Selecting Local Models in Multiple Regression

by Maximizing Power

Chad M. Schafer and Kjell A. Doksum*

October 26, 2006

Abstract

This paper considers multiple regression procedures for analyzing the relationship between a response variable and a vector of d covariates in a nonparametric setting where both tuning parameters and the number of covariates need to be selected. We introduce an approach which handles the dilemma that with high dimensional data the sparsity of data in regions of the sample space makes estimation of nonparametric curves and surfaces virtually impossible. This is accomplished by abandoning the goal of trying to estimate true underlying curves and instead estimating measures of dependence that can determine important relationships between variables. These dependence measures are based on local parametric fits on subsets of the covariate space that vary in both dimension and size within each dimension. The subset which maximizes a signal to noise ratio is chosen, where the signal is a local estimate of a dependence parameter which depends on the subset dimension and size, and the noise is an estimate of the standard error (SE) of the estimated signal. This approach of choosing the window size to maximize a signal to noise ratio lifts the curse of dimensionality because for regions with sparsity of data the SE is very large. It corresponds to asymptotically maximizing the probability of correctly finding non-spurious relationships between covariates and a response or, more precisely, maximizing asymptotic power among a class of asymptotic level α t -tests indexed by subsets of the covariate space. Subsets that achieve this goal are called features. We investigate the properties

*Chad M. Schafer is Visiting Assistant Professor, Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213 (email: cschafer@stat.cmu.edu). Kjell A. Doksum is Professor, Department of Statistics, University of Wisconsin, Madison, WI 53607 (email: doksum@stat.wisc.edu). This research was partially supported by NSF grants DMS-0505651, DMS-0604931, and DMS-0240019. The authors are grateful to Alex Samarov for many helpful comments.

of specific procedures based on the preceding ideas using asymptotic theory and Monte Carlo simulations and find that within a selected dimension, the volume of the optimally selected subset does not tend to zero as $n \rightarrow \infty$ unless the volume of the subset of the covariate space where the response depends on the covariate vector tends to zero.

Keywords: Testing; Efficacy; Signal to noise; Curse of dimensionality; Local linear regression; Variable selection; Bandwidth selection.

1 Introduction

In this paper, we analyze the relationship between a response variable Y and a vector of covariates $\mathbf{X} = (X_1, X_2, \dots, X_d)^T$ by using “signal to noise” descriptive measures of dependencies between Y and $\mathbf{X}$. This approach addresses a dilemma of nonparametric statistical analysis (the curse of dimensionality), that with high dimensional data, sparsity of data in large regions of the sample space makes estimation of nonparametric curves and surfaces virtually impossible. On the other hand, if the goal is to find regions of strong dependence between variables, parametric methods may provide evidence for such relationships. This paper develops hybrid methods between parametric and nonparametric procedures that are designed to lift the curse of dimensionality by uniting the goal of analyzing dependencies between Y and the X ’s with the goal of finding a good model.

These methods are based on measures of dependencies between variables rather than on the estimation of some “true” curve or surface. Such measures can, for instance, be based on the currently available procedures that are defined in terms of tuning parameters. However, instead of asking what value of the tuning parameter will bring us closest to the “truth”, we ask what value of the tuning parameter will give us the best chance of finding a relationship between the variables, if it exists. For instance, consider the case where the tuning parameter is the window size $\mathbf{h} = (h_1, \dots, h_d)^T$ of local regions in $\mathbb{R}^d$ within which we do local parametric fits of

Y to $\mathbf{X}$. Our approach consists of finding the “best” window size $\mathbf{h}$ by maximizing a local signal to noise ratio where the signal is an estimate of a measure of dependence and the noise is the estimated standard error (SE) of that estimate. By dividing by the noise we lift the curse of dimensionality because the noise will be very large if the tuning parameter is such that the local region in the sample space is nearly empty.

Here we focus on the problem of exploring the relationship between the response Y and a *covariate of interest* X_1 while controlling for covariates $X_2, X_3, \dots, X_d$. Applied research abounds with situations where one is interested in the relationship between two variables while controlling for many other factors. The usual approach to dealing with the curse of dimensionality is to assume either a simple parametric relationship (usually linear) or a nonparametric but additive relationship between the covariates and response. Our approach allows for nonparametric modeling of the interaction effects and hence can find more subtle patterns in the data.

Our analysis methods begin with a *reduction step*, a procedure to analyze the multivariate data in a family of subsets of the covariate space, subsets that vary in both dimension and size within each dimension. This step is motivated by the question “Over which subsets of the covariates do these data provide evidence of a significant relationship between the response and the covariate of interest?” Once these subsets (called *features*) are extracted from the data, the analysis can go in different directions depending on the objective. Inspection of the features is itself a useful exploratory data analysis tool. Scale space analysis is possible by varying the minimum feature size.

The idea of using the signal to noise to choose a good procedure is motivated by the result (Pitman (1948), Pitman (1979), Serfling (1980), Lehmann (1999)) that the asymptotic power of asymptotically normal test statistics T_n for testing $H_0: \theta = \theta_0$ versus $H_1: \theta > \theta_0$ when θ is contiguous to θ_0 can be compared by considering their

efficacies, defined as

$$\text{EFF}_0(T_n) = \frac{\frac{\partial}{\partial \theta} E_\theta(T_n)|_{\theta=\theta_0}}{\text{SD}_{\theta_0}(T_n)}, \quad (1)$$

or the pre-limit version

$$\text{EFF}_\theta(T_n) = \frac{E_\theta(T_n) - E_{\theta_0}(T_n)}{\text{SD}_\theta(T_n)}. \quad (2)$$

See Doksum and Schafer (2006), who considered the single covariate case and selected subsets using estimates of the nonparametric efficacy

$$\text{EFF}(T_n) = \frac{E_P(T_n) - E_{H_0}(T_n)}{\text{SD}(T_n)} \quad (3)$$

where H_0 is a specified independence property, P is the probability distribution of $(\mathbf{X}, Y)$, and $\text{SD}(T_n)$ is either $\text{SD}_P(T_n)$ or $\text{SD}_{H_0}(T_n)$. Let $\text{SE}(T_n)$ denote a consistent estimator for $\text{SD}(T_n)$ in the sense that $\text{SE}(T_n)/\text{SD}(T_n) \xrightarrow{P} 1$. We refer to estimates

$$\widehat{\text{EFF}} = \frac{T_n - \widehat{E}_{H_0}(T_n)}{\text{SE}(T_n)} \quad (4)$$

of $\text{EFF}(T_n)$ as a *signal to noise ratio* or *t-statistic*. If $T_n \xrightarrow{P} T(P)$ for some functional $T(P) \equiv \theta$, then $\widehat{\text{EFF}}$ resembles the familiar Wald test statistic

$$\widehat{t}_W = \frac{\widehat{\theta}_n - \theta_0}{\text{SE}(\widehat{\theta}_n)}. \quad (5)$$

By Slutsky's Theorem, $\widehat{\text{EFF}} \xrightarrow{\mathcal{L}} N(0, 1)$ in general, as is the case of $\widehat{t}_W$ for parametric models. We could refer to “ $\widehat{\text{EFF}} \xrightarrow{\mathcal{L}} N(0, 1)$ ” as the *Wald-Slutsky phenomenon*. In some frameworks, if we consider $(\widehat{\text{EFF}})^2$, this is related to the Wilks phenomenon (Fan, Zhang, and Zhang (2001)). By not squaring $\widehat{\text{EFF}}$ we have a much simpler bias problem that makes it possible to address new optimality questions.

Most work in the nonparametric testing literature starts with a class of tests depending on a bandwidth $\mathbf{h}$ (or other tuning parameter) which tends to zero and then finds the fastest possible rate that an alternative can converge to the null hypothesis and still be detectable by one of the tests in the given class. These rates and alternatives depend on $\mathbf{h}$ and the sample size n . See, for instance, the development in

Section 3.2 of Fan, Zhang, and Zhang (2001). We consider alternatives *not* depending on $\mathbf{h}$ and ask what $\mathbf{h}$ maximizes the probability of detecting the alternative. We find that this $\mathbf{h}$ does not tend to $\mathbf{0}$ unless the volume of the set on which $E(Y|\mathbf{X})$ is nonconstant tends to zero as $n \rightarrow \infty$.

This paper is organized as follows. Section 2 defines and motivates our signal to noise criterion and Section 3 develops asymptotic properties. Section 4 describes how we incorporate variable selection into the procedure for bandwidth selection. Section 5 explains how critical values can be approximated via simulating the distribution of the criterion under random permutations of the data. Section 6 shows results from Monte Carlo simulations and analysis of real data.

2 Local Efficacy in Nonparametric Regression

We consider $(\mathbf{X}_1, Y_1), (\mathbf{X}_2, Y_2), \dots, (\mathbf{X}_n, Y_n)$ i.i.d. as $(\mathbf{X}, Y) \sim P$, $\mathbf{X} \in \mathbb{R}^d$, $Y \in \mathbb{R}$, and write $Y = \mu(\mathbf{X}) + \epsilon$ where $\mu(\mathbf{X}) \equiv E(Y|\mathbf{X})$ is assumed to exist and where $\epsilon = Y - \mu(\mathbf{X})$. We will focus on finding dependencies between Y and the covariate of interest X_1 for $\mathbf{X}$ in a neighborhood of a given covariate vector $\mathbf{x}_0$ from a targeted unit (e.g. patient, component, DNA sequence). We want to know if a perturbation of x_1 will affect the mean response for units with covariate values near $\mathbf{x}_0$. Formally, we test $H_0^{(1)}: "X_1 \text{ is independent of } X_2, X_3, \dots, X_d, Y"$ versus $H_1^{(1)}: "\mu(\mathbf{x}) \text{ is not constant as a function of } x_1 \in \mathbb{R}."$

Our test statistics will be based on linear fits for $\mathbf{x}$ restricted to subregions of the sample space. The subregion that best highlights the dependencies between Y and covariates will be determined by maximizing signal to noise ratios. We illustrate such procedures and their properties by first considering a linear model fit locally over the neighborhood $N_{\mathbf{h}}(\mathbf{x}_0)$ of $\mathbf{x}_0$ where $\mathbf{h} = (h_1, h_2, \dots, h_d)$ are bandwidths in each of the d dimensions. The true relationship between Y and $\mathbf{X}$ is such that $E[Y|\mathbf{X} = \mathbf{x}] \equiv \mu(\mathbf{x})$

and $\text{Var}(Y | \mathbf{X} = \mathbf{x}) \equiv \sigma^2$ are unknown. The local linear model gives fitted values as a function of $\mathbf{x} \in N_{\mathbf{h}}(\mathbf{x}_0)$, denoted

$$\widehat{\mu}_L(\mathbf{x}) \equiv \widehat{\beta}_0 + \sum_{j=1}^d \widehat{\beta}_j (x_j - x_{0j}), \quad (6)$$

where $\widehat{\beta}_j = \widehat{\beta}_j(\mathbf{h})$ are sample coefficients depending on $\mathbf{x}_0$ and $\mathbf{h}$. Thus, for coefficients $\beta_j = \beta_j(\mathbf{h})$ depending on $\mathbf{x}_0$ and $\mathbf{h}$,

$$\mu_{\mathbf{h}}(\mathbf{x}) \equiv E[\widehat{\mu}_L(\mathbf{x})] = \beta_0 + \sum_{j=1}^d \beta_j (x_j - x_{0j}). \quad (7)$$

We let $\widehat{\beta}_1(\mathbf{h})$ be the test statistic T_n for testing $H_0^{(1)}$ versus $H_1^{(1)}$ and develop some of the properties of the efficacy and estimated efficacy for this test statistic. Let $\mathcal{V}(\mathbf{h}) \equiv P(\mathbf{X} \in N_{\mathbf{h}}(\mathbf{x}_0))$ and assume throughout that $0 < \mathcal{V}(\mathbf{h}) \leq 1$.

2.1 The Signal

Formally, the *signal* is $E(\widehat{\beta}_1(\mathbf{h})) = \beta_1(\mathbf{h})$ where $\boldsymbol{\beta}(\mathbf{h}) = (\beta_0(\mathbf{h}), \beta_1(\mathbf{h}), \dots, \beta_d(\mathbf{h}))^T$ is the coefficient vector of the best local linear fit to Y . That is

$$\boldsymbol{\beta}(\mathbf{h}) \equiv \arg \min \{E^{\mathbf{h}}[Y - (a + \mathbf{b}^T \mathbf{X})]^2 : a \in \mathbb{R}, \mathbf{b} \in \mathbb{R}^d\} \quad (8)$$

where $E^{\mathbf{h}}$ is expected value for the conditional distribution $P^{\mathbf{h}}$ of $(\mathbf{X}, Y)$ given $\mathbf{X} \in N_{\mathbf{h}}(\mathbf{x}_0)$. With $\text{Var}^{\mathbf{h}}$ and $\text{Cov}^{\mathbf{h}}$ also being conditional,

$$(\beta_1(\mathbf{h}), \dots, \beta_d(\mathbf{h}))^T = \Sigma_{\mathbf{X}}^{-1}(\mathbf{h}) \Sigma_{\mathbf{X}Y}(\mathbf{h}) \quad (9)$$

where $\Sigma_{\mathbf{X}}(\mathbf{h}) \equiv \text{Var}^{\mathbf{h}}(\mathbf{X})$ and $\Sigma_{\mathbf{X}Y}(\mathbf{h}) \equiv \text{Cov}^{\mathbf{h}}(\mathbf{X}, Y)$.

Similarly, $\widehat{\boldsymbol{\beta}}(\mathbf{h})$ is the local linear least squares estimator

$$\widehat{\boldsymbol{\beta}}(\mathbf{h}) = (X_{\mathbf{h}}^T X_{\mathbf{h}})^{-1} X_{\mathbf{h}}^T \mathbf{Y}_{\mathbf{h}} \quad (10)$$

where $X_{\mathbf{h}}$ is the design matrix $(X_{ij}), X_{i0} = 1, 0 \leq j \leq d, i \in \mathcal{I}$ with $\mathcal{I} = \{k : \mathbf{X}_k \in N_{\mathbf{h}}(\mathbf{x}_0)\}$, and $\mathbf{Y}_{\mathbf{h}} = \{Y_i : i \in \mathcal{I}\}$. It follows that conditionally given $\mathbb{X} \equiv (\mathbf{X}_1, \dots, \mathbf{X}_n)$,

$$\text{Var}(\widehat{\boldsymbol{\beta}}(\mathbf{h}) | \mathbb{X}) = \sigma^2 (X_{\mathbf{h}}^T X_{\mathbf{h}})^{-1}. \quad (11)$$

By writing $X_{\mathbf{h}}^T X_{\mathbf{h}} = X_{\mathbf{D}}^T W_{\mathbf{h}} X_{\mathbf{D}}$ where $W_{\mathbf{h}} \equiv \text{diag}(\mathbf{1}[\mathbf{X}_i \in N_{\mathbf{h}}(\mathbf{x}_0)])$, $0 \leq i \leq n$, with $\mathbf{1}[\mathbf{X}_0 \in N_{\mathbf{h}}(\mathbf{x}_0)] \equiv 1$, and $X_{\mathbf{D}}$ is the full design matrix $(X_{ij})_{n \times (d+1)}$, we see that by the law of large numbers,

$$\begin{aligned} n^{-1} X_{\mathbf{h}}^T X_{\mathbf{h}} &\xrightarrow{a.s.} E(X_{\mathbf{D}} X_{\mathbf{D}}^T \mathbf{1}[\mathbf{X} \in N_{\mathbf{h}}(\mathbf{x}_0)]) = \mathcal{V}(\mathbf{h}) E(X_{\mathbf{D}} X_{\mathbf{D}}^T | \mathbf{X} \in N_{\mathbf{h}}(\mathbf{x}_0)) \\ &\equiv \mathcal{V}(\mathbf{h}) \Sigma_{1\mathbf{X}}(\mathbf{h}). \end{aligned} \quad (12)$$

Because $E(\widehat{\beta}_1(\mathbf{h}) | \mathbb{X}) \xrightarrow{P} \beta_1(\mathbf{h})$ as $n \rightarrow \infty$, this shows that the estimated local signal $\widehat{\beta}_1(\mathbf{h})$ is conditionally consistent given $\mathbb{X}$ as $n \rightarrow \infty$ and $d \rightarrow \infty$ when

$$n^{-1} \sigma_{1\mathbf{X}}^{11}(\mathbf{h}) \mathcal{V}^{-1}(\mathbf{h}) \rightarrow 0, \quad (13)$$

where $\sigma_{1\mathbf{X}}^{11}(\mathbf{h})$ is defined by $(\sigma_{1\mathbf{X}}^{jk}(\mathbf{h})) = \Sigma_{1\mathbf{X}}^{-1}(\mathbf{h})$, $0 \leq j \leq d$, $0 \leq k \leq d$.

2.2 The Noise

The denominator of the efficacy, the *noise*, is the asymptotic standard deviation $\sigma_0^2(\widehat{\beta}_1(\mathbf{h}))$ of $\widehat{\beta}_1(\mathbf{h})$ under H_0 . We derive a formula and develop a consistent estimator in what follows. Let $n_{\mathbf{h}}$ be the number of $\mathbf{X}_i$ that fall in $N_{\mathbf{h}}(\mathbf{x}_0)$. Note that

$$\sqrt{n} [\widehat{\beta}(\mathbf{h}) - \beta(\mathbf{h})] = \sqrt{n} \left\{ \widehat{\Sigma}_{1\mathbf{X}}^{-1} \left[n_{\mathbf{h}}^{-1} X_{\mathbf{h}}^T \mathbf{Y}_{\mathbf{h}} - \widehat{\Sigma}_{1\mathbf{X}} \beta(\mathbf{h}) \right] \right\} \quad (14)$$

where $\widehat{\Sigma}_{1\mathbf{X}} = n_{\mathbf{h}}^{-1} (X_{\mathbf{h}}^T X_{\mathbf{h}}) \xrightarrow{P} \Sigma_{1\mathbf{X}}(\mathbf{h})$ by Equation (12) and $[n_{\mathbf{h}}/n \mathcal{V}(\mathbf{h})] \xrightarrow{P} 1$. By Slutsky's Theorem, $\sqrt{n} [\widehat{\beta}(\mathbf{h}) - \beta(\mathbf{h})]$ has the same asymptotic distribution as $\sqrt{n} \{ n_{\mathbf{h}}^{-1} \Sigma_{1\mathbf{X}}^{-1} X_{\mathbf{h}}^T \mathbf{e}_{\mathbf{h}} \}$ where $\mathbf{e}_{\mathbf{h}} \equiv \mathbf{Y}_{\mathbf{h}} - \beta^T(\mathbf{h}) X_{\mathbf{h}}$. Note that by Equation (8), $E^{\mathbf{h}}(X_{\mathbf{h}}^T \mathbf{e}_{\mathbf{h}}) = \mathbf{0}$, and with $e_i = Y_i - \beta^T(\mathbf{h}) \mathbf{X}_i$,

$$n_{\mathbf{h}}^{-1} X_{\mathbf{h}}^T \mathbf{e}_{\mathbf{h}} = (n/n_{\mathbf{h}}) \left\{ n^{-1} \sum_{i=1}^n X_{ij} e_i \mathbf{1}(\mathbf{X}_i \in N_{\mathbf{h}}(\mathbf{x}_0)) : 0 \leq j \leq d \right\}. \quad (15)$$

Thus by Slutsky's Theorem and the Central Limit Theorem

$$\sqrt{n} (n_{\mathbf{h}}^{-1} X_{\mathbf{h}}^T \mathbf{e}_{\mathbf{h}}) \xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \mathcal{V}^{-1}(\mathbf{h}) \Sigma_{\mathbf{X}\mathbf{e}}(\mathbf{h})) \quad (16)$$

where

$$\Sigma_{\mathbf{X}\mathbf{e}}(\mathbf{h}) = (E(e^2 X_j X_k | \mathbf{X} \in N_{\mathbf{h}}(\mathbf{x}_0)))_{(d+1) \times (d+1)} \quad (17)$$

with $e = Y - \beta(\mathbf{h})^T \mathbf{X}$.

We have shown the following.

Proposition 1. Suppose that $\mathbf{h}$ and d are fixed in n , that $\mathcal{V}(\mathbf{h}) > 0$, $0 < \text{Var}^{\mathbf{h}}(Y) < \infty$, and that $\Sigma_{\mathbf{1}\mathbf{X}}^{-1}(\mathbf{h})$ exists; then

$$\sqrt{n} [\hat{\beta}(\mathbf{h}) - \beta(\mathbf{h})] \xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \mathcal{V}^{-1}(\mathbf{h}) \Sigma_{\mathbf{1}\mathbf{X}}^{-1}(\mathbf{h}) \Sigma_{\mathbf{X}\mathbf{e}}(\mathbf{h}) \Sigma_{\mathbf{1}\mathbf{X}}^{-1}(\mathbf{h})). \quad (18)$$

By using Equation (15) and Liapounov's Central Limit Theorem we can allow d and $\mathbf{h}$ to depend on n when we consider the asymptotic distribution of $\sqrt{n}[\hat{\beta}_1(\mathbf{h}) - \beta_1(\mathbf{h})]$. We have shown the following.

Proposition 2. Suppose that $\Sigma_{\mathbf{1}\mathbf{X}}^{-1}(\mathbf{h})$ exists, that $[\sigma^{11}(\mathbf{h})]^2 \mathcal{V}(\mathbf{h})$ is bounded away from zero and that $0 < \text{Var}(Y) < \infty$; then as $n \rightarrow \infty$,

$$\sqrt{n} [\hat{\beta}_1(\mathbf{h}) - \beta_1(\mathbf{h})] \xrightarrow{\mathcal{L}} \mathcal{N}\left(0, \sigma^2(\hat{\beta}_1(\mathbf{h}))\right). \quad (19)$$

Because under $H_0^{(1)}$, X_1 is independent of $X_2, \dots, X_d$ and Y , if we set $\mu_L(\mathbf{x}) = \beta_0 + \sum_{j=2}^d \beta_j X_j$ and $e = Y - \mu_L(\mathbf{x})$, the asymptotic variance of $\sqrt{n}[\hat{\beta}_1(\mathbf{h}) - \beta_1(\mathbf{h})]$ is

$$\sigma_0^2(\hat{\beta}_1(\mathbf{h})) \equiv \frac{E_{H_0}^{\mathbf{h}}[e]^2}{\mathcal{V}(\mathbf{h}) \text{Var}_{H_0}^{\mathbf{h}}(X_1)}. \quad (20)$$

Now $\sigma_0(\hat{\beta}_1(\mathbf{h}))/\sqrt{n}$ is the *noise* part of the efficacy.

The sample variance $s_1^2(h_1)$ calculated using $\{X_{i1} : X_{i1} \in [x_{01} - h_1, x_{01} + h_1]\}$ is our estimate of $\sigma_1^2 = \text{Var}_{H_0}^{\mathbf{h}}(X_1)$. It is consistent whenever $nh_1 \rightarrow \infty$ as $n \rightarrow \infty$. Note that with $\mu_0(\mathbf{X}) = E_{H_0}^{\mathbf{h}}(Y|\mathbf{X})$, the null residual variance $\sigma_e^2(\mathbf{h}^{(-1)}) \equiv E_{H_0}^{\mathbf{h}}(e^2)$ is

$$\sigma_e^2(\mathbf{h}^{(-1)}) = \sigma^2 + E_{H_0}^{\mathbf{h}}[\mu_0(\mathbf{X}) - \mu_L(\mathbf{X})]^2 \equiv \sigma^2 + \sigma_L^2(\mathbf{h}^{(-1)}) \quad (21)$$

where $\mathbf{h}^{(-1)} = (h_2, \dots, h_d)$ and $\sigma_L^2(\mathbf{h}^{(-1)})$ is the contribution to the variance of the error due to lack of linear fit to $\mu_0(\mathbf{X})$ under H_0 . Let $\hat{\mu}_L^{(-1)}(\mathbf{X}) = \hat{\beta}_0 + \sum_{j=2}^d \hat{\beta}_j X_j$ be the locally linear fit based on

$$\mathcal{D}(\mathbf{h}^{(-1)}) \equiv \{(X_{ij}, Y_i) : X_{ij} \in [x_{0j} - h_j, x_{0j} + h_j], 2 \leq j \leq d, 1 \leq i \leq n\}, \quad (22)$$

then a natural estimator for σ_e^2 is

$$s_e^2(\mathbf{h}^{(-1)}) \equiv (n(\mathbf{h}^{(-1)}) - d)^{-1} \sum_{\mathcal{D}(\mathbf{h}^{(-1)})} [Y_i - \hat{\mu}_L^{(-1)}(\mathbf{X}_i)]^2, \quad (23)$$

where $n(\mathbf{h}^{(-1)})$ is the number of data points in $\mathcal{D}(\mathbf{h}^{(-1)})$. The following can be shown (see Appendix).

Lemma 1. If $\sigma^2 > 0$, $\sigma_L^2(\mathbf{h}^{(-1)}) < \infty$, then, with $\mathcal{V}(\mathbf{h}^{(-1)}) = E(n(\mathbf{h}^{(-1)}))/n$,

$$E_{H_0} \left[\sum_{\mathcal{D}(\mathbf{h}^{(-1)})} [Y_i - \hat{\mu}_L^{(-1)}(\mathbf{X}_i)]^2 \right] = [n\mathcal{V}(\mathbf{h}^{(-1)}) - d] \sigma^2 + n\mathcal{V}(\mathbf{h}^{(-1)}) \sigma_L^2(\mathbf{h}^{(-1)}). \quad (24)$$

Thus $s_e^2(\mathbf{h}^{(-1)})$ is a consistent estimator of $\sigma_e^2(\mathbf{h}^{(-1)})$ under $H_0^{(1)}$ whenever $[d/n\mathcal{V}(\mathbf{h}^{(-1)})] \rightarrow 0$ and $\text{Var}_{H_0}[s_e^2(\mathbf{h})] \rightarrow 0$ as $n \rightarrow \infty$ and $d \rightarrow \infty$. It follows that under these conditions, a consistent estimate of $\sqrt{n} \times \text{noise} \equiv \sigma_1(\mathbf{h}) \equiv \sigma_0(\hat{\beta}_1(\mathbf{h}))$ is

$$\hat{\sigma}_1(\mathbf{h}^{(-1)}) = \left\{ s_e^2(\mathbf{h}^{(-1)}) / \mathcal{V}(\mathbf{h}^{(-1)}) s_1^2(h_1) \right\}^{1/2}. \quad (25)$$

Note that because $\hat{\beta}_1 \xrightarrow{P} 0$ under $H_0^{(1)}$, we can use Slutsky's theorem to show that this estimate is asymptotically equivalent to the usual linear model estimate of $\text{SD}^h(\sqrt{n} \hat{\beta}_1(\mathbf{h}))$ based on data $(\mathbf{X}_i, Y_i)$ with $\mathbf{X}_i \in N_{\mathbf{h}}(\mathbf{x}_0)$.

2.3 The Efficacy Criterion

The usual procedures for selecting bandwidths $\mathbf{h}$ involve minimizing mean squared predication or estimation error when predicting the response Y_0 or estimating the mean $E(Y_0|\mathbf{x}_0)$ of a case with covariate vector $\mathbf{x}_0$. The disadvantage is that significant relationships may escape procedures based on $\mathbf{h}$ selected in this way because they will be based on small bandwidths which lead to large noise. Here we propose to choose the bandwidth to minimize the probability of Type II error, that is, maximize power among all t -tests with the same asymptotic significance level. We thereby maximize the probability of finding significant relationships between Y and a specified covariate. This procedure automatically selects $\mathbf{h}$ to keep the noise small.

Because $T_n = \widehat{\beta}_1(\mathbf{h})$ is asymptotically normal and $E_{H_0}(T_n) = 0$, we know (Pitman (1948), Pitman (1979), Serfling (1980), Lehmann (1999)) that for contiguous alternatives the $\mathbf{h}$ that maximizes the asymptotic power is the $\mathbf{h}$ that maximizes the absolute efficacy, where

$$n^{-1/2}\text{EFF}_1(\mathbf{h}, \mathbf{x}_0) = \frac{\beta_1(\mathbf{h})}{\sigma_1(\mathbf{h})}. \quad (26)$$

Because $\beta_1(\mathbf{h}) = \sum_{j=1}^d \sigma^{1j}(\mathbf{h})\sigma_{Yj}(\mathbf{h})$, where $\sigma_{Yj}(\mathbf{h}) = \text{Cov}^{\mathbf{h}}(X_j, Y)$, we can write

$$n^{-1/2}\text{EFF}_1(\mathbf{h}, \mathbf{x}_0) = \sum_{j=1}^d \sigma^{1j}(\mathbf{h})\sigma_{Yj}(\mathbf{h})/\sigma_1(\mathbf{h}). \quad (27)$$

The efficacy optimal $\mathbf{h}$ for testing $H_0^{(1)}$ versus $H_1^{(1)}$ is defined by

$$\mathbf{h}_1^{(0)} \equiv \left(h_{11}^{(0)}, \dots, h_{1d}^{(0)}\right)^T = \arg \max_{\mathbf{h}} \text{EFF}_1(\mathbf{h}, \mathbf{x}_0). \quad (28)$$

For fixed alternatives that do not depend on n with $\beta_1(\mathbf{h}) > 0$, this $\mathbf{h}_1^{(0)}$ will not satisfy $\min\{h_{1j}^{(0)}\} \rightarrow 0$ as $n \rightarrow \infty$ because $\sigma_1(\mathbf{h}) \rightarrow \infty$ as $\min\{h_{1j}^{(0)}\} \rightarrow 0$.

Remark 1. The definition of the efficacy optimal $\mathbf{h}$ makes sense when the relationships between Y and covariates are monotone on $N_{\mathbf{h}}(\mathbf{x}_0)$. To allow for possible relationships such as

$$Y = a + \dots + (X_j - x_{0j})^2 + \dots + \epsilon \quad (29)$$

in Section 3 we will use neighborhoods of the form $[x_{0j} - (1 - \lambda_j)h_j, x_{0j} + (1 + \lambda_j)h_j]$, where $-1 \leq \lambda_j \leq 1$, rather than the symmetric intervals $[x_{0j} - h_j, x_{0j} + h_j]$. The properties of the procedures are not changed much by the introduction of the extra tuning parameters $\lambda_1, \lambda_2, \dots, \lambda_d$. See Doksum and Schafer (2006) for the case $d = 1$.

The estimate of $n^{-1/2}\text{EFF}_1$ is $n^{-1/2}t_1(\mathbf{h}, \mathbf{x}_0)$ where $t_1(\mathbf{h}, \mathbf{x}_0)$ is the t -statistic

$$t_1(\mathbf{h}, \mathbf{x}_0) = \begin{cases} \frac{n^{1/2}(\mathbf{h}^{(-1)})\widehat{\beta}_1(\mathbf{h}^{(-1)})s_1(h_1)}{n^{1/2}s_e(\mathbf{h})} & \text{if } n(\mathbf{h}^{(-1)}) > d + 1 \\ 0 & \text{otherwise.} \end{cases} \quad (30)$$

Write d_n for d to indicate that d may depend on n . Using the results of Sections 2.1 and 2.2, we have the following.

Proposition 3. In addition to the assumptions of Proposition 2, assume that as $n \rightarrow \infty$, $[d_n/n\mathcal{V}(\mathbf{h})] \rightarrow 0$, $[\sigma^{11}(\mathbf{h})]^2\mathcal{V}(\mathbf{h})$ is bounded away from zero, and $\text{Var}_{H_0}[s_{\mathbf{e}}^2(\mathbf{h})] \rightarrow 0$; then $t_1(\mathbf{h}, \mathbf{x}_0)$ is a consistent estimate of EFF_1 , in the sense that

$$\frac{t_1(\mathbf{h}, \mathbf{x}_0)}{\text{EFF}_1(\mathbf{h}, \mathbf{x}_0)} \xrightarrow{P} 1 \text{ as } n \rightarrow \infty. \quad (31)$$

Corollary 1. Suppose $\mathcal{G}$ (a grid) denotes a finite set of vectors of the form $(\mathbf{h}, \mathbf{x}_0)$, then, under the assumptions of Proposition 3, the maximizer of $t_1(\mathbf{h}, \mathbf{x}_0)$ over $\mathcal{G}$ is asymptotically optimal in the sense that

$$\frac{\max \{t_1(\mathbf{h}, \mathbf{x}_0) : (\mathbf{h}, \mathbf{x}_0) \in \mathcal{G}\}}{\max \{\text{EFF}_1(\mathbf{h}, \mathbf{x}_0) : (\mathbf{h}, \mathbf{x}_0) \in \mathcal{G}\}} \xrightarrow{P} 1. \quad (32)$$

Remark 2. Hall and Heckman (2000) used the maximum of local t -statistics to test the global hypothesis that $\mu(\cdot)$ is monotone in the $d = 1$ case. Their estimated local regression slope is the least squares estimate based on k nearest neighbors and the maximum is over all intervals with k at least 2 and at most m . They established unbiasedness and consistency of their test rule under certain conditions.

Because the asymptotic power of the test based on $\widehat{\beta}_1(\mathbf{h})$ tends to one for all $\mathbf{h}$ with $\beta_1(\mathbf{h}) > 0$ and $|\mathbf{h}| > 0$, it is more interesting to consider Pitman contiguous alternatives $H_{1n}^{(1)}$ where $\beta_1(\mathbf{h})$ depends on n and tends to zero at a rate that ensures that the limiting power is between α and 1, where α is the significance level. That is, we limit the parameter set to the set where deciding between $H_0^{(1)}$ and the alternative is difficult. We leave out the cases where the right decision will be reached for large n regardless of the choice of $\mathbf{h}$.

Now the question becomes: For sequences of contiguous alternatives with $\beta_1(\mathbf{h}) = \beta_{1n}(\mathbf{h}) \rightarrow 0$ as $n \rightarrow \infty$, what are the properties of $h_{1n}^{(0)}$? In particular, does the efficacy optimal $h_{1n}^{(0)}$ tend to zero as $n \rightarrow \infty$? The answer depends on the alternative $H_{1n}^{(1)}$ as will be shown below.

2.4 Optimal Bandwidths for Pitman Alternatives with Fixed Support

Under $H_0^{(1)}$,

$$\beta_1(\mathbf{h}) = \sum_{j=1}^d \beta_{1j}(\mathbf{h}) \equiv \sum_{j=1}^d \sigma^{1j}(\mathbf{h}) \text{Cov}^{\mathbf{h}}(X_j, Y) = 0. \quad (33)$$

We consider sequences of contiguous Pitman alternatives with $\beta_{1j}(\mathbf{h}) \propto cn^{-1/2}$, such as

$$Y = \alpha + \gamma_n r(\mathbf{X}) + \epsilon \quad (34)$$

where $\gamma_n \equiv cn^{-1/2}$, $c \neq 0$, and $|\text{Cov}^{\mathbf{h}}(X_j, r(\mathbf{X}))| > b$, for $b > 0$. Here,

$$\text{EFF}_1(\mathbf{h}, \mathbf{x}_0) \longrightarrow \left(\frac{c\mathcal{V}(\mathbf{h}) \text{SD}^{h_1}(X_1)}{\sigma_e^2(\mathbf{h}^{(-1)})} \right) \sum_{j=1}^d \beta_{1j}(\mathbf{h}) \quad (35)$$

with $\beta_{1j}(\mathbf{h}) = \sigma^{1j}(\mathbf{h}) \text{Cov}^{\mathbf{h}}(X_j, r(\mathbf{X}))$. As in the case of fixed alternatives, the maximizer $\mathbf{h}_1^{(0)}$ does not satisfy $\min\{h_{1j}^{(0)}\} \rightarrow 0$ as $n \rightarrow \infty$ because $\text{EFF}_1(\mathbf{h}, \mathbf{x}_0) \rightarrow 0$ as $\min\{h_{1j}^{(0)}\} \rightarrow 0$. This $\mathbf{h}_1^{(0)}$ is asymptotically optimal in the sense of Corollary 1.

2.5 Comparison with Mean Squared Error

There is a large literature on selecting bandwidths by minimizing mean squared error (MSE). Here MSE can be expressed as the following.

Proposition 4. If $0 < \text{Var}(Y) < \infty$, then

$$\begin{aligned} E^{\mathbf{h}}\{[\widehat{\mu}_L(\mathbf{x}_0) - \mu(\mathbf{x}_0)]^2 \mid \mathbb{X}\} &= \frac{\sigma^2}{n_{\mathbf{h}}} + [\mu_L(\mathbf{x}_0) - \mu(\mathbf{x}_0)]^2 \\ &= \frac{\sigma^2}{n\mathcal{V}(\mathbf{h})} + [\mu_L(\mathbf{x}_0) - \mu(\mathbf{x}_0)]^2 + o_P(1/n). \end{aligned} \quad (36)$$

Proof. MSE is variance plus squared bias where the conditional squared bias is as given. The conditional variance of the local least squares estimate $\widehat{\mu}_L(\mathbf{x}_0) = \widehat{\beta}_0(\mathbf{h})$ given $\mathbb{X}$ is $\sigma^2/n_{\mathbf{h}}$ where $(n_{\mathbf{h}}/n) \xrightarrow{P} \mathcal{V}(\mathbf{h})$. $\square$

If finding significant dependencies is the goal, we prefer maximizing Equation (27) to minimizing Equation (36) because the local MSE (36), as well as its global version,

focuses on finding the $\mathbf{h}$ that makes $\widehat{\mu}_L(\mathbf{x})$ close to the true unknown curve $\mu(\mathbf{x})$. Using results of Ruppert and Wand (1994) and Fan and Gijbels (1996), page 302, we can show that under regularity conditions (the Hessian of $\mu(\mathbf{x})$ exists), the bandwidths minimizing (36) tend to zero at the rate $n^{-1/(d+4)}$. By plugging such bandwidths into Equation (27) we see that this will in many cases, make it nearly impossible to find dependencies. Various semiparametric assumptions have brought the rate of convergence to zero of the bandwidths to $n^{-1/5}$ which still makes it likely that dependencies will be missed. By using (27) we have a simple method for finding bandwidths that focus on finding dependencies for any type of alternative. However there are alternatives where $\mathbf{h} \rightarrow \mathbf{0}$ makes sense. We construct these next.

3 Bandwidths for Alternatives with Shrinking Support

We consider sequences of alternatives where the set $A_n = \{\mathbf{x}: |\mu(\mathbf{x}) - \mu_Y| > 0\}$ is a connected region whose volume tends to zero as $n \rightarrow \infty$. One such set of alternatives is given by Pitman alternatives of the form

$$K_n: Y = a + \sum_{j=1}^d \gamma_j W_j \left(\frac{X_j - x_{0j}}{\theta_j} \right) + \epsilon, \quad (37)$$

where $\mathbf{X}$ and ϵ are uncorrelated, ϵ has mean zero and variance σ^2 , and each $W_j(\cdot)$ has support $[-1, 1]$. We assume that each X_j is continuous, in which case the hypothesis holds with probability one when $\gamma_j \theta_j = 0$ for all j . We consider models where $\theta_j = \theta_j^{(n)} \rightarrow 0$ as $n \rightarrow \infty$, and γ_j may or may not depend on θ_j and n . For these alternatives the neighborhood where $E(Y|\mathbf{X})$ is non-constant shrinks to achieve a Pitman balanced model where the power converges to a limit between the level α and 1 as $n \rightarrow \infty$. Note, however, that the alternative does not depend on $\mathbf{h}$. We are in a situation where “nature” picks the neighborhood size $\theta \equiv (\theta_1, \dots, \theta_d)^T$, and the statistician picks the bandwidth $\mathbf{h}$. This is in contrast to Blyth (1993), Fan, Zhang, and Zhang (2001), and Ait-Sahalia, Bickel, and Stoker (2001) who let the alternative

depend on $\mathbf{h}$.

We next show that for fixed $\mathbf{h}$ with $|\mathbf{h}| > 0$, $\text{EFF}_1(\mathbf{h}, \mathbf{x}_0) \rightarrow 0$ as $\max_j \{\gamma_j \theta_j\} \rightarrow 0$ in model (37). First note that

$$\text{Cov}^{\mathbf{h}}(X_j, Y) = \sum_{k=1}^d \gamma_k \text{Cov}^{\mathbf{h}}\left(X_j, W_k\left(\frac{X_k - x_{0k}}{\theta_k}\right)\right). \quad (38)$$

Next note for $h_k > 0$ fixed, γ_k bounded from above, and $\theta_k \rightarrow 0$,

$$\gamma_k \text{Cov}^{\mathbf{h}}\left(X_j, W_k\left(\frac{X_k - x_{0k}}{\theta_k}\right)\right) \rightarrow 0 \quad (39)$$

because $W_k((X_k - x_{0k})/\theta_k)$ tends to zero in probability as $\theta_k \rightarrow 0$ and any random variable is uncorrelated with a constant. This heuristic can be verified (see the appendix) by a change of variable. This result is stated more precisely as follows.

Proposition 5. If $\mathbf{h}$ is fixed with $|\mathbf{h}| > 0$, then as $\max_k \{\gamma_k \theta_k\} \rightarrow 0$ in model (37),

(a) $\text{EFF}_1(\mathbf{h}, \mathbf{x}_0) = O(\max_k \{\gamma_k \theta_k\})$ and

(b) $t_1(\mathbf{h}, \mathbf{x}_0) = O_P(\max_k \{\gamma_k \theta_k\})$.

Proposition 5 shows that for model (37), fixed $\mathbf{h}$ leads to small $\text{EFF}_1(\mathbf{h}, \mathbf{x}_0)$. Thus we turn to the $h_1 \rightarrow 0$ case. If $h_1 > 0$, then observations (X_1, Y) with X_1 outside $[x_{01} - \theta_1, x_{01} + \theta_1]$ do not contribute to the estimation of $\beta_1(\mathbf{h})$. Thus choosing a smaller h_1 may be better even though a smaller h_1 leads to a larger variance for $\beta_1(\mathbf{h})$. This heuristic is made precise by the next results which provide conditions where $h_1 = \theta_1$ is the optimal choice among h_1 satisfying $h_1 \geq \theta_1$. First, define

$$m_j(W_1) \equiv \int_{-1}^1 s^j W_1(s) ds, \quad j = 0, 1, 2. \quad (40)$$

Theorem 1. Assume that X_1 is independent of $X_2, \dots, X_d$, that the density $f_1(\cdot)$ of X_1 has a bounded, continuous second derivative at x_0 , and that $f_1(x_0) > 0$. Then, in model (37), as $\theta_1 \rightarrow 0$, $h_1 \rightarrow 0$ with $h_1 \geq \theta_1$, and $\mathbf{h}^{(-1)}$ fixed with $|\mathbf{h}^{(-1)}| > 0$,

(a) $\text{Cov}^{\mathbf{h}}(X_1, Y) = \gamma_1 \theta_1^2 m_1(W_1) / 2h_1 + o\left(\frac{\gamma_1 \theta_1^2}{h_1}\right)$.

(b) If $m_1(W_1) \neq 0$ and $m_2(W_1) \neq 0$, then

$$n^{-1/2} \text{EFF}_1(\mathbf{h}, \mathbf{x}_0) \propto [\sigma^2 + \sigma_{\mathbf{L}}^2(\mathbf{h}^{(-1)})]^{1/2} \gamma_1 \theta_1^2 h_1^{-3/2} m_1(W_1) f_1^{1/2}(x_{01}) \mathcal{V}^{1/2}(\mathbf{h}^{(-1)}) \quad (41)$$

which is maximized subject to $h_1 \geq \theta_1$ by $h_1 = \theta_1$.

Theorem 2. Suppose the joint density $f(\mathbf{x})$ of $\mathbf{X}$ has a bounded continuous Hessian at $\mathbf{x}_0$, and that $f(\mathbf{x}_0) > 0$, then in model (37), as $|\mathbf{h}| \rightarrow 0$ and $|\boldsymbol{\theta}| \rightarrow 0$ with $h_1 \geq \theta_1$, conclusion (b) of Theorem 1 with $\mathcal{V}^{1/2}(\mathbf{h}^{(-1)})$ replaced by $\{2^{d-1} \prod_{j=2}^d h_j\}^{1/2}$ holds.

The case where $h_1 < \theta_1$ remains; then some of the data in the neighborhood where the alternative holds is ignored and efficacy is reduced. Under suitable conditions on W_j , $j = 1, \dots, d$, the optimal h_1 equals θ_1 ; that is, $2h_1^{(0)}$ is the length of the interval where the X_1 term in the additive model (37) is different from zero. Details of the justification can be constructed by extending the results of Doksum and Schafer (2006) to the $d > 1$ case.

Remark 3. Fan (1992), Fan (1993), and Fan and Gijbels (1996) considered minimax kernel estimates for models where $Y = \mu(X) + \epsilon$ in the $d = 1$ case and found that the least favorable distribution for estimation of $\mu(x_0)$ using asymptotic mean squared error has $Y = \mu_0(X) + \epsilon$ with

$$\mu_0(x) = \frac{1}{2}b_n^2 \left[1 - c \left(\frac{x - x_0}{b_n} \right)^2 \right]_+ \quad (42)$$

where $b_n = c_0 n^{-1/5}$, for some positive constants c and c_0 . Similar results were obtained for the white noise model by Donoho and Liu (1991a,b). Lepski and Spokoiny (1999), building on Ingster (1982), considered minimax testing using kernel estimates of $\mu(x)$ and a model with normal errors ϵ and $\text{Var}(\epsilon) \rightarrow 0$ as $n \rightarrow \infty$. Their least favorable distributions (page 345) has $\mu_0(x)$ equal to a random linear combinations of functions of the form

$$\mu_j(x) = \left(h^{1/2} \int W^2(t) dt \right)^{-1} W\left(\frac{x - t_j}{h} \right). \quad (43)$$

4 Variable Selection

Suppose we want to test $H_0^{(1)}$ that X_1 and Y are unrelated and wonder whether we should keep X_j , $j \geq 2$, in the model when we construct a t -statistic for this testing problem. In experiments where confounding variables are possibly present, it would seem unreasonable to keep or exclude X_j on the basis of power for testing $H_0^{(1)}$ because confounding could lead to dropping X_j in situations where the relationship between X_1 and Y is spurious. For this reason it would seem more reasonable to base the decision about keeping X_j on the strength of its relationship to (X_1, Y) and thus we ask whether X_j contributes to accurate prediction of the response Y conditionally given X_1 . More generally we consider dropping several variables and consider fits to Y over subsets of $\mathbb{R}^d$ that vary both in dimension and size (tuning parameter) within each dimension and we select the dimension and size which maximizes efficacy in the X_1 direction while controlling for spurious correlation.

We define a *feature* $\mathcal{S}_k$ to be a subset of the covariate space which yields the maximal absolute t -statistic for some $\mathbf{x} \in \mathcal{S}_k$. In other words, among all of the subregions over which linear models are fit, we discard those that do not maximize the absolute t -statistic for **some** covariate vector $\mathbf{x}$. The remaining regions are denoted $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_r$. Assume the features are placed in decreasing order of the value of absolute t -statistic for the variable of interest. These features are subsets of $\mathbb{R}^d$ ordered with respect to their relevance to the relationship between X_1 and Y . Define

$$\mathcal{S}'_k \equiv \mathcal{S}_k \cap (\mathcal{S}'_1 \cup \mathcal{S}'_2 \cup \dots \cup \mathcal{S}'_{k-1})^c \quad (44)$$

with $\mathcal{S}'_1 \equiv \mathcal{S}_1$. Then $\cup_k \mathcal{S}'_k = \cup_k \mathcal{S}_k$ and the $\mathcal{S}'_k$ are disjoint.

There are competing goals: We want the model to be parsimonious (not include too many covariates), especially since we are focusing on cases where d is large. But, we don't want to exclude any covariate which, if included, changes the picture of the (X_1, Y) relationship. Thus if X_1 and X_2 are closely related, as are X_2 and Y , there

will also be an apparent relationship between X_1 and Y . To avoid making false claims regarding the strength of the relationship between X_1 and Y , we tentatively include X_2 , and let the analysis show whether that weakens the (X_1, Y) relationship.

This analysis of the (X_1, Y) relationship after correcting for other X 's can be accomplished in the following way: Consider the quantity

$$\gamma(x_1, \hat{\mu}_j^{(\mathcal{S})}) \equiv E \left[\left(Y - \hat{\mu}_j^{(\mathcal{S})}(\mathbf{X}) \right)^2 \mid X_1 = x_1 \right], \quad (45)$$

the expected squared prediction error when $X_1 = x_1$, where $(\mathbf{X}, Y)$ is independent of $(\mathbf{X}_i, Y)$, $1 \leq i \leq n$, and $\hat{\mu}_j^{(\mathcal{S})}(\mathbf{x})$ stands for the model fit for model j based on a subset of covariates that are restricted to subregion $\mathcal{S}$. We seek the subset of the covariates which minimizes this criterion. If there is no relationship between X_2 and Y when $X_1 = x_1$, then X_2 will be excluded based on comparisons of $\gamma(x_1, \hat{\mu}_j^{(\mathcal{S})})$ for different models j ; this could lead to different covariates included for different values of X_1 .

This is implemented as follows. Within each $\mathcal{S}_k$, m different linear models are fit; these models differ in which covariates are included, but the notation is consistent in the sense that “model j ” always refers to the same list of covariates. Let $\hat{\mu}_{jk}(\mathbf{x})$ denote the fitted value at covariate value $\mathbf{x} \in \mathcal{S}_k$ from model j fit over $\mathcal{S}_k$ and let $\hat{\mu}_j(\mathbf{x}) \equiv \hat{\mu}_{jk}(\mathbf{x})$ if $\mathbf{x} \in \mathcal{S}'_k$. Each $\mathbf{x}$ lies in exactly one $\mathcal{S}'_k$, so this is uniquely defined.

The mean squared prediction error for a model fit is commonly approximated by the leave-one-out cross validation score based on the fitted value $\hat{\mu}_{-i}(\mathbf{X}_i)$ from the fit with $(\mathbf{X}_i, Y_i)$ removed. If $\hat{\mu}$ is a linear model fit using least squares then the cross-validation prediction error is

$$(Y_i - \hat{\mu}_{-i}(\mathbf{X}_i))^2 = \left(\frac{Y_i - \hat{\mu}(\mathbf{X}_i)}{1 - h_i} \right)^2 \quad (46)$$

where h_i is the i^{th} diagonal element of the hat matrix from the linear model.

Here, we do not fit a global linear model, but the fitted value $\hat{\mu}_j^{(\mathcal{S})}(\mathbf{X}_i)$ is the result of the fit from **some** linear model. Let h_{ij} denote the diagonal element corresponding

to $\mathbf{X}_i$ of the hat matrix from the linear model fit which gives us $\hat{\mu}_j^{(S)}(\mathbf{X}_i)$. Then define

$$\hat{\gamma}_{ij}^{(S)} \equiv \left(\frac{Y_i - \hat{\mu}_j^{(S)}(\mathbf{X}_i)}{1 - h_{ij}} \right)^2. \quad (47)$$

Note how this automatically adjusts for the differing degrees of freedom in the different models since $\sum(1 - h_{ij}) = n - d - 1$. We can now estimate $\gamma(\cdot, \hat{\mu}_j^{(S)})$ by the smooth $\hat{\gamma}(\cdot, \hat{\mu})$ of $\hat{\gamma}_{ij}^{(S)}$ versus the observed value of X_1 for each data point $\mathbf{X}_i, i = 1, 2, \dots, n$. For a given $X_1 = x_1$, the model j within $\mathcal{S}_k$ with the smallest $\hat{\gamma}(x_1, \hat{\mu}_j^{(S_k)})$ is selected. Figure 5(a) shows an example where four models are compared.

Note that we avoid the curse of dimensionality by using features $\mathcal{S}_k$ with large values of the t -statistics in direction X_1 . If there are a large number of variables, instead of considering all possible combinations of variables, we can use the backward deletion approach commonly used in connection with AIC or SBC (Schwarz's Bayesian Criteria).

To summarize, by selecting the features as described above, we are simultaneously selecting the number of variables to include and the size of candidate neighborhoods for computing t -statistics in a given direction, here X_1 . We select the neighborhoods where the t -statistics are maximized and we use conditional prediction error to select the variables in such a way that we are protected against using models that produce spurious correlations.

Remark 4. A great number of tests of model assumptions and variable selection procedures are available in a nonparametric setting, e.g. Azzalini, Bowman, and Hardle (1989), Raz (1990), Eubank and LaRiccia (1993), Hardle and Mammen (1993), Bowman and Young (1996), Hart (1997), Stute (1997), Lepski and Spokoiny (1999), Fan, Zhang, and Zhang (2001), Polzehl and Spokoiny (2002), Zhang (2003), Samarov, Spokoiny, and Vial (2005), among others. One class of tests of whether the j th variable has an effect looks at the difference between the mean $\mu(\mathbf{x})$ of Y given all the x 's and the mean $\mu_{-j}(\mathbf{x}_{-j})$ of Y given all but the j th variable. Specifically, for

some weight functions $w(\mathbf{X})$, measures of the form

$$m_j = E\{[\mu(\mathbf{X}) - \mu_{-j}(\mathbf{X}_{-j})]^2 w(\mathbf{X})\}$$

are considered (Doksum and Samarov (1995), Ait-Sahalia, Bickel, and Stoker (2001)).

Similar measures compare prediction errors of the form

$$E\{[Y - \mu(\mathbf{X})]^2 w(\mathbf{X})\} \quad \text{and} \quad E\{[Y - \mu_{-j}(\mathbf{X}_{-j})]^2 w(\mathbf{X}_{-j})\}.$$

Our efficacy measure is a version of $\mu(\mathbf{x}) - \mu_{-j}(\mathbf{x}_{-j})$ adjusted for its estimability while our measure of spurious correlation is a conditional version of the above prediction errors.

5 Critical Values

First consider testing $H_0^{(1)}: "X_1 \text{ is independent of } X_2, X_3, \dots, X_d, Y"$ against the alternative $H_1^{(1)} = H_1^{(1)}(\mathbf{x}_0)$ that $\mu(\mathbf{x})$ is not constant as a function of x_1 in a neighborhood of a given covariate vector $\mathbf{x}_0$ from a targeted unit of interest. Our procedure uses the data to select the most efficient t -statistic adjusted for spurious correlation, say $t_1(\mathbf{x}_0, (\mathbf{X}, \mathbf{Y}))$, where $(\mathbf{X}, \mathbf{Y}) = ((X_{ij})_{n \times d}, (Y_i)_{n \times 1})$ are the data. Let $\mathbf{X}_1^*$ be a random permutation of $(X_{11}, \dots, X_{n1})$ and $\mathbf{X}^* = (\mathbf{X}_1^*, \mathbf{X}_2, \dots, \mathbf{X}_d)$ where $\mathbf{X}_j = (X_{1j}, \dots, X_{nj})^T$. Then, for all $c > 0$,

$$P_{H_0}(|t_1(\mathbf{x}_0, (\mathbf{X}, \mathbf{Y}))| \leq c) = P(|t_1(\mathbf{x}_0, (\mathbf{X}^*, \mathbf{Y}))| \leq c).$$

Next, we select B independent random permutations $\mathbf{X}_1^*, \dots, \mathbf{X}_B^*$ and use the $(1 - \alpha)$ sample quantile of $|t_1(\mathbf{x}_0, (\mathbf{X}_k^*, \mathbf{Y}))|$, for $k = 1, 2, \dots, B$, as the critical value. As $B \rightarrow \infty$, this quantile converges in probability to the level α critical value. Figure 1 gives as example of the simulated distribution of $|t_1(\mathbf{x}_0, (\mathbf{X}, \mathbf{Y}))|$ under the null hypothesis for $\mathbf{x}_0 = (0.4, 0.4, 0.4)$. Here $B = 500$, and the model is described in Section 6, stated in Equation (48).

Figure 1: The simulated distribution of the maximum absolute t -statistic at $\mathbf{x}_0 = (0.4, 0.4, 0.4)$ using the model described in Section 6, Equation (48).

Note that although $t_1(\mathbf{x}_0, (\mathbf{X}, \mathbf{Y}))$ is selected using absolute values of efficacy, we can still perform valid one-sided tests by following the above procedures without the absolute values. In this case the alternative is that $\mu(\mathbf{x})$ is increasing (or decreasing) as a function of $\mathbf{x}_1$ in a neighborhood of $\mathbf{x}_0$.

Next consider testing $H_0^{(1)}$ against the alternative $H_{(1)}$ that X_1 is not independent of $X_2, \dots, X_d, Y$. If $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(g)}$ is a set of grid points we can use the sum, sum of squares, maximum, or some other norm, of

$$|t_1(\mathbf{x}^{(1)}, (\mathbf{X}, Y))|, \dots, |t_1(\mathbf{x}^{(g)}, (\mathbf{X}, Y))|$$

to test $H_0^{(1)}$ versus $H_{(1)}$. We may also instead of grid points use $\mathbf{X}_1, \dots, \mathbf{X}_n$ or a subsample thereof. Again the permutation distribution of these test statistics will provide critical values.

Finally, consider the alternative that $\mu(\mathbf{x})$ is monotone in x_1 for all $x_1 \in \mathbb{R}$. We would proceed as above without the absolute values, use a left sided test for

monotone decreasing, and a right sided test for monotone increasing, see Hall and Heckman (2000) who considered the $d = 1$ case.

6 The Analysis Pipeline; Examples

In this section we describe practical implementation of the data analysis pipeline developed in the previous sections, using both simulated and real data. The simulations are built around the following functions, plotted in Figure 2.

$$f_1(x) = 4x - 2 + 5 \exp(-64(x - 0.5)^2)$$

$$f_2(x) = 2.5x \exp(1.5 - x)$$

$$f_3(x) = 3.2x + 0.4$$

Figure 2: The three functions which will be used in the simulations.

The algorithm can be described as follows. Fix a covariate vector $\mathbf{x}_0 \in [0, 1]^3$ and a neighborhood which includes $\mathbf{x}_0$, but is not necessarily centered at $\mathbf{x}_0$. A linear model is to be fit over this region. The slope in the X_1 direction β_1 , is estimated

using least squares, giving us $\hat{\beta}_1$. The t -statistic for $\hat{\beta}_1$ follows easily. One can now imagine repeating this for all possible neighborhoods, and finding that neighborhood which results in the largest absolute t -statistic for $\hat{\beta}_1$. We then imagine doing this for all possible values of $\mathbf{x}_0$. In practice, a reasonable subset of all possible neighborhoods will be chosen, and models fit over these regions. Here, we lay down a grid of values for each of the covariates, based on evenly spaced quantiles of the data. The number of grid points can be different for each covariate; we choose to have a larger number for the covariate of interest (we use 15) than for the other covariates (we use 5). Thus, we have an interior grid made of up of $15 \times 5 \times 5 = 375$ covariate vectors. A neighborhood is formed by choosing two of these points and using them as the corners of the region. Thus, there are a total of $375 \times 374/2 = 70125$ potential neighborhoods; some will be discarded due to sparsity of data.

We consider a simple case first. Set

$$Y = f_1(X_1) + f_2(X_2) + f_3(X_3) + \epsilon \quad (48)$$

where ϵ is normal with mean zero and variance $\sigma^2 = 0.02$. A sample of size $n = 1000$ is taken. The random variables X_1, X_2, X_3 are i.i.d. $\mathcal{U}(0, 1)$. Figure 3(a) shows the raw results of the analysis. For each of the local linear models, there is one line on the plot. The value on the vertical axis gives the t -statistic for $\hat{\beta}_3$. The range of values along the horizontal axis represents the range of values of the variable of interest for that neighborhood. The shade of the line indicates the proportion of the covariate space in the other two variables covered by that neighborhood. For example, the darker lines indicate regions which, regardless of their width in the X_1 direction, cover most of the X_2 and X_3 space.

Figure 3(b) is the same as Figure 3(a) except that the only regions plotted are those that yield the maximum absolute t -statistic for some $\mathbf{x}_0$. A convenient computational aspect of this approach is that wide ranges covariate values x_0 will share the same

Figure 3: Figure 3(a) the plot of t -statistics for each neighborhood fit for the simple case of Equation (48). Figure 3(b) is the same, except it only shows those neighborhoods which maximize the absolute t -statistic for some covariate vector $\mathbf{x}_0$.

region: We can describe fixing x_0 and finding the optimal region for that x_0 , but in fact we only need to fit all of these 70125 candidate models, calculate the t -statistic for each, and then discard any fit which is not the maximum for some x_0 . In this example there are 18 regions remaining following this process. The regions represented in this plot are the previously defined *features*, $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_r$. We call this plot the “ t -statistic plot for variable of interest X_1 .”

A more useful way of plotting the remaining regions is shown in Figure 4. We will refer to this graph as the “feature plot for variable of interest X_1 .” In each plot, there is one light, horizontal line for each feature; note that they are labeled with numbers going from 1 to 18. The vertical axis gives $\hat{\beta}_1$ for that local model. Each dot is an observed data point, and the shade of the dots on one line (i.e., in one region) again represents the extent of that region in the other two covariates.

From this plot, one can pick out dependencies in the data. Consider the regions

labeled “1,” “2,” and “3.” For each of these, the range of values of X_1 in the region is approximately 0.5 to 0.7 (look at the third plot). This is capturing the steep downslope in $f_1(x)$ for x in that range. But, the dependence between X_2 and Y (characterized here by the slope β_2) varies with X_2 . Ignoring this fact when modeling Y as a function of X_1 would mask this downslope. This is seen clearly when we look at these three regions in the second plot in Figure 4. The three regions correspond to different ranges of values of X_2 : Region “1” is approximately 0.5 to 0.75, where $f_2(x)$ is starting to level off, region “2” extends up to 1.0, where $f_2(x)$ is almost flat, and region “3” goes from 0 up to 0.5, where $f_2(x)$ is the steepest.

Since the relationship between X_3 and Y is linear, there is no such pattern to be found in the last of the three plots. The light horizontal lines in this first plot show that our procedure chooses the largest possible bandwidths in the direction X_1 , that is, the response Y is modeled to be linear in this direction, as we know it should be.

Figure 5(a) shows the estimates of the expected squared prediction error $\gamma(\cdot, \hat{\mu}_j)$ for each of four models, relative to this quantity for the “null” model, the model which only uses the mean of Y within that region to predict the response. We call this plot the “CV plot.” In this case, the best choice is to use all three predictors, as evidenced by the solid black line being the lowest for all values of X_1 . Figure 5(b) (the “slope plot”) shows, for each X_1 value in the data set, the estimated slope $\hat{\beta}_1$ for the region $\mathcal{S}'_k$ in which that observation lies. In practice, if the variable selection procedure chose a simpler model for a particular subset $\mathcal{S}_k$, that model would be used in the slope plot.

6.1 Other Cases

We considered alterations to the simple simulation model described in the previous section. First, consider a case with $Y = f_1(X_1) + f_2(X_2) + \epsilon$, so that Y is not a function of X_3 , but now take (X_1, X_3) to be bivariate normal, each with mean 0.5 and SD 1,

and with correlation $\sqrt{0.5}$. X_2 is still $\mathcal{U}(0, 1)$, and independent of X_1 and X_3 . Again, $n = 1000$. Figure 6(a) shows the CV plot for this case. Note that since conditional on X_1 , X_3 and Y are independent, the variable selection procedure is indicating that X_3 could be excluded. Contrast this with the second case where (X_1, X_3) have the same bivariate normal distribution, but now $Y = f_2(X_2) + f_3(X_3) + \epsilon$; see Figure 6(b) for the CV plot. Here, the best choice is to include all three variables since excluding X_3 would lead to misleading conclusions regarding the strength of the relationship between X_1 and Y .

6.2 Analysis of Currency Exchange Data

The original motivation for this study was to address a question in the study of currency exchange rates regarding the nature of the relationship between the volume of trading and the return. Data were obtained on the Japanese Yen to U.S. dollar exchange rate for the period of January 1, 1992 to April 14, 1995, a total of 1200 trading days. The response variable used was today's log volume with three covariates: 1) today's log return, 2) yesterday's log volume, and 3) yesterday's log return. The first of these, today's log return, is set as the covariate of interest. Figure 7 shows the feature plot for this data set. One interesting result is that when today's log return is positive, the coefficient for today's log return is positive; see feature 8 at the top of the plot. And when today's log return is negative, those coefficients are mostly negative; see features 1,3,4, and so forth, on the left side of the plot. This confirms a prediction of Karpoff (1987).

When variable selection is applied, we see that there is evidence that yesterday's log return is not important; see Figure 8(a). Finally, the slope plot (Figure 8(b)) shows again how the estimated slope $\hat{\beta}_1$ for the coefficient for today's log return abruptly switches from negative to positive once that variable becomes positive. This is an important finding that may have been missed under a standard multiple regression fit

to this surface. This becomes clearer when inspecting the *level plot* shown in Figure 8(c). The plot shows one dot for each vector $\mathbf{x}$ in the data set, the horizontal axis gives today's log return, and the vertical axis is the fitted value $\hat{\mu}(x)$. The superimposed solid curve is a smooth of the plot. For comparison, also plotted (as a dashed line) is the level curve for the $d = 1$ case where the only covariate is today's log return, as in Karpoff (1987).

Figure 8(d) is the level plot from the same analysis, except now making yesterday's log volume the response, and using today's log volume as one of the explanatories. Once again, we see some evidence of the "Karpoff effect." An interesting aspect of this plot is that when all three covariates are included in the analysis, the slope on the left of the level curve is significantly smaller than the slope on the right. This seems to imply that there would be a way to exploit yesterday's log volume in an effort to predict whether today's log return will be positive or negative. But, this artifact is removed by excluding today's log volume as one of the covariates: See the level curve corresponding to "Two Covariates." Using only yesterday's information, it now does not seem possible to utilize yesterday's log volume to predict if today's log return will be positive or negative.

7 Appendix

7.1 Proof of Lemma 1

The following implies Lemma 1.

Lemma 2. Let $\text{RSS}^{\mathbf{h}} \equiv \sum^{\mathbf{h}} [Y_i - \hat{\mu}_L(\mathbf{X}_i)]^2$, then

$$E^{\mathbf{h}}(\text{RSS}^{\mathbf{h}}) = [n\mathcal{V}(\mathbf{h}) - d] \sigma^2 + n\mathcal{V}(\mathbf{h}) E^{\mathbf{h}}[\mu(\mathbf{X}) - \mu_L(\mathbf{X})]^2. \quad (49)$$

Proof. If we condition on $\mathbb{X} \equiv (\mathbf{X}_1, \dots, \mathbf{X}_n)$, then $\mathbf{X}_0 \in N_{\mathbf{h}}(\mathbf{x}_0)$ no longer are random

events. We can adapt Hastie (1987), equation (16) to find

$$E^{\mathbf{h}}(\text{RSS}^{\mathbf{h}}|\mathbb{X}) = n_{\mathbf{h}}\sigma^2 + \Sigma^{\mathbf{h}} [\mu(\mathbf{X}_i) - E(\widehat{\mu}_L(\mathbf{X}_i))]^2 - d\sigma^2. \quad (50)$$

Here $E^{\mathbf{h}}(\widehat{\mu}_L(\mathbf{X}_i)) = \mu_L(\mathbf{X}_i)$ and Lemma 2 follows by the iterated expectation theorem.

Lemma 1 is a special case with μ and μ_L computed under the null hypothesis. $\square$

7.2 Proof of Proposition 5

Proof. We compute expected values using the joint density

$$f^{\mathbf{h}}(\mathbf{x}) = f(\mathbf{x}) \mathbf{1}[\mathbf{x} \in N_{\mathbf{h}}(\mathbf{x}_0)] / \mathcal{V}(\mathbf{h}) \quad (51)$$

of $\mathbf{X}$ given $\mathbf{X} \in N_{\mathbf{h}}(\mathbf{x}_0)$. Thus, with $k = j$,

$$\begin{aligned} \text{Cov}^{\mathbf{h}}\left(X_j, W_j\left(\frac{X_j - x_{0j}}{\theta_j}\right)\right) &= \int [x_j - E^{\mathbf{h}}(X_j)] W_j\left(\frac{x_j - x_{0j}}{\theta_j}\right) f^{\mathbf{h}}(\mathbf{x}) d\mathbf{x} \\ &= \int [x_j - E^{\mathbf{h}}(X_j)] W_j\left(\frac{x_j - x_{0j}}{\theta_j}\right) f_j^{\mathbf{h}}(x_j) dx_j, \end{aligned} \quad (52)$$

where

$$f_j^{\mathbf{h}}(x_j) \equiv \int f^{\mathbf{h}}(\mathbf{x}) d\mathbf{x}_{(-j)}, \quad (53)$$

with $\mathbf{x}_{(-j)} = (x_1, \dots, x_{j-1}, x_{j+1}, \dots, x_d)^T$, is the marginal density of X_j given $\mathbf{X} \in N_{\mathbf{h}}(\mathbf{x}_0)$. The change of variables $s_j = (x_j - x_{0j})/\theta_j$ gives

$$\text{Cov}^{\mathbf{h}}\left(X_j, W_j\left(\frac{X_j - x_{0j}}{\theta_j}\right)\right) = \theta_j \int_{-h_j}^{h_j} [\mathbf{x}_{0j} - E^{\mathbf{h}}(X_j) + s_j] W_j(s_j) f_j(x_{0j} + \theta_j s_j) ds_j. \quad (54)$$

A similar argument for $k \neq j$ shows that $\gamma_k \text{Cov}^{\mathbf{h}}(X_j, Y) = O(\max_k \{\gamma_k \theta_k\})$. Now (a) follows because the terms in EFF_1 other than $\gamma_k \text{Cov}^{\mathbf{h}}(X_j, Y)$ are fixed as $\gamma_k \theta_k \rightarrow 0$. Finally, (b) follows from Proposition 3. $\square$

7.3 Proof of Theorem 1

Proof. Because of the independence of the X_1 and $X_2, \dots, X_d$,

$$n^{-1/2} \text{EFF}_1(\mathbf{h}, \mathbf{x}_0) = [\sigma^2 + \sigma_{\mathbf{L}}^2(\mathbf{h}^{(-1)})]^{1/2} \mathcal{V}^{1/2}(h_1) \mathcal{V}^{1/2}(\mathbf{h}^{(-1)}) [\text{SD}^{h_1}(X_1)]^{-1} \text{Cov}^{h_1}(X_1, Y) \quad (55)$$

where $\mathcal{V}(h_1) \equiv P(x_{01} - h_1 \leq X_1 \leq x_{01} + h_1)$. The result now follows from the proof of Theorem 2.1 in Doksum and Schafer (2006). $\square$

7.4 Proof of Theorem 2

Proof. The proof can be constructed by using the fact that for small $|\mathbf{h}|$, $X_1, \dots, X_d$ given $\mathbf{X} \in N_{\mathbf{h}}(\mathbf{x}_0)$ are approximately independent with $\mathcal{U}(x_{0j} - h_j, x_{0j} + h_j)$, $1 \leq j \leq d$, distributions. This can be seen by Taylor expanding $f(\mathbf{x})$ around $f(\mathbf{x}_0)$ and noting that

$$\begin{aligned} P(\mathbf{X} \in A \mid \mathbf{X} \in N_{\mathbf{h}}(\mathbf{x}_0)) &= \int_{A \cap N_{\mathbf{h}}(\mathbf{x}_0)} f(\mathbf{x}) d\mathbf{x} \bigg/ \int_{N_{\mathbf{h}}(\mathbf{x}_0)} f(\mathbf{x}) d\mathbf{x} \\ &= \int_{A \cap N_{\mathbf{h}}(\mathbf{x}_0)} d\mathbf{x} \bigg/ \int_{N_{\mathbf{h}}(\mathbf{x}_0)} d\mathbf{x} + o(|\mathbf{h}|^2). \end{aligned} \quad (56)$$

A similar approximation applies to moments. Now use the proof of Theorem 1 with appropriate small error terms. $\square$

References

- Ait-Sahalia, Y., Bickel, P. J., and Stoker, T. M. (2001), “Goodness-of-fit Tests for Kernel Regression with an Application to Option Implied Volatilities,” *J. of Econ.*, 105, 363–412.
- Azzalini, A., Bowman, A. W., and Hardle, W. (1989), “On the Use of Nonparametric Regression for Model Checking,” *Biometrika*, 76, 1–11.
- Blyth, S. (1993), “Optimal Kernel Weights Under a Power Criterion,” *J. Amer. Statist. Assoc.*, 88, 1284–1286.
- Bowman, A. and Young, S. (1996), “Graphical Comparison of Nonparametric Curves,” *Applied Statistics*, 45, 83–98.

- Doksum, K. and Samarov, A. (1995), “Nonparametric Estimation of Global Functionals and a Measure of the Explanatory Power of Covariates in Regression,” *Ann. Statist.*, 23, 1443–1473.
- Doksum, K. and Schafer, C. (2006), “Powerful Choices: Tuning Parameter Selection Based on Power,” in *Frontiers in Statistics*, eds. Fan, J. and Koul, H., pp. 113–141.
- Donoho, D. L. and Liu, R. C. (1991a), “Geometrizing Rates of Convergence, II,” *Ann. Statist.*, 19, 633–667.
- (1991b), “Geometrizing Rates of Convergence, III,” *Ann. Statist.*, 19, 668–701.
- Eubank, R. and LaRiccia, V. (1993), “Testing for no effect in nonparametric regression,” *Journal of Statistical Planning and Inference*, 36, 1–14.
- Fan, J. (1992), “Design-adaptive Nonparametric Regression,” *J. Amer. Statist. Assoc.*, 87, 998–1004.
- (1993), “Local Linear Regression Smoothers and Their Minimax Efficiencies,” *Ann. Statist.*, 21, 196–216.
- Fan, J., Zhang, C., and Zhang, J. (2001), “Generalized Likelihood Ratio Statistics and Wilks Phenomenon,” *Ann. Statist.*, 29, 153–193.
- Fan, J. Q. and Gijbels, I. (1996), *Local Polynomial Modelling and Its Applications*, London: Chapman & Hall.
- Hall, P. and Heckman, N. E. (2000), “Testing for Monotonicity of a Regression Mean by Calibrating for Linear Functions,” *Ann. Statist.*, 28, 20–39.
- Hardle, W. and Mammen, E. (1993), “Comparing Nonparametric Versus Parametric Regression Fits,” *Ann. Statist.*, 21, 1926–1947.
- Hart, J. (1997), *Nonparametric Smoothing and Lack-of-Fit Tests*, New York: Springer-Verlag.

- Hastie, T. (1987), “A Closer Look at the Deviance,” *Amer. Statist.*, 41, 16–20.
- Ingster, Y. I. (1982), “Minimax nonparametric detection of signals in white Gaussian noise,” *Problems in Information Transmission*, 18, 130–140.
- Karpoff, J. M. (1987), “The Relation Between Price Changes and Trading Volume: A Survey,” *J. of Fin. and Quant. Anal.*, 22, 109–126.
- Lehmann, E. (1999), *Elements of Large Sample Theory*, New York: Springer.
- Lepski, O. V. and Spokoiny, V. G. (1999), “Minimax Nonparametric Hypothesis Testing: The Case of an Inhomogeneous Alternative,” *Bernoulli*, 5, 333–358.
- Pitman, E. (1948), *Lecture Notes on Nonparametric Statistics*, Columbia Univ. Press.
- (1979), *Some Basic Theory for Statistical Inference*, London: Chapman & Hall.
- Polzehl, J. and Spokoiny, V. (2002), “Local likelihood modeling by adaptive weights smoothing,” Tech. Rep. 787, WIAS.
- Raz, J. (1990), “Testing for No Effect When Estimating a Smooth Function by Nonparametric Regression: A Randomization Approach,” *J. Amer. Statist. Assoc.*, 85, 132–138.
- Ruppert, D. and Wand, M. (1994), “Multivariate locally weighted least squares regression,” *Ann. Statist.*, 22, 1346–1370.
- Samarov, A., Spokoiny, V., and Vial, C. (2005), “Component Identification and Estimation in Nonlinear High-Dimensional Regression Models by Structural Adaptation,” *J. Amer. Statist. Assoc.*, 100, 429–445.
- Serfling, R. (1980), *Approximation Theorems of Statistics*, New York: Wiley.
- Stute, W. (1997), “Nonparametric Model Checks for Regression,” *Ann. Statist.*, 25, 613–641.

Zhang, C. M. (2003), “Adaptive Tests of Regression Functions via Multi-scale Generalized Likelihood Ratios,” *Canadian J. Statist.*, 31, 151–171.

Figure 4: The feature plot for variable X_1 , for the simple case.

(a) CV Plot

(b) Slope Plot

Figure 5: More results from the analysis of the simple model. Figure 5(a) is the CV plot, comparing competing models for the purposes of variable selection. Figure 5(b) is the slope plot for variable X_1 .

(a) CV Plot

(b) CV Plot

Figure 6: Examples from analyses of extensions of the simple model. Figure 6(a) is the CV plot for the case where X_1 and X_3 are dependent, but Y is not a function of X_3 , comparing competing models for the purposes of variable selection. Figure 6(b) is the CV plot for the case where X_1 and X_3 are dependent, but Y is a function of X_3 , not of X_1 .

Figure 7: The feature plot for today's log volume for currency exchange data.

(a) CV Plot

(b) Slope Plot

(c) Level Plot

(d) Level Plot

Figure 8: Results from the analysis of the currency exchange data: The CV plot (Figure 8(a)), the slope plot (Figure 8(b)), and the level plot (Figure 8(c)) from the analysis where the response is today's log volume. Figure 8(d) is the level plot for the using yesterday's log volume as the response.