability and thus $P(CV(\hat{f}_{n_1,1}) \leq CV(\hat{f}_{n_1,2}))$ is equal to

$$EP(CV(\hat{f}_{n_1,1}) \leq CV(\hat{f}_{n_1,2})|W) = E\left(\frac{\sum_{\pi \in \Pi} \tau_\pi}{n!}\right).$$

From Theorem 1, under the given conditions, $P(CV(\hat{f}_{n_1,1}) \leq CV(\hat{f}_{n_1,2})) \rightarrow 1$. Thus $E(\sum_{\pi \in \Pi} \tau_\pi/n!) \rightarrow 1$. Since $\sum_{\pi \in \Pi} \tau_\pi/n!$ is between 0 and 1, for its expectation to converge to 1, we must have $\sum_{\pi \in \Pi} \tau_\pi/n! \rightarrow 1$ in probability. Consequently $P(\sum_{\pi \in \Pi} \tau_\pi \geq n!/2) \rightarrow 1$. This completes the proof of Theorem 2.

From the above proof, it is clearly seen that the consistency result also holds for the aforementioned voting-based repeated learning-testing and r -fold CV-v methods.

Note that for the CV-a methods, for each candidate procedure, we av-

erage $CV(\hat{f}_{n_1,j})$ over the different data splittings first and then compare the criterion values to select a winner. Intuitively, since the CV-v only keeps the ranking of the procedures for each splitting, it may have lost some useful information in the data. If so, the CV-v may perform worse than CV-a. However, in terms of the consistency property in selection, it is unlikely that the two versions of CV are essentially different. Indeed, if $P\left(\sum_{\pi \in \Pi} (CV_{\pi}(\hat{f}_{n_1,2}) - CV_{\pi}(\hat{f}_{n_1,1})) > 0\right) \rightarrow 0$, due to symmetry, one expects that for the majority of splittings in Π , with high probability, we have $CV_{\pi}(\hat{f}_{n_1,2}) - CV_{\pi}(\hat{f}_{n_1,1}) > 0$. Then $P\left(\sum_{\pi \in \Pi} I_{(CV_{\pi}(\hat{f}_{n_1,2}) - CV_{\pi}(\hat{f}_{n_1,1}) > 0)} > \frac{|\Pi|}{2}\right)$ is close to 1.

Based on the above reasoning, we conjecture that the two CV methods generally share the same status of consistency in selection (i.e., if one is consistent so is the other). This of course does not mean that they perform similarly at a finite sample size. In the simulations reported in Section 4, we see that for comparing two converging parametric procedures, CV-v performs much worse than CV-a, but when CV-a selects the best procedure with probability closer to 1, their difference becomes small; for comparing a converging nonparametric procedure with a converging parametric one, CV-v often performs better. We tend to believe that the differences between CV-v and CV-a are second order effects, which are hard to quantify in general.

3.4. *Are the sufficient conditions on data splitting also necessary for consistency?* Theorem 2 shows that if a single data splitting ensures consistency, voting based on multiple splittings also works. In the reverse direction, one may wonder if multiple splittings with averaging can rescue an inconsis-

tent CV selection with only one splitting. We give a counter-example below.

In the following example, we consider two simple models which allow us to exactly identify the sufficient and necessary conditions for ensuring consistency in selection by a CV-a method. Model 1 is $Y_i = \varepsilon_i, i = 1, \dots, n$, where ε_i are iid normal with mean zero and variance σ^2 . Model 2 is $Y_i = \mu + \varepsilon_i, i = 1, \dots, n$, with the same conditions on the errors. Clearly, model 1 is a sub-model of model 2. Under model 1, obviously $\hat{f}_{n,1}(x) = 0$, and under model 2, the maximum likelihood method yields $\hat{f}_{n,2}(x) = \bar{Y}$. We consider the multifold CV with splitting ratio $n_1 : n_2$, where again n_1 is the estimation sample size and n_2 is the evaluation size. Note that for the multifold CV-a method, we consider all possible data splittings at the given ratio and the averaging of the prediction errors is done over the splittings. Let S denote the observations in the estimation set. Then

$$CV(j) = \sum_S \left(\sum_{i \in S^c} (Y_i - \hat{Y}_{i,j})^2 \right),$$

where S is over all possible data splittings with the given ratio, S^c denotes the complement of S , $\hat{Y}_{i,1} = 0$, and $\hat{Y}_{i,2} = \frac{1}{n_1} \sum_{l \in S} Y_l$.

Proposition 1: A sufficient and necessary condition for the above CV to be consistent in selection is that the data splitting ratio satisfies: 1) $n_1 \rightarrow \infty$, 2) $n_2/n_1 \rightarrow \infty$.

Note that the conditions in this proposition match those sufficient conditions in Theorem 1. Consequently, we know that for the consistency property, multiple splittings (or even all possible splittings) do not help in this case. Although it seems technically difficult to derive a similar general result, we tend to believe that the example represents the typical situation of comparing nested parametric models.

In the context of variable selection in linear regression, Zhang (1993) showed that for any fixed splitting ratio, none of the multifold CV, repeated learning-testing, and r -fold CV based on averaging is consistent in selection.

Proof of Proposition 1: With some calculations, under the larger model, we have $CV(1) = \binom{n-1}{n_2-1} (n\mu^2 + 2\mu \sum_{i=1}^n \varepsilon_i + \sum_{i=1}^n \varepsilon_i^2)$ and $CV(2)$ equals

$$\begin{aligned} \binom{n-1}{n_2-1} \sum_{i=1}^n \varepsilon_i^2 &- 2 \binom{n-2}{n_1-1} \frac{1}{n_1} \sum_{i=1}^n \varepsilon_i \left(\sum_{j \neq i}^n \varepsilon_j \right) \\ &+ \frac{n_2}{n_1^2} \left(\binom{n-1}{n_1-1} \sum_{i=1}^n \varepsilon_i^2 + 2 \binom{n-2}{n_1-2} \sum_{1 \leq j_1 < j_2 \leq n} \varepsilon_{j_1} \varepsilon_{j_2} \right). \end{aligned}$$

Then with more simplifications, we get that $CV(1) - CV(2)$ equals

$$\binom{n-1}{n_2-1} \left(n\mu^2 + 2\mu \sum_{i=1}^n \varepsilon_i \right) + \sum_{i=1}^n \varepsilon_i^2 \cdot \frac{(n-2)!(n_1+1)}{n_1 n_1! (n_2-1)!} - \left(\sum_{i=1}^n \varepsilon_i \right)^2 \cdot \frac{(n-2)!(n_1+n)}{n_1 n_1! (n_2-1)!}.$$

When $\mu \neq 0$, it is not hard to show that $CV(1) - CV(2)$ is positive with probability tending to one if and only if $n_1 \rightarrow \infty$. When model 1 holds, i.e., $\mu = 0$, $CV(1) - CV(2) > 0$ is equivalent to $\frac{(n-1)n_1}{n} (\sum_{i=1}^n \varepsilon_i)^2 > (n + n_1) \sum_{i=1}^n (\varepsilon_i - \bar{\varepsilon})^2$. Since $\sum_{i=1}^n \varepsilon_i$ is independent of $\sum_{i=1}^n (\varepsilon_i - \bar{\varepsilon})^2$ and they have normal and Chi-square distributions respectively, we know that for the probability of the event to go to zero, we must have $n_1/n \rightarrow 0$, which is also sufficient. This completes the proof of the proposition.

4. Simulation. In this section, we present some simulation results that are helpful to understand the differences of several versions of CV and the effect of data splitting ratio.

We consider estimating a regression function on $[0, 1]$ with three competing regression methods. The true model is $Y_i = f(X_i) + \varepsilon_i$, $1 \leq i \leq n$, where X_i are iid uniform in the unit interval, and the errors are independent of

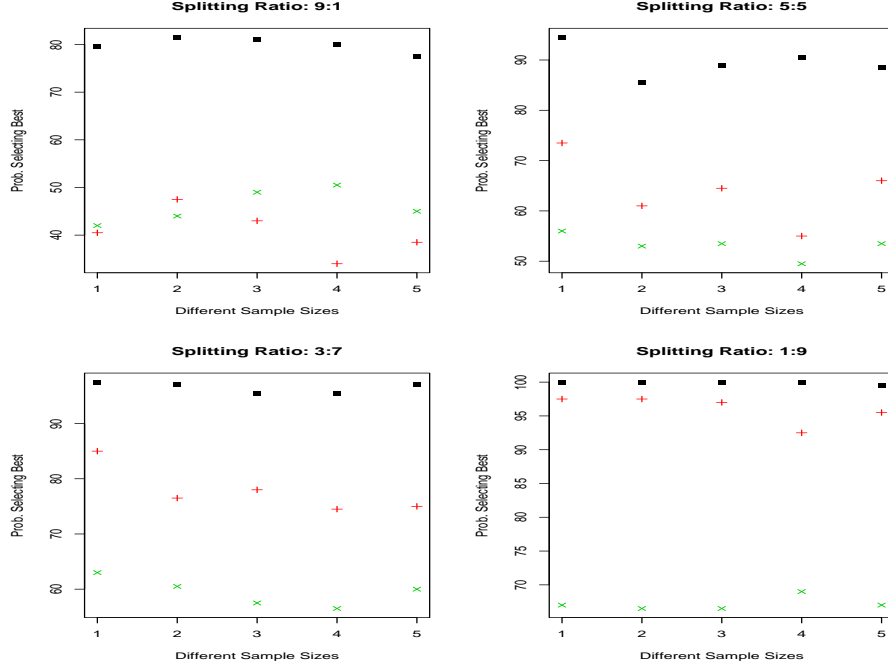


FIG 1. Probability of Selecting the Best Estimator for Case 1. Coding: ■ = Repeated Learning-testing, + = Repeated Sampling with Voting, × = Single Splitting. The Sample Sizes Are 100, 200, 400, 800, 1600.

X_i and are iid normal with mean zero and standard deviation $\sigma = 0.3$. The true function is taken to be one of the following three functions:

$$f_1(x) = 1 + x \quad (\text{Case 1})$$

$$f_2(x) = 1 + x + 0.7(x - 0.5)^2 \quad (\text{Case 2})$$

$$f_3(x) = 1 + x - \exp(-200(x - 0.25)^2) \quad (\text{Case 3}).$$

In all these cases, for σ not too large, a linear trend in the scatter plot is more or less obvious, but when σ is not small, it is usually not completely clear whether the true function is simply linear or not. We consider

three regression methods, simple linear regression, quadratic regression, and smoothing spline. The simulation was conducted using R , where a smoothing spline method is provided. We take the default choice of GCV for smoothing parameter selection.

Clearly, in Case 1, the linear regression is the winner; the quadratic regression is the right one in Case 2 (when the sample size is reasonably large); and the smoothing spline method is the winner in Case 3. Case 1 is the most difficult in comparing the methods because all of the three estimators converge to the true regression function with two converging at the same parametric rate. Case 2 is easier, where the simple linear estimator no longer converges and consequently the task is basically to compare the parametric and the nonparametric estimators (note that the spline estimator converges at a slower order). Case 3 is the easiest, where only the smoothing spline estimator is converging.

We consider the following versions of CV: 1) single splitting CV: the data splitting is done just once; 2) repeated learning-testing (a version of CV-a), denoted as RLT: we randomly split the data into the estimation and evaluation parts 100 times and average the prediction errors over the splittings for each estimator; 3) repeated splittings with voting (a version of CV-v), denoted as RSV: differently from the previous one, we select the best method based on each data splitting and then vote to decide the overall winner. The sample sizes considered are 100, 200, 400, 800 and 1600. In Case 2, the first one or two sample sizes are not considered for the splitting ratios 3 : 7 and 1 : 9 because the smoothing spline method has difficulty in parameter estimation due to the small sample size in the estimation part. In Case 3,

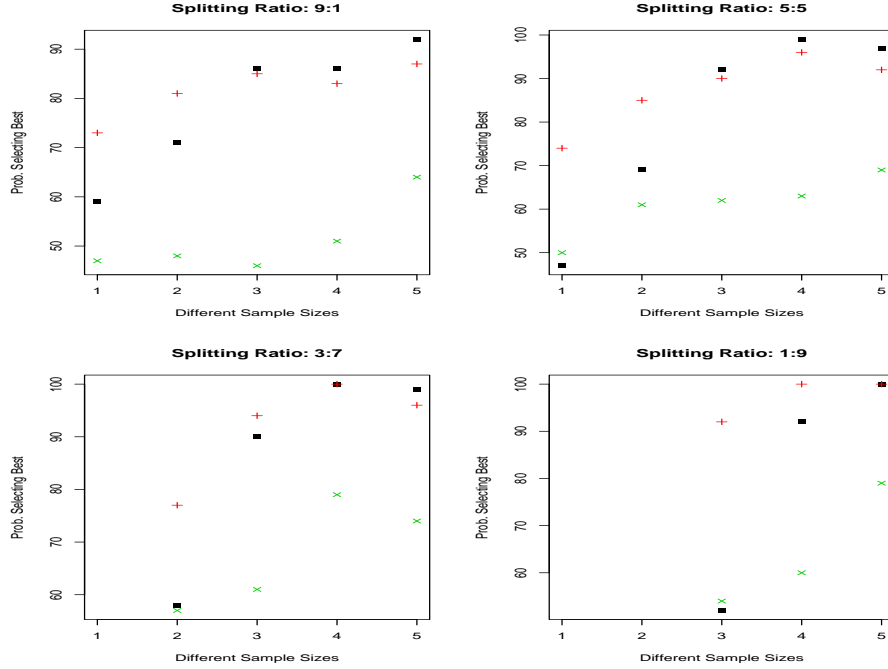


FIG 2. Probability of Selecting the Best Estimator for Case 2. Coding: ■ = Repeated Learning-testing, + = Repeated Sampling with Voting, × = Single Splitting. The Sample Sizes Are 100, 200, 400, 800, 1600.

only the splitting ratios 9 : 1 and 5 : 5 are included because lower splitting ratios make the CV methods perform perfectly. Note also that the first two sample sizes are not included for the splitting ratio 5 : 5 for the same reason as mentioned above.

The results based on 200 replications are presented in Figures 1-3 for the three cases.

From the graphs, we observe the following:

1. In Case 1, at a given splitting ratio, the increase of sample size does not lead to improvement on correct identification of the best estimator. This

nicely matches what is expected from Theorem 1: since the simple linear and the quadratic regression estimators both converge at the parametric rate, no matter how large the sample size is, any fixed proportion is not sufficient for consistency in selection.

2. In Case 2, at a given splitting ratio, when the sample size is increased, the ability of CV to identify the best estimator tends to be enhanced. This is also consistent with Theorem 1: in this case, the two converging estimators converge at different rates and thus a fixed splitting ratio is sufficient to ensure consistency in selection.

3. In Case 3, even at splitting ratio 9 : 1, RLT and RSV have little difficulty in finding the best estimator. From Theorem 1, in this case, since the smoothing spline estimator is the only converging one, the splitting ratio $n_1 : n_2$ is even allowed to go to ∞ without sacrificing the property of consistency in selection.

4. The different versions of CV behave quite differently. Overall, RLT seems to be the best, while the single splitting is clearly inferior. In Case 1, for the splitting ratios that favor the estimation size, RSV did poorly. However, with higher and higher splitting ratio towards evaluation part, it improved dramatically. This is consistent with the understanding that for a sequence of consistent splitting ratios, CV with voting or CV with averaging do not differ asymptotically (although their second order behaviors may be different). In Cases 2 and 3, the CV-v actually performed similarly to or better than the CV-a (the difference is large in Case 2 when the estimation sample size is small (40, 50, 60)).

In summary, the simulation results are very much in line with the under-

standing in Section 3.

5. Concluding remarks. We have shown that under some sensible conditions on the L_2 , L_4 , and L_∞ norms of $\hat{f} - f$ for the competing estimators, with an appropriate splitting ratio of the data for cross-validation, the better model/procedure will be selected with probability converging to 1. Unlike the previous results on CV that focus on either comparing only parametric regression models or on selecting a smoothing parameter in nonparametric regression, our result can be applied generally to compare both parametric and nonparametric models/procedures.

The result also reveals some interesting behaviors of CV. Differently from the parametric model selection case, for comparing two models converging at the same nonparametric rate, it is not necessary for the evaluation size in data splitting to be dominating. Actually, the proportion of the evaluation part can even be of a smaller order than the estimation proportion without damaging the property of consistency in selection. An implication is that it may be desirable to take the characteristics of the regression estimators into consideration for data splitting, which to our knowledge has not been seriously addressed in the literature. Based on our result, for comparing two estimators with at least one being nonparametric, half-half splitting is a good choice.

Delete-one CV is a popular choice in applications. This is usually suitable for estimating the regression function. However, when one's goal is to find which method is more accurate for the data at hand, the proportion of evaluation needs to be much larger. Further research on practical methods for properly choosing the data splitting proportion can be valuable for

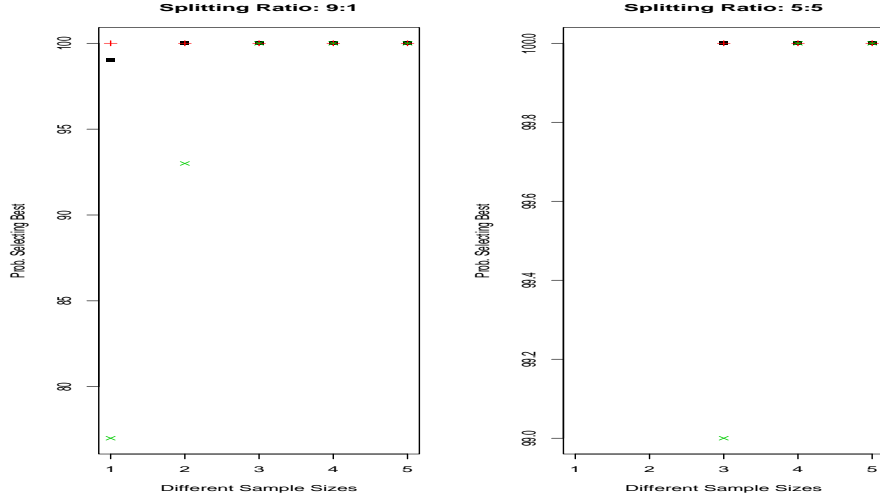


FIG 3. Probability of Selecting the Best Estimator for Case 3. Coding: ■ = Repeated Learning-testing, + = Repeated Sampling with Voting, × = Single Splitting. The Sample Sizes Are 100, 200, 400, 800, 1600.

successful applications of cross validation.

6. Proof of Theorem 1. Without loss of generality, assume that $\hat{f}_{n,1}$

is the asymptotically better estimator by Condition 2. Note that

$$\begin{aligned}
 CV(\hat{f}_{n_1,j}) &= \sum_{i=n_1+1}^n \left(f(X_i) - \hat{f}_{n_1,j}(X_i) + \varepsilon_i \right)^2 \\
 &= \sum_{i=n_1+1}^n \varepsilon_i^2 + \sum_{i=n_1+1}^n \left(f(X_i) - \hat{f}_{n_1,j}(X_i) \right)^2 + 2 \sum_{i=n_1+1}^n \varepsilon_i \left(f(X_i) - \hat{f}_{n_1,j}(X_i) \right).
 \end{aligned}$$

Define for $j = 1, 2$,

$$L_j = \sum_{i=n_1+1}^n \left(f(X_i) - \hat{f}_{n_1,j}(X_i) \right)^2 + 2 \sum_{i=n_1+1}^n \varepsilon_i \left(f(X_i) - \hat{f}_{n_1,j}(X_i) \right).$$

Then $CV(\widehat{f}_{n_1,1}) \leq CV(\widehat{f}_{n_1,2})$ is equivalent to $L_1 \leq L_2$ and thus also equivalent to

$$2 \sum_{i=n_1+1}^n \varepsilon_i \left(\widehat{f}_{n_1,2}(X_i) - \widehat{f}_{n_1,1}(X_i) \right) \leq \sum_{i=n_1+1}^n \left(f(X_i) - \widehat{f}_{n_1,2}(X_i) \right)^2 - \sum_{i=n_1+1}^n \left(f(X_i) - \widehat{f}_{n_1,1}(X_i) \right)^2.$$

Conditional on Z^1 and $X^2 = (X_{n_1+1}, \dots, X_n)$, assuming $\sum_{i=n_1+1}^n \left(f(X_i) - \widehat{f}_{n_1,2}(X_i) \right)^2$ is larger than $\sum_{i=n_1+1}^n \left(f(X_i) - \widehat{f}_{n_1,1}(X_i) \right)^2$, by Chebyshev's inequality, we have

$$\begin{aligned} & P \left(CV(\widehat{f}_{n_1,1}) > CV(\widehat{f}_{n_1,2}) \mid Z^1, X^2 \right) \\ & \leq \min \left(1, \frac{4\bar{\sigma}^2 \sum_{i=n_1+1}^n \left(\widehat{f}_{n_1,2}(X_i) - \widehat{f}_{n_1,1}(X_i) \right)^2}{\left(\sum_{i=n_1+1}^n \left(f(X_i) - \widehat{f}_{n_1,2}(X_i) \right)^2 - \sum_{i=n_1+1}^n \left(f(X_i) - \widehat{f}_{n_1,1}(X_i) \right)^2 \right)^2} \right). \end{aligned}$$

Let Q_n denote the ratio in the upper bound in the above inequality and let S_n be the event of $\sum_{i=n_1+1}^n \left(f(X_i) - \widehat{f}_{n_1,2}(X_i) \right)^2 > \sum_{i=n_1+1}^n \left(f(X_i) - \widehat{f}_{n_1,1}(X_i) \right)^2$.

It follows that

$$\begin{aligned} & P \left(CV(\widehat{f}_{n_1,1}) > CV(\widehat{f}_{n_1,2}) \right) \\ & = P \left(\left\{ CV(\widehat{f}_{n_1,1}) > CV(\widehat{f}_{n_1,2}) \right\} \cap S_n \right) + P \left(\left\{ CV(\widehat{f}_{n_1,1}) > CV(\widehat{f}_{n_1,2}) \right\} \cap S_n^c \right) \\ & \leq E \left(P \left(CV(\widehat{f}_{n_1,1}) > CV(\widehat{f}_{n_1,2}) \mid Z^1, X^2 \right) I_{S_n} \right) + P(S_n^c) \\ & \leq E \min(1, Q_n) + P(S_n^c). \end{aligned}$$

If we can show that $P(S_n^c) \rightarrow 0$ and $Q_n \rightarrow 0$ in probability as $n \rightarrow \infty$, then due to the boundedness of $\min(1, Q_n)$ (which implies that the random variables $\min(1, Q_n)$ are uniformly integrable), we have $P \left(CV(\widehat{f}_{n_1,1}) > CV(\widehat{f}_{n_1,2}) \right)$ converges to zero as $n \rightarrow \infty$. Suppose we can show that for every $\epsilon > 0$,

there exists $\alpha_\epsilon > 0$ such that when n is large enough,

$$(6.1) \quad P \left(\frac{\sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,2}(X_i))^2}{\sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,1}(X_i))^2} \geq 1 + \alpha_\epsilon \right) \geq 1 - \epsilon.$$

Then $P(S_n) \geq 1 - \epsilon$ and thus $P(S_n^c) \rightarrow 0$ as $n \rightarrow \infty$. By the triangle inequality,

$$\begin{aligned} & \sum_{i=n_1+1}^n (\hat{f}_{n_1,2}(X_i) - \hat{f}_{n_1,1}(X_i))^2 \\ & \leq 2 \sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,1}(X_i))^2 + 2 \sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,2}(X_i))^2. \end{aligned}$$

Then with probability no less than $1 - \epsilon$, Q_n is upper bounded by

$$(6.2) \quad \frac{8\bar{\sigma}^2 \left(\sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,1}(X_i))^2 + \sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,2}(X_i))^2 \right)}{\left(\left(1 - \frac{1}{1+\alpha_\epsilon}\right) \sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,2}(X_i))^2 \right)^2} \leq \frac{8\bar{\sigma}^2 \left(1 + \frac{1}{1+\alpha_\epsilon}\right)}{\left(1 - \frac{1}{1+\alpha_\epsilon}\right)^2 \sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,2}(X_i))^2}.$$

From (6.1) and (6.2), to show $P(S_n^c) \rightarrow 0$ and $Q_n \rightarrow 0$ in probability, it suffices to show (6.1) and

$$(6.3) \quad \sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,2}(X_i))^2 \rightarrow \infty \text{ in probability.}$$

Suppose a slight relaxation of Condition 1 holds: for every $\epsilon > 0$, there exists $A_{n_1,\epsilon}$ such that when n_1 is large enough, $P \left(\|f - \hat{f}_{n_1,j}\|_\infty \geq A_{n_1,\epsilon} \right) \leq \epsilon$ for $j = 1, 2$. Let H_{n_1} be the event $\left\{ \max \left(\|f - \hat{f}_{n_1,1}\|_\infty, \|f - \hat{f}_{n_1,2}\|_\infty \right) \leq A_{n_1,\epsilon} \right\}$. Then on H_{n_1} , we have $W_i = \left(f(X_i) - \hat{f}_{n_1,j}(X_i) \right)^2 - \|f - \hat{f}_{n_1,j}\|_2^2$ is bounded between $-(A_{n_1,\epsilon})^2$ and $(A_{n_1,\epsilon})^2$. Notice that conditional on Z^1 and H_{n_1} ,

$$\text{Var}_{Z^1}(W_{n_1+1}) \leq E_{Z^1} \left(f(X_{n_1+1}) - \hat{f}_{n_1,j}(X_{n_1+1}) \right)^4 = \|f - \hat{f}_{n_1,j}\|_4^4,$$

where the subscript Z^1 in Var_{Z^1} and E_{Z^1} is used to denote the conditional expectation given Z^1 . Thus conditional on Z^1 , on H_{n_1} , by Bernstein's inequality (see, e.g., Pollard (1984), page 193), for each $x > 0$, we have

$$\begin{aligned} & P_{Z^1} \left(\sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,1}(X_i))^2 - n_2 \|f - \hat{f}_{n_1,1}\|_2^2 \geq x \right) \\ & \leq \exp \left(-\frac{1}{2} \frac{x^2}{n_2 \|f - \hat{f}_{n_1,1}\|_4^4 + \frac{2(A_{n_1,\epsilon})^2 x}{3}} \right). \end{aligned}$$

Taking $x = \beta_n n_2 \|f - \hat{f}_{n_1,1}\|_2^2$, the above inequality becomes

$$\begin{aligned} & P_{Z^1} \left(\sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,1}(X_i))^2 \geq (1 + \beta_n) n_2 \|f - \hat{f}_{n_1,1}\|_2^2 \right) \\ & \leq \exp \left(-\frac{1}{2} \frac{\beta_n^2 n_2 \|f - \hat{f}_{n_1,1}\|_2^4}{\|f - \hat{f}_{n_1,1}\|_4^4 + \frac{2(A_{n_1,\epsilon})^2 \beta_n}{3} \|f - \hat{f}_{n_1,1}\|_2^2} \right). \end{aligned}$$

Under Condition 2, for every $\epsilon > 0$, there exists $\alpha'_\epsilon > 0$ such that when n is large enough, $P \left(\|f - \hat{f}_{n_1,2}\|_2^2 / \|f - \hat{f}_{n_1,1}\|_2^2 \leq 1 + \alpha'_\epsilon \right) \leq \epsilon$. Take β_n such that $1 + \beta_n = \|f - \hat{f}_{n_1,2}\|_2^2 / \left((1 + \alpha'_\epsilon/2) \|f - \hat{f}_{n_1,1}\|_2^2 \right)$. Then with probability at least $1 - \epsilon$, $\beta_n \geq (\alpha'_\epsilon/2) / (1 + \alpha'_\epsilon/2)$. Let D_n denote this event and let $S1 = \sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,1}(X_i))^2$. Then on D_n , we have

$$\beta_n \geq \alpha'_\epsilon \|f - \hat{f}_{n_1,2}\|_2^2 / \left(2(1 + \alpha'_\epsilon) (1 + \alpha'_\epsilon/2) \|f - \hat{f}_{n_1,1}\|_2^2 \right),$$

$$\begin{aligned}
 & P_{Z^1} \left(S1 \geq (1 + \beta_n) n_2 \|f - \hat{f}_{n_1,1}\|_2^2 \right) \\
 = & P_{Z^1} \left(S1 \geq \frac{n_2}{1 + \alpha'_\epsilon/2} \|f - \hat{f}_{n_1,2}\|_2^2 \right) \\
 \leq & P_{Z^1} \left(S1 \geq \left(1 + \frac{\alpha'_\epsilon \|f - \hat{f}_{n_1,2}\|_2^2}{2(1 + \alpha'_\epsilon)(1 + \alpha'_\epsilon/2) \|f - \hat{f}_{n_1,1}\|_2^2} \right) n_2 \|f - \hat{f}_{n_1,1}\|_2^2 \right) \\
 \leq & \exp \left(- \frac{(\alpha'_\epsilon)^2}{8(1 + \alpha'_\epsilon)^2 (1 + \alpha'_\epsilon/2)^2} \cdot \frac{n_2 \|f - \hat{f}_{n_1,2}\|_2^4}{\|f - \hat{f}_{n_1,1}\|_4^4 + \frac{\alpha'_\epsilon (A_{n_1,\epsilon})^2}{3(1 + \alpha'_\epsilon)(1 + \alpha'_\epsilon/2)} \|f - \hat{f}_{n_1,2}\|_2^2} \right).
 \end{aligned}$$

If we have

$$(6.4) \quad \frac{n_2 \|f - \hat{f}_{n_1,2}\|_2^4}{\|f - \hat{f}_{n_1,1}\|_4^4} \rightarrow \infty \text{ in probability,}$$

$$(6.5) \quad \frac{n_2 \|f - \hat{f}_{n_1,2}\|_2^2}{(A_{n_1,\epsilon})^2} \rightarrow \infty \text{ in probability,}$$

then the upper bound in the last inequality above converges to zero in probability. From these pieces, we can conclude that

$$(6.6) \quad P \left(\sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,1}(X_i))^2 \geq \frac{n_2}{1 + \alpha'_\epsilon/2} \|f - \hat{f}_{n_1,2}\|_2^2 \right) \leq 3\epsilon + \Delta(\epsilon, n),$$

for some $\Delta(\epsilon, n) \rightarrow 0$ as $n \rightarrow \infty$. Indeed, for every given $\epsilon > 0$, when n is

large enough,

$$\begin{aligned}
& P \left(\frac{1}{n_2} \sum_{i=n_1+1}^n \left(f(X_i) - \widehat{f}_{n_1,1}(X_i) \right)^2 \geq \frac{1}{1 + \alpha'_\epsilon/2} \|f - \widehat{f}_{n_1,2}\|_2^2 \right) \\
& \leq P(H_{n_1}^c) + P(D_n^c) \\
& \quad + P \left(H_{n_1} \cap D_n \cap \left\{ \frac{1}{n_2} \sum_{i=n_1+1}^n \left(f(X_i) - \widehat{f}_{n_1,1}(X_i) \right)^2 \geq \frac{1}{1 + \alpha'_\epsilon/2} \|f - \widehat{f}_{n_1,2}\|_2^2 \right\} \right) \\
& \leq 3\epsilon + EP \left(H_{n_1} \cap D_n \cap \left\{ \frac{1}{n_2} \sum_{i=n_1+1}^n \left(f(X_i) - \widehat{f}_{n_1,1}(X_i) \right)^2 \geq \frac{1}{1 + \alpha'_\epsilon/2} \|f - \widehat{f}_{n_1,2}\|_2^2 \right\} \middle| Z^1 \right) \\
& \leq 3\epsilon + E \exp \left(- \frac{(\alpha'_\epsilon)^2}{8(1 + \alpha'_\epsilon)^2(1 + \alpha'_\epsilon/2)^2} \cdot \frac{n_2 \|f - \widehat{f}_{n_1,2}\|_2^4}{\|f - \widehat{f}_{n_1,1}\|_4^4 + \frac{\alpha'_\epsilon(A_{n_1,\epsilon})^2}{3(1 + \alpha'_\epsilon)(1 + \alpha'_\epsilon/2)} \|f - \widehat{f}_{n_1,2}\|_2^2} \right) \\
& \triangleq 3\epsilon + \Delta(\epsilon, n),
\end{aligned}$$

where the expectation in the upper bound of the last inequality above (i.e., $\Delta(\epsilon, n)$) converges to zero due to the convergence in probability to zero of the random variables of the exponential expression and their uniform integrability (since they are bounded above by 1), provided that (6.4) and (6.5) hold. The assertion of (6.6) then follows.

For the other estimator, similarly, for $0 < \widetilde{\beta}_n < 1$, we have

$$\begin{aligned}
& P_{Z^1} \left(\sum_{i=n_1+1}^n \left(f(X_i) - \widehat{f}_{n_1,2}(X_i) \right)^2 \leq (1 - \widetilde{\beta}_n)n_2 \|f - \widehat{f}_{n_1,2}\|_2^2 \right) \\
& \leq \exp \left(- \frac{1}{2} \frac{n_2 \widetilde{\beta}_n^2 \|f - \widehat{f}_{n_1,2}\|_2^4}{\|f - \widehat{f}_{n_1,2}\|_4^4 + \frac{2(A_{n_1,\epsilon})^2 \widetilde{\beta}_n}{3} \|f - \widehat{f}_{n_1,2}\|_2^2} \right).
\end{aligned}$$

If we have

$$(6.7) \quad \frac{n_2 \widetilde{\beta}_n^2 \|f - \widehat{f}_{n_1,2}\|_2^4}{\|f - \widehat{f}_{n_1,2}\|_4^4} \rightarrow \infty \text{ in probability,}$$

$$(6.8) \quad \frac{n_2 \tilde{\beta}_n \|f - \hat{f}_{n_1,2}\|_2^2}{(A_{n_1,\epsilon})^2} \rightarrow \infty \text{ in probability,}$$

then following a similar argument used for $\hat{f}_{n_1,1}$, we have

$$(6.9) \quad P \left(\sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,2}(X_i))^2 \leq (1 - \tilde{\beta}_n) n_2 \|f - \hat{f}_{n_1,2}\|_2^2 \right) \rightarrow 0.$$

From this, if $n_2 \|f - \hat{f}_{n_1,2}\|_2^2 \rightarrow \infty$ in probability and $\tilde{\beta}_n$ is bounded away from 1, then (6.3) holds. If in addition, we can choose $\tilde{\beta}_n \rightarrow 0$, then for each given ϵ , we have $(1 - \tilde{\beta}_n) \|f - \hat{f}_{n_1,2}\|_2^2 > \frac{1+\alpha_\epsilon}{(1+\alpha_\epsilon/2)} \|f - \hat{f}_{n_1,2}\|_2^2$ for some small $\alpha_\epsilon > 0$ when n_1 is large enough. Now for every $\tilde{\epsilon} > 0$, we can find $\epsilon > 0$ such that $3\epsilon \leq \tilde{\epsilon}/3$ and there exists an integer n_0 such that when $n \geq n_0$ the probability in (6.9) is upper bounded by $\tilde{\epsilon}/3$ and $\Delta(\epsilon, n) \leq \tilde{\epsilon}/3$. Consequently when $n \geq n_0$,

$$P \left(\frac{\left(\sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,2}(X_i))^2 \right)^2}{\left(\sum_{i=n_1+1}^n (f(X_i) - \hat{f}_{n_1,1}(X_i))^2 \right)^2} \geq 1 + \alpha_\epsilon \right) \geq 1 - \tilde{\epsilon}.$$

Recall that we needed the conditions (6.4), (6.5), (6.7), and (6.8) for (6.1) to hold. Under Condition 3, $n_2 \|f - \hat{f}_{n_1,2}\|_2^4 / \|f - \hat{f}_{n_1,1}\|_4^4$ is lower bounded in order in probability by $n_2 \|f - \hat{f}_{n_1,2}\|_2^4 / \left(M_{n_1}^4 \|f - \hat{f}_{n_1,1}\|_2^4 \right)$. From all above, since $\hat{f}_{n_1,1}$ and $\hat{f}_{n_1,2}$ converge exactly at rates p_n and q_n respectively under the L_2 loss, we know that for the conclusion of Theorem 1 to hold, it suffices to have the requirements: for every $\epsilon > 0$, for some $\tilde{\beta}_n \rightarrow 0$, we have that each of $n_2 \tilde{\beta}_n^2 M_{n_1}^{-4}$, $n_2 (q_{n_1}/p_{n_1})^4 M_{n_1}^{-4}$, $n_2 \tilde{\beta}_n q_{n_1}^2 (A_{n_1,\epsilon})^{-2}$, $n_2 q_{n_1}^2 (A_{n_1,\epsilon})^{-2}$, $n_2 q_{n_1}^2$ goes to infinity.

Under Condition 1, for every $\epsilon > 0$, there exists a constant $B_\epsilon > 0$ such that $P \left(\|f - \hat{f}_{n_1,j}\|_\infty \geq B_\epsilon A_{n_1} \right) \leq \epsilon$ when n_1 is large enough. That is, for a

given $\epsilon > 0$, we can take $A_{n_1, \epsilon} = O(A_{n_1})$. Therefore if we have $n_2 M_{n_1}^{-4} \rightarrow \infty$ and $n_2 q_{n_1}^2 / (1 + A_{n_1}) \rightarrow \infty$, then we can find $\tilde{\beta}_n \rightarrow 0$ such that the 5 requirements in the previous paragraph are all satisfied. This completes the proof of Theorem 1.

7. Acknowledgment. Seymour Geisser passed away in March 2004. As is well known, he was one of the pioneers of the widely applicable cross-validation methods. This paper is in memory of his great contributions to statistics and to the School of Statistics at the University of Minnesota. The author thanks a referee, the associate editor, and the co-editor for helpful comments. This research was supported by US NSF CAREER grant 0094323.

References.

- [1] Akaike, H. (1973) Information theory and an extension of the maximum likelihood principle. In *Proc. 2nd Int. Symp. Info. Theory*, 267-281, eds. B.N. Petrov and F. Csaki, Akademia Kiado, Budapest.
- [2] Allen, D.M. (1974) The relationship between variable selection and data augmentation and a method for prediction. *Technometrics*, 16, 125-127.
- [3] Breiman, L. Friedman, J.H., Olshen, R.A., Stone, C.J. (1984) *Classification and Regression Trees*. Wadsworth, Belmont, CA.
- [4] Burman, P. (1989) A comparative study of ordinary cross-validation, ν -fold cross-validation and the repeated learning-testing methods, *Biometrika*, 76, 503-514.

- [5] Burman, P. (1990) Estimation of optimal transformations using v -fold cross validation and repeated learning-testing methods, *Sankhya*, A, 52, 314-345.
- [6] Craven, P. and Wahba, G. (1979) Smoothing noisy data with spline functions, *Numerische Mathematik*, 31, 377-403.
- [7] Donoho, D.L. and Johnstone, I.M. (1998) Minimax estimation via wavelet shrinkage, *Ann. Statist.*, 26, 879-921.
- [8] Fan, J. and Gijbels, I. (1996) *Local Polynomial Modelling and Its Applications*, London: Chapman and Hall.
- [9] Geisser, S. (1975) The predictive sample reuse method with applications, *Journal of the American Statistical Association*, 70, 320-328.
- [10] Györfi, L., Kohler, M., Krzyzak, A., Walk, H. (2002) *A Distribution-Free Theory of Nonparametric Regression*. Springer-Verlag, New York.
- [11] Hall, P. and Johnstone, I. (1992) Empirical functionals and efficient smoothing parameter selection (with discussion), *JRSS*, B, 54, 475-509.
- [12] Härdle, W., Hall, P., and Marron, J.S. (1988) How far are automatically chosen regression smoothing parameters from their optimum? (with discussion), *JASA*, 83, 86-95.
- [13] Hart, J.D. (1997) *Nonparametric Smoothing and Lack-of-fit Tests*. Springer-Verlag Inc, New York.
- [14] Li, K.-C. (1984) Consistency for cross-validated nearest neighbor estimates in nonparametric regression, *The Annals of Statistics*, 12, 230-240.

- [15] Li, K.-C. (1987) Asymptotic optimality for C_p , C_L , cross-validation and generalized cross-validation: Discrete index set, *The Annals of Statistics*, 15, 958-975.
- [16] Nemirovski, A. Topics in non-parametric statistics. *Lectures on probability theory and statistics* (Saint-Flour, 1998), 85–277, *Lecture Notes in Math.*, 1738, Springer, Berlin, 2000.
- [17] Opsomer, J., Wang, Y., and Yang, Y. (2001) Nonparametric regression with correlated errors, *Statistical Science*, 16 (2), 134-153.
- [18] Pollard, D. (1984) *Convergence of Stochastic Processes*. Springer, New York.
- [19] Shao, J. (1993) Linear model selection by cross-validation, *JASA*, 88, 486-494
- [20] Shao, J. (1997) An asymptotic theory for linear model selection (with discussion), *Statistica Sinica*, 7, 221-242
- [21] Simonoff, J.S. (1996) *Smoothing methods in statistics*. Springer-Verlag, New York.
- [22] Speckman, P. (1985) Spline smoothing and optimal rates of convergence in nonparametric regression models, *Annals of Statistics*, 13, 970-983.
- [23] Stone, C.J. (1980) Optimal rates of convergence for nonparametric estimators, *Annals of Statistics*, 8, 1348-1360.
- [24] Stone, C.J. (1982) Optimal global rates of convergence for nonparametric regression, *Annals of Statistics*, 10, 1040-1053.

- [25] Stone, M. (1974) Cross-validation choice and assessment of statistical predictions. *JRSS*, B, 36, 111-147.
- [26] van der Laan, M.J. and Dudoit, S. (2003) Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: finite sample oracle inequalities and examples. U.C. Berkeley Division of Biostatistics Working Paper Series (<http://www.bepress.com/ucbbiostat/paper130>).
- [27] van der Laan, M.J., Dudoit, S., van der Vaart, A.W. (2006) The cross-validated adaptive epsilon-net estimator, to appear in *Statistics and Decisions*.
- [28] van der Vaart, A.W., Dudoit, S., van der Laan, M.J. (2006) Oracle inequalities for multi-fold cross-validation, to appear in *Statistics and Decisions*.
- [29] Wahba, G. (1990) *Spline Models for Observational Data*, CBMS-NSF Regional Conference Series in Applied Mathematics 59, SIAM (Philadelphia, PA).
- [30] Wegkamp, M. (2003) Model selection in nonparametric regression, *Annals of Statistics*, 31, 252-273.
- [31] Weisberg, S. (2005) *Applied Linear Regression*, Hoboken NJ: John Wiley.
- [32] Wong, W.H. (1983) On the consistency of cross-validation in kernel nonparametric regression, *Annals of Statistics*, 11, 1136-1141.

- [33] Yang, Y. (2001) Adaptive regression by mixing. *JASA*, 96, 574-588.
- [34] Yang, Y. (2003) Regression with multiple candidate models: selecting or mixing? *Statistica Sinica*, 13, 783-809.
- [35] Zhang, P. (1993) Model selection via multifold cross validation, *Annals of Statistics*, 21, 299-313.

313 FORD HALL
SCHOOL OF STATISTICS
224 CHURCH STREET
MINNEAPOLIS, MN 55455