

Dimension Reduction Using Inverse Regression and Nonparametric Factors with an Application to Stock Returns

Pian Chen

Bates White, LLC
pian.chen@bateswhite.com

and

Aaron Smith

Department of Agricultural and Resource Economics
University of California, Davis
adsmith@ucdavis.edu

Abstract: We present a nonparametric method to reduce the dimension of econometric models that contain multiple explanatory variables and multiple regression functions. To enable nonparametric regression, we reduce the explanatory variables to a few inverse regression variates. We then reduce the set of regression functions to a few nonparametric factors. Not only may these nonparametric factors represent economically meaningful objects, but they also facilitate high dimensional data visualization and parsimonious model specification. We apply this method to financial asset pricing by studying the relationship between the excess returns on 25 stock portfolios and 15 macroeconomic variables. Our method reduces the 15 explanatory variables to three inverse regression variates, and it reduces the 25 regression functions to three nonparametric factors although the first of these factors dominates. Data visualization reveals how credit market conditions and inflationary expectations affect stock returns differently depending on the underlying state of the business cycle.

Key words: curse of dimensionality, variable space, function space, data visualization, model specification, financial asset pricing

1. Introduction

The curse of dimensionality hampers econometric analysis of nonlinear models with multiple explanatory variables and multiple regression functions. To make useful inference in such settings, econometricians aim to reduce the dimension of the data to a parsimonious level while preserving relevant information. This dimension reduction is often carried out by imposing linear parametric structures, excluding variables, or aggregating data. Rather than imposing such restrictions *a priori*, we apply a data-driven method to reduce the dimension in two spaces: the variable space spanned by the explanatory variables and the function space spanned by the regression functions. We reduce the variable space to a few inverse regression variates and reduce the function space to a few nonparametric factors.

We demonstrate how to visualize high dimensional data using estimated inverse regression variates and nonparametric factors. Data visualization can often identify interesting relationships among variables, but a high dimensional system with an unknown joint distribution cannot be fully represented by arbitrary bi-variate plots (Carroll and Li 1992, Cook 1994). However, by reducing the system to a set of orthogonal nonparametric factors and a set of uncorrelated inverse regression variates, we can fully characterize a high dimensional system using a minimum number of bi-variate plots. These plots can reveal underlying structure previously hidden in high-dimensional data and guide parsimonious model specification.

Specifically, we study the nonparametric model

$$(1) \quad Y_i = F_0(X_i) + U_i \quad (i = 1, 2, \dots, n),$$

where Y_i represents a G -vector of dependent variables, X_i denotes J explanatory variables, $F_0(X_i)$ is a G -vector of unknown functions of X_i and $E(U_i | X_i) = \underline{0}$. We reduce the dimension of

the function space by decomposing $F_0(X_i)$ as

$$(2) \quad F_0(X_i) = AH(X_i),$$

where A is a $G \times L$ constant matrix and $H(X_i)$ is an L -vector of unknown functions with $L \leq G$. This decomposition implies that $F_0(X_i)$ falls into a function space spanned by L basis functions (H_1, \dots, H_L) . Following Donald (1997), we call these unknown basis functions nonparametric factors.

Equation (2) arises naturally in economics. For example, in asset pricing models, Y_i denotes a vector of excess stock returns, X_i contains variables that capture aggregate market risk, and each nonparametric basis function represents an aggregate risk factor. The capital asset pricing model (CAPM, Sharpe 1964, Lintner 1965) specifies one factor that is a linear function of the market portfolio; the arbitrage pricing model (APT, Ross 1976) specifies multiple linear factors; Fama and French (1992, 1993) specify size and book-to-market factors in addition to the market portfolio; Chen, Roll, and Ross (1986) specify factors that are linear functions of several macroeconomic variables. Equation (2) generalizes these linear parametric models by allowing the aggregate risk factors to be nonparametric functions of a potentially large set of observed X variables. Other examples of the model in (2) include demand systems and nonparametric instrumental variables (Donald 1997).

To estimate (H_1, \dots, H_L) , we need to estimate $F_0(X_i)$ nonparametrically. In applications with many X variables, such estimation is hindered by the curse of dimensionality. To overcome this problem, we use inverse regression, which is based on the premise that a small number of linear combinations of X can capture the information in X that is relevant for Y

(Li 1991). Therefore, we can replace J -dimensional X with d linear combinations of X , i.e., $F_0(X_i) \equiv F_1(X_i B)$, and redefine the nonparametric factors as unknown functions of $X_i B$, i.e.,

$$(3) \quad F_1(X_i B) = AH(X_i B),$$

where $B = (\beta_1, \dots, \beta_d)$ is a $J \times d$ constant matrix with $d \leq J$.¹ Replacement of X_i with $X_i B$ entails an orthogonal projection that transforms a large set of correlated X variables into a small set of uncorrelated inverse regression variates XB .

The matrix B cannot be uniquely identified from the data, but the space spanned by B is identifiable. Li (1991) proposes the inverse regression method, which generates root- n consistent estimates of this dimension reduction space without knowledge of the link function. Because the inverse regression estimates converge at a faster rate than standard nonparametric statistics, substituting the estimated inverse regression variates for the observed X variables does not affect the asymptotic distribution of nonparametric statistics in general. Consequently, in addition to the function space reduction problem we treat in this paper, inverse regression can be applied to many other testing and estimation problems in nonparametric econometrics. It has been applied successfully in fields such as marketing (e.g., Naik, Wedel and Kamakura 2010), bioinformatics (e.g., Li and Li 2004), and chemometrics (e.g., Amato, Antoniadis, and De Feis 2006).

The rest of the paper is organized as follows. In Section 2, we introduce inverse regression for variable space reduction and compare it with other dimension reduction methods such as principal components and least squares regression. We present a new result that, under the assumption of a linear link function, inverse regression is identical to OLS and canonical correlation. In Section 3, we discuss nonparametric regression with a large number of

¹ For notational convenience, we define X_i as the $1 \times J$ row vector that inhabits the rows of the $n \times J$ matrix X .

explanatory variables, function space reduction, and high-dimensional data visualization. In Section 4, we apply our dimension-reduction method to extract the aggregate risk factors that explain the returns on 25 stock portfolios, where each risk factor is a nonparametric function of 15 macroeconomic variables. We reduce the 15 explanatory variables to three inverse regression variates and we reduce the 25 regression functions to three nonparametric factors. Section 5 concludes the paper.

2. Variable Space Reduction Using Inverse Regression

Inverse regression is a dimension reduction method first proposed by Li (1991). In this section, we review the fundamental theory, outline the estimation procedure, and compare inverse regression to principal component analysis, linear regression, and canonical correlations.

2.1 Inverse Regression: Theory and Estimation

Write the distribution of $Y_i|X_i$ as a function d linear combinations of X_i , i.e.,

$$(4) \quad Y_i = F(X_i\beta_1, X_i\beta_2, \dots, X_i\beta_d, U_i),$$

where $d \leq J$, $B = (\beta_1, \beta_2, \dots, \beta_d)$ are J -dimensional column vectors known as “effective dimension reduction” (*e.d.r.*) directions, the link function F is unknown, X_iB are called inverse regression variates, and U_i is independent of X_i . Without further structure on F , the vectors $\beta_1, \beta_2, \dots, \beta_d$ are only identified up to an orthogonal rotation, i.e., the model would be identical if B were replaced by BC , for some $d \times d$ full rank matrix C . Inverse regression methods therefore focus on estimating the space spanned by B .

If $d=J$, then the representation in (4) places no restrictions on the functional form of the relationship between Y and X . However, in a typical case d is much less than J , and the model achieves parsimony by reducing F from a possibly nonlinear function of a large number of X variables to a possibly nonlinear function of a few inverse regression variates. If $d=1$ and F is linear, then (4) is simply a linear regression model, in which one linear combination of X is sufficient to explain Y . On the other hand, if $d=1$ and F is nonlinear, (4) reduces to the single-index model, which has been studied extensively in econometrics (see, for example, Ichimura 1993, Klein and Spady 1993, Powell, Stock and Stoker 1989). Logit and probit models can be viewed as special cases in which Y is binary and F takes certain parametric functional forms. The model in (4) generalizes these special cases by allowing multiple indexes in the explanatory variables.

Li (1991) shows how to obtain root- n consistent estimates of the space spanned by B without knowledge of the link function F . His key insight comes from the fact that Y is related only to d linear combinations of the X variables. It follows directly that estimating B requires finding the linear combinations of X that relate to Y . Li begins by first estimating nonlinear regressions of each X variable on Y . The resulting regression functions must be linearly dependent because Y is only related to a few linear combinations of X . Thus, rather than regressing Y on X (forward regression), Li estimates B by finding common factors in regressions of each X variable on Y (inverse regression).

To obtain this estimator, Li assumes that the joint distribution of X_i is such that $E(X_i b | X_i \beta_1, X_i \beta_2, \dots, X_i \beta_d)$ is linear in $X_i \beta_1, X_i \beta_2, \dots, X_i \beta_d$ for any vector $b \in R^J$. This assumption is known as the linear design condition, and it holds if the distribution of X_i is elliptically symmetric (e.g., the normal distribution). However, elliptical symmetry is not necessary for the

condition to hold. Hall and Li (1993) prove under general conditions that deviations from the linear design condition converge to zero as the dimension of X_i increases. In other words, the linear design condition holds asymptotically for high-dimensional data. Under the linear design condition and the model defined in (4), Li (1991) shows that the centered inverse regression curve $E(X_i | Y_i) - E(X_i)$ is contained in the linear subspace spanned by $\Sigma_{XX}B$.

To understand this result, consider a vector $b \in R^J$ defined such that $b' \Sigma_{XX} \beta_k = 0$ for all $k = 1, 2, \dots, d$ where Σ_{XX} denotes the covariance matrix of X_i . This vector b is orthogonal to the linear subspace spanned by $\Sigma_{XX}B$. Li shows that $E(b'X_i | Y_i) - E(b'X_i) = 0$ with probability one for all such vectors b , which implies that $b' \text{Cov}[E(X_i | Y_i)]b = 0$. It follows that the covariance matrix $\text{Cov}[E(X_i | Y_i)]$ is degenerate in any direction orthogonal to $(\Sigma_{XX}\beta_1, \Sigma_{XX}\beta_2, \dots, \Sigma_{XX}\beta_d)$. Therefore, $(\Sigma_{XX}\beta_1, \Sigma_{XX}\beta_2, \dots, \Sigma_{XX}\beta_d)$ are the eigenvectors associated with the d nonzero eigenvalues of $\text{Cov}[E(X_i | Y_i)]$. These expressions simplify slightly if the explanatory variables are standardized to $Z_i = \Sigma_{XX}^{-1/2}[X_i - E(X_i)]$. In that case, if $(\eta_1, \eta_2, \dots, \eta_d)$ are the eigenvectors associated with the d nonzero eigenvalues of $\text{Cov}[E(Z_i | Y_i)]$, then the *e.d.r.* directions are $\beta_k = \Sigma_{XX}^{-1/2} \eta_k$ ($k = 1, 2, \dots, d$).

The eigenvalue decomposition of $\text{Cov}[E(Z_i | Y_i)]$ requires an estimate of $E(Z_i | Y_i)$. For univariate Y_i , Li (1991) suggests dividing the range of Y_i into multiple “slices”. Within each slice, the mean value of Z_i provides a crude estimate of the expected value of Z given a value of Y in that slice. In other words, $E(Z_i | Y_i)$ is estimated from Nadaraya-Watson kernel regression using the naïve kernel (see Pagan and Ullah, pg. 8). The sliced inverse regression (SIR) estimates of B are thus obtained from an eigenvalue decomposition of

$$\hat{M}_{SIR} = \sum_{h=1}^H \hat{\rho}_h \hat{m}_h \hat{m}_h'$$

where $\hat{\rho}_h = n^{-1} \sum_{i=1}^n \mathbf{1}(S_{h-1} \leq Y_i \leq S_h)$, $\hat{m}_h = (\hat{\rho}_h n)^{-1} \sum_{i=1}^n \mathbf{1}(S_{h-1} \leq Y_i \leq S_h) Z_i$, $\mathbf{1}(S_{h-1} \leq Y_i \leq S_h)$ is an indicator function, and S_0, S_1, \dots, S_H denotes the cutoff points for the slices. The precision of SIR is determined by the precision of the estimate of $Cov[E(Z_i | Y_i)]$, and it is quite insensitive to the number of slices, i.e., the bandwidth. Changing the bandwidth makes one component of the covariance more precise while making another component less precise. For example, using a small number of slices allows for a large number of observations per slice and therefore generates relatively precise estimates of $E(Z_i | Y_i)$, but leaves few slices over which to estimate the covariance. At the other extreme, a large number of slices produces imprecise $E(Z_i | Y_i)$ estimates but leaves many terms with which to precisely estimate the covariance.

Chen and Li (1998) show that $\hat{\beta}_k$ from SIR is consistent and asymptotically normal with a covariance matrix that can be approximated by $n^{-1} \hat{\lambda}_k^{-1} (1 - \hat{\lambda}_k) \hat{\Sigma}_{XX}^{-1}$, where $\hat{\lambda}_k$ is the k^{th} largest eigenvalue of the matrix \hat{M}_{SIR} . The order of this covariance matrix illustrates the root- n consistency of each $\hat{\beta}_k$. For normally distributed X_i , Li (1991) derives the asymptotic distribution of the summation of the smallest $(J-k)$ eigenvalues of \hat{M}_{SIR} :

$$(5) \quad n \sum_{j=k+1}^J \hat{\lambda}_j \sim \chi^2_{((J-k)(S-k-1))}.$$

This asymptotic distribution enables estimation of the number of significant *e.d.r.* directions by testing $H_0: d=k$ v.s. $H_1: d=k+1$. Specifically, if $n \sum_{j=k+1}^J \hat{\lambda}_j$ exceeds the corresponding $\chi^2_{((J-k)(S-k-1))}$ critical value, then there are at least $k+1$ significant *e.d.r.* directions. Applied users can conduct

these tests sequentially for $k = 0, 1, \dots$, until the null hypothesis is not rejected. For non-normal X , Bura and Cook (2001a) develop a weighted chi-square test.

Slicing provides one method for estimating $E(Z_i | Y_i)$. Other proposed inverse regression estimators estimate $E(Z_i | Y_i)$ by kernel regression (Zhu and Fang 1996), nearest neighbors (Hsing 1999) and parametric regression (Bura and Cook 2001b). Methods to estimate B using the second moment of $X|Y$ have also been proposed, including Sliced Average Variance Estimate (SAVE, Cook and Weisberg 1991) and SIR-II (Li 1991). In addition, SIR- α (Li, 1991) is an attempt to combine the information from both the first and second moments. These second-moment based methods may capture additional information that is orthogonal to the first moment of $X|Y$ (Cook and Weisberg 1991).

SIR is easy to implement if the dependent variable Y_i is univariate. When Y_i is of high dimension, the slicing step of SIR becomes impractical because too many cells are empty. However, Hsing's (1999) nearest neighbor inverse regression (NNIR) proves quite robust to multi-dimensional Y_i .

To expound NNIR, let $i \in \{1, 2, \dots, n\}$ and $i^* \in \{1, 2, \dots, n\} - \{i\}$ be the indices for which

$$(6) \quad \gamma(Y_i, Y_{i^*}) = \min_{j \neq i, 1 \leq j \leq n} \gamma(Y_i, Y_j),$$

where $\gamma(\cdot, \cdot)$ is a metric such as Euclidean distance, i.e., Y_{i^*} is the nearest neighbor of Y_i . Let X_{i^*} be the concomitant of Y_{i^*} . Hsing (1999) estimates

$$(7) \quad Cov[E(Z_i | Y_i)] = 0.5E[Z_i' E(Z_i | Y_i)] + 0.5E[E(Z_i' | Y_i) Z_i]$$

using

$$(8) \quad \hat{M}_{NNIR} = (2n)^{-1} \sum_{i=1}^n (\hat{Z}_i' \hat{Z}_{i^*} + \hat{Z}_{i^*}' \hat{Z}_i),$$

where \hat{Z}_i and \hat{Z}_{i^*} are standardized X_i and X_{i^*} respectively. Although \hat{Z}_{i^*} provides a noisy

estimate of $E(Z_i|Y_i)$, the space spanned by the first d eigenvectors of \hat{M}_{NNIR} is root- n consistent for the *e.d.r.* space under general conditions. We use NNIR in our application (see Section 4).

To our knowledge, the asymptotic distribution of the d smallest eigenvalues of \hat{M}_{NNIR} has not been derived. Thus, hypothesis tests based on a statistic like (5) cannot be used to estimate the number of significant *e.d.r.* directions. However, cross-validation and permutation tests (Cook and Weisberg 1991, Cook and Yin 2001) can be useful for determining d in practice. In this paper, we use cross validation to determine d (see Section 3.3 for details).

2.2 Comparison to Other Dimension Reduction Methods

2.2.1 Principal Components Analysis

Rather than using inverse regression, it is common to apply principal component analysis (PCA) on X and keep the first few principal components for modeling the relationship between Y and X (e.g., Stock and Watson 2002). The main drawback of this approach is that PCA does not take into account the relationship between Y and X when estimating the principal components. Thus, for any two different Y variables, PCA reduces the data to the same linear combinations of X , even if the relationships between Y and X are different across the two Y variables. Consequently, PCA can misspecify the fundamental relationship between Y and X by excluding relevant components and introducing irrelevant components for Y .

A simple example illustrates this point. Consider the linear regression model $Y_i = X_i\beta + U_i$. The first principal component of X is the linear combination that explains as much as possible of the variance of X . In general, there is no correspondence between this principal

component and β . For example, suppose X contains several highly positively correlated variables measuring activity in the real economy and a single monetary variable that is weakly correlated with the real variables. PCA will be successful if the real variables have large effects on Y with coefficients of the same sign. If the single monetary variable matters most for Y , then PCA will miss this relationship because the first principal component loads heavily on the block of real variables. Similarly, if the real variables matter for Y but enter the regression with opposing signs, then the first principal component will perform poorly.

To illustrate this last scenario in a nonlinear context, we generate 200 observations using $Y_i = \sin(X_i\beta_1) + U_i$, where $U_i \sim iid N(0, 0.01^2)$, $\beta_1 = [0.25 \quad -0.35 \quad -0.57 \quad 0.70]'$, and $X_i \sim iid N(\mu, \Sigma_{XX})$, with mean vector $\mu = [5 \quad 5 \quad 5 \quad 5]$ and covariance matrix

$$\Sigma_{XX} = \begin{bmatrix} 1 & \rho & \rho & \rho \\ \rho & 1 & \rho & \rho \\ \rho & \rho & 1 & \rho \\ \rho & \rho & \rho & 1 \end{bmatrix}.$$

In this example, the sine link function is nonlinear in the four X variables, and the four X variables can be reduced to one linear component denoted by $X_i\beta_1$. We experiment with three correlation coefficients $\rho = 0, 0.5, \text{ and } 0.95$ and apply both SIR and PCA to estimate β_1 . We use ten slices to obtain the SIR estimates. In Figure 1, we plot Y against the first SIR variate in the left panel and plot Y against the first principal component in the right panel. It is evident that SIR correctly identifies the nonlinear sine pattern while PCA finds no pattern for all three correlation coefficients. When ρ increases, the X variables become highly collinear, and SIR's ability to identify the *e.d.r.* direction decreases.

2.2.2 OLS and Canonical Correlations

Under the assumption of a linear link function, inverse regression is identical to OLS and canonical correlation. In the following proposition, we formalize this statement, which generates intuition about the theory underlying inverse regression. To our knowledge this result has not been published before. For easy exposition, we standardize X_i to $Z_i = \Sigma_{XX}^{-1/2} [X_i - E(X_i)]$.

Proposition: Consider the classic linear regression model $Y_i = Z_i\beta + u_i$ and the linear inverse regression curve $E(Z_i | Y_i) = \Sigma_{ZY}\Sigma_{YY}^{-1}Y_i$, where $\Sigma_{ZY} = Cov(Z_i, Y_i)$ and $\Sigma_{YY} = Cov(Y_i)$. The coefficient matrix β may be of reduced rank.

- (i) For a univariate Y_i , the estimates of β by OLS and by inverse regression are identical in population.
- (ii) For a multivariate Y_i , the estimates of β by canonical correlation and by inverse regression are identical in population.

We only prove part (ii) because it includes part (i) as a special case. Consider the canonical correlation between $Z_i\eta$, a linear combination of Z_i , and $Y_i\alpha$, a linear combination of Y_i . For identification, we use the normalizations $\eta'\Sigma_{ZZ}\eta = \eta'\eta = 1$ and $\alpha'\Sigma_{YY}\alpha = 1$, which imply that the canonical correlation coefficients η equal the eigenvectors of $\Sigma_{ZY}\Sigma_{YY}^{-1}\Sigma_{YZ}$ (Hamilton 1994, pp. 630-635). Inverse regression calculates the eigenvectors of $Cov[E(Z_i | Y_i)]$. Linearity implies that $E(Z_i | Y_i) = \Sigma_{ZY}\Sigma_{YY}^{-1}Y_i$, so it follows that $Cov[E(Z_i | Y_i)] = \Sigma_{ZY}\Sigma_{YY}^{-1}\Sigma_{YZ}$. Because both methods perform an eigenvalue decomposition of $\Sigma_{ZY}\Sigma_{YY}^{-1}\Sigma_{YZ}$, canonical correlation and inverse regression are identical given a linear relationship between Y and Z .

Moreover, when estimated by maximum likelihood, a reduced rank regression model $Y_i = Z_i\eta\xi' + \varepsilon_i$ with a normally distributed ε_i has the first order condition $(\Sigma_{ZY}\Sigma_{YY}^{-1}\Sigma_{YZ} - \lambda^2)\eta = 0$ (see Camba-Mendez, Kapetanios, Smith, and Weale 2003). This condition is the same as the one for canonical correlation and therefore makes inverse regression identical to maximum likelihood for this case. Although all four methods use information from Y , inverse regression generalizes linear regression, canonical correlation, and reduced rank regression because it imposes no structure on link function and error term. The inverse regression curve is estimated nonparametrically and the error term need not be additive.

3. Function Space Reduction Using Nonparametric Factors

In this section, we shift our attention from the variable space to the function space. First, we discuss how to estimate regression functions after reducing the dimension of the variable space. Second, we introduce our function space reduction and visualization method. The idea of function space reduction is similar to that of variable space reduction. That is, the information in a large number of unknown regression functions may be well captured by a small number of orthogonal basis functions, known as nonparametric factors. Identifying nonparametric factors enables function space visualization, which in turn can reveal structure previously hidden in the data. It could also guide parsimonious model specification.

3.1 Nonparametric Regression after Inverse Regression

Without reducing the dimension of the variable space, nonparametric estimation of $F_0(X_i)$ is subject to the *curse of dimensionality*. For example, consider the Nadaraya-Watson

kernel estimator

$$(9) \quad \hat{F}_0(X_i) = \frac{\sum_{j=1}^n K(h^{-1}(X_j - X_i))Y_j}{\sum_{j=1}^n K(h^{-1}(X_j - X_i))},$$

where K represents a kernel weighting function and X_i is of dimension J . The optimal bandwidth $h_{opt} = O(n^{-1/(J+4)})$ results in the smallest mean squared error (MSE) of $O(n^{-4/(J+4)})$ under the regularity conditions (A2)-(A7) in Pagan and Ullah (1999, pp. 96-104). It follows that the MSE converges to zero slowly when J is large, creating the *curse of dimensionality*.

To mitigate the *curse of dimensionality*, we regress Y_i on the d inverse regression variates X_iB , rather than Y_i on the J -dimensional X_i . The reduced variates X_iB capture all the information in X_i that is relevant for Y_i , i.e., $F_0(X_i) \equiv F_1(X_iB)$, which implies that we can estimate $F_0(X_i)$ using

$$(10) \quad \hat{F}_1(X_iB) = \frac{\sum_{j=1}^n K_{ij} Y_j}{\sum_{j=1}^n K_{ij}},$$

where

$$(11) \quad K_{ij} = \prod_{k=1}^d K((X_i - X_j)\beta_k / h_k)$$

is a product kernel (Hall, Li and Racine 2007). The MSE of this kernel regression is $O(n^{-4/(d+4)})$, which is smaller than $O(n^{-4/(J+4)})$ if $d < J$.

In practice, we do not observe the true *e.d.r.* directions, so we estimate B using inverse regression. Inverse regression methods produce root- n consistent estimates, whose convergence rate is faster than those of nonparametric methods. Therefore, replacing B with \hat{B} does not affect the asymptotic properties of nonparametric estimates of $F_0(X_i)$ in general.

3.2 Identifying Nonparametric Factors

The matrix A and the nonparametric factors in (3) are not separately identifiable because $F_1(X_iB) = AH(X_iB) = A\Pi^{-1} \cdot \Pi H(X_iB) = \hat{A}^* \hat{H}^*(X_iB)$ holds for any $L \times L$ full rank matrix Π . Although $H(X_iB)$ and $\hat{H}^*(X_iB) = \Pi H(X_iB)$ are different functions, they span the same function space because $\hat{H}^*(X_iB)$ is a linear transformation of $H(X_iB)$. Consequently, we could span the function space using either $H(X_iB)$ or $\hat{H}^*(X_iB)$. This fact implies that we can consistently estimate the function space spanned by the nonparametric factors using

$$(12) \quad \hat{H}(X_i\hat{B}) = \Xi \hat{F}_1(X_i\hat{B}),$$

where Ξ denotes an $L \times G$ matrix with full rank L .

To identify the nonparametric factors, we impose the orthonormal restriction

$$(13) \quad \Xi \left(n^{-1} \sum_{i=1}^n \hat{F}_1(X_i\hat{B}) \hat{F}_1(X_i\hat{B})' \right) \Xi' = I_L.$$

This restriction generates orthogonal nonparametric factors, and it implies that we can estimate Ξ using an eigenvalue decomposition. Specifically, we set $\Xi = \Lambda^{-1/2} C'$, where Λ is a diagonal matrix containing the L largest eigenvalues of $n^{-1} \sum_{i=1}^n \hat{F}_1(X_i\hat{B}) \hat{F}_1(X_i\hat{B})'$ and C is a $G \times L$ matrix containing the corresponding eigenvectors.

The factor loadings matrix, $\hat{A} = \Xi'(\Xi\Xi')^{-1}$, measures the sensitivity of the dependent variables to the nonparametric factors. Because (13) implies $n^{-1} \sum_{i=1}^n \hat{H}(X_i\hat{B}) \hat{H}(X_i\hat{B})' = I_L$, the resulting L nonparametric factors are orthogonal and have a unit norm. Consequently, the magnitude of \hat{A} indicates the relative contribution of each nonparametric factor to the variation of $F_1(X_iB)$.

To visualize the function space, we use the $L \times d$ scatter plots of the estimated nonparametric factors, $[\hat{H}_1(X_i\hat{B}), \dots, \hat{H}_L(X_i\hat{B})]$, against the estimated inverse regression variates, $(X_i\hat{\beta}_1, \dots, X_i\hat{\beta}_d)$. Such bi-variate plots can effectively reveal features of the data that would be imperceptible from the bi-variate plots of each Y against each X or from the plots of estimated nonparametric regression functions $\hat{F}_1(X_i\hat{B})$ against $X_i\hat{B}$. We illustrate function space visualization in Section 4.

3.3 Order Selection

Function space visualization requires the user to identify the number of nonparametric factors as well as the number of inverse regression variates. We use cross validation to jointly determine the number of significant nonparametric factors L , the number of significant NNIR variates d , and the optimal bandwidths $h = (h_1, \dots, h_d)$ for kernel regression.

Specifically, we select (d, h, L) by minimizing the estimated prediction error (EPE). We calculate the EPE using

$$(14) \quad EPE(d, h, L) = \frac{\sum_{i=1}^n \omega(X_i\hat{B}) (Y_i - \tilde{A}\hat{H}_{-i}(X_i\hat{B}))' (Y_i - \tilde{A}\hat{H}_{-i}(X_i\hat{B}))}{\sum_{i=1}^n \omega(X_i\hat{B})},$$

where $\omega(X_i\hat{B}) \in \{0, 1\}$ denotes a trimming function that removes outliers in $X_i\hat{B}$ thereby reducing sensitivity of EPE to outliers. We specify this function to equal zero for all observations for which the maximum value of $X_i\hat{B}$ exceeds its mean by more than a pre-set threshold. This approach enables us to trim the same observations for all candidate values of d , h , and L . In our application, we set the threshold to remove one percent of the observations in the first four inverse regression variates.

The terms \tilde{A} and \hat{H}_{-i} in (14) are leave-one-out estimates of factor loadings and nonparametric factors with the i^{th} observation being deleted while performing the estimation. Specifically, we first estimate $F_1(X_i B)$ using the leave-one-out estimator

$$(15) \quad \tilde{F}_{-i}(X_i B) = \frac{\sum_{j \neq i}^n K_{ij} Y_j}{\sum_{j \neq i}^n K_{ij}} .$$

Next, we use an eigenvalue decomposition to estimate $\tilde{\Xi}$ from

$$\tilde{\Xi} \left(n^{-1} \sum_{i=1}^n \tilde{F}(X_i \hat{B}) \tilde{F}(X_i \hat{B})' \right) \tilde{\Xi}' = I_L$$

which is the orthonormal restriction in (13) . Finally, we estimate the factor loadings matrix as $\tilde{A} = \tilde{\Xi}' (\tilde{\Xi} \tilde{\Xi}')^{-1}$ and the nonparametric factors as $\hat{H}_{-i}(X_i \hat{B}) = \hat{\Xi} \hat{F}_{-i}(X_i \hat{B})$.

As advocated by Hall, Li and Racine (2007), we use a product kernel in the nonparametric regression. This kernel allows a different bandwidth for each inverse regression variate and therefore can remove an insignificant variate from the model by selecting a very large bandwidth for it. When $(X_i - X_j) \beta_k$ is divided by a large bandwidth, it is close to zero, and consequently the individual kernel $K((X_i - X_j) \beta_k / h_k)$ is close to a constant that cancels out the corresponding term in the denominator.

The EPE statistic in equation (14) may overestimate d because it does not account explicitly for parameter estimation error in estimating B . The leave-one-out nonparametric regression in (15) uses the same \hat{B} for each observation i and that estimate is obtained using all the data. An alternative would be to leave observation i out of the inverse regression estimate that is used to predict observation i . This alternative is unlikely to make a difference because the order of magnitude of the error in the nonparametric fit exceeds the order of magnitude of the

error in estimating B . In our application, we find no material difference in the EPE or the selected model from using this more computationally expensive alternative.

4. Application to Financial Asset Pricing

Asset pricing theory purports that risk premia depend on the sensitivity of each asset to aggregate risk. The finance literature has proposed and tested numerous measures of aggregate risk. In a seminal paper, Chen, Roll, and Ross (1986) showed that inflation shocks, industrial production, the spread between high and low grade bonds, and the spread between long and short term interest rates were priced in expected returns. More recent additions to the list of priced risk factors include growth in various components of consumption (e.g., Parker and Julliard, 2005; Yogo, 2006; Hansen, Heaton, and Li, 2006), investment growth (e.g., Cochrane, 1996), GDP growth (e.g., Vassalou, 2003), liquidity risk (e.g., Pastor and Stambaugh, 2003; Acharya and Pedersen, 2005), and the consumption-to-wealth ratio (e.g., Lettau and Ludvigson, 2001). For a more comprehensive list of priced macroeconomic risk factors, see Lewellen, Nagel, and Shanken (2010). These papers follow a standard research template by first devising a theory that a particular variable may be a priced risk factor and then testing econometrically whether it is priced. Typically, the proposed factors are tested in isolation or against a small subset of alternatives. We take the opposite approach by allowing the data to choose which macroeconomic variables are most strongly related to the cross section of stock returns.

For asset returns, we use 25 portfolios of monthly stock returns constructed by Fama and French.² These portfolios comprise the intersections of 5 portfolios formed on firm size and 5 portfolios formed on the ratio of book equity to market equity. We model excess returns,

² http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

which we calculate by subtracting the monthly average of the one-month Treasury bill rate from each portfolio return. For the explanatory variables, we use a selection of 15 monthly macroeconomic variables measured over the period 1960:2–2008:12. Table 1 lists these variables and the data source for each one.

We use eight real and seven nominal macroeconomic variables. The first seven variables measure the monthly flow of real economic activity, and the eighth variable, unemployment, provides a snapshot of the employment situation during the month. We also include an inflation variable and several interest rate variables. We create the term spread and Treasury bond yield series as the difference between the yield on the last day of the month and the last day of the previous month. Our credit spread variable equals the difference between Moody's Seasoned Baa Corporate bond yield and the 10-year Treasury bond yield. Because daily corporate bond data go back only to 1983, we use monthly data to calculate credit spreads. The reported monthly data measure average daily yields during the month, whereas we would like to measure the yield at the end of the month, which is the same time we measure stock returns. To account for the fact that this variable measures interest rates with a lag, we include both the contemporaneous and first lead of credit spreads. The final variable is a survey of consumer expectations, for which we also include a one-month lead to account for the potential staleness of each monthly observation by the end of the month.

4.1 Variable and Function Space Reduction

The econometric model is $Y_i = F_1(X_iB) + U_i$, where R_i denotes the vector of excess stock returns, $F_1(X_iB) = AH(X_iB)$, and U_i satisfies $E(U_i | X_i) = \underline{0}$. In this model, excess stock

returns are linear in the nonparametric factors H , which in turn are unknown functions of the observed X variables. This model differs from existing asset pricing models because it includes a large number of X variables, and it allows for the factors to be nonlinear functions of X . For convenience, we divide each X variable by its standard deviation before we reduce dimensions for both variable space and function space.

To reduce the variable space, we estimate the *e.d.r.* directions, B , using the NNIR method introduced in Section 2. To indicate the relative importance of NNIR variates, we report in Table 2 the eigenvalues of \hat{M}_{NNIR} from the inverse regression algorithm. The first two eigenvalues appear notably greater than the next three, which appear notably greater than the remaining eigenvalues. These results suggest that between two and five NNIR variates may be needed to explain the variation in the 25 stock portfolios. We use the cross validation (CV) procedure introduced in Section 3.3 to select jointly the number of significant NNIR variates, d , along with the optimal kernel regression bandwidths, h , and the number of significant nonparametric factors, L .

The CV procedure suggests that the smallest EPE, 581.51, is achieved when $L=3$, $d=3$, and the optimal bandwidths for the first three NNIR variates are $(h_1=1.06, h_2=0.62, h_3=2.31)$. In Table 3, we report detailed CV results when trimming 1% of the data. We obtain similar results when trimming up to 5% of the data. The table reports the fit obtained by each specification, which is a pseudo out-of-sample statistic because EPE is calculated using leave-one-out kernel regression. The model explains about 25 percent of the variation in returns when $d=3$, compared to 14.5 percent for a single index model ($d=1$). Setting $d=2$ achieves a fit of 21 percent (not shown in Table 3). The relatively large increment in fit when increasing d from one to two, along with the relatively large bandwidth on the third NNIR variate, shows that the first

two NNIR variates contribute more to the model than the third. These results show a substantial improvement in fit from moving beyond a single-index model.

The CV procedure minimizes EPE at three nonparametric factors. The eigenvalues in Table 2 show that the first of these factors dominates. A single nonparametric factor explains vastly more of the variation in returns than the remaining 24 factors; the largest eigenvalue equals 208.2 and the next largest is 4.91. The explanatory power of the first nonparametric factor can also be seen in the EPE values in Table 3. In the multi-index model, a single factor causes the EPE to drop from 775.6 to 584.16, but adding a second factor reduces EPE only by a further 2.57 points to 581.59. Nonetheless, EPE does decline as we increase the number of nonparametric factors to three, which suggests that the second and third factors are statistically significant.

To compare the fit of our dimension reduction model with that of canonical correlations, we also apply cross validation to determine the number of significant canonical correlations between Y and X . To make the results comparable to those in Table 3, we trim the same observations when calculating EPE. The EPE values in Table 4 indicate that two canonical correlations are statistically significant. However, even the smallest EPE in Table 4 (616.57) is larger than the EPE of any multi-index model in Table 3, although it is smaller than the EPE for the single index models.

To further assess whether the nonlinearity is statistically significant, we apply Anderson and Vahid's (1998) test for common nonlinear components in multiple time series. The test begins with a candidate nonlinear function of X , for which we use $W_i = [(X_i \hat{\beta}_1)^2, (X_i \hat{\beta}_2)^2, (X_i \hat{\beta}_3)^2]$, where $X_i \hat{\beta}_1$, $X_i \hat{\beta}_2$, and $X_i \hat{\beta}_3$ denote the first three NNIR

variates. The next step is to compute the residuals from regressions of Y on X and W on X to remove the linear components of the relationship between Y and X . In the final step, we calculate the canonical correlations between these two sets of residuals. We obtain p-values of 0.08, 0.01, and 0.00 for tests of the null hypotheses that there are at most 2, 1, and 0 common nonlinear components. Using a 5 percent significance level, these statistics suggest the presence of two common nonlinear components and reinforce the CV result that the link function is nonlinear.

The fact that the canonical correlations model fits significantly worse than the multi-index model implies that significant nonlinearity exists in the relationship between stock returns and the macroeconomic variables. In the next sub-section, we compare the loadings of NNIR variates and canonical correlations to show how they differ.

4.2 Interpretation of NNIR Variates and Comparison with Canonical Correlations

In Figure 2, we plot the estimated betas of the first three NNIR variates as well as corresponding canonical correlations. Because the standardized X variables have standard deviation equal one and each beta vector is constrained to have unit length, the betas are scale independent and therefore indicate the relative importance of each variable. Moreover, because the NNIR variates, XB , are arguments in a nonparametric function, the sign of XB is not identified. The model is the same whether we use XB or $-XB$. Thus, to aid interpretation we impose a positive sign on the consumer expectation variable (#14) as we expect this variable to be positively correlated with returns.

The first NNIR variate is dominated by the 10-year Treasury bond yield (#11) and credit spread (#13), which also featured prominently in the seminal macroeconomic asset pricing

paper by Chen, Roll, and Ross (1986). The lead of the credit spread variable (#13) has a larger coefficient than the contemporaneous value (#12), suggesting that average credit spreads over a month contain much information that is stale for stock returns during the month. The Treasury bond and credit spread variables both measure long-run credit market conditions and reflect inflationary expectations and default risk, respectively.

The consumption variables have the next largest set of coefficients in the first NNIR variate. Consumption of non-durables (#3) has a positive sign and consumption of durables (#2) a negative sign, which implies that substitution from non-durable to durable consumption is negatively associated with this factor. One possible cause of such substitution is an increase in expected inflation. Barsky et. al. (2007) use a New Keynesian model to show theoretically that a monetary expansion causes real consumption to increase in the sector with more flexible prices relative to the sticky-price sector. If durable goods prices are stickier, then the expectation of high prices in the future causes consumers to make more non-durable goods purchases today before prices rise. It follows that a drop in durable goods consumption relative to non-durables may indicate a monetary expansion and higher expected future inflation. Overall, we interpret large values of the first NNIR variate as indicating low credit risk and low inflationary expectations.

The first canonical correlation also loads on the 10-year Treasury bond yield (#11) and credit spread (#13). For the consumption variables, it picks up non-durables (#3) but not durables (#2). It assigns high weights to both contemporaneous and lead of the credit spread variable (#12 and # 13) as well as the consumption-wealth ratio (cay, #1). Overall, the first canonical correlation is not dissimilar to the first NNIR variate.

The second NNIR variate loads positively on four variables that foreshadow future

increases in aggregate output. These variables are consumption of durables (#2), housing starts (#7), term spread (#10), and consumer expectations (#15). Moreover, it loads negatively on the credit spread (#12), which is negatively associated with future output increases. In contrast, this variate loads negatively on three measures of current economic activity: personal income (#5), industrial production (#6), and the inflation shock (#9). It also has a negative coefficient on the consumption-wealth ratio (#1, *cay*, Lettau and Ludvigson, 2004), which measures current consumption of non-durables and services relative to income and wealth. Overall, the second variate loads positively on leading indicators of future macroeconomic growth and negatively on indicators of current macroeconomic activity.

The second canonical correlation picks up consumption of durables (#2), housing starts (#7), term spread (#10), credit spread (#12), personal income (#5) with the same sign and magnitude as those of the second NNIR variate. However, it differs from the second NNIR variate by being weakly related to consumer expectations (#15), industrial production (#6), the inflation shock (#9), and the consumption-wealth ratio (#1, *cay*).

The large CV bandwidth in Table 3 implies that the third NNIR variate relates weakly to returns. This variate is most strongly related to the inflation shock variable (#9) and industrial production (#6). The third canonical correlation is not statistically significant, and it favors the 10-year Treasury bond yield (#11) and the contemporaneous value of the credit spread variable (#12).

4.3 Function Space Visualization

Data visualization is often the first step of preliminary data analysis. If we had only a few variables, it would be convenient to use bivariate scatter plots to examine the relationships

among them. However, plotting 25 returns against 15 macroeconomic variables would require 375 bivariate plots, which are too many to synthesize even if some plots share common features. Reducing the 15 macroeconomic variables to three NNIR variates reduces the number of bivariate plots to 75; reducing the 25 return functions to three nonparametric factors further narrows down the task of visualization to nine bivariate plots. We present these nine scatter plots in Figure 3. Specifically, we plot each estimated nonparametric factor $\hat{H}_l(X_i\hat{\beta}_1, X_i\hat{\beta}_2, X_i\hat{\beta}_3)$ against each NNIR variate $X_i\hat{\beta}_k$ for $l = 1, 2, 3$ and $k = 1, 2, 3$.

The first factor is positively correlated with the first two NNIR variates ($X_i\hat{\beta}_1$ and $X_i\hat{\beta}_2$) and weakly correlated with the third variate $X_i\hat{\beta}_3$. The second factor is also positively correlated with $X_i\hat{\beta}_1$, but is negatively correlated with $X_i\hat{\beta}_2$. The third factor is negatively correlated with $X_i\hat{\beta}_1$, especially when $X_i\hat{\beta}_1$ is negative; it is weakly correlated with the other two variates. The weak relationship of all three factors to $X_i\hat{\beta}_3$ follows from the large bandwidth chosen for it by the CV procedure (see Table 3). As implied by the eigenvalues in Table 2, the first factor explains almost all the variation in equation-by-equation regression functions $\hat{F}_1(X_i\hat{B})$. Specifically, the correlation between the first nonparametric factor, $\hat{H}_1(X_i\hat{B})$, and the 25 columns of $\hat{F}_1(X_i\hat{B})$ ranges from 0.92 to 0.99 with a median of 0.98. Thus, because it explains almost all the variation in the data, we focus our discussion on interpretation of the first nonparametric factor.

The first nonparametric factor is a nonlinear function of $X_i\hat{\beta}_1$ and $X_i\hat{\beta}_2$. In fact, the existence of more than one significant NNIR variate implies nonlinearity, because if the factor were linear in both $X_i\hat{\beta}_1$ and $X_i\hat{\beta}_2$ we would not be able to separately identify $\hat{\beta}_1$ and $\hat{\beta}_2$. To

illustrate this nonlinearity and interpret this factor, we show in Figure 4 two views of the relationship between this factor and the first two NNIR variates. Panel A shows a surface plot of H_1 holding $X_i\hat{\beta}_3$ constant at its mean and varying $X_i\hat{\beta}_1$ and $X_i\hat{\beta}_2$ over a range that avoids the tails of the distribution. Panel B shows time series plots of H_{1it} , $X_i\hat{\beta}_1$, and $X_i\hat{\beta}_2$. For clarity, the figure shows calendar-year averages of these variables, rather than monthly observations. Panel B supports our leading indicator interpretation of $X_i\hat{\beta}_2$, as it tends to lead $X_i\hat{\beta}_1$.

The surface in Panel A of Figure 4 is steepest when both variates are positive and flattest when both variates are close to zero. Thus, the two variates interact positively with each other. Using our interpretation of the variates in Section 4.2, low credit risk and low inflationary expectations have a greater positive effect on stock returns when leading macroeconomic indicators are high relative to current economic activity. Panel B shows that such features arise early in the business cycle as the economy emerges from recession (e.g., 1976, 1982, 1991, and 2003).

As expansions proceed, we observe periods with positive $X_i\hat{\beta}_1$ and negative $X_i\hat{\beta}_2$ (e.g., mid 1960s, mid 1980s, mid-to-late 1990s and mid 2000s). In such periods, the economy is booming relative to its leading indicators, credit markets are loosening and inflationary expectations declining. The high values of the nonparametric factor in these periods reveals these are good times for stocks. Panel A shows that these periods coincide with larger values of the nonparametric factor than periods with $X_i\hat{\beta}_2$ close to zero. In other words, if $X_i\hat{\beta}_1$ is large and positive, then the nonparametric factor has a U-shaped relationship with $X_i\hat{\beta}_2$. Given that the region with positive $X_i\hat{\beta}_1$ and negative $X_i\hat{\beta}_2$ is often visited when the economy is booming,

it may reflect declining risk aversion. Reinforcing this conjecture is the fact that these episodes immediately preceded events such as the 1987 stock market crash, the 1996 “irrational exuberance” speech by Alan Greenspan, the 2000 dotcom crash, and the recent subprime crisis.

We show in Table 5 the factor loading estimates (\hat{A}) that relate the nonparametric factors to the 25 stock returns. The magnitudes of the loadings on the first factor are much greater than those for the other factors, reflecting the fact that the first factor has most of the explanatory power. Small stocks are more sensitive to the first factor than large stocks; the average loading for the smallest quintile is 3.2 compared to 2.2 for the largest quintile. This result is consistent with the notion that small stocks are more risky, and therefore co-vary more strongly with the aggregate risk factor.

Growth stocks, which are represented in the leftmost column, have larger factor loadings than any column to the right. This result is consistent with the notion that growth stocks, which consist mostly of young companies with small current earnings relative to expected future earnings, are more susceptible to aggregate risk than other stocks. The loadings are non-monotonic in value quintile. The loading for the highest quintile exceeds the loading for the second-highest quintile in all cases.

5. Conclusions

In this article, we present a new and convenient framework to analyze high-dimensional nonlinear econometric models. Rather than imposing a parametric model *a priori*, we take a nonparametric approach and apply dimension reduction techniques to evade the *curse of dimensionality*. We reduce dimension in both the variable space and the function space. For the variable space, we reduce a large number of explanatory variables to a few inverse regression

variates. For the function space, we model a large number of regression functions using a small number of nonparametric factors. Our dimension reduction method provides a practical toolkit for nonparametric estimation, data visualization, and parsimonious model specification, thereby enabling a more detailed understanding of complex economic processes.

We apply this framework to 25 stock return portfolios. We use inverse regression to extract from a set of 15 macroeconomic variables three variates that have substantial explanatory power for stock returns. There are substantial gains in fit from moving beyond a single index model. The first two inverse regression variates can naturally be interpreted as representing low credit risk and low inflationary expectations (the first variate) and strong leading economic indicators relative to current economic activity (the second variate). The third variate captures inflation shocks and makes only a small contribution to the model. We show that combining the three inverse regression variates into a single nonparametric factor explains 95 percent of the common variation in the stock returns. This nonparametric factor is nonlinear in the variates, with its most notable nonlinear feature revealing an increasing appetite for risk during the expansion stage of the business cycle.

References

- V. Acharya and L. Pedersen (2005), "Asset Pricing with Liquidity Risk", *Journal of Financial Economics*, Vol. 77 (2), pp.375-410.
- H.M. Anderson, and F. Vahid (1998), "Testing Multiple Equation Systems for Common Nonlinear Components", *Journal of Econometrics*, Vol. 84, pp. 1-36.
- U. Amato, A. Antoniadis, and I. De Feis (2006), "Dimension Reduction in Functional Regression with Applications", *Computational Statistics and Data Analysis*, Vol. 50, pp. 2422-2446.
- R.B. Barsky, C.L. House, and M.S. Kimball (2007), "Sticky-Price Models and Durable Goods", *The American Economic Review*, Vol. 97, pp. 984-998.
- E. Bura and R.D. Cook (2001a), "Extending SIR: The Weighted Chi-Square Test", *Journal of the American Statistical Association*, Vol. 96, pp. 996-1003.
- G. Camba-Mendez, G. Kapetanios, R.J. Smith, and M.R. Weale (2003), "Tests of Rank in Reduced Rank Regression Models", *Journal of Business and Economic Statistics*, Vol. 21, pp. 145-155.
- R.J. Carroll and K.C. Li (1992), "Measurement Error Regression with Unknown Link: Dimension Reduction and Data Visualization", *Journal of the American Statistical Association*, Vol. 87, pp. 1040-1050.
- C.H. Chen and K.C. Li (1998), "Can SIR Be as Popular as Multiple Linear Regression", *Statistica Sinica*, Vol. 8, pp. 289-316.
- N.F. Chen, R. Roll, and S. Ross (1986), "Economic Forces and the Stock Market", *Journal of Business*, Vol. 59, pp. 383-403.

- J. Cochrane (1996), "A Cross-Sectional Test of an Investment-based Asset Pricing Model",
Journal of Political Economy, Vol. 104, 572-621.R.D.
- Cook (1994), "On the Interpretation of Regression Plots", Journal of the American Statistical
Association, Vol. 89, pp. 177-189.
- R.D. Cook and X. Yin (2001), "Dimension Reduction and Visualization in Discriminant Analysis",
Australian & New Zealand Journal of Statistics, Vol. 43(2), pp. 147-199.
- R.D. Cook and S. Weisberg (1991), Discussion of "Sliced Inverse Regression" by K.C. Li, Journal of
the American Statistical Association, Vol. 86, pp. 328-332.
- S. Donald (1997), "Inference Concerning the Number of Factors in a Multivariate Nonparametric
Relationship", Econometrica, Vol. 65, pp. 103-131.
- E. F. Fama and K.R. French (1992), "The Cross-Section of Expected Stock Returns", Journal of
Finance, Vol. 47, pp. 427-465.
- E. F. Fama and K.R. French (1993), "Common Risk Factors in the Returns on Stocks and Bonds",
Journal of Financial Economics, Vol. 33, pp. 3-56.
- P. Hall and K.C. Li (1993), "On Almost Linearity of Low Dimensional Projections from High
Dimensional Data", Annals of Statistics, Vol. 21, pp. 867-889.
- P. Hall, Q. Li and J.S. Racine (2007), "Nonparametric Estimation of Regression Functions in the
Presence of Irrelevant Regressors", Review of Economics and Statistics, Vol. 89, pp. 784-
89.
- J.D. Hamilton (1994), Time Series Analysis, Princeton University Press.

- L. Hansen, J. Heaton, and N. Li (2006), "Consumption Strikes Back? Measuring Long-run Risk", *Journal of Political Economy*, Vol. 116, pp. 260-302.
- T. Hsing (1999), "Nearest Neighbor Inverse Regression", *Annals of Statistics*, Vol. 27, pp. 697-731.
- H. Ichimura (1993), "Semiparametric Least Squares (SLS) and Weighted SLD Estimation of Single-Index Models", *Journal of Econometrics*, Vol. 58, pp. 71-120.
- R.L. Klein and R.H. Spady (1993), "An Efficient Semiparametric Estimator for Binary Response Models", *Econometrica*, Vol. 61, pp. 387-421.
- M. Lettau and S. Ludvigson (2004), "Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption," *American Economic Review*, Vol. 94, pp. 279-299.
- J. Lewellen, S. Nagel and J. Shanken (2010), "A Skeptical Appraisal of Asset Pricing Tests", *Journal of Financial Economics*, Vol. 96 (2), pp. 175-194.
- K.C. Li (1991), "Sliced Inverse Regression for Dimension Reduction", *Journal of the American Statistical Association*, Vol. 86, pp. 316-333.
- L. Li and H. Li (2004), "Dimension Reduction Methods for Microarrays with Application to Censored Survival Data", *Bioinformatics*, Vol. 20, pp. 3406-3412.
- J. Lintner (1965), "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets", *Review of Economics and Statistics*, Vol. 47, pp. 13-37.
- P.A. Naik, M. Wedel, and W. Kamakura (2010), "Multi-Index Binary Response Analysis of Large Data Sets", *Journal of Business and Economic Statistics*, Vol. 28, pp. 67-81.
- A. Pagan and A. Ullah (1999), "Nonparametric Econometrics", Cambridge University Press.

- J. Parker and C. Julliard(2005), "Consumption Risk and the Cross Section of Expected Returns",
Journal of Political Economy, Vol.113, pp. 185-222.
- L.Pastor and R. Stambaugh (2003), "Liquidity Risk and Expected Stock Rreturns", Journal of
Political Economy, Vol.111, pp. 642-685.
- J. Powell, J. Stock, and T. Stocker (1989), "Semiparametric Estimation of Index Coefficients",
Econometrica, Vol. 57, pp. 1403-1430.
- S.A. Ross (1976), "The Arbitrage Theory of Capital Asset Pricing", Journal of Economic Theory,
Vol. 13, pp.341-360.
- W.F. Sharpe (1964), "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of
Risk", Journal of Finance, Vol. 19, pp. 425-442.
- J.H. Stock and M.W. Watson (2002), "Macroeconomic Forecasting Using Diffusion Indexes",
Journal of Business and Economic Statistics, Vol. 20, pp. 147-162.
- M. Vassalou (2003), "News Related to Future GDP Growth as a Risk Factor in Equity Returns",
Journal of Financial Economics, Vol. 68, pp. 47-73.
- M. Yogo (2006), "A Consumption-based Explanation of Expected Stock Returns", Journal of
Finance, Vol. 61, pp. 539-580.

Table 1: Macroeconomic Variables

	Variable Description	Source	Transform	Seas Adj
1	consumption-wealth ratio (cay)	Lettau	2	yes
2	real personal cons exp per capita: durables	FRED(pcedc96/pop)	1	yes
	real personal cons exp per capita: non-durables	FRED(pcendc96/pop)	1	yes
3	durables	FRED(pcendc96/pop)		
4	real personal cons exp per capita: services	FRED(pcesc96/pop)	1	yes
5	real personal income per capita	FRED(dspic96/pop)	1	yes
6	industrial production	FRED(indpro)	1	yes
7	housing starts	FRED(houst)	1	yes
8	unemployment rate	FRED(unrate)	2	yes
9	inflation shock (deviation from 12-month ave.)	FRED(cpiaucsl)	3	yes
10	term spread	FRED(dgs10,dtb3m)	2	no
11	10-year T-bond yield	FRED(dgs10)	2	no
12	credit spread	FRED(baa,gs10)	2	no
13	credit spread (1 month lead)	FRED(baa,gs10)	2	no
14	UM index of consumer expectation	BCD-83/UM	2	no
15	UM index of cons expectation (1 month lead)	BCD-83/UM	2	no

Notes:

(1) We obtained the cay variable from http://faculty.haas.berkeley.edu/lettau/data_cay.html (Lettau's website) and converted it from quarterly to monthly frequency using linear interpolation.

(2) The consumer expectations variable is as used in Stock and Watson (2002) under the moniker "hhsntn". We updated the consumer expectations series using the final column of Table 3 published at <http://www.sca.isr.umich.edu> (University of Michigan Survey Research Center website).

(3) The variables are transformed according to the following code: 1=log difference (growth), 2=first difference, and 3= $\pi_t - \sum_{i=1}^{12} \pi_{t-i} / 12$.

Table 2: Eigenvalues from Variable and Function Space Reduction

Variable Space (NNIR)		Function Space (F'F)	
1	0.34	1	208.20
2	0.27	2	4.91
3	0.19	3	2.95
4	0.16	4	0.72
5	0.13	5	0.51
6	0.07	6	0.31
7	0.04	7	0.27
8	0.02	8	0.20
9	-0.01	9	0.18
10	-0.03	10	0.16
11	-0.06	11	0.14
12	-0.08	12	0.13
13	-0.11	13	0.10
14	-0.15	14	0.10
15	-0.16	15	0.09
		16	0.07
		17	0.06
		18	0.05
		19	0.04
		20	0.03
		21	0.03
		22	0.03
		23	0.02
		24	0.02
		25	0.02

Note: We computed the eigenvalues for function space using the first three NNIR variates, which were selected by cross-validation (see Table 3).

Table 3: Cross Validation for Optimal (d, h, L)

Models	L	EPE	Fit	$h1$	$h2$	$h3$	$h4$
Mean-Only	0	775.6					
Single Index Model	1	663.5	14.4%	0.88	-	-	-
	2	662.2	14.6%	0.88	-	-	-
	3	662.5	14.6%	0.89	-	-	-
	4	662.5	14.6%	0.89	-	-	-
	25	663.1	14.5%	0.89	-	-	-
Double-Index Model	1	617.8	20.4%	1.04	0.56	-	-
	2	615.6	20.6%	1.04	0.56	-	-
	3	615.8	20.6%	1.05	0.56	-	-
	4	616.0	20.6%	1.05	0.56	-	-
	25	618.2	20.3%	1.04	0.62	-	-
Multi-Index Model	1	583.8	24.7%	1.03	0.62	2.31	1000
	2	581.3	25.1%	1.03	0.62	2.31	1000
	3	581.3	25.1%	1.03	0.62	2.31	1000
	4	581.6	25.0%	1.03	0.62	2.31	1000
	25	583.6	24.8%	1.04	0.62	2.31	1000

Note: 1% trimming, Epanechnikov kernel, Fit is percent EPE reduction over mean-only model. Because EPE is calculated using leave-one-out kernel regression, the fit is a pseudo out-of-sample statistic.

Table 4: Cross Validation for Canonical Correlations

Number of CC	EPE	Fit
1	625.4	19.4%
2	625.0	19.4%
3	625.2	19.4%
4	626.8	19.2%
15	647.4	16.5%

Table 5: Nonparametric-Factor Loadings by Portfolio

Loadings on First Nonparametric Factor

	value 1	value 2	value 3	value 4	value 5
size 1	3.65	3.33	3.07	2.92	3.21
size 2	3.47	3.17	2.93	2.88	3.09
size 3	3.32	3.02	2.67	2.68	2.94
size 4	2.99	2.81	2.81	2.57	2.79
size 5	2.41	2.28	2.07	2.17	2.25

Loadings on Second Nonparametric Factor

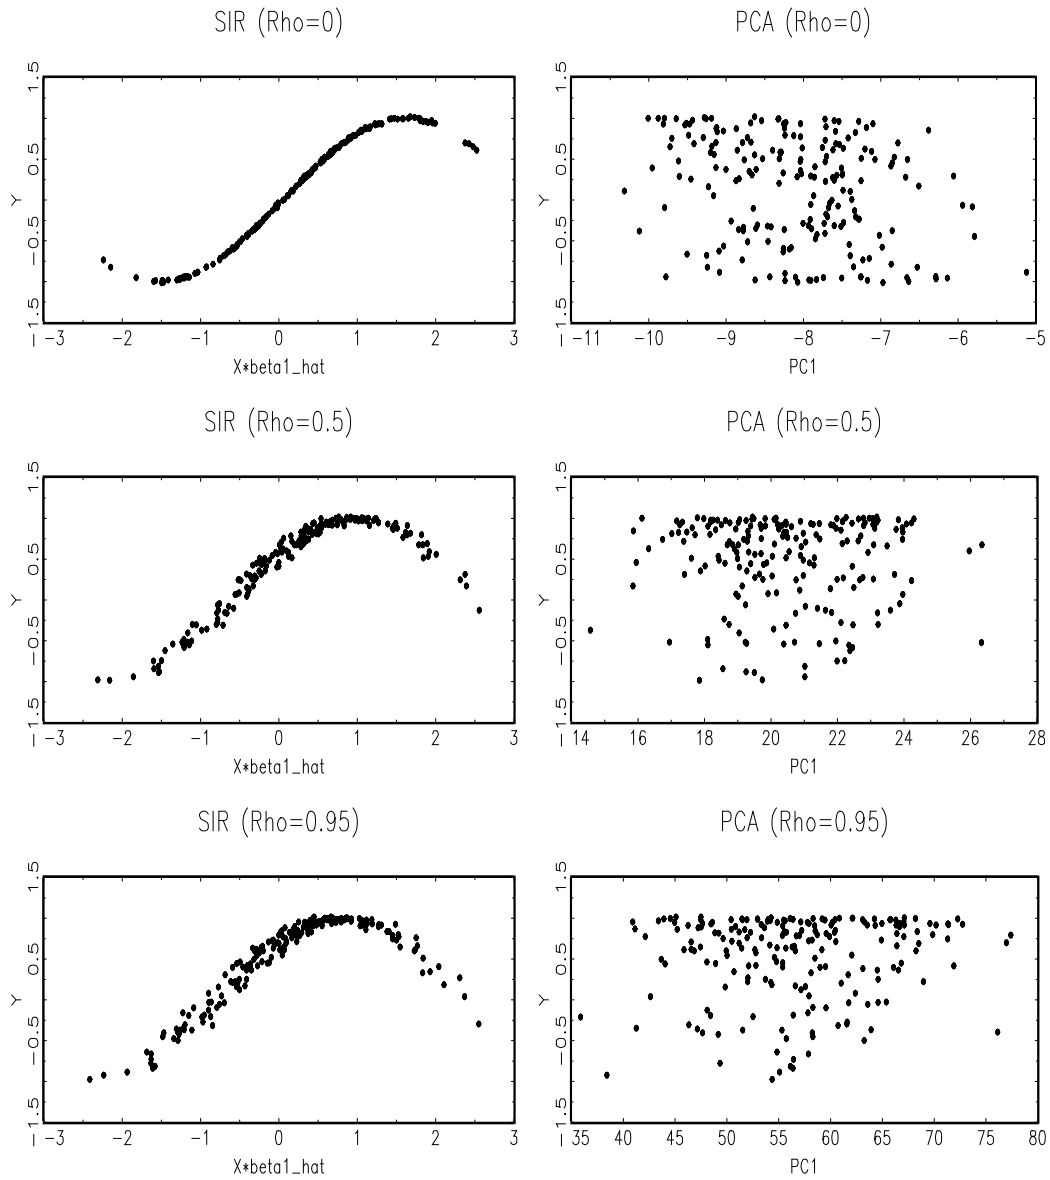
	value 1	value 2	value 3	value 4	value 5
size 1	-0.98	-0.73	-0.43	-0.30	-0.34
size 2	-0.52	-0.29	-0.03	0.14	-0.01
size 3	-0.34	0.05	0.21	0.26	0.20
size 4	-0.02	0.35	0.48	0.48	0.40
size 5	0.42	0.57	0.68	0.65	0.51

Loadings on Third Nonparametric Factor

	value 1	value 2	value 3	value 4	value 5
size 1	-0.43	0.19	0.24	0.39	0.51
size 2	-0.48	-0.02	0.31	0.34	0.44
size 3	-0.56	0.01	0.19	0.21	0.36
size 4	-0.60	-0.25	-0.01	0.13	0.15
size 5	-0.67	-0.33	-0.24	-0.06	0.17

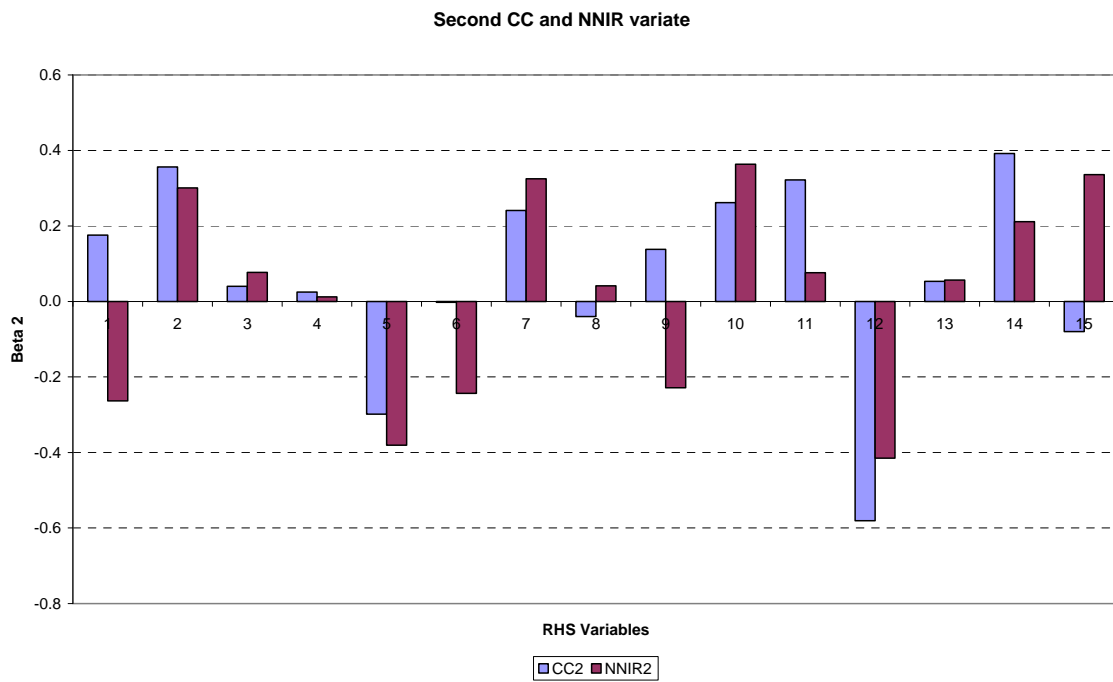
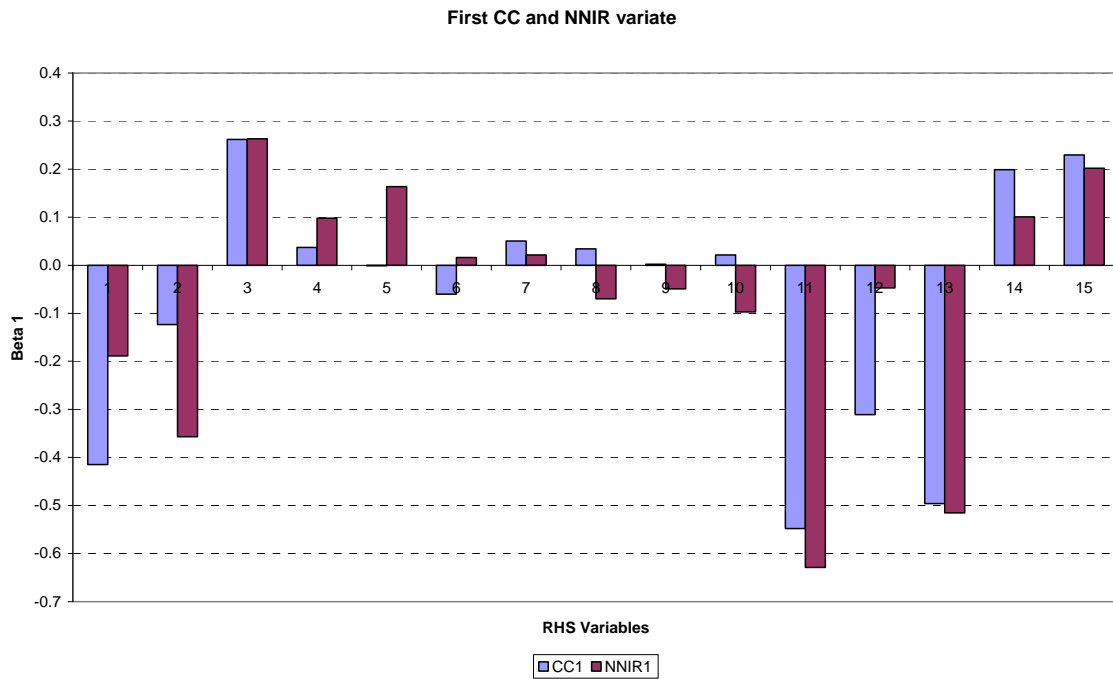
Note: Portfolios formed by intersections of 5 portfolios formed on firm size and 5 portfolios formed on the ratio of book equity to market equity. Sizes 1–5 represent small – large stocks; values 1–5 present growth – value stocks. We calculate excess returns by subtracting the monthly average of the one-month Treasury bill rate from each portfolio return. Sample period: 1960:2–2008:12.

Figure 1: Sliced Inverse Regression versus Principal Component Analysis

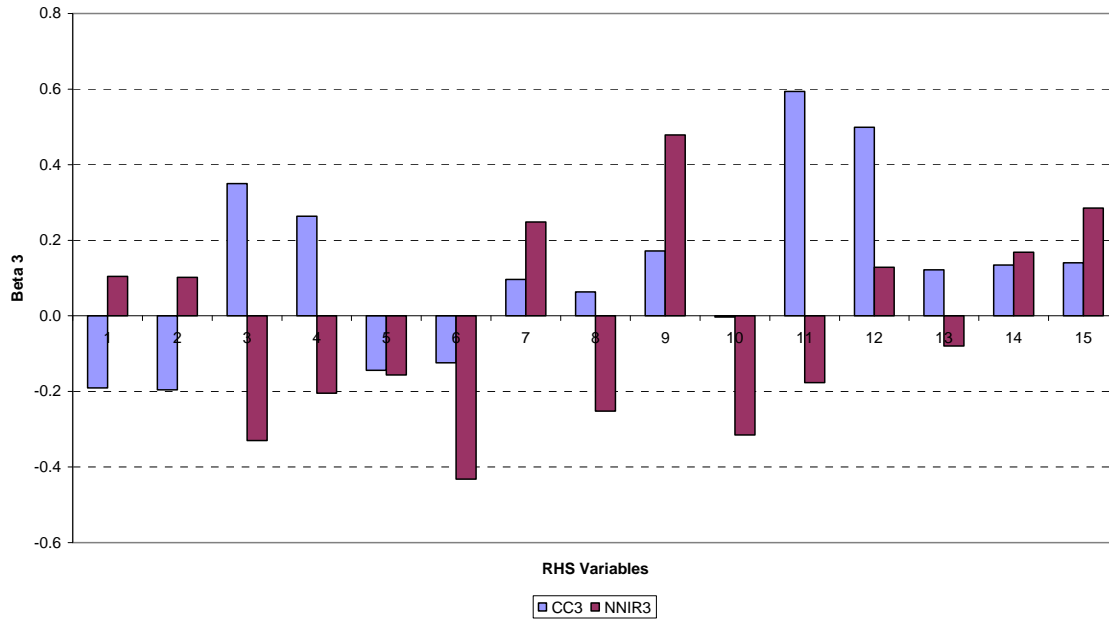


Note: We generate 200 observations using $Y_i = \sin(X_i \beta_1) + U_i$, where $U_i \sim iid N(0, 0.01^2)$, $\beta_1 = [0.25 \quad -0.35 \quad -0.57 \quad 0.70]$, $X_i \sim iid N(\mu, \Sigma_{XX})$, and $Rho \equiv corr(X_{ij}, X_{ik})$ for $j \neq k$. The graphs show Y_i plotted against $X \hat{\beta}$ for $\hat{\beta}$ estimated as the first SIR *e.d.r.* direction and the first principal component.

Figure 2: Loadings on the First Three NNIR Variates and Canonical Correlations




Third CC and NNIR variate



Note: The variables are: 1=cay; 2=PCE-durables; 3=PCE-non-durables; 4=PCE services; 5=personal income; 6=industrial production; 7=housing starts; 8=unemployment rate; 9=inflation shock; 10=term spread; 11=10-yr T-bond yield; 12=credit spread; 13=lead of credit spread; 14=UM index of cons. exp.; 15=lead of UM index of cons. exp. See Table 1 for details.


Figure 3: Scatter Plots of Three Nonparametric Factors against Three NNIR Variates



Note: Nonparametric factors estimated as in (13) with bandwidth as selected by CV and shown in Table 2.

Figure 4: Two Views of the First Nonparametric Factor and the First Two NNIR Variates

Panel A: Surface Plot



Panel B: Time Series Graph of Calendar-Year Averages

