Nonlinearity and Flight-to-Safety in the Risk-Return Tradeoff for Stocks and Bonds

Tobias Adrian   Richard Crump   Erik Vogt

Federal Reserve Bank of New York *

March 30, 2015

Abstract

We document a highly significant, strongly nonlinear dependence of stock and bond returns on past equity market volatility as measured by the VIX. We propose a new estimator for the shape of the nonlinear forecasting relationship that exploits additional variation in the cross-section of returns. The nonlinearities are mirror images for stocks and bonds, revealing flight-to-safety: expected returns increase for stocks when volatility increases from moderate to high levels while they decline for Treasuries. We further demonstrate that these findings are evidence of dynamic asset pricing theories where the time variation of the price of risk is a function of the level of the VIX.

JEL Classification: G12, G17, G01

Keywords: flight-to-safety, risk-return tradeoff, dynamic asset pricing, volatility, nonlinear regressions, intermediary asset pricing, asset management

*Adrian: tobias.adrian@ny.frb.org, Crump: richard.crump@ny.frb.org, Vogt: erik.vogt@ny.frb.org.
The views expressed in this paper are those of the authors and do not necessarily represent those of the Federal Reserve Bank of New York or the Federal Reserve System. The authors would like to thank Daniel Stackman and Rui Yu for outstanding research assistance.

1 Introduction

Investor flight-to-safety is pervasive in times of elevated risk (Longstaff (2004), Beber, Brandt, and Kavajecz (2009), Baele, Bekaert, Inghelbrecht, and Wei (2013)). Economic theories of investor flight-to-safety predict highly nonlinear asset pricing relationships (Vayanos (2004), Weill (2007), Caballero and Krishnamurthy (2008), Brunnermeier and Pedersen (2009)). Such nonlinear pricing relationships are difficult to document empirically as the particular shape of the nonlinearity is model specific, and inference of nonlinear relationships presents econometric challenges.

In this paper, we document an economically and statistically strong nonlinear risk-return tradeoff by estimating the relationship between stock market volatility as measured by the VIX and future returns. The nonlinear risk-return tradeoff features evidence of flight-to-safety from stocks to bonds in times of elevated stock market volatility consistent with the above cited theories. The VIX strongly forecasts stock and bond returns up to 24 months into the future when the nonlinearity is accounted for, in sharp contrast to the insignificant linear relationship.

The nature of the nonlinearity in the risk-return tradeoffs for stocks and bonds are virtually mirror images, as can be seen in Figure 1 on the next page, estimated from a large cross-section of stocks and bonds. Both stock and bond returns have been normalized by their unconditional standard deviation in order to allow plotting them in the same figure. There are three notable regions that characterize the nature of the nonlinear risk-return tradeoff, defined by the VIX median of 18 and the VIX 99.3rd-percentile of 50. When the VIX is below its median of 18, both stocks and bonds exhibit a risk return tradeoff that is relatively insensitive to changes in the VIX. In the intermediate 18-50 percent range of the VIX, the nonlinearity is very pronounced: as the VIX increases above its unconditional median, expected Treasury returns tend to fall, while expected stock returns rise. This finding is consistent with a flight-to-safety from stocks to bonds, raising expected returns to stocks and compressing expected returns to bonds. For levels of the VIX above 50, which has only

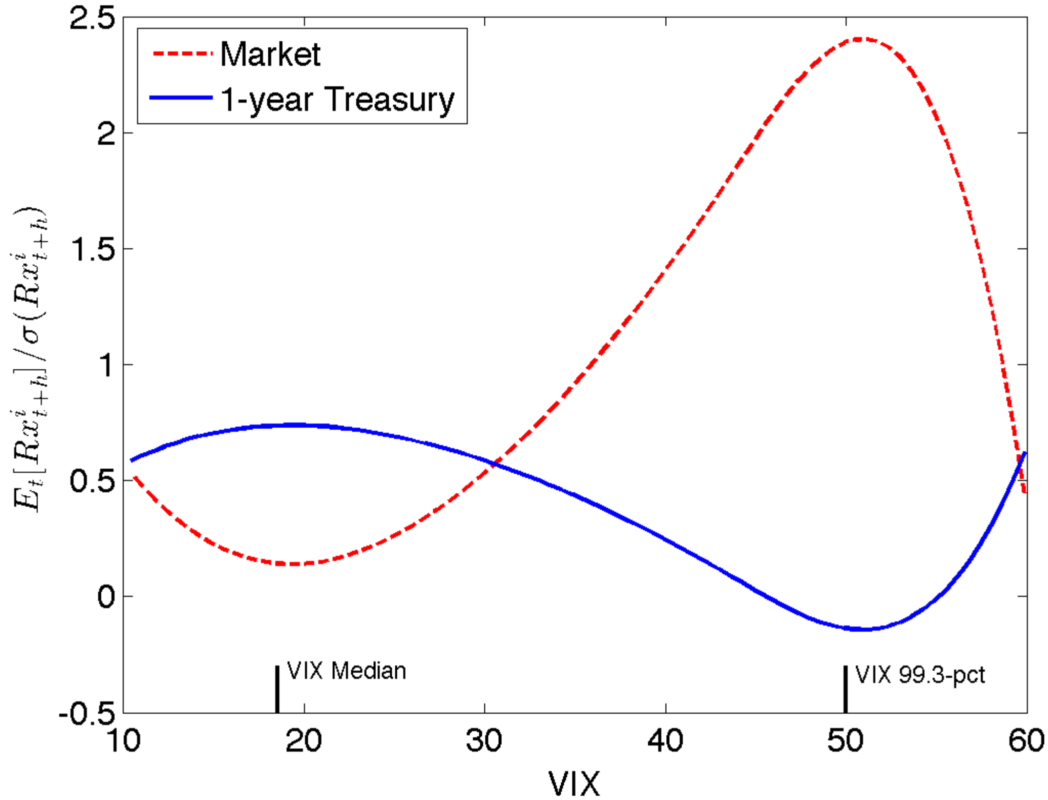

Figure 1: This figure shows the relationship between the six month cumulative equity market return and the six month lag of the VIX in red, as well as the relationship between the six month cumulative 1-year Treasury return and the six month lag of the VIX in blue. Both nonlinear relationships are estimated using reduced rank sieve regressions on a large cross-section of stocks and bonds. The y-axis is expressed as a ratio of returns to the full sample standard deviation. The x-axis shows the VIX.

occurred in the aftermath of the Lehman failure, this logic reverses, and a further increase in the VIX is associated with lower stock and higher bond returns. The latter finding for very high values of the VIX likely reflects the fact that severe financial crises are followed by abysmal stock returns and aggressive interest rate cuts, due to a collapse in real activity.

What is most notable is that a linear regression using the VIX does not forecast stock or bond returns significantly at any horizon. Nonlinear regressions, on the other hand, do forecast stock and bond returns with very high statistical significance and reveal the striking mirror image property of Figure 1. We study the nature of the nonlinearity and mirror image property in a variety of ways, using kernel regressions, polynomial regressions, as well as nonparametric sieve regressions. In all cases and on subsamples, we find pronounced nonlinearity within risky assets and reversed nonlinearities for safe assets, in terms of both statistical and economic magnitudes.

In order to estimate the shape of the nonlinearity in a robust way, we propose a novel way to nonparametrically estimate the shape using a reduced-rank sieve regression on a large cross-section of stock and bond returns. We specify a nonlinear forecasting function $\phi_h(v)$ according to the following set of equations

$$Rx^i_{t+h} = a^i_h + b^i_h \cdot \phi_h(vix_t) + \varepsilon^i_{t+h}, \qquad i = 1, \ldots, n,$$

where h denotes the forecasting horizon and i refers to the individual stock and bond portfolios and Rx are excess returns. The nonlinearity of the function $\phi_h(v)$ is highly significant, and its forecasting power is very strong. Importantly, when we estimate $\phi_h(v)$ separately for stocks and bonds, we obtain statistically indistinguishable functions (up to an affine transformation).

A major advantage of estimating $\phi_h(v)$ from a large panel of stock and bond returns is that it exploits additional cross-sectional variation unavailable in the univariate regressions that are typical in the return forecasting literature. The algebra for the estimator can be described intuitively in two stages. In the first stage, returns to each asset are regressed in the time series on lagged sieve expansions of the VIX. In the second stage, the rank of the matrix of forecasting coefficients is reduced using an eigenvalue decomposition, and only a rank one approximation is retained (see Adrian, Crump, and Moench (2014) for a related derivation). This is a dimensionality reduction that is optimal when errors are conditionally Gaussian and the number of regressors are fixed. The resulting factor $\phi_h(v)$ is a nonlinear function of volatility and is the best common predictor for the whole cross section of stock and bond returns.

The finding that the VIX forecasts stock and bond returns in a nonlinear fashion is robust to the inclusion of standard predictor variables such as the dividend yield, the BAA/10-year

Treasury default spread, the 10-year/3-month Treasury term spread, and the volatility risk premium. Furthermore, we show that the nonlinear relationship is highly significant for the 1990-2007 sample which excludes the 2008-09 financial crisis. Importantly, the shape of the nonlinearity in the 1990-2007 and the 1990-2014 sample resemble each other closely, even though the tail events in those samples are distinct. We also verify that Treasury returns are forecasted only by a nonlinear function of the VIX, not the Treasury implied volatility as measured by the MOVE. The latter result suggests that pricing of risk is proxied by the VIX as a common forecasting variable for stocks and bonds.

The sieve reduced rank regression estimator restricts expected returns of each asset i to be an affine transformation of $\phi_h(v)$ with intercept a_h^i and slope b_h^i. Asset pricing theories predict these coefficients to be determined by risk factor loadings. We take this prediction to the data, estimating the beta representation of a dynamic asset pricing kernel that features the market return, the one year Treasury return, and innovations to $\phi_h(v)$ as cross-sectional pricing factors, and $\phi_h(v)$ as price of risk variable. We show that this asset pricing model performs well in pricing the cross-section of stock, bond, and credit portfolios, and that there is a tight cross-sectional relationship between the forecasting slopes b_h^i and the risk factor loadings.

The dynamic asset pricing results indicate that the pricing of risk over time is related to the level of volatility in a nonlinear fashion. A number of alternative theories are compatible with such a finding, including 1) flight-to-safety theories due to redemption constraints on asset managers, 2) macro-finance models with financial intermediaries, and 3) representative agent models with habit formation. We discuss the extent to which each of these types of theories are compatible with our findings.

Among asset management pricing theories, our findings are particularly in line with the theory of Vayanos (2004), where asset managers are subject to funding constraints that (endogenously) depend on the level of market volatility. When volatility increases, the likelihood of redemptions rises, leading to a decline in the risk appetite of the asset managers. Increases in volatility generate flight-to-safety as managers attempt to mitigate the impact of higher volatility on redemption risk by allocating more to relatively safe assets. In equilibrium, the dependence of risk appetite on volatility generates expected returns with features that are qualitatively similar to our estimated function $\phi_h(v)$. Furthermore, Vayanos (2004) theory gives rise to a dynamic asset pricing kernel that would predict that the forecasting slope b_h^i is cross sectionally related to risk factor loadings, as explained earlier.

We present direct evidence in favor of the flight-to-safety mechanism related to asset managers by estimating the shape of global mutual fund flows' dependence on the VIX. The shape of the function resembles the shape of the return forecasting function $\phi_h(v)$ closely for the range of the VIX from 18 to 50. Furthermore, the signs of the loadings on the flow function exhibits evidence of flight-to-safety. For the VIX below 18, where the slope of the risk-return tradeoff is negative for stocks and positive for bonds, we conjecture that the theory of Buffa, Vayanos, and Woolley (2014) might be of help. For the range of the VIX above 50, where $\phi_h(v)$ is declining for stocks and increasing for bonds, we conjecture that cash flow news, not discount rate news are key.

Our findings are also closely linked to intermediary asset pricing theories. In Adrian and Boyarchenko (2012), intermediaries are subject to value at risk (VaR) constraints that directly link intermediaries' risk taking ability to the level of volatility. Prices of risk are a nonlinear function of intermediary leverage, which has a one-to-one relationship to the level of volatility. A similar nonlinear risk-return tradeoff is also present in the theories of He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014). We also discuss the extent to which our findings are compatible with the habit formation model of Campbell and Cochrane (1999).

The remainder of the paper is organized as follows. Section 2 provides a brief overview of the related literature. Section 3 presents evidence of the nonlinearity in the risk-return tradeoff using polynomial, spline, and kernel regressions for stocks and bonds. Importantly, we develop the sieve reduced rank regression estimator and tests of the stability of the nonlinear shape across asset classes and over time. Section 4 provides an economic analysis of the flight to safety feature in the nonlinear risk-return tradeoff, establishing a link to dynamic asset pricing and resting our findings to theories of flight to safety. Furthermore, we discuss the theoretical literature in light of our findings in detail. Section 5 concludes.

2 Related literature

Economic theory strongly suggests a risk-return tradeoff in the pricing of risky assets (Sharpe (1964), Merton (1973), Ross (1976)). An unexpected increase in riskiness should be associated with a contemporaneous drop in the asset price and an increase in expected returns. While the first half of this logic is readily verified—asset returns and volatility changes tend to be strongly negatively correlated contemporaneously—the latter prediction has been much harder to prove. Indeed, studies that have documented a positive risk return tradeoff in the time series have typically relied on the use of mixed frequency data (see Ghysels, Santa-Clara, and Valkanov (2005)), cross-sectional approaches (see Guo and Whitelaw (2006), Bali and Engle (2010)), or very long historical data (see Lundblad (2007)). A simple regression of asset returns on lagged measures of risk such as the VIX or realized volatility typically do not yield any statistically significant relationship for the risk-return tradeoff (e.g. Bekaert and Hoerova (2014) and Bollerslev, Osterrieder, Sizova, and Tauchen (2013)). In contrast, we show that there is a strong nonlinear relationship between stock and bond returns and lagged equity market volatility.

The nonlinear risk-return tradeoff that we document exhibits evidence of flight-to-safety, as the nature of the nonlinearity is reversed for stocks and bonds. This type of nonlinear relationship is broadly consistent with theories of flight-to-safety which tend to predict highly nonlinear equilibrium asset pricing relationships (Vayanos (2004), Weill (2007), Caballero and Krishnamurthy (2008), Brunnermeier and Pedersen (2009), Vayanos and Woolley (2013)). Our findings can also be rationalized within the context of intermediary asset pricing theories, which generate strongly time varying pricing of risk evolves with aggregate, endogenous volatility (Adrian and Boyarchenko (2012), Brunnermeier and Sannikov (2014), He and Krishnamurthy (2013)).

Inference on nonlinear relationships present an econometric challenge, as standard errors tend to be large, particularly when investigating tail risk in security returns. A technical contribution of this paper is to propose sieve reduced rank regressions, which combine nonparametric regressions with a reduced rank assumptions on panel data. We follow the standard method of sieve approximations to estimate the unknown functions, resulting in simple, closed-form expressions. Our empirical results consist of a number of asset pricing tests that can be implemented using standard critical values. The asymptotic properties of these tests are derived from the results of Chen, Liao, and Sun (2014) with modifications in the spirit of Hodrick (1992) to account for serial correlation in multi-horizon returns.

Our work falls within the vast literature on asset return forecasting. Seminal papers include Campbell and Shiller (1988a,b), Lettau and Ludvigson (2001), Cochrane and Piazzesi (2005), Ang and Bekaert (2007), and are nicely surveyed by Cochrane (2011). The majority of this literature focuses on forecasting returns using financial ratios or yields. While much of that literature employs linear forecasting relationships, some do model nonlinearities. Lettau and Van Nieuwerburgh (2008) present a regime shifting model for stock return forecasting. Pesaran, Pettenuzzo, and Timmermann (2006) present forecasting relationships for US Treasury bonds subject to stochastic breakpoint processes. Rossi and Timmermann (2010) document a nonlinear risk-return tradeoff in equities using boosted regression trees. To the best of our knowledge, no paper has estimated a common nonlinear forecasting relationship across different asset classes. Furthermore, our method is computationally straightforward.

Our finding that expected returns to stocks, Treasury bonds, and credit returns are forecast by a common nonlinear function $\phi_h(v)$ suggests that this function is a price of risk variable in a dynamic asset pricing model. Most markedly, we find that functions of Treasury implied volatility is not forecasting Treasury or equity returns, while functions of the VIX is forecasting both Treasury and stock returns. Based on this evidence, we estimate a dynamic asset pricing model the cross-section of stocks, bonds, and credit, and show that $\phi_h(v_t)$ is a highly significant price of risk variable for the market return, the one year Treasury return, and innovations to $\phi_h(v_t)$. These findings thus point towards joint dynamic asset pricing of stocks and bonds, as explored in linear settings by Mamaysky (2002), Bekaert, Engstrom, and Grenadier (2010), Lettau and Wachter (2010), Ang and Ulrich (2012), Koijen, Lustig, and van Nieuwerburgh (2013), and Adrian, Crump, and Moench (2014).

3 Estimation of the Nonlinear Risk-Return Tradeoff

We start this section by presenting evidence from univariate predictive regressions of stock and bond returns on VIX polynomials, presenting strong evidence of nonlinearity (subsection 3.1). We find evidence of a mirror image property: the shape of the nonlinearity of bonds mirrors inversely the shape of the nonlinearity of stocks. We document that this mirror image property not only holds for polynomial regressions, but also for nonparametric estimators such as kernel regressions or sieve regressions (subsection 3.2). This motivates us to develop a panel estimation method for the shape of the nonlinearity that allows each asset return to be an affine function of a common nonlinear function of market volatility (subsection 3.3). We label this panel estimation method sieve reduced rank regressions. This method exploits cross-section variation in excess returns to estimate the shape of the nonlinearity. We use the sieve reduced rank regressions to document that the nature of the nonlinearity is reversed when the excess return to be predicted is the equity market versus Treasuries, pointing towards flight-to-safety from stocks to bonds as equity market volatility rises above its unconditional median (subsection 3.4). We also document the robustness of the predictive relationships across forecasting horizons, Treasury maturities, and for different measures of implied volatility. Strikingly, we show that the shape of the nonlinearity is statistically indistinguishable whether it is extracted from only bonds or only stocks. We also present results for broader cross sections, including industry sorted portfolios, maturity sorted Treasury returns, and credit returns (subsection 3.6).

3.1 Suggestive Univariate Evidence from VIX Polynomials

To demonstrate the gains that can be obtained by allowing for nonlinearities, we estimate the linear regression

$$Rx_{t+h}^{i} = a_{h}^{i} + b_{h}^{i}\,(vix_{t}) + \varepsilon_{t+h}^{i}, \tag{3.1}$$

the polynomial regression

$$Rx_{t+h}^{i} = a_{h}^{i} + b_{h}^{i}(vix_{t}) + c_{h}^{i}\,(vix_{t})^{2} + d_{h}^{i}\,(vix_{t})^{3} + \varepsilon_{t+h}^{i}, \tag{3.2}$$

$$Rx_{t+h}^{i} = a_{h}^{i} + b_{h}^{i}(move_{t}) + c_{h}^{i}\,(move_{t})^{2} + d_{h}^{i}\,(move_{t})^{3} + \varepsilon_{t+h}^{i}, \tag{3.3}$$

and the augmented polynomial regressions

$$\begin{aligned} Rx_{t+h}^{i} = {} & a_{h}^{i} + b_{h}^{i}(vix_{t}) + c_{h}^{i}\,(vix_{t})^{2} + d_{h}^{i}\,(vix_{t})^{3} \\ & + bm_{h}^{i}(move_{t}) + cm_{h}^{i}\,(move_{t})^{2} + dm_{h}^{i}\,(move_{t})^{3} + \mathbf{f}^{i\prime}\,\mathbf{z}_{t} + \varepsilon_{t+h}^{i} \end{aligned} \tag{3.4}$$

separately for i representing equity market or Treasury excess returns. Here, $\mathbf{z}_t$ is a vector of predictors, and $Rx_{t+h}^{i} = (12/h)[(r_{t+1}^{i} - r_{t}^{f}) + \cdots + (r_{t+h}^{i} - r_{t+h-1}^{f})]$ denotes the continuously compounded h-month holding period return of asset i in excess of the one-month riskfree rate r_t^f (at an annual rate). For comparison, we include both equity market option-implied volatility (VIX) as well as its Treasury counterpart (MOVE).

Table 1 reports t-statistics for the coefficients of regressions (3.1) through (3.4), as well as p-values for the joint hypothesis test under the null of no predictability (H_0: $Rx_{t+h}^{i} = a_{h}^{i} + \varepsilon_{t+h}^{i}$). While the top panel of Table 1 shows results where Rx_{t+h}^{i} represents excess returns on 1-year maturity US Treasuries for forecasting horizons $h = 6$, 12, and 18 months, the bottom panel reports analogous results for excess returns on the CRSP value-weighted US equity market portfolio. Since our sample represents monthly observations from 1990:1 to 2014:9, we follow Ang and Bekaert (2007) and compute standard errors using the Hodrick (1992) correction for multihorizon overlapping observations.

The most striking features of Table 1 are the predictive gains obtained by simply augmenting the VIX with squares and cubes of itself. For 1-year Treasuries, the t-statistic on the VIX coefficient jumps from 1.91 in the linear regression to 4.13 when squares and cubes of VIX are included at the $h = 6$ month forecasting horizon, from 1.86 to 3.60 at the $h = 12$ month horizon, and from 1.13 to 3.21 at the $h = 18$ month horizon. Moreover, the coefficients on the squares and cubes of the VIX are themselves highly statistically significant – an effect that persists even with the inclusion of standard forecasting variables like the BAA/10-year Treasury default spread (DEF), the variance risk premium (VRP) following Bollerslev, Tauchen, and Zhou (2009), the 10-year/3-month Treasury Term Spread (TERM), and the (log) dividend yield (DY). The p-values also suggest strong evidence for the joint predictive content of the VIX polynomial.

Similar gains are obtained for equity market excess returns. While we can confirm the findings of Bekaert and Hoerova (2014) and Bollerslev, Osterrieder, Sizova, and Tauchen (2013) that the VIX itself does not (linearly) forecast excess stock market returns, we document marked improvements in predictability when squares and cubes of the VIX are included: p-values for the joint test of no predictability drop from 0.316 for the linear regression case to 0.007 for the VIX polynomial case at the $h = 6$ month horizon, from 0.460 to 0.032 at the $h = 12$ month horizon, and from 0.439 to 0.088 at the $h = 18$ month horizon. For the short forecast horizon $h = 6$, we note further that evidence for the VIX polynomial's predictability remains even after the inclusion of other forecasting variables, while for longer horizons, the predictive content of VIX polynomials appears to subside.

The MOVE is an analogous portfolio of yield curve weighted options written on Treasury futures. To the extent that some form of segmentation between Treasury and equity markets could give rise to separate pricing kernels for bonds and stocks, one may surmise that excess returns in either market reflect compensation for exposure to different types of volatility or uncertainty risk. Somewhat surprisingly, however, Table 1 shows that this is not the case. Whereas the VIX polynomials exhibit t-statistics at times above five, t-statistics on MOVE polynomials coefficients are struggling to exceed one. The p-values show that regressions on the MOVE cannot be statistically distinguished from regressions on a constant.

A final noteworthy feature of Table 1 are the signs on the constant and coefficients of the VIX polynomials for Treasuries compared to equities. While the coefficients on the VIX, VIX^2, and VIX^3 alternate as $(\hat{b}_h^i > 0)$, $(\hat{c}_h^i < 0)$, and $(\hat{d}_h^i > 0)$ for i = Treasuries, they are exactly the opposite for equities across all forecasting horizons. The same is true for the intercepts: while the intercepts in the Treasury regressions are all negative when VIX polynomials are included, the intercepts for the equity regressions are all positive. In contrast, the linear VIX specification appears to make no signed distinction between Treasury and equity market excess returns and is instead reporting a statistically insignificant relationship between the VIX and future excess returns across all horizons.

To examine predictability for various Treasury maturities and equity market returns across many forecast horizons, Figure 2 plots p-values by h ranging from 1 to 24 months for both the linear regression (3.1) (thin line) and polynomial regression (3.2) (thick line). Several noteworthy features emerge from the figure. First, the predictive gains that result from allowing for VIX nonlinearities, as measured by the distance between the thin and thick lines, are substantial for all horizons h, Treasury maturities, and equity returns. In particular, the polynomial specification dominates the linear one across all Treasury and equity excess returns for horizons $h = 3, \dots, 24$. Second, the polynomial specification is strongly rejected at the 5% level (the thick line falls below the dashed line) for short-maturity Treasuries and for a wide range of forecast horizons h. Third, as Treasury maturities lengthen, VIX polynomial predictability begins to wane as the thick line gradually shifts upward and becomes insignificant for 10-year Treasuries. Fourth, for equity market returns, the null of no predictability is rejected at the 5% level for forecast horizons $h = 3, \ldots, 15$ months.

As a robustness check, we examine to what extent the VIX's predictive results are driven by the 2008 financial crisis. Figure 3 repeats the exercise of Figure 2 with a sample spanning only 1990:1 to 2007:7. On this pre-crisis sample, 1-year Treasuries are still strongly predicted by VIX polynomials across horizons $h = 3, \ldots, 24$ while outperforming the linear specification in all panels. As Treasury maturities increase, VIX polynomial predictability appears stronger than even in the full sample. 10-year Treasuries, in particular, are showing signs of predictability at longer horizons, which contrasts with the result on the sample ending in 2014. On the other hand, equity market return predictability deteriorates and only marginally rejects the null of no predictability in favor of the polynomial VIX specification. In relative terms, however, we again note the gains from allowing the VIX to nonlinearly predict returns are substantial compared to the linear specification across both Treasuries and equities.

3.2 Motivation of Sieve Reduced Rank Regressions

The preceding results showed that polynomials – rather than linear functions – of the VIX have important predictive power for future excess stock and bond returns. But instead of accepting a cubic VIX polynomial as the true data generating process for excess returns, we conjecture that the polynomials provide an approximation to some general nonlinear relationship between equity implied volatility and future excess stock and bond returns. To test this conjecture, we nonparametrically estimate the relationship between the VIX and future excess stock and bond returns via the method of sieves, which facilitates intuitive comparisons to polynomial regressions. To motivate our nonparametric sieve estimation framework, fix asset i and forecast horizon h and consider

$$Rx_{t+h}^{i} = \phi_{h}^{i}\left(v_{t}\right) + \varepsilon_{t+h}^{i}, \tag{3.5}$$

where $v_t = vix_t$. Equation (3.5) effectively replaces the polynomial $(a_h^i + b_h^i v_t + c_h^i v_t^2 + d_h^i v_t^3)$ from before with an unknown function $\phi_h^i(v_t)$.

To estimate the function $\phi_h^i(\cdot)$ nonparametrically, we assume that $\phi_h^i \in \Phi$, where Φ is a general function space of sufficiently smooth functions. In practice, estimation over the entire function space Φ is challenging because it is infinite dimensional. In settings like these, the method of sieves (e.g., Chen (2007)) proceeds instead by estimation on a sequence of m-dimensional approximating spaces $\{\Phi_m\}_{m=1}^{\infty}$. We say that $\{\Phi_m\}_{m=1}^{\infty}$ is a valid sieve for Φ if it is nested (i.e. $\Phi_m \subset \Phi_{m+1} \subset \cdots \subset \Phi$) and eventually becomes dense in Φ (i.e. $\cup_{m=1}^{\infty}\Phi_m$ is dense in Φ). Letting $m = m_T \to \infty$ slowly as the sample size $T \to \infty$, the idea then is that the spaces Φ_{m_T} grow and increasingly resemble Φ, so that the least squares solution

$$\hat{\phi}_{m_T,h}^i \equiv \arg\min_{\phi \in \Phi_{m_T}} \frac{1}{T} \sum_{t=1}^{T} \left(Rx_{t+h}^i - \phi(v_t)\right)^2 \tag{3.6}$$

converges to the true unknown function $\phi_h^i \in \Phi$ in (3.5) in some suitable sense.

For our choice of Φ_m we use the space spanned by linear combinations of m B-splines of the VIX. Thus, any element $\phi_m \in \Phi_m$ may be written as $\phi_m(v) = \sum_{j=1}^{m} \gamma_j \cdot B_j(v)$ where $\gamma_j \in \mathbb{R}$ for $j = 1, \ldots m$, v is a value in the support of the VIX, and B_j is the j^{th} B-spline (see Appendix for further details). B-splines have a number of appealing features such as well-established approximation properties and substantial analytical tractability. This is because for fixed m, the solution to the least squares problem (3.6) is simply the OLS estimator on B-spline coefficients $\gamma^h = \left(\gamma_1^h, \ldots \gamma_m^h\right)'$:

$$\hat{\gamma}^h = \left(X_m X_m'\right)^{-1} X_m \mathrm{Rx},, \tag{3.7}$$

where $\mathrm{Rx}^i = \left(Rx_{1+h}^i, \ldots, Rx_{T+h}^i\right)'$ and X_m is the $(m \times T)$ matrix of predictors with j^{th} row equal to $[B_j(\mathrm{VIX}_t), \ldots, B_j(\mathrm{VIX}_T)]$. Therefore, for fixed m, the solution to (3.6) becomes

$$\hat{\phi}_{m,h}^i(v) = \sum_{j=1}^{m} \hat{\gamma}_j^h \cdot B_j(v). \tag{3.8}$$

Equation (3.8) makes clear that the simple polynomial specification introduced in the previous section may be thought of as an alternative nonparametric estimate of $\phi_h^i(\cdot)$ using powers v^j instead of B-splines $B_j(v)$. However, this approach was informal in the sense that the choice of the maximum degree of polynomial was not made with relation to the sample size. Instead, we think of the number of basis functions $m = m_T$ as growing to infinity at some optimal rate that depends on the sample size T.[1]

The top half of Figure 4 shows various estimates for $\phi_h^i(v)$, where $h = 6$ and i refers to either 1-year Treasury excess returns (dashed line) or equity market excess returns (solid line) over the full sample period from 1990:1 to 2014:9. In the left graph, we show the cross-validated sieve B-spline estimates $\hat{\phi}_{m_T,h}^i(v)$ (equation (3.8)), whereas the middle graph shows the functional form implied by the simple polynomial specification of the previous section. The estimated functional forms in both the left and middle graphs are very similar, implying that the cubic polynomial choice in the previous section provided a reasonable first pass at investigating the nonlinear relationship. As a further robustness check, the right panel shows the estimated function based on a nonparametric kernel regression, which shows qualitatively similar impression of $\phi_h^i(v)$ for stocks and bonds.

Figure 4 also demonstrates another noteworthy empirical regularity. If we compare $\hat{\phi}_h^i(v)$ using either equity returns or bond returns as test assets it appears that they are related by a simple scale and reflection transformation. This could already be deduced from the alternating coefficient signs from the polynomial regressions, and is now additionally confirmed with two nonparametric estimators. Moreover, the bottom panel of Figure 4 shows that the mirror image relationship between $\phi_h^{\text{Treasuries}}(v)$ and $\phi_h^{\text{Stocks}}(v)$ existed prior to the 2008 financial crisis and is therefore not an artifact of a few extreme observations. Instead, the crisis is merely helpful in identifying $\phi_h^i(v)$ for large v.

[1] In particular, it can be shown that m behaves very much like a bandwidth parameter in that it is chosen to optimally trade notions of bias and variance: heuristically, if m is too small, Φ_m is too small relative to Φ, which causes bias, and if m is too big, it results in overfitting. In the remainder of the paper, we follow the existing literature in sieve estimation and choose m_T by leave-one-out cross-validation. See, e.g. Li and Racine (2007).

We interpret this finding as strongly suggestive that equity market and Treasury excess returns load on a *common* $\phi_h(v)$ function, up to location, scale, and reflection transformations. In this case, we show next that $\phi_h(\cdot)$ could then be estimated jointly across assets rather than estimating univariate regressions equation by equation, as was done above. This has the benefit of allowing Treasury returns across multiple maturities as well as the equity market excess returns to jointly inform the estimate of the common $\phi_h(\cdot)$, thereby exploiting information in the cross-section of asset returns.

3.3 Derivation of Sieve Reduced Rank Regressions

In this subsection, we formalize the intuition of a common volatility function $\phi_h(v)$ by introducing a reduced-rank, sieve-based procedure which produces a nonparametric estimate of $\phi_h(v)$ under only weak assumptions. The novelty of our approach is that we use cross-sectional information across assets to better inform our estimate of this function. We label our approach "sieve reduced-rank regression" (SRRR) as it combines the cross-sectional restrictions implied by a reduced-rank assumption with the flexibility of a nonparametric sieve estimator. We will see that the estimator is conveniently available in closed form and hypothesis tests rely on standard critical values.

Suppose we observe excess returns on $i = 1, \ldots, n$ assets that follow

$$Rx_{t+h}^{i} = a_h^i + b_h^i \cdot \phi_h(v_t) + \varepsilon_{t+h}^{i}. \tag{3.9}$$

Here, a_h^i and b_h^i are asset-specific shift and scale parameters, $\phi_h(\cdot)$ is the same for all assets, and $v_t = vix_t$. This specification can be compared with equation (3.5), which held that $Rx_{t+h}^{i} = \phi_h^i(v_t) + \varepsilon_{t+h}^{i}$. Thus in the univariate regressions from the previous section, $\phi_h^i(v_t)$ was estimated separately for each asset i, with no cross-asset restrictions imposed. In contrast, the specification (3.9) implies that the same function $\phi_h(v_t)$ forecasts returns across assets, which amounts to the restriction $\phi_h^i(v_t) = a_h^i + b_h^i \cdot \phi_h(v_t)$.

If we take the same approach as in the univariate specification we can rewrite this equation as

$$Rx_{t+h}^{i} = a_h^i + b_h^i \left(\gamma_h' X_{m,t}\right) + f_h^i z_t + \tilde{\varepsilon}_{t+h}^i \tag{3.10}$$

where

$$\tilde{\varepsilon}_{t+h}^i = \varepsilon_{t+h}^i + b_h^i \cdot \left(\phi_h\left(v_t\right) - \gamma_h' X_{m,t}\right). \tag{3.11}$$

$\tilde{\varepsilon}_{t+h}^i$ in equation (3.11) is composed of two terms. The first term is a standard error term from the original regression equation. The second term represents the approximation error of the true nonlinear function and the best approximation from the space Φ_m. As m grows with the sample size this approximation error vanishes in the appropriate sense.

If we stack equation (3.10) across n assets we obtain

$$\mathrm{Rx}_{t+h} = \mathbf{a}_h + \mathbf{A}_h X_{m,t} + \mathbf{F}_h Z_t + \tilde{\varepsilon}_{t+h}, \qquad A_h = \mathbf{b}_h \gamma_h' \tag{3.12}$$

where $\mathbf{a}_h = (a_h^1, \ldots, a_h^n)'$, $b_h = (b_h^1, \ldots, b_h^n)'$, $\mathrm{Rx}_{t+h} = \left(Rx_{t+h}^i, \ldots, Rx_{t+h}^n\right)'$ and $\tilde{\varepsilon}_{t+h} = \left(\tilde{\varepsilon}_{t+h}^1, \ldots, \tilde{\varepsilon}_{t+h}^n\right)'$. For any fixed m, equation (3.12) is a reduced-rank regression where A_h is assumed to be of rank one.[2] The parameters $\left(\mathbf{a}_h', \mathbf{b}_h', \gamma_h', \mathbf{F}_h'\right)'$ may be estimated in closed form. However, in order to separately identify $\mathbf{a}_h$ and $\mathbf{b}_h$ additional restrictions must be imposed. In our empirical analysis we impose the normalization $\phi_h(0) = 0$ and $b_h^1 = b_h^{\mathrm{MKT}} = 1$. The first restriction allows us to identify the constant term for each asset, while the second implies that the market return is our reference asset.

To describe the estimation procedure, let $\hat{\mathbf{a}}_{h,\mathrm{ols}}$ $(n \times 1)$, $\hat{A}_{h,\mathrm{ols}}$ $(n \times m)$ and $\hat{F}_{h,\mathrm{ols}}$ $(n \times p)$ be the stacked OLS estimates and W_2 a symmetric, positive-definite weight matrix. In our empirical application, we set W_2 to a diagonal matrix that scales excess returns by the inverse of their standard deviation to avoid overweighting high-variance assets in the

[2]See Reinsel and Velu (1998) for a general introduction. Examples of parametric reduced-rank regressions are systems-based cointegration analysis (see e.g. Johansen (1995)), beta representations of dynamic asset pricing models (see e.g. Adrian, Crump, and Moench (2013, 2014)), and bond return forecasting (see e.g. Cochrane and Piazzesi (2008)).

estimation. Then,

$$\hat{\mathbf{b}}_h = \frac{\tilde{\mathbf{b}}_h}{\tilde{b}_h^1}, \qquad \hat{\gamma}_h = \tilde{\gamma}_h \cdot \tilde{b}_h^1, \qquad \left[\hat{\mathbf{a}}_h \;\vdots\; \hat{\mathbf{F}}_h\right] = \left[\hat{\mathbf{a}}_{h,\text{ols}} \;\vdots\; \hat{\mathbf{F}}_{h,\text{ols}}\right] + \left(\hat{A}_{h,\text{ols}} - \tilde{\mathbf{b}}_h \tilde{\gamma}_h'\right) X_m Z' \left(ZZ'\right)^{-1},$$

where $Z = (Z_1, Z_2, \ldots, Z_T)$, $\tilde{\mathbf{b}}^h = W_2^{-1/2} L$, $\tilde{\gamma}^h = \hat{A}_{h,\text{ols}}' W_2 \hat{\mathbf{b}}_h$ and L is the eigenvector associated with the maximum eigenvalue of the matrix $\hat{A}_{h,\text{ols}} \left(X_m M_Z X_m'\right) \hat{A}_{h,\text{ols}}'$ where $M_Z = I_T - Z\left(ZZ'\right)^{-1} Z'$. If it were the case that $\tilde{\varepsilon}_{t+h} \sim_{iid} \mathcal{N}\left(0, W_2^{-1}\right)$ and m was fixed, then $\left(\hat{\mathbf{a}}_h', \hat{\mathbf{b}}_h', \hat{\gamma}_h', \text{vec}\left(\hat{\mathbf{F}}_h\right)'\right)'$ would be the maximum likelihood estimates of $\left(\mathbf{a}_h', \mathbf{b}_h', \gamma_h', \text{vec}\left(\mathbf{F}_h\right)'\right)'$.

In this paper there are three primary hypotheses of interest:

$$\begin{aligned}
&\mathbb{H}_{1,0} : b_h^j\, \phi_h = 0 \qquad &&\mathbb{H}_{1,A} : b_h^j\, \phi_h \neq 0 \\
&\mathbb{H}_{2,0} : \mathbf{b}_h \phi_h = \mathbf{0}_n \qquad &&\mathbb{H}_{2,A} : \mathbf{b}_h \phi_h \neq \mathbf{0}_n \\
&\mathbb{H}_{3,0} : \phi_h\left(\bar{v}\right) = 0 \qquad &&\mathbb{H}_{3,A} : \phi_h\left(\bar{v}\right) \neq 0
\end{aligned} \tag{3.13}$$

The first hypothesis tests the null that ϕ_h does not predict excess returns Rx_{t+h}^j of asset j, while allowing it to predict another asset $i \neq j$. This test replaces t-tests on the loadings b_h^j because the scale of $\mathbf{b}_h$ cannot be determined separately from the scale of ϕ_h, which prompted our normalization $b_h^1 = 1$. This means, in particular, that a test of $b_h^1 = 0$ cannot be conducted, motivating our test on the product $b_h^j\, \phi_h$. In finite samples when the number of sieve expansion terms is fixed at some m, we show below that $\mathbb{H}_{1,0}$ is tested with a standard $\chi^2(m+1)$ test on the product $b_h^j\, (\gamma_h^1, \ldots, \gamma_h^m)$.[3] This represents an additional convenient aspect of the sieve-based nonparametric procedure, since it allows us to test hypotheses about predictability in effectively the same way as a parametric joint test of significance. The second hypothesis is a joint test of significance for whether the whole cross-section of test assets jointly loads on ϕ_h. Finally, the third hypothesis is a comparison of whether the function $\phi_h(\cdot)$ is different from zero at a fixed value $\bar{v}$. By inverting a test of this hypothesis for different values of $\bar{v}$ we are able to construct pointwise confidence intervals

[3] The use of a χ^2 test is a small sample correction. See Crump, Hotz, Imbens, and Mitnik (2008) and the references therein.

for the unknown function.

Proposition 1. *Under regularity conditions given in the Appendix*

$$\left[vec\left(\hat{b}_h^j\hat{\gamma}_h\right)'\ \hat{\mathcal{V}}_1\ vec\left(\hat{b}_h^j\hat{\gamma}_h\right)-(m+1)\right]/\sqrt{m+1}\ \rightarrow\ _{d,\mathbb{H}_{1,0}}\ \mathcal{N}(0,1)$$

$$\left[vec\left(\hat{\mathbf{b}}_h\hat{\gamma}_h\right)'\ \hat{\mathcal{V}}_2\ vec\left(\hat{\mathbf{b}}_h\hat{\gamma}_h\right)-(m+n-1)\right]/\sqrt{m+n-1}\ \rightarrow\ _{d,\mathbb{H}_{2,0}}\ \mathcal{N}(0,1)$$

$$\frac{\hat{\phi}_{h,m}(\bar{v})-\phi(\bar{v})}{\hat{\mathcal{V}}_3}\ \rightarrow\ _{d,\mathbb{H}_{3,0}}\ \mathcal{N}(0,1)$$

as $T\rightarrow\infty$, *where* $\hat{\mathcal{V}}_1$, $\hat{\mathcal{V}}_2$, *and* $\hat{\mathcal{V}}_3$ *are defined in the Appendix and* $\rightarrow_{d,\mathbb{H}_0}$ *signifies convergence in distribution under the hypothesis* $\mathbb{H}_0$.

Proposition 1 provides the appropriate limiting distributions to conduct the asset pricing tests for the paper and relies on extensions of the Hodrick (1992) standard errors to the n asset sieve reduced rank setting.[4] The underlying distribution theory relies on the procedures of Chen, Liao, and Sun (2014). Comprehensive details about the empirical implementation are provided in the Appendix.

3.4 Estimation of Sieve Reduced Rank Regressions

Our main empirical findings using the SRRRs (3.5) for the market return and the maturity sorted bond returns are presented in Table 2. As we had seen in the univariate VIX polynomial regressions, substantial improvements are gained when allowing the cross-section of market returns and maturity sorted bond returns to depend on the VIX nonlinearly: Whereas the VIX does not linearly forecast excess returns in panel (1), the nonlinear forecasting relationship for stocks and bonds is highly significant in panel (2). Moreover, panel (3) shows that the nonlinear forecasting factor is robust to the inclusion of common predictor variables (the default spread DEF, the variance risk premium VRP, the term spread TERM,

[4] An extension of the Hodrick (1992) to our setting was necessary to remove the serial dependence induced by estimating multi-horizon returns with overlapping data. We therefore construct valid standard errors from "reverse" regressions under a weak assumption of covariance stationarity. See also Wei and Wright (2013).

and the log dividend yield DY). Furthermore, the significant predictability is present in the 1990-2007 period which excludes the financial crisis (Table 3).

Examining Table 2 in more detail, we see that the market return is most strongly predicted at the six month horizon at the one percent level. Per construction, the coefficient on the market return is 1. Overall, the strongest predictability appears for shorter-maturity bonds, as the one year bond return is highly significantly predicted at the one percent level for the 6, 12, and 18 month horizons. Interestingly, the significance is unchanged when even the variance risk premium (a volatility measure constructed from the VIX) is included, suggesting that the nonlinear forecasting factor is unrelated to VRP. Longer maturity Treasuries such as the five year or the ten year bond return tend to be somewhat less significant at longer horizons, but their significance is actually aided by the inclusion of the other predictor variables. The sign on all of the Treasury variables is negative whereas the market return is positive, indicating the flight to safety feature that is strongest for liquid short-maturity Treasuries. While individual coefficient significance was tested by $\mathbb{H}_{1,0}$, joint significance for the function $\phi_h(vix_t)$ in the cross-section of excess returns is tested with $\mathbb{H}_{2,0}$. Again the joint test provides strong justification for nonlinearities $\phi_h(vix_t)$ across all forecasting horizons, whereas the linear VIX specification cannot be statistically distinguished from regressions on constants a_h^i.

For the pre-crisis period 1990-2007 presented in Table 3, the equity market is only significant at the longer 18 month horizon, and not at the shorter 6 and twelve months horizons. On the other hand, Treasury returns are again very significant, and result in stronger rejections for shorter maturity Treasuries. In particular, the one-year Treasury is significant at the one percent level across all specifications and forecasting horizons both the pre-crisis and full samples, which include the specifications with common predictor variables. Furthermore, test $\mathbb{H}_{2,0}$ confirms that $\phi_h(vix_t)$ is a strong predictor of excess returns jointly across all test assets and horizons. We highlight again the gains obtained by allowing excess returns to nonlinearly depend on the VIX. Most importantly, the mirror image property between stock and bond returns is revealed in the pre-crisis period as suggested by the coefficient signs, although we are careful to point out that for specification (3), coefficient signs are difficult to interpret when ϕ_h interacts with the control predictors. In sum, we find that the nonlinear sieve reduced rank regressions reveal the mirror image property, which is not manifested for the linear VIX regressions.

3.5 Out-of-Sample Evidence

The preceding tables strongly suggest that expected excess returns for the market and Treasury returns are driven by a common nonlinear function of the VIX, i.e. $E_t[Rx^i_{t+h}] = a^i_h + b^i_h \phi_h(vix_t)$. We next examine the extent to which this relation holds out of sample by studying the risk-adjusted returns of a portfolio that uses the forecasting information $E_t[Rx^i_{t+h}] = a^i_h + b^i_h \phi_h(vix_t)$. Intuitively, the risk-adjusted returns of such a portfolio should reflect the compensation that investors receive for taking on equity market risk and selling Treasuries in moderate- to high volatility periods.

To this end, we split our monthly sample of excess returns $\mathrm{Rx}_t = (Rx^{MKT}_t, Rx^{cmt1}_t, \ldots, Rx^{cmt30}_t)'$ and the VIX for $t = 1, \ldots, T$ into an initial in-sample estimation period $t = 1, \ldots, t^*$, and an out-of-sample forecasting period $t = (t^* + 1), \ldots, T$. Using in-sample data from $t = 1, \ldots, t^*$, we estimate (3.9) via sieve reduced rank regressions and form the out-of-sample joint forecast of excess returns $E_{t^*}[Rx_{t^*+h}]$ for $h = 6$ months ahead. Given the forecast, we form standard risk-weighted portfolios with weights $\omega_{t^*} = V^{-1}_{t^*} E_{t^*}[Rx_{t^*+h}]$, where risk weights $V^{-1}_{t^*}$ are the unconditional variances of excess returns using data from $t = 1, \ldots, t^*$.[5] The portfolio is held for the $h = 6$ months, and then the in-sample period is expanded to $t = 1, \ldots, (t^* + h)$, yielding a new out-of-sample forecast $E_{t^*+h}[Rx_{t^*+2h}]$ and new portfolio weights ω_{t^*+h}. Note that $V^{-1}_{t^*}$ is not rolled forward in order to isolate the effect of the forecast $E_{t^*+h}[Rx_{t^*+2h}]$. The process is iterated every $h = 6$ months, yielding pseudo

[5] Since our interest is in evaluating the forecasting performance of $E_t[Rx_{t+h}]$, we hold $V^{-1}_{t^*}$ constant throughout to facilitate comparisons of alternative ways to compute the "numerator" $E_t[Rx_{t+h}]$. Intuitively, scaling excess return forecasts by the inverse of variance helps put Treasury excess returns (which have small variance) and MKT returns on equal footing. Alternative choices of V_{t^*} yield very similar results.

out-of-sample excess returns for $t = (t^* + 1), \ldots, T$.

Table 4 column (1) shows the annualized Sharpe ratios of the pseudo out-of-sample portfolio returns thus obtained for various in-sample cutoffs $(t^*/T) = 0.4$, 0.5, and 0.6. Remaining columns show the Sharpe ratios of alternative portfolios over the same out-of-sample period. In particular, columns (2) and (3) represent Sharpe ratios of portfolios formed with identical risk weights $V_{t^*}^{-1}$ as above, but with differing excess return forecasts: Column (2) uses linear reduced rank forecasts $E_t[Rx_{t+h}^i] = a_h^i + b_h^i \; vix_t$, while column (3) uses running mean forecasts $E_t[Rx_{t+h}^i] = a_h^i$. Thus, the portfolios in columns (1) through (3) differ only in how the "numerator" $E_t[Rx_{t+h}^i]$ is formed, since the "denominator" $V_{t^*}^{-1}$ is held fixed. In contrast, column (4) shows Sharpe ratios for an equally-weighted portfolio of market and Treasury excess returns. The remaining columns place full weight on the indicated asset.

The results show that the portfolios formed on the sieve reduced rank forecasts $E_t[Rx_{t+h}^i] = a_h^i + b_h^i \phi_h(vix_t)$ receive higher risk-adjusted returns relative to the displayed alternatives. The sole exception is the risk-adjusted return on the 1-year Treasury over the out-of-sample period from 1999 to 2014. However, for the other in-sample cutoffs, the SRRR forecasts yield nontrivially higher Sharpe ratios. Interestingly, the linear VIX forecast and the unconditional excess return forecast yield very similar risk-adjusted returns. This corroborates our in-sample findings of Tables 2 and 3, which showed that the linear VIX specification cannot be statistically distinguished from a regression on a constant. Hence it is perhaps not surprising that using the linear VIX in forecasting excess returns in the portfolio decision problem does not result in higher ex-post risk-adjusted returns than the portfolio formed on the running mean forecast $E_t[Rx_{t+h}^i] = a_h^i$.

It is also helpful to take a conditional view by examining *when* the SRRR based portfolio is earning its risk-adjusted returns. Figure 8 shows the cumulative returns of the SRRR based portfolio over the out-of-sample period $(t^* + 1), \ldots, T$, where (t^*/T) splits the sample in half, alongside the cumulative returns to an investment in the market and a levered investment in the 1-year Treasury portfolio. (All returns in the figure are scaled (levered) to have the same ex-post variance of the market excess return.) In particular, we find that the SRRR based portfolio is earning its highest returns during high-volatility periods. In contrast, during the low-volatility expansionary periods from 2003 to 2007 and 2012 to 2014, the SRRR based portfolio did not accumulate significant returns. This is consistent with the regression functions in Figure 1, which shows a relatively unresponsive relationship between uncertainty and expected returns for low volatility periods, and a very strong relationship between uncertainty and expected returns for the market and Treasuries during high volatility periods. Hence the nonlinear forecasting relationship $E_t[Rx^i_{t+h}] = a^i_h + b^i_h \phi_h(vix_t)$ is most helpful in the portfolio formation stage during high volatility periods.

Finally we note that the gains to the SRRR based portfolio returns are not merely accrued during the 2007-2009 financial crisis. Indeed, the VIX frequently exceeded it's unconditional median of about 18 in the post-crisis period in response to uncertainty about the robustness of the recovery and events surrounding the European sovereign debt crises. Correspondingly, the SRRR based portfolio accrued substantial gains over the 2009-2012 period as well.

3.6 Evidence Using Broader Cross-Sections of Assets

While our results so far have focused on the aggregate stock return and maturity sorted Treasury bond portfolios, we now estimate the SRRR on a broader set of test assets in order to improve the economic insight. We use the 12 industry sorted stock portfolios from Kenneth French's website[6] and the industry and rating sorted investment grade credit returns from Barclays. We also continue to include the maturity sorted Treasury bond portfolios.

Figure 5 displays the results of hypothesis tests $\mathbb{H}_{2,0}$. The height of bar j represents the point estimate of $\hat{b}^j_h$ for $h = 6$. For each $j = 1, \ldots, n$, the color of the bar denotes the significance of the associated $\hat{b}^j_h$ coefficient based on the results in Proposition 1. The figure shows that the majority of stock and Treasury portfolios load significantly on $\phi_h(vix_t)$. Manufacturing, known to be highly procyclical, has a strong positive exposure while the

[6] http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

only equity portfolios that appear invariant to the volatility factor are non-durables, energy, utilities, and healthcare, which are known to represent inelastic sectors. Most strikingly, the only assets with negative exposures to the volatility factor $\hat{\phi}_h(v)$ are Treasury portfolios and AAA corporate bonds, which is consistent with a flight to safety interpretation. Furthermore, the results confirm our previously reported findings, which showed the strongest predictive content for nonlinear VIX functions at shorter maturity Treasuries. For $h = 6$, however, the corporate bond loadings are not statistically significant.

Next, Figure 6 shows $\hat{E}_t[Rx^i_{t+h}] = \hat{a}^i_h + \hat{b}^i_h \hat{\phi}_h(v)$ for each of the $i = 1, \ldots, 26$ portfolios and horizon $h = 6$ months. The dashed lines in blue represent assets with a negative $\hat{b}^i_h$ exposure to $\hat{\phi}_h(v)$, while the solid lines in red denote positive $\hat{b}^i_h$ exposures. To differentiate our reference asset, the black line denotes the market excess return estimate, whereas the gray dashed line represents the short-maturity Treasury return. We note again that all dashed assets (blue and gray) are Treasury returns and the AAA corporate bond return and are distinguished by their negative loadings on $\hat{\phi}_h(v)$.

Finally, we examine the plausibility of the affine structure in (3.9) by estimating $Rx^i_{t+h} = a^i_h + b^i_h \phi_h(v_t) + \varepsilon^i_{t+h}$ by SRRR separately for i ranging equities and i ranging over bonds. That is, we estimate sieve reduced rank regressions, where the left-hand side variables are the excess return on the equity market and 11 industry portfolios, which yields an estimate of $\hat{\phi}^{\text{Stocks}}_h(\cdot)$.[7] Next, we repeat the regression, but where i includes only the seven maturity-sorted Treasury portfolios, yielding an estimate of $\hat{\phi}^{\text{Treas}}_h(\cdot)$. Under the common $\phi_h(v)$ assumption implicit in (3.9), $\phi^{\text{Stocks}}_h(\cdot)$ and $\phi^{\text{Treas}}_h(\cdot)$ should be equivalent in the population when identified separately from stocks and bonds, up to a location and scale parameter. We find the location and scale parameters by regressing $\hat{\phi}^{\text{Stocks}}_h(v)$ on $\hat{\phi}^{\text{Treas}}_h(v)$ for a range of v in the support of the VIX. The result of this exercise is plotted in Figure 7 which shows $\hat{\phi}^{\text{Stocks}}_h(v)$ along with the location- and scale-shifted $\hat{\phi}^{\text{Treas}}_h(v)$. The figure clearly supports our conjecture of a common $\phi_h(v)$ function when estimation error is taken into account.

[7] The 12th industry portfolio is omitted because the market portfolio is included in the regression.

4 Economics of Flight-to-Safety

We now turn to the economic interpretation of the nonlinear risk-return tradeoff. We first establish the link of the SRRR model to dynamic asset pricing theories (subsection 4.1). We empirically show that the cross-sectional dispersion of the forecasting slopes b_h^i from the SRRR are related to risk factor loadings on the market return, the one year Treasury return, and the nonlinear volatility function $\phi_h(v_t)$. This is evidence in favor of a dynamic pricing kernel where $\phi_h(v_t)$ is a price of risk variable. We next turn to asset pricing theories that are giving rise to time varying effective risk aversion as a nonlinear function of market volatility. The first types of theories that we discuss feature flight-to-safety (subsection 4.2) that leads to time varying pricing of risk because asset managers are subject to withdrawal after poor performance. The theory of Vayanos (2004) gives rise to a dynamic pricing kernel that is qualitatively similar to our estimated $\phi_h(v_t)$ relationship. We also verify empirically that global mutual fund flows exhibit the nonlinear (contemporaneous) relationship to the VIX as the expected returns do. We then review intermediary asset pricing theories that link the pricing of risk to the level of volatility, for example due to the VaR constraint of Adrian and Boyarchenko (2012) (subsection 4.3). We also discuss the extent to which habit formation theories might explain the shape of $\phi_h(v_t)$ (subsection 4.4). Finally, we show in subsection 4.5 that the VIX also forecasts macroeconomic activity in a nonlinear fashion when the crisis period is included. We conjecture that this relationship explains the downward slope of $\phi_h(v_t)$ for very high levels of the VIX.

4.1 Dynamic Asset Pricing

Equation (3.9) shows that expected returns are affine functions of $\phi_h(v_t)$ with intercept a^i and slope b^i. Asset pricing theory suggests that these intercepts and slopes are cross-sectionally related to risk factor loadings (see Sharpe (1964), Merton (1973), Ross (1976)). In particular, an equilibrium pricing kernel with affine prices of risk, as for example presented by Adrian, Crump, and Moench (2014), would suggest that

$$E_t[Rx^i_{t+h}] \;=\; \alpha^i_h + \beta^i_h \left(\lambda_0 + \lambda_1 \phi_h(v_t) + \Lambda_2 x_t\right). \tag{4.1}$$

In this expression, β^i_h denotes a $(1 \times K)$ vector of risk factor loadings, λ_0 is the $(K \times 1)$ vector of constants for the prices of risk, λ_1 is the $(K \times 1)$ vector defining how prices of risk vary as a function of $\phi_h(v_t)$, and Λ_2 is the $(K \times p)$ matrix mapping defining how the price of risk depends on p additional risk factors x_t. The expression also allows for a pricing error α^i_h, representing deviations from no-arbitrage. Equation 4.1 is the beta representation of expected returns when the pricing kernel is of an essentially affine form (see Duffee (2002)).

Equation (4.1) has time series and cross-sectional predictions, which in turn can be linked to alternative theories of time varying pricing of risk. We start with an investigation of the cross-sectional predictions. In comparison to the SRRR model of equation (3.9), the asset pricing theories of equation (4.1) put the following constraints on the intercept and slope:

$$a^i_h \;=\; \alpha^i_h + \beta^i_h \lambda_0 \tag{4.2}$$

$$b^i_h \;=\; \beta^i_h \lambda_1. \tag{4.3}$$

In order to show that the cross-sectional dispersion of a^i_h and slope b^i_h is compatible with such asset pricing restrictions, we proceed in two steps. We first estimate the unrestricted panel forecasting relationship $Rx^i_{t+h} = a^i_h + b^i_h \phi_h(v_t) + \varepsilon^i_{t+h}$ by sieve reduced rank regression, yielding a cross-section of parametric estimates of a^i_h and b^i_h and a single nonparametric estimate of $\phi_h(v_t)$. Here, $i = 1, \ldots, n$ ranges over the CRSP market excess return, maturity-sorted Treasury excess returns, industry-sorted portfolio excess returns, and ratings and industry sorted corporate bond excess returns. Next, we estimate prices of risk as well as risk factor exposures according to the dynamic asset pricing restrictions in (4.1). That is, we estimate jointly $Rx^i_{t+h} = (\alpha^i_h + \beta^i_h \lambda_0) + \beta^i_h \lambda_1 \phi_h(v_t) + \beta^i u_{t+h} + \varepsilon^i_{t+h}$, where the estimate of $\phi_h(v_t)$ is taken as given from the unrestricted first step regression, and where $u_{t+h} \equiv Y_{t+h} - E_t[Y_{t+h}]$

represents vector autoregression innovations to the risk factors Y_t consisting of the market return, the one-year Treasury return, and the nonlinear volatility factor $\phi_h(v_t)$. Results of the regressions are given in Table 5, while Figure 9 shows the cross-sectional relationships between a_h^i and $\alpha_h^i + \beta_h^i \lambda_0$ and between b_h^i and $\beta_h^i \lambda_1$ for $h = 1$.[8]

Table 5 shows that the industry portfolios and credit portfolios are highly significantly exposed to equity market risk, while few if any of the Treasury returns loads significantly on the market return. The Treasury and corporate bond returns are significantly exposed to the one year Treasury return and also to innovations to the nonlinear volatility factor, although to a lesser extent. Importantly, the prices of risk λ_1 show that the market return commands a significant positive risk premium, while the Treasury return and nonlinear volatility factor each command a negative risk premium, which is consistent with a flight to safety interpretation. Taken together, the product $\beta_h^i \lambda_1$, representing loadings on the forecasting factor $\phi_h(v_t)$, has negative signs for all bond returns and positive signs for all equity returns.

Figure 9 shows the cross-sectional relationship between the forecasting intercept a_h^i and slope b_h^i and the risk factor exposures. The top left panel shows that the forecasting slope b_h^i is strongly related to the risk factor loadings β_h^i and prices of risk λ_1. Correspondingly, when deviations from arbitrage α_h^i are permitted, the top right panel supports the view that the dynamic asset pricing model that restricts slope coefficients b_h^i to be $\beta_h^i \lambda_1$ results in correct predictions about unconditional excess returns in the cross-section. These predictions are also captured in the unrestricted forecasting regressions (bottom right panel). However, under no-arbitrage in frictionless markets α_h^i is forced to zero, the pricing performance deteriorates in the bottom left panel. [9] Taken together, the results from Table 5 and Figure 9 strongly

[8] Because the right-hand side regressors in $Rx_{t+h}^i = (\alpha_h^i + \beta_h^i \lambda_0) + \beta_h^i \lambda_1 \phi_h(v_t) + \beta^i u_{t+h} + \varepsilon_{t+h}^i$ mix lagged and contemporaneous variables, overlapping regressions that result from using $h > 1$ imply certain parameter restrictions. To avoid the need to model such restrictions, we follow Adrian, Crump, and Moench (2014) and focus on the $h = 1$ case.

[9] The deterioration in pricing performance is attributable to our choice of industry sorted portfolios as test assets for the equities, as those are well known to generate pricing errors relative to risk based explanations. Size or book to market sorted portfolios would improve cross-sectional pricing performance.

suggest the interpretation of $\phi_h(v_t)$ as a price of risk variable in a dynamic asset pricing model with frictions.

Equation (4.1) also has the time series prediction that $\phi_h(v_t)$ might be related to alternative proxies for the time variation in pricing of risk. There is a large literature documenting the extent to which asset prices are predictable. An obvious question is to what extent our estimated return predictor $\phi_h(v_t)$ is related to other forecasting variables that have been shown to be significant. To do so, we plot six commonly used variables together with $\phi_h(v_t)$ in the time series for $h = 6$.[10] In particular, we show the slope of the Treasury yield curve as measured by TERM, the Cochrane and Piazzesi (2005) CP factor, the BAA-AAA credit spread (DEF), the dividend yield (DY) of the S&P500, the CAY factor by Lettau and Ludvigson (2001), and the variance risk premium (VRP) of Bollerslev, Tauchen, and Zhou (2009). It is apparent from Figure 10 that the relationships of $\phi_6(v_t)$ with all six variables is very weak. The only two variables that bear some resemblance are the credit spread DEF and the variance risk premium VRP, both of which increase during the financial crisis in tandem with $\phi_6(v_t)$. However, our earlier SRRR forecasting results showed that DEF and VRP do not impact the significance of $\phi_h(v_t)$ markedly.

4.2 Theories of Flight-to-Safety

The theoretical literature has proposed a number of distinct mechanisms generating flight-to-safety. Vayanos (2004) studies equilibrium asset pricing where volatility is stochastic, and assets are invested by fund managers. Assets' illiquidity arises as trading is subject to fixed transactions costs. Fund managers are subject to withdrawals when fund performance is poor, generating a preference for liquidity that is a time varying function of volatility.

[10] We provide a visual representation of the alternative forecasting variables, since four of the six plotted series were previously shown to be unrelated to the predictive content of $\phi_h(v_t)$ (Tables 1, 2, and 3), and since CAY is only available quarterly. In separate results (not included for brevity), we also included the CP factor as a predictive control and found that it did not affect our results.

Vayanos (2004) derives equilibrium expected returns of the following form[11]

$$E_t\left[Rx^i_{t+1}\right] = \alpha^i\left(v_t\right) + A\left(v_t\right)Cov_t\left(Rx^i_{t+1}, Rx^M_{t+1}\right) + Z\left(v_t\right)Cov_t\left(Rx^i_{t+1}, v_{t+1}\right). \tag{4.4}$$

Expected returns are thus functions of their covariance with the market return and with volatility, where the impact of these covariances on expected returns depends on the endogenously time varying effective risk aversion $A\left(v_t\right)$ and the endogenously time varying volatility risk premium $Z\left(v_t\right)$. Our estimated function $\phi\left(v_t\right)$ corresponds to $A\left(v_t\right)\cdot v_t$ and our estimated exposures on $\phi\left(v_t\right)$ are related to the covariance of asset returns with the risk factors, including market risk and interest rate risk. The solution for $A\left(v_t\right)$ is highly nonlinear, and must be computed numerically. Pricing errors α^i_t are related to trading costs across assets times the withdrawal likelihood, and are also a function of v_t, but they are unrelated to risk factor loadings. We note, however, that the downward sloping $\phi\left(v_t\right)$ for the VIX below its median and above its 99th percentile is not an immediate implication of the theory. We will discuss possible explanations below.

What is key to the framework of Vayanos (2004) is that $A\left(v_t\right)$ is convex. Furthermore, our estimated $\phi\left(v_t\right)$ has to be compared to $A\left(v_t\right)\cdot v_t$, making it more convex. Our estimate of $\phi\left(v_t\right)$ is convex when the VIX is above its median and below the 99th percentile. Within the theory, this is because risk premia are affected by fund managers' concern with withdrawals. Withdrawals are costly to the managers because the managers' fee is reduced, and holding a riskier portfolio makes withdrawals more likely by increasing the probability that performance falls below the threshold. When volatility is low, managers are not concerned with withdrawals and hence the component of the risk premium that corresponds to withdrawals is very small and almost insensitive to volatility. That component starts increasing rapidly, however, when volatility increases, leading the managers' effective risk aversion to increase with volatility.

[11] We are changing notation from continuous to discrete time to make it consistent with the rest of the notation in this paper.

The theory of Vayanos (2004) also features a volatility risk premium, $Z(v_t)$. The dynamic asset pricing setup allows us to separate out $Z(v_t)$ and $A(v_t)$, as we estimate prices of risk for the market, interest rate risk, and volatility risk to be affine functions of $\phi(v_t)$. $Z(v_t)$ captures managers' reduced willingness to hold illiquid assets during volatile times, leading liquidity premia to increase with volatility. We do note, however, that the volatility premium is also convex because when volatility is low, managers are not concerned with withdrawals because the event that performance falls below the threshold requires a movement of several standard deviations. When volatility increases the probability of withdrawals starts increasing rapidly, and so does the volatility premium. Importantly, the theory of Vayanos (2004) also gives rise to αs due to transaction costs. In equilibrium, those transaction costs are also related to the level of volatility, and to the withdrawal intensity of mutual fund investors.

We next investigate direct evidence of flight-to-safety by analyzing global mutual fund flows from the Investment Company Institute (ICI) (Table 6). The table reports a contemporaneous SRRR of the mutual fund flows on the VIX. The regression is contemporaneous, as flows drive expected returns contemporaneously with volatility. The table shows that U.S., world equity, and hybrid funds have positive loading on the fund flow function $\phi^{FF}(v_t)$ while government bond funds exhibit strongly negative loadings. This is direct evidence of flight to safety from stocks to bonds in terms of mutual fund flows. We can also see that the flight-to-safety pattern is present in the mutual fund flows whether the crisis is included or not. It is interesting to note that money market funds are not found to be a safe asset in these regressions, consistent with the fact that investors ran on money market funds. Figure 11 compares (an affine transformation of) the shape of the nonlinearity from the fund flows to the $\phi(v_t)$ function.

While the mutual fund flows line up with the $\phi(v_t)$ function in the intermediate range of the VIX between 18 and 50, the $\phi(v_t)$ function is downward sloping outside of that range. For the lower end of the range, the theory by Buffa, Vayanos, and Woolley (2014) might offer an explanation.

Buffa, Vayanos, and Woolley (2014) augment the Vayanos (2004) model with competition among fund managers. Because of agency frictions, investors make managers' fees more sensitive to performance and benchmark performance against a market index. This makes managers unwilling to deviate from the index and exacerbates price distortions. Because trading against overvaluation exposes managers to greater risk of deviating from the index than trading against undervaluation, agency frictions bias the aggregate market upwards. They can also generate a negative relationship between risk and return because they raise the volatility of overvalued assets.[12] [13] [14]

Caballero and Krishnamurthy (2008) study a very different but complementary mechanism for flight to safety, based on Knightian uncertainty. In their model, agents faced with Knightian uncertainty consider the worst case among the scenarios over which they are uncertain. When the aggregate quantity of liquidity is limited, Knightian agents grow concerned with liquidity shortages and they therefore sell risky financial claims in favor of safe and uncontingent claims, i.e. there is flight to safety. Even though the flight to safety seems prudent from individuals' point of views, it is collectively costly for the macroeconomy

[12]An earlier paper by Vayanos and Woolley (2013) features exogenous fund flows that enter into equilibrium pricing relationships. As a result, expected returns depend on covariance with fund flows as well as covariance with the market return. The price of risk of fund flow risk and market risk is in turn a function of fund flows, which are also correlated with equilibrium volatility. Hence Vayanos and Woolley (2013) also imply a nonlinear relationship between expected returns and market volatility, but only indirectly. However, there is no negative risk-return tradeoff.

[13]A related literature studies the role of market makers for flight-to-safety. Building on work by Gromb and Vayanos (2002), Weill (2007) models financial crises in a setting where market makers provide liquidity by absorbing external selling pressure. Market makers buy when selling pressure is large, accumulate inventories, and sell when the pressure alleviates. Weill (2007) points out that market makers can only provide liquidity when they have sufficient capital, justifying central bank lending in times of financial market disruptions. Equilibrium pricing is highly nonlinear, and related to volatility, but does not lend itself to a reduced form, empirically implementable representation as is the case with Vayanos (2004) model. Hence we cannot map our estimated $\phi(v_t)$ directly into the model of Weill (2007).

[14]Another extension of Gromb and Vayanos (2002) is provided by Brunnermeier and Pedersen (2009), who introduce an explicit role for stochastic volatility within a pricing setting that feature arbitrage capital. Brunnermeier and Pedersen (2009) present a model that links an asset's market liquidity and traders' funding liquidity. Traders' ability to provide funding is in turn linked to asset return volatility via margin setting. As a result, expected returns are linked to market volatility via traders' ability to provide liquidity. Under certain conditions, margins are destabilizing and market liquidity and funding liquidity are mutually reinforcing, leading to liquidity spirals that are associated with flight to safety, and widening expected returns to risky assets. The model of Brunnermeier and Pedersen (2009) gives rise to a strongly nonlinear and potentially discontinuous relationship between future returns and volatility, which is broadly consistent with our estimated $\phi(v_t)$ function.

because scarce liquidity goes wasted. To the extent that a high level of the VIX might trigger Knightian agents to flight to safety, or a high level of the VIX is correlated with an increase in uncertainty, the predictions of Caballero and Krishnamurthy (2008) broadly support our empirical finding that riskier securites tend to load positively on $\phi(v_t)$, while safe securities load negatively on $\phi(v_t)$. More importantly, Caballero and Krishnamurthy (2008) implies that flight to safety might be associated with adverse macroeconomic consequences in extreme tail events. The latter can explain why the $\phi(v_t)$ curve starts to slope down when the VIX is above the VIX's 99th percentile of 50. In those extreme states, flight to safety might have severe consequences on expectations of future cash flows, so that our estimated relationship between expected returns and the VIX no longer reflects the pricing of risk, but rather the expectations about future cash flows in extreme tail events. Our observations for the VIX at these extreme levels are concentrated in the months following the Lehman failure of 2008, which was followed by abysmal macroeconomic performance and cash flows, leading to ever declining stock prices and yields. The decline in yields, in turn, is at least partially attributable to central bank liquidity injections. The reversal of the relationship between expected returns and volatility for very high levels of volatility is fully consistent with the framework of Caballero and Krishnamurthy (2008), as extreme uncertainty leads to extreme flight to safety, causing potentially large macroeconomic costs.

4.3 Intermediary Asset Pricing Theories

Intermediary asset pricing theories model the impact of intermediary balance sheet frictions on the pricing of risk and real activity within dynamic general equilibrium models of the macroeconomy. The strand of literature was pioneered by He and Krishnamurthy (2013), who model an intermediary sector whose ability to raise external equity capital depends on past performance, similar to Vayanos (2004). However, in He and Krishnamurthy (2013), volatility is an endogenous variable, while it is exogenous in Vayanos (2004). He and Krishnamurthy (2013) feature a single state variable, which is the share of intermediary wealth relative to total wealth in the economy. Aggregate volatility is a nonlinear function of the share of that state variable. The approach of Brunnermeier and Sannikov (2014) has many commonalities with He and Krishnamurthy (2013). Even though constraints on the intermediary sector differ, the model of Brunnermeier and Sannikov (2014) also features the wealth share of intermediaries as single state variable that determines the pricing of risk, the level of endogenous volatility, and real activity.

Adrian and Boyarchenko (2012) expand on those intermediary asset pricing theories by introducing intermediary leverage as an additional state variable. The intermediary leverage state variable arises endogenously, as intermediaries are subject to value at risk (VaR) constraints. Furthermore, intermediaries face liquidity shocks in addition to productivity shocks, while the models of He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014) only feature a single shock. The pricing kernel of Adrian and Boyarchenko (2012) is expressed in terms of shocks to output and shocks to leverage, with prices of risk depending on relative intermediary wealth (as in He-Krishnamurthy and Brunnermeier-Sannikov) as well as the level of leverage. This type of pricing kernel is well supported by empirical asset pricing evidence. Adrian, Moench, and Shin (2010) show that intermediary leverage is a strong forecasting factor for asset returns, while Adrian, Etula, and Muir (2014) demonstrate that shocks to intermediary leverage is a priced risk factor. Adrian, Moench, and Shin (2014) combine both of these results and present a dynamic asset pricing model that features leverage as a pricing factor, and a price of risk variable.

The important connection between the present paper, and Adrian and Boyarchenko (2012), is the fact that the VaR constraint directly links aggregate volatility to the leverage of intermediaries. Hence the pricing kernel of Adrian-Boyarchenko could be expressed as a function of volatility—instead of a function of leverage—both for the pricing factor, and the pricing of risk. Intuitively, increases in volatility, which arise endogenously, tighten the leverage constraints on intermediaries, increasing their effective risk aversion, as well as equilibrium pricing of risk. Booms correspond to periods when volatility is endogenously low, pricing of risk is compressed, and intermediary leverage is elevated. When adverse shocks hit, either to productivity or to liquidity, volatility rises endogenously, tightening balance sheet constraints and leading to a widening in the pricing of risk.

In order to gauge the relationship of our pricing of risk function $\phi(v_t)$ to intermediary asset pricing theories, we estimate the relationship between the cross-section of VaRs over time and the VIX, using SRRRs. We obtain the VaRs from Bloomberg for Bank of America, Citigroup, Goldman Sachs, JP Morgan, and Morgan Stanley. VaRs are expressed in dollar terms. We use the aforementioned five banks as those institutions are the main US banking organizations with trading operations that have reported data continuously since 2004. The VaR data is, unfortunately, only available at a quarterly frequency, so that return forecasting regressions have very few observations.[15] Instead, we present in Figure 12 the results of regressing the panel of VaRs contemporaneously on the VIX, using SRRR regressions. The result of this is shown in the lower panel of Figure 12, while the upper panel presents the sum of VaRs together with the VIX. We can see that there is a tight association between the VIX and the VaRs, and that the SRRR is suggestive of a slightly concave relationship.

The empirical results of Figure 12 support the assumption of Adrian and Boyarchenko (2012) that intermediary balance sheet constraints are related to market volatility due to risk management constraints. Furthermore, our earlier finding that expected returns are systematically related to the VIX are compatible with the notion that constraints on intermediary balance sheets matter for the pricing of risk. Of course, this evidence is only suggestive, and more rigorous analysis would require the calibration of intermediary asset pricing models, or an identification strategy for exogenous variation in dealer balance sheet capacity. We leave such research for future work.

[15] We did perform SRRRs of the stock and bond portfolios on the six month lag of the summed VaRs and found significant forecasting ability using the SRRRs. We also uncovered the mirror image property.

4.4 Time Varying Pricing of Risk due to Habit Formation

Besides theories of flight-to-safety and intermediary asset pricing theories, habit formation can give rise to variation in the pricing of risk. In habit formation theories, utility depends not just on the level of current consumption, but rather on current consumption in excess of previously experienced consumption. Asset pricing implications of habit formation were pioneered by Abel (1990), Constantinides (1990), and Sundaresan (1989). Our discussion will focus on the theory by Campbell and Cochrane (1999, 2000), which has synthesized earlier work, and proven successful in explaining asset pricing puzzles.

The key state variable in the Campbell and Cochrane (1999, 2000) setup is the surplus consumption ratio s_t, which is a slow moving, mean reverting function of past shocks to aggregate consumption. The pricing of risk is a function of the surplus consumption ratio. Campbell and Cochrane (1999, 2000) present parameter calibrations that have been shown to be able to explain asset pricing puzzles, both in the cross-section and in the time series.

In order to gauge the plausibility of our estimated $\phi(v_t)$ function within the Campbell and Cochrane (1999, 2000) asset pricing context, we first compute their pricing kernel from the surplus consumption ratio and shocks to consumption:

$$\begin{aligned}
\ln M_{t+1} &= \ln(\delta) - \gamma g - \gamma(\phi - 1)(s_t - s) - \gamma(1 + \lambda(s_t))(c_{t+1} - c_t - g) \\
s_{t+1} &= (1 - \phi)\bar{s} + \phi s_t + \lambda(s_t)(c_{t+1} - c_t - g) \\
\lambda(s_t) &= \begin{cases} \frac{1}{\bar{S}}\sqrt{1 - 2(s_t - \bar{s})} - 1 & s_t \leq s_{\max} \\ 0 & s_t > s_{\max} \end{cases}
\end{aligned}$$

where we use the parameters where $g = 1.89$, $\sigma = 1.5$, $\phi = 0.87$, $\gamma = 2$, $\delta = 0.89$, $\bar{S} = 0.057$, and $S_{\max} = 0.094$ from Campbell and Cochrane (2000). We then consider the conditional Sharpe ratio generated by that pricing kernel,

$$\frac{E_t\left[R^*_{t+1}\right]}{\sqrt{Var_t\left[R^*_{t+1}\right]}} = \sqrt{Var_t(\ln M_{t+1})} = \gamma(1 + \lambda(s_t))\sigma,$$

where R^*_{t+1} corresponds to the returns of the maximum Sharpe ratio portfolio. It is thus a property of the maximum Sharpe ratio portfolio that its Sharpe ratio relates to the volatility of the pricing kernel in a linear fashion. This raises the question how the habit induced maximum Sharpe ratio links to market volatility as measured by the VIX. There can be a difference between the conditional volatility of the maximum Sharpe ratio portfolio and market implied volatility as the S&P500, whose conditional volatility the VIX measures, might not be the maximum Sharpe ratio portfolio.

Figure 13 shows the Campbell-Cochrane implied maximum conditional Sharpe ratio together with the VIX over time. There is some positive correlation among the two series. For example, during the financial crisis, when the VIX increases sharply, the Sharpe ratio from the habit model also increases unusually strongly. However, while the VIX reverts back to lower levels after 2009, the habit Sharpe ratio remains high through 2015. Another striking divergence also occurs during the late 1990s and early 2000s, when the tech boom in the stock market was associated with high volatility, but a low level of the habit Sharpe ratio. Hence overall, the correlation between the VIX and the habit Sharpe ratio is fairly low.

Our takeaway from this finding is that time variation of expected returns derived from the SRRRs on the VIX captures distinct economic mechanism when compared to the time variation of expected returns induced by habit formation. While habit formation is very tightly linked to the growth of aggregate consumption, expected returns derived from the VIX are likely to capture funding constraints on financial institutions such as major banking organizations and fund managers, as we argued above. Hence the two sources of time variation in expected returns and Sharpe ratios capture complementary economic forces.

4.5 Macroeconomic Consequences

Our final investigation concerns macroeconomic aggregates. In general equilibrium, variations in the pricing of risk associated with fund flows or bank balance sheet constraints will translate into distortions of consumption and savings decisions. The work of He and Krishnamurthy (2013), Brunnermeier and Sannikov (2014), and Adrian and Boyarchenko (2012) formalize this link in dynamic macro-finance theories. One would therefore expect the VIX to forecast macroeconomic aggregates. Table 7 reports the output of SRRRs on the five business cycle indicators that receive the largest weight in the Chicago Fed National Activity Index (CFNAI): industrial production (IP), IP manufacturing (IPMFG), manufacturing capacity utilization (CUMFG), change in goods-producing employment (LAGOODA), and total private nonfarm payroll series (LAPRIVA), with the market return included as the reference series.[16] We can see that the VIX strongly forecasts macro activity, with all of the business cycle indicators receiving the same sign. Furthermore, allowing for nonlinearity helps in terms of explanatory power.

Figure 14 shows the relationship between macroeconomic activity, as measured by the CFNAI's 3-month moving average, and $\phi_h(vix_t)$, estimated solely from asset return data as in Section 3.6 using our benchmark 6-month forecast horizon. The figure illustrates that the reversal of $\phi_h(vix_t)$ as VIX rises above its 99th percentile in the fall of 2008 (vertical line) presages the subsequent collapse in real activity. We therefore conjecture that very extreme observations of the VIX lead to a drastic update of expected cash flows, leading to the downward sloping $\phi_h\left(v_t\right)$ seen in prior figures. For Treasury returns, expectations about accommodative monetary policy might lead to the sharp rise in the relation between returns and the VIX, preserving the mirror image property even at extreme events. Because the entire cross-section of asset returns considered is pricing this reversal in expected returns, we emphasize that the shape of $\phi_h\left(v_t\right)$ is still precisely estimated for these extreme values of the VIX.

[16]We use the same stationarity-inducing transformations employed in the actual CFNAI calculation. For more details, see https://www.chicagofed.org/~/media/publications/cfnai/background/cfnai-technical-report-pdf.pdf.

5 Conclusion

We propose sieve reduced rank regressions (SRRR) to extract a common nonlinear function of the VIX that jointly forecasts stock and bond returns at horizons up to two years. We show that the forecasting function $\phi_h(v)$ is the same across diverse sets of stock, bond, and credit returns, up to affine transformations. Intriguingly, the loadings of stock and bond returns on the common forecasting variable switches signs when comparing stocks and bonds. This is evidence of investor flight-to-safety: when the VIX rises above its median value, investors tend to reallocate from stocks to bonds, leading to an increase in expected returns for stocks and a compression of expected returns for bonds. We show that the shape of the functional form is robust across asset classes and across time. We can extract virtually indistinguishable shapes from only stocks or only bonds, or use subsamples of the data that include the 1987 crash but exclude the 2008 crisis, or vice versa.

When we relate $\phi_h(v)$ to common return predictors, we find only very weak relationships, suggesting that the nonlinear function of equity market volatility captures economic forces that are complementary to previously documented forecasting variables. Cross-sectional regressions show that the loadings of future returns on $\phi_h(v)$ are cross-sectionally related risk factor loadings, suggesting that $\phi_h(v)$ is a price of risk variable in a dynamic asset pricing model. Our findings support the nonlinear pricing predictions of the asset management theory by Vayanos (2004) where flight-to-safety is associated with increases in equity market volatility, and the intermediary asset pricing theory of Adrian and Boyarchenko (2012) where value at risk constraints on banks give rise to a tight relationship between the pricing of risk and the level of aggregate volatility. We provide supportive evidence for both theories by analyzing global mutual fund flows and the value at risk constraints of major banking organizations.

Appendix

A.1 Details of Empirical Implementation

Robust Multi-Horizon Inference As discussed in the main text we utilize the "reverse regression" approach of Hodrick (1992) to our multivariate setting to ameliorate concerns about spurious inference induced by the overlapping returns. For a similar derivation for the univariate case see Wei and Wright (2013). Consider the multivariate h-period long-horizon return regression where $R_{t+h}^{(h)} = R_{t+1} + R_{t+2} + \cdots + R_{t+h}$,

$$R_{t+h}^{(h)} = \mathbf{a}_h + \mathbf{A}_h X_t + \varepsilon_{t+h}, \qquad t = 1, \ldots, T,$$

where $\mathbb{E}\left[\varepsilon_{t+h} \otimes X_t\right] = 0$ and $\mathbb{E}\left[\varepsilon_{t+h}\right] = 0$. Then, $\mathbf{A}_h = \mathbb{C}\left(R_{t+h}, X_t\right) \mathbb{V}\left(X_t\right)^{-1}$ so long as $\mathbb{V}\left(X_t\right)$ is of full rank. Similarily, for the multivariate h-period reverse regression,

$$R_{t+1} = \mathbf{a} + \mathbf{A} X_t^{(h)} + \varepsilon_{t+1}, \qquad t = 1, \ldots, T,$$

then $\mathbf{A} = \mathbb{C}\left(R_{t+1}, X_t^{(h)}\right) \mathbb{V}\left(X_t^{(h)}\right)^{-1}$ where $X_t^{(h)} = X_t + X_{t-1} + \cdots + X_{t-h}$. Then, under the assumption of covariance stationarity of the joint process (R_t, X_t) we have that $\mathbb{C}\left(R_{t+h}^{(h)}, X_t\right) = \mathbb{C}\left(R_{t+1}, X_t^{(h)}\right)$ and so

$$\mathbf{A}_h = \mathbb{C}\left(R_{t+h}^{(h)}, X_t\right) \mathbb{V}\left(X_t\right)^{-1} = \mathbb{C}\left(R_{t+1}, X_t^{(h)}\right) \mathbb{V}\left(X_t^{(h)}\right)^{-1} \mathbb{V}\left(X_t^{(h)}\right) \mathbb{V}\left(X_t\right)^{-1} = \mathbf{A} \mathbb{V}\left(X_t^{(h)}\right) \mathbb{V}\left(X_t\right)^{-1}.$$

Thus, the ith row of $\mathbf{A}_h$ has all elements equal to zero if and only if the ith row of $\mathbf{A}$ has all elements equal to zero. Thus, tests of the null hypothesis of no predictability of asset i in the forward regression are equivalent to the same test in the reverse regression which formally justifies our empirical approach.

Proof of Proposition 1: The first and second results follow by the same steps as in Crump, Hotz, Imbens, and Mitnik (2008), except the Berry-Esseen bound used in their Lemma A.3 is replaced by its counterpart for weakly dependent processes, as in Bulinskii and Shashkin (2004). Then for sufficiently smooth $\phi(\cdot)$, the convergence holds. For the second result, to establish that $m+n-1$ is the correct degrees of freedom, note that the rank of the variance matrix given in equation (3.20) of Anderson (1999) can be shown to be $m+n-1$ in our notation, using the rank result of Cline and Funderlic (1979). Finally, the third result follows from Chen, Liao, and Sun (2014).

References

Abel, A. B. (1990): "Asset prices under habit formation and catching up with the Joneses," *American Economic Review*, pp. 38–42.

Adrian, T., and N. Boyarchenko (2012): "Intermediary Leverage Cycles and Financial Stability," *Federal Reserve Bank of New York Staff Reports*, 567.

Adrian, T., R. K. Crump, and E. Moench (2013): "Pricing the term structure with linear regressions," *Journal of Financial Economics*, 110(1), 110–138.

——— (2014): "Regression-based estimation of dynamic asset pricing models," *Journal of Financial Economics*, forthcoming.

Adrian, T., E. Etula, and T. Muir (2014): "Financial Intermediaries and the Cross-Section of Asset Returns," *Journal of Finance*, 69, 2557–2596.

Adrian, T., E. Moench, and H. S. Shin (2010): "Financial Intermediation, Asset Prices, and Macroeconomic Dynamics," Federal Reserve Bank of New York Staff Reports No. 442.

——— (2014): "Dynamic Leverage Asset Pricing," Federal Reserve Bank of New York Staff Reports No. 625.

Anderson, T. (1999): "Asymptotic distribution of the reduced rank regression estimator under general conditions," *The Annals of Statistics*, 27(4), 1141–1154.

Ang, A., and G. Bekaert (2007): "Stock return predictability: Is it there?," *Review of Financial Studies*, 20(3), 651–707.

Ang, A., and M. Ulrich (2012): "Nominal Bonds, Real Bonds, and Equity," Working paper.

Baele, L., G. Bekaert, K. Inghelbrecht, and M. Wei (2013): "Flights to safety," Discussion paper, National Bureau of Economic Research.

Bali, T. G., and R. F. Engle (2010): "The intertemporal capital asset pricing model with dynamic conditional correlations," *Journal of Monetary Economics*, 57(4), 377–390.

Beber, A., M. W. Brandt, and K. A. Kavajecz (2009): "Flight-to-quality or flight-to-liquidity? Evidence from the euro-area bond market," *Review of Financial Studies*, 22(3), 925–957.

Bekaert, G., E. Engstrom, and S. Grenadier (2010): "Stock and Bond Returns with Moody Investors," *Journal of Empirical Finance*, 17(5), 867–894.

Bekaert, G., and M. Hoerova (2014): "The VIX, the variance premium and stock market volatility," *Journal of Econometrics*, forthcoming.

Bollerslev, T., D. Osterrieder, N. Sizova, and G. Tauchen (2013): "Risk and return: Long-run relations, fractional cointegration, and return predictability," *Journal of Financial Economics*, 108(2), 409–424.

Bollerslev, T., G. Tauchen, and H. Zhou (2009): "Expected stock returns and variance risk premia," *Review of Financial Studies*, 22(11), 4463–4492.

Brunnermeier, M. K., and L. H. Pedersen (2009): "Market liquidity and funding liquidity," *Review of Financial studies*, 22(6), 2201–2238.

Brunnermeier, M. K., and Y. Sannikov (2014): "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 104(2), 379–421.

Buffa, A. M., D. Vayanos, and P. Woolley (2014): "Asset management contracts and equilibrium prices," Discussion paper, Working Paper, London School of Economics.

Bulinskii, A., and A. Shashkin (2004): "Rates in the CLT for sums of dependent multi-indexed random vectors," *Journal of Mathematical Sciences*, 122(4), 3343–3358.

Caballero, R. J., and A. Krishnamurthy (2008): "Collective risk management in a flight to quality episode," *Journal of Finance*, 63(5), 2195–2230.

Campbell, J. Y., and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), 205–251.

——— (2000): "Explaining the poor performance of consumption-based asset pricing models," *Journal of Finance*, 55(6), 2863–2878.

Campbell, J. Y., and R. J. Shiller (1988a): "The dividend-price ratio and expectations of future dividends and discount factors," *Review of Financial Studies*, 1(3), 195–228.

——— (1988b): "Stock prices, earnings, and expected dividends," *Journal of Finance*, 43(3), 661–676.

Chen, X. (2007): "Large sample sieve estimation of semi-nonparametric models," *Handbook of Econometrics*, 6, 5549–5632.

Chen, X., Z. Liao, and Y. Sun (2014): "Sieve inference on possibly misspecified semi-nonparametric time series models," *Journal of Econometrics*, 178, 639–658.

Cline, R. E., and R. Funderlic (1979): "The rank of a difference of matrices and associated generalized inverses," *Linear Algebra and its Applications*, 24, 185–215.

Cochrane, J. H. (2011): "Presidential address: Discount rates," *The Journal of Finance*, 66(4), 1047–1108.

Cochrane, J. H., and M. Piazzesi (2005): "Bond Risk Premia," *American Economic Review*, 95(1), 138–160.

——— (2008): "Decomposing the yield curve," Working Paper.

Constantinides, G. M. (1990): "Habit formation: A resolution of the equity premium puzzle," *Journal of Political Economy*, pp. 519–543.

Crump, R. K., V. J. Hotz, G. W. Imbens, and O. A. Mitnik (2008): "Nonparametric tests for treatment effect heterogeneity," *The Review of Economics and Statistics*, 90(3), 389–405.

Duffee, G. (2002): "Term Premia and Interest Rate Forecasts in Affine Models," *Journal of Finance*, 57(1), 405–443.

Ghysels, E., P. Santa-Clara, and R. Valkanov (2005): "There is a risk-return tradeoff after all," *Journal of Financial Economics*, 76(3), 509–548.

Gromb, D., and D. Vayanos (2002): "Equilibrium and welfare in markets with financially constrained arbitrageurs," *Journal of financial Economics*, 66(2), 361–407.

Guo, H., and R. F. Whitelaw (2006): "Uncovering the risk–return relation in the stock market," *Journal of Finance*, 61(3), 1433–1463.

He, Z., and A. Krishnamurthy (2013): "Intermediary Asset Pricing," *American Economic Review*, 103(2), 732–770.

Hodrick, R. J. (1992): "Dividend yields and expected stock returns: Alternative procedures for inference and measurement," *Review of Financial Studies*, 5(3), 357–386.

Johansen, S. (1995): "Likelihood-based inference in cointegrated vector autoregressive models," .

Koijen, R. S. J., H. N. Lustig, and S. van Nieuwerburgh (2013): "The Cross-Section and Time-Series of Stock and Bond Returns," Working Paper.

Lettau, M., and S. Ludvigson (2001): "Consumption, aggregate wealth, and expected stock returns," *Journal of Finance*, 56(3), 815–849.

Lettau, M., and S. Van Nieuwerburgh (2008): "Reconciling the return predictability evidence," *Review of Financial Studies*, 21(4), 1607–1652.

Lettau, M., and J. Wachter (2010): "The Term Structures of Equity and Interest Rates," *Journal of Financial Economics*, forthcoming.

Li, Q., and J. S. Racine (2007): *Nonparametric econometrics: Theory and practice*. Princeton University Press.

Longstaff, F. A. (2004): "The Flight-to-Liquidity Premium in US Treasury Bond Prices," *Journal of Business*, 77(3), 511–526.

Lundblad, C. (2007): "The risk return tradeoff in the long run: 1836–2003," *Journal of Financial Economics*, 85(1), 123–150.

Mamaysky, H. (2002): "Market Prices of Risk and Return Predictability in a Joint Stock-Bond Pricing Model," Working Paper.

Merton, R. C. (1973): "An intertemporal capital asset pricing model," *Econometrica*, pp. 867–887.

Pesaran, M. H., D. Pettenuzzo, and A. Timmermann (2006): "Forecasting time series subject to multiple structural breaks," *The Review of Economic Studies*, 73(4), 1057–1084.

Reinsel, G. C., and R. P. Velu (1998): *Multivariate Reduced-Rank Regression*. Springer, New York.

Ross, S. A. (1976): "The arbitrage theory of capital asset pricing," *Journal of Economic Theory*, 13(3), 341–360.

Rossi, A., and A. Timmermann (2010): "What is the Shape of the Risk-Return Relation?," *Unpublished working paper. University of California–San Diego*.

Sharpe, W. F. (1964): "Capital asset prices: A theory of market equilibrium under conditions of risk," *Journal of Finance*, 19(3), 425–442.

Sundaresan, S. M. (1989): "Intertemporally dependent preferences and the volatility of consumption and wealth," *Review of Financial Studies*, 2(1), 73–89.

Vayanos, D. (2004): "Flight to quality, flight to liquidity, and the pricing of risk," Discussion paper, National Bureau of Economic Research.

Vayanos, D., and P. Woolley (2013): "An institutional theory of momentum and reversal," *Review of Financial Studies*, pp. 1087–1145.

Wei, M., and J. H. Wright (2013): "Reverse Regressions And Long-Horizon Forecasting," *Journal of Applied Econometrics*, 28(3), 353–371.

Weill, P.-O. (2007): "Leaning against the wind," *Review of Economic Studies*, 74(4), 1329–1354.

Figure 2: **P-values by Forecast Horizon: 1990 to 2014**

This figure plots p-values by forecast horizon for linear and polynomial VIX predictive regressions. The regressions $Re^i_{t+h} = a^i_0 + a^i_1 VIX_t + \varepsilon^i_{t+h}$ and $Re^i_{t+h} = b^i_0 + b^i_1 VIX_t + b^i_2 VIX^2_t + b^i_3 VIX^3_t + \varepsilon^i_{t+h}$ are each estimated by OLS for $h = 1, \ldots 24$, where i ranges over 1-year, 2-year, 3-year, 5-year, and 10-year Treasury excess returns and stock market excess returns. p-values for Wald tests of joint significance of slope coefficients $H_0 : b^i_1 = b^i_2 = b^i_3 = 0$ using Hodrick (1992) standard errors are reported. The sample period is 1990:1-2014:9.

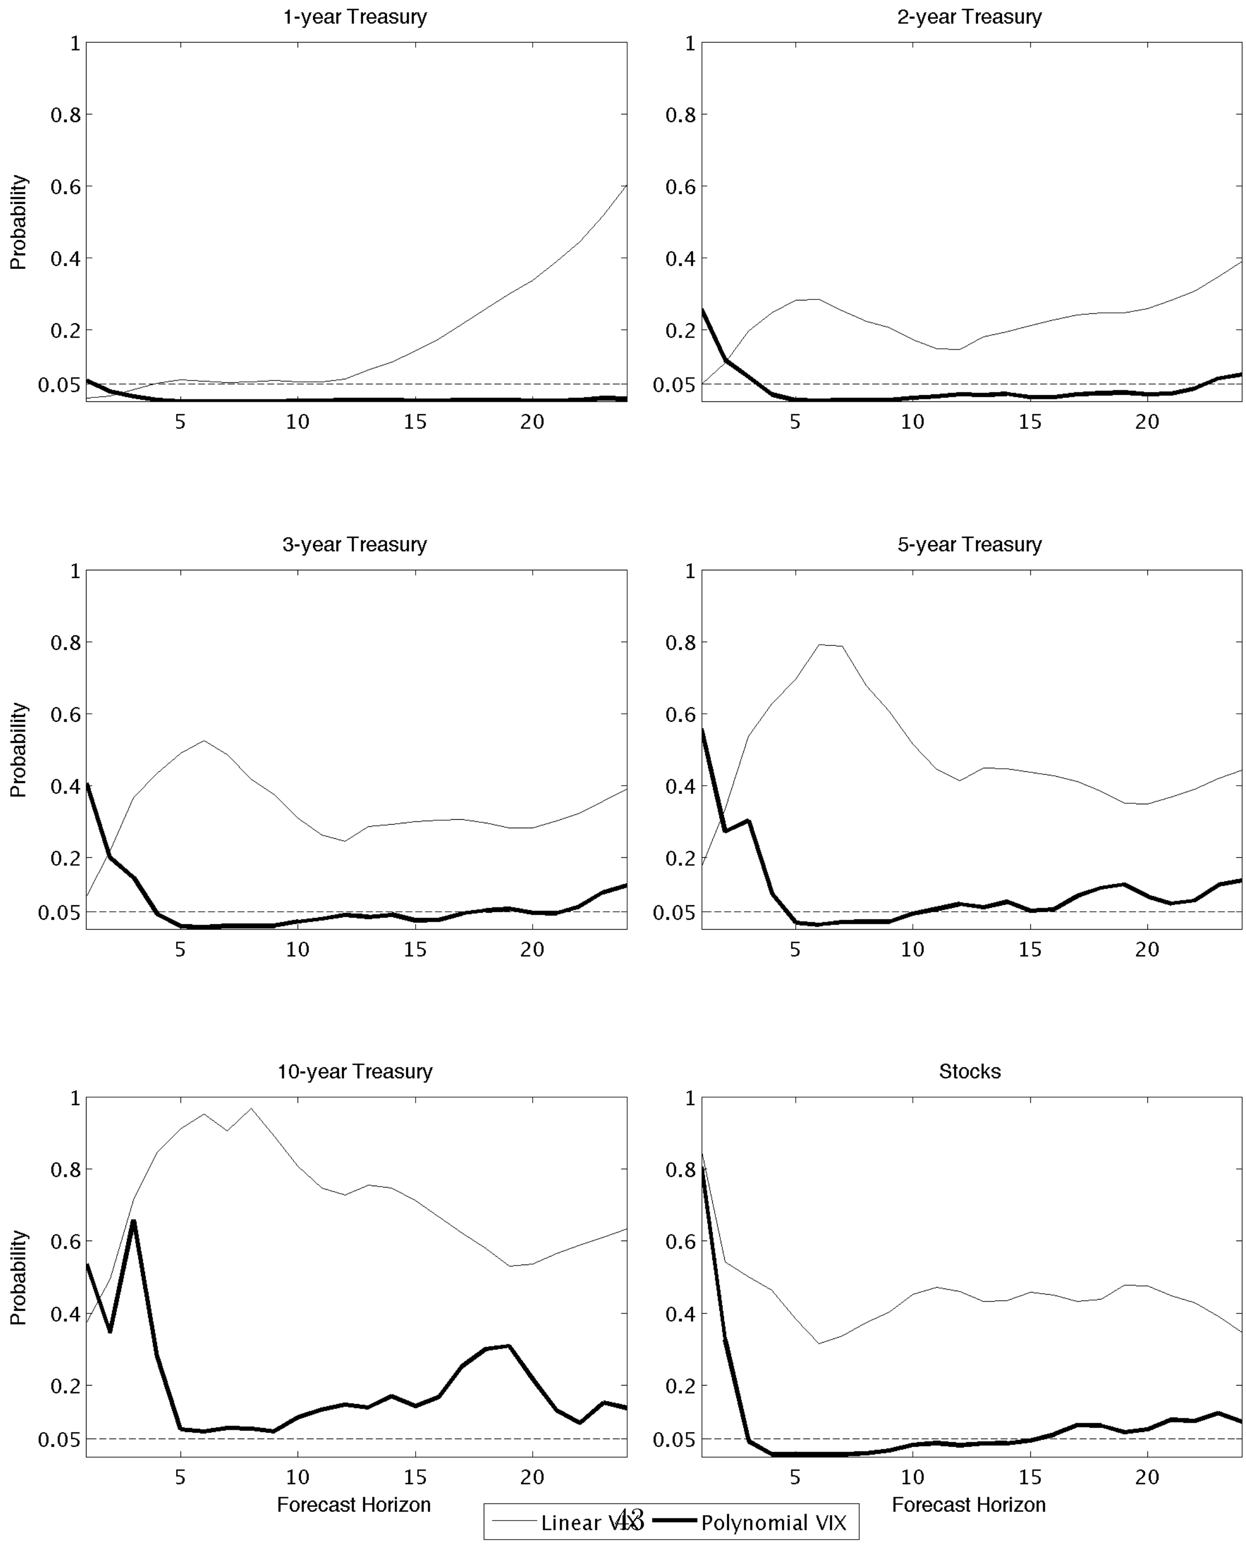

Figure 3: **P-values by Forecast Horizon: 1990 to 2007**

This figure plots p-values by forecast horizon for linear and polynomial VIX predictive regressions. The regressions $Re^i_{t+h} = a^i_0 + a^i_1 VIX_t + \varepsilon^i_{t+h}$ and $Re^i_{t+h} = b^i_0 + b^i_1 VIX_t + b^i_2 VIX^2_t + b^i_3 VIX^3_t + \varepsilon^i_{t+h}$ are each estimated by OLS for $h = 1, \ldots 24$, where i ranges over 1-year, 2-year, 3-year, 5-year, and 10-year Treasury excess returns and stock market excess returns. p-values for Wald tests of joint significance of slope coefficients $H_0 : b^i_1 = b^i_2 = b^i_3 = 0$ using Hodrick (1992) standard errors are reported. The sample period is 1990:1-2007:7.

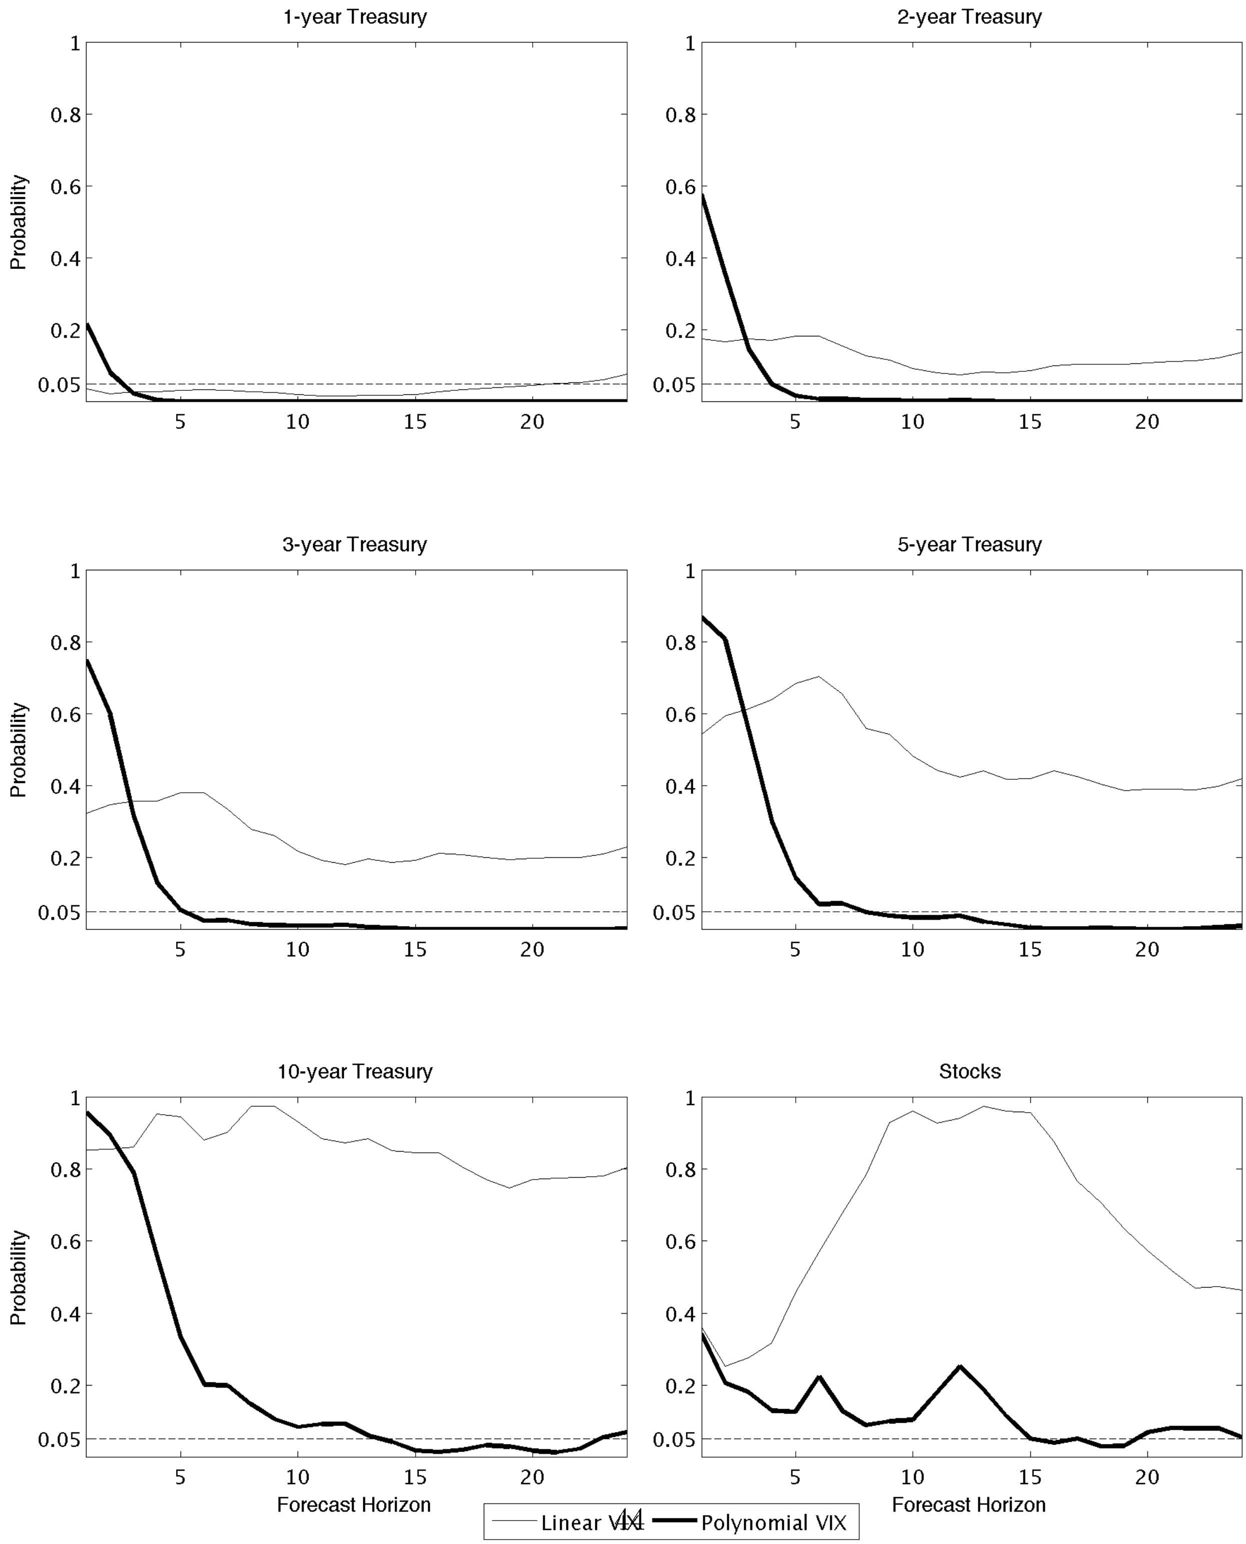

Figure 4: **Univariate Nonparametric and Polynomial Estimates of $\phi_h^i(v)$**

This figure shows sieve, polynomial, and kernel estimates of the nonlinear volatility function $\phi_h^i(v)$ from univariate predictive regressions $Rx_{t+h}^i = \phi_h^i(v_t) + \varepsilon_{t+h}$, where $E_t[Rx_{t+h}^i] = \phi_h^i(v_t)$. The superscript i indexes separate regressions in which the left hand side variable is either equity market excess returns (solid line) or 1-year Treasury excess returns (dashed line). In the top panel, the sample consists of monthly observations on $v_t = VIX_t$ from 1990:1 to 2014:9, whereas in the bottom panel, the sample consists of monthly observations from 1990:1 to 2007:7. In both panels, the forecast horizon plotted is $h = 6$ months. Within each panel, the left plot shows the nonparametric sieve estimate of $\phi_h^i(v_t)$, where the number of B-spline basis functions used in the estimation is chosen by out-of-sample cross validation. The middle plot shows a parametric cubic polynomial regression where $\phi_h^i(v_t) = a_0^i + a_1^i v_t + a_2^i v_t^2 + a_3^i v_t^3$, and the right plot shows $\phi_h^i(v_t)$ estimated by a Nadaraya-Watson kernel regression, where the bandwidth was chosen by Silverman's rule of thumb. The y-axis was rescaled by the unconditional standard deviation of Rx_t^i to display risk-adjusted returns.

$\phi_h^i(VIX)$ 1990:1 to 2014:9

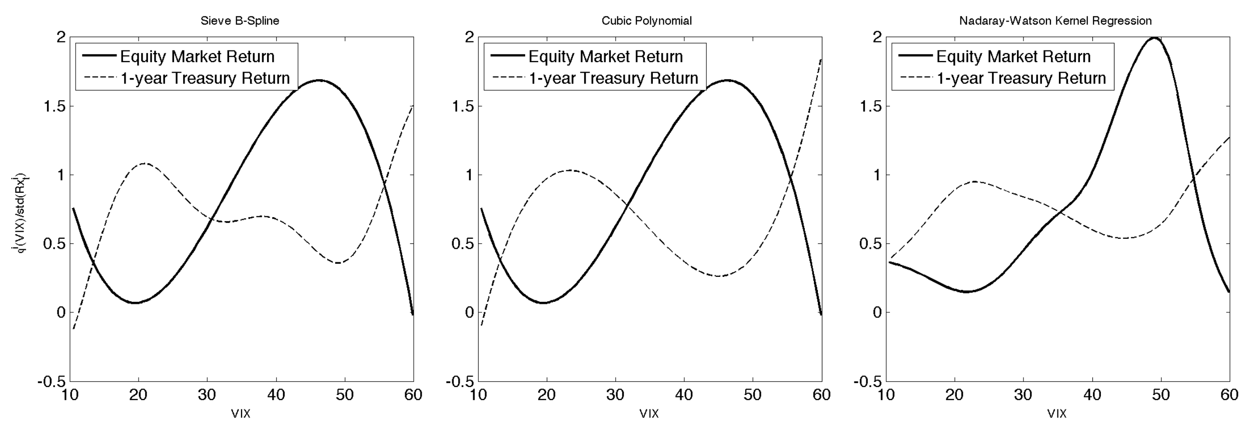

$\phi_h^i(VIX)$ 1990:1 to 2007:7

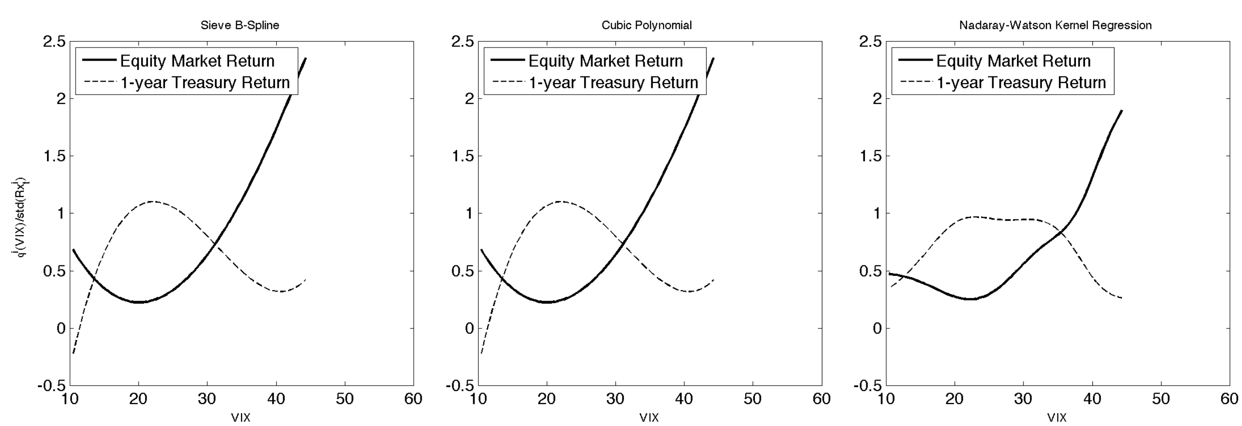

Figure 5: **SRRR Loadings:** $\hat{b}_h^i$

This figure plots SRRR estimated portfolio loadings $\hat{b}_h^i$, where i ranges over the market return (MKT), 11 industry portfolio returns, constant-maturity Treasury returns with 1, 2, 5, 7, 10, 20, and 30 year maturities, and Barclays corporate bond portfolios. The shades of bars denote outcomes of the hypothesis test $\mathbb{H}_{1,0}: b_h^i \phi_h = 0$ of whether asset i loads signficantly on $\phi_h(v)$ (see Proposition 1). The samples consist of monthly observations from 1990:1 to 2014:10.

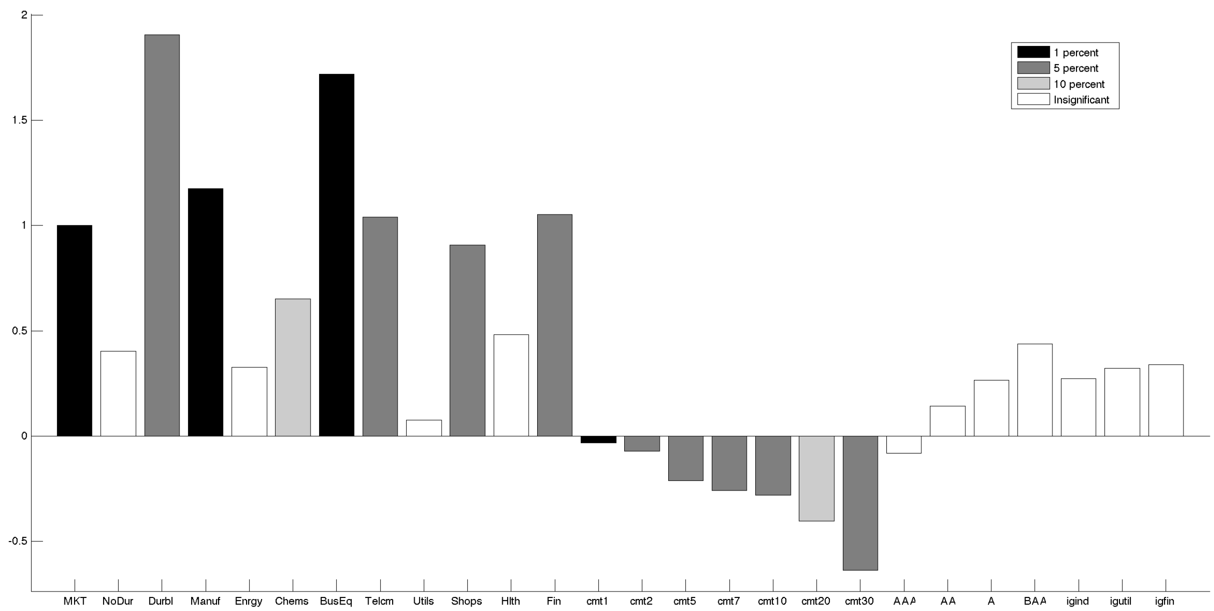

Figure 6: **SRRR: Expected Excess Returns for i^{th} Portfolio**

This figure plots normalized SRRR estimated excess returns on portfolio i, $\hat{E}_t[Rx^i_{t+h}]/\hat{\sigma}(Rx^i_{t+h})$, where $\hat{E}_t[Rx^i_{t+h}] = \hat{\alpha}^i_h + \hat{b}^i_h \hat{\phi}_h(v_t)$, $\hat{\sigma}(Rx^i_{t+h})$ scales by unconditional excess return standard deviation, and where i ranges over the market return (MKT), 11 industry portfolio returns, constant-maturity Treasury returns with 1, 2, 5, 7, 10, 20, and 30 year maturities, and Barclays corporate bond portfolios. Red lines denote estimated excess returns with positive $\hat{b}^i_h$ loadings, and blue dashed lines denote estimated excess returns with negative $\hat{b}^i_h$ loadings. The forecast horizon is $h = 6$, and the sample consists of monthly observations from 1990:1 to 2014:9.

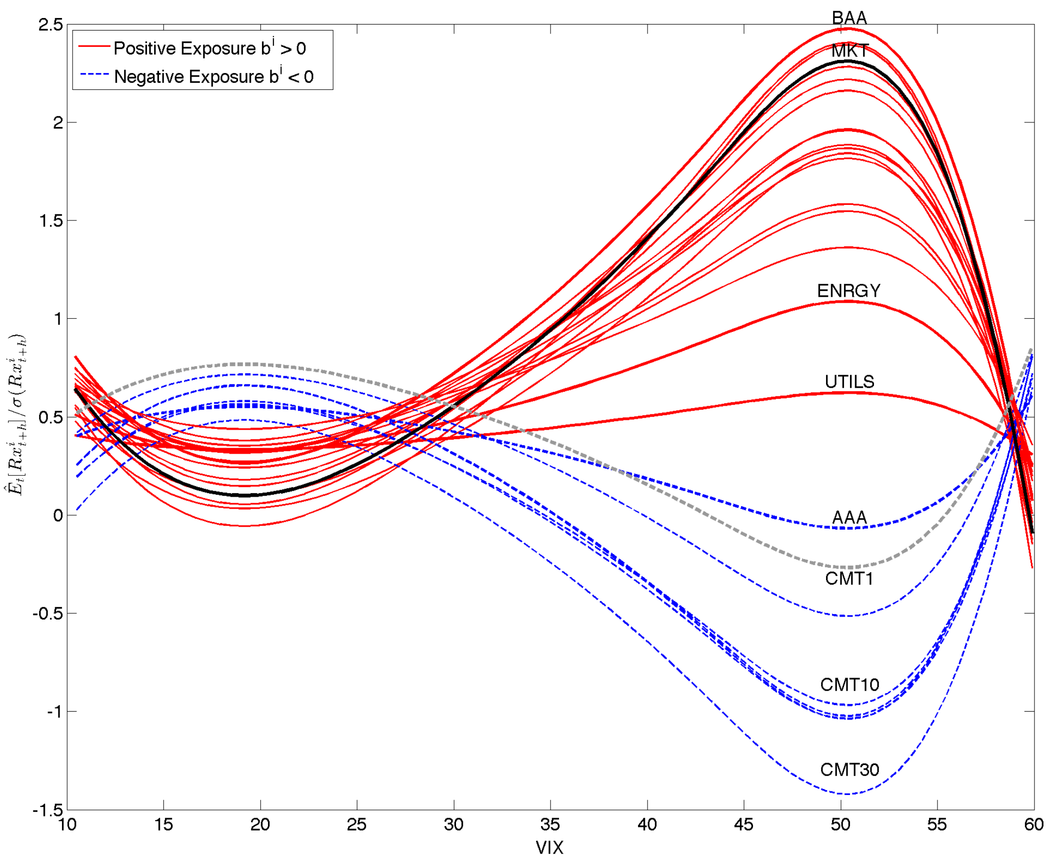

Figure 7: **SRRR $\phi_h(v)$: Separately Estimated from Equities and Treasuries**

This figure plots two versions of SRRR-estimated $\hat{\phi}_h(v)$. The first estimates $\hat{\phi}_h(v)$ from the sieve reduced rank regression $Rx^i_{t+h} = a^i_h + b^i_h \phi_h(v) + \varepsilon^i_{t+h}$ where i ranges over equity industry and market portfolios only (red dashed). The second estimate of $\phi_h(v)$ comes from the same sieve reduced rank regression, but where i ranges over Treasury return portfolios only (blue). The figure examines whether the two resulting nonlinear volatility functions $\hat{\phi}^{\text{Treas}}_h(v)$ and $\hat{\phi}^{\text{Equity}}_h(v)$ differ only by location and scale. This is tested by regressing $\hat{\phi}^{\text{Equity}}_h(v_t)$ on $\hat{\phi}^{\text{Treas}}_h(v_t)$ and a constant, and then plotting $\hat{\phi}^{\text{Equity}}_h(v)$ and $\hat{c}_1 + \hat{c}_2\hat{\phi}^{\text{Treas}}_h(v)$ alongside each other, where $\hat{c}_1$ and $\hat{c}_2$ are the regression coefficients. Dotted lines represent 95-percent confidence intervals from test $\mathbb{H}_{3,0}$ in Proposition 1 in the text. The forecast horizon is $h = 6$, and the sample consists of monthly observations from 1990:1 to 2014:9.

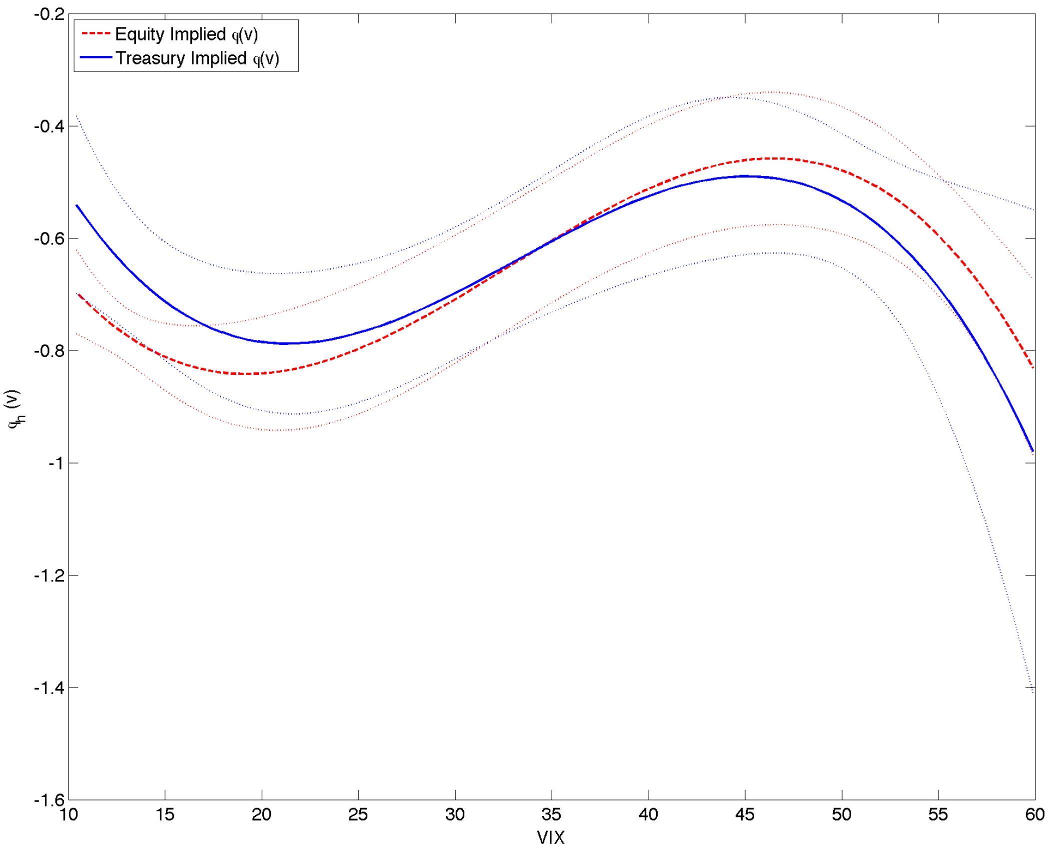

Figure 8: **Out-of-Sample Performance**

This figure plots cumulative excess returns of the CRSP value-weighted market portfolio (MKT), the 1-year constant maturity Treasury portfolio ($cmt1$), and a portfolio of assets formed using pseudo out-of-sample forecasts of the sieve reduced rank regressions $E_t[Rx^i_{t+h}] = a^i_h + b^i_h \phi_h(vix_t)$, for i ranging over MKT and $cmt1, \ldots, cmt30$. That is, given data on excess returns $\mathrm{Rx}_t = (Rx_t^{MKT}, Rx_t^{cmt1}, \ldots, Rx_t^{cmt30})'$ and the VIX for $t = 1, \ldots, T$, the sample is split into an initial in-sample data set $t = 1, \ldots, t^*$ and an out-of-sample data set $t^* + 1, \ldots, T$, where $(t^*/T) = 0.5$ is the fraction of available data reserved for the first out-of-sample forecast. Then, using in-sample data, a pseudo out-of-sample joint forecast $E_{t^*}[\mathrm{Rx}_{t^*+h}]$ is made for $h = 6$ months by sieve reduced rank regression. Based on this forecast, an optimal portfolio with weight $\omega_{t^*} = V_{t^*}^{-1} E_{t^*}[\mathrm{Rx}_{t^*+h}]$ is formed and held for $h = 6$ months, where risk weights $V_{t^*}^{-1}$ are the unconditional variances of excess returns using data from $t = 1, \ldots, t^*$. After $h = 6$ months, the in-sample data set is expanded to $t = 1, \ldots, t^* + h$, yielding a new out-of-sample forecast $E_{t^*+h}[\mathrm{Rx}_{t^*+2h}]$ and new portfolio weights ω_{t^*+h}. The process is repeated, yielding pseudo out-of-sample excess returns from $t^* + 1, \ldots, T$ whose cumulative return is displayed in the figure. All returns were scaled (levered) to have the same ex-post variance as MKT.

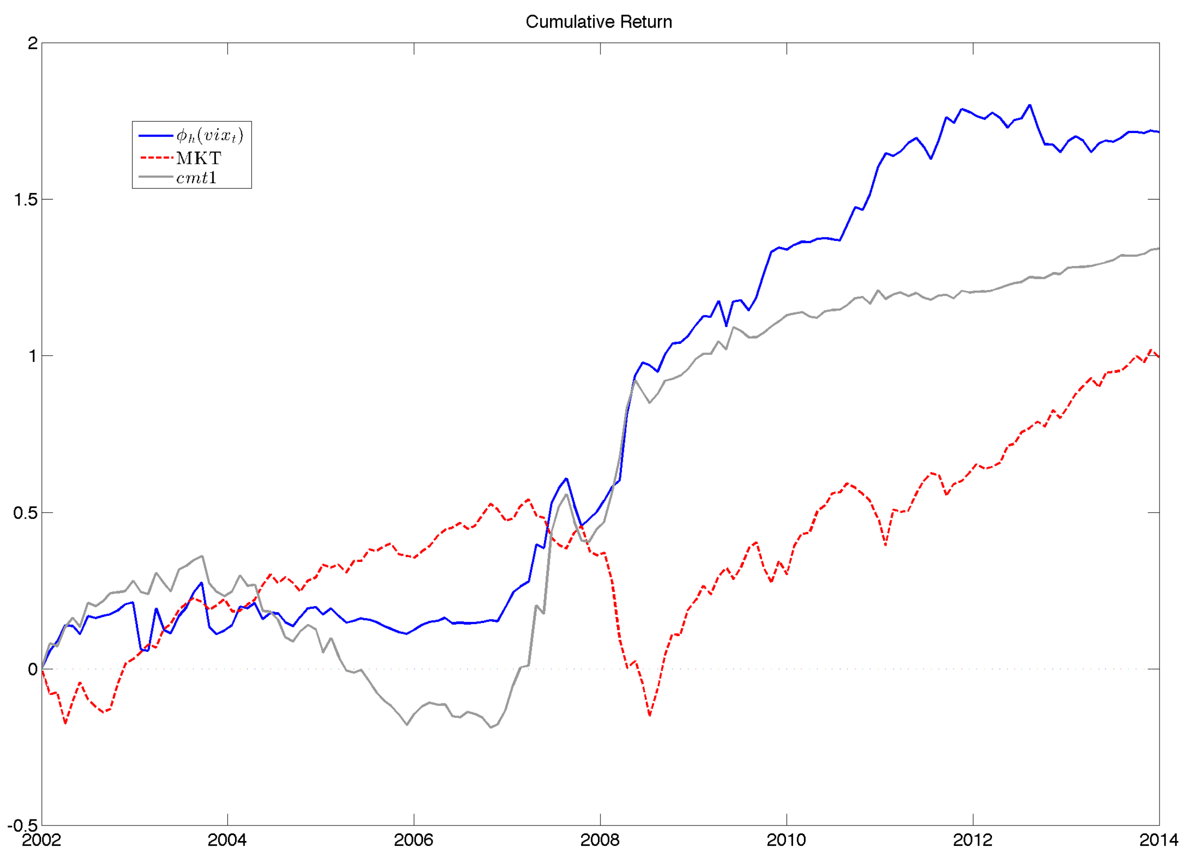

Figure 9: **Cross-Sectional Pricing**

This figure plots the results of the unrestricted joint forecasting regressions $Rx^i_{t+1} = a^i + b^i\phi(v_t) + \varepsilon^i_{t+1}$ against the restricted joint forecasting regressions $Rx^i_{t+1} = (\alpha^i + \beta^i\lambda_0) + \beta^i\lambda_1\phi(v_t) + \beta^i u_{t+1} + \varepsilon^i_{t+1}$, obtained from a dynamic asset pricing model with affine prices of risk. The innovations $u_{t+1} = Y_{t+1} - E_t[Y_{t+1}]$ correspond to the cross-sectional pricing factors $Y_t = (MKT_t, TSY1_t, \phi(v_t))$, where MKT is the return to the CRSP value-weighted equity market return, TSY1 is the return to the one year Treasury, and $\phi(v_t)$ is the nonlinear pricing factor for $v_t = vix_t$. Excess returns over i refer to 11 equity portfolios sorted by industry from Ken French's website, seven maturity-sorted Treasury portfolios, the six Barclay's industry and ratings sorted corporate bond portfolios, and the CRSP market return. To obtain estimates, the unrestricted regression is estimated by sieve reduced rank regression, yielding parametric estimates of a^i and b^i and a nonparametric estimate of $\phi(v_t)$. Then the restricted joint forecasting regression is estimated by taking $\phi(v_t)$ as given, making a^i and b^i from the unrestricted forecasts directly comparable to $(\alpha^i + \beta^i\lambda_0)$ and $\beta^i\lambda_1$ in the restricted regressions. The comparisons are scattered in the plots. The sample consists of monthly observations from 1990:1 to 2014:9.

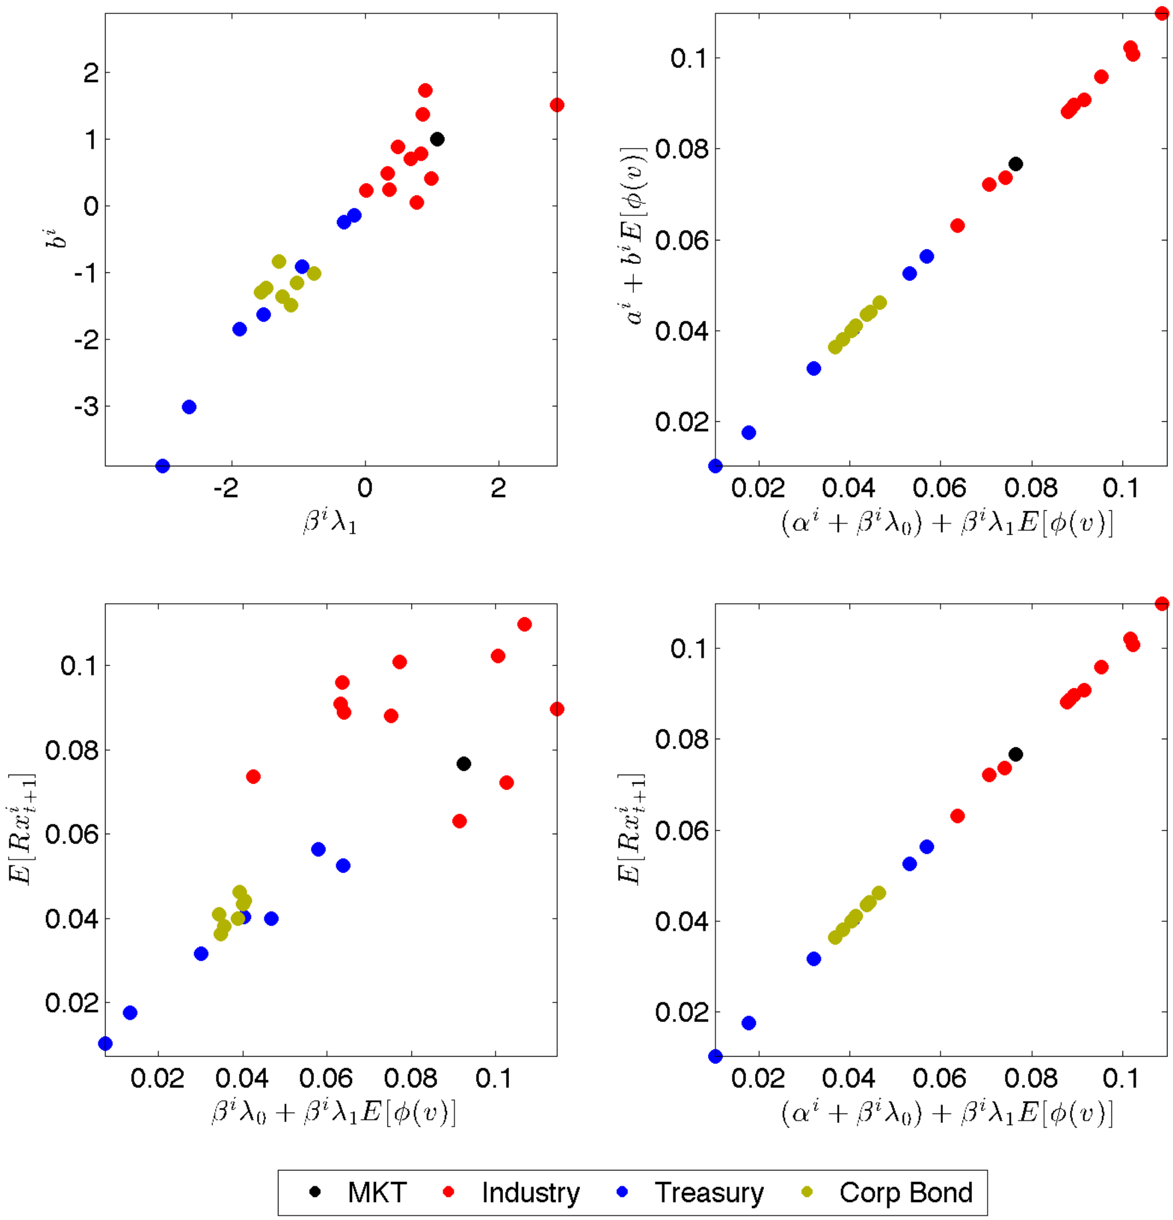

Figure 10: $\widehat{\phi}_6(v_t)$ **and Price of Risk Factors: 1990:3 to 2013:9**

This figure plots a nonlinear function of the VIX, $\phi_6(v_t)$, estimated by sieve reduced rank regression (SRRR) against known price of risk factors: the 10-year Treasury yield (TSY10), the term spread between the 10-year and 3-month Treasury yield (TERM), the Cochrane-Piazzesi Factor (CP), the spread between Moody's Baa-rated corporate bonds and the 10-year Treasury yield (DEF), the log dividend yield (DY), and the Lettau-Ludvigson CAY factor (CAY). To facilitate visual comparisons, all variables are demeaned and scaled by their unconditional standard deviatons.The sample period is 1990:3-2013:9.

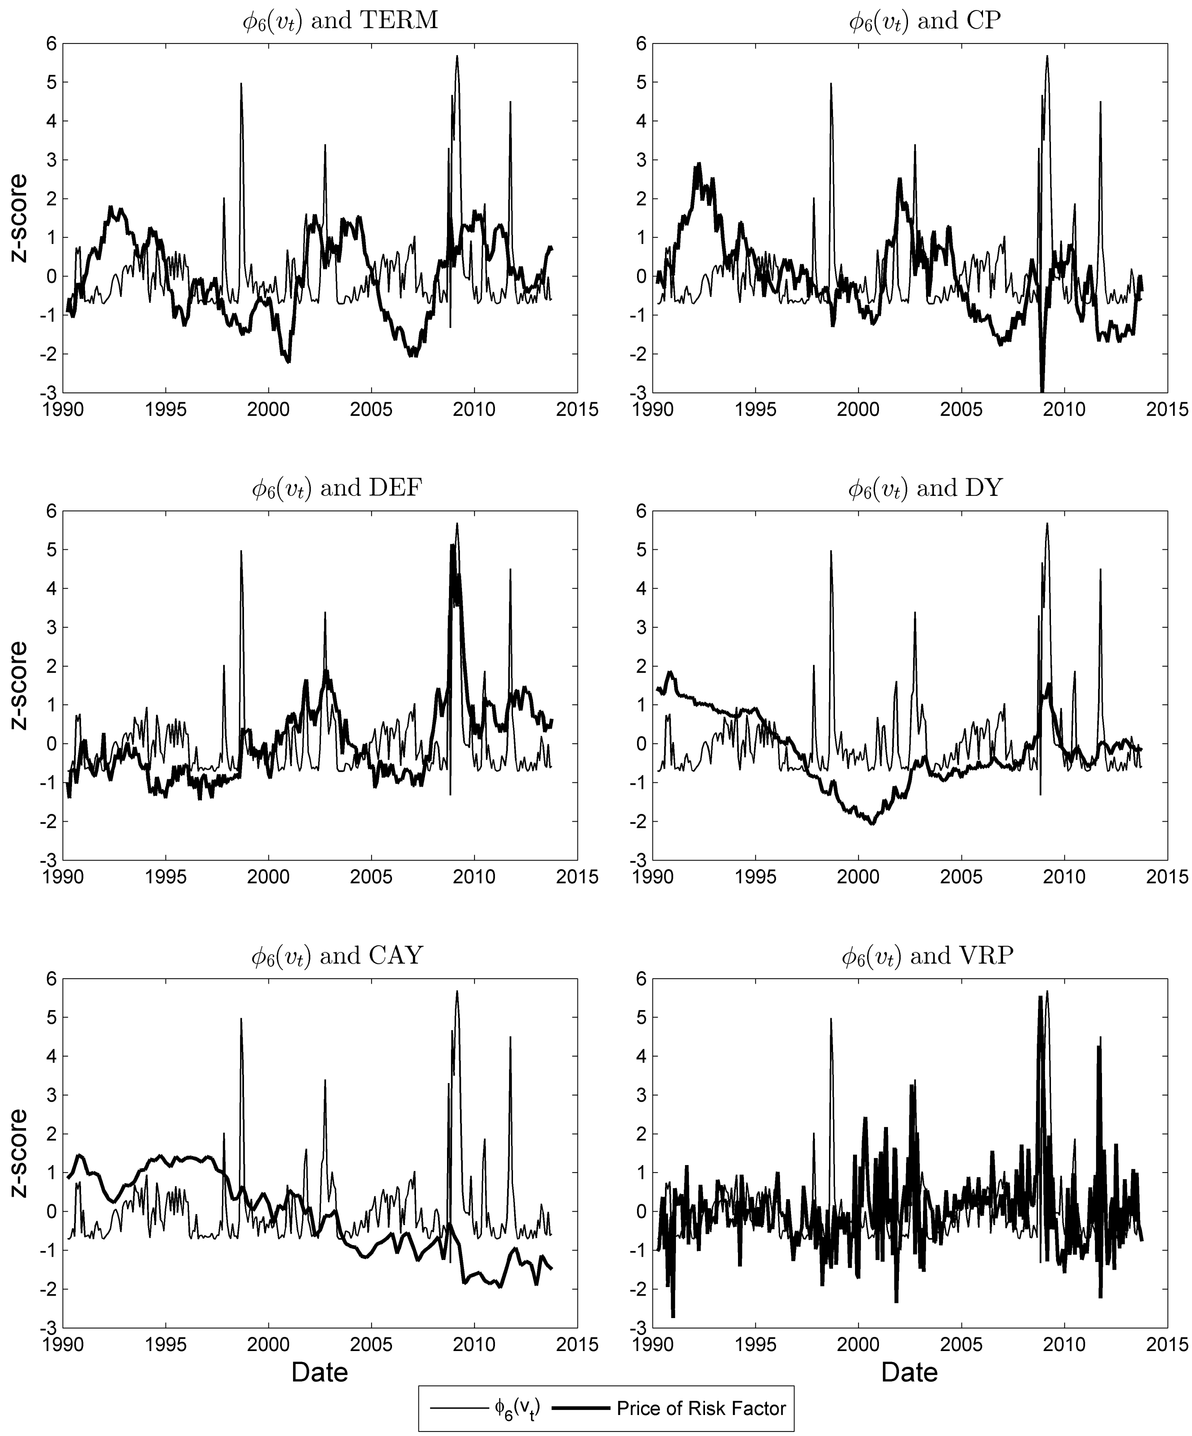

Figure 11: **Fund Flow** $\phi^{FF}(vix_t)$ **vs Market-Treasury** $\phi_h(vix_t)$

This figure plots $\phi^{FF}(v)$ estimated from a panel of mutual fund flows in the sieve reduced rank regressions $Flows_t^i = a^i + b^i \phi^{FF}(vix_t) + \varepsilon_t^i$. Plotted alongside $\phi^{FF}(v)$ is $\phi_h(v)$ estimated separately from sieve reduced rank regressions $Rx_{t+h}^i = a_h^i + b_h^i \phi_h(v_t) + \varepsilon_{t+h}^i$, where $h = 6$ months and $i = 1, \ldots, n$ ranges over the CRSP value-weighted market excess return and the seven CRSP constant maturity Treasury excess returns corresponding to 1, 2, 5, 7, 10, 20, and 30 years to maturity. Given the different units in each regression, $\phi^{FF}(v)$ was affine-translated to the scale of $\phi^{FF}(v)$. The sample period is 1990:1 to 2014:9.

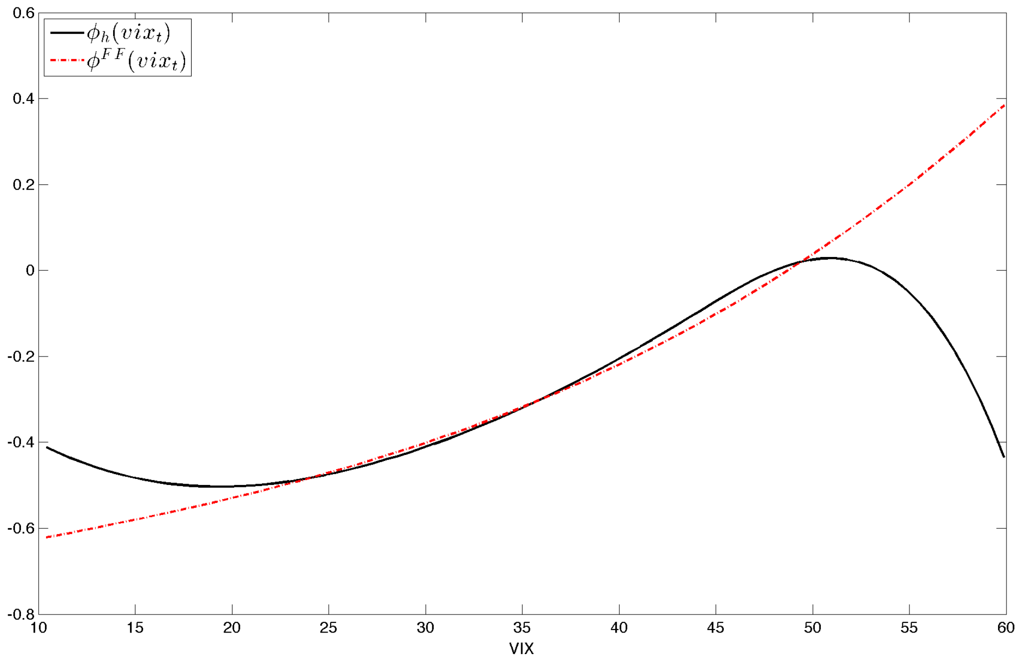

Figure 12: **SRRR: Value-at-Risk Cross-Section and the VIX**

The top panel of this figure plots the VIX (left axis, solid line) next to major dealer banks' summed Value-at-Risks (VaR) (right axis, dashed line) over time. The bottom panel shows the estimated $\phi^{VaR}(vix)$ from the contemporaneous sieve reduced rank regression of dealers' disaggregated VaRs on the VIX, $VaR_t^i = a^i + b^i\phi^{VaR}(vix_t) + \varepsilon_t^i$, where $i = 1, \ldots, 5$ indexes individual VaRs of the dealer banks that comprise the aggregate measure: Bank of America, Citigroup, Goldman Sachs, JP Morgan, and Morgan Stanley. Dotted lines represent 95-percent confidence intervals from test $\mathbb{H}_{3,0}$ in Proposition 1 in the text. The sample consists of quarterly observations from 2004:1 to 2014:4.

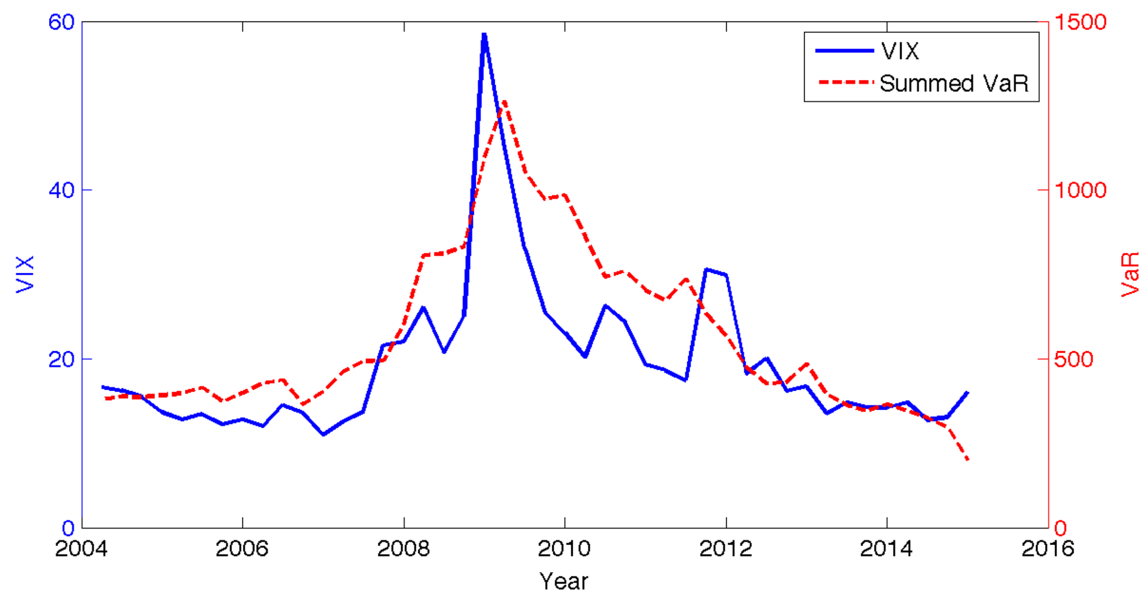

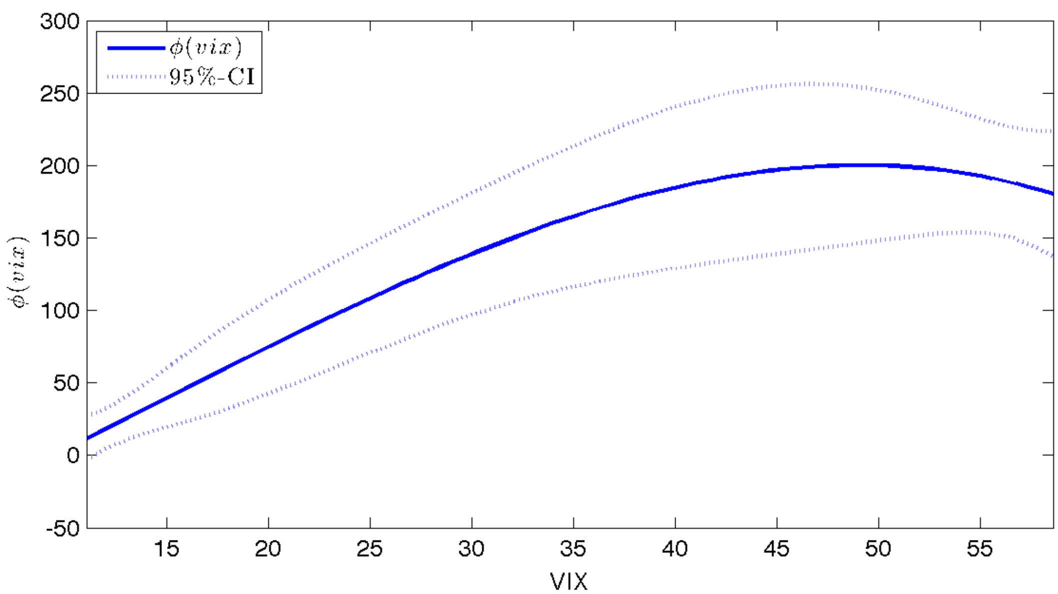

Figure 13: **Relation to Habit-Based Asset Pricing**

The plot shows the VIX together with the maximum Sharpe ratio of the habit formation model by Campbell and Cochrane (1999, 2000). Using NIPA consumption data on non-durable goods and services, a time series s_t of the log surplus consumption ratio was generated, yielding the maximal implied time-varying Sharpe Ratio as a function of s_t (plotted in dashed red). The VIX is plotted in solid blue. The sample consists of quarterly observations from 1990:1 to 2014:4.

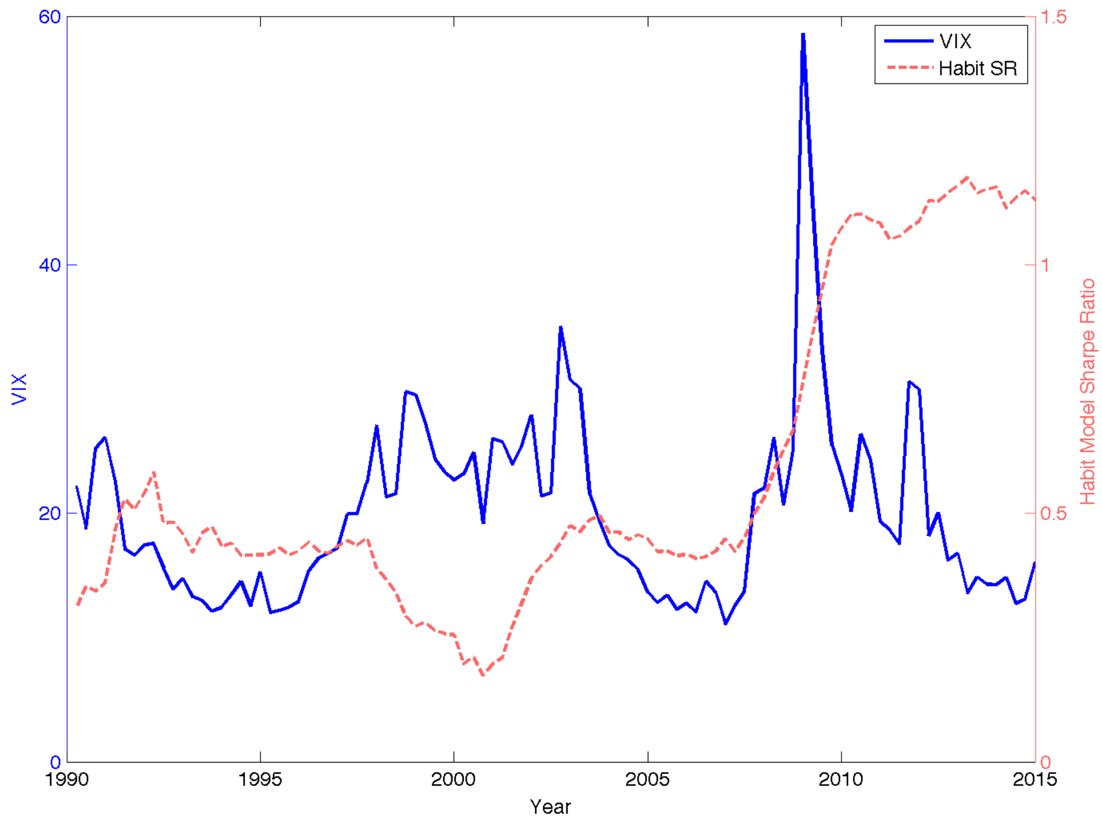

Figure 14: **Real Activity and** $\phi_h(vix_t)$

This plot shows the Federal Reserve Bank of Chicago's measure of real activity, the CFNAI-MA3, alongside $\phi(v_t)$ estimated from stocks and bonds alone. The sample consists of monthly observations from 1990:1 to 2014:9.

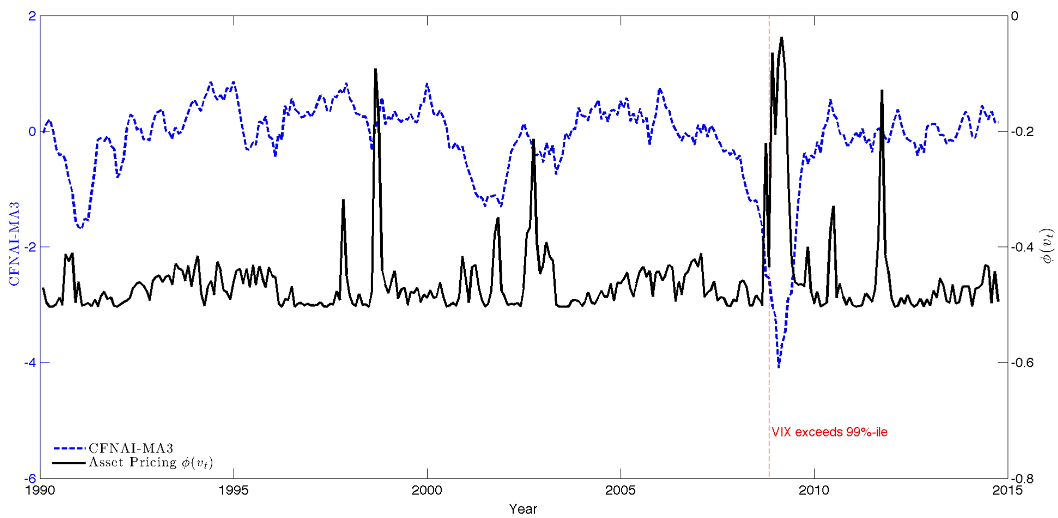

Table 1: **Excess Return Predictability: VIX and MOVE Polynomials: 1990 to 2014**

This table reports Hodrick (1992) t-statistics for coefficients from the regressions $Rx^i_{t+h} = a^i_h + \mathbf{b}^i_h(vix_t, vix^2_t, vix^3_t)' + \mathbf{c}^i_h(move_t, move^2_t, move^3_t)' + \mathbf{f}^i_h\mathbf{z_t} + \varepsilon^i_{t+h}$, where vix_t is the VIX equity implied volatility index at time t, and $move_t$ is the MOVE Treasury implied volatility index. The superscript i indexes separate regressions in which the left hand side variable is either the indicated Treasury excess return (top panel) or the equity market excess return (bottom panel). $\mathbf{z_t}$ consists of control variables representing the default spread between the 10-year Treasury yield and Moody's BAA corporate bond yield, the variance risk premium, the term spread between the 10-year and 3-month Treasury yield, and the log dividend yield. p-values report the outcome of the joint hypothesis test under the null of no predictability. The sample consists of monthly observations from 1990:1 to 2014:9.

1-year Treasury Excess Returns												
	$h = 6$ months				$h = 12$ months				$h = 18$ months			
VIX^1	1.91	4.13		5.01	1.86	3.60		5.02	1.13	3.21		4.73
VIX^2		-4.08		-4.77		-3.61		-4.86		-3.37		-4.71
VIX^3		3.89		4.66		3.51		4.76		3.38		4.64
$MOVE^1$			0.26	-0.57			-0.58	-1.94			-0.23	-2.21
$MOVE^2$			0.17	0.81			1.00	2.14			0.54	2.43
$MOVE^3$			-0.35	-0.97			-1.16	-2.27			-0.69	-2.57
DEF				-0.99				-0.99				-1.23
VRP				2.49				2.84				3.87
TERM				-1.45				-2.76				-3.47
DY				3.34				3.58				3.24
const	1.40	-3.42	-0.14	0.58	1.65	-2.75	1.05	1.57	2.17	-2.19	1.19	1.63
p-value	0.057	0.001	0.089	0.000	0.064	0.005	0.303	0.000	0.260	0.004	0.845	0.000
Stock Excess Returns												
	$h = 6$ months				$h = 12$ months				$h = 18$ months			
VIX^1	1.00	-3.18		-2.68	0.74	-2.47		-1.98	0.78	-1.87		-1.35
VIX^2		3.36		2.90		2.61		2.22		2.01		1.54
VIX^3		-3.24		-2.68		-2.45		-2.15		-1.87		-1.48
$MOVE^1$			0.15	0.27			1.14	1.26			0.13	0.33
$MOVE^2$			-0.06	-0.23			-1.17	-1.35			-0.07	-0.41
$MOVE^3$			-0.01	-0.01			1.20	1.23			0.16	0.40
DEF				-0.58				-0.62				-0.34
VRP				-2.34				-2.03				-2.33
TERM				0.47				0.94				1.03
DY				1.17				1.51				1.63
const	-0.15	3.28	0.13	2.43	0.50	2.68	-0.65	2.06	0.58	2.10	0.15	2.01
p-value	0.316	0.007	0.991	0.018	0.460	0.032	0.674	0.075	0.439	0.088	0.810	0.112

Table 2: **Nonlinear VIX Predictability using the Cross-Section: 1990 - 2014**

This table reports results from three predictive sieve reduced rank regressions (SRRR) for each of $h = 6$, 12, and 18 month ahead forecasting horizons: (1) estimates of a_h^i and b_h^i from the SRRR $Rx_{t+h}^i = a_h^i + b_h^i v_t + \varepsilon_{t+h}^i$ of portfolio i's excess returns on linear $v_t = vix_t$; (2) estimates of a_h^i and b_h^i from the SRRR $Rx_{t+h}^i = a_h^i + b_h^i \phi_h(v_t) + \varepsilon_{t+h}^i$ of portfolio i's excess return on the common nonparametric function $\phi_h(\cdot)$ of $v_t = vix_t$; (3) the same regression augmented with controls $\mathbf{f^i} \cdot (DEF_t, VRP_t, TERM_t, DY_t)'$ representing the default spread (DEF, 10-year Treasury yield minus Moody's BAA corporate bond yield), the variance risk premium (VRP, realized volatility minus VIX), the term spread (TERM, 10-year minus 3-month Treasury yields), and the S&P 500's (log) dividend yield. The index $i = 1, \ldots, n$ ranges over the CRSP value-weighted market excess return and the seven CRSP constant maturity Treasury excess returns corresponding to 1, 2, 5, 7, 10, 20, and 30 years to maturity. The sieve reduced rank regressions are introduced in section 3 in the text. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level for t-statistics on a_i and $\mathbf{f^i}$ and for the χ^2-statistic on $b^i \phi_h(\cdot)$ derived in Proposition 1. The joint test p-value reports the likelihood that the sample was generated from the model where $(b^1, \ldots, b^n) \cdot \phi_h(\cdot) = 0$.

Horizon h = 6

	(1) Linear VIX		(2) Nonlinear VIX		(3) Nonlinear VIX and Controls					
	a^i	b^i	a^i	b^i	a^i	b^i	f^i_{DEF}	f^i_{VRP}	f^i_{TERM}	f^i_{DY}
MKT	−0.01	1.00	1.00*	1.00***	0.31	1.00***	0.05**	−1.42***	−0.01	0.17
cmt1	0.00	0.07*	−0.05*	−0.07***	−0.09**	−0.20***	0.00	0.03*	0.00*	0.02***
cmt2	0.01	0.09	−0.11*	−0.14***	−0.15*	−0.32***	0.00	0.08**	0.00	0.02**
cmt5	0.03	0.04	−0.26	−0.31***	−0.25	−0.60***	−0.02*	0.23**	0.01**	0.01
cmt7	0.04	0.04	−0.31	−0.38**	−0.27	−0.70***	−0.03**	0.32**	0.02**	0.00
cmt10	0.05	−0.08	−0.30	−0.37**	−0.25	−0.66**	−0.03**	0.39**	0.03***	0.01
cmt20	0.08	−0.22	−0.39	−0.49	−0.23	−0.74	−0.05***	0.51*	0.05***	−0.03
cmt30	0.10	−0.52	−0.58	−0.68	−0.29	−0.98	−0.07***	0.70*	0.06***	−0.06
Joint p-val		0.273		0.000		0.000				

Horizon h = 12

	(1) Linear VIX		(2) Nonlinear VIX		(3) Nonlinear VIX and Controls					
	a^i	b^i	a^i	b^i	a^i	b^i	f^i_{DEF}	f^i_{VRP}	f^i_{TERM}	f^i_{DY}
MKT	0.03	1.00	0.64	1.00*	0.09	1.00*	0.03**	−0.70***	0.00*	0.18
cmt1	0.00	0.11*	−0.05	−0.10***	−0.08	−0.36***	0.00	0.03**	0.00**	0.02***
cmt2	0.01	0.22	−0.08	−0.18**	−0.12	−0.57***	0.00	0.05**	0.00	0.03***
cmt5	0.02	0.33	−0.18	−0.39*	−0.21	−1.03**	−0.01	0.08*	0.01	0.03
cmt7	0.02	0.35	−0.23	−0.48*	−0.24	−1.23*	−0.02*	0.12*	0.01**	0.02
cmt10	0.03	0.15	−0.22	−0.47*	−0.25	−1.25*	−0.02**	0.13*	0.02**	0.03
cmt20	0.05	0.12	−0.31	−0.65	−0.27	−1.52	−0.04**	0.16	0.03***	0.00
cmt30	0.06	−0.17	−0.44	−0.88	−0.33	−1.92	−0.05**	0.21	0.04***	−0.01
Joint p-val		0.380		0.002		0.000				

Horizon h = 18

	(1) Linear VIX		(2) Nonlinear VIX		(3) Nonlinear VIX and Controls					
	a^i	b^i	a^i	b^i	a^i	b^i	f^i_{DEF}	f^i_{VRP}	f^i_{TERM}	f^i_{DY}
MKT	0.03	1.00	0.44	1.00	−0.03	1.00	0.02**	−0.59***	0.01	0.18
cmt1	0.01	0.07	−0.04	−0.13***	−0.06	−0.60***	0.00	0.04***	0.00**	0.02***
cmt2	0.01	0.19	−0.07	−0.24**	−0.10	−0.95***	0.00	0.07***	−0.01	0.03***
cmt5	0.02	0.36	−0.14	−0.48*	−0.18	−1.65***	−0.01	0.11**	0.00	0.03**
cmt7	0.02	0.38	−0.16	−0.56	−0.19	−1.87**	−0.01*	0.14**	0.00	0.03*
cmt10	0.03	0.23	−0.15	−0.52	−0.20	−1.90*	−0.01*	0.15**	0.01*	0.04
cmt20	0.04	0.29	−0.19	−0.69	−0.20	−2.16	−0.02*	0.15	0.02**	0.01
cmt30	0.05	0.04	−0.28	−0.91	−0.24	−2.63	−0.03**	0.18	0.03**	0.00
Joint p-val		0.586		0.025		0.000				

Table 3: **Nonlinear VIX Predictability using the Cross-Section: 1990 - 2007**

This table reports results from three predictive sieve reduced rank regressions (SRRR) for each of $h = 6$, 12, and 18 month ahead forecasting horizons: (1) estimates of a_h^i and b_h^i from the SRRR $Rx_{t+h}^i = a_h^i + b_h^i v_t + \varepsilon_{t+h}^i$ of portfolio i's excess returns on linear $v_t = vix_t$; (2) estimates of a_h^i and b_h^i from the SRRR $Rx_{t+h}^i = a_h^i + b_h^i \phi_h(v_t) + \varepsilon_{t+h}^i$ of portfolio i's excess return on the common nonparametric function $\phi_h(\cdot)$ of $v_t = vix_t$; (3) the same regression augmented with controls $\mathbf{f^i} \cdot (DEF_t, VRP_t, TERM_t, DY_t)'$ representing the default spread (DEF, 10-year Treasury yield minus Moody's BAA corporate bond yield), the variance risk premium (VRP, realized volatility minus VIX), the term spread (TERM, 10-year minus 3-month Treasury yields), and the S&P 500's (log) dividend yield. The index $i = 1, \ldots, n$ ranges over the CRSP value-weighted market excess return and the seven CRSP constant maturity Treasury excess returns corresponding to 1, 2, 5, 7, 10, 20, and 30 years to maturity. The sieve reduced rank regressions are introduced in section 3 in the text. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level for t-statistics on a_i and $\mathbf{f^i}$ and for the χ^2-statistic on $b^i \phi_h(\cdot)$ derived in Proposition 1. The joint test p-value reports the likelihood that the sample was generated from the model where $(b^1, \ldots, b^n) \cdot \phi_h(\cdot) = 0$.

Horizon h = 6

	(1) Linear VIX		(2) Nonlinear VIX		(3) Nonlinear VIX and Controls					
	a^i	b^i	a^i	b^i	a^i	b^i	f^i_{DEF}	f^i_{VRP}	f^i_{TERM}	f^i_{DY}
MKT	0.03	1.00	0.72	1.00	0.25	1.00	−0.03	−0.84*	0.00	0.12
cmt1	0.00	0.23**	−0.09	−0.16***	−0.15	−0.59***	0.01	0.03	0.00*	0.03***
cmt2	0.00	0.29	−0.16	−0.28**	−0.22	−0.92***	0.01	0.10*	0.00	0.03***
cmt5	0.01	0.30	−0.31	−0.53*	−0.34	−1.57***	0.00	0.25*	0.01	0.03*
cmt7	0.03	0.18	−0.36	−0.62*	−0.35	−1.75**	−0.01	0.31*	0.02	0.02
cmt10	0.04	−0.23	−0.37	−0.62	−0.32	−1.74*	−0.03	0.36*	0.02*	0.02
cmt20	0.06	−0.16	−0.42	−0.72	−0.31	−1.95	−0.05	0.46**	0.03**	−0.01
cmt30	0.05	−0.27	−0.51	−0.85	−0.35	−2.25	−0.06	0.58**	0.04**	−0.03
Joint p-val		0.058		0.001		0.000				

Horizon h = 12

	(1) Linear VIX		(2) Nonlinear VIX		(3) Nonlinear VIX and Controls					
	a^i	b^i	a^i	b^i	a^i	b^i	f^i_{DEF}	f^i_{VRP}	f^i_{TERM}	f^i_{DY}
MKT	0.09	1.00	0.46	1.00	−0.01	1.00	−0.03	−0.52*	0.00	0.16
cmt1	0.00	−2.09**	−0.08	−0.23***	−0.12	3.25***	0.00	0.03	0.00	0.03***
cmt2	0.00	−3.21*	−0.13	−0.40***	−0.18	5.27***	0.00	0.07*	0.00	0.03***
cmt5	0.00	−4.69	−0.24	−0.72**	−0.31	9.32***	0.00	0.12	0.01	0.04
cmt7	0.01	−3.98	−0.28	−0.84**	−0.33	10.77***	−0.01	0.13	0.01*	0.03
cmt10	0.03	−0.71	−0.27	−0.80*	−0.33	11.27**	−0.03*	0.14	0.02**	0.04
cmt20	0.04	−1.06	−0.33	−1.00*	−0.34	13.10**	−0.04**	0.14	0.03**	0.01
cmt30	0.04	−1.15	−0.40	−1.17*	−0.39	15.27**	−0.06**	0.16	0.04***	0.00
Joint p-val		0.008		0.001		0.000				

Horizon h = 18

	(1) Linear VIX		(2) Nonlinear VIX		(3) Nonlinear VIX and Controls					
	a^i	b^i	a^i	b^i	a^i	b^i	f^i_{DEF}	f^i_{VRP}	f^i_{TERM}	f^i_{DY}
MKT	0.11	1.00	0.57	1.00**	0.13	1.00	−0.03	−0.54**	0.00	0.13
cmt1	0.00	−0.36**	−0.07	−0.17***	−0.11	−0.92***	0.00	0.05**	0.00	0.03***
cmt2	0.00	−0.59	−0.13	−0.30***	−0.19	−1.55***	0.00	0.11**	0.00	0.03***
cmt5	0.00	−0.93	−0.24	−0.55***	−0.35	−2.83***	0.01*	0.19**	0.00	0.05*
cmt7	0.01	−0.86	−0.27	−0.63***	−0.39	−3.28***	0.00**	0.22**	0.01*	0.05
cmt10	0.03	−0.25	−0.25	−0.59***	−0.40	−3.44***	−0.01**	0.24**	0.01**	0.06
cmt20	0.04	−0.45	−0.29	−0.70**	−0.41	−3.81***	−0.02**	0.20*	0.02**	0.03
cmt30	0.03	−0.43	−0.37	−0.85**	−0.50	−4.55***	−0.03**	0.25**	0.03***	0.03
Joint p-val		0.019		0.000		0.000				

Table 4: **Out-of-Sample Performance**

This table examines the forecasting performance of the sieve reduced rank regressions $E_t[Rx^i_{t+h}] = a^i_h + b^i_h \phi_h(vix_t)$ by comparing pseudo out-of-sample Sharpe ratios of portfolios formed using information from the joint forecasts $E_t[Rx^i_{t+h}]$. As in previous tables, i ranges over CRSP excess market return (MKT) and the seven constant maturity Treasury portfolios (CMT) with maturities of 1, 2, 5, 7, 10, 20, and 30 years. Given data on excess returns $\mathrm{Rx}_t = (Rx^{MKT}_t, Rx^{cmt1}_t, \ldots, Rx^{cmt30}_t)'$ and the VIX for $t = 1, \ldots, T$, the sample is split into an initial in-sample data set $t = 1, \ldots, t^*$ and an out-of-sample data set $t^*+1, \ldots, T$, where (t^*/T) is the fraction displayed in the first column. Using in-sample data, a pseudo out-of-sample joint forecast $E_{t^*}[\mathrm{Rx}_{t^*+h}]$ is made for $h = 6$ months by sieve reduced rank regression. Based on this forecast, an optimal portfolio with weight $\omega_{t^*} = V^{-1}_{t^*} E_{t^*}[\mathrm{Rx}_{t^*+h}]$ is formed and held for $h = 6$ months, where risk weights $V^{-1}_{t^*}$ are the unconditional variances of excess returns using data from $t = 1, \ldots, t^*$. After $h = 6$ months, the in-sample data set is expanded to $t = 1, \ldots, t^* + h$, yielding a new forecast $E_{t^*+h}[\mathrm{Rx}_{t^*+2h}]$ and new portfolio weights ω_{t^*+h}. The process is repeated, yielding pseudo out-of-sample excess returns from $t^* + 1, \ldots, T$ whose ex-post Sharpe ratios are displayed in the table. The column labeled "$\phi_h(vix_t)$" uses sieve reduced rank forecasts $E_t[Rx^i_{t+h}] = a^i_h + b^i_h \phi_h(vix_t)$, the column labeled "$vix_t$" uses linear reduced rank forecasts $E_t[Rx^i_{t+h}] = a^i_h + b^i_h\ vix_t$. The column labeled "Uncond." uses running mean forecasts $E_t[Rx^i_{t+h}] = a^i_h$. The column labeled "EQL" represents an equal-weighted portfolio of MKT and $cmt1, \ldots, cmt30$. The remaining columns put all of the portfolio weight on the indicated asset.

In-Sample Cutoff (t^*/T)	**Annualized Sharpe Ratio** (1) $\phi_h(vix_t)$	(2) vix_t	(3) Uncond.	(4) EQL	(5) MKT	(6) cmt1	(7) cmt5	(8) cmt10	(9) cmt30
0.4 (Nov-1999)	1.03	0.87	0.88	0.79	0.24	1.04	0.79	0.61	0.47
0.5 (May-2002)	0.91	0.73	0.72	0.82	0.53	0.71	0.68	0.59	0.45
0.6 (Nov-2004)	1.12	0.77	0.73	0.77	0.46	0.76	0.66	0.58	0.41

Table 5: **Dynamic Asset Pricing: Factor Risk Exposures and Prices of Risk**

This table provides estimates of factor risk exposures and prices of risk from the dynamic asset pricing model $Rx^i_{t+1} = (\alpha^i + \beta^i\lambda_0) + \beta^i\lambda_1\phi(v_t) + \beta^i u_{t+1} + \varepsilon^i_{t+1}$, where i ranges over the test assets in the left column: MKT denotes the CRSP value-weighted market excess return, (NoDur ... Fin) are industry-sorted portfolio excess returns from Ken French's website, (cmt1, ..., cmt30) are constant maturity Treasury portfolio excess returns, and (AAA, ..., igfin) are Barclay's ratings and industry sorted corporate bond excess returns. In a first stage, $\phi(v_t)$ is estimated from a sieve reduced rank regression $Rx^i_{t+1} = a^i + b^i\phi(v_t) + \varepsilon^i_{t+1}$ jointly across i. The factor innovations $u_{t+1} = Y_{t+1} - E_t[Y_{t+1}]$ are then estimated from a VAR on the market (MKT), Treasury (TSY1), and nonlinear volatility factor ($\phi(v_t)$), for $v_t = vix_t$. In a second stage, coefficients are estimated jointly across all $i = 1, \dots, n$ via a reduced rank regression. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level.

Exposures	β^i_{MKT}	β^i_{TSY1}	$\beta^i_{\phi(v)}$	$\beta^i\lambda_1$	$(\alpha^i + \beta^i\lambda_0)$
MKT	1.00***	−0.24***	0.02	1.08***	0.33***
NoDur	0.61***	0.45	0.01	0.49**	0.21***
Durbl	1.19***	−2.24**	1.90**	0.68	0.23
Manuf	1.09***	−0.86	0.85***	0.83***	0.30***
Enrgy	0.71***	−1.11	1.11	0.36	0.18
Chems	0.75***	−0.80	0.20	0.87***	0.30***
BusEq	1.44***	−1.76**	−1.47***	2.88***	0.80***
Telcm	0.94***	0.06	0.27	0.77*	0.25**
Utils	0.39***	0.10	0.59	0.01	0.08
Shops	0.86***	−0.64	0.10	0.99***	0.32***
Hlth	0.69***	1.04	0.12	0.33	0.18**
Fin	1.08***	1.26	−0.24	0.90**	0.30***
cmt1	0.00	0.73***	−0.07	−0.16**	−0.03*
cmt2	−0.01	1.40***	−0.14	−0.32**	−0.06*
cmt5	−0.03*	2.90***	0.16	−0.95***	−0.19***
cmt7	−0.04*	3.55***	0.77*	−1.52***	−0.32***
cmt10	−0.04	3.90***	1.19***	−1.88***	−0.41***
cmt20	−0.08*	4.76***	1.96***	−2.64***	−0.57***
cmt30	−0.12**	5.45***	2.23*	−3.03***	−0.67***
AAA	0.02	2.54***	0.68	−1.11***	−0.23***
AA	0.06***	2.28***	0.71	−1.02***	−0.21***
A	0.09***	2.00***	1.24***	−1.24***	−0.26***
BAA	0.12***	1.55***	1.58***	−1.29***	−0.26***
igind	0.08***	1.90***	1.67***	−1.48***	−0.31***
igutil	0.07**	1.99***	1.73***	−1.56***	−0.33***
igfin	0.11***	1.83***	0.57	−0.76***	−0.14**
Prices of Risk	*MKT*	*TSY1*	$\phi(v_t)$		
λ_1	1.02***	−0.28**	−0.62**		

Table 6: **Fund Flows and Nonlinear VIX**

This table reports results from the contemporaneous panel regressions of mutual fund flows into the funds of indicated type i (left column) on the common nonparametric function $\phi^{FF}(vix_t)$, $Flows_t^i = a^i + b^i \phi^{FF}(vix_t) + \varepsilon_t^i$. The panel regressions were estimated via the sieve reduced rank regressions introduced in section 3 in the text. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level for t-statistics on a^i and $\mathbf{f^i}$ and for the χ^2-statistic on $b^i \phi^{FF}(\cdot)$ derived in Proposition 1. The joint test p-value reports the likelihood that the sample was generated from the model where $(b^1, \ldots, b^n) \cdot \phi^{FF}(\cdot) = 0$.

	Sample: 1990 - 2014		Sample: 1990 - 2007	
	a^i	b^i	a^i	b^i
us equity	0.50*	−1.00***	−0.29	−1.00
world equity	0.88	−1.77***	−1.06	−3.60***
hybrid	0.94*	−1.89***	−1.10	−3.75***
corporate bond	0.16	−0.32	0.02	0.08
HY bond	−0.02	0.05	0.04	0.12
world bond	0.34	−0.68	−0.29	−1.00
govt bond	−0.63*	1.27***	1.08	3.68***
strategic income	0.02	−0.04	0.35	1.18
muni bond	−0.01	0.03	−0.05	−0.16
govt mmmf	−0.42	0.85	0.39	1.33
nongovt mmmf	−0.02	0.03	0.36	1.23
national mmmf	0.00	0.00	0.16	0.56
state mmmf	0.28	−0.56**	0.10	0.35
Joint p-value		0.000		0.000

Table 7: **Business Cycle Panel Predictability**

This table reports results from two cross-sectional predictability regressions of various macro business cycle indicators on $h = 6$ month lagged functions of $v_t = VIX_t$. The left panel is estimated on our full sample of monthly observations from 1990 to 2014, whereas the right panel is estimated on the pre-crisis subsample. The business cycle indicators represent the industrial production index (IP), the IP manufacturing index (IPMFG), the manufacturing capacity utilization index (CUMFG), the change in goods-producing employment (LAGOODA), and the total private nonfarm payroll series (LAPRIVA), which receive the largest weight within the Chicago Fed National Activity index. The leading reference asset is the market excess return. The left panel shows estimates of a_h^i and b_h^i from the sieve reduced rank regression $y_{t+h}^i = a_h^i + b_h^i \phi^{\text{macro}}(v_t) + \varepsilon_{t+h}^i$ of series i on the common nonparametric function $\phi_h^{\text{macro}}(\cdot)$. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level for t-statistics on a_i and for the χ^2-statistic on $b^i \phi_h^{\text{macro}}(\cdot)$ derived in Proposition 1. The joint test p-value reports the likelihood that the sample was generated from the model where $(b^1, \ldots, b^n) \cdot \phi_h(\cdot) = 0$.

	Sample: 1990 - 2014		Sample: 1990 - 2007	
	a^i	b^i	a^i	b^i
MKT	−0.02	1.00	0.10	1.00
IP	0.01	−0.06***	0.00	−0.22***
IPMFG	0.01	−0.07***	0.00	−0.25***
CUMFG	0.00	−0.05**	−0.01	−0.26***
LAPRIVA	0.00	−0.04***	0.00	−0.14***
LAGOODA	0.01	−0.07***	−0.01	−0.26***
Joint p-value		0.000		0.000